

Adjoint-Based Algorithms for Adaptation and Design Optimization on Unstructured Grids

Eric J. Nielsen
NASA Langley Research Center
Hampton, VA 23681
USA

Abstract

Schemes based on discrete adjoint algorithms present several exciting opportunities for significantly advancing the current state of the art in computational fluid dynamics. Such methods provide an extremely efficient means for obtaining discretely consistent sensitivity information for hundreds of design variables, opening the door to rigorous, automated design optimization of complex aerospace configurations using the Navier-Stokes equations. Moreover, the discrete adjoint formulation provides a mathematically rigorous foundation for mesh adaptation and systematic reduction of spatial discretization error. Error estimates are also an inherent by-product of an adjoint-based approach, valuable information that is virtually non-existent in today's large-scale CFD simulations.

An overview of adjoint-based algorithm work at NASA Langley Research Center is presented, with examples demonstrating the potential impact on complex computational problems related to design optimization as well as mesh adaptation.

Introduction

Aerodynamic solvers based on the Euler and Navier-Stokes equations have become increasingly prevalent in realistic aerospace applications during the last few decades. Structured grid solvers have seen widespread use since the early 80's, with overset and unstructured grid discretizations becoming more popular throughout the 90's. Ever more powerful computing hardware has enabled engineers to rapidly compute flowfields over highly complex geometries across the speed regime. In addition to high Reynolds number transonic cruise conditions, computations over complicated high-lift systems are routinely performed, as well as solutions for supersonic applications and hypersonic re-entry problems. Thanks to highly touted successes in such application areas, the field of computational fluid dynamics (CFD) is widely believed to be a mature practice, with little research potential or process improvements left to be had. This is unfortunate, as there remain a number of largely uncharted and potentially highly fruitful research avenues left to be explored.

Today's large-scale CFD solvers are generally aimed at the analysis problem. That is, given a vehicle configuration of interest, the user will discretize the problem using some form of a computational grid and the solver package will then compute the flowfield around the geometry and report quantities of engineering interest, such as lift, drag, or heating. Unfortunately, this practice is highly susceptible to input error and the user's skill level; today's grid generation practices are generally much more of an art than a science. An experienced user may have some feel for where to cluster grid resolution on a known configuration; however, the process can easily break down on geometries where little a priori knowledge exists. Moreover, the engineer typically receives no quantifiable measure of accuracy on the outputs from the solver. Instead, the process is highly subjective, with the user often left to interpret solution quality and judge whether it is "good" or not. The situation is even more disconcerting in the case of an unskilled user without an extensive background in computational methods and fluid dynamics. It is easily conceivable to envision ten separate users analyzing a given configuration and arriving at ten different – possibly all incorrect – solutions. Portrayed in this light, CFD is indeed a very immature field.

While high-fidelity CFD has certainly found a niche in the analysis of new vehicle concepts, it is also becoming increasingly popular in the design stages for such problems. While many of these approaches are heuristic in nature, much success has been demonstrated on a wide range of problems. However, analogous to the shortcomings present in CFD-based analysis, the designer is usually required to have some a priori understanding of the optimal solution that is desired. For example, if the engineer believes a modification to the camber of a wing is needed to improve performance, he will most likely steer the process in that direction, either consciously or subconsciously. However, it is very possible that the optimal solution may lie

elsewhere in the design space, outside the often limited imagination of the user. In these cases, it could prove highly beneficial to have a formal optimization capability that would automatically and efficiently steer the design in such a direction.

The role of adjoint solutions for the governing equations as applied to both analysis and design problems will be covered in the current work. Adjoint solutions provide a very powerful approach to computing output error estimates, as well as systematically adapting grids to reduce spatial discretization error. Results show that the user is no longer required to “guess” at grid clustering strategies, essentially allowing the human to be removed from the analysis once an initial crude grid has been generated. Moreover, adjoint discretizations allow discretely consistent and extremely efficient sensitivity analyses to be performed for complicated problems governed by the Euler or Navier-Stokes equations. Derivatives of cost functions and constraints for hundreds of design variables can be computed at the cost of a single flowfield analysis, opening the door to efficient gradient-based optimization of such problems.

Ultimately, adjoint formulations are being targeted to enable rigorous, automated design optimization of aerospace configurations with known error bounds on the result, where no a priori knowledge of the flowfield nor optimal solution is known. While certainly an extremely high-risk goal, the potential payoff of such a capability could prove tremendous.

A brief overview of adjoint methods and their application to problems of interest at NASA Langley Research Center is given. An exhaustive list of relevant literature is beyond the scope of the current work; instead, the reader is referred to the resources cited in the publications provided at the end of the text.

The Flowfield Adjoint Equation

Consider the vector of discretized residual equations \mathbf{R} for the Euler or Navier-Stokes equations as a function of the design variables \mathbf{D} , computational mesh \mathbf{X} , and flowfield variables \mathbf{Q} . Given a steady-state solution of the form $\mathbf{R}(\mathbf{D}, \mathbf{Q}, \mathbf{X}) = 0$, a Lagrangian function L can be defined as

$$L(\mathbf{D}, \mathbf{Q}, \mathbf{X}, \Lambda_f, \Lambda_g) = f(\mathbf{D}, \mathbf{Q}, \mathbf{X}) + \Lambda_f^T \mathbf{R}(\mathbf{D}, \mathbf{Q}, \mathbf{X}) + \Lambda_g^T (\mathbf{K}\mathbf{X} - \mathbf{X}_{surf}) \quad (1)$$

where $f(\mathbf{D}, \mathbf{Q}, \mathbf{X})$ represents an objective function of interest and Λ_f and Λ_g represent vectors of Lagrange multipliers, or adjoint variables, corresponding to the flowfield and grid, respectively. The third term on the right hand side is the residual of the mesh movement problem, which in the current work is posed as a solution of the linear elasticity relations $\mathbf{K}\mathbf{X} = \mathbf{X}_{surf}$. Differentiating the above expression with respect to \mathbf{D} yields the following:

$$\frac{dL}{d\mathbf{D}} = \frac{\partial f}{\partial \mathbf{D}} + \left[\frac{\partial \mathbf{R}}{\partial \mathbf{D}} \right]^T \Lambda_f + \left[\frac{\partial \mathbf{Q}}{\partial \mathbf{D}} \right]^T \left\{ \frac{\partial f}{\partial \mathbf{Q}} + \left[\frac{\partial \mathbf{R}}{\partial \mathbf{Q}} \right]^T \Lambda_f \right\} + \left[\frac{\partial \mathbf{X}}{\partial \mathbf{D}} \right]^T \left\{ \frac{\partial f}{\partial \mathbf{X}} + \left[\frac{\partial \mathbf{R}}{\partial \mathbf{X}} \right]^T \Lambda_f + \Lambda_g^T \mathbf{K} \right\} - \Lambda_g^T \left[\frac{\partial \mathbf{X}}{\partial \mathbf{D}} \right]_{surf} \quad (2)$$

Since the vector of adjoint variables Λ_f is essentially arbitrary, the coefficient multiplying $\partial \mathbf{Q} / \partial \mathbf{D}$ can be eliminated using the following equation:

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{Q}} \right]^T \Lambda_f = - \frac{\partial f}{\partial \mathbf{Q}} \quad (3)$$

This expression represents the discrete adjoint equation for the flowfield.

Using Adjoint for Error Estimation and Mesh Adaptation

From inspection of Eq. 3, it is apparent that the adjoint variable Λ_f represents the effect of local equation error on the output of interest:

$$\Lambda_f \equiv \frac{\partial f}{\partial \mathbf{R}} \quad (4)$$

While the details of the underlying algebra are beyond the scope of this work (see for example Ref. 1), loosely speaking, if some measure of the spatial residual \mathbf{R} is known, an estimate of the output error can be obtained by taking an inner product with the adjoint variable Λ_f over the domain. Moreover, the solution for Λ_f can also be used to provide a direct relationship between local grid spacing requirements as they pertain to output accuracy. Such an explicit relationship has not been previously available, and the

approach often yields non-intuitive insight into gridding requirements. The methodology will be shown to have profound implications for mesh adaptation examples presented in a subsequent section.

The Mesh Adjoint Equation and Sensitivity Analysis

Once the solution Λ_f in Eq. 3 has been determined, the coefficient multiplying $\partial\mathbf{X}/\partial\mathbf{D}$ in Eq. 2 can be eliminated in a similar fashion by satisfying the following expression:

$$\mathbf{K}^T \Lambda_g = - \left\{ \frac{\partial f}{\partial \mathbf{X}} + \left[\frac{\partial \mathbf{R}}{\partial \mathbf{X}} \right]^T \Lambda_f \right\} \quad (5)$$

This relationship gives an expression for the adjoint variable Λ_g corresponding to the grid movement equations. Once known, the solutions for \mathbf{Q} , Λ_f , and Λ_g can be combined to yield the desired sensitivity vector, primarily a matrix-vector product given by the third term in the following expression:

$$\frac{dL}{d\mathbf{D}} = \frac{\partial f}{\partial \mathbf{D}} + \Lambda_f^T \frac{\partial \mathbf{R}}{\partial \mathbf{D}} - \Lambda_g^T \left[\frac{\partial \mathbf{X}}{\partial \mathbf{D}} \right]_{surf} \quad (6)$$

In summary, the analysis problem can be viewed as three phases: (1) Given \mathbf{D} , determine the corresponding surface grid; (2) Solve the interior grid movement equations based on the new surface grid; and (3) Solve the flowfield equations for the relevant outputs. The sensitivity analysis can be viewed as three similar steps, albeit in reverse: (3) Solve the flowfield adjoint equations; (2) Solve the mesh adjoint equations; and (1) Perform an explicit matrix-vector product over the surface to obtain the desired sensitivity vector. In this manner, a rigorous discrete sensitivity analysis can be performed at a cost equivalent to that of the analysis problem. The impact of this algorithmic efficiency on large-scale optimization problems will be demonstrated below.

Mesh Adaptation Examples

NACA 0012 Airfoils in Supersonic Inviscid Flow

The first problem is taken from Ref. 2. While this test case is largely academic in nature, it provides a simple example to demonstrate the tremendous potential for adjoint-based mesh adaptation. Consider a pair of staggered NACA 0012 airfoils in a Mach 3 freestream. The flowfield for this computation is shown in Fig. 1a. As expected, the flowfield contains a number of complicated shock patterns throughout the domain. The objective for the current mesh adaptation is to compute the drag on the lower airfoil as accurately as possible.

First, the problem is attempted using a traditional feature-based technique, where grid resolution is added to regions of the domain where strong gradients of pressure are found as shown in Fig. 1b. Naturally, this approach adds grid points along each of the shocks present in the flowfield and results in a final mesh consisting of 37,352 grid points and a drag on the lower airfoil of 767 counts.

The same test is also performed using the adjoint approach to guide the adaptation procedure, and the result is shown in Fig. 1c. In this case, the adjoint approach is implicitly aware that not all of the features present in the flowfield necessarily impact the output of interest. Instead, the procedure performs some refinement in the leading and trailing edge regions, as well as in between the two airfoils. Here, the final mesh consists of just 3,810 grid points – an order of magnitude reduction over the traditional approach – and yields essentially the same drag result for the lower airfoil, 766 counts. Placed within the context of three-dimensional design optimization where dozens of analyses may be required, this vastly reduced grid requirement could impact turnaround time by several orders of magnitude.

Turbulent Flow Over High-Lift Airfoil

This example is taken from Ref. 3 and demonstrates the ability of the adjoint approach to account for regions of the flow not necessarily indicated by feature-based approaches. In this test, an initial isotropic grid is constructed for a three-element airfoil configuration. The goal of the adaptation procedure is to predict the lift coefficient as accurately as possible.

The results of the feature-based approach are shown in Fig. 2. Due to the strong flowfield gradients in the immediate vicinity of the geometry, the adaptation procedure adds a number of points in these areas. Quantitative results of the present approach are shown in Fig. 3, where it can be seen that the feature-based

approach results in a lift coefficient of approximately 3, using approximately 52,235 grid points to do so. Also included on the plot are results from an independent uniform mesh refinement study performed using the same solver, as well as experimental data. These two sets of results are in good agreement with one another, yielding a lift coefficient of 3.5. The results from the uniform grid refinement study indicate that the physical modeling of the solver is capable of achieving the correct result, provided a suitable spatial discretization is used. However, the feature-based adaptation fails to recover such a mesh.

Adjoint-based adaptation is also applied to the current test problem. In this case, the adapted results rapidly approach the lift value obtained through uniform grid refinement, eventually reaching it after arriving at a final grid of 24,965 mesh points as also shown in Figs. 2 and 3. Several observations can be made for these results. Despite using half the number of grid points, the adjoint-based approach is able to recover the correct result by performing refinement in those areas which will directly impact the output of interest. Note that considerable refinement has taken place along the incoming stagnation streamline and wake regions, areas that are not identified by the feature-based approach. This highlights the ability of the adjoint method to target not only distinct flow features, but also more subtle, smooth regions of the flowfield which can also be of considerable importance. It is also interesting to observe the incorrect location of the wake in the cove region beneath the slat for the feature-based adaptation. Such behavior has also been observed in other tests of the feature-based technique: a feature may be identified early in the adaptation procedure which may be in an incorrect location. The feature-based method may subsequently refine this region, giving a final solution that may appear more crisp or well-defined; however, the final result may actually be grossly inaccurate.

Supersonic Inviscid Flow Over Double Cone-Cylinder Configuration

The next test case is axisymmetric flow over a double cone-cylinder configuration as shown in Fig. 4 at Mach 1.26. This case is repeated from Ref. 4 and serves as a model problem for validating the adjoint-based adaptation procedure for nearfield sonic boom applications. The objective in this case is to accurately predict the pressure signature at a distance of 6 body lengths away from the geometry. For this test, the objective function is the integral of static pressure over a user-defined surface at this location.

To simplify the computation, the three-dimensional nature of the problem is discretized as a 9-degree wedge of the domain in the axisymmetric sense. The initial grid is shown in Fig. 4 and has been generated with no attention paid to the expected final results. The corresponding pressure signal on the symmetry plane obtained using this initial grid is also included in Fig. 4, where it can be seen that there is clearly insufficient grid density to resolve the shocks at the integration surface.

The resulting mesh and pressure signal after several cycles of adjoint-based adaptation are shown in Fig. 5. The method has successfully propagated the shock structures to the near-field integration surface at 6 body lengths, with the corresponding pressure signal agreeing well with the experimental data at this location.

Supersonic Inviscid Flow Over Wing-Body Configuration

This test case is similar to the previous example and has also been taken from Ref. 4. In this demonstration, the nearfield pressure signal is sought at an integration surface located 2.5 body lengths below the wing-body configuration shown in Fig. 6. For this computation the integration surface is taken as a narrow strip of the domain located immediately adjacent to the symmetry plane, since this location corresponds to available experimental data. It should be noted that the original blunt trailing edge geometry used for this test case has been slightly modified to obtain a sharp trailing edge. This minor planform modification can be seen as the red portion of the surface in Fig. 6 and has been done to eliminate the strong expansion in this area, which is widely known to cause numerical problems for inviscid computations.

The initial grid is shown in Fig. 6 and has been constructed with no consideration for the final flowfield structure. The initial pressure signal is also included in Fig. 6 and clearly illustrates the lack of grid resolution required to achieve agreement with the experimental results. Following several cycles of adjoint-based adaptation, the grid shown in Fig. 7 is obtained. The corresponding pressure signal is also shown, and it can be seen that the results are in relatively good agreement with the experimental data, with the exception of the large expansion towards the aft end of the signal which is not predicted. This test case is a preliminary result for which ongoing work is focused; however, the current results demonstrate great potential for the method and are very encouraging. It is also interesting to note the lack of grid refinement in the region above the vehicle, where the adjoint approach is inherently aware that this region will have little or no impact on the objective function of interest.

Turbulent Flow Over Transonic Wing-Body Configuration

The final example of adjoint-based mesh adaptation is taken from Ref. 4 and stems from participation in the recent AIAA Drag Prediction Workshop series. In the most recent of these workshops, participants were required to compute the transonic turbulent flowfield over a transport wing-body configuration using a series of three successively refined grids. Results of this study are shown in Fig. 8a, where the drag coefficient has been plotted as a function of the mesh density. Note that an extrapolation of these results predicts an infinitely refined grid result of approximately 278 counts for the drag coefficient, while the experimental value is roughly 295 counts. This large discrepancy on a seemingly well-resolved grid of 9 million unknowns is disappointing.

For the current test, the medium grid consisting of 3 million grid points is taken as the starting point for adjoint-based adaptation, with the objective of predicting the drag coefficient as accurately as possible. Due to the lack of robust three-dimensional anisotropic grid adaptation mechanics for the boundary layer region, this region of the domain is frozen throughout the computation.

As is also shown in Fig. 8a, the results from the adaptation procedure deviate from the grid refinement study and immediately recover larger values for the drag coefficient. Since the procedure is limited to mesh adaptation outside the boundary layer, the process eventually stalls as resolution is not being added as requested by the adjoint solution. Nevertheless, the final value of drag is roughly 292 counts. A typical plot of the adaptation parameter throughout the field is shown in Fig. 9a. In this figure, blue regions represent areas where grid resolution is sufficient, while red regions indicate refinement is needed. Similar to results shown earlier, the stagnation streamlines and wakes appear to be grossly underresolved, areas that are not typically clustered in common unstructured grid generation tools. A view of the adapted symmetry plane grid and a slice through the wing are shown in Fig. 9b and reveal that significant refinement has taken place in these regions.

Finally, it should be noted that an additional fine grid consisting of 65 million grid points has recently been generated. The results from this computation are included in Fig. 8b and show that with increasing resolution, the uniform grid refinement study does indeed turn toward the experimental data, albeit much slower than the adjoint-based adaptation procedure, which immediately targets the deficient regions of the grid while not overresolving others.

Design Optimization Examples

Inviscid Lift-Constrained Drag Minimization

This test case is based on an unmanned air vehicle and is taken from Ref. 5. The aircraft is designed to cruise in a conventional configuration and to fold its wings into a stowed position for the dash portion of its mission. This test case focuses on the stowed configuration and is computed using the Euler equations. Similar computations have been performed for viscous flow, however results are not included here. The grid used for this case is shown in Fig. 10a. The freestream Mach number is 0.80 and the angle of attack is 2 degrees. For this test, the outer mold line has been parameterized using 175 design variables to describe the wing and fuselage.

To demonstrate an optimization for the current test case, an objective function based on drag has been used in conjunction with an explicit lift constraint. A subset of 63 of the 175 total shape variables are free to change, subject to side constraints that prohibit thinning of the geometry. The results of the optimization are shown in Fig. 10b, where the drag coefficient has been reduced by 52% and the lift coefficient satisfies the requested constraint value. For this test case, the optimizer requested 26 sensitivity analyses during the course of the computation, which took roughly 24 hours of wallclock time on 16 processors. Based on the cost of a single analysis for this configuration, it is estimated that the same computation would have taken approximately 7 weeks of wallclock time to perform without an adjoint formulation for the sensitivity analysis.

Adjoint for Hypersonic Flows: Extension to General Equation Sets

There has been a significant effort recently within the CFD community to provide accurate and robust hypersonic aerodynamic and aerothermodynamic capabilities within unstructured grid frameworks, and progress to date has resulted in significantly more elaborate flow solution algorithms. In addition to the nonlinear limiter functions necessary for supersonic flows, high-energy flow solvers typically contain curve

fits for transport properties, eigenvalue limiters, a variable number of species and energy equations, and may employ embedded Newton iterations to determine thermodynamic properties or to implement boundary conditions.

Consider the hand-coded discrete adjoint implementation demonstrated above for the perfect gas Reynolds-averaged Navier-Stokes equations. This capability has taken over 5 years to evolve into a mature capability for large-scale problems. Any extension to its basic functionality remains an extremely sobering undertaking, and manually extending it to include the additional complexities required for thermochemical nonequilibrium flows is simply untenable. This issue has instead served as the primary motivation for the work performed in Ref. 6, wherein an automatable, more efficient, and less error-prone procedure for developing a discrete adjoint solver for increasingly complex sets of governing equations has been sought. This alternative approach is based on the use of complex variables and is described briefly below.

Using Complex Variables to Form Discrete Adjoint Operators

A Taylor series with a complex step size ih can be used to derive an expression for the first derivative of a real-valued function $f(x)$:

$$f'(x) = \frac{\text{Im}[f(x+ih)]}{h} + O(h^2) \quad (7)$$

Several observations can be made about Eq. 7. As with real-valued central differencing, the expression is second-order accurate; however, there is no subtraction of neighboring terms involved. This analytical extension allows true second-order accuracy to be realized, where two additional digits of accuracy are obtained for each order of magnitude reduction in the step size h . Moreover, implementation of the method is straightforward: declare all floating point variables complex and apply a complex perturbation to the design variable of interest. Execute the simulation, and upon completion, the imaginary part of the output is the partial derivative with respect to the perturbed variable multiplied by the step size h . The drawbacks to this technique are the need to recompute f for each perturbation and the additional cost of performing complex arithmetic.

Since the complex variable approach is a direct, or forward, mode of differentiation, the method must be applied in a careful manner in order to be computationally efficient when used to form adjoint operators. Such an approach is outlined in Ref. 6. By taking advantage of properties such as known stencil layouts across the mesh, numerous elements of a flux Jacobian, for example, can be obtained within the context of a single complex-valued flux evaluation. The reader is encouraged to refer to Ref. 6 for a detailed description of the complete procedure.

Hypersonic Test Case – Flow Over a Circular Cylinder

A 5-species air laminar flow cylinder test case is included from Ref. 6. Shown in Fig. 11a, the grid used for this case has been derived from a structured grid and contains a single layer of cells in the spanwise direction. Note that the structured grid cells have been diagonalized in a uniform manner so that a severe spanwise bias is present in the computation. This grid topology is being heavily relied upon in related work for “stress-testing” the accuracy of various discretization schemes on tetrahedra.

For this test, the freestream velocity is 5,000 m/s, the freestream density is 0.001 kg/m³, and the temperature is 200 K. These conditions give a Reynolds number of approximately 425,000 based on the cylinder diameter and a freestream Mach number of 17.6. The flowfield is governed by nine conservation equations: five species equations (N₂, O₂, N, O, and NO) and the usual momentum and energy equations. The computation has been performed on eight processors and contours of the Mach number are shown in Fig. 11b, where a strong bow shock can be seen upstream of the cylinder. The temperature downstream of the shock in the leading edge region exceeds 6,300 K, causing the freestream molecules to dissociate; contours of atomic nitrogen and oxygen are shown in Figs. 11c and 11d. Using the complex-variable approach described above, adjoint solutions based on drag and surface heating functions have been computed. Results for these adjoint computations are shown in Figs. 11e and 11f.

References

1. Park, M.A., “Three-Dimensional Turbulent RANS Adjoint-Based Error Correction,” AIAA 2003-3849, 2003.

2. Venditti, D.A., "Grid Adaptation for Functional Outputs of Compressible Flow Simulations," Ph.D. Thesis, Massachusetts Institute of Technology, 2002.
3. Venditti, D.A. and Darmofal, D.L., "Anisotropic Grid Adaptation for Functional Outputs: Application to Two-Dimensional Viscous Flows," *Journal of Computational Physics*, Vol. 187, 2003, pp. 22-46.
4. Lee-Rausch, E.M., Park, M.A., Jones, W.T., Hammond, D.P., and Nielsen, E.J., "Application of Parallel Adjoint-Based Error Estimation and Anisotropic Grid Adaptation for Three-Dimensional Aerospace Configurations," AIAA 2005-4842, 2005.
5. Nielsen, E.J. and Park, M.A., "Using an Adjoint Approach to Eliminate Mesh Sensitivities in Computational Design," AIAA 2005-0491, 2005.
6. Nielsen, E.J. and Kleb, W.L., "Efficient Construction of Discrete Adjoint Operators on Unstructured Grids by Using Complex Variables," AIAA 2005-0324, 2005.

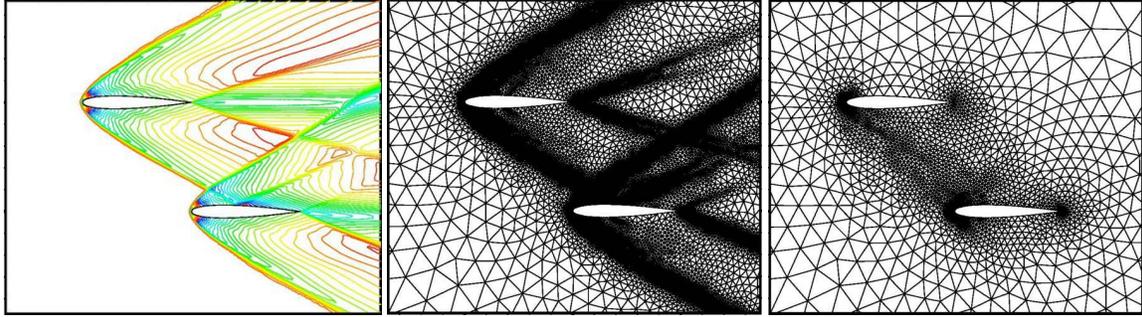


Figure 1. a) Staggered NACA 0012 airfoils in Mach 3 freestream, b) Results of feature-based adaptation, c) Results of adjoint-based adaptation.

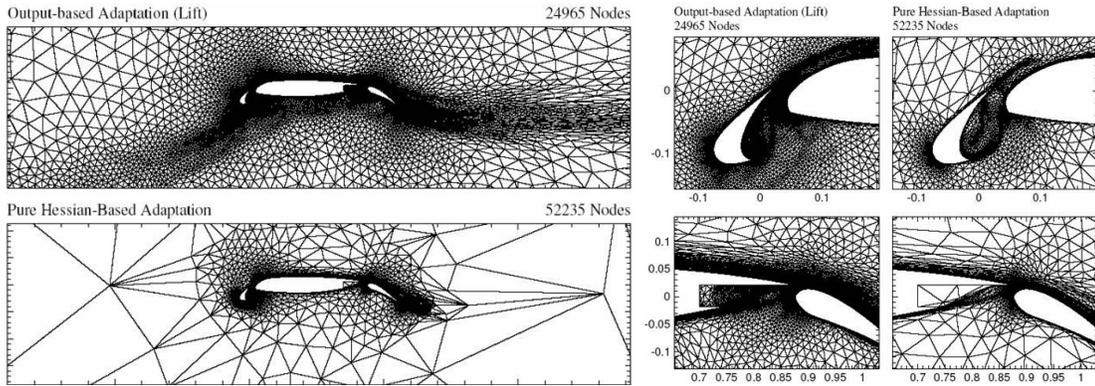


Figure 2. Grids obtained through adaptation for high-lift example.

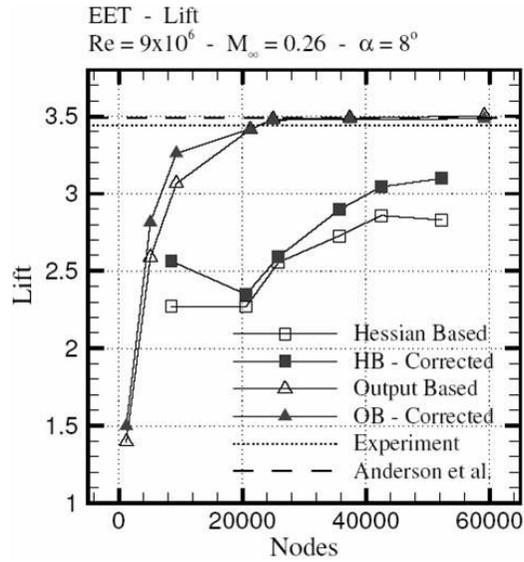


Figure 3. Lift coefficient versus grid density for each approach.

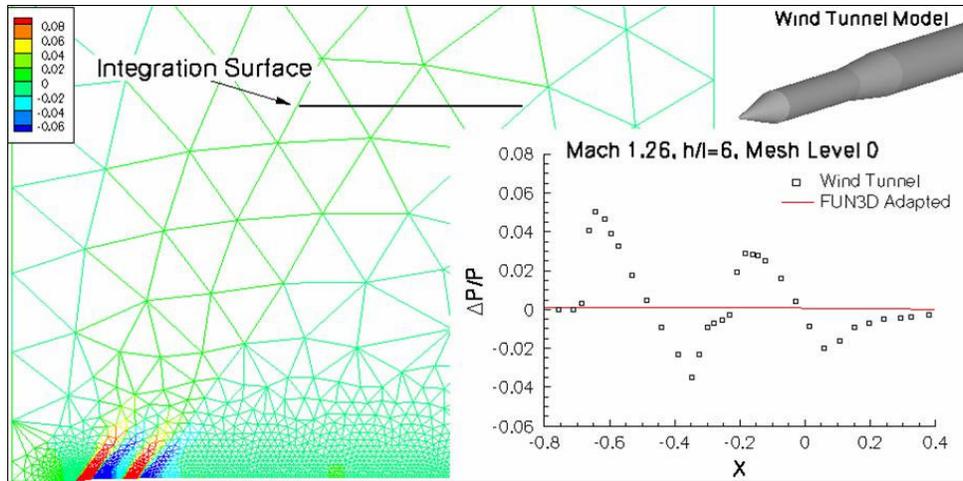


Figure 4. Initial double cone-cylinder grid and solution.

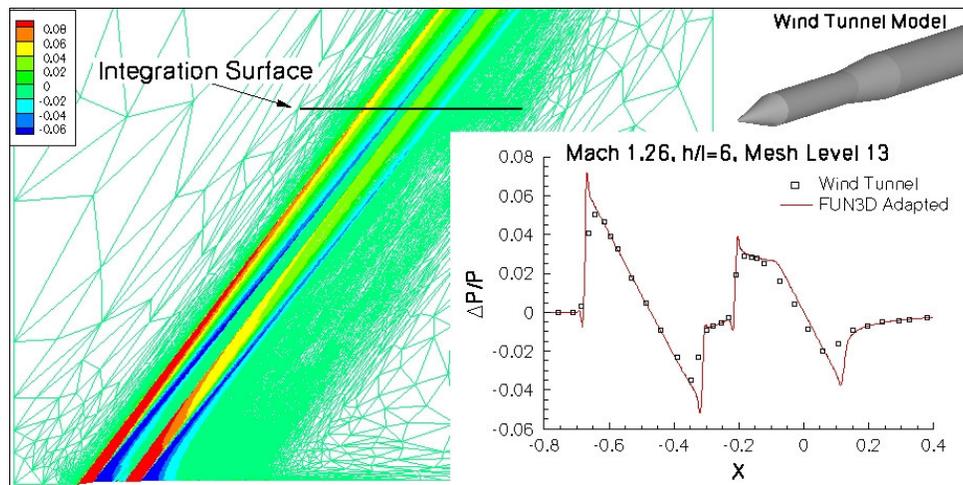


Figure 5. Final grid for double cone-cylinder and pressure signal at integration surface.

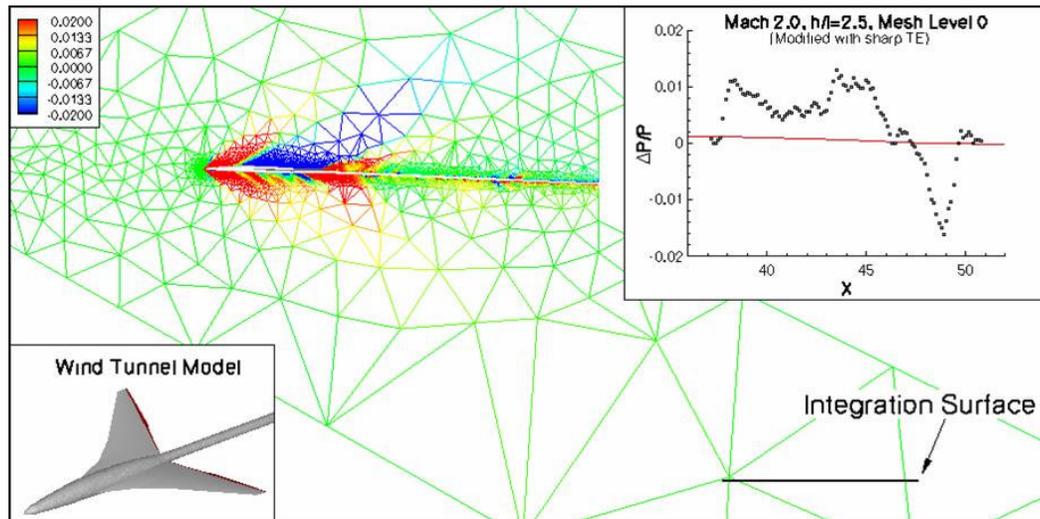


Figure 6. Initial wing-body grid and pressure signal at integration surface.

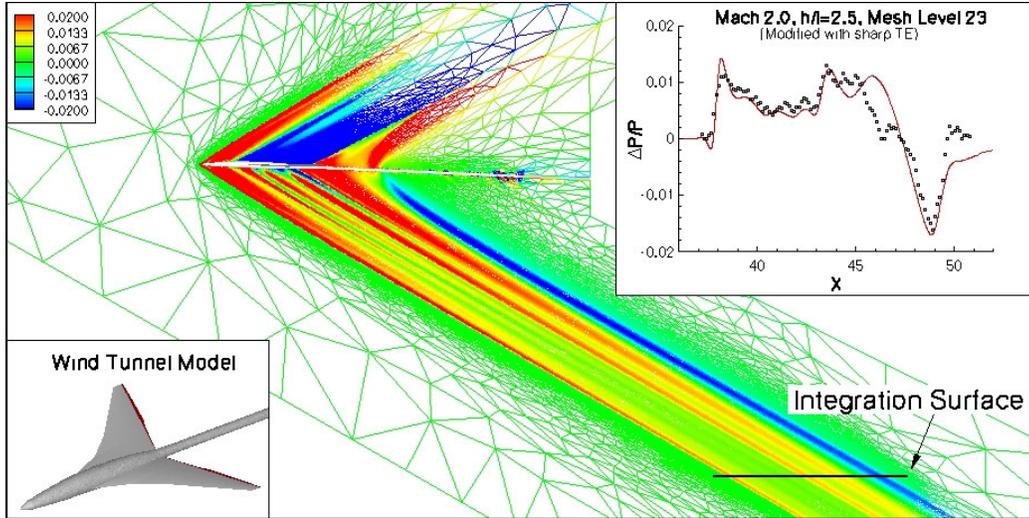


Figure 7. Final wing-body grid and pressure signal at integration surface.

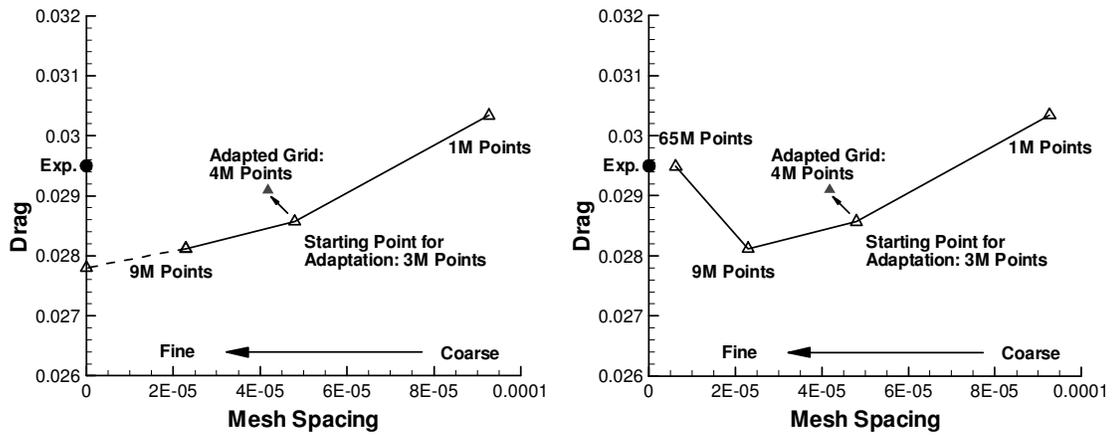


Figure 8. a) Uniform grid refinement and adaptation results for transonic wing-body configuration, b) Results including 65 million-point grid computation.

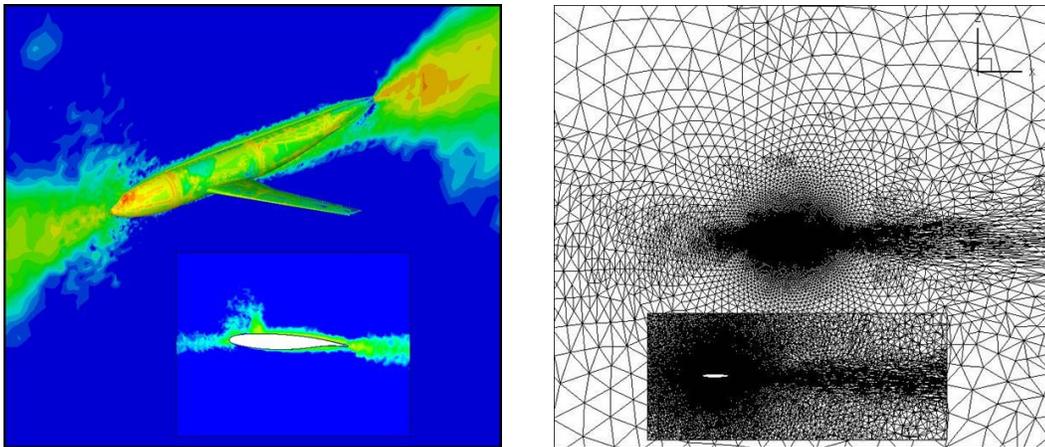


Figure 9. a) Adjoint-based adaptation parameter, b) Adapted grid for transonic wing-body.

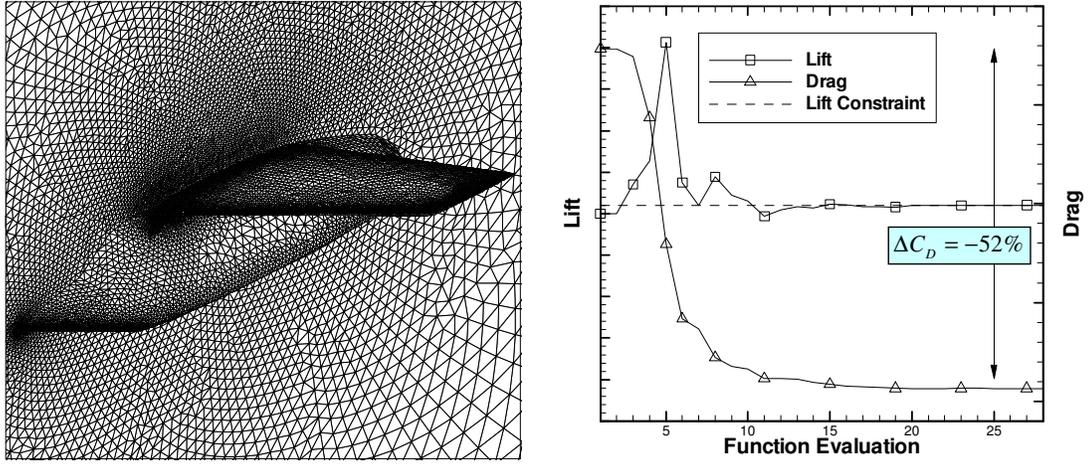


Figure 10. a) Grid for morphing configuration, b) Optimization results.

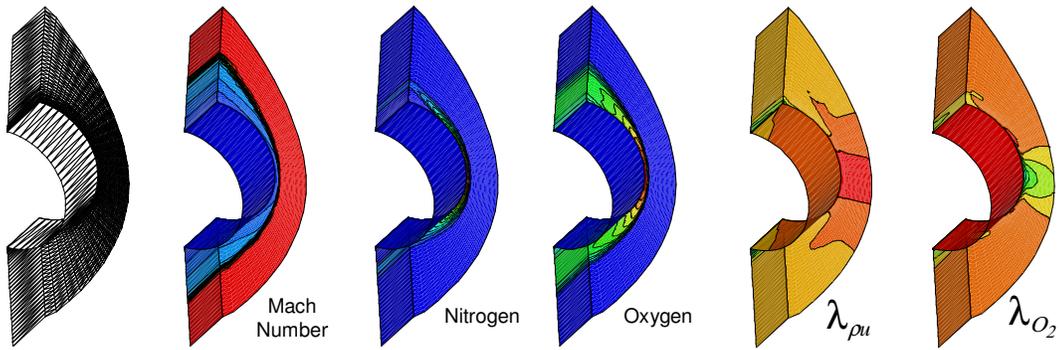


Figure 11. a) Grid for hypersonic cylinder computation, b) Mach number distribution, c) Atomic nitrogen distribution, d) Atomic oxygen distribution, e) Streamwise momentum adjoint solution for drag, f) Molecular oxygen concentration adjoint solution for heating