AeroBest 2025

III ECCOMAS Thematic Conference on Multidisciplinary Design Optimization of Aerospace Systems

Proceedings

22–24 April 2025
Lisboa, Portugal

André C. Marta & Afzal Suleman (editors)

Published by:
IDMEC
Instituto Superior Técnico
Universidade de Lisboa
Portugal
https://aerobest2025.idmec.tecnico.ulisboa.pt/

ISBN: 978-989-53599-5-0

Editors:
André C. Marta
Afzal Suleman

Preface

This Book of Proceedings contains the full manuscripts on the works presented at AeroBest 2025, the third edition of the ECCOMAS Thematic Conference on Multidisciplinary Design Optimization of Aerospace Systems. This document is a testimony of the latest developments in the field, covering a wide range of MDO computational methods, tools and processes and applications spanning civil aviation and space systems.

The extension of this book is a testimony of the growing interest in this thematic, having this third edition gathered almost 80 participants from 14 countries. A total of 57 presenters, including four keynote speakers from academia, research centers and industry, shared their most recent developments in MDO methods and latest applications to complex aerospace design problems. The drive towards more sustainable aerospace technology was present in many works, in which different challenges in aerodynamic efficiency, emissions and noise, advanced materials, manufacturing, robust control and propulsion integration were tackled.

The nature of this conference, having a well defined underlying theme, translated into engaging technical sessions and social breaks throughout the three days of the event, that led to the network enlargement and strengthening among all participants.

The mix of participants from academia, government laboratories and industry allowed for productive and pragmatic discussions, which potentially will help align future developments with the societal challenges, industry demands and technical gaps identified.

A final word of appreciation to all participants that made the conference possible, particularly to those whose work is documented in these proceedings.

André C. Marta
(Conference Chair)

Afzal Suleman
(Conference Chair)



Scientific Committee

Juan Alonso, Stanford University, USA

Pedro Camanho, Universidade do Porto, Portugal

Giuseppe Catalanotti, Kore University of Enna, Italy

Jonathan Cooper, University of Bristol, United Kingdom

Pedro V. Gamboa, Universidade da Beira Interior, Portugal

Graeme Kennedy, Georgia Tech, USA

Alicia Kim, University of California San Diego, USA

André C. Marta, Instituto Superior Técnico, Portugal

Joaquim R. R. A. Martins, University of Michigan, USA

Joseph Morlier, ISAE-SUPAERO / ICA, France

Melike Nikbay, Istanbul Technical University, Turkey

Jacques F. Périaux, CIMNE / UPC Barcelona, Spain

Sergio Ricci, Politecnico di Milano, Italy

Hélder C. Rodrigues, Instituto Superior Técnico, Portugal

Carlos A. M. Soares, Instituto Superior Técnico, Portugal

Craig Steeves, University of Toronto, Canada

Afzal Suleman, University of Victoria, Canada

Sponsors

Mechanical Engineering Institute (*Instituto de Engenharia Mecânica*)

Portuguese Society for Theoretical, Applied and Computational Mechanics (*Associação Portuguesa de Mecânica Teórica, Aplicada e Computacional*)

European Community on Computational Methods in Applied Sciences

Secretariat

IDMEC
Instituto Superior Técnico
Av. Rovisco Pais, 1
1049-001 Lisboa, Portugal

(+351) 218 417 351
aerobest2025@idmec.tecnico.ulisboa.pt
https://aerobest2025.idmec.tecnico.ulisboa.pt/

Keynote Lecturers



Melike Nikbay is a Professor at the Istanbul Technical University, where she heads the Aerospace Multidisciplinary Design Optimization Laboratory. Her research interests include advanced computational design, multidisciplinary analysis and optimization, multi-fidelity modeling, surrogate modeling, machine learning, uncertainty quantification, aeroelasticity and aeroelastic optimization, fluid-structure interaction, sensitivity analysis (direct and adjoint methods), reliability-based design optimization, robust and stochastic optimization, and sonic boom minimization. Dr. Nikbay is the Chair of the Department of Astronautical Engineering since May 2024, and an Executive Board Member of TÜBİTAK's Space Technologies Research Group since 2024. Prior to her full-time academic career in 2004, Dr. Nikbay served as R&D Leader at multiple industries, was a visiting researcher at the USAF Wright-Patterson Multidisciplinary Science and Technology Center in Ohio, and led a project at NASA Langley Research Center on "The Impact of Aeroelastic Uncertainty on Low Sonic Boom Aircraft Shaping Optimization" as part of the Commercial Supersonic Aircraft Technology program. Dr. Nikbay earned her B.S. (1996) and M.S. (1998) degrees in Mechanical Engineering from Boğaziçi University, Türkiye. She later obtained a second M.S. in Aerospace Engineering from the University of Colorado Boulder, USA, in 1999. In 2002, she completed her Ph.D. under the mentorship of Prof. Charbel Farhat. Dr. Nikbay has been a representative of Türkiye in the NATO Science and Technology Organization (STO) Applied Vehicle Technology (AVT) Panel since 2008, contributing extensively through various leadership roles. She currently serves as a member of the Scientific Committee of the NATO STO AVT Panel, Director of the AVT-385 Research Lecture Series on "Multi-Fidelity Methods for Multidisciplinary Design Optimization," and co-chair of the AVT-411-RSM program on "Machine Learning and Artificial Intelligence for Military Vehicle Design." Since 2010, Dr. Nikbay has also served as an External Expert Evaluator for the European Commission's Clean Sky and EISMEA programs in aviation. Her achievements include recognition as an AIAA Associate Fellow (2017) and prestigious awards, such as the NATO STO "Young Contributor" Award (2016) and the NATO STO "Panel Excellence" Award (2024).



Joaquim R. R. A. Martins is the Pauline M. Sherman Collegiate Professor of Aerospace Engineering at the University of Michigan, where he heads the Multidisciplinary Design Optimization Laboratory. His research ranges from developing new MDO techniques to applying those techniques to the design of aircraft and other engineering systems. He is a co-author of "Engineering Design Optimization", a textbook published by Cambridge University Press. Prof. Martins is a Fellow of the American Institute of Aeronautics and Astronautics and a Fellow of the Royal Aeronautical Society. Before joining the University of Michigan faculty in 2009, he was an Associate Professor at the University of Toronto Institute for Aerospace Studies, where from 2002 he held a Tier II Canada Research Chair in Multidisciplinary Optimization. He received his undergraduate degree in Aeronautical Engineering from Imperial College, London, with a British Aerospace Award. He obtained his M.Sc. and Ph.D. degrees from Stanford University, where he was awarded the Ballhaus prize for best thesis in the Department of Aeronautics and Astronautics. He has received the Best Paper Award in AIAA Conferences five times. He was a member of the AIAA MDO Technical Committee (2006–2018) and was the technical co-chair for the 2008 AIAA Multidisciplinary Analysis and Optimization Conference. He has served as Associate Editor for the AIAA Journal, Optimization and Engineering, and Structural and Multidisciplinary Optimization. He is currently an Associate Editor for the Journal of Aircraft.



Anant Grewal has been a Research Officer at NRC's Aerospace Research Centre since 1994, apart from the period from 2000 to 2006 when he worked in industry. In 2015, he was promoted to the grade of Principal Research Officer. From 2012 to 2019, Anant assumed the role of Program Leader of the Working and Travelling on Aircraft program at the NRC. In this capacity he led NRC's research program in the area of the aircraft cabin environment, and cabin human factors, health and safety. He continues as the Air Travel Research Focus Area Leader at NRC. As a Research Officer, he has been leading collaborative R&D and Technical Services projects in the areas of acoustics, structural dynamics, aeroelasticity, multidisciplinary analysis and optimization and smart structures applications with industrial and government partners. In addition to the NRC, Anant also worked as a Senior Designer at Nortel Networks, and as a Defence Scientist with the Canadian Department of National Defence. Anant is currently also an Adjunct Professor at the Toronto Metropolitan University in the Department of Aerospace Engineering. Anant is currently a member of the American Institute of Aeronautics and Astronautics (AIAA) Structural Dynamics Technical Committee, as well as an Associate Fellow of AIAA.



Ögmundur is the eXpert for Airframe Multidisciplinary Design Analysis and Optimization at Airbus Defence and Space in Manching, Germany. He is responsible for the development of the in-house multidisciplinary structural optimisation tool Lagrange and the application of MDO to airframe design tasks in both civil and military aircraft programmes including the A350, for which the team received an award for engineering excellence, and within the FCAS project. He has led the activities of Airbus Defence and Space in numerous national and European research projects on the topic of MDO, most recently within the Clean Aviation Joint Undertaking of the European Commission and Industry. Ögmundur completed his Bachelor's degree in Mechanical Engineering from the University of Iceland and Master's degree at the Technische Universität München in Germany before joining Airbus Defence and Space where he has since worked in the team for Stress Methods and Optimization within the Airframe Centre of Competence. Current research topics focus mostly on the creation of coupled digital design platforms, leveraging optimisation techniques, existing capabilities in disciplinary analysis and recent advances in sensitivity analysis to enable efficient design of coupled systems.

Programme Overview

Tuesday, April 22nd	
8:00 - 8:45	Registration
8:45 - 9:00	Opening Ceremony
9:00 - 10:00	Keynote Lecture I Melike Nikbay, Instanbul Technical University
10:00 - 11:00	SESSION 1 - Design Optimization I
11:00 - 11:20	Coffee Break
11:20 - 13:00	SESSION 2 - Design Optimization II
13:00 - 14:00	Lunch
14:00 - 15:40	SESSION 3 - Multi-Disciplinary Optimization I
15:40 - 16:00	Coffee Break
16:00 - 18:00	SESSION 4 - Discipline Analysis Models I

Wednesday, April 23rd	
9:00 - 10:00	Keynote Lecture II Joaquim Martins, University of Michigan
10:00 - 11:00	SESSION 5 - Design Optimization III
11:00 - 11:20	Coffee Break
11:20 - 13:00	SESSION 6 - Design Optimization IV
13:00 - 14:00	Lunch
14:00 - 15:40	SESSION 7 - Multi-Disciplinary Optimization II
15:40 - 16:00	Coffee Break
16:00 - 18:00	SESSION 8 - Discipline Analysis Models II
20:30 - 22:30	Conference Dinner

Thursday, April 24th	
9:00 - 10:00	Keynote Lecture III Anant Grewal, National Research Council of Canada
10:00 - 11:00	SESSION 9 - Systems Engineering and Integration
11:00 - 11:20	Coffee Break
11:20 - 13:00	SESSION 10 - Aerospace Design and Integrated Systems
13:00 - 14:00	Lunch
14:00 - 15:40	SESSION 11 - Multi-Disciplinary Optimization III
15:40 - 16:00	Coffee Break
16:00 - 17:00	SESSION 12 - Multi-Disciplinary Optimization IV
17:00 - 18:00	Keynote Lecture IV Ögmundur Petersson, Airbus Defence and Space
18:00 - 18:15	Closing Ceremony

Conference Program

Tuesday, April 22nd

Keynote Lecture I – Melike Nikbay, Istanbul Technical University

1 LEVERAGING MULTI-FIDELITY SURROGATES FOR ADVANCED COMPUTATIONAL DESIGN AND MULTIDISCIPLINARY DESIGN OPTIMIZATION IN AEROSPACE ENGINEERING
Melike Nikbay

Session 1 – Design Optimization I

2 MULTI-OBJECTIVE AERODYNAMIC DESIGN OPTIMIZATION USING GRADIENT-ENHANCED SURROGATE MODELING
Emre Özkaya, Long Chen and Nicolas R. Gauger

16 COMPARISON OF CONSTRAINT MANAGEMENT STRATEGIES FOR HIGH DIMENSIONAL BAYESIAN OPTIMISATION USING CLUSTERED ACTIVE SUBSPACES
Maxime Chapron, Christophe Blondeau, Itham Salah El Din, Denis Sipp and Michel Bergmann

17 MULTIFIDELITY BAYESIAN OPTIMIZATION FOR AEROSPACE ENGINEERING DESIGN
Sihmehmet Yildiz and Melike Nikbay

Session 2 – Design Optimization II

37 SHAPE OPTIMIZATION OF AUXETIC UNIT CELLS UNDER DYNAMIC LOADING IN MACROSCOPIC COMPONENTS
Nicolas Grünfelder, Lars Kälber, Navina Waschinsky and Tim Ricken

52 AERODYNAMIC SHAPE OPTIMIZATION OF RAE-2822 AIRFOIL BY PHYSICS-INFORMED PARAMETRIC MODEL EMBEDDING
Damiano Squillace, Umberto Iemma, Domenico Quagliarella, Matteo Diez and Andrea Serani

67 OPTIMAL WING FOR LONG ENDURANCE UAV APPLICATIONS
Abhiyan Paudel, Sharanjeet Kaur, Nikhil Vijay Shende and N. Balakrishnan

84 DEFORMATION STRATEGIES IN FUSELAGE AERODYNAMIC SHAPE OPTIMIZATION USING PAYLOAD VOLUME CONSTRAINTS
Luís D. Pinheiro, Nuno M. B. Matos and André C. Marta

100 OPTIMIZING UAV WINGLETS: A MULTIOBJECTIVE APPROACH WITH SURROGATE MODELS
Mohammed Almutairi, Peter Dunning and Alireza Maheri

Session 3 – Multi-Disciplinary Optimization I

111 MULTIDISCIPLINARY DESIGN AND OPTIMIZATION OF STRUT-BRACED HYBRID-ELECTRIC AIRCRAFT UNDER MULTI-TRAJECTORY SCENARIOS
Pablo Norczyk Simon, Rauno Cavallaro, Danlin Zheng Zhang and Jānis Brižs

133 GREEN AVIATION MANUFACTURING: ADDRESSING ENVIRONMENTAL IMPACTS WITH MDO METHODOLOGIES
Shantanu Sapre, Joseph Morlier and Christian Gogu

149 HIGH-FIDELITY MULTIDISCIPLINARY OPTIMISATION PROCESSES FOR SOLVING AERO-STRUCTURAL PROBLEMS: RECENT DEVELOPMENTS FROM AN INDUSTRIAL PERSPECTIVE
Joel Brezillon, Imane Fadli, Kittinan Thanissaranon and Cian Conlan-Smith

169 MULTIDISCIPLINARY DESIGN OPTIMIZATION OF NONPLANAR LIFTING SYSTEMS FOR UNMANNED COMBAT AERIAL VEHICLES
Berkay Pirlepeli, Sihmehmet Yildiz, Hassan Aleisa, Konstantinos Kontis and Melike Nikbay

193 AIRCRAFT WING STRUCTURAL SIZING COMPUTATIONAL TOOL TAILORED FOR A COLLABORATIVE MULTIDISCIPLINARY DESIGN FRAMEWORK
Pedro F. Albuquerque, Ana M. P. Silva and André N. Pereira

Session 4 – Discipline Analysis Models I

213 BENCHMARKING SURROGATE MODELS OF VARIABLE FIDELITY FOR THE DRAG COEFFICIENT PREDICTIONS OF AIRFOIL AERODYNAMICS WITH DIFFERENT AIRFOIL PARAMETRIZATION TECHNIQUES
Ivo Zell, Maurice Zimmnau and Eike Stumpf

233 ON THE AEROACOUSTIC FEEDBACK LOOP IN SUPERSONIC IMPINGING JETS
António Rebelo, Maxime Fiore and Fernando Lau

251 INVESTIGATING THE INFLUENCE OF A TIP PROPELLER'S VERTICAL OFFSET ON A TRAILING WING'S LIFT AND DRAG DISTRIBUTIONS
Gustavo Padovany da Silva, Shamsheer S. Chauhan, Gustavo L. O. Halila, Joaquim R. R. A. Martins and João Luiz F. Azevedo

269 MISSION-INFORMED SAMPLING STRATEGY FOR AN AERODYNAMIC MULTI-FIDELITY SURROGATE MODEL
Maurice Zimmnau and Eike Stumpf

287 HIGH-FIDELITY NUMERICAL INVESTIGATIONS OF FOLDING WINGTIPS ON HIGH ASPECT RATIO WINGS
Pranesh Chandrasekaran, Emre Güler and Melike Nikbay

302 SMART WING STALL FLUTTER ALLEVIATION
Reza Moosavi

Wednesday, April 23rd

Keynote Lecture II – Joaquim Martins, University of Michigan

303 SHAPING NEXT-GENERATION AIRCRAFT WITH MULTIDISCIPLINARY DESIGN OPTIMIZATION
Joaquim R. R. A. Martins

Session 5 – Design Optimization III

304 A VELOCITY-FIELD LEVEL-SET BASED TOPOLOGY OPTIMIZATION OF STRUCTURES SUBJECTED TO AEROTHERMODYNAMICS LOADS
Apoorva Kanti, Pankil N. Mishra and Abhijit Gogulapati

316 MACH: A HIGH-FIDELITY SURROGATE MODEL DESIGNED FOR GENERALIZATION ACROSS GEOMETRIES AND PHYSICS
Chinmay Shrivastava and Noah Evers

317 COLLABORATIVE TOPOLOGY AND FIBER ORIENTATION OPTIMIZATION SIMPLIFIED VIA MULTI-RESOLUTION ANISOTROPY REPRESENTATION
Shuya Nozawa and Gokhan Serhat

Session 6 – Design Optimization IV

328 CALIBRATION OF TRANSITIONAL MODEL FOR LOW REYNOLDS NUMBER FLOWS AT HIGH ANGLES OF ATTACK USING SURROGATE-BASED OPTIMIZATION FRAMEWORK
Pranesh Chandrasekaran, Dilan Kilic and Melike Nikbay

340 PRECONDITIONING TECHNIQUES FOR GENERAL CONSTRAINED OPTIMIZATION PROBLEMS: APPLICATION TO INDUSTRIAL AEROSPACE COMPONENT DESIGN
Shouvik Bandopadhyay, François Gallard and Simone Coniglio

358 MULTI-PATCH HYBRID OPTIMIZATION OF COMPOSITE WINGS
Dario Zamani, Alfonso Pagani, Marco Petrolo and Erasmo Carrera

374 DEEP LEARNING-BASED FLOW FIELD RECONSTRUCTION FOR M219 CAVITY FLOW
Berkay Pirlepeli, Burak Berkan Bedir and Melike Nikbay

Session 7 – Multi-Disciplinary Optimization II

387 IMPACT OF ADOPTING ACTIVE FLUTTER SUPPRESSION SYSTEMS ON THE OPTIMIZATION OF A TYPICAL LONG-RANGE TRANSPORT AIRCRAFT
Elena Roncolini, Francesco Toffol and Sergio Ricci

407 A COMPARISON OF MDO BI-LEVEL FORMULATION VARIANTS WITH MDF ON A HIGH-FIDELITY AEROELASTIC WING TEST CASE
Gilberto Ruiz Jimenez, Nicolas Roussouly, François Gallard, Anne Gazaix, Joël Brezillon, Kittinan Thanissaranon, Vida Brück and Joaquim R.R.A. Martins

427 AEROSTRUCTURAL OPTIMIZATION OF HIGH ASPECT-RATIO WINGS: NAVIGATING COMPLEX TRADE-OFFS
Ousmane Sy, Joseph Morlier and Emmanuel Benard

444 EFFICIENT GRADIENT-BASED STRUCTURAL OPTIMISATION WITH MODAL AND BUCKLING CONSTRAINTS ASSISTED BY ALGORITHMIC DIFFERENTIATION
Luca Scalia, Rauno Cavallaro, Andrea Cini, Nicolas R. Gauger and Max Sagebaum

474 HIGH ASPECT-RATIO COMPOSITE WING AEROSTRUCTURAL OPTIMISATION OF A SHORT-MEDIUM RANGE AIRCRAFT
Luís M. M. Pacheco, Fabian Volle, Ögmundur Petersson and André C. Marta

Session 8 – Discipline Analysis Models II

494 EXPLOITING POST-BUCKLING NONLINEARITIES FOR GUST LOAD REDUCTION IN HIGH ASPECT-RATIO AIRCRAFT WINGS
Pedro Farinha, Francesco Toffol and Chiara Bisagni

508 SOLVING NONLINEAR PERIODIC AEROELASTIC SYSTEMS USING BOUNDARY VALUE FORMULATION AND CONFORMAL MAPPING
Anshuman Mehta, Pankil N. Mishra and Abhijit Gogulapati

527 PRELIMINARY NUMERICAL AND EXPERIMENTAL COMPARISON OF TWO TYPES OF PLA BCC LATTICE STRUCTURES MANUFACTURED VIA FDM AND SUBJECTED TO IMPACT TESTING
Giuseppe Iacolino, Calogero Orlando, Davide Tumino, Giuseppe Catalanotti, Emilio V. González, Giuseppe Mantegna and Andrea Alaimo

537 EXTENSION OF THE TWIST-KIRCHHOFF ELEMENTS TO THE BUCKLING ANALYSIS OF PLATES
Hugo A.F.A. Santos and Thomas J.R. Hughes

549 INVERSE ESTIMATION METHOD FOR THE CHARACTERISATION OF ELASTIC PROPERTIES OF ANISOTROPIC OPEN-CELL MATERIALS
Huina Mao, Xuefeng Li, Gunnar Tibert and Romain Rumpler

562 OPTIMIZING TRANSITION PROFILES FOR TAILSITTER UAVS: A GRADIENT DESCENT APPROACH
Alexandre Athayde, Alexandra Moutinho and José Raúl Azinheira

Thursday, April 24th

Keynote Lecture III – Anant Grewal, National Research Council of Canada

581 MULTIDISCIPLINARY DESIGN AND OPTIMIZATION RELATED TO NOISE AND VIBRATION: 30 YEARS OF RESEARCH AT NRC
Anant Grewal

Session 9 – Systems Engineering and Integration

583 MAGNETIC-ANOMALY POSITIONING FOR GNSS-INDEPENDENT UAV NAVIGATION
Dominik Pisarski, Rami Faraj, Lukasz Jankowski, Blazej Poplawski and Robert Konowrocki

593 ASSESSING WIND EFFECTS ON THE OPTIMAL TRAJECTORIES OF EXTENDED FORMATION FLIGHTS IN COMMERCIAL AVIATION
M. Cerezo-Magaña, Ernesto Staffetti and Alberto Olivares

607 OBSTACLE AVOIDANCE FOR FIXED-WING AIRCRAFT BY USING CONTROL BARRIER FUNCTIONS
Huriye Nur Toprak and Emre Koyuncu

Session 10 – Aerospace Design and Integrated Systems

608 EXPLORING COUPLING EFFECTS BETWEEN AIRFRAME AND PROPULSION IN PARAMETRIC MODELS FOR ELECTRIC REGIONAL AIRCRAFT WITH RANGE EXTENDER
Didier Brousset-Matheu, Eric Nguyen Van, Sebastien Defoort, Xavier Roboam and Bruno Sareni

628 HYDROGEN-POWERED AIRCRAFT DESIGN: A MULTIDISCIPLINARY APPROACH TO MINIMIZE EMISSIONS AND COST
Raúl Quibén Figueroa, Rauno Cavallaro, Andrea Cini and Manuel Soler Arnedo

645 EFFECTS OF TRAILING EDGE MODIFICATIONS ON NEGATIVE LIFT CHARACTERISTICS OF SYMMETRICAL AIRFOILS AT TRANSITIONAL LOW REYNOLDS NUMBERS
Pranesh Chandrasekaran, Emre Güler and Melike Nikbay

659 DESIGN OPTIMIZATION OF ELASTOMER MOUNTS TO ASSESS WHIRL FLUTTER PHENOMENON IN A WIND TUNNEL
Viresh Wickramasinghe, Amin Fereidooni, Eric Chen, Anant Grewal and Adam Ben Amar

678 AEROELASTIC MODELING AND DESIGN OF A FLEXIBLE BLENDED-WING-BODY (BWB) RESEARCH VEHICLE
Mário Brás, Sid Banerjee and Afzal Suleman

Session 11 – Multi-Disciplinary Optimization III

693 FREQUENCY-AWARE SURROGATE MODELING WITH SMT KERNELS FOR ADVANCED DATA FORECASTING
Nicolas Gonel, Paul Saves and Joseph Morlier

705 AN MDAO FRAMEWORK APPROACH EXTENDED TO UNSTEADY MULTIPHYSICS: KEY INGREDIENTS AND EXAMPLES
Florian Roß, Simon Ehrmanntraut, Constantin Höing, Sebastian Gottfried and Arthur Stück

719 DATA-DRIVEN ACCELERATION OF MULTIDISCIPLINARY SIMULATIONS FOR MULTI-QUERY APPLICATIONS IN AEROSPACE ENGINEERING
Susanna Baars and Ulrich Römer

735 TOWARD DISTRIBUTED AND SCALABLE MULTI-DISCIPLINARY OPTIMIZATION WITH GEMSEO
Jean-Christophe Giret, François Gallard, Antoine Dechaume, Jérôme Fasquel, Jean-François Figué and Anne Gazaix

755 AEROACOUSTIC SENSITIVITY ANALYSIS OF GEOMETRICAL PARAMETERS FOR OPEN-CAVITY FLOWS
Ramazan Kaba and Melike Nikbay

Session 12 – Multi-Disciplinary Optimization IV

775 DEVELOPMENT OF A SOUNDING ROCKET MULTIDISCIPLINARY PRELIMINARY DESIGN OPTIMISATION FRAMEWORK WITH TRAJECTORY OPTIMISATION
Alexandre M. Palaio, André C. Marta and Paulo J. S. Gil

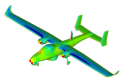
799 A STUDY ON THE MULTIDISCIPLINARY DESIGN OPTIMIZATION OF HYBRID-ROCKET PROPULSION
Joana Matos, Frederico Afonso, Girolamo Musso, Iara Figueiras, Alain Souza, Inês Ribeiro, Fernando Lau and Joseph Morlier

814 MULTI-OBJECTIVE SURROGATE-BASED OPTIMISATION OF ROCKET FIN DESIGN
José Silva, Frederico Afonso and Fernando Lau

Keynote Lecture IV – Ögmundur Petersson, Airbus Defence and Space

830 APPLYING MDO TO AIRCRAFT STRUCTURAL DESIGN AT AIRBUS DEFENCE AND SPACE; FROM PUNCH CARDS TO DIGITAL TWINS
Ögmundur Petersson

833 **Index of Authors**



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

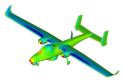
LEVERAGING MULTI-FIDELITY SURROGATES FOR ADVANCED COMPUTATIONAL DESIGN AND MULTIDISCIPLINARY DESIGN OPTIMIZATION IN AEROSPACE ENGINEERING

Melike Nikbay

Faculty of Aeronautics and Astronautics, Astronautical Engineering Department
Istanbul Technical University
ITU Ayazaga Campus, 34469 Maslak/Istanbul, Turkey
nikbay@itu.edu.tr, https://aeromdo.itu.edu.tr/

Abstract. *This keynote lecture provides an overview of advanced computational design methods and multi-fidelity surrogate modeling, along with their applications in multidisciplinary design optimization (MDO) and uncertainty quantification for aerospace engineering systems. The lecture begins with an introduction to the foundational concepts of computational design, followed by detailed discussions on MDO, including problem formulation and recent applications. Uncertainty quantification techniques are then introduced, highlighting their integration into computational design processes to account for randomness in system performance. We also briefly examine surrogate modeling techniques, emphasizing multi-fidelity approaches such as Gaussian process methods and neural networks, which enhance the efficiency and accuracy of design optimization. Next, the lecture explores relevant studies of surrogate modeling-based optimization (SBO) frameworks. The methodologies presented offer a robust framework for addressing the complexities and uncertainties inherent in modern engineering design and optimization. Additionally, the lecture presents an overview of aerospace research projects at the Istanbul Technical University AeroMDO Lab that utilize multi-fidelity surrogate modeling. Technology transfer applications between ITU and GE Vernova are demonstrated, including development of co-Kriging models for high dimensional problems, multi-fidelity deep neural networks, multi-fidelity convolutional neural networks, multi-fidelity support vector regressions, and hyperparameter optimization using high dimensional Bayesian optimization. Next, the outcomes of NATO STO AVT-331 RTG project, Goal-driven, Multi-Fidelity Approaches for Military Vehicle System-Level Design, are also presented. Furthermore, the lecture highlights applications from two major TÜBİTAK projects conducted by ITU AeroMDO Lab which are: Project 218M471, titled "Development of Multi-fidelity and Multidisciplinary Methodologies Integrating Sonic Boom, Aeroelasticity and Propulsion System for Supersonic Aircraft Design" and Project 222M019, titled "Development of Environmentally Friendly Aircrafts with High Aspect Ratio Wings via Multi-fidelity Surrogate Modeling Based Aeroelastic Optimization" are presented.*

Keywords: Aerospace, multidisciplinary design optimization, multi-fidelity models, machine learning, deep neural networks, uncertainty quantification



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTI-OBJECTIVE AERODYNAMIC DESIGN OPTIMIZATION USING GRADIENT-ENHANCED SURROGATE MODELING

Emre Özkaya, Long Chen and Nicolas R. Gauger

Chair for Scientific Computing
University of Kaiserslautern-Landau (RPTU)
Paul-Ehrlich-Strasse 34
67663 Kaiserslautern, Germany
emre.oezkaya, long.chen, nicolas.gauger@scicomp.uni-kl.de, https://scicomp.rptu.de/

Abstract. *We present a multi-objective Bayesian optimization framework that integrates a derivative-enhanced Gaussian Process Regression surrogate model to efficiently explore the Pareto front in aerodynamic design problems. The framework leverages gradient-enhanced surrogate modeling to incorporate arbitrary directional derivatives of multiple objective functions alongside their functional values, significantly improving prediction accuracy and model fidelity. By effectively capturing the underlying trends in the design space, this approach mitigates the issue of over-exploration, a common challenge in Bayesian optimization, leading to a more targeted and efficient search for optimal trade-offs. To demonstrate the effectiveness of this framework, we apply it to the shape optimization of the RAE2822 airfoil, aiming to minimize drag while maximizing lift and satisfying cross-sectional area constraint. The results show that selectively incorporating gradient information into the surrogate model leads to a substantial improvement in the convergence rate and the quality of the identified Pareto front. This highlights the potential of gradient-enhanced surrogate models in accelerating multi-objective aerodynamic design optimization while maintaining robust and accurate Pareto front approximations.*

Keywords: Pareto front, Bayesian optimization, surrogate modeling, design optimization

1 INTRODUCTION

Conflicting criteria are an inherent aspect of engineering design, where trade-offs must be made between competing objectives. This challenge has become even more pronounced with the increasing complexity of modern systems, which demand more sophisticated and comprehensive design approaches. As a result, engineering design problems frequently require balancing multiple performance criteria that inherently conflict with each other. These challenges are typically formulated either as multi-objective optimization problems or by embedding secondary objectives as constraints within a single-objective framework.

Multi-objective design optimization problems can be efficiently addressed using surrogate-based optimization (SBO) techniques, which approximate expensive objective and constraint functions with computationally cheaper models. Among these, Bayesian Optimization (BO) has emerged as a particularly promising approach due to its purely data-driven nature. BO's data-driven formulation lends it great flexibility, allowing objectives or constraints to be modified without structural changes to the underlying optimization algorithm. Furthermore, BO naturally accommodates design constraints through constraint-aware acquisition functions, making it well-suited for real-world engineering problems.

However, a known limitation of Bayesian Optimization is the lack of formal convergence guarantees, which can lead to slow convergence or excessive exploration, especially in high-dimensional or complex design spaces. With the growing availability of adjoint solvers, it is now feasible to compute gradient information efficiently, even for computationally expensive simulations. This has led to efforts to integrate gradient information into surrogate models, such as Gradient-Enhanced Kriging (GEK) [1]. While these methods offer the potential for improved model accuracy and faster convergence, they often suffer from robustness issues, particularly when gradients are noisy or inconsistently scaled, leading to degraded numerical conditioning or model fidelity.

To address these challenges, we develop a robust gradient-enhanced surrogate modeling framework that seamlessly integrates gradient information into the Bayesian optimization loop. This framework synergizes the global exploration strength of Bayesian optimization with the precision of local gradient-based search, resulting in improved convergence speed and solution quality for multi-objective design problems.

2 MULTI-OBJECTIVE DESIGN OPTIMIZATION PROBLEM

In this paper, we concentrate on a constrained design optimization problem with two objective functions:

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^D} \quad & \mathbf{f}(x) = [f_1(x), f_2(x)]^\top \\ \text{subject to} \quad & x_i^l \leq x_i \leq x_i^u, \quad i = 1, \ldots, D \\ & h_j(x) \geq 0, \quad j = 1, \ldots, m \end{aligned}$$

Here, the scalar functions $f_1(x)$ and $f_2(x)$ represent the objective functions to be minimized. The m inequality constraints $h_j(x)$ define additional restrictions on feasible designs. The vectors x^l and x^u specify the box constraints, which enforce lower and upper bounds on the design variables.

These box constraints typically arise from design requirements and ensure that the optimization process considers only physically or practically viable solutions. Tightly defined bounds are essential, as they directly influence the shape of the feasible design space. They help prevent excessive exploration of irrelevant regions and improve optimization efficiency and convergence.

In this context, the **Pareto front** represents the set of optimal trade-offs, where no objective can be improved without negatively affecting the other. A design x is considered **Pareto optimal** if and only if it is **feasible**, and there exists no other feasible design x' such that $f_1(x') \leq f_1(x)$ and $f_2(x') \leq f_2(x)$, with at least one strict inequality. The structure of the Pareto front is strongly influenced by the imposed constraints, including the box constraints, and provides valuable insight for decision-makers in selecting the most appropriate design.

The key challenges in this setting arise from the need to allocate limited computational resources efficiently while addressing the following difficulties:

- **Computationally Expensive Objective Function Evaluations**: Evaluating objective functions $f_1(x), f_2(x)$ may be highly resource-intensive, severely limiting the number of evaluations during optimization. This is particularly relevant in aerodynamic design, where computational budgets often allow for only a very limited number of evaluations.

- **High Computational Cost of Inequality Constraints**: While certain constraints can be expressed analytically, others rely on costly numerical simulations for evaluation. These constraints often arise from multiple disciplines, further increasing the complexity of the optimization problem.

- **Challenges in Gradient Evaluations**: Gradient-based methods enhance convergence by providing directional information. Adjoint solvers offer a cost-effective way to compute gradients, but they are not always available for all functions. Moreover, adjoint solutions may not be applicable to all design configurations, resulting in inconsistent gradient availability.

- **Gradient Computation for Constraints**: Some constraints permit gradient evaluations, but others, due to their complexity or formulation, make gradient computation impractical.

- **Multimodal Nature of the Objective Function**: The objective function may exhibit multiple local minima or flat regions, complicating the search for global optima.

- **Noisy or High-Frequency Behavior of the Objective Function**: Even in the absence of noise, numerical artifacts can introduce high-frequency variations, leading to unreliable gradients that misrepresent global trends.

- **Ensuring Feasibility of Design Configurations**: Identifying feasible designs can be challenging, especially when starting from an infeasible point. The optimizer must often navigate from infeasible to feasible regions.

- **Convexity of the Pareto Front**: The shape of the Pareto front significantly affects optimization difficulty. Convex fronts are easier to explore, while non-convex fronts complicate trade-offs and may lead to poor convergence or balance.

A surrogate-based optimization (SBO) approach provides an effective compromise for tackling these challenges by using computationally efficient models—such as Gaussian processes, radial basis functions, or neural networks—to approximate expensive functions. This significantly reduces the need for direct evaluations and makes SBO particularly attractive in applications with strict computational budgets, such as aerodynamic design.

To further improve SBO efficiency, we integrate gradient information into the surrogate model. This addresses the issue of over-exploration common in standard SBO by enabling more informed search directions. With the availability of adjoint solvers, gradients can be computed at a fraction of the cost of finite-difference methods. Integrating this information helps the surrogate better capture local behavior, reducing the risk of redundant sampling. As a result, the enhanced surrogate model achieves faster convergence and improved robustness, even in the presence of complex constraints or noisy functions.

In multi-objective settings, this gradient-enhanced SBO framework enables efficient approximation of the Pareto front—even when it is non-convex or high-dimensional—thus providing a scalable and practical solution for high-fidelity design optimization problems.

3 GRADIENT ENHANCED SURROGATE MODELING

Kriging, also known as Gaussian Process Regression, is a widely used surrogate modeling technique that provides an efficient way to approximate expensive objective functions. Its primary advantage lies in its ability to construct an interpolating model that not only predicts function values but also quantifies uncertainty in the predictions. The Kriging model is built upon a probabilistic framework, where the function is assumed to be a realization of a Gaussian process, characterized by a mean function and a covariance structure defined by a correlation function. The choice of correlation function, such as the widely used Gaussian radial basis function, governs how training points influence predictions at new locations.

In this section, we do not provide a full introduction to the Kriging method for brevity. Instead, we directly present its mathematical formulation and the proposed gradient-enhanced extension. For a comprehensive introduction to Kriging and its applications in surrogate modeling, the interested reader may refer to Forrester et al. [2] for practical implementations in engineering design.

The first step in surrogate modeling is the Design of Experiment (DoE), which involves strategically selecting a set of sample points in the design space to serve as the training data for the surrogate model. A well-planned DoE ensures that the surrogate captures the essential features of the objective function while minimizing the number of expensive function evaluations. Common DoE strategies include Latin Hypercube Sampling (LHS) and Sobol sequences, each offering a balance between exploration and exploitation of the design space.

Assuming the training data is available from the DoE, we can introduce the Kriging estimator:

$$\tilde{f}(\mathbf{x}) = \beta_0 + \mathbf{r}^T \mathbf{R}^{-1}(\mathbf{y} - \beta_0 \mathbf{1}). \tag{1}$$

Here, $\mathbf{x}$ is the vector of optimization parameters as introduced in the generic formulation. The vector $\mathbf{y} \in \mathbb{R}^N$ contains the function values obtained from N samples in the DoE. The vector $\mathbf{r}$ represents the point-wise correlations between the input vector $\mathbf{x}$ and the training samples $\mathbf{x}^{(i)}, i = 1, \ldots, N$:

$$\mathbf{r} = [\phi(\mathbf{x}^{(1)}, \mathbf{x}), \phi(\mathbf{x}^{(2)}, \mathbf{x}), \ldots, \phi(\mathbf{x}^{(N)}, \mathbf{x})]^T. \tag{2}$$

The matrix $\mathbf{R}$ in Eq. (1) is the correlation matrix:

$$\mathbf{R} = \begin{bmatrix} \phi(\mathbf{x}^{(1)}, \mathbf{x}^{(1)}) & \phi(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}) & \ldots & \phi(\mathbf{x}^{(1)}, \mathbf{x}^{(N)}) \\ \phi(\mathbf{x}^{(2)}, \mathbf{x}^{(1)}) & \phi(\mathbf{x}^{(2)}, \mathbf{x}^{(2)}) & \ldots & \phi(\mathbf{x}^{(2)}, \mathbf{x}^{(N)}) \\ \vdots & \vdots & \ddots & \vdots \\ \phi(\mathbf{x}^{(N)}, \mathbf{x}^{(1)}) & \phi(\mathbf{x}^{(N)}, \mathbf{x}^{(2)}) & \ldots & \phi(\mathbf{x}^{(N)}, \mathbf{x}^{(N)}) \end{bmatrix}. \tag{3}$$

The function $\phi(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$, which appears in both the vector $\mathbf{r}$ and the matrix $\mathbf{R}$, is the correlation function. It measures how strongly a sample $\mathbf{x}^{(i)}$ correlates with another sample $\mathbf{x}^{(j)}$. The function takes values between 0 (no correlation) and 1 (perfect correlation). A typical choice for the correlation function is the Gaussian radial basis function:

$$\phi(\mathbf{x}^{(i)}, \mathbf{x}) = e^{-\sum_{j=1}^{D} \theta_j (x_j^{(i)} - x_j)^2}. \tag{4}$$

The Gaussian radial basis function is an even function centered around the ith sample $\mathbf{x}^{(i)}$. The vector $\boldsymbol{\theta}$ in Eq. (4) is a hyper-parameter vector that controls how rapidly the correlation decreases as we move away from a sample.

We can conceptualize the Kriging model as the sum of two separate models: $\tilde{f}(\mathbf{x}) = \tilde{f}_1(\mathbf{x}) + \tilde{f}_2(\mathbf{x})$. The term β_0 represents the constant model $\tilde{f}_1(\mathbf{x})$, where the estimate is simply a scalar value. This term dominates the estimate if $\mathbf{x}$ is far from all samples in the training data (i.e., when the vector $\mathbf{r}$ has very small values). Since the samples are correlated, the best linear unbiased estimate for β_0 is given by

$$\beta_0 = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}},$$

where $\mathbf{1}$ denotes the vector of ones, $\mathbf{1} = [1, 1, \ldots, 1]^T \in \mathbb{R}^N$.

The second term in Eq. (1) resembles the weighted sum formulation used in radial basis function networks, given by:

$$\tilde{f}_2(\mathbf{x}) = \sum_{i=1}^{N} w_i \phi_i(\mathbf{x}), \tag{5}$$

where $\phi_i(\mathbf{x}) = \phi(\mathbf{x}, \mathbf{x}^{(i)})$ represents the correlation function evaluated between $\mathbf{x}$ and the ith sample.

The weight vector $\mathbf{w}$ is obtained by solving the following system of equations:

$$\tilde{f}_2(\mathbf{x}^{(i)}) = y_i - \beta_0, \quad i = 1, \ldots, N, \tag{6}$$

where the contribution of the first model, β_0, is subtracted from the observed function values y_i in order to fit the second model $\tilde{f}_2$. This leads to the linear system:

$$\mathbf{R}\mathbf{w} = \mathbf{y} - \beta_0 \mathbf{1}, \tag{7}$$

from which the weight vector can be explicitly computed as:

$$\mathbf{w} = \mathbf{R}^{-1}(\mathbf{y} - \beta_0 \mathbf{1}), \tag{8}$$

yielding the result consistent with Eq. 1, as expected.

The second model $\tilde{f}_2(\mathbf{x}^{(i)})$ has a greater influence on the prediction, particularly when $\mathbf{x}$ is close to the training samples—that is, when it exhibits a high correlation with some or all points in the training dataset.

To incorporate gradient information into the surrogate, we focus on enhancing the second model, $\tilde{f}_2(\mathbf{x}^{(i)})$. This is achieved by augmenting the basis with derivatives of the original basis functions, as proposed in [3]. For completeness, we briefly summarize the underlying concept. The central idea is to enrich the model using directional derivatives of the basis function ϕ_i:

$$\mathbf{d} \cdot \nabla \phi_i(\mathbf{x}) = 2\,\mathbf{d} \cdot \begin{bmatrix} \mathbf{x}_1^{(i)} - \mathbf{x}_1 \\ \mathbf{x}_2^{(i)} - \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_D^{(i)} - \mathbf{x}_D \end{bmatrix} \phi_i(\mathbf{x}), \tag{9}$$

where $\mathbf{d}$ represents the direction in which the derivative is taken (row vector). When gradient information is available at training points, a natural and effective choice for $\mathbf{d}^{(i)}$ is the normalized gradient vector itself:

$$\mathbf{d}^{(i)} = \left(\frac{\nabla f(\mathbf{x}^{(i)})}{\|\nabla f(\mathbf{x}^{(i)})\|} \right)^{\top} \tag{10}$$

This choice aligns the basis function with the direction of steepest ascent (or descent), thereby improving the model's ability to capture local behavior around each training point. Given that our goal is optimization, constructing the surrogate to emphasize the dominant direction of change is both intuitive and advantageous.

The gradient-enhanced surrogate model is constructed by augmenting the original basis with directional derivatives of the basis functions. The extended model can be expressed as:

$$\tilde{f}(\mathbf{x}) = \tilde{f}_1(\mathbf{x}) + \sum_{i=1}^{N} w_i \phi_i(\mathbf{x}) + \sum_{i=1}^{N} w_{N+i}\, \mathbf{d}^{(i)} \cdot \nabla \phi_i(\mathbf{x}), \tag{11}$$

where $\mathbf{w}$ is the vector of model weights, now of length $2N$ due to the inclusion of the differentiated basis functions. To determine these weights, we impose interpolation conditions that ensure the model matches both the function values and directional derivatives at all training points. Specifically, the model must satisfy:

$$\tilde{f}(\mathbf{x}^{(i)}) = f(\mathbf{x}^{(i)}), \quad i = 1, \ldots, N, \tag{12}$$

and

$$\mathbf{d}^{(i)} \cdot \nabla \tilde{f}(\mathbf{x}^{(i)}) = \mathbf{d}^{(i)} \cdot \nabla f(\mathbf{x}^{(i)}), \quad i = 1, \ldots, N. \tag{13}$$

These conditions lead to a linear system of size $2N \times 2N$, written compactly as:

$$\dot{\mathbf{R}}\,\dot{\mathbf{w}} = \dot{\mathbf{y}}, \tag{14}$$

where $\dot{\mathbf{R}}$ is the extended system matrix incorporating both the original and differentiated basis functions, and $\dot{\mathbf{y}}$ is the extended right-hand side vector that includes function values and directional derivatives. The matrix $\dot{\mathbf{R}}$ takes the form:

$$\dot{\mathbf{R}} = \begin{bmatrix} \phi_1(\mathbf{x}^{(1)}) & \dots & \phi_N(\mathbf{x}^{(1)}) & \mathbf{d}^{(1)} \cdot \nabla\phi_1(\mathbf{x}^{(1)}) & \dots & \mathbf{d}^{(1)} \cdot \nabla\phi_N(\mathbf{x}^{(1)}) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \phi_1(\mathbf{x}^{(N)}) & \dots & \phi_N(\mathbf{x}^{(N)}) & \mathbf{d}^{(N)} \cdot \nabla\phi_1(\mathbf{x}^{(N)}) & \dots & \mathbf{d}^{(N)} \cdot \nabla\phi_N(\mathbf{x}^{(N)}) \\ \mathbf{d}^{(1)} \cdot \nabla\phi_1(\mathbf{x}^{(1)}) & \dots & \mathbf{d}^{(1)} \cdot \nabla\phi_N(\mathbf{x}^{(1)}) & \mathbf{d}^{(1)} \cdot \nabla(\mathbf{d}^{(1)} \cdot \nabla\phi_1(\mathbf{x}^{(1)})) & \dots & \mathbf{d}^{(1)} \cdot \nabla(\mathbf{d}^{(1)} \cdot \nabla\phi_N(\mathbf{x}^{(1)})) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{d}^{(N)} \cdot \nabla\phi_1(\mathbf{x}^{(N)}) & \dots & \mathbf{d}^{(N)} \cdot \nabla\phi_N(\mathbf{x}^{(N)}) & \mathbf{d}^{(N)} \cdot \nabla(\mathbf{d}^{(N)} \cdot \nabla\phi_1(\mathbf{x}^{(N)})) & \dots & \mathbf{d}^{(N)} \cdot \nabla(\mathbf{d}^{(N)} \cdot \nabla\phi_N(\mathbf{x}^{(N)})) \end{bmatrix}$$

and the corresponding right-hand side vector is:

$$\dot{\mathbf{y}} = \begin{bmatrix} f(\mathbf{x}^{(1)}) - \beta_0 \\ \vdots \\ f(\mathbf{x}^{(N)}) - \beta_0 \\ \mathbf{d}^{(1)} \cdot \nabla f(\mathbf{x}^{(1)}) \\ \vdots \\ \mathbf{d}^{(N)} \cdot \nabla f(\mathbf{x}^{(N)}) \end{bmatrix}.$$

The proposed formulation offers several advantages over the classical GEK approach. First, the size of the coefficient matrix grows more moderately compared to alternative formulations, resulting in improved computational efficiency. Second, the use of a single differentiated basis function per training point leads to a set of basis functions that are nearly orthogonal, which enhances numerical stability and model robustness.

4 OPTIMIZATION FRAMEWORK

4.1 Bayesian Optimization

Bayesian Optimization is a powerful and versatile framework for optimizing expensive black-box functions that lack analytical expressions and gradient information. It builds a probabilistic surrogate model—most commonly a Kriging model (also known as Gaussian Process) to estimate the true objective function and quantify uncertainty in unexplored regions. By balancing exploration and exploitation through an acquisition function, such as Expected Improvement (EI), BO strategically selects new evaluation points to efficiently approach the global optimum. One of the most well-known implementations of BO is the Efficient Global Optimization (EGO) algorithm introduced by Jones et al. [4], which combines GP modeling with the EI criterion to iteratively refine the surrogate and guide the search process. This approach has been successfully applied in engineering design, machine learning hyperparameter tuning, and other domains where function evaluations are costly.

Bayesian Optimization typically begins with a Design of Experiments (DoE) step, where a set of input points is selected using space-filling strategies such as Latin Hypercube Sampling (LHS) or Sobol sequences. These initial samples are evaluated using the simulation model to create a starting dataset. A surrogate model is then trained on this

data, with its hyperparameters estimated via *Maximum Likelihood Estimation (MLE)*. When Kriging is used as the surrogate model, it not only predicts the mean objective value at new inputs but also provides uncertainty estimates. This predictive uncertainty is crucial for defining an acquisition function. A common choice in Bayesian Optimization is the Expected Improvement (EI), which quantifies the expected gain from evaluating a new point $\mathbf{x}$ compared to the current best value f_{best}. Given the predicted mean $\mu(\mathbf{x})$ and standard deviation $\sigma(\mathbf{x})$, the EI is given by:

$$\text{EI}(\mathbf{x}) = (\mu(\mathbf{x}) - f_{\text{best}})\Phi(Z) + \sigma(\mathbf{x})\phi(Z), \quad \text{where} \quad Z = \frac{\mu(\mathbf{x}) - f_{\text{best}}}{\sigma(\mathbf{x})} \tag{15}$$

Here, $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively. This expression captures a balance between exploitation (favoring regions where $\mu(\mathbf{x})$ is high) and exploration (favoring regions where $\sigma(\mathbf{x})$ is large). The next sample to be evaluated is obtained by maximizing the EI function, and the resulting data is added to the training set for the next iteration. This loop—surrogate model update, EI maximization, and objective function evaluation via simulation—continues iteratively, efficiently guiding the search toward the global optimum while minimizing the number of simulations.

Maximizing the EI acquisition function is a crucial step in Bayesian Optimization, as it determines where the next expensive function evaluation will occur. However, the EI landscape is often non-convex and highly multi-modal, especially in high-dimensional problems or when the surrogate model exhibits strong uncertainty in multiple regions. As a result, global optimization methods are typically employed to find the maximizer of EI.

4.2 Extension for Multi-Objective Optimization

Bayesian Optimization can be extended to *multi-objective optimization (MBO)*, where the goal is to simultaneously optimize several conflicting objectives while minimizing the number of expensive function evaluations. In this setting, the solution is not a single optimum but a set of *Pareto-optimal solutions* that represent trade-offs between objectives. One widely used approach is to extend acquisition functions to the multi-objective case—for example, the *Expected Hypervolume Improvement (EHVI)* [5] or *Expected Volume Gain (EVG)* [6], which estimate the expected increase in the dominated volume of the objective space (i.e., the hypervolume) upon evaluating a new point. These methods require modeling each objective with its own Gaussian Process and computing the improvement in terms of Pareto dominance. Alternatively, *scalarization-based approaches* [7], such as the *ε-constraint method* [8], convert the multi-objective problem into a series of constrained single-objective problems by optimizing one objective while treating the others as inequality constraints. This allows the use of standard BO techniques while exploring different regions of the Pareto front. Both approaches aim to efficiently approximate the Pareto front with minimal evaluations, making MBO a powerful tool for complex design and decision-making problems.

The ε-constraint method is a widely used scalarization technique for handling multi-objective optimization problems. Instead of optimizing all objectives simultaneously, this method focuses on a single primary objective while converting the remaining objectives into constraints with specified thresholds. For example, consider a two-objective minimization problem with objective functions $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$. Using the ε-constraint

method, the problem is reformulated as:

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^D} \quad & f_1(\mathbf{x}) \\ \text{subject to} \quad & f_2(\mathbf{x}) \leq \varepsilon \end{aligned} \tag{16}$$

Here, ε is a user-defined threshold that controls the acceptable level of the second objective $f_2(\mathbf{x})$. By varying ε systematically across multiple optimization runs, different trade-off solutions can be obtained, effectively sampling the Pareto front. In a Bayesian Optimization setting, surrogate models are used to model both objectives, and constraint-aware acquisition functions are employed to guide the search within the feasible region. The ε-constraint method is especially attractive due to its simplicity and compatibility with the standard single-objective BO framework, making it a practical choice for many engineering and design problems.

5 RESULTS ON A BENCHMARK TESTCASE

We first evaluate our approach on the welded beam design problem, a widely used benchmark in structural optimization. This problem is particularly suitable for testing optimization algorithms due to its conflicting objectives and multiple engineering constraints, making it an ideal case for assessing the effectiveness of our method.

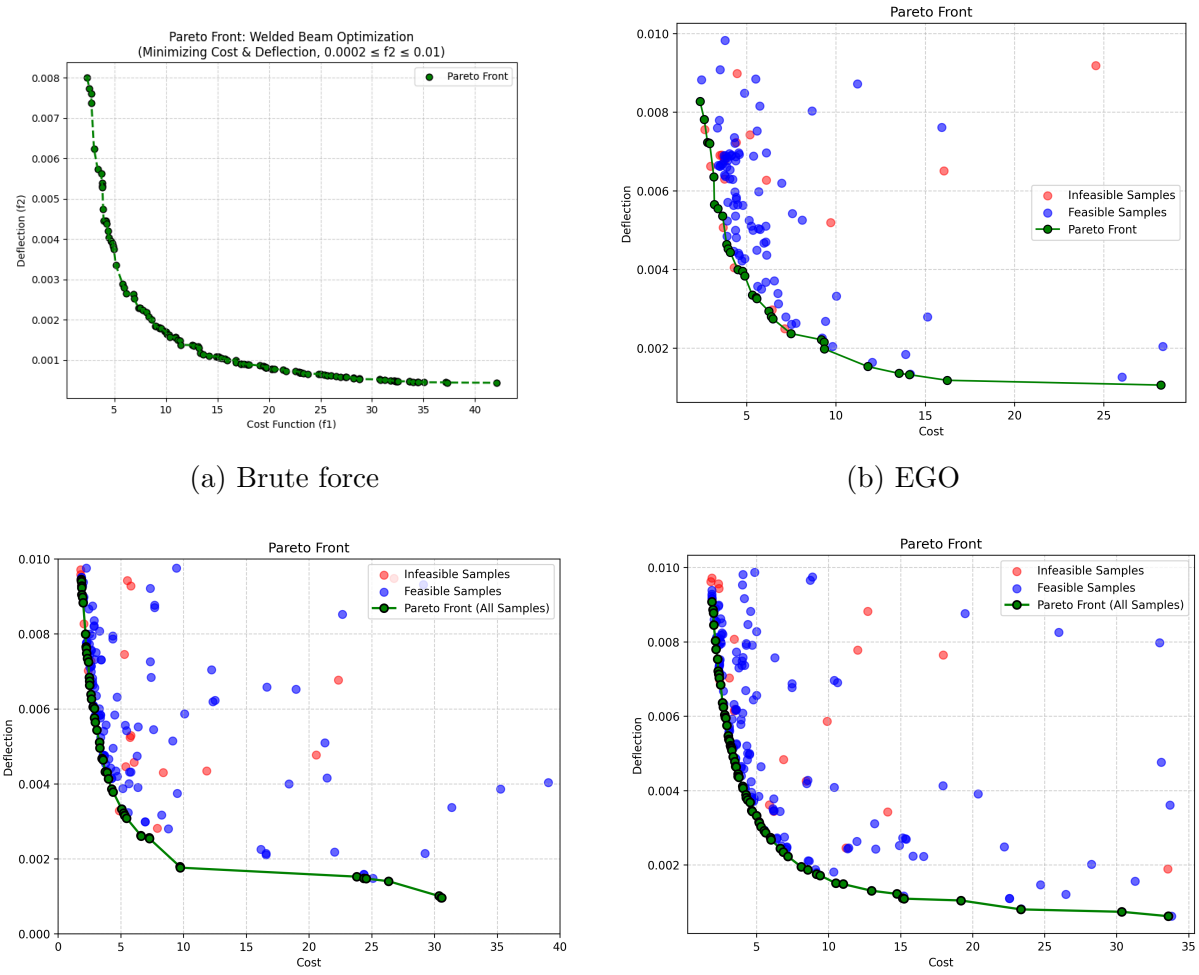
The welded beam problem is a classic test case in multi-objective structural optimization. It involves designing a welded beam that can support a specified load while minimizing two conflicting objectives: the fabrication cost and the maximum end deflection of the beam. The beam is assumed to be welded to a support at one end and subjected to a transverse load at the other end. The design variables typically include the weld thickness, weld length, height of the beam, and width of the beam.

The optimization is subject to several nonlinear constraints such as limits on shear stress, bending stress and buckling load to ensure the structural feasibility and safety of the design. A MATLAB tutorial demonstrating multi-objective optimization on this benchmark is available on the MathWorks website[1].

Figure 1 presents the results of the multi-objective optimization for the welded beam problem. Subfigure 1a shows the Pareto front obtained using a brute-force approach with 1,000,000 random function evaluations. This solution serves as the ground truth for comparisons. In Subfigure 1b, the Pareto front obtained using the gradient-free EGO method is depicted. This optimization was performed using only 350 samples—100 for the initial Design of Experiments (DoE) and 250 during the optimization phase. The EGO method demonstrates good performance, closely approximating the true Pareto front with a relatively small number of evaluations.

Subfigure 1c shows the Pareto front obtained using the gradient-enhanced EGO (GE-EGO) method, using the same total number of samples as EGO, but with only 43 additional gradient evaluations. It can be observed that GE-EGO significantly improves the resolution of the Pareto front and generally yields higher-quality designs. Finally, in Subfigure 1d, the GE-EGO method is applied again, but this time using gradient information only for the first objective (fabrication cost), while omitting gradients for the second objective (deflection). Interestingly, the resulting Pareto front does not show a significant degradation in quality, suggesting that for some problems, omitting gradient

[1] `https://de.mathworks.com/help/gads/multiobjective-optimization-welded-beam.html`



(a) Brute force

(b) EGO

(c) GE-EGO

(d) GE-EGO (only first objective gradient)

Figure 1: Welded beam multi objective optimization results

evaluations for certain objectives may still lead to competitive results, thereby reducing the computational cost.

6 AERODYNAMIC SHAPE OPTIMIZATION RESULTS

The **gradient-enhanced Efficient Global Optimization (GE-EGO)** framework is applied to a benchmark aerodynamic shape optimization problem from the SU2 repository[2].

`SU2` is an open-source software suite for the simulation of fluid flows and other continuum mechanics problems. Initially developed at Stanford University, SU2 is designed to solve partial differential equations using finite volume and finite element methods on unstructured meshes. It supports a wide range of applications, including compressible and incompressible flows, turbulence modeling, heat transfer, and shape optimization. Thanks to its modular architecture and Python interface, SU2 is particularly well-suited for high-fidelity simulations in design optimization workflows.

Using SU2, the goal is to perform multi-objective shape optimization of the RAE2822 airfoil under transonic, viscous, turbulent flow conditions. Specifically, the objective is

[2] `https://su2code.github.io/tutorials/Turbulent_2D_Constrained_RAE2822/`

to simultaneously minimize the drag coefficient (C_D) and maximize the lift coefficient (C_L). To fit within the optimization framework, the second objective is reformulated as the negative lift, i.e., $-C_L$. Additionally, a key geometric constraint in airfoil shape optimization is the area constraint, which ensures that the cross-sectional area of the airfoil remains above a specified threshold. In this study, the constraint is formulated such that the area of the optimized airfoil must be greater than or equal to the cross-sectional area of the reference RAE2822 airfoil (0.0778). This condition helps preserve desirable structural characteristics while still allowing for shape modifications aimed at aerodynamic performance improvement.

The optimization problem can be formulated as:

$$\min_{x \in \mathbb{R}^{38}} \left[C_D(x),\, -C_L(x) \right] \quad \text{subject to} \quad \text{area}(x) \geq 0.0778 \tag{17}$$

Here, x represents the coefficients of Hicks-Henne basis functions used to deform the airfoil surface. A total of 38 basis functions are employed, with box constraints imposed on the design variables ($-0.002 \leq x_i \leq 0.002$ for $i = 1, \ldots, 38$) to prevent unrealistic geometries. These basis functions also ensure that the leading and trailing edges of the airfoil remain fixed.

The objective is to identify a set of Pareto-optimal solutions that reflect the trade-off between minimizing drag and maximizing lift, while maintaining the airfoil's cross-sectional area above a given reference value.

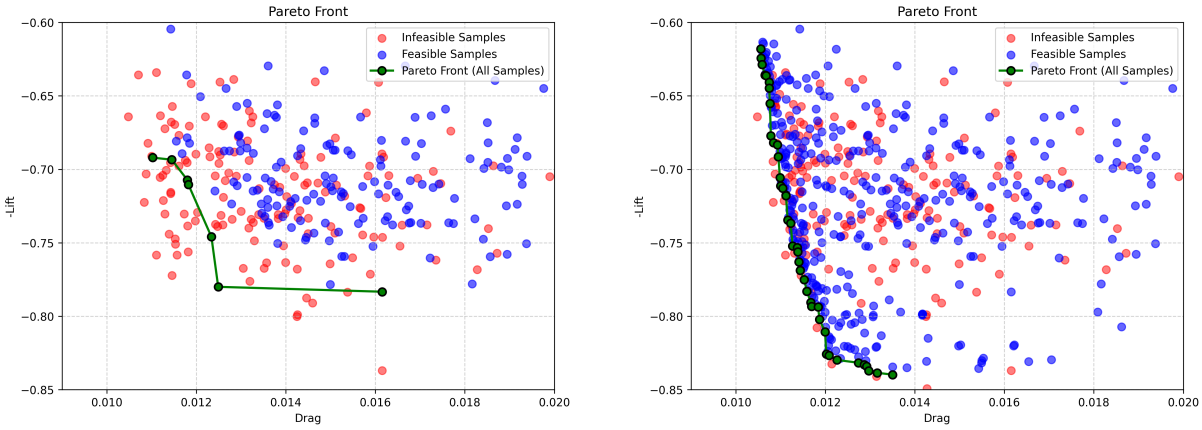
For gradient evaluations, GE-EGO utilizes the discrete adjoint solver in SU2 [9], selected for its efficiency and robustness in computing gradient vectors. Compared to the continuous adjoint method, the discrete variant provides greater reliability, supporting accurate and efficient optimization in complex aerodynamic scenarios.

For the Design of Experiments (DoE), a total of 300 samples are generated using the Latin Hypercube Sampling (LHS) strategy. During this initial phase, only the primal solver is used to evaluate the samples. The gradient vector is computed solely for the Pareto optimal samples, since most samples in this phase are far from optimal and do not benefit from gradient information.

The optimization is subject to a strict computational budget of 300 function evaluations, meaning the CFD solver may be called at most 300 times. During the optimization, gradients are evaluated selectively—only when a candidate point shows significant improvement in the objective space. This policy minimizes reliance on the adjoint solver.

The optimization algorithm combines global and local search strategies. Local refinement is performed using a trust-region approach around the actual Pareto optimal solution for the specified value of ε. Each local search allocates a fixed number of samples and gradually shrinks the trust region. If a local search fails due to unreliable gradient information (e.g., noise or missing values), the algorithm safely falls back to global search based on the classical EGO procedure, ensuring robustness throughout the optimization process.

Figure 2 illustrates the evolution of the Pareto front throughout the optimization process. The subplot on the left shows the Pareto front obtained at the end of the initial Design of Experiments (DoE) phase, after evaluating 300 samples. In contrast, the subplot on the right presents the Pareto front at the end of the complete optimization process, after 600 samples. It can be observed that the multi-objective GE-EGO method successfully generates a well-resolved Pareto front, significantly improving both objectives and



(a) Pareto front after DoE (300 samples) (b) Pareto front at final stage (600 samples)

Figure 2: Evolution of the Pareto front during the multi-objective optimization process.

demonstrating a balanced exploration and exploitation of the design space. In both subplots, red points indicate infeasible samples that violate the geometric constraint, while blue points correspond to feasible solutions.

Figure 3 compares the Pareto fronts obtained using the gradient-enhanced EGO (GE-EGO) and the gradient-free EGO methods. It can be observed that incorporating gradient information significantly improves both objective functions across almost the entire Pareto front. Moreover, GE-EGO achieves a higher resolution of the front, capturing more trade-off designs. In this study, GE-EGO identified 41 Pareto-optimal solutions, whereas the standard EGO approach yielded only 26. Notably, GE-EGO required a total of 13 adjoint evaluations. Considering that each adjoint run has approximately the same computational cost as a primal run, the overall computational overhead of GE-EGO is estimated to be only about 2% higher than that of the gradient-free EGO approach.

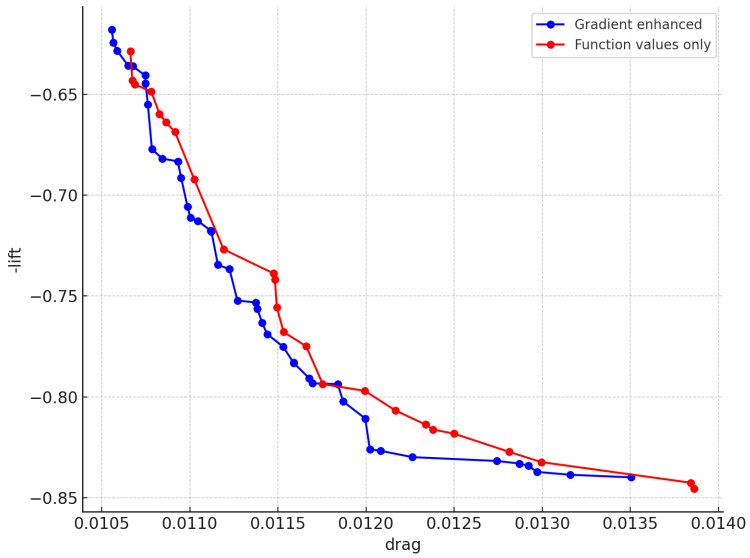


Figure 3: Comparison between EGO and GE-EGO

7 CONCLUSIONS

In this work, we presented a multi-objective Bayesian optimization framework that incorporates a gradient-enhanced Gaussian Process surrogate model to efficiently explore the Pareto front in aerodynamic and structural design problems. By selectively integrating directional derivative information, the proposed method improves surrogate model accuracy, enables faster convergence, and reduces the number of expensive objective function evaluations.

The approach was validated on two benchmark problems: the welded beam design problem and the aerodynamic shape optimization of the RAE2822 airfoil. In both cases, the gradient-enhanced EGO (GE-EGO) method outperformed the classical EGO algorithm in terms of Pareto front resolution and solution quality, while maintaining a low computational overhead. Notably, the results demonstrate that incorporating gradient information—even partially—can significantly improve optimization performance without substantially increasing cost.

The proposed surrogate model formulation also addresses common robustness issues observed in classical gradient-enhanced Kriging, offering a more stable and scalable alternative for high-dimensional design spaces. Furthermore, the framework naturally integrates with existing Bayesian optimization workflows and supports constraint handling, making it highly suitable for practical engineering applications.

REFERENCES

[1] Z.-H. Han, S. Görtz, and R. Zimmermann. Improving variable-fidelity surrogate modeling via gradient-enhanced kriging and a generalized hybrid bridge function. *Aerospace Science and Technology*, 25(1):177–189, 2013. ISSN 1270-9638. doi:https://doi.org/10.1016/j.ast.2012.01.006. URL `https://www.sciencedirect.com/science/article/pii/S127096381200017X`.

[2] A. Forrester, A. Sobester, and A. Keane. *Exploring and Exploiting a Surrogate*, chapter 3, pages 77–107. John Wiley Sons, Ltd, 2008. ISBN 9780470770801. doi:https://doi.org/10.1002/9780470770801.ch3. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470770801.ch3`.

[3] E. Özkaya, J. Rottmayer, and N. R. Gauger. *Gradient Enhanced Surrogate Modeling Framework for Aerodynamic Design Optimization.* doi:10.2514/6.2024-2670. URL `https://arc.aiaa.org/doi/abs/10.2514/6.2024-2670`.

[4] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 12 1998. doi:10.1023/A:1008306431147.

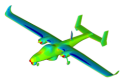
[5] I. Couckuyt, D. Deschrijver, and T. Dhaene. Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *Journal of Global Optimization*, 60:575–594, 2014.

[6] P. Feliot, J. Bect, and E. Vazquez. A bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1):97–133, 2017.

[7] G. Eichfelder. *Adaptive scalarization methods in multiobjective optimization*, volume 436. Springer, 2008.

[8] V. Chankong and Y. Y. Haimes. *Multiobjective decision making: theory and methodology*. Courier Dover Publications, 2008.

[9] T. Albring, M. Sagebaum, and N. Gauger. *A Consistent and Robust Discrete Adjoint Solver for the SU2 Framework—Validation and Application*, pages 77–86. 03 2016. ISBN 978-3-319-27278-8. doi:10.1007/978-3-319-27279-5_7.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

COMPARISON OF CONSTRAINT MANAGEMENT STRATEGIES FOR HIGH DIMENSIONAL BAYESIAN OPTIMISATION USING CLUSTERED ACTIVE SUBSPACES

Maxime Chapron[1*,3], Christophe Blondeau[1], Itham Salah El Din[2], Denis Sipp[2] and Michel Bergmann[3]

1: Department of Aerodynamics, Aeroelasticity, Acoustics (DAAA)
ONERA, Institut Polytechnique de Paris
92320, Châtillon, France
maxime.chapron@onera.fr

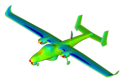
2: Department of Aerodynamics, Aeroelasticity, Acoustics (DAAA)
ONERA, Institut Polytechnique de Paris
92190 Meudon, France

3: MEMPHIS Team
INRIA Bordeaux
200 avenue de la Vieille Tour, 33405, Talence, France

Abstract. *The increasing number of design variables in industrial-scale aerodynamic shape optimisation has prompted a re-evaluation of state-of-the-art optimisation strategies. Bayesian Optimisation (BO) is well-suited for cases where multimodality challenges gradient-based methods and heuristic approaches are computationally intractable. The curse of dimensionality in Bayesian Optimisation can be mitigated using dimension reduction techniques such as Active Subspaces (AS), which enable accurate surrogate modelling within physics-informed reduced-order bases. The Clustered Active Subspaces (CAS) method, previously introduced by these authors, improves upon AS by employing a piecewise approximation to better capture complex high-dimensional objective functions. CAS has also been successfully integrated into a reduced-dimension BO framework, facilitating global optimisation of aerodynamic designs.*

In this work, we compare multiple constraint management strategies for the Clustered Active Subspaces framework, assessing their impact on optimisation performance. Specifically, we examine how expensive nonlinear constraints such as the lift coefficient can be efficiently handled within the reduced-dimension Bayesian Optimisation. The proposed approaches are evaluated on high-dimensional analytical test functions and demonstrated on the constrained global optimisation of a NACA0012 airfoil in transonic viscous flow.

Keywords: Optimization, Dimension Reduction, Clustering, Bayesian Optimization, Active Subspaces



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTIFIDELITY BAYESIAN OPTIMIZATION FOR AEROSPACE ENGINEERING DESIGN

Sihmehmet Yildiz[1*], Melike Nikbay[2]

1: Aeronautical Engineering Department
Istanbul Technical University
34469, Istanbul, Türkiye
yildizsih@itu.edu.tr

2.: Astronautical Engineering Department
Istanbul Technical University
34469, Istanbul, Türkiye
nikbay@itu.edu.tr

Abstract. *Multifidelity Bayesian Optimization (MFBO) has emerged as a powerful tool for tackling computationally expensive engineering design problems by leveraging variable-fidelity approximations to guide the optimization process. This study investigates the application of MFBO to efficiently solve such problems, demonstrating its superior performance compared to single-fidelity Bayesian Optimization (SFBO). As engineering design generally relies on computationally intensive high-fidelity (HF) simulations, MFBO mitigates this burden by integrating information from models of varying fidelity. The methodology uses advanced surrogate modeling technique, specifically Autoregressive Gaussian Processes, to capture the relationships between design variables and performance metrics across different fidelity levels. It also investigates the impact of different acquisition functions on optimization trajectories, with a particular focus on information-theoretic approaches such as Multifidelity Max-value Entropy Search (MF-MES) and Multifidelity General-purpose Information-Based Bayesian OptimizatioN (MF-GIBBON). The performance of these MFBO methods is comparatively evaluated against their single-fidelity counterparts (SF-MES, SF-GIBBON) using both analytical benchmark functions (rosenbrock and shifted-rotated rastrigin) and a constrained aerospace shape optimization problem (a re-entry vehicle). The results highlight that MFBO significantly reduces computational cost, achieves faster convergence, and identifies superior final solutions, particularly for the complex engineering problem and the high-dimensional analytical cases. These findings contribute to research on efficient optimization strategies by highlighting the potential of MFBO to advance aerospace engineering design through an effective balance between accuracy and computational efficiency.*

Keywords: Multifidelity, Bayesian Optimization, Surrogate Modeling, Acquisition Function

1 INTRODUCTION

The design and optimization of complex engineering systems present considerable computational challenges. High performance, often with optimal results, can only be achieved by using high-fidelity (HF) methods such as Computational Fluid Dynamics (CFD), Finite Element Analysis (FEA), or physical experiments. These methods are extremely time-consuming and require significant computational resources [1–3]. As a result, it becomes computationally infeasible to conduct the design space searches or perform optimization studies using HF models within a reasonable computational time. Frequently, HF simulations are treated as black-box models and their gradient information is often unavailable or unreliable [4]. This kind of computational bottleneck necessitates the development of advanced methodologies to accelerate the optimization process while maintaining the quality of the solutions.

Surrogate-Based Optimization (SBO) is an effective approach that substitutes a high-fidelity (HF) model with a more efficient approximation known as a surrogate model or metamodel [3]. Among various SBO techniques, Bayesian Optimization (BO) has become popular for global optimization of costly black-box functions[5, 6]. BO utilizes a probabilistic surrogate model, most often a Gaussian process (GP), to manage the uncertainty related to predictions of the objective function. This probabilistic belief is continuously updated as new data points are obtained. The choice of the next point to evaluate is determined by an acquisition function, which strategically balances the need to explore uncertain parts of the design space and to exploit areas where the surrogate model indicates high potential for performance [6]. This goal-driven learning approach allows Bayesian optimization to effectively explore the design space and find optimal solutions while minimizing the number of HF model evaluations needed.

In standard single-fidelity Bayesian Optimization (BO), the approach typically relies solely on high-fidelity (HF) evaluations. However, models or simulations in various engineering fields are often available at different levels of fidelity, each varying in accuracy and computational cost [1, 2]. Lower-fidelity (LF) models, such as those that use coarser grids, simplified physics, or empirical correlations, are less computationally expensive but also less accurate compared to high-fidelity models. Multifidelity Optimization (MFO) aims to effectively utilize these various sources of information [2]. Specifically, Multifidelity Bayesian Optimization (MFBO) combines the concepts of BO with principles of MFO [1, 2]. It incorporates multifidelity surrogate models that learn the relationships between different fidelity levels, typically using methods like co-kriging or autoregressive schemes [7, 8]. Specialized multifidelity acquisition functions are employed to make decisions not just on where to sample next, but also at which fidelity level to perform the sampling [2]. By judiciously querying cheaper LF models to gain insights about the objective function landscape and guide the more selective querying of the expensive HF model, MFBO aims to achieve significant computational savings compared to single-fidelity BO [1, 2].

This paper investigates the performance of MFBO methods, with a particular focus on information-theoretic acquisition functions. Specifically, we compare the Multifidelity Max-value Entropy Search (MF-MES) [9] and the Multifidelity General-purpose Information-Based Bayesian OptimizatioN (MF-GIBBON) [10] against their single-fidelity counterparts. The evaluation is conducted using both well-established analytical benchmark functions [11] to examine fundamental optimization behavior, and a realistic, constrained aerospace shape optimization problem concerning a re-entry vehicle [12], pro-

viding insight into practical engineering applicability. Performance is assessed in terms of convergence rate, final solution quality, and computational efficiency, particularly the ability of each method to strategically allocate computational resources across low- and high-fidelity evaluations. The overall goal is to contribute to the understanding of efficient optimization strategies for computationally demanding aerospace problems. This work demonstrates the potential of MFBO frameworks to accelerate the design cycle in real-world aerospace engineering contexts by highlighting the trade-offs between accuracy and computational cost.

2 METHODOLOGY

This section details the Bayesian Optimization (BO) methodology employed for optimizing expensive-to-evaluate functions. It introduces Gaussian Process (GP) surrogate models, first describing the single-fidelity version (Section 2.1), then the multi-fidelity extension using data of varying cost and accuracy (Section 2.2), and finally discusses the acquisition functions.

2.1 Single-Fidelity Gaussian Process

The Gaussian process (GP) is a stochastic surrogate modeling technique used to approximate an unknown function $f(\mathbf{x})$, using compromising input-output pairs $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$. Here, y_i represents the potentially noisy observation of $f(\mathbf{x}_i)$. The latent function $f(\mathbf{x})$ is modeled by its mean function $\mu(\mathbf{x})$ and the covariance function (or kernel) $\kappa(\mathbf{x}, \mathbf{x}'; \theta)$, which defines the relationship between any two points in the input space based on the hyperparameters θ. The mean function is often assumed to be zero ($\mu(\mathbf{x}) = 0$), simplifying the formulation:

$$f(\mathbf{x}) \sim \mathcal{GP}(0, \kappa(\mathbf{x}, \mathbf{x}'; \theta)) \tag{1}$$

The covariance function κ plays a crucial role in determining the properties (e.g., smoothness) of the functions sampled from the GP. Common choices for the kernel include the squared exponential kernel, Matérn kernel, and periodic kernel. For instance, the squared exponential kernel is defined as:

$$\kappa(\mathbf{x}, \mathbf{x}'; \theta) = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2l^2}\right) \tag{2}$$

where σ_f^2 is the signal variance and l is the length scale, both part of the hyperparameters θ, controlling the function's amplitude and smoothness. The hyperparameters θ and the noise variance σ_ϵ^2 are typically obtained by minimizing the negative log marginal likelihood (NLML) using the available data $\mathcal{D}$:

$$\text{NLML}(\theta, \sigma_\epsilon^2) = \frac{1}{2}\mathbf{y}^T(\mathbf{K} + \sigma_\epsilon^2\mathbf{I})^{-1}\mathbf{y} + \frac{1}{2}\log|\mathbf{K} + \sigma_\epsilon^2\mathbf{I}| + \frac{N}{2}\log(2\pi) \tag{3}$$

where $\mathbf{y} = [y_1, \ldots, y_N]^T$ is the vector of observed outputs, $\mathbf{I}$ is the $N \times N$ identity matrix, and $\mathbf{K}$ is the $N \times N$ covariance matrix with entries $(\mathbf{K})_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j; \theta)$.

Once the hyperparameters are optimized and the surrogate model is constructed, the predictive distribution for the function value $f(\mathbf{x}_*)$ at a new test point $\mathbf{x}_*$ can be derived.

This distribution is Gaussian, characterized by a posterior mean $\mu(\mathbf{x}_*)$ and posterior variance $\sigma^2(\mathbf{x}_*)$:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*(\mathbf{x}_*)^T(\mathbf{K} + \sigma_\epsilon^2\mathbf{I})^{-1}\mathbf{y} \tag{4}$$

$$\sigma^2(\mathbf{x}_*) = \kappa(\mathbf{x}_*, \mathbf{x}_*; \theta) - \mathbf{k}_*(\mathbf{x}_*)^T(\mathbf{K} + \sigma_\epsilon^2\mathbf{I})^{-1}\mathbf{k}_*(\mathbf{x}_*) \tag{5}$$

where $\mathbf{k}_*(\mathbf{x}_*)$ is the $N \times 1$ vector of covariances between the test point $\mathbf{x}_*$ and the training points $\mathbf{x}_i$, i.e., $(\mathbf{k}_*(\mathbf{x}_*))_i = \kappa(\mathbf{x}_*, \mathbf{x}_i; \theta)$, and $\kappa(\mathbf{x}_*, \mathbf{x}_*; \theta)$ is the prior variance at the test point $\mathbf{x}_*$.

2.2 Multi-Fidelity Gaussian Process Regression

The multi-fidelity Gaussian process approach utilizes a linear auto-regressive information fusion framework introduced by Kennedy and O'Hagan [13]. This method assumes a linear relationship between the outputs of models at consecutive fidelity levels:

$$f^{(m)}(x) = \rho_{m+1} f^{(m+1)}(x) + \delta^{(m+1)}(x) \quad m = M-1, ..., 1 \tag{6}$$

In this equation, ρ_{m+1} represents the scaling factor that quantifies the correlation between the outputs $f^{(m)}(x)$ and $f^{(m+1)}(x)$. The lowest fidelity model, denoted as $f^{(M)} \sim GP(0, \kappa_M(x, x'; \theta_M))$, is a Gaussian process characterized by a mean function $\mu_M = 0$ and a covariance function $\kappa_M(x, x'; \theta_M)$. The term $\delta^{(m+1)} \sim GP(0, \kappa_{m+1}(x, x'; \theta_{m+1}))$ models the discrepancy between consecutive fidelity levels, with θ_{m+1} representing the associated hyperparameters.

The recursive structure defined in Eq. (6) implies a specific conditional independence property, formalized by Kennedy and O'Hagan as:

$$\operatorname{cov}\left[f^{(m)}(x), f^{(m+1)}(x') \,|\, f^{(m+1)}(x)\right] = 0 \quad \forall x \neq x' \tag{7}$$

This means that, conditioned on the value of the higher-fidelity function $(f^{(m)}(x))$ at a point x, the lower-fidelity function $(f^{(m+1)}(x))$ is conditionally independent of the higher-fidelity function $(f^{(m)}(x))$ at any different point. In other words, the lower-fidelity function offers no additional information about the higher-fidelity function beyond what is already known at that point.

This recursive definition constructs a joint Gaussian process prior over all fidelity levels. Given a set of training data $\mathcal{D} = \{(x_i, m_i, y_i)\}_{i=1}^N$, where y_i is the noisy observation $f^{(m)}(x_i)$, we can obtain the posterior distribution of the function value $f^{(m)}(x_*)$ at a new test point x_* and the desired fidelity level m. The posterior is also Gaussian:

$$f^{(m)}(x_*)|\mathcal{D} \sim \mathcal{N}(\mu^{(m)}(x_*), (\sigma^{(m)})^2(x_*)) \tag{8}$$

The posterior mean $\mu^{(m)}(x_*)$ and variance $(\sigma^{(m)})^2(x_*)$ are given by:

$$\mu^{(m)}(x_*) = (\mathbf{k}^{(m)}(x_*))^T(\mathbf{K} + \sigma_\epsilon^2\mathbf{I})^{-1}\mathbf{y} \tag{9}$$

$$(\sigma^{(m)})^2(x_*) = \kappa((x_*, m), (x_*, m)) - (\mathbf{k}^{(m)}(x_*))^T(\mathbf{K} + \sigma_\epsilon^2\mathbf{I})^{-1}\mathbf{k}^{(m)}(x_*) \tag{10}$$

where:

- $\mathbf{y} = [y_1, \ldots, y_N]^T$ is the vector of observed training outputs.

- $\mathbf{K}$ is the $N \times N$ covariance matrix computed across all training data points, where $K_{ij} = \kappa((x_i, m_i), (x_j, m_j))$ using the composite kernel derived from the autoregressive structure.
- $\mathbf{k}^{(m)}(x_*) = [\kappa((x_*, m), (x_1, m_1)), \ldots, \kappa((x_*, m), (x_N, m_N))]^T$ is the $N \times 1$ vector of covariances between the test point (x_*, m) and all training points (x_i, m_i).
- $\kappa((x_*, m), (x_*, m))$ is the prior variance at the test point (x_*, m).
- $\mathbf{I}$ is the $N \times N$ identity matrix.
- σ_ϵ^2 is the variance of the observation noise.

The composite kernel $\kappa((x, m), (x', m'))$ needed to compute $\mathbf{K}$ and $\mathbf{k}^{(m)}(x_*)$ depends on the base kernels and the scaling factors . The hyperparameters are typically estimated by maximizing the logarithmic marginal likelihood of the training data $\mathcal{D}$, similar to the single-fidelity case.

3 ACQUISITION FUNCTIONS

This section presents theoretical information on the single- and multi-fidelity acquisition functions used in the study. The functions are categorized based on their underlying approach: Max-value Entropy Search (MES)-based methods and General-purpose Information-Based Bayesian OptimizatioN (GIBBON)-based methods. The highest fidelity function is consistently denoted as $f^{(1)}(\boldsymbol{x})$, and its maximum value is represented as $f_* = \max_{\boldsymbol{x} \in \mathcal{X}} f^{(1)}(\boldsymbol{x})$.

3.1 Max-value Entropy Search (MES)

Max-value Entropy Search (MES) is an acquisition function approach used in Bayesian optimization, based on information theory [14]. Unlike Predictive Entropy Search (PES) [15] or Entropy Search (ES) [16], which focus on reducing the uncertainty about the location of the global maximizer $\boldsymbol{x}_* = \arg\max_{\boldsymbol{x} \in \mathcal{X}} f^{(1)}(\boldsymbol{x})$, MES aims to maximize the information gained about the maximum value itself, f_*. This shift significantly simplifies the computation, since f_* is a one-dimensional quantity. The core idea is to measure the mutual information between a potential next observation and the maximum value f_*.

3.1.1 Single-Fidelity Max-value Entropy Search (MES)

In the single-fidelity setting, the MES acquisition function measures the mutual information between a potential next observation y at point $\boldsymbol{x}$ and the maximum value f_*, given the current data $D_t = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^t$:

$$\alpha_{\text{MES}}(\boldsymbol{x}) = I(f_*; y|\boldsymbol{x}, D_t) = H(y|\boldsymbol{x}, D_t) - \mathbb{E}_{f_*|D_t}[H(y|\boldsymbol{x}, f_*, D_t)] \tag{11}$$

where $H(\cdot)$ is the differential entropy. The first term, $H(y|\boldsymbol{x}, D_t)$, is the entropy of the predictive Gaussian distribution, determined by the predictive variance $\sigma^2(\boldsymbol{x})$. The second term involves an expectation over the posterior $p(f_*|D_t)$, typically approximated using Monte Carlo (MC) by drawing K samples $\{f_*^{(k)}\}_{k=1}^K$ (e.g., via Gumbel approximation or maximizing posterior samples [14]). The conditional entropy $H(y|\boldsymbol{x}, f_*, D_t)$ corresponds to that of a Gaussian truncated at the sampled $f_*^{(k)}$, which can be efficiently approximated [14].

3.1.2 Multi-fidelity Max-value Entropy Search (MF-MES)

Multi-fidelity MES (MF-MES) [9] extends this principle to leverage cheaper, lower-fidelity evaluations ($m > 1$) to find $f_* = \max_{\boldsymbol{x}\in\mathcal{X}} f^{(1)}(\boldsymbol{x})$ more efficiently. It selects the point-fidelity pair $(\boldsymbol{x}, m)$ maximizing the cost-normalized information gain:

$$a(\boldsymbol{x}, m) = \frac{I(f_*; y^{(m)}|\boldsymbol{x}, m, D_t)}{\lambda^{(m)}} \tag{12}$$

where $y^{(m)}$ is the observation at fidelity m, $\lambda^{(m)}$ is the cost, and D_t contains data from all fidelities $\{(\boldsymbol{x}_i, m_i, y_i)\}_{i=1}^t$. The mutual information $I(f_*; y^{(m)}|\dots)$ retains the structure of Eq. (11), but calculations employ a Multi-Fidelity Gaussian Process (MF-GP) model:

$$I(f_*; y^{(m)}|\boldsymbol{x}, m, D_t) = H(y^{(m)}|\boldsymbol{x}, m, D_t) - \mathbb{E}_{f_*|D_t}[H(y^{(m)}|\boldsymbol{x}, m, f_*, D_t)] \tag{13}$$

The first term $H(y^{(m)}|\boldsymbol{x}, m, D_t)$ depends on the MF-GP predictive variance $(\sigma^{(m)})^2(\boldsymbol{x})$. The expectation $\mathbb{E}_{f_*|D_t}[\cdot]$ is again estimated via MC sampling of f_*, typically using methods adapted for the MF-GP's prediction of the highest fidelity mean $\mu^{(1)}(\boldsymbol{x})$.

The key computational difference lies in evaluating the conditional entropy term $H(y^{(m)}|\boldsymbol{x}, m, f_*, D_t)$ for a sampled f_*:

- If $m = 1$ (highest fidelity), the calculation is analogous to the single-fidelity case, using efficient approximations for the truncated Gaussian entropy [9].
- If $m > 1$ (lower fidelity), the MF-GP structure induces a more complex dependency. Calculating this conditional entropy requires evaluating a one-dimensional numerical integral for each sampled f_* [9].

Consequently, evaluating lower-fidelity candidates ($m > 1$) in MF-MES incurs an additional computational cost per MC sample due to this numerical integration.

3.2 General-purpose Information-Based Bayesian OptimizatioN (GIBBON)

The General-purpose Information-Based Bayesian OptimizatioN (GIBBON) acquisition function [10] builds upon the core idea of MES (maximizing information about $f_* = \max_{\boldsymbol{x}} f^{(1)}(\boldsymbol{x})$) but introduces a novel analytical approximation for the information gain. This makes GIBBON computationally efficient and broadly applicable often without numerical integration [10]. It leverages the mutual information's symmetry and develops an analytical lower bound for the Information Gain (IG) between potential observation and a sampled max-value $z \sim p(f_*|D_t)$.

3.2.1 Single-Fidelity GIBBON

In the standard single-fidelity, non-batch setting, GIBBON approximates the MES objective (Eq. (11)) by replacing the challenging expected conditional entropy term. It achieves this by calculating an analytical lower bound, $IG_{\text{Approx}}(y, z)$, for the information gain $IG(y, z|D_t) = H(y|D_t) - H(y|f(\boldsymbol{x}) \leq z, D_t)$, where z is a sample from $p(f_*|D_t)$ and $f(\boldsymbol{x})$ is the latent function value.

The acquisition function is the Monte Carlo average of this approximation over K samples $\{f_*^{(k)}\}_{k=1}^K$:

$$\alpha_{\text{GIBBON}}(\boldsymbol{x}) \approx \frac{1}{K}\sum_{k=1}^{K} IG_{\text{Approx}}(y, f_*^{(k)}|\boldsymbol{x}, D_t) \tag{14}$$

For a single max-value sample $z = f_*^{(k)}$, the analytical approximation is [10]:

$$IG_{\text{Approx}}(y, z|\boldsymbol{x}, D_t) = \frac{1}{2}\log\left(\frac{\text{Var}(y|\boldsymbol{x}, D_t)}{\text{Var}(y|f(\boldsymbol{x}) \leq z, \boldsymbol{x}, D_t)}\right) \tag{15}$$

where $\text{Var}(y|\boldsymbol{x}, D_t) = \sigma_y^2(\boldsymbol{x})$ is the predictive variance of the observation y. The conditional variance $\text{Var}(y|f(\boldsymbol{x}) \leq z, \dots)$ also has a closed-form expression involving the standard normal PDF (ϕ) and CDF (Φ), the GP predictive mean $\mu(\boldsymbol{x})$ and variance $\sigma^2(\boldsymbol{x})$ of the latent function $f(\boldsymbol{x})$ via $\gamma(z) = (z-\mu(\boldsymbol{x}))/\sigma(\boldsymbol{x})$, and the correlation $\rho = \text{Corr}(f(\boldsymbol{x}), y)$ [10]. This fully analytical form is GIBBON's main advantage. Sampling f_* uses methods like those for MES.

3.2.2 Multi-fidelity GIBBON (MF-GIBBON)

The GIBBON framework's analytical nature allows its direct application to the multi-fidelity setting, handling single or batch ($B \geq 1$) point-fidelity queries $\mathcal{B} = \{(\boldsymbol{x}_j, m_j)\}_{j=1}^B$ while avoiding the numerical integration needed by MF-MES for $m > 1$.

MF-GIBBON approximates the information gain about $f_* = \max_{\boldsymbol{x}} f^{(1)}(\boldsymbol{x})$ from observing the vector of outcomes $\boldsymbol{A} = (y^{(m_1)}, \dots, y^{(m_B)})$ associated with $\mathcal{B}$. Similar to the single-fidelity case (Eq. (14)), the acquisition function is the MC average of the approximate IG over samples $z = f_*^{(k)}$:

$$\alpha_{\text{GIBBON}}(\mathcal{B}) \approx \frac{1}{K}\sum_{k=1}^{K} IG_{\text{Approx}}(\boldsymbol{A}, f_*^{(k)}|\mathcal{B}, D_t) \tag{16}$$

The approximate information gain $IG_{\text{Approx}}(\boldsymbol{A}, m|\mathcal{B}, D_t)$ remains analytical even for batches [10].

For the common sequential case ($B = 1$), selecting a single pair $(\boldsymbol{x}, m)$, the formula simplifies to a structure analogous to Eq. (15), but utilizing the Multi-Fidelity GP (MF-GP) model's predictions:

$$IG_{\text{Approx}}(y^{(m)}, z|(\boldsymbol{x}, m), D_t) = \frac{1}{2}\log\left(\frac{\text{Var}(y^{(m)}|\boldsymbol{x}, m, D_t)}{\text{Var}(y^{(m)}|f^{(1)}(\boldsymbol{x}) \leq z, \boldsymbol{x}, m, D_t)}\right) \tag{17}$$

Here, the variances are derived from the MF-GP. The conditional variance term remains analytical, depending on the MF-GP's prediction of the highest fidelity function $f^{(1)}$ at $\boldsymbol{x}$ (via $\mu^{(1)}(\boldsymbol{x}), \sigma^{(1)}(\boldsymbol{x})$ in $\gamma_1(z)$), the predictive variance of the observation $y^{(m)}$, and their correlation $\rho = \text{Corr}(y^{(m)}, f^{(1)}(\boldsymbol{x}))$.

Crucially, all components remain analytical, relying only on MF-GP predictions and standard functions (ϕ, Φ). This makes MF-GIBBON computationally advantageous compared to MF-MES, particularly when evaluating lower fidelities ($m > 1$) or batches ($B > 1$). The selection criterion is typically maximization of $\alpha_{\text{GIBBON}}(\mathcal{B})/c(\mathcal{B})$, where $c(\mathcal{B})$ is the evaluation cost.

4 BENCHMARK PROBLEMS

This section presents a detailed description of the benchmark problems utilized in the applications.

4.1 Analytical Benchmark Problems

The analytical equations considered in this study include the Rosenbrock function, which exhibits localized behavior, and the shifted-rotated Rastrigin function, which demonstrates multimodal characteristics. Scalability is assessed across various dimensions, with details provided in the following subsections.

4.1.1 Rosenbrock Function

The Rosenbrock function, commonly referred to as the Valley or Banana function, is a well-known non-convex optimization problem in a D-dimensional space, mathematically defined as:

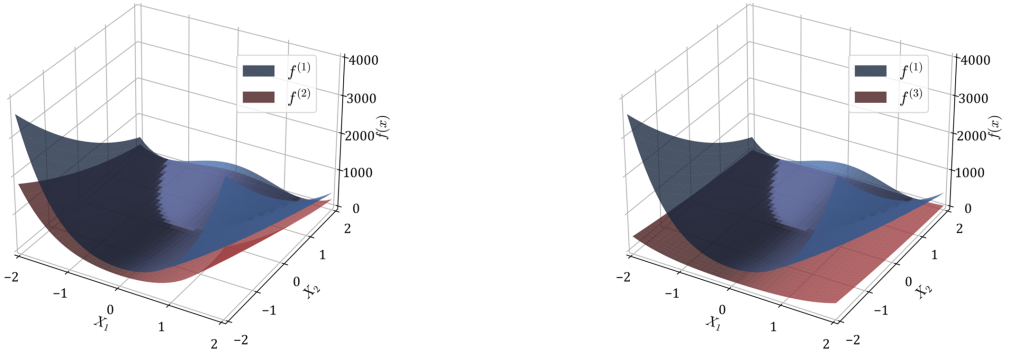
$$f_1(\mathbf{x}) = \sum_{i=1}^{D-1} 100 \left(x_{i+1} - x_i^2\right)^2 + (1 - x_i)^2 . \tag{18}$$

The domain of the variables is typically defined as $x_i \in [-2, 2]$ for $i = 1, 2, \ldots, D$. This function is unimodal, with a global minimum located at $\mathbf{x}^* = \{1, 1, \ldots, 1\}$, where the function value is $f_1(\mathbf{x}^*) = 0$. For multi-fidelity applications, the function is extended through the following formulations. Here, f_2 represents a medium-fidelity level, while f_3 corresponds to the lowest fidelity level:

$$f_2(\mathbf{x}) = \sum_{i=1}^{D-1} 50 \left(x_{i+1} - x_i^2\right)^2 + (-2 - x_i)^2 - \sum_{i=1}^{D} 0.5 x_i. \tag{19}$$

$$f_3(\mathbf{x}) = \frac{f_1(\mathbf{x}) - 4 - \sum_{i=1}^{D} 0.5 x_i}{10 + \sum_{i=1}^{D} 0.25 x_i}. \tag{20}$$

Figure 1 illustrates the Rosenbrock benchmark problem in two dimensions.



(a) High ($f^{(1)}$) and medium ($f^{(2)}$) fidelity

(b) High ($f^{(1)}$) and low ($f^{(3)}$) fidelity

Figure 1: Rosenbrock benchmark problem over the 2 dimensional domain

As seen in the figure, the high-fidelity version preserves the classic valley shape, while the lower fidelity versions introduce increasing levels of distortion.

4.1.2 Shifted-Rotated Rastrigin Function

The shifted-rotated Rastrigin function serves as a well-established benchmark problem designed to simulate real-world optimization challenges. This function is highly multimodal and is mathematically defined as follows:

$$f_1(\mathbf{z}) = \sum_{i=1}^{D} \left(z_i^2 + 1 - \cos(10\pi z_i) \right), \tag{21}$$

where $\mathbf{z}$ is given by:

$$\mathbf{z} = R(\theta)(\mathbf{x} - \mathbf{x}^\star), \quad \text{with} \quad R(\theta) = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}. \tag{22}$$

Here, R represents the rotation matrix with a rotation angle of $\theta = 0.2$. The variable domain is defined as $\mathbf{x} \in [-0.1, 0.2]^D$, and the global minimum is located at $\mathbf{x}^\star = \{0.1, \ldots, 0.1\}$, where the function value is $f(\mathbf{x}^\star) = 0$. For multi-fidelity applications, the fidelity levels are defined based on the following formulation [17]:

$$f_i(\mathbf{z}, \phi_i) = f_1(\mathbf{z}) + e_r(\mathbf{z}, \phi_i), \tag{23}$$

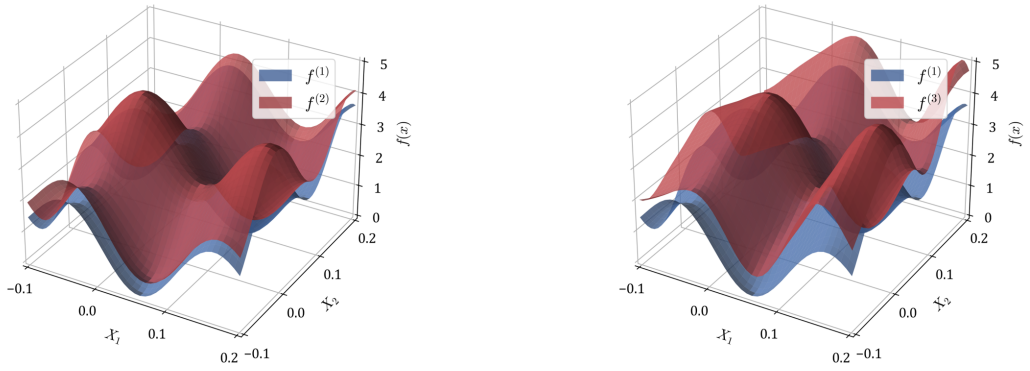
where $e_r(\mathbf{z}, \phi_i)$ represents the resolution error, defined as:

$$e_r(\mathbf{z}, \phi) = \sum_{i=1}^{D} a(\phi) \cos^2(w(\phi) z_i + b(\phi) + \pi). \tag{24}$$

The parameters $\Theta(\phi)$, $a(\phi)$, $w(\phi)$, and $b(\phi)$ are given by:

$$\Theta(\phi) = 1 - 0.0001\phi, \quad a(\phi) = \Theta(\phi), \quad w(\phi) = 10\pi\Theta(\phi), \quad b(\phi) = 0.5\pi\Theta(\phi). \tag{25}$$

Three fidelity levels are considered: high-fidelity ($\phi_1 = 10000$), medium-fidelity ($\phi_2 = 5000$), and low-fidelity ($\phi_3 = 2500$). Additionally, the benchmark problem can be extended to an n-dimensional space using the Aguilera-Perez Algorithm [18].



(a) High ($f^{(1)}$) and medium ($f^{(2)}$) fidelity (b) High ($f^{(1)}$) and low ($f^{(3)}$) fidelity

Figure 2: Rastrigin benchmark problem over the 2-dimensional domain

Figure 2 demonstrates the shifted-rotated Rastrigin function for different fidelity levels in the 2-dimensional case. As illustrated, the increasing discrepancy between fidelity

levels highlights the challenge in accurately modeling the underlying landscape, making this function a rigorous test case for multi-fidelity optimization frameworks.

4.2 Re-Entry Vehicle Design Problem

This problem, referred to as the Space Benchmark problem in the reference study by the NATO STO AVT-331 Research Task Group (RTG) [12], focuses on the aerodynamic shape optimization of a re-entry vehicle inspired by the Intermediate eXperimental Vehicle (IXV) technology demonstrator. The primary objective is to assess the influence of geometric modifications on aerodynamic efficiency while ensuring adherence to structural and thermal constraints. The baseline configuration of the space benchmark, including the initial shape with zero-displacement control points and the internal bay, is illustrated in Figure 3.

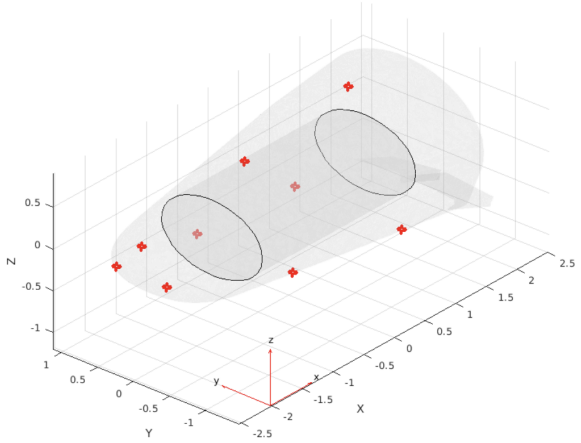


Figure 3: Baseline re-entry vehicle shape with zero-displacement control points and internal bay visualization

The re-entry conditions are prescribed at an altitude of 70 km, a velocity of $6,500$ m/s, and an angle of attack of $10°$. Shape deformations are introduced by morphing of the radial basis function (RBF), with 15 design variables that govern localized modifications in the x, y, and z directions. To preserve geometric feasibility, two volume coefficients are incorporated to prevent excessive deformations that could compromise the internal payload accommodation.

Furthermore, shape modifications are applied using a mesh morphing approach based on MIMMO [19]. The aerothermodynamic analysis is conducted utilizing the *Free Open Source Tool for Re-entry of Asteroids and Debris (FOSTRAD)* [20, 21], which employs *Modified Newtonian Theory* for continuum flow and the *Schaaf and Chambre model* for the free molecular regime. *FOSTRAD* is a multifidelity modeling framework tailored for hypersonic aerodynamic and aerothermodynamic analysis across various flow regimes pertinent to the re-entry of space vehicles, debris, and asteroids. Its predictive accuracy has been validated against Direct Simulation Monte Carlo (DSMC) and Computational Fluid Dynamics (CFD) simulations, ensuring reliability across diverse re-entry scenarios. The calculations in the transition regime are managed through a *bridging function*, facilitating a seamless transition between different flow conditions. Fidelity levels are defined by the resolution of the visibility detection module in the panel method solver, enabling

multi-fidelity analysis of aerodynamic and aerothermodynamic properties across varying flow regimes.

The present benchmark problem, including the aerodynamic and aerothermodynamic modules, as well as the morpher interface, is accessible via the GitLab repository of the AVT-331 working group. The computational framework is designed for execution on Linux operating systems. For an in-depth discussion of the aerodynamic and aerothermodynamics models, the mesh morphing interface, and the overall computational methodology, refer to [12].

5 APPLICATIONS

5.1 Analytical Benchmarks

In this section, the performance of multi-fidelity Bayesian optimization (MFBO) methods is evaluated through analytical benchmark problems. Before applying them to complex engineering application, the effectiveness of the methods is tested on Rosenbrock and Rastrigin functions.

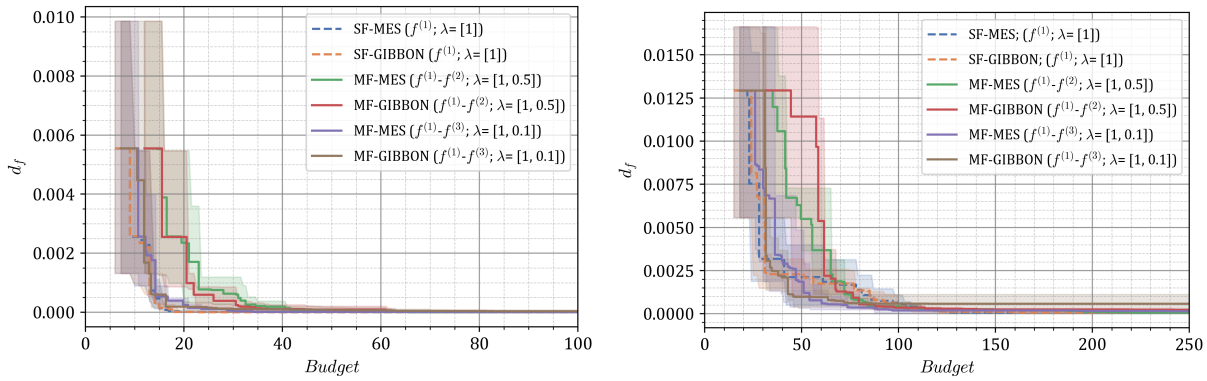
In the benchmark experiments, the total computational budget for each application is set to $50 \times N_{dim}$. As detailed in Section 2.2, the used multifidelity surrogate model relies on high-fidelity analysis results and assumes that modeling errors originate from low-fidelity sources. Therefore, in both single- and multi-fidelity applications, the number of high-fidelity evaluations used for training the initial surrogate models is taken as three times the dimension in all cases. In the two and three fidelity scenarios, the initial surrogate models are trained using low-fidelity data equal to six times problem dimension. Additionally, in the three-fidelity cases, the number of medium-fidelity evaluations is set to four times the dimension. The Latin Hypercube Sampling (LHS) method is used to generate the initial sampling points. To account for variability due to sampling, each experiment is repeated five times with different initial sample sets. The results are evaluated based on the difference error normalized in the input space, as defined in Eq. (26):

$$d_f = \frac{f(\hat{x}^*) - f_{min}}{f_{max} - f_{min}} \tag{26}$$

where $f(\hat{x}^*)$ is the best objective found, and f_{min}, f_{max} are the function's true minimum and maximum.

5.1.1 Examining the impact of fidelity level on results

This section analyzes the performance of single-fidelity (SF) and multi-fidelity (MF) Bayesian optimization strategies, focusing particularly on how the information obtained from the low-fidelity model affects the efficiency of MF approaches. Figures 4 and 5 show the convergence behavior for the Rosenbrock and Rastrigin benchmark functions, respectively. The performance metric plotted is the goal-driven metric (d_f) versus the computational budget consumed. Lower d_f values indicate better solutions. We compare SF-MES and SF-GIBBON, which use only the high-fidelity (HF) model ($f^{(1)}$), with MF-MES and MF-GIBBON, which use a medium-fidelity model ($f^{(2)}$) or a low-fidelity model ($f^{(3)}$) along with the HF model. The relative computational costs (λ) associated with each level of accuracy are indicated in the legend.



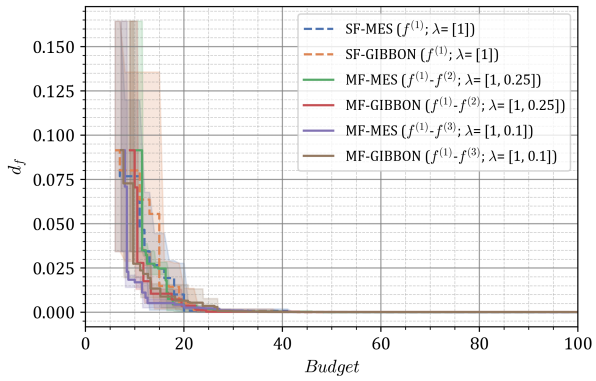
(a) 2D-Rosenbrock function results

(b) 5D-Rosenbrock function results

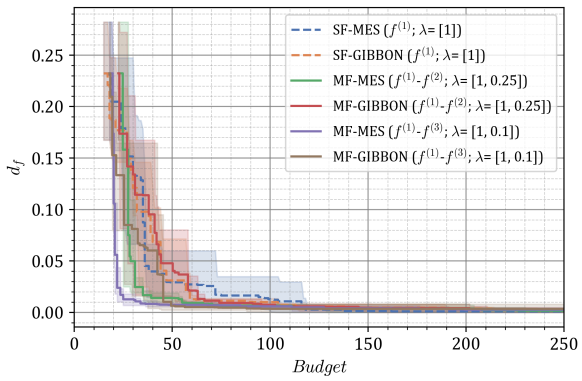
Figure 4: Rosenbrock function results

Considering the two-dimensional Rosenbrock results (Fig. 4a), it is observed that the multi-fidelity strategies using the cheapest low-fidelity model ($f^{(3)}$ with $\lambda = 0.1$) converge to the global minimum within a computational budget similar to the single-fidelity approaches. However, when using the relatively more expensive medium-fidelity model ($f^{(2)}$ with $\lambda = 0.5$), the convergence of the corresponding MF strategies (MF-MES($f^{(1)}$-$f^{(2)}$) and MF-GIBBON($f^{(1)}$-$f^{(2)}$)) seems to be slightly delayed compared to both the SF and MF($f^{(1)}$-$f^{(3)}$) methods. This suggests that the cost of the low-fidelity model may affect the convergence speed. For the two-dimensional Rastrigin function (Fig. 5a), the distinction between single- and multi-fidelity performance is less pronounced. All strategies appear to converge towards the optimal region with similar budget requirements. This indicates that for this particular low-dimensional multi-modal problem, the benefit of adding low-fidelity information to the tested budget is marginal; no significant difference is observed between the convergence plots.

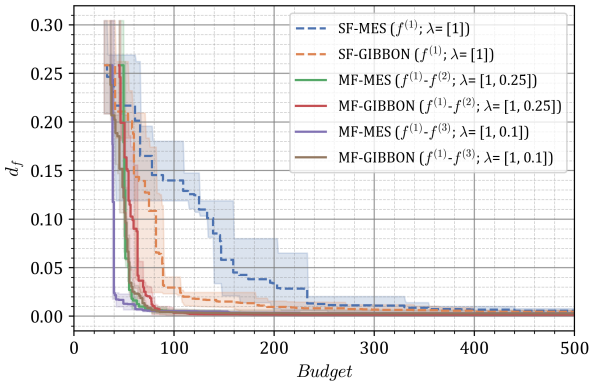
The advantages of the multi-fidelity approach become more apparent in the five-dimensional Rosenbrock case (Fig. 4b). Here, the strategies using the cheapest low-fidelity model ($f^{(1)}$-$f^{(3)}$) clearly outperform the single-fidelity methods and reach the global minimum region with significantly less computational budget. Interestingly, even the multi-fidelity approach using the more expensive medium-fidelity model ($f^{(1)}$-$f^{(2)}$) exhibits a slight advantage over the single-fidelity strategies, converging slightly earlier, but the difference is not significant. When examining the five-dimensional Rastrigin results (Fig. 5b), the multi-fidelity methods (especially MF-MES) show a clear benefit. Both variants of MF-MES appear to converge to significantly faster and potentially better solutions within the allocated budget compared to the single-fidelity methods and MF-GIBBON. This trend becomes even more pronounced in the ten-dimensional Rastrigin case (Fig. 5c). Here, multi-fidelity approaches exhibit a striking acceleration in convergence. Both MF-MES and MF-GIBBON approach the global minimum using approximately 100 units of computational budget, whereas single-fidelity strategies require around 300 units to reach comparable performance. This case highlights an approximate 60% computational cost savings of multi-fidelity methods over single-fidelity ones. Consistent with the 5D Rastrigin results, MF-MES appears to be slightly more sample-efficient than MF-GIBBON in this higher-dimensional, highly multimodal scenario. In both 5D and 10D Rastrigin cases, MF-GIBBON appears to be less effective than MF-MES, which



(a) 2D-Rastrigin function results



(b) 5D-Rastrigin function results



(c) 10D-Rastrigin function results

Figure 5: Rastrigin function results

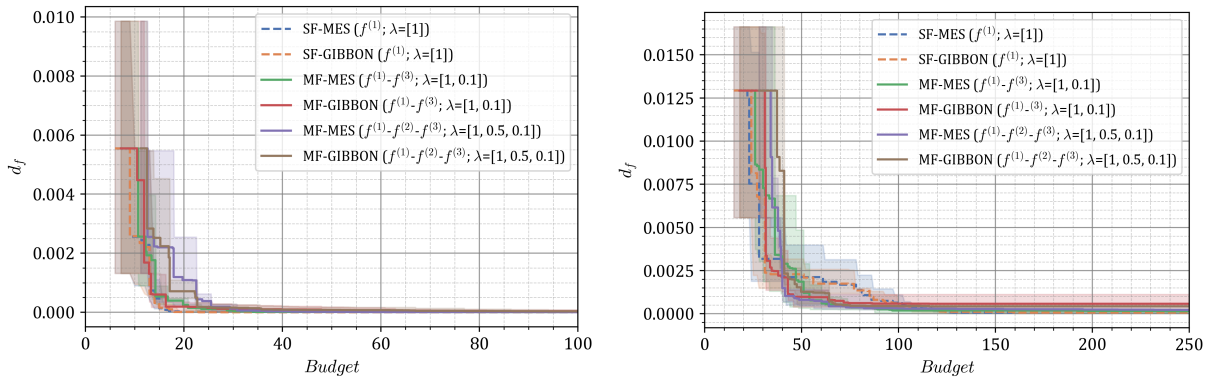
may indicate a less efficient use of multi-fidelity information or a greater emphasis on exploration for this particular function and acquisition function combination.

Results indicate that Multi-Fidelity Bayesian Optimization (MFBO) effectiveness strongly depends on low-fidelity (LF) model cost and problem dimensionality. Advantages are less pronounced in lower dimensions, especially with relatively expensive LF models. However, benefits increase significantly with dimensionality; in higher dimensions, MFBO strategies using cheap LF models (e.g., $f^{(3)}$) can substantially accelerate convergence compared to single-fidelity methods. The relative LF cost is critical, determining the cost-benefit trade-off for guiding the search.

5.1.2 Examining the importance of fidelity numbers

This section delves the impact of the number of fidelity levels used in the multi-fidelity Bayesian optimization framework. Specifically, we compare the performance of using three fidelity levels (high $f^{(1)}$, medium $f^{(2)}$, and low $f^{(3)}$) with only two levels (high $f^{(1)}$ and low $f^{(3)}$) and the standard single-fidelity approach (high $f^{(1)}$ only). Figures 6 and 7 show these comparisons for the Rosenbrock and Rastrigin functions, respectively. The legends indicate the fidelity levels used and their relative computational costs (λ).

Analyzing the results for the 2D Rosenbrock function (Figure 6a), an interesting trend emerges. The three-fidelity strategies exhibit slightly delayed convergence compared to both the single-fidelity approaches and the two-fidelity strategies using the cheapest low-



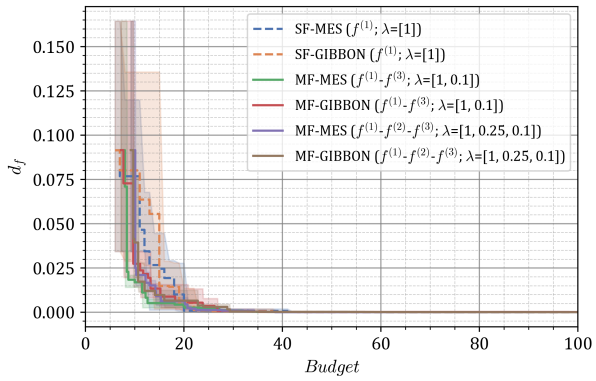
(a) 2D-Rosenbrock function results

(b) 5D-Rosenbrock function results

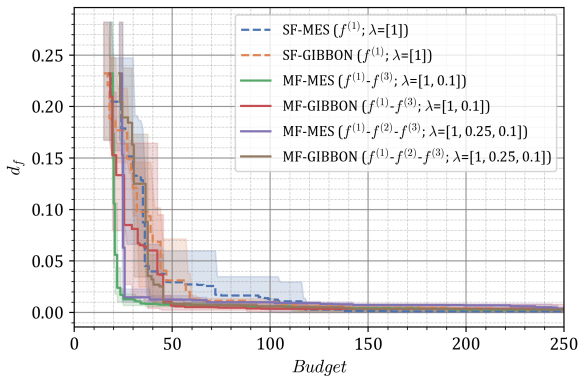
Figure 6: Rosenbrock function results

fidelity model. This suggests that, for this relatively simple, low-dimensional problem, incorporating an intermediate fidelity level ($f^{(2)}$) introduces additional complexity and cost that may not provide enough information gain to justify its inclusion—resulting in marginally slower progress toward the optimum within the given budget. In contrast, for the 2D Rastrigin function (Figure 7b), the performance differences between the single-fidelity, two-fidelity, and three-fidelity approaches are minimal. All strategies exhibit similar convergence behaviors in general and rapidly converge near the optimal solution within a comparable computational budget. This suggests that for this particular low-dimensional, multimodal problem, including low-fidelity information (regardless of whether it is two or three levels) does not offer a significant advantage over the single-fidelity approach under the conditions tested.

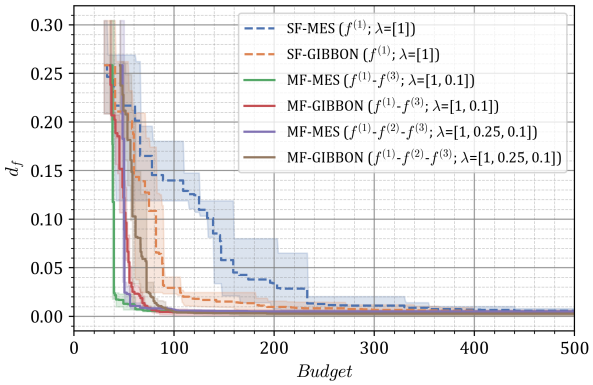
The benefits of multi-fidelity approaches and the nuances of using different numbers of fidelity levels become more pronounced in higher dimensions. For the 5D Rosenbrock function (Figure 6b, both two- and three-fidelity MFBO strategies clearly outperform single-fidelity methods, converging toward the optimal region significantly faster and with a lower computational budget. When comparing the two- and three-fidelity MF approaches, their convergence budgets appear notably similar. In this case, adding a third fidelity level does not yield a substantial speed-up. However, a consistent observation is that MF-MES variants (both two- and three-fidelity) tend to achieve slightly more accurate final solutions (lower average df values) than their MF-GIBBON counterparts for this problem. Finally, in the 5D Rastrigin results (Figure 7b), we again observe that multi-fidelity strategies converge significantly earlier than single-fidelity methods. The convergence behaviors of the two-fidelity and three-fidelity setups are very similar, indicating that adding a third fidelity level does not fundamentally alter the convergence path or required budget. Interestingly, in this particular case, the two-fidelity strategies yield slightly more accurate results after convergence compared to the three-fidelity ones. This might suggest that the intermediate fidelity level ($f^{(2)}$) does not provide information that aligns well with the path toward the global optimum, or that the balance between exploration and exploitation is slightly more effective in the two-fidelity configuration. A similar pattern emerges in the ten-dimensional Rastrigin case (Fig. 7c). Both the two-fidelity and three-fidelity multi-fidelity optimization strategies substantially outperform their single-fidelity counterparts, which exhibit minimal improvement within the allocated



(a) 2D-Rastrigin function results



(b) 5D-Rastrigin function results



(c) 10D-Rastrigin function results

Figure 7: Rastrigin function results

budget. For both MF-MES and MF-GIBBON, the inclusion of the intermediate fidelity level ($f^{(2)}$) does not significantly affect convergence speed or final solution quality.

In summary, while multi-fidelity methods are often advantageous in higher dimensions, adding more fidelity levels does not guarantee further benefit and is highly problem-dependent. Sometimes, using only a well-correlated high-fidelity and a cheap low-fidelity source can be as effective, or even slightly more accurate, than including an intermediate level. The relative costs and information quality of each fidelity level are critical in determining the optimal configuration.

5.2 Shape Optimization Problem

This optimization problem serves as a multi-fidelity design benchmark, originally proposed by the NATO STO AVT-331 team, and is formally defined as Equation (27):

$$\begin{aligned} \min_{\mathbf{x}\in\mathbb{R}^{15}} \quad & -\frac{C_L(\mathbf{x})}{C_D(\mathbf{x})} \\ \text{subject to} \quad & L(\mathbf{x}) \geq L_r \\ & \dot{Q}_{\max}(\mathbf{x}) \leq \dot{Q}_{r,\max} \\ & C_{V,i}(\mathbf{x}) = 0, \quad i = 1, 2. \end{aligned} \tag{27}$$

In this formulation, $\mathbf{x}$ represents the morphing displacements along the x, y, and z

axes, constrained to symmetric deformations. The term L denotes the aerodynamic lift force, with L_r as the minimum required lift, set to 1500 N. The maximum heat flux, $\dot{Q}_{\max}$, is constrained by the reference value $\dot{Q}_{r,\max}$, which is set at 550 kW/m^2. The coefficients $C_{V,1}$ and $C_{V,2}$ account for volume constraints related to an internal bay area, approximated as a cylindrical structure in this benchmark. These volume factors are computed as follows:

$$C_{V,1} = w_1 \frac{n_{\text{exp}}}{n_{\text{tot}}} \qquad C_{V,2} = w_2 \frac{v_{\text{exp}}}{v_{\text{tot}}} \tag{28}$$

where n_{exp} represents the number of voxel center points within the bay that extend beyond the vehicle's surface mesh, and n_{tot} is the total voxel count inside the bay. Similarly, v_{exp} denotes the number of bay or payload surface vertices outside the vehicle mesh, while v_{tot} is the total count of these vertices. The parameters w_1 and w_2 serve as weighting coefficients.

To facilitate numerical optimization, a penalty-based reformulation is introduced, yielding the modified objective function in Equation (29):

$$f_{\text{pen}} = \frac{-C_L}{C_D} + k_1 p^+(L, L_r) + k_2 p^-(\dot{Q}_{\max}, \dot{Q}_{r,\max}) + k_3 p^{\neq}(C_{V,1}) + k_4 p^{\neq}(C_{V,2}) \tag{29}$$

where the penalty coefficients are set as $k_1 = k_2 = 2$ and $k_3 = k_4 = 0.1$ for all fidelity levels. The penalty functions are defined as:

$$p^+(x,y) = \begin{cases} 0, & x \geq y \\ x - y, & x < y \end{cases} \quad p^-(x,y) = \begin{cases} 0, & x \leq y \\ x - y, & x > y \end{cases} \quad p^{\neq}(x) = \begin{cases} 0, & x = 0 \\ x, & x \neq 0 \end{cases} \tag{30}$$

Bayesian optimization, both single- and multi-fidelity, was conducted using a modified objective function integrated with the constraints provided in Eq. (29). The low-fidelity and high-fidelity models exhibit a linear correlation; however, the discontinuities introduced by penalty terms on the response surface pose challenges for traditional modeling approaches, such as Gaussian processes, in achieving accurate modeling.

To perform a preliminary statistical analysis of the optimization results, each optimization process was executed three times. In each repetition, the initial sampling data was generated from scratch. The single-fidelity Bayesian optimization commenced with 45 $(3 * n_{dim})$ high-fidelity analyses, whereas the multi-fidelity Bayesian optimization started with 45 $(3 * n_{dim})$ high-fidelity and 90 $(6 * n_{dim})$ low-fidelity analyses. Both multi-fidelity and single-fidelity Bayesian optimization approaches utilized the same high-fidelity model results. Due to the inclusion of low-fidelity model results in the initial surrogate model, the multi-fidelity Bayesian optimization was initiated with a slightly higher computational budget compared to the single-fidelity Bayesian optimization.

The comparative performance of single- and multi-fidelity Bayesian optimization strategies is illustrated in Figure 8, which plots the convergence history of the best-found objective function value $(-C_L/C_D)$ that satisfy the constraint against the computational budget consumed. The results clearly demonstrate the advantages of the multi-fidelity approach for this challenging aerospace shape optimization problem. Both multi-fidelity

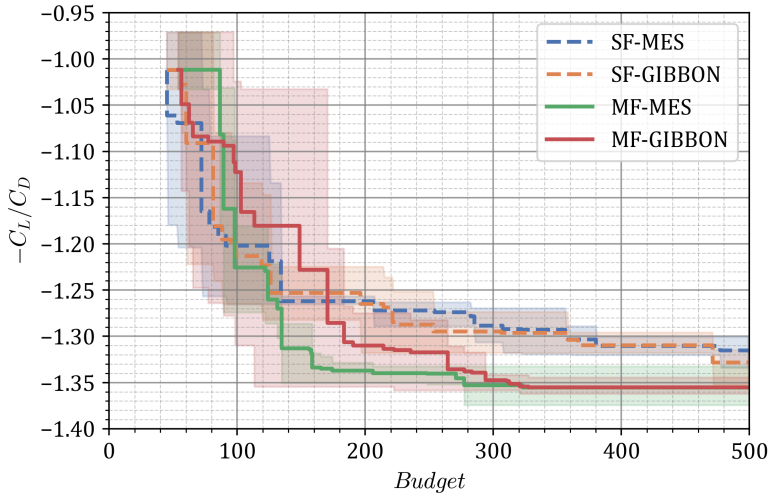


Figure 8: Comparison of Convergence History

methods, MF-MES and MF-GIBBON, exhibit significantly faster convergence compared to their single-fidelity counterparts (SF-MES and SF-GIBBON). Within a budget of approximately 200 units, the MFBO strategies have already identified solutions with objective values around -1.33, a level that the SFBO methods struggle to reach even with the full budget of 500 units. Furthermore, the final solution quality achieved by MFBO is substantially better. Both MF-MES and MF-GIBBON converge to an average objective value of approximately -1.36, whereas SF-MES and SF-GIBBON plateau at significantly worse values, around -1.30 and -1.33 respectively. A key reason for this performance difference lies in the discontinuities introduced into the objective function landscape by the penalty terms defined in Eq. (29). Single-fidelity Bayesian optimization (SFBO), relying solely on a limited number of expensive high-fidelity (HF) data points, struggles to moderately model these sharp changes and the boundaries defined by the constraints. Surrogate models like Gaussian Processes tend to assume smoothness, and these discontinuities degrade model accuracy. Consequently, SFBO may be guided towards exploring regions that appear promising according to the surrogate model but actually incur heavy penalties (i.e., violate constraints). This leads to inefficient spending of expensive HF evaluations in unsuitable areas and hinders convergence to good, feasible solutions. In contrast, multi-fidelity Bayesian optimization (MFBO) overcomes this challenge by utilizing numerous inexpensive low-fidelity (LF) evaluations. The large volume of LF evaluations allows MFBO to build a more detailed map of the objective function landscape, including a better estimate of where the discontinuities (penalty boundaries) lie. This improved understanding enables MFBO to direct its limited HF evaluations more effectively towards regions that are genuinely promising and feasible. This leads to both faster convergence and superior final designs within the given computational constraints.

The efficiency gains observed in Figure 8 can be attributed to the strategic utilization of computational resources within the MFBO framework, as detailed in Figure 9. This figure shows the distribution of computational cost between low-fidelity (LF) and high-fidelity (HF) model evaluations during the optimization process. Both MF-MES and MF-GIBBON allocated a substantial portion of their computational budget – roughly $50 - 60\%$ – to inexpensive LF evaluations. These LF simulations effectively guided the exploration of the design space, allowing the algorithms to identify promising regions

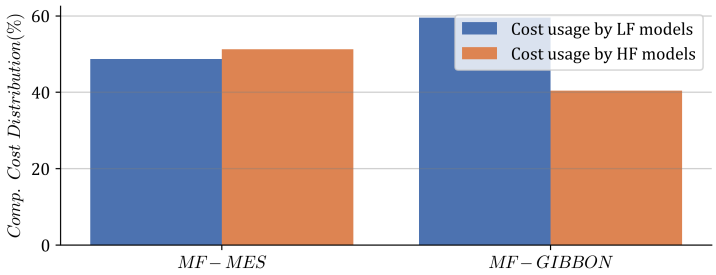


Figure 9: Computational cost distribution used in LF and HF model evaluation (Average of 3 repetitions)

before committing resources to costly HF evaluations (which consumed the remaining $45-60\%$ of the budget). This intelligent allocation, leveraging cheaper models to inform the selection of expensive ones, is the core reason for MFBO's superior performance in terms of both speed and final solution quality compared to SFBO, which relies solely on expensive HF evaluations. Furthermore, an interesting detail emerges when comparing the two MFBO methods internally. While both MFBO strategies exhibit very similar final performance in Figure 8, Figure 9 suggests that MF-GIBBON potentially achieves this performance slightly more efficiently. MF-GIBBON allocated a marginally smaller proportion of its budget ($\sim 40\%$) to high-fidelity model evaluations, whereas MF-MES used a slightly higher proportion ($\sim 52\%$). This might suggest that, for this specific problem, MF-GIBBON was marginally more efficient in its use of HF resources, although the overall performance difference remains small.

6 CONCLUSION

This study investigated Multi-Fidelity Bayesian Optimization (MFBO) methods, particularly those employing information-theoretic acquisition functions such as MF-MES and MF-GIBBON, aiming to accelerate the optimization of aerospace engineering design problems. The performance of these methods was evaluated against their single-fidelity (SF) counterparts and using different numbers of fidelity levels on both analytical benchmark problems (Rosenbrock and Rastrigin) and a re-entry vehicle shape optimization problem. The analytical tests demonstrated that the advantages of MFBO become more pronounced, especially with increasing problem dimensionality. The cost of the low-fidelity (LF) models and their correlation with the high-fidelity (HF) model play a critical role in MFBO's effectiveness; cheaper LF models generally offer a better cost-benefit trade-off. It was observed that the additional benefit of using more than two fidelity levels is problem-dependent and not always guaranteed. In the challenging re-entry vehicle shape optimization problem, the MFBO methods (both MF-MES and MF-GIBBON) exhibited significantly superior performance compared to the single-fidelity (SFBO) methods. MFBO intelligently allocated a substantial portion of the computational budget ($50-60\%$) to inexpensive LF evaluations, enabling a more effective mapping of the complex objective function landscape, particularly where discontinuities arose from penalty functions. This, in turn, allowed for more efficient guidance of the HF evaluations, resulting in much faster convergence and markedly better final designs. Among the two MFBO methods, it was observed that MF-GIBBON might have utilized HF resources marginally more efficiently for this specific problem. Overall, the findings confirm that

MFBO is a robust and effective framework for computationally intensive aerospace optimization tasks. By leveraging inexpensive LF information, MFBO holds the potential to significantly accelerate the design process and discover superior solutions within limited computational resources. This positions MFBO as a valuable tool in modern engineering design cycles.

ACKNOWLEDGEMENTS

The authors would like to express their gratitude to General Electric Vernova, UK, for the research grant titled "Implementation of Bayesian Optimization and Machine Learning Algorithms to Automated Process and Optimization Workbench".

REFERENCES

[1] M. Nardelli. Multifidelity active learning for aerospace design and optimization. Master's thesis, Politecnico di Torino, Turin, Italy, 2022.

[2] B. Do and R. Zhang. Multifidelity bayesian optimization: A review. *AIAA Journal*, page 1–37, Feb. 2025. ISSN 1533-385X. doi:10.2514/1.j063812.

[3] F. Di Fiore, M. Nardelli, and L. Mainini. Active learning and bayesian optimization: A unified perspective to learn with a goal. *Archives of Computational Methods in Engineering*, 31:2985–3013, 2024. doi:10.1007/s11831-024-10064-z.

[4] Y. Zhang, T. N. Hoang, B. K. H. Low, and M. Kankanhalli. Information-based multi-fidelity bayesian optimization. In *NIPS Workshop on Bayesian Optimization*, 2017.

[5] R. Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.

[6] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. doi:10.1109/JPROC.2015.2494218.

[7] C. Park, R. T. Haftka, and N. H. Kim. Remarks on multi-fidelity surrogates. *Structural and Multidisciplinary Optimization*, 55(3):1029–1050, 2017. doi:10.1007/s00158-016-1533-0.

[8] L. Brevault, M. Balesdent, and A. Hebbal. Overview of gaussian process based multi-fidelity techniques with variable relationship between fidelities. *Aerospace Science and Technology*, 107:106339, 2020. doi:10.1016/j.ast.2020.106339.

[9] S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In H. D. III and A. Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9334–9345. PMLR, 13–18 Jul 2020. arXiv:1901.08275.

[10] H. B. Moss, D. S. Leslie, J. Gonzalez, and P. Rayson. GIBBON: General-purpose information-based Bayesian OptimisatioN. *arXiv preprint arXiv:2102.03324*, 2021.

[11] L. Mainini, A. Serani, M. P. Rumpfkeil, E. Minisci, D. Quagliarella, H. Pehlivan, S. Yildiz, S. Ficini, R. Pellegrini, F. Di Fiore, D. Bryson, M. Nikbay, M. Diez, and P. Beran. Analytical benchmark problems for multifidelity optimization methods, 2022. URL `https://arxiv.org/abs/2204.07867`.

[12] D. Quagliarella, D. Clark, D. Bryson, P. Beran, A. Thelen, L. Mainini, S. Yildiz, M. Nikbay, E. Minisci, P. Leyland, A. Serani, and M. Diez. Reproducible industrial multifidelity optimization benchmark problems for air, space, and sea vehicles. Research workshop AVT-354 on Multifidelity methods for military vehicle design, 09 2022.

[13] M. C. Kennedy and A. O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.

[14] Z. Wang and S. Jegelka. Max-value entropy search for efficient Bayesian optimization. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, volume 70 of *Proceedings of Machine Learning Research*, pages 3627–3635. PMLR, 06–11 Aug 2017.

[15] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems 27 (NIPS)*, pages 918–926, 2014.

[16] P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research (JMLR)*, 13(Jun):1809–1837, 2012.

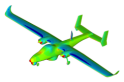
[17] H. Wang, Y. Jin, and J. Doherty. A generic test suite for evolutionary multifidelity optimization. *IEEE Transactions on Evolutionary Computation*, 22(6):836–850, 2018. doi:10.1109/TEVC.2017.2758360.

[18] A. Aguilera and R. Pérez-Aguila. General n-dimensional rotations. In *International Conference in Central Europe on Computer Graphics and Visualization*, 2004.

[19] O. engineering srl. mimmo: Surface manipulation and mesh morphing library, 2015. https://github.com/optimad/mimmo.

[20] P. Mehta, E. Minisci, M. Vasile, A. C. Walker, and M. Brown. An open source hypersonic aerodynamic and aerothermodynamic modelling tool. In *8th European Symposium on Aerothermodynamics for Space Vehicles*, 2015.

[21] A. Falchi, V. Renato, E. Minisci, and M. Vasile. Fostrad: An advanced open source tool for re-entry analysis. In *15th Reinventing Space Conference*, 2017.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

SHAPE OPTIMIZATION OF AUXETIC UNIT CELLS UNDER DYNAMIC LOADING IN MACROSCOPIC COMPONENTS

Nicolas Grünfelder[1*], Lars Kälber[1] , Navina Waschinsky[1] and Tim Ricken[1]

1: Institute of Structural Mechanics and Dynamics in Aerospace Engineering (ISD)
Faculty for Aerospace Engineering and Geodesy
University of Stuttgart
70569
nicolas.gruenfelder@isd.uni-stuttgart.de, https://www.isd.uni-stuttgart.de/

Abstract. *Auxetic metamaterials, characterized by their negative Poisson's ratio, exhibit unique mechanical properties such as enhanced energy dissipation, improved vibration damping, and superior energy absorption. These properties make them particularly attractive for aerospace applications, where lightweight and high-performance structures are crucial. However, integrating auxetic structures into macroscopic components remains challenging due to the complex interplay between mesostructural design and overall structural behavior. This study introduces an optimization methodology that enables independent shape adjustments of auxetic mesostructures within different regions of a macroscopic component. Using a Bayesian optimization algorithm, the unit cell geometry is optimized to maximize energy dissipation, enhance dynamic stiffness, and minimize mass, ensuring a well-balanced trade-off between these competing objectives. A macroscopic cantilever beam composed of reentrant auxetic unit cells serves as a real-world-inspired case study and is analyzed under dynamic loading conditions, demonstrating the effectiveness of the optimized mesostructure in improving structural performance. Beyond the optimized structure itself, this study provides an in-depth analysis of the optimization process, offering valuable insights into the application of auxetic metamaterials in engineering practice. Additionally, a mesh convergence study is conducted to validate numerical accuracy. The results underscore the potential of auxetic metamaterials for aerospace applications, highlighting their performance-driven optimization and real-world applicability.*

Keywords: Auxetics, Shape Optimization, Bayesian Optimization, Lightweight, Energy Dissipation, Aerospace

1 INTRODUCTION

Auxetic structures represent a distinctive subclass of metamaterials named after and mainly characterized by their negative Poisson's ratio, which leads to unique mechanical properties [1]. These include properties such as enhanced energy dissipation [2], high energy absorption capabilities [3] and vibration damping [4] as well as vibration control [5, 6]. These characteristics have garnered significant interest in aerospace engineering, where lightweight, high-performance structures are essential for improving aircraft efficiency and safety across various applications [7]. This is particularly relevant as emerging design and manufacturing techniques necessitate exploring synergies across multiple scales and hierarchical levels [8]. Additionally, auxetic structures present significant opportunities for lightweight construction, enabling applications that were previously infeasible [9]. Addressing the challenges raised by these upcoming tasks requires a comprehensive understanding of how auxetic mesostructures influence macroscopic structural behavior. To date, extensive research has been conducted to identify suitable mesostructure topologies [10–12].

However, a critical aspect, that remains insufficiently looked at, is the bias, that can unintentionally and unknowingly be introduced in this kind of studies. A key challenge in comparing auxetic structures and implementing the right structure in the right macroscopic component lies in ensuring unbiased comparisons between different configurations in order to draw accurate conclusions. While this is not a unique issue for auxetic structures and is commonly encountered when comparing conceptually different shapes or topologies, it becomes particularly evident in auxetic structures due to their strong dependence on mesostructural geometric behavior.

To overcome this limitation, this study introduces an optimization methodology that enables independent shape adjustments of mesostructures within different regions of a macroscopic component. With this method, the most advantageous configuration can be achieved for specific performance objectives under realistic loading conditions. Additionally this method allows for a controlled optimization that in return leaves the possibility of making the results comprehensible for an engineering classification of the results.

Traditionally, optimization of structures consists either of a topology optimization or a hierarchical approach that combines topology and shape optimization [13]. These methods offer the advantage of requiring less initial setup effort while providing a broad design space for structural configurations [14]. However, they also present challenges. First, this study deals with a given mesostructure, so by default with a predefined topology, which cannot be effectively tackled using topology optimization. Second, in complex lattice-structured components, comprehensibility and explainability are often compromised, as the connection between the initial and optimized configurations is not straightforward.

In this study, a macroscopic cantilever beam composed of reentrant auxetic unit cells is analyzed under dynamic loading conditions. The geometry of the unit cells is discretized on the surface, and the resulting nodes are controlled by a Bayesian optimization algorithm, allowing for a high degree of design flexibility. The optimization process focuses on three key objectives: minimizing structural mass, maximizing energy dissipation, and enhancing dynamic stiffness over a specific frequency range. These objectives are selected to imitate practical benefits for real-world aerospace components, rather than purely theoretical improvements. By implementing auxetic structures within an actual component and subjecting them to a realistic load case, this study provides valuable insights into

the practical applications of auxetic metamaterials. The findings contribute to a deeper understanding of how these structures can enhance the performance of aerospace components, paving the way for more efficient, resilient, and adaptive aircraft designs.

2 GEOMETRY

In this study, the investigated and used mesostructure is a re-entrant shaped unit cell that can be seen in Fig.1. The unit cell is divided into quarters to define the relevant parameters. Here, a and b represent the width and height of each quarter. t_1, t_2, and t_3 describe the thicknesses, whereas t_3 describes the thickness into the plain. α indicates the angle of the diagonal strut.

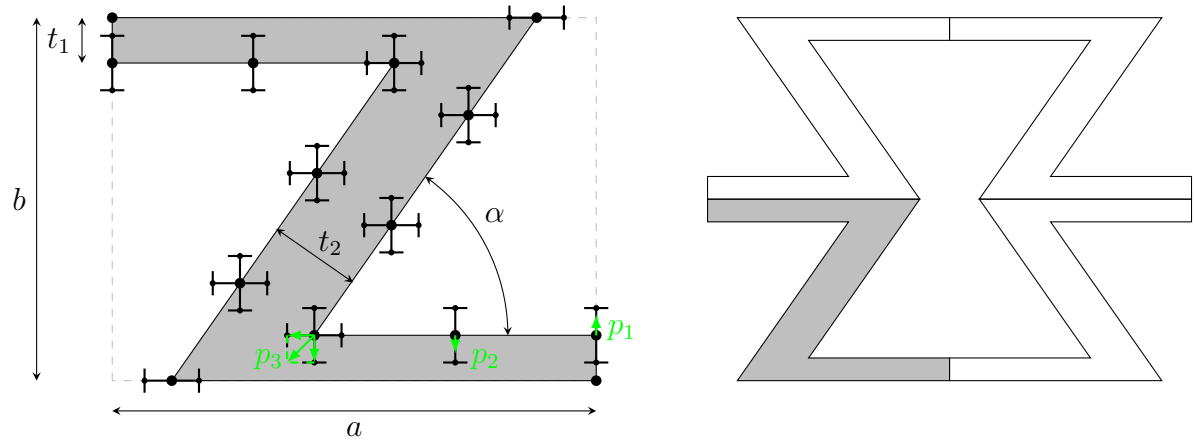


Figure 1: Schematic illustration of a quarter cell (left), as part of the two-dimensional auxetic unit cell (right). On the quarter unit cell the bars indicate the possible shape changes the optimization can leverage.

The points of a quarter-unit cell can move in the x- and y-directions within the intervals $p_i^{x,nom} \pm \Delta t_i^x$ and $p_i^{y,nom} \pm \Delta t_i^y$, respectively, as shown in Fig.1. However, not every point can be moved in every spatial direction, as they serve as connections between adjacent unit cells. Altering these positions indiscriminately could result in overlaps or gaps. Additionally, the maximum displacement of each point is constrained to prevent intersections and to maintain the unit cell's topology. These constraints are implemented as bounds in the optimization algorithm, as described in section 5. The shape modifications applied to a quarter-unit cell are mirrored across the remaining quarters to ensure symmetry. Furthermore, the macroscopic geometry is divided into nine sections, where modifications are applied uniformly to all unit cells within each section. An example of this division is illustrated in figure 2.

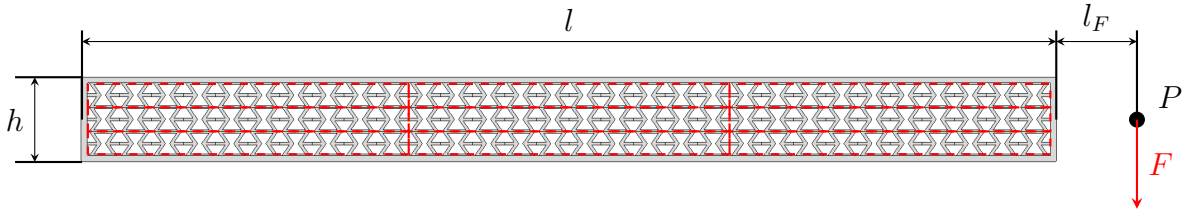


Figure 2: Macroscopic geometry with the optimization area highlighted.

This approach significantly reduces the computational cost of the study while still allowing for an optimization tailored to the deformation behavior. This is important because it has been shown already that the deformation behaviour of the unit cells is highly position dependent, especially for dynamic problems [15]. The macroscopic geometry consists of a lattice structure with three unit cells in the y-direction and 30 unit cells in the x-direction. This lattice is enclosed by a casing with a wall thickness of $2 \cdot t_2$. Since notched geometries are highly susceptible to stress singularities, a fillet with radius r is introduced at all notches within the mesostructure. The specific parameters used for the initial model are listed in Tab.1.

Table 1: Geometry parameters of the unit cell and macroscopic structure.

a	b	t_1	t_2	t_3	α
in mm	in mm	in mm	in mm	in mm	in degree
8.00	6.00	0.75	1.50	10.00	55.00
l	h	l_F	r	columns	rows
in mm	in mm	in mm	in mm		
486.00	42.00	97.20	0.10	30	3

3 SIMULATION

All simulations shown in this study rely on the in-plane stress theory, which allows for a significant reduction in computational costs for the otherwise expensive simulations and the iterations for the optimization, that need to be carried out. As a solver ABAQUS/Explicit 2024 is used. The pre- and postprocessing as well as the optimization are scripted in python, and Abaqus operations are called when needed.
A linear elastic material is chosen as generic steel with a Young's modulus of $E = 210\,000.00\,\frac{\mathrm{N}}{\mathrm{mm}^2}$, a Poisson's ratio of $\nu = 0.3$ and a density of $\rho = 7.85 \cdot 10^{-6}\,\frac{\mathrm{kg}}{\mathrm{mm}^3}$. With this as a basis, the simulation consists of a natural frequency analysis and afterwards a harmonic response analysis. A description of the meshing routine can be found in section 4. In the harmonic response simulation damping is introduced by employing a Rayleigh damping. For this a damping constant of $\zeta = 0.02$ is introduced for a minimal frequency of 50.00 Hz and a maximal frequency of 1800.00 Hz, which leads to the Rayleigh constants $\alpha_R = 1.946$ and $\beta_R = 2.162 \cdot 10^{-5}$ for the mass and stiffness matrices respectively.
As boundary conditions, a fixation is applied as Dirichlet boundary on the left edge of the macroscopic component, e.g. the outer left edge of the encasing. On the right hand side of the macroscopic component a point P is set outside the structure as seen in Fig.2. This point is connected rigid to the outer right edge of the encasing. To this point Neumann boundary conditions are applied. A sinusoidal and frequency-dependent oscillating force in x- and y-direction of both 60.00 N and a sinusoidal and frequency-dependent oscillating moment around the z axis of 10.00 N mm.

4 MESH CONVERGENCE

The meshing of the structures is done in Abaqus. The whole structure is modeled using two-dimensional solid plane stress elements, also called CPS6 elements from the Abaqus standard library. Before the mesh is applied, the structure is partitioned into different

regions. This is done so that the notches of unit cells are meshed fine, while the mesh fades out from there to the global mesh of the unit cells and lastly to the coarse mesh of the encasing. Figure 3 shows these partitions and mesh sizes.

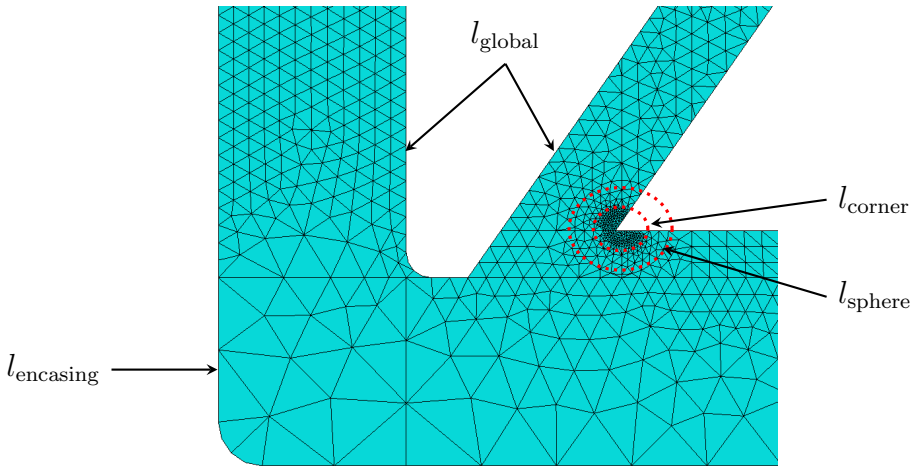


Figure 3: Detail of the mesh structure with the different mesh regions highlighted.

This approach stems from the study's focus on the auxetic mesostructure rather than the casing. To evaluate the influence of this partitioning and the mesh itself, a mesh convergence study is conducted. Therefore, the previously described procedure is applied, with the actual element sizes proportionally scaled. The resulting number of nodes and the corresponding characteristic element lengths are presented in Tab.2. The convergence study is performed on the initial geometry, without any optimization-related modifications.

Table 2: Mesh parameters for the mesh convergence study.

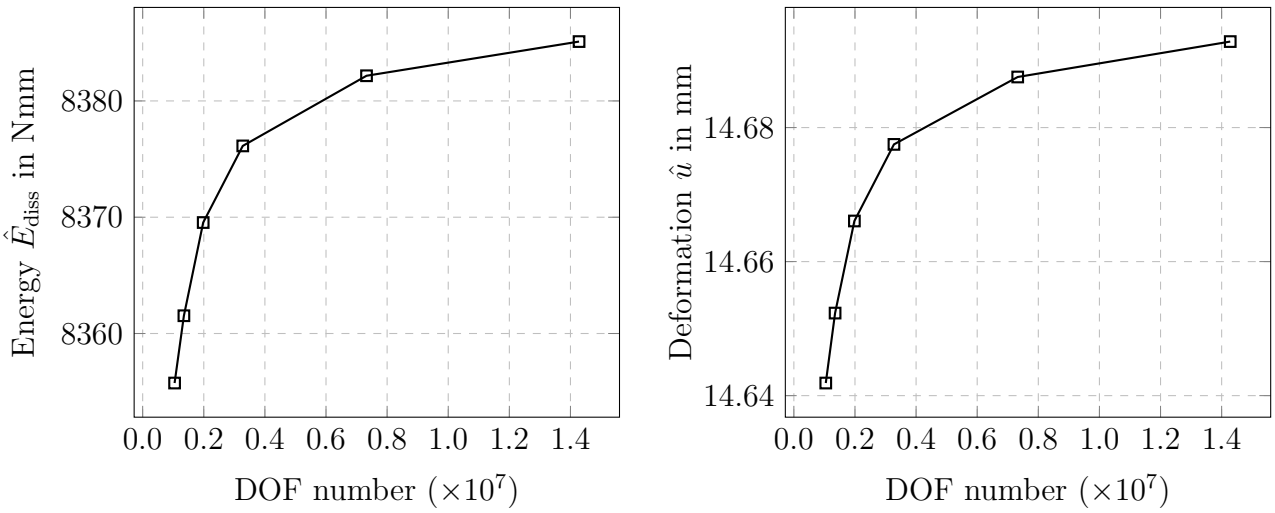
DOF number rounded	l_{global} in mm	l_{corner} in mm	l_{sphere} in mm	l_{encasing} in mm
$1.044 \cdot 10^6$	0.90	0.15	0.36	3.30
$1.349 \cdot 10^6$	0.75	0.125	0.30	2.75
$1.980 \cdot 10^6$	0.60	0.10	0.24	2.20
$3.275 \cdot 10^6$	0.45	0.075	0.18	1.65
$7.325 \cdot 10^6$	0.30	0.05	0.12	1.10
$1.428 \cdot 10^7$	0.21	0.035	0.084	0.77

The main results of these variations are presented in Fig.4. Here, the viscously dissipated energy $\hat{E}_{\text{diss}}$, summed over the frequency band f, and the total deformation magnitude $\hat{u}$, again summed over the frequency band f, at the point of load application are plotted against the number of degrees of freedom (DOF). This is shown in Eq.(1) and Eq.(2).

$$\hat{E}_{\text{diss}} = \sum_{f} E_{\text{diss}}(f) \tag{1}$$

$$\hat{u} = \sum_{f} \left[\sqrt{u^x(f)^2 + u^y(f)^2} \right] \tag{2}$$

These two parameters where chosen because they will also play a significant role in the goal function of the optimization, which is further discussed in section 5. It has to be mentioned here, that the convergence study could not be driven further because of limitations in hardware. And by examining the plots in detail in Fig.4, both show a clear convergence behaviour. Although both are not fully converged, the graph of energy dissipation is slightly more converged than the deformation graph. Despite the fact that both graphs are not fully converged, it must be taken into consideration that the absolute deviation from the finest to the coarsest mesh is very small. In numbers this can be expressed as 0.35% in terms of energy dissipation and 0.34% deviation in terms of deformation. This leads to the conclusion that the introduced fillets prevent the overpowering effect of the notches that usually occurs in this type of lattice structure. For this reason and the high computationally costs and limitations, the mesh with the DOF number $7.325 \cdot 10^6$ is chosen for the optimization simulations.

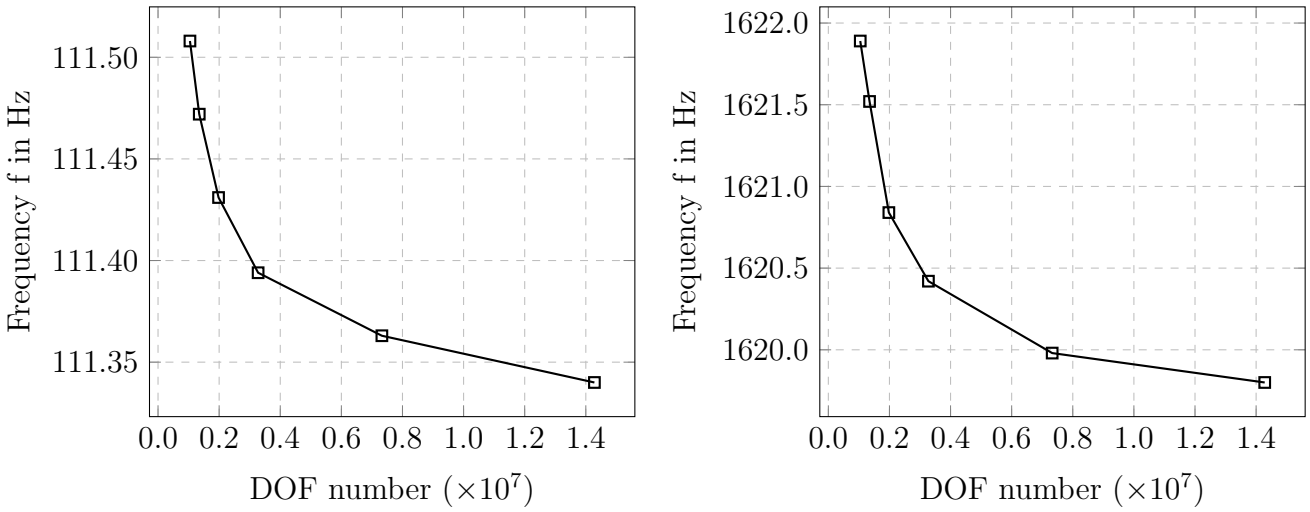


(a) Mesh convergence of the energy dissipation.

(b) Mesh convergence of the magnitude of deformation.

Figure 4: Results of the non-uniform mesh convergence study of the macro structure for the sum of the dissipated energy and the sum of the deformation in the point where the load is applied. The mesh parameters are listed in Tab.2.

In addition to the previous analysis, it has proven useful in past studies to consider the positions of the eigenfrequencies when assessing mesh convergence for these kind of problems. The results for this evaluation are shown in Fig.5 for the first and fifth eigenfrequencies, respectively. These values are selected from a total of six eigenfrequencies within the frequency band described in Section 3. A similar behavior to that observed in Fig.4 can be seen here. While complete convergence is not fully achieved, the absolute deviation between the coarsest and finest meshes remains minimal.



(a) Mesh convergence for the first natural frequency.

(b) Mesh convergence for the fifth natural frequency.

Figure 5: Results of the non-uniform mesh convergence study of the macro structure for the natural frequencies. The mesh parameters are listed in Tab.2.

5 OPTIMIZATION PROBLEM

The objective of this study is to determine an optimal configuration that minimizes a weighted objective function, incorporating displacement, mass, and dissipated energy. The problem can be formulated as follows:

$$\begin{aligned} \min_{p_i \in \mathbb{R}} \quad & G\left(p_i\right), \quad i = 1, 2, \ldots, 126 \\ \text{subject to} \quad & p_1 - \Delta t \leq p_1 \leq p_1 + \Delta t \\ & \vdots \\ & p_{126} - \Delta t \leq p_{126} \leq p_{126} + \Delta t \end{aligned} \tag{3}$$

where

$$G = \delta_u \cdot \sum_f \left[\sqrt{u_i^x(p_i, f)^2 + u_i^y(p_i, f)^2} \right] + \delta_m \cdot m(p_i), + \delta_E \cdot \sum_f E_{\text{diss}}(p_i, f) \tag{4}$$

In Eq.(4), G is the objective function, constructed as the sum over the frequency band of deformations in both spatial directions, u^x and u^y, at point P where the load is applied (cf. Fig.2), the mass m, and the sum over the frequency band of viscous dissipation energy $\hat{E}_{diss}$ of the whole model. The deformation serves as a substitute for the previously mentioned stiffness, providing an easily measurable quantity, one that is supposed to be minimized, rather than maximized. Additionally this also acts as a closer real life case. All terms in the equation are geometry dependent to some capacity and expressed as functions of the position of the geometry points p_i, which the optimization process can modify. These thickness changes are constrained by bounds to prevent the generation of non-physical geometries, such as intersections or holes. These bounds are

represented by the Δt_x, implemented into the geometry generation as shown in Eq.(3). Each of the three terms is assigned a weighting factor to normalize the initially unrelated values to a comparable scale and to ensure an unbiased optimization. The following values are therefore chosen:

$$\begin{aligned} \delta_u &= 600.00\,\frac{1}{\mathrm{mm}} \\ \delta_m &= 11\,000.00\,\frac{1}{\mathrm{kg}} \\ \delta_E &= -1.00\,\frac{1}{\mathrm{N\,mm}} \end{aligned} \tag{5}$$

By selecting the weighting factor of the dissipation energy as $\delta_E = -1.00\,\frac{1}{\mathrm{N\,mm}}$ this term acts as a kind of penalty term and allows for the maximization of this value.
For this optimization the python library Optuna [16] is used. Here, a Gaussian process-based Bayesian optimization is applied to the stated problem and the optimization is carried out in a loop with input and output to and from Abaqus simulations.

6 RESULTS

In the following chapter, the general results of the optimization study are discussed first and then the optimal solution of the problem is considered in detail.

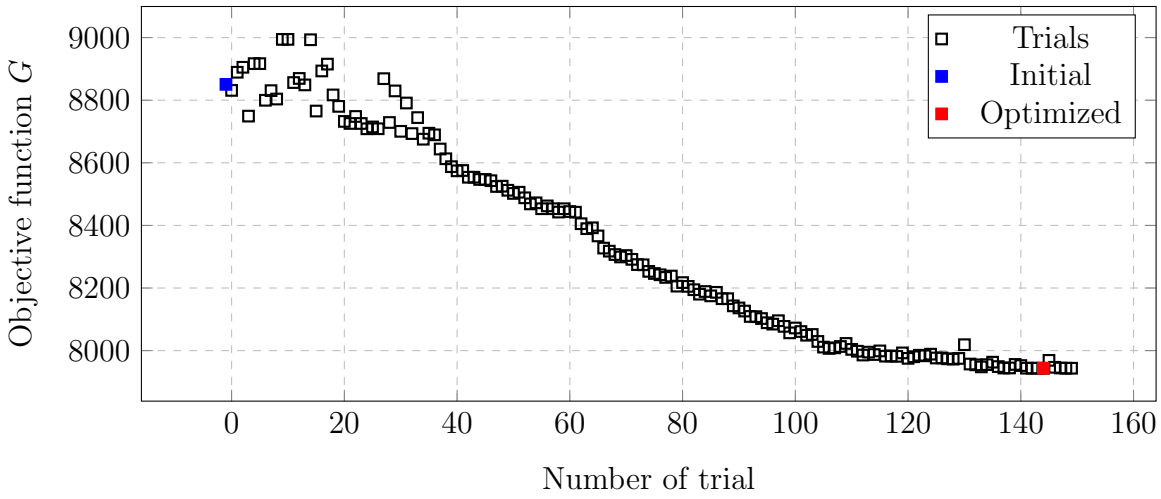
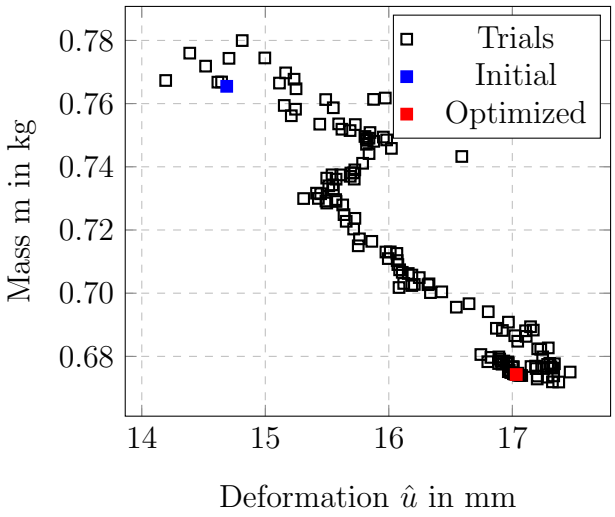


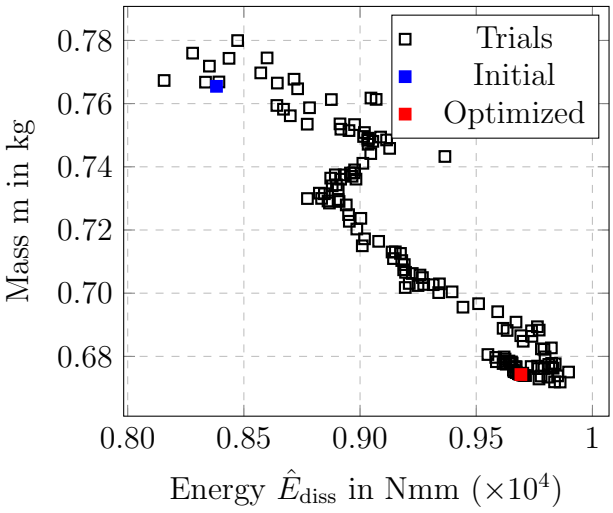
Figure 6: Objective function G plotted over the number of trials. The unchanged result of the Beam is plotted in blue and the best optimized result is plotted in red.

In Fig.6, the objective function G is plotted against the trial number. A total of 150 trials are performed until the optimization converges. In the early stages, the optimization exhibits significant oscillations, but after trial 37, the graph follows an almost linearly decreasing trend until trial 108. Beyond this point, the rate of decrease slows down, and convergence is reached at approximately 140 trials, with the best result occurring at trial 144. This plot clearly demonstrates that the optimized geometry (red) significantly outperforms the initial geometry (blue) in terms of the objective function, as defined

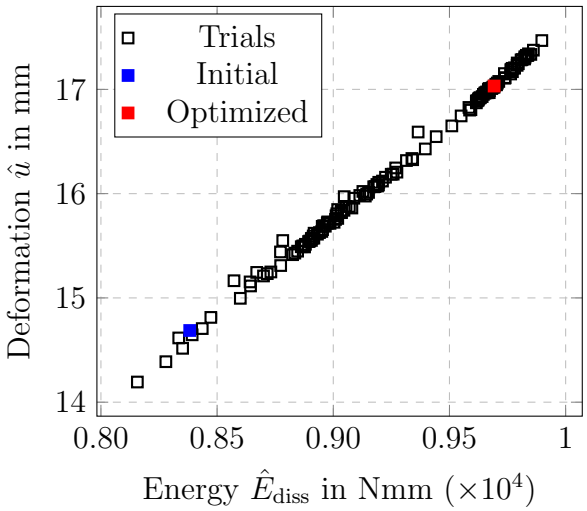
in Eq.(4). To further analyze the solution, Fig.7 presents a detailed view of the three parameters that constitute G. A clear trend can be observed in these results.



(a) Mass plotted over the sum of deformation.



(b) Mass plotted over the sum of viscous dissipation energy.



(c) Sum of deformation plotted over the sum of viscous dissipation energy.

Figure 7: Results of the optimization, divided in the mass, the sum of deformation and the sum of energy dissipation.

Mass is in opposition to both deformation and dissipation energy. As a result, in plots (a) and (b), a similar Pareto front is formed in both cases. In plot (c), an almost linear relationship can be observed between deformation and dissipation energy. By examining Eq.(6) more closely, a correlation between energy dissipation and deformation can be established:

$$E_{diss}(f) = \int_0^t \left(\int_V \sigma^v \cdot \dot{\varepsilon} \, dV \right) d\tau \tag{6}$$

where t is the time of one oscillation, V is the volume, σ^v represents the viscous stress,

and $\dot{\varepsilon}$ denotes the strain rate [17]. From this equation, it becomes evident that an increase in deformation leads to a corresponding increase in energy dissipation. Due to the fact that a greater oscillation amplitude requires covering a larger distance within the same time frame, leading to a higher strain rate and, typically, increased stress values, thereby resulting in greater energy dissipation. Since energy dissipation is included in the optimization as a kind of penalty term or, in other words, as a quantity to be maximized, the best trade-off is found at the lower end of the mass spectrum, where energy dissipation and deformation are relatively high. However, it is important to note that in many real-world applications, excessive deformation is often undesirable. Because the optimization does not inherently favor or constrain any particular parameter, the results do not necessarily reflect the constraints that might be required in an industrial setting.

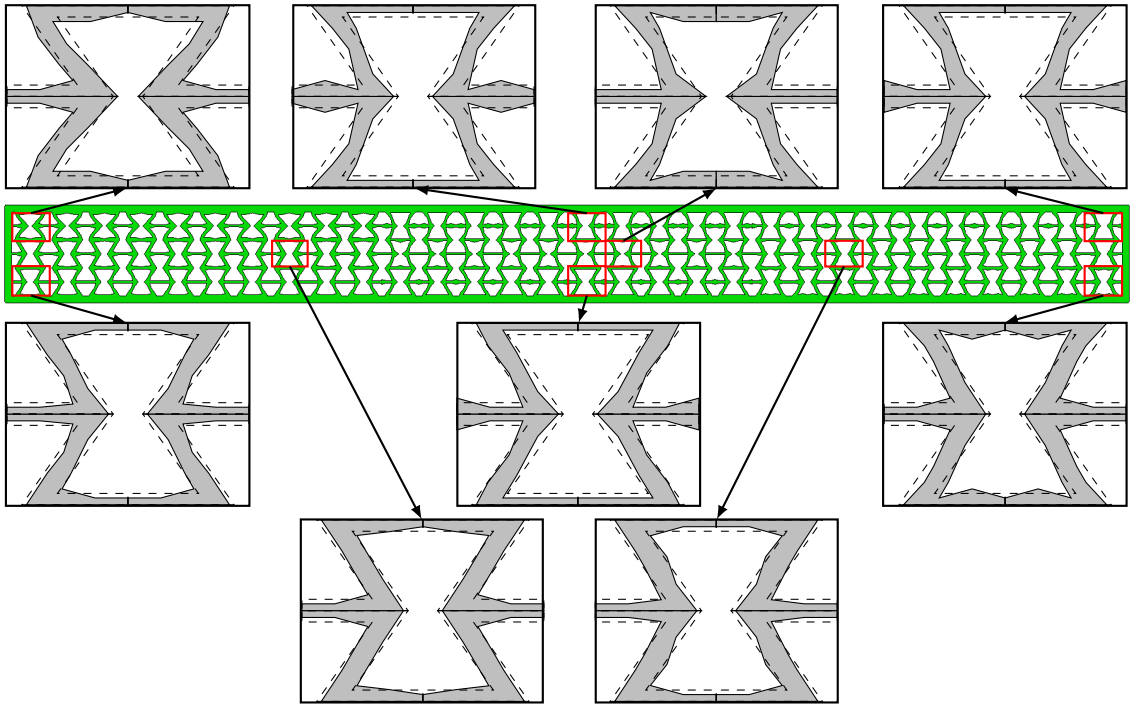


Figure 8: Optimized beam geometry with the nine individually optimized unit cells in detail. The details are in scale and the initial geometry is shown as a dashed line.

Figure 8 displays the optimized geometry of the cantilever beam with the mesostructure in detail. It is evident that the optimal result does not exhibit a homogeneous geometry across the macro component but instead varies locally. This variation is driven by the asymmetric mode shapes of the beam. Since the structure — and particularly the objective functions of energy dissipation and deformation — is strongly influenced by the mode shapes, the optimization naturally reflects these characteristics to some extent. The connection between mode shapes and optimization becomes intuitive when examining Tab.3.

By comparing Fig.8 to Tab.3, several conclusions can be drawn regarding the optimization. In the first column of unit cells neighboring the fixation point, the cells appear predominantly bulky, indicating the necessity for increased stiffness in this region. Based on the mode shapes, even when deflection occurs, no significant compression of the unit cells is observed in this area.

Table 3: Mode shapes of the initial beam and the optimized beam. The color scheme shows a qualitative distribution and not absolute values.

	Initial	Best
Mode-number	Eigenfrequency in Hz	
1	111.36	107.96
2	414.60	415.05
3	833.31	708.87
4	1229.82	1044.92
5	1619.98	1477.24
6	1641.16	1692.88

An outstanding feature in Fig.8 is that the unit cell shape in the middle row changes in terms of thickness, while the struts remain relatively straight. In contrast, in the top and bottom rows of the first column, the diagonal struts of the unit cells exhibit a more or less radially convex shape, whereas in the remaining unit cells of these rows, the diagonal struts display a radially concave shape. This behavior can be explained by the varying stiffness requirements for the optimized beam. It has already been shown that concave struts can, in some cases, have a stiffening effect [18]. However, depending on the parameters of the concave shape, they can also lead to a significant reduction in stiffness. This appears to be the case in this specific area. Referring back to Tab.3, the concave struts

tend to be compressed in the mode shapes, causing the beam thickness to shrink during bending, which significantly influences the mode shapes — particularly in modes three to six. This effect is so pronounced that mode five of the initial beam shifts to mode six in the optimized beam, as it transitions from a purely longitudinal mode to a hybrid mode combining bending and longitudinal vibration. By introducing this cross-sectional change the optimization forces the inclusion of unit cell behaviour in the macroscopic mode shape and therefore enhances the influences in energy dissipation and deformation drastically.

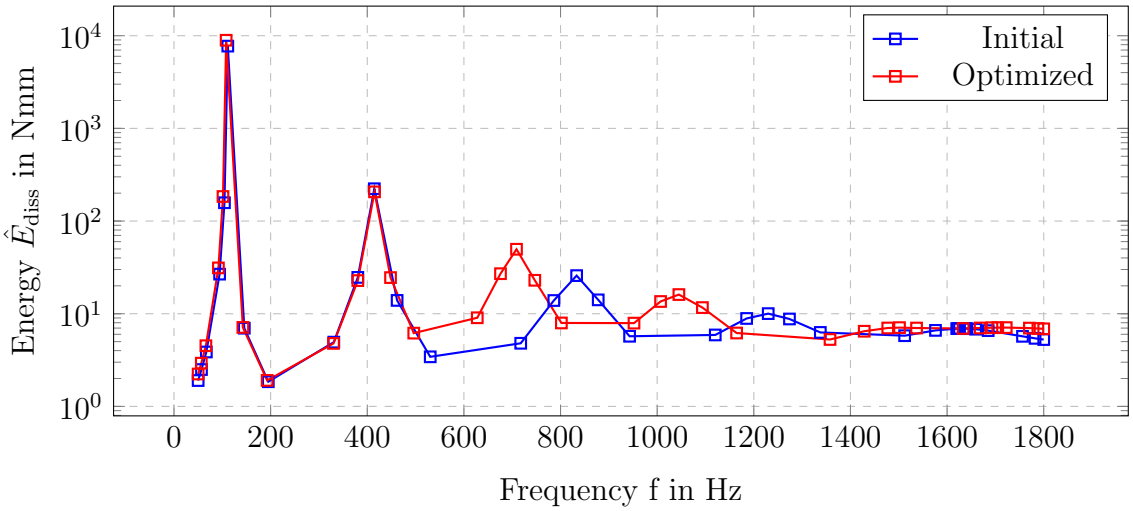


Figure 9: Viscous energy dissipation plotted over the frequency band for the initial and the optimized macro component.

To emphasize this point once again, it can be shown that only modes that allow a cross-sectional change after optimization are influenced by the optimization. Modes that fit this description are modes three to six, whereas modes one and two remain unaffected. By plotting the dissipation energy in the frequency band for both the initial and optimized beam, as exemplified in Fig.9, this theory can be confirmed. No change occurs for the first two modes, while a noticeable increase and shift in eigenfrequency and energy are observed in modes three to six. This aligns with the visual mode shapes in Tab.3, where only these mode shapes activate the mesostructure.

7 CONCLUSIONS

This work proposes a shape optimization method for macro components based on mesostructured auxetic unit cells. The overarching physical problem is formulated as a two-dimensional cantilever beam subjected to a dynamic load case within a frequency band of 50.00 Hz to 1800.00 Hz. The shape of locally specific unit cells is optimized and distributed over the surrounding area to achieve beneficial behavior in a realistic optimization setting. This includes reducing the mass of the component while enhancing its resistance to oscillations by minimizing deformation and maximizing energy dissipation. As would be the case for many aerospace related problems in the field. By doing so, significant improvements are achieved not only in the mode shape but also in the eigenfrequency, as well as in all of the aforementioned objectives. Furthermore, this method reveals a separation in eigenfrequencies and modes between macroscopic modes and mesoscopic — or

at least mesoscopically influenced — modes. These conclusions provide valuable insights into optimization possibilities for auxetic lattice structures in aerospace engineering and various other industrial applications.
Lastly, this study outlines potential directions for future work in the area of auxetic structures and optimization. To further enhance the effectiveness of the optimization, a two-scale optimization approach could be highly beneficial. In this approach, a larger scale could indicate areas requiring uniform optimization while identifying regions where structural differentiation is necessary.

SUPPLEMENTARY MATERIAL

The python code which was used to achieve the results shown in this study can be be accessed at DaRUS (`https://doi:10.18419/DARUS-4777`).

ACKNOWLEDGEMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project ID 498601949 – TRR 364 (Syntrac, Project B06). TR thanks the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for supporting this work via the following projects: 504766766 (Project Hybrid MOR). In addition, TR and LK are supported via the European Union and the German Federal Ministry for Economy and Climate Protection within the framework of the economic stimulus package no. 35c in module b on the basis of a decision by the German Bundestag by the project "DigiTain – Digitalization for Sustainability". TR, NW and NG are also grateful to the Baden-Württemberg Ministry of Science, Research and the Arts for financial support via the project "FLUTTER: Virtual flutter testing of next-generation aircraft prototypes: fundamental research on predictive and data-based methods", Project ID: FKZ Land: MWK32-7531-49/13/7

REFERENCES

[1] R. Lakes. Advances in negative poisson's ratio materials. *Advanced Materials*, 5(4): 293–296, 1993.

[2] L. Boldrin, S. Hummel, F. Scarpa, D. Di Maio, C. Lira, M. Ruzzene, C. Remillat, T. Lim, R. Rajasekaran, and S. Patsias. Dynamic behaviour of auxetic gradient composite hexagonal honeycombs. *Composite Structures*, 149:114–124, 2016. ISSN 0263-8223. doi:https://doi.org/10.1016/j.compstruct.2016.03.044. URL `https://www.sciencedirect.com/science/article/pii/S0263822316302082`.

[3] S. K. Sahu, N. D. Badgayan, S. Samanta, D. Sahu, and P. S. R. Sreekanth. Influence of cell size on out of plane stiffness and in-plane compliance character of the sandwich beam made with tunable pctpe nylon honeycomb core and hybrid polymer nanocomposite skin. *International Journal of Mechanical Sciences*, 148:284–292, 2018. ISSN 0020-7403. doi:https://doi.org/10.1016/j.ijmecsci.2018.08.011. URL `https://www.sciencedirect.com/science/article/pii/S0020740318312864`.

[4] M. Bianchi and F. Scarpa. Vibration transmissibility and damping behaviour for auxetic and conventional foams under linear and nonlinear regimes. *Smart Materials*

and Structures, 22(8):084010, jul 2013. doi:10.1088/0964-1726/22/8/084010. URL https://dx.doi.org/10.1088/0964-1726/22/8/084010.

[5] X. Shiyin, H. Xiuchang, and H. Hongxing. A study on the isolation performance of trichiral lattices with gradient geometry. *Journal of Vibration and Control*, 21(16): 3465–3475, 2015. doi:10.1177/1077546314524261. URL https://doi.org/10.1177/1077546314524261.

[6] Y. Chen, T. Li, F. Scarpa, and L. Wang. Lattice metamaterials with mechanically tunable poisson's ratio for vibration control. *Phys. Rev. Appl.*, 7:024012, Feb 2017. doi:10.1103/PhysRevApplied.7.024012. URL https://link.aps.org/doi/10.1103/PhysRevApplied.7.024012.

[7] C. Lira, F. Scarpa, and R. Rajasekaran. A gradient cellular core for aeroengine fan blades based on auxetic configurations. *Journal of Intelligent Material Systems and Structures*, 22(9):907–917, 2011. doi:10.1177/1045389X11414226. URL https://doi.org/10.1177/1045389X11414226.

[8] N. Khan and A. Riccio. A systematic review of design for additive manufacturing of aerospace lattice structures: Current trends and future directions. *Progress in Aerospace Sciences*, 149:101021, 2024. ISSN 0376-0421. doi:https://doi.org/10.1016/j.paerosci.2024.101021. URL https://www.sciencedirect.com/science/article/pii/S0376042124000472.

[9] M. Shruti, N. Sri Hemanth, N. D. Badgayan, and S. K. Sahu. Compressive behavior of auxetic structural metamaterial for lightweight construction using ansys static structural analysis. *Materials Today: Proceedings*, 38:12–17, 2021. ISSN 2214-7853. doi:https://doi.org/10.1016/j.matpr.2020.05.410. URL https://www.sciencedirect.com/science/article/pii/S2214785320339900. 2nd International Conference on Future Learning Aspects of Mechanical Engineering.

[10] H. G. Menon, S. Dutta, A. Krishnan, H. M. P., and B. Shankar. Proposed auxetic cluster designs for lightweight structural beams with improved load bearing capacity. *Engineering Structures*, 260:114241, 2022. ISSN 0141-0296. doi:https://doi.org/10.1016/j.engstruct.2022.114241. URL https://www.sciencedirect.com/science/article/pii/S0141029622003704.

[11] T. W. Tan, G. R. Douglas, T. Bond, and A. S. Phani. Compliance and longitudinal strain of cardiovascular stents: Influence of cell geometry. *Journal of Medical Devices*, 5(4):041002, 11 2011. ISSN 1932-6181. doi:10.1115/1.4005226. URL https://doi.org/10.1115/1.4005226.

[12] J. C. Álvarez Elipe and A. Díaz Lantada. Comparative study of auxetic geometries by means of computer-aided design and engineering. *Smart Materials and Structures*, 21(10):105004, jul 2012. doi:10.1088/0964-1726/21/10/105004. URL https://dx.doi.org/10.1088/0964-1726/21/10/105004.

[13] N. Bakhtiary, P. Allinger, M. Friedrich, F. Mulfinger, J. Sauter, O. Müller, and M. Puchinger. A new approach for sizing, shape and topology optimization. *SAE*

Transactions, 105:745–761, 1996. ISSN 0096736X, 25771531. URL http://www.jstor.org/stable/44734105.

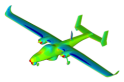
[14] X. Bin, C. Nan, and C. Huajun. An integrated method of multi-objective optimization for complex mechanical structure. *Advances in Engineering Software*, 41(2):277–285, 2010. ISSN 0965-9978. doi:https://doi.org/10.1016/j.advengsoft.2009.07.004. URL https://www.sciencedirect.com/science/article/pii/S0965997809001549.

[15] N. Grünfelder, B. P. Savall, S. M. Seyedpour, N. Waschinsky, and T. Ricken. Exploring the dependencies of poisson's ratio in auxetic structures. *PAMM*, 2024. ISSN 1617-7061. doi:10.1002/pamm.202400073.

[16] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

[17] M. Smith. *ABAQUS/Standard User's Manual, Version 6.9*. Dassault Systèmes Simulia Corp, United States, Providence, RI, 2009.

[18] T. Yuxin, Z. Yifeng, Z. Yilin, and L. Rong. Energy absorption characteristics and auxetic effect of novel elliptic-arc re-entrant honeycomb structures. *Engineering Structures*, 323:119260, 2025. ISSN 0141-0296. doi:https://doi.org/10.1016/j.engstruct.2024.119260. URL https://www.sciencedirect.com/science/article/pii/S0141029624018224.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AERODYNAMIC SHAPE OPTIMIZATION OF RAE-2822 AIRFOIL BY PHYSICS-INFORMED PARAMETRIC MODEL EMBEDDING

Damiano Squillace[1], Umberto Iemma[1], Domenico Quagliarella[2], Matteo Diez[3], and Andrea Serani[3]

[1]Department of Civil, Computer Science and Aeronautical Technologies Engineering
Università Roma tre
Via Vito Volterra, 62, 00146 Roma RM, Italy
{damiano.squillace,umberto.iemma}@uniroma3.it, https://www.uniroma3.it/en/

[2]CIRA — Italian Aerospace Research Centre
Via Maiorise, 1, 81043 Capua CE, Italy
d.quagliarella@cira.it, https://www.cira.it/en/

[3]National Research Council – Institute of Marine Engineering
Via di Vallerano, 139, 00128 Roma RM, Italy
{matteo.diez,andrea.serani}@cnr.it, https://www.inm.cnr.it/

Abstract. *The efficiency of optimization algorithms in engineering design is often hindered by the curse of dimensionality, where increasing problem complexity degrades algorithmic performance due to the high dimensionality of the design space. To address this challenge, dimensionality reduction techniques that minimize the number of variables or constrain their variability are essential. This study introduces physics-informed parametric model embedding (PI-PME), a novel approach specifically designed for aerodynamic shape optimization of transonic airfoils. PI-PME extends the parametric model embedding (PME) framework by integrating physics-based insights, enhancing its capability to guide the optimization process toward physically relevant variations in the design space. By leveraging principal component analysis, PI-PME efficiently reduces dimensionality while maintaining the ability to reconstruct variables in the original space, avoiding the limitations of purely geometry-based formulations. To demonstrate its potential, PI-PME is applied to the optimization of the RAE-2822 airfoil at Mach 0.3, a benchmark case in aerodynamic design. Results are compared against PME and full-domain optimization, showing that PI-PME not only reduces computational complexity but also identifies superior configurations in the reduced design space. These findings underscore the practicality and robustness of PI-PME for industrial applications requiring efficient and reliable optimization strategies.*

Keywords: Curse of Dimensionality, Dimensionality Reduction, Simulation-based Optimization, Aerodynamic Shape Optimization, Principal Component Analysis, Parametric Model Embedding

1 INTRODUCTION

In recent years, a new approach has emerged in the design and development cycle of industrial products, ranging from simple mechanical components to highly innovative aircraft. Whereas the construction of a physical or digital model was once considered the final step in the production process, it is now an integral part of the preliminary study, allowing significant reductions in both economic and time inefficiencies.

Simulation-Based Design Optimization (SBDO) involves defining an algorithm that, once the variables and their domain are set, enables the creation of a simulated version of the final product. By analysing the behaviour of this model under design conditions, key values can be extracted, leading to the optimization process, which focuses on identifying the design that provides the best outcome.

However, this process requires considerable computational time, especially when high-fidelity solvers are used to achieve the most accurate possible representation of physical phenomena. Fortunately, recent technological advancements have significantly reduced these times, enabling the study of problems characterized by an increasing number of variables. The expansion in the number of variables is particularly relevant for two reasons: the need to explore more complex domains in the search for previously unidentified optimal configurations and the possibility of analyzing highly innovative structures for which the necessary number of variables is not known a priori. However, this leads to the well-known Curse of Dimensionality (CoD) [1], where the performance of the optimization algorithm deteriorates as the design space domain increases. For this reason, it is crucial to study and develop techniques capable of reducing the dimensionality of the problem (Dimensionality Reduction, DR) before the optimization process, particularly in industrial contexts where available resources are limited. This challenge can be addressed in two ways: by reducing the total number of variables or by restricting the domain of individual variables.

The main methods currently in use are Sensitivity Analysis (SA) and Principal Component Analysis (PCA) [2]. In this work, the focus is on further developing the technique known as Parametric Model Embedding (PME) [3], which is based on PCA. PME offers the advantage of immediately identifying a correlation between the variables in the reduced space and those in the original space. This makes it a particularly effective tool in optimization processes, as it allows for easily constraining the problem to prevent invalid configurations. This method is also particularly useful for applications such as CAD, which rely on a predefined structure and cannot operate with variables in the reduced space. The objective of this work is to integrate the relevant physical information into PME to obtain a reduced design space that is potentially more effective in identifying the optimal configuration [4].

The paper is structured as follows: the second section introduces and explains the functioning of PME and its physics-informed version (PI-PME). Sections three and four present the case study, the tools used, and the various results obtained. Finally, the fifth section provides the conclusions.

2 DESIGN-SPACE DIMENSIONALITY REDUCTION

Consider a manifold $\mathcal{G}$, which represents the original geometry, whose coordinates in three-dimensional space are given by $\mathbf{g}(\boldsymbol{\xi}) \in \mathbb{R}^3$, where $\boldsymbol{\xi} \in \mathcal{G}$ are the curvilinear coordinates defined on $\mathcal{G}$. Assume that, for the purpose of shape optimization, $\mathbf{g}$ can be

transformed into a deformed shape/geometry $\mathbf{g}'(\boldsymbol{\xi}, \mathbf{u})$ through:

$$\mathbf{g}'(\boldsymbol{\xi}, \mathbf{u}) = \mathbf{g}(\boldsymbol{\xi}) + \boldsymbol{\delta}(\boldsymbol{\xi}, \mathbf{u}) \quad \forall \boldsymbol{\xi} \in \mathcal{G} \tag{1}$$

where $\boldsymbol{\delta}(\boldsymbol{\xi}, \mathbf{u}) \in \mathbb{R}^3$ represents the resulting shape modification vector, which is determined by a chosen shape parameterization or modification technique, and $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^M$ is the design variable vector. Consider the identification of the optimal design through the following SBDO optimization problem.

$$\underset{\mathbf{u}\in\mathcal{U}}{\text{minimize}} f(\mathbf{u}) \tag{2}$$

where $\mathbf{u}$ is treated as an uncertain or random parameter due to epistemic uncertainty. The optimal design is unknown before optimization, but exists. The vector of the design variable $\mathbf{u}$ is assigned a probability density function (PDF) $p(\mathbf{u})$, reflecting the belief in finding the optimal solution within a specific region of the design space. This PDF can be based on prior knowledge or experience; if unavailable, a uniform distribution may be used. Once $p(\mathbf{u})$ is defined, the shape modification vector $\boldsymbol{\delta}$ becomes stochastic and can be analyzed as a random field, for example, using the Karhunen-Loève Expansion (KLE), equivalent to Proper Orthogonal Decomposition (POD).

2.1 From Karhunen-Loève Expansion to Parametric Model Embedding

The KLE provides an optimal basis for the linear representation of a random process in a Hilbert space. The goal is to decompose the vector $\hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) = \boldsymbol{\delta} - \langle \boldsymbol{\delta} \rangle$ into orthonormal basis functions $\boldsymbol{\phi}_k(\boldsymbol{\xi})$, such that the geometric variance is retained. The expansion is expressed as:

$$\hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) \approx \sum_{k=1}^{N} x_k(\mathbf{u}) \boldsymbol{\phi}_k(\boldsymbol{\xi}), \tag{3}$$

where $x_k(\mathbf{u})$ are the projection coefficients obtained by:

$$x_k(\mathbf{u}) = (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho = \int_{\mathcal{G}} \rho(\boldsymbol{\xi}) \hat{\boldsymbol{\delta}}(\boldsymbol{\xi}, \mathbf{u}) \cdot \boldsymbol{\phi}_k(\boldsymbol{\xi}) \, \mathrm{d}\boldsymbol{\xi}. \tag{4}$$

The KLE aims to retain the maximum variance of the random process through the basis functions $\boldsymbol{\phi}_k(\boldsymbol{\xi})$. The total variance σ^2 is given by:

$$\sigma^2 = \sum_{k=1}^{\infty} \langle (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho^2 \rangle. \tag{5}$$

The optimal basis functions $\boldsymbol{\phi}_k$ are solutions to the variational problem, maximizing the functional:

$$\begin{aligned} \underset{\boldsymbol{\phi}\in\mathcal{L}^2_\rho\mathcal{G}}{\text{maximize}} \quad & \mathcal{J}(\boldsymbol{\phi}) = \langle (\hat{\boldsymbol{\delta}}, \boldsymbol{\phi}_k)_\rho^2 \rangle \\ \text{subject to} \quad & (\boldsymbol{\phi}, \boldsymbol{\phi})_\rho^2 = 1 \end{aligned} \tag{6}$$

which leads to the eigenvalue equation:

$$\mathcal{L}\boldsymbol{\phi}(\boldsymbol{\xi}) = \lambda \boldsymbol{\phi}(\boldsymbol{\xi}), \tag{7}$$

where λ are the eigenvalues representing the variance retained for each mode $\boldsymbol{\phi}_k(\xi)$.

Upon discretizing the domain and applying Monte Carlo sampling to the random variables, it is possible to move from the continuous KLE formulation to PCA. In PCA, having $\mathbf{d}(\boldsymbol{\xi}, \mathbf{u})$ as the discretization of $\boldsymbol{\delta}(\boldsymbol{\xi}, \mathbf{u})$ and sampling $\mathcal{U}$ by a statistically convergent number of realizations S, it is possible to define the data matrix $\mathbf{D}$:

$$\mathbf{D} = \begin{bmatrix} \hat{\mathbf{d}}_1^{(1)} & \dots & \hat{\mathbf{d}}_1^{(S)} \\ \vdots & \vdots & \vdots \\ \hat{\mathbf{d}}_n^{(1)} & \dots & \hat{\mathbf{d}}_n^{(S)} \end{bmatrix} = \begin{bmatrix} \hat{d}_{1,\xi_1}(\mathbf{u}_1) & & \hat{d}_{1,\xi_1}(\mathbf{u}_S) \\ \vdots & & \vdots \\ \hat{d}_{L,\xi_1}(\mathbf{u}_1) & & \hat{d}_{L,\xi_1}(\mathbf{u}_S) \\ \vdots & \dots & \vdots \\ \hat{d}_{1,\xi_n}(\mathbf{u}_1) & & \hat{d}_{1,\xi_n}(\mathbf{u}_S) \\ \vdots & & \vdots \\ \hat{d}_{L,\xi_n}(\mathbf{u}_1) & & \hat{d}_{L,\xi_n}(\mathbf{u}_S) \end{bmatrix} \tag{8}$$

where $\hat{\mathbf{d}} = \mathbf{d} - \langle \mathbf{d} \rangle$ (with $\langle \cdot \rangle$ the mean value) and $\hat{d}_{i,\xi_j}$ representing the j-th component of the shape modification vector associated with the element i.

To further improve the model's usability, especially in optimization contexts, the PCA methodology is extended into PME. The challenge in KLE/PCA-based dimensionality reduction lies in transforming back to the original design variables, which is addressed by PME. This is achieved by embedding the original design variables along with the shape modification vector in the data matrix $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \end{bmatrix} \quad \text{with} \quad \mathbf{U} = \begin{bmatrix} \hat{\mathbf{u}}^{(1)} & \dots & \hat{\mathbf{u}}^{(S)} \end{bmatrix} = \begin{bmatrix} \hat{u}_{1,1} & & \hat{u}_{1,S} \\ \vdots & \dots & \vdots \\ \hat{u}_{M,1} & & \hat{u}_{M,S} \end{bmatrix} \tag{9}$$

where $\mathbf{U}$ is the matrix of the original design variables to witch is associated a null weight such that

$$\mathbf{W}_\mathrm{u} = 0 \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_\mathrm{d} & 0 \\ 0 & \mathbf{W}_\mathrm{u} \end{bmatrix} \tag{10}$$

The eigenvalue problem for PME is then formulated as follows:

$$\mathbf{AGWZ} = \mathbf{Z\Lambda} \quad \text{with} \quad \mathbf{A} = \frac{1}{S}\mathbf{P}\mathbf{P}^\top. \tag{11}$$

where

$$\mathbf{G} = \begin{bmatrix} \mathbf{G_d} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \quad \text{and} \quad \mathbf{Z} = \begin{bmatrix} \mathbf{z}_1 & \dots & \mathbf{z}_S \end{bmatrix} \quad \text{with} \quad \mathbf{z}_k = \begin{bmatrix} \mathbf{e}_k \\ \mathbf{v}_k \end{bmatrix} \tag{12}$$

Here, $\mathbf{e}_k$ and $\mathbf{v}_k$ represent the eigenvector components associated to the shape modification $\mathbf{d}$ and design variable $\mathbf{u}$ vectors, respectively. Both $\mathbf{G_d} = \mathrm{diag}(\mathbf{G}_1, ..., \mathbf{G}_n)$ and $\mathbf{W_d} = \mathrm{diag}(\mathbf{W}_1, ..., \mathbf{W}_n)$ are block diagonal matrices with dimensionality $[nL \times nL]$, each block $[L \times L]$ being a diagonal matrix itself.

$$\mathbf{G}_k = \mathrm{diag}\left(\Delta\mathcal{G}_1, \dots, \Delta\mathcal{G}_L\right) \quad \text{and} \quad \mathbf{W}_k = \mathrm{diag}\left(\rho_1, \dots, \rho_L\right) \tag{13}$$

The solutions λ_k and $\mathbf{z}_k$ of Equation (11) are finally used to construct the reduced dimensionality representation of the original parameterization, and the reconstruction of the original variables is achieved as follows

$$\mathbf{u} \approx \hat{\mathbf{u}} = \langle \mathbf{u} \rangle + \sum_{k=1}^{N} x_k \mathbf{v}_k. \tag{14}$$

2.2 Physics Informed Parametric Model Embedding

The incorporation of physics-related information into the optimization problem represents the final step of the study proposed in this paper. Similarly to what was described in the previous section, physics is implemented in PME by defining a new matrix $\mathbf{P}_I$ as:

$$\mathbf{P}_I = \begin{bmatrix} \mathbf{D} \\ \mathbf{U} \\ \mathbf{F} \\ \mathbf{C} \end{bmatrix} \quad \text{with} \quad \mathbf{F} = \begin{bmatrix} | & \vdots & | \\ \hat{\mathbf{f}}_j^{(1)} & \dots & \hat{\mathbf{f}}_j^{(S)} \\ | & \vdots & | \end{bmatrix} \quad \text{and} \quad \mathbf{C} = \begin{bmatrix} | & \vdots & | \\ \hat{c}_j^{(1)} & \dots & \hat{c}_j^{(S)} \\ | & \vdots & | \end{bmatrix} \tag{15}$$

where $\mathbf{F}$ is a matrix containing distributed physical information (with $\hat{\mathbf{f}}_j = \mathbf{f}_j - \langle \mathbf{f}_j \rangle$) and $\mathbf{C}$ contains lumped or scalar quantities of interest $\hat{c}_j = c_j - \langle c_j \rangle$). As seen before, a block-diagonal weight matrix $\mathbf{W}_I$ is introduced:

$$\mathbf{W}_I = \begin{bmatrix} \mathbf{W}_\mathrm{d} & 0 & 0 & 0 \\ 0 & \mathbf{W}_\mathrm{u} & 0 & 0 \\ 0 & 0 & \mathbf{W_f} & 0 \\ 0 & 0 & 0 & \mathbf{W}_c \end{bmatrix} = \mathrm{diag}(\mathbf{W_d}, \underbrace{0}_{\mathbf{u}}, \mathbf{W_f}, \mathbf{W}_c), \tag{16}$$

where the new physical information, distributed $\mathbf{F}$ and scalar $\mathbf{C}$, are weighted respectively through the new matrices $\mathbf{W_f}$ and $\mathbf{W}_c$. Each block has to be set up to normalize its respective data by the inverse of the estimated variance, ensuring that geometric modifications and physical data all contribute comparably to the subsequent PCA procedure. A further formulation of Eq. 11 is than obtained:

$$\mathbf{A}_I \mathbf{G}_I \mathbf{W}_I \mathbf{Z}_I = \mathbf{Z}_I \mathbf{\Lambda}_I \tag{17}$$

where

$$\mathbf{A}_I = \frac{1}{S} \mathbf{P}_I \mathbf{P}_I^\mathsf{T}, \quad \mathbf{G}_I = \begin{bmatrix} \mathbf{G_d} & 0 & 0 & 0 \\ 0 & \mathbf{I} & 0 & 0 \\ 0 & 0 & \mathbf{G_f} & 0 \\ 0 & 0 & 0 & \mathbf{I} \end{bmatrix}, \quad \mathbf{Z}_I = \left[\mathbf{z}_{I,1} \dots \mathbf{z}_{I,S} \right], \quad \mathbf{z}_{I,k} = \begin{bmatrix} \mathbf{e}_k \\ \mathbf{v}_k \\ \boldsymbol{\phi}_k \\ \boldsymbol{\pi}_k \end{bmatrix}_I \tag{18}$$

In this context,$\mathbf{G_f}$ represents the element size of the distributed physical vector, which does not necessarily coincide with $\mathbf{G_d}$. Meanwhile, the identity block is also applied to the lumped scalars. The eigenvector solution $\mathbf{Z}_I$ subsequently provides a reduced representation of all data, where $\boldsymbol{\phi}_k$ and $\boldsymbol{\pi}_k$ serve as the eigenvector components embedding the distributed and lumped physical parameters. As a result, a lower-dimensional yet physics-informed space is obtained, facilitating further design-space exploration and optimization.

3 TEST CASE: OPTIMIZATION OF THE RAE-2822

The chosen problem aimed to optimize the shape of the RAE-2822 airfoil, which represents a classical test case for evaluating the performance of optimization algorithms in the aeronautical field. The operating conditions under which the experiments were conducted are Mach equal to 0.3 and Reynolds equal to 6.5×10^6, while the optimization problem was defined as follows:

$$\begin{aligned} \text{minimize} \quad & C_D(\mathbf{u}) \\ \text{subject to} \quad & \alpha = 0 \\ & C_L(\mathbf{u}) \geq 0.212 \\ & t/c = 0.1211 \\ & t_{85}/c \geq 0.02 \\ & \mathbf{u}_l \leq \mathbf{u} \leq \mathbf{u}_u \end{aligned} \tag{19}$$

where C_D represents the drag coefficient, while c denotes the airfoil chord. The decision was made to impose the constraint on the angle of attack α rather than on C_L due to time constraints related to the high-fidelity solver used (SU2). To still obtain meaningful results, a constraint on C_L was imposed to identify only configurations with aerodynamic performance better than the original airfoil. Additionally, constraints were imposed on the maximum thickness-to-chord ratio (t/c) and the minimum thickness at 85% of the chord.

The design space was defined by 20 variables, each associated with a different shape function applied to the upper and lower surfaces of the airfoil. In this case, the selected functions are: 3 polynomial functions, 6 Hicks-Henne bump functions, and one Wagner function, shown in Figure 1 and described as follows:

$$\begin{cases} x = (0.5x^3 - 1.5x^2 + x) \cdot 0.52 \\ x = -0.4x^2 + 0.x \\ x = (-0.5x^3 + 0.5x) \cdot 0.52 \end{cases}$$

$$\begin{cases} x = 0.888(\sqrt{x}(1-x)/e^{13.28x}) \\ x = 0.57(\sqrt{x}(1-x)/e^{5x}) \\ x = 0.1[sin(\pi x^{0.431})]^3 \\ x = 0.1[sin(\pi x^{0.757})]^3 \\ x = 0.1[sin(\pi x^{1.357})]^3 \\ x = 0.1[sin(\pi x^{3.106})]^3 \end{cases}$$

$$x = 0.87((\theta + sin(\theta))/\pi - sin(\theta/2)^2) \quad where \quad \theta = 2sin\sqrt{x}$$

The modification of the airfoil geometry is performed using a code developed by the Italian Aerospace Research Center (CIRA) called WG2AER, which parameterizes the airfoil as a linear combination of the baseline geometry and the modification functions.

To conduct this study, a well-defined methodological approach was followed. Initially, the necessary database was built to perform PME and PI-PME. To create it, a code was written that, by uniformly sampling each variable within its domain, was able to generate an S number of modified configurations of the original airfoil. Subsequently, these configurations were analyzed using the XFOIL software [6] to obtain crucial information

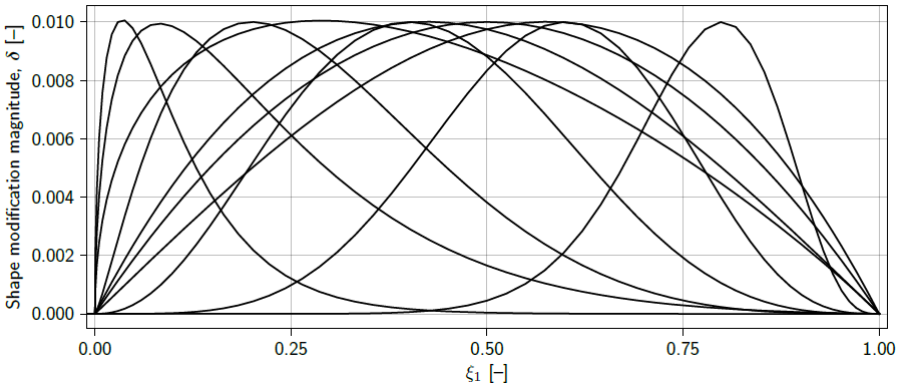


Figure 1: Shape modification functions used [5].

about the physics involved. Once the quality of the obtained dataset was verified, a MATLAB code was developed to perform PME and its physics-informed version. Finally, to optimize the aerodynamic profiles, a code developed by CNR based on Particle Swarm Optimization (PSO) [7] was used. During the optimization process, each configuration identified by individual particles was carefully studied and analyzed through simulations performed using the SU2 software [8]. The number of particles used for every optimization is equal to $4 * (\textbf{Number of variables})$, and the convergence criterion is set as the maximum number of function evaluations, equal to 1500.

3.1 Analysis of the Dimensionality Reduction Processes

Before optimizing the profile by leveraging reduced-dimensionality design spaces, the effect of individual physical information (in this case, C_P and C_D) on the new design space was analyzed.

The first noticeable aspect in the graphs in Figure 2 is that the exclusive use of C_D does not alter the solution of the eigenvalue problem in a way that provides a different number of reduced variables compared to the baseline case (PME).

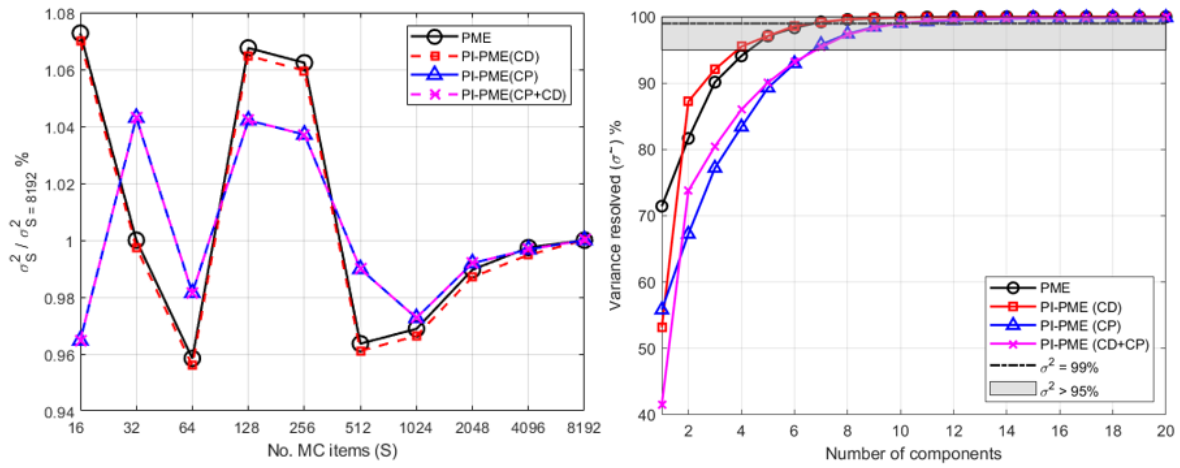


Figure 2: Comparison of the four studies conducted for variance convergence as a function of the MC elements used (left) and the variance retained by the eigenvectors (right).

A similar behavior can also be observed when comparing PI-PME (C_P) and PI-PME

($C_P + C_D$). Although the first eigenvector in the C_D-informed studies is capable of representing a significantly lower percentage of the total variance of the problem compared to its uninformed counterparts, from the second eigenvector onward, it becomes evident that the solution increasingly aligns with the "original" one. Regarding the determination of the number of Monte Carlo samples (S) to be used, referring again to Figure 2, it is clear that meaningful solutions can be obtained using 1024 configurations for any study considered. In this work, the maximum available number of MC samples was chosen since the time required to perform the DR varied imperceptibly as their number increased.

Based on the results obtained, three different optimizations were performed. In addition to the study in the original domain and in the domain obtained using PME, the final optimization was carried out in the reduced domain obtained by leveraging both available physical information. This choice proved to be the most appropriate, not only because the presence of C_D alone did not affect the number of reduced variables but also because it allowed for the analysis of the problem with the maximum amount of available information.

4 ANALYSIS OF OPTIMIZATION RESULTS

Once the optimization domains were selected, the aerodynamic performance of the original unmodified profile was initially evaluated. Subsequently, the three optimizations were performed in the various domains and finally a comparative analysis of the results was performed.

To allow a comparison with the results obtained through the optimization processes, an initial aerodynamic analysis of the unmodified RAE-2822 airfoil was conducted. The simulation performed using SU2 provided various useful outputs, summarized in Table 1, which allowed for the definition of the minimum required constraint C_L for the optimized airfoil, ensuring better efficiency compared to the initial configuration.

Table 1: Data related to the unoptimized RAE-2822 airfoil.

RAE 2822 - Original	
C_L	0.211748
C_D	0.0093607
C_M	-0.06556
E	22.62071
t_{85}/c	0.031766

4.1 Optimization in the original domain

The analysis in the full domain initially involved uniformly sampling the 20 variables between -2 and 2. However, this often resulted in invalid geometries, such as a particularly sharp Leading Edge (LE) or intersecting top and bottom surfaces. These profiles, which could not be evaluated using XFOIL, significantly slowed down the process of building the required database for the subsequent PI-PME. Thus, the range of the 20 variables was adjusted, as shown in Table 2, to minimize the occurrence of these invalid geometries. This adjustment represented the best compromise found to ensure the minimum number of intersecting geometries and an adequate variety of representable profiles during the

dataset creation.

Table 2: Definition of design variables in the full domain.

Variable	Upper bound	Lower bound
x_1	1.800	-1.500
x_2	1.800	-1.500
x_3	1.800	-1.500
x_4	1.800	-1.000
x_5	1.800	-1.000
x_6	1.800	-1.000
x_7	1.000	-0.500
x_8	1.000	-0.500
x_9	1.000	-0.500
x_{10}	1.800	-1.000
x_{11}	1.800	-1.500
x_{12}	1.800	-1.000
x_{13}	1.800	-1.500
x_{14}	1.800	-1.500
x_{15}	1.800	-1.500
x_{16}	1.800	-1.500
x_{17}	1.800	-1.000
x_{18}	1.800	-1.000
x_{19}	1.800	-1.500
x_{20}	1.800	-1.000

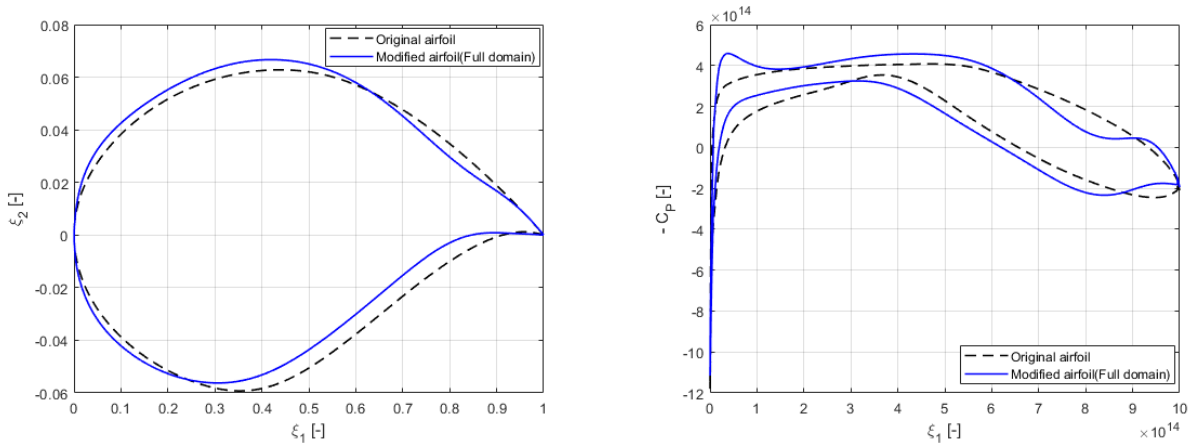
Once the variable domain, the number of particles used (80), and the maximum number of function evaluations were defined, the optimization process was started. After about 250 hours, the process reached the stopping condition.

The aerodynamic characteristics of the new airfoil (shown in Figure 3a) are summarized in Table 3. Notably, the drag coefficient is $C_D = 9.086 \cdot 10^{-3}$ and the lift coefficient is $C_L = 2.273 \cdot 10^{-1}$, representing an improvement of approximately 2.93% in C_D and 10.63% in efficiency.

Among the various results obtained from the simulations, an interesting comparison with the original airfoil is the distribution of the pressure coefficient (C_P) along the profile, as shown in Figure 3b.

Table 3: Data related to the RAE-2822 airfoil optimized using the original domain.

RAE 2822 - Full Domain	
C_L	0.227392
C_D	0.00908617
C_M	-0.0680325
E	25.02621
t_{85}/c	0.022638



(a) Comparison of geometries (full domain) (b) Comparison of C_P (full domain)

Figure 3: Results of the optimization in the full domain.

4.2 Optimization in the reduced domain via PME

The PME was trained using a set of $S = 8192$ MC elements. The initial geometry is described by $L = 364$ nodes, providing a data matrix P of dimensions $[748 \times 8192]$. The coefficients ρ_i that compose the matrices $\mathbf{W}$ were uniformly set to $1/\sigma_g^2$, where σ_g^2 denotes the variance of the shape modification vector.

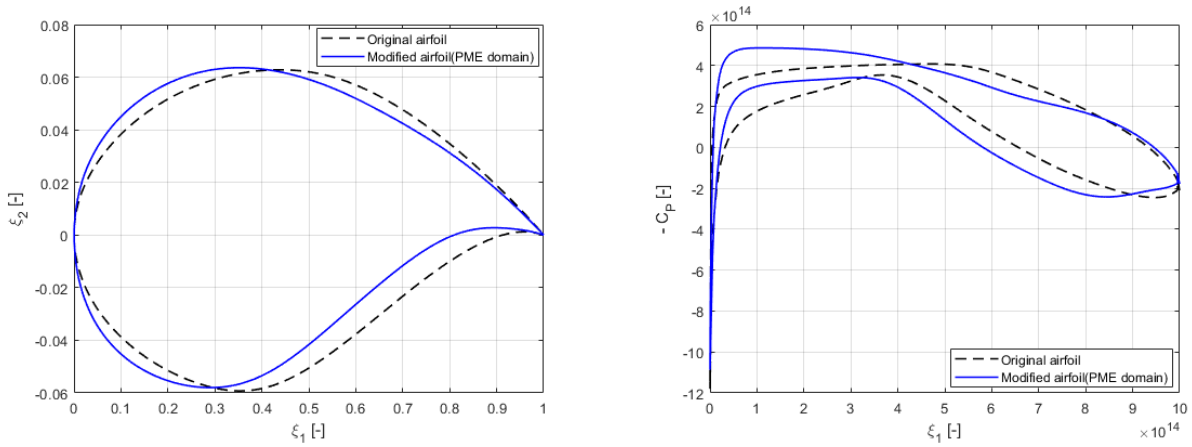
The number of reduced design variables required to retain at least 95% of the original geometric variance is $N = 5$, thus achieving a 75% reduction in dimensionality, whereas $N = 8$ reduced design variables are needed to preserve at least 99% of the geometric variance.

The optimization was configured by setting the number of particles to 20 to form the PSO swarm. The five new variables used are those obtained through PME contained in Table 4

Table 4: Definition of design variables in the reduced domain using PME.

Variable	Upper Bound	Lower Bound
x_1	3.49762	-3.73830
x_2	1.04279	-1.90947
x_3	1.24933	-1.16896
x_4	0.81568	-0.67094
x_5	0.69107	-0.92062

The new optimized geometry, shown in Figure 4a compared to the original configuration, is characterized by a $C_D = 9.058 \cdot 10^{-3}$ and a $C_L = 2.444 \cdot 10^{-1}$, indicating an improvement in C_D of 3.22% and a total efficiency increase of 19.3%. Other relevant parameters are listed in Table 5, while Figure 4b illustrates the distribution of C_P for each point of the profile compared to the original.



(a) Comparison of geometries (PME) (b) Comparison of C_P (PME)

Figure 4: Results of the optimization in the domain reduced via PME.

Table 5: Data related to the RAE-2822 airfoil optimized using PME.

RAE 2822 - PME	
C_L	0.24447
C_D	0.00905892
C_M	-0.0739756
E	26.9867
t_{85}/c	0.022577

4.3 Optimization in the reduced domain via PI-PME

The physical information used was extracted from simulations performed using XFOIL. As in the previous case, the PI-PME was trained using S = 8192 MC elements.

Since C_P is defined for each point of the profile and C_D is unique for each configuration, a matrix B of size [1113 x 8192] is obtained. Similarly to what was done with the matrix $\mathbf{W}$, the matrices $\mathbf{W_f}$ and $\mathbf{W}_c$, associated to C_P and C_D respectively, were also weighted with respect to their variances $\sigma^2_{C_P}$ and $\sigma^2_{C_D}$.

Table 6: Definition of design variables in the reduced space using PI-PME.

Variable	Upper bound	Lower bound
x_1	4.94878	-4.06976
x_2	2.60132	-5.56868
x_3	1.85518	-1.59955
x_4	1.883284	-1.9525
x_5	1.71518	-1.43197
x_6	1.18975	-1.47970
x_7	1.26270	-1.19997

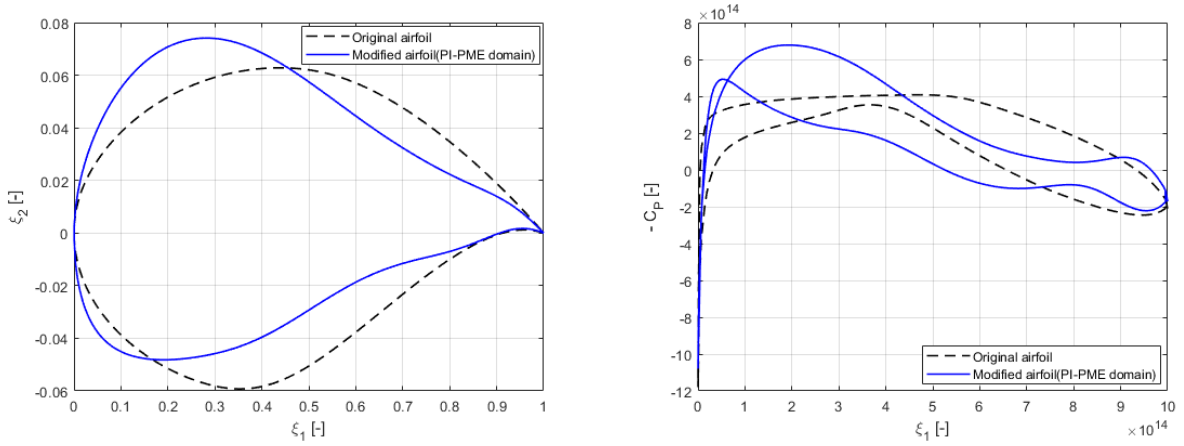
The number of reduced design variables required to retain 95% of the original variance is N = 7, while to retain 99% of the original variance, N = 10 is required.

In this case, the optimization was initiated by setting the number of particles to 28, and the new seven variables were defined as in Table 6.

The newly obtained geometry, visible in Figure 5a, defined by the parameters present in Table 7, is characterized by a $C_D = 8.941 \cdot 10^{-3}$ and a $C_L = 2.308 \cdot 10^{-1}$, which indicate an improvement of C_D by 4.48% and an overall increase in efficiency of 14. 10%. As with previous optimizations, Figure 5b illustrates the trend of C_P.

Table 7: Data related to the RAE-2822 airfoil optimized using PI-PME.

RAE 2822 - PI-PME	
C_L	0.23076
C_D	0.0089409
C_M	-0.054516
E	25.8093
t_{85}/c	0.022577



(a) Comparison of geometries (PI-PME) (b) Comparison of C_P (PI-PME)

Figure 5: Results of the optimization in the domain reduced via PI-PME.

4.4 Comparative Analysis of the Results

Figure 6, which shows the convergence of the optimization algorithms, clearly highlights how the implementation of PI-PME, compared to the study in the full domain (Full) and reduced through PME alone, has proven to be particularly advantageous.

Around the 180th iteration, a better configuration was identified compared to those obtained at the end of the other two tests, indicating a greater speed in identifying optimal geometries for the given problem. Furthermore, the configuration found upon reaching the imposed 1500 function evaluations is significantly better than the two previously obtained. A common feature observed across all three optimization cases, shown in Fig. 7, is the thinning of the lower surface. This suggests that the optimization process systematically reduces the thickness of this region and the extent of this modification varies among the different optimization domains.

In the "FULL" case, the geometric modifications appear to be more evenly distributed along the entire profile. The optimized shape closely follows the original contour but

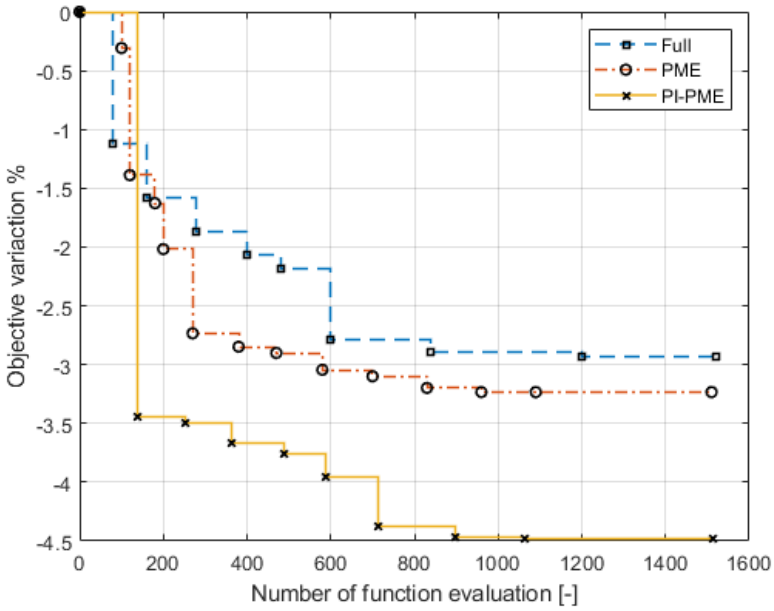


Figure 6: Comparison of convergence towards the optimal configuration identified for the three studies conducted.

with subtle refinements. The "PME" case, on the other hand, exhibits a more localized geometric transformation. While the overall shape remains similar to the original, certain regions show a more pronounced deviation, particularly on the upper surface. The last case displays the most significant alterations in both geometry and pressure distribution. The optimized profile exhibits greater curvature variations, especially along the upper surface, which could imply a more aggressive approach in this optimization domain. The C_P plot further supports this observation, as it presents a noticeable shift in pressure levels, particularly at the leading edge and peak suction regions.

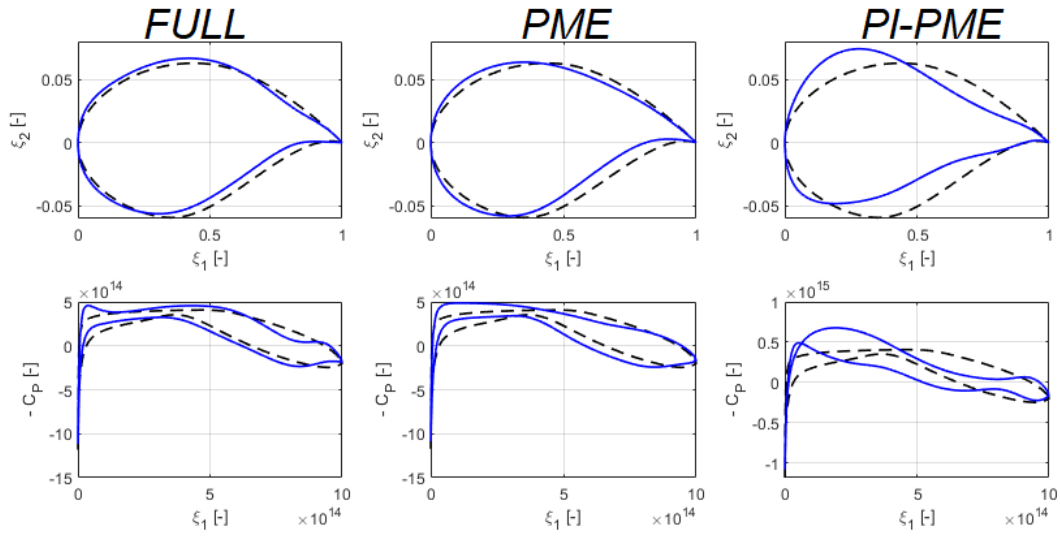


Figure 7: Comparison of the geometry and pressure coefficient for the three studies conducted.

5 CONCLUSIONS

In this work, the problem of shape optimization in engineering has been addressed through the use of DR techniques, taking into account the "Curse of Dimensionality" and presenting a solution based on PI-PME, an extension of PME. Through critical and experimental analysis, the effectiveness of this technique in tackling high-dimensional optimization problems has been demonstrated.

Not only did PI-PME accelerate the convergence of the optimization algorithm, but the final geometry also proved to be significantly better in terms of the objective function compared to those obtained within the complete and reduced domains using PME.

Regarding future developments, several interesting directions can be considered. First, the application of the PI-PME methodology could be extended to test cases with more challenging boundary conditions to further explore the limitations and potential of this technique. Additionally, conducting a study where the initial dataset is created using XFOIL under boundary conditions in which this software is known to produce accurate results would be of great importance. Once the reduced-dimensional space is built with the obtained data, it would be interesting to perform the airfoil optimization under more complex boundary conditions using specialized software such as SU2. This approach would maximize the strengths of XFOIL for generating a reliable starting point while leveraging SU2 for handling more complex flight conditions.

ACKNOWLEDGEMENTS

This work was carried out under the National Recovery and Resilience Plan (PNRR) scholarship on Methods and Models for Sustainable Engineering at the University of Roma Tre, according to the Ministerial Decree (DM) 117/2023, in collaboration with the Italian Aerospace Research Center (CIRA). It was also conducted within the framework of the NATO-AVT-404 activity on 'Enhanced Design Processes of Military Vehicles through Machine Learning Methods'.

REFERENCES

[1] R. Bellman. Dynamic programming. *Cambridge Studies in Speech Science and Communication. Princeton University Press*, 1957.

[2] A. Serani and M. Diez. A survey on design-space dimensionality reduction methods for shape optimization. *arXiv preprint arXiv:2405.13944*, 2024.

[3] A. Serani and M. Diez. Parametric model embedding. *Computer Methods in Applied Mechanics and Engineering*, 404, 2023.

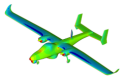
[4] A. Serani, G. Palma, M. Diez, J. Wackers, D. Quagliarella, and S. Gaggero. Enhancing parametric model embedding with physical information for effective vehicle design. In *Proceedings of the AVT-411 Research Workshop on Machine Learning and Artificial Intelligence for Military Vehicle Design*, Washington, USA, May 2025.

[5] A. Serani, M. Diez, and D. Quagliarella. Aerodynamic shape optimization in transonic conditions through parametric model embedding. *Aerospace Science and Technology*, 155:109611, 2024.

[6] M. Drela. XFOIL: An analysis and design system for low reynolds number airfoils. *Low Reynolds Number Aerodynamics*, 54, 1989.

[7] J. Kennedy and R. Eberhart. Particle swarm optimization. *Proc. Int. Conf. Neural Netw*, 4:1942–1948, 1995.

[8] T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAAJ*, 54:828–846., 2016.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

OPTIMAL WING FOR LONG ENDURANCE UAV APPLICATIONS

Abhiyan Paudel[1], Sharanjeet Kaur[1], Nikhil Vijay Shende[2*] and N Balakrishnan[1]

1: Computational Aerodynamics Lab
Department of Aerospace Engineering
Indian Institute of Science
Bangalore 560012, Karnataka, India
nbalak@iisc.ac.in, https://www.iisc.ac.in

2: S & I Engineering Solutions Private Limited
Bangalore 560024, Karnataka, India
nikhil.shende@sandi.co.in, https://www.sandi.co.in

Abstract. *This paper presents a methodology for aerodynamic optimization of a long endurance unmanned aerial vehicle (UAV) wing employing genetic algorithm (GA). It is well known that genetic algorithm–based optimization procedure requires a large number of function evaluations to arrive at a globally optimal configuration. In the present study, the function under consideration is the endurance parameter which depends on the lift and drag coefficients of the wing. Use of Computational Fluid Dynamics (CFD) tools for function evaluation is prohibitively expensive. To address this, a low–fidelity aerodynamic tool based on numerical lifting line theory (NLLT) working in conjunction with high–fidelity sectional aerodynamic data set generated using RANS solver is developed. NLLT significantly reduces the cost of function evaluation without compromising accuracy. The endurance parameter of the optimal wing obtained through a GA–based optimization procedure employing the NLLT tool is compared with CFD and the results show a very close match between them. This clearly brings out the efficacy of the present procedure to obtain an aerodynamically optimal wing for long endurance UAV configuration. In addition, the study also establishes the need for considering a large population for a reliable GA-based optimization process for achieving optimal solution.*

Keywords: Unmanned Aerial Vehicles, Aerodynamic shape optimization, Genetic algorithm, NACA 4–Digit Airfoils, Lifting Line Theory, Reynolds Averaged Navier–Stokes Equations

1 INTRODUCTION

Long endurance low–speed UAV which can deliver long flight duration has a wide range of applications in military, meteorological, maritime and civilian domains. The endurance parameter (E) of a typical UAV using turboprop engine is given as follows:

$$E = \frac{\eta_{pr}}{c} \sqrt{2\rho_\infty S} \frac{C_L^{\frac{3}{2}}}{C_D} \left(W_1^{\frac{-1}{2}} - W_0^{\frac{-1}{2}} \right) \tag{1}$$

where η_{pr} is the propulsive efficiency of the propeller, c is the specific fuel consumption, W_1 is the empty weight of the aircraft, W_0 is its full weight, C_L is lift coefficient, C_D is drag coefficient, S is wing planform area and ρ_∞ is free stream density. In equation (1), it may be seen that $E \propto C_L^{\frac{3}{2}}/C_D$. Therefore, in order to obtain high value of endurance parameter, the aerodynamic quantity $\left(C_L^{\frac{3}{2}}/C_D\right)$ needs to be maximized. This directly points towards choosing an optimal wing to obtain high endurance. The aerodynamic characteristics of a wing are primarily influenced by the choice of aerofoil section, aerodynamic twist, geometric twist and cruise incidence. These factors play a critical role in determining the lift and drag, which in turn impacts the aircraft's endurance. Achieving the right balance between these parameters can be a complex and challenging task. Therefore, this necessitates use of an optimization procedure to design a wing that provides higher endurance. With this motivation, the paper presents a methodology for aerodynamic shape optimization of a rectangular wing employing genetic algorithm (GA). A high aspect ratio wing is defined using six NACA 4–digit airfoil sections. A section is parameterized using maximum camber value and its position. The thickness of aerofoil section at wing root (t_{root}), ratio $r = t_{tip}/t_{root}$, twist at wing tip with respect to the root and cruise angle of attack are other design variables. Altogether, there are 16 design variables defining parametric space. While the parameters presented here are based on aerodynamic considerations, the methodology is general to allow extension for other crucial variables such as wing volume for fuel storage, structural weight etc.

There are a number strategies which have been developed in the field of optimization, but evolutionary algorithms stand out due to their simplicity and effectiveness in handling multiple objectives [1]. In the realm of aerodynamics, these algorithms have been extensively explored and a comprehensive review of these strategies can be found in reference [2]. Building upon these insights, the present work employs a genetic algorithm (GA) based optimization approach for wing design.

It is well known that genetic algorithm–based optimization requires a large number of function evaluations to obtain globally optimal configuration. As mentioned earlier, in this study, the function under consideration is endurance parameter which depends on lift and drag coefficients. Use of high–fidelity RANS–based CFD tool for function evaluation is prohibitively expensive. Hence, a low–fidelity aerodynamic tool based on numerical lifting line theory (NLLT) using a high–fidelity sectional aerodynamic data set generated using a RANS solver is developed. NLLT significantly reduces the cost of function evaluation without compromising the accuracy. The study also establishes the need for considering a large population for a reliable GA–based optimization process for achieving optimal and realistic solution.

2 PROBLEM DEFINITION

A genetic algorithm based optimization is performed using open source software DAKOTA [3] at cruise flight conditions associated with a low–speed UAV. The optimization problem is defined as follows:

$$\begin{aligned} \max_{\mathbf{x}} \quad & E_p\left(\mathbf{x}\right) = \frac{C_L\left(\mathbf{x}\right)^{\frac{3}{2}}}{C_D\left(\mathbf{x}\right)} \\ \text{w.r.t} \quad & \mathbf{x} \in \mathbb{R}^{N_{DV}} \\ \text{s.t.} \quad & C_{L_{min}} - C_L\left(\mathbf{x}\right) \leq 0 \end{aligned} \tag{2}$$

where $\mathbf{x}$ is the vector of design variables and N_{DV} is the number of design variables. The objective is to maximize the endurance parameter (E_p) subject to minimum allowable lift coefficient ($C_{L_{min}}$) equal to 0.8. A high aspect ratio wing, typically suited for long endurance configurations, is defined using NACA 4–digit airfoil sections. An aerofoil section is parameterized using three variables, namely, maximum camber (m), maximum camber position (p) and thickness (t). The values of maximum camber and maximum camber position are specified at six different locations along the wing half span including root and tip sections resulting in 12 design variables. The thickness of the aerofoil section at the wing root (t_{root}) is another design variable. To account for structural considerations, the thickness of the aerofoil section at the wing tip (t_{tip}) is determined by the ratio r such that $r = t_{tip}/t_{root}$, which is yet another design variable. Corresponding to a washout configuration, the twist at the wing tip (tw) with respect to that at the root is also included as a design variable. The cruise angle of attack (AOA) is yet another design variable that will dictate the optimal design. Altogether, there are 16 design variables defining parametric space $\mathbf{x}$. The lower and upper bounds of the design variables are specified in Table 1. While the aforementioned parameters are based on aerodynamic considerations, the beauty of this methodology lies in its generality, allowing for seamless extension to incorporate other crucial variables such as wing volume for fuel storage, structural weight etc.

Table 1: Lower and upper bounds of design variables

	m	p	t_{root}	r	tw	AOA
Lower Bound	0.02	0.2	0.12	0.8	-4^o	0^o
Upper Bound	0.08	0.8	0.24	1.0	0^o	6^o

3 METHODOLOGY

For the purpose of wing optimization, the GA module of DAKOTA is used. This module employs a set of gradient–free optimization algorithms known as COLINY [3]. The upper and lower bounds of all the design parameters are provided to DAKOTA along with population size and number of generations. For the populations created by DAKOTA, the function evaluation sought by the optimizer is provided by NLLT. At the end of the optimization process, DAKOTA provides information about the best function evaluation and its corresponding optimal parameters. Figure 1 shows the flow chart of the whole optimization cycle. The details of the procedures involved from the parameters provided by the optimizer to function evaluation are explained in the subsequent sections.

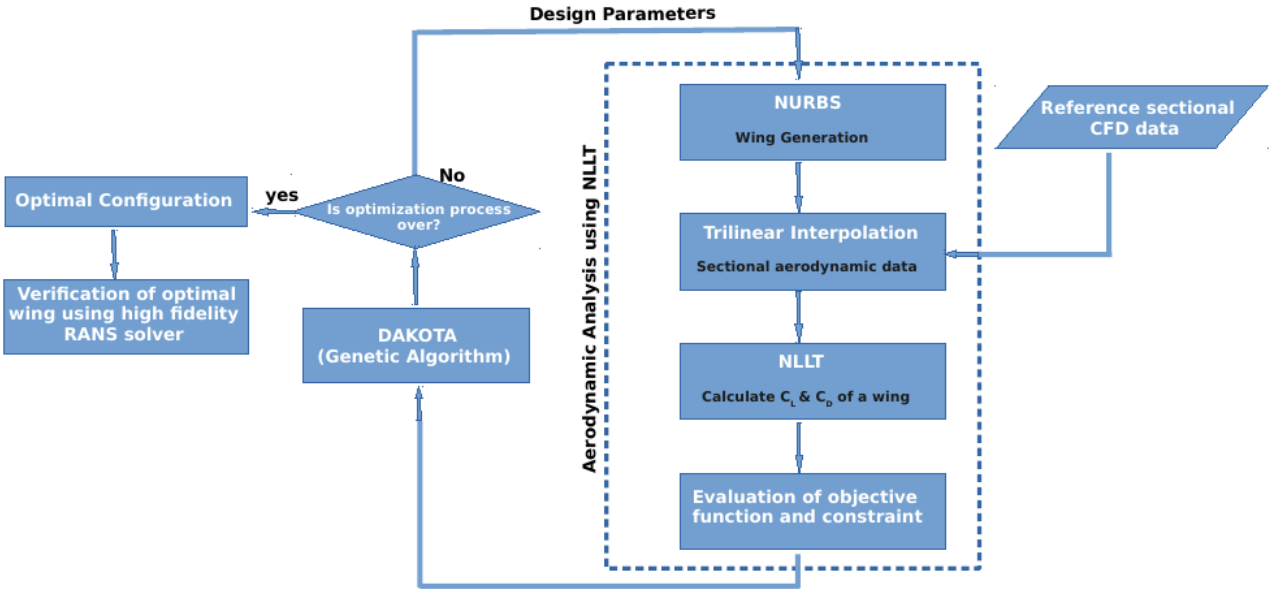


Figure 1: Optimization Flow Chart

3.1 Wing Definition

Once DAKOTA provides the values of maximum camber (m) and maximum camber position (p) at six locations along wing semi–span, these values are used to define Non–Uniform Rational B–Splines (NURBS) [4, 5] for both the parameters. These NURBS then evaluate the values of m and p at different locations, 51 sections in this study, along half wing span. Using t_{root} and r, the value of t_{tip} is calculated and the thickness of airfoil sections in the spanwise direction is distributed linearly along the wing half–span. Similarly, the twist is varied linearly along the wing half–span. These 51 sections along the wing half–span define the wing. A typical half–span wing is pictorially depicted in Figure 2.

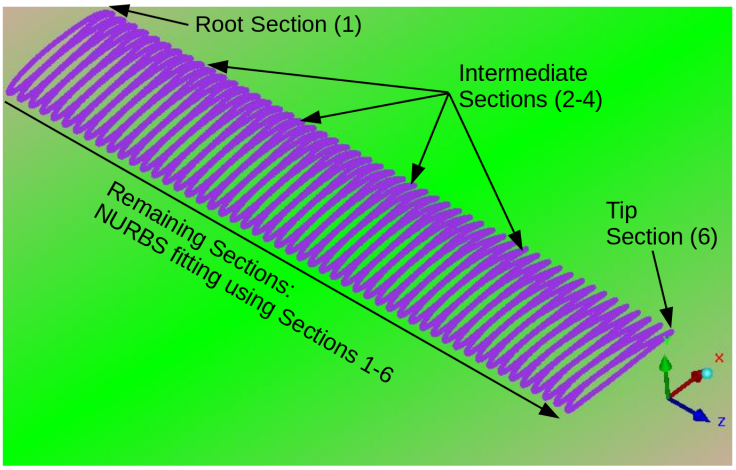


Figure 2: Wing Definition

3.2 Wing Aerodynamic Data Generation

The steps involved in calculating the lift and total drag coefficients for a wing using NLLT are as follows:

1. Simply stated, NLLT is a numerical variant of the classical lifting line theory of Prandtl [6, 7] and the readers are referred to earlier works for details of the methodology. Akin to the classical theory, NLLT also requires sectional lift and drag coefficients at a given effective angle of incidence, for computing the 3D lift and drag coefficients of the wing. The accuracy of the NLLT depends upon the number of such sections included in the analysis and in the present work 51 sections are considered to represent the wing semi–span.

2. At the heart of the present optimization procedure is the sectional aerodynamic data set generated using the RANS solver HiFUN [8] over the parametric space defined in Table 1 and depicted in Figure 3. Altogether sectional aerodynamic data for 45 aerofoils, with a 3–point resolution along m and p directions and 5–point resolution along the t direction are precomputed for $M_\infty = 0.2$, $Re_\infty = 4.65$ million and alpha ranging from -6 to 6 degrees. The sectional aerodynamic coefficients for an arbitrary airfoil sought by DAKOTA are obtained from the precomputed data set using a multi–step interpolation procedure.

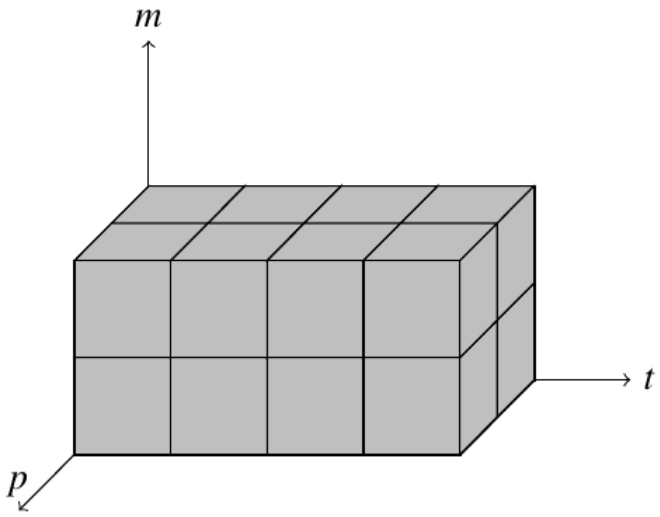


Figure 3: Parametric design space

3. In the multi–step interpolation procedure, the first step is to identify a cuboid in the parametric design space consisting of m, p and t in which the random aerofoil section lies. In the next step, two angles of attack on the C_l–α curve between which the effective angle of attack lies are identified. Then, at each of the two angles of attack, C_l and C_d are calculated for the random aerofoil employing a trilinear interpolation procedure using the C_l and C_d data available for the reference aerofoils at the corners of the cuboid (see Figure 4). Finally, the C_l and C_d values for the random aerofoil at the effective angle of attack are calculated using linear interpolation from the values interpolated at the two adjacent angles of attack obtained using trilinear

interpolation. The details of the trilinear interpolation procedure are explained in the next paragraph.

4. Consider a cuboid as shown in the Figure 4. The eight vertices of the cuboid represent eight design points. Each design point corresponds to a specific airfoil. Let C be the point where the sectional lift and drag coefficients (C_l and C_d) are sought. The point C corresponds to an arbitrary airfoil in the design space. Let (t, m, p) be the coordinates of point C. Let (t_0, m_0, p_0) and (t_1, m_1, p_1) be the coordinates of points having minimum and maximum parameter values respectively as shown in Figure 4. Define t_d, m_d and p_d as

$$t_d = \frac{t - t_0}{t_1 - t_0}, \quad m_d = \frac{m - m_0}{m_1 - m_0}, \quad p_d = \frac{p - p_0}{p_1 - p_0} \tag{3}$$

To begin with, interpolate the values along t direction:

$$C_{00} = P\left[t_0, m_0, p_0\right](1 - t_d) + P\left[t_1, m_0, p_0\right] t_d \tag{4}$$
$$C_{10} = P\left[t_0, m_1, p_0\right](1 - t_d) + P\left[t_1, m_1, p_0\right] t_d \tag{5}$$
$$C_{01} = P\left[t_0, m_0, p_1\right](1 - t_d) + P\left[t_1, m_0, p_1\right] t_d \tag{6}$$
$$C_{11} = P\left[t_0, m_1, p_1\right](1 - t_d) + P\left[t_1, m_1, p_1\right] t_d \tag{7}$$

where $P\left[t, m, p\right]$ corresponds to the value of lift or drag coefficient of the airfoil at the location (t, m, p). Further, interpolate along p direction,

$$C_0 = C_{00}(1 - p_d) + C_{01} p_d \tag{8}$$
$$C_1 = C_{10}(1 - p_d) + C_{11} p_d \tag{9}$$

Finally, interpolate along m direction,

$$V = C_0(1 - m_d) + C_1 m_d \tag{10}$$

where V corresponds to the value of the aerodynamic coefficients at the desired point C.

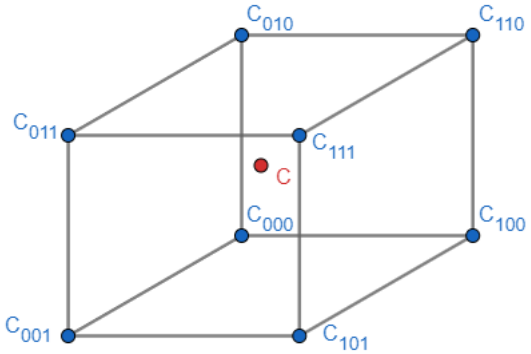


Figure 4: Geometry for Trilinear interpolation

4 RESULTS AND DISCUSSION

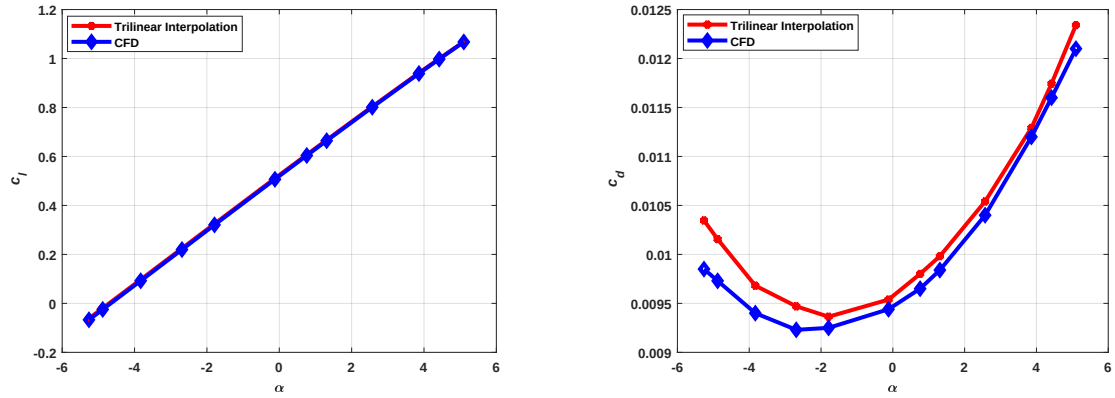
In this section, the results of the optimization study are presented. As discussed earlier, NLLT procedure is employed for function evaluation in the present work. NLLT uses sectional data obtained from a multi–step interpolation strategy. Hence, a validation study demonstrating efficacy of multi–step interpolation strategy and NLLT procedure is undertaken as presented in sections 4.1 and 4.2 respectively. Further, section 4.3 elucidates the effect of number of generations and population size on the evolution of optimal aerodynamic configuration suited for long endurance flight. This analysis establishes the efficacy and accuracy of the proposed methodology strengthening the reliability of present approach. This section also presents aerodynamic parameters pertaining to the optimal configuration.

At this stage, it is worth mentioning that the reference CFD data is obtained using the RANS solver HiFUN [8]. The steady state simulations are carried out using following algorithms:

- Roe [9] scheme for interfacial inviscid flux calculations,
- Green–Gauss theorem based diamond path reconstruction [10] for interfacial viscous flux calculation and solution reconstruction,
- Venkatakrishnan limiter [11] for preserving solution monotonicity,
- first order backward Euler [12] procedure for implicit time integration,
- symmetric Gauss–Seidel (SGS) iterative relaxation procedure [12] for solving system of linear and equations.
- Spalart-Allmaras [13] turbulence model for eddy viscosity computation.

4.1 Multi–step Interpolation Procedure

The multi–step interpolation procedure is validated by considering twenty (20) random aerofoils within the parametric design space. Aerodynamic data is generated for these aerofoils at twelve random angles of attack chosen between -6^o to 6^o using interpolation procedure and the HiFUN solver. The data obtained using both the procedures are compared to study the accuracy of the multi–step interpolation procedure in predicting aerodynamic force coefficients for arbitrary airfoil sections. The freestream Mach and Reynolds number chosen for the analysis are 0.2 and 4.65 million respectively. The $C_l - \alpha$ and $C_d - \alpha$ curves for a representative airfoil are shown in Figures 5–(a) and 5–(b) respectively. The results indicate that the interpolation procedure predicts the sectional lift and drag coefficients reasonably well in comparison to the HiFUN data. Similar accuracy levels are observed across all random aerofoils. This is evident from Figure 6 which depicts average error in lift and drag coefficients, calculated as a difference of corresponding coefficients between interpolated and CFD data, plotted verses the angle of attack. From this figure, it is seen that maximum average error for lift coefficient is less than 14 counts (1 lift count = 0.001) while maximum average error for drag coefficient is about 24 counts (1 drag count = 0.0001).



(a) $C_l - \alpha$ curve (b) $C_d - \alpha$ curve

Figure 5: Prediction of integrated coefficients for a random aerofoil given by $m = 0.0449$, $p = 0.454$ and $t = 0.118$

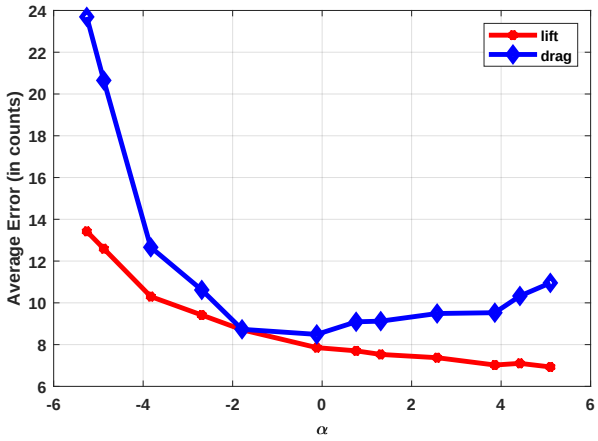


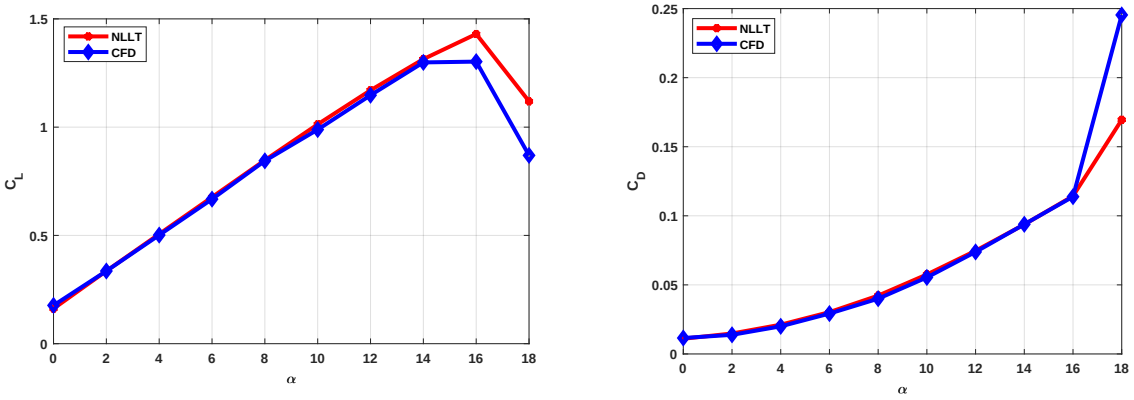
Figure 6: Average error in trilinear interpolation predictions for lift and drag coefficients for 20 random aerofoil sections

4.2 Numerical Lifting Line Theory

This study establishes the efficacy of NLLT by comparing its results with the results obtained using the HiFUN solver. Additionally, two dimensional sectional data required for NLLT is also generated using the HiFUN solver. The analysis focuses on an RAE–2822 wing at a freestream Mach number of 0.2 and Reynolds number of 6.5 million. In Figure 7, the lift and drag coefficients of the RAE–2822 wing are compared between the NLLT and HiFUN. The NLLT data is in good agreement with the HiFUN data, showcasing its capability to accurately predict the lift and drag coefficients for a three dimensional wing. This validates the effectiveness of NLLT approach in capturing the aerodynamic characteristics of a three dimensional wing and its usefulness for function evaluation in the present optimization process.

4.3 Optimal configurations

As remarked earlier, population size as well as accuracy and cost of function evaluation are critical to the success of a GA based optimization procedure. In the present study, varying population sizes are considered to bring out its influence on the optimal solution.



(a) $C_L - \alpha$ curve (b) $C_D - \alpha$ curve

Figure 7: Lift and drag curves of RAE–2822 wing

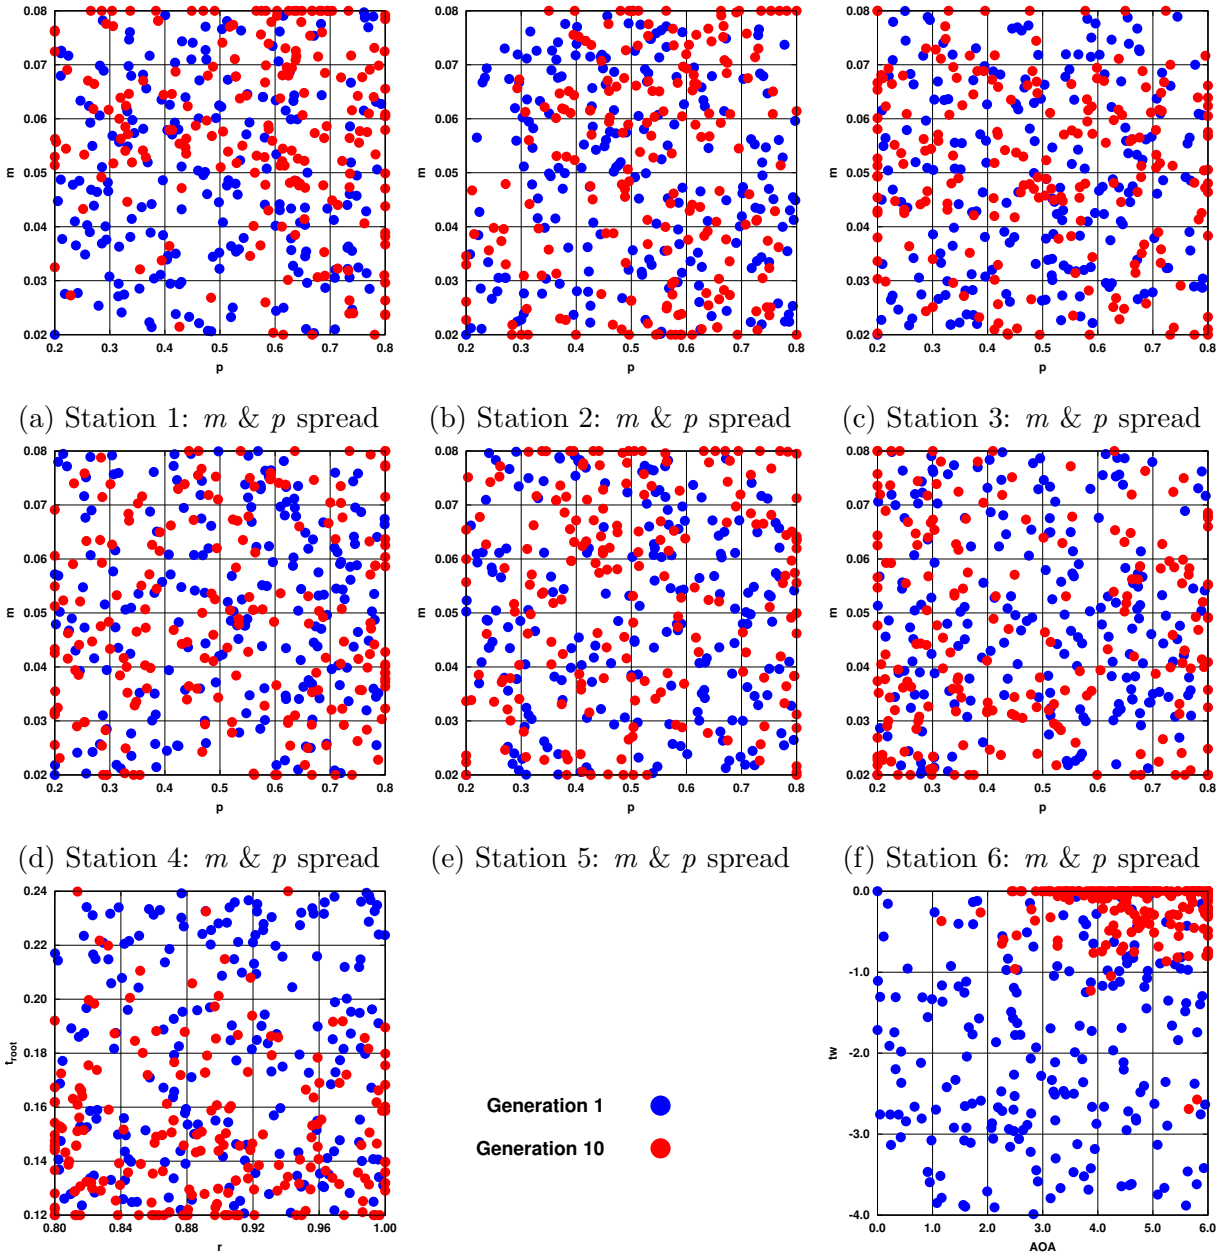
For this, following three different cases are considered:

1. A population size of 200 with 10 generations,
2. A population size of 200 with 100 generations and
3. A population size of 500 with 250 generations.

In a genetic algorithm, during initial generation, the population is randomly spread across the search space to ensure the exploration of a diverse set of potential solutions. Figures 8, 9 and 10 depict the population distribution in the first generation for each case, clearly demonstrating that the population is initially dispersed throughout the search space. Moving to subsequent generations, the genetic algorithm iteratively improves the population through selection, crossover and mutation, gradually converging toward the optimal solution. Analyzing the population evolution, following observations may be drawn:

1. For the first case corresponding to population size of 200 with 10 generations, Figure 8 illustrates that even after 10 generations, the population remains spread across the search space. This indicates that the population is still evolving and an aerodynamically optimal candidate has not yet emerged. Only the populations of tw and AOA show signs of clustering in the optimal region by the 10th generation.
2. For the second case corresponding to population size of 200 with 100 generations, Figure 9 illustrates that by 100th generation, most of the population has converged towards the optimal region, except for the populations of parameters m and p at stations 2, 3 and 4. These populations still exhibit a spread across the search space.
3. For the third case corresponding to population size of 500 with 250 generations, Figure 10 illustrates that the populations of all design parameters tend to cluster towards the optimal region of the search space, indicating a higher level of convergence.

To further evaluate the optimal configurations obtained for aforementioned cases, the distributions of maximum camber, position of maximum camber, thickness and load along



(a) Station 1: m & p spread (b) Station 2: m & p spread (c) Station 3: m & p spread

(d) Station 4: m & p spread (e) Station 5: m & p spread (f) Station 6: m & p spread

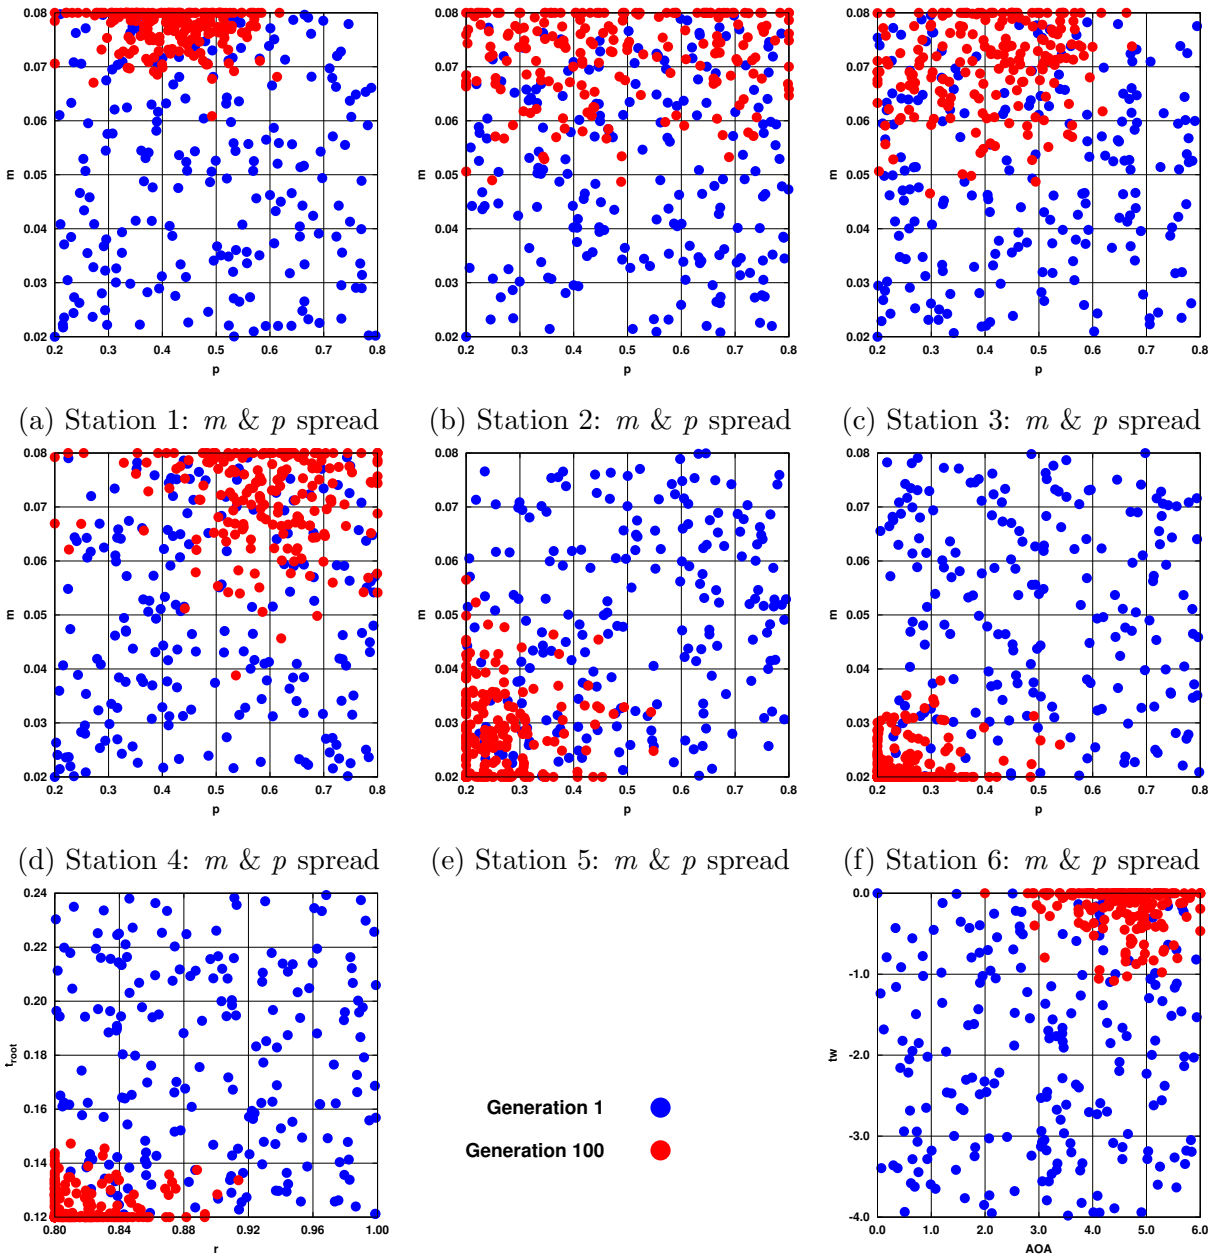
(g) t_{root} & r spread (h) tw & AOA spread

Figure 8: Population explored during the first and 10th generation for the population size of 200

the wing semi–span as depicted in Figures 11, 12, 13 and 14 respectively are examined. Comparing the optimal configurations for these cases, following observations may be drawn:

1. The parameters for optimal wing obtained for Case 1 are considerably different compared to other two cases.

2. Cases 2 and 3 exhibit similar optimal wing parameters with a slight variation in the position of the maximum camber.



(a) Station 1: m & p spread (b) Station 2: m & p spread (c) Station 3: m & p spread

(d) Station 4: m & p spread (e) Station 5: m & p spread (f) Station 6: m & p spread

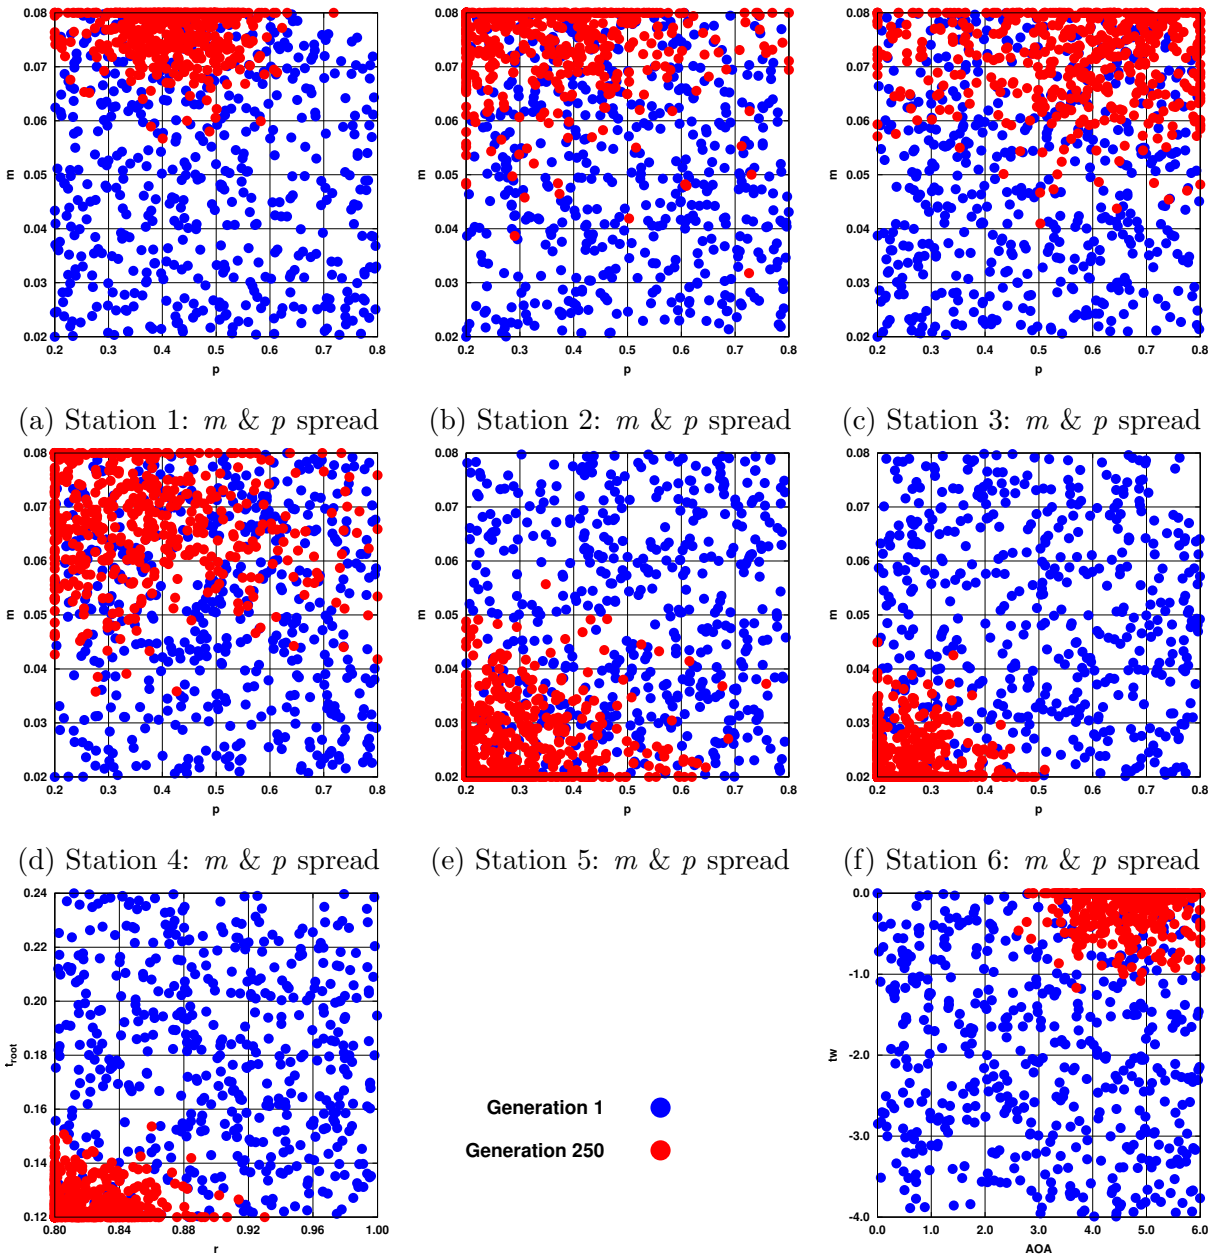
(g) t_{root} & r spread (h) tw & AOA spread

Figure 9: Population explored during the first and 100th generation for the population size of 200

3. The optimal wing for Case 3 exhibits smoother variation in position of maximum camber along semi–span compared to Case 2.

4. The load distribution on optimal wings for cases 2 and 3 is also very close to each other.

The optimal wing obtained for Case 1 is hereafter referred to as Wing–1. Similarly, the optimal wing obtained for Case 3 is hereafter referred to as Wing–2. These optimal wings are further compared in terms of their endurance parameter to establish the need for larger population in GA–based optimization.



(a) Station 1: m & p spread (b) Station 2: m & p spread (c) Station 3: m & p spread

(d) Station 4: m & p spread (e) Station 5: m & p spread (f) Station 6: m & p spread

(g) t_{root} & r spread (h) tw & AOA spread

Figure 10: Population explored during the first and 250th generation for the population size of 500

Table 2 presents the parameters obtained for Wing–1. Table 3 presents aerodynamic force coefficients and endurance parameter obtained at AOA of 3.585^o for Wing–1. Figure 15 presents the variation of endurance parameter with the angle of attack. From Table 3 and Figure 15, it is evident that the prediction of aerodynamic data by NLLT is very close to that of CFD, further demonstrating the accuracy of the NLLT procedure. It is observed that both NLLT and CFD give a maximum endurance parameter of approximately 35 at an AOA of around 5^o, which is higher than the AOA of 3.585^o obtained from the optimization process. This clearly indicates non convergence of the angle of attack (AOA) during optimization process. This observation highlights the effect

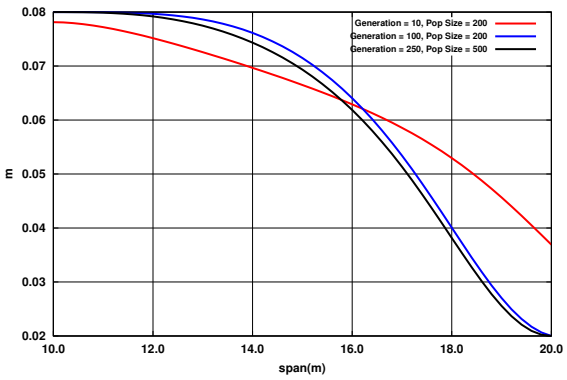


Figure 11: Maximum camber distribution

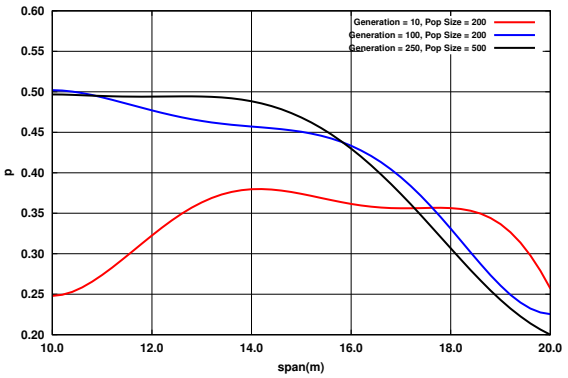


Figure 12: Position of maximum camber distribution

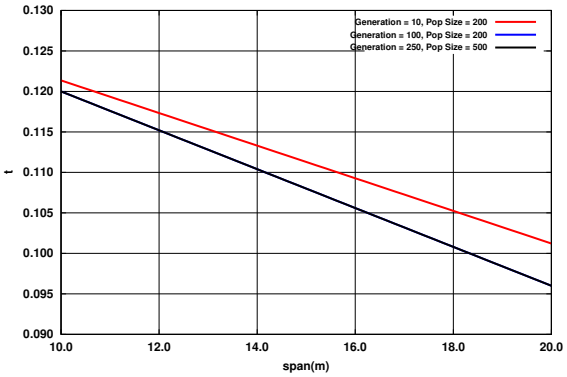


Figure 13: Thickness distribution

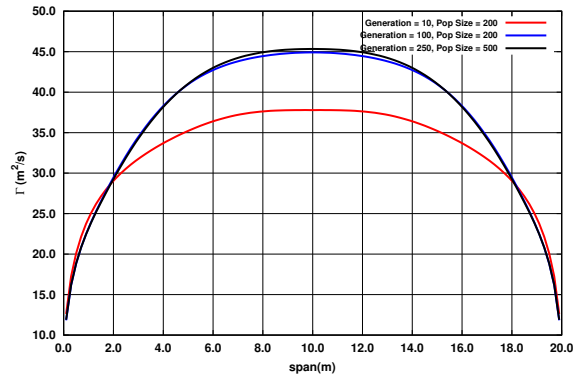


Figure 14: Gamma distribution along wing span

of number of generations and population size on the optimal solution. Increasing the number of generations provides GA with more opportunities to explore the search space, refine the population and ultimately achieve better solutions. Through this study, the importance of carefully selecting the population size and number of generations in a genetic algorithm to effectively converge toward optimal solutions in complex optimization problems is emphasized. This is further established when the results for the optimal wing configuration obtained from a larger population size and a larger number of generations, i.e. Wing–2, are discussed.

Table 2: Parameters of Wing–1

Sections	1	2	3	4	5	6
m	0.0781	0.0645	0.0734	0.0556	0.0622	0.0369
p	0.2479	0.6422	0.3688	0.2000	0.5832	0.2570

t_{root}	r	tw	AOA
0.1214	0.8340	0^o	3.585^o

Table 4 presents the parameters obtained for Wing–2. Table 5 presents aerodynamic force coefficients and endurance parameter obtained at AOA of 4.857^o for Wing–2. Figure 16 presents the variation of endurance parameter with the angle of attack. From

Table 3: Endurance parameter for Wing–1 using NLLT and CFD

	C_L	C_D	$C_L^{\frac{3}{2}}/C_D$
NLLT	0.968	0.0274	34.73
CFD	0.976	0.0278	34.68

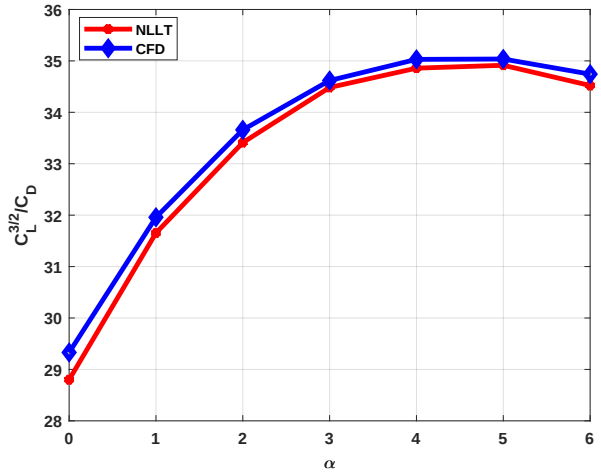


Figure 15: Endurance parameter variation for Wing-1 with angle of attack

Table 5 and Figure 16, it is clear that the prediction of aerodynamic data by NLLT is very close to that of CFD, thereby emphatically demonstrating the accuracy of the NLLT procedure. It is observed that both NLLT and CFD give a maximum endurance parameter of approximately 36 at the same AOA of 4.857^o and is slightly higher than the endurance parameter at AOA of 5^o. This clearly indicates the convergence of the angle of attack (AOA) during the optimization process.

At this stage, it should be remarked that the optimization process leading to Wing–2 involved about 125,000 function evaluations for 3D wings and it would have been prohibitively expensive to carry out such an optimization using a RANS solver. This justifies the use of NLLT which is computationally far less expensive compared to the RANS–based high–fidelity solver like HiFUN. In case of Wing–2, with the NLLT procedure using RANS based high fidelity sectional aerodynamic data employed in the function evaluation, the entire optimization process took approximately 10 hours on a modest desktop system having Intel® processor and 16 GB RAM.

Table 4: Parameters of Wing–2

Sections	1	2	3	4	5	6
m	0.0800	0.0800	0.0718	0.0706	0.0200	0.0200
p	0.4967	0.4476	0.5968	0.3741	0.2817	0.2000

t_{root}	r	tw	AOA
0.12	0.8	0^o	4.857^o

Finally, Figure 17 depicts the comparison of endurance parameter obtained using HiFUN for Wing–2 with a more common wing having Eppler aerofoil section E423 [14]. The aerofoil E423 has a maximum thickness of 12.5% at 23.7% of chord. This analysis involves

Table 5: Endurance parameter of Wing–2 obtained using NLLT and CFD

	C_L	C_D	$C_L^{\frac{3}{2}}/C_D$
NLLT	1.096	0.0319	35.90
CFD	1.110	0.0328	35.59

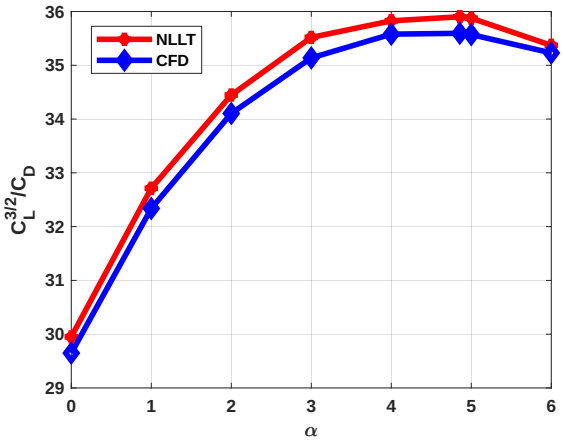


Figure 16: Variation of endurance parameter of Wing–2 with angle of attack

Eppler wing with identical thickness along the span and devoid of any twist. It is worth noting that the maximum endurance parameter of the Eppler wing is approximately 32, which is considerably lower than the value obtained for Wing–2. This outcome reinforces the effectiveness of the present methodology in designing wings capable of achieving long endurance.

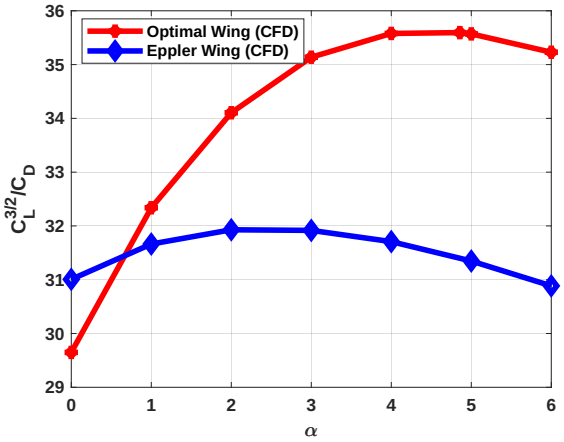


Figure 17: Comparison of endurance parameter of Wing-2 and Eppler wing

5 CONCLUSIONS

This paper presents a cost–effective and reliable methodology for aerodynamic optimization of a long endurance unmanned aerial vehicle (UAV) wing employing genetic algorithm (GA). The work establishes the need for considering a large population size and accurate function evaluation for a reliable GA based optimization process. The accuracy and cost effectiveness of the NLLT working in conjunction with a sectional aerodynamic

data set generated using high fidelity RANS solver is also established.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of SERC for granting them access to their advanced supercomputing facility to generate the reference data for this research project. The authors would like to express their sincere appreciation to Mr. A. Vadivelan, Scientist G at ADE, Bangalore, for his invaluable insights and suggestions regarding the optimization problem addressed in the study. Additionally, the authors extend their thanks to Dhanabal K., CSO at Sandi and Shashank Vissom, former Project Engineer at Sandi for their valuable contributions in generating the necessary grids for the study.

REFERENCES

[1] P. Mondal, A. V. Kanale, N. Balakrishnan, N. V. Shende, and R. P. Wadnerkar. Optimization of flap position using a modified discrete-vortex method. *AIAA Journal*, 55(12):4083–4093, 2017. doi:10.2514/1.J055070.

[2] L. Yongsheng, O. Akira, and M.-S. Liou. Progress in design optimization using evolutionary algorithms for aerodynamic problems. *Progress in Aerospace Sciences*, 46(5-6):199–223, 2010. doi:10.1016/j.paerosci.2009.08.003.

[3] B. M. Adams, W. J. Bohnhoff, K. R. Dalbey, M. S. Ebeida, J. P. Eddy, M. S. Eldred, J. R. Frye, G. Geraci, R. W. Hooper, P. D. Hough, K. T. Hu, J. D. Jakeman, M. Khalil, K. A. Maupin, J. A. Monschke, E. M. Ridgway, A. A. Rushdi, J. A. Stephens, L. P. Swiler, D. M. Vigil, and T. M. Wildey. Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.7 Users Manual. Sandia Technical Report SAND2014-4633, Sandia National Laboratories, 7 2014. updated November 2017.

[4] B. S. Arjun. Motion optimisation of plunging airfoil using swarm algorithm. Master's thesis, Indian Institute of Science, Bangalore, 2008.

[5] L. A. Piegl and W. Tiller. *The NURBS Book*. Springer, 1997.

[6] W. F. Phillips. *Mechanics of Flight*. Wiley, New York, 2004.

[7] W. F. Phillips and D. O. Snyder. Modern adaptation of prandtl's classic lifting line theory. *Journal of Aircraft*, 37(4):662–670, 2000. doi:10.2514/2.2649.

[8] `https://www.sandi.co.in/products`. Last accessed on 26 March 2025.

[9] P. L. Roe. Approximate riemann solvers, parameter vectors and difference schemes. *Journal of Computational Physics*, 43:357–372, 1981.

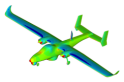
[10] N. Munikrishna. *On viscous flux discretisation procedures for finite volume and meshless solvers*. PhD thesis, Department of Aerospace Engineering, Indian Institute of Science, Bangalore, June 2007.

[11] V. Venkatakrishnan. Convergence to steady state solutions of the euler equations on unstructured grids with limiters. *Journal of Computational Physics*, 118:120–130, 1995.

[12] N. V. Shende and N. Balakrishnan. New migratory memory algorithm for implicit finite volume solvers. *AIAA Journal*, 42(9):1863–1870, 2004.

[13] P. R. Spalart and S. R. Allmaras. A one-equation turbulence model for aerodynamic flows. In *AIAA Paper 92-0439*, 1992.

[14] `http://airfoiltools.com/airfoil/details?airfoil=e423-il`. Last accessed on 26 March 2025.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

DEFORMATION STRATEGIES IN FUSELAGE AERODYNAMIC SHAPE OPTIMIZATION USING PAYLOAD VOLUME CONSTRAINTS

Luís D. Pinheiro[1,2*], Nuno M. B. Matos[1,2] and André C. Marta[2]

1: Research and Development
Tekever UAS
Rua das Minas, 2, 2500-750 Caldas da Rainha, Portugal
{luis.pinheiro, nuno.matos}@tekever.com, https://www.tekever.com

2: IDMEC
Instituto Superior Técnico
Universidade de Lisboa
Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal
{luis.pinheiro, andre.marta}@tecnico.ulisboa.pt, https://mdo.tecnico.ulisboa.pt

Abstract. *The usage of powerful optimization tools is becoming common in solving many engineering problems due to available computational resources and mature numerical algorithms. In this work, an adjoint-based high-fidelity aerodynamic shape optimization framework is used to manipulate a generic aircraft fuselage shape to minimize the total aerodynamic drag with specific payload volume constraints. The external fuselage shape is modified using the free-form deformation (FFD) technique to allow greater flexibility. The impact of using different deformation strategies applied to the displacement of FFD control points is studied, including displacements along the normal direction, the transverse axis directions, and cambering along the longitudinal axis direction. It is demonstrated that the combination of flexible control point displacement in multiple FFD box cross-section planes together with a streamwise cambering can produce a very significant drag reduction, in excess of 40% compared to the selected baseline shape, while still satisfying volume constraints to account for internal or protruding payloads.*

Keywords: aircraft design, gradient-based optimization, adjoint method, free form deformation, high-fidelity analysis, aerodynamic performance

1 INTRODUCTION

In recent years, the aviation industry has been expanding to more diverse solutions, proposing novel aircraft configurations and shapes. In addition, Unmanned Aerial Vehicles (UAVs) have been pivotal in multiple areas, such as agriculture, warfare and commodities built into our society [1–3]. Because of this growing market, aircraft shape optimization plays a crucial role in enhancing aerodynamic efficiency, payload distribution, structural integrity and overall flight performance. As this industry keeps evolving towards high-performance and efficient designs, the optimization of fuselages becomes critical.

Although traditional fuselage designs often present challenges related to drag reduction, weight minimization, and integration of different payloads, such as radars, gimbals or other components with random geometries, Computational Fluid Dynamics (CFD) and Multidisciplinary Design Optimization (MDO) [4] have enabled engineers to push the boundaries of fuselage designs, resulting in significant performance improvements [5–7]. Fuselage optimization is a critical area of modern aircraft design that has a direct influence on aerodynamic efficiency, structural stability, and operational flexibility [5, 8, 9]. Reducing parasitic drag [10], which accounts for a huge percentage of an aircraft's total aerodynamic drag, is one of the main targets in fuselage design. Through fuselage shape optimization, designers can have maximum control over airflow, which enables more efficient intersections with wings and tails, which in turn allows for innovative constructions with high aerodynamic performance [5]. Also, the ability to accommodate random and modular payload configurations efficiently is another critical factor in modern aircraft design, particularly in the UAV industry [11]. Unlike manned aircraft, UAVs do not require pressurized cabins, which offers greater design freedom in shape, weight distribution, and internal component layout. Therefore, fuselage shape optimization potentially offers greater performance gains while being flexible enough to accommodate a range of mission profiles.

In this work, gradient-based techniques will be used to guarantee rapid convergence to an optimized fuselage configuration, as demonstrated in similar cases [12, 13]. The gradients of the cost and constraint functions will be efficiently calculated using the adjoint method [14], making the computational cost independent of the number of design variables, which is expected to be large in shape optimization.

Various parametrization strategies will be explored to efficiently optimize the fuselage shape for best aerodynamic performance for a simple generic fuselage. For this, the free-form deformation (FFD) method, which is one kind of geometry modification technique that parametrizes shape perturbation instead of the shape itself used in aircraft design optimization, will be used [15]. The deformation process will include deformations along the fuselage axis, radial deformations, and deformations along the normals of the control points of the Free-Form Deformation (FFD) box. In addition, local shape modifications, such as cambering specific areas, mainly around the nose or tail area, will be analyzed.

In order to ensure practical feasibility, volume constraints will be included in the optimization framework. These constraints are important to preserve payload volume, as well as internal component allocation, and to avoid aggressive shape distortions that would compromise the aircraft's functional requirements [16, 17]. Through the imposition of volume preservation in key fuselage regions, the optimization will balance aerodynamic gains with real-world design constraints.

2 SHAPE DEFORMATION MODELING

Shape deformation modeling is an important technique in aerodynamic optimization to drive precise geometric deformation to achieve desired performance goals. Various techniques can be employed, including basis vector, domain element, discrete, analytical, Free-Form Deformation (FFD), partial differential equation (PDE), polynomial and spline, and Computer-Aided Design (CAD) [18–20]. FFD and CAD-based approaches are very efficient for complex problems [21]. Nevertheless, CAD systems still have problems parameterizing complex geometries and generating suitable grids for automatic simulations [21], making FFD a promising solution.

The FFD method has been successfully applied in aircraft design problems, such as helicopter fuselage optimization [22], manipulating the control points within a volume to generate smooth deformations. Other examples can be found in the works of Ekici and Juniper [23], Zhao et al. [24]. Moreover, this strategy also supports alterations in triangulated surface nodes and structure surface meshes along with a decrease in design variables [25]. Though FFD is not said to control any part of the design directly, it is still quite effective where solutions to aerodynamics entail huge displacements. The framework used in this study, Mach-Aero [26], incorporates the FFD method [27, 28].

2.1 Fuselage Deformation using FFD

The first study case is illustrated in Fig. 1, where the CFD mesh of the simplified fuselage used in the preliminary studies within the corresponding FFD box is shown. This arrangement sets the baseline geometry, allowing systematic modifications of the fuselage geometry by displacing control points within the FFD box, which subsequently rebuilds the 3D computational mesh.

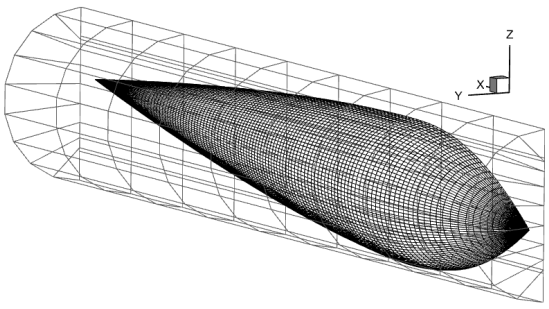
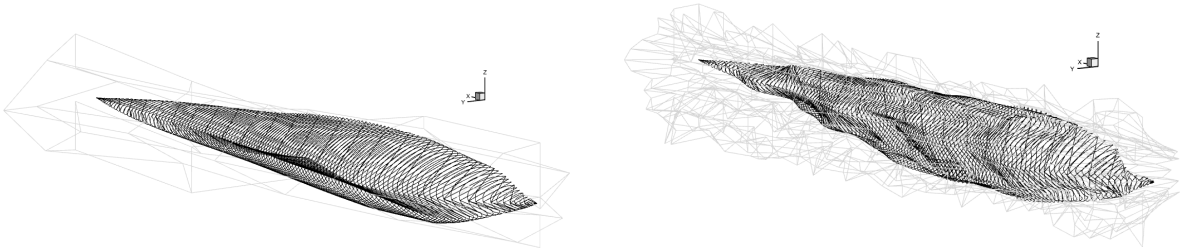


Figure 1: Simplified fuselage CFD mesh with the initial FFD box.

The FFD box resolution, in terms of number of control points, plays an important role in the flexibility of the deformations and must align with the optimization study's goals. A coarser grid with a few control points has a good capacity for preserving global deformations and, therefore, it is well suited to initial design and shape optimization at scales like the one shown in Fig. 2a. However, a finer grid with more control points, as illustrated in Fig. 2b, provides room for local control. The balance between shape degrees of freedom and computational cost has to be treated with great caution when introducing additional shape design variables [29].

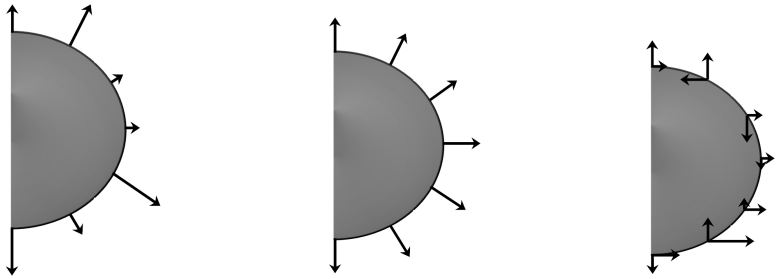


(a) Coarse FFD box. (b) Fine FFD box.

Figure 2: FFD box modeling.

2.2 Strategies for FFD Control Point Displacements

The parametrization of the FFD box and, hence, of the fuselage shape, can be done by several strategies with different advantages, depending on the optimization problem, as illustrated in Fig.4 for three different strategies.



(a) Normal direction (b) Radial direction (c) Y-, Z-direction

Figure 3: Different deformation strategies.

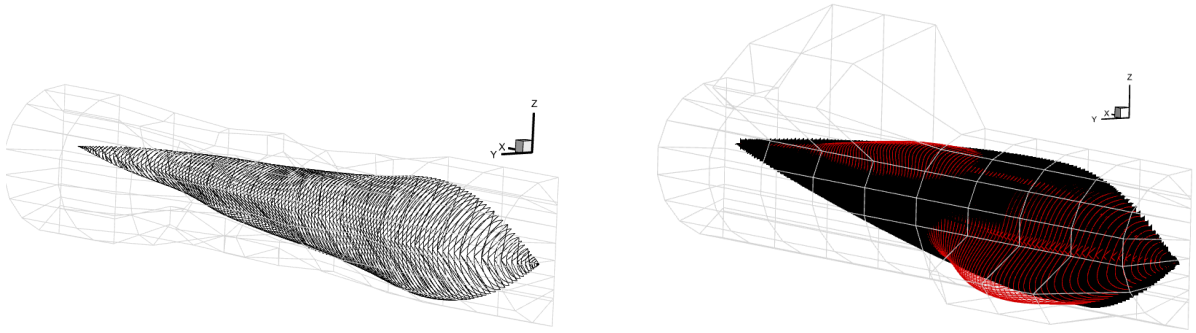
The first parametrization strategies (Fig.3a) uses the normal directions of the FFD box points, hence the deformations are defined in different directions based on the FFD box and fuselage geometry, offering flexibility to the optimizer. However, this strategies might produce less intuitive deformations for certain geometries, as the control points will move in different directions.

The second parametrization strategies (Fig.3b) uses radial deformations, making necessary the definition of a reference axis. By adjusting the fuselage's dimensions through scaling functions, this strategies makes it possible to expand or contract selected areas. However, for the present study, it was not as effective since the goal was to reshape the geometry in study and not reduce its size [28].

The last parametrization strategies (Fig.3c) performs deformations along the X, Y, and Z axes. It makes it convenient to use when scaling, stretching, or compressing along the axes, since it provides controlled shape variation in specific directions and provides the user with a clear understanding of the optimizer choices.

In addition to the three strategies described, it is also possible to define specific regions of the fuselage to be parametrized, spatially restricting the deformation region, with the use of a subset of control points defined by Point Selection. This gives the ability to deform different parts of the fuselage using distinct functions with different values, as observed in Fig. 4, which demonstrates how it is possible to reshape only certain parts of

the geometry in study using the normal vectors of the FFD box.

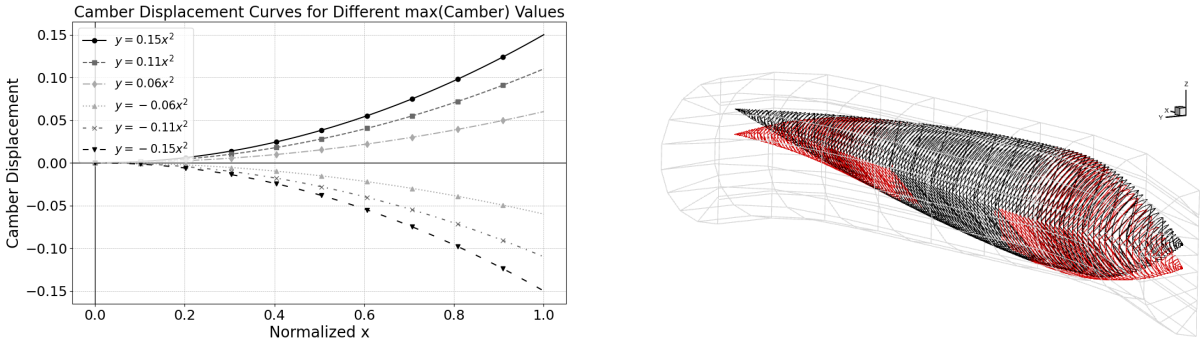


(a) Specific Volume (b) Specific Points

Figure 4: Specified deformations in the FFD box.

2.3 Cambered Fuselages

While the strategies presented in Sec.2.2 offer great freedom, the desire to employ a well-defined engineering parameter led to the definition of fuselage camber. Although fuselages are typically aerodynamically designed to have little contribution to total lift, that can be obtained by cambering the tail and nose sections [30, 31]. Such parametrization of a fuselage using camber was defined in the current framework using a quadratic function, as plotted in Fig. 5a, being the quadratic coefficient used later as a design variable during optimization. An example of such deformation can be observed in Fig. 5b, where both nose and trail camber have been added, where the fuselage length was normalized for the application of the quadratic function along its entire length.



(a) Quadratic camber functions. (b) Nose and tail fuselage deformation.

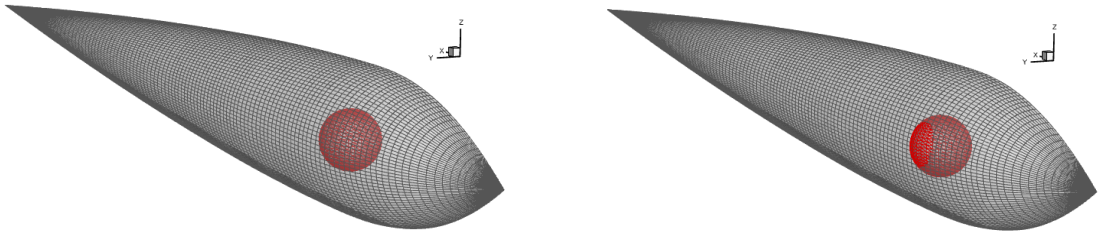
Figure 5: Camber parametrization.

This strategy can be further customized by selecting only certain FFD control points to restrict the deformation to only certain fuselage regions, namely by deforming the nose and tail with different independent curvatures. This selective approach can potentially enhance lift and stability adjustments while minimizing unwanted aerodynamic disturbances, particularly at intersection regions (with wing or tail).

2.4 Volume Constraints

To prevent that the fuselage shape deformations during the optimization still lead to internal fuselage volumes that account for specific payload, such as radar or fuel tank, volume and triangulated surface constraints must be implemented [16].

On one hand, the volume constraint enforces volumes into the fuselages that prevent excessive deformations that could comprise the necessary payload volume. On the other hand, the triangulated surface constraint uses an STL file format to define a given 3D shape that must be encapsulated by the fuselage shape, thus limiting the shape deformation during optimization. The latter constraint facilitates the integration of payloads into aerodynamic problems, giving more freedom in the optimization. Fig. 6 illustrates an example of a possible internal (a) or partially external (b) payload shape (sphere) defined using an STL file.



(a) Internal payload. (b) Partially external payload.

Figure 6: Fuselage volume constraint (in red) defined by triangulated surface.

3 FUSELAGE SHAPE OPTIMIZATION

The baseline fuselage shape shown in Fig.1 will be used as a starting point for the aerodynamic optimization problem aimed at minimizing the drag coefficient, posed in standard form as

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimize}} \quad & C_D(\boldsymbol{x}) \\ \text{subject to} \quad & \text{KS}_{\text{geom}} \leq 0 \end{aligned} \quad , \tag{1}$$

where $\boldsymbol{x}$ is the vector of all design variables (DVs) used in each case, which include the coordinates of the FFD control points and/or the amplitude of the camber function. The constraint function $\text{KS}_{\text{geom}}(\boldsymbol{x})$ represents an aggregated geometric constraint based on the Kreisselmeier–Steinhauser (KS) function, which provides a smooth and conservative approximation of the maximum deviation between triangulated surfaces [16]. All optimization cases were performed in the MACH-Aero framework [14], using ADFlow [32] to evaluate the aerodynamic performance. The gradient-based SLSQP algorithm was used with a convergence tolerance of 10^{-6} and with a maximum of 1000 iterations.

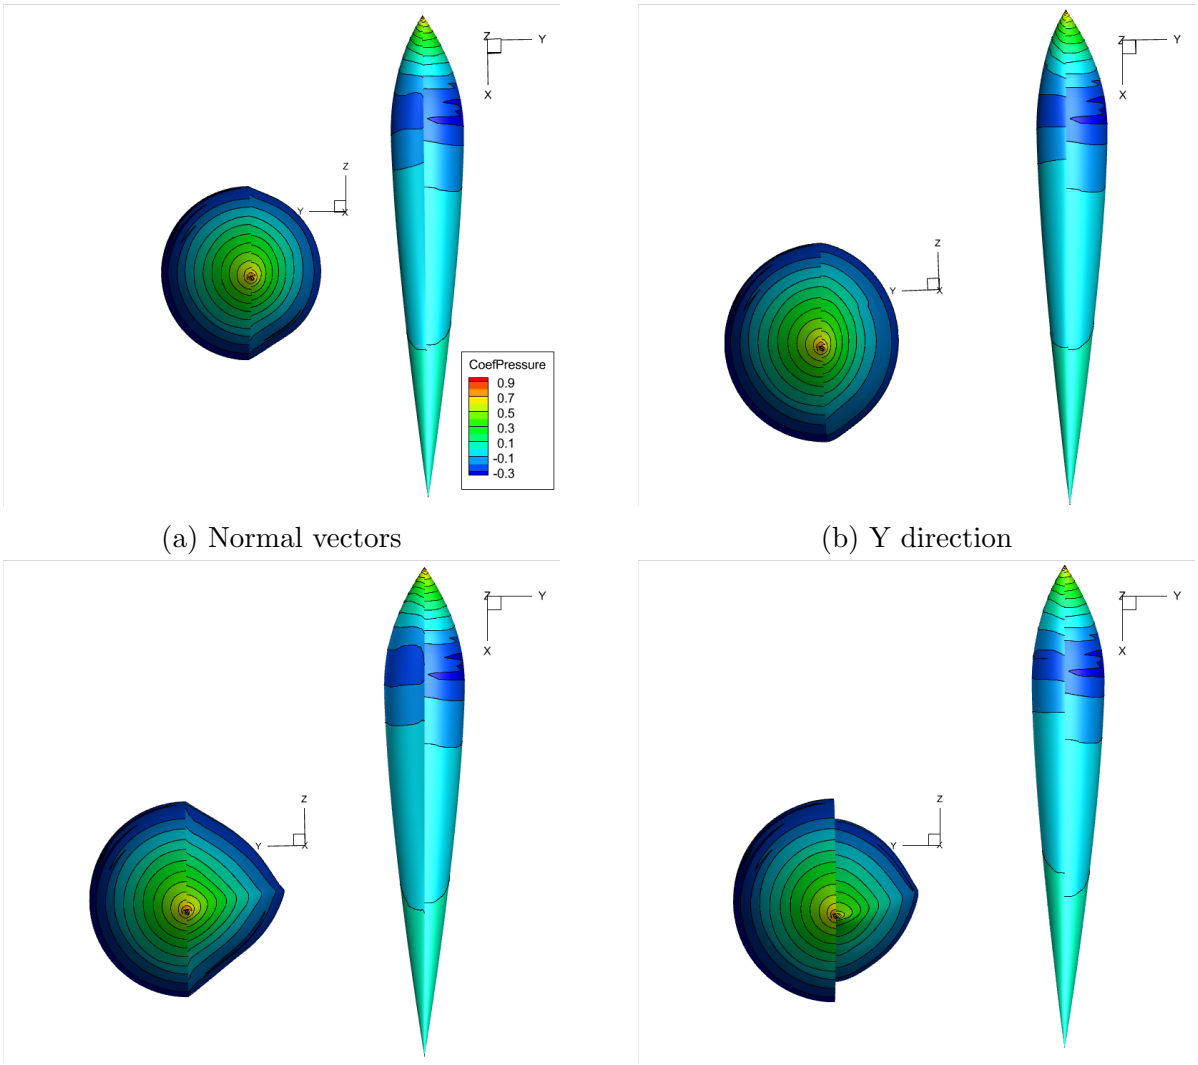
The fuselage operating condition represented a flight at a speed of 34 m/s at Standard Sea Level, with zero angle of attack (flow aligned with X-axis).

The different deformation strategies explained in Sec.2.2 are evaluated and theirs effectiveness compared. The strategies individual considered are normal-based deformations (Fig.7a), directional deformations along a specific axis Y (Fig.7b), Z (Fig.7c) or both (Fig.7d), and cambered fuselage deformations (Figs.9a, 9b and 9c). The combined camber and normal vectors (Fig.9d) or Y-/Z-directions (Fig.9e) are also studied. Lastly, the

effect of volume constraints in the optimal fuselage shape is also studied for both internal and protruding objects (Fig.10). The detailed discussion of all results is included in Sec.4.

3.1 Selected FFD Points

The first set of studies included the normal-based, and the (Y- and Z-) direction-based deformation parametrization strategies. Each strategy required specific considerations to ensure the stability of the optimization process, particularly to prevent the generation of invalid geometries, such as negative volumes, which would cause the mesh to fail and compromise the optimization. The resulting optimized fuselage shapes are illustrated in Fig. 7.



(a) Normal vectors

(b) Y direction

(c) Z direction

(d) Y and Z directions

Figure 7: Optimal shapes using different parametrization strategies.

For the normal-based deformations, Fig. 7a, it was necessary to carefully select which regions of the FFD box could be parametrized. Points located near the symmetry plane ($y = 0$) were excluded, as their displacements could easily produce negative CFD mesh volumes and, subsequently, make the optimization process to fail. Additionally, the points

near the edges of the fuselage, specifically at the nose and tail, were also frozen to prevent geometric inconsistencies at the edges, avoiding similar problems. Furthermore, the freedom given to the FFD points was kept relatively small to avoid, again, the risk of generating invalid mesh volumes. These limitations directly impacted the optimization results, as seen in the corresponding figures and data summarized in Tab.1.

In contrast, the directional deformation strategy, which deforms the geometry along the Y- and Z-axis, was generally easier to control. As expected, the freedom given to the FFD points could be increased without immediately risking mesh failures in these cases. However, this strategy tended to excessively deform certain regions of the fuselage, particularly near the tail, creating irregular geometries, almost like a twist or warp in these regions, as observed in Fig.8. Similar to the normal-based strategy, the points at

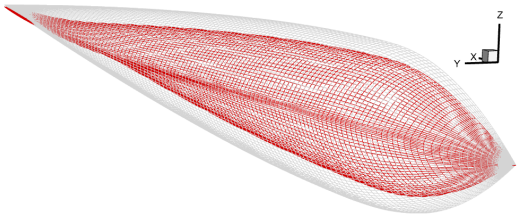


Figure 8: Irregularities encountered using the directional deformation method

the fuselage leading and trailing edges were kept fixed to avoid negative mesh volumes, which limited the flexibility of the surrounding regions, thus aggravating the localized deformations observed.

3.2 Camber Function

The second set of studies was conducted to identify the benefits of adding camber to the fuselage. The precautions when setting up these cases are primarily related to avoid creating mismatches in the intersection regions with components such as wings or tails. To address this issue, freedom was given only at specific longitudinal (X-axis) intervals that excluded intersection regions to prevent the aforementioned problems. Three different cases were tested, applying the camber function only in the nose region, tail region, or both, which produced the optimal shapes shown in Fig. 9,

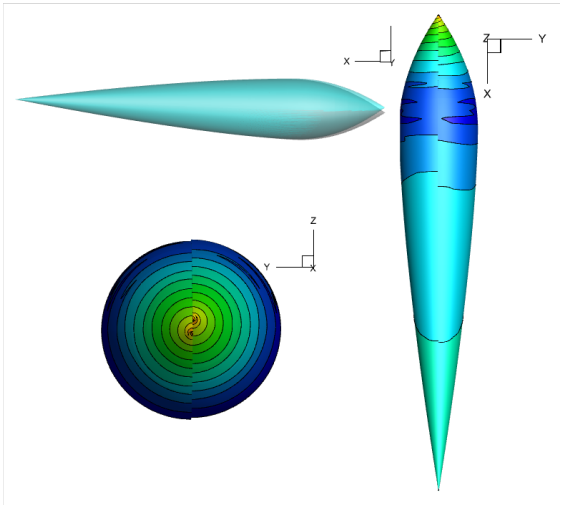
3.3 Combined Deformation Strategies

The last two studies include the combination of the camber function with the normal- or direction-based strategy. The resulting fuselage shapes are shown in Figs.9d and 9e, respectively.

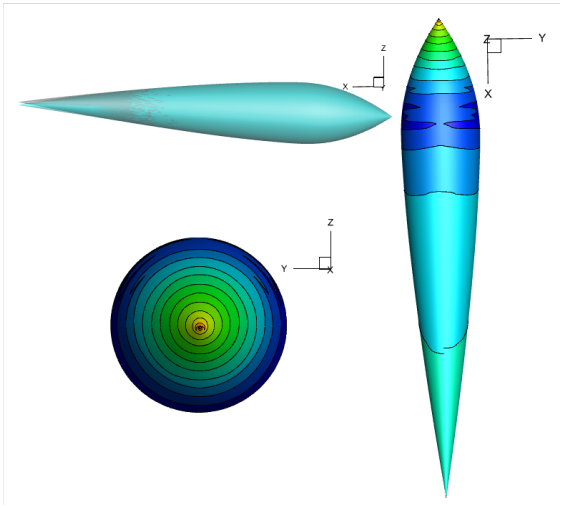
3.4 Volume Constraints

To complete the study, the Triangulated Surface Constraint was tested with the same different shape deformation strategies. The same 3D object, a sphere, was used in these cases to represent any given payload. Two different object locations were experimented: the object inside the baseline fuselage and partially outside, protruding the baseline fuselage shape.

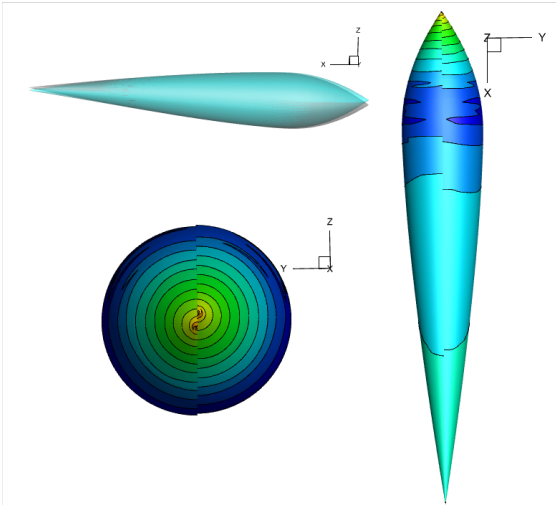
The cross-section of the resulting optimal fuselage shapes obtained with the different deformation strategies, passing through the object, are illustrated in Fig. 10. The so-



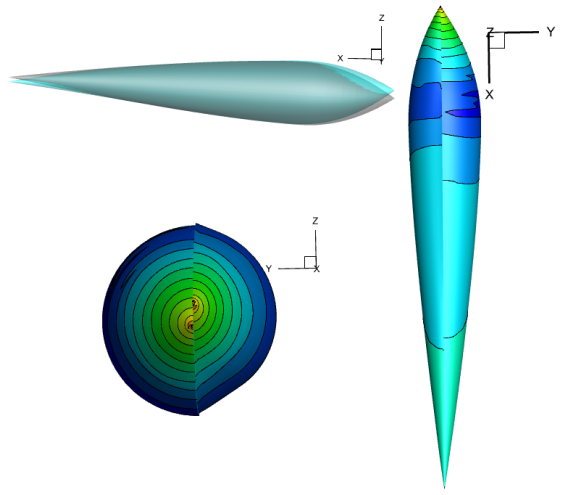
(a) Restricted to nose.



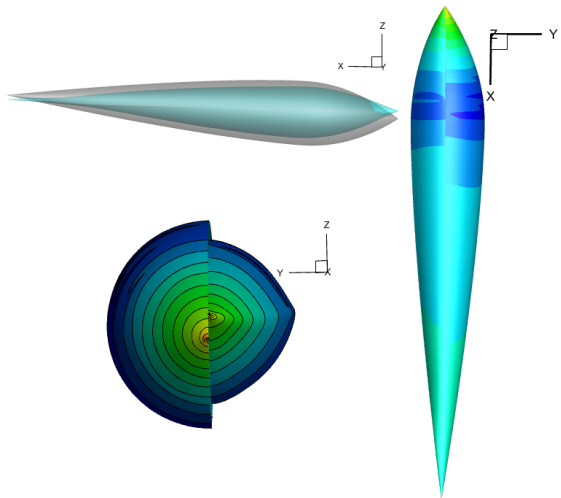
(b) Restricted to tail.



(c) Restricted to nose and tail.



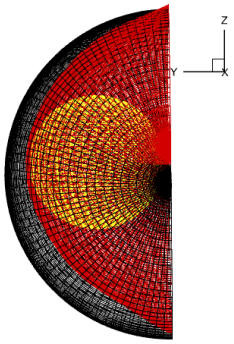
(d) Camber combined with normal vectors.



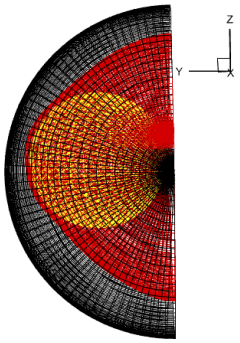
(e) Camber combined with Y and Z directions.

Figure 9: Optimal shapes with camber deformation strategy.

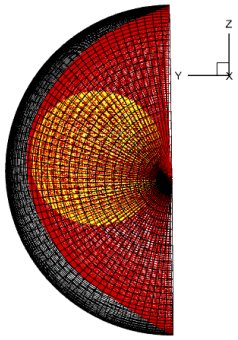
lutions to the problems with the object inside the baseline fuselage include: combined camber and normal-based deformations (Fig.10a), combined camber and deformations along the Y- and Z-axes (Fig.10b), and deformations only along normal-vectors of the FFD box (Fig.10c). In addition, the cases with the object initially protruding the baseline fuselage are shown in Figs.10d and 10e, using normal-vectors, and combined with z-axis deformation, respectively.



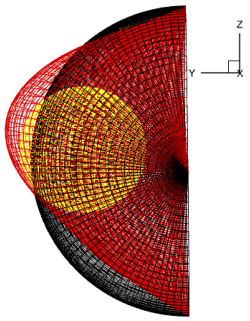
(a) Camber + normal vectors deformations with sphere inside fuselage.



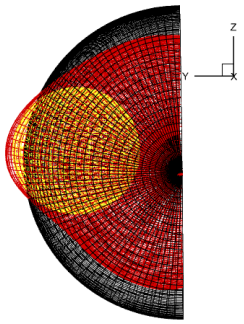
(b) Camber + (Y and Z) deformations with sphere inside fuselage.



(c) Parametrization using normals of FFD points with sphere inside fuselage.



(d) Parametrization using normals of FFD points.



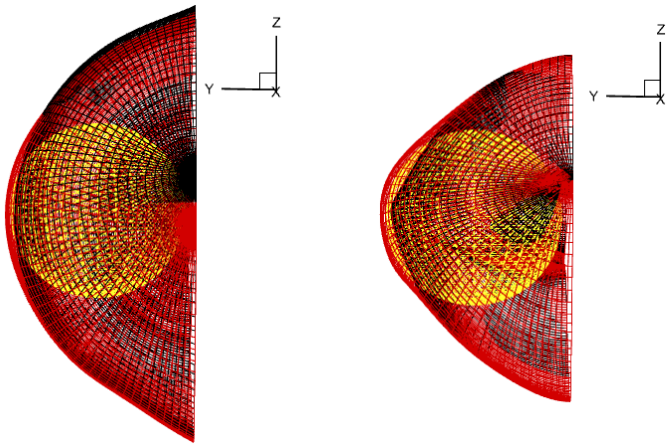
(e) Parametrization along Z-axis, combined with normal vectors of FFD points.

Figure 10: Optimal shapes using the triangulated surface constraint.

As attested in Fig. 10, each of the presented cases made the optimizer manipulate the FFD control points in different ways to ensure that the fuselage could envelop the payload (sphere).

The main differences encountered between the cases shown in Figs.10d and 10e arise from the variation in the selected points. While deforming the shape using the normal vectors of the FFD box points, it was necessary to exclude points near the plane of symmetry ($y = 0$) to prevent mesh issues and negative volumes, which caused the differences. To mitigate this problem, it was given freedom to these points along the Z axis in the second case, which visibly improved the design, as the fuselage was able to shrink near the symmetry plane, resulting in a smoother and more effective deformation.

Finally, the differences between parameterizations with and without the triangulated surface constraint can be observed in Fig.11, where two different cases are illustrated: combined camber and normal-based deformations (Fig.11a) and camber with deformations along the Y and Z axes (Fig.11b), both showing optimized fuselages taking (red) or not (black) into consideration the triangulated surface constraint. It is clear that, despite using the same parameter values, the resulting shapes differ significantly. It is also worth noting that, for the optimized fuselages (black), the optimal design variable values reached their imposed lower bounds, which suggests that the optimizer could have further reduced the fuselage size in the absence of these.



(a) Camber + normal vectors deformations.

(b) Camber + (Y and Z) deformations.

Figure 11: Comparison of optimal shapes with (red) or without (black) triangulated surface constraint.

4 DISCUSSION OF RESULTS

The drag reduction obtained for each of the different deformation strategies studied in Sec.3 is listed in Tab.1.

Through the analysis of optimal shapes presented in Sec.3 and the aerodynamic performance gains presented in Tab.1, it is possible to assess about the effectiveness of the different shape parametrization strategies proposed. Notice, however, that these optimal shapes were obtained from problems that differed slightly in terms of design variable bounds to overcome mesh morphing fails, as such, the comparisons might not extrapolate directly to other fuselage geometries.

In the cases with the deformation parametrization along the normal vectors of the FFD box points, the optimized shapes exhibited fewer irregularities and a more stable CFD mesh morphing, leading to a nearly 17% drag reduction for either the unconstrained and internal volume constrained cases, but less an expressive 7% reduction in the protruding payload case. However, this strategy revealed to be relatively limited and not well adapted to arbitrary fuselage shapes.

Looking at the strategies using deformations along the Y- and Z-directions, they exhibited a much greater drag reduction compared to the normal direction strategy. The Y-direction only strategy, corresponding the lateral fuselage deformations led to consid-

Table 1: Optimization deformation strategy studies and corresponding drag coefficient reduction.

Parametrization strategy	**Constraints**	**Opt. shape**	**Min C_D**	ΔC_D
(baseline)			0.0189	ref
Normals	None	Fig.7a	0.0161	-16.6%
Y-axis	None	Fig.7b	0.0155	-21.6%
Z-axis	None	Fig.7c	0.0168	-11.7%
(Y, Z)-axis	None	Fig.7d	0.0114	-40.4%
Nose Camber	None	Fig.9a	0.0185	-2.1%
Tail Camber	None	Fig.9b	0.0185	-2.1%
(Tail, Nose) Camber	None	Fig.9c	0.0184	-2.3%
Camber + Normals	None	Fig.9d	0.0161	-17.1%
Camber + (Y, Z)-axis	None	Fig.9e	0.0133	-41.6%
Camber + Normals	TS (IS)	Fig.10a	0.0162	-16.9%
Camber + (Y, Z)-axis	TS (IS)	Fig.10b	0.0135	-39.8%
Normals	TS (IS)	Fig.10c	0.0162	-16.4%
Normals	TS (IntS)	Fig.10d	0.0176	-7.0%
Normals + Z-axis	TS (IntS)	Fig.10e	0.0155	-21.3%

TS = Triangulated Surface, IS = Internal Sphere, IntS = Intersected Sphere.

erably better shapes compared to the Z-direction only strategy, and, unsurprisingly given the greater deformation freedom, the combination of both Y- and Z-direction deformations led to an impressive 40% drag reduction.

When considering the deformation using the camber function alone, little impact in the aerodynamic performance was obtained, limited to about 2% reduction in drag, regardless of the regions manipulated (nose, tail or both). However, tuning the fuselage camber has been shown to improve the airflow going through the fuselage to the wing [30, 31], as should not be discarded in more complex cases.

The combination of the camber deformation with the simultaneous Y- and Z-direction deformation resulted in the best overall unconstrained fuselage shape, which exhibited a massive 41.6% drag reduction compared to the baseline. However, some cautions regarding the optimized shape must be taken into consideration since some irregularities were created during the optimization. This combined strategy also produced the best aerodynamic shape for the constrained case with internal payload, with almost 40% drag reduction.

In future studies, to mitigate the referred problems seen in Fig.8, it might be beneficial to make some modifications near the tail region of the fuselage, such as adding a volume constraint might prevent the formation of torsion or warping. Another way to prevent these irregularities could be to incorporate a parametrization strategy, such as those discussed in this study, which, when combined with the Point Selection method, might allow for greater flexibility and deformation control in the affected region.

Referring the the protruding payload cases, only the deformation strategy along the normal vectors was considered to prevent irregular deformations in the fuselage while integrating the triangulated surface constraint with the sphere intersected in the main body. However, in this study, and when comparing the cases represented by Figs.10d and

10e, the normal direction deformation only case led to a 7% drag reduction while the combination with the Z-direction deformation increase the reduction significantly to over 21%, mostly due to the impact of the manipulation of the FFD points near the fuselage vertical symmetry XZ-plane.

5 CONCLUSIONS

Several parametrization strategies for fuselage deformation in aerodynamic shape optimization were proposed, tested and compared. This allowed an understanding of the characteristics of each strategy and how they can be applied in more complex design problems.

Drag reduction ranging from 2.1% to 41.6% were observed, depending on the case study considered. Despite the directional parametrization leading to greater drag reduction, the parametrization along the normal vectors of the FFD box points should not be excluded when opting for a deformation strategy. This strategy combined with a deformation along the vertical Z-axis have produced good results as well, with the added benefit of reducing the surface irregularities if the DV bounds are properly set.

This work highlighted the importance of selecting appropriate parametrization strategies based on both aerodynamic performance and payload constraints (internal to the fuselage or partially protruding). The triangulated surface constraint proved to be a powerful tool in constraining an aerodynamic problem, offering multiple applications to handle the payload positioning and volume required.

While this study focused on specific implementations, future research should explore the impact of varying constraint parameters, such as spatial tolerance [16].

ACKNOWLEDGEMENTS

The authors would like to thank Tekever UAS for supporting the production of this article by offering all the needed data to successfully produce the analyses here shown.

The third acknowledges Fundação para a Ciência e a Tecnologia (FCT) for its financial support via the project LAETA Base Funding (DOI: 10.54499/UIDB/50022/2020).

REFERENCES

[1] K. K. Reddy, S. K. Basha, B. Naveen, and K. Abhilash. The game-changing role of drones in agriculture. *Agri Express*, 02, May 2024. ISSN 2584 - 2498.

[2] Ö. Ö. Kanat. The significance of unmanned aerial vehicles (uavs) in strategic contexts. *Anadolu Strateji Dergisi*, 5(2):75–87, 2023.

[3] H. Shakhatreh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, and M. Guizani. Unmanned aerial vehicles (uavs): A survey on civil applications and key research challenges. *IEEE Access*, February 2019.

[4] J. R. R. A. Martins and A. B. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075, 2013. doi:10.2514/1.J051895.

[5] S. Rohani and A. Jahangirian. Multi-objective aerodynamic optimization of a civil aircraft fuselage. *ICAS, Shanghai, China*, pages 1–6, 2021.

[6] A. Dhara, A. Majumder, N. E. Kumar, M. Dhanunjay, L. T. Dhanumjaya, and J. Muruga Lal Jeyan. Design optimization to minimize wake of wide-body transport aircraft. *Journal of Engineering Research*, 12(2):200–209, 2024. ISSN 2307-1877. doi:https://doi.org/10.1016/j.jer.2023.08.014.

[7] J. R. Martins. Aerodynamic design optimization: Challenges and perspectives. *Computers & Fluids*, 239:105391, 2022. ISSN 0045-7930. doi:10.1016/j.compfluid.2022.105391.

[8] J. Welstead, B. Reitz, and G. Crouse. Modeling fuselage aerodynamic effects in aircraft design optimization. In *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 01 2012. ISBN 978-1-60086-936-5. doi:10.2514/6.2012-394.

[9] J. Mieloszyk and T. Goetzendorf-Grabowski. Introduction of full flight dynamic stability constraints in aircraft multidisciplinary optimization. *Aerospace Science and Technology*, 68:252–260, 2017. ISSN 1270-9638. doi:10.1016/j.ast.2017.05.024.

[10] D. P. Raymer. *Aircraft Design: A Conceptual Approach*. AIAA Education Series. American Institute of Aeronautics and Astronautics, Inc., 2 edition, 2012. ISBN 0-930403-51-7.

[11] S. A. H. Mohsan, N. Q. H. Othman, L. Yanlong, M. Alsharif, and M. Khan. Unmanned aerial vehicles (uavs): Practical aspects, applications, open challenges, security issues, and future trends. *Intelligent Service Robotics*, 01 2023. doi:10.1007/s11370-022-00452-4.

[12] S. Roy and W. Crossley. Hybrid multi-objective combinatorial optimization technique with improved compatibility between ga and gradient-based local search. In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, September 2012. doi:10.2514/6.2012-5568.

[13] S. Khosravi and D. W. Zingg. A numerical optimization study on winglets. In *15th AIAA/ISSMO multidisciplinary analysis and optimization conference*, 2014. doi:10.2514/6.2014-2173.

[14] G. K. W. Kenway, C. A. Mader, P. He, and J. R. R. A. Martins. Effective adjoint approaches for computational fluid dynamics. *Progress in Aerospace Sciences*, 110: 100542, October 2019. doi:10.1016/j.paerosci.2019.05.002.

[15] J. Reid and P. He. *Free-Form Deformation (FFD)*. DAFoam, February 2021. Accessed on January 25, 2025.

[16] B. J. Brelje, J. L. Anibal, A. Yildirim, C. A. Mader, and J. R. R. A. Martins. Flexible formulation of spatial integration constraints in aerodynamic shape optimization. *AIAA Journal*, 58(6):2571–2580, 2020. doi:10.2514/1.J058366.

[17] G. Burgreen and O. Baysal. Three-dimensional aerodynamic shape optimization of supersonic delta wings. In *5th Symposium on Multidisciplinary Analysis and Optimization*, September 1994. doi:10.2514/6.1994-4271.

[18] J. A. Samareh. Survey of shape parameterization techniques for high-fidelity multidisciplinary shape optimization. *AIAA Journal*, 39(5):877–884, 2001. doi:10.2514/2.1391.

[19] A. Hahn. Vehicle sketch pad: A parametric geometry modeler for conceptual aircraft design. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2010. doi:10.2514/6.2010-657.

[20] D. Rodriguez and P. Sturdza. *A Rapid Geometry Engine for Preliminary Aircraft Design*. American Institute of Aeronautics and Astronautics, January 2006. doi:10.2514/6.2006-929.

[21] J. A. Samareh. Aerodynamic shape optimization based on free-form deformation. *Collection of Technical Papers - 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 6:3672–3683, 2004. doi:10.2514/6.2004-4630.

[22] M. Wentrup. An adjoint based optimization chain for complex helicopter fuselage parts using a free form deformation or cad based parameterization method. In *41st European Rotorcraft Forum*, September 2015.

[23] E. Ekici and M. P. Juniper. Adjoint based shape optimization for thermoacoustic stability of combustors using free form deformation. In *Turbo Expo: Power for Land, Sea, and Air*, volume 87943, page V03AT04A058. American Society of Mechanical Engineers, 2024.

[24] H. Zhao, J. T. Hwang, and J.-S. Chen. Open-source shape optimization for isogeometric shells using fenics and openmdao, 2024. URL `https://arxiv.org/abs/2410.02225`.

[25] N. R. Secco, J. P. Jasa, G. K. W. Kenway, and J. R. R. A. Martins. Component-based geometry manipulation for aerodynamic shape optimization with overset meshes. *AIAA Journal*, 56(9):3667–3679, 2018. doi:10.2514/1.J056550.

[26] MDO Lab. MDO Lab Mach-Aero Documentation. URL `https://mdolab-mach-aero.readthedocs-hosted.com/en/latest/`. Accessed: March 10, 2025.

[27] G. K. W. Kenway, G. J. Kennedy, and J. R. R. A. Martins. A CAD-free approach to high-fidelity aerostructural optimization. In *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, September 2010. AIAA 2010-9231.

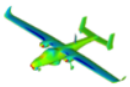
[28] H. M. Hajdik, A. Yildirim, E. Wu, B. J. Brelje, S. Seraj, M. Mangano, J. L. Anibal, E. Jonsson, E. J. Adler, C. A. Mader, G. K. w. Kenway, and J. R. r. a. Martins. pygeo: A geometry package for multidisciplinary design optimization. *Journal of Open Source Software*, 8(87):5319, 2023. doi:10.21105/joss.05319.

[29] L. A. Schwarz. Non-rigid registration using free-form deformations. Master's thesis, Technical University of Munich (TUM), May 2007.

[30] M. Z. bin Md Shah, M. R. bin Abu Bakar, and B. Basuno. The aerodynamics analysis on cambered fuselage model. In *Advances in Mechanical, Materials and Manufacturing Engineering*, volume 660 of *Applied Mechanics and Materials*, pages 492–497, 11 2014. doi:10.4028/www.scientific.net/AMM.660.492.

[31] K. D. Harris. An experimental investigation of the subsonic drag and pitching moment characteristics of slender cambered bodies with pointed noses and tails. Technical report, The College of Aeronautics, Cranfield University, September 1958.

[32] C. A. Mader, G. K. W. Kenway, A. Yildirim, and J. R. R. A. Martins. ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization. *Journal of Aerospace Information Systems*, 17(9):508–527, Sept. 2020. doi:10.2514/1.I010796.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

Optimizing UAV Winglets: A Multiobjective Approach with Surrogate Models

M. Almutairi[1*], P. Dunning[1], and A. Maheri[1]

1: School of Engineering, Fraser Noble Building

University of Aberdeen, Aberdeen, United Kingdom

m.almutairi.21@abdn.ac.uk, https://www.abdn.ac.uk/people/m.almutairi.21
peter.dunning@abdn.ac.uk, alireza.maheri@abdn.ac.uk

Abstract. *This study presents a multi-objective optimization approach to enhance the aerodynamic performance of unmanned aerial vehicle winglets using high-fidelity computational analysis and surrogate modeling. Kriging and polynomial surrogate models, trained on high-fidelity computational fluid dynamics (CFD) simulations, are integrated with a multi-objective genetic algorithm (NSGA-II) to explore the design space and identify optimal trade-offs. The optimization focuses on two multi-objective problems: maximizing the lift-to-drag ratio while minimizing the root bending moment and maximizing lift whilst minimizing drag. A comparative analysis of the surrogate models reveals that Kriging excels with smaller datasets, particularly in RBM prediction, while polynomial regression improves with larger datasets, demonstrating superior drag prediction accuracy. Results indicate that using 50 samples provides the optimal balance between prediction accuracy and computational efficiency. The findings highlight the effectiveness of Kriging for predicting structural performance and polynomial regression for aerodynamic performance. These insights contribute to UAV and aerospace applications, emphasizing the importance of integrating both aerodynamic and structural objectives for effective winglet design.*

Keywords: UAV, Winglet, CFD, Surrogate Modelling, Multiobjective Optimization

1 INTRODUCTION

Unmanned Aerial Vehicles (UAVs) play a vital role in surveillance, monitoring, and search-and-rescue missions. Maximizing their endurance and efficiency requires optimizing aerodynamic performance, which directly influences fuel consumption and flight range. Winglets, a key design feature, enhance the lift-to-drag (L/D) ratio and reduce induced drag, making them critical for UAV performance [1-8]. However, to achieve optimal benefits, winglets must be carefully tailored to specific UAV designs and flight conditions, balancing aerodynamic and structural considerations [9]. While wind tunnel testing is traditionally used for aerodynamic assessments, computational approaches offer a cost-effective and efficient alternative, particularly Reynolds-Averaged Navier-Stokes (RANS) CFD simulations, which can capture complex flow phenomena, making it widely used for winglet design [10].

Studies indicate that winglet parameters such as cant angle, sweep angle, and airfoil shape strongly influence aerodynamic performance. Liu et al. [11] demonstrated that optimized winglets reduce drag by 6.5% and increase range by 7.5%. Beechook & Wang [12] validated RANS CFD simulations for winglet analysis, while Elham & van Tooren [13] employed multi-objective optimization, achieving 3.8% fuel savings and 1.7% lower operational costs for aircraft like the Boeing 747.

Winglet optimization has also benefited wind turbine blades. Lawton & Crawford [14] demonstrated a 2% increase in power extraction using CFD-based winglet optimization. Ribeiro et al. [15] and Elfarra et al. [16] further enhanced wind turbine winglets using genetic algorithms and surrogate models, achieving a 9% increase in power production while reducing computational time.

Surrogate models significantly accelerate optimization by efficiently exploring design variations without the computational burden of extensive full-scale CFD simulations [17].

This paper builds upon prior research by comparing surrogate models and analyzing the impact of key winglet parameters on UAV wing performance (lift, drag, L/D ratio, and root bending moment). Using RANS CFD, surrogate modeling, and multiobjective optimization, we aim to enhance UAV aerodynamic efficiency. The methodology is detailed in Section 2, followed by results and discussion for a test case in Section 3, and conclusions in Section 4.

2 METHODOLOGIE

Winglet optimization requires balancing aerodynamic and structural efficiency. In this paper, the initial focus is on aerodynamic performance, where the two objectives are to maximize lift, whilst minimizing drag. This allows us to observe how winglet design parameters effect aerodynamic performance. Next, structural considerations are integrated, by maximizing the L/D ratio, whilst minimizing the root bending moment (RBM) to reduce wing loading, leading to a reduction in wing mass.

To solve the multiobjective optimization problems, this study employs the NSGA-II genetic algorithm (GA) [18], which identifies Pareto-optimal solutions that optimize both objectives. This is supported by surrogate models that are trained on samples generated from hi-fidelity 3D RANS CFD simulations.

In the following section, details of the test case, automated geometric modeling, optimization formulation, CFD analysis and surrogate modeling are explained.

2.1 Wing-Winglet Geometry Definition

To enable automatic winglet design optimization, a flexible and robust geometry modeler is developed. The modeler automatic creates the geometry of a single wing with winglet from defining aerodynamic parameters such as span, twist, chord length, sweep, and tip-chord ratio, for both the wing and winglet. It can also use different airfoils along the wing and winglet and smoothly blends the wing into the winglet.

The test case used in this paper is taken from [4]. The winglet, positioned at the tip of the wing, consists of a single aerofoil, PSU94-097, and the baseline design is taken from [4] and summarized in Table 1, along with the upper and lower bounds for optimization. The main wing spans 3.15 m and utilizes the NASA NLF(1)-1015 airfoil, with a root chord of 1.1 m, a tip chord of 0.5 m, a sweep angle of 1.5°, and a twist angle of 0°.

Table 1. Design variables for winglet design.

Parameter	*Baseline Configuration*	*Lower Bound*	*Upper Bound*
Span (m)	*0.345*	*0.335*	*0.375*
Tip Chord (m)	*0.2*	*0.15*	*0.25*
Sweep Angle (°)	*41*	*0*	*50*
Cant Angle (°)	*60*	*40*	*70*
Twist Angle (°)	*0*	*-5*	*5*
Toe Angle (°)	*0*	*-1*	*1*

2.2 Aerodynamic Modelling and CFD Simulation

Once the geometry is automatically generated by a Matlab script, meshing for CFD analysis is done using ICEM meshing software. Localized refinements are used in critical areas, such as the winglet and wing junction, while inflation layers are applied at the airfoil surface to enhance boundary layer resolution. To balance computational efficiency with accuracy, a coarse mesh is applied to far-field regions, reducing processing costs while maintaining precision in key areas.

A mesh convergence study is conducted on the baseline design, and a mesh size of 0.01m near the wing is identified as the optimal balance between computational efficiency and accuracy, with minimal changes in lift, drag, and RBM beyond this refinement. For the baseline design, the mesh comprised 5,446,299 elements and 1,551,915 nodes, see Figure 1.

The aerodynamic performance of the wing with and without the winglet is analyzed at 1.5° angle of attack. The RANS equations are solved using ANSYS Fluent with the Spalart-Allmaras (SA) turbulence model, which is preferred for external flow problems due to its robustness and stability in solution convergence [19]. The inlet velocity is set to 38.9 m/s, while altitude conditions are adjusted to 2,000 m. The convergence criteria for all residual equations are set to 10^{-4} to ensure solution accuracy. The baseline design aerodynamic forces and root bending moment (RBM) at 1.5° angle of attack are summarized in Table 2.

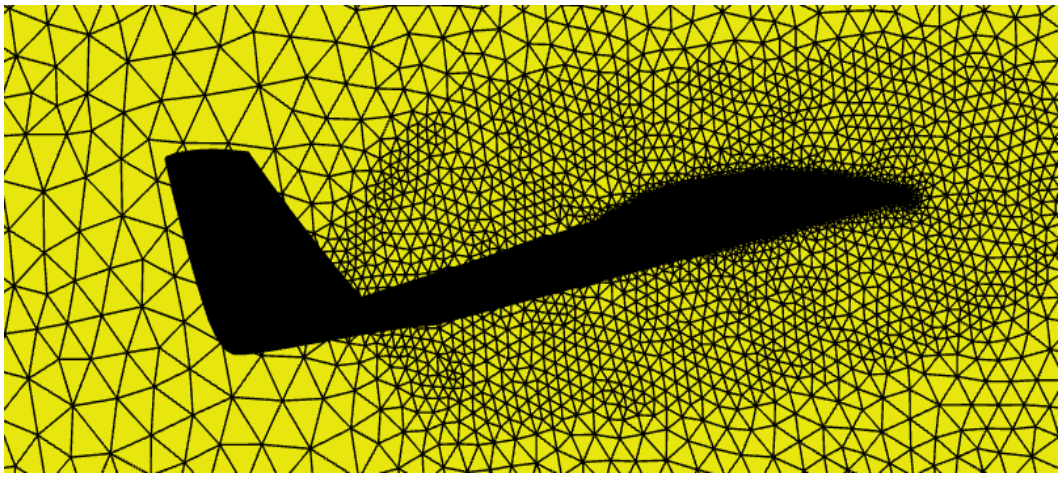


Figure 1: Illustration of the meshing strategy employed in this study

Table 2. Baseline aerodynamic forces and RBM at 1.5° angle of attack.

Configuration	***Drag (N)***	***Lift (N)***	***CD -***	***CL -***	***RBM (Nm)***	***L/D -***
Baseline wing (no winglet)	*47.88*	*1017*	*0.025*	*0.5300*	*1318*	*21.25*
Baseline design with winglet	*47.20*	*1126*	*0.0228*	*0.5452*	*1611*	*23.85*

2.3 Surrogate Modelling and Optimization

Using hi-fidelity CFD analysis directly with a GA optimizer would be extremely computationally expensive, as it would require 1000s of hi-fidelity simulations. Thus, to reduce computational cost, surrogate models are employed to interpolate a sample of CFD results.

First, Latin Hypercube Sampling (LHS) is used to generate 75 diverse winglet configurations to ensure broad design space exploration. These are then analyzed using the automated process described above. Next, sub-sets of these samples (25, 50 and 75) are used to train two different types of surrogate model: Kriging and polynomial regression, using the Surrogate Modeling Toolbox implemented in Matlab [20]. This allows us to investigate which type of surrogate model works best for each objective and how many samples are needed for an accurate model.

The multiobjective optimization problems are solved using the surrogate models and NSGA-II configured with a population size of 200, a maximum of 100 generations, a crossover probability of 0.9, a mutation probability of 0.5, and a mutation strength of 0.2 to ensure efficient convergence towards optimal designs.

The overall optimization process is illustrated in Figure 2. It starts with sampling the design space using LHS. Next, the geometry for each sample is automatically generated by a Matlab script. The geometry files are then imported into ICEM to generate mesh files, which are in turn imported into ANSYS fluent for CFD analysis. The results are then read back into Matlab for training the surrogate models. These are then used to solve the multiobjective optimization problems using NSGA-II in Matlab.

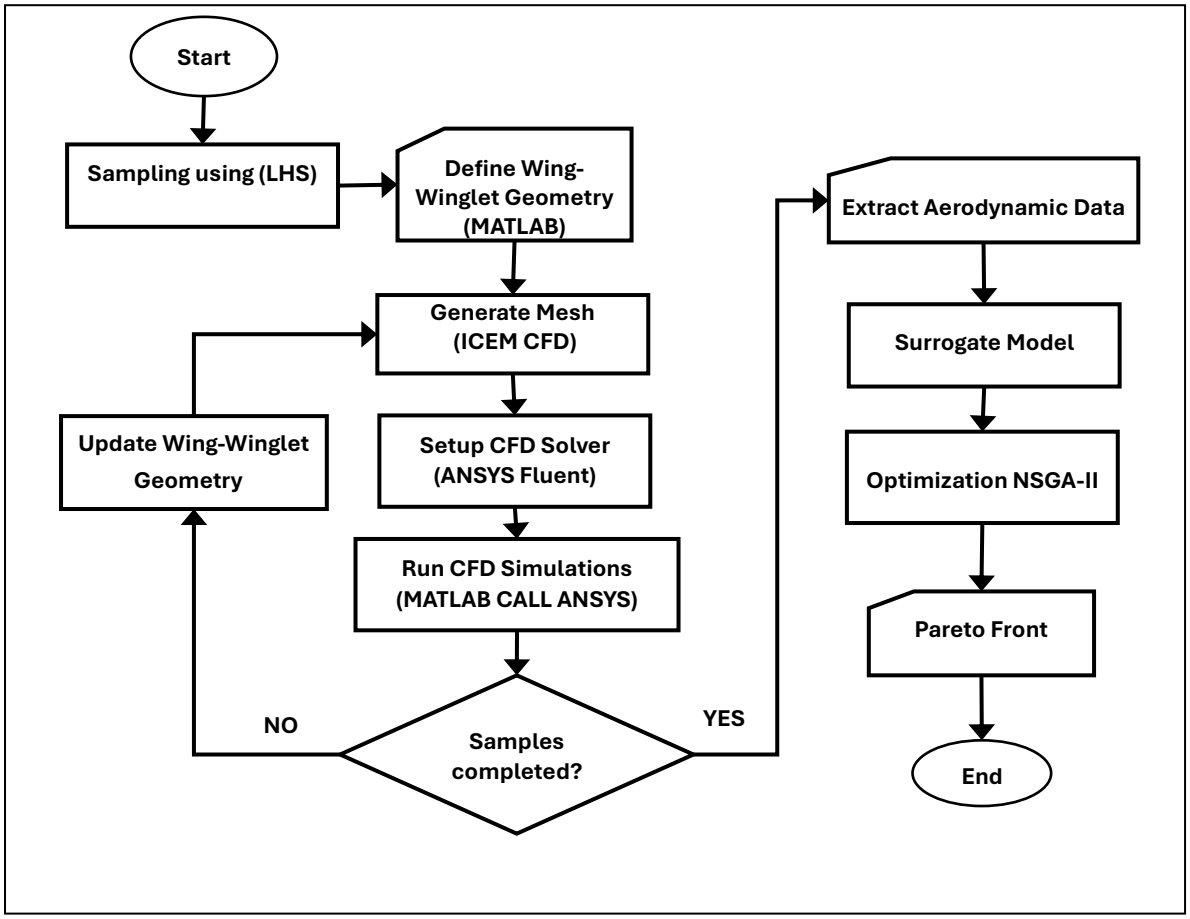


Figure 2: Flowchart and the main MATLAB scripts of aerodynamic optimization

3 RESULTS AND DISCUSSION

3.1 Surrogate Model Performance

To evaluate the accuracy of Kriging and polynomial regression models the performance metrics Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and R^2 are calculated using six representative points selected from the NSGA-II Pareto Front. These points are then tested in Fluent, and the results compared against those predicted by the surrogate models to evaluate model accuracy, see Table 3. Note that polynomial regression is tested up to order 6, with the 2nd-order polynomial proving best for lift and drag.

Based on the results presented in Table 3, 2nd-order polynomial regression demonstrates significant improvement at 50 samples, achieving better predictive performance overall. At this sample size, it attains high accuracy in both lift and drag predictions, outperforming Kriging in terms of R^2 while keeping error values relatively low. This makes polynomial regression with 50 samples the most effective balance between prediction accuracy and computational efficiency in this case. In addition, the position of the Pareto front does not significantly change when increasing the number of samples further, as shown in Figure 3.

At 75 samples, both models achieve high accuracy, with polynomial regression continuing to show superior predictive performance, particularly for drag. Kriging may still be considered but with reduced accuracy in comparison to polynomial regression when optimizing aerodynamic performance.

Table 3. Kriging and polynomial regression model errors for lift and drag prediction.

Model	*Drag MAE*	*Drag RMSE*	*Drag R^2*	*Lift MAE*	*Lift RMSE*	*Lift R^2*
Polynomial (25 samples)	*0.0606*	*0.0980*	*0.6229*	*1.5458*	*2.7234*	*0.4163*
Polynomial (50 samples)	*0.0483*	*0.0620*	*0.9744*	*0.5517*	*0.7146*	*0.9931*
Polynomial (75 samples)	*0.0393*	*0.0484*	*0.9877*	*0.5783*	*0.6837*	*0.9975*
Kriging (25 samples)	*0.0833*	*0.0995*	*0.8135*	*1.1717*	*1.5004*	*0.9375*
Kriging (50 samples)	*0.0567*	*0.0589*	*0.9032*	*0.6667*	*0.8489*	*0.9856*
Kriging (75 samples)	*0.0433*	*0.0535*	*0.9461*	*0.4567*	*0.5366*	*0.9947*

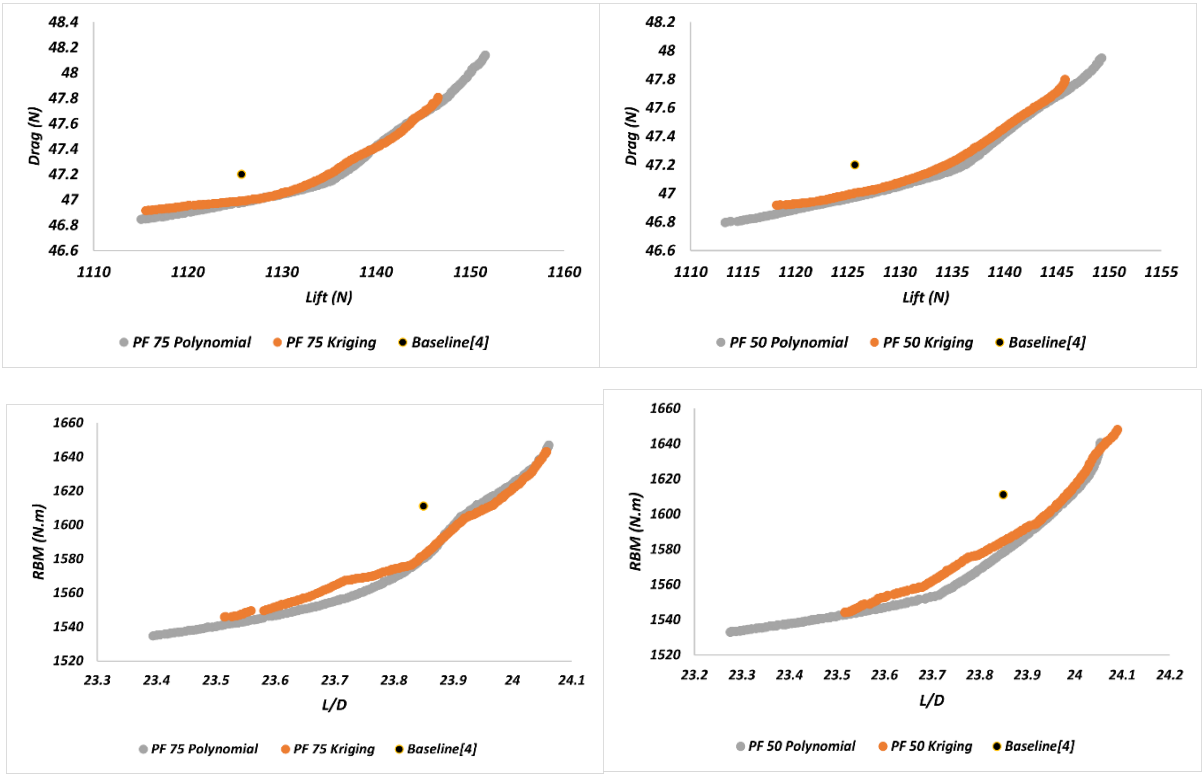


Figure 3: Pareto front using a surrogate model, with 50 and 75 samples

The performance of the surrogate models is also calculated for UAV winglet optimization when maximizing L/D ratio while minimizing RBM across different sample sizes. For this problem, Kriging with 50 samples is chosen as the best trade-off between accuracy and efficiency. As shown in Table 4, polynomial regression struggles with RBM predictions, whereas Kriging is significantly more accurate for all sample sizes. However, L/D predictions remain most effective with a 2nd-order polynomial.

A possible reason for the poor performance of polynomial regression for RBM is that it fluctuates between choosing 1st- and higher-order as dataset size increases, sometimes leading

to overfitting. For 25 samples, both L/D and RBM are best modelled with a 1st-order polynomial. At 50 samples, L/D shifts to a 2nd-order polynomial, while RBM varies between 1st, 2nd, and 3rd order. With 75 samples, L/D remains at 2nd order, while RBM fluctuates between 1st- and 6th-order models depending on the sampling, highlighting dataset sensitivity.

Table 4. Kriging and polynomial regression model errors for L/D and RBM prediction.

Model	*RBM MAE*	*RBM RMSE*	*RBM R^2*	*L/D MAE*	*L/D RMSE*	*L/D R^2*
Polynomial (25 samples)	*4.2379*	*5.6985*	*0.9569*	*0.0493*	*0.0644*	*0.8546*
Polynomial (50 samples)	*11.2739*	*11.8836*	*0.7311*	*0.0161*	*0.0182*	*0.9814*
Polynomial (75 samples)	*9.5618*	*10.8712*	*0.8618*	*0.0128*	*0.014*	*0.9931*
Kriging (25 samples)	*1.2417*	*1.965*	*0.9952*	*0.0371*	*0.0414*	*0.9466*
Kriging (50 samples)	*1.696*	*1.8715*	*0.994*	*0.034*	*0.038*	*0.9199*
Kriging (75 samples)	*1.5283*	*1.6824*	*0.9959*	*0.018*	*0.021*	*0.9814*

These results underscore the importance of selecting appropriate polynomial orders based on aerodynamic and structural objectives. Polynomial regression is ideal for L/D modelling, while Kriging remains superior for RBM, ensuring reliable structural predictions across different sample sizes (as shown in Table 4).

3.2 Aerodynamic Optimization

This section examines the optimization results to identify key design trends when optimizing for aerodynamic performance. A summary of designs and their performance from the Pareto front is shown in Table 5. We observe that a larger cant angle generally reduces induced drag whilst also reducing lift, depending on the winglet configuration. However, optimizing the cant angle, along with a larger span (0.375 m), can improve the overall L/D ratio, enhancing aerodynamic efficiency. Negative toe angles (θ) and twist angles (β) enhance lift generation while reducing drag by improving flow attachment and delaying separation. These aerodynamic adjustments optimize the pressure distribution over the winglet, leading to increased aerodynamic efficiency by maximizing lift while minimizing drag. The sweep angle (22.95° to 50°) influences vortex shedding, with higher values reducing drag and mitigating interference. Tip chord variations (0.15–0.25 m) impact lift generation, where larger values (0.20 m) maximize lift, and smaller values (0.17 m) minimize drag.

Table 5. Sample of Pareto front designs for maximizing lift and minimizing drag.

	Design Variables						Objectives	
Pareto Solution	**Span (m)**	**Tip chord (m)**	**Twist angle (β)**	**Toe angle (θ)**	**Sweep angle (λ)**	**Cant angle (γ)**	**Pareto Front Lift (N)**	**Pareto Front Drag (N)**
1	0.347	0.170	2.715	0.458	50	67.318	1117.658	46.914
2	0.349	0.169	2.827	0.472	50	67.302	1117.867	46.914
...	...	...	...	...	...	...	...	...
15	0.352	0.165	3.016	0.22	50	65.247	1120.271	46.927
46	0.355	0.173	5	0.565	50	40	1125.164	46.992
77	0.364	0.176	5	-0.201	50	40.954	1130.005	47.075
112	0.368	0.179	5	-0.941	50	41.511	1133.971	47.184
128	0.374	0.187	5	-0.97	47.658	41.854	1135.406	47.236
179	0.375	0.196	1.174	-1	36.612	43.505	1140.065	47.46
...	...	...	...	...	...	...	...	...
244	0.373	0.211	-5	-1	22.582	43.267	1145.855	47.795
245	0.373	0.212	-5	-1	22.955	43.53	1145.869	47.802

3.3 Structural Optimization

Table 6 presents the multi-objective optimization results, balancing aerodynamic efficiency (maximizing L/D) with structural performance (minimizing RBM). The winglet span varies between 0.335 m and 0.374 m, demonstrating that larger spans enhance L/D but contribute to increased RBM.

The twist angle (β) and toe angle (θ) play a critical role in load management. The highest L/D values, correspond to a twist angle of 4.86° and a toe angle of -0.96°, whereas the lowest RBM occurs at $\theta = 1°$ with a reduced span of 0.335 m.

The cant angle ranges from 40° to 70°. Higher cant angles can enhance aerodynamic efficiency by increasing lift but may also increase structural mass due to higher bending moments. The sweep angle, varying from 0° to 49.32°, redistributes aerodynamic loads, reducing RBM while maintaining aerodynamic stability.

These results demonstrate how including structural performance (RBM) in the optimization problem changes optimal winglet design. For example, the designs that maximize aerodynamic performance (L/D) (see designs 242 onwards in Table 6) have similar characteristics to designs in the middle of the Pareto front for aerodynamic performance only (e.g. see design 128 in Table 5). However, designs that minimize RBM (hence reduce wing mass) have high cant angles, with small sweep angles and spans, which is not seen in Table 5.

Table 6. Sample of designs from Pareto front for maximizing L/D and minimizing RBM.

	Design Variables						Objectives	
Pareto Solution	***Span (m)***	***Tip chord (m)***	***Twist angle (β)***	***Toe angle (θ)***	***Sweep angle (λ)***	***Cant angle (γ)***	***Pareto Front L/D (-)***	***Pareto Front RBM (N.m)***
1	*0.336*	*0.162*	*4.86*	*1.000*	*0.00*	*70.00*	*23.52*	*1544*
2	*0.336*	*0.162*	*4.86*	*1.000*	*0.20*	*70.00*	*23.52*	*1544*
...	*...*	*...*	*...*	*...*	*...*	*...*	*...*	*...*
65	*0.345*	*0.150*	*-5.00*	*0.928*	*0.00*	*70.00*	*23.70*	*1562*
127	*0.336*	*0.150*	*4.78*	*0.963*	*49.32*	*64.65*	*23.85*	*1585*
147	*0.337*	*0.150*	*4.86*	*1.000*	*49.22*	*58.25*	*23.90*	*1593*
181	*0.354*	*0.150*	*4.86*	*1.000*	*49.32*	*41.85*	*23.96*	*1605*
207	*0.374*	*0.150*	*4.86*	*0.927*	*47.89*	*46.07*	*24.00*	*1616*
242	*0.374*	*0.213*	*4.86*	*-0.852*	*49.32*	*52.37*	*24.06*	*1639*
...	*...*	*...*	*...*	*...*	*...*	*...*	*...*	*...*
256	*0.374*	*0.239*	*4.86*	*-0.973*	*49.22*	*49.04*	*24.09*	*1647*
257	*0.374*	*0.238*	*4.86*	*-0.963*	*49.32*	*47.15*	*24.09*	*1648*

4 CONCLUSION

This study presents a multi-objective optimization framework for UAV winglet design, balancing aerodynamic efficiency and structural performance. The methodology integrates high-fidelity CFD simulations, surrogate modelling, and NSGA-II optimization, systematically evaluating two key trade-offs: maximizing lift while minimizing drag and maximizing the L/D ratio while minimizing the RBM.

A key outcome of this research is the identification of an optimal sample size for accurate surrogate modelling while maintaining computational efficiency. Second-order polynomial regression with 50 samples emerges as the best approach for aerodynamic optimization, achieving high-accuracy predictions for lift and drag with minimal error. Conversely, Kriging with 50 samples is the preferred method for structural performance optimization, excelling in RBM prediction whilst providing similar accuracy for L/D prediction, compared with polynomial regression. These findings confirm that 50 samples provide the best trade-off between computational cost and prediction accuracy.

The design parameter analysis confirms that increasing winglet span and cant angles improve lift generation, but increase RBM, which may result in increased wing mass. Negative twist and toe angles help redistribute loads and mitigate excessive structural stresses. Sweep angle optimization plays a crucial role in reducing RBM, further reinforcing the importance of structural considerations in aerodynamic optimization. The Pareto front analysis highlights that a well-balanced approach is necessary to achieve high aerodynamic efficiency without compromising structural efficiency.

Overall, Kriging is recommended for structural-driven optimization with fewer computational resources, while polynomial regression is more suitable for aerodynamic optimization at larger datasets. Future work will refine surrogate models and expand the methodology to diverse UAV configurations for enhanced performance predictions.

ACKNOWLEDGEMENTS

The first author gratefully acknowledges the scholarship from the Saudi Arabia government.

REFERENCES

[1] Mueller, T. J., & DeLaurier, J. D. (2003). Aerodynamics of small vehicles. *Annual Review of Fluid Mechanics*, *35*, 89–111.

[2] Kontogiannis, S. G., & Ekaterinaris, J. A. (2013). Design, performance evaluation and optimization of a UAV. Aerospace science and technology, 29(1), 339-350.

[3] Panagiotou, P., Kaparos, P., Salpingidou, C., & Yakinthos, K. (2016). Aerodynamic design of a MALE UAV. *Aerospace Science and Technology*, *50*, 127-138.

[4] Panagiotou, P., Kaparos, P., & Yakinthos, K. (2014). Winglet design and optimization for a MALE UAV using CFD. *Aerospace Science and Technology*, *39*, 190–205.

[5] Takenaka, K., Hatanaka, K., Yamazaki, W., & Nakahashi, K. (2008). Multidisciplinary design exploration for a winglet. *Journal of aircraft*, *45*(5), 1601-1611.

[6] ALI. I., & Kurtulus, D. F. (2017). Winglet design and analysis for low-altitude solar-powered UAV. *International Journal of Sustainable Aviation*, *3*(1), 64-86.

[7] Makgantai, B., Subaschandar, N., & Jamisola, R. S. (2021, June). Design optimization of wingtip devices to reduce induced drag on fixed-wings. In *2021 International Conference on Unmanned Aircraft Systems (ICUAS)* (pp. 1459-1465). IEEE.

[8] Moghimi Esfandabadi, M. H., & Djavareshkian, M. H. (2024). Investigating selected methods to improve aerodynamic coefficients and better performance of UAV. *Technology in Aerospace Engineering*, *8*(3), 1-16.

[9] Pfeiffer, N. (2004). Numerical winglet optimization. In *42nd AIAA Aerospace Sciences Meeting and Exhibit* (p. 213).

[10] Peerlings, B. (2018). A review of aerodynamic flow models, solution methods and solvers and their applicability to aircraft conceptual design. Delft University of Technology: Delft, The Netherlands.

[11] Liu, X., Song, L., Jin, D., Hu, T., & Li, S. (2015, December). Numerical Simulation of Winglet to Optimization Design. In 2015 4th International Conference on Mechatronics, Materials, Chemistry and Computer Engineering (pp. 575-578). Atlantis Press.

[12] Beechook, A., & Wang, J. (2013, September). Aerodynamic analysis of variable cant angle winglets for improved aircraft performance. In 2013 19th International Conference on Automation and Computing (pp. 1-6). IEEE.

[13] Elham, A., & van Tooren, M. J. (2014). Winglet multi-objective shape optimization. *Aerospace Science and Technology*, *37*, 93-109.

[14] Lawton, S., & Crawford, C. (2014, June). Investigation and optimization of blade tip winglets using an implicit free wake vortex method. In Journal of Physics: Conference Series (Vol. 524, No. 1, p. 012033). IOP Publishing.

[15] Ribeiro, A. F. P., Awruch, A. M., & Gomes, H. M. (2012). An airfoil optimization technique for wind turbines. Applied Mathematical Modelling, 36(10), 4898-4907.

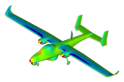
[16] Elfarra, M. A., Sezer-Uzol, N., & Akmandor, I. S. (2014). NREL VI rotor blade: numerical investigation and winglet design and optimization using CFD. Wind Energy, 17(4), 605-626.

[17] Madsen, M. H. A., Zahle, F., Sørensen, N. N., & Martins, J. R. (2019). Multipoint high-fidelity CFD-based aerodynamic shape optimization of a 10 MW wind turbine. Wind Energy Science, 4(2), 163-192.

[18] Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. A. M. T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE transactions on evolutionary computation, 6(2), 182-197.

[19] Kostić, Č. (2015). Review of the Spalart-Allmaras turbulence model and its modifications to three-dimensional supersonic configurations. Scientific Technical Review, 65(1), 43-49.

[20] Forrester, A., Sobester, A., & Keane, A. (2008). *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTIDISCIPLINARY DESIGN AND OPTIMIZATION OF STRUT-BRACED HYBRID-ELECTRIC AIRCRAFT UNDER MULTI-TRAJECTORY SCENARIOS.

Pablo Norczyk Simon[1*], Rauno Cavallaro[2], Danlin Zheng Zhang[3], Jānis Briž[4]

1 - Barcelona Supercomputing Centre (BSC), Spain. pnorczyk@bsc.es

2 - Department of Aerospace Engineering, Universidad Carlos III de Madrid, Spain
rauno.cavallaro@uc3m.es

3 - CRIDA A.I.E., Spain. dzheng@e-crida.enaire.es

4 - Riga Airport, Latvia. J.Brizs@riga-airport.com

Abstract. *Hybrid-electric aircraft offer a promising pathway for reducing aviation emissions, but integrating novel powerplant components and their mission-level control presents significant design challenges. This study employs a GEMSEO-based multidisciplinary optimization framework, to perform comprehensive sizing and optimization of the INDIGO aircraft—characterized by a high aspect ratio wing and distributed hybrid-electric propulsion optimized for reduced landing and takeoff (LTO) impact. The research evaluates optimization strategies across three scenarios: (1) single trajectory analysis, (2) multiple trajectories with a common control strategy, and (3) multiple trajectories with mission-specific control strategies. Results highlight the critical importance of multi-point analysis for achieving robust and reliable aircraft design. Although mission-specific control strategies provided more optimal outcomes, they also increased mission planning complexity. Furthermore, the study examines aircraft performance under five distinct failure scenarios, offering insights into the resilience and adaptability of hybrid-electric propulsion architectures. These findings contribute to a deeper understanding of optimization methodologies for emerging hybrid-electric aircraft configurations.*

Keywords: Multidisciplinary Optimization, INDIGO, Large Aspect Ration Wings, Local Air Quality, Noise, Distributed Hybrid Electric Aircraft

1 INTRODUCTION

Aviation stands at a critical juncture, facing the urgent need to decarbonize operations amid escalating environmental concerns. As of 2023, the sector accounted for approximately 2.5% of global CO_2 emissions [1], exhibiting the fastest emission growth rates among all transportation modes [2]. By 2050, the International Civil Aviation Organization (ICAO) projects that, without decisive intervention, international aviation emissions could triple compared to 2015 levels [3]. This environmental urgency drives the exploration of innovative aircraft configurations aimed at minimizing aviation's impact on emissions, Local Air Quality, and Noise (LAQN). The EU-funded INDIGO project [4] exemplifies such efforts, integrating Large Aspect Ratio Wings (LARW) with Distributed Hybrid Electric Propulsion (DHEP) to enhance efficiency and reduce environmental impact.

However, developing novel configurations is often a complex endeavor; with a lack of historical data and many disciplines often competing with each other, traditional methods are often unable to leverage the new technologies. Multidisciplinary Analysis and Optimization (MDAO) has emerged here as a transformative technology, enabling holistic design synthesis by simultaneously evaluating various aircraft systems and eliminating the need for manual iteration through complex trade-offs.

Although MDAO can substantially accelerate the aircraft design process, optimizing for single-point or single-trajectory scenarios frequently yields non-robust configurations that, while optimal at their design points, exhibit inferior performance under off-design conditions or prove unreliable for diverse mission profiles. This challenge becomes particularly critical when considering real-world aircraft operations, where inherent variability in trajectories—especially during initial climb and final descent phases across airports with different procedures—can result in designs over-fitted to specific scenarios.

This paper investigates the effects of various trajectories on hybridization and distributed propulsion strategies to perform a sensitivity analysis comparing single-mission versus multi-mission optimization approaches. The paper is structured as follows: first, we present the methodology for modeling the aircraft architecture, followed by an overview of the MDA framework. We then formulate the MDO problem statement, analyze the resulting data, and draw conclusions regarding optimal design strategies for robust hybrid aircraft configurations.

2 METHODOLOGY

In this paper, `TOPAZ`, a GEMSEO-based [5] MDAO framework for aircraft synthesis, is utilized. The framework develops aircraft synthesis capabilities comparable to other established codes such as OpenConcept [6], FastOAD [7], MOTIVATION [8] or Aviary [9]. Originally created with a special focus on mass distribution and internal volume allocation for novel configurations such as hydrogen-powered aircraft [10], `TOPAZ` performs comprehensive aircraft sizing through mission-performance evaluation—flying nominal and extended missions while assessing takeoff performance under various failure scenarios. This performance evaluation employs well-established handbook formulas and surrogate models. Additionally, the platform employs JAX [11] Automatic Differentiation within select disciplines through the GEMSEO-JAX plugin. The following section details the models used for aircraft optimization and the assessment of control strategies across various mission profiles.

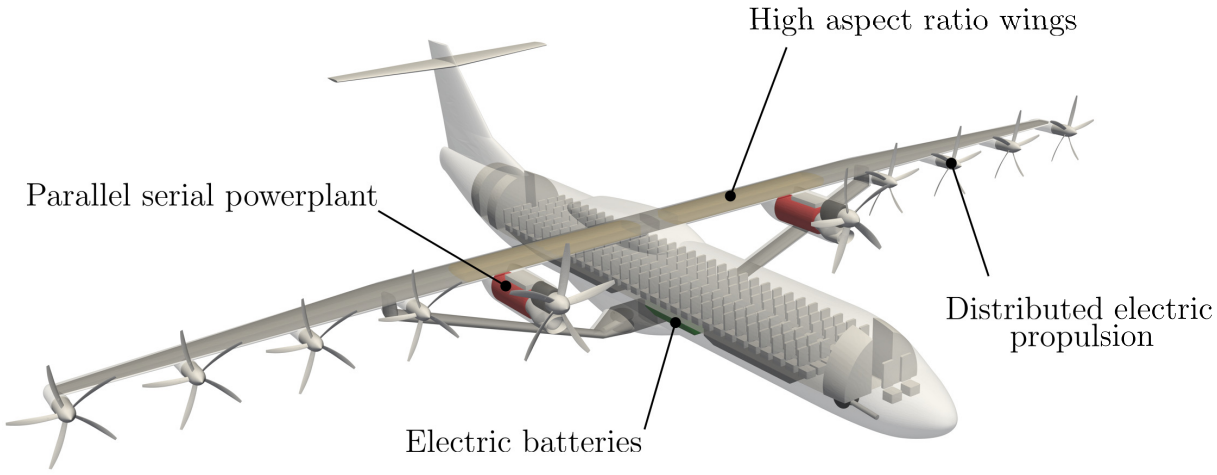


Figure 1: Conceptual model of the INDIGO aircraft architecture.

2.1 The INDIGO Architecture

Funded by HORIZON Europe, the INDIGO project aims to improve Local Air Quality and Noise (LAQN) by developing a novel mid-range aircraft designed for 1000 nm route operations. The concept, illustrated in Fig. 1, achieves these improvements through aerodynamic enhancements featuring strut-braced, high aspect ratio wings, and a cutting-edge parallel-serial propulsion architecture. This architecture combines turboshaft engines operating in parallel with distributed electric propulsion units arranged along the wing in a serial configuration.

The Airbus A320neo serves as the reference aircraft for the INDIGO project. Accordingly, the INDIGO aircraft is designed to carry an equivalent payload weight, W_{PL}, of 20 tons with a maximum takeoff weight (MTOW) of 79 tons. The cabin layout replicates that of the A320neo, accommodating 180 passengers in a single-class economy configuration [12].

2.1.1 Aerodynamic performance estimation

For aerodynamic performance estimation, the platform employs a classical drag polar with linearized lift and pitching moment coefficients (Eq.1a, Eq.1b, Eq.1c), comprising 11 coefficients that capture the first and second-order dependencies with angle of attack α (denoted by the subscripts α and $\alpha\alpha$), first and second-order dependencies with elevator deflection δ_e (denoted by the subscripts δ_e and $\delta_e\delta_e$) and values at zero incidence (denoted by the subscript 0). These aircraft-level coefficients are derived by aggregating individual fitted surfaces (wing and tail for lift and drag, with fuselage, fairings, and vertical tail contributing to drag).

$$C_L = C_{L_0} + \alpha \cdot C_{L_\alpha} + \delta_e \cdot C_{L_{\delta_e}} \tag{1a}$$

$$C_D = C_{D_0} + \alpha \cdot (C_{D_\alpha} + \alpha \cdot C_{D\alpha\alpha}) + \delta_e \cdot (C_{D_{\delta_e}} + \delta_e \cdot C_{D_{\delta_e\delta_e}}) \tag{1b}$$

$$C_M = C_{M_0} + \alpha \cdot C_{M_\alpha} + \delta_e \cdot C_{M_{\delta_e}} \tag{1c}$$

The polar coefficients for the wing are obtained through a surrogate model of a strut-braced wing trained with optimal wing designs. The training dataset is generated by performing aerodynamic shape optimizations using VSPAERO [13]. Each optimization

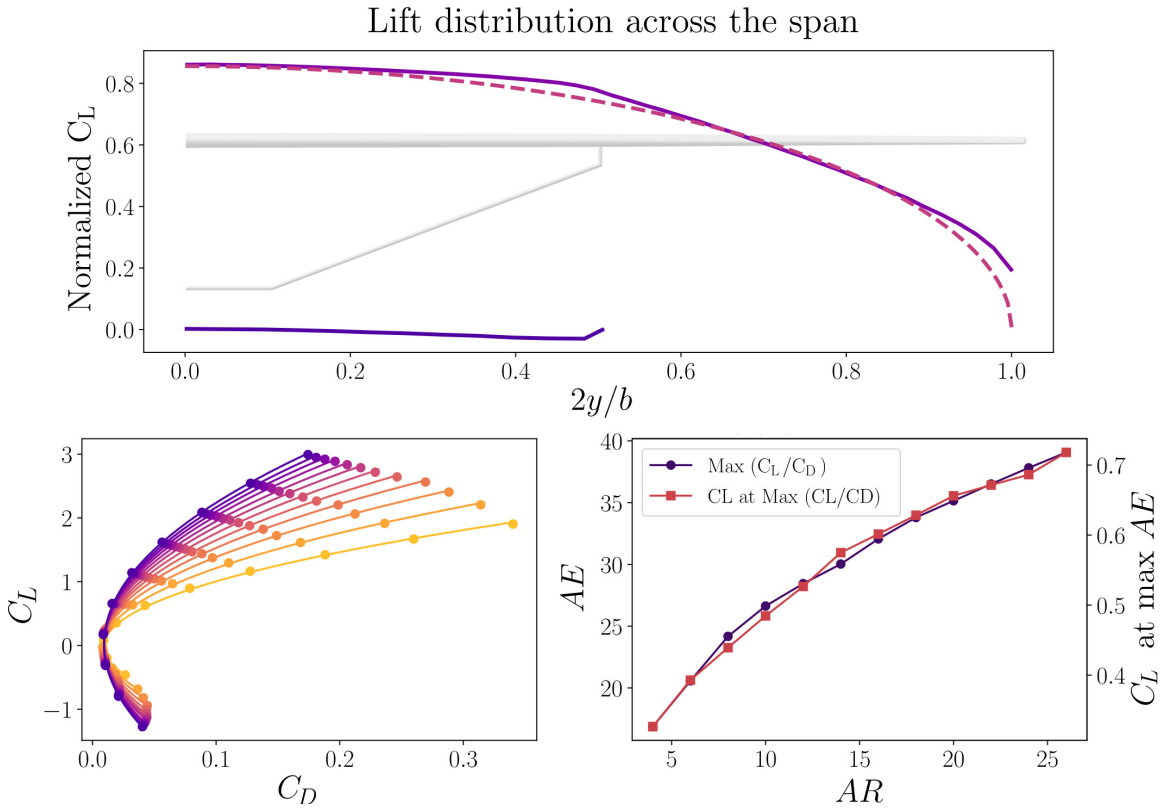


Figure 2: Top: normalized lift distribution across the wingspan for an $AR = 22$. Bottom left: reconstructed aerodynamic polars from $AR = 4$ to $AR = 26$ in $\Delta AR = 2$ increments; in marker, values at direct VSPAERO executions. Bottom right: maximum aerodynamic efficiency (AE) and C_L at that point obtained from the reconstructed polars.

is performed at a fixed aspect ratio (AR) using a parametric model composed of a main wing and a supporting strut. The main wing features a constant chord up to the mid-section ($y = b/4$), beyond which it tapers at a constant taper ratio (λ_w). The strut has a constant chord equal to 0.5 times the wing root chord. Design variables for these discipline-level optimizations include wing twists at the root $\theta_{w,y=0}$, mid $\theta_{w,y=b/4}$ and tip section $\theta_{w,y=b/2}$, the strut twist the root $\theta_{w,y=0}$ and the outer wing tapper ratio λ_w. Subsequently, the coefficients C_{L_0}, C_{L_α}, C_{D_0}, C_{D_α}, $C_{D_{\alpha\alpha}}$ from each optimized configuration are fitted across AR for system-level optimization. This methodology emulates a bi-level approach, with the inner discipline-level design variables remaining hidden from the system-level optimization, which only manages the AR. Fig.2 illustrates the resultant optimized lift distributions and aerodynamic polars of each optimized wing.

For the horizontal tail, vortex drag is calculated using classical methods such as DATCOM [14]. Parasite drag for the fuselage, horizontal tail, and vertical tail is determined using drag breakdown coefficients from established sources [15–18]. Downwash estimation assumes an elliptical lift distribution to decouple wing and tail coefficients. While this assumption may not apply to all wing platforms, the near-elliptical distributions achieved through optimization (Fig. 2) support its validity. To ensure proper tail sizing, the model calculates the neutral point and static margin at each flight phase, following the methodology previously implemented in [10].

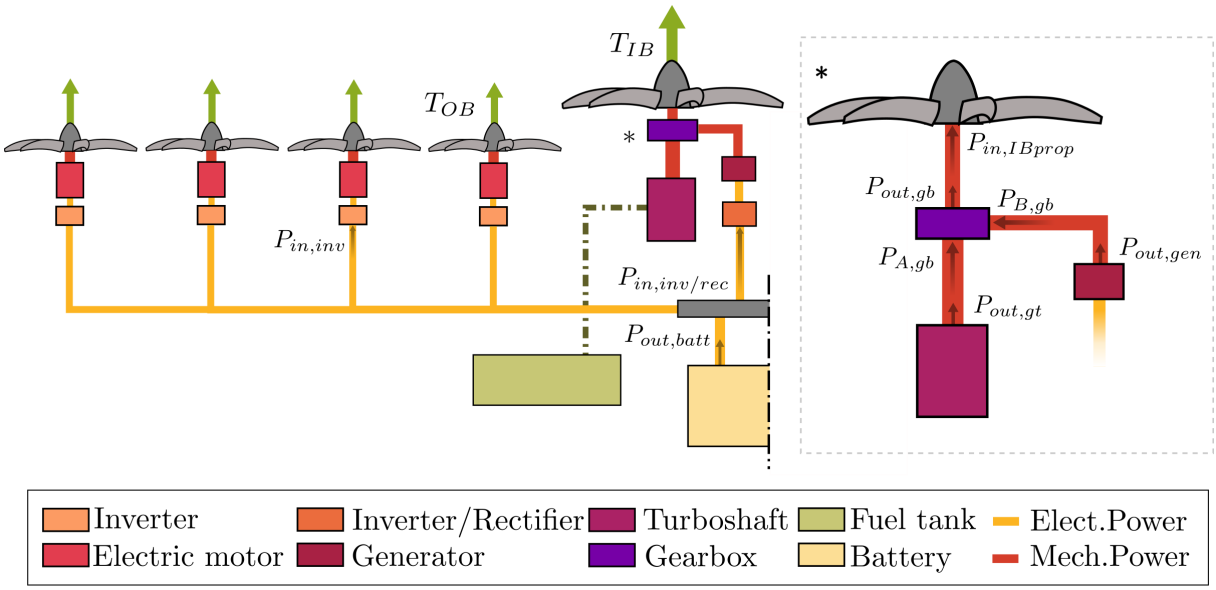


Figure 3: Definition of the parallel-serial configuration for the INDIGO aircraft.

2.1.2 Powerplant component modeling

The INDIGO powerplant architecture comprises two components: an inboard parallel system and an outboard serial electric propulsion system. The inboard parallel system features an inboard propeller connected to a gearbox operated by a turboshaft engine and an electric machine. This electric machine can function both as an electric motor (supplying power to the gearbox) and as a generator (drawing power from the gearbox). To manage the power split between the gas turbine and the electric machine, a local hybridization factor HF^* is introduced. The powers provided by the turboshaft and the electric machine are expressed as shown in Eq. 3. For positive hybridization values ($HF^* \in [1, 0[$), the electric machine acts as a motor, delivering additional power to assist the propeller. For negative values ($HF^* \in [0, -\infty]$), the electric machine acts as a generator, extracting mechanical power from the turboshaft to support the electrical system

$$P_{A,gb} = P_{out,gb} \cdot (1 - HF^*) \quad , \qquad P_{B,gb} = P_{out,gb} \cdot HF^* \tag{2}$$

For the outboard serial system, eight Electric Propulsion Units (EPUs) are incorporated, four per wing. Each EPU consists of an electric motor with its inverter, driving an outboard propeller. All EPUs are connected to an electric bus, which distributes power from the inboard electric machine. Consequently, the total power required from the battery is the sum of the power demands from all components linked to the bus.

To regulate thrust distribution between the inboard and outboard propulsion systems, a split factor ψ is introduced. The thrust produced by each inboard propeller, T_{IB}, and by each outboard propeller, T_{OB}, are defined as:

$$T_{IB} = T_{total} * (1 - \psi)/2 \quad , \qquad T_{OB} = T_{total} * \psi/8 \tag{3}$$

For the component models, the performance of the gas turbine is estimated using a surrogate model of the multi-spool turboshaft implemented in pyCycle [19]. This model is then "rubberized" by normalizing the output power and interpolating across simulated altitudes and Mach numbers (Fig. 4).

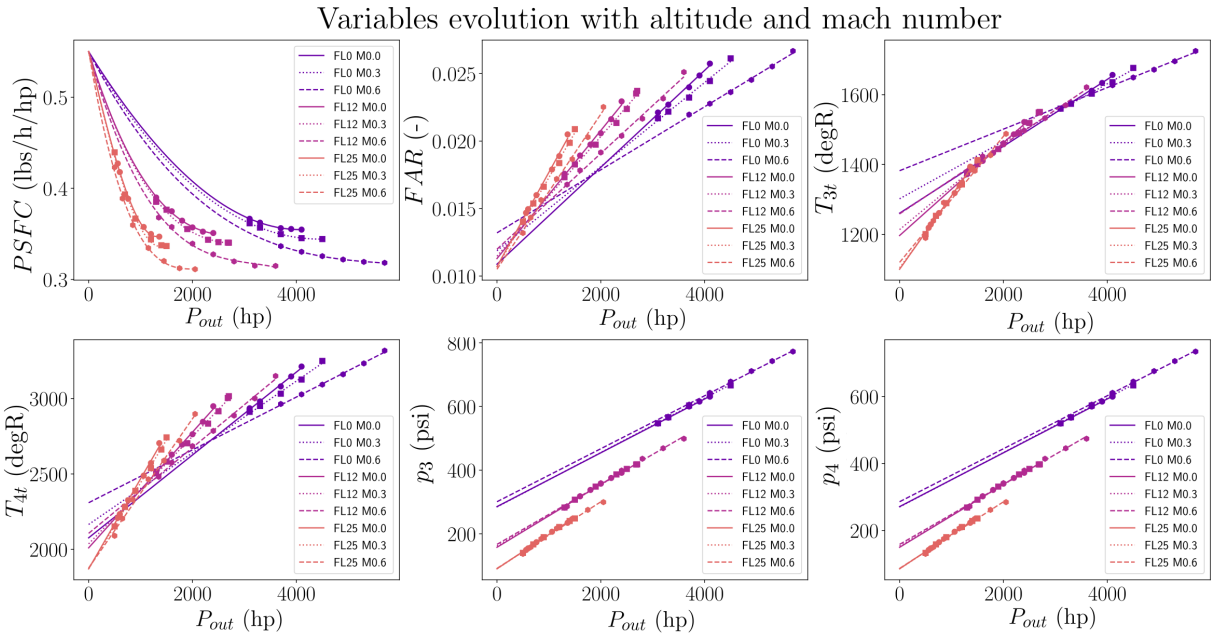


Figure 4: Engine stage variables for a kerosene multi-spool turboshaft sized at $M = 0.6$, $h = 25000$ ft, and $W_{shp} = 1800$ hp. The lines and curves represent fitted responses, while the markers correspond to the obtained data.

Propeller performance is assessed using a surrogate model generated from XROTOR [20] by producing $J - C_P$ propeller maps. To account for different propeller designs, the propellers are sized at $h = 6000$ m, with a forward velocity of 0.6 M, operating at 800 rpm. Disk loadings range from 286 N/m^2 to 1145 N/m^2, represented by a sizing factor ϖ in increments of $\Delta\varpi = 0.25$. This approach allows the optimizer to design propellers capable of operating at higher power coefficients (higher ϖ) at the cost of reduced propeller efficiency η (see Fig. 5).

Clearance constraints are imposed on the propellers to ensure a minimum roll angle of 8° (Ψ_{roll}) and a clearance of 0.1 m for both the inboard and outboard propellers, denoted as $\Delta Clear_{IB}$ and $\Delta Clear_{OB}$, respectively (see ig. 6). For all optimizations, the inboard propeller is assumed to be positioned at $0.25b/2$, while the outboard propellers are equally spaced from $0.5b/2$ to $b/2$.

For the electric components, fixed efficiencies of $\eta = 0.95$ are assumed for the electric motor, inverter, rectifier, and electric machine. These efficiency values were provided by partners from the University of Strathclyde [21]. For the battery, a pack-level specific energy of 350 Wh/kg and a pack-level specific power of 1000 W/kg are assumed, based on projections from [22].

2.1.3 Aicraft weight estimation

For weight estimation, empirical formulas were applied. Wing mass, furnishings, nacelles, landing gear, and horizontal and vertical tail weights were calculated using Raymer's formulas [15], while the fuselage mass was determined using the formula from Torenbeek [18]. All values were corrected based on the weight breakdown of the A320-200 presented in [23]. Additional adjustments for the T-tail were made using weight breakdown data from the ATR42 [23] and ATR72 [24]. To properly account for the effects

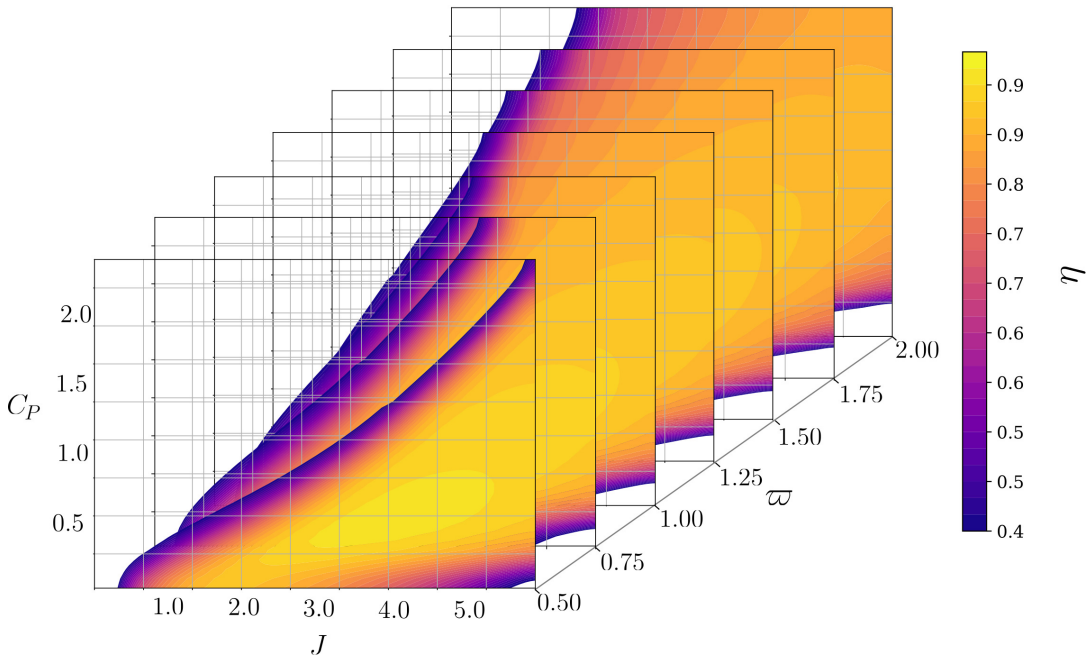


Figure 5: Propeller $J - C_P$ maps at different design disk loadings.

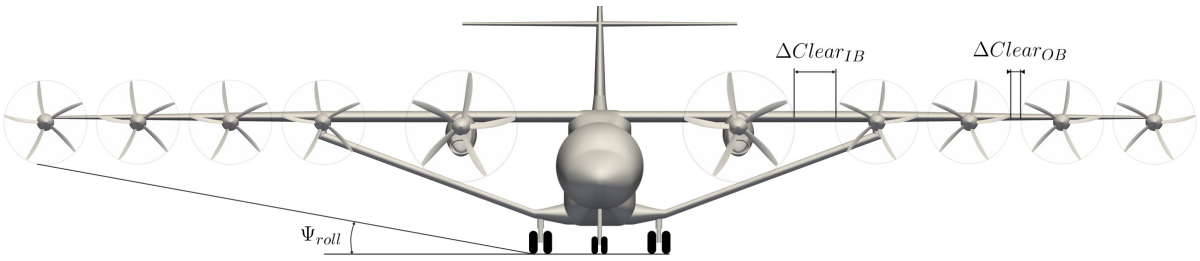


Figure 6: Definition of the inboard clearance constraint $\Delta Cler_{IB}$, outboard clearance constraint $\Delta Cler_{OB}$, and roll angle clearance constraint Ψ_{roll}

of aspect ratio, a modified version of Raymer's dependency on AR was employed, based on second-order fitted curves from [25]. The evolution of the normalized[1] wing weight is illustrated in Fig. 7, showing that the classical Raymer equation tends to underestimate the increase in wing weight with rising AR. For this study, the curve corresponding to the strut-braced CFRP wing will be applied. 5

For the powerplant weight estimation, adjusted specific power values are applied to various components. A specific power of 2.6 kW/kg is considered for the turboshaft dry core and auxiliary systems. For the gearbox, a specific power of 16.4 kW/kg is assumed [26], while electric machines (including electric motors and generators) are assigned a specific power of 15 kW/kg [27]. Additionally, inverters and rectifiers are characterized by a specific power of 23 kW/kg [28]. For the payload, passengers are assumed to be homogeneously distributed across the cabin, occupying all available seats, with an average weight of 75.4 kg [29]. It is further assumed that 1800 kg of luggage is stored in the overhead compartments, with the remaining 4628 kg placed in the lower fuselage cargo compartments.

Additional fudge weights for avionics, power distribution subsystems, and operational

[1]With an aspect ratio of $AR = 10$

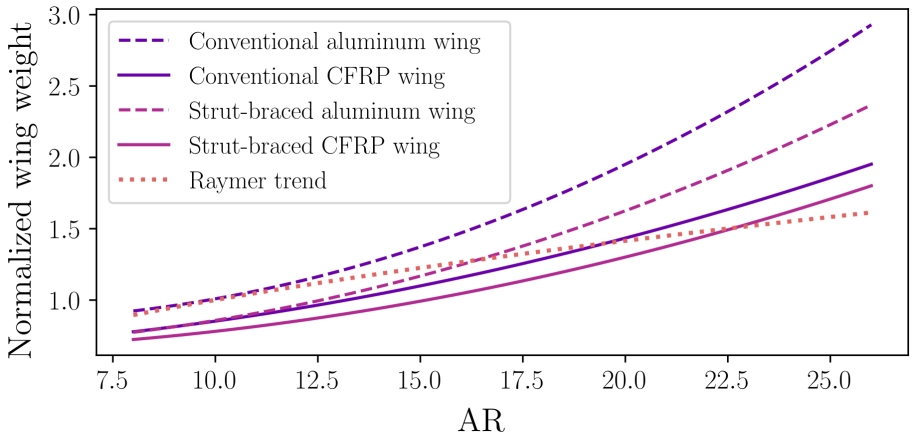


Figure 7: Tendency of the normalized wing weight for various wing configurations.

items were introduced based on the A320 weight breakdown [23] to accurately adjust the operational empty weight.

2.2 Flight Mission

The model evaluates a nominal mission of 1,200 nm with a 400 nm diversion segment, including a 30-minute holding pattern. A cruise altitude of 6 km and a flight speed of Mach 0.6 are selected for the analysis.

To account for the inherent variability of low-altitude procedures (below 900 m), 10 distinct procedures are incorporated, each with its corresponding initial climb and final descent profiles. These procedures are derived from flight data collected at four European airports: LEBL (Barcelona El Prat), LEMD (Adolfo Suárez Madrid–Barajas), RIX (Riga), and DTM (Dortmund). The selected trajectories were extracted from a dataset generated through extensive research conducted by CRIDA [30], representing the most common flight paths as well as best and worst case scenarios in terms of Local Air Quality (LAQ) and noise impacts. For this study, the most common trajectories were selected along with six additional distinct trajectories to ensure sufficient variability in the analysis.

To ensure trajectory continuity, each low-altitude procedure was integrated with the remainder of the mission by interpolating speed and flight path angles using spline functions. All flight phases are discretized into 10 nodes, and for the multiple-trajectory cases, only the first trajectory is connected to the extended mission. The resulting mission profiles are illustrated in the upper panel of Figure 8, with detailed views of the climb and descent segments shown in the lower panel. This approach introduces a controlled source of uncertainty for low-altitude procedures in the optimization process while maintaining nearly identical cruise segments (the cruise duration varies slightly due to differences in climb and descent durations).

2.2.1 Low-speed failure modes

To adapt the One Engine Inoperative (OEI) condition to the parallel-serial hybrid-electric architecture, five distinct failure cases are considered. For each failure case, the Balanced Field Length (BFL) is calculated following the methodology presented in [31].

- **FC1**: This case involves the simultaneous failure of the left inboard propeller and its adjacent outboard propeller. Despite these failures, the left turboshaft remains

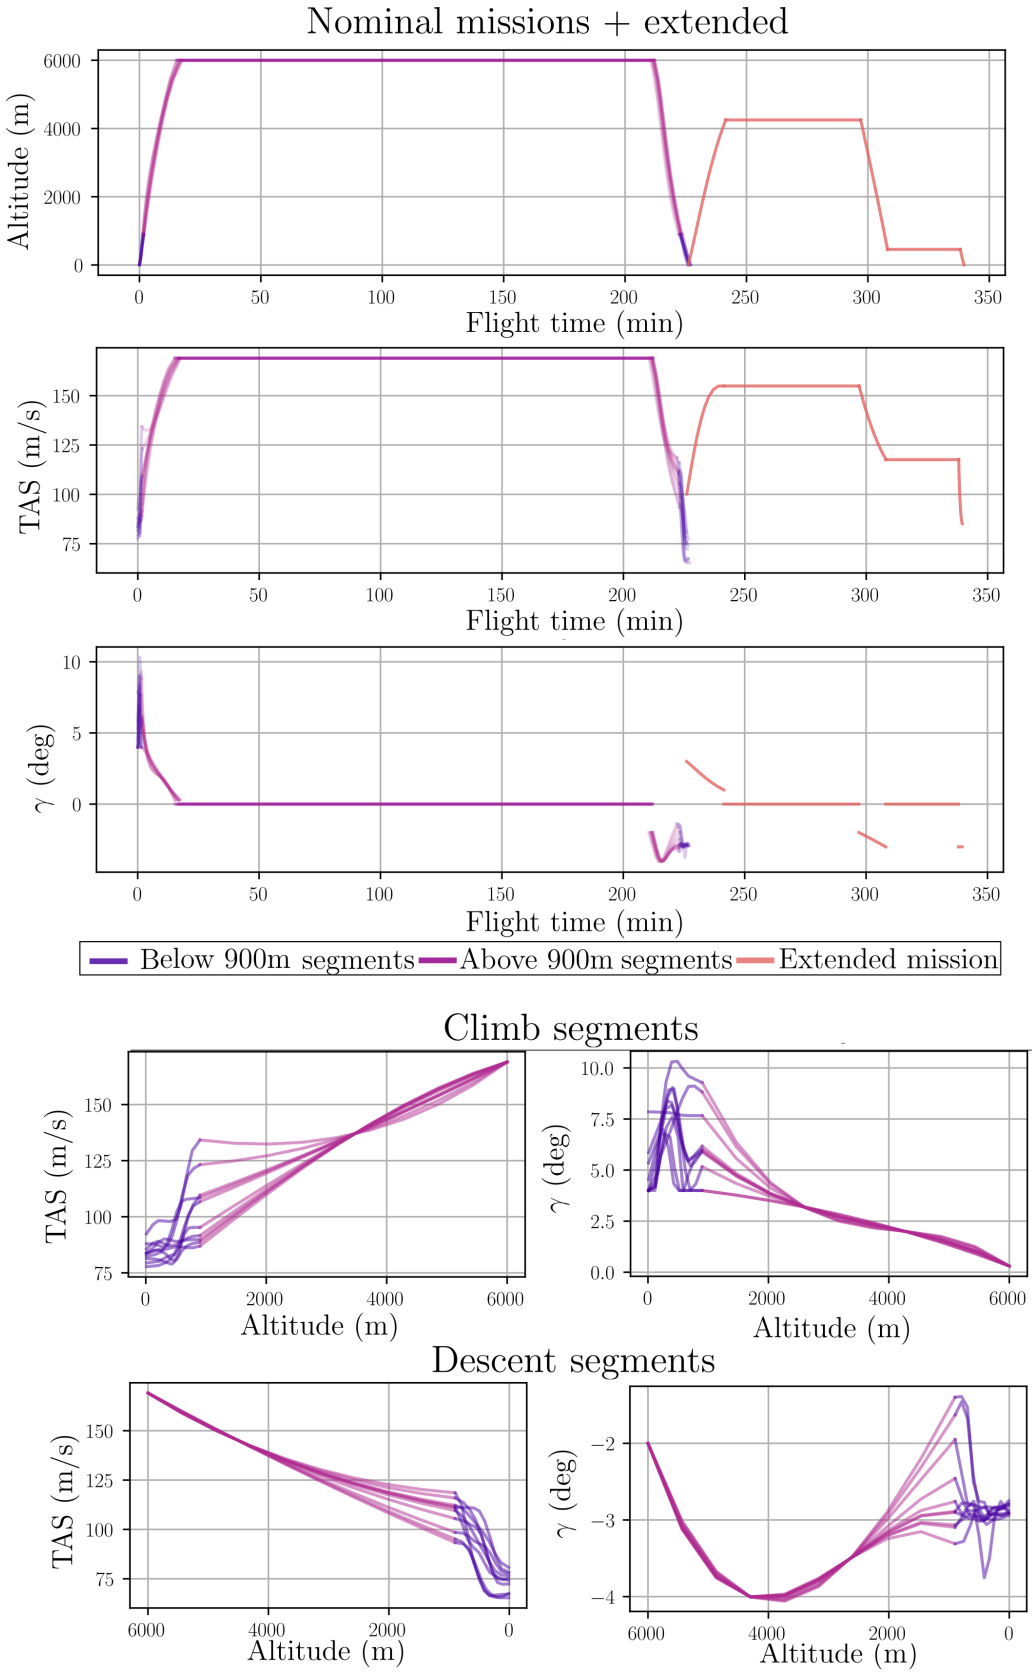


Figure 8: Top three figures: vertical altitude profile, true airspeed profile and flight path angle γ profile for the nominal missions and the extended mission. Bottom four figures: closeup of the TAS and the γ as a function of the altitude for the climb and descent segments.

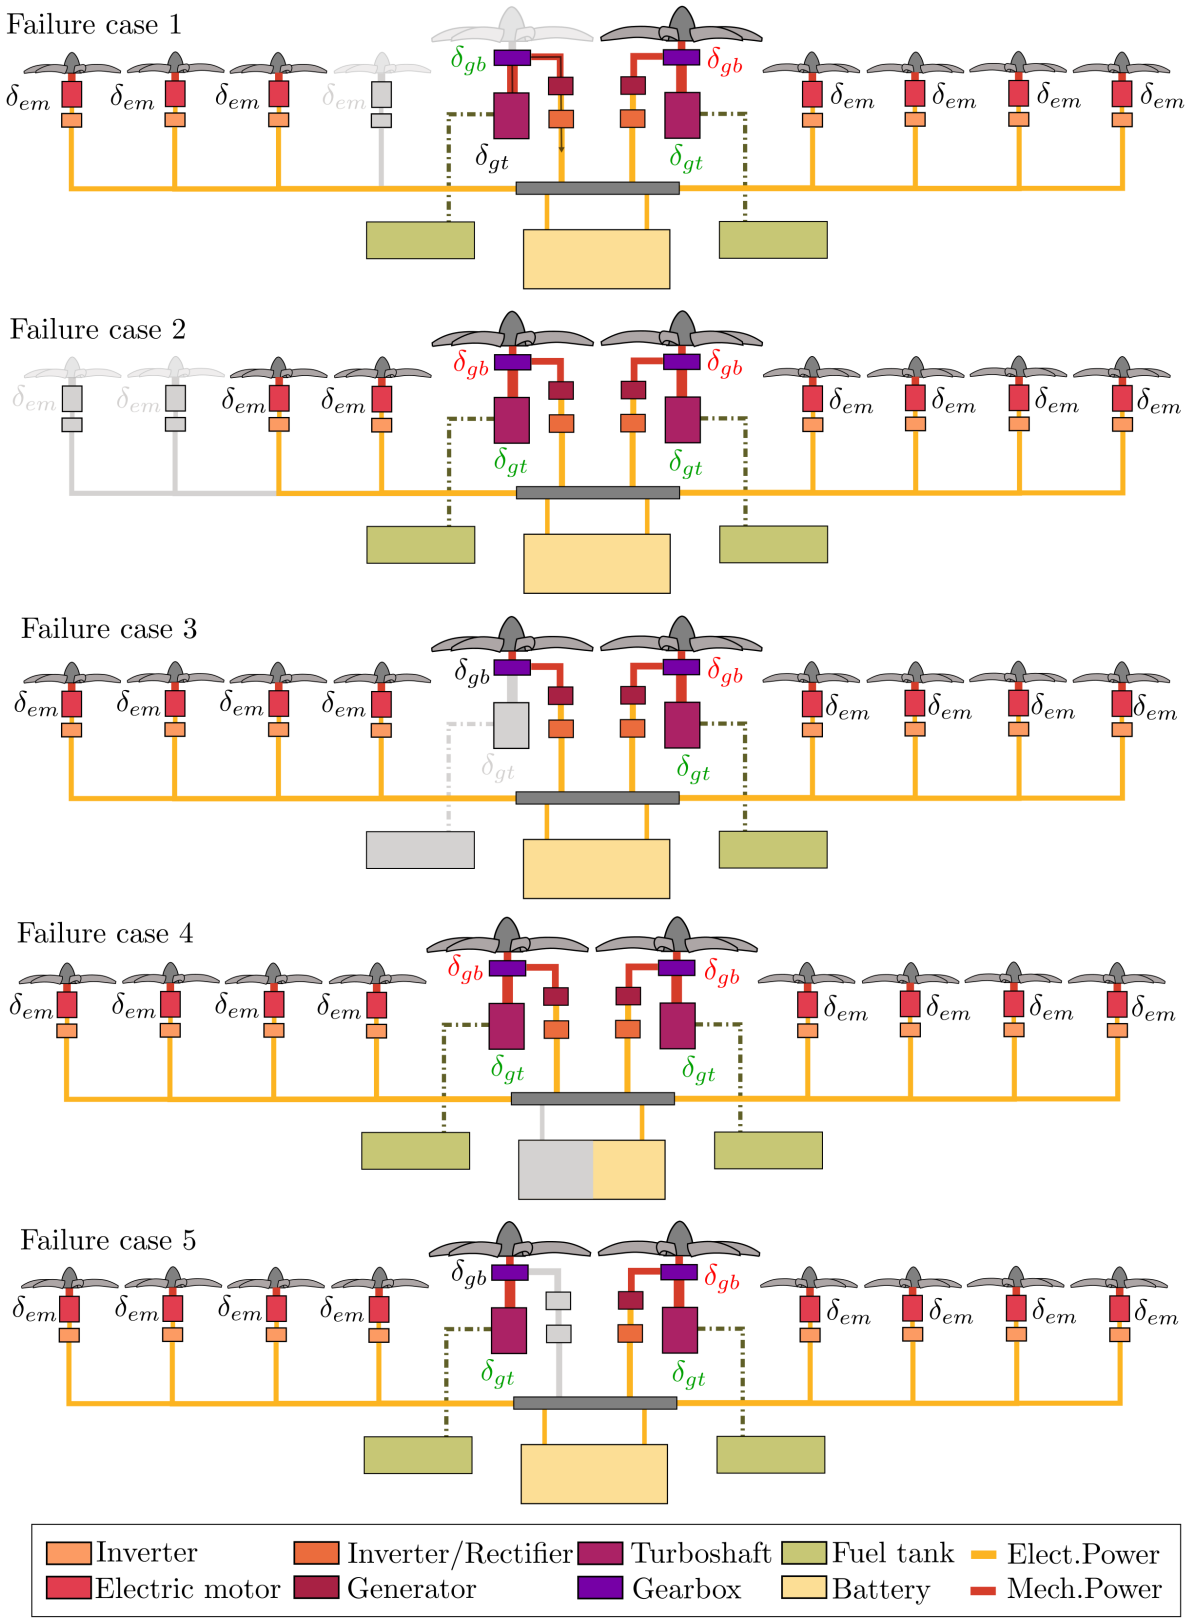


Figure 9: Failure case scenarios for the hybrid electric powerplant. Controls in red indicate fixed values, controls in black indicate optimization design variables and controls in green indicate state variables.

operational as a turbogenerator at a prescribed gas turbine throttle setting (δ_{gt}), which is treated as a design variable. Consistent with OEI conditions, the right gearbox is working at full power ($\delta_{gb} = 1$), while the outboard electric motor throttle (δ_{em}) remains a design variable. The right gas turbine throttle is determined by the local hybridization factor, which may differ from the nominal hybridization factor under non-failure conditions.

- **FC2**: This case examines the failure of both outermost outboard propeller units, representing the most critical scenario in terms of yawing moment due to asymmetric propulsion. As in FC1, both gearbox throttles are fixed at $\delta_{gb} = 1$, while the remaining outboard propeller throttles are treated as design variables.

- **FC3**: In this scenario, the left gas turbine fails while the inboard propeller remains functional. The left propeller operates at a reduced gearbox throttle (δ_{gb}), treated as a design variable—with power supplied exclusively by the electric machine. The right propulsion system operates with a gearbox throttle fixed at $\delta_{gb} = 1$, while electric motor throttles and the right local hybridization factor are treated as design variables.

- **FC4**: This case simulates the failure of half of the battery pack. Both gearbox throttles are fixed at $\delta_{gb} = 1$, with a new local hybridization factor introduced as a design variable. Outboard electric throttles are also treated as design variables to optimize power distribution under these conditions.

- **FC5**: This case mirrors FC3 but with the failure of the left inboard electric component instead of the gas turbine. The gearbox operates at a reduced throttle (as a design variable), with power provided solely by the turboshaft engine. All other control parameters follow the same methodology as in FC3.

For failure case FC2, the yawing moments produced by asymmetric thrust during the rolling phase were analyzed, along with the maximum counteracting yawing moment achievable with a sideslip angle of $\beta = 10°$ and a rudder deflection of $\delta_r = 10°$. The aerodynamic coefficients for sideforce (C_{Fy}) and yawing moment (C_{Mz}) were derived from a VSPAero regression model incorporating the fuselage, vertical tail, and horizontal tail configurations across various vertical tail areas. To ensure directional controllability during this critical failure mode, a constraint was implemented requiring the aerodynamic moment to exceed the asymmetric thrust moment.

3 THE MDA

The Multidisciplinary Design Analysis (MDA) problem is formulated using a Multidisciplinary Feasible (MDF) architecture, implemented with a fixed-point Gauss-Seidel algorithm. Conceptually, the MDA involves several interrelated disciplines (Fig. 10). First, given the parametric aircraft design variables, the system computes component dimensions, center of mass locations, and component weights, while providing the optimizer with constraint information related to clearances and mass consistency. Next, a BFL analysis is conducted across multiple failure scenarios to determine required takeoff distances and evaluate powerplant component performance. This process continues with mission-specific analysis, computing powerplant performance throughout each flight

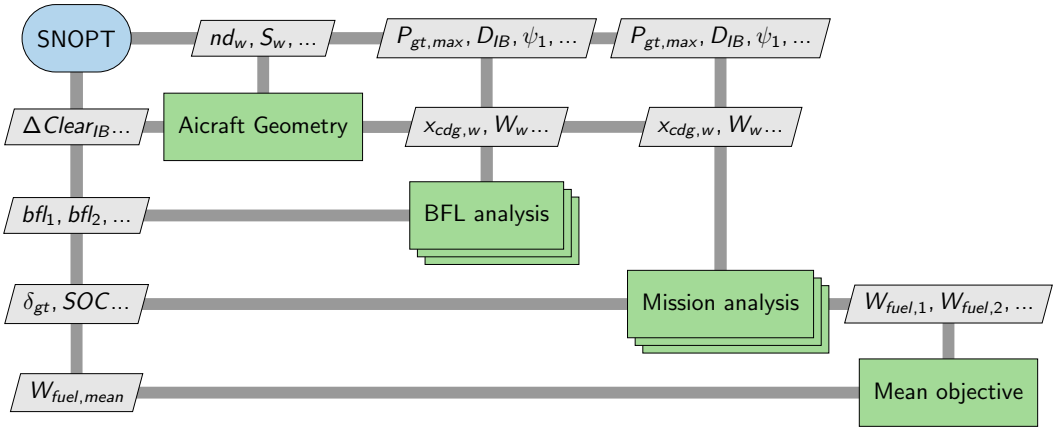


Figure 10: XDSM of the main performed optimization

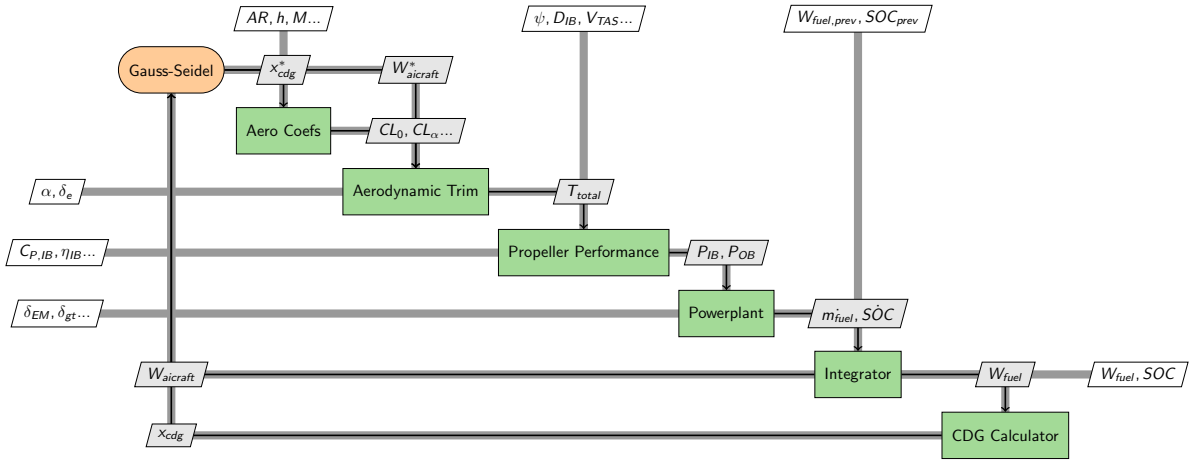


Figure 11: MDA performed on each phase of the flight mission

phase. Finally, these mission-specific results are integrated to construct a mean fuel consumption objective function.

Examining the interconnection within aircraft performance disciplines, Fig.11 illustrate the multidisciplinary analysis performed at each flight phase to determine aircraft performance. The analysis begins by calculating aerodynamic coefficients based on the current aircraft geometry and flight conditions. These coefficients are then used to trim the aircraft under steady-state aerodynamic conditions, as described by Eq.4:

$$q \cdot S_{ref} \cdot CL(\alpha, \delta_e) - W \cdot \cos(\gamma) = 0 \tag{4a}$$

$$T_{total} - q \cdot S_{ref} \cdot CD(\alpha, \delta_e) - W \cdot \sin(\gamma) = 0 \tag{4b}$$

$$C_M(\alpha, \delta_e) = 0 \tag{4c}$$

After trimming, the total thrust requirement T_{total} is distributed between the inboard and outboard propellers. A Newton solver determines the input power P required for each propeller to generate the necessary thrust T at the current advance ratio J. AThese power requirements feed into the powertrain model, which calculates the powerplant's operating point and provides the fuel mass flow rate $\dot{m}_{fuel}$ and battery discharge rate $\dot{SOC}$. These

values are integrated and linked to data from previous flight phases. The updated fuel mass is then fed into the center of gravity (CG) calculator, which adjusts the CG position accordingly. The implicit relationship between CG position and aircraft mass affecting aerodynamic coefficients and trim is solved using the Gauss-Seidel method.

For the takeoff analysis, a similar process is employed, with the Gauss-Seidel method now handling the speeds of each segment and the equations of motion being solved with a nested Newton-solver (for more detail see [10]).

4 THE TESTCASES

In this study, three optimization analyses are proposed to investigate the effects of control strategy and the impact of variability in low-altitude procedures:

- **Case 1:** Optimization considering a single mission profile.
- **Case 2:** Optimization considering multiple trajectories under a unified control strategy.
- **Case 3:** Optimization considering multiple trajectories, each with its own trajectory-specific control strategy.

For all analyses, the objective function to be optimized is the average block fuel, representing the fuel consumed to fly the nominal part of the mission. This is expressed as:

$$W_{\text{block fuel,mean}} = \frac{1}{N} \sum_{i=1}^{N} W_{\text{block fuel},i} \tag{5}$$

In the case of a single trajectory (Case 1), the mean is computed from a single value.

The design variables for the optimization are presented in Tab. 1 (left). These variables include: Propeller design parameters, such as the diameter, tangential tip Mach number[2], and disk loading factor; tail surface sizing (vertical and horizontal tail surfaces[3]); wing characteristics (aspect ratio and its position along the fuselage); powerplant components (power rating, representing the maximum peak power values at which components can operate; and battery weght).

The control variables ψ and HF^* are defined using discrete nodes to ensure consistency across different flight phases. They remain constant during cruise and procedures below 900 m, while varying linearly along the nominal climb and descent segments of the nominal mission profile, as well as throughout the extended mission. For the single-mission optimization (Case 1) and the multiple-mission optimization with unified control (Case 2), a common control profile is applied across all missions. In contrast, for the multiple-mission optimization with customized control (Case 3), individual control profiles are tailored specifically for each nominal mission. For low-speed operations, the control variables are defined to remain constant under All Engines Operative (AEO) conditions, with

[2]A constraint is enforced over the Mach number obtained with the velocity triangle to ensure the propeller does not exceed the high-transonic regime.

[3]the aspect ratio and taper ratios for the tail are pre-selected based on configurations similar to the ATR72 and Dash-8 400.

the flexibility to adjust during failure scenarios. When operating under asymmetric conditions during failure cases, additional control variables are introduced: the gas turbine control parameter δ_{gt} for FC1, and the gearbox control parameter δ_{gb} for FC3 and FC5.

The optimization problem includes constraints detailed in Tab. 1 (right). These constraints govern the throttle level δ of each component, defined as the ratio of operating power to power rating. During normal operation, components must maintain 90% throttle (maximum continuous operating power), while reaching 100% during failure scenarios. For gas turbines, throttle calculations account for operating power relative to the maximum available power at the given altitude and operating point. Additional constraints include mass margin, requiring total component weight to remain below MTOW (Eq.6), propeller linear and angular clearances, residual yawing moment $\mathcal{R}(M_z)$, battery state of charge, aircraft aerodynamic static margin, and balanced field length for each failure mode.

$$\text{Mass Margin} = MTOW - \left(\sum_{i \in \text{components}} W_i \right) - W_{\text{fuel}} - W_{\text{PL}} \tag{6}$$

Design variable (units)	Lower	Upper	Constraint (units)	Lower	Upper
IB diameter (m)	1	6	δ gas turbine (-)	—	0.9/1
IB M_{tang} (-)	0.4	0.8	δ gearbox (-)	—	0.9/1
IB factor (-)	0.5	2	δ electric motor (-)	—	0.9/1
OB diameter (m)	1	6	δ inverter (-)	—	0.9/1
OB M_{tang} (-)	0.4	0.8	δ electric machine (-)	—	0.9/1
OB factor (-)	0.5	2	δ battery (-)	—	0.9/1
Sref HTP (m^2)	1	30	Propeller Cp (-)	—	3
Sref VTP (m^2)	17	40	Propeller M_{true} (-)	—	0.85
AR_w (-)	4	26	Clearance IB (m)	0.1	—
nd_w (-)	0.3	0.6	Clearance OB (m)	0.1	—
Turboshaft rating (MW)	1	20	Ψ_{roll} (deg)	10	—
Gearbox rating (MW)	0.05	20	$\mathcal{R}(M_z)$ (N·m)	0.01	—
Elec motor rating (MW)	0.05	20	Mass margin (m)	0.01	—
Inverter rating (MW)	0.05	20	BFL (m)	—	2000
Elec machine rating (MW)	0.05	20	SOC (-)	0.2	1
Battery Weight (kg)	100	20000	Static margin (-)	0.3	—
ψ (-)	0.1	0.9			
HF^* (-)	-30	1			
Takeoff δ_{em} (-)	0.1	0.9			
Takeoff HF (-)	-30	1			
FC1 δ_{gt} (-)	0.01	0.99			
FC3 δ_{gb} (-)	0.01	0.99			
FC4 δ_{gb} (-)	0.01	0.99			

Table 1: Design variables and bounds (left), and associated constraints with limits (right). Green values apply across all mission phases following the aforementioned logic. The values in Red implement one value for AEO and one value for each FC. Values in Pink are applied for all phases. Values in Orange apply for each FC.

Case 1 comprises 53 design variables with 2219 constraints. Case 2 retains the same variables but expands to 6719 constraints due to additional powerplant, State of Charge (SOC), and static margin constraints applied across all mission nodes and phases. Case 3 significantly increases complexity to 143 design variables (as control design variables are applied to each nominal mission) while maintaining the 6719 constraints from Case 2.

The selected optimizer is SNOPT (Sparse Nonlinear OPTimizer) [32], is a gradient-based optimization software using Sequential Quadratic Programming (SQP) to solve large-scale, constrained nonlinear problems with sparse gradients; SNOPT it is widely used for MDO applications due to its efficiency and robustness.

5 RESULTS

The optimized aircraft configurations for all three cases are presented in Fig.12, with their technical specifications detailed in Tab.2.

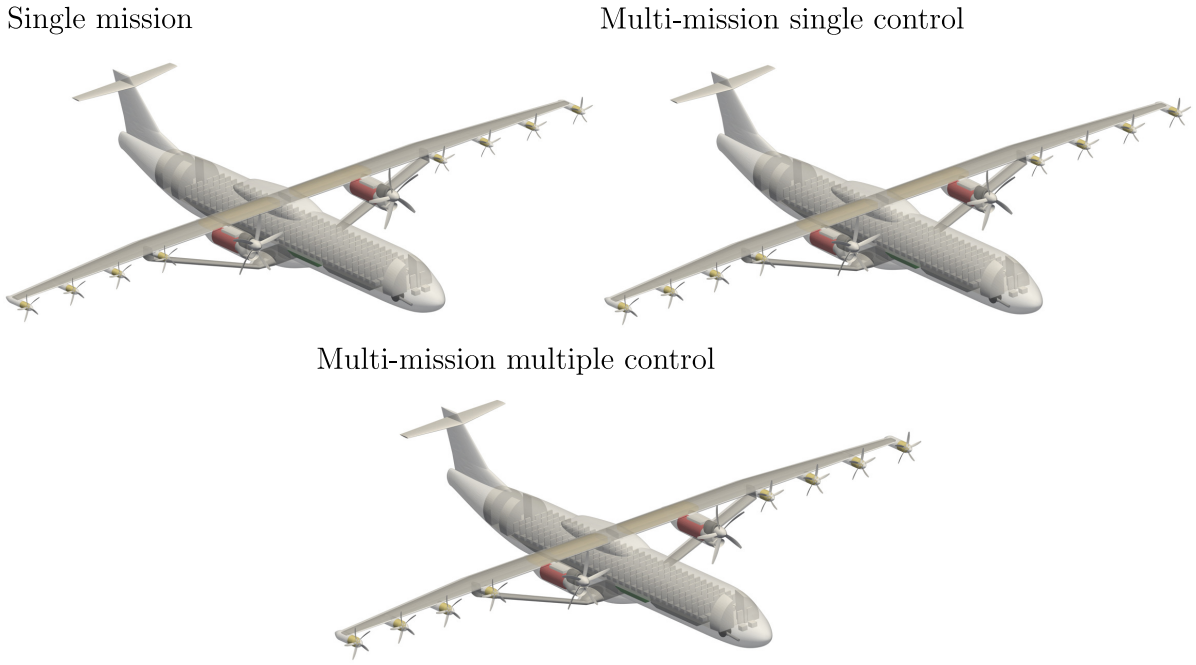


Figure 12: CAD model of the optimized configurations for the three optimization cases.

The three designs exhibit remarkable similarity, with only minor variations in propeller dimensions and electric motor power ratings. This stems primarily from the fact that while low-altitude procedures introduce some variability, the majority of mission parameters remain consistent across all cases. Consequently, the control strategies shown in Fig. 13 demonstrate nearly identical approaches across all scenarios, with the outboard propellers providing additional thrust during climb and contributing approximately $30 - 37\%$ of the thrust during cruise. Meanwhile, the electric machine predominantly functions as a generator, diverting $27 - 34\%$ of the turboshaft power to the electrical system. It is important to note that variations are more pronounced between different cases than between missions within the same case (Case 3), indicating that these differences are primarily driven by the specific aircraft configuration identified by the optimizer rather than mission requirements.

Variables (units)	Single Trajectory	Single Control	Multiple Control
IB diameter (m)	4.8451	4.7568	4.6767
IB Tang Mtip (-)	0.6587	0.6652	0.6515
IB factor (-)	1.021	0.8348	0.9379
OB diameter (m)	2.1850	2.3811	2.558
OB Tang Mtip (-)	0.5426	0.5011	0.553
OB factor (-)	0.7634	0.7601	1.0135
Sref HTP (m^2)	11.5382	11.6815	12.7385
Sref VTP (m^2)	**17**	**17**	**17**
AR_w (-)	22.0134	22.0112	21.9986
nd_w (-)	0.4748	0.4738	0.4715
Turboshaft rating (MW)	5.8761	5.9039	5.8003
Gearbox rating (MW)	4.8147	4.8703	4.6331
Electric power rating (MW)	0.8638	1.121	1.7608
Inverter rating (MW)	0.8589	1.1146	1.8442
Elec machine rating (MW)	2.683	2.5987	2.8199
Battery Weight (kg)	11982.9	11112.5	11040.7
Total fuel (Nom+Ext) (kg)	6638.6	6691.5	6780.11

Table 2: Optimization results for the optimization of the aircraft in the three optimization cases.

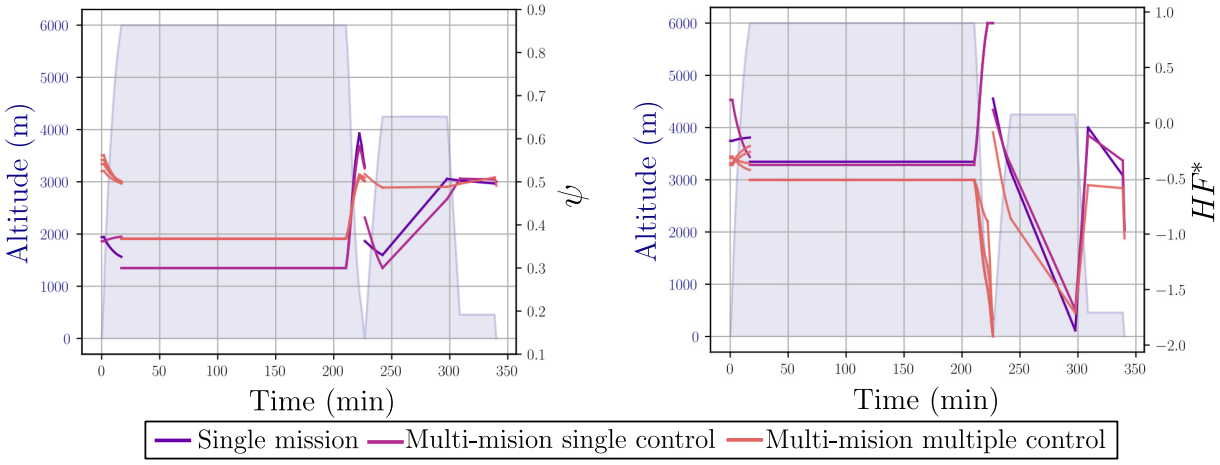


Figure 13: Optimized HF^* and ψ controls for the nominal and extended missions for the three optimization cases.

Examining the performance-defining variables across the nominal mission, Fig. 14 (top) illustrates the propeller efficiency η for both inboard and outboard propellers. The results confirm that the optimizer effectively sizes the propellers for cruise efficiency, with only Case 3 exhibiting lower outboard efficiencies due to the selection of a higher thrust-producing outboard propeller design, corresponding to higher cruise split factors ψ.

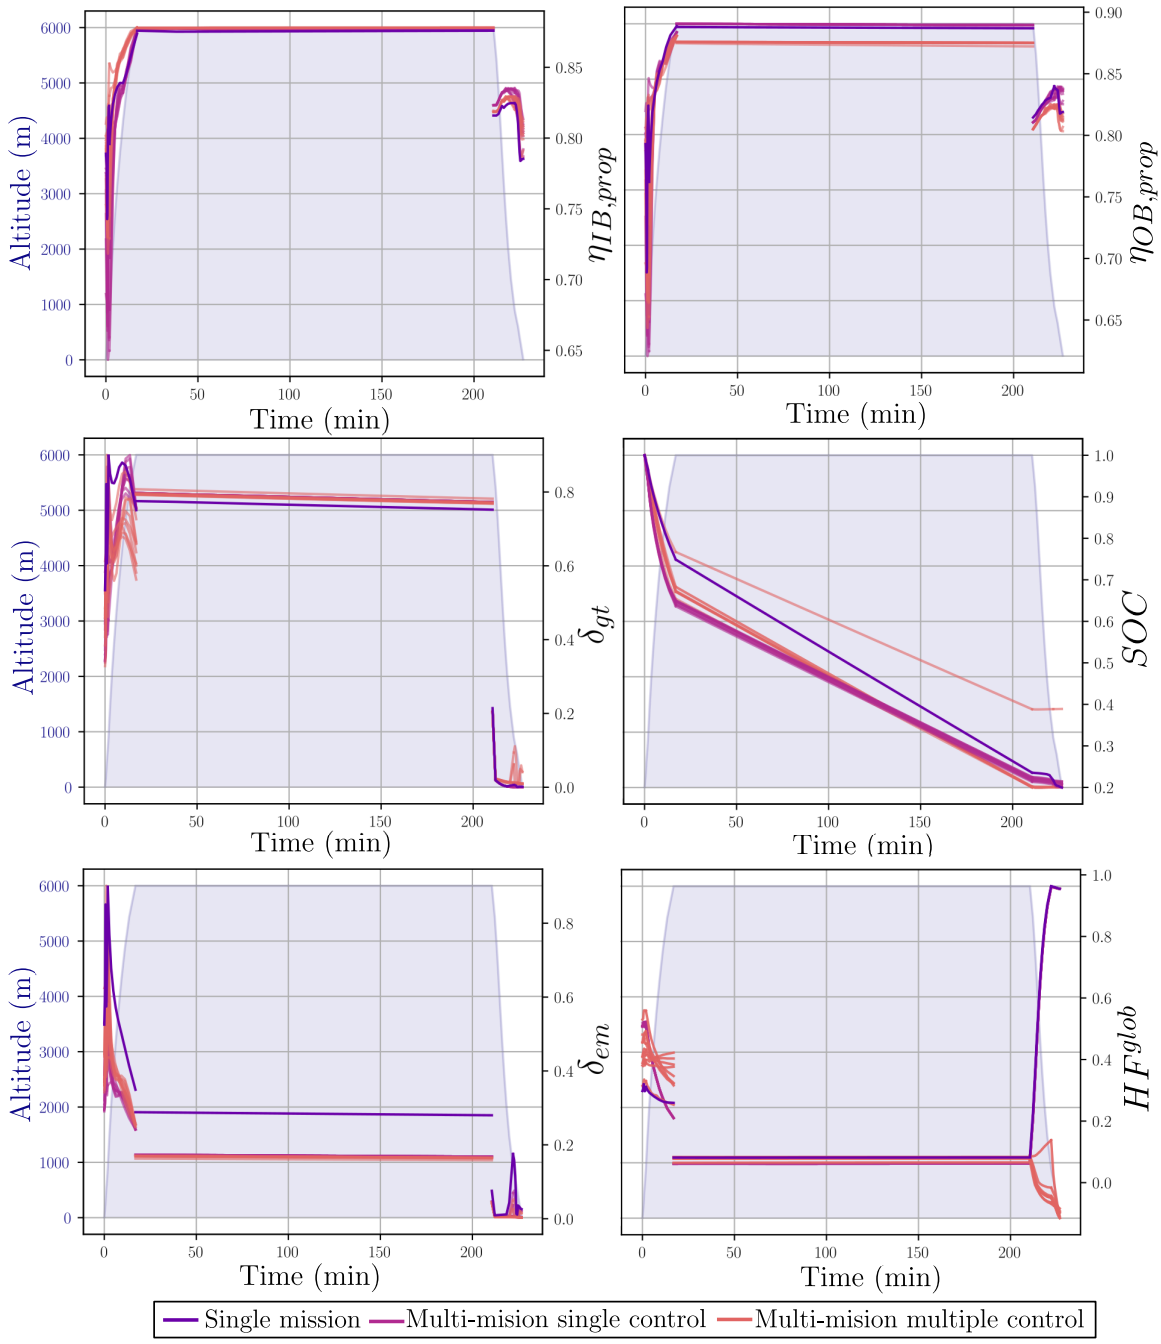


Figure 14: Aircraft performance variables during nominal mission for the three optimization cases.

Regarding powertrain sizing, Fig.14 (middle) shows that both electric motors and turboshafts are sized to meet the demands of lower-altitude climb procedures, with the turboshaft subsequently operating at approximately $\delta_{gt} \approx 78\%$ during the cruise segment. The battery usage pattern, depicted in Fig.14 (down), illustrates the evolution of the state of charge (SOC), showing similar discharge profiles across all missions.[4] To quantify the hybrid power distribution, a global hybridization factor is defined as the ratio of battery

[4]With the exception of one mission-specific control case, which appears to be an outlier resulting from optimizer convergence issues, as evidenced by the absence of battery discharge.

power to the total energy provided by both the battery and turboshafts:

$$HF^{global} = \frac{P_{battery}}{P_{battery} + 2P_{gt}} \tag{7}$$

This metric reveals that batteries primarily serve as power-boosting devices, with hybridization factors upwards of 57% during climb, followed by reduced cruise hybridization to approximately 10%.

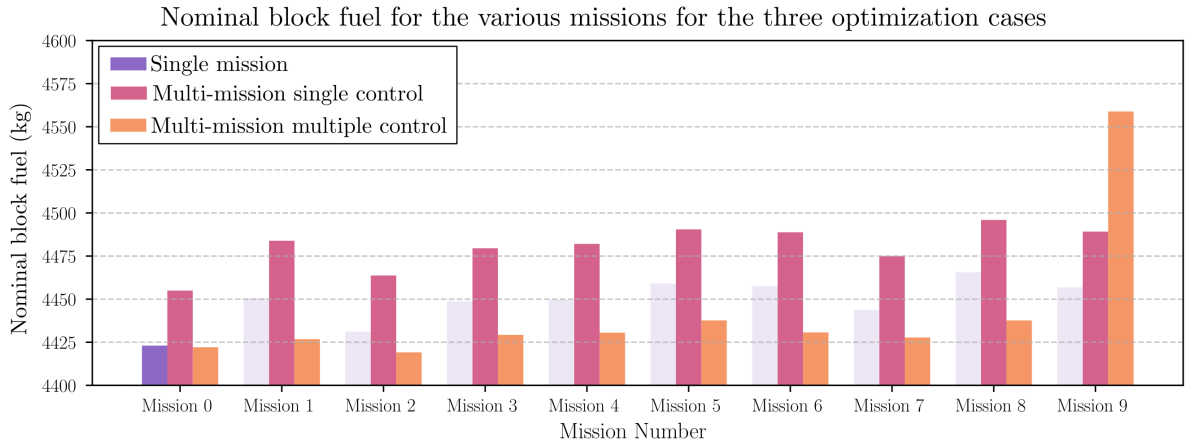


Figure 15: Block fuel for nominal missions across the three optimization cases. Lower opacity bars for the single mission indicate MDA results using optimized design variables.

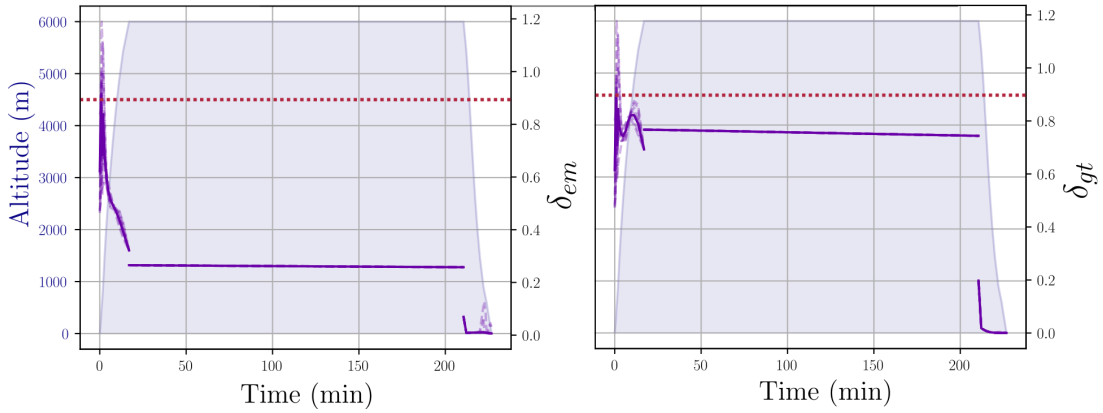


Figure 16: Throttles of the gas turbine and electric motors for all the missions using the aircraft and HF^*, ψ control profiles from Case 1.

Finally, considering the effects of multi-mission and mission-specific control, Fig.15 presents the block fuel consumption for the nominal mission across the three optimization cases. As observed, the multi-mission analysis with mission-specific control consistently outperforms the multi-mission approach with a monolithic control strategy, except for the outlier case with undischarged batteries previously identified in Fig.14. This result is expected, as mission-specific controls enable more precise energy management tailored to the constraints of each mission phase. For the single-mission case, an MDA was performed to assess performance across other missions (represented by bars with lower opacity in Fig.14). While the initial analysis may suggest that the single-mission approach

outperforms the multi-mission monolithic control, a detailed examination of powerplant components during the mission reveals that the powerplant operates beyond the designed power limits (Fig.16), underscoring the importance of multi-mission analysis for achieving reliable designs.

	Case 1	**Case 2**	**Case 3**
$\mu_{\textbf{block fuel}}$ **(kg)**	4448.57	4480.35	4429.07
$\sigma_{\textbf{block fuel}}$ **(kg)**	13.03	12.74	6.19

Table 3: Mean and sample standard deviation of nominal block fuel missions.

Examining the performance statistics across the various optimized cases, Tab. 3 presents the mean and standard deviation of each optimization case, excluding the outlier from Case 3. As expected, the multi-mission analysis with mission-specific control outperforms the monolithic control approach in terms of both mean block fuel and deviation, achieving a 1.14% reduction in mean block fuel. While this fuel reduction can be significant over the long term, additional operational uncertainties may diminish these savings, potentially offsetting the benefits of the increased operational complexity.

6 CONCLUSIONS

This paper has shown the application of the in-house MDO software `TOPAZ` to optimize a high aspect ratio, strut-braced aircraft with hybrid-electric distributed propulsion, considering various mission and sizing scenarios. The findings indicate that multi-mission analysis is essential to achieve a robust design capable of addressing uncertainties associated with low-altitude procedures. While incorporating mission-specific control results in a design with lower mean block fuel, the improvements are minimal. The mission-specific controls provide slight adjustments for each mission, resulting in a mean block fuel reduction of 1.14%.

7 ACKNOWLEDGEMENTS

The activities described in this paper have been carried out under the project INDIGO (Integration and Digital Demonstration of Low-emission Aircraft Technologies and Airport Operations). INDIGO project [4] has received funding from the European Climate, Infrastructure and Environment Executive Agency (CINEA) under the HorizonEurope programme under grant agreement No 10109605.

REFERENCES

[1] International Energy Agency. Aviation – Energy System, Transport, 2025. URL `https://www.iea.org/energy-system/transport/aviation`. Last updated on 16 January 2025.

[2] European Commission. Reducing emissions from aviation, 2025. URL `https://climate.ec.europa.eu/eu-action/transport/reducing-emissions-aviation_en`. Accessed on 2025-02-21.

[3] International Civil Aviation Organization. Trends in Emissions that affect Climate Change, 2025. URL https://www.icao.int/environmental-protection/Pages/ClimateChange_Trends.aspx. Accessed on 2025-02-21.

[4] Innovative aircraft and propulsion technologies to improve air quality near airports. URL https://indigo-sustainableaviation.eu/.

[5] F. Gallard, C. Vanaret, D. Guénot, V. Gachelin, R. Lafage, B. Pauwels, P.-J. Barjhoux, and A. Gazaix. Gems: A python library for automation of multidisciplinary design optimization process generation. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2018.

[6] B. J. Brelje and J. R. R. A. Martins. Development of a conceptual design model for aircraft electric propulsion with efficient gradients. In *Proceedings of the AIAA/IEEE Electric Aircraft Technologies Symposium*, Cincinnati, OH, July 2018. doi:10.2514/6.2018-4979.

[7] C. David, S. Delbecq, S. Defoort, P. Schmollgruber, E. Benard, and V. Pommier-Budinger. From fast to fast-oad: An open source framework for rapid overall aircraft design. *IOP Conference Series: Materials Science and Engineering*, 1024(1):012062, jan 2021. doi:10.1088/1757-899X/1024/1/012062. URL https://dx.doi.org/10.1088/1757-899X/1024/1/012062.

[8] R. Quiben Figueroa, R. Cavallaro, and A. Cini. Feasibility studies on regional aircraft retrofitted with hybrid-electric powertrains. *Aerospace Science and Technology*, 151: 109246, Aug. 2024. ISSN 1270-9638. doi:10.1016/j.ast.2024.109246.

[9] J. Gratz, J. Jasa, E. Aretskin-Hariton, K. Moore, K. Marfatia, J. Kirk, and C. Recine. Aviary: An Open-Source Multidisciplinary Design, Analysis, and Optimization Tool for Modeling Aircraft with Analytic Gradients. In *AIAA Aviation Forum*, Las Vegas, NV, United States, July 2024. American Institute of Aeronautics and Astronautics. URL https://ntrs.nasa.gov/citations/20240007608. NASA Document ID: 20240007608. Public Use Permitted.

[10] P. N. Simon, R. Q. Figueroa, A. Cini, and R. Cavallaro. *Multidisciplinary Design and Optimization of H2-Powered Regional Aircraft Architectures*. doi:10.2514/6.2025-0360. URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-0360.

[11] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/jax-ml/jax.

[12] Airbus S.A.S. *A320 Aircraft Characteristics – Airport and Maintenance Planning*. Airbus Customer Services, Technical Data Support and Services, 31707 Blagnac Cedex, France, revision no. 44 edition, June 2024. Originally issued Sep 30, 1985. Proprietary document intended for planning purposes.

[13] R. A. McDonald and J. R. Gloudemans. *Open Vehicle Sketch Pad: An Open Source Parametric Geometry and Analysis Tool for Conceptual Aircraft Design*. doi:10.2514/6.2022-0004. URL https://arc.aiaa.org/doi/abs/10.2514/6.2022-0004.

[14] R. D. Finck. USAF Stability and Control DATCOM. Technical Report AF33(616)-6460, F33615-76-C-3061, McDonnell Douglas Corporation, Douglas Aircraft Division, Wright-Patterson Air Force Base, Ohio, April 1978. Originally published October 1960. Revised April 1978. Project No. 8219, Task No. 821901.

[15] D. P. Raymer. *Aircraft design: A conceptual approach*. AIAA education series. American Institute of Aeronautics and Astronautics, Reston, Va., 2006. ISBN 1563478293. URL http://www.loc.gov/catdir/toc/ecip068/2006004706.html.

[16] D. Howe. *Basic Lift, Drag and Mass Representations*, chapter 6, pages 139–164. John Wiley and Sons, Ltd, 2000. ISBN 9781118903094. doi:https://doi.org/10.1002/9781118903094.ch6. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9781118903094.ch6.

[17] M. H. Sadraey. *Aircraft design: A systems engineering approach*. Aerospace Series. John Wiley and Sons, Chichester, 2013. ISBN 9781118352700. doi:10.1002/9781118352700. URL http://dx.doi.org/10.1002/9781118352700.

[18] E. Torenbeek. *Synthesis of Subsonic Airplane Design*. Delft University Press, Delft, 1988.

[19] E. S. Hendricks and J. S. Gray. pycycle: A tool for efficient optimization of gas turbine engine cycles. *Aerospace*, 6(87), August 2019. doi:10.3390/aerospace6080087.

[20] M. Drela and H. Youngren. XROTOR: An interactive program for the design and analysis of ducted and free-tip propellers and windmills. https://web.mit.edu/drela/Public/web/xrotor/, 2022. URL https://web.mit.edu/drela/Public/web/xrotor/. Version 7.69, Accessed: 2025-03-31.

[21] K. F. P. N. C. Jones. Safety-driven baseline of hybrid electric aircraft electrical power system architectures. In *ICAS Proceedings: 34th Congress of the International Council of the Aeronautical Sciences*, number $\mathrm{ICAS2024}_1027, Florence, Italy, October 2024. ICAS$.

[22] H. Kühnelt, F. Mastropierro, N. Zhang, S. Toghyani, and U. Krewer. Are batteries fit for hybrid-electric regional aircraft? *Journal of Physics: Conference Series*, 2526(1):012026, jun 2023. doi:10.1088/1742-6596/2526/1/012026. URL https://dx.doi.org/10.1088/1742-6596/2526/1/012026.

[23] E. Obert. *Aerodynamic design of transport aircraft*. IOS Press, Netherlands, 2009. ISBN 978-1-58603-970-7.

[24] L. Attanasio, A. De Marco, F. Nicolosi, and P. Vecchia. Development of a java-based framework for aircraft preliminary design and optimization. 06 2015. doi:10.2514/6.2015-2651.

[25] G. P. Chiozzotto. Conceptual design method for the wing weight estimation of strut-braced wing aircraft. In *5th CEAS Air & Space Conference*, Delft, The Netherlands, September 2015. URL https://elib.dlr.de/98248/. Conference presentation; full text not available from repository.

[26] *Next Generation Turboprop Gearboxes*, volume Volume 2: Aircraft Engine; Marine; Microturbines and Small Turbomachinery of *Turbo Expo*, 04 1982. doi:10.1115/82-GT-236. URL https://doi.org/10.1115/82-GT-236.

[27] C. L. Pastra, C. Hall, G. Cinar, J. Gladin, and D. N. Mavris. Specific power and efficiency projections of electric machines and circuit protection exploration for aircraft applications. In *2022 IEEE Transportation Electrification Conference & Expo (ITEC)*, pages 766–771, 2022. doi:10.1109/ITEC53557.2022.9813927.

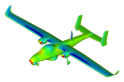
[28] C. Hall, C. L. Pastra, A. Burrell, J. Gladin, and D. N. Mavris. Projecting power converter specific power through 2050 for aerospace applications. In *2022 IEEE Transportation Electrification Conference & Expo (ITEC)*, pages 760–765, 2022. doi:10.1109/ITEC53557.2022.9813991.

[29] H. Berg, P. Macedo, B. Picinatti, M. Adler, S. Pel, and K. Jacobs. Review of standard passenger weights: Final report. Technical report, Lufthansa Consulting GmbH and SEO Amsterdam Economics, Frankfurt/M., Germany, September 2022. Report submitted to the European Union Aviation Safety Agency (EASA) under Contract Reference EASA.2021.C24.

[30] D. Zheng Zhang, D. Gómez López, J. A. López Sánchez, C. Xia, X. Wang, and A. García Fernández. Local air quality and noise assessment for landing and take-off operations in future airport environment. *Engineering Proceedings*, 90(1), 2025. ISSN 2673-4591. doi:10.3390/engproc2025090074. URL https://www.mdpi.com/2673-4591/90/1/74.

[31] P. Norczyk and R. Cavallaro. *Local Air Quality and Noise Improvements via Optimization of Strut-Braced Wings with Distributed Electric Propulsion*. doi:10.2514/6.2025-0456. URL https://arc.aiaa.org/doi/abs/10.2514/6.2025-0456.

[32] P. E. Gill, W. Murray, and M. A. Saunders. Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12(4):979–1006, 2002. doi:10.1137/S1052623499350013. URL https://doi.org/10.1137/S1052623499350013.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

GREEN AVIATION MANUFACTURING: ADDRESSING ENVIRONMENTAL IMPACTS WITH MDO METHODOLOGIES

Shantanu Sapre[1*], Joseph Morlier[2] and Christian Gogu[2]

1: ISAE-SUPAERO
Université de Toulouse
Toulouse, France
shantanu-rajendra.sapre3@isae-supaero.fr

2: ICA, ISAE-SUPAERO, MINES ALBI, UPS, INSA, CNRS
Université de Toulouse
Toulouse, France
{joseph.morlier,christian.gogu}@isae-supaero.fr

Abstract. *The aeronautical industry faces increasing pressure to reduce its environmental footprint while meeting the growing demand for aircraft. Although the operational phase remains the most significant contributor to environmental impact, the manufacturing phase may gain prominence as aviation transitions to alternative fuels and innovative designs. This study addresses the need for sustainability in aircraft production by integrating environmental considerations into key stages of the aircraft lifecycle, particularly during preliminary design. The study evaluates aircraft mass and material composition using a bottom-up approach, considering buy-to-fly ratios, greenhouse gas emissions, energy consumption, and waste generation of the legacy aircraft. To mitigate environmental impact, this research explores a multi-objective optimization technique for selecting environmentally friendly material composition while maintaining performance standards. A key focus is the role of multidisciplinary design optimization (MDO) in balancing sustainability with functional and economic requirements during the overall aircraft design. Ultimately, this study highlights the need for introducing material composition selection as a new parameter within the classical MDAO framework.*

Keywords: Environmental impact, Material selection, Eco-design, Multiobjective Optimization

1 INTRODUCTION

While the need for air travel is beginning to be debated, it remains the essential mode of transport for changing continents, and finding an alternative would be hard to imagine in today's interconnected world. The aviation industry is a major contributor to global greenhouse gas (GHG) emissions, accounting for approximately 2-3% of global CO2 emissions and 12% of transport-related emissions [1]. The lifecycle assessment approach helps in quantifying the environmental impacts of various phases in the product's life, however, very few studies have assessed the aeronautics sector's impacts in the manufacturing phase. Amongst the studies that were reviewed, it was found that the aeronautical manufacturing industry accounts for 1-3% of aviation sector emissions, making it a tiny part of the world's greenhouse gas emissions [1, 2]. Today, the focus of aviation industry is towards reducing the emissions from the operational phase with methods such as introduction of Sustainable Aviation Fuels (SAF), derived from bio-based and synthetic sources, which is expected to reduce the carbon dioxide (CO_2) emissions by up to 80% [3]. The industry is also exploring the possibilities of breakthrough innovations such as new configurations in aircraft design, hybrid-electric propulsion system, advanced flight control and energy management, and new generation of engines which claim to reduce the tailpipe emissions between the range of 10% - 50% [4]. Given these improvements in the technology, the environmental impacts of the manufacturing phase may increase in future. The prospective assessment done by Thonemann et al [2] on ecosystem quality show a rise in the contribution of manufacturing phase from less than 1% for conventional aircraft to 8% or even higher in case of hybrid electric aircraft.

Current research in the manufacturing field to reduce the environmental impact is focused on additive manufacturing - for reducing wastage, energy consupmtion and making topology optimized parts [5], bio-based material development and raw material circularity. Although these methods will potentially help in transitioning towards green manufacturing, their current technology readiness level (TRL) lies between 4 and 7 [6]. Being expensive than the conventional materials and manufacturing techniques, it may also lead to significant increase in the upfront cost of the product. Some key materials, particularly structural metals remain irreplaceable due to their strength, fatigue resistance, and certification history. It would therefore be interesting to find and evaluate ways of reducing the impact of the aeronautical manufacturing sector considering conventional materials and optimizing their selection, combination, and hybridization with sustainable alternatives. This study attempts to make a general assessment of the environmental impact of the aeronautical manufacturing sector, focusing on the two dominant civil aircraft manufacturers - Airbus and Boeing having a combined market share of about 89% in the civil aviation industry. This combined market share is projected to have same magnitude for years to come, thus making these two aerospace giants the biggest contributors (directly or indirectly) to the environmental impact of the civil aviation industry. In the first part of the study, a bottom up approach as followed by Pierrat et al. [7] was used to estimate the energy consumption, CO_2 emissions and waste generation from the commercial aircraft in production. Having identified the common patterns in the category of aircraft based on its material composition, the next phase consisted of developing a methodology for formulating a multiobjective optimization problem for material selection, and the results were compared with the existing material composition of aircraft.

2 METHODOLOGY

2.1 Environmental Impact Assessment

The overall assessment focuses on the "cradle to gate" level - thus the system boundary of this study is limited from the mineral extraction stage to the assembly stage. Figure 1 is a simplified flowchart derrived from Pierrat et al. [7] methodology which illustrates the key stages in the lifecycle of an aircraft.

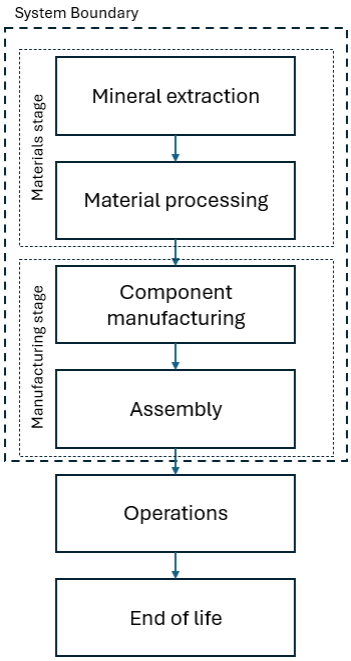


Figure 1: Key stages in aircraft lifecycle

The materials used in an airframe can be grouped into 4 main categories of materials: Aluminium, Titanium, Steel and Carbon Fibre Reinforced Polymer (CFRP). Data from the aircraft manufacturers as well as references suggest that more than 90% of the airframe comprises of theses 4 material categories and these materials are matured in the aerospace industry for their mechanical and thermal properties[7–9].

The environmental impact of materials and manufacturing stage was assessed on 3 axes - Energy consumption, CO_2 emissions and waste generation. The bottom-up approach focuses on estimating annual demand for materials on the basis of the material composition of aircraft, average environmental impact in producing unit quantity of material, average buy-to-fly ratio associated with the material and number of aircraft delivered in a year. The estimation of waste generation is linked with material losses during manufacturing considering the buy-to-fly (BtF) ratios - thus creating a scope for assessing recyclable manufacturing waste. The energy required for production of material and the associated CO_2 emissions depend on the quality of the ore, process of production and the region where it is produced. Since accurate information on the material suppliers and their energy mix remain confidential to the manufacturers, data available from the literature as well as LCA tools was compiled and the average value for impact per unit quantity of material produced was considered for the assessment [10]. The input data is summarized in Tab.1.

It also includes the Density and Cost values which have been used later in the assessment.

Table 1: Material Properties

Property	Unit	Aluminium	Steel	Titanium	CFRP
BtF	kg/kg	5	6	10	1.5
$Density$	kg/m^3	2780	7800	4430	1600
$Energy$	MJ/kg	163	20	536	514
CO_2	kg/kg	14.9	1.9	20	28
$Cost$	$\$/kg$	2.5	2.4	21	90

The buy-to-fly (BtF) ratio is a crucial parameter in manufacturing, and it directly affects cost efficiency and environmental sustainability. The manufacturing technique employed significantly influences the BtF ratio. Metals often undergo subtractive manufacturing, which involves removing material from a solid block to achieve the desired shape, resulting in high BtF ratios. Formative manufacturing techniques, such as molding and forging, typically result in lower BtF ratios because they shape materials into near-final forms with minimal excess waste. Thus, CFRP which undergoes formative manufacturing has a lower BtF ratio as compared to the metals. The carbon dioxide (CO_2) emissions associated with material production are closely linked to the energy required for the process. Materials that undergo energy-intensive extraction and processing methods tend to have higher CO_2 emissions, following a similar trend as energy consumption.Material composition and empty weights of aircraft produced by Airbus and Boeing were obtained from the grey literature as well as references available in public domain [7, 11–13]. This input data is listed in Tab.2.

Table 2: Material Composition and empty weights of selected aircraft

Aircraft	**OWE (t)**	**Aluminium**	**Steel**	**Titanium**	**CFRP**
A300	86.2	67%	12%	4%	9%
A320	40	72%	7%	6%	10%
A330	124	58%	19%	8%	9%
A380	277	61%	4%	6%	22%
A350	135	19%	7%	14%	53%
B737	41.7	80%	8%	8%	3%
B747	184.6	81%	13%	4%	1%
B767	84.5	78%	13%	2%	5%
B777	168.7	70%	11%	7%	12%
B787	128.8	20%	10%	15%	50%

To estimate the Energy requirement and CO_2 emissions, Eq.(1) was used where $W_{aircraft}$ is the empty weight of aircraft, e_m is the environmental impact (CO_2 or Energy) per unit material produced, w_m is the weight fraction of material in aircraft and BtF_m is the average BtF ratio of the material.

$$\text{Environmental Impact for one aircraft} = W_{\text{aircraft}} \times \sum (e_m \times w_m \times BtF_m) \quad (1)$$

For calculating the waste generated from manufacturing (scrap material), Eq.(2) was utilized. This equation determines the waste as the difference between the total material procured and the material that ultimately remains in the final aircraft structure. This method provides a clear assessment of material efficiency, highlighting the portion of metal that can be recycled and the amount of CFRP waste that is likely to end up in landfills. Understanding these waste streams is essential for developing sustainable recycling strategies and minimizing environmental impact.

$$\text{Manufacturing Waste} = W_{\text{aircraft}} \times \sum (w_m \times (BtF_m - 1)) \quad (2)$$

It is also important to note that not all metal scrap generated during manufacturing is suitable for recycling. Some of it is lost as tailings or ends up in landfills due to contamination or processing limitations. The fraction of non-recyclable metal waste varies depending on the specific manufacturing process. For aluminum and steel, approximately 5% [14] of the scrap was considered non-recyclable, for titanium, this value was higher at 15% [15], and for CFRP, the entire 100% of the waste was assumed to be non-recyclable due to the lack of effective recycling methods [16]. Thus, out of this total waste generated, the non-recyclable waste can be calculated from Eq.(3), where the term 'nrw_m' represents the fraction of scrap that is non-recoverable.

$$\text{Non-recyclable Waste} = W_{\text{aircraft}} \times \sum (w_m \times (BtF_m - 1) \times nrw_m) \quad (3)$$

Finally, these indicators were multiplied by the number of deliveries of the respective aircraft in production to estimate the total impact in one year and to identify the aircraft programs which needs attention from the manufacturing environmental impact perspective.

2.2 Optimization of material composition for reduced environmental impact

The selection of materials for structural components is conventionally performed using the Ashby method, which prioritizes material performance factors such as strength, fatigue life, and durability while simultaneously minimizing weight and cost [17]. This inherently presents a multi-objective optimization problem. When environmental considerations, such as energy consumption and waste generation, are incorporated, the complexity of the optimization further increases.

In this study, rather than selecting a single optimal material, we address a higher-level problem concerning the overall material composition of an aircraft structure. To achieve this, we conceptualize the material system as an "effective material," wherein the composition is defined by the weight fractions of its constituent materials. The optimization process involves adjusting these weight fractions to minimize key objectives, including effective density (weight), cost, energy required for production, and the amount of non-recyclable waste generated during manufacturing. By formulating the material composition problem in this manner, we apply multi-objective optimization techniques to

identify a set of Pareto-optimal material compositions. These optimal solutions represent trade-offs among the competing objectives, enabling a systematic comparison with the existing material compositions currently employed in aircraft structures.

Let the weight fractions of the constituent materials be:

$$w_1, w_2, w_3, w_4$$

and the corresponding volume fractions be:

$$x_1, x_2, x_3, x_4$$

The volume fraction of each material is determined using the Eq.(4). Here the numbers 1 - 4 represent the individual constituent materials:

$$x_i = \frac{w_i/\rho_i}{\sum_{j=1}^{4}(w_j/\rho_j)}, \quad i = 1, 2, 3, 4 \tag{4}$$

The densities of the individual materials are known hence, the effective density of the effective material can be determined using Eq.(5):

$$\rho_{\text{eff}} = x_1\rho_1 + x_2\rho_2 + x_3\rho_3 + x_4\rho_4 \tag{5}$$

After establishing the fundamental relationships between the variables, we formulate this as a discrete optimization problem, where the final outcome is the optimal material composition derived from pure materials. The objectives to be minimized include cost, energy consumption during material production, non-recyclable waste generation, and the resulting composite density, which directly correlates with weight minimization. The optimization is subject to the following constraints:

- The sum of individual weight fractions must equal 1.
- The quantity of titanium in the composition must be at least 1/4 of the CFRP content, considering CFRP's mechanical and chemical dependence on titanium [15]. This relationship is based on observed material compositions in aircraft, particularly the trend where increased CFRP usage leads to higher Titanium content.
- Lower bounds for individual materials must be respected, defined as the minimum values found in the material compositions of the selected aircraft from the reference Tab.2.

The energy consumption and cost are calculated by multiplying the unit cost of materials by their respective weight fractions and the BtF ratio. Similarly, non-recyclable waste generation is computed by multiplying the total waste generated by the percentage of material that cannot be recycled, as previously discussed. Thus, the mathematical representation of this optimization framework is given by Eq.(6).

$$\begin{aligned}
&\min_{\mathbf{w}} \quad \{\text{Cost}, \text{Energy}, \text{Waste}, \rho_{\text{eff}}\} \\
&\text{subject to:} \\
&\sum_{i=1}^{4} w_i = 1 \\
&w_1 > 0.15, \quad w_2 > 0.01, \quad w_3 > 0.01, \quad w_4 > 0.01 \\
&w_3 > 0.25 \times w_4 \\
&\text{Energy}_m = w_m \times e_m \times BtF_m \qquad (6a) \\
&\text{Cost}_m = w_m \times c_m \times BtF_m \qquad (6b) \\
&\text{Waste}_m = w_m \times (BtF_m - 1) \times nrw_m, \quad m = 1, \ldots, 4 \qquad (6c) \\
&\text{Total impact} = \sum_{i=1}^{4} Parameter_i \quad Parameter = \{Energy, Cost, Waste\} \\
&\rho_{\text{eff}} = \sum_{i=1}^{4} x_i \rho_i \qquad (6d)
\end{aligned}$$

The optimization code was developed using the Pymoo Python package [18]. The eXtended Design Structure Matrix (XDSM) diagram, illustrating the logical workflow, is shown in Fig.2. The process begins with discrete property data for four materials provided as input. During initialization, the weight fractions of the first three materials are assigned, while the fraction of the fourth material is determined by subtracting the sum of the first three from 1, to ensure that the total composition sums to 100%. To solve this multi-objective optimization problem, the NSGA-II (Non-Dominated Sorting Genetic Algorithm II) was selected due to its efficiency in handling conflicting objectives that are chosen while maintaining diversity in solutions. It maintains the diverse set of solutions which would help in plotting wide range of material composition outputs. Using the formulae for calculating energy consumption (Eq.(6a)), cost(Eq.(6b)), waste generation(Eq.(6c)), and density(Eq.(6d)), NSGA-II iteratively evolves the population of material compositions. It ultimately produces a set of optimal solutions plotted on Pareto fronts which enable users to analyze trade-offs and select the most suitable material composition based on the priorities whether minimizing cost, environmental impact or balancing multiple factors.

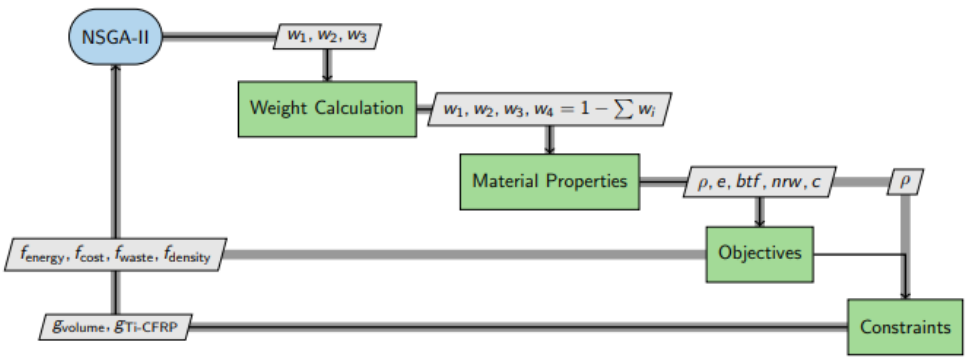


Figure 2: XDSM of optimization for finding pareto optimal material compositions

3 RESULTS

3.1 Environmental Impact

The energy consumption and CO_2 emissions during the materials stage were plotted as functions of the aircraft's empty weight Fig.3 and Fig.4 respectively. A linear trend is observed for both environmental indicators, suggesting that the impact is directly proportional to the aircraft's weight. Notably, CFRP-intensive aircraft, such as the A350 and B787, lie above the trendline, indicating that their material compositions have a greater energy consumption compared to similar aircraft made primarily of Aluminum.

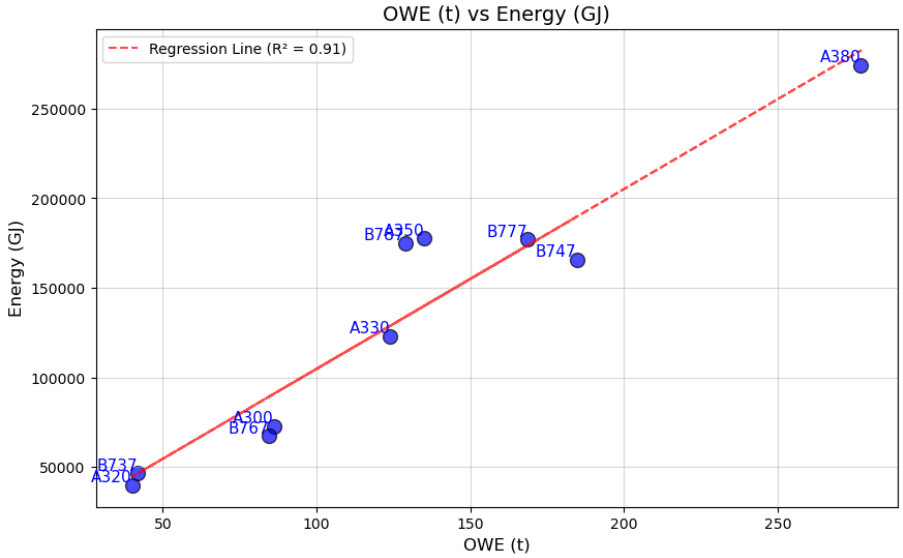


Figure 3: Aircraft weight vs energy in material stage. Energy required for producing materials increases linearly with the mass of aircraft to be produced

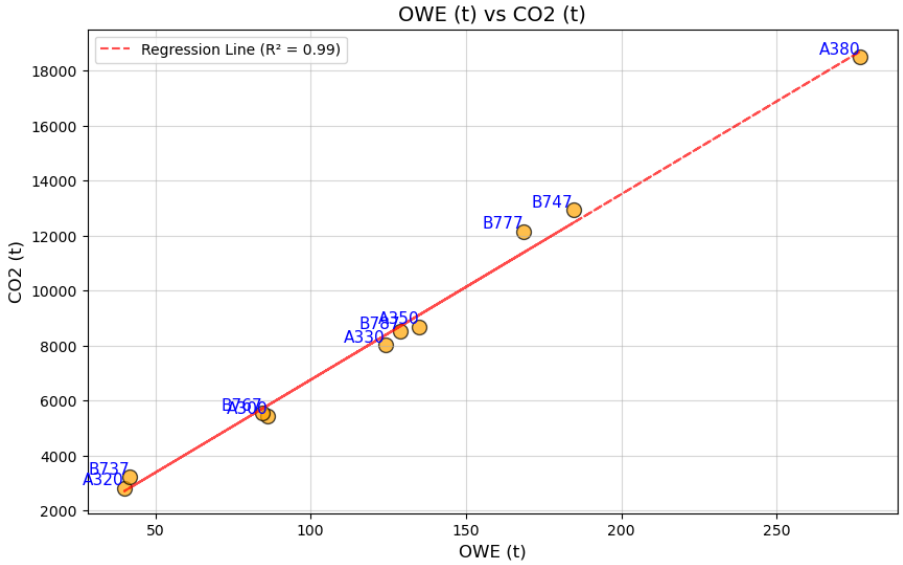


Figure 4: Aircraft weight vs CO2 emissions in material stage. Linear trend, similar as that of energy consumption is observed signifying the relation between energy and the CO2 emissions

The increased use of CFRP also necessitates a higher proportion of Titanium for design-related reasons. Since both CFRP and Titanium are energy-intensive materials, their extensive use results in a higher overall energy consumption. Additionally, as CO_2 emissions are closely linked to the energy consumed during processing, both plotted trends exhibit similar patterns. This correlation suggests that reducing energy consumption in the manufacturing process would lead to a decrease in CO_2 emissions. Due to this reason, only Energy parameter was chosen as one of the objectives for the optimization in the later stage.

Table 3 presents the approximate amount of manufacturing waste (in tons) generated per material per aircraft. Metals such as Aluminum, steel, and Titanium produce scrap during manufacturing, but a significant portion of this waste can be recycled if properly segregated and reintegrated into the supply chain. This aligns with raw material circularity principles from an ecodesign perspective, reducing the overall environmental impact of metal-based waste. However, CFRP waste, especially from thermoset CFRP poses a far greater challenge. Unlike metals, recycling technologies for thermoset CFRP are not yet mature, meaning that most of this waste currently ends up in landfills. This is particularly concerning for modern aircraft with high CFRP content, such as the A350 and B787, as it significantly increases non-recyclable waste. Due to this issue, "Non-recyclable waste" was chosen as a key parameter to be minimized in the optimization process.

Table 3: Manufacturing waste generated in tons per material per aircraft

Aircraft	Aluminium	Steel	Titanium	CFRP
A300	231	52	31	4
A320	115	18	22	2
A330	288	118	89	6
A380	676	55	150	30
A350	103	47	170	36
B737	133	17	30	1
B747	598	120	66	1
B767	264	55	15	2
B777	472	93	106	10
B787	103	64	174	32

These values, derived from the assessment of individual aircraft, were multiplied by the respective deliveries in 2023 to evaluate the environmental impact over one year and identify areas for improvement. The results for energy consumption are presented in Fig.5. The numbers in brackets below each aircraft name represent the total deliveries for the year 2023 [19]. Single-aisle aircraft, such as the A320 and B737, consume more energy than other aircraft programs. However, despite having fewer than one-eighth the deliveries of the A320, the CFRP intensive A350 requires nearly half the total energy needed for A320 material production. This highlights the significant influence of material selection, even when scaled by varying production volumes.

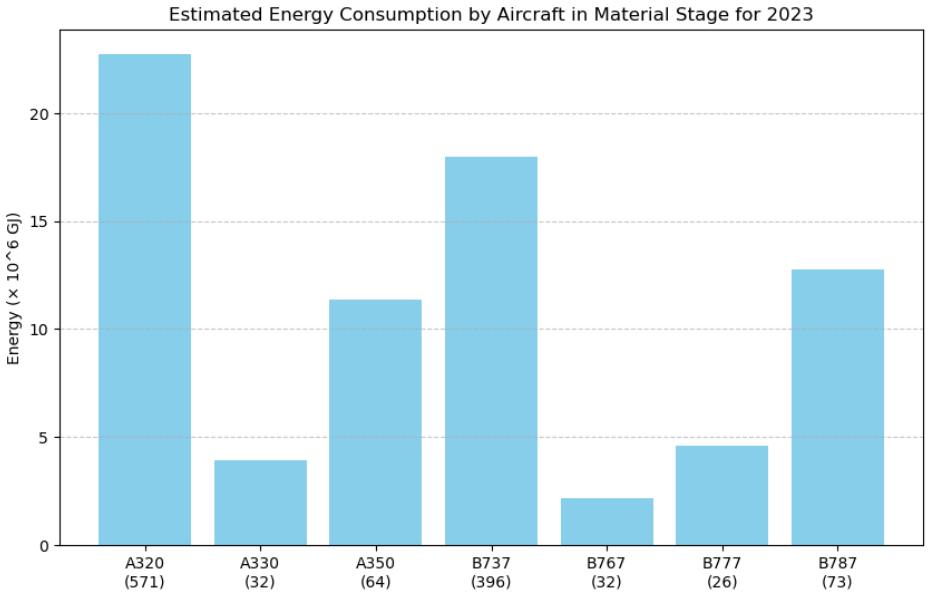


Figure 5: Energy consumption by aircraft in material stage in the year 2023. Single aisle aircraft (A320/B737) have significantly higher energy consumption which needs which creates a scope for implementing eco-design approach

A similar trend is observed in waste generation during manufacturing. While the total waste from the A320 and B737 is substantially higher than that of other programs, the amount of non-recyclable waste, as estimated by this analysis, remains nearly the same across A320, A350, B737 and B787 aircraft. This underscores the need for a multidisciplinary optimization approach to minimize environmental impacts in the early stages of aircraft life cycle effectively.

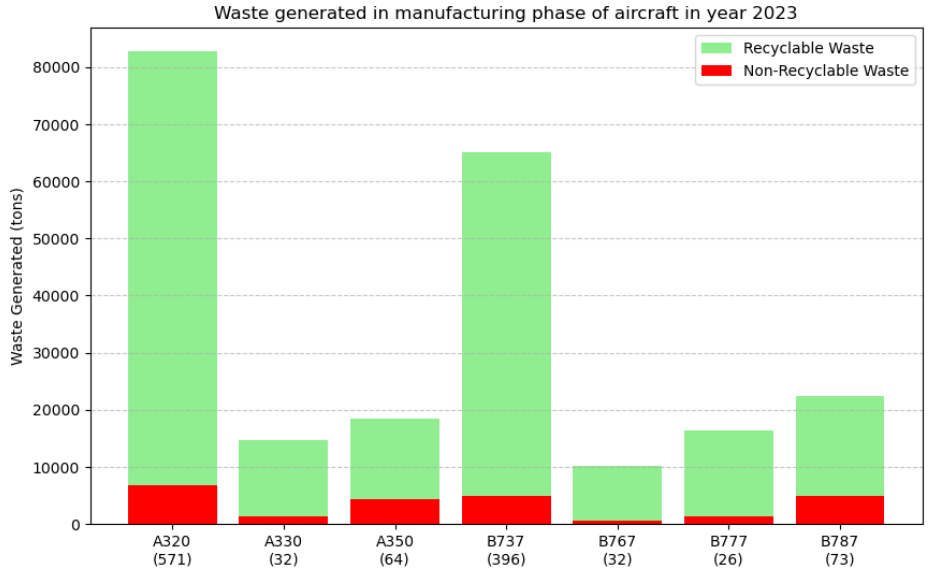


Figure 6: Waste generated in manufacturing phase of aircraft in year 2023. Despite the difference in volume of total waste generated by single aisle and new aircraft (A350/B787), the amount of net non-recyclable waste is comparable

3.2 Optimization for material composition as a design variable of a multiobjective optimization

The optimization results are represented through three Pareto front plots, each showing trade-offs between effective density which would affect the weight of aircraft and in turn performance, and a key environmental or economic objective: energy consumption in material production, non-recyclable waste generated during manufacturing, and cost of material purchased by the manufacturer. To assess how existing aircraft compare to the optimized solutions, the material compositions of legacy aircraft were input into the program, and their corresponding parameters were plotted on these graphs.

In the plot of Energy Consumption vs. Effective Density (Fig.7), the top-left region corresponds to steel-heavy compositions, which have a higher effective density but relatively low energy consumption. In contrast, the bottom-right region, dominated by CFRP-intensive compositions, exhibits significantly lower effective density but much higher energy consumption for material production. This is due to the energy-intensive nature of CFRP production and the additional titanium usage required in CFRP-based designs. The trend reinforces that while CFRP reduces weight, it comes at the cost of greater environmental impact during manufacturing. The region in the central part represents Aluminium heavy material compositions where old generation aircraft can be seen.

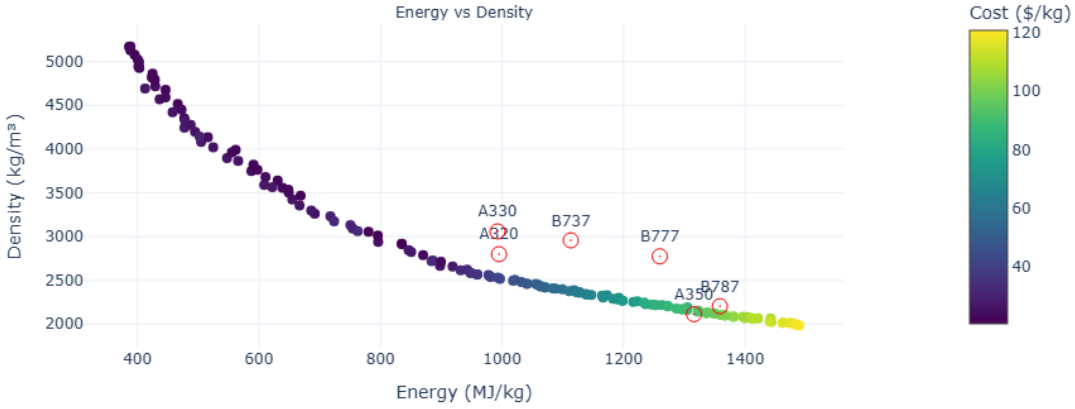


Figure 7: Energy in material production vs Effective density of material composition

The plot of non-recyclable waste in production vs. effective density (Fig.8) highlights the environmental burden associated with waste generation in manufacturing stage of aircraft. Again, steel-heavy compositions in the top-left region have higher effective density but generate mostly recyclable waste, making their environmental footprint lower in terms of landfill contribution. In contrast, CFRP-heavy configurations in the bottom-right region generate significantly more non-recyclable waste, which raises concerns about waste management in modern CFRP-intensive aircraft.

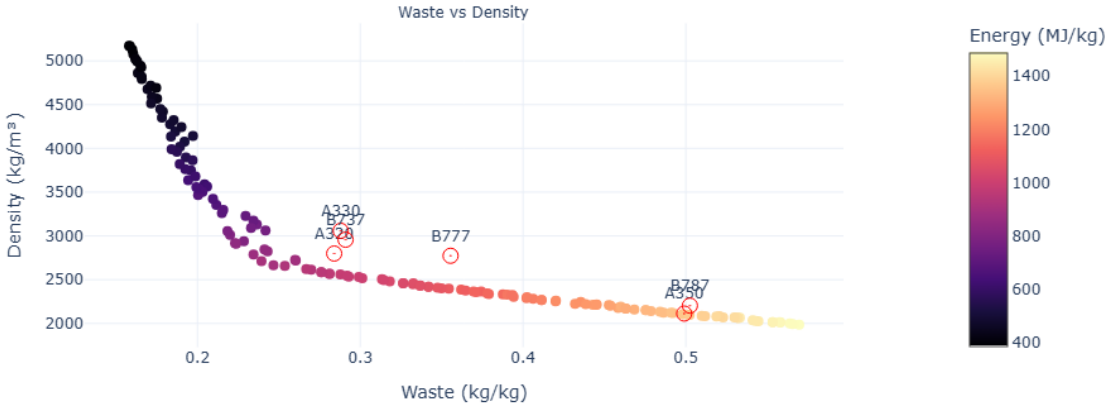


Figure 8: Non-Recyclable waste in manufacturing vs Effective density of material composition

The economic trade-offs are visualized in the Cost vs. Effective Density plot (Fig.9), where metal-heavy compositions (steel as well as Aluminium) populated on the left region exhibit lower costs due to the affordability of steel and Aluminum compared to CFRP and Titanium. Conversely, CFRP-heavy configurations in the bottom-right region are significantly more expensive due to the high production costs of CFRP and Titanium. The Pareto front indicates that achieving a low-density structure typically comes with increased material costs, making cost reduction a challenging objective in weight-efficient aircraft designs.

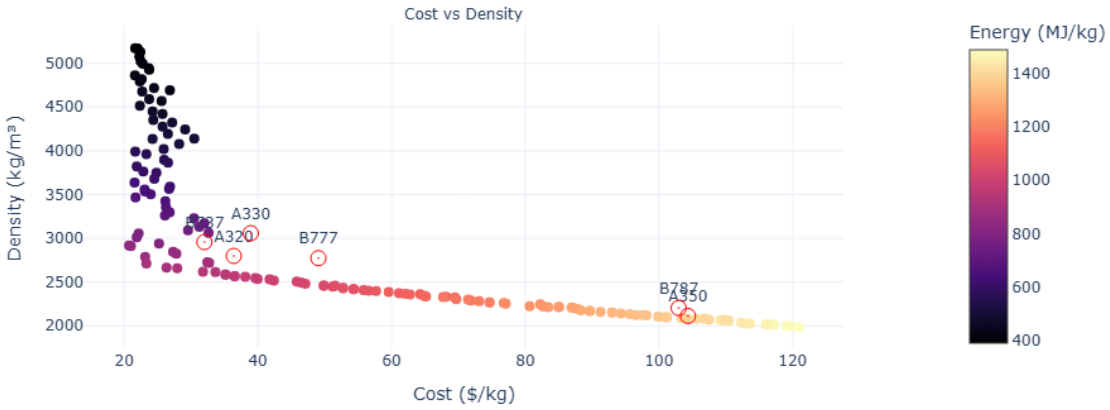


Figure 9: Material cost vs Effective density of material composition

In all the plots, it can be observed that older aircraft are positioned farther from the Pareto front, indicating that their material choices are suboptimal in terms of environmental and economic efficiency. However, modern aircraft like the A350 and B787 are observed to be much closer to the Pareto front, suggesting that their material configurations strike a better balance between weight efficiency, cost, and environmental impact.

3.2.1 Example: Selecting material composition on pareto front

As an example, we consider a scenario in which the A320 which is currenlty away from the optimal point is redesigned with an optimized material composition to reduce manufacturing environmental impact and material costs while maintaining the same effective

density and consequently, the same aircraft weight. To achieve this, a point on the Pareto front corresponding to the same effective density was selected, as shown in Fig.10. The optimized composition, detailed in Tab.4 under the column "New Composition", represents an environmentally and cost-efficient alternative. This optimized redesign would result in a 14% reduction in energy consumption during the material production stage and in non-recyclable waste generated during manufacturing, all while preserving the aircraft's weight.

Alternatively, if an optimal point just below the existing configuration were selected on the Energy Consumption vs. Effective Density plot, it would lead to a lighter material composition while maintaining the same energy consumption in production. This reduction in weight could potentially lower emissions during aircraft operations. However, this alternative comes with trade-offs, including higher non-recyclable waste generation during manufacturing and increased upfront material costs due to the greater reliance on CFRP and Titanium.

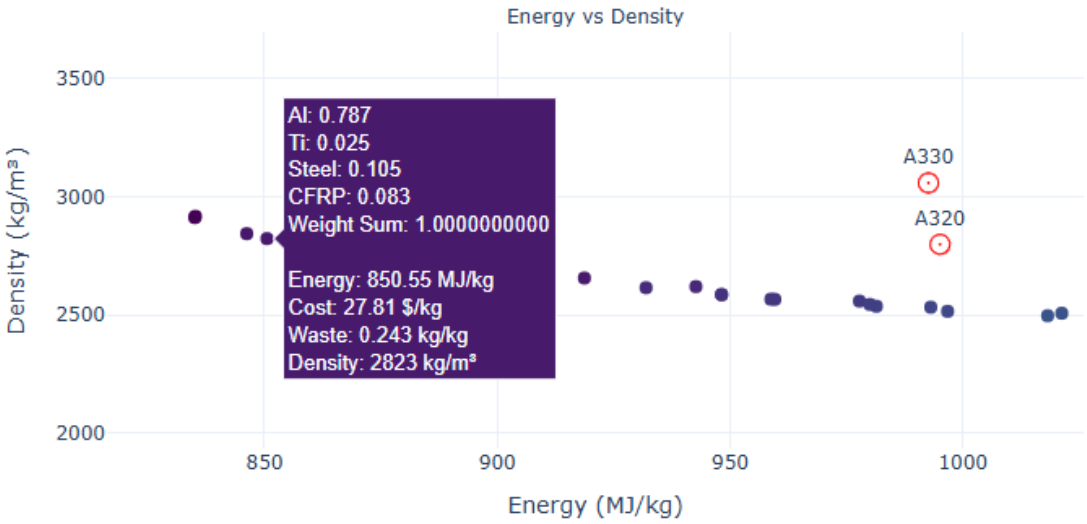


Figure 10: Example- Selecting material composition for reduced environmental impact

Table 4: Comparison- A320 existing material composition vs Pareto optimal solution

Parameter	Unit	Existing composition	New Composition	% Change
Aluminium	%	72	78.7	+9%
Titanium	%	6	2.5	-58%
Steel	%	9	10.5	+17%
CFRP	%	10	8.3	-17%
Energy	GJ	39804	34040	-14%
Waste(NR)	ton	11.4	9.73	-14%
Cost	M$	1.46	1.11	-24%

However, it is important to note that the methodology does not account for the strength properties of materials, which play a vital role in material selection. Different materials

require varying quantities to achieve the same structural performance due to differences in mechanical properties such as tensile strength, stiffness, and fatigue resistance. In future iterations of this methodology, material strength should be explicitly included to ensure that the optimized material compositions remain structurally viable. This would allow for more realistic material substitutions by considering not only weight and environmental impact but also the actual amount of material required to manufacture each component.

4 CONCLUSIONS

The aviation industry's current efforts to reduce environmental impact during the operational phase have highlighted the need to address manufacturing emissions as well. A comprehensive approach to decarbonize the entire aircraft life cycle—from material production to end-of-life disposal is essential for achieving true sustainability in the sector. This study provides a low fidelity approach in estimating the environmental impact of the aircraft manufacturing stage which can be used for other industries related to production as well. The methodology tried to set-up a ground work for material selection in the early design stage of aircraft considering environmental impact in the cradle to gate phase of the life cycle. The recently developed aircraft like A350 or B787 are close to the pareto optimal material composition, however, the single aisle aircraft like A320 or B737 which have a larger environmental impact because of their huge volumes need an optimized composition to reduce their environmental impact in materials and manufacturing phase. The current LCA approach provides an effective estimation of environmental impact, but increasing the fidelity of LCA modeling would improve accuracy. Future work should incorporate more detailed datasets, including alloy based variation in energy consumption, transportation impacts, and end-of-life material recovery processes. Additionally, refining energy consumption estimates based on specific manufacturing techniques (e.g., machining vs. additive manufacturing) would result in a more precise evaluation of a material's environmental footprint in the manufacturing stage.

ACKNOWLEDGEMENTS

We would like to acknowledge the contributions of the previous student project titled 'Étude de l'impact Environnemental de l'industrie Manufacturière Aéronautique et des Moyens de Réduction de cet Impact' conducted by C. Atin, S. Dura, A. Lasserre, B. Roiron, R. Zuppo at ISAE-SUPAERO. Their work provided valuable insights and foundational research that contributed to this study.

REFERENCES

[1] L. Rupcic, E. Pierrat, K. Saavedra-Rubio, N. Thonemann, C. Ogugua, and A. Laurent. Environmental impacts in the civil aviation sector: Current state and guidance. *Transportation Research. Part D: Transport & Environment*, 119, 2023. doi:10.1016/j.trd.2023.103717. URL `https://doi.org/10.1016/j.trd.2023.103717`.

[2] N. Thonemann, E. Pierrat, K. M. Dudka, K. Saavedra-Rubio, A. L. S. T. Dragsdahl, and A. Laurent. Towards sustainable regional aviation: Environmental potential of hybrid-electric aircraft and alternative fuels. *Sustainable Production and Con-*

sumption, 45:371–385, 2024. ISSN 2352-5509. doi:10.1016/j.spc.2024.01.013. URL https://doi.org/10.1016/j.spc.2024.01.013.

[3] I. A. T. A. (IATA). Iata annual review, 2017. URL https://www.iata.org/contentassets/c81222d96c9a4e0bb4ff6ced0126f0bb/iata-annual-review-2017.pdf.

[4] M. Rendón, C. S. R., J. G. M., et al. Aircraft hybrid-electric propulsion: Development trends, challenges and opportunities. *Journal of Control, Automation, and Electrical Systems*, 32:1244–1268, 2021. doi:10.1007/s40313-021-00740-x. URL https://doi.org/10.1007/s40313-021-00740-x.

[5] R. Huang, M. Riddle, D. Graziano, J. Warren, S. Das, S. Nimbalkar, J. Cresko, and E. Masanet. Energy and emissions saving potential of additive manufacturing: the case of lightweight aircraft components. *Journal of Cleaner Production*, 135: 1559–1570, 2016. ISSN 0959-6526. doi:10.1016/j.jclepro.2015.04.109. URL https://doi.org/10.1016/j.jclepro.2015.04.109.

[6] K. A. Henschel. *Biocomposites in Aviation Structures on the Example of Flax and its Hybrids*. TUM.University Press, 2020. doi:10.14459/2019md1488689.

[7] E. Pierrat, L. Rupcic, M. Z. Hauschild, and A. Laurent. Global environmental mapping of the aeronautics manufacturing sector. *Journal of Cleaner Production*, 297:126603, 2021. ISSN 0959-6526. doi:10.1016/j.jclepro.2021.126603. URL https://doi.org/10.1016/j.jclepro.2021.126603.

[8] A. Murphy. Materials in aircraft structures. *The Journal of the Royal Aeronautical Society*, 70(661):114–119, 1966. doi:10.1017/S0001924000094021. URL https://doi.org/10.1017/S0001924000094021.

[9] R. Boyer, J. Cotton, M. Mohaghegh, and R. Schafrik. Materials considerations for aerospace applications. *MRS Bulletin*, 40(12):1055–1066, 2015. doi:10.1557/mrs.2015.278. URL https://doi.org/10.1557/mrs.2015.278.

[10] S. Sapre and J. Morlier. Environmental impact manufacturing calculations, 2024. URL https://github.com/mid2SUPAERO/pLCA-MDO/blob/main/Environmental_Impact_Manufacturing_Calculations.xlsx.

[11] J. V. de Oliveira Fernandes Lopes. *Life Cycle Assessment of the Airbus A330-200 Aircraft*. PhD thesis, Instituto Superior Técnico, Universidade de Lisboa, 2010.

[12] A. L. Leao, I. Cesarino, M. Chanes, E. C. Botelho, O. A. T. Dias, and M. Jawaid. *Ecologically Enhanced Natural/Synthetic Polymer Hybrid Composites for Aviation-Interior and Secondary Structures*, pages 43–59. Springer Nature Singapore, Singapore, 2023. ISBN 978-981-99-1583-5. doi:10.1007/978-981-99-1583-5_4. URL https://doi.org/10.1007/978-981-99-1583-5_4.

[13] D. R. Lloyd Jenkinson, Paul Simpkin. Data a: Aircraft data file. In *Civil Jet Aircraft Design*. Elsevier, 1999. URL https://booksite.elsevier.com/9780340741528/appendices/data-a/default.htm.

[14] S. Al-Alimi, N. K. Yusuf, A. M. Ghaleb, M. A. Lajis, S. Shamsudin, W. Zhou, Y. M. Altharan, H. S. Abdulwahab, Y. Saif, D. H. Didane, I. S.T.T, and A. Adam. Recycling aluminium for sustainable development: A review of different processing technologies in green manufacturing. *Results in Engineering*, 23:102566, 2024. ISSN 2590-1230. doi:10.1016/j.rineng.2024.102566. URL https://doi.org/10.1016/j.rineng.2024.102566.

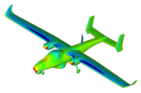
[15] O. Takeda, T. Ouchi, and T. Okabe. Recent progress in titanium extraction and recycling. *Metallurgical and Materials Transactions B*, 51:1315–1328, August 2020. doi:10.1007/s11663-020-01898-6. URL https://doi.org/10.1007/s11663-020-01898-6.

[16] C.-H. Chen, C.-L. Chiang, J.-X. Wang, and M.-Y. Shen. A circular economy study on the characterization and thermal properties of thermoplastic composite created using regenerated carbon fiber recycled from waste thermoset cfrp bicycle part as reinforcement. *Composites Science and Technology*, 230:109761, 2022. ISSN 0266-3538. doi:10.1016/j.compscitech.2022.109761. URL https://doi.org/10.1016/j.compscitech.2022.109761.

[17] M. Ashby and D. Cebon. Materials selection in mechanical design. *J. Phys. IV France*, 03(C7):C7–1–C7–9, 1993. doi:10.1051/jp4:1993701. URL https://doi.org/10.1051/jp4:1993701.

[18] J. Blank and K. Deb. pymoo: Multi-objective optimization in python. *IEEE Access*, 8:89497–89509, 2020.

[19] F. International. Airbus and boeing report december and full-year 2023 commercial aircraft orders and deliveries, 2024. URL https://flightplan.forecastinternational.com/2024/.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

HIGH-FIDELITY MULTIDISCIPLINARY OPTIMISATION PROCESSES FOR SOLVING AERO-STRUCTURAL PROBLEMS: RECENT DEVELOPMENTS FROM AN INDUSTRIAL PERSPECTIVE.

J. Brezillon[1]*, I. Fadli[1], K. Thanissaranon[1] and C. Conlan-Smith[2]

[1]: Flight Physics Capability Development
Airbus Operations S.A.S.
316 route de Bayonne
31060 Toulouse Cedex 09
France
[2]: Flight Physics Capability Development
Airbus Operations GmbH
Airbus-Allee 1
28199 Bremen
Germany

{Joel.Brezillon, Imane.Fadli, Kittinan.Thanissaranon, Cian.Conlan-Smith}@airbus.com,
http://www.airbus.com

Abstract *The purpose of this paper is to present the Multidisciplinary Design Optimisation (MDO) capabilities developed at Airbus within the EU-funded Horizon Europe NEXTAIR project (multidisciplinary digital-enablers for NEXT-generation AIRcraft design and operations), with applications to the DLR-F25 configuration, a small-to-medium-range aircraft featuring a high-aspect ratio wing. A simplified but realistic use case around an aero-structural wing optimisation problem has been defined by Airbus and proposed to partners in the project, to demonstrate the feasibility and maturity of MDO processes based on medium to high-fidelity simulations in early design trades. From an industrial viewpoint, the selection of the MDO architecture to solve such a problem is key: while the Multi-Disciplinary Feasible approach is certainly proven to efficiently solve aero-structural MDO problems, it turns out to be difficult or impractical to operate in an industrial context where monodisciplinary capabilities are already well established and mature. Alternative distributed MDO architectures offer the opportunity to reuse these disciplinary processes and to scale up MDO complexity. Finally, the bi-level formulation, originally proposed by IRT-Saint Exupéry and available in the workflow orchestrator GEMSEO, has been down-selected to solve the aero-structural optimisation problem. GEMSEO is used here to orchestrate and connect the in-house aerodynamic shape optimisation capability PODShaperOpt with the loads-structure sizing suite cpacs-MONA developed by DLR through a HTTPS client-server. Initial results on the DLR-F25 configurations highlight the possibilities and difficulties of this bi-level formulation for solving aero-structural wing design problems.*

Keywords: MDO, Aero-structure, Optimisation, Distributed Architecture, Bi-level MDO formulation, Adjoint

1. INTRODUCTION

The aircraft industry is working on the limitation of carbon emissions by among other measures reducing energy consumption during aircraft operations. Aircraft with a high span or with a High-Aspect-Ratio Wing promise to achieve new aerodynamic performance standards and to reduce aircraft energy consumption. However, they require a re-think of the design capabilities, notably in the early-phase to ensure the right level of performance. For example, increasing wing aspect ratio typically leads to more flexible structural designs which creates a stronger coupling between disciplines, such as structure and aerodynamics or handling qualities. Due to these strong multidisciplinary couplings, high-fidelity based Multidisciplinary Optimisation (MDO) is seen as necessary in early-phase design in order to secure these performance targets.

In this paper, we investigate the applicability of high-fidelity MDO methods for industrial design, re-using existing well-established monodisciplinary optimisation processes. The paper has the following outline: Section 2 details the test case, from the chosen configuration to the formulation of the MDO problem. Different MDO architectures and their applicability to industrial design are discussed in Section 3. Section 5 includes details of the specific analysis and optimisation methods used. Monodisciplinary optimisation results are presented in Section 6 and multidisciplinary results in Section 7. Finally, conclusions are discussed in Section 8.

2. MDO TEST CASE DESCRIPTION

2.1 DLR-F25 Model

In the frame of the LuFo project VirEnfREI, DLR recently developed in close collaboration with Airbus, the F25 configuration, a short-to-medium-range High-Aspect Ratio Wing Transport Aircraft[1, 2, 3].

This configuration is well suited for MDO activities for the following main reasons:

- high-aspect ratio wing is a promising technology to improve aerodynamic efficiency to be applied to the next generation aircraft, which surely will be a short-to-medium-range aircraft following the recent introductions of new developments in the long-range sector;
- the DLR-F25 configuration has today a pretty “young” design maturity, representing a quite realistic challenge for demonstrating the benefit of high-fidelity MDO capabilities in early design phase;
- the DLR-F25 configuration has been selected in the NEXTAIR [4] and Clean Aviation UpWing [5] European projects, which among other addresses the development of digital enablers for efficient aircraft ecosystem design. It was agreed between partners (DLR, IRT, ONERA, NTUA and Dassault Aviation) to exchange results to maximize synergies and to enable fast and easy transfer of findings between both projects.

The baseline three-view drawing is presented in Fig. 1.

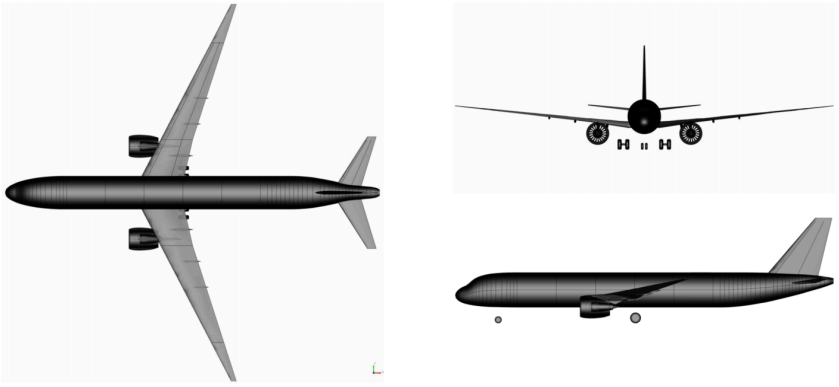


Figure 1. Three-view drawing of the DLR-F25 [1].

The Top-Level Aircraft Requirements and Assumptions for this configuration are based on the Airbus A321neo configuration [6], and summarized in Table 1.

Table 1. Top Level Aircraft Requirements and Assumptions.

Requirement/Assumption	Unit	Value
Design Range	NM	2’500
Evaluation Range	NM	800
Design PAX (single class)	-	239
Mass per PAX	kg	95
Max. Payload	kg	25’000
Cruise Mach number	-	0.78
Max. operating Mach number	-	0.80
Max. operating altitude	ft	40’000

The Key Aircraft Characteristics for the DLR-F25 used here, are summarized in Table 2.

Table 2. Key Aircraft Characteristics of the DLR-F25 configuration.

Aircraft Characteristics	Unit	Value
Mass		
Max. Take-Off Mass (MTOM)	kg	86’000
Max. Landing Mass (MLM)	kg	74’800
Max. Zero Fuel Mass (MZFM)	kg	72’000
Geometry		
Wing Span	m	45
Wing Aspect Ratio	-	15.6
Wing MAC	m	3.535
Wing Reference Area	m^2	130.12

2.2 Trade and Design Variables

From an industrial perspective, one of the objectives of the design task is to demonstrate the usability of high-fidelity MDO to support the aircraft design process in its early-phase when decisions on the configuration are not frozen. In this context, two major types of variables are considered:

- **Design variables**, which are typical parameters directly controlled by the optimizer to find the optimal solution;
- **Trade variables**, which are of interest to help understand differences between configurations. They should answer questions like "what is the impact on a quantity of interest if this variable is modified". They are generally explored in a discrete manner, defined either by the designers or through more advanced DOE techniques.

The advantage of combining trade and design variables is to help designers to better understand their problem and capture the rate exchange between different configurations of similar maturity thanks to the use of optimisation techniques, which is key to support design decisions in the early design phase. Of course, trade variables can also be considered directly by the optimizer to find the overall optimal solution. However, in this case, designers will have access to much less information to understand why this is the optimal solution, to quantify the robustness of this solution, or to assess the impact of other disciplines - like the cost for manufacturing - not always available in the original MDO problem.

In this context, one of the objectives of the exercise is to demonstrate how MDO could be used to support high-speed wing design trade-offs, and Table 3 below presents a typical list of trade and design variables. The "Thickness Scale" trade parameter is a global factor to be applied to the wing thickness in spanwise direction, assuming that the parameterisation of wing sections enables the control of either the maximum wing thickness or the minimum thickness at key positions like the front and rear spars.

Table 3. Trade and Design Variables..

Variable name	Type	Variable name	Type
Wing planform		**Moveables**	
Area	Trade	Ailerons spanwise position	Design
Span	Trade	Ailerons spanwise size	Design
Aspect Ratio	Trade	Ailerons chordwise size	Design
Sweep	Trade	Spoilers spanwise position	Design
Thickness Scale	Trade	Spoilers spanwise size	Design
Sectional		Spoilers chordwise size	Design
Twist	Design		
Camber	Design		
Structural parameters	Design		

2.3 Common MDO problem

In the frame of the European project NEXTAIR, Airbus defined a common MDO test case to enable fair and comprehensive comparison of methodologies used by partners working on this configuration. This simplified test case is not intended to be fully representative of an industrial design exercise, but rather to demonstrate the feasibility, maturity and benefit of the medium to high-fidelity simulations to solve an aero-structural MDO problem.
The objective for this task is to improve the fuel burned in cruise of the aircraft for a typical 800 NM mission as sketched in Fig. 2, with a 20'000 kg of payload. Only the passenger and their baggage are considered as payload, there is no cargo. The mission altitude is 36'000 ft and is considered fixed during the optimisation.

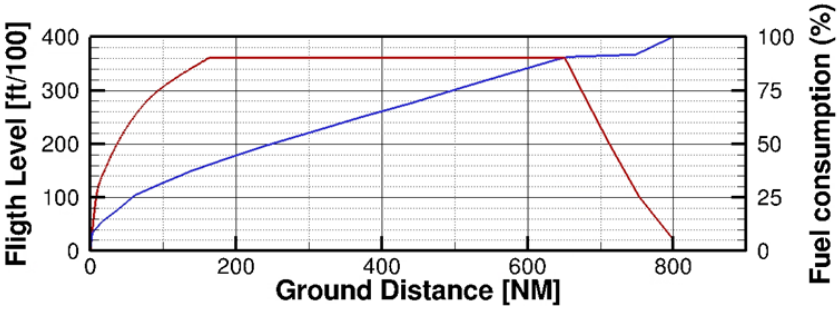


Figure 2. Mission profile of the baseline DLR-F25 configuration with flight level (red) and fuel consumption (blue).

In this exercise, minimising the fuel mass (FM) needed to perform the 800 NM cruise mission is considered as the objective function. This objective function is derived from the Breguet equation, see Eq. (1), where the aerodynamic performance is computed at a single flight cruise condition, i.e. CL=0.59 and FL=360. Constant parameters used in Eq. (1) are shown in Table 4. A correction factor (α) is applied to the Breguet range equation to take into account climbing and landing phases. This correction factor is based on the evolution of the fuel consumption during the full mission, with a value tuned for this exercise. The Take-Off Mass (TOM) is also selected to ensure a lift coefficient around CL=0.59 in cruise. It is proposed that any variation of mass due to the structural optimisation process will modify the take-off mass accordingly. This mass variation will also impact the target lift coefficient in cruise, which can be in fact compensated by adapting the altitude accordingly. It is however expected that the mass variation will be small enough during the optimisation to neglect the impact of the variation of altitude on drag. As a consequence, the cruise condition (CL=0.59 and FL=360) is kept constant during the optimisation exercise, regardless of the mass variation.

Based on the Breguet range equation, the range in cruise is

$$R = \alpha \cdot \frac{CL}{CD} \cdot \frac{V}{g \cdot T_{\mathrm{SFC}}} \cdot \ln\left(\frac{TOM}{TOM - FM}\right) \tag{1}$$

and the fuel mass needed for this mission is:

$$FM = TOM \cdot \left(1 - e^{-\frac{1}{\alpha} \cdot \frac{CD}{CL} \cdot \frac{R \cdot g \cdot T_{\mathrm{SFC}}}{V}}\right) \tag{2}$$

Table 4. Key Aircraft Characteristics of the DLR-F25 configuration.

Parameter	Unit	Value
Mission Range (R)	NM	800
Breguet range correction factor (α)	-	0.974
Cruise Mach number (M)	-	0.78
Cruise altitude (FL=360)	ft	36'000
Cruise speed (V)	$m.s^{-1}$	230.23
Lift coefficient in cruise (CL)	-	0.59
Thrust-specific fuel consumption (T_{SFC})	kg/s/kN	0.0135

Drag coefficient and hypothesis for aerodynamic flow computations in cruise

In order to enable a fair comparison between partners, it is proposed to use the baseline drag coefficient as reference ($CD_{Baseline}$) and to apply drag variation obtained during the optimisation (at CL=0.59 and Cmy=0), such to start with the baseline drag at the beginning of the optimisation. The drag coefficient used during the optimisation is defined as:

$$CD=CD_{Baseline}+\Delta CD \quad (3)$$

The parameters in Table 5 are taken into consideration to perform aerodynamic flow computations:

Table 5. Key Aircraft Characteristics of the DLR-F25 configuration.

Parameter	Unit	Value
Cruise altitude	ft	36'000
Air flow density	kg/m^3	0.365
Static pressure	Pa	22'729
Reynolds number	m^{-1}	5'908'190
Lift coefficient	-	0.59
Baseline Drag Coefficient ($CD_{Baseline}$)	-	0.0301

Mass coefficients and hypothesis for structural sizing

Similarly, to the aerodynamic discipline, any mass variation obtained during the structural sizing process is directly added to the Aircraft Take-Off Mass of the baseline configuration ($TOM_{Baseline}$). The Aircraft Take-Off Mass used is then updated as:

$$TOM=TOM_{Baseline}+\Delta Mass \quad (4)$$

Table 6 presents masses obtained on the baseline configuration (with a Lift to Drag ratio of LoD=19.6 at CL=0.59), which should be the same for all partners at the beginning of the optimisation.

Table 6. Mass distribution on the baseline DLR-F25 configuration.

Baseline Mass Distribution	Unit	Value
A/C Take-Off Mass ($TOM_{Baseline}$)	kg	79'985
A/C Mass in cruise (CL=0.59, FL=360)	kg	75'746
A/C Landing Mass (LM)	kg	73'153
Operating Empty Mass (OEM)	kg	48'466
Fuel Mass ($FM_{Baseline}$)	kg	6'832

Fuel Mass sensitivities with respect to structural mass and drag level

Table 7 gives sensitivities of fuel mass in cruise, when reducing the structural mass by 1% of OEM (i.e. 485kg) and the drag by 1% of total baseline drag (i.e. 0.0003). It can also be observed that a structural mass reduction of 255 kg gives the same fuel mass benefit as would give a 0.0001 drag reduction.

Table 7. Fuel sensitivities on structural mass and drag level.

Quantities	Level of variation	ΔFM
Structural Mass	-1% (-485 kg)	-41 kg
Structural Mass	-255 kg	-22 kg
Drag Coefficient	-1% (-0.0003)	-65 kg
Drag Coefficient	-0.0001	-22 kg

3. MDO ARCHITECTURE

This section describes three formulations to solve an aero-structural MDO problem, namely: Multi-Disciplinary Feasible (MDF), the Embedded and the bi-level formulations. The first formulation is considered the state-of-the-art and taken as a reference formulation, while the two others are alternative multi-level formulations, originally proposed by DLR in [7] and by IRT-Saint Exupery [8] respectively.

3.1 Multi-Disciplinary Feasible (MDF)

The MDF formulation is considered here as the reference MDO formulation. This approach consists of a single optimiser that controls all the disciplines at the same level. This optimiser determines all the design variables values that will minimize the objective function while fulfilling all the constraints and provides design variables to each discipline for the evaluation. The XDSM (eXtended Design Structure Matrix) of the MDF formulation is shown below in Fig. 3.

The MDF formulation requires a fully integrated multidisciplinary simulation process including cross-sensitivities of each discipline which capture sensitivities with respect to design variables owned by other disciplines. These cross-sensitivities are not always feasible or available in industrial context, and simply neglecting them can lead to convergence issues in the optimisation process. This creates a need to explore alternative MDO formulations more in line with the disciplinary way of working in industry, while converging to a true multidisciplinary solution that is able to capture trade-offs between disciplines as we expect from the MDF formulation.

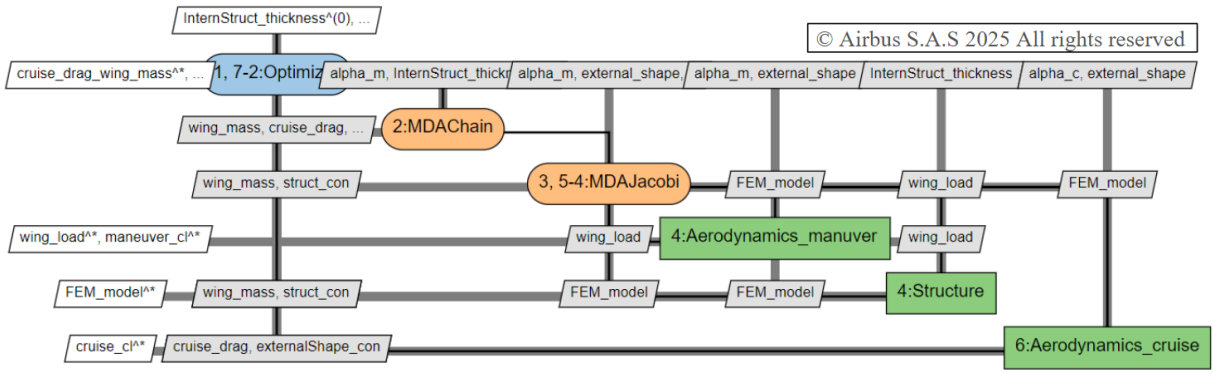


Figure 3. Generic xDSM diagram of the aero-structural MDF formulation.

3.2 The Bi-Level Embedded Formulation

Within the numerous range of bi-level formulations, the bi-level-embedded (BL-EMB) formulation originally proposed by DLR [7] (also called a "two-optimizer formulation" by the authors) is an appealing MDO architecture. Unlike the MDF formulation, 2 optimisers are used where the structural optimisation is embedded within the aerodynamic optimisation and still reuses the classical structural sizing process. The xDSM below (Fig. 4) shows an example of the optimisation process, where the optimiser at the upper level controls the external shape along with the angle of attack at cruise. For a given external shape, there is an optimisation embedded process (rounded rectangle in yellow) called `Structure_Opt` providing the structural model needed for aero-elastic evaluation in cruise condition (Aerodynamics_cruise). The standard BL-EMB approach neglects cross-sensitivities, like the drag sensitivities with respect to structural elements, but strategies exist to recover them as proposed in [7].

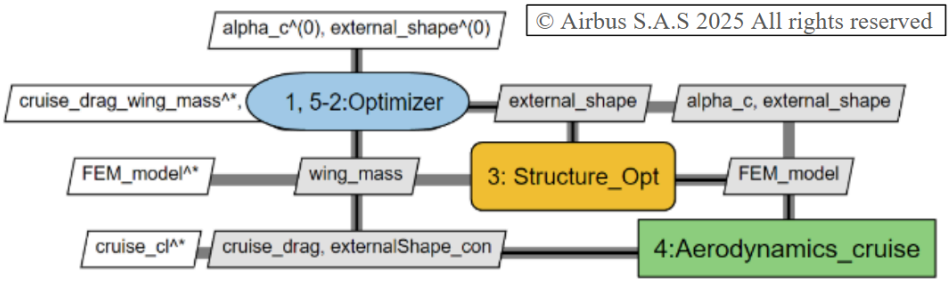


Figure 4. Generic xDSM diagram of the Embedded formulation for an aero-structural optimisation.

3.3 The IRT-Bi-Level Formulation

The bi-level formulation, as proposed by IRT [8, 9], consists of an aerodynamic and structural optimisation processes at the lower-level and a system-level optimiser at the upper-level to explore multidisciplinary trade-offs between aerodynamics and structure. Several approaches exist to select the design parameters at system level, and the wing root rolling moment in cruise is retained in this study. For each rolling moment value, the optimised wing design is determined through a balance between two disciplinary subprocesses: StructOpt and AeroOpt.

StructOpt minimizes the wing mass by optimising wing structural thicknesses, while satisfying stress constraints from multiple load conditions. On the other hand, AeroOpt minimizes the aircraft's drag coefficient during the cruise case by optimising the external wing shape, ensuring compliance with the wing rolling moment specified by the system-level optimiser. AeroOpt incorporates the wing's flexibility by using the structural configuration from the latest StructOpt run.
The interaction between these two subprocesses is managed by the block coordinate descent (BCD) loop [10], which alternates between StructOpt and AeroOpt, monitoring changes in local design variables to ensure convergence and achieve a balanced, optimised wing design.

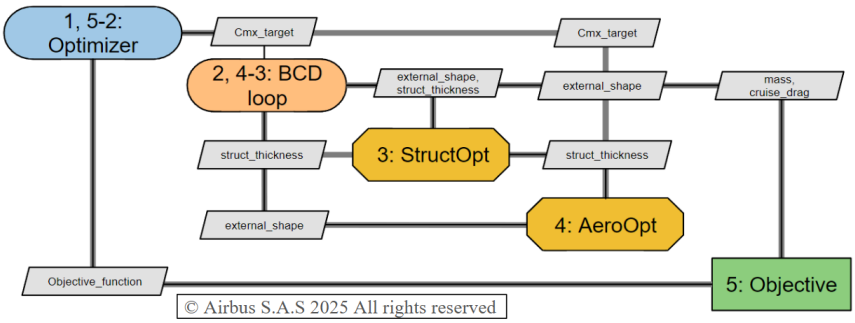


Figure 5. Generic xDSM diagram at system level of the bi-level formulation for the aero-structural optimisation.

4. MDO PROCESS

While Airbus has tested all three MDO architectures on a simplified wing-only testcase described in [8], the IRT-bi-level formulation has been retained here since it better mimics the disciplinary approaches used at Airbus and does not require cross-coupled terms while providing good multidisciplinary trades. Referring to Fig. 5, one can see 4 main sub-processes: Optimiser (at System Level), StructOpt, AeroOpt, and Objective.

4.1 System Level Optimiser

The system level optimiser is in charge of converging the multidisciplinary optimisation problem. The objective function is the Fuel Mass needed to perform the "evaluation mission" and is a function of the aircraft mass and drag in cruise as mentioned in the previous section.
The only design variable managed at this level is the target wing rolling moment. In fact, the aircraft rolling moment (Cmx) has been retained, because the wing rolling moment is today not available as an adjoint variable needed as constraint for the disciplinary aeroshape optimisation. The reason for having Cmx as a design variable at the system level is so to encourage exploration of the whole design space to promote the emergence of the multidisciplinary solution between structural and aerodynamic disciplines similar to that of an MDF formulation.
The bi-level formulation is orchestrated by GEMSEO [11], an MDAO python library developed by IRT. GEMSEO controls the sequence and information transfer between sub-processes as well as determining the calculations convergence during the BCD loop.

4.2 AeroOpt: Aerodynamic for Performance Process

The simulation tool used behind the AeroOpt subprocess is an Airbus tool called PODShaperOpt [12], where the CFD solver Tau [13] will be used for solving the Reynolds-averaged Navier–Stokes (RANS) equations. The objective is to minimize the aircraft drag in cruise, considering wing flexibility of a given structural model. The AeroOpt process is built around 4 main pillars: CAD-Rom approach for the shape parametrisation, CFD-CSM coupled computations for the aerodynamic states, the adjoint approach for efficient computations of the sensitivities [14, 15] and a second-order gradient-based constrained optimiser.

CAD-Rom approach

The CAD-Rom approach consists in capturing the behaviour of a shape parametrisation, based on snapshots. These snapshots are in fact a sampling of the design space described by the shape parameterisation and are used to create a surrogate model based on a proper orthogonal decomposition (POD) [12].

The resulting POD model acts as a surrogate of the CAD kernel where the inputs are the current set of design variables and outputs are deformation vectors on the surface of the wing. These deformation vectors can then be provided to the mesh deformation tool to update the CFD mesh for the current design. There are two main advantages to this approach, the first is that it completely removes the CAD software from the optimisation process. The second advantage is that the surrogate model is fully differentiable and thus can provide design gradients compatible with the adjoint process, which is not available in most commercial CAD softwares.

While Airbus has developed their own wing shape parameterisation, it has been agreed to rely on the parametrisation and snapshots provided by DLR, in order to ease future comparisons between partners. The wing shape parametrisation consists of controlling the spanwise twist distribution at 5 sections, the airfoil cambers are controlled in 6 sections along the span with 8 control points per section - the airfoil thicknesses are unchanged. Two parameters are added to control the angle of attack of the aircraft and the horizontal tail plane (HTP) rotation. In total, 55 design variables (53 external shapes + 2 angles) are used in AeroOpt process.

Flexible CFD/CSM computations

As a preprocessing step to the AeroOpt discipline a static condensation is performed on the structural sizing model provided by StructOpt (as NASTRAN inputs). This process reduces the number of degrees of freedom to a subset of points, each with six degrees of freedom. These points are placed at the centre of the fuselage and wingboxes throughout their entire length, see Fig. 6a. The stiffness or Eigen modes/vectors of the structure can then be expressed in-terms of this subset of points. There are two main motivating factors for using this condensation process rather than the full FEM model. The first is that it simplifies the coupling between the aero surface and structure as well as ensuring a smooth mesh deformation in the CFD mesh. The second is that we assume the stiffness of the structure is fixed throughout the AeroOpt process and hence no change in stiffness is captured when the wing changes shape throughout the aerodynamic shape optimisation. It is also assumed that the center of the wingbox does not change significantly throughout the shape optimisation and

as such by considering only structural nodes at the wingbox center there are no updates required to the structural mesh throughout the aero optimisation which also simplifies the coupled gradient calculations.

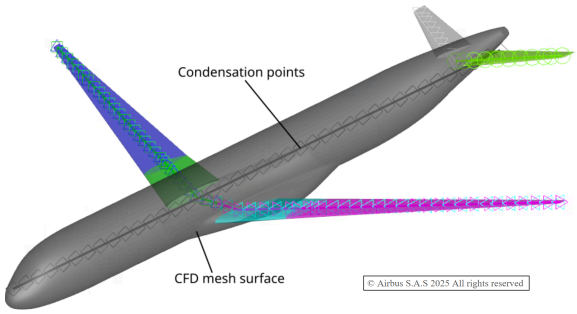


Figure 6a. The load-displacement transfer for CFD-CSM coupling is defined using the condensation points and CFD mesh surfaces shown here.

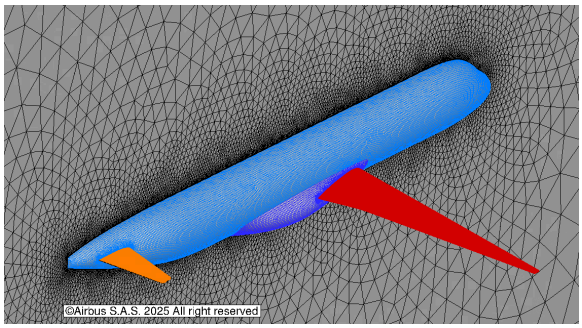


Figure 6b. CFD Grid Mesh (half-configuration) on the DLR Wing-Body-HTP configuration.

The stiffness matrix and Eigen modes/vectors are also calculated via NASTRAN in the preprocessing. These quantities are provided as an input to the flow solver and since the stiffness is fixed throughout the aero optimisation they are only calculated once at the start of the AeroOpt discipline. As all stiffness quantities are calculated as a preprocessing step, we avoid the need for additional calls to NASTRAN throughout the CFD/CSM coupling. For the studies presented in this work all deformations are captured using a static modal analysis (using precalculated Eigen modes/vectors) with free-free boundary conditions rather than relying on stiffness matrices as this would require modifying the structural models from StructOpt to include clamped boundary conditions.

Interpolation functions for the load-displacement transfer are defined by joining condensation points into dummy 1D-beam elements, where displacements from these elements are transferred to the CFD mesh using the perpendicular distance from the element to the CFD node. The load transfer is then defined as the transpose of the displacement transfer. Loads applied to the structure include pressure loads from the aero surface as well as gravity loads from the structure. The CFD mesh used here was generated by DLR and shared to all partners involved in the task. The mesh is presented in Fig. 6b and consists of about 10 Mio points. While it does not contain an engine or vertical tail plane (VTP) geometries, the gravity loads due to these components are still accounted for in the CFD/CSM coupling. A Gauss-Seidel method is used to converge the coupled CFD/CSM problem and the coupled adjoint problem.

Flexible and Fixed-flexible adjoint

Two approaches have been investigated for solving aeroshape optimisation problems accounting for flexibility. The first approach accounts for full coupling where the CFD/CSM problem is solved on each design iteration, followed by solving a coupled adjoint problem. The second approach solves the coupled CFD/CSM problem only on the first iteration and freezes these displacements throughout the optimisation processes where only a CFD problem

is solved on each subsequent iteration. This approach corresponds to a rigid aero optimisation in the flight shape where the flight shape bending is defined at the start of the AeroOpt process, and as such there is no need to conduct a coupled adjoint analysis (only CFD/Rigid adjoint). Because the CFD/CSM coupling is not recalculated there may be a mismatch between the flight shape bending that is modelled and the true bending, which may impact the optimised design as demonstrated in [14]. However, it is predicted that the true bending on the first and final iteration of AeroOpt will converge as the BCD loop also converges.

4.3 StructOpt: Structural Sizing Process

This subprocess is in charge of minimising the structural mass for a given external wing shape. It determines the thickness of the structural elements within the wing. The result obtained from this optimisation process defines the sized structure under a set of representative load conditions, where the external wing shape obtained from the previous AeroOpt run is taken as an input.

DLR developed cpacs-MONA [16], an automated aeroelastic structural design and analysis tool, and made it at disposal to perform the StructOpt process. The tool automates the creation of a parametric finite element model based on the Common Parametric Aircraft Configuration Schema (CPACS) input, a structured description of main aircraft characteristics. This model is then employed for load analysis and subsequent structural optimisation of major aircraft components such as the wing. Cpacs-MONA includes an extensive loads analysis module where manoeuvre and gust loads for various flight conditions and mass configurations are computed, resulting in over 600 load cases. While cpacs-MONA is able to automatically search for critical load envelopes for structural sizing, it also provides the option to utilize preselected load cases. In this study a set of preselected critical load cases has been compiled on various wing configurations with various twist distributions. From this preliminary study, about 45 load conditions associated with 6 different mass cases have been identified, and retained during the whole MDO campaign. It is worth mentioning that the aerodynamic model can additionally consider aileron effects for maneuver load alleviation. All these load conditions have been shared to all partners to enable fair and consistent comparisons in the future. Following the load analysis, cpacs-MONA performs a structural sizing of wing ribs, spars, skin panels, etc. based on stress, strain and buckling constraints, using the NASTRAN solver. This assists in the development of lightweight yet robust structures capable of withstanding the applied loads while adhering to aeroelastic requirements and avoiding aileron reversal. Once the load analysis and sizing are completed, a subsonic flutter check is performed in cpacs-MONA. This capability allows designers to analyse the interplay between the aircraft's structural flexibility and aerodynamic forces, enabling the design of structures that can effectively resist these instabilities.

Finally, cpacs-MONA employs parametric modelling techniques. This means that with only an update of the wing's external shape, the analysis models and structural optimisations are automatically performed, facilitating the efficient exploration of different design possibilities.

4.4 Objective

This subprocess is a simple analytical function which takes the resulting mass and drag coefficient from StructOpt and AeroOpt respectively to compute the Fuel Mass, based on the Breguet’s range equation in Eq (1).

4.5 HTTPS connection with GEMSEO

From an implementation perspective, the calculation for each subprocess will be carried out on dedicated computing infrastructure, depending on the underlying simulation tool, as detailed in Table 8 below. Note that AIRBUS SC, AIRBUS HPC and AIRBUS PSO refer to the Airbus Scientific Computer, Airbus High-Performance Computing infrastructures and Airbus PODShaperOpt, respectively. It is evident that the entire optimisation process is executed on three different operating systems, two of which are located on Airbus premises and the third within the DLR server ecosystem. Consequently, a dedicated communication system has been implemented to ensure streamlined and automatic data communication.

Table 8. Subprocesses within the bi-level formulation along with their corresponding simulation tools and execution infrastructure.

bi-level Subprocess	Simulation Tool	Execution platform
System Level Optimiser	IRT-GEMSEO	AIRBUS SC
StructOpt	DLR-MONA	DLR Cluster
AeroOpt	AIRBUS-PSO	AIRBUS HPC
Objective	IRT-GEMSEO	AIRBUS SC

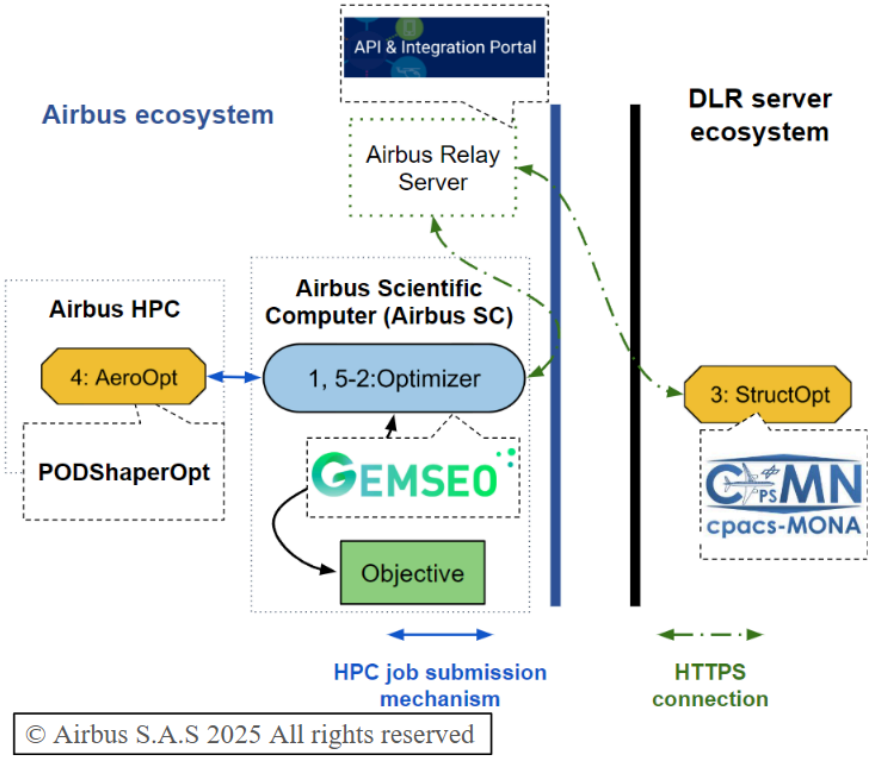


Figure 7. The means for data transfer between subprocesses.

Figure 7 above illustrates the overall view of the subprocesses and their corresponding

simulation tools and execution infrastructure. A HTTPS connection was chosen to ensure secure communication with the DLR server. IRT, in collaboration with DLR, set up the server on the DLR side, enabling Airbus and other partners to access cpacs-MONA via an HTTPS connection. Airbus, on its part, configured an additional process called 'Airbus Relay Server' to ensure a secure connection between the Airbus and DLR ecosystems. The data transfer and remote execution of cpacs-MONA via HTTPS was implemented by IRT in GEMSEO [17] and made available to Airbus during the project.

5. MONODISCIPLINARY OPTIMISATIONS RESULTS

Before launching the multidisciplinary process, some trade studies have been performed to assess the robustness and behaviour of the aeroshape optimisation and structural sizing processes.

5.1. Flexible and Fixed-Flexible aeroshape optimisations

In aerodynamics, the adjoint approach is considered as a must-have solution to efficiently perform shape optimisations. The flexible adjoint enables the fast, accurate and efficient computation of the aerodynamic sensitivities with respect to the shape parameters considering the aero-structural effects. It has been implemented at Airbus [14, 15] and is now in the hands of the aero designers for their daily tasks. Its main benefits are the introduction of a unique "jig" shape for both the aerodynamic and structure parametrisation, ensuring consistency on the shape during the optimisation process, and the consideration of the "true" flight shapes in all cruise flight phases, which may be an important aspect to consider on high-aspect ratio wings subject to large flexibilities. This approach however requests additional cost compared to the rigid adjoint, since coupled CFD-CSM computations have to be performed on the state and all adjoint equations at every optimisation steps. A second drawback is that the structural model is considered as fixed during the aerodynamic shape optimisation, the structure is not updated when the external shape is modified. In order to do so, cross-sensitivities are needed, like the sensitivities of the aerodynamic coefficient with respect to the structural design parameters, and the ambition of the study is to avoid this costly development.
Romain Olivanti observed in his thesis [14] that a cost-effective alternative of the flexible adjoint consists in "fixing" the structural deformation during the shape optimisation process. More precisely, a CFD/CSM coupled computation is performed only once, at the beginning of the optimisation, and this jig-to-flight deformation is simply reused during the rest of the optimisation. In this way, all adjoint states are computed considering rigid shapes, reducing the computational effort, especially since an adjoint systems need to be solved for each cost function and constraint.
Both approaches have been tested on the DLR-F25 wing-body-HTP configuration sketched in Fig. 8a. The objective is to minimize the drag, while keeping the lift and pitching moment at CL=0.59 and Cmy=0.0 respectively. The shape parametrisation has been shared by DLR, and consists of 55 design variables in total, as explained in Section 4.2. The convergence histories of the optimisation process are presented in Fig. 8b, where both strategies present very similar behaviour. Since coupled problems are not systematically solved in the fixed-flexible adjoint approach there is significant cost saving of roughly a factor 3 in run time compared to the fully coupled flexible approach.

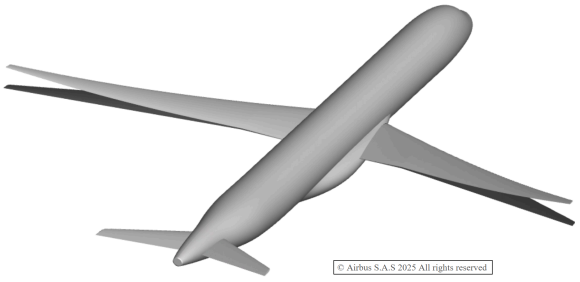


Figure 8a. DLR-F25 wing-body-HTP configuration with wing in jig and flight shapes.

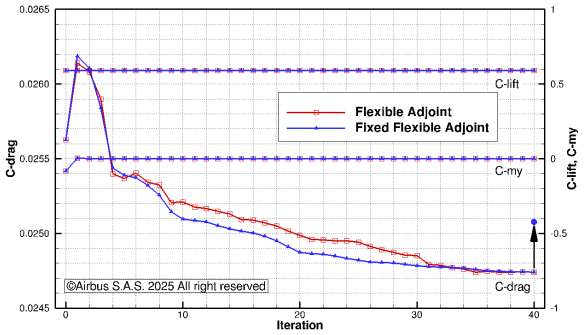


Figure 8b. Convergence plots of the aerodynamic shape optimisations with different formulations. The arrow shows the change in drag when recalculating the true bending for the converged fixed flexible design.

The drag coefficients are similar at the end of both optimisations, with a wave drag distribution in spanwise direction reduced to its minimum level and almost identical local lift distributions, as presented in Fig. 9a. But it is worth noting that the structural flexibility is not updated in the fixed-flexible adjoint approach. Once this optimal shape is updated, the resulting flight shape is no longer optimal with an increase of about 3 drag counts, mainly due to a new flight-twist distribution, as presented in Fig. 9b. The small curve break observed on the optimal twist distribution with the flexible adjoint is due to the twist parametrisation which purposefully does not ensure C2 continuity at the wing Yehudi break.

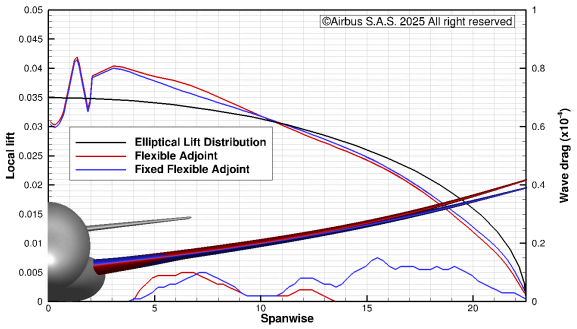


Figure 9a. Final lift and wave drag distribution in spanwise direction on optimised configuration.

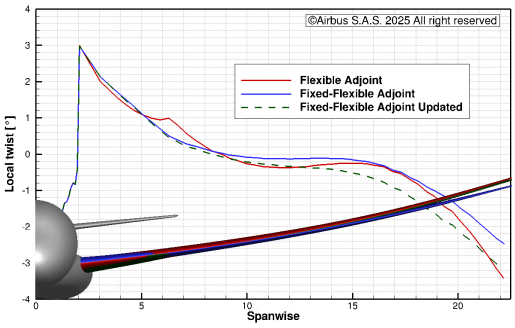


Figure 9b. Local wing twist of optimal shapes.

5.2. Fixed-Flexible aeroshape optimisations at different rolling moment

One of the major key aspects of the bi-level formulation is to drive the tight coupling between disciplines at the system level. Based on the experience made in close collaboration with IRT [8], the rolling moment in cruise is also used in this study to drive the overall optimisation at the system level. This implies the ability to perform aeroshape optimisations at a given rolling moment (Cmx).

In a first design attempt, 5 drag minimisations at constant lift and pitching moment have been performed at different rolling coefficients in a range between 0.19 and 0.25. Figure 10a presents the resulting optimisations, performed with the fixed-flexible approach and

considering the baseline structural model. A relatively smooth drag behaviour is observed, with a minimum located at around Cmx=0.4 corresponding as expected to the rolling moment obtained in the previous optimisation obtained without constraints on the rolling moment. This small study confirms the feasibility and the robustness of the process to perform systematic shape optimisations at different rolling moments.

5.3. Structural Sizing at different rolling moment

Similarly to the aeroshape optimisation process, the DLR structural sizing process cpacs-MONA is here evaluated at fixed cruise rolling moment. While this constraint can not be directly handled by MONA, the idea here is to perform the structural sizing on the optimised shapes obtained previously at a given rolling moment. A clear advantage can be noticed here, as performing the aeroshape parameterisation directly on the jig shape enables the outer model line to be easily provided to the structural discipline, without considering a flight-to-jig process.

Once the sizing process is performed on 5 different shapes, Fig. 10b synthesises the evolution of the resulting mass for the aircraft (OEM) and for its major components (Wing, Vertical Tail Plane and Horizontal Tail Plane). The fuselage mass is considered constant in this study. The major variation is coming from the wing structure, with a mass delta of about 600 kg at the two extreme rolling moments considered.

As expected, the behaviour of masses is smooth and increases smoothly and steadily with respect to the cruise rolling moment, confirming the relevance in considering this parameter at the system level to control both disciplines.

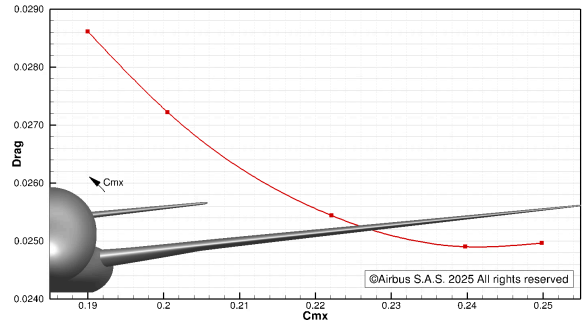


Figure 10a. Drag coefficients after shape optimisations at different rolling moments constraints. Fixed flexible adjoint considering baseline structural model.

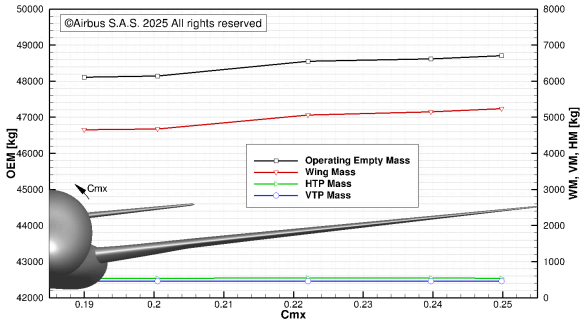


Figure 10b. Weight distributions after structural sizing on optimal shapes obtained at different rolling moments.

6. MULTIDISCIPLINARY AERO-STRUCTURAL OPTIMISATION

6.1 The need for the Block Coordinate Descent approach

The standard bi-level formulation made a strong assumption about the separability of the disciplines, which is at stake on tightly coupled aero-structural optimisation problems. As shown in the previous section, the structural sizing is performed on the aerodynamic optimal shape and as a consequence, the aerodynamic state associated with this geometry has to be re-assessed with the new structural model through CFD-CSM coupled computations. There is however no guarantee that the resulting flight shape remains optimal, nor that the rolling

moment constraint is satisfied: a new aerodynamic shape optimisation has to be performed with the new sized structural model to ensure consistency and optimality between both disciplines. This procedure has to be performed until no major change is observed between any two consecutive iterations. This is in simple words the purpose of the block coordinate descent, originally introduced by IRT in his bi-level formulation [10] and successfully applied on a simplified aero-structural problem in [8].

The recommended way to assess the convergence of the BCD loop is to observe the evolution of the local design variables. In this study, the evolution of the aerodynamic (drag, lift, pitching- and rolling-moments) and mass (Operating Empty Mass, Wing Mass, HTP Mass and VTP Mass) observables has been preferred to monitor the convergence since they give a more physical interpretation and understanding for engineers than local structural thicknesses or airfoil camber variations. Examples of the convergence of these observables are presented in Fig 11a and Fig. 11b for the aerodynamic coefficients and structural masses respectively. It should be noted that each point indicates the result of an aerodynamic shape optimisation and a structural sizing at a given BCD loop, except at the zeroth iteration corresponding to the baseline configuration as given by DLR. It can be seen that 5 BCD loops are needed to get convergence of observables.

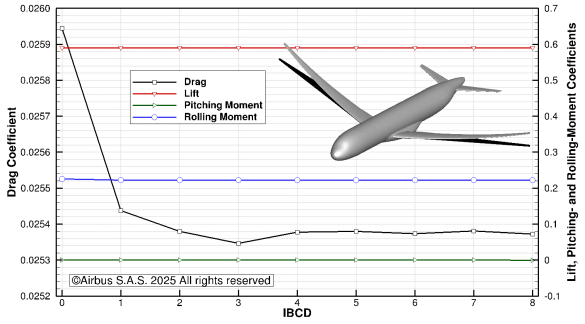


Figure 11a. Convergence of the aerodynamic coefficients during the BCD iterations - M=0.78, CL=0.59, Cmy=0.0, Cmx=0.222.

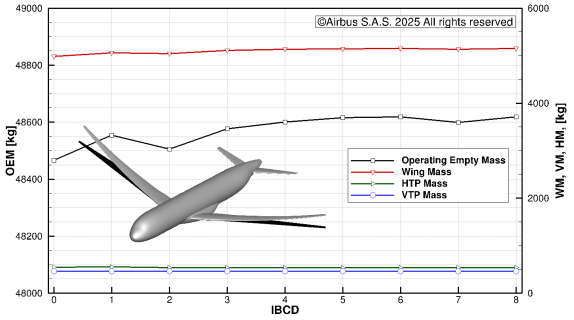


Figure 11b.Convergence of the weight distributions during the BCD iterations - M=0.78, CL=0.59, Cmy=0.0, Cmx=0.222.

6.2 Aero-structural optimisation of the DLR-F25 configuration

In order to capture the trade-off between the aerodynamic and structure disciplines, about 5 BCD runs have been performed concurrently at different rolling-moments in cruise. Each BCD run took about 4 to 6 loops to converge.

In total, about 25 structural sizings and 300 CFD computations are performed to capture the trade, all within 1 week wall-clock time. The latter represents an interesting turnaround time for medium-to-high-fidelity sensitivities in the early design phase. Figure 12a presents all the feasible solutions obtained during the study, i.e. trimmed solutions in lift and pitching moment with consistency between aerodynamic and structure states - which was not the case in Fig. 10. These configurations can now be plotted in a Mass vs. Drag diagram, see Fig. 12b: as expected, a Pareto front between wing mass and aircraft drag is observed, where a decrease of

0.0010 in drag coefficient leads to an increase of 200 kg in the wing mass. The Fuel Mass, which is the common multidisciplinary objective, as defined in Eq (1), is also plotted with colored symbols. It can be observed that the minimum fuel mass corresponds to the lowest drag coefficient, which is somehow expected considering the sensitivities of the fuel mass with respect to the drag coefficient and masses as presented in Table 7.

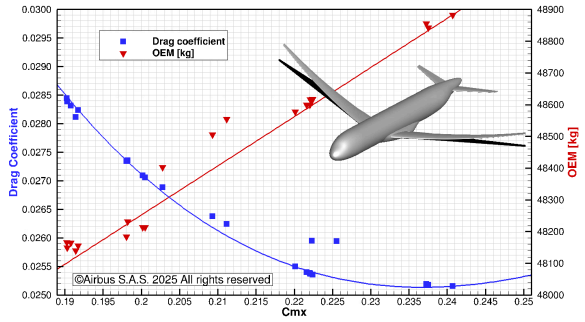


Figure 12a. Evolution of the Aircraft Drag (in blue) and Operating Empty Mass (in red) according to aircraft rolling moment in cruise.

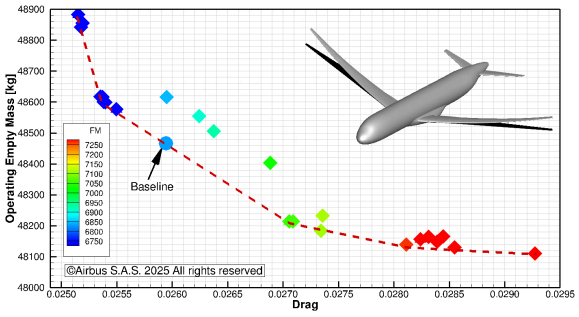


Figure 12b. Pareto front between OEM and aircraft drag - Symbols coloured by Fuel Mass (FM) level.

7. CONCLUSION

This paper has presented a multidisciplinary optimisation process, detailing the MDO test case description and the selected MDO architecture employing a bi-level with block coordinate descent formulation. The results obtained from both mono- and multidisciplinary optimisations demonstrate the potential for achieving good outcomes through this approach. It is however acknowledged that this approach searches for optimal multidisciplinary solutions resulting from a compromise between lowest drag and the lightest structure under a single constraint at the system level, and further developments are needed to explore more advanced approaches like the aeroelastic wing tailoring.

Future work includes a comparative analysis with our project partners, expanding the process to consider a wider range of cruise conditions and to incorporate control of airfoil thicknesses at the system level, and extending the analysis to a full aircraft configuration with the engine and vertical tail plane. This will enable a more comprehensive evaluation of the multidisciplinary optimisation process in a realistic design scenario.

8. ACKNOWLEDGMENTS

The DLR-F25 is a conceptual design of an advanced short-to-medium-range airliner provided by DLR. It features an initial shape of a very high-aspect-ratio wing and was developed as a challenge for high-fidelity disciplinary and multidisciplinary early design. The development was funded by the German Federal Ministry for Economic Affairs and Climate Action as part of the LuFo project VirEnfREI.

This work is part of the NEXTAIR project funded by the European Union under Grant Agreement No. 101056732. The views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or REA. Neither the European Union nor the REA can be held responsible for them.

The authors would like to particularly thank our colleagues at IRT and DLR for their invaluable support in developing this MDO framework within the NEXTAIR project. Their significant contributions were essential to the success of this work.

9. REFERENCES

[1] S. Wöhler, J. Häßy, V. Kriewall, Establishing the DLR-F25 as a Research Baseline Aircraft for the Short-Medium Range Market in 2035, icas2024_0398_paper.pdf, ICAS Congress 2024, Florence.

[2] S. Daehne et al., Structural aspects of the high aspect ratio wing: DLR-F25, EASN conference 2023, Salerno.

[3] C. Ilic et al., Phased High-Fidelity Aerodynamic Design from Scratch of a Very High-Aspect Ratio Narrow-Body Airliner, AIAA-2024-3838,, AIAA Aviation 2024, Las Vegas. https://doi.org/10.2514/6.2024-3838.

[4] https://www.nextair-project.eu/

[5] https://www.clean-aviation.eu/up-wing

[6] Airbus S.A.S., "A321 aircraft characteristics airport and maintenance planning," manual, Airbus, 31707 Blagnac Cedex, France, 2020. Rev Apr 01/20.

[7] M. Abu-Zurayk et al., Comparing Two Multidisciplinary Optimization Formulations of Trimmed Aircraft Subject to Industry-relevant Loads and Constraints, In: AIAA AVIATION 2021 FORUM. Reston, Virginia: American Institute of Aeronautics and Astronautics, 2021. isbn: 978-1-62410-610-1. https://doi.org/10.2514/6.2021-3075.

[8] G. Ruiz Jimenez, N. Roussouly, F. Gallard, A. Gazaix, J. Brezillon, K. Thanissaranon, V. Brück and J R.R.A. Martins, A Comparison of MDO Bi-Level Formulation Variants With MDF on a High-Fidelity Aeroelastic Wing Test Case, Third ECCOMAS Thematic Conference on Multidisciplinary Design Optimization of Aerospace Systems Conference, Lisbon, Portugal, April 22-24 2025

[9] A. Gazaix, F. Gallard, V. Ambert, D. Guénot, M. Hamadi, S. Grihon, P. Sarouille, T. Y. Druot, J. Brézillon, V. Gachelin, et al. Industrial application of an advanced bi-level mdo formulation to aircraft engine pylon optimization. In AIAA Aviation 2019 Forum, page 3109, 2019.

[10] Y. David, F. Gallard, and A. Rondepierre. Locally Convergent Bi-Level MDO Architectures Based on the Block Coordinate Descent Algorithm. working paper or preprint, Nov. 2024. https://hal.science/hal-04758286.

[11] Gallard, F., Vanaret, C., Guénot, D, et al., GEMS: A Python Library for Automation of Multidisciplinary Design Optimization Process Generation. In : 2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference. 2018. p. 0657. https://arc.aiaa.org/doi/10.2514/6.2018-0657

[12] K. Bobrowski, E. Ferrer, E. Valero, and H. Barnewitz. Aerodynamic Shape Optimization Using Geometry Surrogates and Adjoint Method. AIAA Journal, 55(10):3304–3317, 2017.

[13] D. Schwamborn, T. Gerhold and R. Heinrich, "The DLR TAU-Code: Recent Applications in Research and Industry," ECCOMAS CFD 2006: Proceedings of the European Conference on Computational Fluid Dynamics, European Community on Computational Methods in Applied Sciences (ECCOMAS), Sept. 2006.

[14] R. Olivanti, Multi-Fidelity Aerodynamic Shape Optimization Considering Flexible Adjoint at Multiple Operating Conditions, Diss. Toulouse, ISAE, Dec. 2021, https://depozit.isae.fr/theses/2021/2021_Olivanti_Romain.pdf

[15] R. Olivanti, J. Brézillon, On the Benefits of Engaging Coupled-Adjoint to Perform High-Fidelity Multipoint Aircraft Shape Optimization, AIAA AVIATION 2021 FORUM, 2021

[16] T. Klimmek, M. Schulze, M. Abu-Zurayk, C. Ilic and A. Merle, cpacs-MONA – An independent and in high fidelity based MDO tasks integrated process for the structural and aeroelastic design for aircraft configurations. International Forum on Aeroelasticity and Structural Dynamics 2019, IFASD 2019. IFASD 2019 - International Forum on Aeroelasticity and Structural Dynamics, 10.-13. June 2019, Savannah, GA (USA) – https://elib.dlr.de/128099/

[17] J.C. Giret, F. Gallard, A. Dechaume, J. Fasquel, J.F. Figué and A. Gazaix, Toward Distributed and Scalable Multidisciplinary Optimization with GEMSEO, Third ECCOMAS Thematic Conference on Multidisciplinary Design Optimization of Aerospace Systems Conference, Lisbon, Portugal, April 22-24 2025.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTIDISCIPLINARY DESIGN OPTIMIZATION OF NONPLANAR LIFTING SYSTEMS FOR UNMANNED COMBAT AERIAL VEHICLES

Berkay Pirlepeli[1], Sihmehmet Yildiz[1], Hassan Aleisa[2], Konstantinos Kontis[2], and Melike Nikbay[3]

1: Faculty of Aeronautics and Astronautics
Istanbul Technical University
34469, Istanbul, Türkiye
pirlepeli18@itu.edu.tr, yildizsih@itu.edu.tr
2: Aerospace Engineering Department
University of Glasgow
Glasgow, G12 8QQ, UK
2501058A@student.gla.ac.uk, kostas.kontis@glasgow.ac.uk
3: Astronautical Engineering Department
Istanbul Technical University
34469, Istanbul, Türkiye
nikbay@itu.edu.tr

Abstract. *This paper provides an overview of the AVT-SP-002 research project collaborated between Istanbul Technical University and the University of Glasgow with support from the NATO Science and Technology Organization. The project focused on the design and optimization of non-planar lifting systems for potential applications in civil and military aircraft using both computational and experimental methods across three phases. In the initial phase, preliminary design of a non constant sweep angle lambda wing Unmanned Combat Air Vehicle (UCAV) configuration is aerodynamically optimized using surrogate based low fidelity method at subsonic flow conditions. In the second phase, non planar wingtip extensions are aerodynamically optimized using high fidelity simulations. During the final phase, a multidisciplinary design optimization framework is implemented integrating aerodynamic analysis with radar cross-section prediction to develop non-planar UCAV geometries characterized by low observability and improved aerodynamic efficiency. The project also involved wind tunnel testing of an aerodynamically optimized non-planar wing model to understand the flow physics associated. Wind tunnel tests are conducted by University of Glasgow. The complex vortical flow structures governing UCAV aerodynamic behavior were analyzed across various pitch angles, providing insights into leading-edge vortex formation, development, and breakdown phenomena. The optimized non-planar UCAV model exhibited enhanced aerodynamic performance in terms of the lift and drag characteristics compared to the baseline UCAV design.*

Keywords: Nonplanar lifting systems, UCAV, Multidisciplinary design, Aerodynamic optimization, Wind tunnel testing, Design optimization, CFD

1 INTRODUCTION

The aerodynamic performance and maneuverability of UCAVs have become increasingly important in modern military aviation. UCAVs typically employ tailless flying wing configurations characterized by low to moderate aspect ratios, internal weapon carriage, concealed propulsion systems, and moderate to high leading-edge sweep angles. These design choices aim to balance aerodynamic efficiency with stealth characteristics, presenting significant challenges during the aerodynamic design optimization.

The performance of swept tailless flying wings at medium to high angles of attack (AOA) is heavily influenced by complex three-dimensional vortical flow structures. Unlike conventional two-dimensional airfoils that typically stall at angles of 10-15 degrees, flying wing configurations such as delta, diamond, and lambda wings can maintain aerodynamic performance at significantly higher AOA values [1]. This extended performance envelope is primarily attributed to leading-edge vortex formations that enhance lift generation through local suction pressure near the leading edge[2].

However, the benefits of swept flying wing vortices are limited; at greater AOA values, vortex breakdowns result in an abrupt drop in lift due to stalling. This nonlinear aerodynamic behavior presents both challenges and opportunities for design optimization. The leading-edge sweep angle emerges as a critical parameter affecting the development, strength, and breakdown characteristics of these vortical structures. While a higher sweep angle delays stall inception and enhances lift at higher AOA values, it may deteriorate aircraft controllability at low speeds. Conversely, reduced sweep angles increase leading-edge vortex strength but can trigger vortex breakdown at lower AOA values.

Recent research has increasingly focused on variable planform parameters for lambda wings, particularly non-constant leading-edge sweep configurations. This approach represents a departure from conventional constant leading-edge sweep designs exemplified by demonstrators such as ical studies were performed on UCAV lambda wing models such as UCAV 1303 [3–7], SACCON [8–13] and MULDICON [13–17]. Non-constant leading-edge sweep configurations exhibit complex mixed attached/separated flow fields at high lift conditions, requiring extensive investigation to fully understand their aerodynamic characteristics.

The majority of UCAV designs employ tailless configurations characterized by swept planar wing forms with relatively low to moderate aspect ratios[18]. Recent research indicates that beyond sweep angle optimization, nonplanar wing extensions demonstrate remarkable potential for enhancing aerodynamic efficiency compared to conventional planar designs[19, 20]. These configurations achieve improved performance through reduced total drag coefficients while maintaining equivalent span and lift characteristics. Various nonplanar wing extensions have been proposed across different aircraft classifications, suggesting widespread applicability of this approach[21, 22]. The integration of optimized nonplanar wingtip designs thus represents a promising way for further advancements in UCAV aerodynamic performance.

This project employs a multidisciplinary approach combining computational fluid dynamics (CFD) simulations using open-source software, experimental wind tunnel testing, and optimization methodologies to develop and validate aerodynamically enhanced UCAV configurations. The research addresses questions regarding the influence of constant versus non-constant leading-edge sweep angles, the benefits of nonplanar wingtip extensions, and the complex vortical flow physics governing UCAV aerodynamic behavior.

This investigation spans three interconnected phases that advance our understanding of UCAV aerodynamics and optimization. Phase I examines the fundamental influence of leading-edge sweep angles on vortex flow structures, comparing constant and non-constant leading-edge sweep configurations through surrogate-based optimization and wind tunnel validation. Phase II presents a surrogate-based optimization methodology for designing non-planar wingtips to enhance aerodynamic performance and reports on experimental and computational investigations of a Non-planar Wingtip-Section Optimized UCAV Model at various pitch angles. Finally, Phase III outlines a multidisciplinary design optimization (MDO) framework integrating aerodynamic analysis with radar cross-section (RCO) prediction to create a novel non-planar UCAV geometry with enhanced stealth properties and provides experimental and numerical analyses, validating the aerodynamic characteristics through wind tunnel testing and CFD simulations.

2 PHASE I: NON-CONSTANT SWEEP WING DESIGN OPTIMIZATION

The first phase of the AVT-SP-002 project aims to develop a conceptual configuration of a non-constant leading edge (LE) sweep angle flying wing with optimization to enhance aerodynamic performance. Initially, a scaled down conceptual design is obtained via aerodynamic sizing and planform selection for wind tunnel testing. This scaled model is optimized for low-speed conditions using a Kriging-based surrogate model with vortex-lattice methods to maximize lift-to-drag ratio. A hybrid rapid prototype low-cost wind tunnel model is produced to validate the numerical model and to study how various wingtip deflections affect aerodynamic characteristics. Testing was conducted in the De Havilland low-speed wind tunnel at the University of Glasgow, complemented by numerical simulations to better characterize leading-edge vortex flow patterns, with results validated against the experimental data.

2.1 Methodology

2.1.1 Geometry Definition and Parametrization

Initial aerodynamic sizing and planform selection led to a conceptual UCAV design with a non-constant leading-edge flying planar wing configuration. This full-scale design was subsequently scaled down at a 1:20 ratio to create a wind tunnel model suitable for experimental evaluation. The baseline wing, shown in Fig. 1 is composed of three parts: the fuselage, the mid-section, and the tip-section. The trailing edge (TE) sweep angles (Λ_{TE}) were kept fixed at 40^o for all sections to maintain a fixed span (b). The NACA 64_1212 airfoil is used along the span and kept fixed during optimization.

The wind tunnel experiments are conducted at the University of Glasgow's De Havilland Wind Tunnel facility, which features a test section measuring 2.65m × 2.1m × 5.4m. The facility is capable of generating airflow up to 70 m/s with Reynolds numbers (Re) reaching 4×10^6, though the tests are conducted at 30 m/s. Manufacturing constraints influenced the design decisions. Using Fused Deposition Modeling (FDM) 3D printing technology with a maximum working envelope of 400 mm × 400 mm × 350 mm, the span of the model is limited to 0.75 m. The baseline geometry dimensions are given in Table 1.

The design parameters for optimization are the LE inboard sweep angle (Λ_{LEI}) and the LE outboard sweep angle (Λ_{LEO}) shown in Fig. 1. The geometry was parameterized

Table 1: Baseline Geometry Dimension

Parameter	Description	Value
S, [m^2]	Wing Area	0.155
b, [m]	Wing Span	0.75
l_f, [m]	Length of Fuselage	0.525
c_r, [m]	Wing Root Chord	0.116
Λ_{LEI}, [deg]	LE inboard sweep angle	60
Λ_{LEO}, [deg]	LE outboard sweep angle	40
Λ_{TE}, [deg]	TE sweep angles (all sections)	40

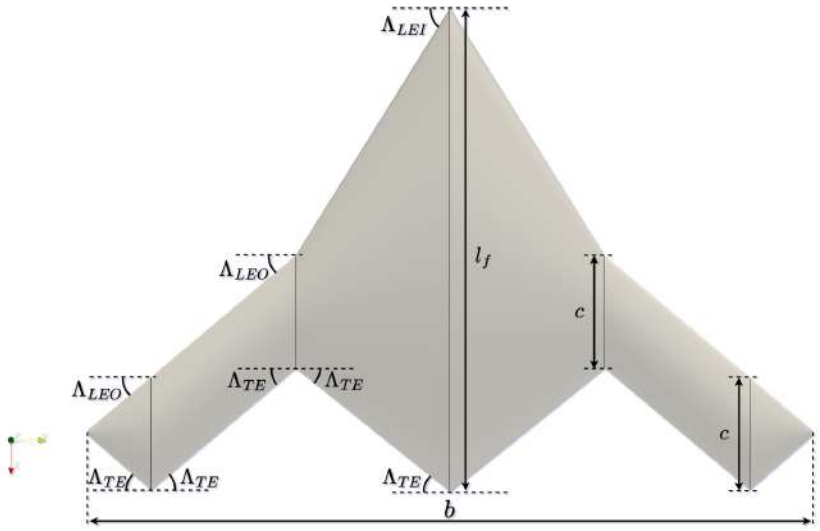


Figure 1: Baseline geometry

and generated using OpenVSP [23].

2.1.2 Aerodynamic Analysis

Aerodynamic performance is evaluated using a low-fidelity linear VLM solver, VSPAERO, developed at NASA Ames Research Center [24] which assumes incompressible and inviscid flow, neglecting thickness effects. The wing surfaces were discretized into 53 panels chord-wise and 29 panels spanwise, with leading-edge clustering of 0.25 and trailing-edge clustering of 0.05. The flow conditions are Ma = 0.087 (approximating incompressible flow), $Re = 10^6$, and altitude $h = 0m$.

2.1.3 Optimization Problem

A surrogate-based optimization approach was employed using Kriging surrogate model for the objective functions with the sample data generated via the Latin Hypercube Sampling (LHS) method [25]. A total of 150 design points were evaluated using the VLM analysis. The surrogates for C_L/C_D and $b/\sqrt{S}$ were cross-validated using 10% of the data, with reported errors shown in Table 2. The constrained multi-objective optimization problem was solved using the Non-dominated Sorting Genetic Algorithm [26], implemented in the Pymoo library [27].

Table 2: Surrogate model error for C_L/C_D and $b/\sqrt{S}$

Quality metric	C_L/C_D	$b/\sqrt{S}$
Root mean squared (Training Points)	1.2090×10^{-14}	4.0087×10^{-14}
Root mean squared (Test Points)	1.4064×10^{-06}	3.6504×10^{-05}

The optimization problem aimed to improve aerodynamic efficiency, represented by maximizing C_L/C_D and minimizing the span factor $b/\sqrt{S}$. The formulation is given by:

$$\begin{aligned} &\min_{\boldsymbol{\Lambda_{LEI},\Lambda_{LEO}}} \quad -\frac{C_L(s)}{C_D(s)} \\ &\min_{\boldsymbol{\Lambda_{LEI},\Lambda_{LEO}}} \quad \frac{b}{\sqrt{S}} \\ &\textbf{subject to} \quad h_1(s) = \frac{b}{b_{ref}} - 1 = 0, \\ &\qquad\qquad h_2(s) = \frac{l_f}{l_{f_{ref}}} - 1 = 0 \end{aligned} \tag{1}$$

where s represents the design vector $[\Lambda_{LEI}, \Lambda_{LEO}]$, $b_{ref} = 0.75\ m$ and $l_{f_{ref}} = 0.525\ m$. The design variables and their bounds are given in Table 3.

Table 3: Optimization Parameters

Parameters	Description	Interval	Initial value
Λ_{LEI} [deg]	LE inboard sweep angle	[50, 60]	60
Λ_{LEO} [deg]	LE outboard sweep angle	[30, 50]	40

2.1.4 Experimental Method

The experimental studies were carried out in the De Havilland low-speed wind tunnel facility at the University of Glasgow (Fig. 2). It is a closed return atmospheric low-speed wind tunnel with a maximum speed and Re of 70 m/s and 4×10^6, respectively. The test model was mounted on a sting that adjusted the yaw, incidence, and roll angles and could be operated and remotely controlled from the wind tunnel control room. It comprises an octagonal working section with a contraction ratio of 5:1. The inlet turbulence intensity is less than 1%. The wind tunnel has a test section of 2.65 m × 2.1 m × 5.4 m and a platform of 6-axes load cells beneath the working section with various capacity load cells. The experimental test was conducted on the UCAV model at the AOA flow range of 0^o – 40^o with data being recorded every 5^o, and the tunnel speed was set at 30 m/s for the balance measurements with a Re of 5.78×10^5 based on the model mean aerodynamic chord. This research was performed to estimate the aerodynamic characteristics of a UCAV model experimentally.

2.2 Results

The sweep angle optimization identified optimal leading edge sweep angles of $\Lambda_{LEI} = 58^o$ and $\Lambda_{LEO} = 40^o$. The low-fidelity VLM optimization predicted a modest 1.3% im-

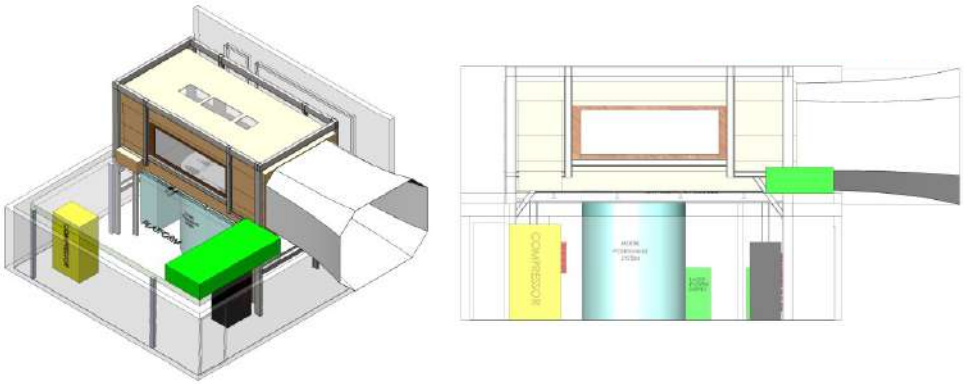


Figure 2: Isometric view (left) and side view of the de Havilland wind tunnel (right)

provement in the lift-to-drag ratio (C_L/C_D). Table 4 summarizes the C_L/C_D improvements. Fig. 3 shows a comparison between the baseline and optimized wing geometries.

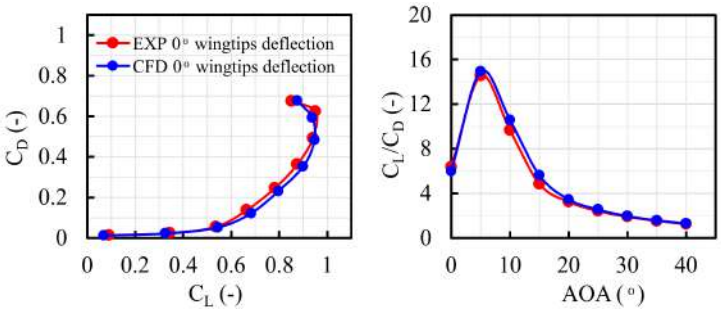
Table 4: Optimization results

Configurations	C_L/C_D
Baseline	4.8221
Optimized	4.8868
Improvement [%]	1.3

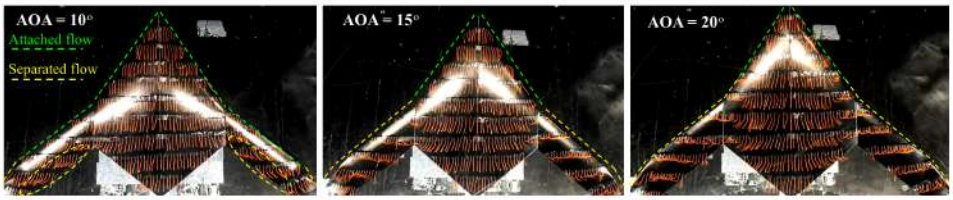


Figure 3: Comparison between baseline wing geometry (black) and optimized wing geometry (white)

Experimental validation of the Phase I non-constant sweep design ($\Lambda_{LEI} = 58°$, $\Lambda_{LEO} = 40°$) was performed using a 1:20 scale FDM model (0.75 m span) in the University of Glasgow's De Havilland low-speed wind tunnel. Aerodynamic loads and surface flow patterns were investigated at 30 m/s (Re $\approx 5.78 \times 10^5$) across an AoA range of 0° to 40° using a balance system and tuft visualization, respectively.



(a) Validation of aerodynamic coefficients[28].



(b) Tuft flow visualization at various AoA.

Figure 4: Experimental results for the Phase I optimized non-constant sweep UCAV[28].

The experimental results, summarized in Fig. 4, provided crucial validation for the optimized design. The measured aerodynamic coefficients (Fig. 4(a)) confirmed the predicted performance trends, aligning well with computational data. Concurrently, tuft visualization (Fig. 4(b)) offered qualitative insight into the flow physics, illustrating the progression from attached flow to flow separation on the outboard wing sections at increasing AoA (approx 15° - 20°). These tests validated the Phase I non-constant sweep concept and established a verified baseline for subsequent development phases. For further details on experimental methodology, readers are referred to previous work [28]

3 PHASE II: OPTIMIZATION OF A NON-PLANAR WINGTIP CONFIGURATION FOR THE LAMBDA WING UCAV

This phase focuses on designing non-planar wingtip extensions for the lambda wing UCAV optimized in Phase I, aiming to further enhance aerodynamic efficiency. The constrained optimization problem solved via surrogate-based derivative-free genetic optimization method.

3.1 Methodology

3.1.1 Geometry Definition and Parametrization

The baseline planar wing geometry for this phase is the optimized non-constant leading-edge sweep flying wing from Phase I ($\Lambda_{LEI} = 58^o$, $\Lambda_{LEO} = 40^o$, NACA 64_1212 airfoil) shown in Fig. 3.

The parametric non-planar wingtip model is generated with using OpenVSP [23]. The design parameters are selected as tip sweep angle, tip dihedral and tip blending strength. The design variables with bounds are shown in Table 5.

Table 5: Design Variables

Parameter	Initial	Interval
Tip Sweep Angle [deg]	40	[30, 60]
Tip Dihedral Angle [deg]	0	[-30, 30]
Tip Blending Strength	1	[0.5, 1.5]

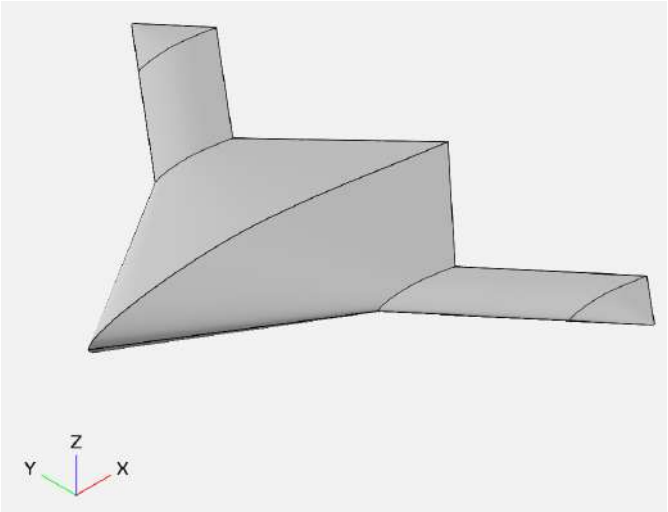


Figure 5: Baseline geometry of the planar UCAV

3.1.2 Aerodynamic Analysis

Aerodynamic performance is analyzed using the SU2 CFD solver [29]. RANS simulations were performed using the Shear-Stress Transport turbulence model (SST) k-ω turbulence model. A semi-span model was used within a semi-spherical computational domain (diameter 10 times root chord length). To capture near-wall effects, 30 prism layers were employed with an initial normal wall spacing of 1×10^{-5}m and a layer expansion ratio of 1.2 (Fig. 6). Boundary conditions included far-field, symmetry plane, and constant heat flux no-slip walls on the UCAV surface. Convective fluxes were discretized using the second-order Lax-Friedrich central scheme, and turbulent convective terms used a second-order scalar upwind technique.

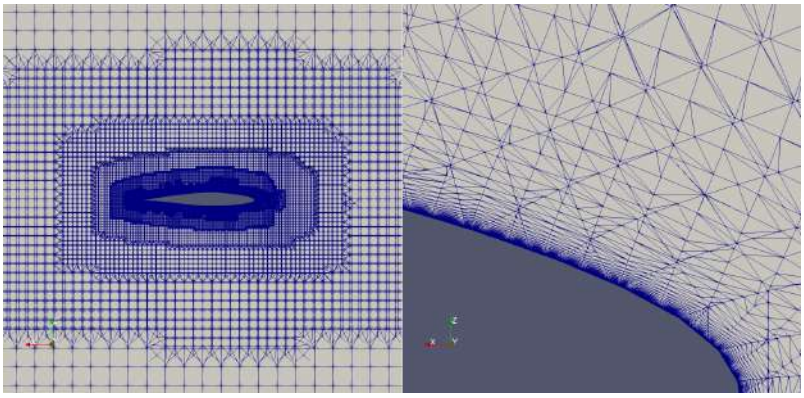


Figure 6: Aerodynamic mesh

The flow conditions simulated the wind tunnel testing environment, listed in Table 6.

Table 6: Flow conditions

Parameter	Value
Mach Number [-]	0.0874
Altitude [m]	0
Dynamic Pressure $[kg/ms^2]$	1488.68
Reynolds Number [-]	1.5×10^6

A mesh sensitivity analysis was conducted using three meshes (coarse, medium, fine) characterized by average cell length h. Richardson extrapolation was used to estimate the zero-grid-spacing solution. Relative errors between meshes were less than 2%. The medium mesh was found to offer sufficient accuracy for the objective function assessment with moderate computational cost (See Table 7).

Table 7: L/D values of three different mesh

	h	# of Cells	r	L/D	Wall Clock Time[hr]
Fine	0.3385	25008043	1.74	5.8678	5
Medium	0.5893	16275991	1.34	5.7706	2
Coarse	0.7889	8883917		4.8340	1.30

3.1.3 Optimization Algorithm

A surrogate model-based optimization procedure is employed. Incremental Latin Hypercube sampling was used to generate data points (50 samples total, 45 for training, 5 for validation). A Kriging surrogate model (Equations 2-5) was constructed to approximate the objective function (C_L/C_D) based on the high-fidelity CFD results. The process involved hyperparameter optimization and validation (Fig. 7).

$$\hat{f} = \mathbf{w}^T \boldsymbol{\psi} \tag{2}$$

$$\psi_i = \exp\left(-\sum_{j=1}^{k} \theta_j |x_j^i - x_j|^{p_j} \right) \tag{3}$$

$$\mathrm{cor}[Y(x^i), Y(x^k)] = \exp\left(-\sum_{j=1}^{k} \theta_j |x_j^i - x_j^k|^{p_j} \right) \tag{4}$$

$$\mathrm{Cov}[\mathbf{Y}, \mathbf{Y}] = \sigma^2 \boldsymbol{\Psi} \tag{5}$$

The parameters θ_j and p_j were optimized to maximize the likelihood of the evaluations. Parametric geometry generation was automated using the OpenVSP scripting API (Angelscript). An automated workflow was developed using Python scripts to integrate model generation, meshing, CFD evaluation (SU2), and surrogate model construction. A derivative-free genetic optimization algorithm was used to find the optimal set of design variables (Tip Sweep, Dihedral, Blending Strength).

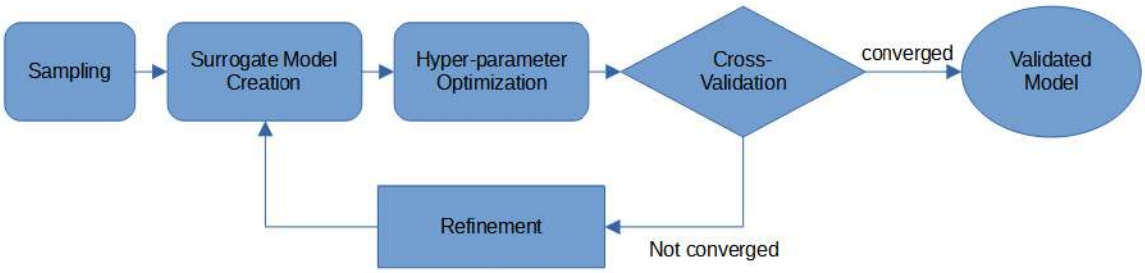


Figure 7: Surrogate model generation process

The optimization aimed to maximize aerodynamic efficiency (C_L/C_D) subject to geometric constraints (fixed span and fuselage length). The formulation is:

$$\begin{aligned} \min_{s \in S} \quad & -\frac{C_L(s)}{C_D(s)} \\ \text{subject to} \quad & h_1(s) = \frac{b}{b_{ref}} - 1 = 0, \\ & h_2(s) = \frac{l_f}{l_{f_{ref}}} - 1 = 0 \end{aligned} \tag{6}$$

where s is the design vector (Tip Sweep, Tip Dihedral, Tip Blending Strength), S is the feasible design space defined by the bounds in Table 5, and b_{ref}, $l_{f_{ref}}$ are the reference span and length from the baseline geometry.

3.1.4 Experimental Method

Experimental validation is conducted for the non-planar wingtip UCAV configuration developed in Phase II. This involved testing a physical model, which integrated the FDM-printed optimized wingtips onto the optimum geometry from Phase I, within the University of Glasgow's de Havilland low-speed wind tunnel. Utilizing the wind tunnel's balance system for aerodynamic load measurements and surface tufts for flow visualization, the model was assessed at 30 m/s across a range of angles of attack (AoA).

3.2 Results

The optimization process resulted in a winglet-type non-planar wingtip extension. The optimized geometric parameters are given in Table 8. This optimized configuration achieved an approximately 4% improvement in aerodynamic efficiency (L/D) compared to the planar baseline wing, as shown in Table 9. Fig. 8 shows the optimized wingtip geometry.

The optimized wingtip features a significantly increased sweep angle (from 40° to 56.4°), maximum upward dihedral (30°, reaching the constraint boundary), and reduced blending strength (0.62). This non-planar extension effectively mitigates the pressure gradient at the wingtip. The fact that the dihedral angle reached its upper bound suggests that exploring a larger interval for this parameter might yield further improvements.

Figure 9 presents the experimentally measured lift-to-drag ratio (CL/CD) comparison between the planar baseline model and the optimized non-planar wingtip configura-

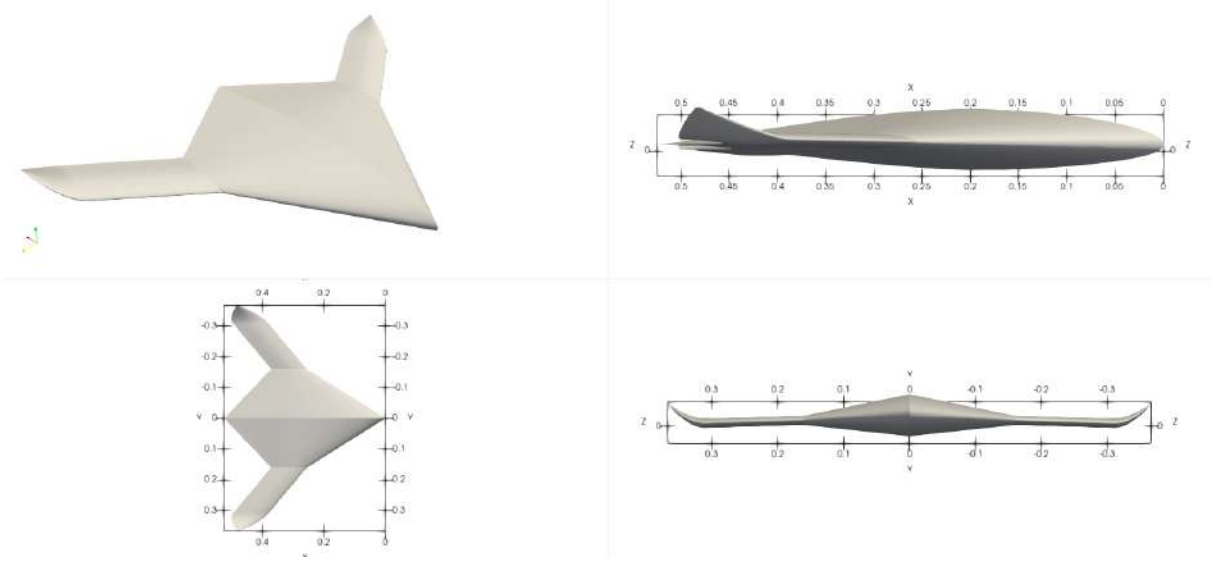


Figure 8: Optimized wingtip geometry

Table 8: Best parameters

Parameter	Initial	Interval	Best
Tip Sweep Angle [deg]	40	[30,60]	56.4
Tip Dihedral [deg]	0	[-30,30]	30
Tip Blending Strength	1	[0.5,1.5]	0.62

Table 9: Optimization Results

Configuration	L/D
Baseline	4.88
Optimized wingtip	5.08
Improvement (%)	4

tion. Enhanced aerodynamic efficiency achieved with the non-planar tips, particularly noticeable at lower and moderate angles of attack, aligning with the optimization goal of approximately 4% improvement reported computationally.

The surface flow patterns at various AoA are provided by the tuft flow visualizations shown in Fig. 10. At lower AoA (up to 10-15°), the flow appears largely attached across most of the wing. As AoA increases (e.g., 20° and beyond), regions of flow separation become evident, typically initiating near the wingtips and potentially progressing inboard, consistent with the performance characteristics observed in the CL/CD curves. These experimental investigations validated the aerodynamic improvements predicted for the Phase II non-planar wingtip design under low-speed wind tunnel conditions. For further details on experimental methodology, readers are referred to previous work [30]

4 PHASE III: MULTI-DISCIPLINARY DESIGN OPTIMIZATION OF A NOVEL NON-PLANAR UCAV MODEL

In this section, the latest phase of the AVT-SP-002 project which is based on the multi-disciplinary design optimization process for a novel UCAV geometry. The study aims to present a non-planar UCAV geometry with low-observable features using the

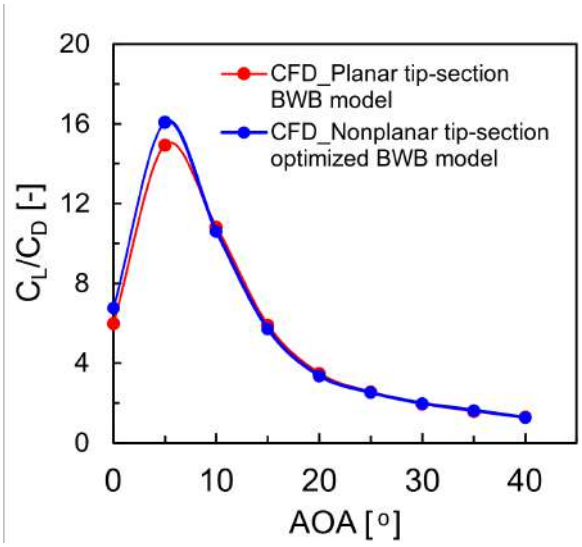


Figure 9: Experimental CL/CD comparison[30].

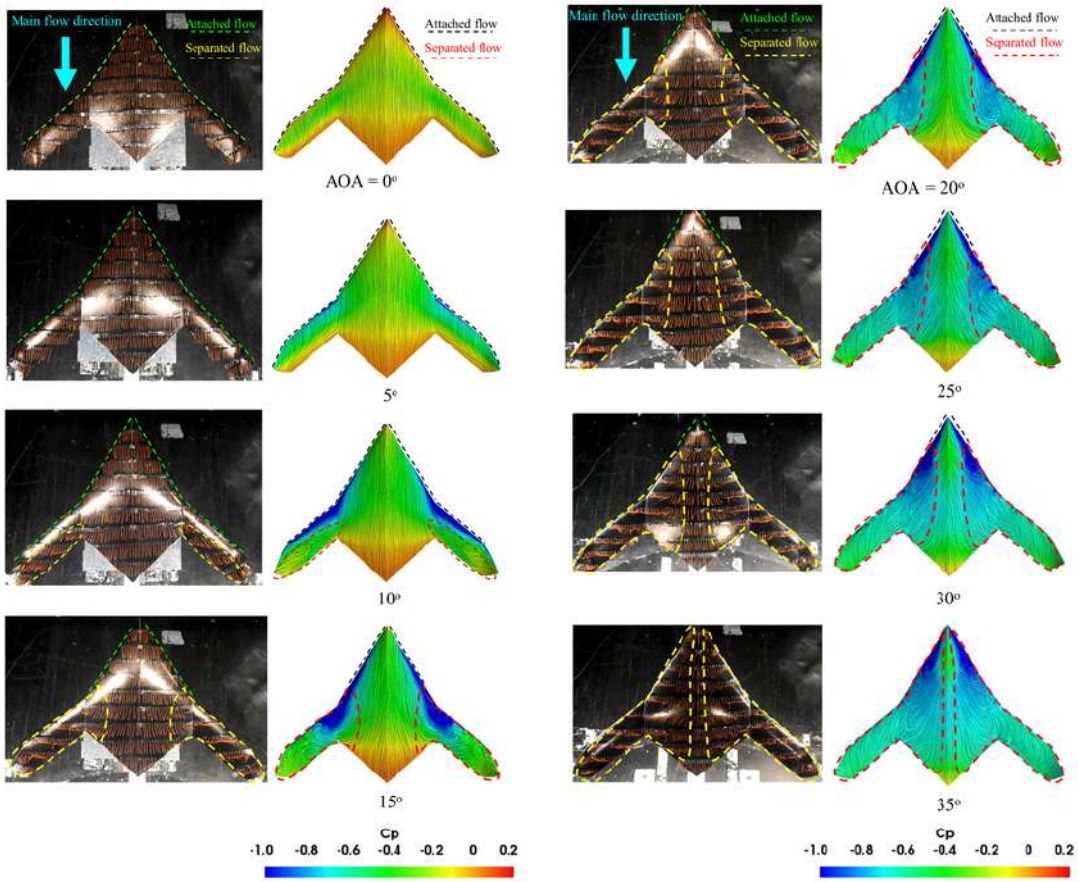


Figure 10: Tuft flow visualization illustrating surface flow patterns at various AoA[30].

fully non-planar flying wing concept.

4.1 Methodology

This section outlines the tools used for geometry generation, meshing, radar cross-section computation, and aerodynamic analysis. The software module was developed in

Bash and Python for their functionality and seamless task integration. Figure 11 presents the workflow, illustrating the sequential steps of the computational process.

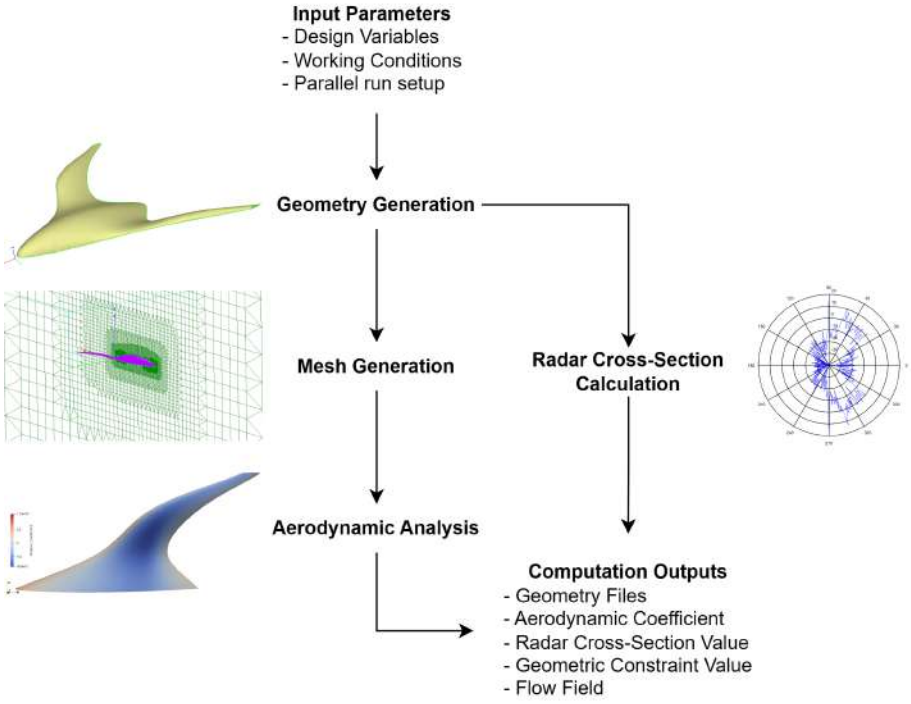


Figure 11: Flow diagram for analysis process

4.1.1 Geometry Definition and Parametrization

In the study, the baseline geometry was derived from the conceptual design developed during the initial phase of the project.

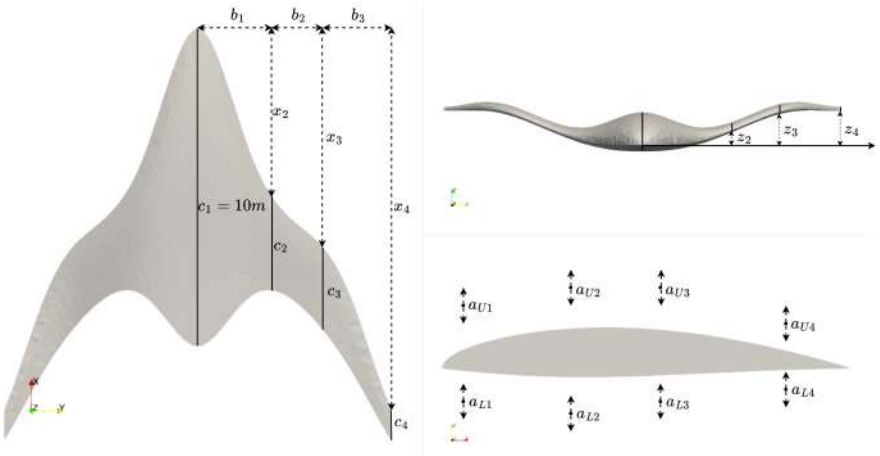


Figure 12: Non-Planar Wing Parameters

The aircraft geometry is parameterized with 20 variables, using a base geometry as reference. To achieve sufficient lift with a smaller wing area, 8 variables define the airfoil

shape via the Kulfan method. The remaining 12 variables control non-planarity by setting the chord length and section positions along the x, y, and z axes. Figure 12 illustrates the geometric parameters used during optimization.

Table 10: Design variable bounds

Description	Symbol	Interval	Units
Control points for airfoil upper surface	a_{U1}, a_{U2}, a_{U3}, a_{U4}	[0.01, 0.35]	[−]
Control points for airfoil lower surface	a_{L1}, a_{L2}, a_{L3}, a_{L4}	[-0.01, -0.15]	[−]
Chord Lengths	c_2	[2.0, 5.0]	$[m]$
	c_3	[2.0, 4.0]	$[m]$
	c_4	[0.2, 1.0]	$[m]$
Span Lengths	b_1	[1.5, 3.5]	$[m]$
	b_2	[1.5, 5.0]	$[m]$
	b_3	[0.5, 2.5]	$[m]$
x-Location of Sections	x_2	[2.0, 8.0]	$[m]$
	x_3	[2.0, 10.0]	$[m]$
	x_4	[2.0, 12.0]	$[m]$
z-Location of Sections	z_2	[-0.5, 0.5]	$[m]$
	z_3	[-1.0, 1.0]	$[m]$
	z_4	[-2.0, 2.0]	$[m]$

Table 10 shows the lower and upper limits of the 20 design variables used for geometry parametrization, ensuring the initial geometry stays within the defined design space.

4.1.2 Aerodynamic Analysis

In this study, the SU2 code [29], developed by the Aerospace Design Laboratory (ADL) at Stanford University, is used as the high-fidelity analysis solver.

Table 11: Flow conditions

Parameter	Value
Altitude [m]	11000
Mach Number [-]	0.8
Reynolds Number [-]	14×10^6

For the simulations, SU2 version 7.5.0 is used, with the SST turbulence model to accurately capture turbulence effects. The flow conditions during cruise flight, provided in Table 11. To ensure accurate calculation of the conservation of quantities, a second-order Monotonic Upstream-centered Scheme for Conservation Laws (MUSCL) reconstruction method is applied. The convective fluxes are computed using the Lax-Friedrich scheme, while the FGMRES (Flexible Generalized Minimal RESidual) algorithm is employed to solve the linear system of equations. A numerical error tolerance of 10^{-10} is set as the convergence criterion.

4.1.3 Radar Cross Section Calculation

Radar cross section (RCS) quantifies an object's detectability by measuring how effectively it reflects incoming electromagnetic waves back to a radar receiver. RCS values depend on factors such as the object's size, shape, surface properties, and the characteristics of the radar waves. Higher RCS values increase detectability, whereas lower RCS values enhance low observability capabilities[31].

The POFACETS software (version 4.5) from the US Naval Postgraduate School is used for RCS signature simulations[32]. POFACETS uses the physical optics approximation technique to calculate the RCS of complex geometries. This method determines the directional scattered power from an illuminated target by normalizing it against the incident field's power density. Monostatic analyses are performed at 10 GHz, with surface roughness neglected, a region length set to 10^{-5}, and 5 Taylor series terms used in the calculations.

In radar terminology, the elevation angle (θ) is defined as the angle between the radar and the target[31]. An elevation angle of 0° indicates that the target is directly above the radar, while an angle of 90° means the radar and target are in the same horizontal plane. Elevation angles greater than 90° occur when the radar is positioned above the target. For optimization, the elevation angle was fixed at 90° to simplify RCS calculations and concentrate on parameter analysis.

4.1.4 Optimization Algorithm - Constrained Efficient Global Optimization

Efficient Global Optimization (EGO) is a Bayesian optimization algorithm commonly used for black-box optimization problems. Its main goal is to identify the most promising locations based on an infill sampling criterion. This iterative process uses a probabilistic surrogate model to make predictions and estimate uncertainties. Sampling continues iteratively until a predefined stopping criterion is met. Each selected point is evaluated using the black-box model, and after each evaluation, the surrogate model is updated with the newly obtained training data (see Fig. 13).

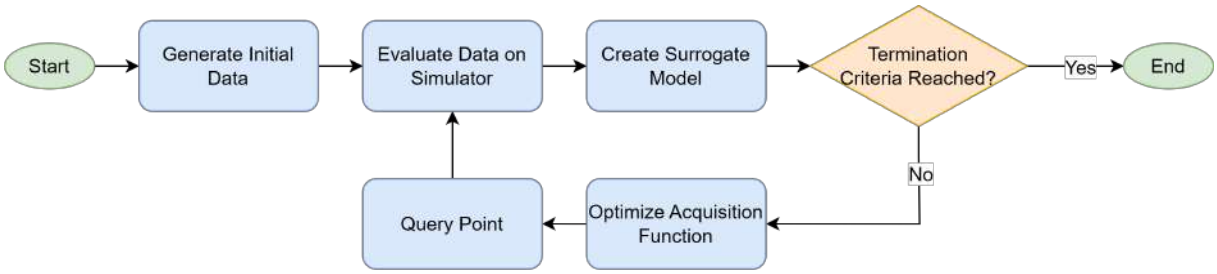


Figure 13: Flow diagram of optimization algorithm

The surrogate model and infill sampling criterion are the two key components of the EGO algorithm. The following subsections explain the surrogate model and the infill sampling criterion.

Kriging Partial Least Squares Surrogate Modeling

Kriging Partial Least Squares (KPLS), introduced by Bouhlel et al. [33], combines Kriging with Partial Least Squares (PLS) to reduce the number of hyperparameters while maintaining high estimation accuracy. KPLS is especially useful for large-dimensional

problems, leveraging the PLS structure created by the kernel. The method uses a modified correlation function of inputs, with PLS coefficients calculated through the linear dependence between input and output. The KPLS is a stochastic surrogate modeling method, similar to the Kriging method, used to approximate an unknown computational model ($M(x)$) by utilizing a set of input-output pairs $D = \{(x_i, y_i)\}_{i=1}^{N}$. The computational model $M(x)$ is modeled by a Gaussian process with zero mean and r specific variances, as shown in Eq. (7).

$$M(x) \sim \mathcal{GP}\left(0, r\left(x, x'; \theta\right)\right) \tag{7}$$

In Equation (7), r is the covariance function and modified covariance function for KPLS is expressed in Eq. (8).

$$r\left(x, x'; \theta\right) = \alpha^2 \prod_{l=1}^{h} \prod_{i=1}^{d} \exp\left(\frac{w_i^l \left(x_i - x_i'\right)^2}{\beta_l^2}\right) \tag{8}$$

where $\theta = (\alpha, \beta)$ are hyper-parameters. Here, $d,\ h,\ w$ show the problem dimension, PLS component number and, PLS coefficient respectively. The hyperparameters are obtained by minimizing the negative log likelihood function given in Eq. (9), where R is the covariance matrix and expressed as $\mathbf{R} = \boldsymbol{r}(\mathbf{x}, \mathbf{x}; \boldsymbol{\theta}) + \boldsymbol{\sigma}^2\mathbf{I}$.

$$\boldsymbol{NLL}(\theta) = \frac{1}{2}\mathbf{y}^T\mathbf{R}^{-1}\mathbf{y} + \frac{1}{2}\log\left|\mathbf{R}\right| + \frac{N}{2}\log(2\pi) \tag{9}$$

Employing the hyper-parameters estimated using the training dataset, the surrogate model's prediction and predictive variance values are then determined at a new test point (x^*) using Eq. (10) and (11).

$$\widehat{\mathrm{M}}\left(x^*\right) = r\left(x^*, \mathbf{x}\right)\mathbf{R}^{-1}\mathbf{y} \tag{10}$$

$$\hat{\sigma}^2(x^*) = r\left(x^*, x^*\right) - r\left(x^*, \mathbf{x}\right)\mathbf{R}^{-1}r\left(\mathbf{x}, x^*\right) \tag{11}$$

Acquisition Function

Jones et al. [34] introduced the Expected Improvement (EI) acquisition function as a key element of the Efficient Global Optimization (EGO) framework, demonstrating its effectiveness in navigating complex design spaces [35]. The EI function balances exploration and exploitation by evaluating the design space relative to the best observed design, accounting for both predicted values and uncertainty:

$$EI(\boldsymbol{x}) = \left(y_{\min} - \hat{y}(x)\right)\Phi\left(\frac{y_{\min} - \widehat{M}(\boldsymbol{x})}{\hat{\sigma}(\boldsymbol{x})}\right) + \hat{\sigma}(\boldsymbol{x})\phi\left(\frac{y_{\min} - \widehat{M}(\boldsymbol{x})}{\hat{\sigma}(\boldsymbol{x})}\right) \tag{12}$$

Here, ϕ and Φ are the standard normal probability density function (PDF) and cumulative distribution function (CDF), respectively, $y_{\min}$ is the best observed value, and $\hat{\sigma}(x)$ is the surrogate model's estimated standard deviation.

For constrained optimization problems, the Expected Improvement with Constraints (EIC) function refines EI by incorporating the probability of satisfying constraints:

$$EIC(\boldsymbol{x}) = EI(\boldsymbol{x}) \prod_{k=1}^{K} \Pr\left(g_k(\mathbf{x}) \leq 0\right) \tag{13}$$

Here, $\Pr(g_k(\boldsymbol{x}) \leq 0)$ estimates the likelihood of satisfying each constraint. This adjustment enables the algorithm to focus on promising, feasible regions, enhancing efficiency in constrained design spaces [36].

4.2 Multi-disciplinary Design Optimization Study for Non-planar UCAV

This section presents a multidisciplinary design optimization study for a novel non-planar UCAV based on a flying wing configuration with low-observable capabilities. The following subsections will discuss the optimization problem, the dataset generation for the base surrogate model, and conclude with the results obtained from Constraint Efficient Global Optimization (CEGO).

4.2.1 Optimization Problem Formulation

In this optimization study, the main goal is to design a novel non-planar geometry with the highest possible aerodynamic efficiency. However, the final design should not have lower lift or higher drag than the base geometry. It is also important for the optimal geometry to maintain low radar observability. Additionally, to ensure the manufacturability, the airfoil thickness must be at least 0.12. Therefore, the optimization study considers the lift and drag of the base geometry, airfoil thickness, and radar cross-section as constraints, as shown in Eq. (14). The optimization problem is defined using the optimization variables $\boldsymbol{s}$ as follows:

$$
\begin{aligned}
\underset{\boldsymbol{s}\in\boldsymbol{S}}{\textbf{min}} \quad & -\frac{C_L(\boldsymbol{s})}{C_D(\boldsymbol{s})} \\
\textbf{subject to} \quad & g_1(\boldsymbol{s}) = D(\boldsymbol{s}) - D_{base} \leq 0, \quad g_1(\boldsymbol{s}) \in \Re \\
& g_2(\boldsymbol{s}) = L_{base} - L(\boldsymbol{s}) \leq 0, \quad g_2(\boldsymbol{s}) \in \Re \\
& g_3(\boldsymbol{s}) = RCS(\boldsymbol{s}) - 0.01 \leq 0, \quad g_3(\boldsymbol{s}) \in \Re \\
& g_4(\boldsymbol{s}) = 0.12 - t/c(\boldsymbol{s}) \leq 0, \quad g_4(\boldsymbol{s}) \in \Re \\
& S = \{\boldsymbol{s} \in \Re,\ \boldsymbol{s}_L \leq s \leq \boldsymbol{s}_U\}
\end{aligned}
\tag{14}
$$

4.2.2 Optimization Results

Constrained Efficient Global Optimization was carried out using an initial surrogate model trained with 500 data points. The optimization process ran for 100 iterations. In each iteration, the candidate geometry chosen by the Constraint Expected Improvement Criterion (EIC) was added to the training dataset. The convergence plot of the algorithm is shown in Fig. 14.

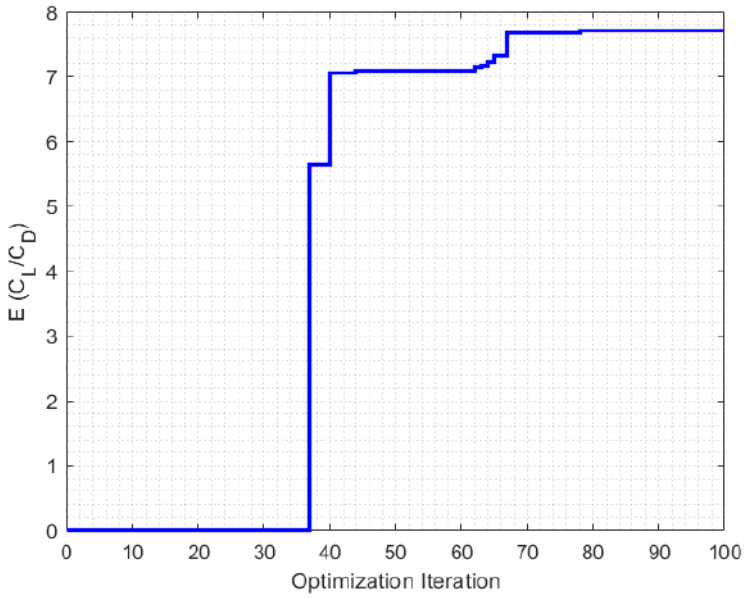


Figure 14: Convergence history

Since the initial training dataset lacked geometries that satisfied the constraints, the CEGO algorithm included non-conforming geometries in the dataset for the first 36 iterations to explore the design space and constraints. After this exploration phase, the CEGO algorithm found a feasible solution in the 37^{th} iteration and gradually improved the geometry until the 78^{th} iteration. However, no further improvement was observed from the 78^{th} to the 100^{th} iteration. Considering that the optimization algorithm likely reached the global optimum by the 100^{th} iteration, the optimization process was terminated.

Table 12: Optimum and base geometry results

Metric	**Base Geometry**	**Optimum Geometry**	**Improvement (%)**
$E\ (C_L/C_D)$	5.3394	7.7182	44.55
C_L	0.2958	0.3384	14.40
C_D	0.0554	0.0438	20.93
$RCS(m^2)$	0.0005	0.0082	< 0.01

The performance values of the optimal geometry and the initial geometry from the optimization study are shown in Table 12. The optimal geometry meets the radar visibility constraint and shows a 44.55% improvement in aerodynamic efficiency compared to the initial geometry. Fig. 15 displays different views of the optimal geometry.

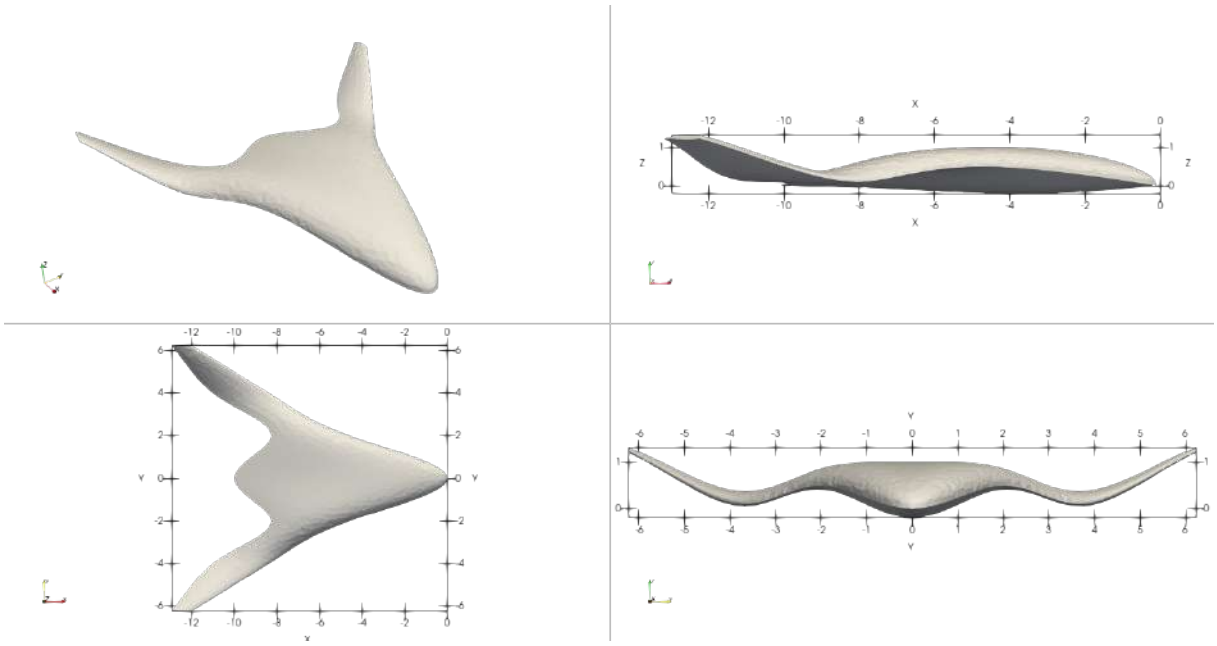


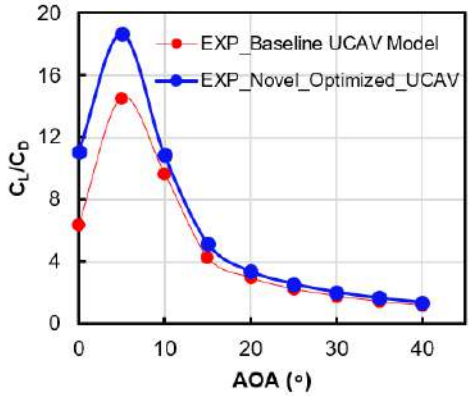
Figure 15: Optimum geometry

4.3 Experimental Results

The experimental investigations were carried out to validate the computational predictions as well as to assess the performance of the final optimized non-planar UCAV geometry developed in Phase III. A rapid prototype model constructed according to the optimized design was fabricated using Fused Deposition Modeling and subsequently tested in the de Havilland low-speed wind tunnel facility at the University of Glasgow. The experimental campaign involved two main test methodologies. First, quantitative measurements of the aerodynamic loads were taken with an accurate sinking balance system; secondly, qualitative surface flow visualization was captured using the tufting technique. The experiments were carried out at a freestream of 30 m/s for a large range of AoAs and yaw angles to simulate relevant subsonic flight.

The key findings from the experimental campaign are summarized in Fig. 16. The aerodynamic load measurements, including the lift-to-drag ratio (CL/CD) comparison shown in Fig. 16(a), provide strong validation for the MDO framework. The data demonstrates that the optimized non-planar UCAV geometry exhibits an enhanced aerodynamic efficiency compared to the baseline UCAV model across the entire tested AoA range. This improvement is particularly pronounced at moderate angles of attack, which are often critical for mission performance.

In addition to the force measurements, tuft flow visualization (shown in Fig. 16(b)) offered insights into the complicated 3D flow physics. The tuft patterns illustrated the evolution and movement of leading-edge vortices, areas of attached and separated flow, and the initiation of flow breakdown at high angles of attack. Understanding these flow phenomena is significant as they relate directly to the measured lift, drag, and moment characteristics. The visualization confirmed the complicated vortical flow structures that are inherent to such highly swept, non-planar configurations and provided very useful qualitative data which supports the quantitative measurements.



(a) Experimental CL/CD comparison.



(b) Wind tunnel setup and tuft visualization.

Figure 16: Experimental validation of the Phase III optimized non-planar UCAV: (a) Comparison of lift-to-drag ratio against the baseline model, showing enhanced performance. (b) Model installed in the wind tunnel showing setup for aerodynamic load measurement and tuft flow visualization.

5 CONCLUSION

The paper presented an overview of the research that was conducted as part of the NATO STO AVT-SP-002 project that was a collaboration between Istanbul Technical University and the University of Glasgow and the MDO of non-planar lifting systems for UCAVs. The study integrated CFD using several levels of fidelity and experimental wind tunnel testing during three phases to develop and validate advanced UCAV configurations. The initial phase concentrated on optimizing a non-constant leading-edge sweep lambda wing UCAV using low-fidelity surrogate-based methods, demonstrating a significant aerodynamic efficiency (L/D) improvement of approximately 6% confirmed via RANS and establishing a validated baseline through experimental tests involving force measurements and flow visualization, which provided crucial insights into the complex vortical flow structures at varying angles of attack.

In phase two, the development of a non-planar wingtip extension was based upon the geometric optimization in phase one, specifically the aerodynamic optimization of non-planar wingtip extensions with high-fidelity RANS simulations in conjunction with a surrogate-based optimization process. The result of the aerodynamic optimization was a winglet-type non-planar configuration with a higher tip sweep and higher dihedral, which yielded an additional computational aerodynamic performance improvement of approximately 4% beyond the optimized planar baseline. The computational aerodynamic predictions were validated experimentally in the de Havilland wind tunnel at the University of Glasgow, and the experiment confirmed the performance improvement of the non-planar wingtips, especially at low to moderate angles of attack for mission performance, and helped to clarify the changes in the flow physics due to the influence of the non-planar wingtips using tuft visualization.

The final phase initiated a MDO framework, coupling high-fidelity aerodynamic analysis (SU2) with RCS prediction (POFACETS) to develop a novel, non-planar UCAV geometry focused on both high lift to drag ratio and low observable characteristics. The MDO process applied a CEGO approach utilizing a KPLS surrogate model to effectively

reach to an optimal design that met or exceeded defined aerodynamic, geometric, and RCS constraints. The final optimized non-planar UCAV showed 44.55% improvement in aerodynamic efficiency over our initial baseline geometry for this final phase and met the defined low observability requirements. Additionally, experimental evaluation of the final design confirmed that the novel non-planar UCAV geometry significantly improved aerodynamic performance across a broader range of angles of attack, when compared to the baseline UCAV model under low-speed conditions. Overall, this research demonstrates the efficacy of integrated computational and experimental MDO frameworks to design next generation UCAVs with enhanced aerodynamic performance and optimally tailored low-observable features through the use of non-planar lifting system concepts.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the NATO Science and Technology Organization's Support Project AVT-SP-002 titled "Turbulence and the Aerodynamic Optimization of Non-planar Lifting Systems."

Computing resources used in this work were provided by the National Center for High Performance Computing of Turkey (UHeM) under grant number 1008262020.

REFERENCES

[1] H. Aleisa, K. Kontis, and M. Nikbay. Predictions of the low-speed aerodynamic characteristics of generic ucavs. In *AIAA SCITECH 2022 Forum*, 2022. doi:10.2514/6.2022-0428. URL `https://arc.aiaa.org/doi/abs/10.2514/6.2022-0428`.

[2] C. W. P. Ford and H. Babinsky. Lift and the leading-edge vortex. *Journal of Fluid Mechanics*, 720:280–313, 2013.

[3] S. McParlin, R. Bruce, A. Hepworth, and A. Rae. Low speed wind tunnel tests on the 1303 ucav concept. In *24th AIAA Applied Aerodynamics Conference*. American Institute of Aeronautics and Astronautics, June 2006. doi:10.2514/6.2006-2985. URL `http://dx.doi.org/10.2514/6.2006-2985`.

[4] K. Petterson. Cfd analysis of the low-speed aerodynamic characteristics of a ucav. In 44^{th} *AIAA Aerospace Sciences Meeting and Exhibit, 09 - 12 January 2006, Reno, Nevada*, page 1259, 2006, AIAA Paper 2006-1259. doi:https://doi.org/10.2514/6.2006-1259.

[5] K. Petterson. Low-speed aerodynamic and flowfield characteristics of a ucav. In 24^{th} *Applied Aerodynamics Conference. San Francisco, California 5-8 June*, AIAA Paper 2006-2986. doi:https://doi.org/10.2514/6.2006-2986.

[6] S. Woolvin. Ucav configuration & performance trade-offs. In 44^{th} *AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, 09 - 12 January, 2006*, page 1264, AIAA Paper 2006-1264. doi:https://doi.org/10.2514/6.2006-1264.

[7] H. Shim and S. O. Park. Low-speed wind-tunnel test results of a bwb-ucav model. *Procedia Engineering*, 67:50–58, 2013. ISSN 1877-7058. doi:https://doi.org/10.1016/j.proeng.2013.12.004.

[8] R. Cummings and A. Schuette. An integrated computational/experimental approach to ucav stability &; control estimation: Overview of nato rto avt-161. In *28th AIAA Applied Aerodynamics Conference*. American Institute of Aeronautics and Astronautics, June 2010. doi:10.2514/6.2010-4392. URL `http://dx.doi.org/10.2514/6.2010-4392`.

[9] T. Loeser, D. Vicroy, and A. Schütte. Saccon static wind tunnel tests at dnw-nwb and 14 x22 nasa larc. In 28^th *AIAA applied aerodynamics conference, Chicago, Illinois, 28 June 2010 - 01 July 2010*, page 4393, AIAA Paper 2010-4393. doi:https://doi.org/10.2514/6.2010-4393.

[10] D. Vicroy, T. Loeser, and A. Schütte. Saccon dynamic wind tunnel tests at dnw-nwb and 14´x22´ nasa larc. In 28^{th} *AIAA Applied Aerodynamics Conference, 2010*, page 4394, AIAA Paper 2010-4394. doi:https://doi.org/10.2514/6.2010-4394.

[11] R. M. Cummings and A. Schütte. Integrated computational/experimental approach to unmanned combat air vehicle stability and control estimation. *Journal of Aircraft*, 49(6):1542–1557, 2012. doi:https://doi.org/10.2514/1.C031430.

[12] C. M. Liersch, K. C. Huber, A. Schütte, D. Zimper, and M. Siggel. Multidisciplinary design and aerodynamic assessment of an agile and highly swept aircraft configuration. *CEAS Aeronautical Journal*, 7(4):677–694, 2016. doi:https://doi.org/10.1007/s13272-016-0213-4.

[13] R. K. Nangia, J. Coppin, and M. Ghoreyshi. A ucav wing design, assessment and comparisons. In *2018 Applied Aerodynamics Conference, Atlanta, Georgia, 25-29 June, 2018*, page 2842, AIAA Paper 2018-2842. doi:https://doi.org/10.2514/6.2018-2842.

[14] C. M. Liersch, R. Cummings, A. Schütte, et al. Nato sto/avt-251: A joint exercise in collaborative combat aircraft design. In *AVT-324 Specialists' Meeting on Multidisciplinary design approaches and performance assessment of future combat aircraft, Braunschweig and online, 28-30 September 2020*, pages 1–18, NATO Science and Technology Organization, NATO STO/AVT-251, 2020. URL `https://elib.dlr.de/138971/`.

[15] C. M. Liersch and G. Bishop. Conceptual design of a 53 deg swept flying wing ucav configuration. In *2018 Applied Aerodynamics Conference, Atlanta, Georgia, June 25-29, 2018*, page 2839, AIAA Paper 2018-2839. doi:https://doi.org/10.2514/6.2018-2839.

[16] A. Schuette, J. Vormweg, R. G. Maye, and T. Jeans. Aerodynamic shaping design and vortical flow design aspects of a 53deg swept flying wing configuration. In *2018 Applied Aerodynamics Conference*. American Institute of Aeronautics and Astronautics, June 2018. doi:10.2514/6.2018-2841. URL `http://dx.doi.org/10.2514/6.2018-2841`.

[17] R. M. Cummings, C. Liersch, and A. Schuette. Multi-disciplinary design and performance assessment of effective, agile nato air vehicles. In *2018 Applied Aerodynamics*

Conference, Atlanta, Georgia, June 25-29, 2018, page 2838, AIAA Paper 2014-2838, June 2018. doi:https://doi.org/10.2514/6.2018-2838.

[18] R. Nangia and M. Palmer. A comparative study of ucav type wing planforms-aero performance & stability considerations. In *23rd AIAA Applied Aerodynamics Conference*, page 5078, 2005.

[19] I. Kroo. Nonplanar wing concepts for increased aircraft efficiency. *VKI lecture series on innovative configurations and advanced concepts for future civil aircraft*, pages 6–10, 2005.

[20] P. W. Jansen, R. E. Perez, and J. R. R. A. Martins. Aerostructural optimization of nonplanar lifting surfaces. *Journal of Aircraft*, 47(5):1490–1503, 2010.

[21] C. Conlan-Smith, N. Ramos-García, and C. Schousboe Andreasen. Aerodynamic shape optimization of highly nonplanar raised and drooped wings. *Journal of Aircraft*, 59(1):206–218, 2022.

[22] Z. Lyu and J. R. R. A. Martins. Aerodynamic design optimization studies of a blended-wing-body aircraft. *Journal of Aircraft*, 51(5):1604–1617, Sept. 2014. ISSN 1533-3868. doi:10.2514/1.c032491. URL `http://dx.doi.org/10.2514/1.C032491`.

[23] Openvsp, openvsp is a parametric aircraft geometry tool. `OpenVSP, OpenVSPisaparametricaircraftgeometrytool,http://openvsp.org/,2021`, 2021.

[24] R. A. McDonald and J. R. Gloudemans. Open vehicle sketch pad: An open source parametric geometry and analysis tool for conceptual aircraft design. In *AIAA SciTech 2022 Forum*, page 0004, 2022. doi:https://doi.org/10.2514/6.2022-0004.

[25] M. D. McKay, R. J. Beckman, and W. J. Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.

[26] K. Deb and A. Pratap. Agarwal and t. *Meyarivan,"A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II", IEEE_J_EVC*, pages 192–197, 2002.

[27] pymoo. `https://pymoo.org/`, 2022.

[28] H. Aleisa, K. Kontis, B. Pirlepeli, and M. Nikbay. Conceptual design of a nonconstant swept flying wing unmanned combat aerial vehicle. *Journal of Aircraft*, 60(6):1872–1888, 2023. doi:10.2514/1.C037257. URL `https://doi.org/10.2514/1.C037257`.

[29] T. D. Economon, F. Palacios, S. R. Copeland, T. W. Lukaczyk, and J. J. Alonso. SU2: An open-source suite for multiphysics simulation and design. *AIAA Journal*, 54(3):828–846, 2016. ISSN 00011452. doi:10.2514/1.J053813.

[30] H. Aleisa, K. Kontis, B. Pirlepeli, and M. Nikbay. Low-subsonic aerodynamic analyses of a nonplanar bwb model: An experimental and cfd study. In *34th Congress of the International Council of the Aeronautical Sciences*, 2024. URL `https://www.icas.org/icas_archive/icas2024/data/preview/icas2024_0232.htm`.

[31] E. F. Knott, J. F. Schaeffer, and M. T. Tulley. *Radar cross section*. SciTech Publishing, 2004.

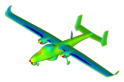
[32] D. Jenn. Pofacets4.5, 2025. URL `https://www.mathworks.com/matlabcentral/fileexchange/50602-pofacets4-5`.

[33] M. Bouhlel, N. Bartoli, J. Morlier, and A. Otsmane. Improving kriging surrogates of high-dimensional design models by partial least squares dimension reduction. *Structural and Multidisciplinary Optimization*, 53:935–, 01 2016.

[34] D. Jones, M. Schonlau, and W. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 12 1998. doi:10.1023/A:1008306431147. URL `https://doi.org/10.1023/A:1008306431147`.

[35] D. Jones. A taxonomy of global optimization methods based on response surfaces. *J. of Global Optimization*, 21:345–383, 12 2001. doi:10.1023/A:1012771025575. URL `https://doi.org/10.1023/a:1012771025575`.

[36] M. Schonlau, W. J. Welch, and D. R. Jones. *Global versus local search in constrained optimization of computer models*, page 11–25. Institute of Mathematical Statistics, 1998. doi:10.1214/lnms/1215456182. URL `http://dx.doi.org/10.1214/lnms/1215456182`.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AIRCRAFT WING STRUCTURAL SIZING COMPUTATIONAL TOOL TAILORED FOR A COLLABORATIVE MULTIDISCIPLINARY DESIGN FRAMEWORK

Pedro F. Albuquerque[1], Ana M. P. Silva[1] and André N. Pereira[1]

1: CEiiA - Centre of Engineering & Product Development
CEiiA — PACT - Herdade da Barba Rala, Rua Luís Adelino Fonseca Lote 1 A, 7005-841 Évora, Portugal
ana.silva@ceiia.com, https://www.ceiia.com/

Abstract. *This article presents two innovative tools,* ***WinG3N*** *and* ***WingSizer****, designed to streamline the traditional process of hand-made model development for aircraft wings.* ***WinG3N*** *enables the automatic generation of a wing finite element model within seconds, facilitating additional trade-off and sensitivity studies that would be impractical with manual methods. Complementarily,* ***WingSizer*** *automates the sizing process for static and buckling analyses, significantly reducing the time required compared to standard procedures, making it a valuable tool for efficient design. The* ***WingSizer****'s capability to compare metallic and composite structures allows for comprehensive trade-off studies, offering a quantitative assessment of the advantages and disadvantages of using composites. This feature enables a thorough evaluation of the potential benefits of altering the material composition of the entire wing assembly or specific components. While these tools are tailored for aircraft wing analysis, their applicability extends to the initial sizing of horizontal and vertical stabilizers, provided they feature no more than two longitudinal spars. The* ***WingSizer*** *versatility by seamlessly integrating into a multidisciplinary optimization framework has been tested during the 2023 International Forum for Aviation Research (IFAR) Young Engineers Challenge and presented at the 2024 International Council of the Aeronautical Sciences Congress. Two case studies are shown in the current article, one featuring the impact of buckling on a fully metallic wing and another comparing the structural mass of a fully metallic wing against a composite skin with metallic skeleton and a fully composite wing. The results corroborate prior research on the topic.*

Keywords: Aircraft Wing, Composite Aircraft, Structural Sizing and Optimization, Collaborative Design, Multidisciplinary Design

1 INTRODUCTION

Decisions taken early in the aircraft design process have great impact on the aircraft's performance and total project development cost with, as per Zhao et al [1], around 85% of the total Life Cycle Cost (LCC) being determined by design choices made at the end of the preliminary design stage. Furthermore, James et al [2] carried a study at an aerospace company on two programs covering 7 years of development and determined that early design decisions account for 86% of the cost impact of all design decisions.

It is therefore a differentiating factor to have state-of-the art methods and tools ready to facilitate and foster as many sensitivity analyses, trade-offs, and optimisation studies as possible during the conceptual and preliminary design phases. The more comprehensive these tools become, the more informed designers are from the earliest stages of the project development, thus enhancing the decision-making process.

Finite element models are a fundamental instrument in aircraft design by providing a high-fidelity prediction of the aircraft structural behaviour. Nonetheless, their development with a reasonable degree of detail can take, typically, several days to create and some more to be post-processed using conventional structural analysis methods. Consequently, this design approach is usually replaced with lower-fidelity equivalent models at early design stages.

However, the improvement of computational performance has created an opportunity for the employment of higher-fidelity models throughout the entire aircraft development process. Recent works have focused on the development of automated structural design and optimization frameworks using both low and high-fidelity analysis methods.

Chintapalli et al [3] have presented a methodology for the preliminary design optimization of a skin-stringer panel of an aircraft wing-box. The ability to resist the compressive load is assessed through a stability study to compute the critical buckling load of the stiffened panel while the ability to resist the tensile load is assessed through damage tolerance analysis of the lower wing panels. Nasoulis et al [4] have used an in-house Python-based tool and OpenVSP to generate the aircraft initial geometry, taking into consideration user defined Top-Level Aircraft Requirements, and the BETA CAE suite to generate, solve and post-process a wing box finite element model, all parts of an automated sizing procedure executed using a Python script. The design framework is applied to the optimization of a hybrid-electric aircraft wing considering both component position and thickness variables and resulting in an estimated mass reduction of 29%.

Additionally, aircraft structural design is coupled with other disciplines, such as aerodynamics and loads. A design optimization should consider the impact of structural changes in the remaining aircraft disciplines to avoid discrepancies and changes to the project in latter design stages. Incorporating multidisciplinary variables into structural optimization falls within the scope of Multidisciplinary Design Optimization (MDO). Benaouali and Kachel [5] have presented a fully automated framework dedicated to the high-fidelity multidisciplinary design optimization of an aircraft wing. Aiming at maximizing the aircraft range, given by the Breguet equation, while maintaining the lift coefficient and the structural performance. The case study resulted in a 8.9% increase in the aircraft range by considering shape and structural design variables. Kilimtzidis and Kostopoulos [6] presented a multidisciplinary optimization framework of high-aspect ratio composite aircraft wings combining variable fidelity computational tools. The input aerodynamic loads are calculated using a low-fidelity tool in the initial iteration and a high-fidelity tool in

subsequent loops, followed by flutter and static analysis. Throughout the iterations, a significant mass reduction is obtained by employing high-fidelity CFD solutions.

To enhance aircraft performance and efficiency the aerospace industry interest in incorporating composite materials into aircraft design and in improving manufacturing costs and consistency of composite parts has grown during the last decades. Several recent works have targeted structural optimization of aircraft wings considering composite materials.

Kennedy and Martins [7] have devised a wing aerostructural optimization framework and compared metallic to composite structures. They have shown that the composite wing designs are between 34% and 40% lighter than the equivalent metallic wings. Ebadi et al [8] have benchmarked a typical aluminum structure against CFR carbon composite to evaluate the potential mass savings and the environmental effects associated with the move from metallic to composite structures. They have come up to a 37% reduction in the mass of the composite structure compared to that of aluminum.

The current research aimed to develop a wing sizing tool using finite element models, **WingSizer**, to expedite the conceptual design and preliminary sizing of an aircraft wing and provide an early insight on the advantage of using composite materials for selected structural components. **WingSizer** includes another standalone development, the **WinG3N** tool [9], which automatically generates a wing finite element model from its geometric parameters, thus reducing the model development time to a few seconds. These tools were firstly presented during the 2024 International Council of the Aeronautical Sciences (ICAS) Congress [10].

The developed sizing framework was designed to allow for an easy and simple integration within a multidisciplinary design optimization schema. This integration was tested during the 2023 International Forum for Aviation Research (IFAR) Young Engineers Challenge, in which 18 international institutions contributed with design tools addressing different design disciplines, including aerodynamics, propulsion, loads, structure, performance, mass, cost and air traffic management.

The developed framework was further reduced in the number of involved partners (DLR [1], CEiiA [2], ONERA [3], JAXA [4] and TUWien [5]), and considering a wider scope of optimization cases. The results obtained were presented in another ICAS 2024 publication [11].

2 WING3N - WING GFEM GENERATOR

WinG3N [9] was developed to speed-up the generation of a wing finite element model by allowing the user to define its most relevant design parameters, including its span, aspect ratio, sweep and dihedral angles, front and rear spar location, stringers and ribs spacings, among others. Once these inputs are defined, the routine creates a Nastran coded bulk data file with the finite element model of the wing, ready to run a linear static analysis. If the user intends to conduct a parametric study evaluating the impact of changing some of these variables, it is possible to decide to generate as many wing

[1]German Aerospace Centre, Germany
[2]Engineering & Product Development Centre, Portugal
[3]Aerospace Lab, France
[4]Japan Aerospace Centre, Japan
[5]Technical University of Wien, Austria

skeletons as required for a given wing loft. This way, **WinG3N** enables the creation of different wing geometries, as shown in Fig. 1, in a few minutes.

Although this tool has been tailored to the development of aircraft fixed wing structures, it is easily adaptable to other lifting surfaces, namely the horizontal and vertical stabilisers, with the single limitation of not allowing the existence of more than two longitudinal spars.

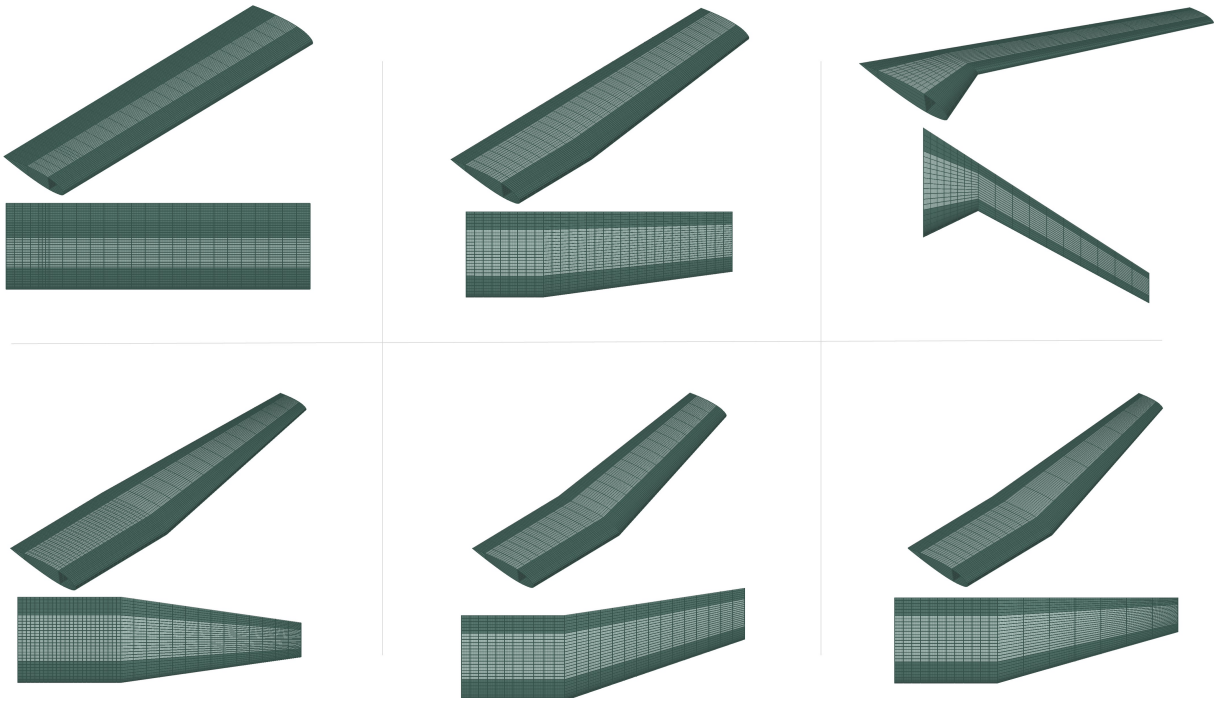


Figure 1: Different wing layouts, all of which have been generated using **WinG3N**.

3 WINGSIZER - WING SIZING TOOL

The wing sizing tool, **WingSizer**, uses **WinG3N** to create a baseline wing model which is run using Nastran and post-processed using Altair Hyperview. After running the baseline model, **WingSizer** will initiate a sizing loop using the Fully Stressed Design (FSD) method where, for each component, the maximum stresses are compared with their respective allowable and each component thickness or ply lay-up is updated until a positive margin of safety is obtained for all considered failure modes. In each loop, the updated model is run using a Nastran linear static solution and post-processed using Altair Hyperview. The tool mainframe, presented in Fig. 2, is developed in Python and uses an auxiliary TCL script to interact with Altair software and collect the maximum structural stresses for each component under analysis. The TCL script creates an Hyperview session, imports the model and the respective results file, creates an extreme envelope with the results of all load cases and loops through all model components to extract the maximum von mises stress, the minimum longitudinal compression stress, the maximum shear stress or the failure index, depending on whether the component is metallic or composite.

WingSizer evaluates the structural behaviour of wing skins, ribs, spars, leading and trailing edges. For metallic components, the tool is focused on two of the main concerns regarding wing structural analysis: static failure and instability, which tends to size a

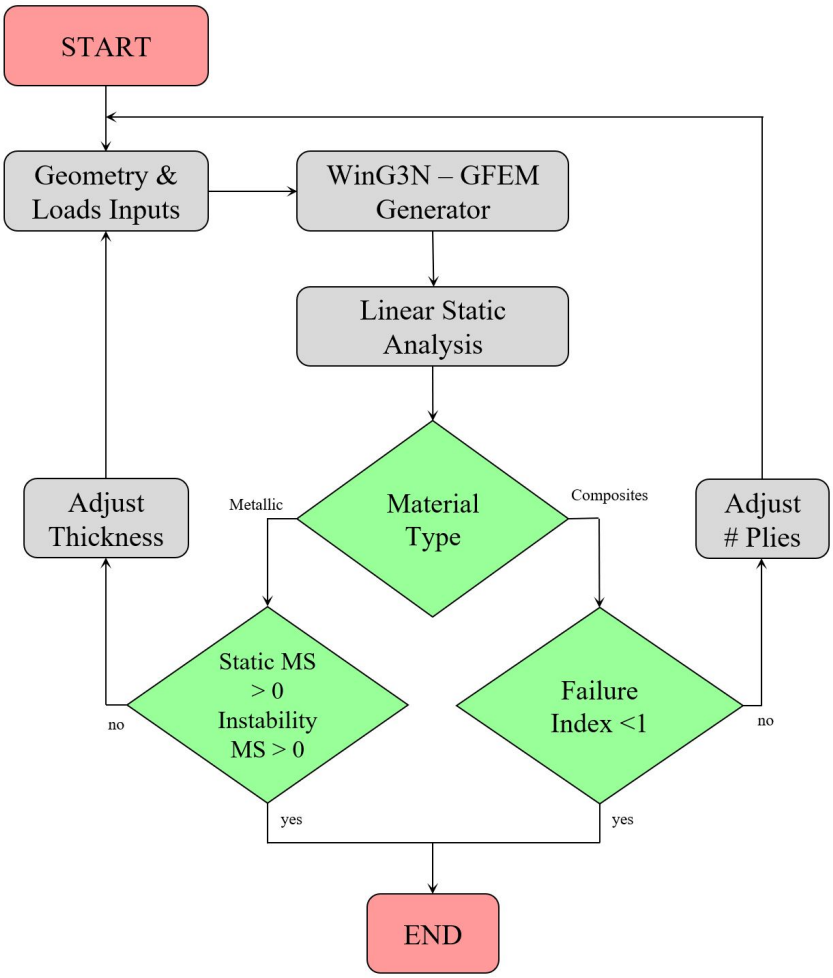


Figure 2: **WingSizer** framework schematic.

significant portion of the structural components exposed to the highest compressive loads. Instability analysis is based on NACA experiments from the 1940s [12], [13], aimed at evaluating the behaviour of flat and curved panels when subjected to pure shear, pure compression and a combination of these two types of loads. Some of these studies made it possible to draw curves for several buckling coefficients for metallic panels with different boundary conditions and understand how the buckling coefficients varied with the panels' geometry. For now, the instability assessment is applied to the wing central box skin panels only by evaluating compressive buckling, shear buckling and combined (shear and panel longitudinal stress) buckling.

To evaluate static failure the resulting von-mises stress (f_{vms}) is compared against the material ultimate tensile allowable (F_{tu}), as per Equation (1).

$$\mathrm{MS} = \frac{F_{tu}}{f_{\mathrm{vms}}} - 1 \tag{1}$$

The compressive buckling stress allowable ($F_{c_{cr}}$) is computed using Equation (2).

$$F_{c_{cr}} = \frac{\pi^2 k_c E}{12(1-\nu_e^2)} \left(\frac{t}{b}\right)^2 \tag{2}$$

, where:

- b is the panel width
- E is the skin material Young's modulus
- k_c is the axial compressive buckling coefficient (considering simply-supported edges)
- t is the panel thickness
- ν_e is elastic Poisson's coefficient

The respective MS is computed using Equation (3).

$$\mathrm{MS}_{\mathrm{comp}} = \frac{F_{c_{cr}}}{\sigma} - 1 \tag{3}$$

, where:

- σ is the acting compressive stress

Likewise, the shear stress can also cause a panel to buckle. The shear stress buckling allowable is computed from Equation (4).

$$F_{s_{cr}} = \frac{\pi^2 k_s E}{12(1-\nu_e^2)} \left(\frac{t}{b}\right)^2 \tag{4}$$

, where:

- k_s is the axial shear buckling coefficient, considering hinged edges

The respective MS is computed using Equation (5).

$$\mathrm{MS}_{\mathrm{shear}} = \frac{F_{s_{cr}}}{\tau} - 1 \tag{5}$$

, where:

- τ is the acting shear stress

The combined action of shear and compression can decrease the resultant component buckling allowable, whereas the action of shear and tension can increase it. The combined shear and compression/tension buckling margin of safety ($\mathrm{MS}_{\mathrm{comb}}$) is computed using Equation (6).

$$\mathrm{MS}_{\mathrm{comb}} = \frac{2}{R_L + \sqrt{R_L^2 + 4R_S^2}} - 1 \tag{6}$$

$$R_S = \frac{\tau}{F_{s_{cr}}} \tag{7}$$

$$R_L = \frac{\sigma}{F_{c_{cr}}} \tag{8}$$

For metallic components, **WingSizer** increases the component thickness until all applicable margins of safety, calculated using Equations (1), (3), (5) and (6), are positive.

As for composite components, only a static analysis is devised in the current study, where the structural integrity of the component is assessed using the laminate failure index. As it is common for composite materials, the Tsai-Wu failure criteria has been adopted, as per Equation (9).

$$FI = \left(\frac{1}{X_t} - \frac{1}{X_c}\right)\sigma_1 + \left(\frac{1}{Y_t} - \frac{1}{Y_c}\right)\sigma_2 + \frac{\sigma_1^2}{X_t X_c} + \frac{\sigma_2^2}{Y_t Y_c} + \frac{\sigma_{12}^2}{S^2} + 2F_{12}\sigma_1\sigma_2 = 1 \quad (9)$$

, where:

- σ_1 is the ply longitudinal stress
- σ_2 is the ply transverse stress
- σ_{12} is the ply shear stress
- X_t is the tensile stress allowable in the principal x (or 1) direction of the material
- X_c is the compressive stress allowable in the principal x (or 1) direction of the material
- Y_t is the tensile stress allowable in the principal y (or 2) direction of the material
- Y_c is the compressive stress allowable in the principal y (or 2) direction of the material
- S is the shear stress allowable in the principal material system
- F_{12} is the interaction strength constant, experimentally determined

The respective MS is computed using Equation (10).

$$\text{MS} = \frac{1}{F.I.} - 1 \quad (10)$$

If the failure index ($F.I.$) is higher than unity, the number of composite plies is increased until the MS becomes positive.

Further than sizing a structural component composite laminate, the **WingSizer** aims at benchmarking metallic with composite structures with the same geometry by providing a first-hand assessment of the mass difference between the two material categories. That can be done for individual components (e.g. wing spars, wing ribs) or for the wing as whole.

3.1 Assumptions

The forthcoming sections present results derived from an engineering modeling approach, underpinned by several key assumptions outlined in the current section:

- The hence referenced stringer pitch is the distance between stringers at the wing root, which is progressively reduced proportionally to the wing taper ratio as one moves towards the wing tip.

- The wing sizing procedure presented below relies solely on Nastran linear static analysis. Nevertheless, it should be noted that for metallic materials, an empirically based instability assessment facility has been devised.
- In order to guarantee, as much as possible, a fair comparison between the structural performance of composite versus metallic wings, the instability analysis of the metallic wing components was disabled, since there wasn't an analogous feature for the composite components.
- The stringers section and material (metallic) remain unchanged regardless of the wing components material choice.
- The model is clamped at the wing root using NASTRAN single point constraints (SPC) with all degrees of freedom restrained. Therefore, the first bay (wing root) components have been disregarded from the analysis due to the existence of unrealistic stress concentration points. Their thickness was set to 3mm for the metallic wing and the number of plies was set to 4 plies for the composite wing.

It should be noted, regardless of the assumptions listed above, the conclusions drawn from the results presented below come without loss of generality.

3.2 Material Properties

For benchmarking purposes, Aluminum 2024-T3 has been chosen as the reference metallic material. Its properties are presented in Tab. 1.

Table 1: Metallic material properties (Al2024-T3, QQ-A-250/4)) [14].

Symbol	Value	Units	Description
E	72395	MPa	Young's modulus
ν	0.33	–	Poisson coefficient
ρ	2.768×10^{-6}	kg/mm^3	Density

The physical and mechanical properties considered for the composite material [15] are presented in Tab. 2. The values with asterisk on Tab. 2 were not available in Hexcel properties datasheet and have been assumed by the authors based on their experience.

3.3 Case Studies

The case studies presented hereinafter have considered the Airbus A320 wing loft/geometry. Given the absence of data regarding the wing skeleton or the wing airfoil, typical values have been assumed. The front and rear spars have been assumed to seat at 25% and 79% of the mean aerodynamic chord, respectively. Likewise, the spacing between ribs ranges between 500mm and 800mm. One-dimensional elements have been used to represent the stringers, spar caps and rib caps, which have not been analysed for the current report. In the results that follow, these cross-sections are assumed to remain unchanged, being the same regardless of having metallic or composite structures.

Two design variables are particularly relevant in what concerns to wing structural sizing or optimization: the wing material (either metallic, composite or a combination of the two) and the spacing between stringers (skin panels longitudinal reinforcements).

Table 2: Composite material and mechanical properties, Hexply 8552 AGP 280-5H [15]

Symbol	Units	Value	Description
E_{12}, E_{21}	MPa	67000, 66000	Elasticity modulus
G^*	MPa	5280	Shear modulus
ν^*	-	0.07	Poisson coefficient
ρ	kg/mm^3	1.57×10^{-6}	Density
X_t	MPa	876	Tensile stress allowable X-direction
X_c	MPa	924	Compression stress allowable X-direction
Y_t	MPa	800	Tensile stress allowable Y-direction
Y_c^*	MPa	844	Compressive stress allowable Y-direction

In order to grasp the relevance of each of these variables on the wing under study, three pitches between stringers (180, 210 and 240mm) have been considered. These stringer pitches were combined with three different material arrangements, namely:

- Fully metallic wing structure
- Composite skins with a metallic skeleton
- Fully composite wing structure

Three fully metallic finite element models have been created, with the only difference between them being the stringer pitch: 180mm, 210mm and 240mm, respectively.

A baseline thickness of 1.0mm has been considered for all metallic components. The **WingSizer** tool evaluates the static failure modes for all components. The metallic skins are further evaluated in compressive, shear and combined (shear + compressive) buckling modes. In all components featuring positive margins of safety the thickness will remain unchanged while the remaining components' thicknesses will be progressively increased until positive margins of safety are attained.

Conversely, a baseline stack of 4 plies is used for the composite components, being that, whenever the failure index of a given component is greater than unity, the number of plies is incremented so that the stacking remains symmetrical and as close to balanced as possible. Fig. 3 shows the stacking sequence as the number of plies is increased, with the baseline stacking shown in gray and the added plies shown in white.

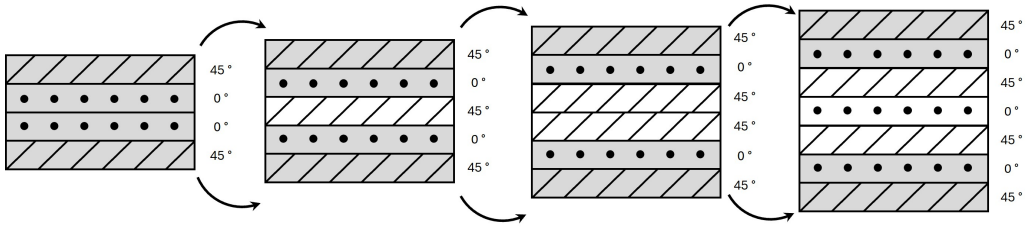


Figure 3: Composite Stacking sequence.

The aerodynamic loads for a 2g's positive load factor of the Airbus A320 in its maximum take-off mass (78,000 kg) configuration have been considered in all case studies.

4 RESULTS & DISCUSSION

4.1 Fully Metallic Wing Structure

In order to understand how the evaluation of skin instability affects the overall wing structural mass a comparative study with the skin instability evaluation activated and deactivated has been devised.

As expected, there is a slight mass penalty associated with the consideration of skin buckling. Moreover, while the most efficient wing, without buckling consideration, is the one with the greatest pitch between stringers, considering buckling, the most efficient wing becomes the one with the lowest pitch between stringers. The results are presented in Fig. 4 and Tab. 3.

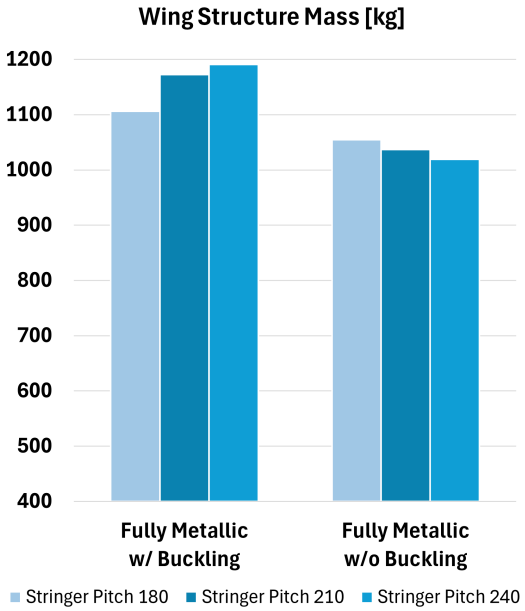


Figure 4: Metallic Wing Structure Mass [kg] as a function of the pitch between stringers, considering buckling and not considering buckling.

Table 3: Wing Structure Mass [kg].

Pitch Between Stringers [mm]	Fully Metallic considering buckling mass [Kg]	Fully Metallic without considering buckling mass [kg]
180	1106.3	1054.9
210	1172.6	1037.2
240	1191.1	1019.3

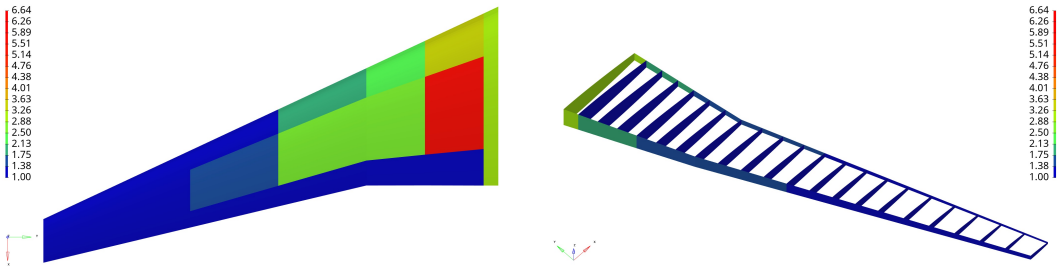
It is thus apparent that it is crucial to consider the evaluation of skin buckling while conducting trade-off studies during conceptual and preliminary design stages, given that skin buckling may significantly affect the optimum wing configuration.

Furthermore, it is noticeable that the wing's structural mass depends on the stringer spacings under consideration. The results suggest that stringer pitch can have meaningful

impact and should be carefully evaluated. In order to fully profit from the **WingSizer** potential, a broader study should be carried out, considering a broader range of stringer pitches.

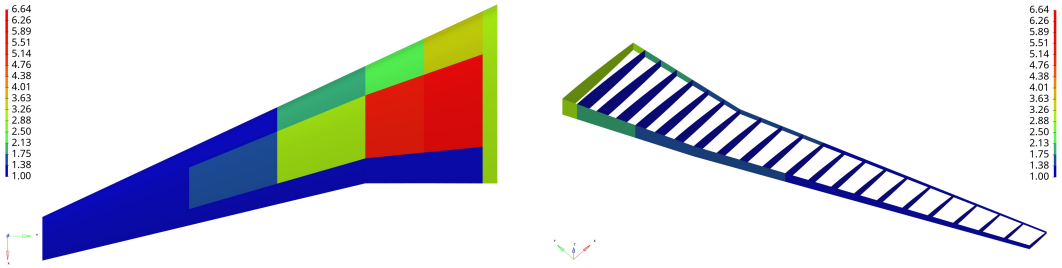
4.1.1 Static Analysis + Buckling

The components' thicknesses for the fully metallic wing, considering buckling, are shown in Fig. 5 for spacing between stringers of $180mm$, $210mm$ and $240mm$, respectively. The legend of each plot has been patterned so that the reader can relate the colors between plots. It is noticeable that the thicknesses change depending on the pitch between stringers.



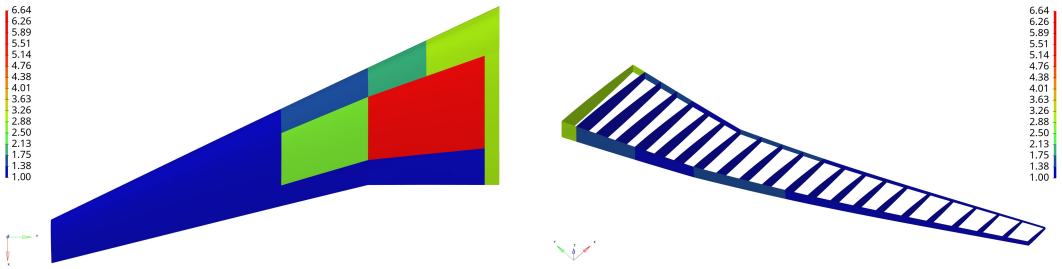
(a) Skins. (b) Skeleton (ribs and spars).

Stringer pitch: 180mm.



(c) Skins. (d) Skeleton (ribs and spars).

Stringer pitch: 210mm.



(e) Skins. (f) Skeleton (ribs and spars).

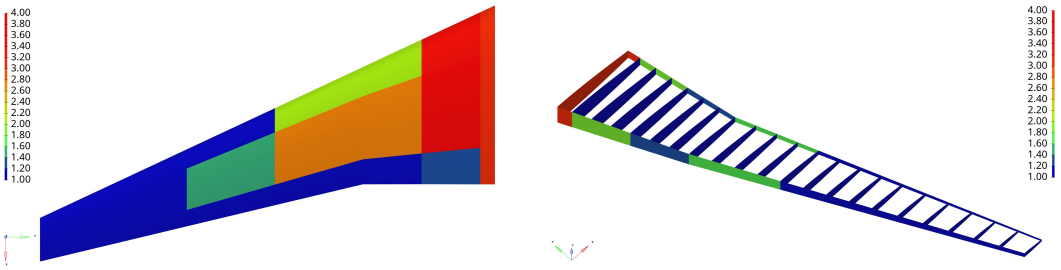
Stringer pitch: 240mm.

Figure 5: Fully metallic wing, considering buckling.

4.1.2 Only Static Analysis

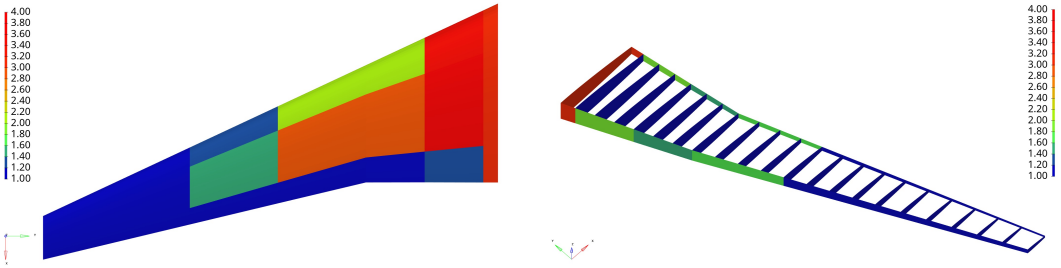
The components' thicknesses for the fully metallic wing, without considering buckling, are shown in Fig. 6 for spacing between stringers of $180mm$, $210mm$ and $240mm$, respectively.

The legend of each plot has been likewise patterned so that the reader can relate the colors between plots. It is noticeable that the thicknesses change depending on the pitch between stringers.



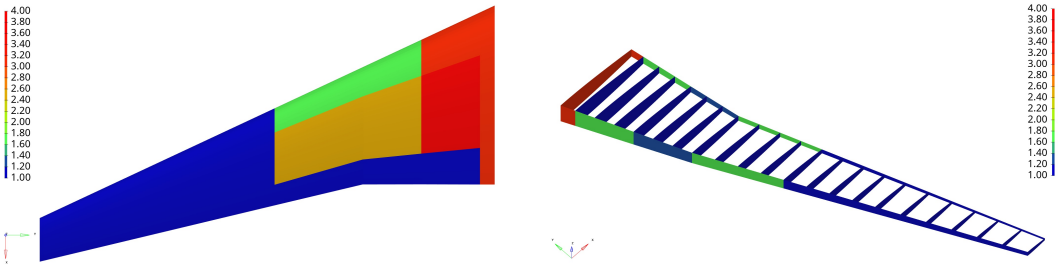
(a) Skins. (b) Skeleton (ribs and spars).

Stringer pitch: 180mm.



(c) Skins. (d) Skeleton (ribs and spars).

Stringer pitch: 210mm.



(e) Skins. (f) Skeleton (ribs and spars).

Stringer pitch: 240mm.

Figure 6: Fully metallic wing, without considering buckling.

4.2 Metallic versus Composite Wing Structure Benchmark

While the **WingSizer** tool allows for an instability evaluation of metallic skins, it wouldn't be equitable to compare composite solutions without also considering skin instability. Therefore, in the current section, the wings encompassing composite materials will be assessed against a fully metallic wing without considering buckling.

The results presented in Tab. 4 and Fig. 7 show that from a structural mass standpoint, the more composite materials are used, the more efficient the wing will be. That is in line with what would be expected.

As for the 9 simulations which have disregarded the skin buckling, the larger the pitch between stringers, the lower the wing structural mass.

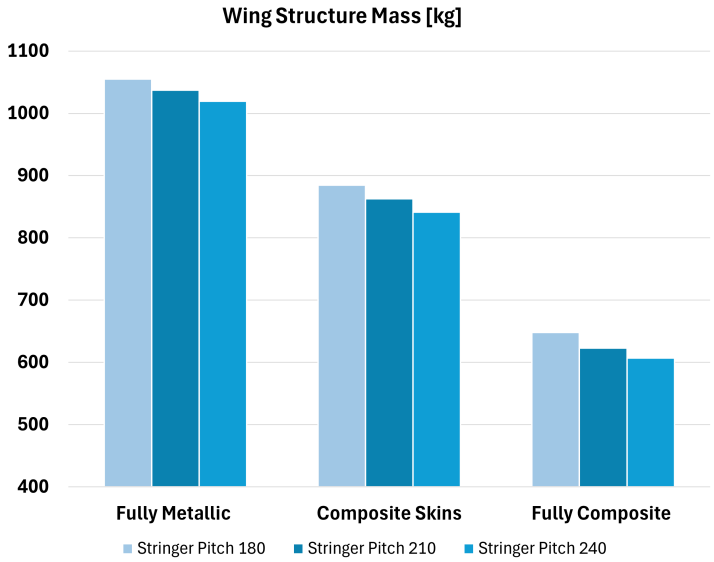


Figure 7: Metallic, Composite and Hybrid Wings Structure Mass [kg] as a function of the pitch between stringers.

The results presented in Fig. 7 have been summarized in Tab. 4.

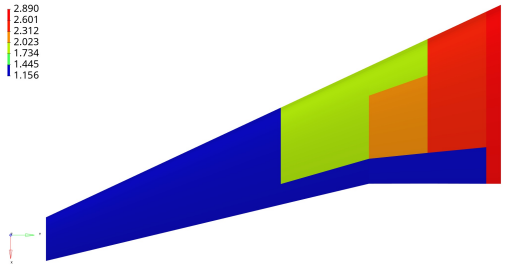
Table 4: Wing Structure Mass [kg] without considering skin buckling.

Pitch Between Stringers [mm]	Fully Metallic considering buckling mass [Kg]	Fully Metallic without considering buckling mass [kg]
180	1106.3	1054.9
210	1172.6	1037.2
240	1191.1	1019.3

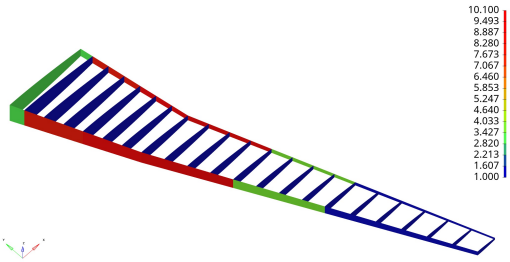
A mass saving between 39% and 40% was achieved by changing all wing components material from metallic to composite material. Despite the underlying assumptions in the WingSizer, which aims to be a preliminary sizing tool, these results are in line with the mass savings obtained by Kennedy and Martins [7].

4.2.1 Composite Wing Skin with Metallic Skeleton

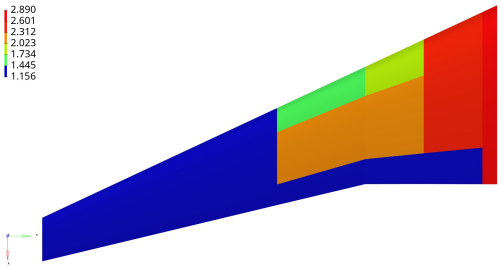
The components' thicknesses for a wing featuring composite wing skins and a metallic skeleton, without considering buckling, are shown in Fig. 8 for spacing between stringers of $180mm$, $210mm$ and $240mm$, respectively. The legend of each plot has been likewise patterned so that the reader can relate the colors between plots. Each 1.156mm refers to 4 plies (for the composite skins). It is noticeable that the thicknesses change depending on the pitch between stringers.



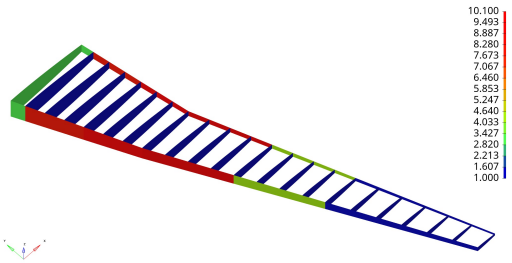
(a) Composite Skins.



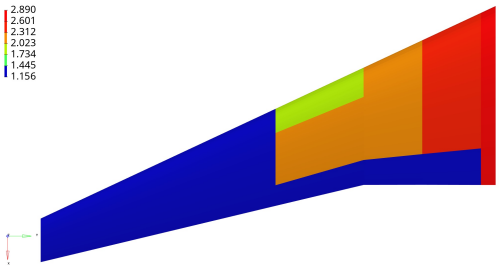
(b) Metallic skeleton (ribs and spars). Stringer Pitch: 180mm.



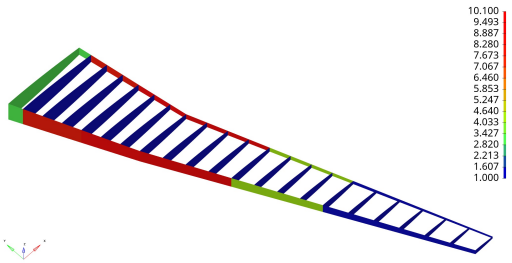
(c) Composite Skins.



(d) Metallic skeleton (ribs and spars). Stringer Pitch: 210mm.



(e) Composite Skins.



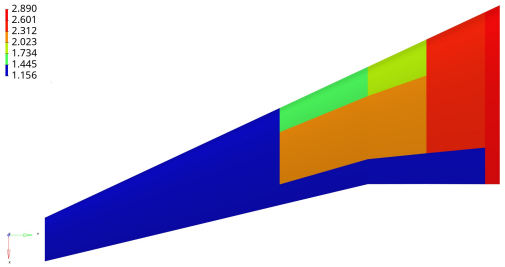
(f) Metallic skeleton (ribs and spars). Stringer Pitch: 240mm.

Figure 8: Composite wing skin, disregarding buckling, wing components thickness [mm].

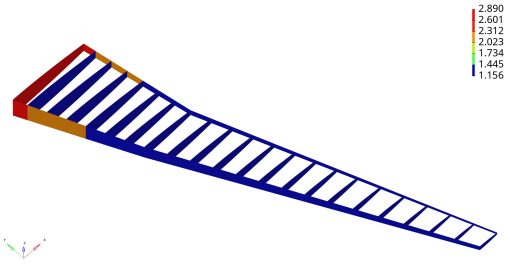
4.2.2 Fully Composite Wing Structure

The components' thicknesses for a fully composite wing, without considering buckling, are shown in Fig. 9 for spacing between stringers of $180mm$, $210mm$ and $240mm$, respectively. The legend of each plot has been likewise patterned so that the reader can

relate the colors between plots. Each 1.156mm refers to 4 plies. It is noticeable that the thicknesses change depending on the pitch between stringers.

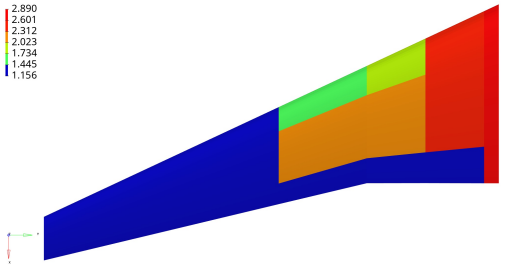


(a) Skins.

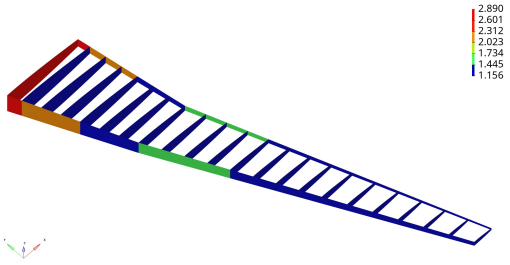


(b) Skeleton (ribs and spars).

Stringer Pitch: 180mm

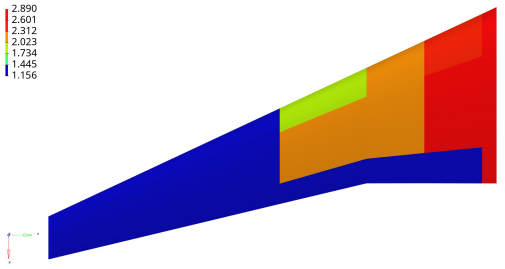


(c) Skins.

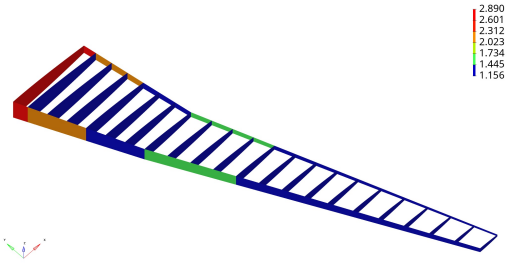


(d) Skeleton (ribs and spars).

Stringer Pitch: 210mm



(e) Skins.



(f) Skeleton (ribs and spars).

Stringer Pitch: 240mm

Figure 9: Composite wing, disregarding buckling, wing components thickness [mm].

5 WING CONFIGURATIONS DISPLACEMENTS

Lastly from the case studies presented in the previous sections, Fig. 10 shows the wing deformation for a stringer pitch of 180mm.

It is clear that the fully composite wing yields a much larger deformation than the fully metallic wing. As expected, the composite skin with a metallic skeleton seats somewhere in between the fully metallic and fully composite solutions. It is noticeable that, when buckling was considered (fully metallic wing), the wing became structurally stiffer, which

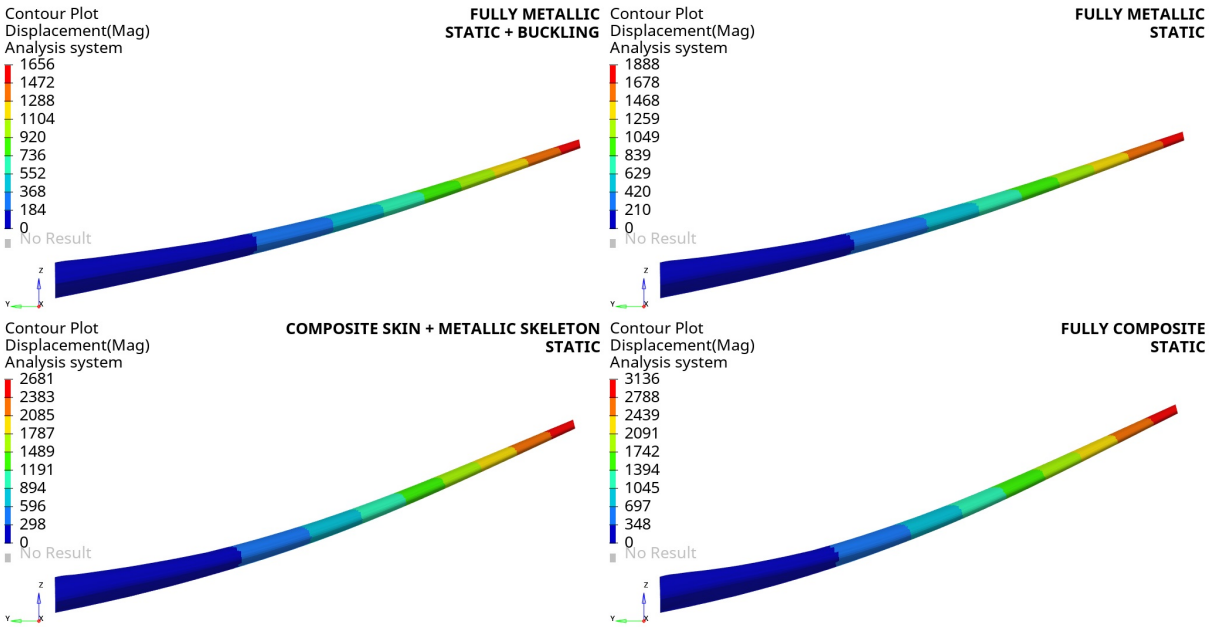


Figure 10: Wing displacements, depending on the wing structure material and sizing considerations, for a stringer pitch of 180mm.

resulted in a smaller deformation with respect to the homologue metallic wing where only static failure was considered.

6 WINGSIZER TOOL PERFORMANCE

As a preliminary design tool, the **WingSizer** should swiftly deliver low-fidelity analysis results in order to drive the designer's choice concerning key design variables throughout the early stages of aircraft design. The WingSizer can be an invaluable tool to evaluate the worthiness of using composite materials, or the extent to which they should be used, optimizing the wing skeleton (spars, ribs and stringers spacings) and even evaluating how changing the stringers inertia may modify these conclusions.

Fig. 11 and 12 show the simulation running time and convergence loops for each of the simulations run for the current article, respectively.

Although there is margin for greater computational performance, namely by revising the components thickness gain function, the simulations have taken anywhere from 12min up to 2h15min, which is quite acceptable to perform a large number of trade-off and preliminary sizing studies in a rather efficient way. It should also be considered that several simulations can be run in parallel thus enhancing the tool's output performance.

Despite its advantages, the WingSizer tool results should not be regarded as a full wing sizing, even when buckling is activated, because most failure modes are not grasped by the underlying methodology. Nevertheless, both the tool results and performance figures indicate that it can be extremely helpful to drive better design decisions from the earliest stages of wing design.



Figure 11: WingSizer running time until convergence [min].

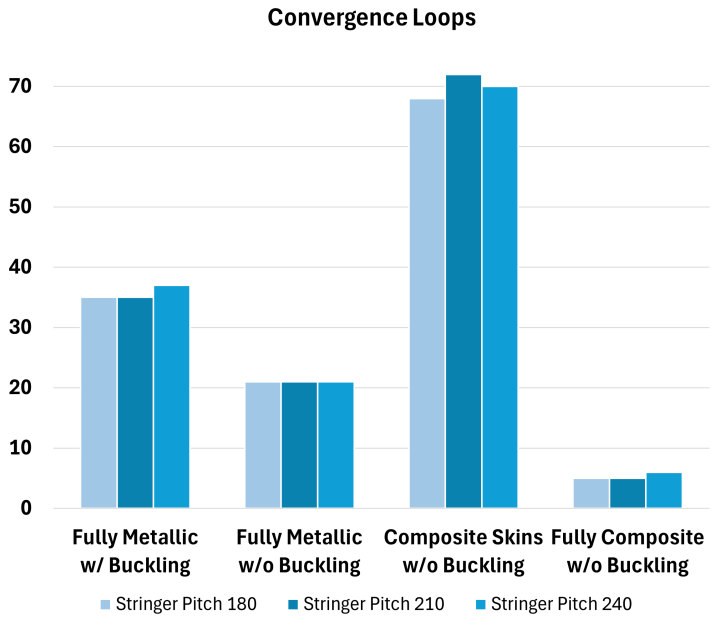


Figure 12: WingSizer simulations number of loops until convergence.

7 INTEGRATION WITHIN A MDAO WORKFLOW

The **WingSizer** integration within a multidisciplinary design optimization schema was tested during the 2023 International Forum for Aviation Research (IFAR) Young Engineers Challenge.

All contributing tools were integrated within a collaborative Multidisciplinary Design Analysis and Optimization (MDAO) Workflow Design Accelerator (MDAx) developed by the DLR [6]. The exchange of data between the different disciplines/partners was done using

[6]German Aerospace Centre, Germany

a common xml based parametric scheme known as Common Parametric Aircraft Configuration Schema (CPACS) [11]. CPACS is therefore a driver for multidisciplinary and multi-fidelity design in distributed environments and widens the chances of collaboration by eliminating the need to share proprietary tool information.

The **WingSizer** tool received as inputs the wing loft (from ONERA) and the wing aerodynamic loads distribution (from ILOT [7] – Polish Institute of Aviation) and provided as outputs the thickness per component and the wing structural mass (to A*STAR [8], Singapore Institute for High Performance Computing).

The developed framework was further reduced in the number of involved partners (DLR, CEiiA, ONERA, JAXA and TUWien), but considering a wider scope of optimization cases in another ICAS 2024 publication [11].

8 CONCLUSIONS

The standard hand-made wing finite element model can take from a few days to some weeks to develop, the **WinG3N** tool makes it possible to automatically generate a wing finite element model within a few seconds, thus unfolding a number of additional trade-off and sensitivity studies which would be unfeasible otherwise.

Likewise, a first sizing of a wing finite element model (for static and buckling) analysis may take several days as well. With the sizing framework implemented in **WingSizer**, not only does that process become automatic but also much faster than the traditional standard procedure, once more saving from a few hours to up to several days, depending on the model dimension and complexity.

The results have shown the worthiness of this tool in the automation of the preliminary sizing of an aircraft wing main structural components. It contributes to make the conceptual and preliminary design stage more efficient and effective due to the swiftness inherent to the automation process and the consequent much larger number of trade-off studies and parametric optimisation analysis that can be made that would otherwise be limited due to their cost.

Further to the preliminary sizing of the structure, **WingSizer** makes it possible to benchmark the performance of metallic against composite materials in selected components of the wing assembly. This ability can be used to perform trade-off studies which can provide a quantitative assessment of the advantages or disadvantages of using composites instead of metallic structures enabling the assessment of changing the whole wing assembly material or only some of its components.

The case study results presented have shown that considering skin instability (for metallic materials) can have a great impact on the wing's structural mass and the best stringer pitch is not the same whether wing sizing is made considering buckling or only considering static failure.

An additional comparative study was undertaken in order to compare a fully metallic wing against a composite skin with a metallic skeleton and a fully composite wing. The results indicate a potential wing mass saving ranging between 39% and 40%, which is in line with literature. These results were achieved with the buckling option being deactivated for all three wing material arrangements, for the sake of fairness, due to the methodology not supporting buckling on the composite components.

[7] Institute of Aviation, Poland

[8] Agency for Science, Technology and Research, Singapore

Although the two tools developed (**WinG3N** and **WingSizer**) have been tailored to the development and analysis of aircraft wings, they may be likewise used for the initial sizing of aircraft horizontal and vertical stabilizers, provided they feature no more than two longitudinal spars.

In addition to its standalone features, **WingSizer** was designed to allow for an easy and simple integration within a multidisciplinary design optimization framework, as it has been shown in the context of the IFAR-X 2023 Young Engineers challenge and the ICAS 2024 congress [11], where it interacted with other partner tools and design disciplines using a standard common xml based parametric scheme to exchange inputs and outputs, further enhancing its scope.

9 ACKNOWLEDGEMENTS

The authors would like to thank Witold Klimczyk from ILOT [9] for contributing with the aerodynamic load distribution used in Section 4.

REFERENCES

[1] X. Zhao, W. J. Verhagen, and R. Curran. Estimation of aircraft component production cost using knowledge based engineering techniques. *Advanced Engineering Informatics*, 29(3):616–632, 2015. ISSN 1474-0346. doi:https://doi.org/10.1016/j.aei.2015.05.004.

[2] T. J. Y. James, K. Otto, and K. Wood. A comparison of design decisions made early and late in development. *ICED*, 08 2017.

[3] S. Chintapalli, M. S. Elsayed, R. Sedaghati, and M. Abdo. The development of a preliminary structural design optimization method of an aircraft wing-box skin-stringer panels. *Aerospace Science and Technology*, 14(3):188–198, 2010. ISSN 1270-9638. doi:https://doi.org/10.1016/j.ast.2009.12.007. URL `https://www.sciencedirect.com/science/article/pii/S1270963809001072`.

[4] C. Nasoulis, P. Tsirikoglou, and A. Kalfas. Structural optimization of the wing box for a hybrid-electric commuter aircraft. 04 2022. doi:10.33737/gpps21-tc-211.

[5] A. Benaouali and S. Kachel. Multidisciplinary design optimization of aircraft wing using commercial software integration. *Aerospace Science and Technology*, 92:766–776, 09 2019. doi:10.1016/j.ast.2019.06.040.

[6] S. Kilimtzidis and V. Kostopoulos. Multidisciplinary structural optimization of novel high-aspect ratio composite aircraft wings. *Structural and Multidisciplinary Optimization*, 66, 06 2023. doi:10.1007/s00158-023-03600-1.

[7] G. Kennedy and J. Martins. A comparison of metallic and composite aircraft wings using aerostructural design optimization. 09 2012. doi:10.2514/6.2012-5475.

[8] S. Ebadi, K. Shahbazi, and E. Anbarzadeh. Investigation of aluminum and composite aircraft wings under the influence of aerodynamic forces and their effects on envi-

[9] Institute of Aviation, Poland

ronmental impacts. *Journal of Environmental Friendly Materials*, 5(1):13–21, 2021. ISSN 2538-3620.

[9] A. N. Pereira. Airplane wing skeleton structural efficiency optimisation. 2023. Master Thesis Dissertation - University of Aveiro, Portugal.

[10] P. F. Albuquerque, A. M. P. Silva, and A. N. Pereira. Aircraft wing structural sizing computational tool tailored for a collaborative multidisciplinary design framework. 2024. International Council of the Aeronautical Sciences Congress.

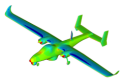
[11] S. Garg, P. S. Prakasha, A. M. P. Silva, P. F. Albuquerque, E. N. Van, K. Kimura, Y. Sugioka, and M. de R. Jacinto. Ifar-x: Collaborative engineering framework for next generation aerospace engineers working on efficient and faster aircraft design. 2024. International Council of the Aeronautical Sciences Congress.

[12] M. Schildcrout and M. Stein. Critical combinations of shear and direct axial stress for curved rectangular panels. 1949.

[13] G. Gerard and H. Becker. Handbook of structural stability part i: Buckling of flat plates. 1957. doi:10.2172/4343548.

[14] F. A. Administration. Metallic materials properties development and standardization (mmpds). 2010.

[15] Hexcel properties datasheet - hexply 8552 epoxy matrix (180^{o}c curing matrix).



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

BENCHMARKING SURROGATE MODELS OF VARIABLE FIDELITY FOR THE DRAG COEFFICIENT PREDICTIONS OF AIRFOIL AERODYNAMICS WITH DIFFERENT AIRFOIL PARAMETRIZATION TECHNIQUES

Ivo Zell, Maurice Zimmnau, Eike Stumpf

Institute of Aerospace Systems
RWTH Aachen University
Wuellnerstrasse 7, 52062 Aachen, Germany
ivo.zell@rwth-aachen.de, zimmnau@ilr.rwth-aachen.de, stumpf@ilr.rwth-aachen.de

Abstract. *This contribution examines various existing surrogate models, including variable-fidelity models, to evaluate their suitability within the conceptual aircraft design, and compares them against one another. The drag coefficient calculation for two-dimensional airfoil sections serves as a benchmark. The benchmarking framework developed includes the integration of sampling strategies, automatic generation of modified airfoils, and calculation of the training datasets for surrogate model training. The geometric representation of airfoils is achieved through the parametrizations PARSEC and NURBS, offering varying levels of input dimensionality. Low-fidelity (LF) data is computed using xFoil, while higher-fidelity (HF) data and test datapoints are generated with MSES. The Surrogate Modeling Toolbox (SMT) is used to analyze different types of surrogate models. In addition to single-fidelity models, three types of variable-fidelity models are tested. Training datasets contain up to 8,192 data points with 8 to 23 input variables and the drag coefficient as target variable. The models are assessed based on prediction error, training time, and memory usage. The results reveal significant trends in surrogate model performance across different dataset sizes and configurations. For smaller training datasets, the Kriging model exhibits the highest accuracy, while Artificial Neural Networks deliver the lowest prediction error for larger datasets. Among variable-fidelity models, the direct incorporation of LF data as a mean function of the surrogate model achieves the lowest error for small (8-64 points) HF datasets, making them ideal for applications where HF data is sparse or expensive to generate. For medium-sized datasets (128–512 points), Co-Kriging achieves the best results. In cases with large datasets (1,024–8,192 points), single-fidelity ANN models provide the lowest prediction errors. The study also shows that LF data can be used to reduce complexity and therefore allow easier surrogate model building when the underlying correlation structure of the data is suitable for the model.*

Keywords: Surrogate models, Machine Learning, Kriging, Artificial Neural Network, NURBS, PARSEC

1 INTRODUCTION

In conceptual aircraft design one of the objectives is to examine large design spaces and to investigate numerous different aircraft concepts. Within this conceptual phase, on the one hand, fast design results are necessary. On the other hand, the accuracy of the results shall be reliable. The design methods in this phase are often, but not exclusively, based on semi-empirical methods [1]. These semi-empirical methods are valid, mostly for conventional, well known, aircraft configurations. As we move from conventional to unconventional configurations, the semi-empirical methods lack in accuracy [2]. To enhance this accuracy, more physics-based, methods of higher fidelity are necessary. Higher-fidelity methods, however, come with their higher computational cost. To keep this rise of computational cost acceptable, surrogate modeling is a prominent technique [3]. Especially, if there exist methods of different fidelity, surrogate models can be used to combine these models and further reduce the computational cost, with a minimal lack of accuracy [4]. However, the choice of a suitable surrogate model depends, among others, on the data set and density, on the design of experiment, on the dimensionality of input and output, and on the model training time.

This contribution, in particular, investigates and compares various surrogate model approaches and model types, with an emphasis on accuracy, efficiency, and dependency on input dimensionality and dataset size. To simplify the complexity and computational demands of data generation, this study uses the calculation of drag on 2D-airfoils as a representative case study. This example provides insight into surrogate model behavior and performance without the intensive resource requirements associated with full-scale aircraft models.

The benchmarking framework developed in this contribution includes the integration of sampling strategies, automatic generation of modified airfoils, and calculation of the training datasets for surrogate model training. Therefore, the airfoils are parametrized with the PARSEC and NURBS methods, and the drag calculated using the xFoil and MSES software. Based on an initial comparison of surrogate model libraries, the Surrogate Modeling Toolbox (SMT) is selected, and preliminary parameter tuning is conducted. The study then assesses multiple surrogate model types, VF methods, and parametrization based on accuracy, wall-clock time, and memory requirements.

2 METHODS AND BENCHMARKING PROCESS CHAIN

In the last decades the use of surrogate models became popular, in the engineering field and in aerospace as well [3]. Consequently there exist numerous of different surrogate modeling techniques from simple regression models to the prominent neural networks, which all fit also into the term machine learning [5]. To benchmark a small subset of these surrogate modeling techniques in the following subsections the used methods are presented. At first the process chain for the data generation and evaluation process is described in detail (Fig.1).

The benchmark process starts with the generation of airfoil sets. We start with a so-called base airfoil, which we use to define a starting point for the design space. As depicted, the design space for the geometrical parameters is then constrained to guarantee a feasible set of airfoils. These constraints are used afterwards as an input for the sampling strategy. Output of the sampling strategy is a set of airfoils which then in turn is parametrized by chosen parametrization techniques. Furthermore, aerodynamic analysis of the resulting

airfoils is performed with two methods (xFoil and MSES) of different fidelity. The resulting data is combined into a large dataset. This dataset is divided into training and test set. The training data set is used to train a subset of the surrogate models present in the SMT. The benchmark process ends with the quantification of the deviations in the predictions of the surrogate models at the chosen test points. All calculations are performed on the RWTH HPC [6].

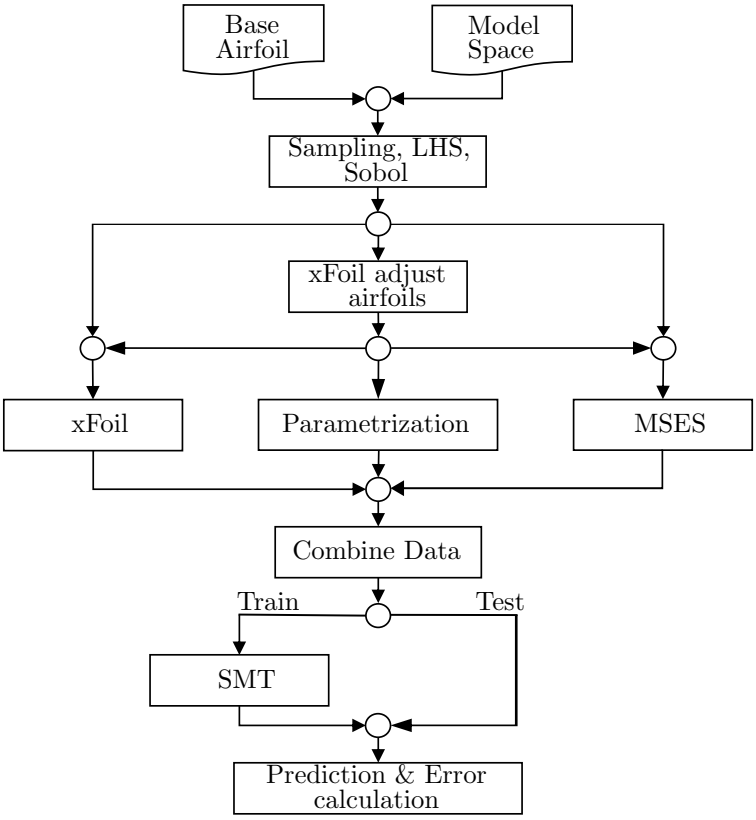


Figure 1: Process chain for the benchmarking of surrogate models, including test and training data generation, airfoil paramatrization, and error evaluation.

2.1 Sampling strategies and design space definition

Surrogate models are trained with data sets. In literature there exist different methods for the generation of training data sets, so called samples. These sampling strategies are called Design of Experiment. Prominent representatives are Latin Hypercube Sampling (LHS) and Sobol sequences. [7]

LHS is based on a random process and ensures a space-filling design. LHS divides each dimension of the design space into sub-spaces with

$$[0,1) = [0,k_0), \cap [k_0,k_1), \cap \cdots \cap [k_{n-1},k_n)], \quad k_i = \frac{i}{n_u} \tag{1}$$

These sub-spaces are filled randomly with samples until each sub-space is sampled once.

Low-discrepancy sequences are methods designed to generate well-distributed, deterministic sequences of points across a given domain. These sequences maintain a low discrepancy as new points are added, meaning that the points remain evenly spread and

avoid clustering, ensuring that the sample density remains consistent throughout the design space. Many of these sequences are generalizations of the one-dimensional van der Corput sequence, adapted to higher dimensions. The Sobol sequence, which uses a base-2 structure to minimize correlations between dimensions. Sobol sequences are highly effective in high-dimensional spaces. This makes them particularly valuable for complex simulations and optimization problems. [8, p. 384 ff]

As depicted in the process chain Fig.1, a design space must be defined in order to generate feasible airfoil candidates. The design space of the airfoils is defined by the flight conditions Mach number Ma, Reynolds number Re and the angle of attack AOA, and the geometry parameters thickness t, maximum thickness position x_t, camber c, position of maximum camber x_c. and leading edge radius $lera$. We use the RAE2822 airfoil as an origin for the design space. We then open the design space to a certain percentage of deviation from this origin, again to restrict the design space to feasible airfoils. The design space definition is depicted in Tab.1. For each parameter there is a defined start value, the variation of this parameter as well as the resulting boundaries of the design space.

Table 1: Design space definition

Parameter	Start Value	± Variation	Min. value	Max. Value
Ma	0.6	±0.15	0.45	0.75
Re	$10e^{10}$	$\pm 2.5e^{10}$	$7.5e^{10}$	$12.5e^{10}$
AOA	2°	±1.5°	0.5°	3.5°
t	12.1%	±2.5%	9.6%	14.6%
x_t	37.8%	±10%	27.8%	47.8%
c	1.26%	±0.5%	0.76%	1.76%
x_c	75.7%	±10%	65.7%	85.7%
$lera$	1	±0.5	0.5	1.5

With the defined design space LHS and Sobol methods are used to generate the desired number of sampling points, each representing a unique combination of parameters. For the training data set, the Sobol sampling method from the SCIPY.STATS.QMC toolbox [9] is used. To ensure balance in the sequence, the number of sampling points must be set to $n = 2^m$, $n, m \in \mathbb{N}$. A total of $2^{13} = 8192$ points are sampled, with the option to use smaller subsets for model training. The number of dimensions is defined and *scrambling* is set to *False* to generate a reproducible and deterministic DoE with the Sobol method. For LHS, the integrated method of SMT [10] is used, with the option set to maximize the minimum distance between the samples. The total number of samples for the test set is set to 500. From the given parameters, airfoils are constructed with XFOIL [11].

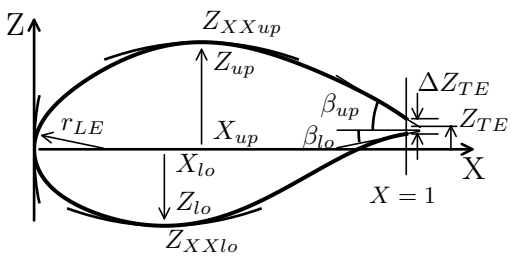
2.2 Airfoil parametrization

The purpose of geometric parametrization is to provide a flexible and structured way to describe, modify, and manipulate geometric shapes. The goal is to describe a wide range of possible configurations while maintaining physical feasibility. Airfoils can be described with different parametrizations. Plain parametrization with X- and Z-coordinates nearly

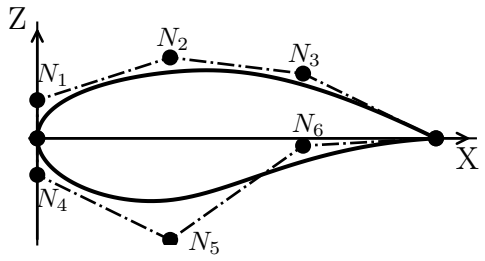
always exceeds 100 variables. This is too large for (most) surrogate models. Therefore, another description is required.

Prominent geometry parametrizations are PARametric SECtion (PARSEC) method and Non-uniform rational B-splines (NURBS). The PARSEC airfoil parametrization system was developed in [12] as part of his work to create a universal method for parametrizing both wings and fuselages by describing their cross-sectional shapes. In this system, an airfoil is characterized by eleven coefficients that define key geometric features, which are depicted in Fig.2a.

NURBS are a generalization of B-splines and Bezier curves [13] and are commonly used for airfoil parameterization, although they were not specifically developed for this purpose. Originally designed for more general tasks in computer graphics and CAD, NURBS offer flexibility in representing curves and surfaces, making them suitable not only for airfoil design but also for wing planforms and fuselage parametrization. An example airfoil parametrization with 6 NURBS control points is shown in Fig.2b.



(a) PARSEC airfoil geometry defined by 11 parameters [12]



(b) NURBS example with six control points N_1 to N_6

Figure 2: Airfoil parametrization methods used in this contribution

In this contribution, a comparison of NURBS parametrizations with 6, 10, and 20 control points are investigated. For PARSEC a 9 parameter representation is used[1]. For the automated transformation of the seed airfoils to the parametrizations, the values of the parameters for both PARSEC and NURBS must be determined through an optimization process. This is achieved using a metaheuristic optimization function to fit the parameters to the airfoil geometry. The optimizer, developed by Mirjalili et al. [14], has been shown to be robust and reliable in previous projects, delivering sound results. For this contribution, the Particle Swarm Optimizer (PSO) is selected.

2.3 Aerodynamic airfoil calculation methods

For the aerodynamic analysis of airfoils, in the literature, there exist different methods. Commonly used methods are based on potential flow, e.g., panel methods, as they are fast methods, where Computational Fluid Dynamics (CFD) are used for a detailed analysis with higher accuracy and higher computation times.[15]

The linear-vorticity panel method is based on potential flow theory and is widely used to calculate aerodynamic properties of the airfoil such as lift and drag coefficients, as well

[1]The trailing edge position is fixed at (1,0) $Z_{TE} = 0$. The trailing edge thickness is set to $\Delta Z_{TE} = 0$. This reduces the number of parameters from 11 to 9. (Fig.2a)

as pressure distributions. Since the potential flow theory assumes inviscid, irrotanional and incompressible flow, it is suitable for subsonic flows. For the application at higher subsonic Mach numbers, where compressibility effects become significant, corrections, like the Kármán-Tsien compressibility correction must be applied.

CFD is a large research field, where the objective is to solve the Navier-Stokes equations with certain assumptions or simplifications. The Navier-Stokes equations are a set of coupled differential equations describing the conservation of mass, energy and momentum. Since directly solving these equations is computationally too demanding, commonly the turbulence is modeled by equations, which results in the so-called Reynolds-averaged Navier-Stokes equations (RANS) . Further simplifications are e.g., assuming inviscid flow, neglecting heat convection terms and shear stresses, which results in the Euler-equations.

In this contribution, for the application case of 2D airfoil aerodynamics we use a linear-velocity panel method, provided by XFOIL [16], as the low-fidelity solver. For the higher-fidelity aerodynamic analysis we use the MSES CFD solver, which solves the Euler-equations.

2.4 Surrogate models

In this subsection the surrogate models used are briefly introduced. We mainly focus on Gaussian Process Regression (GPR) also called Kriging and a modified Kriging variant for higher-dimensional problems, and on Artificial Neural Networks (ANN). Furthermore, these methods are used in a variable and multi-fidelity framework (VF). For benchmarking purposes, Radial basis functions (RBF) and inverse distance weighting (IDW) are analyzed as well. For the benchmark process, with one variable-fidelity model exception we use the models implemented in the Surrogate Modeling Toolbox (SMT) [10].

2.4.1 Gaussian Process Regression (GPR) / Kriging (KRG)

GPR / KRG can be interpreted as a linear predictor, i.e. the function observation is a linear combination of basis functions, each centered at a training point:

$$\hat{f}(\mathbf{x}) = \sum_{i=1}^{n} \alpha_i k(\mathbf{x_i}, \mathbf{x}_*) \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^m$ is each training point, m denoting the input dimension, $\mathbf{x}_*$ being a test location for prediction and $\alpha = (K + \sigma_n^2 I)^{-1}\mathbf{y}$. $\mathbf{y}$ denotes the vector of observations. [17]

The basis functions $k(\cdot, \cdot)$ are also called kernel functions or covariance functions, which express a distance-based correlation between two sample locations. In this contribution we use the squared-exponential covariance function, which has the following form for one dimension:

$$k(x_p, x_q) = \sigma_f^2 \exp(-\frac{1}{2l^2}(x_p - x_q)^2) + \sigma_n^2 \delta_{pq} \tag{3}$$

σ_f^2 and σ_n^2 are signal and noise variance terms, and $l \in \mathbb{R}^m$ is the length scale parameter vector, with one parameter per dimension. These parameters are so-called hyper-parameters, which can be optimized, such that the GPR fits the observations.

2.4.2 Kriging partial least squares (KPLS)

KPLS is an advanced surrogate modeling technique that combines Kriging with Partial Least Squares regression. This method is particularly effective for high-dimensional problems where traditional Kriging approaches encounter problems due to the amount of required hyperparameters. The KPLS uses Partial Least Squares [18] to reduce the dimensionality of the input space, creating a set of uncorrelated latent variables that best explain the variance in both the input and output data. A Kriging model is then built using these latent variables instead of the original high-dimensional input space. [19]

While the normal Kriging uses one hyperparameter l for the scaling of the correlation function in the direction of each input dimension, the KPLS applies fewer hyperparameter l_k and orientates them in the direction of maximum covariance.

2.4.3 Artificial neural network (ANN)

ANN are biologically inspired surrogate models. They aim to replicate the structure and function of the brain consisting of layers of interconnected so-called nodes or neurons. Each node processes input data and applies a non-linear activation function to produce an output. [20]

ANNs typically consist of three types of layers: the input layer, hidden layers, and the output layer. One significant challenge in using ANNs as surrogate models is determining the appropriate topology of the network, that is, the number of layers and the number of nodes per layer. The SMT provides a Gradient-Enhanced Neural Network (GENN) model which is a Jacobian-Enhanced Neural Network (JENN). As this benchmarking problem does not contain gradient information, the SMT gradient parameter $\gamma = 0$ is set for all analysis. The most important tuning factor is the number of hidden layers and the number of nodes of each layer. In the scope of this work different rules for chosing a topology are compared. The default topology of the SMT toolbox $\langle 12, 12 \rangle$ means that the ANN consists of two hidden layers with 12 nodes each.

In [21] 2 hidden layers are proposed with $n_{1,2} = \sqrt{n_i * n_o}$ nodes in each hidden layer. n_i and n_o are the number of nodes in the input and output layer. For $n_i = 8$ and $n_o = 1$, this results in a $\langle 3, 3 \rangle$-topology. In [22] only one hidden layer with $n_1 = 2n_i + 1$ nodes is suggested. For 8 input dimensions, this results in a $\langle 17 \rangle$ topology. In [23] a two hidden layer case $\langle$Huang$\rangle$ is proposed in which the number of nodes is given by the following equations for a sufficient result. For the first hidden layer the number of nodes is $n_1 = \sqrt{(n_o + 2)nt} + 2\sqrt{nt/(n_0 + 2)}$ and in the second it is $n_2 = n_o\sqrt{nt/(n_o + 2)}$ with the number of training points nt.

These proposals for topologies are used in combination with variations over other parameters to identify a parameter set used for the actual analysis.

2.4.4 VF surrogate models – Variable-fidelity modeling (VFM)

For the utilization of samples produced with models of two (or more) different fidelities, so-called variable or multi-fidelity surrogate models can be used. For this purpose different strategies exist, such as additive or multiplicative methods, where a bridge function γ is used:

$$\hat{y}_{HF}(\mathbf{x}_*) = \hat{y}_{LF}(\mathbf{x}_*) + \gamma(\mathbf{x}_*) \tag{4}$$

Furthermore, a variable-fidelity surrogate model can be constructed, using the low-fidelity model as a mean function. The GPR prediction formulation from eq. (2) can be reformulated to:

$$\hat{f}(\mathbf{x}) = \sum_{i=1}^{n} \{ m(\mathbf{x_i}) + \left[k(\mathbf{x_i}, \mathbf{x_i}) + \sigma_n^2 \right]^{-1} \left[m(\mathbf{x_i}) - \mathbf{y_i} \right] \cdot k(\mathbf{x_i}, \mathbf{x}_*) \} \tag{5}$$

$m(\cdot)$ denotes the low-fidelity mean function and $m(\mathbf{x_i}) - \mathbf{y_i}$ the difference between the higher (HF) and the lower-fidelity (LF) model training points.

2.4.5 VF surrogate models – scaling with LF data

The scaling with LF data as mean model is not directly available within the SMT library but is straightforward to implement and offers specific advantages for the benchmarking use-case presented here. The ratio of LF to HF computational cost, often termed cost ratio $CR = time_{HF}/time_{LF}$ typically falls between 3 and 10 for many examples [24, 25]. However, in a conceptual aircraft design study example presented in [26], the cost ratio between a lifting-line model and a CFD model is 480. In the case study presented in this thesis, the CR is calculated at 26.3. Given the relatively low cost of LF data, direct incorporation of LF data into the surrogate model is feasible and can be advantageous.

The manual implementation of this approach involves computing either the difference or ratio between HF and LF training data to obtain ΔC_d. This dataset containing ΔC_d is then used to train basic surrogate models from the SMT library. For $\hat{C}_d$ prediction, the $\hat{\Delta} C_d$ prediction is either added to (for additive scaling) or multiplied by (for multiplicative scaling) the LF prediction generated using xFoil.

2.4.6 VF surrogate models – Multi-Fidelity Kriging (MFK)

The third variable-fidelity approach evaluated is the SMT Multi-Fidelity Kriging (MFK) model, which is a Co-Kriging model. It is applicable to the KRG, KPLS, and KPLSK Kriging variants. In this contribution, the MFK model is tested specifically with the KRG surrogate model. Three scaling options, constant, linear, and quadratic, are applied for the scaling/correlation function and compared to assess performance.

2.5 Evaluation metrics

Errors are computed as the difference between the surrogate model predicted $\hat{C}_d$ values and those calculated by MSES for corresponding test points. These HF test $C_{d,HF}$ values are stored within the test dataset. The final test dataset is constructed using LHS with 500 independent sample points, ensuring a robust and consistent evaluation framework. The primary error metric is the median of all absolute differences between predicted and actual values (MedAE). Furthermore, the average of the absolute differences between predicted and actual values (MeanAE) and the Root Mean Square Error (RMSE) are examined.

Wall clock time is chosen as a practical metric for evaluating training duration for the surrogate models. While a range of computational effort metrics exist, wall clock time is commonly used due to its simplicity. However, it has certain limitations, such as dependency on the hardware specifications and the system's current state (e.g. processor load, temperature, internal settings).

3 EVALUATION

Before analyzing the surrogate models, it is essential to examine the datasets used for training and testing. The quality and structure of the data plays a crucial role in determining how different surrogate models process information and impact prediction accuracy. In the latter part of this section, the results of the surrogate model evaluation are presented.

3.1 Accuracy of airfoil parametrizations

To assess the accuracy of the airfoil geometry parametrizations, the cumulative RMSE of the differences in the Z-direction between the seed airfoil and the parametrization is calculated at 160 tessellation points. The results for the 8192 training data seed airfoils are presented in Fig.3.

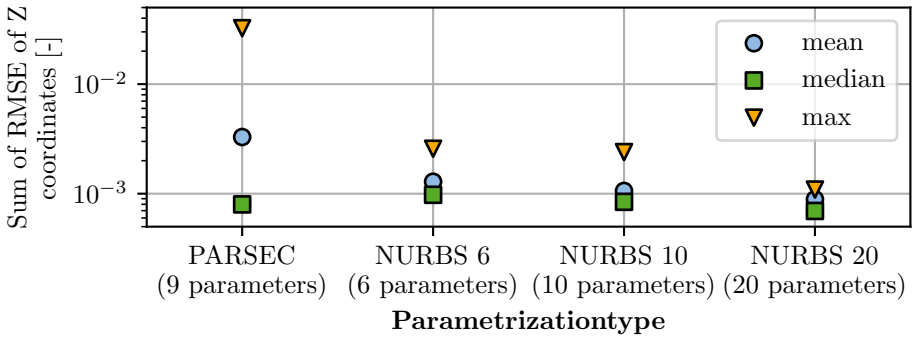


Figure 3: Cumulative RMSE of the differences in Z-direction at all coordinate points between the seed airfoil and the parametrization for the 8192 training data seed airfoils.

The results indicate that using more NURBS control points leads to better reproduction of the airfoils. On average, PARSEC performs worse than all NURBS configurations in terms of the mean and maximum error; however, its median error is lower than all others. This suggests that while a large portion of PARSEC results are highly accurate, a smaller portion performs very poorly. It is important to note, however, that PARSEC typically produces aerodynamically sound airfoils, whereas NURBS does not always guarantees this.

Even the worst result of all parametrizations corresponds to an RMSE per point in the Z-direction equivalent to 0.02% of the airfoil chord. Therefore, it can be assumed that the parametrizations are sufficiently close to the seed airfoils that the same aerodynamic drag calculations can be used reducing the computational effort.

3.2 Correlation of airfoil parameterizations

The presented PARSEC and NURBS parametrizations describe the sample airfoils using 6 to 20 geometric parameters while the initial seed airfoils are varied by only 5 parameters. For the evaluation of the surrogate models, it is important to understand the correlation between the parameters of the geometric parameterizations. The results of a correlation analysis are presented in Fig.4.

The first subplot 4a shows the correlation of the initial Sobol sequence for the 5 geometric parameters (Tab.1). No significant correlations are found across the samples,

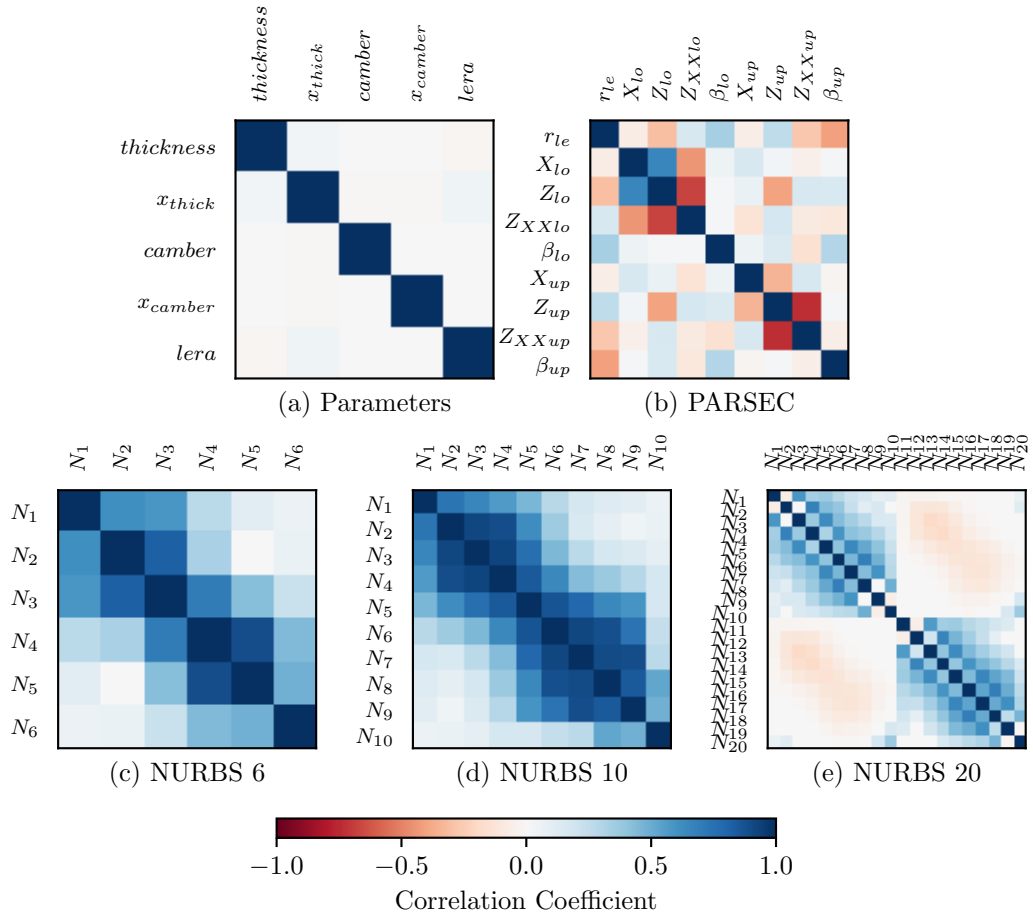


Figure 4: Correlation matrices for the variables of different geometric airfoil parametrization methods maping the 8192 Sobol seed airfoils. N_x indicates the corresponding NURBS parameters.

demonstrating that the Sobol sampling method generates a sampling set with a low discrepancy.

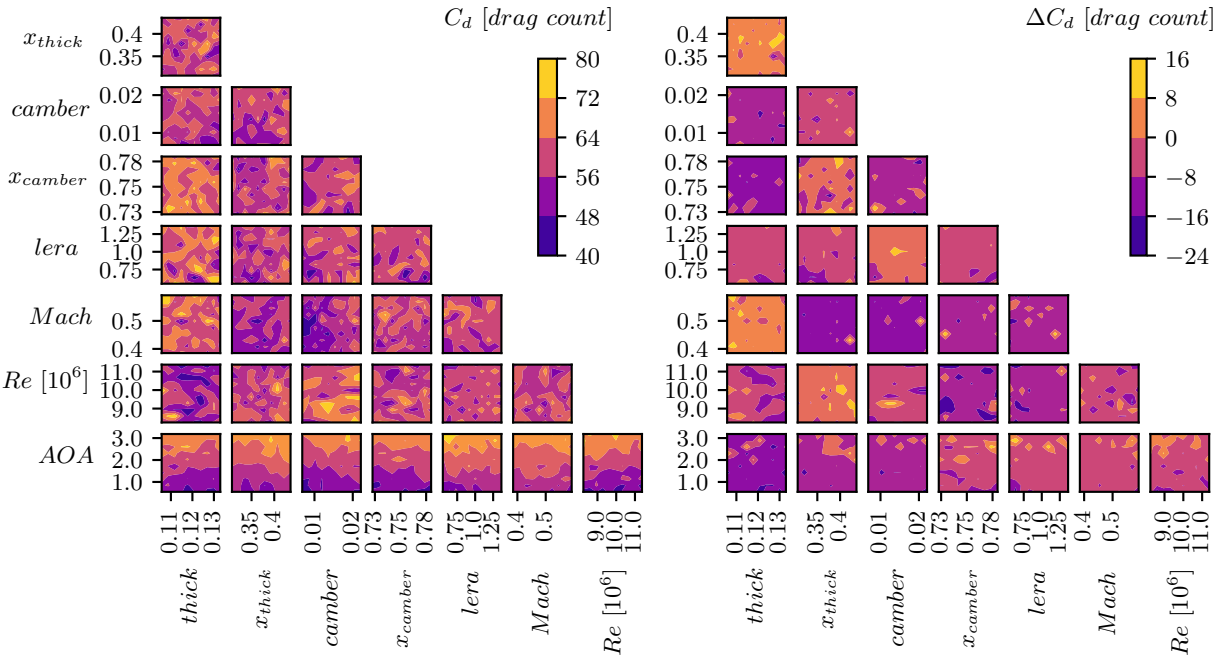
In the case of the PARSEC parameterization in subfigure 4b, strong negative correlations are observed between the upper and lower surface curvatures (Z_{XXup} and Z_{XXlo}), as well as between the trailing edge angles (β_{up} and β_{lo}).

For NURBS (subfigures 4c, 4d, and 4e), there are clearly distinct correlations for the upper and lower surface parameters (the first half of the parameters describes the upper surface and the second half the lower surface of the airfoil), with higher correlations in the upper-left and lower-right quadrants.

3.3 Correlation analysis of datasets

In Fig.5, grids of contour plots illustrate the distribution of $C_{d,HF}$ and ΔC_d within the training datasets. Each subplot shows the distribution of $C_{d,HF}$ (resp. ΔC_d) result for one combination of two input parameters. The parameter for each row and column are indicated on the left and bottom edge of the plot.

In subfig. 5a the direct values from the HF calculations are plotted. The absence of discernible structure in the subplots highlights the inherent complexity of this bench-



(a) $C_{d,HF}$ over all combinations of input parameters

(b) ΔC_d between $C_{d,HF}$ and $C_{d,LF}$ over all combinations of input parameters

Figure 5: Grid of contour plots mapping $C_{d,HF}$ and ΔC_d. The input parameter are defined for each row and column on the left and bottom edge of the plot. The contour values are subdivided in discrete color values for better clarity.

marking problem. An example of identifiable sensitivity and therefore relation can be seen in the subplots in which data is aranged by AOA (the last row in the grid). The lack of correlation between $C_{d,HF}$ and the other parameter combinations as well as the low correlation between the parameters themselves emphasizes the challenge posed by the data's unstructured nature for accurate model prediction.

Subfigure 5b provides contour plots of the drag coefficient difference, ΔC_d, between HF data (calculated with MSES) and LF data (calculated with xFoil). Notably, the variability in ΔC_d is significantly reduced compared to the broader range of variations in C_d itself. This reduced variability is particularly evident in the plots as a function of Mach number. xFoil does not account for Mach effects, whereas the HF data captures these dependencies accurately. Thus, while xFoil omits certain complex aerodynamic interactions, it effectively captures the primary influences of other variables. Incorporating xFoil data, therefore, simplifies the problem space, creating a more structured and interpretable dataset for surrogate model training and making it interesting for the incorporation of variable-fidelity models.

3.4 Basic model comprision

In Fig.6 the comparison of the basic surrogate models of single fidelity (QP, RBF, KRG, KPLS, GENN $\langle 12, 12 \rangle$, and GENN $\langle$Huang$\rangle$) is depicted for varying training sample sizes. The input-output relation for this evaluation is

$$\hat{C}_d = f(Ma,\ Re,\ AOA,\ t,\ x_t,\ x,\ x_c,\ lera). \tag{6}$$

The axis of abscissas represent the number of training points on a logarithmic scale. The plot shows the absolute drag count difference (equivalent to a C_d of 0.0001) between the predictions from surrogate models and the corresponding HF results at these points. The black line indicate the median absolute error (MedAE) for the reference LF data.

The second subplot shows the wall-clock time required for model training, with a base-10 logarithmic scale on the y-axis and a base-2 logarithmic scale on the axis of abscissas. The black lines indicate the required total wall clock time for the calculation of the corresponding number of LF and HF training points.

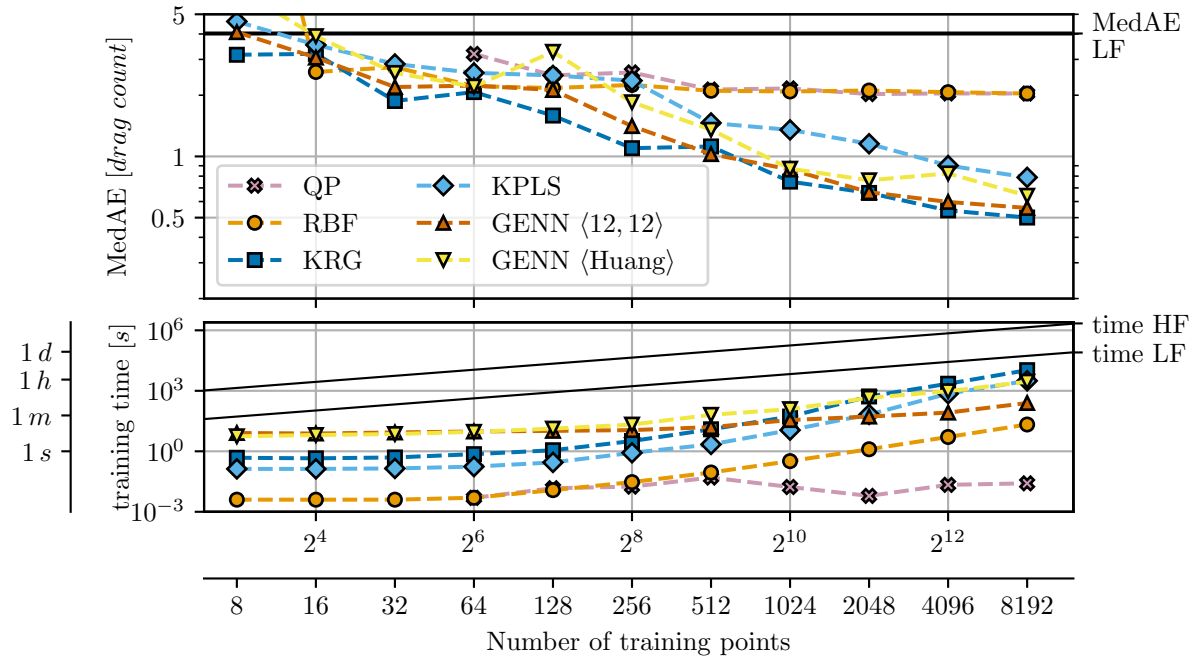


Figure 6: Comparision of the basic surrogate models. The error is between $\hat{C}_d$ prediction and HF test data $\tilde{C}_{d,HF}$. The references indicated on the right side are from the xFoil LF and MSES HF data.

For the QP model, data is only shown for 64 or more training points because its least-squares equations require a sufficient number of variables to avoid under-determination. As observed in preliminary studies, QP and RBF models show high deviations in predictions.

In terms of MedAE, the KRG model is only outperformed by the GENN $\langle 12, 12 \rangle$ model with 16 and 512 training points. KRG's MedAE is, on average, 0.7 drag counts (34%) better than that of the KPLS model. This relative improvement increases with larger sample sizes, indicating that KPLS is less capable of capturing fine details in the data due to its lower hyperparameter count.

For mean absolute error (MeanAE), the GENN $\langle 12, 12 \rangle$ model is the most accurate at and above 512 training points, averaging a 7.3% smaller error than KRG. The GENN $\langle$Huang$\rangle$ model outperforms the GENN $\langle 12, 12 \rangle$ in MedAE and RMSE only at 64 and 1028 points.

For datasets up to approximately 1024 training points, the ANN models require the most training time, whereas Kriging models become more time-intensive beyond this point. QP and RBF have the shortest training times due to their minimal parameter

requirements and smaller matrix sizes.

KRG and KPLS models scale in wall-clock time approximately with an order of $\mathcal{O}(n^{2.5})$, and in memory usage with an order of $\mathcal{O}(n^{1.6})$, as the number of training points increases. While KPLS scales similarly, it requires approximately three times less wall-clock time and 30% less RAM. However, KPLS requires 6 times as many training points for comparable prediction error, resulting in a longer effective total training time.

The GENN $\langle 12, 12 \rangle$ model scales linearly in wall-clock time $\mathcal{O}(n^1)$ and less than linearly in memory ($\mathcal{O}(n^{0.75})$). For the GENN $\langle$Huang$\rangle$ model, wall-clock time scales with $\mathcal{O}(n^{1.4})$ and RAM with $\mathcal{O}(n^{1.15})$, reflecting increased number of hidden nodes as training data grows.

Overall, the KRG and GENN $\langle 12, 12 \rangle$ models provide the best performance. KRG performs better with up to approximately 256 training points in terms of MeanAE and MedAE, while GENN $\langle 12, 12 \rangle$ is preferable for larger datasets in terms of all metrics. The QP and RBF models perform significantly worse than other models and are excluded from subsequent analysis for clarity.

3.5 Variable-fidelity surrogate models comparison

In this section the performance of several variable-fidelity surrogate models is compared. The input-output relation is the same as in Eq. (6). The definition of the surrogate model is extended by the additional LF data:

$$\hat{y} = f(\tilde{x},\ x_{HF},\ x_{LF},\ y_{HF},\ y_{LF}). \tag{7}$$

Figure 7 shows results for the SMT VFM and MFK models. All models are trained with all $2^{13} = 8192$ LF points from the training dataset. The HF data training set size is varied along the axis of the abscissa.

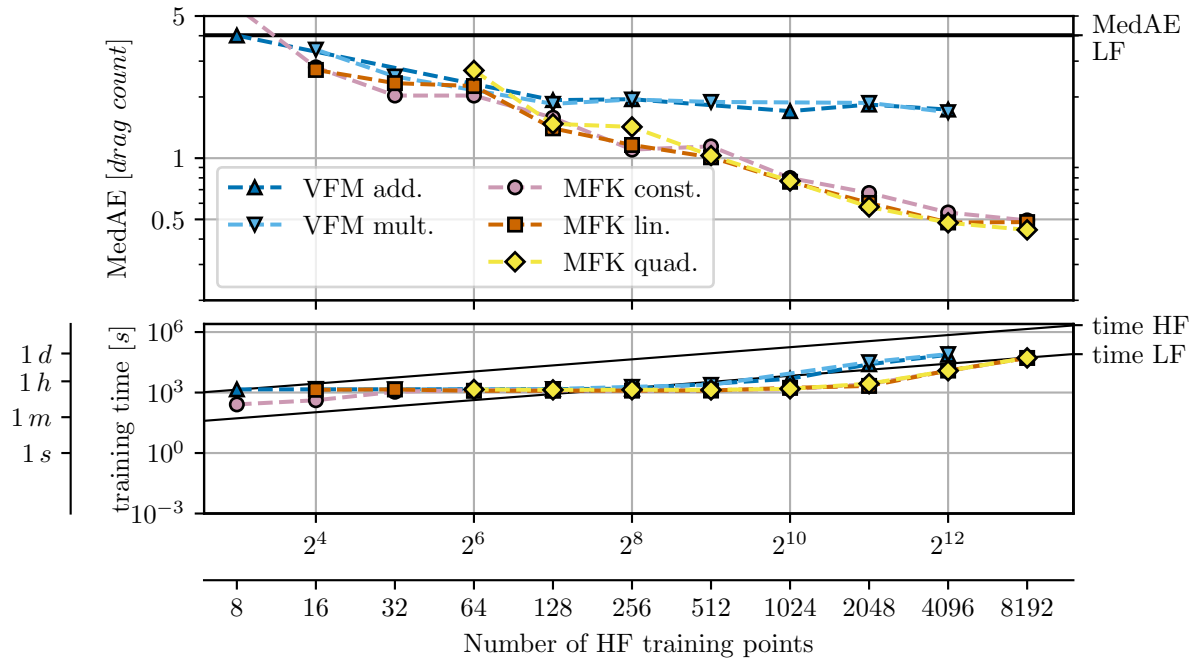


Figure 7: Comparision of the SMT VFM and SMT MFK model. The error is between $\hat{C}_d$ prediction and HF test data $\tilde{C}_{d,HF}$.

For the SMT VFM model the RBF model is used in this case as the LF surrogate because it is identified in preliminary studies as the best simple model. Tests with other simple LF functions resulted in larger prediction errors without improving the number of successful surrogate model trainings while KRG as LF function failed in most cases. The KRG model serves as the bridge model. ANNs are not available as LF or bridge models in this evaluation. Both additive and multiplicative bridge functions are tested.

The overall performance of the SMT VFM model is limited, with frequent instances of no successful results across test cases. Despite the substantial computational costs, the model's prediction error generally exceeds that of the basic KRG model.

The SMT MFK model is configured with the *squared exponential* correlation function (Eq. (3)) between LF and HF data. Three data series are compared, each using different correlation functions for $p(x)$: *constant*, *linear*, and *quadratic*.

For small training datasets (8–64 points), the MFK models with constant correlation factors show similar performance to the basic KRG model, with more consistent gains as dataset size increases. Between 128 and 512 HF training points, the MFK model with linear correlation factor offers an average improvement of 9% over the basic KRG model and 16% over the basic GENN $\langle 12, 12 \rangle$ model.

The primary drawback of the MFK model is its high computational cost, even with smaller HF datasets, due to the inclusion of the large LF dataset in each training session. The wall clock time scales approximately with $\mathcal{O}(n^{1.3})$ for HF dataset size and $\mathcal{O}(n^5)$ for the combined LF and HF training points.

Figure 8 shows the results for the VF scaling model using LF data as the mean. The basic SMT surrogate models are trained on the differences ΔC_d between HF and LF points, predicting $\hat{\Delta} C_d$, which is then added to the LF data from xFoil at the prediction point to obtain $\hat{C}_d$ (Eq. (5)).

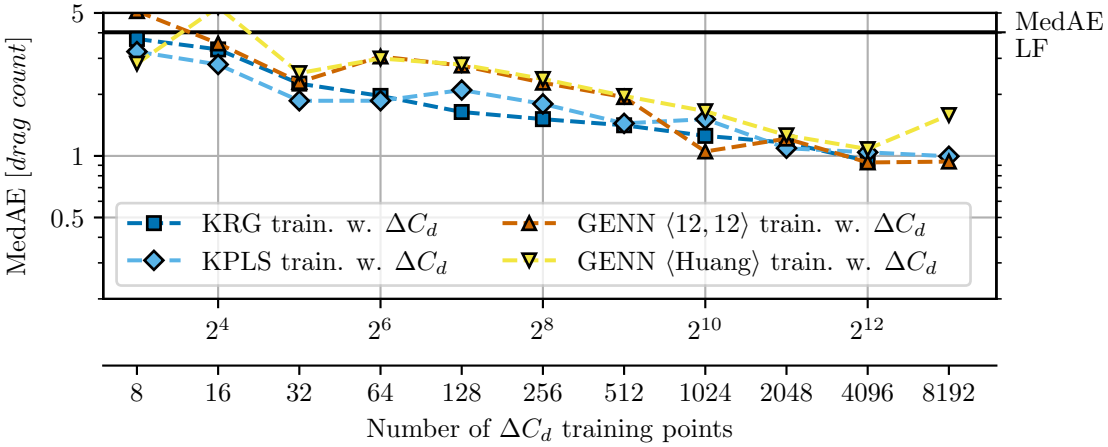


Figure 8: Comparision of VF additive scaling with LF data as mean surrogate models. The error is between $\hat{C}_d$ prediction and HF test data $\tilde{C}_{d,HF}$.

For datasets up to 64 training points, the KPLS model demonstrates the best performance in MedAE, MeanAE, and RMSE. Between 128 and 512 training points, the KRG model becomes the best-performing model across all metrics, while the GENN models only achieve comparable performance at 1024 points and beyond.

For smaller datasets (up to 64 points), VF KRG method performs similarly to basic KRG. In contrast, the GENN $\langle 12, 12 \rangle$ and $\langle$Huang$\rangle$ models do not show prediction error improvements with using the LF data as mean. For larger training sets of 1024 points and beyond, the prediction errors for all models are approximately 1.5 to 2 times worse than those of the basic models, likely due to the increased influence of the xFoil's model's inaccuracies in the final predictions.

Compared to the basic KPLS model the KPLS with LF as mean shows up to 512 training points an improvement in terms of MedAE and up to 64 training points an improvement in terms of MeanAE and RMSE. For smaller datasets of up to 64 training points, it surpasses the basic KRG model by approximately 10%, and it exhibits a consistent improvement with increasing training points. The reduced number of hyperparameters in the KPLS model aligns well with the simplified training data shown in the contour plots of Fig.5b. By training on the ΔC_d between LF and HF data, the complexity of certain variable relationships is reduced, offering a suitable basis for KPLS, which only employs distinct hyperparameters for the most significant data variances.

The required wall clock time and RAM usage are similar to those of the basic models, as the training data structure remains constant, differing only in target values. The initial calculation of ΔC_d has minimal impact on computation time.

The comparison of the three VF models with each other and with the basic surrogate models trained solely on HF data leads to the following conclusions: The SMT VFM model proved ineffective across all tested configurations, consistently performing worse than other models. Figure 9 shows a comparison of the best models for each range of number of (HF) training points.

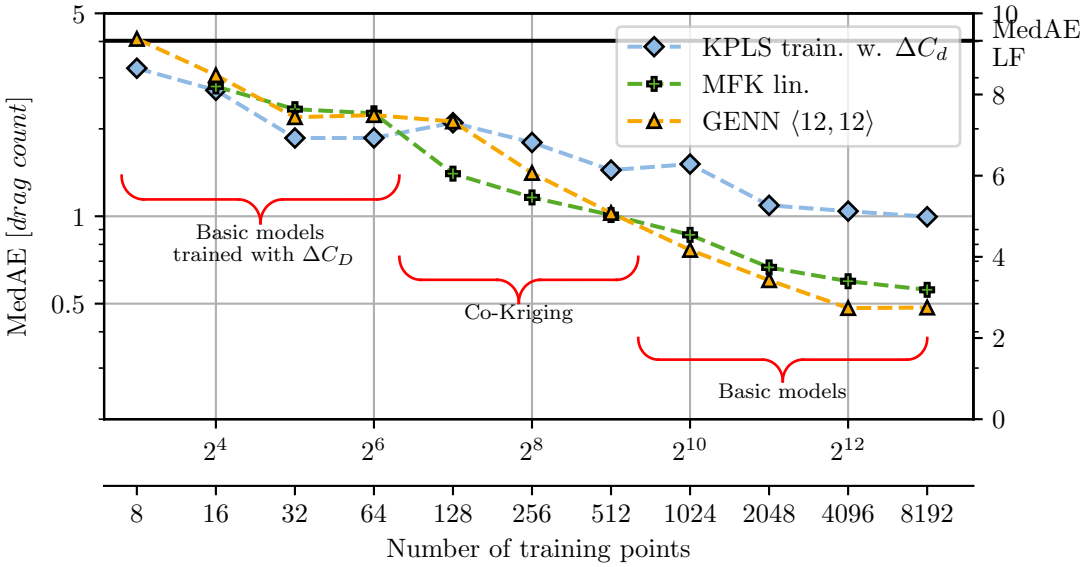


Figure 9: Categorization of best surrogate model types depending on the size of the HF training dataset.

For small HF training datasets (8–64 data points), directly incorporating LF data as a mean function achieved the best results. With sparse HF data, the LF data's ability to approximate relationships among variables becomes an asset, allowing it to capture key trends. For medium HF training datasets (128–512 data points), the co-Kriging SMT

MFK model deliveres the best outcomes. Here, the LF data trend is incorporated via the surrogate model, enhancing the HF data without introducing direct inaccuracies from the LF data itself. For large HF training datasets (1024–8192 data points), directly training basic surrogate models on the HF data alone provided the best balance between prediction accuracy and computational costs.

3.6 Comparison of geometric airfoil parametrizations

This section presents and analyzes the results of basic surrogate models trained with alternative geometric parametric representations of airfoils. Within this analysis, the single-fidelity surrogate models from the basic model comparison (subsection 3.4) are used. The input-output relations are

$$C_d = f(x) = f(\mathrm{Ma},\ \mathrm{Re},\ AOA,\ t,\ x_t,\ x,\ x_c,\ lera), \tag{8}$$

for the initial parametrization from the sampling strategy. For the PARSEC parametrization it is

$$C_d = f(\mathrm{Ma},\ \mathrm{Re},\ AOA,\ P1,\ P2,\ P3,\ P4,\ P5,\ P6,\ P7,\ P8,\ P9), \tag{9}$$

with the PARSEC parameters P_1 to P_9. The input-output relation for the NURBS parametrizations are

$$C_d = f(\mathrm{Ma},\ \mathrm{Re},\ AOA,\ N_1,\ N_2, \cdots,\ N_n), \tag{10}$$

with the NURBS parameters N_i, and $n = 6,\ 10,\ 20$. The number of HF training points for all evaluations is $nt_{HF} = 2^n$ with $n = 3, \cdots, 13$. Figure 10 shows the results of applying different parametrizations to the airfoils and training different single-fidelity models in the four subplots.

Compared to the initial parametrization, the KRG model with PARSEC parametrization shows a 32% increase in MedAE over the full range of training dataset sizes, while KPLS shows a 27% increase. The error differences between KRG and KPLS models are less pronounced with increasing number of NURBS parameters. According to the correlation analysis in Fig.4, several variables are strongly correlated, which supports the effective hyperparameter reduction in the KPLS model.

The GENN $\langle 12, 12\rangle$ model trained with the PARSEC parametrization performs 48% worse in MedAE compared to its performance with the initial parametrization. The GENN $\langle$Huang$\rangle$ model, however, shows a more modest reduction in accuracy, with a 24% decrease in MedAE. Notably, the GENN $\langle$Huang$\rangle$ model's relative deterioration diminishes as the training dataset size increases, unlike the other models. With NURBS 20, the number of training points at which GENN $\langle$Huang$\rangle$ matches GENN $\langle 12, 12\rangle$ increases from 64 (in the initial parametrization) to 512, highlighting that prediction performance varies with input dimensionality and suggesting the possibility of optimal topologies other than the $\langle 12, 12\rangle$ structure for different dimensional and dataset configurations. Therefore, additional optimization iterations would be required to find the best topology for each number of input dimensions.

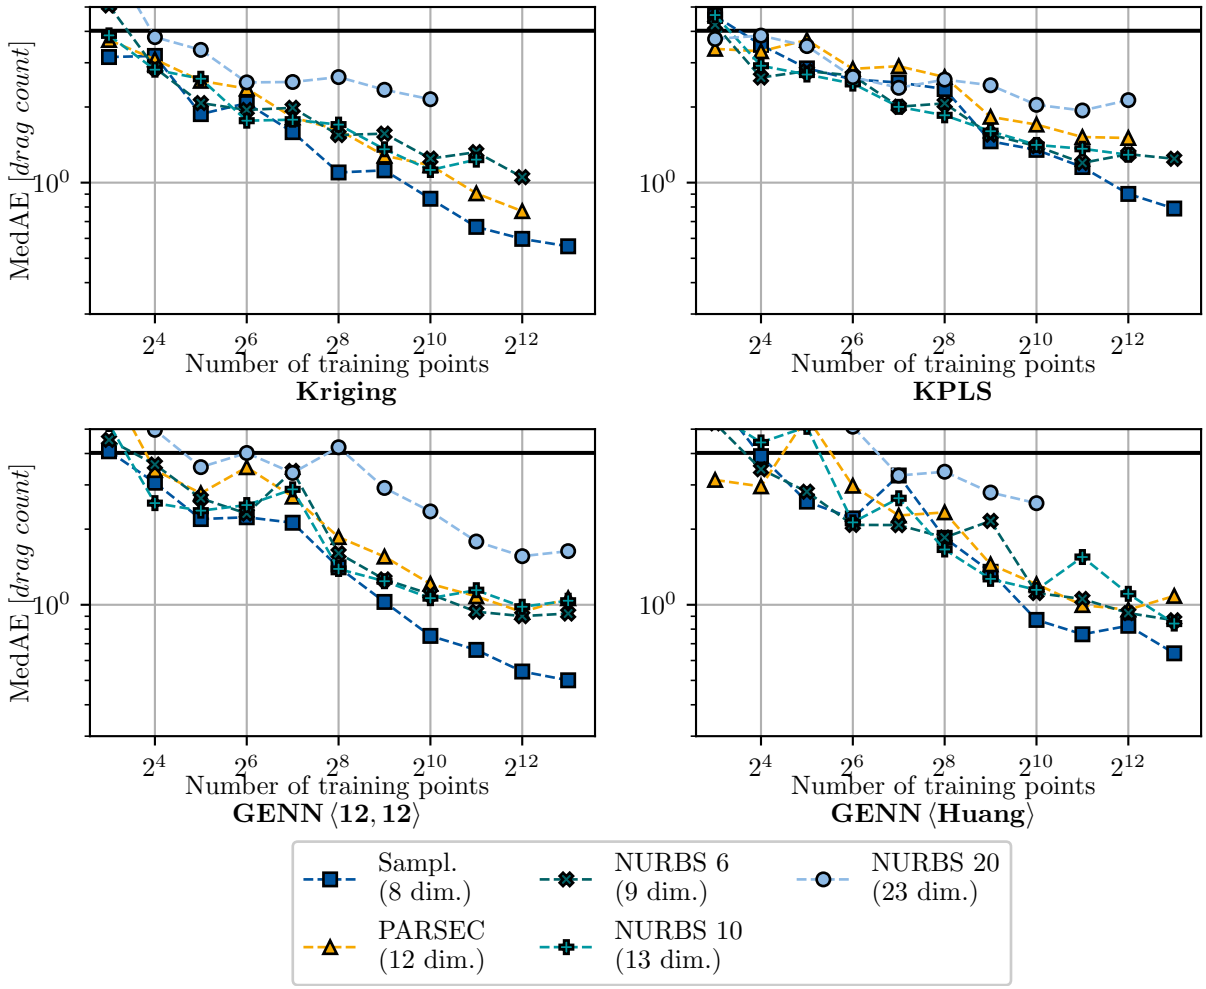


Figure 10: Comparision of for four single-fidelity surrogate models trained with different parametrized airfoils. The error is between $\hat{C}_d$ prediction and HF test data $\tilde{C}_{d,HF}$.

Overall NURBS 6 yields similar prediction errors and computational costs to the initial parametrization, though slightly higher due to 9 input dimensions instead of 8. The NURBS 10 parametrization provides prediction accuracy and computational costs comparable to those of the PARSEC parametrization, with a slight increase due to 13 input dimensions instead of 12. Thus, no definitive recommendation can be made between the PARSEC and NURBS parametrizations.

These different parametrizations also allow for an analysis of the computational scaling of surrogate models with varying input dimensions. For the KRG model, wall clock time scales approximately with $\mathcal{O}(n^6)$ relative to the number of input dimensions, while RAM scales with $\mathcal{O}(n^{0.8})$, indicating a less-than-linear increase. KPLS shows similar behavior, although the time-saving factor of KPLS relative to KRG increases with dimensionality. For the GENN $\langle 12, 12 \rangle$ model, computation time scales with $\mathcal{O}(n^{2.5})$ and RAM with $\mathcal{O}(n^{1.2})$. In the GENN $\langle$Huang$\rangle$ model, time scales with $\mathcal{O}(n^4)$, while RAM scales with $\mathcal{O}(n^{0.8})$.

Overall, increasing the number of input dimensions negatively impacts prediction accuracy and significantly raises computational costs for a fixed number of training points. However, as demonstrated with the KPLS model, combining a higher dimensionality than minimal required with a simpler model can help mitigate these disadvantages. The results for the two GENN models indicate that ANN models offer advantages over Kriging models, especially with larger training datasets. However, for an increasing number of input dimensions KRG and KPLS demonstrate to be superior in terms of MedAE to ANNs for sparse to medium dense training data ($< 2^7$ training points).

4 CONCLUSION

The analysis of surrogate modeling approaches for airfoil design highlights key insights regarding the influence of model complexity, parametrization, and training data requirements on both prediction accuracy and computational efficiency.

For ANN models such as GENN $\langle 12, 12 \rangle$ and GENN $\langle$Huang$\rangle$, large training datasets are advantageous, particularly in high-dimensional cases, though they require careful topology selection to balance model complexity with training time and prediction accuracy. Despite this, ANN models perform optimally only when substantial training datasets are available, whereas Kriging models exhibit more robust performance with smaller HF training sets. This finding underlines the potential value of Kriging for cases where higher-fidelity data is sparse and expensive to generate. By contrast, with larger datasets, the greater flexibility in ANN architectures generally provides superior results in both prediction and computational efficiency. As demonstrated with KPLS, simpler Kriging models can yield significant time savings when training time is limited. Although KPLS's training time is three to six times faster than that of KRG for the same number of points, it requires a larger dataset to achieve comparable accuracy, which leads to additional time costs.

Increasing the number of input dimensions, as seen with parametrizations like PARSEC and NURBS, generally has a negative effect on prediction accuracy and significantly raises computational costs for a given number of training points. However, simpler surrogate models like the KPLS model can effectively leverage high-dimensional parametrizations, helping to offset these drawbacks by reducing the number of hyperparameters, which lowers both model complexity and computational load.

The effectiveness of various VF models varies depending on the size of the HF training set. For small HF datasets, VF models that use LF data as a mean perform best, while for medium-sized datasets, Co-Kriging with the SMT MFK model proves most effective. Direct training on HF data is optimal for larger datasets, suggesting that the classification of datasets into small, medium, and large in this study could be adjusted based on application-specific factors such as input dimensionality, though the overall trends are likely broadly applicable.

ACKNOWLEDGEMENTS

Simulations were performed with computing resources granted by RWTH Aachen University under the project thes1664.

REFERENCES

[1] E. Torenbeek. *Synthesis of subsonic airplane design: an introduction to the preliminary design of subsonic general aviation and transport aircraft, with emphasis on layout, aerodynamic design, propulsion and performance*. Springer Science & Business Media, 2013.

[2] D. Böhnke, B. Nagel, and V. Gollnick. An approach to multi-fidelity in conceptual aircraft design in distributed design environments. In *2011 Aerospace Conference*, pages 1–10. IEEE, 2011.

[3] A. Forrester, A. Sobester, and A. Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

[4] T. Maxim and L. Jae-Woo. Efficient multi-response adaptive sampling algorithm for construction of variable-fidelity aerodynamic tables. *Chinese Journal of Aeronautics*, 32(3):547–558, 2019.

[5] G. Cerulli. *Fundamentals of supervised machine learning: with applications in Python, R, and Stata*. Springer Nature, 2023.

[6] RWTH Aachen University IT Center. Rwth hochleistungsrechnen, 2024. URL `https://www.hpc.itc.rwth-aachen.de/cms/~bcenat/hpc/`.

[7] H. Schiefer and F. Schiefer. Statistische versuchsplanung, design of experiments (doe). *Statistik für Ingenieure: Eine Einführung mit Beispielen aus der Praxis*, pages 1–20, 2018.

[8] J. R. R. A. Martins and A. Ning. *Engineering design optimization*. Cambridge University Press, Cambridge, 2022. ISBN 9781108833417. doi:10.1017/9781108980647. URL `https://zbmath.org/1480.74002`.

[9] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL `http://www.scipy.org/`.

[10] P. Saves, R. Lafage, N. Bartoli, Y. Diouane, J. Bussemaker, T. Lefebvre, J. T. Hwang, J. Morlier, and J. R. R. A. Martins. SMT 2.0: A surrogate modeling toolbox with a focus on hierarchical and mixed variables gaussian processes. *Advances in Engineering Sofware*, 188:103571, 2024. doi:https://doi.org/10.1016/j.advengsoft.2023.103571.

[11] M. Drela. Xfoil: An analysis and design system for low reynolds number airfoils. In *Low Reynolds Number Aerodynamics: Proceedings of the Conference Notre Dame, Indiana, USA, 5–7 June 1989*, pages 1–12. Springer, 1989.

[12] H. Sobieczky. Parametric airfoils and wings. In *Notes on Numerical Fluid Mechanics 1998*, volume 68, pages 71–88. Vieweg Verlag, 1998.

[13] J. A. Cottrell, T. J. Hughes, and Y. Bazilevs. *Isogeometric analysis: toward integration of CAD and FEA*. John Wiley & Sons, 2009.

[14] S. Mirjalili, S. M. Mirjalili, and A. Lewis. Grey wolf optimizer. *Advances in engineering software*, 69:46–61, 2014.

[15] H. Schlichting and E. A. Truckenbrodt. *Aerodynamik des Flugzeuges: Erster Band Grundlagen aus der Strömungsmechanik Aerodynamik des Tragflügels (Teil I)*. Springer-Verlag, 2013.

[16] M. Drela. Mses overview, 2015. URL `https://web.mit.edu/drela/Public/web/mses/`.

[17] C. E. Rasmussen. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, Cambridge, Mass., 2006. ISBN 978-0-262-25683-4. doi:10.7551/mitpress/3206.001.0001?locatt=mode:legacy.

[18] H. O. A. Wold. Soft modelling by latent variables: The non-linear iterative partial least squares (nipals) approach. In *Perspectives in probability and statistics*, volume 12, pages 520–540. Routledge, 1977. doi:10.1080/0031383770210103.

[19] M. A. Bouhlel, N. Bartoli, A. Otsmane, and J. Morlier. An improved approach for estimating the hyperparameters of the kriging model for high-dimensional problems through the partial least squares method. *Mathematical Problems in Engineering*, 2016:1–11, 2016. ISSN 1024-123X. doi:10.1155/2016/6723410.

[20] S. Freund. *Metamodelle mehrerer Genauigkeitsstufen zur Stabilitätsanalyse versteifter Schalen aus Faserverbundwerkstoffen*. Phd thesis, Technischen Universität Carolo-Wilhelmina zu Braunschweig, Braunschweig, 2020.

[21] H. S. Sachdev. Choosing number of hidden layers and number of hidden neurons in neural networks, 2020. URL `https://www.linkedin.com/pulse/choosing-number-hidden-layers-neurons-neural-networks-sachdev/`.

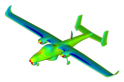
[22] R. Hecht-Nielsen. Kolmogorov's mapping neural network existence theorem, 1987.

[23] G.-B. Huang. Learning capability and storage capacity of two-hidden-layer feedforward networks. *IEEE transactions on neural networks*, 14(2):274–281, 2003. ISSN 1045-9227. doi:10.1109/TNN.2003.809401.

[24] X. Ruan, P. Jiang, Q. Zhou, J. Hu, and L. Shu. Variable-fidelity probability of improvement method for efficient global optimization of expensive black-box problems. *Structural and Multidisciplinary Optimization*, 62(6):3021–3052, 2020. ISSN 1615-147X. doi:10.1007/s00158-020-02646-9.

[25] L. Shu, P. Jiang, and Y. Wang. A multi-fidelity bayesian optimization approach based on the expected further improvement. *Structural and Multidisciplinary Optimization*, 63(4):1709–1719, 2021. ISSN 1615-147X. doi:10.1007/s00158-020-02772-4.

[26] L. Messner, M. Zimmnau, and E. Stumpf. *Recherche und Analyse zu aktuellen Methoden zur Erstellung von Ersatzmodellen variabler Genauigkeit für die aerodynamische Analyse im konzeptionellen Flugzeugvorentwurf*. Bachelor thesis, RWTH Aachen University, Aachen, 2022.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

ON THE AEROACOUSTIC FEEDBACK LOOP IN SUPERSONIC IMPINGING JETS

António Rebelo[†1], Maxime Fiore[‡2] and Fernando Lau[†3]

†: Instituto Superior Técnico
Universidade de Lisboa
1049-001 Lisboa, Portugal

‡: ISAE-SUPAERO
Université de Toulouse
31055 Toulouse, France

Abstract. *An Aeroacoustic Feedback Loop (AFL) generates intense tonal noise in supersonic impinging jet flows, potentially impacting the integrity of structures, noise emissions and even aircraft performance. This work investigates the AFL closure mechanism by analysing wave propagation data through a numerical model and a stability simulation of the supersonic jet column, aiming to identify the wave mode responsible for closing the feedback loop. A two-stage methodology is applied: first, a one-dimensional jet containing a mixing region is modelled by coupling a differential equation with an analytical far-field solution, reformulating the problem as an Initial Value Problem. Second, a bi-global (2D) stability analysis using an in-house Finite Element Method solver, adapted for compressible jet flows, is performed on an analytical 2D jet flow, with wave mode stability examined through a complete post-processing framework that identifies upstream and downstream-propagating jet modes. Results confirm that upstream-propagating modes responsible for AFL closure (so-called "closing modes") are guided by the jet rather than propagating purely in the quiescent medium or within the jet column. Stability analysis reveals the presence of a high number of unstable eigenvalues, attributed to boundary condition artefacts, indicating the need for an update on the simulation's numerical setup. The predicted frequency ranges for the closing modes align well with experimental data, supporting the validity of the approach. The developed post-processing framework for stability studies,*

Keywords: Aeroacoustics, Aeroacoustic Feedback Loop, Supersonic Jet, Global Stability Analysis and Jet Column Modelling.

[1]Master in Aerospace Engineering; antonio.rebelo@tecnico.ulisboa.pt

[2]Associate Professor at *Département Aérodynamique, Énergétique et Propulsion* (DAEP); maxime.fiore@isae-supaero.fr

[3]Associate Professor at *Centro de Ciências e Tecnologias Aeronáuticas e Espaciais* (CCTAE); lau@tecnico.ulisboa.pt

1 INTRODUCTION

Aeroacoustics is the study of noise generation by fluid flows, particularly when interacting with solid surfaces, and its propagation through a medium. Since its emergence in the 1950s, with the rise of jet propulsion, Aeroacoustics has remained an important area of research, driven by the need to understand and reduce the noise produced by high-speed fluid flows. The noise discussed in this manuscript is of tonal nature and can originate from various aircraft components, such as wing flaps and shallow or deep cavities (e.g., airframe gaps or landing gear bays). It can also result from the interaction between aircraft-generated flows and the surrounding environment, including impinging jet columns and jet shock cells. All of these aircraft-generated tones share a common underlying mechanism, often referred to as the Aeroacoustic Feedback Loop (AFL).

As the name suggests, this self-sustained cycle results from the interaction between aerodynamic flow and acoustic waves under certain conditions. As aerodynamic fluctuations can produce sound [1], and in turn, these sound waves can generate (more) aerodynamic perturbations, at a certain point, the phenomenon ties both mechanisms in a resonant state of the system, characterized by the emergence of intense acoustic tones. These tones ought to be well understood as, apart from the intense Sound Pressure Level (SPL) produced, they can induce cyclic vibrations on structures leading to unwanted loading. The feedback mechanism has also been linked to high losses of lift (as much as 40%) and increases in near-field noise (6 to 8 dB) in Short Take-Off and Vertical Landing aircraft [2].

The study of the intense tones produced in jet impingement "took off" with the work of Powell [3] in 1953. Powell took upon the efforts of previous authors to explain the physical mechanism behind the AFL and, unsatisfied with the lack of quantitative evidence surrounding the topic, he presented a formulation for the appearance of these tones. In jets, the feedback loop was separated into two different stages. Firstly, aerodynamic disturbances are convected downstream with the jet flow and are amplified within the system, eventually generating acoustic waves. These waves are generated at the system's singularity — a wall in the jet impingement case — following the theory proposed by Sir James Lighthill [1] that momentum fluctuations across a fixed surface convert some of the flow's kinetic energy into sound waves. In the second stage, these acoustic waves propagate upstream and interact with the jet flow, creating more aerodynamic disturbances, thus closing the loop. This interaction allows the system to eventually stabilize at a limit cycle — a state of stable oscillation where the acoustic tones maintain a constant frequency $\omega = 2\pi f$ over time.

Figure 1 illustrates the steps of the AFL: (1) An aerodynamic perturbation convects along the jet's shear layer, amplified by Kelvin-Helmholtz instabilities [3], until (2) it interacts with the system's instability, converting part of its energy into acoustic waves. (3) These acoustic waves then propagate upstream along the jet column at a frequency f until reaching the nozzle, where (4) they trigger new perturbations in the shear layer at the same frequency. After reaching the limit cycle state, the resulting tones are recorded at this same frequency.

With his formulation, Powell [3] defined two criteria that must be met for the AFL to be self-sustained: one of phase and another of amplitude. The phase criterion is written

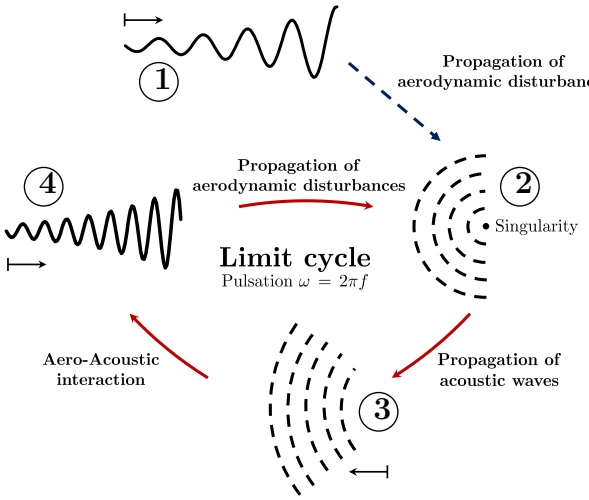


Figure 1: Scheme of the AFL mechanisms.

as

$$\frac{N+\phi}{f} = \frac{L}{U_h} + \frac{L}{U_a}\,, \tag{1}$$

stating that the time it takes for an aerodynamic instability to arrive at the singularity, propagating through the nozzle-to-singularity distance L at velocity U_h, plus the time taken for an acoustic wave to leave the singularity and reach the nozzle (L/U_a) is the period T_0 of a complete AFL cycle, which can be rewritten as $T_0 = NT$. Here, the integer N is the number of aero or acoustic (not both) disturbances generated during T_0 and $T = 1/f$ is the period of the AFL oscillation. The lag term ϕ is added since the energy conversion, wave propagation and excitation processes are not instant.

The amplitude criterion states that, for the feedback loop to be steadily sustained, the amplitude of a newly generated disturbance must match with the previous. This implies that when the system reaches its limit cycle, the product of the gains from all steps in Fig. 1 must equal one, such that:

$$\eta_4\,\eta_2\,\eta_3\,\eta_{AA} = 1\,, \tag{2}$$

where terms η are the efficiencies of each step in Fig. 1 (η_{AA} is the efficiency of the aero-acoustic interaction at the nozzle).

As a consequence of having to satisfy both Eqs. (1) and (2), a characteristic of the AFL is the tendency to exhibit *staging* behaviour, where small changes in operating conditions lead to large jumps in the frequency of the intense tones. Since N must be an integer value, and as at least some of the terms in Eq. (2) are frequency dependent, at some conditions a maximization of Eq. (2) will be achieved by a change in N, and thus a sudden change in frequency, according to Eq. (1). This behaviour is observed in Fig. 2 from the work of Edgington-Mitchell [4], where the intense tones (patches of pink) show large jumps in frequency when the nozzle-to-singularity distance changes around $z/D_j = \{2.15, 2.8\}$, with z being the jet's streamwise direction, measured from the nozzle and D_j the jet's nozzle diameter. Changes in jet Mach number (M) or nozzle radius R_j can also lead to *staging*.

Since the AFL is extremely sensitive to the problem's parameters, the development of a post-processing tool such as the one later presented in §3.2 is crucial for capturing and

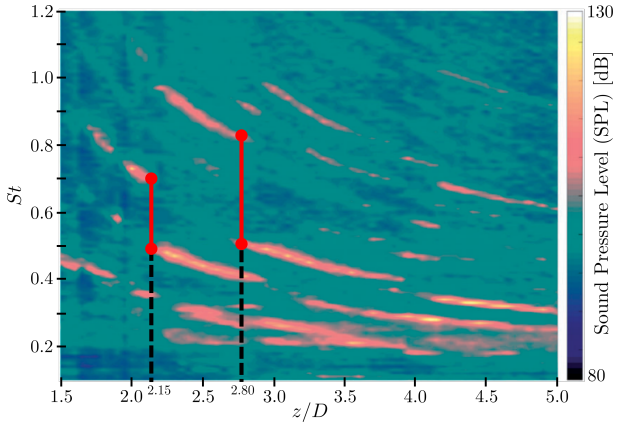


Figure 2: SPL as a function of frequency and distance [4].

analysing staging events, enabling their incorporation into optimization frameworks for aerospace systems, where discrete frequency transitions must be predicted and managed effectively.

1.1 State of the Art

Since Powell's seminal work in 1953, research on the Aeroacoustic Feedback Loop (AFL) has expanded across multiple fields, including mathematical modelling, experimental studies, computational fluid dynamics (CFD), and, more recently flow stability analysis. These efforts have sought to clarify the mechanisms behind intense tone generation in the AFL, with this work specifically focusing on the closure of the feedback loop.

Mathematical modelling of the AFL advanced significantly with the work of Tam and Hu [5], who introduced the vortex-sheet model (VSM). This analytical model of the jet column suggested that upstream-propagating waves, capable of closing the loop, were intrinsic to the jet but propagated partially over the free-stream, rather than just propagating in the free-stream, outside of the jet column, contradicting the earlier hypothesis. In the same work, Tam and Hu [5] developed a numerical model for a jet with a mixing layer (the VSM did not have one). This "more realistic" finite-thickness model (FTM) retrieved very similar results, suggesting that indeed the waves of interest (so called "closing modes") propagate in the free-stream region but are guided upstream by the jet.

Experimental studies have played a key role in validating these findings. Lepicovsky and Ahuja [6] provided evidence by showing that feedback frequencies remained insensitive to the outer flow velocity, suggesting that closure occurs within the jet column. Further pressure measurements and particle image velocimetry studies confirmed that coherent structures in the shear layer (hydrodynamic part of the AFL) convect at approximately 60% of the jet velocity.

The advent of CFD allowed researchers to investigate the AFL at a greater level of detail. Initially, simulations focused on replicating experimental observations, affirmed CFD as a tool for AFL studies. More recently, high-fidelity Large Eddy Simulations (LES) have provided deep insights into flow dynamics, particularly in supersonic impingement cases. Studies Bogey and Gojon [7] confirmed that Tam and Hu's closing modes are valid in supersonic regimes and also confirmed the convection velocity of hydrodynamic instabilities previously obtained ($\approx 60\%$). Some other CFD studies successfully obtained

a lower bound in Mach number for the appearance of the intense tones ($M \approx 0.75$). LESs also reproduced staging effects, convective velocities of shear-layer structures, and frequency distributions consistent with experimental findings.

Analysis of jet flow stability, both local (1D) and bi-global (2D/3D), have further enhanced the understanding of the AFL in impinging jets. These methods identify which disturbances are amplified within the flow, predicting dominant wave modes and their growth rates. Recent bi-global stability studies have supported the presence of closing modes, reinforcing their role in AFL closure [8, 9]. However, open questions remain regarding mode interactions, nonlinear effects, and the coexistence of multiple tones, requiring further understanding.

These advances have significantly improved the understanding of the AFL in supersonic impinging jets, yet challenges persist in fully characterizing the wave dynamics responsible for feedback-loop closure. This is primarily due to the AFL's strong dependence on the problem's specific parameters, which complicates its study and generalization.

1.2 Research Context

This work is the result of a collaborative effort between IST Lisboa and ISAE-SUPAERO. While the work was developed at both institutions, through a research internship in Toulouse, ISAE-SUPAERO provided key scientific resources and research infrastructure, playing a crucial role in the work development.

2 WAVE ANALYSIS

The closure of the Aeroacoustic Feedback Loop (AFL), is first studied through a type of jet model presented by Tam and Hu [5], where only the jet column is modelled. The objective is to understand which waves can propagate inside the jet column and the conditions required for their existence. First, a model without a mixing region in the jet column is introduced as a reference: the vortex-sheet model (VSM). This purely analytical model developed by Tam and Hu [5], depicted in Fig. 3, is represented by the green step-like velocity profile, where the transition from the jet flow to the surrounding quiescent medium occurs abruptly.

The second model considered in this section of the work, is an identical model of the jet column but with the presence of a finite-thickness mixing region: finite-thickness model (FTM). This more advanced model was also developed by Tam and Hu [5] but, since it is not analytical, an algorithm was developed to solve the problem and obtain the results through numerical processes. The jet modelled in the FTM is represented in Fig. 3 as three different blue velocity profiled with different shear-layer thicknesses.

In both models, the singularity and the nozzle are not included, yet their presence is connected to the existence of some of the waves propagating within the jet column.

The two jet models consider a supersonic jet column system which evolves from its initial state ("base flow") when perturbed by an external periodic forcing of frequency ω. The generated small-amplitude waves are then studied through: 1) the analysis of their dispersion; and 2) their radial structure $\hat{p}(r)$. What distinguishes the two models is in fact the base flow. For the VSM, it is built by the jet column of constant thickness (equal to the jet nozzle diameter D_j) travelling at the jet exit velocity U_j and by the medium at rest when $r > R_j$. The FTM base flow assumes a slight velocity gradient between these two sections of the flow.

Assuming homogeneity in the streamwise (z), azimuthal (θ) and temporal (t) directions, the pressure is decomposed into normal modes as

$$p(r,\theta,z,t) = \sum_m \hat{p}(r)\exp\left\{i(\alpha z + m\theta - \omega t)\right\}\,, \tag{3}$$

with $\hat{p}(r)$ being the radial pressure amplitude and α and ω being respectively the wavenumber and angular frequency of the normal mode. Pressure is decomposed into its azimuthal components (m) with each one being responsible for a different oscillation geometry in the θ direction: $m = 0$ corresponds to an axisymmetric mode; $m = 1$ to a helical mode and so on. In this work, only the first two azimuthal modes are being studied since they are responsible for the production of the most dominant tones in jet impingement [2, 7].

2.1 Vortex-Sheet Model (VSM)

The construction of the VSM starts from the linearized continuity, momentum and energy equations of a compressible inviscid fluid. The pressure associated with the waves propagating inside the column is obtained through the superimposition of small-amplitude disturbances on the pre-defined base flow, obtaining two separate equations, naturally split by the discontinuity at the shear-layer ($r = R_j$). These pressure equations (p_{in} describing the pressure "inside" the jet and p_{out} describing the pressure "outside" of the jet) are the convective wave equation for the jet column and the wave equation for the outside section of the jet, as presented by Tam and Hu [5]. At the vortex-sheet (transition between solutions), dynamic and kinematic boundary conditions are defined, using ζ as the vortex-sheet displacement, measured from its resting position ($r = R_j$), assuming the same decomposition as in Eq. (3). Thus, at $r = R_j$:

$$p_{in} = p_{out}\,, \tag{4a}$$

$$-\frac{1}{\rho_{out}}\frac{\partial p_{out}}{\partial r} = \frac{\partial^2 \zeta}{\partial t^2}\,, \tag{4b}$$

$$-\frac{1}{\rho_{in}}\frac{\partial p_{in}}{\partial r} = \left(\frac{\partial}{\partial t} + U_j\frac{\partial}{\partial z}\right)^2 \zeta\,, \tag{4c}$$

with ρ being the density of the fluid either inside or outside the jet column.

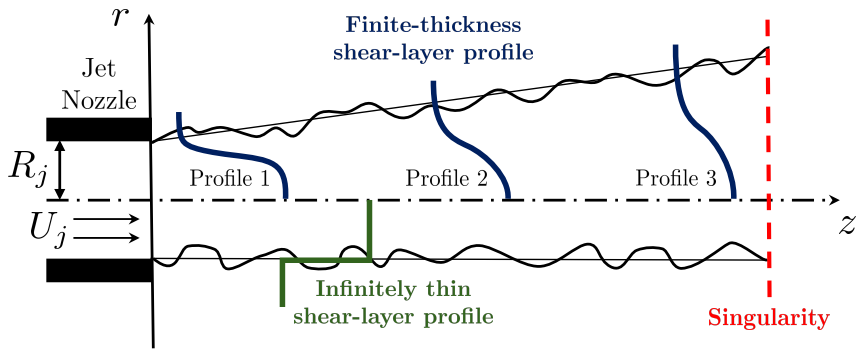


Figure 3: Scheme of the jet models implemented.

Tam derived an analytical pressure solution that satisfies both the system's governing equations and boundary conditions, expressed as:

$$\begin{bmatrix} p_{out} \\ p_{in} \end{bmatrix} = \begin{bmatrix} B\, H_m^{(1)}(i\eta_o r) \\ A\, H_m^{(1)}(i\eta_o R_j)\ J_m(\eta_i r)/J_m(\eta_i R_j) \end{bmatrix} \exp\left\{i(\alpha z + m\theta - \omega t)\right\} . \tag{5}$$

This solution, proved to be bounded at $r = 0$ and $r \to \infty$, is dependent on the Bessel (J_m) and Hankel ($H_m^{(1)}$) functions of the first kind and on quantities $\eta_o = (\alpha^2 - \omega^2/a_{out}^2)^{\frac{1}{2}}$ and $\eta_i = ((\omega - U_j\alpha)^2/a_{in}^2 - \alpha^2)^{\frac{1}{2}}$. Branch cuts are done to these two quantities so that only physical solutions of the jet column are considered:

$$-\frac{\pi}{2} < \arg(\eta_o) \leq \frac{\pi}{2}\,, \quad 0 \leq \arg(\eta_i) < \pi\,. \tag{6}$$

In Eq. (5), A and B are arbitrary constants, taken to be 1.

The dispersion relation which identifies all the waves propagating in the jet problem is given by the expression

$$D_m(\alpha,\omega) \equiv \frac{i\eta_o}{\rho_o\omega^2} J_m(\eta_i R_j) H_m^{(1)'}(i\eta_o R_j) - \frac{\eta_i}{\rho_j(\omega - U_j\alpha)^2} H_m^{(1)}(i\eta_o R_j) J_m'(\eta_i R_j) = 0\,, \tag{7}$$

which is obtained from the coupling of the boundary conditions with the presented pressure solution. Each azimuthal mode (m) has its own dispersion relation, meaning that the waves are identified for each type of azimuthal oscillation. For each pair (α, ω) that satisfies Eq. (7), a wave mode propagates inside the jet system with a certain radial structure $\hat{p}(r)$. The dispersion relation is shown after the more complex FTM model is presented.

2.2 Finite-Thickness Model (FTM)

The finite-thickness model (FTM) also stems from the linearized continuity, momentum and energy equations of a compressible inviscid fluid, but, instead of the fluctuations being superimposed on a base flow with two distinct regions (in/out of the column), the base flow contains a mixing layer, allowing for a smooth transition from the high-speed jet core up to the resting medium surrounding the main flow. This superimposition generates a set of waves propagating in the system with a radial structure given in Tam's work by:

$$\frac{d^2\hat{p}}{dr^2} + \left[\frac{1}{r} - \frac{1}{\overline{\rho}}\frac{d\overline{\rho}}{dr} + \frac{2\alpha}{\omega - \overline{u}\alpha}\frac{d\overline{u}}{dr}\right]\frac{d\hat{p}}{dr} + \left[\frac{(\omega - \overline{u}\alpha)^2}{\overline{a}^2} - \frac{m^2}{r^2} - \alpha^2\right]\hat{p} = 0\,. \tag{8}$$

All the quantities with a bar superscript are related to the base flow ($\bar{Q}$). In the FTM, the mixing layer is defined by an hyperbolic tangent profile already implemented by other authors of the literature. This profile is used so that comparisons are easily done. The velocity along the radial direction is then written as:

$$\frac{u(r)}{U_j} = \frac{1}{2}\left\{1 + \tanh\left[\frac{R_j}{4\theta_R}\left(\frac{R_j}{r} - \frac{r}{R_j}\right)\right]\right\}, \tag{9}$$

with θ_R being the momentum loss boundary layer thickness. Temperature is defined, from the velocity profile, with the Crocco-Busemann Law and density is obtained with the

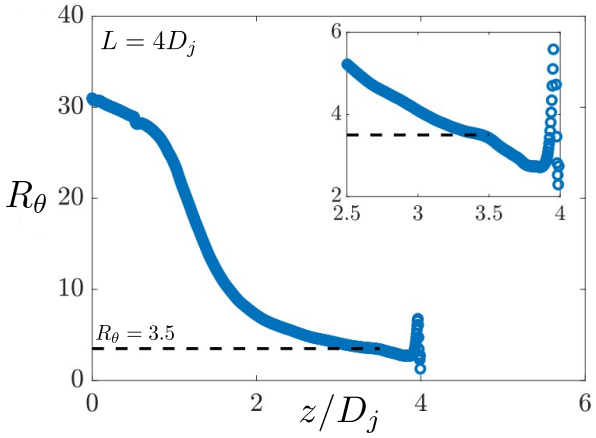


Figure 4: Computation of R_θ for an impinging jet with $L = 4D_j$ [10].

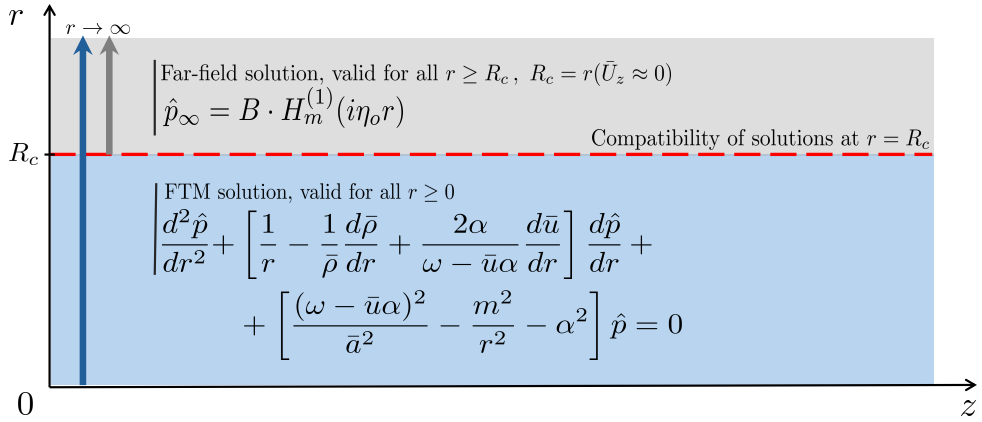


Figure 5: Domain used for the FTM's pressure solution computation.

Ideal Gas Law after computing the temperature. The "thickness parameter" of the FTM is given by R_j/θ_R, following the treatment applied by Maia [10], and will be henceforth referred to as R_θ. The author studied jets simulated with Large Eddy Simulations and found that the profile of Eq. (9) was a good fit for the simulated jet. Maia's key finding was that, to accurately identify the waves propagating inside a real jet with a given nozzle-to-surface length L, the shear-layer thickness in the FTM should match that of the jet's velocity profile near the reflection point at $z = L$. Figure 4 shows the thickness of the shear layer computed by Maia for a supersonic impinging jet from a LES (from Bogey and Gojon [7]). For this jet with $L = 4D_j$, the author chose a shear-layer thickness of $R_\theta = 3.5$ to accurately model the wave propagation with the FTM. In his work, a value of $R_\theta = 5$ is chosen for the same jet but with $L = 3D_j$ and $R_\theta = 2.5$ for both the jets $L = 5D_j$ and $L = 6D_j$. These values of shear-layer thickness are adopted in this work. The jet exit velocity is fixed at $M = 1.5$.

The dispersion pairs (α, ω) and the radial pressure distribution $\hat{p}(r)$ are both obtained from Eq. (8). Since it does not have a known exact solution, an iterative method is built to solve the ODE. The equation is solved as an IVP where only the axial position ($r = 0$) needs a boundary condition. To ensure the solution remains bounded at $r \to \infty$, Eq. (8) is coupled to the far-field analytical pressure solution of the VSM, as shown in Fig. 5. The same kinematic and dynamic boundary conditions are applied at the position of

solution compatibility ($r = R_c$) to ensure solution smoothness. A new quantity is derived from all the conditions, which is given by k:

$$\frac{d\hat{p}}{dr}(R_c) - \left(1 - U_j\frac{\alpha}{\omega}\right)^2 \frac{\hat{p}(R_c)}{H_m^{(1)}(i\eta_o R_c)} \frac{\partial\left(H_m^{(1)}(i\eta_o r)\right)}{\partial r}\bigg|_{r=R_c} = k = 0\,. \tag{10}$$

This "compatibility factor" needs to be zero for both the pressure solutions to be compatible and the field be well solved. In practice this is achieved by imposing a threshold for $k = 0$.

An iterative process is implemented to find the pairs (α, ω) that satisfy the model conditions and governing equations. A guess is given for this pair and the wavenumber (α) is iterated through the Broyden's Method (quasi-Newton method) until function k is below the defined threshold. This is done for a large range of dispersion pairs.

The dispersion results, of both the VSM (black curves) and the FTM (blue curves) are present on Fig. 6. The waves being identified in the figure belong to the axisymmetric azimuthal mode $m = 0$ and were obtained for a jet with $R_\theta = 3.5$. For $m = 0$, and for the selected frequency (Strouhal number, St) and wavenumber ranges, three dispersion branches are obtained, with each one of the them being related to a radial order (shape of $\hat{p}(r)$). The inclined dashed line identifies the region of the dispersion map where waves would travel with a velocity (phase velocity, $v_p = \omega/\alpha$) equal to the sound speed and upstream the jet column (negative group velocity, $v_g = \partial\omega/\partial\alpha < 0$). The dispersion of the VSM, came from Eq. (7) and contains waves travelling both upstream ($v_g < 0$) and downstream ($v_g > 0$) the modelled jet column. The modes of interest ("closing modes") — upstream-propagating wave modes with a phase velocity close to the speed of sound a_0 — are, for each one of the dispersion branches, all the pairs (α, ω) in the figure with negative group velocity (negative slope). For each branch they start at the line $\omega/\alpha = -a_0$ and go up to the top of the branch where $v_g = 0$. For the FTM, these frequencies which allow the existence of the "closing modes", are included in the grey bands. The computation of the limiting frequencies (orange and red points for the FTM) is done through simple numerical processes.

When comparing the two models, the effect of including a mixing layer becomes clear: the frequency range of the closing modes expands, with both limits shifting downward, suggesting that the more realistic model predicts the intense tones over a broader range and at lower average frequencies.

The radial pressure distribution obtained for the closing modes with the FTM confirms one important finding from the VSM: the closing modes cannot propagate only inside of the jet column but they also cannot propagate just in the quiescent medium. Instead, propagate out of the column but are guided by the jet on their way upstream. This result is framed with the literature further down the text.

Finite Thickness Model (FTM) Verification

The verification process of the FTM is conducted by assuming a base flow with a very thin shear-layer in order to establish a direct comparison with the analytical model (VSM). The two models should output the same results since they are built over the same governing equations. Figure 7 shows that the VSM results are well recovered for the

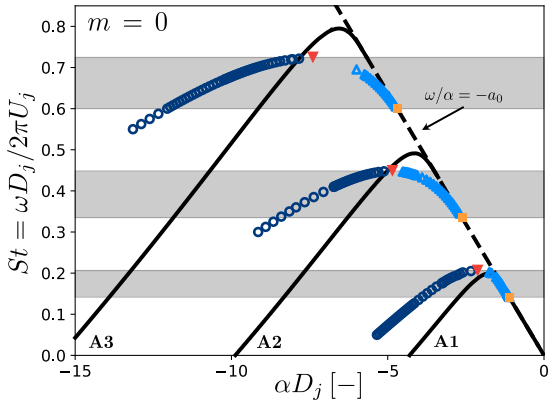


Figure 6: Dispersion for the VSM (black solid lines) and FTM (blue dots). Axisymmetric mode ($m = 0$) and $R_\theta = 3.5$.

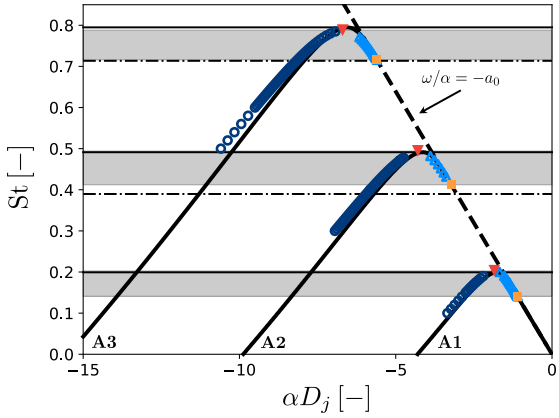


Figure 7: Dispersion of the FTM with a thin shear-layer: comparison against the FTM (black curves). $m = 0$ and $R_\theta = 30$.

FTM with a very thin shear-layer, with only minor discrepancies arising from the shear-layer being finite, though very small, rather than infinitely thin. For the verification, the thickness was chosen to be $R_\theta = 30$, which corresponds to that of a velocity profile right at the edge of the nozzle ($z = 0$ in Fig. 4), where the shear-layer starts to develop. The frequency limits for the existence of closing modes (grey bands) are also well recovered. In Fig. 7, the limits obtained for the FTM are represented by the horizontal black lines: solid lines for upper limits and dashed for lower. Their quantitative comparison is present in Tab. 1, where the relative error between the thin-shear-layer FTM and the VSM is presented (δ [%]).

3 STABILITY ANALYSIS

The second part of the work brings a new understanding to the closure of the AFL by connecting the spatial structure of the modes propagating in the column with their temporal stability. Since the AFL is a frequency-centred phenomenon, accessing the temporal behaviour of each wave mode is crucial to understand if they can sustain the intense acoustic tones.

Table 1: Frequency limits for the existence of closing modes in the FTM. Results for $R_\theta = 30$.

	Symmetric ($m = 0$)			Helical ($m = 1$)	
	A1	A2	A3	H1	H2
St^{max}	0.202	0.490	0.787	0.350	0.646
$\delta(St^{max})$	+1.5%	-0.2%	-1.0%	+1.4%	0.0%
St^{min}	0.141	0.412	0.715	0.339	0.618
$\delta(St^{min})$	-	+5.6%	+0.1%	+6.3%	-1.0%

Table 2: Characteristic dimensions of the computational domain.

$L = 7D_j$	$N_z = 201$	$\Delta z = 7/201D_j$	$R_{max} = 4D_j$	$N_r = 101$	$\Delta r = 4/101D_j$	D_j
0.014 [m]	–	$\approx 6.97E-5$ [m]	0.008 [m]	–	$\approx 7.92E-5$ [m]	0.002 [m]

The same linearized continuity, momentum and energy equations of a compressible fluid are still used, with the difference that the Sutherland's Law is now applied to characterize viscosity. These equations are developed with the decomposition typical of stability analyses, where the base flow is separated from the perturbed flow: $q(r, \theta, z, t) = \bar{Q}(r, z) + q'(r, \theta, z, t)$; and the perturbed flow is taken to be homogeneous along the azimuth θ and time t: $q'(r, \theta, z, t) = \hat{q}(r, z) \exp\{i(m\theta - \omega t)\}$. The governing equations are therefore written in terms of the base flow $\bar{Q}$, the amplitude $\hat{q}$ of the perturbed flow and their azimuthal wavenumber m and angular frequency ω. The system is assembled in matrix form:

$$(A - \omega B)X = 0\,, \tag{11}$$

with A containing all the terms not dependant on ω and B all the dependent terms. The vector of unknowns $X = \hat{\mathbf{q}}(r, z)$ and the frequency ω are regarded as complex quantities, with the imaginary frequency being the growth factor of each wave mode: $\omega = \omega_r + i\,\omega_i$.

The base flow is defined in two dimensions (r and z) by a streamwise extrusion of the profile represented in Eq. (9), meaning that the base velocity field only has a streamwise component: $\bar{\mathbf{U}} = (0, \bar{U}_z)$. The jet is made to be supersonic, with an axial velocity of $M = 1.5$. Temperature and density are computed, for the base flow, following the same process as in the FTM.

The governing equations are applied to a rectangular computational domain with its dimensions present in Tab. 2. L and R_{max} correspond, respectively, to the width and height of the domain. The parameters Δz (element width), N_z (number of nodes), Δr and N_r, on the other hand, identify the Cartesian mesh where the governing equations are discretized. For the discretization, each rectangular element was divided across its diagonal and Taylor-Hood triangular elements (Th) were then used to compute the temperature, density and velocity. The Th element combines a $\mathcal{P}_1$ Lagrange element for temperature and density and a $\mathcal{P}_2$ element for velocity.

To approximate the stability result to that of the mathematical modelling of the FTM (1D jet column), no boundary condition is imposed on the side boundaries (left and right). At the top, a Dirichlet condition is applied (quantity is forced to be zero) and at the axis (only a half-jet is simulated due to the assumption of symmetry) the following conditions

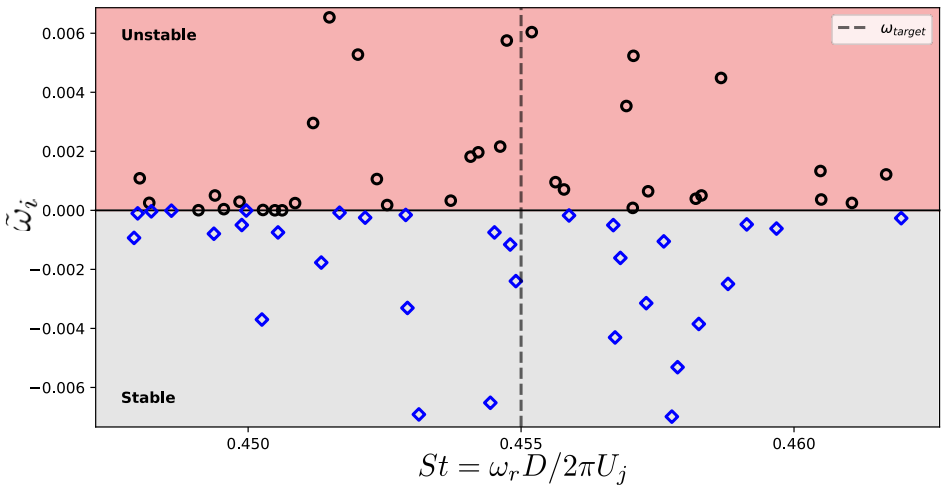


Figure 8: Temporal spectrum obtained from *GIFIE*. The target frequency is represented by the vertical line.

are applied:

$$\frac{\partial \hat{\rho}}{\partial r} = \frac{\partial \hat{u}_z}{\partial r} = \frac{\partial \hat{T}}{\partial r} = \hat{u}_r = \hat{u}_\theta = 0\,. \tag{12}$$

These axial boundary conditions are dependent on the azimuthal order (m) of the solution. Equation (12) reflects the shape of the solution at $r = 0$ for the axisymmetric mode since that is mode being simulated ($m = 0$). As later commented, the condition applied at the top boundary needs to be reviewed since it introduced reflective spurious wave modes in the problem's domain.

After applying the boundary conditions, the discretized system is read with Python and converted into a format compatible with SLEPc (an eigenproblem solver library) where the resolution of the eigenvalue problem is preformed in parallel. Because the discretized matrices are sparse, the eigenvalues are projected onto a lower dimension space (Krylov subspace) and the problem is solved based on a shift-invert transformation where the algorithm looks for eigenvalues close to a target frequency, ω_{target}. All steps in this stability analysis are processed in the Global Stability Finite Element Solver (*GIFIE*), an in-house (ISAE-SUPAERO) solver used in global stability studies. *GIFIE* works through a framework developed in FreeFEM++.

3.1 Primary Results

The temporal spectrum of Fig. 8 is obtained for the target frequency of $St = 0.455$, with 66 eigenvalues being computed for the pre-defined Krylov subspace size. The growth factor is normalized as $\tilde{\omega}_i = \omega_i D_j/2\pi U_j$, with 32 of the eigenvalues being stable ($\tilde{\omega}_i < 0$) and 34 are unstable. The system is globally unstable since at least one eigenvalue is unstable. Global instability is not commonly observed in the literature [8] for the type of jet column simulated, suggesting that while some unstable modes could be physical, most—if not all—likely result from interactions between physical and spurious modes introduced by the numerical setup.

The pressure eigenfunction of the most neutrally stable eigenvalue (#34, ranked from

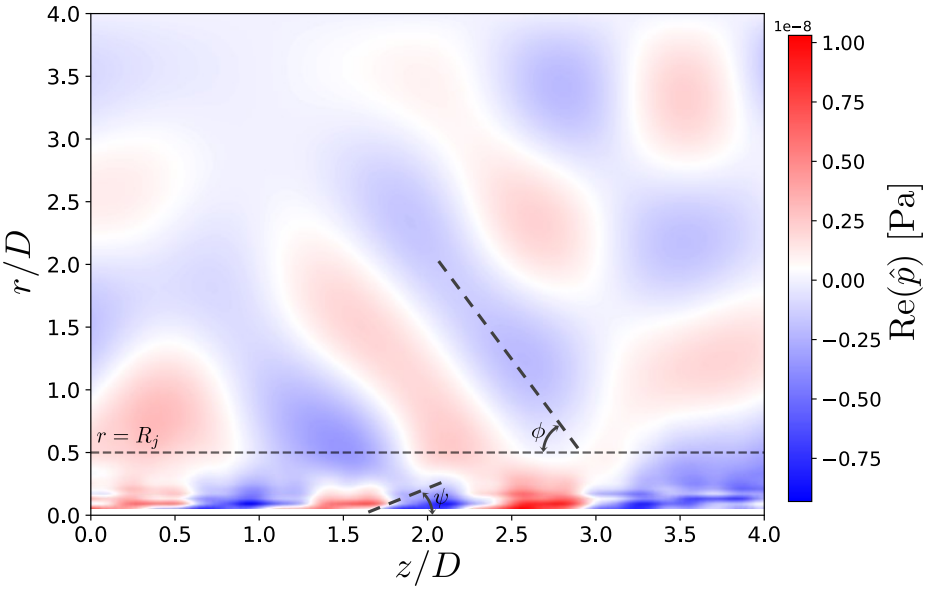


Figure 9: Pressure eigenfunction obtained for eigenvalue #34.

most to least unstable) is obtained, from the temperature and density computed fields, using the Ideal Gas Law. This wave mode present in Fig. 9, is believed to be able to sustain an intense tone in time, given its temporal neutrality. Two distinct structures can be identified: those concentrated in the axial and shear regions of the jet column (associated with angles ϕ and ψ) and those forming a "chequered" pattern across the rest of the computational domain. The former align with the analytical jet model [5], experimental observations [2], and high-fidelity numerical simulations [7]. The latter are likely artifacts caused by spurious mode reflections at the rigid Dirichlet top boundary, which prevents full wave dissipation and feeds disturbances back into the flow.

3.2 Post-Processing

For the waves of interest to be identified in the eigenfunction $\hat{p}(r, z)$, a post-processing tool was developed. It consists on frequency-decomposing the pressure field through the application of a Fast Fourier Transform (FFT), making it possible for waves to be detected from their propagation direction (upstream/downstream) and velocity (phase velocity given by $v_p = \omega/\alpha$); the two quantities needed to define closing modes in the AFL.

First, the results are split along the radial direction to obtain N_r one-dimensional complex pressure signals, $\hat{p}(z)$. Since the FFT requires periodicity, each signal undergoes windowing to minimize spectral leakage. To enhance frequency resolution for later integral computations, the windowed signals are then zero-padded before the FFT is applied. Applying the FFT to the zero-padded signals yields the spectral decomposition $\mathcal{F}(\alpha_b, r)$, representing the amplitude distribution of each frequency bin α_b across the radial direction. The power spectrum $\mathcal{S} = |\mathcal{F}(\alpha_b, r)|^2$ is then computed, allowing the identification of dominant wave modes propagating within the jet column. The normalized power spectrum, $\tilde{\mathcal{S}}(\alpha_b, r) = \mathcal{S}(\alpha_b, r)/\max_{\alpha_b, r} \mathcal{S}(\alpha_b, r)$, is shown in Fig. 10, where the whitest regions

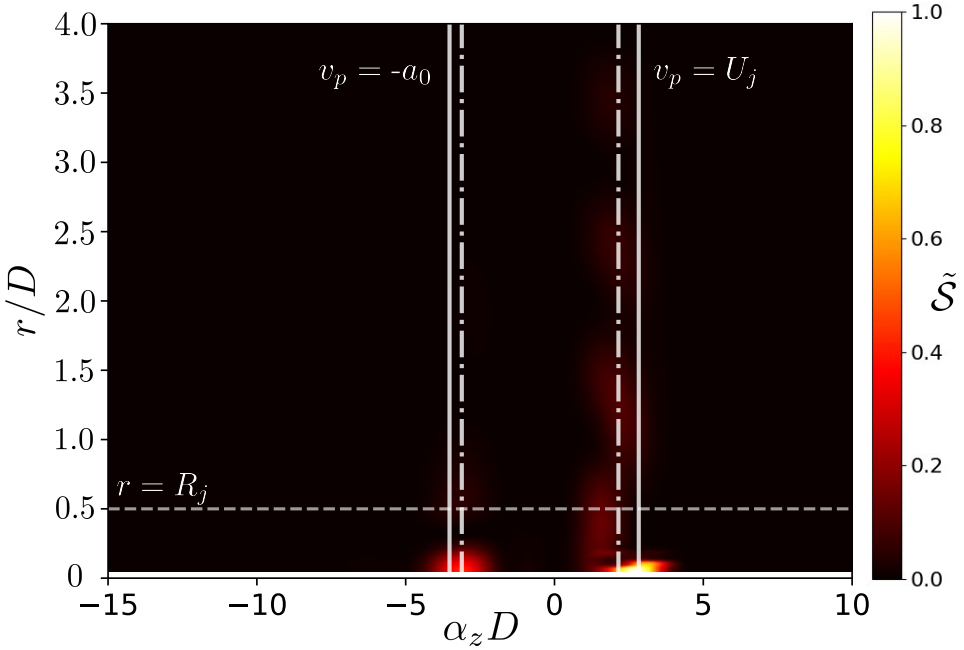


Figure 10: Power spectrum obtained for eigenvalue #34.

indicate the highest spectral power. Solid vertical lines mark the expected locations of downstream-propagating hydrodynamic modes ($v_p = U_j$) and upstream-propagating acoustic modes ($v_p = a_0$), while dashed lines represent their average detected frequencies (weighted average). Ideally, both sets should align, but spurious interactions have introduced unphysical wave propagation, with acoustic waves exceeding a_0 and hydrodynamic waves surpassing U_j. These speeds follow $v_p = \omega/\alpha$, where lower α at fixed ω results in higher v_p.

Since this work is mostly focused on analysing the closure of the AFL through the upstream-propagation of acoustic waves, a radial pressure eigenfunction is obtained from the frequency bin closest to $\alpha = -a_0/\omega$, such that $\hat{p}(r) = \tilde{\mathcal{S}}\left(\alpha_b \approx \frac{\omega}{-a_0},\ r\right)$. For each eigenvalue in Fig. 8, the presence of the closing mode is quantified using:

$$\mathrm{P}_{cm} = \int_0^{R_{max}} |\hat{p}(r)|\ dr\,. \tag{13}$$

This metric is then mapped back onto the temporal spectrum, highlighting the eigenvalues with upstream-propagating acoustic waves capable of closing the AFL, as shown in Fig. 11.

The distribution of this quantity in the temporal spectrum suggests that most unstable modes lack upstream-propagating acoustic waves, linking them instead to either hydrodynamic modes or spurious numerical artifacts. Since hydrodynamic modes typically form organized branches in the stable region [8], the scattered distribution of unstable eigenvalues in Fig. 11 supports the hypothesis that they primarily stem from spurious modes.

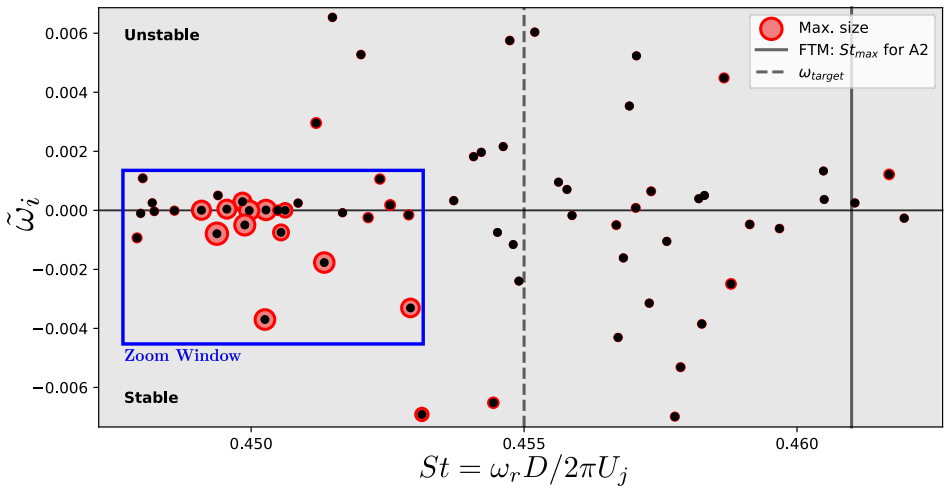


Figure 11: Acoustic presence of each eigenvalue projected onto the original spectrum. "Max. size" corresponds to the highest recorded P_{cm} value (eigenvalue #50).

4 FRAMING THE RESULTS

4.1 Pressure Eigenfunction

The pressure eigenfunctions $\hat{p}(r)$ obtained from different methodologies are compared to assess the propagation of the closing mode in the supersonic jet. They are presented in Fig. 12. Five approaches are considered: the vortex-sheet model (VSM) [5], two finite-thickness models (FTM) solved through different numerical methods (those presented here and Maia's "LoSA" [10]), the bi-global stability analysis (*GIFIE*), and a Large Eddy Simulation (LES) [7]. Results show strong agreement between the FTM and LES, with minor discrepancies due to the complexity of extracting pressure fields from LES data [10]. The *GIFIE* eigenfunction, taken for the eigenvalue with highest P_{cm}, while capturing the wave mode's shape, exhibits lower intensity, most likely due to numerical damping from spurious reflections at the domain top boundary. Despite the discrepancies, all methods confirm that closing modes propagate upstream via the jet's guiding effect rather than solely within the jet core or quiescent medium. This consistency reinforces the validity of the models and strengthens the base foundation of the Aeroacoustic Feedback Loop (AFL), the VSM.

4.2 Frequency Range for closing modes

The recorded impingement tones from experimental [2, 11–15] and numerical [7] studies are compared against the closing modes frequency limits predicted with the VSM [5] and FTM. Results for $m = 0$ and $R_\theta = 3.5$ in Fig. 13 show that finite-thickness effects significantly impact frequency predictions, with both upper and lower limits shifting downward compared to the VSM. The FTM (blue shaded regions) and Maia's "LoSA" [10] model (blue dashed lines, identified with black arrows) exhibit strong agreement, though a discrepancy appears in the lower frequency limit, where the methoded developed in this article captures more experimental points (only in helical $m = 1$ modes; not shown here due to page requirements). Most experimental tones fall within the finite-thickness predictions, confirming the necessity of including a shear layer for accurate frequency estimation.

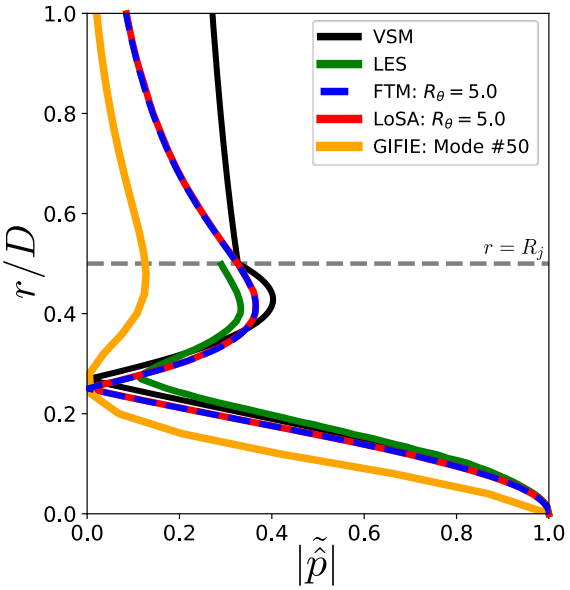


Figure 12: Pressure eigenfunction for the closing wave mode obtained through different methods.

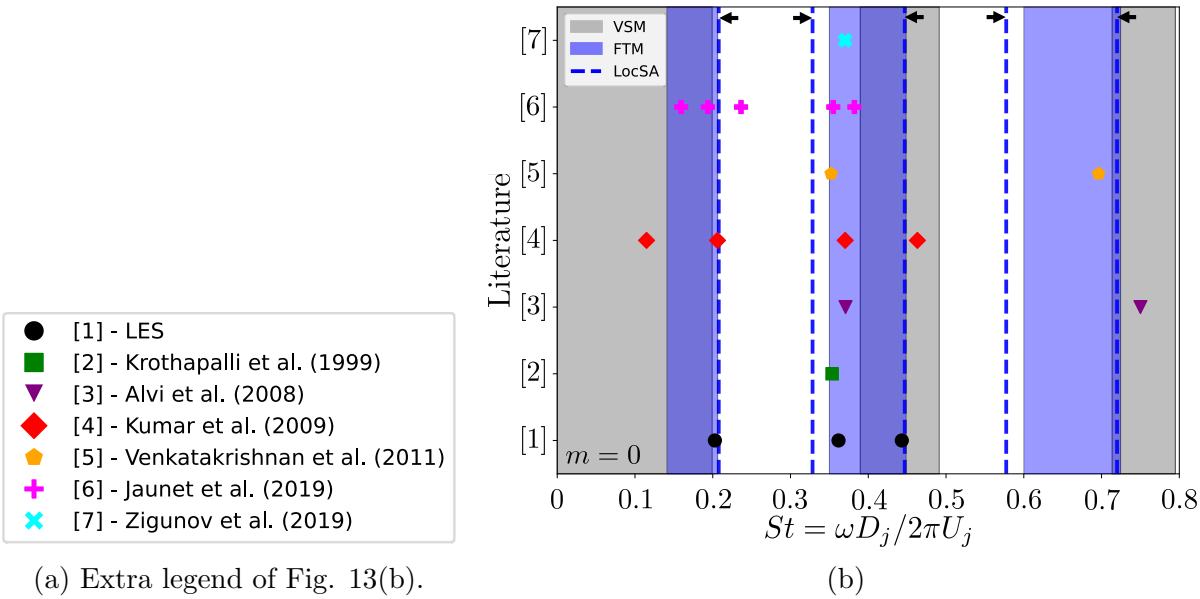
Some deviations are attributed to nonlinear modal interactions, previously identified by Schmidt [9], which should be further examined using advanced spectral decomposition techniques.

5 CONCLUSIONS

This work studied the Aeroacoustic Feedback Loop (AFL) in supersonic impinging jets, focusing on the upstream propagation of acoustic modes responsible for closing the loop. Analytical, and numerical perspectives were considered through the Finite-Thickness Model (FTM) and the bi-global stability analysis (*GIFIE*) implemented. Results confirmed that closing modes are guided upstream by the jet shear-layer while also propagating along the medium at rest. The method used to solve the FTM was successful, as it not only reproduced results consistent with existing literature but was also verified against the analytical vortex-sheet model, confirming its accuracy. Regarding the stability simulation, interesting results were obtained, but the numerical setup requires further refinement. The main focus of this second method was the development of a post-processing tool, which successfully identified all the propagating modes in the supersonic jet column predicted by the first method (mathematical modelling of the jet column). This tool provides a powerful framework for analysing the interplay between temporal and spatial dynamics, making it a valuable asset for design optimization in aerospace systems, where precise characterization of flow instabilities is essential for performance-driven design.

ACKNOWLEDGEMENTS

The DAEP department at ISAE-SUPAERO is gratefully acknowledged for providing the numerical tools, including the *GIFIE* solver, and the computational resources that were instrumental to the successful completion of this work.



(a) Extra legend of Fig. 13(b). (b)

Figure 13: Predicted frequency range for the appearance of impingement tones overlaid with experimental and numerical data points.

REFERENCES

[1] M. J. Lighthill. On sound generated aerodynamically. i. general theory. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, (1107), 1952.

[2] A. Krothapalli, E. Rajkuperan, F. Alvi, and L. Lourenco. Flow field and noise characteristics of a supersonic impinging jet. *Journal of Fluid Mechanics*, 1998.

[3] A. Powell. On edge tones and associated phenomena. *Acustica*, 3:233–243, 1953.

[4] D. Edgington-Mitchell. Aeroacoustic resonance and self-excitation in screeching and impinging supersonic jets – a review. *International Journal of Aeroacoustics*, 18, 2019. doi:10.1177/1475472X19834521. URL `https://doi.org/10.1177/1475472X19834521`.

[5] C. K. W. Tam and F. Q. Hu. On the three families of instability waves of high-speed jets. *Jornal of Fluid Mechanics*, pages 447–483, 1989.

[6] J. Lepicovsky and K. K. Ahuja. Experimental results on edge-tone oscillations in high-speed subsonic jets. *AIAA Journal*, 23:1463–1468, 1985. doi:10.2514/3.9111.

[7] C. Bogey and R. Gojon. Feedback loop and upwind-propagating waves in ideally expanded supersonic impinging round jets. *Journal of Fluid Mechanics*, 823:562–591, 7 2017. ISSN 14697645. doi:10.1017/jfm.2017.334.

[8] J. W. Nichols, S. K. Lele, and D. P. Moin. Global mode decomposition of supersonic jet noise. 2009.

[9] O. T. Schmidt, A. Towne, T. Colonius, A. V. Cavalieri, P. Jordan, and G. A. Brès. Wavepackets and trapped acoustic modes in a turbulent jet: Coherent structure

eduction and global stability. *Journal of Fluid Mechanics*, 825:1153–1181, 8 2017. ISSN 14697645. doi:10.1017/jfm.2017.407.

[10] I. A. Maia, M. Fiore, and R. Gojon. Tones and upstream-traveling waves in ideally expanded round impinging jets. *Physical Review Fluids*, 9:083904, Aug 2024. doi:10.1103/PhysRevFluids.9.083904. URL `https://link.aps.org/doi/10.1103/PhysRevFluids.9.083904`.

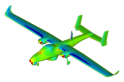
[11] R. Kumar, S. Lazic, and F. S. Alvi. Control of high-temperature supersonic impinging jets using microjets. *AIAA Journal*, 47(12):2800–2811, 2009. doi:10.2514/1.39061.

[12] F. S. Alvi, H. Lou, C. Shih, and R. Kumar. Experimental study of physical mechanisms in the control of supersonic impinging jets using microjets. *Journal of Fluid Mechanics*, 613:55–83, 2008. doi:10.1017/S0022112008003091.

[13] L. Venkatakrishnan, A. Wiley, R. Kumar, and F. Alvi. Density field measurements of a supersonic impinging jet with microjet control. *AIAA Journal*, 49(2):432–438, 2011. doi:10.2514/1.J050511. URL `https://doi.org/10.2514/1.J050511`.

[14] V. Jaunet, M. Mancinelli, P. Jordan, A. Towne, D. M. Edgington-Mitchell, G. Lehnasch, and S. Girard. Dynamics of round jet impingement. In *25th AIAA/CEAS Aeroacoustics Conference, 2019*. American Institute of Aeronautics and Astronautics Inc, AIAA, 2019. ISBN 9781624105883. doi:10.2514/6.2019-2769.

[15] F. Zigunov, P. Sellappan, and F. Alvi. Instability modes of millimeter-scale supersonic jets. *Experiments in Fluids*, 60, 4 2019. ISSN 07234864. doi:10.1007/s00348-019-2707-4.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

INVESTIGATING THE INFLUENCE OF A TIP PROPELLER'S VERTICAL OFFSET ON A TRAILING WING'S LIFT AND DRAG DISTRIBUTIONS

Gustavo Padovany da Silva[1], Shamsheer S. Chauhan[2], Gustavo L. O. Halila[1*], Joaquim R. R. A. Martins[2], and João Luiz F. Azevedo [3]

1: Mechanical Engineering Graduate Program (PPGEM)
Pontifícia Universidade Católica do Paraná (PUCPR)
Curitiba, PR, Brazil
guspadovany@gmail.com
*gustavo.halila@pucpr.br

2: Department of Aerospace Engineering
University of Michigan
Ann Arbor, MI, USA
{sschau,jrram}@umich.edu

3: Instituto de Aeronáutica e Espaço
São José dos Campos, SP, Brazil
joaoluiz.azevedo@gmail.com

Abstract. *Wingtip-mounted tractor propeller configurations are receiving attention across the aerospace industry because of their aerodynamic benefits. However, results displaying spanwise lift and drag variations due to changes in the propeller-wing vertical offset for such configurations are scarce in the literature. In this paper, we investigate the changes in a rectangular wing's spanwise lift and drag coefficient distributions by varying the propeller-wing vertical offset of a wingtip propeller. For this, we perform steady Reynolds-averaged Navier–Stokes computational fluid dynamics simulations using the open-source flow solver ADflow with the Spalart–Allmaras turbulence model and an actuator-disk approach. We present results for two angles of attack with a range of propeller-wing vertical offsets for both inboard-up and outboard-up propeller rotation. We include plots for spanwise lift and drag coefficient distributions, integrated lift and drag coefficients, and lift-to-drag ratios. Our results show that changing the propeller-wing vertical offset can significantly change a trailing wing's lift and drag distributions, which strongly depend on the direction of the vertical offset. Additionally, this work paves the way for further studies that benefit from ADflow's efficient gradient-computation capabilities.*

Keywords: Propeller-wing interference, propeller-wing integration, distributed propulsion

NOMENCLATURE

b	=	Wing span
C_D	=	Integrated drag coefficient of the wing
C_d	=	Sectional drag coefficient
C_L	=	Integrated lift coefficient of the wing
C_l	=	Sectional lift coefficient
D	=	Wing drag
$D_\varnothing$	=	Propeller diameter
F	=	Axial-force distribution parameter
f_x	=	Axial force per unit radius
f_θ	=	Tangential force per unit radius
J	=	Propeller advance ratio, $V_\infty/(\nu D_\varnothing)$
L	=	Wing lift
n	=	Propeller-model shape parameter
P	=	Propeller pitch
R	=	Propeller radius
r	=	Radial coordinate
$\hat{r}$	=	Normalized radius, $(r - r_{\text{in}})/(R - r_{\text{in}})$
r_{in}	=	Inner radius
T	=	Thrust
T_C	=	Thrust coefficient, $T/(\rho V_\infty^2 D_\varnothing^2)$
V_∞	=	Freestream velocity
μ	=	Dynamic viscosity
ν	=	Rotation rate (rotations per unit time)
ρ	=	Fluid density

1 INTRODUCTION

Propeller-wing interaction has been a topic of investigation since the dawn of aviation [1–4]. Recent interest in developing electric aircraft [5–9] has led to renewed interest in modeling propeller-wing interaction using computational tools [10–13]. When a propeller is located near a wing, the flow induced by the propeller affects both the lift and the drag of the wing. The nature of the influence depends on the propeller's location and rotation direction [1, 14–17]. When located directly in front of the wing, the propeller's axial and tangential induced-velocity components create spanwise variations in the wing's local angles of attack and effective freestream speeds. Additionally, when a propeller is vertically offset relative to a wing, the pressure and velocity distributions associated with the streamtube contraction produce upwash or downwash effects, even when the slipstream does not impinge directly on the wing [1, 17, 18]. In this paper, we investigate how these effects combine to produce visibly different lift and drag coefficient distributions when the propeller-wing vertical offset is varied for a wingtip propeller configuration. There is a lack of published studies that show computational or experimental results for this.

Prandtl [1] and Veldhuis [17] investigated the effect of propeller-wing vertical offset on a trailing wing's lift and drag using experimental methods. They presented integrated lift and drag coefficient (C_L and C_D) results for a range of positive and negative vertical offsets, and they attributed the non-monotonic trends to two primary factors. One of

these two factors is the pressure and velocity distributions associated with streamtube contraction. This factor is responsible for their observations that when the propeller is located close to the wing but not close enough for the wing to be inside the streamtube, the lift is greater for a positive vertical offset (above the wing) than for a negative vertical offset. The other factor is a combination of the non-uniform radial distribution of dynamic pressure in a propeller's streamtube and the varying extent of the wing inside the propeller's streamtube. However, their results are limited to integrated lift and drag coefficients and do not include results showing the effect of propeller-wing vertical offset on spanwise lift and drag distributions.

In a recent paper [18], we investigated the changes in a rectangular wing's spanwise lift and drag coefficient distributions by varying the propeller-wing vertical offset of an inboard propeller configuration. For that, we performed steady Reynolds-averaged Navier–Stokes (RANS) computational fluid dynamics (CFD) simulations using the open-source flow solver ADflow [19, 20] with the Spalart–Allmaras (SA) [21] turbulence model and an actuator-disk propeller model [13, 22]. We presented spanwise lift and drag coefficient distributions, integrated lift and drag coefficients, lift-to-drag ratios, and chordwise sectional pressure-coefficient distributions. These results show how the two aforementioned factors combine with an additional factor to produce visibly different lift and drag coefficient distributions when the propeller-wing vertical offset is varied. The additional factor is a combination of the varying magnitude of the upwash and downwash components of the propeller swirl velocities [23] and the varying extent of the wing inside the propeller's streamtube. We name this factor *Factor 1* (swirl related), and we name the two aforementioned factors *Factor 2* (streamtube contraction related) and *Factor 3* (dynamic pressure related). For the present work, we use the same methods, software tools, and baseline geometry as that paper but with the propeller located at the wingtip instead of inboard.

There are many approaches of varying complexity and fidelity that can be used to model propeller-wing interaction. These include vortex-lattice approaches [17, 24–28] and CFD approaches [10, 13, 17, 29–34]. The vortex-lattice approaches have limitations related to viscous effects, thickness effects, and streamtube modification. The CFD approaches can capture those effects more directly but have significantly higher computational costs. The approach of using steady CFD with an actuator disk has been shown to provide accurate time-averaged predictions for a wing located behind a propeller [17, 22, 30, 32, 33]. This approach provides a tractable balance of fidelity and cost that is suitable for studies requiring large numbers of simulations. Therefore, we use an actuator-disk CFD approach for this work.

In this paper, we investigate the changes in a rectangular wing's spanwise lift and drag coefficient distributions by varying the propeller-wing vertical offset of a wingtip propeller configuration. We present results for two angles of attack and a range of propeller-wing vertical offsets for both inboard-up and outboard-up propeller rotation. We include plots for spanwise sectional lift and drag coefficient distributions, integrated lift and drag coefficients, and lift-to-drag ratios.

2 COMPUTATIONAL TOOLS

We use open-source modules of the MDO of Aircraft Configurations with High Fidelity (MACH) framework[1] [35] for this work. The tools and methodology described here are

[1] `www.github.com/mdolab/MACH-Aero`

the same as those used for our recent paper investigating the influence of an inboard propeller's vertical offset on a trailing wing's lift and drag distributions [18].

2.1 Flow Solver

For this study, we use the open-source CFD solver ADflow [19, 20] with modifications for actuator-disk propeller-modeling capabilities [13, 22]. ADflow has options to solve the Euler, laminar Navier–Stokes, and RANS equations. It is a second-order-accurate finite-volume CFD solver that solves the flow equations on structured multiblock meshes, and it also supports overset meshes by using implicit hole cutting [36–38].

For this study, we solve the steady RANS equations with the SA-noft2 variant of the SA turbulence model [21]. This variant is described in the NASA Turbulence Modeling Resource[2], and the production term is computed using strain. We use the Jameson–Schmidt–Turkel scheme with scalar dissipation [39] for inviscid fluxes, and the viscous-flux gradients are calculated using the Green–Gauss approach. A pseudo-transient continuation strategy is used in ADflow to converge the flow equations, and multiple algorithms are available. For this work, we use the approximate Newton–Krylov solver [19] and converge the total residual norm by eight orders of magnitude. The simulations for this work were performed using the high-performance computing resources of the Center for Mathematical Sciences Applied to Industry (CeMEAI).

2.2 Mesh Generation

For this study, we use MACH's pySpline[3] and pyGeo[4] modules to generate the wing surface geometry. We generate the wing surface mesh and the propeller-region volume mesh using the commercial package ICEM-CFD. To generate the wing volume mesh and the background volume mesh, we use MACH's pyHyp[5] [40] module. The pyHyp module uses a hyperbolic mesh marching method [41] to produce high-quality volume meshes by extruding structured surface meshes. The primary user-defined parameters are the height of the first layer of cells, the extrusion distance, and the number of layers. These volume meshes are structured multiblock meshes, and we combine them using the open-source tool cgnsUtilities[6] to obtain an overset mesh. We also use cgnsUtilities to change the location of the propeller-region volume mesh for the different propeller locations studied in this work. This overset approach allows for the effortless generation of new meshes for the different propeller locations.

2.3 Propeller Model

We model the propeller using an actuator-disk approach in which terms corresponding to propeller forces are added to the momentum and energy equations in the manner of body forces [22, 29]. We use the following model [42] for the distributions of the axial and tangential loading. The radial distribution of the axial force is given by

$$f_x = F\hat{r}\,(1-\hat{r})^n\,, \tag{1}$$

[2] https://turbmodels.larc.nasa.gov/spalart.html
[3] https://github.com/mdolab/pyspline
[4] https://github.com/mdolab/pygeo
[5] https://github.com/mdolab/pyhyp
[6] https://github.com/mdolab/cgnsutilities

where f_x is the axial force per unit radius, $\hat{r}$ is defined by Eq. (2), n is a parameter that controls the shape of the distribution, and F is a value computed by integrating this expression over the radius and equating it to the total thrust T. The normalized radius $\hat{r}$ is defined as

$$\hat{r} = \frac{r - r_{\text{in}}}{R - r_{\text{in}}}, \tag{2}$$

where r is the radial distance from the axis of rotation, r_{in} is the inner radius of the propeller, and R is the outer radius of the propeller. The radial distribution of the tangential force is given by

$$f_\theta = f_x \left(\frac{P/D_\varnothing}{\pi(r/R)} \right), \tag{3}$$

where f_θ is the tangential force per unit radius and $P/D_\varnothing$ is the propeller pitch-to-diameter ratio.

3 PROBLEM DESCRIPTION

The methodology, baseline geometry, baseline meshes, and operating conditions described here are the same as those used for our recent paper investigating the influence of an inboard propeller's vertical offset on a trailing wing's lift and drag distributions [18]. The primary difference here is the propeller location.

3.1 Geometry and Specifications

The baseline configuration is based on the PROWIM (Propeller-Wing Interference Model) configuration experimentally studied by Veldhuis [17], which is depicted in Fig. 1. It consists of a rectangular half-wing with a tractor propeller. The wing is not swept, twisted, or tapered, and the propeller axis is aligned with the wing section's chord line and has no vertical offset relative to it. The wing airfoil is a NACA 64_2-A015.

In this paper, we use a modified version of the PROWIM configuration by locating the propeller at the wingtip (at $b/2$ where b is the full-wing span) instead of at the original inboard location (at 47% $b/2$). We fit splines to airfoil coordinates from the online UIUC database [43] for the wing sections. We do not model the nacelle, nor do we round the leading edge of the wingtip.

The Reynolds number and freestream dynamic pressure reported for the experimental cases are 0.8×10^6 and 1500 Pa, respectively [17]. For the fluid properties, we use standard sea-level values. The propeller operates at an advance ratio of $J = V_\infty/(\nu D_\varnothing) = 0.85$ and a thrust coefficient of $T_C = T/(\rho V_\infty^2 D_\varnothing^2) = 0.168$ [17]. Here, V_∞ is the freestream velocity, ν is the rotation rate, $D_\varnothing$ is the propeller diameter, ρ is the freestream density, and T is the propeller thrust.

We previously presented validation and mesh refinement results in an earlier paper [22] for the inboard propeller configuration using the same baseline component meshes and propeller model used here. For the validation study [22], we compared our results with experimental results [17] for 0 and 4 deg angle-of-attack (AoA) cases. For this work, we use the same two angles of attack. We change the angle of attack in these simulations by changing the direction of the freestream velocity. This maintains the relative orientations of the actuator disk and the wing.

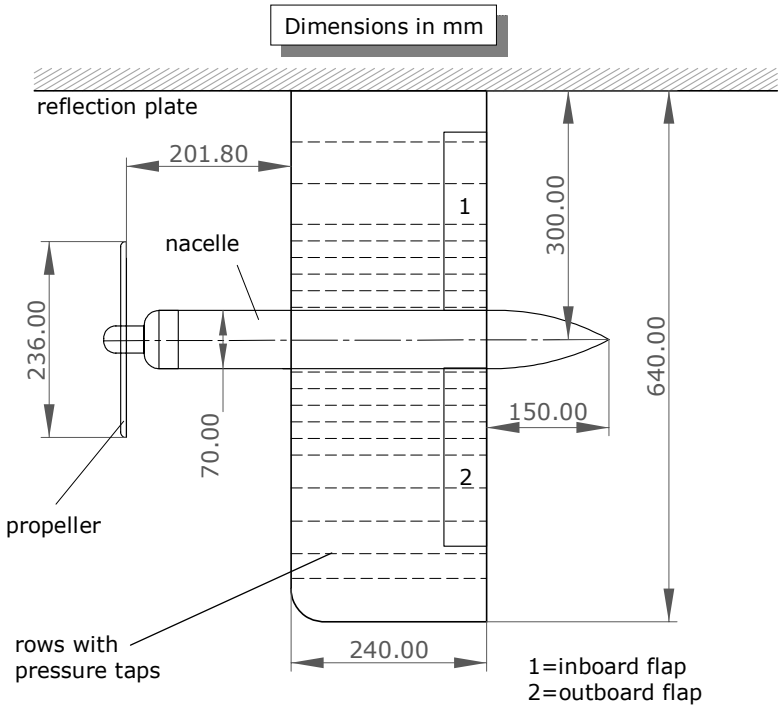
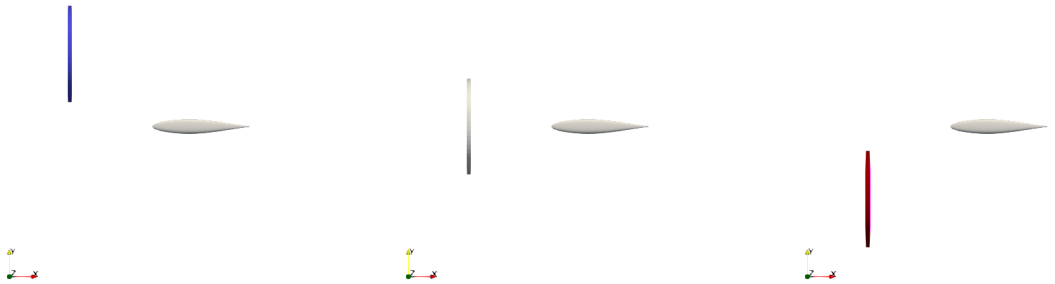


Figure 1: Original geometry of the baseline configuration (half-wing planform area is 0.15 m^2; aspect ratio is 5.3; reproduced with permission from Veldhuis [17])

3.2 CFD Volume Meshes

The overset volume meshes used for this work are composed of a wing volume mesh, a cylindrical volume mesh for the propeller region, and a background volume mesh. Force terms are applied to the cells in the cylindrical volume mesh that lie inside a user-specified disk surface. Figure 2 shows side views of the disk surfaces used to specify the actuator zone in the baseline and two offset positions. The wing surface mesh is generated using the commercial package ICEM-CFD. The wing volume mesh is generated by extruding the wing surface mesh using MACH's pyHyp [40] module.



(a) Vertical offset of $+1.5R$ (b) Baseline propeller location (vertical offset = 0) (c) Vertical offset of $-1.5R$

Figure 2: Side views of the wing and disk surfaces for the baseline and extreme propeller locations

Each background mesh consists of a Cartesian volume mesh encapsulating the wing and propeller-region meshes. A top view is shown in Fig. 3. The outer surfaces of these Cartesian meshes are extruded by 45 chord lengths to yield hemispherical domains. We use far-field boundary conditions for the outer hemispherical surface of the domain and symmetry boundary conditions for the plane coincident with the root of the wing. In order to allow oversetting with the background mesh, the propeller-region volume mesh is larger than the disk of cells to which the actuator forces are applied. ADflow uses an implicit hole-cutting method for connecting the overset meshes [36–38].

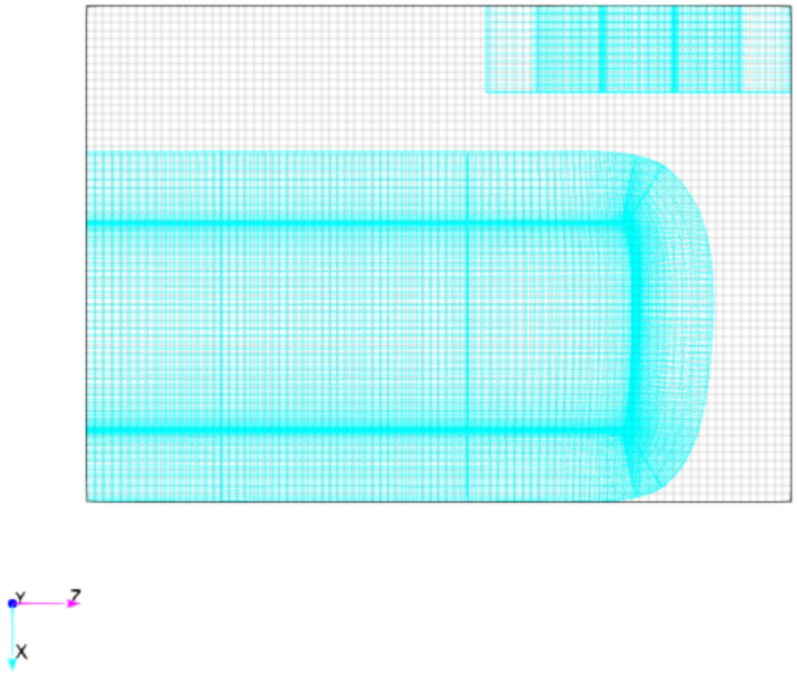


Figure 3: Top view of the outlines of the blocks of the wing volume mesh, propeller-region volume mesh, and Cartesian portion of the background volume mesh; hemispherical mesh extending from the Cartesian mesh to the farfield not shown

To change the vertical offset of the propeller for this study, we reposition the propeller-region volume mesh and regenerate the background mesh. The midplanes of the cells that are selected for the actuator zones are located 0.2 m in front of the wing. The radius of the actuator zone is the propeller radius of 0.12 m, and its thickness is 9 mm (one layer of cells).

We generate the wing volume mesh by extruding the wing surface mesh by 0.4 chord lengths using pyHyp. The wing volume mesh has 40 layers of cells, and the maximum growth ratios of these layers in the off-wall directions range from 1.2 to 1.3. The cell thickness of the first off-wall layer is set to a flat-plate $y+$ value of 1. We previously presented validation and mesh refinement results for the original inboard propeller configuration in an earlier paper [22], and this work re-uses the Level-2 wing volume and propeller-region meshes from that work.

3.3 Propeller Model Inputs

For the original experimental configuration propeller [17], $P/D_{\varnothing}$ ranges from 0.9 at $r/R = 0.45$ to 1.1 at $r/R = 0.95$ [44]. For the propeller model, we use $P/D_{\varnothing} = 1.05$, which corresponds to the $0.75R$ location. The tangential-loading model,given in Eq. (3), is based on the assumptions that the propeller has a radially constant pitch and that the forces on the blade are locally perpendicular to the blade chords, which are approximations.

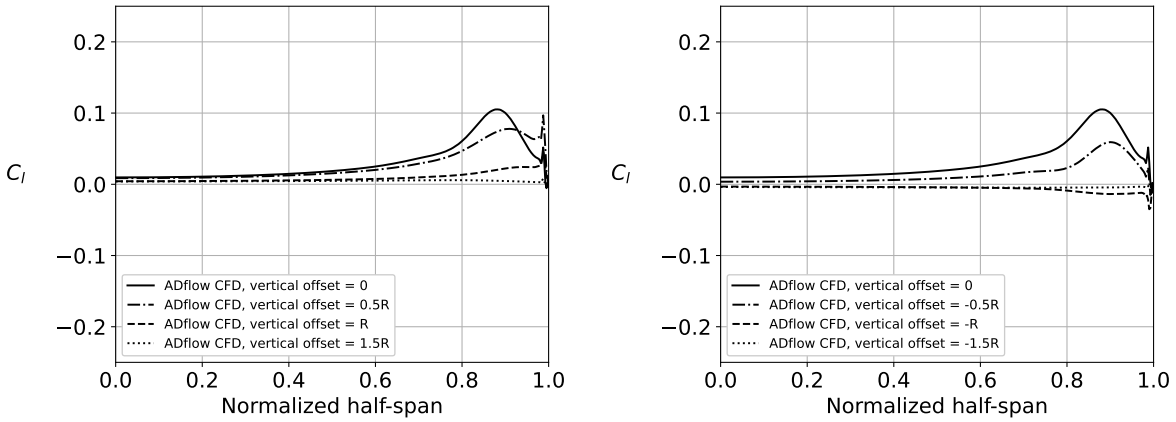
CFD simulations by Stokkermans [33] for the original propeller geometry show that the blade root, that is, the region from $r = 0.35R$ to $r = 0.15R$, generates negative thrust due to the thick profiles and flow separation. Therefore, we use $r_{\text{in}} = 0.35R$ in the axial-loading model (Eq. (1)), which results in negative values between $r = 0.35R$ and $r = 0.15R$. Additionally, we multiply the axial forces from Eq. (1) by 0.25 for this root portion, as the results of Stokkermans [33] show a reduction in slope after the axial force becomes zero at approximately $r = 0.35R$. We set the axial forces at the spinner portion, *i.e.*, for $r < 0.15R$, to zero. We use $n = 0.2$ for the adjustable parameter n, which gives a thrust distribution with the peak near the tip.

For the tangential force distribution, we use Eq. (3) between $r = R$ and $r = 0.35R$ and set the forces to zero for $r < 0.35R$. The F parameter is set to a calculated value that makes the total axial force applied to the region between r_{in} and R equal to the specified thrust.

4 RESULTS

4.1 Angle of Attack = 0 deg

Figures 4 and 5 show the wing sectional lift coefficient, C_l, and sectional drag coefficient, C_d, distributions, respectively, for the 0 deg AoA cases with inboard-up propeller rotation. Figures 6 and 7 show the same for outboard-up propeller rotation. For a rectangular wing, the shapes of the C_l and C_d distributions are also the shapes of the lift and drag distributions, respectively.

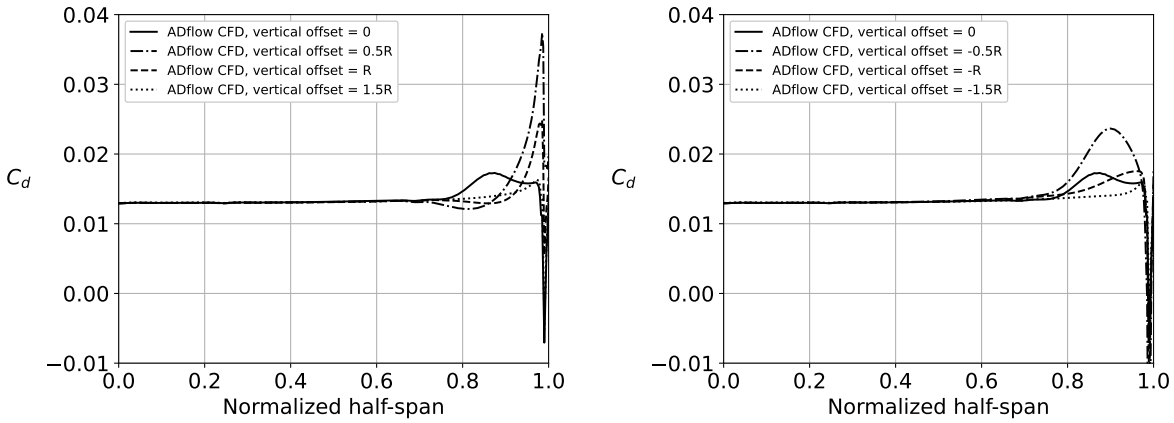


(a) Positive vertical offset

(b) Negative vertical offset

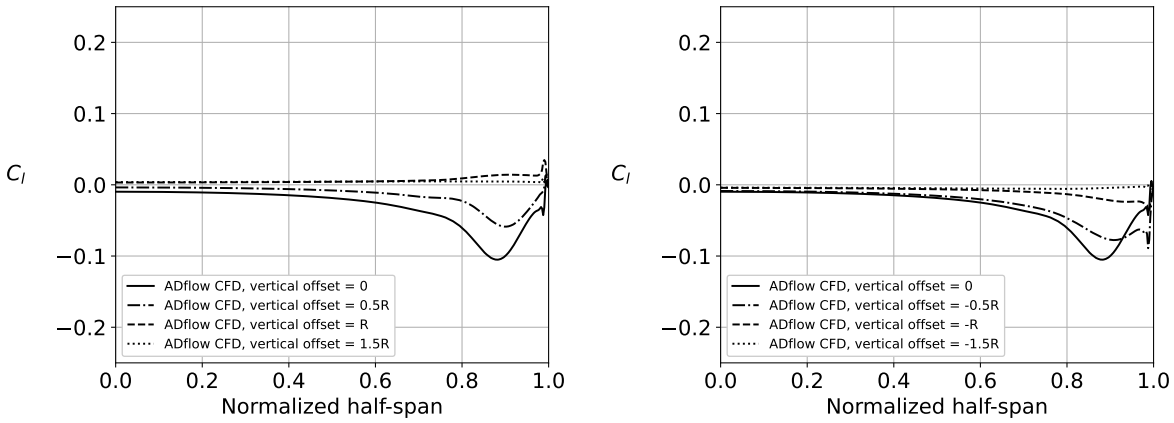
Figure 4: Wing C_l distributions for 0 deg AoA with inboard-up propeller rotation

For the inboard-up cases, when the propeller-wing vertical offset is zero, there is a peak in the C_l distributions for the portion corresponding to the upgoing-blade region.



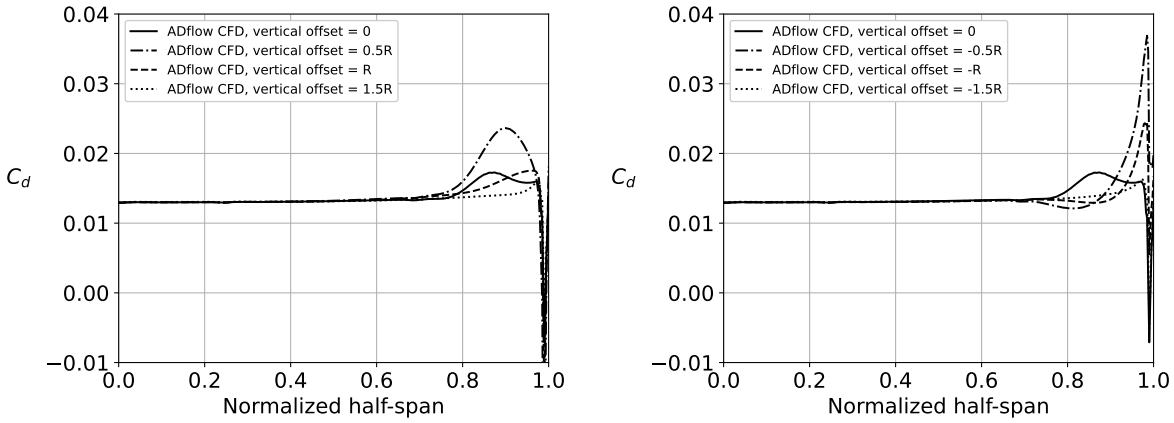
(a) Positive vertical offset (b) Negative vertical offset

Figure 5: Wing C_d distributions for 0 deg AoA with inboard-up propeller rotation



(a) Positive vertical offset (b) Negative vertical offset

Figure 6: Wing C_l distributions for 0 deg AoA with outboard-up propeller rotation



(a) Positive vertical offset (b) Negative vertical offset

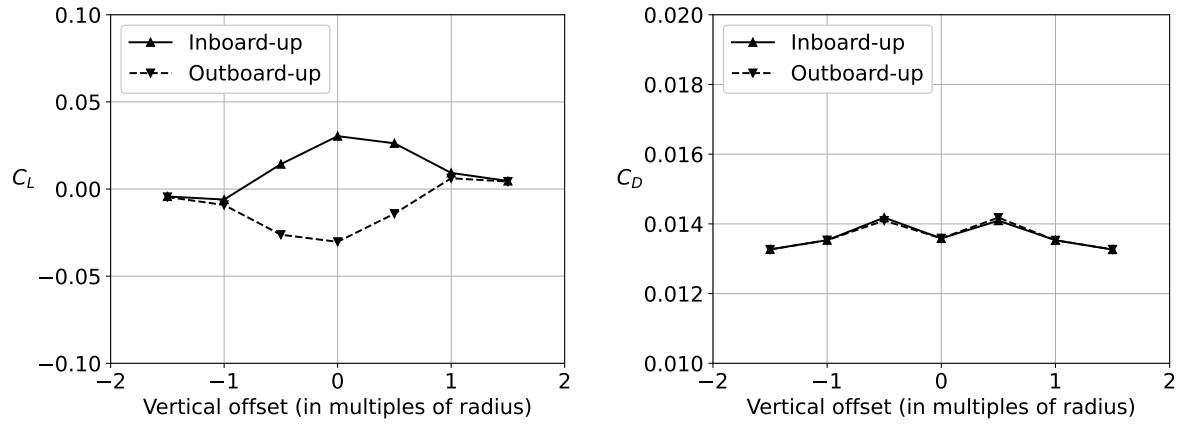
Figure 7: Wing C_d distributions for 0 deg AoA with outboard-up propeller rotation

This is expected because of the upwash from the propeller swirl. When the offset is $\pm 0.5R$, there is a reduction in the amplitude of this peak, for both positive and negative offsets. However, there are some visible differences in the distributions. The C_l is greater across the entire span for $+0.5R$ than for $-0.5R$, particularly at the wingtip portion between 0.9 $b/2$ and 1.0 $b/2$. The amplitude reduction for both positive and negative offsets can be attributed to Factor 1, and the greater lift across the span for the positive offset can be attributed to Factor 2. Recall that Factor 1 is a combination of the varying magnitude of the upwash or downwash components of the propeller swirl velocities and the varying extent of the wing inside the propeller's streamtube with respect to vertical offset. Factor 2 is the pressure and velocity distributions associated with streamtube contraction, which create upwash (for positive vertical offset) or downwash (for negative vertical offset) effects.

When the offset is increased to $+1.0R$, the C_l across the span becomes significantly lower in magnitude and is positive, and when the offset is $-1.0R$, the C_l across the span is similar in magnitude but negative. This can be attributed to Factor 2. Very little of the propeller streamtube impinges directly on the wing when the offset is $\pm 1.0R$. When the offset is $\pm 1.5R$, the propeller has negligible influence on the wing. The features in the C_l distributions for the outboard-up cases, shown in Figs. 6 and 7, can also be explained in this manner.

The C_d distributions also have striking features. For example, for the inboard-up cases, when the offset is $\pm 0.5R$, the C_d distributions are significantly different in shape from each other. For the positive offset, there is a trough followed by a steep rise. For the negative offset, there is a smooth peak without a trough. The trough for $+0.5R$ and the lack of a trough for $-0.5R$ can be attributed to Factor 2. From a lifting-line perspective, induced drag depends on the product of circulation and downwash, a positive lift and upwash (i.e., negative downwash) result in negative induced drag, which results in the trough for $+0.5R$. For $-0.5R$, Factor 2 results in downwash effects, which with the positive lift produce positive induced drag and, therefore, there is no trough for this case. The opposite is true for $+0.5R$ and $-0.5R$ with outboard-up rotation.

Results for the integrated lift coefficients, C_L, and drag coefficients, C_D, for the wing are shown in Figs. 8a and 8b, respectively.



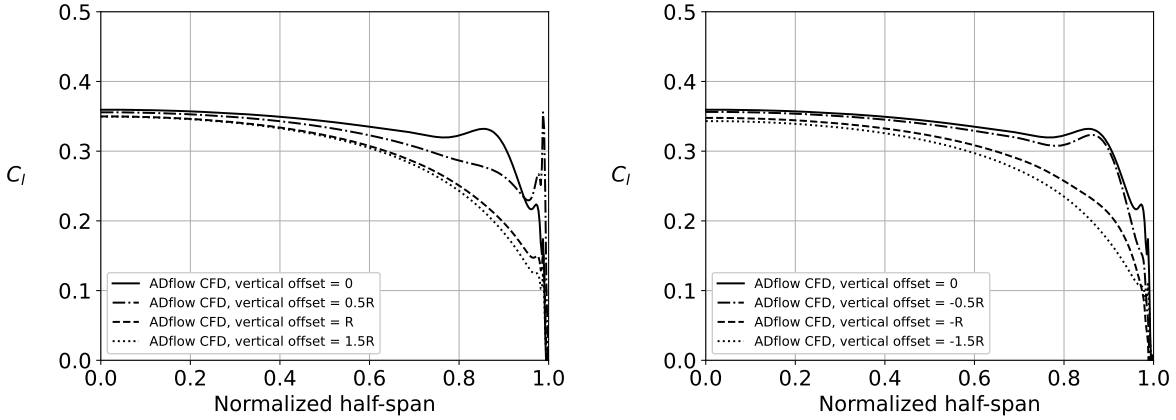
(a) Wing C_L (b) Wing C_D

Figure 8: ADflow CFD results for wing C_L and C_D at 0 deg AoA

Figures 8a and 8b show non-monotonic anti-symmetric (for C_L) and symmetric (for C_D) trends with respect to vertical offset. The gap between the C_L values for the inboard-up and outboard-up cases, for $0R$ and $\pm 0.5R$, is significantly greater than what we observed in our earlier study [18] for an inboard propeller configuration. This is because only half of the streamtube impinges directly on the wing for a wingtip propeller configuration.

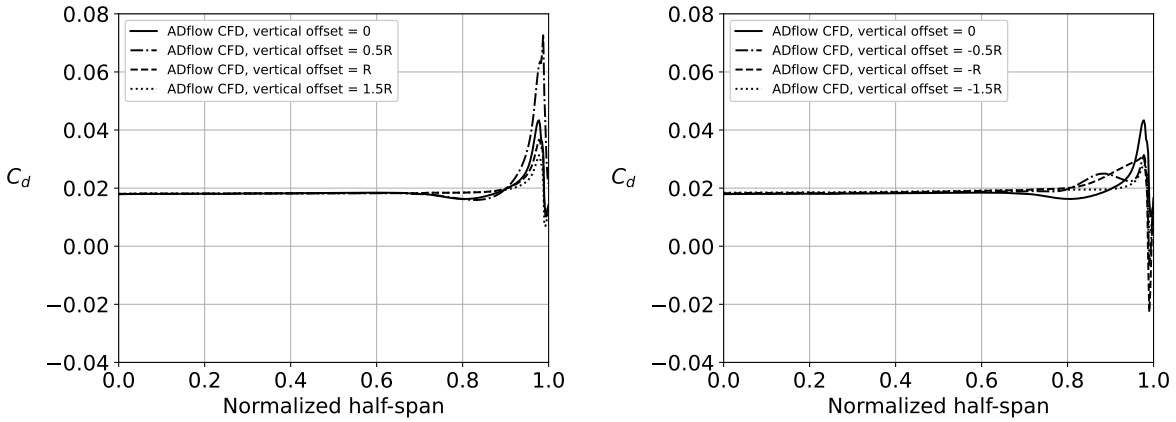
4.2 Angle of Attack = 4 deg

Figures 9 and 10 show the wing C_l and C_d distributions, respectively, for the 4 deg AoA cases with inboard-up propeller rotation. Figures 11 and 12 show the same for outboard-up propeller rotation.



(a) Positive vertical offset (b) Negative vertical offset

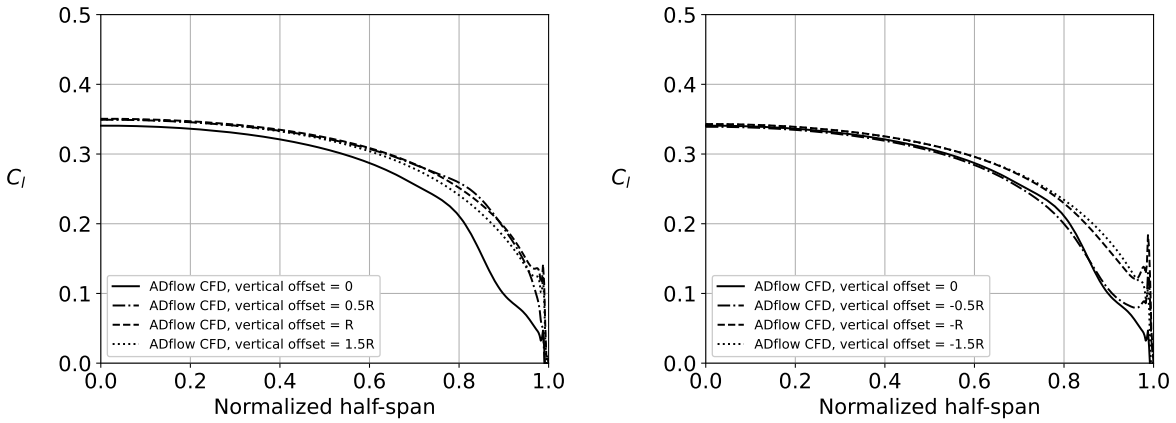
Figure 9: Wing C_l distributions for 4 deg AoA with inboard-up propeller rotation



(a) Positive vertical offset (b) Negative vertical offset

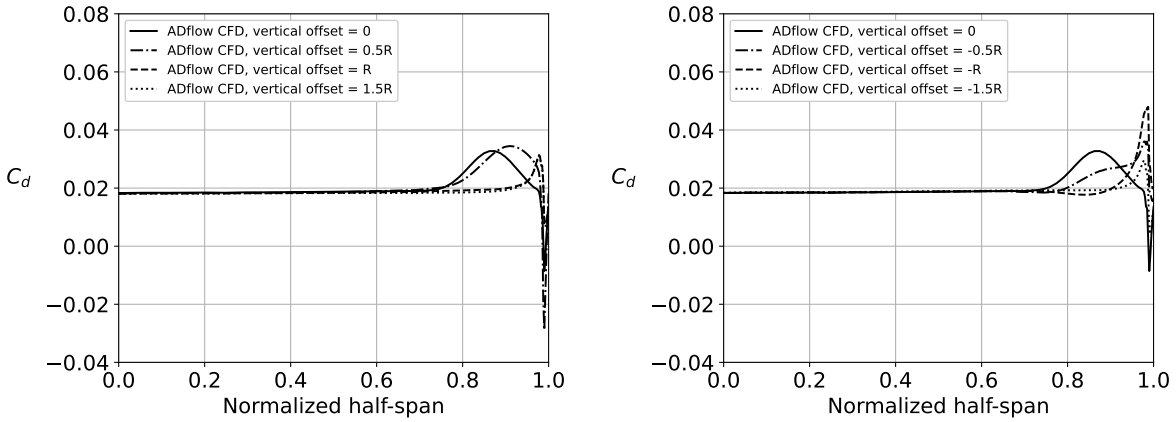
Figure 10: Wing C_d distributions for 4 deg AoA with inboard-up propeller rotation

For the inboard-up cases, when the propeller-wing vertical offset is zero, there is a peak in the C_l distributions for the portion corresponding to the upgoing-blade region. This is expected because of the upwash from the propeller swirl. When the offset is $\pm 0.5R$, there are more pronounced differences in the C_l distributions here than for the 0 deg cases.



(a) Positive vertical offset (b) Negative vertical offset

Figure 11: Wing C_l distributions for 4 deg AoA with outboard-up propeller rotation



(a) Positive vertical offset (b) Negative vertical offset

Figure 12: Wing C_d distributions for 4 deg AoA with outboard-up propeller rotation

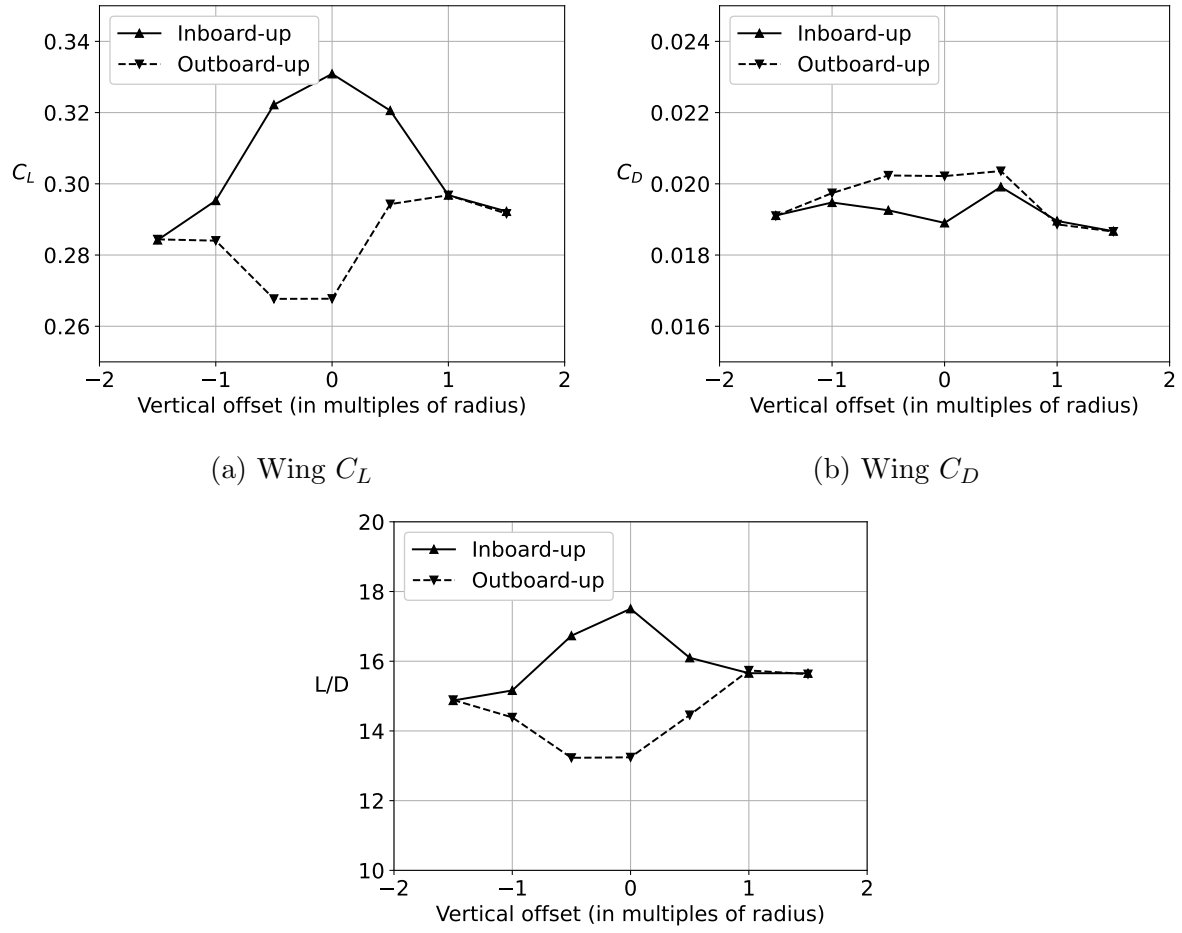
For $-0.5R$, the C_l distribution is almost unchanged. However, for $+0.5R$, the peak for the portion corresponding to the upgoing-blade region is no longer there. This can be attributed to Factor 1 and the propeller streamtube convecting with the freestream flow for this positive angle of attack. For the positive offset, the streamtube convects away from the wing, and for the negative offset, it convects towards the wing. This affects the extent of the wing inside the streamtube.

For the outboard-up cases, when the offset is zero, there is a lift reduction in the C_l distributions for the portion corresponding to the downgoing-blade region. This is expected because of the downwash from the propeller swirl. When the offset is $\pm 0.5R$, there are again more pronounced differences in the C_l distributions here than for the 0 deg cases. For $-0.5R$, the C_l distribution is almost unchanged. For $+0.5R$, the C_l distribution resembles that of an undisturbed wing. Again, this can be attributed to Factor 1 and the propeller streamtube convecting with the freestream flow for this positive angle of attack.

The C_d distributions also have striking features and dissimilarities between the curves corresponding to the positive and negative offsets. For example, for the inboard-up cases,

there is little difference between the curves for $0R$ and R, except for the absence of the shallow trough for R. However, the differences are more noticeable between the curves for $0R$ and $-R$. The drag increase for $-R$ can be attributed to a combination of Factor 2 (downwash effect), Factor 3 (increased dynamic pressure), and the freestream convecting the streamtube closer to the wing.

Figures 13a and 13b show the C_L and C_D of the wing as a function of the vertical offset for both inboard-up and outboard-up propeller rotation. The C_L is highest at $0R$ for the inboard-up cases and lowest at $0R$ and $-0.5R$, coincidentally, for the outboard-up cases. The C_D is highest at $0.5R$ and lowest at $-1.5R$ for both rotation directions. The corresponding lift-to-drag ratios, L/D, plotted in Fig. 13c show that L/D is highest at $0R$ for the inboard-up cases and lowest $0R$ and $-0.5R$, coincidentally, for the outboard-up cases.



(a) Wing C_L

(b) Wing C_D

(c) Wing L/D

Figure 13: ADflow CFD results for wing C_L, C_D, and L/D at 4 deg AoA

5 CONCLUSIONS

In this paper, we present simulation results for a rectangular wing with a tractor propeller located at the wingtip using steady RANS CFD and an actuator-disk approach. These results are for 0 and 4 deg AoA and propeller-wing vertical offsets ranging from

$-1.5R$ to $+1.5R$, for both inboard-up and outboard-up propeller rotation.

The results show that changing the propeller-wing vertical offset can significantly change a trailing wing's lift and drag distributions, and these distributions have a strong dependence on the direction of the vertical offset. The results also show an interplay of three primary factors:

1. Factor 1 is a combination of the varying magnitude of the upwash and downwash components of the propeller swirl velocities and the varying extent of the wing inside the propeller's streamtube with respect to vertical offset.

2. Factor 2 is the pressure and velocity distributions associated with streamtube contraction, which create upwash (for positive vertical offset) or downwash (for negative vertical offset) effects.

3. Factor 3 is a combination of the non-uniform radial distribution of dynamic pressure in a propeller's streamtube and the varying extent of the wing inside the propeller's streamtube with respect to vertical offset.

These factors are related to the proximity of the wing to the streamtube, which depends not only on the propeller-wing vertical offset but also on the convection of the streamtube with the freestream for non-zero angles of attack. The interplay of these factors results in trends that are non-monotonic and asymmetric with respect to propeller-wing vertical offset.

There are many possibilities for future work. One possibility is repeating this study to compare C_l and C_d distributions, as well as L/D values, for fixed C_L instead of fixed AoA. The efficient adjoint-based gradient-computation method of ADflow enables carrying out such a study in an automated and efficient manner by using a gradient-based optimizer to adjust the AoA to achieve the target C_L. Another possibility is carrying out gradient-based wing-shape optimization studies for a range of propeller-wing vertical offsets.

ACKNOWLEDGMENTS

The authors acknowledge the support of Fundação de Amparo à Pesquisa do Estado de São Paulo, FAPESP, under Research Grants No. 2013/07375-0 and 2021/11258-5, and of Conselho Nacional de Desenvolvimento Científico e Tecnológico, CNPq, under Research Grant No. 315411/2023-6.

REFERENCES

[1] L. Prandtl. Mutual influence of wings and propeller. Technical Note NACA-TN-74, NACA, December 1921. An English translation of an extract from the First Report of the Göttingen Aerodynamic Laboratory, Chapter 4, Section 6, pp. 112–118.

[2] R. E. Kuhn and J. W. Draper. An investigation of a wing-propeller configuration employing large-chord plain flaps and large-diameter propellers for low-speed flight and vertical take-off. Technical Note NACA-TN-3307, NACA, January 1954.

[3] R. E. Kuhn and J. W. Draper. Investigation of the aerodynamic characteristics of a model wing-propeller combination and of the wing and propeller separately at angles of attack up to 90°. Technical Report NACA-TR-1263, NASA, January 1956.

[4] B. W. McCormick. *Aerodynamics of V/STOL Flight*. Academic Press, 1 edition, 1967.

[5] M. D. Moore and B. Fredericks. Misconceptions of electric aircraft and their emerging aviation markets. In *52nd Aerospace Sciences Meeting*, number AIAA-2014-0535. AIAA, January 2014. doi:10.2514/6.2014-0535.

[6] A. M. Stoll, J. Bevirt, P. P. Pei, and E. V. Stilson. Conceptual design of the Joby S2 electric VTOL PAV. In *14th AIAA Aviation Technology, Integration, and Operations Conference*, number AIAA-2014-2407. AIAA, June 2014. doi:10.2514/6.2014-2407.

[7] C. Courtin, A. Mahseredjian, A. J. Dewald, M. Drela, and J. Hansman. A performance comparison of eSTOL and eVTOL aircraft. In *AIAA AVIATION 2021 Forum*, number AIAA-2021-3220. AIAA, 2021. doi:10.2514/6.2021-3220.

[8] R. de Vries, M. Brown, and R. Vos. Preliminary sizing method for hybrid-electric distributed-propulsion aircraft. *Journal of Aircraft*, 56(6):2172–2188, 2019. doi:10.2514/1.C035388.

[9] S. S. Chauhan and J. R. R. A. Martins. Tilt-wing eVTOL takeoff trajectory optimization. *Journal of Aircraft*, 57(1):93–112, January 2020. doi:10.2514/1.C035476.

[10] K. A. Deere, S. A. Viken, M. B. Carter, J. K. Viken, J. M. Derlaga, and A. M. Stoll. Comparison of high-fidelity computational tools for wing design of a distributed electric propulsion aircraft. In *35th AIAA Applied Aerodynamics Conference*, number AIAA-2017-3925. AIAA, June 2017. doi:10.2514/6.2017-3925.

[11] G. Droandi, M. Syal, and G. Bower. Tiltwing multi-rotor aerodynamic modeling in hover, transition and cruise flight conditions. In *AHS Forum 74*, number 74-2018-1267. AHS, May 2018.

[12] E. J. Alvarez and A. Ning. Meshless large-eddy simulation of propeller–wing interactions with reformulated vortex particle method. *Journal of Aircraft*, 61(3):811–827, 2024. doi:10.2514/1.C037279.

[13] S. S. Chauhan and J. R. R. A. Martins. RANS-based aerodynamic shape optimization of a wing considering propeller–wing interaction. *Journal of Aircraft*, 58(3):497–513, May 2021. doi:10.2514/1.C035991.

[14] M. H. Snyder, Jr. and G. W. Zumwalt. Effects of wingtip-mounted propellers on wing lift and induced drag. *Journal of Aircraft*, 6(5):392–397, 1969. doi:10.2514/3.44076.

[15] I. Kroo. Propeller-wing integration for minimum induced loss. *Journal of Aircraft*, 23(7):561–565, 1986. doi:10.2514/3.45344.

[16] L. R. Miranda and J. E. Brennan. Aerodynamic effects of wingtip-mounted propellers and turbines. In *4th Applied Aerodynamics Conference*, number AIAA-1986-1802. AIAA, June 1986. doi:10.2514/6.1986-1802.

[17] L. L. M. Veldhuis. *Propeller Wing Aerodynamic Interference*. PhD thesis, Delft University of Technology, 2005.

[18] S. S. Chauhan, G. Padovany da Silva, G. L. O. Halila, J. R. R. A. Martins, and J. L. F. Azevedo. Investigating the influence of an inboard propeller's vertical offset on a trailing wing's lift and drag distributions. In *34th Congress of the International Council of the Aeronautical Sciences*, number ICAS2024-0124. ICAS, October 2024.

[19] A. Yildirim, G. K. W. Kenway, C. A. Mader, and J. R. R. A. Martins. A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations. *Journal of Computational Physics*, 397:108741, November 2019. ISSN 0021-9991. doi:10.1016/j.jcp.2019.06.018.

[20] C. A. Mader, G. K. W. Kenway, A. Yildirim, and J. R. R. A. Martins. ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization. *Journal of Aerospace Information Systems*, 17(9):508–527, September 2020. doi:10.2514/1.I010796.

[21] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*, number AIAA-1992-439, June 1992. doi:10.2514/6.1992-439.

[22] S. S. Chauhan and J. R. R. A. Martins. RANS-based aerodynamic shape optimization of a wing with a propeller in front of the wingtip. *Aerospace*, 11(7), 2024. ISSN 2226-4310. doi:10.3390/aerospace11070512.

[23] S. S. Chauhan. Estimating propeller-induced tangential velocities by hand. In *AIAA SciTech 2025 Forum*, number AIAA-2025-1265. AIAA, January 2025. doi:10.2514/6.2025-1265.

[24] D. P. Witkowski, A. K. H. Lee, and J. P. Sullivan. Aerodynamic interaction between propellers and wings. *Journal of Aircraft*, 26(9):829–836, 1989. doi:10.2514/3.45848.

[25] R. M. Ardito Marretta, G. Davi, A. Milazzo, and G. Lombardi. Wing pitching and loading with propeller interference. *Journal of Aircraft*, 36(2):468–471, 1999. doi:10.2514/2.2455.

[26] J. T. Hwang and A. Ning. Large-scale multidisciplinary optimization of an electric aircraft for on-demand mobility. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number AIAA-2018-1384, January 2018. doi:10.2514/6.2018-1384.

[27] C. Alba, A. Elham, B. J. German, and L. L. M. Veldhuis. A surrogate-based multi-disciplinary design optimization framework modeling wing-propeller interaction. *Aerospace Science and Technology*, 78:721–733, July 2018. doi:10.1016/j.ast.2018.05.002.

[28] K. R. Moore and A. Ning. Takeoff and performance trade-offs of retrofit distributed electric propulsion for urban transport. *Journal of Aircraft*, 56(5):1880–1892, 2019. doi:10.2514/1.C035321.

[29] D. L. Whitfield and A. Jameson. Euler equation simulation of propeller-wing interaction in transonic flow. *Journal of Aircraft*, 21(11):835–839, 1984. doi:10.2514/3.45052.

[30] F. Moens and P. Gardarein. Numerical simulation of the propeller/wing interactions for transport aircraft. In *19th AIAA Applied Aerodynamics Conference*, number AIAA-2001-2404. AIAA, June 2001. doi:10.2514/6.2001-2404.

[31] E. W. M. Roosenboom, A. Stürmer, and A. Schröder. Advanced experimental and numerical validation and analysis of propeller slipstream flows. *Journal of Aircraft*, 47(1):284–291, 2010. doi:10.2514/1.45961.

[32] A. Gomariz-Sancha, M. Maina, and A. J. Peace. Analysis of propeller-airframe interaction effects through a combined numerical simulation and wind-tunnel testing approach. In *53rd AIAA Aerospace Sciences Meeting*, number AIAA-2015-1026. AIAA, January 2015. doi:10.2514/6.2015-1026.

[33] T. C. A. Stokkermans, N. van Arnhem, T. Sinnige, and L. L. M. Veldhuis. Validation and comparison of RANS propeller modeling methods for tip-mounted applications. *AIAA Journal*, 57(2):566–580, 2019. doi:10.2514/1.J057398.

[34] H. U. Koyuncuoglu and P. He. Simultaneous wing shape and actuator parameter optimization using the adjoint method. *Aerospace Science and Technology*, 130, 2022. doi:10.1016/j.ast.2022.107876.

[35] J. R. R. A. Martins. Aerodynamic design optimization: Challenges and perspectives. *Computers & Fluids*, 239:105391, 2022. doi:10.1016/j.compfluid.2022.105391.

[36] Y. Lee and J. Baeder. Implicit hole cutting - a new approach to overset grid connectivity. In *16th AIAA Computational Fluid Dynamics Conference*, number AIAA-2003-4128, June 2003. doi:10.2514/6.2003-4128.

[37] B. Landmann and M. Montagnac. A highly automated parallel Chimera method for overset grids based on the implicit hole cutting technique. *International Journal for Numerical Methods in Fluids*, 66(6):778–804, 2011. doi:10.1002/fld.2292.

[38] G. K. W. Kenway, N. Secco, J. R. R. A. Martins, A. Mishra, and K. Duraisamy. An efficient parallel overset method for aerodynamic shape optimization. In *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number AIAA-2025-1265, January 2017. doi:10.2514/6.2017-0357.

[39] A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge–Kutta time stepping schemes. In *14th Fluid and Plasma Dynamics Conference*, number AIAA 1981-1259, 1981. doi:10.2514/6.1981-1259.

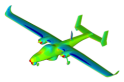
[40] N. Secco, G. K. W. Kenway, P. He, C. A. Mader, and J. R. R. A. Martins. Efficient mesh generation and deformation for aerodynamic shape optimization. *AIAA Journal*, 59(4):1151–1168, April 2021. doi:10.2514/1.J059491.

[41] W. M. Chan and J. L. Steger. Enhancements of a three-dimensional hyperbolic grid generation scheme. *Applied Mathematics and Computation*, 51(2–3):181–205, 1992. doi:10.1016/0096-3003(92)90073-A.

[42] M. Hoekstra. A RANS-based analysis tool for ducted propeller systems in open water condition. *International Shipbuilding Progress 53*, 53:205–227, 2006.

[43] M. Selig. University of Illinois Urbana-Champaign Airfoil Coordinates Database. `https://m-selig.ae.illinois.edu/ads/coord_database.html`.

[44] T. Sinnige, R. de Vries, B. D. Corte, F. Avallone, D. Ragni, G. Eitelberg, and L. L. M. Veldhuis. Unsteady pylon loading caused by propeller-slipstream impingement for tip-mounted propellers. *Journal of Aircraft*, 55(4):1605–1618, 2018. doi:10.2514/1.C034696.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MISSION-INFORMED SAMPLING STRATEGY FOR AN AERODYNAMIC MULTI-FIDELITY SURROGATE MODEL

Maurice Zimmnau, Eike Stumpf

Institute of Aerospace Systems
RWTH Aachen University
Wuellnerstrasse 7, 52062 Aachen, Germany
research@ilr.rwth-aachen.de

Abstract. *In this contribution, we propose an adaptive sampling strategy employing Gaussian process regression for computationally expensive aerodynamic predictions. In contrast to standard adaptive sampling strategies such as expected improvement – which rely on the Gaussian process statistics and function evaluation values – this strategy also integrates information from the conceptual aircraft design process, i.e., the mission simulation. The relative frequencies of the flight conditions are derived from the mission information, and a kernel density estimation is performed. Both the kernel density estimate and the posterior standard deviation of the Gaussian process inform the selection of the next sample location. The primary objective of this strategy is not to find an optimum, but rather to improve the accuracy of the surrogate model at locations of interest. We compare the proposed sampling strategy with commonly used adaptive sampling approaches. The effects of Gaussian process hyperparameter choices on the predicted sample locations are analyzed. Different bandwidth selection methods for the kernel density estimation are evaluated. The applicability and limitations of this approach are examined in detail. A special focus is placed on cases where the kernel density estimate's optima arise at the boundary of the design space.*

Keywords: Gaussian process regression, surrogate models, multi-fidelity, aerodynamics, conceptual aircraft design

1 INTRODUCTION

In the early conceptual aircraft design phase, one of the objectives is to explore the design space and successively narrow it down, ultimately choosing a final design for the preliminary and detailed design phases [1]. This design phase is the cheapest but of great importance. Incorrect design decisions made during this phase are costly to rectify in later phases. Therefore, obtaining reliable results in the conceptual design phase is crucial. Typically, conceptual design methods are of low fidelity and based on semi-empirical regression models [2]. These methods are computationally fast and well-suited for exploring large design spaces applicable to conventional aircraft configurations. However, they often lack accuracy, particularly regarding so-called unconventional aircraft configurations. On the other hand, computationally expensive methods such as FEM for structures or CFD for aerodynamics are often inappropriate due to time constraints. A prominent approach to reducing computational demand is utilizing surrogate models ,also known as machine learning models. These models predict the functional behavior of an unknown function, using sample locations (inputs) and observations (function evaluations). Gaussian Process Regression (GPR), also referred to as Kriging, is a common method for surrogate modeling when evaluating the objective function is costly, or when dealing with sparse datasets. Within this contribution, a surrogate model based on GPR is constructed using a limited number of CFD samples. When computational resources are limited, a pertinent question arises: where should we conduct our next expensive function evaluation? Since GPR is a statistical method, various approaches – so-called adaptive sampling strategies – exist, that leverage these statistical properties to propose optimal sample locations. However, these strategies rely solely on the GPR statistics. Given that we possess *engineering information* from the conceptual aircraft design process, we can integrate this information into our sampling strategy to direct it toward regions of interest. To effectively utilize this available information, it must be converted into appropriate format so that it can influence GPR statistics. Therefore, using relative frequencies obtained from simulated flight missions, we generate a kernel density estimate (KDE), which can be multiplied by a posterior standard deviation derived from GPR. This product yields a candidate for the next sample location.

This contribution is structured as follows: In Sec.2, GPR is briefly introduced. Next, adaptive sampling strategies, along with kernel density estimation will be presented. This section concludes with an introduction to the proposed mission-informed sampling strategy. In Section Sec.3 the proposed sampling strategy is applied subsequently to a one-dimensional, a two-dimensional and a three-dimensional application case. In Section Sec.4 the model sensitivities and robustness are analyzed concerning varying hyperparameters and modifications of the proposed sampling strategy. This contribution concludes with final conclusions found in Section Sec.5.

2 METHODS

As discussed earlier, the surrogate model employed in this work relies on Gaussian Process Regression (GPR) [3]. A Gaussian process generalizes the normal (Gaussian) distribution and defined by a mean function $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$ for

an unknown function $f(\mathbf{x})$ as follows:

$$m(\mathbf{x}) = \mathbb{E}[f(\mathbf{x})] \tag{1}$$
$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))] \tag{2}$$

For the covariance function we employ the squared-exponential kernel, expressed in one dimension for two sample points x, x' as:

$$k(x, x') = \sigma_f^2 \exp\{-\frac{1}{2l^2}(x - x')^2\} + \sigma_n^2 \delta_{\cdot,'} \tag{3}$$

To predict a function value at new locations X_* for a model trained on training points X, we use the predictive mean of the GPR given by:

$$\bar{\mathbf{f}}_* = \mathbf{m}(X_*) + K(X_*, X)\left[K(X, X) + \sigma_n^2 I\right]^{-1}(\mathbf{y} - \mathbf{m}(X)) \tag{4}$$

The posterior covariance at locations X_* is given by:

$$\text{cov}(\mathbf{f}_*) = K(X_*, X_*) - K(X_*, X)\left[K(X, X) + \sigma_n^2 I\right]^{-1} K(X, X_*) \tag{5}$$

2.1 Adaptive sampling strategies for Gaussian processes

Based on GPR (or Kriging), various adaptive sampling strategies [4] have been proposed, which pursue exploration, exploitation, or a balance of both within the design space. The first most straightforward approach is to minimize the posterior covariance across the design space – or in other words, to sample at the location with the highest posterior variance (predicted model uncertainty). Sampling at locations with highest variance is commonly referred to as pure exploration:

$$\mathbf{x}_{n+1} = \arg\max_{\mathbf{x} \in \mathcal{X}} \hat{s}^2 \tag{6}$$

with $\hat{s}^2 = \text{diag}(\text{cov}(\mathbf{f}_*))$. Another sampling strategy that balances exploration and exploitation is Expected Improvement (EI) [5]. This strategy is widely used in the context of Bayesian Optimization (BO) and Efficient Global Optimization (EGO) [6].

$$EI_{y^*}(\mathbf{x}) = \mathbb{E}\left[[y^* - f(\mathbf{x})]_+\right] = \sigma(\mathbf{x}) h\left(\frac{y^* - \mathbf{m}(\mathbf{x})}{\sigma(\mathbf{x})}\right) \tag{7}$$
$$= \sigma(\mathbf{x})\phi\left(\frac{y^* - \mathbf{m}(\mathbf{x})}{\sigma(\mathbf{x})}\right) + (y^* - \mathbf{m}(\mathbf{x}))\Phi\left(\frac{y^* - \mathbf{m}(\mathbf{x})}{\sigma(\mathbf{x})}\right) \tag{8}$$

with ϕ, Φ denoting the standard density and the distribution function and $[\cdot]_+$ being the $\mathbf{max}(0, \cdot)$ operator. y^* is the current minimum of the observations.

The two previously discussed adaptive sampling strategies are by no means the only approaches found in literature. Many other strategies exist, such as probability of improvement, upper/lower confidence bound, goal seeking, and constrained expected improvement [7]. However, for this work, only the prominent sampling strategy (EI) is used for comparison with the later proposed approach.

2.2 Kernel density estimation

Various techniques exist to estimate a probability density function (PDF) from observed data, including parametric and nonparametric density estimations. Parametric density estimations are typically applied to known distributions – such as the Gaussian distribution – parameterized by its mean μ and variance σ^2. The distributions covered by parametric density estimates are usually unimodal (with a global maximum) [8]. In practical applications, the distributions are often not unimodal, but multimodal. For these applications, nonparametric density estimations are more suitable. In this contribution, we focus on nonparametric kernel density estimators.

Let $f(\mathbf{x})$ be a probability density function satisfying the Riemann-density properties $f(\mathbf{x}) \geq 0$ and $\int_{-\infty}^{\infty} f(\mathbf{x}) = 1$ [9]. For the univariate case, the kernel density estimation of $f(\mathbf{x})$ is denoted by:

$$\hat{f}(x,h) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{x - X_i}{h}\right) = \frac{1}{n}\sum_{i=1}^{n} K_h(x - X_i) \tag{9}$$

with h being the bandwidth. Usually the kernels have the form of a normal distribution $\mathcal{N}(0,1)$ [8]:

$$K\left(\frac{x - X_i}{h}\right) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{(x - X_i)^2}{2h^2}\right) \tag{10}$$

and the data points can be interpreted as the mean values of the normal distribution:

$$\hat{f}(x,h) = \frac{1}{n}\sum_{i=1}^{n} \mathcal{N}(x; X_i, h^2) \tag{11}$$

An example for a kernel density estimation $\hat{f}(x)$, estimated from samples x_i for a dummy function $f(x)$, is depicted in Fig.1.

A nonparametric kernel density estimator requires an estimation of the bandwidth, also called smoothing parameter. In literature, there exist different methods for estimating the bandwidth parameter, so-called bandwidth selectors. Three major types of selectors can be distinguished: *rule-of-thumb* (ROT), *cross-validation* (CV), and *plug-in* (PI) selectors. [8] Two ROT methods are *Silverman* [10] and *Scott* [11]. These methods yield good results for univariate distributions, however, tend to oversmooth multivariate density estimations [12]. In this contribution, for the basic KDE, we use the ROT method from *Scott*.

2.3 Proposed sampling strategy

The sampling strategy for this contribution was presented for the one-dimensional case in [13], for the three-dimensional case in [14] and in [15]. In this contribution, we especially focus on the analysis of the sampling strategy and compare it to already introduced off-the-shelf sampling strategies.

The concept of the sampling strategy is to use the posterior standard deviation $\hat{s}$ of the GPR and bias it by a KDE. This ensures a weighted sampling at locations of interest and at high model uncertainty. The posterior covariance can be denoted by the diagonal entries of eq. (5):

$$\hat{s}^2 = \mathrm{diag}\{\mathrm{cov}(\mathbf{f}_*)\} \tag{12}$$

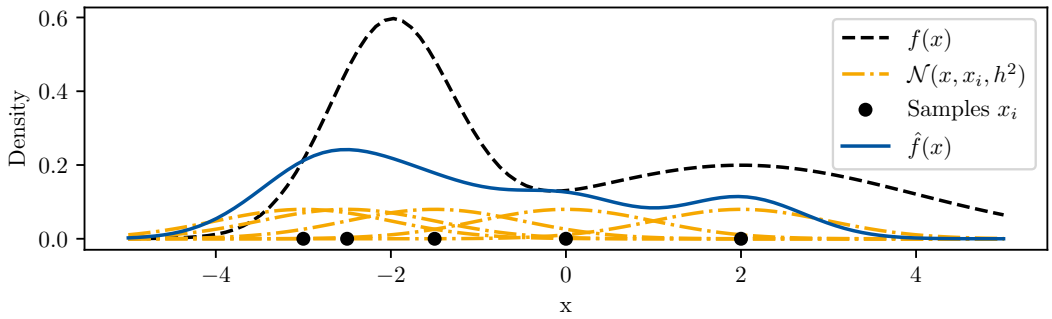


Figure 1: Exemplary kernel density estimate $\hat{f}(x)$ constructed from samples x_i

The objective is to minimize the biased uncertainty *globally* for a next sample with the following acquisition function:

$$\mathbf{x_{n+1}} = \arg\min_{\mathbf{x}\in\mathcal{X}} \int_{\mathcal{X}} \mathrm{kde}(r) \cdot \hat{s}_{+x}(r) dr \tag{13}$$

The investigation of interest is how the KDE parameters and methods, as well as the GPR parameters affect the next sample location.

3 APPLICATION

The application case for this contribution is the aerodynamic drag prediction of a Blended Wing Body (BWB) for a 3000 NM mission. For the sampling procedure, the relative frequencies of the mission flight conditions are given via a mission simulation[1], where for each time step the equations of motion are solved, assuming the aircraft to be a point mass. For each time step, there are a lift coefficient C_L, a flight altitude h, and a Mach number Ma. The proposed sampling strategy uses these relative frequencies of the flight conditions to estimate a kernel density.

The sampling procedure is embedded within the training of a GPR, where the mean function (lower fidelity – LF) for the aerodynamics for the respective flight condition is given via the *aerodynamic_analysis* module of the conceptual aircraft design environment *UNICADO* [16]. CFD simulations (higher fidelity – HF) are emulated[2], only for the samples derived from the presented sampling strategy, resulting in a surrogate model of variable fidelity (VF).

3.1 1D application

Within the one-dimensional case we investigate the sampling strategy for the functional relation $C_D = f(C_L)$, i.e. the drag coefficient as a function of the lift coefficient. For this, the relative frequencies of C_L during the simulated mission are of particular interest.

For the investigated mission, the C_L values over mission time are depicted in Fig.2 (left). It is important to note that we are investigating only the clean configuration, therefore excluding the take-off and landing phases. It can be observed that during the mission, one step climb is performed. In Fig.2 (right) a histogram is depicted showing

[1]performed with the *mission_analysis* module of the conceptual aircraft design environment *UNICADO* [16]

[2]i.e. dummy samples are produced by adding noise to the LF results

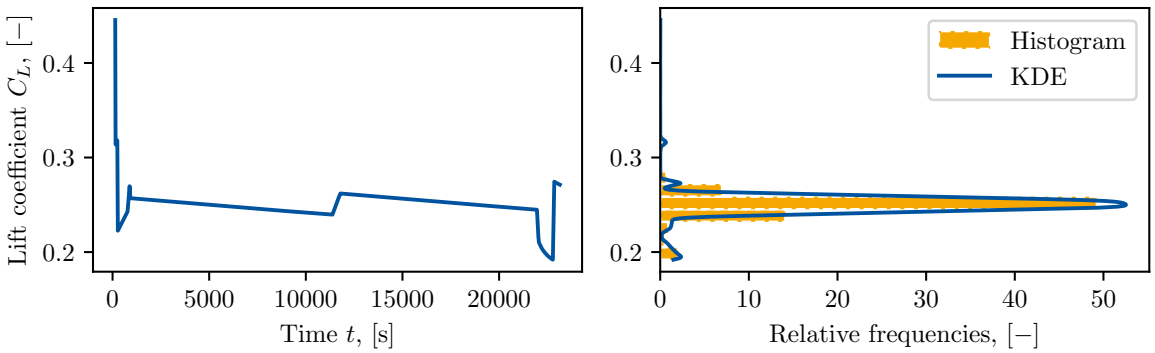


Figure 2: Occurring lift coefficients C_L over time during a simulated 3000 NM mission (left) and the relative frequencies of C_L values (right)

the relative frequencies of the C_L values. Furthermore, a KDE is plotted applying the ROT method from *Scott*, utilizing the module *gaussian_kde* of the open source library SciPy [12]. We can deduce from the distribution of C_L values throughout the mission that in the majority of the mission time, a similar value is present, peaking at $C_L = 0.253$.

If we assume that the computational resources are very limited, we aim to limit the CFD simulations to a number of 1 to 3. Thus, the first sample location for CFD simulation should return a maximum information gain or a maximum uncertainty reduction at locations of interest. If we analyze both sampling strategies EI, and the proposed mission-informed sampling strategy (MIS), a major difference is that EI uses the observations. Instead MIS uses no observations but a KDE which contains the locations of interest. The acquisition function of MIS is depicted in Fig.3 (left).

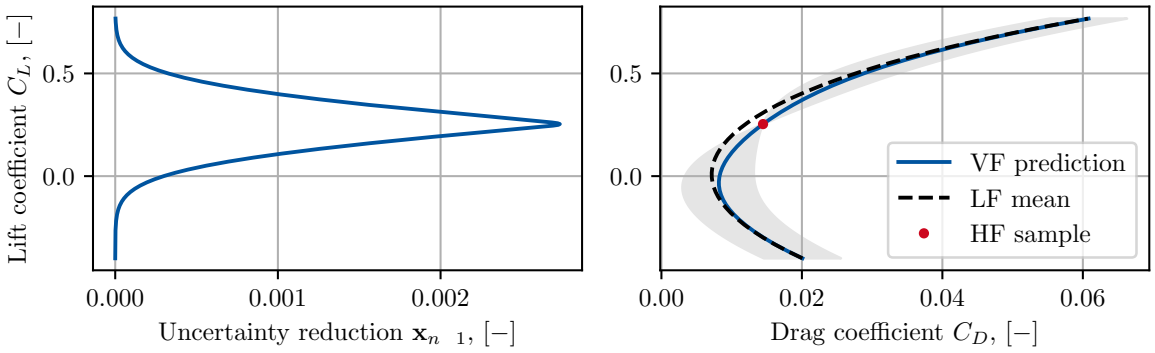


Figure 3: Acquisition function MIS for the first sample location for a CFD simulation (left) and a VF prediction where the LF polar is enhanced by an artificial HF sample

The acquisition function shows, for each C_L location, the extent to which an additional sample would reduce the model uncertainty biased by the mission information distribution. The sample with the maximum uncertainty reduction is located at $C_L = 0.253$. In Fig.3 (right) an artificial HF sample is inserted at the proposed location and the GPR is trained to fit this observation.

In comparison, as already mentioned, EI needs observations to propose a next sample location. Furthermore, the strategy uses the *optimal* observation. If there are no HF observations available, the optimum of the mean function can be used for the first sample.

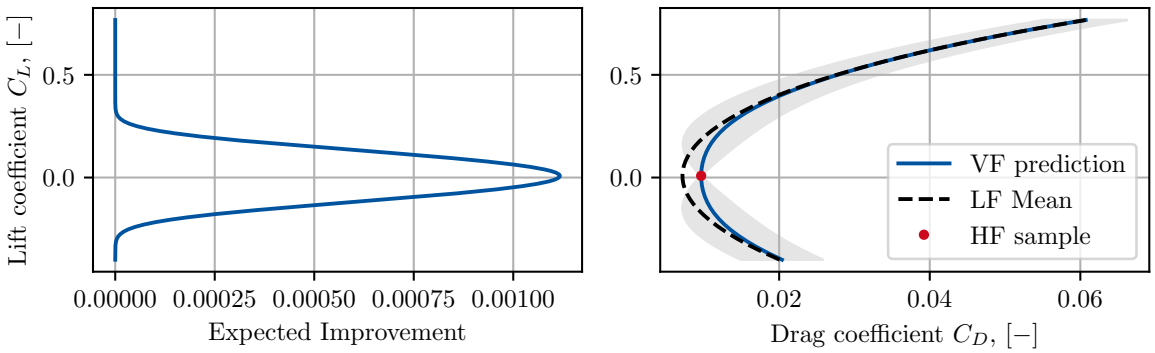


Figure 4: Acquisition function EI for the first sample location for a CFD simulation (left) and a VF prediction where the LF polar is enhanced by an artificial HF sample

The acquisition function is depicted in Fig.4 (left). Therefore, the first expensive evaluation is at $C_L \approx 0.0$. Note that the uncertainty, denoted by the gray tube, is increased in the vicinity of the flight condition of interest at $C_L = 0.253$.

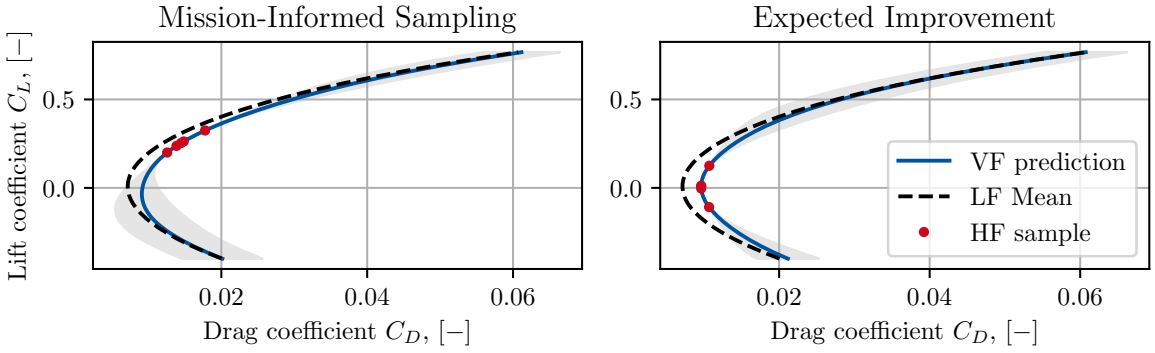


Figure 5: Comparison of the adaptive sampling strategies MIS and EI for five HF samples

For five HF samples, the locations of those samples are depicted in Fig.5 for both sampling strategies. For MIS, the samples are concentrated around the mission cruise flight condition, where for EI, the samples are concentrated near the minimum drag coefficient at $C_D \approx 0.0$.

3.2 2D application

In the two-dimensional application, we will investigate two combinations for functional relations, namely: $C_D = f(C_L, h)$ and $C_D = f(C_L, \mathrm{Ma})$. The flown altitude over time during the simulated mission is depicted in Fig.6 (left). A step climb is performed during the mission from the initial cruise altitude of FL330 (10 058 m) to FL350 (10 668 m). The one-dimensional histogram and KDE (Fig.6 (right)) have two local optima for the relative frequencies at the two mentioned flight levels.

The Mach number over time during the mission is depicted in Fig.7. After reaching the initial cruise altitude, the Mach number remains constant at $\mathrm{Ma} = 0.82$ throughout the mission until descent. The KDE has its maximum at this cruise Mach number.

When applying the KDE with ROT from *Scott* on the two-dimensional design space, the

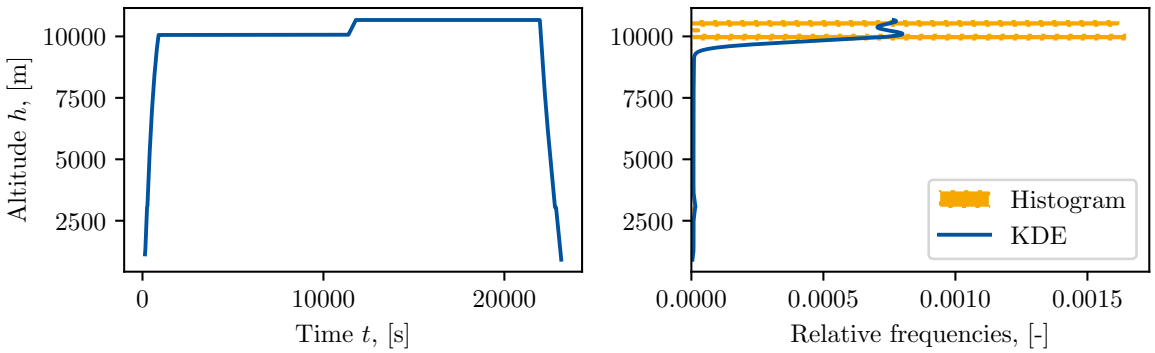


Figure 6: Occurring flight altitudes h over time during a simulated 3000 NM mission (left) and the relative frequencies of h values (right)

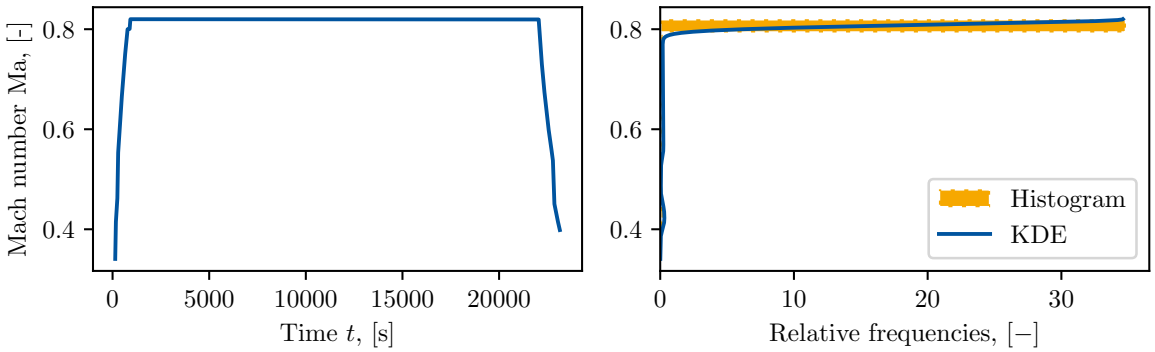


Figure 7: Occurring Mach numbers Ma over time during a simulated 3000 NM mission (left) and the relative frequencies of Ma values (right)

resulting contour plots are shown in Fig.8. Note that for the $C_D = f(C_L, h)$ application, the Mach number is fixed at $\mathrm{Ma} = 0.82$. For the $C_D = f(C_L, \mathrm{Ma})$ application, the altitude range is fixed at $h \in [10\,058\,\mathrm{m}, 10\,668\,\mathrm{m}]$. Fixing a Mach number or allowing only a certain range of altitudes reduces the dataset of relative frequencies of flight conditions. This also accounts for optima represented by the KDE, which could diffuse if a KDE is constructed with all available flight conditions (Subsec.3.3). In the contour plot Fig.8(left), the KDE clearly represents the two main flight conditions at FL330 and FL350 with the highest relative frequencies. For FL330 the lift coefficient C_L ranges between 0.24 and 0.258, while for FL350 the C_L range is slightly increased with a range between 0.245 and 0.262. It can be recognized, that the relative frequencies of FL350 are on the boundary of the design space, i.e. on the simulated mission there do not exist higher flight altitudes than FL350.

For the application case $C_D = f(C_L, \mathrm{Ma})$, the contour plot seems to be empty, where the color bar of the relative frequencies state, that there are regions with high relative frequencies. As previously described, the Mach number has its maximum at $\mathrm{Ma} = 0.82$. Like in the design space of the flight altitudes, the Mach number of 0.82 is the highest Mach number flown, which marks the boundary of the design space. Therefore, the highest relative frequencies are depicted on the upper boundary. The C_L values with the highest relative frequencies range between 0.24 and 0.261 again.

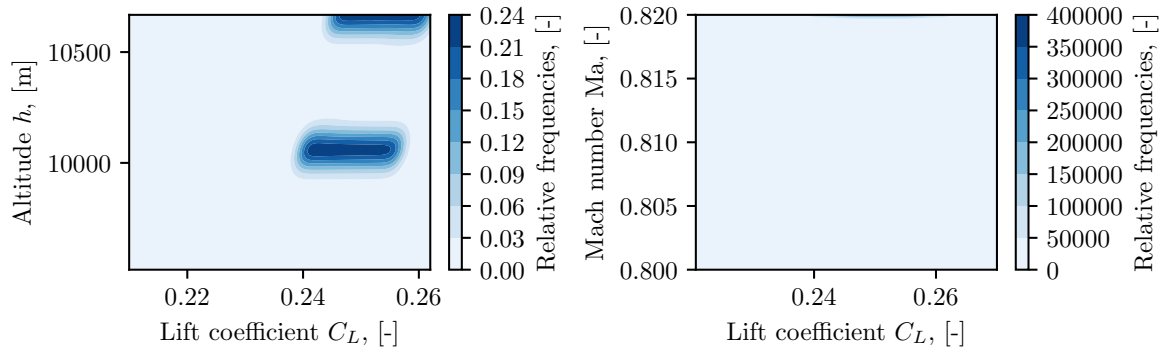


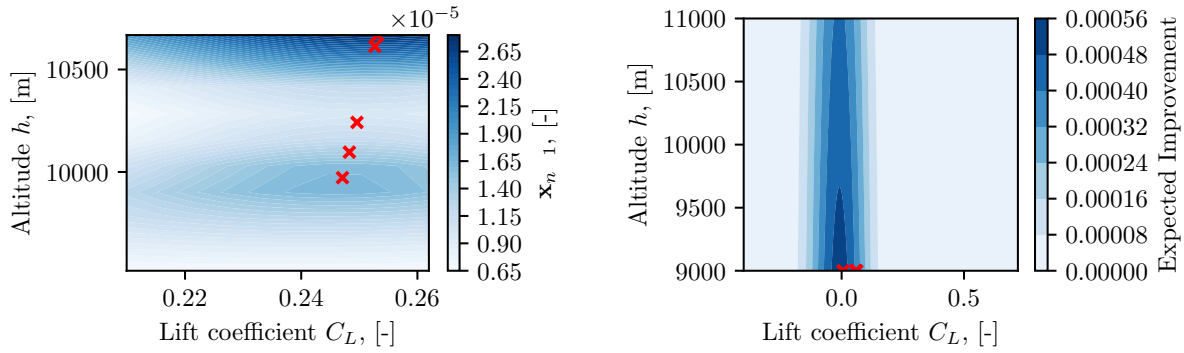
Figure 8: Contours of the estimated relative frequencies of occurring flight conditions C_L, h (left) and C_L, Ma (right) for a simulated 3000 NM mission

For both two-dimensional relations again the proposed MIS strategy is compared to EI. The comparison is made with five sample points, which are added to the LF-trend function. The HF dummy samples are again simulated by the LF evaluation with added random noise with $\mathcal{N}(0, 0.005)$, indicating a standard deviation of 50 drag counts to the LF predictions. Since the sample locations can be located anywhere in the given design space and we emphasize fast predictions of the LF function, we train a new GPR model for the LF prediction from three polars evenly spaced from 9000 m to 11 000 m at a fixed Mach number of $\mathrm{Ma} = 0.82$ for the case $C_D = f(C_L, h)$ and for a Mach numbers from 0.78 to 0.82 for a fixed altitude at $h = 10\,668\,\mathrm{m}$ for the case $C_D = f(C_L, \mathrm{Ma})$. The mean function in Eq.(4) is implemented here as a GPR as-well. As each polar is represented densely with > 1000 function evaluations by *aerodynamic_analysis*, a hyperparameter tuning is performed, to fit the observed data. Note that, in contrast to the multi-fidelity GPR model, the LF GP model is constructed with a 0-mean trend function. It is also important to note, that we only need these LF evaluations, because EI needs function evaluations. However MIS does not need any function evaluations, only the mission information and the posterior variance $\mathrm{cov}(\mathbf{f}_*)$, Eq.(5).

For comparison of MIS and EI in Fig.9 the contours of both acquisition functions for $C_D = f(C_L, h)$ are depicted, after adding five sample points. The five sample points are indicated by red crosses. The acquisition functions have been evaluated sequentially for each added sample, i.e. it is not a batch procedure. For EI (Fig.9b), the five samples are not clearly distinguishable, since they are all located in the vicinity of $(C_L = 0.0, h = 9000\,\mathrm{m})$, where also for the fifth sample the contour yields a maximum of the acquisition function in this region.

For MIS (Fig.9a) the samples are clearly in the region with the highest relative frequencies of flight conditions with $C_L \in [0.245, 0.255]$ and $h \in [10\,000\,\mathrm{m}, 10\,668\,\mathrm{m}]$. The contours suggest, that the most information gain or uncertainty reduction can be achieved, when sampling at the flight altitudes with the highest relative frequencies, namely FL330 and FL350. However there exist also samples in between, because the MIS with the integral in its acquisition function, samples at locations, where the uncertainty is reduced *globally* at the locations of interest.

For comparison of EI and MIS of the case $C_D = f(C_L, \mathrm{Ma})$, the contours of both acquisition functions with five samples are depicted in Fig.10. In contrast to the previous case, the MIS strategy now has five indistinguishable sample points at $(C_L = 0.25, \mathrm{Ma} =$

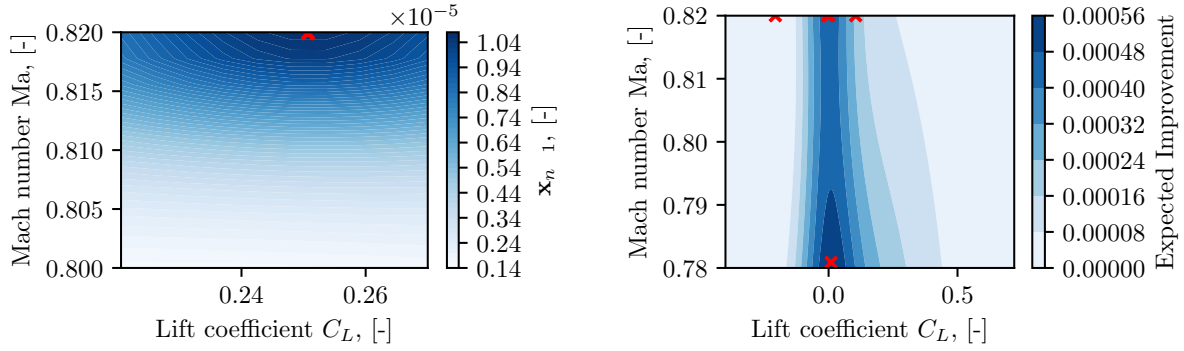


(a) MIS (b) EI

Figure 9: MIS and EI with five added samples for the case $C_D = f(C_L, h)$

0.82), i.e. at the global maximum of the relative frequencies for this application case. This result has two indications: Either one sample is sufficient to reduce the uncertainty at this location of interest, or the KDE has the higher weight, compared to the posterior variance. For the latter indication a weighting factor can be introduced between the KDE and the posterior variance, to balance the ratio between exploration and exploitation, which is dealt with in Subsec.4.3.

The EI strategy has five samples, of which 4 are distinguishable. All are located in the vicinity of $C_L = 0.0$, where the minimum of the LF polar is located. However, at $\mathrm{Ma} = 0.82$, there are three distinguishable samples, $(-0.2, 0.82)$, $(0.0, 0.82)$ and $(0.1, 0.82)$ which are at the Mach number of interest but not near the C_L of interest.



(a) MIS (b) EI

Figure 10: MIS and EI with five added samples for the case $C_D = f(C_L, \mathrm{Ma})$

3.3 3D application

In the three-dimensional application, we investigate the sampling for the functional relationship $C_D = f(C_L, h, \mathrm{Ma})$. Since the relative frequencies are depicted in great detail in the previous subsections, in this subsection, only the results of the proposed samples for both, MIS and EI are described briefly.

We expect the MIS strategy to sample at the locations with the highest relative frequencies, i.e. at the locations of interest. However, the data (Tab.1) shows, that there

is more expected information gain or uncertainty reduction not exactly at the expected locations. E.g. there is none of the five samples at FL350 nor is there a sample at exactly Ma = 0.82. Still the samples are in the region of interest with the highest relative frequencies.

For EI, as in the previous cases, we observe samples mainly near the minimum Drag at $C_D \approx 0.0$ with one sample in the C_L region of interest. Furthermore, the majority of samples is located at low flight altitudes with $h \in [1259\,\mathrm{m}, 1561\,\mathrm{m}]$ which are occurring flight conditions, but of course not the ones with the highest relative frequencies. However one flight altitude is at $h = 9869\,\mathrm{m}$ and therefore, in the region of interest. The Mach numbers are all in the region of interest. But in summary EI samples mostly occur at locations where a minimum drag is expected and rarely at the regions of our interest.

Table 1: MIS and EI with five added samples for the case $C_D = f(C_L, h, \mathrm{Ma})$

	MIS			EI		
#	C_L	h	Ma	C_L	h	Ma
1	0.250	9986 m	0.807	-0.011	1543 m	0.784
2	0.253	10 435 m	0.810	0.256	1277 m	0.82
3	0.246	9481 m	0.796	-0.153	1561 m	0.82
4	0.256	10 014 m	0.781	-0.007	9869 m	0.82
5	0.242	9054 m	0.778	0.007	1259 m	0.82

4 ANALYSIS

In the previous section two sampling strategies are compared, namely MIS and EI. It can be concluded that for the target of enhancing the accuracy of the aerodynamic predictions with a surrogate model based on GPR with a LF trend function and HF CFD simulations with restricted computational resources, the MIS strategy is superior to EI, with regard to sampling the most informative locations. Therefore, in this section, only the MIS strategy is further investigated.

For the one-dimensional case $C_D = f(C_L)$, as well as for the two-dimensional case $C_D = f(C_L, h)$, MIS sampled at locations of interest robustly, i.e. sampling at the maximum of relative frequencies, not sampling at the same location, staying in the region of interest. However, for the cases $C_D = f(C_L, \mathrm{Ma})$ and $C_D = f(C_L, h, \mathrm{Ma})$, the former samples got stuck in the same location at $(0.251, 0.82)$, where the information gain at an already sampled location shall be nearly 0. The latter case is not exactly sampling at the locations of highest frequencies, but in the vicinity. But we expect that the *optima* of the KDE are sampled.

There are four possible reasons for this deviation. The first reason is the choice of the hyperparameters of the Gaussian Process. These parameters influence the posterior covariance of the Gaussian process, which in turn has influence on the next sample location. The second reason is the location of the optima of relative frequencies of flight conditions with regard to the overall design space since the Mach number and the flight altitude FL350 are on the boundary of the design space. With the MIS strategy we evaluate the uncertainty reduction *integrated*. A sample on the boundary of the design returns only *half* of the information which a sample in the center of the design space would return. The third reason, which is proven to be wrong, is that the optimizer does not find the true

optimum of the acquisition function. The fourth reason is the selection of the bandwidth parameter for the KDE. As stated in Subsec.2.2, the ROT bandwidth selection with the rule from *Scott* tends to oversmooth multivariate, multimodal distributions.

4.1 Choice of hyperparameters

Since the posterior covariance $\mathrm{cov}(\mathbf{f}_*)$ is included in the acquisition function of MIS, the hyperparameters of the Gaussian Process have influence on the next sample location. We started our initial hyperparameter choice by intuition or engineering guess, since hyperparameter tuning is not possible without CFD HF samples in the gaussian process. Furthermore, with one to five samples, which are all located at a close range, the likelihood for overfitting is high. The chosen hyperparameters are depicted in Tab.2, for the case $C_D = f(C_L, \mathrm{Ma})$. The choice of the initial hyperparameters can be stated as follows: We expect a prior standard deviation for the LF drag prediction of 10%, which results in the given $\sigma_f = 0.0028$. Furthermore, we expect the CFD sample to have an uncertainty / measurement error (standard deviation) of 5 drag counts, i.e. $\sigma_n = 0.0005$. For the length-scale parameters, we expect for the lift coefficient a rather global shift with a new HF sample l_{C_L}, where we expect the mach number to have a more local influence due to transsonic effects with $l_{\mathrm{Ma}} = 0.05$.

Table 2: hyperparameter choice for the Gaussian Process for the case $C_D = f(C_L, \mathrm{Ma})$

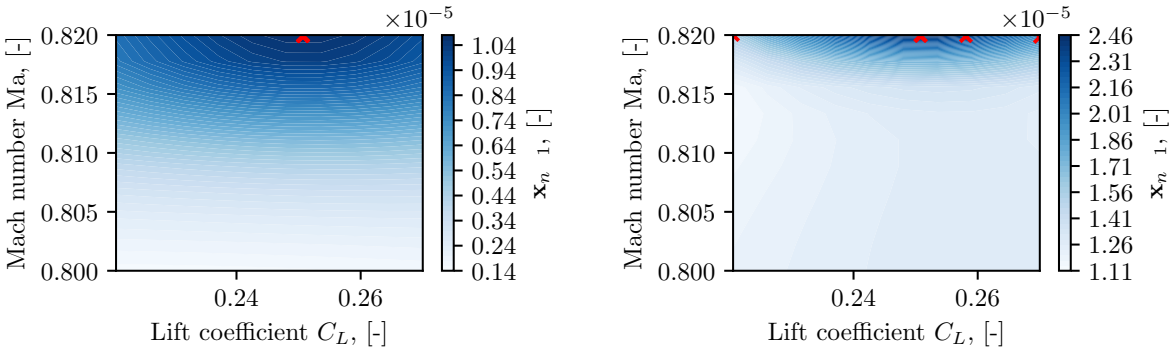
	σ_f	σ_n	l_{C_L}	l_{Ma}	ϵ
Initial parameters	0.0028	0.0005	0.25	0.05	1e−8
Tuned parameters	0.1303	1e−6	0.48	0.27	1e−6

Within this contribution, we also trained a GP with the LF data. An interesting approach for the hyperparameter choice is to use the hyperparameters from the LF GP for the multi-fidelity GP, as we expect the same global trends of the functional relation, i.e. the drag polar is a nearly quadratic function. The likelihood of overfitting is reduced on the one hand, by the amount of (cheap) data for the LF process and on the other hand, by a multi-start optimization procedure for the hyperparameter tuning. The hyperparameters estimated by the tuning process are depicted in Tab.2 as well, depicted as *Tuned parameters*. The regularization parameter[3] ϵ is chosen larger, just like the prior variance σ_f^2 and both length scale parameters, indicating larger expected uncertainty at far distance from samples, but at the same time, a global trend for the observations of the C_L value and the Mach number. The noise variance is expected to be smaller, assuming a higher precision of the CFD computations. However, the model cannot estimate an accuracy of the CFD method; it can only set the hyperparameters to fit the observations optimally.

For the tuned parameter set, the contour plot is depicted in Fig.11b. In comparison to the initial hyperparameter set (Fig.11a), where the samples are located at the same position, the samples with the tuned parameter set are spread at a constant Mach number but for varying C_L values. Only two samples are very close to the same location at $(0.251, 0.82)$.

For the three-dimensional case, the hyperparameter sets are compared as well in Tab.3.

[3]Small value, added to the diagonal entries of the matrices in the GP, ensuring stability for matrix inversion. Avoiding singularity. Has a similar effect as the noise variance σ_n^2



(a) MIS with the initial hyperparameter set (b) MIS with the tuned hyperparameter set

Figure 11: Comparison of MIS with different hyperparameter sets for the Gaussian Process with five added samples for the case $C_D = f(C_L, \mathrm{Ma})$

Compared to the previous two-dimensional case, the prior variance σ_f^2 as well as the length scale parameter for C_L are even larger. The noise variance σ_n^2 and the regularization parameter are the same as before. The Mach number length scale l_{Ma} decreased slightly in comparison to the two-dimensional case.

Table 3: hyperparameter choice for the Gaussian Process for the case $C_D = f(C_L, h, \mathrm{Ma})$

	σ_f	σ_n	l_{C_L}	l_h	l_{Ma}	ϵ
Initial parameters	0.0028	0.0005	0.25	500	0.05	1e−8
Tuned parameters	0.2211	1e−6	0.92	2.02e5	0.198	1e−6

We can observe that, in comparison to the initial parameter set, the sample locations of the tuned parameter set for all three input parameters C_L, h, Ma, are spread more in the design space, i.e. have greater distance to the location of interest (Tab.4). On the other hand, there are no samples at the same location. We can conclude, that the hyperparameters of the Gaussian Process have also great influence on the sample location of MIS. There is a tendency, that the hyperparameters can be manually adjusted, such that the locations of interest are sampled. Note, that a adjustment of the hyperparameters influences the prediction accuracy of the Gaussian Process. We assume, that the parameters from the parameter tuning, fit the data and reflects the functional relation, better than a manual choice.

Table 4: Comparison of MIS with different hyperparameter sets for the Gaussian Process with five added samples for the case $C_D = f(C_L, h, \mathrm{Ma})$

	Initial parameters			Tuned parameters		
#	C_L	h	Ma	C_L	h	Ma
1	0.250	9986 m	0.807	0.250	9987 m	0.806
2	0.253	10 435 m	0.810	0.216	5852 m	0.640
3	0.246	9481 m	0.796	0.270	2556 m	0.453
4	0.256	10 014 m	0.781	0.243	8854 m	0.757
5	0.242	9054 m	0.778	0.314	10 668 m	0.784

4.2 Effect of the integration criterion

The MIS strategy uses an integration criterion which shall be minimized, if another sample is added. For the three-dimensional case we integrate over a volume. The objective is to add a sample, which then in turn minimizes the *volume of biased uncertainty*. This criterion however *can* perform poorly on the boundaries of the design space. Instead of integrating the products, an option is to directly use the product. The acquisition function can be reformulated in the following way:

$$\mathbf{x_{n+1}} = \arg\max_{\mathbf{x}\in\mathcal{X}} \begin{cases} \text{kde}, & \text{if } X \in \emptyset, \\ \text{kde} \cdot \hat{s}_{+x}, & else \end{cases} \tag{14}$$

In this formulation, if no training points X are available, the location of the maximum of the KDE is used as the first sample. After this sample is sampled, for a next sample the product of KDE and posterior standard deviation $\hat{s}_{+x}$ is evaluated and again, the location where the maximum value is located is selected as a next sample. Note, that this procedure also reduces the uncertainty at already sampled locations, because $\hat{s}_{+x} = \sqrt{\text{diag}(\text{cov}(\mathbf{f}_*))}$ is trained to fit also each new observation X_*, which in turn reduces the product quantity at these locations.

In Tab.5, the results of the utilization of this direct product sampling strategy is depicted, again for the initial parameter set and for the tuned parameters. It can be noted, that the locations with the highest relative frequencies are robustly sampled within the first five samples. In this case, the initial parameter set returns sample locations, which are in a smaller range in the region of interest. For the tuned parameters, there are also samples, which are in greater distance from the location of interest. This is an indicator, that the length scale parameters for the altitude h and the Mach number Ma are chosen too large.

Table 5: Comparison of the direct product method with different hyperparameter sets for the Gaussian Process with five added samples for the case $C_D = f(C_L, h, \text{Ma})$

	Initial parameters			Tuned parameters		
#	C_L	h	Ma	C_L	h	Ma
1	0.253	10 668 m	0.820	0.253	10 668 m	0.820
2	0.248	10 083 m	0.820	0.248	9595 m	0.801
3	0.248	9693 m	0.806	0.194	4231 m	0.580
4	0.253	10 473 m	0.810	0.243	10 083 m	0.820
5	0.253	10 668 m	0.820	0.273	1597 m	0.413

4.3 Weighted Mission-Informed Sampling

Another method to shift samples into a potential region of interest is weighting or balancing. In the Bayesian optimization context, there exist the acquisition functions *Upper Confidence Bound* (UCB) or *Lower Confidence Bound* (LCB), where the ratio between exploitation and exploration can be set via a parameter we call $\kappa \in [0, 1]$. We can use this parameter to balance the exploitation of the KDE optima and the exploration

by means of reducing the predicted model variance. We can reformulate Eq.(13) to:

$$\mathbf{x_{n+1}} = \arg\min_{\mathbf{x}\in\mathcal{X}} \int_{\mathcal{X}} \{\mathrm{kde}(r)^{\kappa} \cdot \hat{s}_{+x}(r)^{(1-\kappa)}\}^2 dr \tag{15}$$

Note, that we use κ as an exponent, since KDE and $\hat{s}_{+x}$ are multiplied. Furthermore, we square the product, to get the initial product for $\kappa = 0.5$. Results on the one hand show the expected: If we put weight on the variance, i.e. $\kappa = 0.2$, the sampling procedure is largely sampling the design space, but without many samples in the region of interest. On the other hand, even with forcing exploitation of the KDE optima with a chosen $\kappa = 0.99$ the samples are in the region of interest (Tab.6), but still not sampling the maxima robustly, with both hyperparameter settings. This leads to two possible conclusions: 1. The integration criterion expects most of the information gain not at the optima but somewhere in the vicinity. 2. The chosen bandwidth strategy for KDE oversmooths.

Table 6: Comparison of the weighted MIS, with $\kappa = 0.99$ and different hyperparameter sets for the Gaussian Process with five added samples for the case $C_D = f(C_L, h, \mathrm{Ma})$

	Initial parameters			Tuned parameters		
#	C_L	h	Ma	C_L	h	Ma
1	0.253	10 668 m	0.814	0.250	10 136 m	0.811
2	0.247	9387 m	0.801	0.250	9500 m	0.790
3	0.253	10 624 m	0.813	0.206	9080 m	0.815
4	0.247	9057 m	0.796	0.445	10 668 m	0.82
5	0.250	10 175 m	0.811	0.192	6132 m	0.784

4.4 Adaptive KDE

As stated in Subsec.2.2, there are different methods to choose the bandwidth. One method is the adaptive kernel density estimation (AKDE) [17]. This method has several advantages in comparison to ROT or plugin-bandwidth selectors. It avoids oversmoothing of the optima, it can deal with boundary bias and can adapt the bandwidth *locally*. This KDE method is applied again for the initial parameter set and the tuned parameter set. Note, that the AKDE is automatically enlarging the design space. Which in turn results in an increased information gain at the optima, which were previously on the boundary of the design space. In Tab.7 the results for MIS using AKDE are depicted. The optima at FL330, Ma = 0.82 and FL350, Ma = 0.82, with the corresponding C_L values are both detected and sampled for the initial parameter set. Again we can observe, that for the tuned parameter set, the samples are spread in the design space, even not detecting the optima anymore.

Table 7: Comparison of the MIS using AKDE for different hyperparameter sets for the Gaussian Process with five added samples for the case $C_D = f(C_L, h, \mathrm{Ma})$

	Initial parameters			Tuned parameters		
#	C_L	h	Ma	C_L	h	Ma
1	0.249	10 193 m	0.820	0.251	10 366 m	0.819
2	0.253	10 681 m	0.820	0.216	6381 m	0.643
3	0.247	9859 m	0.819	0.277	2165 m	0.438
4	0.253	10 817 m	0.82	0.242	10 065 m	0.770
5	0.252	10 431 m	0.801	0.179	8531 m	0.819

5 CONCLUSIONS

In this contribution, a mission-informed sampling strategy (MIS) for Gaussian processes for use with a multi-fidelity model is presented. This sampling strategy uses information from a flight-mission simulation in order to propose a next sample location for an expensive CFD simulation for the drag coefficient prediction. In the application section, the MIS strategy is compared to the well-known Expected Improvement (EI) sampling strategy for a one-dimensional application case $C_D = f(C_L)$, two two-dimensional application cases $C_D = f(C_L, h)$, $C_D = f(C_L, \mathrm{Ma})$ and a three-dimensional application case $C_D = f(C_L, h, \mathrm{Ma})$. MIS samples in the region of interest, whereas EI samples of locations in the vicinity of the minimum drag. Therefore, the proposed samples from MIS are more *informative* when the computational resources are very limited. Another advantage is that MIS does not need function evaluations to generate a next sample candidate in contrast to EI. However, especially for the three-dimensional case and for cases where the highest relative frequencies are on the boundary of the design space, MIS samples in the vicinity but not directly at the optima. To investigate this issue, in the analysis subsection different enhancement strategies for MIS are presented, like modifying the hyperparameters of the Gaussian process, different adaptations to MIS, such as omitting the integral, or balancing the exploration and exploitation. Lastly, a different bandwidth selection and thus a different kernel density estimate – namely the adaptive kernel density estimation – is implemented and tested.

The results show that the hyperparameter choice has a great influence on the sample candidates. These hyperparameters in turn largely depend on the optimization and the optimization in turn depends on the data set quality. Directly using the product of KDE and the posterior standard deviation $\hat{s}_{+x}$ samples the optima reliably for the initial hyperparameter set. The weighted Mission-Informed Sampling strategy was tested with different weighting factors κ. The method sampled in the region of interest, but not directly the expected optima. The adaptive KDE (AKDE) also samples the optima for the initial hyperparameter set, but also deviates for the tuned hyperparameter set.

We can conclude that using the product of KDE and posterior standard deviation $\hat{s}_{+x}$ returns the most reliable results, also for different hyperparameter sets for the Gaussian process. Furthermore, there is a trade-off between a manual choice of hyperparameters shifting the samples to the region of interest and the prediction accuracy of the Gaussian Process constructed by these hyperparameters.

REFERENCES

[1] D. Raymer. *Aircraft design: a conceptual approach*. American Institute of Aeronautics and Astronautics, Inc., 2012.

[2] E. Torenbeek. *Synthesis of subsonic airplane design: an introduction to the preliminary design of subsonic general aviation and transport aircraft, with emphasis on layout, aerodynamic design, propulsion and performance*. Springer Science & Business Media, 2013.

[3] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

[4] A. Forrester, A. Sobester, and A. Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

[5] J. Mockus. The application of bayesian methods for seeking the extremum. *Towards global optimization*, 2:117, 1998.

[6] S. Ament, S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy. Unexpected improvements to expected improvement for bayesian optimization. *Advances in Neural Information Processing Systems*, 36:20577–20612, 2023.

[7] J. N. Fuhg, A. Fau, and U. Nackenhorst. State-of-the-art and comparative review of adaptive sampling methods for kriging. *Archives of Computational Methods in Engineering*, 28:2689–2747, 2021.

[8] A. Gramacki. *Nonparametric kernel density estimation and its computational aspects*, volume 37. Springer, 2018.

[9] E. Cramer and U. Kamps. *Grundlagen der Wahrscheinlichkeitsrechnung und Statistik: Eine Einführung für Studierende der Informatik, der Ingenieur-und Wirtschaftswissenschaften*. Springer-Verlag, 2020.

[10] B. W. Silverman. *Density estimation for statistics and data analysis*. Routledge, 2018.

[11] D. W. Scott. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015.

[12] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17: 261–272, 2020. doi:10.1038/s41592-019-0686-2.

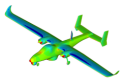
[13] M. Zimmnau, D. Stenger, A. Uhl, and E. Stumpf. Aerodynamic variable fidelity surrogate model for the conceptual aircraft design with UNICADO. In *Deutscher Luft- und Raumfahrtkongress 2022*, Dresden, 2022. Deutsche Gesellschaft für Luft- und Raumfahrt - Lilienthal-Oberth e.V.

[14] M. Zimmnau, A. Lindicke, A. Uhl, and E. Stumpf. Multidimensional aerodynamic variable fidelity surrogate model for the conceptual aircraft design with UNICADO. In *Deutscher Luft- und Raumfahrtkongress 2023*, Stuttgart, 2023. Deutsche Gesellschaft für Luft- und Raumfahrt - Lilienthal-Oberth e.V.

[15] M. Zimmnau, S. Cremer, A. Uhl, and E. Stumpf. Towards the integration of an aerodynamic variable fidelity surrogate model into a conceptual aircraft design environment. In *AIAA SCITECH 2025 Forum*, page 0003, 2025.

[16] F. Schültke and E. Stumpf. UNICADO : Development and Establishment of a University Conceptual Aircraft Design Environment. In *Deutscher Luft- und Raumfahrtkongress 2021*, Bremen, 2021. Deutsche Gesellschaft für Luft- und Raumfahrt - Lilienthal-Oberth e.V., RWTH Aachen University. doi:10.18154/RWTH-2021-09317.

[17] Z. Botev, J. Grotowski, and D. Kroese. Kernel density estimation via diffusion. *Ann. Statist.*, 38(1):2916–2957, 2010.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

HIGH-FIDELITY NUMERICAL INVESTIGATIONS OF FOLDING WINGTIPS ON HIGH ASPECT RATIO WINGS

Pranesh Chandrasekaran[1*], Emre Güler[1] and Melike Nikbay[1]

1: AeroMDO Lab, Faculty of Aeronautics and Astronautics
Istanbul Technical University
İTÜ Ayazaga Campus, 34469, Maslak/Istanbul
chandrasekaran@itu.edu.tr, gulere22@itu.edu.tr, nikbay@itu.edu.tr

Abstract. *The growing demand for sustainable aviation solutions has led to an increased focus on innovative wing designs. Of interest to us in this work is the high aspect ratio (HAR) wings, which can offer a notable reduction in lift-induced drag. One of the inherent difficulties in designing HAR wings is satisfying the dimensional constraints related to the runway and taxiway separation, which places severe demands on the maximum allowable wingspan. In this regard, the concept of folding wingtips offers a unique solution to mitigate the ground handling difficulties faced by the HAR wings. The aim of the present work is to investigate the effect of folding wingtips on the aerodynamic performance of representative HAR wings. Most previous studies on HAR wings with folding wingtips have employed low-fidelity tools to demonstrate the performance benefits associated with this concept. In contrast, the present study employs a high-fidelity open-source CFD tool, SU2, for computing and analyzing the flow field around these wing geometries. A rectangular wing with an aspect ratio of 17.93 is considered, and the impact of various geometrical parameters such as fold angle and location of the hinge line from the wing root are studied. The configurations examined include wingtip fold angles of 15°, 30°, and 45°, as well as the hinge locations at around 75% and 85% of the wingspan. Results revealed that the effect of fold angle on aerodynamic parameters is relatively more pronounced at high angles of attack, while its effect was found to be negligible at low angles of attack.*

Keywords: aerodynamics, folding wingtips, computational fluid dynamics, high aspect ratio wings, hinge location, fold angle

1 INTRODUCTION

The impacts of climate change are becoming increasingly evident in our daily lives. The rapid expansion of air travel, and the continuous increase in air traffic, has significantly contributed to the rising greenhouse gas emissions and air pollutants. Considering the fact that the global aviation sector is responsible for $\approx 2\%$ of all man-made carbon dioxide (CO2) emissions [1], together with the ever rising demand for air travel, there exists a strong need to come up with environmentally friendly aircraft designs with improved fuel efficiency and reduced emissions. Driven by this need, the aviation industry has focused its efforts on developing new design concepts for the next generation of commercial civil transport aircraft. Some of these efforts are aimed at transitioning to alternative fuels and propulsion systems [2], while the remainder is focused on improving the flight performance by exploring novel wing configurations such as the blended wing body, strut-braced wing, and truss-based wing [3–5]. Furthermore, cutting-edge technologies in the fabrication of lightweight composite structures can play an important role in reducing the environmental impact.

Even though the initial drive for reduced fuel consumption in aviation sector was motivated by economic incentives, the current push, however, focuses on achieving sustainability through lower emission and enhanced environmental responsibility in the aeronautical industry. Various initiatives like ICAO CORSIA, Horizon 2020, CLEANSKY2, and Flight Path 2050 have set ambitious targets for the reductions of CO2, NOx, and noise by the year 2050 [6–9]. In addition, the National Aeronautics and Space Administration (NASA) and the European Commission have consistently supported these initiatives by funding projects like SUGAR [10], AGILE [11], RHEA [12], and U-HARWARD [13]. One of the critical aspects of these sustainability efforts is the reduction of lift-induced drag, which can account for nearly 40% of a typical civil transport aircraft's total drag in cruise condition [14]. Hence, any effort in reducing the induced drag can result in an improved lift-to-drag ratio and reduced fuel consumption. This in turn will enhance the aircraft's capabilities in terms of the flight range, thus positively impacting the environment. Of particular interest in this context is the design of High Aspect Ratio (HAR) wings, which can offer a significant reduction in lift-induced drag.

Traditionally, HAR wings have been successfully used in unmanned air vehicles (UAV), specifically in high-altitude long-endurance (HALE) aircraft. Recently, there has been an increase in interest regarding their potential applications within the civil transport sector. While increasing the aspect ratio (AR) can often be achieved by increasing the wingspan, this approach presents serious challenges in the form of weight penalties and increased bending moment at the wing root. The design of HAR wings also brings the challenge of increased flexibility and significant geometric deformations under aerodynamic loading. This, in turn, may lead to aeroelastic instabilities in flight and could even result in structural failure of the wing. Additionally, the need to comply with dimensional constraints related to the runway and taxiway separation places severe demand on the maximum allowable wing span of HAR wing designs. In this regard, the concept of folding wingtips offers a unique solution to mitigate the ground handling difficulties faced by HAR wings. More recently, the Boeing 777X has achieved a 7 m extension in wingspan by successfully implementing the folding wingtip concept [15]. A configuration of this nature allows the aircraft to reduce its wingspan while taxiing, parking, or navigating on the ground, thus making them compatible with the existing airport infrastructures.

Some of the recent studies have demonstrated that folding wingtips can also be used as a means to facilitate load alleviation and improve roll performance [16–20]. Most of these studies, however, have relied on low-fidelity tools to estimate the aerodynamic properties. There exists a clear need for more comprehensive investigations into the detailed flow field characteristics associated with this configuration. The present work aims to address this gap by employing a high-fidelity open-source CFD tool, SU2, to compute and analyze the complicated flow features of HAR wings with folding wingtips.

2 DETAILS OF THE GEOMETRICAL MODEL

The test case geometry used in the present study is a rectangular wing with a chord length of 0.15 m and a semi-span of 1.345 m. This configuration results in an aspect ratio of 17.9, indicative of its high-aspect ratio design. The wing is modeled with a constant NACA 0015 cross-sectional profile, meaning that the airfoil shape remains the same at every spanwise location. This configuration will be referred to as the baseline or clean wing case. In addition to this baseline case, wings with folding wingtips are also studied to investigate the effect of various parameters, such as fold angle and location of hinge axis, on the separation characteristics and the aerodynamic properties of these high-aspect ratio wings. Two different wing geometries, with wingtip hinge axis located at 1 m and 1.14325 m from the wing root are considered. The total wetted area of the wingtips in these configurations represents 25% and 15% of the total area, respectively. For both configurations, three different fold angles (15°, 30° and 45°) are investigated. The choice of the baseline geometry is motivated by the availability of experimental data for high-aspect ratio wings, particularly those featuring folding wingtips. Although the experimental model of Cheung et al. [19] involved a highly flexible wing (see Fig.1), in the present work the wing is modeled as a rigid body. This is because the chosen wing was experimentally tested in a wind tunnel at speed of 12 to 18 m/s, corresponding to a Reynolds number (Re) of 123,226 to 184, 838, essentially placing them in the transitional low Re regime. At such Reynolds numbers, the flow is highly prone to flow separation, even at small adverse pressure gradients at low angles of attack, and is extremely sensitive to even small changes in geometry, Reynolds number, angle of attack, freestream turbulence intensity, and surface roughness. Hence, it was decided to first study the rigid wing case before examining the effects of wing flexibility on these laminar separation bubbles and their corresponding impact on aerodynamic properties.

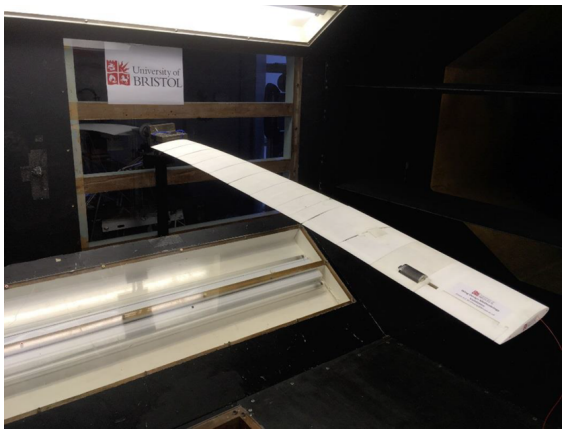
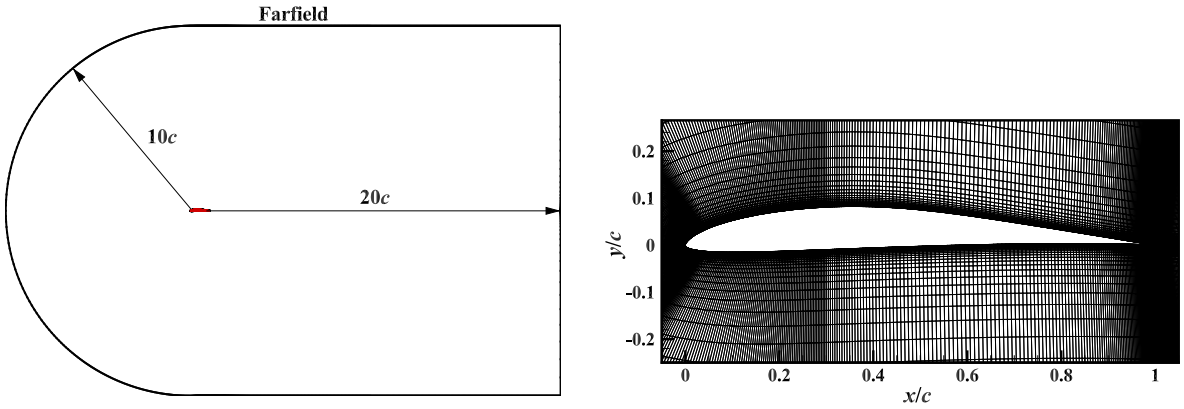


Figure 1: Experimental model investigated by Cheung et al. [19].

3 COMPUTATIONAL PROCEDURE

In order to establish the reliability of the computational procedure, the results from the present work must be compared with those available in the literature. For this purpose, the well-known low Re Eppler 387 airfoil operating at a Reynolds number of 200,000 and at a freestream turbulence intensity (Tu) of 0.1% is considered. The dimensions of the computational domain and the boundary conditions applied at the domain boundaries are shown in Fig.2a, while the mesh close to the airfoil surface is shown in Fig.2b. Enough care was taken to cluster the grids in the near wall region to accurately capture the boundary layer features.



(a) (b)

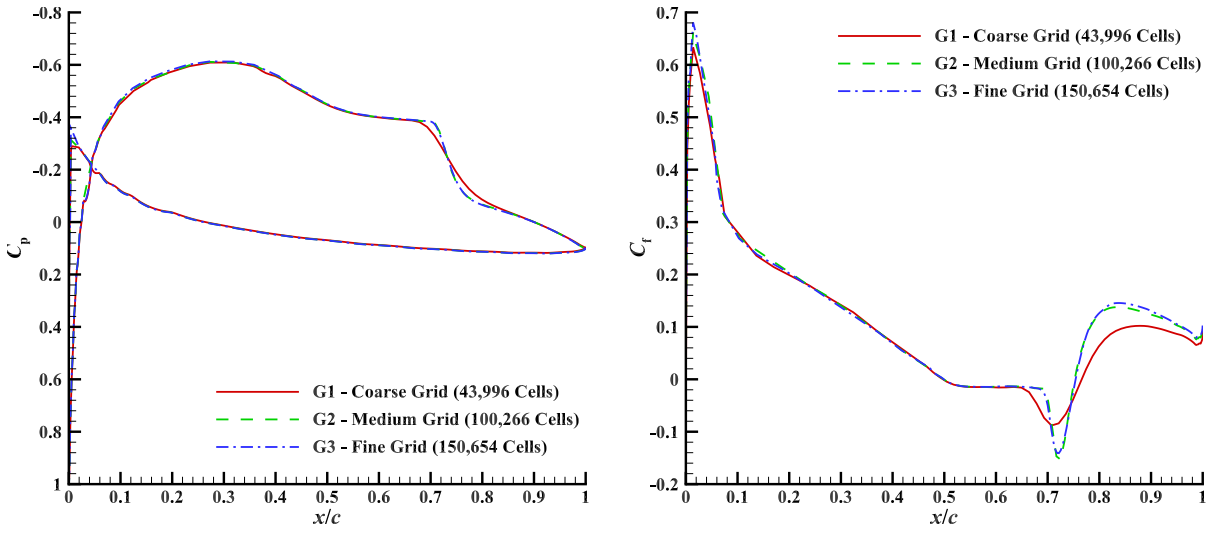
Figure 2: (a) Details of the computational domain and (b) Mesh close to the airfoil surface.

The sensitivity of the aerodynamic parameters to the number of computational cells is studied by performing a careful grid independence study. Three grids namely G1, G2, and G3 with increasing mesh sizes were generated. The lift and drag coefficients obtained using these three grids are listed in Tab.1, while the C_p and C_f distributions are plotted in Fig.3a and 3b, respectively. As can be seen, increasing the number of computational cells from 100,266 to 150,654 did not result in significant changes in both global and local aerodynamic properties. Hence, the medium grid with 100,266 cells was chosen for all further two-dimensional airfoil analysis.

Table 1: Grid independence study of Eppler 387 airfoil at $\alpha = 0°$, Re = 200,000 and Tu of 0.1%.

Grid	No. of Cells	C_l	Absolute Difference	C_d	Absolute Difference
G1	43,996	0.3666	0.0020	0.0103	5.68E-05
G2	100,266	0.3686	0.0006	0.0103	4.66E-06
G3	150,654	0.3693		0.0103	

Having established the grid independent nature of the results, the next step is to validate them against the experimental [21] and computational data [22] available in the literature. For this, computations were carried out for a wide range of angles of attack



(a) (b)

Figure 3: (a) C_p distribution over the Eppler 387 airfoil and (b) C_f distribution over the upper surface of Eppler 387 airfoil; $\alpha = 0°$, $Re = 200{,}000$ and $Tu = 0.1\%$.

(α) between 0° to 12°. In the present work, a finite volume-based open-source flow solver, SU2, is utilized to solve the governing equations of fluid flow, essentially the incompressible Reynolds-Averaged Navier-Stokes (RANS) equations. The implementation of the γ-Re_θ transitional model within SU2 is detailed in one of the recent papers by Rausa et al. [22]. While the use of SU2 has gained popularity over the last decade, the transitional model's implementation is relatively new. Therefore, it is essential to compare the results obtained from SU2 with those from the well-established commercial software ANSYS FLUENT. Simulations were conducted using both FLUENT and SU2, and the results are presented in Fig.4.

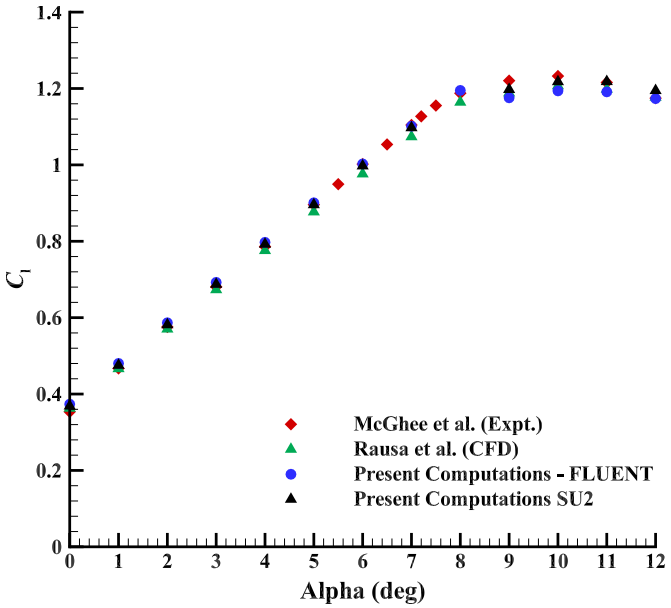


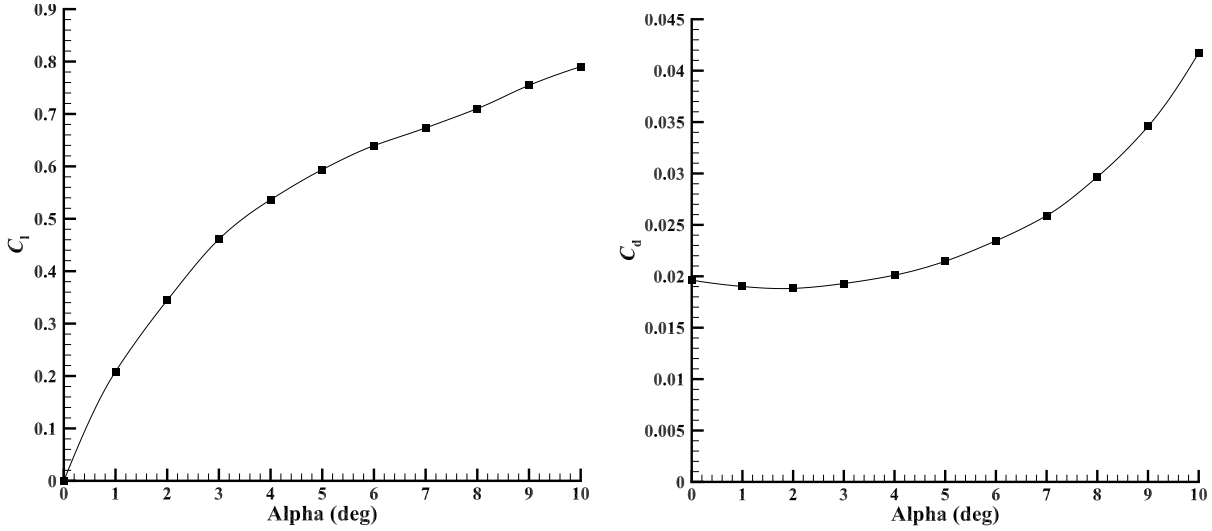
Figure 4: Validation of computational procedure; Eppler 387 airfoil at $Re = 200{,}000$ and $Tu = 0.1\%$.

As can be seen, the results from SU2 are in excellent agreement with both the experimental and computational results available in the literature, as well as those obtained using ANSYS FLUENT across the entire range of angles of attack considered. Hence, for all further two-dimensional computations, only SU2 was used.

4 RESULTS AND DISCUSSION

4.1 Flow Around NACA 0015 Airfoil

In this section, the flow over NACA 0015 airfoil is computed at a Reynolds number of 123,226 and $Tu = 0.08\%$. A structured grid with C-H topology, with the same number of computational cells as that of Grid G2 discussed in the previous section, was generated. The chosen grid had around 600 points on the airfoil surface, 700 points in the wake and 50 points in the radial direction. Simulations were performed for various angles of attack and the corresponding lift and drag coefficients are presented in Fig.5.



(a) (b)

Figure 5: Lift and drag curves of NACA 0015 airfoil at $Re = 123{,}226$ and $Tu = 0.08\%$.

It may be noted that the lift curve does not exhibit the linear characteristics typically observed at high Reynolds numbers. Instead, the lift curve is non-linear, with an apparent change in slope at moderate angles of attack. Such a phenomenon is attributed to the boundary layer characteristics and the complex flow features, such as the formation of laminar separation bubbles, associated with the transitional low Re regime. A closer look at the axial-velocity contours revealed that flow separation occurs at an angle of attack as low as 0° and on both upper and lower surfaces; see Fig.6. The separated boundary layer eventually reattaches, leading to the formation of laminar separation bubble (LSB). As the angle of attack increases, the LSB shrinks in size and shifts closer to the leading edge on the upper surface, while on the lower surface the presence of LSB ceases to exist at $\alpha > 4^\circ$. In other words, the flow remains completely attached to the lower surface at high angles of attack. It is at this angle of attack ($\alpha = 4^\circ$) we see a notable change in the lift curve slope. Having computed the flow around a two-dimensional airfoil, the next step is to simulate the flow around the three-dimensional wing geometries.

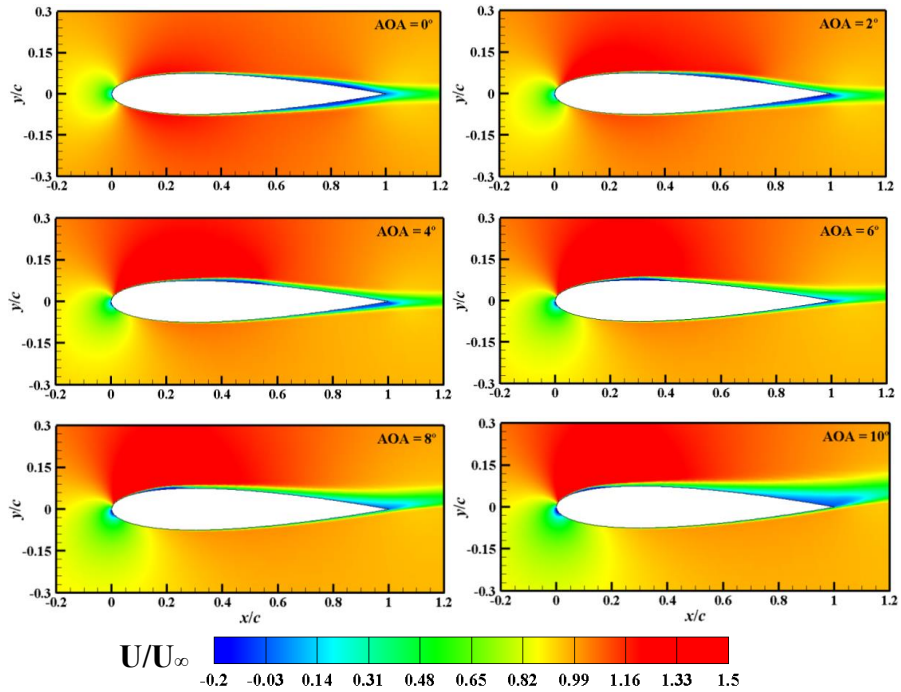


Figure 6: Contours of axial velocity around the NACA 0015 airfoil at $Re = 123{,}226$ and at $Tu = 0.08\%$.

4.2 Flow Around the NACA 0015 Wing

In this section, the flow field around a three-dimensional rectangular wing with a NACA 0015 cross-sectional profile and an AR of 17.9 is computed at various angles of attack. The wingtip fold angle (FA) is set to 0°, representing the clean wing case. A structured three-dimensional grid is utilized in the present study; the mesh on the wing surface and on the symmetry plane is shown in Fig.6. Like in the two-dimensional case, the computational domain extends $10c$ ahead of the wing leading edge and $20c$ behind it;

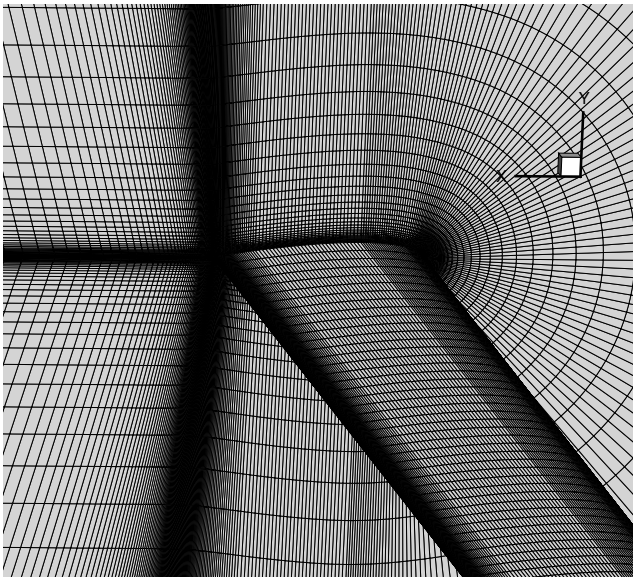


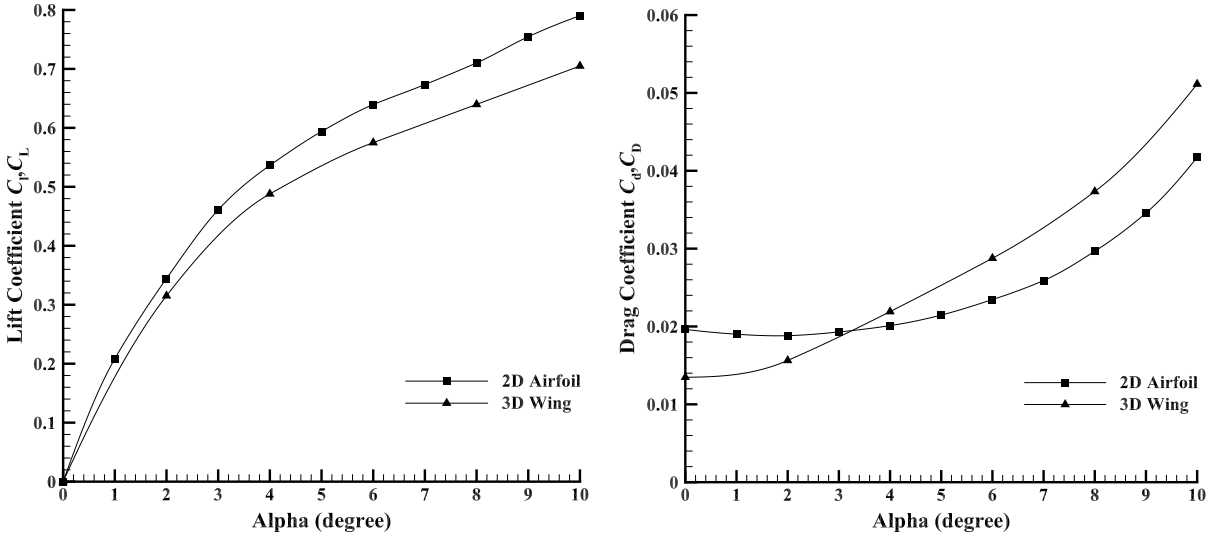
Figure 7: Mesh on the wing surface.

c here represents the root chord. The spanwise distance from the wingtip to the domain boundary is equal to the wing span. A no-slip boundary condition is applied to the wing surface, while a symmetry boundary condition is applied at the domain boundary adjacent to the wing root. The farfield boundary condition was applied elsewhere. Three different grids namely G1 (≈1.6 million), G2 (≈3.6 million), and G3 (≈6.2 million) with increasing mesh size were generated. Results revealed that the medium grid G2 with around 3.6 million cells had reasonably low discretization error (see Tab.2), and was chosen for all further computations.

Table 2: Grid independence study of NACA 0015 wing at $\alpha = 6°$, Re = 123,226 and Tu of 0.08%.

Grid	No. of Cells	C_l	Absolute Difference	C_d	Absolute Difference
G1	1,676,412	0.5300	0.0449	0.0255	0.0032
G2	3,661,290	0.5749	0.0188	0.0288	0.0009
G3	6,284,964	0.5937		0.0297	

Simulations were performed over an angle of attack range from 0° to 10°, with a $\Delta\alpha$ of 2°. The variation of lift and drag coefficients as a function of angle of attack is presented in Fig.7. The lift and drag curves of the two-dimensional airfoil case are also included for the sake of comparison. The influence of three-dimensional effects on the aerodynamic performance is evident in the results. As expected, the lift coefficients for the three-dimensional wing are lower than those for the airfoil case, which can be attributed to the effect of wing tip vortices. In terms of drag, the values for the three-dimensional wing are less than those of airfoil at low angles of attack. However, as angle of attack increases, the three-dimensional drag coefficient is observed to be greater than that of the airfoil case.



(a) (b)

Figure 8: Lift and drag curves of clean wing at Re = 123,226 and Tu = 0.08%; 2-D results are also plotted for comparison.

We resort to the contours of skin-friction distribution on the wing surface to investigate the flow features associated with the clean wing configuration. Figure 9 shows the contours of C_{fx} on the upper surface of the wing for six different angles of attack. Like in the two-dimensional case, the flow is observed to be separated even at $\alpha = 0°$. It may be observed that the LSB shrinks in size and moves towards the leading edge, with increasing angles of attack. At high angles of attack ($\alpha \geq 8°$), apart from the presence of LSB, a trailing edge recirculation region is also observed.

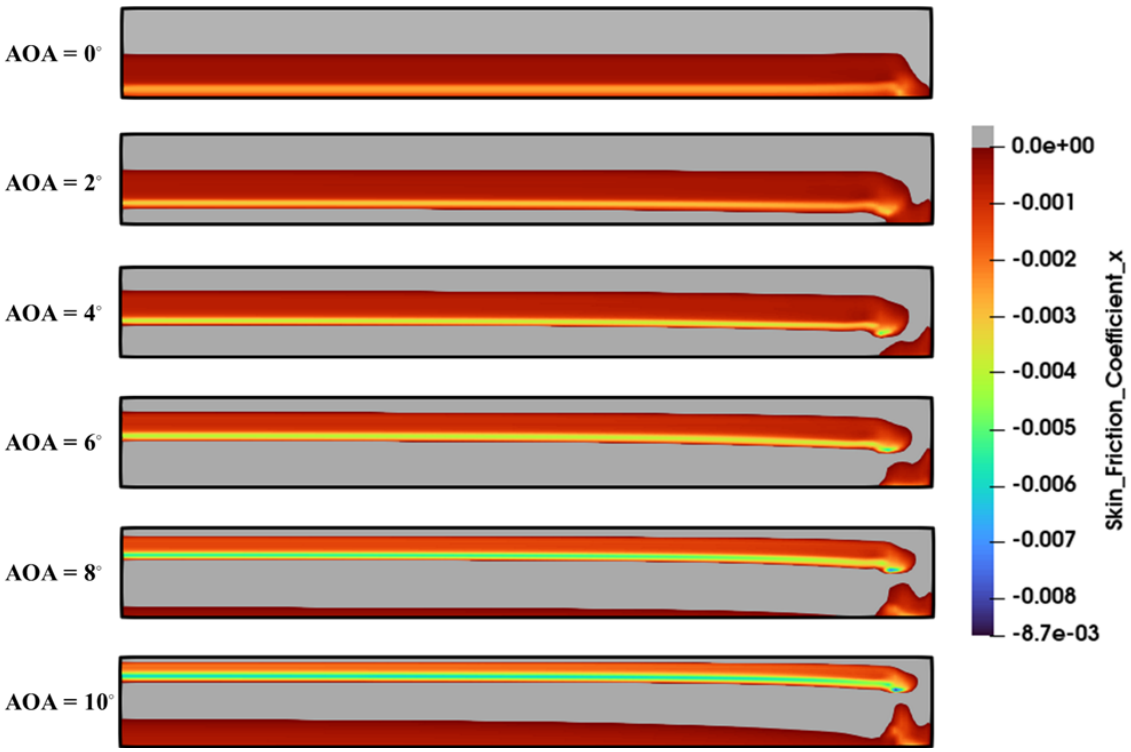


Figure 9: Contours of skin-friction distribution, highlighting the separated region over the upper surface of the clean wing. Wing root is on the left, wingtip is on the right and the flow direction is from the top.

4.3 Effect of Fold Angle on the Aerodynamic Coefficients

In order to investigate the effect of wingtip fold angle on the aerodynamic properties, simulations were carried out on wings with three different fold angles 15°, 30°, and 45° (see Fig.10). In accordance to the experimental model investigated by Cheung et al. [19], the hinge axis was positioned at 1 m from the wing root.



Figure 10: Geometries Studied. Viewed from behind.

Simulations were performed across an angle of attack range of 0° to 10°. The variation of lift, drag and lift-to-drag ratio as a function of angle of attack for these configurations are presented in Fig.11. The results from clean wing are also plotted for comparison. General observations from this study indicate that, for any given angle of attack, an increase in the wingtip fold angle resulted in a monotonic and marginal decrease in lift coefficient. In contrast, the drag coefficients and the lift-to drag ratios do not exhibit a consistent trend across the tested configurations. Increasing the fold angle resulted in a reduction in drag at low to moderate angles of attack compared to the clean wing case. However, for the highest angle of attack considered in this study, the clean wing configuration exhibited the lowest drag value. Comparison of skin-friction contours on the upper surface of the wing revealed that the flow over the main wing remains relatively unaffected by the changes in the fold angle. However, in the wingtip region, the transition and reattachment locations

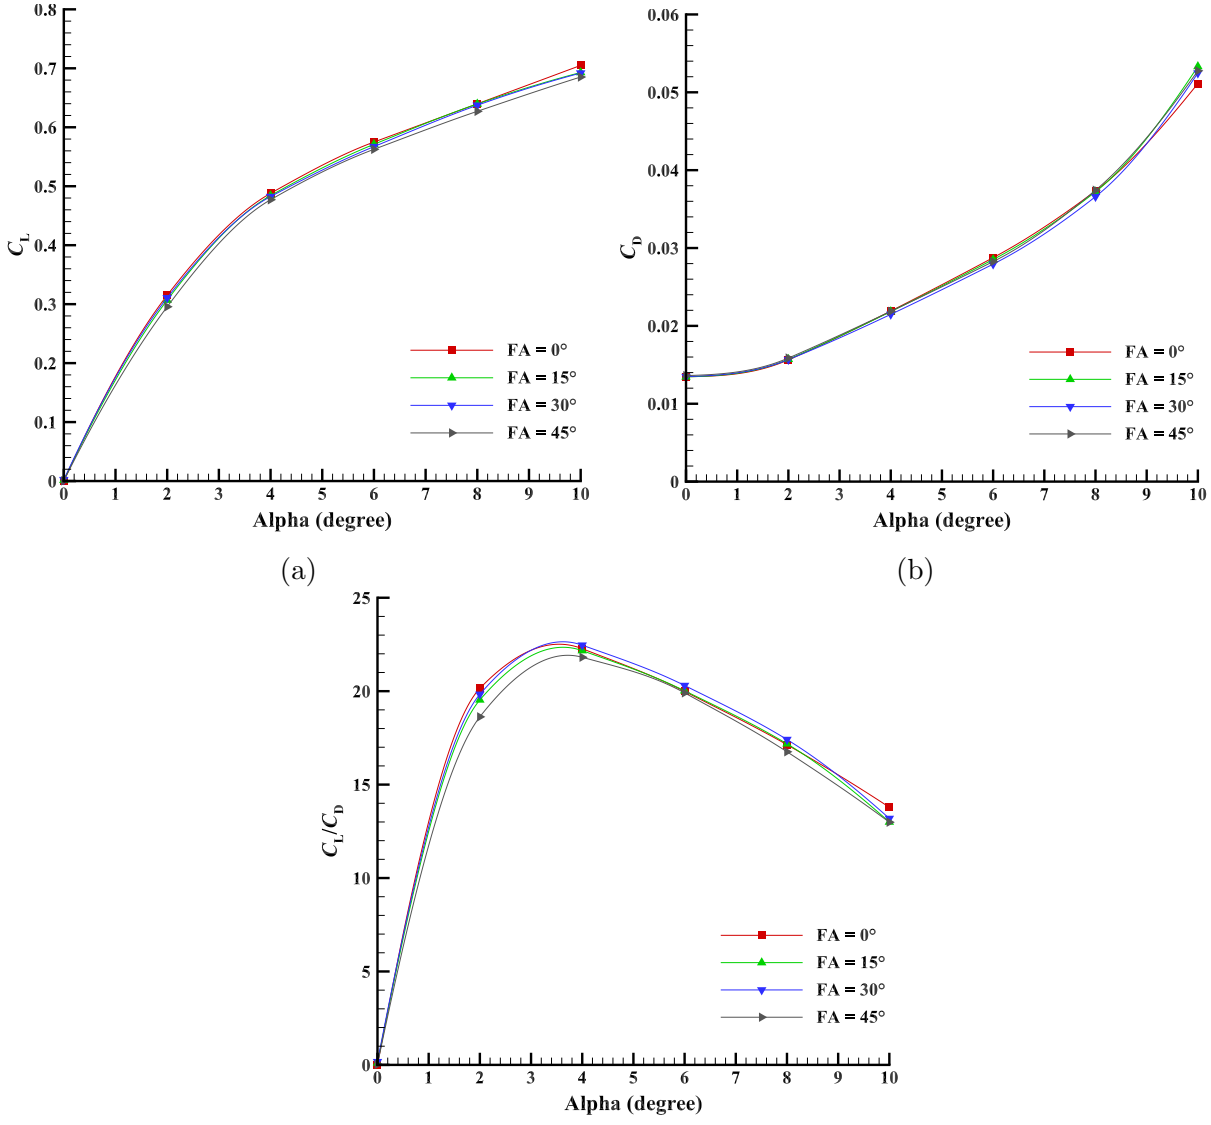


Figure 11: Effect of fold angle on the global aerodynamic properties. (a) lift coefficient (b) drag coefficient and (c) lift-to-drag ratio at Re = 123,226 and Tu = 0.08%.

are found to be delayed, while the separation points remains relatively unaffected. This indicates that the wingtip region experiences a relatively longer LSB as compared to the main wing region.

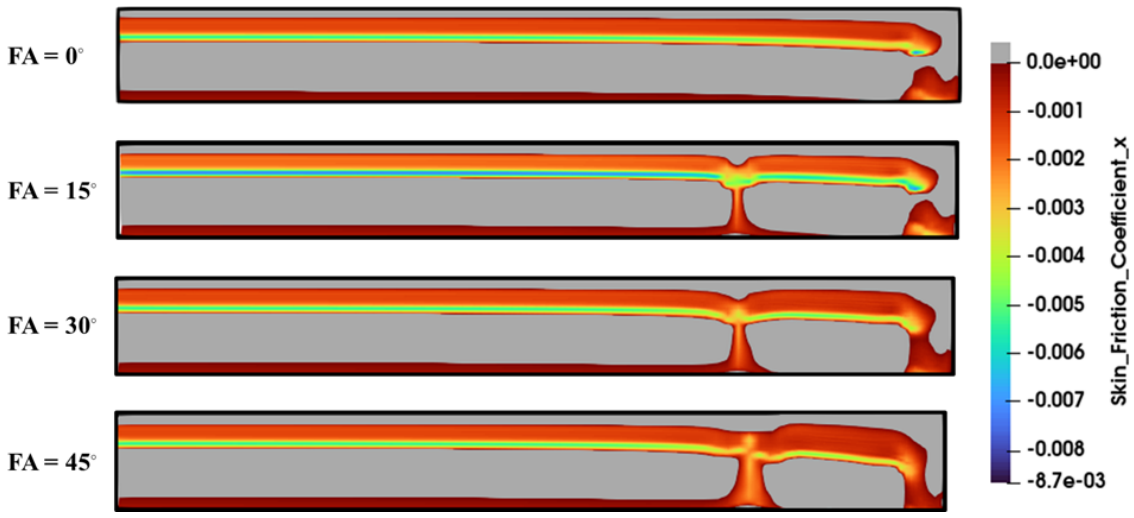


Figure 12: Effect of fold angle on the separated region over the upper surface; $\alpha = 8°$, $Re = 123{,}226$ and $Tu = 0.08\%$.

4.4 Effect of Hinge Location on the Aerodynamic Coefficients

In order to investigate the effect of hinge location on the aerodynamic parameters, simulations were now carried out on wings whose hinge locations are positioned at 1.14325 m (See Fig.13). In such a configuration, the area of the wingtip constitute 15% of the original area of the clean wing case. Again, three different fold angles (15°, 30°, and 45°) were considered. The variation of lift, drag, and lift-to-drag ratio as a function of angle of attack is presented in Fig.14. Unlike the previous configuration, the effect of fold angle is more pronounced only at high angles of attack ($\alpha \geq 8°$) and is negligible at low angles of attack. For the highest angle of attack considered, the clean wing case had the lowest drag value, whereas the wing with fold angle 45 degree had the highest value of drag. The size of LSB observed in the wingtip region is relatively bigger than the ones observed over the main wing (See Fig. 15). These observations align with the results from the previous configuration.



Figure 13: Geometries Studied. Viewed from behind.

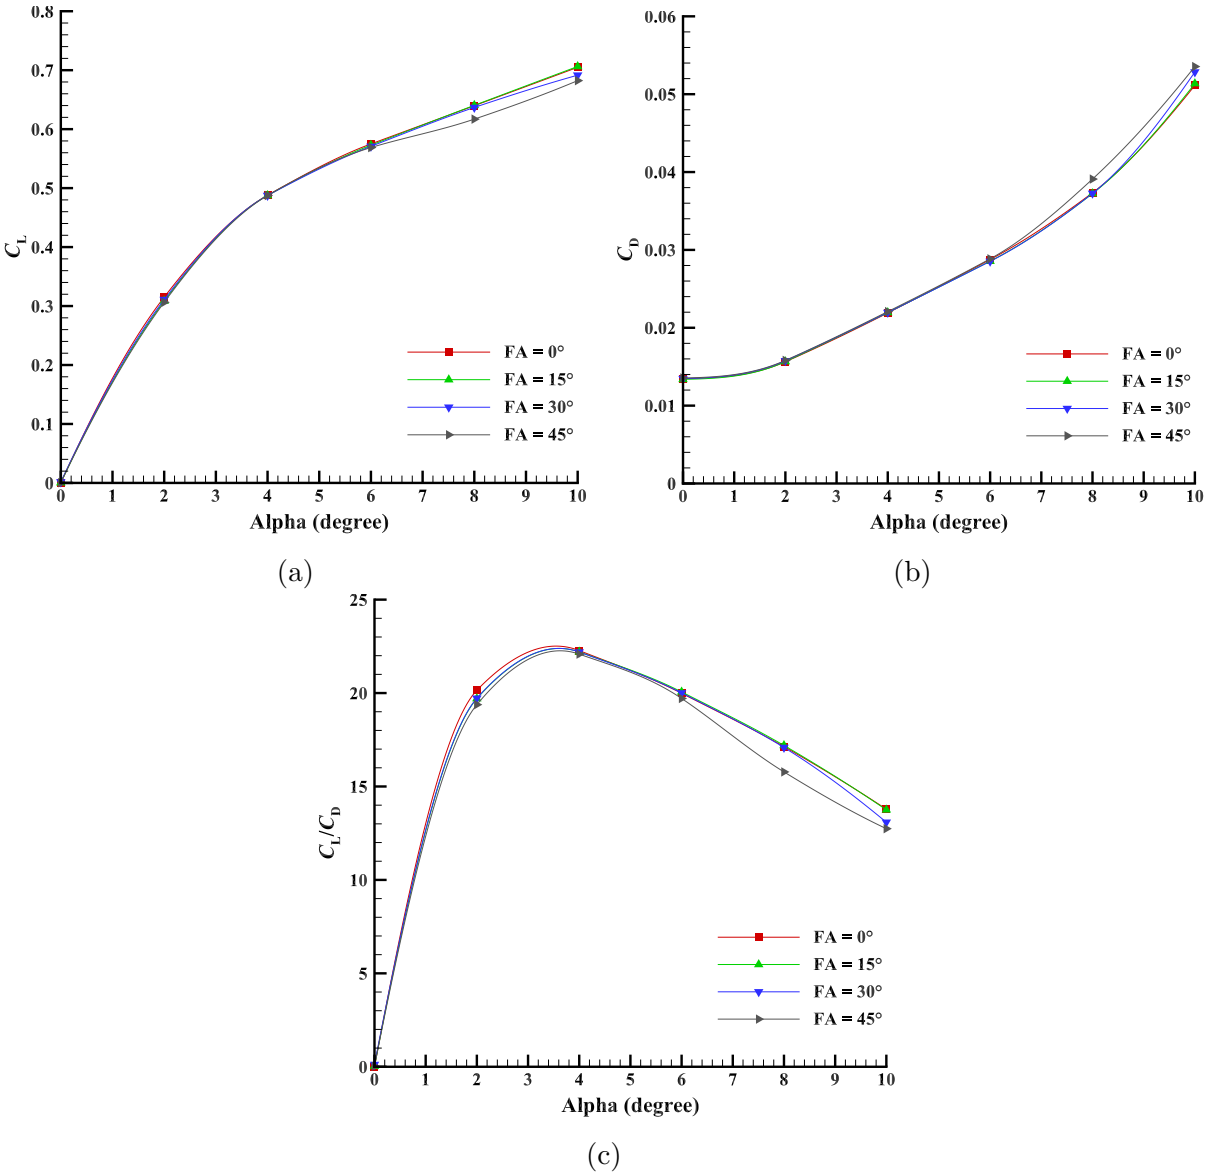


Figure 14: Effect of fold angle on the global aerodynamic properties; hinge is located at 1.14325 m from wing root. (a) lift coefficient (b) drag coefficient and (c) lift-to-drag ratio at Re = 123,226 and Tu = 0.08%.

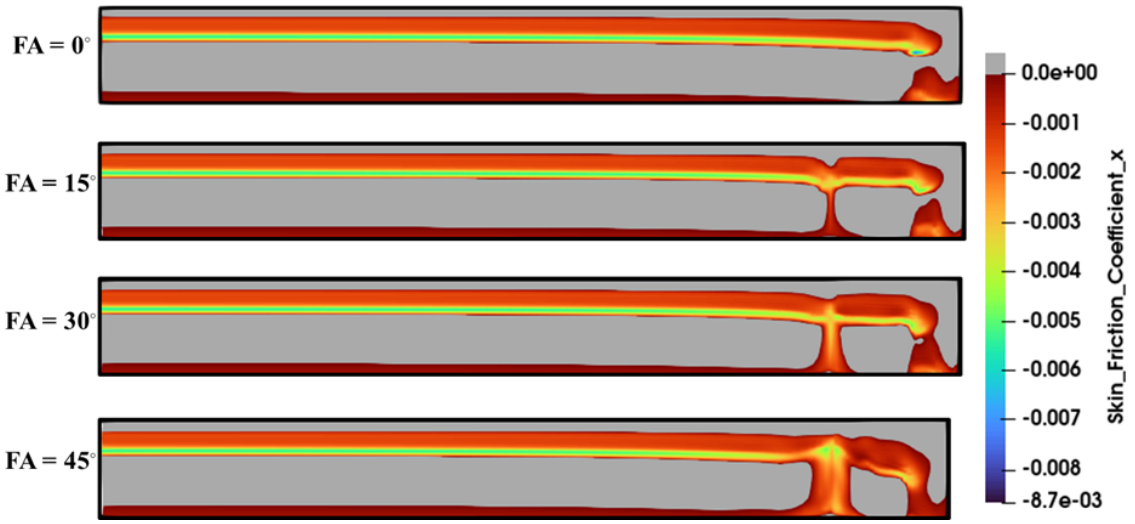


Figure 15: Effect of fold angle on the separated region over the upper surface for the configuration where the wingtip constitutes 15% of the main wing; $\alpha = 8°$, $Re =$ 123,226 and $Tu = 0.08\%$.

5 CONCLUSIONS AND FUTURE WORK

This study provides critical insigths into the folding wingtip concept, focusing on the effects of wingtip fold angle and hinge location on aerodynamic coefficients. Two configurations were computationally investigated, with the folding wingtip constituting, 15% and 25% of the main wing, and three different fold angles were examined: 15°, 30°, and 45°. The results from the present work, align well with the experimental and computational data available in the literature, demonstrating the reliability of our high-fidelity CFD simulations, using an open-source software-SU2, in analyzing the folding wingtip concept at transitional low Re regime. For the configuration with a 15% wingtip, the fold angle's impact was significant only at high angles of attack. In contrast, for the configuration with a 25% wingtip, the effect is relatively more pronounced across the entire range of angles considered. A common observation across both configurations was that increasing the fold angle led to a slight decrease in lift and drag coefficients at low angles of attack. However, at high angles of attack, the clean wing demonstrated lower drag compared to the folded wingtip configuration.

Building on the established reliability of the computational procedure for steady-state flow around a rigid wing at transitional low Reynolds numbers, several research directions have been identified. Future investigations will include both static and dynamic fluid-structure interactions of the folding wingtip concept, assessing its potential as a load alleviation device; the impact of wing flexibility will also be considered. Additionally, the concept of flared-folding wingtips, where the hinge axis is oriented at an angle to the flow direction, will be investigated. While the present study focused on representative high-aspect ratio wing, the analysis will be extended to configurations that closely resemble the existing civil transport aircraft. Such a progression will further enhance the understanding of the practical applications of this concept and its implications on the future aircraft designs.

ACKNOWLEDGEMENTS

All authors would like to express their gratitude to TÜBİTAK for the research grant provided under the 222M019 TÜBİTAK 1001 project titled as "Development of Environmentally Friendly Aircraft With High Aspect Ratio Wings via Multi-Fidelity Surrogate Modeling Based Aeroelastic Optimization." The computing resources used in this work was provided by the National Center for High Performance Computing of Turkey (UHeM) under grant number 5020292024.

REFERENCES

[1] D. Eisenhut, N. Moebs, E. Windels, D. Bergmann, I. Geiß, R. Reis, and A. Strohmayer. Aircraft requirements for sustainable regional aviation. *Aerospace*, 8 (3):61, 2021.

[2] Anon. Aircraft technology roadmap to 2050. Technical report, IATA, 2019. Technical report.

[3] P. Okonkwo and H. Smith. Review of evolving trends in blended wing body aircraft design. *Progress in Aerospace Sciences*, 82:1–23, 2016.

[4] C. K. Droney, A. J. Sclafani, N. A. Harrison, A. D. Grasch, and M. D. Beyar. Subsonic ultra green aircraft research: Phase iii–Mach 0.75 transonic truss-braced wing design. Technical report, Boeing Research and Technology, Huntington Beach, California, 2020.

[5] M. Sohst, J. L. do Vale, F. Afonso, and A. Suleman. Optimization and comparison of strut-braced and high aspect ratio wing aircraft configurations including flutter analysis with geometric non-linearities. *Aerospace Science and Technology*, 124:107531, 2022.

[6] F. Fichert, P. Forsyth, and H.-M. Niemeier. Aviation and climate change. *Economic Perspectives on Greenhouse Gas Reduction Policies*, 2020.

[7] E. Marzi, M. Morini, and A. Gambarotta. Analysis of the status of research and innovation actions on electrofuels under horizon 2020. *Energies*, 15(2):618, 2022.

[8] S. Ameduri, A. Concilio, I. Dimino, R. Pecora, and S. Ricci. AIRGREEN2-clean sky 2 programme: Adaptive wing technology maturation, challenges and perspectives. *Smart Materials, Adaptive Structures and Intelligent Systems*, 51944:V001T04A023, 2018.

[9] R. M. Arnaldo Valdes, S. Burmaoglu, V. Tucci, L. M. Braga da Costa Campos, L. Mattera, and V. F. Gomez Comendador. Flight path 2050 and acare goals for maintaining and extending industrial leadership in aviation: a map of the aviation technology space. *Sustainability*, 11(7):2065, 2019.

[10] M. K. Bradley and C. K. Droney. Subsonic ultra green aircraft research. Technical report, Technical report, NASA, 2011.

[11] P. D. Ciampa, P. S. Prakasha, F. Torrigiani, J.-N. Walther, T. Lefebvre, N. Bartoli, H. Timmermans, P. Della Vecchia, L. Stingo, D. Rajpal, et al. Streamlining cross-organizational aircraft development: results from the AGILE project. In *AIAA Aviation 2019 Forum*, page 3454, 2019.

[12] Y. Ma, S. Karpuk, and A. Elham. Conceptual design and comparative study of strut-braced wing and twin-fuselage aircraft configurations with ultra-high aspect ratio wings. *Aerospace Science and Technology*, 121:107395, 2022.

[13] S. Ricci, N. Paletta, S. Defoort, E. Benard, J. E. Cooper, and P. Barabinot. U-HARWARD: a CS2 EU funded project aiming at the design of ultra high aspect ratio wings aircraft. In *AIAA Scitech 2022 Forum*, page 0168, 2022.

[14] N. Beck, T. Landa, A. Seitz, L. Boermans, Y. Liu, and R. Radespiel. Drag reduction by laminar flow control. *Energies*, 11(1):252, 2018.

[15] A. Castrichini, V. H. Siddaramaiah, D. E. Calderon, J. E. Cooper, T. Wilson, and Y. Lemmens. Preliminary investigation of use of flexible folding wing tips for static and dynamic load alleviation. *The Aeronautical Journal*, 121(1235):73–94, 2017.

[16] A. Castrichini, V. Hodigere Siddaramaiah, D. Calderon, J. E. Cooper, T. Wilson, and Y. Lemmens. Nonlinear folding wing tips for gust loads alleviation. *Journal of Aircraft*, 53(5):1391–1399, 2016.

[17] A. Castrichini. *Parametric assessment of a folding wing-tip device for aircraft loads alleviation*. PhD thesis, University of Bristol Bristol, England, UK, 2017.

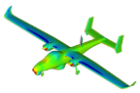
[18] R. C. Cheung, D. Rezgui, J. E. Cooper, and T. Wilson. Testing of a hinged wingtip device for gust loads alleviation. *Journal of Aircraft*, 55(5):2050–2067, 2018.

[19] R. C. Cheung, D. Rezgui, J. E. Cooper, and T. Wilson. Testing of folding wingtip for gust load alleviation of flexible high-aspect-ratio wing. *Journal of Aircraft*, 57(5): 876–888, 2020.

[20] F. J. Healy and F. J. Healy. The impact of geometric nonlinearities on the behaviour of floating wingtips. *Aircraft Design*, 82:1–23, 2023.

[21] R. J. McGhee. *Experimental results for the Eppler 387 airfoil at low Reynolds numbers in the Langley low-turbulence pressure tunnel*, volume 4062. National Aeronautics and Space Administration, 1988.

[22] A. Rausa, A. Guardone, and F. Auteri. Implementation of γ-Re_θ transition model within SU2: model validation and verification. In *AIAA Scitech 2023 Forum*, page 1570, 2023.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

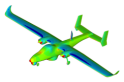
SMART WING STALL FLUTTER ALLEVIATION

R. Moosavi

School of Physics, Engineering & Computer Science
Department of Engineering & Technology, Centre for Engineering Research
University of Hertfordshire
Hertfordshire, Hatfield, AL10 9AB, UK
m.moosavi@herts.ac.uk.ac, https://www.herts.ac.uk/

Abstract *One of classical aeroelastic phenomenon seriously affecting the performance of flexible wings at high angle of attack is stall flutter. This work shows how the stall flutter phenomenon can be alleviated by implementing a piezoelectric patch. The dynamic responses are calculated to analyze the nonlinear characteristics of the aeroelastic system. Parametric study has been performed for different parameters. The main contribution of this work is to represent a way to alleviate the stall flutter on a regular wing by using a passive aeroelastic control including piezoelectric patches and shunt circuits in which the size of the required inductance is small. The results indicate a considerable improvement in dynamic aeroelastic behavior of a wing.*

Keywords: Smart Wing, Stall Flutter, Piezoelectric Material



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

SHAPING NEXT-GENERATION AIRCRAFT WITH MULTIDISCIPLINARY DESIGN OPTIMIZATION

Joaquim R. R. A. Martins

Department of Aerospace Engineering
University of Michigan
FXB Building, 1320 Beal Avenue, Ann Arbor, MI 48109-2140, USA
jrram@umich.edu, https://mdolab.engin.umich.edu/

Abstract. *Air transportation's quest for enhanced efficiency demands groundbreaking advancements in aircraft design, particularly in the optimization of aircraft configurations, including aerodynamic shape, structural sizing, and propulsive performance. This talk examines the latest developments in computational fluid dynamics (CFD)-based design optimization, addressing the critical challenges and innovative solutions in shaping aircraft components for maximum performance and operational efficiency. We explore the integration of CFD with structural mechanics to optimize wing aerodynamics, enhance lift-to-drag ratio, and reduce structural weight. The discussion extends to the aerodynamic optimization of propellers and turbofans, highlighting how advanced computational methods facilitate the design of propulsion systems that integrate optimally with the airframe. By using multidisciplinary optimization (MDO) approaches, we advance the design of these key aircraft components, considering complex interactions and design constraints. We focus on the adjoint method for efficient gradient computation, which is essential for handling the large number of shape and sizing variables required in high-fidelity aircraft design. The computational tools and methods developed through this research enable unprecedented fidelity and scale in design optimization problems. The talk presents case studies demonstrating the successful application of these methods in optimizing aircraft wings and innovative designs of propellers and turbofans. Finally, we discuss our vision for advancing computational design optimization methods to shape the future of high-performance aerospace engineering and other complex engineered systems.*

Keywords: MDO, CFD, aerodynamic shape optimization, aircraft design

AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

A VELOCITY-FIELD LEVEL-SET BASED TOPOLOGY OPTIMIZATION OF STRUCTURES SUBJECTED TO AEROTHERMODYNAMICS LOADS

Apoorva Kanti[1*], Pankil N. Mishra[2] and Abhijit Gogulapati[3]

1: B.Tech. student (210010013@iitb.ac.in)
2: PhD candidate (pankilmishra@iitb.ac.in)
3: Assistant Professor (abhijit@aero.ac.in)

Department of Aerospace Engineering
Indian Institute of Technology Bombay
IIT Bombay, Powai, Mumbai 400 076, Maharashtra, India

Abstract. *Level-set-based topology optimization is used to identify optimal material layout in compliant structures operating in high-speed flow. The static stability of a two-dimensional panel, described using a thermal buckling metric, is used as the objective function. The panel has fixed-fixed boundary conditions and is subjected to hypersonic flow over its top surface. The aerothermal loads due to flow are computed using a combination of the piston theory and Eckert's reference enthalpy. The topology optimization is based on a velocity field level set method (VFLSM). This approach describes the level set using a signed distance function, with the solid boundary as the zero-level contour. The velocity field is defined over the design domain using piecewise basis functions. An in-house modular and scalable Python code has been developed to perform the topology optimization of the panel under consideration. The sensitivity analysis is performed analytically using an adjoint-based method. The stepping is based on the method of moving asymptotes (MMA), which allows efficient handling of multiple constraints and additional types of design variables. The case study presented explores a coupled aerothermodynamic scenario with a non-uniform temperature and a fixed solid top surface. The study is a precursor to VFLSM-based topology optimization of coupled aero-thermo-elastic systems.*

Keywords: velocity field level-set method, topology optimization, thermal buckling, coupled aerothermal loads, fluid-thermal-structural interactions

1 INTRODUCTION

Studying flexible aerospace structures experiencing combined aerodynamic pressure and thermal loads is crucial for designing and advancing hypersonic vehicles operating in extreme environments. The LSM of topology optimization is a powerful computational approach that helps design lightweight and high-performance aerospace structures by optimizing material distribution within a given design domain.

The LSM employs an implicit geometry representation for topology optimization, enabling smooth material boundaries and precise control over shape-sensitive design constraints. This makes it particularly suitable for applications involving additive manufacturing and multiphysics interactions. Owing to these capabilities, the LSM has been effectively applied to structural systems subjected to pronounced thermal effects, where smooth interfaces and accurate thermal-structural coupling are critical. The approach can be used to identify optimal configurations when coupled thermal, aerodynamic, and structural interactions are present. However, topology optimization of aero-thermoelastic systems has received limited attention [1–4].

Researchers have developed optimization frameworks using LSM for thermoelastic systems, where thermal expansion, heat conduction, and stress redistribution are critical considerations [5–9]. [6] systematically evaluated the performance of different objective functions—compliance, strain energy, and stress—in thermoelastic topology optimization using LSM, demonstrating its flexibility and robustness in balancing thermal and structural performance. Similarly, in [7], the authors explored the design of architected materials tailored for thermoelastic macrostructures, emphasizing LSM's capability to capture intricate topological features while accommodating thermal expansion effects. Together, these studies underscore the effectiveness of LSM in addressing the unique challenges posed by thermoelastic behavior in advanced structural systems.

The LSM has also been increasingly applied to aeroelastic systems, where fluid-structure interactions and instability phenomena such as flutter and divergence significantly influence structural performance. Its ability to represent smooth, evolving boundaries makes it well-suited for aeroelastic topology optimization, enabling tight coupling between aerodynamic and structural responses. Studies such as [10, 11] demonstrate the method's utility in capturing optimal load-bearing configurations under aeroelastic constraints. Furthermore, the spectral level set method has been employed to optimize aircraft aeroelastic response [12]. The applicability of the LSM for realistic configurations by using a common research wing model is highlighted in [13]. Collectively, these works affirm the LSM's strength in addressing multiphysics constraints inherent in aeroelastic design.

Subsequent research on LSM has yielded several enhancements and improvements from a computational implementation point of view [14, 15]. The velocity field level set method (VFLSM), introduced in [16], updates structural shape and topology using a velocity field defined over the whole design domain. This velocity field is discretized using basis functions. This approach maps the functional variational optimization problem into a finite-dimensional design space, enabling mathematical programming algorithms to efficiently handle multiple constraints and additional types of design variables. While retaining the

implicit representation benefits, the method advances material boundaries (zero level-set contour) by solving the Hamilton-Jacobi equation. This approach has been implemented in compliance minimization and active vibration control problems as demonstrated in[17–19].

Although significant progress has been made in applying the LSM to both thermoelastic and aeroelastic systems independently, its implementation in aero-thermoelastic topology optimization remains largely unexplored, to the best of the author's knowledge. This gap highlights the need for a unified framework to fully exploit LSM's strengths, particularly VFLSM, in capturing smooth boundaries and accommodating evolving interfaces under coupled thermal and aerodynamic constraints.

The specific objective of the present study is to explore the following in the context of a coupled aero-thermoelastic system using the VFLSM:

- Develop and verify a modular framework that employs VFLSM for topology optimization of structures subjected to combined aerodynamic and thermal loads, with the aim of enhancing thermal buckling resistance.

In this study, a one-way coupling is considered: the influence of aerothermal loads on structural response is modeled, while the reverse influence of structural deformation on aerothermal loads is not included. The structure of the paper is as follows: Section 2 describes the panel setup, the computational tools used in the analysis, a brief insight into the VFLSM, and the topology optimization framework. Section 3 presents and discusses the results obtained from a case study. Lastly, Section 4 concludes the study by summarizing the main outcomes and suggesting potential avenues for future research.

2 CONFIGURATION AND MODELING

The structural system under consideration consists of a panel with a length L and a thickness h, as illustrated in Fig.1. The panel is assumed to extend infinitely in the third dimension. To ensure computational feasibility and effectively analyze changes in the panel's internal topology, the ratio L/h is set to 25. Details regarding the panel's geometric and material properties are provided in Table 1. Additionally, the top surface is subjected to hypersonic flow conditions, as described in Table 2.

The aero-thermal response is determined by integrating finite element-based structural and thermal solvers with pressure and heat loads. These loads are computed using piston theory and Eckert's reference enthalpy methods, respectively, as discussed in detail in [1].

The optimization framework provided in [17] has been modified for the present study. It was written in MATLAB specifically for compliance minimization problems. The following changes were incorporated into the MATLAB framework for the present study:

1. Python is used as the programming language for the present framework. It has a modular structure so that new features can be incorporated easily in the future to expand the framework's capability and versatility.

2. The Matlab code allowed for only a particular mesh size, i.e., the design domain size was the mesh size. A decoupling was performed to allow for any mesh size for any given design domain. This had to be ensured at every step where element size was taken as 1 by 1. Additionally, it was non-dimensional; for the current study, appropriate dimensionalization has been implemented.

3. *Gauss_Phi* represents the interpolated level set values at Gauss integration points within finite elements. It is computed by mapping the level set function ϕ from the grid to element nodes and then interpolating to Gauss points using bilinear shape functions. *Gauss_Phi* determines the material state at each point and is essential for evaluating stiffness and sensitivities. The FE mesh and *Gauss_Phi* calculations had to be modified to be in compliance with the FEA model given in [1].

4. The evolution and re-initialization function for ϕ have been modified from the Euler scheme in [17] to the 4th order Runge-Kutta scheme for better accuracy in obtaining level set function from the design variables.

The present framework has been verified using all the cases of compliance minimization problems presented in [17].

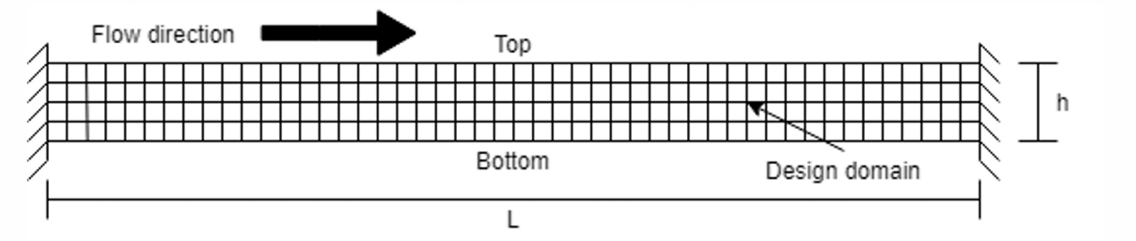


Figure 1: Panel Geometry

Table 1: Geometric and Material properties of the panel

Property	**Value**
Length	0.125 m
Thickness	0.005 m
Young's Modulus	103 GPa
Density	4520 kg/m^3
Poisson's ratio	0.3
Coefficient of Thermal Expansion	4.7x10^{-6}/Deg C
Conductivity	7.1 $W/m-K$
Specific Heat	463.0 $J/kg-K$

2.1 Velocity Field Level Set Method (VFLSM)

In the VFLSM [16, 17], the structural topology evolves through a normal velocity field, $V_n(x,t)$, which governs the advancement of the level-set function $\Phi(x,t)$. Generally, $V_n(x,t)$ is the steepest descent direction obtained during the optimization process.

Table 2: Flow Conditions

Property	**Value**
Free-stream Mach Number (M_∞)	7.0
Altitude	30.0 km
Free-stream Pressure (P_∞)	1090 Pa
Free-stream Temperature (T_∞)	227 K
Upstream Transition to Turbulence	1.0 m

However, it is formulated here by discretizing the velocity field onto the level-set grid and defining the velocity at the discretized nodes as the design variable for optimization. This mapping process represents the core concept of the Velocity Finite Level-Set (VFLS) method. The mapping is expressed by:

$$V^n(\mathbf{x}) = \sum_{j=1}^{N} \beta_j p_j(\mathbf{x}) \tag{1}$$

Here, variable $\beta_j (j = 1, 2, ..., N)$ is equal to the velocity value V_j^n at specific velocity points, and $p_j(\mathbf{x})$ are the basis function.

When the basis functions $p_j(\mathbf{x})$ satisfy the Kronecker delta property (e.g., FE shape function, Shepard basis function, or piecewise basis function), the velocity design variables β_j are equal to the velocity values V_j^n at specific velocity points (i.e., $V_j^n = V^n(\mathbf{x}_j)$). The present study utilizes the piecewise basis, so the velocity values are used as velocity design variables. The material boundary evolution is achieved by solving the Hamilton-Jacobi equation:

$$\frac{\partial \Phi}{\partial t} - V_b^n |\nabla \Phi| = 0 \tag{2}$$

where V_b^n is the normal boundary velocity that drives the shape and topology toward the corresponding direction. The level set function is usually defined as a signed distance function, and to preserve the signed distance property, the following re-initialization step is usually required:

$$\frac{\partial \Phi}{\partial t} - sign(\Phi_o)(|\nabla \Phi| - 1) = 0 \tag{3}$$

where Φ_o is the level set function before re-initialization, the constructed velocity field is merged with the Hamilton-Jacobi equation (2) to achieve boundary evolution. The signed distance property of the level set function is retained by solving (3).

Instead of optimizing a continuous implicit boundary velocity function, the above process simplifies the problem by utilizing the explicit discretized velocity design variables. This allows the use of general mathematical optimization algorithms to solve the optimization problem. While these algorithms may take slightly more time per iteration than the steepest descent method used in the standard level set method, they are better at handling multiple constraints and diverse design variables. Additionally, it retains the benefits of implementing an implicit representation for the material boundary since the LS function and the evolution scheme are the same as those implemented in the standard LSM.

2.2 Sensitivity Analysis

Sensitivity analysis involves analytically computing design derivatives to efficiently handle large-scale design variables (V_n). The objective function is formulated using thermal buckling eigenvalues(λ_b) and is given by λ^*:

$$\lambda^* = min(\lambda_b) \tag{4}$$

The detailed sensitivity analysis combines the formulations given in [1] and [17] to obtain the derivative of objective and constraint functions with respect to the design variables.

2.3 Topology Optimization Model

The FE analysis module contains the models described in [1]. The simulation concludes upon reaching thermal buckling or the designated simulation time, N_s seconds. The primary goal of topology optimization is to maximize the thermal buckling metric λ^* at the end of the analysis. Essentially, the objective is to achieve a topology that remains stable without buckling under the given boundary conditions. The topology optimization problem is solved using the Method of Moving Asymptotes (MMA) [20], which explores the design space based on sensitivity analysis discussed in 2.2. The topology optimization problem based on the VFLS method is formulated as follows:

$$\begin{cases} \text{Find } \mathbf{V}^n = [V_1^n, V_2^n, \ldots, V_N^n]^T, \\ \text{Maximize: } J = \lambda^* \\ \text{Subject to:} \\ \quad G_V \equiv \int_\Omega d\Omega - f_V V_D \leq 0, \\ \quad \underline{V}^n \leq V_j^n \leq \overline{V}^n, \quad j = 1, 2, \ldots, N. \end{cases} \tag{5}$$

Where:

- J denotes the objective function (i.e., the thermal buckling metric).
- $\mathbf{D}$ denotes the material elasticity tensor.
- Γ_D is the Dirichlet boundary.
- G_V denotes the material volume constraint.
- V_D is the volume of the design domain D.
- f_V is the given volume fraction.
- $\underline{V}^n$ and $\overline{V}^n$ are, respectively, the lower and upper bounds of the velocity design variables, which are consistent with the Courant-Friedrichs-Lewy (CFL) condition [21].

For the present case study, it is essential to designate certain specific regions of the panel as a non-design solid domain, wherever it is exposed to the hypersonic flow. Otherwise, the temperature values would escalate uncontrollably, resulting in a singularity, especially since the load models employed—piston theory and Eckert's reference enthalpy method—are not equipped to handle discontinuities such as holes or cavities on the top surface. There is more than one way to set a non-design domain in the level set method, and an easily implemented one is chosen [17]. The level set function in the non-design domain is firstly set to a positive constant and then modified as the signed distance function

by re-initialization by satisfying (3). Here, the first 3 rows from the top of the panel in the level-set grid are fixed as per the above-mentioned procedure. The initial topology is chosen in such a way that there is no freely floating material in the system. It is obtained by specifying the number of holes in the x and y direction ($xnum$ and $ynum$ respectively), and the hole radius (cr) [17].

The optimization algorithm terminates either due to the maximum number of iterations (N_I) or the percentage change in the objective function being less than or equal to 0.1%, or the relative deviation of the current material volume fraction from the required volume fraction must be less than or equal to 1%.

3 PRELIMINARY RESULTS AND DISCUSSION

Several parameters influence the outcome of the topology optimization model, including the initial conditions, evolution, re-initialization of the level-set function (ϕ), and the *move* parameter. The MMA algorithm used to solve the problem defined in (5) includes the move parameter, which limits the maximum step size of design variable updates. To evaluate its impact, case studies were conducted across varying design domains and boundary conditions. Many configurations failed to converge for extreme move values, while the results presented here correspond to optimally converged cases.

Table 3 lists the key parameters used in the optimization framework, where LeX and LeY refer to the element length and height, respectively. As the analysis focuses on aerothermal-structural coupling, pressure loads are omitted by assuming flow over an undeformed (flat) panel. The thermal loads are obtained from hypersonic flow over the surface. The top surface of the panel is fixed as a solid non-design domain; the thermal loads resulting from aerodynamic heating and radiation are also applied to the top surface, while adiabatic boundary conditions are imposed on the remaining surfaces for the performed cases. Several cases were run to check for accuracy and performance; one is presented in the following subsection.

Table 3: Key parameters taken initially

Parameter	**Value**
Number of elements	2500
Maximum iterations (N_I)	150
Min. change	1e-3
Volume fraction (f_v)	0.35
Max. and min. V_n value ($\overline{V}^n, \underline{V}^n$)	±0.5*min.(LeX,LeY)
Simulation time (N_s)	3 seconds
Thermal time step (Δt_{AT})	0.1 seconds
Reference Temperature (T_{ref})	300 K

3.1 Case Study: Coupled aerothermal system

The parameters mentioned in Table 3 are considered here. The initial design domain obtained by fixing $xnum = 16, ynum = 1, cr = 2.0$ having a volume fraction of 0.4734321

is shown in Fig.2. The evolution and re-initialization functions are evaluated using the 4th-order Runge-Kutta method with a step size of 0.5. Fig.3 shows the final topology, objective function, and volume evolution for a move value of 0.02.

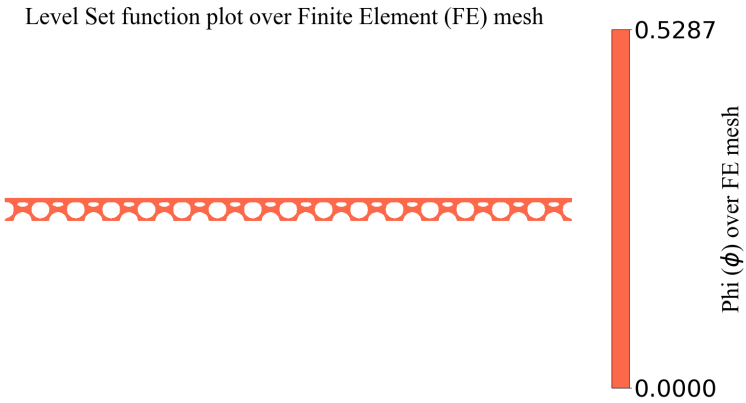


Figure 2: Case study: Phi ϕ field plot at iteration 0, i.e., initial design configuration

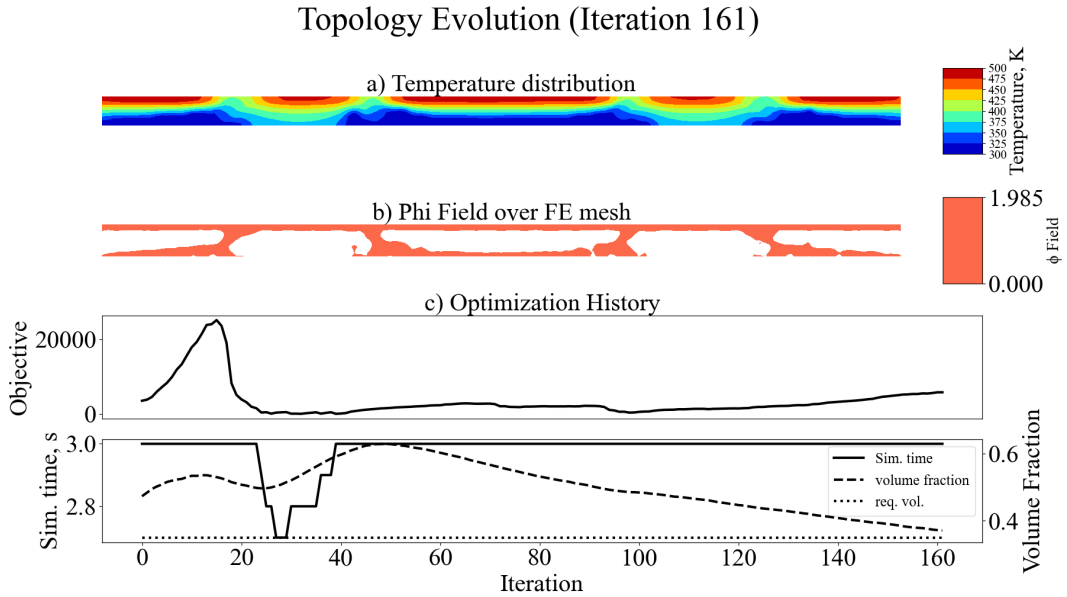


Figure 3: Case study: Final configuration and results

3.1.1 Discussion

The optimization algorithm demonstrates an interplay between thermal buckling metric, volume constraint, and structural evolution, with stringer configurations emerging as critical components in balancing these competing objectives. Fig.3 consists of three subplots (a, b, and c), each illustrating different aspects of the topology optimization process. Subfigure (a) presents the temperature distribution across the spatial domain using *contourf*, with the x and y axes representing the spatial extent and the color gradient showing temperature values in Kelvin (K), as indicated by the color bar. Subfigure (b) visualizes the ϕ field over the finite element (FE) mesh. It represents material distribution, with positive values indicating material presence and negative values denoting voids.

Subfigure (c) provides the optimization history of three metrics, which is divided into two plots. The top plot shows the objective function evolution over iterations, with the x-axis representing the iteration number and the y-axis displaying objective values. The bottom plot tracks simulation time (solid line) and volume fraction (dashed line) during optimization. The x-axis again represents the iteration number, while the left y-axis shows simulation time in seconds, and the right y-axis shows the volume fraction. The dotted line represents the required volume, i.e., 0.35.

The simulation optimized at iteration 176; however, minor discontinuities in the stringer configuration introduced instabilities, leading to a decline in the objective value as the volume continued to decrease. This trend is attributed to the system's increased susceptibility to thermal buckling under reduced volume conditions. Consequently, the results are presented up to the iteration corresponding to the peak objective value prior to its decline. Notably, this optimal value surpasses the initial objective, indicating that the optimization algorithm successfully identified an improved configuration. This final configuration exhibits a higher buckling performance metric while simultaneously satisfying the prescribed volume constraint.

The objective function's trajectory reveals a three-phase response:

- **Initial Growth Phase:** The metric rises steadily as the algorithm prioritizes thermal buckling resistance, redistributing material to form preliminary stringer geometries that delay buckling onset.
- **Decline Phase:** After peaking, the function declines as the optimizer aggressively reduces volume to meet constraints, inadvertently weakening buckling resistance. This is reflected in reduced simulation times (falling below 3s), where the panel buckles prematurely.
- **Recovery Phase:** The algorithm counteracts instability by strategically increasing volume, restoring the simulation time to 3s. Subsequent iterations refine the topology to maintain buckling resistance while incrementally reducing volume.

This adaptive loop highlights the framework's ability to retain structural viability while satisfying the volume constraint.

To enhance the current optimization algorithm, the stopping criteria can be refined by decreasing the value of the optimization parameter Min. Change (Table 3). This will allow the optimization algorithm to traverse further in the design space. Observing the change between the topology obtained for iteration 161 and those obtained in the subsequent iterations shows that its evolution is sensitive to the optimization parameters. Based on this, it was found that the maximum and minimum V_n value, RK step size dt, and re-initialization step size dt_{re} affect the overall stability of topology evolution. Conduction is the only mode of heat transfer considered in the present study, which is entirely based on the solid material distribution, i.e., topology obtained from ϕ. Hence, the disappearance of thin structural links significantly alters the temperature distribution of the panel, resulting in premature buckling. It was found that the improper tuning of parameters in the Hamilton Jacobi update and the re-initialization process results in such

panel configurations. With additional computational resources and time, these variations can be explored and utilized to avoid such instabilities. It will enable the framework to provide robust and thermally stable configurations.

4 CONCLUSIONS

This study presented a modular and scalable implementation of the VFLSM for topology optimization of structures subjected to coupled aerodynamic and thermal loads. A one-way coupled aero-thermal model was developed and integrated with a modified Python-based framework that incorporates analytical sensitivities and a Runge-Kutta-based evolution scheme for improved accuracy. Case studies demonstrated the algorithm's capability to evolve topologies that enhance thermal buckling resistance while satisfying volume constraints. The key observation is the sensitivity of the solution to design parameters such as velocity bounds, re-initialization steps, move parameters, and simulation time. Observed discontinuities in optimized configurations—linked to thermal conduction pathways—emphasize the importance of robust link retention in high-speed thermal environments. The current framework lays the foundation for an extension to coupled aero-thermoelastic systems, where there is an interplay of flow, structural deformation, and thermal loads. Such integration would enable a more accurate representation of multiphysics interactions, facilitating the design of robust high-speed aerospace structures.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the Department of Aerospace Engineering, IIT Bombay, for providing the necessary computational resources and infrastructure that enabled this work.

REFERENCES

[1] P. N. Mishra and A. Gogulapati. Topology optimization of fully coupled aerothermoelastic system. In *Proceedings of the international conference on multidisciplinary design optimization of aerospace systems (AEROBEST 2021)*, volume 192, 2021.

[2] B. Stanford and P. Beran. Aerothermoelastic topology optimization with flutter and buckling metrics. *Structural and Multidisciplinary Optimization*, 48:149–171, 2013.

[3] D. J. Munk, G. A. Vio, and G. P. Steven. Aerothermoelastic structural topology optimisation for a hypersonic transport aircraft wing. In *11th World Congress on Structural and Multidisciplinary Optimization. Sydney, Australia. https://web. aeromech. usyd. edu. au/WCSMO2015/papers/1013_paper. pdf*, 2015.

[4] L. J. Smith, L. J. Halim, G. Kennedy, and M. J. Smith. A high-fidelity coupling framework for aerothermoelastic analysis and adjoint-based gradient evaluation. In *AIAA Scitech 2021 Forum*, page 0407, 2021.

[5] H. Chung, O. Amir, and H. A. Kim. Level-set topology optimization considering nonlinear thermoelasticity. *Computer Methods in Applied Mechanics and Engineering*, 361:112735, 2020.

[6] D. J. Neiferd, R. V. Grandhi, J. D. Deaton, and P. S. Beran. Level-set topology optimization of thermoelastic structures-a comparison of compliance, strain energy,

and stress objectives. In *2018 multidisciplinary analysis and optimization conference*, page 3577, 2018.

[7] L. Li, Z. Du, and H. A. Kim. Design of architected materials for thermoelastic macrostructures using level set method. *JOM*, 72(4):1734–1744, 2020.

[8] Q. Xia and M. Y. Wang. Topology optimization of thermoelastic structures using level set method. *Computational Mechanics*, 42:837–857, 2008.

[9] H. Chung, O. Amir, and H. A. Kim. Nonlinear thermoelastic topology optimization with the level-set method. In *AIAA Scitech 2019 Forum*, page 1470, 2019.

[10] B. Stanford, P. Beran, and M. Bhatia. Aeroelastic topology optimization of blade-stiffened panels. *Journal of Aircraft*, 51(3):938–944, 2014.

[11] P. D. Dunning, B. Stanford, and H. A. Kim. Level-set topology optimization with aeroelastic constraints. In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, page 1128, 2015.

[12] A. Gomes and A. Suleman. Optimization of aircraft aeroelastic response using level set methods. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, page 1892, 2005.

[13] P. D. Dunning, B. Stanford, and H. A. Kim. Aerostructural level set topology optimization for a common research model wing. In *10th AIAA Multidisciplinary Design Optimization Conference*, page 0634, 2014.

[14] M. Yulin and W. Xiaoming. A level set method for structural topology optimization and its applications. *Advances in Engineering software*, 35(7):415–441, 2004.

[15] S. Townsend and H. A. Kim. A level set topology optimization method for the buckling of shell structures. *Structural and Multidisciplinary Optimization*, 60(5): 1783–1800, 2019.

[16] Y. Wang and Z. Kang. A velocity field level set method for shape and topology optimization. *International Journal for Numerical Methods in Engineering*, 115(11): 1315–1336, 2018.

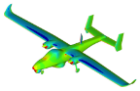
[17] Wang, y., kang, z. (2021). matlab implementations of velocity field level set method for topology optimization: an 80-line code for 2d and a 100-line code for 3d problems. structural and multidisciplinary optimization, 64(6), 4325-4342.

[18] Y. Wang, Z. Kang, and X. Zhang. A velocity field level set method for topology optimization of piezoelectric layer on the plate with active vibration control. *Mechanics of Advanced Materials and Structures*, 30(7):1326–1339, 2023.

[19] W. Cheng, X. Zhang, T. Hu, J. Li, and Y. Wang. Stress-constrained topology optimization using the velocity field level set method. *Computers & Structures*, 305: 107577, 2024.

[20] K. Svanberg. The method of moving asymptotes—a new method for structural optimization. *International journal for numerical methods in engineering*, 24(2):359–373, 1987.

[21] S. Osher, R. Fedkiw, and K. Piechor. Level set methods and dynamic implicit surfaces. *Appl. Mech. Rev.*, 57(3):B15–B15, 2004.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

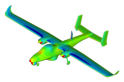
MACH: A HIGH-FIDELITY SURROGATE MODEL DESIGNED FOR GENERALIZATION ACROSS GEOMETRIES AND PHYSICS

Chinmay Shrivastava[1]* and Noah Evers[1]

1: CompLabs
{chinmay, noah}@complabs.co, www.complabs.co

Abstract *High-fidelity simulations remain a computational bottleneck in aerospace design. Traditional surrogate models are often used to accelerate simulations but lack generalization to novel geometries, struggle with nonlinear physics, and require extensive data and manual engineering. We introduce Mach, a family of pretrained geometric foundation models for 3D simulation and optimization. Trained on a large corpus of 3D geometries, Mach learns scalable and generalizable latent representation of shapes, enabling efficient fine-tuning for simulations like CFD. It has demonstrated strong generalization in predicting simulation results on unseen geometries and scales with additional data and compute. Mach is not generative, but rather enables localized optimization via semantic beam search—starting from an existing geometry and refining it within a high-dimensional, structured latent space. This allows for expressive design updates, including subtle morphing and nonparametric edits, without handcrafted parameterization. Mach offers a path toward high-fidelity simulation and optimization workflows that are more computationally efficient, integrate tightly with modern design tools, and accelerate iteration in real-world engineering pipelines.*

Keywords: Foundation models, surrogate modeling, large geometric model, physics-informed machine learning



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

COLLABORATIVE TOPOLOGY AND FIBER ORIENTATION OPTIMIZATION SIMPLIFIED VIA MULTI-RESOLUTION ANISOTROPY REPRESENTATION

Shuya Nozawa and Gokhan Serhat

Department of Mechanical Engineering
KU Leuven
Spoorwegstraat 12, 8200 Bruges, Belgium
{shuya.nozawa, gokhan.serhat}@kuleuven.be

Abstract. *The use of fiber-reinforced composites in the aerospace industry has been continuously increasing over the last decades due to their exceptional stiffness-to-weight ratio. Consequently, the development of efficient design methodologies for composites has become crucial. Topology optimization (TO), the geometrical design technique with the highest degree of freedom, offers great utility in this context. The performance of composite structures can be further enhanced by exploiting the material's anisotropic nature through the optimization of spatial fiber orientations. These two design approaches have been often used together in recent works since the advances in additive manufacturing have enabled straightforward fabrication of composite parts with complex geometries and curved fibers. However, combined optimization of structural topology and fiber angle distribution frequently causes numerical instability and local optimality problems due to the large number of design variables having different nature. To circumvent these issues, we propose a novel anisotropic TO methodology in this study. Our approach relies on initiating the optimization iterations from isotropic design and decoupling the update processes for the topological density and fiber orientation distributions. In addition, different from the conventional methods, we assign auxiliary design variables controlling fiber directions to the nodes of an unstructured triangular mesh created over the density distribution. The fiber angle value in each finite element is interpolated based on the information from the surrounding control nodes. The effectiveness of the proposed approach is demonstrated through case studies involving different geometrical constraints and load cases.*

Keywords: Fiber-reinforced composites, Stiffness tailoring, Topology optimization, Fiber path design, Additive manufacturing

1 INTRODUCTION

In recent decades, fiber-reinforced composites have garnered significant interest across various industries (e.g., aerospace, automotive, construction, etc.) due to their outstanding mechanical properties. The exceptional stiffness- and strength-to-mass ratios of these materials make them essential for lightweight structures.

A key challenge in the design of fiber-reinforced components is the need to consider material anisotropy. Instead of using the composite in a quasi-isotropic setting, fiber angles across the structure can be tailored to obtain directional elastic properties for different purposes such as maximizing stiffness [1, 2] or fundamental frequency [3, 4]. Even higher structural performance can be attained using variable stiffness (VS) composite involving spatially varying fiber orientations [5, 6]. However, VS designs may suffer from manufacturing defects such as fiber breakage [7] and tow gaps/overlaps [8]. These issues are mitigated by advances in composite additive manufacturing, which enable increasingly higher printing resolution and fiber curvatures [9].

In the context of lightweight design, topology optimization (TO) for anisotropic materials has been an intensive research topic as it can algorithmically generate optimal composite structures [10]. Nevertheless, anisotropic TO has multiple intrinsic characteristics that hinder its wide adoption. For example, it requires handling numerous design variables that represent different physical properties [11]. This condition often causes numerical instabilities that lead to optimized designs with unclear topological density distribution.

Another problematic aspect of anisotropic TO is its high susceptibility to local optimality problems. Such issues arise primarily from the strong influence of the initial fiber direction field on the optimization process [12]. Changes in anisotropy information at early optimization steps strongly affect the update direction for the density distribution, leading to significant variations in the final design. The large number of design variables also increases the likelihood of getting trapped in locally optimal solutions.

Optimal designs hold no value unless they meet manufacturability requirements. When anisotropic composites are considered, the fiber orientations must gradually vary across the structure to ensure the continuity constraints. Failure to satisfy this condition can lead to unexpected stress concentrations and fiber breakage [13].

As mentioned above, combined topology and fiber path optimization of composites involves a number of inherent difficulties. In this work, we address these hurdles by proposing a novel design technique that can manipulate the geometry and spatial anisotropy to carry the applied loads efficiently. The presented preliminary results highlight that our approach can provide stiff and lightweight structures that are also suitable for 3D printing.

2 COLLABORATIVE TOPOLOGY AND FIBER ORIENTATION OPTIMIZATION TECHNIQUE

This section covers the key components of the introduced anisotropic TO method. We summarize the strategies implemented to handle the aforementioned challenges regarding the design process.

2.1 Density-based topology optimization

The proposed method relies on the density-based approach for the TO. Specifically, we utilize the solid isotropic material with penalization (SIMP) method [14], which involves parameterizing material properties via density penalization. In this strategy, the effective stiffness matrix ($\mathbf{C}$) is computed by power-law interpolation as

$$\mathbf{C} = \rho^p \mathbf{C}^0 \tag{1}$$

where p, ρ, and $\mathbf{C}^0$ are the penalization parameter, normalized density, and anisotropic stiffness matrix for the material, respectively. Under plane stress state, $\mathbf{C}^0$ can be described as follows:

$$\mathbf{C}^0 = \begin{pmatrix} C_{11} & C_{12} & 0 \\ C_{12} & C_{22} & 0 \\ 0 & 0 & C_{66} \end{pmatrix} \tag{2}$$

The entities of the above matrix are defined as

$$\begin{aligned} C_{11} &= E_1/(1-\nu_{12}\nu_{21}) \\ C_{12} &= \nu_{21}E_1/(1-\nu_{12}\nu_{21}) \\ C_{22} &= E_2/(1-\nu_{12}\nu_{21}) \\ C_{66} &= G_{12} \end{aligned} \tag{3}$$

where E_1 and E_2 are the elastic moduli along and perpendicular to the fibers, respectively; ν_{12} and ν_{21} are the major and minor Poisson's ratios, respectively; and G_{12} is the shear modulus.

In the SIMP method, the normalized elemental density (ρ_e) assigned to each e-th finite element is used as a design variable. Since anisotropic TO is considered in this paper, the elemental fiber orientation (θ_e) is introduced as an additional parameter. Thus, the effective stiffness matrix ($\mathbf{C}_e$) on the e-th element is expressed as follows [15]:

$$\mathbf{C}_e = \rho_e^p[\mathbf{T}(\theta_e)\mathbf{C}^0\mathbf{T}^T(\theta_e)] \tag{4}$$

Here, $\mathbf{T}$ is the transformation matrix defined as

$$\mathbf{T} = \begin{pmatrix} c^2 & s^2 & -2cs \\ s^2 & c^2 & 2cs \\ cs & -cs & c^2-s^2 \end{pmatrix} \tag{5}$$

where c and s denote $\cos(\theta_e)$ and $\sin(\theta_e)$, respectively.

In this study, we focus on compliance minimization. For this prevalent design objective, the TO problem can be formulated as follows:

$$\begin{aligned} \min_{\rho_e} \quad & \mathbf{U}^T\mathbf{K}\mathbf{U} \\ \text{subject to} \quad & \rho_{\min} \le \rho_e \le 1 \\ & \textstyle\sum_{e=1}^{N} \gamma_e \rho_e \le \bar{V} \end{aligned} \tag{6}$$

where $\mathbf{U}$ and $\mathbf{K}$ are the global displacement and stiffness matrices, respectively; and $\rho_{\min}$, N, γ_e, and $\bar{V}$ are the minimum normalized density value, number of finite elements in the mesh, elemental volume fraction, and maximum total volume fraction, respectively.

The implementation of density-based TO also requires spatial regularization to prevent designs with checkerboard patterns [16]. Thus, we use a density filtering [17] to obtain manufacturable structures with clear geometric layouts.

2.2 Vector-based representation of fiber angles

Explicit representation of the fiber angles poses numerical issues in the fiber orientation optimization (FOO) process due to its periodicity, which results in multi-modal responses [12, 18]. As a remedy, we employ vector-based fiber orientation representation in Cartesian coordinates [19, 20]. In this strategy, the fiber orientation (θ) is replaced by the vector $[u, v]^T$ defined as follows:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} \tag{7}$$

Using this double-valued anisotropy description not only helps mitigate the periodicity problem but also expands the design space, enabling the attainment of structures with higher performance.

2.3 Collaborative optimization strategy

In anisotropic TO, where both global topology and local fiber orientations are controlled, optimization instability is a common issue. This problem arises from the need to concurrently manage a large number of variables with varying scales and physical nature. To address this challenge, the optimization of topology and fiber angles are decoupled in our strategy. Such methods can be classified as either sequential or collaborative. The former type relies on performing TO and FOO only once in sequence, while the latter involves executing these operations iteratively by passing the results from one step to the other [21]. We adopt the collaborative approach since it is methodologically closer to the simultaneous optimization (which can theoretically provide structures with the highest performance, although it is more unstable) and is compatible with the anisotropy increment method described in the following subsection.

2.4 Gradual increase of anisotropy

Composite TO often produces locally optimal solutions due to its dependence on the initial fiber orientation distribution. One effective way to mitigate this issue is to start the optimization from the isotropic state. Using this principle, Ichihara et al. [22] introduced the concept of the gradual emergence of orthotropy via penalization. However, this treatment has drawbacks as it depends on an artificial penalization parameter that may be difficult to tune. By contrast, our approach relies on controlling the level of orthotropy directly by adjusting the fiber volume fraction (V_F) in the composite. This physical parameter offers several advantages, such as interpretability and straightforward implementation. In our strategy, the fiber volume fraction is updated at the specified iteration intervals according to the following equation:

$$V_F{}^{n+1} = \min\left(V_F{}^n + \Delta V_F, 0.5\right) \tag{8}$$

Here, $V_F{}^n$ is the fiber volume fraction used at the n-th step, and ΔV_F is the applied increment. The resulting value ($V_F{}^{n+1}$) represents the degree of anisotropy at the next step, and it is used to calculate E_1 and E_2 in Eq.(3) based on the rule of mixtures [15]. The iterations continue until reaching the desired fiber volume fraction, which is defined as 0.5 in this study.

2.5 Anisotropy parametrization via interpolation between sparse control points

In TO, neater designs are more likely to be obtained by reducing the number of variables involved. Considering this fact, we adopt an interpolation technique to represent spatially varying fiber orientations with fewer design variables concerning anisotropy. Guo and Zhou [23] made a similar attempt, but their technique involved densely and uniformly distributed sample points, limiting the potential level of problem reduction. In contrast, we achieve an enhanced downscaling capability by performing the interpolation via a triangular mesh that is generated over the structure and adaptively modified at each FOO phase.

The procedure involves the following stages: First, the solid region (defined as the area with a normalized density exceeding 0.5) is extracted from the design obtained at the last TO step. Then, this area is represented using a signed distance function, where the boundary is defined as the zero-level set. Next, we select a small positive level value to create an isosurface, which represents an outline that encloses the original density distribution while preserving the gaps. This contour is then smoothed using a Helmholtz PDE-based method [24], and a triangular mesh is generated over the refined isosurface. Finally, fiber direction information at the structured TO grid elements is obtained by linearly interpolating between the orientation vectors at the nodes of each associated triangular element. This process is illustrated in Fig.1, where the triangular mesh in light blue corresponds to the density distribution whose symmetric counterpart is shown on the upper side.

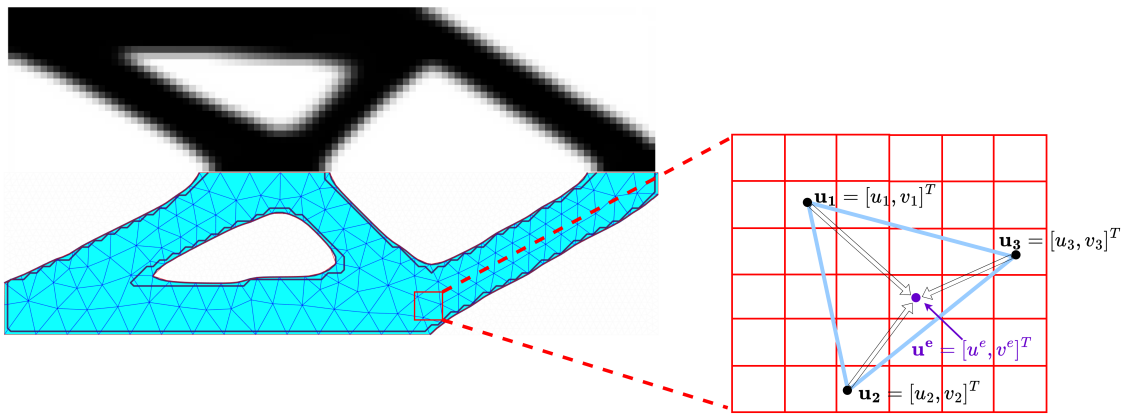


Figure 1: A sample result from the adaptive discretization (left) and the conceptual diagram of the fiber orientation interpolation (right).

The interpolation weights used for calculating the fiber orientations are based on the distances between the finite element center and the nodes of the corresponding auxiliary triangular element. Thus, the elemental anisotropy information $[u_{we}, v_{we}]^T$ is obtained as

$$\begin{bmatrix} u_{we} \\ v_{we} \end{bmatrix} = \begin{bmatrix} (w_{1e}u_1 + w_{2e}u_2 + w_{3e}u_3)/(w_{1e} + w_{2e} + w_{3e}) \\ (w_{1e}v_1 + w_{2e}v_2 + w_{3e}v_3)/(w_{1e} + w_{2e} + w_{3e}) \end{bmatrix} \tag{9}$$

where u_i and v_i are the components of the direction vector placed on the i-th node of the triangular element, and w_{ie} represents the inverse of Euclidean distance between the

node and center of the e-th finite element enclosed by the triangle. Since the physical orientation vector entities u^e and v^e must satisfy the Pythagorean identity, they can be expressed as follows:

$$\begin{bmatrix} u^e \\ v^e \end{bmatrix} = \begin{bmatrix} u_{we}/(\sqrt{u_{we}^2 + v_{we}^2 + \epsilon}) \\ v_{we}/(\sqrt{u_{we}^2 + v_{we}^2 + \epsilon}) \end{bmatrix} \tag{10}$$

Here, ϵ is a small positive value added to avoid division by zero. Finally, we can compute the effective stiffness matrix $\mathbf{C}_e$ of the e-th finite element by substituting u^e and v^e for c and s in Eq.(5), respectively.

2.6 Fiber orientation optimization with the augmented objective function

Our method inherently poses a challenge regarding manufacturability. Since the fiber orientation distribution within each triangular element is controlled using only the three vectors defined at the nodes, even minor variations in design variables propagate excessively through the domain. Consequently, the global orientation field can involve abrupt changes, which make the design difficult to fabricate. To encourage achieving similar orientation vectors within the same triangular element, we augment the objective function with a penalty term P defined as follows:

$$P = \sum_{t=1}^{M} \sum_{k,l=1,k<l}^{3} \frac{(u_k^t - u_l^t)^2 + (v_k^t - v_l^t)^2}{d_{kl}^t} \tag{11}$$

Here, M is the total number of triangular elements in the mesh, u_k^t and v_k^t are the orientation vector components on the k-th node of the t-th triangular element, and d_{kl}^t is the Euclidean distance between the k-th and l-th nodes of the t-th triangular element. The augmented objective function for FOO is the weighted sum of the penalty term and original quantity of interest, compliance. Hence, the FOO problem can be expressed as

$$\begin{array}{ll} \min_{[u_k^t, v_k^t]^T} & \mathbf{U}^T\mathbf{K}\mathbf{U} + \alpha P \\ \text{subject to} & -1 \leq u_k^t \leq 1 \\ & -1 \leq v_k^t \leq 1 \end{array} \tag{12}$$

where α is the penalty coefficient used to balance the two objectives with different scales.

2.7 Propagation of anisotropy field

Our strategy relies on constructing the triangular mesh at each FOO stage based on the density distribution obtained at the preceding TO step. As the topology is updated, some nodes of the new mesh can fall outside of the solid region used to produce the previous mesh. At such nodes, estimating the fiber orientation coherently with the existing distribution would be more logical than simply using isotropic properties. For this purpose, we propagate the fiber orientation field to the entire design domain at the end of the FOO process. This operation is performed using the global interpolation technique based on Poisson's equation [25]:

$$\begin{cases} \nabla^2\phi(x,y) = \mathbf{u}(x,y), & (x,y) \in \Omega \text{ and } |\mathbf{u}(x,y)| = 0 \\ \phi(x,y) = \mathbf{u}(x,y), & (x,y) \in \Omega \text{ and } |\mathbf{u}(x,y)| \neq 0 \\ \phi(x,y) = 0, & (x,y) \in \partial\Omega \end{cases} \tag{13}$$

In these equations, ϕ is the vector field obtained by the global interpolation, $\mathbf{u}$ is the fiber orientation distribution from the last FOO stage, and Ω is the design domain. Once we obtain the field ϕ, it is passed to the next TO process.

3 CASE STUDY

This section covers the sample results for the compliance minimization of a cantilever beam. First, we describe the details of the problem and numerical settings used in the optimization process. Then, we present and discuss the optimal structure obtained with the introduced methodology, which is also compared with a design attained via a conventional approach.

3.1 Problem definition

The schematics of the considered cantilever beam problem are depicted in Fig.2. The domain has a height of $h = 1$ and a width of $w = 2$. The left edge is fully clamped, while a unit downward force F is applied at the center of the right edge. The finite element method is employed for the numerical analysis, and the structure is discretized into a 50×100 grid of square elements. The elastic moduli of the fiber and matrix materials are set to $E_f = 1$ and $E_m = 1/15$, respectively. Using these values and the fiber volume fraction obtained via Eq.(8), the elastic moduli E_1 and E_2 are calculated based on the rule of mixtures.

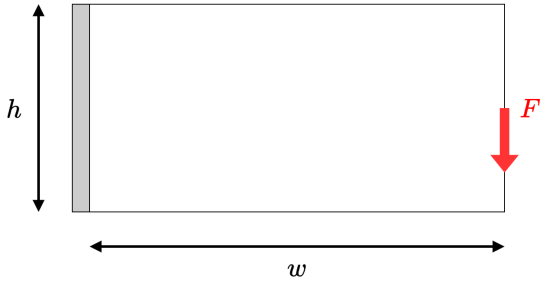


Figure 2: Schematics of considered cantilever beam under a downward tip force.

For the TO, we set the penalization parameter used in the SIMP method as $p = 3$. The filtering radius of the density filter is chosen to be twice the width of the finite element. The TO process ends and transitions to FOO after three internal updates are performed on the design variables. We use $\rho_{\min} = 0.001$, $\gamma_e = 1/5000$, and $\bar{V} = 0.5$ for the minimum normalized density value, elemental volume fraction, and maximum total volume fraction, respectively.

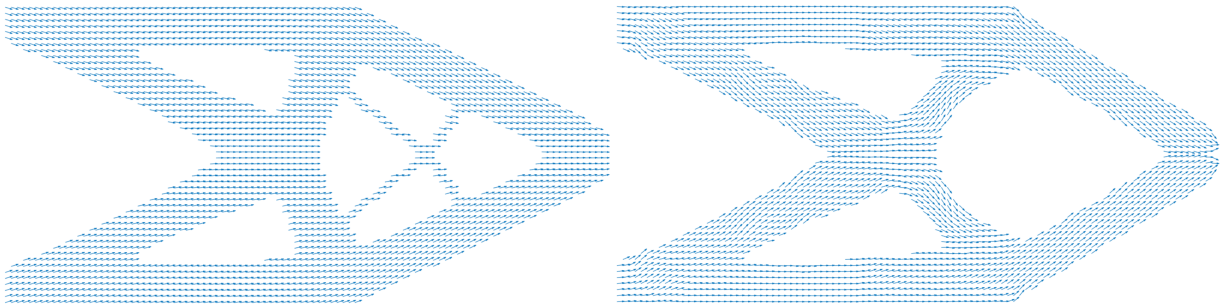
For the FOO iterations, a move limit is necessary because even small changes within the single fiber orientation vector can disrupt the smoothness of the resulting field. Accordingly, we apply a move limit with the initial value of 0.5, which is reduced by 3% after each FOO stage to improve convergence. We use the Matlab function **generateMesh** for the triangular mesh creation by setting the maximum segment size to $H_{\max} = 5$. This operation is conducted on the region inside the isosurface with the level set value of -1. We choose $\alpha = 0.3$ as the weighting factor of the extended objective function. After the design variables are internally updated three times, the FOO ends, and the anisotropy field propagation routine starts.

The overall optimization process is terminated when TO and FOO are performed 55 times each. The fiber volume fraction is initially set to zero, and it is gradually increased by $\Delta V_F = 0.05$ at every five transitions from TO to FOO. In both procedures, the method of moving asymptotes [26] is utilized as the optimizer. Considering the problem symmetricity, only the upper half of the design domain is optimized, and the results are reflected in the lower half.

3.2 Optimization results

In the following, we present the composite cantilever beam obtained with the proposed optimization method. For a brief performance comparison, we attained another design using the conventional concurrent technique involving the fiber orientation vectors defined in each finite element. The number of optimization iterations in the standard approach was set to match the number of finite element analyses used within our method to fairly evaluate the effectiveness of the introduced optimization strategy.

The optimal configurations obtained using the two methods are shown in Fig.3. The design obtained with the standard approach involves certain regions where fiber orientations are aligned perpendicularly to the corresponding structural members, reducing the load-carrying efficiency. In addition, the right side of the structure contains a thin X-shaped formation, which would have a minor contribution to the overall stiffness but notably increase the manufacturing complexity. In contrast, our technique resulted in a more rational design featuring fiber orientations aligned with the principal stress directions (such arrangement was shown to facilitate stiffness maximization [27]). Besides, a simpler density distribution that is more suitable for fabrication was achieved.



(a) (b)

Figure 3: Optimal topology and fiber orientations of the cantilever beams designed with (a) a conventional approach and (b) the proposed technique.

In Tab.1, the compliance values attained with the standard approach and our method are tabulated. Comparing these results demonstrates the effectiveness of the proposed technique, which yielded 15% higher performance compared to the conventional approach.

4 CONCLUSIONS

A novel anisotropic topology optimization strategy was introduced in this study. Starting from an isotropic design and decoupling the update processes for topology and fiber orientation, the proposed method suppresses the common numerical instability and local optimality issues. Additionally, the number of fiber angle variables is reduced by using a novel interpolation approach that involves generating an adaptive triangular mesh over

Table 1: Compliance values of the optimal designs obtained with a conventional approach and the developed technique.

Method	**Compliance value**
Conventional approach	188.0
Proposed technique	162.9

the solid domain. Together, these techniques enhance the efficiency and robustness of the optimization process. The covered case study concerning the design of a cantilever beam demonstrated the potential superiority of the developed methodology over conventional ones in terms of design neatness and attainable structural performance levels.

In future work, the capabilities and limitations of the presented method will be comprehensively investigated through multiple case studies involving diverse domain geometries, boundary conditions, solid volume fractions, and fiber volume ratios. In addition, the computational efficiency of the proposed approach will be rigorously assessed. In the present study, using the same termination criteria based on the number of optimization iterations resulted in nearly identical computation times for the conventional approach and our technique. Convergence could be alternatively determined by the objective function's rate of change, which could be more appropriate for runtime comparison. Moreover, the fiber volume fraction, which was assumed to be uniform across the entire domain in the current study, can be included as an extra design variable. With this addition, the proposed technique can be extended to design functionally graded composites with spatially varying fiber density.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the financial support from KU Leuven - Internal Funds (Grant No. STG/22/041).

REFERENCES

[1] G. Serhat. Concurrent lamination and tapering optimization of cantilever composite plates under shear. *Materials*, 14:2285, April 2021. doi:10.3390/ma14092285.

[2] S. Verstraete, G. Serhat, S. Debruyne, and F. Desplentere. Process optimization to maximize bonding performance of injection-moulded short-fibre composites. *Composites Part A: Applied Science and Manufacturing*, 193:108867, March 2025. doi:10.1016/j.compositesa.2025.108867.

[3] G. Serhat. Design of circular composite cylinders for optimal natural frequencies. *Materials*, 14:3203, June 2021. doi:10.3390/ma14123203.

[4] W. Wang, Q. Wang, R. Zhong, L. Chen, and X. Shi. Stacking sequence optimization of arbitrary quadrilateral laminated plates for maximum fundamental frequency by hybrid whale optimization algorithm. *Composite Structures*, 310:116764, April 2023. doi:10.1016/j.compstruct.2023.116764.

[5] Z. Gurdal and R. Olmedo. In-plane response of laminates with spatially varying fiber orientations - Variable stiffness concept. *AIAA Journal*, 31:751–758, April 1993. doi:10.2514/3.11613.

[6] P. Khandar Shahabad, M. Rafiei Anamagh, G. Serhat, I. Basdogan, and B. Bediz. Advanced lamination parameter interpolation and extrapolation methods for designing manufacturable variable stiffness laminates. *Composite Structures*, 326:117608, December 2023. doi:10.1016/j.compstruct.2023.117608.

[7] H. Zhang, J. Chen, and D. Yang. Fibre misalignment and breakage in 3d printing of continuous carbon fibre reinforced thermoplastic composites. *Additive Manufacturing*, 38:101775, February 2021. doi:10.1016/j.addma.2020.101775.

[8] P. Khandar Shahabad, M. Rafiei Anamagh, G. Serhat, I. Basdogan, and B. Bediz. Design and fabrication of variable stiffness composite laminates with enhanced dynamic behavior. In *Proceedings of the 21st European Conference on Composite Materials*, volume 3, pages 524–531. European Society for Composite Materials, July 2024. doi:10.60691/yj56-np80.

[9] X. Tian, A. Todoroki, T. Liu, L. Wu, Z. Hou, M. Ueda, Y. Hirano, R. Matsuzaki, K. Mizukami, K. Iizuka, A. V. Malakhov, A. N. Polilov, D. Li, and B. Lu. 3D printing of continuous fiber reinforced polymer composites: Development, application, and prospective. *Chinese Journal of Mechanical Engineering: Additive Manufacturing Frontiers*, 1:100016, March 2022. doi:10.1016/j.cjmeam.2022.100016.

[10] S. Nozawa and G. Serhat. Topology and fiber path optimization of composite structures: A critical review. *Materials & Design*, 251:113699, March 2025. doi:10.1016/j.matdes.2025.113699.

[11] J. Zhao, T. Qi, and C. Wang. Efficient GPU accelerated topology optimization of composite structures with spatially varying fiber orientations. *Computer Methods in Applied Mechanics and Engineering*, 421:116809, March 2024. doi:10.1016/j.cma.2024.116809.

[12] M. P. Schmidt, L. Couret, C. Gout, and C. B. W. Pedersen. Structural topology optimization with smoothly varying fiber orientations. *Structural and Multidisciplinary Optimization*, 62:3105–3126, July 2020. doi:10.1007/s00158-020-02657-6.

[13] Y. Huang, X. Tian, Z. Zheng, D. Li, A. V. Malakhov, and A. N. Polilov. Multiscale concurrent design and 3D printing of continuous fiber reinforced thermoplastic composites with optimized fiber trajectory and topological structure. *Composite Structures*, 285:115241, April 2022. doi:10.1016/j.compstruct.2022.115241.

[14] M. P. Bendsøe. Optimal shape design as a material distribution problem. *Structural Optimization*, 1:193–202, December 1989. doi:10.1007/bf01650949.

[15] T. W. Clyne and D. Hull. *An Introduction to Composite Materials*. Cambridge University Press, 2019. doi:10.1017/9781139050586.

[16] M. P. Bendsoe and O. Sigmund. *Topology optimization*. Springer, 2011. doi:10.1007/978-3-662-05086-6.

[17] B. Bourdin. Filters in topology optimization. *International Journal for Numerical Methods in Engineering*, 50:2143–2158, February 2001. doi:10.1002/nme.116.

[18] T. Nomura, A. Kawamoto, T. Kondoh, E. M. Dede, J. Lee, Y. Song, and N. Kikuchi. Inverse design of structure and fiber orientation by means of topology optimization with tensor field variables. *Composites Part B: Engineering*, 176:107187, November 2019. doi:10.1016/j.compositesb.2019.107187.

[19] J. Lee, D. Kim, T. Nomura, E. M. Dede, and J. Yoo. Topology optimization for continuous and discrete orientation design of functionally graded fiber-reinforced composite structures. *Composite Structures*, 201:217–233, October 2018. doi:10.1016/j.compstruct.2018.06.020.

[20] C. Luo, F. Ferrari, and J. K. Guest. Design optimization of advanced tow-steered composites with manufacturing constraints. *Composites Part B: Engineering*, 284: 111739, September 2024. doi:10.1016/j.compositesb.2024.111739.

[21] G. Li and Y. Chen. Highly efficient and extensible parallel topology optimization for continuous fiber reinforced composites. *Computer Methods in Applied Mechanics and Engineering*, 436:117713, March 2025. doi:10.1016/j.cma.2024.117713.

[22] N. Ichihara, M. Ueda, and T. Yokozeki. Penalized anisotropy: Controlling anisotropy growth in concurrent optimization of topology and fiber orientation for orthotropic composite materials. *Journal of Composite Materials*, 58:677–688, January 2024. doi:10.1177/00219983241230379.

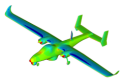
[23] X. Guo and K. Zhou. Topology optimization of continuous fiber-reinforced composites using Shepard interpolation and its design variable reduction. *Mathematics and Mechanics of Solids*, 29:1746–1763, June 2024. doi:10.1177/10812865241253520.

[24] A. Kawamoto, T. Matsumori, S. Yamasaki, T. Nomura, T. Kondoh, and S. Nishiwaki. Heaviside projection based topology optimization by a PDE-filtered scalar function. *Structural and Multidisciplinary Optimization*, 44:19–24, August 2010. doi:10.1007/s00158-010-0562-2.

[25] J. Wong and E. Kuhl. Generating fibre orientation maps in human heart models using poisson interpolation. *Computer Methods in Biomechanics and Biomedical Engineering*, 17(11):1217–1226, December 2012. doi:10.1080/10255842.2012.739167.

[26] K. Svanberg. The method of moving asymptotes—A new method for structural optimization. *International Journal for Numerical Methods in Engineering*, 24:359–373, February 1987. doi:10.1002/nme.1620240207.

[27] P. Pedersen. On optimal orientation of orthotropic materials. *Structural Optimization*, 1:101–106, June 1989. doi:10.1007/bf01637666.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

CALIBRATION OF TRANSITIONAL MODEL FOR LOW REYNOLDS NUMBER FLOWS AT HIGH ANGLES OF ATTACK USING SURROGATE-BASED OPTIMIZATION FRAMEWORK

Pranesh Chandrasekaran[1*], Dilan Kilic[1] and Melike Nikbay[1]

1: AeroMDO Lab, Faculty of Aeronautics and Astronautics
Istanbul Technical University
ITU Ayazaga Campus, 34469, Maslak/Istanbul
chandrasekaran@itu.edu.tr, kilicd15@itu.edu.tr, nikbay@itu.edu.tr

Abstract. *Growing popularity of micro-air vehicles (MAVs) has resulted in an increased interest in the study of aerodynamics at low Reynolds numbers (10^4–10^5). At such low* Re*, the lifting surfaces of these vehicles may experience regions of laminar, transitional, and turbulent flow, all of which could significantly impact their aerodynamic characteristics. Therefore, in order to obtain solutions that are adequately close to experimental results, it is imperative for any computational technique to model each of these regions accurately. The aim of the present work is to calibrate the model coefficients of the transitional γ-Re_θ model, for making improved predictions at high angles of attack, using a surrogate model-based optimization framework. The reliability of the computational procedure at low angles of attack, along with the challenges in achieving good agreement with experimental results at high angles of attack, is demonstrated by simulating the flow around a NACA 0012 airfoil at $Re = 50{,}000$, $Tu = 0.1\%$, and within the angle of attack range of $0\,^\circ \leq \alpha \leq 15\,^\circ$. A one-at-a-time sensitivity analysis is conducted to identify the ten most influential model coefficients that affect the performance of the transitional model. The identified model coefficients then serve as the design variables in the optimization studies, which aim to minimize the absolute difference between the numerical results from the present work and the experimental results available in literature. A surrogate-based optimization framework is developed to identify the optimal combination of model coefficients to enhance the model's prediction accuracy. Kriging-based approach is used for constructing surrogate models while a population-based Genetic Algorithm is used as the optimizer. An exploitation-based infill strategy is used to update the surrogate at every design iteration. Finally, the performance of the calibrated model is validated at different angles of attack that were not involved during the optimization process.*

Keywords: computational fluid dynamics, turbulence modeling, boundary layer transition, low Reynolds number, calibration, surrogate-based optimization

1 INTRODUCTION

The flow field around airfoils and wings at low Reynolds numbers (Re) is significantly influenced by viscous effects and is characterized by boundary layer features such as laminar separation, flow transition to turbulence, and reattachment, resulting in the formation of a Laminar Separation Bubble (LSB). Efficient modeling of these features plays a key role in accurately predicting the lift and drag forces at low Re. Additionally, low Re flows are known to be highly sensitivity to several parameters; hence, obtaining reliable data at low Re has been difficult both computationally and experimentally [1]. In particular, achieving good agreement with experimental results at high angles of attack remains to be challenging [2, 3]. Some of the notable efforts in low-cost transition modeling include the development of semi-empirical e^{N} model [4] and correlation-based transitional models such as the γ-Re_θ model [5]. The aim of the present work is to improve the predictive capabilities of the γ-Re_θ model at high angles of attack, particularly for cases post the bubble burst phenomenon. To overcome the limitations of the existing models, a novel approach is proposed that involves calibrating the model coefficients of the turbulence model using formal optimization techniques. While the calibration of turbulence models for specific applications is not a new concept, the present work distinguishes itself by integrating the model coefficients within a surrogate-based optimization (SBO) framework specifically tailored for a novel application of transitional low Re flows. What follows is the discussion of past works in the field of calibration of turbulence models.

Da Ronch et al. [6] used adaptive design of experiments and machine learning approach for calibrating the model coefficients of the Spalart-Allmaras model. The calibration was performed using the data generated by simulating the flow around RAE 2822 airfoil. They validated their calibrated model on several test cases, not limited to, such as ONERA M6 wing, axisymmetric transonic bump, and NACA 0012 airfoil. Employing machine learning techniques helped them mitigate the limitations associated with the manual fine-tuning process. Yildizeli and Cadirci [7] calibrated the k-ω SST model for accurately simulating the flow over a backward facing step. The skin friction distribution obtained from the DNS data available in the literature was used as the reference in the calibration process. OAT sensitivity analysis was carried out to identify the six most influential parameters to participate in the calibration process. They employed a deep neural network-based multi-objective optimization framework to improve the predictive accuracy of the k-ω SST model. Results revealed that the closure coefficients pertaining to both diffusion and production terms played a role in improving the accuracy of the model.

Barrouillet et al. [8] calibrated the transitional k-ω-γ-Re_θ turbulence model by customizing the constants in the intermittency and turbulent kinetic energy equations. They also replaced the local turbulent intensity with a more consistent representation. The calibration of the model was performed based on the data of well-known test cases such as the flat plates of T3 series and the NACA 0012 airfoil. NLF-0416 and S809 airfoil cases were used for validation purposes. Klavaris et al. [9] used neural network-based adjoint optimization framework to tune the generalized k-ω model for gas turbine combustion applications. The LES data available in the literature was used as the reference data in the calibration process. The results obtained using the calibrated model were observed to be in qualitative and quantitative agreement with both LES and experimental results available in the literature. Validation studies were carried out for different flow conditions, including a wide range of Re.

Zhang et al. [10] performed the uncertainty analysis and calibration of SST turbulence model for accurately capturing the development and the reattachment of the free shear layer in supersonic cavity-ramp flows. A non-intrusive polynomial chaos (NIPC) method was used for uncertainty quantification, while the Sobol indices were used to analyze the sensitivity of the results to the model parameters. Having identified the most influential model parameters, they employed a Bayesian inference approach to calibrate the SST model. The calibration process resulted in significant improvements in the predictive accuracy of the model.

Recently, integration of data-driven methods into turbulence modeling has gained significant interest. For instance, Shan et al. [11] employed data assimilation and machine learning techniques with deep neural networks for achieving satisfactory predictive accuracy in modeling largely separated flows at high Reynolds numbers. In their approach, the DNN model essentially maps the mean flow variables to the eddy viscosity. Their results indicated that the mean relative error of the lift coefficients in the stall region had decreased by over 57% for all the airfoils considered.

Ray et al. [12] used the Bayesian inverse problem formulation to calibrate the k-ϵ model for accurate jet-in-crossflow simulations. A total of three different model constants were used in the calibration process and the results revealed substantial improvement in the prediction of the entire flow field. Furthermore, the calibrated model was found to offer improvement for interactions at other Mach numbers and jet strengths. Mauro et al. [13] had calibrated the transitional turbulence model to accurately simulate the flow around airfoils for use in wind turbine applications. The calibration was carried out using a micro-genetic algorithm and is based on the data generated using the well-known NREL's S809 profile in the stall regime. The calibrated model parameters were then employed to test their performance in simulating the flow around NACA 0018 airfoil. The results were found to be satisfactory, ensuring the reliability of their calibration approach.

A detailed literature review clearly indicates that data-driven optimization approaches can play a vital role in improving the predictive capabilities of the turbulence models. The present study employs a Kriging-based surrogate model in combination with population-based genetic algorithms to arrive at optimal model parameters for accurately simulating the low Re flows at high angles of attack.

2 METHODOLOGY

The present work focuses on calibrating the closure coefficients of the transitional γ-Re_θ model to enhance its prediction accuracy at high angle of attack (α). The manuscript is organized as follows. Initially, a brief overview of the mathematical formulation of the transitional model is presented, followed by a comparison of numerical results from this study with existing experimental data in the literature. This comparison emphasizes the need for calibration at high α, where agreement is poor. Next, the results of a one-at-a-time sensitivity analysis are discussed, highlighting how variations in turbulence model closure coefficients impact the output quantities of interest, in this case the lift coefficients of the airfoil. Following this, the details of the surrogate-based optimization framework employed in the calibration process are presented. This framework is central to refining the closure coefficients to minimize discrepancies between numerical results and experimental data at transitional low Re regime. Lastly, the applicability of the calibrated model under various flow conditions that were not included in the optimization phase is assessed.

3 COMPUTATIONAL PROCEDURE

3.1 Transitional Model

In the present study, the numerical results are obtained by solving the governing equations of fluid flow, the Reynolds-Averaged Navier-Stokes (RANS) equations, using a finite volume method-based code, ANSYS Fluent. To model flow transition effectively, the SST k-ω model is used in conjunction with the γ-Re_θ model. As suggested by Menter et al. [5], the equations are solved using a second-order numerical scheme and in double-precision arithmetic. The correlation-based transition model used in this work is based on solving transport equations for intermittency (γ) and transition onset Reynolds number ($Re_{\theta t}$), and is given by

$$\frac{\partial(\rho\gamma)}{\partial t} + \frac{\partial(\rho U_j\gamma)}{\partial x_j} = P_\gamma - E_\gamma + \frac{\partial}{\partial x_j}\left[\left(\mu + \frac{\mu_t}{\sigma_f}\right)\frac{\partial\gamma}{\partial x_j}\right] \tag{1}$$

$$\frac{\partial(\rho\widetilde{R}e_{\theta t})}{\partial t} + \frac{\partial(\rho U_j\widetilde{R}e_{\theta t})}{\partial x_j} = P_{\theta t} + \frac{\partial}{\partial x_j}\left[\sigma_{\theta t}(\mu+\mu_t)\frac{\partial\widetilde{R}e_{\theta t}}{\partial x_j}\right]. \tag{2}$$

The transition and relaminarization sources in Eq.(1) are defined as:

$$P_\gamma = F_{\text{length}}c_{a1}\rho S[\gamma F_{\text{onset}}]^{0.5}(1-c_{e1}\gamma) \tag{3}$$

$$E_\gamma = c_{a2}\rho\Omega\gamma F_{\text{turb}}(c_{e2}\gamma - 1), \tag{4}$$

while the source terms in Eq.(2) are given by

$$P_{\theta t} = c_{\theta t}\frac{\rho}{t}(Re_{\theta t} - \widetilde{R}e_{\theta t})(1-F_{\theta t}). \tag{5}$$

For cases involving separation-induced transition the intermittency is modified as

$$\gamma_{\text{sep}} = \min\left(s_1 \max\left[0, \left(\frac{Re_v}{3.235Re_{\theta c}}\right) - 1\right]F_{\text{reattach}}, 2\right)F_{\theta t} \tag{6}$$

$$\gamma_{\text{eff}} = max(\gamma, \gamma_{\text{sep}}). \tag{7}$$

The model constants for these equations are:

$c_{a1} = 2.0$, $c_{e1} = 1.0$, $c_{a2} = 0.06$, $c_{e2} = 1.0$, $\sigma_f = 1.0$, $c_{\theta t} = 0.03$, $\sigma_{\theta t} = 2.0$, and $\mathrm{s}_1 = 2$.

The transition model interacts with the SST turbulence model [26] as follows:

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho u_j k)}{\partial x_j} = \tilde{P}_k - \tilde{D}_k + \frac{\partial}{\partial x_j}\left[(\mu + \sigma_k\mu_t)\frac{\partial k}{\partial x_j}\right] \tag{8}$$

$$\frac{\partial(\rho\omega)}{\partial t} + \frac{\partial(\rho u_j\omega)}{\partial x_j} = \alpha\frac{P_k}{v_t} - D_\omega + Cd_\omega + \frac{\partial}{\partial x_j}\left[(\mu+\sigma_k\mu_t)\frac{\partial k}{\partial x_j}\right] \tag{9}$$

$$\tilde{P}_k = \gamma_{\text{eff}}P_k; \quad \tilde{D}_k = \min(\max(\gamma_{\text{eff}}, 0.1), 1.0)D_k. \tag{10}$$

The model constants for these equations are:

$\alpha^* = 1$; $\alpha = 0.52$; $\beta^* = 0.09$; $a_1 = 0.31$; $\beta_1 = 0.075$; $\beta_2 = 0.0828$; $\sigma_{k1} = 1.176$; σ_{k2}= 1; $\sigma_{\omega1} = 2.0$; $\sigma_{\omega2}$= 1.168, and $c_1 = 10$.

For more detailed description about the formulation of the transitional model, the readers are directed to Langtry [14].

3.2 Flow Around the NACA 0012 Airfoil

The flow field around the NACA 0012 airfoil at $Re = 50{,}000$ and at a Tu of 0.1% was computed for a wide range of angles of attack (0°-15°) using the transitional γ-Re_θ model. The choice of these flow condition was influenced by the availability of experimental results from Ohtake et al. [15]. The computational grid employed in the present study is identical to the one used by Pranesh et al. [3], and is shown in Fig.1. The computational domain boundary is placed sufficiently far away from the airfoil surface to rule out any possible influence on the accuracy of the obtained results. A careful grid independence study, with three different grids namely G1 (51,120 cells), G2 (115,236 cells) and G3 (259,848), was conducted to determine the sufficient number of elements necessary for accurately predicting the flow field. Based on the results, the G2 grid with 115,236 cells was selected for all further computations, as it had a relatively low value of discretization error.

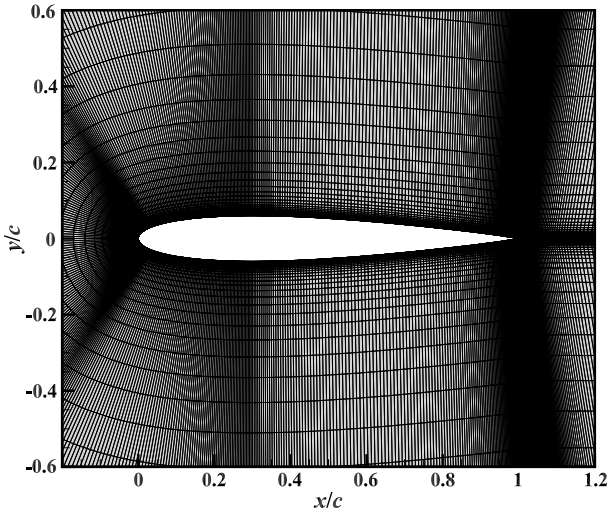


Figure 1: Mesh around the NACA 0012 airfoil.

The variation of lift coefficients as a function of angle of attack is illustrated in Fig.2, with experimental data from Ohtake et al. [15] also included for comparison. The results from the present study show good agreement with the experimental data up to $\alpha = 8°$. However, for $\alpha > 8°$, a significant deviation from the experimental data is observed. To investigate the reasons behind these discrepancies at high angles of attack, streamlines around the airfoil were examined, as depicted in Fig.3. Based on boundary layer behavior and separation characteristics, inferred through contours of axial velocity, intermittency and C_f distribution, several distinct flow regimes were identified on the upper surface of the airfoil (summarized in Tab.1). For angles between 0° and 1°, the flow exhibits laminar separation without transition or reattachment. As the angle of attack increases to 1.5°-3°, the separated laminar boundary layer undergoes flow transitions to turbulence but fails to reattach. At angles from 4° to 5°, the flow is characterized by laminar separation with transition and subsequent reattachment, forming a Laminar Separation Bubble (LSB). Further increasing the angle to $6° \leq \alpha \leq 8°$ introduces a second separated region close to the trailing edge, termed as trailing edge recirculation (TER). Finally, for angles ranging from 9° to 15°, the LSB merges with the TER, indicating a complex flow interaction. It is in this range of α the numerical results deviate substantially from experiments measurements.

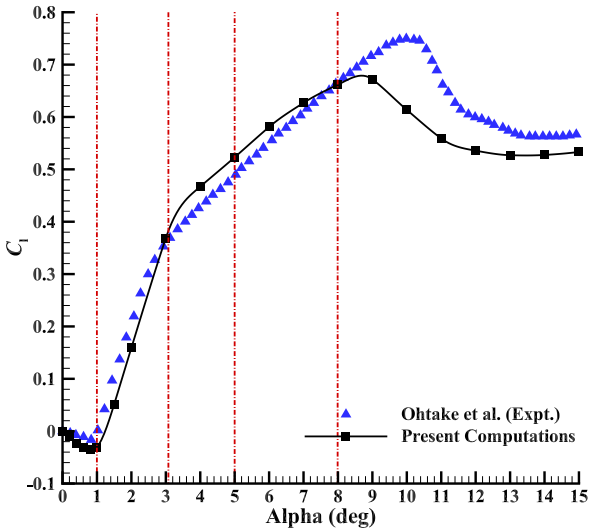
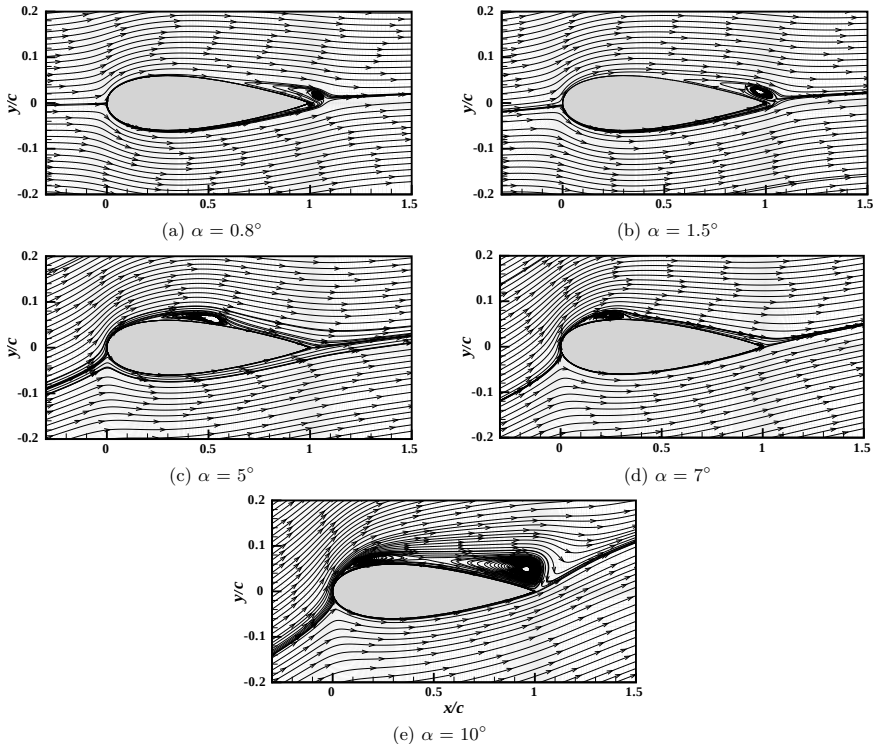


Figure 2: Comparison of lift coefficients of the NACA 0012 airfoil at $Re = 50{,}000$.

Table 1: Various flow modes observed over the upper surface of the NACA 0012 airfoil.

S. No.	Alpha Range	Flow Mode
1	$0° \leq \alpha \leq 1°$	Laminar separation without transition and without reattachment
2	$1.5° \leq \alpha \leq 3°$	Laminar separation with transition and without reattachment
3	$4° \leq \alpha \leq 5°$	Laminar Separation with transition and with reattachment Bubble (LSB)
4	$6° \leq \alpha \leq 8°$	LSB with Trailing Edge Recirculation (TER)
5	$9° \leq \alpha \leq 15°$	Merging of LSB and TER



(a) $\alpha = 0.8°$ (b) $\alpha = 1.5°$

(c) $\alpha = 5°$ (d) $\alpha = 7°$

(e) $\alpha = 10°$

Figure 3: Streamlines around the NACA 0012 airfoil at $Re = 50{,}000$ and $Tu = 0.1\%$, highlighting various flow regimes.

The discrepancies between numerical results and experimental data at high angles of attack have also been reported by Chen et al. [2] and Tank et al. [1]. These findings clearly highlight the need for enhancing the performance of the transitional model in this critical α regime. Improving the accuracy of this model is crucial for advancing the understanding of high-angle-of-attack aerodynamics in the transitional low Re regime. The present work aims to address this gap by calibrating the closure coefficients of the transitional model through formal optimization techniques.

4 ONE-AT-A-TIME SENSITIVITY ANALYSIS

As discussed in Section 3.1, the transitional SST model involves 19 closure coefficients. A one-at-a-time sensitivity analysis was conducted to determine the ten most influential coefficients that affect the model's performance. The value of lift coefficient is selected as the parameter of interest to facilitate this analysis. The analysis began by establishing a bound of ±15% around the default values for all 19 model coefficients. A total of 38 CFD simulations were then performed to assess how variations in each closure coefficient affected the value of lift coefficient. Figure 4 illustrates the results of this sensitivity analysis. Constants appearing in the transport equations for intermittency (γ) and $Re_{\theta t}$ demonstrated minimal impact on lift coefficient values, whereas the closure coefficients of the SST k-ω model showed significant influence. Among the nine coefficients from the transport equations of γ and $Re_{\theta t}$, only the constant s_1, which appears in Eq.6, exhibited notable sensitivity. Based on these findings, the ten most influential parameters were selected as design variables for the subsequent optimization procedure.

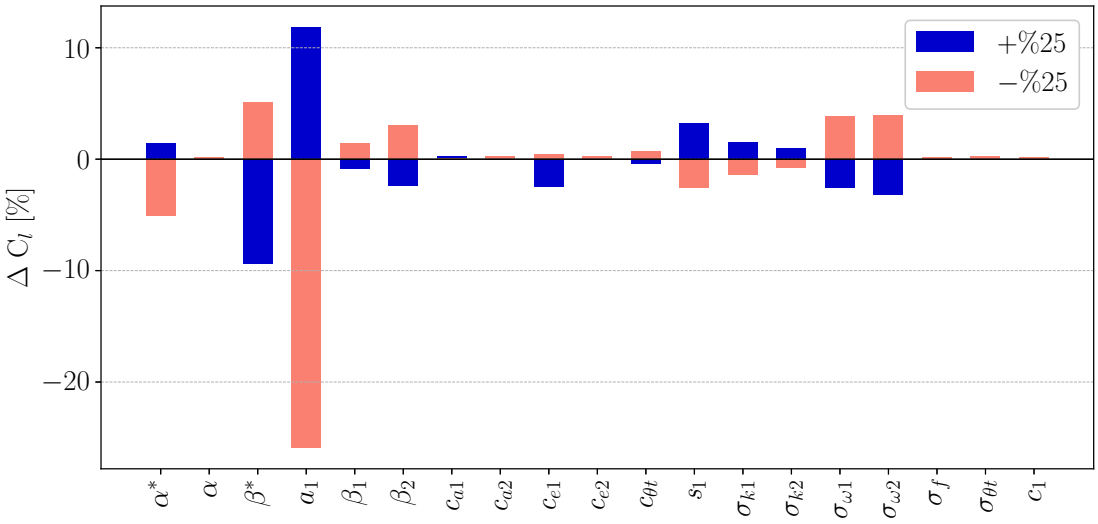


Figure 4: Sensitivity analysis results of lift coefficient parameter.

4.1 Calibration of Closure Coefficients

In the present work, calibration of transitional model's closure coefficients is achieved using a surrogate-based optimization approach that minimizes the absolute difference between experimental and numerical lift coefficients. An angle of attack of 10° is chosen as the design point, where the discrepancy between numerical and experimental values

is most significant; see Fig.2. The ten most influential model coefficients identified in the previous section are chosen as design variables, with the absolute difference between experimental and numerical results serving as the objective function. Figure 5 illustrates the general framework of a surrogate-based optimization process. The main intention of using such an approach is to create a low-cost approximation of the expensive true objective function by building a quick-to-evaluate surrogate model. This in turn enables the application of exhaustive global search techniques directly on the surrogate model, thereby making the optimization process to be highly cost effective. In the present work, a Kriging-based approach is used to construct the surrogate model, with a population-based Genetic Algorithm (GA) serving as the optimizer. Brief description regarding the Kriging surrogate and GA implementation are as follows:

4.2 Kriging

In this study, the computationally expensive objective function is approximated using a Kriging-based surrogate model [16], which approximates the response as a combination of known functions and a realization of a stochastic process:

$$\hat{y}(\mathbf{x}) = \sum_{i=1}^{k} \beta_i f_i(\mathbf{x}) + Z(\mathbf{x}), \tag{11}$$

where β_i are regression coefficients, and $Z(\mathbf{x})$ is a Gaussian process with zero mean and spatial covariance. The covariance between two samples is defined as:

$$\text{cov}\left[Z(\mathbf{x}^{(p)}), Z(\mathbf{x}^{(q)})\right] = \sigma^2 R(\mathbf{x}^{(p)}, \mathbf{x}^{(q)}), \tag{12}$$

where σ^2 is the process variance and $R(\mathbf{x}^{(p)}, \mathbf{x}^{(q)})$ is the correlation function. In this work, the squared exponential (Gaussian) correlation is used:

$$R(\mathbf{x}^{(p)}, \mathbf{x}^{(q)}) = \prod_{i=1}^{n_z} \exp\left(-\theta_i \left(x_i^{(p)} - x_i^{(q)}\right)^2\right), \tag{13}$$

where θ_i are the hyper-parameters controlling the influence of each input variable. These hyper-parameters are estimated via maximum likelihood estimation. All input variables are normalized prior to model fitting to improve numerical stability.

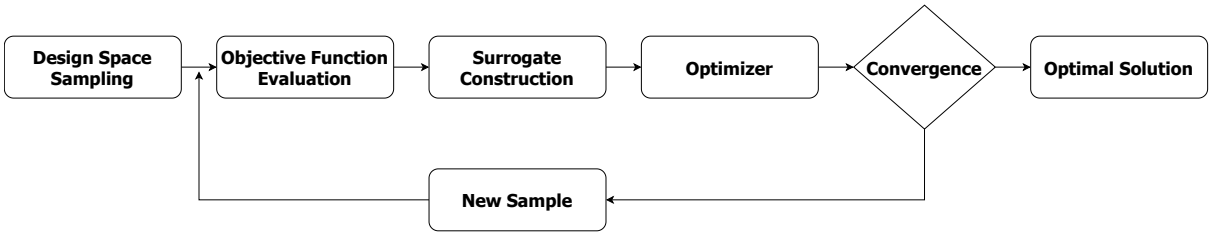


Figure 5: SBO framework.

4.3 Genetic Algorithm

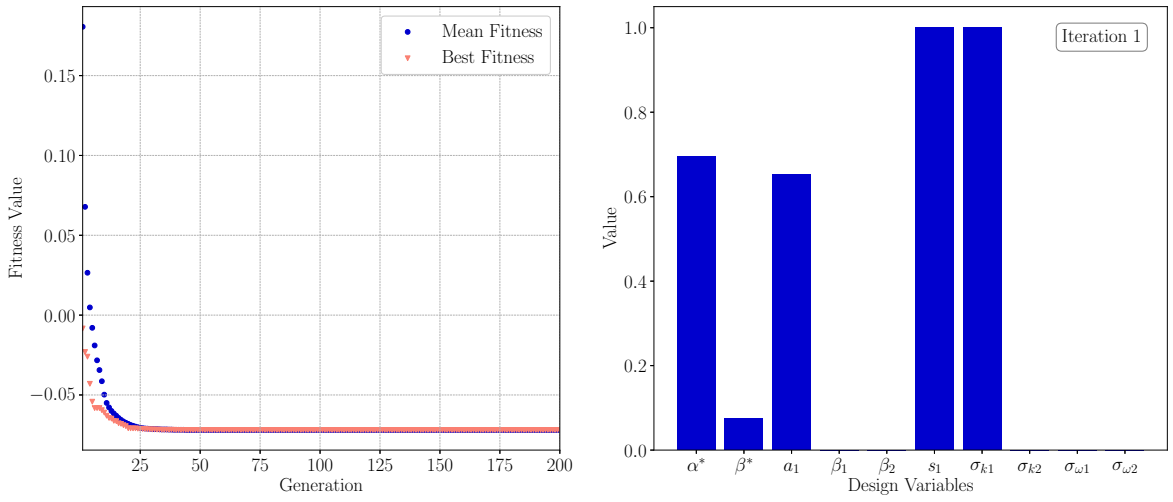
Genetic Algorithm (GA) is integrated with the Kriging model to perform design optimization. It is a population-based stochastic search method inspired by natural selection.

It operates by evolving a population of candidate solutions through selection, crossover, and mutation operators. The surrogate-based optimization loop iteratively evaluates the Kriging model predictions to guide the GA toward the optimal design, greatly reducing the computational cost associated with the optimization procedure.

4.4 Implementation of Surrogate-Based Optimization

The optimization process began by establishing a bound of ±25% around the default values for all ten selected model coefficients. An initial sample plan of 100 points within the design space was generated using an optimized Latin hypercube sampling technique. High-fidelity CFD simulations were conducted at these sample points to evaluate the true objective function values, which were then used to construct a Kriging-based surrogate model. The GA was subsequently exercised on this surrogate to identify the optimal solution. A high-fidelity CFD simulation is conducted at the identified optima to determine the true objective function value. This point is then appended to the initial samples and the surrogate model is reconstructed followed by the estimation of new optima. Such a process is often referred to as exploitation-based infill process. The entire sequence was repeated until the true objective function value does not change significantly between two consecutive design iterations.

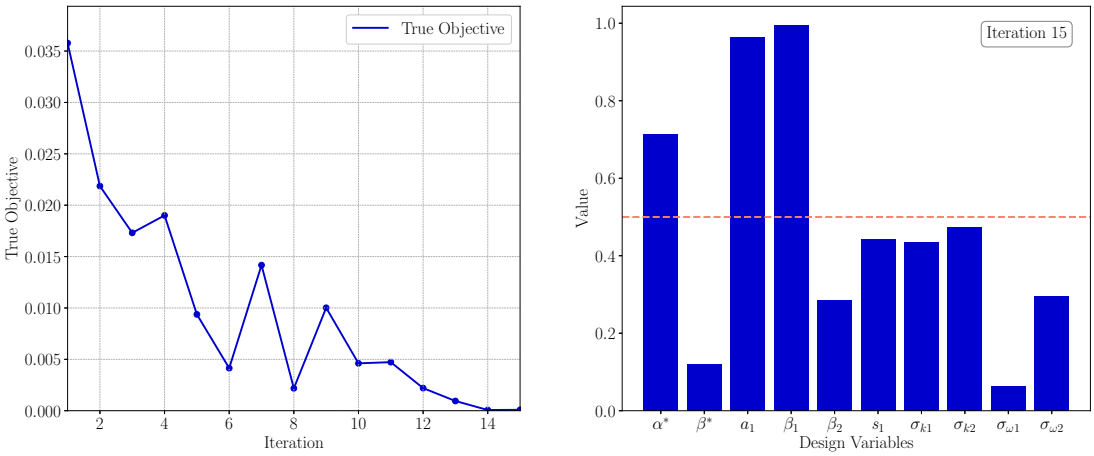
The quality of the initial surrogate was estimated using a cross-validation technique and quantified using an r^2 metric. The correlation coefficient r^2 of the initial surrogate was estimated to be 0.96, which was deemed fit for optimization purposes. The optimization process began with a random initial population of 200 individuals. These individuals are then allowed to evolve and progress to subsequent generations. At each generations, the individuals were ranked based on their fitness functions. A function tolerance of $1e^{-6}$ was set as the convergence criteria. Figure 6a illustrates the evolution of the objective function based on the implementation of GA on the constructed initial surrogate, while Fig.6b displays the value of optimal variables. It should be noted that the design variables are presented on a normalized scale of 0 to 1. This completes the first design iteration.



(a) (b)

Figure 6: a) Implementation of GA on the surrogate model and b) optimal solution.

The optimization process was continued using an exploitation-based infill strategy. Figure 7a illustrates the convergence of the true objective function value at the surrogate-predicted optima as a function of design iterations, while Fig.7b presents the value of model coefficients of the final optimal candidate. As observed in Fig.7a, the optimization procedure terminated after 15 iterations, indicating that a total of 115 high-fidelity evaluations were required to arrive at the final optimal solution. The objective function value and lift coefficient estimated through CFD simulation at the final optimal solution were $8.76e^{-5}$ and 0.7496, respectively. These results demonstrates excellent agreement with the experimental lift coefficient of 0.7495 at the chosen angle of attack.



(a) (b)

Figure 7: a) Evolution of true value of the objective function with respect to design iteration and b) final optimal solution.

Figure 8 presents the performance of the calibrated model at various angles of attack surrounding the optimization design point. It may be inferred that the current optimization process has resulted in an improved prediction capability of the transitional model.

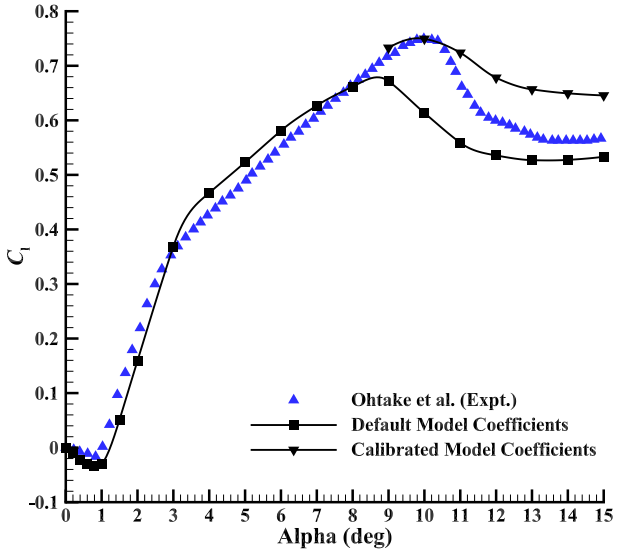


Figure 8: Performance of the calibrated model: comparison of lift coefficients.

Notably, it correctly identifies the stall angle of attack. The linear regime of the C_l-α curve is also extended and shows good agreement with experimental data. However, the discrepancies between numerical and experimental results persists at off-design conditions (i.e. $\alpha \geq 12°$). This suggests that while the current optimal solution performs well up to the design point, there exists a scope for further refinement to enhance accuracy across a wider range of angles of attack. Such improvements could be achieved through multi-point optimization involving several design points (for e.g., 10°, 15°, and 20°). It should be noted that increasing the number of design points would necessitate a substantial increase in the number of initial design samples required to build a reasonably accurate surrogate model. To address this computational challenge, a multi-fidelity based approach can be implemented in future work to extend the current methodology.

5 CONCLUSIONS

This study presented a systematic approach for calibrating transitional model closure coefficients with specific focus on improving predictions at high angles of attack, where the transitional model with default model coefficients typically resulted in poor agreement with experimental data. Through a surrogate-based optimization technique, the present methodology successfully reduced the discrepancy between numerical predictions and experimental measurements at the critical design point ($\alpha = 10°$), achieving a lift coefficient of 0.7496 compared to the experimental value of 0.7495. Sensitivity analysis revealed that while most coefficients in the intermittency and the momentum thickness Reynolds number transport equations had minimal impact on the lift coefficients, several SST k-ω model parameters significantly influenced the results. Particularly, the optimized model demonstrated substantial improvement in accurately predicting the stall angle of attack. The implementation of Kriging surrogates in conjunction with Genetic Algorithm proved to be computationally efficient, requiring only 115 high-fidelity CFD simulations to arrive at the optimal solution. While the calibrated model performs exceptionally well near the design point, further enhancement could be achieved through multi-point optimization strategies incorporating a broader range of flow conditions. The present work establishes a foundation for developing more reliable transitional models that balance computational efficiency with prediction accuracy for complex aerodynamic configurations at challenging flight regimes.

REFERENCES

[1] J. Tank, L. Smith, and G. Spedding. On the possibility (or lack thereof) of agreement between experiment and computation of flows over wings at moderate Reynolds number. *Interface focus*, 7(1):20160076, 2017.

[2] Z. Chen, N. Qin, and A. Nowakowski. Three-dimensional laminar-separation bubble on a cambered thin wing at low Reynolds numbers. *Journal of Aircraft*, 50(1):152–163, 2013.

[3] C. Pranesh, M. Sivapragasam, and H. K. Narahari. Multi-fidelity aerodynamic shape optimization of an airfoil at transitional low Reynolds number. In *National Conference on Multidisciplinary Design, Analysis, and Optimization, IISc, Bangalore*, 2018.

[4] J. Van Ingen. The e^N method for transition prediction. Historical review of work at TU delft. In *38th Fluid Dynamics Conference and Exhibit*, page 3830, 2008.

[5] F. R. Menter, R. B. Langtry, S. Likki, Y. B. Suzen, P. Huang, and S. Völker. A correlation-based transition model using local variables—part I: model formulation. *Journal of turbomachinery*, 128(3):413–422, 2006.

[6] A. Da Ronch, M. Panzeri, J. Drofelnik, and R. d'Ippolito. Sensitivity and calibration of turbulence model in the presence of epistemic uncertainties. *CEAS Aeronautical Journal*, 11(1):33–47, 2020.

[7] A. Yildizeli and S. Cadirci. Calibration of the k-ω sst turbulence model for backward facing step problem using multi-objective optimization. In *ASME International Mechanical Engineering Congress and Exposition*, volume 87660, page V009T10A031. American Society of Mechanical Engineers, 2023.

[8] B. Barrouillet, É. Laurendeau, and H. Yang. Calibration of the transitional k-ω-γ-$Re_{\theta t}$ turbulence model. *AIAA Journal*, 60(7):4140–4148, 2022.

[9] G. Klavaris, M. Xu, S. Patwardhan, I. Verma, S. Orsino, and P. Nakod. Tuning of generalized k-omega turbulence model by using adjoint optimization and machine learning for gas turbine combustor applications. In *Turbo Expo: Power for Land, Sea, and Air*, volume 86960, page V03BT04A049. American Society of Mechanical Engineers, 2023.

[10] K.-l. Zhang, Y.-t. Zhao, Q. Wang, J.-p. Li, F.-z. Zeng, and C. Yan. Uncertainty analysis and calibration of sst turbulence model for free shear layer in cavity-ramp flow. *Acta Astronautica*, 192:168–181, 2022.

[11] X. Shan, Y. Liu, W. Cao, X. Sun, and W. Zhang. Turbulence modeling via data assimilation and machine learning for separated flows over airfoils. *AIAA Journal*, 61(9):3883–3899, 2023.

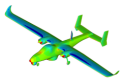
[12] J. Ray, S. Lefantzi, S. Arunajatesan, and L. Dechant. Bayesian parameter estimation of a $k-\varepsilon$ model for accurate jet-in-crossflow simulations. *AIAA Journal*, 54(8):2432–2448, 2016.

[13] S. Mauro, R. Lanzafame, M. Messina, and D. Pirrello. Transition turbulence model calibration for wind turbine airfoil characterization through the use of a micro-genetic algorithm. *International Journal of Energy and Environmental Engineering*, 8:359–374, 2017.

[14] R. B. Langtry. A correlation-based transition model using local variables for unstructured parallelized cfd codes. 2006.

[15] T. Ohtake, Y. Nakae, and T. Motohashi. Nonlinearity of the aerodynamic characteristics of naca0012 aerofoil at low Reynolds numbers. *Japan Society of Aeronautical Space Sciences*, 55(644):439–445, 2007.

[16] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. Martins. A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, 135:102662, 2019.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

PRECONDITIONING TECHNIQUES FOR GENERAL CONSTRAINED OPTIMIZATION PROBLEMS: APPLICATION TO INDUSTRIAL AEROSPACE COMPONENT DESIGN

Shouvik Bandopadhyay[1], François Gallard[1], and Simone Coniglio[1]

1: Methods and Tools for the Development of Complex Systems
MDO Competence Centre
IRT Antoine de Saint Exupéry
CS34436, 3 Rue Tarfaya, 31400 Toulouse
shouvik.bandopadhyay@irt-saintexupery.com

Abstract. *This research addresses a typical source of inefficiency in Multidisciplinary Design and Optimization (MDO) when the problems are highly non-linear. Ill-conditioning of the objective function or constraints' Hessian matrices is known to hinder the convergence of optimizers and increase the computational cost of finding the solution. It leads to well-known zigzag convergence paths, even for quasi-Newton methods, because the typical BFGS Hessian approximations have difficulties in capturing the range of eigenvectors and eigenvalues, resulting in excessive intermediate evaluation points. Preconditioning, a method of transforming numerical problems to improve stability and convergence, has been extensively developed for solving linear algebraic systems. Optimization problems, such as constrained multidisciplinary design optimization (MDO), can be viewed as generalizations or extensions of these linear systems (eg. via residual minimization). However, despite this close relationship, the use of preconditioning in constrained MDO remains relatively under-explored. This study compares two preconditioning strategies: a static external preconditioner, built from previous optimizations and applied before the optimization starts, and a dynamically updated preconditioner, computed during sub-optimization cycles of the Augmented Lagrangian method. They are solved using the P-L-BFGS-B algorithm, an extension of the L-BFGS-B optimizer for unconstrained preconditioned problems, and the SNOPT algorithm, which allows black-box preconditioning thanks to its exact linear constraint management. The methods are first validated on analytical test cases. Further validation is performed on a representative industrial application, the sizing optimization of the wing structure of the Airbus XRF-1 research aircraft configuration, where significant speed-up of 66.8% and 60% were observed for the dynamic preconditioning strategy and the static external preconditioning strategy respectively. These methods have been integrated into GEMSEO, an open-source MDO library, making them available for wider industrial and research applications.*

Keywords: optimization, Preconditioning, Algorithms, Aerospace, Industrial MDO

1 INTRODUCTION

Multidisciplinary Design Optimization (MDO) is a powerful technique for supporting and achieving industrial design through the use of physics-based simulations that capture the necessary details of real-world phenomena. These simulations are typically integrated into MDO frameworks that cover all steps of the design process [1, 2]. A typical overview of the aero-structural optimization can be found in Fig.1.

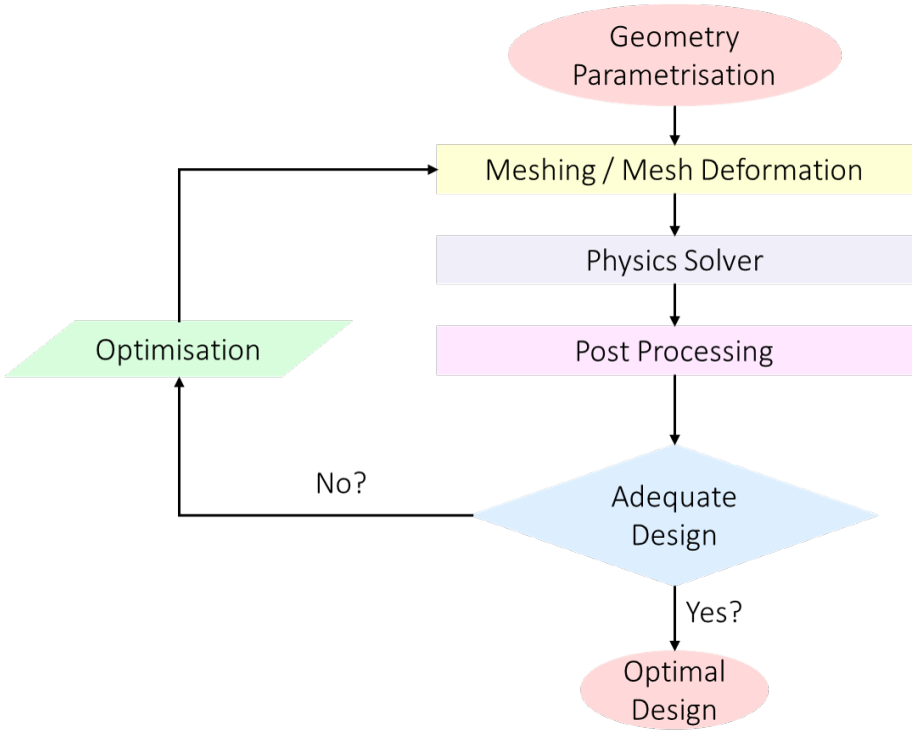


Figure 1: Conventional MDO workflow for industrial component design

The objective function(s) of this minimization problem is subject to both equality and inequality constraints, and to bound constraints on the design variables, as in Eq.(1).

$$\begin{aligned} \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{subject to:} \quad & \\ & g_i(\mathbf{x}) \leq 0, \quad i = 1, \ldots, m, \\ & h_j(\mathbf{x}) = 0, \quad j = 1, \ldots, p, \\ & a \leq \mathbf{x} \leq b. \end{aligned} \tag{1}$$

In the above Eq.(1), it is assumed that the functions f, g, & h are continuously differentiable at least once, computationally expensive to evaluate, have efficiently computable gradients, and separable [3]. Various robust optimization algorithms exist to solve such problems, and quasi-Newton algorithms [4] are widely used due to their ability to exploit gradient information and approximating the Hessian to converge to the local minimum more rapidly [5, 6]. However, when this Hessian matrix of the objective function or constraints is ill-conditioned, the convergence is slowed down. This phenomenon can be

observed on convergence history plots: in case of an ill-conditioned Hessian matrix, a typical zigzag path is observed from the initial design point to the optimal design point, with many extra intermediate points that are not smoothly aligned. Each of these extra intermediate points corresponds to a new design that triggers the full cycle of the MDO workflow, i.e., each step in the workflow is individually executed for every new design the optimizer proposes. As a result, it makes the MDO workflow computationally intensive. The presence of additional constraints further adds to the convergence difficulties.

The condition number of the Hessian at the optimum can be used to quantify the ill-conditioning of the optimization problem. The condition number is the ratio of the highest to the lowest eigenvalue of the Hessian matrix of the objective function, which can also be defined geometrically as the ratio of the highest to the lowest curvature. A high condition number generates the zigzag phenomenon in the optimization history plot as described above. The closer the condition number is to 1, the more well-conditioned the optimization problem. Similarly to linear system resolution, a 'preconditioning operator' can be applied to the entire optimization problem to reduce the otherwise high condition number of objective functions [7]. Preconditioning transforms $\mathbf{x}$ into $\mathbf{y}$ using $\mathbf{x} = \mathbf{P}\mathbf{y}$, where $\mathbf{P}$ is the preconditioning operator, usually a matrix, but could also be a non-linear operator. The optimization problem can then be represented as Eq.(2).

$$\begin{aligned} &\min_{\mathbf{y}} \quad f(\mathbf{P}\mathbf{y}) \\ &\text{subject to:} \\ &\qquad g_i(\mathbf{P}\mathbf{y}) \leq 0, \quad i = 1, \ldots, m, \\ &\qquad h_j(\mathbf{P}\mathbf{y}) = 0, \quad j = 1, \ldots, p, \\ &\qquad a \leq \mathbf{P}\mathbf{y} \leq b. \end{aligned} \tag{2}$$

This transformation aims at improving the numerical stability by reducing the condition number and, by consequence, reduces the number of intermediate points between the starting point and the optimum. An illustrative summary of the preconditioning process can be found in Fig.2.

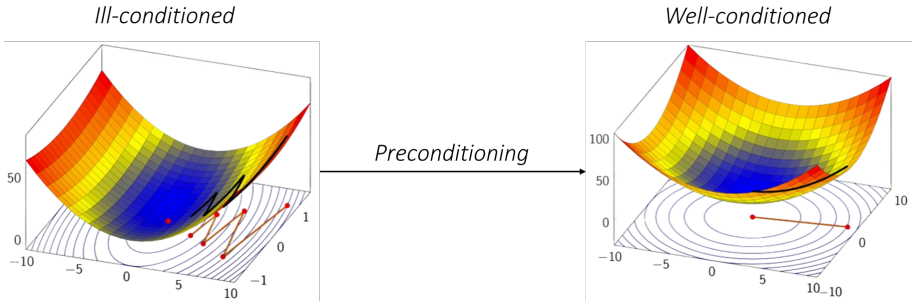


Figure 2: Illustration of the concept of preconditioning

Several methods exist to construct preconditioning operators, each with their own advantages and limitations. Incomplete factorization [8] and sparse approximate inverses [8] offer balanced trade-off between efficiency and robustness but struggle with scalability and sensitivity to parameter tuning. Multi-grid and algebraic multilevel methods [9] effectively handle elliptic PDE problems with optimal convergence but can be challenging

to generalize or parallelize efficiently. Domain decomposition methods [10] are scalable and suitable for parallel computing yet depend heavily on effective coarse-grid correction strategies. Hessian-based operators [10] exploit matrix structure and allows modular updates, performing best when constraints are fewer than variables. Selecting between structured and algebraic approaches remains crucial, balancing performance, general applicability, and implementation complexity.

Although preconditioning has shown promise in improving optimization efficiency, its application to constrained problems remains under explored, particularly in the context of industrial-scale Multidisciplinary Design optimization (MDO). In industrial MDO, problems often involve large-scale non-linear systems with heterogeneous constraints, coupling variables across disciplines, and complex interactions that create highly ill-conditioned and structurally diverse linear systems. Effective preconditioning typically exploits clear mathematical properties of the Hessian, such as sparsity or symmetry; however, in multidisciplinary design optimization (MDO), the Hessian structure is often complex, making it difficult to develop efficient preconditioning operators. Additionally, in the presence of general constraints, it is unsure which Hessian should be exploited to construct the preconditioner - Objective function's Hessian? Hessian of the Lagrangian? Or Hessian of the constraint functions? Furthermore, industrial-scale MDO problems frequently exhibit dynamic active sets because of constraints that become active/inactive throughout the process of optimization. These changes to the constraints alter the structure and condition number of the overall problem and, thus, further complicate the construction of stable and scalable preconditioners that account for the constraints, such as the ones based on the Lagrangian. Hence, tailored yet generalizable preconditioning strategies remain challenging and understudied. Particularly, there is a high potential for such strategies to construct preconditioning operators in multi-level MDO, solving the overall optimization problem involves addressing numerous disciplinary sub-optimization problems, often requiring repetitive sub optimizations across similar problem structures. This provides a significant opportunity to reuse information, such as gradients, Hessians, or preconditioning matrices, between runs to improve efficiency.

In this research, the authors bridge the gap mentioned above by exploring novel preconditioning strategies that handle general constraints. The present work is structured as follows:

- A brief explanation of the P-L-BFGS-B algorithm [11], an adaptation of the classical L-BFGS-B algorithm [6] to handle optimization with bound constraints in preconditioned spaces is presented.
- Two preconditioning strategies for general constrained optimization problems are proposed and compared:
  (a) a static external preconditioner built from previous optimizations and applied before the optimization process starts. SNOPT algorithm [12], which allows exact linear constraint management, is used to solve the optimization problem as a sequence of sub-problems.
  (b) a dynamically updated preconditioner, computed during sub-optimization cycles of the Augmented Lagrangian Method [13], where the P-L-BFGS-B algorithm is used to solve the optimization problem.

- The proposed strategies are validated on analytic benchmark optimization problems with non-linear inequality constraints and on a aircraft wing sizing optimization test case using the generic research configuration of the Airbus XRF-1 transport aircraft as reference geometry.

This enhanced framework aims to provide a clearer understanding of how preconditioning can be effectively integrated into complex optimization workflows, thus streamlining the design process in aerospace engineering and other fields.

2 METHODOLOGY

This section details the methods and processes followed in the direction of the development of a preconditioning framework for general constrained optimization problems in the frame of the R-EVOL [14] project.

2.1 Overview of the P-L-BFGS-B algorithm

In a nutshell, P-L-BFGS-B accelerates bound-constrained optimization by transforming the space, using prior Hessian approximations to guide convergence, thus reducing the number of expensive objective evaluations required. Extensive formulations and assumptions of the P-L-BFGS-B algorithm can be found in an earlier work of Gurol et. al.[11], that has been embedded in a complex aero-structure aircraft MDO process [15]. The P-L-BFGS-B algorithm is an enhancement of the L-BFGS-B algorithm to work in a preconditioned space defined by a user defined matrix change of variables [11]. A brief overview of the algorithm is provided in Fig.3.

2.2 Static Preconditioning with SNOPT method

It is reiterated that the optimization problem can be transformed to the "preconditioned space" by performing a linear transformation $\mathbf{x} = \mathbf{P}\mathbf{y}$ as shown in Eq.(2). The $\mathbf{P}$ is the preconditioning matrix, and the lower and upper bounds $\mathbf{a}$ & $\mathbf{b}$ are applied to the transformed variable $\mathbf{P}\mathbf{y}$ instead of $\mathbf{y}$ alone. The problem can be further simplified as an objective function with general constraints and linear constraints, as illustrated in Eq.(3). It should be noted that $\mathbf{P}$ is obtained from an earlier optimization of a similar problem belonging to the same family (from previous iterations of a bi-level MDO process for instance), to ensure a similar Hessian structure. Typically, in [11], a BFGS approximation square root inverse is built from the optimization history. In this work, to extend the approach to the general constraint, the BFGS approximation is built on the gradient of the Lagrange function. The transformation is applied before the optimization process begins.

$$
\begin{aligned}
&\min_{\mathbf{y}} \quad f(\mathbf{P}\mathbf{y}) \\
&\text{subject to:} \\
&\qquad g_i(\mathbf{P}\mathbf{y}) \leq 0, \quad i = 1, \ldots, m, \\
&\qquad h_j(\mathbf{P}\mathbf{y}) = 0, \quad j = 1, \ldots, p, \\
&\qquad a - \mathbf{P}\mathbf{y} \leq 0, \\
&\qquad \mathbf{P}\mathbf{y} - b \leq 0.
\end{aligned}
\tag{3}
$$

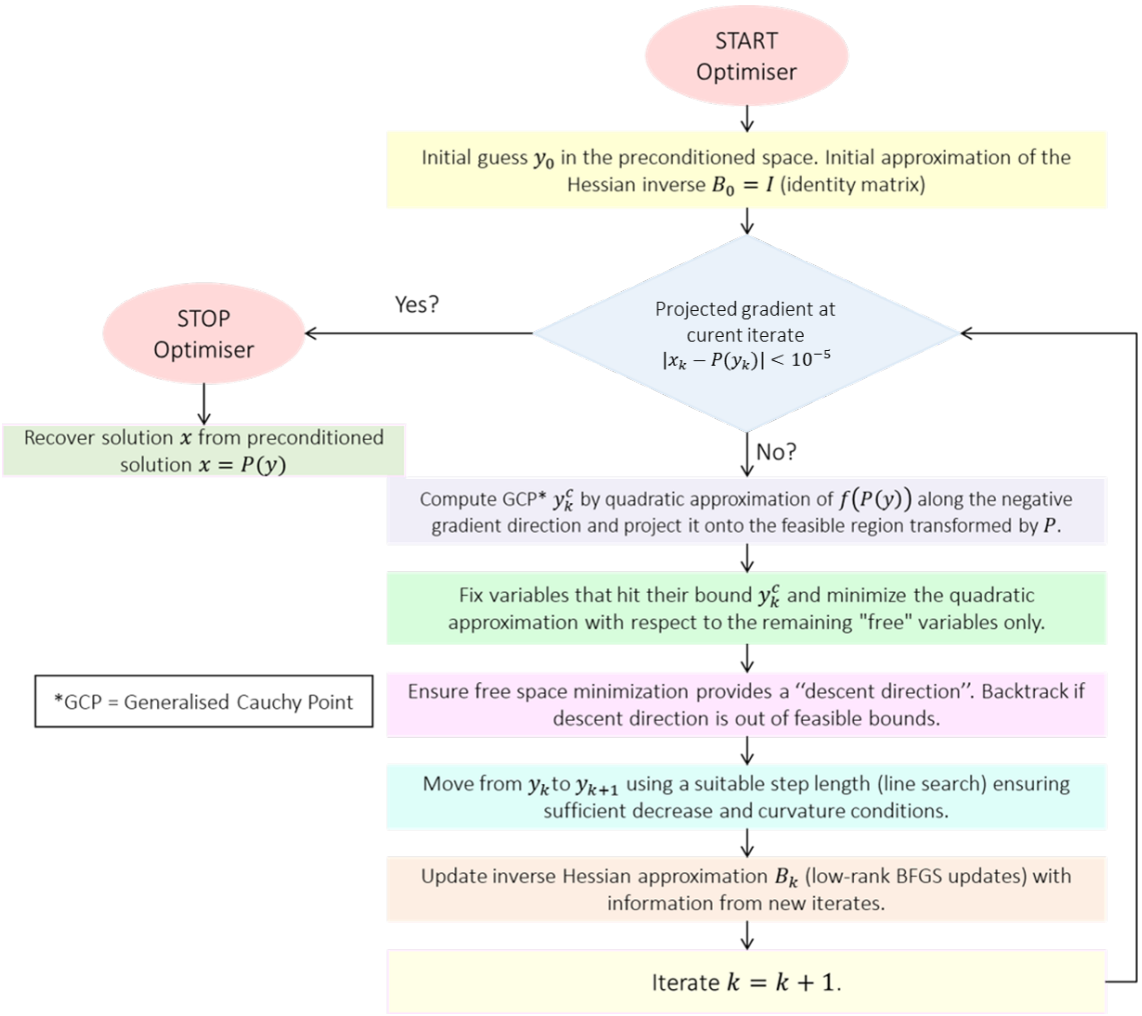


Figure 3: P-L-BFGS-B Algorithm Summary

SNOPT algorithm [12] is able to efficiently solve a class of problems similar to Eq.(3) by converting them into a series of simpler sub-problems. It employs quadratic models to first approximate the objective function and the constraints & then solves these approximations iteratively using Sequential Quadratic Programming (SQP). Similarly to P-L-BFGS-B, SNOPT maintains sparse Hessian approximations using limited-memory quasi-Newton updates, making it scalable for large-scale problems. In contrast to P-L-BFGS-B, the SQP strategy allows SNOPT to handle inequality and equality constraints directly. It should be noted that $\mathbf{P}$ remains unchanged across all the sub-problems until convergence is achieved.

It is crucial to acknowledge that solving quadratic approximations inherently introduces approximation errors that affect the accuracy of intermediate solutions. Nevertheless, the iterative nature of the SQP method typically ensures convergence to an accurate solution despite these intermediate approximations. Additionally, the use of a fixed preconditioner matrix $\mathbf{P}$ throughout the iterative process can introduce further approximation errors. If the preconditioner does not accurately capture the essential structure or properties of the Hessian, this mismatch can degrade the quality of the quadratic approximations and potentially slow convergence or reduce the accuracy of the final solution. Therefore, it is of utmost importance that the static preconditioner matrix $\mathbf{P}$ be constructed keeping

the above concerns in mind. This can be accomplished through the implementation of robust initialization techniques, regularization methods, and continuous monitoring of the optimization convergence. An alternative approach involves the utilization of a filter to incorporate spectral information, such as eigenvectors and Ritz vectors, thereby ensuring the accurate capture of critical modes. A summary of the static preconditioning method can be found in Fig.4.

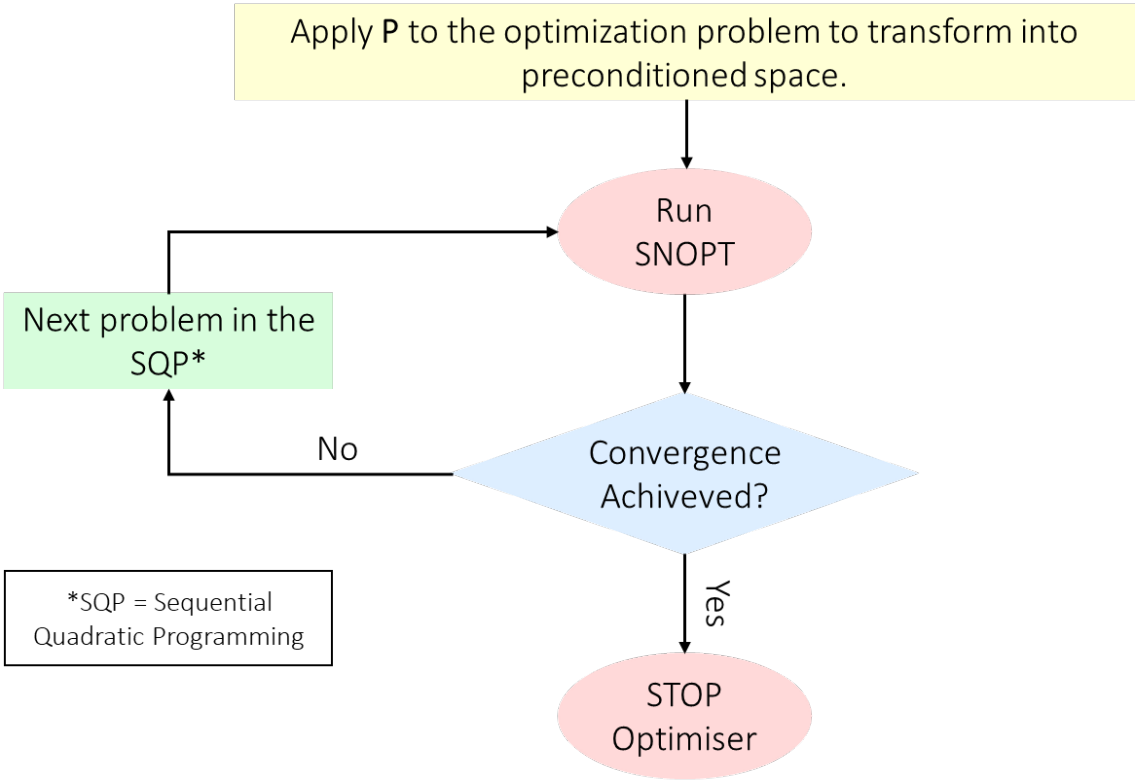


Figure 4: External Static Preconditioning with SNOPT Algorithm

2.3 Dynamic Preconditioning with Augmented Lagrangian Method

P-L-BFGS-B method can only solve optimization problems with bound constraints. It does not have an explicit mechanism to handle the general constraints as seen in the Eq.(1). The Augmented Lagrangian Method [13], hereby abbreviated as ALM for the rest of the paper, can convert it into a bound constraint optimization problem as seen in Eq.(4).

$$\begin{aligned} \text{Minimize} \quad & f(\mathbf{x}) + \sum_{j=1}^{p} \lambda_j h_j(\mathbf{x}) + \frac{\rho}{2} \sum_{j=1}^{p} h_j(\mathbf{x})^2 + \frac{\rho}{2} \sum_{i=1}^{m} \left[\max\left(0, g_i(\mathbf{x}) + \frac{\mu_i}{\rho}\right) \right]^2 \\ \text{w.r.t.} \quad & \mathbf{x} \\ \text{subject to} \quad & a_k \leq x_k \leq b_k, \quad k = 1, \ldots, n \end{aligned} \tag{4}$$

In the event that the constraints are violated, the objective is increased or "augmented" by adding a penalty term. In the above Eq.(4), ρ is the penalty parameter that is iteratively updated. The $\lambda_\mathbf{j}$ and $\mu_\mathbf{i}$ are Lagrange multipliers for the equality and inequality

constraints, respectively. The function of these multipliers is to aid in constraint satisfaction without the need for excessive penalties. These multipliers are iteratively updated by ALM, ensuring rapid convergence to a solution that precisely satisfies all constraints.

Similar to the SNOPT algorithm, ALM too converts the optimization problem into a series of smaller sub-optimization problems but, the sub-problems created by ALM are unconstrained. Here, the preconditioner matrix is constructed $\mathbf{P}$ by using the Hessian matrix of the first sub-optimum. Thereafter, this $\mathbf{P}$ is applied to the next sub-optimization problem and an updated $\mathbf{P}$ is obtained at each sub-optimum, until convergence is achieved. It is different from the external static preconditioning method in the sense that the $\mathbf{P}$ matrix changes dynamically during the optimization process. The central idea behind changing the preconditioner dynamically is to continuously adapt to the changes in the problem curvature (Hessian) and constraints structure arising due to turning on/off of constraints during the optimization, and thereby, eliminating sensitivity to initial Hessian approximation, reduce approximation errors, improve stability and convergence. A summary of the method used to construct this so-called 'dynamic preconditioner' is illustrated in Fig.5.

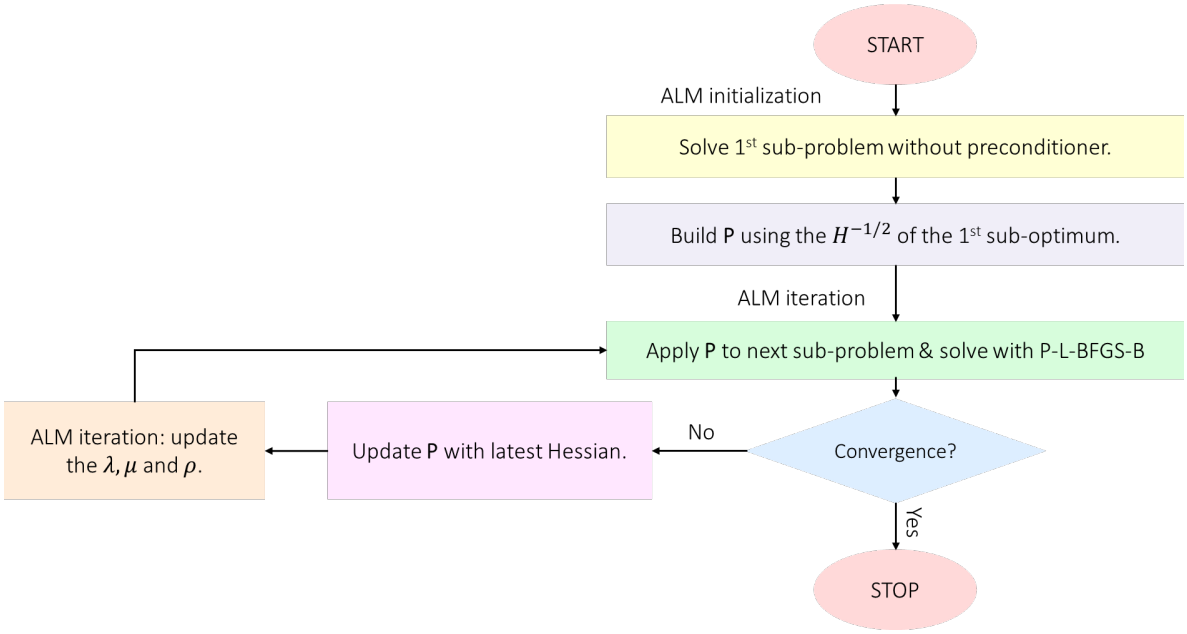


Figure 5: Dynamic Preconditioning with ALM and P-L-BFGS-B Algorithm

2.4 Overview of XRF-1 Wing Sizing optimization Test Case

The "static external preconditioning" and the "dynamic preconditioning" approaches are applied on the structural sizing case of the Airbus XRF-1 research aircraft wing. The optimization problem at hand can be seen in the Eq.(5).

$$\begin{aligned} \min_{t \leq thickness \leq \bar{t}} \quad & Weight(thickness, x_0, x_{90}) \\ \text{s.t.} \quad & g_{buckling}(thickness, x_0, x_{90}) \leq 0 \\ & g_{strains}(thickness, x_0, x_{90}) \leq 0 \end{aligned} \tag{5}$$

A summary of the dimensions of the optimization problem is tabulated in Table 1.

In the context of wing sizing, "super stringers", tabulated in the table below, refer to structural components that add stiffness and strength to the wing structure.

Table 1: Summary of optimization problem dimensions

Parameter	**Quantity**
Design variables	2268
Super Stringers	772
Load cases	5
Failure mode per super stringer	16
Total constraints	57760 ($16 \times 722 \times 5$)
FEM degrees of freedom (DOFs)	72324
Free-body outputs	43320 ($4 \times 3 \times 722 \times 5$)
FEM inputs	3732

A summary of the global optimization framework for this test case is illustrated in Fig.6.

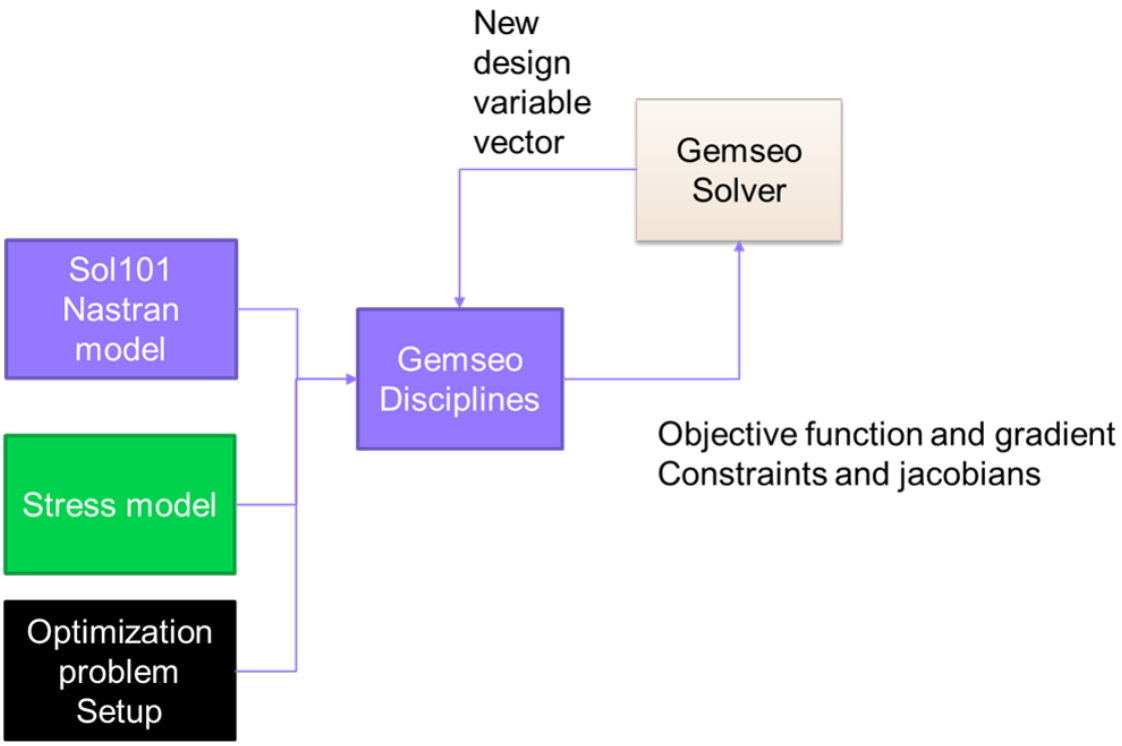


Figure 6: Global optimization Framework for the XRF-1 Wing Sizing Test Case

However, a full-scale optimization cycle as proposed in Fig.1 would lead to several calls to the structural solver, which in this case is Nastran, as shown above Fig.6. To save time while obtaining the structural analysis data, the following 2-step strategy is proposed:-

1. **Sequential Partially Linear Formulations**: The Finite Element Analysis is replaced by a process of iterative linear developments using 1^{st} order Taylor expansions, while all other disciplines are kept unchanged. This allows for a lesser number of Finite Element calls and stepwise linearization formulation of the optimization

problem, thus creating a sequence of simplified linearized sub-optimization problem. The above is achieved using a trust-region based algorithm called the Sequential Partially Linear Programming (SPLP), which is a trust-region-based optimization algorithm to solve non-linear optimization problems with general constraints. The "trust-region" is the region where the linear approximations are still valid. The size of this so-called trust-region is increased or decreased based on each iteration's success, quantified by a merit function. The job of the merit function is to balance the reduction of the objective and satisfying the constraints.

2. **Sequential Quadratic Programming**: Each linearized optimization problem also contains general constraints and thus can be solved using a non-linear programming optimization algorithm like SNOPT or ALM. The preconditioner construction can be turned on and off whenever needed while solving this sub-optimization problem. In the event that ALM is used, the P-L-BFGS-B algorithm is used to solve the sub-sub-optimization problem obtained due to penalization of the linearized objective function.

A summary of the approach is described in Fig.7. Although preconditioning is expected to reduce the design space exploration and optimization time globally, in the present work, the impact of the preconditioning strategies described in sections 2.2 and 2.3 respectively is evaluated, in the zone within the red border, i.e., the gains are seen only for 1 optimization cycle as opposed to solving the full problem.

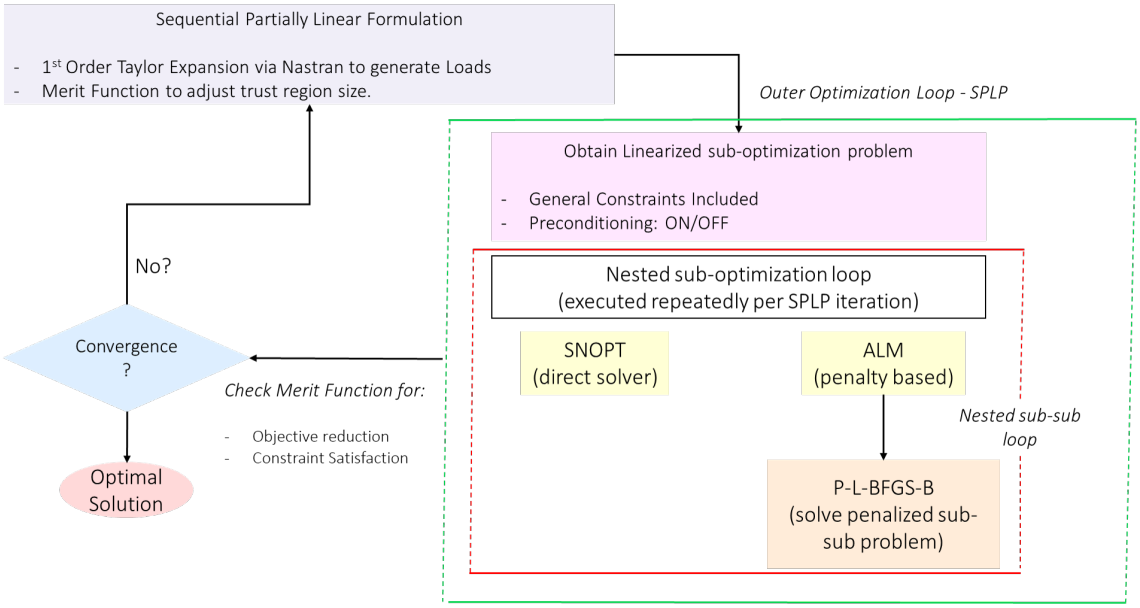


Figure 7: XRF-1 Wing Sizing Optimization: 2 Step Strategy

As the focus of the current work is strictly to observe numerical gains and not physical analyses, the optimization is performed on a "section" of the XRF-1 wing, sectioned from near the pylon of the XRF-1 research aircraft. This section of the wing is henceforth referred to as the "simplified model". The problem dimensions of the simplified model are summarized in Table 2.

Table 2: Summary of Simplified Optimization Problem Dimensions

Parameter	Quantity
Design variables	259
Super Stringers	72
Load cases	5
Failure mode per super stringer	16
Total constraints	5760 ($16 \times 72 \times 5$)
FEM degrees of freedom (DOFs)	72324
Free-body outputs	4320 ($4 \times 3 \times 72 \times 5$)
FEM inputs	3732

It must be understood with clarity that each SPLP iteration corresponds to multiple iterations of SNOPT or ALM, depending on the choice for solving the sub-problem. It is re-emphasized that the preconditioner is applied only on the sub-problems (represented by the red border in Fig.7 and not the problems being solved by SPLP and therefore, the history of the sub-optimization, i.e., number of iterations taken by SNOPT or ALM to reach the optimum is of broader interest. Therefore, computations are performed only for 1 iteration of SPLP wherein, the authors intend to observe any gains in terms of number of sub problem iterations when solved with and without the preconditioner. 1 SPLP iteration corresponds to 1 Nastran run. It is crucial to acknowledge that while the solution at the end of just 1 SPLP iteration may not be feasible in the context of the full optimization problem, it is still accepted as long as the solution at the end of the sub-optimizations is a feasible solution. For clarification, solution at the end of the sub-optimization refer to the optimization result at the end of solving the sub-problems and before forced start and termination of the next SPLP iteration.

3 RESULTS AND DISCUSSION

3.1 Impact of Preconditioning on Benchmark General Constraint Problems

A similar approach as described in the MDA-MDO project [11] is followed, wherein, a few general constrained problems from the Constrained and Unconstrained Testing Environment (CUTE) database [16] are shortlisted. ALM is used to convert the general constraint problems into bound constraint problems with a penalty term, and the sub-problems are solved with the P-L-BFGS-B method twice, once with dynamic preconditioning OFF, tabulated in the ALM column, and once with Dynamic Preconditioning ON, tabulated in the "Prec ALM" column in the Fig.8.

maxcor = 5									
PROBLEM	n	ALM			Prec ALM			nfg gain	f - f_prec
		nact	nfg	time	nact	nfg	time		
S277-280	8	1	23	0.1	1	31	0.1	-35%	6e-05
ARGALALE	200	1	67	0.3	1	71	0.3	-6%	-4e-07
ARGBLLE	200	1	21373	140.7	1	99	1.1	100%	3e-04
ARGCLCLE	200	1	10011	67.9	1	141	1.4	99%	7e-11
HS24	2	1	457	1.4	1	58	0.3	87%	-7e-03
EIGENA2	50	1	89	3.5	1	66	54.1	26%	-6e-05
DEGENQP	50	1	6	64.2	1	6	62.1	0%	-9e-08
FERRISDC	200	1	2	30.3	1	2	31.9	0%	1e-02
MIN								-34.78%	-7.22e-03
MEAN								33.82%	4.12e-04
MAX								99.54%	1.00e-02

Figure 8: Comparison of ALM and Dynamically Preconditioned ALM (Prec ALM) on CUTEst Problems With General Constraints

Different columns contain the following: "n" contains the number of variables, "nact" contains the number of active constraints, "nfg" contains the number of objective evaluations, "time" contains the time of execution for each problem, "nfg gain" contains the gain in terms of objective evaluations, and "f-f_prec" contains the gain in terms of objective value relative to the value obtained without the presence of a preconditioner. A limited memory size "maxcor" = 5 is set for both "ALM" and "Prec ALM" evaluations. It can be observed that gain in terms of number of evaluations is mostly positive, and even reaches as high as ***99.54%***. Globally, a mean gain of ***33.92%*** and $\mathbf{4.12e^{-04}}$ are observed in terms of number of function evaluations and objective values respectively. Prima facie, it can be inferred that a Dynamic Preconditioner yields excellent accelerations for general constraint optimization problems.

However, it is essential to account for the fact that the factor by which the penalty parameter ρ increases in the Eq.(4), shall greatly influence the penalty term during sub-optimum evaluations by the ALM and by consequence, the total number of function evaluations. Henceforth, this increment factor shall be denoted as γ.

To account for the influence of this γ while accessing the gains made due to a Dynamic Preconditioner, a modular version of the Sellar problem [17] is solved, wherein, the number of problem dimensions, denoted by $\mathbf{n_x}$, can be manipulated. The "Best γ" is defined as the value of γ that results in the least number of function evaluations without the presence of a preconditioner, hereby denoted as $\mathbf{nfg_1}$, while number of function evaluations with a Dynamic Preconditioner shall be denoted as $\mathbf{nfg_2}$. Throughout the present work, the "Best γ" is calculated using via trial and error. $\mathbf{f_{opt1}}$ & $\mathbf{f_{opt2}}$ is used to represent the objective value with and without the preconditioner and the notations for the gains for number of function evaluations and objective value remains the same as before. Numerical experiments were performed for $\mathbf{n_x}$ 5 to 100, in steps of 5, and the results are summarized in Fig.9.

n_x	Best γ	ALM nfg_1	ALM f_{opt1}	Prec ALM nfg_2	Prec ALM f_{opt2}	$nfg\ gain$	$f - f_{prec}$
5	2.5	88	3.18	190	3.18	-115.91%	-6.47e-05
10	2.5	91	3.18	217	3.18	-138.46%	2.36e-05
20	2.5	126	3.18	107	3.18	15.08%	-4.28e-06
25	1.5	113	3.18	140	3.18	-23.89%	-8.73e-04
30	2.0	92	3.18	30	3.18	67.93%	1.9e-05
35	2.5	211	3.18	115	3.18	45.50%	1.33e-04
40	1.5	110	3.18	99	3.18	10.00%	-2.69e-07
45	2.5	156	3.18	121	3.18	22.44%	-4.81e-08
50	2.0	219	3.18	171	3.18	21.92%	-1.79e-06
55	3.0	153	3.18	160	3.18	-4.58%	-7.51e-04
60	3.0	132	3.18	99	3.18	25.00%	-3.74e-05
65	1.5	97	3.18	106	3.18	-9.28%	4.9e-02
70	3.5	89	3.18	91	3.18	-2.25%	5.1e-07
75	1.5	124	3.18	99	3.18	20.16	-2.87e-09
80	2.5	105	3.18	104	3.18	0.95%	-7.21e-05
85	1.5	103	3.18	132	3.18	-28.16%	-6.31e-06
90	2.5	139	3.18	64	3.18	53.95%	1.9e-05
95	3.0	111	3.18	91	3.18	18.02%	7.4e-03
100	3.0	123	3.18	33	3.18	73.17%	-5.42e-07
MIN						-138.46%	-7.51e-04
MEAN						2.69%	2.88e-03
MAX						73.17%	4.9e-02

Figure 9: Comparison of ALM and Dynamically Preconditioned ALM (Prec ALM) on Modular Sellar Problems

For the sake of ease of visualization, positive gains are coloured in green and negative gains are coloured in red throughout the course of the present work. Similarly to

the CUTEst problems, a mean positive gain w.r.t. number of function evaluations and objective values are observed, with much more conservative values ***2.69%*** & $\mathbf{2.88e^{-03}}$ respectively. Interestingly, it is observed that the gains overall improve as $\mathbf{n_x}$ increases, as a maximum gain of ***73.17%*** is seen for $\mathbf{n_x}$ in terms of the number of function evaluations.

In the static preconditioning approach, optimization history is utilized to construct a BFGS approximation, and the preconditioner is the square root of the approximate inverse of the Hessian matrix at the optimum. The successive optimizations of the Sequential Partially Linear formulation offer previous optimization histories to do so.

$$P = H^{-1/2} \tag{6}$$

H being positive definite since it is constructed with the BFGS formula, this always exists and can be easily computed with a singular value decomposition [11].

The SNOPT method converges to an optimum very rapidly for both the CUTEst problems and the Modular Sellar problems shown above, with a maximum of 1 to 2 function evaluations, even in the absence of preconditioners; hence, a larger problem is used to validate the static preconditioning approach. In the present context, a "larger problem" simply implies a problem with higher number of dimensions and constraints. A classical non-convex scalable benchmark is used: the Rosenbrock problem with a generalized spherical inequality constraint [18] with problem dimensions starting from 500 to 1000 in steps of 100. The results of the same are summarized in the Fig.10, where $\mathbf{nfg_1}$ and $\mathbf{nfg_2}$ represent the number of iterations without and with the presence of a static preconditioner respectively.

n_x	nfg_1	nfg_2	$nfg\ gain$	$f - f_{prec}$
500	61	58	4.91%	-4.7e-05
600	61	54	11.47%	-3.7e-05
700	45	78	-73.34%	-6.1e-04
800	79	66	16.45%	-1.7e-05
900	48	92	-91.66%	-3.1e-06
1000	57	36	36.84%	-8.3e-03
MIN			-91.66%	-8.3e-03
MEAN			-15.89%	-1.5e-03
MAX			36.84%	-3.1e-06

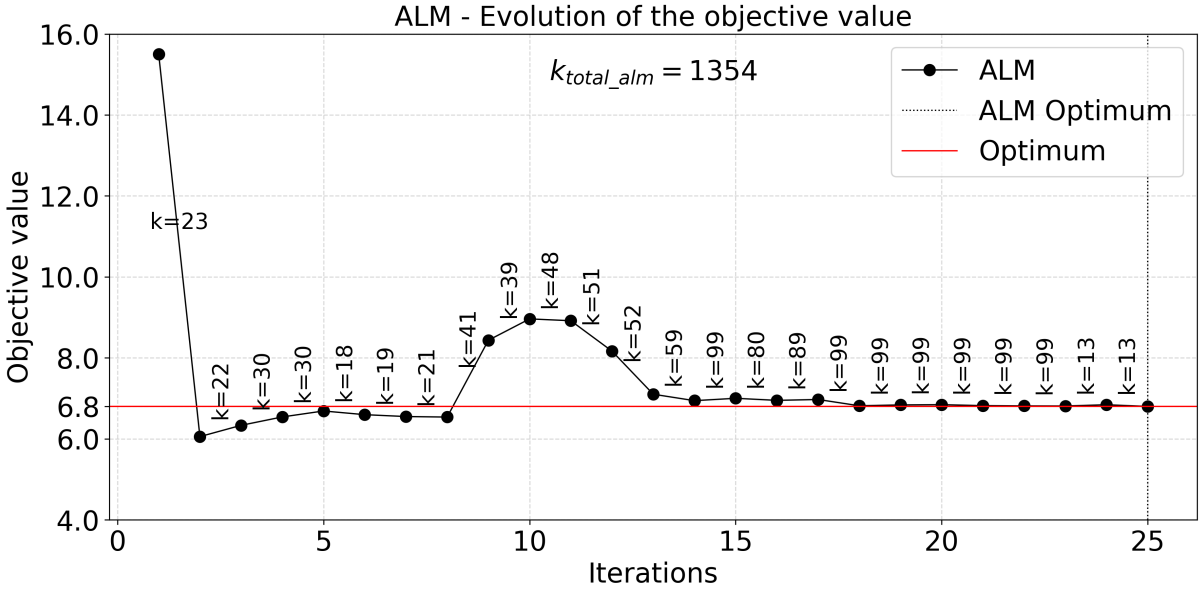
Figure 10: Numerical Experiments using Static preconditioner on the Rosenbrock Problem with Generalized Spherical Inequality Constraint

Similarly to the CUTEst and Sellar problems, positive gains are highlighted in green and negative gains are highlighted in red. It can be observed that the static preconditioner is able to decrease the number of iterations to optimum and thus has positive gains with the exception of $\mathbf{n_x}$=700 & $\mathbf{n_x}$=900, where negative gains are observed, particularly for the case of $\mathbf{n_x}$=900, where the negative gain is as high as ***-91.66%***.

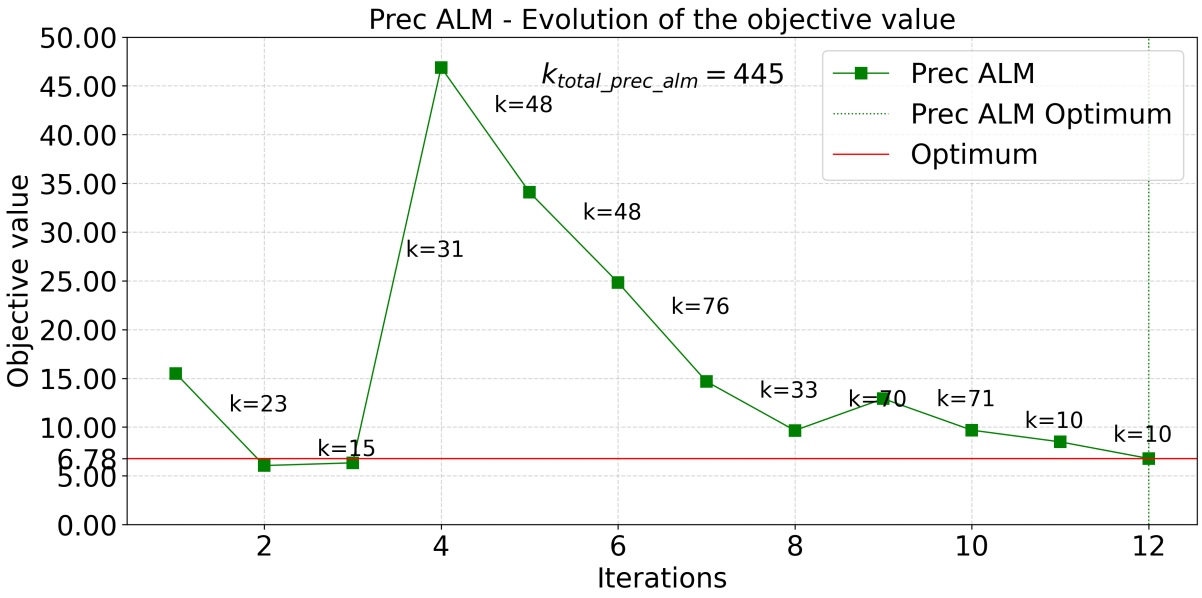
All the methods presented here are integrated in the open source MDO library called GEMSEO (Generic Engine for Multidisciplinary Scenarios, Exploration and Optimization) [2].

3.2 Comparison of preconditioning strategies on the XRF-1 Wing Sizing Test Case

Fig.11 illustrates the impact of the dynamic preconditioning approach on the simplified model. The ALM takes 25 iterations to converge to the optimum without preconditioning and 12 iterations with dynamic preconditioning, as observed in Fig.11a and Fig.11b respectively.



(a) *ALM with Preconditioning OFF*



(b) *ALM with Dynamic Preconditioning: ON*

Figure 11: Dynamic Preconditioning: Numerical Experimentation on Simplified Model

Notably, a steep rise in the objective value is seen in the preconditioned case. This is due to the preconditioner allowing "larger-than-needed" steps and thus, pushing the

solution into regions that are far from the optimum. The P-L-BFGS-B iterations between two successive ALM iterations are annotated as k in the figures. The total number of P-L-BFGS-B iterations is also seen to reduce from 1354 in the reference case to 445 in the dynamically preconditioned case. A net positive gain of *66.81%* is recorded in terms of the total number of calls to the disciplines, i.e., sum of all iterations of ALM, P-L-BFGS-B and SPLP when the dynamic preconditioning is active. In terms of the objective value, a positive gain of $\mathbf{2.5e^{-04}}$ is recorded. The total time of computation for the reference case and the preconditioned case are 29 min 16 sec and 12 min 51 sec respectively.

As a limitation, a lack of robustness was observed in the ALM solver utilized in the present work, whose convergence depends on the initial parameters. This could be attributed to a more basic implementation of the ALM in the present work as opposed to the formulations in the works of X. Jia et al. [13]. More specifically, the computation of the Lagrange multipliers in the current implementation [2] is based on fixed tolerances for active constraint detection as opposed to an active approach based on both the constraint value and its gradient contribution to the KKT system (Karush-Kuhn-Tucker conditions) as articulated in the work of X. Jia et al. [13]. Furthermore, in the current implementation [2], the KKT system is solved without scaling, which could lead to poor conditioning of the KKT system itself. Therefore, although the preconditioning strategy yields favourable gains, further improvements are required in the ALM implementation before applying to actual industrial design problems.

Fig.12 demonstrates the impact of the external static preconditioner on the simplified model, where the reference solution is obtained with the SNOPT algorithm.

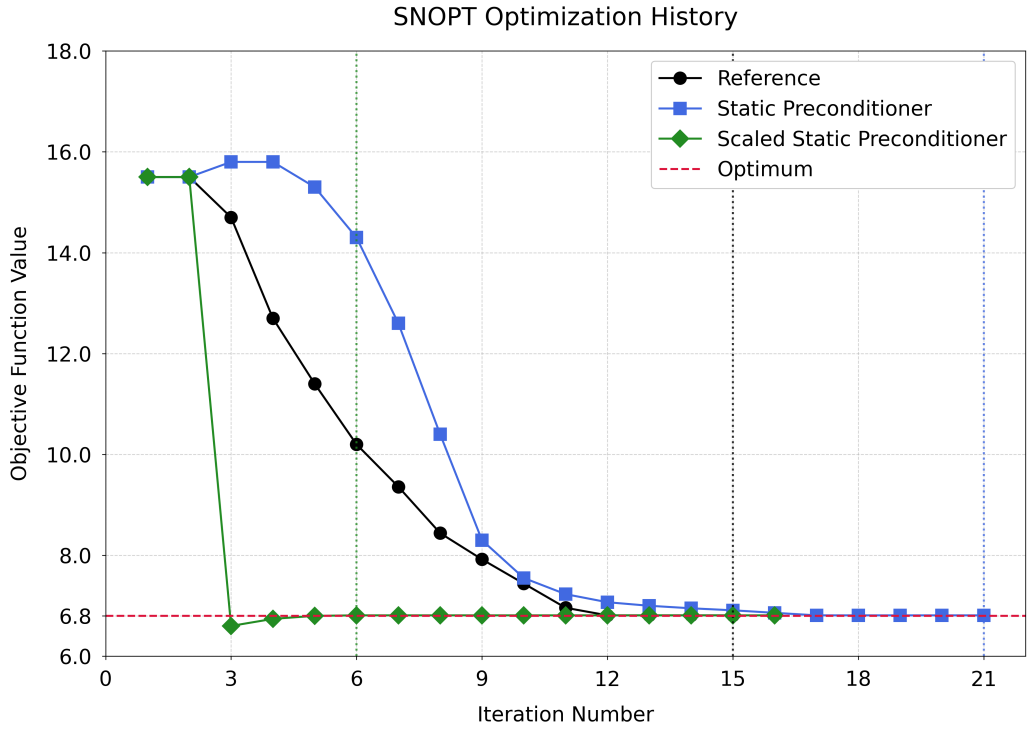


Figure 12: Static External Preconditioning with SNOPT: Numerical Experimentation on Simplified Model

To trigger a "warm-start" [19], the "final preconditioner matrix" obtained from the

dynamic preconditioning with ALM, which is symmetric positive definite, was used as an initial approximate of the square root inverse hessian matrix. The final static preconditioner matrix $\mathbf{P}$ was further built with splitting technique [20] using the Eq.(6) before implementing it in the problem transformation approach with SNOPT as defined in the section 2.2. It is observed that while the reference optimization process, represented in black converges to the optimum in 15 iterations, the preconditioned process, represented in blue, takes 21 iterations. The rise in the number of iteration in the preconditioned case can be attributed to the fact that unlike a dynamic preconditioner, the static preconditioner doesn't account for curvature changes between the each iteration of the SQP. This is verified by the optimization process of the so-called "Scaled Static Preconditioner" represented in green, which converges in 6 iterations. In the present context, scaling simply refers to the transformation of the static preconditioner matrix $\mathbf{P}$ to $\mathbf{K}.\mathbf{P} + \tau.\mathbf{I}$ where, $\mathbf{K}$ & τ are scalar factors and $\mathbf{I}$ is an identity matrix of the same size as $\mathbf{P}$. The idea of this scaling is similar to the works of [8, 9] wherein, they attempt to use scaling in the event their Hessian inverse approximate doesn't yield a favourable eigenvalue distribution for the linear system of interest. The optimum is represented in the figure as a horizontal dotted red line and the dotted vertical lines denote the arrival at the optimum for each optimization process. a computational gain of $\boldsymbol{60\%}$ and a negligible objective difference of $-\mathbf{4e^{-08}}$ are observed.

4 CONCLUSIONS

The present work aimed to explore the potential of different preconditioning strategies and to assess their applicability to industrial MDO problems. These are typically minimization problems and the objective function(s) of the of this minimization problem is subjected to general inequality, equality and bound constraints. In reviewing the state-of-the-art, it was found that preconditioning for constrained optimization remains under explored due to the complexity of Hessian structures and variable constraint activation. To address this, two strategies have been proposed: a static external preconditioner (fixed prior to optimization) and a dynamic preconditioner (updated during optimization of the Augmented Lagrangian solver).

Numerical experiments were performed on analytic benchmark problems for validation, and on a sizing optimization problem on a simplified model of the wing of the Airbus XRF-1 configuration. On the wing test case, the dynamic preconditioning achieved a $\boldsymbol{66.8\%}$ speed-up. The static preconditioning with SNOPT achieved up to $\boldsymbol{60\%}$ speed-up when properly scaled.

Dynamic preconditioning exhibited superior performance but was sensitive to initial ALM settings. One important limitation of the present work is the robustness of the ALM implementation, which required many iterations and influences the results. Static preconditioning, while effective post scaling, needed manual tuning and historical data. Both methods showed trade-offs between convergence speed and implementation complexity.

Upon obtaining favourable gains with preconditioning for the simplified optimization problem, the next planned step the application of the strategies to the full optimization problem whose dimensions are specified in Table 1, for a full-cycle optimization, within a larger bi-level MDO process [14]. Finally, the proposed strategies demonstrate the potential of preconditioning for a wide variety of industrial MDO problems within the aerospace industry and beyond.

ACKNOWLEDGMENTS

This work was supported by the French government under the France 2030 programme under the reference ANR-10-AIRT-01. We also wish to acknowledge the industrial members of the IRT Saint Exupéry project R-Evol: Airbus, Liebherr, Altran Technologies, Capgemini DEMS France, CENAERO, Expleo Group and CERFACS for their support, financial funding, and own knowledge.

REFERENCES

[1] J. Gray, J. Hwang, J. Martins, K. Moore, and B. Naylor. Openmdao: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 04 2019. doi:10.1007/s00158-019-02211-z.

[2] F. Gallard, C. Vanaret, D. Guénot, V. Gachelin, R. Lafage, B. Pauwels, P.-J. Barjhoux, and A. Gazaix. Gems: A python library for automation of multidisciplinary design optimization process generation. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 01 2018. doi:10.2514/6.2018-0657.

[3] L. Etman, A. Groenwold, and J. Rooda. First-order sequential convex programming using approximate diagonal qp subproblems. *Structural and Multidisciplinary Optimization*, 45, 04 2012. doi:10.1007/s00158-011-0739-3.

[4] N. Andrei. *Quasi-Newton Methods*, pages 261–314. 10 2022. ISBN 978-3-031-08719-6. doi:10.1007/978-3-031-08720-2_6.

[5] D. F. Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of Computing*, 24:647–656, 1970. doi:10.1090/S0025-5718-1970-0274029-X.

[6] R. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16, 02 2003. doi:10.1137/0916069.

[7] A. Toper, K. Kabongot, R. Stewart, and A. Daehnke. The mechanism, optimization and effects of preconditioning. *Journal of The South African Institute of Mining and Metallurgy*, 100:7–15, 01 2000.

[8] M. Benzi. Preconditioning techniques for large linear systems: A survey. *Journal of Computational Physics*, 182:418–477, 11 2002. doi:10.1006/jcph.2002.7176.

[9] A. J. Wathen. Preconditioning. *Acta Numerica*, 24:329–376, 2015. doi:10.1017/S0962492915000021.

[10] T. Choi and C. Kelley. Estimates for the nash–sofer preconditioner for the reduced hessian for some elliptic variational inequalities. *Siam Journal on Optimization - SIAM J OPTIMIZATION*, 9, 04 1999. doi:10.1137/S1052623497323364.

[11] S. Gürol. On the use of limited memory preconditioners for geosciences and aerodynamic shape optimization. Presentation Slides, September 2020. URL `http://www.cerfacs.fr`. CERFACS.

[12] P. Gill, W. Murray, M. Saunders, A. Drud, and E. Kalvelagen. Snopt: An sqp algorithm for large-scale constrained optimization. *SIAM Review*, 47, 04 2001. doi:10.1137/S0036144504446096.

[13] X. Jia, C. Kanzow, P. Mehlitz, and G. Wachsmuth. An augmented lagrangian method for optimization problems with structured geometric constraints. *Mathematical Programming*, 199, 09 2022. doi:10.1007/s10107-022-01870-z.

[14] A. Gazaix, V. Ambert, M. Guillen, N. Roussouly, F. Gallard, T. Druot, S. Coniglio, M. Hamadi, V. Drouet, J.-C. Giret, R. Hollinger, J.-F. Figue, A. Dechaume, M. De Lozzo, and G. Ruiz-Jimenez. Industrialization of mdo methods, software and processes, from mbse to high fidelity multi-components wing, pylon and nacelle optimisation. 07 2024. doi:10.2514/6.2024-4404.

[15] A. Gazaix, F. Gallard, V. Ambert, D. Guénot, M. Hamadi, S. Grihon, P. Sarouille, T. Druot, J. Brezillon, V. Gachelin, J. Plakoo, N. Desfachelles, N. Bartoli, T. Lefebvre, S. Gürol, B. Pauwels, C. Vanaret, and R. Lafage. Industrial application of an advanced bi-level mdo formulation to aircraft engine pylon optimization, 06 2019.

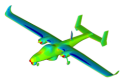
[16] I. Bongartz, A. Conn, N. Gould, and P. Toint. Cute: Constrained and unconstrained testing environment. *ACM Trans. Math. Softw.*, 21:123–160, 03 1995. doi:10.1145/200979.201043.

[17] J. Gray, K. Moore, T. Hearn, and B. Naylor. Standard platform for benchmarking multidisciplinary design analysis and optimization architectures. *AIAA Journal*, 51: 2380–2394, 10 2013. doi:10.2514/1.J052160.

[18] Y. Shang and Y. Qiu. A note on the extended rosenbrock function. *Evolutionary Computation*, 14:119–126, 03 2006. doi:10.1162/106365606776022733.

[19] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004. ISBN 978-0-521-83378-3. URL `https://web.stanford.edu/~boyd/cvxbook/`.

[20] T. F. Chan and T. P. Mathew. Domain decomposition algorithms. *Acta Numerica*, 3:61–143, 1994.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTI-PATCH HYBRID OPTIMIZATION OF COMPOSITE WINGS

D. Zamani[1*], A. Pagani[1], M. Petrolo[1] and E. Carrera[1]

1: MUL2 Lab
Department of Mechanical and Aerospace Engineering
Politecnico di Torino, Italy
Corso Duca degli Abruzzi 24, 10129 Turin, Italy
dario.zamani@polito.it, http://www.mul2.com

Abstract. *The demand for lightweight and structurally efficient composite aerospace structures continues to stimulate the advancement of refined optimization methodologies. This paper proposes a novel multi-patch, multi-objective optimization framework to minimize mass and strain energy in composite wing-box structures. The optimization problem is inherently mixed-integer, as it involves discrete design variables alongside continuous variables, necessitating tailored solution strategies. Two distinct optimization approaches are compared: the traditional multi-objective genetic algorithm NSGA-II and a hybrid method integrating evolutionary and gradient-based optimization. The gradient computations are directly performed at the fundamental nuclei level derived within the Carrera Unified Formulation (CUF), significantly enhancing computational efficiency and accuracy. We provide a scalable low- to high-fidelity methodology for gradient computation, where the CUF framework enables seamless transitions between different levels of accuracy by adjusting the expansion order. Each wing-box segment, modeled as an independent patch, is optimized separately, enabling tailored fiber angle distribution. The effectiveness of the proposed hybrid optimization strategy is demonstrated through numerical case studies, highlighting superior convergence characteristics and improved optimality compared to the NSGA-II standalone approach. The sensitivity of the optimal design solutions to varying structural modeling approaches within the CUF framework is also briefly discussed, providing valuable insights into the practical implementation of multi-patch gradient-enhanced composite structure optimization.*

Keywords: Composite Structures, Multi-objective Optimization, Mixed-Integer Optimization, Carrera Unified Formulation, Multi-patch Optimization

1 INTRODUCTION

The need for enhanced performance and reduced weight in aerospace components has driven substantial research on advanced composite materials and optimization strategies. A key challenge in designing these structures is identifying optimal stacking sequences that minimize mass yet achieve the required mechanical performance. Early studies on laminate optimization addressed straight-fiber configurations, encoding discrete ply angles into genetic algorithms (GA) and demonstrating the feasibility of integer-coded design variables [1–3]. Although such approaches allow direct inclusion of manufacturing constraints [4], the resulting search space is frequently large and nonconvex, making convergence toward globally optimal solutions both computationally intensive and complex. To mitigate these difficulties, an alternative paradigm based on lamination parameters has gained traction [5]. By representing the laminate's stiffness in a continuous variable space, lamination parameters reduce the problem's dimensionality and facilitate more efficient gradient-based or hybrid searches. This concept has been widely employed in buckling and layup optimization studies, including for long cylindrical shells under combined loads [6]. Additional research has explored how approximate feasible regions for lamination parameters can enhance the efficiency of structural optimization [7] and how strength-based failure criteria can be embedded within the lamination parameter design space [8]. Such advancements enable the formulation of composite structure optimization with greater fidelity while keeping the computational burden manageable. For instance, a two-step "slice and swap" procedure for large-scale wing structures demonstrates how continuous optimization over lamination parameters can be followed by a discrete refinement stage that adjusts ply angles per practical design rules [9]. Despite these methodological advancements, purely evolutionary or gradient-based algorithms exhibit certain limitations. GAs excel at global exploration but often require high computational costs, mainly when high-fidelity structural analyses are performed at every iteration. Conversely, given reliable sensitivities, gradient-based methods converge quickly but can struggle with non-convexities or large discrete design spaces. Consequently, a growing body of work focuses on hybrid frameworks integrating global GA searches with local gradient-based refinements. For example, Yang and Fan [10] proposed a multi-objective GA with deterministic gradient steps to accelerate convergence in design problems. Recent work has introduced integer-based programming and advanced learning-based methods to broaden the scope of hybrid optimization in composite design. For example, Borwankar et al. [11] present a two-level weight optimization approach for laminates that leverages integer programming to streamline the design space. Meanwhile, Taymaz [12] investigates how optimizing the structural distribution of composite layers can reduce helicopter rotor blade vibrations, emphasizing the potential of hybrid frameworks in dynamic applications. Finally, Singh and Kapania [13] propose an active learning-based multi-objective optimization (ALMO) method to accelerate constrained evolutionary algorithms, illustrating how machine-learning-enhanced techniques can reduce computational costs while preserving solution quality.

Carrera Unified Formulation (CUF) has emerged as a powerful tool for analyzing composite structures with variable fidelity [14–16]. By selecting the expansion order through the thickness, CUF allows scaling from simplified, low-fidelity to high-fidelity simulations under a unified framework. Prior applications have demonstrated the ability of CUF to capture buckling, flutter, and other complex behaviors in composite wings [17–19], while

proving effective in optimization frameworks [20–23]. Crucially, performing gradient evaluations directly at the level of the fundamental nuclei yields accurate sensitivities without resorting to finite-difference approximations, thus making high-fidelity hybrid optimizations more practical.

This paper builds on recent developments by proposing a multi-patch, multi-objective optimization approach for composite wing-box structures. Each patch is modeled and optimized separately, utilizing discrete and continuous design variables to minimize mass and strain energy while adhering to multiple constraints. We compare two methodologies: the standard NSGA-II and a hybrid variant that combines heuristic optimization with CUF-based gradient refinements. The numerical results demonstrate that combining GA-driven exploration with gradient-based tuning can significantly improve the efficiency of high-fidelity composite optimization.

2 STRUCTURAL MODELING

This study analyzes composite laminates using refined 2D models based on CUF [24]. Within the CUF framework, the three-dimensional displacement field is described by an arbitrary expansion in the thickness direction, $F_\tau(z)$, representing the generalized displacement components in the x-y plane. This formulation allows for:

$$\mathbf{u}(x, y, z) = F_\tau(z)\mathbf{u}_\tau(x, y) \qquad \tau = 1, ..., M \tag{1}$$

Following Einstein's notation, the subscript τ indicates summation over the M expansion terms, while $\mathbf{u}_\tau(x, y)$ denotes the generalized displacements in the x-y plane. This study uses a Taylor expansion (TE) to develop an Equivalent Single Layer (ESL) model. Conversely, a Lagrange expansion (LE) is applied to individual layers to construct Layer-Wise (LW) models. A Lagrange displacement (LD) based LW approach is adopted to ensure the continuity of displacements between layers' interfaces.
In this context, the notations TEn and LDn correspond to the use of Taylor and Lagrange polynomials of the n-th order, respectively. Furthermore, XLDn specifies LW models employing X Lagrange polynomials of n-th order per layer. By integrating the Finite Element (FE) method with CUF, the following expression for the three-dimensional displacement field is obtained:

$$\mathbf{u}(x, y, z) = N_i(x, y)F_\tau(z)\mathbf{q}_{\tau i} \quad i = 1, ..., N_n \quad \tau = 1, ..., M \tag{2}$$

in which $\mathbf{q}_{\tau i}$ denotes the unknowns vector, and N_n is the number of nodes in each element. The present study consideres bi-quadratic FEs, referred to as Q9, as the $N_i(x, y)$ shape functions.

The governing equations are derived through the principle of virtual displacements (PVD). First, the principle underlying PVD is applied to describe the problem in terms of virtual work, i.e.:

$$\delta\mathcal{L}_{int} = \delta\mathcal{L}_{ext} \tag{3}$$

The virtual variation of the internal strain energy is defined as:

$$\delta\mathcal{L}_{int} = \int_V \delta\varepsilon^T \sigma dV \tag{4}$$

where V is the volume of the body. The virtual work due to external loadings includes the contribution of surface loads, $\mathbf{P}_S$, line loads, $\mathbf{P}_l$, and point loads, $\mathbf{P}$; and can be written in the following forms:

$$\delta\mathcal{L}_{ext} = \int_S \delta\mathbf{u}^T\mathbf{P}_S dS + \int_l \delta\mathbf{u}^T\mathbf{P}_l dl + \delta\mathbf{u}^T\mathbf{P} \tag{5}$$

Equation (4) can be reformulated by employing Eq. (2), the constitutive equation $\sigma = \mathbf{C}\varepsilon$, and the geometrical relations between strains and displacements, resulting in:

$$\delta\mathcal{L}_{int} = \delta\mathbf{q}_{sj}^T \left[\int_V \mathbf{D}^T(N_jF_s)\widetilde{\mathbf{C}}\mathbf{D}(N_iF_\tau)\,dV\right]\mathbf{q}_{\tau i} = \delta\mathbf{q}_{sj}^T\mathbf{k}^{ij\tau s}\mathbf{q}_{\tau i} \tag{6}$$

in which $\mathbf{k}^{ij\tau s}$ is the 3×3 Fundamental Nucleus (FN) of the stiffness matrix, which is invariant regardless of the order of the $N_i(x,y)$ shape functions and the $F_\tau(z)$ expansion. The differential operator matrix containing the geometrical relations is expressed as $\mathbf{D}(\cdot)$. For orthotropic materials with a generic fiber orientation in space, the global components of the elasticity tensor $\widetilde{\mathbf{C}}$ are computed using the rotation matrix $\mathbf{T}$ without altering the definition of the FNs. As described in [25], $\widetilde{\mathbf{C}}$ is the material stiffness matrix in the global reference frame, calculated as $\widetilde{\mathbf{C}} = \mathbf{T}\mathbf{C}\mathbf{T}^T$. Finally, the global assembled stiffness matrix $\mathbf{K}$ is constructed by iterating over the indices i, j, τ, and s. The assembly of the stiffness matrix, $\mathbf{K}$, varies based on the chosen modeling approach. As mentioned, two commonly utilized strategies for analyzing multilayered structures are ESL and LW methods. In ESL modeling, the properties of individual layers are homogenized and then combined to compute the stiffness matrix. However, as explained in [26], the ESL approach does not satisfy the C_z^0 continuity requirements. On the other hand, the LW method models each layer independently, expanding the displacement field within each lamina. To ensure C_z^0 continuity, displacement continuity must be enforced at the interfaces between layers.

Once the global stiffness matrix $\mathbf{K}$ has been assembled according to the CUF framework, the strain energy stored in the structure can be expressed as:

$$E \;=\; \frac{1}{2}\int_V \boldsymbol{\varepsilon}^T\boldsymbol{\sigma}\,dV = \frac{1}{2}\,\mathbf{u}^T\,\mathbf{K}\,\mathbf{u} \tag{7}$$

where $\boldsymbol{\sigma}$ and $\boldsymbol{\varepsilon}$ denote, respectively, the stress and strain fields within the volume V. Within the FE framework, the strain energy formulation directly follows the quadratic expression in Eq. (7).

3 OPTIMIZATION FRAMEWORK

The optimization problem addressed in this study aims to minimize the structural mass and the global strain energy of a composite wing, by simultaneously tuning the stacking sequence in multiple spanwise patches. The design vector is defined as $\boldsymbol{\alpha} = \{n_1, n_2, \ldots, n_P, \theta_1, \ldots, \theta_{N_{\max}}\}$, where P is the number of patches. The integer variables $n_k \in [4, 12]$ represent the number of plies in the k-th patch, while the continuous variables $\theta_i \in [-90°, 90°]$ define the fiber orientation angles. The resulting multi-objective optimization problem can be formally stated as:

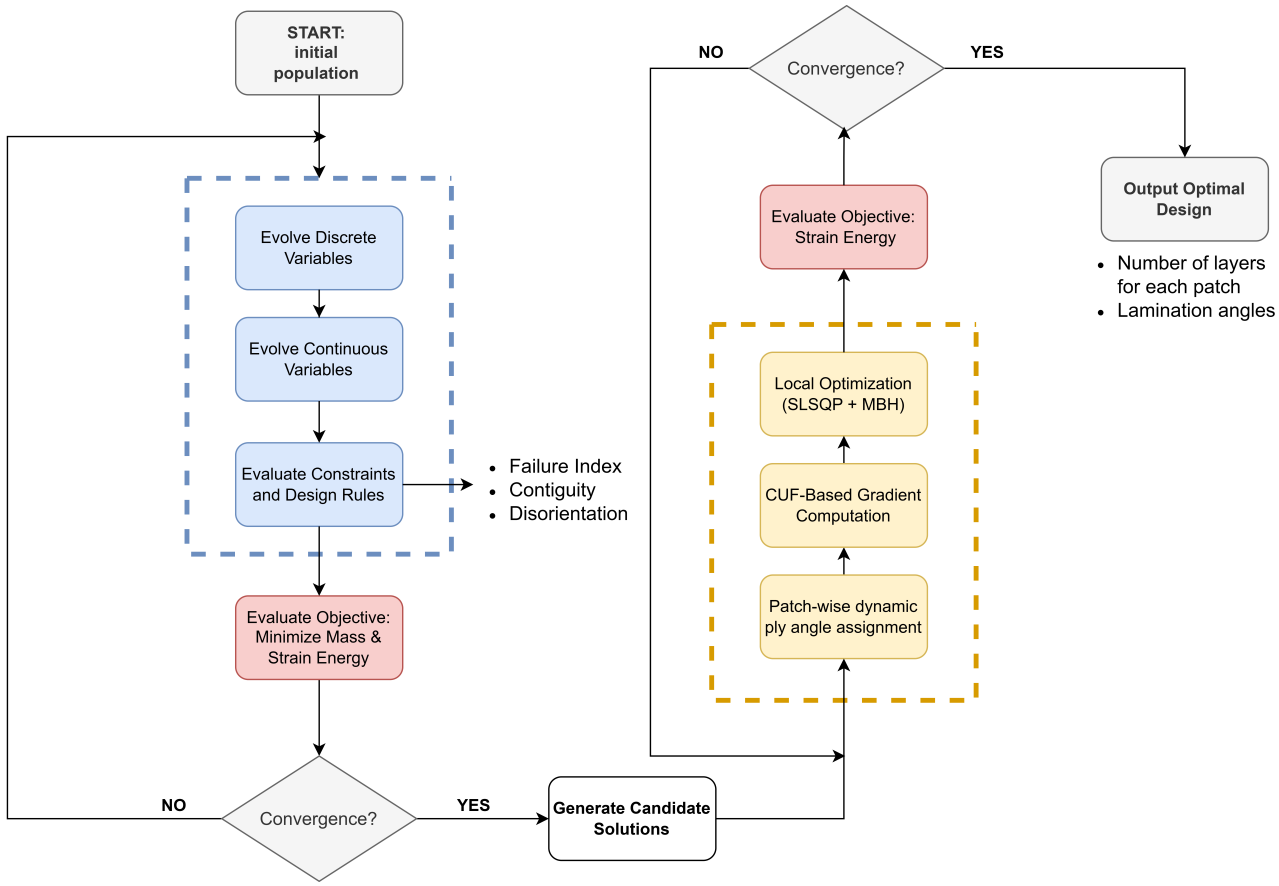


Figure 1: Overview of the proposed hybrid optimization strategy for composite laminates. The left branch represents the global exploration phase based on SADE. Feasible designs are passed to the local refinement phase (right) using SLSQP with MBH. The gradient of the strain energy is computed analytically via CUF, as defined in Eq. (12), enabling patch-wise sensitivity-based angle updates.

$$\begin{aligned} &\text{Minimize} && \Big[f_1(\boldsymbol{\alpha}),\ f_2(\boldsymbol{\alpha})\Big], \\ &\text{w.r.t.} && \boldsymbol{\alpha} \ = \ \big\{\, n_1,\, n_2,\, \ldots,\, n_P,\, \theta_1,\, \ldots,\, \theta_{N_{\max}} \big\}, \\ &\text{subject to} && \mathbf{K}(\boldsymbol{\alpha})\,\mathbf{u}(\boldsymbol{\alpha}) = \mathbf{F}, \\ & && C(\boldsymbol{\alpha}) \ \leq \ 0\,. \end{aligned} \tag{8}$$

Here, $f_1(\boldsymbol{\alpha})$ denotes the total mass of the structure, while $f_2(\boldsymbol{\alpha})$ corresponds to the global strain energy computed under prescribed loading and constraints. The displacement field $\mathbf{u}$ satisfies the linear elastic equilibrium $\mathbf{Ku} = \mathbf{F}$, with $\mathbf{K}$ being the stiffness matrix assembled from the CUF formulation. The constraint function $C(\boldsymbol{\alpha})$ aggregates all feasibility conditions, including structural integrity (via failure indices) and manufacturability requirements such as contiguity and disorientation. The number of fiber angles $N_{\max}$ is determined dynamically as the maximum number of plies across all patches, ensuring that each layer can be assigned an independent orientation.

A key set of design rules is associated with the stacking sequence design, specifically contiguity and disorientation. The contiguity requirement limits the number of consecutive plies with the same fiber orientation to a maximum of three, while the disorientation

condition enforces that the angular difference between two adjacent plies does not exceed 45°. These constraints are crucial to ensure smooth variation of mechanical properties through the thickness and to mitigate the risk of delamination and localized stress concentrations.

The structural response is further evaluated by computing failure indices (FI) derived from the Hashin 3D criteria [27]. Four separate modes are considered: fiber tension (FT), fiber compression (FC), matrix tension (MT), and matrix compression (MC), each expressed in terms of the local stress components σ_{ij} and the material strengths $X_T, X_C, Y_T, Y_C, S_{ij}$. The respective values of FT, FC, MT, and MC are computed at through-the-thickness points in all critical regions of the structure, including the wing root and the interfaces between adjacent patches. These indices are then aggregated using a Kreisselmeier–Steinhauser (KS) function:

$$\mathrm{KS}_\rho(\mathbf{f}) \;=\; \max(\mathbf{f}) \;+\; \frac{1}{\rho}\,\ln\Big(\textstyle\sum_i \exp\{\rho\,\big(f_i - \max(\mathbf{f})\big)\}\Big), \quad \rho = 100, \tag{9}$$

which produces a single measure of how close each failure mode is to its allowable limit. The global failure index FI is then obtained by combining these modes:

$$\mathrm{FI} \;=\; \min\Big(1, \frac{1}{\mathrm{KS}_\rho(\mathrm{FI}_{FT})}\Big)\,\min\Big(1, \frac{1}{\mathrm{KS}_\rho(\mathrm{FI}_{FC})}\Big)\,\min\Big(1, \frac{1}{\mathrm{KS}_\rho(\mathrm{FI}_{MT})}\Big)\,\min\Big(1, \frac{1}{\mathrm{KS}_\rho(\mathrm{FI}_{MC})}\Big). \tag{10}$$

Finally, to unify all manufacturing and structural requirements, the aggregated feasibility measure is defined as:

$$C \;=\; \mathrm{FI} \;\cdot\; \mathrm{Contiguity} \;\cdot\; \mathrm{Disorientation}, \tag{11}$$

where Contiguity and Disorientation both equal 1 if their respective rules are satisfied (and 0 otherwise), while FI assesses the overall safety margin with respect to multiple failure modes across critical thickness locations.

Two different optimization strategies are compared. The first employs *NSGA-II* to solve the mixed-integer bi-objective problem directly, producing a Pareto front of $[f_1, f_2]$. An *early stopping* mechanism terminates NSGA-II if the hypervolume (HV) improvement remains below 10^{-3} for three consecutive generations. The second approach is a *hybrid* method that combines Self-Adaptive Differential Evolution (SADE) with a gradient-based local refinement (SLSQP). In this hybrid method, a weighted sum of f_1 and f_2 provides a single scalar objective for SADE, followed by a multistart basin-hopping routine with SLSQP. Early stopping in the SADE phase occurs when the improvement in best objective is less than 10^{-3} for five consecutive generations, while SLSQP is limited to a maximum of 100 function evaluations in each local run.

Figure 1 illustrates the hybrid optimization framework employed in this work. The process begins with a global evolutionary phase (left side), where a multi-objective genetic algorithm explores the discrete (number of plies per patch) and continuous (ply angles) design space. The candidate designs are evaluated in terms of mass and strain energy, while enforcing manufacturability constraints such as contiguity and disorientation. Upon convergence, the best design is passed to a local refinement phase (right side), which minimizes the strain energy via a gradient-based algorithm (SLSQP) wrapped in a Monotonic Basin Hopping (MBH) strategy.

Table 1: Key parameters of the optimization algorithms.

Parameter	**NSGA-II**	**Hybrid**
Population size	52	52
Max generations	60	60
Early stopping criterion	$\Delta\text{HV} < 10^{-3}$	$\Delta f < 10^{-3}$
Stall generations	5	5
Gradient usage	None	SLSQP refinement
Scalarization weights	—	$w_1 = 0.5$, $w_2 = 0.5$
Max SLSQP evaluations	—	100

The gradient of the strain energy with respect to the lamination angles is computed directly within the CUF. Specifically, the derivative is obtained at the level of the fundamental nuclei of the global stiffness matrix and reads:

$$\frac{dE}{d\theta} = \frac{1}{2}\mathbf{u}^T \frac{d\mathbf{K}}{d\theta}\mathbf{u} = \frac{1}{2}\mathbf{u}^T \mathbf{D}_K \mathbf{u}, \tag{12}$$

where $\mathbf{u}$ is the global displacement vector and $\mathbf{D}_K$ denotes the derivative of the global stiffness matrix with respect to the fiber angle θ. This formulation is compatible with both ESL and LW CUF-based models, and enables efficient computation of sensitivities for patch-wise angle optimization.

By contrasting NSGA-II and this two-phase hybrid method, it becomes possible to assess the effectiveness of local refinement in mixed-integer, multi-objective composite laminate design. The next section presents numerical results and discusses the relative performance of these approaches.

4 NUMERICAL RESUTS

The composite structure under investigation (see Fig. 2) consists of a flat wing-like laminate with a semi-span length of $L = 0.305\,\text{m}$ and a semi-chord of $c = 0.0762\,\text{m}$. Each lamina has a uniform thickness of $t_i = 0.134\,\text{mm}$, and the total thickness depends on the number of plies defined per patch. The laminate is divided into an arbitrary number of spanwise patches ($P = 1, 2, 3$), which are uniformly distributed along the span. In each patch, both the number of plies and the ply angles are independently optimized.

The structure is fully clamped at the left edge, mimicking a root attachment to a fuselage or support frame, while a uniformly distributed pressure load of $10^3\,\text{Pa}$ is applied to the upper surface. FE discretization is performed using bi-quadratic Q9 elements over a 5×20 mesh, which was determined to yield converged results through a preliminary convergence study. Both ESL-TE1 and LW-LD3 models are employed in this Section for consistent comparison across all test cases. The material used is a unidirectional carbon/epoxy composite [28], whose elastic and strength properties are summarized in Table 2.

4.1 Single-Patch Optimization Results

In this section, we present the results of the single-patch optimization, where the number of plies and fiber angles are simultaneously optimized over the entire wing domain.

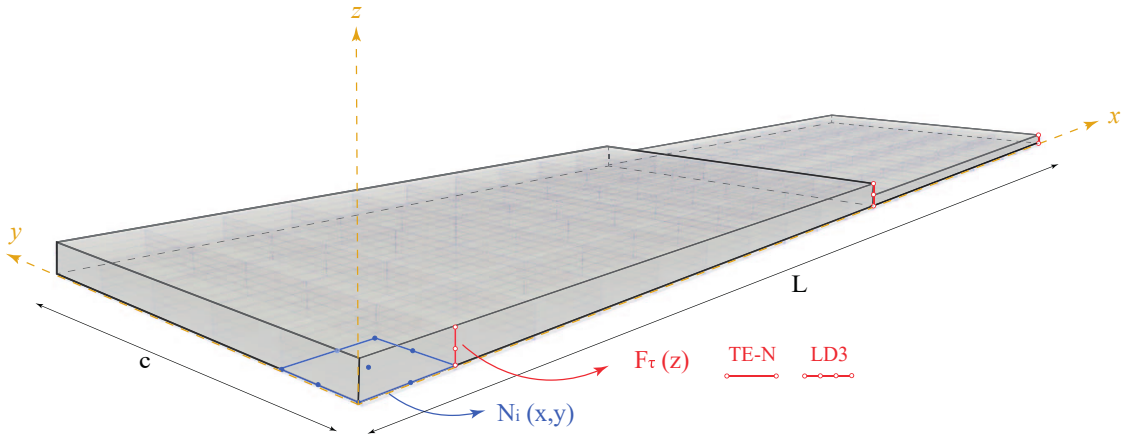


Figure 2: Schematic representation of the composite laminate model. The structure is clamped at the root and subjected to a uniform pressure load on the upper surface. Spanwise patches are introduced along the length for localized design optimization.

Table 2: Material properties and strength characteristics of the composite laminate.

Elastic properties		**Strength properties**	
E_1 [GPa]	98.0	X_T [MPa]	1690.0
E_2, E_3 [GPa]	7.90	X_C [MPa]	800.0
G_{12}, G_{13} [GPa]	5.60	Y_T [MPa]	28.0
G_{23} [GPa]	5.60	Y_C [MPa]	100.0
ν_{12}, ν_{13}	0.28	S_{12} [MPa]	75.0
ν_{23}	0.28	$S_{23} = S_{13}$ [MPa]	75.0

Two structural models are employed: an ESL based on a first-order Taylor expansion (TE1) and an LW model using third-order Lagrange polynomials (LD3) per layer. Each model is tested using two optimization strategies: a standard multi-objective genetic algorithm (NSGA-II) and a hybrid two-step approach that combines global evolutionary search with local gradient-based refinement.

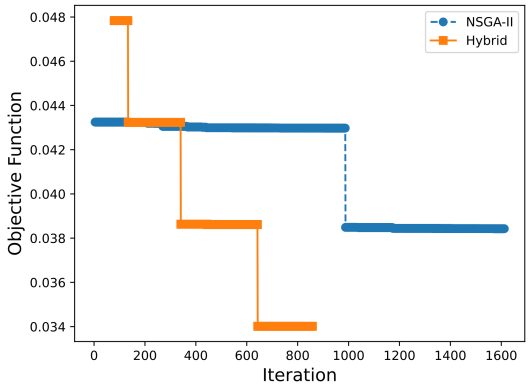
Table 3 summarizes the optimal single-patch designs for the ESL-TE1 model. NSGA-II converges to an 8-ply laminate with a minimum strain energy of 0.564 J at a mass of 37.9 g, whereas the hybrid solution adopts 7 plies, reducing the mass to 33.1 g but increasing the strain energy to 0.881 J. From a computational perspective, the hybrid approach proves significantly faster, converging in approximately 2.4 times less time than NSGA-II.

The convergence behavior of the two strategies is illustrated in Fig. 3a. The hybrid method demonstrates a steep initial descent in the objective function, reaching viable solutions early in the process, while NSGA-II requires more iterations to achieve comparable performance. However, as shown in the final Pareto front in Fig. 3b, NSGA-II exhibits superior spread and dominance in the low-strain energy region, indicating a more effective exploration of the design space in terms of structural performance.

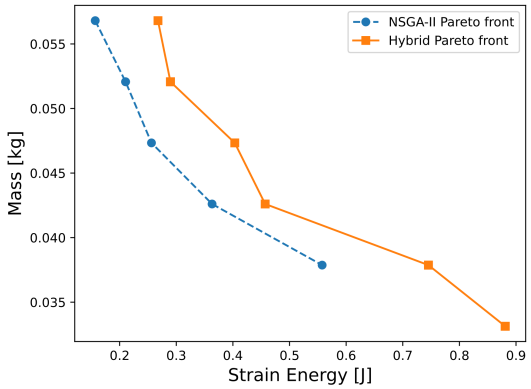
Table 4 reports the results for the single patch for the LW-LD3 model. NSGA-II identifies a 9-ply solution (42.6 g) with the lowest strain energy (0.366 J), while the hybrid design uses 8 plies (37.9 g) and a higher strain energy (0.638 J). As in the ESL-TE1 case,

Table 3: Single-patch optimum results for ESL-TE1.

	NSGA-II	**Hybrid**
No. of plies	8	7
Mass [kg]	0.0379	0.0331
Strain [J]	0.5639	0.8809
Angles [°]	2.49, -7.06, 5.40, -7.67, 30.80, -57.51, -8.88, -9.06	-8.87, 17.99, -67.09, -18.03, -8.32, -7.38, -8.51



(a) Objective function convergence for the ESL-TE1 model.



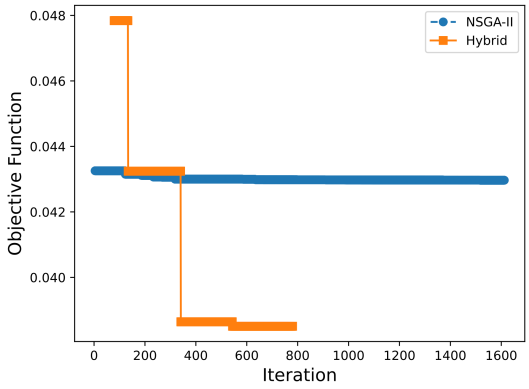
(b) Pareto front comparison for the ESL-TE1 model.

Figure 3: (a) NSGA-II vs. Hybrid objective convergence and (b) final Pareto front for the single-patch ESL-TE1 configuration. The hybrid method converges about 2.4 times faster, although NSGA-II explores more low-strain designs.

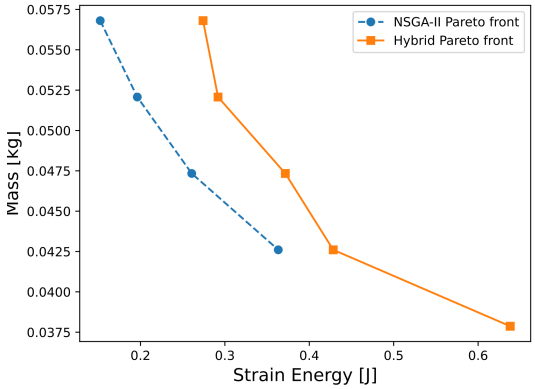
the hybrid algorithm exhibits a faster convergence trend, which can be seen in Fig. 4a showing a rapid stabilization of the objective function. However, Fig. 4b reveals that NSGA-II maintains a better spread of solutions, particularly in the region of low strain energy, confirming its superior exploratory capabilities.

Table 4: Single-patch optimum results for LW-LD3.

	NSGA-II	**Hybrid**
No. of plies	9	8
Mass [kg]	0.0426	0.0379
Strain [J]	0.3663	0.6380
Angles [°]	-0.92, -2.03, 8.44, -14.53, 57.22, -50.30, -10.38, -9.76, 5.79	2.73, -18.89, -55.05, 3.19, 18.51, 28.97, -39.61, -2.34



(a) Objective function convergence for the LW-LD3 model.



(b) Pareto front comparison for the LW-LD3 model.

Figure 4: (a) NSGA-II vs. Hybrid objective convergence and (b) final Pareto front for the single-patch LW-LD3 configuration.

4.2 Multiple-Patch Optimization Results

This section investigates how partitioning the laminate into multiple spanwise patches affects the trade-off between mass and strain energy. As in the single-patch case, two structural models (ESL-TE1 and LW-LD3) and two optimization strategies (NSGA-II and Hybrid) are considered. The best designs are evaluated for both two- and three-patch configurations, highlighting the relative performance of each approach.

4.2.1 Two-Patch Configurations

Table 5 reports the optimum solutions for the two-patch ESL-TE1 configuration. NSGA-II converges to a heavier but stiffer design, using 12 and 6 plies in Patch I and II, respectively, reaching a minimum strain energy of 0.192 J at a mass of 42.60 g. In contrast, the Hybrid solution reduces the ply count to (8, 4), resulting in a much lighter structure (28.40 g), but with a higher strain energy of 0.989 J.

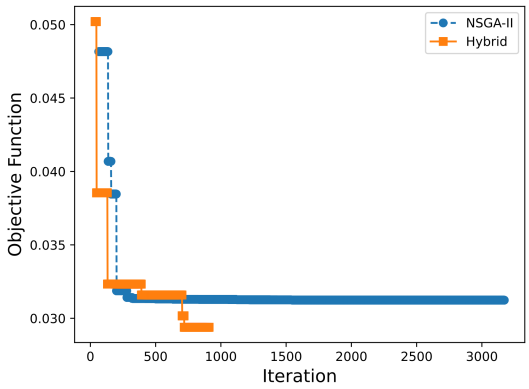
The convergence behavior in Fig. 5a highlights the advantage of the Hybrid approach in terms of computational speed, reaching near-optimal solutions in significantly fewer iterations. However, the Pareto front in Figure 5b shows that NSGA-II produces a broader front, with superior performance in the low-strain region, confirming its effectiveness in exploring a more diverse solution space.

The LW-LD3 model shows similar trends in performance, as summarized in Table 6. NSGA-II achieves the lowest strain energy (0.149 J) using 12 plies in both patches, while the Hybrid design reduces the total ply count substantially (6, 4), achieving a mass of just 23.67 g but a much higher strain energy (1.497 J).

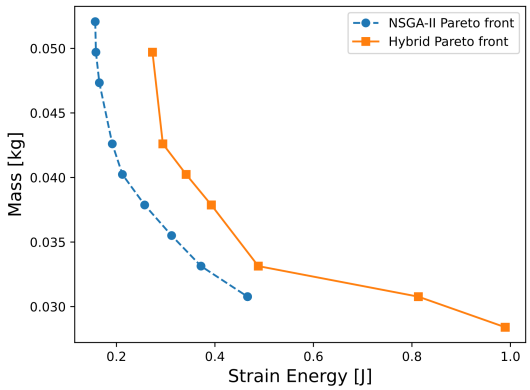
In terms of convergence, Fig. 6a confirms that the Hybrid method converges faster, though Fig. 6b demonstrates that NSGA-II again dominates in the low-strain energy region of the Pareto front. These results reinforce the complementary nature of the two approaches: Hybrid excels at minimizing weight efficiently, while NSGA-II uncovers a richer set of trade-offs.

Table 5: Two-patch optimum results for ESL-TE1.

	NSGA-II	**Hybrid**
No. of plies (Patch I)	12	8
No. of plies (Patch II)	6	4
Mass [kg]	0.0426	0.0284
Strain [J]	0.1923	0.9894
Angles [°]	-0.40, -3.75, 12.82, -1.60, 4.08, -19.92, 9.31, -5.31, 9.67, 0.07, -20.47, 2.61	-9.79, -15.70, 75.18, -43.13, -2.30, -6.14, 14.52, 4.72



(a) Objective function convergence for the two-patch ESL-TE1 case.



(b) Pareto front comparison for the two-patch ESL-TE1 case.

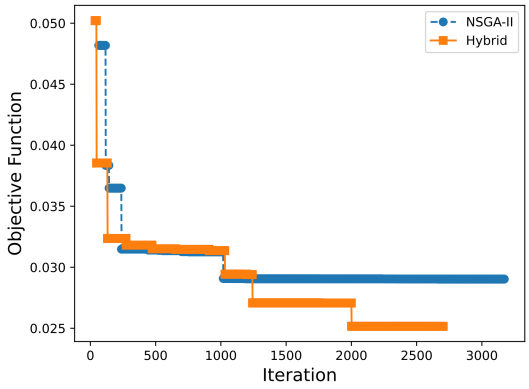
Figure 5: (a) NSGA-II vs. Hybrid convergence and (b) final Pareto front for the two-patch ESL-TE1 configuration.

Table 6: Two-patch optimum results for LW-LD3.

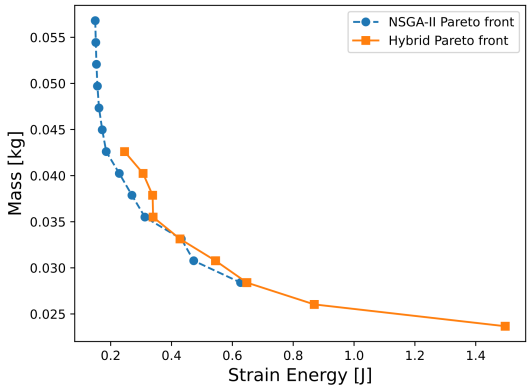
	NSGA-II	**Hybrid**
No. of plies (Patch I)	12	6
No. of plies (Patch II)	12	4
Mass [kg]	0.0568	0.0237
Strain [J]	0.1494	1.4968
Angles [°]	-0.54, -2.61, 5.18, -1.08, -31.19, 10.74, -20.77, 0.88, 4.96, -0.34, 0.70, 0.31	-6.76, 14.63, 36.17, 31.67, -8.60, 3.08

4.2.2 Three-Patch Configurations

The performance comparison for the three-patch ESL-TE1 configuration is provided in Table 7. NSGA-II selects a 12–5–4 laminate layout, achieving a strain energy of 0.433 J and a total mass of 33.14 g. The Hybrid algorithm opts for a lighter 7–4–4 solution



(a) Objective function convergence for the two-patch LW-LD3 case.



(b) Pareto front comparison for the two-patch LW-LD3 case.

Figure 6: (a) NSGA-II vs. Hybrid convergence and (b) final Pareto front for the two-patch LW-LD3 configuration.

(23.67 g) at the expense of a significantly increased strain energy (1.496 J). Interestingly, while the Hybrid algorithm typically converges faster, in this case its runtime marginally exceeds that of NSGA-II, likely due to the expanded dimensionality of the search space, as can be seen in Fig. 7a. The Pareto front of Fig. 7b, where NSGA-II maintains a clear advantage in identifying low-strain designs, while the Hybrid front is clustered in the lightweight region.

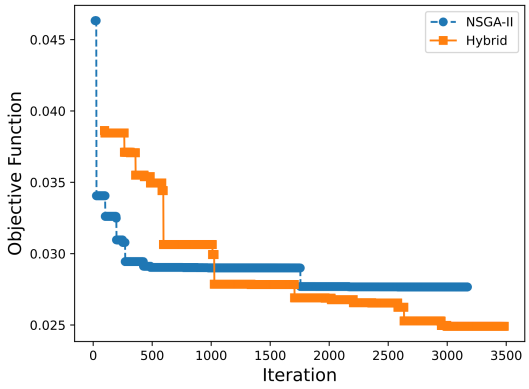
Table 7: Three-patch optimum results for ESL-TE1.

	NSGA-II	**Hybrid**
No. of plies (Patch I)	12	7
No. of plies (Patch II)	5	4
No. of plies (Patch III)	4	4
Mass [kg]	0.0331	0.0237
Strain [J]	0.4334	1.4955
Angles [°]	-5.08, 2.15, 5.85, -4.92, 13.84, -2.08, 3.47, 1.27, -3.05, -2.26, 3.84, 1.60	-19.64, 11.70, 51.10, 6.80, -21.86, -8.22, -1.98

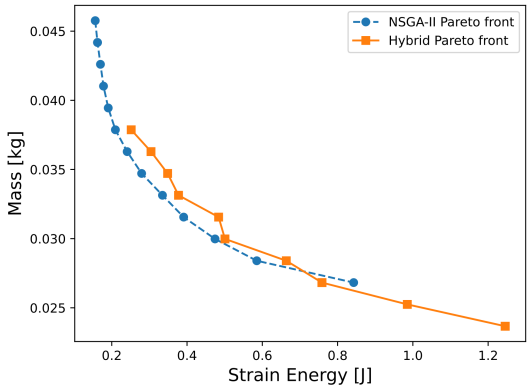
The LW-LD3 model follows similar trends. As shown in Table 8, the best NSGA-II design employs 11–7–4 plies and achieves the lowest strain energy (0.279 J) at a moderate mass of 34.71 g. The Hybrid counterpart again prioritizes weight, selecting a 7–4–4 layout with a significantly reduced mass (23.67 g), but with a strain energy more than four times higher (1.245 J).

The performance plots in Figure 8 reaffirm the fact that in this case NSGA-II converges faster (Fig. 8a); while the hybrid algorithm is more capable of exploring the low-mass design regions (Fig. 8b)

Overall, increasing the number of patches enhances local tailoring capability but also



(a) Objective function convergence for the three-patch ESL-TE1 case.



(b) Pareto front comparison for the three-patch ESL-TE1 case.

Figure 7: (a) NSGA-II vs. Hybrid convergence and (b) final Pareto front for the three-patch ESL-TE1 configuration.

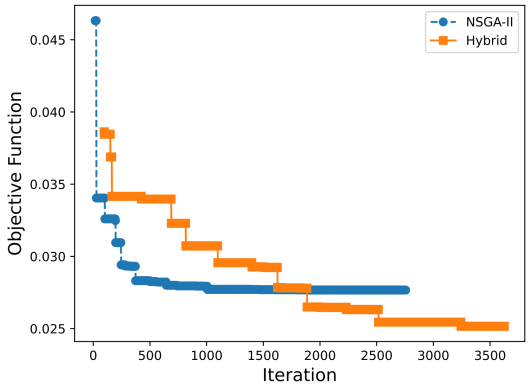
Table 8: Three-patch optimum results for LW-LD3.

	NSGA-II	**Hybrid**
No. of plies (Patch I)	11	7
No. of plies (Patch II)	7	4
No. of plies (Patch III)	4	4
Mass [kg]	0.0347	0.0237
Strain [J]	0.2794	1.2453
Angles [°]	-0.77, 3.44, 1.83, -4.93, 8.23, -3.45, 0.69, -2.00, 2.75, 1.01, -1.80	-1.78, 2.17, 49.70, -0.07, -2.05, -11.69, 7.59

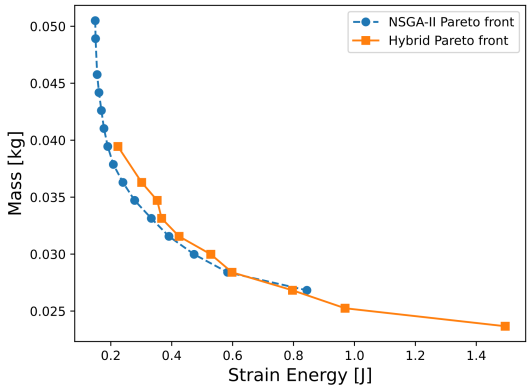
introduces a more complex optimization landscape. Across all cases, the Hybrid method consistently reaches feasible designs with significantly lower mass in reduced computational time, while NSGA-II excels in producing low-strain solutions and a broader Pareto front, demonstrating the intrinsic trade-off between convergence speed and exploratory capacity in multi-objective optimization of composite structures.

5 CONCLUSIONS

A hybrid optimization strategy has been presented for designing composite wing-box structures, coupling evolutionary search methods with a gradient-based local refinement. Central to this approach is the CUF, which provides an analytical path for evaluating the sensitivity of the objective function with respect to the lamination parameters. Specifically, the gradient of the strain energy is computed by differentiating the fundamental nuclei of the stiffness matrix, enabling the construction of the $\mathbf{D}_K$ matrix from a single run of the CUF-based solver rather than relying on multiple finite element simulations for finite-difference approximations. This feature reduces computational overhead and facilitates the easy switch between ESL and LW modeling, thereby scaling from low- to



(a) Objective function convergence for the three-patch LW-LD3 case.



(b) Pareto front comparison for the three-patch LW-LD3 case.

Figure 8: (a) NSGA-II vs. Hybrid convergence and (b) final Pareto front for the three-patch LW-LD3 configuration.

high-fidelity structural models.

The numerical results demonstrated that the hybrid method converges faster than a pure evolutionary scheme, often favoring lighter laminates at the expense of higher strain energy. Meanwhile, NSGA-II explores a more diverse region of the design space and can discover configurations with lower strain energy, albeit with additional computational cost. Extending the optimization from single-patch to multi-patch laminates revealed further trade-offs in the mass-strain performance, underscoring the importance of localized tailoring in composite design. Across all cases, the direct gradient computation via CUF proved instrumental in delivering rapid and accurate sensitivity data, supporting both ESL and LW formulations under a unified implementation. This flexibility suggests that the proposed framework can be readily scaled to more complex geometries and higher-fidelity analyses.

ACKNOWLEDGEMENTS

This work is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 850437). The support of the Italian Ministry of University and Research through the program FARE-III Edition (project LOUD, No. R20EENHZEJ) is also acknowledged.

REFERENCES

[1] R. T. Haftka and J. L. Walsh. Stacking-sequence optimization for buckling of laminated plates by integer programming. *AIAA Journal*, 30(3):814–819, March 1992. doi:10.2514/3.10989.

[2] R. L. Riche and R. T. Haftka. Optimization of laminate stacking sequence for buckling load maximization by genetic algorithm. *AIAA Journal*, 31(5):951–956, May 1993. doi:10.2514/3.11710.

[3] R. L. Riche and R. Haftka. Improved genetic algorithm for minimum thickness

composite laminate design. *Composites Engineering*, 5(2):143–161, February 1995. doi:10.1016/0961-9526(95)90710-S.

[4] F. Irisarri, A. Lasseigne, F. Leroy, and R. L. Riche. Optimal design of laminated composite structures with ply drops using stacking sequence tables. *Composite Structures*, 107:559–569, January 2014. doi:10.1016/j.compstruct.2013.08.030.

[5] H. Fukunaga and H. Sekine. Stiffness design method of symmetric laminates using lamination parameters. *AIAA Journal*, 30(11):2791–2793, 1992.

[6] G. Serhat and I. Basdogan. Lamination parameter interpolation method for design of manufacturable variable-stiffness composite panels. *AIAA Journal*, 57(7):3052–3065, July 2019. doi:10.2514/1.J057902.

[7] M. A. Nik, K. Fayazbakhsh, D. Pasini, and L. Lessard. Surrogate-based multi-objective optimization of a composite laminate with curvilinear fibers. *Composite Structures*, 94(8):2306–2313, August 2012. doi:10.1016/j.compstruct.2012.03.021.

[8] M. A. Nik, K. Fayazbakhsh, D. Pasini, and L. Lessard. Optimization of variable stiffness composites with embedded defects induced by automated fiber placement. *Composite Structures*, 107:160–166, January 2014.

[9] G. Silva, A. P. do Prado, P. H. Cabral, R. D. Breuker, and J. Dillinger. Tailoring of a composite regional jet wing using the slice and swap method. *Journal of Aircraft*, 56(3):990–1004, March 2019. doi:10.2514/1.C035094.

[10] D. Yang and Q. Fan. Gradient-based hybrid method for multi-objective optimization problems. *Expert Systems with Applications*, 272:126675, 2025. doi:10.1016/j.eswa.2025.126675.

[11] P. Borwankar, W. Zhao, R. K. Kapania, and M. Bansal. Two-level weight optimization of composite laminates using integer programming. *AIAA Journal*, 60(11): 6436–6446, 2022. doi:10.2514/1.J061560.

[12] H. A. Taymaz. Helicopter rotor blade vibration reduction with optimizing the structural distribution of composite layers. *Journal of Measurements in Engineering*, 10 (1):27–37, 2022. doi:10.21595/jme.2022.22337.

[13] K. Singh and R. K. Kapania. ALMO: Active learning-based multi-objective optimization for accelerating constrained evolutionary algorithms. *Applied Sciences*, 14 (21):9975, 2024. doi:10.3390/app14219975.

[14] E. Carrera and E. Zappino. One-dimensional finite element formulation with node-dependent kinematics. *Computers Structures*, 192:114–125, 2017. doi:10.1016/j.compstruc.2017.07.008.

[15] E. Zappino, G. Li, A. Pagani, and E. Carrera. Global-local analysis of laminated plates by node-dependent kinematic finite elements with variable esl/lw capabilities. *Composite Structures*, 172:1–14, 2017. doi:10.1016/j.compstruct.2017.03.057.

[16] E. Carrera, A. Pagani, and S. Valvano. Multilayered plate elements accounting for refined theories and node-dependent kinematics. *Composites Part B: Engineering*, 114:189–210, 2017. doi:10.1016/j.compositesb.2017.01.022.

[17] M. Filippi and E. Carrera. Aerodynamic and mechanical hierarchical aeroelastic analysis of composite wings. *Mechanics of Advanced Materials and Structures*, 23(9): 997–1004, 2016. doi:10.1080/15376494.2015.1121561.

[18] E. Carrera and A. Pagani. Accurate response of wing structures to free-vibration, load factors, and nonstructural masses. *AIAA Journal*, 54(1):227–241, 2016. doi:10.2514/1.J054164.

[19] M. Petrolo. Flutter analysis of composite lifting surfaces by the 1d carrera unified formulation and the doublet lattice method. *Composite Structures*, 95:539–546, 2013. doi:10.1016/j.compstruct.2012.06.021.

[20] A. Pagani, A. R. R. Sánchez-Majano, D. Zamani, M. Petrolo, and E. Carrera. Fundamental frequency layer-wise optimization of tow-steered composites considering gaps and overlaps. *Aerotecnica Missili Spazio*, 104:135–151, 2025. doi:10.1007/s42496-024-00212-w.

[21] A. Pagani, A. R. Sánchez-Majano, and D. Zamani. Optimization of variable stiffness composites considering gaps/overlaps and unified structural theories. In *ASME 2024 Aerospace Structures, Structural Dynamics, and Materials Conference*, volume V001T01A016, Renton, Washington, USA, April 29–May 1 2024. doi:10.1115/SSDM2024-121550.

[22] A. R. Sánchez-Majano and A. Pagani. Buckling and fundamental frequency optimization of tow-steered composites using layerwise structural models. *AIAA Journal*, 61 (9):4149–4163, 2023. doi:10.2514/1.J062976.

[23] D. Zamani. Least-weight design of tow-steered composites. In *Materials Research Proceedings*, volume 42, pages 51–55, 2024.

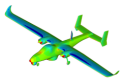
[24] E. Carrera, M. Cinefra, M. Petrolo, and E. Zappino. *Finite Element Analysis of Structures through Unified Formulation*. Wiley & Sons, Hoboken, NJ, 2014.

[25] J. N. Reddy. *Mechanics of Laminated Composite Plates and Shells: Theory and Analysis*. CRC Press, Boca Raton, Florida, USA, 2004.

[26] E. Carrera. Cz0 requirements: models for the two dimensional analysis of multilayered structures. *Composite Structures*, 37(3-4):373–383, 1997. doi:10.1016/S0263-8223(98)80005-6.

[27] Z. Hashin. Failure criteria for unidirectional fiber composites. *Journal of Applied Mechanics*, 47(2):329–334, 1980. doi:10.1115/1.3153664.

[28] M. Kameyama and H. Fukunaga. Optimum design of composite plate wings for aeroelastic characteristics using lamination parameters. *Computers & Structures*, 85: 213–224, 2007.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

DEEP LEARNING-BASED FLOW FIELD RECONSTRUCTION FOR M219 CAVITY FLOW

Berkay Pirlepeli[1*], Burak Berkan Bedir[1] and Melike Nikbay[2]

1: AeroMDO Lab
Faculty of Aeronautics and Astronautics
Istanbul Technical University
34469
pirlepeli18@itu.edu.tr, bedirb16@itu.edu.tr

2: Astronautical Engineering Department
Faculty of Aeronautics and Astronautics
Istanbul Technical University
34469
nikbay@itu.edu.tr

Abstract. *Modern military aircraft weapon bays require accurate flow field and acoustic predictions for design optimization. While the computational fluid dynamics methods provide high-fidelity simulations, they demand significant computational resources and time. Conventional reduced-order modeling approaches such as proper orthogonal decomposition and dynamic mode decomposition can accelerate predictions but often struggle to capture complex nonlinear dynamics and unsteady characteristic of cavity flows. This study presents a deep learning framework to accelerate these predictions while maintaining accuracy, specifically designed for sensitivity analysis and uncertainty quantification studies of transonic cavity flows. The present framework is developed using the M219 cavity configuration by validating a Detached Eddy Simulation model implemented in OpenFOAM against the experimental results available in the literature. The presented architecture learns nonlinear mappings between limited pressure measurements enabling more accurate reconstruction of flow field and preserving coherent structures and turbulent fluctuations compared to traditional modal decomposition methods. The model is trained using a comprehensive dataset spanning different Mach numbers and cavity geometries to ensure robustness. The framework provides faster flow field predictions for uncertainty quantification and sensitivity analysis compared to full computational fluid dynamics simulations, while maintaining higher accuracy than conventional reduced-order models.*

Keywords: Deep learning, Cavity flow, M219 cavity, Flow reconstruction, Surrogate model

1 INTRODUCTION

Cavity flow has been an important engineering challenge for both military and commercial aircraft since the 1940's [1]. Cavity flow represents a complex fluid dynamics phenomenon characterized by the interaction between airflow and cavity geometries. When air passes over a cavity area, such as weapon bays or landing gear, creates distinctive flow patterns with significant aerodynamic and acoustic implications. When flow passes over a cavity, the shear layer separates at the leading edge and reattaches at the trailing edge. This creates a complex flow environment with vortical structures that impinge on the trailing edge, generating acoustic waves that propagate upstream and trigger further instabilities at the leading edge. This complex interaction generates intense turbulence, pressure fluctuations, and acoustic resonance effects. At transonic and supersonic speeds, these phenomena intensify through self-sustaining feedback mechanisms, where acoustic waves generated at the trailing edge propagate upstream to influence shear layer development, creating a cycle that amplifies noise production, induces structural vibrations, and potentially compromises aircraft performance and operational capabilities. [1, 2].

Experimental investigation of cavity flows presents significant challenges, including prohibitive costs and complex testing environments. Among the significant contributions to this field, the M219 cavity model tests, conducted in 2000 [3]. In addition to experimental approaches, Computational Fluid Dynamics (CFD) offers a valuable alternative that enhances design efficiency while providing deeper insights into underlying flow physics. However, even numerical approaches demand substantial computational resources, particularly for high-fidelity simulations. Various turbulence modeling strategies exist for these complex flows, ranging from Reynolds-Averaged Navier-Stokes (RANS) to Direct Numerical Simulation (DNS), with selection criteria balancing accuracy against computational demands. In practical aerospace applications, RANS and hybrid methods like Detached Eddy Simulation (DES) are frequently employed due to their favorable accuracy-to-cost ratio. The DES approach, especially when coupled with Spalart-Allmaras or two-equation RANS formulations, has demonstrated excellent correlation with experimental data while requiring significantly fewer computational resources than more intensive Large Eddy Simulation (LES) approaches.[4, 5].

Machine learning and particularly deep learning have emerged as transformative technologies in CFD research and applications related to cavity flows [6]. These advanced computational approaches can capture the complex, nonlinear relationships between flow parameters that govern aeroacoustic phenomena in cavity configurations. By leveraging extensive training datasets, deep learning models can efficiently predict unsteady flow characteristics within cavities—including velocity fields, pressure distributions, and turbulence patterns—with significantly reduced computational expense compared to traditional high-fidelity simulations [7]. This predictive capability represents a substantial advance in aerospace engineering, where accurate flow behavior analysis and noise source identification have historically challenged conventional numerical methods, especially in transonic regimes [8]. The application spectrum extends beyond pure analysis; machine learning algorithms trained on experimental or flight test acoustic data provide rapid estimation of noise levels across multiple aircraft configurations, facilitating design optimization [9]. Deep learning frameworks have successfully addressed cavity geometry optimization, weapon bay door configuration refinement, and optimal placement of flow control devices—all critical aspects of noise mitigation in military aircraft [10, 11]. The

technology further enhances store separation prediction accuracy by efficiently modeling the highly complex aerodynamic interactions within cavity environments [12], while simultaneously enabling sophisticated control strategies through intelligent actuation systems that dynamically respond to flow conditions [13].

Proper Orthogonal Decomposition (POD) [14, 15] has established itself as a foundational technique in fluid dynamics for dimensionality reduction of complex flow data. By extracting the most energetically significant spatial patterns from high-dimensional flow snapshots, POD creates an optimal basis for efficient representation of the underlying flow physics. This modal decomposition approach has evolved through various innovative applications: Qamar and Sanghi demonstrated POD's effectiveness in reconstructing parametrically-varying supersonic flow fields with minimal computational overhead [16]; Troshin et al. extended POD capabilities to domains with dynamic boundaries through volume-preserving transformations, validating their methodology on oscillating ellipse configurations [17]; while Nakamura et al. further advanced POD methodology by developing a unified framework that enables flow reconstruction across varying geometries by mapping diverse computational grid data into a standardized space, demonstrating remarkable predictive accuracy for unsteady flows around both elliptical bodies and aerodynamic profiles [18].

Flow field reconstruction from sparse or noisy measurements represents a fundamental challenge in fluid dynamics applications. Traditional reconstruction approaches based on least-squares regression often produce suboptimal results when confronted with measurement limitations or signal contamination. Callaham et al. advanced this field by implementing sparse representation techniques that significantly enhance reconstruction accuracy and robustness from limited sensor data, while introducing an innovative patch-based methodology particularly effective for multiscale flow phenomena [19]. Erichson et al. developed a shallow neural network (SNN) architecture that eliminates the need for prior physical knowledge by learning direct mappings from raw sensor measurements to complete flow fields demonstrating superior performance compared to conventional modal decomposition techniques while requiring significantly fewer measurement points [20].

This study presents an application of the SNN methodology for cavity flow field reconstruction using limited sensor data. The research workflow encompasses numerical model validation, data generation, and neural network-based flow reconstruction. Using the M219 cavity geometry benchmark case, the reconstruction model is trained snapshots from high-fidelity DES. The numerical approach is first validated against established experimental data from Henshaw's comprehensive M219 cavity experiments [3]. The resulting pressure field data serves dual purposes: training a SNN decoder for efficient reconstruction and establishing baseline performance through POD modes, enabling comparative analysis of reconstruction accuracy with limited sensor inputs.

2 METHODOLOGY

This section presents the methodology for aeroacoustic cavity flow reconstruction using pressure sensor data. The workflow utilizes the first 700 snapshots from high-fidelity Improved Delayed Detached Eddy Simulations (IDDES) for model training, while reserving the remaining 265 snapshots for independent performance evaluation. Pressure field data generated through these simulations serves dual purposes: training a SNN decoder architecture for efficient flow reconstruction, and POD models for comparative perfor-

mance assessment. Both approaches enable reconstruction of complete flow fields from limited sensor inputs, with comparative analysis highlighting the relative strengths of each method under various sensing configurations and flow conditions.

2.1 M219 Cavity Geometry

The M219 cavity represents a widely studied benchmark geometry in aeroacoustic research, providing standardized experimental data essential for validating computational approaches. Conducted in 1991 at the Aircraft Research Association wind tunnel in Bedford through a collaborative BAe/DERA initiative, these experiments examined fluid-acoustic interactions within a precisely dimensioned rectangular cavity configuration. The test setup featured a rectangular cavity with dimensions of 20 inches × 4 inches and two depth configurations: shallow (2 inches) and deep (4 inches) [3]. Unsteady pressure measurements were systematically captured using 10 Kulite transducers positioned along the cavity ceiling centerline shown in Fig. 2. These measurements were sampled at 6000 Hz under specified flow conditions, in Table 1. The experiments encompassed three critical flow regimes—subsonic (Mach 0.6), transonic (Mach 0.85), and supersonic (Mach 1.35)

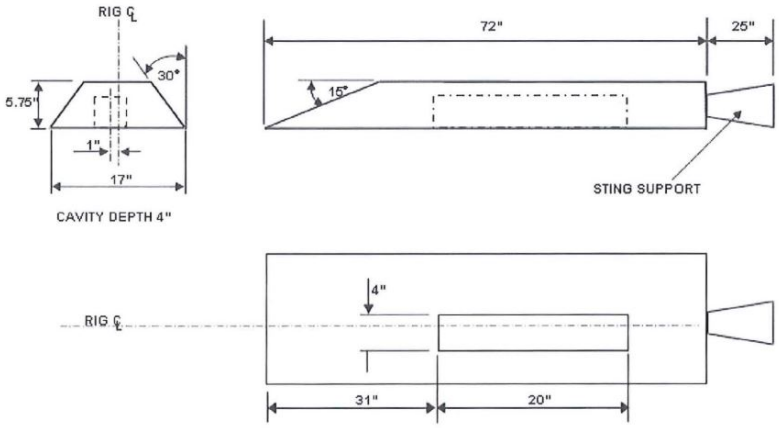


Figure 1: Experimental Rig of the M219 Cavity Model [3]

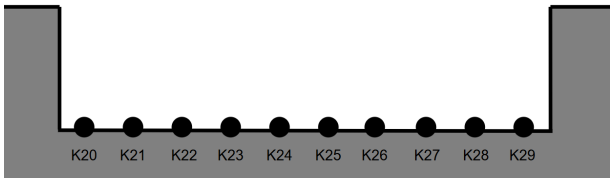


Figure 2: 10 Transducer Location in cavity floor surface

This study focuses on the deep cavity configuration (L/D = 5) at a Mach number of 0.85, representing transonic flow conditions. The locations of the transducers are given in Table 2.

2.2 Numerical Simulation

The computational domain for cavity flow simulation, shown in Fig. 3, encompasses a region extending five cavity lengths vertically above the plate and two cavity lengths

Table 1: Flow Conditions

Description	Value	Unit
Mach	0.85	
Total Pressure	99612.06	[Pa]
Pressure	62059	[Pa]
Total Temperature	305.06	[K]
Density	0.811741	$[kg/m^3]$
Velocity	278.197	$[m/s]$
Reynolds Number	6.84×10^6	

Table 2: Kulite Transducer Positions

Transducer Number	Distance from LE [m]
K20	0.0254
K21	0.0762
K22	0.1270
K23	0.1778
K24	0.2286
K25	0.2794
K26	0.3302
K27	0.3810
K28	0.4318
K29	0.4826

horizontally on each side. To accurately capture viscous effects, no-slip boundary conditions are enforced on all solid surfaces of both the cavity and the test rig. The domain boundaries utilize symmetry conditions, ensuring accurate flow physics. The numerical simulations for transonic cavity flow are performed using OpenFOAM, a finite-volume solver capable of handling both transient and steady-state flow computations [21]. For turbulence modeling, the computations utilize the IDDES methodology coupled with the k-ω SST turbulence model which provides an effective balance between computational efficiency and accuracy for cavity flows. To ensure numerical accuracy, the spatial discretization employs second-order gradient schemes based on the least-square method for calculating non-viscous fluxes, capturing the complex flow physics while maintaining solution stability.

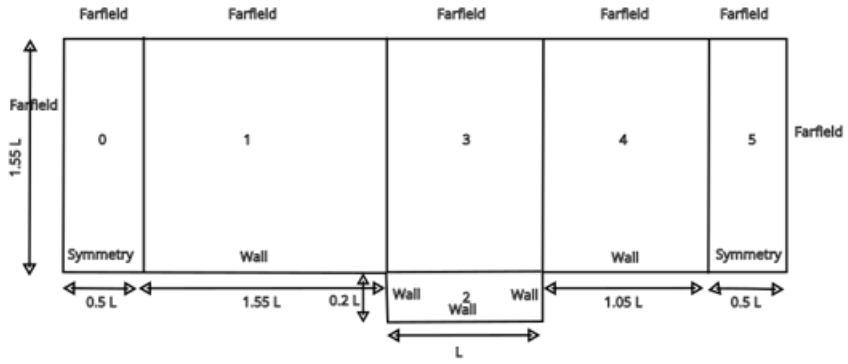


Figure 3: Computational Domain with Boundary Conditions

The aeroacoustic performance of the 2D cavity is evaluated based on the pressure signals recorded at probe locations along the cavity floor. The Overall Sound Pressure Level (OASPL) is used to quantify the total noise level across all frequencies of the pressure signal. Following numerical simulations, the root mean square of the pressure fluctuations (p'_{rms}) and the OASPL at the probe locations are computed as follows:

$$p'rms = \sqrt{\frac{1}{N}\sum p'^2}, \qquad OASPL = 20\log 10\left(\frac{p_{rms}}{p_{ref}}\right) \tag{1}$$

where p_{ref} is the reference pressure, set at 2×10^{-5} Pa. The spatial discretization of the computational domain, illustrated in Figure 3, is generated using the blockMesh tool in OpenFOAM. The mesh is predominantly unstructured hexahedral, with 15 prism layers in the boundary layer region, ensuring a y^+ value of less than 1. The estimated first-layer height is 6.84×10^{-6} m.

To validate the model, three different mesh resolutions are generated to ensure grid independence, with element counts summarized in Table 3. The simulation timestep is determined through careful consideration of both numerical stability and acoustic resolution requirements. To properly capture the high-frequency content of cavity acoustics, particularly the 4th Rossiter mode, a target frequency of 5 kHz is established—approximately five times greater than the mode frequency. Applying the Nyquist criterion [22], a sampling frequency of 10 kHz results in a timestep of 1×10^{-4} seconds. To ensure the accuracy of the CFD simulations, the timestep is further refined to resolve at least five points per wavelength, leading to a final timestep of 2×10^{-5} seconds.

Table 3: Number of Elements

	Number of Cells	Cells in Cavity
Mesh1 (Coarse)	41832	14874
Mesh2 (Fine)	150016	38138
Mesh3 (Very Fine)	384464	287455

A computational mesh comprising 151,278 pressure points was constructed for the flow domain to capture the complex aeroacoustic phenomena within the cavity. From the numerical simulations, 925 temporal snapshots were collected at regular intervals, providing a comprehensive dataset of the time-evolving flow field. To develop a robust predictive model, 700 snapshots (approximately 75% of the data) were allocated for training purposes, while the remaining 225 snapshots were reserved for testing and validation. As shown in Fig. 2, ten discrete transducer points were used as limited measurement locations, mirroring the experimental setup and providing the input data for the machine learning model.

2.3 Proper orthogonal decomposition

POD is one of the most established techniques for reconstructing high-dimensional fluid flow fields from sparse sensor measurements[20]. The method creates a low-dimensional approximation of the complete flow field using a set of spatially orthogonal modes extracted through Singular Value Decomposition (SVD) of the training dataset. The fundamental approach involves two key steps: first, computing the dominant spatial modes that capture the essential flow physics; and second, estimating the time-dependent coefficients for these modes using the limited sensor data. By projecting these sensor measurements onto the POD basis, the full flow field can be reconstructed with reasonable accuracy, though the performance depends significantly on sensor placement and the complexity of the underlying flow structure.

The POD modes are derived by applying SVD to a mean-centered training dataset $X = (x_1, \ldots, x_n) \in \mathbb{R}^{m \times n}$, where $n \leq m$:

$$X = U \Sigma V^T \tag{2}$$

where $U \in \mathbb{R}^{m\times n}$ contains the left singular vectors (POD spatial modes), $V \in \mathbb{R}^{n\times n}$ consists of the right singular vectors, $\Sigma \in \mathbb{R}^{n\times n}$ is a diagonal matrix with singular values. The reduced-order approximation modes are selected as:

$$\Phi := U_k \tag{3}$$

where U_k consists of the top k singular vectors, typically chosen based on approximation modes.

In the standard POD approach, the high-dimensional flow field x is approximated using the reduced basis Φ and corresponding mode coefficients ν. Given a linear measurement operator $H : \mathbb{R}^m \to \mathbb{R}^p$, which relates the full state to sensor measurements:

$$s = Hx \tag{4}$$

the field approximation is given by:

$$x \approx \hat{x} = \Phi\nu \tag{5}$$

where ν is obtained by solving the least-squares problem:

$$\nu = \arg\min_{\tilde{\nu}} \|s - H\Phi\tilde{\nu}\|_2^2 \tag{6}$$

The optimal solution, using the Moore-Penrose pseudo-inverse, is:

$$\nu = (H\Phi)^+ s \tag{7}$$

Thus, the final reconstructed state is:

$$\hat{x} = \Phi\nu \tag{8}$$

2.4 Shallow Decoder

In this study, a Shallow Decoder (SD) architecture is employed based on the approach proposed by Erichson et al [20]. In Figure 4, the SD model employs a SNN architecture with only a few hidden layers, mapping sensor inputs to the flow field through a simple structure.

$$s \to fist\ hidden\ layer \to second\ hidden\ layer \to \hat{x}$$

Mathematically, the SD can be formulated as a nested set of functions:

$$F(s; W) := \Omega(W_K R(W_{K-1} R(W_1 s + b_1) + b_{K-1}) + b_K)$$

where $s \in \mathbb{R}^p$ represents the sensor measurements, $W = \{W_k\}_{k=1}^K$ are weight matrices with appropriate dimensions, b_k are bias terms, and $R(\cdot)$ is a nonlinear activation function (typically ReLU, defined as $R(z) := \max(z, 0)$). For the SD, K is typically 2 or 3, indicating a shallow architecture. Tthe output layer Ω, which is linear and maps the latent representation to the high-dimensional flow field $\hat{x} \in \mathbb{R}^m$. This can be expressed as:

$$\hat{x} = \Omega(z_\nu) := \Phi z_\nu + b_\Phi$$

where Φ represents the learned modes and z_ν the coefficients. The SD is trained by minimizing the mean squared error between the estimated and true flow fields:

$$\min_W \sum_{i=1}^{n} \|x_i - F(s_i; W)\|_2^2 + \lambda \|W\|_2^2$$

where λ is a regularization parameter.

The SD offers several advantages over traditional flow reconstruction techniques. First, it features a linear last layer which provides a supervised joint learning framework for the low-dimensional approximation space of the flow field and the mapping from measurements to this space. This allows the approximation basis to be tailored not only to the state space but also to the associated measurements, preventing observability issues common in traditional methods. Second, the SD allows for flexibility in measurements, which don't necessarily have to be linearly related to the state. Finally, the SD produces interpretable features of the dynamics, potentially improving on classical POD modes by not constraining the features to be orthogonal.

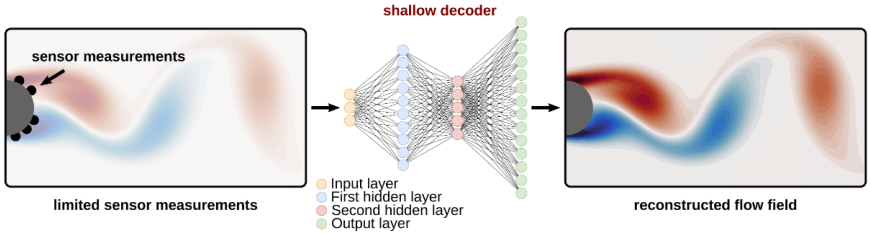


Figure 4: Shallow Neural Network for Fluid Flow Reconstruction with Limited Sensors [20].

3 RESULT

The OASPL values along cavity floor, at 10 probe locations along the cavity floor shown in Fig. 5. The increasing trend in pressure fluctuations from leading to trailing edge, consistent with known cavity flow physics where acoustic waves intensify toward the rear wall. All three mesh resolutions follow similar patterns, with the finest mesh (Mesh3) providing closest agreement with experimental data, though all simulations overpredict noise levels by approximately 3-4 dB. This systematic offset likely results from turbulence modeling limitations or idealized boundary conditions, but the consistent shape of the distribution confirms that our computational approach successfully captures the fundamental flow structures and acoustics phenomena.

Figure 6 presents the first nine POD modes extracted from the snapshot matrix of the pressure field of train data set. These modes represent the most energetic spatial structures in the data, ordered by their contribution to the total energy. The dominant flow features are captured in the first few modes, with higher modes capturing increasingly finer-scale structures and less energetic components. This modal decomposition provides a compact and interpretable basis for reduced-order modeling and flow reconstruction.

Figure 7 presents the relative reconstruction error for training and test datasets as a function of the number of modes used in POD model. The error initially decreases with the number of modes, indicating improved reconstruction. However, a sharp increase is

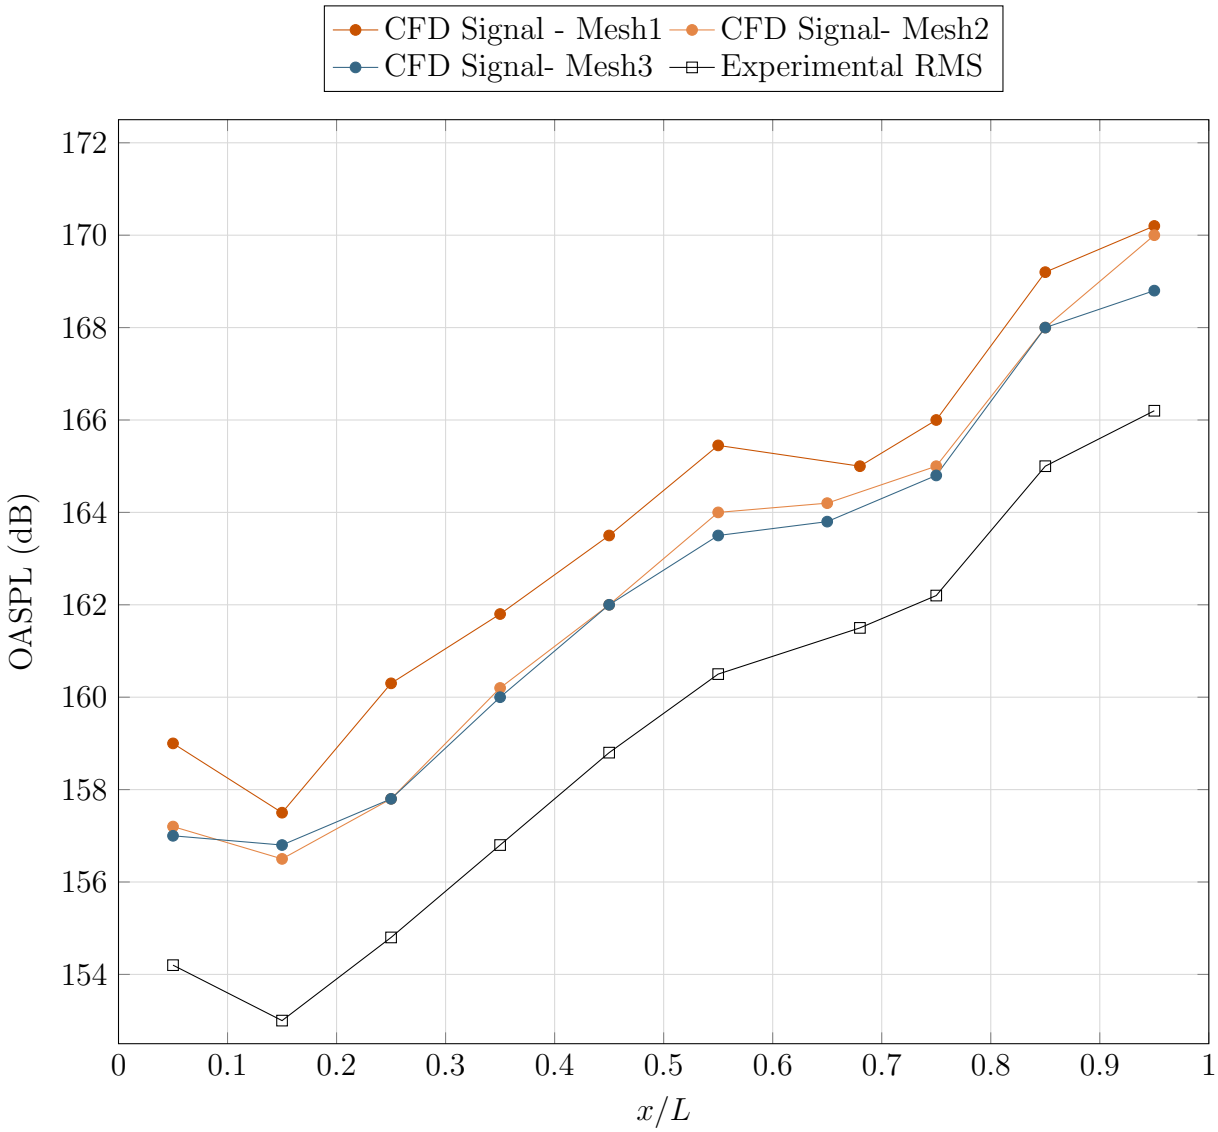


Figure 5: OASPL Along Cavity Floor (M=0.85).

observed at 8 modes and the system becomes ill-posed near sensor number taken as 10.

To evaluate the POD and Shallow Decoder models, the same training and testing datasets were used. Additionally, identical transducer (sensor) locations in Figure 2 were employed for both models. To evaluate the effectiveness of different reconstruction methods, we computed the relative reconstruction error for the Shallow Decoder, and POD methods. The results are summarized in Table 4.

Table 4: Relative reconstruction error for different methods.

	Relative Error
Shallow Decoder	0.0396
POD	0.1855

Figure 8 indicates that POD fails to capture high-order structures, leading to larger reconstruction errors, while the Shallow Decoder effectively retains more spectral information, enabling more accurate state estimation.

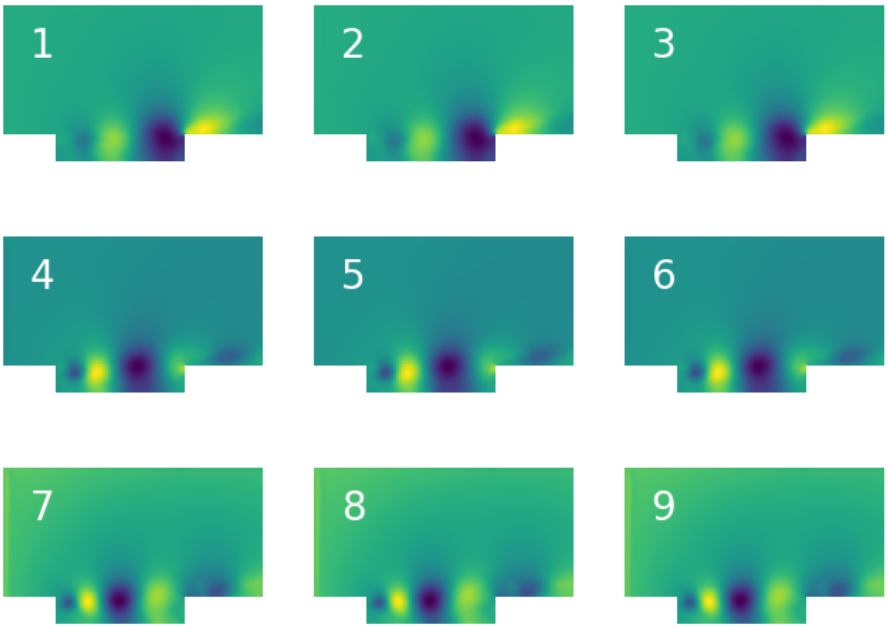


Figure 6: First 9 POD modes of Cavity Flow

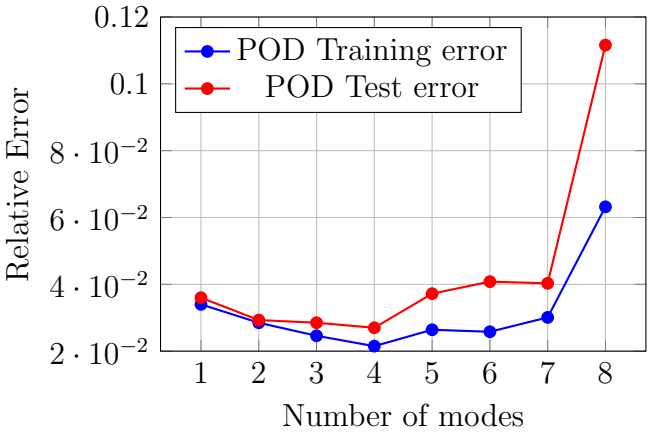


Figure 7: Relative errors of POD training and test data as a function of the number of modes.

4 CONCLUSION

This study demonstrated the efficacy of shallow neural networks for reconstructing complex cavity flow fields from limited sensor measurements. By implementing and comparing both POD and SD architectures on the benchmark M219 cavity dataset, the machine learning approaches offer significant advantages in flow field reconstruction accuracy. While conventional POD methods showed adequate performance with a relative error of 0.1855, the analysis of spectral information retention revealed their limitations in capturing higher-order flow structures. In contrast, the Shallow Decoder architecture, with its ability to learn nonlinear mappings between sensor measurements and flow fields without enforcing orthogonality constraints, demonstrated reconstruction accuracy with a relative error of 0.0396.

The singular spectrum analysis results provided critical insights into why the SD archi-

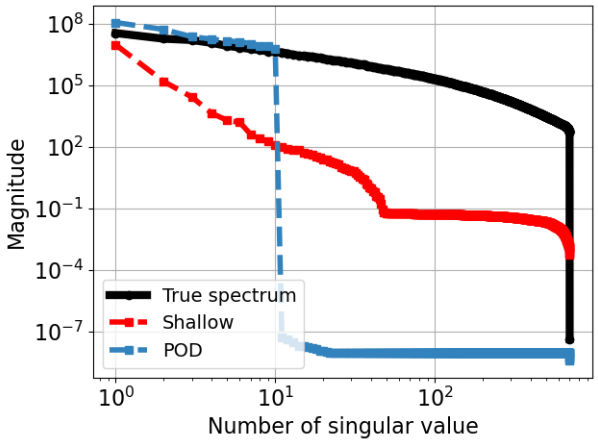


Figure 8: Singular spectrum analysis of reconstructed data with POD and Shallow Decoder

tecture maintains higher fidelity across a broader range of flow structures. This capability is particularly valuable for transonic cavity flows, where accurate prediction of complex aeroacoustic phenomena is essential for aerospace applications. Future work will focus on extending this methodology to more complex flow conditions, optimizing sensor placement strategies, and exploring the integration of physics-informed constraints to further enhance reconstruction accuracy.

Acknowledgments

All authors would like to thank Turkish Aerospace Industries (TAI) for the research grant provided under "ITU1773 TTO TM1061 - Modeling Unsteady Flow Characteristics of Supersonic Cavity Flow via Deep Learning".

All authors would like to thank ITU UHEM National Center for High Performance Computing for the Project with Grant Number 5019802024, titled "Modeling Unsteady Flow Characteristics of Supersonic Cavity Flow via Deep Learning".

REFERENCES

[1] S. Lawson and G. N. Barakos. Review of numerical simulations for high-speed, turbulent cavity flows. *Progress in Aerospace Sciences*, 47(3):186–216, 2011.

[2] C. O. H. Smith, N. Lawson, and G. A. Vio. History, review and summary of the cavity flow phenomena. *European Journal of Mechanics-B/Fluids*, 2024.

[3] M. Henshaw. M219 cavity case. 2000. URL `https://api.semanticscholar.org/CorpusID:107486625`.

[4] S. Lawson and G. Barakos. Assessment of passive flow control for transonic cavity flow using detached-eddy simulation. *Journal of aircraft*, 46(3):1009–1029, 2009.

[5] S. A. Syed and K. Hoffmann. Numerical investigation of 3-d open cavity with and without cover plates. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, page 551, 2009.

[6] S. Le Clainche, E. Ferrer, S. Gibson, E. Cross, A. Parente, and R. Vinuesa. Improving aircraft performance using machine learning: A review. *Aerospace Science and Technology*, 138:108354, 2023.

[7] Z. Zhang, J. Wang, R. Huang, R. Qiu, X. Chu, S. Ye, Y. Wang, and Q. Liu. Data-driven turbulence model for unsteady cavitating flow. *Physics of Fluids*, 35(1), 2023.

[8] S. Le Clainche, E. Ferrer, S. Gibson, E. Cross, A. Parente, and R. Vinuesa. Improving aircraft performance using machine learning: A review. *Aerospace Science and Technology*, 138:108354, 2023.

[9] Y. Zhang, B. Chen, K. Zhao, X. Tang, X. Yang, and G. Hu. Intelligent prediction of acoustic performance of landing gear using deep learning. *Physics of Fluids*, 35(7), 2023.

[10] J. Li, X. Du, and J. R. Martins. Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134:100849, 2022.

[11] W. Wang and J. Ma. A review: Applications of machine learning and deep learning in aerospace engineering and aero-engine engineering. *Advances in Engineering Innovation*, 6:54–72, 2024.

[12] N. Peters. A data driven modeling approach for store distributed load and trajectory prediction. 2022.

[13] Y. Dong, J. Tao, Y. Zhang, W. Lin, and J. Ai. Deep learning in aircraft design, dynamics, and control: Review and prospects. *IEEE Transactions on Aerospace and Electronic Systems*, 57(4):2346–2368, 2021.

[14] G. Berkooz, P. Holmes, and J. L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics*, 25(1):539–575, 1993.

[15] P. Holmes. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge university press, 2012.

[16] A. Qamar and S. Sanghi. Steady supersonic flow-field predictions using proper orthogonal decomposition technique. *Computers & fluids*, 38(6):1218–1231, 2009.

[17] V. Troshin, A. Seifert, D. Sidilkover, and G. Tadmor. Proper orthogonal decomposition of flow-field in non-stationary geometry. *Journal of Computational Physics*, 311: 329–337, 2016.

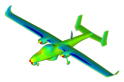
[18] Y. Nakamura, S. Sato, and N. Ohnishi. Application of proper orthogonal decomposition to flow fields around various geometries and reduced-order modeling. *Computer Methods in Applied Mechanics and Engineering*, 432:117340, 2024.

[19] J. Callaham, K. Maeda, and S. L. Brunton. Robust flow field reconstruction from limited measurements via sparse representation. *arXiv preprint arXiv:1810.06723*, 2018.

[20] N. B. Erichson, L. Mathelin, Z. Yao, S. L. Brunton, M. W. Mahoney, and J. N. Kutz. Shallow neural networks for fluid flow reconstruction with limited sensors. *Proceedings of the Royal Society A*, 476(2238):20200097, 2020.

[21] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. A tensorial approach to computational continuum mechanics using object-oriented techniques. *Computer in Physics*, 12(6):620–631, 11 1998. ISSN 0894-1866. doi:10.1063/1.168744. URL `https://doi.org/10.1063/1.168744`.

[22] H. Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928. doi:10.1109/T-AIEE.1928.5055024.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

IMPACT OF ADOPTING ACTIVE FLUTTER SUPPRESSION SYSTEMS ON THE OPTIMIZATION OF A TYPICAL LONG-RANGE TRANSPORT AIRCRAFT

Elena Roncolini*, Francesco Toffol and Sergio Ricci

Department of Aerospace Science and Technology
Politechnico di Milano
Via La Masa 34, Milano - ITALY

*elena.roncolini@polimi.it, http://www.aero.polimi.it

{francesco.toffol, sergio.ricci}@polimi.it, http://www.aero.polimi.it

Abstract. *The increasing pressure to reduce the environmental impact of future aircraft is driving the adoption of more innovative aerostructural solutions, such as unconventional configurations and the use of active control systems to ensure high overall efficiency across the entire flight envelope. Implementing high aspect ratio wings to decrease induced drag, which are generally more flexible, requires advanced active control systems to mitigate potential aeroelastic issues. Active Wing Control systems are likely to be integrated with technologies currently utilized in existing fly-by-wire control systems. However, Active Flutter Suppression systems (AFS) are relatively new, although with a long history behind them, and their expected performance and technical feasibility need to be evaluated in terms of safety, robustness, and reliability. Despite this, two aircrafts with AWC have been certified under special conditions: the B747-800 and the B787-10, highlighting the necessity for further investigation into the adoption and certification of these systems. Setting aside the intricate certification considerations, a key design question remains: assuming stability can be assured with suitable AFS systems, what potential weight savings could be achieved by designing the structure without the explicit constraint of a minimum flutter velocity requirement? This paper focuses on a generic long-range, twin-aisle transport aircraft and presents comparisons based on optimizing the wing structure with various design constraints, including or excluding the flutter constraint, to draw general conclusions.*

Keywords: Active Flutter Suppression (AFS), aeroelastic optimization

1 INTRODUCTION

Flutter is a dynamic aeroelastic instability that can lead to divergent oscillatory vibrations, posing a significant risk to aircraft performance and safety, potentially resulting in aircraft loss [1]. In the traditional three-phase design process - conceptual, preliminary, and detailed - flutter analysis is typically conducted once a sufficiently accurate aeroelastic model is available [2]. However, conventional conceptual design tools often rely on statistical methods that are unable to accurately capture unsteady aerodynamic and structural stiffness effects, limiting their efficacy in flutter prediction [3].

In industrial practice, flutter analysis is usually performed once the Global Finite Element Model (GFEM) of the aircraft is developed. This model facilitates the creation of a reduced structural representation (stick model) for flutter calculations within the modal space [4]. If the design fails to meet flutter requirements, redesign efforts become necessary, often leading to increased costs and extended development timelines.

To mitigate these challenges, integrating flutter analysis early in the aircraft design process - particularly in wing design - can streamline development by enabling more informed trade-offs between flutter constraints, performance objectives, and other design considerations. Therefore, there is growing interest in incorporating flutter constraints into standard Multidisciplinary Design Optimization (MDO) frameworks [5]. This approach allows for the simultaneous optimization of aerodynamic and structural parameters, enabling structural weight and fuel consumption minimizations, while ensuring compliance with flutter requirements [6].

Despite ongoing research, integrating flutter constraints into optimization cycles remains a challenge [3]. Flutter is inherently complex and nonlinear, though often linearized for practical purposes, making it difficult to express it as an explicit constraint [7]. The computational model must balance complexity, being detailed enough to capture critical aeroelastic effects while remaining efficient for iterative design processes [8, 9]. Additionally, optimization methodologies vary widely, encompassing gradient-based and other methods, integer and continuous variables, and local versus global solution strategies.

Further considerations involve the inclusion of post-flutter behavior as a design requirement [10], or the role of Active Flutter Suppression (AFS) systems. While various AFS solutions have been explored over the years, certification regulations have historically adopted a conservative stance, prohibiting the certification of aircraft exhibiting flutter within the flight envelope, due to safety concerns [11].

Recently, two key developments have renewed interest in these topics.

First, the push for reduced environmental impact in future civil aircrafts has led to the exploration of unconventional and innovative configurations. High Aspect-Ratio Wings (HARW), designed to minimize induced drag during cruise, exhibit significant structural flexibility, making them more susceptible to aeroelastic phenomena, including flutter and associated nonlinearities [12]. Moreover, many proposed configurations incorporate distributed or unconventional propulsion systems, leading to an increased modal density in the wing structure compared to traditional designs. These challenges must be addressed from the earliest design stages.

Second, for the first time in years, there is recognition of the need to update certification processes to accommodate active flutter mitigation or suppression systems. Rather than being seen solely as a corrective measure for unexpected flutter behavior discovered during

certification, AFS is now being considered a fundamental design opportunity.

Active Wing Control (AWC) systems are likely to be integrated with technologies currently utilized in existing fly-by-wire control systems. However, AFS are relatively new, although they have a long history behind them [1], and their expected performance and technical feasibility need to be evaluated in terms of safety, robustness and reliability. Despite this, two AWC aircrafts have been certified under special conditions: the B747-800 and the B787-10 (see [13, 14]), highlighting the need for further investigation into the adoption and certification of these systems.

Recent research efforts have focused on the actuation, detection, redundancy, and validation aspects of AFS. The main objective is to confirm the feasibility of these systems in addressing unforeseen situations that arise during the final design phase, where they become the sole systems ensuring the minimum aeroelastic damping required for successful certification. Essentially, AFS systems are designed to rectify unintentional instability cases that cannot be resolved otherwise.

Setting aside the intricate certification considerations, a key design question remains: assuming stability can be assured with suitable AFS systems, what potential weight savings could be achieved by designing the structure without the explicit constraint of a minimum flutter velocity requirement?

This paper focuses on a generic long-range, twin-aisle transport aircraft and presents comparisons of wing structure optimizations with various design constraints, including or excluding flutter, to draw general conclusions. Exploring the integration of flutter constraints in modern aircraft design and optimization methodologies, the present work does not answer completely the many open issues, but it wants to open the discussion among the researchers in the field.

2 USED TOOLS

Recognizing the importance of integrating aeroelastic considerations from the initial design phase - particularly given the increasing flexibility of modern aircrafts - POLIMI started developing the NeoCASS design suite in 2007 [15–17]. NeoCASS was among the first physics-based conceptual design tools, aiming to provide:

1. a simplified yet realistic structural model, capable of capturing key aeroelastic characteristics through rapid analyses;

2. automated structural model generation, seamlessly integrated into conceptual design workflows;

3. a transparent user experience, minimizing the need for expertise in finite element and aeroelastic analysis.

NeoCASS employs a simplified finite element model - stick model with beam elements - along with a Vortex Lattice (VLM) or Doublet Lattice (DLM) model for steady and unsteady calculations, respectively.

Beyond the aeroelastic predictive capabilities, the early use of stick finite element models is mainly what diversifies NeoCASS from conventional industrial approaches. Indeed, while in industry one typically relies on statistical mass estimates during conceptual design

only introducing a GFEM, after selecting candidate configurations, NeoCASS integrates structural sizing from the outset. By using stick models for the full load set required by certification regulations, it offers physics-based structural weight estimation, eliminating reliance on statistical data often unavailable for novel technologies or unconventional aircraft configurations and aligning with aeroelastic requirements. The minimum structural mass is then determined solely by the material required to sustain applied loads. This contrasts with traditional workflows, where aeroelastic assessments occur late in the process, often necessitating design revisions that lead to weight penalties.

Over the past decade, several conceptual and preliminary design tools have emerged, including PROTEUS (TU-Delft), ATLASS (Gulfstream), Bombardier's proprietary methodology, and EMWET, among others. Many of them depend on GFEMs or commercial finite element software, limiting their applicability at the conceptual design stage, where hundreds or thousands of configurations must be evaluated. Differently, NeoCASS automatically generates stick models at the beginning of the design loop from a limited set of geometric inputs, sizing the aircraft based on real load conditions dictated by certification rules (e.g., CS23, CS25) [18, 19]. As anticipated, this approach enables a realistic mass and stiffness distribution for structural elements, while relying on statistical models for non-structural masses.

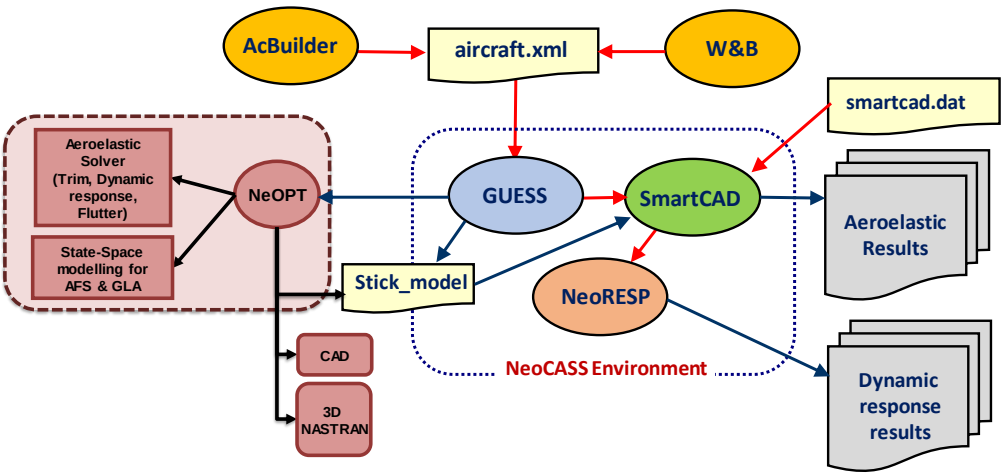


Figure 1: NeoCASS Overview

To accelerate the automatic sizing process, NeoCASS incorporates simplifications, limiting structural models to isotropic or quasi-isotropic materials (e.g. composites are treated as "black aluminum"). This restricts the ability to capture fiber-induced structural coupling. Additionally, for unconventional configurations like Blended-Wing Bodies (BWB) or Strut-Braced Wings (SBW), a beam-only model cannot fully capture structural behavior, necessitating 3D modeling capabilities. To address these issues, NeoCASS underwent a major revision, introducing the NeOPT optimization module and enhancing solvers, particularly for composite wing structures. The integration of NeoCASS and

NeOPT creates a unified design suite, where wingbox sizing is conducted via a semi-analytical meta-model representation. This enables comprehensive structural and aeroelastic assessments, including failure and buckling analyses, flutter, control effectiveness, and gust response - with and without active control technologies (e.g. maneuver and gust load alleviation systems) [20, 21].
Figure 1 provides a schematic overview of NeoCASS, which consists of three core modules:

1. GUESS: performs initial structural sizing based on limited geometric inputs and user-selected structural design philosophy. The sizing process adheres to certification rules or user-defined maneuver sets.

2. SmartCAD: the aeroelastic kernel, comprising solvers for aeroelastic trim, flutter, and control reversal analyses. Both GUESS and SmartCAD employ frozen maneuvers to solve trim problems algebraically, without time-dependent load evolution.

3. NeoRESP: solves generic response problems, including frequency-domain gust response analysis and state-space model generation for active control system design, such as maneuver and gust load alleviation systems.

The code includes two additional modules: AcBuilder, for parametric aircraft geometry definition; Weight & Balance (W&B), for initial mass estimation via statistical formulas. The latter provides a preliminary mass distribution, which serves as an input for structural sizing in GUESS. A simplified but complete finite element aeroelastic model, supported by dedicated structural and aerodynamic solvers, is then generated. The structural model relies on a nonlinear three-node finite volume beam approach, while aerodynamic computations leverage Vortex Lattice and Doublet Lattice methods for steady and unsteady conditions, respectively.

In summary, GUESS enables a physics-based, fully stressed aircraft design with key simplifications to expedite the process and facilitate automatic aeroelastic model generation. These include assuming a rectangular and symmetric wingbox, limiting materials to isotropic types, conducting rigid trim-based load envelope analyses, and omitting direct load sensitivity calculations with respect to design variables. These challenges are then addressed by the new NeOPT module, which refines the wingbox sized by GUESS by means of optimization. A Meta-Model approach is adopted, with a semi-analytical representation of the wingbox that organizes its 3D geometrical and physical properties into a unified framework. This allows different analysis models to be managed using a single set of information, enabling seamless result transfer between domains. The wingbox is modeled with higher fidelity without increasing FEM model complexity. Additionally, the sizing process leverages all NeoCASS solvers, incorporating dynamic load conditions such as gusts and control inputs, as well as flutter analysis to ensure the aero-servo-elastic stability of the aircraft.

In this paper, NeoCASS and NeOPT have been coupled in an integrated MDO framework tailored for flutter optimization of a typical long-range transport aircraft.

3 FLUTTER OPTIMIZATION AND CERTIFICATION RULES

The FAA Part 25 (Title 14, Code of Federal Regulations Part 25) [22] and EASA CS-25 (Certification Specifications for Large Aeroplanes) [18] contain specific requirements

related to stability, control, and flutter for commercial transport aircrafts. From the stability and flutter perspective, the key certification requirements include the following:

- the aircraft must be free from aeroelastic instability (flutter, divergence, control reversal) for all speeds up to 1.15 times the dive speed (V_D).
- Flutter clearance must be demonstrated considering structural damping, aerodynamic forces, and mass properties.
- Damping requirements: all aeroelastic modes must be sufficiently damped across the entire flight envelope.
- Mass balancing of control surfaces must prevent flutter.
- Ground Vibration Testing (GVT) and analytical validation are required to confirm that the structural and aerodynamic models accurately predict flutter behavior.

The same requirements can be coded in a graphical way in what is called Method 1, where no structural damping is included in the model before computing flutter analysis:

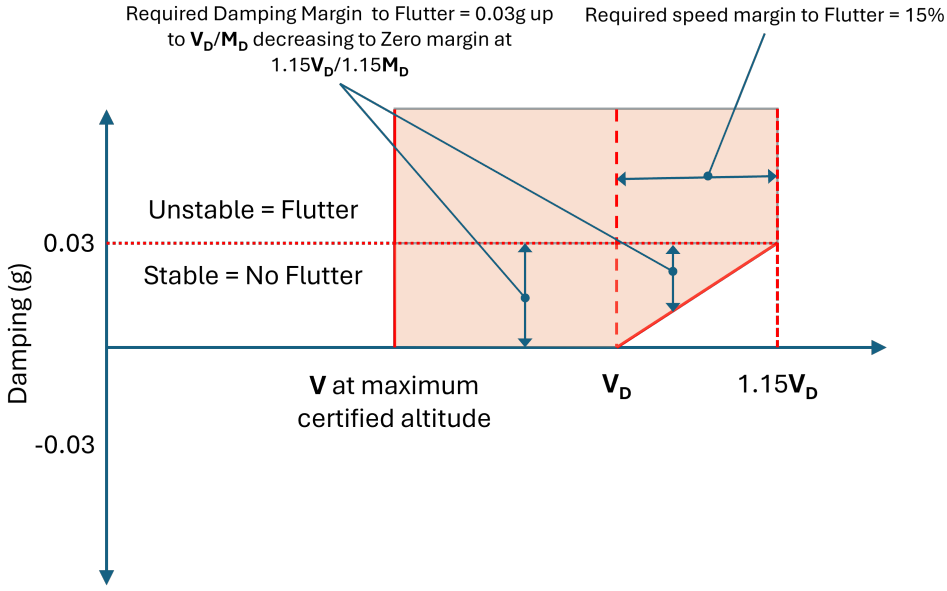


Figure 2: Method 1 - flutter requirements in shaded region

By analyzing Figure 2, we can determine that the aircraft exhibits natural or passive stability and complies with CS-25 regulations if a damping margin of g=0.03 is guaranteed up to V_D, and linearly reduced up to 1.15 V_D. Flutter stability is hence guaranteed if the tracked modal dampings do not cross g=0 line up to V_D and g=0.03 line up to 1.15 V_D. This is valid in case no structural damping is included in the model.

The literature presents various examples of aeroelastic optimization with flutter constraints, implemented using different minimization algorithms, such as gradient-based and stochastic methods, but all adhering to the same regulatory requirements in different ways. Some studies apply an explicit constraint on velocity or minimum damping values, while others impose broader constraints on entire rejection curves to prevent potential hump modes from entering the rejection zone [23, 24].

It is important to note that in a more realistic nonlinear approach, flutter constraints can vary with flight parameters, as they depend on the specific conditions at which the problem is linearized. Traditionally, to achieve a balance between computational cost and accurate problem representation, flutter constraints are often imposed in frequency domain. This is done by coupling medium-complexity finite element structural models with unsteady aerodynamic methods, typically analyzed using the Doublet-Lattice Method (DLM).

The approach of designing a naturally or passively flutter-free aircraft relies primarily on optimizing the distribution of stiffness and mass along the wing, demonstrating analytically - then experimentally - that the imposed regulatory conditions are satisfied.

Interest in AFS systems, which ensure the necessary damping through active control technologies, has long captured the imagination of researchers. Since the late 1970s, various solutions have been proposed, transitioning from purely passive optimization to hybrid, or aeroservoelastic, optimization. Nontheless, this approach has never been a viable design option due to regulatory limitations, as well as concerns regarding safety and robustness [12].

As anticipated, the situation has recently evolved due to two key factors: the drive for more efficient, lighter aircrafts and the certification of two aircrafts that seemingly contradict the traditional principle of flutter-freedom up to 1.15 V_D, without relying on dedicated active control systems. Specifically, the cases at hand are Boeing 747-800 [13] and Boeing 787-10 [14], which have been certified by the FAA with explicit reference to the use of onboard active control systems to ensure sufficient damping margins. Moreover, for the former aircraft, to account for potential failures of Active Modal Suppression system (called F0VMS), the rejection curve from Figure 2 was modified as reported in Figure 3. The initial requirement of minimum damping margin is lowered to 0.015 in case of inoperative control systems.

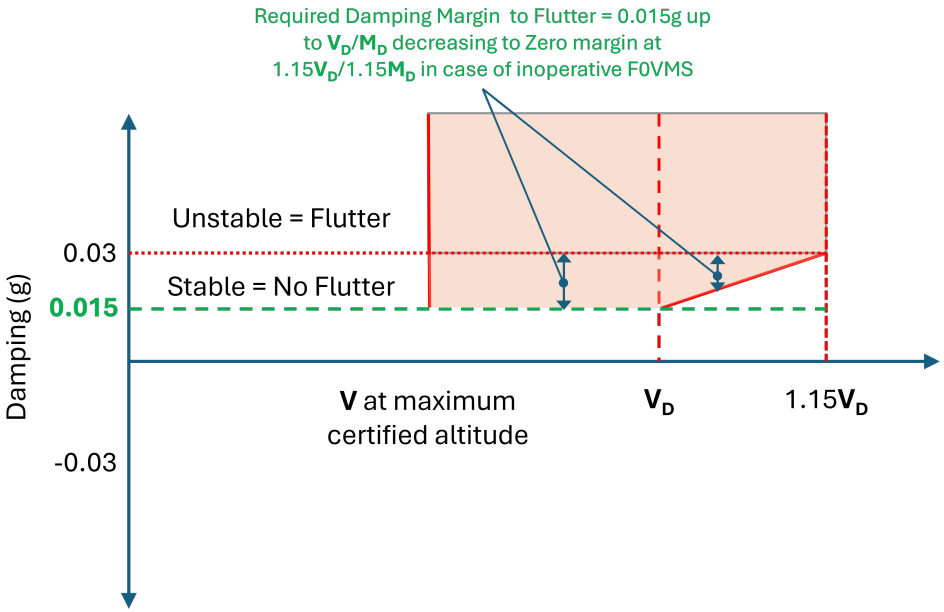


Figure 3: Method 1 for inoperative F0VMS - flutter requirements in shaded region

For the two examples, flutter-related issues emerged during the certification phase, when design modifications were no longer feasible. As often occurs, such phenomena are case-dependent and cannot be generalized. However, these instances have reignited interest in the feasibility and benefits of certifying aircrafts that meet flutter requirements using active control systems.

From a design perspective, the key question is how much structural weight savings can be achieved by leveraging existing onboard control systems to ensure compliance with flutter certification requirements while maintaining adequate safety margins. This remains an open debate, which has led to extensive specialized literature, dedicated research projects, and focused discussions at industry conferences.

The authors want to stress that this article does not aim to provide definitive answers to these complex questions, but rather contribute to the discussion with a specific case study. In particular, the work does not focus on the development of an active control system to guarantee stability, but instead seeks to quantify the potential structural weight savings associated with such an approach.

4 MODEL DESCRIPTION

The aircraft considered in the present work has been obtained in [21] starting from the publicly available data of the Airbus A330-300, whose three views are reported in Figure 4.

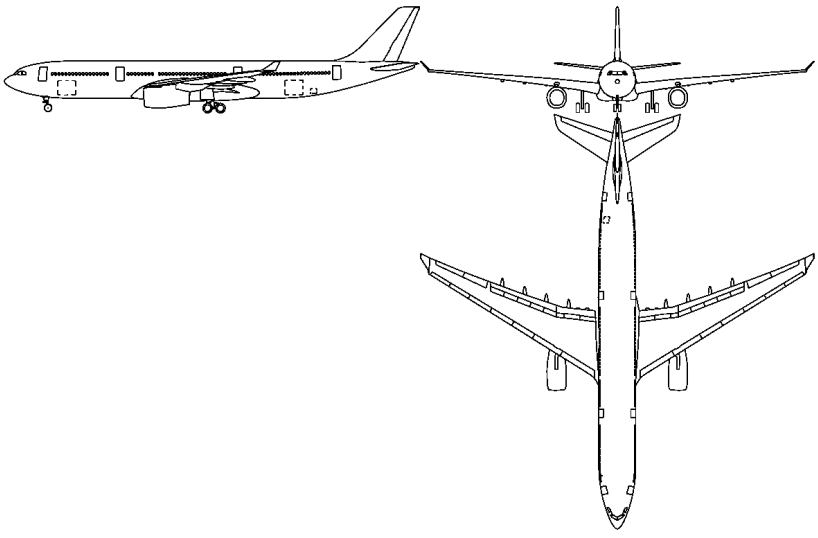


Figure 4: Views of A330-300 [21]

Since some changes have been included to produce the studied model, to avoid any confusion between the reference aircraft (A330-300) and the final one, the latter has been renamed Twin-Aisle Long-Haul (TALH) transport aircraft - from now on, simply TALH. In particular, the geometrical updates considered while creating the TALH model

in AcBuilder concern wing aspect ratio, which has been increased to create a HARW. Table 1 summarizes the main requirements for such aircraft.

Table 1: TALH main requirements

Number of passengers	300 [-]
Wing Area	366.7 [m^2]
Wing Span	67 [m]
Wing Aspect Ratio	12.24 [-]
Max range	10200 [km]
Max cruise altitude	13100 [m]
Max Operating Mach number	0.87
Max Operating Speed (EAS)	175 [m/s]
Cruise speed at FL 380 (TAS)	244.91 [m/s]
Cruise speed (EAS)	127.5 [m/s]
Max lift coefficient in take-off configuration	2.1 [-]
Max lift coefficient in landing configuration	2.5 [-]
Max lift coefficient in clean wing configuration	1.75 [-]
Clean lift curve slope	5.7 [1/rad]

Recalling that GUESS allows to consider only isotropic materials, the entire model was firstly made of aluminum alloy, with the properties reported in Table 2.

Table 2: TALH material properties for GUESS sizing

Material	$E_{11,22}$ [GPa]	G_{12} [GPa]	ν_{12} [-]	ρ [kg/m^3]	σ_{max} [MPa]
AL7075-T6	72	27.69	0.3	2800	430

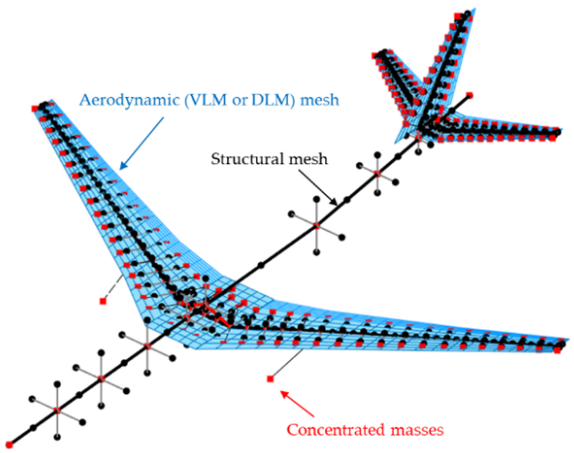


Figure 5: TALH model in NeoCASS

Starting from the requirements defined in Table 1, the entire set of manoeuvres compliant with the CS-25 regulations have been derived (Table 3) and used to generate the

sized model. Figure 5 shows NeoCASS stick model, in conjunction with aerodynamic panels in correspondence of the lifting surfaces. Additionally, also the concentrated masses representing the engines are reported.

Table 3: Database of sizing manoeuvers compliant with CS25 regulations

ID	Type	N	z [m]	Mach
1-4	Pull-up	2.5	0/8061/ 0/12250	0.514/0.87/ 0.375/0.87
5-8	Push-down	-1.0	0/8061/ 0/12250	0.514/0.87/ 0.375/0.87
9	Flap max down	2.0	0	0.329
10	Static Gust for landing configuration	1.0	0	0.329
11-16	Aft movement of pitch control at VA, VC, VD URDD5 (positive and negative)	1.0	0	0.357/0.357/ 0.514
17/20/ 23/26	Sideslipe max rudder (positive and negative)	1.0	0	0.357/0.226
18/21/ 24/27	Sideslipe overswing (positive and negative)	1.0	0	0.357/0.226
19/22/ 25/28	Sideslipe maneuver 15° (positive and negative)	1.0	0	0.357/0.226
29/31	Aileron abrupt maximum deflection at VA $=\sqrt{N_{max}}$*VS (positive and negative)	1.0	0	0.357
30/32	Aileron abrupt deflection + 2/3 max load factor (positive and negative)	1.7	0	0.357

Once obtained the sized model with NeoCASS, it is possible to proceed with NeOPT. In particular, NeOPT just focuses on the wing, meaning that the optimization problem will just update the structural thicknesses of that component.

5 MDO PROBLEM DEFINITION

The optimization problem solved in NeOPT can be represented as follows:

$$\min_{x \in S} f(x), S = \{x \in \mathbb{R}^{\kappa}, g(x) < 0 \cap x_{lb} \leq x \leq x_{ub}\}$$

which reads "find the design variables x belonging to the space S, that minimize the objective function $f(x)$, such that the constraints $g(x)$ are satisfied and the design variables x stay between the lower and upper bounds".

Let's define the design variables, the objective function and the constraints for the problem at hand.

5.1 Design variables

It has been anticipated that through NeOPT one can overcome the simplifications of having a symmetric isotropic wing cross-section. This is possible thanks to a 2D cross-sectional solver that provides the 6x6 beam stiffness matrix for any empty section composed of isotropic or orthotropic material. In particular, depending on the cross-section description chosen by the user, the matrices will be fully-populated, accounting for different couplings. As reported in Figure 6, the section is represented, accordingly to the semimonocoque theory, as the combination of stringers, spar caps, skins and webs (spars).

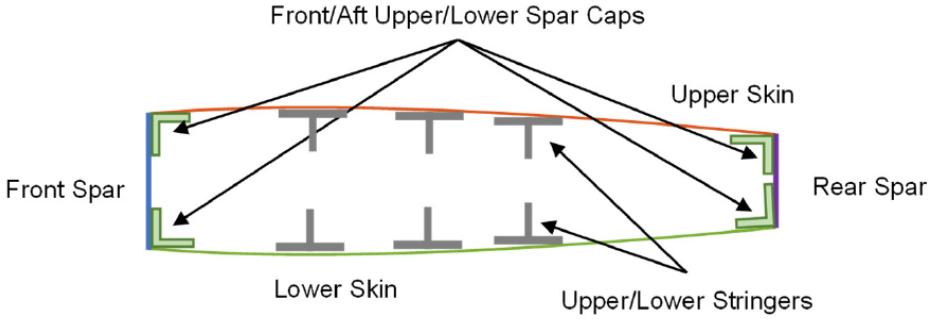


Figure 6: NeOPT cross-section description [21]

The design variables of the optimization problem, i.e. the unknowns that are determined once the problem reaches convergence, are the thicknesses of these elements. In the most general case, they are free to change independently, leading to 10 design variables per section: upper and lower stringer thickness (1-2), upper and lower front and rear spar caps thickness (3-4-5-6), upper and lower skin thickness (7-8), front and rear web thickness (9-10).

It is possible to reduce the number of unknowns by expliciting a sectional variable linking. In the case of the present study, for instance, just 7 design variables have been chosen to describe each wing section, because all the spar caps have been considered to have the same thickness. Moreover, for what concerns the material, aluminum has been substituted by composite with the stacking sequence reported in Figure 7 and the Carbon Fiber Uni-Directional (CFUD) properties reported in Table 4.

Table 4: TALH wing Carbon Fiber Uni-Directional properties

Material	E_{11} [GPa]	E_{22} [GPa]	G_{12} [GPa]	ν_{12} [-]	ρ [kg/m^3]	σ_{max} [MPa]
CFUD	135	10	5	0.3	1500	1500

The second variable linking that is computed is spanwise. Indeed, the code requires the user to provide informations on the wing spanwise description too. In this way, it is possible to divide the component in patches, assigning to each finite element that falls inside a patch the very same section description. As one can see from Figure 8, in this case 5 patches have been selected, leading to a total of 7x5=35 design variables.

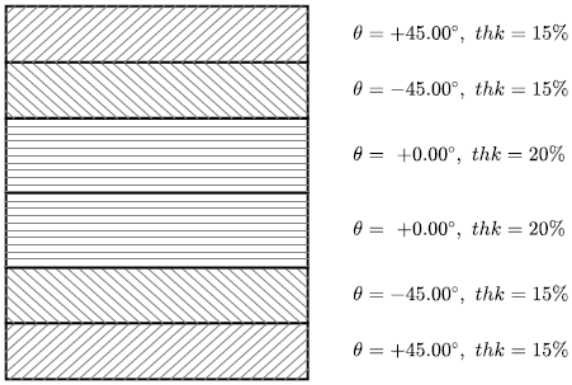


Figure 7: TALH wing composite stacking sequence

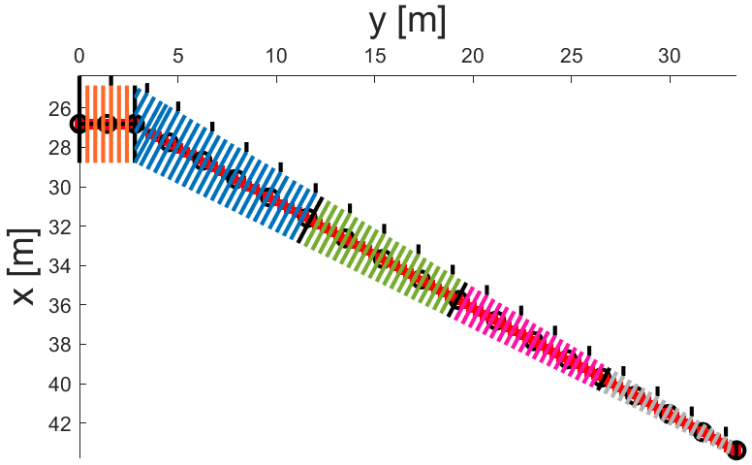


Figure 8: TALH wing patches

5.2 Objective function and constraints

As reported at the beginning of this section, the optimization problem solved in NeOPT is a minimization. The objective function is the weight of the semi-wingbox. The code, indeed, allows to just work on half wing, thanks to model symmetry. In the end, the obtained converged value can be compared either to the starting one, or to other optimization solutions, as it will be the present case.

For what concerns the constraints, two categories have been considered:

- strength requirements, i.e. no failure nor buckling;
- no flutter inside the flight envelope.

To satisfy the former set, aeroelastic trim analyses are computed in the flight conditions used to previously size the model, i.e. the ones reported in Table 3. The internal loads,

and hence the stresses, are evaluated in each section component, multiplied by a safety factor of 1.5 and compared to the critical values. Tsai-Hill criterion is used for failure considerations. Once the iterative process is completed, if no local optimum is found, the code prints some informations concerning maxima constraint violations in terms of manoeuvre ID, failure or buckling, and the location at which it occurs, such that the user knows which are the weakest spots on the wing and the associated most critical manoeuvres.

Flutter constraint has been included computing a linear flutter analysis, by means of continuation method, at the corner point (z=8061 [m], Mach=0.87), requiring that for each tracked mode the maximum allowable damping was g<0, up to V_D, complying with the requirements represented in Figure 2. Actually, Method 1 requires to guarantee the absence of flutter up to a velocity of 1.15 V_D. However, this condition translates into a Mach number very close to 1, which poses a limit for DLM - the method used to solve aerodynamics. It has thus been chosen to not rely on aerodynamic loads that were too close to Mach=1, limiting the flutter requirements just to dive condition, without any safety margin.

5.3 Optimization problems

Since the aim of the present work is to estimate what happens in terms of mass when either flutter is explicitly excluded from the flight envelope through a passive approach, or when the structure is sized without any requirement on flutter, knowing that it can be treated *a posteriori* by means of active control systems, two different optimization problems have been solved.

- Optimization 1: failure and buckling constraints. If flutter would occur, it would be taken care of by active control systems.
- Optimization 2: failure, buckling and flutter constraints included in the problem.

The two objective functions can be compared to derive some conclusions on how treating flutter by means of passive or active approach may affect structural sizing.

6 RESULTS

6.1 Optimization 1: failure and buckling constraints

Figure 9 shows the resulting design variables for the first optimization, where only strength constraints have been considered. The initial solution has been generated adapting the semi-wingbox preliminary sizing, produced by GUESS, to the new wingbox description of NeOPT, updating the material and including the variable linking.

Looking at the web thicknesses, one can notice that the initial solution produced an unrealistic value in correspondence of the fourth patch. This can be easily solved by the optimization process, which ensures not only the satisfaction of the constraints, but also that the design variables fall between lower and upper bounds defined by the user.

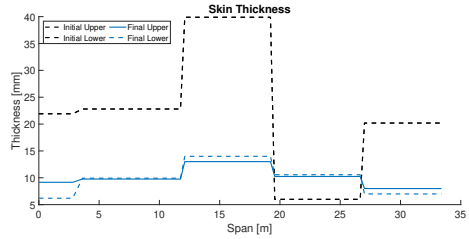
In terms of masses, a comparison between the initial and the optimized solutions is given in Table 5.

Figure 10 shows a comparison between the initial and the optimized models root loci, computed at z=8061 [m] and V_D, to keep Mach<1. As one can notice, the initial model does not present flutter, while this occurs for the optimized one. This is not a surprise,

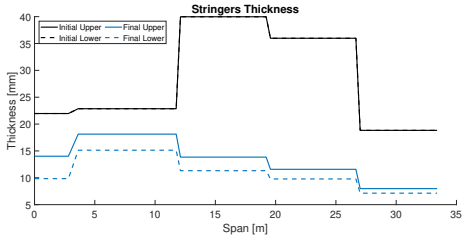
considering that the optimization problem has been defined without taking care of flutter, and hence accepting that it could have fallen inside the envelope.

Table 5: TALH semi-wingbox results - Optimization 1

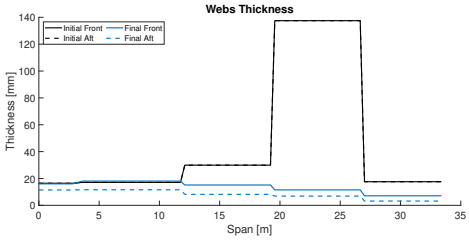
	Initial	Final	Δ
Semi-wingbox Mass	24795.31 [kg]	6417.27 [kg]	-74%



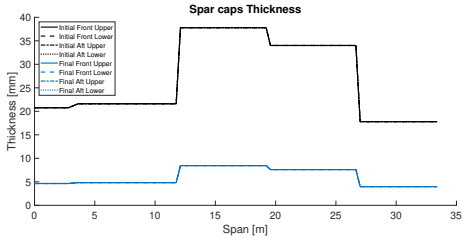
(a) Skins



(b) Stringers

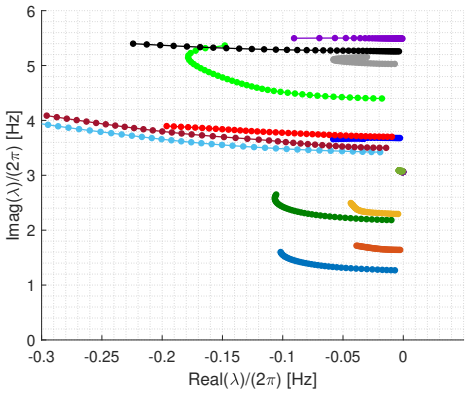


(c) Webs

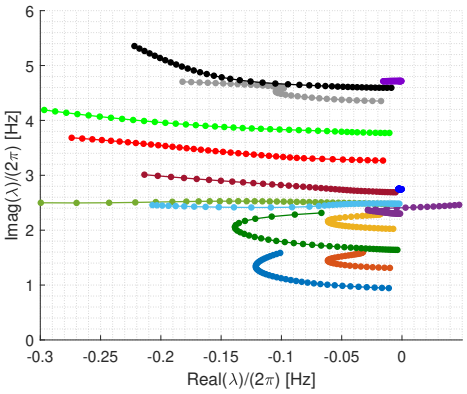


(d) Sparcaps

Figure 9: TALH thicknesses - Optimization 1



(a) Initial

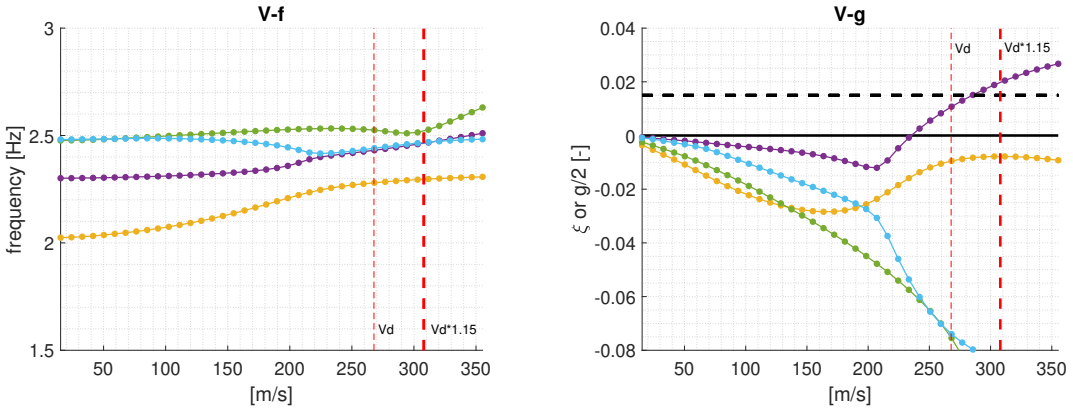


(b) Final

Figure 10: TALH root loci - Optimization 1

To characterize the mechanism, V-f and V-g plots are reported in Figure 11 for the optimized model, isolating the modes that are closer to flutter frequency. The coalescing modes are the fifth and the sixth deformable ones. In both subfigures the lines correspondent to V_D and 1.15 V_D have been reported highlighting the limits that must be

considered from the regulatory point of view. Moreover, in the subfigure of damping also an horizontal line correspondent to g<0.03 (or ξ <0.015) has been added to delimit the region of unstable flutter accordingly to Figure 2. In this optimization no active control system was implemented yet to counteract flutter.



(a) Frequency (b) Damping

Figure 11: TALH V-f and V-g - Optimization 1

To further investigate the phenomenon, one can look at the resulting flutter mode, which is shown in Figure 12. It is given by an interaction between inplane and out-of-plane wing symmetric bending and wing symmetric torsion. However, looking at the tail plane, it seems that a participation of the horizontal tail and the fuselage are also present. Since NeOPT acts only on the wing, if an AFS system would be designed and included inside the aero-servo-elastic optimization, it could occur that flutter won't be canceled, due to the uncontrollable participation of fuselage and tail. This issue has not been addressed yet and will be treated in future works.

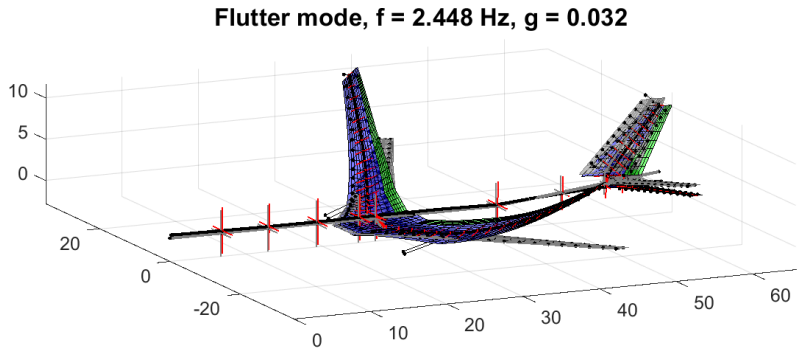
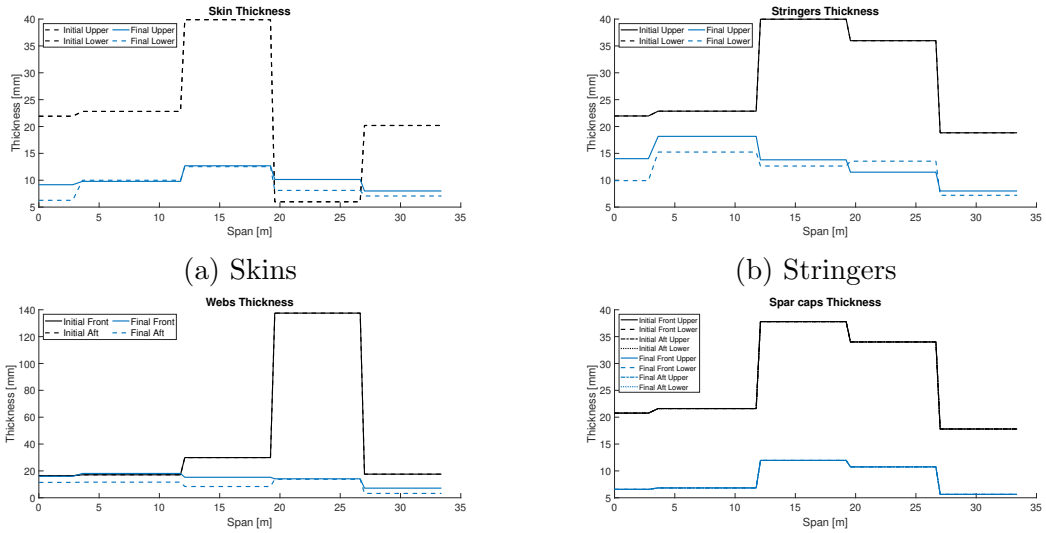


Figure 12: TALH flutter mode - Optimization 1

6.2 Optimization 2: failure, buckling and flutter constraints

For the second problem, also the flutter constraint has been included inside the optimization, leading to the results shown in Figure 13.



(a) Skins (b) Stringers

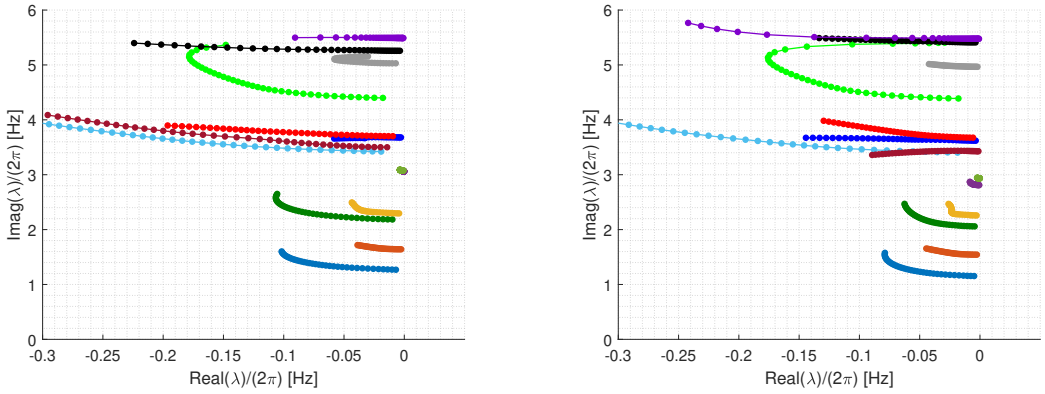
(c) Webs (d) Sparcaps

Figure 13: TALH thicknesses - Optimization 2

Table 6 contains the results in terms of sizing. As one can notice, the final mass, i.e. the one of the optimized model, is higher in this case with respect to what has been obtained from the first optimization.

Table 6: TALH semi-wingbox results - Optimization 2

	Initial	Final	Δ
Semi-wingbox Mass	24795.31 [kg]	6676.83 [kg]	-73%

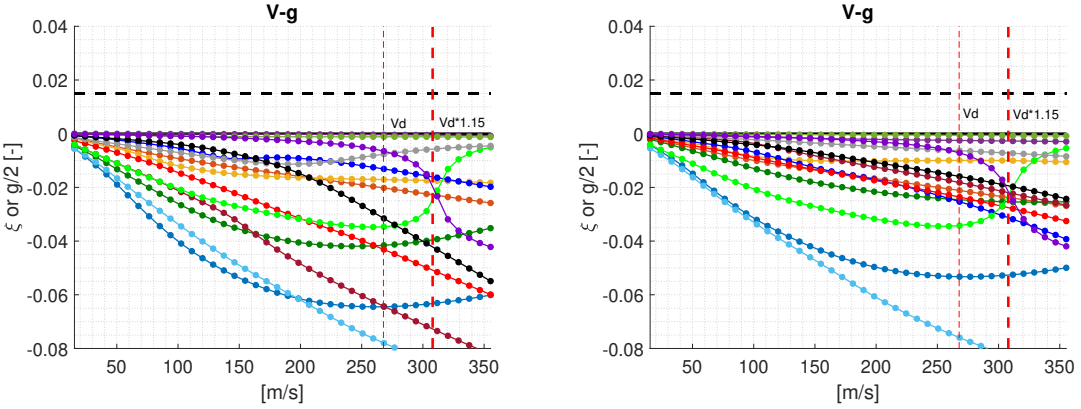


(a) Initial (b) Final

Figure 14: TALH root loci - Optimization 2

For what concerns flutter analysis, simply by looking at the root loci, computed also in this case at z=8061 [m] and V_D, one can notice that no flutter occurs, as shown in Figure 14.

For sake of completeness, also the V-g plots for both the initial and the optimized model are reported in Figure 15, including all the deformable modes to be tracked, to show that none of them crosses the flutter constraint line.

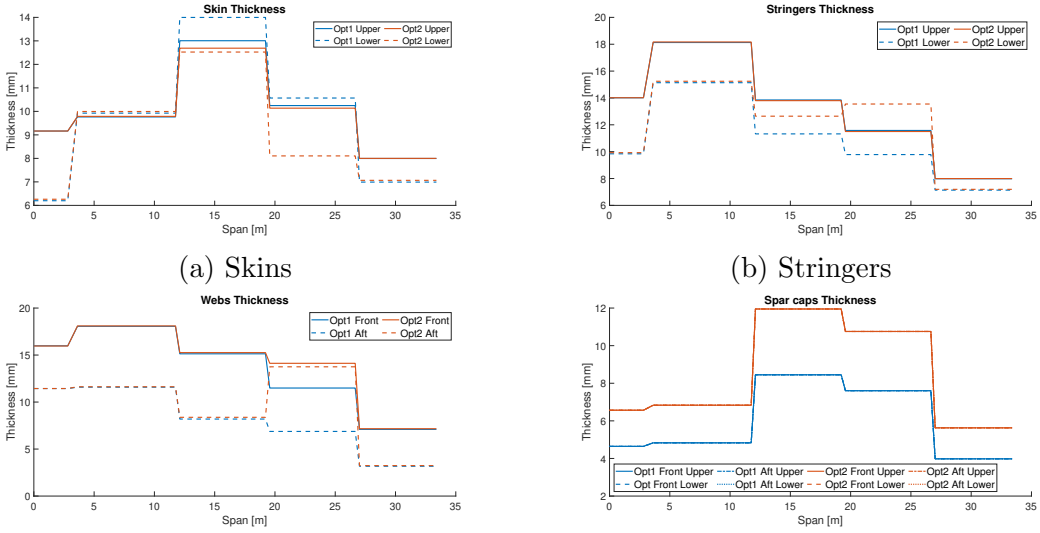


(a) Initial (b) Final

Figure 15: TALH V-g plot - Optimization 2

6.3 Sizing results comparison

A comparison of sizing results in terms of thicknesses and semi-wingbox mass is reported in Figure 16 and Table 7, respectively. Not including flutter constraint inside the optimization process leads to a mass gain of around 200 [kg] - 3.9% - per half wing.



(a) Skins (b) Stringers

(c) Webs (d) Sparcaps

Figure 16: TALH results comparison

Table 7: TALH semi-wingbox results comparison

	Optimization 1	Optimization 2	Δ
Semi-wingbox Mass	6417.27 [kg]	6676.83 [kg]	+3.88%

7 CONCLUSIONS

Flutter is one of the most dangerous phenomena that an aircraft can experience. For safety reasons, aircrafts are required to prove flutter-freedom throughout the entire flight envelope, with a due safety margin. Despite the effectiveness of using Active Flutter Suppression systems to manage flutter, as demonstrated by the vast literature, this option has not been supported by certification rules. However, two recent cases highlight certification possibilities with active control systems, sparking new discussions on studying and enhancing such systems.

From a structural standpoint, this means taking advantage of the many active control systems already present in future aircraft to control flutter and simultaneously reduce structural weight, aligning with modern efforts to reduce emissions. While flutter is case-dependent, making it difficult to derive a general rule, this paper attempts to quantify the potential structural weight advantage in the analyzed case if the flutter constraint is considered from the beginning of the design process, allowing active control systems to handle flutter suppression. Two optimization problems were solved using the coupling of NeoCASS and NeOPT, explicitly including or excluding the flutter constraint. The results indicate a potential reduction in wingbox mass of around 200 [kg] per semi-wingbox. This weight advantage needs to be weighed against the mass of new control systems required to implement the AFS functionality. Additionally, issues related to reliability and safety aspects must be addressed to complete the certification process.

ACKNOWLEDGEMENTS

The U-HARWARD Project has received funding from the Clean Sky 2 Joint Undertaking, under the European's Union Horizon 2020 research and innovation Program under Grant Agreement number: 886552-H2020-CS2-CFP10-2019-01. The content of this document reflects only the author's view. The European Commission and Clean Sky 2 Joint Undertaking (CS2JU) are not responsible for any use that may be made of the information it contains.

REFERENCES

[1] R. L. Bisplinghoff, H. Ashley, and R. L. Halfman. *Aeroelasticity*. Dover Publications, Mineola, NY, 2 edition, 1996.

[2] M. M. Opgenoord, M. Drela, and K. M. Willcox. Influence of Transonic Flutter on the Conceptual Design of Next-Generation Transport Aircraft. In *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Kissimmee, FL, 2018. doi:10.2514/6.0948.

[3] E. Livne. Future of Airplane Aeroelasticity. *Journal of Aircraft*, 40(6):1066–1092, 2003. doi:10.2514/2.7218.

[4] E. Garrigues. A Review of Industrial Aeroelasticity Practices at Dassault Aviation for Military Aircraft and Business Jets. *Aerospace Lab*, 14:1–34, 2018. doi:10.12762/2018.AL14-09.

[5] J. R. Martins and A. B. Lambe. Multidisciplinary Design Optimization: A Survey of Architectures. *AIAA Journal*, 51(9):2049–2075, 2013. doi:10.2514/1.J051895.

[6] Z. Lyu, G. K. W. Kenway, and J. Martins. Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark. *AIAA Journal*, 53(4): 968–985, 2015. doi:10.2514/1.J053318.

[7] G. Dimitriadis. *Introduction to Nonlinear Aeroelasticity*. John Wiley and Sons, 2017. doi:10.1002/9781118756478.

[8] B. Stanford and P. Beran. Direct Flutter and Limit Cycle Computations of Highly Flexible Wings for Efficient Analysis and Optimization. *Journal of Fluids and Structures*, 36:111–123, 2013. doi:10.1016/j.jfluidstructs.2012.08.008.

[9] G. J. Kennedy, G. K. W. Kenway, and J. Martins. Towards Gradient-Based Design Optimization of Flexible Transport Aircraft with Flutter Constraints. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, 2014. doi:10.2514/6.2726.

[10] E. Jonsson, C. Rico, C. Lupp, C. Cesnik, and R. Martins, Joaquim. Flutter and Post-Flutter Constraints in Aircraft Design Optimization. *Progress in Aerospace Sciences*, 109:100537, 2019.

[11] E. Livne. Aircraft Active Flutter Suppression: State of the Art and Technology Maturation Needs. *Journal of Aircraft*, 55(1):410–452, 2017. doi:10.2514/1.C034442.

[12] C. Cesnik, R. Palacios, and E. Reichenbach. Reexamined Structural Design Procedures for Very Flexible Aircraft. *Journal of Aircraft*, 51(5):1580–1591, 2014. doi:10.2514/1.C032464.

[13] Proposed Special Condition for Installation of Flutter Suppression System Applicable to Boeing B 747-8F/-8.

[14] Special Condition on Flaps-Up Vertical Modal-Suppression System Applicable to Boeing 787-10. URL `https://www.easa.europa.eu/en/downloads/23783/en`.

[15] L. Cavagna, S. Ricci, and L. Travaglini. NeoCASS: An Integrated Tool for Structural Sizing, Aeroelastic Analysis and MDO at Conceptual Design level, 2011. ISSN 03760421.

[16] L. Cavagna, S. Ricci, and L. Riccobene. Structural Sizing, Aeroelastic Analysis, and Optimization in Aircraft Conceptual Design. *Journal of Aircraft*, 48:1840–1855, 2011. ISSN 0021-8669. doi:10.2514/1.C031072.

[17] L. Cavagna, S. Ricci, and L. Travaglini. Aeroelastic Analysis and Optimization at Conceptual Design Level Using NeoCASS Suite. In *Collection of Technical Papers - AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2011. doi:10.2514/6.2011-2079.

[18] EASA. Certification Specifications for Large Aeroplanes CS25, 2024. URL https://www.easa.europa.eu/en/document-library/certification-specifications/group/cs-25-large-aeroplanes.

[19] EASA. CS-23 Normal, Utility, Aerobatic and Commuter Aeroplanes, 2025. URL https://www.easa.europa.eu/en/document-library/certification-specifications/group/cs-25-large-aeroplanes.

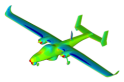
[20] F. Toffol and S. Ricci. A Meta-Model for Composite Wingbox Sizing in Aircraft Conceptual Design. *Composite Structures*, 306, 2023. ISSN 02638223. doi:10.1016/j.compstruct.2022.116557.

[21] F. Toffol and S. Ricci. Preliminary Aero-Elastic Optimization of a Twin-Aisle Long-Haul Aircraft with Increased Aspect Ratio. *Aerospace*, 10(4), 2023. ISSN 22264310. doi:10.3390/aerospace10040374.

[22] Federal Aviation Administration. AC 25.629-1B - Aeroelastic Stability Substantiation of Transport Category Airplanes. 10 2014. URL https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_25_629-1B.pdf.

[23] A. S. Thelen, B. Stanford, and K. Jacobson. Flutter and stochastic gust constraint studies of the undeflected crm using mphys. In *AIAA Scitech 2024 Forum*, page 2774, 2024.

[24] A. Lunghitano, F. Afonso, and A. Suleman. A study on the surrogate-based optimization of flexible wings considering a flutter constraint. *Applied Sciences*, 14(6): 2384, 2024.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

A COMPARISON OF MDO BI-LEVEL FORMULATION VARIANTS WITH MDF ON A HIGH-FIDELITY AEROELASTIC WING TEST CASE

Gilberto Ruiz Jimenez[1*], Nicolas Roussouly[1], François Gallard[1], Anne Gazaix[1,2], Joël Brezillon[2], Kittinan Thanissaranon[2], Vida Brück[2] and Joaquim R.R.A. Martins[3]

1: IRT Saint-Exupéry
{gilberto.ruiz-j,nicolas.roussouly,francois.gallard,anne.gazaix}@irt-saintexupery.com

2: Airbus S.A.S.
{joel.brezillon,kittinan.thanissaranon,vida.brueck, anne.gazaix}@airbus.com

3: University of Michigan
jrram@umich.edu

Abstract. *In multidisciplinary design optimization (MDO), the selection of the formulation (or architecture) constitutes a critical aspect for industrial applications. While the MDF formulation has been successfully applied to various cases, its usage in a large-scale industrial context using segmented numerical software is limited. Consequently, distributed formulations are actively investigated by industrial actors as better adapted candidates. In particular, several variants of bi-level formulations have been recently developed to address use cases where the number of shared design variables is moderate, the number of coupling variables is large, and when the Jacobians of the disciplines are partially available. The bi-level IRT formulation was originally developed to take advantage of gradient-based disciplinary optimizations within an MDO process. To this end, the disciplinary optimizations are embedded within a system optimization that handles the shared variables using a gradient-free optimizer. Two multidisciplinary analysis (MDA) are utilized before and after the sub-optimizations to satisfy the high-dimensional couplings. However, the separation of the disciplinary optimizations can result in convergence degradation. The bi-level block coordinate descent (BCD) formulation has been proposed as an enhancement, offering more robust mathematical foundations, to overcome this issue by using the BCD algorithm, which iterates on the disciplinary optimizations. This paper presents a comparison of the bi-level IRT and the bi-level BCD formulations with MDF as a reference, on an aeroelastic aircraft wing model, using RANS CFD and computational solid mechanics (CSM) solvers. The MACH framework is selected because it has a state of the art MDF implementation with the coupled adjoint. The comparison highlights the merits and limitations of the bi-level formulations with respect to performance and accuracy. Encouraging results were obtained with the IRT bi-level formulation although a lack of robustness was observed with respect to the initial point of the optimization, the BCD enhancement handles this discrepancy at the expense of computational cost.*

Keywords: Multidisciplinary design optimization, bi-level, multidisciplinary feasible, block coordinate descent, aerostructural optimization

1 INTRODUCTION.

In the domain of multidisciplinary design optimization (MDO), the selection of the formulation (or architecture) constitutes a critical aspect for industrial applications. While the multidisciplinary feasible (MDF) formulation [1] has been successfully applied to various cases, its usage in a large-scale industrial context using segmented numerical software is limited. Consequently, distributed formulations are actively investigated by industrial actors as better adapted candidates.

Several variants of bi-level formulations have been recently developed by IRT Saint Exupéry in the IRT projects MDA-MDO [2] and R-Evol [3], to address use cases where the number of shared design variables is moderate, the number of coupling variables is large, and when the Jacobians of the disciplines are partially available. The bi-level IRT formulation was originally developed to take advantage of gradient-based disciplinary optimizations within an MDO process. To this end, the disciplinary optimizations are embedded within a system optimization that handles the shared variables using a gradient-free optimizer. Two MDAs are utilized before and after the sub-optimizations to satisfy the high-dimensional couplings. However, the separation of the disciplinary optimizations can result convergence degradation. The bi-level block coordinate descent (BCD) formulation has been proposed as an enhancement, offering more robust mathematical foundations, to overcome this issue by using the BCD algorithm which iterates on the disciplinary optimizations.

In the following sections, these two bi-level formulations will be applied and compared to solve an optimization problem based on an aircraft wing model known as the MACH tutorial model. Section 2 introduces the context and the objectives of this analysis and Section 3 describes the optimization problem to be solved. Section 4 gives the details of the IRT bi-level formulation applied to the wing model and Section 5 presents several results obtained with the MACH framework tools. Finally, Section 6 presents the bi-level BCD approach with results on the wing model use case.

This work has been carried out in the framework of a tripartite IRT Saint Exupéry-Airbus-University of Michigan collaboration.

2 CONTEXT AND OBJECTIVES.

The objective of the current work is to demonstrate the ability of two bi-level formulations, namely IRT bi-level and bi-level BCD, with high-fidelity models to support early design trades on an aeroelastic wing model. We also aim to assess them with respect to the reference MDF formulation. In particular, we will focus on the aero-structure coupling by taking into account:

- Aerodynamics for performance: drag computation in cruise flight.
- Aerodynamics for loads: computation of loads applied to the structure through maneuver flight conditions.
- Structural sizing optimization: computation of the mass with failure criteria due to the loads.

Our interest is to find the optimal design through an automated MDO process, capturing the trade-offs between the aerodynamic and the structure effects, especially the fact that:

- The change of the aerodynamic shape affects the drag, as well as the structural mass and failure constraints through the repartition of the loads.
- The change of the structural stiffness affects both the failure constraints and the drag due to the deformation of the wing.

The well-known MDF formulation efficiently finds the aforementioned trade-offs, but it requires the computation of the coupled Jacobian for the entire process. While several academic developments from last decades have shown a real interest for this approach, it is generally not affordable in an industrial context since disciplinary tools that are used do not allow exposing all the required terms. Another issue with the MDF formulation is that it is only relevant on a limited number of constraints, since the cost of the coupled gradient computation may be unaffordable. While the constraint aggregation strategy reduces the number of constraints, it introduces some approximations that are not allowable to some industrial applications.

The advantage of the bi-level formulation in industrial settings is that it can benefit from specialized disciplinary optimizers at the sub-scenario level, while using the system level to optimize the shared variables. The downside is that it is not always easy to properly split a highly coupled problem in a way that preserves the ability of the system to control the main trade-offs between the sub-scenarios. In the following, we will show different bi-level formulations and compare the results that they yield with the reference computed with MDF.

3 THE OPTIMIZATION PROBLEM OF INTEREST.

The MACH framework [4] developed by the University of Michigan is a state-of-the-art implementation of the MDF formulation that aimed to solve aerostructural problems. As such, we chose it as the framework to both compute our reference case and wrapping in our bi-level formulations. The MACH framework includes an aerostructural optimization tutorial [5] that solves the optimization problem (Opt. Problem 1) with the MDF formulation. It is considered as our optimization problem of interest.

Minimize	$2Mass + D_{cruise}$	(Opt. Problem 1)
w.r.t.	Twist (7), shape (96), 2 angles of attack, 108 panel thicknesses	
subject to	Cruise lift constraint $L_{cruise} - W_{ref} = 0$	
	Maneuver lift constraint $L_{man} - 2.5W_{ref} = 0$	
	Stress constraints $KS_i \leq 1.0,\ i = 0, \ldots, N$	
	Minimum volume constraint	
	100 minimum thickness constraints	
	16 fixed leading and trailing edge constraints	

As described, the design variables will be limited to the airfoils shape of the wing, the twist of the wing, and the thicknesses of the structure while the planform variables of the wing will be considered as fixed for the first applications described in the following sections. Two flight conditions are considered in the previous problem: a cruise flight point that allows assessing the aerodynamic performance, and a maneuver flight point that allows assessing the structure integrity through the computation of stress constraints.

4 A BI-LEVEL REFORMULATION OF THE PROBLEM.

4.1 Description.

As briefly described in section 1, and illustrated in Fig. 1, the IRT bi-level formulation consists in dividing the problem into:

- A system level optimization that solves shared design variables z (i.e. variables that influence all the disciplines) and subject to all of the problem constraints $\{g_i\}_{i\in[0,n]}$;
- Local level disciplinary optimizations that solve local design variables x_i (i.e. variables that influence a single discipline) and subject to disciplinary constraints g_i.

Two MDA stages are introduced in order to solve the coupling variables, before and after the sub-optimizations. As shown in Fig. 1, the process starts with initial values for the local design variables x^0 and the shared design variables z^0. These values are passed to the first MDA, denoted $y = \Psi(x_1, \ldots, x_n, z)$, to compute the coupling variables necessary for the sub-optimization blocks that run in parallel. The inner disciplinary optimizations minimize the objective function subject to their local constraints g_i. In a general approach, if the objective function f cannot be minimized in the sub-problem i, another function f_i can be considered as long as its decay is consistent (monotone decreasing) with the decay of the main objective function f. The decomposition is such that sub-problem constraints are assumed to depend on other block variables $x_{\neq i}$ only through the couplings that remain fixed during the sub-optimizations. Once an optimum value is found for each of them, the second MDA stage, $y^* = \Psi(x_1^*, \ldots, x_n^*, z)$, solves the couplings again to find the equilibrium and send the objective and constraint values to the system level optimizer so it can propose a new point of the shared design variables. The sub-optimization blocks do not exchange any information when they are solved in parallel, which means that the synchronization is ensured by the two MDAs and the system iterations, which warm start each block with the previous optimal values of the local variables. This process continues until convergence at the system level.

It is important to note that since the Jacobian matrix of the inner problem with respect to shared design variables cannot be computed (due to post-optimal sensitivity not available), a gradient-free algorithm is used at the system level (which is affordable due to the reasonable number of shared variables). On the other hand, a gradient-based algorithm is used at disciplinary levels since an efficient, analytic, computation of disciplinary Jacobian matrices is generally available with industrial tools.

The process described in Fig. 1 presents a general baseline of how the bi-level formulation has been thought. In practice, several variants can be implemented, as we show in section 5.2.1. These variants often depend on how the couplings between the disciplines can be managed.

4.2 Discussion of sub-problem separation.

The main difficulty related to the application of the bi-level formulation is how the set of design variables must be split with respect to the different sub-problems. This separability question is not always straightforward since all variables can have a non-negligible influence over all the disciplines. Practically, the sub-problems separation is often driven by the capability of the tools especially regarding the gradient computation. In the aero-structural (Opt. Problem 1), classical capabilities of industrial tools are:

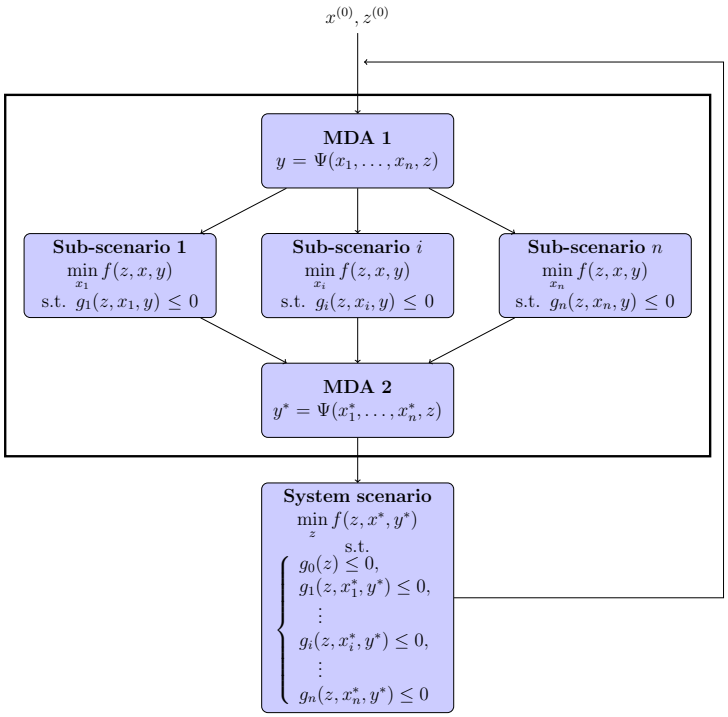


Figure 1: Simplified view of the IRT bi-level formulation.

- The structure solver computes the derivatives of its outputs (failure criteria, mass...) with respect to structural thicknesses;
- The aerodynamics solver computes the derivatives of its outputs (lift, drag) with respect to the shape of the airfoils and the twist of the wing.

It is known that structural thicknesses affect aerodynamics outputs through the stiffness and that the airfoils and the twist of the wing affect the structural outputs through the loads that are generated by the flow but also through the external shape that defines a bounding box of the structure geometry. The IRT bi-level formulation allows accounting for those coupling effects since the MDAs compute the coupling variables throughout system iterations and all disciplinary constraints are returned back to the system level.

However, when the couplings are very strong, local disciplinary design variables x_i have non-negligible effects on outputs of other disciplines $g_{\neq i}$ which can lead to constraints violation for the system level optimizer that are not correctly controlled by it [6]. In that case, the violation of the constraints returned back at the system level depends on the couplings while the system level optimizer directly acts on shared design variables, and not on couplings, which can prevent it to converge. This means that the main trade-offs are driven by the couplings while this information is not captured by any of the sub-problems. This situation can be the source of difficulty when separating the MDO problem into sub-problems.

Additionally, significant coupling effects mean that the optimal solution x^* is sensitive to the initial guess x and therefore that for same values of shared variables z, different solutions x^* can be obtained. Therefore, the objective function of the bi-level formulation is not a deterministic function from the point of view of the system, since the evaluation of the lower problem is dependent on the previous values of the local and coupling variables [6].

4.3 Proposed approach for highly-coupled sub-problems.

The problem of interest (Opt. Problem 1) presents a complex separation problem as described above especially in case of an aero-structure wing model. In particular, the shape of the wing defined by the twist and the airfoils which are efficiently managed into a gradient-based aerodynamic optimization, with the adjoint method, has a strong influence on the structure failure criteria since it directly affects the loads applied on the structure. To deal with this strong influence, it is proposed to introduce target variables on the loads at the system level that would be enforced at the disciplinary level through additional consistency constraints in a similar way as Analytic Target Cascading formulations [1]. To limit the number of those targets, a more global quantity, such as a force resultant, is considered. As a first approach to solve the (Opt. Problem 1), an interesting one that is classically used in industrial design is the bending moment of the wing around the fuselage axis (generally the x axis), denoted C_{mx} . However, on more complex applications, depending on the complexity of the loads that are detrimental for the structure failure criteria, the number of resultant forces would need to be adjusted.

The bi-level reformulation of the (Opt. Problem 1) with the introduction of the C_{mx}^{target} is referred as the C_{mx} -based approach. It leads to the new (Opt. Problem 2), implemented in the next section.

System scenario	(Opt. Problem 2)
Minimize	$2Mass + D_{cruise}$
w.r.t.	C_{mx}^{target}
subject to	Cruise lift constraint $L_{cruise} - W_{ref} = 0$
	Maneuver lift constraint $L_{man} - 2.5W_{ref} = 0$
	Stress constraints $KS_i \leq 1.0,\ i = 0, \ldots, N$
Structure sub-scenario	
Minimize	$Mass$
w.r.t.	Maneuver angle of attack, 108 panel thicknesses
subject to	Maneuver lift constraint $L_{man} - 2.5W_{ref} = 0$
	Stress constraints $KS_i \leq 1.0,\ i = 0, \ldots, N$
Aerodynamic sub-scenario	
Minimize	D_{cruise}
w.r.t.	Twist (7), shape (96), cruise angle of attack
subject to	$C_{mx} = C_{mx}^{target}$
	Cruise lift constraint $L_{cruise} - W_{ref} = 0$
	Minimum volume constraint
	100 minimum thickness constraints
	16 fixed leading and trailing edge constraints

5 APPLICATION ON A SIMPLIFIED AERO-STRUCTURE WING MODEL.

5.1 Model description.

The previously described IRT bi-level formulation with the C_{mx} -based approach is applied here to a simplified wing model. The model is known in the literature as the MACH tutorial model, it is relatively similar to the one considered in [7]. Although simplified compared to real industrial models, it is nevertheless considered a high-fidelity model with a shell finite element model for the structure part and a RANS solver for the aerodynamic part.

As previously explained, the MACH framework has been developed for last decades in order to solve problems similar to (Opt. Problem 1) with the MDF formulation [4]. It is used here for several purposes:

- First, to solve the (Opt. Problem 1) with the MDF formulation that is considered as the reference;
- Second, to implement disciplinary sub-optimizations of the proposed bi-level approach (Opt. Problem 2) to be compared to the MDF solution.

The bi-level approach (Opt. Problem 2) is implemented through the GEMSEO [8] framework; only sub-optimization scenarios are implemented with MACH and considered as black-box disciplines.

5.1.1 Structural model description.

The structural model is a conventional wingbox arranged like a grid, with ribs intersecting spars and stringers at perpendicular junctions (see Fig. 2a). When creating the model, we defined 19 rib stations and 10 spar stations, which means the wingbox surfaces are currently split into 18 span wise zones and 9 chord wise zones. In TACS (the structural solver of the MACH framework), we usually refer to each zone as a *component*.

For the purpose of the structural optimization problem, the design variables will be the 108 panel thicknesses that correspond to each of the components that we previously defined. The design variable components are shown by color in Fig. 2b.

5.1.2 Aerodynamic model.

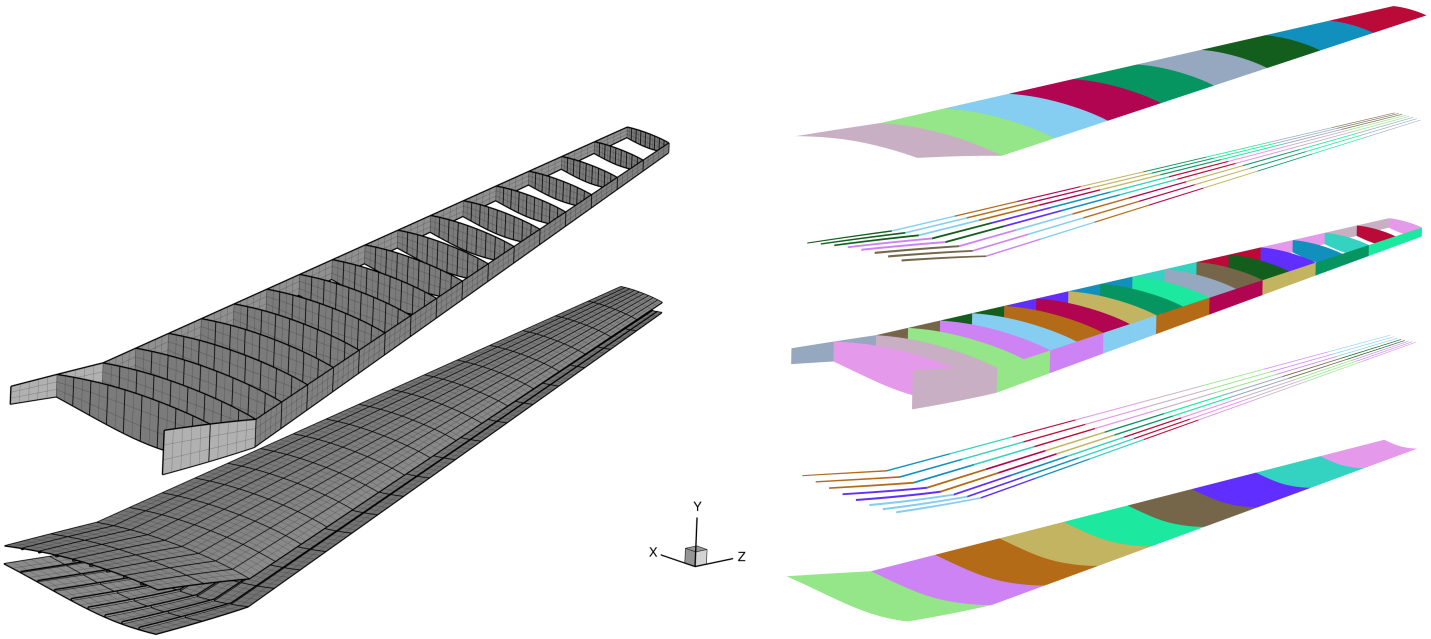
The parametrization of the aerodynamic model is based on a free-form deformation (FFD) method. The mesh for the aerodynamic analysis is composed of 193536 cells (see Fig. 3).

5.2 IRT bi-level application.

5.2.1 Process description.

The IRT bi-level variant with C_{mx} -based approach described in (Opt. Problem 2) is implemented on the previously described model. Fig. 4a shows the XDSM of the system level process while Fig. 4b and Fig. 4c show the XDSM of the structure and aerodynamic sub-scenarios described in (Opt. Problem 2).

Compared to the general description of the formulation in section 4.1 and Fig. 1, a few differences can be observed with the current implementation. First, the sub-scenarios



(a) Structural finite element model [5]. (b) Visualization of the structural design variables [5].

Figure 2: Structure model.

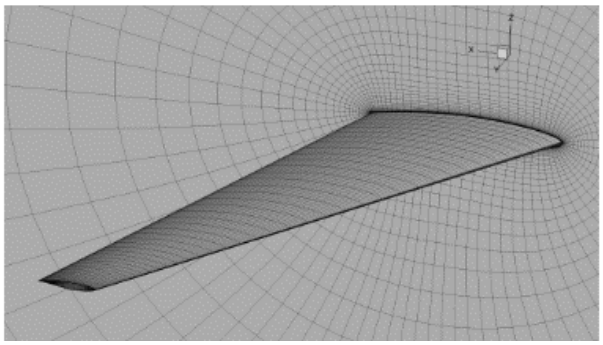
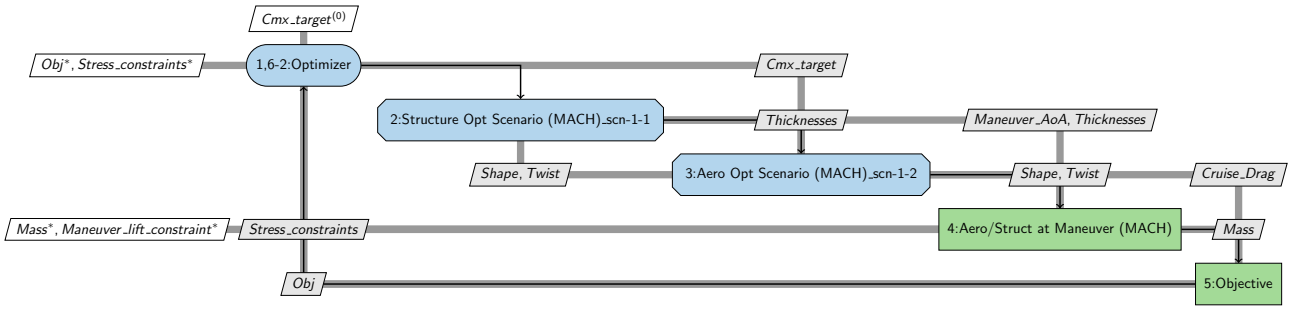


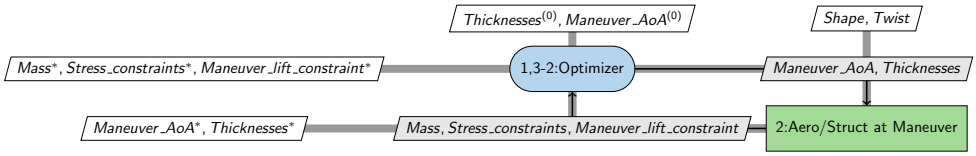
Figure 3: Aerodynamic mesh.

are sequential rather than parallel in the general approach; this is because the MPI-based solvers in MACH (aerodynamic and structure) could not be easily handled with the Python-based multiprocessing approach developed in GEMSEO, with our available infrastructure. However, several tests were performed to replicate the parallel behavior (i.e. taking for the block i the local variables $x^*_{\neq i}$ at their previous values from system iterations) and similar conclusions can be drawn from the following results. Second, there is no need for the MDA1 at the beginning of the inner problem because both the aerodynamic and structure optimization sub-scenarios update their own coupling variables (loads and displacements). Indeed the sub-scenario presented in Fig. 4b that corresponds to the structure sub-optimization problem in (Opt. Problem 2) compute the aero-structure MDA at maneuver conditions (denoted as `Aero/Struct at Maneuver`) while the sub-scenario Fig. 4c that corresponds to the aerodynamic sub-optimization problem computes the aero-structure MDA at cruise conditions (denoted as `Aero/Struct at Cruise`). However, the MDA2 (denoted as `Aero/Struct at Maneuver (MACH)` in

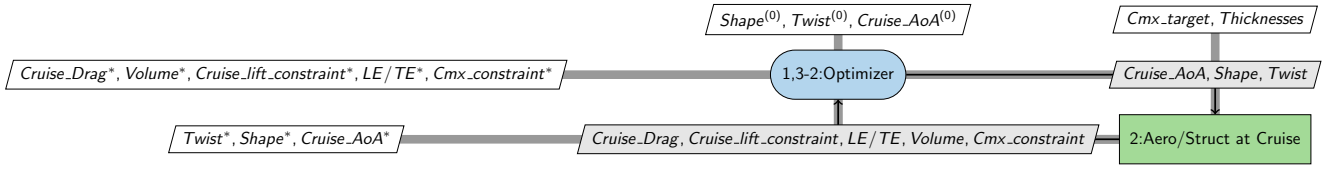
Fig. 4a) remains since we need to re-compute the structural mass and constraints with the optimal aerodynamic variables to return back to the system level optimizer. It must be noted in the system level XDSM Fig. 4a that all blocks including `(MACH)` characters mean that they directly wrap MACH executions as black-boxes. A few details about the implementation are given in the next section.



(a) System scenario.



(b) Structure optimization scenario.



(c) Aerodynamic optimization scenario.

Figure 4: XDSM diagrams of the IRT bi-level formulation with the C_{mx} -based approach.

5.2.2 Implementation description.

To interface the MACH solvers and optimizers with GEMSEO, we converted the scripts from the tutorials provided with the software into Python functions that take as input the values for the variables that we wish to set and return the results of either an aerostructural MDA (`Aero/Struct at Maneuver (MACH)` in Fig. 4a), a structural optimization (`Aero/Struct at Maneuver` in Fig. 4b), or an aerodynamic optimization (`Aero/Struct at Cruise` in Fig. 4c). The data exchange is done in memory; some result files and logs are written to unique numbered folders. Our wrapped MACH disciplines use the MPI protocol inside a Docker image to execute the optimization or analysis in parallel, taking advantage of the MACH framework's native support for MPI.

All the computations (including the reference results and the bi-level configurations that will be presented in section 6) ran using 4 nodes, each of them with two Intel Xeon

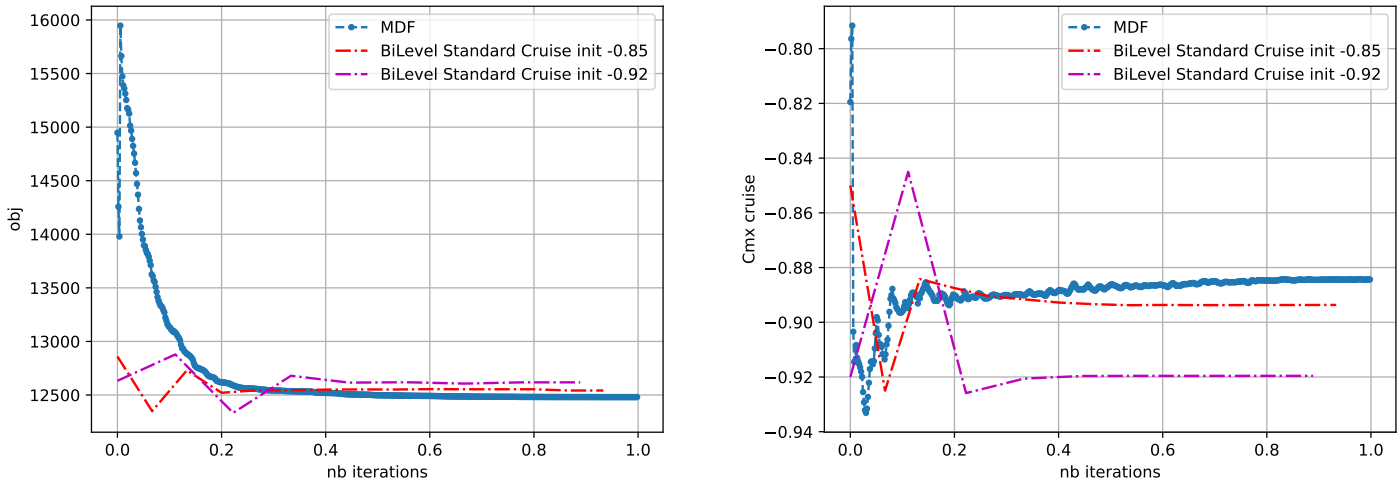
Gold 6140 processors (18 cores skylake processor at 2.3 Ghz) and 96 GB DDR4 memory. Each node ran 18 MPI processes.

5.2.3 Results.

To assess the robustness of the IRT bi-level approach with respect to the initial value of design variables, two results are compared with the MDF solution:

- The first result is obtained when initializing the C_{mx}^{target} at system level with -0.85
- The second result is obtained when initializing the C_{mx}^{target} with -0.92

In both cases, the initial values of local disciplinary variables (thicknesses, airfoils shape and twists) have been obtained from a preliminary sequence of disciplinary optimizations. The results of the evaluation of the objective function are presented in Fig. 5a. The evolution of the C_{mx} value is presented in Fig. 5b. In both cases we see that the convergence value is different when the initial C_{mx}^{target} is not the same. This difference illustrates the sensitivity of the solution with respect to the initial values of the shared variable, as briefly explained in Section 4.2. This sensitivity motivates an extension to solve the lower bi-level problem more accurately; this new approach is presented in the next section.



(a) Objective function history.

(b) C_{mx} value history.

Figure 5: Optimization history results for the IRT bi-level with C_{mx} -based approach. Comparison with MDF.

6 BI-LEVEL ENHANCEMENT: BLOCK COORDINATE DESCENT STRATEGY.

6.1 Description of the bi-level BCD.

Previously, we saw that the results for the IRT bi-level formulation were different if we started from different values of C_{mx}^{target} . The reason for this is that from the system level standpoint, the outputs of the inner problem have some discrepancies with respect

to the system level variables because they implicitly depend on the local variables [6]. This affects convergence properties.

The bi-level BCD adds more robustness and stability with respect to the previous bi-level formulation regarding the discrepancy of the lower problem resolution. This new approach is shown in Fig. 6 where it can be seen that the lower problem is solved by a Block Coordinate Descent (BCD) method, also known as Block Gauss-Seidel (BGS) method, which means that each block is sequentially optimized within an iterative loop until convergence. This additional iterative loop is not present in the previous IRT bi-level (Fig. 1). As a consequence, each block i is updated at every BCD iterations with $x^*_{\neq i}$ until a fixed point x^* is found, regardless the initial guess x, which drastically reduces the discrepancy of the lower problem solutions.

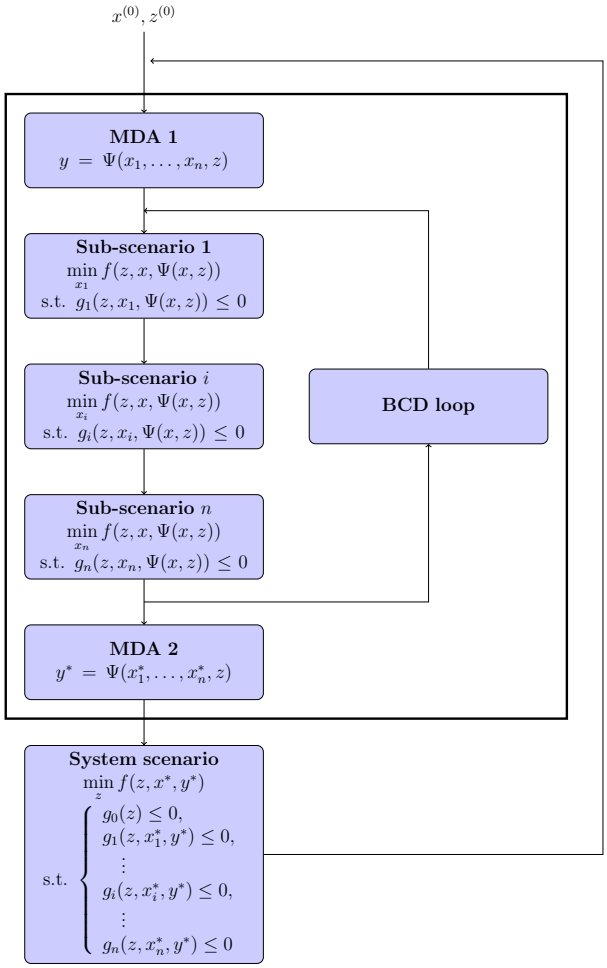


Figure 6: Simplified view of the BCD bi-level formulation.

David [6] used Sobieski's super-sonic business jet (SSBJ) test case is used to assess the sensitivity of the IRT bi-level and bi-level BCD with respect to the choice of the starting point. The authors plotted the performance profile, shown in Fig. 7, for 150 different starting points. Two different bi-level BCD variants are compared with the standard IRT bi-level and MDF. The two bi-level BCD differ according to the way in which the couplings are solved:

- BCD-MDF means that couplings are solved by running MDAs within each sub-

optimization ;

- BCD-WK means that each sub-optimization no longer solves the whole coupling vector but only its own block of coupling variables, which are then exchanged through the BCD loop.

In the current aerostructural application only the BCD-MDF is considered since the aero-structure couplings (forces and displacements) are always computed within the sub-scenarios. In Fig. 7, it can be seen that the IRT bi-level formulation (IRT on the plot legend) converges less than 40% of the time regardless of the amount of disciplinary calls, while the BCD versions converge more than 90% of the time when the number disciplinary calls is high enough.

A down side of the introduction of the BCD is that every system iteration becomes more costly because it requires several iterations over the sub-optimizations to converge, as shown in the convergence history (Fig. 7). This additional cost may be larger during the initial system optimizations because the system variables vary more and thus make the sub-scenarios produce local design variables that are very different from one system iteration to the next.

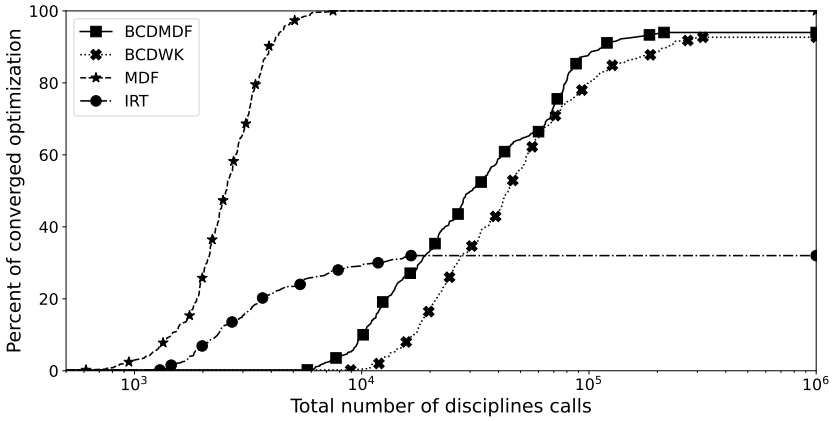


Figure 7: Performance profile of the Sobieski's SSBJ problem with 150 random starting points [6]. Two BCD variants are compared to MDF and the standard IRT bi-level.

6.2 Application to the aerostructural optimization problem.

The bi-level BCD strategy was applied to our aero-structure wing model. The XDSM of the process is shown in Fig. 8. Note that the structure and aerodynamic sub-scenarios remain exactly the same as with the previous IRT bi-level, as presented in Fig. 4b and Fig. 4c; they are not shown again here. In the XDSM Fig. 8, the BCD loop represented by an `MDAGaussSeidel` object in GEMSEO is added between the two sub-scenarios. Within this loop all optimal values of local variables are exchanged.

6.2.1 Results.

After running the bi-level BCD formulation with different initial values for the C_{mx}^{target} , we showed that the results are similar regardless of the starting point. However, a few differences can still be observed when comparing to the MDF reference solution.

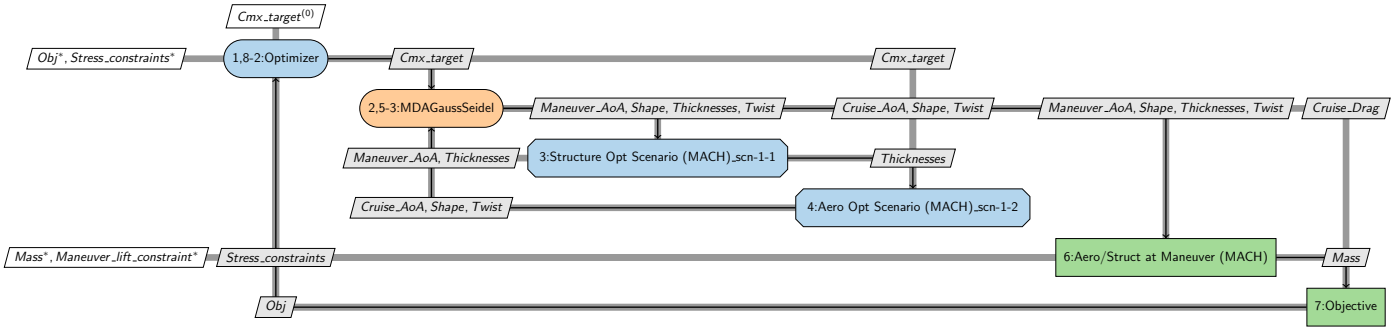


Figure 8: XDSM of the bi-level BCD formulation.

In terms of the objective value (Fig. 9a), the bi-level BCD runs get to almost the same value at convergence regardless of the initial C_{mx}^{target} . The plot also shows the solutions for the same initial C_{mx}^{target} without the BCD MDA (from the previous section), notice how the solution for $C_{mx}^{target} = -0.92$ gets closer to the reference with the BCD approach. The same behavior is observed in terms of the C_{mx} history on Fig. 9b.

If we look at the main components of the objective function, both BCD solutions are closer to the reference in terms of cruise drag (Fig. 9c) and mass (Fig. 9d) than the IRT bi-level. Another interesting point is that both kinds of bi-level formulations tend to privilege drag minimization at the expense of mass increments.

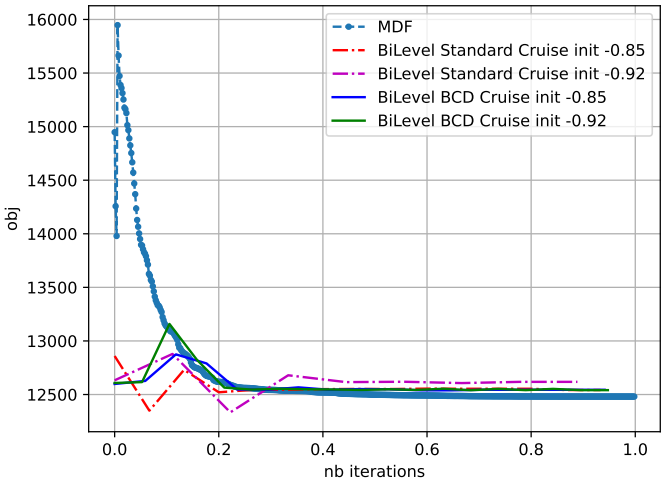
Turning our attention to the lift distribution (Fig. 10c), the BCD runs produce a non-elliptical distribution that is close to the reference (MDF). In case of the IRT bi-level formulations, we see that the run with $C_{mx}^{target\ 0} = -0.85$ captures the non-elliptical distribution whereas the run with $C_{mx}^{target\ 0} = -0.92$ does not, which illustrates once more the sensitivity of the IRT bi-level to the choice of the initial point.

Regarding the twist distribution (Fig. 10a), we do not manage to capture the same behavior as the MDF solution, even with the bi-level BCD approach. However, we do see that the gap between the solutions of two different starting points gets significantly reduced if the bi-level BCD is used. The thickness/chord (Fig. 10b) does not capture the MDF behavior either, although in this case all the bi-level formulations give practically the same result.

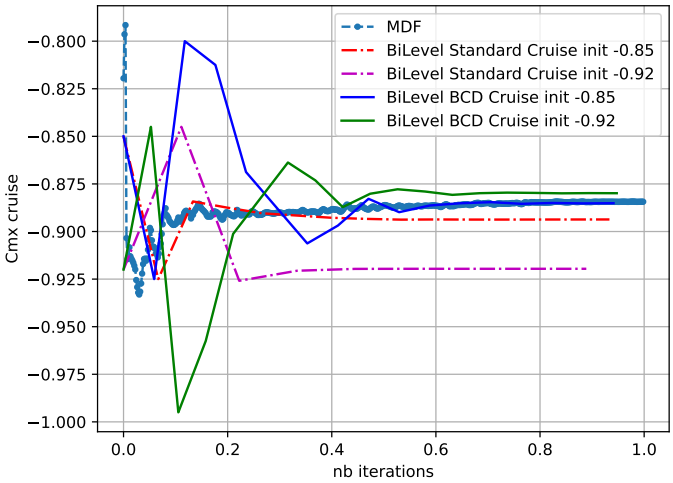
Figures 11 show the evolution of structural thicknesses along the span for both bi-level formulations with an initial $C_{mx}^{target\ 0} = -0.85$, along with the reference solution (runs with $C_{mx}^{target\ 0} = -0.92$ are not shown for clarity but the conclusion is similar to the one drawn from previous results). Both bi-level solutions tend to have thicker elements, which is consistent with the larger overall mass observed in Fig. 9d.

Table 1 presents a summary of results for both bi-level formulations, with the two different starting points, and the reference case. In addition to the aspects that we have commented before (objective value, mas, drag, C_{mx}) we also compare the execution time for each case. The IRT bi-level formulation takes around 4 times more to converge than MDF; this is an expected outcome because the bi-level formulation requires the computation of two costly sub-optimizations nested into a gradient-free system optimization. The bi-level BCD formulation takes around 11 times more to converge than MDF; this large increase in computational cost comes from the iterative solution of the BCD loop.

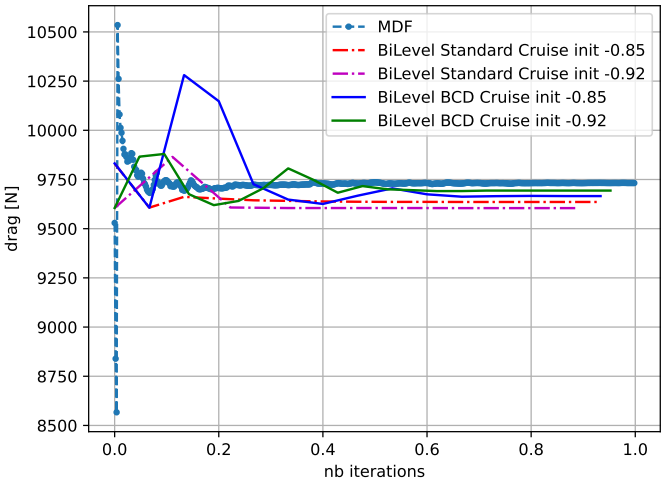
While these results highlight some remaining challenges, the BCD formulation represents a significant step forward in providing a formulation that overcomes the implemen-



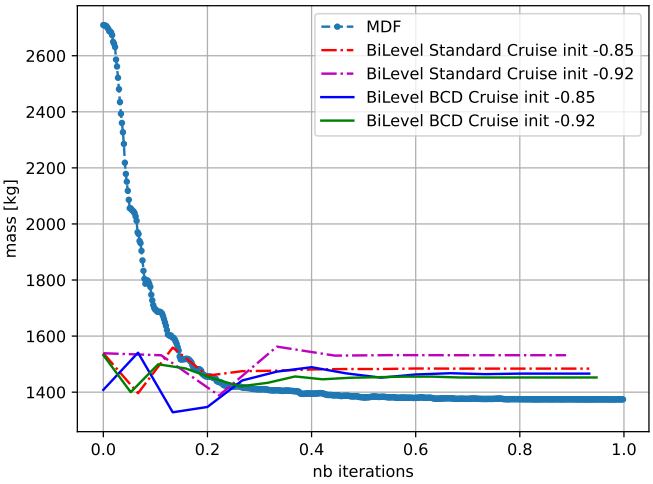
(a) Objective function.



(b) C_{mx} .



(c) Cruise Drag.

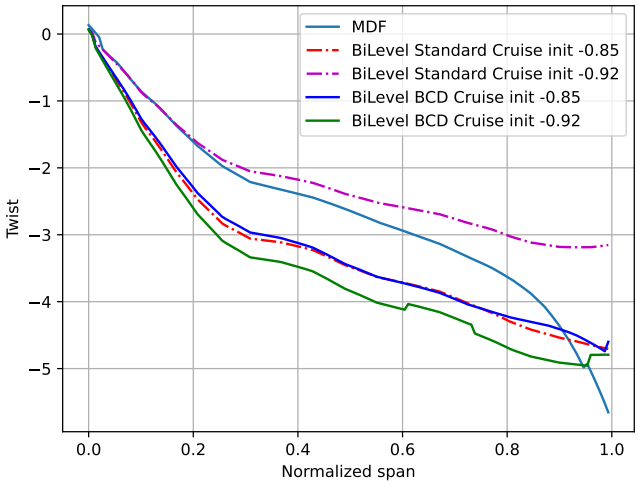


(d) Mass.

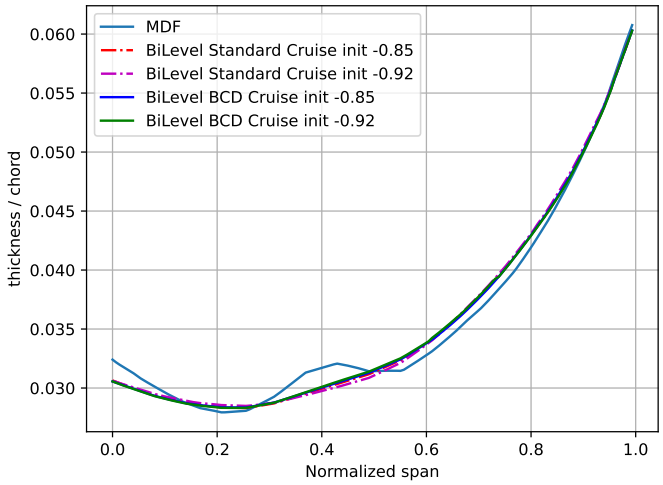
Figure 9: Optimization convergence history for the bi-level BCD with C_{mx} -based approach. Comparison with MDF and IRT bi-level.

tation issues that limit the applicability of MDF in industrial contexts. Also, this work is the first application of the bi-level BCD formulation and we think that the overall restitution time could be improved by better adjusting the different numerical settings involved within the different algorithms (multiple MDAs, BCD loop, and multiple optimizers) and by using MDA acceleration methods.

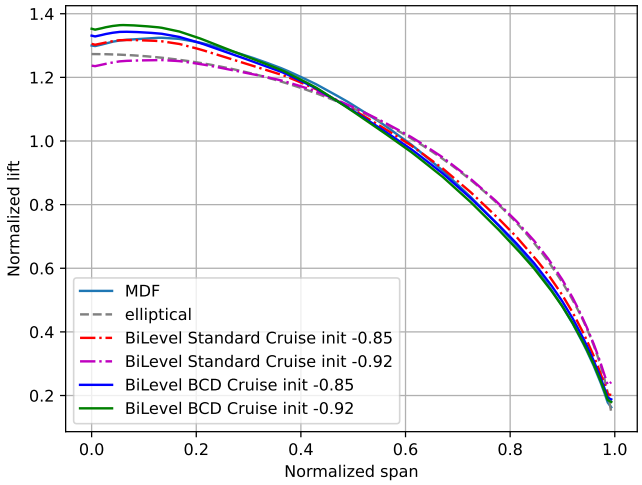
Regarding the accuracy of bi-level solutions, it can be seen that the optimal values of design variables are close but still different than the MDF ones. This makes the authors think that, whatever the bi-level formulation presented, standard IRT or BCD, the good trade-off between the aerodynamic and structure discipline is still not fully captured. In order to better understand why, the next section investigates a new analysis.



(a) Twist distribution.



(b) Airfoil thickness over chord.



(c) Lift distribution.

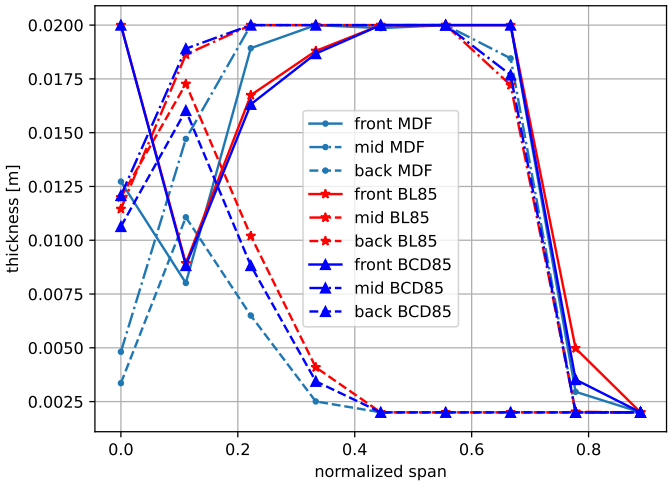
Figure 10: Optimal geometric solutions for the bi-level BCD with C_{mx} -based approach. Comparison with MDF and IRT bi-level.

Table 1: Result summary for both bi-level formulations.

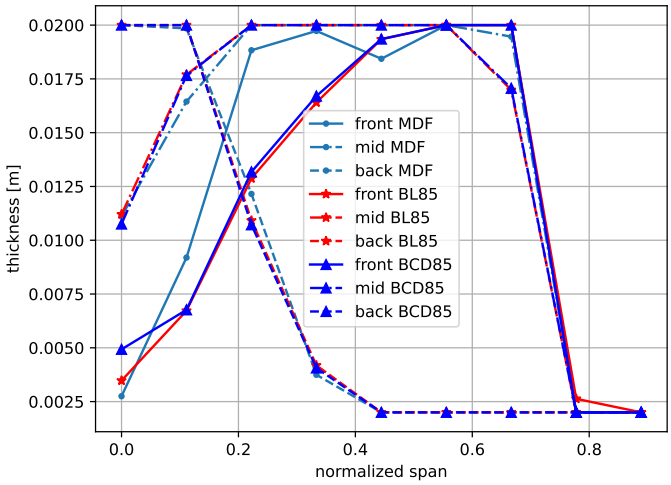
Optimization Case	**Obj. [N]** $\times 10^{-4}$	**Drag [N]**	**Mass [kg]**	**Opt** C_{mx} **[-]**	**Exec Time**
MDF (Reference)	1.2481	9731.99	1374.35	-0.8844	3 h 48 m
IRT bi-level $C_{mx}^{target\ 0} = -0.85$	1.2542	9635.75	1484.01	-0.8907	12 h 23 m
IRT bi-level $C_{mx}^{target\ 0} = -0.92$	1.2618	9604.52	1531.82	-0.9195	11 h 48 m
Bi-level BCD $C_{mx}^{target\ 0} = -0.85$	1.2536	9665.69	1466.16	-0.8851	33 h 12 m
Bi-level BCD $C_{mx}^{target\ 0} = -0.92$	1.2548	9693.64	1452.38	-0.8826	35 h 25 m

6.2.2 Results of the bi-level BCD formulation without aerodynamic shape optimization.

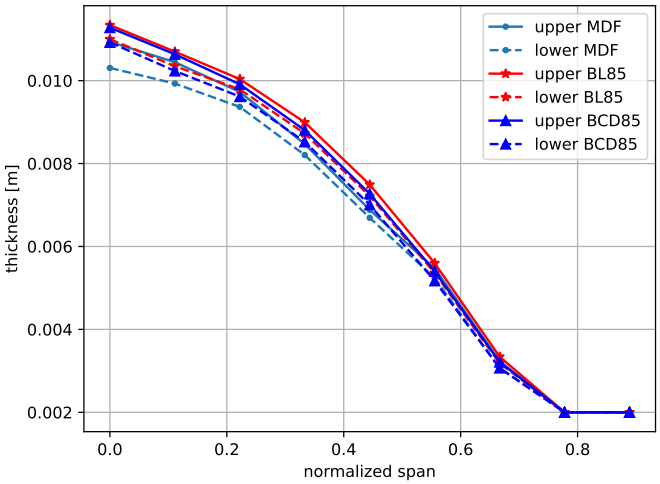
This new analysis focuses on a simplified version of the bi-level (Opt. Problem 2), solved with the bi-level BCD, where the shape design variables of the aerodynamic sub-



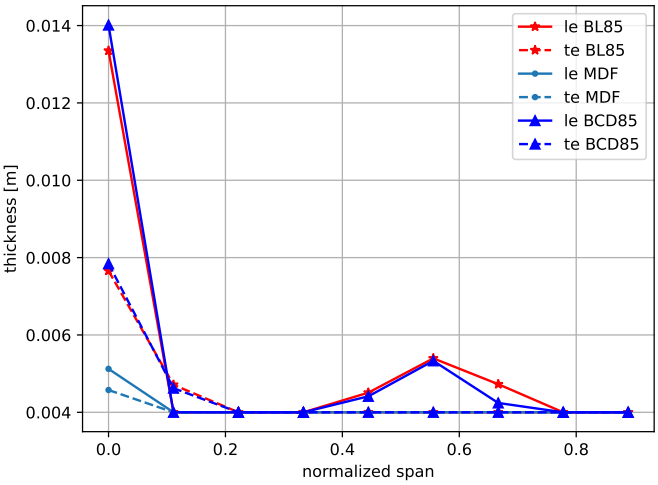
(a) Upper stringer thickness distributions.



(b) Lower stringer thickness distributions.



(c) Skin thickness distributions.

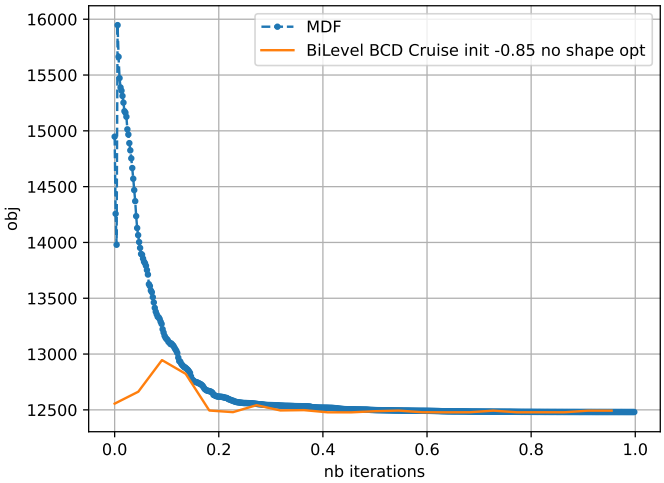


(d) Spar thickness distributions.

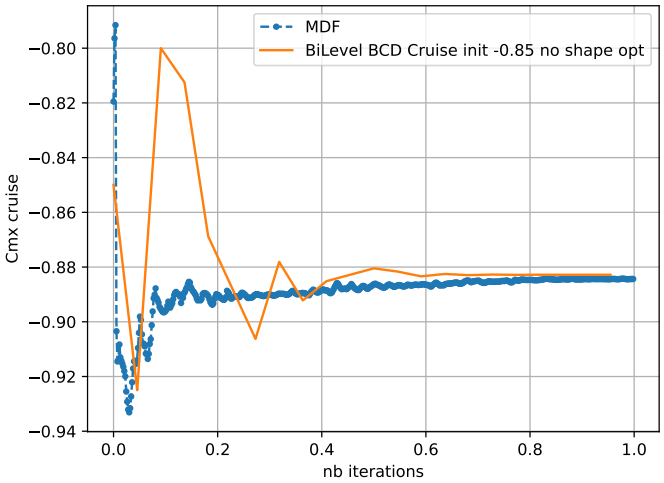
Figure 11: Optimal structure variables for the bi-level BCD with C_{mx} -based approach. Comparison with MDF and IRT bi-level.

optimization remain fixed, with values taken from the MDF solution. This means that only twist variables are optimized within the aerodynamic sub-problem, all the other sub-problems (system and structure) remaining unchanged. The reason for this simplification is that it is thought that the influence of the shape of the wing on the structure sub-problem through the wing box size (airfoil thicknesses) may not be correctly captured.

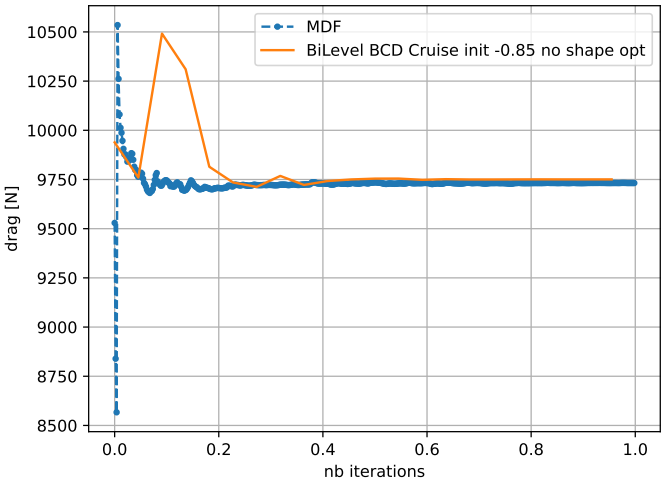
Fig. 12 shows the optimization history of the objective function, the C_{mx} , the cruise drag and the mass where a better agreement between the bi-level BCD and MDF is observed, especially on the optimal mass. Fig. 13 shows the same trend with a twist distribution (Fig. 13a) that has the same shape than the MDF solution despite a few differences. Fig. 13b matches exactly as expected since the airfoil shapes (thus thicknesses and chords) are forced to the optimal MDF solution. Finally, the optimal structural values



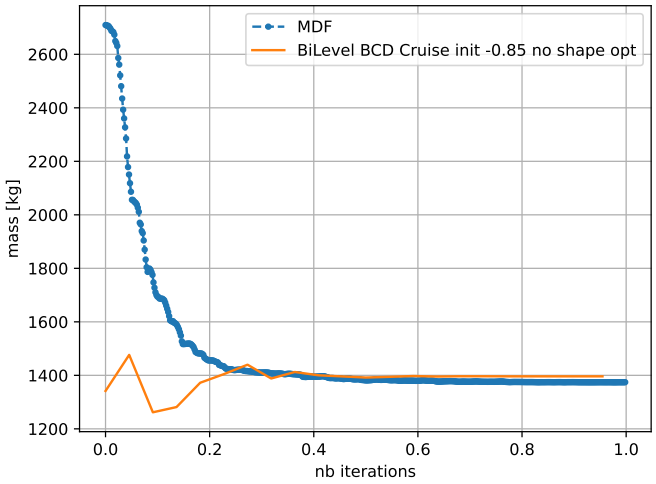
(a) Objective function.



(b) C_{mx} .



(c) Cruise Drag.



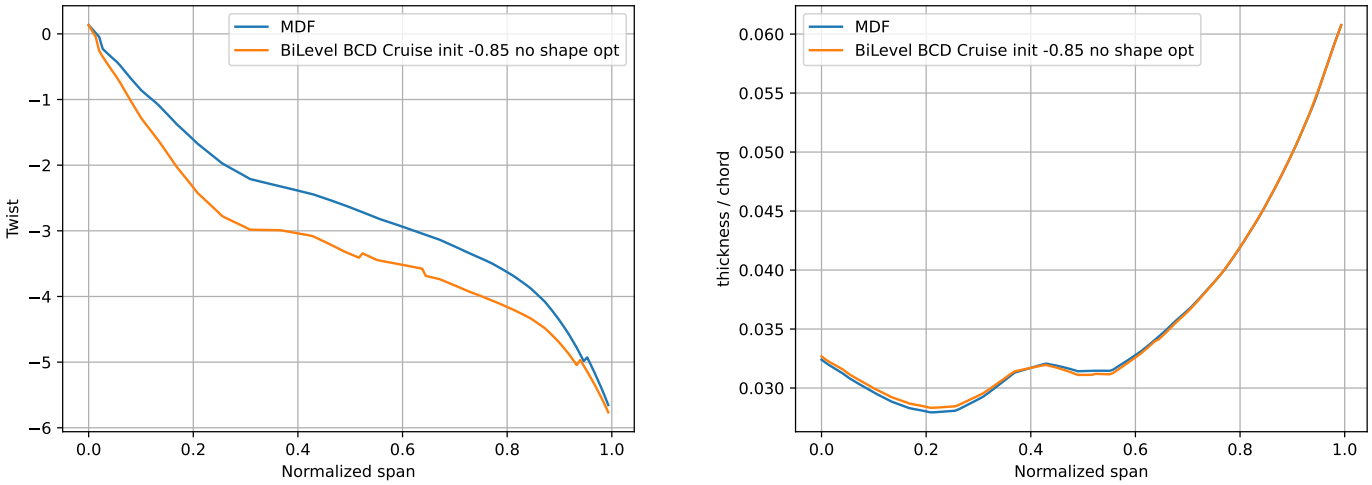
(d) Mass.

Figure 12: Optimization history results of the bi-level BCD with C_{mx} -based approach without shape optimization. Comparison to the MDF formulation.

presented in Fig. 14 also show a better agreement to the MDF solution than the previous bi-level solutions. This leads us to the conclusion that there is a need to investigate more deeply a way to account for the overall influence that airfoils shapes have on both disciplinary sub-optimizations.

7 CONCLUSIONS

The IRT bi-level formulation has been proposed to solve large scale aero-structure coupled problems, taking into account industrial constraints (no coupled adjoint, large number of constraints, etc.). The proposed C_{mx}^{target} approach allows to better handle the separability problem between sub-scenarios and helps finding interesting trade-offs. Encouraging results were obtained although a lack of robustness was observed with respect



(a) Twist distribution.

(b) Airfoil thickness over chord.

Figure 13: Optimal geometric parameters of the bi-level BCD with C_{mx} -based approach without shape optimization. Comparison to the MDF formulation.

to the initial point of the optimization due to an inherent discrepancy of the inner problem outputs.

The bi-level BCD has been presented as an enhancement to handle the aforementioned discrepancy. A BCD loop on the sub-optimizations solves the local design variables, increasing the robustness of the solution regardless of the initial point of the optimization.

We were not able to find the exact same solution of the MDF formulation with either of the bi-level formulations, but the results are encouraging. Further work is needed on the problem separation approach to better capture the trade-offs of the aero-structural optimization, especially regarding the management of airfoil thicknesses.

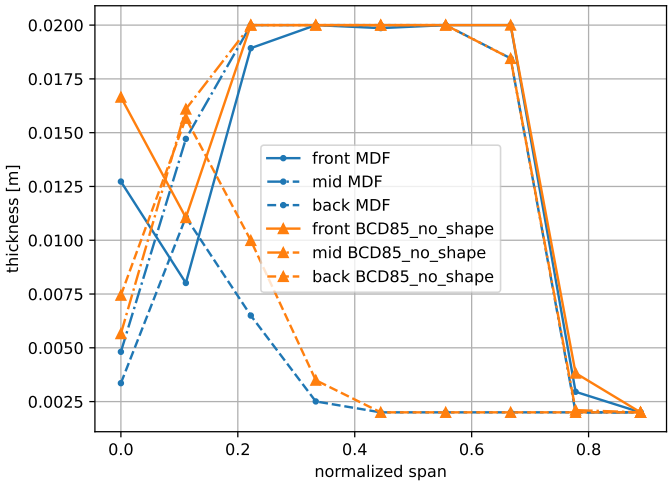
The authors think that there is still room for improvement in terms of the computational cost of the bi-level formulation. These include, but are not limited to: parameters tuning and acceleration methods on the BCD loop, and exploration of other gradient-free optimization algorithms. Another perspective would be working on accounting for other planform variables.

ACKNOWLEDGEMENTS

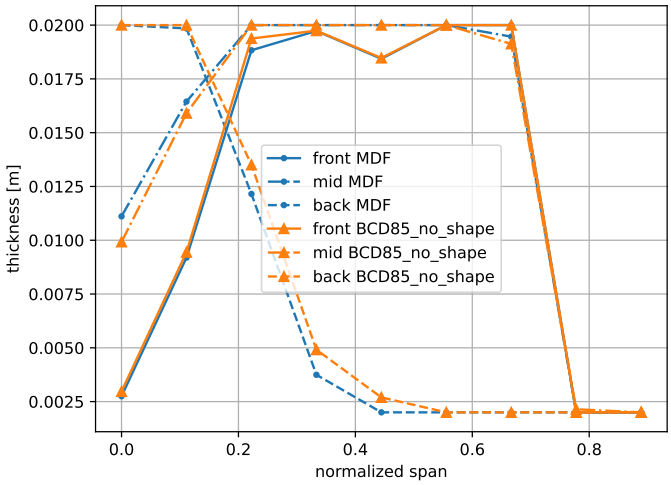
IRT Saint Exupéry wishes to acknowledge the industrial members of the IRT Saint Exupéry project R-Evol: Airbus, Liebherr, Altran Technologies, Capgemini DEMS France, CENAERO, Expleo Group and CERFACS for their support, financial funding and knowledge in this work.

This work was supported by the French government under the France 2030 program under the reference ANR-10-AIRT-01.

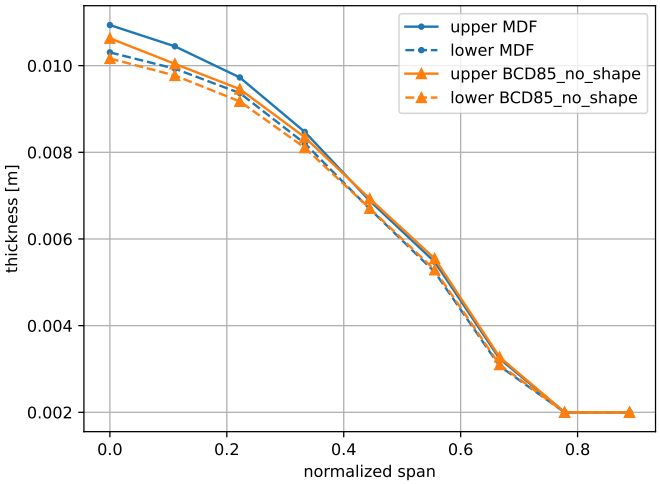
In addition, some of the work of the IRT Saint Exupéry was undertaken as part of the project Ultra Performance Wing (UP Wing, project number: 101101974) supported by the Clean Aviation Joint Undertaking and its members co- Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Clean Aviation Joint Undertaking.



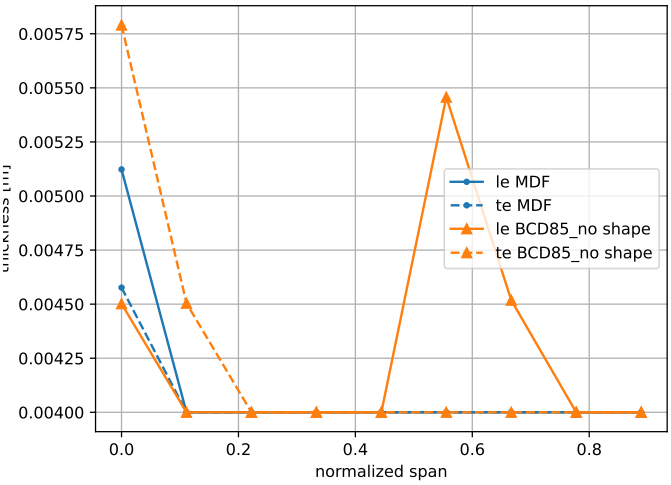
(a) Upper stringer thickness distributions.



(b) Lower stringer thickness distributions.



(c) Skin thickness distributions.



(d) Spar thickness distributions.

Figure 14: Optimal structure variables of the bi-level BCD with C_{mx} -based approach without shape optimization. Comparison to the MDF formulation.

Neither the European Union nor the granting authority can be held responsible for them.

CLEAN AVIATION Co-funded by the European Union

REFERENCES

[1] J. R. R. A. Martins and A. B. Lambe. Multidisciplinary design optimization: a survey of architectures. *AIAA journal*, 51(9):2049–2075, 2013.

[2] A. Gazaix, F. Gallard, V. Ambert, D. Guénot, M. Hamadi, S. Grihon, P. Sarouille, T. Y. Druot, J. Brézillon, V. Gachelin, et al. Industrial application of an advanced bi-level mdo formulation to aircraft engine pylon optimization. In *AIAA Aviation 2019 Forum*, page 3109, 2019.

[3] A. Gazaix, V. Ambert, M. Guillen, N. Roussouly, F. Gallard, T. Druot, S. Coniglio, M. Hamadi, V. Drouet, J.-C. Giret, et al. Industrialization of mdo methods, software and processes, from mbse to high fidelity multi-components wing, pylon and nacelle optimisation. In *AIAA AVIATION FORUM AND ASCEND 2024*, page 4404, 2024.

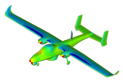
[4] G. J. Kennedy and J. R. R. A. Martins. A parallel aerostructural optimization framework for aircraft design studies. *Structural and Multidisciplinary Optimization*, 50: 1079–1101, 2014.

[5] MDOLab. Mach tutorial. `https://mdolab-mach-aero.readthedocs-hosted.com/en/latest/machAeroTutorials/intro.html`. [Accessed 19-03-2025].

[6] Y. David, F. Gallard, and A. Rondepierre. Locally convergent bi-level MDO architectures based on the block coordinate descent algorithm. working paper or preprint, Nov. 2024. URL `https://hal.science/hal-04758286`.

[7] A. C. Gray and J. R. R. A. Martins. A proposed benchmark model for practical aeroelastic optimization of aircraft wings. In *AIAA SciTech 2024 Forum*, page 2775, 2024.

[8] F. Gallard, C. Vanaret, D. Guénot, V. Gachelin, R. Lafage, B. Pauwels, P.-J. Barjhoux, and A. Gazaix. Gems: a python library for automation of multidisciplinary design optimization process generation. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, page 0657, 2018.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AEROSTRUCTURAL OPTIMIZATION OF HIGH-ASPECT-RATIO WINGS: NAVIGATING COMPLEX TRADE-OFFS

Ousmane Sy[1,2*], Joseph Morlier[2] Emmanuel Benard [1]

1: Fédération ENAC ISAE-SUPAERO ONERA
Université de Toulouse
31400 Toulouse,France
ousmane.sy2@isae-supaero.fr, emmanuel.benard@isae-supaero.fr

2: ICA, ISAE-SUPAERO, MINES ALBI, UPS, INSA, CNRS
Université de Toulouse
Toulouse, France
joseph.morlier@isae-supaero.fr

Abstract. *A promising innovation path in aircraft design aimed at improving fuel efficiency and reducing emissions is the adoption of concepts featuring increased wing aspect ratios. This study investigates the aerostructural optimization of high aspect ratio (HAR) wings using the uCRM-13.5 benchmark and the OpenAeroStruct framework. The impact of model fidelity, objective metrics, analysis methods, and material selection on optimal wing design is examined. Results showed that neglecting viscous and compressibility effects leads to unrealistic designs, with fuel burn underestimated by up to 45%. Mid-fidelity models, which combine a vortex lattice method-based aerodynamic model with estimations of viscous and compressibility drag, and a 1D beam finite element structural model, achieve results within 6% for fuel burn and 4.85% for wing weight compared to high fidelity models, while reducing computational time by over 97%. Objective metric selection was found to significantly influence the optimum design. Pareto fronts analysis revealed that extreme weighting cause disproportionate trade-offs between fuel burn, structural weight and take-off gross weight (TOGW). For instance, a 0.6% improvement in structural weight led to a 21.28% increase in fuel burn, while a 0.17% improvement in TOGW came with a 2.26% increase in fuel burn. Optimal designs lie between these extremes, balancing direct operating, manufacturing and acquisition costs. Different analysis methods were evaluated across mission ranges. While single- and multi-point approaches performed adequately for long-range missions, they underestimated fuel burn by up to 50% for short missions. The sequential method, which includes the climb phase, offered more accurate fuel predictions at lower computational cost compared to the multi-point method.Material exploration showed that switching from aluminium to composite materials reduced fuel burn and structural weight by 18% and 13%, respectively, for fuel burn-focused design and 31% and 2.24%, respectively, for structural weight focused design. However, this improvement comes at the cost of increased CO_2 emissions from material manufacturing.*

Keywords: Aircraft design, High aspect ratio wings, MDO, Aerostructural optimization

1 INTRODUCTION

The exponential growth in air transport [1] poses a real problem for the sustainability of aviation [2]. To address this situation, NASA [3] and EASA [4] have established ambitious sustainable aviation goals , aligning with the general objectives of the Paris agreement. This environmental context, along with economic demands, calls for new and much more efficient aircraft concepts. A promising innovation path in aircraft design that aims to improve fuel efficiency and reduce emissions is the adoption of concepts that feature increased wing aspect ratios. Quantitative research studies [5–8] show that increasing the wingspan can result in a substantial decrease in fuel consumption, mainly by reducing the induced drag and therefore lowering emissions.

However, the high aspect ratio (HAR) wings present some drawbacks, making their design very tedius. One of the drawbacks of increasing the aspect ratio is the corresponding structural weight due to the maximum bending moment. Unconventional HAR wing concepts, such as the strut-braced [9] and twin-fuselage [10] concepts, have been proposed to overcome this problem and allow for a larger wingspan. Furthermore, the extended flexibility of HAR aircrafts makes them particularly susceptible to aeroelastic phenomena such as flutter and divergence.

On the design methodology, these downsides make the traditional aircraft design methods and tools often inadequate for such wing concepts. High aspect-ratio wing designs require the simultaneous consideration of both aerodynamics and structure, as these disciplines are tightly coupled. This coupling is addressed through aerostructural optimization, which concurrently manages structural sizing and aerodynamic shape and/or planform optimization. This is essential not only to achieve the optimal balance between low weight and high aerodynamic performances but also to address aeroelastic phenomena. However, the aerostructural optimization of high aspect ratio (HAR) wings is highly challenging, and each aspect must be carefully considered as it significantly impacts the optimal design.

Aerostructural optimization is performed in the early design process, where many different designs are evaluated and optimized to determine the best one. Therefore, computational cost is crucial. In contrast, precision with a drag prediction accuracy of one drag count is needed [11], which is equivalent to the weight of one passenger [12]. Hence, the choice of models fidelity is critical and based on an optimal trade-off between accuracy and computation time. Due to their computational cost, low-fidelity modeling methods, which combines a lifting- line model [13] or vortex lattice method [8, 14, 15] with a simple finite element model remain popular and useful for preliminary aerostructural wing design. Nevertheless, methods with lower fidelity trade off precision and design responsiveness for reduced computational expense. The mid-fidelity methods [16, 17], are used to enhance the low-fidelity models accuracy while maintaining the computational cost lower. These levels of fidelity may not be able to capture all the inherent phenomena associated with unconventional concepts like HAR wings. Thus, the use of high-fidelity methods, which combine Euler CFD [18, 19] or RANS CFD [20–22] models with a detailed 3D wingbox finite element structural model, becomes a necessity to achieve the best optimal design, with the drawback of high computational cost.

The choice of the objective metric driving the optimization is not straightforward and has an significant impact on the optimum design. Therefore it must be carefully chosen to accurately reflect the design intent of the aircraft. Early aerostructural optimization were based on drag and structural weight minimization or a combination of both [13, 18, 23].

The fuel cost is a significant part of the direct operating cost and represent a trade-off between the structural weight and aerodynamic efficiency. This, combined with the goal of reducing gas emissions, makes fuel burn the most practical objective function for aerostructural optimization. However, the trade-off between acquisition and operating costs varies by airlines, aircraft type, and fuel prices. To develop a commercially attractive aircraft, an optimal trade-off of these costs is needed [22, 24, 25].

Another important aspect of aerostructural optimzation is the way of computing the objective function, which can significantly impact the optimal design. The method can be single-point [8, 26], multi-point [27–29] or mission-based [30, 31]. Single-point method compute the metrics at a single cruise flight condition, which make it prone to exhibit poor-off design performances. Multi-point aerostructural optimization overcome this by considering multiple cruise flight points. The objective function is a weighted average of the performance at the multiple design points. This method improve the average of the performance across the stencil of flight points with the cost of resulting to intermittent performance. However, both approaches focus solely on the cruise phase and may not be suitable for shorter mission aircrafts. Instead, mission-based analysis provides a more meaningful evaluation but comes at the cost of high computational demand.

The exploration of materials plays an important role in designing unconventional aircraft concepts. The trade-off between structural weight and drag differs from one material to another, and thus its selection significantly influences the optimal aerostructural design of the wing. Compared to the metallic materials, composite material have the potential to limit the wing mass increase due to aspect ratio extension [32, 33], thanks to their lightweight and strength properties. And more recently, with the ability to manufacture tow-steered composites material, there has been growing interest in exploring their benefits through an aerosrtructural optmization [21, 33].

The objective of this work is to investigate the consideration of each of these aspects on the aerostructural optimal design of a high aspect ratio wing concept.

The remainder of the paper is organized as follows : Section 2 describes the aerostructural analysis and optimization framework used. Section 3 presents the reference HAR wing model. Section 4 outlines the methodology and problem formulations. Section 5 presents the optimization results and discussions.

2 COMPUTATIONAL FRAMEWORK

We use the OpenAeroStruct tool to perform the studies. OpenAeroStruct [34] is an open-source, aerostructural analysis and gradient-based optimization framework. It couples the vortex-lattice aerodynamic model (VLM) and 1-D beam FEA model using six degree-of-freedom (DOF) elements that account for axial, bending, and torsional stiffness. A typical transport-aircraft wingbox model [35] is implemented to compute the cross-sectional properties. The aerodynamic and structural meshes share the same spanwise discretization, ensuring consistent and conservative load displacement transfer, as no interpolation is necessary for the transfer.

The viscous drag is estimated by using skin-fraction-coefficient (Eq.(1)) and form-factor (Eq.(2)) formulas from Raymer [36]:

$$C_f = (C_{f,\text{lam}} - C_{f,\text{turb, trans}}) \times k_{\text{lam}} + C_{f,\text{turb}} \tag{1}$$

where C_f represent the combined skin fraction coefficient, $C_{f,\text{lam}}$ is the laminate skin

friction coefficient,$C_{f,\text{turb}}$ the turbulent skin friction coefficient, $C_{f,\text{trans}}$ the transitional skin friction coefficient and k_{lam} represents the percentage of chord with laminar flow.

$$k_{FF} = 1.34M^{0.18}\left(1 + 0.6\frac{t/c}{c_{\max t}} + 100\left(\frac{t}{c}\right)^4\right) \tag{2}$$

where k_{ff} is the form- factor correction, M the Mach number, t/c the thickness-to-chord ratio, $c_{\max t}$ the chord length where the maximum thickness occurs.

To account the compressibility effects, the wave drag is computed based on the cres-critical Mach number, determined using the Korn equation [37]:

$$M_{crit} = \frac{\kappa_A}{\cos\Lambda} - \frac{t/c}{\cos^2\Lambda} - \frac{C_L}{10\cos^3\Lambda} - (\frac{0.1}{80})^{\frac{1}{3}} \tag{3}$$

where κ_A is a technology factor that we set to κ_A= 0.95, Λ the wing sweep angle, $\frac{t}{c}$ the thickness-to-chord ratio and C_L the lift coefficient of the wing. The wave drag coefficient is then computed using:

$$C_{D,wave} = 20(M_{crit} - M) \tag{4}$$

where M is the flight mach number.

3 BASELINE AIRCRAFT MODEL

We use the undeflected Common Research Model (uCRM-13.5) as the baseline aircraft model for this work [38]. The uCRM-13.5 model is a high aspect ratio variant of the CRM developed to study flexible wing designs in transonic transport aircraft. The model was created through multi-point aerostructural optmization incorporating constraints such as buffet onset to enhance aerodynamic performances. It serves as a widely used benchmark model for aerostructural analysis and design of high-aspect ratio wing configurations. Figure 1 shows the uCRM-13.5 wing model used. The key planform design parameters are listed in Tab.1, while Tab.2 provides the design specifications.

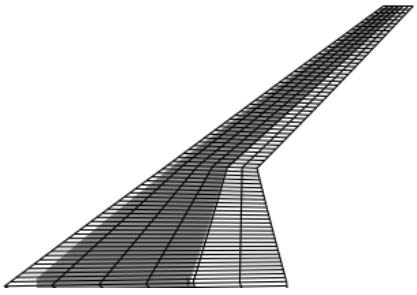


Figure 1: uCRM-13.5 VLM and wingbox model illustration

4 METHODOLOGY & PROBLEM DESCRIPTION

4.1 Models fidelity

To study the impact of the model fidelity choice, we compared the optimum design of five levels of fidelity: 3 three low-fidelity, one mid-fidelity, and one high-fidelity models as

Table 1: uCRM-13.5 wing planform specifications

Parameter	Value	Units
Aspect ratio	13.5	-
Span	72.00	m
Root chord	16.02	m
Side of body chord	13.85	m
Yehudi chord	5.36	m
MAC	5.36	m
Tip chord	1.915	m
Wingpress reference area	384.05	m^2
Gross area	465.22	m^2
Exposed area	377.45	m^2
1/4 chord sweep	35	deg
Taper ratio	0.250	-
Gear post depth	0.648	m

Table 2: Design optimization specifications

Parameter	Value	Units
Initial cruise altitude	37 000	ft
Initial MTOW	297 500	kg
Initial MLW	213 180	kg
Initial MZFW	195 040	kg
Initial OEW	138 100	kg
Design range	7 725	nm
Design payload	34 000	kg
Reserve fuel	15 000	kg
Initial wing weight	30 286	kg
Fixed weight	114 000	kg
TSFC	0.53	lb/(lbf·h)

given in the Tab.3.

Table 3: Modeling fidelities

Name	Aerdoyanamic model	Structural model
LF1	VLM	1-D beam FEA
LF2	VLM + Korn wave drag	1-D beam FEA
LF3	VLM + Viscous drag	1-D beam FEA
MF	VLM + Korn wave drag + Viscous drag	1-D beam FEA
HF	RANS CFD	3D wingbox FEA

The three low-fidelity and the mid-fidelity share the same structural model,with only the aerodynamic model varying. The VLM mesh includes 45 aerodynamic panels across the semi- span and 7 aerodynamic panels along the chord. The structural model's dimension match the spanwise dimension of the VLM mesh. The wingbox model's structural sizing includes adjustable skin and spar thicknesses along the span. The skin thicknesses are the same on the top and bottom surfaces, while the spar thicknesses are identical for the front and back spars. The wingbox extends from 10% chord at the front to 60% chord at the rear, with its shape defined by a NASA SC2-0612 airfoil. For The high fidelity model, we directly use the results from the work of Brooks et al. [38] at the University of Michigan. They performed high-fidelity optimization of the uCRM-13.5 model by coupling the ADflow [39] solver for RANS CFD aerodynamic analysis and TACS [40] solver for the realistic 3D wingbox analysis.

The optimization problem I, listed inTab.4, closely replicates the one used for high-fidelity optimization. The objective is to minimize fuel burn, while satisfying lift-equals-weight and structural failure constraints. The fuel burn is computed at a nominal cruise

lift coefficient CL = 0.5, using the Breguet range equation:

$$\mathrm{FB} = \mathrm{LGW}\left[\exp\left(\frac{RC}{V_\infty(L/D)}\right) - 1\right] \tag{5}$$

where LGW is the aircraft landing weight,R is the range, C is the thrust-specific fuel consumption of the engine, V_∞ is the flight speed and (L/D) is the lift-to-drag ratio at cruise condition.

The structure is sized at 2.5 g maneuver load condition. We allow the optimizer to control wing twist and thickness-to-chord distributions, the wing span, the cruise, and 2.5g maneuver angles of attack, the skin-thickness distribution (same for the upper and lower skins), and the spar-thickness distribution (same for the front and rear spars).

The gradient-based SLQSP optimizer is selected through pyOptSpare framework [41] and a convergence criterion of 10^{-4} is used for both optimality and feasibility.

Table 4: Optimization problem I

Category	Name	Lower	Upper	Units	Quantity
Objective function	fuel burn			kg	1
Design variables	wing twist	-15	15	degrees	4
	thickness-to-chord ratio	0.07	0.02		4
	wing span		100	m	1
	α_{cruise}	-10	10	degrees	1
	$\alpha_{2.5\text{g}}$	-15	15	degrees	1
	spar thickness	0.003	0.1	m	4
	skin thickness	0.003	0.1	m	4
	Total				**20**
Constraints	$\mathrm{C_{L,cruise}}$	0.5	0.5		1
	$\mathrm{lift_{2.5g}}$	$\mathrm{wieght_{2.5g}}$	$\mathrm{wieght_{2.5g}}$	N	1
	$\sigma_{\text{von Mises,2.5g}}$		$\frac{\sigma_{\text{mat}}}{1.5}$	MPa	1
	Total				**3**

The material properties and the flight conditions are listed in the Tab.5 and Tab.6 repectively.

Table 5: Material properties (AL7000)

Paramter	Value	Unit
Density	2780	$\mathrm{kg/m^3}$
Young's Modulus	73100	MPa
Poisson Ratio	0.33	
Yield Stress	420	MPa

Table 6: Flight conditions

Name	Mach	$\mathrm{C_L}$/lift	Altitude
cruise	0.84	0.5	37000 ft
2.5g maneuver	0.64	2.5xMTOW	0

4.2 Objective metrics

To assess the trade-offs between fuel burn (direct operating cost), structural weight (manufacturing cost), and take-off gross weight (acquisition cost), we plot two Pareto

fronts representing the set of all feasible designs. We define two objective functions: a linear combination of fuel burn and zero fuel weight, and a linear combination of fuel burn and Takeoff Gross Weight, respectively :

$$f_1 = \beta \mathrm{FB} + (1-\beta)\mathrm{ZFW} \tag{6}$$

$$f_2 = \beta \mathrm{FB} + (1-\beta)\mathrm{TOGW} \tag{7}$$

where FB is the fuel burn, ZFW is the zero-fuel weight (aircraft structural weight and payload), TOGW is the take-off gross weight and β is the weighting parameter ($0 \leq \beta \leq 1$).

The design points in the Pareto front are obtained by a multi-objective optimizations with different weighting parameters.

The zero fuel weight is computed by

$$\begin{aligned} ZFW = 1.25 \times \mathit{Wing\ Weight} + \mathit{Fixed\ Weight} + \mathit{Payload} \\ + \mathit{Reserve\ fuel} + \mathit{Secondary\ Wing\ Weight} \end{aligned} \tag{8}$$

and the take-off gross weight is computed by

$$TOGW = ZFW + FB \tag{9}$$

where the wing weight is the wingbox weight, the fixed weight is the aircraft structural weight excluding the wing, the secondary wing weight represents the nonstructural system such as controls and high-lift devices and FB is the fuel burn.

Only the objective function is changed in the optimization problem I in section 4.1.

4.3 Analysis method

We compare three optimization methods and their impact on the optimal design :

- **Single-point method:** Uses a single cruise point to compute fuel burn.
- **Multi-point method:** Minimizes the average fuel burn at the cruise conditions [20] listed in Tab.7.

  $$\mathrm{Obj} = \sum_{i=1}^{5} T_i \times \mathrm{FB}_i, \tag{10}$$

  with the weighting of each condition, $T_i = 1/5$.
- **Sequential method**: Uses a modified version of the Breguet range equation to account for the effects of nonzero flight path angles, such as during climb. This approach follows the inverse direction of the method developed by Adler et al. [31], whose work demonstrated that it achieves most of the benefits of mission-based optimization at a lower computational cost. The new objective function is the sum of the fuel burn from the climb and cruise segments.

  $$FB = W_f \left\{ \exp\left[\frac{R \cdot C}{V} \left(\frac{1}{L/D} + \gamma \right) \right] - 1 \right\} \tag{11}$$

  where W_f is the weight at the end of the phase and γ is the flight path angle.

For each method, several optimizations are performed for different design ranges, varying from 300 nm to 2500 nm.

Table 7: uCRM-13.5 flight condition stencil.

Point	**Condition**	**Mach**	C_L	**Altitude (ft)**
1	Cruise	0.85	0.549	37000
2	Cruise	0.85	0.524	37000
3	Cruise	0.85	0.574	37000
4	Cruise	0.84	0.562	37000
5	Cruise	0.86	0.536	37000

4.4 Material exploration

To explore the impact of material usage in the aerostructural optimization of HAR wings, we compare the metallic material AL7000 and the composite material CFRP. For both materials, fuel burn and structural weight optimizations are performed, and for each design, the optimal metrics and CO_2 impact are compared.

Regarding CO_2 impact, we consider emissions from fuel consumption on one hand and material manufacturing on the other.

The CO_2 impact due to fuel consumption for completing the design mission is computed as:

$$CO_{2,fuel} = 3.155 \times FB \tag{12}$$

where FB is the fuel burn weight. To determine the contribution per passenger, we divide the CO_2 weight by the number of passenger.

The CO_2 emitted by the structure is computed as the product of the mass of the wing (W_{wing}) by the CO_2 footprint ($CO_{2,mat}$) of the material used.

$$CO_{2,structure} = CO_{2,mat} \times W_{wing} \tag{13}$$

The Table 8 lists the material properties.

Table 8: uCRM-13.5 flight condition stencil.

Material	AL7000	CFRP
Young's Modulus (GPa)	73.1	150
Shear Modulus (GPa)	27.48	5.5
Density (kg/m^3)	2780	1600
Allowable Yield stress (MPa)	420	1200
C0$_{2,\mathrm{mat}}$ (kg CO$_2$/kg)	8.66	48.1

5 RESULTS AND DISCUSSIONS

5.1 Models fidelity

The optimized performance metrics resulting from the optimization with the five levels of fidelity, described in Section 4.1 are shown in the Table 9.

Neglecting the viscous and compressibility effects, which account for 28% and 1% of the total drag of the optimum design lead to an unrealistic design. Indeed, this drag underestimation corresponds to an overestimation of the aerodynamic efficiency by approximately 62%, resulting in an underestimation of fuel burn and wing mass by 45% and 14%, respectively. The mid-fidelity optimization, which combines a vortex lattice method with statistical estimations of the viscous and compressibility effects of the wing, provides acceptable accuracy compared to the high-fidelity optimization, which integrates a RANS CFD model with a 3D wingbox FEM model. The high-fidelity optimized aerodynamic efficiency is 0.56% greater than that of the mid-fidelity optimum, resulting in an error of -0.6% in fuel burn and 5% in wing mass. The computational cost of the mid-fidelity optimization is less than 1 hour, compared to approximately 48 hours for the high-fidelity optimization.

Table 9: Optimized performance metrics for the different levels of fidelity

Level of fidelity	Fuel burn		Wing weight		C_D		L/D	
	kg	Δ (%)	kg	Δ (%)	counts	Δ (%)	-	Δ (%)
LF1	49993	-44.14	24472	-14.55	135.52	-38.25	36.89	61.94
LF2	83704	-6.48	19674	-31.30	217.72	-0.79	22.96	0.80
LF3	63575	-28.97	35110	22.58	158.20	-27.92	31.60	38.72
MF	89500	Ref.	28641	Ref.	219.46	Ref.	22.78	Ref.
HF [20]	84072	-6.07	30032	4.85	-	-	22.91	0.56

Figure 2 compares the optimized distributions obtained using the mid-fidelity and high fidelity. The cruise and 2.5 g maneuver lift distribution from the wing-fuselage junction to the tip are more or less the same. However, the overall distribution trends of the twit and the thickness-to-chord ratio differ, although their behavior appears similar in the outboard section of the wing.

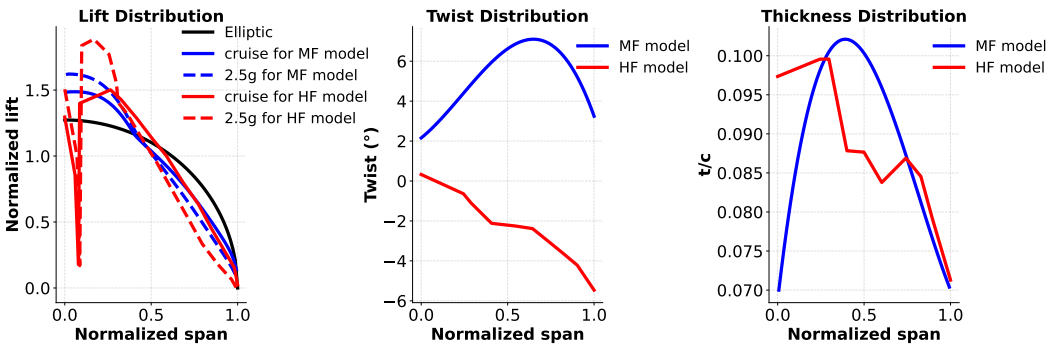


Figure 2: Optimized distributions of the MF and HF models

5.2 Objective metrics

Performing optimization for different weightings in the combined fuel burn and structural weight metric, on one hand, and in the fuel burn and takeoff gross weight metric, on the other hand, resulted in the Pareto fronts shown in Figures 3 and 4 respectively. The choice of the objective function has a significant impact on nearly every aspect of the optimized design.

As expected, the fuel burn for the fuel burn-focused design ($\beta = 1$)is lower than that of the TOGW-focused design, which in turn is lower than the ZFW-focused design. This is because fuel burning is a component of TOGW, which influences the overall optimization outcome. The non-linearity observed at the extremities of the Pareto fronts indicates that, in these regions, small improvement in one objective metric leads to a significant penalty in the other. For the fuel burn and zero fuel weight combination, we observe that comparing the designs with $\beta = 0.7$ and $\beta = 1$ results in a fuel burn reduction of only 0.3% accompanied by 6.08% increase in structural weight. From the design with beta=0.1 to the design with beta=0.0, an improvement of 0.6% in structural weight comes at cost of a substantial 21.28% increase the fuel burn. The same observations apply to the fuel burn and takeoff gross weight combination. From the design with $\beta = 0.1$ to design $\beta = 0$, a slight decrease of only 0.17% in takeoff gross results in a 2.26% increase in fuel burn. Thus, it is clear that the optimal design lies between the extremes, representing the best trade-off between Direct operating costs, manufacturing costs, and acquisition costs.

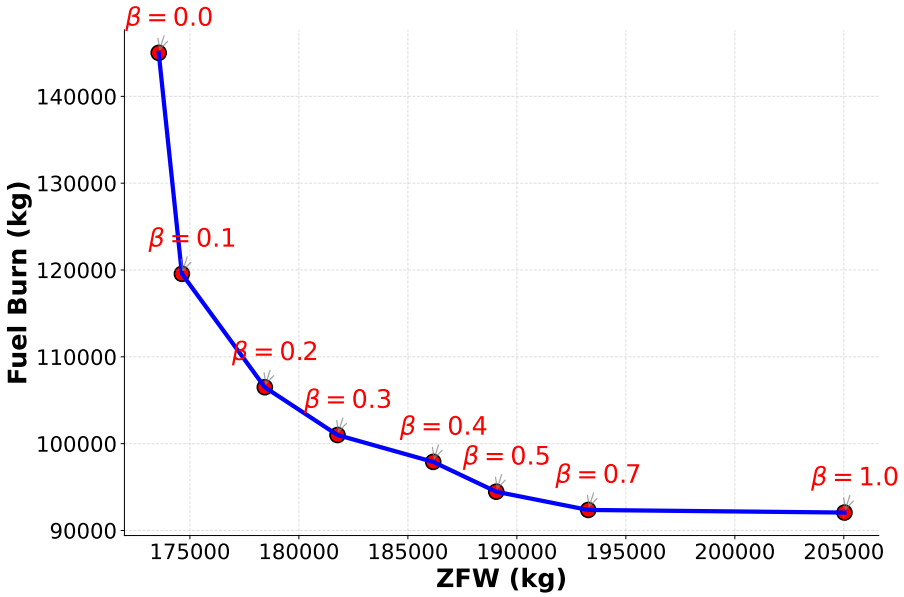


Figure 3: Pareto front for fuel burn and structural weight

Figures 5 and 6 show the relative trends for fuel burn, gross take-off weight (MTOW), structural wing weight, and L/D metrics compared to fuelburn-focused design. As β increase, the optimal L/D increases because of the growing importance of aerodynamic performance for the fuel-burn optimized design. Trade-off with TOGW offers better aerodynamic efficiency and, consequently, lower fuel burn compared to the trade-off with structural weight.

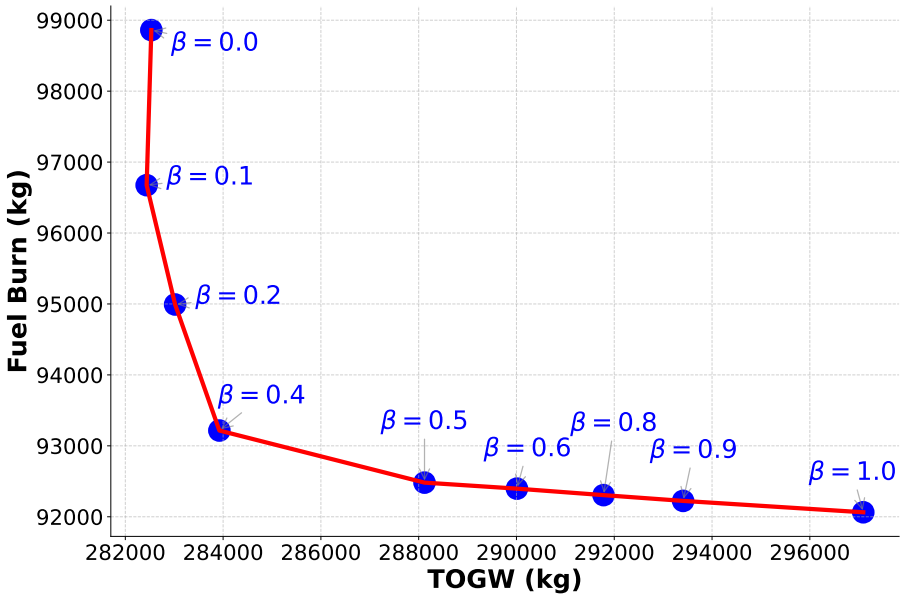


Figure 4: Pareto front for fuel burn and TOGW

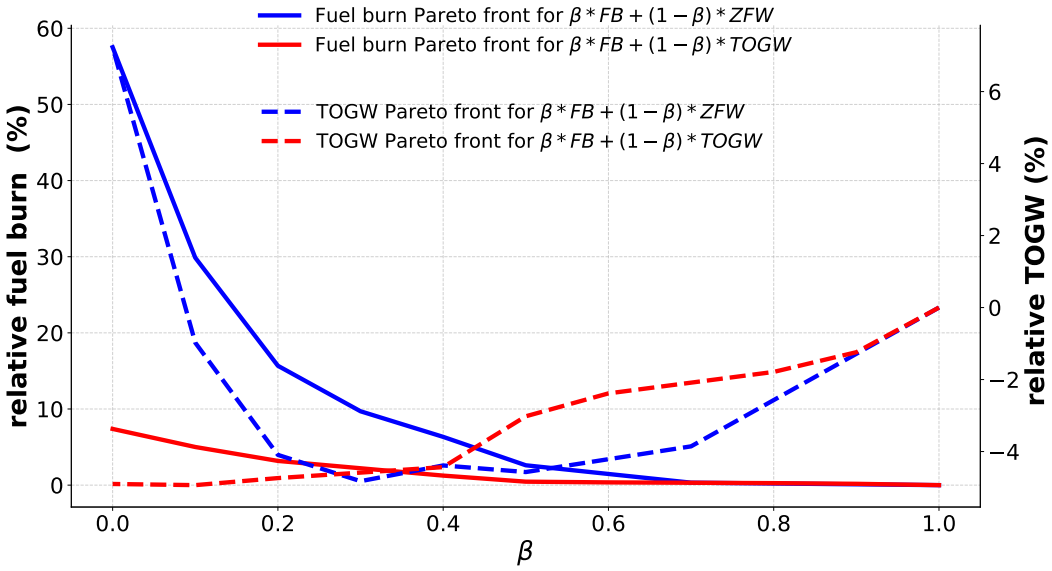


Figure 5: fuel burn and TOGW Pareto front trendlines for varying β

5.3 Analysis method

The results obtained using the three optimization methods for different ranges, from 300 to 2500 nm, is shown in Figures 7.It becomes clear that, while traditional methods estimate fuel burn reasonably well for long ranges, they become unsuitable for short missions. . In fact, for short ranges, the fuel burn from the climb phase represents up to 50% of the total fuel burn of the optimal design. To achieve a realistic design, it is essential to consider the climb phase and, more generally, all other segments of the mission. The most significant difference in actual fuel burn and wing weight between the single point and multi-point optimized design is only 1.7% and 11.9% respectively. The computational cost of the multi-point method is 200 times that of the single point method. Compared to the multipoint method, the sequential method offer better aerodynamic efficiency and

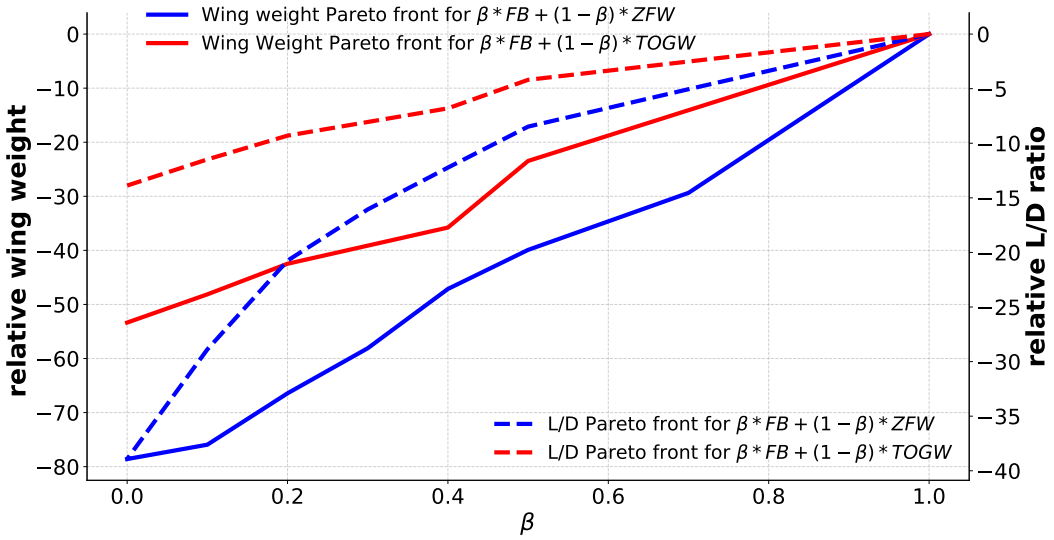


Figure 6: wing weight and L/D ratio Pareto front trendlines for varying β

a lighter wing at lower computational cost.

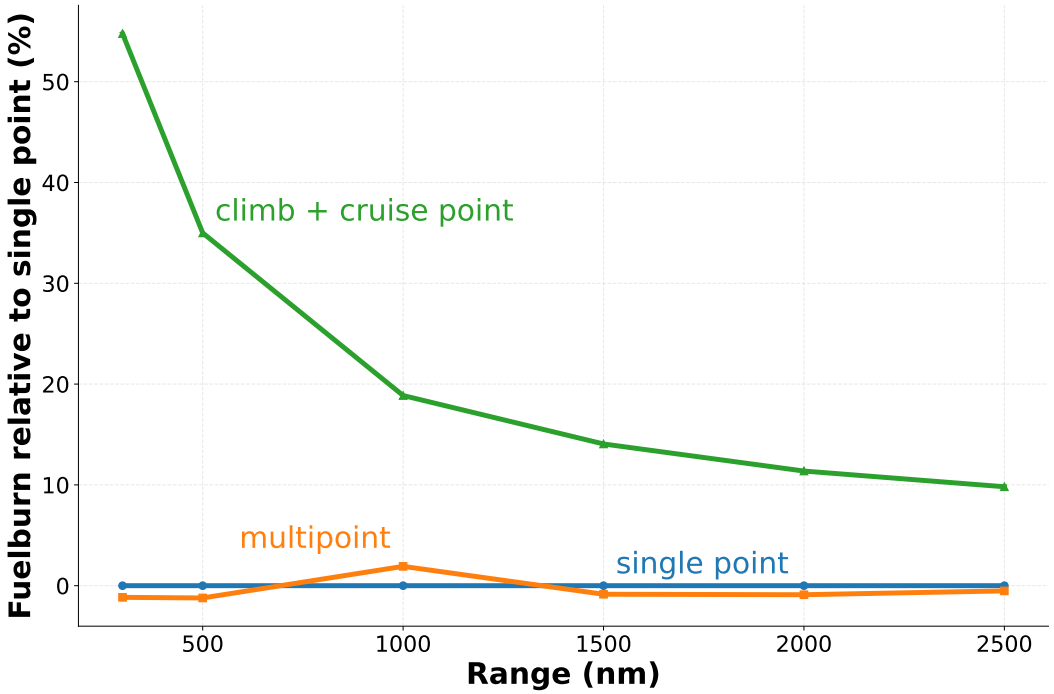


Figure 7: Fuel burn error compared to single point method

5.4 Material exploration

Table 10 presents the optimized data for fuel burn and zero fuel weight (ZFW) optimizations for both metallic and composite materials. It appears clearly that the use of advanced material outperforms the aluminum material, leading to simultaneous fuel burn and structural weight reduction. For the minimum fuel burn design, switching from aluminum to composite material reduces fuel burn and CO2 emissions per passenger from

fuel consumption by 18% and the structural weight by 13%. For the minimum structural weight design, the switch from aluminum to composite material reduces the structural weight by 2.24% but increases the CO2 emissions from material production by 203%.

Table 10: Optimized metrics comparison for metallic and composite materials

Performance metrics	Fuel burn design		ZFW design	
	AL700 (Ref)	CFRP (Δ%)	AL700 (Ref)	CFRP (Δ%)
Fuelburn (kg)	92063	75349(-18.2%)	145010	99013 (-31.72%)
Zero Fuel Weight (kg)	205029	178293 (-13.04%)	173568	169682(-2.24%)
Wing weight (kg)	32023	10634 (-66.79%)	6854	3745.83(-45.34%)
L/D	22.64	23.82 (+5.21%)	13.82	18.27 (+32.20%)
CO_2 from fuel (g $CO_2/PAX/NM$)	85.45	69.94 (-18.15%)	134.60	91.90 (-31.72%)
CO_2 from structure (tons CO_2)	277.32	511.50 (+84.44%)	59.36	180.55 (+203%)

6 CONCLUSIONS

In this paper, we investigate the complex aerostructural trade-offs of high-aspect ratio (HAR) wings. The research highlights how each design aspect—model fidelity, objective metrics, analysis methods, and material selection— significantly influences the optimal design.

While low-fidelity models provide rapid solutions, mid-fidelity and high-fidelity methods yield more realistic design and become essential when addressing transonic wing configurations.Mid-fidelity mod- els, which combine a vortex lattice method-based aerodynamic model with estimations of viscous and compressibility drag, and a 1D beam finite element structural model, achieve results within 6% for fuel burn and 4.85% for wing weight compared to high fidelity models, while reducing computational time by over 97%.

The choice of objective metric plays a critical role and optimum trade-off between the direct operating cost and the acquisition cost should be carefully selected based on the mission and operational goals of the aircraft.

The study also underscores the importance of selecting an appropriate analysis method. While the single and multi-point methods are suitable for long-range missions, they underestimate the fuel burn by up to 50% for short-range missions. As demonstrated in this study, mission-based or enhanced approaches – such as the proposed sequential method – are essential for more accurate performance prediction in those scenarios.

Furthermore, the material plays a key role in aerostructural optimization. The transition from aluminum to composites materials results in notable reductions in fuel burn and structural weight, although it comes with increased CO2 emissions from material manufacturing. This highlights a promising direction encouraging the exploration of more advanced materials such as tow-steered composites for future aircraft structures.

ACKNOWLEDGEMENTS

The authors would like to especially thank Yoan Le-Lamer from ISAE-SUPAERO for his valuable support, particularly in cross-checking and validating the results presented in this work.

REFERENCES

[1] Boeing. Commercial market outlook 2021-2040. Technical report, Boeing, 2021.

[2] A. Sgueglia. *Methodology for sizing and optimising a Blended Wing-Body with distributed electric ducted fans*. PhD thesis, Doctorat Université de Toulouse, 2019.

[3] E. Greitzer and al. N+3 aircraft concept designs and trade studies, final report. Technical report, NASA Glenn Research Center, 2010.

[4] E. Commission. Flightpath 2050: Europe's vision for aviation. Technical report, Advisory Council for Aeronautics Research in Europe, 2011.

[5] Marchetti, F. Toffol, S. Ricci, J. Beretta, and N. Paletta. Aeroelastic optimization of high aspect ratio wings for environmentally friendly aircraft. *AIAA SCITECH 2022 Forum*, 2022. doi:https://doi.org/10.2514/6.2022-0166.

[6] Toffol and S. F. Ricci. Preliminary aero-elastic optimization of a twin-aisle long-haul aircraft with increased aspect ratio. *Aerospace 2023*, 2023. doi:https://doi.org/10.3390/aerospace10040374.

[7] G. Carrier and al. Multidisciplinary analysis and design of strut-braced wing concept for medium range aircraft. *AIAA SCITECH 2022 Forum*, 2022. doi:http://arc.aiaa.org/ DOI: 10.2514/6.2022-0726.

[8] Y. L. Lamer, J. Morlier, E. Benard, and P. He. Aeroelastic analysis of high aspect ratio and strut-braced wings. *33th Congress of the International Council of the Aeronautical Sciences*, 2022. doi:ffhal-03888120f.

[9] W. Pfenninger. Passive aeroelastic control in truss-braced wings using vibration suppression. *UNIVERSITY OF BRISTOL*, 2019.

[10] V. Y.V., V. Chepiga, M. V.P., K. E.A., and G. V.G. Configuration of the medium-haul twin-fuselage passenger aircraft. *International Journal of Applied Engineering Research 12*, 12:414–421, 01 2017.

[11] H. Deconinck, K. Sermeus, and C. Dam. *CFD-based aircraft drag prediction and reduction*. von Karman Institute for Fluid Dynamics, 02 2003.

[12] M. PT. *Viscous phenomena affecting high-lift systems and suggestions for future CFD development*. AGARD TR-94- 18415- 04–01, 1993.

[13] Haftka and R. T. Optimization of flexible wing structures subject to strength and induced drag constraints. *AIAA Journal*, 15(8):1101–1106, 1977. doi:10.2514/3.7400.

[14] P. W. Jansen, R. E. Perez, and J. R. R. A. Martins. Aerostructural optimization of nonplanar lifting surfaces. *Journal of Aircraft*, 47(5):1490–1503, 2010. doi:10.2514/1.44727.

[15] G. Kennedy and J. Martins. Parallel solution methods for aerostructural analysis and design optimization. *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference 2010*, 09 2010. doi:10.2514/6.2010-9308.

[16] G. E. Fujiwara and N. T. Nguyen. Aeroestructural design optimization of a subsonic wing with continuous morphing trailing edge. *35th AIAA Applied Aerodynamics Conference*, 2017. doi:10.2514/6.2017-4218.

[17] A. Elham and M. Van tooren. Coupled adjoint aerostructural wing optimization using quasi-three-dimensional aerodynamic analysis. *Structural and Multidisciplinary Optimization*, 54:889–906, 10 2016. doi:10.1007/s00158-016-1447-9.

[18] J. Martins, J. Alonso, and J. Reuther. High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41:523–530, 05 2004. doi:10.2514/1.11478.

[19] S. Khosravi and D. W. Zingg. Aerostructural optimization of drooped wings. *Journal of Aircraft*, 55:1261–1268, 2017.

[20] T. Brooks, G. Kenway, and J. Martins. Benchmark aerostructural models for the study of transonic aircraft wings. *AIAA Journal*, 56:2840–2855, 06 2018. doi:10.2514/1.J056603.

[21] T. Brooks, J. Martins, and G. Kennedy. Aerostructural tradeoffs for tow-steered composite wings. *Journal of Aircraft*, 57:1–13, 04 2020. doi:10.2514/1.C035699.

[22] A. Christison Gray and J. Martins. A proposed benchmark model for practical aeroelastic optimization of aircraft wings. *AIAA SciTech Forum*, 01 2024. doi:10.2514/6.2024-2775.

[23] G. K. W. Kenway, G. J. Kennedy, and J. R. R. A. Martins. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA Journal*, 52(5):935–951, 2014. doi:10.2514/1.J052255.

[24] N. Beck, T. Landa, A. Seitz, and L. Boermans. Aerostructural design optimization of the d8 aircraft using active aeroelastic tailoring. *AIAA Scitech 2020 Forum*, 2020. doi:https://doi.org/10.2514/6.2020-1967.

[25] G. Kennedy and J. Martins. A parallel aerostructural optimization framework for aircraft design studies. *Structural and Multidisciplinary Optimization*, 50, 12 2014. doi:10.1007/s00158-014-1108-9.

[26] G. Kenway, G. Kennedy, and J. R. R. A. Martins. Aerostructural optimization of the common research model configuration. *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2014. doi:10.2514/6.2014-3274.

[27] G. Kenway and J. Martins. Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft 51*, 2014. doi:https://doi.org/10.2514/1.C032150.

[28] Z. Lyu, G. K. W. Kenway, and J. R. R. A. Martins. Aerodynamic shape optimization investigations of the common research model wing benchmark. *AIAA Journal*, 53 (4):968–985, 2015. doi:10.2514/1.J053318.

[29] G. K. W. Kenway and J. R. R. A. Martins. Multipoint aerodynamic shape optimization investigations of the common research model wing. *AIAA Journal*, 54(1): 113–128, 2016. doi:10.2514/1.J054154.

[30] M. Clarke, R. Erhard, J. Smart, and J. Alonso. Aerodynamic optimization of wing-mounted propeller configurations for distributed electric propulsion architectures. *AIAA Aviation 2021 Forum*, 08 2021. doi:10.2514/6.2021-2471.

[31] E. J. Adler and J. R. R. A. Martins. Efficient aerostructural wing optimization considering mission analysis. *Journal of Aircraft*, 60(3):800–816, 2023. doi:10.2514/1.C037096.

[32] G. Kennedy and J. Martins. A comparison of metallic and composite aircraft wings using aerostructural design optimization. *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSM 17-19*, 09 2012. doi:10.2514/6.2012-5475.

[33] J. R. R. A. Martins, G. Kennedy, and G. K. Kenway. High aspect ratio wing design: Optimal aerostructural tradeoffs for the next generation of materials. *52nd Aerospace Sciences Meeting*, 2014. doi:10.2514/6.2014-0596.

[34] J. Jasa, J. Hwang, and J. Martins. Open-source coupled aerostructural optimization using python. *Structural and Multidisciplinary Optimization*, 57, 02 2018. doi:10.1007/s00158-018-1912-8.

[35] S. Chauhan and J. Martins. Low-fidelity aerostructural optimization of aircraft wings with a simplified wingbox model using openaerostruct. *6th International Conference on Engineering Optimization. EngOpt 2018*, 09 2018. doi:10.1007/978-3-319-97773-7_38.

[36] D. P. Raymer. *Aircraft Design: A Conceptual Approach*. AIAA, 5th ed, 2012.

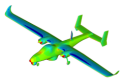
[37] B. Malone and W. H. Mason. Multidisciplinary optimization in aircraft design using analytic technology models. *Journal of Aircraft*, 32(2):431–438, 1995. doi:10.2514/3.46734.

[38] T. Brooks. Undeflected common research model (ucrm): An aerostructural model for the study of high aspect ratio transport aircraft wings. *35th AIAA Applied Aerodynamics Conference*, 2017. doi:https://doi.org/10.2514/6.2017-4456.

[39] P. Lyu, G. Kenway, C. Paige, and J. Martins. Automatic differentiation adjoint of the reynolds-averaged navier-stokes equations with a turbulence model. *21st AIAA Computational Fluid Dynamics Conference*, 06 2013. doi:10.2514/6.2013-2581.

[40] G. J. Kennedy and J. R. Martins. A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elements in Analysis and Design*, 87:56–73, 2014. ISSN 0168-874X. doi:https://doi.org/10.1016/j.finel.2014.04.011.

[41] E. Wu, G. Kenway, C. A. Mader, J. Jasa, and J. R. R. A. Martins. pyoptsparse: A python framework for large-scale constrained nonlinear optimization of sparse systems. *Journal of Open Source Software*, 5(54):2564, 2020. doi:10.21105/joss.02564.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

EFFICIENT GRADIENT-BASED STRUCTURAL OPTIMISATION WITH MODAL AND BUCKLING CONSTRAINTS ASSISTED BY ALGORITHMIC DIFFERENTIATION

Luca Scalia[1], Rauno Cavallaro[1], Andrea Cini[1], Nicolas R. Gauger[2], Max Sagebaum[2]

1: Department of Aerospace Engineering
Universidad Carlos III de Madrid (UC3M)
lscalia@pa.uc3m.es
rcavalla@ing.uc3m.es
acini@ing.uc3m.es

2: Chair for Scientific Computing
University of Kaiserslautern-Landau (RPTU)
nicolas.gauger@scicomp.uni-kl.de
max.sagebaum@scicomp.uni-kl.de

Abstract. *The pursuit of sustainable aviation has led to the exploration of unconventional aircraft configurations, such as high-aspect-ratio wings, blended-wing bodies, and strut-braced wings. Performance assessment of novel layouts cannot rely solely on conventional design approaches built upon legacy experience. Instead, high-fidelity analysis tools are required from the early design stages.*

Building upon a previously developed framework for efficient gradient-based aero-structural optimisation enhanced by algorithmic differentiation, this work extends the structural module to enable modal and buckling analyses through the solution of generalised eigenvalue problems, along with the development of their associated sensitivities. Designed for industrially relevant applications, the framework is implemented with parallel computing capabilities to ensure scalability and efficiency.

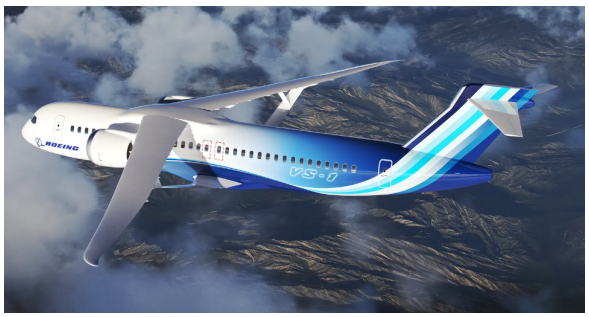
The suggested tool capability is demonstrated on two structural optimisation testcases, involving, enforcing isolated and combined, stress, frequency, and buckling constraints. Results highlight the critical role of simultaneously imposing multiple constraint types in ensuring robust and optimal wing designs.

Keywords: Structural optimisation, algorithmic differentiation, finite element method, gradient-based optimisation, adjoint method, buckling constraint

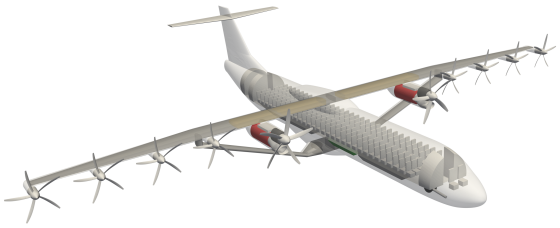
1 INTRODUCTION

The aerospace industry is increasingly focused on achieving sustainable aviation by reducing its environmental impact. Despite the challenges posed by the supply chain issues, this commitment remains strong. Recent analyses by the International Civil Aviation Organization (ICAO) revealed that, by the end of 2023, air traffic had recovered to—or even exceeded—pre-pandemic levels on most routes [1]. Additionally, the 2022 European Aviation Environmental Report [2], issued by the European Aviation Safety Agency (EASA), projects that the number of flights at European Union and EFTA airports will reach 12.2 million annually by 2050. Without the implementation of environmentally focused countermeasures, CO_2 emissions could rise to 188 million tonnes.

In this context, improving aerodynamic and structural efficiency has emerged as a key strategy to reduce emissions and fuel consumption. One way of pursuing this is studying novel airframe concepts such as Large Aspect-Ratio Wings (LARW), Strut-Braced and Truss-braced Wings (SBW and TBW) (Fig.1) [3–7]. These advancements necessitate appropriate design and optimisation strategies to ensure both performance and safety, while simultaneously meeting sustainable aviation objectives. The structural optimisation pro-



(a) Boeing's Truss-braced Wing concept [4].



(b) Strut-braced wing and distributed electric propulsion concept studied in INDIGO's project [5].

Figure 1: Two examples of novel aircraft concepts.

cess seeks to enhance structural efficiency while satisfying constraints related to stiffness, integrity, and manufacturability. In the case of novel configurations, new challenges may emerge that cannot be addressed through traditional design experience; therefore, higher-fidelity digital models are required from the outset. An inadequately managed design and optimisation process can lead to inefficient redesign cycles, longer time-to-market, and suboptimal performance.

Several competing requirements characterise the structural optimisation process applied to novel designs. On one hand, high-fidelity analyses are essential to ensure accuracy, but they are computationally demanding. On the other hand, it is crucial not to overly restrict the design space a priori and to maintain a reasonably high number of design variables to fully exploit novel configurations. This combination makes optimisation problems particularly challenging to solve. A viable strategy to mitigate this issue is the use of gradient-based optimisation, supported by adjoint methods [8–11] to efficiently compute sensitivities at an acceptable computational cost. Among the various techniques available for computing sensitivities, algorithmic differentiation [12] has proven particularly effective and is the method adopted in this study.

1.1 State of the art

The adjoint method is a well-established technique in structural optimisation. One of the first applications in aerospace design is from Bartholomew [13, 14] who introduced the code *STAR*, applied to both civil and military aircraft, highlighting the crucial role of Computer-Aided Design (CAD) in the design process.

Soares et al. [15], focused on the optimisation of composite plates, underscoring the increasing relevance of these materials in the aerospace sector.

Subsequent decades have seen a marked shift towards developing high-fidelity models characterised by significantly larger numbers of degrees of freedom (DOFs) and design variables, driven by the need for greater accuracy and robustness in the design process. Akgun's work [8] exemplifies this trend, comparing the efficiency of the adjoint and direct methods for gradient computation across both small (533 DOFs) and large (at least 43,000 DOFs) High-Speed Civil Transport (HSCT) models. Results demonstrated the superiority of the adjoint method as the problem size increases, establishing its position as the preferred approach for large-scale problems where the computational cost of direct differentiation becomes prohibitive. Moreover, the study underscored the need of aggregation techniques to maintain the computational efficiency of the adjoint method when handling large numbers of optimisation constraints.

The importance of constraint aggregation is further addressed by Lambe [16], who compares various aggregation techniques in the context of wing-box structural optimisation. This work highlights the dependency of optimisation results on parameters such as the number of aggregation areas, which can be reduced from 680 to as few as 17 through appropriate aggregation strategies. Such reductions are essential for preserving the computational efficiency of the adjoint method, especially when multiple constraints are simultaneously enforced. Notably, Kennedy [17] introduced TACS (Toolkit for the Analysis of Composite Structures), an integrated parallel FE analysis tool specifically designed for gradient-based optimisation, where sensitivities are computed via hand-coded routines. TACS exemplifies the broader trend towards developing specialised software frameworks to efficiently address the increasing complexity of structural optimisation problems

More recent efforts have increasingly embraced Multidisciplinary Design Optimisation (MDO) to enhance the robustness and feasibility of designs by considering multiple disciplines concurrently, particularly structural, aerodynamic, and aeroelastic performance. This shift reflects a growing recognition that optimising each discipline independently often leads to suboptimal solutions when the interactions between them are significant. Adler [18], for instance, demonstrates the benefits of coupling mission analysis with aerostructural optimisation, using a detailed fuel burn objective function derived from numerical integration of the fuel flow across the mission profile. This approach, employing a vortex lattice aerodynamic model coupled with a 1D FE structural model, illustrates the effectiveness of leveraging simplified aerodynamic models while maintaining a high-fidelity structural representation.

The importance of considering complex aeroelastic phenomena is further demonstrated by Jonsson [19], who tackles the challenge of non-linear flutter constraints in conjunction with stress constraints. The study features a reduced-order model, condensing a high-fidelity FE model to a beam representation to facilitate the computation of aeroelastic sensitivities. This approach highlights the ongoing trend towards developing efficient and accurate methodologies for handling coupled aerostructural and aeroelastic sensitivities

through the adjoint method, even when the underlying physical phenomena are inherently nonlinear.

The solution of eigenvalue problems and the associated sensitivity of eigenpairs is required for aeroelastic analyses as well as structural modal and buckling problems. Cardani [20] proposes a method simultaneously determining the sensitivity of both the eigenvalue and the eigenvector. This approach is based on solving a linear system for each eigenpair, derived from the direct hand-differentiation of the eigenvalue problem. Mateus [21] examines the sensitivity of buckling loads in thin, multi-layered, angled-ply composite structures by employing a combination of adjoint and analytical methods. The issue of sensitivity analysis in case of multiple eigenvalues and their corresponding eigenvectors is addressed by Lee [22]. In He [23], the authors present an application of based on analytical formulas for AD reverse mode. The method is extended to compute derivatives of eigenpairs for complex matrices in [24]. Schwalbach [25] integrates the AD library *CoDiPack*[26] into a FE code to compute eigenvalue sensitivities. A tangent and an adjoint approach, relying respectively on AD forward and reverse modes, are compared. The work of Torrigiani [27] focuses on implementing an aeroelastic stability module featuring both wing planform and structural design variables. The author applies the method described in Cardani [20] to obtain the modal shape derivatives, which are essential when wing planform design variables are integrated into the optimisation.

In pursuit of developing accurate and efficient optimisation frameworks, AD has emerged as a critical tool, enabling the efficient computation of gradients for increasingly complex problems. Its application, initially limited to relatively simple models involving trusses or plates with tens or hundreds of design variables [28–30], has progressively expanded to more sophisticated scenarios. For instance, Tadjouddine [31] demonstrated the effectiveness of AD in structural dynamic optimisation of a satellite boom, highlighting its clear runtime advantages over traditional Finite Differences (FD) approaches. Recent applications exist also in the field of structural shape [32] and topology optimisation [33], although they do not relate to aerospace applications.

In Bombardieri [34], AD is leveraged for computing coupled sensitivities within a gradient-based aero-structural optimisation framework for wings. Building on the work of Albring et al. [35, 36], the approach employs a fixed-point iteration strategy for sensitivity computation. The structural analysis is handled by the in-house solver *pyBeam* for beam nonlinear FE analysis which is differentiated with the AD library *CoDiPack*. The fixed-point formulation facilitates the decoupling of the framework modules. Information exchange between solvers is managed via a high-level *python* wrapper, with the CFD analysis conducted by *SU2*. Limitations to a wing stick model and the absence of structural code parallel computing capability restrict the number of available design variables preventing analysis of higher-fidelity industrial models. This work aims at overcoming such limitations by presenting *AUGUSTO++*, a tool for gradient-based structural optimisation of highly flexible aeronautical structures, featuring up to millions of degrees of freedom and thousands of design variables. The implementation is supported by AD to obtain discrete-consistent gradients and by parallel computing techniques enhancing computational performance in both the primal and adjoint phases. The design philosophy favours a flexible and modular code structure that facilitates modifications, such as replacing an external library for solving linear systems or eigenvalue problems, as well as incorporating new design variables and responses.

1.2 Objective and paper organisation

This work builds upon the research by Scalia et al.[37], which introduced *AUGUSTO++* and demonstrated its application to structural optimisation under stress constraints. The ultimate goal is to develop a gradient-based, coupled aero-structural optimisation framework capable of handling aeroelastic constraints, such as flutter, at industrial scales. Few features are missing in [34, 37], to fulfil that goal, and their implementation is this paper subject.

The main objective is to extend the capabilities of the aforementioned framework to support modal and linearised buckling analyses, both of which are particularly relevant in the aeronautical field. In addition, the framework is enhanced to provide gradients of responses related to natural vibration frequencies and critical loads. This is achieved entirely through the use of AD, combined with an MPI-based implementation to efficiently handle large-scale problems.

The paper is organised as follows. Section 2 provides a technical background overview, including the formulation of static, modal, and linear buckling analyses, as well as the methodology for computing the derivatives of the different structural responses. Section 3 describes the finite element (FE) solver and its implementation. Section 4 presents the results of mass minimisation for a flat plate and a wing model, considering constraints on stress, frequency, and critical buckling load, both individually and in combination. Finally, Section 5 summarises the main conclusions and outlines directions for future research.

2 TECHNICAL BACKGROUND

This section outlines the technical background underpinning the formulation of the structural optimisation problem. It presents the problem formulation alongside the Finite Element (FE) governing equations for static, modal, and buckling analyses. Additionally, an AD outline is provided, including the workaround employed to differentiate code solving a linear system (LS) or a generalised eigenvalue problem (GEP). The discussion then turns to the response sensitivity computation procedure, which integrates the adjoint method and AD.

2.1 Structural optimisation problem

The structural optimisation problem is formulated as follows:

$$\begin{aligned} \text{Minimise} \quad & m \\ \text{w.r.t.} \quad & \boldsymbol{\alpha}_s \\ \text{subject to} \quad & c_i \leq 0 \qquad i = 1, ..., n_c \end{aligned} \tag{1}$$

The objective is to minimise the structural mass m by varying the set of selected design variables $\boldsymbol{\alpha}_s$, while ensuring compliance with a defined set of optimisation constraints. The constraint responses c_i pertain to structural stress, modal frequencies, and critical load, with their formulations introduced in Sec.(2.2.1), Sec.(2.2.3), and Sec.(2.2.3). The optimisation framework relies on FE analyses for static, modal, and buckling problems, as detailed in the following subsections.

Gradient-based constrained optimisation algorithms require the computation of the gradients of both the objective function m and the constraint responses c_i with respect to

$\boldsymbol{\alpha}_s$. The methodologies employed to compute these derivatives are discussed in Sec.(2.4).

2.2 The primal problem

The *primal problem* is defined as the system of equations and boundary conditions whose solution yields a specific output of interest. Depending on the type of analysis, this output may correspond to structural displacements due to an applied load, modal frequencies, or the buckling critical load along with the associated eigenvectors. Further postprocessing of these results can provide additional structural responses, such as the stress within a single finite element or the aggregated stress over a group of elements.

2.2.1 Static analysis

The structural static FE formulation is developed under the assumption of large displacements and small strains:

$$\boldsymbol{\mathcal{S}}(\boldsymbol{u}_s, \boldsymbol{f}_s) = \boldsymbol{f}_s - \boldsymbol{f}_{int}(\boldsymbol{u}_s) = \boldsymbol{0} \tag{2}$$

where $\boldsymbol{\mathcal{S}}$ represents the nodal residual vector, while $\boldsymbol{f}_{int}$, $\boldsymbol{f}_s$, and $\boldsymbol{u}_s$ correspond to the internal forces, applied external loads, and generalised displacements vectors, respectively. However, in this work, the analysis is confined to the linear elastic case, reducing the equilibrium equations to:

$$\boldsymbol{f}_s = \boldsymbol{f}_{int} = \boldsymbol{K}_{el}\boldsymbol{u}_s \tag{3}$$

where the linear elastic component of the global stiffness matrix, $\boldsymbol{K_{el}}$, depends solely on the structural design variables $\boldsymbol{\alpha_s}$. Subsequent post-processing of the displacements $\boldsymbol{u}_s$, obtained by solving the LS in Eq.(3), provides the *Von Mises* stress (σ_{VM}). For brevity, the subscript VM is omitted hereafter.

Stress constraints are defined as:

$$g_j = \frac{\sigma_j}{\sigma_{ADM}} - 1 \leq 0 \qquad j = 1, ..., n \tag{4}$$

where σ_j represents the stress in each of the n elements for which compliance is evaluated and σ_{ADM} is the material allowable stress . To reduce the design response number, constraint aggregation techniques are commonly employed in aerospace design and optimisation. One widely used approach is the *Kreisselmeier–Steinhauser* (KS) function [38], which consolidates stress constraints from multiple elements into a single expression:

$$J_\sigma = KS(g_j) = \frac{1}{\rho_{KS}} log \sum_{j=1}^{n_{KS}} e^{\rho_{KS}\, g_j} \tag{5}$$

where ρ_{KS} is an aggregation parameter, acting similarly to a penalty factor in constrained optimisation, and n_{KS} represents the number of aggregated elements. Structural integrity against a maximum stress is expressed by:

$$J_\sigma \leq 0 \tag{6}$$

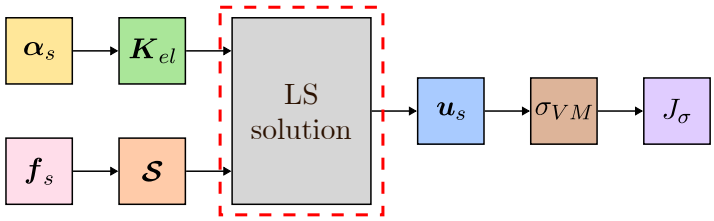


Figure 2: Computational graph of the static analysis.

Figure 2 shows the computational path for evaluating the aggregated stress response.

2.2.2 Modal analysis

The FE equations governing the free vibration response of a structure are given by:

$$\boldsymbol{M}\ddot{\boldsymbol{u}}_s + \boldsymbol{K}_{el}\boldsymbol{u}_s = \boldsymbol{0} \tag{7}$$

where $\boldsymbol{M}$ is the global mass matrix, depending on the design variables $\boldsymbol{a}_s$. Substituting the trial solution $\boldsymbol{u}_s = \boldsymbol{\phi}_k e^{i\omega_k t}$ into Eq.(7) transforms the system into a Generalised Eigenvalue Problem (GEP):

$$(\boldsymbol{K}_{el} - \lambda_k \boldsymbol{M})\boldsymbol{\phi}_k = 0 \tag{8}$$

where $\lambda_k = \omega_k^2$ represents the k^{th} modal eigenvalue (squared natural frequency) and $\boldsymbol{\phi}_k$ is the corresponding eigenvector.

For a system with n_{dof} Degrees of Freedom (DoFs), equation Eq.(8) yields an equal number of eigenpairs $(\lambda_k, \boldsymbol{\phi}_k)$, although only a subset of low-frequency modes are of primary interest, and therefore computed, in structural engineering applications.

The frequency optimisation constraint considered in this work is

$$(\omega_j - \omega_i) \geq \rho_{mod}(\omega_j - \omega_i)_0 \tag{9}$$

stating that the distance between two frequencies of the optimised design $(\omega_j - \omega_i)$ must be larger than that of the initial design $(\omega_j - \omega_i)_0$, scaled by a factor ρ_{mod}. Notice that this kind of constraint could be used as a simplistic approach to maintain a certain margin between coalescing frequencies inducing flutter. However, such a constraint is not robust, as during optimisation, other modes will shift in frequency. Consequently, a different mode may converge to the same frequency as the ith and jth modes, causing a mode-switching phenomenon. This, in turn, leads to a non-smooth constraint function, undesirable in gradient-based optimisation.

Rearranging Eq.(9), a structural response related to the modal analysis is obtained:

$$J_\omega = \rho_{mod}\frac{(\omega_j - \omega_i)_0}{(\omega_j - \omega_i)} - 1 \leq 0 \tag{10}$$

which is negative or null in a successfully optimised structure.

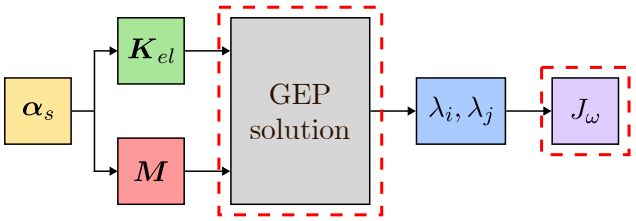


Figure 3: Computational graph of the modal analysis. Only the two eigenvalues of interest are indicated as output of the GEP solution.

Figure 3 illustrates the computational path for evaluating the frequency constraint response. Computation of the gradient vector $\frac{dJ_\omega}{d\boldsymbol{a}_s}$, which in turns depends on the derivatives of the two eigenvalues $\frac{d\lambda_i}{d\boldsymbol{a}_s}$ and $\frac{d\lambda_j}{d\boldsymbol{a}_s}$, is discussed in Sec.(2.4.2).

2.2.3 Linear Buckling analysis

The linear buckling problem in FE analysis follows a two-step procedure: an initial linear static step, where a reference load $\boldsymbol{f}_{ref}$ is applied to the structure, and a subsequent step in which the following GEP is solved for the critical eigenvalue:

$$(\boldsymbol{K}_{el} + \lambda_k \boldsymbol{K}_g)\boldsymbol{\phi}_k = 0 \tag{11}$$

where $\boldsymbol{K}_g$ represents the geometric component of the stiffness matrix, computed from the stress field resulting from the application of $\boldsymbol{f}_{ref}$. As in modal analysis, solving Eq.(11) for a system with n_{dof} degrees of freedom yields n_{dof} critical eigenpairs $(\lambda_k, \boldsymbol{\phi}_k)$. To maintain consistency in notation, the same symbols λ_k and $\boldsymbol{\phi}_k$ are used for the solutions of both GEPs in Eq.(8) and Eq.(11), despite their numerical and physical differences in the two problems.

Multiplying $\boldsymbol{f}_{ref}$ by the obtained λ_k yields the set of critical loads for the structure. Typically, and in this study, only the first (lowest) critical eigenvalue is of interest, denoted hereafter simply as λ_{cr}:

$$\lambda_{cr} \doteq \lambda_{k=1} \tag{12}$$

A buckling constraint is formulated as:

$$\lambda_{cr} \geq \rho_{buck}\, \lambda_{cr,0} \tag{13}$$

stating that the critical eigenvalue of the optimised structure must be greater than that of the initial design $(\lambda_{cr,0})$ by a factor ρ_{buck}, analogous to ρ_{mod}. The corresponding buckling response is:

$$J_{\lambda_{cr}} = \rho_{buck} \frac{\lambda_{cr,0}}{\lambda_{cr}} - 1 \leq 0 \tag{14}$$

which is negative or zero in a successfully optimised structure.

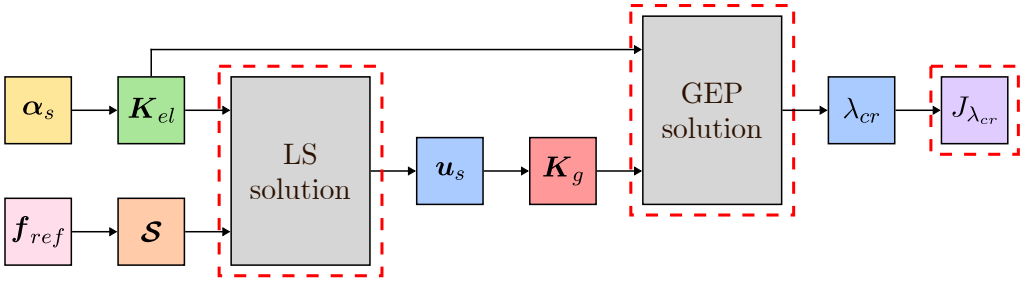


Figure 4: Computational graph of the buckling analysis. Only the first critical eigenvalue is indicated as output of the GEP solution.

Figure 3 illustrates the computational process for evaluating the buckling constraint response. Computation of the gradient vector $\frac{dJ_{\lambda_{cr}}}{d\boldsymbol{a}_s}$, which in turns depends on the gradient of the critical eigenvalue $\frac{d\lambda_{cr}}{d\boldsymbol{a}_s}$, is discussed in Sec.(2.4.3).

2.3 Algorithmic Differentiation

Algorithmic Differentiation is a powerful technique for accurately and efficiently computing the derivative of a function implemented as a computer program. It enables differentiation with respect to selected input parameters without requiring a derivative explicit symbolic representation. This subsection provides an overview of AD and addresses the challenges of differentiating a code that relies on an external library for computations. Further details on AD can be found in [39–41].

2.3.1 Overview

The fundamental principle of AD is that a computer program can be decomposed into a sequence of elementary mathematical operations, such as additions, multiplications, and function calls to elementary functions (e.g., $sin(x)$), whose derivatives can be readily computed with respect to given inputs. The sequence of these fundamental operations, where the output of one operation serves as the input for subsequent operations, forms a *computational graph*, which can be visually represented (Fig.5).

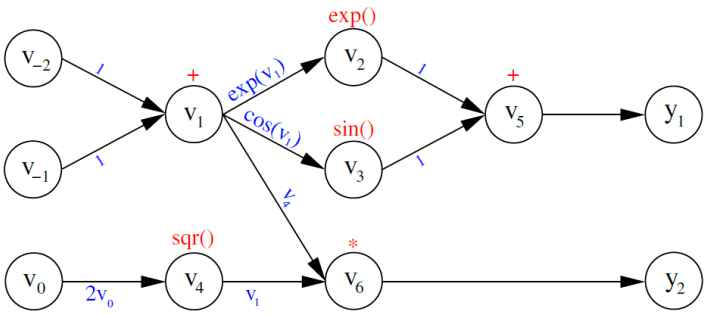


Figure 5: Example of a computational graph, from [39]. v_{-2}, v_{-1}, and v_0 are the inputs of a computer function that returns y_1 and y_2 as outputs.

The computation of the output gradient with respect to the inputs is achieved by

applying the chain rule throughout the computational graph. A key advantage of AD is that it obtains results up to machine precision and is not affected by truncation errors [41], typical of other techniques such as Finite Differences (FD).

Two AD paradigms exist, depending on the direction of derivative propagation (or accumulation). In AD *forward mode*, the derivative is propagated from the inputs to the outputs simultaneously with the evaluation of the computer function. This approach enables the computation of the gradient of an arbitrary number of outputs with respect to a single input and, therefore, requires as many evaluations of the computer function as the number of inputs of interest. Conversely, in AD *reverse mode*, the derivative is propagated from the outputs to the inputs. This process first requires the evaluation of the computer function (*direct path*), which is *recorded* by storing in memory the values of all intermediate variables between the inputs and the outputs. During the subsequent execution of the *reverse path*, the derivative of a single output y with respect to each intermediate variable v_j between the output and the input is computed:

$$\bar{v}_j \doteq \frac{\partial y}{\partial v_j} \tag{15}$$

where v_j and $\bar{v}_j$ denote the primal and the corresponding adjoint variable, respectively. A single run of the reverse path is required for each output y of interest, but it returns the gradient with respect to all the specified inputs.

For aerospace optimisations, where the number of design variables significantly exceeds the number of responses requiring gradients, the reverse mode of AD is preferable. Moreover, it is often coupled with the adjoint method due to their shared efficiency characteristics.

Consider a function $\boldsymbol{func()}$ with generic inputs $\boldsymbol{x} \in \boldsymbol{R}^n$ and outputs $\boldsymbol{y} \in \boldsymbol{R}^m$. In AD reverse mode, the primal and adjoint statements are defined as:

$$\boldsymbol{y} = \boldsymbol{func(x)} \qquad \text{(primal statement)} \tag{16a}$$

$$\bar{\boldsymbol{x}} += \left[\frac{\boldsymbol{\partial func(x)}}{\boldsymbol{\partial x}}\right]^{\boldsymbol{T}} \bar{\boldsymbol{y}} \qquad \text{(adjoint statement)} \tag{16b}$$

In Eq.(16b), the adjoint variable $\bar{\boldsymbol{y}}$ represents the *seed* of the output $\boldsymbol{y}$. The reverse path is initialised by "seeding" the output $\boldsymbol{y}$ of the computer program with 1 (i.e., the derivative of each component of $\boldsymbol{y}$ with respect to itself is equal to 1). This approach enables AD reverse mode to compute the rows of the matrix $\left[\frac{\boldsymbol{\partial func(x)}}{\boldsymbol{\partial x}}\right]$, ultimately yielding the full vector of sensitivities $\bar{\boldsymbol{x}}$.

The tool integrated into the in-house software *AUGUSTO++* for implementing AD is *CoDiPack* [26], a header-only library designed for differentiating industrial-scale software while, according to the High-Performance Computing (HPC), optimising both runtime and memory consumption. *CoDiPack* is paired with *MeDiPack*, an *MPI* wrapper that manages AD for the Message Passing Interface, leveraged by *AUGUSTO++* for its parallel computing capabilities. *CoDiPack* and *MeDiPack* have been developed at the Chair of Scientific Computing of the Technical University of Kaiserslautern - Landau. Based on *Jacobian Taping*, *Expression Templates*, and *Operator Overloading*, the integration of their functions into a computer code is straightforward and does not require significant modifications. The library usage consists of *registering the inputs* $\boldsymbol{x}$ at the beginning of the recording of the call to $\boldsymbol{func()}$ and *registering the output* $\boldsymbol{y}$.

2.3.2 Differentiation of a code involving LS and GEP solutions

The primal problems described in Sec.(2.2) require solving either a linear system, a generalised eigenvalue problem, or both. Differentiating LS and GEP solvers is generally either infeasible or computationally inefficient and is therefore avoided in practice [25, 37, 42]. Calculating gradients through the solution of these problems may not be possible, as their resolution is managed by external libraries which are not internally differentiated. Moreover, differentiating an LS or GEP solver would be highly inefficient. In the case of a direct linear solver, for instance, storing the entire computational graph associated with the factorisation phase of the system matrix would impose an unsustainable memory burden as the FE model size increases. This issue is even more pronounced in eigenvalue solvers, as they are typically iterative (e.g., *Lanczos* or *Arnoldi* methods) and often require the solution of linear systems within their algorithms [43].

To address this issue, Eq.(16) is considered. Let $\boldsymbol{func()}$ denote the function solving either the LS or GEP, with $\boldsymbol{x}$ and $\boldsymbol{y}$ representing generic inputs and outputs, respectively. Operationally, the recording of the computational graph is halted before calling $\boldsymbol{func()}$ and resumed immediately after its execution. As a result of this interruption, the gradient $\left[\frac{\boldsymbol{\partial func(x)}}{\boldsymbol{\partial x}}\right]$ cannot be directly obtained during the execution of the reverse path. Consequently, adjoint accumulation in Eq.(16b), which propagates derivatives from outputs to inputs, is only effective up to the computation of $\bar{\boldsymbol{y}}$. To circumvent this limitation, the computation of $\bar{\boldsymbol{x}}$ is instead performed by explicitly solving the adjoint statements associated with the LS or GEP. These adjoint formulations rely on analytical expressions for matrix operations. Further details on matrix derivative results for both forward and reverse AD modes can be found in [44], while the adjoint statements for GEP solutions are specifically outlined in sections (2.4.2) and (2.4.3).

2.4 Gradients computation

The derivative computation with respect to the design variables of the design responses ,i.e. both the objective function m and the constraints J_ω and $J_{\lambda_{cr}}$ are detailed in following subsections. For the sake of conciseness, the gradient computation for J_σ is omitted. Readers may refer to [37] for further details.

2.4.1 Gradient of the mass

The computation of the mass gradient is straightforward, as it does not involve solving any LS or GEP. The process consists of recording the computational path for mass evaluation as a function of $\boldsymbol{\alpha}_s$, which is registered as input. Each process calculates the mass of its allocated mesh portion, following the parallelisation strategy outlined in Sec.(3.2.1), and the total structural mass m, registered as output, is obtained through an *MPI reduction* operation. Reverse path execution yields the gradient vector $\frac{dm}{d\boldsymbol{\alpha}_s}$.

2.4.2 Gradient of the modal frequency, J_ω

The response J_ω sensitivity with respect to $\boldsymbol{\alpha}_s$ is computed using the discrete adjoint gradient evaluation procedure for vibration analysis described in [25]. The computational path of the modal analysis problem is recorded (Fig.3), excluding the GEP solution and

the evaluation of J_ω, with $\boldsymbol{\alpha}_s$ registered as the input and the two eigenvalues of interest, λ_i and λ_j, as outputs. Two reverse path runs are performed: first, λ_i is seeded with 1, returning $\frac{d\lambda_i}{d\boldsymbol{\alpha}_s}$, and then λ_j is seeded with 1, returning $\frac{d\lambda_j}{d\boldsymbol{\alpha}_s}$. The gradient $\frac{dJ_\omega}{d\boldsymbol{\alpha}_s}$ is derived by differentiating Eq.(10) as follows:

$$\frac{dJ_\omega}{d\boldsymbol{\alpha}_s} = -\rho_{mod}\frac{(\omega_j-\omega_i)_0}{(\omega_j-\omega_i)^2}\left(\frac{1}{2\omega_j}\frac{d\lambda_j}{d\boldsymbol{\alpha}_s} - \frac{1}{2\omega_i}\frac{d\lambda_i}{d\boldsymbol{\alpha}_s}\right) \tag{17}$$

This expression is evaluated after computing $\frac{d\lambda_i}{d\boldsymbol{\alpha}_s}$ and $\frac{d\lambda_j}{d\boldsymbol{\alpha}_s}$.

The adjoint model for the GEP, which allows retrieving $\bar{\boldsymbol{K}}_{el}$ and $\bar{\boldsymbol{M}}$, is given by [25]:

$$\bar{\boldsymbol{K}}_{el} \mathrel{+}= \bar{\lambda}_k \boldsymbol{\phi}_k \boldsymbol{\phi}_k^T \tag{18a}$$

$$\bar{\boldsymbol{M}} \mathrel{+}= -\lambda_k \bar{\boldsymbol{K}}_{el} \tag{18b}$$

These expressions hold under the assumption that the eigenvectors are mass-normalised, i.e., $\boldsymbol{\phi}_k^T \boldsymbol{M} \boldsymbol{\phi}_k = 1$. The adjoint model is executed during each reverse path iteration ($k = i$ and $k = j$), with $\bar{\lambda}_k = 1$, as the eigenvalues are seeded with this value.

2.4.3 Gradient of the buckling critical load, $J_{\lambda_{cr}}$

Similarly to the previous procedure, the computational path of the linear buckling analysis problem is recorded (Fig.4), excluding the LS and GEP solutions as well as the evaluation of $J_{\lambda_{cr}}$ and registering $\boldsymbol{\alpha}_s$ as input and λ_{cr} as output. A single execution of the reverse path, where λ_{cr} is seeded with 1, yields $\frac{d\lambda_{cr}}{d\boldsymbol{\alpha}_s}$. This derivative is then substituted into the hand-differentiated expression of Eq.(14) to compute $\frac{dJ_{\lambda_{cr}}}{d\boldsymbol{\alpha}_s}$:

$$\frac{dJ_{\lambda_{cr}}}{d\boldsymbol{\alpha}_s} = -\rho_{buck}\frac{\lambda_{cr,0}}{\lambda_{cr}^2}\frac{d\lambda_{cr}}{d\boldsymbol{\alpha}_s} \tag{19}$$

Since in the primal problem the LS is solved before the GEP, in the reverse path, the adjoint problem for the GEP is solved before that of the LS. The adjoint model of the GEP, under the assumption of $\boldsymbol{K}_g$-normalisation of the eigenvectors (i.e., $\boldsymbol{\phi}_k^T \boldsymbol{K}_g \boldsymbol{\phi}_k = 1$), is given by:

$$\bar{\boldsymbol{K}}_{el} \mathrel{+}= -\bar{\lambda}_k \boldsymbol{\phi}_k \boldsymbol{\phi}_k^T \tag{20a}$$

$$\bar{\boldsymbol{K}}_{g} \mathrel{+}= \lambda_k \bar{\boldsymbol{K}}_{el} \tag{20b}$$

This formulation is identical to Eq.(18), except that $\bar{\boldsymbol{K}}_g$ replaces $\bar{\boldsymbol{M}}$, and there is a sign change due to the different placement of the matrices $\boldsymbol{M}$ and $\boldsymbol{K}_g$ in Eq.(8) and Eq.(11).

The adjoint statement related to the solution of a LS is [25, 37]:

$$\boldsymbol{K}_{el}^T \boldsymbol{s} = \frac{\partial O}{\partial \boldsymbol{u}_s} \tag{21a}$$

$$\bar{\boldsymbol{K}}_{el} \mathrel{+}= -\boldsymbol{s} \cdot \boldsymbol{u}_s^T \tag{21b}$$

$$\bar{\boldsymbol{\mathcal{S}}} \mathrel{+}= \boldsymbol{s} \tag{21c}$$

where O is the output to be differentiated, i.e., in this case, $O = \lambda_{cr}$. Since $\boldsymbol{K}_{el}$ serves as input to both the LS and GEP solvers, the reverse accumulation of $\bar{\boldsymbol{K}}_{el}$ occurs in two steps: first through Eq.(20a), followed by Eq.(21b).

3 THE FE SOLVER

This section provides an overview of the in-house software *AUGUSTO++*, which is designed to solve the primal problem, compute sensitivities for gradient-based structural optimisation, and offer an interface for integrating with the *python* module for structural optimisation.

3.1 The Code

The code is written in *C++*, following the Object-Oriented Programming paradigm to ensure modularity, reusability, and ease of maintenance. Parallel computing capabilities are implemented via the *MPI* library [45], while AD is provided by *CoDiPack* alongside *MeDiPack*. *AUGUSTO++* computes the response of highly flexible structures under the assumption of small strains and linear elastic material within an Updated Lagrangian framework. For large displacement analyses, where the FE formulation includes geometric nonlinearities, strains are evaluated using a co-rotational approach [46]. The user can choose to link the code with either *MUMPS* [47, 48] or *PETSC* [49–51] for solving LS, and with *SLEPC* [52] for solving GEP. In this work, *PETSC* and *SLEPC* are used.

The code FE library includes:

1. *4-node bilinear isoparametric shell element*, based on the *Nastran* CQUAD4 element [53]. The computation of the geometric contribution to the stiffness matrix, $\boldsymbol{K}_g$, follows the formulation presented in [54, 55], treating large rotations according to [56] and [57]. Selective Reduced Integration (SRI) and Assumed Strains Method (ASM) mitigate the problem of shear locking in the linear and non-linear case, respectively;

2. *3-nodes Constant Strain Triangle (CST) and Discrete Kirchhoff Triangle (DKT) elements* based on [58].

3. *2-nodes linear isoparametric bar element* with constant cross section and Selective Reduced Integration. The section can be offset from the nodes line to model aeronautical stiffened panels. The element is based on *Nastran* CBAR formulation [59].

4. *Nonlinear Euler-Bernoulli beam element* with Hermite interpolation [58].

5. *Rigid connector element RBE2*

3.2 Parallelisation strategy

3.2.1 Primal problem

The primal problem parallelisation strategy distributes the computational mesh of the structural model among the running *MPI* processes. Given a total of N_{FE} finite elements (FEs), each of the N processes is assigned (approximately) N_{FE}/N elements, forming a contiguous set of element IDs. Consequently, each process computes and stores only a subset of the matrices $\boldsymbol{K}_{el}$, $\boldsymbol{K}_g$, and $\boldsymbol{M}$. Similarly, the residual vector $\boldsymbol{\mathcal{S}}$ is distributed among processes according to the degrees of freedom (DOFs) associated with the assigned FEs. The LS and GEP solutions in equations (3), (8), and (11) are computed in parallel using *PETSC* and *SLEPC* routines. The solution vectors $\boldsymbol{u}_s$ and $\boldsymbol{\phi}_k$ are retrieved as distributed vectors and can be gathered either on the master process or across all processes

as required. Quantities registered as outputs by *CoDiPack* are allocated on the master process, while value of zero is assigned to other processes. The rationale of this approach is discussed in Sec.(3.2.2).

3.2.2 Gradients computation

As already mentioned, computation of gradients with *CoDiPack* relies on the registration of inputs and outputs, which is organised among the running processes as follows:

- *Input registration:* Each design variable (DV) in $\boldsymbol{\alpha}_s$ is registered as an input by all processes. In the *AUGUSTO++* current implementation, only section property and material DVs are available (e.g., shell thickness or material properties). Each process initially loads all property cards from the FE input file along with the corresponding DV objects.

- *Output registration:* The quantities m, J_σ, and eigenvalues λ_k are registered as outputs by all processes. Unlike input registration, output information is collectively communicated to all running processes. This ensures that when a single process selects an output to be differentiated, the information becomes available to the entire *MPI* communicator. However, this can lead to an incorrect gradient computation due to an implicit multiplication factor depending on the number of processes. To prevent this, the registered output value is allocated exclusively on the master process, while a value of zero is assigned all the other processes, as discussed in Sec.(3.2.1).

Regarding the reverse path, the adjoint computations associated with the LS and GEP solutions are also executed in parallel. In the LS Eq.(21a), both $\boldsymbol{K}_{el}$ (available from the direct path) and $\dfrac{\partial O}{\partial \boldsymbol{u}_s}$ are distributed among processes. Similarly, the adjoint matrices and vectors computed in equations (21), (18), and (20) follow the same distribution scheme as their corresponding primal counterparts.

At the end of the reverse path, the gradients are extracted by each running process. Due to the mesh partitioning and the reverse path parallel execution, each process P_i computes a portion of the gradient vector, denoted as $\left(\dfrac{dJ_\sigma}{d\boldsymbol{\alpha}_s}\right)_{P_i}$. To obtain the complete derivative values, a summation across all processes is performed at the end of the algorithm. For instance, considering the derivative of J_σ:

$$\frac{dJ_\sigma}{d\boldsymbol{\alpha}_s} = \sum_{i=1}^{N} \left(\frac{dJ_\sigma}{d\boldsymbol{\alpha}_s}\right)_{P_i} \tag{22}$$

where the summation is executed via an *MPI All-reduction* operation.

3.3 Code wrapping with *SWIG*

The integration of *AUGUSTO++* within a *python* environment is achieved by wrapping its high-level functions with *SWIG*. These functions handle key analysis operations, including mesh parsing and preprocessing, solving the primal problem, extracting desired responses, and computing sensitivities. Additionally, higher-level *python* scripts facilitate

communication with other *SWIG*-wrapped modules, such as *SU2* [11, 60, 61] and spline modules referenced in [34, 37] for aerostructural optimisation, as well as the numerical optimisation library *SciPy* [62].

4 RESULTS

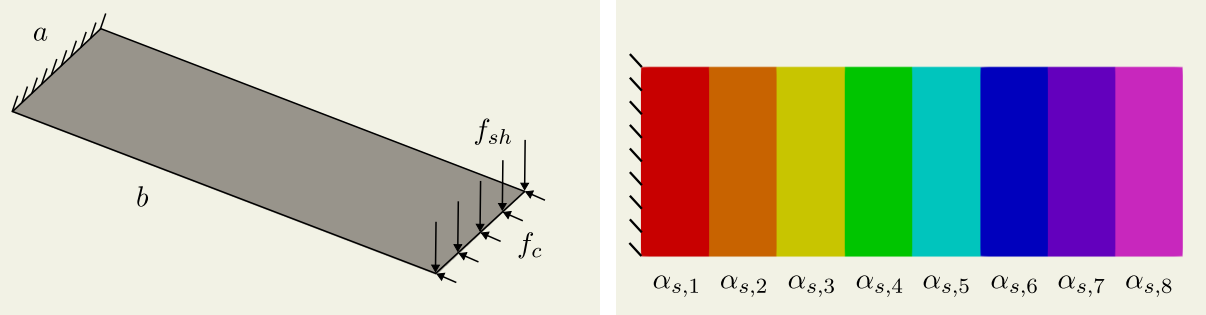
This section presents the application of the structural optimisation framework to two test cases: a flat plate, and a wingbox from the High-Fidelity Aeroelastic Optimization Benchmark [63]. Several weight minimisation problems are solved for each test case, considering the constraints described above, either individually or in combination. The optimisation algorithm used is the Sequential Least Squares Quadratic Programming (SLSQP) [64], available in the *python SciPy* library.

4.1 Flat plate

4.1.1 Testcase

The first test case considered is a rectangular flat plate, modelled using bilinear shell elements. The structure is clamped along one of its short edges and subjected to a distributed in-plane and out-of-plane load along the opposite edge. Fig.6a shows the test case, including the applied load, boundary conditions, and dimensions, while Tab.1 summarises the relevant mesh, material, and physical parameters. Stress constraints are aggregated over all mesh elements.

DVs are defined by grouping the shell elements into eight independent property cards, dividing the structure into an equal number of longitudinal stripes, as shown in Fig.6b. The thickness of each stripe shell property can be independently adjusted by the optimiser.



(a) Problem definition.

(b) Position and identifiers of the DVs.

Figure 6: Flat plate testcase.

Stress constraints are imposed under the loading condition involving only shear forces f_{sh}, whereas the buckling constraint is assessed under in-plane compressive forces f_c.

Four optimisation runs were performed: one for each individual constraint (stress, frequency, and buckling), and a fourth one where all three constraints were simultaneously enforced. An admissible equivalent stress of 270 MPa was considered for the stress constraint. For the frequency constraint, the gap between the first two natural frequencies was monitored and required to remain above 120% of its initial value. The buckling

Table 1: Flat plate data.

Parameter	Description	Value
N_{FE}	Number of FEs	179,101
N_n	Number of nodes	180,000
N_{dofs}	Number of DOFs	1,080,000
a	Short edge length	700.0 mm
b	Long edge length	2000.0 mm
t_0	Initial thickness	5.0 mm
m_0	Initial mass	18.90 kg
E	Young's modulus	70,000 MPa
ν	Poisson's ratio	0.3
ρ	Density	2700.0 kg/m^3
f_{sh}	Shear load	0.125 N/mm
f_c	Compressive load	1.0 N/mm

load was enforced to be at least 20% higher than its initial value. Tab.2 lists all the optimisation parameters.

4.1.2 Optimisation results

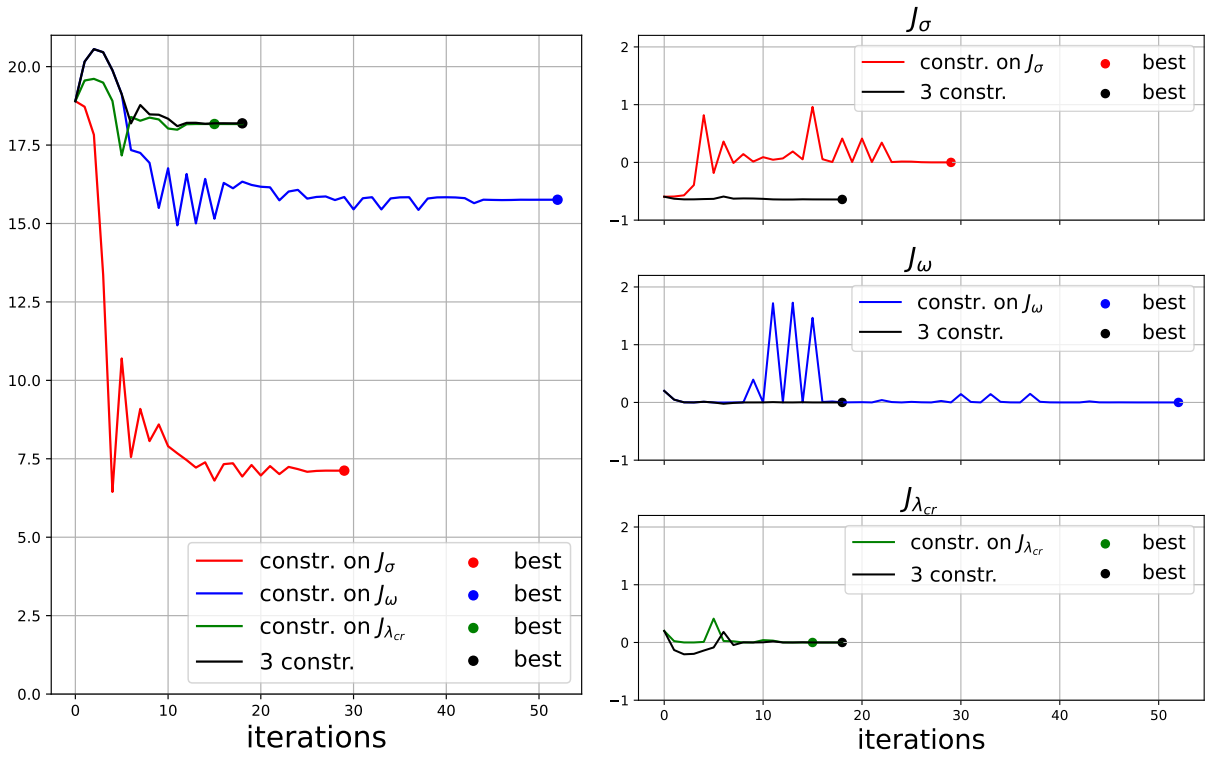
Table 2: Flat plate optimisation parameters.

parameter	value
ρ_{KS}	50.0
ρ_{mod}	1.2
ρ_{buck}	1.2
σ_{ADM}	270.0 MPa
$(\omega_2 - \omega_1)_0$	1.021 rad/s
$\lambda_{cr,0}$	0.4602

Fig.7a shows the evolution of the structural mass throughout the optimisation iterations, while Fig.7b illustrates the behaviour of each constraint for the four carried out optimisations with individual and combined constraints. Fig.8a depicts the design variable evolutions during the optimisation, while the final values of the plate thicknesses and mass are summarised in Tab.3.

Among the single-constraint optimisations, the one involving the frequency constraint required significantly more iterations than those involving stress or buckling constraints. This behaviour is ascribed to mode swapping, triggered by variations in the design variables, which complicates convergence for the optimiser. Fig.9a supports this interpretation, showing how, during the optimisation, the frequencies of modes 2 and 3 become

nearly identical, leading to a swap: from one iteration to the next, the second frequency is associated with a physically different mode.

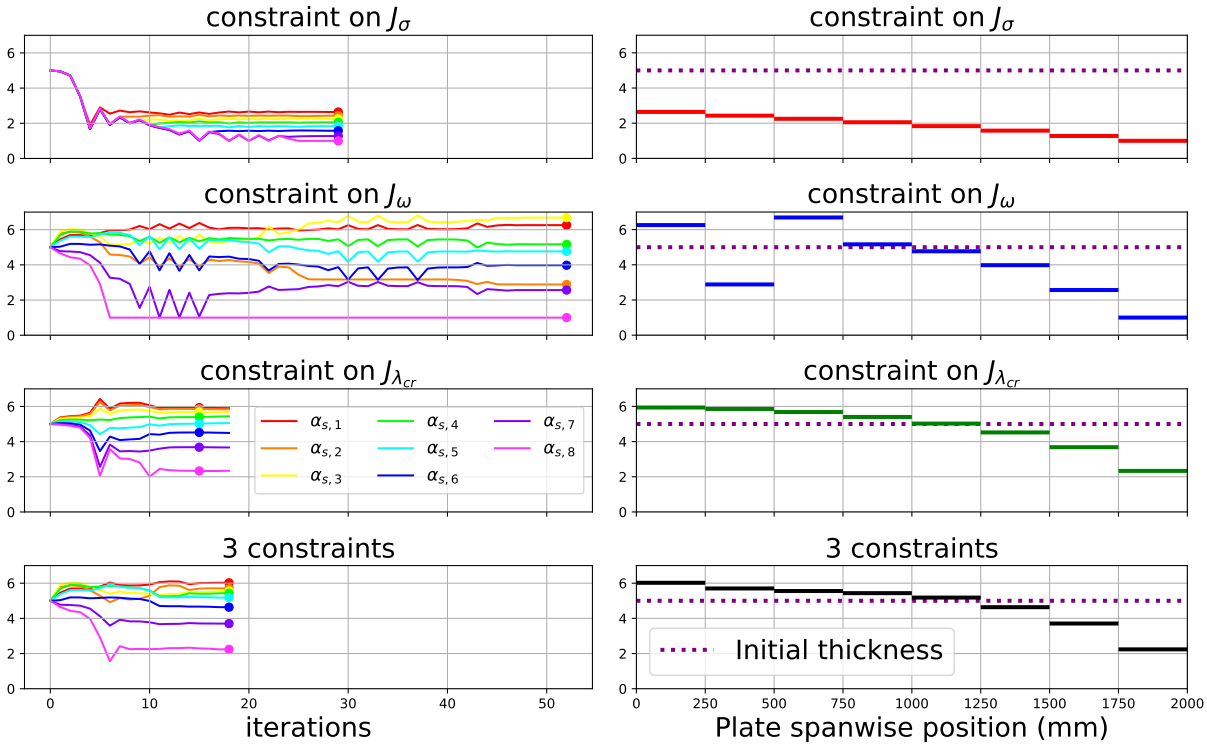


(a) Mass (Kg).

(b) Constraints.

Figure 7: Mass and constraints vs optimisation iterations.

It is evident that the mass reduction achieved when frequency or critical load constraints are imposed is smaller than when only the stress constraint is applied. Moreover, when all constraints are simultaneously imposed, the optimisation process is primarily governed by the buckling and frequency constraints. As can be noticed from Fig.8b, in this last case the thickness trend is similar to the case in which only buckling is enforced. However, small differences exist in order not to violate the frequency constraint. This can be appreciated in Fig.9b, showing the first 5 natural frequencies for the optimisation case featuring the buckling constraint only. For this optimal design, the value of the constraint on the frequency is slightly below the admissible one (1.225 rad/s).

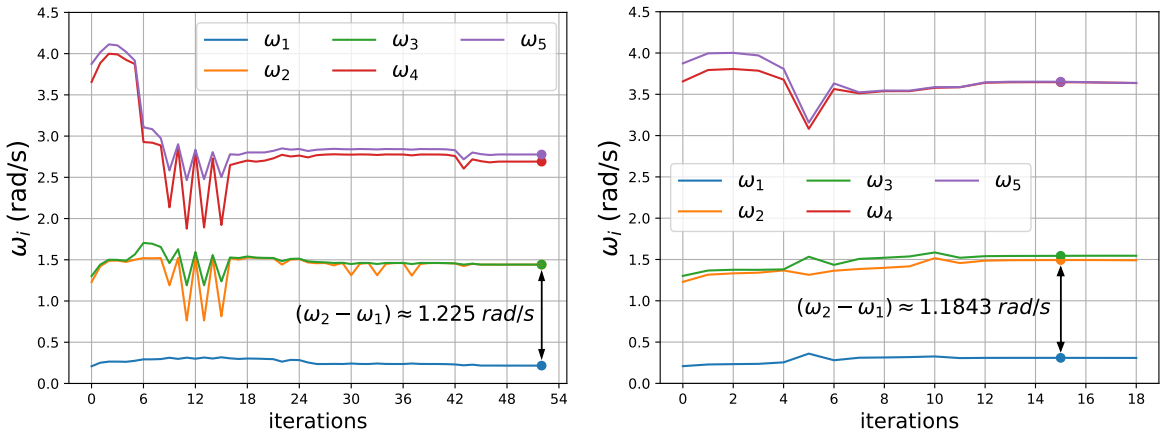


(a) Design variables (mm). (b) Thickness distribution along the plate (mm).

Figure 8: DVs evolution during the optimisation and thickness distribution along the plate of the best designs.

Table 3: Plate DVs: initial values, optimisation boundaries and optimum values.

DV	$\alpha_{s,0}$	$\alpha_{s,min}$	$\alpha_{s,max}$	constr. J_σ	constr. J_ω	constr. $J_{\lambda_{cr}}$	3 constr.
$\alpha_{s,1}$	5.0	1.0	10.0	2.64	6.26	5.94	6.02
$\alpha_{s,2}$	5.0	1.0	10.0	2.43	2.88	5.85	5.70
$\alpha_{s,3}$	5.0	1.0	10.0	2.25	6.69	5.68	5.55
$\alpha_{s,4}$	5.0	1.0	10.0	2.05	5.16	5.40	5.43
$\alpha_{s,5}$	5.0	1.0	10.0	1.83	4.77	5.02	5.18
$\alpha_{s,6}$	5.0	1.0	10.0	1.58	3.98	4.52	4.64
$\alpha_{s,7}$	5.0	1.0	10.0	1.27	2.57	3.68	3.71
$\alpha_{s,8}$	5.0	1.0	10.0	1.00	1.00	2.33	2.24
m (Kg)				7.12	15.76	18.17	18.19
Δm (%)				-62.33	-16.61	-3.86	-3.76
J_σ				**9.0e-6**	-	-	-0.641
J_ω				-	**8.0e-7**	-	**-6.0e-7**
$J_{\lambda_{cr}}$				-	-	**8.0e-6**	**5.0e-6**



(a) Case: constraint on frequencies

(b) Case: constraint on buckling.

Figure 9: First 5 modes frequencies for the different optimisation cases.

4.2 High-Fidelity Aeroelastic Optimisation Benchmark

4.2.1 Testcase

The second testcase is the wing-box benchmark model for high-fidelity aeroelastic optimisation developed at the Multidisciplinary Design Optimization Laboratory of the University of Michigan [65]. The wing is modelled with bilinear shell elements, and its mesh is displayed in Fig.10a. It is constrained at both the root and the wing-fuselage intersection, as shown in Fig.10b. Table 4 lists the relevant mesh, material, and physical parameters. A DV is assigned to each wingbox panel, corresponding to a rib or forming part of the front/rear spars or the top/bottom skin. The DVs comprises the groups the shell elements associated to the panel. The distribution of these DVs, amounting to a total of 111, is illustrated in Fig.10c. All rib nodes are subjected to a uniform nodal force of 20.4 N, resulting in a wing tip vertical displacement of approximately 12.5% of the semi-span.

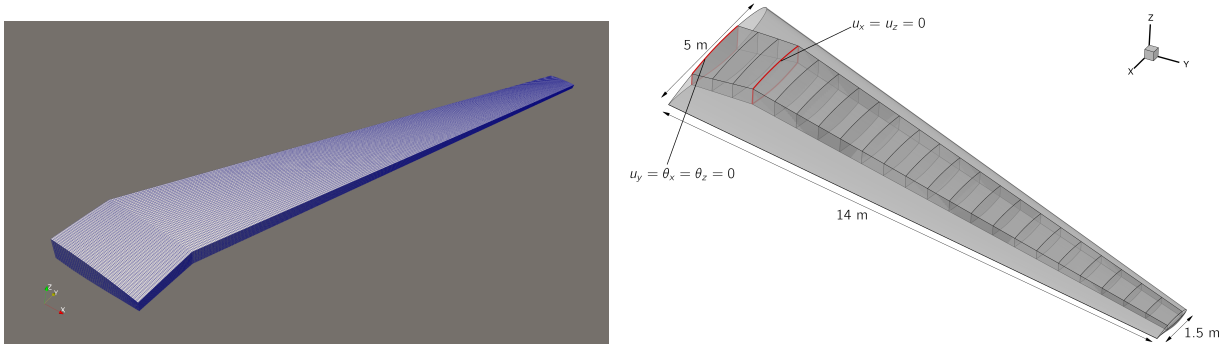
Three optimisation runs were performed, corresponding to two distinct constraint cases: stress constraints and frequency constraints plus a third case combining both sets simultaneously. A brief description of each case is provided below:

- Stress constraints set (*constr.* J_σ): the stress is aggregated separately across the five components constituting the wing: the ribs, front spar, rear spar, top skin, and bottom skin (Fig.10d). This results in five distinct stress responses: J_σ^{RIBS}, J_σ^{FSPAR}, J_σ^{RSPAR}, J_σ^{TSKIN}, and J_σ^{BSKIN}, all of which are constrained simultaneously. The admissible stress value is 320 MPa.

- Frequency constraints set (*constr.* J_ω): the frequencies separations, $(\omega_2 - \omega_1)$ and $(\omega_3 - \omega_2)$, are maintained above their respective initial design values, i.e., $\rho_{mod,(2-1)}$ and $\rho_{mod,(3-2)}$ are set to 1. Additionally, a condition requires that ω_1 remains above 115% of its initial value. According to Eq.(10), this condition corresponds to $\omega_j = \omega_1$, $\omega_i = 0$, and $\rho_{mod,(1)} = 1.15$.

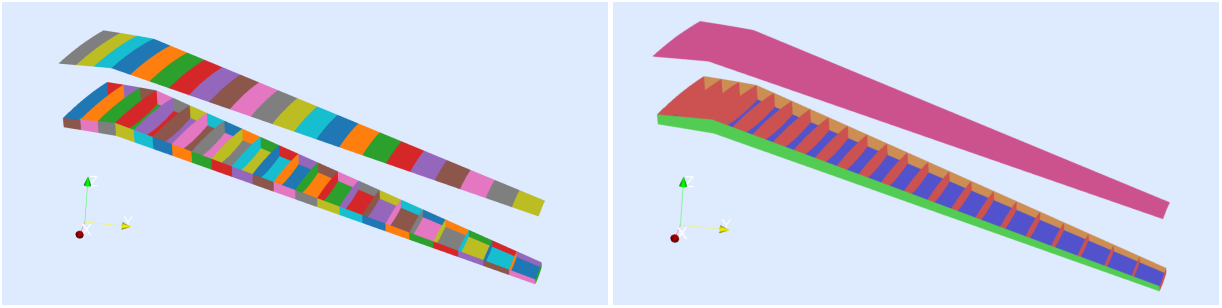
- Mixed constraints set (*constr.* J_σ & J_ω): This case involves the simultaneous application of the stress and frequency constraint sets previously described , for a total of eight constraints.

Tab.5 summarised the optimisation parameters.



(a) Mesh.

(b) Applied constraints (from [63]).



(c) Wing DVs. Each separately coloured panel represents an independent DV.

(d) Wing Components: ribs (red), front spar (orange), rear spar (green), top skin (magenta) and bottom skin (blue).

Figure 10: High-Fidelity Aeroelastic Optimisation Benchmark testcase.

4.2.2 Optimisation results

Figure (11a) shows the evolution of the structural mass throughout the optimisation iterations, while Figure (11b) reports the behaviour of each constraint across the three optimisation cases. Figure (12) illustrates the optimal thickness distributions of the five wing components, with the optimal design masses listed in Tab.6.

Unlike the plate test case, the optimisation involving frequency constraints required significantly fewer iterations compared to those with stress constraints. This is ascribed to the absence of mode swapping affecting the constrained frequencies ω_1, ω_2, and ω_3. Indeed, Fig.13 shows that the first five modal frequencies exhibit considerably smoother variations throughout the optimisation process.

The mass reduction achieved when applying the frequency constraint set is more pronounced than when enforcing the stress constraint set, indicating the stress constraint as the primary driver, in these particular optimisations.

The optimisation with the J_σ set, shows all five constraints active. The optimal wing thickness distribution follows the expected trend, with decreasing values towards the tip. Minor deviations occur near the very outboard region for the ribs and spars, and close to the fuselage-wing junction for the front spar.

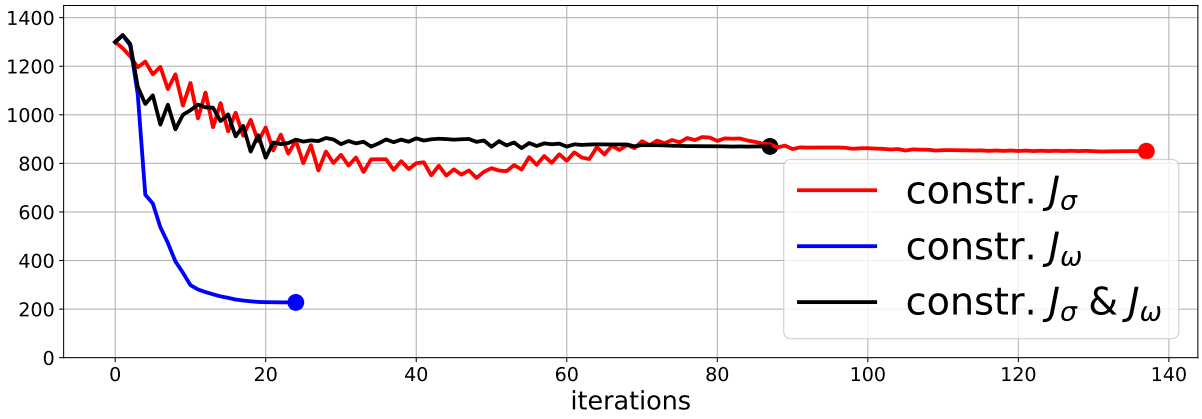
In the optimisation with the J_ω set, two constraints are active for the optimal design. With the rib and spar thicknesses at their lower bounds, the optimiser is unable to further

Table 4: Wing data.

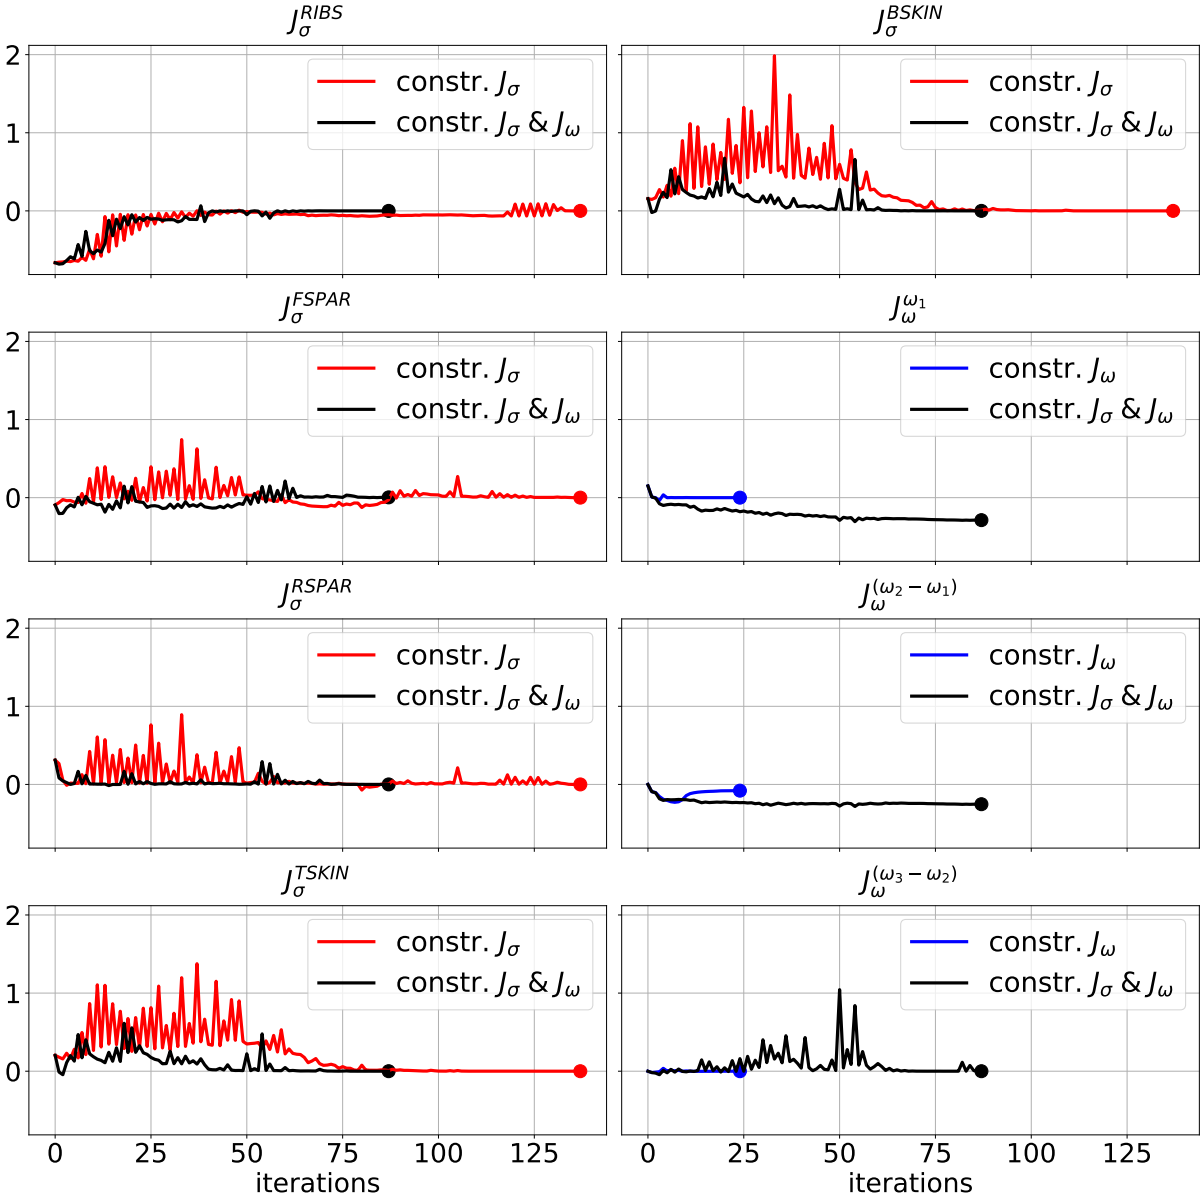
Parameter	Description	Value
N_{FE}	Number of FEs	71,200
N_n	Number of nodes	69,963
N_{dofs}	Number of DOFs	419,778
b	Wing semi-span	14.0 m
t_0	Panels initial thickness	7.0 mm
m_0	Initial mass	1299.47 kg
E	Young's modulus	70 GPa
ν	Poisson's ratio	0.33
ρ	Density	2700.0 kg/m^3

Table 5: Wing optimisation parameters.

parameter	value
ρ_{KS}	50.0
$\rho_{mod,(1)}$	1.15
$\rho_{mod,(2-1)}$	1.0
$\rho_{mod,(3-2)}$	1.0
σ_{ADM}	320.0 MPa
$\omega_{1,0}$	0.6995 rad/s
$(\omega_2 - \omega_1)_0$	1.3607 rad/s
$(\omega_3 - \omega_2)_0$	0.7030 rad/s
$\alpha_{s,0}$	7.0 mm
$\alpha_{s,min}$	1.0 mm
$\alpha_{s,max}$	20.0 mm



(a) Mass (kg).



(b) Constraints.

Figure 11: Mass and constraints vs optimisation iterations.

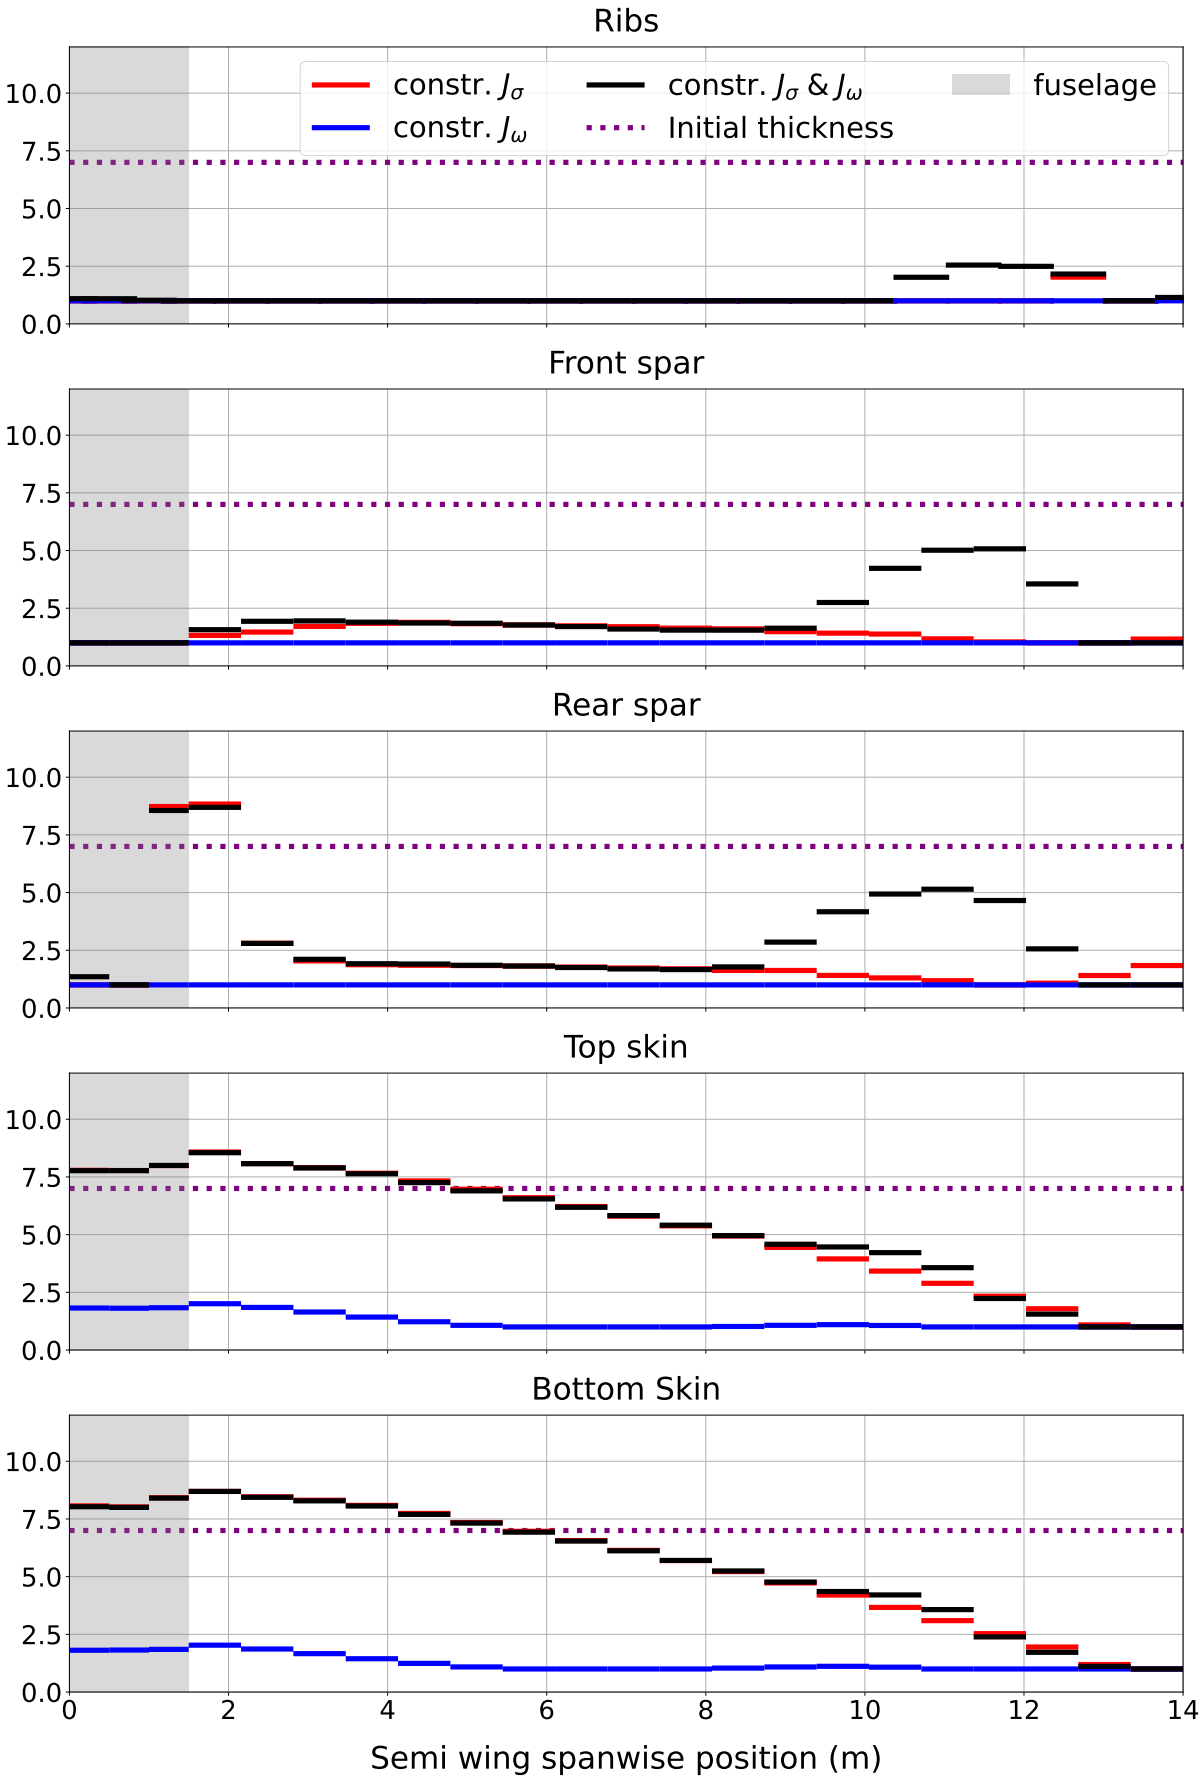


Figure 12: Optimal thickness distribution along the span.

Table 6: Wing: optimum values.

	constr. J_σ	constr. J_ω	constr. J_σ & J_ω
m (Kg)	850.17	227.58	869.91
Δm (%)	-34.58	-82.49	-33.06
J_σ^{RIBS}	**3.1e-5**	-	**8.7e-6**
J_σ^{FSPAR}	**3.1e-4**	-	3.7e-3
J_σ^{RSPAR}	**1.4e-4**	-	**2.4e-4**
J_σ^{TSKIN}	**1.1e-5**	-	**1.4e-4**
J_σ^{BSKIN}	**6.9e-6**	-	**1.4e-4**
$J_\omega^{\omega_1}$	-	**4.0e-7**	-0.2877
$J_\omega^{(\omega_2-\omega_1)}$	-	-0.07966	-0.2531
$J_\omega^{(\omega_3-\omega_2)}$	-	**2.4e-5**	**5.9e-5**

reduce the wing mass while attempting to selectively activate the constraint on the $\omega_2 - \omega_1$ separation. Fig.13 illustrates the trend of the first five natural frequencies throughout the

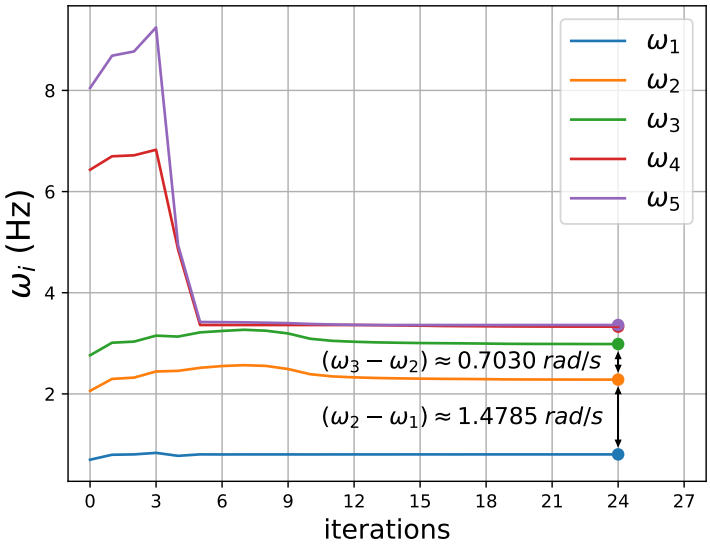


Figure 13: First 5 modes frequencies for constraint on frequencies case.

optimisation process.

When all constraints are imposed simultaneously, the optimisation process is primarily governed by the stress constraints set. However, the constraint on the $\omega_3 - \omega_2$ separation is also active and influences the design, resulting in an additional 20 kg of mass compared to the optimisation considering only the J_σ set. This is observed from Fig.12, showing the thickness distributions of the ribs, top skin, and bottom skin almost perfectly aligned along the first half to two-thirds of the semi-wing span. In the outboard region of the wing, the spars and ribs exhibit a wavy-like thickness increase, necessary to maintain the desired separation between the second and third frequencies.

5 CONCLUSIONS

This paper presents the development of an in-house finite element solver for gradient-based structural optimisation of aircraft wings, assisted by automatic differentiation. The tool features a modular architecture, enabling straightforward replacement of external libraries and seamless integration of additional design variables and responses. Furthermore, a *python* wrapper allows the use of high-level functions within a *python* environment, facilitating efficient interfacing with other modules, particularly those related to computational fluid dynamics for aerostructural optimisation.

The solver technical background is then introduced, including the formulation of the primal problem for static, modal, and buckling analysis. A brief overview of AD is provided, alongside a discussion of the challenges posed by partial recording when external libraries are involved. The methodology for computing derivatives of selected responses is detailed, employing the combination between the adjoint method and AD.

The development of *AUGUSTO++* is exhaustively examined with particular attention paid to the parallel strategy implemented for both the primal and sensitivity problems.

The optimisation tool capabilities are tested on both a flat plate and a wing-box model. The solver successfully identifies feasible designs under various combinations of stress, frequency, and buckling constraints applied individually or simultaneously. Interactions between constraints of differing nature occasionally direct the optimiser towards thickness distributions which appear less intuitive from an engineering perspective, warranting further investigation.

Future developments will focus on implementing a modal tracking algorithm, such as the Modal Assurance Criterion (MAC), to mitigate the risk of modal swapping and support future aeroelastic analyses. Additionally, the inclusion of aerodynamic design variables and responses, alongside structural ones, will facilitate the coupled optimisation of high-fidelity wing models.

6 ACKNOWLEDGEMENTS

The research work is carried out as part of the research project OPERATIONAL (TED2021-131119A-I00), funded by MCIN/AEI/10.13039/501100011033 and the European Union NextGenerationEU/PRTR.

REFERENCES

[1] ICAO. https://www.icao.int/Newsroom/Pages/Passenger-air-traffic-surpasses-pre-pandemic-levels.aspx, .

[2] EASA. https://www.eurocontrol.int/publication/european-aviation-environmental-report-2022#:~:text=While%20the%20number%20of%20flights,emissions%20potentially%20rising%20to%20188, 2022.

[3] R. Cavallaro and L. Demasi. Challenges, ideas, and innovations of joined-wing configurations: A concept from the past, an opportunity for the future. *Progress in Aerospace Sciences*, 87:1 – 93, 2016. ISSN 0376-0421. doi:http://dx.doi.org/10.1016/j.paerosci.2016.07.002. URL http://www.sciencedirect.com/science/article/pii/S0376042116300471.

[4] https://simpleflying.com/boeing-ttbw/, .

[5] Innovative aircraft and propulsion technologies to improve air quality near airports, . URL https://indigo-sustainableaviation.eu/.

[6] Innovative aircraft and propulsion technologies to improve air quality near airports, . URL https://cordis.europa.eu/project/id/101096055.

[7] P. Norczyk, A. Cini, and R. Cavallaro. Local air quality and noise improvements via optimization of strut-braced wings with distributed electric propulsion. Number AIAA 2025-0456. SciTech2025, Orlando, Florida, 6-10 January 2025. doi:10.2514/6.2025-0456.

[8] M. A. Akgun, R. T. Haftka, K. C. Wu, J. L. Walsh, and J. H. Garcelon. Efficient structural optimization for multiple load cases using adjoint sensitivities. *AIAA journal*, 39(3):511–516, 2001.

[9] G. Kenway, G. Kennedy, and J. R. Martins. Aerostructural optimization of the common research model configuration. In *15th AIAA/ISSMO multidisciplinary analysis and optimization conference*, page 3274, 2014.

[10] E. J. Adler and J. R. R. A. Martins. Efficient aerostructural wing optimization considering mission analysis. *Journal of Aircraft*, 60(3), May 2023. ISSN 1533-3868. doi:10.2514/1.c037096.

[11] T. Albring, M. Sagebaum, and N. Gauger. Efficient aerodynamic design using the discrete adjoint method in su2. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2016.

[12] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon. Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124 (1-2):171–190, 2000.

[13] P. Bartholomew and H. Wellen. Computer-aided optimization of aircraft structures. *Journal of Aircraft*, 27(12):1079–1086, 1990.

[14] P. Bartholomew and H. Wellen. Computer-aided optimum design of aircraft structures. In *Control and Dynamic Systems*, volume 59, pages 225–263. Elsevier, 1993.

[15] C. M. Soares, V. F. Correia, H. Mateus, and J. Herskovits. A discrete model for the optimal design of thin composite plate-shell type structures using a two-level approach. *Composite Structures*, 30(2):147–157, 1995.

[16] A. B. Lambe, G. J. Kennedy, and J. R. Martins. An evaluation of constraint aggregation strategies for wing box mass minimization. *Structural and Multidisciplinary Optimization*, 55:257–277, 2017.

[17] G. J. Kennedy and J. R. Martins. A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elements in Analysis and Design*, 87:56–73, 2014. ISSN 0168-874X. doi:https://doi.org/10.1016/j.finel.2014.04.011. URL https://www.sciencedirect.com/science/article/pii/S0168874X14000730.

[18] E. J. Adler and J. R. Martins. Efficient aerostructural wing optimization considering mission analysis. *Journal of Aircraft*, 60(3):800–816, 2023.

[19] E. Jonsson, C. Riso, B. B. Monteiro, A. C. Gray, J. R. Martins, and C. E. Cesnik. High-fidelity gradient-based wing structural optimization including geometrically nonlinear flutter constraint. *AIAA Journal*, 61(7):3045–3061, 2023.

[20] C. Cardani and P. Mantegazza. Calculation of eigenvalue and eigenvector derivatives for algebraic flutter and divergence eigenproblems. *AIAA Journal*, 17(4):408–412, 1979.

[21] H. Mateus, C. M. Soares, and C. M. Soares. Buckling sensitivity analysis and optimal design of thin laminated structures. *Computers & structures*, 64(1-4):461–472, 1997.

[22] T. H. Lee. Adjoint method for design sensitivity analysis of multiple eigenvalues and associated eigenvectors. *AIAA journal*, 45(8):1998–2004, 2007.

[23] S. He, E. Jonsson, and J. R. Martins. Derivatives for eigenvalues and eigenvectors via analytic reverse algorithmic differentiation. *AIAA Journal*, 60(4):2654–2667, 2022.

[24] S. He, Y. Shi, E. Jonsson, and J. R. Martins. Eigenvalue problem derivatives computation for a complex matrix using the adjoint method. *Mechanical Systems and Signal Processing*, 185:109717, 2023.

[25] M. Schwalbach, T. Verstraete, and N. R. Gauger. Discrete adjoint gradient evaluations for linear stress and vibration analysis. *Computing and Visualization in Science*, 21:23–31, 2019.

[26] N. G. M. Sagebaum, T. Albring. High-performance derivative computations using codipack. *ACM Transactions on Mathematical Software (TOMS)*, 45(4), 2019. URL `https://dl.acm.org/doi/abs/10.1145/3356900`.

[27] F. Torrigiani, B. Nagel, and R. Cavallaro. Development of an aeroelastic stability module for wing planform optimization. In *19th International Forum on Aeroelasticity and Structural Dynamics, IFASD 2022*, 2022.

[28] I. Ozaki and T. Terano. Applying an automatic differentiation technique to sensitivity analysis in design optimization problems. *Finite elements in analysis and design*, 14 (2-3):143–151, 1993.

[29] J. F. M. Barthelemy and L. E. Hall. Automatic differentiation as a tool in engineering design. *Structural optimization*, 9:76–82, 1995.

[30] D. T. Nguyen, R. Qamar, and H. Runesha. Automatic differentiation for design sensitivity analysis of structural systems using parallel-vector processors. *Advances in engineering software*, 29(3-6):375–382, 1998.

[31] M. Tadjouddine, S. A. Forth, and A. J. Keane. Adjoint differentiation of a structural dynamics solver. In *Automatic Differentiation: Applications, Theory, and Implementations*, pages 309–319. Springer, 2006.

[32] G. Wu. A framework for structural shape optimization based on automatic differentiation, the adjoint method and accelerated linear algebra. *Structural and Multidisciplinary Optimization*, 66(7):151, 2023.

[33] A. Chandrasekhar, S. Sridhara, and K. Suresh. Auto: a framework for automatic differentiation in topology optimization. *Structural and Multidisciplinary Optimization*, 64(6):4355–4365, 2021.

[34] R. Bombardieri, R. Cavallaro, R. Sanchez, and N. R. Gauger. Aerostructural wing shape optimization assisted by algorithmic differentiation. *Structural and Multidisciplinary Optimization*, 64:739–760, 2021.

[35] T. Albring, M. Sagebaum, and N. R. Gauger. Development of a consistent discrete adjoint solver in an evolving aerodynamic design framework. 16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2015. ISBN 978-1-62410-368-1.

[36] R. Sanchez, T. Albring, R. Palacios, N. R. Gauger, T. D. Economon, and J. J. Alonso. Coupled adjoint-based sensitivities in large-displacement fluid-structure interaction using algorithmic differentiation. *International Journal for Numerical Methods in Engineering*, 113(7):1081–1107, 2018.

[37] L. Scalia, A. Cini, and R. Cavallaro. Development of a fe code for adjoint-based coupled aerostructural optimisation. In *AeroBest 2023 – II ECCOMAS Thematic Conference on Multidisciplinary Design Optimization of Aerospace Systems*. Lisbon, 19-21 July 2023, Portugal, 2023. ISBN –.

[38] J. Martins and N. M. Poon. On structural optimization using constraint aggregation. In *VI World Congress on Structural and Multidisciplinary Optimization WCSMO6, Rio de Janeiro, Brasil*. Citeseer, 2005.

[39] A. H. Gebremedhin and A. Walther. An introduction to algorithmic differentiation. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10 (1):e1334, 2020.

[40] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.

[41] J. R. Martins and J. T. Hwang. Review and unification of methods for computing derivatives of multidisciplinary computational models. *AIAA journal*, 51(11):2582–2599, 2013.

[42] P. Gomes and R. Palacios. Pitfalls of discrete adjoint fixed-points based on algorithmic differentiation. *AIAA Journal*, 60(2):1251–1256, 2022.

[43] J. Roman, C. Campos, E. Romero, and A. Tomás. Slepc users manual–scalable library for eigenvalue problem computations. Technical report, Tech. Report 3.14, Universitat Polit ‘ecnica de Valencia, Spain, 2016.

[44] M. B. Giles. Collected matrix derivative results for forward and reverse mode algorithmic differentiation. In *Advances in Automatic Differentiation*, pages 35–44. Springer, 2008.

[45] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R. H. Castain, D. J. Daniel, R. L. Graham, and T. S. Woodall. Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, pages 97–104, Budapest, Hungary, September 2004.

[46] T. Belytschko, W. K. Liu, B. Moran, and K. Elkhodary. *Nonlinear finite elements for continua and structures*. John wiley & sons, 2014.

[47] P. Amestoy, I. S. Duff, J. Koster, and J.-Y. L'Excellent. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications*, 23(1):15–41, 2001.

[48] P. Amestoy, A. Buttari, J.-Y. L'Excellent, and T. Mary. Performance and Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures. *ACM Transactions on Mathematical Software*, 45:2:1–2:26, 2019.

[49] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. M. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, J. Faibussowitsch, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. C. McInnes, R. T. Mills, L. Mitchell, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang. PETSc Web page. `https://petsc.org/`, 2024. URL `https://petsc.org/`.

[50] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, J. Faibussowitsch, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. C. McInnes, R. T. Mills, L. Mitchell, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, H. Suh, S. Zampini, H. Zhang, H. Zhang, and J. Zhang. PETSc/TAO users manual. Technical Report ANL-21/39 - Revision 3.22, Argonne National Laboratory, 2024.

[51] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

[52] V. Hernandez, J. E. Roman, and V. Vidal. SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362, 2005.

[53] R. H. Macneal. A simple quadrilateral shell element. *Computers & Structures*, 8(2): 175–183, 1978. ISSN 0045-7949. doi:https://doi.org/10.1016/0045-7949(78)90020-2.

[54] K. Kim, G. Lomboy, and S. Han. A co-rotational 8-node assumed strain shell element for postbuckling analysis of laminated composite plates and shells. *Computational Mechanics*, 30:330–342, 2003.

[55] K.-D. Kim, C.-S. Lee, and S.-C. Han. A 4-node co-rotational ans shell element for laminated composite structures. *Composite Structures*, 80(2):234–252, 2007.

[56] J. Argyris. An excursion into large rotations. *Computer methods in applied mechanics and engineering*, 32(1-3):85–155, 1982.

[57] C. C. Rankin and F. A. Brogan. An Element Independent Corotational Procedure for the Treatment of Large Rotations. *Journal of Pressure Vessel Technology*, 108 (2):165–174, 05 1986. ISSN 0094-9930. doi:10.1115/1.3264765. URL `https://doi.org/10.1115/1.3264765`.

[58] R. Levy and W. Spillers. *Analysis of geometrically nonlinear structures*. Number v. 1. Kluwer Academic Publishers, Dordrecht, Netherlands, 2003. ISBN 9781402016547.

[59] *Element Library Reference*. Siemens Industry Software, Granite Park One 5800 Granite Parkway Suite 600 Plano, TX 75024 USA.

[60] F. Palacios, T. D. Economon, A. D. Wendorff, and J. J. Alonso. Large-scale aircraft design using su2. In *53rd AIAA Aerospace Sciences Meeting*, 2015.

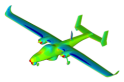
[61] M. Pini, S. Vitale, P. Colonna, G. Gori, A. Guardone, T. Economon, J. Alonso, and F. Palacios. Su2: the open-source software for non-ideal compressible flows. *Journal of Physics: Conference Series*, 821(1):012013, 2017.

[62] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL `http://www.scipy.org/`.

[63] A. C. Gray and J. R. R. A. Martins. A proposed benchmark model for practical aerostructural optimization of aircraft wings. In *AIAA SciTech Forum*, Orlando, FL, January 2024. doi:10.2514/6.2024-2775.

[64] D. Kraft. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, DLR German Aerospace Center – Institute for Flight Mechanics, Koln, Germany., 1988.

[65] Multidisciplinary design optimization laboratory at the university of michigan. URL `https://mdolab.engin.umich.edu/`.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

HIGH ASPECT-RATIO COMPOSITE WING AEROSTRUCTURAL OPTIMISATION OF A SHORT-MEDIUM RANGE AIRCRAFT

Luís M. M. Pacheco[1,2*], Fabian Volle[3], Ögmundur Petersson[2] and André C. Marta[1*]

1:IDMEC
Instituto Superior Técnico, Universidade de Lisboa
Av. Rovisco Pais, 1, 1049-001 Lisboa, Portugal
* {luis.miguel.martins.pacheco, andre.marta}@tecnico.ulisboa.pt,
https://mdo.tecnico.ulisboa.pt

2: Airbus Defence and Space
Department of Stress Methods and Optimisation
Rechliner Straße, 85077, Manching, Germany

3: DLR
Institute of Aerodynamics and Flow Technology
Lilienthalplatz 7, 38108, Braunschweig, Germany

Abstract. *A key pathway towards climate-neutral aviation lies in the adoption of high aspect-ratio wings, which reduce lift-induced drag. A recently developed high-fidelity multidisciplinary design optimisation framework, incorporating DLR's aerodynamic TAU solver and Airbus' structural Lagrange solver, is used for the aeroelastic structural sizing optimisation of two high aspect-ratio composite medium-range transport aircraft wings. Three distinct objective functions are considered: i) the classical minimisation of mass, ii) the maximisation of aerodynamic efficiency, and iii) the maximisation of Breguet range. A gradient-based algorithm with direct sensitivity analysis is used. The design variables include structural sizing parameters such as the thickness or cross-sectional area of the skin, spars, and stringers. Constraints reflect industry requirements, encompassing structural strength, buckling stability, and manufacturing criteria. For the objective functions and design points analysed, the findings highlight the advantages of high aspect-ratio wings in reducing drag and improving aerodynamic efficiency, despite the increase in structural weight. Optimising for Breguet range results in a trade-off between structural and aerodynamic efficiency, demonstrating the benefits of considering a multidisciplinary performance-based objective for structural sizing in the preliminary design phase.*

Keywords: Multidisciplinary Design Optimisation, Structural Sizing, Composite High Aspect-Ratio Wings, Breguet Range, Aerodynamic Efficiency

1 INTRODUCTION

The environmental impact of the aviation industry has become an increasingly pressing concern, particularly in the context of climate change. This problem is poised to intensify, as the current pace of technological advancements fails to offset the rise in emissions caused by air traffic growth [1]. To address this challenge, the development of novel, environmentally sustainable aircraft configurations is imperative.

Among various proposed designs, increasing the aspect ratio (AR) of wings has emerged as a particularly effective concept to improve aircraft efficiency. This reduces induced drag, which accounts for about 40% of total drag in cruise for transport aircraft [2]. However, it comes with a structural weight penalty due to higher wing root bending moments.

This gives rise to a tightly coupled and conflicting set of design objectives across the aerodynamic and structural domains. As such, a multidisciplinary design optimisation (MDO) approach becomes essential to explore these inherent trade-offs in a highly automated manner and identify the optimal balance between AR and structural integrity that maximises performance.

These trade-offs have been explored with a combined shape parameter study and sizing optimisation of a composite-wing UAV [3] to determine the optimal aspect ratio and structural sizing for maximum range. Five UAV variants, with aspect ratios ranging from 12 to 24, were structurally sized using a gradient-based aeroelastic optimisation. The results indicated that an aspect ratio of 16.8 was optimal, highlighting the fact that the performance benefits of reducing induced drag by increasing aspect ratio are eventually offset by the additional mass required to satisfy structural constraints.

The influence of metallic and composite materials, as well as the choice of objective function, on the trade-off between weight and drag has also been investigated [4], using gradient-based aerostructural optimisations considering aerodynamic shape and structural sizing, with strength and buckling constraints. The resulting Pareto front of fuel burn and take-off gross weight showed distinct optimal planforms, with fuel burn minimisation favouring higher aspect ratios and span. Furthermore, composite wings were found to be up to 40% lighter, resulting in fuel burn savings of up to 8%. These findings highlight the benefits of using composite materials for greener aircraft.

The present study pursues a twofold objective. Firstly, it aims to investigate the impact of increasing the aspect ratio on the performance of a transport aircraft. To this end, two variants of the DLR-F25 [5, 6] with different aspect ratios were structurally sized using a gradient-based algorithm with direct sensitivity analysis. The optimisations considered strength, buckling, and manufacturing constraints, reflecting industry requirements. Secondly, the influence of different objective functions on the structural sizing is explored, namely: minimisation of weight, maximisation of lift-to-drag ratio, and maximisation of the Breguet range. The objective is not only to compare the final designs but also to assess whether the use of computational fluid dynamics (CFD), necessary to compute the latter two objectives, is justified during the preliminary design stage.

The remainder of this paper is structured as follows: first, a brief description of the aerostructural optimisation process and disciplines is done. Subsequently, both DLR-F25 variants and the analysis models used in this study, are introduced. The results of the structural sizing for the different objective functions are then presented, followed by a discussion of the impact of AR and objective function on overall performance. Finally, the conclusions of the study are drawn, and suggestions for future research are provided.

2 METHODOLOGY

Three distinct objectives are proposed for the structural sizing optimisations: i) the classical minimisation of mass, ii) the maximisation of lift-to-drag ratio, and iii) the maximisation of Breguet range. In all optimisations, an aerostructural analysis based on linear aerodynamics is used to determine the loads for the structural sizing. Potential flow methods are chosen due to their sufficient accuracy in lift load distribution estimates in the preliminary design stage at a relatively low computational cost [7]. Moreover, they exhibit robustness to the large wingtip displacements induced by structurally critical loads, a feature often lacking in higher-fidelity models.

While the mass objective can be directly obtained from the structural model during the structural sizing process without the need for additional calculations, the other two objectives require additional computations to estimate drag. Since drag can only be reliably predicted with high-fidelity aerodynamics, a distinct aerostructural analysis model from that used to drive the structural sizing is required in these cases.

As a result, two distinct optimisation processes are put in place, whose XDSM diagrams [8] are presented in Figures 1 and 2. The first focuses on the mass objective and relies exclusively on the linear aerostructural analysis, while the second addresses the drag-dependent objectives by incorporating an additional expensive high-fidelity aerostructural analysis.

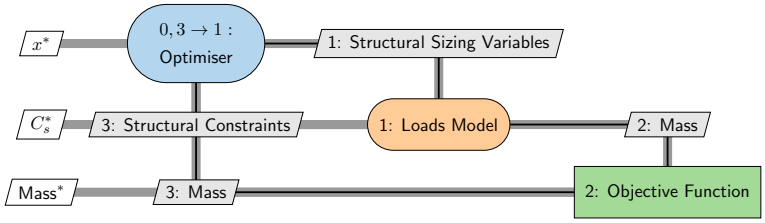


Figure 1: Structural sizing optimisation process for mass minimisation.

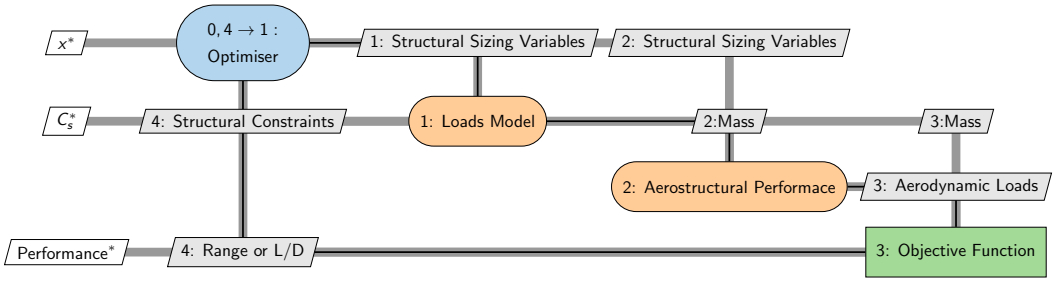


Figure 2: Structural sizing optimisation process for lift-to-drag or Breguet range maximisation.

All optimisations are conducted using a high-fidelity multidisciplinary design optimisation framework [9–11]. This framework couples Airbus's in-house MDO suite, Lagrange [12, 13], with DLR's flow solver, TAU [14, 15], through the high-performance computing software integration platform, FlowSimulator [16].

These two aerostructural design processes include two major analysis modules: the

loads model (1), which directly drives the structural sizing, and the higher-fidelity counterpart (2), needed to evaluate the drag-dependent objective functions, described next.

2.1 Loads Model

The aerostructural coupling is handled in the form of a multidisicplinary analysis (MDA), posed as an iterative process that includes the fluid flow solution, load interpolation, structural deformation, and displacement transfer, as illustrated in Figure 3.

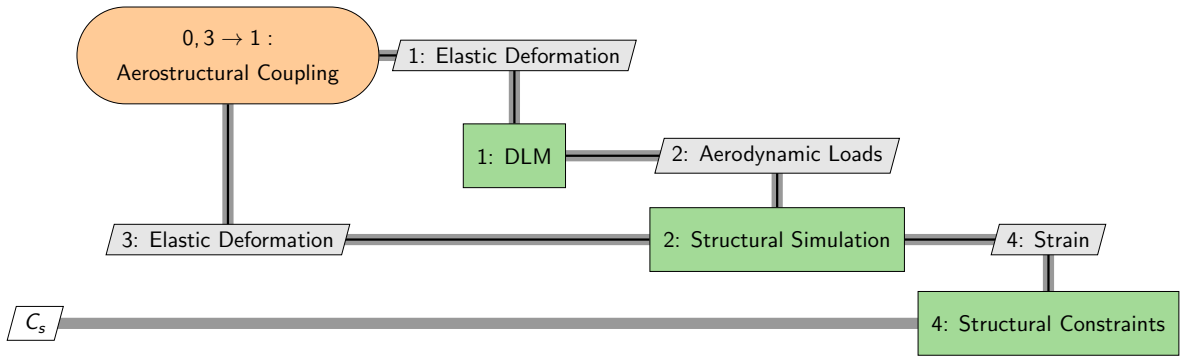


Figure 3: Coupled aerostructural analysis driving the loads model.

The process starts with the linear aerodynamic solution. The resulting loads are interpolated to the structural model using the infinite plate spline method, outlined in Section 2.5. Using these aerodynamic forces and adding the inertial loads, the structural problem is solved, yielding the elastic displacement field. These structural deformations are applied to the aerodynamic model, altering the lift distribution on the wing and starting over the iterative loop. After the coupling has converged, the structural constraints are evaluated.

2.2 Aerostructural Performance

The high-fidelity aerostructural analysis needed to evaluate the drag-dependent objective functions is based on a three-field formulation of the coupled aerostructural problem, where the mesh deformation is incorporated as an additional discipline alongside aerodynamic and structural analysis, as illustrated in Figure 4.

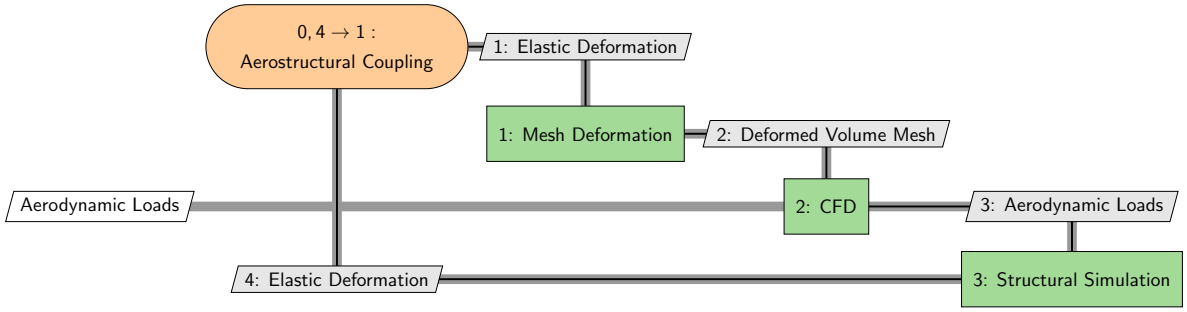


Figure 4: Coupled aerostructural analysis for the drag-dependent objective functions.

The analysis begins with a CFD simulation, solving the governing flow equations. The aerodynamic forces are then interpolated from the aerodynamic domain onto the structural mesh using the moving least squares method, outlined in Section 2.5. Considering these aerodynamic forces and the inertial loads, the structural deformations are computed and interpolated onto the aerodynamic surface mesh. These deformations are applied as boundary conditions in the mesh deformation problem, which is solved using a linear elasticity analogy method, presented in Section 2.6. The outcome of this aerostructural analysis is the aerodynamic loads used to compute the lift-to-drag and Breguet range objective functions.

The convergence criterion of this coupling process is defined by the L^2 norm of the variation in elastic deformation at the fluid-solid interface. The rate of convergence for the aerostructural coupling is improved through a dynamic under-relaxation technique using the standard Aitken Δ^2 method [17].

These two processes share the same structural model but use different aerodynamic models, as described next.

2.3 Structural Model

In both processes, the structural problem is governed by Hooke's law of elasticity. The structural residuum R_S can be expressed as

$$\mathcal{R}_S = K y_S - F_S = 0\,, \tag{1}$$

where K is the symmetric stiffness matrix, y_S is the state variables vector representing the structural displacements, and F_S is the sum of aerodynamic and inertia forces acting on the aircraft. The structural problem is solved using the built-in finite element solver of Airbus's MDO suite Lagrange [12, 13].

2.4 Aerodynamic Models

Two levels of fidelity are used for the aerodynamic discipline: high-fidelity aerodynamics is used to calculate the loads needed to evaluate the drag-dependent objective functions, whereas a low-fidelity linear aerodynamic model is used for the structural sizing process.

The linear aerodynamic model is the doublet lattice method (DLM) [18]. In this method, the aerodynamic system is solved once in advance to generate the complex-valued aerodynamic influence coefficient (AIC) matrix. This matrix is then provided to Lagrange, which incorporates a linear aerodynamic analysis tool capable of computing the aerodynamic loads acting on the structure [12].

In the high-fidelity model, the fluid problem is governed by the compressible Reynolds-averaged Navier-Stokes (RANS) equations, coupled with the Spalart-Allmaras (SA) one equation turbulence model in its negative formulation. In differential form, the residuals of the flow governing equations are expressed as

$$\mathcal{R}_F = \frac{\partial y_F}{\partial t} + \nabla \cdot (\phi_c - \phi_v) = 0\,, \tag{2}$$

where ϕ_c and ϕ_v are the convective and diffusive fluxes, respectively, and the state variables y_F denote the conserved quantities of the flow. These equations are discretised with the finite volume method and solved using DLR's CFD solver TAU [14, 15].

2.5 Loads And Displacement Transfer

Due to the distinct domain discretisations, the meshes at the fluid-structure interface typically do not match, preventing direct information exchange.

For the aerostructural analysis of the loads model, aerodynamic loads and elastic displacements are interpolated through a set of structural nodes using the infinite plate spline (IPS) method [19]. For the aerostructural analysis used to calculate the drag-dependent objectives, a mesh-free approach based on the moving least squares (MLS) method [20] is used. This method to transfer forces and displacements between the structural and aerodynamic boundary meshes is both conservative and consistent.

2.6 Aerodynamic Mesh Deformation

When the DLM method is used in the loads model, no volume mesh deformation is required. In contrast, the aerostructural analysis used to calculate the drag-dependent objectives requires that the aerodynamic volume mesh must adapt to the deformation of its surface mesh induced by the structural displacement. The mesh deformation method used is based on the linear elasticity analogy [21], whereby the fluid flow mesh is considered analogous to a volumetric structure problem. The governing equation for the mesh deformation problem is given by

$$\mathcal{R}_M = K_M y_M - F_M(u) = 0\,, \tag{3}$$

where K_M is a symmetric stiffness matrix constructed by assigning stiffnesses to each element of the fluid flow mesh, inversely proportional to the element's volume, and F_M is a fictitious force imposing the Dirichlet boundary condition of the structural displacements at the aerodynamic surface mesh.

3 DLR-F25 MODELS

In order to investigate the impact of increasing the aspect ratio on aircraft performance, two aircraft models were structurally sized: the high aspect-ratio DLR-F25 and a newly developed variant with an even greater aspect ratio.

3.1 Baseline Geometric Model

The DLR-F25 is a single-aisle, narrow-body aircraft model with a high aspect-ratio wing, designed for the short-to-medium range market segment [5, 6]. It has been primarily developed by the DLR and it is based on the Airbus A321neo. This research aircraft model is currently used in the UPWing project [22] as a reference platform for assessing and validating innovative technologies.

The baseline DLR-F25 wing planform is divided by five sections - centre, root, kink, mid, and tip - as depicted in Figure 5.

The key characteristics of the baseline DLR-F25 wing are summarised in Table 1. Its aspect ratio of 15.6 is significantly higher than that of conventional transport aircraft, such as the Airbus A321-100, which has an aspect ratio of 9.4 [23]. The DLR-F25 wing has a tip chord of 0.6 metres and a taper ratio of 0.12, resulting in a considerably thinner and narrower wing compared to a typical transport aircraft, such as those of the Airbus A320 family, whose tip chord measures approximately 1.5 metres and have taper ratios of 0.24 [23, 24].

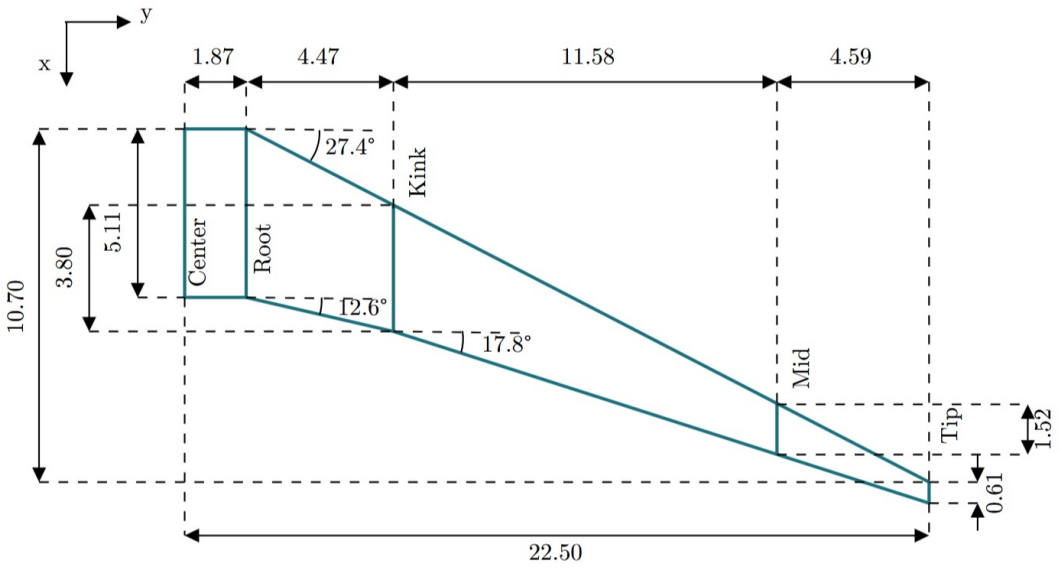


Figure 5: DLR-F25 wing planform (dimensions in metres) [6].

The DLR-F25 features a carry-through wingbox extending through the fuselage. The wing comprises a two-spar design with 31 ribs per half-span and 11 stringers, which progressively taper as the wing narrows towards the tip. The ribs and stringers have a minimum pitch of 800 mm and 220 mm, respectively, to prevent buckling from becoming a dominant design constraint.

3.2 Higher Aspect-Ratio Variant Geometric Model

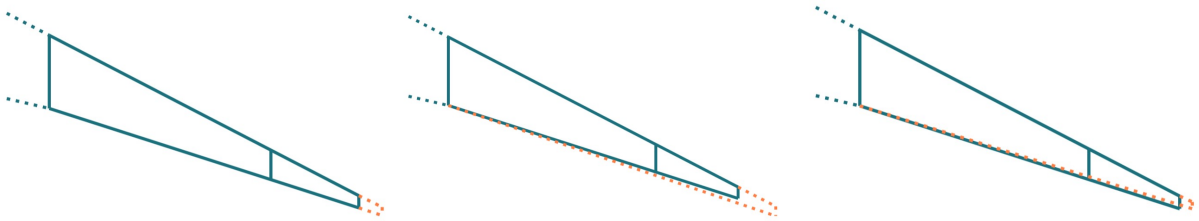
The DLR-F25 baseline aircraft was modelled using a parametric geometry representation in CPACS [25, 26], wherein the positions of structural components were defined relative to the wing's coordinate system. This approach enables the automatic adjustment of component positions in response to shape variations, such as an increase in wingspan.

Leveraging that parametric geometry, a higher aspect-ratio variant of the DLR-F25 was generated by modifying the baseline wing shape using Airbus's in-house tool Descartes [27, 28]. Descartes is a pre-processing tool capable of generating a parametric geometry model, from which it can derive the necessary input data for structural sizing, including structural and aerodynamic models, as well as the optimisation model itself.

In the higher aspect-ratio variant, it was decided that the central part of the wing, up to the kink section, remained unaltered to preserve the wing-fuselage junction and the original pylon attachment configuration. Consequently, only the outer wing sections, aft of the engine position, were modified.

Several methods exist in the literature for increasing an aircraft's aspect ratio, involving changing or preserving parameters such as taper ratio, trailing and leading-edge sweep, or wing surface area, depending on the design objectives [6]. Some of these methods are illustrated in Figure 6.

For this study, the area method depicted in Figure 6b was adopted. This method maintains the leading-edge sweep and taper ratio constant but increases the surface area and trailing-edge sweep. This approach was selected to prevent unintended aerodynamic effects that could arise from varying the taper ratio, which might obscure or counteract the impact of increasing the aspect ratio. Maintaining the baseline taper ratio was particularly



(a) Span variation. (b) Area variation. (c) Tip chord variation.

Figure 6: Schematic representation of different aspect-ratio variation methods [6].

critical, given that the DLR-F25's taper ratio was already small, and further reduction would pose significant aerodynamic challenges at the wingtip. Furthermore, this method has been shown to provide a longitudinally stable aircraft [6] and, among the methods shown in Figure 6, it is the most effective in reducing induced drag.

Using this approach, the new higher aspect-ratio DLR-F25 variant was generated in Descartes by stretching the baseline´s outer wing segments, increasing the overall wingspan by 10%. This resulted in a new aircraft model with a 11.5% higher aspect ratio of 17.4 and a 7.5% larger wing surface area. The vertical and horizontal tailplanes remained unchanged and the maximum take-off weight was considered a top-level aircraft requirement and, therefore, kept constant. Henceforth, the baseline DLR-F25 and its higher aspect-ratio variant will be referred to as F25-AR15 and F25-AR17, respectively. Table 1 summarises the main geometrical differences between the model's wings.

Table 1: DLR-F25 wing key characteristics: baseline and higher aspect-ratio wings.

Parameter		**F25-AR15**	**F25-AR17**
Aspect ratio	AR	15.6	17.4
Wingspan	b	44.60 m	49.04 m
Wing area	S	129.59 m^2	139.21 m^2
Sweep at $\frac{1}{4}$ chord	$\Lambda_{c/4}$	24.43°	24.69°
Taper ratio	λ	0.12	0.12

The substantial increase in aspect ratio for the F25-AR17 altered the baseline rib pitch. This resulted in significantly larger buckling fields for the F25-AR17, making buckling a critical design constraint and precluding a fair direct comparison between the two aspect ratio variants. Consequently a topological modification to the F25-AR17 wing structure was performed, where two additional ribs were incorporated into the middle section of the wing and one near its tip.

3.3 Structural Analysis Models

The finite element (FE) structural models were generated using Descartes' internal meshing functionality. The skins, spars, and ribs, were modelled using shell elements, specifically CQUAD4 and CTRIA3, whilst the stringers and spar caps were represented using one-dimensional elements, namely CBAR and CROD, respectively. The FE structural model for the F25-AR15 is depicted in Figure 7.

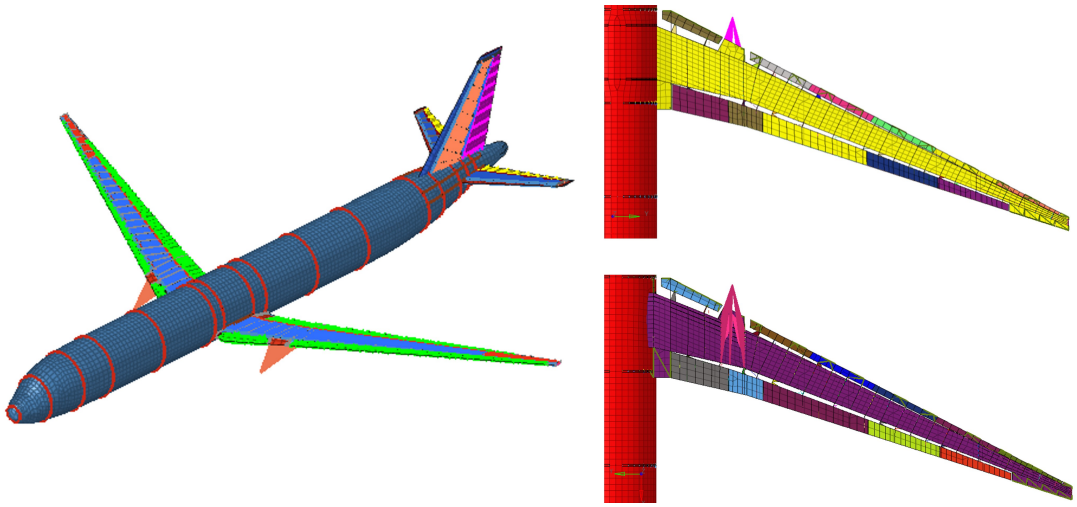


Figure 7: FE structural model of the F25-AR15 variant.

The F25-AR17 model follows a comparable discretisation approach, as both models were constructed using the same methodology. Table 2 presents a comparative summary of the refinement and discretisation used in the structural modelling of each aircraft variant.

Table 2: Structural discretization and elements of F25-AR15 and F25-AR17 wings.

Parameter	**F25-AR15**	**F25-AR17**
Nodes	13,152	13,331
CQUAD4 elements	13,567	13,815
CTRIA3 elements	1,191	1,201
CROD elements	5,812	5,906
CBAR elements	3,175	3,331

The DLR-F25 model considers two distinct mass configurations: the maximum take-off weight (MTOW) of 81,656 kg and the maximum zero fuel weight (MZFW) of 69,322 kg. The MTOW configuration serves as the driver for structural sizing, whereas the MZFW configuration is used to determine the aerodynamic loads required for the drag-dependent objective functions. Irrespective of the mass configuration, the aircraft weight is divided into two components: the structural weight of the composite wing and the combined weight of the fuel, payload, passengers, and the remaining aircraft structure, including engines and pylons.

The weight of the composite wing is evaluated from the size and material properties of its finite elements. The wing skin and spar webs are modelled using a symmetric and balanced 24-ply carbon fibre reinforced polymer (CFRP) laminate with four ply orientations (0°, 90°, and ±45°), and a material density of 1580 kg/m^3. The T-shaped stringers and spar caps are modelled using homogenised CFRP properties, assuming a ply distribution of 70% at 0°, 20% at ±45°, and 10% at 90°, with a density of 1750 kg/m^3.

The remaining weight is represented by 131 distributed concentrated mass points, implemented as CONM2 elements and connected to the FE model via rigid body elements

(RBE), as illustrated in Figure 8.

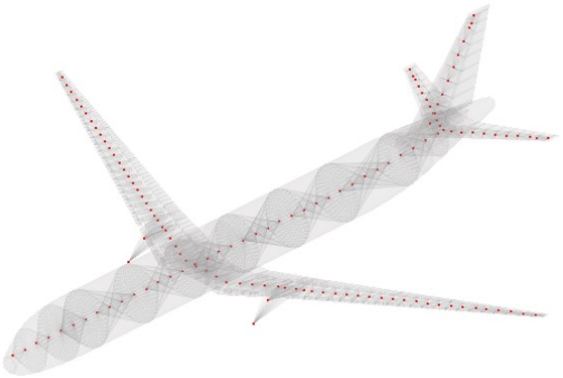
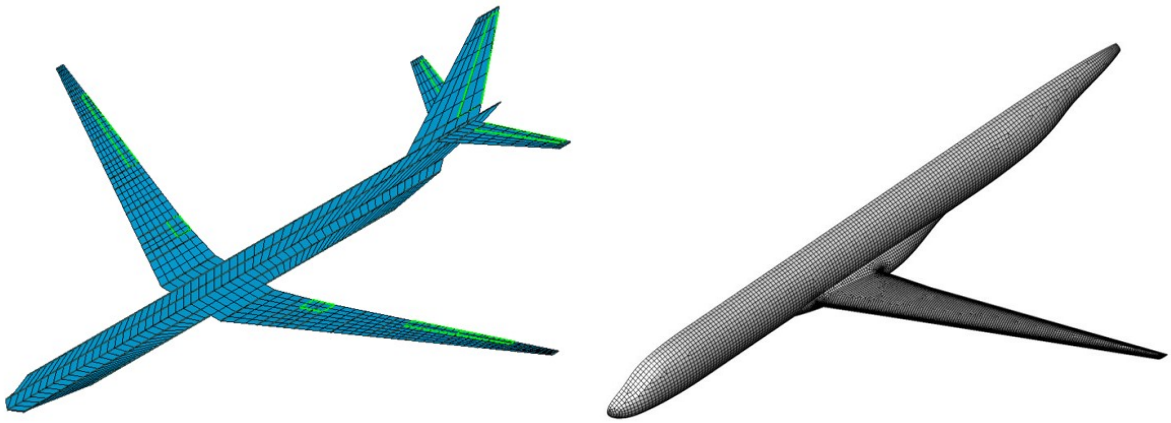


Figure 8: F25-AR15 CONM2 mass point distribution.

3.4 Aerodynamic Models

As detailed in Section 2.4, this study employs two levels of fidelity for the aerodynamic discipline. High-fidelity aerodynamics is used for the evaluation of objective functions, whereas the structural sizing optimisation relies on a linear aerodynamic model.

The DLM meshes used in the loads model were generated from the parameterised geometric models using Descartes. The entire aircraft, including the wing, fuselage, and vertical and horizontal tailplanes, was discretised, as presented in Figure 9a for the F25-AR15. The wing was discretised into 7 panels in the chordwise direction and 43 panels in the spanwise direction. The discretisation for the F25-AR17 follows the same approach.



(a) Doublet lattice method mesh. (b) High-fidelity aerodynamic surface mesh.

Figure 9: F25-AR15 aerodynamic mesh.

The high-fidelity aerodynamic mesh used in the aerostructural performance analysis, is based on a wing-body configuration of the DLR-F25. To reduce the computational cost associated with CFD analyses, a half-model of the aircraft is used, leveraging the symmetry of the configuration, as illustrated in Figure 9b for the F25-AR15. This computational grid, comprising 1.02 million nodes, was developed by DLR. The aerodynamic model of

the F25-AR17 was generated by morphing the existing F25-AR15 mesh, therefore both meshes are topologically identical with the same number of nodes and elements.

3.5 Coupling Models

The IPS method used in the aerostructural analysis of the loads model requires a set of nodes. These nodes are strategically arranged in a diamond pattern at the intersections of the wing's upper skin with the spars and ribs. Beyond the wingbox, this approach is extended to the fuselage and the horizontal and vertical tailplanes. Figure 10 illustrates the nodal distribution for the F25-AR15 wing. The coupling model for the F25-AR17 follows the same principles.

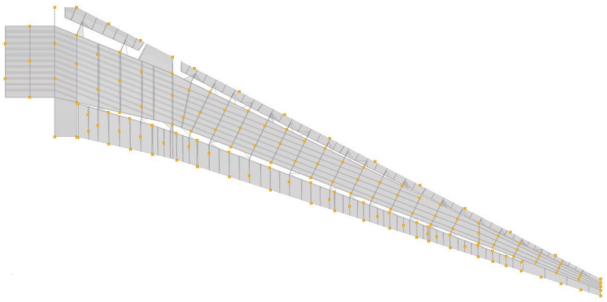


Figure 10: F25-AR15 wing coupling for the loads model.

In the aerostructural analysis used to compute the drag-dependent objectives, the MLS method was used, which does not rely on a predefined set of nodes.

3.6 Optimisation Problem

The three proposed objectives for the structural sizing optimisations are the minimisation of the wingbox mass, and the maximisation of lift-to-drag ratio and Breguet range.

All optimisations are performed using the gradient-based NLPQL optimisation algorithm [29] with a direct sensitivity analysis. The optimisations converge based on a Karush-Kuhn-Tucker (KKT) criterion [30] of 10^{-5}. An active-set strategy is implemented, restricting the sensitivity analysis to the 20,000 most violated constraints.

The optimisation focused exclusively on the structural sizing of the composite wingbox, specifically the skin, spars, and stringers. The design variables are defined within patches. For the skin, each patch is bounded by two ribs and two stringers, whereas patches for the stringers and spars are segmented by ribs. These patches are symmetrically linked to ensure both sides of the wing were identical. Within each skin and spar web patch, plies of the same orientation are linked to preserve laminate symmetry. To ensure a balanced laminate, the thicknesses of the $+45°$ and $-45°$ plies are also linked. Consequently, three independent design variables per skin and spar web patch control the thickness of the ply orientations. Each spar cap patch is governed by a single design variable that determines its cross-sectional area. The T-shaped stringers are characterised by three design variables, which define the web height, foot width and thickness. The thickness of the foot and web are linked in a 1:2 ratio. The bounds and initial values for the design variables are presented in Table 3.

The structural sizing criteria for the composite wing focus primarily on optimising the wing covers, as they are the main driver for aeroelastic tailoring. Constraints on strength,

Table 3: Design variables and their bounds for structural sizing.

Component	Section	Design Variable	Lower Bound	Initial Value	Upper Bound
Skin	Upper skin	Ply thickness [mm]	0.164	0.254	2
	Lower skin root		0.254	0.508	2
	Lower skin kink		0.254	1.016	2.5
	Lower skin tip		0.254	0.381	2
Spar	Inner & outer	Web ply thickness [mm]	0.164	0.254	2
		Cap area [mm^2]	80	100	1000
Stringer	Inner	Web height [mm]	51.5	66.5	81.5
		Foot length [mm]	80	100	120
		Foot thickness [mm]	3	3	9
	Outer	Web height [mm]	51.5	66.5	68.5
		Foot length [mm]	80	100	120
		Foot thickness [mm]	3	3	6

buckling, and manufacturability are incorporated into the sizing optimisations. Strength constraints are uniformly applied to the skin, spars, and stringers, with maximum allowable material strains of 5000 $\mu\varepsilon$ in tension and 3500 $\mu\varepsilon$ in compression and a safety factor of 1.5 incorporated. The skin and spar buckling panels are modelled as biaxially loaded, simply supported flat plates with anisotropic material properties. The critical buckling loads are determined using analytical methods [31–33]. Stringer buckling is evaluated by modelling the stringer and attached sheet as a super-stiffener, with the critical buckling strength determined via the Johnson-Euler formula [34]. A correction factor of 0.95 is applied to the skin buckling field size to account for the idealised representation of the stringers as one-dimensional elements. To ensure manufacturability, thickness and ply share constraints are imposed. Thickness variations between adjacent skin patches are limited by a ramp rate of 1:20 in the spanwise direction and 1:10 in the chordwise direction. Additionally, continuity constraints prevent abrupt ply share variations by restricting thickness differences between adjacent plies to 1/10 of the ply thickness. Ply share percentages in the skin are constrained to a range of 10 to 63% for 0° and 90° plies and 20 to 80% for ±45° plies. Additionally, a minimum thickness of 4 mm was enforced for the upper skin. On the lower skin, a minimum thickness of 20 mm was imposed in the pylon attachment area, with a 6 mm and 8 mm minimum thicknesses constraint outboard and inboard of this region.

3.7 Load Cases

Pull-up and push-over manoeuvres were defined to be structurally design driving. The MTOW configuration was used in the manoeuvres to generate the most critical aerostructural loads for the structural constraints. These load cases were trimmed to balance the pitching moment and ensure equilibrium between aerodynamic and inertial forces, with the angle of attack and elevator deflection serving as trimming variables. A summary of structurally design driving load cases is presented in Table 4.

The aerostructural analysis used to obtain the aerodynamic loads for the drag-dependent objectives (Section 2.2) was conducted under cruise conditions with the MZFW configu-

Table 4: Load cases considered in the structural optimisation.

Load case	**Load factor [g]**	**Mach**	**Altitude [m]**
Pull-up	2.5	0.81	11,000
Push-over	-1		

ration, at a Mach number of 0.78 and an altitude of 10,363 metres, corresponding to a Reynolds number of 22 million. The aircraft was trimmed to achieve a target lift force using a gradient-based approach with the angle of attack as the trimming variable.

4 RESULTS

Prior to the structural sizing optimisations, both F25-AR15 and F25-AR17 had generic uniform material thickness with an initial MTOW of 81,656 kg. As expected, neither variant satisfied the imposed criteria prior to structural sizing, with significant constraint violations, particularly near the wing root.

4.1 Structural Sizing For Mass Minimisation

Having infeasible baseline designs as a starting point, both the F25-AR15 and F25-AR17 underwent structural sizing for mass minimisation using the optimisation process depicted in Figure 1. The structural sizing reduced the maximum constraint violation by six orders of magnitude. Table 5 summarises the characteristics of interest for both variants following the convergence of the conventional structural sizing. The results of the F25-AR15 serve as reference for all following optimisations.

Table 5: Key characteristics of both variants after conventional structural sizing.

	F25-AR15	**F25-AR17**
MTOW [kg]	81,618	81,836
Normed L/D [-]	1.0000	1.0216
Normed root bending moment [-]	1.0000	1.0597

The MTOW for the optimised F25-AR17 is 218 kg (0.27%) heavier than that of the F25-AR15. This difference can be attributed to the additional material required due to the 7.5% larger wing surface area of the F25-AR17, as well as the increased thickness of structural components at the wing root necessary due to the 5.97% higher bending moments associated with the increased aspect ratio. Compared to the initial MTOW, the mass gains may not appear significant, with the MTOW of the F25-AR17 even slightly heavier than its baseline configuration. However, the optimisation ensures a design that satisfies all imposed constraints, a condition not met by the initial configurations.

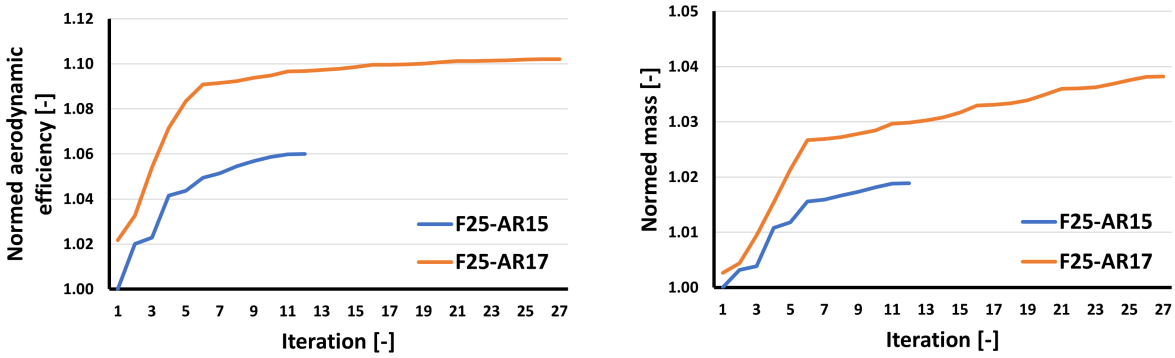
The F25-AR17 has a 2.16% higher aerodynamic efficiency, compared to the F25-AR15. This can be attributed to the reduction in drag associated with the higher aspect-ratio. In fact, even though the friction drag of the F25-AR17 is 4.43% higher than that of the F25-AR15, due to the higher aspect-ratio larger wetted wing area, its pressure drag decreased by 6.74%.

The structural sizing with mass minimisation as the objective has yielded feasible designs, which serve as the starting points for the subsequent two optimisation cases.

4.2 Structural Sizing For Aerodynamic Efficiency

Following the structural sizing for mass minimisation, two additional structural sizing optimisations with distinct objectives are conducted for each aspect ratio. The optimisation process illustrated in Figure 2, which integrates the aerostructural performance model outlined in Section 2.2 into the conventional structural sizing process, is used.

Aerodynamic efficiency is selected as one of the objective functions of interest to provide a contrast with the conventional structural sizing approach, which focuses on minimising structural mass. The convergence history for the structural sizing optimisations for aerodynamic efficiency is illustrated in Figure 11.



(a) Aerodynamic efficiency (b) Aircraft MTOW

Figure 11: Convergence history of the structural sizing for aerodynamic efficiency.

Starting from the mass-optimal designs, the optimisations converge quickly and smoothly. However, only the final iterations yielded feasible designs.

The aerodynamic efficiency of the F25-AR15 increases 6.00%, albeit at the expense of a 1.89% increase in MTOW. Similarly, the F25-AR17 aerodynamic efficiency improved 7.86% at the cost of a 3.54% increase in MTOW, compared to its mass-optimal design. These results highlight the distinct design characteristics that emerge when aerodynamic efficiency is prioritised over structural weight.

The improvements in aerodynamic efficiency, relative to the conventional structural sizing, are due to a reduction in pressure drag. This reduction arises from a shift in the lift distribution, with the centre of lift moving outboard by more than 10%, as illustrated in Figure 13.

As a consequence of this shift, the wing root bending moment increases by more than 10%, directly affecting the mass. In the conventional structural sizing, minimising wing root bending moments through passive load alleviation is crucial for mass reduction, however, when the design objective shifts towards aerodynamic performance, this consideration is deprioritised.

When comparing the two aspect ratio configurations, the F25-AR17 exhibits a 3.98% higher aerodynamic efficiency due to its lower drag. This emphasises the potential of higher aspect-ratio wings to improve aerodynamic performance at the investigated conditions.

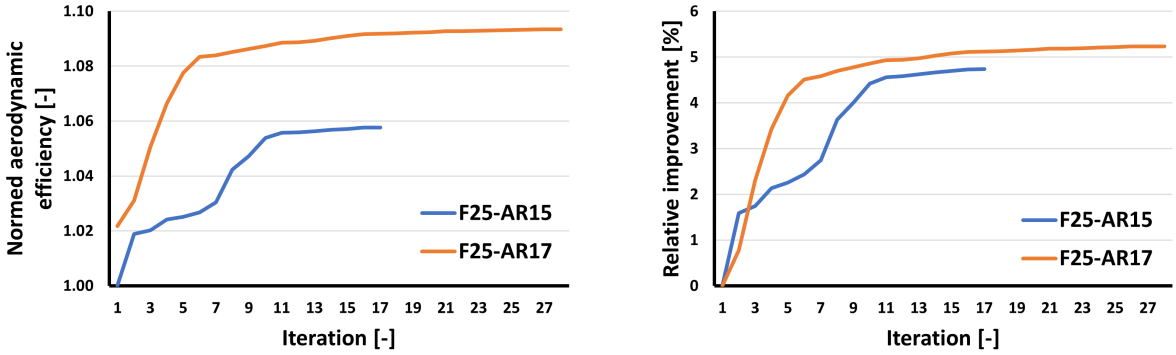
4.3 Structural Sizing For Breguet Range

Following the structural sizing optimisations for structural and aerodynamic performance metrics, Breguet range was selected as the final objective function. Breguet range is a comprehensive measure of overall aircraft performance, providing a trade-off between structural and aerodynamic efficiency in a physically meaningful manner. The Breguet range of the aircraft can be estimated for a simplified cruise segment as [35],

$$R_{Br} = \frac{aM}{\mathrm{TSFC}} \frac{L}{D} \ln\left(\frac{W_{\mathrm{initial}}}{W_{\mathrm{final}}}\right) , \tag{4}$$

where the Mach number M and the speed of sound a are functions of altitude and velocity; TSFC denotes the thrust-specific fuel consumption, which is determined by the propulsion system; the lift-to-drag ratio denotes the aerodynamic efficiency; and W_{initial} and W_{final} correspond to the total aircraft mass at the beginning and end of the cruise segment.

Figure 12 presents the convergence history for the structural sizing optimisation for Breguet range. The starting point for this optimisation was also the designs obtained from the conventional structural sizing.



(a) Aerodynamic efficiency (b) Breguet range improvement

Figure 12: Convergence history of the structural sizing for Breguet range.

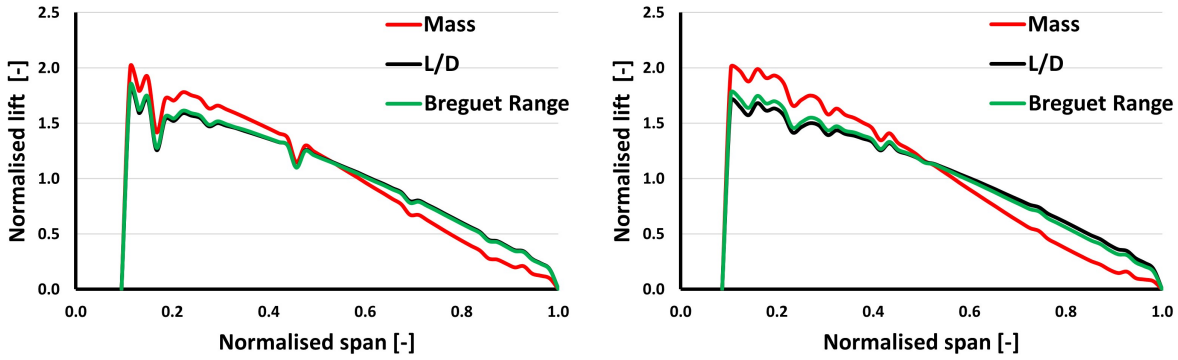
During the optimisations, the Breguet range improved by 4.73% and 5.23% for the F25-AR15 and F25-AR17, respectively. These optimised designs maintain an aerodynamic efficiency similar to that obtained for the L/D objective, with only a marginal difference of 0.23% and 0.80%, while simultaneously achieving a substantial mass reduction of 0.94% and 1.89% for the F25-AR15 and F25-AR17, respectively. The final masses of the designs optimised for Breguet range are between those obtained through structural sizing for minimal mass and those derived from aerodynamic efficiency objectives. The wing root bending moments are very similar compared to the optimisations for aerodynamic efficiency, as a result of the near-identical lift distributions.

The drag reduction associated with the higher aspect-ratio wing is also observed in the optimisations for Breguet range. Despite having a greater wetted surface area, the F25-AR17 exhibits a 3.38% lower drag, highlighting the aerodynamic advantages conferred by the higher aspect ratio in this context.

The optimisations for Breguet range took longer to convergence compared to the aerodynamic efficiency objective, due to the conflicting nature of aerodynamic and structural performance. The average runtime per iteration for the primal solution was approximately

1 hour and 20 minutes, with approximately half of this time allocated to the loads model and the other half to the high-fidelity aerostructural analysis. The sensitivity analysis required approximately 2 hours per iteration, again with an even distribution between the loads and the performance analysis model.

Figure 13 illustrates the lift distributions after the structural sizing optimisations considering mass, aerodynamic efficiency, and Breguet range as objectives.



(a) F25-AR15 (b) F25-AR17

Figure 13: Lift distributions for the structural sizing optimisations for mass, aerodynamic efficiency, and Breguet range.

All optimisation cases exhibit lift distributions that closely resemble the bell-shaped distribution proposed by Prandtl in 1933 as the aerostructurally optimal lift distribution for wing designs with variable span and constant weight [36]. For both aspect ratios, the structural sizing optimisations aimed at reducing finite element mass shift the aerodynamic loads inboard to reduce the wing root bending moment, which is one of the main contributors to structural mass in wing design, through passive load alleviation. The optimisations for aerodynamic efficiency and Breguet range produce lift distributions that are very similar to each other and start tending towards an elliptical profile, highlighting the prioritisation of drag reduction in these objectives.

Figure 14 presents a comparison between the F25-AR15 and F25-AR17 Pareto fronts. These Pareto fronts were estimated based on the three optimal design points, corresponding to the minimisation of structural mass, maximisation of aerodynamic efficiency, and the optimisation of Breguet range.

When comparing the results of the first two optimisation cases, which focused on a single discipline objective, with those obtained by optimising for the Breguet range, it becomes evident that the latter achieves a better trade-off between structural and aerodynamic efficiency, thereby resulting in superior overall aircraft performance in the investigated design points.

The F25-AR15 and F25-AR17 offer markedly different characteristics in terms of mass and lift-to-drag ratio. For short-range missions, the F25-AR15 configuration is more favourable due to its lower structural weight. Conversely, for longer missions, the F25-AR17 offers superior performance as a result to its improved aerodynamic efficiency.

The optimisations presented in this work do not take into account constraints imposed by overall aircraft design or by disciplines beyond structural mechanics. Nonetheless, the results presented herein may serve as a valuable reference for aircraft designers in

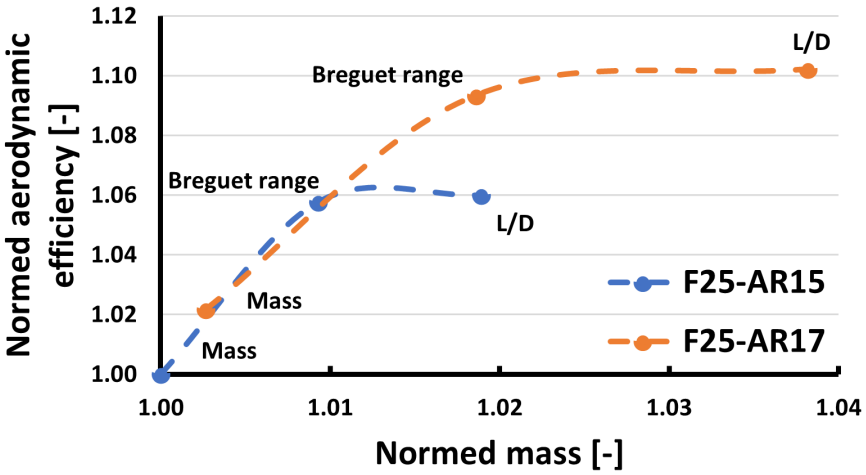


Figure 14: Pareto fronts for the F25-AR15 and F25-AR17.

informing the early stages of the design process.

5 CONCLUSIONS

This study investigated the aeroelastic structural sizing optimisation of two high aspect-ratio composite wings, using three distinct objective functions. The first two objectives -the minimisation of structural mass and the maximisation of aerodynamic efficiency - focused on a single discipline, whilst the third - the maximisation of the Breguet range - provided a multidisciplinary trade-off between the structural and aerodynamic domains. The optimisations were conducted using a high-fidelity multidisciplinary design optimisation framework, integrating DLR's CFD solver TAU and Airbus' structural solver Lagrange. A gradient-based algorithm with direct sensitivity analysis was used. The design variables, which are in the order of the 3,000, included structural sizing parameters, such as the thickness or cross-sectional area of the skin, spars, and stringers, while the constraints encompassed structural strength, buckling stability, and manufacturing criteria.

The results capture the effects of increasing the aspect ratio in modern transport aircraft, particularly the reduction in drag and consequent increase in aerodynamic efficiency, and the associated increase in structural weight. These findings may assist designers in evaluating the impact of aspect ratio on overall aircraft performance. Furthermore, the study demonstrates that using a multidisciplinary performance-based objective such as the Breguet range in the structural sizing optimisations during the preliminary design stage facilitates a trade-off between structural and aerodynamic efficiency.

Future work should expand the current shape parameter study to more aspect ratios, exploiting the opportunity to run these optimisations in parallel. Furthermore, shape variables, such as jig twist and aerofoil geometry, should be integrated into the design space. Moreover, a broader range of load cases, such as gust encounters and aeroelastic phenomena like flutter should be added to yield more realistic and robust design outcomes.

ACKNOWLEDGEMENTS

The project Ultra Performance Wing (UP Wing, project number: 101101974) is supported by the Clean Aviation Joint Undertaking and its members. Co-Funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or Clean Aviation Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

[1] Clean Aviation Joint Undertaking. Strategic Research and Innovation Agenda. Technical Report CAJU-GB-2021-12-16-SRIA, European Commision, December 2021.

[2] I. Kroo. Drag due to lift: concepts for prediction and reduction. *Annual review of fluid mechanics*, 33:587–617, 01 2001. doi:10.1146/annurev.fluid.33.1.587.

[3] F. Daoud, S. Deinert, R. Maierl, and Ö. Petersson. Integrated multidisciplinary aircraft design process supported by a decentral MDO framework. In *16th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Dallas, Texas, June 2015. doi:10.2514/6.2015-3090.

[4] G. J. Kennedy and J. R. R. A. Martins. A Comparison of Metallic and Composite Aircraft Wings Using Aerostructural Design Optimization. In *14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Indianapolis, Indiana, September 2012. doi:10.2514/6.2012-5475.

[5] S. Wöhler, J. Häßy and V. Kriewall. Establishing the DLR-F25 as a research baseline aircraft for the short-medium range market in 2035. In *34th Congress of the International Council of the Aeronautical Sciences,*. Florence, Italy, September 2024.

[6] M. Schulze, F. Müller, S. Cumnuantip, V. Handojo, T. Hecken, T. Klimmek, M. Ritter, and M. Zimmer. The Effect of Aspect Ratio Variation on the Preliminary Aeroelastic Assessment of a Mid-Range Transport Aircraft. In *International Forum on Aeroelasticity and Structural Dynamics*, The Hague, Netherlands, June 2024.

[7] S. Deinert. *Shape and Sizing Optimization of Aircraft Structures with Aeroelastic and Induced Drag Requirements*. PhD thesis, Technische Universität München, 2016.

[8] A. B. Lambe and J. R. R. A. Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Structural and Multidisciplinary Optimization*, 46:273–284, 2012. doi:10.1007/s00158-012-0763-y.

[9] F. Volle. Aero-Structural Optimization of the AIAA MDO Benchmark Model With TAU and Lagrange. In *AIAA SCITECH 2025 Forum*, Orlando, Florida, January 2025. doi:10.2514/6.2025-2815.

[10] F. Volle, M. Abu-Zurayk, A. U. Balani, S. Deinert and Ö. Petersson. Simultaneous Shape and Sizing Optimization for Improving Aircraft Performance Using the Aero-Structure Adjoint. In *Deutsche Luft- und Raumfahrtkongress - DLRK*, Stuttgart, Germany, September 2023.

[11] A. A. Gastaldi, F. Volle, F. Daoud, C. Breitsamter. Simultaneous Airframe Shape and Sizing Optimization Using FlowSimulator and the Lagrange MDO Suite. In *Deutsche Luft- und Raumfahrtkongress - DLRK*, Dresden, Germany, September 2022.

[12] G. Schuhmacher, F. Daoud, Ö. Petersson and M. Wagner. Multidisciplinary Airframe Design Optimization. In *28th International Congress of the Aeronautical Sciences*. Brisbane, Australia, September 2012.

[13] F. Daoud, Ö. Petersson, S. Deinert and P. Bronny. Multidisciplinary Airframe Design Process: Incorporation of steady and unsteady aeroelastic loads. In *12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conferences*. Indianapolis, Indiana, September 2012.

[14] D. Schwamborn, T. Gerhold, R. Heinrich. The DLR TAU-code: recent applications in research and industry. In *ECCOMAS CFD Conference*, Delft, Netherlands, 2006.

[15] S. Langer, A. Schwöppe, and N. Kroll. The DLR Flow Solver TAU - Status and Recent Algorithmic Developments. In *52nd Aerospace Sciences Meeting*. Maryland, USA, January 2014.

[16] M. Meinel and G. Ó. Einarsson. The FlowSimulator framework for massively parallel CFD applications. In *Applied parallel and scientific computing 10th International Conference*. Reykjavik, Iceland, June 2010.

[17] U. Küttler and A. W. Wall. Fixed-point fluid-structure interaction solvers with dynamic relaxation. *Computational mechanics*, 43(1):61–72, 2008.

[18] E. Albano and W. P. Rodden. A doublet-lattice method for calculating lift distributions on oscillating surfaces in subsonic flows. *AIAA journal*, 7(2):279–285, 1969. doi:10.2514/3.5086.

[19] R. L. Harder and R. N. Desmarais. Interpolation using surface splines. *Journal of aircraft*, 9(2):189–191, 1972. doi:10.2514/3.44330.

[20] A. Schuster, L. Reimer and J. Neumann. A Mesh-Free Parallel Moving Least-Squares-based Interpolation Method for the Application in Aeroelastic Simulations With the Flow Simulator. *New Results in Numerical and Experimental Fluid Mechanics X*, 132:573–583, 2016. doi:10.1007/978-3-319-27279-5_50.

[21] A. Rempke. Netzdeformation mit Elastizitätsanalogie in multidisziplinärer FlowSimulator-Umgebung. In *Jahresbericht 2016 zum 20. DGLR-Fach-Symposium der STAB*. Braunschweig, Germany, November 2016.

[22] A. Eberle, B. Stefes, and D. Reckzeh. Clean Aviation Ultra-Performance Wing (UP-Wing). In *AIAA SCITECH 2024 Forum*, Orlando, Florida, January 2024. doi:doi.org/10.2514/6.2024-2109.

[23] Á. R. Triviño. Horizontal double bubble aircraft design for the middle of the market. Master's thesis, Universidad Europea De Madrid, 2024.

[24] Airbus. A321 aircraft characteristics - airport and maintenance planning. Technical report, Airbus S.A.S., 2005.

[25] C. M. Liersch and M. Hepperle. A distributed toolbox for multidisciplinary preliminary aircraft design. *CEAS Aeronautical Journal*, 2:57–68, 2011.

[26] B. Nagel, D. Böhnke, V. Gollnick, P. Schmollgruber, A. Rizzi, G. La Rocca, J. J. Alonso. Communication in aircraft design: Can we establish a common language. In *28th International Congress of the Aeronautical Sciences*. Brisbane, Australia, September 2012.

[27] R. Maierl, Ö. Petersson, F. Daoud. Automated Creation of Aeroelastic Optimization Models from a Parameterized Geometry. In *15th International Forum on Aeroelasticity and Structural Dynamics*, June 2013.

[28] S. Deinert, Ö. Petersson, F. Daoud and H. Baier. Aircraft loft optimization with respect to aeroelastic lift and induced drag loads. In *10th World Congress on Structural and Multidisciplinary Optimization*, May 2013.

[29] K. Schittkowski. NLPQL: A FORTRAN subroutine solving constrained nonlinear programming problems. *Annals of operations research*, 5:485–500, 1986. doi:10.1007/BF02022087.

[30] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[31] P. M. Weaver. On Optimisation of Long Anisotropic Flat Plates Subject to Shear Buckling Loads. In *45th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics & Materials Conference*. Palm Springs, California, April 2004. doi:10.2514/6.2004-2053.

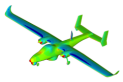
[32] P. M. Weaver. Design Formulae for Buckling of Biaxially Loaded Laminated Rectangular Plates with Flexural/twist Anisotropy. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. Austin, Texas, April 2005. doi:10.2514/6.2005-2105.

[33] P. M. Weaver. Approximate analysis for buckling of compression loaded long rectangular plates with flexural/twist anisotropy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 462(2065):59–73, 2006. doi:10.1098/rspa.2005.1552.

[34] D. E. Calderon, J. E. Cooper, M. Lowenberg, S. A. Neild and E. B. Coetzee. Sizing High-Aspect-Ratio Wings with a Geometrically Nonlinear Beam Model. *Journal of Aircraf*, 56:1455–1470, 2019. doi:10.2514/1.C035296.

[35] D. Raymer. *Aircraft Design: A Conceptual Approach*. American Institute of Aeronautics and Astronautics, 2012. ISBN 0-930403-51-7. doi:10.2514/4.869112.

[36] L. Prandtl. Über Tragflügel kleinsten induzierten Widerstandes. In *Zeitschrift für Flugtechnik und Motorluftschiffahrt*, volume 24, pages 305–306, 1933.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

EXPLOITING POST-BUCKLING NONLINEARITIES FOR GUST LOAD REDUCTION IN HIGH ASPECT RATIO AIRCRAFT WINGS

Pedro Farinha[1*,2], Francesco Toffol[2] and Chiara Bisagni[2]

1: Universidade de Lisboa
Instituto Superior Técnico
Av. Rovisco Pais 1, 1049-001, Lisboa, Portugal
pedro.lourenco.farinha@tecnico.ulisboa.pt, https://tecnico.ulisboa.pt

2: Department of Aerospace Science and Technology
Politecnico di Milano
via la Masa 34, 20156, Milano, Italy
{francesco.toffol,chiara.bisagni}@polimi.it, https://www.polimi.it

Abstract. *The ERC Advanced Grant Project NABUCCO (New Adaptive and BUCkling-driven COmposite aerospace structures) seeks to develop innovative concepts for adaptive and buckling-driven composite structures for next-generation aircraft. By exploiting the buckling phenomena as a design opportunity, the project focuses on utilizing the stiffness redistribution in the post-buckling regime. This study, part of the NABUCCO Project, exploits the structural nonlinearities in the post-buckling regime to mitigate peak loads experienced by high aspect ratio wings, particularly under gust responses. A novel methodology is being developed to incorporate localized nonlinearity introduced by buckling, while maintaining the efficiency of conventional approaches for dynamic load computation. This work focuses on the structural analysis of a wingbox. First, a parametric analysis is conducted to assess the impact of various parameters on the buckling load of a wing section near the root. Based on the results, a new configuration for this section is proposed. Subsequently, a nonlinear analysis is performed on a simplified model to evaluate the stiffness reduction associated with the new configuration. This updated configuration is integrated into the aeroelastic model, and the corresponding results are presented. A structural analysis is carried out to evaluate composite failure, as well as bending and torsion effects on the wing. Preliminary results indicate that the post-buckling response of certain wing components can effectively reduce peak loads caused by gusts, promoting more sustainable aviation practices.*

Keywords: Buckling, post-buckling, gust loads alleviation, structural analysis

Acknowledgements: Funded by the European Union (ERC Advanced Grant, NABUCCO, project number 101053309). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The author Pedro Farinha thanks the Erasmus+ Internship Program for funding this opportunity and Professor Afzal Suleman from Instituto Superior Técnico for supervising this work.

1 INTRODUCTION

With air traffic expected to rise to 25 million flights annually in Europe by 2050, it is pivotal to develop new ways to decrease the impact of aviation on the environment and society. The European Union objectives, defined in the FlightPath2050 document [1], call for a 75% reduction in CO_2 emissions per passenger kilometer, a 90% reduction in nitrogen oxide (NOx) emissions, and a reduction of 65% of the perceived noise.

Advancements in materials, such as the use of high-strength composites, have led to a decrease in weight [2]. Additionally, increasing the wing aspect ratio reduces aerodynamic-induced drag, improving the lift-to-drag ratio and range, but also introduces greater structural flexibility and higher stress levels at the wing root. Reinforcing the wing's internal structure mitigates these issues but increases weight, and reduces fuel efficiency. Moreover, airports and maintenance facilities limit commercial aircraft wingspan [3].

Higher flexibility of the wings poses aeroelastic challenges, where aerodynamic loads cause wing deflection, further altering the loads and possibly leading to catastrophic results [4]. Wind gusts and turbulence increase the Wing Root Bending Moment (WRBM), affecting airframe loads, which can be analyzed using aeroelastic models [5]. Research in Gust Load Alleviation (GLA) techniques has been increasing with passive techniques leveraging structural deformation and active methods using controlled systems to mitigate loads [6].

Various passive load alleviation techniques have been researched. Hahn and Haupt [7] explored the effect of nonlinear behavior of the wing design to trigger downward twist and camber line or profile change. Runkel et al. [8] present a shape adaptation technique for wing structures that enables passive variations in twist. Hahn et al. [9, 10] also suggest a load alleviation concept which incorporates a component inside the wing structure which buckles at a critical load higher than cruise lift leading to a upwards deformation of the trailing edge, resulting in a reduction of the effective angle of attack and a reduction of lift. Ahmadi et al. [11] introduced a new method for the design and optimization of high aspect ratio composite wings that are equipped with passive control devices. while Vaneck et al. [12] invented a stiffness-adaptive wing structure for passive load alleviation which consists of alternating rigid and inflatable spar sections.

Beyond passive load alleviation techniques, numerous active load alleviation approaches have also been developed recently. Toffol and Ricci [13] summarize the work developed within the Clean Sky 2 AIRGREEN2 project, where an innovative wing tip with an additional control surface was developed, achieving between 4 to 17% reduction in WRBM. Stanford [14] performed an optimization of the control surfaces of a high aspect ratio wing for Maneuvre Load Alleviation (MLA) and unsteady GLA.

The ERC Advanced Grant Project NABUCCO (New Adaptive BUCkling-driven COmposite aerospace structures) [15] seeks to develop adaptive buckling-driven structural concepts for the next generation of aircraft. By leveraging buckling as a design opportunity, NABUCCO aims to control structural responses under critical conditions and enable shape variation through non-linear stiffness redistribution induced by the post-buckling regime.

The aim of the research here presented is to reduce WRBM by inducing local buckling at certain locations of the wing, and introduce structural nonlinearity, which can reduce peak loads in cases of gust responses, into the conventional framework for dynamic aeroelastic response calculations. Most current numerical approaches rely on linear aeroelastic

analysis tools, which are inadequate for handling nonlinear stiffness changes induced by buckling. To address this, a novel methodology is being developed that incorporate localized nonlinearity introduced by buckling without significantly increasing computation cost [16, 17].

This paper focuses on the structural analysis of the wingbox of a high aspect ratio wing. The goal is to design a configuration that induces post-buckling in the composite wing skin under critical gust loads. A non-linear analysis is then used to evaluate stiffness redistribution due to buckling, which is subsequently integrated into the aeroelastic framework.

2 NUMERICAL MODEL OF A HIGH ASPECT RATIO WING

The numerical model of the wing analyzed in this study was provided by the Department of Aerospace Science and Technology of the Politecnico di Milano. This model is part of the Ultra High Aspect Ratio Wing Advanced Research and Designs (U-HARWARD) CS2JU founded project [18]. It represents a wingbox similar to that of an Airbus A321, with an aspect ratio (AR) of 15.

This 3D finite element model was generated using the NeoCASS OPtimization Toolbox (NeOPT) [19], an optimization module of the Next generation Conceptual Aero-Structural Sizing Suite (NeoCASS)[20, 21], developed at the Department of Aerospace Engineering at Politecnico di Milano. Using NeOPT, a detailed NASTRAN Finite Element [FE] model of the wingbox was generated from an aeroelastic stick model and subsequently converted to ABAQUS, the solver used for the structural analysis presented in this study.

2.1 Mesh

The FE model consists of composite shell elements (S4R) to represent the skins, while beam elements (B31) are used to model the stringers and spar caps. The ribs are also modeled using shell elements (S4R). The model is divided into five sections along the span, as indicated by the change in skin color shown in Fig. 1. Each section consists of a set of stringers, ribs, and composite skins. Figure 2 provides a close-up view of the first and second sections of the wing, with the upper skin panel removed.

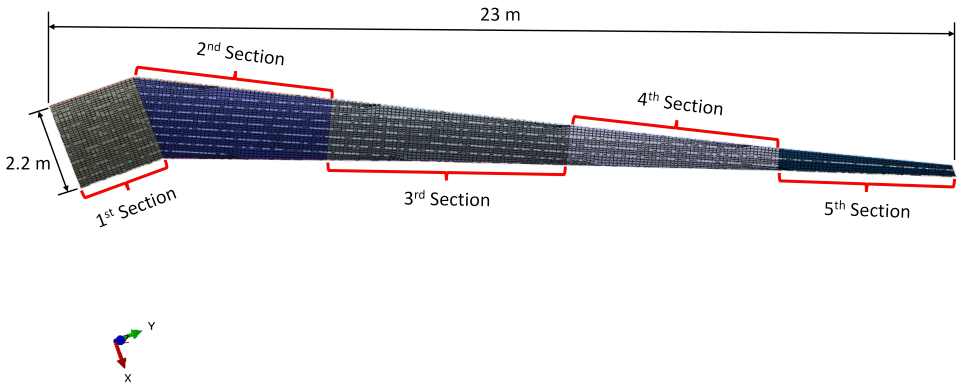


Figure 1: Top view of the whole wingbox.

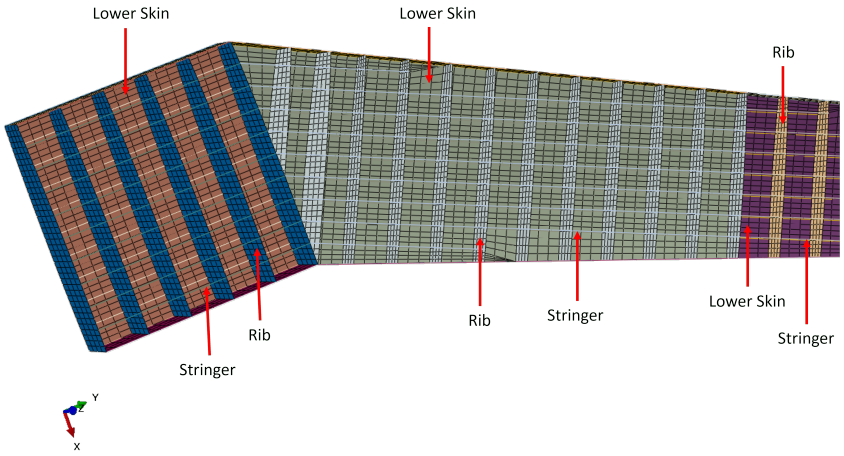


Figure 2: Closeup view of the first and second sections of the wingbox.

2.2 Loads and boundary conditions

The applied loads are derived using NeoCASS's capability to generate a static load set equivalent to a frozen time instant of the dynamic gust. Each load is extracted from the NeoCASS aeroelastic stick model and applied to reference nodes in ABAQUS, which are connected to nodes centered on the ribs' surfaces through kinematic coupling constraints, with all degrees of freedom constrained. The load is then distributed to every node of the rib from this surface node using a coupling constraint. This constraint, shown in yellow in Fig. 3a, enables a weighted distribution of the loads applied at the reference node to the rib nodes, based on a weight factor associated with each coupling node, which depends on its distance from the reference node. Boundary conditions are applied to all nodes of the root rib, constraining all translational degrees of freedom (U1 = U2 = U3 = 0). Figure 3b illustrates the boundary conditions in orange, with forces and moments represented by yellow and purple arrows, respectively. Figure 3 displays only the first and second sections of the wing, however, this method is applied across the entire wingspan.

The loads applied are representative of the most critical case of the wing's response to a gust load, which corresponds to the one that produces the highest wing root bending moment.



(a) Coupling constraints applied on each rib of the model, in yellow.

(b) Loads (arrows) and boundary conditions, in orange, applied on the model.

Figure 3: Closeup views on the coupling constraints, loads, and boundary conditions.

2.3 Materials

The wing skin consists of a carbon fiber composite laminate. Table 1 reports the engineering properties of the laminate. These properties closely resemble those of the

IM7/8552 carbon fiber [22–25]. Table 2 presents the ply strengths of the wing's composite skin. The ribs and stringers are made of an aluminum alloy, with properties listed in Tab. 3 and Tab. 4, respectively.

Table 1: Laminate properties of the wing's skin.

E_1 [MPa]	E_2 [MPa]	ν_{12}	G_{12} [MPa]	G_{13} [MPa]	G_{23} [MPa]	ρ $[kg/m^3]$
161000	11400	0.32	5200	3740	2710	1570

Table 2: Ply strengths of the wing's skin.

X_T [MPa]	X_C [MPa]	Y_T [MPa]	Y_C [MPa]	SL [MPa]
2724	1690	111	264	120

Table 3: Material properties of the ribs.

E [MPa]	ν_{12}	ρ $[kg/m^3]$
72000	0.33	1540

Table 4: Material properties of the stringers.

E [MPa]	G [MPa]	ν_{12}	ρ $[kg/m^3]$
67000	27006.5	0	1570

2.4 Section properties

Table 5 presents the thickness of each wing skin panel, the thickness of the ribs, and the stringer areas for each section.

Table 5: Section properties.

	Section				
	First	Second	Third	Fourth	Fifth
Upper Panel Thickness [mm]	11.67	11.47	14.00	13.47	9.70
Lower Panel Thickness [mm]	11.13	12.20	13.60	12.80	12.80
Ribs Thickness [mm]	12.00	6.00	4.00	3.00	3.00
Upper Stringers Area $[mm^2]$	1840	1940	1630	1120	937
Lower Stringers Area $[mm^2]$	1710	1640	1620	1360	901

All composite wing skin laminates are quasi-isotropic and share the same ply stacking sequence, $[0, 45, -45, 90]_s$. Table 6 presents the thickness distribution for each ply orientation:

Table 6: Thickness by ply orientation.

	Orientation			
	0°	45°	−45°	90°
Thickness (%)	30	30	30	10

2.5 Linear static and buckling analyses

Linear static and buckling eigenvalue analyses were performed on the wing model. Applying the gust load, the linear static analysis yields a maximum vertical displacement at the wing tip of $\Delta_z = 2.040\ m$, and the buckling analysis results in a first eigenvalue, corresponding to the first buckling mode, of $\lambda = 3.0438$. Thus, in the original model, buckling occurs only under load conditions three times higher than those applied. Moreover, buckling is localized in a very small region of the wing, specifically at a stringer interruption.

3 SENSITIVITY STUDY OF THE WING

The results of the linear buckling analysis indicate that the original wing model does not present any buckling with the gust load. A sensitivity study using linear static and buckling analyses provides valuable insight into the structural response to design changes, allowing buckling in some parts of the structure.

3.1 Upper panel skin thickness

The first parameter analyzed is the thickness of the upper panel in the wing's second section, while preserving the thickness distribution shown in Tab. 6.

Results showed that reducing the upper panel skin thickness significantly decreases the first eigenvalue. However, an eigenvalue below 1 is only reached at a 50% reduction, which may lead to issues such as ply failure, ultimately compromising the structure's ability to withstand the applied loads. Furthermore, ABAQUS results indicate that the buckling shape remains confined to a very small area, leading to minimal stiffness reduction in the post-buckling regime.

3.2 Stringers area

The influence of the stringers area in the upper and lower skins of the second wing section on the results was also examined. Reducing the upper panel stringers area resulted in small reduction of the eigenvalue, decreasing by only 14.40% with a 50% decrease in stringers area. Furthermore, the buckling region remained small. However, the maximum vertical displacement at the tip was comparable to that observed when varying the upper skin thickness.

The effect of the lower panel stringer area was negligible, therefore this parameter was excluded from further studies.

3.3 Skin thickness and upper stringers area

Given that varying the upper skin thickness and upper stringer area yielded favorable results, the next step in this sensitivity analysis is to assess the combined effect of modifying both parameters simultaneously.

A greater reduction in the first buckling eigenvalue and a higher maximum vertical displacement was achieved compared to only varying the skin thickness, reaching the post-buckling regime with a 40% decrease in skin thickness and upper stringer area. A higher maximum vertical displacement highlights the role of the upper stringers in preventing bending.

3.4 Number of stringers

In an effort to expand the buckling region, the next step in the sensitivity analysis was to assess the effect of the number of stringers in the wing's second section on the results. Since buckling occurred at the trailing edge, the stringers were progressively removed one by one, starting from the trailing edge and moving toward the leading edge.

Removing the first two stringers near the trailing edge reduces the first buckling eigenvalue by nearly 68%, bringing it below 1. In contrast, removing all remaining stringers results in only a 16% further reduction. The removal of stringers expands the buckling area but confines it between ribs, forming elongated chord-wise buckling zones. Aerodynamically, this is undesirable since buckling near the leading edge disrupts laminar flow. Ideally, buckling should occur near the trailing edge, where upward deflection reduces the effective angle of attack, thereby decreasing lift and alleviating gust loads [10].

3.5 Number of ribs

The influence of rib count on structural behavior was also examined. Initially, the 12^{th} rib was removed, as it is adjacent to the primary buckling zone. Subsequently, the 13^{th} rib was removed, creating a buckling-prone region between adjacent ribs. In the following analyses, two ribs were removed at a time, one from each side of the region. Since loads are applied at the ribs, their removal redistributes the loads to adjacent ribs through distributing coupling constraints. To preserve the load distribution, the outermost ribs of the second section of the wing were retained. ABAQUS results showed that as more ribs were removed, the buckling area expanded, and the eigenvalue dropped significantly, while the tip displacement remained largely unchanged.

A larger buckling area was achieved, demonstrating that rib removal is crucial for this purpose. However, achieving a first buckling eigenvalue close to 0.7 required removing at least 8 ribs, which is too extreme and leads to other structural issues.

3.6 Configuration study

Based on the previous analyses, it was concluded that rib removal is crucial for creating a significant skin area in buckling. Therefore, the 11^{th} and 12^{th} ribs of the model were removed, creating a new section without ribs, hereafter referred to as the *buckling section* as shown in Fig. 4.

A parametric study was conducted on this model, in which the 11^{th} and 12^{th} ribs were removed, all stringers were retained, and the upper panel skin thickness in the *buckling section* was the parametric variable. When the skin thickness was reduced by 50%, the first eigenvalue fell below 1; however, the buckling area remained too small. Therefore, it was necessary to remove stringers and ribs from the model.

After several iterations, two final configurations were developed. In both configurations, the 11^{th} and 12^{th} ribs were removed. In *Configuration 1*, within the *buckling section*, the

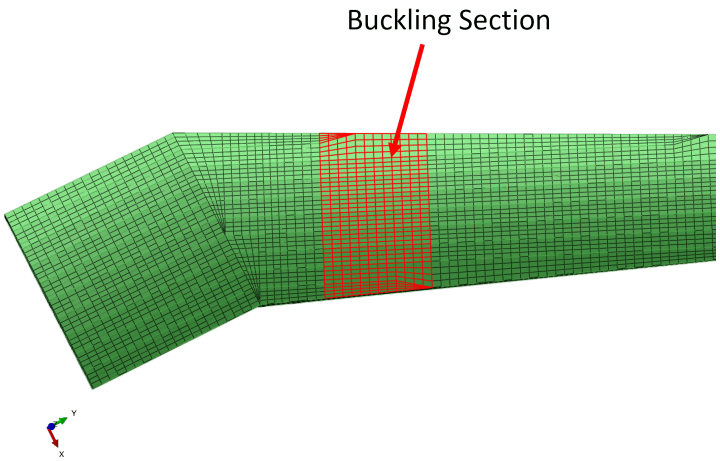


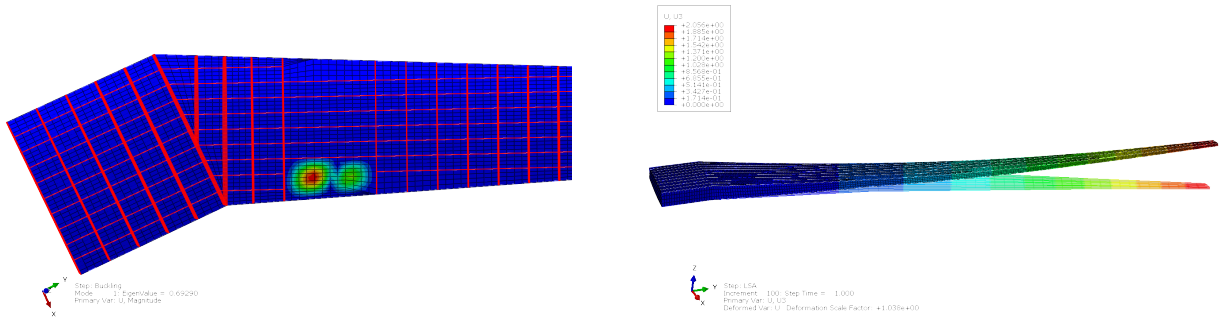
Figure 4: Buckling section region.

two stringers closer to the trailing edge were removed, and the skin thickness was reduced by 10%. In *Configuration 2*, within the *buckling section*, the three stringers closer to the trailing edge were removed, and the skin thickness was increased by 10%. Table 7 presents the results of the linear buckling and static analyses.

Table 7: Linear buckling and linear static analyses results for the two configurations.

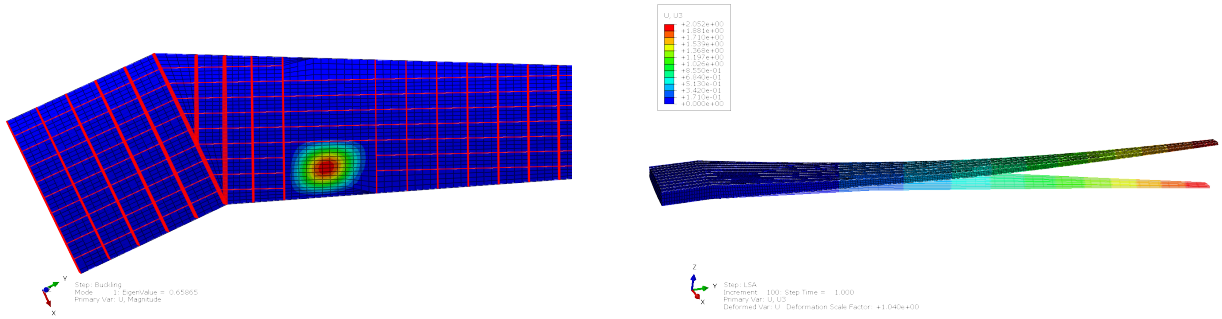
Configuration	Eigenvalue	Maximum Tip Vertical Displacement [m]
Configuration 1	0.693	2.056
Configuration 2	0.659	2.052

Both configurations resulted in similar first buckling eigenvalues and tip displacements, differing primarily in the first buckling mode, which influences post-buckling stiffness. Figures 5 and 6 illustrate the results of the linear buckling and static analyses for the two configurations. *Configuration 1* shows a half-wave buckling mode with a larger area than *Configuration 2*.



(a) First buckling shape with $\lambda = 0.693$. (b) Deformed shape.

Figure 5: Linear buckling and linear static analyses of Configuration 1.



(a) First buckling shape with $\lambda = 0.659$. (b) Deformed shape.

Figure 6: Linear buckling and linear static analyses of Configuration 2.

4 NON-LINEAR ANALYSIS

Both configurations obtained from the parametric study were analyzed in a non-linear analysis. Non-linear analysis allows observation of the skin's displacement behavior over time, enabling the extraction of data from both the pre- and post-buckling regimes.

Although the first buckling eigenvalue of *Configuration 1* was lower than 1, the non-linear analysis revealed that the stiffness variation in post-buckling fiel was minimal, resulting in almost no change in the wing's vertical displacement compared to the linear static solution. The largest change in maximum vertical displacement was only 0.9%. *Configuration 2* exhibited the same behavior, leading to the conclusion that the buckling areas of both configurations were too small to significantly affect the wing's stiffness. Consequently, both configurations were discarded from further studies, and new designs with larger buckling areas were developed.

A new design was developed to achieve a larger buckling region while maintaining a first eigenvalue similar to that of the previous configurations. In this new configuration, the 10^{th}, 11^{th}, 12^{th}, and 13^{th} ribs were removed. Additionally, the four stringers closer to the trailing edge were removed. The thickness of the *buckling section*, now comprising the upper panel skin from the 9^{th} to the 14^{th} rib, was increased by 35%.

However, due to the larger buckling area, it is crucial to ensure that there is no skin failure. Therefore, the first ply-failure criterion, using Tsai-Hill, was evaluated through the non-linear analysis, resulting in a change of the laminate ply orientation to $[-45, 45, 90, 0]_s$, while maintaining the thickness distribution shown in Tab. 6. Figure 7 shows the contour of the maximum Tsai-Hill criterion values across all plies. The peak value obtained is equal to 0.834, ensuring there is no ply failure.

Figure 8 illustrates the buckling shape of this new configuration, with a first mode eigenvalue of $\lambda = 0.63294$.

By comparing the maximum vertical displacements at different spanwise positions, as shown in Tab. 8, both linear and non-linear analyses provide insight into the extent of stiffness reduction due to buckling.

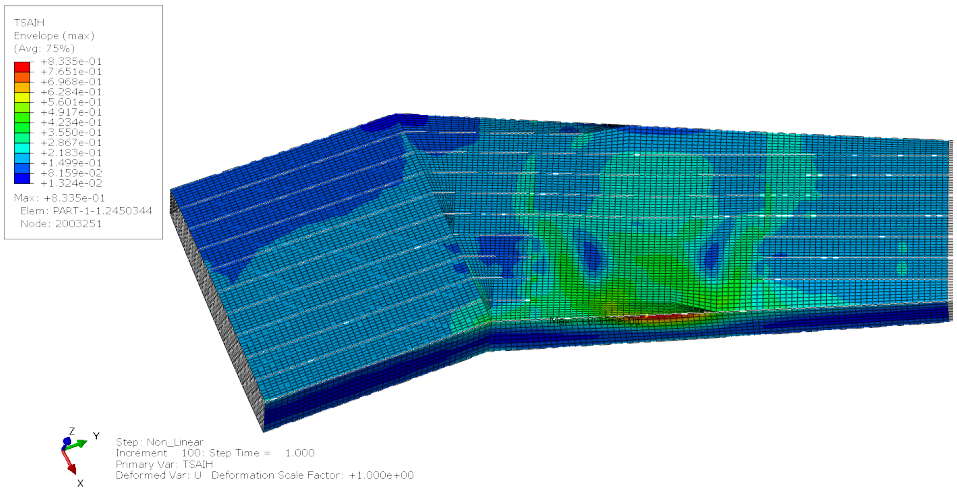


Figure 7: Contour of maximum Tsai-Hill.

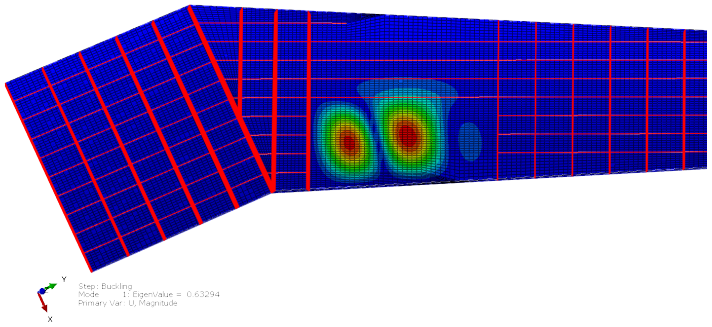


Figure 8: First buckling shape with $\lambda = 0.63294$.

Table 8: Maximum vertical displacement in linear and non-linear analysis.

Node number	Spanwise position [m]	Maximum vertical displacement [m]		Variation [%]
		Linear analysis	Non-linear analysis	
1	1.04	0.00332	0.003340	0.567
2	2.07	0.0109	0.0110	1.58
3	3.91	0.0450	0.0475	5.62
4	5.75	0.101	0.111	9.97
5	6.91	0.155	0.172	10.42
6	8.07	0.224	0.246	9.93
7	9.79	0.354	0.385	8.67
8	11.52	0.519	0.558	7.50
9	13.25	0.718	0.764	6.48
10	14.98	0.952	1.005	5.57
11	16.70	1.217	1.276	4.82
12	18.25	1.477	1.540	4.27
13	19.79	1.755	1.822	3.81
14	21.34	2.038	2.109	3.45

Analysis of the results reveals that nodes near the buckling region experience a more significant stiffness reduction due to the non-linearities of the post-buckling regime, with the maximum vertical displacement increasing by up to 10%.

Figure 9 illustrates the relationship between the total applied force and the vertical displacement of a node at a spanwise position of 6.91 m from the wing root for both linear static and non-linear analyses. The non-linear analysis results begin to diverge from the linear static solution at 55% of the applied gust load, indicating a stiffness reduction and an increase in vertical displacement. This suggests that the skin enters the post-buckling regime at 55% of the gust load whereas the linear buckling analysis predicted a buckling load at 63%. Furthermore, the figure includes linear interpolation equations for both the linear static response and the post-buckling regime in the non-linear analysis. The slopes of these interpolations represent the wing's stiffness at this node allowing the calculation of the stiffness ratio between the pre- and post-buckling regimes. For the node located at a spanwise position of 6.91 m from the wing root, the post-buckling regime results in a 17.4% reduction in stiffness.

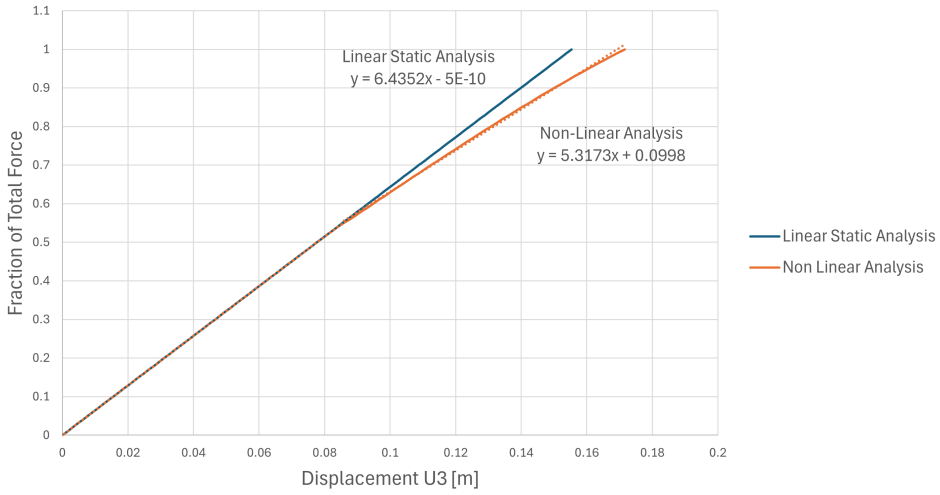


Figure 9: Vertical displacement of a node at a spanwise position of 6.91 m from the wing root in function of total force applied.

5 IMPLEMENTATION INTO AEROELASTIC FRAMEWORK

The post-buckling results obtained from the non-linear analysis can be integrated into the aeroelastic framework to assess the gust load alleviation potential of the proposed configuration, quantified by the reduction in WRBM.

When comparing the non-linear analysis results obtained for the proposed configuration with the linear analysis of the original high aspect ratio wing, an 82% reduction of the buckling load was obtained along with an 11.7% reduction of the post-buckling stiffness. According to the study by Toffol and Bisagni [26], which focuses on the same aircraft, such levels of buckling load and stiffness reduction can lead to an approximate 0.5% decrease in WRBM.

6 CONCLUSIONS

This work presents the structural analysis of a high aspect ratio carbon fiber wing, focusing on leveraging post-buckling behavior in wingbox components to mitigate dynamic peak loads. To achieve this, design modifications were made to induce the wing skin to enter the post-buckling regime.

Buckling analysis carried out in ABAQUS revealed that the buckling load was three times higher than the gust load. A sensitivity study was conducted, to find a configuration with a buckling load lower than the gust load, and a buckling area high enough to produce significant stiffness reduction in the post-buckling regime.

The final configuration achieved a buckling load of approximately 55% of the critical gust load, with the proposed configuration registering an 82% decrease in the buckling load and an 11.7% reduction of the overall stiffness, compared to the original high aspect ratio wing. It was obtained by removing 4 ribs, 4 stringers, and by increasing the skin thickness by 35% and changing the laminate ply sequence to ensure there was no ply failure.

Preliminary results indicate that a WRBM reduction of 0.5% can be obtained with the proposed configuration. For future work, the latest design will be integrated into the aeroelastic stick model framework to assess whether the achieved stiffness reduction effectively mitigates gust loads dynamically. These advancements aim to develop feasible, robust buckling-driven designs for aerospace structures, contributing to structural weight reduction, lower dynamic loading, and improved fuel efficiency.

REFERENCES

[1] European Commission and Directorate-General for Mobility and Transport and Directorate-General for Research and Innovation. *Flightpath 2050 – Europe's vision for aviation – Maintaining global leadership and serving society's needs*. Publications Office, 2011. doi:10.2777/50266.

[2] C. Soutis. Fibre reinforced composites in aircraft construction. *Progress in Aerospace Sciences*, 41(2):143–151, 2005.

[3] F. Afonso, J. Vale, É. Oliveira, F. Lau, and A. Suleman. A review on non-linear aeroelasticity of high aspect-ratio wings. *Progress in Aerospace Sciences*, 89:40–57, 2017.

[4] J. Bussemaker. Wing optimization with active load control. *Master's Thesis, Faculty of Aerospace Engineering, TU Delft, Delft, The Netherlands*, 2018.

[5] D. Balatti, H. Haddad Khodaparast, M. I. Friswell, M. Manolesos, and M. Amoozgar. The effect of folding wingtips on the worst-case gust loads of a simplified aircraft model. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 236(2):219–237, 2022.

[6] C. L. Bottasso, F. Campagnolo, A. Croce, and C. Tibaldi. Optimization-based study of bend–twist coupled rotor blades for passive and integrated passive/active load alleviation. *Wind Energy*, 16(8):1149–1166, 2013.

[7] D. Hahn and M. Haupt. Exploration of the effect of wing component post-buckling on bending-twist coupling for nonlinear wing twist. *CEAS Aeronautical Journal*, 13 (3):663–676, 2022.

[8] F. Runkel, U. Fasel, G. Molinari, A. F. Arrieta, and P. Ermanni. Wing twisting by elastic instability: A purely passive approach. *Composite Structures*, 206:750–761, 2018.

[9] D. Hahn, M. Haupt, and S. Heimbs. Load alleviation capabilities of wings with nonlinear structural behavior in stationary and dynamic load cases. In *International Forum on Aeroelasticity and Structural Dynamics, IFASD, Madrid, Spain*, pages 642–646, 2022.

[10] D. Hahn, M. Haupt, and S. Heimbs. Passive load alleviation by nonlinear stiffness of airfoil structures. In *AIAA SciTech 2022 Forum*, paper AIAA 2022-0318, 2022.

[11] M. Ahmadi, T. Farsadi, and H. H. Khodaparast. Enhancing gust load alleviation performance in an optimized composite wing using passive wingtip devices: Folding and twist approaches. *Aerospace Science and Technology*, 147:109023, 2024.

[12] T. Vaneck and R. Parks. Wing load alleviation structure. *Patent US8720822*, 2014.

[13] F. Toffol and S. Ricci. Development of an active wingtip for aeroelastic control. *Aerospace*, 10(8):693, 2023.

[14] B. Stanford. Optimal aircraft control surface layouts for maneuver and gust load alleviation. In *AIAA SciTech 2020 Forum*, paper AIAA 2020-0448, 2020.

[15] ERC NABUCCO - New Adaptive and BUCkling-driven COmposite aerospace structures. `https://www.aero.polimi.it/it/progetti-di-ricerca/erc-nabucco-new-adaptive-and-buckling-driven-composite-aerospace-structures`. Accessed: 25-03-2025.

[16] F. Toffol and C. Bisagni. Gust load passive alleviation by means on nonlinear, buckling driven, structural response. In *ICAS Proceedings*. International Council of the Aeronautical Sciences, paper ICAS2024_0147, 2024.

[17] C. Bisagni. Nabucco take-off: Multi-stable panels for an adaptive wing. In *ICAS Proceedings*. International Council of the Aeronautical Sciences, paper ICAS2024_1230, 2024.

[18] U-HARWARD PROJECT. Ultra High Aspect Ratio Wing Advanced Research and Designs. `https://www.u-harward-project.eu/project/`. Accessed: 25-03-2025.

[19] F. Toffol and S. Ricci. Neopt: an optimization suite for the aeroelastic preliminary design. In *18th International Forum on Aeroelasticity and Structural Dynamics (IFASD 2019)*, paper IFASD-2019-020, 2019.

[20] L. Cavagna, S. Ricci, and L. Travaglini. Neocass: an integrated tool for structural sizing, aeroelastic analysis and mdo at conceptual design level. *Progress in Aerospace Sciences*, 47(8):621–635, 2011.

[21] L. Cavagna, S. Ricci, and L. Travaglini. Structural sizing and aeroelastic optimization in aircraft conceptual design using neocass suite. In *13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, paper AIAA 2010-9076, 2010.

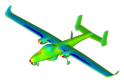
[22] HEXCEL. HexPly 8552. `https://www.hexcel.com/user_area/content_media/raw/HexPly_8552_eu_DataSheet(1).pdf`. Accessed: 12-03-2025.

[23] Carbon Fiber Fabric 4 Harness Satin Intermediate Modulus 6k Hexcel IM7 Model #: F-1226-49. `https://compositeenvisions.com/wp-content/uploads/2018/03/F-1226-49-TDS-Hexcel-IM7.pdf`. Accessed: 12-03-2025.

[24] A. E. Lovejoy, S. Scotti, S. Miller, P. Heimann, and S. Miller. Characterization of IM7/8552 thin-ply and hybrid thin-ply composites. In *AIAA SciTech 2019 Forum*, paper AIAA 2019-0773, 2019.

[25] B. P. Justusson, M. T. Molitor, J. S. Iqbal, M. Rassaian, T. M. Ricks, and R. K. Goldberg. Overview of coupon testing of an IM7/8552 composite required to characterize high-energy impact dynamic material models. Technical report, NASA - 2020-220498, 2020.

[26] F. Toffol and C. Bisagni. Passive method for dynamic loads alleviation using structural buckling in a composite wing. In *AIAA SciTech 2025 Forum*, paper AIAA 2025-0425, 2025.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

SOLVING NONLINEAR PERIODIC AEROELASTIC SYSTEMS USING BOUNDARY VALUE FORMULATION AND CONFORMAL MAPPING

Anshuman Mehta, Pankil N. Mishra and Abhijit Gogulapati

Department of Aerospace Engineeing, IIT Bombay, Powai, Mumbai, Maharashtra, India - 400076

Abstract. *Traditional approaches for simulating nonlinear aeroelastic systems typically rely on initial value problems (IVPs), which do not fully exploit the inherent periodicity of these systems. Our research addresses this limitation by reformulating the problem as a periodic boundary value problem (BVP), leveraging numerical optimization techniques to enhance computational efficiency and system insights. This study focuses on the optimization of nonlinear periodic aeroelastic systems using a BVP formulation.*

A trust-region-based optimization framework was implemented to solve the periodic BVP, incorporating nonlinear constraints and periodic boundary conditions. The aerodynamic analysis employed Joukowsky conformal mapping-based inviscid aerodynamic models to simulate unsteady flows around airfoils undergoing periodic kinematics, including pitch, plunge, and lead motions. The formulation also accounts for leading-edge vortices, ensuring accurate representation of the complex wake dynamics and key physical conditions.

The proposed framework was validated through numerical simulations of oscillating airfoils, demonstrating its robustness across a range of reduced frequencies and dynamic configurations. A systematic comparison of various formulations, based on choices of objective functions and nonlinear constraints, highlighted significant reductions in computation time while maintaining solution accuracy. Additionally, the study identified the required wake length for accurately capturing the periodic system behavior, which significantly reduces computational power compared to traditional approaches.

This research presents a novel methodology for the analysis and design of nonlinear aeroelastic systems, offering enhanced computational efficiency and deeper insights into system behavior. The integration of mid-fidelity aerodynamic models and advanced optimization algorithms establishes a robust framework for multidisciplinary aerospace design, particularly for applications involving periodic phenomena.

Keywords: Boundary value problem (BVP), trust-region optimization, Joukowsky conformal mapping, inviscid aerodynamic models, periodic kinematics, leading-edge vortices

1 INTRODUCTION

Periodicity is crucial in analyzing dynamical systems, particularly nonlinear ones. It reveals repeating patterns in behavior, with the stability and bifurcation of periodic solutions being vital for grasping the system's long-term dynamics.[1].

1.1 Nonlinear Periodic Dynamics

Nonlinear periodic aeroelastic systems display complex dynamics and regular patterns due to their interaction of aerodynamic forces, structural dynamics, and inertia. Typically, analyzing these systems requires solving an initial value problem (IVP) with time-forward integration to find a stable periodic solution. However, this process can be inefficient, often necessitating many cycles for transient decay and careful timestep adjustment to locate the correct periodic solution. Moreover, since IVP simulations naturally align with stable attractors, capturing unstable or multiple periodic solutions poses difficulties. Thus, exploring a method that inherently integrates the system's periodicity into its analysis is recommended.

Early studies of periodic solutions in nonlinear vibrating systems date back to the early 20th century, with contributions from pioneers such as Poincaré [2] and Hopf [3], focusing initially on single-degree-of-freedom systems and employing approximation theory based on smoothness without discontinuities. Insight was generally confined to basic cases. Poincaré recognized periodic solutions as fundamental to pure mathematics. Analyses with early numerical methods, including analog computers evaluating Duffing and Van Der Pol oscillators, achieved limited progress. The 1980s saw numerical algorithms transform the computation of periodic solutions. Researchers like Seydel [4] and Holodniok and Kubíček [5] utilized the multiple shooting method, originally developed by Bulirsch, Stoer, and Deuflhard [6], addressing boundary value problems with unknown periods. However, the challenge of defining boundary value problems due to the non-uniqueness of periodic solutions persisted. Recasting the search for periodic solutions as a boundary-value problem (BVP) requires an effective optimization strategy to manage the resultant nonlinear equations efficiently. Unlike IVP approaches, which rely on protracted time integration, the BVP method applies periodicity constraints directly, transforming the issue into a global root-finding or constrained optimization task.

Another critical aspect of accurately predicting periodic aerodynamic loads is the fidelity of the inviscid unsteady aerodynamic model employed.

1.2 Inviscid Unsteady Aerodyanmic models

The equations of fluid flow, particularly the 3D Navier–Stokes equations, are notoriously unsolvable analytically. Numerical solutions are possible, but turbulence makes direct simulations impractical for high Reynolds numbers, as they exceed even modern supercomputers' capabilities. To model practical problems, simplified equations—such as averaged or filtered Navier–Stokes relations—are used [7]. The fastest solutions come from potential flow equations, which assume inviscid, irrotational, and isentropic (or incompressible) flow. While these assumptions discard significant physics, potential flow methods remain valid for many aeronautical applications under nominal conditions, such as aircraft aeroelastic calculations. They are computationally efficient, requiring calculations only on the body's surface and wake, and can even approximate separated flows

by shedding vortices into the wake. Potential flow methods for steady aerodynamics are well-documented in textbooks like Katz and Plotkin [8], while Gülçat [9] covers unsteady aerodynamics across flow regimes.

To model 2D inviscid unsteady flow with mid-fidelity, numerical methods can be employed to more accurately represent the wake behind oscillating airfoils or airfoils with non-negligible thickness with camber. Two such established methods are presented here, each with its advantages and disadvantages. These methods are used to demonstrate the physics of wake formation behind oscillating airfoils and the mechanisms of thrust generation. Comparisons of predicted aerodynamic loads with experimental results highlight how mid-fidelity numerical methods can capture the physics of real phenomena.

A discrete vortex method for numerical modeling of unsteady aerodynamics is a 2D Joukowsky conformal mapping based given by [10, 11] and is based on [12]. It employ a conformal mapping between the airfoil and a circle with a known potential flow solution. An image vortex, with strength equal and opposite to the fluid vortex, is strategically placed to enforce the non-penetration and Kutta conditions.

To expand the range of wing kinematics, release vortex elements from both edges of the wing, and allow the elements to move according to classical vortex dynamics. The increased number of elements enables these models to capture a much richer set of transient phenomena, such as the formation of leading-edge vortices and the emergence of coherent vortex structures. Unfortunately, these models no longer have analytical solutions, and flow response must be numerically computed. The interaction between vortex elements requires computing an n-body problem (with n being the number of elements) at every simulation time step. Even using fast summation methods that can solve n-body problems in $O(n)$, the computational complexity of these models is at best quadratic in time, as new elements must be regularly introduced into the fluid to model vorticity production. This quadratic growth of time complexity makes these vortex models impracticable for real-time applications [13].

Another aspect is that the vorticity in the wake at any given time is generated on and shed from the wing in the recent past, affecting the flow near the wing. However, as the wake vorticity is transported downstream, its influence diminishes, resulting in a fading memory. Considering a finite part of the wake is motivated by the fact that fading memory is related to the notion of a unique steady state [14].

1.3 Optimization framework in boundary-value formulation

The optimization framework employed in this study builds upon trust-region methods [15], which are well-suited for solving large-scale, nonlinear systems with constraints. The trust-region approach ensures numerical stability and efficient convergence by controlling the step size in parameter updates, allowing for adaptive refinement of the solution. Specifically, the solver minimizes the residual norm of the periodicity constraints while simultaneously satisfying aerodynamic governing equations.

1.4 Periodic Aerodynamic system in boundary- Value formulation

This shooting method based approach has been extensively employed in the helicopter community to compute the wake trailing behind the blades [16]. In this method, the strengths and positions of the trailed vortex elements attached to the blades are specified as an initial condition, and the periodicity of the positions and strengths is enforced as

a boundary condition. These methods, known as wake relaxation methods [17] in rotor aerodynamics, can be further classified as influence coefficient-based [18, 19], semi-implicit [20], and pseudo-implicit methods [17, 21]. Details on implementing boundary conditions can be found in Ref [17].

1.5 Objectives

This study reformulates the pursuit of periodic solutions as a boundary value problem (BVP). By directly applying periodicity conditions as time constraints, rather than waiting for transients to subside, optimization techniques are used to resolve the limit-cycle oscillation (LCO) or periodic orbit. The global benefits of this method include:

- achieving periodic steady-state directly without transient simulation,
- accessing unstable or subtle periodic solutions through smart initial guesses or continuation,
- enhanced numerical stability and insight by satisfying global constraints, avoiding long-term integration.

In this present work, we integrate the recently developed boundary value problem formulation by authors [22, 23] and particular numerical vortex methods to allow for

1. arbitrary body shapes (not necessarily flat plate),
2. free wake or wake deformation, and
3. leading edge wake along with trailing edge wake,
4. exploration of the convergence of the model alongwith computational cost
5. development of an optimization framework to compute periodic solutions.

hence making it useful for many applications in aeronautical engineering.

2 METHODOLOGY

The traditional approaches are formulated as an initial-value problem, in that initial conditions are known and the flow is solved for all subsequent times using a time-marching algorithm.

In the unsteady motion of an aerofoil, the kinematic boundary condition induces velocities over the aerofoil surface. These upset the Kutta-Joukowski condition that the trailing edge must be a stagnation point because nature abhors the infinite velocity that must exist there if the flow did not stagnate. A similar condition holds for the leading edge where separation is also observed in insect-like flapping wings. The Kutta-Joukowski condition is, therefore, enforced by the 'creation' of bound circulation. A pair of vortices are shed at every time-step, one each from the leading and trailing edges in accordance with Kelvin's law and to satisfy the Kutta-Joukowski conditions there. Over the duration of the time-step, these vortices convect with the local Kirchhoff velocity, so that at the next time-step the Kutta-Joukowski condition is disturbed once again. In addition, the kinematics of the wing and the associated boundary condition may have changed so that

the velocity distribution on the wing is further altered. The Kutta-Joukowski condition is restored by shedding another pair of vortices from the respective edges. This proces repeats over time and wakes from both separation edges evolve. A spin-off of the method is that a time-history of the evolution of the wake is automatically generated which is useful in demonstrating the physics of the flow.

Governing equations for inviscid vortex models with rigid body kinematics are described by the potential-flow Laplace equation, $\nabla^2\phi = 0$, where ϕ represents the velocity potential. However, the existence and uniqueness of solutions for boundary value problems (BVPs) are more complex than for initial value problems (IVPs). To ensure uniqueness, five boundary conditions (BCs) are implemented:

1. **Far-field BC**: This condition limits the domain of the independent variables (x and y) in the numerical approximation of the differential equations. Mathematically, as x and y tend to infinity, ϕ approaches zero. By choosing a suitable vortex model, the far-field BC is automatically satisfied, and the domain size can be limited without significantly affecting the solution.

    - **Fading Memory in x direction**: A fading memory relationship exists between vorticity and induced velocity [24]. The vorticity in the wake at any given time is generated on and shed from the wing in the recent past, affecting the flow near the wing. However, as the wake vorticity is transported downstream, its influence diminishes, resulting in a fading memory. Considering a finite part of the wake is motivated by the fact that fading memory is related to the notion of a unique steady state [14].

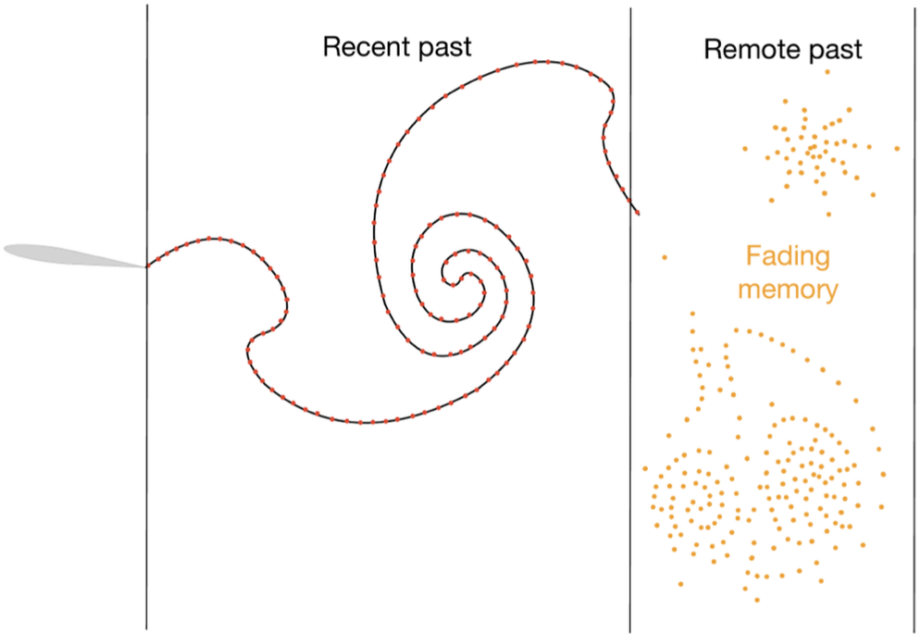


Figure 1: Fading memory

    - **Bendixson-Poincare Theorem in y direction**: The solution's boundedness leads to the Bendixson-Poincare theorem [25], which states that any bounded solution of a differential equation converges to:

  - a fixed point
  - a periodic orbit
  - a connected set composed of a finite number of fixed points together with homoclinic and heteroclinic orbits containing these.

2. **Non-penetration condition**: In the flow past a solid body moving, the non-penetration condition ensures that the fluid velocity normal to the surface $(\nabla\phi\cdot\hat{n})\cdot\hat{n}$ is equal to the solid's velocity. This condition is a Neumann-type boundary condition applied at the solid body surface to prevent fluid penetration.

3. **Kelvin-Circulation Theorem**: The Kelvin-Circulation theorem states that the circulation around a closed curve consisting of fluid particles remains constant with time. It is an invariant of the inviscid aerodynamic model and is applied as a boundary condition.

4. **Kutta Condition**: The Kutta condition states that the trailing edge of an airfoil should not be a source of circulation but rather a conduit for passing circulation from the body into the fluid. It ensures the proper representation of the flow behavior near the trailing edge.

5. **Wake Transport Equation**: The wake transport equation accounts for the interaction between wake vortices and the body, capturing the changes in vortex positions and strengths.

These boundary conditions ensure the uniqueness of the solution within the free wake methods. By considering fading memory, implementing the non-penetration condition, incorporating the Kelvin-Circulation theorem, enforcing the Kutta condition, and including the wake transport equation, the wake modeling framework achieves reliable and accurate results.

2.1 Lighthill Vorticity Creation Mechanism

Different formulations of inviscid potential-flow methods [8, 26, 27]require the time evolution of a state vector that fully describes the instantaneous state of the fluid-body system. This state vector, denoted as $\mathbf{u}(t)$, typically consists of the positions and strengths of vortices, along with the configuration of the body. The time evolution of the state vector follows a two-step sequence known as the Lighthill vorticity creation mechanism [28].

In the case of unsteady aerodynamics with periodic kinematics, this state vector comprises vortex elements' positions and strengths along with the configuration of the body. The time evolution of the state vector is given by a set of ordinary differential equations, which generally can be written as:

$$\frac{d\mathbf{u}}{dt} = \mathbf{F}(\mathbf{u}, t), \qquad \mathbf{u}(0) = \mathbf{u}_0$$

Flow is developed as a sequence of infinitesimal steps consisting of a new wake vorticity creation, wake vortex sheet development, and wake convection and diffusion. Each time step $(\mathbf{u}^n, t^n) \mapsto (\mathbf{u}^{n+1}, t^{n+1})$ consists of the following two-part sequence:

1. **Wake Advancement** $(\mathbf{u}^n, t^n) \mapsto (\mathbf{u}_*^{n+1}, t_*^{n+1})$: It accounts for existing wake vortices' convection and placement of new wake vortex element. The new vortex element increases the length of the state vector by the number of degrees of freedom associated with the new element. The existing wake may be displaced following the local flow velocity, as described by the free-wake model. The propagation equation for the wake in discrete time is given by:

$$\zeta_w^{t+\Delta t} = \zeta_w^t + \int_t^{t+\Delta t} \left(B_{w_1 b_v} \left(\zeta_b^t, \zeta_w^t \right) \Gamma_b^t + B_{w_1 w_2} \left(\zeta_w^t \right) \Gamma_w^t + V_\infty \right) dt \tag{1}$$

where V_∞ is velocity due to freestream conditions and $B_{w_1 b_v}$ and $B_{w_{v1} w_{v2}}$ are defined as influence coefficients for the interaction between wake vortices and bound vortices, and between wake vortices themselves, respectively.

$$\begin{aligned} B_{w_{v1} w_{v2}} \left(\zeta_w^t \right) &= \vec{q}_{w_{v1} w_{v2}} && \text{for} \quad w_{v1}, w_{v2} = 1, \ldots, N_w \\ B_{w_{v1} b_v} \left(\zeta_b^t, \zeta_w^t \right) &= \vec{q}_{w_{v1}, b_v} \hat{\eta}_c && \text{for} \quad b_v = 1, \ldots, N_b \quad \text{and} \quad w_{v1} = 1, \ldots, N_w. \end{aligned} \tag{2}$$

2. **Wake Modification** $(\mathbf{u}_*^{n+1}, t_*^{n+1}) \mapsto (\mathbf{u}^{n+1}, t^{n+1})$ In this step, the positions of wake vortices have been updated, and the strengths of wake vortices are recalculated. As the circulation of the wing changes, new rows of vortex rings are shed into the wake from the trailing edge of each lifting surface.

   The equations for quasi-steady (wake-free) circulation are:

$$\begin{aligned} \gamma_0(\theta, t) = & \frac{1}{R} \left[-A_1 \cos\theta - \left(A_2 + \frac{1}{2} A_7 \right) \cos 2\theta + A_3 \sin\theta \right. \\ & + \left(A_4 - \frac{1}{2} A_5 + \frac{1}{2} A_6 \right) \sin 2\theta - A_8 \sin\theta \cos 2\theta \\ & + A_9 \sin\theta \sin 2\theta - A_{10} \cos\theta \cos 2\theta + A_{11} \cos\theta \sin 2\theta \\ & \left. + \frac{\Gamma_0(t)}{2\pi} \right] - 2U_\infty \sin(\theta - \alpha(t)) \end{aligned} \tag{3}$$

$$\begin{aligned} \Gamma_0(t) = & 2\pi \left[2R \left(\left(\dot{l}(t) - U_\infty \right) \sin\alpha(t) + \dot{h}(t) \cos\alpha(t) \right) \right. \\ & \left. + \dot{\alpha}(t) \left(\frac{1}{2}\tau^2 + \frac{1}{2}\sigma^2 - 2R(R+a) \right) \right] \end{aligned} \tag{4}$$

   where the symbols $A_i, i = 1 \ldots 11$ are mentioned in reference [10].

   Elemental vorticity from the above equation can be converted to elemental circulation (which is the same in both circle and physical planes) using the relation:

$$\delta\Gamma = \delta\gamma_z \delta Z \quad (= \delta\gamma_\zeta \delta\zeta) \tag{5}$$

   Z and ζ correspond to the circle and physical planes respectively.

   The expression for wake-induced vorticity is discretised as follows:

$$\begin{aligned}\gamma_1(\theta,t)|_{wake+lev} = &-\frac{\Gamma_0(t)}{2\pi R} \\ &-\frac{1}{2\pi R}\left[\sum_{i=1}^{n_{wake}} \Re\left(\frac{Z_{wake_i}+Re^{i\theta}}{Z_{wake_i}-Re^{i\theta}}\right)\delta\Gamma_{wake_i}\right. \\ &\left.+\sum_{i=1}^{n_{lev}} \Re\left(\frac{Z_{lev_i}+Re^{i\theta}}{Z_{lev_i}-Re^{i\theta}}\right)\delta\Gamma_{lev_i}\right]\end{aligned} \tag{6}$$

where the subscripts $_{wake}$ and $_{lev}$ refer to the trailing-edge wake and leading-edge vortex respectively, and n refers to the number of shed vortices.

Wake modification step is computed by simultaneously solving the below discretised equations:

(a) **Condition 1**:

$$\begin{aligned}\Gamma_0(t) = -&\left[\sum_{i=1}^{n_{wake}} \Re\left(\frac{Z_{wake_i}+R}{Z_{wake_i}-R}\right)\delta\Gamma_{wake_i}\right. \\ &\left.+\sum_{i=1}^{n_{lev}} \Re\left(\frac{Z_{lev_i}+R}{Z_{lev_i}-R}\right)\delta\Gamma_{lev_i}\right]\end{aligned} \tag{7}$$

(b) **Condition 2**:

$$\begin{aligned}\frac{1}{R}\left[A_1-\left(A_2+\frac{1}{2}A_7\right)+A_{10}\right] - 2U_\infty \sin\alpha(t) = \\ \frac{1}{2\pi R}\left[\sum_{i=1}^{n_{wake}} \Re\left(\frac{Z_{wake_i}-R}{Z_{wake_i}+R}\right)\delta\Gamma_{wake_i}\right. \\ \left.+\sum_{i=1}^{n_{lev}} \Re\left(\frac{Z_{lev_i}-R}{Z_{lev_i}+R}\right)\delta\Gamma_{lev_i}\right]\end{aligned} \tag{8}$$

In order to solve for the latest wake and LEV circulations, the equations can be rearranged as follows:

$$\begin{aligned}\Re\left(\frac{Z_{wake_n}+R}{Z_{wake_n}-R}\right)\delta\Gamma_{wake_n} + \Re\left(\frac{Z_{lev_n}+R}{Z_{lev_n}-R}\right)\delta\Gamma_{lev_n} \\ = -\Gamma_0(t) - f_1(wake) - g_1(lev)\end{aligned} \tag{9}$$

$$\begin{aligned}\Re\left(\frac{Z_{wake_n}-R}{Z_{wake_n}+R}\right)\delta\Gamma_{wake_n} + \Re\left(\frac{Z_{lev_n}-R}{Z_{lev_n}+R}\right)\delta\Gamma_{lev_n} \\ = 2\pi R\Upsilon_0(t) - f_2(wake) - g_2(lev)\end{aligned} \tag{10}$$

where:

$$f_1(wake) = \sum_{i=1}^{n-1} \Re\left(\frac{Z_{wake_i}+R}{Z_{wake_i}-R}\right)\delta\Gamma_{wake_i} \tag{11}$$

$$g_1(lev) = \sum_{i=1}^{n-1} \Re \left(\frac{Z_{lev_i} + R}{Z_{lev_i} - R} \right) \delta\Gamma_{lev_i} \tag{12}$$

$$f_2(wake) = \sum_{i=1}^{n-1} \Re \left(\frac{Z_{wake_i} - R}{Z_{wake_i} + R} \right) \delta\Gamma_{wake_i} \tag{13}$$

$$g_2(lev) = \sum_{i=1}^{n-1} \Re \left(\frac{Z_{lev_i} - R}{Z_{lev_i} + R} \right) \delta\Gamma_{lev_i} \tag{14}$$

$$\Upsilon_0(t) = \frac{1}{R} \left[A_1 - \left(A_2 + \frac{1}{2} A_7 \right) + A_{10} \right] + 2U_\infty \sin\alpha(t) \tag{15}$$

where $i = n$ refers to the latest shed vortex and $i = 1 \ldots n-1$ to all previous shed vortices. The subscripts $wake$ and lev have been dropped from n for clarity.

These are two linear simultaneous equations of the form $Ax + By = C$, which can be solved together for the circulations of the latest shed vortices $\delta\Gamma_{wake_n}$ and $\delta\Gamma_{lev_n}$. They still have two unknowns — the positions of the latest shed vortices (Z_{wake_n} and Z_{lev_n}). These are estimated from the movement of the shedding edge and the position of the previous shed vortex.

2.2 Nonlinear constraints

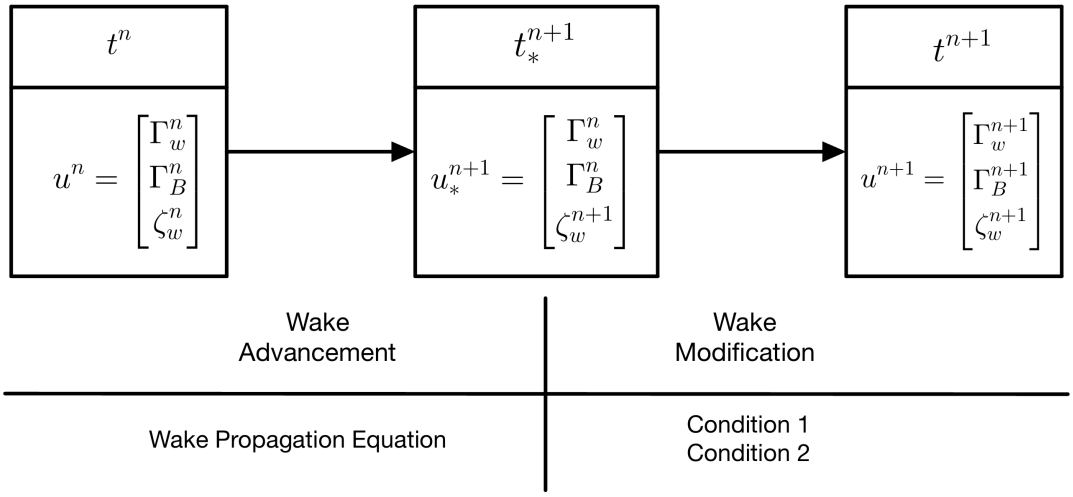


Figure 2: Nonlinear constraints

Nonlinear constraints can then be summarised as follows:

1. **Step 1: Wake Advancement** : At time instance n propagate the positions of the wake vortices using wake propagation eq. (Eq. (1)) and shed new vortices $\Gamma_{w_1}(t_n)$ at position (x_{w_1}, y_{w_1}) at trailing edge and $\Gamma_{l_1}(t_n)$ at position (x_{l_1}, y_{l_1}) at leading edge, respectively.

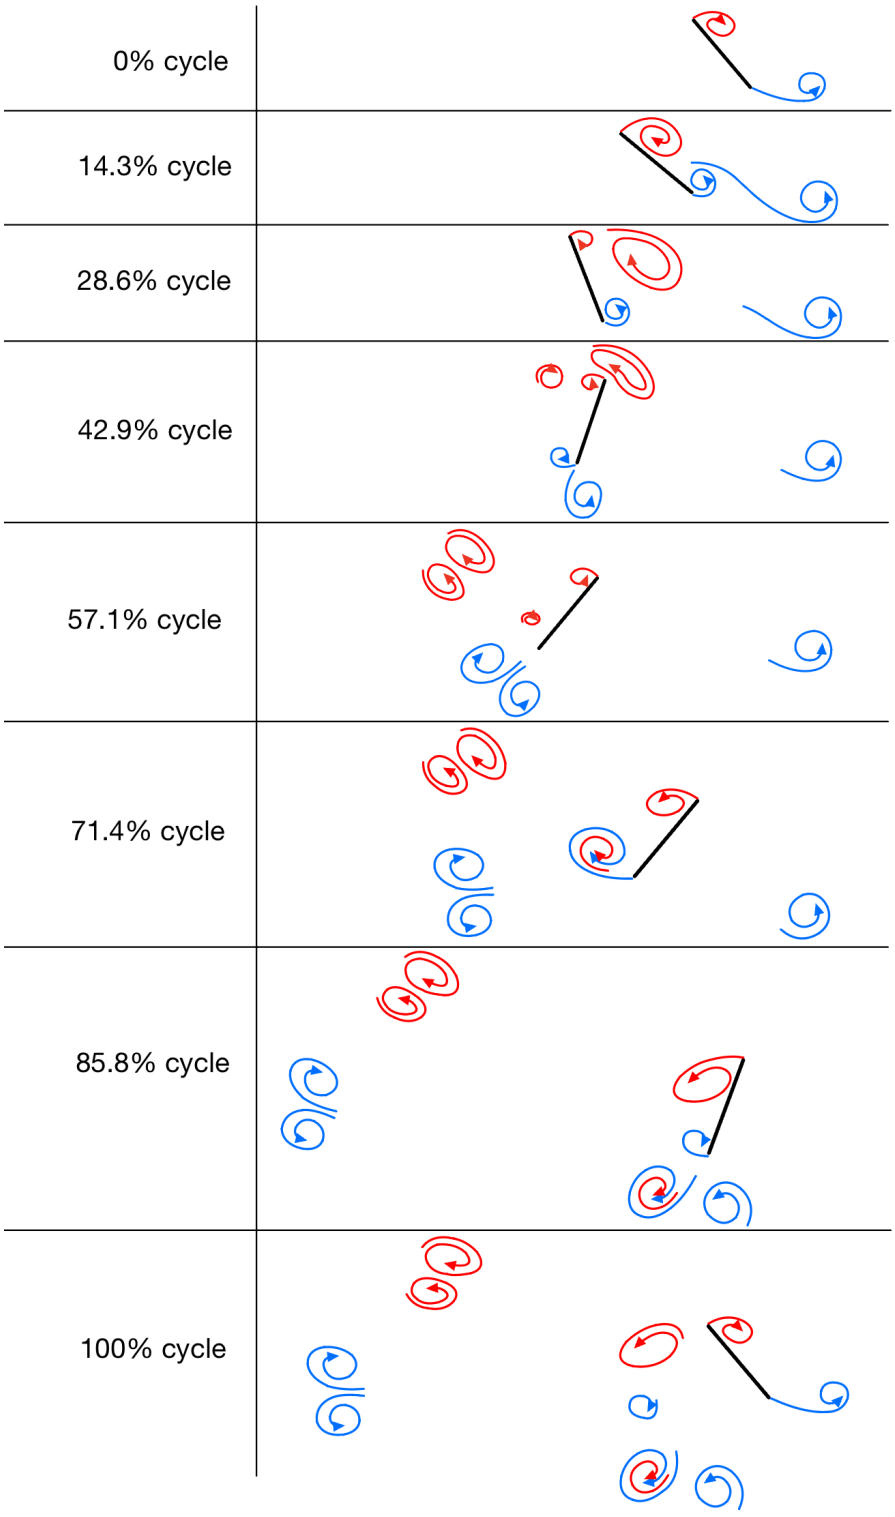


Figure 3: Initial value formulation

2. **Step 2: Wake Modification** :Solve Conditions 1 (Eq. (7)) and 2 (Eq. (8)) simultaneously for the strengths of the bound vortices $\Gamma_b(t_n)$ and that of the latest shed vortices $\Gamma_{w_1}(t_n)$ and $\Gamma_{l_1}(t_n)$.

3. **Step 3** :Increment time index n and go back to step 1 . Stop when n has reached

the desired value.

2.3 Trust Region Framework

We employ a trust region based optimization algorithm (4). This approach allows us to iteratively update the design variables by computing the gradient of the objective function with respect to the design variables. The optimization algorithm modifies the design variables in the direction that minimizes the objective function. The optimization framework includes design variables, an objective function and nonlinear constraints. By minimizing the objective function, the optimization algorithm aims to converge towards a wake configuration that is equivalent to the periodic convergence.

Python serves as our primary programming language within our optimization framework. In this framework, we employ the nonlinear equality constraints Large-scale Optimization method. This approach allows us to efficiently minimize the objective function while ensuring that all constraints are satisfied.

One noteworthy feature of our framework is its utilization of Python's parallel computing capabilities. This aspect significantly enhances the overall efficiency of our optimization process, leading to faster and more scalable results. Python's capacity for parallel computing enables the distribution of tasks across multiple processors, making it a valuable asset when dealing with computationally intensive optimization problems. Furthermore, our framework integrates the use of sparse matrices for the gradient and Hessian matrices as we are providing analytic expression for gradient. By doing so, we effectively reduce memory consumption, an essential aspect when dealing with large-scale optimization tasks. This optimized memory usage allows our framework to handle complex and extensive optimization problems with remarkable efficiency.

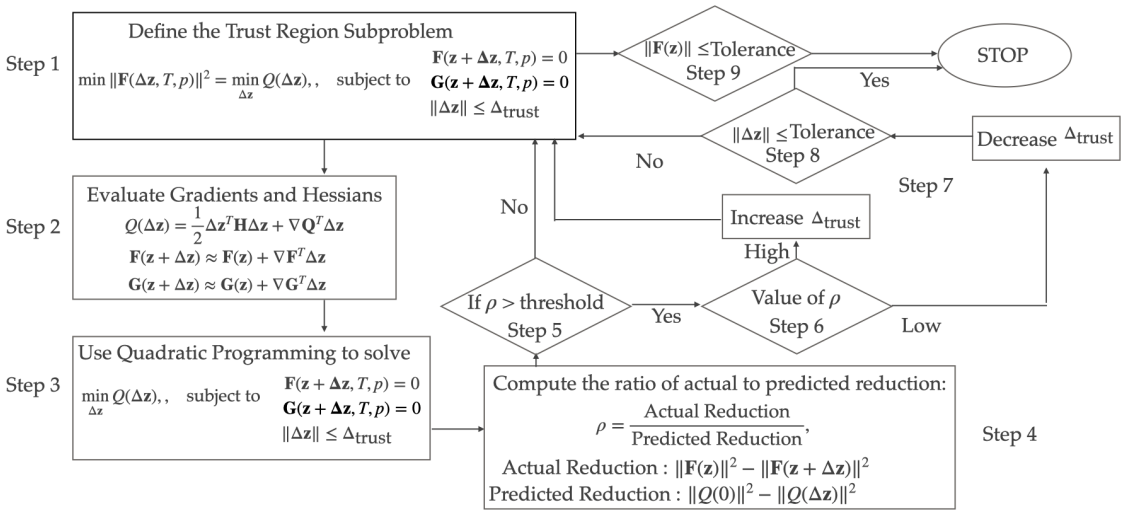


Figure 4: Trust-region algorithm

The design variables are wake strengths and positions of shed wake and wake strength of bound vortex. The choice of optimization methodology for the wake modeling framework depends on several factors, including the complexity of the wake model, the dimensionality of the optimization problem, and the desired computational efficiency. Different optimization algorithms can be employed different methods for various configurations of

objective functions and nonlinear constraints and found out the optimal objective function and nonlinear constraints which satisfy the computational efficiency and solution quality (6).

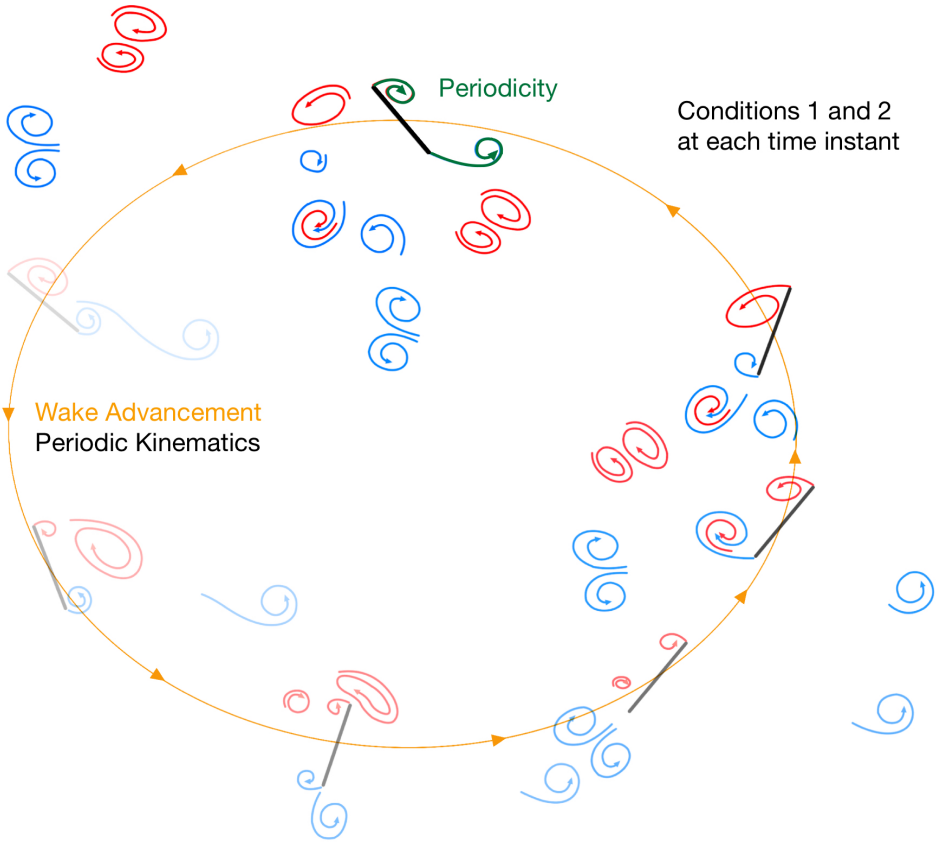


Figure 5: Periodicity and manifold

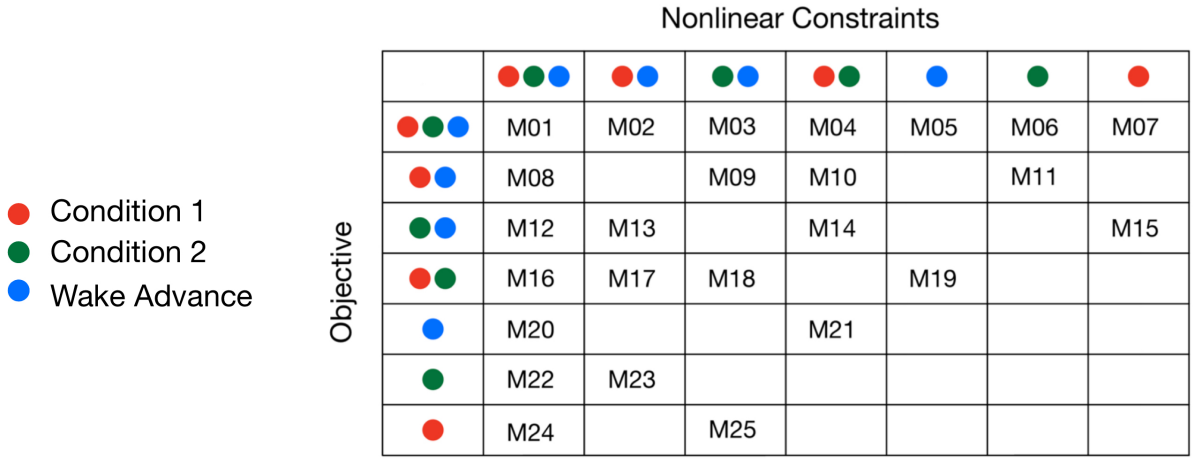


Figure 6: Various Methodologies based on Choice of nonlinear constraints and objective functions

3 RESULTS

The following section presents key findings from the validation and comparative analysis of the aerodynamic simulation methods explored in this study.

3.1 Validation of the Aerodynamic IVP Models

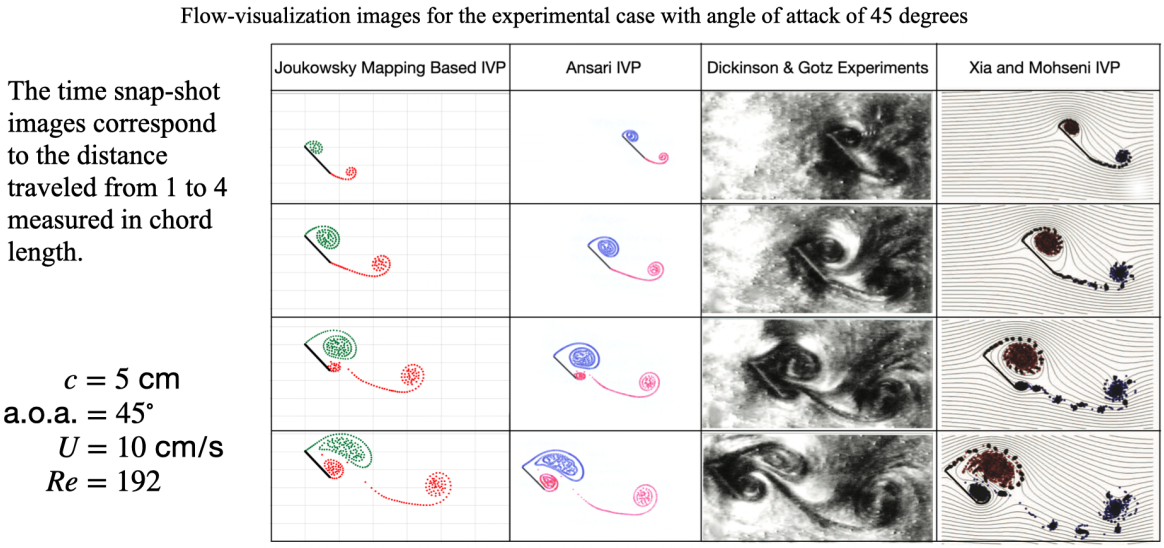


Figure 7: Validation of Joukouwsky conformal mapping based IVP model

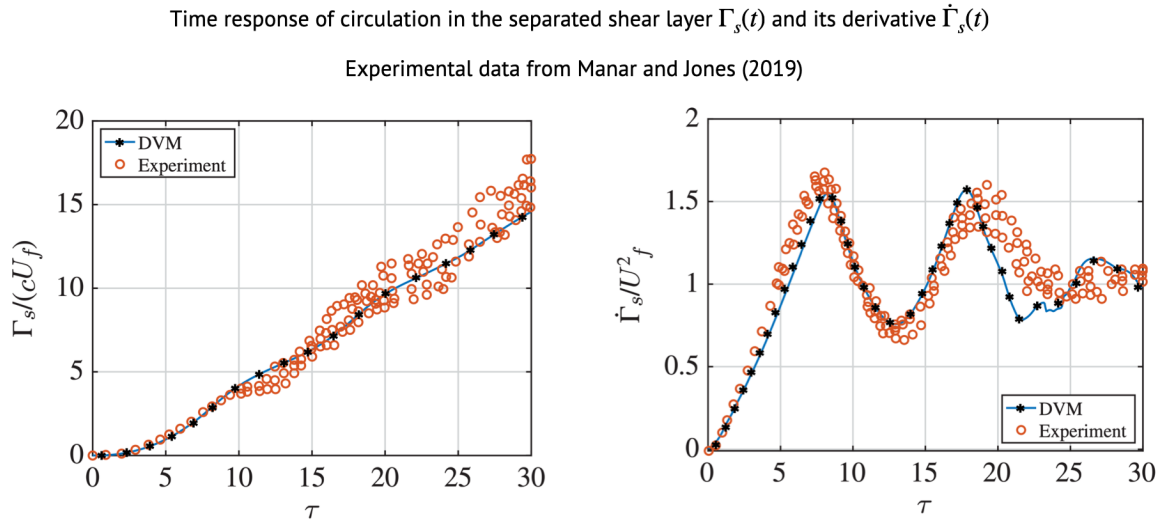


Figure 8: Validation of Joukouwsky conformal mapping based IVP model

The unsteady aerodynamic model based on the initial value problem (IVP) formulation was validated against experimental data from Manar and Jones (2019), as shown in 7 and 8. The time evolution of circulation $\Gamma_s(t)$ and its time derivative $\dot{\Gamma}_s(t)$ within the separated shear layer demonstrates excellent agreement between the Discrete Vortex Method (DVM) simulations and experimental measurements.

3.2 Computational Cost IVP vs BVP

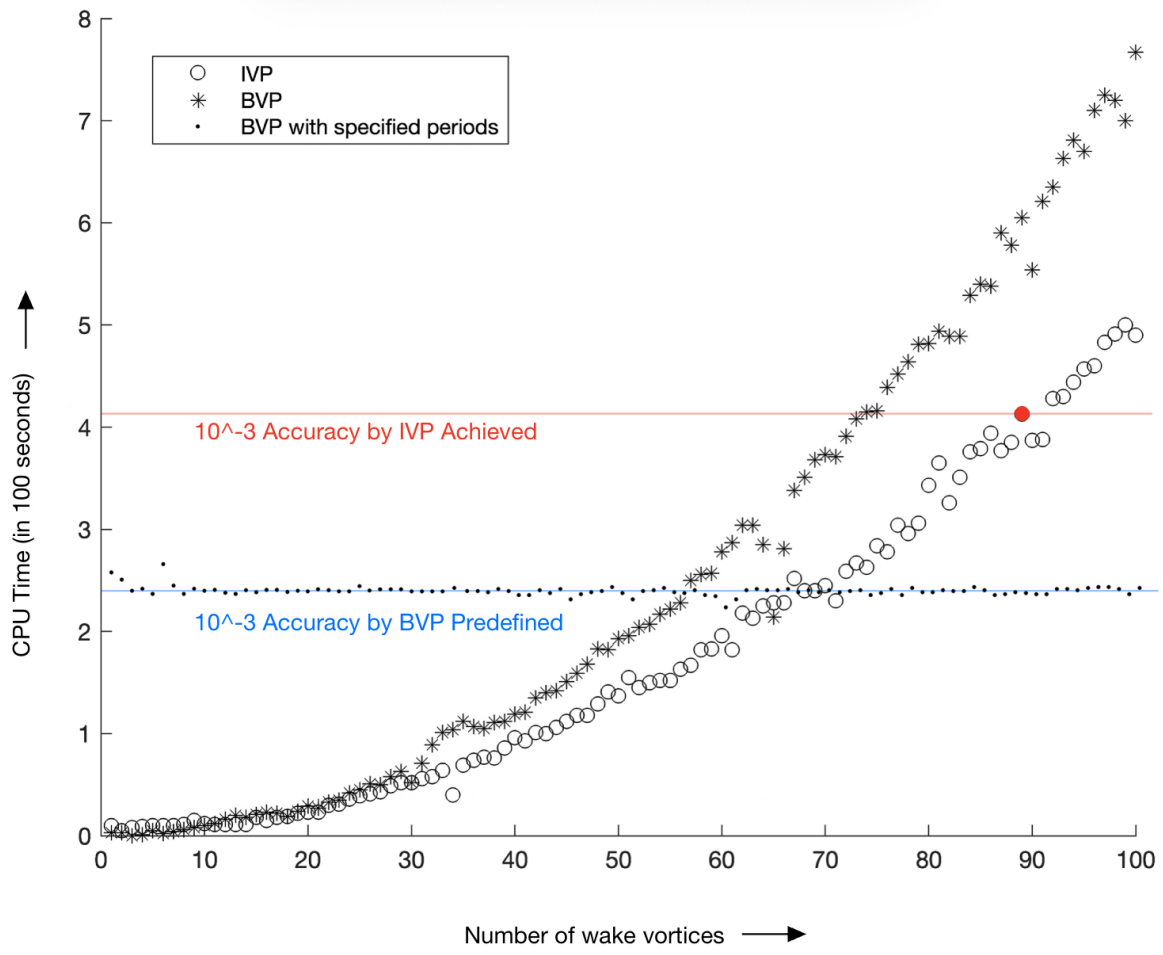


Figure 9: Computational cost - IVP vs BVP

The figure 9 clearly demonstrates the computational performance comparison between the initial value problem (IVP), standard boundary value problem (BVP), and BVP with specified periods for systems with increasing numbers of wake vortices.

While the IVP method (open circles) shows a moderate rise in CPU time with the number of vortices, it only achieves the target accuracy of 10^{-3} at around 90 vortices, requiring over 400 seconds of CPU time (marked by the red horizontal line and point). The standard BVP approach (asterisks), although accurate, becomes even more computationally expensive, showing a steep rise in CPU time beyond 60 vortices and exceeding the IVP cost at high vortex counts.

In contrast, the BVP with specified periods (small black dots) consistently maintains a nearly flat and low computational cost—around 250 seconds—regardless of the number of vortices, as indicated by the blue horizontal line. This demonstrates that the BVP with specified periods not only meets the desired 10^{-3} accuracy but does so with significantly improved scalability and robustness, making it the most efficient method for large-scale periodic vortex simulations where both precision and computational efficiency are crucial.

3.3 Wake Structure Comparison: IVP vs BVP

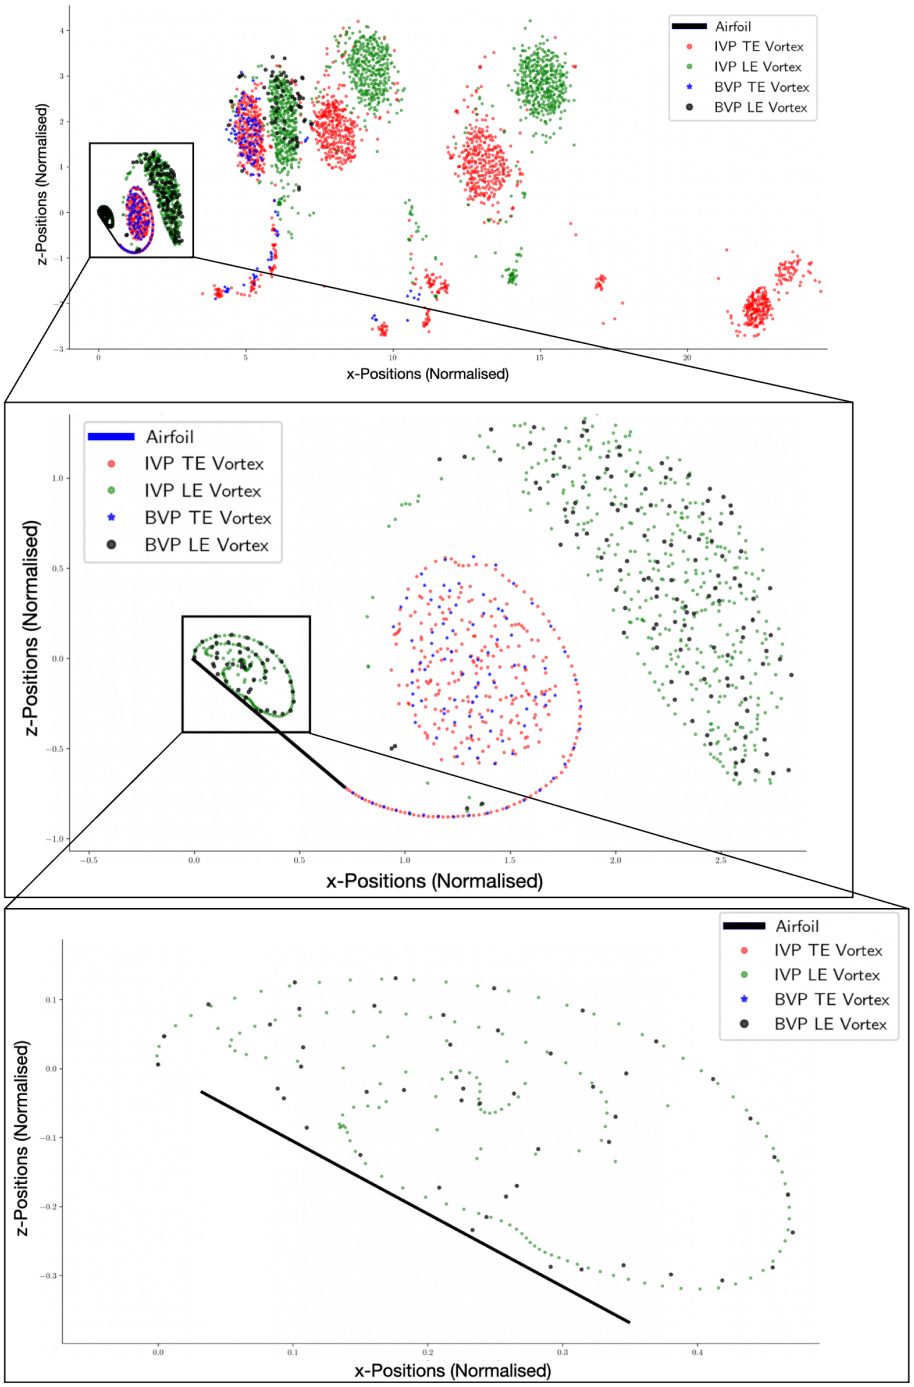


Figure 10: IVP vs BVP - Joukowsky conformal mapping based inviscid aerodynamic model - wake plots

Wake evolution and vortex shedding patterns were compared between IVP and BVP formulations using visualizations shown in 10. The normalized positions of leading-edge (LE) and trailing-edge (TE) vortices are plotted for both approaches. While both IVP and BVP capture the expected alternating vortex patterns and coherent shedding, the BVP solutions demonstrate slightly more compact and consistent wake structures. This

suggests that BVP inherently enforces periodicity with higher spatial symmetry, whereas IVP results may retain minor residual transients unless integrated over long durations. In particular, the BVP-generated wake forms show a tighter clustering of vortices indicative of faster convergence to a periodic limit cycle.

3.4 Lift coefficient (C_l comparison: IVP vs BVP)

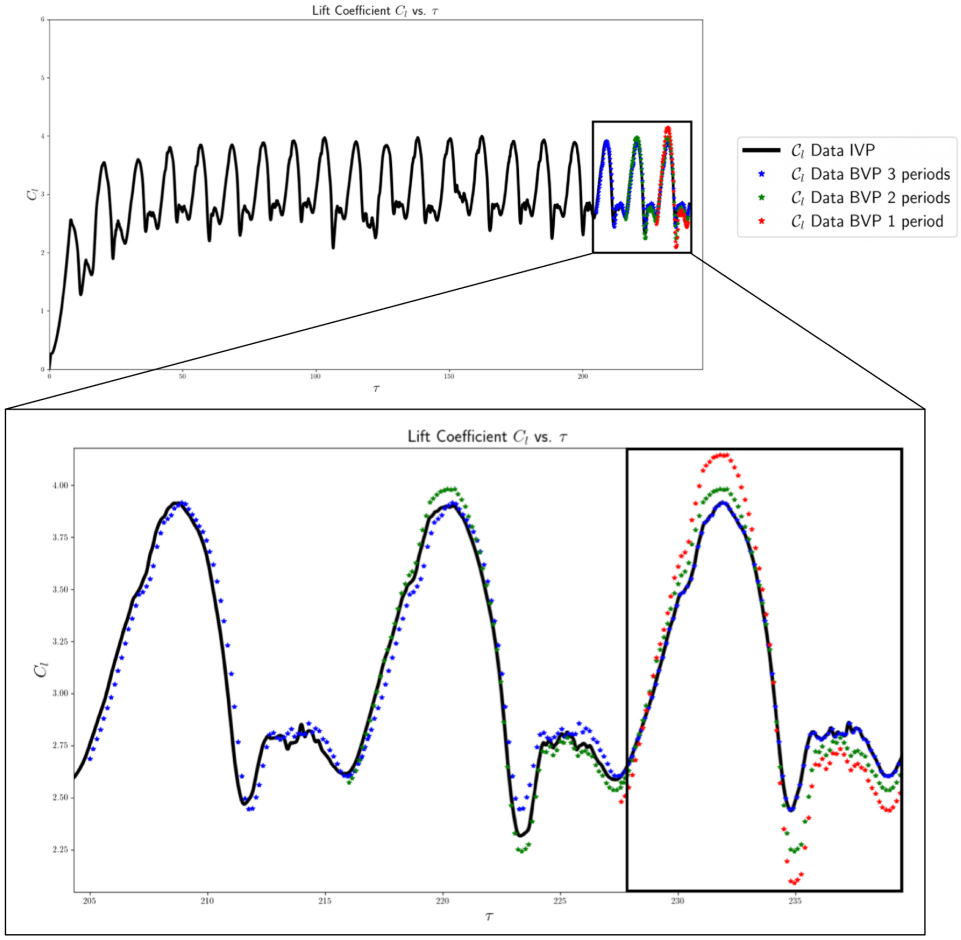


Figure 11: IVP vs BVP - Joukowsky conformal mapping based inviscid aerodynamic model - C_l values

The lift coefficient history, shown in 11, provides a quantitative comparison of unsteady aerodynamic forces from the IVP and BVP approaches. The IVP simulation converges to a periodic lift response over several cycles, while BVP solutions obtained with 1, 2, and 3-period window constraints show varying levels of accuracy. The 3-period BVP data aligns most closely with the IVP reference, replicating the amplitude, phase, and waveform of the C_l signal. Notably, even the 1-period BVP captures the essential features of the waveform, though with slight discrepancies in peak values. This reinforces the suitability of the BVP framework in replicating long-time periodic responses with significantly reduced computational burden.

4 SUMMARY AND CONCLUSIONS

In this work, we have developed and demonstrated a boundary value problem of nonlinear inviscid unsteady aerodynamic systems with periodic boundary conditions. The parameters of the model - model dimensionality (k_0) and dynamic richness (p_0) are chosen to obtain a rich taxonomy of vortex element models. Key findings are as follows:

1. The existence and uniqueness of the periodic solution depend on the implementation of all boundary conditions and the elimination of translational invariance, respectively.

2. Periodic solution is determined by solving for only one time period.

3. A new general starting and stopping scheme is devised by finding k_0 and p_0. This enables the balance between model dimensionality k_0, dynamic richness p_0 and accuracy.

4. Computational savings are achieved by preventing the number of elements from growing indefinitely.

5. Periodic matrix is the state transition matrix for nonlinear inviscid unsteady aerodynamic systems

The proposed procedure restricts the number of computational elements required to capture the dynamics of vortex formation. It can capture the development and subsequent shedding of the wake vortices with lesser wall-clock time and remain sufficiently low-dimensional to simulate real-time applications.

We have also found good quantitative agreement with the theoretical prediction for small-amplitude oscillatory pitching, heaving and leading and the correct qualitative and quantitative behaviour for larger amplitudes of this motion. Furthermore, the present results closely agree with those of time marching methods. Although we have focused on a specific inviscid aerodynamic model in this conference paper, the proposed methodology works for any initial value formulation due to the generality of this method.

One can distinguish several advantages using this periodicity based vortex method. First, its Lagrangian approach with particle discretization makes it intrinsically close to the flow physics it aims to model and provides a direct scheme to solve the transport of particles and the quantities they carry, based on ODEs systems, without numerical dissipation. Secondly, the vorticity field is a compact field compared to the velocity. It is therefore much more efficient to directly solve and simulate the vorticity field. Furthermore, in this vortex method the boundary condition at infinity can be automatically satisfied contrary to most grid-based approaches where external flows must be solved within a finite simulation domain. Finally, the stability condition provided by the vortex method scheme allows larger time steps compared to grid-based methods.

ACKNOWLEDGEMENTS

I wish to record a deep sense of gratitude to the Department of Aerospace Engineering, IIT Bombay, Prime Minister Research Fellowship (PMRF), Ministry of Education (formerly MHRD), and Anusandhan National Research Foundation (ANRF) Fellowship for their invaluable support throughout the course of my Ph.D. studies and related research.

REFERENCES

[1] E. Reithmeier. *Periodic solutions of nonlinear dynamical systems: Numerical computation, stability, bifurcation and transition to chaos*. Springer, 2006.

[2] H. Poincaré. *Les méthodes nouvelles de la mécanique céleste*, volume 2. Gauthier-Villars et fils, imprimeurs-libraires, 1893.

[3] E. Hopf. Abzweigung einer periodischen lösung von einer. *Selected Works of Eberhard Hopf with Commentaries: With Commentaries*, 17:91, 2002.

[4] R. Seydel. Branch switching in bifurcation problems for ordinary differential equations. *Numerische Mathematik*, 41:93–116, 1983.

[5] M. Holodniok and M. Kubíček. Derper—an algorithm for the continuation of periodic solutions in ordinary differential equations. *Journal of Computational Physics*, 55(2): 254–267, 1984.

[6] R. Bulirsch, J. Stoer, and P. Deuflhard. Boundsol, a fortran code for a numerical solution of boundary-value problems of odes. *Technische Universität, München*, 1977.

[7] G. Dimitriadis. *Unsteady Aerodynamics: Potential and Vortex Methods*. John Wiley & Sons, 2023.

[8] J. Katz and A. Plotkin. *Low-speed aerodynamics*, volume 13. Cambridge university press, 2001.

[9] Ü. Gülçat. *Fundamentals of modern unsteady aerodynamics*, volume 2010. Springer, 2010.

[10] S. A. Ansari, R. Żbikowski, and K. Knowles. Non-linear unsteady aerodynamic model for insect-like flapping wings in the hover. part 2: implementation and validation. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 220(3):169–186, 2006.

[11] S. Ansari, R. Żbikowski, and K. Knowles. Non-linear unsteady aerodynamic model for insect-like flapping wings in the hover. part 1: methodology and analysis. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 220(2):61–83, 2006.

[12] T. H. von Karman and W. R. Sears. Airfoil theory for non-uniform motion. *Journal of the Aeronautical Sciences*, 5(10):379–390, 1938.

[13] D. Darakananda and J. D. Eldredge. A versatile taxonomy of low-dimensional vortex models for unsteady aerodynamics. *Journal of Fluid Mechanics*, 858:917–948, 2019.

[14] S. Boyd and L. Chua. Fading memory and the problem of approximating nonlinear operators with volterra series. *IEEE Transactions on circuits and systems*, 32(11): 1150–1161, 1985.

[15] A. R. Conn, N. I. Gould, and P. L. Toint. *Trust region methods*. SIAM, 2000.

[16] G. J. Leishman. *Principles of helicopter aerodynamics with CD extra*. Cambridge university press, 2006.

[17] A. Bagai and J. G. Leishman. Rotor free-wake modeling using a pseudoimplicit relaxation algorithm. *Journal of Aircraft*, 32(6):1276–1285, 1995.

[18] W. Johnson. A general free wake geometry calculation for wings and rotors. In *AHS, Annual Forum, 51 st, Fort Worth, TX*, pages 137–153, 1995.

[19] T. R. Quackenbush, D. B. Bliss, C. C. Ong, and C. O. Ching. Free wake analysis of hover performance using a new influence coefficient method. Technical report, NASA, 1990.

[20] W. MILLER and D. BLISS. Direct periodic solutions of rotor free wake calculations by inversion of a linear periodic system. In *AHS, Annual Forum, 46 th, Washington, DC*, pages 757–769, 1990.

[21] A. Bagai and J. G. Leishman. Free-wake analysis of tandem, tilt-rotor and coaxial rotor configurations. *Journal of the American Helicopter Society*, 41(3):196–207, 1996.

[22] A. Mehta and A. Gogulapati. An approach to determine periodic solutions of nonlinear inviscid unsteady aerodynamic systems. In *AIAA SCITECH 2023 Forum*, page 0999, 2023.

[23] A. Mehta, P. Mishra, and A. Gogulapati. A continuation method based solution for periodic nonlinear aeroelastic problems. In *AIAA SCITECH 2025 Forum*, page 1202, 2025.

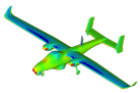
[24] B. A. Roccia, S. Preidikman, J. C. Massa, and D. T. Mook. Modified unsteady vortex-lattice method to study flapping wings in hover flight. *AIAA journal*, 51(11): 2628–2642, 2013.

[25] I. Bendixson. Sur les courbes définies par des équations différentielles. *Acta Mathematica*, 24:1–88, 1901.

[26] J. D. Eldredge. *Mathematical modeling of unsteady inviscid flows*, volume 50. Springer, 2019.

[27] M. Drela. *Flight vehicle aerodynamics*. MIT press, 2014.

[28] M. J. Lighthill. Introduction boundary layer theory. *Laminar boundary layers*, 1963.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

PRELIMINARY NUMERICAL AND EXPERIMENTAL COMPARISON OF TWO TYPES OF PLA BCC LATTICE STRUCTURES MANUFACTURED VIA FDM AND SUBJECTED TO IMPACT TESTING

G. Iacolino[1]*, C. Orlando[1], D. Tumino[1], G. Catalanotti[1], E.V. González[2], G. Mantegna[1] and A. Alaimo[1]

1: Department of Engineering and Architecture
Kore University of Enna
Cittadella Universitaria, 94100, Enna (EN), Italy
*: giuseppe.iacolino@unikorestudent.it
2: AMADE, Polytechnic School, Universitat de Girona
Campus Montilivi s/n, Girona, 17071, Spain

Abstract *Lattice structures have gained increasing attention due to their exceptional mechanical properties combined with significant weight reduction. Recently, their application in engineering has expanded, driven by the ease of production enabled by innovative technologies in Additive Manufacturing such as Fused Deposition Modelling (FDM). Additionally, Polylactic Acid (PLA) offers excellent mechanical properties alongside environmental sustainability, making it an ideal material for aerospace, automotive, and other industries requiring efficient energy absorption and lightweight solutions. In this regard, this study investigates the impact response of PLA body-centered cubic (BCC) lattice structures manufactured via FDM. Specifically, a preliminary numerical-experimental comparative analysis was conducted on two BCC configurations with strut diameters of 1.5 mm and 2 mm. The experimental tests were performed using a drop-tower setup, while the numerical simulations were carried out in Ansys LS-DYNA. The findings demonstrate the suitability of FDM-printed PLA BCC lattice structures for applications requiring lightweight energy-absorbing components. Moreover, the integration of numerical and experimental approaches confirms the accuracy of simulations in predicting structural behaviour under dynamic loading conditions.*

Keywords: Lattice Structures, Additive Manufacturing, PLA, Drop-Tower Impact Testing, Numerical-Experimental Comparison

1. INTRODUCTION

Cellular solids represent a specific macro-area of structures employed in the engineering field. In [1], Tao et al. classify these structures into two categories: stochastic structures (foams), and non-stochastic structures (lattices) reported in Fig. 1. As highlighted in [2], they can be employed in sectors such as aerospace, automotive, infrastructural, and biomedical industries. The distinctive feature of these structures lies in their ability, unlike other types, to exhibit excellent mechanical properties [3]. The literature comprises numerous studies investigating the mechanical properties of BCC structures, such as [4], and [5]. Furthermore, others investigate their performance under impact testing. For instance, Acanfora et al. [6], analyzed impact absorbers with BCC cores manufactured via additive manufacturing (AM) and compared them to honeycomb structures. The results demonstrate that BCC cores, combined with composite skins, absorb more energy per unit mass than honeycomb cores. Bernard et al. [7] numerically examined 24 lattice topologies under dynamic loading and highlighted the moderate but competitive crashworthiness of BCC structures, driven by combined bending and stretching mechanisms. They concluded that adjusting the overall relative density of the lattice — by uniformly varying the strut dimensions throughout the structure — can further enhance energy absorption.

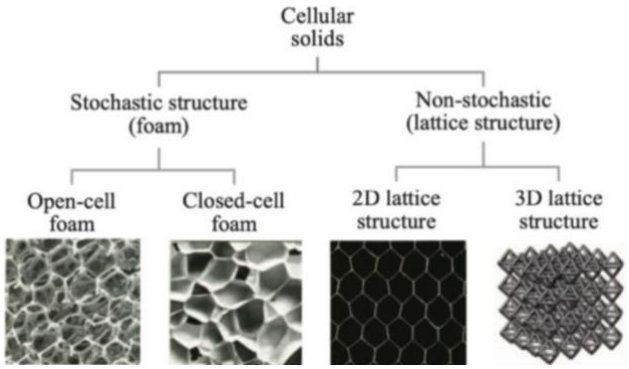


Figure 1. Cellular Solids [1].

From a manufacturing standpoint, AM technologies have become increasingly popular due to their ability to fabricate complex geometries with minimal material waste [8]. In this study, the Fused Deposition Modelling (FDM) technique was utilized to produce BCC structures, featuring a periodic arrangement of single units within the *X-Y* plane. Also, it is worth mentioning that the core optimization process is not included in this work, which solely focuses on a numerical-experimental comparison.

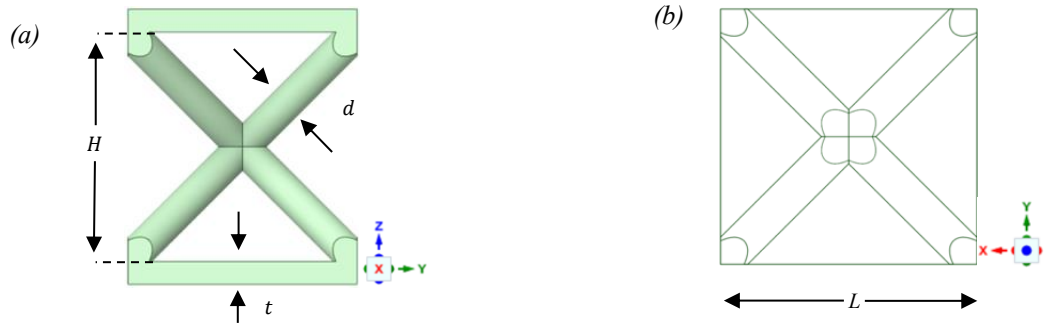


Figure 2. *(a)* lateral view of a single unit; *(b)* translucent top view of a single unit.

In Fig. 2(a), the single unit cell is provided in a lateral view. In Fig. 2(b), the top view is shown with a translucent skin to reveal the internal struts.
Each unit has a total thickness of 12 mm, consisting of a 10 mm cell thickness (H) and two skins (t), each 1 mm thick. Moreover, each unit features a side length of $L = H = 10$ mm. This study employs two types of BCC structures, which differ in their struts' diameter (d). Specifically, the acronym BCC1.5 denotes a configuration composed of units with $d = 1.5$ mm, whereas BCC2 represents a configuration with $d = 2$ mm.

2. EXPERIMENTAL TESTS

The samples employed in this experimental campaign were printed at the Kore University of Enna, Italy, using two FDM printers: *Flashforge Dreamer* and *Flashforge Adventurer M5 Pro.* The printing angle in the X-Y plane, as shown in Fig. 2, was set at ±45°. The material used is PLA, whose mechanical properties are listed in Tab. 1. In order to obtain these properties, three dog-bone specimens were tested under tensile test conditions in accordance with ASTM D638, with an average weight of 10.26 g and a standard deviation of 0.5 g.

Table 1. Mechanical Properties of PLA at ±45°.

E [MPa]	Average	3232.0
	STDV	145.05
σ_{yield} [MPa]	Average	33.90
	STDV	2.03
σ_{max} [MPa]	Average	42.39
	STDV	1.77
ε_{break}	Average	0.0528
	STDV	0.0110

Figs. 3 and 4 illustrate the two specimens employed in the tests (one specimen impacted for each BCC type), which comply with ASTM D7136 standard [20] and have dimensions of 150×100 mm. Before the impact test, each specimen was measured at multiple locations using an electronic caliper. Mean values for width W_m, length L_m and thickness t_m were then calculated and are reported in Tab. 2.

Table 2. Mean geometric values.

Mean values [mm]	BCC1.5	BCC2
W_m	100.23	99.99
L_m	150.06	150.20
t_m	12.10	12.20

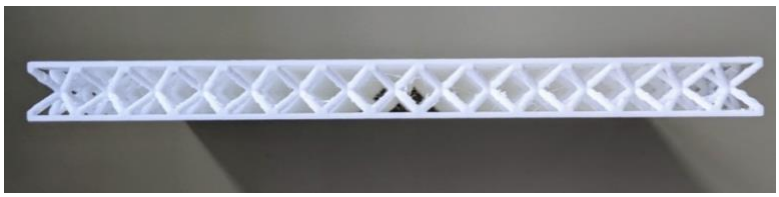
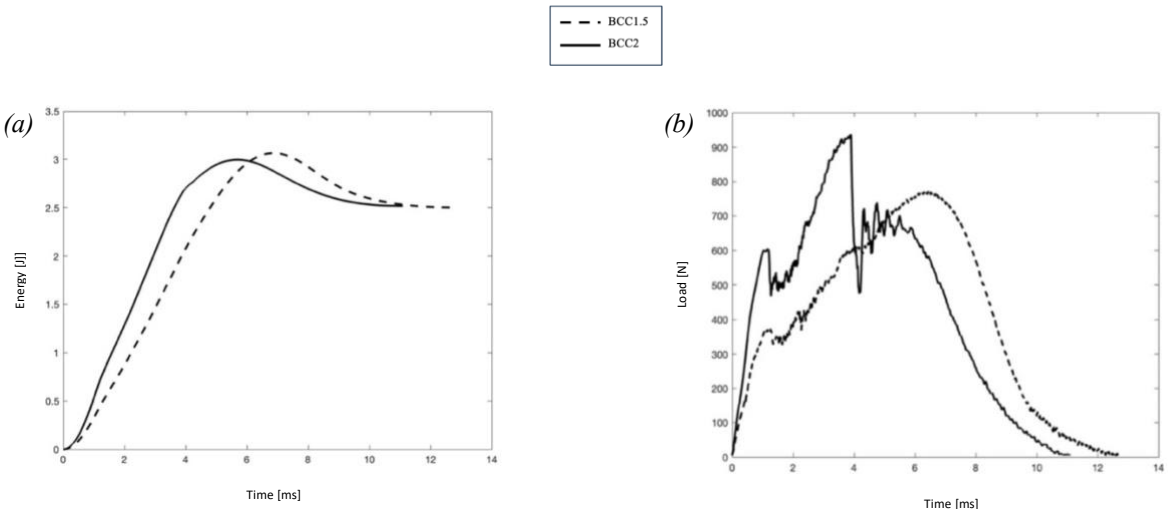


Figure 3. BCC1.5.



Figure 4. BCC2.

The two experimental tests were conducted at the AMADE research center of the University of Girona, Spain, using the *CEAST Fractovis Plus* impact tester. The specimens were secured using mechanical clamps, and an impact test were subsequently conducted under identical conditions for each BCC type. In both cases, a semi-spherical impactor with a diameter of 16 mm and a mass m of 2 kg was dropped from a height h of 150 mm. The clamps are symmetrically represented as cylindrical elements with a diameter of 14 mm and a height of 4 mm. Additionally, the supporting base includes a central rectangular void, ensuring that the specimen is supported solely along its edges. The external and internal rectangular contours of this void measure 200×150 mm and 125×75 mm. To avoid a second impact, the impactor was rebounded and caught immediately after the first contact. Following the experimental tests, the time–energy, time–load, time–displacement, and load–displacement diagrams were obtained (Fig. 5). The images of the structures after the impact are presented in Fig. 6.



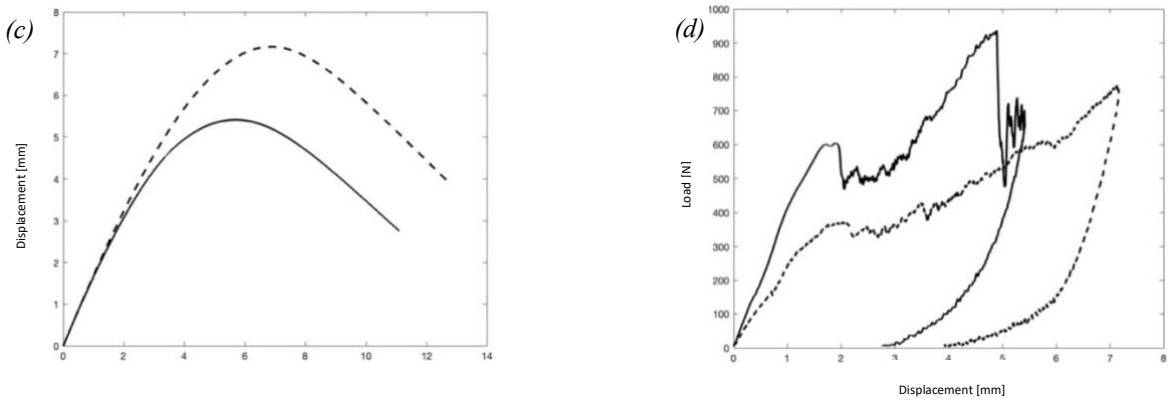


Figure 5. Experimental *(a)* Time – Energy, *(b)* Time – Load, *(c)* Time – Displacement, *(d)* Displacement – Load.

Following the data analysis shown in Fig. 5, it was observed that both BCC1.5 and BCC2 exhibit a similar energy plateau (2.50 J). Furthermore, BCC1.5 shows higher displacement (7.16 mm) compared to BCC2 (5.41 mm).
However, BCC2 withstands a greater maximum load of 937.17 N, exceeding the 772.81 N recorded for BCC1.5.

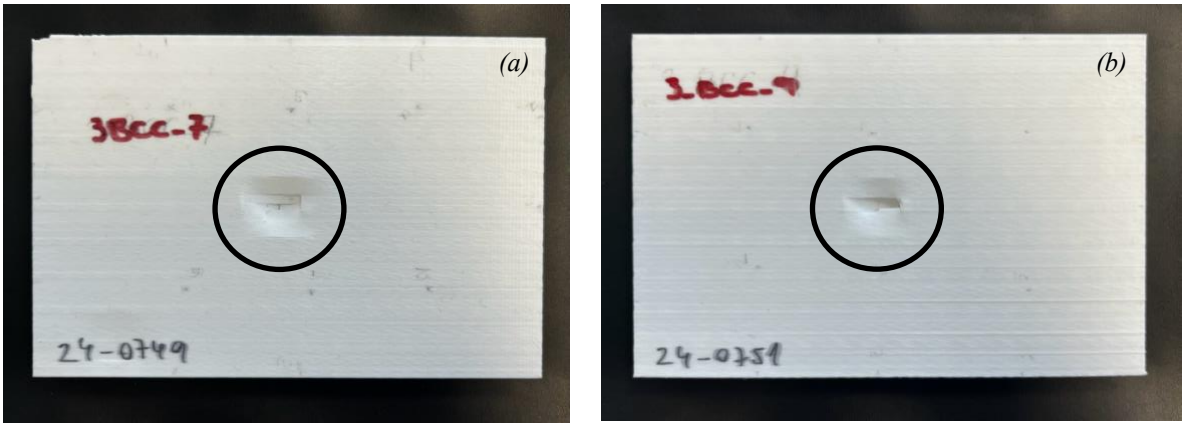


Figure 6. Damage after impact comparison between *(a)* BCC1.5 and *(b)* BCC2.

3. NUMERICAL TESTS

The numerical simulations were conducted using Ansys LS-DYNA. Fig. 7 illustrates the numerical setup, which is identical for both BCC1.5 and BCC2 structures and reproduces the experimental low velocity impact clamping system of tests described in Section 2. Specifically, the figure highlights the following components: clamps, impactor, supporting base, and the lattice structure, whether BCC1.5 or BCC2. Both the structures are made of PLA, with their mechanical properties provided in Tab. 1. Additionally, a bilinear material model with a zero tangent modulus was implemented, corresponding to an elastic-perfectly plastic behavior. The material for the impactor was selected to ensure that its geometric characteristics would yield a

body with a mass compatible with the experimentally chosen value of 2 kg. Both the support base and the clamps were modeled as rigid. Specifically, all contacts were defined as frictional with a coefficient of friction set to 0.2, which was chosen purely based on numerical fitting considerations. Given the role of each component in the simulation, the mesh was designed with a variable element size depending on the object. Also, before conducting the analysis, appropriate boundary conditions were defined as follows: the base was fully constrained, a displacement perpendicular to the structure's upper skin was applied to the impactor, and a force of 50 N was applied to each clamp. Figs. 8 and 9 present the total displacement following the analysis in the BCC1.5 and BCC2 samples respectively, while Fig. 10 provides a comparison between experimental and numerical results.

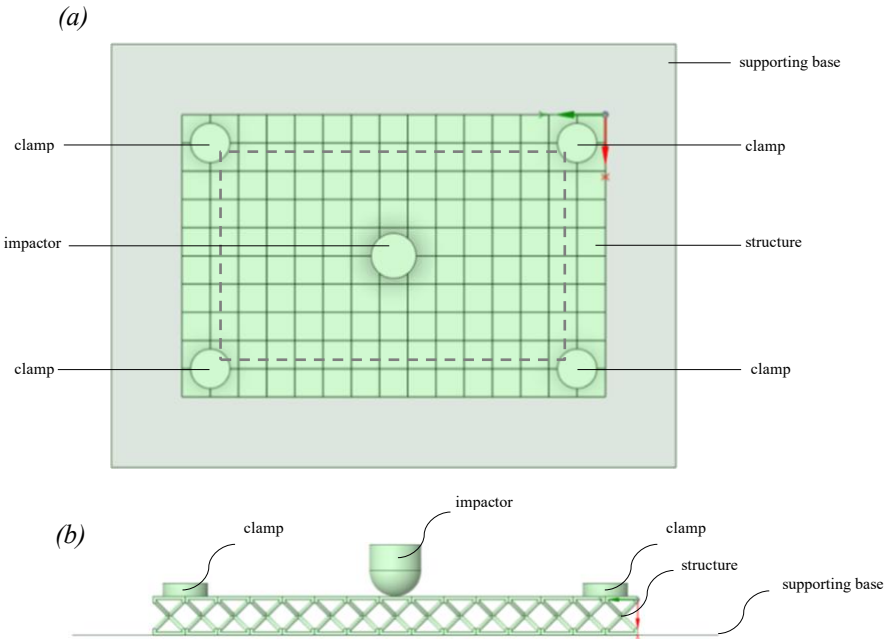


Figure 7. *(a)* Top view of the configuration *(b)* later view of the configuration

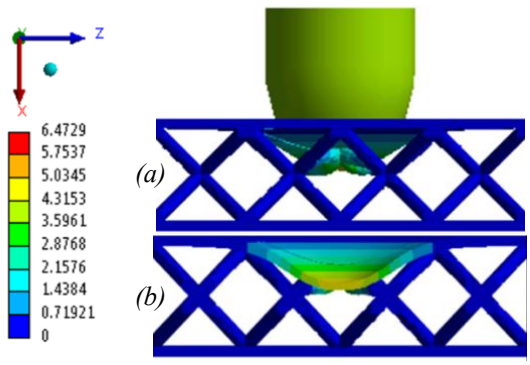


Figure 8. Total Displacement BCC1.5; lateral view *(a),* section view *(b).*

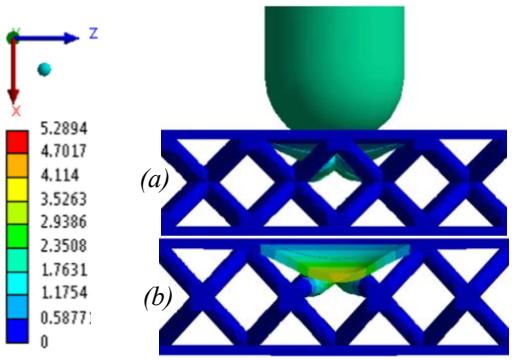


Figure 9. Total Displacement BCC2; lateral view *(a),* section view *(b).*

As observed in Fig. 10, BCC1.5 exhibits a slightly higher energy in the plateau region, reaching 2.20 J, compared to 2.19 J for BCC2. In addition, BCC2 demonstrates a greater maximum load of 1030.9 N, exceeding the 852.18 N recorded for BCC1.5. In terms of displacement, BCC1.5 experiences a higher value of 6.47 mm, whereas BCC2 exhibits a lower displacement of 5.28 mm.

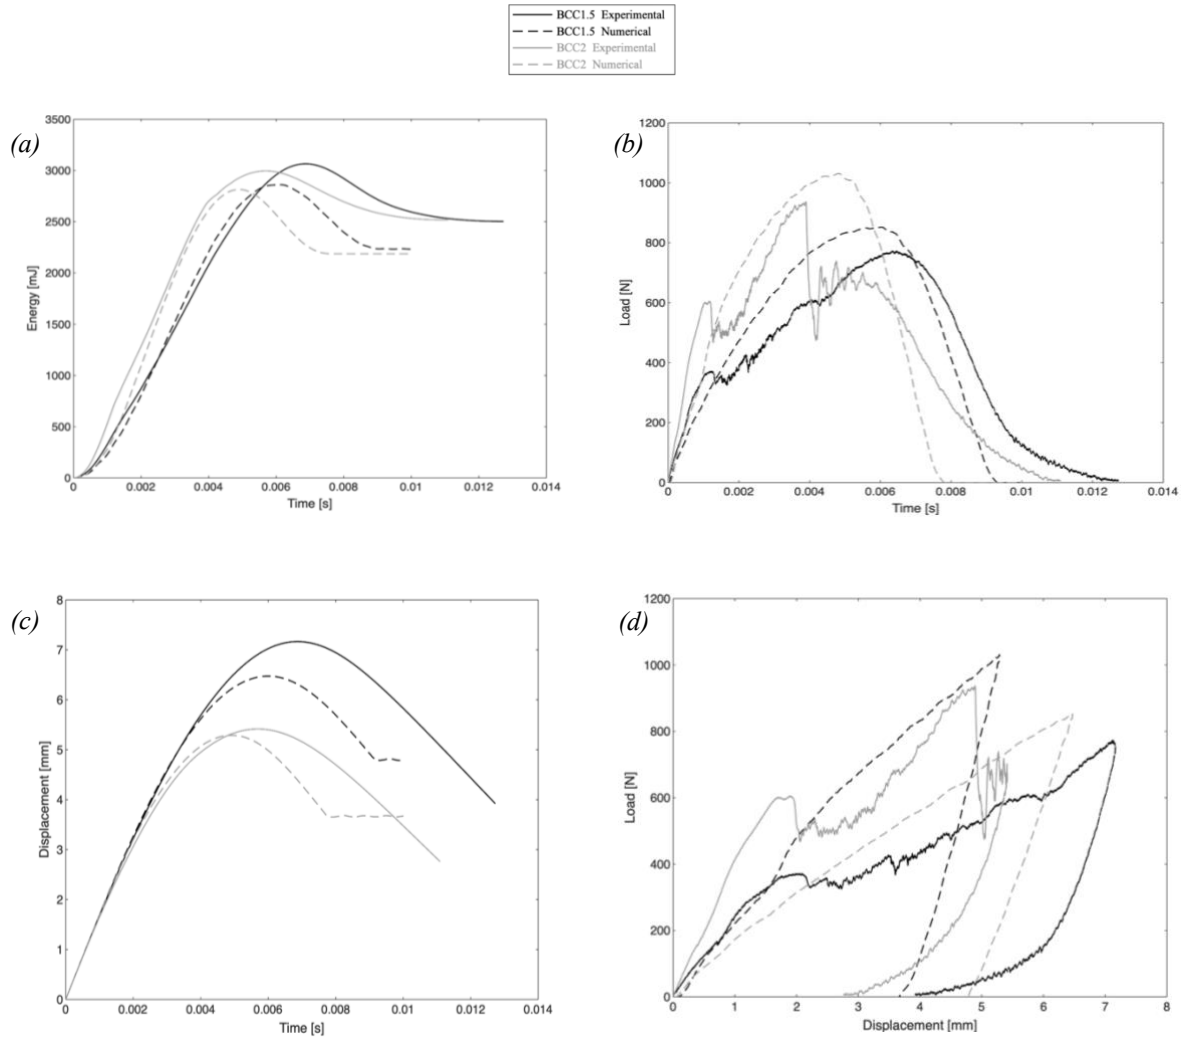


Figure 10. Numerical and Experimental (*a)* Time – Energy, (*b*) Time – Load, (*c*) Time – Displacement, (*d*) Displacement – Load.

4. DISCUSSION

This preliminary study analyzed two periodic lattice structures, BCC1.5 and BCC2, with strut diameters of 1.5 mm and 2 mm respectively, by integrating experimental drop-tower tests and numerical simulations in Ansys LS-DYNA. The load–displacement diagrams, both experimental and numerical, revealed a progressive collapse mechanism, ensuring steady energy dissipation during impact. Also, time–load diagrams showed an initial sharp peak corresponding to the impact load, followed by a gradual decline as the structure deformed. In

most cases, the BCC2 configuration exhibits lower percentage errors than BCC1.5 across all analyzed parameters, with the experimental values serving as the reference for error calculation.

For maximum load, BCC2 showed a slightly lower error (10.00%) than BCC1.5 (10.27%). In terms of maximum displacement, BCC2 outperformed BCC1.5 with an error of 2.40% versus 9.63%. For the plateau energy, which represents the energy dissipated by the specimen once the impact event is completed, BCC1.5 exhibited a lower error (12%) compared to BCC2 (12.4%).

The percentage error in the results obtained for the individual structures remains within acceptable limits. Above all, the numerical time-energy and time-displacement diagrams aligned reasonably well with experimental results. Furthermore, the time–load and load–displacement diagrams in the simulations appear smooth and linear. Conversely, the experimental time–load and load–displacement curves display pronounced oscillations and irregularities. These fluctuations can be attributed to physical factors, such as material imperfections, localized buckling, and dynamic phenomena like vibrations during testing. Moreover, as reported in Fig. 6(b), based on the visual inspection of post-impact damage, the BCC2 specimen exhibits a more pronounced hole compared to BCC1.5 in Fig. 6(a). In contrast, the damage in BCC1.5 appears to be more evenly distributed, with less concentrated rupture, indicating a better ability to spread the impact forces across the structure. This observation implies that BCC1.5 might perform slightly better in terms of mitigating localized damage under impact conditions, whereas BCC2 might show higher susceptibility to puncture or concentrated failure under the same conditions.

However, this work has certain limitations. From an experimental perspective, only one specimen for each BCC configuration was tested. The limited sample size restricts the ability to account for variability in manufacturing defects and material inconsistencies, which could influence the results. From a numerical standpoint, the absence of a material fracture model could represent a limitation, as it reflects an idealized computational environment.

5. CONCLUSIONS

This study confirms that BCC lattice structures exhibit a progressive collapse mechanism enabling steady energy dissipation. BCC2 generally shows lower percentage errors, while BCC1.5 ensures a more uniform damage distribution, highlighting different trade-offs based on impact requirements. Despite the limited sample size and the absence of a fracture model, the results affirm the reasonably good predictability of the numerical model, together with the applicability of these lattices for energy absorption purposes.

Future work, including larger test sets and refined simulations, will further optimize their design for aerospace, automotive, and protective systems.

REFERENCES

[1] W. Tao and M. C. Leu, “DESIGN OF LATTICE STRUCTURE FOR ADDITIVE MANUFACTURING,” in *Proceedings of ISFA2016*, 2016.

[2] M. I. Chibinyani, T. C. Dzogbewu, M. Maringa, and A. Muiruri, "Lattice Structures Built with Different Polygon Hollow Shapes: A Review on Their Analytical Modelling and Engineering Applications," *Applied Sciences*, vol. 14, no. 4, p. 1582, Feb. 2024, doi: 10.3390/app14041582.

[3] Y. Li, D. Jiang, R. Zhao, X. Wang, L. Wang, and L.-C. Zhang, "High Mechanical Performance of Lattice Structures Fabricated by Additive Manufacturing," *Metals (Basel)*, vol. 14, no. 10, p. 1165, Oct. 2024, doi: 10.3390/met14101165.

[4] G. Mantegna, "Mechanical properties of additively manufactured lattice structures through numerical BCC cell characterisation," Oct. 2023, pp. 177–184. doi: 10.21741/9781644902677-26.

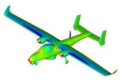
[5] A. Alaimo *et al.*, "Modified beam modeling of powder bed fusion manufactured lattice structures," *Int J Mech Sci*, vol. 259, p. 108599, Dec. 2023, doi: 10.1016/j.ijmecsci.2023.108599.

[6] V. Acanfora, R. Castaldo, and A. Riccio, "On the Effects of Core Microstructure on Energy Absorbing Capabilities of Sandwich Panels Intended for Additive Manufacturing," *Materials*, vol. 15, no. 4, p. 1291, Feb. 2022, doi: 10.3390/ma15041291.

[7] A. R. Bernard and M. S. A. ElSayed, "Crashworthiness of 3D Lattice Topologies under Dynamic Loading: A Comprehensive Study," *Materials*, vol. 17, no. 7, p. 1597, Mar. 2024, doi: 10.3390/ma17071597.

[8] G. Mantegna, "Numerical characterisation of additively manufactured components," Jul. 2024, pp. 188–192. doi: 10.21741/9781644903193-41.

.pdftk



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

EXTENSION OF THE TWIST-KIRCHHOFF ELEMENTS TO THE BUCKLING ANALYSIS OF PLATES

Hugo A.F.A. Santos[1*] and Thomas J.R. Hughes[2]

1: Departamento de Engenharia Mecânica, UnIRE, CIMOSM
Instituto Superior de Engenharia de Lisboa
Instituto Politécnico de Lisboa
Rua Conselheiro Emídio Navarro, 1, Lisboa, 1959-007, Portugal
hugo.santos@isel.pt

2: Computational Mechanics Group
Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
201 East 24th Street, 1 University Station C0200, Austin, TX 78712, USA
hughes@oden.utexas.edu

Abstract. *We extend the first two members of the twist-Kirchhoff family of arbitrary quadrilateral elements to the linear buckling analysis of plates. While standard polynomial C^0-continuous approximation functions are adopted for the transverse displacements, Raviart-Thomas vector-field approximations are considered for the rotations. Continuity of the normal components of the rotation vector across mesh edges is preserved by using the Piola transformation to map the rotations from the parent domain to the physical domain. The elements possess a unique combination of efficiency and robustness in that minimal quadrature rules are sufficient to guarantee stability without rank deficiency. We numerically study the convergence and accuracy of the elements on square, circular and skew plate buckling problems.*

Keywords: Thin plates, twist-Kirchhoff theory, buckling, quadrilateral finite elements, Raviart-Thomas vector fields

1 INTRODUCTION

Plate-like structures (thin-walled structures) are one of the most widely used structural elements in advanced engineering design, particularly in the aerospace industry. Some advantages of thin-walled structures are high strength coupled with the ease of manufacturing and the relative low weight.

Plates are often subjected to normal compressive and shearing loads acting in their middle plane (in-plane loads). When the stored compressive membrane strain energy reaches a critical level, such loadings can lead to buckling or elastic instability [1, 2]. This phenomenon is characterised by sudden and disproportionate large displacements that could lead to structural failure. Buckling of plates is, therefore, of great practical importance and plate buckling presents an integral part of the general analysis of a structure.

In order to efficiently study such structures, different plate theories have been proposed during the years. In classical plate theory (also known as Poisson-Kirchhoff) the shear deformation through the plate thickness is neglected under the assumption that the normal to the middle surface remains normal after deformation. In shear deformable plate theory (also referred to as Reissner-Mindlin), it is assumed that the normal-to-the-middle surface will not necessarily remain perpendicular after deformation, adding rotations of the normal as extra unknowns into the partial differential equations. While the former theory is valid for thin plate analysis, the latter can be used also for thick plates.

Particularly in aerospace applications, plates generally have ratios of L/t a number of factors greater than 20 and the effect of shear deformation on their elastic buckling strength is therefore relatively unimportant. Furthermore, the critical buckling load of plates is over-predicted using thin plate theory. The classical thin plate theory buckling predictions become increasingly more conservative as the thickness of the plate increases. From a structural design point of view, it is therefore reasonable to consider that thin plate buckling models are of utmost importance.

A new family of finite elements was proposed for the analysis of plate problems in [3, 4]. This new family of elements was developed on the framework of the so-called twist-Kirchhoff theory of plates, which conceptually lies in between the classical Reissner-Mindlin and Poisson-Kirchhoff theories. While standard polynomial C^0-continuous approximation functions are adopted for the transverse displacements, Raviart-Thomas vector-field approximations are considered for the rotations, such that continuity of the normal components across mesh edges are preserved. The elements possess a unique combination of efficiency and robustness in that minimal quadrature rules are sufficient to guarantee stability without rank deficiency.

It is the goal of this work to present the extension of the twist-Kirchhoff elements to the linear buckling analysis of thin plates. Several numerical examples involving uniform plane stress states are presented to illustrate the accuracy and convergence of the proposed finite elements on square, circular and skew plate buckling problems. The obtained results are compared, where possible, to exact solutions and to results presented by other authors. Although the numerical examples presented apply specifically to cases of uniform plane stresses, the finite element analysis can also be used to calculate the buckling loads of plates subjected to more complex in-plane loadings.

2 TWIST-KIRCHHOFF THEORY FOR LINEAR BUCKLING ANALYSIS OF THIN PLATES

Let $\Omega \subset \mathcal{R}^2$ denote the mid-plane of a plate undergoing rotation $\boldsymbol{\theta} = [\theta_x \ \theta_y]^T$ and transverse deflection w as a result of applied in-plane (membrane) loading N_x, N_y and N_{xy}. Further, let $\partial\Omega = \Gamma_N \cup \Gamma_D$ denote the boundary of Ω, with Γ_N and Γ_D the Neumann and Dirichlet parts, respectively, such that $\Gamma_N \cap \Gamma_D = \emptyset$. The plate is assumed to be linear elastic, homogeneous and isotropic with Young's modulus E, Poisson's ratio ν and constant thickness t.

In the classical Reissner-Mindlin formulation, the strain energy of the plate is given as follows

$$U(\boldsymbol{\theta}, w) = \int_\Omega \frac{1}{2}\left(\boldsymbol{\kappa}^T \boldsymbol{D}_M \boldsymbol{\kappa} + \boldsymbol{\gamma}^T \boldsymbol{D}_V \boldsymbol{\gamma}\right) \, d\Omega \tag{1}$$

with $\boldsymbol{\kappa}$ and $\boldsymbol{\gamma}$ the curvature tensor and shear strain vector of the plate given by

$$\boldsymbol{\kappa} = \begin{bmatrix} \kappa_{xx} \\ \kappa_{yy} \\ \kappa_{xy} \end{bmatrix} = \begin{bmatrix} \theta_{x,x} \\ \theta_{y,y} \\ \theta_{x,y} + \theta_{y,x} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\gamma} = \begin{bmatrix} \gamma_x \\ \gamma_y \end{bmatrix} = \begin{bmatrix} w_{,x} - \theta_x \\ w_{,y} - \theta_y \end{bmatrix} \tag{2}$$

and $\boldsymbol{D}_M$ and $\boldsymbol{D}_V$ the bending/twisting and shear constitutive matrices defined as

$$\boldsymbol{D}_M = D \begin{bmatrix} 1 & \nu & 0 \\ \nu & 1 & 0 \\ 0 & 0 & \frac{1-\nu}{2} \end{bmatrix} \quad \text{and} \quad \boldsymbol{D}_V = kGt \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{3}$$

where $D = Et^3/(12(1-\nu^2))$, $G = E/(2(1+\nu))$ and k is the shear correction factor, usually taken to be $5/6$.

The *twist-Kirchhoff theory* was proposed for thin plates in [3]. This theory lies in between the classical Reissner-Mindlin and Poisson-Kirchhoff theories. The main ingredient of this theory consists in replacing the Reissner-Mindlin twist component of curvature, defined in terms of first-order derivatives of the rotation components, with the classical Kirchhoff twist component of curvature, defined in terms of the second-order cross derivative of the transverse displacement, while retaining the definitions of the bending curvatures. In a rectangular Cartesian frame, this amounts to setting

$$\boldsymbol{\kappa} = \begin{bmatrix} \theta_{x,x} \\ \theta_{y,y} \\ 2w_{,xy} \end{bmatrix} \tag{4}$$

Based on this theory, a new family of rectangular displacement/rotation-based plate elements was proposed for the bending analysis of thin plates [3]. The generalization of this theory to the arbitrary quadrilateral element case was accomplished in [4].

In this work, we extend this theory and its associated elements (TK elements) to the buckling analysis of thin plates.

Within the framework of the linear buckling analysis, the total potential energy of a plate, $\Pi_p : \mathcal{U} \to \mathcal{R}$, corresponds to the strain energy given in Eq.(1) added to the potential energy of the in-plane loads as follows

$$\Pi_p(\boldsymbol{\theta}, w) = \int_\Omega \frac{1}{2}\left(\boldsymbol{\kappa}^T \boldsymbol{D}_M \boldsymbol{\kappa} + \boldsymbol{\gamma}^T \boldsymbol{D}_V \boldsymbol{\gamma}\right) \, d\Omega + \int_\Omega \frac{1}{2}\left(N_x w_{,x}^2 + N_y w_{,y}^2 + 2N_{xy} w_{,x} w_{,y}\right) \, d\Omega \tag{5}$$

with $\mathcal{U}$ the kinematically admissible space as

$$\mathcal{U} = \{(\boldsymbol{\theta}, w) \in \mathcal{H}^1(\Omega)^2 \times \mathcal{H}^1(\Omega) |\ \boldsymbol{\theta} = \mathbf{0} \text{ and } w = 0 \text{ on } \Gamma_D\} \tag{6}$$

Note that, here, the so-called linear classical buckling analysis is followed where pre-buckling deformations of the plate are either ignored or are zero.

Equilibrium of the loaded plate requires that its total potential energy given by Eq.(5) has a stationary value, *i.e.*, $\delta\Pi_p = 0$. This requirement gives the eigenvalue-eigenvector problem governing the twist-Kirchhoff plate buckling model. The TK elements for plate buckling that will be presented next are constructed on the basis of this variational statement.

3 FIRST- AND SECOND-ORDER TK FINITE ELEMENTS

Let us consider a non-degenerate mapping from a reference square $\widehat{K} = [-1,1] \times [-1,1]$ to a physical domain of arbitrary quadrilateral shape K, where $\phi : \widehat{K} \to K$ is a regular (orientation preserving) C^1 mapping, and (x, y) and (ξ, η) denote Cartesian and curvilinear coordinates, respectively.

Further, let $\widehat{\boldsymbol{\theta}}$ be a rotation vector field on $\widehat{K}$. The Piola transform of $\widehat{\boldsymbol{\theta}}$ is given by

$$\boldsymbol{\theta} = \frac{1}{J}\boldsymbol{F}\widehat{\boldsymbol{\theta}} \tag{7}$$

where

$$\boldsymbol{\theta} = \begin{bmatrix} \theta_x \\ \theta_y \end{bmatrix}, \quad \widehat{\boldsymbol{\theta}} = \begin{bmatrix} \widehat{\theta}_\xi \\ \widehat{\theta}_\eta \end{bmatrix} \tag{8}$$

$\boldsymbol{F}$ is the Jacobian matrix of ϕ defined by

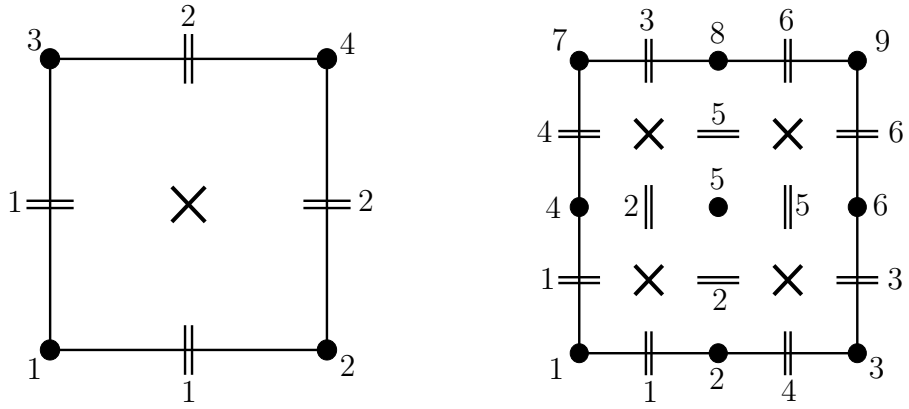
$$\boldsymbol{F} = \begin{bmatrix} x_{,\xi} & x_{,\eta} \\ y_{,\xi} & y_{,\eta} \end{bmatrix} \tag{9}$$

and $J = x_{,\xi}y_{,\eta} - x_{,\eta}y_{,\xi}$ is its determinant. The Piola transformation preserves continuity of the normal components of the rotation vector across mesh edges.

The first- and second-order elements, defined in the parent domain, are schematically illustrated in Fig.1.

In the parent domain, the first- and second-order elements are based on the following sets of approximations: the rotations are represented by $H(\mathrm{div})$-conforming Raviart-Thomas vector fields of order 0 and 1, whereas the transverse displacement is represented by standard C^0-continuous, piecewise bi-Lagrange functions of order 1 and 2, respectively, as follows

$$\begin{aligned}
\widehat{\theta}^h_\xi &= \sum_{i=1}^{2} \widehat{N}^{\theta_\xi}_i \theta^\xi_i & \widehat{\theta}^h_\xi &= \sum_{i=1}^{6} \widehat{N}^{\theta_\xi}_i \theta^\xi_i \\
\widehat{\theta}^h_\eta &= \sum_{i=1}^{2} \widehat{N}^{\theta_\eta}_i \theta^\eta_i \quad \text{and} & \widehat{\theta}^h_\eta &= \sum_{i=1}^{6} \widehat{N}^{\theta_\eta}_i \theta^\eta_i \\
\widehat{w}^h &= \sum_{i=1}^{4} \widehat{N}^{w}_i w_i & \widehat{w}^h &= \sum_{i=1}^{9} N^{w}_i w_i
\end{aligned} \tag{10}$$



(a) TK1 (b) TK2

Figure 1: First- and second-order TK elements

●	w^h degree-of-freedom
=	θ^h_ξ degree-of-freedom
‖	θ^h_η degree-of-freedom
×	Quadrature points

where θ_i^ξ and θ_i^η represent mid-side rotational degrees-of-freedom about the η- and ξ-constant coordinate lines, respectively, whereas w_i represent vertex displacement degrees-of-freedom, as illustrated in Fig.1. The first-order element (TK1) possesses, therefore, 8 degrees-of-freedom, whereas the second-order element (TK2) possesses 21 degrees-of-freedom. Also note that, in the physical domain, the rotational degrees-of-freedom become the normal components of the rotation on the element boundaries.

The shape functions for the rotations, $\widehat{N}_i^{\theta_\xi}$ and $\widehat{N}_i^{\theta_\eta}$, and transverse displacement N_i^w are given in [4].

Numerical stability requires the use of a reduced quadrature rule; see [3] for further details. For TK1 element we use a 1-point Gaussian quadrature rule, whereas for TK2 element we use a 2×2-point Gaussian quadrature rule.

The geometrical map of the elements utilizes the same basis functions as those for the transverse displacement.

4 LINEAR BUCKLING ANALYSIS

On insertion of the adopted approximation functions into the potential energy functional given in Eq.(5) gives

$$\Pi_p^h = \frac{1}{2}\boldsymbol{d}^T\boldsymbol{K}_M\boldsymbol{d} + \frac{\lambda}{2}\boldsymbol{d}^T\boldsymbol{K}_G\boldsymbol{d} \tag{11}$$

The effects of membrane forces are accounted for by matrix $\boldsymbol{K}_G$, which augments the global material stiffness matrix $\boldsymbol{K}_M$. Matrix $\boldsymbol{K}_G$ is called the global geometric stiffness matrix and it is defined by the elements' geometry, displacement fields, and states of

stress. $\boldsymbol{d}$ is the global vector of degrees-of-freedom. The first variation of Π_p^h with respect to $\boldsymbol{d}$ gives the following global system of equilibrium equations

$$(\boldsymbol{K}_M + \lambda \boldsymbol{K}_G)\boldsymbol{d} = \boldsymbol{0} \tag{12}$$

which defines a generalized eigenvalue problem. The resulting values for λ represent the buckling loads of the TK plate model and the associated vectors $\boldsymbol{d}$ are the corresponding buckling shape modes. The smallest solution for λ is usually referred to as the critical load parameter. The magnitude of $\boldsymbol{d}$ is indeterminate. Therefore, $\boldsymbol{d}$ identifies shape but not amplitude.

5 NUMERICAL RESULTS

For all computations, a shear correction factor of 5/6 is utilized. All the examples refer to plates with Young's modulus and Poisson's ratio taken as $E = 10^7$ and $\nu = 0.3$. The specification of essential boundary conditions is identical to that of Poisson-Kirchhoff theory.

The convergence characteristics and accuracy of the elements are demonstrated comparing the results to those obtained by other authors and, where possible, to exact solutions.

5.1 Simply-Supported Square Plate

We first consider the simply-supported thin square plate illustrated in Fig.2. We computed solutions for a thickness $t = 1/10000$ and length $L = 1$. The plate is first assumed to be under a uniform axial compressive load in the x-direction, N_x, and, afterwards, the plate is considered to be subjected to a uniform shear loading, N_{xy}. The adopted refinement scheme consists in uniform $n_{el} \times n_{el}$ square meshes.

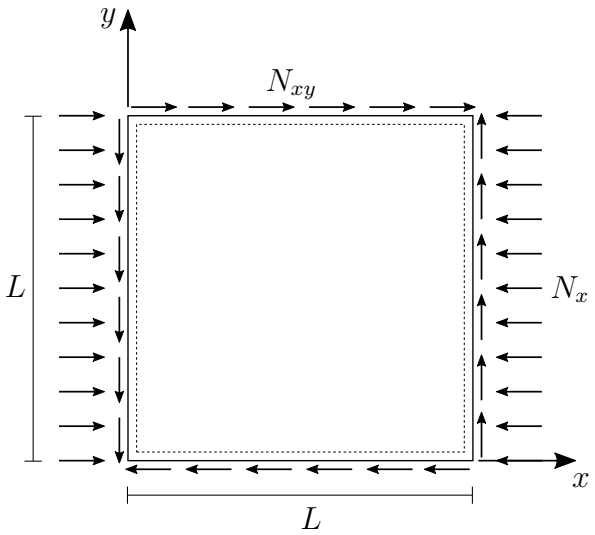


Figure 2: Simply-supported square plate model

5.1.1 Uniform and Uni-Axial Compressive Load N_x

We first analyse the plate under a uniform axial compressive load in the x-direction. The exact analytical buckling loads for the Poisson-Kirchhoff (thin or classical) plate

theory can be found in [5]. They are as follows

$$N_x^{mn} = \frac{\pi^2 D}{L^2}\left(m + \frac{n^2}{m}\right)^2 \tag{13}$$

with $m,\ n = 1,\ 2,\ 3, ...$

We study the convergence of the first and second buckling loads. Tab.1 displays the obtained results for TK1 and TK2 elements. The exact first and second buckling loads result from setting $m = n = 1$ and $n = 1,\ m = 2$, respectively, which, dividing by $\frac{\pi^2 D}{L^2}$, leads to 4 and 6.25, respectively. Comparing our results with the reference Poisson-

Table 1: First and second buckling loads ($N_x \times L^2/(\pi^2 D)$) for the simply-supported square plate model under uniform and uni-axial compressive load N_x. Exact thin plate solutions are 4 and 6.25, respectively.

	1st		2nd	
Mesh	TK1	TK2	TK1	TK2
2×2	6.48456	4.06022	NA	7.10538
4×4	4.45030	4.00408	8.90059	6.30905
8×8	4.10510	4.00026	6.73951	6.25399
16×16	4.02584	4.00002	6.36478	6.25026
32×32	4.00643	4.00000	6.27825	6.25002
64×64	4.00161	4.00000	6.25704	6.25000
Ref. Sol. [5]	4		6.25	

Kirchhoff solutions, we find the twist-Kirchhoff buckling loads of this model converge monotonically to the thin plate solutions from above.

5.1.2 Uniform Shear-Load N_{xy}

We now study the convergence of the first buckling load of the simply-supported square plate under uniform shear loading. Tab.2 displays the obtained results for TK1 and TK2 elements. As it can be observed, the obtained solutions are in close agreement with the

Table 2: First buckling load ($N_{xy} \times L^2/(\pi^2 D)$) for the simply-supported square plate model under uniform shear load N_{xy}. Reference thin plate solution is 9.34 [1].

Mesh	TK1	TK2
2×2	NA	11.15744
4×4	14.70559	9.46239
8×8	10.38424	9.33390
16×16	9.57476	9.32512
32×32	9.38621	9.32455
64×64	9.33988	9.32452
Ref. Sol. [1]	9.34	

reference solution.

5.2 Simply-Supported Circular Plate

Let us now consider an isotropic simply-supported circular plate of radius $R = 1$ and thickness $t = 1/10000$ under a uniform radial compressive load, N, such that $N_x = N_y = N$ and $N_{xy} = 0$, as illustrated in Fig.3. The exact solutions can be found in [5].

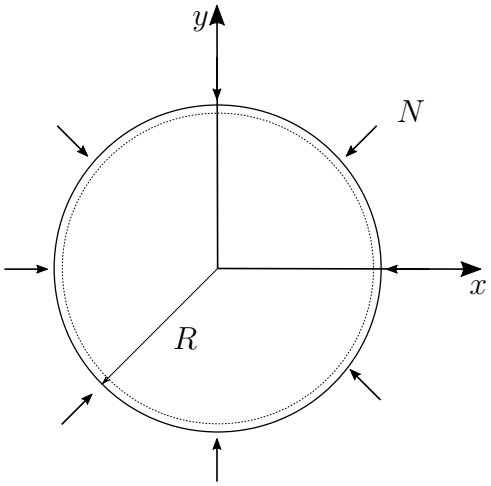
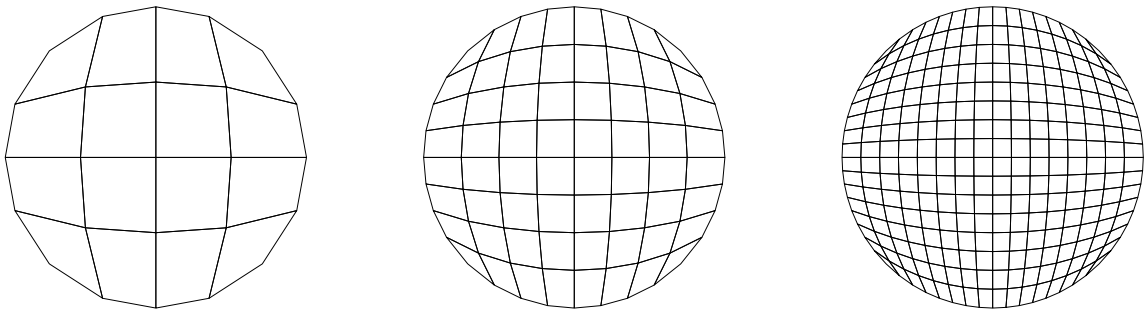


Figure 3: Simply-supported circular plate model

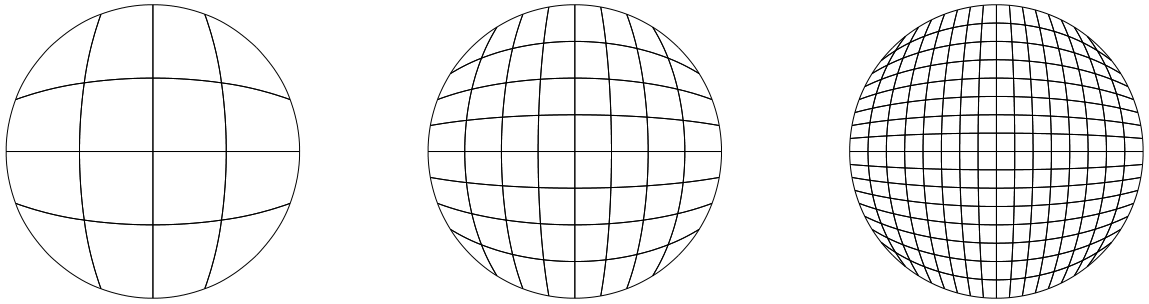
The adopted mesh refinement scheme is based on a sequence of meshes obtained by mapping a square domain onto a circle as illustrated in Fig.4 and Fig.5 for TK1 and TK2 elements, respectively.



(a) $n_{el} = 2$ (b) $n_{el} = 4$ (c) $n_{el} = 8$

Figure 4: Meshes with bi-linear mapping for circular plate problem

We study the convergence of the first and second buckling loads. Tab.3 displays the obtained results for TK1 and TK2 elements. The exact solutions were computed using the analytical expressions provided in [5] for the axisymmetric (first) and asymmetric (second) shape modes of Poisson-Kirchhoff plates. As previously, comparing our results with the exact Poisson-Kirchhoff solutions, we find the twist-Kirchhoff buckling loads of this model converge monotonically to the thin plate solutions from above.



(a) $n_{el} = 2$ (b) $n_{el} = 4$ (c) $n_{el} = 8$

Figure 5: Meshes with bi-quadratic mapping for circular plate problem

Table 3: First and second buckling loads ($N \times (2R)^2/(\pi^2 D)$) for the simply-supported circular plate model under a uniform radial compressive load N. Exact thin plate solutions are 1.70130 and 5.32467, respectively [5].

	1st		2nd	
Mesh	TK1	TK2	TK1	TK2
2×2	4.22238	1.75262	NA	6.57251
4×4	2.06195	1.70711	8.92227	5.41171
8×8	1.78321	1.70168	5.96887	5.33067
16×16	1.72132	1.70132	5.47463	5.32505
32×32	1.70628	1.70130	5.36152	5.32469
64×64	1.70254	1.70130	5.33384	5.32467
Ref. Sol. [5]	1.70130		5.32467	

5.3 Simply-Supported Skew Plate

Let us now consider an isotropic simply-supported skew plate under a uniform axial compressive load in the x-direction as illustrated in Fig.6. We computed solutions for thicknesses $t = 1/100$ and $t = 5/100$, and length $L = 1$. As is well known, numerical techniques used to solve skew plate problems become more onerous and require numerical models with larger numbers of degrees-of-freedom as the angle of skew of the plate increases [6]. This is due to the bending moment singularities that occur at the obtuse vertices of skew plate geometries [7].

The adopted refinement scheme consists in uniform $n_{el} \times n_{el}$ parallelogram-shape-regular meshes as shown in Fig.7 for the first three meshes.

We study the convergence of the first buckling load. Tab.4 displays the obtained results for TK1 and TK2 elements. There are no known analytical solutions available for this problem. However, for the thickness-to-length ratio $t/L = 0.05$, we found numerical solutions, given in [8], that we will use as reference solutions. These solutions were computed using the Rayleigh-Ritz method applied to a Reissner-Mindlin (thick) plate model, being 5.7784 and 5.4182 for hard and soft simply-supported boundary conditions, respectively. As it can be observed, the obtained solutions converge monotonically. It can

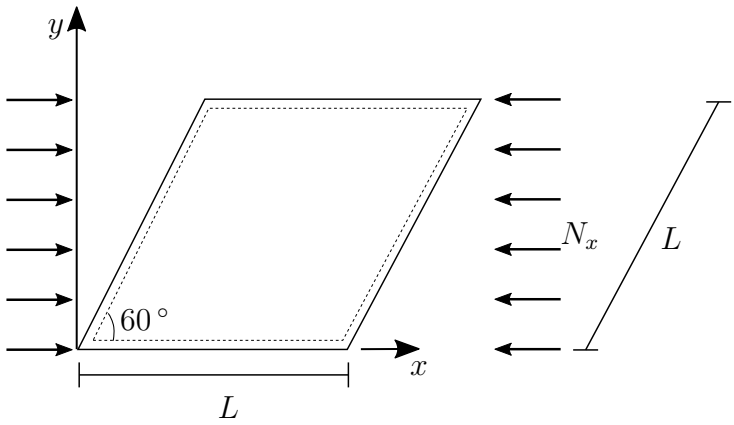
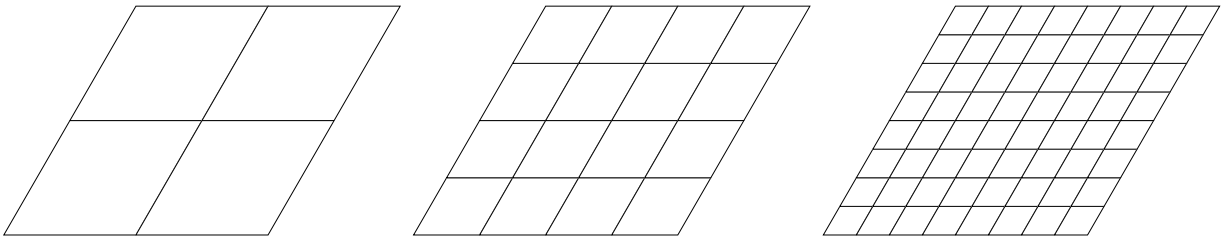


Figure 6: Simply-supported skew plate model



(a) $n_{el} = 2$ (b) $n_{el} = 4$ (c) $n_{el} = 8$

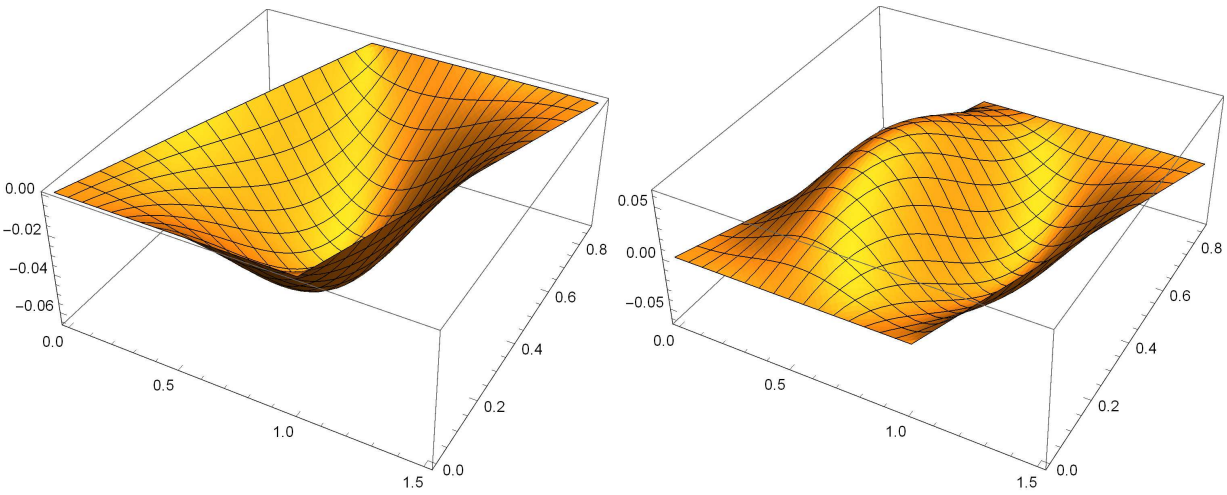
Figure 7: Parallelogram meshes for the skew plate problem

Table 4: First buckling loads ($N_x \times L^2/(\pi^2 D)$) for the simply-supported skew plate model under uniform and uni-axial compressive load N_x.

	$t/L = 0.01$		$t/L = 0.05$	
Mesh	TK1	TK2	TK1	TK2
2×2	10.49466	6.01022	10.31540	5.94965
4×4	6.82290	5.85708	6.74159	5.77513
8×8	6.09762	5.84859	6.00495	5.75765
16×16	5.91130	5.84760	5.81817	5.75646
32×32	5.86345	5.84752	5.77181	5.75639
64×64	5.85149	5.84751	5.76024	5.75638
Ref. Sol. [8]			5.7784	(Hard SS)
Ref. Sol. [8]			5.4182	(Soft SS)

also be seen that, the converged solution for $t/L = 0.05$ (obtained on the 64×64 mesh with TK2 elements) is in close agreement with the reference Reissner-Mindlin solutions, lying in between hard and soft simply-supported solutions.

The first and second buckling shape modes obtained with TK2 on the 16×16 element mesh are shown in Fig.8.



(a) First buckling mode

(b) Second buckling mode

Figure 8: First and second buckling shape modes obtained with TK2 on the 16×16 element mesh for the skew plate problem

6 CONCLUSIONS

We have extended the first two members of the twist-Kirchhoff family of elements to the buckling analysis of plates. Problems studied include linearly elastic square and circular simply-supported plates, and a simply-supported skew plate problem. We computed the first few buckling loads and buckling shape modes of the plates using the first- and second-order TK elements. The obtained solutions were compared, where possible, to exact solutions and to results presented by other authors. Good overall convergence behaviour was realized.

REFERENCES

[1] S. P. Timoshenko and J. M. Gere. *Theory of elastic stability*. McGraw-Hill, 1961.

[2] E. Ventsel and T. Krauthammer. *Thin Plates and Shells: Theory: Analysis, and Applications*. CRC Press, 2001. doi: 10.1201/9780203908723.

[3] F. Brezzi, J. A. Evans, T. J. R. Hughes, and L. D. Marini. New rectangular plate elements based on twist-Kirchhoff theory. *Computer Methods in Applied Mechanics and Engineering*, 200(33-36):2547–2561, 2011.

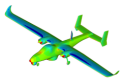
[4] H. A. F. A. Santos, J. A. Evans, and T. J. R. Hughes. Generalization of the twist-Kirchhoff theory of plate elements to arbitrary quadrilaterals and assessment of convergence. *Computer Methods in Applied Mechanics and Engineering*, 209:101–114, 2012.

[5] C. M. Wang, C. Y.Wang, and J. N. Reddy. *Exact Solutions for Buckling of Structural Members*. CRC Press, 2005.

[6] T. Mizusawa and T. Kajita. Vibration and buckling of skew plates with edges elastically restrained against rotation. *Computers & structures*, 22(6):987–994, 1986.

[7] A. W. Leissa. Singularity considerations in membrane, plate and shell behaviors. *International journal of solids and structures*, 38(19):3341–3353, 2001.

[8] S. Kitipornchai, Y. Xiang, C. M. Wang, and K. M. Liew. Buckling of thick skew plates. *International journal for numerical methods in engineering*, 36(8):1299–1310, 1993.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

INVERSE ESTIMATION METHOD FOR THE CHARACTERISATION OF ELASTIC PROPERTIES OF ANISOTROPIC OPEN-CELL MATERIALS

Huina Mao[1*], Xuefeng Li[2], Gunnar Tibert[1] and Romain Rumpler[2,3]

1: Aerospace Engineering
Dpt of Engineering Mechanics
KTH Royal Institute of Technology
SE-100 44, Stockholm, Sweden
{huina,tibert}@kth.se, `www.kth.se`

2: The Marcus Wallenberg Laboratory for Sound and Vibration Research (MWL)
Dpt of Engineering Mechanics
KTH Royal Institute of Technology
SE-100 44, Stockholm, Sweden
{xuefengl,rumpler}@kth.se, http://www.kth.se - MWL

3: Digital Futures and The Centre for ECO2 Vehicle Design,
KTH Royal Institute of Technology
SE-100 44, Stockholm, Sweden
{rumpler}@kth.se, www.digitalfutures.kth.se, www.eco2vehicledesign.kth.se

Abstract. *Open-cell architected materials exhibit highly tunable elastic properties governed by both the geometry of their cellular architecture and the intrinsic behaviour of their solid frame. This work presents a simulation-based framework for the characterisation and design of anisotropic stiffness in such materials. A parametrised geometric model based on modified Kelvin-Cell structures is introduced to explore the impact of twist, tilt, and stretch transformations on effective elastic behaviour. Directional moduli and stiffness surfaces are computed using beam-based finite element models under periodic boundary conditions, revealing systematic trends in anisotropy. Complementing this, an inverse estimation method is employed to identify the anisotropic Hooke's tensor of the solid phase from macroscopic mechanical responses. The method formulates a constrained optimisation problem linking simulated displacement fields to homogenised stiffness responses, enabling identification of the solid's orthotropic elastic parameters. Comparison between geometric predictions and inversely estimated tensors highlights the central role of microstructural distortion in shaping anisotropic material properties. Together, the results demonstrate a robust and generalisable approach for the analysis and design of architected cellular materials with directional stiffness characteristics, relevant to aerospace structures and acoustic metamaterials.*

Keywords: Anisotropic cellular materials, inverse characterisation, stiffness tensor, Kelvin structure, homogenisation, lattice modelling

1 INTRODUCTION

Open-cell cellular materials are widely used in structural and acoustic engineering due to their low density, mechanical resilience, and tunable functional properties [1, 2]. Their effective elastic behaviour results from a complex interplay between the properties of the solid frame material and the geometric architecture of the cellular network [3, 4]. In engineering applications, particularly those involving lightweight structural or vibroacoustic components, the ability to design, control, and characterise anisotropic stiffness is of critical importance [5].

In [6], an inverse characterisation method was introduced to estimate the full static elastic stiffness tensor of the solid frame of anisotropic open-cell materials. The method relies on the formulation of a constrained optimisation problem linking numerically simulated responses to homogenised stiffness tensors obtained under periodic boundary conditions. This approach provides a non-destructive, simulation-based pathway for reconstructing the underlying Hooke's tensor from a limited set of mechanical responses, and extends previous work on inverse identification techniques applied to anisotropic porous solids [5, 7].

A complementary investigation into the geometric origin of anisotropy was presented in [8], where a parametric family of distorted Kelvin-type unit cells was constructed through affine and non-affine transformations, including controlled twist, tilt, and stretch. This approach builds on earlier studies of idealised foam geometries [9, 10], and enables systematic exploration of the relationship between geometric distortion and induced elastic anisotropy [11].

The present work brings these developments into a common framework, with an emphasis on the inverse estimation of stiffness tensors in architectured materials exhibiting geometric anisotropy. The methodological coupling of simulated mechanical loading responses with a calibrated inverse identification procedure is particularly relevant for design tasks where anisotropic properties must be both predicted and targeted. The framework exemplifies a route towards efficient numerical characterisation and model-based design of cellular materials with directional stiffness, relevant to applications in aerospace structures and acoustic metamaterials [12–14].

2 LATTICE-BASED MODELLING OF ANISOTROPIC CELLULAR MATERIALS

The geometrical and mechanical complexity of open-cell foams can be effectively captured using periodic lattice-based unit cell models [2, 10, 15]. In this work, a modified Kelvin cell architecture is used to investigate the elastic anisotropy resulting from controlled geometrical transformations of the base cell. The framework builds upon the classical tetrakaidecahedral topology introduced by Lord Kelvin [16] and later generalised for cellular materials [9, 17]. The original cell consists of 36 ligaments and 24 nodes, forming a unit cell that is topologically periodic in all directions.

Each ligament may be modelled as a Timoshenko beam with uniform cross-section and isotropic solid material properties. Anisotropy at the cellular level is thus purely geometric, induced by systematic transformations applied to the isometric reference cell. This approach enables the systematic exploration of structure–property relationships, relevant for the design of lightweight mechanical and acoustic metamaterials [8, 11, 18].

Geometric parameterisation and deformation modes

Three transformation modes are defined to control the cell geometry:

- **Twist** — a rotation of the top face of the cell relative to the bottom around the vertical axis, or, in a similar way, of the side and opposite faces;
- **Tilt** — a rotation of the boundary square faces around one of the principal, in-plane axes of these faces;
- **Stretch** — an elongation or compression of the cell along one or several of the three direction linking opposite square faces, preserving ligament connectivity.

Each mode can be applied independently or in combination, enabling a wide range of anisotropic cell shapes. To define the rotations, local coordinate systems associated with each of the cell square faces, denoted $\hat{x}\hat{y}\hat{z}_j$ where subscript j refers to the cell face as explained below, are defined to be parallel to the reference geometry unit cell global coordinate system xyz, with its origin defined at the center of the face, Fig. 1a. The transformation between the local and the global coordinate systems is defined in terms of Euler rotation angles $\theta^{\hat{z}}$, $\theta^{\hat{y}}$, and $\theta^{\hat{x}}$ with successive, in order rotations around the axes $\hat{z}$, $\hat{y}$, and $\hat{x}$. To identify a particular angle with a certain face of the cell, the following notations are used. For each angle, θ^i_j, the cell face is identified with subscript j, where $j = x_-, x_+, y_-, y_+, z_-, z_+$. The rotation axis is identified with a superscript i, such that $i = \hat{x}, \hat{y}, \hat{z}$. The direction cosine matrix in terms of the $(\hat{z}\hat{y}\hat{x})$ Euler rotations is given by

$$[\hat{T}]_j = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta^{\hat{x}}_j) & -\sin(\theta^{\hat{x}}_j) \\ 0 & \sin(\theta^{\hat{x}}_j) & \cos(\theta^{\hat{x}}_j) \end{bmatrix} \begin{bmatrix} \cos(\theta^{\hat{y}}_j) & 0 & \sin(\theta^{\hat{y}}_j) \\ 0 & 1 & 0 \\ -\sin(\theta^{\hat{y}}_j) & 0 & \cos(\theta^{\hat{y}}_j) \end{bmatrix} \begin{bmatrix} \cos(\theta^{\hat{z}}_j) & -\sin(\theta^{\hat{z}}_j) & 0 \\ \sin(\theta^{\hat{z}}_j) & \cos(\theta^{\hat{z}}_j) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{1}$$

The point $[\hat{x}, \hat{y}, \hat{z}]$ of the reference KC face will be $[\hat{x}', \hat{y}', \hat{z}']$ after the rotational transformation, and given as

$$\begin{bmatrix} \hat{x}' & \hat{y}' & \hat{z}' \end{bmatrix}^T_j = [\hat{T}]_j \begin{bmatrix} \hat{x} & \hat{y} & \hat{z} \end{bmatrix}^T_j . \tag{2}$$

Table 1 shows the distorted Kelvin cells used in the numerical cases. To identify each of the KC geometries, and their corresponding equivalent elastic Hooke's tensors, they will be referred to as; pure **T**wist: $\mathbf{NTw}_1$ and $\mathbf{NTw}_2$; pure **T**ilt: $\mathbf{NTi}_1$ and $\mathbf{NTi}_2$; and combined **T**wist and **T**ilt: $\mathbf{NTT}_1$ and $\mathbf{NTT}_2$.

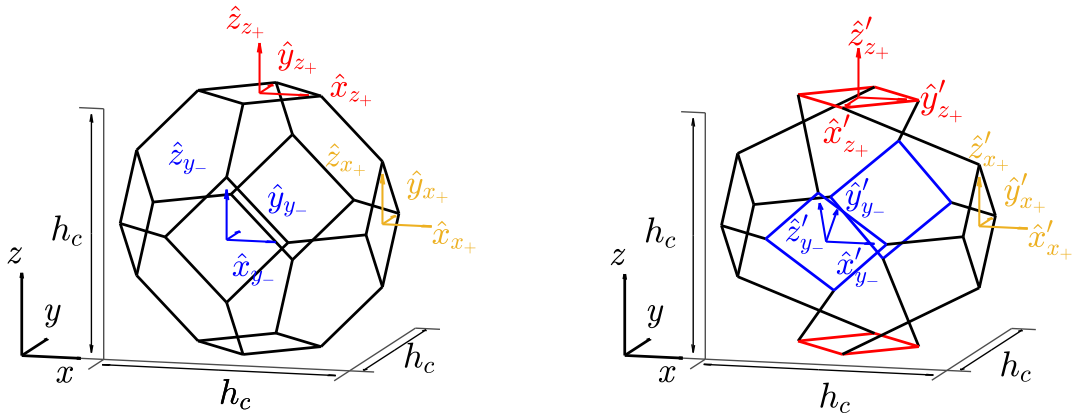
Similar deformation modes have been considered in previous parametric studies of Kelvin-like structures [19, 20]. The resulting microstructures remain periodic (with the exception of tilt transformation which require the definition of a $2 \times 2 \times 2$ supercell), which allows for homogenisation of their effective elastic properties under imposed strain fields. Fig. 1 shows the geometry of the isometric Kelvin Cell in (a), and an example of transformed unit cell in (b).

Computation of effective elastic properties

For each configuration considered among the vast range of possibilities, the homogenised stiffness tensor is computed using strain-driven simulations on the unit cell under periodic boundary conditions [8]. The cell is subjected to six independent macroscopic strain

Table 1: Distorted Kelvin cell geometries due to rotations according to Eq. (1) and (2) for each square cell face indicated and where ϕ corresponds to a twist angle and ψ corresponds to a tilt angle.

Angles	$[\theta_j^{\hat{x}}, \theta_j^{\hat{y}}, \theta_j^{\hat{z}}]$	face x_+ $j = x_+$	face x_- $j = x_-$	face y_+ $j = y_+$	face y_- $j = y_-$	face z_+ $j = z_+$	face z_- $j = z_-$
Twist	**NTw**$_1$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, \phi]$	$[0, 0, \phi]$
	NTw$_2$	$[\phi, 0, 0]$	$[\phi, 0, 0]$	$[0, \phi, 0]$	$[0, \phi, 0]$	$[0, 0, \phi]$	$[0, 0, \phi]$
Tilt	**NTi**$_1$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, \psi]$	$[0, 0, \psi]$	$[0, 0, 0]$	$[0, 0, 0]$
	NTi$_2$	$[0, 0, \psi]$	$[0, 0, \psi]$	$[\psi, 0, 0]$	$[\psi, 0, 0]$	$[0, \psi, 0]$	$[0, \psi, 0]$
Twist & Tilt	**NTT**$_1$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, \psi]$	$[0, 0, \psi]$	$[0, 0, \phi]$	$[0, 0, \phi]$
	NTT$_2$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, 0]$	$[\psi, 0, \phi]$	$[\psi, 0, \phi]$



(a) (b)

Figure 1: (a) Reference isometric single KC, in unit cell Cartesian coordinate system xyz, and local square faces coordinate systems $\hat{x}\hat{y}\hat{z}$ with subscripts identifying to which of the square faces it refers; and (b) illustration of a transformed KC geometry with +90 degree twists around $\hat{z}'_{z_+}$ of the top and bottom square faces (in red), and tilt of the side square faces around $\hat{x}'_{y_-}$ (in blue).

fields, and the corresponding volume-averaged stresses are computed. This yields the full 6×6 stiffness matrix $\tilde{\mathbf{H}}$ under the assumption of linear elasticity.

The geometric transformations induce a wide range of orthotropic and fully anisotropic stiffness tensors. For example, twist introduces shear-extension coupling, tilt alters symmetry in the off-diagonal terms, and stretch predominantly affects stiffness ratios between axial and transverse directions. Similar findings were reported for anisotropic elastic response in transformed Kelvin structures [11, 21].

Insights on controllability and design

The parameterised Kelvin cell framework enables a direct link between microstructure geometry and macroscopic elastic response [8]. One notable outcome is the ability to achieve a broad range of stiffness anisotropy while maintaining the same base material and

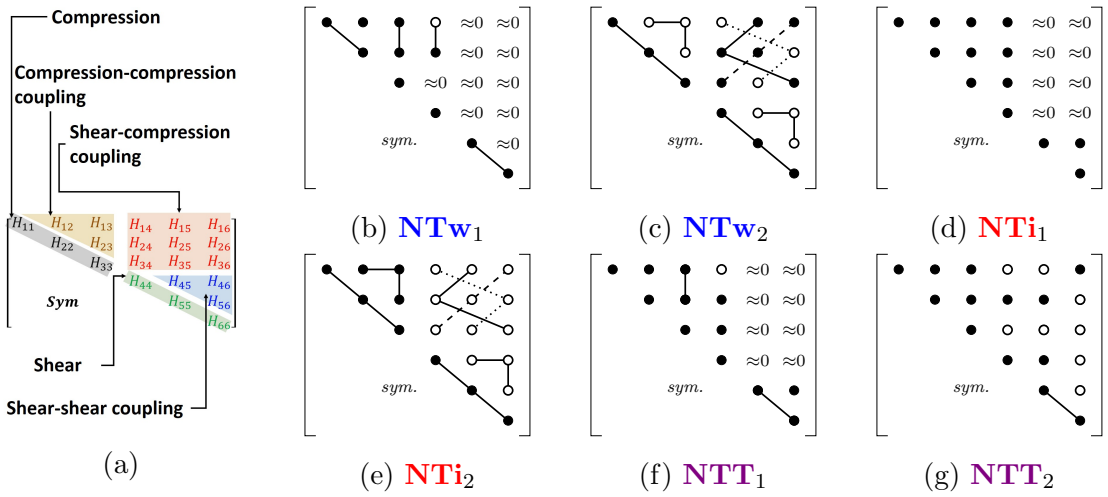


Figure 2: Symbolic representation of estimated equivalent Hooke's tensors, using a representation from [24]. ●: positive value, ○: negative value, tensor elements connected with lines have the same amplitude.

volume fraction. This opens pathways for the rational design of acoustic and mechanical metamaterials where directional stiffness control is required [18, 22].

The results also show that extreme values of twist and tilt can break symmetry conditions and lead to significant deviations from orthotropic behaviour. However, geometric limits exist beyond which ligament overlap or self-contact occurs, thereby constraining the design space.

The use of beam-based lattice models provides a computationally efficient means of exploring these trends [2, 10]. Similar design strategies and parameterisation frameworks have been developed for curved and composite lattice structures [13, 23], demonstrating the potential of wave-based identification techniques for mechanical characterisation of architected media with anisotropy. Nevertheless, the fidelity of such models depends on assumptions regarding ligament behaviour, particularly in the context of shear deformation and joint compliance.

Figure 2 shows, in symbolic representation of the elasticity tensor, the resulting symmetries as well as the existing couplings associated with diverse transformations considered, for the estimated equivalent elastic Hooke's tensor. All six configurations introduced in Table 1 are included.

3 INVERSE CHARACTERISATION OF ANISOTROPIC ELASTIC PROPERTIES

The inverse characterisation method used is the one developed in [6], targetting the estimation of the effective elastic stiffness tensor $\tilde{\mathbf{H}}$ of the solid phase of open-cell porous materials, as presented in the previous sections. This approach is motivated by the difficulty of directly measuring the intrinsic properties of the solid frame in such micro-architectured systems, where the material phase is challenging to access experimentally and its geometry complex to model explicitly. Experimental efforts in this domain have included modal analysis [25, 26], digital image correlation [27], and resonance testing [5], yet these techniques often lack the ability to identify the full anisotropic stiffness tensor.

The methodology builds on previous inverse estimation approaches [5, 7] and combines

finite element simulations with an optimisation scheme to match simulated responses with experimental data from the porous material. A notable feature of the method is its independence from full-field strain measurements or detailed knowledge of the local geometry. An overview of the inverse estimation process is shown in Fig. 3. Further

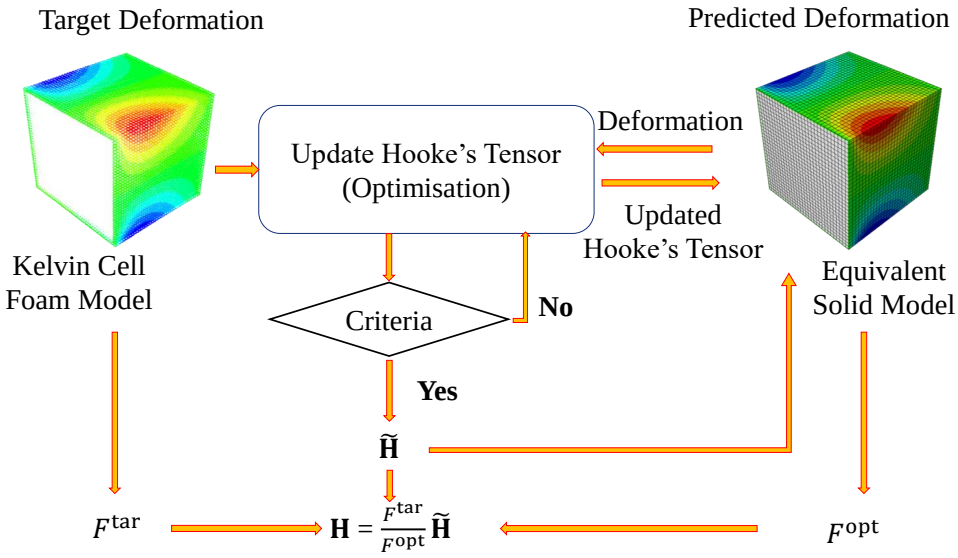


Figure 3: Inverse estimation flowchart: minimisation of the displacement difference on free faces (color contour plot) between the target and predicted models in the final inverse estimation.

details may be found in [6], for the interested reader.

3.1 Parameterisation of anisotropic elasticity

In its general form, the solid phase of an open-cell material may exhibit orthotropic or transversely isotropic behaviour [4, 28]. To account for this, the effective Hooke's tensor is parameterised according to the implementation proposed in [6], where the elasticity tensor is parameterised using 20 dimensionless scaling parameters ξ_i, applied to selected components of the symmetric Hooke's matrix $\tilde{\mathbf{H}}$, the H_{33}-normalized version of $\mathbf{H}$,

$$\mathbf{H} = H_{33}\begin{bmatrix} \tilde{H}_{11} & \tilde{H}_{12} & \tilde{H}_{13} & \tilde{H}_{14} & \tilde{H}_{15} & \tilde{H}_{16} \\ & \tilde{H}_{22} & \tilde{H}_{23} & \tilde{H}_{24} & \tilde{H}_{25} & \tilde{H}_{26} \\ & & 1 & \tilde{H}_{34} & \tilde{H}_{35} & \tilde{H}_{36} \\ & & & \tilde{H}_{44} & \tilde{H}_{45} & \tilde{H}_{46} \\ & & sym & & \tilde{H}_{55} & \tilde{H}_{56} \\ & & & & & \tilde{H}_{66} \end{bmatrix}. \tag{3}$$

The first eight parameters correspond to the principal and shear moduli:

$$\xi_{1-8} = [\tilde{H}_{11}, \tilde{H}_{22}, \tilde{H}_{12}, \tilde{H}_{13}, \tilde{H}_{23}, \tilde{H}_{44}, \tilde{H}_{55}, \tilde{H}_{66}]$$

while the remaining twelve parameters account for coupling terms:

$$\xi_{9-20} = [\tilde{H}_{14}, \tilde{H}_{24}, \tilde{H}_{34}, \tilde{H}_{15}, \tilde{H}_{25}, \tilde{H}_{35}, \tilde{H}_{45}, \tilde{H}_{16}, \tilde{H}_{26}, \tilde{H}_{36}, \tilde{H}_{46}, \tilde{H}_{56}]$$

This formulation allows the inverse estimation procedure to explore a broad space of admissible stiffness tensors while imposing compliance with symmetry and stability constraints. Since the boundary conditions in the forward simulations prescribe displacements, the stiffness tensor is identified only up to a scalar factor. To recover its absolute scale, a global scaling constant C_{H} is computed based on the reaction forces in a reference load case, with $\tilde{H}_{33}$ subsequently normalised to 1 MPa, see Eq. (3).

A key modelling choice concerns the imposition of orthotropic symmetry, which was validated in [6] by a symmetry classification procedure based on spectral decomposition of the estimated tensor [29]. Attention must also be paid to boundary artefacts, as linear deformations near the edges of cellular specimens may locally violate assumptions of macroscopic homogeneity [30].

3.2 Forward simulations and synthetic response data

The inverse estimation framework is applied to synthetic datasets generated through forward simulations on architectured porous solids. These simulations rely on a homogenised continuum representation of the porous material, for which a series of boundary value problems are solved using finite element analysis. The material response is computed under six independent macroscopic loading conditions—three axial and three shear strains—providing displacement fields on the specimen surface.

The synthetic displacement fields serve as input data for the inverse procedure, mimicking virtual experiments while avoiding experimental uncertainties. This approach enables a fully controlled investigation of the identifiability of elastic parameters, and facilitates quantitative comparison with known reference tensors derived from beam-based lattice models.

The boundary conditions in the forward simulations prescribe displacements, resulting in stiffness identification up to a scalar factor. This scalar is retrieved by calibrating the global scaling of the elasticity tensor against reaction forces in a reference simulation, after normalising the estimated $\tilde{H}_{33}$ to 1 MPa.

3.3 Optimisation procedure and convergence

3.3.1 Convergence for orthotropic targets

A first validation is performed for a known orthotropic target material, whose effective elasticity tensor is obtained from a 30×30×30 Kelvin cell model. Synthetic displacement data are generated under six loading cases, and the inverse estimation is applied using equivalent solid models of varying resolution: from coarse (5^3) to fine (30^3) meshes. Figure 4 shows the evolution of selected scaling parameters ξ_i over iterations, comparing three optimisation runs using different initial guesses and mesh refinement strategies. These results demonstrate that a multi-level approach, starting from a coarse model and refining progressively, may yield faster and more stable convergence.

3.3.2 Convergence for fully anisotropic targets

A second example targets a fully anisotropic material with strong coupling between shear and normal responses. In this case, the reference displacement fields are generated from a solid finite element model with no underlying lattice structure. The inverse estimation is initiated using an orthotropic guess, and the multi-level procedure is applied

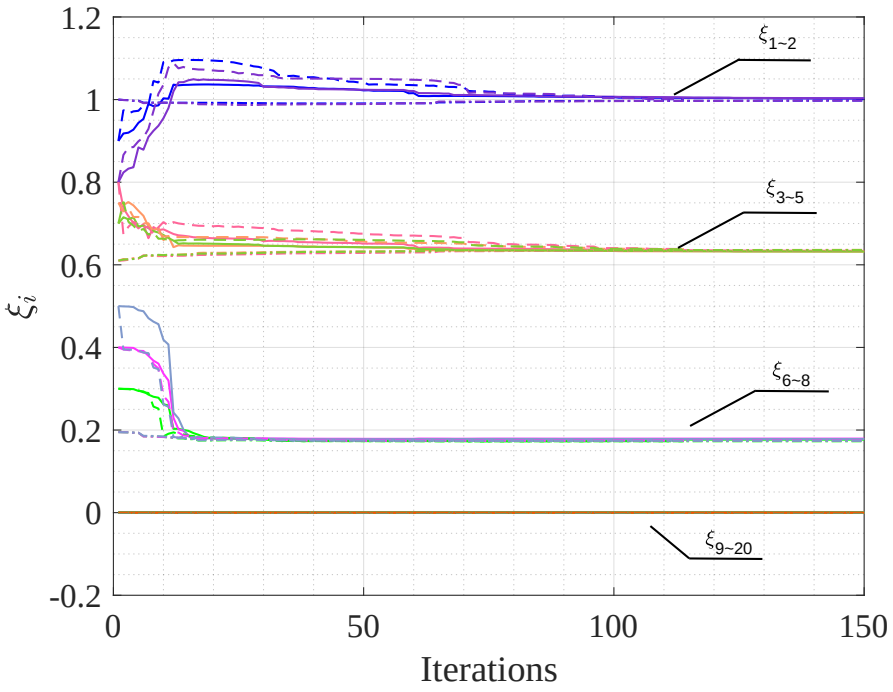


Figure 4: Evolution of the 20 scaling parameters ξ_i over the course of the optimisation. Dash-dotted, dashed, and solid lines correspond to Cases $\text{Case}_{30,5}^{\text{ave, }1}$, $\text{Case}_{30,5}^{\text{ave, }2}$, and $\text{Case}_{30,20}^{\text{ave, }3}$, respectively.

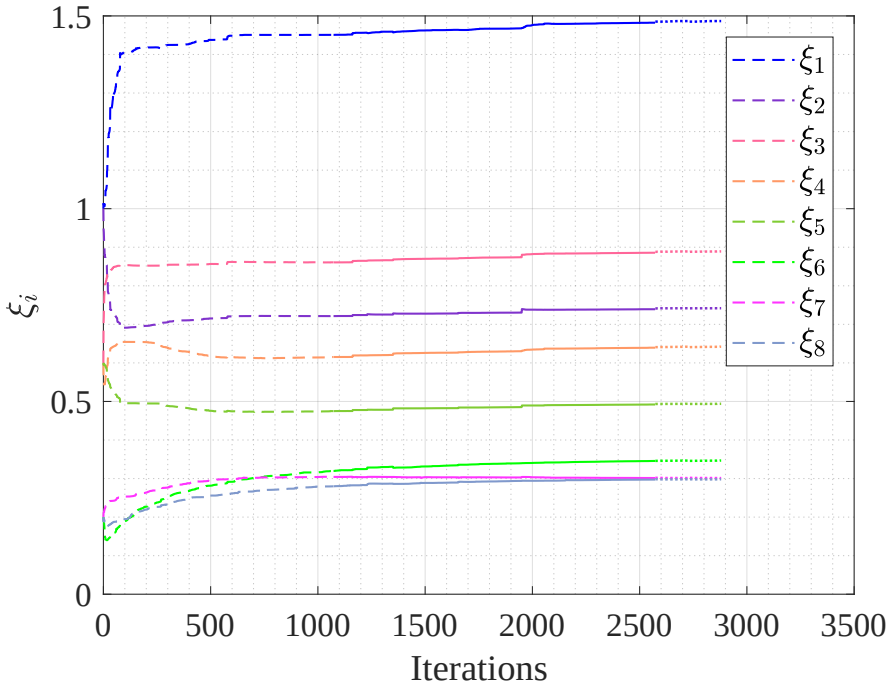
with increasingly refined meshes (5^3, 10^3, and 20^3). As shown in Fig. 5, the primary moduli parameters ξ_{1-8} converge relatively rapidly, whereas the coupling terms ξ_{9-20} require significantly more iterations. These results confirm the method's applicability to complex anisotropic materials and highlight the benefit of hierarchical refinement in managing computational cost.

4 DISCUSSION AND CONCLUDING REMARKS

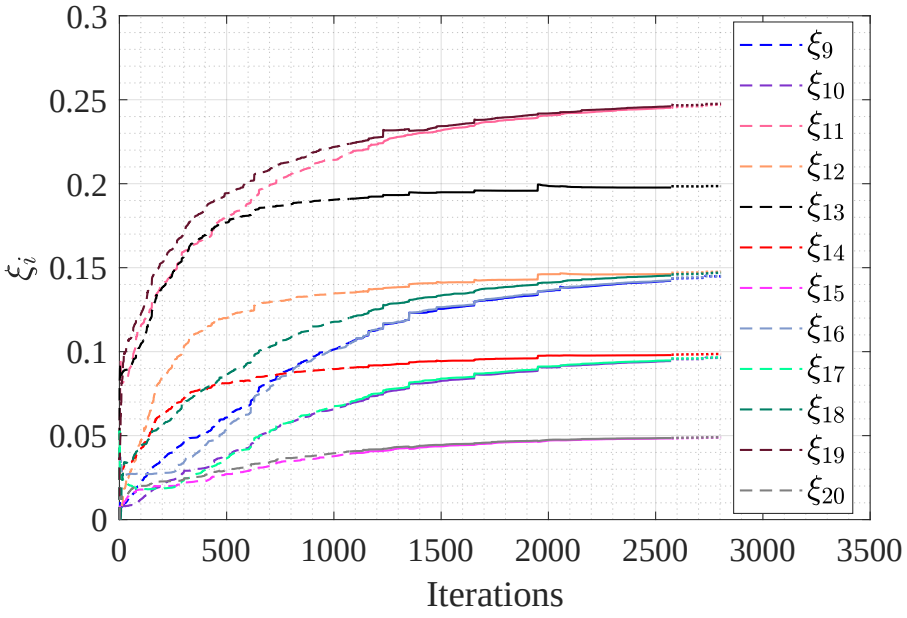
This work presents a simulation-based framework for the characterisation and design of anisotropic open-cell materials, integrating two complementary methodologies: geometric modelling through parametric cell distortions, and inverse estimation of the elastic stiffness tensor from synthetic displacement fields. Together, these methods support both forward design and numerical identification tasks in the development of architectured materials with directional stiffness properties.

The geometric modelling strategy introduced in [8] enables systematic distortion of Kelvin unit cells through twist, tilt, and stretch transformations, each inducing characteristic patterns of anisotropy. These low-dimensional shape transformations provide intuitive design handles for stiffness tuning, while maintaining ligament connectivity and periodicity. Beam-based finite element models are used to efficiently compute homogenised stiffness tensors under periodic boundary conditions, revealing systematic structure–property trends relevant to mechanical and acoustic metamaterials [10, 11, 18].

Complementing this, the inverse characterisation method developed in [6] enables identification of the anisotropic Hooke's tensor of the solid phase from displacement responses simulated on homogenised solid models. The method formulates a constrained optimi-



(a)



(b)

Figure 5: Iteration history of parameters ξ_i for an artificial anisotropic material with three refinements in the multi-level optimisation approach: dashed line, solid line and dotted line represent 5^3, 10^3, 20^3-element meshed solid model respectively.

sation problem operating in a reduced parameter space of 20 dimensionless elasticity components, ensuring compliance with symmetry and thermodynamic constraints. The use of synthetic input data enables controlled validation of the method, while a hierarchical refinement strategy significantly improves convergence efficiency, particularly in fully anisotropic settings.

A key insight from this combined approach is the strong agreement observed between geometrically induced stiffness anisotropy and the stiffness tensors recovered through inverse estimation. In particular, twist and tilt transformations reproduce coupling terms and symmetry-breaking features that are accurately captured by the inverse procedure. This coherence highlights the potential of combining geometric modelling and inverse estimation in a feedback loop for design verification and calibration, even in the absence of detailed microstructural data.

Several limitations and future directions merit consideration. The inverse method assumes uniform boundary conditions and may require regularisation for bending-dominated or under-constrained configurations. While the present study focuses on linear elasticity, extensions to frequency-dependent or viscoelastic behaviour could broaden its applicability, particularly for noise control and damping materials [12, 31]. Further, coupling the current framework with wave-based identification strategies [23], homogenisation-informed design [13], or model order reduction techniques [14] may enhance scalability for large-scale or multiscale optimisation tasks.

In conclusion, the presented methodology offers a coherent path from parametric microstructural transformations to quantitative characterisation of anisotropic stiffness. By bridging geometric modelling and numerical identification, it provides a versatile and computationally efficient basis for the analysis and design of architectured materials with directional mechanical properties.

ACKNOWLEDGEMENTS

This work was partially funded through the following projects: AERIALIST project no 723367 European Union H2020; the Centre for ECO2 Vehicle Design; the Swedish Research Council (VR Grant No. 2021-05791); the European Union's Horizon Europe Programme under the Marie Skłodowska-Curie Actions grant agreement No. 101149748 (project MuInSE).

REFERENCES

[1] L. J. Gibson and M. F. Ashby. *Cellular Solids: Structure and Properties*. Cambridge University Press, 1997. doi:10.1017/CBO9781139878326.

[2] H. Zhu, J. Knott, and N. Mills. Analysis of the elastic properties of open-cell foams with tetrakaidecahedral cells. *Journal of the Mechanics and Physics of Solids*, 45(3): 319–343, 1997. doi:10.1016/S0022-5096(96)00090-7.

[3] L. J. Gibson and M. F. Ashby. The mechanics of three-dimensional cellular materials. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 382(1782):43–59, 1982. doi:10.1098/rspa.1982.0081.

[4] S. C. Cowin. Bone poroelasticity. *Journal of Biomechanics*, 28(12):1371–1385, 1995. doi:10.1016/0021-9290(95)00087-9.

[5] J. Cuenca and P. Göransson. Inverse estimation of the elastic and anelastic properties of the porous frame of anisotropic open-cell foams. *The Journal of the Acoustical Society of America*, 132(2):621–629, 2012. doi:10.1121/1.4731222.

[6] H. Mao, R. Rumpler, and P. Göransson. An inverse method for characterisation of the static elastic hooke's tensors of solid frame of anisotropic open-cell materials. *International Journal of Engineering Science*, 147:103198, 2020. doi:10.1016/j.ijengsci.2019.103198.

[7] C. V. der Kelen, J. Cuenca, and P. Göransson. A method for characterisation of the static elastic properties of the porous frame of orthotropic open-cell foams. *International Journal of Engineering Science*, 86:44–59, 2015. doi:10.1016/j.ijengsci.2014.10.005.

[8] H. Mao, R. Rumpler, M. Gaborit, P. Göransson, J. Kennedy, D. O'Connor, D. Trimble, and H. Rice. Twist, tilt and stretch: From isometric kelvin cells to anisotropic cellular materials. *Materials Design*, 193:108855, 2020. doi:10.1016/j.matdes.2020.108855.

[9] W. E. Warren and A. M. Kraynik. Linear elastic behavior of a low-density kelvin foam with open cells. *Journal of Applied Mechanics*, 64(4):787–794, 1997. doi:10.1115/1.2788983.

[10] R. M. Sullivan, L. J. Ghosn, and B. A. Lerch. A general tetrakaidecahedron model for open-celled foams. *International Journal of Solids and Structures*, 45(6):1754–1765, 2008. doi:10.1016/j.ijsolstr.2007.10.028.

[11] P. Kumar, F. Topin, and L. Tadrist. Geometrical characterization of kelvin-like metal foams for different strut shapes and porosity. *Journal of Porous Media*, 18(6): 637–652, 2015. doi:10.1615/jpormedia.v18.i6.70.

[12] H. Mao, M. Gaborit, E. Lundberg, R. Rumpler, B. Yin, and P. Göransson. Dynamic behaviour of low- to high-density anisotropic cellular materials. *Journal of Sound and Vibration*, page 117137, 2022. doi:10.1016/j.jsv.2022.117137.

[13] X. Li, H. Mao, M. Ichchou, R. Rumpler, L. Shao, and P. Göransson. A new wave-based structural identification framework for estimating material properties of honeycomb sandwich structural components. *Engineering Structures*, 322:105278, 2025. doi:10.1016/j.engstruct.2024.119042.

[14] R. Rumpler and Q. Aumann. Mwcawe: A multivariate wcawe approach for parametric model order reduction, and a sampling strategy for the bivariate case. *Journal of Computational Physics*, 490:112319, 2023. doi:10.1016/j.jcp.2023.112319.

[15] W. E. Warren and A. M. Kraynik. Linear elastic behavior of a low-density kelvin foam with open cells. *Journal of Applied Mechanics*, 55(2):341–346, 1988. doi:10.1115/1.3173676.

[16] W. T. L. Kelvin). LXIII. on the division of space with minimum partitional area. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 24(151):503–514, 1887. doi:10.1080/14786448708628135.

[17] R. M. Sullivan, L. J. Ghosn, and B. A. Lerch. A finite element model for evaluating the effective properties of open-cell foams with random geometry. *Journal of Cellular Plastics*, 45(6):487–503, 2009. doi:10.1177/0021955X09345746.

[18] G. Campoli, M. Borleffs, S. A. Yavari, R. Wauthle, H. Weinans, and A. A. Zadpoor. Mechanical properties of open-cell metallic biomaterials manufactured using additive manufacturing. *Materials & Design*, 49:957–965, 2013. doi:10.1016/j.matdes.2013.02.063.

[19] W. Y. Jang and S. Kyriakides. On the compressive strength of open-cell metal foams with kelvin and random topology. *International Journal of Solids and Structures*, 47 (21):2872–2883, 2010. doi:10.1016/j.ijsolstr.2010.06.010.

[20] T. Storm, C. Gonzalez, and J. Llorca. Topology, anisotropy and inelastic deformation of open-cell foams: Micromechanics and finite element modeling. *Mechanics of Materials*, 135:109–121, 2019. doi:10.1016/j.mechmat.2019.04.007.

[21] L. Gong, S. Kyriakides, and W.-Y. Jang. Compressive response of open-cell foams. part i: Morphology and elastic properties. *International Journal of Solids and Structures*, 42(5-6):1355–1379, 2005. doi:10.1016/j.ijsolstr.2004.07.023.

[22] J. Wu, A. Mcloren, L. Parry, and M. Y. Wang. Anisotropic lattice structure design via density-based topology optimization. *Computer Methods in Applied Mechanics and Engineering*, 344:602–622, 2019. doi:10.1016/j.cma.2018.10.030.

[23] X. Li, H. Mao, M. Ichchou, R. Rumpler, L. Shao, and P. Göransson. Generalized algebraic k-space identification technique for multidimensional signals: Application to wave and energy propagation characterization of curved structures. *Mechanical Systems and Signal Processing*, 225:112304, 2025. doi:10.1016/j.ymssp.2025.112304.

[24] D. Royer and E. Dieulesaint. *Elastic waves in solids I, II*. Springer-Verlag Berlin Heidelberg, 2000.

[25] M. Mélon, R. Panneton, and N. Atalla. Determination of the bulk properties of sound absorbing materials using an impedance tube. part i. theory. *The Journal of the Acoustical Society of America*, 104(5):2871–2878, 1998. doi:10.1121/1.423873.

[26] S. Sahraoui, E. Mariez, and M. Etchessahar. Linear elastic properties of anisotropic open-cell foams. *The Journal of the Acoustical Society of America*, 110(1):635–637, 2001. doi:10.1121/1.1378351.

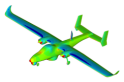
[27] R. Guastavino and P. Göransson. A 3d displacement measurement methodology for anisotropic porous cellular foam materials. *Polymer Testing*, 26(6):711–719, 2007. doi:10.1016/j.polymertesting.2007.02.008.

[28] S. C. Cowin and M. M. Mehrabadi. On the identification of material symmetry for anisotropic elastic materials. *The Quarterly Journal of Mechanics and Applied Mathematics*, 40(4):451–476, NOV 1987. doi:10.1093/qjmam/40.4.451.

[29] P. Kowalczyk. A review of symmetry-based elasticity tensor classification methods. *Archives of Computational Methods in Engineering*, 16(4):515–548, 2009. doi:10.1007/s11831-009-9037-z.

[30] H. Mao, R. Rumpler, and P. Göransson. A note on the linear deformations close to the boundaries of a cellular material. *Mechanics Research Communications*, page 103657, 2021. doi:10.1016/j.mechrescom.2021.103657.

[31] A. Geslain, O. Dazel, J.-P. Groby, S. Sahraoui, and W. Lauriks. Influence of static compression on mechanical parameters of acoustic foams. *The Journal of the Acoustical Society of America*, 130(2):818–825, 2011. doi:10.1121/1.3605535.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

OPTIMIZING TRANSITION PROFILES FOR TAILSITTER UAVS: A GRADIENT DESCENT APPROACH

Alexandre Athayde[1*], Alexandra Moutinho[1] and José Raúl Azinheira[1]

1: IDMEC - Instituto de Engenharia Mecânica
Instituto Superior Técnico
Universidade de Lisboa
1049-001, Lisbon, Portugal
alexandre.athayde@tecnico.ulisboa.pt, alexandra.moutinho@tecnico.ulisboa.pt,
jose.raul.azinheira@tecnico.ulisboa.pt

Abstract. *Tailsitter uncrewed aerial vehicles (UAV) operate in two primary flight modes: hovering, where the UAV weight is sustained only by the thrust of the motors, and aerodynamic flight, where lift generated by the wing contributes to weight support, similar to fixed-wing aircraft. Transitioning between these modes presents a significant challenge due to the limited control authority of the elevons at low airspeeds, as their effectiveness depends on the slipstream from the propellers. Conventional transition strategies rely on simplistic methods, such as accelerating until sufficient airspeed is reached for effective control surface actuation, or by changing the reference attitude from one mode to the other. Naturally, these approaches often lack precision and can lead to inefficient or even unstable transitions.*

To address the specific challenge of tailsitter transitions, the research conducted in this article proposes an optimized strategy based on a comprehensive mapping of stable trim points across a range of airspeeds and climb angles. Using an iterative gradient descent optimization method, a complete envelope of steady-flight trim conditions is obtained, and is then analysed to determine regions of viable operation and evaluate control limitations. From this analysis, a structured transition profile is derived, specifying optimal pitch angles and airspeeds for both forward (vertical to aerodynamic flight) and backward (aerodynamic to vertical) transitions. The proposed profiles are then tested using MATLAB/Simulink based simulations, incorporating an incremental control approach to regulate the transition process.

The obtained results demonstrate that the optimized transition profiles allow for smooth and continuous transitions with low acceleration levels, minimizing abrupt changes in dynamics. Additionally, a trade-off is identified between transition smoothness and altitude gain, providing valuable insights for future implementations in experimental flight scenarios. This methodology offers a systematic and efficient approach to improving the manoeuvrability of tailsitters, contributing to enhanced stability and performance in transitioning flight conditions.

Keywords: tailsitter; UAV; VTOL; transition optimization; gradient descent

1 INTRODUCTION

Tailsitter uncrewed aerial vehicles (UAVs) are a class of VTOL (Vertical Take-Off and Landing) [1] aircraft distinguished by their ability to takeoff vertically and then transition to aerodynamic flight by significantly pitching forward [2], being capable of operating in these two fundamentally different flight modes [2, 3]. In vertical flight, the aircraft's weight is sustained entirely by the thrust generated by its propellers, allowing for hover-like behavior and low-speed maneuverability. In contrast, during aerodynamic flight, the wings generate lift, significantly contributing to the balance of forces and enabling efficient high-speed operation similar to conventional fixed-wing aircraft [4]. These two conditions represent well-known equilibrium states that enable controlled, constant-altitude flight in both modes [5].

Traditionally, tailsitters are interpreted as aircraft that operate mainly in these two primary equilibrium states, switching rapidly between one and the other. However, if the constraint of maintaining a constant altitude is removed - allowing for controlled variations in altitude at a constant rate - a broader spectrum of equilibrium states are obtainable [6], bridging the gap between vertical and aerodynamic flight. By systematically identifying these equilibrium points across a range of airspeeds and climb angles [5], it becomes possible to construct a complete trimmed flight envelope. This envelope provides critical insights into the aircraft's operational limits, the effectiveness of its control surfaces, and the role of different aerodynamic forces throughout its flight regime.

Naturally, a particularly important application of this analysis is the transition phase between the aforementioned hovering and constant-altitude aerodynamic flight modes. The transition from vertical to aerodynamic flight (denoted in this article as "forward transition") and the reverse maneuver back to vertical flight (referred to as "backward transition") are inherently complex, as they involve significant changes in aerodynamic forces, influencing significantly the control effectiveness [4, 7]. By analysing the trimmed flight envelope, it is possible to delineate a relationship between the airspeed and the aircraft's pitch that ensures stability of the UAV. This enables the aircraft to follow a smooth and controlled trajectory through its transition phases, reducing abrupt changes in dynamics and thus ensuring a more robust guidance strategy. Addressing these topics, the main contributions of the research conducted for this article are:

- The generation of a trimmed flight envelope for a tailsitter UAV, covering airspeeds from 0 to 10 m/s and climb angles from 0 to 90 degrees, using an iterative gradient descent optimization approach.
- The development of optimized transition profiles using the obtained envelope that facilitate quasi-stable transitions between vertical and aerodynamic flight and vice-versa, ensuring control authority and maneuvering precision.
- The validation of these transition profiles in a simulated environment, evaluating the impact of the transition time in the overall stability of the UAV, control action effort and altitude variation.

By improving our understanding of these equilibrium states and their role in transition dynamics, this work provides a foundation for the development of more efficient, stable, and reliable tailsitter UAV guidance laws, ultimately contributing to advancements in VTOL aircraft performance and autonomy.

2 TAILSITTER TRIM ENVELOPE USING GRADIENT DESCENT

For the objective of generating a trimmed envelope of a tail-sitter UAV, two components are necessary: firstly, a model of such aircraft should be introduced, namely by describing its equations of motion (EOM); and then the application of the trimming process to this model ought to be outlined. Both these topics are addressed in this section, concluding with the presentation of the trimmed envelope for pre-defined ranges of climb angles and airspeed values.

2.1 UAV model and simulator

The aircraft model is based on the X-Vert VTOL, a bi-rotor and control surface equipped, 200-gram radio-controlled tailsitter manufactured by E-Flite/Horizon Hobby, which has since been discontinued. Such model is described in detail in [7], being largely based in [4], and has been integrated in a MATLAB/Simulink environment. The fundamentals of this model are the following four vector equations of motion, mathematically describing the time-derivatives of the state vector $\mathbf{x} = [\mathbf{v}_g^B, \mathbf{w}_g^B, \mathbf{p}^{NED}, \mathbf{q}^{NED}]^T$ [5, 8]:

$$\dot{\mathbf{v}}_g^B = \frac{1}{m_{uav}}\mathbf{f}^B - (\mathbf{w}_g^B \times \mathbf{v}_g^B) \tag{1}$$

$$\dot{\mathbf{w}}_g^B = \mathbf{J}_{uav}^{-1}(\mathbf{m}^B - \mathbf{w}_g^B \times \mathbf{J}_{uav}\mathbf{w}_g^B) \tag{2}$$

$$\dot{\mathbf{p}}^{NED} = \mathbf{R}_B^{NED}\, \mathbf{v}_g^B \tag{3}$$

$$\dot{\mathbf{q}}^{NED} = \frac{1}{2}\Big(\mathbf{S}_{\mathbf{q}}^{\mathbf{w}}\, \mathbf{q}^{NED}\Big) \tag{4}$$

where $\mathbf{v}_g^B = [u_g, v_g, w_g]^T$ and $\mathbf{w}_g^B = [p_g, q_g, r_g]^T$ are the linear and angular velocities with respect to the ground, expressed in metres per second and radians per second, respectively, whilst the $\mathbf{p}^{NED} = [N, E, D]^T$ and $\mathbf{q}^{NED} = [q_0, q_1, q_2, q_3]^T$ denote the position in metres and the quaternion-based orientation with the respect to the fixed North-East-Down frame. A quaternion-attitude representation is preferred over the conventional Euler angles due the the wide-range of attitude variations tailsitter encounter. Equation (3) and (4) show two auxiliary matrices, $\mathbf{R}_B^{NED}$ and $\mathbf{S}_{\mathbf{q}}^{\mathbf{w}}$ that express the kinematic relationships between the aforementioned velocities and the position and orientation of the UAV [5]. Lastly, m_{uav} and $\mathbf{J}_{uav}$ represent the mass and inertia of the aircraft, and the overall force and moment vectors are denoted by $\mathbf{f}^B = [F_x, F_y, F_z]^T$ and $\mathbf{m}^B = [M_x, M_y, M_z]^T$.

When it comes to the actuators, the X-Vert is equipped with two elevons and two proprotors, denoting combined functions of aileron/elevator and propeller/rotor, respectively. The input vector consists on the aileron and elevator deflections, expressed in radians, and the rudder and throttle signals, taken as dimensionless between 0 (zero) and 1, leading to $\mathbf{u} = [\delta_a, \delta_e, \tau_r, \tau_t]^T$.

The complexity of this tailsitter model arises from different aspects that contribute to $\mathbf{f}^B$ and $\mathbf{m}^B$, namely accouting for large angle-of-attack (AOA) aerodynamics [9, 10], non-axial flow on the propellers [11], and the elaborate slipstream effects on the elevons due to the positioning of these behind the proprotors [12], which play a substantial part in the control authority of these surfaces. For a detailed explanation of the UAV model and its underlying components, the reading of [7] and [4] is recommended as it was considered that a thorough description of this model would be redundant with these works and would also fall out of the scope of this article.

Lastly, for the purposes of this work, some simplifications for the model were taken: firstly, the absence of wind is assumed, and therefore the ground and air velocity vectors are identical, $\mathbf{v}_g^B = \mathbf{v}_a^B$ and $\mathbf{w}_g^B = \mathbf{w}_a^B$; constant environmental properties are taken, namely the acceleration of gravity, $g = 9.8065$ m/s^2, and air density, $\rho = 1.2250$ kg/m^3; and no contact forces and moments were considered in the trim process, since in normal flight conditions these are absent.

2.2 Gradient descent application in aircraft trim

Gradient descent is an unconstrained optimization method that uses the gradient of a given function $\mathfrak{f}(\mathbf{x}, \mathbf{u})$, represented by $\nabla\mathfrak{f}$, as an opposite direction for the minimization of such function. Equation (5) illustrates the progression of the solution to find the optimal point $\mathbf{p}$ from iteration k to $k+1$ [13]

$$\mathbf{p}_{k+1} = \mathbf{p}_k - \mu\nabla\mathfrak{f} \tag{5}$$

where μ denotes the step size and $\mathfrak{f}$ is taken to be a function of $\mathbf{p}$.

Several different applications of gradient descent exist with numerous variations, namely differing in the choice of search direction, although usually somewhat based on $\nabla\mathfrak{f}$, or step size μ. In this work, the Quasi-Newton Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm [14–17] was used, as it consists on a computationally-effective optimization method that not only uses the information from the gradient, but also from an approximate Hessian matrix at each iteration, $\mathbf{H}_k$, in order to improve the convergence to the solution [13], which is obtainable from:

$$\mathbf{H}_{k+1} = \mathbf{H}_k + \frac{\mathbf{q}_k\mathbf{q}_k^T}{\mathbf{q}_k^T\mathbf{s}_k} - \frac{\mathbf{H}_k\mathbf{s}_k\mathbf{s}_k^T\mathbf{H}_k^T}{\mathbf{s}_k^T\mathbf{H}_k\mathbf{s}_k} \tag{6}$$

where $\mathbf{s}_k = \mathbf{x}_{k+1} - \mathbf{x}_k$ and $\mathbf{q}_k = \nabla\mathfrak{f}(\mathbf{x}_{k+1}) - \nabla\mathfrak{f}(\mathbf{x}_k)$. Despite the improvement that the Quasi-Newton BFGS provides over more traditional gradient-based optimization methods, it still suffers from the major drawback of these: the possibility for the convergence process to stop at a local minima, instead of the global minimum. To ensure robustness and accuracy when using such methods, the initial estimate for the solution should be reasonably close to the desired solution, which may many times be impractical to compute. The method for addressing this problem when generating a complete trim envelope is described in the beginning of section 2.3. Nevertheless, some considerations must first be provided when it comes to the cost function to be minimized, and how the constraints can be included in it, forming an unconstrained minimization problem, allowing for the application of the Quasi-Newton BFGS algorithm.

- Trim process considerations
  A trim point of an aircraft consists on a state where it is flying with certain constant aspects, namely constant airspeed and/or constant altitude [5], corresponding to equilibrium conditions. The research conducted in this article focuses on the trim as a longitudinal problem, targeting wings-level and constant airspeed flight conditions, where the lateral variables should take null values. For this purpose, longitudinal subsets of the state and input vectors can be defined as $\mathbf{x}_{long} = [u_g, w_g, q_g, N, D, q_2]^T$ and $\mathbf{u}_{long} = [\delta_e, \tau_t]^T$, allowing for a more compact notation, and where q_0 is omitted from $\mathbf{x}_{long}$ as it becomes redundant with q_2 under these conditions. The lateral

variables can also be grouped in $\mathbf{x}_{lat} = [v_g, p_g, r_g, E, q_1, q_3]^T$ and $\mathbf{u}_{lat} = [\delta_a, \tau_r]^T$, but remain irrelevant for the trim process.

For achieving flight at constant velocity, it makes sense to assume that the derivatives of all components of $\mathbf{x}_{long}$ should be zero, with the exceptions of $\dot{N}$ and $\dot{D}$. In fact, the conditions when $\dot{u}_g = \dot{w}_g = \dot{q}_g = \dot{q}_2 = 0$ can correspond to three different situations depending on the time derivatives of D and N: if $\dot{D} = 0$ and $\dot{N} \neq 0$, a constant altitude and airspeed flight is being performed; vertical flight at constant airspeed is defined by $\dot{N} = 0$ and $\dot{D} \neq 0$; and when both $\dot{D}$ and $\dot{N}$ are zero, the UAV is either hovering or completely stopped.

It is evident from Eq.(1), Eq.(2) and Eq.(4) that $\dot{u}_g = \dot{w}_g = \dot{q}_g = \dot{q}_2 = 0$ when $\mathbf{f}^B$, $\mathbf{m}^B$ and $\mathbf{w}_g^B$ are also zero. Furthermore, the aircraft should fly at a certain AOA and airspeed, defined by Eq.(7) and Eq.(8):

$$\alpha = \arctan\left(\frac{u_g}{w_g}\right) \tag{7}$$

$$V_t = \sqrt{u_a^2 + v_a^2 + w_a^2} \tag{8}$$

Depending on the pitch of the UAV, obtainable from the quaternion representation using $\theta = \arcsin\left(2(q_0 q_2 - q_1 q_3)\right)$ (or simply $\arcsin\left(2(q_0 q_2)\right)$ when $\mathbf{x}_{lat}$ is assumed to be zero), it can be either climbing, descending or maintaining altitude, according to the sign of the climb angle γ:

$$\gamma = \theta - \alpha \tag{9}$$

The desired trim conditions of a steady-state, constant velocity and climbing rate can be summarized by combining Eq.(7), Eq.(8) and Eq.(9) with the transformations from the air frame to the body referential [5], resulting in the following equality, where the asterisk is used to denote the values of the variables at such equilibrium:

$$\begin{bmatrix} \dot{u}_g \\ \dot{w}_g \\ \dot{q}_g \\ \dot{N} \\ \dot{D} \\ \theta \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ V_t^* \cos\left(\alpha^*\right) \\ V_t^* \sin\left(\alpha^*\right) \\ \theta^* \end{bmatrix} \tag{10}$$

Inspecting Eq.(10), it is clear that the first three derivatives are readily available to be included in a cost function, as they should be minimized to zero. Nevertheless, if possible, it is useful to replace them directly by the forces and moments, F_x, F_z and M_y, allowing to bypass the equations of dynamics whilst still representing these equilibrium conditions. In addition to this, the constraints related to V_t, α and θ must be formulated in a way that allows the employment of unconstrained

optimization methods, resulting in the following vector

$$\mathbf{v}_c = \begin{bmatrix} F_x \\ F_z \\ M_y \\ V_t - V_t^* \\ \gamma - \gamma^* \\ q_0^2 + q_2^2 - 1 \end{bmatrix} \tag{11}$$

where the restriction of unitary quaternion norm was also included. Additional constraints could have been included in Eq.(11), namely those related with the limits of the actuators, as the elevons will have maximum and minimum deflections, and the throttle signal sent to each motor is taken from 0 (zero) to 1. However, for obtaining the previously described trim envelope, a decision was made to drop such constraints and allowing the trim process to find solutions that could go beyond these, and then signal if such limitations were violated or not.

From Eq.(11), it is possible to obtain a quadratic cost function J_c, which can be formulated so that it includes a diagonal matrix $\mathbf{K}_c$ that provides different weights to each component in its minimization process. The final cost function used in the addressed tailsitter trim problem is:

$$J_c = \mathbf{v}_c^T \cdot \mathbf{K}_c \cdot \mathbf{v}_c \tag{12}$$

The obtained solution of Eq.(12) for a certain trim point will correspond to a set of values for $\mathbf{x}^*$ and $\mathbf{u}^*$, once again using the asterisk to denote that these correspond to an equilibrium condition. The relevant information for a trim point of these two vectors can be summarized into a more compact vector $\mathbf{z}^* = [\alpha^*, \delta_e^*, \tau_t^*]^T$, as it should be unique for the flight conditions specified for $V_t = V_t^*$ and $\gamma = \gamma^*$.

2.3 Trimmed envelope for constant airspeed and climbing rate

In order to generate the trimmed envelope for the described conditions, each equilibrium point must be obtained individually by minimizing Eq.(12) for the specified V_t^* and γ^* values. For this end, the lower-bound in γ is taken as 0 (zero) because negative values at low velocities will correspond to a descending, almost vertical flight, which leads to AOAs near π or $-\pi$ radians, in which the aerodynamic models may lose validity. The upper boundary of V_t was taken as 10 m/s, as it was expected to be near the limits of this UAV [4]. Therefore, the considered ranges for these variables were:

$$\gamma^* \in \left[0; \frac{\pi}{2}\right] \quad , \quad V_t^* \in \Big[0; 10\Big]$$

Using the ranges in Eq.(13) and defining the maximum number points for γ and V_t respectively as N_γ and N_{V_t}, γ^*_{range} and $V^*_{t,range}$ can be defined as vectors that have values for these variables spaced by $(\gamma_{max} - \gamma_{min})/N_\gamma$ and $(V_{t_{max}} - V_{t_{max}})/N_{V_t}$. Using i and j as iteration variables, an iterative procedure can sweep across these vectors considering $\gamma_i^* = \gamma^*_{range}(i)$ and $V^*_{t,j} = V^*_{t,range}(j)$ in order to find the different trim points in these ranges. Such procedure can delineated so that it uses the previously-obtained solution

closest to the trim point to be obtained as an initial solution to compute it, addressing directly the previously-described limitation of the gradient-based optimization methods requiring a good initial estimate of the solution to facilitate convergence. Nonetheless, an overall starting point is still required to initialize this algorithm, $\mathbf{z}^*_{1,1}$, and it was chosen as hovering flight, with $\gamma^*_1 = \frac{\pi}{2}$ and $V^*_{t,1} = 0$. This point can be obtained analytically since the thrust is equal to the weight of the UAV and the aerodynamic forces and moments are null under these conditions, resulting in $\mathbf{z}^*_{1,1} = [0, 0, 0.8405]^T$. It should be noted that, with hovering flight being the overall starting equilibrium condition, the values for γ^* will be decreasing, meaning that the spacing of γ^*_{range} will in fact be $(\gamma_{min} - \gamma_{max})/N_\gamma$.

With this iteration process, each trim point is addressed as an individual unconstrained optimization problem that starts with the guess solution $\mathbf{z}^0_{i,j}$, and aims to find the respective solution $\mathbf{z}^*_{i,j}$, and the overall iterative process is summarized in Algorithm 1.

Algorithm 1 Iterative method for finding trim points across the specified ranges of V^*_t and γ^*

$z_{1,1} = [0, 0, 0.8405]^T$ ▷ Initial solution for hovering flight
for $i = 1, ..., N_\gamma$ **do**
  $\gamma^* = \gamma_{i+1}$ and $\mathbf{z}^0_{i+1,1} = \mathbf{z}^*_{i,1}$ ▷ Iterate on γ^*_{range}.
  for $j = 1, ..., N_{V_t}$ **do**
    $V^*_t = V_{t,j+1}$ and $\mathbf{z}^0_{i+1,j+1} = \mathbf{z}^*_{i+1,j}$ ▷ Iterate on $V^*_{t,range}$
    Compute $\mathbf{z}^*_{i+1,j+1}$
      ▷ Run unconstrained optimization problem for conditions defined by γ^*_{i+1} and $V^*_{t,j+1}$
  end for
end for

For the purposes of this article, $N_\gamma = 50$ different points for γ^* and $N_{V_t} = 100$ for V^*_t were considered, respectively corresponding to increments of -0.0321 radians and 0.1010 m/s, and totaling 5000 different equilibrium points to be computed. Figure 1 illustrates the obtained trim points using the Quasi-Newton BFGS unconstrained optimization method described in section 2.2 combined with Algorithm 1. In the analysis of this figure, as well as the remaining figures of this section, the asterisk is omitted, but it should be noted that the obtained variables correspond to equilibrium conditions.

In Fig.1, a colour code was applied to the feasible trim points for visualization convenience: blue points correspond to high γ and low V_t values, meaning conditions near vertical flight; green points represent high airspeed values with low climb angles, when the UAV is performing aerodynamic flight; and when both of these have low values, a continuous transition to cyan is evident, denoting low-airspeed, near-vertical flight conditions. For all the unfeasible trim points, the colour used was red in order to denote conditions in which the actuators were beyond their respective limitations. Analyzing the feasible regions of Fig.1, it is clear that the UAV cannot keep a stable constant altitude when the airspeed is between 4 and 7 metres per second and γ is low, corresponding to situations where the elevons are at their mechanical limit, and the aircraft must be climbing in order to be at equilibrium. Similarly, the upper-right zone of Fig.1 shows another unfeasible zone, denoting that the higher the climb angle, the lower the maximum airspeed, and corresponding to conditions characterized by throttle saturation.

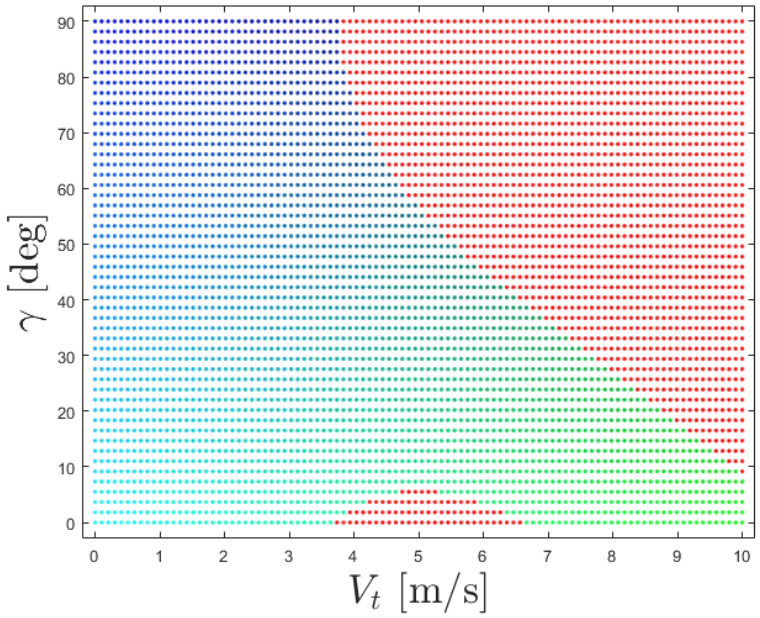


Figure 1: Envelope of trim points obtained for $\gamma^* \in [0; \frac{\pi}{2}]$ rads and$V_t^* \in [0; 10]$ m/s.

From this set of trim conditions obtained for the specified airspeeds and climb angles, the same points can be remapped into the θ/V_t plane, as shown in Fig.2, where the colour code was maintained, also including a black line that corresponds to the minimum velocity within actuator limits for each pitch angle.

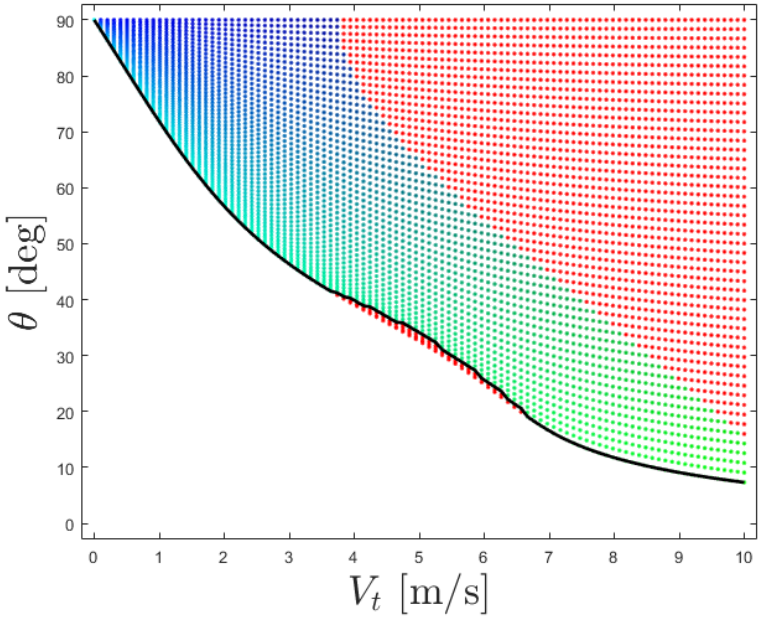


Figure 2: Remapping of the trim points of Fig.1 into the θ/V_t plane, with the line for the minimum feasible values of airspeed at each pitch angle shown in black.

Similarly to the previous figure, Fig.2 shows the region between 4 and 7 metres per second where the trim is unfeasible for lower airspeed values, and the previously-described unfeasible region present when throttle saturation occurs. Nevertheless, this figure hints that it is possible to draw a relationship between V_t and θ, which is discussed next.

- Transition profile
  For the purpose of creating a guidance law that combines V_t and θ, a polynomial regression can be made using the points that correspond to the minimum V_t at each θ that ensures stability, which are represented the black line shown in Fig.2. Although this could be performed directly on these points, it is useful to consider a small margin for the airspeed, aiming to include a safety factor that could help avoid the negative influence of any eventual model mismatches. With this in mind, the values of V_t that correspond to the black line in Fig.2 were increased by $V_{t,margin} = 0.5$ m/s, and the trim points were recomputed accounting for this margin. However, the first and last points were kept at their original values, as shown in Fig.3, since these correspond to the desired hovering and aerodynamic flight conditions which must be present before and after the transitions.

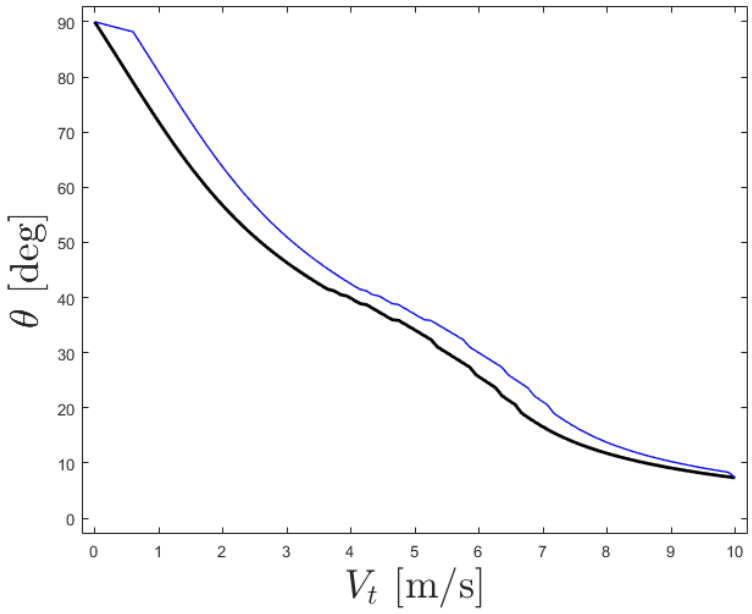


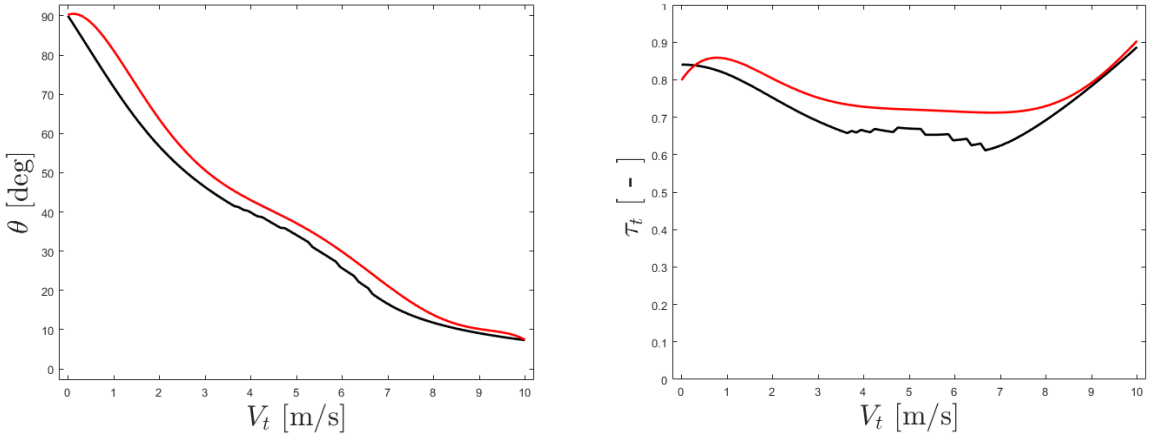
Figure 3: Original minimum feasible airspeed trim points, shown in black, and subsequent recomputation of these including a 0.5 m/s in airspeed margin, in blue.

Using these pairs of V_t/θ as shown in Fig.3, a polynomial regression can be made in order to provide a continuous law that states θ as a function of V_t. For such end, a 6-th degree polynomial was chosen, as it was verified that it captured the overall shape of the desired profile without overfitting it, and its expression is:

$$\begin{aligned}\theta\left(V_t\right) \quad &= (-6.398 \cdot 10^{-5})V_t^6 + 0.0021V_t^5 - 0.0255V_t^4 \\ &+0.1442V_t^3 - 0.3571V_t^2 0.0771V_t + 1.5754\end{aligned} \tag{13}$$

Furthermore, it is also beneficial to obtain an analogous polynomial for τ_t, with the goal of ensuring that the elevons are granted enough authority during the transitions [4]. This results in an additional equation for the lower-bound of the throttle input that must be respected when transitioning:

$$\begin{aligned}\tau_{t,min}\left(V_t\right) \quad &= (-2.050 \cdot 10^{-5})V_t^6 + (7.05 \cdot 10^{-4})V_t^5 - 0.0092V_t^4 \\ &+0.0581V_t^3 - 0.1742V_t^2 + 0.1816V_t + 0.7988\end{aligned} \tag{14}$$



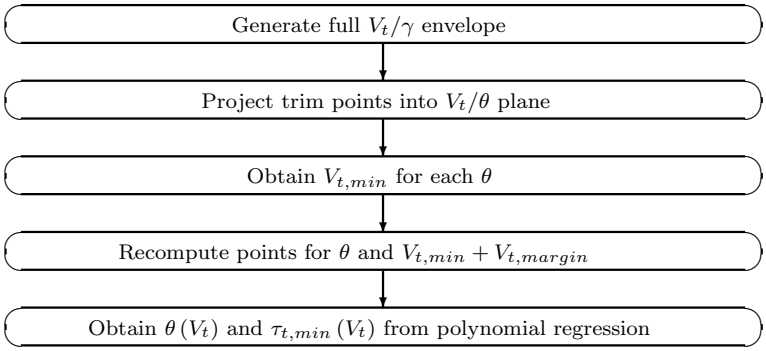
(a) Profile for $\theta\,(V_t)$, in red. (b) Obtained profile for $\tau_{t,min}\,(V_t)$, in red.

Figure 4: Plots for the polynomials obtained from regression of the trim points including a 0.5 m/s margin in airspeed, in red, and the original trim points, in black.

Figure 4 shows the graph of Eq.(13) and Eq.(14) for the considered airspeed interval. From its analysis, it is clear that the polynomials of Eq.(13) and Eq.(14) manage to capture the general curvature of the profile including the aforementioned airspeed margin. When it comes specifically to the graphical representation of $\tau_{t,min}\,(V_t)$, there is a crossing with the black, noting the absence of margin for low airspeed values, and an overlap when it is high, but both these situations correspond to well-known and stable flight conditions, and therefore it is assumed that these discrepancies are not prejudicial for the design of a transition profile.

To summarize this section, Algorithm 2 provides an overview of the different steps followed to design a transition profile starting from the complete V_t/γ envelope obtained in Fig.1.

Algorithm 2 Method for obtaining a transition profile from the mapping of V_t/γ trim points



3 SIMULATION OF FORWARD AND BACKWARD FLIGHT TRANSITIONS

Having been obtained the transition profile based on stable trim points in the V_t/θ map, a validation of this profile can be attempted using the simulated model of the aircraft described in section 2.1. For this objective, some considerations related with the starting conditions and the choice of control strategy are provided first, followed by the results of the simulations for the forward and backward transitions, addressed in sections 3.1 and 3.2 respectively. It should be noted that, although the aircraft trim and analysis were performed only for the longitudinal states, the simulations account for the complete model of the UAV.

- Equilibrium conditions for transitioning
  For the simulations, the forward transition is defined as starting in hovering flight and then transition to a constant altitude and airspeed of 10 m/s aerodynamic flight, and the backward transition corresponds to the inverse situation. These two conditions correspond to specific trim points, and are represented by the state and input vectors in Eq.(15) and Eq.(16) respectively, where an arbitrary altitude of $D = -20$ metres was considered for its starting point of the simulations.

$$\begin{aligned} \mathbf{x}^{hover} &= [0,0,0,0,0,0,0,0,0,-20,0.7071,0,0.7071,0]^T \\ \mathbf{u}^{hover} &= [0,0,0,0.8405]^T \end{aligned} \tag{15}$$

$$\begin{aligned} \mathbf{x}^{aero,10} &= [9.9182,0,1.2767,0,0,0,0,0,0,-20,0.9980,0,0.0640,0]^T \\ \mathbf{u}^{aero,10} &= [0,-0.2991,0,0.8874]^T \end{aligned} \tag{16}$$

- Incremental control for attitude and velocity
  The Incremental Nonlinear Dynamics Inversion (INDI) is the strategy chosen for controlling the attitude and velocity of the UAV, as it is robust to model imperfections and only requires partial knowledge of the model, namely the influence of its actuators in the states to be tracked [18–20]. The control law has an incremental form, representing how the control action at instant k, $\mathbf{u}_k^{INDI}$, is calculated from the previous one, as illustrated by Eq.(17), which was expanded from [21] to include airspeed regulation [22]:

$$\mathbf{u}_k^{INDI} = \begin{bmatrix} \delta_a \\ \delta_e \\ \tau_r \\ \tau_t \end{bmatrix}_k = \lambda \mathbf{G}^{-1} \left(\begin{bmatrix} \dot{\mathbf{w}}_{g,d}^B \\ \dot{V}_{t,d} \end{bmatrix}_k - \begin{bmatrix} \dot{\mathbf{w}}_g^B \\ \dot{V}_t \end{bmatrix}_k \right) + \begin{bmatrix} \delta_a \\ \delta_e \\ \tau_r \\ \tau_t \end{bmatrix}_{k-1} \tag{17}$$

  in which $\mathbf{G}$ is the matrix that represents the actuator authority in the states to control, $\lambda = 0.8$ is an Input Scaling Gain (IGS) [23], $\dot{\mathbf{w}}_{g,d}^B$ and $\dot{V}_{t,d}$ denote the desired values for these variables, and $\dot{\mathbf{w}}_g^B$ and $\dot{V}_t$ represent the estimated variables, obtained by a second-order derivative filter [24]. $\dot{\mathbf{w}}_g^B$ and $\dot{V}_{t,d}$ are defined in terms of the error variables as depicted by Eq.(19):

$$\begin{aligned} \dot{\mathbf{w}}_{g,d}^B &= \mathbf{K_w}(\mathbf{K_q}\mathbf{q}_{err,1:3} + \mathbf{w}_g^B) \\ \dot{\mathbf{V}}_t &= K_{V_t} V_{t,err} \end{aligned} \tag{18}$$

where $\mathbf{K_w} = \text{diag}(5, 10, 5)$ and $\mathbf{K_q} = \text{diag}(2, 10, 2)$ are diagonal gain matrices, and $K_{V_t} = 1$ is a scalar gain, all of these being represented in their respective SI units. The altitude was not directly regulated since its variation was taken an aspect to compare in the simulations.

The last aspect of Eq.(17) to address is the $\mathbf{G}$ matrix: although it could be deduced analytically for an equilibrium condition [7], this matrix can also be obtained numerically by linearizing a model into the $\dot{\mathbf{x}} = \mathbf{Ax} + \mathbf{Bu}$ [8, 25] form at a specific operating point and then selecting the components of the $\mathbf{B}$ matrix that are relevant. Taking the equilibrium condition as the hovering flight and assuming that δ_a, δ_e and τ_r affect separately each component of $\mathbf{w}_g^B$, and that the influence of τ_t in V_t can be assumed to be the same as in u_g, the $\mathbf{G}$ can be computed:

$$\mathbf{G} = \begin{bmatrix} \frac{\partial \dot{p_g}}{\delta_a} & 0 & 0 & 0 \\ 0 & \frac{\partial \dot{q_g}}{\delta_e} & 0 & 0 \\ 0 & 0 & \frac{\partial \dot{r_g}}{\tau_r} & 0 \\ 0 & 0 & 0 & \frac{\partial \dot{u_g}}{\tau_t} \end{bmatrix} = \begin{bmatrix} -37.491 & 0 & 0 & 0 \\ 0 & -141.772 & 0 & 0 \\ 0 & 0 & -167.305 & 0 \\ 0 & 0 & 0 & 17.822 \end{bmatrix} \tag{19}$$

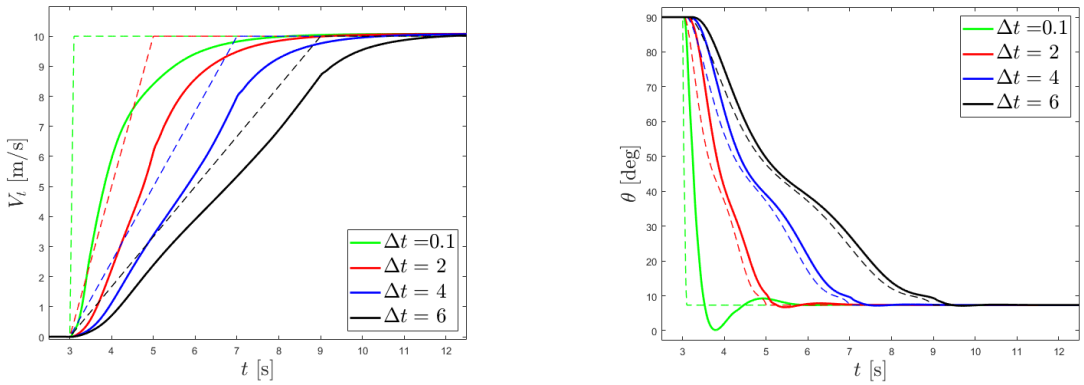
- Simulation process and evaluation metrics
  The simulations were run using Simulink 10.7 (MATLAB 2023a) for 4 different transition intervals Δt: 0.1 (considered instantaneous), 2, 4 and 6 seconds. In the first 3 seconds of simulation time, the controlled UAV model was allowed to run at the starting conditions to ensure stationary conditions, and then the reference airspeed would increase linearly in order to achieve its final value after Δt seconds. The reference attitude and throttle lower bound were provided according to $\theta\left(V_t\right)$ and $\tau_{t,min}\left(V_t\right)$ in equations Eq.(13) and Eq.(14) respectively.

  The metrics to evaluate were the maximum absolute angular acceleration, $\overline{\dot{q}_{g_{max}}}$, the altitude variation ΔD taken in relation to its original value at the end of the simulation, and the possible saturation of the actuators.

3.1 Forward transition results

For the simulations of the forward transition, the simulated UAV started at the equilibrium state defined by $\mathbf{x}^{hover}$, and then performed a controlled transition to $\mathbf{x}^{aero,10}$ within the prespecified Δt interval.

Figure 5 presents the evolution of airspeed and pitch angle over time, illustrating the evolution of the transition process for the different transition intervals. Inspecting firstly Fig.5a, the airspeed plot clearly demonstrates that the transition process follows a time-scaled behavior across the different time intervals. The velocity controller tracks the desired velocity profile, but it is clear that the rapid transitions have a negative impact on its performance. Regarding Fig.5b, the different curves highlight the shape of the transition profile that was obtained in section 2.3, with the INDI controller maintaining precise tracking throughout the transition. However, a noticeable overshoot occurs in the case of $\Delta t = 0.1$s, represented by the green line. Despite this, the overall performance remains robust, with the remaining simulations achieving near-perfect regulation for $\Delta t = 2$, $\Delta t = 4$ and $\Delta t = 6$ seconds.

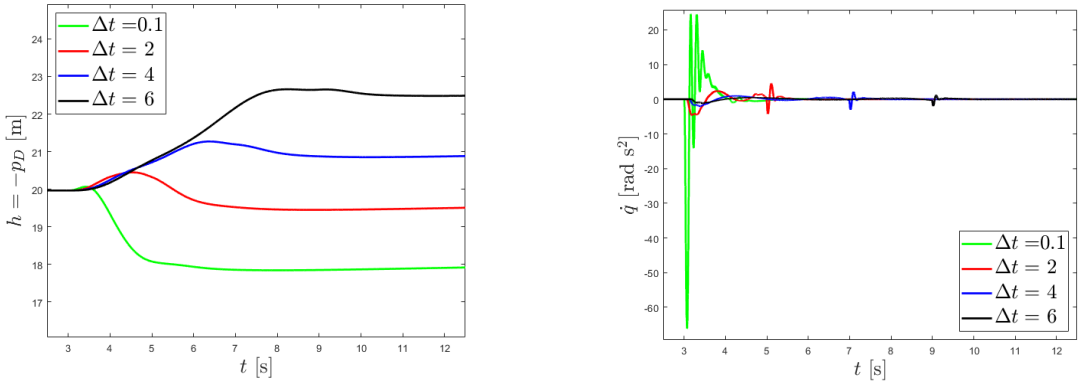


(a) Airspeed reference $V_{t,ref}$ and actual values V_t for simulations with different Δt interval.

(b) Pitch angle reference θ_{ref} and actual values θ for simulations with different Δt intervals.

Figure 5: Airspeed (in the left) and pitch angle (in the right) plots in time for the forward transition,with the references represented in dashed lines, and the actual variables in solid lines, for $\Delta t = 0.1$s (in green), $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

Moving on to the altitude and pitch angular acceleration, respectively shown in Fig.6a and Fig.6b, it is evident by immediate inspection that the $\Delta t = 0.1$s transition provides significantly worse results, by indicating a loss of altitude of 2.2 metres, and an peak of -66.106 rad/s^2 in the pitch acceleration. In contrast, the transition in $\Delta t = 6$s (in black) shows an altitude increase of more than 2 metres, but minimal angular acceleration. Providing a middle-ground between these two results, the $\Delta t = 2$s and $\Delta t = 4$s transition times show minimal altitude and angular acceleration variations.



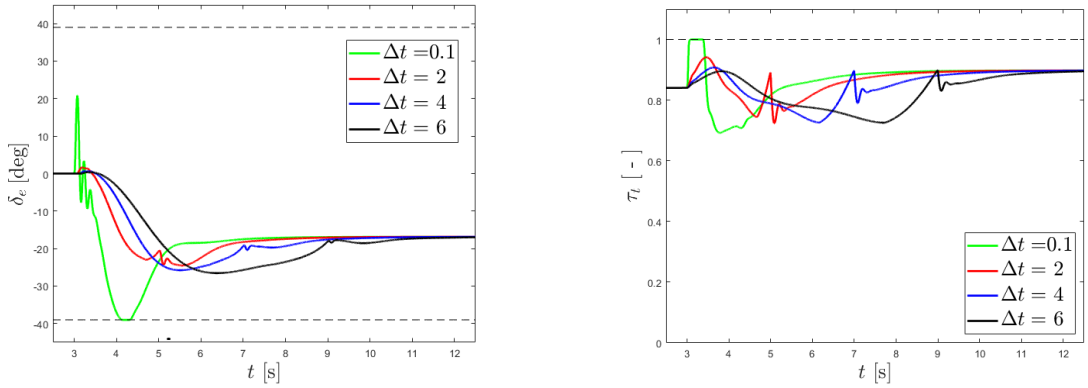
(a) Altitude $h = -D$ for simulations with different Δt intervals.

(b) Pitch angular acceleration $\dot{q}_g$ for simulations with different Δt intervals.

Figure 6: Altitude (in the left) and pitch angular acceleration (in the right) plots in time for the forward transition, for $\Delta t = 0.1$s (in green), $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

The last aspect to inspect is the control input during the simulation, and the plots for its two longitudinal components are shown in Fig.7a and Fig.7b, providing insight into the control effort required for different transition durations. In Fig.7a, the impact of the aforementioned angular acceleration peak for $\Delta t = 0.1$s is evident, with the green line exhibiting significant oscillation at approximately 3 seconds, which is the transition start mark, and saturating near the 4-second mark. Analogous conclusions can be obtained

from Fig.7b, where τ_t saturates in its upper bound for $\Delta t = 0.1$s. From both these figures, the remaining transition times allow controlled actuation levels, demonstrating that longer transition durations ensure smoother and more stable elevator and throttle control, with minimal oscillation.



(a) Elevator input δ_e for simulations with different Δt intervals.

(b) Throttle input τ_t for simulations with different Δt intervals.

Figure 7: Elevator (in the left) and throttle (in the right) plots in time for the forward transition, for $\Delta t = 0.1$s (in green), $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

Table 1 summarizes the results obtained from the forward transition simulations with different times to transition, providing support to the considerations provided in the analysis with quantifiable metrics.

Table 1: Summary of results for the forward transition.

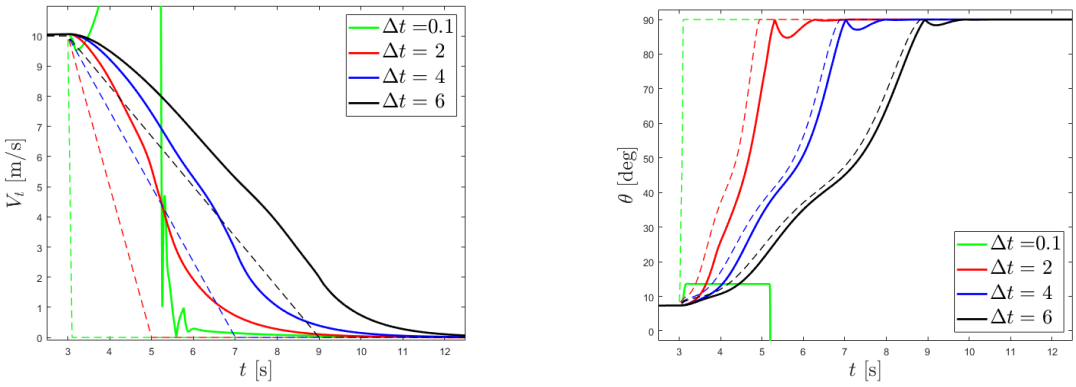
	$\Delta t = 0.1$s	$\Delta t = 2$s	$\Delta t = 4$s	$\Delta t = 6$s
ΔD [m]	2.155	0.548	-1.271	-2.660
$\overline{\dot{q}_{g_{max}}}$ [rad/s^2]	66.106	4.615	3.063	2.017
sat(δ_e)	yes	no	no	no
sat(τ_t)	yes	no	no	no

3.2 Backward transition results

The second part of the simulation results is the backward transition, with the UAV starting at an aerodynamic flight defined by $\mathbf{x}^{aero,10}$, and then transitioning to the hovering flight within the prespecified Δt interval. This transition is known to be significantly challenging for tailsitters to perform since the control strategy must ensure sufficient airspeed in the control surfaces to perform the maneuver, whilst also reducing the overall airspeed of the UAV.

The results for the backward transition are shown in Fig.8a and Fig.8b, which respectively show the airspeed and pitch angle evolution during the simulations for the same Δt intervals. The overall behavior resembles that of the forward transition, with the UAV effectively tracking the desired velocity and attitude profiles across most transition durations. Nevertheless, a critical difference is observed for $\Delta t = 0.1$s, represented by the

green line, where the aircraft fails to perform the transition, and crashes, as evidenced by the loss in altitude shown in Fig.9. This indicates that an excessively rapid backward transition poses greater stability challenges compared to the forward case, likely due to the aforementioned control authority changes when returning to vertical flight, and longer transition durations remain well-regulated and controlled.



(a) Airspeed reference $V_{t,ref}$ and actual values V_t for simulations with different Δt intervals.

(b) Pitch angle reference θ_{ref} and actual values θ for simulations with different Δt intervals.

Figure 8: Airspeed (in the left) and pitch angle (in the right) plots in time for the backward transition,with the references represented in dashed lines, and the actual variables in solid lines, for $\Delta t = 0.1$s (in green), $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

The altitude plot in Fig.9 also shows that all the remaining simulations present a significant gain in height, with the greatest being the one that corresponds to $\Delta t = 6$s and the lowest to $\Delta t = 2$s, which is to be expected. In the remaining plots, the lines that were generated by the $\Delta t = 0.1$s simulation are omitted in order to not overburden the images unnecessarily, as it corresponds to a crash scenario.

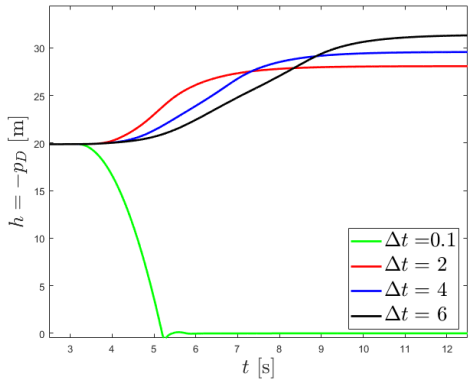


Figure 9: Altitude plot in time for the backward transition, for $\Delta t = 0.1$s (in green), $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

The results for the pitch angular acceleration during the backward transition are shown in Fig.10, where it can be seen that the $\Delta t = 2$s line corresponds to some significant oscillations and peaks, whilst $\Delta t = 4$s and $\Delta t = 6$s show a more robust behavior. Although it appears that these results for the pitch acceleration are in the same order of magnitude of the forward transition when compared to Fig.6a, it is clear that the aforementioned oscillation in $\Delta t = 2$s had a negative impact on the overall control behaviour, since elevator

saturation occurs during almost 1 second, as it can be seen in Fig.11a.

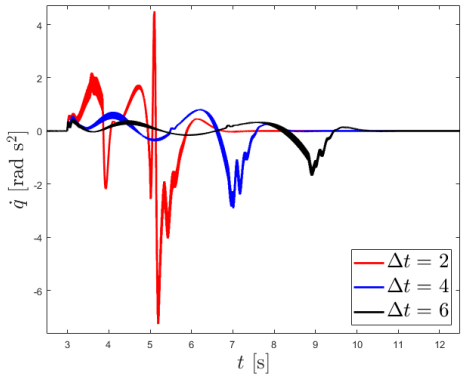
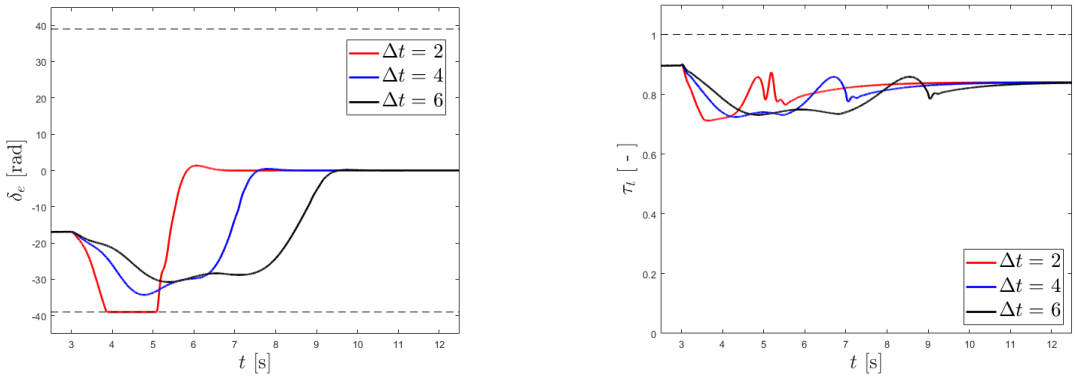


Figure 10: Pitch angular acceleration plot in time for the backward transition, for $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

On the other hand, the throttle input remains controlled in Fig.11b, despite still having the small oscillations that had also been noted for the forward transition.



(a) Elevator input δ_e for simulations with different Δt intervals.

(b) Throttle input τ_t for simulations with different Δt intervals.

Figure 11: Elevator (in the left) and throttle (in the right) plots in time for the forward transition, for $\Delta t = 2$s (in red), $\Delta t = 4$s (in blue) and $\Delta t = 6$s (in black).

The overall results for the backward transition are summarized in Tab.2, where the column corresponding to $\Delta t = 0.1$s is not filled due to the simulated crash.

Table 2: Summary of results for the backward transition.

	$\Delta t = 0.1$s	$\Delta t = 2$s	$\Delta t = 4$s	$\Delta t = 6$s
ΔD [m]	-	-8.068	-9.556	-11.308
$\overline{\dot{q}_{g_{max}}}$ [rad/s^2]	-	7.262	2.911	1.688
sat(δ_e)	-	yes	no	no
sat(τ_t)	-	no	no	no

4 CONCLUSIONS

This paper aimed to generate a trimmed flight envelope for a tailsitter UAV, covering airspeeds from 0 to 10 m/s and climb angles from 0 to 90 degrees, utilizing an iterative gradient descent optimization approach. Additionally, optimized transition profiles were developed to enable quasi-stable transitions between vertical and aerodynamic flight, aiming to enhance control authority and maneuvering precision. These objectives were successfully met, with the flight envelope generated and corresponding transition profiles having been derived.

The validation of these transition profiles was conducted in simulation using MATLAB/Simulink, where the impact of transition time on UAV stability, altitude gain and control effort was thoroughly analyzed. The results indicated that rapid transitions may lead to actuator saturation and even cause subsequent instability, contrasting with the simulations for longer transitions, which provide a more controlled response with reduced oscillations but at the cost of higher altitude variations. This analysis revealed the trade-offs between transition speed and system stability, offering valuable insights for fine-tuning tailsitter guidance and control methods in future experimental scenarios.

Future work is suggested to focus on extending the model to include descending vertical flight, enabling the creation of a trimmed envelope for negative climb angles. Additionally, an interesting aspect to develop would be to apply the method of envelope generation followed by the design of a suitable transition profile to other types of VTOL UAVs, although with the necessary adaptations, aiming to test the versatility of the approach. Lastly, when it comes to the performance during transitions, a comparative study between the developed transition profiles and existing methods for performing tailsitter transitions should be conducted, contributing to the optimization and refinement of transition strategies for improved UAV performance.

ACKNOWLEDGEMENTS

This work was supported by National Funds by FCT—Fundação para a Ciência e Tecnologia, I.P. under project LAETA Base Funding (DOI: 10.54499/UIDB/50022/2020) and under the scholarship with reference 2024.03868.BD.

REFERENCES

[1] A. Misra, S. Jayachandran, S. Kenche, A. Katoch, A. Suresh, E. Gundabattini, S. K. Selvaraj, and A. A. Legesse. A review on vertical take-off and landing (vtol) tiltrotor and tilt wing unmanned aerial vehicles (uavs). *Journal of Engineering*, 2022 (1):1803638, 2022.

[2] A. S. Saeed, A. B. Younes, C. Cai, and G. Cai. A survey of hybrid unmanned aerial vehicles. *Progress in Aerospace Sciences*, 98:91–105, 2018.

[3] G. J. Ducard and M. Allenspach. Review of designs and flight control techniques of hybrid and convertible vtol uavs. *Aerospace Science and Technology*, 118:107035, 2021.

[4] R. Chiappinelli and M. Nahon. Modeling and control of a tailsitter uav. In *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 400–409. IEEE, 2018.

[5] R. W. Beard and T. W. McLain. *Small unmanned aircraft: Theory and practice*. Princeton university press, 2012.

[6] A. D. Marco, E. Duke, and J. Berndt. A general solution to the aircraft trim problem. In *IAA Modeling and Simulation Technologies Conference and Exhibit*, 2007. doi:10.2514/6.2007-6703.

[7] A. Athayde, A. Moutinho, and J. R. Azinheira. Experimental nonlinear and incremental control stabilization of a tail-sitter uav with hardware-in-the-loop validation. *Robotics*, March 2024. doi:10.3390/robotics13030051.

[8] B. L. Stevens, F. L. Lewis, and E. N. Johnson. *Aircraft control and simulation: dynamics, controls design, and autonomous systems*. John Wiley & Sons, 2015.

[9] W. Khan and M. Nahon. Real-time modeling of agile fixed-wing uav aerodynamics. In *2015 international conference on unmanned aircraft systems (ICUAS)*, pages 1188–1195. IEEE, 2015.

[10] W. Khan and M. Nahon. Modeling dynamics of agile fixed-wing uavs for real-time applications. In *2016 international conference on unmanned aircraft systems (ICUAS)*, pages 1303–1312. IEEE, 2016.

[11] W. Khan and M. Nahon. Toward an accurate physics-based uav thruster model. *IEEE/ASME Transactions on Mechatronics*, 18(4):1269–1279, 2013.

[12] B. W. McCormick. *Aerodynamics, aeronautics, and flight mechanics*. John Wiley & Sons, 1994.

[13] A. C. Marta. Multidisciplinary design optimization of aircraft. `https://fenix.tecnico.ulisboa.pt/homepage/ist31052/optimizacao-multidisciplinar-de-aeronaves`, September 2024. PhD course notes, Intituto Superior Técnico, Universidade de Lisboa.

[14] C. G. BROYDEN. The convergence of a class of double-rank minimization algorithms 1. general considerations. *IMA Journal of Applied Mathematics*, 6(1):76–90, 03 1970. ISSN 0272-4960. doi:10.1093/imamat/6.1.76. URL `https://doi.org/10.1093/imamat/6.1.76`.

[15] R. Fletcher. A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317–322, 01 1970. ISSN 0010-4620. doi:10.1093/comjnl/13.3.317. URL `https://doi.org/10.1093/comjnl/13.3.317`.

[16] D. Goldfarb. A family of variable metric updates derived by variational means," mathematics of computing. *Mathematics of Computing*, 24:23–26, 1970.

[17] D. F. Shanno. Conditioning of quasi-newton methods for function minimization. *Mathematics of Computation*, 24(111):647–656, 1970. ISSN 00255718, 10886842. URL `http://www.jstor.org/stable/2004840`.

[18] J. R. Azinheira and A. Moutinho. Hover control of an uav with backstepping design including input saturations. *IEEE Transactions on Control Systems Technology*, 16 (3):517–526, 2008. doi:10.1109/TCST.2007.908209.

[19] E. J. Smeur, Q. Chu, and G. C. De Croon. Adaptive incremental nonlinear dynamic inversion for attitude control of micro air vehicles. *Journal of Guidance, Control, and Dynamics*, 39(3):450–461, 2016.

[20] P. Acquatella, E. van Kampen, Q. P. Chu, et al. Incremental backstepping for robust nonlinear flight control. *Proceedings of the EuroGNC*, 2013, 2013.

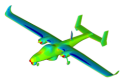
[21] A. Athayde, A. Moutinho, and J. R. Azinheira. Incremental nonlinear dynamics inversion and incremental backstepping: Experimental attitude control of a tail-sitter uav. *Robotics*, June 2024. doi:10.3390/act13060225.

[22] N. M. Guerreiro and A. Moutinho. Robust incremental backstepping controller for the attitude and airspeed tracking of a commercial airplane. In *2019 IEEE 10th International Conference on Mechanical and Aerospace Engineering (ICMAE)*, pages 607–611. IEEE, 2019.

[23] R. A. Cordeiro, A. S. Marton, J. R. Azinheira, J. R. Carvalho, and A. Moutinho. Increased robustness to delay in incremental controllers using input scaling gain. *IEEE Transactions on Aerospace and Electronic Systems*, 58(2):1199–1210, 2021.

[24] R. A. Cordeiro, J. R. Azinheira, and A. Moutinho. Cascaded incremental backstepping controller for the attitude tracking of fixed-wing aircraft. In *5th CEAS Conference on Guidance, Navigation and Control*, 2019.

[25] R. C. Nelson et al. *Flight stability and automatic control*, volume 2. WCB/McGraw Hill New York, 1998.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTIDISCIPLINARY DESIGN AND OPTIMIZATION RELATED TO NOISE AND VIBRATION: 30 YEARS OF RESEARCH AT NRC

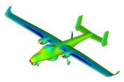
Anant Grewal

Flight Research Laboratory
National Research Council Canada
Ottawa, ON K1A 0R6, Canada
anant.grewal@nrc-cnrc.gc.ca, https://nrc.canada.ca/en/

Abstract. *The Aeroacoustics and Structural Dynamics Group at National Research Council Canada (NRC) has been engaged in noise and vibration research over the past five decades. Over the past 30 years, much of this work was performed collaboratively with partners from Industry, and Canadian government departments and agencies to develop solutions for noise and vibration issues plaguing operational aircraft. The challenge of addressing noise and vibration issues often begins with a choice of active control or passive system, and in the case of the former, designing optimal or near-optimal actuation, sensing and control systems that are capable of suppressing the undesired disturbances. At the outset three decades ago, a common approach to arriving at near-optimal solutions was the adoption of heuristic methods, where the application of engineering judgment, experience and intuition was the primary guiding principle. An example of this, namely the design of an actuation system based on segmented piezoelectric actuators for structural control of cabin noise due to propeller aircraft, is presented. A primary motivation in this case is the tailoring of modal control authority through distributed actuation. Cases of structural or acoustic sensing within both feedback and feedforward control architectures are explored. Significant noise and vibration reduction is achieved in all cases. Of note is the case when only structural sensing is employed whereby significant noise reduction is also obtained. Next, a more formal optimization approach using genetic algorithms to design optimal actuator configurations is presented. The problem addressed here involves placing elements into several actuator groups through combinatorial optimization to minimize the vibration and sound levels. The advantage of this approach is the possibility of identifying optimum configurations that heuristic approaches may miss altogether. The ability to reach optimum or near-optimum configurations by sampling only a small portion of design space is demonstrated. In the next step of the evolution of formal methods of multidisciplinary design optimization at NRC, the case of the multi-objective design optimization of the blade of a ducted fan for eVTOL application is explored. The design variables are the blade skew parameters, the blade angle, and the number of blades. Two separate sets of design objectives were investigated, the first whereby thrust is maximized, while torque is minimized, and the second case when far-field noise is minimized and aerodynamic efficiency is maximized. The workflow encompasses parametric geometry and mesh generation using Salome, CFD modelling with OpenFOAM, surrogate modelling using Kriging, and a Multi-Objective Genetic Algorithm. The Dakota open-source optimization platform is used for the surrogate modelling and optimization. The results yield*

a number of pareto-optimal solutions, which are then validated using the full-order solver. There is a significant improvement in aerodynamic efficiency (i.e. 3%) across a range of advance ratios, and a reduction in far-field noise of 10 dB at the target advance ratio. Finally, a similar multi-objective optimization framework is employed to study the use of acoustic metamaterials design using layered connected resonant cavities and Helmholtz resonators for broadband cabin noise reduction applications.

Keywords: Noise reduction, vibration suppression, active control, passive system, heuristic methods, multi-objective optimization



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MAGNETIC-ANOMALY POSITIONING FOR GNSS-INDEPENDENT UAV NAVIGATION

D. Pisarski, R. Faraj, Ł. Jankowski, B. Popławski, R. Konowrocki

Department of Intelligent Technologies
Institute of Fundamental Technological Research
Polish Academy of Sciences
Pawińskiego 5b, 01-106, Warsaw
dpisar@ippt.pan.pl, http://www.ippt.pan.pl

Abstract. *The paper introduces an original method for accurate positioning of an unmanned aerial vehicle (UAV) operating within a magnetic anomaly field. To estimate the UAV's position in GNSS-denied environments, the proposed positioning method relies on particle filtering with an enhanced velocity propagation model. This approach employs data fusion of kinematic measurements from the Inertial Measurement Unit (IMU) and magnetic field readings from a scalar magnetometer, which are correlated with a magnetic anomaly map. Notably, this method utilizes magnetic field measurements in two distinct ways. Firstly, the rate of change of the magnetic field along the UAV's trajectory is fed to the propagation model where the Bayesian inference is used to fuse the IMU and magnetometer measurements to estimate the particles velocity. Then, the absolute magnetic field values are compared with a magnetic anomaly map to recompute the particles' weights and refine the position estimate. The method was tested offline using in-flight recorded data. The analyses demonstrated the high performance of the particle filter with the enhanced velocity propagation model compared to dead reckoning navigation. As a result, the UAV's position error was significantly reduced. The study also includes a discussion on the method's robustness under different levels of IMU errors and the accuracy of magnetic field measurements.*

Keywords: magnetic anomaly navigation, Particle Filter, Bayesian inference, unmanned aerial vehicle, sensor fusion.

1 INTRODUCTION

A broad class of navigation systems has been developed to operate reliably in GNSS-denied environments, such as dense urban areas, tunnels, and confined spaces, while also offering protection against intentional disruptions like spoofing or jamming. One potential solution involves vision-based systems, which determine position by analyzing terrain features. These systems provide relatively high spatial resolution and are cost-effective in terms of hardware components, likely contributing to their current popularity as an alternative or supplement to Inertial Navigation Systems (INS) [1–4]. This is particularly evident in UAV technology, where encoders commonly used in robotics are not applicable, and traditional dead reckoning based on Inertial Measurement Unit (IMU) readings alone is insufficient. However, vision-based solutions cannot be applied in all possible environments. For instance, they struggle in environments such as sea areas, where reference points are absent, making visual analysis comparable to relying on inertial instrumentation. Similarly, their effectiveness diminishes over flat deserts or dense forests. Additionally, vision-based systems face several limitations, including high computational demands, reduced accuracy at higher flight altitudes, dependence on lighting conditions, and vulnerability to rapid motion or adverse environmental factors such as fog, rain, snow, or dust [5].

When vision-based navigation is not feasible, an attractive alternative is the implementation of a magnetometric measurement system, which utilizes crustal magnetic anomalies as a source of information correlated with an absolute coordinate system. Due to the significant variability of magnetic anomalies across most terrains, this approach provides a reliable reference for positioning, even in environments that appear uniform to camera-based systems. Furthermore, the compact and lightweight design of magnetometers makes them suitable for integration into ultralight UAV platforms. Numerous methods and techniques have been developed for magnetic positioning and navigation, each tailored to specific environmental characteristics. For indoor applications, Simultaneous Localization and Mapping (SLAM) plays a crucial role. Originally developed for robotic systems, SLAM involves both mapping the ambient magnetic field and positioning within the identified environment [6–8]. For precise estimation of an object's position, SLAM employs Kalman filters [9] or Particle Filters [10, 11] to efficiently fuse data from magnetometers and onboard inertial sensors. Such methods have been demonstrated in various studies. For example, in [12], an extended Kalman filter was used to fuse data from a magnetometer and an IMU installed on a ship model. Similarly, in [13], a gradient-based Kalman filter and a dynamic system were applied to reduce noise in magnetic field measurements, enhancing the performance of position estimation for a foot-mounted IMU-based system. Additionally, in [14], a sequential batch fusion algorithm was introduced for indoor magnetic map matching. Data fusion of IMU and magnetometer measurements was also explored in [15], where Bayesian inference was implemented to improve the computational efficiency of the positioning algorithm. Beyond indoor applications, outdoor implementations of magnetic anomaly-based navigation have also been extensively studied in the scientific literature. The airborne magnetic anomaly navigation described in [16, 17] successfully enabled the determination of an aircraft's absolute position. To achieve precise magnetic field measurements, scalar sensors are employed, and map-matching techniques are utilized to mitigate drift in inertial navigation systems. This approach has demonstrated the ability to achieve accuracy within tens of meters for flights lasting over an hour

[18]. Recent studies have also highlighted the effective application of machine learning for in-flight magnetic navigation [19].

This paper focuses on magnetic field-based navigation for aerial systems, employing the Particle Filter as the core methodology. While the Particle Filter is a well-established approach, the original contribution of this research lies in the enhanced velocity input, which is estimated using Bayesian inference. In the following sections, this concept is introduced, and its mathematical framework is discussed. Later in the paper, the method is analyzed using in-flight data collected from a Vertical Take-off and Landing (VTOL) platform equipped with a scalar magnetometer, conducting an autonomous flight over a sea area.

2 PARTICLE FILTERING WITH IMU VELOCITY ENHANCEMENT

2.1 System setting and assumptions

To estimate the position of the UAV, we developed a Particle Filter-based approach that leverages magnetic field measurements in two distinct ways. First, the magnetometer readings are compared with the magnetic anomaly map to aid in positioning. Second, and less intuitively, the magnetic field values are incorporated into the velocity estimation process, which also serves as an input to the Particle Filter (see Fig. 1). A detailed mathematical description of the method is presented in the following section. For now, we will focus on the primary objectives and assumptions. The first objective of our method is to achieve accurate estimation of the UAV's velocity by using Bayesian inference to integrate inertial measurements with variations in the magnetic field. The second objective is to compute the UAV's position within the Particle Filter framework, leveraging magnetic field values, the magnetic anomaly map, and a propagation model driven by the enhanced velocity input.

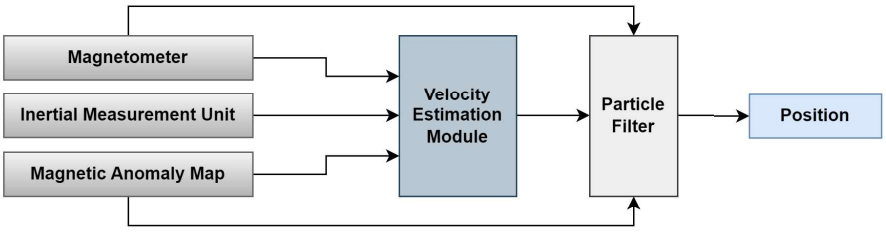


Figure 1: Schematic of the magnetic anomaly-based positioning system using a Particle Filter with enhanced velocity input.

To obtain desired operation of the proposed positioning method a number of assumptions must be satisfied:

- the drift of velocity estimated from the IMU is sufficiently small,
- the magnetic anomaly field is stationary (constant in time),
- the magnetic anomalies have been previously mapped with sufficient accuracy,
- the magnetic anomaly field possesses continuous spatial derivatives,
- the quality of the magnetometer measurements is sufficient for obtaining smooth time derivative of the magnetic anomaly field.

2.2 Particle Filter algorithm with enhanced IMU velocity integration

The Particle Filter iteratively estimates the distribution p of the state X_t at time t given the observation z_t:

$$p(X_t|z_t) = \sum_{i=1}^{n_t} w_t^i \, \delta(X_t - \mathbf{x}_t^i). \tag{1}$$

In Eq. (1), δ denotes the Dirac delta function, n_t represents the number of particles, while $\mathbf{x}_t^1, ..., \mathbf{x}_t^{n_t}$ and $w_t^1, ..., w_t^{n_t}$ correspond to the particle locations and their weights, respectively. In the standard implementation of the Particle Filter, the computation of the distribution $p(X_t|z_t)$ relies on the three major steps:

Step 1 Calculation of the weights w_t^i, $i = 1, ..., n_t$, which are proportional to the measurement likelihood $p(z_t|X_t)$.

Step 2 Perform re-sampling, where particles with lower significance are discarded, and new particles are generated based on the measurement likelihood $p(z_t|X_t)$.

Step 3 Update the positions of the particles for time $t + 1$ using a propagation model.

In our framework, the state X_t represents the position, while the observation z_t represents the measured magnetic field. The propagation model in Step 3 relies on the estimated velocity, which, in UAV applications, is generally obtained by integrating acceleration data from the IMU. However, as velocity estimates based solely on IMU data are prone to substantial drift, this study enhances the propagation model's performance by integrating a magnetic field correction into the velocity estimation process. In our method, the updates to the particle locations in Step 3 are calculated using a procedure akin to Euler's numerical integration scheme. While this first-order scheme is not particularly accurate, it allows for decoupling of the location update from the velocity update. Below, we outline the sub-algorithm designed for Step 3, which integrates the magnetic field-based correction into the IMU-derived velocity and executes the Particle Filter's propagation model.

Step 3.1 For the assumed time increment Δt update the particle location using the expected value of the velocity $\mathbf{V}_t$:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \Delta t \, \boldsymbol{\mu}_{\mathbf{V}_t}.$$

Step 3.2 Fuse the velocity information provided by $\bar{\mathbf{V}}_{t+1} = \mathbf{V}_t + \Delta \mathbf{V}_{t+1}^{\mathrm{IMU}}$, and $\mathbf{V}_{t+1}^{\mathrm{IMU}}$ into a single random variable $\mathbf{W}_{t+1}$. Under certain simplifying independence assumptions:

$$\mathbf{W}_{t+1} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\mathbf{W}_{t+1}}, \boldsymbol{\Sigma}_{\mathbf{W}_{t+1}}\right),$$

where

$$\boldsymbol{\Sigma}_{\mathbf{W}_{t+1}} = \left(\boldsymbol{\Sigma}_{\bar{\mathbf{V}}_{t+1}}^{-1} + \boldsymbol{\Sigma}_{\mathbf{V}_{t+1}^{\mathrm{IMU}}}^{-1}\right)^{-1},$$
$$\boldsymbol{\mu}_{\mathbf{W}_{t+1}} = \boldsymbol{\Sigma}_{\mathbf{W}_{t+1}} \left(\boldsymbol{\Sigma}_{\bar{\mathbf{V}}_{t+1}}^{-1} \boldsymbol{\mu}_{\bar{\mathbf{V}}_t t+1} + \boldsymbol{\Sigma}_{\mathbf{V}_{t+1}^{\mathrm{IMU}}}^{-1} \boldsymbol{\mu}_{\mathbf{V}_{t+1}^{\mathrm{IMU}}}\right).$$

Step 3.3 Fuse the information provided by $\mathbf{W}_{t+1}$ and the time derivative of the magnetic field, which is estimated during flight as $G_{t+1} \sim \mathcal{N}\left(g_{t+1}, \sigma_{\mathrm{G}}^2\right)$, into the final velocity

$\mathbf{V}_{t+1}$. Let the updated location be $\mathbf{x}_{t+1}$. Assuming the independence of $\mathbf{W}_{t+1}$ and G_{t+1}, the final velocity $\mathbf{V}_{t+1}$ is normally distributed:

$$\mathbf{V}_{t+1} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\mathbf{V}_{t+1}}, \boldsymbol{\Sigma}_{\mathbf{V}_{t+1}}\right),$$

where

$$\boldsymbol{\Sigma}_{\mathbf{V}_{t+1}} = \left(\boldsymbol{\Sigma}^{-1}_{\mathbf{W}_{t+1}} + \frac{\nabla M(\mathbf{x}_{t+1})\left(\nabla M(\mathbf{x}_{t+1})\right)^{\mathrm{T}}}{\sigma^2_{\mathrm{G}}}\right)^{-1}, \tag{2}$$

$$\boldsymbol{\mu}_{\mathbf{V}_{t+1}} = \boldsymbol{\mu}_{\mathbf{W}_{t+1}} + \frac{g_{t+1} - \boldsymbol{\mu}^{\mathrm{T}}_{\mathbf{W}_{t+1}} \nabla M(\mathbf{x}_{t+1})}{\sigma^2_{\mathrm{G}}} \boldsymbol{\Sigma}_{\mathbf{V}_{t+1}} \nabla M(\mathbf{x}_{t+1}). \tag{3}$$

The output data consist of the updated position $\mathbf{x}_{t+1}$ (computed in Step 3.1) and the corrected velocity $\mathbf{V}_{t+1}$ (determined in Step 3.3).

3 ALGORITHM VALIDATION UTILIZING UAV PLATFORM

The developed algorithm was tested using a specially designed platform – a VTOL (Vertical Take-Off and Landing) aircraft equipped with a set of high-precision scalar magnetometers (see Figure 2). At the initial stage, the magnetic measurement system was calibrated to mitigate disturbances caused by changes in aircraft orientation. To accomplish this, ground-based measurements were carried out using a non-magnetic tripod to gather the necessary data. A neural network was then employed to model the disturbance as a function of yaw, pitch, and roll angles, which was subsequently subtracted from the magnetic field measurements. Additionally, a ground station magnetometer system was utilized to monitor and account for magnetic field fluctuations associated with solar activity.



Figure 2: VTOL aircraft utilized for algorithm testing, featuring a magnetometer system with high-precision scalar cesium units (Geometrics MFAM) installed on the nose.

The calibrated magnetometer system was subsequently employed to generate a magnetic field map of the terrain spanning longitudes of 16.64–16.69 [deg] and latitudes of 54.565–54.586 [deg] (Figure 3). A series of carefully planned flight paths were executed to ensure comprehensive coverage of the selected area, achieving the desired precision and resolution. The magnetic field map was constructed using the collected data and Bayesian regression (see [15]). The average discrepancy between the map readings and the collected data points was 0.12 [nT].

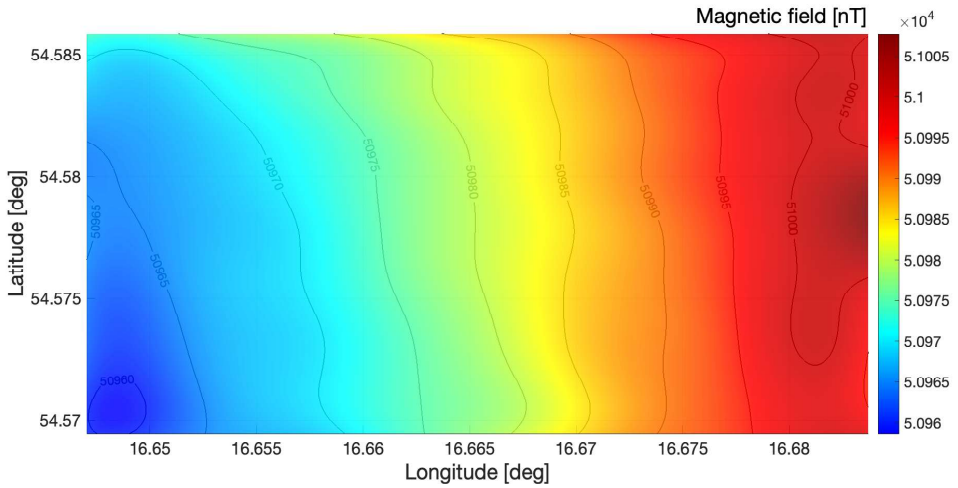


Figure 3: Magnetic field map generated from the collected data using Bayesian regression.

The test flight followed a predefined trajectory shaped by several waypoints, as illustrated in Fig. 6 (yellow curve). The total flight duration was 500 seconds. The IMU readings incorporated a deterministic velocity drift that increased linearly over time, resulting in a position error of 417.7 meters by the end of the mission. The Particle Filter was activated at $t = 50$ [s]. The algorithm (Steps 1–3) was initialized with 100 particles, distributed according to a normal distribution centered on the estimated starting position, with a standard deviation of 0.0001 [deg]. The Particle Filter's output was updated periodically with a time step of $\Delta t = 1$ [s]. At each update time $t \geq 50$, the weight calculation (Step A) was performed using a Gaussian function, where the argument was derived from the difference between the measured magnetic reading z_t and the corresponding value from the magnetic field map. During the resampling process (Step 2), particles with normalized weights below 0.9 were discarded. New particles were generated by sampling from a normal distribution centered on the remaining particles, with a standard deviation of 0.00003 [deg]. Additionally, the total number of particles was maintained within a range of 50 to 200. For the velocity correction in Step 3.3, the time derivative of the magnetic field in Eq. (3) was calculated using the increment of the magnetic readings, expressed as $g_{t+1} = (z_{t+1} - z_t)/\Delta t$. For the variance of the time derivative of the magnetic field measurement (Eq. (2) and (3)), we chose $\sigma_G^2 = 0.5$ [nT^2/s^2], which aligns with the magnetic field measurement error.

3.1 Analysis of results

To assess the positioning accuracy of the proposed method, position errors were analyzed along the assumed flight trajectory using the Particle Filter with both uncorrected and corrected IMU-based velocity. Each configuration was tested 100 times, and the results for uncorrected and corrected velocity are presented in Figures 4 and 5, respectively.

By analyzing the mean positioning errors, comparable performance is observed during the initial phase of the flight ($t < 150$ [s]), with errors for the uncorrected and corrected velocities similar. A noticeable divergence appears for $t > 250$ [s], where, in the case of uncorrected velocity, the mean position error steadily increases, reaching 80.7 [m] at time step $t = 486$ [s]. In contrast, the developed IMU-based velocity correction results

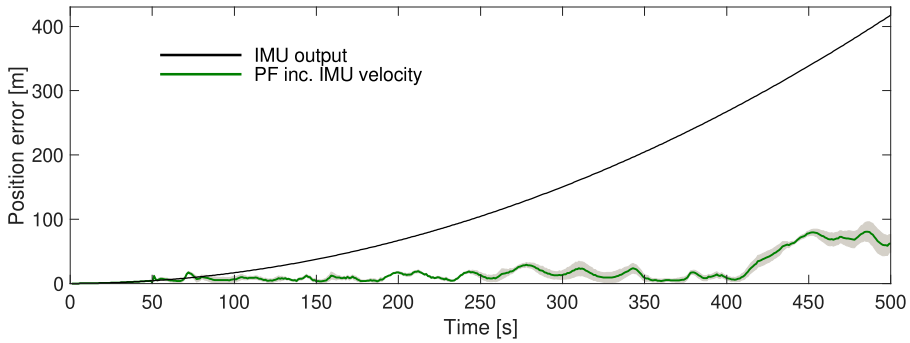


Figure 4: Positioning accuracy using Particle Filter (PF) with uncorrected IMU-based velocity. The green curve represents the mean position errors from 100 simulations, while the grey shaded area indicates the distribution within one standard deviation. The results are compared to the position error estimated exclusively by the IMU (black curve).

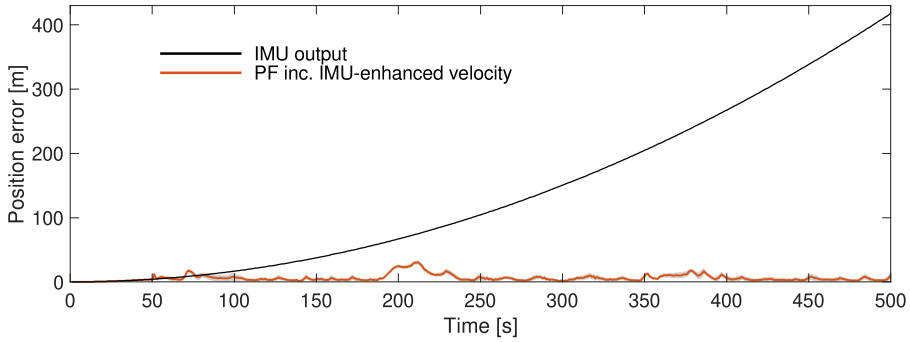


Figure 5: Positioning accuracy using Particle Filter with corrected IMU-based velocity. The orange curve represents the mean position errors from 100 simulations, while the grey shaded area indicates the distribution within one standard deviation. The results are compared to the position error estimated exclusively by the IMU (black curve).

in a substantial reduction in the mean position error. For $t > 215$ [s], the mean error stays below 20 meters. The peak position error occurs at time $t = 212$ [s], reaching 30.3 [m], which is 62.4% lower than the maximum error observed in the uncorrected velocity case. Additionally, the overall variation in position error is significantly reduced when the proposed velocity correction is applied. These observations are further supported by analyzing the trajectories generated by the evaluated positioning methods, as shown in Figures 6 and 7. Figure 8 illustrates the evolution of particle distributions at selected time steps, as generated by the Particle Filter with corrected velocity estimation. The characteristics of the distributions remain consistent throughout the entire flight, resulting in smooth positioning output, which can be advantageous for operating autopilot systems.

4 CONCLUSIONS

This study presented a novel method for localizing an object in GNSS-denied environments, utilizing magnetic anomaly measurements integrated into a Particle Filter at two levels. At the velocity estimation level, the time derivative of the magnetic field measurements is fused with velocity data from an Inertial Measurement Unit using Bayesian

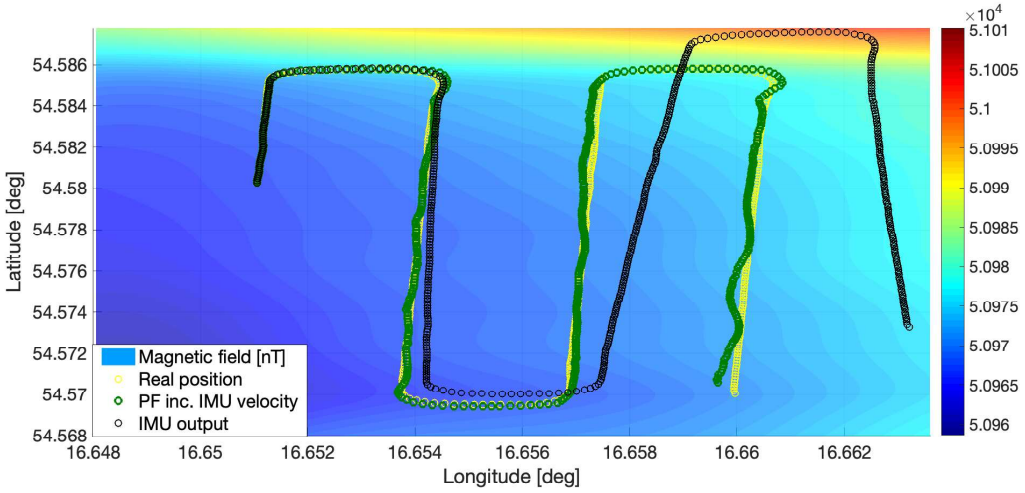


Figure 6: Comparison of trajectories generated by the evaluated positioning methods. For the Particle Filter utilizing uncorrected IMU-based velocity, the trajectory shown is derived from randomly selected simulations out of a total of 100 runs.

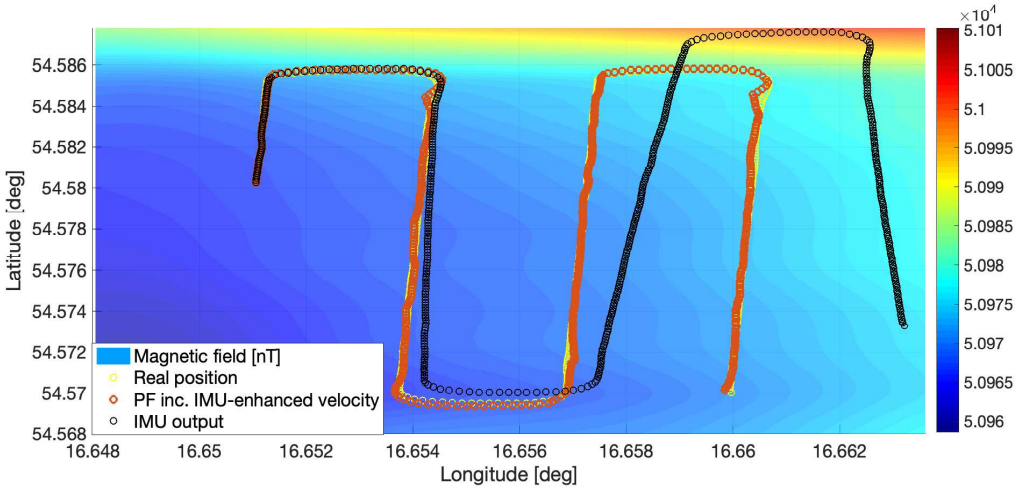


Figure 7: Comparison of trajectories generated by the evaluated positioning methods. For the Particle Filter utilizing corrected IMU-based velocity, the trajectory shown is derived from randomly selected simulations out of a total of 100 runs.

inference. At the position estimation level, the Particle Filter generates the posterior distribution by combining the measured magnetic field values with a reference magnetic anomaly map. The method was validated through a test flight using a custom-designed UAV platform equipped with high-precision cesium magnetometers. The results demonstrated that the Particle Filter, enhanced with the proposed magnetic field-based velocity correction, significantly improves positioning accuracy when compared to traditional IMU-based velocity approach.

ACKNOWLEDGEMENTS

This research was funded by the National Centre for Research and Development (NCBR), Poland, under the grant agreement DOB-SZAFIR/01/B/018/04/2021.

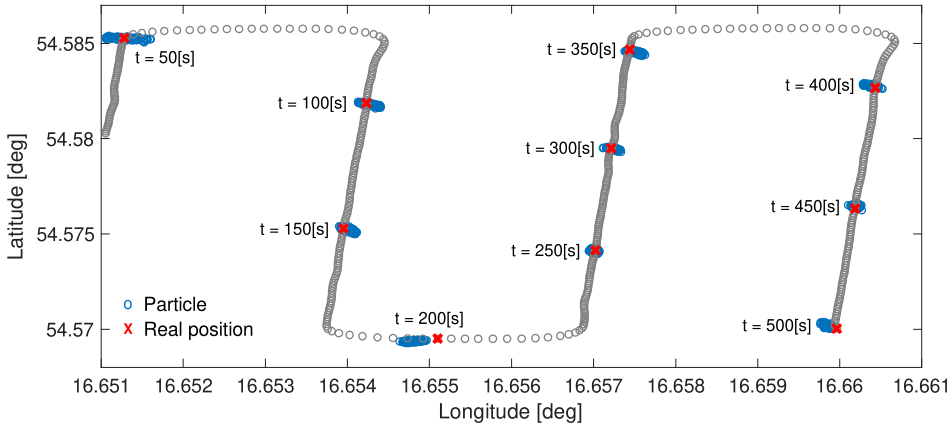


Figure 8: Particle distributions at selected time steps for corrected IMU-based velocity, with the reference trajectory depicted in the background (grey circles with red crosses indicating the real positions at the analyzed time steps). Results are based on randomly selected simulations from a set of 100.

REFERENCES

[1] G. Conte and P. Doherty. Vision-based unmanned aerial vehicle navigation using geo-referenced information, eurasip. *Journal on Advances in Signal Processing*, 2009 (1): 387308, 2009. doi:https://doi.org/10.1155/2009/387308.

[2] V. L. R. Mebarki and B. Siciliano. Nonlinear visual control of unmanned aerial vehicles in gps-denied environments. *IEEE Transactions on Robotics*, 31 (4):1004–1017, 2015. doi:https://doi.org/10.1109/TRO.2015.2451371.

[3] Y. Lu, Z. Xue, G. S. Xia, and L. Zhang. A survey on vision-based uav navigation. *Geo-Spatial Information*, 21(1):21–32, 2018. doi:https://doi.org/10.1080/10095020.2017.1420509.

[4] A. V. R. Katkuri, H. Madan, N. Khatri, A. S. H. Abdul-Qawy, and K. S. Patnaik. Autonomous uav navigation using deep learning-based computer vision frameworks: A systematic literature review. *Array*, 23:100361, 2024. doi:https://doi.org/10.1016/j.array.2024.100361.

[5] C. Zhu, M. Meurer, and C. Gunther. Integrity of visual navigation–developments, challenges, and prospects. *NAVIGATION: Journal of the Institute of Navigation*, 69(2):1–27, 2022. doi:https://doi.org/10.33012/navi.518.

[6] I. Vallivaara, J. Haverinen, and A. Kemppainen. Magnetic field-based slam method for solving the localization problem in mobile robot floor-cleaning task. *15th International Conference on Advanced Robotics (ICAR*, 21(1):198–203, 2011. doi:https://doi.org/10.1109/ICAR.2011.6088632.

[7] B.Li, T. Gallagher, A. Dempster, and C. Rizos. How feasible is the use of magnetic field alone for indoor positioning? *International Conference on Indoor Positioning and Indoor Navigation (IPIN), IEEE, Sydney, Australia*, pages 1–9, 2012. doi:https://doi.org/10.1109/IPIN.2012.6418880.

[8] G. Ouyang and K. Abed-Meraim. A survey of magnetic-field-based indoor localization. *Electronics*, 11(6):864, 2022. doi:https://doi.org/10.3390/electronics11060864.

[9] J. Aulinas, Y. R. Pétillot, J. Salvi, and X. Lladó. The SLAM problem: a survey, in: Artificial Intelligence Research and Development. *Frontiers in Artificial Intelligence and Applications*, pages 363–371, 2008. doi:https://doi.org/10.3233/978-1-58603-925-7-363.

[10] I. Vallivaara, J. Haverinen, A. Kemppainen, and J. Röning. Simultaneous localization and mapping using ambient magnetic field. *2010 IEEE Conference on Multisensor Fusion and Integration, IEEE, Salt Lake City, USA*, pages 14–19, 2010. doi:https://doi.org/10.1109/MFI.2010.5604465.

[11] X. Wang, C. Zhang, F. Liu, Y. Dong, and X. Xu. Exponentially weighted particle filter for simultaneous localization and mapping based on magnetic field measurements. *IEEE Transactions on Instrumentation and Measurement*, 66(7):1658–1667, 2017. doi:https://doi.org/10.1109/TIM.2017.2664538.

[12] F. Viseta, R. Helmons, and M. Kok. An extended kalman filter for magnetic field slam using gaussian process regression. *Sensors*, 22:2833, 2022. doi:https://doi.org/10.3390/s22082833.

[13] M. Zmitri, H. Fourati, and C. Prieur. Magnetic field gradient-based ekf for velocity estimation in indoor navigation. *Sensors*, 20:5726, 2020. doi:https://doi.org/10.3390/s20205726.

[14] Y. H. Kim, H. J. Kim, J. H. Lee, S. H. Kang, E. J. Kim, and J. W. Song. Sequential batch fusion magnetic anomaly navigation for a low-cost indoor mobile robot. *Measurement*, 213:112706, 2023. doi:https://doi.org/10.1016/j.measurement.2023.112706.

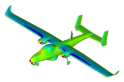
[15] Ł. Jankowski, D. Pisarski, R. Konowrocki, B. Popławski, and R. Faraj. Efficient real-time positioning using bayesian analysis and magnetic anomaly field. *Measurement*, 233:114738, 2024. doi:https://doi.org/10.1016/j.measurement.2024.114738.

[16] A. Canciani and J. Raquet. Absolute positioning using the Earth's magnetic anomaly field. *NAVIGATION: Journal of the Institute of Navigation*, 63(2):111–126, 2016. doi:https://doi.org/10.1002/navi.138.

[17] A. Canciani and J. Raquet. Airborne magnetic anomaly navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1):67–80, 2017. doi:https://doi.org/10.1109/TAES.2017.2649238.

[18] T. N. Lee and A. J. Canciani. MagSLAM: Aerial simultaneous localization and mapping using Earth's magnetic anomaly field. *Navigation*, 67(1):95–107, 2020. doi:https://doi.org/10.1002/navi.352.

[19] A. Gnadt. Machine learning-enhanced magnetic calibration for airborne magnetic anomaly navigation. *AIAA SCITECH 2022 Forum*. doi:https://doi.org/10.2514/6.2022-1760.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

ASSESSING WIND EFFECTS ON THE OPTIMAL TRAJECTORIES OF EXTENDED FORMATION FLIGHTS IN COMMERCIAL AVIATION

M. Cerezo-Magaña[1*], E. Staffetti[2], A. Olivares[2]

1: Department of Aerospace Engineering
Universidad Carlos III de Madrid
Leganés, Spain
2: Department of Signal Theory and Communicatios
Universidad Rey Juan Carlos
Fuenlabrada, Spain
*mcerezo@ing.uc3m.es, ernesto.staffetti@urjc.es, alberto.olivares@urjc.es

Abstract. *This study presents a qualitative analysis to evaluate the impact of wind on the design of formation flights. The formation mission design problem is formulated as follows: given several commercial flights, three in this study, along with weather forecasts and fuel-saving schemes for formation flight, the objective is to determine the optimal trajectories that minimize operating costs. To integrate wind effects into the problem formulation, the mission is modeled as a switched dynamical system, where aircraft can operate in two modes: solo flight and formation flight. Switching between these modes is governed by logical constraints in disjunctive form. To address the computational challenges, an embedding approach is employed to convert the logical constraints into smooth inequality and equality constraints, resulting in a conventional optimal control problem. This transformation eliminates the need for multiple phases or integer variables, thus avoiding combinatorial complexity and allowing the problem to be efficiently solved using traditional optimal control techniques. To analyze the influence of wind, various case studies were performed, covering different days throughout 2024. Wind data was obtained from the ECMWF Reanalysis v5- ERA5 dataset, ensuring accuracy and reliability. Initially, the algorithm's efficiency in solving the formation flight problem was validated using high-fidelity aircraft models and meteorological data. The results indicate that wind patterns substantially affect the optimal formation topology, influencing both formation flight benefits and configuration. Significant variations in optimal topology were observed, including shifts from three-aircraft formations to two-aircraft formations, with one aircraft flying solo throughout the journey. These findings highlight the importance of accounting for wind in the formulation of the problem, as wind significantly affects trajectory predictability and cost savings. Incorporating high-fidelity models is essential to accurately capture wind influence and maximize formation flight efficiency.*

Keywords: Optimal Control; Trajectory optimization; Formation flight; Wind effects; Switched dynamical system

1 INTRODUCTION

Aircraft formation flight has attracted increasing interest in recent years due to its potential to reduce fuel consumption and mitigate the environmental impact of aviation. By flying in coordinated formations, aircraft can take advantage of aerodynamic interactions to minimize induced drag, leading to significant fuel savings [1]. Formation flight can be categorized based on longitudinal spacing between aircraft [2]. Close formations involve separations of less than 10 wingspans between the leader and trailing aircraft, offering the greatest drag reduction and fuel efficiency benefits. Extended formations, with aircraft spaced between 10 and 40 wingspans apart, provide comparatively lower fuel savings but enhance operational safety, a key factor in commercial aviation. Due to these safety advantages, this study focuses on extended formations.

The formation flight planning problem is generally formulated as a two-stage optimization problem. As shown in Figure 1, the first stage involves the optimal partner allocation, which focuses on selecting flights that are most likely to form profitable formations in terms of overall Direct Operating Cost (DOC). The second stage addresses the mission design problem, which aims to determine the optimal trajectories for the selected formations. This paper concentrates solely on the formation mission design problem, which can be defined as follows: given a set of commercial flights, including their departure and arrival locations, departure schedules, weather forecasts, and the fuel savings scheme for formation flight, the goal is to determine how to group the flights, i.e. either in formations or solo, and find the optimal trajectories to minimize the overall cost.

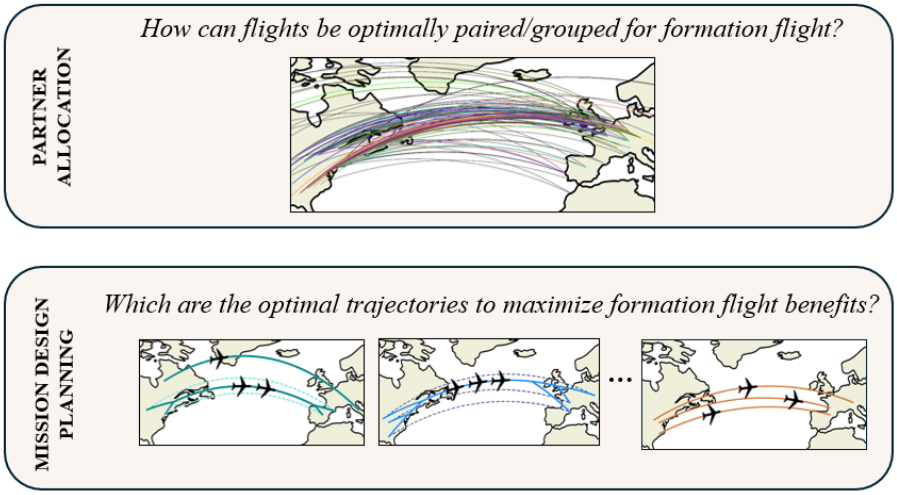


Figure 1: The two-stage formation planning problem approach, including the partner allocation and the mission design optimization problems.

Given the inherent combinatorial complexity of the formation planning problem, it has traditionally been addressed using highly simplified methods and approximations to make the problem more tractable. In contrast, the mission design problem has been studied using high-performance methodologies, allowing for a more detailed formulation while minimizing simplifications. In particular, in [3, 4], the mission design problem was approached by leveraging simplified models such as the Breguet range equation and the Fermat point extension, respectively. Alternatively, [5–7] utilized geometric approaches. These methods introduce significant approximations to streamline the partner allocation

process, thereby reducing the combinatorial complexity. However, such simplifications often exclude the use of detailed dynamic aircraft models and fail to incorporate essential meteorological data. Despite the reduced complexity, the implementation of precise dynamic models is critical to ensure feasible solutions, and integrating wind forecast data is vital for improving trajectory accuracy and optimizing fuel consumption predictions.

In last years, several studies have addressed the problem using higher-fidelity methodologies, in particular, following approaches base on optimal control [8–11]. In particular, in [8, 9], the mission design problem has been studied for two- and three-aircraft formations using a multi-phase optimal control approach. While this approach enables the use of accurate models, it requires the formation structure to be predefined. As a result, solving the mission design problem involves evaluating multiple combinations of solo and formation flights, solving the corresponding optimal control problems, and comparing their objective function values to determine the best solution. For example, a group of three flights already requires solving at least thirteen optimal control problems. As the number of flights increases, the number of possible combinations grows rapidly, making this method difficult to scale. A more general framework for designing formation missions by solving a single optimal control problem is presented in [10, 11]. These works employ the embedding approach to enhance scalability by eliminating the need for predefined switching assumptions. Unlike traditional methods, this approach does not require integer variables to model transitions between discrete states, and switching times between modes are determined implicitly rather than being treated as unknowns in the optimization process. As a result, computational efficiency is significantly improved.

The importance of incorporating detailed dynamic aircraft models and meteorological data has been demonstrated in these latter studies that avoid large simplifications. These works show that including wind data improves trajectory predictability, as evidenced by comparisons with wind-free scenarios. However, despite these advancements, a detailed assessment of how wind field variations impact formation flight has yet to be conducted. To what extent can wind affect formation flight? Does it only impact the fuel savings obtained through formation, or can it also influence the formation topology itself, for example? Is it possible that the same group of flights, under different wind conditions, could result in completely different formation patterns? Are there specific wind scenarios that significantly hinder formation benefits? These questions underline the need for a deeper investigation into the influence of wind on formation flight, aiming to optimize fuel efficiency and improve operational planning.

To answer these questions, in this study, we aim to address this gap by analyzing the impact of wind field variations on formation flight. Specifically, we assess how changes in wind conditions influence formation topology, including rendezvous and splitting locations, timing adjustments, and the overall fuel savings achieved. By conducting a detailed evaluation, we seek to provide a more comprehensive understanding of the role of wind in formation flight planning and optimization.

The paper is organized as follows. Section 2, firstly, describes the formulation of a Switched Optimal Control Problem (SOCP) for switched dynamical systems. Then, the methodology used to handle this problem is introduced tailored for the formation mission design problem. Section 3 presents and analyzes the results of the numerical experiments. Finally, Section 4 draws some conclusions.

2 METHODOLOGY

In this section, following [10], the general formulation of a Switched Optimal Control Problem (SOCP) for switched dynamical systems is introduced. The specific characteristics of the formation flight planning problem are then outlined to illustrate how the embedding approach is applied to address the SOCP in the context of this study. This approach enables the transformation of the original SOCP into an Embedded Optimal Control Problem (EOCP), resulting in a smooth Optimal Control Problem (OCP). The smooth formulation can then be tackled using traditional NonLinear Programming (NLP) techniques. A schematic overview of the methodology presented in this section is provided in Figure 2.

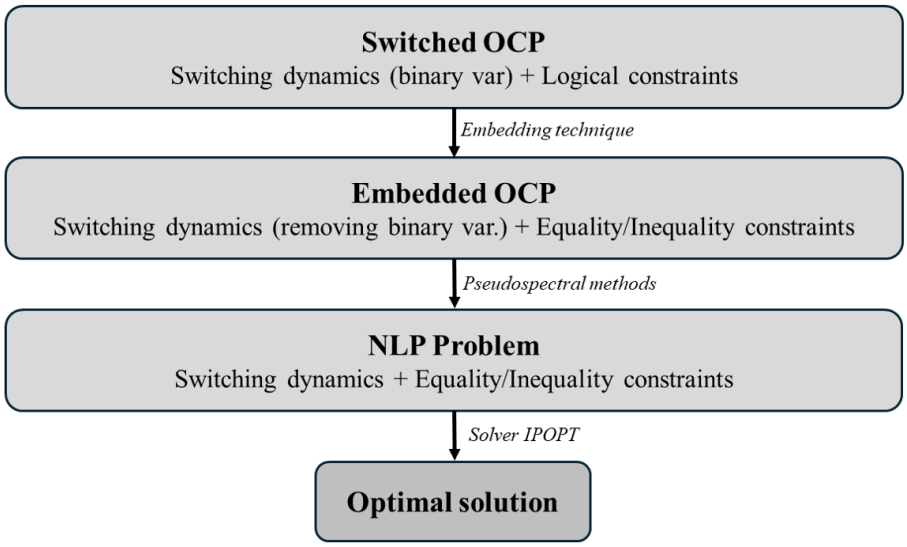


Figure 2: Schematic block diagram of the methodology used in this study.

2.1 The Switched Optimal Control Problem

Let $x_S(t)$ be the continuous state, $u_S(t)$, and $v_S(t)$, the discrete control input, namely, the switching control.

The dynamical model of a switched dynamical system is defined by

$$\dot{x}_S(t) \quad = \quad f_{v_S(t)}(t, x_S(t), u_S(t)), \tag{1}$$

where $f_{v_S(t)}(t, x_S(t), u_S(t))$ for each $v_S(t)$ is a piecewise $\mathcal{C}^1$ function, thus, the switching control $v_S(t)$ selects the function that defines the evolution of the continuous state. The boundary conditions, $x_S(t_I) = x_I$ and $x_S(t_F) = x_F$, are defined by the initial time t_I, the initial state $x_S(t_I)$, the final time t_F, and final state $x_S(t_F)$. Note that, a switching function could also be integrated into the cost functional J_S, in particular, in the Lagrange term. By incorporating the cost associated with transitions between discrete states within the integral term, the optimization framework captures both the running cost in each state and the penalties or benefits related to the switching behavior. This is not included in detail here, as it is not applicable to the specific problem of formation flight under study.

Therefore, the SOCP can be stated as follows: minimize the cost functional J_S subject to the switching dynamics given in Eq. (1), endpoint constraints, and path constraints, if applicable.

2.2 Specification of the formation flight problem

In the formation flight problem, aircraft are not allowed to fly in formation during take-off, climb, approach, or landing.Therefore, this study focuses solely on the cruise phase, where a two-degree-of-freedom dynamic model is used to capture the aircraft's flight behavior, accounting for mass variation over time [12]. The analysis is conducted within the horizontal plane, with the assumption of a constant cruise altitude.

Additionally, flight envelope constraints have been included to ensure aircraft's performance remains within safe operational limits, preventing risks associated with exceeding its design capabilities. These boundaries encompass various factors, including altitude, load factor, and airspeed, among others. The Base of Aircraft Data (BADA) of Eurocontrol, version 3.6 [13], has been used to include these constraints.

As said in the introduction, only extended formations have been considered in this study. Therefore, a longitudinal separation of more than 10 wingspans between aircraft have been set for safety reasons. Additionally, formation benefits are considered negligible for longitudinal separations of more than 20 wingspans. To simplify the analysis, the reduction in induced drag for trailing aircraft is directly modeled as a proportional reduction in fuel consumption, $\mathcal{R}_{\text{fuel}}$. Based on findings from key experimental studies on formation flight [14, 15], $\mathcal{R}_{\text{fuel}} = 0.1$ has been chosen.

Flying in formation requires deviations from the optimal flight paths of individual aircraft, leading to increased travel distances and extended flight duration. As a result, both fuel consumption and flight time should be incorporated into the objective function, with appropriate weighting factors to balance their contributions. Although the proposed methodology allows for the inclusion of such an objective function, this study focuses solely on minimizing fuel consumption. This approach is chosen to facilitate comparisons across various wind field scenarios and to evaluate results against low-fidelity models, where only fuel consumption can be optimized.

2.3 Conversion of the SOCP into an EOCP

This section outlines the process of converting the SOCP into an Embedded Optimal Control Problem (EOCP) using the embedding approach, specifically tailored for the formation flight problem. The goal is to transform the stated SOCP into a smoother optimal control problem OCP, facilitating the use of traditional control techniques to solve it.

The key idea behind the conversion is the introduction of continuous variables, $v_{Ei}(t) \in [0, 1]$, which replace the integer variables $v_S(t)$ typically used to switch between different dynamic states (e.g., solo flight or formation flight). This allows the model to represent the aircraft's behavior in formation or solo flight without the need for mixed integer programming.

The first step in the conversion is to introduce a set of continuous variables, $v_{Ei}(t)$, that always sum to exactly one. These variables represent the different possible dynamic behaviors for each aircraft. The new system of differential equations is then defined as a sum of the different dynamic models. The introduction of the continuous variables $v_{Ei}(t) \in [0, 1]$ to replace the integer variable $v_S(t)$ is the key aspect of the embedding approach.

Notice that, the obtained EOCP is a smooth OCP without integer variables. Therefore, classical techniques from optimal control theory can be applied to solve it, avoiding the need to use mixed integer programming methods.

2.4 Specification to the formation flight problem

In the particular study of the formation mission design problem examined in this paper, the only dynamic equation of the aircraft affected by switching dynamics is the mass flow rate equation. Specifically, the mass flow rate differs depending on whether the aircraft is flying solo, without any formation benefits, or participating in a formation as a trailing or intermediate aircraft, where fuel savings are achieved.

To model this switching behavior, the mass flow rate equation for an aircraft p when it does not obtain formation-related savings is combined with the equation representing the fuel consumption of an aircraft when it benefits from formation flight. This is achieved using the selection variable $v_{E_p}(t)$ for each aircraft p, as follows:

$$\dot{m}_p(t) = (1 - v_{E_p}(t))\,(-T_p(t) \cdot \eta_p(t)) + v_{E_p}(t) \cdot \left(-(1 - \mathcal{R}_{\text{fuel}_p}) T_p(t) \cdot \eta_p(t)\right), \tag{2}$$

The mass flow rate for each aircraft is represented as a function of $v_{Ei}(t)$, where:

- When $v_{Ep}(t) = 1$, aircraft p benefits from formation flight (i.e., as a trailing or intermediate aircraft),
- When $v_{Ep}(t) = 0$, aircraft p operates in solo flight or is the leader aircraft.

This approach allows the dynamic behavior of each aircraft to smoothly transition between different states without the need for binary switching, making the problem easier to solve with standard optimal control methods.

Additionally, in the formation planning problem, the switching control variables $v_{Ep}(t)$ are subject to continuous state-dependent constraints, which are based on the geographical positions of the aircraft. For example, formation flight is not feasible if the distance between two aircraft exceeds a certain threshold, such as 20 wingspans. These constraints, along with the switching dynamics ensure that the system behaves within acceptable limits.

The approach proposed in [16] has been employed to transform logical constraints in disjunctive form into inequality and equality constraints that involve only continuous auxiliary variables and, therefore, can be incorporated in an OCP.

Hence, at any time t, the distance $\mathcal{D}_{pq}(t)$ between aircraft p and q satisfies

$$\begin{aligned} 10b \leq \;& \mathcal{D}_{pq}(t) \leq 20b \qquad \text{(formation savings)}, \\ \text{or} \quad & \mathcal{D}_{pq}(t) \geq 20b \qquad \text{(no formation savings)}. \end{aligned} \tag{3}$$

The logical constraints for each pair of aircraft, p and q, in Eq.(3) can be expressed as

$$\begin{aligned} \forall p, q \in \{1, \ldots, N_a\},\ p < q: \quad & \alpha_{pq}(t)\,(\mathcal{D}_{pq}(t) - 20b) \leq 0, \\ \text{and} \quad & (1 - \alpha_{pq}(t))\,(20b - \mathcal{D}_{pq}(t)) \leq 0, \\ \text{and} \quad & 0 \leq \alpha_{pq}(t) \leq 1, \\ \text{and} \quad & \mathcal{D}_{pq}(t) \geq 10b. \end{aligned} \tag{4}$$

where $\alpha_{pq}(t) \in [0, 1]$. It is straightforward to observe that:

- When $\mathcal{D}_{pq}(t) \leq 20b \rightarrow \alpha_{pq}(t) = 1$, formation flight benefits will be realized.
- When $\mathcal{D}_{pq}(t) > 20b \rightarrow \alpha_{pq}(t) = 0$, no formation flight savings can be achieved.

It is important to note that the variables $v_{E_p}(t)$ and $\alpha_{pq}(t)$ have a similar nature. However, the former controls the switching behavior in the aircraft's dynamic model, while the latter is used to select between the inequality constraints added to the formulation. Specifically, as previously mentioned, the switching control variables $v_{E_p}(t)$ are subject to continuous state-dependent constraints, and this dependency is expressed by the relation $\alpha_{pq}(t) = v_{E_p}(t)$, which must be satisfied for each aircraft p.

2.5 Solution of the Switched Optimal Control Problem

As previously mentioned, after reformulating the SOCP into an EOCP and converting the logical constraints from disjunctive form into a set of inequality and equality constraints without integer variables through the embedding approach, the problem can be approached as a classical OCP. This transformation enables the use of conventional optimal control methods to solve it. In this work, the knotting pseudospectral method [17] is employed to convert the OCP into a NLP problem. The resulting NLP is then modeled with Pyomo [18] and solved using the IPOPT solver [19].

3 NUMERICAL RESULTS

In this section, several numerical experiments have been conducted to evaluate the impact of wind field variations on formation flight. For this purpose, the following North Atlantic flights have been selected:

Table 1: Latitudes and longitudes of the departure and arrival airports for the three flights.

Flight	Departure Airport	Arrival Airport
Flight A	Toronto (43.68, -79.63)	Madrid (40.49, -3.68)
Flight B	Atlanta (33.64, -84.42)	Paris (48.85, 2.35)
Flight C	NY JFK (40.64, -73.78)	Roma (41.77, 12.24)

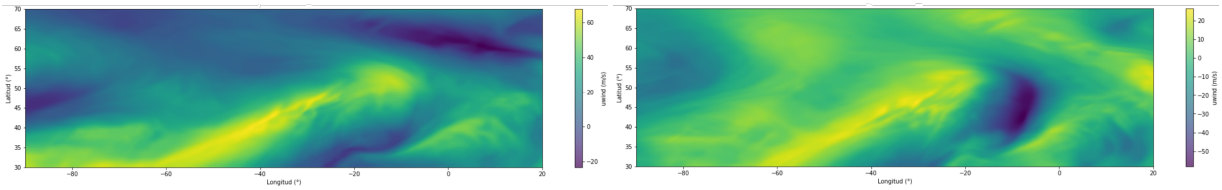
Note that, in this study, only the cruise phase is considered. Therefore, the initial and final locations of the cruise phase have been assumed to correspond to the latitudes and longitudes of the departure and arrival airports at cruise altitude, respectively. All other flight phases are completely disregarded.

For each flight, Airbus A330-200 aircraft models from the BADA dataset have been used. The initial mass of each aircraft is predefined, along with the initial and final velocities, which are set to typical cruise values for this aircraft type. Additionally, the initial heading angle is determined based on the great-circle route connecting the departure and arrival points. Due to the strong dependence of flight time on the wind field, where differences of more than two hours can occur when covering the same distance on a transoceanic flight, both the departure and arrival times have been left free. This allows the algorithm to determine the optimal initial and final times for each flight.

To enhance computational efficiency, initial guesses have been incorporated. Specifically, the latitude, longitude, and heading angle estimates are based on the great-circle route between each flight's departure and arrival points. Additionally, typical cruise speeds and fuel consumption rates for the selected aircraft model have been used to approximate velocity and mass variations throughout the flight.

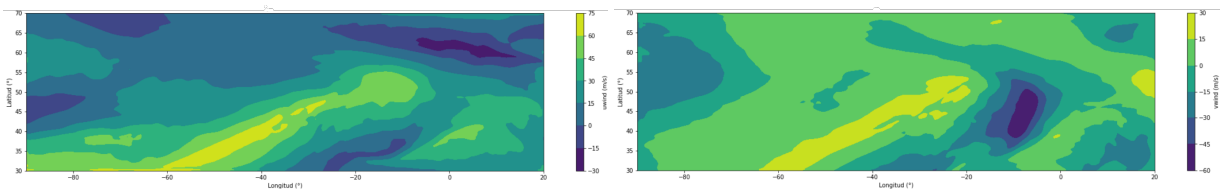
In this study, wind data obtained from the ECMWF Reanalysis v5 (ERA5) dataset [20] has been utilized to model the wind field variations influencing the formation flight. To integrate the wind effects into the aircraft's kinematics, we employed Radial Basis Functions (RBFs) as an interpolation method. RBFs are particularly suitable for approximating spatially varying fields, as they can provide smooth approximations of the wind field across the entire flight trajectory. The wind data, which includes both horizontal and vertical components, is used to modify the aircraft's kinematic equations by incorporating wind velocity as an external force acting on the aircraft during flight. This approach allows for the dynamic adjustment of the optimal flight paths, accounting for the time-varying wind field along the route. By using RBFs to approximate the wind field, we ensure a high degree of accuracy and computational efficiency, enabling the effective integration of wind effects into the optimal control problem for formation flight. This technique improves the prediction of aircraft trajectories and the calculation of fuel savings, offering a more realistic representation of flight conditions compared to models that neglect wind variations.

In Figure 3, the original U and V components obtained from ERA5 at 8:00 am on January 1, 2024 at 225 hPa of pressure level, are plotted. This day has been chosen as an example to plot the original wind components and the approximations made in this work. However, although it is shown for this specific day, the same procedure has been applied to each of the selected wind fields in the different scenarios. The U component represents the zonal wind velocity (east-west direction), while the V component represents the meridional wind velocity (north-south direction). Then, in Figure 4, both approximated wind components obtained using RBFs are shown to easily appreciate the accuracy of the approximation. Note that this accuracy could be improved by increasing the number of nodes considered in the approximation. However, in this work, due to the trade-off with computational cost, this level of accuracy has been considered sufficient, as the pattern and intensity of the wind field are clearly maintained.



(a) Original U-component. (b) Original V-component.

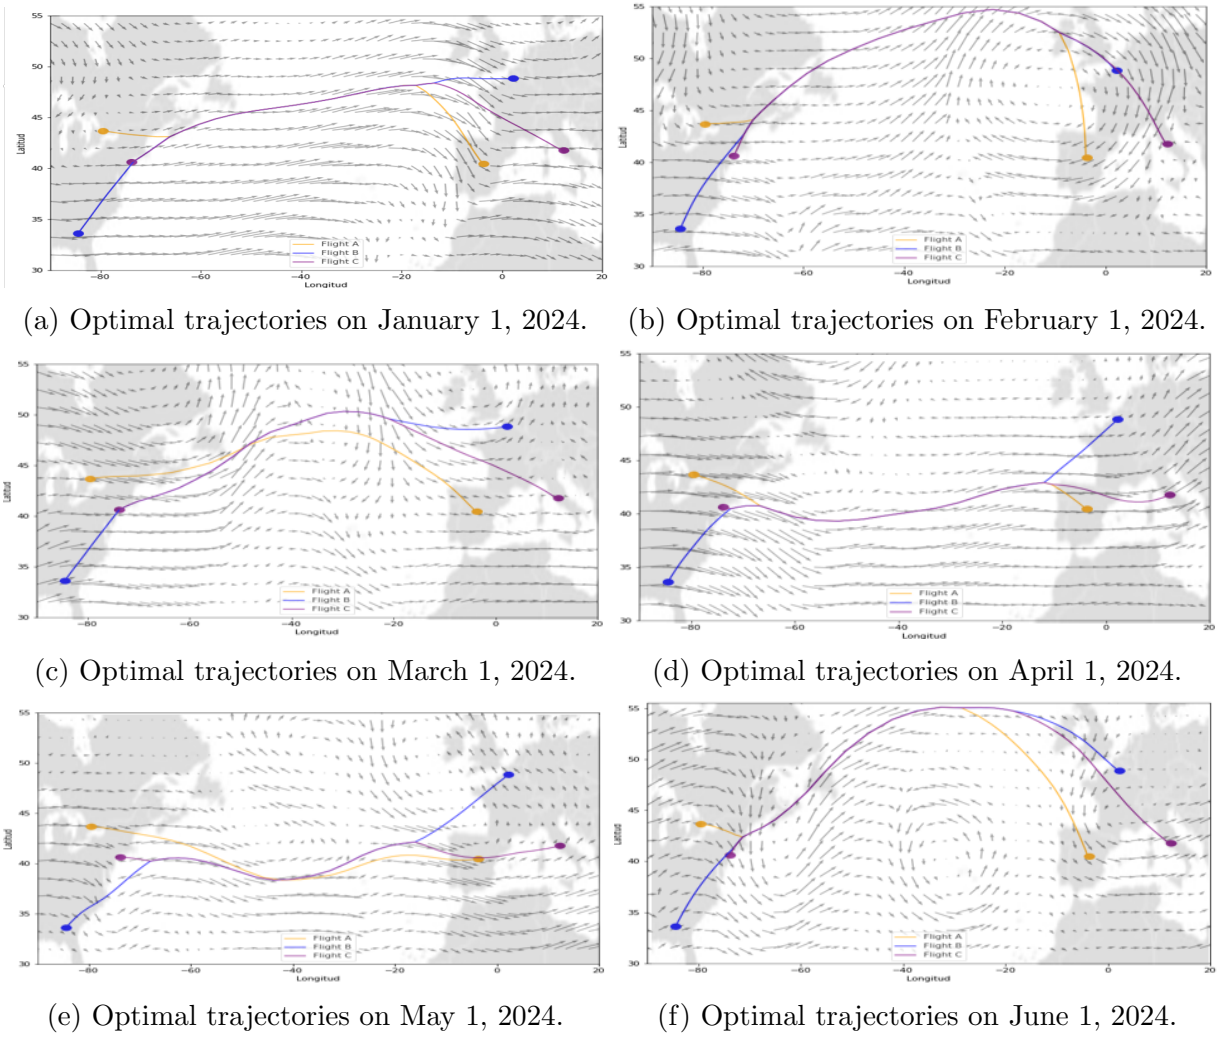
Figure 3: Original U and V components of the wind on January 1, 2024.



(a) U-component approximation. (b) V-component approximation.

Figure 4: Approximated U and V components of the wind on January 1, 2024, obtained by RBFs.

The same three flights have been analyzed under varying wind conditions throughout the year 2024. To present the results clearly, the optimal trajectories obtained for six selected days are shown, specifically, the first day of each month from January to June. Figure 5 shows the optimal trajectories calculated for each of these days.



(a) Optimal trajectories on January 1, 2024.

(b) Optimal trajectories on February 1, 2024.

(c) Optimal trajectories on March 1, 2024.

(d) Optimal trajectories on April 1, 2024.

(e) Optimal trajectories on May 1, 2024.

(f) Optimal trajectories on June 1, 2024.

Figure 5: Optimal trajectories obtained for the three aircraft, along with the wind field, for different days.

Figure 5 clearly highlights the differences in formation topology across the selected days. These differences are not only evident in the specific rendezvous and splitting locations and times but also in the overall structure of the formation itself.

On one hand, there are four days, i.e. January 1 (Figure 5a), February 1 (Figure 5b), April 1 (Figure 5d), and June 1 (Figure 5f), where the formation phases share a similar pattern. Initially, all aircraft fly solo, followed by a brief phase where two of them form a formation pair while the third remains solo. Subsequently, the three aircraft join together to form a formation for most of the ocean crossing. Later, one aircraft leaves the formation, while the remaining two continue flying in a two-aircraft formation, and finally, each aircraft resumes solo flight to their respective destinations. Although the specific trajectories differ, such as on January 1, 2024, where the paths do not exceed

a latitude of 50^o (Figure 5a), compared to February 1, where all three aircraft reach latitudes of 55^o (Figure 5b), the general structure of the formations remains consistent across these four days.

On the other hand, the trajectories for March 1 (Figure 5c) and May 1 (Figure 5e) display a completely different formation topology. In these cases, only a two-aircraft formation is formed, while Flight A remains solo for the entire duration of the flight. This significant difference in formation patterns highlights the variability in optimal flight configurations under changing wind conditions.

In addition to the trajectory and formation pattern changes, Table 2 presents the fuel consumption for the three flights when flying solo and when formation flight is allowed, across six different days. The table also shows the reduction in fuel consumption achieved when formation flight is utilized compared to solo flights.

Table 2: Fuel consumption comparison: solo flights vs. formation flights.

Day	Fuel consumption (solo flights) [kg]	Fuel consumption (formation flight allowed) [kg]	Reduction (%)
Jan 1	122722.48	119430.10	2.68
Feb 1	129425.41	125649.76	2.92
Mar 1	128882.16	126603.36	1.77
Apr 1	129025.36	124016.82	3.88
May 1	134369.82	132329.85	1.52
Jun 1	135887.92	131043.04	3.57

The results presented in Table 2 indicate that allowing formation flight consistently reduces fuel consumption across the selected days, with reductions ranging from approximately 1.52% to 3.88%. The mean reduction is around 2.72% with a standard deviation of 0.86%, i.e. $2.72 \pm 0.86\%$, highlighting the efficiency gains obtained from formation flight despite variations in wind conditions.

These findings demonstrate that even though formation patterns and trajectories may vary significantly depending on the day, formation flight generally leads to lower fuel consumption, emphasizing its potential for enhancing fuel efficiency in aviation.

Similarly, the formation flight planning problem for the three selected flights has also been analyzed without considering any wind field. In this case, the optimal control methods used are similar to those applied in the original problem. However, when no wind is taken into account, the problem can be more efficiently addressed using simplified low-fidelity methods, such as geometric approaches. The calculated fuel consumption for the no-wind scenario is 135,446.91 kg. Comparing this result with the optimal solutions obtained when including wind effects reveals an interesting trend. In both scenarios with winds, whether formation flight is allowed or solo flights are imposed, the total fuel consumption is generally lower than in the no-wind case. This is because the selected flights in this study are all eastbound, and the typical jet stream pattern over the North Atlantic region generally favors eastward flights. Therefore, the presence of wind typically results in reduced fuel consumption compared to the no-wind scenario. It is expected that if westbound flights (from Europe to the USA) had been considered, the opposite trend would likely have been observed, with wind conditions increasing fuel consumption.

However, there is a noteworthy exception observed for the trajectories obtained on June 1, 2024. In this case, the no-wind condition results in a fuel consumption comparable to or even slightly lower than the wind-influenced scenarios. This anomaly highlights the variability of wind effects and their potential impact on formation flight efficiency.

Table 3: Distances flown for the wind formation flight scenario for the three flights.

Day	Flight A [km]	Flight B [km]	Flight C [km]
Jan 1	6340.91	7188.29	6943.69
Feb 1	6669.74	7089.48	6982.02
Mar 1	6153.41	7166.72	6947.91
Apr 1	6513.97	7741.80	7288.20
May 1	6497.31	7730.56	7406.35
Jun 1	6568.83	7145.07	7032.27

Table 3 shows the distances flown by each of the three aircraft (Flight A, B, and C) on the selected days. The distances are measured in kilometers, and there is noticeable variability in the distances flown by each aircraft on different days. For instance, Flight B consistently flies the longest distance compared to Flights A and C, with a notable peak of 7741.80 km on April 1, which is nearly 700 km more than the distance flown on February 1, 2024. While the distances do not fluctuate drastically across the days, minor changes in the trajectories are expected as wind conditions influence the optimal flight paths. These variations demonstrate the role of wind in shaping flight trajectories, even if the overall trend remains relatively stable.

Table 4: Flight times for the wind formation flight scenario for the three flights.

Day	Flight A [h]	Flight B [h]	Flight C [h]
Jan 1	6.81	7.69	7.53
Feb 1	7.47	7.82	7.84
Mar 1	6.87	7.89	7.86
Apr 1	6.85	8.34	7.80
May 1	7.03	8.65	8.23
Jun 1	7.45	8.15	8.19

Finally, Table 4 shows the flight times for each of the three aircraft for the same days in the wind formation scenario. The flight times are provided in hours, and they show the time taken for each aircraft to complete the journey. Flight B has the longest flight times across all days, with a notable 8.65 hours on May 1. Flight A, on the other hand, generally takes less time to complete the flight compared to Flight B, while Flight C's flight times are in between those of Flight A and B. It can be observed that the flight times slightly changes during the different days, due to the differences on the wind field and on the flown distances. For instance, Flight C's time on May 1 increases to 8.23 hours compared to 7.53 hours on January 1.

Overall, tables 3 and 4 highlight the impact of different wind conditions on both the distance flown and the time required for each flight in the formation flight scenario, highlighting the importance of considering wind dynamics for accurate optimization in formation flights.

4 CONCLUSIONS

The study highlights the importance of accounting for wind conditions when planning formation flights, as wind variations significantly impact both optimal formation topologies and fuel consumption. The results demonstrate that optimal flight configurations are highly sensitive to changing wind patterns, underscoring the need for high-fidelity models that accurately capture wind effects. By incorporating wind data into the problem formulation, it is possible to achieve more efficient flight trajectories and significant fuel savings, as demonstrated by the reduction of approximately $2.72\% \pm 0.86\%$ when formation flight is allowed. Consequently, the integration of high-fidelity models becomes essential for accurately optimizing formation flight strategies,

The study highlights the importance of accounting for wind conditions when planning formation flights. The numerical experiments conducted reveal that wind variations significantly affect both the optimal formation pattern and the resulting fuel consumption. The differences observed in formation patterns across various days demonstrate that optimal flight configurations are highly sensitive to wind changes, highlighting the complexity of formation flight optimization in real-world scenarios. In particular, the results show that accounting for wind effects, formation flight can lead to consistently lower fuel consumption compared to solo flights, with an average reduction of approximately $2.72\% \pm 0.86\%$ when formation flight is allowed.

This reduction is especially relevant for transoceanic flights, where wind fields, such as the jet stream over the North Atlantic, can have a substantial impact on flight efficiency. Additionally, the analysis shows that the specific formation topology may vary depending on the wind conditions of each day, influencing not only the rendezvous and splitting points but also the overall structure of the formation, the distances flown and the flight times.

Therefore, the use of high-fidelity models that accurately capture wind effects is essential to develop effective and fuel-efficient formation flight strategies. Simplified low-fidelity methods may fail to account for these critical variations, leading to suboptimal solutions. Consequently, the integration of high-fidelity wind models into the problem formulation is fundamental to accurately optimize formation flight, ultimately contributing to more sustainable and cost-effective aviation operations.

ACKNOWLEDGEMENTS

This research was partially funded by Universidad Carlos III de Madrid through the grant “Ayudas para la Actividad Investigadora de los Jóvenes Doctores, del Programa Propio de Investigación de la UC3M.”

REFERENCES

[1] A. Ning, T. C. Flanzer, and I. M. Kroo. Aerodynamic performance of extended formation flight. *Journal of aircraft*, 48(3):855–865, 2011.

[2] G. J. Durango, C. Lawson, and A. Z. Shahneh. Formation flight investigation for highly efficient future civil transport aircraft. *The Aeronautical Journal*, 120(1229): 1081–1100, 2016.

[3] J. Xu, S. A. Ning, G. Bower, and I. Kroo. Aircraft route optimization for formation flight. *Journal of Aircraft*, 51(2):490–501, 2014.

[4] T. E. Kent and A. G. Richards. Analytic approach to optimal routing for commercial formation flight. *Journal of Guidance, Control, and Dynamics*, 38(10):1872–1884, 2015.

[5] L. Benvenuti and A. de Santis. A design methodology for commercial aircraft formation flight plans with minimal fuel consumption. *IEEE Transactions on Aerospace and Electronic Systems*, 60(4):4157–4169, 2024.

[6] W. B. Blake and T. C. Flanzer. Optimal routing for drag-reducing formation flight: A restricted case. *Journal of Guidance, Control, and Dynamics*, 39(1):173–176, 2016.

[7] G. C. Bower, T. C. Flanzer, and I. M. Kroo. Formation geometries and route optimization for commercial formation flight. In *Proceedings of the 27th AIAA Applied Aerodynamics Conference*, San Antonio, TX, USA, 2009.

[8] S. Hartjes, M. E. van Hellenberg Hubar, and H. G. Visser. Multiple-phase trajectory optimization for formation flight in civil aviation. In *Advances in Aerospace Guidance, Navigation and Control*, pages 389–405. Springer, 2018.

[9] S. Hartjes, H. G. Visser, and M. E. van Hellenberg Hubar. Trajectory optimization of extended formation flights for commercial aviation. *Aerospace*, 6(9):100, 2019.

[10] M. Cerezo-Magaña, A. Olivares, and E. Staffetti. Formation mission design for commercial aircraft using switched optimal control techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 57(4):2540–2557, 2021.

[11] M. Cerezo-Magaña, A. Olivares, and E. Staffetti. A stochastic switched optimal control approach to formation mission design for commercial aircraft. *IEEE Transactions on Aerospace and Electronic Systems*, 58(5):4342–4360, 2022.

[12] D. G. Hull. *Fundamentals of Airplane Flight Mechanics*. Springer, Berlin Heidelberg New York, 2007.

[13] Eurocontrol. User manual for the Base of Aircraft Data (BADA). Technical Report 130416, Eurocontrol, 2013.

[14] S. R. Bieniawski, R. W. Clark, S. E. Rosenzweig, and W. E. Blake. Summary of flight testing and results for the formation flight for aerodynamic benefit program. In *Proceedings of the 52nd Aerospace Sciences Meeting, AIAA SciTech Forum*, National Harbor, MD, USA, 2014.

[15] T. C. Flanzer, S. R. Bieniawski, and W. B. Blake. Operational analysis for the formation flight for aerodynamic benefit program. In *Proceedings of the 52nd Aerospace Sciences Meeting, AIAA SciTech Forum*, National Harbor, MD, USA, 2014.

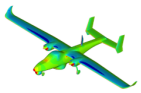
[16] S. Wei, M. Zěfran, and R. A. DeCarlo. Optimal control of robotic systems with logical constraints: Application to UAV path planning. In *Proceedings of the 2008 IEEE International Conference on Robotics and Automation*, Pasadena, CA, USA, 2008.

[17] I. M. Ross and F. Fahroo. Pseudospectral knotting methods for solving optimal control problems. *Journal of Guidance, Control, and Dynamics*, 27(3):397–405, 2004.

[18] W. E. Hart, C. D. Laird, J.-P. Watson, D. L. Woodruff, G. A. Hackebeil, B. L. Nicholson, and J. D. Siirola. *Pyomo-Optimization Modeling in Python*. Springer, 2012.

[19] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

[20] H. Hersbach, B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, et al. ERA5 hourly data on single levels from 1940 to present, Copernicus Climate Change Service (C3S) Climate Data Store (CDS)[data set], Accessed on 02-04-2025.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

OBSTACLE AVOIDANCE FOR FIXED-WING AIRCRAFT BY USING CONTROL BARRIER FUNCTIONS

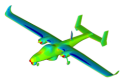
Huriye Nur Toprak[1]*, Emre Koyuncu[2]

1: Aviation Institute
Istanbul Technical University
toprakh16@itu.edu.tr

2: Aviation Institute
Istanbul Technical University
emrekoyuncu@itu.edu.tr

Abstract *With the improvements in aviation industry and increased manufacturing capabilities, aircraft are now being deployed in more complex environments and mission profiles. However, ensuring autonomous flight safety is still a challenging problem, particularly in environments with both static and dynamic obstacles. Among the various methods for ensuring safe operation, Control Barrier Functions (CBF) have emerged as a promising approach due to its ability to enforce system constraints both geometrically and dynamically, providing a robust framework for safe navigation. CBFs have been widely applied in robotics and autonomous vehicle systems to enhance obstacle avoidance. With the help of the geometric representation of the potential collision region, CBF can be defined. While there are many studies on CBF have been successfully applied to unmanned ground vehicles and quadrotors, its adaptation to fixed-wing aircraft remains challenging due to their nonholonomic motion constraints and highly nonlinear dynamic characteristics. This study aims to extend the CBF based obstacle avoidance approach to fixed-wing aircraft operating in environments containing both static and dynamic obstacles. The methodology involves dynamically constructing a collision region based on the aircraft's relative velocity with respect to obstacles and incorporating it into a CBF framework for real-time avoidance. To accommodate the specific dynamics and constraints of fixed-wing aircraft appropriate to CBF, an affine control system representation is considered. To evaluate the minimum control inputs satisfying the obstacle avoidance, CBF Quadratic Program (CBF-QP) is defined and solved by numerical optimization. Furthermore, the environment is modeled using Digital Terrain Elevation Data (DTED) to accurately represent real-world navigation conditions. By integrating CBF with fixed-wing aircraft dynamics, this study extends the usage of CBF based obstacle avoidance beyond ground vehicles and quadrotors, shows a feasible application for safe fixed-wing aircraft navigation in environments with obstacles.*

Keywords: CBF, CBF-QP, obstacle avoidance, fixed-wing aircraft



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

EXPLORING COUPLING EFFECTS BETWEEN AIRFRAME AND PROPULSION IN PARAMETRIC MODELS FOR ELECTRIC REGIONAL AIRCRAFT WITH RANGE EXTENDER

Didier Brousset-Matheu[1*2*3], Eric Nguyen Van [1*2], Sebastien Defoort [1*2], Xavier Roboam[3] and Bruno Sareni[3]

1: DTIS, ONERA,
Université de Toulouse
31400, Toulouse, France
2: Fédération ENAC ISAE-SUPAERO ONERA
Université de Toulouse
31400, Toulouse, France
{didier.brousset-matheu, eric.nguyen, sebastien.defoort}@onera.fr
3: Laplace (UMR CNRS-Toulouse INP-UPS)
Université de Toulouse
ENSEEIHT 2 Rue Camichel, 31071 Toulouse, France
{xavier.roboam, bruno.sareni}@laplace.univ-tlse.fr

Abstract.

Recently, the Hybrid Aircraft: academic reSearch on Thermal and Electrical Components and Systems (HASTECs) project developed an advanced serial hybrid electric propulsion system, for regional transport aircraft, with a parametric model in each propulsion component. This approach integrated main physical phenomena covering several fields (electrical, mechanical, thermal) for sizing and evaluation of device losses and weights. The study highlighted the necessity for propulsive system optimization rather than component optimization, intending to achieve an optimal balance between efficiency and mass. However, the project decoupled the airframe from the propulsion system, resulting in the inability to ascertain if this compromise was reflected in the overall aircraft optimization. Concerning the Overall Aircraft Design (OAD) front, the FAST-OAD software has been used in several studies to analyze the propulsion-airframe integration but it relied on low-fidelity models for the hybrid propulsion system. This approach introduces significant uncertainties in the results, especially due to the lack of consideration for certain electrical phenomena, such as partial discharges, making it difficult to satisfy design constraints.

The objectives of this study are twofold: first, to present a new OAD process by coupling HASTECs and FAST-OAD software; second, to structure the oad process with the integration of different topologies of aircraft (turboelectric, serial-hybrid, full electric), the number of motors, and a method to determine the maximum possible range in pure electric mode. Finally, the objective is to ensure the robustness and relevance of the final OAD results despite strong interactions with an acceptable computation cost facing the system complexity for further optimizations. The case study for the new design process is the design exploration of an electric regional transport aircraft with a range extender.

Keywords: Hybrid, aircraft, Electric, powertrain, system, multidisciplinary, design, optimization

1 INTRODUCTION

In 2021, the International Air Transport Association (IATA) proposed a plan [1] to reach net zero emissions in 2050. A possible solution is to use new zero CO_2 emission energy sources, such as electric power, but its low specific energy density does not provide sufficient range for medium and long haul [2], As a consequence, the research focuses on the regional sector.

In the context of regional aviation, particularly with the ATR72, the design range is approximately 750NM (1390km) [3]. The state-of-the-art in battery technology does not currently permit, with only the use of battery, the achievement of the ATR72's maximal range. Consequently, numerous studies have been conducted on integrating battery packs with conventional turboshafts to create hybrid aircraft [4–7] with various configurations such as parallel hybrid, serial hybrid, etc.

These novel energy systems deviate from conventional aircraft design, necessitating a re-evaluation of optimal architecture. The overall aircraft design aims to identify the most efficient solution. In 2023, Novelli et al [8], drew up a state-of-the-art report on hybrid electric architectures. Although the parallel hybrid allows a simpler integration in terms of structural parts, the results show interest only for a very short range. The serial hybrid configuration, whilst necessitating greater structural modification, has the potential to reduce energy consumption. The modeling employed in this report and the cited papers are based on a low-fidelity propulsion model developed using specific energy and power density, introducing considerable uncertainty into the output process.

In order to facilitate the study of novel aircraft concepts during the preliminary design phase, Onera and Supaero developed a software program designated FAST-OAD [9]. The software is based on OpenMDAO [10], which is an open-source framework for multidisciplinary design, analysis, and optimization (MDAO). It provides a modular structure for defining computational models with efficient derivative computation. The FAST-OAD facilitates the adaptation of the modular structure of OpenMDAO for the aircraft preliminary design, incorporating specific models such as drag polar computation, temporal mission simulation, fuselage, and wing geometric models, among others.

Between 2018 and 2021, the HASTECS project [11] optimized an electric hybrid propulsion powertrain with physical parametric models for electric motors, power converters, etc. These models were integrated into a global powertrain system, from the propeller to the battery pack. This approach has integrated the consideration of diverse physical phenomena (electric, mechanical, thermal) such as voltage source inverter cooling, motor cooling, and partial discharges due to high voltage. The aircraft system part was decoupled from the hybrid system part by a linear equation (equation 1), based on reference thrust and mass extracted from ATR72 data. This equation integrates, with simplification, the "snowball effect" created by the structure and powertrain coupling implying that a higher powertrain increased the weight of the aircraft, which in turn increases the weight of the powertrain, resulting in higher thrust and therefore a higher propulsion system mass, etc. The aforementioned equation was replaced by a supplementary design variable ($RThrust$) and an additional constraint linked to final mass, resulting in equations 2 and 3.

$$T_{new} = T_{ref} * \frac{MTOW_{new}}{MTOW_{ref}} \tag{1}$$

$$T_{new} = T_{ref} * RThrust \tag{2}$$

$$MTOW_{new} \leq RThrust * MTOW_{ref} \tag{3}$$

with T_i the aircraft thrust, $MTOW_i$ the maximum take off weight and $RThrust$ the new design variable.

During the optimization process, the optimizer has converged on a compromise between efficiency and component mass, demonstrating an interest in global optimization rather than component optimization. Unfortunately, uncoupling the aircraft system and hybrid system cannot confirm if the compromise found by the optimizer was reflected in the overall aircraft design.

Consequently, there is a significant interest in integrating models of the electric propulsion chain from the HASTECS project within the FAST-OAD environment. This will lead to a multidisciplinary design capable of integrating electrical, aerodynamic, and structural couplings with greater fidelity.

Although the ATR72 has a range of 750NMI, 91% of the flights made by turboprops are less than 435NMI (804 km) [12]. The idea is here to develop a series hybrid aircraft capable of reaching the 750NMI but also capable of reaching a certain distance with only electric battery power (related to the so-called "pure electric mission"). This solution authorizes the flexible mission of an actual ATR72 but decarbonizes a total segment of regional transport. Examples of such concepts are presented in the IMOTHEP project [4] and company Elysian which claims 200 NM in full electric mode with a cell density of 300 Wh/kg.

The paper employs a structured methodology to address the challenges of multidisciplinary aircraft design. The paper's methodology is initiated with the Hypothesis and Multidisciplinary Design Analysis, which defines the core assumptions and the multidisciplinary design analysis. Subsequently, the MDA Subgroups and Models are explained in detail, highlighting how different disciplinary models interact within the design process. Finally, the Optimization Problem and Design of Experiment are presented, describing the optimization problem, conducting a design space exploration, and evaluating tendencies.

2 HYPOTHESIS AND MULTIDISCIPLINARY DESIGN ANALYSIS

The incorporation of the HASTECS powertrain model within the FAST-OAD software necessitates a series of modifications with respect to initial models and hypotheses.

2.1 Hypothesis

The aircraft under consideration involves a series hybrid architecture, equipped with identical electric motors for each propeller, ensuring independence in case of motor failure. Each powertrain is composed by a propeller, a gearbox, a motor, and an inverter, all located in a nacelle that will be referred to as a powertrain nacelle. On each side, the powertrain nacelle nearest to the fuselage will be grouped with a turbogenerator in a bigger nacelle, termed as a "POD". At the rear, it is possible to have two turbogenerators positioned inside the fuselage in a nacelle, which we designate as an "energy nacelle". A maximum of four turbogenerators is assumed, distributed either under the wing or at the rear (figure 1). The battery packs can be arranged in three positions: inside the wing, in the fuselage between the cockpit and the passengers, and inside the POD. At the

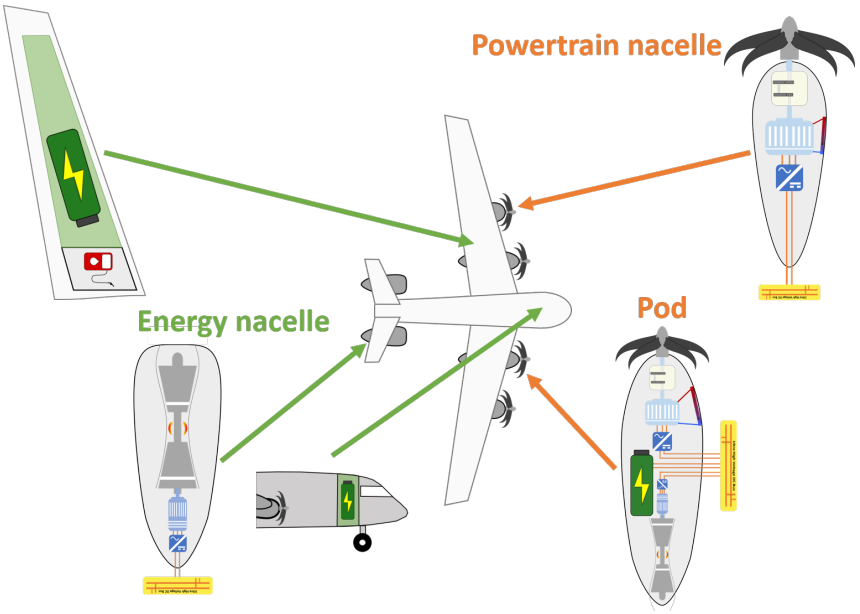


Figure 1: Placement of each energy source system in the hybrid aircraft

moment there is no distinction between placing the battery in POD or nacelle. Therefore for simplicity, the battery is considered in the pod only. All cables are connected to a central point (i.e. the "power center") located in the aircraft center between the wing. Although this hypothesis does not necessarily constitute the most optimal solution for cable management, it does offer greater flexibility in the event of multiple topologies.

2.2 Multidisciplinary Design Analysis

The Multidisciplinary Design Analysis (MDA) involves the design of a hybrid aircraft, which is equipped with a turbogenerator and battery sources for 72 passengers, 6840 kg of payload, and a mission with 750 nautical miles (NMI) plus a mission of diversion. The MDA gives a sized converged aircraft that respects the Top Level Aircraft Requirements depending on the design variables. Furthermore, it calculates the aircraft's range when operating on battery power alone, called the "pure electric range".

The MDA is structured using the eXtended Design Structure Matrix (XDSM), as illustrated in figure 2. This "top-level XDSM" offers a synthetic but global vision of the design at the aircraft level while lower (more detailed) levels of XDSM are provided in section 3. The inputs are represented vertically, while the outputs are represented horizontally. When the variables are positioned in the same line as the blue oval block labeled "MDA," it means that the variables are global inputs/outputs. The following modules are part of the MDA:

- Setup: Determine the power ratio (i.e. the power source participation ratio, see Eq.6) of each turbogenerator and battery pack regarding the aircraft topology selected.
- Aircraft sizing: Size the geometry of the aircraft, and determine its weights and aerodynamic characteristics.
- Safety performance: Determine the thrust necessary in failure condition.
- Hybrid system: Size the powertrain and each source based on the thrust and mission

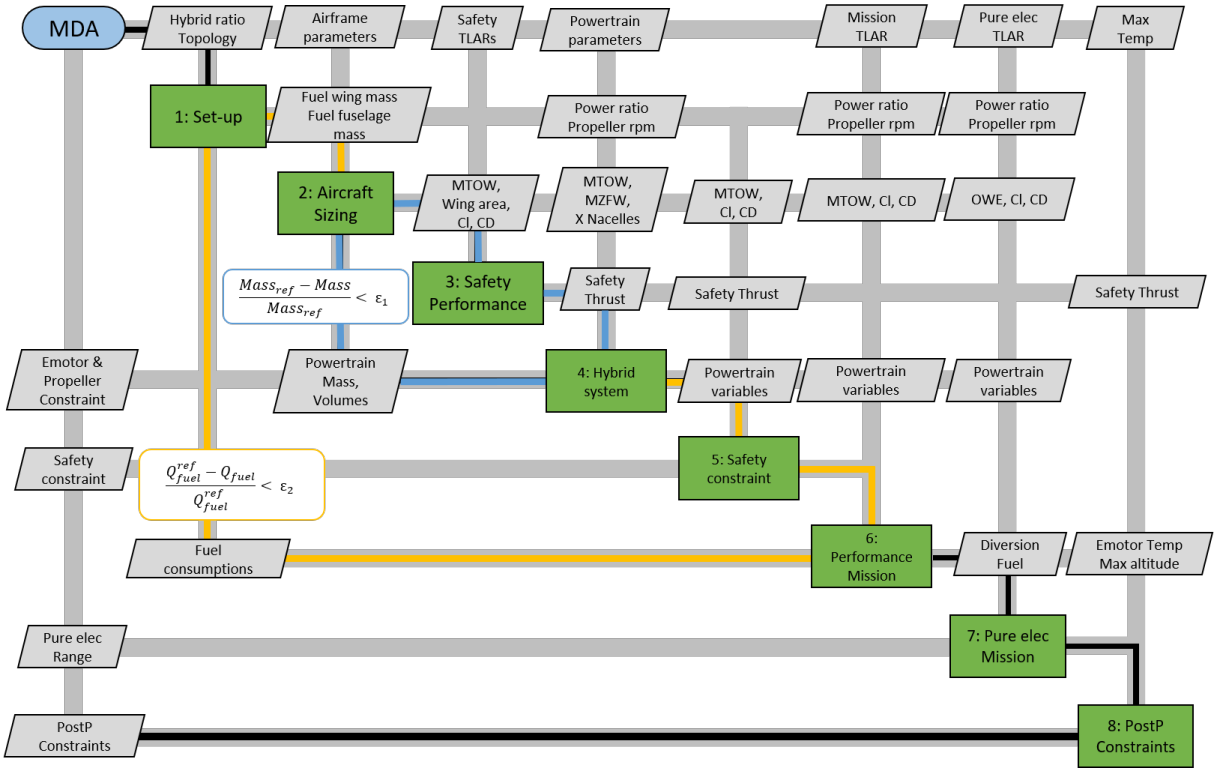


Figure 2: Top Level XDSM for aircraft design process

profile.

- Safety constraint: Verify if the other sources have enough power to complete the safety mission in case of failure in one turbogenerator or battery pack.
- Performance mission: Simulate a temporal mission with the design range (750 NM) plus a mission of diversion. During the simulation, the fuel and electric consumptions are to be calculated.
- Pure elec mission: Determine the maximal possible range using only battery energy.
- PostP constraints: Verify if the motor exhibits any indications of partial discharges or attains its maximum operating temperature during the execution of the three missions.

The analysis is composed of two loops. The first loop ensures convergence between aircraft sizing and the hybrid system, allowing the aircraft to complete the mission. The second loop is made to converge the fuel mass and battery energy.

The analysis is composed of three different missions: “hybrid”, “pure electric”, and “safety”.

The “hybrid” mission is the sizing aircraft mission with the design range of the aircraft. It must be satisfied in hybrid mode with parametric power distribution between battery and turbogenerator sources for each segment. The mission is a temporal simulation with a fuel mass loop that converges when the fuel mass necessary to achieve the mission is equal to the mass at the beginning of the mission.

The objective of the “pure electric” mission is to determine the maximum range attainable by the hybrid aircraft using solely battery power. Firstly, a temporal mission composed of a taxi-out, takeoff, climb, descent, and taxi-in is simulated to determine the energy quantity consumed during these phases. Secondly, a Breguet equation for electric aircraft [2] is calculated to know the maximal range possible (at cruise operation) in

pure electric mode by all the battery packs. This calculation is made by subtracting the energy required to perform the phases simulated during the temporal mission from the total energy stored in batteries. The fuel mass during the pure electric mission remains constant during the main mission because it is exclusively utilized during the diversion mission. This constant mass during the flight allows us to eliminate the convergence loop during the mission simulation and replace it with a temporal simulation of the climb and descent plus a Breguet equation for electric aircraft to estimate the cruise range.

The "safety" mission is composed of a take-off, an initial climb to 450 meters, ten minutes of cruise, and a descent. This simulates an emergency return-to-field procedure after a failure case during the take-off. Two safety cases are integrated: the first is the failure of one or more electric motors, and the second is the failure of one energy source. The CS-25 amendment [13], does not consider the scenario of motor failure if the number of motors is greater than four. However, with electric motors, it can be interesting to have more than four motors. To address this, a novel rule has been devised that accounts for the motor failure case contingent on the motor number. The number of inoperative motors is determined by the probability of motor failure multiplied by the number of motors (equation 4). This equation integrates all motor numbers. It takes into account that an aircraft with a large number of motor may have two or more failures at the same time.

$$nb_{failure} = \lfloor X * nb_{mot} + 1 \rfloor \tag{4}$$

with $nb_{failure}$ the number of failure motor, X the probability of failure for one motor and nb_{mot} the number of aircraft motor.

The second safety criterion is to ensure that if the largest power source fails, all other sources have sufficient power to complete the safety mission. It is assumed that each battery pack will have the same state of charge during the mission. Therefore, it is not necessary to certify the amount of power in the case of source failure, because if the aircraft can perform a pure electric mission that is larger than the safety mission, it has sufficient power to complete the safety mission without a single turbogenerator.

3 MDA SUBGROUPS AND MODELS

This section explains the most important modules and models of the MDA using the color code for group of variables below:

Orange: Outputs variables from Aircraft system
Blue: Outputs variables from Hybrid propulsion system
Yellow: Performance mission
Green: Design variables used at the inputs of the Design of Experiments
Violet: Inputs Variable fixed before the MDA execution
Black: Variables exchanged between two models inside a same module

3.1 Setup module

The setup module facilitates the delineation of the allocation of each source power through the minimization of design variables. This section elucidates the methodology for reducing the number of variables.

Is it possible to derive the power equilibrium equation at the node Ultra high voltage DC (UHVDC), figure 4, to obtain:

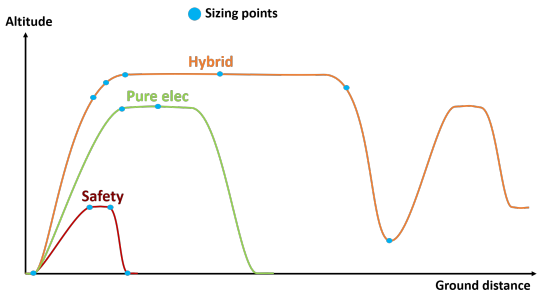


Figure 3: Sizing point for each mission

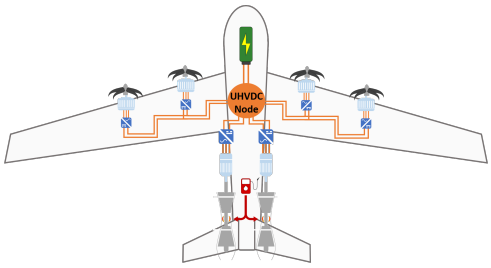


Figure 4: Position of the UHVDC node

$$P_c^s = P_{twl}^s + P_{twr}^s + P_{tfl}^s + P_{tfr}^s + P_{bw}^s + P_{bf}^s + P_{bn}^s \tag{5}$$

with P_c^s the power required by the powertrain and auxiliary systems at each segment s (takeoff, climb, cruise, etc.), P_i the power of the source i with the subscript:

- t for turbogenerator
- b for batteries
- r for the component at the aircraft right
- l for the component at the aircraft left
- w for the component under the wing for tubogenerator and in the wing for battery pack
- f for the battery pack or turbogenerator in the fuselage
- n for the battery pack inside the pod

It is assumed that the turbogenerator does not recharge the batteries. The power can be categorized into two distinct flows: the battery and turbogenerator. With these considerations, the two distinct flows can be divided by the power required (P_c^s) to obtain the hybrid ratio at each segment.

$$1 = Hr_{battery}^s + Hr_{turbogenerator}^s \text{ with } Hr_i^s = \frac{P_i^s}{P_c^s} \tag{6}$$

For each segment, the battery hybrid ratio will be a design variable.

It is possible to express each hybrid ratio in function of the power ratio of each component

$$Hr_{turbogenerator}^s = Pr_{twr}^s + Pr_{twl}^s + Pr_{tfr}^s + Pr_{tfl}^s \tag{7}$$

$$Hr_{battery}^s = Pr_{bw}^s + Pr_{bf}^s + Pr_{bn}^s \tag{8}$$

$$\tag{9}$$

with $Pr_i^s = \frac{P_i^s}{P_r^s}$ the Power ratio of the turbogenerator i

The tubogenerators can be below the wing or at the rear and the number can be one or two. The variable T_p designates the position of the turbogenerators with $Tp = 1$ if the turbogenerators are below the wing else $Tp = 0$. The variable S_r designates if the number of turbogenerators is one or two with $S_r = 0.5$ for two turbogenerators and $S_r = 0$ for only one. In the case of one turbogenerator below the wing, the turbogenerator will be at the right of the aircraft, coupled with the nearest powertrain nacelle in a larger nacelle called POD. In the case of one

turbogenerator at the rear, it will be collocated at the center of the fuselage and the batteries pack nacelle will be split into two parts in the nearest to the fuselage powertrain nacelles called POD. With these new variables, it is possible to write:

$$Pr^s_{twr} = (1 - Hr^s_{battery}) * Tp * S_r \tag{10}$$
$$Pr^s_{twl} = (1 - Hr^s_{battery}) * Tp * (1 - S_r) \tag{11}$$
$$Pr^s_{tfr} = (1 - Hr^s_{battery}) * (1 - Tp) * S_r \tag{12}$$
$$Pr^s_{tfl} = (1 - Hr^s_{battery}) * (1 - Tp) * (1 - S_r) \tag{13}$$
$$\tag{14}$$

To satisfy the battery energy constraint in case of failure of one battery pack during the flight, it is assumed that during the mission the battery State of Charge (SoC) of each battery pack will be the same. This assumption allows us to keep the power relative ratio of each battery pack independent of the segment and the equations:

$$Pr^s_{bw} = Hr^s_{battery} * Rr_{bw} \tag{15}$$
$$Pr^s_{bf} = Hr^s_{battery} * Rr_{bf} \tag{16}$$
$$Pr^s_{bn} = 1 - Pr_{bw} - Pr_{bf} \tag{17}$$
$$\tag{18}$$

with $Rr_i = \frac{P_{bi}}{P_{bw}+P_{bf}+P_{bn}}$ the relative power ratio of the battery i compared to all batteries packs

The integration of two categorical parameters enables the simulation of four distinct aircraft configurations. The hybrid ratio management is achieved through the utilization of only two parameters, in conjunction with eleven hybrid ratio variables, derived from the eleven segments that constitute the hybrid mission.

The SoC enables the maintenance of the battery's relative power ratio during both the pure electric mission and the hybrid mission.

3.2 Hybrid system module

The hybrid system module (figure 5) is composed of 6 sub-modules within the HASTECS based models [11] of the Propulsion and Battery sizing sub-modules.

3.2.1 Cable length

This submodule determines the length of each cable relative to the aircraft geometry and the locations of turbogenerator, batteries and powertrain. Additionally, it determines the position of each motor on the wing. The position of each powertrain is determined by placing the first motor at the nearest point to the fuselage, ensuring a safety distance, and positioning the subsequent powertrain at a distance equivalent to one propeller diameter plus an additional safety margin.

3.2.2 Points of Interests mission

It is necessary to size all the components of the powertrain. The motor model needs the maximal rotational speed and geometry parameters to be sized. The inverter requires the voltage and the maximal RMS current. The gearbox model needs the maximal torque. Each sizing condition corresponds to different mission points (figure 3) and it is impossible to know before simulating the performance of the others components, which point will be the sizing point because it is dependent of the others component efficiency. To solve this problem, a list of potential points will be selected.

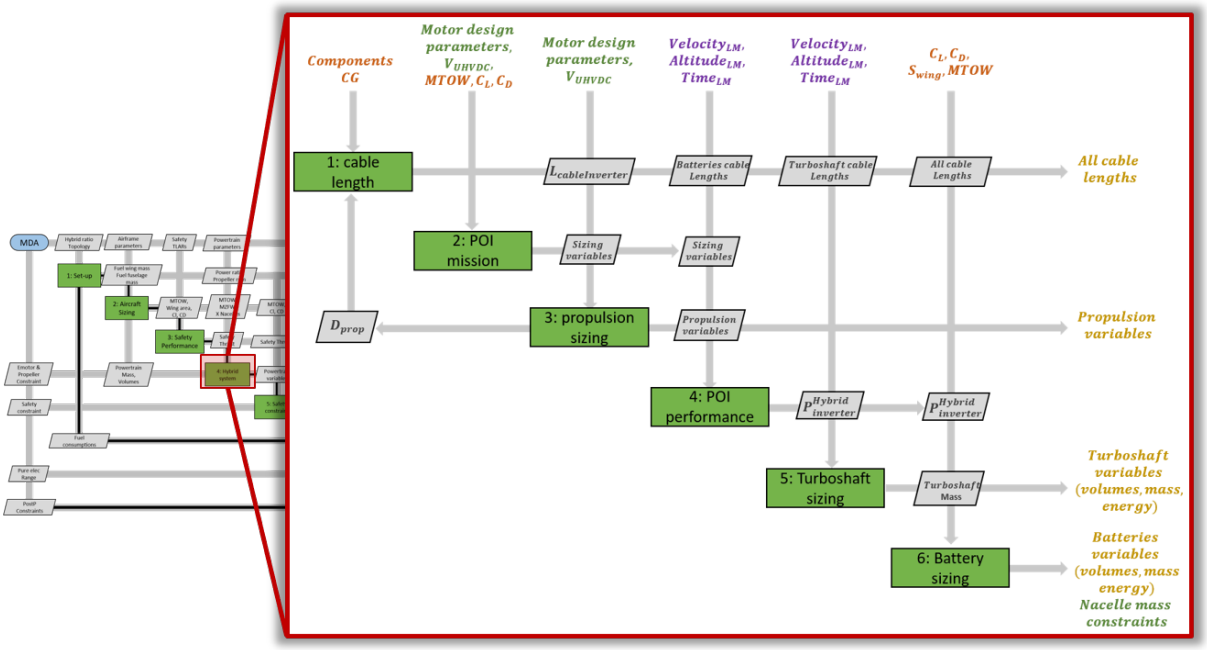


Figure 5: XDSM of the Hybrid system module

The sizing mission points are composed of two points for the takeoff, initial climb, end climb, and cruise points for each mission (hybrid, safety, pure electric). The sizing mission is further refined by incorporating three additional points at 3000 meters, 4000 meters, and 5000 meters, thereby ensuring a more precise alignment between the sizing and performance missions. This sizing mission point will be called "the point of interest mission" (POI mission). Subsequent to the selection of points of interest, the model determines the requisite thrust at each point, taking into account the new aircraft mass and drag, as well as various criteria, including slope angle and velocity. The outputs of this process are a set of points of interest (potentially sizing) containing the required thrust, altitude, and speed.

3.2.3 Powertrain sizing

The powertrain is sized moving from propeller to inverter. It begins with the propeller sizing and finishes with the cable definition from the inverter to the central Ultra High DC Voltage node. Each component is sized according to its sizing criteria (Maximal torque, Maximal RMS current, etc.) and its parameters with design variables as inputs. After, the performance is simulated to determine the next component input data. The outputs of this model are the volumes, mass, and variables needed to size the aircraft structure and simulate the performance for each component of the powertrain (figure 6).

3.2.4 Points of interest performance

This sub-module selects the hybrid mission points to determine the power necessary by the powertrain at each point of the mission.

3.2.5 Turboshaft sizing

The turboshaft model comes from the thesis of J.Thauvin [14] and is a model based on data regression from turboprop and helicopter engines which have high specific powers. The performance model is based on response surfaces with three inputs: the static power at sea

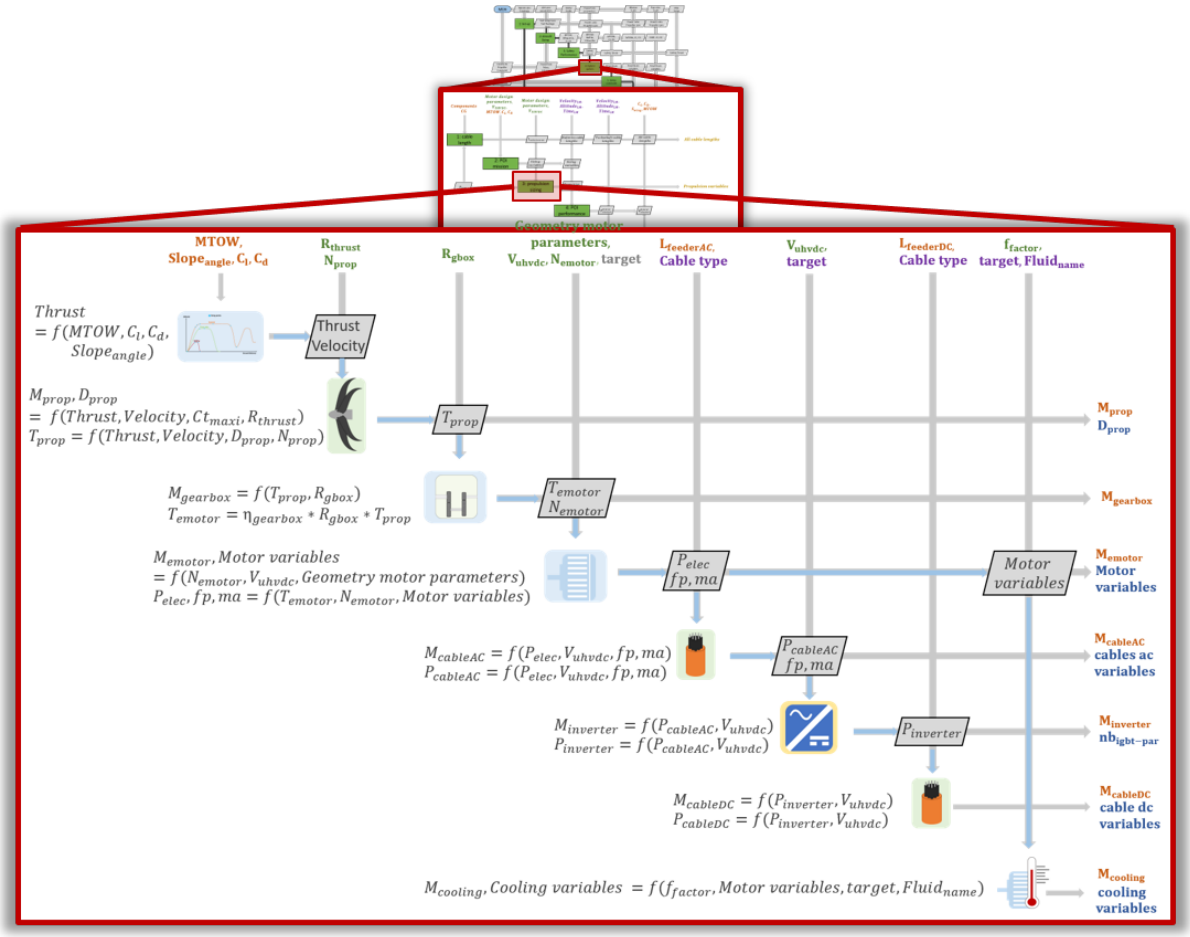


Figure 6: Flowchart of the powertrain sizing

level, the ratio between total and static pressures at sea level and the ratio between total and static temperature at sea level.

3.2.6 Battery sizing

During the HASTECS project, a battery maturity study was conducted to project future energy and power density at the 2035 horizon (tables 1 and 2). The study identified three promising technology solutions: Nickel Manganese Cobalt (NMC), Lithium Sulfur (LiS), and Lithium Titanate Oxide (LTO/TNO). A form factor of two was selected to represent the transition from cell scale to pack scale. The model determines the mass necessary to meet the power and energy requirements of each technology and identifies the most lightweight solution among the selected options.

Specific energy at cell scale (Wh/kg)	2025	2035
High Power: LTO/TNO	140	180
High Energy: NMC	250	350
Very High Energy: LiS	500	600

Table 1: Gravimetric energy density assessments (Wh/kg) for the selected technologies at the cell scale

Capabilities (Crate)	All targets
LTO/TNO	10
NMC	2
LiS	1

Table 2: Capability assessments for the selected technologies at the cell scale

3.3 Aircraft sizing module

The aircraft system models, figure 7 (structure, geometry, aerodynamics) is based on a collection of semi-empirical models adapted from several aircraft design books [15, 16]. the models have been adjusted to reproduce a regional transport aircraft, such as the ATR 72, during the V.Palladino thesis [17].

The geometry sizing model uses the method and models developed by V.Palladino during his Thesis [17] modified to adapt to our new hybrid propulsion system. In this section, only the modification will be presented.

3.3.1 Fuselage geometry

In order to integrate the battery pack dedicated to the fuselage, the surface of the battery section is considered to be a cylinder with a diameter equal to the inner diameter of the fuselage at the level of the cockpit. Knowing the battery energy necessary to complete the hybrid mission, and based on its energy volumetric density (in kWh/l), the length of the cylinder is determined. This length is added to the fuselage length. In the case of turbogenerator located at the rear, a similar process is employed for the fuel tanks, which are positioned in the rear fuselage. If turbogenerators are under the wings, the fuel tanks will be in the wing with the battery wing pack with the methodology explain in the section 3.3.3.

3.3.2 Nacelles geometry

The Nacelles powertrain geometry is defined by considering the diameter as the diameter of the motor more than the diameter of the cooling system tube added by a free space distance. The length is determined by the volume necessary to collocate all the powertrain elements in the nacelle multiplied by an oversizing factor, called cooling form factor. The cooling motor system is an active closed-loop circuit using water-glycol liquid (figure 9). The heat exchanger is

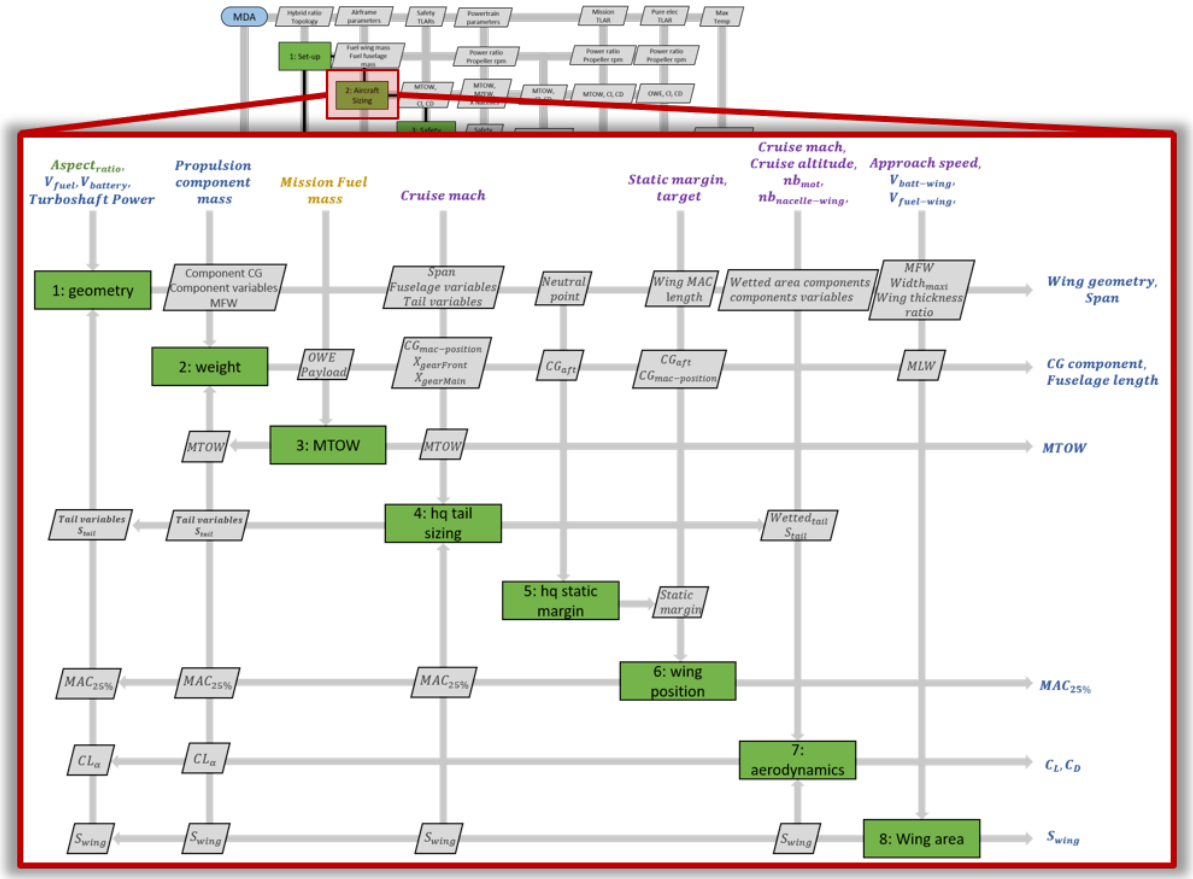


Figure 7: XDSM of the Aircraft sizing module

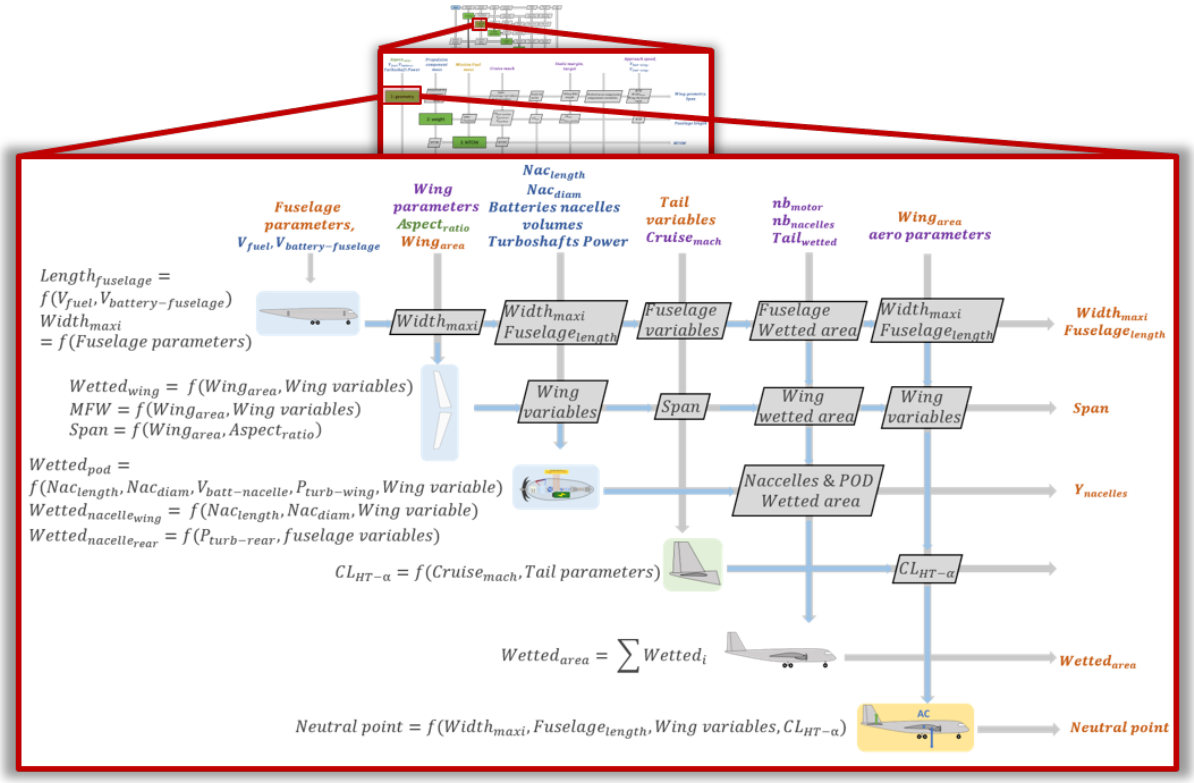


Figure 8: Flowchart of the aircraft geometry sizing

completely part of the nacelle skin. This oversizing factor allows us to increase the heat surface exchange.

3.3.3 Wing area

The calculation of the wing area is contingent upon the constraints imposed by the necessity to incorporate the fuel volume and the requisite landing velocity for the aircraft. This equation is further augmented by the constraint on the battery volumes required within the wing.

To accomplish this, the volume in the wing is divided into two parts: one from the root to the kink, which is designated as the base volume, and the other from the kink to the tip, which is referred to as the tip volume (figure 11). For the cross-sectional area of the first wing volume, an ellipse is considered, with the width being the wing chord and the height being the thickness of the wing at the "kink". The area available within the wing surface is between the two wing stiffeners. This procedure is repeated for the second wing volume, considering the area at the "tip" (10).

Initially, the volume occupied by the fuel in these two volumes is determined, subject to the constraint that fuel is only introduced into the peak volume if the bass volume is already completely filled with fuel.

For the battery volume, the largest rectangle is determined, starting from the stiffener closest to the center. Then, the largest rectangle running from the stiffener furthest from the center to the edge of the first rectangle is determined. This calculation provides the total section surface of the battery, which is multiplied by the length, subtracting the length occupied by the fuel. The calculation is performed for both, the base volume and the peak volume to determine the possible battery volume in the wing.

The module is an implicit component that iterates until the residual between the requested

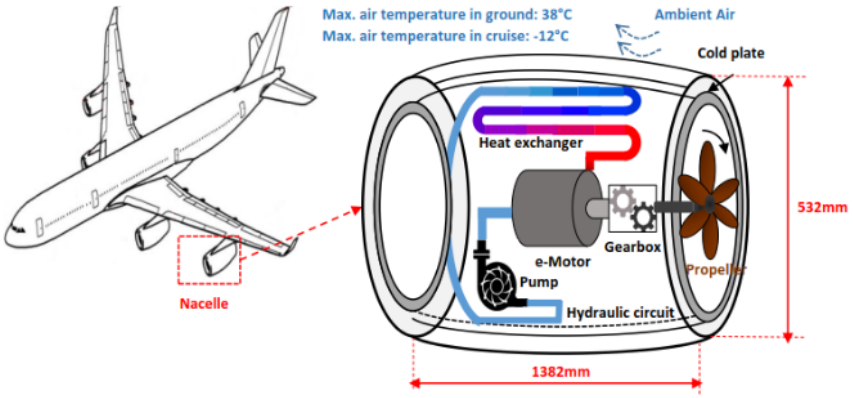


Figure 9: Nacelle architecture (extract from M.Pettes-Duler thesis [18])

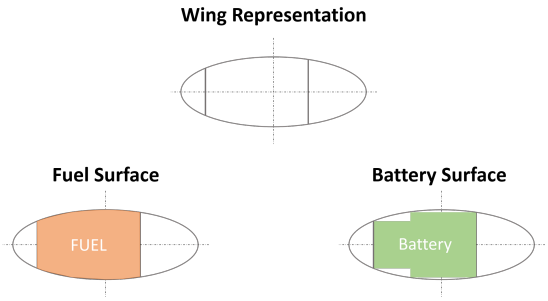


Figure 10: Representation of usable surfaces in the wing

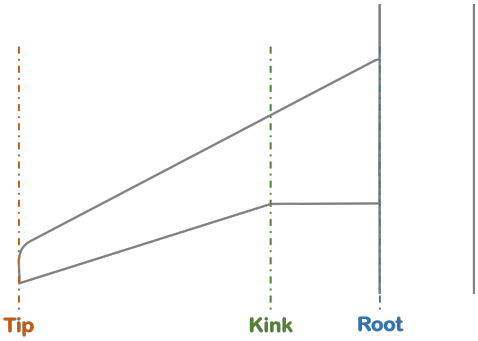


Figure 11: Breakdown of the wing into two parts

battery volume and the available battery volume is zero, then returns the required wing area.

4 CALIBRATION AND DESIGN OF EXPERIMENTS

4.1 Calibration of a conventional aircraft is considered the baseline

In order to validate the MDA, a case study with a comparison between a baseline aircraft and a regional aircraft like ATR-72 from V.Palladino thesis [17] is presented. The powertrain input variables are the results of the optimization obtained in the HASTECS project [18] for the 2025 target.

The baseline aircraft based on our design models is close in terms of mass to an actual regional aircraft like ATR-72 as shown in the table 3. The sizing block fuel is lower than the regional aircraft because the fuel consumption is based on the interpolation of different engines from Thauvin's thesis [14].

4.2 Design of Experiments

Once the optimization problem is set up (table 5), there are 38 potential design variables with 32 constraints to satisfy. The substantial number of variables and the constrained design space, resulting from a large number of constraints, make it difficult or even impossible to use optimization without a prior study to understand the impact of each variable and try to reduce its bounds. Consequently, the Design of Experiments (DOE) is executed with the battery hybrid ratio at each segment as an input variable with the segment decomposition explained in

Table 3: Comparison between the baseline aircraft and the regional aircraft from Palladino thesis data [17]

Outputs variables	Regional aircraft	Baseline
Design Range (NM)	750	750
Number of Passenger	72	72
Cruise Mach	0.45	0.45
Payload (kg)	6840	6840
MTOW (kg)	23593	23006
OWE (kg)	14155	13863
Wing area (m^2)	61	60
Sizing block fuel (kg)	2641	2403

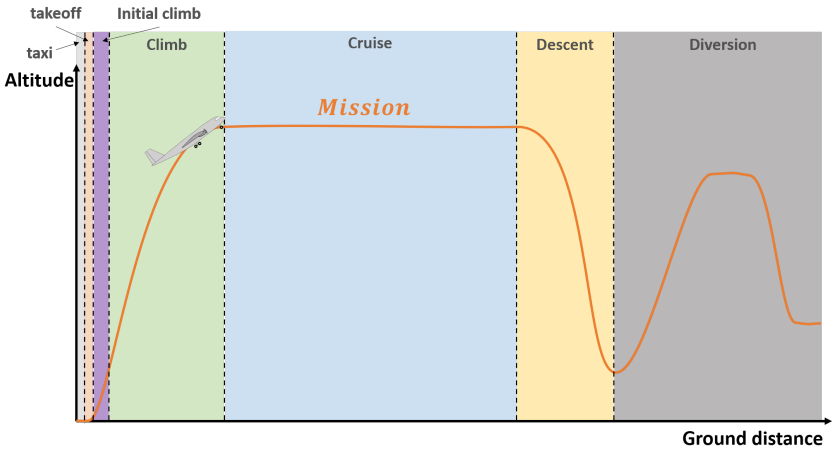


Figure 12: Mission segment descomposition

figure 12. The outputs of this DOE include the pure electric range, the maximum takeoff weight (MTOW), the wing area, and the pure electric mission validation constraint. The constraint returns a discrete value of 3 if the battery does not have enough energy to complete the pure elec temporal mission, or a value of 0 if the aircraft can perform the mission. For each DOE, the input variable repartition are a Latin Hypercube Sampling (LHS). The objectives are to determine the input variables with the greater impact, reduce the bounds, and determine if there are input variables that do not affect the outputs.

4.2.1 Hybrid ratio variation with higher battery pack density

To maximize the number of convergence points without altering trends, the energy pack density at the pack level for the very high energy solution (LiS) is determined to be 600 Wh/kg, which, at the present moment, is not a feasible technological solution within the timeframe of the horizon 2035. The inputs to the model include the initial climb, climb, and cruise hybrid ratio, with bounds from 0 to 1 for each variable, except the cruise variable, which ranges from 0 to 0.5.

The Design of Experiments (figure 14), indicates a strong effect of the cruise hybrid ratio on the pure electric range (orange lines). The pure electric range and MTOW as a function of cruise hybrid ratio present the lowest vertical dispersion, thereby indicating the importance of this variable over the others. Conversely, a low cruise hybrid ratio results in insufficient energy to complete a pure electric mission, as depicted within the green surroundings.

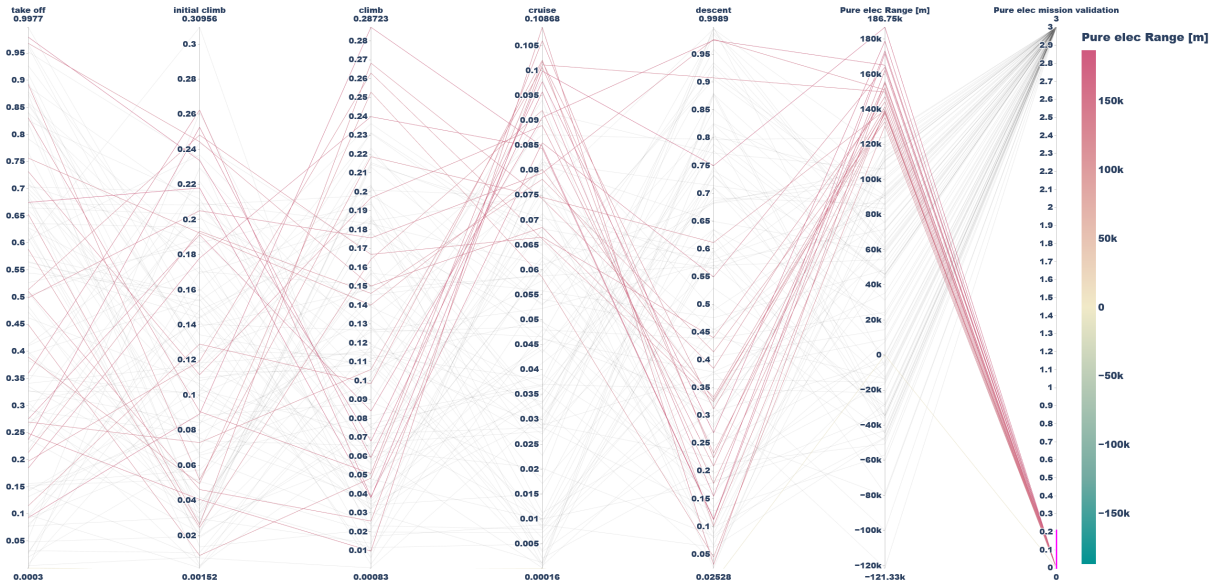


Figure 13: Parallel plot with takeoff, initial climb, climb, cruise, and descent hybrid ratio variation (hypothesis f=2)

4.2.2 Hybrid ratio variation during four segments of the mission

The present DOE was executed following the HASTECS technology hypothesis, encompassing the three climb segments in conjunction with the takeoff, cruise, and descent segments. A previous 1000-sample DOE, not shown in this paper, was performed to determine the non-convergence zone of each variable to reduce the boundaries of each hybrid ratio segment. For the present DOE, the cruise segment bounds are set from 0 to 0.15, while the initial climb and climb bounds are established from 0 to 0.4. The takeoff and descent hybrid ratios are set from 0 to 1. In figure 13 are represented the results of the DoE in a parallel plot where only the valid configurations are highlighted. It shows that an adequate cruise ratio, between 0.07 and 0.1, is essential to meet the mission validation constraint and the convergence of the MDA. This adequate cruise ratio enables the achievement of the higher pure elec range, which is the function to be maximized. As illustrated in the parallel plot, an elevated initial climb ratio is frequently accompanied by a diminished climb ratio. This phenomenon indicates the achievement of a maximum feasible battery mass that can be onboard on the aircraft. When a higher initial climb is chosen, the climb ratio must be lower to have a quasi-constant global hybrid ratio. It should be noted that the take-off and descent hybrid ratio does not have a tangible effect on the output variables. It is imperative to acknowledge that this DOE has been constructed with fixed motor and aircraft variables, which leads to constraint violation. The next section aims to determine the tendency of these constraints.

4.2.3 Hybrid system variables

The current DOE (figure 15) is for a hybrid ratio fixed for each segment, and the variables are the DC voltage (V uhvdc), the gearbox ratio, the motor bore(R_alesage), the motor maximal speed (Maximal speed), and the oversizing factor of the powertrain nacelle (cooling form factor). It is important to note that the DOE exclusively illustrates the points that satisfy the pure elec mission validation constraint. The DOE does not disclose a particular tendency concerning the pure electric range or its constraints. However, the high amplitude of the pure electric range highlights the influence of the variables.

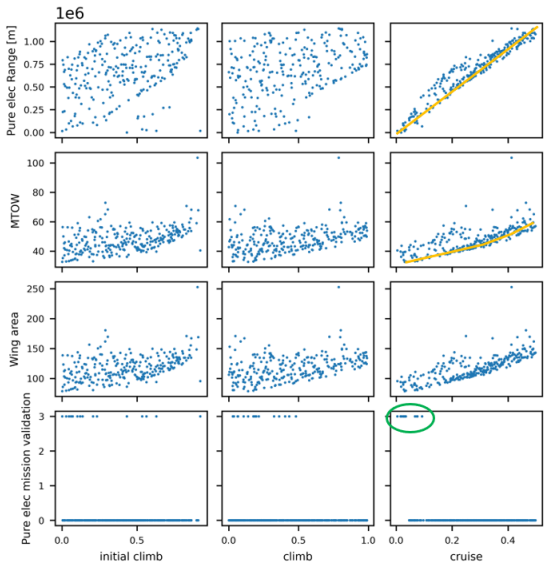


Figure 14: DOE with initial climb, climb, and cruise hybrid ratio variation (hypothesis: f=1)

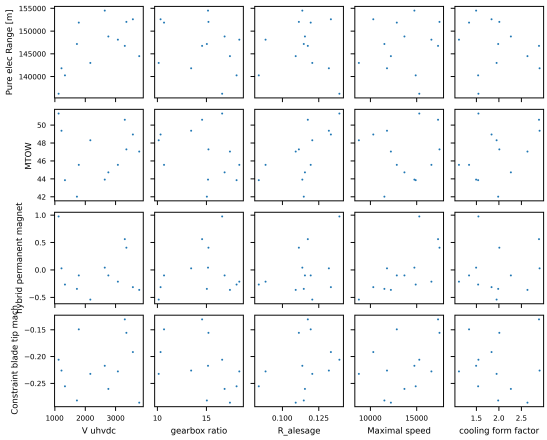


Figure 15: DOE with four motor variables

5 CONCLUSIONS

The implementation of the HASTECS propulsion system in the FAST-OAD aircraft design software has been demonstrated. This work represents an attempt to reduce process output uncertainties by integrating the couplings between propulsion, structure, and aerodynamics.

This novel method transcends the limitations of HASTECS and FAST-OAD taken separately, thereby facilitating the exploration of a more expansive design space through enhanced hybridization, as well as greater flexibility in topology, mission performance, and wing design.

The Design of the Experiments (DOE) underscores the significance of the cruise hybrid ratio. The maximum pure electric range, corresponding to a ratio of approximately 0.15, marks the point preceding the non-convergence of the MDA. While the motor and aircraft variables demonstrate no discernible tendency, the substantial variation in the pure electric range indicates a significant impact.

The primary challenges associated with the analysis of MDA results are primarily due to the difficulty of obtaining a converged MDA result that satisfies all constraints and the significant impact of input variables on the pure electric range without having a tendency. One potential approach to address these challenges is the implementation of optimization methods, such as Bayesian optimization, or the utilization of importance sampling to increase the number of feasible points.

Table 4: Optimization problem to maximize the pure elec mission range

Optimization problem	
Maximize: Range in pure electric	
With respect to: (38 variables)	
Motor number ($n_{b_{mot}}$)	($R_{hr_{ral}}$)
Hybrid ratio battery ($Hr^{s}_{battery}$) × 11 segments	Ratio between air gap thickness and stator inner radius ($R_{g_{ral}}$)
Topology variable (T_p)	Ratio between magnet thickness and stator inner radius ($R_{pm_{ral}}$)
Symmetry variable (A_r)	Maximum yoke flux density (B_{yoke})
Battery wing and nacelle relative power ratio (Rr_{bw})	Maximum teeth flux density (B_{teeth})
Battery nacelle relative power ratio (Rr_{bn})	Number of conductors per slot (N_{ce})
Nacelle form factor (f_f)	Number of slots per pole and per phase (n_{epp})
Motor maximal rotational speed (V^{max}_{motor})	Number of pole pairs (p)
Motor cruise relative speed ($\%speed_{cruise}$)	Gearbox ratio (R_{gbox})
Motor take-off relative speed ($\%speed_{takeoff}$)	Aspect ratio (A_R)
Ultra High DC Voltage (V_{uhvdc})	Virtual Taper ratio (V_{tp})
Inner radius of the stator ($R_{alesage}$)	Kink span ratio (Sp_{kink})
Ratio between rotor diameter and active length ($R_{drot_{lm}}$)	Propeller oversized factor (P_{of})
Ratio between stator slot and inner radius	Battery oversized factor (B_{of})
	Turbogenerator oversized factor (T_{of})
Subject to: (33 constraints)	
Nacelle left mass = Nacelle right mass	Pure elec climb time < Maximal climb time
Blade tip Mach < 1	Pure elec maximal altitude ≥ maximal altitude
Outside propeller position < Wing length	Pure elec mission validation == 0
Discharge fill fit in the slot × 3 missions types	Power battery i during pure elec mission < Maximal power battery i
Discharge flag == 0 × 3 missions types	Minimal airgap < Emotor airgap
Hybrid & Pure elec & Safety mission:	Centrifugal pressure < Maximal centrifugal pressure
Maximal flux density < maximal possible flux density	Peripheral speed < Maximal peripheral speed
Head windings temperature < Maximal head windings temperature	Minimal shaft radius < Shaft radius
Permanent magnet temperature < Maximal permanent magnet temperature	max($P_{sources}$(failure condition)) > max($P_{required}$)
Stator temperature < Maximal stator temperature	

ACKNOWLEDGEMENTS

ONERA / LAPLACE collaboration, co-financed by the Occitanie region under the number 23003064

REFERENCES

[1] H. Dodd. Waypoint 2050 an air transport action group project. URL https://atag.org/resources/waypoint-2050-2nd-edition-september-2021.

[2] G. E. Wroblewski and P. J. Ansell. Hybrid-electric breguet range equation: Development & validation against a hybrid-electric aircraft simulation. In *AIAA Propulsion and Energy 2020 Forum*. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-602-6. doi:10.2514/6.2020-3579. URL https://arc.aiaa.org/doi/10.2514/6.2020-3579.

[3] ATR 72-600 aircraft | ATR aircraft. URL https://www.atr-aircraft.com/our-aircraft/atr-72-600/.

[4] P. Novelli. IMOTHEP_final-brochure, 2024. URL https://www.imothep-project.eu/.

[5] V. Marciello, M. Di Stasio, M. Ruocco, V. Trifari, F. Nicolosi, M. Meindl, B. Lemoine, and P. Caliandro. Design exploration for sustainable regional hybrid-electric aircraft: A study based on technology forecasts. 10(2):165, February 2023. ISSN 2226-4310. doi:10.3390/aerospace10020165. URL https://www.mdpi.com/2226-4310/10/2/165.

[6] R. Jansen, C. C. Kiris, T. Chau, L. M. Machado, J. C. Duensing, A. Mirhashemi, J. Chapman, B. D. French, L. Miller, J. S. Litt, C. L. Denham, M. N. Lynde, R. L. Campbell, B. R. Hiller, N. J. Blaesser, and N. Heersema. Subsonic single aft engine (SUSAN) transport aircraft concept and trade space exploration. In *AIAA SCITECH 2022 Forum*. American Institute of Aeronautics and Astronautics, January 2022. ISBN 978-1-62410-631-6. doi:10.2514/6.2022-2179. URL https://arc.aiaa.org/doi/10.2514/6.2022-2179.

[7] G. Cinar, Y. Cai, M. V. Bendarkar, A. I. Burrell, R. K. Denney, and D. N. Mavris. System analysis and design space exploration of regional aircraft with electrified powertrains. 60(2): 382–409, March 2023. ISSN 0021-8669, 1533-3868. doi:10.2514/1.C036919. URL https://arc.aiaa.org/doi/10.2514/1.C036919.

[8] P. Novelli, S. Defoort, N. Tantot, D. Zimmer, D. Varchetta, C. Viguier, and P. Iodice. IMOTHEP european project: an investigation of hybrid electric propulsion for commercial aircraft. 2023.

[9] C. David, S. Delbecq, S. Defoort, P. Schmollgruber, E. Benard, and V. Pommier-Budinger. From FAST to FAST-OAD: An open source framework for rapid overall aircraft design. 1024:012062, January 2021. ISSN 1757-899X. doi:10.1088/1757-899X/1024/1/012062. URL https://iopscience.iop.org/article/10.1088/1757-899X/1024/1/012062.

[10] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor. OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization. 59(4): 1075–1104, April 2019. ISSN 1615-147X, 1615-1488. doi:10.1007/s00158-019-02211-z. URL http://link.springer.com/10.1007/s00158-019-02211-z.

[11] X. Roboam, S. Touhami, N. Erroui, A. Zeaiter, F. Accorinti, P. Collin, and M. Pettes-Duler. Hastecs: Hybrid aircraft: research on thermal and electric components and systems, 2021. URL https://www.laplace.univ-tlse.fr/portfolio/hastecs/. PhD Thesis Manuscript.

[12] T. Marien, K. Antcliff, M. D. Guynn, D. P. Wells, S. J. Schneider, M. J. Tong, A. Trani, N. Hinze, and S. Dollyhigh. Short-haul revitalization study final report. 2018. doi:10.13140/RG.2.2.10821.37608. URL http://rgdoi.net/10.13140/RG.2.2.10821.37608. Publisher: NASA.

[13] EASA. cs-25_amendment_25, 2020.

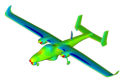
[14] J. Thauvin. Exploring the design space for a hybrid-electric regional aircraft with multidisciplinary design optimisation methods, 2018. URL http://www.theses.fr/2018INPT0092/document. PhD Thesis Manuscript.

[15] W.-P. Dupont and C. Colongo. *Avant Projet d'un avion de transport commercial*. Cours ISAE-SUPAERO.

[16] E. Torenbeek. *Synthesis of subsonic airplane design: an introduction to the preliminary design, of subsonic general aviation and transport aircraft, with emphasis on layout, aerodynamic design, propulsion, and performance*. Delft University Press ; Nijhoff ; Sold and distributed in the U.S. and Canada by Kluwer Boston. ISBN 978-90-247-2724-7.

[17] V. Palladino. Méthode de conception pluridisciplinaire appliquée à un avion régional à faibles émissions, 2023. URL http://www.theses.fr/2023ESAE0020/document. PhD Thesis Manuscript.

[18] M. Pettes-Duler. Conception intégrée optimale du système propulsif d'un avion régional hybride électrique, 2021. URL http://www.theses.fr/2021INPT0043/document. PhD Thesis Manuscript.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

HYDROGEN-POWERED AIRCRAFT DESIGN: A MULTIDISCIPLINARY APPROACH TO MINIMIZE EMISSIONS AND COST

Raúl Quibén Figueroa[1*], Rauno Cavallaro[1], Andrea Cini[1] and Manuel Soler Arnedo[1]

1: Department of Aerospace Engineering
Universidad Carlos III de Madrid
Av. de la Universidad 30, 28911 Leganés, Madrid
{rquiben, rcavalla, acini, masolera}@ing.uc3m.es

Abstract. *Hydrogen fuel is increasingly recognized as a transformational solution to address the environmental impact of traditional fossil fuel-based propulsion systems in aviation. This disruptive technology holds the potential to revolutionize the industry, but its successful industrialization demands a comprehensive evaluation of both its economic feasibility and environmental benefits. This study exploits a novel, integrated design and optimization tool tailored for conceptualizing hydrogen-powered aircraft. The framework, named MOTIVATION and built on OpenMDAO, employs gradient-based optimization and total derivatives computation for complex models, enabling the efficient evaluation of the problem regarding all involved flight performance and climate-related disciplines. Different modules are designed to comply with CS/FAR 25 requirements and airworthiness constraints, while estimating specific aircraft responses and evolution, particularly suited for hydrogen-powered systems. Special attention is given to direct and indirect emissions, ensuring comprehensive environmental impact evaluations. The presented optimization campaign assesses the feasibility of retrofitting an A320neo with hydrogen propulsion, focusing on the associated uncertainties owing to the lack of established reference designs for such innovative systems. Research key aspects include analyzing various configurations for hydrogen storage while ensuring structural and operational feasibility through equilibrium and stability evaluations. Comparative analyses are conducted to quantify the costs and emissions, including carbon dioxide (CO_2), nitrogen oxides (NO_x) and water vapor (H_2O). Parametric studies further investigate the influence of mission profiles and advancements in hydrogen technology, identifying critical viability thresholds and significant climate impact mitigation opportunities. This study provides valuable insights into the potential pathways for integrating hydrogen into sustainable aviation.*

Keywords: hydrogen, climate impact, sustainability, multidisciplinary, design, optimization, aerospace

1 INTRODUCTION

Air traffic is recognized as a major contributor to anthropogenic global warming [1]. In response, the aviation industry has undertaken several initiatives to mitigate its climate impact, with hydrogen (H_2) emerging as one of the most promising yet challenging alternatives. Hydrogen-powered systems offer the potential to mitigate direct CO_2 emissions, while overcoming the range and scalability limitations associated with battery-electric propulsion [2].

The introduction of such vehicles still needs studies to assess their feasibility and maximize their potential benefits. Design optimization platforms have demonstrated their effectiveness in aircraft design by integrating physics-based principles into the decision-making process [3, 4]. Automated search algorithms facilitate early-stage design exploration, leveraging all the potential of the new technology. These tools are particularly beneficial during conceptual design, especially when assessing novel aircraft configurations that require extensive trade-off analyses.

Beyond CO_2 reduction, the environmental implications of hydrogen combustion, particularly in terms of non-CO_2 emissions such as water vapor and nitrogen oxides (NO_x), must be carefully evaluated. These emissions can influence atmospheric conditions, contributing to the formation of contrails and cirrus clouds [5]. Hence, hydrogen-powered aircraft design must also be approached as both an emissions management and a climate impact optimization problem.

Recent studies have examined the feasibility and performance of hydrogen-powered mid-long range aircraft from an environmental impact perspective [5–7]. Moreover, alternatives, such as the dorsal tank layout [8], the Blended Wing Body (BWB) design [9], and novel concepts like the box-wing configuration [10] have been investigated as well.

This study focuses on the retrofitting of existing aircraft to facilitate a more seamless market entry for hydrogen propulsion technology. Specifically, the feasibility of modifying commercial aircraft to meet certification standards such as CS-25 and FAR-25 is assessed. The in-house platform, MOTIVATION [7, 11], is used to retrofit the A320neo aircraft, estimating the impact of hydrogen propulsion on performance, emissions, costs, and operational efficiency. The methodology, test case, and results are presented to offer insights into the practical application of hydrogen-powered systems compared against modern commercial aviation.

2 METHODOLOGY

The present study has been conducted using the MOTIVATION platform [7], a computational framework designed for the conceptual design and optimization of novel aircraft configurations. This platform is derived from OpenRegional [3], which, in turn, originates as a fork from OpenConcept [4]. All of these tools are developed upon OpenMDAO library [12], an open-source framework for multidisciplinary analysis and optimization.

The core modules of this platform are introduced below. These modules facilitate the mission definition and the analysis of aircraft performance and environmental impact. Specifically, they enable the continuous monitoring of the aircraft state throughout the mission, ensuring a comprehensive evaluation of its operational characteristics.

2.1 Mission and flight plan

The mission module defines both the nominal and diversion flight segments, ensuring compliance with CS-25 and FAR-25 airworthiness regulations [13]. The nominal mission follows standard flight phases (Fig. 1), which include taxi, takeoff, climb, cruise, descent, holding, and landing. In contrast, the diversion mission accounts for extended operations by incorporating additional fuel reserves and including climb, cruise, descent, and a 30-minute loiter at an altitude of 1,500 feet, as required by regulatory standards [13].

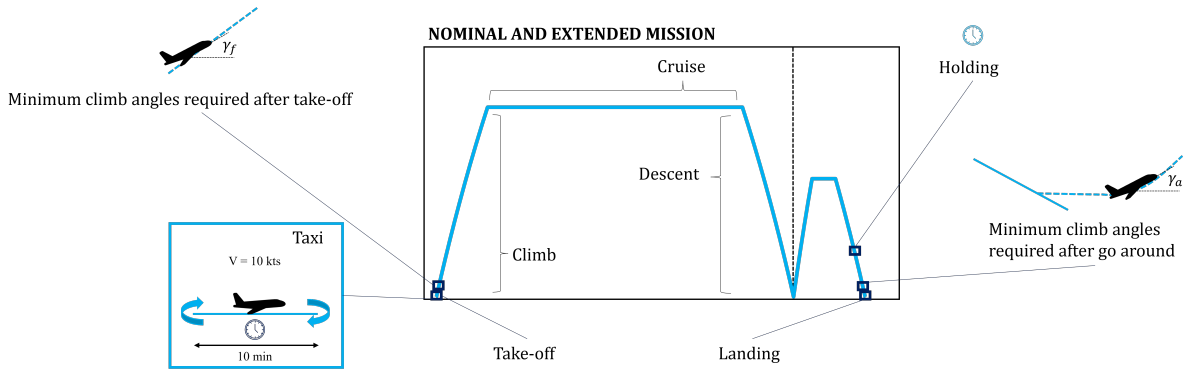


Figure 1: Phases of the reference flight profile

The aircraft dynamics throughout the mission are governed by the fundamental equations of motion. During airborne phases, the forces acting on the aircraft, namely lift (L), weight (mg), thrust (T), and drag (D), must satisfy the equilibrium equations:

$$\begin{cases} L = mg\cos\gamma, \\ T - D = mg\sin\gamma, \end{cases} \tag{1}$$

where γ represents the climb angle. In this formulation, thrust is assumed to be aligned with the velocity vector in all flight phases.

During ground operations, the aircraft acceleration (dV/dt) is governed by the following equation:

$$m\frac{dV}{dt} = T - D - \mu(mg - L), \tag{2}$$

where μ is the runway friction coefficient, accounting for the friction between the aircraft tires and the runway surface.

The mission simulation iteratively adjusts throttle settings to achieve steady-state conditions during each flight phase. For takeoff and landing, the equations governing ground acceleration (Eq. 2) are numerically integrated to compute the Balanced Field Length (BFL), which represents the minimum runway length required for a safe takeoff. Additionally, minimum climb gradient requirements are enforced following takeoff and during go-around maneuvers, ensuring regulatory compliance, as thoroughly explained in [3].

A taxi phase has been incorporated into all mission profiles to account for fuel consumption and worst-case scenarios related to cryogenic fuel boil-off in hydrogen fuel tank. The taxi phase is modeled with a ground speed of 10 knots and a duration of 10 minutes. However, it is excluded from BFL calculations, as it does not influence the worst-case takeoff performance scenario.

2.2 Aerodynamics model

The aerodynamic module employs the classic drag polar method to compute drag, lift, and stall speed [14]:

$$C_D = C_{D0} + KC_L^2 \quad \left(K = \frac{1}{\pi ARe}\right) \tag{3}$$

where C_{D0} represents the parasitic drag coefficient (or zero-lift drag coefficient), K is the induced drag factor, AR is the wing aspect ratio, and e is the Oswald efficiency factor. The Oswald factor in clean configuration is calculated using Obert's method [15], while parasitic drag is estimated via the equivalent skin-friction method [14].

To account for variations in take-off and landing configurations caused by the high-lift devices, the aerodynamic module applies corrections to the drag polar. Ground effect is incorporated by modifying K according to [3, 14].

At higher flight speeds, compressibility effects become significant, introducing wave and profile drag components that greatly impact aerodynamic efficiency. These effects, primarily associated with transonic and localized supersonic regimes, contribute to an increase in the total drag coefficient (C_D), particularly through an added compressibility drag term within the parasitic drag component (C_{D0}). Estimating this contribution solely based on skin-friction models is challenging, as it depends on multiple factors, including Mach number (M), local flow conditions, and airfoil geometry. To account for these effects, this study introduces a fitted function of the form $C_{Dw} = f(c_n, M) = c_0 + c_1 M + c_2 M^2 + ...$, which is calibrated against empirical aerodynamic models, the reference for the A320 family, using Poll-Schumamm [16] and BADA [17] methods.

2.3 Storage and power-plant configuration

2.3.1 Hydrogen storage

Given the low volumetric energy density of liquid hydrogen (LH_2), large tank volumes are required, introducing structural and thermal management challenges [18]. The efficiency of hydrogen storage is assessed using the gravimetric index (η_{grav}), defined as:

$$\eta_{grav} = \frac{m_{LH2}}{m_{LH2} + m_{structure}}, \tag{4}$$

where m_{LH2} is the mass of stored hydrogen, and $m_{structure}$ is the structural and auxiliary system mass of the tank.

Maintaining stable pressure within cryogenic LH_2 tanks is essential for safe operation, as pressure variations arise due to heat-induced boil-off and fuel consumption fluctuations [19, 20]. Ideally, hydrogen consumption matches evaporation rates to minimize boil-off losses; however, real-world deviations require controlled venting.

Pressure dynamics are modeled using the cryogenic storage framework of Lin et al. [21], incorporating thermodynamic and mass conservation principles. The pressure evolution equation is given by:

$$\frac{dP}{dt} = 2\frac{\Phi}{V_t H2}\left[\dot{Q} - \dot{m}H_{lv}\left(1 + \frac{\rho_g}{\rho_g - \rho_l}\right)\right], \tag{5}$$

where $\dot{Q}$ represents incoming heat , $\dot{m}$ is the hydrogen mass flow rate, H_{lv} is the latent heat of vaporization, and ρ_g and ρ_l denote the gaseous and liquid hydrogen densities,

respectively. The model iteratively adjusts tank pressure through mass venting, ensuring safe operation without exceeding structural limits.

2.3.2 Turbofan propulsion

Surrogate turbofan model is developed using pyCycle [22], an OpenMDAO-based thermodynamic cycle modeling tool inspired by NASA's NPSS [23]. Engine performance data is generated across throttle, Mach number, and altitude ranges, covering the full flight envelope. A surrogate model, built with Kriging interpolation [24], enables rapid predictions of thrust, fuel flow, fuel-to-air ratio, compressor discharge pressure (P_3), and compressor discharge temperature (T_3).

2.4 Longitudinal stability

Longitudinal stability is a fundamental aspect of aircraft controllability, referring to an aircraft ability to return to its original attitude after a perturbation. This stability is primarily influenced by the position of the center of gravity (CoG) and the aerodynamic forces acting on the aircraft [25]. The integration of hydrogen propulsion systems introduces additional complexities due to the presence of large hydrogen tanks at the far ends of the fuselage. As these tanks are located far from the center of gravity, fuel consumption during flight induces significant shifts in mass distribution and CoG, requiring continuous monitoring and management of longitudinal stability throughout all flight phases.

A key requirement for static longitudinal stability is that the derivative of the pitching moment coefficient about the center of gravity with respect to the angle of attack must be negative [25]. This can be expressed as:

$$\frac{\partial C_{M_{cg}}}{\partial \alpha} = C_{M\alpha} < 0 \tag{6}$$

where $C_{M_{cg}}$ is the pitching moment coefficient about the center of gravity, α is the angle of attack, and $C_{M\alpha}$ is the stability derivative, which must be negative for stability. To satisfy this condition, the aircraft mass must be properly balanced, taking into account the CoG location and the placement of wing and horizontal stabilizer as surfaces that generate the lift force.

The total aircraft longitudinal moment, which influences the aircraft stability, is given by:

$$\begin{aligned} M_{\text{total}} = {} & M_{\text{wings}} + M_{\text{tail}} + M_{\text{fuselage}} + M_{\text{engines}} + M_{\text{landing gear}} \\ & + M_{\text{payload}} + M_{\text{H2 tanks}} + M_{\text{systems}} + M_{\text{fuel systems}} \\ & + M_{\text{furnishing}} + M_{\text{operational items}} \end{aligned} \tag{7}$$

The center of gravity is then determined by:

$$x_{\text{CoG}} = \frac{M_{\text{total}} - M_{\text{consumed fuel}}}{W_{\text{total}}} \tag{8}$$

where M_{total} is the total longitudinal moment, $M_{\text{consumed fuel}}$ represents the moment contribution from the burned fuel across the mission, and W_{total} is the total weight of the aircraft. As fuel is consumed, the CoG shifts, which in turn affects the aircraft stability.

The aircraft neutral point (x_n) is the location where longitudinal stability becomes neutral [25]. It is approximated by:

$$\overline{x_n} = \overline{x_{ac}} + C_{HT} \frac{C^*_{L_{H\alpha}}}{C_{L_{WB\alpha}}} \left(1 - \frac{\partial \varepsilon}{\partial \alpha}\right) \tag{9}$$

where $\overline{x_{ac}}$ is the aerodynamic center location, C_{HT} is the horizontal tail volume coefficient, $C^*_{L_{H\alpha}}$ is the lift coefficient derivative of the horizontal tail (with respect to the local angle of attack), $C_{L_{WB\alpha}}$ is the lift coefficient derivative of the main wing and body, and $\frac{\partial \varepsilon}{\partial \alpha}$ is the downwash gradient.

The static margin (SM), key indicator of stability, is defined as:

$$SM = \overline{x_n} - \overline{x_{cg}} \tag{10}$$

where $\overline{x_n}$ is the non-dimensional neutral point location, $\overline{x_{cg}}$ is the non-dimensional center of gravity location, and the overline denotes normalization by the mean aerodynamic chord (MAC). If $\overline{x_{cg}} < \overline{x_n}$, then $SM > 0$ and the aircraft is statically stable. If $SM < 0$, the aircraft is unstable. Modern fly-by-wire systems allow for relaxed stability conditions by augmenting stability through control algorithms, maintaining safe flight even when the static margin is near zero.

2.5 Emissions and climate impact

Computational models that assess economic and environmental costs are essential for evaluating the feasibility of hydrogen as an alternative fuel in novel aircraft designs. Table 1 lists a comparison of the key properties of conventional Jet A-1 fuel and two hydrogen variants: blue hydrogen and green hydrogen [26].

		Jet A-1	**Blue H_2**	**Green H_2**
LHV	MJ/kg	43.1	118.35	118.35
EI_{CO_2}	kg_{CO_2}/kg_{Fuel}	3.76	3.689	0.054
EI_{H_2O}	kg_{H_2O}/kg_{Fuel}	1.2	8.937	8.937
EI_{NO_x} (SL)	kg_{NO_x}/kg_{Fuel}	ICAO Databank [27]	$A \cdot P_3^m \cdot \exp(B \cdot T_3 + C)$ [7]	
EI_{NO_x} (FL)	kg_{NO_x}/kg_{Fuel}	BFFM [28]	P_3-T_3 method [29]	

Table 1: Comparison of Jet A-1 and hydrogen Fuels, including Low Heating Value (LHV) and emission index (EI) for H_2O, CO_2 and NO_x

The emission of nitrogen oxides (NO_x) is a key factor in assessing environmental impact. For Jet A-1, NO_x emissions at sea level (SL) are obtained from the ICAO Engine Emissions Databank [27] and adjusted for flight level (FL) conditions using the Boeing Fuel Flow Methodology (BFFM) [28].

Unlike Jet A-1, hydrogen combustion does not produce NO_x via prompt or fuel-nitric mechanisms due to the absence of hydrocarbon radicals [30]. However, thermal NO_x is still generated as a result of high combustion temperatures. These emissions are estimated using a semi-empirical model based on compressor discharge pressure (P_3) and temperature (T_3) [7], where A, B, C, and m are empirical coefficients.

Analogous to the BFFM approach, which adjusts SL emissions to FL conditions, the P_3–T_3 methodology [29] is applied to extrapolate thermal NO_x emissions to actual flight conditions.

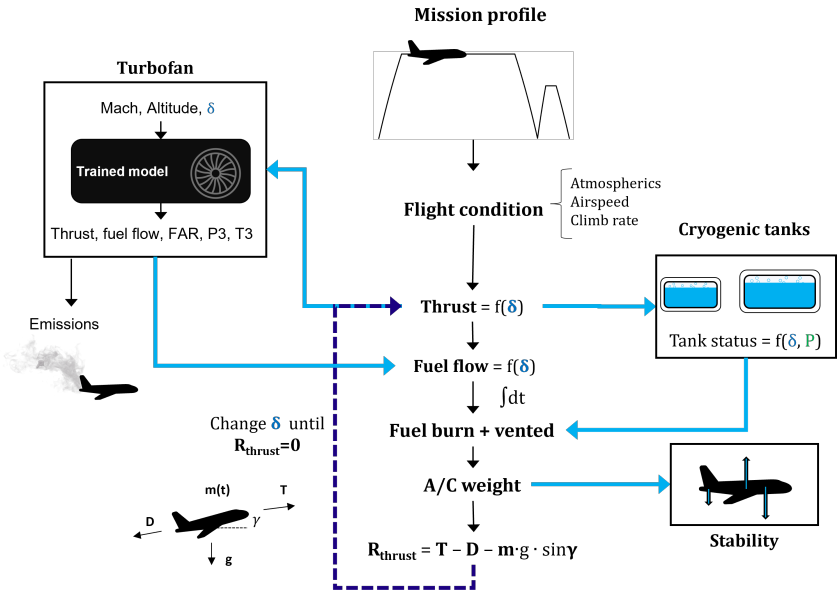


Figure 2: Analysis workflow during mission segments

2.6 MDA and MDO

The multidisciplinary analysis (MDA) ensures dynamic equilibrium throughout the mission via nested iterative cycles. The outer cycle adjusts segment durations for range consistency, while the inner cycle solves for throttle (δ) to satisfy thrust equilibrium:

$$\boldsymbol{R}_{\text{thrust}} = \boldsymbol{T}(\delta) - \boldsymbol{D} - m\boldsymbol{g}\sin\gamma \tag{11}$$

The Newton solver of OpenMDAO minimizes the residual ($\boldsymbol{R}_{\text{thrust}} = 0$).

Tank pressure is continuously monitored throughout the mission, and excess hydrogen is vented when necessary to maintain safety and structural limits. Fuel and energy consumption, whether from jet fuel or hydrogen, is tracked to ensure sufficient reserves are available for all mission segments. By resolving the aircraft inertial state over time, the evolution of mass distribution can be predicted, enabling a continuous assessment of static stability.

Optimization enforces airworthiness constraints (e.g., climb gradients) and fuel sufficiency, using a hybrid MDF-SAND approach [31]. OpenMDAO analytic derivatives enhance gradient-based optimization. Fig.2 shows the complete analysis workflow per mission segment

3 TEST CASE

This work analyzes the feasibility of hydrogen conversion for the Airbus A320neo, a widely used mid-range turbofan airliner. To support the study, an extensive data collection process is conducted, focusing on geometric, weight, and performance parameters of this aircraft. Relevant information is gathered from various literature sources [32, 33], with reasonable estimates applied when explicit data is unavailable.

A conventional A320neo reference model is developed within the platform to enable comparative analysis with the hydrogen-powered variant. The selected propulsion system is the CFM-56 turbofan, for which a publicly available surrogate model [34] is employed.

Following data acquisition, the weight estimation module computes the operating empty weight (OEW), resulting in a 9% discrepancy from the assumed reference value of 43,200 kg [7]. This discrepancy is addressed through the application of a correction factor, ascribed to the omission of composite materials in the semi-empirical equations typically provided in handbooks.

3.1 H2-retrofit aircraft

The selected hydrogen turbofan is defined in [7] as an adaptation of a CFM56-sized model, operating on hydrogen with assumed technological advancements projected by the 2035 timeframe. The surrogate model is constructed using MOTIVATION, based on performance results obtained from the pyCycle platform.

The integration of hydrogen tanks into the airframe represents a critical design step. Fuselage-integrated tanks are preferred [19] due to their favorable area-to-volume ratio, which enhances thermal efficiency and minimizes weight penalties. However, this configuration reduces the passenger cabin length (l_{cabin}), thereby limiting payload capacity, a significant economic challenge for the aviation industry. A potential mitigation strategy is to elongate the fuselage to accommodate the hydrogen tank while preserving payload capacity, as illustrated in Fig. 3. Here, l_{fus} denotes the total fuselage length, and l_{LH2} represents the section allocated for hydrogen storage.

Variations in l_{fus} directly impact the OEW and the aircraft aerodynamic polar. More critically, positioning a heavy tank at the rear of the fuselage introduces substantial challenges for aircraft stability, particularly affecting longitudinal stability and control.

3.1.1 Single rear tank configuration

A single rear-mounted tank is proposed as, theoretically, the most efficient configuration. Increasing the tank size in this manner, reduces the surface-to-volume ratio, thereby enhancing insulation and thermal inertia [35].

The combined weight of the liquid hydrogen and its tank, W_{TLH2}, is a key design variable in the optimization process. The required fuselage length to accommodate the tank, l_{LH2}, is derived in Eq. 12. The occupied volume, V_{TLH2}, is calculated from W_{TLH2} using the gravimetric index η_{grav} and the liquid hydrogen density ρ_{LH2}. This calculation also accounts for the initial tank fill fraction η_{fill} and volume losses due to subsystems and structural components, η_{st}. The resulting volume extends within the internal fuselage section, with a radius R smaller than that of the fuselage.

$$l_{LH2} = \frac{V_{TLH2}}{\pi \cdot R^2} = \frac{W_{TLH2} \cdot \eta_{grav}}{\rho_{LH2} \cdot \eta_{fill} \cdot \eta_{st} \cdot \pi \cdot R^2} \tag{12}$$

The maximum payload capacity of this aircraft is expressed as a linear function of the cabin length, as shown below (see also Fig. 3):

$$\text{Payload}_{\text{max}} = \text{Payload}_{\text{maxA320neo}} + \frac{\text{Payload}_{\text{maxA320neo}}}{l_{cabin0}} \cdot (\Delta l_{fus} - l_{LH2}) \tag{13}$$

3.1.2 Double tank configuration

Airworthiness and stability constraints may necessitate the addition of a forward fuel tank located between the cockpit and the cabin. Incorporating redundant systems is

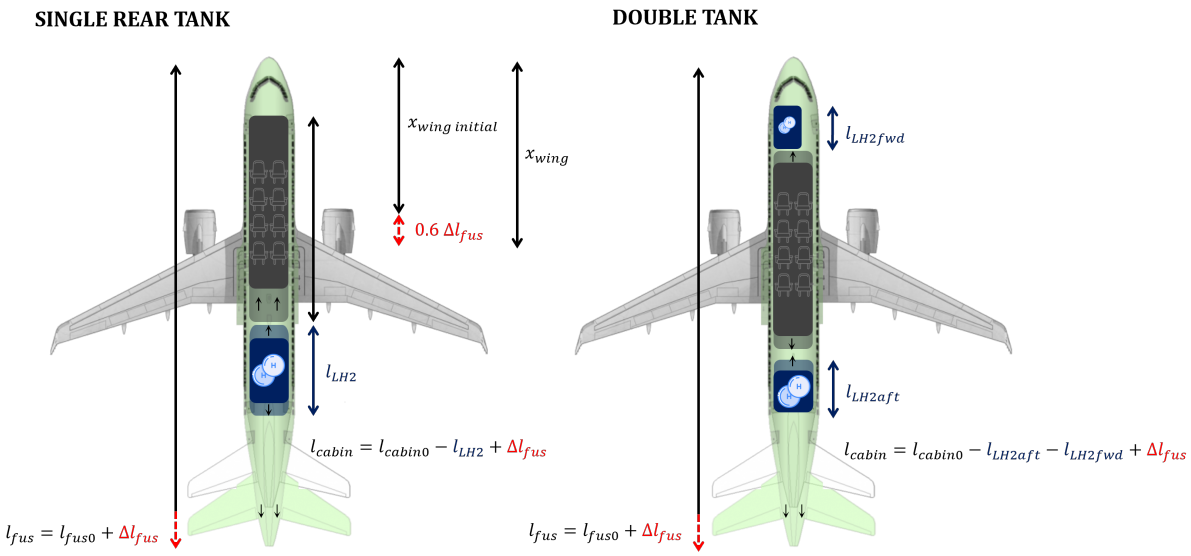


Figure 3: Single tanks (left) and double tank (right) configurations scheme

essential for safety, while managing fuel distribution between tanks is critical for maintaining balance. This forward tank is designed with a smaller radius to allow sufficient space for a corridor. For this configuration, Eqs.12 and 13 are modified accordingly, as can be deduced from Fig.3.

A new parameter, termed hydrogen feeding ratio ($H2FT$), is introduced to represent the split of mass flow from each tank to the engines. This parameter is used to control the depletion rate of the tanks, enabling continuous monitoring of balance throughout the flight.

$$H2FT = \frac{\dot{m}_{H2aft}}{\dot{m}_{H2aft} + \dot{m}_{H2fwd}} \tag{14}$$

3.2 Optimization campaign

The H2 fuel consumption during a nominal trip is the selected objective for minimization. The primary design variables include the weight of the tanks, comprising both the

Objective function
f = Fuel consumption during nominal trip

Table 2: Selected objective function

structural weight and the mass of the stored liquid hydrogen (W_{TLH2} or $W_{TLH2fwd}, W_{TLH2aft}$), as well as the fuselage length (l_{fus}).

The position of the wing shifts proportionally with a variation in l_{fus}, with the wing displacement amounting to 60% Δl_{fus}. However, a variation in the initial wing position ($x_{\text{wing initial}}$) is introduced as an additional degree of freedom to accommodate the stability constraint (see Figure 3). The actual wing position is given by

$$x_{\text{wing}} = x_{\text{wing initial}} + 0.6 \cdot \Delta l_{\text{fus}} \tag{15}$$

This also implicitly affects the positioning of the main landing gear and engines.

Given that the MDA does not ensure weight consistency, the takeoff weight (TOW) is introduced as an additional synthetic design variable. TOW is restricted by the maximum takeoff weight (MTOW = 79 tons) and must align with the sum of all aircraft weight components ($R_{TOW} = 0$).

The payload is constrained by an upper bound corresponding to the maximum payload ($Payload_{max}$) defined in Eq.13.

Performance constraints ensure takeoff and landing field lengths (TOFL, LFL) do not exceed those of the A320neo, while climb gradients (γ_i) comply with airworthiness regulations [3]. A minimum hydrogen reserve ($H2RES_{\text{end}}$ or $H2RES_{\text{end fwd}}, H2RES_{\text{end aft}}$) of 2% is mandated for safety at the end of an extended mission. Additionally, maximum throttle limits ($\delta_{\text{OEI TO}}$ for one-engine inoperative take-off maneuvers and $\delta_{\text{rest of segments}}$ for other segments) are imposed.

Table 3 summarizes the optimization parameters.

Design Variables			
Variable	**Units**	**Min**	**Max**
Single Tank: W_{TLH2} Double Tank: $W_{TLH2,fwd}, W_{TLH2,aft}$	kg	10	12000
TOW	kg	40000	-
l_{fus}	m	36.57	-
$x_{\text{wing initial}}$	m	11	16
Double Tank: $H2FT$	-	0	1
Constraints			
Constraint	**Units**	**Min**	**Max**
Single tank: $R_{TOW} = TOW - W_{TLH2} - OEW - Payload$ Double tank: $R_{TOW} = TOW - W_{TLH2,fwd} - W_{TLH2,aft} - OEW - Payload$	kg	0	0.1
$\Delta TOW = MTOW - TOW$	kg	0	-
$\Delta Payload = Payload_{max} - Payload$	kg	0	-
TOFL	m	-	1828
γ_i	-	[0% - 3.2%] [3]	-
LFL	m	-	1400
Single Tank: $H2RES_{\text{end}}$ Double Tank: $H2RES_{\text{end fwd}}, H2RES_{\text{end aft}}$	-	0.02	-
$\delta_{\text{OEI TO}}$	-	-	1.0
$\delta_{\text{rest of segments}}$	-	-	0.9
SM	-	0.05	0.30

Table 3: Optimization: Design Variables and Constraints

4 RESULTS AND DISCUSSION

Simulations were performed over a range of mission profiles with a cruise altitude of $h = 34,000$ ft, ranges varying from 1,000 to 2,500 nautical miles (NM), and payloads ranging from 14 to 20 tons. The mean gravimetric index (η_{grav}) of the cryogenic hydrogen tanks was assumed to be 0.5. The performance of the three configurations was compared: the conventional A320neo, the hydrogen-retrofitted single rear tank variant (AH_2), and the hydrogen-retrofitted dual-tank configuration (AH_2DT).

Fig. 4 presents the TOW computed by the optimization across various range-payload conditions. The higher energy density of hydrogen compared to Jet-A, coupled with the increased efficiency of hydrogen-powered engines, enhances the operational range of the H_2-retrofitted aircraft. For the conventional A320neo, feasible mission profiles are

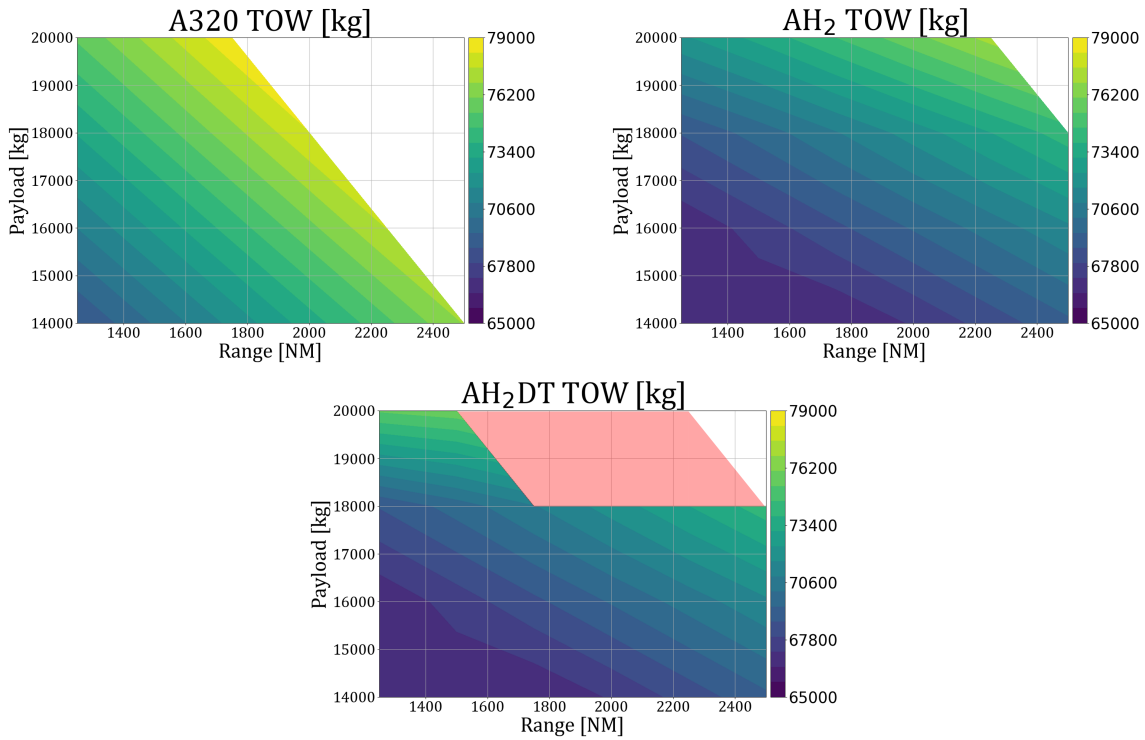


Figure 4: TOW over the range-payload diagram for the different solutions.

constrained by the MTOW of 79 tons. For the AH_2, the operational limit is reached when W_{TLH2} approaches its upper bound. Missing data points for the AH_2DT configuration, which exhibits an irregular pattern before reaching its expected operational limit, are attributed to unidentified numerical errors encountered during the simulations.

Considering block fuel, a direct mass comparison between consumed kerosene and hydrogen is not meaningful due to their differing Low Heating Values (LHVs) (see Table 1). An energy-based comparison reveals that the H_2-optimized model achieves over a 15% reduction in energy consumption at each point compared to the A320neo model, as shown in Fig.5. The non-uniform reduction in energy consumption across the payload-range diagram stems from the ability of the AH_2 configuration to optimize engine and fuselage design for specific payload and range combinations. The payload-range regions resulting in the highest energy reduction, approximately 17% as shown in Fig.5, are associated with shorter fuselage designs (see Fig.7). This drag reduction significantly impacts performance, particularly over longer distances.

Fig.6 shows the consumed fuel cost difference as function of range and payload. It was assumed a 5.60\$/kg price for H_2 and 0.77\$/kg for Jet A [26]. Despite the lower consumption, the significantly higher costs associated with hydrogen production and transportation result in the fuel cost per flight being more than twice the one of Jet-A fuel.

4.1 H2 aircraft geometrical configurations

The structural adaptations of the AH_2 and AH_2DT models are analyzed, revealing similar trends in both configurations. The fuselage length (l_{fus}) varies to accommodate payload requirements, reaching a minimum when payload capacity is exceeded, as illustrated in Fig.7 and Fig.8.

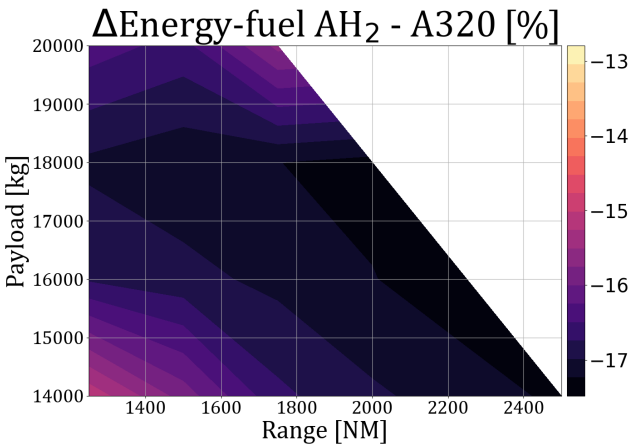


Figure 5: Fuel consumption difference in energetic terms for AH_2 and A320neo models.

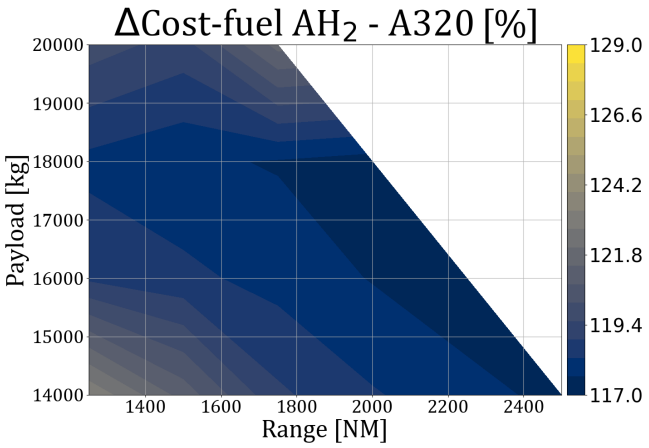


Figure 6: Cost of consumed fuel for AH_2 and A320neo models.

As explained in Section 3.2, increasing the fuselage length proportionally shifts the wing position. The design variable $x_{\text{wing initial}}$ allows for a relative adjustment of the wing position to satisfy stability constraints. Fig.7 and Fig.8 illustrate a rearward shift of the wing position resulting from the increased weight of the hydrogen tanks. This effect, combined with fuselage length contraction at lower payloads, causes $x_{\text{wing initial}}$ to adjust accordingly, compensating for the tank size increase at longer ranges, as shown in Fig.7.

The AH_2DT configuration consistently favors rear tank expansion, employing the larger cross-sectional area of rear tank to minimize tank cabin space occupation. However, in certain payload-range regions, the optimization results indicate forward tank volume increases, as shown in Fig.8. This results in a smaller displacement of x_{wing} compared to the AH_2 configuration.

4.2 Emissions

Emitted pollutants from the AH_2 aircraft are compared against the A320neo baseline. While hydrogen combustion eliminates direct CO_2 emissions, indirect CO_2 emissions arise from hydrogen production. Based on $EICO_2$ values (Table 1), in an optimistic scenario using green hydrogen, emissions reductions of up to 99% can be achieved (Fig. 9). For

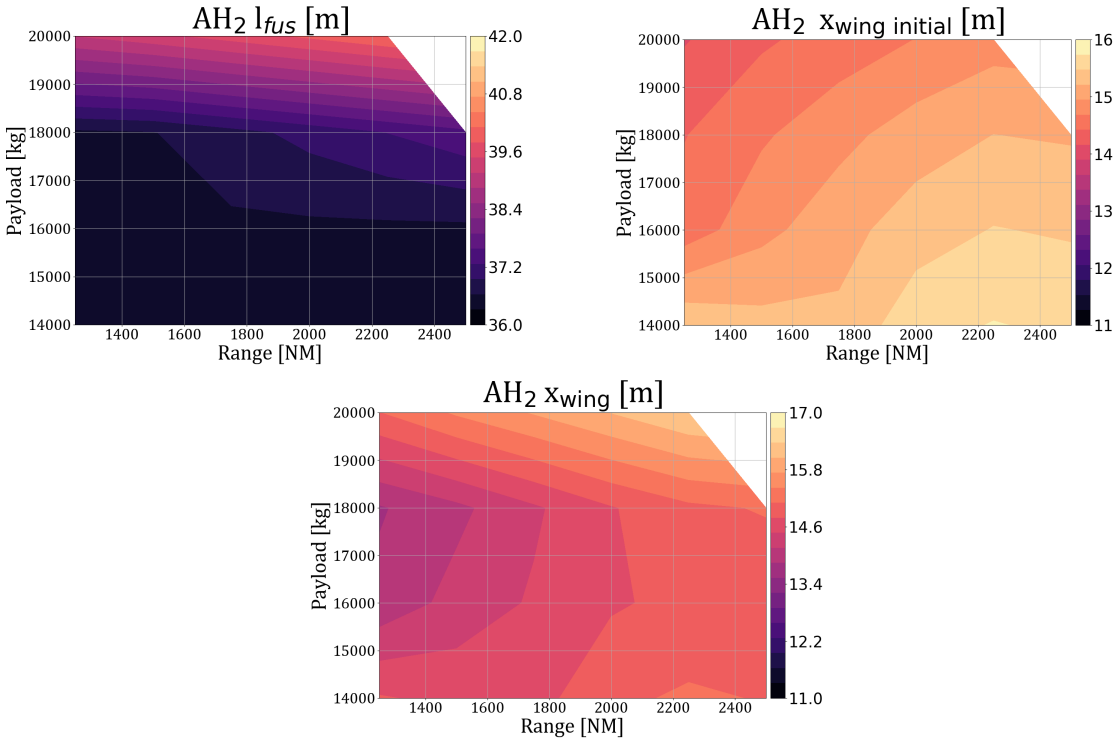


Figure 7: AH_2: (top-left) fuselage length, (top-right) wing relative aft shift, (bottom) wing absolute position.

blue hydrogen, this reduction is limited to approximately 70%.

Additionally, Fig. 10 shows a reduction of more than 60% in emitted NO_x. However, as expected, hydrogen combustion results in increased H_2O emissions, reaching up to twice the amount produced by conventional kerosene combustion.

5 CONCLUSIONS AND FUTURE WORK

This study employed the MOTIVATION platform to assess the feasibility of hydrogen-retrofitted Airbus A320neo variants featuring single and double tank configurations. The results indicate that hydrogen-powered configurations can achieve notable energy efficiency improvements and emissions reductions, particularly when using green hydrogen. Energy consumption was reduced by over 15% compared to the conventional A320neo across a range of payload and mission profiles. Furthermore, green hydrogen effectively eliminated direct CO_2 emissions while reducing NO_x emissions by more than 60%. However, the increase in water vapor emissions requires further investigation due to its potential impact on contrail formation and climate forcing.

Future work will focus on improving numerical stability, especially for dual-tank configurations, integrating more accurate models for non-CO_2 climate impacts and assessing the the impact of gravimetric index (η_{grav}) on overall system performance.

ACKNOWLEDGEMENTS

The research work is carried out as part of the research project OPERATIONAL (TED2021-131119A-I00), funded by MCIN/AEI/10.13039/501100011033 and the European Union NextGenerationEU/PRTR

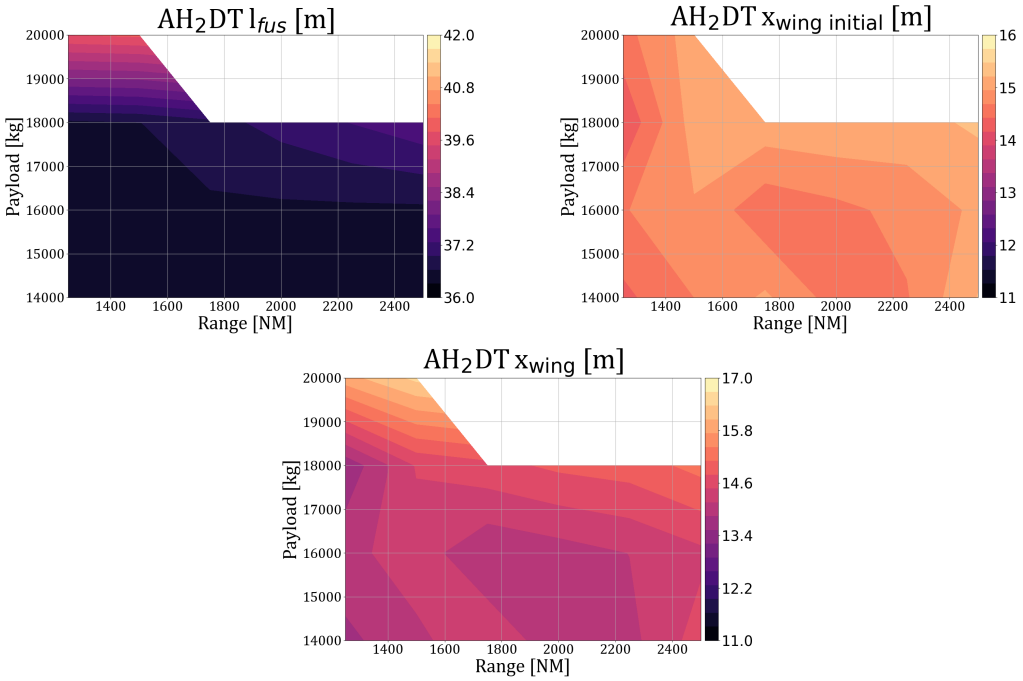


Figure 8: AH_2DT: (top-left) fuselage length, (top-right) wing relative aft shift, (bottom) wing absolute position.

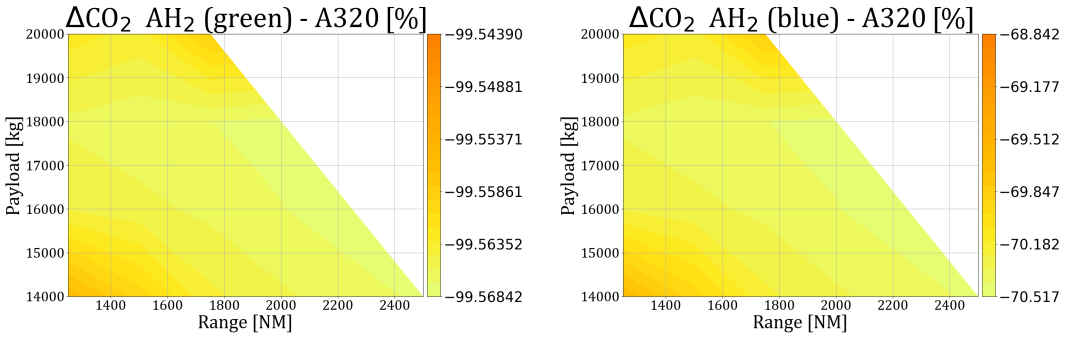


Figure 9: Difference in CO_2 emissions between the AH_2 and A320neo aircraft using green hydrogen (left) and blue hydrogen (right).

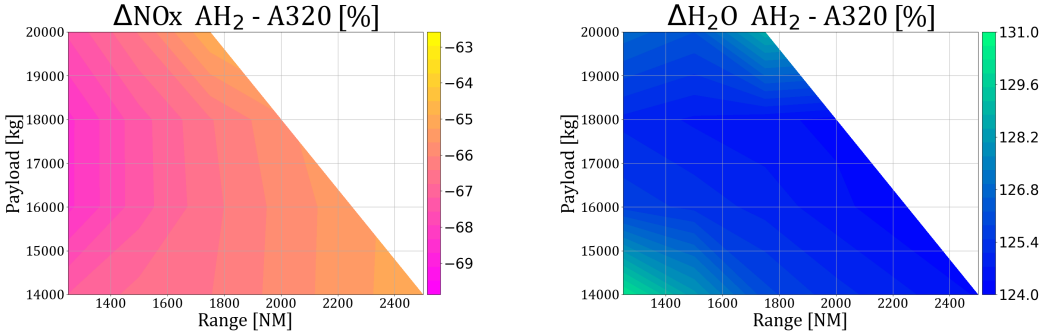


Figure 10: Differences in NO_x (left) and H_2O (right) emissions between the AH_2 and A320neo aircraft.

REFERENCES

[1] D. Lee, D. Fahey, A. Skowron, M. Allen, U. Burkhardt, Q. Chen, S. Doherty, S. Freeman, P. Forster, J. Fuglestvedt, A. Gettelman, R. De León, L. Lim, M. Lund, R. Millar, B. Owen, J. Penner, G. Pitari, M. Prather, R. Sausen, and L. Wilcox. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244:117834, Jan. 2021. ISSN 1352-2310. doi:10.1016/j.atmosenv.2020.117834.

[2] E. J. Adler and J. R. Martins. Hydrogen-powered aircraft: Fundamental concepts, key technologies, and environmental impacts. *Progress in Aerospace Sciences*, 141: 100922, Aug. 2023. ISSN 0376-0421. doi:10.1016/j.paerosci.2023.100922.

[3] R. Quiben Figueroa, R. Cavallaro, and A. Cini. Feasibility studies on regional aircraft retrofitted with hybrid-electric powertrains. *Aerospace Science and Technology*, 151: 109246, Aug. 2024. ISSN 1270-9638. doi:10.1016/j.ast.2024.109246.

[4] B. J. Brelje and J. R. Martins. Development of a conceptual design model for aircraft electric propulsion with efficient gradients. In *2018 AIAA/IEEE Electric Aircraft Technologies Symposium*. American Institute of Aeronautics and Astronautics, July 2018. doi:10.2514/6.2018-4979.

[5] D. I. Barton, C. A. Hall, and M. K. Oldfield. Design of a hydrogen aircraft for zero persistent contrails. *Aerospace*, 10(8):688, July 2023. ISSN 2226-4310. doi:10.3390/aerospace10080688.

[6] D. Silberhorn, K. Dahlmann, A. Görtz, F. Linke, J. Zanger, B. Rauch, T. Methling, C. Janzer, and J. Hartmann. Climate impact reduction potentials of synthetic kerosene and green hydrogen powered mid-range aircraft concepts. *Applied Sciences*, 12(12):5950, June 2022. ISSN 2076-3417. doi:10.3390/app12125950.

[7] R. Quibén Figueroa, R. Cavallaro, A. Cini, and M. Soler Arnedo. Motivation (mdao for sustainable aviation) - framework development for the design and optimization of h2 powered aircraft. In *AIAA AVIATION 2023 Forum*. American Institute of Aeronautics and Astronautics, June 2023. doi:10.2514/6.2023-4151.

[8] Y. Guo, J. Chow, H. Chen, X. Qiao, and C. Fan. Conceptual design and optimization of a hydrogen-powered aircraft with a dorsal tank layout. *Journal of Physics: Conference Series*, 2882(1):012022, Nov. 2024. ISSN 1742-6596. doi:10.1088/1742-6596/2882/1/012022.

[9] E. J. Adler and J. R. R. A. Martins. Blended wing body configuration for hydrogen-powered aviation. *Journal of Aircraft*, 61(3):887–901, May 2024. ISSN 1533-3868. doi:10.2514/1.c037582.

[10] G. Palaia, K. Abu Salem, and E. Carrera. Preliminary performance analysis of medium-range liquid hydrogen-powered box-wing aircraft. *Aerospace*, 11(5):379, May 2024. ISSN 2226-4310. doi:10.3390/aerospace11050379.

[11] P. Norczyk, A. Cini, and R. Cavallaro. Local air quality and noise improvements via optimization of strut-braced wings with distributed electric propulsion. Number AIAA 2025-0456. SciTech2025, Orlando, Florida, 6-10 January 2025. doi:10.2514/6.2025-0456.

[12] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor. Openmdao: an open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, Mar. 2019. ISSN 1615-1488. doi:10.1007/s00158-019-02211-z.

[13] E. U. A. S. A. (EASA). *Certification Specifications for Large Aeroplanes (CS-25)*. European Union Aviation Safety Agency (EASA), Cologne, Germany, latest revision edition, 2025. URL `https://www.easa.europa.eu/en/document-library/certification-specifications/group/cs-25-large-aeroplanes#cs-25-large-aeroplanes`. Accessed: 2025-03-26.

[14] D. P. Raymer. *Aircraft Design a Conceptual Approach*. American Institute of Aeronautics Astronautics, 2018. ISBN 9781624104909. doi:10.2514/4.104909.

[15] E. Obert. Calculation of the oswald factor for clean configuration using obert's method. Technical Report 19930081097, NASA Technical Reports Server (NTRS), 1993. URL `https://ntrs.nasa.gov/citations/19930081097`. Accessed: 2025-03-26.

[16] D. Poll and U. Schumann. An estimation method for the fuel burn and other performance characteristics of civil transport aircraft during cruise: part 2, determining the aircraft's characteristic parameters. *The Aeronautical Journal*, 125(1284):296–340, Dec. 2020. ISSN 2059-6464. doi:10.1017/aer.2020.124.

[17] E. E. Centre. Base of aircraft data (bada). Technical report, EUROCONTROL, 2024. URL `https://www.eurocontrol.int/model/bada`. Last accessed: March 2025.

[18] D. Verstraete, P. Hendrick, P. Pilidis, and K. Ramsden. Hydrogen fuel tanks for subsonic transport aircraft. *International Journal of Hydrogen Energy*, 35(20):11085–11098, Oct. 2010. ISSN 0360-3199. doi:10.1016/j.ijhydene.2010.06.060.

[19] D. Verstraete. Long range transport aircraft using hydrogen fuel. *International Journal of Hydrogen Energy*, 38(34):14824–14831, Nov. 2013. ISSN 0360-3199. doi:10.1016/j.ijhydene.2013.09.021.

[20] C. Winnefeld, T. Kadyk, B. Bensmann, U. Krewer, and R. Hanke-Rauschenbach. Modelling and designing cryogenic hydrogen tanks for future aircraft applications. *Energies*, 11(1):105, Jan. 2018. ISSN 1996-1073. doi:10.3390/en11010105.

[21] C. S. Lin, N. T. Van Dresar, and M. M. Hasan. Pressure control analysis of cryogenic storage systems. *Journal of Propulsion and Power*, 20(3):480–485, May 2004. ISSN 1533-3876. doi:10.2514/1.10387.

[22] E. S. Hendricks and J. S. Gray. pycycle: A tool for efficient optimization of gas turbine engine cycles. *Aerospace*, 6(8):87, Aug. 2019. ISSN 2226-4310. doi:10.3390/aerospace6080087.

[23] L. J. Stauber and C. G. Naiman. Numerical propulsion system simulation (npss): An award winning propulsion system simulation tool. Technical report, NASA Glenn Research Center, 2002. URL https://ntrs.nasa.gov/citations/20050214739. Last accessed: March 2025.

[24] G. Wahba. *Spline Models for Observational Data*. Society for Industrial and Applied Mathematics, Jan. 1990. ISBN 9781611970128. doi:10.1137/1.9781611970128.

[25] S. A. Brandt, R. J. Stiles, J. J. Bertin, and R. Whitford. *Introduction to Aeronautics: A Design Perspective*. AIAA, 2004. ISBN 978-1563476146.

[26] J. Mukhopadhaya, D. Rutherford, L. Ranasinghe, D. Morgan, O. Weiß, F. Keles, F. Troeltsch, A. Murphy, V. Simon, G. Flis, S. Searle, Y. Zhou, M. Bernard, S. Zheng, B. Graver, and G. Gardner. Performance analysis of evolutionary hydrogen-powered aircraft. 2022.

[27] E. U. A. S. A. (EASA). Icao aircraft engine emissions databank, 2025. URL https://www.easa.europa.eu/en/domains/environment/icao-aircraft-engine-emissions-databank. Last accessed: March 2025.

[28] D. DuBois and G. C. Paynter. Fuel flow method2 for estimating aircraft emissions. *Journal of Aircraft*, 38(5):791–795, 2001. doi:10.2514/2.3020.

[29] N. Chandrasekaran and A. Guha. Study of prediction methods for nox emission from turbofan engines. *Journal of Propulsion and Power*, 28(1):170–180, Jan. 2012. ISSN 1533-3876. doi:10.2514/1.b34245.

[30] L. E. Bakken and L. Skogly. Parametric modeling of exhaust gas emission from natural gas fired gas turbines. *Journal of Engineering for Gas Turbines and Power*, 118(3):553–560, July 1996. ISSN 1528-8919. doi:10.1115/1.2816683.

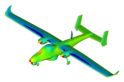
[31] J. R. A. Martins and A. Ning. *Engineering Design Optimization*. University of Cambridge ESOL Examinations, 2021. ISBN 9781108833417.

[32] Airbus. A320 family - passenger aircraft. https://www.airbus.com/en/products-services/commercial-aircraft/passenger-aircraft/a320-family, 2024. Accessed: 2024-03-30.

[33] A. S.A.S. *A320 Aircraft Characteristics: Airport and Maintenance Planning*. AIRBUS S.A.S. Customer Services, Technical Data Support and Service, rev: mar 01/22 edition, 2022. Original issue: Sep 30, 1985.

[34] M. Lab. Openconcept: A framework for conceptual aircraft design. https://github.com/mdolab/openconcept, 2024. Accessed: 2024-03-30.

[35] D. Verstraete. On the energy efficiency of hydrogen-fuelled transport aircraft. *International Journal of Hydrogen Energy*, 40(23):7388–7394, June 2015. ISSN 0360-3199. doi:10.1016/j.ijhydene.2015.04.055.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

EFFECTS OF TRAILING EDGE MODIFICATIONS ON NEGATIVE LIFT CHARACTERISTICS OF SYMMETRICAL AIRFOILS AT TRANSITIONAL LOW REYNOLDS NUMBERS

Pranesh Chandrasekaran[1*], Emre Güler[1] and Melike Nikbay[1]

1: AeroMDO Lab, Faculty of Aeronautics and Astronautics
Istanbul Technical University
ITU Ayazaga Campus, 34469, Maslak/Istanbul
chandrasekaran@itu.edu.tr, gulere22@itu.edu.tr, nikbay@itu.edu.tr

Abstract. *Numerical investigations on the flow over blunt trailing edge airfoils, obtained by symmetrically increasing the trailing edge thickness of NACA 0012 and NACA 0015 airfoils, were carried out for various angles of attack at a chord-based Reynolds number (Re) of 40,000 and at a turbulence intensity (Tu) of 0.1%. The aim of the present work is to investigate the effects of trailing edge modifications on the aerodynamic characteristics, particularly the non-intuitive negative lift phenomenon, of symmetrical NACA airfoils at transitional low Reynolds number regime. The geometries were generated and initially analyzed using XFLR5, a low-fidelity tool that employs integral boundary layer formulation and is capable of modeling flow transition. These preliminary results are then compared with high-fidelity results obtained using ANSYS FLUENT, allowing for an in-depth assessment of the aerodynamic performance. A detailed analysis of the flow field was conducted by plotting the pressure distribution on the airfoil surface and visualizing the streamlines around these airfoils. The locations of separation, transition and reattachment are also presented to facilitate the identification of various flow regimes, thereby offering valuable insights into the unusual aerodynamic behavior of these airfoils. Results indicated that the separated region over the upper surface of the airfoils created a negative camber near the trailing edge, leading to negative lift at low positive angles of attack (α). In the case of NACA 0012 airfoil, increasing the trailing edge thickness resulted in a decrease in the magnitude of negative lift and a reduction in the range of α at which such a phenomenon occurred. Conversely, for the NACA 0015 airfoil, the opposite effect was observed. These findings are non-intuitive and highlight the intricate role of thickness distribution on the aerodynamic characteristics of airfoils at transitional low Reynolds numbers.*

Keywords: negative lift, transitional flow, low Reynolds number, trailing-edge, NACA 0012, XFLR5

1 INTRODUCTION

The study of aerodynamics at low Reynolds numbers is of great importance in designing small sized flying machines such as a micro-air vehicle (MAV). The lifting surfaces of MAVs are known to suffer significant performance degradation due to their increased tendency to flow separation. At low Reynolds numbers, the airfoils and wings experience a laminar boundary layer, which is highly unstable and may result in flow transition to turbulence, and in formation of complicated flow feature such as a Laminar Separation Bubble (LSB). The presence of LSB, depending on their size and location, may have detrimental effects on the aerodynamic performance. Considerable efforts have been made to utilize active and passive flow control strategies [1, 2], formal optimization methods [3, 4], and boundary layer control via geometric modifications [5], to improve the aerodynamic performance of airfoils and wings operating at transitional low Reynolds numbers (Re $\mathcal{O}$ 10^4–10^5).

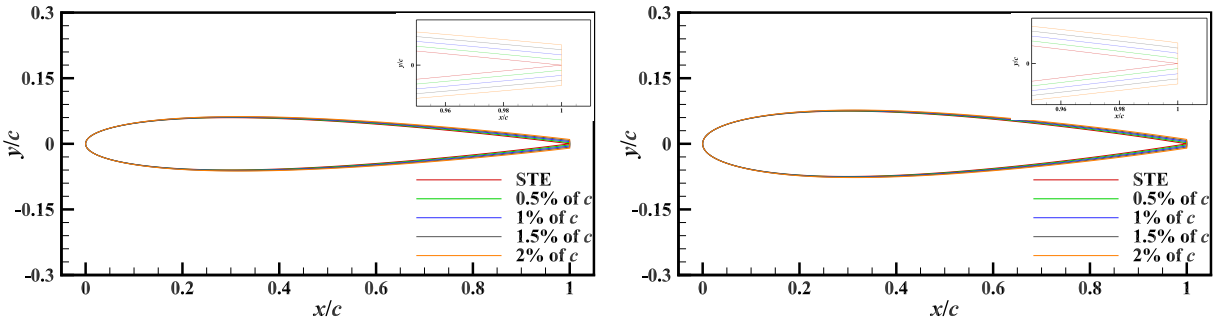
The concept of modifying geometrical parameters for performance enhancement is not new, and has found extensive applications in many engineering sectors, including aviation and energy industries. For instance, morphing wing technologies have shown great potential for reducing aircraft fuel consumption, and therefore represent a promising solution for enhancing overall efficiency [6]. Recently, some investigators, inspired by biomimetics, have proposed a wide range of geometry modifications, such as leading-edge tubercles and wavy surfaces, to harness potential aeroacoustic benefits [7–9]. Wang et al. demonstrated that the use of surface ridges and serrated edges can result in trailing edge noise reduction with very little change in the overall drag [10]. Similarly, the use of blunt trailing edges (also known as flatback profiles) can offer significant structural benefits for small wind turbine applications. The study by Smith and Schaefer in the 1950's at NASA represents one of the earliest investigations on the potential benefits associated with flatback profiles [11]. They found that airfoils with blunt trailing edges resulted in a better maximum lift coefficient compared to those with sharp trailing edges under identical flow conditions.

Investigations on the effect of trailing edge modifications on the aerodynamic performance of lifting surfaces have garnered special interest in the transitional low Re regime. Thomareis and Papadakis [12] performed Direct Numerical Simulations to investigate the flow field around a NACA 0012 airfoil with three different trailing edge shapes (sharp, blunt and serrated) at Re = 50,000 and at two different angles of attack. Leloudas et al. [13] modified the original RG15 airfoil to accommodate increased thickness-to-chord ratio and rounded trailing edge to better suit horizontal-axis wind turbine applications. Xu et al. [14] analyzed a series of blunt trailing edge airfoils at a Reynold number of 3×10^6, and inferred that increasing the thickness of trailing edge extended the linear region of lift and also improved the maximum value of lift coefficients. Similar observations were also made by Pei et al. [15], who investigated the flow around a blunt NWT600 airfoil with different trailing edge thicknesses.

While most of the previous studies on blunt trailing edge airfoils focused on exploring the aerodynamic performance of profiles that are specifically designed for wind turbine applications, there exists a gap in the literature that focuses on challenging NACA profiles, which are predominantly used in aerospace applications. Symmetrical NACA profiles at transitional low Re regime are known to exhibit unusual aerodynamic characteristics, thus making them an intriguing choice [16]. The present work aims to modify the symmetrical NACA profiles to incorporate a blunt trailing edge, compute the flow field around them, and investigate the effect of trailing edge thickness on their aerodynamic performance.

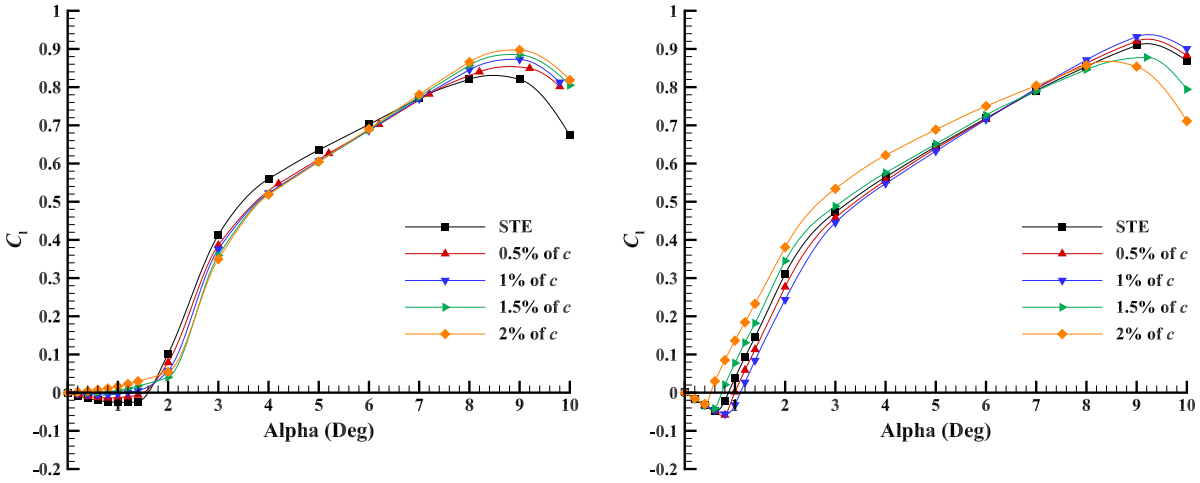
2 GEOMETRIES STUDIED

The blunt trailing edge geometries for the present study were created by symmetrically increasing the trailing edge thickness on both pressure and suction side of the traditional NACA 0012 and NACA 0015 airfoils; see Fig.1. Four trailing edge thicknesses (T_{TE}): 0.5%c, 1%c, 1.5%c, and 2%c were considered, where c is the airfoil chord length. These geometries were generated using XFLR5, an open-source low-fidelity tool that is capable of modeling flow transition [17]. It is a widely used tool for making rapid predictions of aerodynamic coefficients at the transitional low Re regime. Initially, low-fidelity analyses were carried out on all these geometries at Re = 40,000 and at a freestream turbulence intensity Tu of 0.1%. The variation in lift coefficients as a function of angle of attack for both family of airfoils are presented in Fig.2. Some of the general observations include: lift curves are highly nonlinear, negative values of lift coefficients (C_l) are observed at low positive angles of attack, change of sign in C_l, sudden change of slope at moderate angles of attack, and finally a decrease in C_l at high angles of attack.



(a) (b)

Figure 1: Geometries studied; a) NACA 0012 airfoil and its variants and b) NACA 0015 airfoil and its variants. A zoomed in view of the trailing edge is also shown.



(a) (b)

Figure 2: Variation of lift coefficients as a function of angle of attack computed using XFLR5; a) NACA 0012 airfoil and its variants and b) NACA 0015 airfoil and its variants.

With respect to the effect of trailing edge thickness T_{TE} on the lift coefficients of the NACA 0012 family of airfoils, it is observed that increasing T_{TE} leads to a decrease in the magnitude of negative lift and a reduction in the range of α at which such a phenomenon occurred. However, for the NACA 0015 family of airfoils, the opposite effect is observed. Such non-intuitive and interesting observations warrant a detailed investigation of the flow field using high-fidelity CFD simulations, with the aim of understanding the underlying physical mechanisms and gaining deeper insight into the complex flow features associated with these behaviors.

3 COMPUTATIONAL PROCEDURE

The reliability of computational procedure is established by comparing the results from the present computations with the experimental and computational data available in the literature. For this purpose, the flow over Eppler 387 airfoil at $Re = 200{,}000$ and $\alpha = 2°$ is considered. A two-dimensional structured grid with C-type mesh topology is employed. The mesh close to the airfoil geometry is displayed in Fig.3. A no-slip boundary condition is applied on the airfoil surface, while a velocity inlet and a pressure outlet boundary conditions are applied at the domain boundaries. The distance from the airfoil leading edge to domain inlet and domain outlet are set to $10c$ and $20c$, respectively.

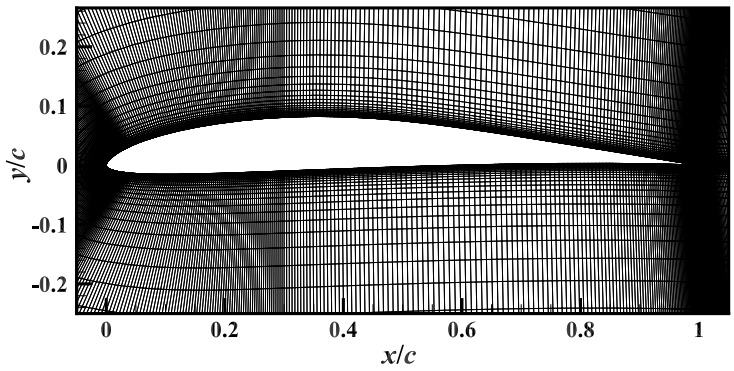


Figure 3: Mesh around the Eppler 387 airfoil.

The high-fidelity numerical simulations were carried out using a commercial CFD solver ANSYS FLUENT. Since the chosen value of the Reynolds number lies within the transitional flow regime, steady-state Reynolds-Averaged Navier Stokes (RANS) simulations were performed using the four-equation transitional k-ω SST model. Further information about this transitional model, also popularly known as γ-Re_θ model, can be found in [18]. A Second-order spatial discretization scheme was employed for both flow and turbulence variables to ensure numerical accuracy. A grid independence study was carried out by performing computations using three different grids with increasing mesh sizes. The values of lift and drag coefficients obtained from these three grids are tabulated in Tab.1. As can be seen, a grid with 100,266 cells had a reasonably low discretization error and was therefore used for all further computations.

Figure 4 presents the C_p distribution over the E387 airfoil at the chosen flow conditions. It may be noted that the results from the present work is in excellent agreement with the experimental [19] and the computational results [20] available in the literature, thus establishing the reliability of the present computational procedure.

Table 1: Grid Independence Study; E387 Airfoil at Re = 200,000 and Tu = 0.1%.

Grid	No. of. Cells	C_l	Absolute Difference	C_d	Absolute Difference
G1	43,996	0.5789	0.0026	0.0116	0.0001
G2	100,266	0.5815	0.0003	0.0117	0.0001
G3	150,654	0.5818		0.0116	

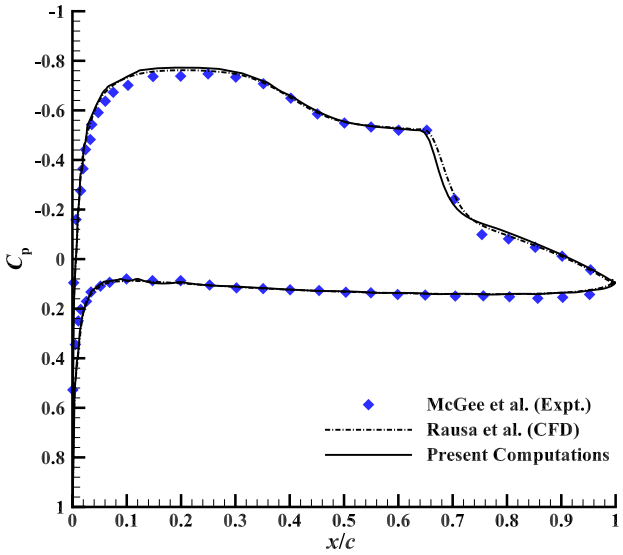


Figure 4: Comparison of C_p Distribution over the E387 airfoil at Re = 200,000 and Tu = 0.1%.

4 RESULTS AND DISCUSSION

4.1 NACA 0012 Airfoil

The flow around a NACA 0012 airfoil with a sharp trailing edge (STE) is computed at Re = 40,000 and at a Tu = 0.1%. Figure 5 illustrates the variation of lift coefficients as a function of angle of attack; experimental results from Ohtake et al. [21] and low-fidelity results from XFLR5 are also plotted for comparison. The high-fidelity results are seen to be in better agreement with the experimental data than the low-fidelity results, which consistently overestimate the value of lift coefficients. Interestingly, both low- and high-fidelity simulations were able to predict the negative values of lift coefficients as observed in the experiments at low positive angles of attack. An investigation of streamlines near the trailing edge at these angles revealed that the flow separates from the lower surface, turns around the trailing edge, and flows into the separated region on the upper surface. The thickness of the separated region on the upper surface was found to be greater than that on the lower surface, effectively creating a negative camber that results in negative lift. In such cases, contrary to typical behavior, the pressure on the upper surface is observed to be greater than that on the lower surface for a significant part of the chord length, thus resulting in negative values of lift coefficients; see Fig. 6. Previous studies on symmetrical airfoils at low Reynolds numbers have also reported this negative lift phenomenon at low positive angles of attack [22, 23]. As noted by Tank et al. [24] and Chen et al. [25], achieving good agreement with experimental results remains to be challenging at high α.

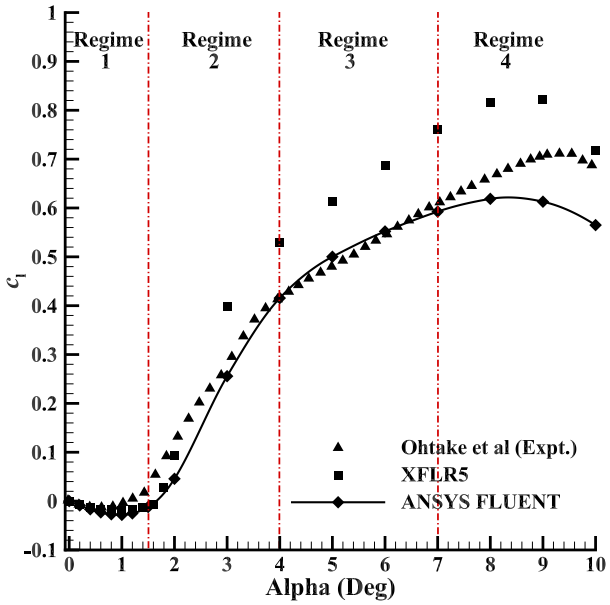
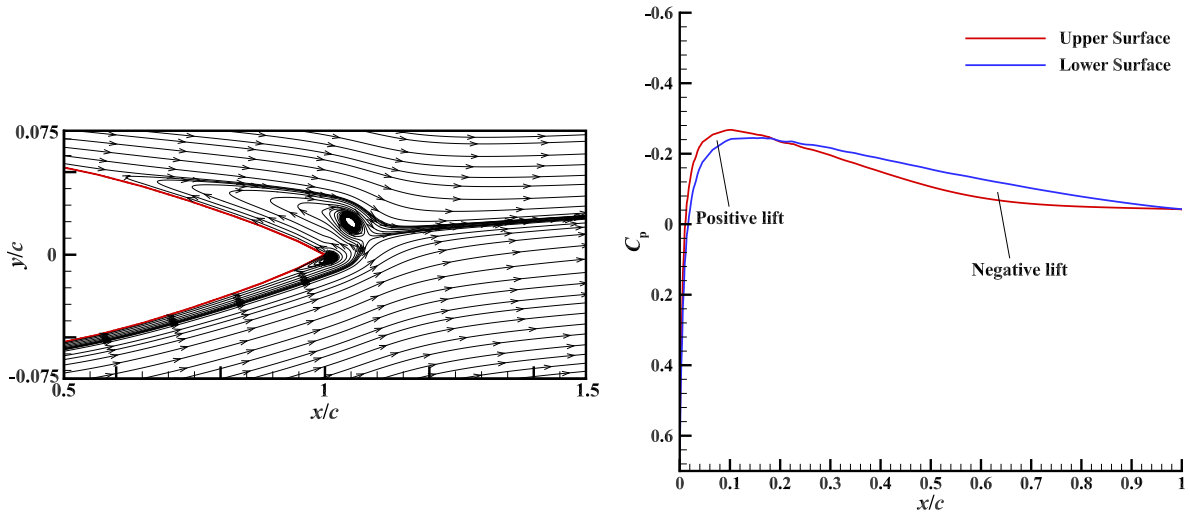


Figure 5: Variation of lift coefficients as a function of angle of attack for NACA 0012 airfoil with a sharp trailing edge; $Re = 40{,}000$ and $Tu = 0.1\%$.



(a) (b)

Figure 6: a) Streamlines close to the trailing edge and b) C_p distribution on the NACA 0012 airfoil surface; $\alpha = 1°$, $Re = 40{,}000$ and $Tu = 0.1\%$.

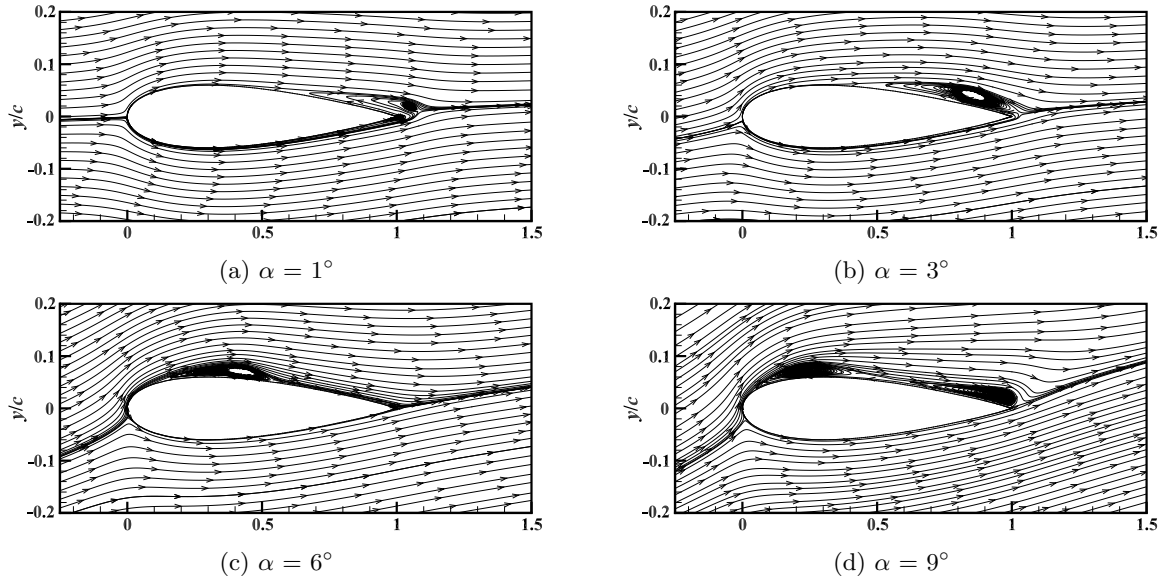
The locations of separation (x_s), transition (x_t), and reattachment (x_r) on the upper surface of the airfoil are summarized in Tab.2, revealing four distinct flow regimes. Flow separation occurs at an angle of attack as low as 0°, with the separation point on the upper surface observed at $x/c = 0.68$ for $\alpha = 0°$. As the angle of attack increases, this separation point gradually moves towards the leading edge, while on the lower surface it shifts towards the trailing edge. Up to $\alpha = 1.5°$, flow separation occurs without transitioning to turbulence or reattachment. It is in this range of α where the negative values of C_l are observed. In the range of $2° \leq \alpha \leq 4°$, the separated flow transitions to turbulence but does not reattach. For angles between 5° and 7°, the separated flow not only

transitions to turbulence but also reattaches resulting in the formation of LSB. Additionally, a turbulent recirculation region is also observed near the trailing edge, characterized by turbulent separation and turbulent reattachment. Finally, for angles of attack in the range of $8° \leq \alpha \leq 10°$, the leading edge LSB merges with the trailing edge recirculation region resulting in a large separated region that almost cover the entire chord length.

Table 2: Summary of separation (x_s), transition (x_t) and reattachment (x_r) locations on the upper surface of the airfoil; $Re = 40{,}000$ and $Tu = 0.1\%$.

Regime	α	C_l	Flow regime	Laminar separation bubble				Trailing edge recirculation		
				x_s/c	x_t/c	x_r/c	L_{LSB}	x_{sT}/c	x_{sR}/c	L_{TER}
1	0	0.0000	Laminar separation without transition and without reattachment	0.68	-	Flow does not reattach		Turbulent separation was not observed		
	0.2	-0.0088		0.66	-					
	0.4	-0.0168		0.62	-					
	0.6	-0.0231		0.59	-					
	0.8	-0.0270		0.57	-					
	1	-0.0282		0.55	-					
	1.2	-0.0252		0.52	-					
	1.5	-0.0120		0.49	-					
2	2	0.0457	Laminar separation with flow transition without reattachment	0.45	0.99					
	3	0.2557		0.34	0.88					
	4	0.4155		0.24	0.73					
3	5	0.5002	Laminar separation with transition and with reattachment (LSB) and trailing edge recirculation (TER)	0.15	0.58	0.76	0.61	0.99	1.00	0.01
	6	0.5520		0.09	0.45	0.61	0.52	0.96	1.00	0.04
	7	0.5926		0.05	0.34	0.51	0.46	0.88	1.00	0.11
4	8	0.6187	Merging of LSB and TER	0.04	0.26	-	-	-	1.00	-
	9	0.6128		0.03	0.22	-	-	-	1.00	-
	10	0.5650		0.03	0.22	-	-	-	1.00	-

Streamlines over the airfoil at select angles of attack is presented in Fig.7 to highlight the presence of the four distinct flow regimes.



(a) $\alpha = 1°$ (b) $\alpha = 3°$

(c) $\alpha = 6°$ (d) $\alpha = 9°$

Figure 7: Streamlines over the NACA 0012 airfoil at $Re = 40{,}000$ and $Tu = 0.1\%$, highlighting various flow regimes.

4.2 Effect of Trailing Edge Thickness on NACA 0012 Airfoil

To investigate the effect of trailing edge thickness on the aerodynamic characteristics of symmetrical NACA profiles, high-fidelity CFD simulations were carried out on four modified NACA 0012 airfoils, each with a different value of trailing edge thickness (T_{TE}). The variation of lift coefficients as a function of angle of attack for these geometries is illustrated in Fig.8. The lift curve of the traditional NACA 0012 airfoil with a sharp trailing edge (STE) is also plotted for comparison. Based on the impact of T_{TE} on C_l, three different α regimes were identified. The impact is most evident at low (0° to 2°) and high angles of attack (4° to 10°), whereas it is insignificant at moderate angles of attack (2° to 4°). With respect to the effect of T_{TE} on the negative values of lift coefficient, it may be observed that, increasing the trailing edge thickness resulted in a decrease in the magnitude of negative lift and also a reduction in the range of α at which such a phenomenon occurred. At high angles of attack, increasing the trailing edge thickness led to an increase in the maximum lift coefficient. Such an observation was also reported in [14, 15].

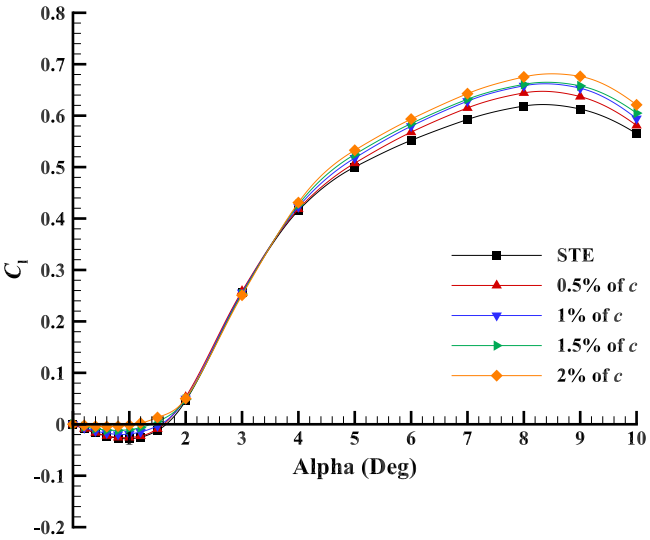


Figure 8: Variation of lift coefficients as a function of angle of attack for NACA 0012 airfoil with sharp and blunt trailing edge shapes; Re = 40,000 and Tu = 0.1%.

A careful examination of C_p and C_f distributions on these airfoils, at various angles of attack, can yield valuable insights into the impact of T_{TE} on both the global aerodynamic properties and the separation characteristics of these geometries. Figure 9 compares the C_p distribution of the NACA 0012 airfoil with a STE to that of an airfoil featuring a blunt trailing edge, which has a thickness of 2% of the chord. This comparison is made at an angle of attack of 1°, where negative values of C_l were observed. It may be noted that the location at which the C_p curves cross each other has shifted aft for the blunt trailing edge geometry; also see Fig.10a. This essentially signifies that the percentage of chord contributing to the negative lift decreases with an increase in trailing edge thickness, thus resulting in a decrease in the magnitude of negative lift. Furthermore, an analysis of the C_f distribution on these geometries revealed that increasing the trailing edge thickness delays flow separation on both upper and lower surfaces; see Fig.11. This delay ultimately leads to improved lift coefficient values compared to the standard sharp trailing edge case.

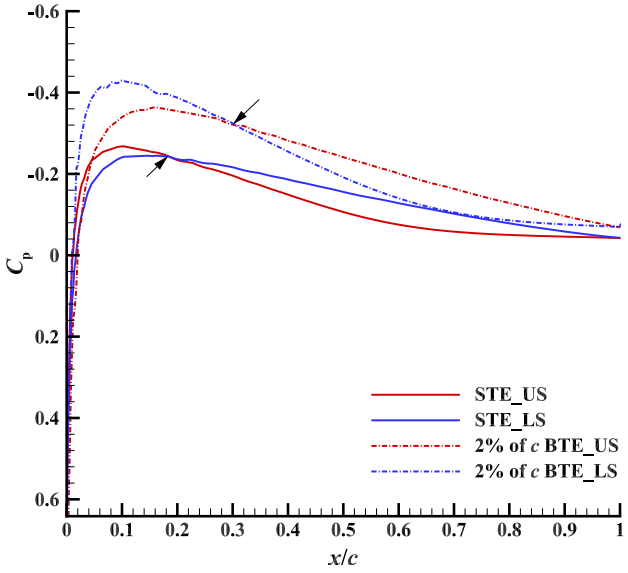
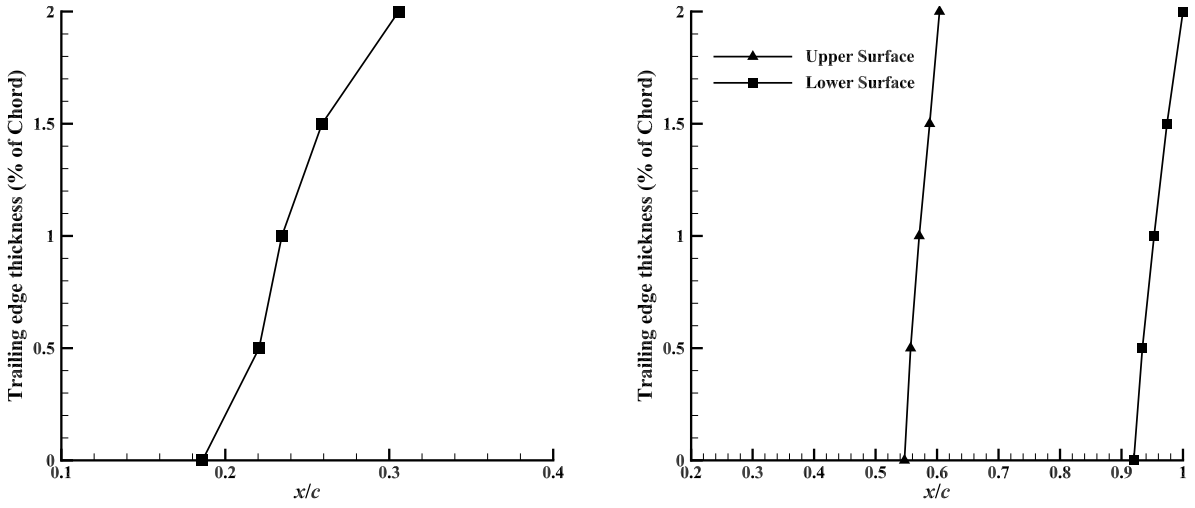


Figure 9: Effect of trailing edge thickness on the C_p distribution of NACA 0012 airfoils with sharp and blunt trailing edge shapes; Re = 40,000 and Tu = 0.1%.



(a) (b)

Figure 10: Effect of trailing edge thickness on a) the chordwise location at which the C_p curves crosses each other and b) the separation location on the upper and lower surface of the airfoils; Re = 40,000 and Tu = 0.1%.

The analysis of flow behavior at $\alpha = 7°$ revealed key insights regarding the separation, transition, and reattachment locations on the upper surface of the geometries considered. It may be inferred from Tab.3 that increasing the value of T_{TE} has led to a slight delay in flow separation and an early reattachment, resulting in a shorter laminar separation bubbles. The transition location remains unaffected to the modifications in the geometry. Furthermore, with respect to the trailing edge recirculation, increasing the T_{TE} delayed turbulent separation, with all cases showing reattachment exactly at the trailing edge. In other words, the length of the trailing edge recirculation region (L_{TER}) decreases with an increase in T_{TE}. Such observations highlight the reasons behind the better lift coefficient values of the blunt tailing edge geometries.

Table 3: Summary of separation characteristics of NACA 0012 airfoils with sharp and blunt trailing edge shapes; $Re = 40{,}000$ and $Tu = 0.1\%$.

Airfoil	Laminar Separation Bubble				Trailing Edge Recirculation		
	x_s	x_t/c	x_r/c	L_{LSB}	x_{sT}/c	x_{sR}/c	L_{TER}
STE	0.049	0.343	0.506	0.457	0.885	0.999	0.114
0.5% of c	0.050	0.322	0.488	0.438	0.898	1.000	0.102
1% of c	0.049	0.322	0.484	0.435	0.935	1.000	0.065
1.5% of c	0.048	0.322	0.482	0.434	0.959	1.000	0.041
2% of c	0.049	0.322	0.481	0.432	0.985	1.000	0.015

4.3 Effect of Trailing Edge Thickness on NACA 0015 Airfoil

In continuation of the analysis presented in Section 4.2, this section delves into the effect of trailing edge thickness on the aerodynamic characteristics of NACA 0015 airfoil. Simulations were performed on four modified NACA 0015 profiles, each featuring a distinct value of T_{TE}. Figure 11 presents the variation of C_l as a function of α for these geometries. The lift curve of the NACA 0015 airfoil with STE is also plotted. It may be inferred that the impact of T_{TE} is felt across the entire range of α considered. While the NACA 0012 airfoil with STE exhibited negative values of C_l at low positive angles of attack, the NACA 0015 airfoil with STE maintains positive C_l throughout the entire range of α considered. Unlike the NACA 0012 family of airfoils, where three distinct α regimes were identified, the NACA 0015 airfoil family, exhibit only two regimes. For α up to 3°, the airfoils with sharp trailing edge demonstrated superior lift coefficients compared to those with blunt trailing edges. Notably, as the trailing edge thickness increases, the lift coefficient decreases, with the airfoil featuring a blunt trailing edge thickness of 2% of c

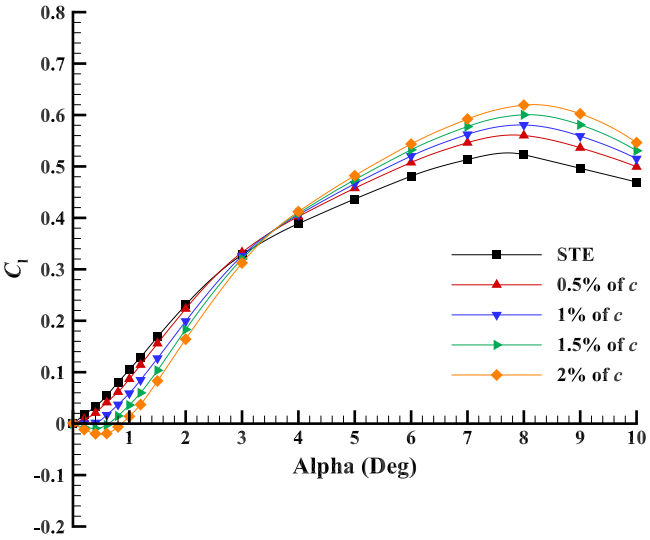
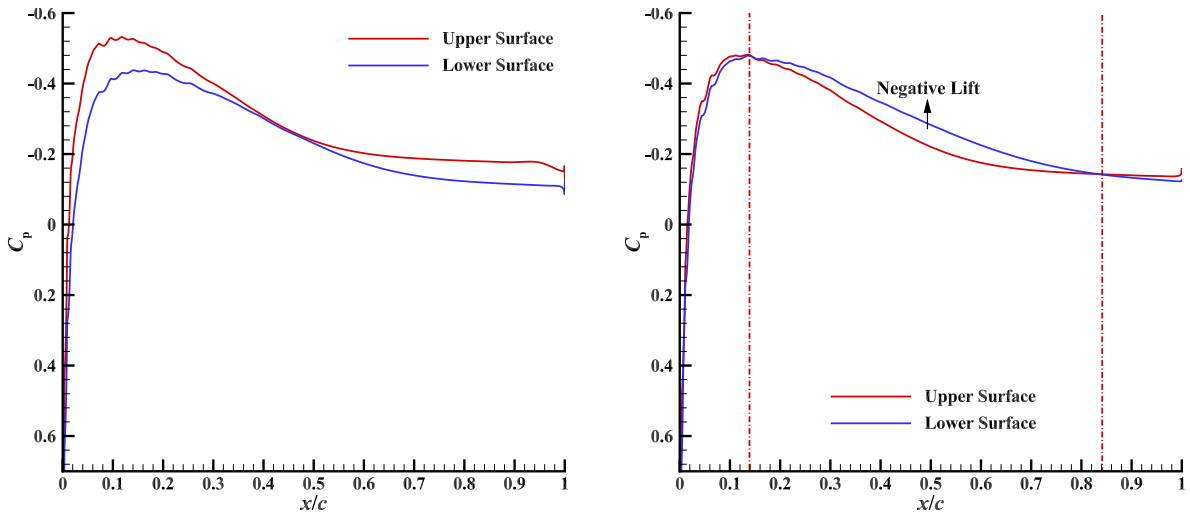


Figure 11: Variation of lift coefficients as a function of angle of attack for NACA 0015 airfoil with sharp and blunt trailing edge shapes; $Re = 40{,}000$ and $Tu = 0.1\%$.

yielding the lowest lift coefficients across this range. Conversely, for $\alpha > 3°$, an opposite trend is observed. Here, increasing the trailing edge thickness results in enhanced lift coefficients, aligning with the findings from the NACA 0012 airfoil family. In both cases, higher the value of T_{TE}, higher is the maximum value of lift coefficients. Such contrasting observations between the two family of airfoils illustrate the complexity of aerodynamic behavior of lifting surfaces at low Reynolds numbers and emphasize that generalizing these findings may not be appropriate for airfoils with different maximum thicknesses. This highlights the necessity for thorough investigations to fully understand the intricate relationships governing the aerodynamic performance across various designs and geometric parameters.

The effect of T_{TE} on the c_p distribution of these geometries at $\alpha = 0.6°$ is shown in Fig.12. It may be noted that the C_p distribution for the upper and lower surfaces do not cross each other for the airfoil with a STE (see Figure 12a), indicating a positive value of C_l as observed in Figure 11. However, for the airfoil with a blunt trailing edge thickness of 2%, the C_p curves cross each other at two locations along the chord, resulting in three distinct regions of lift contribution; the front, middle and rear sections (see Figure 12b). In the cases of 0.5% and 1% thickness (not shown here), the positive lift generated by the front and rear section dominates the negative lift from the middle portion, yielding a net positive C_l. Conversely, for the 1.5% and 2% thickness cases, the negative lift from the middle portion dominates over the positive lift from the front and rear portions, leading to a net negative C_l.



(a) STE

(b) BTE with a thickness of 2%c

Figure 12: Effect of trailing edge thickness on the C_p distribution of NACA 0015 airfoils with sharp and blunt trailing edge shapes; $Re = 40{,}000$ and $Tu = 0.1\%$.

Comparison of C_p distribution of these geometries at $\alpha = 7°$ revealed that an increase in trailing edge thickness resulted in a better suction peak as compared to the airfoil with a sharp trailing edge, thus resulting in a better value of lift coefficients. These observations align with the characteristics observed in the NACA 0012 family of airfoils.

5 CONCLUSION

The impact of trailing edge thickness on the lift characteristics of symmetrical airfoils, specifically the NACA 0012 and NACA 0015 families, was computationally investigated at a transitional low Reynolds number of 40,000 and at a Tu of 0.1%. Four different values of trailing edge thicknesses (T_{TE}): 0.5%c, 1%c, 1.5%c, and 2%c, were considered. The analyses revealed distinct differences in how the trailing edge thickness influenced lift performance across the two families of airfoils. For the NACA 0012 family, there exists a narrow range of α (2° to 4°) where the T_{TE} had minimal impact on the lift coefficients, with blunt trailing edge airfoils performing better as compared to the airfoil with a shape trailing edge across the entire range of α considered. Increasing trailing edge thickness reduced both the magnitude of negative lift and the range of angles at which it occurred. In contrast, the NACA 0015 family exhibited the opposite trend; an airfoil with sharp trailing edge did not produce negative lift at low angles of attack, but increasing the trailing edge thickness led to negative values of lift coefficients. Despite these differences at low angles of attack, both airfoil families showed similar behavior at moderate to high angles of attack, where increasing trailing edge thickness led to an improvement in the maximum value of lift coefficients. This suggests that while the NACA 0012 and NACA 0015 families respond differently to trailing edge modifications, they both benefit from increased thickness in terms of lift performance at higher angles of attack. The present study underscores the importance of careful consideration in generalizing findings, as the aerodynamic performance may vary significantly for airfoils with different maximum thicknesses. It also highlights the need for further investigations to fully understand the complex relationships between geometrical parameters and the aerodynamic performance.

REFERENCES

[1] M. S. Selig, J. Guglielmo, A. Broern, and P. Giguere. Experiments on airfoils at low Reynolds numbers. In *34th Aerospace Sciences Meeting and Exhibit*, page 62, 1996.

[2] A. Gopalarathnam, B. A. Broughton, B. D. McGranahan, and M. S. Selig. Design of low Reynolds number airfoils with trips. *Journal of Aircraft*, 40(4):768–775, 2003.

[3] J. Li, M. Zhang, C. M. J. Tay, N. Liu, Y. Cui, S. C. Chew, and B. C. Khoo. Low-Reynolds-number airfoil design optimization using deep-learning-based tailored airfoil modes. *Aerospace Science and Technology*, 121:107309, 2022.

[4] C. Pranesh, M. Sivapragasam, and H. K. Narahari. Multi-fidelity aerodynamic shape optimization of an airfoil at transitional low Reynolds number. In *National Conference on Multidisciplinary Design, Analysis, and Optimization, IISc, Bangalore*, 2018.

[5] P. Bulat, P. Chernyshov, N. Prodan, and K. Volkov. Control of aerodynamic characteristics of thick airfoils at low Reynolds numbers using methods of boundary layer control. *Fluids*, 9(1):26, 2024.

[6] S. Barbarino, O. Bilgen, R. M. Ajaj, M. I. Friswell, and D. J. Inman. A review of morphing aircraft. *Journal of Intelligent Material Systems and Structures*, 22(9): 823–877, 2011.

[7] D. Miklosovic, M. Murray, L. Howle, and F. Fish. Leading-edge tubercles delay stall on humpback whale (megaptera novaeangliae) flippers. *Physics of Fluids*, 16(5): L39–L42, 2004.

[8] K. L. Hansen, R. M. Kelso, and B. B. Dally. Performance variations of leading-edge tubercles for distinct airfoil profiles. *AIAA Journal*, 49(1):185–194, 2011.

[9] T. Smith and C. Klettner. Airfoil trailing-edge noise and drag reduction at a moderate Reynolds number using wavy geometries. *Physics of Fluids*, 34(11), 2022.

[10] J. Wang, C. Zhang, Z. Wu, J. Wharton, and L. Ren. Numerical study on reduction of aerodynamic noise around an airfoil with biomimetic structures. *Journal of Sound and Vibration*, 394:46–58, 2017.

[11] H. A. Smith and R. F. Schaefer. Aerodynamic characteristics at Reynolds numbers of 3.0 x 10^6 and 6.0 x 10^6 of three airfoil sections formed by cutting off various amounts from the rear portion of the NACA 0012 airfoil section. Technical report, NASA Report, 1950.

[12] N. Thomareis and G. Papadakis. Effect of trailing edge shape on the separated flow characteristics around an airfoil at low Reynolds number: A numerical study. *Physics of Fluids*, 29(1), 2017.

[13] S. N. Leloudas, A. I. Eskantar, G. N. Lygidakis, and I. K. Nikolos. Low Reynolds airfoil family for small horizontal axis wind turbines based on RG15 airfoil. *SN Applied Sciences*, 2(3):371, 2020.

[14] H. Xu, W. Shen, W. Zhu, H. Yang, and C. Liu. Aerodynamic analysis of trailing edge enlarged wind turbine airfoils. In *Journal of Physics: Conference Series*, volume 524, page 012010. IOP Publishing, 2014.

[15] Z. Pei, H.-Y. Xu, L. Deng, and L.-X. Li. Influence of the blunt trailing-edge thickness on the aerodynamic characteristics of the very thick airfoil. *Wind*, 3(4):439–458, 2023.

[16] C. Pranesh, M. Sivapragasam, M. Deshpande, and H. Narahari. Negative lift characteristics of NACA 0012 aerofoil at low Reynolds numbers. *Sādhanā*, 44(1):21, 2019.

[17] A. Deperrois. Analysis of foils and wings operating at low Reynolds numbers: guidelines for xflr5. *Recuperado de https://www. xflr5. tech/xflr5. htm*, 2009.

[18] F. R. Menter, R. B. Langtry, S. Likki, Y. B. Suzen, P. Huang, and S. Völker. A correlation-based transition model using local variables—part i: model formulation. *Journal of turbomachinery*, 128(3):413–422, 2006.

[19] R. J. McGhee. *Experimental results for the Eppler 387 airfoil at low Reynolds numbers in the Langley low-turbulence pressure tunnel*, volume 4062. National Aeronautics and Space Administration, 1988.

[20] A. Rausa, A. Guardone, and F. Auteri. Implementation of γ-Re_θ transition model within SU2: model validation and verification. In *AIAA Scitech 2023 Forum*, page 1570, 2023.

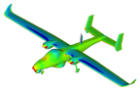
[21] T. Ohtake, Y. Nakae, and T. Motohashi. Nonlinearity of the aerodynamic characteristics of NACA 0012 aerofoil at low Reynolds numbers. *Japan Society of Aeronautical Space Sciences*, 55(644):439–445, 2007.

[22] K. Yonemoto, K. Takato, H. Ochi, and S. Fujie. Kutta condition violation in two-dimensional NACA 0012 airfoil at low Reynolds number. In *26th AIAA Applied Aerodynamics Conference*, page 6399, 2008.

[23] J. Aguilar-Cabello, P. Gutierrez-Castillo, L. Parras, C. Del Pino, and E. Sanmiguel-Rojas. On the onset of negative lift in a symmetric airfoil at very small angles of attack. *Physics of Fluids*, 32(5), 2020.

[24] J. Tank, L. Smith, and G. Spedding. On the possibility (or lack thereof) of agreement between experiment and computation of flows over wings at moderate Reynolds number. *Interface Focus*, 7(1):20160076, 2017.

[25] Z. Chen, N. Qin, and A. Nowakowski. Three-dimensional laminar-separation bubble on a cambered thin wing at low Reynolds numbers. *Journal of Aircraft*, 50(1):152–163, 2013.



AeroBest 2025
Lisboa, 22-24 April 2025
©ECCOMAS, Portugal

DESIGN OPTIMIZATION OF ELASTOMER MOUNTS TO ASSESS WHIRL FLUTTER PHENOMENON IN A WIND TUNNEL

V. Wickramasinghe[1], A. Fereidooni, E. Chen[2], A. Grewal[2], and A. Ben Amar[2]

1: Flight Research Laboratory
National Research Council Canada
K1A 0R6
Viresh.Wickramasinghe @nrc-cnrc.gc.ca, https://nrc.canada.ca/en

2: Flight Research Laboratory
National Research Council Canada
K1A 0R6
{ Amin.Fereidooni, Eric.Chen, Anant.Grewal, Adam.BenAmar}@nrc-cnrc.gc.ca,
https://nrc.canada.ca/en

Abstract *With the growing use of electric motors in propeller aircraft, concerns around the propeller designs and their potential safety implications due to the whirl flutter phenomenon have become increasingly important. Whirl flutter is an instability that arises when the motion of an elastically supported rotating propeller interacts with the propeller aerodynamic forces. This coupling leads to a decrease in the overall dynamic stability of the propeller, engine and pylon system, potentially leading to catastrophic failure of aircraft components. Previous work in the literature has demonstrated that multiple parameters affect whirl flutter, namely, the propeller rotational speed, the structural damping and the stiffness of the power plant assembly in the pitch and yaw planes. Based on this knowledge, the current study investigated the optimization of the elastomeric motor mounts to assess the whirl flutter phenomenon in a wind tunnel environment. For the preliminary mount design, an analytical tool has been developed with the capability of providing accurate analytical predictions of the whirl flutter behavior. This tool uses two different approaches to generate results, namely, the Ribner approach which is focused on the structural damping and the Houbolt approach which is focused on the viscous damping. Using these two numerical models has resulted in the improved prediction of the whirl flutter and enhanced the design of the elastomeric mounts. The study concluded that the wind tunnel test and the results from the analysis are able to assess the whirl flutter phenomenon and optimize the elastomer-based engine mount designs*

Keywords: Whirl flutter, stiffness, damping ratio, aeroelastic, wind tunnel, HEAT aircraft

1. INTRODUCTION

This paper describes the design of a wind tunnel model of a Cessna 337 aircraft to enable experimental validation of whirl flutter prediction methodologies developed at the NRC. The goal is to replicate as much as feasible of the modified Cessna 337 aircraft, known the HEAT (Hybrid-Electric Aircraft-testbed) to be tested in NRC's low speed wind tunnel with test section size 9.1 m high x 9.1 m wide x 24 m long which has a maximum air speed of 55 m/s to identify the point of incipient whirl flutter. To achieve these goals, the stiffnesses of the elastomeric mounts that support the pusher propeller of the aircraft must be varied over a range of values to the point of flutter onset. The NRC in-house whirl flutter prediction model will be employed to assist in selecting the mount stiffnesses.

The whirl flutter phenomenon is known for being a practical engineering safety concern for propeller driven aircrafts. These concerns were triggered by two accidents involving the Lockheed Electra aircraft in 1959 and 1960. It was later discovered that whirl flutter could have been the cause of these accidents. Whirl flutter is an instability that arises when the rotating propeller motion interacts with the propeller aerodynamic forces. This coupling leads to a decrease in the overall stability of the propeller system and potentially leading to catastrophic failure of aircraft components. With the growing adoption of electric motors in propeller driven aircraft, the study of whirl flutter has become more important regarding the design of the engine mounts to avoid the instability. This study focuses on investigating the optimization of the elastomer engine mounts in the propeller system to assess the whirl flutter phenomenon in a wind tunnel. To do so, a numerical model was developed to accurately predict whirl flutter behaviour. Thereafter, the model was used to identify the required characteristics of the mounts to confidently predict flutter onset on the HEAT aircraft propeller operating in the wind tunnel.

Research around whirl flutter parameters have been conducted in the past. For example, Koch [1] conducted studies to determine key whirl flutter parameters, and showed that high wind speeds, increasing the size of the propeller or even reducing the distance between the propeller and the pivot point all have an impact on the whirl flutter onset. To follow up on this work, Koch and Koert [2] studied the impact of blade elasticity for frequency domain whirl flutter predictions and found that including blade elasticity improves the accuracy of the flutter prediction. Other studies were conducted to investigate the effect of other relevant parameters such as the work of Cecrdle [3], who determined the critical stiffness for the propeller engine attachment using an optimized analytical approach. This work also emphasized the benefits of these models in aircraft certification and documented the process to determine the whirl flutter stability boundaries for certification speed. This study concluded that such numerical models developed in the early phase of the design could continue to be used with sufficient confidence without requiring experimental results. Methodological differences also contributed to this field of research. Koch [4] presented a new method to asses whirl flutter using the Transfer Matrix method. This new method was able to assess the complex time domain propeller aerodynamics and blade elasticity. It compared the new transfer matrix method to already existing classical methods to validate their proposed methodology. Cocco [5] also developed a model for whirl flutter stability analysis in a wind tunnel environment. This study used methods such state linearization, Matrix Pencil Method and Periodic Operational Modal Analysis to validate their whirl flutter model.

Wind tunnel studies of propellers and their interactions with the aircraft nacelle and wings started in 1930s. However, the focus in that period was mainly the aerodynamic characteristics of the propellers. Bland and Bennett [6] conducted the first set of wind tunnel tests at the Langley Transonic Dynamics Tunnel (TDT) to evaluate the whirl flutter characteristics of a scaled down propeller model. The analysis of wind tunnel measurements revealed that he experimentally measured whirl flutter boundaries for the isolated propeller were in "very good agreement" with the analytical calculation in which the experimentally extracted aerodynamic derivatives were used. The use of theoretically derived aerodynamic derivatives resulted in a more conservative flutter boundary than the experimentally measured ones. The limited measurements obtained on the wing mounted propeller revealed that for these specific tests, the wing had little effects on the measurements but other studies have shown a more significant impact due to the dynamics of the wing structure.

In another test campaign conducted by Abbott et al., [7] which occurred in parallel with the analytical development of whirl flutter by Reed and Bland, a scaled-down version of a four-engine turboprop transport airplane (L-188 C Electra II aircraft) was tested. The test focused on seven important parameters that will be mentioned below alongside the important conclusions drawn in each case. Note that due to the importance of wording of these conclusions, they are quoted verbatim. Also note that the term "autoprecession" used by Abbott et al., [7] is synonymous with whirl flutter:

- Propeller rotational speed: Overspeeding of the propellers lowers the dynamic pressure for autoprecession. The detrimental effect of propeller overspeeding is more pronounced at the higher stiffnesses, but is little changed by damping.

- Stiffness of the power-plant assembly in the pitch and yaw planes: Reduction of stiffness in a power plant lowers the dynamic pressure at which the boundary for autoprecession occurs.

- Ratio of pitch stiffness to yaw stiffness of the power-plant assembly: At the higher stiffnesses investigated, a configuration with a stiffness ratio of 1.5 is more susceptible to autoprecession than a configuration with the same stiffness but having a stiffness ratio of 1.0.

- Structural damping of the power-plant assembly in the pitch and yaw planes: The dynamic pressure for autoprecession is very sensitive to changes in damping; small reductions in structural damping caused large reductions in the dynamic pressure for autoprecession.

- Simulated fuel load in the wings: Reduction of the fuel loading lowers the dynamic pressure for autoprecession.

- Location and/or number of autoprecessing power-plant assemblies: An inboard power plant is more susceptible to autoprecession than an outboard power plant. The autoprecession boundary is considerably lower for configurations with two or more power plants simultaneously having about the same reduced level of stiffness than it is for configurations in which the stiffness is reduced to this level in only one power plant.

• Wing effect: Propeller precession motion is accompanied at about the same frequency by wing response in bending and torsion, and an increase in amplitude of the precession motion is accompanied by an increase in amplitude of the wing motion.

2. HEAT AIRCRAFT WINDTUNNEL TEST CONFIGURATION

NRC has been using its Cessna 337 Skymaster as a Hybrid-Electric Aircraft-Testbed (HEAT) aircraft to perform research in low-emission in aviation. This twin-engine aircraft was selected because the centreline engine configuration made it easier to replace the rear engine with an electric motor while retaining the standard front engine to ensure safety of flight. The modified aircraft shown in Figure 1 has its rear piston engine replaced with a high-power-density electric motor powered by an 84-kWh battery module placed in the rear cabin connected using a customized high-voltage electrical system. The first fully electric flight of the HEAT aircraft was performed in February 2022. The goal of this facility is to develop the requisite technical, airworthiness and operational knowledge to support the Canadian aerospace industry to develop and validate electrified propulsion systems. Furthermore, the data gathered by NRC is continually shared with Transport Canada to inform the regulator regarding the suitable certification criteria needed for electric and hybrid aircraft.



Figure 1(a) Modified Cessna 337 Skymaster aircraft (b) rear electric propeller

The wind tunnel test article consists of a retired Cessna 377 Skymaster fuselage with the rear propeller replaced with an MTV-18/192-53 propeller operated by an EMRAC 348 electric motor powered by batteries to replicate the HEAT aircraft. Similar to the design of the HEAT aircraft, the rear electrically powered propeller will be mounted to the main fuselage through four elastomeric mounts. The propellers will be operated in both windmilling and powered conditions while the wind speed and mount stiffnesses will be varied as required to identify the whirl flutter onset. For the preliminary mount design, an analysis software has been developed with the capability of providing accurate analytical predictions of the whirl flutter behaviour. The data extracted from this wind tunnel test will be used to further validate the numerical model developed below.

2. WHIRL FLUTTER THEORY

The present whirl flutter prediction model is developed based on the approach developed by Reed and Bland [8]. The key aspects of the method include:

- modelling the power plant and propeller as a rigid body (with mass and inertia) gimballed about a pivot point with rotational springs in the pitching and yawing directions;

- modelling the aerodynamic forces using a method developed by Ribner [9], which considers the forces and moments acting on a propeller subject to non-axial (or yawed) flow.

For the present analysis, the engines mounts are assumed to be the dominant sources of flexibility. This is however a simplification because in reality the truss structure to which the motor is mounted may also provide considerable flexibility to the system. Furthermore, the flexibility of the fuselage itself will also contribute to the motion of the propeller hub and therefore contribute to the whirl flutter characteristic of an aircraft.

Mathematical analyses were performed in order to estimate the whirl flutter margin for the Cessna 337 aircraft with the new rear propeller and electric motor. As was mentioned earlier, whirl mode flutter can arise when the propeller and engine or motor assembly, which is usually mounted to the aircraft structure via elastomeric mounts to reduce the vibration transmitted from the propeller/power-plant system to the aircraft structure, undergoes yawing and pitching motion. Since the dynamic system is gyroscopic due to the spinning propeller and the electric motor, yawing and pitching motions becomes coupled leading to precession. Furthermore, the yawing and pitching motion produces unsteady aerodynamic forces and moments that are potentially destabilizing to the pitch and yaw dynamics. Specifically, when the propeller axis is deflected in pitch (yaw), a lift (side) force, along with a yawing (pitching) moment arises. These aerodynamic forces and moments are, for small deflections, linearly dependent on the pitch (yaw) angle. Note that aerodynamic forces and moments which are proportional to angular velocities are also present [10]. As with other aeroelastic systems, the dependence of aerodynamic forces and moments on linear or angular displacements, and the rates of change of these displacements, alter the modal stiffness and damping of the system. For combinations of operating conditions (airspeed, propeller rotational speed, air density, etc.), the system can become unstable. This is referred to as flutter in the aeroelastic community.

For the present case, certain characteristics of the aircraft and configuration simplify the modelling of the whirl flutter problem:

1. The pusher or the rear propeller and electric motor are fuselage mounted, and therefore the wing structural dynamics need not be included.
2. A preliminary ground vibration test (GVT) conducted to measure the stiffness and damping characteristics of the mount revealed that the lowest frequency structural modes of the propeller, power-plant and aft fuselage were dominated by the *localized* motion of the mounts in pitch and yaw. Hence the model need not incorporate the fuselage structural dynamics modes. A more detailed GVT was also conducted later, and this test revealed that there may be significant contributions from the rear truss structure to which the motor is mounted. This aspect is discussed later in this paper.

The definition of the coordinate systems, relevant forces, moments and orientation (Euler angles) are depicted in the Figure 2 below.

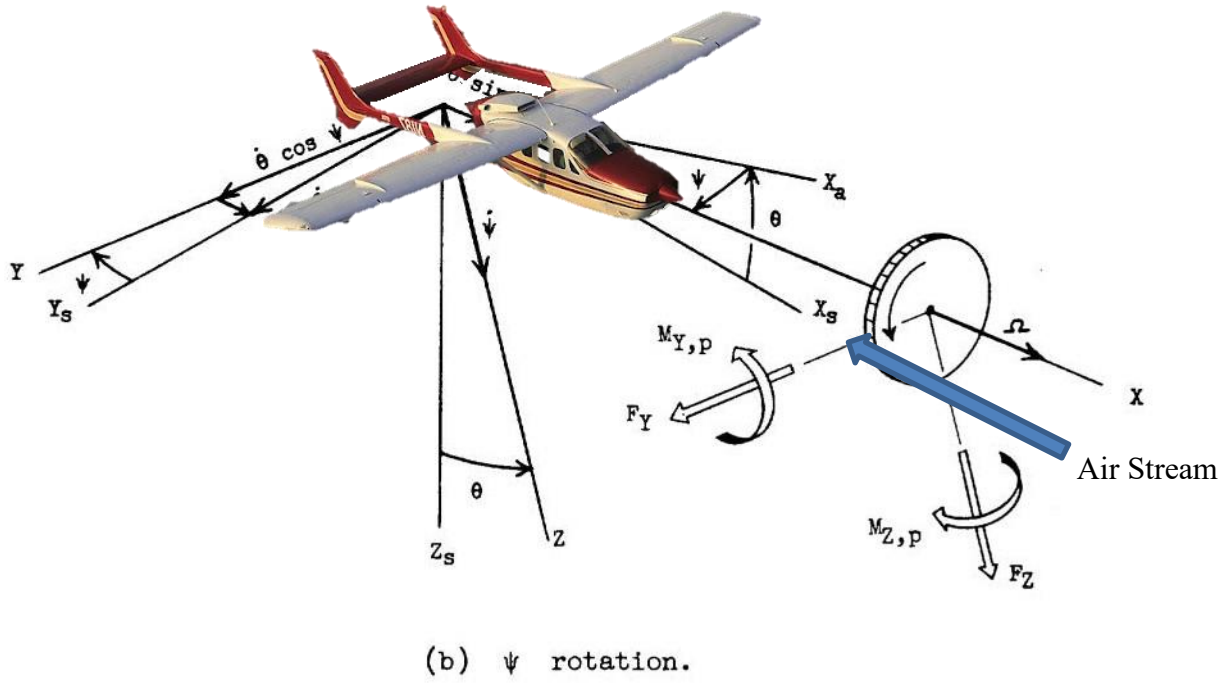


Figure 2 Coordinate Systems used in the analysis

The dynamical model for whirl flutter analysis is developed by applying Lagrange's Equations on the kinetic, T, and potential, U, energies of the system, as well as considering the dissipation (damping), D and the generalized forces, Q_i, acting on the system:

$$\frac{d}{dt}\left(\frac{\partial T}{\partial \dot{q}_i}\right) - \left(\frac{\partial T}{\partial q_i}\right) + \left(\frac{\partial U}{\partial q_i}\right) + \left(\frac{\partial D}{\partial \dot{q}_i}\right) = Q_i \quad [1]$$

where q_i and $\dot{q}_i$ are the generalized coordinates and velocities, respectively. The kinetic energy is expressed with respect to the pitching, yawing and rolling (spinning) motion of the sprung system, consisting of the motor and propeller; the potential energy and dissipation functions are expressed in terms of the pitch and yaw displacements, and velocities, respectively; and the generalized forces are expressed in terms of the unsteady aerodynamics.

The resulting two degree of freedom 2nd order ordinary differential equations are given by:

$$\begin{aligned} I_y\ddot{\theta} + I_x\Omega\dot{\psi} + C_\theta\dot{\theta} + S_\theta\theta &= M_Y + \ell F_Z \\ I_z\ddot{\psi} - I_x\Omega\dot{\theta} + C_\psi\dot{\psi} + S_\psi\psi &= M_Z + \ell F_Y \end{aligned} \quad [2]$$

The unsteady motion-induced aerodynamic forces and moments, i.e., F_y, F_z, M_y and M_z can be expressed in term of the pitch and yaw motion as follows:

$$F_Z = \frac{1}{2}\rho V^2 S'\left(C_{Z_\theta}\bar{\theta} + C_{Z_\psi}\bar{\psi} + C_{Z_r}\frac{\vec{\psi}R}{V}\right) \quad [3]$$

$$M_Y = \rho V^2 S' R\left(C_{m_\psi}\bar{\psi} + C_{m_q}\frac{\vec{\theta}R}{V}\right)$$

$$F_Y = \frac{1}{2}\rho V^2 S' \left(C_{Y_\psi}\bar{\psi} + C_{Y_\theta}\bar{\theta} + C_{Y_q}\frac{\dot{\bar{\theta}}R}{V} \right)$$

$$M_Z = \rho V^2 S' R \left(C_{n_\theta}\bar{\theta} + C_{n_r}\frac{\dot{\bar{\psi}}R}{V} \right)$$

where the effective pitch and yaw, $\bar{\theta}$ and $\bar{\psi}$, are related to: the geometric pitch and yaw, θ and ψ; the pitch and yaw rates, $\dot{\theta}$ and $\dot{\psi}$; the airspeed V; and the distance ℓ between the rotation axis (elastomeric mount location) and the propeller plane by the following relations:

$$\bar{\theta} = \theta - \frac{\ell\dot{\theta}}{V} \qquad [4]$$

$$\bar{\psi} = \psi - \frac{\ell\dot{\psi}}{V}$$

The aerodynamic coefficients can be estimated using two classical methods namely the Ribner [9], and Houbolt and Reed [10] approaches that have been known to provide results to good engineering accuracy and have been validated with experimental methods. The analysis described herein focuses mainly on the Ribner approach. With the dynamic model expressed as a system of two 2nd order Ordinary Differential Equations, the stability of the system with respect to the operating conditions was determined based on standard eigenvalue analysis. The following expressions are simplified approximation for frequency (λ) and structural damping ($\tilde{g}_\theta$).

$$\lambda_{1,2} = 1 \pm \frac{1}{2}\frac{I_x\Omega}{I_y\omega_\theta},\ \tilde{g}_{\theta_{2,1}} = \frac{\kappa}{k_\theta}\left[\lambda_{2,1}a_1 \pm \frac{b_0}{k_\theta}\right], a_1 = c_{m_q} + \frac{1}{2}\left(\frac{l}{R}\right)^2 C_{z_\theta},\ b_0 = c_{m_\psi} - \frac{1}{2}\left(\frac{l}{R}\right)^2 C_{z_\psi}$$

Of course, the various parameters of the model need to be determined before this step can proceed.

3. WHIRL FLUTTER MODEL VALIDATION

Prior to analyzing the NRC Cessna 337 HEAT II aircraft using this whirl flutter prediction model, the model and methodology employed was validated by replicating analyses conducted by NASA [8], which were in turn validated with experimental testing. A comparison of the current analysis of the NASA case with the NASA analysis (shown as Figure 7 in [8]) is shown below in Figure 8. The upper frame of Figure 8 depicts the variation of the forward and backward whirl frequencies as a function of the reduced velocity, i.e., $V/R\omega_\theta$, where ω_θ represents the fundamental frequency in pitch for the elastomeric mount (note that in cases considered in this report for the NASA test case or the HEAT II aircraft analysis, the pitch and yaw inertia, elastic and damping parameters are identical). The lower frame presents the required structural damping for whirl-flutter free operation. For this NASA test case, it is evident that for a wide range of advance ratios (i.e. 1.8 to 4.2) and reduced velocities (i.e. 0 to 5) the NRC analysis match the NASA results very well.

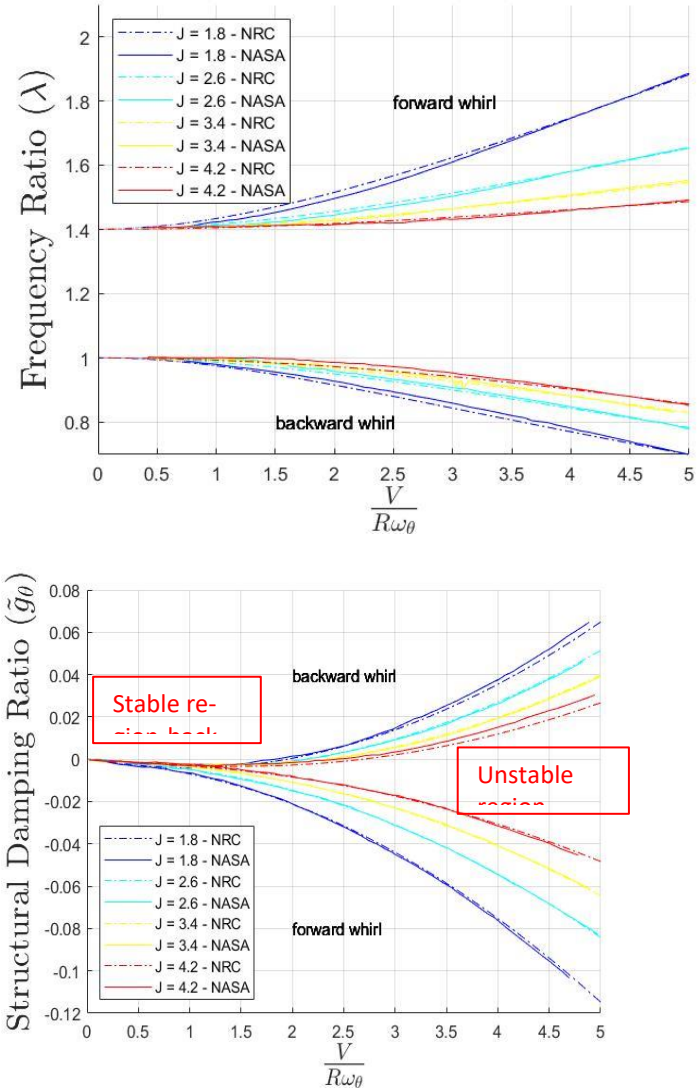


Figure 3: Validation of Current Whirl Flutter Analysis using NASA Test Case

4. ELASTOMERIC MOUNT DESIGN FOR WIND TUNNEL TESTING

Using the analytical method developed, the total bending stiffness of in-house developed elastomers was calculated that resulted in a whirl flutter speed less that the maximum U70 wind tunnel speed of 55 m/s. In light of this objective, we select multiple target whirl flutter speed, and calculated total bending stiffnesses. For the present analysis, a structural damping value of 1.5% was assumed, which results in a prediction of flutter if the damping ratio computed by the model falls below this value. Knowing this, the study was able to test various total bending stiffness and determine at what speed the required damping ratio to maintain stability increases above 1.5%. The bending stiffnesses associated with various whirl flutter speeds are shown in Table 1. The required damping ratios versus wind speed for three flutter speeds of 10 m/s, 30 m/s and 50 m/s are shown in Figure 4, Figure 5 and Figure 6, respectively.

Table 1 Estimation of Bending Stiffness for Selected Whirl Flutter Speeds

Whirl Flutter Speed (m/s)	Total Bending Stiffness (kN.m/rad)	Assigned Name
10	0.2998	NRC-010
15	0.7035	NRC-015
20	1.2610	NRC-020
25	1.9542	NRC-025
30	2.8901	NRC-030
35	3.9219	NRC-035
40	5.3203	NRC-040
45	6.5682	NRC-045
50	8.4817	NRC-050
55	10.4728	NRC-055

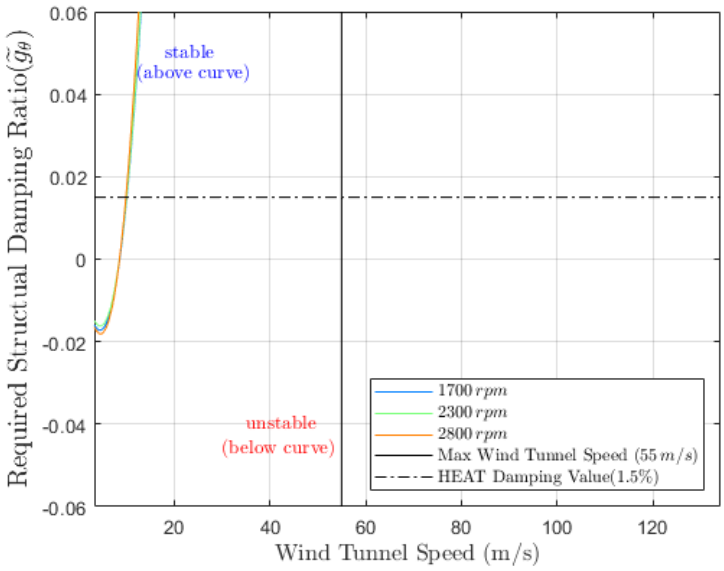


Figure 4 Wind tunnel speed versus required structural damping: Stiffness of 0.2998 kN-m/rad

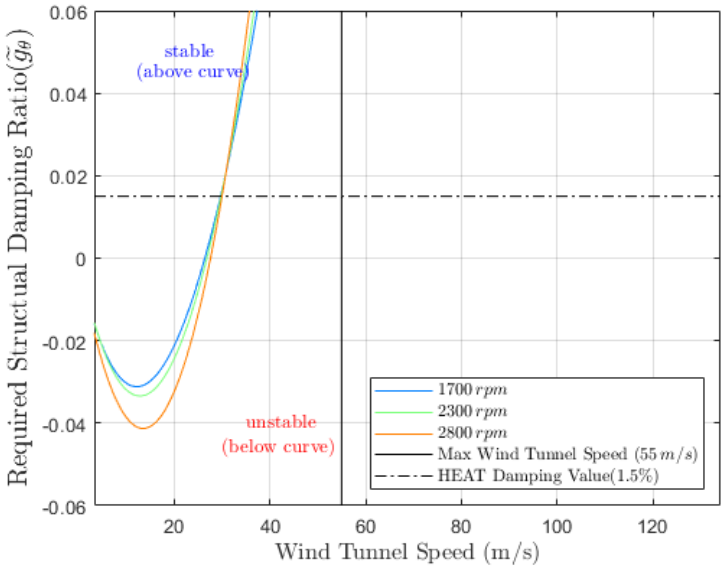


Figure 5 Wind tunnel speed versus required structural damping: Stiffness of 2.8901 kN-m/rad

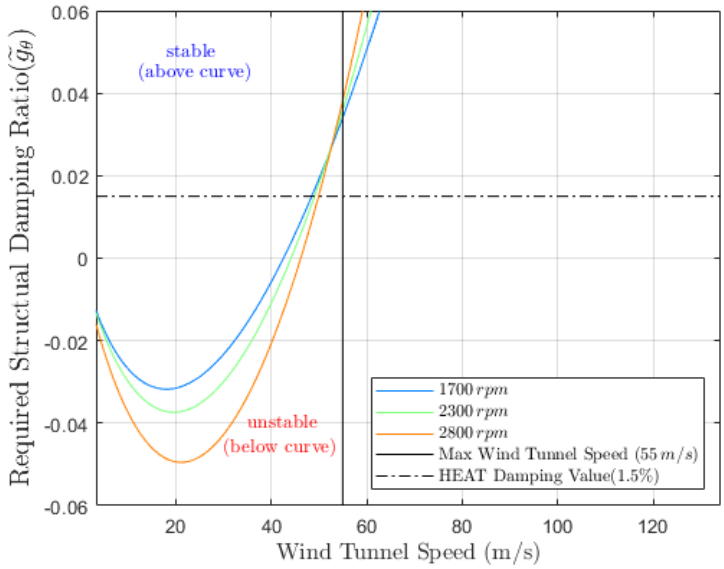


Figure 6 Wind tunnel speed versus required structural damping: Stiffness of 8.4817 kN-m/rad

Now with the total bending stiffness know for each in-house made elastomeric mounts and using a linear assumption we can extrapolate the required axial stiffness needed per mount. This means that we assumed that the stiffness of the potential in-house made elastomeric mounts lie on a line connecting the commercial Lord Corporation two-piece mount (CB2201-1) value to the origin. Its important to note that these estimations are based on assumption and because of this, the values presented are estimate of what one would anticipate in the actual wind tunnel testing. These values need to be used carefully and should be used as a rough guideline when characterizing the in-house made elastomeric mounts.

Table 2 Estimation of Axial Stiffness for Selected Whirl Flutter Speeds

In-house Elastomeric Mount	Required Axial Stiffness per Mount (kN/m)
NRC-010	2.506
NRC-015	5.880
NRC-020	10.539
NRC-025	16.333
NRC-030	24.155
NRC-035	32.778
NRC-040	44.466
NRC-045	54.895
NRC-050	70.888
NRC-055	87.529

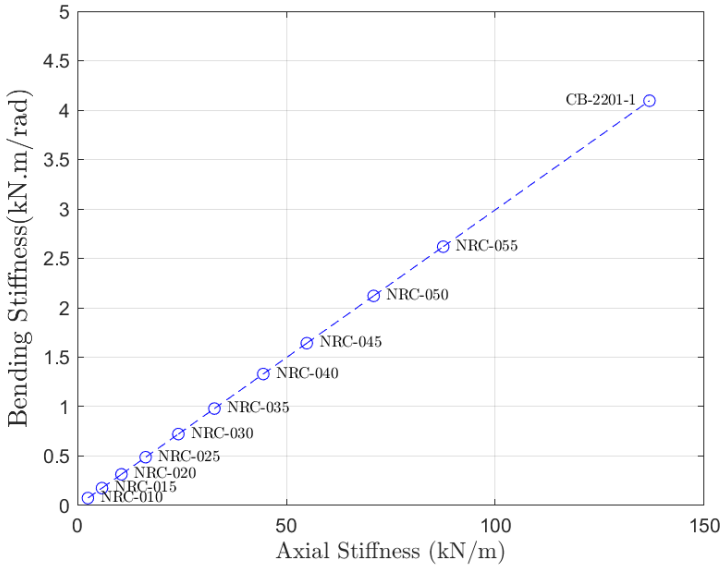


Figure 7 Estimation of Axial Stiffness for a Potential In-house Made Elastomeric Mounts

5. FUTURE WHIRL FLUTTER MODELLING

5.1. Background and Rationale

To overcome the limitation of the modelling presented in Section 4, a larger portion of the fuselage would need to be incorporated into the structural dynamics model. This could be done through development of a Finite Element Model of a portion of, or the complete, aircraft. Unfortunately, such a finite element model is not available. Another approach is to develop a modal model of the aircraft through the use of a ground vibration test to extract modal parameters of the aircraft. Such a test was conducted on the Cessna 337 HEAT aircraft in November to December 2024. The following section provides further details of the test campaign.

5.2. Ground Vibration Test

Since aircraft structural components are relatively flexible, they may be prone to a variety of aeroelastic instabilities. Ground Vibration Tests (GVT) are used to extract the modal characteristics of such flexible structures. This information is used to plan operations of the aircraft and interpret vibration data from flight tests. As such, a GVT was performed in order to extract the modal characteristics for the HEAT aircraft structures. A variety of equipment including structural exciters, multiple sensors and data recording systems were employed to conduct the GVT.

In the GVT, a PCB 086D02 impact hammer and a PCB Model 2100E11 100-lbf modal shaker were used to excite the aircraft structural modes of interest in separate trials. The purpose was to verify that the modes of interest have been effectively excited and extracted, and also verify that the modal parameters were consistent in both excitation methods.

In the hammer impact testing, the impact load was applied at the upper-right corner of the metallic mount plate on the motor side in the X, Y and Z directions, respectively (see Figure 8). Note that the four Corporation two-piece (CB2201-1) rubber mounts were connected in parallel. When the PCB modal shaker was used for excitation, the shaker load was applied at the same location., with the shaker stinger oriented in 0/45/45 degrees relative to the X, Y and Z axis of the HEAT aircraft to excite all directions. The PCB modal shaker was suspended from a portable crane using bungee cords, and controlled to provide low levels swept sine and broadband random excitation energy between 3 and 120 Hz through a stringer in series with an PCB Model 288D01 impedance head.



Figure 8 Shaker and stinger attachment to HEAT motor mounting plate

The HEAT aircraft was positioned on its landing gear inside the NRC aircraft hangar, and there was no occupant or external load on the aircraft. In order to assess boundary condition effects,

the GVT was performed under two tire pressure configurations: in the first case, all three tires were inflated to 100% nominal pressure, while in the second case the tires were deflated to 50% of the nominal pressure. This arrangement was used to determine if aircraft tire pressure would impact the derived modal parameters of the HEAT aircraft structures. Ideally, the aircraft would be suspended elastically from it wings to render a near free-free boundary condition consistent with flight. However, this support condition was not possible at the time of test.

A total of 29 PCB ICP type tri-axial accelerometers were attached to the HEAT aircraft structures at selected locations to measure the structural responses (see Figure 9). The hammer/shaker excitation load and the accelerometer responses were recorded synchronously using a Siemens SCADAS III front end frame and a laptop running the Siemens Test.Lab Spectral Testing Module (R18). The measured input load from the impedance head and vibration responses from accelerometers were used to derive the required data for modal analysis, which included the Frequency Response Functions (FRFs), coherence curves, etc. For data acquisition purposes, a frequency bandwidth of 1024 Hz was used, and the resulting frequency resolution was 0.125 Hz for spectral analysis. In addition, the time traces of the input load and accelerometer responses were also recorded to enable additional post analysis, if necessary.

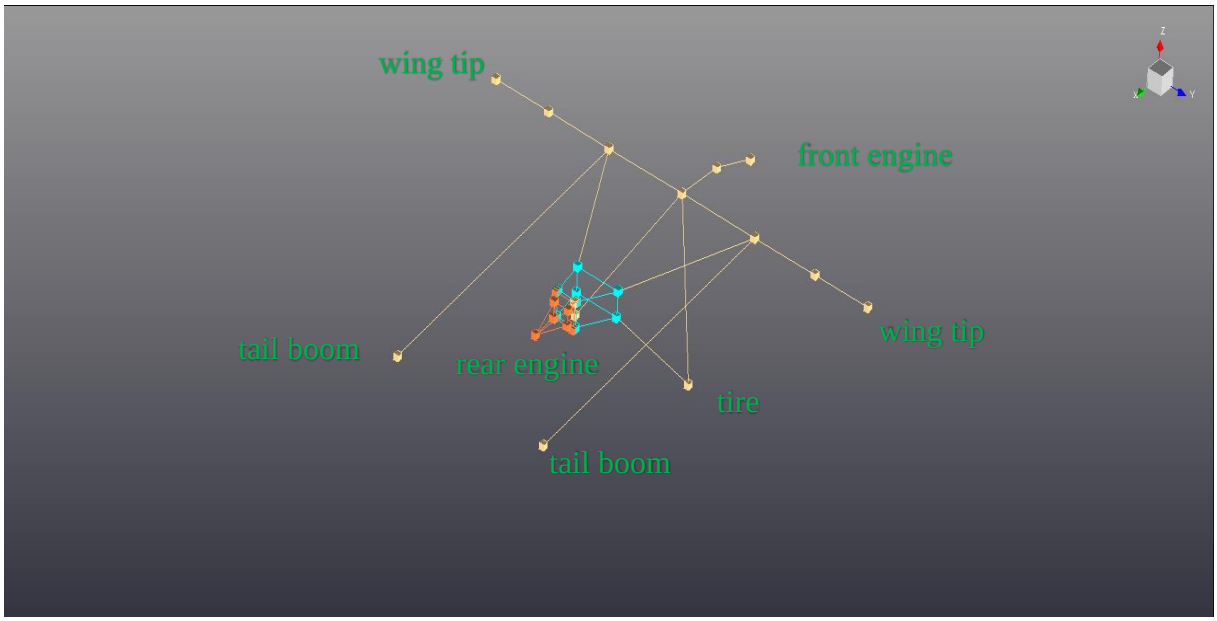


Figure 9 Accelerometer locations on the HEAT aircraft geometry model

Using the GVT data, major structural modes of the HEAT aircraft propeller and motor assembly were identified as per standard modal analysis procedures. Modal parameters including modal frequencies, damping ratios, modal shapes, modal masses and stiffnesses were extracted for each mode using the Test.Lab Polymax Plus Modal Analysis Module (R18). The identified aircraft wing mode and first bending mode of the rear engine at full tire pressure configuration are shown in Figure 10 and Figure 11 respectively. Modal parameters extracted for major aircraft modes between 5 and 60 Hz are listed in Table 3.

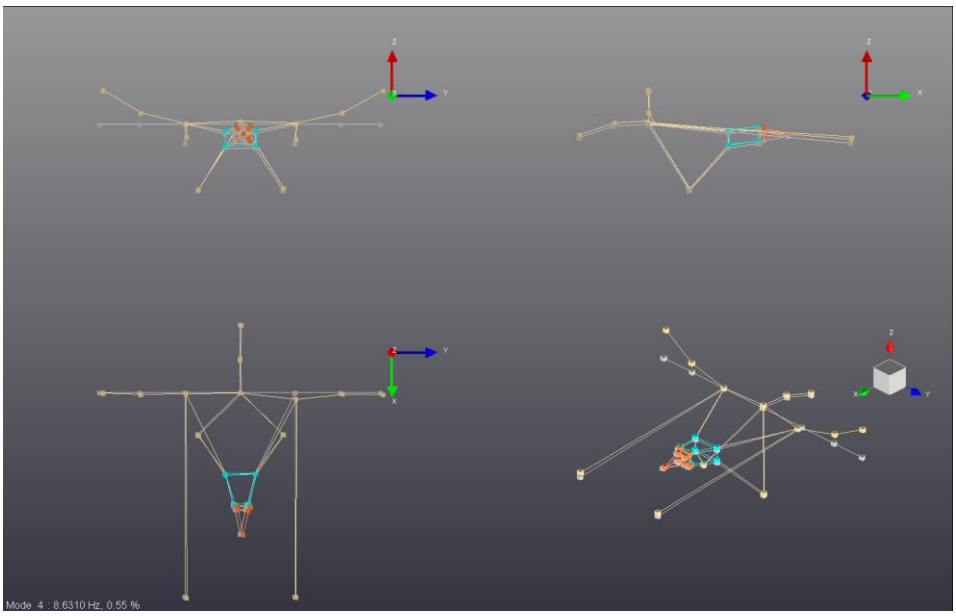


Figure 10 HEAT aircraft first symmetric wing bending mode (8.6Hz)

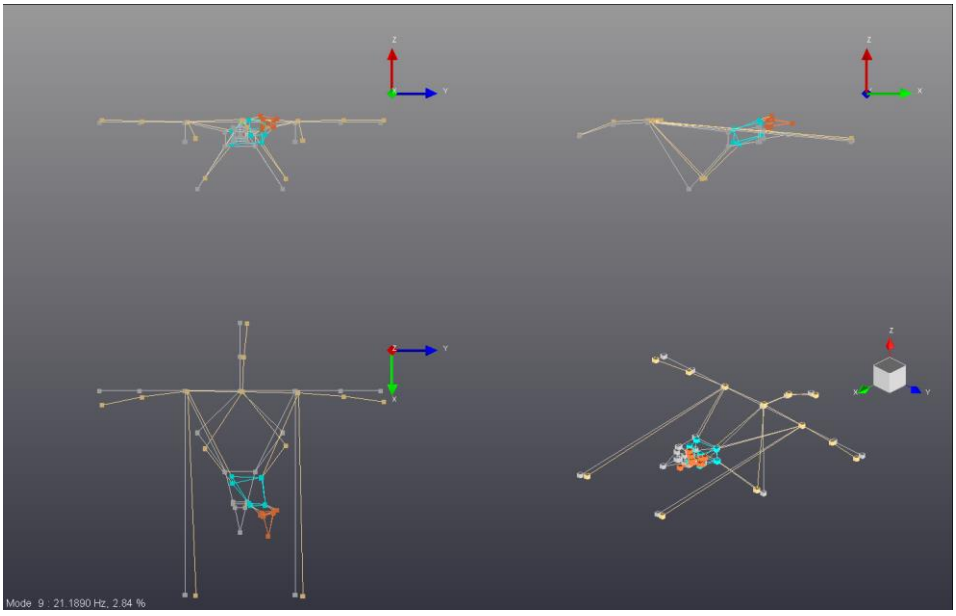


Figure 11 Rear motor mount mode (21.2Hz) related to rubber mount bending

Table 3 Derived modal parameters for major aircraft modes

Mode	Frequency (Hz)	Damping Ratio (%)
1	6.2	1.16
2	6.5	1.02
3	6.9	1.60

4	8.6	0.55
5	11.4	3.79
6	12.2	2.86
7	14.1	1.20
8	16.3	3.36
9	17.6	2.17
10	18.6	2.20
11	21.2	2.93
12	21.8	2.67
13	24.7	3.37
14	27.7	1.29
15	28.9	3.30
16	29.5	1.68
17	33.3	1.11
18	34.8	2.28
19	36.5	1.72
20	40.4	2.98
21	43.7	1.75
22	48.4	2.59
23	49.8	0.65
24	50.8	1.90
25	52.8	1.12

5.3 Updated Modelling Methodology

Expressed in terms of conventional aerodynamic coefficients, as well as translational and rotational motion at the hub, the unsteady motion-induced aerodynamic forces and moments, i.e., F_Y, F_Z, M_Y and M_Z acting on the propeller hub are expressed as follows [11]:

$$F_z = \frac{1}{2}\rho V^2 S' \left[C_{z_\theta}\left(\theta + \frac{\dot{z}}{V} - \frac{w}{V}\right) + C_{z_\psi}\left(\psi - \frac{\dot{y}}{V} + \frac{v}{V}\right) + C_{z_q}\left(\frac{\dot{\theta}R}{V}\right) + C_{z_r}\left(\frac{\dot{\psi}R}{V}\right)\right] \quad [5]$$

$$M_y = \rho V^2 S' R \left[C_{m_\theta}\left(\theta + \frac{\dot{z}}{V} - \frac{w}{V}\right) + C_{m_\psi}\left(\psi - \frac{\dot{y}}{V} + \frac{v}{V}\right) + C_{m_q}(\dot{\theta}R/V) + C_{m_r}(\dot{\psi}R/V)\right]$$

$$F_y = \frac{1}{2}\rho V^2 S' \left[C_{y_\theta}\left(\theta + \frac{\dot{z}}{V} - \frac{w}{V}\right) + C_{y_\psi}\left(\psi - \frac{\dot{y}}{V} + \frac{v}{V}\right) + C_{y_q}(\dot{\theta}R/V) + C_{y_r}(\dot{\psi}R/V)\right]$$

$$M_z = \rho V^2 S' R \left[C_{n_\theta}\left(\theta + \frac{\dot{z}}{V} - \frac{w}{V}\right) + C_{n_\psi}\left(\psi - \frac{\dot{y}}{V} + \frac{v}{V}\right) + C_{n_q}(\dot{\theta}R/V) + C_{n_r}(\dot{\psi}R/V)\right]$$

Where: ρ is the air density; V is the freestream velocity; S' is the propeller disk area; v and w are the side-wash and down-wash, respectively, induced by lifting surfaces; y and z are the hub movements in the lateral and vertical directions, respectively; and θ and ψ, are the pitch

and yaw angles of the hub, respectively. The aerodynamic coefficients for the propeller are defined using double-subscripts where the first subscript denotes the force or moment direction, i.e. y and z are the lateral and vertical forces, respectively, and m and n the pitching and yawing moments, respectively. The second subscript refers to parameters with respect to which derivatives are evaluated. In addition to the previously defined orientation angles, q and r are the pitch and yaw rates, respectively.

Following the approach taken by Koch [4], the hub forces can be written as:

$$\frac{1}{\pi\rho V^2 R^3}\begin{Bmatrix} F_y \\ F_z \\ M_y \\ M_z \end{Bmatrix} = \begin{bmatrix} 0 & 0 & \frac{C_{y_\theta}}{2R} & \frac{C_{y_\psi}}{2R} \\ 0 & 0 & \frac{C_{z_\theta}}{2R} & \frac{C_{z_\psi}}{2R} \\ 0 & 0 & C_{m_\theta} & C_{m_\psi} \\ 0 & 0 & C_{n_\theta} & C_{n_\psi} \end{bmatrix}\begin{Bmatrix} y \\ z \\ \theta \\ \psi \end{Bmatrix} + \begin{bmatrix} -\frac{C_{y_\psi}}{2RV} & \frac{C_{y_\theta}}{2RV} & \frac{C_{y_q}}{2V} & \frac{C_{y_r}}{2V} \\ -\frac{C_{z_\psi}}{2RV} & \frac{C_{z_\theta}}{2RV} & \frac{C_{z_q}}{2V} & \frac{C_{z_r}}{2V} \\ -\frac{C_{m_\psi}}{V} & \frac{C_{m_\theta}}{V} & \frac{C_{m_q}R}{V} & \frac{C_{m_r}R}{V} \\ -\frac{C_{n_\psi}}{V} & \frac{C_{n_\theta}}{V} & \frac{C_{n_q}R}{V} & \frac{C_{n_r}R}{V} \end{bmatrix}\begin{Bmatrix} \dot{y} \\ \dot{z} \\ \dot{\theta} \\ \dot{\psi} \end{Bmatrix} \quad [6]$$

In the frequency domain these equations may be expressed as [4]:

$$\vec{F}_{prop}(\vec{X}, s, V, \Omega) = \begin{Bmatrix} F_y \\ F_z \\ M_y \\ M_z \end{Bmatrix} = \left[\overline{\overline{K}}_{prop}(V, \Omega) + s\overline{\overline{D}}_{prop}(V, \Omega)\right]\vec{X} \quad [7]$$

Where:

$$\overline{\overline{K}}_{prop}(V, \Omega) = \pi\rho V^2 R^3 \begin{bmatrix} 0 & 0 & \frac{C_{y_\theta}}{2R} & \frac{C_{y_\psi}}{2R} \\ 0 & 0 & \frac{C_{z_\theta}}{2R} & \frac{C_{z_\psi}}{2R} \\ 0 & 0 & C_{m_\theta} & C_{m_\psi} \\ 0 & 0 & C_{n_\theta} & C_{n_\psi} \end{bmatrix} \quad [8]$$

$$\overline{\overline{D}}_{prop}(V, \Omega) = \pi\rho V^2 R^3 \begin{bmatrix} -\frac{C_{y_\psi}}{2RV} & \frac{C_{y_\theta}}{2RV} & \frac{C_{y_q}}{2V} & \frac{C_{y_r}}{2V} \\ -\frac{C_{z_\psi}}{2RV} & \frac{C_{z_\theta}}{2RV} & \frac{C_{z_q}}{2V} & \frac{C_{z_r}}{2V} \\ -\frac{C_{m_\psi}}{V} & \frac{C_{m_\theta}}{V} & \frac{C_{m_q}R}{V} & \frac{C_{m_r}R}{V} \\ -\frac{C_{n_\psi}}{V} & \frac{C_{n_\theta}}{V} & \frac{C_{n_q}R}{V} & \frac{C_{n_r}R}{V} \end{bmatrix}$$

$$\vec{X} = \begin{Bmatrix} y \\ z \\ \theta \\ \psi \end{Bmatrix}$$

Note that s is the Laplace variable.

One approach to determining the elements of the aerodynamic coefficients found in the above matrices is through the use of simplified propeller theories, such as by the one proposed by Ribner [9]. Other more accurate methods, include using a transfer matrix approach, outlined in Koch [2], [4] can also be employed. Koch uses a multibody simulation method to determine time domain responses of propeller forces and moments to perturbations in translational and rotational displacement applied at the propeller hub. Note that this analysis is performed on the isolated propeller. The time histories are subsequently transformed to the frequency domain. In the current study, the Ribner or Houbolt approach will be pursued in the next phase of this project.

Next, the propeller force model is coupled with the aircraft structural dynamics equations of model, which are obtained from the previously described modal analysis of the ground vibration tests performed on the Cessna 337 HEAT aircraft. Also included is the generalized aerodynamic force matrix $\bar{\bar{Q}}_{hh}$, which captures the unsteady effects from the (non-propeller) portions of the aircraft.

$$\bar{\bar{M}}_{mod}\ddot{\vec{q}} + \bar{\bar{K}}_{mod}\vec{q} = \left\{\bar{\bar{\Phi}}_{prop}^{T}\left[\bar{\bar{K}}_{prop}(V,\Omega) + s\bar{\bar{D}}_{prop}(V,\Omega)\right]\bar{\bar{\Phi}}_{prop} + \frac{\rho V^2}{2}\bar{\bar{Q}}_{hh}(s)\right\}\vec{q} \quad [9]$$

$\bar{\bar{M}}_{mod}$, $\bar{\bar{K}}_{mod}$ and $\bar{\bar{\Phi}}_{prop}^{T}$ are the modal mass, modal stiffness and mode shape matrices estimated from the modal analysis procedure. $\vec{q}$ is the vector of generalized coordinates. This system of equations is solved to determine the whirl flutter boundary through an eigenvalue analysis. This process is complicated by the frequency dependent character of both the propeller transfer function matrix $s\bar{\bar{D}}_{prop}$, as well as the generalized aerodynamic force matrix $\bar{\bar{Q}}_{hh}(s)$. Moreover, due to the implicit character of the problem, a number of iterative solution techniques that are well established are available to solve for the eigenvalues of the system which reveal the natural frequency and damping ratios of the modes. One such example is the p-k method, which yields the damping and frequency variation of the system modes with airspeed.

5. CONCLUSION

The objective of this study was to optimize the elastomeric mounts of the propeller and electric motor configuration in preparation of the study of the whirl flutter phenomenon in a wind tunnel. The primary purpose of the wind tunnel testing is to validate a whirl flutter analytical model that is required to ensure whirl flutter free operation of the NRC HEAT hybrid electric aircraft testbed, and possibly other flight vehicles. The wind tunnel testing will be performed on a Cessna 337 fuselage modified to represent the NRC HEAT aircraft. Due to the air speed limit of NRC wind tunnel (i.e. 55 m/s versus a flight speed of approximately 85 m/s), the motor mounts employed in the wind tunnel model will necessarily be significantly more compliant compared to the mounts used in flight.

The current analytical model adopted the Ribner aerodynamics modelling approach, and ultimately estimates the level of structural damping required to operate at the whirl flutter boundary. This modelling approach was validated using a previously published NASA data set, which was itself validated using experimental data. The model was then used to estimate the onset of whirl flutter instability under various parameters. This study concluded that the wind tunnel test and the results from the analysis are able to assess the whirl flutter phenomenon and optimize the elastomer-based motor mount designs. This research can provide a guideline to help in the selection of elastomeric mount stiffnesses for further work. Also presented was an extension to this work which will use the experimentally measured mode shapes and other modal parameters in a dynamic formulation, which will be a further improvement of the simplified torsional spring-rotary inertia structural dynamics model employed in the current study.

REFERENCES

[1] Koch, C., “Parametric Whirl Flutter Study Using Different Modelling Approaches,” CEAS Aeronautical Journal, Vol. 13, No. 1, 2021, pp. 57–67. https://doi.org/10.1007/s13272-021-00548-0

[2] Koch, C., and Koert, B., “Including Blade Elasticity into Frequency-Domain Propeller Whirl Flutter Analysis" https://doi.org/10.2514/1.C037501

[3] Cecrdle, J.: Analysis of Twin Turboprop Aircraft Whirl-Flutter Stability Boundaries. J. Aircr. 49(6), 1718–1725 (2012).

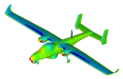
[4] Koch, C., “Whirl Flutter Stability Assessment Using Rotor Transfer Matrices,” Proceedings of International Forum on Aeroelasticity and Structural Dynamics (IFASD), Univ. Carlos III Madrid, Paper IFASD-2022-040, 2022.

[5] Cocco, A., Mazzetti, S., Masarati, P., Hoff, S.V and Timmerman, B., “Numerical Whirl–Flutter analysis of a tiltrotor semi-span wind tunnel Model” CEAS Aeronautical Journal (2022) 13:923–938 https://doi.org/10.1007/s13272-022-00605-2

[6] Bland, S.R., and Bennett, R.M., “Wind-tunnel Measurement of Propeller Whirl-flutter Speeds and Static-stability Derivatives and Comparison with Theory,” Report NASA-TN-D-1807, NASA Langley Research Center, August 1963.

[7] Abbott, F.T, Sr., Kelly, H.N., and Hampton, K.D., “Investigation of Propeller-Power Plant Autoprecession Boundaries for a Dynamic-Aeroelastic Model of a Four-Engine Turbprop Transport Airplane,” Report NASA-TN-D-1806, NASA Langley Research Center, August 1963.

[8] Reed III, W. H. and Bland, S. R., "An analytical treatment of aircraft propeller precession instability," Report NASA-TN-D-659, NASA Langley Research Center, January 1961.
[9] Ribner, H. S., "Propellers in Yaw," National Advisory Committee for Aeronautics, Report NACA-TR-820, 1943.
[10] Houbolt, J.C., and Reed, W.H. III, "Propeller-Nacelle Whirl Flutter," NASA Langley Research Center, Presented at the Aeroelasticty Session, IAS 29th Annual Meeting, New York, Jan 23-24, 1961.
[11] Rodden, W.P., and Rose, T.L., "Propeller/Nacelle Whirl Flutter Addition to MSC/NASTRAN," Proceedings of the 1989 MSC World User's Conference, Universal City, CA, USA, 1989.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AEROELASTIC MODELING AND DESIGN OF A FLEXIBLE BLENDED-WING-BODY (BWB) RESEARCH VEHICLE

Mario Bras[1*], Sid Banerjee[2] and Afzal Suleman[1,3]

1: Center for Aerospace Research
Department of Mechanical Engineering
University of Victoria
3800 Finnerty Rd, Victoria BC, V8P 5C2, Canada
{mbras,suleman}@uvic.ca, https://www.uvic-cfar.com

2: Advanced Design, Research & Technology
Bombardier
400 Boulevard Côte-Vertu West, Dorval QC, H4S 1Y9, Canada
Sid.Banerjee@aero.bombardier.com, https://bombardier.com

3: IDMEC
Instituto Superior Técnico
Universidade de Lisboa
Av. Rovisco Pais, No. 1, 1049-001, Lisboa, Portugal
suleman@tecnico.ulisboa.pt, https://www.idmec.tecnico.ulisboa.pt

Abstract. *The blended-wing-body (BWB) aerodynamic configuration has been shown to reduce drag, improve lift, and decrease weight and fuel consumption, making it a promising solution for future sustainable aviation. Understanding and controlling the aeroelastic behavior of BWB configurations with high-aspect ratio wings is a key challenge where scaled flexible-wing flight test demonstrators can provide a low-cost, low-risk alternative to full-scale aircraft flights.*

These research vehicles can be designed to exhibit aeroelastic instabilities at low speeds to help characterize the aeroelastic behavior of the full-scale aircraft. However, studies on the flutter mechanisms of BWB aircraft are scarce and these configurations can exhibit flutter modes typical of flying wing and conventional wing configurations. Characterizing their aeroelastic behavior is of extreme importance if such aircraft designs are to be flown in the future.

Towards that end, this paper presents a structural design study performed on a reduced-scale BWB flight demonstrator where the aeroelastic response of the aircraft is tailored to exhibit flutter at low speeds while guaranteeing structural integrity throughout the flight envelope. Aeroelastic models of the aircraft are developed in NASTRAN and ASWING. Different internal structural configurations are defined followed by a sweep of parametric studies. Aeroelastic modes, as well as static wing displacements for different load cases, are identified for each data point in the study. Finally, a comparative study between the two models is performed.

Keywords: aeroelasticity, flexible wing, blended-wing-body (BWB), flutter mechanisms

1 INTRODUCTION

The Center for Aerospace Research (CfAR) at the University of Victoria and Bombardier have been collaborating in the area of greening aviation and advanced aircraft design for the past ten years. The partnership has led to the development and successful testing of a large scale blended-wing-body flight demonstrator in the framework of the Bombardier EcoJet Research Project (Fig.1). The objective of the program is to design and manufacture scaled flight test platforms to generate performance data through flight tests. This data has been used to validate the flightworthiness of the novel configuration. The program drew on extensive testing of the aircraft concept both in wind tunnel as well as in free-air using scaled flight test vehicles. These collaborative efforts are expected to yield to a reduction in aircraft emissions and lead the way to sustainable aviation in Canada



Figure 1: Bombardier BWB aircraft.

Findings from this project highlighted the need to fully evaluate new control laws developed and assess the aeroservoelastic performance of the BWB configuration. With this in mind, this work proposes to evaluate the feasibility of developing a flexible aeroelastic demonstrator through modification of the existing rigid composite wing of the scaled vehicle.

The Bombardier BWB flight demonstrator, a 5.4 m span vehicle (see Tab.1 and Fig.2), is a group 3 unmanned aircraft system (UAS) as defined by the U.S. Department of Defense (DoD). This class of vehicles are typical of medium-sized, tactical vehicles, weighing between 55 lb and 1320 lb (see Tab.2). They are used for a variety of missions, including intelligence, surveillance, and reconnaissance (ISTAR), and can operate at altitudes up to 18 000 ft.

Table 1: Aircraft geometry and mass parameters.

Parameter	Value
Span (m)	5.4
AR	6
MAC (m)	0.9
MTOW (kg)	200

The BWB concept represents a transformative approach to subsonic transport design, promising significant advancements in aerodynamic efficiency, structural integration, and operational performance over conventional aircraft configurations. This design integrates the fuselage and wings into a seamless airframe, and has been the subject of extensive

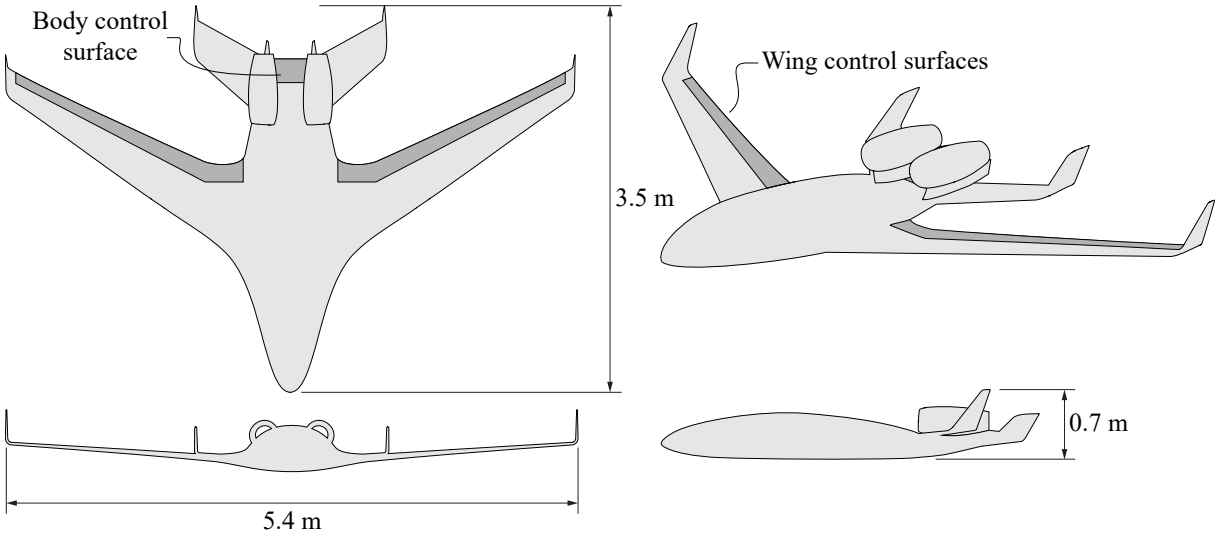


Figure 2: BWB aircraft demonstrator 3-view.

research and development since its inception in the early 1990s.

In the early work of Potsdam et al. [1], the critical design challenges of this configuration where identified, such as inboard wing design with thick transonic airfoils, a complex kink region blending inboard and outboard wings, and the need for a highly integrated multidisciplinary design approach to address cruise trim and stability. Subsequent work by Liebeck et al. [2] builds on this foundation, presenting a refined BWB concept for an 800-passenger commercial transport with a 7000 NM range. This study underscores the configuration's advantages—such as reduced wetted area, boundary layer ingestion by aft-mounted engines, and relaxed static stability—yielding substantial improvements in fuel efficiency compared to a conventional configuration. By 2003, the BWB concept had matured further, with a focus on design challenges related to commonality, speed, and flight mechanics [3]. Liebeck [4] emphasizes the BWB's adaptability, proposing a family of transports (200–600 passengers) with high parts commonality to enhance manufacturing efficiency. The concept also supports dual military and commercial applications, potentially reducing development costs through shared programs, showcasing its versatility beyond commercial transport.

Table 2: Categories of unmanned aircraft systems.

UAS Category	Max. Gross Takeoff Weight (lb)	Nominal Op. Altitude (ft)	Speed (kt)
Group 1	0 to 20	<1200 AGL	100
Group 2	21 to 55	<3500 AGL	<250
Group 3	<1320	<18 000 MSL	<250
Group 4	>1320	<18 000 MSL	any
Group 5	>1320	>18 000 MSL	any

The BWB offers notable efficiency and performance improvements over traditional aircraft, driven by its integrated design. However, aeroelastic challenges persist, requiring

ongoing research to fully realize its potential for future transport applications.

For a BWB configuration, which blends the fuselage and wings into a single lifting surface, flutter mechanisms might combine elements of both tube-and-wing and flying wing designs. Classical flutter, or bending-torsion flutter (BTF), is typical of tube-and-wing configurations and involves the coupling of bending and torsional modes of a wing, leading to divergent oscillations at critical speeds [5]. Limit cycle oscillation (LCO) is a bounded, non-divergent oscillatory response often linked to nonlinear aerodynamic effects or control system interactions, as observed in the F-16 aircraft [6]. Shock-induced oscillation (SIO), evident in the B-2's residual pitch oscillation (RPO), arises from unsteady transonic shock movements that couple with vehicle dynamics, producing sustained oscillations sensitive to Mach number and angle of attack [7]. Body freedom flutter (BFF) results from the coalescence of the rigid-body short-period mode with the wing's first bending mode, a phenomenon exacerbated by low pitch inertia and flexible structures [8–11]. This flutter mechanism is prevalent in tailless configurations, like flying wings, due to the absence of a traditional empennage which reduces pitch damping and static stability. The BWB's distributed lift and reduced wetted area could amplify transonic shock effects akin to the B-2's RPO, while its flexible, high-aspect-ratio wings might predispose it to body freedom flutter (BFF), especially if center of gravity placement and structural stiffness are not optimized. This was recently demonstrated by the works of Shi et al. [12] and Gu et al. [13] where a scaled BWB BFF flight demonstration vehicle was designed and tested successfully, revealing strong coupling between rigid-body motion and flexible wing deformation during preliminary flight tests. It was also shown that conventional aircraft configurations with very short fuselages may also be prone to BFF instability.

To better understand the flutter mechanisms specific to the Bombardier BWB configuration, the design space of wing structural properties was explored in this work with the goal of generating internal wing designs flexible enough to demonstrate flutter at low speeds. The following sections highlight the methods used to generate and study the flexible wing configurations followed by a discussion of the results. Finally, some conclusions are drawn from the work.

2 METHODOLOGY

Two aeroelastic models of the rigid aircraft were developed in NASTRAN and ASWING. The wing portion of the models was made flexible and parametrized to allow for the parametric studies to be conducted.

In NASTRAN, there are four solution sequences that relate to aeroelasticity. `SOL 144` solves quasi-steady aeroelastic problems. It can be used to obtain aerodynamic forces and pressures on the aerodynamic elements for different trim conditions. Aerodynamic loads are first calculated using the Doublet Lattice Method (DLM), although other aerodynamic solvers are available within NASTRAN. Spline functions transfer these loads onto the structural model and a linear finite element analysis calculates structural displacements. This solution sequence is also capable of static aeroelastic divergence analysis. NASTRAN also provides `SOL 145` and `SOL 146`, to solve aeroelastic flutter problems and dynamic aeroelasticity problems, respectively. Finally, the design sensitivity and optimization solution sequence `SOL 200` can be used for multidisciplinary analysis and design [14].

A NASTRAN generalized finite element model (GFEM) of the rigid wing was first converted into a hybrid GFEM-stick model where the composite wing elements were

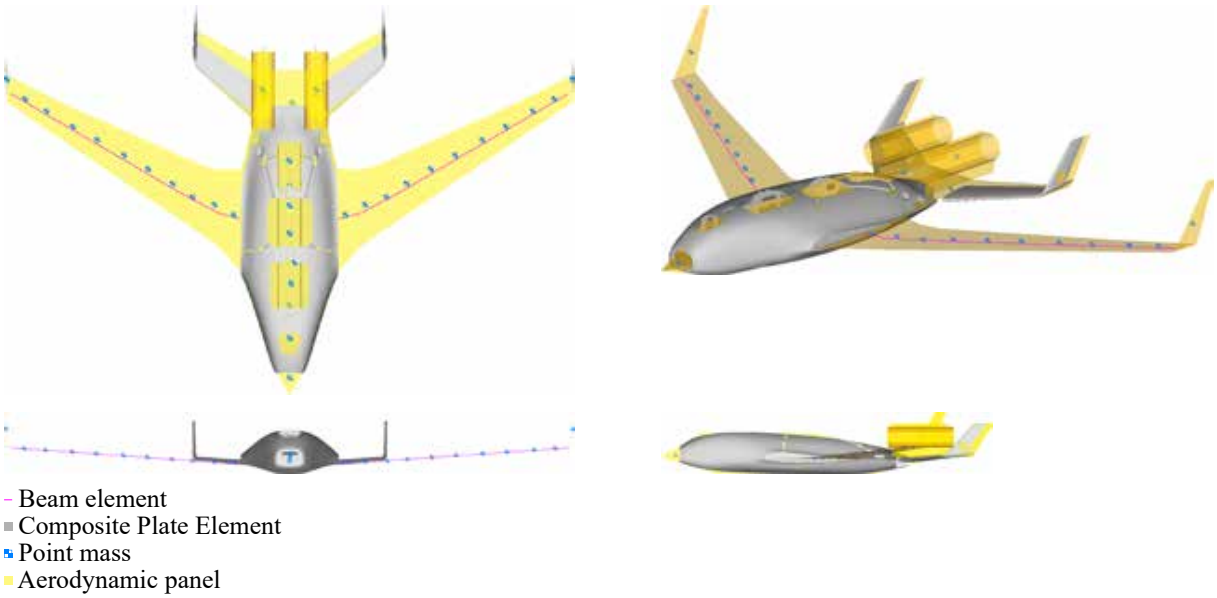


Figure 3: Aeroelastic NASTRAN model of the flexible BWB aircraft.

condensed into beam elements while the composite fuselage section remained unchanged. The model was also extended with aerodynamic panels and `SOL 145` was used to run flutter analyses. Fig.3 depicts the aeroelastic NASTRAN model of the aircraft.

A second aeroelastic model was built using ASWING. ASWING was developed by Drela at MIT for preliminary aircraft design of flexible aircraft [15] and can model nonuniform nonlinear connected beams allowing arbitrarily large deflections, with arbitrary loads [16, 17]. The aerodynamic formulation is a lifting-line model, with wind-aligned trailing vorticity discretized by a vortex-lattice type approach. Volume effects of slender bodies are modeled with source and doublet lines. Prandtl-Glauert compressibility corrections are also applied in wind-aligned axes. Arbitrary gust velocity fields can be imposed [18, 19].

The aerodynamic and structural formulations are fully coupled, with the aerodynamic loads causing structural deflections, and the structural deflections and twists affecting the aerodynamic loads.

The ASWING model of the BWB aircraft depicted in Fig.4 closely follows the NASTRAN model, in terms of its geometry, node distribution, material and section properties. The aerodynamics are inherently different, given the differences in the formulation used in ASWING and NASTRAN. However, the geometry of the ASWING lifting surfaces follows the NASTRAN aerodynamic panels as close as possible.

The model was calibrated to the NASTRAN model in terms of its elastic modal response. The cross-section properties of the flexible wing beam elements in ASWING were varied until its elastic response matched that of the NASTRAN model. A zero-speed modal analysis was set up and the fuselage was anchored to simulate boundary conditions of a cantilevered wing.

In order to obtain the required wing flexibility, this work proposes to replace the carbon fiber wingbox design of the existing rigid wing with a single aluminum spar spanning the entire length of the wing, and taylor its cross-section geometry to give more or less flexibility to the wing. To prevent the skin from adding stiffness to the carefully sized aluminum spar, the skin can be segmented into spanwise sections, which can then be fastened to the spar at single points along its span. This solution is akin to the one

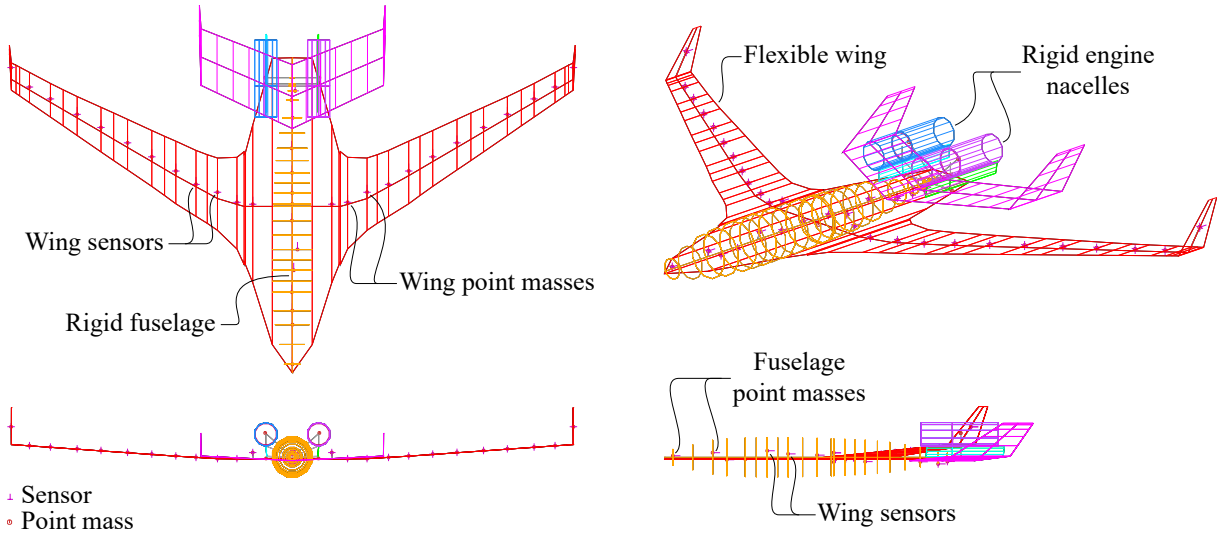


Figure 4: Aeroelastic ASWING model of the flexible BWB aircraft.

studied by Bras et al. [20] and a concept is depicted in Fig.5.

A sweep of parametric studies on the geometry of the spar cross-section as well as the mass of wingtip- and winglet-located mass points was conducted on the ASWING model. The geometric and mass parameters considered in the studies are shown in Fig.6.

A wrapper library for ASWING was developed in the form of a Python module that simplifies the process of loading models, setting up static trim and modal analyses, and post-processing of result text files. The library can automatically generate input scripts for ASWING based on a description of the analysis steps the user provides, runs the ASWING executable, writes the output to a temporary folder, parses the result files and finally populates appropriate Python data structures with the results.

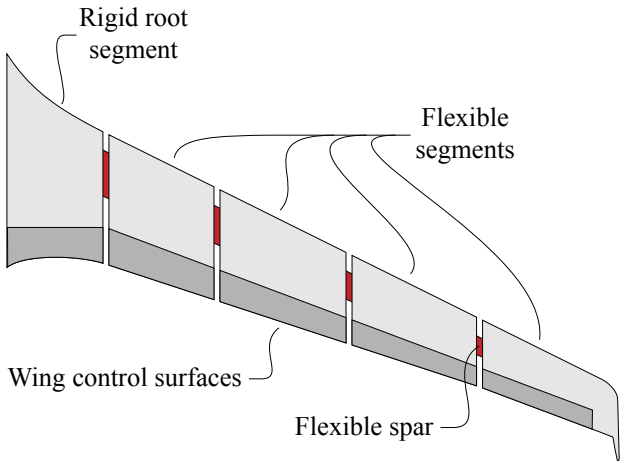
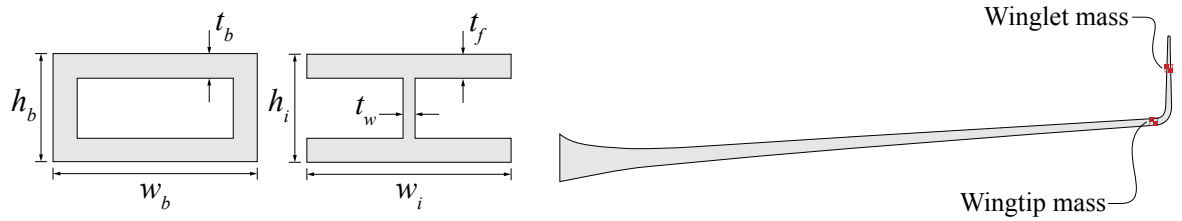


Figure 5: Segmented flexible wing design concept.

A second module was developed to ease the process of working with NASTRAN input and output files. This library abstracts a NASTRAN `BDF` model as a Python class and makes it easy to modify an existing model or create one from scratch. It currently implements tens of different NASTRAN bulk data cards, but the user can easily extend this by writing one data structure and two functions in a seperate file. At this time, the library



(a) Cross-section geometric parameters (b) Wingtip and winglet trim mass parameters

Figure 6: Flexible wing design parameters.

is capable of parsing output from `SOL 101`, `SOL 103`, and `SOL 145` solution sequences in the `F06` file and has been tested with NX and MSC flavors of NASTRAN.

Finally, an orchestrator script and a set of extra utilities in a separate Python module were developed to automate the parametric studies and aid in the post-processing of results. One such utility is a mode tracker and filter for ASWING and NASTRAN that aids in the generation of V–g–f and root locus plots.

The parametric studies focused on changing the size of different cross-section geometries for the spar, as well as changing the mass of wingtip- and winglet-located weights. An alternative would have been to change the spar cross-section bending and torsional stiffnesses directly (EI and GJ) instead of its geometry, but this would not have taken into account the internal size constraints of the wing. By changing the geometry parameters directly, it was possible to set limits for the parameter sweeps based on the actual available space inside the wing.

The studies were separated into two groups: the first group used the box spar design of Fig.6a on the left and a wingtip mass as shown in Fig.6b; the second group used an I spar design (Fig.6a on the right) and a winglet mass. A summary of the design parameters considered as well as their range of values is listed in Tab.3.

Table 3: Summary of design parameters for box- and I-shaped spar geometries.

Parameter	Range
Root width (w_b)	19 cm to 38 cm
Root height (h_b)	2.8 cm to 3.5 cm
Wingtip mass (m_{wingtip})	0 kg to 1 kg

(a) Box-shaped spar parameters

Parameter	Range
Root width (w_i)	19 cm to 38 cm
Root height (h_i)	2.1 cm to 3.5 cm
Flange thickness (t_f)	8 mm to 10 mm
Winglet mass (m_{winglet})	0 kg to 1 kg

(b) I-shaped spar parameters

The goal of the first group of sweeps was to explore the design space of parameters that are able to lower the first bending frequency of the wing enough for the mode to couple

with the highest-frequency rigid-body mode, usually the pitch, or short-period, mode, yielding a body-freedom flutter coupling mechanism. The box-shaped spar design was chosen for this set of studies because it keeps the torsional stiffness high while allowing for the bending stiffness to be reduced by changing its height, width, and thickness values. A wingtip trim mass aligned with the wing elastic axis was also chosen because it has the potential for lowering the bending frequency while keeping the torsional inertia unchanged.

The second group of parameter sweeps focused on exploring the design space of parameters that can lower the first torsional frequency of the wing enough for the mode to couple with the the first out-of-plane bending mode, yielding a bending-torsion flutter coupling mechanism. The I-shaped spar design was chosen because it has a low torsional stiffness due to its open section profile, and allows for both torsional and bending stiffnesses to be tweaked with the width, height, web and flange thickness parameters. An off-axis winglet trim mass was also chosen because it has the potential for increasing the torsional inertia, and decreasing the torsional frequency of the wing.

It should be noted that the box spar thickness and the I spar web thickness were kept constant in these studies.

3 DISCUSSION AND RESULTS

The ASWING model was first calibrated to the rigid NASTRAN model in terms of elastic modes. A comparison of the first five frequencies is presented in Tab.4 where a maximum error of $-5.0\,\%$ was observed for the second wing bending mode (W2B) and a minimum error of $-0.2\,\%$ was observed for the first wing bending (W1B) mode.

Table 4: Comparison of elastic modal frequencies between NASTRAN and ASWING models for cantilevered boundary conditions.

Mode	NASTRAN (Hz)	ASWING (Hz)	Error (%)
1 (W1B)	5.62	5.61	−0.2
2 (W1T)	18.11	18.42	1.7
3 (W2B)	22.16	21.06	−5.0
4 (W2T)	32.74	31.67	−3.3
5 (W3B)	42.46	41.97	−1.2

The matrix of parametric sweep studies performed is summarized in Tab.5a for the box-shaped spar, and Tab.5b for the I-shaped spar. For the box-shaped spar, three sets of sweeps—box width w_b (sweep 1), box height h_b (sweep 2), and wingtip mass m_{wingtip} (sweep 3)—were performed, whereas for the I-shaped spar, the sweeps were comprised of section width w_i (sweep 4), section height h_i (sweep 5), winglet mass m_{winglet} (sweep 6), and flange thickness t_f (sweep 7). For the box beam, the thickness was kept constant at 1.5 mm while for the I beam the web thickness was kept constant at 1 mm.

Fig.7 through Fig.9 show the linear and angular displacement results obtained from static aeroelastic analyses, V–g–f plots and aeroelastic root locus results obtained from aeroelastic modal analyses, for width, height, and wingtip mass sweep studies of the box-shaped spar.

The linear and angular displacements shown in the figures are the maximum values

Table 5: Summary of parametric sweeps.

Parameter	Sweep 1			Sweep 2			Sweep 3		
w_b (cm)	19	28.5	38	19	19	19	19	19	19
h_b (cm)	3.5	3.5	3.5	2.8	3.15	3.5	3.5	3.5	3.5
m_{wingtip} (kg)	0	0	0	0	0	0	0	0.5	1

(a) Box-shaped spar sweeps

Parameter	Sweep 4			Sweep 5			Sweep 6			Sweep 7		
w_i (cm)	19	28.5	38	19	19	19	19	19	19	19	19	19
h_i (cm)	3.5	3.5	3.5	2.1	2.8	3.5	3.5	3.5	3.5	3.5	3.5	3.5
m_{winglet} (kg)	0	0	0	0	0	0	0	0.5	1	0	0	0
t_f (mm)	10	10	10	10	10	10	10	10	10	8	9	10

(b) I-shaped spar sweeps

across the entire span, which occur at the wing tip, and are presented for positive and negative limit load cases. For the box width and height parameters, w_b and h_b, respectively, the sweeps show a small increase in displacement and twist angle with a decrease in parameter values (Fig.7a, Fig.8a). In the case of wingtip mass, m_{wingtip}, the static results show a negligible variation of displacement and twist angle (Fig.9a).

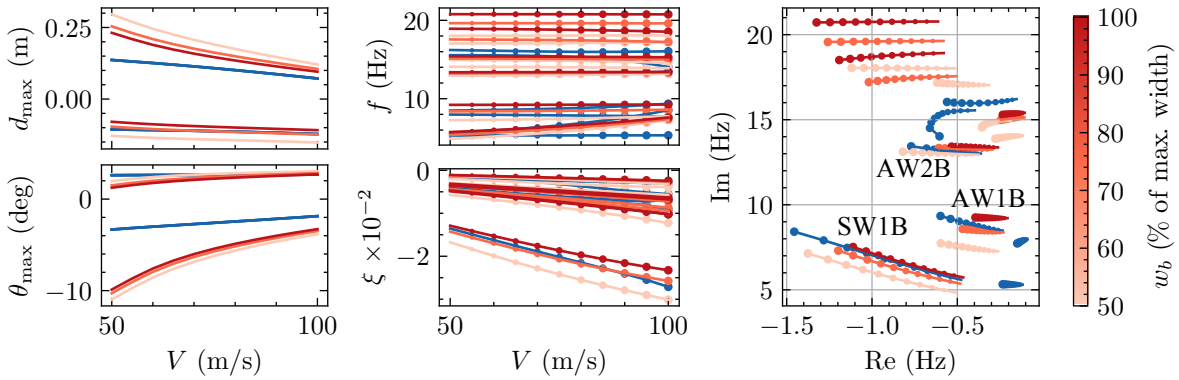
In terms of aeroelastic modes, both V–g–f and root locus plots show that a decrease in both width and height parameters brings the frequency and damping of the first wing symmetric (SW1B) and anti-symmetric (AW1B) bending modes down while an increase in wingtip mass appears to decrease the frequency of both modes while keeping the damping ratio of the symmetric mode constant.

Despite the decrease in frequency of mode SW1B with variation of all three parameters, no coupling with the pitch, or short-period (SP) mode was observable.

For the second group of parameter sweeps, Fig.10 through Fig.12 show the linear and angular displacement results, V–g–f plots, and root locus of aeroelastic modes for width, height, and winglet mass sweep studies of the I-shaped spar.

For the flange width parameter, w_i, the static results show a small increase in displacement and a small decrease in twist angle with a decrease in flange width (Fig.10a). For the cross-section height parameter, h_i, the static results show a larger increase in displacement and twist angle with a decrease in cross-section height (Fig.11a), similarly to the box cross-section height sweep results. In the case of winglet mass, m_{winglet}, the static results show negligible variation of wing displacement and twist angle, similarly to the wingtip mass sweep results (Fig.12a)

From the modal analyses, both V–g–f and root locus plots show that a decrease in cross-section width decreases the frequency and damping of the first wing symmetric (SW1T) and anti-symmetric (AW1T) torsion modes, while an a decrease in cross-section height has a similar effect as for the box-shaped spar, i.e., a decrease in the frequency and damping of the first wing symmetric (SW1B) and anti-symmetric (AW1B) bending modes. Finally, an increase in winglet mass appears to be very effective in bringing the frequency of the first wing symmetric (SW1T) and anti-symmetric (AW1T) torsion modes

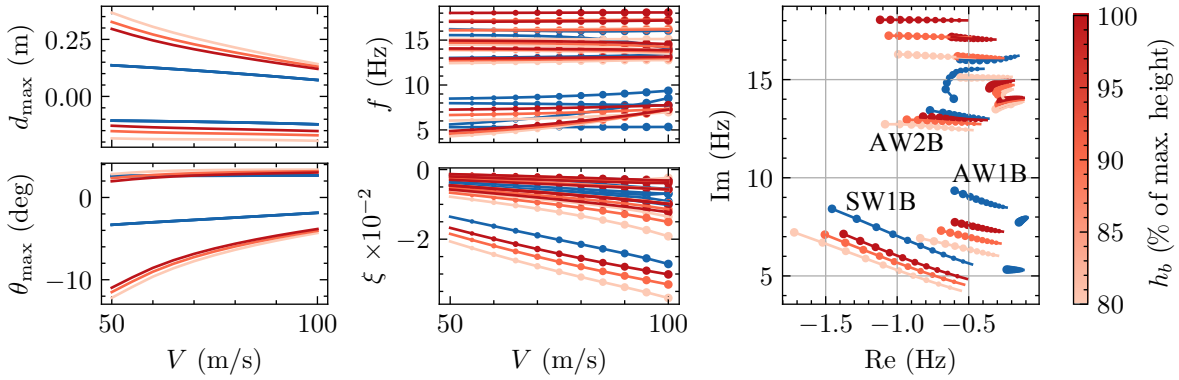


(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 7: Box-shaped spar with fixed height and thickness and varying width. Results in blue are baseline values for the rigid wing.

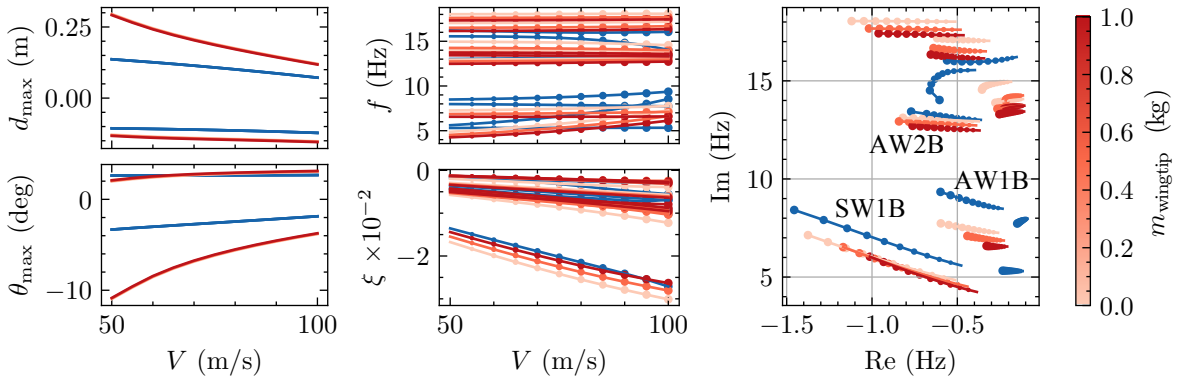


(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 8: Box-shaped spar with fixed width and thickness and varying height. Results in blue are baseline values for the rigid wing.

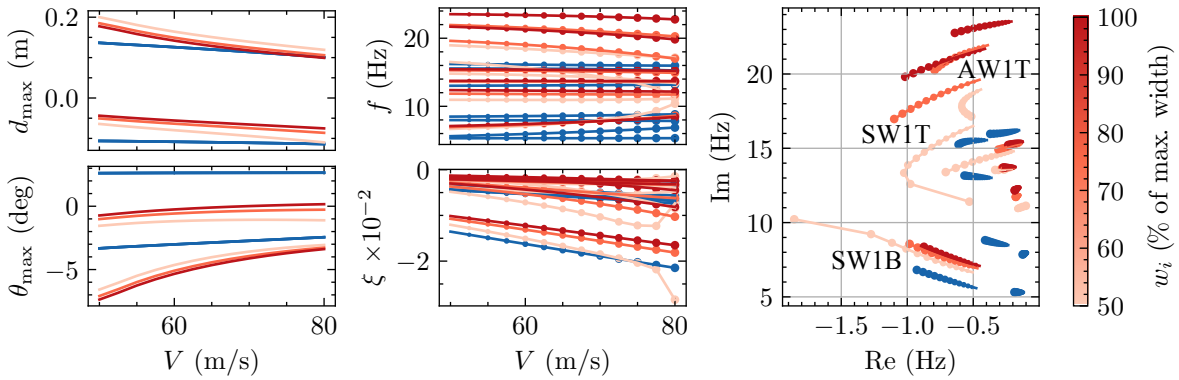


(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 9: Box-shaped spar with fixed width, height, and thickness and varying wingtip mass. Results in blue are baseline values for the rigid wing.

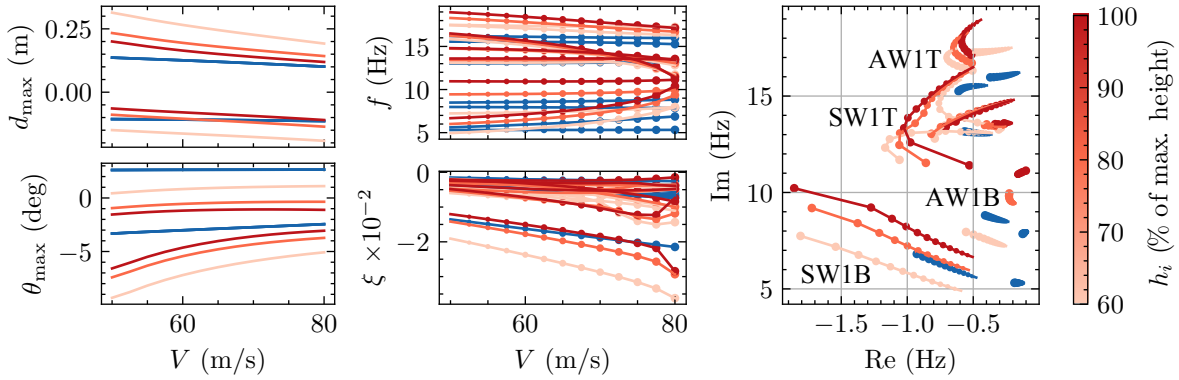


(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 10: I-shaped spar with fixed height, web and flange thicknesses and varying width. Results in blue are baseline values for the rigid wing.

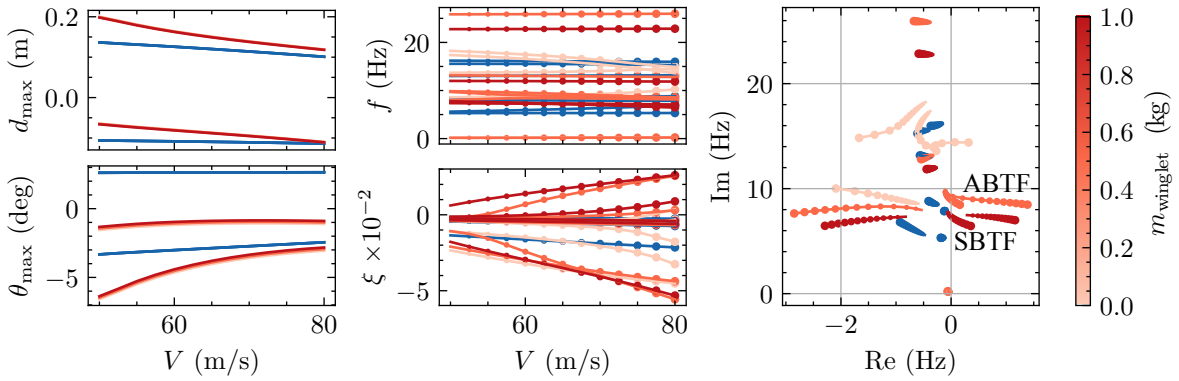


(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 11: I-shaped spar with fixed width, web and flange thicknesses and varying height. Results in blue are baseline values for the rigid wing.



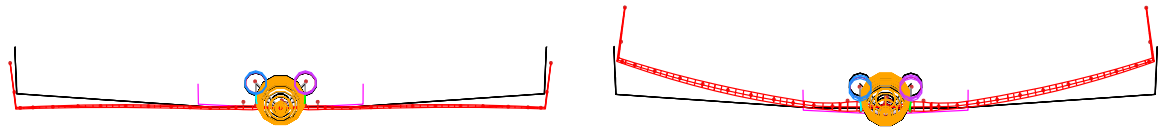
(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 12: I-shaped spar with fixed width, height, web and flange thicknesses and varying winglet mass. Results in blue are baseline values for the rigid wing.

down to the point where they coalesce with bending modes and form symmetric (SBTF) and antisymmetric (ABTF) bending-torsion flutter modes. The shape of the formed BTF flutter mode is depicted in Fig.13 for the symmetric case.

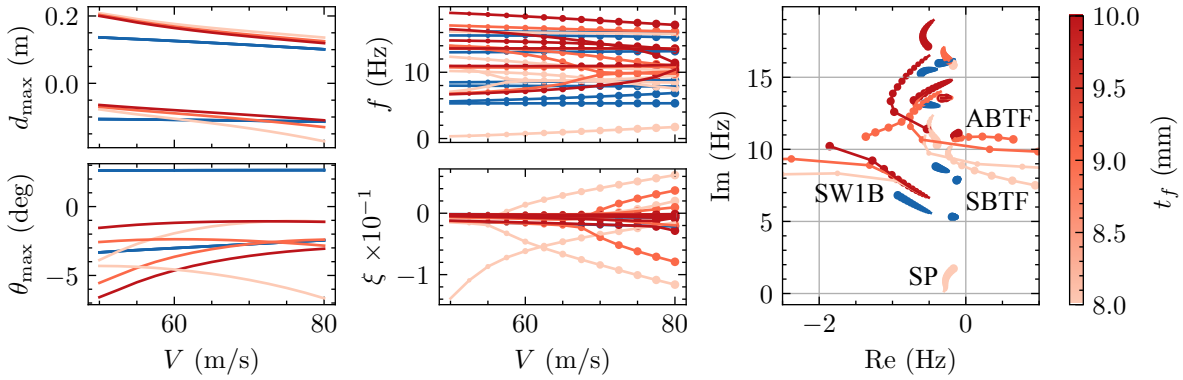


$\phi = 90°$ $\phi = 270°$

Figure 13: BTF mode shape at two different phase angles ϕ.

The last set of parameter sweeps in this group consisted of varying the flange thickness t_f of the I beam cross section. The results are shown in Fig.14 for both the static and modal aeroelastic analyses carried out. The static results show a small increase in displacement, in an absolute sense, with a decrease in flange thickness, for both load cases. A flange thickness decrease appears to cause a decrease in maximum twist angle, i.e., twist angle becomes less negative, for the positive load case. For the negative load case, the maximum twist angle increases in an absolute sense, i.e., becomes more negative, with a decrease in flange thickness.

In terms of aeroelastic modes, both V–g–f and root locus plots show that two bending-torsion flutter modes appear—the symmetric (SBTF) and the anti-symmetric (ABTF) one.



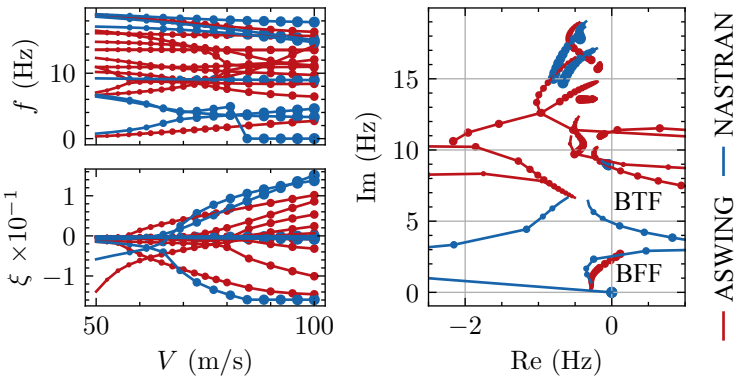
(a) Max. linear and angular displacements

(b) Freq. and damp. ratios of aeroelastic modes

(c) Root locus of aeroelastic modes

Figure 14: I-shaped spar with fixed width, height, and web thickness and varying flange thickness. Results in blue are baseline values the rigid wing.

One interesting aspect of these sweeps is the locus of the pitch, or short-period (SP), aeroelastic mode. The mode is visible in Fig.14c, as well as in Fig.14b where the frequency increases linearly with speed starting from 0 Hz—this is characteristic of a BFF mode. While the mode is stable within this range of speeds, it appears it would become unstable at higher speeds. To test this, another set of runs was prepared. This time, the flange thickness was kept at 8 mm, its lowest value, and the range of speeds for the run was increased to 100 m/s. In addition, a similar wing configuration was prepared in NASTRAN and a flutter (`SOL 145`) analysis was conducted for the same range of speeds. The results are shown in Fig.15.



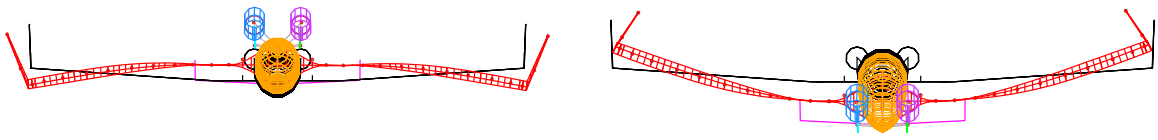
(a) Freq. and damp. ratios of aeroelastic modes

(b) Root locus of aeroelastic modes

Figure 15: Comparison between ASWING and NASTRAN aeroelastic modal results for an I-shaped spar with reduced flange thickness.

Comparing the results from the V–g–f and root locus plots there are some evident differences in the behavior of the two models. While both ASWING and NASTRAN predict the first wing symmetric and anti-symmetric bending frequencies with agreement at approximately 6 Hz, the first torsional frequency prediction shows more error (approximately 9 Hz in ASWING and 11 Hz in NASTRAN). Also, the number of predicted modes in ASWING up to 15 Hz is much higher compared to NASTRAN and the flutter mechanisms the two models predict differ as well—ASWING shows a BTF mechanism at approximately 60 m/s while NASTRAN does not.

Despite these differences, there is an interesting phenomenon present in both models—both predict the onset of what appears to be a BFF mode (NASTRAN shows the onset to occur at approximately 60 m/s while ASWING predicts it at a higher speed of approximately 90 m/s). It's also interesting that this mode did not manifest in the first group of sweeps, i.e., when the first bending frequency of the wing was being lowered while keeping its torsional frequency high. Only when both torsional and bending frequencies of the wing are being lowered does the mode manifest. The modeshape as predicted from ASWING past the flutter onset point is depicted in Fig.16 for two phase angles. It's possible to see both bending and torsional motion of the wing coupling with plunging and pitching of the fuselage. This differs from the typical BFF mode of flying wings where only the wing bending motion couples with aircraft plunging (and possible pitching, depending on sweep angle), highlighting the need to conduct further studies to better understand the flutter mechanisms of the BWB configuration.



$\phi = 90°$ $\phi = 270°$

Figure 16: BFF mode shape at two different phase angles ϕ.

4 CONCLUSIONS

In this work, aeroelastic models of a flexible BWB flight demonstrator were developed using NASTRAN and ASWING to investigate flutter mechanisms at low speeds. Parametric studies on wing spar geometry and wingtip and winglet masses were conducted to tailor the wing structural properties to induce specific flutter types: body-freedom flutter and bending-torsion flutter. The emergence of what appears to be a BFF mode in certain configurations highlights the unique aeroelastic challenges of BWB designs. Comparisons between ASWING and NASTRAN revealed discrepancies in predicted flutter speeds and mechanisms, emphasizing the need for model validation. This work is a step forward in the understanding of BWB aeroelasticity and the design of flexible flight demonstrators, enabling safe, low-speed flutter testing to support the advancement of sustainable aviation.

ACKNOWLEDGEMENTS

The authors acknowledge Fundação para a Ciência e a Tecnologia (FCT), through IDMEC, under LAETA, project UIDB/50022/2020. A. Suleman also acknowledges the NSERC Canada Research Chair and Discovery funding programs. The authors are also grateful to Bombardier in the framework of a NSERC Alliance program ALLRP 578448-22.

REFERENCES

[1] M. Potsdam, M. Page, and R. Liebeck. Blended wing body analysis and design. In *15th Applied Aerodynamics Conference*, 1997. doi:10.2514/6.1997-2317.

[2] R. Liebeck, M. Page, and B. Rawdon. Blended-wing-body subsonic commercial transport. In *36th AIAA Aerospace Sciences Meeting and Exhibit*, 1998. doi:10.2514/6.1998-438.

[3] R. Liebeck. Blended wing body design challenges. In *AIAA International Air and Space Symposium and Exposition: The Next 100 Years*, 2003. doi:10.2514/6.2003-2659.

[4] R. Liebeck. Design of the blended wing body subsonic transport. *Journal of Aircraft*, 41(1):10–25, 2004. doi:10.2514/1.9084.

[5] W. Bollay and C. Brown. Some experimental results on wing flutter. *Journal of the Aeronautical Sciences*, 8(8):313–318, 1941. doi:10.2514/8.10724.

[6] P. Chen, D. Sarhaddi, and D. Liu. Limit-cycle oscillation studies of a fighter with external stores. In *39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit*, 1998. doi:10.2514/6.1998-1727.

[7] S. Jacobson, R. Britt, D. Freim, and P. Kelly. Residual pitch oscillation (RPO) flight test and analysis on the b-2 bomber. In *39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit*, 1998. doi:10.2514/6.1998-1805.

[8] J. Banerjee. Flutter characteristics of high aspect ratio tailless aircraft. *Journal of Aircraft*, 21(9):733–736, 1984. doi:10.2514/3.45022.

[9] J. Banerjee. Flutter modes of high aspect ratio tailless aircraft. *Journal of Aircraft*, 25(5):473–476, 1988. doi:10.2514/3.45607.

[10] L. Niblett. The fundamentals of body-freedom flutter. *The Aeronautical Journal*, 90 (899):373–377, 1986. doi:10.1017/S0001924000015979.

[11] M. Love, P. Zink, P. Wieselmann, and H. Youngren. Body freedom flutter of high aspect ratio flying wings. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2005. doi:10.2514/6.2005-1947.

[12] P. Shi, F. Liu, Y. Gu, and Z. Yang. The development of a flight test platform to study the body freedom flutter of bwb flying wings. *Aerospace*, 8(12):390, 2021. doi:10.3390/aerospace8120390.

[13] Y. Gu, P. Shi, X. Tian, and Z. Yang. Body freedom flutter of scaled vehicles: From blended wing body to conventional configuration. In *International Forum on Aeroelasticity and Structural Dynamics (IFASD) 2024*, 2024.

[14] M. S. Corporation. *MSC Nastran 2018.2: Aeroelastic Analysis User's Guide*, 2018.

[15] M. Drela. Integrated simulation model for preliminary aerodynamic, structural, and control-law design of aircraft. In *40th Structures, Structural Dynamics, and Materials Conference and Exhibit*, 1999. doi:10.2514/6.1999-1394.

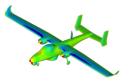
[16] P. Minguet and J. Dugundji. Experiments and analysis for composite blades under large deflections part I: Static behavior. *AIAA Journal*, 28(9):1573–1579, 1990. doi:10.2514/3.25255.

[17] P. Minguet and J. Dugundji. Experiments and analysis for composite blades under large deflections part II: Dynamic behavior. *AIAA Journal*, 28(9):1580–1588, 1990. doi:10.2514/3.25256.

[18] M. Drela. *ASWING 5.99 Technical Description: Steady Formulation*, 2015.

[19] M. Drela. *ASWING 5.99 Technical Description: Unsteady Extension*, 2015.

[20] M. Bras, S. Warwick, and A. Suleman. Aeroelastic evaluation of a flexible high aspect ratio wing UAV: Numerical simulation and experimental flight validation. *Aerospace Science and Technology*, 122:107400, 2022. doi:10.1016/j.ast.2022.107400.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

FREQUENCY-AWARE SURROGATE MODELING WITH SMT KERNELS FOR ADVANCED DATA FORECASTING

N. Gonel[1,2*], P. Saves[3], J. Morlier[4]

1: Fédération ENAC ISAE-SUPAERO ONERA,
Université de Toulouse,
31000, Toulouse, France.

2: IFP Energies nouvelles,
Institut Carnot IFPEN Transports Energie,
92852 Rueil-Malmaison, France.
nicolas.gonel@ifpen.fr

3: IRIT,
Université de Toulouse, CNRS, Toulouse INP, UT3, UT Capitole,
Toulouse, France.
paul.saves@irit.fr

4: ICA,
ISAE–SUPAERO, INSA, CNRS, MINES ALBI, UPS,
Université de Toulouse,
31400 Toulouse, France.
joseph.morlier@isae-supaero.fr

Abstract. *This paper introduces a comprehensive open-source framework for developing correlation kernels, with a particular focus on user-defined and composition of kernels for surrogate modeling. By advancing kernel-based modeling techniques, we incorporate frequency-aware elements that effectively capture complex mechanical behaviors and time-frequency dynamics intrinsic to aircraft systems. Traditional kernel functions, often limited to exponential-based methods, are extended to include a wider range of kernels such as exponential squared sine and rational quadratic kernels, along with their respective first- and second-order derivatives. The proposed methodologies are first validated on a sinus cardinal test case and then applied to forecasting Mauna-Loa Carbon Dioxide (CO_2) concentrations and airline passenger traffic. All these advancements are integrated into the open-source Surrogate Modeling Toolbox (SMT 2.0), providing a versatile platform for both standard and customizable kernel configurations. Furthermore, the framework enables the combination of various kernels to leverage their unique strengths into composite models tailored to specific problems. The resulting framework offers a flexible toolset for engineers and researchers, paving the way for numerous future applications in metamodeling for complex, frequency-sensitive domains.*

Keywords: Frequency correlation kernels, Gaussian process, Aircraft data prediction, Surrogate Modeling Toolbox, Open-Source Software

1 INTRODUCTION

Over the past few decades, the aeronautical industry has witnessed significant advancements driven primarily by incremental improvements in aircraft design optimization [1]. The design process for aeronautic systems is inherently complex and multidisciplinary, requiring a careful balance between various domains such as aerodynamics, structural integrity, geometry, performance, weight management, and dynamic stability [2]. This integrated process, referred to as Multidisciplinary Design Analysis (MDA), lays the foundation for the iterative improvements carried out in Multidisciplinary Design Optimization (MDO) [3, 4]. In our approach to MDO, we adopt the multidisciplinary design feasible strategy, which ensures that the equilibrium between disciplines is re-evaluated and maintained, through a fixed point algorithm, at every modification of the design variables during optimization. However, this approach significantly increases computational demands, posing a major challenge.

To address these challenges, research has increasingly turned to surrogate modeling techniques, which offer an efficient alternative to high-fidelity simulations [5–7].Sofware such as the Surrogate Modeling Toolbox (SMT) have been developed to reduce computational costs by leveraging Gaussian Process (GP) interpolation and Bayesian Optimization (BO) methods [8, 9]. However, traditional kernel functions used in GP models—primarily based on exponential formulations—often struggle to capture the complex cyclic and high-dimensional phenomena encountered in advanced aerospace applications [10].

This paper introduces a comprehensive open-source framework for the development of advanced correlation kernels, specifically designed for surrogate modeling in frequency-sensitive environments. Our framework extends conventional methods by incorporating frequency-aware elements into the kernel design. In doing so, we broaden the list of available kernels to include not only standard forms but also alternatives such as exponential squared sine and rational quadratic kernels, along with their first and second-order derivatives. These enhancements enable the efficient optimization of surrogate models that address intricate mechanical behaviors and time-frequency dynamics, which are crucial for analyzing vibrational and aerodynamic characteristics of aircraft structures.

The proposed methodologies are validated through a series of test cases, beginning with a sinus cardinal example and extending to real-world applications such as forecasting Mauna-Loa CO_2 concentrations and airline passenger traffic. All these advancements are encapsulated within SMT 2.0, offering a versatile platform that supports both standard and user-defined kernel configurations. By allowing the combination of multiple kernel types, our approach enables the construction of composite models that leverage the unique strengths of each kernel to address specific modeling challenges.

This extended framework not only enhances predictive capabilities in the preliminary design phase of aerospace engineering but also paves the way for future applications in diverse fields where complex, frequency-sensitive phenomena are of interest. The remainder of the paper is organized as follows: Section 2 reviews the background of surrogate modeling and kernel methods in aerospace applications; Section 3 details the theoretical underpinnings and implementation of our advanced kernel framework; Section 4 presents validation results on benchmark cases and real-world applications; and Section 5 concludes with a discussion of future directions and potential extensions.

2 GAUSSIAN PROCESSES

A Gaussian Process (GP) is a stochastic process, *i.e.* a collection of random variables, where any finite collection of those variables has a joint Gaussian distribution. Hereinafter, we detail the mathematical aspects of GP interpolation, also known as Kriging [11]. First, the GP distribution is completely determined by its mean and covariance functions [12]. Formally, let f denote the (scalar) target function we wish to estimate:

$$\begin{array}{rccc} f & : & \mathcal{D} \subseteq \mathbb{R}^d & \rightarrow & \mathbb{R} \\ & & \mathbf{x} & \mapsto & f(\mathbf{x}) \end{array} \tag{1}$$

where $\mathbf{x}$ is a vector of input variables from the input set $\mathcal{D}$ called the *design space*. In the function-space view, we assume that $f(\mathbf{x})$ is a realization of a GP Y with mean function m and covariance function k:

$$Y \sim \mathcal{GP}(m, k) \tag{2}$$

Indeed, the GP defines a probability distribution (called the *prior*) over possible functions that fit the set of observations. The main advantage of the GP modeling is that any finite set of n input points $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\} \subset \mathcal{D}$, the corresponding observations follow a multivariate Gaussian distribution:

$$\underbrace{\begin{bmatrix} f(\mathbf{x}_1) \\ f(\mathbf{x}_2) \\ \vdots \\ f(\mathbf{x}_n) \end{bmatrix}}_{\text{outputs vector } \mathbf{f}} \sim \mathcal{N}\left(\underbrace{\begin{bmatrix} m(\mathbf{x}_1) \\ m(\mathbf{x}_2) \\ \vdots \\ m(\mathbf{x}_n) \end{bmatrix}}_{\text{mean vector } \mathbf{m}}, \underbrace{\begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \cdots & k(\mathbf{x}_1, \mathbf{x}_n) \\ k(\mathbf{x}_1, \mathbf{x}_2)^T & k(\mathbf{x}_2, \mathbf{x}_2) & \cdots & k(\mathbf{x}_2, \mathbf{x}_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_1, \mathbf{x}_n)^T & k(\mathbf{x}_2, \mathbf{x}_n)^T & \cdots & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix}}_{\text{covariance matrix } \mathbf{K}} \right). \tag{3}$$

We can also write it more compact using the shorthands $\mathbf{f} = f(\mathbf{X})$, $\mathbf{m} = m(\mathbf{X})$ and $\mathbf{K} = k(\mathbf{X}, \mathbf{X})$:

$$p(f(\mathbf{x}_1), ..., f(\mathbf{x}_N)) = p(\mathbf{f} \mid \mathbf{X}) = \mathcal{N}(\mathbf{m}, \mathbf{K}). \tag{4}$$

We highlight the fact that the covariance matrix $\mathbf{K}$ (also called *Gram matrix*) is a Positive Semi-Definite (PSD) and symmetric $n \times n$ matrix [13] whose components are given by the covariance function k: $K_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. An example of a valid (stationary) covariance function is the anisotropic[1] Squared Exponential (SE) kernel[2]:

$$K^{\theta_i}(\mathbf{x}_i, \mathbf{x}_j) = \sigma^2 \prod_{i=1}^{n} \exp\left(-\theta_i \left(\mathbf{x}_i - \mathbf{x}_j\right)^2\right). \tag{5}$$

where σ^2 is the process variance parameter and $\boldsymbol{\theta} = \{\theta_1, ..., \theta_n\}$ is the vector of inverse length-scale parameters. We call these quantities the hyperparameters of the GP that are needed to be estimated in order to best fit the model to our target function.

To make predictions $\mathbf{f}_*$ for a new set of n_* inputs $\mathbf{X}_*$ (independent of $\mathbf{X}$), we first write the joint distribution of $\mathbf{f}$ and $\mathbf{f}_*$ which also follows a multivariate Gaussian distribution:

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{m} \\ \mathbf{m}_* \end{bmatrix}, \begin{bmatrix} \mathbf{K} & \mathbf{K}_* \\ \mathbf{K}_*^\top & \mathbf{K}_{**} \end{bmatrix} \right). \tag{6}$$

[1] In the isotropic version, θ is scalar so the length-scale is the same for each dimension.

[2] This kernel is also denoted as the Radial Basis Function (RBF) kernel

where $\mathbf{K}_* = k(\mathbf{X}, \mathbf{X}_*)$ denotes the $n \times n_*$ matrix of covariance between all pairs of training and test points, and similarly $\mathbf{K}_{**} = k(\mathbf{X}_*, \mathbf{X}_*)$ is the covariance matrix of the test set.

In practice, the true function value value $f(\mathbf{x})$ is often not directly observable. Instead, we observe a noisy output $y(\mathbf{x}) = f(\mathbf{x}) + \varepsilon$. We assume that this additive noise is an independent and identically distributed (i.i.d.) Gaussian noise [14] and we denote its variance η^2. Using the previous shorthands and letting $\mathbf{y}$ denotes the vector of noisy observations at inputs $\mathbf{X}$, the joint distribution (6) then becomes:

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{m} \\ \mathbf{m}_* \end{bmatrix} , \begin{bmatrix} \mathbf{K} + \eta^2 \mathbf{I} & \mathbf{K}_* \\ \mathbf{K}_*^\top & \mathbf{K}_{**} \end{bmatrix} \right). \tag{7}$$

Thus, we compute the *posterior* distribution over functions by conditioning the joint prior (7) with the observed data. Using the properties of a multivariate normal distribution, we obtain the predicted distribution:

$$\mathbf{f}_* \mid \mathbf{y}, \mathbf{X}, \mathbf{X}_* \sim \mathcal{N}\left(\mathbf{m}_* + \mathbf{K}_*^\top \left[\mathbf{K} + \eta^2 \mathbf{I}\right]^{-1} (\mathbf{y} - \mathbf{m}), \mathbf{K}_{**} - \mathbf{K}^\top \left[\mathbf{K} + \eta^2 \mathbf{I}\right]^{-1} \mathbf{K}_* \right). \tag{8}$$

For the ease of notations, we will then only consider centered GPs, *i.e.* zero mean functions: $\mathbf{m} = \mathbf{m}_* = \mathbf{0}$, but equations can be generalized for more general cases with constant (simple Kriging), linear (ordinary Kriging) or polynomial (universal Kriging) priors, adding hyperparameters to the model [12].
Using the usual formalism of GP, we can rewrite the conditional distribution $\mathbf{f}_* \mid \mathbf{y}$ as:

$$Z_n \triangleq \tilde{Y} \left| \left\{ \tilde{Y}(\mathbf{x}_i) = y(\mathbf{x}_i) \ , \ 1 \leq i \leq n \right\} \sim \mathcal{GP}(\mu_n, c_n). \right. \tag{9}$$

where the posterior mean and covariance functions are defined as follows:

$$\mu_n(\mathbf{x}) \triangleq k(\mathbf{x}, \mathbf{X}) \left[\mathbf{K}^{-1} + \eta^2 \mathbf{I}\right] \mathbf{y}. \tag{10}$$

$$c_n(\mathbf{x}, \mathbf{x}') \triangleq k(\mathbf{x}', \mathbf{x}') - k(\mathbf{x}', \mathbf{X}) \left[\mathbf{K} + \eta^2 \mathbf{I}\right]^{-1} k(\mathbf{X}, \mathbf{x}'). \tag{11}$$

An alternative way of formulating the above functions, implemented in Surrogate Modeling Toolbox (SMT) [8, 15], is to introduce the correlation function R and to consider a unique GP variance σ^2 that is an homoscedastic process:

$$\mathrm{Cov}\left(Y(\mathbf{x}), Y(\mathbf{x}')\right) = k(\mathbf{x}, \mathbf{x}') = \sigma^2 R(\mathbf{x}, \mathbf{x}'). \tag{12}$$

With this new notation, the equations (10) and (11) can be rewritten as follows [16]:

$$\mu_n(\mathbf{x}) = \mathbf{r}_n^\top(\mathbf{x}) \mathbf{R}_n^{-1} \mathbf{y}. \tag{13}$$

$$c_n(\mathbf{x}, \mathbf{x}') = \sigma^2 \left[r(\mathbf{x}, \mathbf{x}') - \mathbf{r}_n^\top(\mathbf{x}) \mathbf{R}_n^{-1} \mathbf{r}_n(\mathbf{x}') \right]. \tag{14}$$

where $\mathbf{y} = \left[y(\mathbf{x}_i), ..., y(\mathbf{x}_n)\right]^\top$ is the observations vector, $\mathbf{r}_n(\mathbf{x}) = k(\mathbf{X}, \mathbf{x}) = \left(k(\mathbf{x}_i, \mathbf{x})\right)_{1 \leq i \leq n}$ is the cross-covariance vector and $\mathbf{R}_n = \mathbf{K} + \eta^2 \mathbf{I} = \left(k(\mathbf{x}_i, \mathbf{x}_j)\right)_{1 \leq i,j \leq n}$ is the covariance matrix Note that if we consider noise-free observations (*i.e.* $\eta^2 = 0$), we simply have $\mathbf{R}_n = \mathbf{K}$. The conditional process defined by Eq.(12) is then used as a surrogate model to approximate our target function f.

To conclude this short presentation of the GP regression model, we address the optimization of the hyperparameters mentioned above. Indeed, it is necessary to tune the

model parameters $\sigma^2, \boldsymbol{\theta}$ and η^2 on the training set $\mathbf{X}$ before making predictions at new inputs $\mathbf{X}_*$. Here, everything being Gaussian and our prior being an uniform distribution between given bounds, we only optimize the *marginal likelihood* $p(\mathbf{y})$ to obtain our optimal hyperparameters [17]. This likelihood is given by:

$$p(\mathbf{y}) = \mathcal{N}(\mathbf{0}, \mathbf{K} + \eta^2\mathbf{I}) = \frac{1}{(2\pi)^{N/2}\det(\mathbf{K} + \eta^2\mathbf{I})^{1/2}} \exp\left(-\frac{1}{2}\mathbf{y}^\top \left(\mathbf{K} + \eta^2\mathbf{I}\right)\mathbf{y}\right). \tag{15}$$

The optimal set of hyperparameters is then obtained by maximizing the log marginal likelihood:

$$\log p(\mathbf{y}) = -\frac{1}{2}\log\left(\det\left(\mathbf{K} + \eta^2\mathbf{I}\right)\right) - \frac{1}{2}\mathbf{y}^\top \left(\mathbf{K} + \eta^2\mathbf{I}\right)^{-1}\mathbf{y} - \frac{N}{2}\log(2\pi). \tag{16}$$

We recall that the marginal likelihood $p(\mathbf{y})$ is the integral of the *likelihood* $p(\mathbf{y}|\mathbf{f}, \mathbf{X})$ times the prior $p(\mathbf{f}|\mathbf{X})$. Thus, given the prior $\mathbf{f}|\mathbf{X} \sim \mathcal{N}(0, \mathbf{K})$ and the likelihood $\mathbf{y}|\mathbf{f} \sim \mathcal{N}(\mathbf{f}, \eta^2\mathbf{I})$, the computation leads to (15). We saw that to define the GP, we need to define its kernel first. This kernel will give the structure of the GP. We already showed the squared exponential kernel but other kernels exists in SMT such as the absolute exponential, the Matern kernel or the power exponential kernel.

3 A NOVEL USER-DEFINED KERNEL LIBRARY FOR GAUSSIAN PROCESS

The SMT 2.0 interface is more user-friendly and offers improved and detailed documentation for both users and developers. Hosted publicly, SMT 2.0 can be directly imported within Python scripts, is released under the New BSD license, and runs on Linux, macOS, and Windows operating systems. Furthermore, regression tests are automatically executed on each supported platform whenever changes are committed to the repository, ensuring robustness and compatibility. To enhance the modeling capabilities of Gaussian Processes (GP), we have expanded SMT by implementing two new kernels:

- **Rational Quadratic Kernel**: This kernel is particularly effective for modeling functions that exhibit irregularities. It can be interpreted as a scale mixture of squared exponential kernels, thereby allowing for varying length scales. Its formula is given by:

$$K^{\theta_l,\theta_k}(x, x') = \left(1 + \frac{(x - x')^2}{\theta_l}\right)^{-\theta_k} \tag{17}$$

  where θ_l controls the length scale and θ_k adjusts the weighting across scales.

- **Periodic Kernel**: Useful for modeling periodic data, this kernel captures cyclic patterns inherent in many physical and time-series processes. Its formula is defined as:

$$K^{\theta_l,\theta_k}(x, x') = \exp\left(-\frac{\sin^2((x_l - x'_l)\theta_l)}{\theta_k}\right) \tag{18}$$

  where θ_l controls the periodicity and θ_k determines the smoothness of the periodic oscillations.

In order to further expand the range of functions that can be modeled with GP, it is essential to leverage the property of kernel composition. This property states that the sum and product of PSD kernels are also valid kernels. Adding kernels models the superposition of independent functions, while multiplying them accounts for interactions between different phenomena [18]. However, SMT previously did not offer a native mechanism for combining kernels, so we introduce a novel architecture that supports flexible kernel composition.

The optimal structure for managing a wide variety of kernel functions is a class-based architecture. To address the need for flexibility and ease of use, we have defined a new class structure within SMT that can accommodate any kernel form. The architecture, illustrated in Fig. 1, is composed of three main elements:

1. **Base Kernel Class**: An abstract class that provides a template for all kernel implementations. Any new kernel can inherit from this class to ensure consistency and compatibility.

2. **Standard Kernel Implementations**: Existing kernels in SMT have been refactored into class-based structures. This includes not only the traditional kernels (such as the squared exponential or Matern kernels) but also the newly added rational quadratic and periodic kernels.

3. **Kernel Composition Utilities**: To simplify the process of creating complex kernels, we provide built-in tools that allow users to combine kernels using basic Python operators (e.g., addition and multiplication) as the product and sum of SDP kernels still results in a kernel [19]. These utilities enable the effortless formation of composite kernels, which can model both the independent superposition and interactions between different data features.

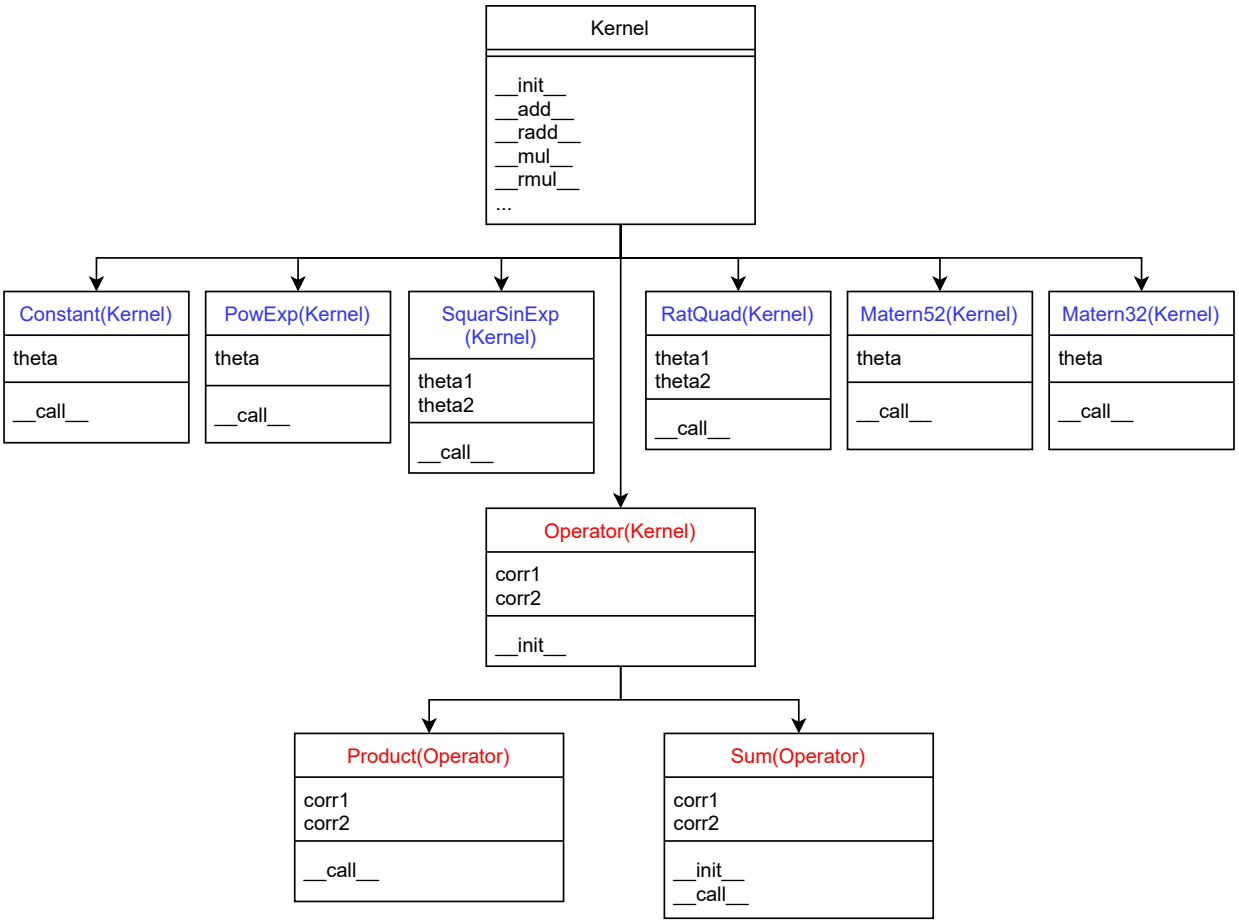


Figure 1: Class diagram for kernels in SMT. In black, the main class, in blue the basic kernels already implemented, and in red the operator.

The new architecture empowers users to construct expressive, hybrid kernels tailored to specific applications. For example, in modeling the oscillatory behavior of an aircraft wing subject to aerodynamic forces, one might combine a periodic kernel to capture cyclic oscillations with a squared exponential kernel to model broader trends. By integrating these advancements, SMT 2.0 significantly extends the flexibility and power of Gaussian Process modeling. This novel kernel library not only supports standard kernel configurations but also facilitates the development of custom, composite models tailored to the unique challenges of complex, frequency-sensitive applications as shown in the following section.

4 RESULTS

To test our framework, we evaluated its ability to model a range of functions and datasets, and assess its capacity to uncover underlying structures and extrapolate beyond the given data. Below, we provide results from different case studies, each with its own characteristics and challenges.

Cardinal Sine function We begin by testing our framework with the well-known cardinal sine (sinc) function. This function, although not strictly periodic, exhibits a periodic component with a varying amplitude. The sinc function is a good candidate for testing, as it has a periodicity that is modified by amplitude fluctuations, which we aim to capture with our kernel model. We employed a kernel formed by the multiplication of a squared exponential (SE) kernel and a periodic kernel:

$$K(x, x') = K_{SE}^{\theta_1}(x, x') \cdot K_{Periodic}^{\theta_2, \theta_3}(x, x') \tag{19}$$

This kernel combination allows us to model both the periodic nature of the function and the amplitude variations. As shown in Fig. 2, the model perfectly captures the periodic structure and the largest amplitude, providing a good representation of the sinc function over the given input domain.

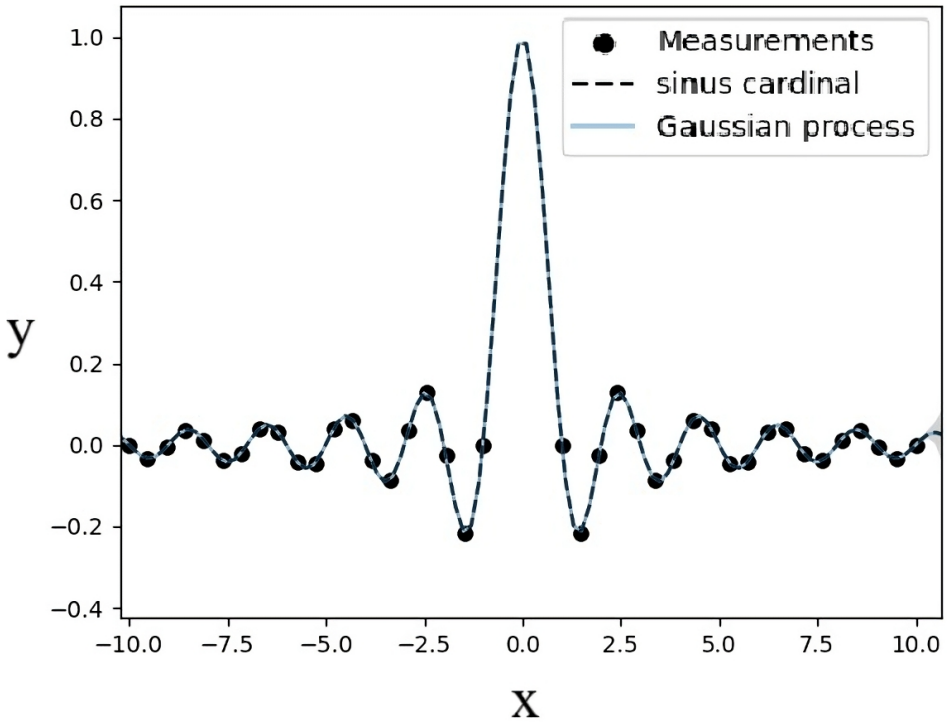


Figure 2: Posterior GP with kernel of Eq.(19) on the sinc function

Mauna-Loa CO_2 observations Next, we apply our framework to the Mauna-Loa CO_2 dataset [20], which contains historical data on the concentration of atmospheric CO_2 since 1958. The Mauna-Loa dataset is often used to test models for environmental data, as it contains a combination of long-term trends, seasonal oscillations, and noise. The primary goal of this case study is to test our model's ability to capture these components, particularly the periodic oscillations and the irregularities introduced by noise.

To model this dataset, we employed a kernel that combines multiple kernel functions to address the different behaviors in the data. Specifically, we use a periodic kernel to model the oscillations, a squared exponential kernel to model the small variations around the periodicity, another squared exponential kernel to account for the linear growth trend,

and a rational quadratic kernel to handle irregularities:

$$K(x,x') = K_{SE}^{\theta_1}(x,x') \cdot K_{Periodic}^{\theta_2,\theta_3}(x,x') + K_{SE}^{\theta_4}(x,x') + K_{rational_quadratic}^{\theta_5,\theta_6}(x,x') \tag{20}$$

The results in Fig. 3 show that the model effectively captures the periodicity and growth, with the interpolation fitting the observed data closely. Additionally, the extrapolation provides reasonable forecasts, with the identified components appearing to repeat in the future, as expected from the underlying structure of the data.

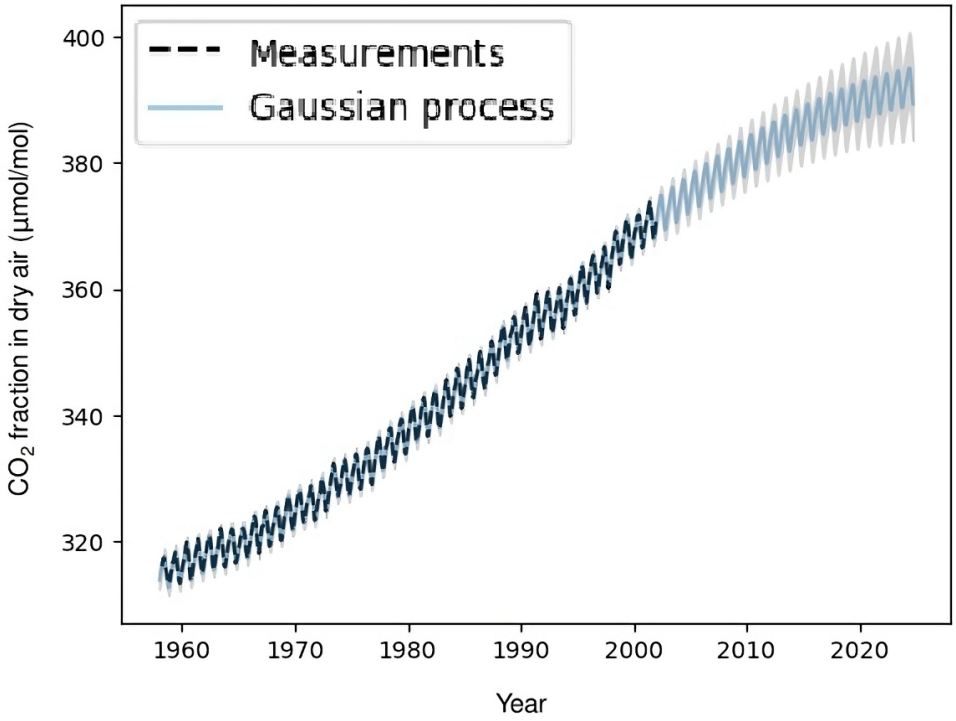


Figure 3: CO_2 forecasting with GP kernel of Eq.(20).

Airlines Annual Passengers Traffic Another case we examine is a dataset containing monthly international airline passenger counts from January 1949 to December 1960 [21]. This dataset is commonly used for testing forecasting models due to its long-term trend, seasonal fluctuations, and sudden changes in direction. Similar to the Mauna-Loa dataset, this data exhibits three primary components: a long-term trend, periodic oscillations, and sharp shifts in direction, such as a sudden increase in the number of passengers during specific periods.

To model this dataset, we apply the same kernel we used for the Mauna-Loa data, given its similar structure:

$$K(x,x') = K_{SE}^{\theta_1}(x,x') \cdot K_{Periodic}^{\theta_2,\theta_3}(x,x') + K_{SE}^{\theta_4}(x,x') + K_{rational_quadratic}^{\theta_5,\theta_6}(x,x') \tag{21}$$

The results shown in Fig. 4 indicate that the kernel successfully captures the underlying structure of the data. The interpolation is smooth and accurate, and the extrapolated

part is plausible, although the variance is higher, possibly reflecting the uncertainty in predicting sudden changes in the future. This suggests that while the model captures the general trends, future data points may have a high degree of uncertainty.

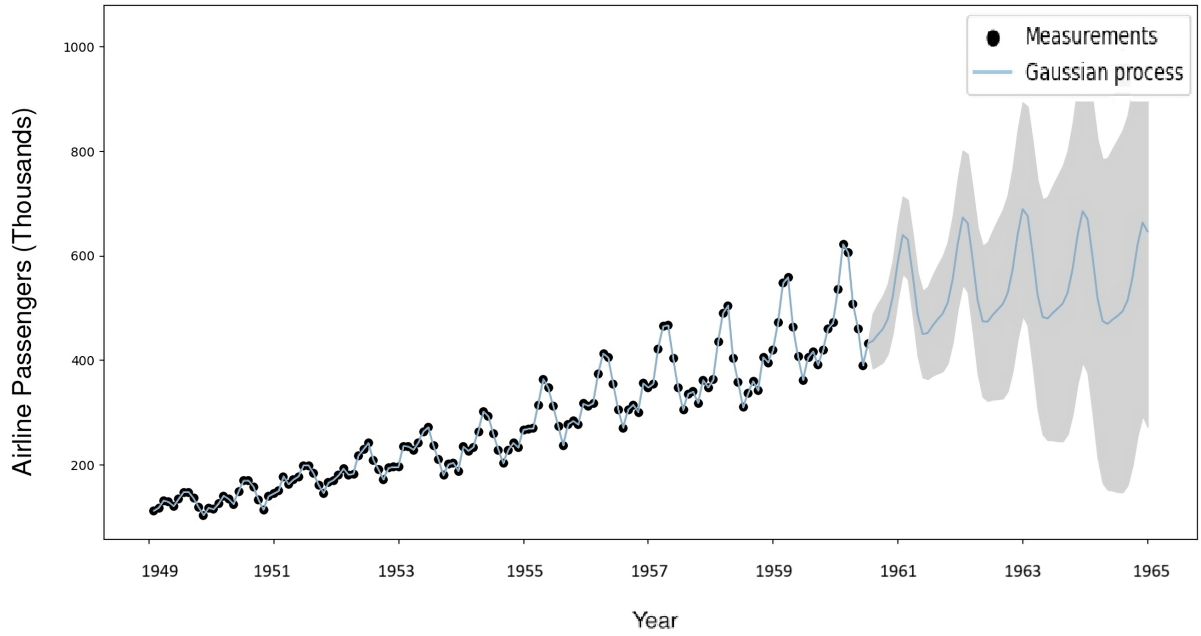


Figure 4: GP approximation with a kernel given by Eq.(21) on passengers database

5 CONCLUSION

In this paper we presented a functional framework to create and combine new kernels for GP in the python library SMT. This expands the scope of functions that can be modeled with this method and also provides a way to use Kriging for prediction and extrapolation with the use of the periodic kernel, which is usually not possible. However, there are still some gray areas. For complex data, finding the right kernel can be a complex endeavor and we can note that all the tests done here are in 1 dimension only, as the complexity increases a lot for complex kernel with the dimension. Indeed the number of hyperparameters in the kernel is multiplied by the dimension making the duration of modelisation quickly prohibitive.

The development of this user-defined kernel library opens several avenues for future work. First this work was limited to sum and product of kernels and can be extended to any PSD matrix manipulation such as ANalysis Of VAriance (ANOVA) [19]. Also, it was limited to spatial classical kernels whereas some more complex ones are better tailored for timeseries forecasting such as the Mercer or spectral kernels [22, 23]. Future updates may include automated techniques for selecting the most appropriate kernel configuration based on data characteristics, utilizing methods such as Bayesian model selection or evolutionary algorithms. Enhancing kernel hyperparameter tuning through gradient-based methods and automatic differentiation will further improve computational efficiency, especially for large-scale problems. The new framework can be directly applied to complex aeroelastic analyses of high aspect ratio wings, such as HALE (High Altitude, Long Endurance) solar aircraft subjected to gust perturbations.

ACKNOWLEDGEMENTS

This work is part of the activities of ONERA - ISAE - ENAC joint research group. The research presented in this paper has been performed within the framework of the COLOSSUS project (Collaborative System of Systems xEploration of Aviation Products, Services and Business Models) and has received funding from the European Union Horizon Programme under grant agreement n°101097120. The authors acknowledge the research project MIMICO funded in France by the Agence Nationale de la Recherche (ANR, French National Research Agency), grant number ANR-24-CE23-0380.

REFERENCES

[1] A. Chan, A. F. Pires, and T. Polacsek. Trying to elicit and assign goals to the right actors. In *Conceptual Modeling: 41st International Conference, ER 2022*, 2022.

[2] N. Bartoli, T. Lefebvre, S. Dubreuil, M. Panzeri, R. D'Ippolito, K. Anisomov, and A. Savelyev. Robust nacelle optimization design investigated in the agile european project. In *19th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2018. doi:10.2514/6.2018-3250.

[3] A. Lambe and J. R. R. A. Martins. A unified description of mdo architectures. In *9th World Congress on Structural and Multidisciplinary Optimization*, 2011.

[4] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59:1075–1104, 2019. doi:10.1007/s00158-019-02211-z.

[5] P. Saves, Y. Diouane, N. Bartoli, T. Lefebvre, and J. Morlier. A general square exponential kernel to handle mixed-categorical variables for gaussian process. In *AIAA Scitech 2022 Forum*, 2022.

[6] A. F. Hernandez and M. A. Grover. Stochastic dynamic predictions using gaussian process models for nanoparticle synthesis. *Computers & Chemical Engineering*, 34 (12):2042–2052, 2010.

[7] N. R. Costa and J. Lourenço. Gaussian process model – an exploratory study in the response surface methodology. *Quality and Reliability Engineering International*, 32 (3):935–944, 2016.

[8] P. Saves, R. Lafage, N. Bartoli, Y. Diouane, J. H. Bussemaker, T. Lefebvre, J. T. Hwang, J. Morlier, and J. R. R. A. Martins. SMT 2.0: A surrogate modeling toolbox with a focus on hierarchical and mixed variables Gaussian processes. *Advances in Engineering Sofware*, 188:1–15, 2024.

[9] P. Saves, Y. Diouane, N. Bartoli, T. Lefebvre, and J. Morlier. A mixed-categorical correlation kernel for gaussian process. *Neurocomputing*, 532:556–568, 2023.

[10] R. Priem, Y. Diouane, N. Bartoli, S. Dubreuil, and P. Saves. High-dimensional bayesian optimization using both random and supervised embeddings. *AIAA Journal*, 63(1):162–173, 2025.

[11] D. G. Krige. A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139, 1951.

[12] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006.

[13] H. Lee. *Gaussian Processes*, chapter 5, pages 575–577. Springer Berlin Heidelberg, 2011.

[14] R. Conde Arenzana, A. López-Lopera, S. Mouton, N. Bartoli, and T. Lefebvre. Multi-fidelity Gaussian process model for CFD and wind tunnel data fusion. In *ECCOMAS Aerobest*, 2021.

[15] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. R. A. Martins. A python surrogate modeling framework with derivatives. *Advances in Engineering Software*, 135, 2019. doi:10.1016/j.advengsoft.2019.03.005.

[16] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

[17] R. Rossi. *Mathematical statistics: an introduction to likelihood based inference*. John Wiley & Sons, 2018.

[18] D. Duvenaud, J. Lloyd, R. Grosse, J. Tenenbaum, and G. Zoubin. Structure discovery in nonparametric regression through compositional kernel search. In *International Conference on Machine Learning*, pages 1166–1174. PMLR, 2013.

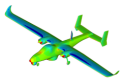
[19] O. Roustant, E. Padonou, Y. Deville, A. Clément, G. Perrin, J. Giorla, and H. Wynn. Group kernels for gaussian process metamodels with categorical inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):775–806, 2020.

[20] C. D. Keeling. The influence of mauna loa observatory on the development of atmospheric co2 research. *In Mauna Loa Observatory. A 20th Anniversary Report.(National Oceanic and Atmospheric Administration Special Report*, pages 36–54, 1978.

[21] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.

[22] A. Wilson and R. Adams. Gaussian process kernels for pattern discovery and extrapolation. *International Conference on Machine Learning*, 2013.

[23] A. Gorodetsky and Y. Marzouk. Mercer kernels and integrated variance experimental design: connections between gaussian process regression and polynomial approximation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):796–828, 2016.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AN MDAO FRAMEWORK APPROACH EXTENDED TO UNSTEADY MULTIPHYSICS: KEY INGREDIENTS AND EXAMPLES

Florian Roß, Simon Ehrmanntraut, Constantin Höing, Sebastian Gottfried and Arthur Stück

German Aerospace Center (DLR)
Institute of Software Methods for Product Virtualization
Nöthnitzer Straße 46b, 01187 Dresden
florian.ross@dlr.de, simon.ehrmanntraut@dlr.de,
constantin.hoeing@dlr.de,sebastian.gottfried@dlr.de, arthur.stueck@dlr.de
https://www.dlr.de/en/sp

Abstract. *State-of-the-art MDAO frameworks like* OpenMDAO *and* GEMSEO *facilitate the setup and the gradient-enabled solution/optimization of large multidisciplinary problems, which are predominantely of steady-state type. Time-dependent MDAO problems, however, pose additional challenges, particularly in conjunction with cost-intensive high-fidelity simulations. This work demonstrates how advanced MDAO-framework capabilities can be transferred from the steady-state to the unsteady domain. To this end, an MDAO framework approach is presented integrating* OpenMDAO *with the HPC ecosystem* FlowSimulator *including high-fidelity simulation plugins like the CFD-software* CODA*. Among the key capabilities of the time-accurate framework approach are a complete system representation on the framework level resolving all input- and output dependencies to yield a monolithic time-resolved view, and checkpointing-enabled reverse gradient accumulation over time. A gradient-based optimization was conducted in the framework for an elementary Van der Pol oscillator case carrying out an unsteady reverse-in-time sensitivity analysis. Moreover, it was successfully shown and verified that the framework is capable of performing unsteady adjoint viscous/turbulent CFD computations for a NACA0012 and the CRM/DPW-5 wing-body configuration, paving the way for large adjoint-enabled unsteady MDAO applications.*

Keywords: Unsteady MDAO, MDAO Frameworks, Sensitivity Analysis, Checkpointing

1 INTRODUCTION

The past decade has seen the development of multidisciplinary design analysis and optimization (MDAO) frameworks, like e.g. the open-source frameworks *OpenMDAO* [1] or *GEMSEO* [2]. Such MDAO frameworks offer features like forward- and reverse gradient accumulation over multiple disciplines using nested multidisciplinary solvers and preconditioners, MPI-parallelism, and reusability of integrated components. These features make the frameworks attractive to the community in academia and industry for use in different levels of fidelity. In steady-state analyses, this has caused MDAO to already have become state of the art. To model time-resolved problems in MDAO frameworks, different kind of approaches were tried. One way is to implement the whole time-resolved problem as a pseudo-steady system [3], and/or by vectorizing the considered ordinary differential equation system, as done by the open-source multidisciplinary optimal control software *DYMOS* [4]. However, these approaches yield a system wide state that contains the information of all simulated quantities of all time steps at once, which quickly becomes prohibitive for problems of high spatial resolution and many physical time-steps as it is the case in conjunction with time-accurate high-fidelity CFD simulations.

To remedy this problem, the open-source extension *RKOpenMDAO* [5] was developed which extends *OpenMDAO* towards time integration via Runge–Kutta methods [6]. This extension performs the time integration for a considered ODE by modeling a single time stage of the ODE as an inner problem to avoid the allocation the system state over all time steps and stages. Currently, explicit and diagonally-implicit Runge–Kutta schemes are supported.

This paper investigates the sensitivity analysis of time-resolved MDAO problems with *RKOpenMDAO*. For this, The second section considers both theoretical and implementation aspects of passing derivatives through time using *RKOpenMDAO*. The third section contains numerical experiments conducted with *RKOpenMDAO* using the mentioned capabilities. This includes both an academic mathematical optimization example and high-fidelity CFD analyses of passing derivatives through time using the CFD software *CODA* [7], which is developed in collaboration between the French Aerospace Lab *ONERA*, *Airbus*, the German Aerospace Center (*DLR*) and their research partners in Europe.

2 SENSITIVITY ANALYSIS IN *RKOPENMDAO*

For time integration in *OpenMDAO* [1], the extension *RKOpenMDAO* was used. As explained in [5], this extension allows Runge–Kutta methods to be applied on *OpenMDAO* problems such that the capabilities of the framework can be extended to the unsteady domain.

Consider a multidisciplinary first-order ODE system in residual form:

$$0 = F(t, x(t), x'(t)), \quad x(t_0) = x_0 \ , \tag{1}$$

where $x = (x^{\text{Disc } 1}, x^{\text{Disc } 2}, \ldots, x^{\text{Disc } m})$ is the multidisciplinary state vector, x' its time derivative, t is time, and F is a function modeling the (possibly implicit) interaction between t, x and x'. The formulation in Eq.(1) also allows the inclusion of functionals and time independent inputs into the ODE system by using transformations similar to the ones in [8]. For independent inputs, assume a system of the form:

$$0 = G(t, x(t), x'(t), a), \quad x(t_0) = x_0 \ ,$$

where a is a vector of independent inputs. Then by adding the ODE $\tilde{a}'(t) = 0$ with initial value $\tilde{a}(t_0) = a$, and defining $\tilde{x} = (x, \tilde{a})$, the system can be brought into the form:

$$\begin{aligned} 0 &= \begin{pmatrix} G(t, x(t), x'(t), \tilde{a}(t)) \\ \tilde{a}'(t) \end{pmatrix} \\ &= \tilde{G}(t, \tilde{x}(t), \tilde{x}'(t)) \\ \tilde{x}(t_0) &= \begin{pmatrix} x(t_0), \\ \tilde{a}'(t_0) \end{pmatrix} = \begin{pmatrix} x_0, \\ a \end{pmatrix} \ , \end{aligned}$$

which fits into the form of Eq.(1).

Similarly, assume there is a functional:

$$J = \int_{t_0}^{t_f} j(x(t))dt$$

for the ODE system:

$$0 = H(t, x(t), x'(t)), \quad x(t_0) = x_0 \ ,$$

where t_f is the end time of the ODE system. Then by introducing the ODE $\hat{J}'(t) = j(x(t))$ with initial values $\hat{J}(t_0) = 0$, and defining $\hat{x} = (x, \hat{J})$, there is a system of form Eq.(1):

$$\begin{aligned} 0 &= \begin{pmatrix} H(t, x(t), x'(t)) \\ \hat{J}'(t) - j(x(t)) \end{pmatrix} \\ &= \hat{H}(t, \hat{x}(t), \hat{x}'(t)) \\ \hat{x}(t_0) &= \begin{pmatrix} x(t_0), \\ \hat{J}'(t_0) \end{pmatrix} = \begin{pmatrix} x_0, \\ 0 \end{pmatrix} \ , \end{aligned}$$

in which $J \equiv \hat{J}(t_f)$ by the fundamental theorem of analysis.

Applying a diagonally implicit Runge–Kutta (DIRK) scheme on Eq.(1) yields:

$$\begin{aligned} x_{n+1} &= x_n + \Delta t \sum_{i=1}^{S} b_i k_i \quad n = 0, \dots, N-1 \\ 0 &= F\left(t_n + \Delta t c_i, x_n + \Delta t (s_i + a_{ii} k_i), k_i\right) \quad i = 1, \dots, s \ , \end{aligned} \tag{2}$$

where x_n and x_{n+1} are the states of x at the current and next time step respectively, k_i is the stage-wise update, Δt is the step size, t_n is the time of the current step, and a_{ii}, b_i, c_i are parameters from the Butcher tableau of the used Runge–Kutta scheme. However, this form is inconvenient for the use in MDAO frameworks like *OpenMDAO*, as the second equation of Eq.(2) has a changing number of variables depending on the current stage of the Runge–Kutta scheme. By substituting $s_i = \sum_{j=1}^{i-1} a_{ij} k_j$, the following system is obtained:

$$\begin{aligned} x_{n+1} &= x_n + \Delta t \sum_{i=1}^{S} b_i k_i \quad n = 0, \dots, N-1 \\ s_i &= \sum_{j=1}^{i-1} a_{ij} k_j \quad i = 1, \dots, s \\ 0 &= F\left(t_n + \Delta t c_i, x_n + \Delta t (s_i + a_{ii} k_i), k_i\right) \quad i = 1, \dots, s \ . \end{aligned} \tag{3}$$

With this substitution, the last equation of Eq.(3) has the same number of variables independent of the time stage which is a requirement for an implementation of that equation system into an *OpenMDAO* model, while the first two equations can be computed once per stage or step, respectively.

For the sensitivity analysis, the matrix vector products $\frac{dx_N}{dx_0}dx_0$ and $\frac{dx_N}{dx_0}^T\overline{x}_N$ (for forward or reverse mode respectively) are of interest. Due to a single time step only depending on the previous step and its own intermediate data, these matrix vector products can be expanded into:

$$\frac{dx_N}{dx_0}dx_0 = \frac{dx_N}{dx_{N-1}}\frac{dx_{N-1}}{dx_{N-2}}\dots\frac{dx_1}{dx_0}dx_0\,, \tag{4}$$

respectively

$$\frac{dx_N}{dx_0}^T\overline{x}_N = \frac{dx_1}{dx_0}^T\frac{dx_2}{dx_1}^T\dots\frac{dx_N}{dx_{N-1}}^T\overline{x}_N\,. \tag{5}$$

This means that the sensitivities can easily be calculated, given that Eq.(3) can be differentiated in forward and reverse mode. Since the calculations of x_{n+1} and s_i are explicit and linear calculations, their differentiation can be done on the framework level without many problems. The main work resides in the differentiation of F wrt. x_n, s_i and k_i. Given that F with its derivatives is implemented on the component level via an *OpenMDAO* model, the framework time integrator then can provide forward sensitivities without further complications. The calculations of the necessary matrix vector products can take place right after the calculation of the stage or step in the non-differentiated time integration, since at that point the necessary linearization points are available without additional costs.

In reverse, the calculation of the involved derivatives have to be backward propagated through time both in time stages and time steps. On the level of a single time step, it is feasible to keep intermediate values of the stages in memory, since Runge–Kutta schemes used in practice usually have only single-digit numbers of stages. Using these intermediate values as linearization points for the backwards pass, the reverse derivative of a single step can be obtained. Storing the data for each time step, especially for high-fidelity simulations with thousands of time steps, would result in prohibitive memory requirements, defeating the purpose of *RKOpenMDAO*. Instead, *RKOpenMDAO* employs checkpointing and stores in the forward pass of an non-differentiated time integration a limited number of intermediate values of the time steps. These results act as a starting point for the backwards pass for the reverse sensitivity analysis. Instead of recalculating the whole time trajectory, a restart is done from the latest checkpoint, reducing the computational effort to get to the step whose reverse derivative is to be computed. In the implementation of *RKOpenMDAO*, binomial checkpointing with the algorithm *revolve* [9] is used, via the Python wrapper *PyRevolve* [10] on the level of the multidisciplinary system. Binomial checkpointing is proven to be optimal in the sense that it minimizes the number of recomputed steps wrt. to the number of allowed checkpoints [9].

3 NUMERICAL EXPERIMENTS

3.1 Van der Pol Oscillator – an Unsteady Optimization Problem

The Van der Pol Oscillator [11] was selected for validation and verification purposes, since it is an well known example for stiff ODEs. It is a second-order ODE defined by

$$\ddot{x} - \epsilon(1 - x^2)\dot{x} + x = 0,$$

with initial values $x(0) = x_0$, $\dot{x}(0) = y_0$ (see Fig.1 for an example trajectory). By introducing $y = \dot{x}$, this can be transformed into a first order ODE system

$$\begin{aligned} \dot{y} - \epsilon(1 - x^2)y + x &= 0, \\ \dot{x} - y &= 0, \end{aligned}$$

with initial values $x(0) = x_0$, $y(0) = y_0$. Note that for $\epsilon = 0$, the ODE system descibes a

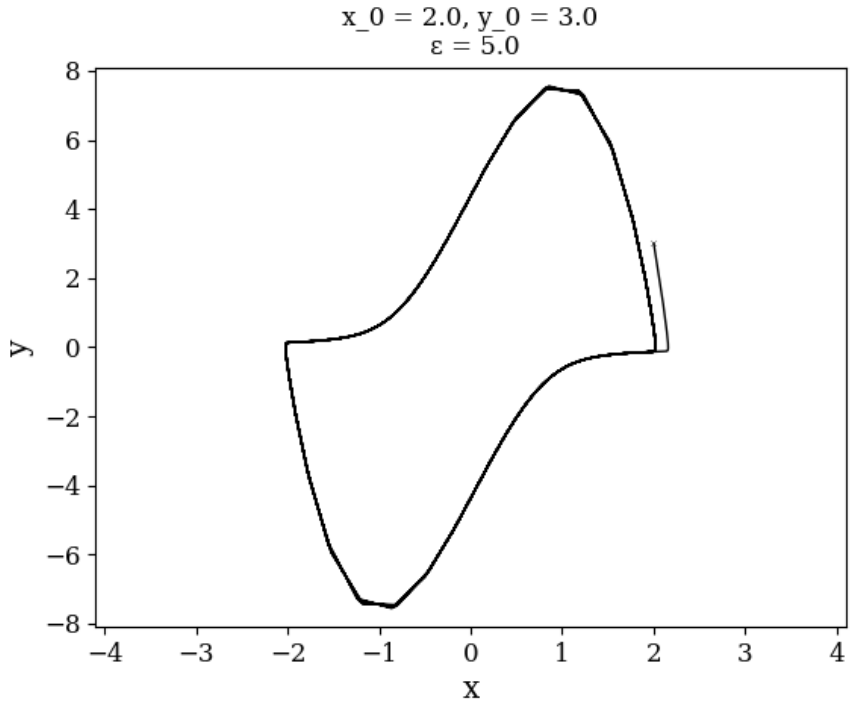


Figure 1: Example trajectory of the Van der Pol oscillator.

harmonic oscillator.

For verification purposes, the problem was artificially turned into an optimization problem, in which a gradient-based optimization based on the reverse-in-time sensitivity analysis is performed. Along these lines, the following functional is defined:

$$J = \int_0^{t_f} (x(t)^2 + y(t)^2 - 1) \cdot \tanh(x(t)^2 + y(t)^2 - 1)dt. \tag{6}$$

This functional consists of two parts:

- The parametrization of the unit circle $x^2 + y^2 - 1 = 0$ for $x, y \in \mathbb{R}$

- An differentiable approximation of the absolute value $|x| \approx x \cdot \tanh x$

Thus, the functional can be interpreted as measuring the distance of the curve of the Van der Pol oscillator from the unit circle over time.

As a first verification case for the sensitivity-analysis capabilities of *RKOpenMDAO*, the following optimization problem is considered:

$$\begin{aligned} \text{Minimize} \quad & J \\ \text{w.r.t.} \quad & \epsilon, x_0, y_0 \\ \text{subject to} \quad & \dot{y} - \epsilon(1 - x^2)y + x = 0 \\ & \dot{x} - y = 0 \\ & x(0) = x_0 \\ & y(0) = y_0 \end{aligned}$$

The functional is minimized w.r.t. the parameter epsilon and the initial values of the ODE. Analytically, this it optimal if and only if $\epsilon = 0$ and x_0, y_0 lie on the unit circle. In that case, $J = 0$ holds:

1. The solution will be of form $x(t) = x_0 \cos(t) + y_0 \sin(t)$, $y(t) = -x_0 \sin(t) + y_0 \cos(t)$, as for $\epsilon = 0$ the ODE system describes a harmonic oscillator.

2. It follows that $x(t)^2 + y(t)^2 = x_0^2 + y_0^2 \equiv 1$, the first equality holding due to the Pythagorean trigonometric identity, the second since x_0, y_0 reside on the unit circle.

3. It follows that $(x(t)^2 + y(t)^2 - 1) \cdot \tanh(x(t)^2 + y(t)^2 - 1) \equiv 0$, and thus $J = 0$.

As $(x^2 + y^2 - 1) \cdot \tanh(x^2 + y^2 - 1) \geq 0$ holds for all $x, y \in \mathbb{R}$, this extends to $J \geq 0$, making a configuration with $J = 0$ optimal. Note that the minimum being $J = 0$ means that this problem also could have been solved using a root-finding algorithm (e.g. the Newton method). However, the choice to solve the problem using optimization was done for verification purposes.

Numerically, the ODE involved in the problem is discretized using a third-order four-stage ESDIRK scheme from [12] with $\gamma = 0.4358665$, using *RKOpenMDAO* for time integration with a step size of $\Delta t = 0.05$, starting at $t_0 = 0$ until $t_f = 10$. The involved optimization is done via *OpenMDAOs* `SciPyOptimizeDriver` using the SLSQP method with a tolerance of 10^{-6}. As initial values for the design variables, $\epsilon = 0.3$ and $x_0 = y_0 = 1.5$ were chosen. Under these circumstances, a minimum was found after 15 iterations, using 32 function evaluations. See Fig.2 for a visualization of the trajectories. The minimal functional value obtained is $J = 1.23638251 \cdot 10^{-7}$, with the final values of the design variables being $x_0 = -0.32984621$, $y_0 = -0.9441123$ and $\epsilon = -4.11990118 \cdot 10^{-5}$. Taking into account numerical errors and tolerances, these results reflect the analytical minimum, as $x_0^2 + y_0^2 = 1.00008338$ being close to 1, and ϵ and J being close to 0.

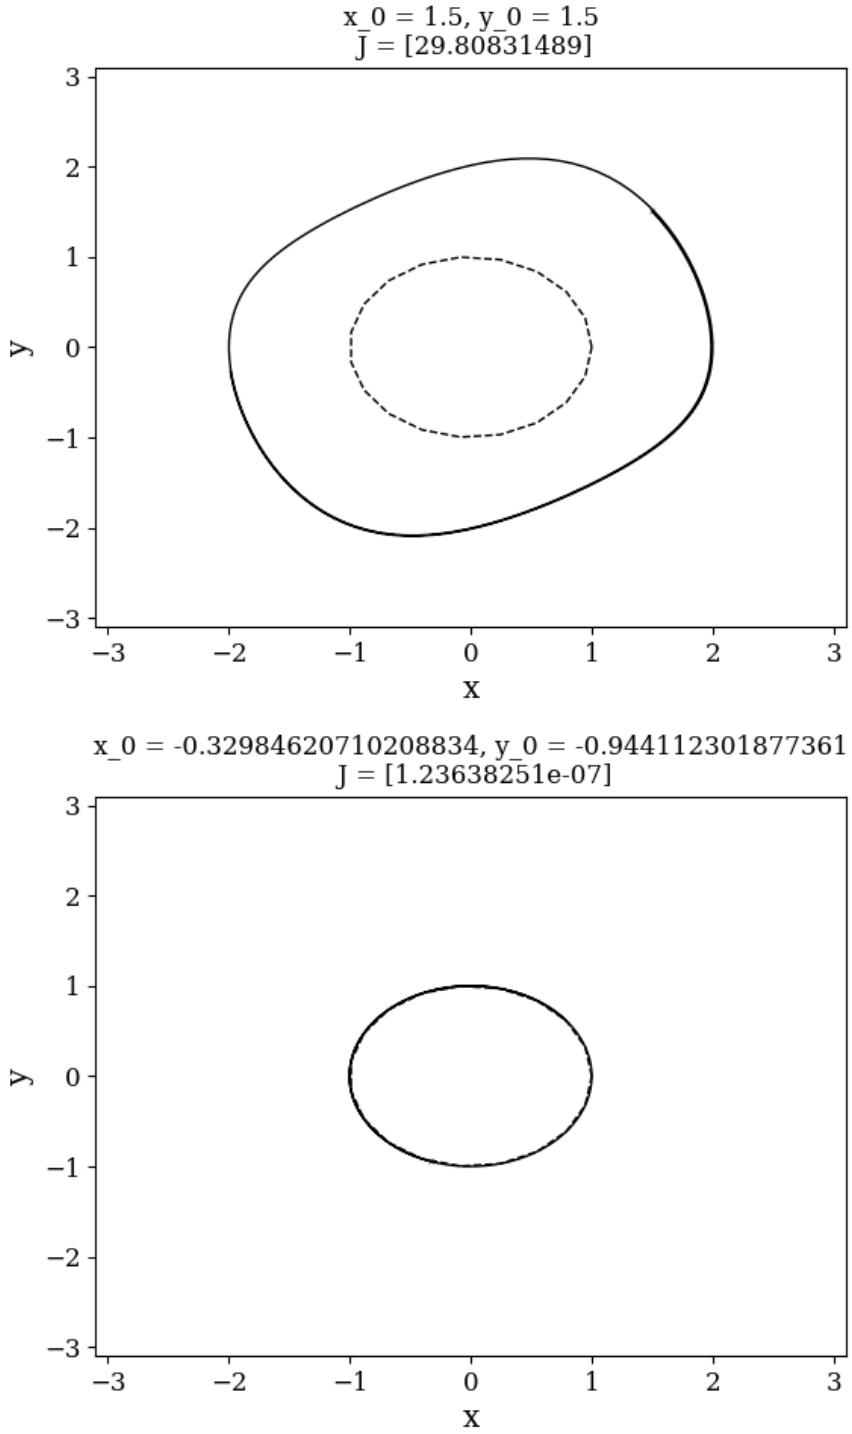


Figure 2: Initial (up) and final (down) run of the time integration during the optimization of the Van der Pol oscillator. The dotted circle marks the optimal solution, cf. bottom.

3.2 Sensitivity Analysis Reverse-in-Time with the CFD Software CODA

To challenge the presented MDAO framework approach, high-fidelity unsteady CFD cases are considered in conjunction with the CFD software *CODA*. Independent of the exact configuration, the following equations need to be iterated in time to obtain the time trajectory of the flow state W:

$$\frac{\partial W}{\partial t} + R(W) = 0, \tag{7}$$

with $R(W)$ denoting the stationary residual contributions from the spatial discretization.

In the first example, a NACA0012 airfoil is considered in laminar Navier–Stokes flow conditions. A second-order finite volume approach is used to discretize the equations in space. A source of momentum Q is introduced as a control term, yielding a slightly modified set of equations:

$$\frac{\partial W}{\partial t} + R(W) + Q = 0, \tag{8}$$

Q is a (constant in time, non-dimensional) source of momentum, which was introduced as a simple control term in this study. In the subsequent unsteady adjoint sensitivity analysis, the unsteady residual equation had to be differentiated with respect to that source term. The momentum in the direction of the drag, being the time integral of the drag coefficient was considered as objective functional:

$$J = \int_0^{t_f} C_D(W(t))dt$$

where C_D is the drag coefficient of the airfoil. After time discretization, a time step of the problem reads:

$$\begin{aligned}
W_{n+1} &= W_n + \Delta t \sum_{i=1}^{s} b_i k_i^W \\
I_{n+1}^D &= I_n^D + \Delta t \sum_{i=1}^{s} b_i k_i^J \\
s_i^W &= \sum_{j=1}^{i-1} a_{ij} k_j^W \\
0 &= \frac{W_n^i - W_n - \Delta t s_i^W}{\Delta t a_{ii}} + R(W_n^i) + Q \\
0 &= k_i^W - \frac{W_n^i - W_n - \Delta t s_i^W}{\Delta t a_{ii}} \\
k_i^J &= C_D(W_n^i)
\end{aligned}$$

with $i = 1, ..., s$.

The following simulations use a Mach number $Ma = 0.1$ and a Reynolds number $Re = 10000$. The third-order four-stage ESDIRK method given in [13] was used with a non-dimensional step size $\Delta t = 0.1$, while the simulated time corresponds to $0.04 \times T_{\text{conv}}$

with $T_{\text{conv}} = s_{\text{chord}}/U_\infty$. Initial values W_0 are the corresponding steady state with the Mach and Reynolds number at an angle of attack $\alpha = 0°$. 5 time steps with $Q = 0$ result in $J = 0.01602079$. The sensitivities derivatives $\overline{Q}^T = dJ/dQ$, computed a reverse pass through time, indicate where to most effectively add or remove momentum in order to change J, as can be seen in Fig.3. For example, with the leading edge of the airfoil at $(0, 0)$ and the trailing edge at $(1, 0)$, setting the horizontal part of Q to -0.1 at around $(0.37, 0.47)$ increases the functional to $J = 0.01804821$, whereas setting it to 0.1 decreases it to $J = 0.01401113$.

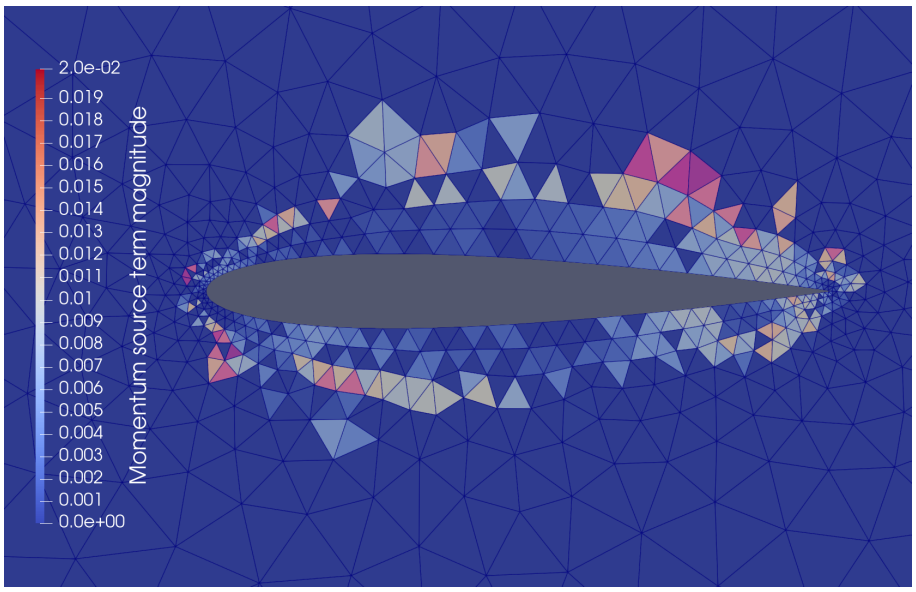


Figure 3: Euclidean norm of the sensitivity vector distribution $\overline{Q}^T = dJ/dQ$ for the NACA0012 airfoil.

Subsequently, the Common Research Model CRM/DPW5 [14] depicted in Fig.4 was considered in (U)RANS flow using the negative extension of the Spalart–Allmaras turbulence model (SA-neg). The governing equations Eq.(7) were solved with a standard second-order finite-volume discretization using a Roe upwind scheme for the convective fluxes on a computational mesh of approx. 2.8 million control volumes. The objective functional is defined to be the momentum,

$$J = \int_0^{t_f} C_{F_z}(W(t))dt\,,$$

wherein C_{F_z} is the non-dimensional coefficient of the aerodynamic force acting on the aircraft in vertical direction. The time integration is performed by means of the multidisciplinary implicit Runge–Kutta integrator implemented in *RKOpenMDAO*. The time-

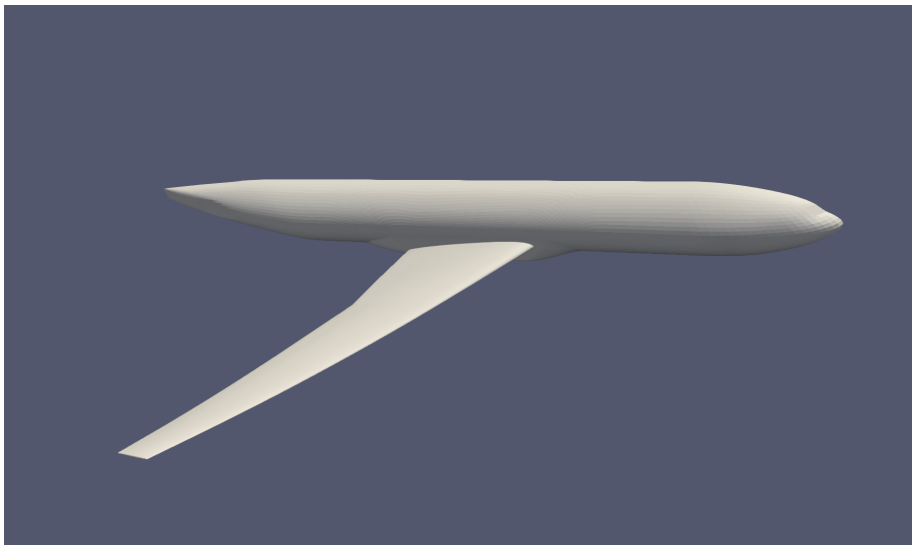


Figure 4: Surface plot of the CRM/DPW5.

stepping operations performed on the level of the framework are

$$
\begin{aligned}
W_{n+1} &= W_n + \Delta t \sum_{i=1}^{s} b_i k_i^W \\
I_{n+1}^z &= I_n^z + \Delta t \sum_{i=1}^{s} b_i k_i^J \\
s_i^W &= \sum_{j=1}^{i-1} a_{ij} k_j^W \\
0 &= \frac{W_n^i - W_n - \Delta t s_i^W}{\Delta t a_{ii}} + R(W_n^i) \\
0 &= k_i^W - \frac{W_n^i - W_n - \Delta t s_i^W}{\Delta t a_{ii}} \\
k_i^J &= C_{F_z}(W_n^i)\,,
\end{aligned}
$$

with $i = 1, ..., s$.

All simulations for the CRM/DPW-5 were carried out at a Mach number of $Ma = 0.85$ and a Reynolds number of $Re = 5 \cdot 10^6$. The unsteady simulation were started from free-stream initial conditions W_0 at an angle of attack of $\alpha = 0°$. The third-order four-stage ESDIRK method by [13] was used with a non-dimensional step size $\Delta t = 0.001$.

First, the derivatives $\frac{dJ}{dW_0}$ computed on-the-fly by forward-mode differentiation were verified against finite differences. In Fig.5, the analytical sensitivity $\frac{dJ}{dW_0} dW_0$ with $dW_0 = 1$ is compared against forward and central finite differences in a simulation with 5 time steps. In both cases, the approximation error was reduced down to ϵ of 10^{-11} for forward FD, and 10^{-9} for central FD, respectively. From that point on, the truncation error dominated the FD approximations.

Subsequently, a duality test of the objective functional derivative was performed to verify consistency of the sensitivity derivatives obtained from forward and reverse-mode time integration for arbitrary choices of the initial state perturbation dW_0 and the reverse-mode seeding $\overline{J} = 1$. The duality check was carried out for different numbers of time steps

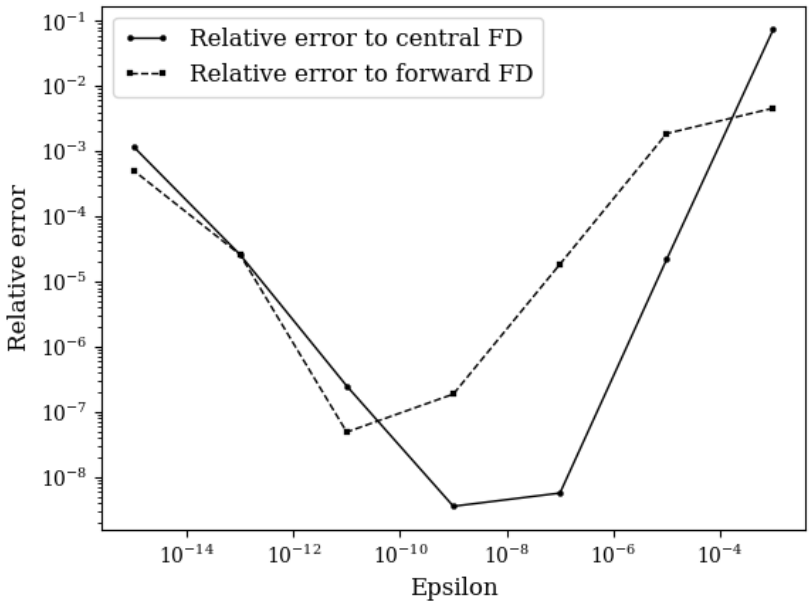


Figure 5: Comparison of forward-mode sensitivities against forward and central finite differences of a simulation with 5 time steps. As expected, central differences yield a lower error and a steeper convergence.

Table 1: Forward and reverse mode derivatives and their errors for different perturbations and numbers of time steps.

# time steps	Initial state perturbation dW_0	Absolute error
5	1	$1.31838984 \cdot 10^{-16}$
15	1	$9.29811783 \cdot 10^{-16}$
15	density field filled with ones, else zeros	$5.87460429 \cdot 10^{-7}$
15	vertical momentum field filled with ones, else zeros	$3.4399488 \cdot 10^{-7}$

N, with the absolute deviation between forward and reverse-mode sensitivity derivatives summarized in Tab.1.

In the next step, a first performance study was conducted for $N = 15$ and $dW_0 = 1$, were the performance of the time integration and its forward and reverse sensitivities were compared on 8, 16 and 32 nodes with 64 cores per node. For the reverse calculation, the *revolve* algorithm was allowed to use 5 checkpoints. According to Tab.2, a reverse pass takes almost twice the run-time of a forward-linear pass. This factor is only expected to increase once the number of time steps is increased without increasing the number of checkpoints accordingly, which is unfeasible in the long run due to memory constraints. However, note that the initialization flow state (considered as the control parameter here with respect to which the differentiation is performed) consists of more than $17 \cdot 10^6$ degrees of freedom. If complete gradients were to be computed instead of directional derivatives, a calculation of the gradient in forward mode would be much more costly than the reverse-mode version. In this case, the forward-mode version would take as many passes through

Table 2: Performance comparison of time integration and its derivatives. Here t_{Int} is the wall clock time necessary for time integration (including the collection of checkpoints), t_{fwd} is the time required for a forward pass of the sensitivities, and t_{rev} the time of the backwards pass of sensitivities. S_{t_k} is the speedup of t_k of the computation compared to the previous line (i.e. for doubling the number of cores).

# nodes (# cores)	t_{Int}	$S_{t_{Int}}$	t_{fwd}	$S_{t_{fwd}}$	t_{rev}	$S_{t_{rev}}$
8 (512)	1356.46	n.A.	3888.14	n.A.	5344.70	n.A.
16 (1024)	653.29	2.076	1432.13	2.715	2699.51	1.980
32 (2048)	428.05	1.526	821.27	1.744	1731.55	1.559

the time integrator as there are degrees of freedom, whereas the reverse-in-time mode only takes one per functional. For the case of 16 nodes, the influence of the number of checkpoints was investigated additionally. Fig.6 shows the reduction in runtime for different numbers of checkpoints. Note that the runtime reduction is offset compared to

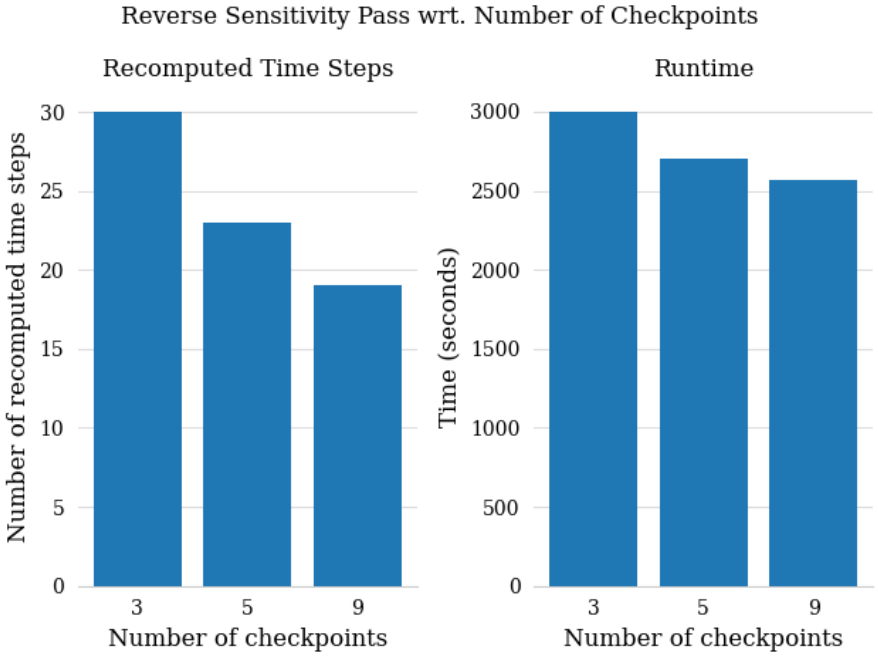


Figure 6: Proof of concept: influence of the number of checkpoints on the performance of the reverse-mode derivative calculation for a total of 15 physical time steps.

the reduction of recomputed time steps, since the runtime of the reverse pass is composed of both the recalculation of time steps and the adjoint sensitivity calculation per original time step. Changing the number of allowed checkpoints only has influence on the first of these two parts. It must be kept in mind that the options for checkpointing are very limited in this case, due to the very small number of 15 physical time steps performed. The situation is expected to look different for real-world unsteady simulation cases with many time steps due to long simulation times or small step sizes. However, Fig.6 indicates that the checkpointing approach implemented here, performs as expected.

4 CONCLUSIONS AND OUTLOOK

A gradient-enabled MDAO framework approach is presented for multidisciplinary, high-fidelity unsteady sensitivity analyses and optimizations. For reverse-in-time gradient computations, checkpointing schemes are provided so that memory can be traded for CPU runtime. First, the extension *RKOpenMDAO* was presented focusing on the ODE representation and key features for reverse mode computations including checkpointing. Numerical verifications – such as finite differences and duality checks – and experiments were carried out with the suggested MDAO framework approach for (a) an elementary coupled analytical test case and (b) the CRM/DPW-5 wing-body configuration being a large-scale CFD case. In all cases exact forward and reverse gradients were obtained for the unsteady problems, demonstrating the capabilities of the framework implementation. Mind that the verification problems considered here with framework approach were only of single-disciplinary type. However, mind that the unsteady extension *RKOpenMDAO* to *OpenMDAO* retains and supports the full multiphysics capabilities of the base-line version of *OpenMDAO* that is usually being used for steady-state MDAO problems.

In the future, the suggested time-accurate MDAO framework approach is to be applied to carry out gradient-based optimizations of large coupled configurations including aeroelastic MDAO of aircraft. Furthermore, the framework is to be extended to offer further time integration schemes and coupling methods. Moreover, it is intended to use *RKOpenMDAO* for high-fidelity unsteady design and control problems with adaptive time stepping and higher-order methods in time and space.

Acknowledgments The authors gratefully acknowledge the scientific support and HPC resources provided by the German Aerospace Center (DLR). The HPC system CARA is partially funded by "Saxon State Ministry for Economic Affairs, Labour and Transport" and "Federal Ministry for Economic Affairs and Climate Action".

REFERENCES

[1] J. Gray, J. Hwang, J. Martins, K. Moore, and B. Naylor. Openmdao: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 04 2019. doi:10.1007/s00158-019-02211-z.

[2] F. Gallard, C. Vanaret, D. Guenot, V. Gachelin, R. Lafage, B. Pauwels, P.-J. Barjhoux, and A. Gazaix. Gems: A python library for automation of multidisciplinary design optimization process generation. *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, 2018. doi:10.2514/6.2018-0657.

[3] J. Hwang and D. Munster. Solution of ordinary differential equations in gradient-based multidisciplinary design optimization. 01 2018. doi:10.2514/6.2018-1646.

[4] R. Falck, J. S. Gray, K. Ponnapalli, and T. Wright. dymos: A python package for optimal control of multidisciplinary systems. *Journal of Open Source Software*, 6 (59):2809, 2021. doi:10.21105/joss.02809. URL `https://doi.org/10.21105/joss.02809`.

[5] F. Roß, A. Büchner, S. Gottfried, and A. Stück. A high-fidelity framework approach enabling high-order implicit time stepping for mdao. *9th European*

Congress on Computational Methods in Applied Sciences and Engineering, 2024. doi:10.23967/eccomas.2024.196.

[6] J. Butcher. *Numerical Methods for Ordinary Differential Equations*. Wiley, 2016. ISBN 9781119121503.

[7] T. Leicht, D. Vollmer, J. Jägersküpper, A. Schwöppe, R. Hartmann, J. Fiedler, and T. Schlauch. Dlr-project digital-x: Next generation cfd solver 'flucs'. 01 2016.

[8] A. Sandu and L. Zhang. Discrete second order adjoints in atmospheric chemical transport modeling. *Journal of Computational Physics*, 227(12):5949–5983, 2008. ISSN 0021-9991. doi:https://doi.org/10.1016/j.jcp.2008.02.011. URL `https://www.sciencedirect.com/science/article/pii/S0021999108000983`.

[9] A. Griewank and A. Walther. Algorithm 799: revolve: an implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Trans. Math. Softw.*, 26(1):19–45, Mar. 2000. ISSN 0098-3500. doi:10.1145/347837.347846. URL `https://doi.org/10.1145/347837.347846`.

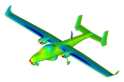
[10] N. Kukreja, J. Hückelheim, M. Lange, M. Louboutin, A. Walther, S. Funke, and G. Gorman. High-level python abstractions for optimal checkpointing in inversion problems. 01 2018. doi:10.48550/arXiv.1802.02474.

[11] E. Hairer and G. Wanner. *Solving ordinary differential equations II*. Springer Berlin, Heidelberg, 1996. doi:https://doi.org/10.1007/978-3-642-05221-7.

[12] C. Kennedy and M. Carpenter. Diagonally implicit runge-kutta methods for ordinary differential equations. a review. 03 2016.

[13] R. Alexander. Diagonally implicit runge–kutta methods for stiff o.d.e.'s. *SIAM Journal on Numerical Analysis*, 14(6):1006–1021, 1977. doi:10.1137/0714068.

[14] NASA. Common research model. http://commonresearchmodel.larc.nasa.gov/, 2012.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

DATA-DRIVEN ACCELERATION OF MULTIDISCIPLINARY SIMULATIONS FOR MULTI-QUERY APPLICATIONS IN AEROSPACE ENGINEERING

Susanna Baars[1*] and Ulrich Römer[1]

1: Institute for Acoustics and Dynamic
Technische Universität Braunschweig
Langer Kamp 19, 38106 Braunschweig, Germany
s.baars@tu-braunschweig.de, https://www.tu-braunschweig.de/inad

Abstract. *Aviation design and analysis demand precise multiphysical predictions, exemplified by fluid-structure interaction (FSI). For complex problems, partitioned approaches are often the only viable option, requiring strong coupling to handle feedback loops. Significant progress has been achieved in efficiently simulating strongly coupled systems using data-driven approaches. Key developments include quasi-Newton methods such as the matrix-free interface quasi-Newton method (IQN-ILS). Recently, surrogate-based enhancements for these methods have been explored, where the surrogate provides both an initial guess and a supplementary Jacobian. In multi-query applications like optimization and uncertainty quantification, information from previous parametric evaluations can further accelerate computations using data-driven surrogates, although this area remains underexplored. Our goal is to address this gap by developing a scheme for parametric multidisciplinary problems. We use previous parametric evaluations to refine the Jacobian approximations and data-driven surrogates for initial estimates of coupling variables and illustrate the approach with two numerical examples.*

Keywords: Multidisciplinary analysis, Surrogate model, Quasi-Newton method, Multi-query application

1 INTRODUCTION

The analysis and design of aircraft systems includes multiple disciplines, such as aerodynamics, structural mechanics, control systems, and thermodynamics [1]. Coupled problems can be solved either monolithically or partitioned, with the latter often being the only feasible option for highly complex problems. In many cases, feedback loops cannot be ignored, necessitating a strong coupling [2]. The required coupling iterations increase the computational expense. Additionally, complexity and computational demands increase significantly when multiphysical simulations need to be run repeatedly. However, multi-query applications such as optimization and uncertainty quantification often require repetitive evaluations.

Simulations for strongly coupled systems have become more efficient through data-driven strategies. Advances in quasi-Newton methods contribute significantly, such as the matrix-free interface quasi-Newton method (IQN-ILS), which employs a least-squares model to approximate the inverse Jacobian. Recently, efforts have been made to explore surrogate-based enhancements for these methods. The strategies include the use of previous time steps, system coarsening and physics simplification. Thereby, the surrogate provides an initial guess and a supplementary Jacobian.

However, most of the applications considered so far are single-query applications. In multi-query applications parametric data-driven surrogates are integral to gradient-free optimization and uncertainty quantification, reducing the number of required evaluations by predicting responses for unknown parameters using existing data. This article examines the intersection of these two areas: instead of using surrogates to predict the response in multi-query applications, surrogates are used to reduce the number of coupling interactions, thereby reducing the computational expense without compromising accuracy. We will examine cost savings in multi-query applications through numerical examples.

We begin with a literature review of existing quasi-Newton methods for accelerating multidisciplinary simulations in section 2. Next, we present our approach for accelerating parametric multidisciplinary systems in section 3 and conduct numerical experiments in section 4. Finally, we conclude and provide an outlook in section 5.

2 DATA-DRIVEN ACCELERATION OF QUASI-NEWTON METHODS

Partitioned multidisciplinary systems can be accelerated with data-enhanced quasi-Newton methods. The current research status of this approach is discussed in section 2.1. To the best of our knowledge, incorporating data from neighboring parametric evaluations has not been explored in the literature. However, utilizing data from previous time steps represents the state of the art and is covered in section 2.2.

2.1 Quasi-Newton methods

The quasi-Newton method presented in this section is a black-box approach that can in principle be applied to various partitioned multidisciplinary systems. However, FSI is a prototypical example of a multiphysical system. Most of the test cases used in the literature and the test cases in this work involve FSI. Therefore, we start by introducing an abstract version of the FSI equations.

An FSI problem involves two submodels: a fluid dynamical model represented by a map $\mathcal{F} : \mathbb{R}^{n_\mathrm{S}} \to \mathbb{R}^{n_\mathrm{F}}$, from interface displacements to interface pressure/traction compo-

nents, and a structural model $\mathcal{S} : \mathbb{R}^{n_\mathrm{F}} \to \mathbb{R}^{n_\mathrm{S}}$. The common interface is obtained as the intersection $\Gamma_\mathrm{I} = \Gamma_\mathrm{F} \cap \Gamma_\mathrm{S}$ of fluid and structural boundaries respectively. This setup is known as Dirichlet-Neumann partitioning [3]. In the following, the displacements and the pressure and traction components at the interface are denoted as $\mathbf{x}$ and $\mathbf{y}$, yielding

$$\tilde{\mathbf{y}} = \mathcal{F}(\mathbf{x}), \tag{1}$$

$$\tilde{\mathbf{x}} = \mathcal{S}(\mathbf{y}). \tag{2}$$

The fluid solver is executed first in the coupling iterations, leading to the following residual map

$$\mathcal{R}(\mathbf{x}) = \mathcal{S} \circ \mathcal{F}(\mathbf{x}) - \mathbf{x} = \tilde{\mathbf{x}} - \mathbf{x}. \tag{3}$$

Following [3], we denote the residual as $\mathcal{R}(\mathbf{x}) = \mathbf{r}$ in this work.

The nonlinear system of equations defined by $\mathcal{R}(\mathbf{x}) = \mathbf{0}$ can be solved using quasi-Newton methods. When the Jacobian is explicitly constructed, the quasi-Newton update takes the form

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \left(\delta^k_\mathbf{x}\mathbf{r}\right)^{-1}(-\mathbf{r}^k). \tag{4}$$

Here, $\mathbf{x}^k$ is the input at the current iteration, and $\mathbf{x}^{k+1}$ is the updated input for the next iteration. The term $-\mathbf{r}^k$ corresponds to the difference between the desired residual, which is zero, and the current residual $\mathbf{r}^k$. Moreover, $\delta^k_\mathbf{x}\mathbf{r}$ denotes the approximate Jacobian evaluated at $\mathbf{x}^k$.

The family of approaches studied, for example in [4] and [5], differ from the classical formulation by computing the update as

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \delta^k_\mathbf{r}\tilde{\mathbf{x}}(-\mathbf{r}^k) + \mathbf{r}^k. \tag{5}$$

Instead of constructing the Jacobian with respect to the input $\mathbf{x}^k$, this method builds the inverse Jacobian with respect to the output $\tilde{\mathbf{x}}^k$. To compensate, $\mathbf{r}^k$ is added to the update. A detailed derivation and the numerical motivations behind this approach are provided in [3]. By defining $\Delta\tilde{\mathbf{x}}^k = \delta^k_\mathbf{r}\tilde{\mathbf{x}}(-\mathbf{r}^k)$, the update becomes

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \Delta\tilde{\mathbf{x}}^k + \mathbf{r}^k. \tag{6}$$

To approximate the inverse Jacobian $\delta^k_\mathbf{r}\tilde{\mathbf{x}}$, we collect the residual vectors $\mathbf{r}^i$ and the output vectors $\tilde{\mathbf{x}}^i$ from all iterations i up to the current one and compute the differences

$$\delta\mathbf{r}^i = \mathbf{r}^{i+1} - \mathbf{r}^k \qquad \text{for} \quad i = 0, \ldots, k-1, \tag{7}$$

$$\delta\tilde{\mathbf{x}}^i = \tilde{\mathbf{x}}^{i+1} - \tilde{\mathbf{x}}^k \qquad \text{for} \quad i = 0, \ldots, k-1. \tag{8}$$

These difference vectors are stored in the matrices $\mathbf{R}^k$ and $\tilde{\mathbf{X}}^k$. Since the method requires at least two evaluations to start, a constant relaxation factor is used for the second iteration [4]. This method is referred to as IQN-ILS, an abbreviation for Interface Quasi-Newton technique with an approximation for the Inverse of the Jacobian from a Least-Squares model.

We construct a linear regression model using the calculated differences as training data, with $\delta\mathbf{r}^i$ and $\delta\tilde{\mathbf{x}}^i$ serving as input and output variables, respectively. The model coefficients correspond to the inverse Jacobian when explicitly computed. Because the dimensionality

of $\delta\mathbf{r}^i$ typically exceeds the number of observations, we perform the regression in dual form [6] and we arrive at

$$\Delta\tilde{\mathbf{x}}^k = \mathbf{f}^k(-\mathbf{r}^k) = \tilde{\mathbf{X}}^k \left(\left(\mathbf{R}^k\right)^T \mathbf{R}^k \right)^{-1} \left(\mathbf{R}^k\right)^T (-\mathbf{r}^k), \tag{9}$$

where $\mathbf{f}^k$ represents the linear model. Since the linear regression problem is underdetermined, numerical issues may arise.

Although the resulting update is independent of the order in which the expression is evaluated, it significantly affects the memory consumption. Explicitly computing the approximated Jacobian $\delta^k_{\mathbf{r}}\tilde{\mathbf{x}}$ results in a quadratic matrix of size $d_i \times d_i$, where d_i is the dimensionality of the interface variable and should thus be avoided.

2.2 Including information from past time steps into quasi-Newton methods

In [3], Delaisse et al. explore various methods to incorporate lower-fidelity information into the Jacobi matrix approximation discussed above. In this section, we focus on integrating evaluations from previous time steps, as this strategy is similar to incorporating information from neighboring parametric evaluations, which is the focus of our study.

Since the Jacobian remains similar across nearby time steps, evaluations from previous time steps can be incorporated into the data-driven approximation. A straightforward approach is to stack the difference matrices from past time steps with those from the current iteration as

$$\begin{aligned} \mathbf{R}^k_{\mathbf{m}=(m_c-q+1,\dots,m_c)} &= \begin{bmatrix} \mathbf{R}_{m_c-q+1} & \dots & \mathbf{R}_{m_c-1} & \mathbf{R}^k_{m_c} \end{bmatrix} \\ \tilde{\mathbf{X}}^k_{\mathbf{m}=(m_c-q+1,\dots,m_c)} &= \begin{bmatrix} \tilde{\mathbf{X}}_{m_c-q+1} & \dots & \tilde{\mathbf{X}}_{m_c-1} & \tilde{\mathbf{X}}^k_{m_c} \end{bmatrix}. \end{aligned} \tag{10}$$

Here, m_c is the current time step corresponding to the current time t_c and q is the number of time steps included in the Jacobian approximation, including the current one. The regression described in eq. (9) is then performed using the matrices $\mathbf{R}^k_{\mathbf{m}=(m_c-q+1,\dots,m_c)}$ and $\tilde{\mathbf{X}}^k_{\mathbf{m}=(m_c-q+1,\dots,m_c)}$ instead of $\mathbf{R}^k$ and $\tilde{\mathbf{X}}^k$. In the matrices from previous time steps, the superscript is omitted to indicate that every evaluation up to convergence is used.

A major advantage of this approach is its simplicity. However, a drawback is the amplification of numerical issues caused by multicollinearity from conflicting information. Various methods to mitigate these issues have been proposed and investigated [7]. The most commonly employed method uses QR decomposition to identify and remove near-linearly dependent columns. In this process, data points from earlier time-steps are discarded, as they are assumed to be less relevant.

An approach to exploit the hierarchy of fidelities more explicitly is to apply a recursive scheme, see [5], as

$$\mathbf{f}^k_m(-\mathbf{r}^k) = \mathbf{f}_{m-1}(-\mathbf{r}^k) + \delta^k_m(-\mathbf{r}^k) \qquad \text{for} \quad m = m_c - q + 1, \dots, m_c. \tag{11}$$

Here, $\delta^k_m(x)$ is a linear model that corrects discrepancies between the predictions of older linear models and the observations. The inverse Jacobian can be constructed either explicitly or implicitly [8].

In addition to using surrogates to generate a low-fidelity approximation of the Jacobian, [3] also uses simplified physics or coarser meshes to produce initial estimates of the coupling variables to initialize the fixed-point iteration. In the following, we denote converged

displacements with $\mathbf{w}$ instead of $\tilde{\mathbf{x}}$ to distinguish them from displacements computed in an arbitrary iteration. For a stationary model, the coupling variables are initialized as $\hat{\mathbf{w}}$, with the hat indicating that these are surrogate-generated estimates. In the case of a time-dependent simulation, it is preferable to apply the update

$$\check{\mathbf{w}}_{m_c} = \mathbf{w}_{m_c-1} + (\hat{\mathbf{w}}_{m_c} - \hat{\mathbf{w}}_{m_c-1}) \tag{12}$$

to eliminate drift.

3 ACCELERATING PARAMETRIC MULTIDISCIPLINARY SYSTEMS

Our objective is to develop a framework that accelerates parametric multidisciplinary simulations for multi-query applications. Building on the ideas from [3], we aim to incorporate parametric evaluations in both the estimation of the Jacobian matrix and the derivation of an initial estimate. Developing such a framework involves addressing three key challenges:

1. **Memory management**: To enhance the quasi-Newton method with data from parametric evaluations, it is essential to store not only the $\mathbf{R}_m$ and $\tilde{\mathbf{X}}_m$ matrices from the last few time steps but also those from all time steps across neighboring evaluations. In the following, we will use the notation $\mathbf{R}_{m,n}$ and $\tilde{\mathbf{X}}_{m,n}$ to indicate dependence not only on the time step m but also on the index n of the parameter $\mathbf{p}_n$. In addition, the converged interface data $\mathbf{w}_{m,n}$ must be stored in order to generate surrogates for the initial predictions. This significantly increases the need for efficient memory management.

2. **Absence of a clear model hierarchy**: Unlike the time reuse surrogate, the hierarchy of information fidelity is less clear. An exception occurs when stationary models are optimized using gradient-based methods, allowing the most recent evaluations to be considered closest to the upcoming ones, an assumption that does not hold for other multi-query applications. This ambiguity is problematic because both approaches to incorporating surrogate information rely on a clear information hierarchy. While this issue is particularly evident for the additive hierarchical approach in eq. (11), the the simple approach in eq. (10) also requires a hierarchy by discarding past information during filtering and setting a cutoff for the number of time steps used.

3. **Multiple sources of information for time-dependent simulation**: In the case of stationary models, parametric evaluations are the only option for data-driven acceleration, unless models of lower fidelity are available. However, in time-dependent parametric simulations, two different sources of information can improve both the data-driven acceleration of quasi-Newton iterations and the generation of initial estimates. Besides schemes that rely solely on previous time steps or parametric evaluations for generating initial guesses and enhancing the Jacobian approximation, hybrid schemes can be developed.

   This raises the question of how to combine the sources most effectively. A major disadvantage of solely reusing previous time steps is the reliance on extrapolation. However, depending on the sensitivity of the model to changing parameter values, nearby evaluations may diverge more than the differences between successive time steps.

3.1 Setup for Numerical Experiments

In the following, we will present our basic ideas for a famework for data-driven acceleration of parametric multidisciplinary simulations. We keep the framework simple while addressing the challenges and questions raised above. Figure 1 schematically illustrates the basic idea.

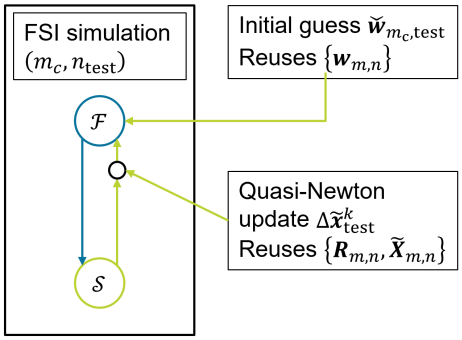


Figure 1: Schematic illustration of how previous time steps and parametric evaluations are used to accelerate an FSI simulation at time step m_c and the parameter index n_{test}.

Our framework consists of the following steps:

1. **Training data generation**: The initial step involves generating training data by evaluating the model at a limited set of parameter points. This is achieved using a linear derivative-based predictor and the IQN-ILS method. For time-dependent cases, information from previous time steps is reused according to eq. (10). After each time step, the converged coupling variables $\mathbf{w}_{m,n}$, along with the matrices $\mathbf{R}_{m,n}$ and $\tilde{\mathbf{X}}_{m,n}$, are stored in a HDF5 file.

2. **Dimension reduction**: In addition to using HDF5 files, we use dimensionality reduction via proper orthogonal decomposition (POD) to achieve efficient memory management (item 1). For this purpose, the matrices are aggregated into snapshot matrices as

$$\begin{aligned} \mathbf{R}_{\text{snap}} &= \left[\mathbf{R}_{m=0,n=0} \quad \dots \quad \mathbf{R}_{m=M,n=N}\right], \\ \tilde{\mathbf{X}}_{\text{snap}} &= \left[\tilde{\mathbf{X}}_{m=0,n=0} \quad \dots \quad \tilde{\mathbf{X}}_{m=M,n=N}\right], \\ \mathbf{W}_{\text{snap}} &= \left[\mathbf{w}_{m=0,n=0} \quad \dots \quad \mathbf{w}_{m=M,n=N}\right]. \end{aligned} \tag{13}$$

   Here, M denotes the final time step and N the highest parametric index. Before applying Singular Value Decomposition (SVD), the mean is subtracted only from the snapshot matrix $\mathbf{W}_{\text{snap}}$, while the matrices $\mathbf{R}_{\text{snap}}$ and $\tilde{\mathbf{X}}_{\text{snap}}$ remain unaltered. This selective mean removal ensures that a zero entry in the original data stays zero in the reduced-order space, which is crucial for maintaining the desirable convergence properties of the IQN-ILS method. Subsequently, SVD is performed, and insignificant modes are discarded, yielding the reduced matrices $\mathbf{W}_{\text{red};m,n}$, $\mathbf{R}_{\text{red};m,n}$, and $\mathbf{X}_{\text{red};m,n}$. The application of POD in the context of quasi-Newton methods has already been proposed several times, e.g. in [7, 9]. If the amount of data is too large to load into memory at once, incremental dimensionality reduction can be used [10].

3. **Surrogate model**: We construct a data-driven surrogate that produces an initial estimate similar to [11], which addresses uncertainty quantification for multidisciplinary systems. In this approach, a multidisciplinary model is evaluated at an initial design of experiment, and Gaussian processes predict the coupling variables at new sampling points.

   While nonlinear models such as Gaussian processes may be more appropriate for larger numbers of evaluations and models with high nonlinearity, we opt for a linear model for basic conceptual studies. The reduced converged coupling variables are predicted as a function of the parameters:

   $$\hat{\mathbf{w}}_{\mathrm{red};m,n} = \mathbf{f}_{\mathrm{pred};m}(\mathbf{p}_n). \tag{14}$$

   This prediction within the reduced space constitutes an application of the Proper Orthogonal Decomposition and Interpolation (PODI) approach [12]. For time-dependent models, the update is calculated using eq. (12).

4. **Test execution**: We generate the combined matrices as described in eq. (10) during execution, with their construction dependent on the parameter q. To selectively also include the training data from the parametric evaluations, we define a parameter v. As the other parametric simulations have already been run, we can include both previous and subsequent evaluations. Thus, we obtain the merged matrices $\mathbf{R}^k_{\mathbf{m}(n),\mathbf{n}}$ and $\tilde{\mathbf{X}}^k_{\mathbf{m}(n),\mathbf{n}}$, where the multi-index $\mathbf{m}(n)$ is set to

   $$\mathbf{m}(n) = \begin{cases} \{\, m_c \,\} & \text{if } q = 1 \wedge n = n_{\mathrm{test}}, \\ \{m_c - q + 1,\ \ldots,\ m_c\} & \text{if } q > 1 \wedge n = n_{\mathrm{test}}, \\ \emptyset & \text{if } v = 0 \wedge n \in \{0, \ldots, N\}, \\ \{\, m_c \,\} & \text{if } v = 1 \wedge n \in \{0, \ldots, N\}, \\ \left\{m_c - \frac{v-1}{2},\ \ldots,\ m_c,\ \ldots,\ m_c + \frac{v-1}{2}\right\} & \text{if } v > 1 \wedge n \in \{0, \ldots, N\}, \end{cases}$$

   and $\mathbf{n} = (0, \ldots, N, n_{\mathrm{test}})$.

   We opt for ridge regression, also known as Tikhonov regularization, instead of the conventional QR-based filtering. This choice is based on its simplicity and the fact that, unlike QR-based filtering, ridge regression does not require a fidelity hierarchy. Consequently, the update in eq. (9) is modified to

   $$\Delta\tilde{\mathbf{x}}^k = \mathbf{f}^k(-\mathbf{r}^k) = \tilde{\mathbf{X}}^k_{\mathbf{m}(n),\mathbf{n}} \left(\left(\mathbf{R}^k_{\mathbf{m}(n),\mathbf{n}}\right)^\top \mathbf{R}^k_{\mathbf{m}(n),\mathbf{n}} + \lambda \mathbf{I} \right)^{-1} \left(\mathbf{R}^k_{\mathbf{m}(n),\mathbf{n}}\right)^\top (-\mathbf{r}^k), \tag{15}$$

   where λ is a user-defined constant that controls the regularization strength.

   Note that no effort was made to eliminate or prioritize specific parametric data points in these conceptual studies. As a result, one of the aforementioned challenges, the lack of a clear model hierarchy, is not fully addressed.

5. **Evaluation**: For evaluation, the average number of iterations per time step is considered.

4 NUMERICAL EXPERIMENTS

Two complementary numerical experiments are conducted to evaluate the methodology outlined in section 3.1. The first test case is the MOK benchmark, a well-established transient 2D FSI test case that has previously been used to benchmark quasi-Newton methods [9]. The second experiment involves a stationary 3D model: a flexible wing inspired by the ONERA M6 model. This example demonstrates the application of the approach to a stationary problem with a larger number of coupling variables, representative of multidisciplinary simulations in the aerospace context.

4.1 MOK example

The MOK test case is a classic FSI benchmark first described by Daniel P. Mok and further studied, for example, in [13]. Figure 2 shows the benchmark consisting of a 2D convergent fluid channel containing a flexible wall structure attached to the bottom. The files required to replicate the nominal case using the Kratos Multiphysics framework [14] as well as a description are available in [15]. We disable the relative convergence criterion by setting the tolerance to zero. This change ensures that the convergence evaluation remains unbiased when improved initial values are provided. To generate the training data, two parameters are made stochastic: The beam's Young's modulus, E, and a newly introduced scaling factor, ϕ, for the inlet velocity profile, with the mean values $\bar{E} = 2.3 \cdot 10^6\,\mathrm{Pa}$ and $\bar{\phi} = 1$.

The stochastic inlet velocity profile becomes

$$v_x = \begin{cases} \phi \times 2 \times 0.06067 \times \left(1 - \cos\left(\dfrac{\pi t}{10}\right)\right) \times y \times (1-y) & \text{if } t < 10 \\ \phi \times 4 \times 0.06067 \times y \times (1-y) & \text{otherwise.} \end{cases} \tag{16}$$

In both cases, we assume a normal distribution where the standard deviation is 2% of the mean value, as

$$\begin{aligned} E_{\mathrm{ri}} &\sim \mathcal{N}(2.3 \cdot 10^6\,\mathrm{Pa}, (4.6 \cdot 10^4\,\mathrm{Pa})^2), \\ \phi_{\mathrm{ri}} &\sim \mathcal{N}(1, 0.02^2). \end{aligned} \tag{17}$$

We generate 10 sample points using Latin hypercube sampling and evaluate the model to produce the training data as described above. Next, we retain 20 modes, which is the maximum number retained for the MOK benchmark reported in [9], where quasi-Newton methods were combined with dimensionality reduction, and apply POD. The user-defined regularization constant is set to $\lambda = 10^{-12}$.

Subsequently, we generate two test points, which are the nominal case with $E_{\mathrm{test},0} = 2.3{\cdot}10^6\,\mathrm{Pa}$ and $\phi_{\mathrm{test},0} = 1$, and an offset case with $E_{\mathrm{test},1} = 1.02{\cdot}2.3{\cdot}10^6\,\mathrm{Pa} = 2.346{\cdot}10^6\,\mathrm{Pa}$ and $\phi_{\mathrm{test},1} = 1.02$ to study the robustness. Figure 3 shows the design of experiment for generating the training data and the two test cases. We then evaluate the model for different values of q and v with and without the PODI predictor.

The results for the test case 1 are shown in table 1. The best performance is achieved using the PODI predictor and setting $q = 6$ and $v = 1$. In this configuration, an average of 1.892 iterations is required; this amounts to 35% of the average iterations without any reuse, which is 5.46. If only past time steps are reused, the minimum average number of

iterations is 2.156. Although this difference is modest, note that the average is already very low and cannot drop below 1.

Table 2 presents the results for test case 2. The best performance is achieved with $q = 11$, $v = 0$ and the PODI predictor, which results in an average of 1.94 evaluations per time step and is nearly equal to the best performance of test case 1.

Considering both test cases, the PODI predictor delivers a clear advantage, while incorporating parametric evaluations into the Jacobian estimation has a marginal effect compared to only reusing previous time steps. For a final evaluation, a more complex example should be used, one in which the average number of iterations is well above two even when these methods are applied.

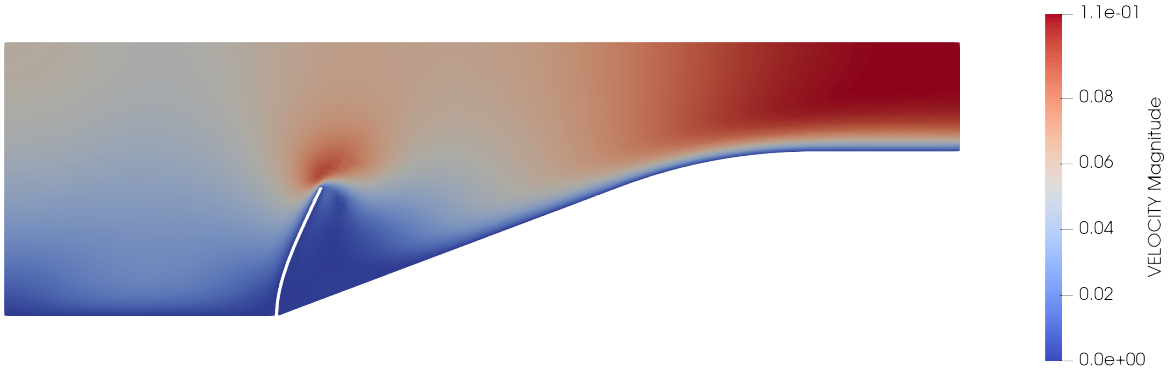


Figure 2: Velocity profile of the fluid model for the Mok benchmark at $t = 8\,\mathrm{s}$

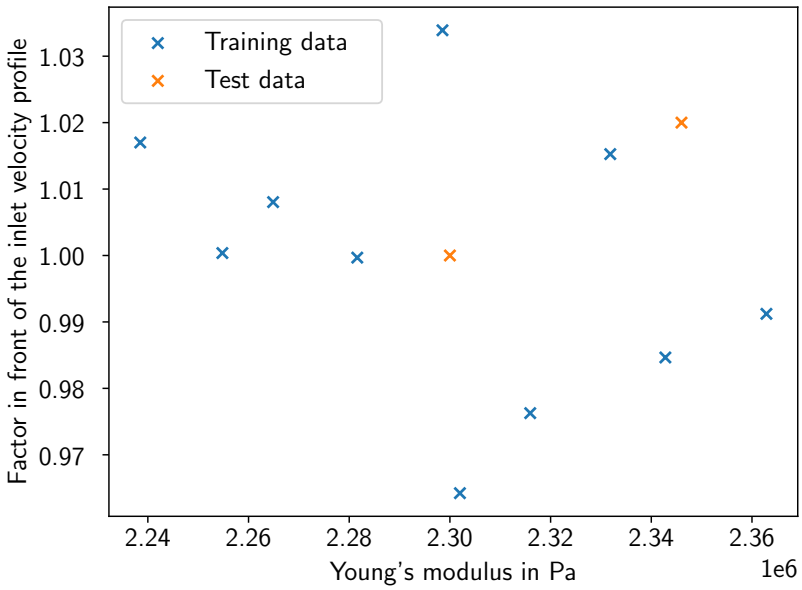


Figure 3: Training/test samples for the two-dimensional MOK benchmark generated with the Latin hypercube sampling.

T	APT	PODI Pred.	Perf. Rank	Avg. Iter.
1	0	no	20	5.46
1	0	yes	19	4.112
1	1	no	18	3.584
1	1	yes	16	2.616
1	7	no	17	3.284
1	7	yes	14	2.184
1	11	no	21	not converged
1	11	yes	15	2.332
6	0	no	11	2.156
6	0	yes	3	1.924
6	1	no	9	2.108
6	1	yes	1	1.892
6	7	no	10	2.128
6	7	yes	4	1.944
6	11	no	21	not converged
6	11	yes	7	2.056
11	0	no	13	2.172
11	0	yes	2	1.896
11	1	no	8	2.072
11	1	yes	6	2.048
11	7	no	10	2.128
11	7	yes	21	not converged
11	11	no	12	2.164
11	11	yes	5	2.024

Table 1: Performance of the MOK example at $\mathrm{E} = 2.3 \times 10^6\,\mathrm{Pa}$ and $f_v = 1$. **T**: Reused time steps of the current parametric evaluation (q); **APT**: Reused time steps of all other parametric evaluations (v); **PODI Pred.**: Using PODI surrogate (yes/no); **Perf. Rank**: Relative performance ranking; **Avg. Iter.**: Average iterations per time step.

T	APT	PODI Pred.	Perf. Rank	Avg. Iter.
1	0	no	19	5.624
1	0	yes	18	4.012
1	1	no	16	3.34
1	1	yes	15	3.052
1	7	no	17	3.556
1	7	yes	14	2.352
1	11	no	20	not converged
1	11	yes	20	not converged
6	0	no	12	2.204
6	0	yes	3	1.956
6	1	no	10	2.18
6	1	yes	7	2.056
6	7	no	9	2.172
6	7	yes	6	2.0
6	11	no	20	not converged
6	11	yes	20	not converged
11	0	no	11	2.192
11	0	yes	1	1.94
11	1	no	20	not converged
11	1	yes	2	1.948
11	7	no	8	2.168
11	7	yes	4	1.972
11	11	no	13	2.228
11	11	yes	5	1.984

Table 2: Performance of the MOK example at E = 2.346×10^6 Pa and $f_v = 1.02$. **T**: Reused time steps of the current parametric evaluation (q); **APT**: Reused time steps of all other parametric evaluations (v); **PODI Pred.**: Using PODI surrogate (yes/no); **Perf. Rank**: Relative performance ranking; **Avg. Iter.**: Average iterations per time step.

4.2 Flexible wing

The Onera M6 wing is a standard CFD validation test case for the transonic regime and has been used to validate a transonic potential flow solver in Kratos Multiphysics with far-field boundary conditions [16]. For our study, we use an example inspired by this validation case, but we render the wing flexible. Note that this model is not designed to reproduce aeroelastic phenomena with high accuracy; rather, it serves as an illustrative example of how the method can be applied within an aviation context. While the physical values of the fluid model are chosen to be realistic, the mesh is coarse. The material values for the solid model are chosen so that the wing deforms enough to require feedback coupling, but not too much to allow the linear elasticity assumption to hold. Figure 4 displays both the deformed and undeformed wing for the nominal setup. The presented interface consists of 36533 nodes, resulting in a 109599 dimensional interface variable. The fluid and solid domains are modeled as follows:

- Fluid: Transonic potential flow as described in [16]
  - Angle of attack (α): $3°$
  - Free stream Mach number (Ma): 0.82
- Solid: Linear elastic
  - Density (ρ_s): $7.850 \cdot 10^3 \, \frac{\mathrm{kg}}{\mathrm{m}^3}$
  - Young's modulus (E): $3 \cdot 10^9 \, \mathrm{Pa}$
  - Poisson ratio (ν): 0.3

We select the angle of attack and Mach number as stochastic parameters. As with the MOK example, we assume they are normally distributed with a standard deviation of 2% of the nominal value

$$\begin{aligned} \alpha_\mathrm{ri} &\sim \mathcal{N}(3°, (0.06°)^2), \\ \mathrm{Ma}_\mathrm{ri} &\sim \mathcal{N}(0.82, 0.0164^2). \end{aligned} \tag{18}$$

This time, we generate 12 sample points using Latin hypercube sampling and evaluate the model to create the training data. One of the sample points does not converge in the first iteration of the fluid solver. To ensure that any suboptimal results are not due to an overly strict cutoff value, we retain all components when applying POD. With the remaining 11 sample points, each converging within 4 iterations, we obtain 11 modes for the interface data and 33 modes for the quasi-Newton method. Again, we set the user-defined regularization constant to $\lambda = 10^{-12}$.

Two test cases are set up analogous to those for the MOK example. The first corresponds to the nominal example with $\alpha_{\mathrm{test},\,0} = 3°$ and $\mathrm{Ma}_{\mathrm{test},\,0} = 0.82$, while the second is offset by adding the standard deviation such that $\alpha_{\mathrm{test},\,1} = 3.06°$ and $\mathrm{Ma}_{\mathrm{test},\,1} = 0.8364$. However, the number of iterations required is identical for both test cases and all setups, as shown in table 3. The best performance is achieved by reusing the parametric evaluations within the quasi-Newton method. Since the quasi-Newton method is applied only in the second iteration, the obtained 2 iterations are the optimal result. Initializing the displacement with the PODI surrogate prediction saves one iteration. Interestingly,

combining both methods worsens performance, increasing the number of iterations from 4 to 6. Since changing the regularization constant does not affect this behavior, it is unlikely to be due to ill-conditioning. Instead, it is more likely caused by the nonlinearity of the model. During the first iteration, initializing displacements to zero causes them to change significantly, which in turn has a large influence on the implicit estimation of the Jacobian.

One possible solution is to apply an additive multifidelity approach, as shown in eq. (11), at least for the current evaluation to ensure it receives priority over all other parametric evaluations.

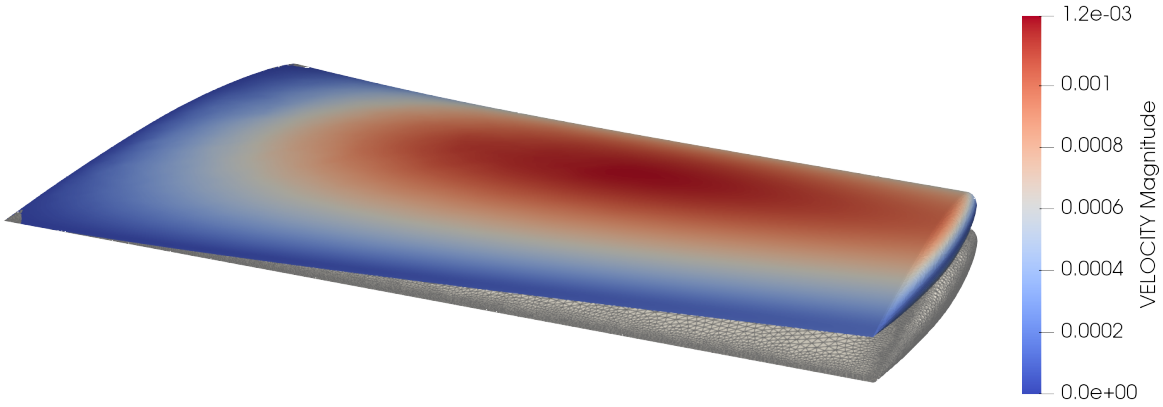


Figure 4: Undeformed and deformed ONERA M6

AP	PODI Pred.	Perf. Rank	Iter.
no	no	1	4
yes	no	2	2
no	yes	3	3
yes	yes	4	6

Table 3: Performance of the ONERA M6 example. The results are identical for ($\alpha = 3°$, $Ma = 0.82$) and ($\alpha = 3.06°$, $Ma = 0.8364$). **AP**: Reuse all parametric evaluations for the Jacobian approximation (yes/no); **PODI Pred.**: Using PODI surrogate (yes/no); **Perf. Rank**: Relative performance ranking; **Iter.**: Number of iterations.

5 CONCLUSIONS AND OUTLOOK

We demonstrated that simulating parametric multidisciplinary systems can be significantly accelerated by reusing data from neighboring evaluations. The data is utilized in two distinct ways: first, to create an initial guess, and second, to accelerate the fixed-point iterations.

Both applications require storing multiple vectors for all time steps, with each vector's dimensionality matching that of the interface variable. To address the memory requirements, we use the HDF5 format and apply POD as a postprocessing step. The initial

guess for each time step can be generated by interpolating over the reduced interface variables (PODI).

By reusing previous parametric evaluations, we observed a significant reduction in the number of evaluations required in two distinct test cases. For the transient example, the PODI predictor proved particularly effective compared to linear derivative-based extrapolation. In the stationary example, each method reduced evaluations, although their combination failed. It should be noted, however, that more complex test cases should be selected, as the test cases used in this study often achieved near-optimal performance.

However, to accelerate parametric problems in real-world multi-query applications, the proposed scheme will require further optimization. The method should work out of the box, without user tuning. To avoid explicitly specifying time steps and parameter evaluations for Jacobian approximation, a fidelity hierarchy should be established automatically. This could be achieved by ranking model fidelity using a distance metric based on the similarity of converged coupling variables.

To resolve the conflict of having to choose between the PODI predictor and extrapolation for time-dependent systems, the residuals between the actual observation and the PODI predictions should be extrapolated. In addition, it would be advantageous to use Gaussian processes when high nonlinearity is present.

In summary, the potential for data-driven acceleration of parametric multi-query applications has been demonstrated. However, to implement this approach in a black-box multiphysics framework such as Kratos Multiphysics, strategies must be devised to minimize the need for manual parameter tuning.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2163/1 - Sustainable and Energy Efficient Aviation – Project-ID 390881007

REFERENCES

[1] L. Brevault, M. Balesdent, and J. Morio. *Aerospace System Analysis and Optimization in Uncertainty*, volume 156 of *Springer Optimization and Its Applications*. Springer International Publishing. ISBN 978-3-030-39125-6 978-3-030-39126-3. doi:10.1007/978-3-030-39126-3. URL `http://link.springer.com/10.1007/978-3-030-39126-3`.

[2] J. Degroote. Partitioned simulation of fluid-structure interaction: Coupling black-box solvers with quasi-newton techniques. 20(3):185–238. ISSN 1134-3060, 1886-1784. doi:10.1007/s11831-013-9085-5. URL `http://link.springer.com/10.1007/s11831-013-9085-5`.

[3] N. Delaissé, T. Demeester, D. Fauconnier, and J. Degroote. Surrogate-based acceleration of quasi-newton techniques for fluid-structure interaction simulations. 260:106720. ISSN 00457949. doi:10.1016/j.compstruc.2021.106720. URL `https://linkinghub.elsevier.com/retrieve/pii/S004579492100242X`.

[4] J. Degroote, R. Haelterman, S. Annerel, P. Bruggeman, and J. Vierendeels. Performance of partitioned procedures in fluid–structure interaction. 88(7):446–457.

ISSN 00457949. doi:10.1016/j.compstruc.2009.12.006. URL https://linkinghub.elsevier.com/retrieve/pii/S0045794909003022.

[5] A. Bogaers, S. Kok, B. Reddy, and T. Franz. Quasi-newton methods for implicit black-box FSI coupling. 279:113–132. ISSN 00457825. doi:10.1016/j.cma.2014.06.033. URL https://linkinghub.elsevier.com/retrieve/pii/S0045782514002199.

[6] J. C. Ye. *Linear, Logistic, and Kernel Regression*, volume 37, pages 45–59. Springer Nature Singapore. ISBN 978-981-16-6045-0 978-981-16-6046-7. doi:10.1007/978-981-16-6046-7_3. URL https://link.springer.com/10.1007/978-981-16-6046-7_3. Series Title: Mathematics in Industry.

[7] R. Haelterman, A. Bogaers, K. Scheufele, B. Uekermann, and M. Mehl. Improving the performance of the partitioned QN-ILS procedure for fluid–structure interaction problems: Filtering. 171:9–17. ISSN 00457949. doi:10.1016/j.compstruc.2016.04.001. URL https://linkinghub.elsevier.com/retrieve/pii/S004579491630164X.

[8] T. Spenke, N. Hosters, and M. Behr. A multi-vector interface quasi-newton method with linear complexity for partitioned fluid–structure interaction. 361:112810. ISSN 00457825. doi:10.1016/j.cma.2019.112810. URL https://linkinghub.elsevier.com/retrieve/pii/S0045782519307029.

[9] R. Zorrilla and R. Rossi. A memory-efficient MultiVector quasi-newton method for black-box fluid-structure interaction coupling. 275:106934. ISSN 00457949. doi:10.1016/j.compstruc.2022.106934. URL https://linkinghub.elsevier.com/retrieve/pii/S0045794922001948.

[10] D. A. Ross, J. Lim, R.-S. Lin, and M.-H. Yang. Incremental learning for robust visual tracking. 77(1):125–141. ISSN 0920-5691, 1573-1405. doi:10.1007/s11263-007-0075-7. URL http://link.springer.com/10.1007/s11263-007-0075-7.

[11] A. Chaudhuri, R. Lam, and K. Willcox. Multifidelity uncertainty propagation via adaptive surrogates in coupled multidisciplinary systems. 56(1):235–249. ISSN 0001-1452, 1533-385X. doi:10.2514/1.J055678. URL https://arc.aiaa.org/doi/10.2514/1.J055678.

[12] H. V. Ly and H. T. Tran. Modeling and control of physical processes using proper orthogonal decomposition. 33(1):223–236. ISSN 08957177. doi:10.1016/S0895-7177(00)00240-5. URL https://linkinghub.elsevier.com/retrieve/pii/S0895717700002405.

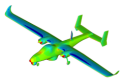
[13] J. G. Valdés Vázquez, J. Miquel Canet, and E. Oñate Ibáñez De Navarra. Nonlinear analysis of orthotropic membrane and shell structures including fluid-structure interaction. URL http://hdl.handle.net/2117/94177.

[14] V. M. Ferrándiz, P. Bucher, R. Zorrilla, S. Warnakulasuriya, A. Cornejo, R. Rossi, C. Roig, jcotela, J. Maria, tteschemacher, M. Masó, G. Casas, M. Núñez, P. Dadvand, S. Latorre, I. de Pouplana, J. I. González, AFranci, F. Arrufat, riccardotosi, A. Ghantasala, K. B. Sautter, P. Wilson, dbaumgaertner, B. Chandra, A. Geiser,

I. Lopez, lluís, jgonzalezusua, and J. Gárate. KratosMultiphysics/kratos: v10.1. URL https://zenodo.org/doi/10.5281/zenodo.3234644. Language: en.

[15] P. Bucher. KratosMultiphysics examples – co_simulation/validation/fsi_mok. URL https://github.com/KratosMultiphysics/Examples/tree/master/co_simulation/validation/fsi_mok.

[16] I. P. López Canalejo. A finite-element transonic potential flow solver with an embedded wake approach for aircraft conceptual design. URL https://mediatum.ub.tum.de/1633175.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

TOWARD DISTRIBUTED AND SCALABLE MULTI-DISCIPLINARY OPTIMIZATION WITH GEMSEO

Jean-Christophe Giret[1*], François Gallard[1], Antoine Dechaume[1], Jérôme Fasquel[1], Jean-François Figué[1] and Anne Gazaix[1]

1: IRT Saint Exupéry
4, Rue Tarfaya CS34436 Toulouse, France
jean-christophe.giret@irt-saintexupery.com, http://www.irt-saintexupery.com

Abstract. *Multidisciplinary design optimization (MDO) workflows are often executed locally on a single computer node. However, industrial-scale or collaborative MDO projects frequently require processes that span multiple machines. This need arises from the demand for specialized resources such as specific software, operating systems, hardware, or increased computing power through horizontal scaling. Additionally, multi-partner collaborations often prioritize intellectual property protection by obfuscating internal models while exposing only interfaces and input-output contracts.*

Although enabling technologies like job schedulers, SSH and HTTP client and servers, and cloud platforms are widely available, their adoption in MDO workflows remains limited. Implementing these technologies presents significant challenges for both end-users and MDO process integrators.

In the frame of the HORIZON European Projects NEXTAIR and RECET4RAIL, efforts have been put to bring a range of platform capabilities to selectively and seamlessly offload parts of the MDO process to remote machines, thus enabling extended-enterprise MDO: - A job scheduler discipline to delegate computations to HPC with job schedulers such as LSF or SLURM. - An SSH discipline that enables the delegation of the computation of a discipline to a distant node accessible via the SSH protocol - An HTTP discipline that provides a client and a server that enables the exposure of GEMSEO disciplines as web services.

These features introduce new challenges to the decentralized MDO process, particularly regarding observability and reliability. For reliability, given the risk of hardware failures or network instability in remote computing, specific developments have been made to selectively address errors and exceptions that may arise from hardware or network issues. Finally, capabilities have been developed to enable observability across the remote components of the distributed MDO processes.

Keywords: MDO, SSH, HTTP, Cloud computing, HPC, RPC, Job Scheduler, Data Management

1 INTRODUCTION

Multi-Disciplinary Optimization (MDO) is now recognized as a key method to obtain better design solutions while reducing the number of iterations in the design phase. This field is flourishing in many fields in the aerospace industry for a large panel of design and/or operation optimizations[1, 2], but also in other domains[3].

Many software and libraries have been developed to enable such MDO processes, alleviating the need for the user to integrate simulation software and implement MDO processes. They can be divided into two groups:

- Commercial, off-the-shelf software (COTS). Examples of such software suites are ModelCenter[4], ModeFrontier[5], Isight[6], pSeven[7], Optimus[8], and HEEDS[9]. Most of the time, they are part of a larger portfolio of applications and transverse tools that manage processes and data (Simulation Process and Data Management - SPDM) and handle the capability to carry out distributed MDO processes.

- Open-Source libraries: they enable end users to build custom MDO processes. OpenMDAO[10], GEMSEO[11] and CosApp[12] can be cited as open-source MDO libraries. All three are implemented in Python and heavily rely on the scientific Python ecosystem (Numpy, Scipy, Matplotlib, ...). They are toolboxes that provide the elementary bricks to implement MDO processes with more ease.

The current work focuses on the second group of open-source libraries. Very early[13–16], the need to scale these MDO processes appeared. By scalability, it is understood:

- the increase of the number of disciplines and software involved: MDO processes are using an increasing number of disciplines, thus increasing the infrastructure needs in terms of resources (CPU, RAM, HDD) but also in terms of complexity (HPC, GPU, Cloud infrastructure, OSes, ...). This increase in resource consumption is mainly found in high-fidelity disciplines.

- the increase of the number of domains and partners involved in an MDO process, with the scattering of the disciplines over several domains and/or external partners. This can be internal to an organization over several engineering domains. It can also involve several partners (for instance, in the frame of a system design program or the frame of a research project) in designing a system of systems while preserving intellectual property.

Yet, this involves distributed resources and/or partners that are heterogeneous in terms of hardware, software, communication protocols, security and access rules. It also introduces new challenges, particularly regarding reliability and observability: network communications can be unstable and challenge the overall reliability of distributed processes. Finally, distributed processes inherently generate distributed data, which can be difficult for the end user to access and use.

Numerous solutions[15] have been documented in the literature, most of the time leveraging available and ubiquitous technologies at creation time, ranging from CORBA[14, 17] to gRPC[18], including HTTP(S) and SSH[19] and relying on standard formats such as XML, JSON or Protobuf. Lupp et *al.* recent work[18] provides an exhaustive list of distributed MDO solutions carried out in the United States of America and provides an

implementation of generic interfaces for MDO leveraging the gRPC framework [20] and the protobuf protocol. Work has also been done in Europe, mainly in European projects, to build capabilities that enable remote MDO processes. For instance, the open-source RCE framework[19] using the SSH protocol and its usage in the AGILE[21] and AGILE 4.0[22] projects can be cited.

The GEMSEO library, developed to ease the creation and maintenance of industrial and large-scale MDO problems, lacked generic services to scale MDO processes in distributed environments. In the NEXTAIR European project frame, GEMSEO has been selected to orchestrate some of the partners' MDO processes. The latter are distributed internally in the partners' infrastructure and require other partners' external resources, as in the work of Brezillon *et al*[23]. After evaluating the existing open-source solutions, it turned out that none were generic enough to be integrated out-of-the-box in GEMSEO and provide the expected capabilities with full compatibility with its reconfigurable features. In light of this, several services have been developed specifically for GEMSEO to enable its scalability in a distributed environment, and, more specifically, to allow the realization of the NEXTAIR distributed MDO processes.

The paper is organized as follows: first, the extension of GEMSEO towards distributed MDO processes is presented. The technological and architectural choices made with respect to other solutions are also highlighted, and an overall view of the different services and their synergies is given. Then, a thorough description of each service is provided, with examples to illustrate their usage. Along with the conclusion, perspectives are proposed to pursue this work.

2 EXTENDING THE GEMSEO LIBRARY TOWARDS DISTRIBUTED MDO PROCESSES

2.1 GEMSEO: an engine enabling reconfigurable MDO processes

The GEMSEO (Generic Engine for Multidisciplinary Scenarios, Exploration and Optimization) open-source software has been developed by IRT Saint Exupéry since 2015 to support the deployment of MDO in the industry. GEMSEO is a Python library that can automatically generate MDO processes according to a catalog of MDO formulations and easily explore different numerical resolution strategies.

GEMSEO is responsible for managing the MDO scenario, which is the overall organization of the process, including the optimization problem (objective function, constraints, coupling variables, design variables), coupling procedures and optimization or design of experiment algorithms. While GEMSEO offers the most common monolithic MDO formulations (Multidisciplinary Design Feasible and Individual Discipline Feasible, see[24]), it mainly focuses on multi-level formulations that divide the optimization problem into sub-optimization problems. These formulations intrinsically offer interesting flexibility properties, preserve autonomy for the different disciplines involved in the problem, and thus are well-suited to meet industrial constraints.

2.2 Distributed MDO processes in GEMSEO: requirements and challenges

Following the above description of GEMSEO, the requirements for the distributed MDO services are as follows:

- As a key feature of GEMSEO is its reconfigurable aspect, based on a selected config-

uration, it is paramount to keep this feature for the distributed services. Practically, these distributed MDO services shall integrate seamlessly into an existing MDO process without hassle.

- These services shall be optional and not exclusive, and it shall be possible to associate them with a selected part of the process while being able to keep other parts locally.
- They shall implement natively the execution and linearization capabilities of GEMSEO and the standard GEMSEO caching mechanism.
- They shall handle back-and-forth file transfers if needed by the process.
- They shall be composable: it shall be possible to nest services together.
- They shall be simple to deploy and easy to use for both the client and the server. Boilerplate code should be kept to a minimum.
- They shall ensure the security of the data transmitted and, if relevant, implement role-based access control (RBAC).

This increase in scalability through the distribution of the MDO process and all these requirements does not come without challenges:

- It often involves horizontal scalability of resources, e.g. increasing the number of computing nodes and infrastructures involved in the process. Contrary to vertical scalability (e.g., increasing the computation power of a single node), horizontal scalability requires managing these resources, which are often heterogeneous in terms of hardware or software.
- It brings another layer of complexity for the MDO process orchestration, as it implies remote disciplinary computations, with data and file transfers to be done over the network, which can be unstable.
- Distant infrastructure may be reachable only by a specific protocol. For instance, SSH is often blocked in companies for obvious security reasons, and the only protocol available is HTTP(s), often through relays and proxies that need to be configured.
- Distributed MDO processes yield, in essence, scattered data (simulation results and artefacts, log files, ...) that are paramount for the end user to understand the results and fine-tune the process. The data management and the observability of such MDO processes, trivial in a single machine, become more complex in a distributed environment.

In light of all these requirements and challenges, it is evident that there is no *"one-size-fits-all"* service to handle all these capabilities. Existing solutions have been evaluated to interface with GEMSEO. Although they can be wrapped into GEMSEO, they do not meet the abovementioned requirements. Therefore, it has been chosen to develop a line of platform services for GEMSEO, either implemented in GEMSEO itself or as plugins. Each of these services enables fulfilling one part of the requirements described above.

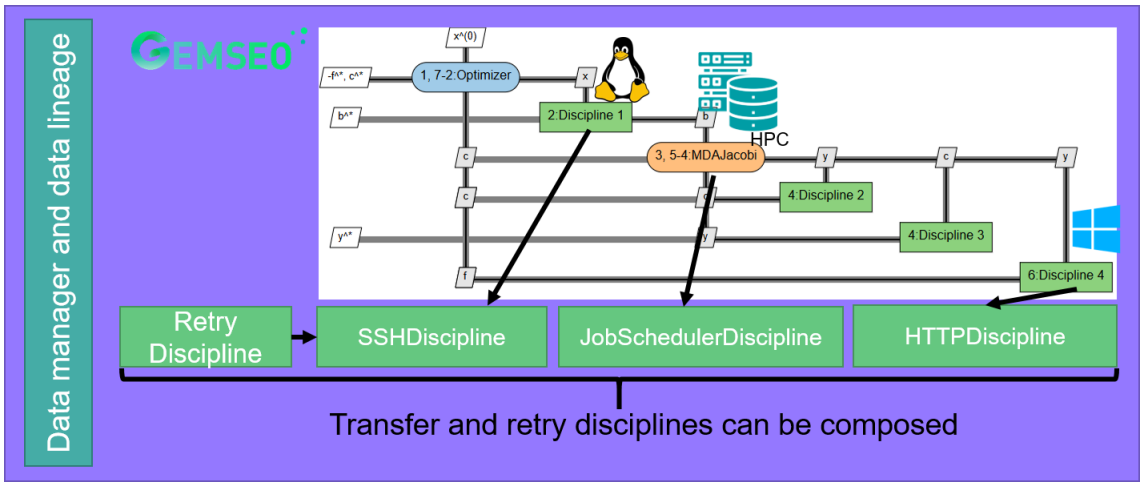


Figure 1: GEMSEO extensions for distributed MDO processes.

2.3 A line of GEMSEO extensions for distributed MDO processes

To fulfill the requirements and constraints described in Sec. 2.2, a line of GEMSEO transfer disciplines, instead of a single one, has been created (see Fig. 1):

- The **JobSchedulerDisciplineWrapper**: this discipline enables the delegation of the computation of an existing discipline instance to a Job Scheduler, typically used to manage HPC resources. Its detailed implementation can be found in 3.1.
- The **SSHDiscipline**: this discipline enables the delegation of the computation of an existing discipline to a distant node that is reachable through SSH. Implementation details can be found in Sec. 3.2.
- The **HTTPDiscipline**: this service brings a GEMSEO discipline, the **HTTPDiscipline**, a client to its server counterpart, an HTTP REST server that exposes GEMSEO disciplines as web services. Communications and data transfers are done through HTTP requests. A thorough description of both the client and the server is done in Sec. 3.3.

These three GEMSEO disciplines can be composed together; for instance, the **SSHDiscipline** can wrap a **JobSchedulerDisciplineWrapper** that enables it to have custom behavior. Also, these transfer disciplines can be applied to a selected part of the initial MDO process, for instance, to a sole discipline or an MDA, which is still a discipline in GEMSEO. It enables the user to apply transfer disciplines to specific parts of the process where there is an opportunity for delegation, while some disciplines may stay locally. It can also be stressed that other transfer disciplines could be added in the future to support other protocols, such as gRPC with Philote[18], RPyC[25], or RCE[19].

To add to these transfer disciplines, a **RetryDiscipline** has been developed to increase the resilience of distributed MDO scenarios in case of network or hardware errors. It enables the retry of the underlying process for specific exceptions. As for the previous transfer disciplines, the **RetryDiscipine** wraps another discipline and can be composed with transfer disciplines. Implementation details can be found in Sec. 4.

Finally, data lineage and data management capabilities have been investigated to improve traceability across the remote components of the MDO process. It provides a client

and a server, relying on the Openlineage[26] standard and implementation. The client part enables hooking easily into the process disciplines and instruments them to send data to the server. An implementation description is done in Sec.5.

3 GEMSEO TRANSFER DISCIPLINES

3.1 Job scheduler discipline

3.1.1 Motivation

Job scheduler software, such as LSF[27], SLURM[28], or OpenPBS[29], is often present in HPC premises to manage resources and ensure fair usage for all users. Yet, MDO processes usually run outside HPC and only delegate parts of the process to HPC jobs, such as high-fidelity disciplines or MDAs that need significant computing resources. In the frame of GEMSEO usage, users typically developed their *ad-hoc* solutions to access the HPC infrastructure, *e.g.* writing and maintaining their own logic in dedicated GEMSEO disciplines. Although this solution is compelling, the following drawbacks arise:

- A lack of genericity: *ad-hoc* solutions are not generic and may not work for another job scheduler.
- A non-standard use of GEMSEO: the implementation could not follow the GEMSEO patterns and may skip some intrinsic features such as the linearization of the disciplines.
- A maintenance burden that is only supported by the developer of the *ad-hoc* solutions, often leading sooner or later to deprecation.

Following these statements, the principle of the **JobSchedulerWrapperDiscipline** is to delegate the execution of a given discipline locally implemented to a Job Scheduler software, to be eventually run on an HPC premise. The logic is meant to be generic enough to handle a large spectrum of job scheduler implementations, such as LSF, SLURM or PBS. This feature has been integrated into GEMSEO 5.0 and higher, and supported by the GEMSEO development team.

3.1.2 Architecture and implementation

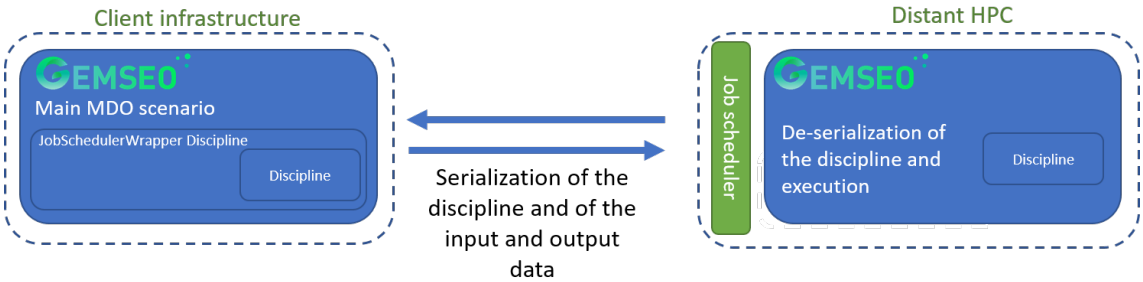


Figure 2: **JobSchedulerDisciplineWrapper** architecture.

Figure 2 gives a schematic view of using the **JobSchedulerDisciplineWrapperr** discipline. It is a discipline that contains the part of the MDO process (i.e. a GEMSEO

discipline) that will be offloaded and eventually executed on the HPC infrastructure. In terms of design, it implements a composite design pattern.

As it is not expected to rely on a distant service (except the HPC job scheduler and a distant GEMSEO installation) to run the discipline to be offloaded, it has been chosen to carry out the underlying discipline transfer through serialization using the standard pickle mechanism. It allows the serialization and storage of the discipline on the disk. It also enables its transfer to the HPC infrastructure, using either a shared folder mechanism or explicit remote data transfer using, for instance, the **SSHDiscipline** (see Sec . 3.2).

The need for genericity leads to the usage of Job Scheduler templates. Job Schedulers are typically used with running scripts, which can be very specific to the Job Scheduler software and the cluster's configuration. The trade-off found here is to provide the user with the option to give a template of job scheduler running scripts, where it can provide its static options.

In the end, the job scheduler script will run a GEMSEO command that will de-serialize the discipline from the disk and execute it. The output data will be serialized and stored on the disk, either from the direct computation or the linearization. Finally, the **JobSchedulerDisciplineWrapper** discipline will load the underlying discipline output data on disk and update its local state from that data.

It can also be noted that the serialization mechanism enables recovery of the state of the underlying discipline that has been remotely executed. It includes exceptions that may have been raised distantly, and that would be thrown back in the primary MDO process by the **JobSchedulerDisciplineWrapper** discipline. This serialization mechanism then enables the continuity of the process, including the propagation of errors and exceptions to the upper-level process.

3.1.3 Usage

```
from gemseo.wrappers.job_schedulers.schedulers_factory import
    SchedulersFactory
from gemseo import create_discipline, create_scenario, create_mda
from gemseo.problems.sellar.sellar_design_space import SellarDesignSpace
disciplines = create_discipline(['Sellar1', 'Sellar2', 'SellarSystem'])
mda = create_mda(disciplines)
wrapped_mda = wrap_discipline_in_job_scheduler(mda,
                                               scheduler_name="SLURM",
                                               workdir_path="workdir",
                                               cpus_per_task=24)
scenario = create_scenario(mda,
                           "DisciplinaryOpt",
                           "obj",
                           SellarDesignSpace(),
                           scenario_type="DOE")

scenario.execute(algo="lhs",
                 n_samples=100,
                 algo_options={"n_processes":10})
```

Listing 1: Usage of the **JobSchedulerDisciplineWrapper**

The example in the Listing 1 demonstrates the use of the **JobSchedulerWrapperDiscipline**. The wrapping uses a helper function and returns an instance of the **JobSched-**

ulerWrapperDiscipline. As, by inheritance, it complies with the interface of the GEMSEO **Discipline** class, this instance can be used in a GEMSEO process. Its execution will trigger the delegation of the computation of the underlying discipline, here an MDA of the well-known Sellar problem, to the HPC. The *wrapped_mda* is a GEMSEO discipline and can be integrated in a larger MDO process or used directly in a DoE or MDO scenario, as shown in the example.

3.2 SSH discipline

3.2.1 Motivation

Industrial and/or complex MDO workflows often need to run on more than one computational node. This may stem from the following reasons:

- Need for external resources, which could be software or hardware.
- Need to reach an external service, such as an HPC gateway.
- Need to run software on another operating system.
- Need to have more computational resources (CPU, RAM, GPU).

SSH is a protocol often used internally in private networks (universities, companies) and sometimes authorized for external enterprise communications. In case of blocked SSH communication, a natural fallback is the HTTP protocol, detailed in Sec. 3.3.

Similarly to the need to access Job schedulers (see Sec. 3.1), GEMSEO users used to implement *ad-hoc* solutions to handle this use-case. Yet, this solution lacks genericity and may not conform to the GEMSEO development best practices. It has then been developed as a generic open-source plugin, the *gemseo-ssh*[30] plugin, that aims to ease the delegation of the computation or linearization of a discipline to a distant node through an SSH channel. Its main features are:

- The distant execution of a GEMSEO discipline (that could be a process discipline such as a chain or an MDA) through an SSH connection.
- The possibility of cross-platform execution (*e.g.*, client and server operating on a different operating system).
- The possibility to transfer files from the client to the server, and the opposite.

3.2.2 Architecture and implementation

The software implementation and architecture of the **SSHDiscipline** are very close to the ones used for the **JobSchedulerDisciplineWrapper** discipline (see Sec. 3.1):

- The **SSHDiscipline**, at the top level, uses a composite design pattern and stores as an attribute the underlying discipline, whose execution will be delegated to the distant node through an SSH connection.
- The data flow, consisting of the transfer back and forth of the discipline, input and output data, is done through a serialization mechanism that is the same as for the **JobSchedulerDisciplineWrapper**.

- The discipline execution on the remote node is done through a GEMSEO entry-point that de-serializes the serialized discipline, executes it, and returns the results to the client.

The main difference is in the transfer protocol: the **SSHDiscipline** opens an SSH connection to the remote node accessible with an SSH connection, and carries out the complete execution or linearization of the underlying discipline. SSH connections are open in the plugin implementation thanks to the Python Paramiko library[31].

In addition, some use cases may need the transfer of files, either from the client to the server or vice versa. File transfers have been implemented in the *gemseo-ssh*[30] plugin, relying again on Paramiko and the sFTP protocol.

Finally, it has to be noted that this feature has been developed as an open-source plugin, as it needs new dependencies such as Paramiko. This feature then requires the installation of the plugin *gemseo-ssh*[30] in the same environment as the current GEMSEO installation, both on the client and on the server side. The installation an

3.2.3 Usage example

```
from gemseo import create_discipline
from gemseo import create_mda
from gemseo_ssh import wrap_discipline_with_ssh

disciplines = create_discipline(["SobieskiPropulsion",
                                 "SobieskiAerodynamics",
                                 "SobieskiMission",
                                 "SobieskiStructure"])
mda = create_mda("MDAChain", disciplines)
remote_discipline = wrap_discipline_with_ssh(
    discipline=mda,
    hostname="remote_hostname",
    local_workdir_path = ".",
    remote_workdir_path="~/test_ssh",
    username="my_username",
    pkey="path/to/my/private/ssh/key",
)

couplings = remote_discipline.execute()
```

Listing 2: Usage of the **SSHDiscipline**

The Listing 2 shows an example of using the **SSHDiscipline**. In the example, an MDA is created locally with the Sobieski disciplines and then sent to a distant node named *remote_hostname*. Some options related to the local and distant workdir shall be provided, as well as the username and the password. For the latter, it is also possible to provide a private key, assuming that the matching public key is on the server side, to avoid putting the password in the script. The remote execution is straightforward with the execution of the *remote_discipline*. The data transfer, back and forth, is automatically handled by the **SSHDiscipline** and the *remote_discipline* can be integrated *as-is* in a larger MDO process.

3.3 HTTP Discipline

3.3.1 Motivation

The main motivation of the **HTTPDiscipline**, as for the **SSHDiscipline**, is to be able to delegate the computation of a given discipline to a distant infrastructure. As described in Sec.3.2, some limitations may hamper the use of an SSH connection to offload the computation of a remote discipline:

- The distant node may not expose an SSH server.
- SSH connections and tunnels may be forbidden for security.
- The distant discipline provider does not want to expose the discipline logic and does not provide the discipline code.

It is also remarkable that the HTTP protocol is ubiquitous nowadays, and ports 80 and 443 are more prone to be open than port 22, which is related to an SSH connection. In light of this, exposing GEMSEO disciplines as a web application is attractive. The advantages are as follows:

- Web applications using REST API are a mature way to expose services and discipline computations and linearizations.
- The benefit of a large, mature and testable ecosystem of Python libraries (FastAPI, Pydantic, SQLAlchemy, SQLModel), task managers (Celery, Huey) and brokers (SQLite, Redis, RabbitMQ, ...)
- The capability to opacify the distant server code: the user only relies on a contract (endpoints and input/output data schemas) and does not have a view on the logic of the distant discipline. This mechanism is appealing in case of IP issues or if there is any willingness to hide some implementation details.
- The mechanism does not rely on the serialization of disciplines and enables the exposure of disciplines that cannot be serialized or would need a lot of development effort to make them serializable. An example is a GEMSEO discipline that wraps compiled code such as C, C++, Rust, Fortran or Julia libraries.

All these points motivated the development of another platform feature for GEMSEO: an HTTP server and client that enables seamless exposure, call, and linearization of a GEMSEO discipline through a web REST API and a web client.

3.3.2 Architecture and implementation

The open-source GEMSEO *gemseo-http*[32] plugin is composed of two elements:

- The server side exposes GEMSEO disciplines through a REST web application, enabling executing and linearizing them remotely. It needs to be exposed as a standard web service, as described in Fig.3.
- The client side is a GEMSEO discipline, the **HTTPDiscipline**, that enables hooking to a server exposing a target GEMSEO discipline and seamlessly configures itself based on the target discipline grammar.

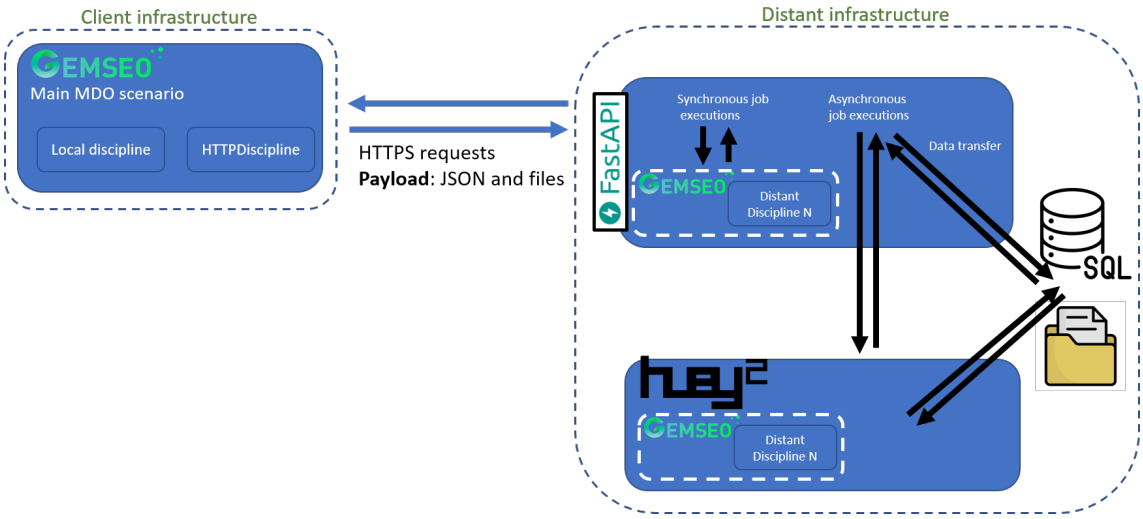


Figure 3: **HTTPDiscipine** client and server architecture.

Server-side On the server side, the FastAPI[33] web framework exposes different endpoints (GET/PUT/POST/DELETE HTTP methods) that enable the execution of the disciplines and their linearization. Some endpoints expose the grammar of the disciplines, which allows setting up the grammar of the **HTTPDiscipline** client. Also, as the **HTTPDiscipline** supports transferring files between the client and the server, some endpoints are dedicated to managing files. They enable the upload of some files to the distant server to be later used by the execution of a distant discipline, and also the possibility to download some files that may be the product of the execution of a discipline.

The server side relies on HTTP requests that are not intended to be long-lived connections. Depending on the expected duration time of the discipline execution and linearization, they can be done by choosing one of these two modes:

- Synchronous mode: if the execution and possibly the linearization of the discipline is not aimed to last more than the timeout limit of the web server, then the execution job can be done through a synchronous call: the output of the discipline execution will be returned at the end of the first job call.

- Asynchronous mode: if the execution and possibly the linearization of the discipline is aimed to last more than the web server timeout limit, then the execution can be done through an asynchronous call: the output of the discipline execution is not returned at the end of the initial execution call, but a job id. The job is queued and executed by the Huey[34] task manager. Input and output data are stored in a broker, which is in this case a SQLite database. The job id is later used to poll the API and check the job status. Depending on its status, the polling mechanism will continue or will exit. In the latter case, it will retrieve either the result of the execution and the linearization or the error in the form of a traceback raised by the distant code.

Finally, it has to be noted that the server enables nominative authentication through the classical OAuth2 scheme and JWT tokens. An SQL database allows storing the users and their hashed passwords, the inputs and outputs of the job run, and the file information transferred to the distant server. The server can be used on its own, e.g., by

making HTTP requests on the server using whatever client language is used. However, the common usage is to use the **HTTPDiscipline** client in a GEMSEO process to reach the server. This enables seamless discipline integration by delegating the computations and/or linearizations to a distant GEMSEO HTTP server.

Client-side The **HTTPDiscipline** is a GEMSEO discipline, which delegates its executions to a distant GEMSEO HTTP server. Its interest lies in its capacity to auto-configure itself based on information such as the server distant URL, the user login and password or an authentication token, and the name of the distant discipline to instantiate. Based on this set of information, the **HTTPDiscipline** configures its input and output grammar schema by requesting those of the distant requested library.

Then, when executed, the **HTTPDiscipline** delegates its execution to the distant server with either a synchronous or asynchronous execution mechanism. It will also automatically upload the files needed for the distant execution. In an asynchronous execution, the **HTTPDiscipline** will poll the server, using a standard or long-polling strategy. When finished, the client will synchronize its state with the results obtained from the execution. In a distant failure, the traceback will be retrieved and the exception will be thrown again. Finally, the output files will be retrieved locally from the server.

3.3.3 Usage

```python
from pathlib import Path
from numpy import array
from gemseo_http.app.job_creation import ExecutionMode
from gemseo_http.disciplines.http_discipline import HTTPDiscipline

discipline = HTTPDiscipline(
    name="RemoteSellar2",
    distant_discipline_name="Sellar2",
    service_base_url="http://my_distant_server",
    port=8000,
    user="my_user",
    password="my_password",
    execution_mode=ExecutionMode.asynchronous,
)
data = {"x_shared": array([1.0, 1.0]), "y_1": 1.0}
data = discipline.execute(data)
```

Listing 3: Usage of the **HTTPDiscipline**

The Listing 3 highlights an example of the **HTTPDiscipline**. As described previously, its execution needs the presence and execution of a *gemseo-http* server, which can be reached on a specific IP or URL, on a specific port. In this example, the remote execution of the **Sellar2** discipline that is bundled in any GEMSEO installation is demonstrated. The execution of the instance of the **HTTPDiscipline** will delegate through HTTP requests the execution to the discipline distantly. Moreover, in this example, the execution is asynchronous, meaning that the **Sellar2** discipline will be executed in the background, and the server will poll the distant server to query the state of the discipline execution. Once successfully executed, the results will be retrieved from the server to the client discipline, and eventually to the client MDO process.

3.3.4 Limitations

The GEMSEO HTTP server and its client counterpart, the **HTTPDiscipline**, have limitations:

- An HTTP server based on a REST architecture is stateless: a stateless protocol is used between the client and the server. This implies that each request to the server should not suppose any specific state: the server's state must be recovered by using the information passed by the client and possibly requests to a database. Although very convenient for horizontal scalability (e.g., the number of server instances can be increased with linear scalability), it also implies that a state must be re-created at each API call. A caching mechanism could alleviate this issue, but that is sometimes not possible in the case of non-serializable objects. It yields a loss of performance with respect to its SSH counterpart.
- A GEMSEO HTTP server needs to be exposed so that some clients can reach it. Although the HTTP(S) protocol is ubiquitous nowadays, opening some ports or passing the traffic through an HTTP proxy can be necessary in some specific cases.

4 ERROR HANDLING IN DISTRIBUTED MDO PROCESSES

4.1 The need for specific error-handling logic

Distributed MDO processes do not come without drawbacks. In particular, they rely on resources that may not be reliable, such as the network, the local and distant disk drives, services, or, more generally, any resource that is not controlled by the MDO user or integrator. As MDO processes often rely on iterative processes such as optimization algorithms, non-deterministic errors could break the entire process and prevent a normal MDO process execution. It becomes then evident that transfer disciplines need an extra and orthogonal mechanism to the transfer disciplines to ensure the correct handling of errors:

- in the presence of deterministic errors, such as non-computable points, the MDO process should end gracefully, and the error should be propagated upstream to enable the user to identify and correct the issue.
- in the presence of non-deterministic errors, such as network instabilities, some restarts shall be allowed to make the MDO process more robust and enable it to come to its normal term.

4.2 Architecture and implementation

The error handling has been implemented in the **RetryDiscipline**. This discipline leverages the composite design pattern by enclosing the part of the process that can be unreliable. In the underlying process, errors shall be handled by plain Python exceptions. If not caught in the underlying discipline logic, a raised exception will propagate up to the **RetryDiscipline**. The latter can be configured upfront to determine which exceptions shall be handled or shall be propagated:

- If the caught exception has been set as a parameter of the **RetryDiscipline**, then the exception will not be propagated and will be handled by an execution retry.

The number of retries and the overall retry timeout are also parameters of the **RetryDiscipline** that can be set.

- Otherwise, the exception is not handled by the **RetryDiscipline**. It will be propagated through the call stack and eventually to the program entry point if not handled in between by another logic.

For the latter occurrence, it has to be stressed that the exception could be handled by the upstream algorithm, which may handle non-computable points. This point emphasizes the need to handle the proper errors by Python exceptions in the right place for proper usage. The **RetryDiscipline** mechanism is integrated as a GEMSEO core feature and is available in GEMSEO 6.1 and higher.

4.3 Usage

```python
class FictiveDiscipline(Discipline):
    """Discipline to be executed several times.

    - The first 2 times, raise a RuntimeError,
    - and finally, succeed.
    """

    def __init__(self) -> None:
        super().__init__()
        self.attempt = 0

    def _run(self, input_data: StrKeyMapping) -> StrKeyMapping:
        self.attempt += 1
        LOGGER.info("attempt: %s", self.attempt)
        if self.attempt < 3:
            raise RuntimeError
        return {}

test_n_trials = FictiveDiscipline()
retry_disc = RetryDiscipline(test_n_trials, n_trials=4)

retry_disc.execute()

#    INFO - 11:26:38: attempt: 1
#    INFO - 11:26:38: attempt: 2
#    INFO - 11:26:38: attempt: 3
#
# {}
```

Listing 4: Usage of the **RetryDiscipline**

The Listing 4 shows a practical usage of the **RetryDiscipline**. Firstly, a **FictiveDiscipline** is defined as a mock of an unstable service. It will fail the first two execution attempts and will succeed at the third execution call. In the example, the discipline is instantiated and wrapped in a **RetryDiscipline**, with a maximum number of trials set to 4. Then, the execution of the **RetryDiscipline** instance shows that the first two attempts fail, while the third one is successful.

5 MDO PROCESS OBSERVABILITY AND DATA MANAGEMENT

5.1 Motivation

As discussed in Sect. 2.2, large and distributed MDO processes may need observability and data management capabilities:

- Process observability enables the traceability of discipline executions and associated data generated during the execution of the MDO workflow. A service allows the end-user to observe the data generated by the overall process (which may be distributed). Observing the process can be done during the process execution or after the execution is finished.
- A data manager capability enables the storage of process data in a centralized storage (local filer, network-shared filer, S3-like object storage service). Associated with the observability mechanism, it enables users to retrieve files of interest without directly accessing the hardware resource used to generate them.

The current work focuses on the observability capability, which can be considered a prerequisite for the second item, which has not yet been addressed and can be considered as a perspective.

5.2 Architecture and implementation

A survey of the current state of the art of open source observability and data lineage solutions has been conducted. Two satisfying solutions have been identified: OpenMetaData[35], which provides data lineage among other features, and OpenLineage[26], which provides a platform standard, as well as implementations of them, both for the client and server parts.

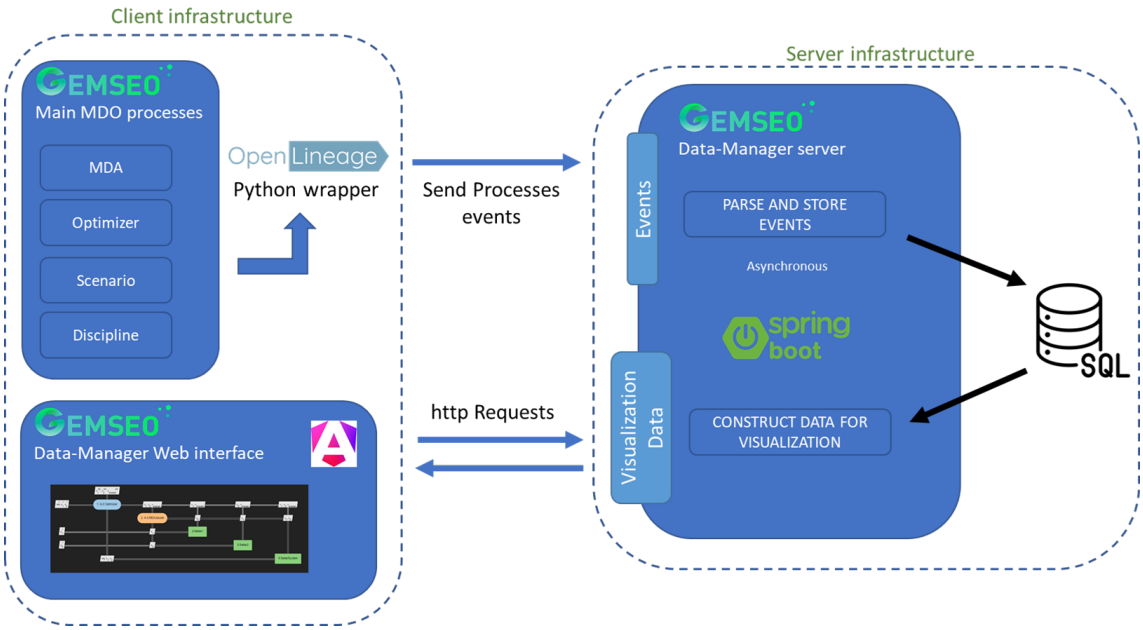


Figure 4: Architecture of the observability and data management services.

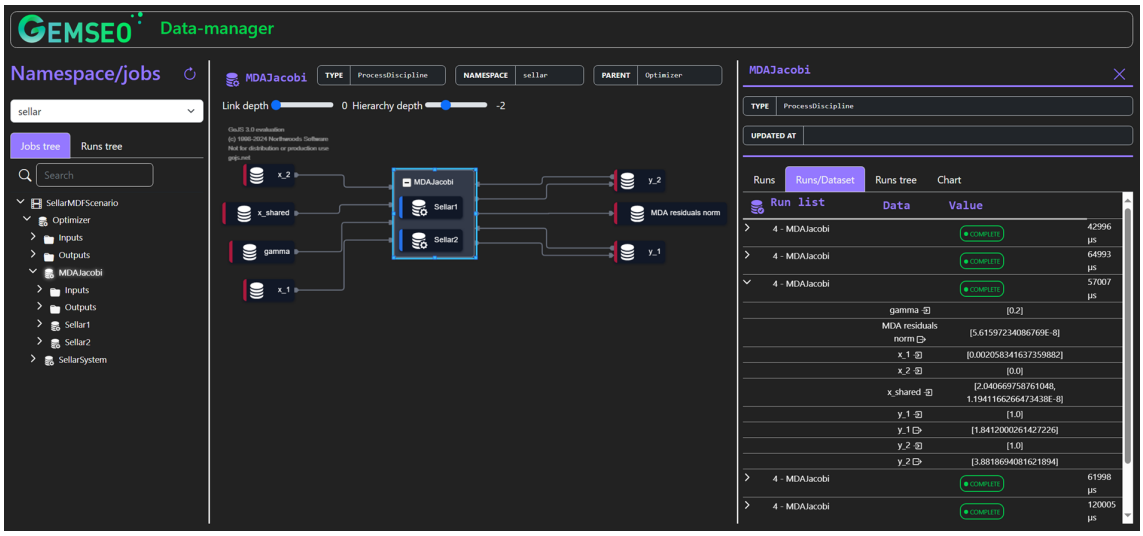


Figure 5: Graphical user interface of the data observer.

To reduce development costs, an existing solution has been selected and OpenLineage has been initially chosen as a lighter alternative to OpenMetaData. It provides a standard for data lineage, flexibility in the metadata definition, and an implementation of a Python client. A GEMSEO plugin allows to hook MDO processes with this client.

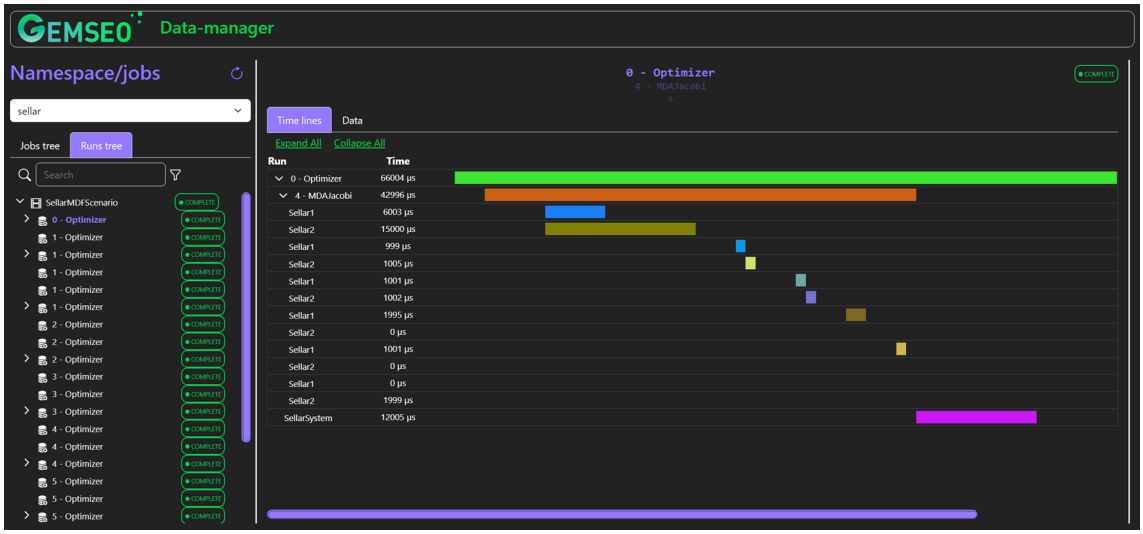


Figure 6: Gantt chart view of an MDO process.

Marquez[36], a server-side implementation of the OpenLineage standard, has also been initially chosen, as it provides a web-based graphical user interface. Yet, it has been quickly found that Marquez shows undesirable behaviors:

- It does not allow traceability of the inputs.
- It does not fully handle hierarchical relationships.
- It has performance latencies which were not circumvented by using asynchronous requests on the client side.

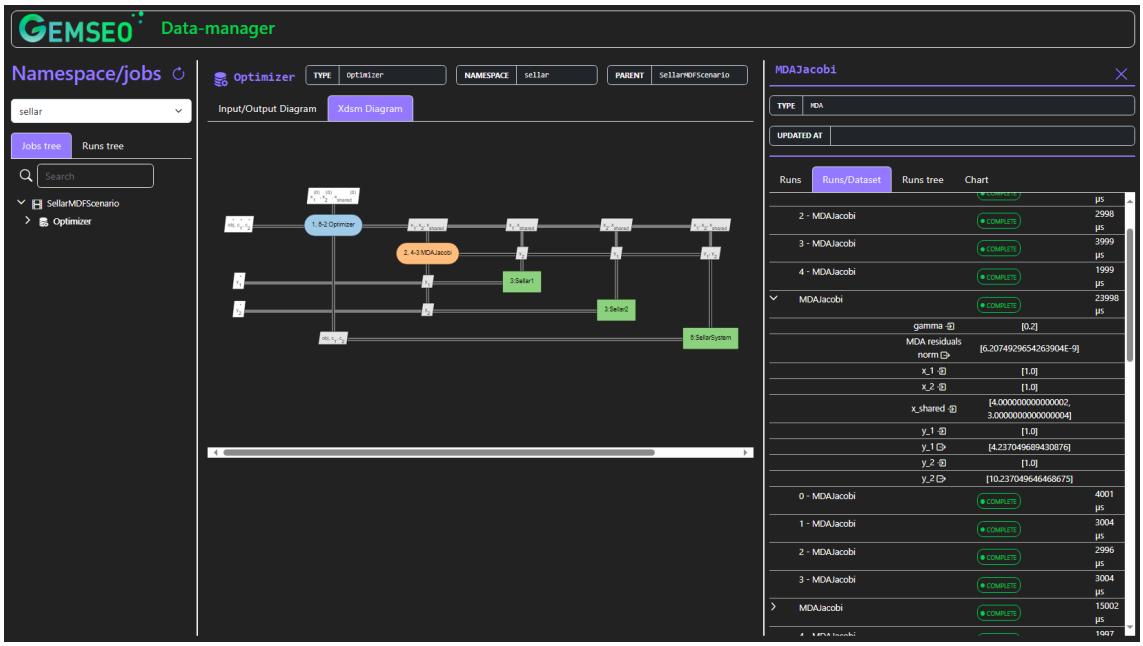


Figure 7: XDSM view of an MDO process.

As a result, it has been decided to re-implement the server and the database schema in a Java Spring Boot web service, while keeping the interface proposed by Marquez to comply with the OpenLineage standard. A schematic view of the services architecture is given in Fig. 4.

Also, the web-based user interface was found to be too generic and did not provide the specific views that MDO users may expect, such as a Gantt chart view or an XDSM view of the process. A prototype of a graphical user interface has been re-implemented in JavaScript using the Angular framework. It queries the server, enabling the user to observe the overall process and the discipline executions, as well as the metadata associated with them. It currently provides the following capabilities:

- A global view of the process, using a tree-view or a block-based view (See Fig. 5).
- An XDSM view (See Fig. 7), allowing to interactively view each execution of the disciplines.
- A view of the iterations of iterative processes, such as optimizers and MDAs, and a view of each evaluation of the disciplines for the current iteration.
- A Gantt chart view of the MDO process (see Fig.6), to identify hot spots in the MDO process.

This back-end is still in development, both in terms of features and maturity. Also, the integration of data observability with the aforementioned transfer disciplines has not yet been tested.

6 CONCLUSIONS AND PERSPECTIVES

The current work has presented recent developments in and around GEMSEO to enable distributed MDO processes and ease their integration. These developments are done

through transfer disciplines that enable delegating the computation of disciplines distantly and mechanisms to robustify and observe distributed MDO processes. This work follows the guidelines for developing and using the GEMSEO library, which considers the industrial constraints of MDO processes and provides generic and validated solutions.

Some perspectives have been identified that can be used to carry on this work. The first would be the observability and data manager service completion and robustification. Also, new transfer disciplines could be added. For instance, a transfer discipline targeting a Philote server would enable interoperability between GEMSEO MDO processes and OpenMDAO disciplines, but other protocols, such as RPyC, may be targeted. Higher-level GEMSEO objects, such as MDO Scenarios, could also be delegated. Another development perspective would be to push toward cloud-centric services and mechanisms for disciplines and overall MDO scenario executions. These are becoming increasingly ubiquitous and central to the "move-to-cloud" trend observed for scientific applications.

ACKNOWLEDGMENTS

This study has been supported by the NEXTAIR and RECET4RAIL projects, which have received funding from the European Union's Horizon Europe research and innovation program under grant agreement No 101056732 and 101015423, respectively. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

REFERENCES

[1] A. Gazaix, F. Gallard, V. Ambert, D. Guénot, M. Hamadi, S. Grihon, P. Sarouille, T. Y. Druot, J. Brézillon, V. Gachelin, et al. Industrial application of an advanced bi-level mdo formulation to aircraft engine pylon optimization. In *AIAA Aviation 2019 Forum*, page 3109, 2019.

[2] S. Kaneko and J. R. Martins. Mdo formulations for simultaneous design and trajectory optimization. In *AIAA SCITECH 2024 Forum*, page 2231, 2024.

[3] T. Ashuri, M. B. Zaaijer, J. R. Martins, G. J. Van Bussel, and G. A. Van Kuik. Multidisciplinary design optimization of offshore wind turbines for minimum levelized cost of energy. *Renewable energy*, 68:893–905, 2014.

[4] Ansys. ModelCenter. URL `https://www.ansys.com/products/connect/ansys-modelcenter`.

[5] ESTECO. ModeFrontier. URL `https://engineering.esteco.com/modefrontier/`.

[6] Dassault Systèmes. SIMULIA ISight. URL `https://www.3ds.com/products/simulia/isight`.

[7] pSeven. pSeven. URL `https://www.pseven.io/product/pseven/`.

[8] Noesis. Optimus. URL `https://www.noesissolutions.com/our-products/optimus`.

[9] Siemens Digital Industries Software. SimCenter HEEDS. URL https://plm.sw.siemens.com/en-US/simcenter/integration-solutions/heeds/.

[10] J. S. Gray, J. T. Hwang, J. R. Martins, K. T. Moore, and B. A. Naylor. Openmdao: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, 2019.

[11] F. Gallard, C. Vanaret, D. Guénot, V. Gachelin, R. Lafage, B. Pauwels, P.-J. Barjhoux, and A. Gazaix. Gems: a python library for automation of multidisciplinary design optimization process generation. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, page 0657, 2018.

[12] É. Lac, G. De Spiegeleer, A. Delsalle, F. Collonval, D.-T. Lê, and M. Malandain. Cosapp: a python library to create, simulate and design complex systems. *Journal of Open Source Software*, 9(94):6292, 2024.

[13] A. Morris, H. Syamsudin, J. Fielding, M. Guenov, K. Payne, P. Deasley, S. Evans, and J. Thorne. *MACRO - A tool to support distributed MDO*. AIAA. doi:10.2514/6.2000-4702. URL https://arc.aiaa.org/doi/abs/10.2514/6.2000-4702.

[14] R. Ratcliff, S. LeDoux, and W. Herling. A modern corba-based approach to ad hoc distributed process orchestrations applied to mdo. In *Infotech@ Aerospace*, page 7143. AIAA, 2005.

[15] B. Chen, D. Liu, B. Mahdavi, Q. Zhou, D. Bouhemhem, A. Ndiaye, F. Guibault, B. Ozell, D. Pelletier, and J.-Y. Trepanier. A data-centric distributed framework for mdo management. In *Proceedings of the Sixth International Conference on Computer Supported Cooperative Work in Design (IEEE Cat. No. 01EX472)*, pages 279–284. IEEE, 2001.

[16] Y. D. Wang, W. Shen, and H. Ghenniwa. Webblow: a web/agent-based multidisciplinary design optimization environment. *Computers in Industry*, 52(1):17–28, 2003. ISSN 0166-3615. doi:https://doi.org/10.1016/S0166-3615(03)00066-6. URL https://www.sciencedirect.com/science/article/pii/S0166361503000666. Knowledge Sharing in Collaborative Design Environments.

[17] S. LeDoux, W. Herling, G. Fatta, and R. Ratcliff. Mdopt-a multidisciplinary design optimization system using higher order analysis codes. In *10th AIAA/ISSMO multidisciplinary analysis and optimization conference*, page 4567, 2004.

[18] C. A. Lupp and A. Xu. Creating a Universal Communication Standard to Enable Heterogeneous Multidisciplinary Design Optimization. In *AIAA SCITECH 2024 Forum*, page 1799, 2024.

[19] B. Boden, J. Flink, R. Mischke, K. Schaffert, A. Weinert, A. Wohlan, C. Ilic, T. Wunderlich, C. M. Liersch, S. Goertz, et al. Distributed multidisciplinary optimization and collaborative process development using rce. In *AIAA aviation 2019 forum*, page 2989, 2019.

[20] A. Kumar, J. Kolhe, S. Ghemawat, and L. Ryan. gRPC Protocol. Internet-Draft draft-kumar-rtgwg-grpc-protocol-00, Internet Engineering Task Force, July 2016. URL `https://datatracker.ietf.org/doc/draft-kumar-rtgwg-grpc-protocol/00/`. Work in Progress.

[21] P. D. Ciampa, E. Moerland, D. Seider, E. Baalbergen, R. Lombardi, and R. D'Ippolito. A collaborative architecture supporting agile design of complex aeronautics products. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, page 4138, 2017.

[22] J. H. Bussemaker, P. D. Ciampa, J. Singh, M. Fioriti, C. Cabaleiro De La Hoz, Z. Wang, D. Peeters, P. Hansmann, P. Della Vecchia, and M. Mandorino. Collaborative design of a business jet family using the agile 4.0 mbse environment. In *AIAA Aviation 2022 Forum*, page 3934, 2022.

[23] J. Brezillon, I. Fadli, K. Thanissaranon, and C. Conlan-Smith. High-fidelity multidisciplinary optimisation processes for solving aero-structure problems: recent developments from an industrial perspective. In *Proceedings of Aerobest 2025*. ECCOMAS.

[24] A. Lambe and J. Martins. A unified description of mdo architectures. In *9th world congress on structural and multidisciplinary optimization*, volume 124, 2011.

[25] RPyC. URL `https://rpyc.readthedocs.io/en/latest/`.

[26] Openlineage, . URL `https://openlineage.io/`.

[27] IBM. Spectrum LSF. URL `https://www.ibm.com/fr-fr/products/hpc-workload-management`.

[28] SchedMD. SLURM. URL `https://www.schedmd.com/`.

[29] OpenPBS, . URL `https://www.openpbs.org/`.

[30] gemseo-ssh, . URL `https://gitlab.com/gemseo/dev/gemseo-ssh`.

[31] Paramiko. URL `https://www.paramiko.org/`.

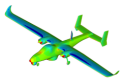
[32] gemseo-http, . URL `https://gitlab.com/gemseo/dev/gemseo-http`.

[33] FastAPI. URL `https://fastapi.tiangolo.com/`.

[34] Huey. URL `https://huey.readthedocs.io/en/latest/`.

[35] OpenMetadata, . URL `https://open-metadata.org/`.

[36] Marquez. URL `https://marquezproject.ai/`.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

AEROACOUSTIC SENSITIVITY ANALYSIS OF GEOMETRICAL PARAMETERS FOR OPEN-CAVITY FLOWS

Ramazan Kaba[1*] and Melike Nikbay[2]

1: Modelling and Simulation Department
Turkish Aerospace
Istanbul, Türkiye
ramazan.kaba@tai.com.tr
2.: Astronautical Engineering Department
Istanbul Technical University
University
34469, Istanbul, Türkiye
nikbay@itu.edu.tr

Abstract. *Open cavity flows commonly encountered in aerospace applications such as weapon housing areas and compartments for landing gear. These flows can generate a substantial amount of noise and may reduce the aircraft's stability, especially at high speeds near the speed of sound, known as the transonic regime. This study examines how changes in the design can affect the noise produced by a specific type of open cavity, named the M219 generic open cavity (with a length-to-depth ratio of L/D=5), when flying at a speed of Mach 0.85. Researchers employed computational methods, specifically using Computational Fluid Dynamics (CFD) simulations with the Spalart-Allmaras turbulence model on a simpler mesh in the OpenFOAM-HISA solver, to test 100 different shape variations. They systematically modified aspects such as the slant angle of the rear wall (θ), and the length (l), angle (β), and thickness (h) of a baffle at the leading edge to determine their impact on noise levels, measured by the Overall Sound Pressure Level (OASPL). To efficiently predict how these geometric changes affect noise, a mathematical model known as a Kriging surrogate model was constructed. This approach allowed researchers to understand the relationship between shape changes and noise without testing each one comprehensively. A Sobol' sensitivity analysis was conducted to identify both individual and combined effects of these shape modifications on noise levels. The results revealed that it is possible to significantly reduce noise, by about 30 decibels, mainly by adjusting the baffle angle positively. The analysis highlighted the baffle angle (β) as the key factor most influencing noise levels. Although the baffle length (l) had little direct effect by itself, it had a notable impact when combined with the baffle angle. The impact of the slant angle of the rear wall (θ) and the baffle thickness (h) was minor within the tested ranges. This study provides useful insights into achieving quieter aircraft designs by focusing on these critical shape adjustments.*

Keywords: Cavity flow, aero-acoustic, sensitivity analysis, surrogate modeling

1 INTRODUCTION

Cavity flow is a fluid mechanics problem that engineers in many different fields may encounter. In aviation, weapon (store) and main landing gear bays can be represented by this problem. Recent developments have led to requirements such as low radar visibility and supersonic/transonic flight continuity in fighter aircraft. These requirements make designers use embedded weapon bay configurations. Thus, operations, especially in the transonic speed regime, bring about a cavity problem. The local flow field due to cavity flow leads to the following undesirable conditions.

- During store separation, aero elastic damage may occur in the surrounding structures (pylons, wings, etc.) due to periodic and random oscillations in the bay (cavity).
- High noise levels generated inside the cavity cause acoustic fatigue phenomena and avionic equipment damage.
- The flow distorted by the cavity can create vertical pressure distributions. This, negatively affects the trajectory of separated store such as weapon and external fuel tank and can cause them to hit the aircraft.

Due to all these potential drawbacks, the cavity has become a critical design issue and it is necessary to analyze and optimize with high accuracy.

1.1 Cavity Flow Physics

The flow through the cavity consists of two main parts: the shear layer and the vortex region(s). The shear layer that starts to form at the leading edge of the cavity separates the flow into inner and outer parts. Generally, rectangular cavities can be classified as open, closed and transitional regions depending on the length/depth ratio (L/D). If the L/D ratio is more than 13 or 14, it is called a closed cavity, and if it is less than 10, it is called an open cavity. A value between these is called a transitional cavity (Figure 1).

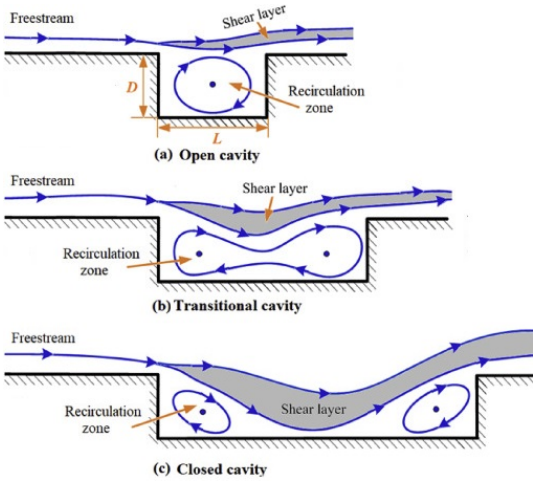


Figure 1: The types of flow through a cavity [1].

In open cavities, which have a simpler flow compared to transitional and closed cavities, the free stream flow separates from the leading edge and forms a shear layer. The shear layer spans through cavity and the separated flow can reattach to the rear wall. With the increase of the L/D ratio, the energy that transmits the fluid up to the aft wall cannot be provided. Therefore, in the transitional cavity, the shear layer partially enters inward, while in the closed cavity it completely enters. As a result, two separate vortex regions are formed (Figure 2).

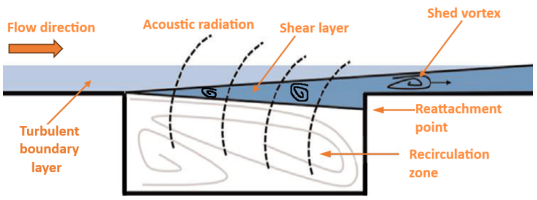


Figure 2: Schematic view of the cavity flow field oscillation mechanism.

The pressure fluctuations resulting from the periodic repeated shock waves formed by the separation of the flow that impinges on the cavity create a time-dependent situation called the feedback mechanism. The pressure oscillations resulting from this mechanism make the reattachment region the source of acoustic radiation.

The practical effects of these flow properties are shaped according to the pressure distribution on the lower wall. An almost homogeneous static pressure distribution occurs in the open cavity. This provides a relatively safer environment for stores. On the other hand, in a closed cavity, low pressure occurs in the regions close to the front walls and high pressure occurs in the regions close to rear walls. This pressure gradient causes a pitching moment on the store.

1.2 Flow control studies

Various flow control methods have been developed to minimize the damage caused by cavity-induced noise. The methods are generally aimed at eliminating the mechanisms that are the source of the noise (flow separation, shear layer, vortex regions, etc.). Flow control methods can be grouped under two main headings: passive and active. Active control requires an additional structural element to manipulate the flow or change the cavity geometry. These applications include sloping the cavity walls, adding spoilers or fixed baffle plates. Active methods involve an external energy input to the flow. The energy input can be achieved by processes such as oscillating flaps, mass injection and jet blowing.

Rossiter [2], Nightingale [3] and Ross [4] have conducted comprehensive wind tunnel experiments on the performance of spoilers as passive flow control devices. In summary, the following modifications were applied in the experiments. [5] [6] [7] [8] [9] [10] [11] [12] [13].

- Transverse rods on leading edge (Fig. 3a).
- Saw-tooth spoiler on leading edge (Fig. 3b).
- Pins steps or angled baffles on leading edge (Fig. 3c).

- Porous spoiler (Fig. 3d).
- Chamfer modifications on trailing edge (Fig. 3e).
- Slanted aft or front wall (Fig. 3f).

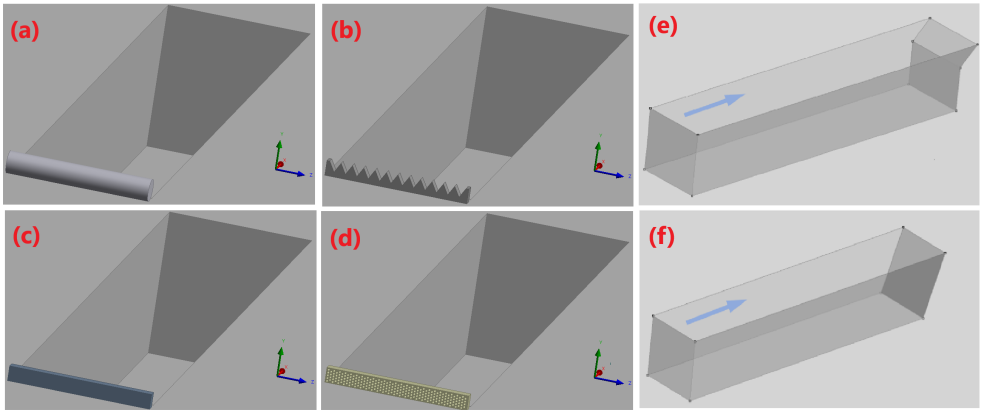


Figure 3: Passive flow control on cavity. a) Transverse rod, b) Saw-tooth spoiler, c) Flaptop spoiler, d) Porous spoiler, e) Fillet or chamfer, f) Slanted

As experimental studies have shown, spoiler applications generally reduce the broadband low-amplitude noise in the transonic regime, but are not effective at high amplitudes. It has been an effective method at all broadband sound amplitudes in supersonic speed regimes. Different locations and sizes of rod-shaped spoilers have been conducted, and noise reduction has been provided in all cases. The cavity leading edge is the most convenient position in terms of aero-acoustics. It has been observed that aero-acoustic recovery is maximized when the spoiler diameter is 30-40% of the boundary layer thickness. Ross and Peto [4] and Heller and Bliss [14] have shown that chamfering and slanting operations on the trailing edge distort the feedback mechanism. In addition, the 48° slope of the rear wall is the most sufficient to reduce noise [3].

Research on baffle applications has increased in recent years. A baffle study [15] showed that long plates (0.3L) create self-oscillating frequencies as the flow area approached the trailing edge. It was observed that the greater slope of the rear wall provided evacuation and dissipation of flow turbulence energy. Overall, the trailing edge modification, the 0.25L sloped plate and jet blowing to the trailing and leading edges, provided a reduction in acoustic tone suppression of approximately 27.5%. A similar study was conducted [16] with using Large Eddy Simulations and Lighthill's analogy to predict the noise induced by the flow of a turbulent boundary layer over an open cavity (L/D=1). As the lid length increases, the turbulence intensity inside the cavity is suppressed. The acoustic sources are concentrated at the downstream edge. In another study [17], flow-generated noise was calculated using the hybrid DES-Curle method. The results show that the change in the shape of the cavity edges has an effect on the sound generation. Change of downstream edge from perpendicular to rounded allowed to reduce the sound pressure level by approx. 10 dB.

2 THEORETICAL BACKGROUND

2.1 Measurement Methods of Experimental Time-Dependent Data

Sound is an energy generated by vibrations and transmitted by pressure alterations. Damaging (high) levels of sound is called noise. While noise can be measured accurately in time-dependent quantities such as speed, acceleration, and pressure, there is a significant loss of useful information in the time-dependent measurement of vibration signals. Since this loss occurs at certain frequencies, the location of the vibration can be narrowed down or amplitude can be changed. Fast Fourier Transform (FFT) is a mathematical calculation that converts a time-domain waveform into a series of discrete sine waves in the frequency domain. For time-dependent pressure data, the FFT can be expressed as follows;

$$p(f) = \int_{-\infty}^{+\infty} p(t) e^{-i(2\pi f)t}\, dt \tag{1}$$

FFTs are efficient in analyzing vibration when there is a limited number of dominant frequency components; however, power spectral densities (PSD) are used to characterize random vibration signals. A typical PSD function displays the power of the measured quantities as a function of frequency. In the context of this aero-acoustic study, the quantity is pressure. Therefore, PSD can be expressed mathematically as follows, expressed as Pa^2/Hz as a unit;

$$PSD(f_n) = \frac{1}{F_s N} |p(f_n)|^2 \tag{2}$$

The terms here are as follows:
p(t): Time-dependent pressure signal,
p(f): Fourier transform of time-dependent pressure signal,
$p(f_n)$: Frequency series of time-dependent pressure signal obtained by FFT,
F_s: Sampling frequency,
N: Data (sample) number.

Sound level may be perceived differently by different people. Therefore, there is a requirement for a means to obtain an objective measurement of sound level expressed in numerical terms. This is defined as Sound Pressure Level (SPL). Sound pressure (p) is the average change in atmospheric pressure caused by sound. SPL can be represented mathematically as follows.($P_{ref} = 20\mu Pa = 2.10^{-5} Pa$)

$$SPL = 10 \left[\log \left(\frac{PSD}{p_{ref}^2} \right) \right] \tag{3}$$

$$OASPL = 20 \left[\log \left(\frac{p_{rms}^{'}}{p_{ref}} \right) \right] \tag{4}$$

$$p_{rms}^{'} = \sqrt{(p'^2)_{average}} \tag{5}$$

In general, SPL is the equivalent SPL obtained by integrating all frequencies. OASPL is the logarithmic representation of all discrete spectral sound pressure levels into a single number.

2.2 Numerical Methods

One of the primary challenges in computational aero-acoustics simulations is the scale problem. Although turbulence can be solved with the applied CFD method, acoustic fluctuations smaller than turbulence cannot be solved completely. In order to obtain the smallest vortex regions numerically, mesh models with a large number of cells are needed. To overcome this requirement, the conservation equations are modified according to the Reynolds average. However, the modified equations add additional unknowns to the problem, and additional turbulence models are required to find these unknowns. Currently, there is no single turbulence model that can solve all turbulent-flow problems. The appropriate model is selected according to the available computing power, time, desired accuracy, and flow physics.

Acoustic analyses have high temporal and spatial requirements. In addition, sensitivity or optimization studies that require a large number of analyses with an iterative approach are very costly in terms of time and resource. Especially in optimization, the data that will feed the surrogate model is obtained with high and low fidelity. Low fidelity analyses are usually performed with potential flow assumption. In some cases (for example airfoil design studies), panel methods are applied. However, in studies based on an acoustic, the simplifications mentioned above cannot be applied. The noise obtained as a result of the disturbances caused by very small pressure fluctuations in the regions close to the wall cannot be resolved with a inviscid assumption. On the other hand, the panel method cannot be applied in open geometries such as the cavity used in this study. As a result, Spalart Allmalas, which can achieve a certain level accuracy but is a relatively low-cost model, was selected.

The Spalart-Allmalas single equation method is a less costly method compared to methods such as DES, LES in terms of time and solver power requirements. In the model, the velocity measure is determined by the transport equation of the kinematic turbulence viscosity variable, while the length measure is determined by algebraic equations. The S-A turbulence model follows a single dynamic equation that describes a kinematic viscosity-like variable called the S-A variable:

$$\begin{aligned}\frac{\partial(\overline{v})}{\partial t}+\frac{\partial(\overline{v})}{\partial x_j}=c_{b1}\left(1-f_{t2}\right)\overline{S}\overline{v}-\left[c_{w1}f_w-\frac{c_{b1}}{\kappa^2}f_{t2}\right]\left(\frac{\overline{v}}{d}\right)^2+\\ \frac{1}{\sigma}\left[\frac{\partial}{\partial x_j}\left(\left(v+\overline{v}\right)\frac{\partial\overline{v}}{\partial x_j}\right)+c_{b2}\frac{\partial\overline{v}}{\partial x_i}\frac{\partial\overline{v}}{\partial x_i}\right]\end{aligned} \tag{6}$$

The purpose of this equation is to determine the turbulent eddy viscosity, which is determined by the following formula:

$$\begin{aligned} \mu &= \rho v f_{v1} \\ f_{v1} &= \frac{\chi^3}{\chi^3 + c_{v1}^3} \\ \chi &= \frac{\overline{v}}{v} \end{aligned} \tag{7}$$

In these equations, molecular kinematic viscosity is $v = \mu/\rho$, fluid density is ρ and μ is the molecular dynamic viscosity of the fluid. Other parameters in the model are given in terms of the distance from the cell center to the nearest wall d and the vortex size.

The purpose of this equation is to determine the turbulent eddy viscosity, which is determined by the following formula:

$$W_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right) \tag{8}$$

The definitions in Equations 2.53, 2.54 and 2.55 are as follows:

$$\begin{aligned} \overline{S} &= \Omega + \frac{\overline{v}}{\kappa^2 d^2} f_{v2} \\ f_{v2} &= 1 - \frac{\chi}{1 + \chi f_{v1}} \\ f_w &= g\left[\frac{1 + c_{w3}^6}{g^6 + c_{w3}^6}\right] \\ g &= r + c_{w2}\left(r^6 - r\right) \\ r &= min\left[\frac{\overline{v}}{\overline{S}\kappa^2 d^2}, 10\right] \end{aligned} \tag{9}$$

Many constants can be seen in the definitions above. These constants are defined as follows;

$$\begin{aligned} \kappa = 0.41, \quad c_{b1} = 0.1355, \quad \sigma = 2/3, \quad c_{b2} = 0.622, \quad c_{w1} = \frac{c_{b1}}{\kappa^2} + \frac{1 + c_{b2}}{\sigma} \\ c_{w2} = 0.3, \quad c_{w3} = 2, \quad c_{v1} = 7.1, \quad c_{t3} = 1.2, \quad c_{t4} = 0.5 \end{aligned} \tag{10}$$

2.3 Turbulent Boundary Layer And The Law Of The Wall

In the case of a fully turbulent flow without separation over a flat plate at a Reynolds number higher than the critical Reynolds number, the boundary layer follows a behavior (profile) defined by u^+ (dimensionless velocity) and y^+ (dimensionless distance from the wall) (Fig. 4). In this respect, y^+ value is an indication of how well the mesh resolves the turbulent boundary layer.

$$y^+ = \frac{y.u_\tau}{\nu} \tag{11}$$

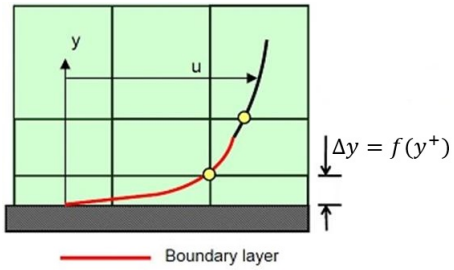


Figure 4: Schematic view of boundary layer.

$$u^+ = \frac{u}{u_\tau} \tag{12}$$

y is the dimensional distance away from the wall, ν is the kinematic viscosity of the fluid, u_τ, the friction velocity, u is the fluid velocity, τ_w is the wall shear stress, ρ is the fluid density.

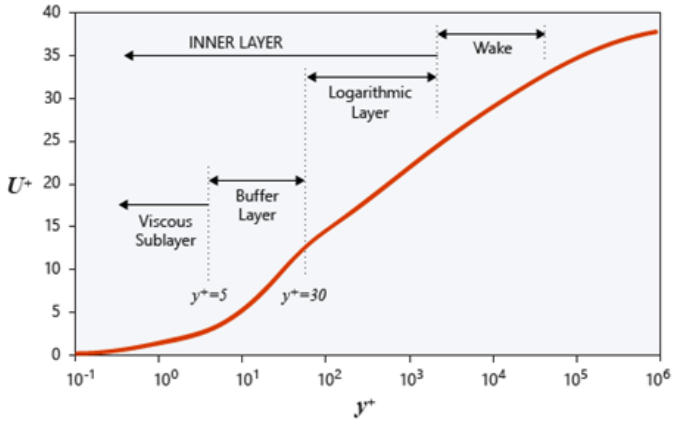


Figure 5: The range of y^+ values.

The turbulent boundary layer can be divided into four regions (Fig. 5):

- Viscous sub-layer: This is the region of the turbulent boundary layer closest to the wall, up to $y^+ < 5$, and is given by the following equation:

$$u^+ = y^+ \tag{13}$$

- Log law layer: This region exists between $30 < y^+ < 300$, and the relationship between u^+ and y^+ is given by the following equation:

$$u^+ = \frac{1}{\kappa} ln(y^+) + C \tag{14}$$

  Where k and C are constants, with widely reported values of $k \sim 0.4$ and $C \sim 0.5$ for a smooth wall, determined from experimental data of turbulent boundary layers. This is known as the law of the wall.

- Buffer layer: This is the region between the viscous sublayer and the log-law region, typically existing between $5 < y^+ < 30$.

$$\begin{aligned} u^+ &\neq y^+ \\ u^+ &\neq \frac{1}{\kappa} ln(y^+) + C \end{aligned} \tag{15}$$

- Outer layer: Beyond $y^+ > 300$, the effects of the freestream become important, and the turbulent boundary layer loses its universality.

Low surface y+ values ($y^+ < 1$): When $y^+ < 1$, the first cell is located inside the viscous sub-layer, resulting in a 'wall-resolution' simulation. This means that there are enough elements to fully resolve the turbulent boundary layer. Medium surface y^+ values ($5 < y^+ < 30$): The first cell is located in the buffer layer. This is generally not recommended. High surface y+ values ($y^+ > 30$): The first cell is in the logarithmic law region of the boundary layer, where turbulence is fully developed. In this case, wall functions are often used to approximate the behavior of the region near the wall (viscous sub-layer and buffer layer). Using a mesh with surface $y^+ > 30$ is computationally less expensive, as fewer computational cells are needed and a high level of accuracy is still achieved using the wall functions. However, wall functions provide a certain amount of inaccuracy when flow separation occurs.

2.4 Surrogate-Assisted Sensitivity Analysis

This section presents a surrogate-based sensitivity analysis method in which a surrogate model is utilized to approximate a complicated computational model. The following sensitivity analysis uses Sobol' indices to quantify the contribution of input variables to the model's output variance. The subsections describe the Kriging-based surrogate modeling technique as well as the generation of Sobol' indices to analyze the influence of individual and interaction effects of the input variables.

2.4.1 Kriging Surrogate Modeling

The Kriging method, originally developed by Krige [18] in the 1950s and mathematically formalized by Matheron [19] in the 1960s, aims to determine the best linear unbiased estimator. Universal Kriging combines a polynomial term for capturing global trends with a Gaussian process term to model local deviations. The stochastic computational model $M(\mathbf{X})$ is approximated as:

$$M(\mathbf{X}) \approx \boldsymbol{\beta}^T f(\mathbf{X}) + Z(\mathbf{X}) \tag{16}$$

In this equation, $\boldsymbol{\beta}^T f(\mathbf{X})$ represents the trend function, where $f(\mathbf{X})$ denotes the regression functions and $\boldsymbol{\beta}$ their corresponding coefficients. The term $Z(\mathbf{X})$ is a Gaussian process with zero mean, and its covariance matrix is given by:

$$Cov(Z(x_i), Z(x_j)) = \sigma^2 R(x_i, x_j; \boldsymbol{\theta}) \tag{17}$$

Here, σ^2 is the variance of $Z(\mathbf{x})$, and $R(x_i, x_j; \boldsymbol{\theta})$ represents the correlation between $Z(x_i)$ and $Z(x_j)$, controlled by hyperparameters $\boldsymbol{\theta}$. The correlation function quantifies the relationship between input points, based on their pairwise distances. In this study, the Gaussian correlation function is employed:

$$R\left(x_i, x_j; \theta\right) = \exp\left[-\frac{1}{2}\left(\frac{|x_i - x_j|}{\boldsymbol{\theta}}\right)^2\right] \tag{18}$$

where x_i and x_j are training points in a dataset of size N_{tr}, and $\boldsymbol{\theta}$ is the length-scale parameter. The correlation matrix $R_{N_{tr}\times N_{tr}}$ is constructed using these pairwise correlations:

$$R_{ij} = r(x_i, x_j; \theta), \quad i, j = 1 \ldots N_{tr} \tag{19}$$

Using this correlation matrix, the optimal parameters $\boldsymbol{\beta}$ and σ are estimated as follows:

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{F}^T\mathbf{R}^{-1}\mathbf{F}\right)^{-1}\mathbf{F}^T\mathbf{R}^{-1}\mathbf{Y_{tr}} \tag{20}$$

$$\hat{\sigma}^2 = \frac{1}{N_{\text{tr}}}\left(\mathbf{Y_{tr}} - \mathbf{F}\hat{\boldsymbol{\beta}}\right)^T\mathbf{R}^{-1}\left(\mathbf{Y}_{\text{tr}} - \mathbf{F}\hat{\boldsymbol{\beta}}\right) \tag{21}$$

where $\mathbf{Y_{tr}}$ denotes the response values at training points $\mathbf{X}_{tr}$, and F (with entries $F_{ij} = f_j(x(i))$) is the Vandermonde matrix. These equations depend on the hyperparameter $\boldsymbol{\theta}$, which is estimated by minimizing the negative log-likelihood function:

$$\hat{\theta} = \arg\min\left(\frac{1}{2}\log(\det(\mathbf{R})) + \frac{N_{\text{tr}}}{2}\log\left(2\pi\sigma^2\right) + \frac{N_{\text{tr}}}{2}\right) \tag{22}$$

Using the estimated hyperparameters, the predicted mean and variance at a new point x are computed as:

$$\hat{y}(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x})^\top\boldsymbol{\beta} + \boldsymbol{r}(\boldsymbol{x})^\top\mathbf{R}^{-1}(Y - \mathbf{F}\boldsymbol{\beta}) \tag{23}$$

$$\sigma_{\hat{y}}^2(\boldsymbol{x}) = \sigma^2\left(1 - \left\langle\boldsymbol{f}(\boldsymbol{x})^\top\boldsymbol{r}(\boldsymbol{x})^\top\right\rangle\begin{bmatrix}\mathbf{0} & \mathbf{F^T}\\ \mathbf{F} & \mathbf{R}\end{bmatrix}^{-1}\begin{bmatrix}\boldsymbol{f}(\boldsymbol{x})\\ \boldsymbol{r}(\boldsymbol{x})\end{bmatrix}\right) \tag{24}$$

These equations allow the Kriging model to make predictions at new points while quantifying uncertainty in the estimations.

2.4.2 Sobol' Sensitivity Indices

Sobol' indices are variance-based sensitivity measures that estimate the contribution of individual input variables (or their interactions) to the output variance of a computational model. For a model $y = M(\mathbf{x})$, where $\mathbf{x} = (x_1, x_2, \ldots, x_d)$ is a vector of d independent random inputs, the total variance D of the output y can be decomposed as shown in Eq. (25).

$$D = \sum_{i=1}^{d} D_i + \sum_{1\le i<j\le d} D_{ij} + \cdots + D_{12\ldots d}, \tag{25}$$

In this decomposition:

- $D_i = \mathrm{Var}_{x_i}\left(\mathbb{E}_{\mathbf{x}_{\sim i}}[y \mid x_i]\right)$ is the first-order effect of x_i,
- $D_{ij} = Var_{x_i,x_j}\left(\mathbb{E}_{\mathbf{x}_{\sim ij}}[y \mid x_i, x_j]\right) - D_i - D_j$ is the second-order interaction effect between x_i and x_j,
- Higher-order terms (e.g., D_{ijk}) represent interactions among three or more variables,
- $\mathbf{x}_{\sim i}$ denotes all variables except x_i, and $\mathbf{x}_{\sim ij}$ excludes x_i and x_j.

In this context, the first-order Sobol' index S_i, the second-order Sobol' index S_{ij}, and the total Sobol' index S_i^T are defined as:

$$S_i = \frac{D_i}{D}, \quad S_{ij} = \frac{D_{ij}}{D}, \quad S_i^T = \frac{D_{T_i}}{D} \tag{26}$$

where D_{T_i} includes all terms involving x_i, and is defined as:

$$D_{T_i} = D_i + \sum_{j \neq i} D_{ij} + \sum_{\substack{j<k \\ j,k \neq i}} D_{ijk} + \cdots + D_{12\ldots d}. \tag{27}$$

In this way, Sobol' indices can be used to analyze in detail the contribution of input variables to the output of the model.

3 NUMERICAL MODEL

Nightingale's wind tunnel test data [3] is accepted as a validation study in the literature. Generic M219 geometry was used in the experiments with applying passive control methods such as slanted walls, transverse rod, flat top spoiler, saw-tooth spoilers. A total of 10 transducers (K20–K29) were located at the center line of the rig in the test study (Fig.6a). The boundary conditions and model dimensions are given in Fig. 6b.

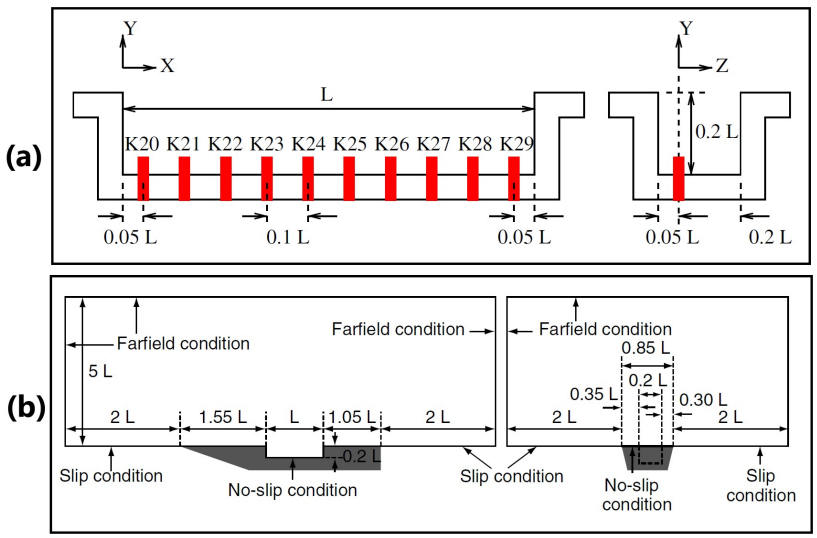


Figure 6: (a) Locations of the pressure transducers. (b) CFD Domain and the boundary conditions [20] [21]

The characteristic length (L) of the M219 generic model was 0.508 m. The L/D and W/D ratios were 5 and 1, respectively. The cavity test was conducted in the transonic regime (M=0.85) for Reynolds number $Re = 1.0 \times 10^6$. Additionally, the operating conditions were set at P_∞=1.0097 bar and T_∞=309.3 K. Static values of temperature, pressure and velocity were derived from the total conditions [20] using isentropic flow equations for a Mach number of 0.85. OASPL results were obtained from tests conducted on the clean geometry, which did not employ any passive or active control methods.

Several temporal and spatial parameters influence the fidelity of the analysis. The temporal parameters are time step (Δt), overall running time (t_{tot}), while spatial parameters are first layer height in wall boundary layer (L_1), cell sizes for x, y and z direction ($\Delta x, \Delta y, \Delta z$) and cell to cell expansion ratio (growth rate). Furthermore, the cell sizes differ across various regions: far field (ff) and cavity (c) region (Fig.7).

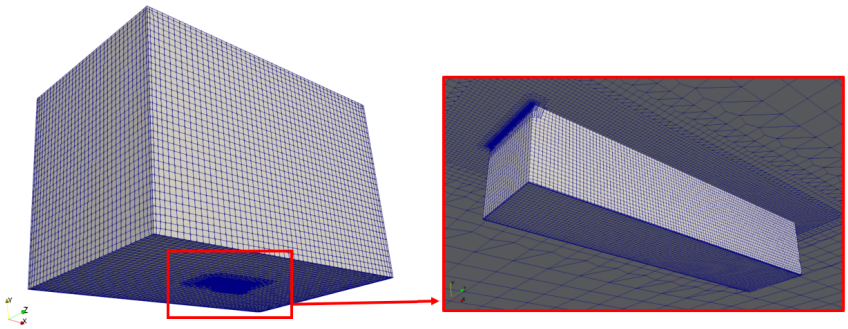


Figure 7: View of mesh model for base case.

In this study, we utilized a time step of 1.84×10^{-6} s, which was essential for resolving high frequencies in the flow. According to the Nyquist theorem, the highest frequency required to obtain an accurate numerical solution for a time step of f = 543 kHz is half of this frequency (272 kHz). The theorem is applied a second time to sample the data, which means that the highest sampling frequency for achieving transient pressure measurements from the flow is 136 kHz [21]. Additionally, the sampling rate of all numerical data is reduced to match the sampling frequency of the experimental data. The total running time (t_{tot}) of the simulation was 0.17 seconds, where the first 0.02 seconds were discarded, resulting in approximately 0.15 seconds of usable data. Simulations are performed via HISA open-source program. It is a fully-coupled density-based solver which mainly consists of OpenFOAM libraries. For numerical schemes, AUSM$^+$-up and Crank-Nicolson is employed for shock-capturing (flux splitting) and temporal discretization respectively.

3.1 Mesh Model For Low Fidelity

Mesh independence study was conducted using the models derived from the parameter alterations listed in Table 1. The growth rate was set to 1.2, and the S-A turbulence model was employed for all cases.

An optimum mesh parameter should be determined for LF simulations. Considering the resource capability and the overall duration of sequential run, according to the law of the wall, it would be a reasonable strategy to keep the value of y^+ between 30 and 150. A mesh model was prepared at an appropriate accuracy level in this range and the "LF-m3" was selected for the LF sensitivity study.

The number of cells downstream (Δx) and the first layer height (L_1) play an important

Table 1: Mesh models for low fidelity.

Model No	L_1 (mm)	Δz_c (mm)	Δx_c (mm)	Δx_{ff} (mm)	Δy_{ff} (mm)	Δz_{ff} (mm)	y^+	Total Size ($\times 10^6$)	*Run* Duration (hour)
LF-m1	0.30	9.2	42.6	306.2	441.0	212.6	200	0.322	8
LF-m2	0.10	5.0	5.0	75.0	75.0	75.0	30	2.691	30
LF-m3	0.70	5.0	5.0	199.4	287.1	413.5	140	0.756	12

role in capturing the flow at the desired accuracy. This result is confirmed when the OASPL plots are examined (see Figure 8). It is also observed that the LF-m1 case cannot resolve sufficient acoustic data, while the other cases reach the desired (for LF) accuracy level.

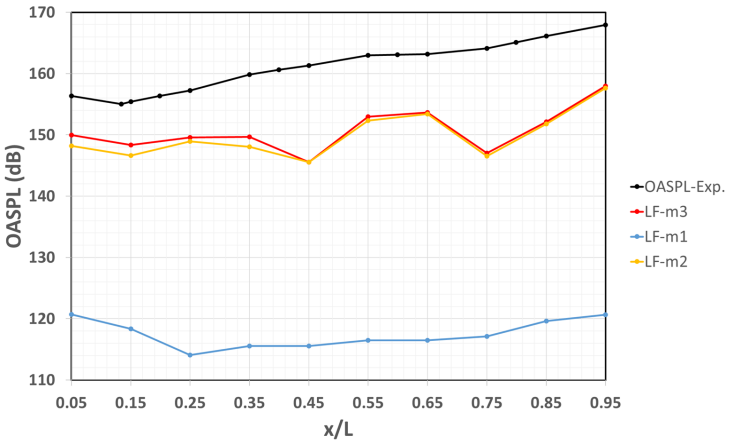


Figure 8: OASPL results for each LF mesh models.

3.2 Geometric Parametrization

According to the results obtained from the experimental and numerical studies conducted to reduce the cavity acoustic penalty, it has been observed that the rear wall and its angle are more dominant. In addition, it is seen that the equipment such as spoilers and baffles added to the leading edge are more efficient and applicable for active control. The parameters determined for the sensitivity analysis performed are (θ) (rear wall angle), (l) (baffle length) and (β) (baffle angle) and h (baffle thickness), as seen in Figure 9. It is assumed that the baffle thickness is limited with the value of 10 percent of the cavity (D=0.1016m) depth. Similarly, the maximum value of the baffle length is limited to 10 percent of the cavity stream-wise length (L=0.508m). The rear wall is allowed to slant a maximum of 10°.

The parametric workflow was carried out with PYTHON. The flow is launched with the definition of the base geometry, parameters, and their limits. The geometries obtained are taken to the OpenFOAM folder directory and geometric (boundary condition) definition, mesh model setting (via CfMesh) and flow analysis are performed for each of them respectively (Figure 10). In this automated study, all settings are applied in the same way for each geometry, and results are obtained separately. The workflow is fulfilled by performing sensitivity analysis with the aerodynamic coefficients (c_f, c_p) and noise level.

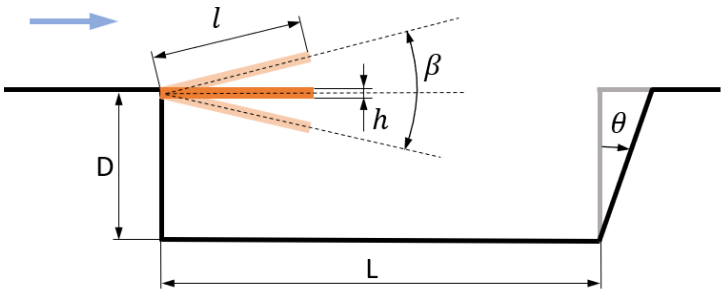


Figure 9: Specified parameters

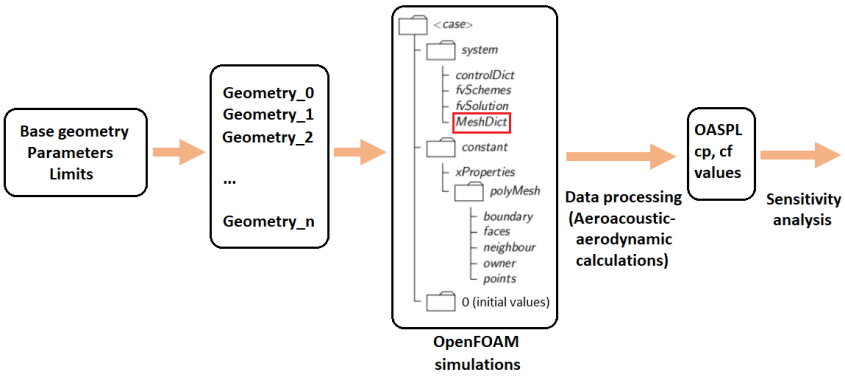


Figure 10: Workflow of PYTHON coding.

4 RESULTS

4.1 CFD Analysis Results

The results of the CFD analysis for a total of 100 geometries were obtained by converting the pressure data from 10 stations on the lower wall to OASPL. In addition, the last 3 probe data close to the aft wall were averaged to obtain a single noise value for each geometry.

Table 2: Parameters for 3 geometries with minimum OASPL.

Sampling No	$l(m)$	$h(m)$	$\beta(^\circ)$	$\theta(^\circ)$
Base case	0.00000	0.00000	0.0	0.0
sampling_Geo_69	0.06328	0.00325	17.4	89.0
sampling_Geo_99	0.07734	0.00200	19.4	81.5
sampling_Geo_89	0.06016	0.00697	16.2	88.2

Table 2 shows the parameters of the 3 cases in which the base geometry and the minimum OASPL were obtained. In Table 3, the values of the aerodynamic coefficients and OASPL values are given. A decrease of approximately 30 dB was observed in the OASPL. Furthermore, the lowest noise levels were obtained at $17^\circ - 19^\circ$ and $80^\circ - 90^\circ$ for β and θ, respectively. It is concluded that the feedback mechanism that generates noise at positive baffle angle values is broken.

Table 4 shows the parameters of the 3 cases in which the base geometry and the maximum OASPL were obtained. In Table 5, the values of the aerodynamic coefficients and OASPL values are given. An increase of approximately 10 dB was observed in the

Table 3: Results for 3 geometries with minimum OASPL.

Sampling No	$OASPL(dB)$	c_p	c_f
Base case	152.6	-0.009815	0.000620
sampling_Geo_69	120.6	-0.120165	0.000787
sampling_Geo_99	122.0	-0.133706	0.000878
sampling_Geo_89	124.4	-0.117788	0.000733

OASPL. Furthermore, the highest noise levels were obtained at $(-2)°-(-3)°$ and $80°-90°$ for β and θ, respectively.

Table 4: Parameters for 3 geometries with maximum OASPL.

Sampling No	$l(m)$	$h(m)$	$\beta(°)$	$\theta(°)$
Base case	0.00000	0.00000	0.0	0.0
sampling_Geo_27	0.08438	0.00160	-3.7	89.2
sampling_Geo_7	0.08750	0.00491	-2.4	80.2
sampling_Geo_37	0.06406	0.00386	-0.5	83.9

Noise negatively affects the cavity when the baffle angle does not change or at small negative values. This situation is due to the flow still clinging to the baffle at these angle values. Thus, the shear layer can easily transmit the flow energy to the rear wall (noise source).

Table 5: Results for 3 geometries with maximum OASPL.

Sampling No	$OASPL(dB)$	c_p	c_f
Base case	152.6	-0.009815	0.000620
sampling_Geo_27	165.3	-0.080374	0.000413
sampling_Geo_7	163.3	-0.050041	0.000396
sampling_Geo_37	163.3	-0.030883	0.000355

When OASPL graphs are plotted for all geometries, the effect of the baffle and the rear wall slope on the cavity can be seen. Except for some cases, it tends to reduce the noise in the system (Fig.11).

The OASPL graphs of the 3 geometries that have the greatest effect on noise (negatively and positively) are given in Fig.12 in comparison with the base case and experimental data.

The best demonstration of the effect on the mechanism of noise generation is LIC visulation. Line integral convolution (LIC) is a method to visualize a vector field at high spatial resolutions.

According to the M219 clean geometry results (base case), in the 27th geometry the flow clings to the baffle and increases its effect on the rear wall. On the other hand, in Geometry-69 the baffle managed to keep the flow away from the cavity (Fig.13).

4.2 Sensivity Analysis Results

Global sensitivity analysis is conducted to identify key uncertain parameters influencing aeroacoustic noise. The statistical properties of the variables used in the analysis are

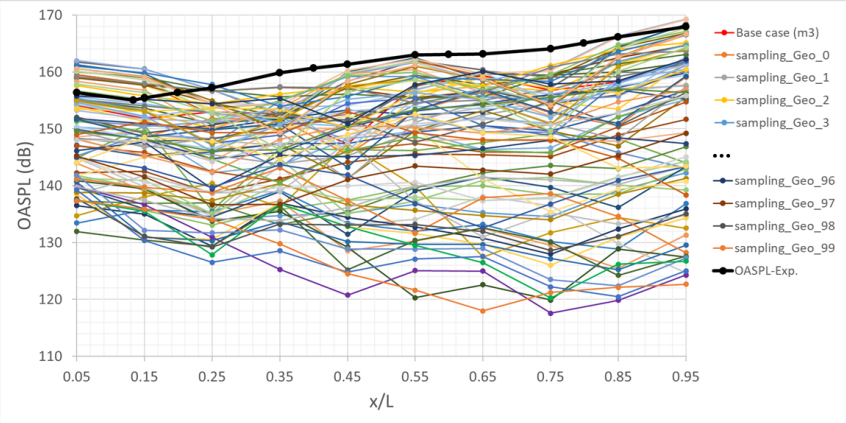


Figure 11: OASPL results for all geometries.

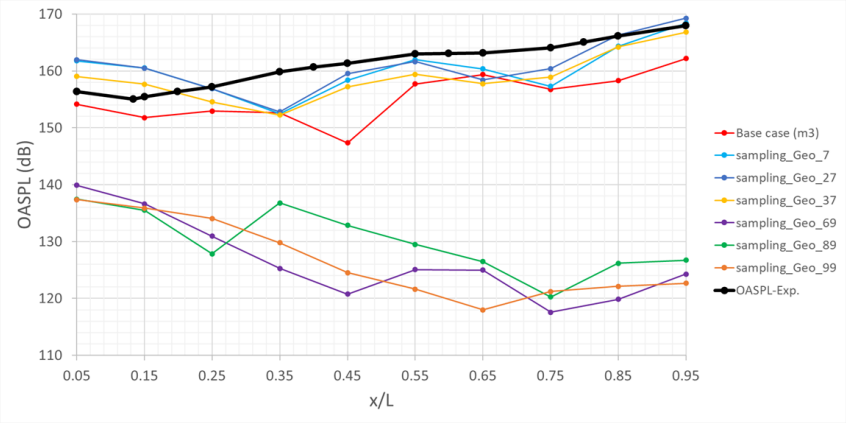


Figure 12: OASPL results for 3 geometries with maximum and minimum values.

presented in Table 6.

Table 6: Statistics of uncertain variable.

Uncertain Variable	Distribution	Interval
l (m)	Uniform	[0, 0.1]
h (m)	Uniform	[0.015, 0.075]
β (°)	Uniform	[-20, 20]
θ (°)	Uniform	[0, 10]

The results of the Sobol' sensitivity analysis have been examined to assess the sensitivity of model parameters to the OASPL value. Figure 14 presents the first- and second-order Sobol' indices, while Figure 15 illustrates the total Sobol' indices.

The first-order Sobol' indices represent the direct effect of each variable on the model output. The results indicate that the β parameter exhibits by far the highest sensitivity. In contrast, the other variables, l, h, and θ, have significantly lower first-order effects. This suggests that the model output is predominantly influenced by the β variable. The second-order Sobol' indices reveal the interaction effects between pairs of variables. The results show that l has strong interactions with both β and h. Notably, the highest second-order index is observed for the $l - \beta$ pair, implying that these two variables collectively have a significant impact on the model output. Other interactions exhibit comparatively

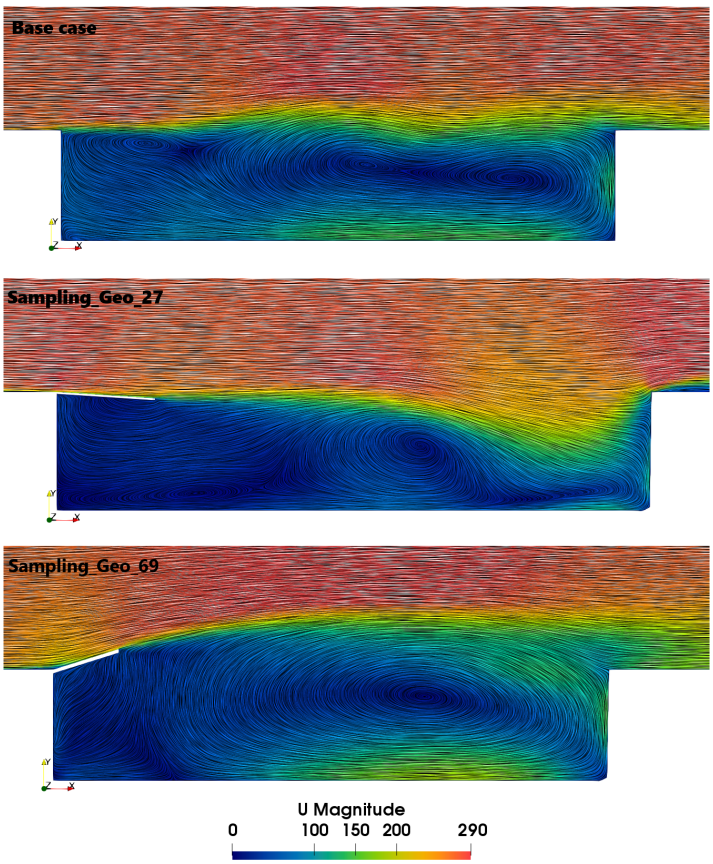
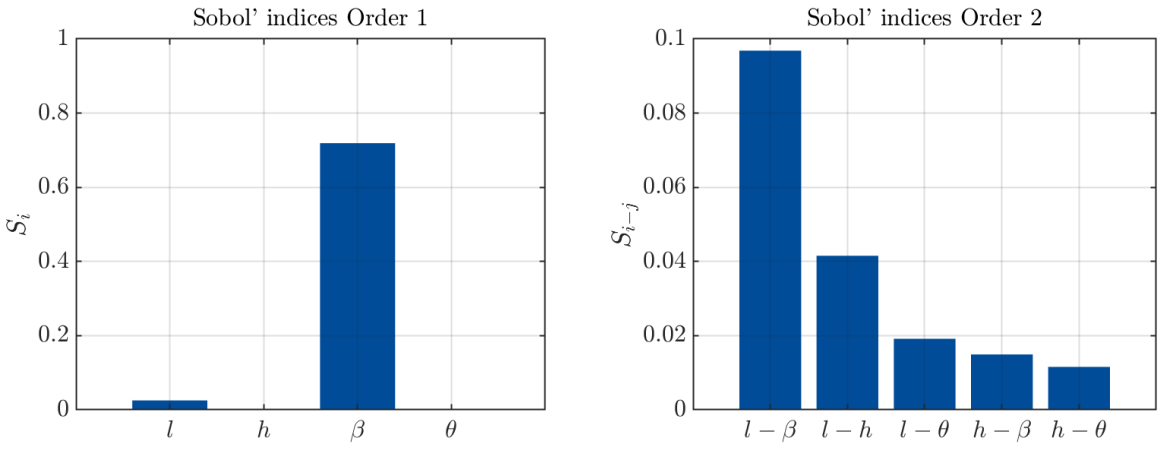


Figure 13: LIC visualization of the U vector field by U magnitude at y=0 plane.



(a) A subfigure (b) Another subfigure

Figure 14: First and second order Sobol indices

lower values.

The total Sobol' indices capture the overall impact of each variable on the model output, considering both their individual contributions and interactions with other variables. The results confirm that β remains the most dominant factor. However, the l variable has a higher total effect compared to its first-order Sobol' index, indicating that its contributions primarily arise from interactions with other input variables. In contrast, the total effects

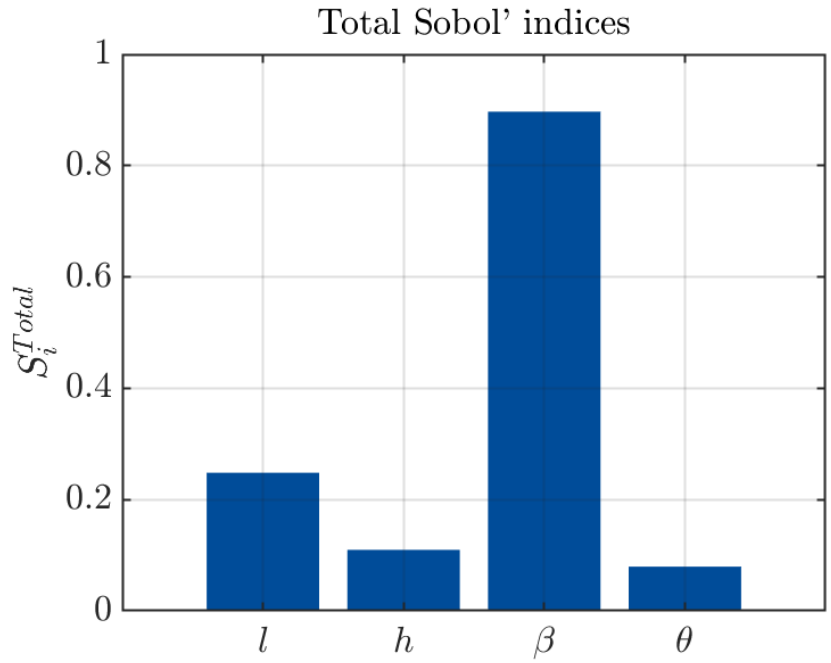


Figure 15: Total Sobol indices.

of h and θ remain relatively low, suggesting their limited influence on the model output.

This analysis highlights that β is the most critical parameter in terms of sensitivity, with the model output being largely dependent on this variable. Furthermore, the relatively high total Sobol' index of l underscores the importance of considering its interactions with other variables for model performance. In contrast, the contributions of h and θ remain relatively minor. These findings provide valuable insights for uncertainty analysis and model optimization.

5 CONCLUSIONS

This study has successfully utilized a computational methodology which consisted of CFD simulations, Kriging surrogate modeling, and Sobol' sensitivity analysis to assess the effects of design changes on the created aeroacoustic noise of the M219 open cavity at transonic speed (M=0.85). The study looked at the sensitivity of the Overall Sound Pressure Level (OASPL) to the rear wall slant angle (θ), and changes to the dimensions (length l, thickness h) and orientation (angle β) of a leading edge baffle.

The simulations using CFD, which stands for Computational Fluid Dynamics, were done on 100 different shapes. These simulations showed that the noise levels varied a lot. When the baffles, which are barriers or plates used for controlling airflow, had angles between 17° and 19° and were tilted in a positive direction, the noise was reduced significantly. Specifically, the noise dropped by about 30 decibels when compared to a design without baffles, often called a "clean cavity." This reduction suggests that these positive angles are effective at interfering with the usual patterns of airflow that increase noise. However, when the baffle angles were slightly negative, the noise levels went up by as much as 10 decibels. This increase in noise probably happens because the stream of air stays stuck to the baffle and hits the back wall more directly, resulting in more sound.

The Sobol' sensitivity analysis measured the importance of different parameters. It

found that the baffle angle (β) was the most important factor, having the biggest impact on the output variance (OASPL). The baffle length (l) on its own didn't affect much, but it played a big role due to its strong interactions with other factors, mostly with the baffle angle (β) and somewhat with the baffle thickness (h). The rear wall slant angle (θ) tested between 0° and 10°, and the baffle thickness (h) tested between 1.5% and 7.5% of the cavity depth, contributed only slightly to changes in noise levels.

This research highlights the importance of the shape and angle (β) of the leading-edge baffle in passively controlling noise in open cavities. The sensitivity to the angle and its interaction with other factors provides valuable guidance for future design improvements to reduce noise. We used a surrogate-based sensitivity analysis framework, which helped identify important design factors from a wide range of options through cost-effective, simpler computer simulations. These insights improve our understanding of cavity air flow and aid in creating quieter aerospace vehicles.

ACKNOWLEDGEMENTS

Acknowledge persons and institutions that supported the work.

REFERENCES

[1] G. Guo and Q. Luo. Flowfield structure characteristics of the hypersonic flow over a cavity:from the continuum to the transition flow regimes. *Acta Astronautica*, 161:87 – 100, 2019.

[2] J. Rossiter and A. Kurn. Wind tunnel measurements of the unsteady pressures in and behind a bomb bay (canberra). *Technical Note- AERO.2845*, 1962.

[3] D. Nightingale, J. Ross, and G. Foster. Cavity unsteady pressure measurements examples from wind tunnel tests. *Technical Report Version 3*, 2005.

[4] J. Ross and J. Peto. The effect of cavity shaping, front spoilers and ceiling bleed on loads acting on stores, and on the unsteady environment within weapon bays. *Technical Report DERA/AS/HWA/CR97010/1*, 1997.

[5] D. A. Buell. An experimental investigation of the airflow over a cavity with antiresonance devices. *NASA Technical Report D-6205*, 1971.

[6] L. Shaw, R. Clark, and D. Talmadge. F-111 generic weapons bay acoustic environment. *Journal of Aircraft*, 25(2):147–53, 1988.

[7] R. Smith, E. Gutmark, and K. Schadow. Mitigation of pressure oscillations induced by supersonic flow over slender cavities. *Journal of Aircraft*, 29(6):999–1004, 1992.

[8] D. MacManus and D. Doran. Passive control of transonic cavity flow. *Journal of Fluids Engineering*, (130: 064501(4pp.)), 2008.

[9] B. Smith, T. Welterlen, B. Maines, S. Linn, M. Stanek, and L. Grove. Weapons bay acoustic suppression from rod spoilers. *40th AIAA aerospace sciences meeting and exhibit*, 2002.

[10] L. Ukeiley, M. Ponton, J. Seiner, and B. Jansen. Suppression of pressure loads in cavity flows. *AIAA Journal*, 42(1):70–9, 2004.

[11] R. Schmit, D. Schwartz, V. Kibens, G. Raman, and J. Ross. High and low frequency actuation comparison for a weapons bay cavity. *43st AIAA aerospace sciences meeting and exhibit*.

[12] J. Rossiter. A preliminary investigation into armament bay buffet at subsonic and transonic speeds. *Technical Memorandum AERO.679*, 1960.

[13] A. Charwat, J. Roos, J. F. Dewey, and J. Hitz. An investigation of separated flows—part i: the pressure field. *AIAA Journal*, 28(6):457–70, 1961.

[14] H. Heller and D. Bliss. Aerodynamically induced pressure oscillations in cavities: Physical mechanisms and suppression concepts. *Technical Report-AFFDL-TR-74-133*, 1975.

[15] O. M. Oyewola and M. J. Labiran. Control of acoustic tone formation using shape geometry modifications and jet blowing in a cavity experiencing supersonic flow. *International Journal of Thermofluids*, 24(100978), 2024.

[16] S. J. Kim, W.-X. Huang, and H. J. Sung. The reduction of noise induced by flow over an open cavity. *International Journal of Heat and Fluid Flow*, 82(108560), 2020.

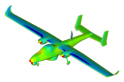
[17] P. Lojeki, K. S. Debska, and M. Mach. Influence of cavity edges shape on flow induced noise. *Vibrations in Physical Systems*, 32 (1)(2021110), 2021.

[18] D. Krige. *A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand*. Chemical, Metallurgical and Mining Society of South Africa, 1951. URL `https://books.google.com.tr/books?id=Mu6SnQAACAAJ`.

[19] G. Matheron. Principles of geostatistics. *Economic Geology*, 58(8):1246–1266, 12 1963. ISSN 0361-0128. doi:10.2113/gsecongeo.58.8.1246. URL `https://doi.org/10.2113/gsecongeo.58.8.1246`.

[20] S. Lawson and G. Barakos. Review of numerical simulations for high-speed, turbulent cavity flows. *Progress in Aerospace Sciences*, 47(1):186–216, 2011.

[21] S. Lawson and G. Barakos. Understanding cavity flows using proper orthogonal decomposition and signal processing. *Jornal of Algorithms & Computational Technology*, 4(1):1–100, 2009.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

DEVELOPMENT OF A SOUNDING ROCKET MULTIDISCIPLINARY PRELIMINARY DESIGN OPTIMISATION FRAMEWORK WITH TRAJECTORY OPTIMISATION

Palaio, A. M.[1*], Marta, A. C.[2] and Gil, P. J. S.[2]

1: Academia da Força Aérea
Força Aérea Portuguesa
Granja do Marquês , 2715-021, Sintra, Portugal
ampalaio@academiafa.edu.pt, https://academiafa.edu.pt/

2: IDMEC, Instituto Superior Técnico
Universidade de Lisboa
Av. Rovisco Pais 1, 1049-001, Lisboa, Portugal
{andre.marta, paulo.gil}@tecnico.ulisboa.pt, https://mdo.tecnico.ulisboa.pt/

Abstract. *The design of rockets is known to be a complex task, not only due to the harsh operating conditions but also the strong coupling among disciplines. A multidisciplinary optimisation (MDO) framework was developed, aimed at providing preliminary designs of a single-stage solid propellant rocket. The choice of the optimiser algorithm, MDO architecture and discipline models, namely, mass and sizing, flight dynamics, aerodynamics, propulsion, structural and atmospheric, were such that the developed numerical tool has a very low computational cost while being able to meet a set of pre-established mission requirements. The resulting design framework solved a co-design optimisation problem, due to the coupling between the trajectory and rocket sizing optimization processes. The capabilities of the design framework were tested for different sets of design variables and multiple missions, with increasing complexity, for an optimisation problem aimed at minimizing the total mass of the rocket while imposing a minimum altitude constraint, with a prescribed payload capacity. First, a case study with 10 geometric design variables achieved a total mass reduction of 27.7 % when compared to a real rocket, REXUS 2. Then, sensitivity studies of the payload and the minimum altitude confirmed that the rocket sizing is significantly impacted by both. The modularity of the framework allows a straightforward extension to other types of rockets, such as multi-stage or liquid-propellant.*

Keywords: MDO, Trajectory, Co-design, Sounding Rocket, Modularity

1 INTRODUCTION

Over the last two decades, a new generation of entrepreneurs has made an unprecedented investment in Space, completely changing the paradigm. Currently, private initiatives play an important role in the future of the space industry, which is no longer controlled by the political agendas of a few superpowers [1]. As a consequence, space exploration, space tourism and space infrastructure are now the main focus of such private and semi-private initiatives, which have completely revived the global space economy. As of 2022, 78% comes from commercial space products, services, infrastructure and support industries and only 22% from government budgets [2].

As Science has always been the major beneficiary from space human endeavours [3], it is expected that this new interest in space affairs will award scientific groups with new lines of investment across a wide range of applications, such as, Research and Development (R&D) on new Launch Vehicle (LV) designs capable of accomplishing their assigned goals in compliance with the most demanding mission requirements.

Currently, the scientific research on modern Multidisciplinary Design Optimization (MDO) methods applied to the design process of LV is a hotspot in the aerospace industry, in an effort to further minimize the material usage, manpower, cost and time, while maximising the reliability, operability and safety of such systems [4].

The main goal of this work is, then, to develop and validate an MDO framework coupled with trajectory optimisation capable of conducting the preliminary design of sounding rockets with a minimum payload capacity of 98 kg and 100 km minimum peak altitude, so that the results may be compared to well known and documented rockets, namely, the Rocket borne Experiments for University Students program (REXUS) [5].

2 ROCKET FUNDAMENTALS

2.1 Mass and Sizing

A model, subdivided into six smaller subcomponents, each of them related to a main rocket part, was created from a set of analytical equations to estimate the masses and component sizing. The targeted rocket parts were: nose cone, modules, fins, nozzle, body tube, and SRM, sorted by the model execution order. Additionally, a final component was also created to calculate a few general properties, namely, the initial mass, empty mass, structural mass, and structural factor of the rocket, as it is illustrated in Fig. 1.

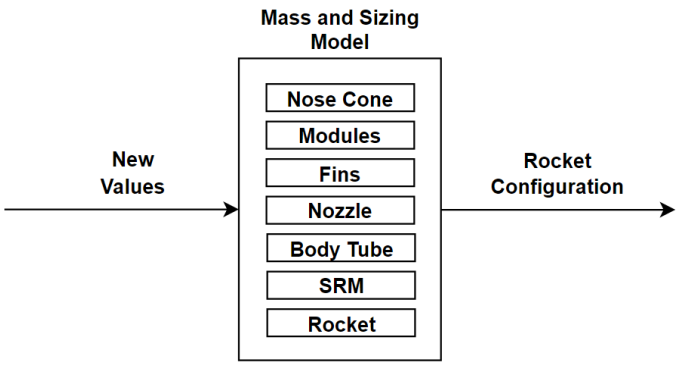


Figure 1: Mass and sizing model simplified schema highlighting the inpust, outputs and inner components.

2.2 Aerodynamics

An aerodynamic model was designed to estimate the aerodynamic behaviour of the rocket at each operating state. The C_d profile of the rocket was calculated based on three main drag sources: nose cone, base, and fins, as stated in the Equation 1.

$$Cd_0 = Cd_{nc} + Cd_b + Cd_f. \tag{1}$$

A compressible flow correction was applied for better accuracy under compressible flow regimes [6].

At subsonic speed ($M_a < 0.8$), the compressible flow correction for the aerodynamic coefficient is defined as

$$C_d = \frac{C_i}{\sqrt{1 - M_a^2}}, \tag{2}$$

where M_a is the free stream Mach number. At the transonic region ($0.8 \leq M_a \leq 1.1$), the corrected aerodynamic coefficient is given by

$$C_d = \frac{C_i}{\sqrt{1 - (0.8)^2}}. \tag{3}$$

At supersonic speed ($M_a > 1.1$), the corrected aerodynamic coefficient is

$$C_d = \frac{C_i}{\sqrt{M_a^2 - 1}}. \tag{4}$$

In addition, a recovery system contribution was also integrated in the model to simulate the behavior of parachute deployment during the descent phase of the flight profile, namely drogue and main parachutes [7].

As the rocket descends and passes the reference altitude of drogue parachute deployment, an additional component of induced drag is added to the overall C_d profile calculated using the following expression:

$$D_{drogue} = \frac{1}{2}\rho v^2 C_{D\,drogue} S_{drogue}\ , \tag{5}$$

where D_{drogue} is the drag component due to the drogue parachute, $C_{D\,drogue}$ is the drag coefficient of the drogue parachute, and S_{drogue} is the drogue parachute area.

Similarly, after the main parachute descent altitude is reached, the additional component of induced drag is calculated as

$$D_{main\,parachute} = \frac{1}{2}\rho v^2 C_{D\,main\,parachute} S_{main\,parachute}\ , \tag{6}$$

where $D_{main\,parachute}$ is the drag component due to the main parachute, $C_{D\,main\,parachute}$ is the drag coefficient of the main parachute, and $S_{main\,parachute}$ is the main parachute area.

2.3 Propulsion

A propulsion model was developed to accurately predict the behaviour of the main physical properties of a Solid Rocket Motor (SRM) under real operating conditions.

Based on several authors [8–12], a set of analytical equations was assembled to model the grain burnback, i.e., the propellant regression rate and respective propellant burning areas over time,

$$m_{pi} = m_{p0} - (m_0 - m_i) \ , \tag{7}$$

$$R_p = \sqrt{\frac{-m_{pi}}{\rho_p \pi L_{grain} + R_i^2}} \ , \tag{8}$$

$$A_b = 2\pi R_p L_{grain} \ , \tag{9}$$

where m_{pi} is the instantaneous propellant mass, m_{p0} is the initial propellant mass, m_0 is the rocket lift-off mass, m_i is instantaneous rocket mass, R_p is the port radius, ρ_p is the propellant density, L_{grain}, is the grain length, R_i is the grain inner radius, and A_b is the instantaneous burning area.

Then, using the propellant burning area as the main input, a second set of equations was assembled to model the internal ballistic behaviour of the motor, namely the combustion chamber pressure and the thrust, using one-dimensional isentropic flow equations [13], as

$$T_e = T_t(1 + \frac{\gamma - 1}{2} M_e^2)^{-1} \ , \tag{10}$$

$$P_t = \left[\frac{a \rho_p A_b}{C_D A_t}\right]^{\frac{1}{1-n}} \ , \tag{11}$$

$$P_e = P_t(1 + \frac{\gamma - 1}{2} M_e^2)^{\frac{-\gamma}{\gamma - 1}} \ , \tag{12}$$

$$V_e = M_e \sqrt{\gamma R T_e} \ , \tag{13}$$

$$\dot{m} = -C_D A_t P_t \ , \tag{14}$$

$$Thrust = \dot{m} V_e + (P_e - P) A_e \ , \tag{15}$$

where a is the burn rate coefficient, C_D is the nozzle discharge coefficient, A_t is the nozzle throat area, n is the propellant ballistic exponent, T_e is the nozzle exit temperature, T_t is the nozzle throat temperature, γ is the gas specific heat ratio, M_e is the nozzle exit mach number, P_e is the nozzle exit pressure, P_t is the internal casing total pressure, V_e is the nozzle exit velocity, R is the universal gas constant, and P is the atmospheric pressure.

2.4 Structures

A model was created to assess the structural integrity of the rocket along the flight profile. First, given the sum of drag, thrust and weight forces, the model calculates the resulting compressive loading at the body tube cross-sectional area. Then, it compares it with the linearised critical buckling stress of a thin elastic cylindrical shell in order to evaluate if, at any moment in time, the rocket was subject to such a loading condition for buckling to occur, through the expression:

$$\sigma_{crit} = \frac{\gamma E}{\sqrt{3(1-v^2)}} \left(\frac{th}{R}\right) \ , \tag{16}$$

where E is the cylindrical shell Young modulus, R its radius, th its thickness, v the material Poisson's ratio and γ a multiplication factor given by

$$\gamma = 1 - 0.901\left(1 - e^{-\phi}\right) \quad \text{with} \quad \phi = \frac{1}{16}\sqrt{\frac{R}{th}} \ . \tag{17}$$

.

Additionally, using the Flutter Boundary Equation [14], the fin flutter velocity is also monitored throughout the entire flight profile in order to evaluate if the structural integrity of the fins remains unharmed, as this is the pivotal rocket component for stability,

$$V_f = a\sqrt{\frac{G_E}{\frac{YAR^3}{(t/c_r)^3(AR+2)}\left(\frac{\lambda+1}{2}\right)\left(\frac{P}{P_0}\right)}} \ . \tag{18}$$

2.5 Atmosphere

An atmospheric model adapted from the *OpenAero-Struct Python* library was also developed to provide with the main atmospheric properties throughout the trajectory. First, arrays with the values of each atmospheric parameter, retrieved from the standard atmosphere convention tables [15, 16], were created.

By interpolating the altitude value (model input) using the *Akima1DInterpolator* class imported from the *Scipy Python* library [17, 18], it was possible to find the respective values of all atmospheric parameters for each altitude, namely, temperature T, pressure P_a, density ρ, speed of sound c, gravitational acceleration g, dynamic viscosity μ and kinematic viscosity k. These atmospheric property values are the outputs of the model, which will be fed to the other models within the trajectory group, as it is ilustrated in Fig. 2.

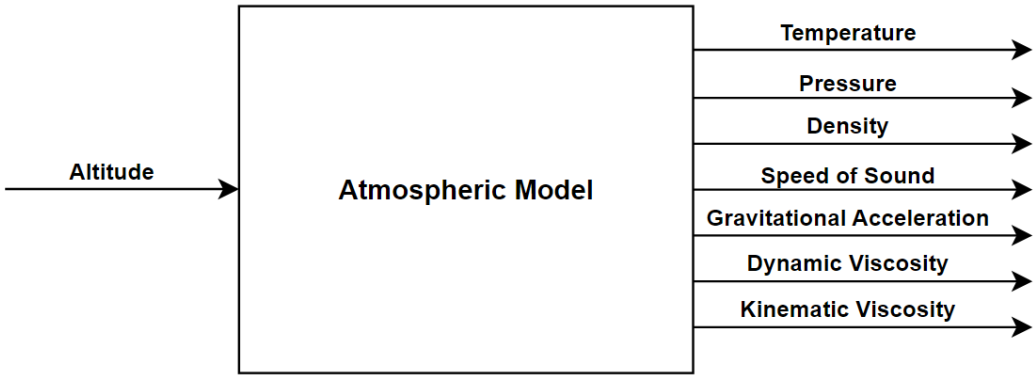


Figure 2: Atmospheric model diagram highlighting inputs and outputs.

2.6 Flight Dynamics

A set of flight dynamics equations was used, capable of translating the complex interactions between the rocket, the atmosphere, and any other external factors with active

influence on the rocket [19].

In order to reduce the number of state variables for simplicity and computational cost efficiency, a 2 degrees of freedom (DoF) trajectory defined in a vertical plane was chosen over other more complex dynamics systems, with higher DoF.

The flight dynamics of the rocket can be reduced to the following set of equations:

$$\dot{V} = \frac{T}{m}\cos\alpha - \frac{D}{m} - g\sin\gamma, \tag{19}$$

$$\dot{\gamma} = -\left(\frac{g}{V} - \frac{V}{R_e + h}\right)\cos\gamma + \frac{T}{m}\sin\alpha, \tag{20}$$

$$\dot{x} = V\cos\gamma, \tag{21}$$

$$\dot{h} = V\sin\gamma. \tag{22}$$

where $\dot{v}$ is the rate change of the velocity, $\dot{\gamma}$ is the rate change of the pitch angle, $\dot{h}$ the rate change in altitude, and $\dot{x}$ the rate change in downrange [19].

2.7 Trajectory

In order to implement the trajectory model, a high-level group was created, with 5 coupled models within, namely, Flight Dynamics, Atmospheric, Propulsion, Aerodynamics and Structural, as illustrated in Fig. 3.

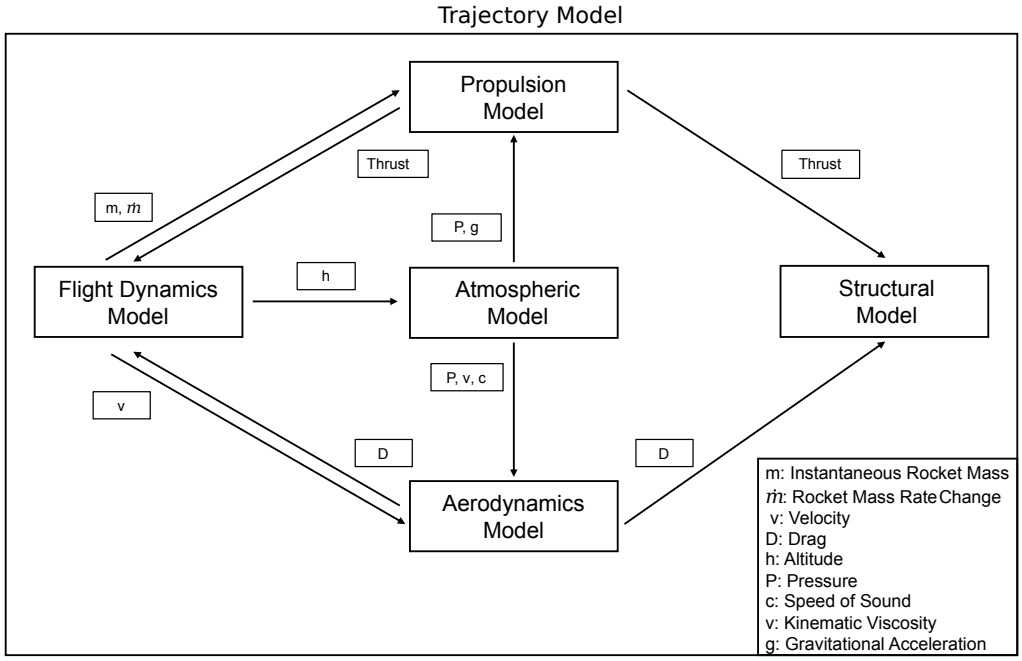


Figure 3: Overview of the trajectory model.

This model was then integrated in a top-level group with the mass and sizing model in order to create a framework capable of conducting an MDO process coupled with trajectory optimisation.

3 MULTIDISCIPLINARY DESIGN OPTIMISATION

Recent advances in technology have improved accessibility to higher computational power at gradually lower costs. Consequently, modern computer-based engineering systems capable of conducting complex MDO processes superseded the traditional concurrent engineering philosophy-based systems, where Disciplinary Design Optimisation (DDO) was conducted.

3.1 MDO Architectures

The main MDO architectures currently in use by the aerospace industry can be classified into two different groups: single-level (or monolithic) and multi-level (or distributed), according to the number of optimisers used in each architecture (single or multiple optimisers, respectively). The monolithic MDO architectures solve a single optimisation problem, while the distributed architectures decompose the original problem into a set of smaller optimisation subproblems which provide the exact same solution.

Single-level architectures are characterised by only using an optimiser in the top level of the multidisciplinary system, which is the governing level responsible for enforcing multidisciplinary feasibility [20].

The multidisciplinary feasible (MDF) architecture, solves the optimisation problem by implementing a system-level optimiser which calls a multidisciplinary analysis (MDA) responsible for solving all governing equations at the subsystem/component level until the coupling variables converge within the specified tolerance limits [21].

As an alternative approach, the individual discipline feasible (IDF) architecture adds additional independent variables to the problem to ensure that each discipline can be solved separately, while interdisciplinary equilibrium is maintained by a set of optimisation constraints that ensure the overall feasibility of the design once the optimisation convergence is achieved [21]. IDF potentially solves the high computational cost opened by the MDF architecture by conducting each discipline feasibility analysis independently and, in parallel, favouring speed and efficiency, at the cost of introducing additional variables and optimisation constraints, which increases the overall complexity of the original problem and might pose scalability issues for larger applications [22].

In contrast to the single-level architectures, multilevel architectures divide the original optimisation problem into a system-level optimisation problem and several sub-system level problems, according to the number of levels. The basic idea is for the system-level optimisation problem to coordinate the smaller sub-level problems, which in turn will be solved locally. The four most common architectures of this sort are: Collaborative Optimisation (CO), Concurrent SubSpace Optimisation (CSSO), Bi-Level Integrated System Synthesis and Analytical Target Cascading (ATC) [22].

After analysis, it was defined that the most suitable architecture for the developed framework was a single-level MDF architecture, as it is capable of solving the optimisation problem using a system-level optimiser that directly handles all the design variables and constraints, relaying on an MDA block to ensure multidisciplinary feasibility at each iteration, balancing simplicity in the hierarchical build of the design, efficiency of the data flows, and computational time.

3.2 Optimisation Algorithms

Optimisation algorithms are numerical methods designed to systematically search for the variable values that optimise the objective function [23]. They can be divided into two major groups: combinatorial (or discrete) or continuous, depending on whether the variables are discrete or continuous quantities, respectively. Discrete optimisation algorithms are hardly suitable for rocket design applications due to the continuous nature of the majority of the design variables involved, which typically represent physical properties (continuous in their essence) [23].

Continuous optimisation algorithms can be further divided into two other groups: linear and nonlinear.

Linear Programming (LP) algorithms are particularly designed for the minimisation (or maximisation) of a linear objective function subject to linear constraints.

Nonlinear Programming algorithms (NLP) are suitable for nonlinear, yet smooth, objective functions with at least continuous first partial derivatives in the target regions of the design space where the optimisation solution might be located [23]. By nature, the objective function, inequality, and equality constraints have a nonlinear behaviour in the rocket design environment with variables having quadratic, cubic, exponential or otherwise nonlinear relationship. Consequently, NLP algorithms need to be used in this work.

One of the most efficient methods for constrained nonlinear optimisation problems is Sequential Quadratic Programming (SQP), regarding function evaluations and computation cost [24]. Some of the most interesting characteristics are: Linear constraints and bounds remain satisfied; For n active constraints, SQP methods can achieve local convergence with quadratic convergence rate; Local convergence speed is superlinear; A large number of constraints can be treated by an active set strategy and the computation of gradients for inactive restrictions can be omitted.

In essence, Sequential Least Squares Quadratic Programming (SLSQP) is an optimisation method within the SQP wider family in which the constraints are linearized about the current point and a quadratic approximation of the objective function is defined [24].

Its formulation can be posed in standard form as

$$\min_{y \in \mathbb{R}^n} \quad f^k(y) \tag{23}$$

$$\text{subject to} \quad g^k(y) \leq 0, \tag{24}$$

where

$$f^k(y) = \frac{1}{2}(y - x_k)^T B_k (y - x_k) + \nabla f(x_k)^T (y - x_k) + f(x_k), \tag{25}$$

$$g_j^k(y) = \nabla g_j(x_k)^T (y - x_k) + g_j(x_k), \quad j = 1, \ldots, m. \tag{26}$$

Then, the Least Squares mathematical method is used to solve iteratively a set of Quadratic Programming subproblems, starting in a given vector of parameters, x^0, until a $(k+1)^{th}$ iterate, x^{k+1}, is reached in which the objective function converges within a

specific tolerance condition, in compliance with all equality and inequality constraints [24].

In each iteration k, the optimiser needs to evaluate the function and constraint gradients, Δf and Δg, respectively, to determine a search direction d^k. Then, a line search is performed along that direction to find the step length α^k that minimises the $f(x)$, and a new iteration then follows at [24]:

$$x^{k+1} := x^k + \alpha^k d^k \ , \tag{27}$$

where d^k is the search direction within the k^{th} step and α^k is the step length.

3.3 Trajectory Optimisation

Trajectory optimisation problems are a part of the larger optimal control theory branch of mathematics, which specifically seeks to find the optimal control law of a dynamic system that satisfies a set of constraints while minimising a cost function.

A general mathematical problem definition can be defined as follows:

$$\begin{aligned}
\text{Optimal Trajectory:} \quad & \{x^*(t), u^*(t)\} && (28)\\
\text{System Dynamics:} \quad & \dot{x} = f(t, x, u) && (29)\\
\text{Constraints:} \quad & c_{\min} < c(t, x, u) < c_{\max} && (30)\\
\text{Boundary Conditions:} \quad & b_{\min} < b(t_0, x_0, t_f, x_f) \\
& < b_{\max} && (31)\\
\text{Cost Functional:} \quad & J = \phi(t_0, x_0, t_f, x_f) + \\
& \int_{t_0}^{t_f} g(t, x, u)\, dt && (32)
\end{aligned}$$

where x is for the state variables, u is for control variables, $f(t, x, u)$ are the system dynamics functions, $c_{\min}$, $c_{\max}$ and $c(t, x, u)$ are the lower, upper bounds and boundary function, respectively, $b_{\min}$, $b_{\max}$, $b(t_0, x_0, t_f, x_f)$ are the lower, upper bounds and boundary function, respectively, and, finally, J is the cost function.

3.3.1 Direct vs Indirect Collocation

Generally speaking, collocation methods belong to a broader transcription family of methods, in which differential equations governing the rocket system dynamics are enforced in a grid of points discretised from an initial continuous time interval, called collocation nodes, ensuring that the discretised approximations at these points are faithful to the continuous dynamics [25].

Collocation methods can be formulated in two different approaches: direct or indirect. Direct methods first discretise and then optimise while indirect methods optimise and then discretise [26], as illustrated in Fig. 4.

Indirect collocation methods, first establish the necessary and sufficient conditions for optimality, thus forming a Hamiltonian boundary-value problem (HBVP) which is analytically derived by applying the Pontryagin's Minimum Principle (PMP). Then, the newly created differential equations governing the adjoint variables, the control equation, and the boundary conditions form a new Two Point Boundary Value Problem (TPBVP).

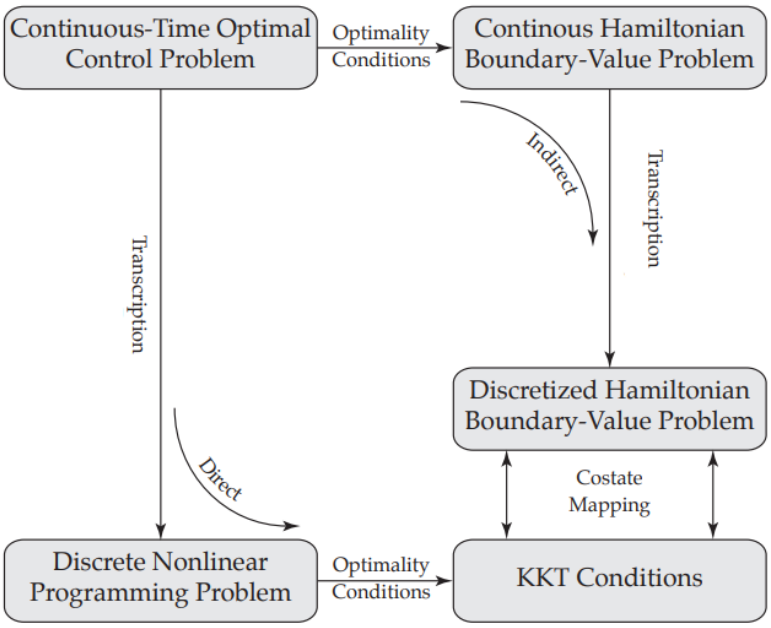


Figure 4: Comparison between direct and indirect collocation methods [27].

Then, TPBVP is discretised using a collocation method, such as, Hermite-Simpson, for example, transforming the continuous-time problem into a finite-dimensional nonlinear programming problem (NLP), which is numerically solved through the application of optimisation solvers, such as, gradient-based methods or sequential quadratic programming (SQP), until the Karush-Kuhn-Tucker (KKT) optimality conditions are met [28].

In contrast, direct collocation methods are the most used in the context of trajectory optimisation due to their simplicity, robustness, and range of application [28]. These methods are characterised by first discretising a continuous time interval into a grid of collocation points. Then, the state and control variables are also discretised at the collocation points, in which dynamics are enforced. Lastly, a nonlinear program is formulated from the discretised points and solved [28].

In comparison with the latter, indirect methods are commonly more accurate, providing stronger solutions with reliable error estimates due to analytically deriving the necessary and sufficient conditions in the early stages of the problem formulation, at the cost of requiring a better initialisation as they tend to have smaller convergence regions [26].

Therefore, at the preliminary design level, for a single-stage suborbital trajectory optimisation process, the direct collocation methods are the better choice because they have proven to be simpler, computationally faster, and accurate enough, while avoiding potential convergence issues for problems with increased complexity.

3.3.2 Pseudo-spectral Methods

Pseudo-spectral methods have gained traction in the trajectory optimisation field in recent years as a powerful, highly efficient alternative for the already well-established direct collocation methods to solve continuous nonlinear constrained optimal control problems with smooth functions, such as single-phase rocket trajectory optimisation problems. Highly complex applications of this method range from low-thrust orbit transfers, impul-

sive orbit transfers, ascent guidance, reentry trajectory design, spacecraft attitude control, among others [29].

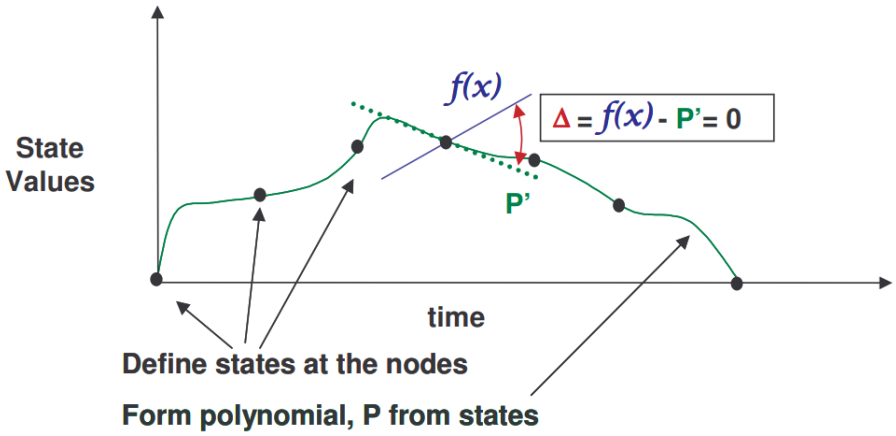


Figure 5: Pseudo-spectral procedure [30].

The basic idea behind a pseudospectral method is to build a high-order polynomial so that its time derivative values match the values of the system dynamics differential equations (state and control variable differential equations) at all collocation points across the entire time interval of the trajectory. By evaluating both the polynomial time derivatives and the physical time derivatives for a well-distributed representative number of discretisation nodes, it is possible to use numerical methods (Legendre-Gauss, Legendre-Gauss-Radau, Legendre-Gauss-Lobatto or Chebyshev-Gauss-Lobatto) to minimise the existing defects until a preset maximum tolerance limit is satisfied [28].

The major difference between direct collocation methods and pseudo-spectral resides in the fact that the first typically divides the trajectory into multiple segments and independently attempts to find a low-order polynomial that suits well with the system dynamics differential equations at the collocation points, facing the necessity of setting continuity constraints between segments and additional interior nodes within segments, whereas the latter is based on building a one segment high-order polynomial whose time derivatives match the system dynamics differential equations for all the collocation nodes, which suits well only for problems with smooth flight dynamics without significant function discontinuities [28].

Given that pseudospectral collocation methods are particularly powerful and highly efficient methods for solving continuous nonlinear constrained optimal control problems when compared to other direct collocation methods, these were the methods selected for the framework to solve the trajectory optimisation problem. Particularly, the high-order Gauss-Lobatto quadrature rules, as higher order polynomials offer an improved accuracy to the collocation method due to the finite precision, and, the number of parameters solved by the NLP problem is potentially lower in comparison to other lower order polynomials.

4 ROCKET DESIGN FRAMEWORK

4.1 MDO *Python* Libraries

In order to implement a multidisciplinary system for the current rocket design optimisation problem, it was necessary to search for an available software framework with the

following characteristics:

- ability for handling with a system with multiple coupled disciplines integrated with trajectory optimisation;
- support a wide range of optimisers so that a suitable option can be chosen according to specific optimisation requirements of the problem;
- an open-source framework with proven capabilities to handle the optimisation problem at hand;
- a modular environment for easier model construction;
- and, lastly, a good metadata and data handling capabilities for less advanced and non database specialised users.

We selected the *OpenMDAO Python* library [31] for multidisciplinary optimisation integrated with the *Dymos Python* library [32] for the trajectory optimisation end of the problem, thus producing a coupled approach to the preliminary rocket design.

4.1.1 OpenMDAO

OpenMDAO is an open-source object-oriented software framework crafted for multidisciplinary design, analysis and optimisation applications, programmed mainly in the *Python* language (for scripting convenience) and completely capable of interacting with other compiled languages, such as, SWIG, Cython, C and C++, among others.

Since it was first introduced for NASA's next-generation advanced single-aisle civil transport project in 2008 at the NASA Glenn Research Center (based in Cleveland, USA) [33], it has been under continuous development with several compelling use cases across a wide range of applications: from a Cubesat MDO problem for maximised data download capabilities [34], to a low-order aerostructural wing optimisation [35], to a structural topology optimisation [36], etc.

4.1.2 Dymos

Dymos is an open-source software tool built on top of the OpenMDAO framework designed to solve optimal control problems, such as trajectory optimisation. The combination of a framework built from an OpenMDAO optimisation architecture integrated with a Dymos trajectory optimisation opens the possibility to solve co-design optimisation problems with high computational efficiency even for complex use cases. The proposed framework will allow the implementation of a static system model within each optimisation cycle (a mass and sizing model, for example), which will receive new design variable values from the optimiser, and then sends the rocket sizing as its outputs (for example, the length and mass of the rocket) to a trajectory group capable of conducting all the necessary dynamic calculations through Ordinary Differential Equations (ODE) or Differential-Algebraic Equations (DAE) [32]. A standard architecture of this framework is shown in Fig. 6

In terms of trajectory optimisation processes, Dymos allows for the implementation of direct transcription methods, particularly, pseudospectral (high-order Gauss-Lobatto and Radau) [32].

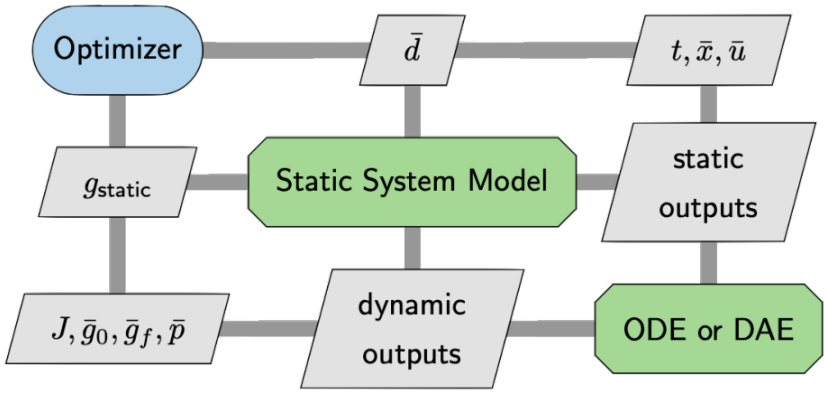


Figure 6: XDSM diagram of a standard coupled co-design problem, i.e., a MDO problem coupled with trajectory optimisation (OpenMDAO base framework integrated with Dymos) [32].

4.2 MDO Framework Implementation

In terms of the hierarchical structure of the framework, it was created a top-level group containing the optimiser and two other main groups: the mass and sizing group, which essentially is the mass and sizing model, and the trajectory group, which essentially is the trajectory model presented in Fig. 3. For each set of design variables x, directly handled by the optimiser, the mass and sizing generates a new rocket configuration, from which a few main parameters are fed within the trajectory model and a new objective function evaluation value is sent back to the optimiser.

At the trajectory level, the flight dynamics model handles four state variables (downrange x, altitude h, velocity v, pitch angle γ, and also their time derivatives, respectively, $\dot{x}$, $\dot{h}$, $\dot{v}$ and $\dot{\gamma}$. A remaining state variable time derivative, $\dot{m}$, is handled by the propulsion model.

These state variables are particularly important in the trajectory integration process because they mark the state values of the trajectory, i.e., the progress of the trajectory at each point in time, as well as the time progress of other models.

At the end of each trajectory simulation, the final altitude and the smallest difference between the critical stress of the body tube and the applied compressive stress are sent back to the optimiser, which does a constraint defect analysis, gradient evaluation, and a new iteration begins after a linesearch process.

Figure 7 illustrates the XDSM diagram of the MDO framework implementation.

5 ROCKET OPTIMAL DESIGN

5.1 Problem Formulation

In order to get a first assessment of the capabilities of the developed MDO framework in the context of a real problem, it is important to formally define it.

The chosen optimisation objective f was to minimise the rocket lift-off total mass subject to a minimum peak altitude constraint of 100 km, using the SLSQP optimisation method.

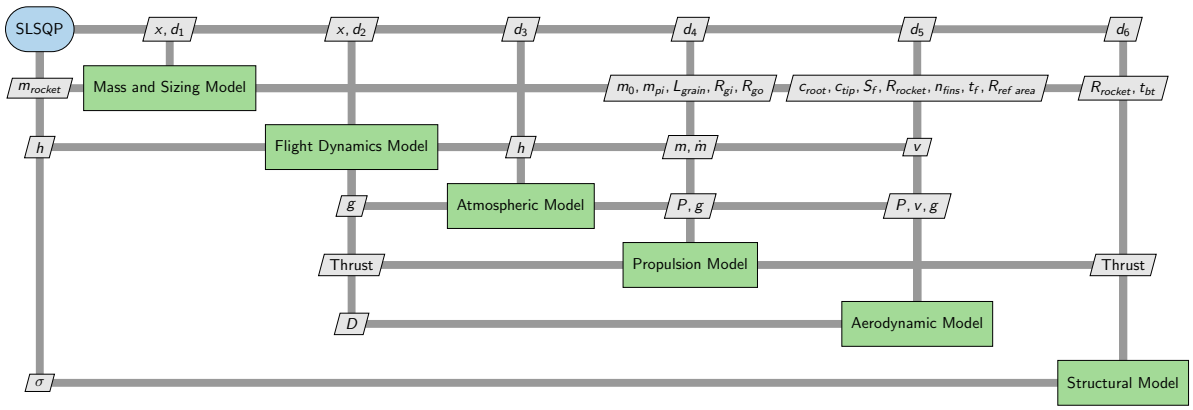


Figure 7: XDSM diagram of the framework highlighting the optimizer SLSQP (blue), the models (green) and design, coupled, local, and static variables (grey).

5.2 Parametric study of optimiser parameters

As the framework was thought for a quick preliminary rocket design application, it is of the greatest importance to use the optimisation setup which provides the most computationally cost-efficient solution. To that end, a parametric study on the impact of the optimiser tolerance level, as well as, the step size of the finite-difference method was conducted.

Generally speaking, it was observed that lower tolerance levels provide with more accurate results in higher computational costs. Reducing the tolerance levels from 10^{-3} to 10^{-6} while maintaining the step size, would require more function and gradient evaluations only to achieve slightly better results. Thus, the benefits from reducing the tolerance levels of the optimisation process were far outweighted by the increase in computational cost.

Regarding the impact of the step size on the optimisation process, it was observed that the system's convergence times were under 30 minutes for most cases although it could take several hours, depending on the quality of the initial design points and the chosen tolerance levels (using a computer operating Windows 10 with a 2.8 GHz, 4 cores processor, 16 GB of RAM and a 128 SSD storage unit).

This parametric study suggests that the best optimisation setting was to use a tolerance level of 10^{-3} combined with a step size of 10^{-3}, as this proved to be the most cost-efficient solution.

5.3 REXUS 2 Case Study

As a first case study, the optimisation goal was set to achieve the minimum lift-off mass for a suborbital flight with a minimum peak altitude of 100 km and a fixed payload of 98 kg, using 10 design variables, namely: the rocket diameter, overall length to diameter ratio, nosecone length to diameter ratio, body tube, SRM casing and propellant grain thicknesses, and the nozzle's expansion ratio, nozzle angle and nozzle convergent section angle. The main parameters of the found solution were then compared with several known masses and dimensions of the REXUS 2, an European suborbital single-staged solid propelled rocket [5], as portrayed in Table 1.

It was observed that these results represent an average percentual deviation of 26.6 % when compared to REXUS 2. Although a 27.7% decrease in total mass was observed,

Table 1: Comparison between the REXUS 2 and the optimised rocket configuration [5, 37].

Parameter	Unit	REXUS 2	Optimised Rocket	Deviation
Length	[m]	5.620	4.709	- 19.3%
Diameter	[m]	0.356	0.268	- 24.7%
Total Mass	[kg]	514.000	371.5	- 27.7%
Propellant Mass	[kg]	290.0	196.5	- 32.2%
Structural Mass	[kg]	126.0	77.0	- 38.9%
SRM Length	[m]	2.800	2.295	- 18.0%
Fin Root Chord	[m]	0.590	0.429	- 27.3%
Fin Tip Chord	[m]	0.400	0.300	- 25.0%

which is extremely positive, these results were treated with great caution as it is believed there might have been an oversimplification of the optimisation modules, as well as, underappreciation of several fixed masses (avionics and recovery system, for example). If adjusted to more realistic values, an increase in the propellant mass, structural mass and size of the optimised rocket would be observed, thus diminishing the percentage of total mass reduction.

In terms of the flight profile, a general agreement between both REXUS 2 and the optimised rocket was noticed, (please refer to Fig. 8) although the former sustains slightly higher altitudes in the first 100 seconds of the ascent phase which are evened out in the last 50 seconds. This flight profile discrepancies are the end result of the significant differences observed in terms of overall weight, size, and performance (velocity, drag, and thrust profiles, for example).

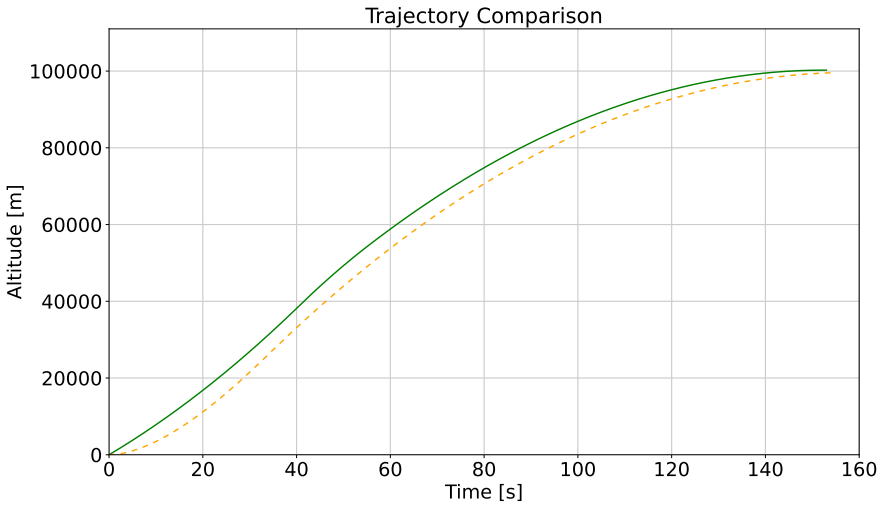


Figure 8: Flight profile comparison between both rockets. The green line represents the Rexus 2 flight profile and the orange dashed line the flight profile of the optimised rocket.

Finally, a 2D visual comparison of both rockets is presented in Fig. 9 where a few known measurements of the Rexus 2 are compared with the optimised rocket.

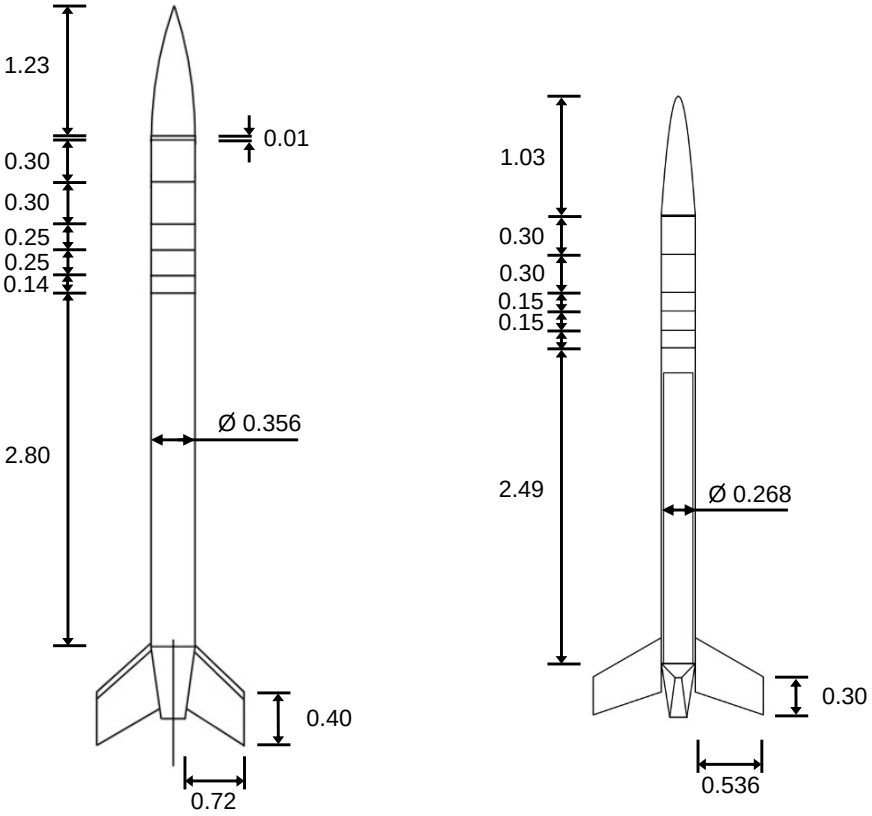


Figure 9: Visual comparison of the Rexus 2 (left) and the optimised rocket (right) highlighting the measurements of a few known parameters.

5.4 Payload and Altitude Sensitivity Analysis

In the Section 5.3, the optimisation behaviour of the framework was tested by comparing the obtained optimisation solution with a real rocket (Rexus 2). Ten design variables, an optimisation goal and specific mission requirements matching with those of the Rexus 2 were defined.

In this section, a sensitivity analysis of two parameters, payload and altitude, will be conducted in order to assess their impact in rocket design. Figure 10 shows how the payload and altitude were varied and four additional optimisation problems were studied.

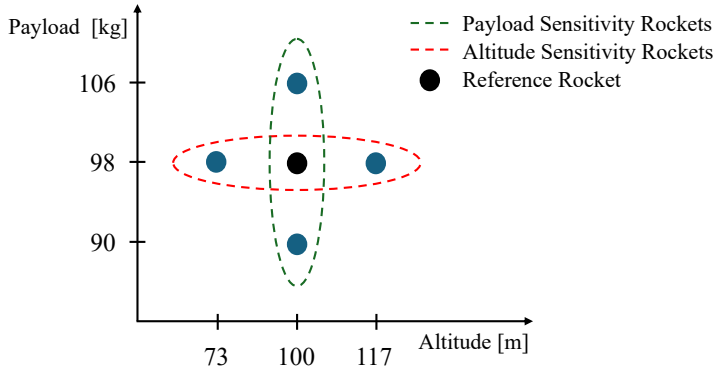


Figure 10: Payload and altitude sensitivity analysis rockets.

5.4.1 Altitude Sensitivity Analysis

In order to assess the altitude sensitivity, two new optimisation problems were created with different minimum peak altitudes, 73 and 117 km while maintaining a constant payload mass of 98 kg. The obtained solutions were then compared with the optimised rocket of the previous case study, which now served as a reference. Table 3 portrays a comparison of a few main parameters of these rockets:

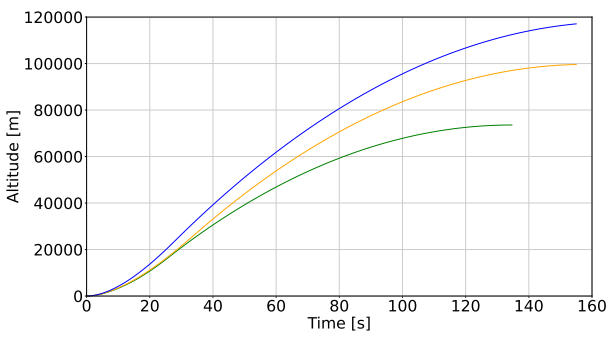
Table 2: Comparison between the altitude sensitivity optimised rockets.

Parameter	**Unit**	**73 km**	**Deviation**	**100 km**	**Deviation**	**117 km**
Length	[m]	4.647	- 1.3 %	4.709	+ 3.3 %	4.868
Diameter	[m]	0.290	+ 8.2 %	0.268	- 2.6 %	0.261
Total Mass	[kg]	353.2	- 4.9 %	371.5	+ 3.6 %	384.8
Payload Mass	[kg]	98.0	0.0 %	98.0	0.0 %	98.0
Propellant Mass	[kg]	178.2	- 9.3 %	196.5	+ 5.0 %	207.1
Structural Mass	[kg]	76.9	- 7.7 %	77.0	+ 3.4 %	79.6
SRM Length	[m]	2.030	-11.6 %	2.295	+ 12.6 %	2.584

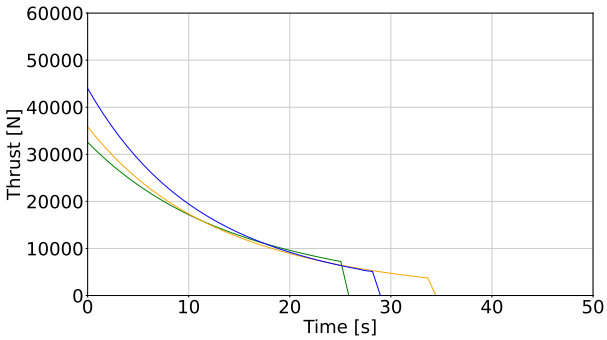
From Table 3, it can be seen that by reducing the minimum peak altitude from 100 km to 73 km while maintaining the same payload, it was possible to reduce the overall mass of the rocket in 5.0 %, the propellant mass in 9.3 %, and also, the structural mass in 7.7 %. On the other hand, increasing the target altitude to 117 km, demanded increasing the propellant mass in 5 %, the structural mass in 3.4 %, and consequently, the total mass in 3.6 %.

Furthermore, this data confirms the patterns expected from a purely theoretical analysis. As higher target altitudes are set, the propellant mass needed increases, thus increasing the SRM length and the overall rocket length. Additionally, the loading on the structure also tends to increase which leads to an increase in the structural mass of the rocket and, ultimately, an increase in the total lift-off mass.

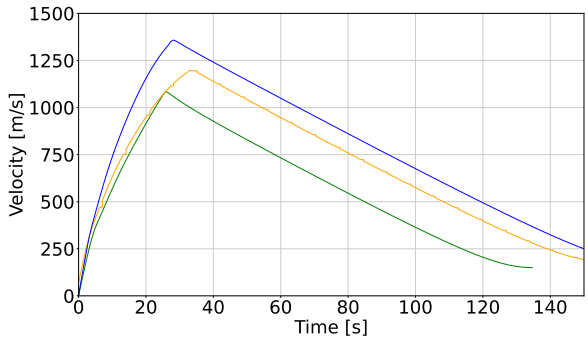
In terms of the flight profile comparison, it was possible to observe that in order to achieve higher altitudes the thrust profile increases in magnitude, as higher velocities are required.



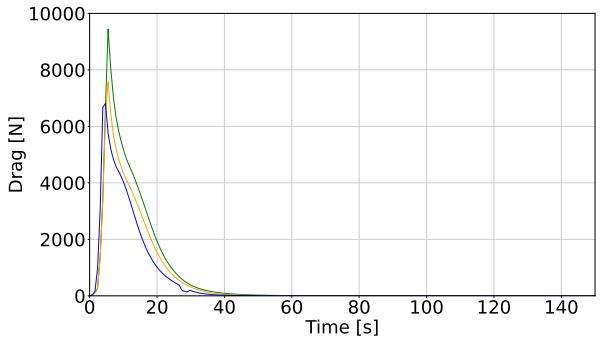
(a) Altitude profile comparison.



(b) Thrust profile comparison.



(c) Velocity profile comparison.



(d) Drag profile comparison.

Figure 11: Flight data analysis of four main parameters: altitude, thrust, velocity and drag. The green, orange and blue lines represent the rockets carrying 98 kg of payload aimed at 73, 100 and 117 km minimum peak altitudes, respectively

On the other hand, under increased velocities, the effects of the aerodynamic loading on the structure, namely, the drag become ever more prominent, for which the optimiser tends to reduce the cross-sectional area of the rocket as much as possible by reducing its diameter.

5.4.2 Payload Sensitivity Analysis

In order to assess the payload sensitivity, another two new optimisation problems were created with different payloads, 90 and 106 kg while maintaining a constant minimum peak altitude, fixed at 100 km. Similarly to the altitude sensitivity analysis, the obtained solutions were then compared with the optimised rocket from the first case study, which served as a reference. Table 3 presents a comparison of a few main parameters of these rockets.

Table 3: Comparison between the payload sensitivity optimised rockets.

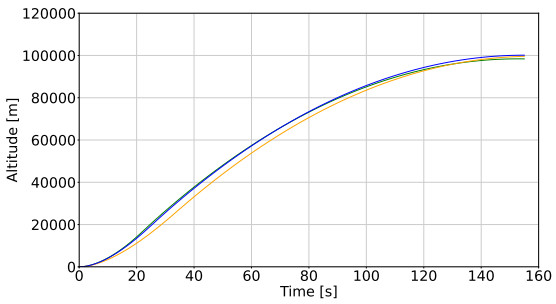
Parameter	**Unit**	**90 kg**	**Deviation**	**98 kg**	**Deviation**	**106 kg**
Length	[m]	4.689	- 0.4 %	4.709	+ 2.709 %	4.846
Diameter	[m]	0.261	-2.6 %	0.268	- 2.2 %	0.262
Total Mass	[kg]	347.6	- 6.4 %	371.5	+ 4.8 %	389.5
Payload Mass	[kg]	90.0	0.0 %	98.0	0.0 %	106.0
Propellant Mass	[kg]	180.7	- 8.1 %	196.5	+ 3.3 %	203.1
Structural Mass	[kg]	76.9	- 0.1 %	77.0	+ 4.4 %	80.4
SRM Length	[m]	2.423	+ 5.6 %	2.295	+ 11.7 %	2.563

This results show that a payload reduction of 8 kg translates in 6.4 % decrease in total mass while a payload increase in the same amount translates in an increase of 4.8 %. The observed variations in total mass are the end result of the partial variations of propellant mass and structural mass.

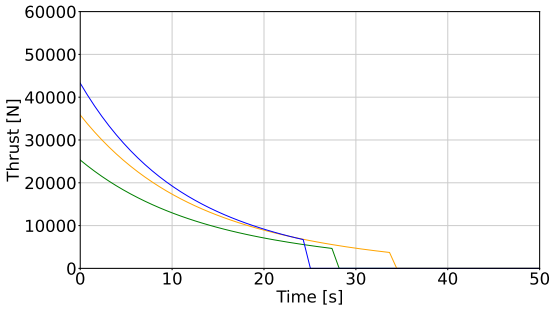
Looking at the flight data, presented in Fig. 12, it is possible to observe significant

differences in the thrust and velocity profiles. As higher payload is carried, more propellant is needed in order to achieve high enough velocities in the early stages of the ascent phase to free the rocket from the atmosphere and reach the desired target altitude.

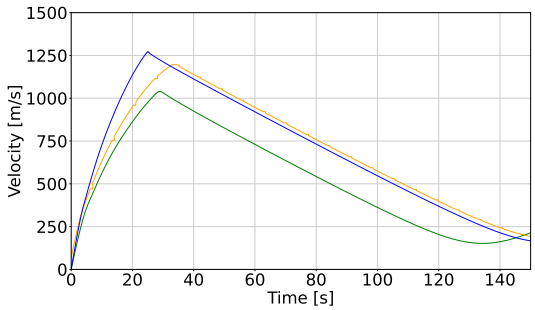
Furthermore, higher velocities involve higher accelerations, higher propellant burn rates and, consequently, higher thrusts which lead to the design of rockets with increased thicknesses and, thus, increased structural mass in order to withstand higher aerodynamic loading conditions.



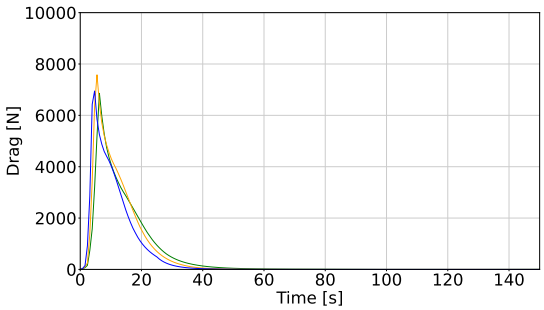
(a) Altitude profile comparison.



(b) Thrust profile comparison.



(c) Velocity profile comparison.



(d) Drag profile comparison.

Figure 12: Flight data analysis of four main parameters: altitude, thrust, velocity and drag. The green, orange and blue lines represent the rockets aimed at 100 km minimum peak altitude carrying payload of 90, 98 and 106 kgs, respectively.

Finally, a 2D visual comparison of all the optimised rockets obtained in this sensitivity analysis is presented in Fig. 13.

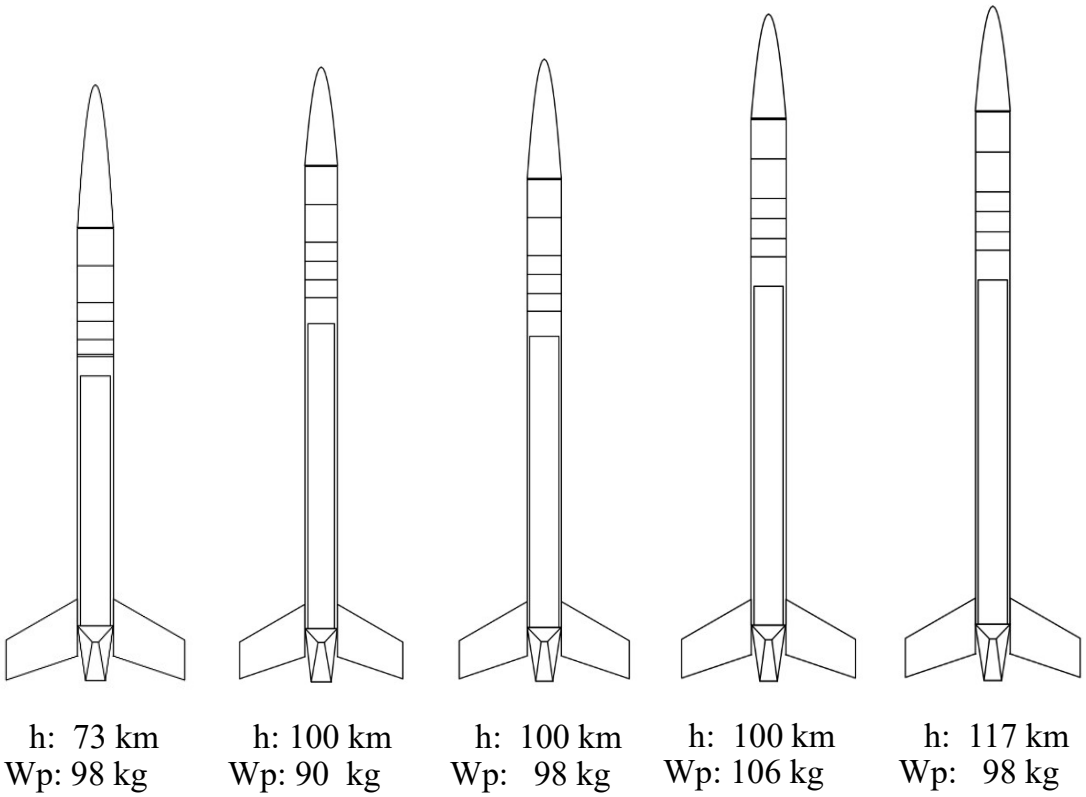


Figure 13: Rocket 2D final comparison.

6 CONCLUSIONS

In this work, a low computational cost multidisciplinary optimisation (MDO) framework capable of solving co-design optimization problems in the context of preliminary design of single-stage solid propellant rockets was developed. Six disciplinary models were successfully developed and integrated within an MDF architecture. As for the optimizer, a gradient-based SLSQP optimization algorithm was selected and successfully integrated in the framework. In addition, the Gauss-Lobatto pseudospectral method was selected to solve the trajectory optimisation problem.

The developed framework underwent several tests in order to assess its optimisation capabilities. First, parametric studies of two main optimisation parameters, the tolerance level of the SLSQP method and the step size of the finite difference, were conducted. From this tests, it was concluded that the best setup for optimisation was using a tolerance level of 10^{-3} and a step size of the finite difference method of 10^{-3}.

An initial case study was conducted to assess the accuracy of the framework against a real rocket. The results showed a 27.7 % total mass reduction which might indicate that the optimisation modules might have been oversimplified and some fixed masses, such as, avionics, recovery system, among others, might have been underappreciated. If adjusted a smaller, more realistic total mass reduction is expected.

Afterwards, a sensitivity analysis allowed to conclude that the payload and minimum altitude greatly influence the behaviour of the optimisation process, with the results showing that reducing the minimum peak altitude from 100 km to 73 km allows for a 4.9 % total mass reduction and increasing it from 100 km to 117 km leads to a 3.6 % increase.

Furthermore, it was observed that varying the payload mass from 98 kg to 90 kg allowed for a 6.4 % total mass reduction against a 4.8 % increase when varied from 98 kg to 106 kg.

Overall, the developed framework shows good signs of being capable of performing the design optimisation of a single stage sounding rocket at a preliminary level. Given its great modularity, a straightforward extension to a larger spectrum of applications is expected, such as multi-stage or liquid-propellant, upon additional development.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the full support of IDMEC and also Academia da Força Aérea throughout the development of this work. A. C. Marta and P. J. S. Gil acknowledge Fundação para a Ciência e a Tecnologia (FCT) for its financial support via the project LAETA Base Funding (DOI: 10.5449/UIDB/50022/2020).

REFERENCES

[1] G. Genta. Private space exploration: A new way for starting a spacefaring society? *Acta Astronautica*, 104(2):480–486, nov 2014. ISSN 0094-5765. DOI:10.1016/j.actaastro.2014.04.008.

[2] S. F. E. Team. Space Foundation Releases the space report 2023 Q2, showing annual growth of global space economy to $546B, 2023. Available: `https://www.spacefoundation.org/2023/07/25/the-space-report-2023-q2/`.

[3] G. Gurney. *Space Technology Spinoffs*. An Impact book. Watts, 1979. ISBN 9780531022900.

[4] A. Papageorgiou and J. Olvander. The role of multidisciplinary design optimization (MDO) in the development process of complex engineering products. In *Proceedings of the 21st International Conference on Engineering Design (ICED17)*, volume 4, 2017.

[5] M. Persson, O. & Hörschgen. REXUS 2 - the first EuroLaunch project. In *17th ESA Symposium on European Rocket and Balloon Programmes and Related Research*, pages 389 – 393, Sandefjord, Norway, 30 May - 2 June 2005. Swedish Space Corporation (SSC) & Deutsches Zentrum für Luft- und Raumfahrt (DLR), ESA Publications Division. ISBN 92-9092-901-4.

[6] B. C. M. H. H. Box, S. Estimating the dynamic and aerodynamic parameters of passively controlled high power rockets for flight simulation. *Cambridge Rocketry*, 2009. Available: `https://cambridgerocket.sourceforge.net/AerodynamicCoefficients.pdf`.

[7] L. Pepermans, E. Menting, M. Rozemeijer, B. Koops, N. S. Dahl, N. Suard, S. khurana, F. V. Marion, F. Kuhnert, and M. Serman. Comparison of Various Parachute Deployment Systems for Full Rocket Recovery of Sounding Rockets. In *Proceedings of the 8th European Conference for Aeronautics and Space Sciences. Madrid, Spain, 1-4 july 2019*, 2019. DOI:10.13009/EUCASS2019-411.

[8] J. M. Tizó. Burnback Analysis of Solid Propellant Rocket Motors. *arXiv preprint*, 2023. DOI:10.48550/ARXIV.2301.04559.

[9] R. Hartfield, R. Jenkins, J. Burkhalter, and W. Foster. A Review of Analytical Methods for Solid Rocket Motor Grain Analysis. In *39th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*. American Institute of Aeronautics and Astronautics, June 2003. DOI:10.2514/6.2003-4506.

[10] W. Miller. *Solid Rocket Motor Performance Analysis and Prediction*. NASA SP 8039. National Aeronautics and Space Administration, May 1971.

[11] Brooks, W.T. *Solid Propellant Grain Design and Internal Ballistics*. NASA SP 8076. National Aeronautics and Space Administration, March 1972.

[12] G. Henry, R. Humble, W. Larson, U. S. D. of Defense, U. S. N. Aeronautics, and S. Administration. *Space Propulsion Analysis and Design*. College custom series. McGraw-Hill, 1995. ISBN 9780070313293.

[13] K. Kuo and M. Summerfield. *Fundamentals of Solid-propellant Combustion*. Progress in astronautics and aeronautics. American Institute of Aeronautics and Astronautics, 1984. ISBN 9780915928842.

[14] D. J. Martin. Summary Of Flutter Experiences As A Guide To The Preliminary Design of Lifting Surfaces On Missiles. Report No. 685, NACA, Langley Aeronautical Laboratory Langley Field, Va, 1940.

[15] United States. National Oceanic and Atmospheric Administration and United States Committee on Extension to the Standard Atmosphere. *U.S. Standard Atmosphere, 1976*. NOAA - SIT 76-1562. National Oceanic and Amospheric Administration, 1976.

[16] United States Committee on Extension to the Standard Atmosphere and United States. National Aeronautics and Space Administration. *U.S. Standard Atmosphere, 1962: ICAO Standard Atmosphere to 20 Kilometers; Proposed ICAO Extension to 32 Kilometers; Tables and Data to 700 Kilometers*. U.S. Government Printing Office, 1962.

[17] H. Akima. A New Method of Interpolation and Smooth Curve Fitting Based on Local Procedures. *Journal of the ACM*, 17(4):589–602, Oct. 1970. ISSN 1557-735X. DOI:10.1145/321607.321609.

[18] Akima1DInterpolator - SciPy v1.14.1 Manual. Available: `https://docs.scipy.org/doc/scipy/reference/generated/scipy.interpolate.Akima1DInterpolator.html`.

[19] P. Sforza. *Theory of Aerospace Propulsion*. Aerospace Engineering. Elsevier Science, 2016. ISBN 9780128096017. DOI:10.1016/c2009-0-61051-5.

[20] P. Dépincé, B. Guédas, and J. Picard. Multidisciplinary and multiobjective optimization: Comparison of several methodss. In *7th World Congress on Structural and Multidisciplinary Optimization)*, 2007.

[21] E. J. Cramer, J. E. Dennis, Jr., P. D. Frank, R. M. Lewis, and G. R. Shubin. Problem Formulation for Multidisciplinary Optimization. *SIAM Journal on Optimization*, 4 (4):754–776, Nov. 1994. DOI:10.1137/0804044.

[22] N. Tedford and J. R. R. A. Martins. On the Common Structure of MDO Problems: A Comparison of Architectures. In *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. American Institute of Aeronautics and Astronautics, June 2006. DOI:10.2514/6.2006-7080.

[23] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006. ISBN 9780387400655.

[24] D. Kraft. *A Software Package for Sequential Quadratic Programming*. Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt Köln: Forschungsbericht. Wiss. Berichtswesen d. DFVLR, 1988.

[25] M. P. Kelly. Transcription Methods for Trajectory Optimization: a beginners tutorial, 2017. DOI:10.48550/arxiv.1707.00284.

[26] O. von Stryk and R. Bulirsch. Direct and indirect methods for trajectory optimization. *Annals of Operations Research*, 37(1):357–373, Dec. 1992. ISSN 1572-9338. DOI:10.1007/bf02071065.

[27] A. Rao, D. Benson, G. Huntington, B. Origin, L. Seattle, C. Wa, C. Francolin, M. Darby, and M. Patterson. User's Manual for GPOPS Version 2.1: A MATLAB R Package for Dynamic Optimization Using the Gauss Pseudospectral Method. *Gainesville, FL, USA*, 2009. Available: `https://gpops2.com/resources/gpops2UsersGuide.pdf`.

[28] F. Topputo and C. Zhang. Survey of Direct Transcription for Low-Thrust Space Trajectory Optimization with Applications. *Abstract and Applied Analysis*, 2014: 1–15, 2014. ISSN 1687-0409. DOI:10.1155/2014/851720.

[29] I. M. Ross and F. Fahroo. Pseudospectral Knotting Methods for Solving Nonsmooth Optimal Control Problems. *Journal of Guidance, Control, and Dynamics*, 27(3): 397–405, May 2004. ISSN 1533-3884. DOI:10.2514/1.3426.

[30] S. Paris, J. Riehl, and W. Sjauw. Enhanced procedures for direct trajectory optimization using nonlinear programming and implicit integration. In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. American Institute of Aeronautics and Astronautics, June 2006. DOI:10.2514/6.2006-6309.

[31] J. S. Gray, J. T. Hwang, J. R. R. A. Martins, K. T. Moore, and B. A. Naylor. OpenMDAO: an open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, Mar. 2019. ISSN 1615-1488. DOI:10.1007/s00158-019-02211-z.

[32] R. Falck, J. Gray, K. Ponnapalli, and T. Wright. Dymos: A Python package for optimal control of multidisciplinary systems. *Journal of Open Source Software*, 6 (59):2809, Mar. 2021. ISSN 2475-9066. DOI:10.21105/joss.02809.

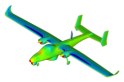
[33] C. Heath and J. Gray. OpenMDAO: Framework for Flexible Multidisciplinary Design, Analysis and Optimization Methods. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 2012*, 04 2012. ISBN 978-1-60086-937-2. DOI:10.2514/6.2012-1673.

[34] J. T. Hwang, D. Y. Lee, J. W. Cutler, and J. R. R. A. Martins. Large-Scale Multidisciplinary Optimization of a Small Satellite's Design and Operation. *Journal of Spacecraft and Rockets*, 51(5):1648–1663, Sept. 2014. ISSN 1533-6794. DOI:10.2514/1.a32751.

[35] J. P. Jasa, J. T. Hwang, and J. R. R. A. Martins. Open-source coupled aerostructural optimization using Python. *Structural and Multidisciplinary Optimization*, 57(4): 1815–1827, Feb. 2018. ISSN 1615-1488. DOI:10.1007/s00158-018-1912-8.

[36] H. Chung, J. T. Hwang, J. S. Gray, and H. A. Kim. Implementation of topology optimization using openMDAO. In *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. American Institute of Aeronautics and Astronautics, Jan. 2018. DOI:10.2514/6.2018-0653.

[37] M. S. *Rexus User Manual: Rocket Experiments for University Students*. Number 7.11. EuroLaunch, JAN 2014. Available: `https://www.zarm.uni-bremen.de/fileadmin/user_upload/teaching/study_programs/rexus_bexus/RX_UserManual_v7-11_08Jan14.pdf`.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

A STUDY ON THE MULTIDISCIPLINARY DESIGN OPTIMIZATION OF HYBRID-ROCKET PROPULSION

Joana Matos[1*], Frederico Afonso[1], Girolamo Musso[1], Iara Figueiras[1], Alain Souza[1], Inês Ribeiro[1], Fernando Lau[1] and Joseph Morlier[2]

1: IDMEC, Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
{joanafmatos, frederico.afonso, girolamo.musso, iara.figueiras, alain.souza, ines.ribeiro, lau}@tecnico.ulisboa.pt

2: ICA, Université de Toulouse, ISAE-SUPAERO, INSA, CNRS, MINES, ALBI, UPS, Toulouse, France
joseph.morlier@isae-supaero.fr

Abstract. *Hybrid rocket engines are considered a promising option for reducing the environmental impact of microsatellite launchers. They offer higher safety than solid propulsion systems and better manufacturing conditions compared to liquid propulsion. However, their lower regression rate and combustion efficiency have hindered their widespread adoption. This paper aims to explore different types of potentially printable fuels combined with N_2O to maximize performance while minimizing environmental impact, all within the context of a predefined sounding rocket geometry. To achieve this, a Multi-Objective Optimization (MOO) approach is developed using an in-house framework created for launcher design. The design variables include the volume of oxidizer tank, the initial oxidizer to fuel mass ratio, total mass of propellant (fuel and oxidizer), throat diameter, nozzle area ratio (expansion ratio), and effective injector area. This approach integrates propulsion, thermochemical equilibrium, initial sizing of the propulsion system, and a Life Cycle Assessment (LCA) that considers both production and use phases.*

Keywords: Eco-design, hybrid-rocket propulsion, printable fuels, life cycle assessment, multi-objective optimization

1 INTRODUCTION

The accelerated pace of satellite launches [1] has contributed to growing environmental concerns. To address this issue, alternative propulsion systems are being explored, namely hybrid rocket engines [2, 3]. They offer improved safety over solid propulsion systems and more favorable manufacturing conditions compared to liquid engines [4].

Although hybrid rocket engines traditionally present challenges such as lower regression rates and combustion instabilities, which reduce their combustion efficiency compared to both solid and liquid engines [5], these drawbacks can be addressed through careful engine design, as highlighted by Wei et al. [4].

A promising approach to achieve this goal is through numerical optimization, which can identify the optimal design parameters that yield engine configurations with maximized specific impulse, thereby enhancing combustion efficiency. At the same time, environmental impact metrics should be integrated alongside propulsion performance indicators to guide the design of hybrid engines that balance these often conflicting objectives.

While there are numerous studies on single-objective optimization in hybrid rocket engines in the literature, primarily focused on maximizing propulsion performance [6–8], as well as studies that integrate propulsion performance with additional disciplinary analyses through Multidisciplinary Design Optimization (MDO) [9–12], there is limited research addressing Multi-Objective Optimization (MOO) that simultaneously considers both performance and environmental goals [13, 14]. Casalino et al. [13] estimated the Global Warming Potential (GWP) of four combustion byproducts (namely CO_2, water vapor, black carbon, and submicrometer alumina particles) using emission indices along a vehicle trajectory. In contrast, Musso et al. [14] applied the ReCiPe 2016 method to evaluate environmental impacts, accounting for both the production and consumption of the fuel and oxidizer.

With this in mind, the current paper builds upon previous work by the research team [14] and explores the use of different fuels for hybrid rocket engine design, addressing both propulsion performance and environmental considerations through multiple Life Cycle Assessment (LCA) metrics, which are increasingly gaining attention within the space community [15–20]. In recent years, there has been growing interest in fuels other than paraffin-based ones [21], particularly those that can be additively manufactured [22]. Although this manufacturing process is not considered in this paper, the fuels explored could potentially be manufactured using this technique.

2 METHODOLOGY

The application case of this paper is a hybrid rocket engine for small sounding rockets, with the goal of maximizing specific impulse (I_{SP}) and minimizing Global Warming Potential (GWP), Ozone Depletion (ODEPL), Abiotic Resource Depletion - fossil fuels (ADEPLf), and a Single Score metric. The main parameters of the remainder of the sounding rocket are described in a previous article by the research group [14], which also provides detailed descriptions of the two models employed in the current paper: one for propulsion and another for LCA, although the latter has suffered some modifications. The propulsion model is a physics-based thermodynamic framework coupled with NASA's CEA software [23], which estimates I_{SP} by computing the thermochemical equilibrium. The LCA model uses pre-generated lookup tables from SimaPro® [24], utilizing the EcoInvent and ESA LCA Databases. These models are integrated using a sequential

optimization architecture.

2.1 Propulsion Model

The propulsion model [14] is a steady-state, zero-dimensional framework that includes several key assumptions. It models a self-pressurized oxidizer tank, where saturation is assumed before liquid depletion [25], while the Span–Wagner equation of state [26] is used only in the vapor regime to account for real gas behavior. Combustion is treated under ideal gas and thermochemical equilibrium assumptions, with a propellant-specific regression rate formulation. NASA's CEA software [23] is used to evaluate combustion performance. The nozzle is modeled as an ideal isentropic, convergent-divergent nozzle with choked flow at the throat. An illustration of the hybrid rocket engine is shown in Figure 1.

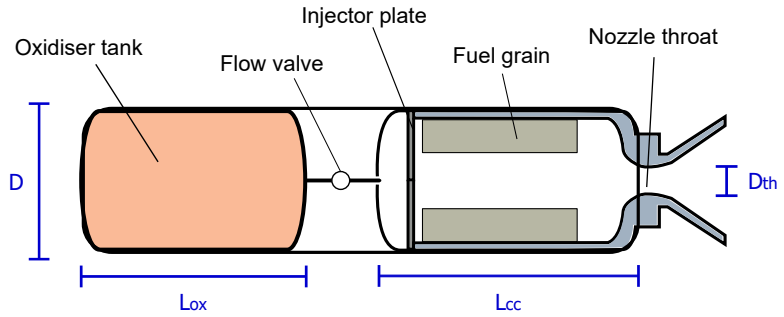


Figure 1: Illustration of the hybrid rocket engine.

In this study, N_2O is used as the oxidizer to maintain a self-pressurized tank, combined with various fuel grains, namely Paraffin wax, Hydroxyl-Terminated PolyButadiene (HTPB), High-Density Polyethylene (HDPE), and Aluminized Paraffin (ParaffinAl8020). Table 1 contains the chemical formulas and some other properties of the fuels tested in hybrid propulsion systems using N_2O as the oxidizer.

Table 1: Fuels input parameters: density (ρ) and ballistic coefficients (a and n) when combined with N_2O.

Fuel	Chemical Formula	ρ [kg/m^3]	a [mm/s]	n	Sources
Paraffin	$C_{32}H_{66}$	930	0.155	0.500	[27]
HTPB	$C_{654}H_{981}O_6$	930	0.098	0.430	[28][29]
HDPE	C_2H_4	948	0.116	0.331	[30][31]
Aluminized Paraffin	$C_{32}H_{66}$+Al(87.5%)80:20	1069	0.175	0.500	[27][32]

The diameter D and the combined length of the oxidizer tank L_{ox} and combustion chamber L_{cc} are kept constant throughout the optimization process. The nozzle throat diameter D_{th} and expansion ratio A_e/A_{th} are allowed to vary, along with the total propellant mass m_p, the starting oxidizer-to-fuel ratio (O/F), volume of the oxidizer tank V_{OT}, and the effective injector area $A_{inj}C_{D_{inj}}$. The optimization process and the implementation of these variables are discussed later in Section 2.3.

As outputs of this model, we obtain the specific impulse I_{SP}, a vector with the mass of each emitted species $\mathbf{m_{es}}$ produced during combustion, the mass of oxidizer m_o and fuel

m_f, the pressure at the combustion chamber p_{cc} and oxidizer tank p_o, and the maximum oxidizer mass flux G_m. Among the thirty-four different collected chemical species emitted in the combustion, the ones found most relevant to the environmental impact are Al_2O_3, $Al(OH)_3$, CO, CO_2, C, H_2, H_2O, HO_2, NO, NO_2, and OH.

2.2 LCA Model

The goal of the implemented LCA model is to assess the environmental impact of different hybrid rocket propellants combinations, considering as system boundaries both production and use phases, more specifically, the burning of different propellant and the corresponding emissions into the atmosphere. These elements are directly linked to the inputs (design variables) and outputs of the propulsion model, namely the O/F, the mass of fuel m_f, the mass of oxidizer m_o (through O/F and m_p), and the mass of each emitted specie during combustion $\mathbf{m_{es}}$. The latter are calculated by interpolating CEA Software data to obtain the corresponding mass fractions, which are subsequently multiplied by the total exhaust mass. The usage of the propulsion system until either the oxidizer or fuel (or both) runs out is set as functional unit.

As already mentioned, the EcoInvent and ESA LCA Databases are also used as Background Life Cycle Inventory (LCI) data, and for some of the chemical species emitted, characterization factors for aviation and space are used instead [33–35]. When no data is available, databases proxies are considered. Environmental impacts are quantified on a per-unit-mass basis across multiple categories using the SimaPro® [24] software, in order to be integrated into the framework. Additionally, for the Life Cycle Impact Assessment (LCIA), a method contained in the ESA LCA Database, which is an adaptation of the Environmental Footprint 3.0 method, is employed to support the comparative assessment between propellant alternatives. The resulting metrics are used to generate lookup tables, which are then resorted during the optimization process to evaluate the environmental footprint of each candidate hybrid propellant combination and engine design. The metrics studied include GWP100 and ADEPL, as they are particularly relevant to propellant production; ODEPL, as it is particularly relevant to combustion emissions; and a Single Score metric, as defined in the aforementioned method, which enables multiple impact categories to be considered simultaneously.

A flowchart of the LCA model is illustrated in Figure 2.

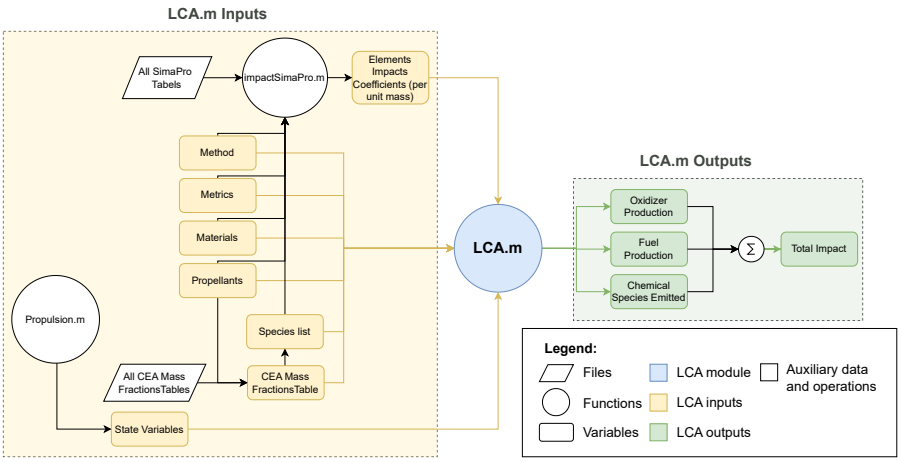


Figure 2: Flowchart of the LCA model.

2.3 MOO Framework

The models are integrated using a sequential MOO architecture, as depicted in Figure 3 following a XDSM diagram [36]. In this process, the design variables $\mathbf{x} = \{\mathrm{O/F}, m_p, V_{OT}, D_{th}, A_e/A_{th}, A_{inj}C_{D_{inj}}\}$ are first passed to the propulsion model to compute a vector of propulsion state variables $\mathbf{u_{prop}}$. This vector includes I_{SP}, $\mathbf{m_{es}}$, m_o, m_f, p_{cc}, p_o, and G_m. Subsequently, the x vector is passed to the LCA model, along with the $\mathbf{u_{prop}}$ vector, to estimate the environmental impacts (GWP, OD, ADEPLf, and Single Score metric), which form a vector of LCA state variables $\mathbf{u_{LCA}}$. While both state vectors contribute to the objective function f, only the propulsion-related state vector is used to define the constraints $\mathbf{g}$.

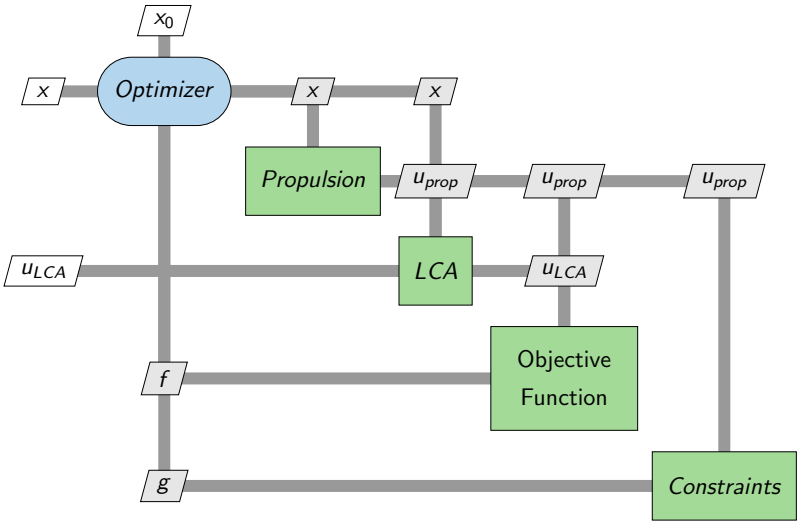


Figure 3: XDSM diagram of the MOO architecture.

Reminding the simultaneous goal of maximizing I_{SP} and minimizing different LCA metrics (Metric), the MOO problem can be mathematically posed as:

$$\begin{aligned} \text{Minimize} \quad & f(\mathbf{x}) = -\alpha\ I_{SP} + (1-\alpha)\,\text{Metric}, \quad \alpha \in [0,1] \\ \text{w.r.t.} \quad & \mathbf{x} \\ \text{subject to} \quad & g_1(\mathbf{x}) = p_{cc} - 0.8\ (p_o - \Delta P_{inj}) \leq 0 \\ & g_2(\mathbf{x}) = G_m - 650 \leq 0\,. \end{aligned} \tag{1}$$

To maintain a feasible engine, the first constraint g_1 is set to prevent back-flow while the second one g_2 limits the oxidizer mass flux to ensure a stable combustion. The weight α will be swept during the optimization by the NSGA-II [37] algorithm chosen for solving the problem, namely the implementation available in [38]. Based on experience a population size of 500, 100 generations, and a mutation rate of 30% were chosen. The lower and upper boundaries established for the design variables are shown in Table 2.

3 RESULTS

The MOO problem was run for each fuel and environmental metric to generate Pareto fronts, enabling a comprehensive comparison between different hybrid rocket engines. The

Table 2: Upper and lower boundaries of the design variables.

Design Variable	Upper Boundary	Lower Boundary	Units
V_{OT}	0.06	5×10^{-6}	m^3
O/F	22.5	0.5	-
m_p	20.298	5×10^{-6}	kg
D_{th}	0.11712	0.01000	m
A_e/A_{th}	15.0	1.5	-
$A_{inj}C_{D_{inj}}$	6×10^{-5}	5×10^{-6}	m^2

resulting Pareto fronts are shown in Figure 4, where it is evident that HDPE consistently yields larger maximum specific impulses and lower environmental impacts for most of the designs along the Pareto front. However, Paraffin outperforms HDPE in terms of ozone depletion minimization.

Moreover, when the NSGA-II prioritizes the reduction of environmental metrics, the comparison becomes less straightforward, with significantly reduced propellant masses and design variables such as D_{th} and $A_{inj}C_{D_{inj}}$ approaching or even reaching their lower boundaries. To quantify the relative differences between designs for each environmental metric, we selected a compromise between objectives: the minimization of the squared root of the sum of the normalized objectives, as used in [14], alongside the maximum specific impulse.

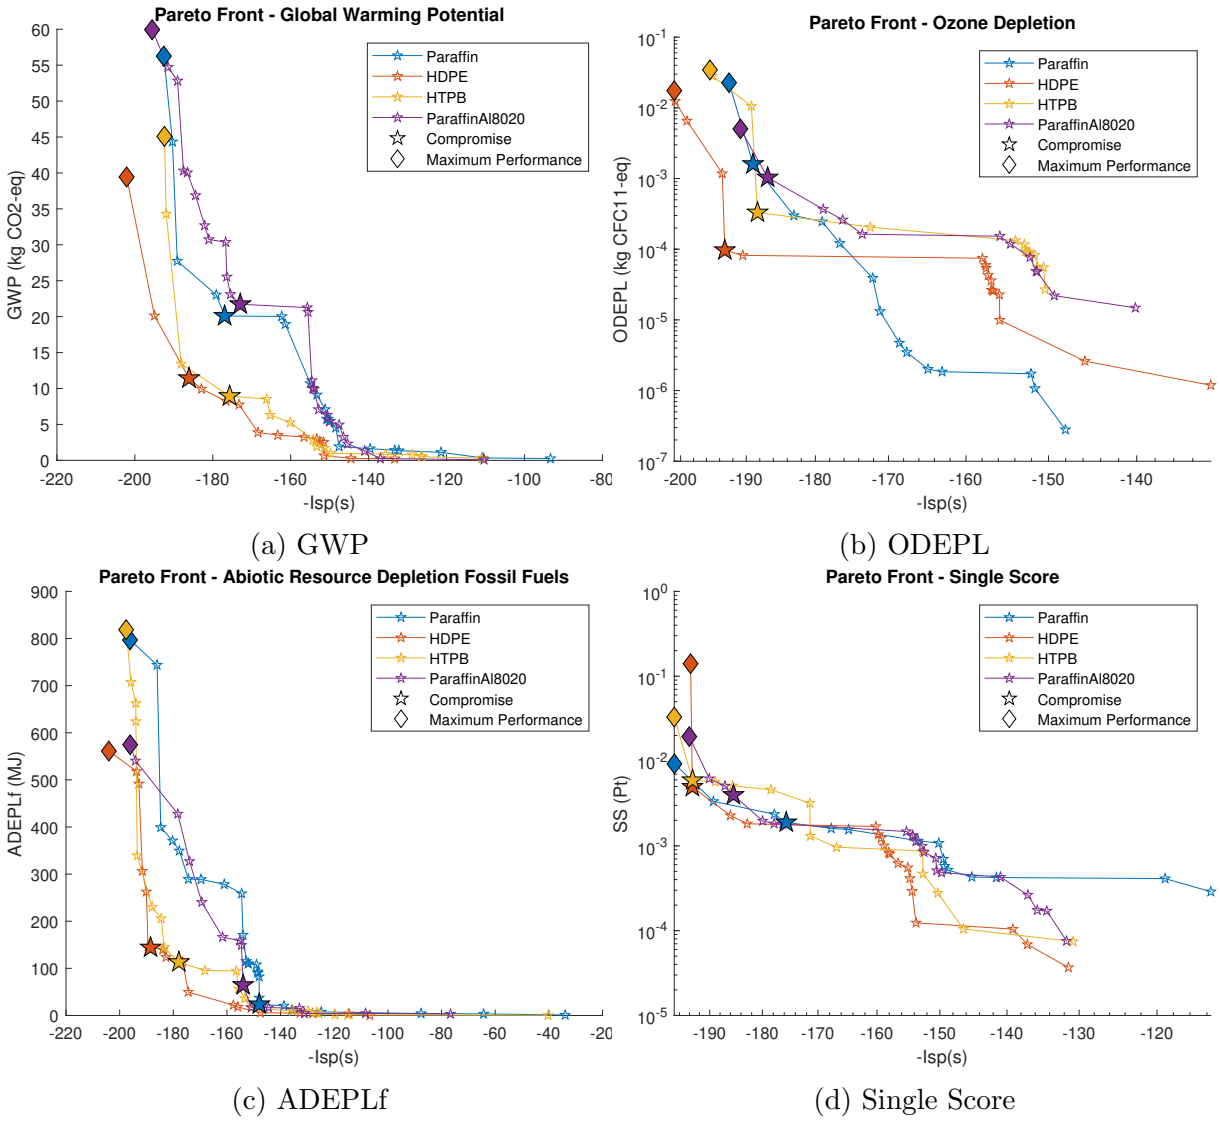
Only the results for the compromise and maximum performance points will be presented. This choice is justified by the fact that the results become less meaningful for minimum impact points, as the mass of the propellant is considerably reduced, ranging from approximately 100 to 600 g, along with other notably low variables, such as burning time, A_e/A_{th}, and $A_{inj}C_{D_{inj}}$, which severely compromise system performance, resulting in unfeasible designs. The maximum performance point is also included, since the concept of Eco-Design often aims to minimize environmental impacts without compromising performance.

The results for each of these points and for each metric are shown in Tables 3, 4, 5, and 6. The best results obtained for each objective, i.e. the lowest environmental impacts and the highest specific impulse I_{sp}, are highlighted in bold.

Starting with GWP (for a 100-year time frame), HDPE provides a 3.3% higher I_{sp} than the second-best performing fuel in terms of I_{sp}, Aluminized Paraffin, while simultaneously reducing the GWP by 34.2% (as shown in Table 3). However, when selecting the compromise design, HTPB performs 21.3% better in terms of environmental impact compared to HDPE, with a slight reduction of 5.6% in I_{sp}.

Regarding ODEPL, as previously mentioned, Aluminized Paraffin achieves a 71.4% lower ozone depletion impact than HDPE, with only a 5% reduction in I_{sp}. However, when selecting a compromise design, HDPE yields the highest I_{sp} while also having a lower ozone depletion impact, as shown in Table 4.

For ADEPLf (Table 5), HDPE again offers the best performance when maximizing specific impulse, with a 3.3% higher I_{sp} than the second-best fuel (HTPB) and a 2.3% lower ADEPLf impact than the second-best solution (ParaffinAl8020). However, when choosing a compromise design, Paraffin achieves an 83.5% reduction in ADEPLf impact, at the cost of a 21.5% decrease in I_{sp}.



(a) GWP (b) ODEPL

(c) ADEPLf (d) Single Score

Figure 4: Pareto Fronts for each fuel (Paraffin, HDPE, HTPB, and ParaffinAl8020) and considering the 4 different LCA metrics as objectives, i.e. Global Warming Potential (a), Ozone Depletion (b), Abiotic Resource Depletion - fossil fuels (c), and Single Score (d). A logarithmic scale was used for ODEPL and Single Score for a better visualization of the results.

Finally, when considering the Single Score metric (Table 6), which accounts for all impact categories, HDPE results in a 72% lower environmental impact than HTPB, which has the highest I_{sp} among the compromise designs (only 2.3% higher than HDPE). Additionally, HDPE outperforms Paraffin in specific impulse (1.4% higher) while also achieving a 1.1% lower overall environmental impact when the performance goal is prioritized.

It should be noted that some of the results were more similar than expected, particularly with regard to Aluminized Paraffin, which was initially expected to have higher impacts and a larger difference from the other combinations. Upon further analysis, it can be observed that in all cases, the oxidizer is depleted earlier than the fuel and, for the Paraffin and ParaffinAl8020 combinations, there is a higher percentage of fuel that remains unused,

Table 3: Characteristics of the optimized hybrid rocket engines, considering the GWP100 metric.

Fuel	Compromise		Maximum I_{sp}	
	GWP100 [kg CO_2-eq]	I_{sp} [s]	GWP100 [kg CO_2-eq]	I_{sp} [s]
Paraffin	20.09	176.96	56.26	192.57
HDPE	11.43	**186.09**	**39.44**	**202.11**
HTPB	**8.95**	175.66	45.09	192.41
ParaffinAl8020	21.73	172.95	59.93	195.57

Table 4: Characteristics of the optimized hybrid rocket engines, considering the OD metric.

Fuel	Compromise		Maximum I_{sp}	
	ODEPL [kg CFC11-eq]	I_{sp} [s]	ODEPL [kg CFC11-eq]	I_{sp} [s]
Paraffin	0.0016	188.10	0.0227	192.57
HDPE	$\mathbf{9.67\times10^{-5}}$	**193.21**	0.0175	**201.00**
HTPB	3.31×10^{-4}	188.33	0.0345	195.47
ParaffinAl8020	0.0010	186.84	**0.0050**	190.87

Table 5: Characteristics of the optimized hybrid rocket engines, considering the ADEPLf metric.

Fuel	Compromise		Maximum I_{sp}	
	ADEPLf [MJ]	I_{sp} [s]	ADEPLf [MJ]	I_{sp} [s]
Paraffin	**23.72**	147.96	796.88	196.14
HDPE	144.15	**188.47**	**561.11**	**204.03**
HTPB	113.23	177.93	818.81	197.58
ParaffinAl8020	64.56	153.97	574.57	196.10

Table 6: Characteristics of the optimized hybrid rocket engines, considering the Single Score metric.

Fuel	Compromise		Maximum I_{sp}	
	Single Score [Pt]	I_{sp} [s]	Single Score [Pt]	I_{sp} [s]
Paraffin	0.0019	175.74	**0.0093**	197.03
HDPE	**0.0016**	189.02	0.0094	**199.88**
HTPB	0.0059	**193.38**	0.0330	197.06
ParaffinAl8020	0.0040	185.44	0.0194	194.04

leading to fewer combustion emissions being released. A possible explanation lies in the current configuration, which involves the usage of a small rocket engine, with constraints that probably need to be revised, particularly regarding the O/F ratio and the proportion between the tank and combustion chamber sizes.

These results can be further explored by analyzing, for each metric, the breakdown of these impacts across the three sections considered: Oxidizer production, Fuel production and combustion Emissions released.

As far as GWP is concerned, there are no significant differences in the distribution of impacts, which remains relatively similar across all compromise and maximum performance points for all the combinations. As a result, the largest fraction corresponds to Oxidizer production (around 75%), followed by Combustion Emissions released (around 20%), and lastly, with the lower impact, Fuel production (around 5%). One of the combinations that exhibits the largest variation from this distribution is ParaffinAl8020 at the point of maximum performance: it has lower values for Oxidizer production (65.8%) and higher values for combustion Emissions (22.1%) and considerably higher values for Fuel production (12%). In fact, for this impact category, the production of the Aluminized Paraffin has, inherently, substantially greater impacts than the other fuels.

Regarding the ODEPL category, its impacts arise almost entirely from combustion Emissions, with the fractions attributed to Fuel and Oxidizer production being close to zero. The usage of this impact category therefore essentially aims to minimize the emissions released into the atmosphere during combustion.

In the case of ADEPLf, since combustion Emissions are not characterized in this category, the impacts derive exclusively from Fuel and Oxidizer production. The usage of this impact category therefore essentially aims to minimize the impact of propellant production. Overall, the breakdown of these impacts is similar for almost all the points and all the combinations, with around 88% corresponding to the Oxidizer and 12% to the Fuel. However, it should be noted that the points with the greatest deviation occur at the maximum performance of HDPE, with around 95.7% attributed to the oxidizer and 4.3% to the fuel, and also at the maximum performance of Aluminized Paraffin, with 81.5% and 18.5% corresponding to the oxidizer and the fuel, respectively.

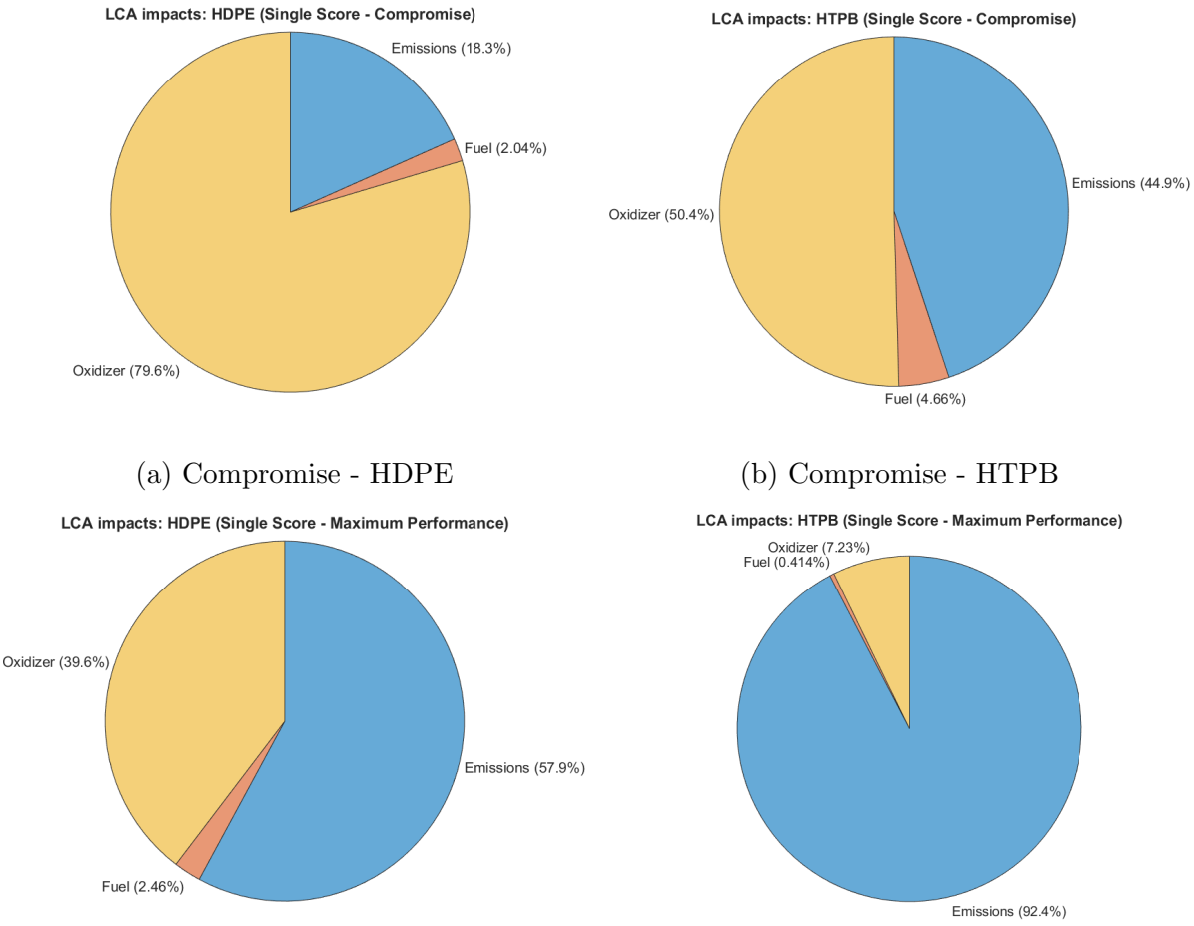
Finally, there remains the metric that encompasses several impact categories, which are subjected to a normalization and weighting: the Single Score method. This is the metric with the most noticeable differences in results, both in terms of the points analyzed and the fuel considered.

Concerning the compromise point, for all combinations, the impacts of the Oxidizer are the highest, followed by Emissions, and finally, Fuel. As for the maximum performance points, the impacts of Emissions increase significantly, exceeding those of the Oxidizer for all combinations.

Among the four combinations studied, there is one that is markedly different from the others: HTPB. In this case, Emissions increase significantly compared to the other fuels, with this increase translating into a much higher percentage attributed to the Emissions than in the other combinations. A closer scrutiny of combustion emissions revealed a high release of NOx, chemical species that have a considerable impact not only on GWP and ODEPL, but also on the other categories considered by this Single Score method.

For a clear visualization of the Single Score results, Figure 5 presents the distribution of impacts for HTPB and HDPE, the fuel that has been standing out, showing lower overall impacts and higher performance in all the metrics analyzed, compared to the other combinations tested. The distributions for Paraffin and ParaffinAl8020 are similar to those for HDPE, with the exception of the slightly higher Emissions impacts (around 71%) and the slightly lower Oxidizer impacts (around 25%) at the maximum performance

point.



(a) Compromise - HDPE

(b) Compromise - HTPB

(c) Maximum Performance - HDPE

(d) Maximum Performance - HTPB

Figure 5: Single Score impacts distribution for HDPE and HTPB and considering the two analyzed points, namely Compromise point - HDPE (a), Compromise point - HTPB (b), Maximum Performance point - HDPE (c), and Maximum Performance point - HTPB (d).

The results obtained across the metrics analyzed show that, in most cases, the largest fraction of impacts corresponds to the Oxidizer and, in some cases, to the Emissions. However, the smallest percentage always corresponds to the Fuel. Although the focus of this study was to test different fuels, and despite the production of these fuels not exhibiting significant differences in the final results, it is important to emphasize that different combinations of fuel not only affect the chemical species released but also the performance achieved, reflecting the multidisciplinary nature of the propulsion system design.

With specific regard to some of the fuels tested, a few additional observations can be drawn. Considering the overall results, it can be seen that, in fact, Aluminized Paraffin would probably have even higher results, both in terms of performance and impacts, if there were not such a significant amount of unburned fuel remaining in the combustion chamber.

Considering all the results presented under the conditions of this study, it can be concluded that the combination with the best trade-off between maximizing performance and minimizing environmental impacts is HDPE. Overall, it reveals greater specific impulses and lower impacts, even considering that it is one of the combinations in which most of the fuel is burned, and therefore leads to larger additional impacts with respect to combustion Emissions.

It should also be mentioned that this propellant offers the additional advantage of being recyclable and can also be used in its recycled form, with studies indicating that it can be recycled up to 10 times [2]. The same reference suggest that injection molding and shredding techniques do not significantly alter its properties upon reuse. It has been proposed that recycled fuel plastic be melted and purified before being cast into solid grain for use as rocket fuel, in order to meet quality standards and reduce unwanted collateral emissions. The combination of Nitrous Oxide and HDPE thus emerges as a promising option that should be considered in the future development of this framework.

4 CONCLUDING REMARKS

In this work, an LCA module was developed that quantifies the impacts of propellants in a hybrid propulsion system, and which was successfully integrated into an MOO framework alongside a Propulsion module. One of the main contributions to the state of the art is the LCA module, which presents broader system boundaries in terms of emissions, as well as an LCI based on specific datasets for space systems. This module, according to a given metric, either a Single Score or a specific Impact Category, accounts for the impacts of propellant production and includes a more extensive range of combustion emissions than what is typically considered in the literature.

Nevertheless, this study shows potential for further development of this framework. The constraints may need to be revised, modified or adjusted depending on the propellant combination under consideration. Additional sensitivity analyses are required for the MOO framework and the respective parameters used in the optimizer. Other performance metrics, beyond specific impulse, could also be explored. Furthermore, the framework should be applied to larger propulsion systems, and the inclusion of a pressurization system should be evaluated in order to test and compare results for different oxidizers. Regarding the LCA module, some improvements can be made. More particularly, at the LCI stage, a 3D printing method for manufacturing the fuel grain should be tested. More broadly, the system boundaries should be expanded to consider the materials and manufacturing processes of the entire propulsion system, including tank, combustion chamber, nozzle and other components, as well as to adopt a cradle-to-grave approach by incorporating waste and end-of-life impacts.

ACKNOWLEDGEMENTS

The authors acknowledge Fundação para a Ciência e a Tecnologia (FCT), through IDMEC, under project LAETA Base Funding (doi:10.54499/UIDB/50022/2020). Part of this work was carried out within the scope of the "New Space Portugal" project (Agenda 2022-C05i0101-02) funded by the European Union's NextGenerationEU program under the Recovery and Resilience Programme (PRR), through Agência para a Competitividade e Inovação (IAPMEI) contract N.º 11.

REFERENCES

[1] C. Christensen, R. Puleo, C. R. Herrera, and C. Mullins. Smallsats by the Numbers: Growing Smallsat Activity and Expanding Commercial Capabilities. In *AIAA AVIATION FORUM AND ASCEND 2024*, Las Vegas, Nevada, USA, July–August 2024. doi:10.2514/6.2024-4804.

[2] F. Barato. Review of Alternative Sustainable Fuels for Hybrid Rocket Propulsion. *Aerospace*, 10(7):643, 2023. doi:10.3390/aerospace10070643.

[3] C. Glaser, J. Hijlkema, and J. Anthoine. Bridging the Technology Gap: Strategies for Hybrid Rocket Engines. *Aerospace*, 10(10):901, 2023. doi:10.3390/aerospace10100901.

[4] S.-S. Wei, M.-C. Li, A. Lai, T.-H. Chou, and J.-S. Wu. A Review of Recent Developments in Hybrid Rocket Propulsion and Its Applications. *Aerospace*, 11(9):739, 2024. doi:10.3390/aerospace11090739.

[5] A. Okninski, W. Kopacz, D. Kaniewski, and K. Sobczak. Hybrid rocket propulsion technology for space transportation revisited - propellant solutions and challenges. *FirePhysChem*, 1(4):260–271, 2021. doi:10.1016/j.fpc.2021.11.015. Progress in Hybrid Rocket Propulsion.

[6] M. Xiao, H. Zhu, Z. Du, Y. Gao, H. Tian, and G. Cai. Design optimization of velocity-controlled cruise vehicle propelled by throttleable hybrid rocket motor. *Aerospace Science and Technology*, 115:106784, 2021. doi:10.1016/j.ast.2021.106784.

[7] L. Casalino, F. Masseni, and D. Pastrone. Optimal Design of Hybrid Rocket Small Satellite Launchers: Ground Versus Airborne Launch. *Journal of Spacecraft and Rockets*, 59(6):2084–2093, 2022. doi:10.2514/1.A35428.

[8] H. Zhu, M. Xiao, J. Zhang, and G. Cai. Uncertainty design and optimization of a hybrid rocket motor with mixed random-interval uncertainties. *Aerospace Science and Technology*, 128:107791, 2022. doi:10.1016/j.ast.2022.107791.

[9] P. Wang, H. Tian, H. Zhu, and G. Cai. Multi-disciplinary design optimization with fuzzy uncertainties and its application in hybrid rocket motor powered launch vehicle. *Chinese Journal of Aeronautics*, 33(5):1454–1467, 2020. doi:10.1016/j.cja.2019.11.002.

[10] A. Zavoli, P. Maria Zolla, L. Federici, M. Tindaro Migliorino, and D. Bianchi. Surrogate Neural Network for Rapid Flight Performance Evaluation of Hybrid Rocket Engines. *Journal of Spacecraft and Rockets*, 59(6):2003–2016, 2022. doi:10.2514/1.A35369.

[11] T. L. Messinger, M. S. Corbiell, and C. T. Johansen. Optimization and parametric studies of two-stage-to-orbit liquid oxygen/paraffin hybrid rocket vehicles. *Aerospace Science and Technology*, 140:108495, 2023. doi:10.1016/j.ast.2023.108495.

[12] A. Souza, P. T. Gonçalves, F. Afonso, F. Lau, N. Rocha, and A. Suleman. On the Multidisciplinary Design of a Hybrid Rocket Launcher with a Composite Overwrapped Pressure Vessel. *Journal of Composites Science*, 8(3):109, 2024. doi:10.3390/jcs8030109.

[13] L. Casalino, A. Ferrero, F. Masseni, and D. Pastrone. Emission-driven hybrid rocket engine optimization for small launchers. *Aerospace*, 9(12):807, 2022. doi:10.3390/aerospace9120807.

[14] G. Musso, I. Figueiras, H. Goubel, A. Gonçalves, A. L. Costa, B. Ferreira, L. Azeitona, S. Barata, A. Souza, F. Afonso, I. Ribeiro, and F. Lau. A Multidisciplinary Optimization Framework for Ecodesign of Reusable Microsatellite Launchers. *Aerospace*, 11(2):126, 2024. doi:10.3390/aerospace11020126.

[15] A. R. Wilson, M. Vasile, C. Maddock, and K. Baker. Implementing life cycle sustainability assessment for improved space mission design. *Integrated Environmental Assessment and Management*, 19(4):1002–1022, 12 2022. doi:10.1002/ieam.4722.

[16] E. Tormena, A. R. Silva Teixeira, B. Brevault, M. Balesdent, and A. Urbano. Multidisciplinary design analysis and optimization of launch vehicles including environmental impact. In *9th European Conference for Aeronautics and Space Sciences (EUCASS)*, Lille, France, June 2022.

[17] L. Miraux, A. R. Wilson, and G. J. Dominguez Calabuig. Environmental sustainability of future proposed space activities. *Acta Astronautica*, 200:329–346, 2022. doi:10.1016/j.actaastro.2022.07.034.

[18] T. Bellier, C. Bil, J. Morlier, A. Urbano, and A. Pudsey. Impact of Geographical and Logistics Considerations on the Life Cycle Assessment of Launch Vehicles. In S. Fu, editor, *2023 Asia-Pacific International Symposium on Aerospace Technology (APISAT 2023) Proceedings*, Singapore, 2024. Springer Nature Singapore.

[19] G. J. Dominguez Calabuig, A. Wilson, S. Bi, M. Vasile, M. Sippel, and M. Tajmar. Environmental life cycle assessment of reusable launch vehicle fleets: Large climate impact driven by rocket exhaust emissions. *Acta Astronautica*, 221:1–11, 2024. doi:10.1016/j.actaastro.2024.05.009.

[20] J. Vila, A. Ott, J. Hassin, and C. Bonnal. Maiaspace: Sustainability applied to the development of a semi-reusable solution for space mobility. *Acta Astronautica*, 226: 31–38, 2025. doi:10.1016/j.actaastro.2024.11.027.

[21] A. Mazzetti, L. Merotto, and G. Pinarello. Paraffin-based hybrid rocket engines applications: A review and a market perspective. *Acta Astronautica*, 126:286–297, 2016. doi:10.1016/j.actaastro.2016.04.036.

[22] C. Nguyen and J. C. Thomas. Performance of additively manufactured fuels for hybrid rockets. *Aerospace*, 10(6):500, 2023. doi:10.3390/aerospace10060500.

[23] S. Gordon and B. J. Mcbride. Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications. Part 1: Analysis. Technical report, NASA, 1994. NASA Reference Publication 1311, Accessed on: 10/04/2025.

[24] LCA software for informed changemakers. https://simapro.com/. Last accessed: 2025-04-11.

[25] E. W. Lemmon, I. H. Bell, M. L. Huber, and M. O. McLinden. Thermophysical Properties of Fluid Systems, 1998. URL https://doi.org/10.18434/T4D303.

[26] J. E. Zimmerman, B. S. Waxman, B. Cantwell, and G. Zilliac. Review and Evaluation of Models for Self-Pressurizing Propellant Tank Dynamics. In *49th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, San Diego, CA, USA, July 2013. doi:10.2514/6.2013-4045.

[27] A. McCormick, E. Hultgren, M. Lichtman, J. Smith, R. C. Sneed, and S. Azimi. Design, Optimization, and Launch of a 3" Diameter N_2O/Aluminized Paraffin Rocket. In *41st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Tucson, Arizona, USA, July 2005. doi:10.2514/6.2005-4095.

[28] J. C. Thomas, J. M. Stahl, A. J. Tykol, F. A. Rodriguez, and E. L. Peterse. Hybrid rocket studies using htpb/paraffin fuel blends in gaseous oxygen flow. In *7th European Conference for Aeronautics and Space Sciences (EUCASS)*, 2015. doi:10.13009/EUCASS2017-251.

[29] X. Sun, H. Tian, Y. Li, N. Yu, and G. Cai. Regression rate behaviors of HTPB-based propellant combinations for hybrid rocket motor. *Acta Astronautica*, 119:137–146, 2016. doi:10.1016/j.actaastro.2015.11.015.

[30] C. C. Jones. Performance and analysis of vortex oxidizer injection in a hybrid rocket. Technical report, U.S. Naval Academy Annapolis, 2009.

[31] E. Doran, J. Dyer, K. Lohner, Z. Dunn, B. Cantwell, and G. Zilliac. Nitrous Oxide Hybrid Rocket Motor Fuel Regression Rate Characterization. In *43rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, Cincinnati, OH, USA, July 2007. doi:10.2514/6.2007-5352.

[32] F. K. Leverone, M. J. Brooks, J.-F. P. Pitot de la Beaujardiere, and K. L. Veale. Performance sensitivity study on a blowdown nitrous oxide paraffin wax hybrid sounding rocket. *Acta Astronautica*, 160:230–239, 2019. doi:10.1016/j.actaastro.2019.04.043.

[33] D. Lee, D. Fahey, A. Skowron, M. Allen, U. Burkhardt, Q. Chen, S. Doherty, S. Freeman, P. Forster, J. Fuglestvedt, A. Gettelman, R. De León, L. Lim, M. Lund, R. Millar, B. Owen, J. Penner, G. Pitari, M. Prather, R. Sausen, and L. Wilcox. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244, 2021. doi:10.1016/j.atmosenv.2020.117834.

[34] G. J. Dominguez Calabuig, L. Miraux, A. Wilson, A. Pasini, and A. Sarritzu. Ecodesign of future reusable launchers: insight into their life cycle and atmospheric impact. In *9th European Conference for Aeronautics and Space Sciences (EUCASS)*, June–July 2022. doi:10.13009/EUCASS2022-7353.

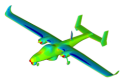
[35] G. J. Dominguez Calabuig, A. Wilson, S. Bi, M. Vasile, M. Sippel, and M. Tajmar. Environmental life cycle assessment of reusable launch vehicle fleets: Large

climate impact driven by rocket exhaust emissions. *Acta Astronautica*, 221, 2024. doi:10.1016/j.actaastro.2024.05.009.

[36] A. Lambe and J. Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Structural and Multidisciplinary Optimization*, 42:273–284, 2012. doi:10.1007/s00158-012-0763-y.

[37] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6 (2), 2002. doi:10.1109/4235.996017.

[38] J. Persson. Constrained NSGA-II with seeding enabled. `https://www.mathworks.com/matlabcentral/fileexchange/75338-constrained-nsga-ii-with-seeding-enabled`, 2021. Last accessed: 2025-04-14.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

MULTI-OBJECTIVE SURROGATE-BASED OPTIMISATION OF ROCKET FIN DESIGN

José Silva[1*], Frederico Afonso[2] and Fernando Lau[2]

1: Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
jose.goulao@tecnico.ulisboa.pt, https://tecnico.ulisboa.pt/

2: IDMEC, Instituto Superior Técnico, Universidade de Lisboa, Lisboa, Portugal
{frederico.afonso,lau}@tecnico.ulisboa.pt, https://www.idmec.tecnico.ulisboa.pt/

Abstract. *This work presents a multi-objective optimisation approach for rocket fin design, focusing on drag coefficient reduction and lift coefficient increase, while ensuring rocket stability and avoiding fin flutter. To achieve these goals, a multi-fidelity surrogate model trained with high and low fidelity three-dimensional computational fluid dynamics simulations is used. The surrogate model, specifically an artificial neural network, is capable of predicting both the lift and drag coefficients, enabling a broader study of fin configurations that would otherwise be computational demanding. Genetic algorithms are then applied to identify the optimal fin shapes and dimensions. Geometric parameters, flutter velocity, and static margin constraints are imposed beforehand to ensure that the optimised fin designs meet the respective requirements. The optimisation was performed across diverse flow conditions, at multiple Mach numbers and corresponding altitudes according to a preliminary flight profile, and at* 0 *and* 5 *angles of attack. The model yielded significant improvements in both aerodynamic coefficients, indirectly contributing to a more sustainable rocket design.*

Keywords: Fin multi-objective optimisation, Multi-fidelity artificial neural network, Genetic algorithm, Multi-disciplinary constraints

1 INTRODUCTION

The space industry has long faced environmental critiques due to its energy and resource consumption, reliance on fossil fuels, and the consequent emission pollution. However, the New Space era emergence, led by the private sector, forced rapid innovation, costs reduction, and the implementation of more sustainable approaches [1]. This work will focus on enhancing sustainability of a microsatellite launcher, capable of deploying microsatellites into low Earth orbit, by indirectly reducing fuel consumption through optimisation of the fins design.

In recent years, remarkable progress has been made in the optimisation of wing and fin designs with surrogate models. Du et al. [2] developed a framework that allows rapid airfoil design optimisation for subsonic and low transonic conditions, utilising an artificial neural network (ANN) capable of predicting aerodynamic coefficients. Liu et al. [3] employed a Kriging surrogate model to optimised aircraft wings' airfoils, resulting in a 4.47% improvement in the lift-to-drag ratio in hypersonic regimes, and 3.19% in transonic regimes. Han et al. [4] also applied surrogate-based optimisation to aircraft wing airfoils in transonic flow, decreasing drag. In another study, Zhang et al. [5] developed a multi-fidelity deep ANN for airfoil optimisation in transonic flow, yielding a 114.07% increase in the lift-to-drag ratio, a superior outcome compared to results obtained with a single-fidelity ANN. Priyadarshi et al. [6] performed a multi-disciplinary multi-objective optimisation of sounding rocket fins, incorporating stability, structural and rocket trajectory concerns, by successfully employing a genetic algorithm in their approach. These studies demonstrate the effectiveness and reliability of surrogate-based optimisation techniques. However, it is noticeable that most of the optimisations focused on the airfoil section of wings, and not on other geometric parameters. As a consequence, factors such as flutter avoidance and stability issues are often neglected in these models. Even if stability was considered, the studies typically target aircraft wings rather than launchers fins. In this sense, a gap in the literature exists, specifically regarding the development of comprehensive multi-objective optimisation of rocket fin designs across multiple flow conditions, with the use of a multi-fidelity ANN, and with multi-disciplinary constraints.

This work aims to develop a multi-objective optimisation framework for rocket fin design, focusing on the minimisation of the drag coefficient and maximisation of the lift coefficient. For this purpose, genetic algorithms are employed with multi-disciplinary constraints to ensure rocket stability and flutter avoidance. This optimisation tool is used based on the aerodynamic coefficients predictions of an ANN trained with multi-fidelity three-dimensional computational fluid dynamics simulations (CFD). Additionally, this work seeks to demonstrate the viability of this approach in the design process of rocket fins, which indirectly enhances sustainability while contributing to mission safety and success.

2 LAUNCHER DESCRIPTION

2.1 Rocket Specifications

As of the development of this work, limited information regarding the New Space Portugal rocket's specifications is available. For instance, its dimensions are still under analysis. Therefore, to employ the estimations for the centre of pressure (CP) location required for the static margin (SM), an existing rocket's dimensions need to adopted as a reference.

The Space-borne Payload Assist Rocket- Kauai (SPARK), by Sandia National Laboratories, Aerojet, and the University of Hawaii, was selected as the reference rocket. SPARK is actually a three-stage small-satellite launcher with a maximum payload capacity of 250 kg to a 400 km sun-synchronous orbit [7, 8]. It was also chosen due to the scarcity of available data on other launcher's dimensions with similar missions to that of New Space Portugal.

2.2 Mission Profile

Table 1 presents the expected Mach number of a rocket at five distinct altitudes. For the purpose of this work, which aims to minimise drag to reduce fuel consumption, the most critical flow condition to analyse would be the phase during which most fuel is used. Nevertheless, at the current stage of this research, no information about the rocket's engine nor fuel specifications are available, but it is clear from Table 1 that the rocket will cover more distance during its ascent in supersonic and hypersonic velocities.

Table 1: New Space Portugal expected profile.

Altitude [km]	Mach number
0.5	0.5
2	1
5	3
35	6

Therefore, to conduct a more complete optimisation work, including the majority of the rocket's flight, and not just the predicted Mach numbers and altitudes, and due to the limited nature of the available data, an extrapolation is required to extend the analysis to other Mach numbers. To obtain the variations of air temperature, pressure, and density, as well as the speed of sound, with altitude, the International Standard Atmosphere model, was employed.

2.3 Initial Fin Considerations

Various types of fins exist, each offering distinct advantages depending on the flow conditions, rocket design, and on the specific parameter to be optimised. For this reason, to ensure the model presented in this work is as comprehensive as possible, it should include all various different shapes. The clipped delta fin is commonly used in microsatellite launchers, as a consequence this will be considered the baseline shape in this study. Moreover, most rockets use four fins, a configuration also adopted. Regarding the fins' cross section, a smooth symmetrical airfoil, similar to those from the NACA 00XX series, is considered.

3 METHODOLOGY

The proposed method follows the schematic illustrated in Figure 1. It initiates with fin sampling, together with the imposition of restrictions on fin dimensions, and the exclusion of fins whose flutter speeds are deemed too low or do not ensure rocket stability, which are covered in Section 4. Thereafter, all the preparations for training the surrogate model are conducted as described in Section 5. These include domain meshing and CFD simulations of a number of the sampled fins - comprising the training dataset -, as well as the actual

training of the model, so it is capable of predicting the aerodynamic coefficients of a given fin design, in a given condition. Subsequently, in Section 6, the trained surrogate model's predictions are used by a genetic algorithm to identify optimal fin configurations.

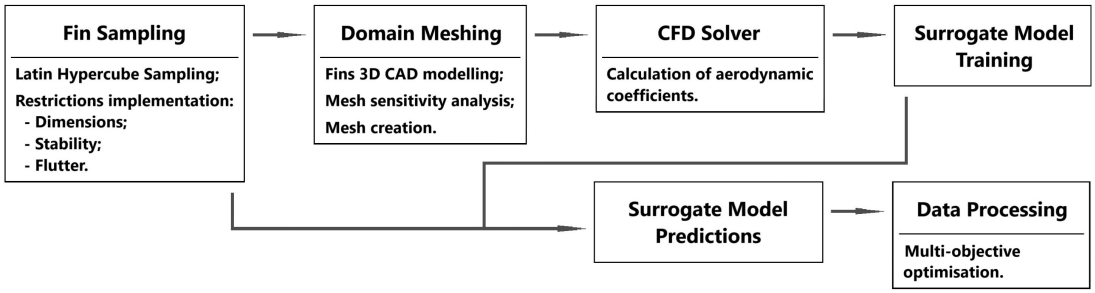


Figure 1: Schematic of the methodology.

4 FIN SAMPLING

4.1 Latin Hypercube Sampling

Sampling is a fundamental step in the development of surrogate models, as it involves generating data points that can constitute both training inputs and initial input data for the optimisation tool. In this context, a common technique is employed, the Latin hypercube sampling (LHS) method. With this strategy the design space is divided into equally sized squares, within which sampled points are generated [9]. This ensures that each dimension (or design variable) is sampled evenly, ensuring that all possible shapes are explored. For this purpose, the MATLAB® function *lhsdesign(n,p)* was employed, where n is the number of desired samples, and p represents the number of variables [10]. In this case: semi-span (S); root chord (R_{chord}); tip chord (T_{chord}); sweep angle (Λ); and thickness-to-chord ratio (t/c). However, any sampled fins that do not meet size nor flutter speed nor the SM range limitations are discarded, as explained subsequently.

4.2 Dimensional Restrictions

The first imposed constraints are dimensional restrictions, so that abnormally small or large fins, and those with odd shapes, are immediately discarded. These thresholds should not be overly strict, so as to still allow for a wide variety of designs to be analysed. Table 2 outlines these values, where ⌀ represents the rocket's outer diameter.

Table 2: Bounds of the fins' design variables.

Variable	Lower bound	Upper bound
S	0.01 m	3⌀
R_{chord}	0.01 m	3⌀
T_{chord}	0 m	R_{chord}
Λ	0°	70°
t/c	0.001	0.25

The semi-span, root chord and thickness-to-chord ratio can never equal 0, as it would result in a non-existent fin. Therefore, a small conservative lower bound of 0.01 m is applied

to S and R_{chord}, and 0.001 for the ratio's bound. This approach allows for the generation of symmetrical or asymmetrical trapezoidal fins, as well as clipped delta, tapered swept, delta (when $T_{chord} = 0$ m), and rectangular designs (when $T_{chord} = R_{chord}$ and $\Lambda = 0°$).

With respect to the upper bounds, the semi-span and root chord values are both set as $3\varnothing$, a conservative limit based on the bound of $1.35\varnothing$ in [6], and $2\varnothing$ in [11]. In addition, fins with a larger tip chord than its root chord are also rejected, which is equivalent to the taper ratio not being larger than 1, as considered in [6, 12]. As for the sweep angle's upper bound, it is 70°, which is based on the maximum of 60° considered in [12], 65° in [6] and 75° in [13]. Lastly, the upper bound for t/c, 0.25, is a conservative value grounded on the maximum of 0.2 in [6].

4.3 Flutter Restrictions

Subsequently, the fins that are not discarded due to dimensional restrictions, must now undergo flutter velocity analysis. The calculation of flutter velocity is complex thanks to its interdisciplinary nature, so Dennis Martin [14] further developed an empirical method for its use in missiles' fins, which was employed in this work. If any fin has a flutter speed below a certain limit, it ought to be eliminated. This boundary is determined by applying a 15% safety factor, as recommended by the National Aeronautics and Space Administration (NASA) in [15], above the rocket's predicted velocity along its entire flight profile. Given that the maximum planned Mach number is 6, as presented in Table 1, the minimum allowable Mach flutter velocity for the fins should be $6 \times 1.15 = 6.9$. However, at an altitude of 35 km, the flutter speed increases substantially, making fins with a flutter speed below 6.9 very rare. In fact, the limit found to be most restrictive was at 5 km, where the Mach number threshold is $3 \times 1.15 = 3.45$, thus being the restriction imposed on all fins.

Furthermore, the flutter speed is influenced by the shear modulus (GE) of the fins' material. Since the rocket's materials, including the fins, have not yet been determined, an assumption must be made. To adopt a conservative approach and avoid selecting a material with a high GE (such as a titanium alloy), which would effectively reduce the number of fins that are not rejected at this stage, Al 6061-T6 was chosen. This is a tempered hardened aluminium alloy commonly used in the aerospace industry, with a GE of 26 GPa [16].

4.4 Stability Restrictions

The stability restrictions are imposed at the final stage. Primarily, the acceptable SM range must be established. The interval selected for this work is between 1.5 and 3.5 for both the subsonic and supersonic regimes.

James S. Barrowman developed a set of equations, known as the Barrowman equations, to predict a rocket's CP location [17]. This model is commonly used in rocketry for its low computational demands and acceptable accuracy, as discussed by Schoch [18]. Galejs [19] proposed extensions to the Barrowman approach, referred to as the body lift correction, in which the aerodynamic forces produced by the nose cone and the fuselage are added, in turn also conferring a dependency in the rocket's CP with the angle of attack (AoA). Furthermore, Pitts et al. [20] presented a procedure for calculating the lift and the CP of wing-body-tail configurations at supersonic speeds. In this model, two vital parameters can be calculated by another algorithm from Barrowman, entitled the FIN code [21].

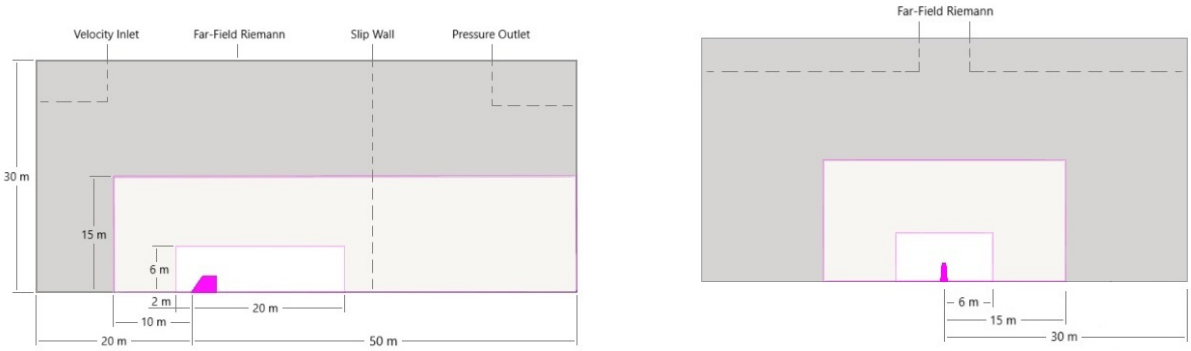
For the subsonic case, the previously mentioned methods are applied to all fins that satisfy the previous constraints. The body lift correction introduces a linear variation with the rocket's angle of attack. Therefore, each fin is evaluated for an AoA of 0° and 10° (the angle up to which this correction model remains satisfactory [19]), and if the SM for either angle falls outside the acceptable range, the fin is discarded.

Subsequently, for the supersonic regime, the algorithm Barrowman developed in the FIN code [21] is applied for each fin across all velocities up to Mach 6, at every AoA until 10°. Thereafter, the SM and all other values needed for its computation are determined utilising Pitts' method [20]. Given that the relationship between SM and Mach number is non-linear, only the minimum and maximum SM values attained across all Mach numbers are verified against the acceptable range. If at least one of these values does not comply, that fin is eliminated.

5 NEURAL NETWORK CREATION AND TRAINING

5.1 Setup for Training

The initial step in the training process involves the development of 3D models, which are integrated in the meshes, in turn required for the CFD simulations that train the ANN. For this purpose, the Open Vehicle Sketch Pad (OpenVSP) software was utilised. The mesh creation process was then conducted on BaramMesh, a BARAM application for mesh generation [22]. Mesh sizes were based on a study by Hosseini [23], and adjusted to suit the simulations in this work. Moreover, near-fields were added to allow for a more accurate analysis on the regions adjacent to the fins, and along its wake, without requiring the same refinement on outer regions. To better visualise the domains' dimensions and the surfaces' boundary conditions, Figure 2 is presented.



(a) Side view.

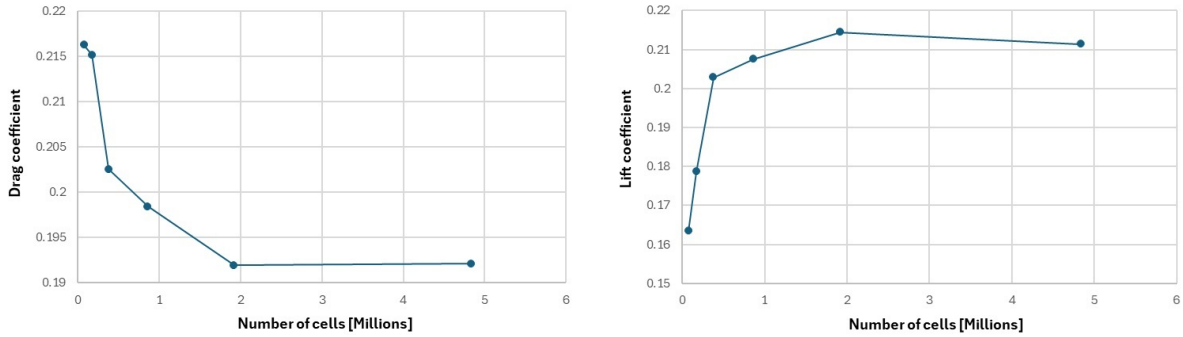
(b) Front view.

Figure 2: Domain dimensions and surfaces' boundary conditions.

Prior to actually utilising the meshes in the CFD simulations, it was necessary to conduct a mesh sensitivity analysis to ensure the accuracy of the simulations[1]. In this sense, to evaluate the meshes in the most extreme cases, six meshes, each more refined than the previous, were created for a fin with near maximum semi-span and thickness-to-chord ratio. These meshes were compared based on drag and lift coefficients with a 5° AoA, at Mach 1 and 7 at the respective altitudes.

[1]All simulations were performed by a computer with an Intel(R) Core(TM) i7-8565U CPU @ 1.80GHz processor, and 12.0 GB of RAM.

A consistent trend was consistently detected: as the number of cells increases, the results appear to converge asymptotically, with minor oscillations in some cases, as seen in Figure 3 depicting the evolution of the lift coefficient at Mach 1.



(a) Evolution of C_D.

(b) Evolution of C_L.

Figure 3: Impact of the meshes' number of cells on the aerodynamic coefficients, at Mach 1.

Another key factor impacting the selection of mesh is the velocity of the corresponding simulations. In this context, Table 3 summarises the time, measured in seconds, for a single CFD simulation iteration, for both flow conditions. Noting that the cell count ratios between the meshes are no less than 2.197, as recommended by Celik et al. [24]. As anticipated, higher cell counts lead to longer simulation times, and the iteration times seen at Mach 7 are consistently slightly greater compared to those at Mach 1.

Table 3: Simulation iteration time for each mesh in both flow conditions.

	Iteration time [s]	
Number of cells	Mach 1	Mach 7
4 835 049	32.088	32.174
1 917 998	11.050	11.279
861 869	4.463	4.468
378 124	2.062	2.228
171 866	0.848	0.857
76 650	0.374	0.380

This analysis holds particular importance for this work, as the distinction between low and high fidelity is determined by mesh refinement. Accordingly, the high-fidelity dataset will consist of simulations whose meshes contain approximately 2 million cells, since its results fall within the converging asymptotic range. Each iteration requires approximately 2.9 times less than the most refined mesh, and has a relative error compared to that mesh of around 1% for both Mach numbers 1 and 7 (extreme cases) and for both the drag and lift coefficients. The low fidelity simulations will employ the next coarser mesh, containing just under 1 million cells. The relative difference between this and the higher fidelity meshes averages around 3%, while the iteration time is less than half (approximately 0.4 times), making this selection reasonable for allowing considerably faster simulations, while

maintaining acceptable accuracy.

The CFD solver used was BARAM, in which the application BaramFlow employs OpenFOAM® solvers modified by NEXTFOAM [22]. The Menter k-ω shear stress transport model (SST), turbulence model was chosen, as it leverages the advantages of the k-ω model for the inner parts of the boundary layer, with the SST formulation benefiting from the robustness of the k-ε model in free-stream conditions, making it well-suited for these simulations [25]. For the numerical setting on turbulence, a second order upwind scheme was utilised, as it is generally recommended for accuracy, even in structured grids aligned with the flow [26]. Concerning the air viscosity, rather than assuming a constant value, Sutherland's law was enforced, as it offers a more accurate relationship between the dynamic viscosity and absolute temperature [27].

Each fin was analysed at two AoA: 0° and 5°. Understanding drag production at 0° AoA is crucial, as this is the expected rocket's orientation during most of the flight. Additionally, a small non-zero angle is also vital to analyse the evolution of drag, but also lift generation, as lift is only produced with a non null angle due to the symmetrical airfoil design. Moreover, each fin was examined at two velocities for each AoA, at the expected altitude. More specifically, two arbitrary Mach numbers between 1 and 7 were generated via LHS, with the respective altitude, as well as air temperature, pressure and density determined via extrapolation of the mission profile, as previously discussed. This approach allows the model to be trained in the high transonic, supersonic and low hypersonic regimes. The exclusion of subsonic flow represents a limitation, preventing fin optimisation under these condition. However, priority was given to higher Mach numbers, since these are the regimes in which the rocket covers the most distance during its predicted flight profile.

In addition, it is important to recall that the simulations were performed on a singular isolated fin. Consequently, interference drag with the rocket's fuselage and with the other fins is not considered, nor is the impact of the wake generated by the nose cone and other upstream structures. Therefore, this simplification presents clear limitations in the model, but were necessary due to the lack of detailed information on the rocket configuration to be used. Furthermore, performing entire rocket simulations would have significantly increased computational costs, reducing the number of training points available to train the ANN.

Therefore, with the approach taken, the network will specialise in high transonic, supersonic and low hypersonic regimes for isolated fins, so they all are compared under these same constraints. In total, 144 high fidelity simulations, corresponding to 36 distinct fins, and 460 low fidelity simulations of 115 individual fins were performed. The ratio of high to low fidelity data, of 31.3%, was determined based on the approximate 30% ratios used by Zhang et al. [24], and adjusted for the specific case of this work's ANN.

5.2 Training the ANN

The function *feedforwardnet(hiddenSizes, trainFcn)*, from MATLAB®, was employed to generate a multilayer perceptron (MLP), which is a feedforward ANN with hidden layers, and, in this case, with a sigmoid activation function. Furthermore, the network was trained with the function *train(net,X,T)*.

With respect to the architecture of the ANN, its characteristics were determined through extensive experimental training, where individual parameters were varied while the others were maintained constant. As a result, the final ANN, depicted in Figure 4, is fully

connected and made up of four layers: the input layer, with seven neurons corresponding to the five dimensions describing a unique fin, as well as the angle of attack and the Mach number; two hidden layers with the first having six neurons and the second three; and the output layer, with two neurons corresponding to the drag and lift coefficients.

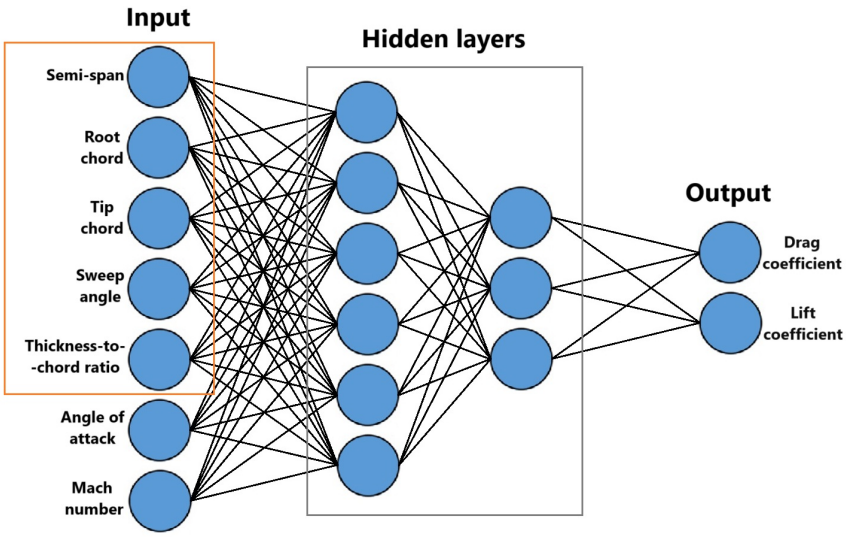


Figure 4: Schematic of the ANN.

For the first training, a Bayesian regularisation backpropagation function was found to best train the MLP. It automatically performs regularisation helping to avoid over-fitting, particularly relevant for this work where the size of the training dataset is considered small. In this sense, 90% of the data was allocated for training, and 10% for testing, with no data being used for validation. Moreover, tuning revealed that the best results were attained with a maximum number of epochs of 5000, and with a minimum gradient of 10^{-8}.

Regarding the second training, since now the goal is to fine-tune, different training conditions are needed. Indeed, the scaled conjugate gradient backpropagation function was selected, with 80 maximum epochs, a minimum gradient of 10^{-6}, and 10 validation sets, where 70% of the data is allocated to training, 15% for validation, and 15% for testing, with the remaining parameters kept at default values. With this second training the fear is that the model becomes over-fitted to the high fidelity data, thus losing the generalisation capabilities it gained in the first training with the low fidelity data.

The training resulted in a validation mean squared error of 1.441×10^{-5}. To further assess the accuracy of the ANN, a mean relative error was calculated, specifically the averaged relative error across all predictions made by the trained ANN when inputting the same fins under identical conditions as used in training. For the drag coefficient in both 0° and 5° AoA, the overall error was 1.79%, with specific errors of 2.51% and 1.55% for the high and low fidelity datasets, respectively. As for the lift coefficient at 5° AoA, the error was 3.10% for the combined fidelities, with individual errors of 5.47% and 2.30% for the high and low fidelity data, respectively. In the particular case of the lift coefficient at 0° AoA, the predicted values averaged 2.448×10^{-6}, 6.715×10^{-5} and -1.950×10^{-5} for the combined, high, and low fidelity datasets, respectively. These errors indicate that the second training phase did not lead to over-fit on the high fidelity data, and that it preserved generalisation.

6 RESULTS AND DISCUSSION

6.1 Baseline Solution

A clipped delta shape was chosen for the baseline solution. In this sense, through the sampling method described in Section 4, a clipped delta shaped fin was attained. Its dimensions are as follows: $S = 2.207$ m; $R_{chord} = 3.315$ m; $T_{chord} = 1.777$ m; $\Lambda = 34.87°$; $t/c = 0.2328$. Table 4 summarises the fin's aerodynamics coefficients - the drag coefficient, C_D, and the lift coefficient, C_L, - both with 0° and 5° AoA, for every Mach number from 1 to 7, at the respective altitudes. Noting that these results were obtained via CFD simulations, with an identical setup as the high fidelity simulations.

Table 4: Baseline aerodynamic coefficients.

	0° AoA		5° AoA	
Mach number	C_D	C_L	C_D	C_L
1	0.1863	0	0.1932	0.1869
2	0.1856	0	0.1914	0.1121
3	0.1725	0	0.1789	0.0740
4	0.1638	0	0.1691	0.0561
5	0.1599	0	0.1647	0.0461
6	0.1586	0	0.1632	0.0405
7	0.1774	0	0.1820	0.0394

6.2 Optimisation

To identify the optimal fin, genetic algorithms had to be employed. Initially, the NSGA-II algorithm [28] was used across the entire design space, but it required considerable computational resources to converge to a final result for all flow conditions. Therefore, by focusing the search to a smaller region computational demands are reduced, enabling a more thorough analysis of this area. To identify this region, 25517 unique fins, created through the sampling method, were analysed in both 0° and 5° AoA, under every Mach number from 1 to 7 at the extrapolated altitude from the predicted flight profile.

With a deeper understanding on the evolution of the aerodynamic coefficients with fin designs, and the most favourable regions for this optimisation, a genetic algorithm (GA) can now be used. The GA had multi-disciplinary constraints identical to the ones present in the sampling method (dimensional, flutter and stability restrictions). For the case of 0° AoA, the initial population consisted of the fins with the lowest t/c out of the previous 25517 fins. The parameters defining the algorithm, such as the crossover rate and mutation rate, were selected after a sensitivity analysis, based on a balance between exploration and exploitation. Employing the GA across all Mach numbers led to the identification of the optimal fin design, which was the same across all conditions, and can be visualised in Figure 5.

The resulting optimal fin design has a low t/c, a semi-span length corresponding to the upper bound outlined in Table 2, and a sweep angle reaching the lower bound. The minimisation of drag drives the decrease of t/c, but simply minimising this ratio without adjusting other parameters would result in fins not capable of stabilisation. Therefore, the

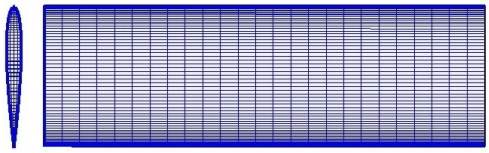


Figure 5: Model of the identified optimal fin at a 0° AoA.

increase in semi-span and decrease in sweep angle force the CP forward, ensuring that the SM remains within the acceptable range. Moreover a unitary taper ratio is also preferred to minimise C_D, thus resulting in a quasi-rectangular design, which was an unexpected result. This optimised design resulted in a substantial decrease in C_D, as summarised in Table 5, with an average reduction of 62.4% across all conditions, when compared to the baseline fin.

Table 5: Aerodynamic coefficients and corresponding improvements over the baseline fin for the identified optimal fin, across all Mach numbers at 0° AoA.

Mach number	C_D	Decrease [%]
1	0.0773	58.5
2	0.0704	62.1
3	0.0643	62.7
4	0.0593	63.8
5	0.0574	64.1
6	0.0596	62.4
7	0.0655	63.1

The optimisation of the fins at a 5° AoA was performed with the NSGA-II algorithm with the same parameters used in the GA for the 0° AoA case. The initial population consisted of all the fins present in the Pareto fronts determined for the 25517 fins, as well as fins with similar dimensions to the ones in the respective fronts. The resulting Pareto fronts can be visualised in Figure 6, in which it is interesting to examine both the C_D minimising and the C_L maximising fins from each front.

With respect to the fins that minimise the drag coefficient within each Pareto front, the same fin design consistently arises. This design is also identical to the one identified as optimal at a 0° AoA. With this design, at a 5° AoA, the drag coefficient is decreased by 57.3%, and the lift coefficient enhanced by 35.8%, averaged across the seven regimes considered, when compared to the baseline fin (Table 6).

As for fin that maximises the lift coefficient within each Pareto front, the same configuration is identified across all conditions, and is depicted in Figure 7. When compared to the fin in Figure 5, two parameters are essentially unaltered: the semi-span and the sweep angle, remaining at the respective upper and lower bound. The most substantial changes occur in the thickness of the fin, as well as in the root and tip chord, both increasing in similar proportions, so the fin retain a quasi-rectangular design. This design actually increases the drag coefficient by 49.2% on average, when compared to the baseline fin, and it increases the lift coefficient by 113.8% (Table 7).

The observation that the optimal fins derived with this approach exhibit a rectangular configuration indicates that a prior assumption in the framework may have been too strict.

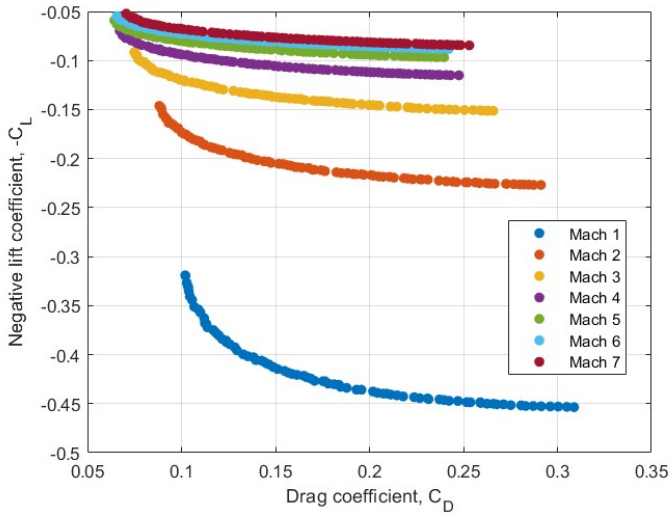


Figure 6: Pareto fronts for all considered Mach numbers, at a 5° AoA.

Table 6: Aerodynamic coefficients and corresponding improvements over the baseline fin for the C_D-minimising fin, across all Mach numbers at 5° AoA.

Mach number	C_D	Decrease [%]	C_L	Increase [%]
1	0.1021	47.2	0.3221	72.3
2	0.0883	53.8	0.1478	31.9
3	0.0750	58.1	0.0925	24.9
4	0.0671	60.3	0.0697	24.2
5	0.0643	61.0	0.0592	28.6
6	0.0656	59.8	0.0547	34.8
7	0.0706	61.2	0.0527	33.6

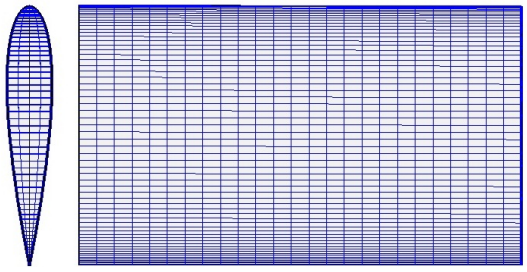


Figure 7: Model of the identified C_L-maximising fin across all Mach numbers at 5° AoA.

To evaluate this hypothesis, the SM range was adjusted to 1 up to 5, followed by a shorter optimisation process using the genetic algorithm, conducted solely at an AoA of 0° and at Mach 3. The resulting fin design is depicted in Figure 8. When compared to the previous optimal fin under identical flow conditions, the root and tip chords are similar, as well as the thickness. However, notable differences include a reduced semi-span, which decreases the fin's area and aspect ratio. Crucially, the sweep angle is increased in this configuration. This design not only aligns more closely with those commonly found in existing microsatellite launchers (as the previously mentioned SPARK), it also exhibits a

Table 7: Aerodynamic coefficients and corresponding improvements over the baseline fin for the C_L-maximising fin, across all Mach numbers at 5° AoA.

Mach number	C_D	Decrease [%]	C_L	Increase [%]
1	0.3178	−64.4	0.4534	142.5
2	0.2912	−52.2	0.2269	102.4
3	0.2659	−48.6	0.1513	104.4
4	0.2476	−46.4	0.1151	105.2
5	0.2397	−45.5	0.0967	109.8
6	0.2421	−48.4	0.0882	117.5
7	0.2531	−39.1	0.0846	114.5

lower drag coefficient, of 0.0553, when compared to the prior optimal configuration, with a C_D of 0.0668. Therefore, these results support the hypothesis that the initial static margin range was overly restrictive.

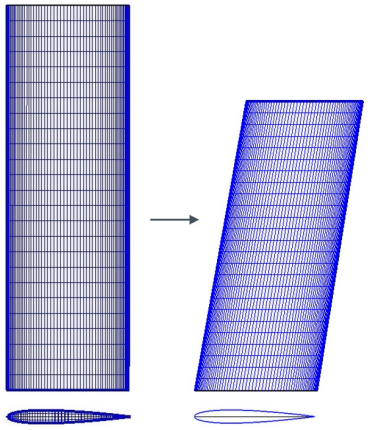


Figure 8: Model of optimal fin at a 0° AoA and Mach 3 with a SM range of 1.5 to 3.5 (left) compared to the fin with a range of 1 to 5 (right).

7 CONCLUSIONS

This work implemented a methodology for optimising the fin design for the microsatellite launcher under the New Space Portugal initiative. Using a multi-fidelity surrogate model integrated with a genetic algorithm with multi-disciplinary constraints, significant aerodynamic improvements were achieved, when compared to the baseline fin.

An additional achievement was the development of a multi-fidelity ANN capable of prediction drag and lift coefficients across a broad range of fin designs and flow conditions, thus reducing the need for extensive CFD simulations. In this sense, the integration of machine learning with aerodynamic shape optimisation was favourably performed, demonstrating the viability of this method in rocket fin design processes. The methodology's strength lies in its versatility, adapting to different launchers, mission profiles, and changing constraints, such as dimensional limits, flutter restrictions, or, as demonstrated, the static margin.

Despite the progress achieved in this work, several topics would benefit from further studies. Airfoil shape optimisation would enhance the current model, which relies on a

symmetric airfoil with a constant thickness-to-chord ratio. Additionally, expanding the analysis to include subsonic conditions would allow for a more complete optimisation across the entire rocket's flight. Moreover, if the fuel consumption rate is known, a weighted average could be performed between the optimal fin's dimensions to achieve a more balanced final design. Further improvements could involve introducing the interaction between fins and other components, for more realistic CFD simulations, as well as incorporating structural and economical factors, for a more balanced and sustainable fin design.

ACKNOWLEDGMENTS

The authors acknowledge Fundação para a Ciência e a Tecnologia (FCT) for its financial support via the project LAETA Base Funding (doi:10.54499/UIDB/50022/2020).

REFERENCES

[1] F. Del Canto Viterale. Transitioning to a New Space Age in the 21st Century: A Systemic-Level Approach. *Systems*, 11(5):232, 2023. doi:10.3390/systems11050232.

[2] X. Du, P. He, and J. R. Martins. Rapid airfoil design optimization via neural networks-based parameterization and surrogate modeling. *Aerospace Science and Technology*, 113:106701, 2021. doi:10.1016/j.ast.2021.106701.

[3] F. Liu, Z.-H. Han, Y. Zhang, K. Song, W.-P. Song, F. Gui, and J.-B. Tang. Surrogate-based aerodynamic shape optimization of hypersonic flows considering transonic performance. *Aerospace Science and Technology*, 93:105345, 2019. doi:10.1016/j.ast.2019.105345.

[4] Z.-H. Han, J. Chen, K.-S. Zhang, Z.-M. Xu, Z. Zhu, and W.-P. Song. Aerodynamic Shape Optimization of Natural-Laminar-Flow Wing Using Surrogate-Based Approach. *AIAA Journal*, 56(7):2579–2593, 2018. doi:10.2514/1.J056661.

[5] X. Zhang, F. Xie, T. Ji, Z. Zhu, and Y. Zheng. Multi-fidelity deep neural network surrogate model for aerodynamic shape optimization. *Computer Methods in Applied Mechanics and Engineering*, 373:113485, 2021. doi:10.1016/j.cma.2020.113485.

[6] P. Priyadarshi, M. Alam, and K. Saroha. Multi-disciplinary multi-objective design optimization of sounding rocket fins. *International Journal of Advances in Engineering Sciences and Applied Mathematics*, 6(3-4):166–182, January 2015. doi:10.1007/s12572-015-0121-6.

[7] N. Brügge. Super Strypi - SPARK. Available online at: `https://b14643.eu/Spacerockets_3/Sandia-Lab/Description/Frame.htm`. Accessed on: 06/09/2024.

[8] G. D. Krebs. Super Strypi (SPARK) / Bolt. Available online at: `https://space.skyrocket.de/doc_lau/super-strypi.htm`. *Gunter's Space Page*. Accessed on: 06/09/2024.

[9] R. L. Iman. Latin Hypercube Sampling. In *Encyclopedia of Quantitative Risk Analysis and Assessment*. John Wiley & Sons, Ltd, 2008. ISBN 9780470061596. doi:10.1002/9780470061596.risk0299.

[10] The MathWorks Inc. lhsdesign: Latin hypercube sample. The MathWorks Inc., 2024. URL https://www.mathworks.com/help/stats/lhsdesign.html. Accessed on: 08/09/2024. MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See mathworks.com/trademarks for a list of additional trademarks.

[11] STAR Public. Sizing fins. Available online at: https://rocketry.gitbook.io/public/tutorials/airframe/sizing-fins. Aaccessed on: 08/09/2024.

[12] P. Żurawka, N. Sahbon, D. Pytlak, M. Sochacki, A. Puchalski, and S. Murpani. Multi-objective optimization of a fin shape for a passive supersonic rocket stage. In *2023 IEEE Aerospace Conference*, Big Sky, MT, USA, March 2023. doi:10.1109/AERO55745.2023.10115859.

[13] A. Pektas, U. Haciabdullahoglu, N. Ejder, Z. Demircan, and C. Tola. Effects of different fin shapes on apogee and stability of model rockets. In *2019 9th International Conference on Recent Advances in Space Technologies (RAST)*, pages 193–199, June 2019. doi:10.1109/RAST.2019.8767439.

[14] D. J. Martin. Sumamary of flutter experiences as a guide to the preliminary design of lifting surfaces on missiles. Technical Note 4197, NACA, Langley Aeronautical Laboratory, Langley Field, Washington, Va., February 1958.

[15] NASA. Flutter, buzz, and divergence. SP 8003, Space Vehicle, Design Criteria, Scientific and Technical Information Facility, Bethesda, Maryland, July 1964.

[16] ASM Material Data Sheet. Aluminum 6061-T6; 6061-T651. ASM Aerospace Specification Metals Inc., March 2013.

[17] J. Barrowman. Calculating the center of pressure of a model rocket. Technical Information Report 33, Centuri Engineering Company, Inc., Phoenix, Arizona, 1970.

[18] A. Schoch. Rocket stability assessment. Bachelor thesis, ZHAW School of Engineering, June 2023.

[19] R. Galejs. Wind instability, what barrowman left out. URL https://argoshpr.ch/joomla1/articles/pdf/sentinel39-galejs.pdf.

[20] W. C. Pitts, J. N. Nielsen, and G. E. Kaattari. Lift and center of pressure of wing-body-tail combinations at subsonic, transonic, and supersonic speeds. Report 1307, NACA, Ames Aeronautical Laboratory, Moffett Field, California, February 1959.

[21] J. S. Barrowman. FIN - a computer program for calculating the aerodynamic characteristics of fins at supersonic speeds. N66 30346, NASA, Goddard Space Flight Center, Greenbelt, Maryland, April 1966.

[22] BARAM. Available online at: https://baramcfd.org/. Accessed on: 24/09/2024.

[23] E. Hosseini. CFD analysis of the aerodynamic characteristics of biconvex airfoil at compressible and high Mach numbers flow. *SN Applied Sciences*, 1(10):1283, 2019. doi:10.1007/s42452-019-1334-2.

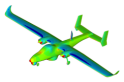
[24] I. Celik, U. Ghia, P. Roache, C. Freitas, H. Coloman, and P. Raad. Procedure of Estimation and Reporting of Uncertainty Due to Discretization in CFD Applications. *J. Fluids Eng.*, 130:078001, July 2008. doi:10.1115/1.2960953.

[25] C. Rumsey. Turbulence Modeling Resource - The Menter Shear Stress Transport Turbulence Model. Available online at: `https://turbmodels.larc.nasa.gov/sst.html`. *NASA, Langley Research Center*. Last accessed on: 26/09/2024.

[26] F. Menter, R. Sechner, and A. Matyushenko. Best Practice: RANS Turbulence Modeling in Ansys CFD. *Ansys*.

[27] W. Sutherland. LII. The viscosity of gases and molecular force. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 36(223):507–531, 1893. doi:10.1080/14786449308620508.

[28] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2): 182–197, 2002. doi:10.1109/4235.996017.



AeroBest 2025
Lisboa, 22–24 April 2025
©ECCOMAS, Portugal

APPLYING MDO TO AIRCRAFT STRUCTURAL DESIGN AT AIRBUS DEFENCE AND SPACE; FROM PUNCH CARDS TO DIGITAL TWINS

Ögmundur Petersson

Airbus Defence and Space
Department of Stress Methods and Optimisation
Rechliner Straße, 85077, Manching, Germany
oegmundur.petersson@airbus.com, https://www.airbus.com/en/

Abstract. *MDO has a long history at Airbus Defence and Space dating back to the origins of the Eurofighter programme in the mid-1980s. In the past four decades, MDO has been applied successfully to numerous programmes at Airbus DS. For the structural analysis and optimisation task, an in-house tool named Lagrange has been in continuous use and development during all this time, providing some unique capabilities; combining as it does a finite element solver, numerous analysis disciplines, a collection of structural sizing criteria and powerful optimisation algorithms. Full access to the source code enables efficient determination of design sensitivities for gradient based optimisation algorithms, necessary for high-dimensional problems. For the structural sizing task, a direct parameterisation of the properties of the elements of the finite element model has been chosen as this is both flexible and straight forward to implement. The global finite element model of the aircraft provides the internal loads on the structure from the external forces and is an essential part of the development process. Stiffened skin panels are usually represented by two-dimensional shell elements reinforced by one dimensional elements, resulting in a set of patches closely matching the geometry of the buckling fields of the actual structure. Our optimisation process uses the GFEM without any modifications ensuring consistency in the design process; these patches form the basis of the parameterisation of the optimisation problem, both for the definition of the optimisation variables and for the structural sizing constraints. Linking of design variables within a patch and based on geometry can be used to enforce design rules such es e.g., symmetric structures and balanced and symmetric composite layups. The disadvantage of this parameterisation is, nevertheless, that it does not closely represent the manufacturing processes and may thus easily lead to optimisation results that can not be manufactured. To counteract this, Lagrange contains several manufacturing constraints forcing the optimiser towards manufacturable designs. Most recently these have been extended to return discrete optimised layups fulfilling stacking rules. The size of the resulting optimisation problem is mainly determined by the granularity of the parameterisation and the number of load cases. The latter, especially, has increased significantly in recent years, leading to an increase in the optimisation problem size far outpacing the increase in single core processing speed. This has been tackled by using active set optimisation algorithms and taking advantage of parallel processing and multi-threading on high-performance clusters to evaluate load cases and individual element results and sensitivities in parallel. A drive towards greater use of MDO in early design phases has highlighted a need for a parametric design capability realised in a new tool*

named Descartes. It combines a standardised parametric description of the aircraft with a geometry kernel and meshing capabilities to efficiently create parametric geometries and simulation models for various disciplines, allowing consistent updates of these for parameter studies or shape optimisation tasks. This is an essential enabler for design studies in early design phases based on numerical simulation. More recently, advances in sensitivity analysis for coupled problems have allowed the creation of an optimisation framework coupling high-fidelity aerodynamic analysis and the structural sizing process. Full access to the source code enables calculation of the coupled flexible adjoint as has been demonstrated by aeroelastic tailoring of a high aspect-ratio transport configuration showing the trade-off between structural mass and aerodynamic performance. The MDO capabilities at Airbus DS are in constant development and will continue to provide a valuable capability during a trend towards increasing complexity and levels of interaction in aircraft design.

Keywords: Industrial perspective, structural sizing, parameterisation, high-fidelity MDO, manufacturing constraints

Index of Authors

––/ A /––

Afonso, F 799, 814
Alaimo, A 527
Albuquerque, P F 193
Aleisa, H 169
Almutairi, M 100
Athayde, A 562
Azevedo, J L 251
Azinheira, J R 562

––/ B /––

Baars, S 719
Balakrishnan, N 67
Bandopadhyay, S 340
Banerjee, S 678
Bedir, B B 374
Ben Amar, A 659
Benard, E 427
Bergmann, M 16
Bisagni, C 494
Blondeau, C 16
Brück, V 407
Brezillon, J 149, 407
Brižs, J 111
Brousset-Matheu, D 608
Brás, M 678

––/ C /––

Carrera, A 358
Catalanotti, G 527
Cavallaro, R 111, 444, 628
Cerezo-Magaña, M 593
Chandrasekaran, P . . . 287, 328, 645
Chapron, M 16
Chauhan, S S 251
Chen, E 659
Chen, L 2
Cini, A 444, 628
Coniglio, S 340
Conlan-Smith, C 149

––/ D /––

Dechaume, A 735
Defoort, S 608
Diez, M 52
Dunning, P 100

––/ E /––

Ehrmanntraut, S 705
Evers, N 316

––/ F /––

Fadli, I 149
Faraj, R 583
Farinha, P 494
Fasquel, J 735
Fereidooni, A 659
Figueiras, I 799
Figué, J-F 735
Fiore, M 233

––/ G /––

Güler, E 287, 645
Gallard, F 340, 407, 735
Gauger, N R 2, 444
Gazaix, A 407, 735
Gil, P J S 775
Giret, J-C 735
Gogu, C 133
Gogulapati, A 304, 508
Gonel, N 693
González, E V 527
Gottfried, S 705
Grünfelder, N 37
Grewal, A 581, 659

––/ H /––

Höing, C 705
Halila, G L O 251
Hughes, T J R 536

––/ I /––

Iacolino, G 527
Iemma, U 52

––/ J /––

Jankowski␣L 583

––/ K /––

Kälber, L 37
Kaba, R 755
Kanti, A 304
Kaur, S 67
Kilic, D 328
Konowrocki, R 583
Kontis, K 169
Koyuncu, E 607

––/ L /––

Lau, F 233, 799, 814
Li, X 549

––/ M /––

Maheri, A 100
Mantegna, G 527
Mao, H 549
Marta, A C 84, 474, 775
Martins, J R R A . . 251, 303, 407
Matos, J 799
Matos, N M B 84
Mehta, A 508
Mishra, P N 304, 508
Moosavi, R 302
Morlier, J 133, 427, 693, 799
Moutinho, A 562
Musso, G 799

––/ N /––

Nikbay, M . . 1, 17, 169, 287, 328, 374, 645, 755
Norczyk Simon, P 111
Nozawa, S 317

––/ O /––

Oezkaya, E 2
Olivares, A 593
Orlando, C 527

––/ P /––

Pacheco, L M M 474
Pagani, A 358
Palaio, A M 775
Paudel, A 67
Pereira, A N 193
Petersson, Ö 474, 830
Petrolo, M 358
Pinheiro, L D 84
Pirlepeli, B 169, 374
Pisarski, D 583
Poplawski, B 583

––/ Q /––

Quagliarella, D 52
Quibén Figueroa, R 628

––/ R /––

Römer, U 719
Rebelo, A 233
Ribeiro, I 799
Ricci, S 387
Ricken, T 37
Roboam, X 608
Roncolini, E 387
Roussouly, N 407
Roß, F 705
Ruiz Jimenez, G 407
Rumpler, R 549

––/ S /––

Sagebaum, M 444
Salah El Din, I 16
Santos, H A F A 536
Sapre, S 133
Sareni, B 608

Saves, P 693
Scalia, L 444
Serani, A 52
Serhat, G 317
Shende, N V 67
Shrivastava, C 316
Silva, A M P 193
Silva, G P 251
Silva, J 814
Sipp, D 16
Soler Arnedo, M 628
Souza, A 799
Squillace, D 52
Stück, A 705
Staffetti, E 593
Stumpf, E 213, 269
Suleman, A 678
Sy, O 427

––/ T /––

Thanissaranon, K 149, 407
Tibert, G 549
Toffol, F 387, 494
Toprak, H N 607
Tumino, D 527

––/ V /––

Van, E N 608
Volle, F 474

––/ W /––

Waschinsky, N 37
Wickramasinghe, V 659

––/ Y /––

Yildiz, S 17, 169

––/ Z /––

Zamani, D 358
Zell, I 213
Zhang, D Z 111
Zimmnau, M 213, 269

