Application of Parallel Adjoint-Based Error Estimation and Anisotropic Grid Adaptation for Three-Dimensional Aerospace Configurations

E. M. Lee-Rausch [*]   M. A. Park [†]   W. T. Jones [‡]   D. P. Hammond [§]   E. J. Nielsen [¶]

I.   Abstract

This paper demonstrates the extension of error estimation and adaptation methods to parallel computations enabling larger, more realistic aerospace applications and the quantification of discretization errors for complex 3-D solutions. Results were shown for an inviscid sonic-boom prediction about a double-cone configuration and a wing/body segmented leading edge (SLE) configuration where the output function of the adjoint was pressure integrated over a part of the cylinder in the near field. After multiple cycles of error estimation and surface/field adaptation, a significant improvement in the inviscid solution for the sonic boom signature of the double cone was observed. Although the double-cone adaptation was initiated from a very coarse mesh, the near-field pressure signature from the final-adapted mesh compared very well with the wind-tunnel data which illustrates that the adjoint-based error estimation and adaptation process requires no a priori refinement of the mesh. Similarly, the near-field pressure signature for the SLE wing/body sonic boom configuration showed a significant improvement from the initial coarse mesh to the final adapted mesh in comparison with the wind tunnel results. Error estimation and field adaptation results were also presented for the viscous transonic drag prediction of the DLR-F6 wing/body configuration, and results were compared to a series of globally refined meshes. Two of these globally refined meshes were used as a starting point for the error estimation and field-adaptation process where the output function for the adjoint was the total drag. The field-adapted results showed an improvement in the prediction of the drag in comparison with the finest globally refined mesh and a reduction in the estimate of the remaining drag error. The adjoint-based adaptation parameter showed a need for increased resolution in the surface of the wing/body as well as a need for wake resolution downstream of the fuselage and wing trailing edge in order to achieve the requested drag tolerance. Although further adaptation was required to meet the requested tolerance, no further cycles were computed in order to avoid large discrepancies between the surface mesh spacing and the refined field spacing.

II.   Introduction

As advanced computational fluid dynamics (CFD) codes have been developed to capture increasingly more complex flow features, assuring that the computational meshes are adequate to resolve these flow features has become a challenge. In aerospace applications, this can be especially difficult due to the wide range of speeds and conditions that must be simulated. A primary goal of CFD simulation is to accurately predict and understand wind tunnel and full-scale tests. Validation efforts have been applied to two-dimensional (2-D) and three-dimensional (3-D) problems to quantify current capabilities and expose deficiencies in current CFD simulation and wind tunnel testing practices. Both 2-D and 3-D validation efforts have been hampered by a number of factors including adequate mesh resolution, unquantified modeling errors, and geometric uncertainty. In 2-D calculations, the mesh can more easily be globally

---

[*]Senior Member AIAA, Research Engineer NASA Langley Research Center(LaRC), Hampton, Virginia.
[†]Member AIAA, Research Scientist NASA LaRC.
[‡]Senior Member AIAA, Computer Engineer NASA LaRC.
[§]Senior Member AIAA, Computer Engineer NASA LaRC.
[¶]Senior Member AIAA, Research Scientist NASA LaRC.

refined to capture the required flow features and prove mesh convergence. However, providing validation through comparison with an experiment can be problematic even for 2-D test cases especially for situations with large separated flow regions. Under these conditions, the experimental flow field is simply no longer dominated by 2-D structures.[1]

Even with significant advances in computer memory and speed, it is difficult to generate a series of globally refined meshes that rigorously prove mesh convergence for a realistic 3-D configuration.[2,3] The resolution of the meshes for the vast majority of these cases is specified manually. In many cases, it takes years of experience to understand the subtleties of complex flow field mesh resolution requirements. New code users must be trained, and experienced users must train themselves to analyze new configurations. Through the manual specification of mesh resolution, the experience of the user is a significant factor in the uncertainty associated with CFD calculations. Nevertheless, the aerospace industry has been successful in many applications using this approach to mesh placement. However, when the CFD results do not agree with an experiment or other CFD methods, the source of the discrepancies can be difficult to understand.[4,5] The lack of demonstrated mesh convergence for 3-D cases makes it difficult to discern how discretization errors, modeling errors, or both are contributing to the discrepancies between computation and experiment.

One approach to the problem of mesh placement is to use mesh adaptation to increase mesh resolution where local discretization error estimates are large.[6–15] This type of solution adaptive process is often referred to as feature-based adaptation because it focuses on the strong features that are present in the flow solution (e.g., shocks and boundary layers). However, the manual specification and feature-based adaptation methods miss the connection between the impact of local errors on global output quantities (i.e., lift and drag) because these local errors can be transported through the solution to degrade the accuracy of the calculation. Also, they lack a rigorous termination criteria. There are a number of documented feature-based adaptation results which have converged to incorrect answers.[7,16,17] Additionally, uniformly reducing the errors associated with all features of the flow may not be optimal from an engineering context in which a few output functionals (i.e. lift and drag) may be of primary importance. An alternative method to feature-based adaptation is an adjoint method that constructs an error estimate to improve the calculation of a specified engineering output functional. This output-based approach has been applied to finite element discretizations and includes applications to error bounds on outputs and adaptation.[18–25]

Pierce and Giles[26] extended the error correction methodology to discretizations other than finite elements. Müller and Giles[27] applied an adaptive technique to reduce the error correction term for finite volume solutions. Venditti and Darmofal[16,17] demonstrated an adaptive technique for compressible 2-D inviscid and viscous flow solutions, which improves the error correction of an output function. Their method for anisotropic adaptation of 2-D turbulent flows was demonstrated on a multielement airfoil. Reference 16 demonstrates significant improvement in the reliability and output accuracy of the adjoint-based method in comparison with pure Hessian-based (feature-based) adaptation. This output-based approach provides a natural termination criteria that is based on a user-specified functional error tolerance. An added benefit of an output-based adaptive process is that it holds the promise of lowering quality requirements on the initial mesh generation process because the output-based adaptation is able to find the correct solution given coarse initial meshes.[16,17,28,29] These initial meshes may be unable to resolve important features that are present in the final, adapted mesh. However, after a number of iterations, the adjoint-adapted meshes have implicitly positioned and resolved these features. This lower resolution and quality requirement directly increases the robustness and reduces the work-hours associated with the initial mesh generation process. The initial mesh resolution specification may also be formulated to automatically satisfy these reduced resolution requirements. Output based adaptation has the potential to reduce the expense and increase the robustness and accuracy of CFD solutions for 3-D problems. However, the mesh mechanics to realize this process improvement for 3-D turbulent applications are not complete.

Park[28,29] extended the error-estimation methods in Ref. 16 to 3-D inviscid, laminar, and turbulent flows. Reference 29 demonstrated automated output (adjoint) anisotropic surface and field adaptation on an inviscid ONERA M6 wing and manual adaptation of a turbulent, extruded NACA 0012 airfoil. The error estimation procedure in Ref. 29 required the sequential computation of the flow and adjoint residual on an h-refined embedded mesh. This sequential execution process limited the size of the problems that could be analyzed on distributed computer clusters. The current paper demonstrates the extension of the error estimation and adaptation methods in Refs. 28 and 29 to parallel computations, which has allowed for larger, more realistic aerospace applications and the quantification of discretization errors for complex 3-D solutions.

The current paper focuses on two challenging applications of current interest to the aerospace community: sonic boom prediction and transonic drag prediction. Sonic boom signature prediction is particularly challenging for many traditional CFD methods because of the requirement that the near-field pressure signature be accurately predicted for propagation to the ground. Solution adaptive methods have been shown to accurately and efficiently predict the near-field pressure disturbance.[15,30–33] The current adjoint-based error estimation and adaption procedure is applied

to to the 3-D Mach 1.26 inviscid flow over two tandem cones connected by a cylinder and a wind tunnel sting. This is a case taken from a 1965 NASA Langley Research Center experiment and analysis reported in Ref. 34. The objective is to perform (output) adjoint-based adaptation on this configuration where the output function of the adjoint is pressure integrated over a half cylinder with a radius of six model lengths. The adjoint-based error-estimation is used to automatically adapt a coarse initial mesh through multiple cycles of adaptation. The near-field pressure signal is extracted and compared to the wind-tunnel data of Ref. 34. Following validation of the method for near-field sonic boom pressure signatures, the technique was applied to Mach 2.0 flow over a wing/body configuration. The segmented leading edge (SLE) model also tested at NASA Langley Research Center.[35] was chosen for the application. The output function for the SLE is defined in a manner consistent with the previous case, but at a distance of 2.5 body lengths where test data was available. The adjoint-based error-estimation is used to automatically adapt a coarse initial mesh through multiple cycles of adaptation. The near-field pressure signal is extracted and compared to the wind-tunnel data of Ref. 35.

The final application of the adjoint-based error estimation and adaptation procedure is the drag prediction on a wing/body transport configuration, the DLR-F6. The DLR-F6 transonic cruise condition was the focus of the AIAA Second Drag Prediction Workshop (DPW-II).[2,36] A notable outcome of DPW-II was that the statistical uncertainty of predicting transonic cruise drag significantly decreased over that of the First AIAA Drag Prediction Workshop (DPW-I). The estimated code-to-code population standard deviations of total drag for the nested solutions was ± 7.3 counts for the wing/body, ± 11.4 counts for the wing/body/nacelle/pylon and ± 8 counts for the nacelle/pylon increment.[3] The mesh resolution studies conducted as a part of the DPW-II were also useful but still led to a consensus that many of the provided meshes were not adequate. As a follow-up to the DPW-II workshop, many of the participants reported individual code results at a special session of the AIAA 42nd Aerospace Sciences Meeting and Exhibit.[37–41] In Ref. 41, an in-depth comparison was made between three distinctly different unstructured Navier-Stokes flow solvers exercised on a similar family of tetrahedral meshes created for the workshop. The expectation was to see less code-to-code variation in the forces and moments than was achieved given the considerable effort to create comparable meshes.[41] Although the variation in the constant-lift wing/body total drag was less than observed for the workshop, the increasing variation in angle-of-attack and pitching moment with mesh refinement was surprising. In the current work, the adjoint-based error estimation and field adaptation procedure is used to estimate the remaining error in the wing/body total drag for the FUN3D solutions shown in Ref. 41. Comparisons are made between results from the field adapted meshes and the globally refined meshes. The adjoint-based adaptation parameter is used to identify the areas of the mesh that are not sufficiently resolved for accurate drag prediction.

III. Flow Solver

The fully unstructured Navier-Stokes three-dimensional[42–44](FUN3D) suite of codes is employed in this study. The compressible flow solver employs an unstructured finite-volume tetrahedral method for conserved variables, $\mathbf{Q}$, i.e.,

$$\mathbf{Q} = [\rho \; \rho u \; \rho v \; \rho w \; E]^T \tag{1}$$

where ρ is density, u, v, and w are velocity components, and E is total energy per unit volume. The node-based variables, $\mathbf{Q}$, are computed by driving the flow equation residual $\mathbf{R}$ to steady-state with an implicit point-iterative method or an implicit line relaxation scheme.[45] FUN3D is able to solve incompressible, Euler, and Reynolds-averaged Navier-Stokes (RANS) flow equations, either tightly or loosely coupled to the Spalart-Allmaras[46] (S-A) one-equation turbulence model. When the S-A model is included, the turbulence model quantity $\tilde{\nu}$ is included in $\mathbf{Q}$. The present study employs the Euler and RANS equations coupled to the S-A turbulence model. The solution of $\mathbf{Q}$ allows the calculation of integral quantities f (e.g., lift and drag).

IV. Adjoint Solver

After the flow solution is known, the discrete adjoint equations[43,44] are solved to complete the dual problem. The first step is to linearize the flow equation residual (including the turbulence model) $\mathbf{R}$ and output function f with respect to the flow solution $\mathbf{Q}$. After this linearization, an adjoint variable λ is solved for each of the flow equations and the turbulence model. An abbreviated derivation, adapted from Taylor et al.,[47] is given below. The chain rule for the linearized output function is

$$\frac{\partial f}{\partial \mathbf{Q}} = \left(\frac{\partial f}{\partial \mathbf{R}} \right)^T \frac{\partial \mathbf{R}}{\partial \mathbf{Q}} \tag{2}$$

American Institute of Aeronautics and Astronautics Paper 2005-4842

The adjoint variable λ is defined as the linearized effect of a flow residual source term on the output function:

$$\frac{\partial f}{\partial \mathbf{R}} = \lambda \tag{3}$$

This set of linear equations is solved to find λ:

$$\left(\frac{\partial \mathbf{R}}{\partial \mathbf{Q}}\right)^T \lambda = \left(\frac{\partial f}{\partial \mathbf{Q}}\right)^T \tag{4}$$

After the flow solution is known, this set of linear equations is solved with a point-iterative time-marching method or an implicit line relaxation scheme.[45,48] To speed execution, the global problem domain is decomposed into multiple sub-domains, and the flow and the adjoint problems are solved with a parallel execution scheme, which communicates via the message passing interface (MPI) standard.

V. Error Correction

The error prediction and correction scheme is taken from Refs. 16 and 17. With a solution on a mesh of reasonable size $\mathbf{Q}^0$, it is desirable to predict the value of an output function evaluated with a solution on a much finer mesh $f(\mathbf{Q}^*)$. This prediction can be computed without the solution on this finer mesh when the adjoint solution on this reasonable mesh λ^0 is used. The full derivation of the error correction term is available in Refs. 16, 17, 26, and 27. An abbreviated derivation is presented by expanding the Taylor series about $f(\mathbf{Q}^0)$, i.e.,

$$f(\mathbf{Q}^*) = f(\mathbf{Q}^0) + \left(\left.\frac{\partial f}{\partial \mathbf{R}}\right|_0\right)^T \left(\mathbf{R}(\mathbf{Q}^*) - \mathbf{R}(\mathbf{Q}^0)\right) + \ldots \tag{5}$$

Employing Eq. (3) and assuming that the residual on the much finer mesh is zero yields an improved estimate for the functional of interest f_{est}:

$$\left.\frac{\partial f}{\partial \mathbf{R}}\right|_0 = \lambda^0 \tag{6}$$

$$\mathbf{R}(\mathbf{Q}^*) = 0 \tag{7}$$

$$f(\mathbf{Q}^*) \approx f_{est} = f(\mathbf{Q}^0) - (\lambda^0)^T \mathbf{R}(\mathbf{Q}^0) \tag{8}$$

To improve the prediction of the functional f_{est}, $\mathbf{Q}^0$ and λ^0 can be interpolated to an embedded mesh. Interpolation is performed in two ways for this study: a linear interpolation ($\mathbf{Q}^L$,λ^L) and a higher order interpolation ($\mathbf{Q}^H$,λ^H). Details of this interpolation and the construction of this embedded mesh are in Ref. 29, including special considerations required for interpolating λ for problems with strong boundary conditions.[16] Substituting the higher order interpolant into Eq. (8) yields the higher order functional estimate f_{est}^H:

$$f_{est}^H = f(\mathbf{Q}^H) - (\lambda^H)^T \mathbf{R}(\mathbf{Q}^H) \tag{9}$$

VI. Adaptation Parameter

The adaptation parameter, also from Ref. 16 and 17, is intended to specify a mesh spacing modification to reduce the uncertainty in the computed error prediction. The mesh is not modified to directly reduce the computed error prediction (as in Ref. 27) because an accurate estimate for the functional, including this error term, can be computed with Eq. (8). Instead, targeting the uncertainty in this computed quantity is more effective and improves the robustness of the adaptive process. The error correction (Eq. (8)), including the uncertainty in the dual solution, is

$$f(\mathbf{Q}^0) - f(\mathbf{Q}^*) \approx (\lambda^0)^T \mathbf{R}(\mathbf{Q}^0) + (\lambda^* - \lambda^0)^T \mathbf{R}(\mathbf{Q}^0) \tag{10}$$

The uncertainty in the computed error correction is

$$f_{est} - f(\mathbf{Q}^*) \approx (\lambda^* - \lambda^0)^T \mathbf{R}(\mathbf{Q}^0) \tag{11}$$

The relation of the primal and dual problems[16,17,26] yields another expression for the error correction uncertainty

$$(\lambda^* - \lambda^0)^T \mathbf{R}(\mathbf{Q}^0) = \mathbf{R}_\lambda(\lambda^0)(\mathbf{Q}^* - \mathbf{Q}^0)^T \tag{12}$$

where $\mathbf{R}_\lambda(\lambda)$ is the residual of the dual problem:

$$\mathbf{R}_\lambda(\lambda) = \left(\frac{\partial \mathbf{R}}{\partial \mathbf{Q}}\right)^T \lambda - \left(\frac{\partial f}{\partial \mathbf{Q}}\right)^T \tag{13}$$

A computable term is found by using the interpolation error of λ to replace $(\lambda^* - \lambda^0)$, and the interpolation error of $\mathbf{Q}$ to replace $(\mathbf{Q}^* - \mathbf{Q}^0)$. The higher order interpolant for $\mathbf{Q}^0$ and λ^0 is employed for the residual calculations to improve prediction in place of the linear interpolant in Ref. 16 and 17. The linear interpolant yields a more conservative estimate of the remaining error; the high-order interpolant generally results in a more accurate estimate of the remaining error on the final adapted mesh. The high-order interpolant estimate may underestimate the remaining error on very coarse initial meshes, but the estimation improves as the adaptation process proceeds. The interpolation error is expressed as the difference in the high-order and linear interpolated values:

$$(\lambda^* - \lambda^0) \approx (\lambda^H - \lambda^L) \tag{14}$$

$$(\mathbf{Q}^* - \mathbf{Q}^0) \approx (\mathbf{Q}^H - \mathbf{Q}^L) \tag{15}$$

The average of the absolute values of the two uncertainty terms in Eq. (12) yields the adaptation intensity $\mathbf{I}$, which is computed for each equation on each embedded node:

$$\mathbf{I} = \frac{\left|(\lambda^H - \lambda^L)^T \mathbf{R}(\mathbf{Q}^H)\right| + \left|\mathbf{R}_\lambda(\lambda^H)(\mathbf{Q}^H - \mathbf{Q}^L)^T\right|}{2} \tag{16}$$

The intensity $\mathbf{I}$ is therefore the average of the absolute values of two terms. The first term is a dual solution interpolation error weighted with the primal residual. The second term is the dual problem residual weighted with a primal solution interpolation error. This form of the adaptation intensity (which includes weighed interpolation errors) focuses on the nonlinear contributions to the function error, which increases robustness of the adaptation method. The global sum of the adaptation parameter (Eq. 16) can be used also as an estimate of the error remaining after the error correction term is applied. Reference 29 details how $\mathbf{I}$ is combined with a user-specified output error tolerance and Mach Hessian to compute an anisotropic metric specification.

VII. Error Correction and Adaptation Process

The error correction process begins with an initial tetrahedral mesh, which can come from any mesh generation system. This initial mesh is partitioned to allow parallel execution. The state variables are computed in parallel as the nonlinear solution to the flow equations on the initial mesh, and then the adjoint variables are computed also in parallel with the linearized flow equations at the converged flow solution. To compute the error prediction and the adaptation parameter, a globally embedded, or h-refined mesh, is created with interpolated primal and dual solutions. In the previous work of Ref. 29, the global problem domain was reconstructed to facilitate the creation of a finer, embedded mesh with interpolated primal and dual solutions. However, in the current work, the computation of the error prediction and the adaptation parameter has been implemented with a parallel execution scheme. Once the adaptation parameter for a given mesh has been computed, this parameter can be used to specify an isotropic or anisotropic adaptation metric to a suite of mesh adaptation modules that also have a distributed memory, parallel execution scheme.[29,49] The FUN3D suite of codes is able to directly interface with the adaptation modules to avoid disk operations that can be costly on commodity clusters.

The adaptation modules use local mesh operators to remove edges by collapsing those that are too small as compared to the specified anisotropic metric. Edges that are too long as compared to the specified anisotropic metric are split. Grid element quality is improved by swapping edges to new configurations. Node locations are also smoothed to improve the quality of the worst incident element. This smoothing is performed in the interior and boundary of the domain. The new nodes that are inserted by an edge split on the boundary are moved to the correct position on a high fidelity boundary representation with a global linear elasticity mesh movement scheme. The interface with high fidelity boundary representation is facilitated by the GridEx framework.

The GridEx[50,51] framework is currently being developed to link various mesh generation and adaptation strategies

to geometry through Computational Analysis Programming Interface[52–54] (CAPRI). CAPRI is a CAD vendor-neutral package that provides a common interface to many native CAD kernel application programming interfaces (API). Interrogating CAD parts with their native kernel removes the difficulties of translation and associated pitfalls. The GridEx framework utilizes CAPRI to store the discretized edge and face parameters. The framework also augments the underlying CAD geometry kernels to improve the robustness of surface parameter resolution. While the actual CAD solid model is preferred, the adaptation modules provide a method of specifying analytical surface geometry. This surrogate geometry source allows the adaptation of meshes that lack the discretized edge and face parameters of a CAD solid model.

For inviscid meshes and applications, both anisotropic surface and field adaptation can be prescribed. As in Ref. 29, the mechanics to automatically adapt 3-D meshes with highly anisotropic regions near curved boundaries is still underway. Therefore, in high Reynolds number viscous cases, the mesh can only be automatically adapted in the field outside of the advancing layer structure.

VIII.   Computational Results

A.   Double-Cone Sonic Boom Configuration

The first application of the adjoint-based error estimation and adaptation procedure is a 3-D Mach 1.26 inviscid flow over two tandem cones connected by a cylinder and a wind tunnel sting at 0.0 degrees angle of attack. The tandem cone configuration is shown in Fig. 1. This is a case taken from a 1965 experiment and analysis reported in Ref. 34. The goal is to reduce the uncertainty of a calculated pressure signal in the near field of the target configuration. The pressure signal is a non-dimensional quantity defined by

$$\mathbf{S}\left(\mathbf{x}\right) = \frac{\left(\mathbf{P}\left(\mathbf{x}\right) - \mathbf{P}_\infty\right)}{\mathbf{P}_\infty} \tag{17}$$

A scalar output function is formulated by integrating Eq. (17) over a cylindrical surface with a radius equal to six model lengths. A model length is defined as the distance from the vertex of the leading cone to the base of the trailing cone. The six model length radius was chosen to match the location of experimental measurements. The cylindrical integration surface was trimmed to the specific area of interest. In this case, the cylinder extends slightly ahead and behind the expected shocks. The output function formulation, while sufficient for reducing the uncertainty in the calculated pressure signal, does not provide an intuitive means of specifying an absolute error tolerance. In contrast to a drag formulated output function where the tolerance can be defined from experience, say $\pm\mathbf{1}$ count, there is insufficient experience to judge how accurately the signal must be resolved for satisfactory propagation to the ground. For the purposes of this report, the acceptable tolerance was simply lowered repeatedly until the mesh was refined to a point beyond available resources.
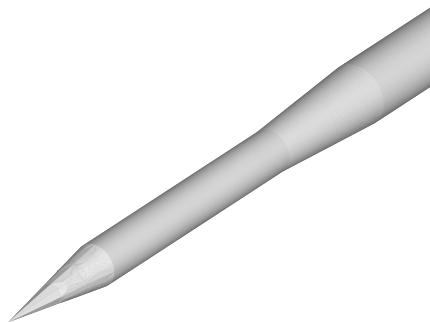

Figure 1. Perspective view of double-cone sonic boom configuration.

Due to the axis-symmetric nature of the configuration, only a 9 degree wedge of the configuration is modeled in the analysis. The initial mesh and associated pressure signal on the symmetry plane are shown in Fig. 2 (a). The initial mesh has approximately 4,600 nodes. The resulting symmetry plane mesh and associated non-dimensional pressure distribution after 13 cycles of error estimation and adaptation are shown in Fig. 2 (b). The final mesh has approximately 2.2M nodes. Note that this is a very challenging application for an adaptation scheme based solely on
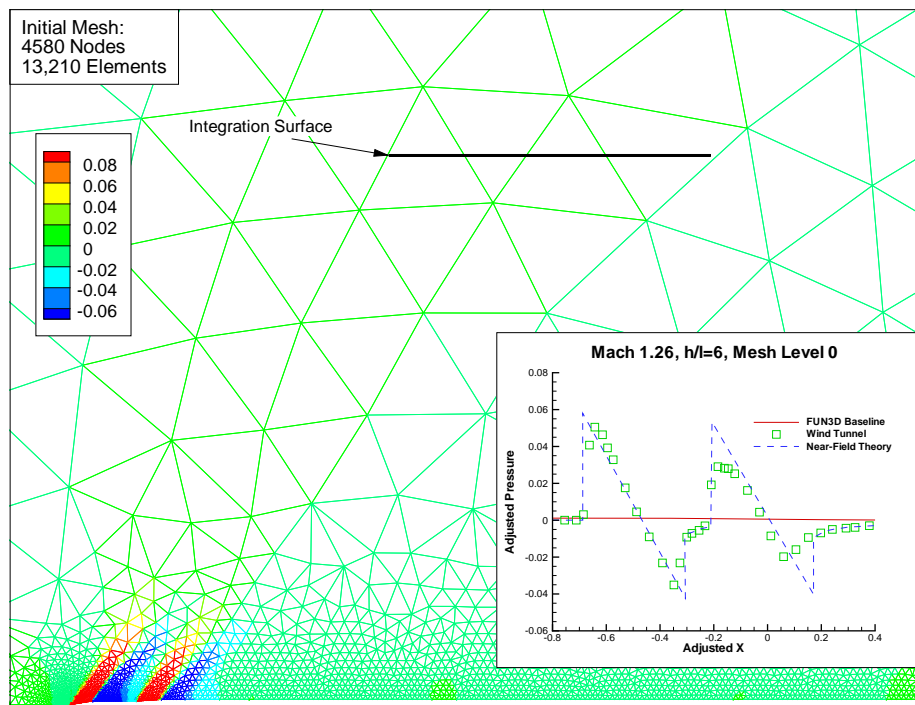
(a) Baseline mesh and solution
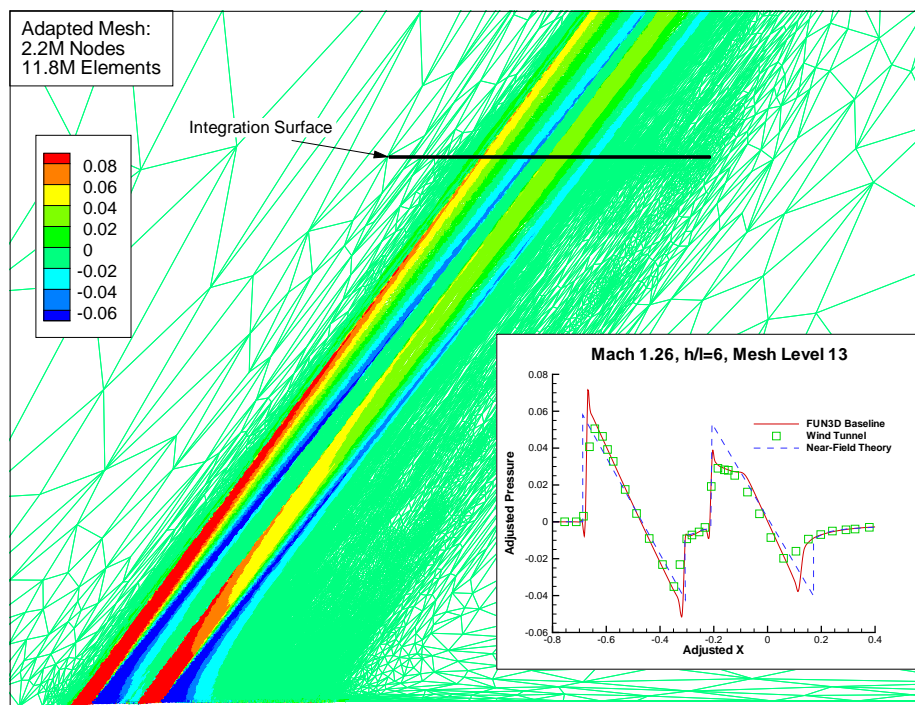
(b) Adapted mesh and solution

Figure 2. Inviscid computations on the double-cone sonic boom configuration at $M = 1.26$**.**

American Institute of Aeronautics and Astronautics Paper 2005-4842

pressure or Mach gradients because the coarse initial mesh does not resolve the extended shocks to provide a cue for refinement. All operations after initial mesh generation (flow solution, adjoint solution, error estimation, adaptation, and load balancing) are performed in an automated parallel fashion.

The insets in Fig. 2 compare the FUN3D computed shock pressure signal at six model lengths with the wind tunnel data and near-field theory. (Note that the experimental data and near-field theory was sampled from a hard copy of the upper left corner of Fig. 4 (h) on page 20 of Ref. 34.) The initial computed pressure signal in Fig.2 (a) shows no indication of the presence on the double-cone configuration at six model lengths from the body. As the error estimation and adaptation process progress, the pressure signal increases in magnitude and has sharper features. The signal at the 13th adaptation cycle is in good agreement with the experimental data as seen in the inset of Fig. 2 (b). For these computations, the second-order reconstruction for the inviscid fluxes are not limited. The overshoots and undershoots shown near the shocks in Fig. 2 (b) are thought to be due to this unlimited reconstruction.

B. Segmented Leading Edge Wing/Body Sonic Boom Configuration

Following validation of the method for near-field sonic boom pressure signatures, the technique was applied to Mach 2.0 flow over a wing/body configuration. The segmented leading edge (SLE) model shown in Fig. 3 was chosen for application. It should be noted that the model definition was modified to eliminate the blunt wing trailing edge. (See small red trailing edge highlights in Fig. 3.) The modification was necessary to resolve solution instabilities due to the strong expansion associated with the inviscid supersonic flow solution in the region of downstream facing geometry. The Roe flux scheme is used with the Venkatakrishnan[55] reconstruction limiter for the SLE study. Such instabilities have been well-documented by other researchers[56] and alternative approaches that avoid the need for geometry modifications are being actively investigated. The upper and lower wing surfaces were extended along surface tangents to the point of intersection resulting in the slight planform changes seen in the figure.
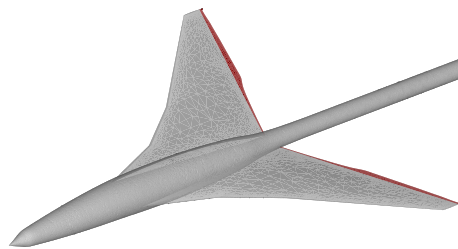

Figure 3. Perspective view of segmented leading edge (SLE) configuration.

The output function was defined in a manner consistent with the previous case, but at a distance of 2.5 body lengths where test data was available.[35] A SLE body length is defined as the distance from the tip of the nose to the aft most portion of the wing trailing edge. Half plane symmetry was assumed and the pressure signatures shown are along the body centerline (symmetry plane). The initial volume mesh had approximately 69,000 nodes. The symmetry plane mesh and associated non-dimensional pressure distribution, along with the pressure signature, are shown in Fig. 4 (a). Following 22 automated cycles of error estimation and adaptation, a mesh with just over 2 million nodes results with the associated pressure signature shown in Fig. 4 (b). Emphasis is placed on the fact that no a priori knowledge of the flow field was involved and, from the initial mesh, no human intervention was required to generate the adapted mesh of Fig. 4 (b).

The current test case is considered an initial result because of the geometry modification described above and the adaptation process has not been fully converged to the specified level. As the mesh is refined, instabilities in the adjoint system of equations have prevented further adaptation of the flowfield. The implementation of the discrete adjoint system previously described in Refs. 43, 57 and 44 has been performed using exact linearizations of the discretized residual; however, for the SLE supersonic test case, flux limiting must be used to avoid non-physical results during the higher-order reconstruction of the solution at control volume interfaces. The linearization of these flux limiters has not been accounted for in the adjoint discretization. The unstable behavior observed in the present test case is similar in nature to that observed in previous work, where subtle inaccuracies in the linearization of the discretized residual were
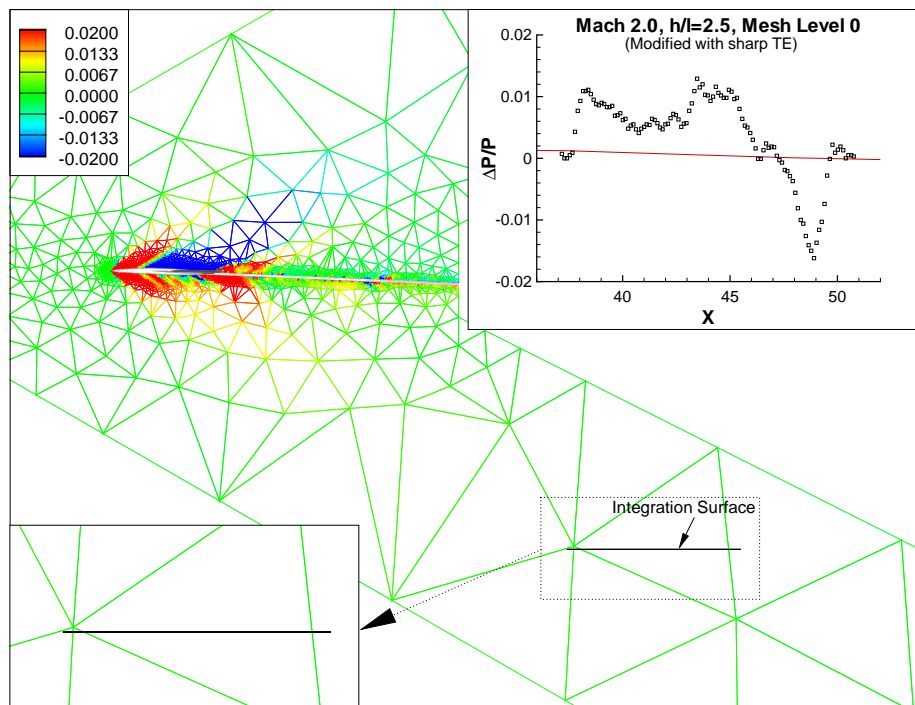
(a) Baseline mesh and solution
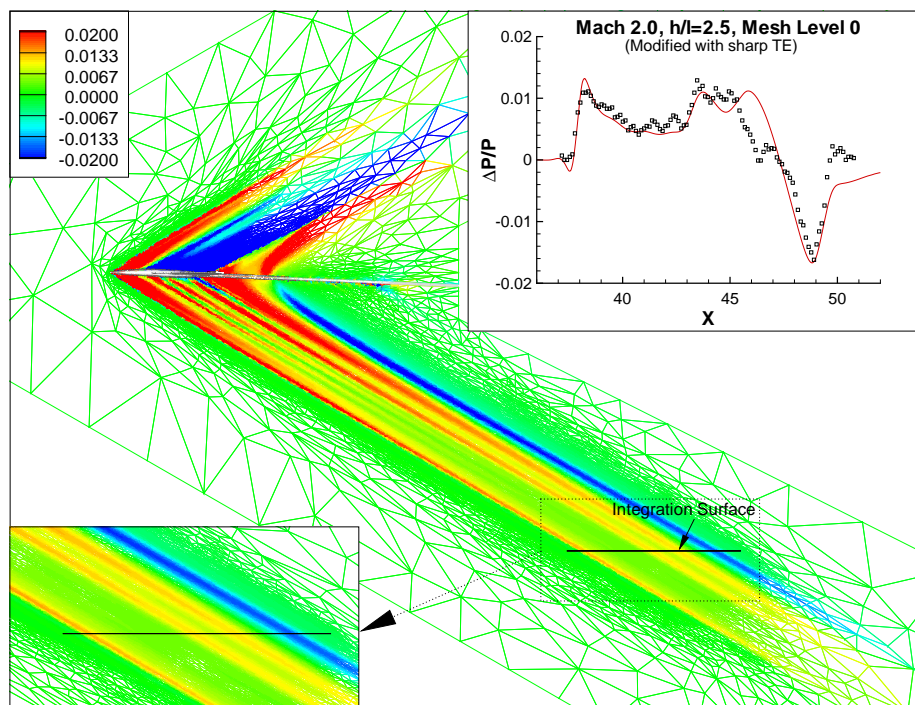
(b) Preliminary adapted mesh and solution

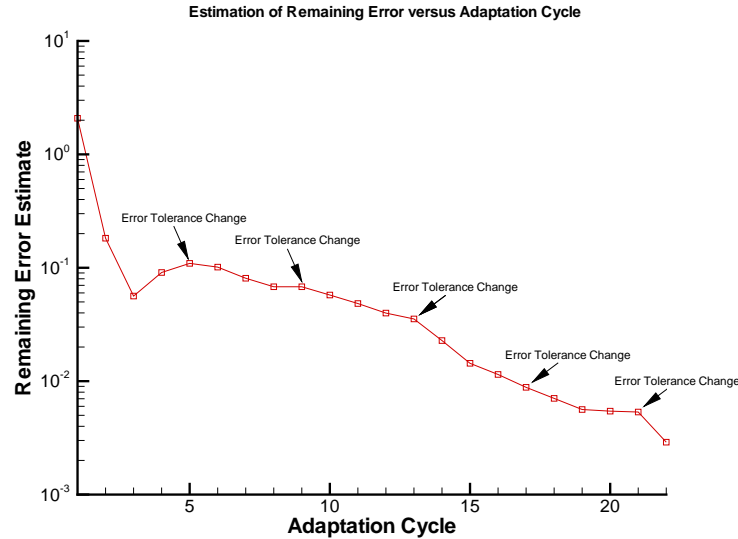**Figure 4. Inviscid computations on the SLE sonic boom configuration at** $M = 2.0$**.**

Figure 5. SLE adaptation convergence.

responsible for divergence of the adjoint solution procedure. The relevant contributions from the flux limiting scheme are currently being implemented so that the present test case can be performed to completion.

However, as seen in Fig. 4, much improvement has been gained by the automated output-based adaptation. In fact, some characteristics of the signal missing from the current solution appear to be initiated and are expected to resolve with continued adaptation. Note the two compressions located at approximately X = 41.4 and X = 42.4 on the 2-D plot of Fig. 4 that currently propagate out to roughly one-third the desired distance away from the body on the symmetry plane. This propagation is expected to extend as adaptation continues.

As mentioned in the last section, the determination of the error tolerance on the sonic boom output function is non-intuitive. As such, a predefined schedule has driven both cases whereby the tolerance is given as some percentage of the calculated remaining error estimate on the baseline mesh. Since the output-based adaption is self terminating, the tolerance is periodically reduced following several cycles of adaptation to converge the overall process. Figure 5 shows an example of the convergence for the SLE adaptation. Here the error tolerance was reduced by half every four adaptation cycles. The reduction is scripted and, thereby, does not destroy the process automation. Also, given that the tolerance is based on a percentage of the calculated error estimation, the schedule is not based on a priori knowledge of the flow field.

C. DLR-F6 Wing/Body Transonic Cruise Configuration

The DLR-F6 is representative of a modern twin-engine wide-body transport aircraft. This configuration was the focus of the DPW-II.[2,36] Multiple engine geometries and installation locations were tested.[58] However, only the clean wing/body was analyzed in this study. The design cruise Mach number for the DLR-F6 is $M_\infty = 0.75$, and the lift coefficient is $C_L = 0.500$. The aspect ratio is 9.5 and the leading edge sweep is $27.1°$. The computational geometry as defined by the DPW-II committee is used for the surface definition of all computational meshes.[36] The experimental data used for comparison in this paper were also provided by the DPW-II organizing committee.[36] In the experiment, the wing boundary layer transition was tripped on the lower surface at 25 percent of chord and on the upper surface at 5 percent of chord at the root, 15 percent at the crank, 15 percent at $\eta = 0.844$, and 5 percent at the tip. However, the current analysis did not model the transition.

The current analysis focuses on the constant lift case (Case 1) from the DPWII workshop: $M_\infty = 0.75$, $C_L = 0.500$, fully turbulent. In Ref. 41, Case 1 was computed on three globally refined meshes: a one million node mesh (coarse), a 3 million node mesh (medium), and a 9 million node mesh (fine). Results from the Ref. 41 calculations are shown in Fig. 6 in terms of total drag, angle of attack, pressure drag, viscous drag, and pitching moment versus $N^{-2/3}$,

where N is the number of nodes in the mesh. Thus, results using fine mesh appear to the left in the plots, and results using coarser mesh appear to the right. The experimental values are included for reference. In the asymptotic range, one would expect a linear variation in forces or moments with $N^{-2/3}$ for a second order scheme. (Table 1 shows a summary of Case 1 calculations along with the experimental data for reference.) The dashed line in Fig. 6 shows a Richardson's extrapolation computed in Ref. 41, assuming a second order convergence rate between the medium and fine mesh. Based on this extrapolation, the infinite-mesh total drag is 278. For the total drag force, the observed order of convergence[59] is 3.4. Using the observed convergence rate, the infinite-mesh total drag is 280.

Nodes	Mesh	Specified Transition	α	C_L	C_D	C_{Dp}	C_{Dv}	C_M
Exp.	—	Yes	0.52	0.500	0.0295	—	—	-0.1211
1.1*	Coarse	No	0.102	0.500	0.03034	0.01812	0.01221	-0.1309
3.0*	Medium	No	0.201	0.500	0.02857	0.01646	0.01210	-0.1269
3.0*	Medium	Yes	0.059	0.500	0.02747	0.01586	0.01161	-0.1331
9.1*	Fine	No	0.263	0.500	0.02812	0.01600	0.01212	-0.1254
65	Finest	No	0.827	0.499	0.0295	0.01790	0.01159	-0.1072

* Reference 41

Table 1. Summary of global mesh refinement results for the DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.

In order to provide an additional validation point for the global refinement study, a globally refined mesh of 65 million nodes was generated with VGRIDns[60] for the DLR-F6 wing/body. This mesh, which is topologically similar to the coarse, medium and fine meshes, has one million nodes on the wing/body surface. This mesh has over double the maximum number of total unknowns that were considered in the DPW-II workshop for the wing/body configuration, which was 23 million.[2] The results from the 65 million nodes (finest) mesh are also included in Fig. 6 and Table 1. The finest mesh results illustrate that although the coarse, medium, and fine mesh results from Ref. 41 appeared to be asymptotic, they were not yet in the linear range. Figure 7 shows the mesh convergence of the wing chordwise surface pressure distributions for the coarse, medium, fine, and finest meshes. The computational results are shown at seven of the eight experimental span locations, along with the experimental pressure coefficients for reference. These pressure distributions indicate two of the relevant flow features at this lift condition: a separation bubble near the trailing edge of the upper wing-root juncture and a mild normal shock across the span of the upper wing near the quarter chord. In Ref. 41, the computations showed that a decrease in the upper-surface suction on the inboard span station resulted from an increase in the area of wing-root juncture separation. The comparison shown in Fig. 7 indicates that the area of wing-root juncture separation is much larger on the finest mesh than on the coarse, medium and fine meshes. Note that the coarse, medium and fine mesh pressure distributions are very similar. The fine mesh solution indicates slightly more inboard separation and slightly stronger shocks than the coarse and medium mesh solutions. The loss of lift because of the flow separation in the finest mesh solutions requires a larger angle of attack to maintain the constant lift condition for the finest mesh as shown in Fig. 6. The increase in angle of attack, in turn, strengthens the shock and moves it aft. The pressure distributions from the finest mesh solution compare more favorably with the experimental data on the outboard section of the wing. However, the pressures on the inboard station show a much larger separation than indicated by the experiment.

For the current work, the workshop medium and fine meshes were used as a starting point for the error estimation and field adaptation process with total drag as the output function of interest. The term "field adaptation" is used to emphasize the fact that the advancing-layer mesh near the model was not modified. Tables 2 and 3 show a summary of the results of five cycles of error estimation and adaptation. For each new mesh, the angle of attack is readjusted to maintain the $C_L = 0.500$. The total drag resulting from the flow solution C_D is compared with the total drag corrected by the adjoint-based error corrected drag, C_{Dcorr}. The estimate of the remaining error in the total drag after the error correction is also included. The user-specified error tolerance t starts at 10 counts and is reduced to 5 counts in the second adaptation cycle. For the fine-mesh adaptation, the error tolerance is further reduced to 3 counts in the third adaptation cycle. Note in Tables 2 and 3 that the corrected drag computed on the baseline meshes is lower than the uncorrected drag. These results are consistent with the infinite-mesh Richardson's extrapolation in the uniform mesh refinement study. However, after several cycles of adaptation, the uncorrected and corrected drag values are higher
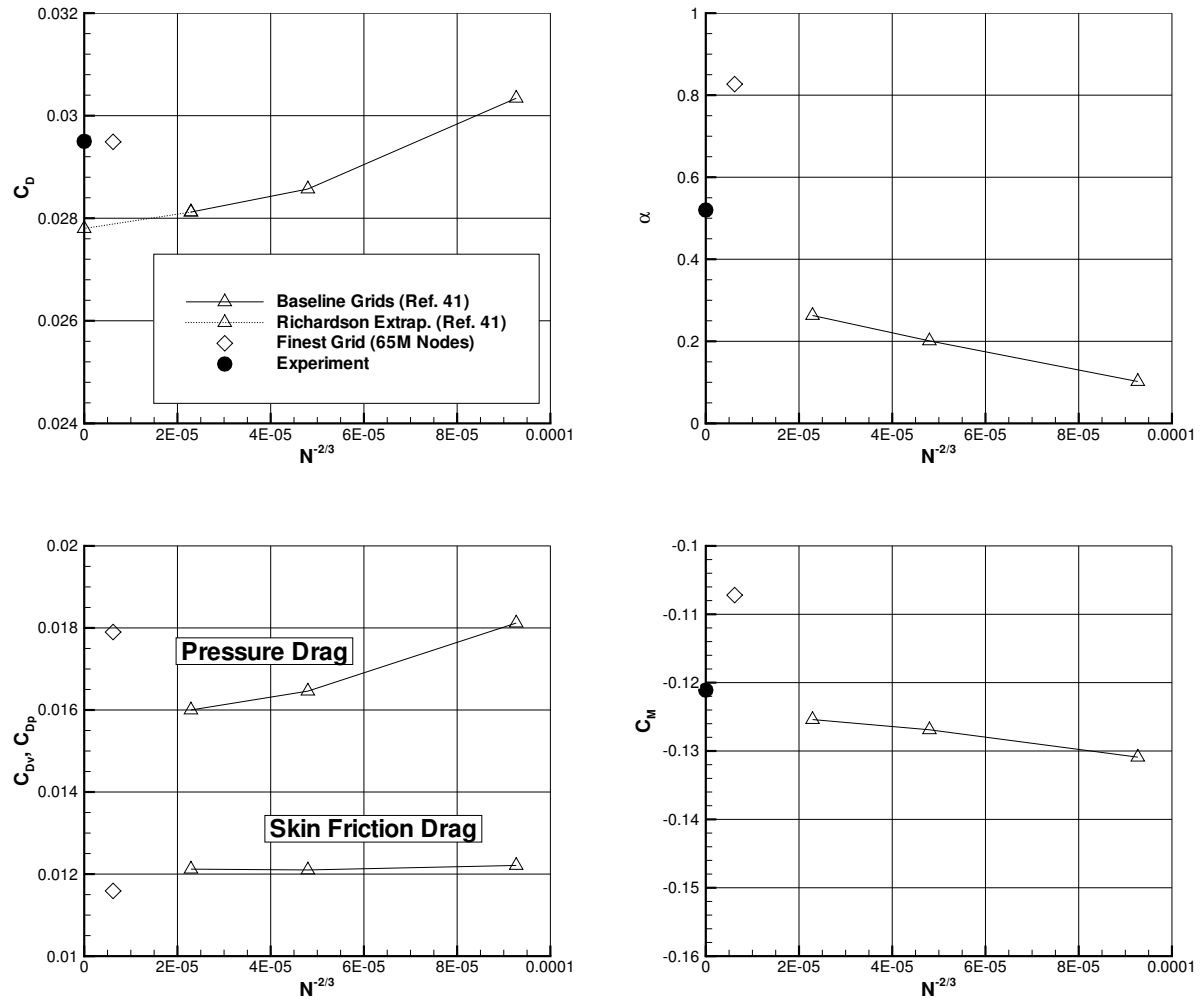
Figure 6. Uniform mesh refinement results for the DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.
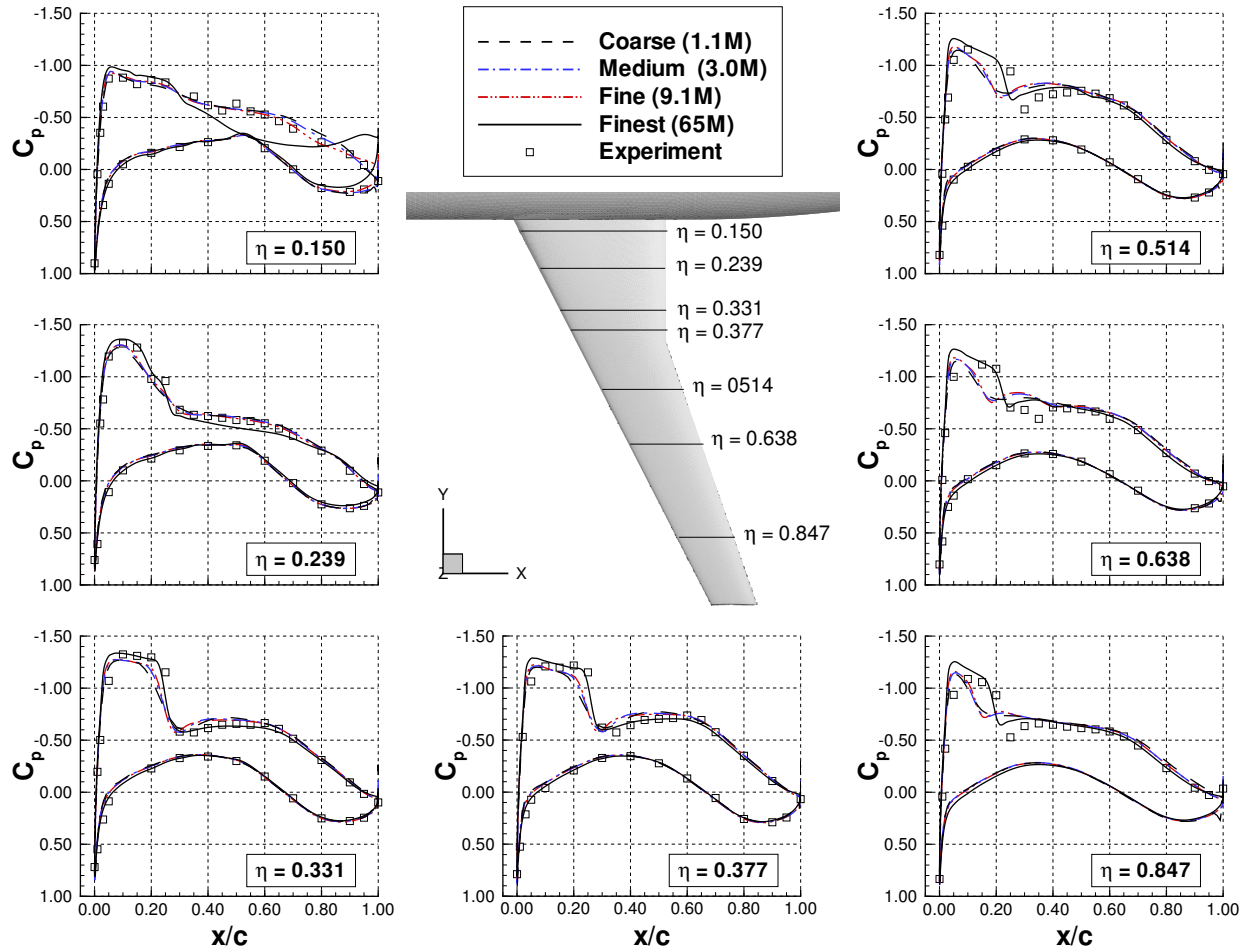
Figure 7. Chordwise pressure distributions for the DLR-F6 wing/body uniform mesh refinement study at $M = 0.75$, $C_L = 0.500$.

than the baseline predictions, and they are increasing toward the finest mesh value. Also, the estimate of remaining error has dropped to almost half the value estimated on the baseline mesh. For the final medium-mesh adaptation, the final corrected drag is 291 counts, and the estimate of the remaining drag error is ±16 counts of drag. For the final fine-mesh adaptation, the final corrected drag is 288 counts, and the estimate of the remaining drag error is ±10 counts of drag. Although further adaptation is required to meet the requested tolerance, no further cycles were computed to avoid large discrepancies between the surface mesh spacing and the refined field spacing because of the adaptation.

Mesh	Number of Nodes (Millions)	α (deg.)	C_L	C_D (counts)	C_{Dcorr} (counts)	Specified Tol. (counts)	Est. Error (counts)
Baseline	3.0	0.201	0.500	286	278	10	±25
First Adapted	2.9	0.269	0.500	286	281	10	±24
Second Adapted	2.8	0.350	0.500	286	281	5	±26
Third Adapted	3.3	0.462	0.500	291	290	5	±21
Fourth Adapted	3.7	0.443	0.500	291	291	5	±16

Table 2. Summary of medium mesh field adaptation results for the DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.

Mesh	Number of Nodes (Millions)	α (deg.)	C_L	C_D (counts)	C_{Dcorr} (counts)	Specified Tol. (counts)	Est. Error (counts)
Baseline	9.1	0.263	0.500	281	276	10	±20
First Adapted	8.2	0.342	0.500	286	286	10	±13
Second Adapted	7.4	0.342	0.501	287	288	5	±11
Third Adapted	7.9	0.325	0.500	287	288	3	±10
Fourth Adapted	9.4	0.325	0.501	287	288	3	±10

Table 3. Summary of fine mesh field adaptation results for the DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.

Figure 8 shows a comparison of the total drag, angle of attack, pressure drag, viscous drag, and pitching moment versus $N^{-2/3}$ between the baseline unadapted meshes and the final adapted meshes. The corrected drag and the estimate of the remaining drag error for the baseline meshes, as well as the final adapted meshes, are also included in the total drag plot. The estimate of the remaining drag error is shown as an error bar on the corrected drag value. Figure 8 indicates that the total drag, corrected drag, and angle-of-attack values for the final adapted meshes are approaching the finest mesh values. The final adapted-medium mesh value of drag is closer to the finest mesh value than the final adapted-fine mesh value. However,the final adapted-fine mesh indicates a much lower remaining error. Figure 8 shows that the finest mesh total drag value does lie within the estimate of the error bounds of all of the corrected drag values. It is interesting to note that the pressure drag, viscous drag, and pitching moment values for the adapted meshes do not tend to approach the finest mesh values. The process of adapting the mesh and changing the angle of attack to maintain constant lift makes it difficult to determine the influence of adaptation on the separate effects of the wing-root juncture separation and the normal shock strength and location. An adaptation study at a constant angle of attack could add some insight into the trends.

Figure 9 shows a comparison of the chordwise pressure distributions for the adapted and unadapted medium meshes. The finest mesh results and experimental data are included for reference. As the medium mesh is field adapted and the angle of attack is increased to maintain the constant lift condition, the upper-surface shock on the outboard portion of the wing slightly strengthens and moves aft as the inboard pressure distribution shows a decrease in the suction, which has been found to indicate increased wing-root juncture separation. The final medium-adapted mesh computation does not predict as much wing-root juncture separation as the finest mesh computation. The final medium-adapted mesh computation also does not predict as strong a shock as the finest mesh computation. A similar trend is shown in Fig. 10 for the the adapted and unadapted fine meshes. However, the fine mesh field adaptation does not have as much of an impact on the normal shock strength and location.
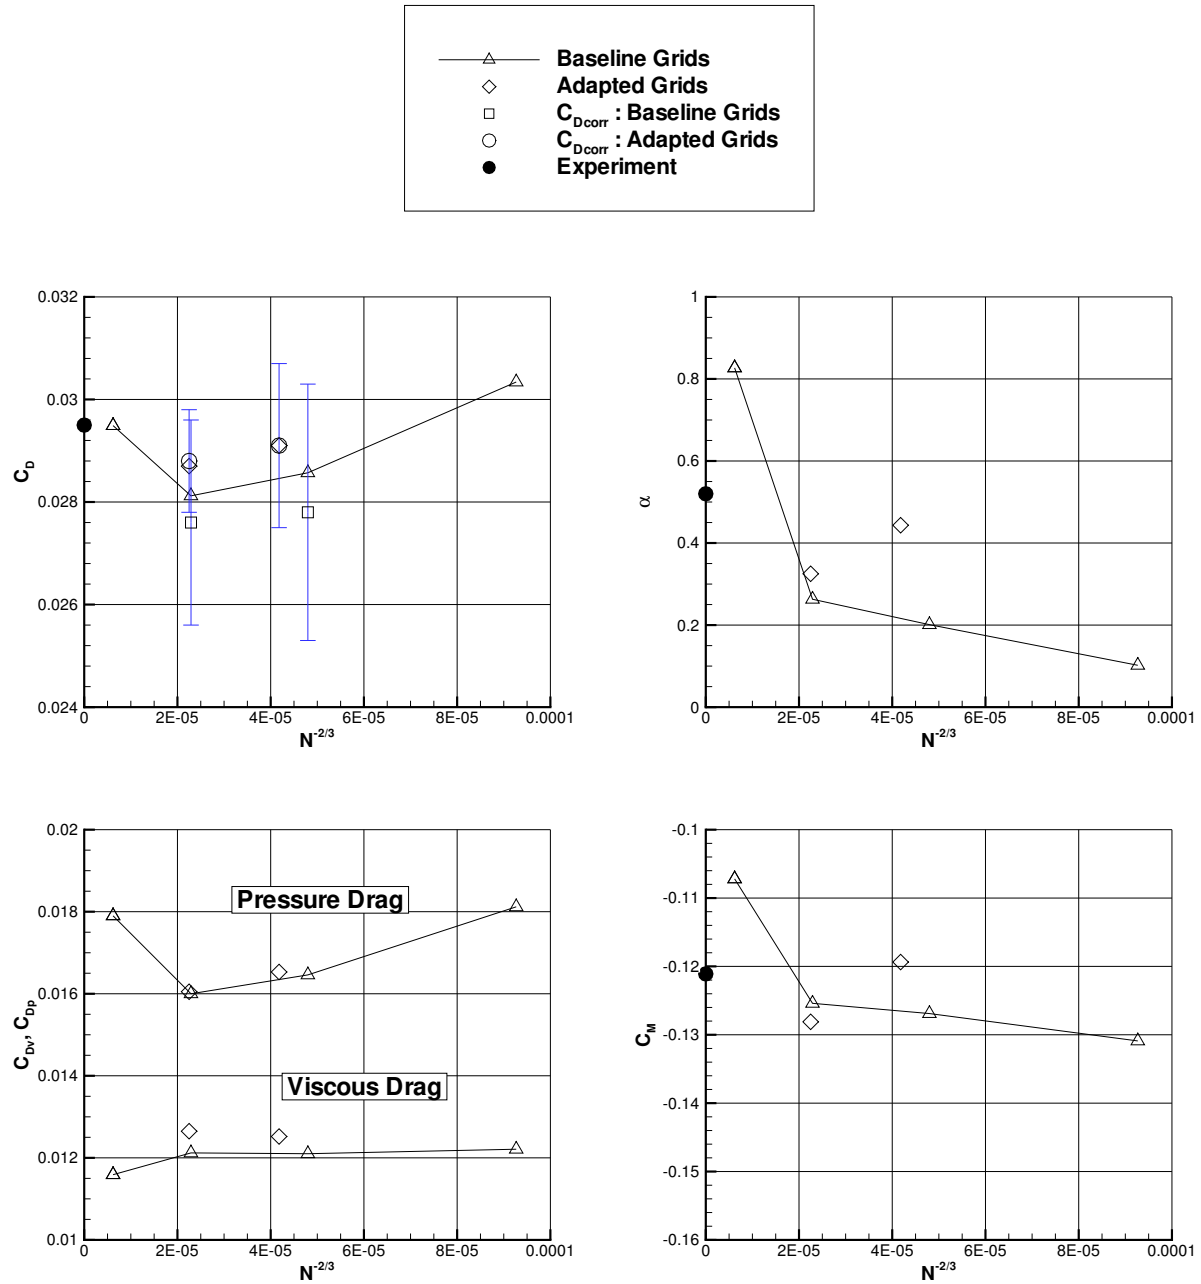
Figure 8. Comparison of adjoint-based field adaptation results with uniform mesh refinement study for the DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.
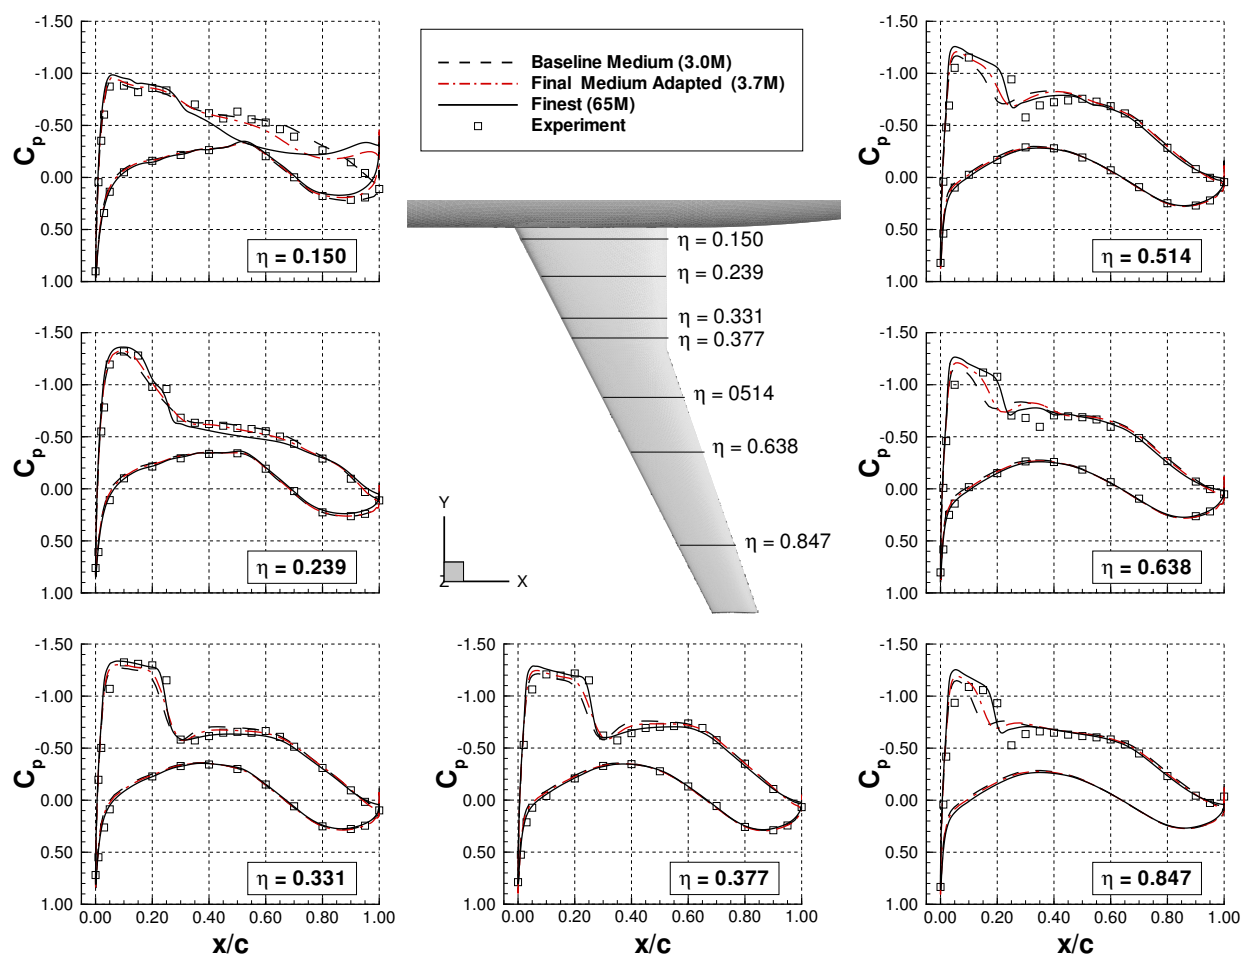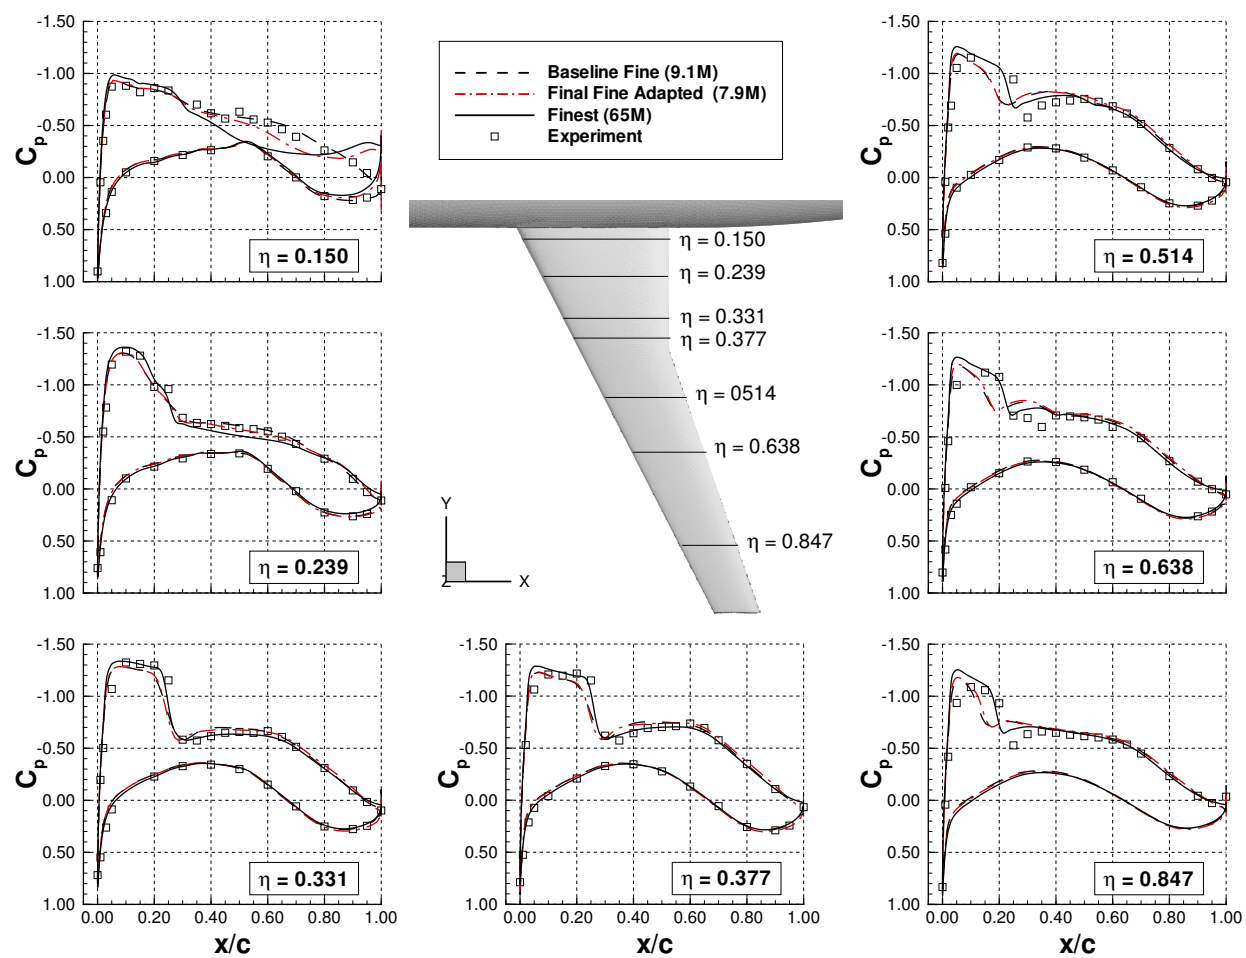
Figure 9. Effect of adjoint-based field adaptation on the medium mesh chordwise pressure distributions for DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.

American Institute of Aeronautics and Astronautics Paper 2005-4842

Figure 10. Effect of adjoint-based field adaptation on the fine mesh chordwise pressure distributions for DLR-F6 wing/body at $M = 0.75$, $C_L = 0.500$.

American Institute of Aeronautics and Astronautics Paper 2005-4842

Figure 11 shows the adjoint-based adaptation parameter for the baseline medium mesh solution on the wing/body, symmetry plane and mid-span plane (inset). The corresponding symmetry plane and mid-span plane meshes are shown in Fig. 12. The scale is an indication of how much the current mesh needs to be refined (or coarsened) to achieve the desired user-specified tolerance in the output function (total drag). Values less than one indicate a need for refinement and values greater than one indicate an allowance for coarsening. The scale in Fig. 12 is capped at one to highlight the areas of refinement. The scale in Fig. 11 indicates that surface mesh is in need of refinement overall but especially in areas on the fuselage. Also there is need for refinement in wake region of the fuselage as well as the near wake of the wing trailing edge (see inset).

Figure 13 shows the adjoint-based adaptation parameter for the final medium-mesh adapted solution on the wing/body, symmetry plane and mid-span plane (inset). The corresponding symmetry plane and mid-span plane meshes are shown in Fig. 14. The scale in Fig. 13 indicates that surface mesh is in need of even more refinement than requested in the baseline results due in part to the reduction of the user-specified error tolerance from 10 count to 5 counts. Many of the areas requiring refinement on the symmetry plane field have been reduced in Fig. 13, but the refinement area in the wake region behind the wing trailing edge (see inset) has extended further downstream. In Fig. 14, the formation of a refined area behind the fuselage and wing trailing is evident. The corresponding adaptation parameter and resulting adapted mesh for the final fine-mesh adapted solution are shown in Figs. 15 and 16. Note the user-specified error tolerance has been reduced to 3 counts of drag. The scale in Fig. 15 indicates additional surface mesh refinement for the specified error tolerance, especially on the fuselage. Also of interest is the circumferential band of requested refinement along the span of the fuselage. These bands correspond to the IGES surface definition and VGRIDns patch boundaries. In Fig. 16, the formation of a refined area behind the fuselage and wing trailing is again evident.

IX.   Summary

This paper demonstrated the extension of the error estimation and adaptation methods in Refs. 28 and 29 to parallel computations enabling larger, more realistic aerospace applications and the quantification of discretization errors for complex 3-D solutions. Results were shown for an inviscid sonic-boom prediction about a double-cone configuration and a wing/body segmented leading edge (SLE) configuration where the output function of the adjoint was pressure integrated over a part of the cylinder in the near field. After multiple cycles of error estimation and surface/field adaptation, a significant improvement in the inviscid solution for the sonic boom signature of the double cone was observed. Although the double-cone adaptation was initiated from a very coarse mesh, the near-field pressure signature from the final-adapted mesh compared very well with the wind-tunnel data which illustrates that the adjoint-based error estimation and adaptation process requires no a priori refinement of the mesh. Similarly, the near-field pressure signature for the SLE wing/body sonic boom configuration showed a significant improvement from the initial coarse mesh to the final adapted mesh in comparison with the wind tunnel results. The SLE results were considered to be initial because of the geometry modification needed for flow solver stability, and also because the adaptation process has not been fully converged to the desired level. As the mesh was adapted, instabilities in the adjoint system of equations from the absence of flux-limiting in the adjoint formulation prevented further refinement of the flowfield. Both issues of flow solver instability and adjoint instability are the focus of current work.

Error estimation and field adaptation results were also presented for the viscous drag prediction of the DLR-F6 wing/body configuration, and results were compared to a series of globally refined meshes. Two of these globally refined meshes were used as a starting point for the error estimation and field-adaptation process where the output function for the adjoint was the total drag. The field-adapted results showed an improvement in the prediction of the drag in comparison with the finest globally refined mesh and a reduction in the estimate of the remaining drag error. However, in comparison with the wind tunnel data, the CFD analysis showed an over-prediction in wing-root juncture separation with mesh refinement. The prediction of the wing upper-surface normal shock location was improved with mesh refinement. Future investigation of the effects of the turbulence model and/or transition could improve our understanding of the modeling error associated with the over-prediction of the wing-root juncture separation. The adjoint-based adaptation parameter showed a need for increased resolution in the surface of the wing/body as well as a need for wake resolution downstream of the fuselage and wing trailing edge in order to achieve the requested drag tolerance. Although further adaptation was required to meet the requested tolerance, no further cycles were computed in order to avoid large discrepancies between the surface mesh spacing and the refined field spacing. The capability to refine the surface and the highly-stretched mesh cells just off the surface is a focus of current work.
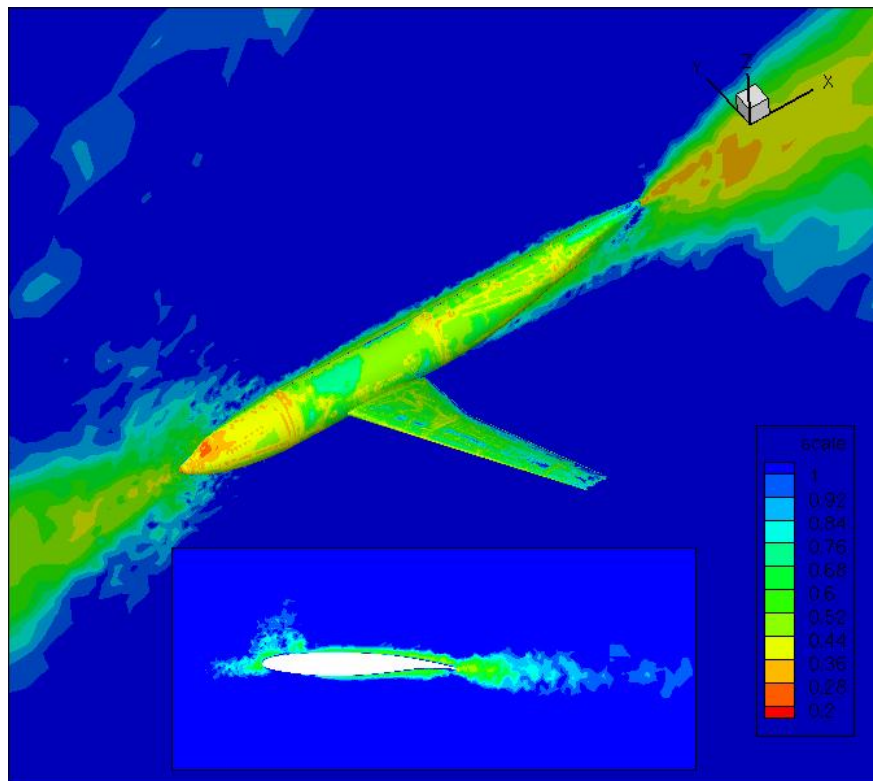
Figure 11. Adjoint-based adaptation parameter for the baseline medium mesh solution on the wing/body, symmetry plane and mid-span plane (inset).
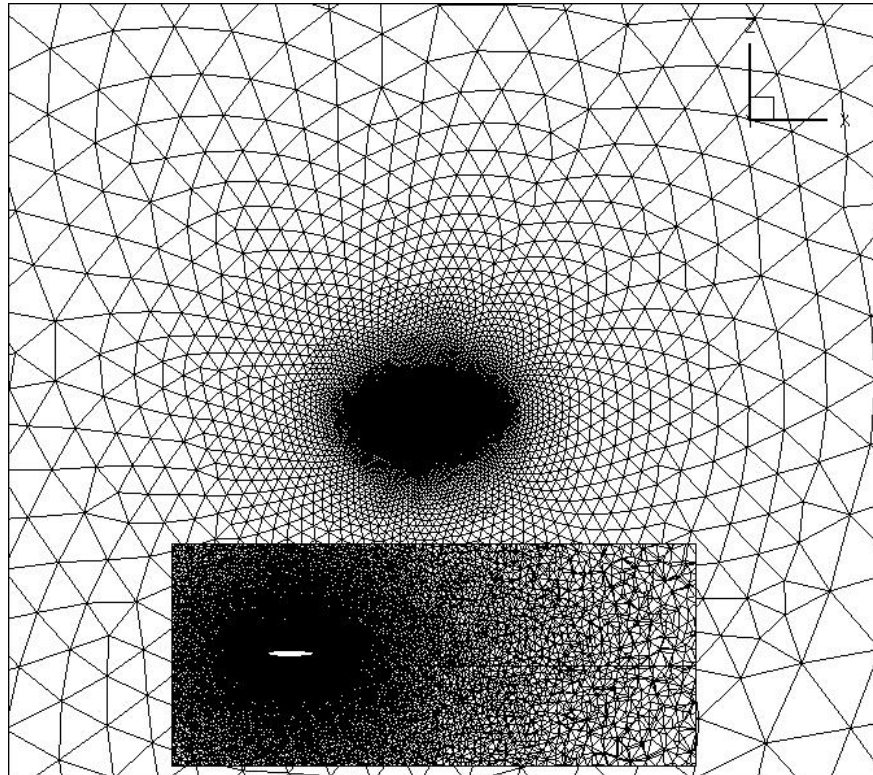

Figure 12. Baseline medium mesh symmetry plane and mid-span plane (inset).

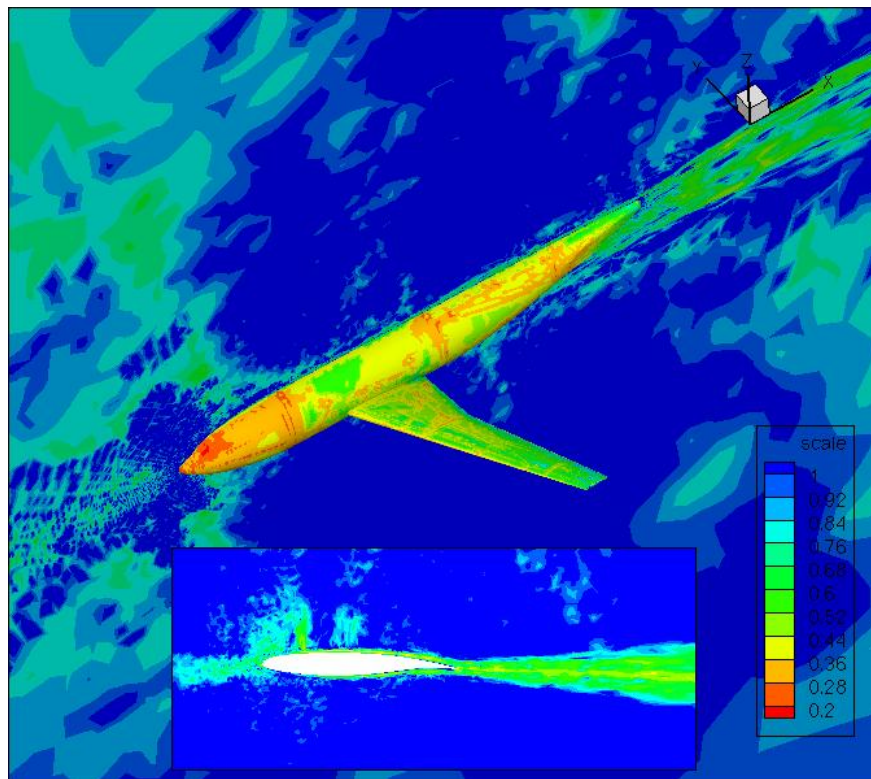American Institute of Aeronautics and Astronautics Paper 2005-4842

Figure 13. Adjoint-based adaptation parameter for the final medium-adapted mesh solution on the wing/body, symmetry plane and mid-span plane (inset).
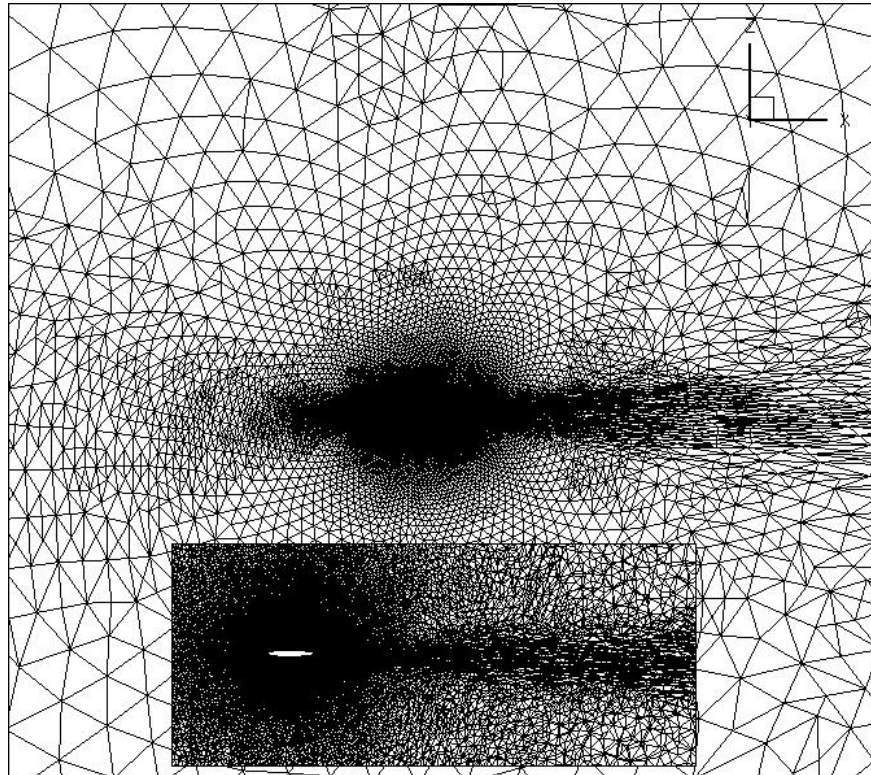

Figure 14. Final medium-adapted mesh symmetry plane and mid-span plane (inset).
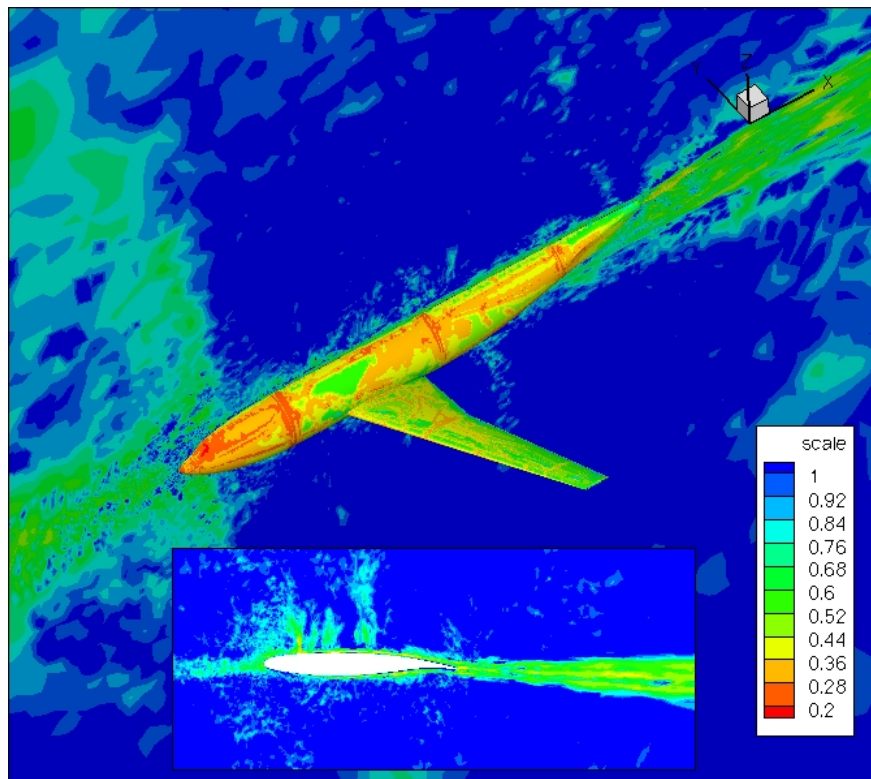
Figure 15. Adjoint-based adaptation parameter for the final fine-adapted mesh solution on the wing/body, symmetry plane and mid-span plane (inset).
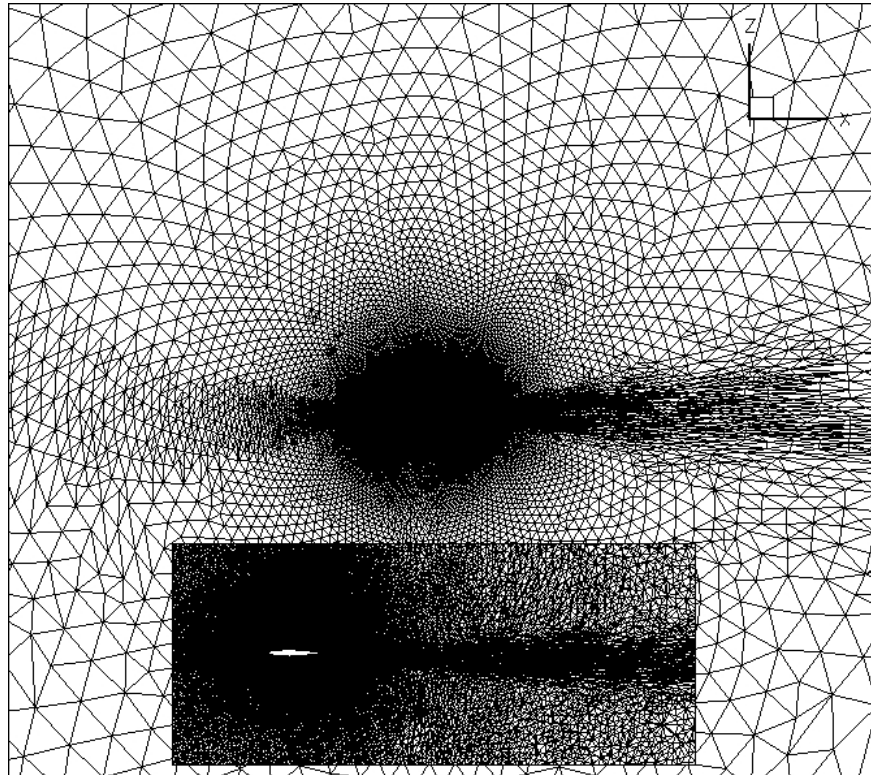

Figure 16. Final fine-adapted mesh symmetry plane and mid-span plane (inset).

X.   Acknowledgments

References

[1]Rumsey, C. L., Lee-Rausch, E. M., and Watson, R. D., "Three-Dimensional Effects on Multi-Element High Lift Computations," *Computers and Fluids*, Vol. 32, No. 5, 2003, pp. 631–657, See also AIAA Paper 2002–0845.

[2]Laflin, K. R., Vassberg, J. C., Wahls, R. A., Morrison, J. H., Brodersen, O., Rakowitz, M., Tinoco, E. N., and Godard, J.-L., "Summary of Data from the Second AIAA CFD Drag Prediction Workshop," AIAA Paper 2004-0555, Jan. 2004.

[3]Hemsch, M. J. and H., M. J., "Statistical Analysis of CFD Solutions from 2nd Drag Prediction Workshop (Invited)," AIAA Paper 2004-0556, Jan. 2004.

[4]Levy, D. W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R. A., Pirzadeh, S., and Hemsch, M. J., "Data Summary from the First AIAA Computational Fluid Dynamics Drag Prediction Workshop," *AIAA Journal of Aircraft*, Vol. 40, No. 5, 2003, pp. 875–882, See also AIAA Paper 2002–0841.

[5]Hemsch, M. J., "Statistical Analysis of Computational Fluid Dynamics Solutions from the Drag Prediction Workshop," *AIAA Journal of Aircraft*, Vol. 41, No. 1, 2004, pp. 95–103, See also AIAA Paper 2002–0842.

[6]Peraire, J., Vahdati, M., Morgan, K., and Zienkiewicz, O. C., "Adaptive Remeshing for Compressible Flow Computations," *Journal of Computational Physics*, Vol. 72, 1987, pp. 449–466.

[7]Warren, G. P., Anderson, W., Thomas, J. L., and Krist, S. L., "Grid Convergence for Adaptive Methods," Tech. Rep. AIAA-91-1592, June 1991.

[8]Baker, T. J., "Mesh Adaptation Strategies for Problems in Fluid Dynamics," *Finite Elements in Analysis and Design*, Vol. 25, No. 3–4, 1997, pp. 243–273.

[9]Pirzadeh, S. Z., "A Solution-Adaptive Unstructured Grid Method by Grid Subdivision and Local Remeshing," *AIAA Journal of Aircraft*, Vol. 37, No. 5, September–October 2000, pp. 818–824, See also AIAA Paper 99–3255.

[10]Mavriplis, D., "Adaptive meshing techniques for viscous flow calculations on mixed element unstructured meshes," *International Journal for Numerical Methods in Fluids*, Vol. 34, No. 2, 2000, pp. 93–111, See also AIAA Paper 97-0857.

[11]Aftosmis, M. J. and Berger, M. J., "Multilevel Error Estimation and Adaptive h-Refinement for Cartesian Meshes with Embedded Boundaries," AIAA Paper 2002-0863, 2002.

[12]Habashi, W. G., Dompierre, J., Bourgault, Y., Ait-Ali-Yahia, D., Fortin, M., and Vallet, M., "Anisotropic mesh adaptation: towards user-independent, mesh-independent and solver-independent CFD. Part I: general principles," *International Journal of Numerical Methods in Fluids*, Vol. 32, 2000, pp. 725–744.

[13]Castro-Diáz, M. J., Hecht, F., Mohammadi, B., and Pironneau, O., "Anisotropic Unstructured Mesh Adaptation for Flow Simulations," *International Journal of Numerical Methods in Fluids*, Vol. 25, 1997, pp. 475–491.

[14]Pelletier, D., Turgeon, E., Lacasse, D., and Borggaard, J., "Adaptivity, Sensitivity, and Uncertainty: Toward Standards of Good Practice in Computational Fluid Dynamics," *AIAA Journal*, Vol. 41, No. 10, October 2003, pp. 1925–1933.

[15]Choi, S., Alonso, J. J., and Van der Weide, E., "Numerical and Mesh Resolution Requirements for Accurate Sonic Boom Prediction of Complete Aircraft Configurations," AIAA Paper 2004–1060, 2004.

[16]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.

[17]Venditti, D. A., *Grid Adaptation for Functional Outputs of Compressible Flow Simulations*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 2002.

[18]Monk, P. and Süli, E., "The Adaptive Computation of Far-Field Patterns by A Posteriori Error Estimation of Linear Functionals," *SIAM Journal on Numerical Analysis*, Vol. 36, 1998, pp. 251–274.

[19]Paraschivoiu, M., Peraire, J., and Patera, A., "A Posteriori Finite Element Bounds for Linear-functional Outputs of Elliptic Partial Differential Equations," *Computer Methods in Applied Mechanics and Engineering*, Vol. 150, 1997, pp. 289–312.

[20]Braack, M. and Rannacher, R., "Adaptive Finite Element Methods for Low-Mach Flows with Chemical Reactions," Vol. 3 of *30th Computational Fluid Dynamics*, von Karman Institute, 1999, pp. 1–93.

[21]Rannacher, R., "Adaptive Galerkin Finite Element Methods for Partial Differential Equations," *Journal of Computational and Applied Mathematics*, Vol. 128, 2001, pp. 205–233.

[22]Becker, R. and Rannacher, R., "An optimal control approach to error control and mesh adaptation," *Acta Numerica 2001*, 2001.

[23]Giles, M. B. and Süli, E., "Adjoint methods for PDEs: a posteriori error analysis and postprocessing by duality," *Acta Numerica*, Vol. 11, 2002, pp. 145–236.

[24]Peraire, J. and Patera, A. T., "Asymptotic a posteriori finite element bounds for the outputs of noncoercive problems: the Helmholtz and Burgers equations," *Computer Methods in Applied Mechanics and Engineering*, Vol. 171, 1999, pp. 77–86.

[25]Barth, T. and Deconinck, H., editors, *Error Estimation and Adaptive Discretization Methods in Computational Fluid Dynamics*, Number 25 in Lecture Notes in Computational Science and Engineering, Springer-Verlag, 2002.

[26]Pierce, N. A. and Giles, M. B., "Adjoint Recovery of Superconvergent Functionals from PDE Approximations," *SIAM Review*, Vol. 42, No. 2, 2000, pp. 247–264.

[27]Müller, J. D. and Giles, M. B., "Solution Adaptive Mesh Refinement Using Adjoint Error Analysis," AIAA Paper 2001–2550, 2001.

[28]Park, M. A., "Adjoint–Based, Three–Dimensional Error Prediction and Grid Adaptation," *AIAA Journal*, Vol. 42, No. 9, September 2004, pp. 1854–1862, See also AIAA Paper 2002–3286.

[29]Park, M. A., "Three-Dimensional Turbulent RANS Adjoint-Based Error Correction," AIAA Paper 2003-3849, June 2003.

[30]Morgenstern, J. M., Arslan, A., Lyman, V., and Vadyak, J., "F-5 Shaped Sonic Boom Demonstrator's Persistence of Boom Shaping Reduction through Turbulence," AIAA Paper 2005–12, 2005.

[31]Kandil, O. A., Yang, Z., and Bobbitt, P. J., "Prediction of Sonic Boom Signature Using Euler-Full Potential CFD with Grid Adaptation and Shock Fitting," AIAA Paper 2002–2542, 2002.

[32]Lyman, V. and Morgenstern, J. M., "Calculated and Measured Pressure Fields for an Aircraft Desgined fro Sonic-boom Alleviation," AIAA Paper 2004–4846, 2004.

[33]Djomehri, M. J. and Erickson, L. L., *An Assessment of the Adaptive Unstructured Tetrahedral Grid, Euler Flow Solver Code FELISA*, NASA/TP-1994-3525, 1994.

[34]Carlson, H., Mack, R., and Morris, O., *A Wind-Tunnel Investigation of the Effect of Body Shape on Sonic-Boom Pressure Distributions*, NASA/TN-1965-3103, 1965.

[35]Mack, R. J. and Kuhn, N., *Determination of Extrapolation Distance With Measured Pressure Signatures From Two Low-Boom Models*, NASA/TM-2004-213264, 2004.

[36]Anon., "2nd AIAA CFD Drag Prediction Workshop. Orlando, FL." //http:aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw, June 2003.

[37]Brodersen, O., Rakowitz, M., Amant, S., Larrieu, P., Destarac, D., and Sutcliffe, M., "Airbus, ONERA, and DLR Results from the 2nd AIAA Drag Prediction Workshop," AIAA Paper 2004–0391, 2004.

[38]Sclanfani, A. J., DeHaan, M. A., and Vassberg, J. C., "OVERFLOW Drag Prediction for the DLR-F6 Transport Configuration; A DPW-II Case Study," AIAA Paper 2004–0393, 2004.

[39]Rumsey, C. L., Rivers, S. M., and Morrison, J. H., "Study of CFD Variations on Transport Configurations from the Second Drag-Prediction Workshop," AIAA Paper 2004–0394, 2004.

[40]Wurtzler, K. E. and Morton, S., "Accurate Drag Prediction Using COBALT," AIAA Paper 2004–0395, 2004.

[41]Lee-Rausch, E. M., Frink, N. T., Mavriplis, D. J., Rausch, R. D., and Milholen, W. E., "Transonic Drag Prediction on a DRL-F6 Transport Configuration Using Unstructured Grid Solvers," AIAA Paper 2004-0554, Jan. 2004.

[42]Anderson, W. K. and Bonhaus, D. L., "An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids," *Computers and Fluids*, Vol. 23, No. 1, 1994, pp. 1–22.

[43]Anderson, W. K., Rausch, R. D., and Bonhaus, D. L., "Implicit/Multigrid Algorithm for Incompressible Turbulent Flows on Unstructured Grids," *Journal of Computational Fluids*, Vol. 128, 1996, pp. 391–408.

[44]Nielsen, E. J., *Aerodynamic Design Sensitivities on an Unstructured Mesh Using the Navier-Stokes Equations and a Discrete Adjoint Formulation*, Ph.D. thesis, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1998.

[45]Nielsen, E. J., Lu, J., Park, M. A., and Darmofal, D. L., "An Implicit, Exact Dual Adjoint Solution Method Implicit, Exact Dual Adjoint Solution Method for Turbulent Flows on Unstructured Grids," *Computers and Fluids*, Vol. 33, No. 9, 2004, pp. 1131–1155, See also AIAA Paper 2003–0272.

[46]Spalart, P. R. and Allmaras, S. R., "One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 92–0429, 1992.

[47]Taylor, A. C., Green, L. L., Newman, P. A., and Putko, M. M., "Some Advanced Concepts in Discrete Aerodynamic Sensitivity Analysis," AIAA Paper 2001–2529, 2001.

[48]Giles, M., Duta, M., Müller, J.-D., and Pierce, N., "Algorithm Developments for Discrete Adjoint Methods," *AIAA Journal*, Vol. 41, No. 2, 2003, pp. 198–205, See also AIAA Paper 2001–2596.

[49]Park, M. A., "Parallel, Distributed Memory Execution of Nearest-Neighbor Grid Adaptation Operations: refine," Class Project 21W.783, Massachusetts Institute of Technology, Cambridge, MA, 2005.

[50]Jones, W. T., "An Open Framework for Unstructured Grid Generation," AIAA Paper 2002-3192, June 2002.

[51]Jones, W. T., "GridEx – An Integrated Grid Generation Package for CFD," AIAA Paper 2003–4129, 2003.

[52]Haimes, R., "CAPRI: Computational Analysis PRogramming Interface," Massachusetts Institute of Technology, Cambridge, MA, 2002.

[53]Haimes, R., "Automatic Generation of CFD-Ready Surface Triangulations from CAD Geometry," AIAA Paper 99–0776, 1999.

[54]Haimes, R., "CAPRI: Computational Analysis PRogramming Interface," *Proceedings of the 6th International Conference on Numerical Grid Generation in Computational Field Simulations*, July 1998.

[55]Venkatakrishnan, V., "Convergence to Steady State Solutions of the Euler Equations on Unstructured Grids with Limiters," *Journal of Computational Physics*, Vol. 118, No. 1, 1995, pp. 120–130.

[56]Linde, T. and Roe, P. L., "On Multidimensional Positively Conservative High-Resoluion Schemes," *Barriers and Challenges in Computational Fluid Dynamics—1998*, edited by V. Verkatakrishnan, M. D. Salas, and S. R. Chakravarthy, Kluwer Academic Publ., 1998.

[57]Nielsen, E. J. and Anderson, W. K., "Recent Improvements in Aerodynamic Design Optimization on Unstructured Meshes," AIAA Paper 2001–0596, 2001.

[58]Brodersen, O. and Stürmer, "Drag Prediction of Engine-Airframe Interference Effects using Unstructured Navier-Stokes Calculations," AIAA Paper 2001-2414, June 2001.

[59]Roache, P. J., *Verification and Validation in Computational Science and Engineering*, Hermosa Publishers, 1998.

[60]Prizadeh, S., "Three-Dimensional Unstructured Viscous Grids by the Advancing Front Method," *AIAA Journal*, Vol. 34, No. 1, Jan. 1996, pp. 43–49.