

Aerodynamic Shape Optimization of an Oblique Wing using the ADjoint Approach

Charles A. Mader¹, Andre C. Marta², and Joaquim R. R. A. Martins¹

¹ *Institute for Aerospace Studies, University of Toronto, Toronto, ON, M3H 5T6, Canada*

² *Stanford University, Stanford, CA, 94305, U.S.A.*

Email: *mader@utias.utoronto.ca*

ABSTRACT

The ADjoint method is applied to a three-dimensional Computational Fluid Dynamics (CFD) solver to generate the sensitivities required for aerodynamic shape optimization. The ADjoint approach selectively uses Automatic Differentiation (AD) to generate the partial derivative terms in the discrete adjoint equations. By selectively applying AD techniques, the computational cost and memory overhead incurred by using AD are significantly reduced, while still allowing for the accurate treatment of arbitrarily complex governing equations and boundary conditions. Once formulated, the discrete adjoint equations are solved using the Portable Extensible Toolkit for Scientific computation (PETSc). With this approach, the adjoint vector is used to compute the total sensitivities required for aerodynamic shape optimization of a complete aircraft configuration. The resulting sensitivities are compared with finite difference derivatives to verify accuracy.

1 INTRODUCTION

Adjoint methods for sensitivity analysis involving partial differential equations (PDE's) have been known and used for over three decades. They were first applied to solve optimal control problems and thereafter used to perform sensitivity analysis of linear structural finite-element models. The first application to fluid dynamics is due to Pironneau [16]. The method was then extended by Jameson to perform airfoil shape optimization [7]. This method was also used, in conjunction with a Newton-Krylov solver, for airfoil shape optimization by Nemec and Zingg [15]. Since then, the adjoint method has been developed for more complex problems, leading to its application to the design optimization of complete aircraft configurations considering aerodynamics alone [17, 18], as well as

aerodynamic and structural interactions [9]. The adjoint method has also been generalized for multidisciplinary systems [10]. However, because of the complexity inherent in developing a manually differentiated adjoint, these methods have not made their way into widespread general use.

2 METHODOLOGY

To remove this impediment, the ADjoint (Automatic Differentiation Adjoint) method has been developed. As described by Martins et al. [11, 12], the ADjoint method uses reverse mode automatic differentiation to accurately and efficiently compute the partial derivative terms required by the discrete adjoint equations. Once formulated from these partial derivatives, the discrete adjoint equations are solved using the Portable Extensible Toolkit for Scientific Computation (PETSc), which has several built-in capabilities including a sparse matrix solver.

In this work, the ADjoint method is applied to NSSUS, a flow solver developed at Stanford University through the ASC turbomachinery program [1]. NSSUS is a vertex-centered flow solver capable of solving the Reynolds-averaged Navier-Stokes equations with a variety of boundary conditions and turbulence models. It has also been developed to handle magnetohydrodynamics (MHD) equations [8]. However, in this work only the Euler and MHD equations have been considered. Total sensitivities will be developed for the force coefficients C_L and C_D and the moment coefficients C_{M_x} , C_{M_y} and C_{M_z} with respect to the coordinates of the volume grid. The sensitivities can then be combined with geometry software to obtain the derivatives with respect to a variety of geometric design variables including sweep, span, twist and the wing surface shape.

3 MOTIVATION

Multidisciplinary Design Optimization (MDO) is a method that can be used to find better design solutions than those produced by existing design methods. In 2004, Martins et al. [9] showed that MDO can find improve design solutions as compared to traditional design methods. However, because of the computational cost of a multidisciplinary analysis, particularly the aerodynamic portion thereof, efficient sensitivity analysis is a necessary prerequisite for gradient based optimization.

The usefulness of the adjoint method lies in the fact that it is an extremely efficient approach to compute the sensitivity of one function of interest with respect to many parameters. However, as mentioned in the introduction, the complexity of implementing an adjoint in CFD codes has prevented it from coming into widespread use. In this research we show a method for implementing adjoint sensitivities that provides excellent accuracy and significantly reduces the complexity of such process. A method that will make the use of adjoint methods accessibly to a larger range of users.

4 BACKGROUND

We will now derive the adjoint equations for the particular case of our flow solver. The governing equations for the three-dimensional Euler equations are,

$$\frac{\partial w}{\partial t} + \frac{\partial f_i}{\partial x_i} = 0, \quad (1)$$

where x_i are the coordinates in the i^{th} direction, and the state and the fluxes for each cell are

$$w = \begin{bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho u_3 \\ \rho E \end{bmatrix}, \quad f_i = \begin{bmatrix} \rho u_i \\ \rho u_i u_1 + p \delta_{i1} \\ \rho u_i u_2 + p \delta_{i2} \\ \rho u_i u_3 + p \delta_{i3} \\ \rho u_i H \end{bmatrix} \quad (2)$$

It must be noted that this derivation is presented here for the Euler equations, but since our approach is only based on the existence of code that computes the *residual* of the governing equations, the procedure can be extended to the full Reynolds-averaged Navier–Stokes equations without modification. Note that the code for the residual computation is assumed to include the application of the required boundary conditions (however complex they may be) and any artificial dissipation terms that may need to be added for numerical stability.

A coordinate transformation to computational coordinates (ξ_1, ξ_2, ξ_3) is used. This transformation is de-

finied by the following metrics,

$$K_{ij} = \begin{bmatrix} \partial X_i \\ \partial \xi_j \end{bmatrix}, \quad J = \det(K), \quad (3)$$

$$K_{ij}^{-1} = \begin{bmatrix} \partial \xi_i \\ \partial X_j \end{bmatrix}, \quad S = JK^{-1}, \quad (4)$$

where S represents the areas of the face of each cell projected on to each of the physical coordinate directions.

The Euler equations in computational coordinates can then be written as,

$$\frac{\partial Jw}{\partial t} + \frac{\partial F_i}{\partial \xi_i} = 0, \quad (5)$$

where the fluxes in the computational cell faces are given by $F_i = S_{ij} f_j$.

In semi-discrete form the Euler equations are,

$$\frac{dw_{ijk}}{dt} + \mathcal{R}_{ijk}(w) = 0, \quad (6)$$

where \mathcal{R} is the residual described earlier with all of its components (fluxes, boundary conditions, artificial dissipation, etc.).

Thus, the adjoint equations can be written for this flow solver as,

$$\left[\frac{\partial \mathcal{R}}{\partial w} \right]^T \Psi = - \frac{\partial I}{\partial w}, \quad (7)$$

where Ψ is the *adjoint vector*. The total sensitivity in this case is,

$$\frac{dI}{dx} = \frac{\partial I}{\partial x} + \Psi^T \frac{\partial \mathcal{R}}{\partial x}. \quad (8)$$

We propose to compute the partial derivative matrices $\partial \mathcal{R} / \partial w$, $\partial I / \partial w$, $\partial I / \partial x$ and $\partial \mathcal{R} / \partial x$ using automatic differentiation instead of using manual differentiation or using finite differences. Where appropriate we will use the reverse mode of automatic differentiation.

5 IMPLEMENTATION

To compute the desired sensitivities, we need to form the discrete adjoint equations (7), solve them and then use the total sensitivity equation (8). We will use automatic differentiation to generate code that computes each of the partial sensitivity matrices present in these equations.

For the purpose of demonstration in this paper, we compute the sensitivity of the five force and moment coefficients, C_L , C_D , C_{M_x} , C_{M_y} and C_{M_z} with respect to the mesh coordinates, i.e., $I = C_i$ and $x = X(i, j, k)$.

The NSSUS solver is a new finite-difference, higher-order solver that has been developed at Stanford University under the Advanced Simulation and Computing (ASC) program sponsored by the Department of Energy [1]. It is a generic node-centred, multi-block, multi-processor solver, currently only for the Euler equations, but soon to be extended to the Reynolds-averaged Navier–Stokes. The finite-difference operators and artificial dissipation terms follow the work by Mattsson [13, 14] and the boundary conditions are implemented by means of penalty terms, according to the work by Carpenter [5, 6]. Addition magnetic source terms were also included in this solver so that MHD computations could be performed. Despite being capable of performing computations up to eight-order accuracy, the implementation of the adjoint solver was restricted to second-order for simplicity. The extension should be straightforward to accomplish. Also note that the method described here does not depend on the form or contents of the governing equations, it simply requires that there be a code that computes the *residual* of the governing equations. Therefore it will be relatively straight forward to extend it to governing equations (such as the Reynolds-Averaged Navier-Stokes equations).

5.1 Brief summary of the discretization

This section summarizes the spatial discretization. Further details can be found in Mattsson [13, 14] and Carpenter [5, 6]. The spatial part of equation (5) is discretized on a block-by-block basis. The internal discretization is straightforward and only requires the first neighbours in each coordinate direction for the inviscid fluxes and the first and second neighbours for the artificial dissipation fluxes, see figure 1. However the boundary treatment needs to be explained in more detail. As the finite-difference scheme only operates on the nodes of a block, one-sided difference formulae are used near block boundaries (be it a physical or an internal boundary). Consequently, the nodes on the interface of internal boundaries are multiply defined, see figure 2.

These multiple instances of the same physical node are driven to the same value by means of a penalty term, i.e. an additional term is added to the residual \mathcal{R} which is proportional to the difference between the instances. This reads

$$\mathcal{R}_{\text{blockA}}^i = \mathcal{R}_{\text{blockA}}^i + \tau(w_{\text{blockB}}^i - w_{\text{blockA}}^i) \quad (9)$$

and a similar expression for $\mathcal{R}_{\text{blockB}}^i$. In equation (9), τ controls the strength of the penalty and is a combination of a user defined parameter and the local flow

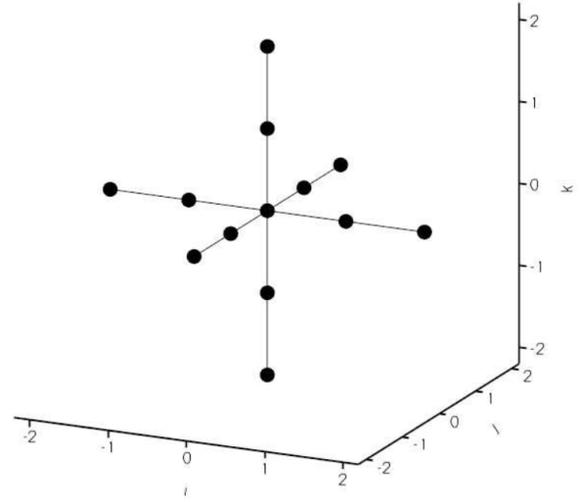


Figure 1: Stencil for the vertex-centred residual computation.

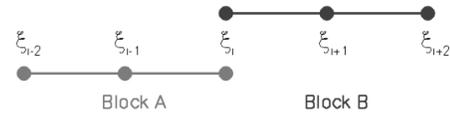


Figure 2: Block-to-block boundary stencil.

conditions, see Mattsson [13] for more details. Hence, the residual of nodes on a internal block boundary is a function of its local neighbours in the block and the corresponding instance in the neighbouring block.

The boundary condition (BC) treatment is very similar to the approach described above except that the penalty state used in equation (9) is now determined by the boundary conditions.

5.2 Computation of $\partial R / \partial w$

To simplify the following discussion, we define the following numbers,

N_n : The number of nodes in the domain. For three-dimensional domains where the Navier–Stokes equations are solve, this can be $O(10^6)$.

N_s : The number of nodes in the stencil whose variables affect the residual of a given node. In our case, when considering inviscid and dissipation fluxes, the stencil is as shown in Figure 1 and $N_s = 13$.

N_w : The number of flow variables (and also residuals)

for each node. In our case $N_w = 5$.

The flux Jacobian, $\partial R/\partial w$, is independent of the choice of function or design variable: it is simply a function of the governing equations, their discretization and the problem boundary conditions. To compute it we need to consider the routines in the flow solver that, for each iteration, compute the residuals based on the flow variables, w . In the following discussion we note that the residual computations are carried out within the context of the NSSUS flow solver. The computation of the residual in NSSUS can be summarized as follows,

1. Compute inviscid fluxes: For our inviscid flux discretization the only flow variables, w , that influence the residual at a node are the flow variables at that node and at the six nodes directly adjacent to the node.
2. Compute dissipation fluxes: For each of the N_n nodes in the domain, compute the contributions of the flow variables on the residual at that node. For this portion of the residual, the solution at the current node and the 12 adjacent nodes in each of the three directions need to be considered.
3. (To be implemented) Compute viscous fluxes (with similar stencil implications: only the nodes directly surrounding the nodes in question need to be considered).
4. Compute magnetic source terms: This only depends on the flow variables in the current node.
5. Apply boundary conditions: Additional penalty terms are added to enforce the BCs, see section 5.1. Note that internal block boundaries are also considered as boundaries.

Note that to compute the residuals over the domain, three nested loops (in each of the three directions) are used and that the correct value of the residual for any given node is only obtained at the end, when all contributions have been accounted for. Using Automatic Differentiation (AD) on this original routine containing several loop would entail several unnecessary calculations. Thus, to make the implementation of the discrete adjoint solver more efficient, it was necessary to re-write the flow residual routine such that it computed the residual for a single specified node. The re-engineered residual routine is a function with the residuals at a given node, r_{Adj} , returned as an output argument and the stencil of flow variables, w_{Adj} , that affect the residuals at that node provided as an input argument.

Now we have a routine that computes N_w residuals in a given cell. These residuals get contributions from all $(N_w \times N_s)$ flow variables in the stencil. Thus there are $N_w \times (N_w \times N_s)$ sensitivities to be computed for each cell, corresponding to N_w rows in the $\partial R/\partial w$ matrix. Each of these rows contains no more than $(N_w \times N_s)$ nonzero entries. The analog in the forward mode would be to consider all of the residuals affected by the states in a single cell. Theoretically, one could compute the derivatives of the $N_w \times N_s$ residuals affected by a single state, thus for a single cell, one would again obtain $N_w \times (N_w \times N_s)$ derivatives. However, because of the one way dependence of the residual on the states, this does not prove to be the case. In the reverse mode, all of the derivatives in the stencil can be calculated from one residual calculation. All of the information for that calculation is contained in a single stencil. For the forward mode, one would need to calculate all of the $N_w \times N_s$ residuals in the inverse stencil. This requires the addition of all of the states in those stencils as well. Thus, rather than having a single calculation with $N_w \times N_s$ states involved to get $N_w \times N_s$ derivative values, one requires N_s residual calculations and many extra states to get the same number of derivative components. Thus, it quickly becomes apparent that the reverse mode is much more efficient in this case.

5.3 Computation of $\partial C_i/\partial w$

The RHS of the adjoint equations (7) for the functions of interest $C_D, C_L, C_{M_x}, C_{M_y}$ and C_{M_z} are easily computed for this flow solver. Because this specific flow solver works with primitive variables $w = (\rho, u, v, w, p)$ and since, for inviscid flow, $C_D, C_L, C_{M_x}, C_{M_y}$ and C_{M_z} are simple surface integrations of the pressure, the derivatives $\partial C_D/\partial w$ and $\partial C_L/\partial w$ are always zero except for $w_5 (= p)$. Therefore, it became trivial to derive analytically the expression for these partial derivatives from the flow solver routine that calculated the functions of interest.

5.4 Adjoint Solver

The adjoint equations (7) can be re-written for the example case as follows,

$$\begin{bmatrix} \partial \mathcal{R} \\ \partial w \end{bmatrix}^T \Psi = -\frac{\partial C_D}{\partial w}. \quad (10)$$

With the two partial derivatives in this equation computed, the adjoint vector, Ψ can now be computed. As we have pointed out, both the flux Jacobian and the right hand side in this system of equations are very

sparse. To solve this system efficiently, and having in mind that we want to have a parallel adjoint solver, we decided to use PETSc [3, 2, 4]. PETSc is a suite of data structures and routines for the scalable, parallel solution of scientific applications modeled by PDEs. It employs the message passing interface (MPI) standard for all interprocessor communication. Using PETSc's data structures, $\partial R/\partial w$ and $-\partial C_D/\partial w$ as sparse entities. Once the sparse matrices were setup, PETSc's built in, parallel, Krylov solver was used to compute the adjoint solution.

5.5 Computation of $\partial \mathcal{R}/\partial X(i, j, k)$

The computation of the partial sensitivity of the residuals with respect to the mesh coordinates, $\partial \mathcal{R}/\partial X(i, j, k)$, was accomplished using an extension of the method used to compute the flux jacobian, $\partial \mathcal{R}/\partial w$. The stencil based approach still applies, however in this situation, the metric transformations need to be taken into account in the residual computation. Once again the first two layers of adjacent nodes in each of the three coordinate directions are required for each nodal residual. This leads to the same size and shape stencil that was present in the flux jacobian computation. Therefore, the same code can be used for both computations. Once again, the matrix is very sparse, so the PETSc data structures are used to store the matrix.

5.6 Computation of $\partial C_i/\partial X(i, j, k)$

The final partial derivative term required for the total sensitivity equation is the explicit effect of the mesh coordinates on the force and moment coefficients. This effect shows up as a change in direction of the normal vectors and associated areas for the surface nodes in the mesh. However, because the force and moment coefficients are a sum over the entire grid, the small stencil used to compute the flux jacobian is no longer valid. In this case, the stencil must become the entire mesh. On the other hand, this does provide some other benefits. For example, now we have a situation where a single reverse mode differentiation will return all of the sensitivities required for one force or moment coefficient. While this does lead to slightly higher memory costs for this computation, it does lead to very fast derivatives. Once again, due to the sparsity of the matrix, the derivative values are stored in PETSc.

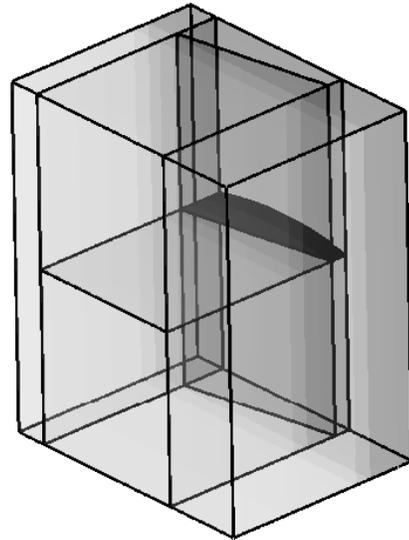


Figure 3: Hypersonic half-body

5.7 Total Sensitivity Equation

The total sensitivity (8) in this case can be written as,

$$\frac{dC_i}{dX(i, j, k)} = \frac{\partial C_i}{\partial X(i, j, k)} + \psi^T \frac{\partial \mathcal{R}}{\partial X(i, j, k)}, \quad (11)$$

where the objective function C_i represents the i^{th} coefficient of interest and the independent variable $X(i, j, k)$ represents the mesh coordinates at location i, j, k . With all four of the partial derivatives matrices computed, and the adjoint equation solved, all that remains is to multiply the terms together as shown in equation 11. This will leave us with a single vector of length $3 \times N_n$ for each force or moment coefficient of interest.

6 RESULTS

The following results are based on three distinct test cases. The first test case, used to verify the sensitivity results, is a half body model of simple hypersonic vehicle (figure 3). This is a six block test case and has been run at Mach=3.0. The remaining two test cases, shown in figures 4 and 5 were used mostly for timing purposes. The first is a more complex hypersonic vehicle, intended as an analog to the NASA X-43 test plane, simulated at Mach = 5. While the last test case is an oblique flying wing modeled at Mach = 1.5. All cases use Euler wall boundary conditions for the wing surface.

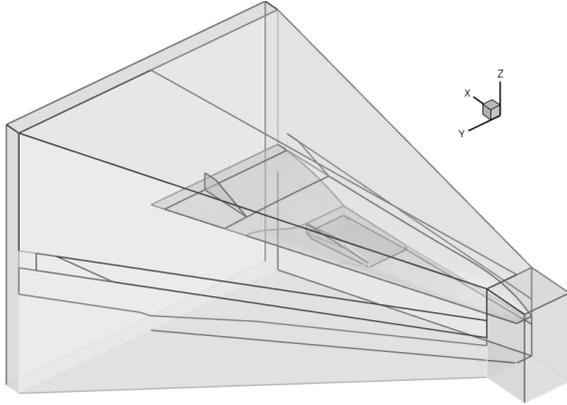


Figure 4: X-43 Analog

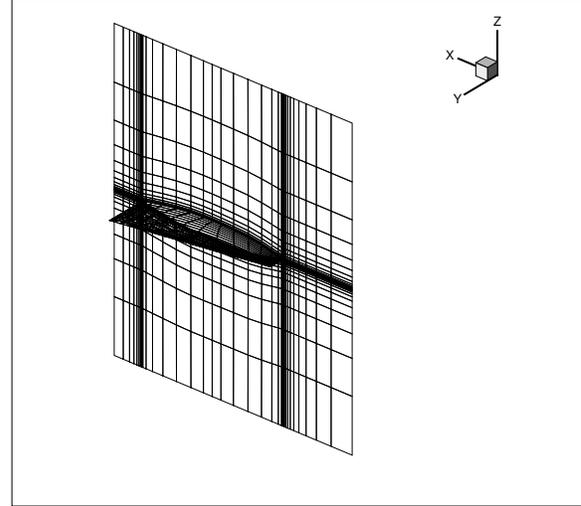


Figure 6: Mesh for hypersonic testcase

Table 1: Mesh coordinate sensitivity verification: $\frac{dJ}{dX(i,j,k)}$ for hypersonic test case, block 3, index 2,7,13.

Control Node	Cost Fun. J	Adjoint	Finite-Diff. (1×10^{-4})	Δ
1	C_L	-0.001015	-0.000910	13.7%
	C_D	-0.000218	-0.000230	5.02%
Control Node	Cost Fun. J	Adjoint	Finite-Diff. (1×10^{-3})	Δ
1	C_L	-0.001015	-0.001051	0.39%
	C_D	-0.000218	-0.000222	1.37%

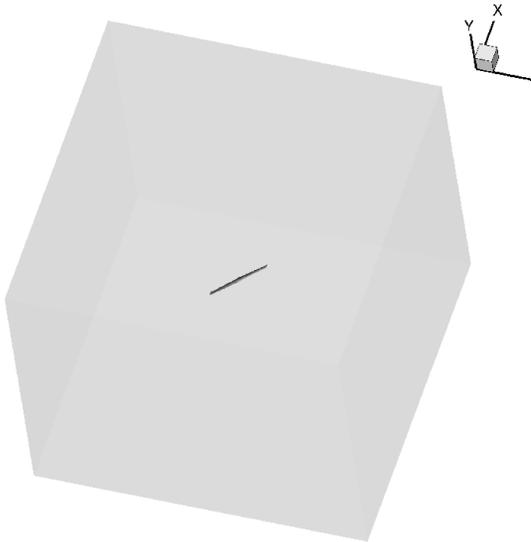


Figure 5: Oblique flying wing

Figures 6 through 11, show the meshes and flow solutions for the three test cases.

Accuracy results for the hypersonic test case are shown in Table 1. This comparison to a forward finite difference shows acceptable accuracy. The relative error is less than 1.5% for the larger step size. However, because of the sensitivity of the finite difference result to step size, it is uncertain which of the two results is more correct. We plan to compare these results against complex step results in the near future and expect this comparison to show an accuracy of 7 or more digits. This expectation is based on the results of a previous comparison done with a single block version of the ADjoint implemented on the Sumb flow solver [12].

The timing results for the oblique wing and X-43 test cases are shown in Table 2. As can be seen for both cases, the ADjoint solution is less expensive than the flow solution, varying from 2/3 to 1/50 of the flow solution time, depending on the test case.

When it comes to comparing the performance of the

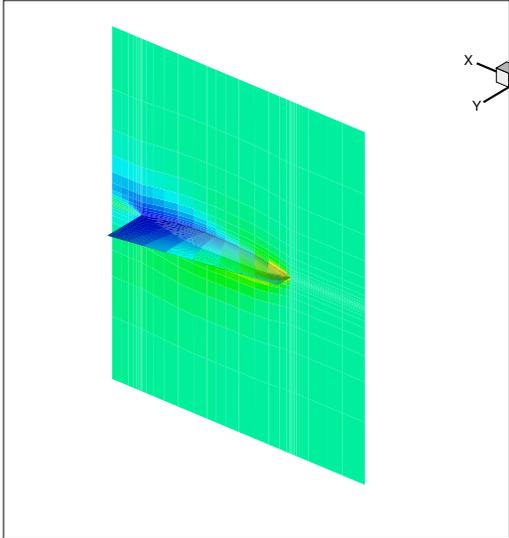


Figure 7: Contour plot of pressure

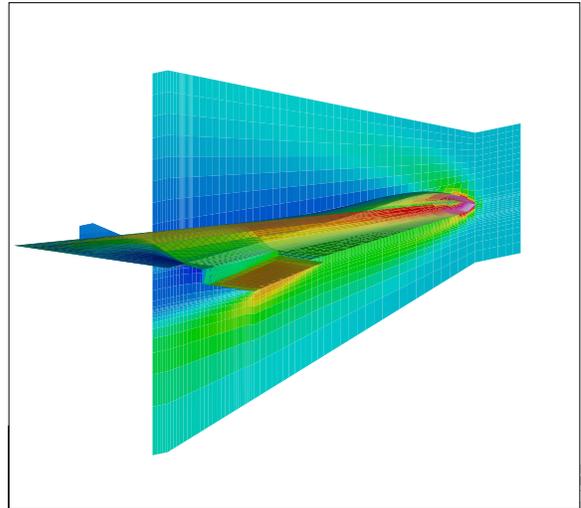


Figure 9: Contour plot of pressure

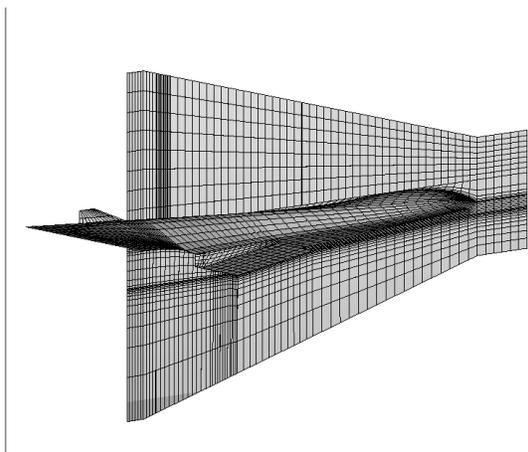


Figure 8: Mesh for hyperplane testcase

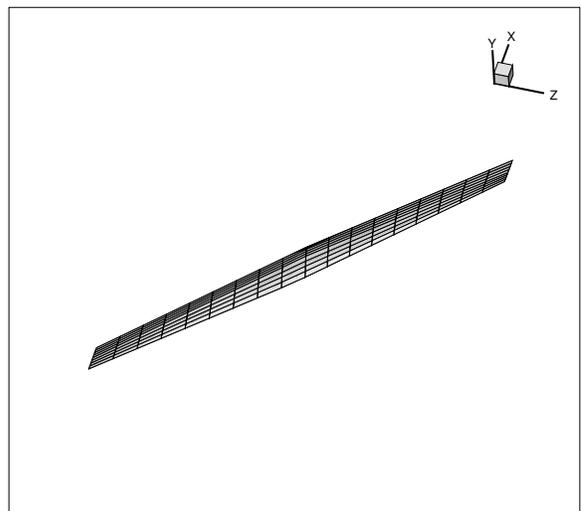


Figure 10: Mesh for oblique wing testcase

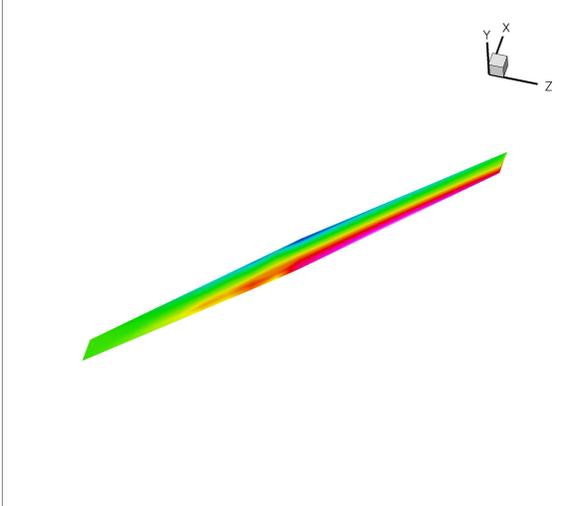


Figure 11: Contour plot of pressure

	Oblique Wing	Hyperplane
No.Processors	1	2
No. Nodes	21820	50776
Flow solution	63.98	1771.29
ADjoint	40.17	46.58
Breakdown:		
Setup PETSc Variables	0.08	0.07
Compute flux Jacobian	7.80	15.85
Compute grid partial	15.76	20.76
Compute RHS	0.00	0.00
Solve the adjoint equations	16.33	9.62
Compute the total sensitivity	0.20	0.28

Table 2: ADjoint computational cost breakdown (times in seconds)

various components in the adjoint solver, we found that most of the time was spent in the solution of the adjoint equations and thus all the automatic differentiation sections performed very well. The costliest of the automatic differentiation routines are the computation of the flux Jacobian and the computation of the mesh partial, $\partial \mathcal{R} / \partial X(i, j, k)$. When one takes into consideration the number of terms in these matrices, spending less than 25% of the flow solution time in this computation is quite impressive.

Also, while we have not done rigorous testing on the memory requirements of this code, preliminary observations indicate that the memory required for the ADjoint code is approximately ten times that required for the original flow solver. Given the pattern of use on most parallel computers, this is well within the acceptable limits for an adjoint code.

The last set of results shown here are the lift and drag sensitivities of the oblique wing case, shown in figures

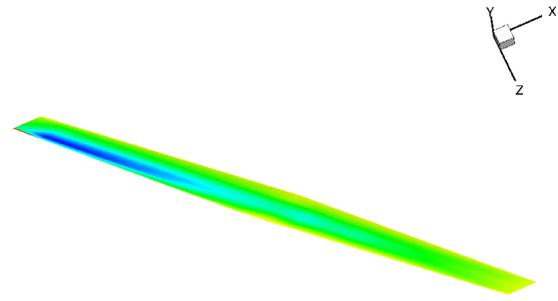


Figure 12: Lift sensitivities

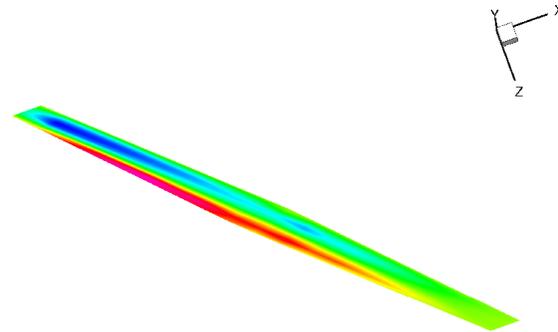


Figure 13: Drag Sensitivities

12 and 13.

As can be seen in the figures, there is a higher sensitivity to both lift and drag on the forward swept portion of the wing. Also as we would expect, there is a high sensitivity to drag along certain portions of the wing leading edge. These are all characteristics that should lead to interesting optimization results.

7 CONCLUSIONS

In this work we have applied the ADjoint method to the NSSUS flow solver to generate the mesh coordinate sensitivities required to perform aerodynamic shape optimization. We have validated the resulting sensitivities through comparison to finite difference results, though we expect to do a more thorough comparison of the results against results from the complex step method in the near future. The implementation has also been shown to be very efficient, with the total ADjoint solution taking less time than the flow solver. Finally, we have shown some overall sensitivity dis-

tributions for the oblique wing case, as an example of the results that the ADjoint implementation on NSSUS can generate. These results also form the basis of the sensitivities required for aerodynamic shape optimization of an oblique wing.

ACKNOWLEDGMENTS

The first two authors are grateful for the funding provided by the Canada Research Chairs program and the Natural Sciences and Engineering Research Council.

REFERENCES

- [1] ASC. Advanced simulation and computing. <http://www.llnl.gov/asc>, 2007.
- [2] S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.
- [3] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. PETSc Web page, 2001. <http://www.mcs.anl.gov/petsc>.
- [4] 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.
- [5] M. H. Carpenter, D. Gottlieb, and S. Abarbanel. Time-stable boundary conditions for finite-difference schemes solving hyperbolic systems: Methodology and application to high-order compact schemes. *Journal of Computational Physics*, 111(2):220–236, Apr. 1994.
- [6] M. H. Carpenter, J. Nordström, and D. Gottlieb. A stable and conservative interface treatment of arbitrary spatial accuracy. *Journal of Computational Physics*, 148(2):341–365, Jan. 1999.
- [7] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3):233–260, sep 1988.
- [8] A. C. Marta and J. J. Alonso. High-speed mhd flow control using adjoint-based sensitivities. In *Proceedings of the 14th AIAA/AHI Space Planes and Hypersonic Systems and Technologies Conference*, Canberra, Australia, 2006. AIAA 2006-8009.
- [9] J. R. R. A. Martins, J. J. Alonso, and J. J. Reuther. High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41(3):523–530, 2004.
- [10] J. R. R. A. Martins, J. J. Alonso, and J. J. Reuther. A coupled-adjoint sensitivity analysis method for high-fidelity aero-structural design. *Optimization and Engineering*, 6(1):33–62, March 2005.
- [11] J. R. R. A. Martins, J. J. Alonso, and E. van der Weide. An automated approach for developing discrete adjoint solvers. In *Proceedings of the 2nd AIAA Multidisciplinary Design Optimization Specialist Conference*, Newport, RI, 2006. AIAA 2006-1608.
- [12] J. R. R. A. Martins, C. A. Mader, and J. J. Alonso. Adjoint: An approach for rapid development of discrete adjoint solvers. In *Proceedings of the 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Portsmouth, VA, 2006. AIAA 2006-7121.
- [13] K. Mattsson and J. Nordström. Summation by parts operators for finite difference approximations of second derivatives. *Journal of Computational Physics*, 199(2):503–540, Sept. 2004.
- [14] K. Mattsson, M. Svärd, and J. Nordström. Stable and accurate artificial dissipation. *Journal of Scientific Computing*, 21(1):57–79, Aug. 2004.
- [15] M. Nemeć and D. W. Zingg. Newton-krylov algorithm for aerodynamic design using the navier-stokes equations. *AIAA Journal*, 40(6):1146–1154, 2002.
- [16] O. Pironneau. On optimum design in fluid mechanics. *Journal of Fluid Mechanics*, 64:97–110, 1974.
- [17] J. Reuther, J. J. Alonso, A. Jameson, M. Rimplinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part I. *Journal of Aircraft*, 36(1):51–60, 1999.
- [18] J. Reuther, J. J. Alonso, A. Jameson, M. Rimplinger, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers: Part II. *Journal of Aircraft*, 36(1):61–74, 1999.