

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

Masayuki Yano\*, James M. Modisette<sup>†</sup> and David L. Darmofal<sup>‡</sup>

*Aerospace Computational Design Laboratory, Massachusetts Institute of Technology*

This work presents an adaptive framework for a higher-order discretization of the Reynolds-averaged Navier-Stokes (RANS) equations. The adaptation strategy is based on an output-based error estimate and explicit control of the degrees of freedom. Adaptation iterates toward the generation of simplex meshes that equidistribute local errors throughout the domain and provide anisotropic resolution 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 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 simplex meshes are generated using the elasticity and the cut-cell techniques, and the competitiveness of the cut-cell method is demonstrated in terms of accuracy per degree of freedom.

## 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 discretizations used for steady state aerodynamic simulations in the industry largely remain second-order accurate. A common perception for the use of a higher-order discretization 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 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 mesh adaptation, the benefit of higher-order discretization can be achieved in low-cost, low-fidelity simulations in which discretization error dominates the modeling error. Mesh adaptation extends the envelope of the application of higher-order methods to problems encountered in the engineering environment by improving the computational affordability of RANS simulations.

Mesh adaptation that aggressively refines relevant flow features is key to realizing the benefit of higher-order discretizations at a 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 penalty of solution irregularities is reduced and the benefit of the high-order discretization is recovered. Anisotropic refinement of boundary layers, shocks, and wakes is also a key to efficiently resolve these features. Constructing the optimal mesh that accounts for solution regularity and anisotropy is a formidable task, especially for high-order discretizations, which are more sensitive to suboptimal mesh grading.

In order to meet the stringent requirements to generate optimal meshes for high-order discretizations, this work relies on autonomous error estimation and adaptation strategies. In particular, an output-based error estimation and adaptation method<sup>1,2</sup> is used to estimate the error in an engineering quantity of interest,

---

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

<sup>†</sup>Doctoral candidate, AIAA student member, 77 Massachusetts Ave. 37-435, Cambridge, MA, 02139, jmmodi@mit.edu

<sup>‡</sup>Professor, AIAA associate fellow, 77 Massachusetts Ave. 33-207, Cambridge, MA, 02139, darmofal@mit.edu

such as lift or drag, and to refine the regions contributing large errors. Then, a Riemannian metric based adaptation strategy with an explicit degree of freedom (dof) control, inspired by work on the continuous-discrete mesh duality,<sup>3</sup> drives toward the dof-optimal mesh—the mesh that minimizes the error for a given degree of freedom. 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.

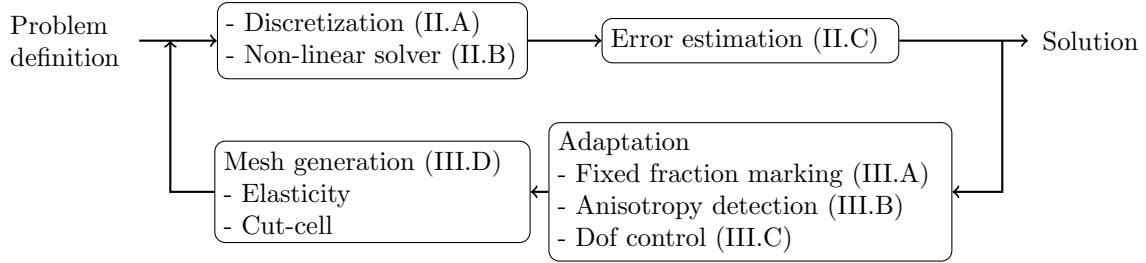


Figure 1. The information flow for the adaptive algorithm. The numbers in parentheses correspond to the section numbers.

Section IV details the key findings and demonstrates the capability of the adaptive algorithm. First, the 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 error estimate level of as high as 10 drag counts—the error level at which the discretization error is deemed more dominant than modeling error. Two high-order mesh generation strategies, based on elastically-curved elements and cut-cells, are comparable in terms of solution efficiency, i.e. accuracy per degree of freedom. Finally, the ability of 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 numerical solution  $u_{h,p}$  asymptotically converges at a rate of  $h^{p+1}$  measured in the  $L^2$  norm. The inviscid flux of the Navier-Stokes equation is discretized using Roe’s approximate Riemann solver,<sup>4</sup> and the viscous flux uses the second discretization of Bassi and Rebay.<sup>5</sup> This work uses the RANS equations with the Spalart-Allmaras (SA) turbulence model<sup>6</sup> 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<sup>7,8</sup> are incorporated.

Shock capturing is performed using the PDE-based artificial viscosity model from Barter and Darmofal.<sup>9</sup> 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 augments 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  $\epsilon$  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 area,  $d$  is the physical dimension, and  $S_K$  is the shock indicator based on the jump in the speed of sound across element faces. The artificial viscosity is added to the Navier-Stokes equations (excluding the SA equation) using a physical viscosity model.<sup>10</sup> 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,  $V_{h,p}$ , 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

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

where  $R_t(U)$  is the unsteady residual,  $M^{\text{tw}}$  is the time-weighted mass matrix, 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,  $\phi^{(m)}$ , i.e.  $[R_s(U)]_m = \mathcal{R}_{h,p}(u_{h,p}, \phi^{(m)})$ . This work uses local time stepping such that the time step and the elemental time-weighted matrix for  $K$  are given by

$$\Delta t_K = \text{CFL} \frac{h_K}{\lambda_K} \quad \text{and} \quad M_K^{\text{tw}} = \frac{1}{\Delta t_K} M_K,$$

where  $h_K$  is the minimum altitude of the element,  $\lambda_K$  is the maximum convective wave speed over the element, and  $M_K$  is the elemental mass matrix. In the pseudo-time continuation algorithm, the CFL number acts as the global continuation parameter. A single step of Newton’s method is used to approximately solve (2) at each time step such that<sup>11</sup>

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

Computation of the state update,  $\Delta U$ , requires the solution of a large linear system with a block-sparse structure. The linear system in this work is solved with restarted GMRES.<sup>12</sup> In order to improve the convergence of the GMRES algorithm, the linear system is preconditioned with an in-place block-ILU(0) factorization<sup>13</sup> with minimum discarded fill ordering and a coarse  $p = 0$  multigrid correction.<sup>14</sup>

The solution process is advanced in time until the two-norm of the spatial residual,  $\|R_s(U^n)\|_2$ , 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 incorporated to increase the reliability of the solution procedure by requiring a decrease of the unsteady residual with each solution update as

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

In other words, the step length is halved until the unsteady residual decreases after taking the step. This requirement aims to prevent the non-linear solver divergence due to the lack of temporal integration accuracy by explicitly controlling the unsteady residual. In addition, the physicality check requires that both the density,  $\rho$ , and internal energy,  $\rho e$ , are limited to changes of less than 10%. Table 1 summarizes the limiting applied to the CFL number.

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.

### II.C. Error Estimation

Once the solution to Eq. (1) is obtained, the quality of the solution is assessed through an error estimation process. The dual-weighted residual (DWR) method<sup>15, 16</sup> is used to estimate the error in a functional output,  $\mathcal{J}(u)$ , and to localize the error to each element. In the DWR framework, the residual of the flow equation (primal equation) is weighted by the error in the dual solution. The dual solution measures the sensitivity of the output to local perturbations. For error estimation, the perturbations are the discretization error of the primal problem. A review of recent developments in DWR error estimators for aerospace applications is provided by Fidkowski and Darmofal and the references therein.<sup>17</sup> The local error indicator used in this work is

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

where  $\psi_{h,p} \in V_{h,p}$  is the adjoint solution satisfying the adjoint equation

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

and  $\mathcal{J}'_{h,p}[u_{h,p}](v_{h,p})$  and  $\mathcal{R}'_{h,p}[u_{h,p}](v_{h,p}, \psi_{h,p})$  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 time steps of pseudo-time continuation, where each step involves single Newton iteration, Eq. (3). The authors have found that a simpler block smoothing scheme<sup>7</sup> 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 iterations was deemed justifiable. The dual truth surrogate solution,  $\tilde{\psi}_{h,p+1} \in V_{h,p+1}$ , is obtained by exactly solving Eq. (5) linearized about  $\tilde{u}_{h,p+1}$ . In evaluating Eq. (4), the residual is computed about  $p$  instead of  $p + 1$  such that the resulting estimate is both globally and locally convergent.<sup>18</sup> The error in the output of interest is estimated by summing the local error contributions,

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

Since  $\eta_K$  defined by Eq. (4) is strictly positive, the error estimate employed in this work does not permit inter-elemental cancellation of local errors and results in a conservative estimate of the output error.

## 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. Riemannian metric based anisotropic adaptation is used to efficiently resolve features such as shocks, wakes, and boundary layers with arbitrary orientations

### III.A. Fixed-Fraction Marking

To drive adaptation based on the local error estimate, Eq. (4), a fixed-fraction adaptation strategy is adopted in the context of unstructured, metric-based adaptation. The objective of adaptation is to create a metric request field,  $\mathcal{M}_r$ , that equidistributes the elemental error indicator,  $\eta_K$ , throughout the domain. To this end, the top  $f_r$  fraction of the elements with the largest error are marked for refinement and the bottom  $f_c$

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  $\alpha_r$  is the refinement rate which is 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  $f_r = f_c = 0.2$ ,  $\alpha_r = 1/4$  for refinement, and  $\alpha_r = 2$  for coarsening.

The distinguishing feature of the fixed-fraction marking strategy 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. As described in Section III.C, the final requested metric is scaled such that the discrete mesh would have the 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 the 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 changes in the element size.

### III.B. Anisotropy Detection

To efficiently resolve shocks, boundary layers, and wakes encountered in the aerodynamics applications, the element orientation and stretching must be aligned with the anisotropic features in the flow. The anisotropy detection in this work is based on the work by Venditti and Darmofal,<sup>19</sup> which was extended to higher-order methods by Fidkowski and Darmofal.<sup>20</sup> 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 (orthogonal to the dominant direction) is selected to equidistribute the interpolation error in the principal directions. Assuming the interpolation error of the Mach number converges at the rate of  $r = p+1$ , the principal lengths,  $\{h_1, h_2\}$ , implied by the anisotropy request metric tensor,  $\mathcal{M}_{\text{ani}}$ , should satisfy

$$\left( \frac{h_2(\mathcal{M}_{\text{ani}})}{h_1(\mathcal{M}_{\text{ani}})} \right)^r = \frac{M_{\max}^{(p+1)}}{M_{\perp}^{(p+1)}}, \quad (6)$$

where  $M_{\perp}^{(p+1)}$  is the  $p+1$  derivative of the Mach number in the direction orthogonal to  $M_{\max}^{(p+1)}$ . Without loss of generality,  $\mathcal{M}_{\text{ani}}$  is specified to have a determinant of unity. The choice of the Mach number as the scalar quantity representing the solution behavior follows from previous work.<sup>19,21</sup> While the choice of the Mach number is somewhat arbitrary, the authors have found it works well in practice for flows with shocks, wakes, and boundary layers. The convergence rate,  $r$ , is nominally set to  $p+1$ ; however, the rate is reduced to  $r = 1$  when the shock is detected, and the derivative quantities used in Eq. (6) are replaced by the first derivatives,  $M_{\max}^{(1)}$  and  $M_{\perp}^{(1)}$ . In the case of vanishing  $M_{\perp}^{(p+1)}$ , the maximum aspect ratio is limited to 1,000. 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 iterations 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,  $\mathcal{M}_{\text{ani}}$ , the anisotropic metric request  $\mathcal{M}_r$  is constructed as

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

However, even with the more robust truth surrogate reconstruction, both the error estimate and the anisotropy detection sometimes suffer from noise 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 restricted such that length change in any direction is confined to a given factor, i.e.

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

where  $\mathcal{M}_{\bar{r}}$  is the limited requested metric tensor and the admissible changes for refinement and coarsening are set to  $\Delta_r = 1/4$  and  $\Delta_c = 2$ . The detailed information is provided by Modisette.<sup>22</sup>

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,<sup>3</sup> the degrees of freedom of a mesh conforming to a Riemannian metric field,  $\mathcal{M}_{\bar{r}}$ , is approximated by

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

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

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

The metric scaling algorithm provides an explicit control of degrees of freedom independent of the number of adaptation cycles.

### 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,<sup>23</sup> which generates linear anisotropic meshes (i.e. straight edged elements). However, linear meshes are unsuitable for higher-order discretizations, and the meshes must be modified to capture higher-order geometry information of curved surfaces. Simply curving the elements on the boundary is not a viable option for highly anisotropic meshes used to resolve boundary layers, as the surface may intersect the opposing interior face as shown in Figure 2. This work uses two methods to overcome this problem: elastically-curved meshes and cut-cell meshes.




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

### III.D.1. *Elastically-Curved Meshes*

The first option is to globally curve a linear boundary conforming mesh by elasticity.<sup>7,24,25</sup> In this method, the mesh is thought of as an elastic material, and the entire mesh is deformed as the boundary faces are moved to conform to curved surfaces. This method requires a mesh generator that can generate a linear, anisotropic, boundary-conforming mesh, a difficult task in three dimensions.

### III.D.2. *Cut-Cell Meshes*

The second mesh generation technique used is a simplex cut-cell method first presented by Fidkowski and Darmofal<sup>20</sup> and extended by Modisette and Darmofal.<sup>26</sup> 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.

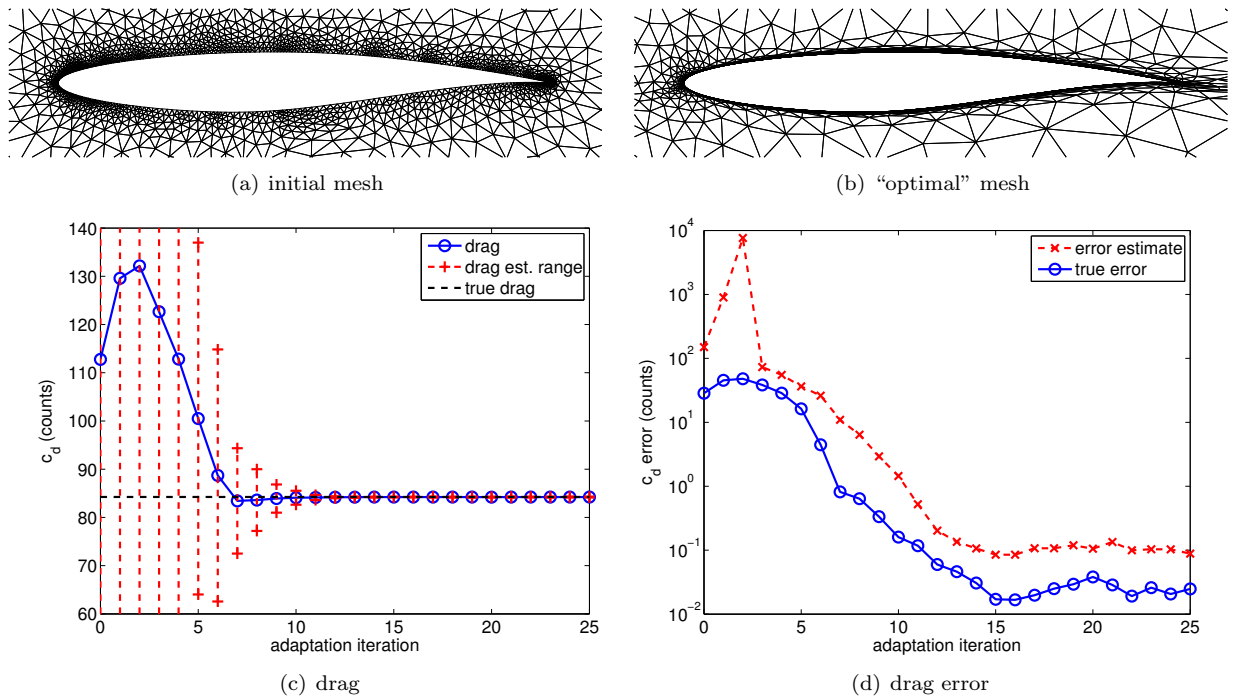
## III.E. Key Features of the Adaptation Strategy

One of the key features of the proposed adaptation strategy is the dof-control mechanism, discussed in Section III.C. As will be shown in Section IV.A, realizing the full potential of 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 a fixed degree 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, since the equidistribution of the local errors is a necessary but not sufficient condition for true optimality. While the size distribution is optimized in the proposed procedure, the shapes of the elements may not be optimal when the Mach number does not serve as the best representation of the anisotropic behaviors of both the primal and dual solutions. Nevertheless, the Mach number based anisotropy detection appears to work well for aerodynamics applications.

To illustrate the behavior of the fixed-dof adaptation algorithm, the algorithm is applied to a  $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$  turbulent flow over an RAE2822 airfoil at  $\alpha = 2.31^\circ$ . The RANS equations are discretized using the  $p = 3$  DG method (4th order accurate), and the adaptation is performed while fixing the degrees of freedom to 40,000. The initial mesh, shown in Figure 3(a), consists of 4,005 isotropic elements and is not suitable for an accurate RANS simulation. After 15 adaptation cycles, the algorithm produces the “optimal” mesh consisting of 3,983 anisotropic elements, shown in Figure 3(b). Figure 3(c) shows the behavior of the drag output, and Figure 3(d) shows the associated error estimate and true error over the adaptation cycles. The true error is computed relative to the drag coefficient obtained on adapted  $p = 3$ , dof = 160,000 mesh. The adaptation leads to an error reduction of over three orders of magnitude, even though the degrees of freedom is unaltered. The adaptation strategy achieves this by redistribution of the element sizes and employment of highly anisotropic elements in the boundary layers and the wake.

Figure 3(d) shows that, after 15 adaptation cycles, the mesh is optimized for this configuration, and the  $c_d$  error estimate fluctuates around 0.1 drag counts. The stagnation of the error estimate indicates that the fixed-dof algorithm has reached stationarity, and the mesh is “optimal”. Note that the number of adaptation cycles required to achieve optimality is dependent on the quality of the initial mesh; a typical case requires fewer adaptation cycles than this case, because the initial mesh typically comes from an optimized mesh at lower degrees of freedom or a lower discretization order. The error estimate convergence history also shows that the fixed-dof adaptation results in a generation of a family of “optimal” meshes that have similar error levels, in this case generated between the 15th and 25th adaptation iterations. All of these meshes have similar metric fields but slightly different triangulations. The family of slightly different meshes arise from the non-uniqueness of the meshes that realize a given metric field<sup>3</sup> and the non-stationarity of the fixed-fraction marking strategy. To account for this fluctuation in the output quantity, the result section reports the envelope of the output quantities computed using five “optimal” meshes for each degrees of freedom. Similarly, the error estimate values are obtained by averaging the five error estimate values.

The output error history in Figure 3(d) also shows that the error estimate is indicative of the true error behavior. While the DWR error estimate does not produce a strict bound, our experience suggests that it provides a conservative estimate of the true error in most cases. Due to the difficulty in obtaining the “true” solution for more complex flows, the results section uses the error estimate along with the range of fluctuation in the output values to assess the quality of the solutions.



**Figure 3.** An example of the initial and “optimal” meshes, the output history, and the error convergence history for the fixed-dof adaptation applied to the subsonic RAE2822 RANS-SA flow ( $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$ ,  $\alpha = 2.31^\circ$ ,  $p = 3$ ,  $\text{dof} = 40,000$ ).

The explicit degree of freedom control is also 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 similar parameter value and maintain a particular degree of freedom count. The effectiveness of the adaptation in performing a parameter sweep is demonstrated in Section IV.C.

## 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 aerodynamic applications, the error convergence behaviors of uniform and adaptive refinements are compared. The first problem considered is a  $M_\infty = 0.5$  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 each. 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 fixed-dof adaptation procedure at 10,000 and 20,000 degrees of freedom. The test is performed for  $p = 1$  and  $p = 3$  solution spaces to assess the impact of adaptation for different discretization orders. Note that the true drag for this problem, posed on the finite domain, is nonzero due to the point vortex influencing the flow angle at the farfield (though not accounted for in the boundary conditions).

The error estimates in Figure 4(a) show that adaptation makes a greater impact for the  $p = 3$  discretization than for the  $p = 1$  discretization. The  $p = 3$  discretization is less efficient than the  $p = 1$  discretization in the absence of adaptive refinement at 10,000 degrees of freedom. The suboptimal convergence rate of uniform refinement is due to the presence of a corner singularity at the trailing edge; the convergence rate is limited by the solution regularity, rather than the interpolation order. With a proper mesh grading, the optimal convergence rate of  $\mathcal{E} \sim h^{2p+1} \sim (\text{dof})^{p+1/2}$  for the output quantity is recovered.

The drag coefficient history, shown in Figure 4(b), confirms the behavior of the error estimate. As



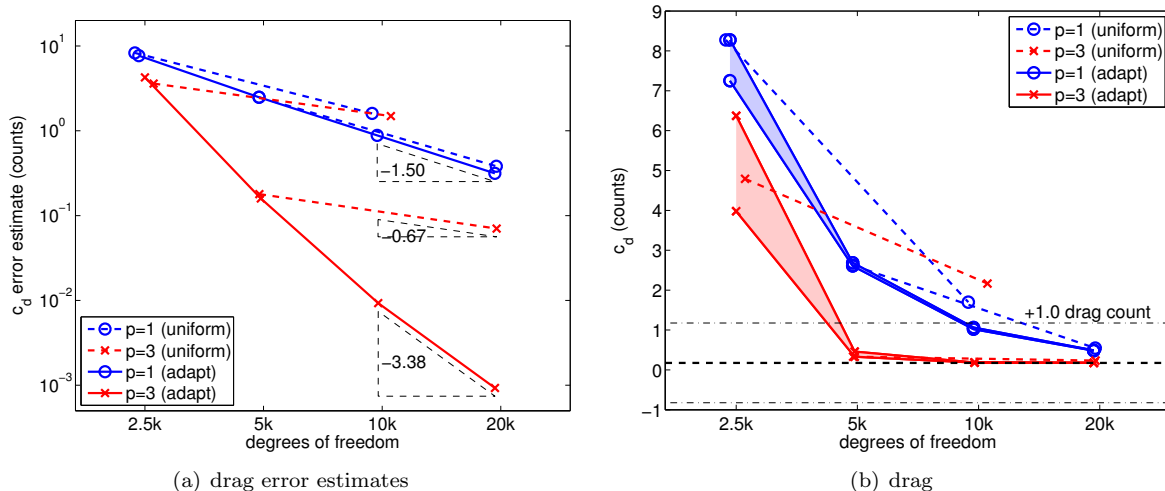


Figure 4. Comparison of the error convergence for uniform and adaptive refinements for the subsonic NACA0012 Euler flow ( $M_\infty = 0.5$ ,  $\alpha = 2.0^\circ$ ).

discussed in Section III.E, the output values computed from the family of “optimal” meshes exhibit small fluctuation; the shaded region in the figure indicates the range of output values computed using the five “optimal” meshes for each of the select degrees of freedom. The uniform refinement results are obtained by subdividing one of these five meshes, and thus the 2,500 and 5,000 dof uniform refinement drag values fall within the shaded region. Note that the fluctuation in the output value decreases with the mesh refinement in general. With adaptation, the  $p = 3$  discretization quickly converges to the true drag; however, in the absence of adaptation, the  $p = 3$  discretization with 10,000 degrees of freedom is less accurate than the adaptive  $p = 1$  discretization. These two behaviors are consistent with that captured by the error estimate, shown in Figure 4(a).

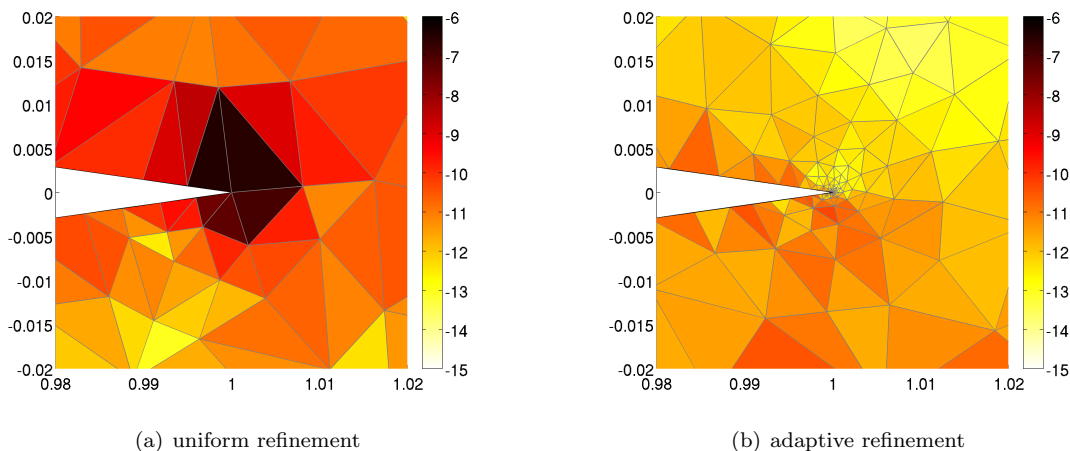


Figure 5. Comparison of the trailing edge mesh grading and error indicator distribution of the  $p = 3$ , dof = 20,000 meshes obtained from uniform and adaptive refinements for subsonic NACA0012 Euler flow ( $M_\infty = 0.5$ ,  $\alpha = 2.0^\circ$ ). The color scale is in  $\log_{10}(\eta_K)$ .

To understand the differences in the error convergence behaviors of the uniform and adaptive refinements, the element size and local error distribution near the trailing edge singularity are analyzed. Figure 5 shows the error indicator distribution at the trailing edge for the  $p = 3$ , dof = 20,000 meshes obtained after uniform 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 due to the inefficient element size distribution. The adaptive refinement targets the corner elements dominating the error and produces a strongly graded mesh that nearly equidistributes the error. The diameter of the trailing edge element is approximately  $3 \times 10^{-5}c$  for the adapted mesh, whereas that

for the uniform mesh is  $h \approx 10^{-2}c$ .

The difference in the mesh grading required to achieve optimality for the  $p = 1$  and  $p = 3$  discretizations is also important to understand. Figure 6 shows the variation in the element size,  $h$ , as a function of the distance from the trailing edge,  $r$ , for  $p = 1$  and  $p = 3$  optimized meshes. Each mesh has approximately 2,000 elements. The solid regression lines are produced by performing least-squares fit in the  $\log(h)$ - $\log(r)$  space for the elements in  $r < 0.1c$ . The  $p = 3$  optimal mesh has a grading factor of 0.86 and employs elements of diameter  $3 \times 10^{-5}c$  on the trailing edge. In comparison, the  $p = 1$  optimal mesh has a weaker grading factor of 0.66, and its trailing edge elements are of diameter  $3 \times 10^{-3}c$ . The higher-order discretization requires a stronger grading toward the corner singularity to equidistribute the error. In addition, the higher-order discretization is more sensitive to suboptimal  $h$  distribution as the error scales with 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.

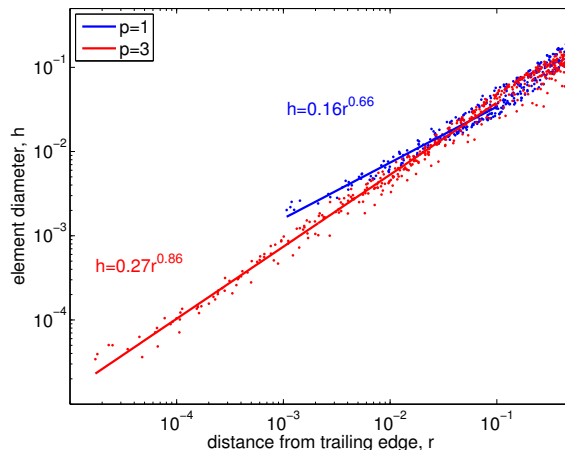


Figure 6. Element size distributions in the vicinity of the trailing edge of the  $p = 1$  and  $p = 3$  optimized meshes for the subsonic NACA0012 Euler flow ( $M_\infty = 0.5$ ,  $\alpha = 2.0^\circ$ ).

#### 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 a  $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$  turbulent flow over an RAE2822 airfoil at  $\alpha = 2.31^\circ$ . Following the procedure for the Euler NACA0012 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.

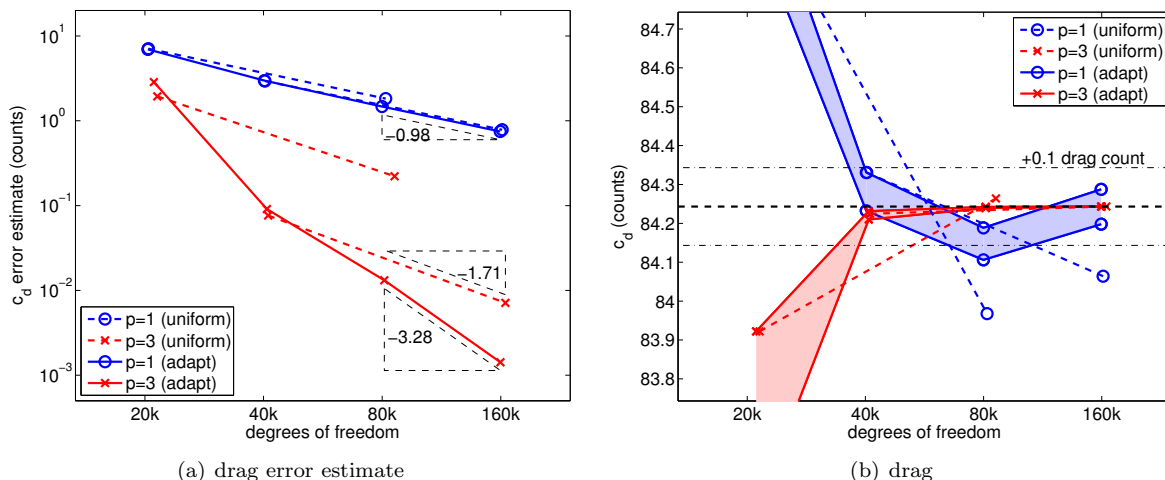
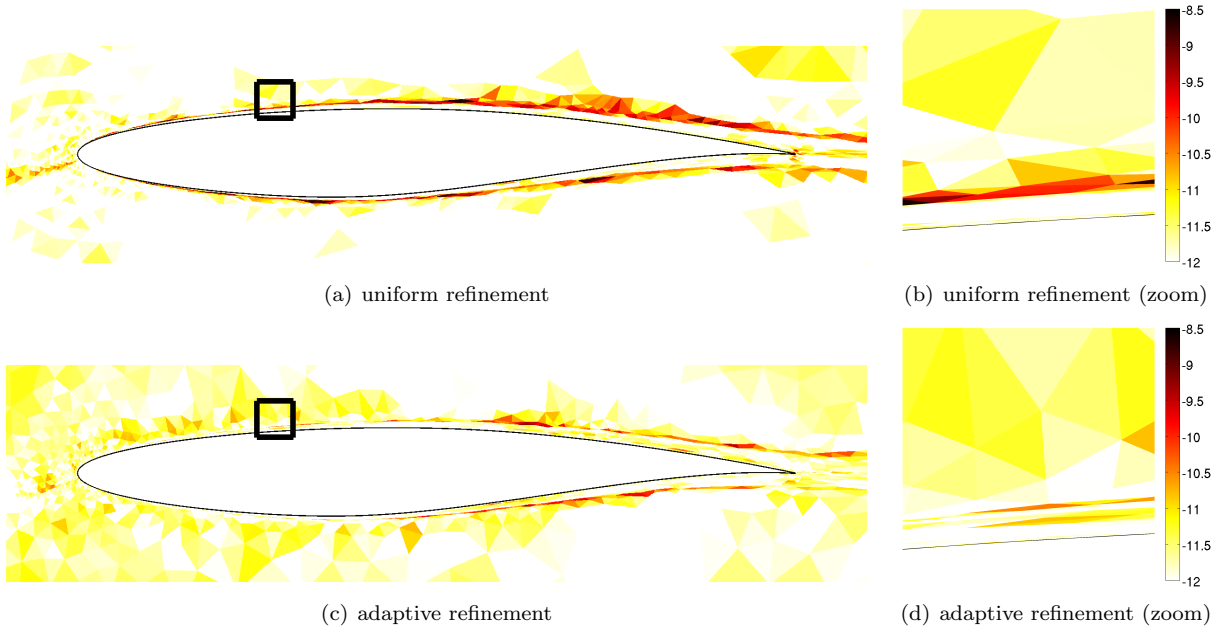


Figure 7. Comparison of the error convergence for uniform and adaptive refinements for the subsonic RAE2822 RANS-SA flow ( $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$ ,  $\alpha = 2.31^\circ$ ).

The error estimate and drag history are shown in Figure 7. Similar to the Euler 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  $\mathcal{E} \sim (\text{dof})^{2p/d}$  for the output quantity is recovered. The “optimal”  $p = 3$  mesh achieves a drag error estimate of approximately 2 counts using 20,000 degrees of freedom (2,000 elements), whereas the optimal  $p = 1$  mesh requires 80,000 degrees of freedom (27,000 elements) to achieve the same fidelity. The drag values in Figure 7(b) show that the adaptive  $p = 3$  discretization converges to the stationary value quicker than its  $p = 1$  counterpart, confirming the effectiveness of the higher-order discretization when combined with adaptation.

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 8. Due to the singularity in the SA equation on the outer edge of the boundary layer,<sup>8</sup> elements in the region are deemed to have high error. The adaptive refinement correctly identifies the region and makes the necessary adjustments to remove these high error elements. The boundary layer edge singularity is unique to the RANS-SA equations and is an example of a flow feature that is hard to locate *a priori*. However, this subtle flow feature limits the convergence rate of higher-order discretizations and exemplifies the need for adaptation driven by an *a posteriori* error estimate for higher-order methods.



**Figure 8.** Comparison of the error indicator distributions of  $p = 3$ ,  $\text{dof} = 160,000$  meshes obtained from uniform and adaptive refinements of the  $p = 3$ ,  $\text{dof} = 40,000$  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  $\log_{10}(\eta_\kappa)$ .

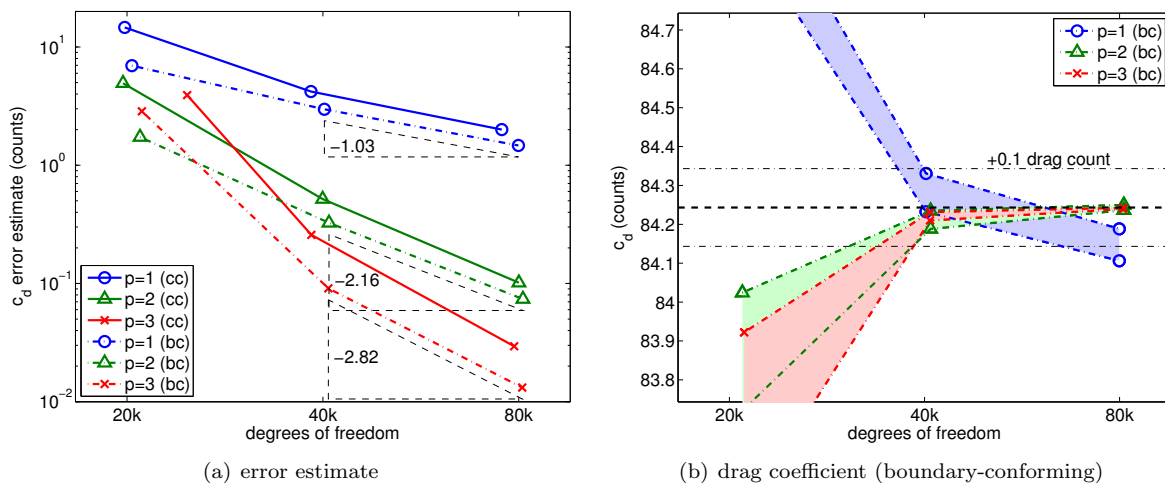
## IV.B. Adaptive High-Order Discretization of RANS-SA Flows

In this section, the adaptive, high-order DG discretization is applied to subsonic, transonic, and supersonic RANS-SA flows. Recalling the importance of adaptation for high-order discretization discussed in Section IV.A, this section only presents adaptive results. The results for elastically-curved and cut-cell meshes are presented for each case.

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

The first case considered is the same RANS-SA flow over an RAE2822 airfoil as Section III.E and IV.A.2 ( $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$ ,  $\alpha = 2.31^\circ$ ). Figure 9(a) shows that the higher-order discretizations ( $p = 2, 3$ ) are superior to the second order discretization ( $p = 1$ ) for high-fidelity simulations requiring a drag error estimate of less than 1 count. In fact, on the highly graded meshes, the higher-order discretization is more efficient at an error estimate level of as high as a few drag counts. Comparing the two high-order mesh

generation strategies, the elastically-curved boundary-conforming (bc) meshes are more efficient than the cut-cell (cc) meshes; however, the effectiveness of the adaptive higher-order discretization can be realized using either strategy.

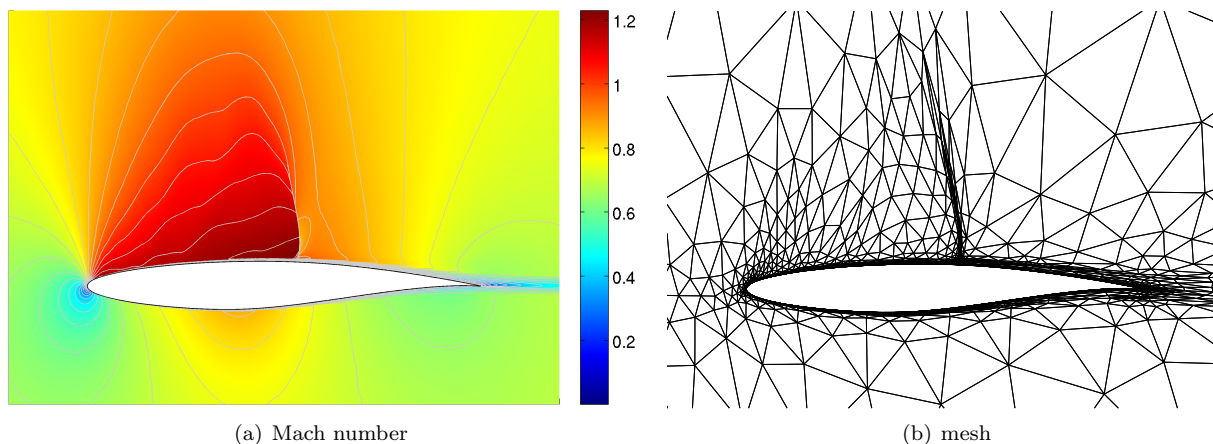


**Figure 9.** The drag coefficients and  $c_d$  error estimates for the subsonic RAE2822 RANS-SA flow ( $M_\infty = 0.3$ ,  $Re_c = 6.5 \times 10^6$ ,  $\alpha = 2.31^\circ$ ).

The drag convergence behavior in Figure 9(b) supports the claim of the high-order efficiency, with the  $p = 2$  and 3 discretizations exhibiting faster convergences to the stationary  $c_d$  value with rapidly decreasing fluctuation. The drag convergence for the cut-cell meshes is similar to that of boundary-conforming meshes and are not shown for brevity. This subsonic RANS case shows that, if the singularity in the SA equation on the boundary layer edge is handled appropriately, the benefit of higher spatial accuracy to resolve the boundary layer can be realized at lower degrees of freedom.

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

Having shown the effectiveness of the higher-order method for the subsonic RANS-SA flow, the next two cases demonstrate the benefit of higher-order discretizations even in the presence of shocks when the mesh is aggressively graded toward singular features. The first case is a transonic flow over an RAE2822 airfoil, with the flow condition  $M_\infty = 0.729$ ,  $Re_c = 6.5 \times 10^6$ , and  $\alpha = 2.31^\circ$ . The output of interest is drag.



**Figure 10.** 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.

The Mach number distribution and the drag-adapted mesh obtained using a  $p = 3$ , dof = 40,000 discretization are shown in Figure 10. The mesh is graded aggressively toward the airfoil surface, the boundary layer edge, and the shock. In contrast, the high-order discretization uses large elements to capture

smoothly varying features away from the airfoil. The Mach contour indicates that the combination of anisotropic grid refinement and the PDE-based shock capturing algorithm enables sharp resolution of the shock.

The  $c_d$  error estimate in Figure 11(a) shows that, even at the error estimate level of 10 drag counts, the  $p = 2$  discretization is more efficient than the  $p = 1$  discretization. For higher-fidelity simulations requiring tighter error tolerance, the  $p > 1$  discretizations are clearly more efficient. The convergence of drag coefficient, shown in Figure 11(b), confirms the superior convergence of the higher-order discretizations. The higher accuracy is achieved by the reduction in the degrees of freedom required to resolve the boundary layer when employing a higher-order approximation. Even though the higher-order method does not improve the efficiency of shock resolution, overall the method is more efficient for transonic turbulent problems with thin boundary layers when 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.

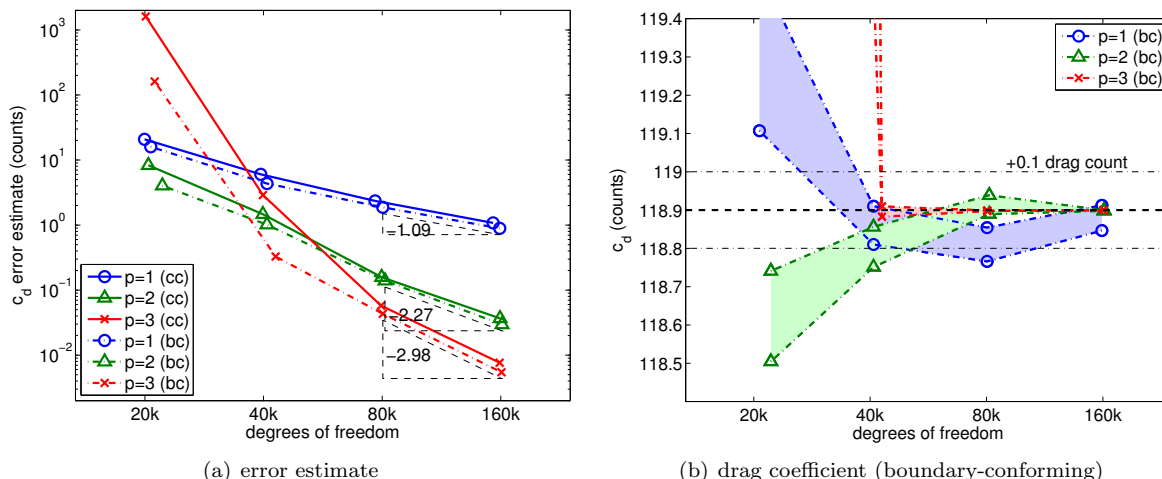


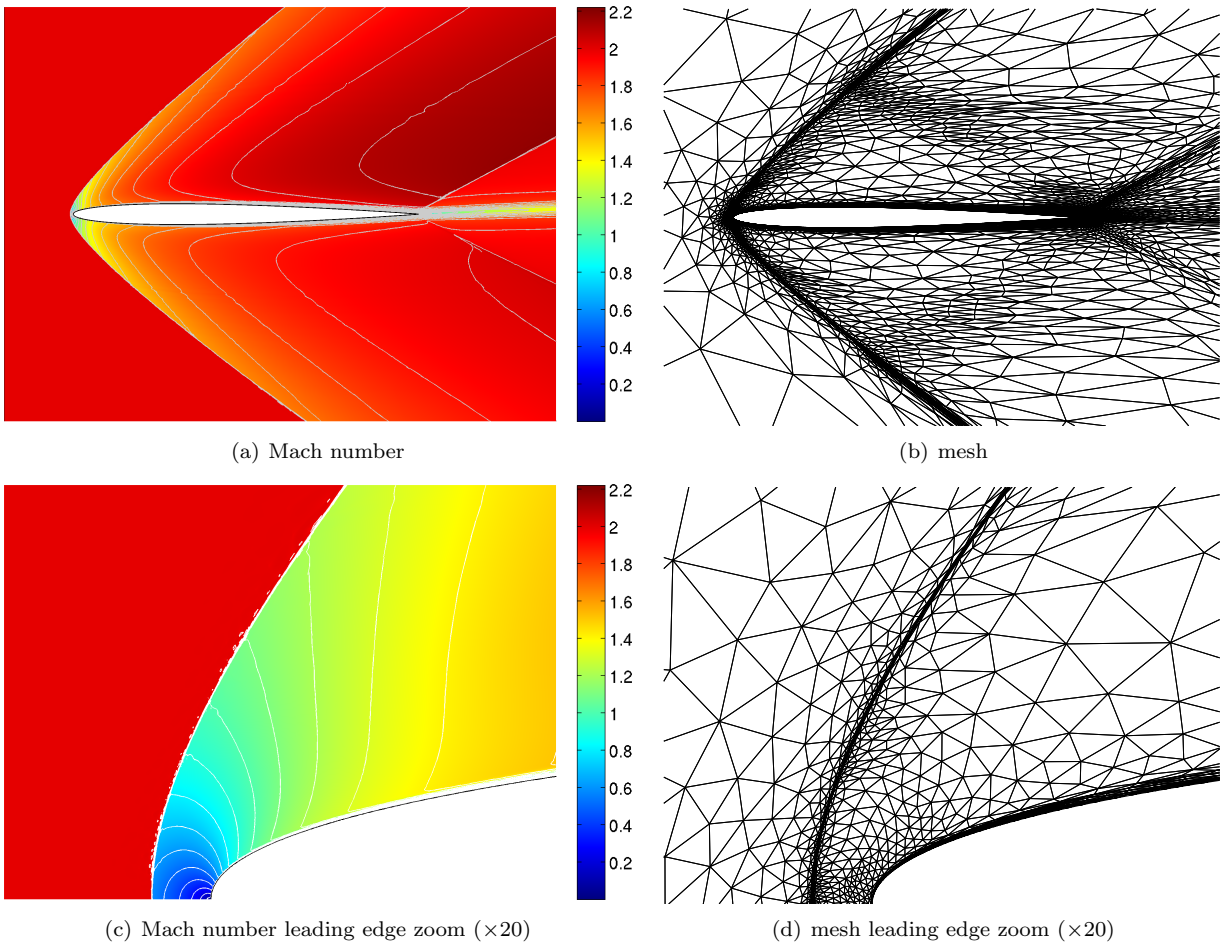
Figure 11. The drag coefficients for the transonic RAE2822 RANS-SA flow ( $M_\infty = 0.729$ ,  $Re_c = 6.5 \times 10^6$ ,  $\alpha = 2.31^\circ$ ).

For this transonic case, the  $p = 3$  discretization does not work effectively at 20,000 degrees of freedom, which only permits the use of 2,000 elements. The small number of elements causes the discretization to fail to achieve the asymptotic performance even with the aggressive mesh refinement. While the adaptation extends the applicability of the higher-order discretization to degrees of freedoms with higher-errors, a limit exists to where its effectiveness can be extended.

#### IV.B.3. NACA0006 Supersonic RANS-SA : $M_\infty = 2.0$ , $Re_c = 1.0 \times 10^6$ , $\alpha = 2.0^\circ$

The second shock problem is a supersonic flow over a NACA0006 airfoil with the flow condition  $M_\infty = 2.0$ ,  $Re_c = 1.0 \times 10^6$ , and  $\alpha = 2.0^\circ$ . 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 dual 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 strongly refined by the adaptation.

Figure 13(a) shows the convergence of the drag error estimate for the problem. Due to the low regularity of the dominant flow feature, the benefit of the  $p = 2$  and  $p = 3$  discretizations is difficult to realize even with the adaptive algorithm. Similar to the transonic RAE2822 flow, the higher-order discretizations reduce the number of elements required to resolve the boundary layer sufficiently so that the methods are more efficient than the  $p = 1$  discretization asymptotically. This shock problem also uncovers the limit of the effectiveness of higher-order discretizations at low degrees of freedoms. At 40,000 degrees of freedom the  $p = 3$  discretization is not able to compete with the  $p = 1$  and  $p = 2$  discretizations due to the inability to resolve the shock with a sufficient number of elements. In fact, the increase in element count for the  $p = 2$  discretization at 80,000 and 160,000 degrees of freedom is more effective than the higher-spatial accuracy of



**Figure 12.** The Mach number distribution and the mesh for the supersonic NACA0006 RANS-SA flow ( $M_\infty = 2.0$ ,  $Re_c = 1.0 \times 10^6$ ,  $\alpha = 2.0^\circ$ ) obtained using a  $p = 2$ , dof = 80,000 discretization. The Mach contour lines are in 0.1 increments.

the  $p = 3$  discretization at the same degree of freedom count.

The differences between the boundary-conforming and cut-cell meshes are smaller for this case, as the dominant feature, i.e. the bow shock, does not interact with the airfoil surface, where the two mesh generation strategies behave differently. The  $c_d$  values in Figure 13(b) confirm the superior performance of the adaptive  $p = 2$  discretization, showing faster convergence toward the stationary value, but also illustrate the struggles of the  $p = 3$  discretization for this flow.

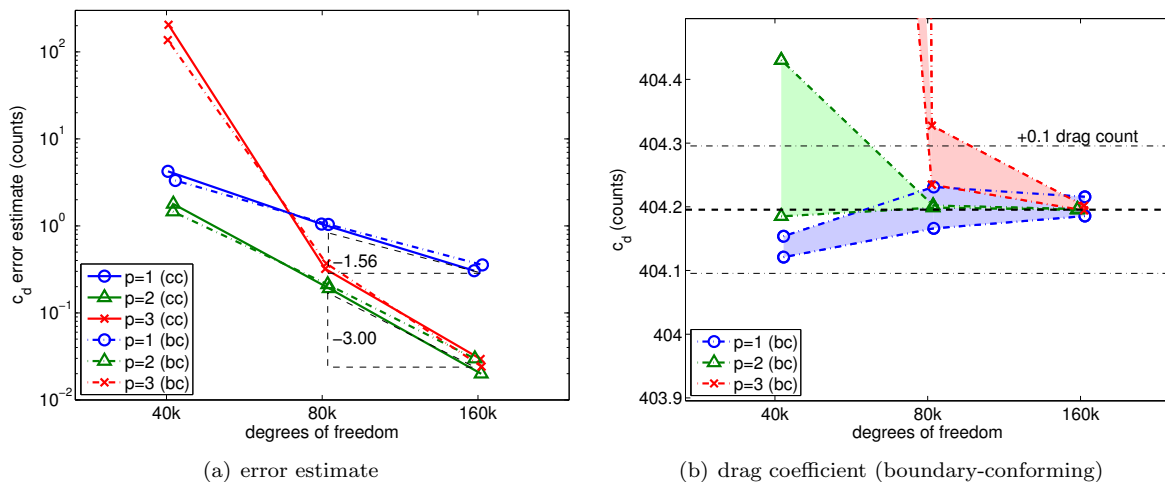


Figure 13. The  $c_d$  error estimate convergence for the supersonic NACA0006 RANS-SA flow ( $M_\infty = 2.0$ ,  $Re_c = 1.0 \times 10^6$ ,  $\alpha = 2.0^\circ$ ).

#### 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 effectively perform a parameter sweep, the lift curve for the McDonnell Douglas Aerospace (MDA) three-element airfoil (30P-30N)<sup>27</sup> is constructed. For this high-lift airfoil test case, the freestream Mach number and Reynolds number are set to  $M_\infty = 2.0$  and  $Re_c = 9 \times 10^6$ , and the angle of attack is varied from  $0.0^\circ$  to  $24.5^\circ$ . All results are obtained using  $p = 2$  polynomials at 90,000 degrees of freedom.

To generate the adaptive meshes, the fixed-dof adaptation is performed at each angle of attack using the mesh optimized for the previous angle of attack as the starting mesh. For example, the fixed-dof adaptation is performed starting from the  $8.1^\circ$ -optimized mesh to generate the  $12^\circ$ -optimized mesh. As the flow features do not change significantly from one angle of attack to the next, only a few fixed-dof adaptation iterations are necessary to generate the “optimal” mesh at a new angle of attack.

Figure 14(a) shows the lift curves obtained by 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 14(b) shows that the error estimate correctly identifies the lack of confidence in the solution for the high angle of attack cases, producing  $c_l$  error estimates on the order of 10. The elemental error estimate indicates that the elements on the upper surface of the slat dominate the fixed grid  $c_l$  errors for angles of attack above  $20^\circ$ . 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 15 shows the Mach number distributions obtained for the  $\alpha = 23.28^\circ$  flow on the fixed and adapted grids. The Mach number distribution indicates that the fixed grid lacks resolution on the front side of the 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 on the rest of the flow, which leads to the local error estimate being high. With proper mesh resolution, the adapted grid eliminates the numerical dissipation induced separation. The shock on the

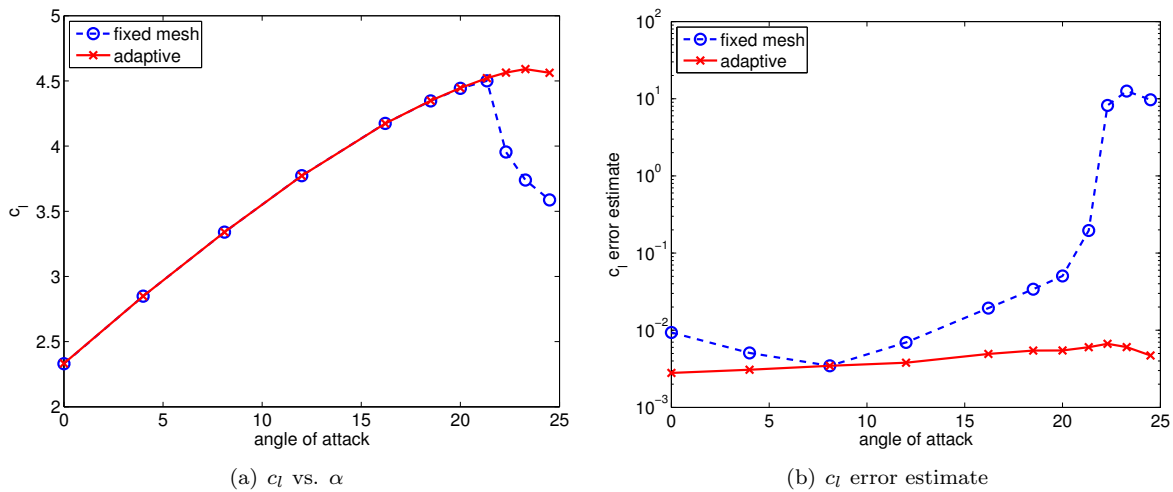


Figure 14. The lift curve and the  $c_l$  error obtained using the fixed mesh and adaptive meshes for the three-element MDA airfoil.

front side of the slat is captured sharply by a combination of the shock capturing mechanism and aggressive refinement.

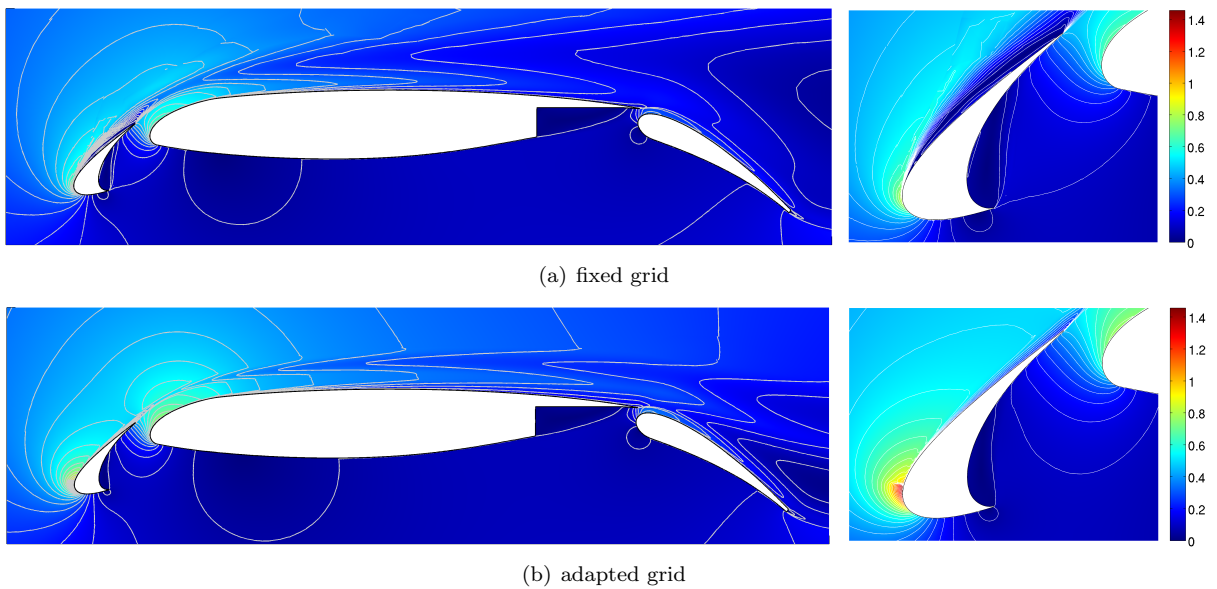


Figure 15. The Mach number distribution for the three-element MDA airfoil at  $\alpha = 23.28^\circ$  obtained on the  $8.10^\circ$  optimized mesh and the  $23.28^\circ$  optimized mesh. The Mach contour lines are in 0.05 increments.

Figure 16 shows the initial mesh and the adapted meshes for select angles of attack. The meshes exhibit different flow features as the angle of attack varies. At lower angles of attack, the flow separates from the bottom side of the slat, and the wake must be captured to account for its influence on the main element. At  $\alpha = 23.28^\circ$ , capturing the acceleration on the upper side and the shock becomes important for accurate calculation of lift.

## V. Conclusions

This paper presents an adaptation strategy that iterates toward dof-“optimal” meshes using an output-based error estimate and explicit control of degrees of freedom. With the adaptation strategy, “optimal” meshes are generated that can realize the benefit of higher-order discretizations at low degrees of freedom. The key features of these meshes are strong grading toward singularities and anisotropic resolutions for






Figure 16. The initial and lift-adapted grids for the three-element MDA airfoil at selected angles of attack.

boundary layers, wakes, and shocks. The numerical experiments demonstrate that uniform refinement is insufficient to attain the benefits of high-order discretizations even starting from “optimal” coarse meshes. With 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 estimate level of as high as 10 drag counts—the level at which the discretization error dominates the modeling error. Furthermore, cut-cell meshes are competitive with elastically-curved meshes in terms of solution efficiency, making the method an attractive choice for complex, three-dimensional geometries. The fixed-dof adaptation allows for an effective parameter sweep by adjusting the meshes to different flow features, as shown for a high-lift, transonic configuration. The combination of higher-order methods, robust nonlinear solver, and output-based adaptation algorithm provides a solution strategy that is a step closer to making a fully automated CFD a reality.

## 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.

## References

- <sup>1</sup>Rannacher, R., “Adaptive Galerkin finite element methods for partial differential equations,” *Journal of Computational and Applied Mathematics*, Vol. 128, 2001, pp. 205–233.
- <sup>2</sup>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.
- <sup>3</sup>Loseille, A. and Alauzet, F., “Continuous Mesh Model and Well-Posed Continuous Interpolation Error Estimation,” INRIA RR-6846, 2009.
- <sup>4</sup>Roe, P. L., “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes,” *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372.
- <sup>5</sup>Bassi, 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.
- <sup>6</sup>Spalart, P. R. and Allmaras, S. R., “A one-equation turbulence model for aerodynamics flows,” AIAA 1992-0439, Jan. 1992.
- <sup>7</sup>Oliver, 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.
- <sup>8</sup>Oliver, T. and Darmofal, D., “Impact of turbulence model irregularity on high-order discretizations,” AIAA 2009-953, 2009.
- <sup>9</sup>Barter, 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.
- <sup>10</sup>Persson, P.-O. and Peraire, J., “Sub-Cell Shock Capturing for Discontinuous Galerkin Methods,” AIAA 2006-0112, 2006.
- <sup>11</sup>Gropp, 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.
- <sup>12</sup>Saad, Y., *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 1996.
- <sup>13</sup>Diosady, L. and Darmofal, D., “Discontinuous Galerkin Solutions of the Navier-Stokes Equations Using Linear Multigrid Preconditioning,” AIAA 2007-3942, 2007.
- <sup>14</sup>Persson, 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.
- <sup>15</sup>Becker, 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.
- <sup>16</sup>Becker, 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.
- <sup>17</sup>Fidkowski, K. and Darmofal, D., “Output error estimation and adaptation in computational fluid dynamics: Overview and recent results,” AIAA 2009-1303, 2009.
- <sup>18</sup>Yano, M. and Darmofal, D., “On dual-weighted residual error estimates for  $p$ -dependent discretizations,” ACDL Report TR-11-1, Massachusetts Institute of Technology, 2011.
- <sup>19</sup>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, No. 1, 2003, pp. 22–46.

<sup>20</sup>Fidkowski, 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.

<sup>21</sup>Fidkowski, 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.

<sup>22</sup>Modisette, J. M., *An Automated Reliable Method for Two-Dimensional Reynolds-averaged Navier-Stokes Simulations*, PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2011.

<sup>23</sup>Hecht, F., “BAMG: Bidimensional Anisotropic Mesh Generator,” 1998, <http://www-rocq1.inria.fr/gamma/cdrom/www/bamg/eng.htm>.

<sup>24</sup>Nielsen, 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.

<sup>25</sup>Persson, P.-O. and Peraire, J., “Curved mesh generation and mesh refinement using Lagrangian solid mechanics,” AIAA 2009-0949, 2009.

<sup>26</sup>Modisette, J. M. and Darmofal, D. L., “Toward a Robust, Higher-Order Cut-Cell Method for Viscous Flows,” AIAA 2010-721, 2010.

<sup>27</sup>Klausmeyer, 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.