

The Importance of Mesh Adaptation for Higher-Order Discretizations of Aerodynamic Flows

Masayuki Yano*, James M. Modisette[†] and David L. Darmofal[‡]

Aerospace Computational Design Laboratory, Massachusetts Institute of Technology

This work presents an adaptive framework that realizes the true potential of a higher-order discretization of the Reynolds-averaged Navier-Stokes (RANS) equations. The framework is based on an output-based error estimate and explicit degree of freedom control. Adaptation works toward the generation of meshes that equidistribute local errors and provide anisotropic resolution aligned with solution features in arbitrary orientations. Numerical experiments reveal that uniform refinement limits the performance of higher-order methods when applied to aerodynamic flows with low regularity. However, when combined with aggressive anisotropic refinement of singular features, higher-order methods can significantly improve computational affordability of RANS simulations in the engineering environment. The benefit of the higher spatial accuracy is exhibited for a wide range of applications, including subsonic, transonic, and supersonic flows. The higher-order meshes are generated using the elasticity and the cut-cell techniques, and the competitiveness of the cut-cell method in terms of accuracy per degree of freedom is demonstrated.

I. Introduction

For decades, the potential of higher-order discretizations as a means to improve computational efficiency for aerodynamics simulations has been discussed. Higher-order discretizations have become a widely accepted tool for applications with high fidelity demands, such as acoustic simulations and Large Eddy Simulations. However, the solvers used for steady state aerodynamics simulations in industry largely remain at most second order accurate. A common perception for the use of a higher-order discretizations in steady state aerodynamics simulations is that the benefit of high spatial accuracy can only be realized for high-cost, high-fidelity simulations that are unaffordable in the current engineering environment. This is particularly true for the Reynolds-averaged Navier-Stokes (RANS) simulations, in which the inaccuracies in the turbulence model is thought to diminish the value of such costly simulations. This work demonstrates that, with aggressive mesh refinements, the benefit of higher-order discretization can be achieved at much lower cost. The mesh adaptation extends the envelope of applications of higher-order methods to those encountered in the engineering environment by improving the computational affordability of RANS simulations.

Mesh adaptation that aggressively refines relevant features in the flow is the key to realizing the benefit of higher order discretizations at lower cost. When uniform mesh refinement is applied to aerodynamic flows, the accuracy is often limited by the solution regularity rather than the discretization order. However, when the mesh is aggressively graded toward singular features, the benefit of the high order discretization is recovered. The anisotropic refinement of boundary layers, shocks, and wakes are also the key to efficiently resolve these features. Constructing the optimal mesh that accounts for solution regularity and anisotropy is a formidable task even for a meshing expert, especially for high order discretizations, which are more sensitive to suboptimal mesh grading and for which experience is limited.

In order to meet the stringent requirements placed to generate optimal meshes for high order discretizations, this work relies on an autonomous error estimation and adaptation strategy. In particular, an output-based adaptation method based on the Dual-Weighted Residual^{1,2} (DWR) is used to estimate the error in an engineering quantity of interest, such as lift and drag, and to identify the regions to refine. Then,

*Doctoral candidate, AIAA student member, 77 Massachusetts Ave. 37-442, Cambridge, MA, 02139, myano@mit.edu

[†]Doctoral candidate, AIAA student member, 77 Massachusetts Ave. 37-435, Cambridge, MA, 02139, jmmodi@mit.edu

[‡]Professor, AIAA associate fellow, 77 Massachusetts Ave. 33-207, Cambridge, MA, 02139, darmofal@mit.edu

a Riemannian metric based adaptation strategy with an explicit degree of freedom (dof) control, inspired by the work on the continuous-discrete mesh duality,³ is used to drive toward the dof-optimal mesh. The previous applications of output-based adaptation include predictions of lift and drag for two-dimensional and three-dimensional flows,^{4–6} sonic boom propagations,^{7,8} and forces on re-entry vehicles.⁹ The information flow of the adaptive algorithm used in this work is summarized in Figure 1, and the key components of the algorithm are described in Section II and III.

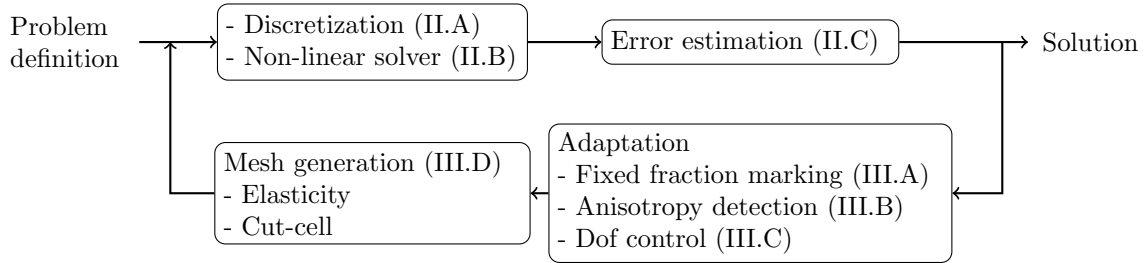


Figure 1. The information flow for the adaptive algorithm. The numbers in parenthesis correspond to the section numbers.

Section IV details the key findings and demonstrates the capability of the adaptive algorithm. First, deficiency of uniform refinement applied to higher-order discretizations is quantified for inviscid and viscous flows, and the ability of adaptive refinement to recover optimal convergence rate is verified. Then, the high-order discretization, when combined with adaptation, is shown to be superior to second-order discretization for subsonic, transonic, and supersonic RANS problems even at the drag level of as high as 10 counts—the error level at which the discretization error is deemed more dominant than modeling error. The competitiveness of the cut-cell meshes in terms of solution efficiency, accuracy per degree of freedom, is also demonstrated. Finally, the ability of the fixed-dof adaptation to efficiently perform a design parameter sweep is demonstrated for a transonic, high-lift airfoil.

II. Solution Strategy

II.A. Discretization

All flow equations considered in this work are steady-state conservation laws of the form,

$$\nabla \cdot \mathcal{F}_i(u) - \nabla \cdot \mathcal{F}_v(u, \nabla u) = S(u, \nabla u),$$

where u is the state, \mathcal{F}_i is the inviscid flux, \mathcal{F}_v is the viscous flux, and S is the source term characterizing a particular governing equation. The conservation law is discretized using a high-order discontinuous Galerkin (DG) finite element method, resulting in the weak form : Find $u_{h,p} \in V_{h,p}$ such that

$$\mathcal{R}_{h,p}(u_{h,p}, v_{h,p}) = 0 \quad \forall v_{h,p} \in V_{h,p}, \quad (1)$$

where $V_{h,p}$ is the space of discontinuous, p -th order piecewise polynomial functions. For elliptic problems with a smooth solution, the solution $u_{h,p}$ asymptotically converges at the rate of h^{p+1} measured in the L^2 norm, i.e. the discretization is $p+1$ order accurate in space. The inviscid flux of the Navier-Stokes equation is discretized using Roe’s approximate Riemann solver,¹⁰ and the viscous flux uses the second discretization of Bassi and Rebay.¹¹ This work uses the RANS equations with the Spalart-Allmaras (SA) turbulence model¹² in the fully turbulent mode. In order to improve the robustness of the solver, the modifications to the original SA model proposed by Oliver and Darmofal^{13,14} are incorporated. Multiple references have previously applied the DG discretizations to the RANS equations.^{15–20}

Shock capturing is performed using the PDE-based artificial viscosity model from Barter and Darmofal.²¹ In this model, a shock indicator that measures the local regularity of the solution is used as the forcing term of an elliptic PDE, which in turn generates a smooth artificial viscosity field. The artificial viscosity PDE, which augment the original conservation law, is given by

$$\frac{\partial \epsilon}{\partial t} = \frac{\partial}{\partial x_i} \left(\frac{C_2}{\tau} (\mathcal{M}^{-1})_{ij} \frac{\partial \epsilon}{\partial x_j} \right) + \frac{1}{\tau} \left[\frac{\bar{h}}{p} \lambda_{\max}(u) S_K(u) - \epsilon \right]$$

where ϵ is the artificial viscosity, $\tau = h_{\min}/(C_1 p \lambda_{\max}(u))$ is the time scale based on the maximum wave speed and the element size, \mathcal{M} is the smooth Riemannian metric tensor field, $\bar{h} = (\det(\mathcal{M}))^{-1/(2d)}$ is the average length scale based on the element volume, and S_K is the shock indicator based on the jump in a scalar quantity across element face. Unlike Barter's original equation that used axis aligned bounding boxes to measure the local element sizes, a Riemannian metric tensor is used in this work to measure the local length scale for the PDE. The new formulation provides consistent propagation of artificial viscosity independent of the coordinate system and enables sharper shock capturing on highly anisotropic elements with arbitrary orientations.

II.B. Non-linear Solver

Upon selecting suitable basis functions for the approximation space $V_{h,p}$, the problem of solving Eq. (1) becomes a discrete, root-finding problem. The solution is obtained using a non-linear solver based on pseudo-time continuation where the unsteady terms of the governing equations are retained. Solving the system of equations in pseudo-time improves the robustness of the solver, particularly through initial transients. A first-order backward Euler method is used for time integration. Given a discrete solution, U^n , the solution after one time step, U^{n+1} , is given by solving

$$R_t(U^{n+1}) \equiv \frac{1}{\Delta t} M(U^{n+1} - U^n) + R_s(U^{n+1}) = 0, \quad (2)$$

where $R_t(U)$ is the unsteady residual and $R_s(U)$ is the spatial residual. The m -th entry of $R_s(U)$ is the residual evaluated against the m -th basis function, i.e. $[R_s(U)]_m = \mathcal{R}_{h,p}(u_{h,p}, \phi^{(m)})$. A single step of Newton's method is used to approximately solve (2) at each time step,²²

$$U^{n+1} - U^n \approx \Delta U \equiv - \left(\frac{1}{\Delta t} M + \frac{\partial R_s}{\partial U} \Big|_{U^n} \right)^{-1} R_s(U^n), \quad (3)$$

where the current time step, Δt , is based on a CFL number. In this work local time stepping is used such that the time step on each element is

$$\Delta t_\kappa = CFL \frac{h_\kappa}{\lambda_\kappa},$$

where h_κ is a measure of element size and λ_κ is the maximum convective wave speed over the element. Computation of the state update, ΔU , requires the solution of a linear system. The linear system in this work is solved with restarted GMRES.²³ In order to improve the convergence of the GMRES algorithm, the linear system is preconditioned with an in-place block-ILU(0) factorization²⁴ with minimum discarded fill ordering and a coarse $p = 0$ multigrid correction.²⁵

The solution process is advanced in time until $\|R_s(U^n)\|_2$ is less than a specified tolerance. In the solution procedure, the CFL number is updated at each time step based on a physicality check and a line search over the unsteady residual, $R_t(U)$. The line search is included to increase the reliability of the solution procedure by requiring a decrease of the unsteady residual with each solution update as

$$\begin{aligned} \eta &= 1, \quad \tilde{U} = U^n + \eta \Delta U \\ \text{while } &\left(\|R_t(\tilde{U})\| > \|R_t(U^n)\| \ \&\& \ (\eta > \eta_{\min}) \right) \\ \eta &\leftarrow \frac{\eta}{2}, \quad \tilde{U} = U^n + \eta \Delta U \end{aligned}$$

Physicality Check and Line Search	CFL Change	Solution Update
$(\Delta\rho, \Delta\rho e < 10\%) \ \&\& \ (\eta = 1)$	$\text{CFL} \leftarrow 2 \cdot \text{CFL}$	full update
$(\Delta\rho, \Delta\rho e > 100\%) \ \ (\eta < \eta_{\min})$	$\text{CFL} \leftarrow \text{CFL}/10$	no update
otherwise	CFL unchanged	partial update

Table 1. Summary of solution update limiting for the non-linear solver.

The physicality check requires that both the density and internal energy, $\rho e = \rho E - \frac{1}{2}\rho(u^2 + v^2)$, are limited to changes of less than 10%. Table 1 summarizes the limiting applied to the CFL number.

II.C. Error Estimation

Once the solution to Eq. (1) is obtained, the quality of the solution is assessed through the error estimation process. The dual-weighted residual (DWR) method^{26,27} is used to estimate the error in a scalar output, $\mathcal{J}(u)$, and to localize the error to drive adaptation. In the DWR framework, the residual of the flow equation (primal equation) is weighted by the error in the dual solution, which measures the sensitivity of the output to local perturbations. The perturbations, in this case, are the discretization error of the primal problem. The DWR method enables estimation of the error in an output quantity of engineering interest, such as lift and drag, and identification of important regions for computing the output accurately. A review of recent developments in DWR error estimators for aerospace applications is provided by Fidkowski and Darmofal and the references therein.²⁸ The local error indicator used in this work is of the form

$$\eta_K = \frac{1}{2} \left[\left| \mathcal{R}_{h,p}(u_{h,p}, \tilde{\psi}_{h,p+1}|_K) \right| + \left| \mathcal{R}_{h,p}^\psi[\tilde{u}_{h,p+1}](\tilde{u}_{h,p+1}|_K, \psi_{h,p}) \right| \right] \quad (4)$$

where $u_{h,p} \in V_{h,p}$ is the solution satisfying the primal flow equation, Eq. (1), and $\psi_{h,p} \in V_{h,p}$ is the adjoint solution satisfying the adjoint equation

$$\mathcal{R}_{h,p}^\psi[u_{h,p}](v_{h,p}, \psi_{h,p}) \equiv \mathcal{J}'_{h,p}[u_{h,p}](v_{h,p}) - \mathcal{R}'_{h,p}[u_{h,p}](v_{h,p}, \psi_{h,p}) = 0 \quad \forall v_{h,p} \in V_{h,p}, \quad (5)$$

where $\mathcal{J}'[u_{h,p}](v_{h,p})$ and $\mathcal{R}'[u_{h,p}](v_{h,p}, \psi_{h,p})$ are the Fréchet derivative of the output functional and the residual semilinear form, respectively, evaluated about $u_{h,p}$ in the direction of $v_{h,p}$. The truth surrogate solution $\tilde{u}_{h,p+1}$ is sought from the $p+1$ order piecewise-polynomial space, which is obtained by solving Eq. (1) approximately on $V_{h,p+1}$ using 10 iterations of Newton smoothing. The authors have found that a simpler block smoothing scheme¹³ sometimes results in an unreliable error estimation for problems with shocks and separation. As the objective was to enable robust and automated adaptation, the additional cost of Newton smoothing was deemed justifiable. The dual surrogate solution, $\psi_{h,p+1} \in V_{h,p+1}$, is obtained by solving Eq. (5) linearized about $\tilde{u}_{h,p+1}$ exactly. The error in the output of interest is estimated by summing the local error contributions, i.e.

$$\mathcal{J}(u) - \mathcal{J}_{h,p}(u_{h,p}) \approx \sum_{K \in \mathcal{T}_h} \eta_K.$$

III. Adaptation Strategy

The objective of mesh adaptation is to generate a mesh that realizes the smallest output error for a given cost. In this work, the output error is estimated using the DWR framework discussed in Section II.C and the degrees of freedom is used as a cost metric. In order to efficiently resolve features such as shocks, wakes, and boundary layers with arbitrary orientations, Riemannian metric based anisotropic adaptation is used in this work.

III.A. Fixed-Fraction Marking

In order to drive adaptation based on the local error estimate Eq. (4), a fixed-fraction adaptation strategy is adopted to the context of unstructured, metric-based adaptation. The objective of adaptation is to create a metric request field, \mathcal{M}_r , that equidistribute the elemental error indicator, η_K , throughout the domain. To this end, the top f_r fraction of the elements with the largest error are marked for refinements and the top f_c fraction of the elements with the smallest error are marked for coarsening. The requested element size, A_r , is specified by

$$A_r = \alpha_r A_c$$

where A_c is the current element size, and α_r is the refinement rate which is set based on whether the element is marked for refinement, coarsening, or no change. Note that the fixed-fraction adaptation strategy is only used to determine the area of the element; the shape of element is determined by the anisotropy detection mechanism described in Section III.B. For this work, the parameters are set to $f_r = f_c = 0.2$ and $\alpha_r = 1/4$ and 2 for refinement and coarsening, respectively.

The distinguishing feature of the fixed-fraction marking strategy used in this work compared to traditional fixed-fraction adaptation based on hierarchical subdivision of elements is that the refinement request is not

directly connected to an increase in the degrees of freedom, i.e. the cost. As described in Section III.C, the final requested metric is scaled such that the discrete mesh would have a desired degrees of freedom. Thus, the fixed-fraction strategy should be thought of as a means to redistribute element sizes and to equidistribute local errors. The continuous scaling of degrees of freedom is enabled because the metric-based adaptation permits continuous variation of element size; this is contrary to the traditional hierarchical fixed-fraction method, in which the local subdivisions only permit discrete change in the element size.

III.B. Anisotropy Detection

In order to efficiently resolve shocks, boundary layers, and wakes encountered in the aerodynamics applications, the element orientation and stretching must be adapted to the anisotropic features in the flow. The anisotropy detection used in this work is based on the work by Venditti and Darmofal,²⁹ which was extended to higher-order methods by Fidkowski and Darmofal.³⁰ The framework attempts to minimize the interpolation error of the solution by first aligning the dominant principal direction of the Riemannian metric tensor with the direction of the maximum $p + 1$ derivative of the Mach number, $M_{\max}^{(p+1)}$. Then, the length scale in the second principal direction is selected to equidistribute the interpolation error in the principal directions. Assuming the Mach number converges at the rate of r , the principal lengths, $\{h_1, h_2\}$, implied by the anisotropy request metric tensor, \mathcal{M}_{ani} , should have the property

$$\left(\frac{h_2(\mathcal{M}_{\text{ani}})}{h_1(\mathcal{M}_{\text{ani}})} \right)^r = \frac{M_{\max}^{(p+1)}}{M_{\perp}^{(p+1)}}$$

where $M_{\perp}^{(p+1)}$ is the derivative of the Mach number in the direction orthogonal to $M_{\max}^{(p+1)}$. Without loss of generality, \mathcal{M}_{ani} is specified to have a determinant of unity. The choice of the Mach number as the scalar quantity representing the solution behavior follows from the previous works.^{29,31} While the choice is somewhat arbitrary, the authors have found it works well for the flows with shocks, wakes, and boundary layers in practice. The convergence rate, r , is nominally set to $p + 1$; however, the rate is reduced to $r = 1$ when the shock indicator is on, and the derivative quantities used in the formulation are replaced by the first derivatives, $M_{\max}^{(1)}$ and $M_{\perp}^{(1)}$. The $p + 1$ derivative of the Mach number is obtained from the truth surrogate solution of the DWR; this is another reason for performing the 10 Newton smoothing to obtain a robust $p + 1$ approximate of the true solution.

III.C. Metric Request Construction and Explicit Degree of Freedom Control

Using the area request A_r and the anisotropy metric \mathcal{M}_{ani} , the anisotropic metric request \mathcal{M}_r is constructed as

$$\mathcal{M}_r = A_r^{-2/\text{dim}} \mathcal{M}_{\text{ani}}.$$

However, even with the more robust truth surrogate reconstruction, both the error estimate and the anisotropy detection sometimes suffer from the noises in the surrogate solutions, particularly on coarse meshes encountered in the earlier stages of adaptation. In order to remedy the problem, the changes in the element size and shape are limited from one adaptation iteration to another. More specifically, the requested metric tensor is limited such that length change in any direction is limited by a given factor, i.e.

$$\frac{\sqrt{e^T \mathcal{M}_{\tilde{r}} e}}{\sqrt{e^T \mathcal{M}_c e}} \in [\Delta_r, \Delta_c] \quad \forall e \in \mathbb{R}^{\text{dim}}$$

where $\mathcal{M}_{\tilde{r}}$ is the limited requested metric tensor and the admissible changes for refinement and coarsening are set to $\Delta_r = 1/4$ and $\Delta_c = 2$, respectively.

Finally, the requested Riemannian metric is scaled to achieve the desired degrees of freedom, which is a measure for the cost of obtaining a solution. In the view of the discrete-continuous mesh duality proposed by Loseille and Alauzet,³ the degrees of freedom of a mesh conforming to a Riemannian metric field, $\mathcal{M}_{\tilde{r}}$, is approximated by

$$\text{dof}(\mathcal{M}_{\tilde{r}}) = \int_{\Omega} C_{p,K} \sqrt{\det(\mathcal{M}_{\tilde{r}})} dx$$

where $C_{p,K}$ is the constant depending on the solution order and the element shape. For example, for a p -th order polynomial simplex element in two dimension, $C_{p,K} = (2/\sqrt{3})(p+1)(p+2)$. The final dof-controlled metric field is

$$\mathcal{M}_{r,\text{final}} = \left(\frac{\text{dof}_{\text{target}}}{\text{dof}(\mathcal{M}_{\bar{r}})} \right)^{2/\text{dim}} \mathcal{M}_{\bar{r}}.$$

The metric scaling algorithm provides an explicit control of degrees of freedom independent of the number of adaptation cycles. The ability to perform multiple adaptation iterations at a fixed degrees of freedom serves two important features to our adaptation strategy.

As will be shown in the results section, achieving the optimality for a higher-order discretization requires significant grading of the mesh toward singularities; the dof-control mechanism allows the fixed-fraction marking strategy to produce this grading through a series of adaptation steps while maintaining degrees of freedom. This, in turn, enables generation of dof-“optimal” meshes, in which the local errors are equidistributed. Note that the “optimal” mesh generated in this manner does not necessarily minimize the error for a given degree of freedom, as the equidistribution of the local errors is a necessary but not sufficient condition for the true optimality. In particular, while the size distribution is optimized in the proposed procedure, the shapes of the elements may not be optimized when the Mach number does not serve as the best representation of the local primal and adjoint solution behaviors. Nevertheless, the Mach number based anisotropy detection appears to work well for aerodynamics applications, and the mesh obtained after several fixed-dof adaptation iterations is referred to as the “optimal” mesh in the result section. We also note that the fixed-dof adaptation results in a generation of a family of “optimal” meshes, all having similar metric distributions but slightly different triangulations. The family of meshes arise from the non-uniqueness of the meshes that realize a given metric field.³ In the result section, we

Second, the explicit degrees of freedom control is useful for performing parameter sweeps, e.g. constructing a lift curve slope or a drag polar. In this case, the explicit control allows the adaptation for a given parameter to start from an optimized mesh for the flow with a close parameter and maintain a particular degree of freedom count.

III.D. Higher-Order Mesh Generation

From the final dof-controlled metric field, an anisotropic mesh is generated using a metric-conforming mesh generator. All meshes used in this work were generated using BAMG,³² which generates linear anisotropic mesh (i.e. straight edged elements). However, the linear mesh is unsuitable for higher-order discretizations, and the mesh must be modified to capture higher-order geometry information of curved surfaces. Simply curving the elements on the boundary is not a valid option for highly anisotropic meshes used to resolve boundary layers, as the surface may push through the opposing face as shown in Figure 2. This work uses two methods to overcome the problem : elastically-curved meshes and cut-cell meshes.

III.D.1. Elastically-Curved Meshes

The first option is to globally curve a linear boundary conforming mesh by elasticity.^{13,33,34} In this method, the mesh is thought of as an elastic material, and the entire mesh is deformed as the boundary is moved to conform to the curved surface. This method requires a mesh generator that can generate a linear, anisotropic, boundary-conforming mesh, which is a difficult task in three dimensions. However, elastically-curved meshes offer benefits in terms of solution efficiency when compared to cut-cell meshes, as shown in the result section.

III.D.2. Cut-Cell Meshes

The second mesh generation technique used is a simplex cut-cell method first presented by Fidkowski and Darmofal³⁰ and extended by Modisette and Darmofal.³⁵ In this method, a linear, anisotropic background mesh is intersected with the geometry, producing “cut” elements on the geometry surface. The cut-cell method does not require a boundary conforming mesh, significantly simplifying the mesh generation process. This is particularly important for autonomous mesh generation for three-dimensional, complex geometry, where the true benefit of the adaptive method would be realized.

In terms of accuracy per degree of freedom, the cut-cell meshes are not as efficient as the elastically-curved meshes. As the cut-cell mesh uses linear background elements, degrees of freedom are wasted near




Figure 2. Diagram of the options for higher-order mesh generation.

curved geometry compared to elastically-curved meshes. However, this paper demonstrates that the solution efficiency loss in using the cut-cell meshes is relatively small, and the ease of mesh generation makes it a competitive method.

IV. Results

IV.A. Comparison of Uniform and Adaptive Refinement

IV.A.1. NACA0012 Subsonic Euler : $M_\infty = 0.5$, $\alpha = 2.0^\circ$

In order to demonstrate the importance of adaptive refinement for higher-order methods applied to aerodynamics applications, the error convergence behaviors of uniform and adaptive refinements are compared. The first problem considered is $M_\infty = 0.5$ Euler flow over a NACA0012 airfoil at $\alpha = 2^\circ$. To perform the comparison, adaptive refinement is first performed at fixed degrees of freedom of 2,500 and 5,000, generating “optimal” meshes for these two degrees of freedom. Then, for uniform refinement, each element is divided into four elements and the solution is obtained on the refined mesh having 10,000 and 20,000 degrees of freedom, respectively. The adaptive refinement results are obtained by continuing the adaptation procedure at 10,000 and 20,000 degrees of freedom. The test is performed for $p = 1$ (2nd order accurate) and $p = 3$ (4th order accurate) polynomials to assess the impact of adaptation for different discretization orders.

The result of the comparison is shown in Figure 3(a). The figure shows that adaptation has a larger impact for the $p = 3$ discretization than for the $p = 1$ discretization. In fact, the $p = 3$ discretization is less efficient than the $p = 1$ discretization in the absence of adaptive refinement. The suboptimal convergence rate of uniform refinement is due to the presence of a corner singularity at the trailing edge, which results in the convergence rate being limited by the regularity of the solution rather than the interpolation order. However, with a proper mesh grading, the optimal convergence rate of $\mathcal{E} \sim (\text{dof})^p \sim h^{2p}$ for the output quantity can be recovered.

Figure 4 shows the error indicator distribution near the trailing edge singularity for the $p = 3$, dof = 20,000 meshes obtained after a uniform refinement and adaptive refinements. After a step of uniform refinement, the error contribution of the trailing edge element is several orders of magnitude larger than that for the other elements, indicating the mesh is suboptimal. The adaptive refinement targets the corner elements dominating the error, and produces a strongly graded mesh that nearly equidistribute the error. The diameter of the trailing edge element is less than $8 \times 10^{-5}c$ for the adapted mesh. In comparison, the $p = 1$ adaptation having the same number of elements produces trailing edge elements of $h \approx 5 \times 10^{-3}c$. Thus, the higher-order discretization requires a stronger grading toward the corner singularity to equidistribute the

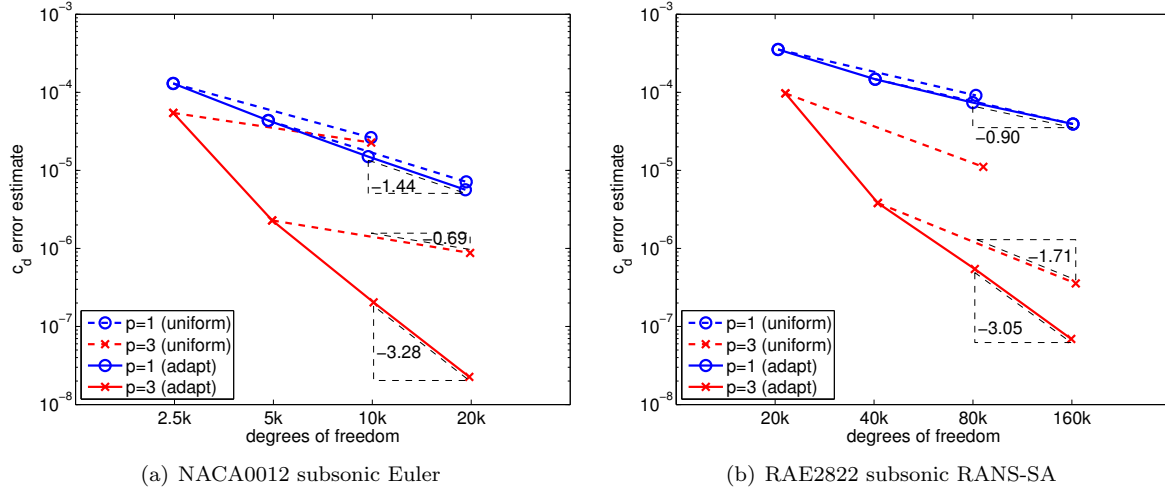


Figure 3. Comparison of the error convergence for uniform and adaptive refinements for the subsonic NACA0012 Euler flow ($M_\infty = 0.5$, $\alpha = 2.0^\circ$) and the subsonic RAE2822 RANS-SA flow ($M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$).

error. In addition, the higher-order discretization is more sensitive to suboptimal h distribution as the error scales as a higher power of h . Thus, h -adaptation is indispensable to achieve the full benefit of higher-order discretizations in the presence of corner singularities encountered in Euler flows.

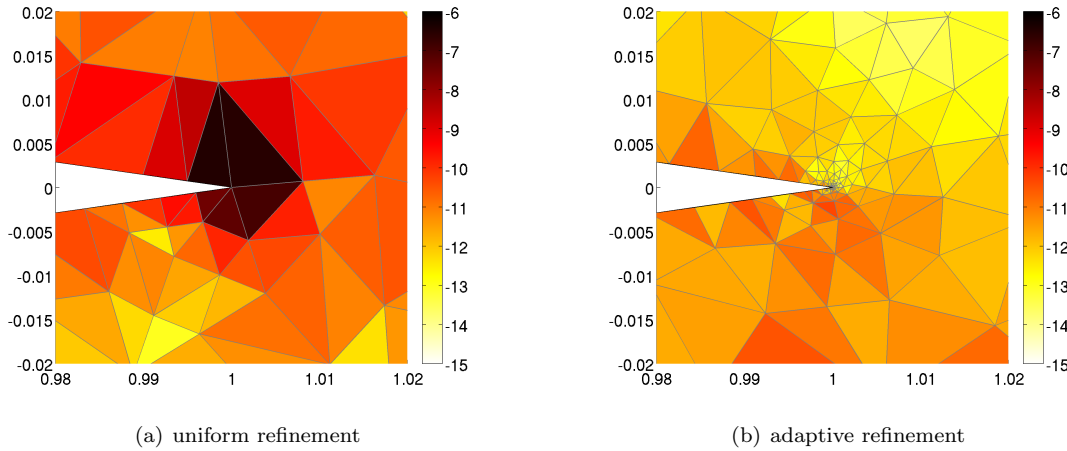


Figure 4. Comparison of the trailing edge mesh gradings of the $p = 3$, $dof = 20,000$ meshes obtained from uniform and adaptive refinements of the $p = 3$, $dof = 5,000$ optimized mesh for the subsonic NACA0012 Euler flow ($M_\infty = 0.5$, $\alpha = 2.0^\circ$). The color scale is in $\log_{10}(\mathcal{E})$.

IV.A.2. RAE2822 Subsonic RANS-SA : $M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$

The second problem considered is $M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$ turbulent flow over an RAE2822 airfoil at $\alpha = 2.31^\circ$. The RANS-SA equations are solved in the fully turbulent mode. Similar to the previous case, adaptation is first performed at 20,000 and 40,000 degrees of freedom to generate “optimal” meshes, and then uniform and adaptive refinements are started from those meshes.

The error convergence result is shown in Figure 3(b). Similar to the previous case, the adaptive refinement makes little difference for the $p = 1$ discretization. The convergence rate of the $p = 3$ discretization is limited by the solution regularity when the mesh is uniformly refined; however, with the adaptive refinement, the optimal convergence rate of $\mathcal{E} \sim (dof)^{2p/dim}$ for the output quantity is recovered. Also note that the “optimal” $p = 3$ mesh achieves drag error of approximately 1 count using just 20,000 degrees of freedom (i.e. 2,000 elements), whereas the optimal $p = 1$ mesh requires 80,000 degrees of freedom to achieve the same fidelity.

In order to understand the region limiting the performance of uniform refinement, the error indicator

distribution obtained after a step of uniform refinement from the 40,000 dof-optimized mesh is shown in Figure 5. Due to the singularity in the SA equation on the outer edge of the boundary layer,¹⁴ elements in the region are deemed to have high error. The adaptive refinement correctly identifies the region and makes the necessary adjustment to remove these high error elements.

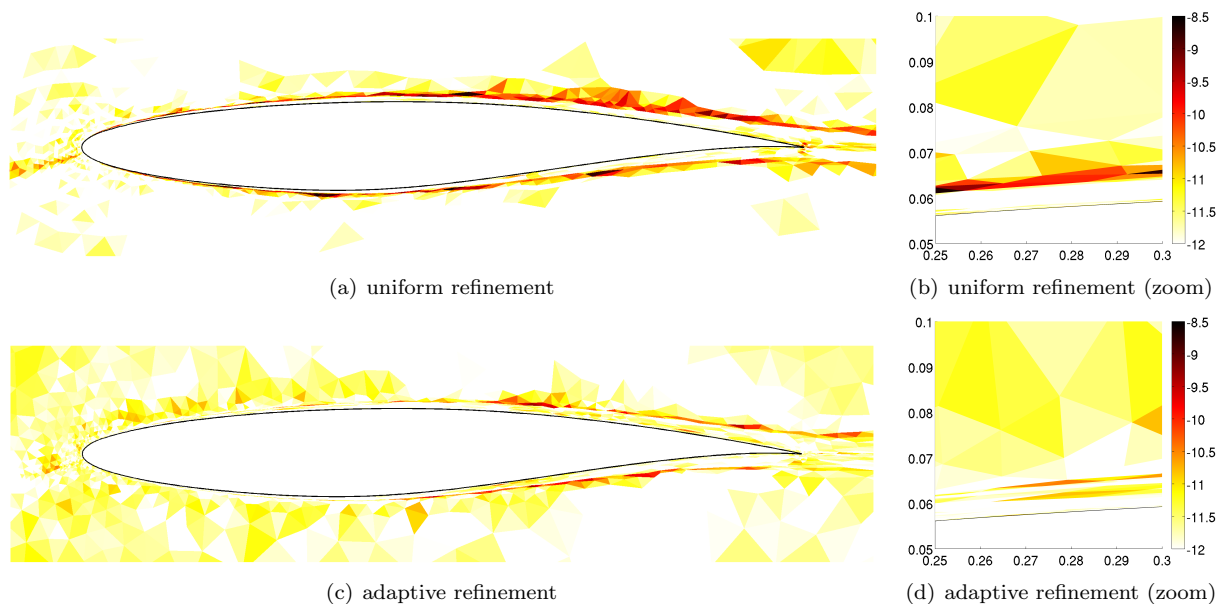


Figure 5. Comparison of the error indicator distributions of $p = 3$, dof = 160,000 meshes obtained from uniform and adaptive refinements of the $p = 3$, dof = 40,000 optimized mesh for the subsonic RAE2822 RANS-SA flow ($M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$). The color scale is in $\log_{10}(\mathcal{E})$.

IV.B. High-Order Discretization of RANS-SA Flows: Comparison of Boundary-Conforming and Cut-Cell Solution Efficiency

A consequence of using cut-cell meshes is a reduction in solution efficiency compared to elastically-curved meshes. As Figure 6 demonstrates, for an anisotropic two-dimensional layer on curved boundary, the elastically-curved mesh requires only two elements to resolve a length of δ . On the other hand, the cut-cell mesh is limited by the linear background elements, and three to four elements are required to resolve the same length. An increases in the degrees of freedom leads to the decrease in the local geometry curvature relative to the $p + 1$ derivative of the solution. Thus, asymptotically, the solution efficiency losses of the cut-cell meshes vanishes.




Figure 6. Diagrams of elastically-curved and cut-cell meshes required to resolve an anisotropic layer of thickness δ .

IV.B.1. *RAE2822 Subsonic RANS-SA* : $M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$

To demonstrate the difference in solution efficiency between cut-cell and elastically-curved meshes, the same RANS-SA flow over an RAE2822 airfoil as section IV.A.2 is considered ($M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$). Figure 8 shows the convergence in the c_d error estimate for “optimal” cut-cell and elastically-curved meshes. The largest solution efficiency gap exists for the $p = 3$, dof = 20,000 mesh. However, there are only 2,000 elements in these meshes with high grading toward the boundary layers, so the accuracy gap is understandable. It is important to note that the cut-cell meshes achieve a similar rate of convergence, and the gap in solution efficiency decreases as the total number of degrees of freedom increases.

In terms of the solution fidelity, the higher-order discretizations ($p = 2, 3$) are clearly superior to the second order discretization ($p = 1$) for high-fidelity simulation requiring the drag error of less than 1 count. In fact, on the highly graded meshes, the higher-order discretization is more efficient at an error level of as high as a few drag counts. Thus, if the singularity in the SA equation on the boundary layer edge is handled appropriately, the benefit of higher spacial accuracy to resolve the boundary layer can be realized at lower degrees of freedom.

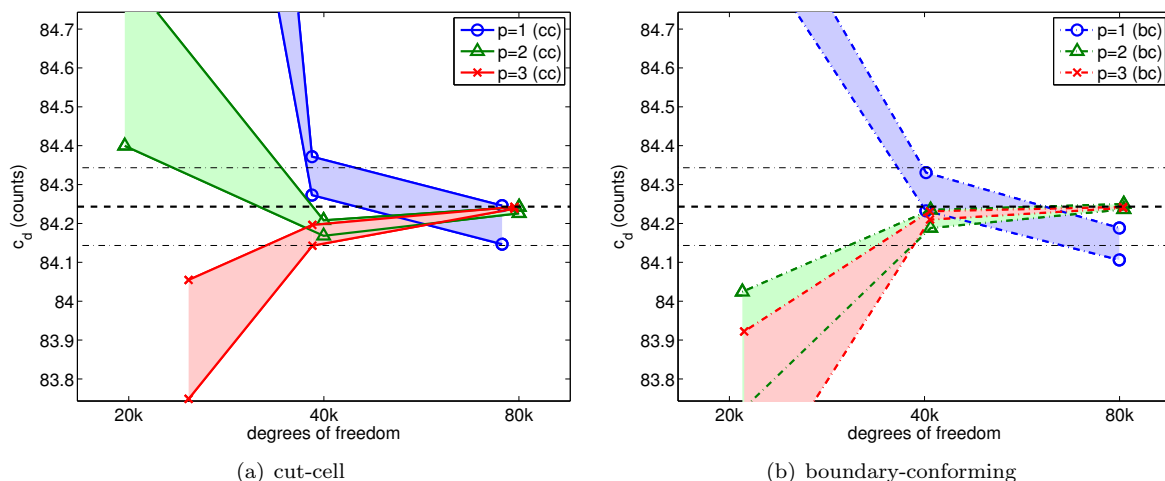


Figure 7. The drag coefficients for the subsonic RAE2822 RANS-SA flow ($M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$).




Figure 8. The c_d error estimate convergence for the subsonic RAE2822 RANS-SA flow ($M_\infty = 0.3$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$).

IV.B.2. *RAE2822 Transonic RANS-SA* : $M_\infty = 0.729$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$

This section demonstrates the benefit of higher-order discretizations even in the presence of shocks when the mesh is aggressively graded toward the singular features. First, transonic flow over an RAE2822 airfoil,

with the flow condition $M_\infty = 0.729$, $Re_c = 6.5 \times 10^6$, and $\alpha = 2.31^\circ$, is considered. The output of interest is drag.

The Mach number distribution and the drag-adapted mesh obtained using the $p = 3$, $dof = 40,000$ discretization are shown in Figure 9. The mesh is graded aggressively toward the airfoil surface, the boundary layer edge, and the shock. The refinement in the sonic pocket indicates the presence of the characteristics in the adjoint equation emanating from the shock root and reflecting off the sonic line. The Mach contour indicates that the combination of anisotropic grid refinement and the shock capturing algorithm enables sharp shock capturing.

Figure 11(a) shows the convergence of the drag error indicator obtained using the $p = 1, 2$, and 3 discretizations. Even at the c_d error level of as large as 10 drag counts, the $p = 2$ discretization is more efficient than the $p = 1$ discretization. For a higher-fidelity solution, the $p > 1$ discretizations are clearly more efficient. This is due to the significant reduction in the degrees of freedom required to capture the boundary layer for a given tolerance when a higher-order approximation is employed. Thus, even though the higher-order method does not improve the efficiency of resolving the shock, overall the method is more efficient for transonic turbulent problems with thin boundary layers given the mesh is appropriately graded. Note, similar to the subsonic RAE2822 case, the cut-cell meshes achieve a similar level of solution efficiency as the elastically-curved meshes.

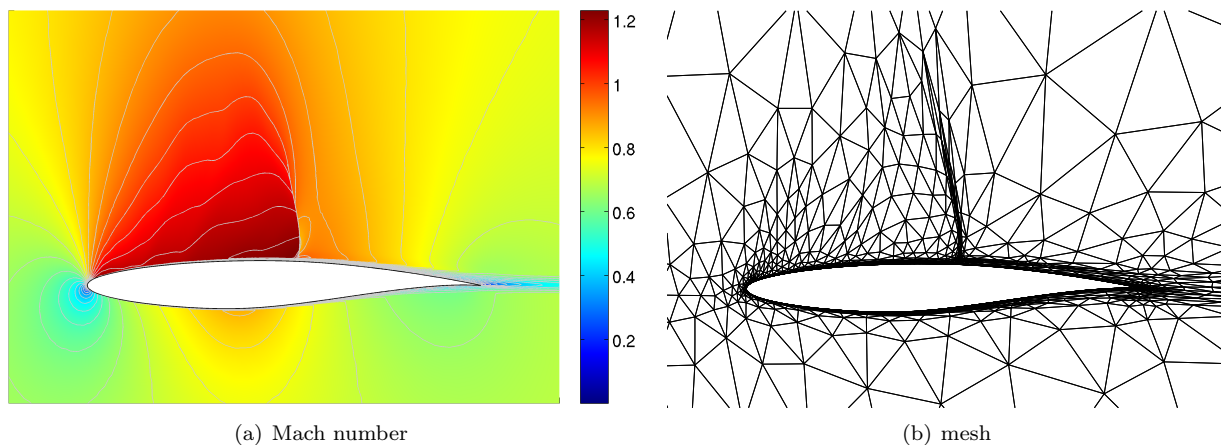


Figure 9. The Mach number distribution and the mesh for the transonic RAE2822 RANS-SA flow ($M_\infty = 0.729$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$) obtained using a $p = 3$, $dof = 40,000$ discretization. The Mach contour lines are in 0.05 increments.

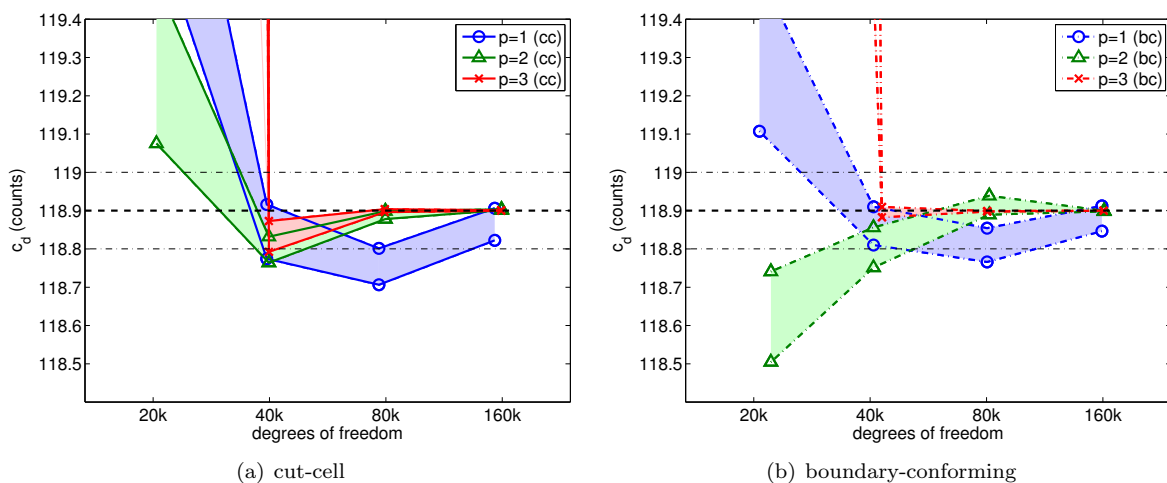


Figure 10. The drag coefficients for the transonic RAE2822 RANS-SA flow ($M_\infty = 0.729$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$).

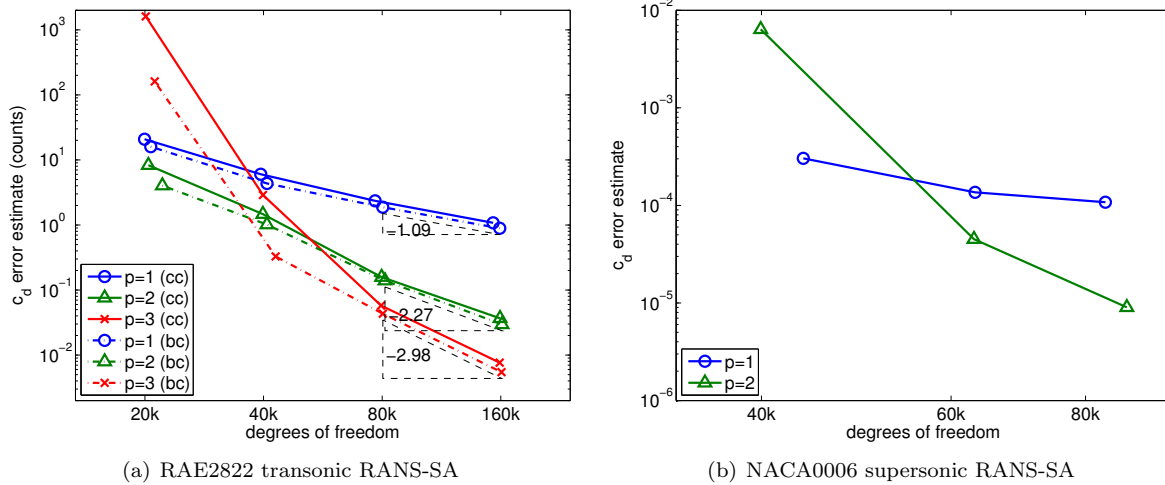


Figure 11. The c_d error estimate convergence for the transonic RAE2822 RANS-SA flow ($M_\infty = 0.729$, $Re_c = 6.5 \times 10^6$, $\alpha = 2.31^\circ$) and the supersonic NACA0006 RANS-SA flow ($M_\infty = 2.0$, $Re_c = 2.0 \times 10^7$, $\alpha = 2.5^\circ$).

IV.B.3. NACA0006 Supersonic RANS-SA : $M_\infty = 2.0$, $Re_c = 2.0 \times 10^7$, $\alpha = 2.5^\circ$

The second shock problem considered is a supersonic flow over a NACA0006 airfoil with the flow condition $M_\infty = 2.0$, $Re_c = 2.0 \times 10^7$, and $\alpha = 2.5^\circ$. The output of interest is drag. The Mach number distribution and the mesh obtained for the $p = 2$ discretization having 80,000 degrees of freedom is shown in Figure 12. The aggressive refinement toward the singularities is evident from the mesh. The bow shock inside the adjoint Mach cone emanating from the trailing edge is captured sharply using highly anisotropic elements. The mesh resolution drops quickly outside of the adjoint Mach cone, as the solution outside of the cone is irrelevant for the calculation of the drag. For the same reason, the shock originating from the trailing edge is not targeted by the adaptation.

Figure 11(b) shows the convergence of the drag error estimate for the problem. Due to the irregularity of the dominant feature in the flow, it is difficult to realize the benefit of $p = 2$ discretization using only 40,000 degrees of freedom (i.e. 6700 elements) even with the adaptive algorithm. However, for a higher-fidelity simulation requiring the tolerance of less than 1 drag count, the $p = 2$ discretization is more efficient than the $p = 1$ discretization due to the significant saving in resolving the boundary layer.

IV.C. Parameter Sweep using Fixed-Dof Adaptation

IV.C.1. Three-element MDA high-lift airfoil : $M_\infty = 0.2$, $Re_c = 9 \times 10^6$

To demonstrate the ability of the fixed-dof adaptation to efficiently perform a parameter sweep, the adaptation strategy is used to construct the lift curve for the McDonnell Douglas Aerospace (MDA) three-element airfoil (30P-30N).³⁶ For this high-lift airfoil test case, the freestream Mach number and Reynolds number are set to $M_\infty = 2.0$ and $Re_c = 9 \times 10^6$, respectively, and the angle of attack is varied from 0.0° to 24.5° . All results are obtained using the RANS-SA equations in fully turbulent mode, discretized by $p = 2$ polynomials at 90,000 degrees of freedom.

Figure 13(a) shows the lift curves obtained using a fixed grid and adaptive grids. The fixed grid used for comparison is optimized for $\alpha = 8.1^\circ$. The lift curve shows that the fixed grid closely matches the adaptive result for $0^\circ < \alpha < 20^\circ$; however, for $\alpha > 20^\circ$, the lift calculation on the fixed grid becomes unreliable and the c_l is significantly underestimated. Figure 13(b) shows that the error indicator correctly identifies the lack of confidence in the solution for the high angle of attack cases, producing the c_l error estimate on the order of 10. The local error estimate indicates that the elements on the upper surface of the slat dominates the fixed grid c_d errors for angles of attack above 20° . On the other hand, when adaptive refinement is performed for each angle of attack, the c_l error estimate remains less than 0.01 for the entire range of angles of attack considered despite the adaptive grid using the same degrees of freedom as the fixed grid.

Figure 14 shows the Mach number distribution obtained for the $\alpha = 23.28^\circ$ flow on the fixed grid optimized for the $\alpha = 8.1^\circ$ and the grid adapted for the $\alpha = 23.28^\circ$ flow. The Mach number distribution

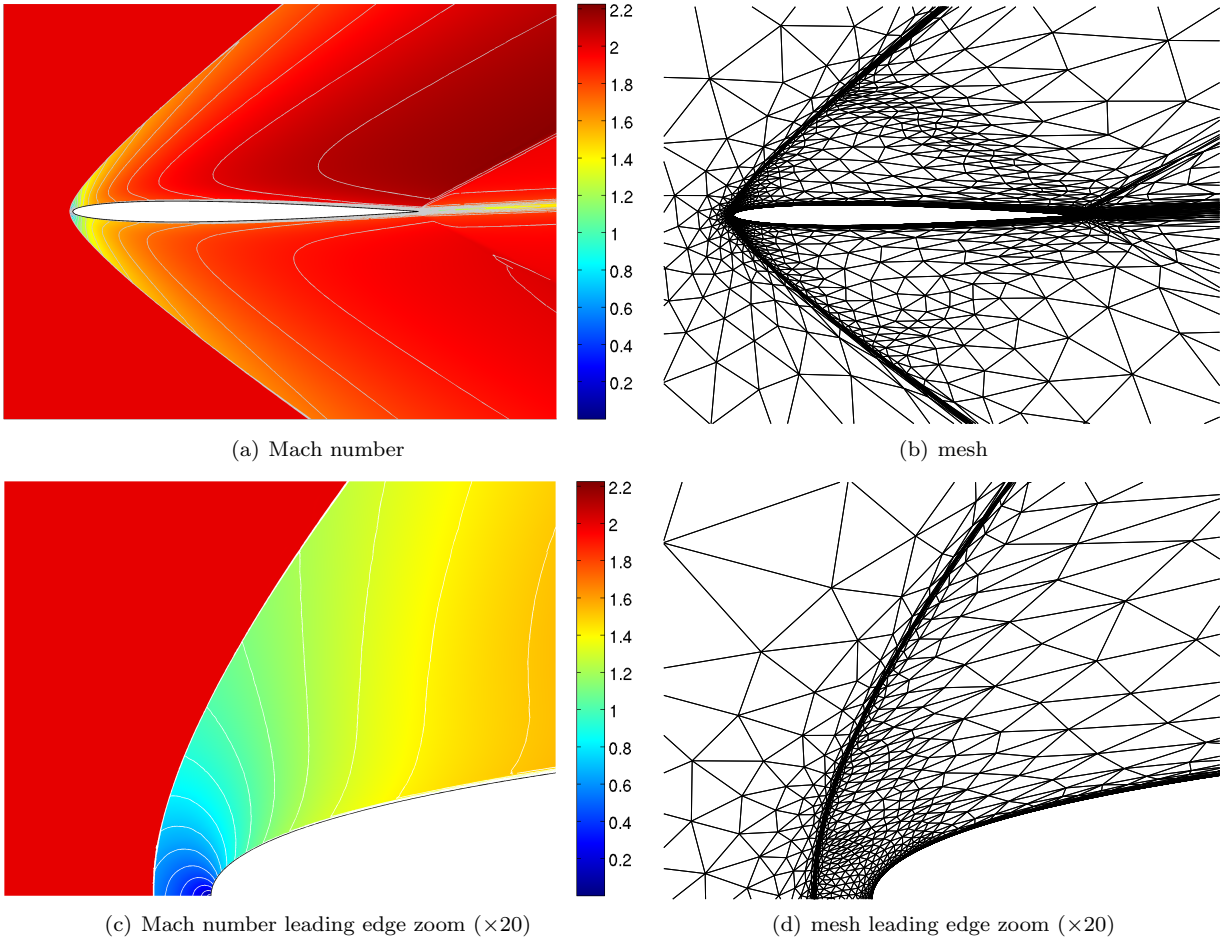


Figure 12. The Mach number distribution and the mesh for the supersonic NACA0006 RANS-SA flow ($M_\infty = 2.0$, $Re_c = 2.0 \times 10^7$, $\alpha = 2.5^\circ$) obtained using a $p = 2$, dof = 80,000 discretization. The Mach contour lines are in 0.1 increments.

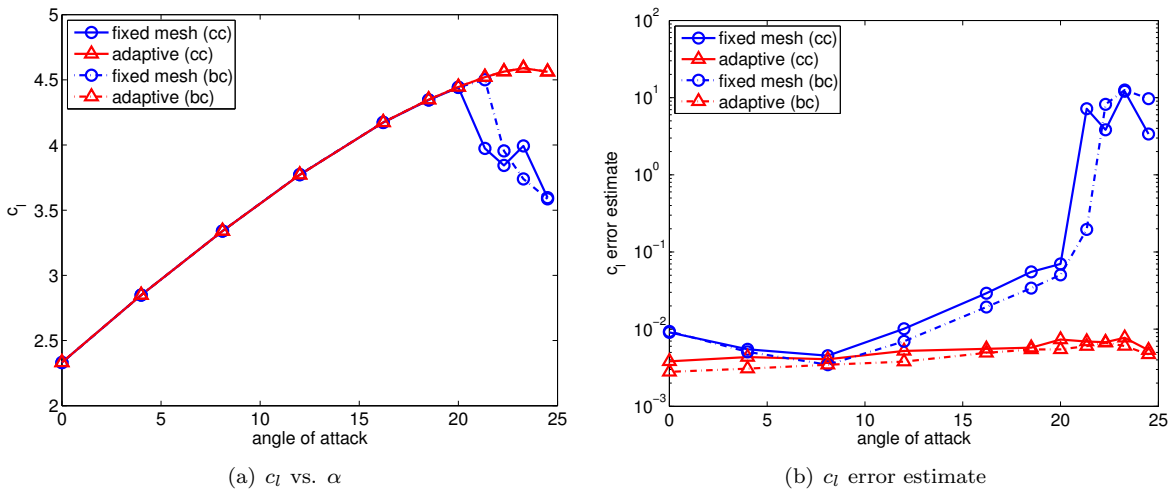


Figure 13. The lift curve and the c_l error obtained using the fixed mesh and adaptive meshes for the three-element MDA airfoil.

indicates that the fixed grid lacks resolution on the front side of the leading edge slat, causing the extra numerical dissipation to induce separation on the upper surface of the slat. The lack of acceleration is evident from the absence of the sonic pocket on the slat. The adjoint captures the impact of the region to the rest of the flow, which leads to the local error estimate being high in the region. With proper mesh resolution, the adapted grid eliminates the numerical dissipation induced separation. The shock in the front side of the slat is captured sharply by combination the shock capturing mechanism and aggressive refinement.

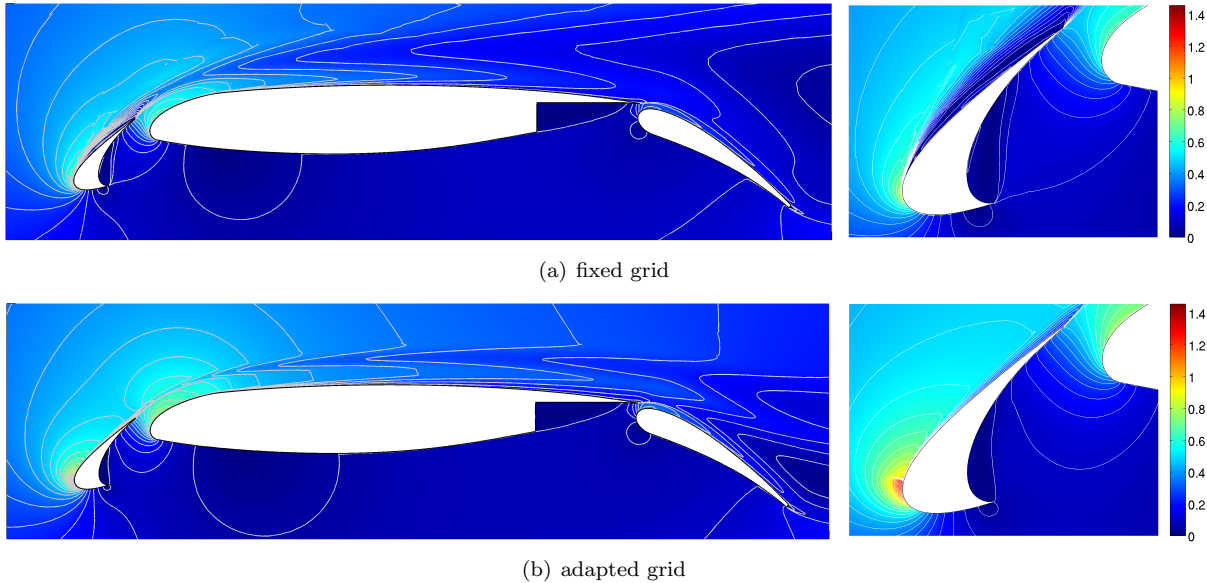


Figure 14. The Mach number distribution for the three-element MDA airfoil at $\alpha = 23.28^\circ$ obtained on the 8.10° optimized mesh and the 23.28° optimized mesh. The Mach contour lines are in 0.05 increments.

Figure 15 shows the initial mesh and the adapted meshes for selected angles of attack. The different flow features that the problem exhibits as the angle of attack varies are evident from the meshes. At lower angle of attack, the flow separate from the bottom side of the slat, and the wake must be captured to account for its influence on the main element. At $\alpha = 23.28$, capturing the acceleration on the upper side and the shock becomes important for accurate calculation of lift.

V. Conclusions

This paper presented an adaptation strategy that works toward generation of dof-optimal meshes using an output-based error indicator with explicit control of degrees of freedom. Using the framework, “optimal” meshes that can realize the benefit of a higher-order discretization at low degrees of freedom were generated. The key features of these meshes were strong grading toward the singularities and anisotropic resolutions for the boundary layers, wakes, and shocks. The numerical experiments demonstrated that uniform refinement is insufficient to attain the benefits of high-order discretizations. With a proper mesh selection, the higher-order methods are shown to be superior to lower-order methods for RANS-SA simulations of subsonic, transonic, and supersonic flows. In addition, the results show that the advantage of the higher-order discretization can be achieved at an error level of as high as 10 drag counts—the level at which the discretization error dominates the modeling error. Furthermore, the cut-cell meshes are shown to be competitive to the elastically-curved meshes in terms of solution efficiency, making the method an attractive choice for complex, three-dimensional geometries. Finally, the potential of using the fixed degrees of freedom adaptation for design parameter sweeps were demonstrated for high-lift, transonic flows.

Acknowledgments

The authors would like to thank Dr. Steven Allmaras for the insightful discussions throughout the development of this work. This work was supported by funding from The Boeing Company with technical monitor Dr. Mori Mani.




Figure 15. The initial and lift-adapted grids for the three-element MDA airfoil at selected angles of attack.

References

- ¹Rannacher, R., “Adaptive Galerkin finite element methods for partial differential equations,” *Journal of Computational and Applied Mathematics*, Vol. 128, 2001, pp. 205–233.
- ²Giles, 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.
- ³Loseille, A. and Alauzet, F., “Continuous Mesh Model and Well-Posed Continuous Interpolation Error Estimation,” INRIA RR-6846, 2009.
- ⁴Venditti, D. A. and Darmofal, D. L., “Grid adaptation for functional outputs: application to two-dimensional inviscid flows,” *Journal of Computational Physics*, Vol. 176, No. 1, 2002, pp. 40–69.
- ⁵Park, M. A., “Adjoint-based, Three-dimensional Error Predication and Grid Adaptation,” *AIAA Journal*, Vol. 42, No. 9, 2004, pp. 1854–1862.
- ⁶Fidkowski, K. J. and Darmofal, D. L., “Output-based adaptive meshing using triangular cut cells.” M.I.T. Aerospace Computational Design Laboratory Report. ACDL TR-06-2, 2006.
- ⁷Lee-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 AIAA-2005-4842, 2005.
- ⁸Jones, W. T., Nielsen, E. J., and Park, M. A., “Validation of 3D Adjoint Based Error Estimation and Mesh Adaptation for Sonic Boom Prediction,” AIAA 2006-1150, 2006.
- ⁹Nemec, M., Aftosmis, M. J., and Wintzer, M., “Adjoint-based adaptive mesh refinement for complex geometries,” AIAA 2008-725, 2008.
- ¹⁰Roe, P. L., “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes,” *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372.
- ¹¹Bassi, F. and Rebay, S., “GMRES discontinuous Galerkin solution of the compressible Navier-Stokes equations,” *Discontinuous Galerkin Methods: Theory, Computation and Applications*, edited by K. Cockburn and Shu, Springer, Berlin, 2000, pp. 197–208.
- ¹²Spalart, P. R. and Allmaras, S. R., “A one-equation turbulence model for aerodynamics flows,” AIAA 1992-0439, Jan. 1992.
- ¹³Oliver, T. A., *A Higher-Order, Adaptive, Discontinuous Galerkin Finite Element Method for the Reynolds-averaged Navier-Stokes Equations*, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2008.
- ¹⁴Oliver, T. and Darmofal, D., “Impact of turbulence model irregularity on high-order discretizations,” AIAA 2009-953, 2009.
- ¹⁵Bassi, F., Crivellini, A., Rebay, S., and Savini, M., “Discontinuous Galerkin solution of the Reynolds averaged Navier-Stokes and $k-\omega$ turbulence model equations,” *Computers & Fluids*, Vol. 34, May-June 2005, pp. 507–540.
- ¹⁶Bassi, F., Crivellini, A., Ghidoni, A., and Rebay, S., “High-order discontinuous Galerkin discretization of transonic turbulent flows,” AIAA 2009-180, 2009.
- ¹⁷Nguyen, N., Persson, P.-O., and Peraire, J., “RANS solutions using high order discontinuous Galerkin methods,” AIAA 2007-0914, 2007.
- ¹⁸Landmann, B., Kessler, M., Wagner, S., and Krämer, E., “A parallel, high-order discontinuous Galerkin code for laminar and turbulent flows,” *Computers & Fluids*, Vol. 37, No. 4, 2008, pp. 427 – 438.
- ¹⁹Burgess, N., Nastase, C., and Mavriplis, D., “Efficient solution techniques for discontinuous Galerkin discretizations of the Navier-Stokes equations on hybrid anisotropic meshes,” AIAA 2010-1448, 2010.
- ²⁰Hartmann, R. and Houston, P., “Error estimation and adaptive mesh refinement for aerodynamic flows,” *VKI LS 2010-01: 36th CFD/ADIGMA course on hp-adaptive and hp-multigrid methods, Oct. 26-30, 2009*, edited by H. Deconinck, Von Karman Institute for Fluid Dynamics, Rhode Saint Genèse, Belgium, 2009.
- ²¹Barter, G. E. and Darmofal, D. L., “Shock capturing with PDE-based artificial viscosity for DGFEM: Part I, Formulation,” *Journal of Computational Physics*, Vol. 229, No. 5, 2010, pp. 1810–1827.
- ²²Gropp, W., Keyes, D., McInnes, L. C., and Tidriri, M. D., “Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD,” *International Journal of High Performance Computing Applications*, Vol. 14, No. 2, 2000, pp. 102–136.
- ²³Saad, Y., *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 1996.
- ²⁴Diosady, L. and Darmofal, D., “Discontinuous Galerkin Solutions of the Navier-Stokes Equations Using Linear Multigrid Preconditioning,” AIAA 2007-3942, 2007.
- ²⁵Persson, P.-O. and Peraire, J., “Newton-GMRES Preconditioning for Discontinuous Galerkin discretizations of the Navier-Stokes Equations,” *SIAM Journal on Scientific Computing*, Vol. 30, No. 6, 2008, pp. 2709–2722.
- ²⁶Becker, R. and Rannacher, R., “A feed-back approach to error control in finite element methods: Basic analysis and examples,” *East-West Journal of Numerical Mathematics*, Vol. 4, 1996, pp. 237–264.
- ²⁷Becker, R. and Rannacher, R., “An optimal control approach to a posteriori error estimation in finite element methods,” *Acta Numerica*, edited by A. Iserles, Cambridge University Press, 2001.
- ²⁸Fidkowski, K. and Darmofal, D., “Output error estimation and adaptation in computational fluid dynamics: Overview and recent results,” AIAA 2009-1303, 2009.
- ²⁹Venditti, D. A. and Darmofal, D. L., “Anisotropic grid adaptation for functional outputs: Application to two-dimensional viscous flows,” *Journal of Computational Physics*, Vol. 187, No. 1, 2003, pp. 22–46.
- ³⁰Fidkowski, K. J. and Darmofal, D. L., “A triangular cut-cell adaptive method for higher-order discretizations of the compressible Navier-Stokes equations,” *Journal of Computational Physics*, Vol. 225, 2007, pp. 1653–1672.

³¹Fidkowski, K. J., *A Simplex Cut-Cell Adaptive Method for High-Order Discretizations of the Compressible Navier-Stokes Equations*, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2007.

³²Hecht, F., "BAMG: Bidimensional Anisotropic Mesh Generator," 1998, <http://www-rocq1.inria.fr/gamma/cdrom/www/bamg/eng.htm>.

³³Nielsen, E. J. and Anderson, W. K., "Recent improvements in aerodynamic design optimization on unstructured meshes," *AIAA Journal*, Vol. 40, No. 6, 2002, pp. 1155–1163.

³⁴Persson, P.-O. and Peraire, J., "Curved mesh generation and mesh refinement using Lagrangian solid mechanics," AIAA 2009-0949, 2009.

³⁵Modisette, J. M. and Darmofal, D. L., "Toward a Robust, Higher-Order Cut-Cell Method for Viscous Flows," AIAA 2010-721, 2010.

³⁶Klausmeyer, S. M. and Lin, J. C., "Comparative Results From a CFD Challenge Over a 2D Three-Element High-Lift Airfoil," NASA Technical Memorandum 112858, 1997.