



Drag Prediction Using Adaptive Discontinuous Finite Elements

Marco Ceze* and Krzysztof J. Fidkowski†
University of Michigan, Ann Arbor, Michigan 48109

DOI: 10.2514/1.C032622

This paper presents the results obtained with an adaptive finite-element solver for the Fifth Drag Prediction Workshop. The discontinuous Galerkin finite-element method is used for the spatial discretization of the Reynolds-averaged Navier–Stokes equations with a modified version of the Spalart–Allmaras turbulence model. Drag convergence is sought via mesh adaptation driven by an adjoint-weighted residual method. Results are presented for the drag polar of the NACA 0012 airfoil under subsonic flow conditions and for the Common Research Model wing–body geometry under transonic flow conditions and fixed lift. The angle of attack that yields the desired lift is obtained via a Newton solve using the lift adjoint.

Nomenclature

$b(i)$	=	benefit function
$c(i)$	=	cost function
C	=	convective flux
C_D	=	coefficient of drag
C_L	=	coefficient of lift
C_p	=	coefficient of pressure
\mathcal{D}	=	diffusive flux
E	=	mass-specific total energy
f^{adapt}	=	fraction of elements to adapt
J	=	output of interest
M	=	mass matrix
M	=	Mach number
$m(i)$	=	merit function
q	=	geometric polynomial order
$R(\cdot)$	=	discrete residual operator
$\mathbb{R}(\cdot, \cdot)$	=	weak form of the residual
Re	=	Reynolds number
S	=	source term for Spalart–Allmaras model
T^H	=	computational mesh
U	=	discrete state
\mathbf{u}	=	state vector
$\mathbf{u}^{H,p}$	=	elementwise polynomial representation of the state
u_s	=	state component s
$\mathcal{V}^{H,p}$	=	space of p -order polynomials with support over elements of T^H
\tilde{v}	=	working variable for Spalart–Allmaras model
v_i	=	i th component of velocity
$\mathbf{w}^{H,p}$	=	vector of weight functions
α	=	angle of attack
ϵ_{tol}	=	trimming tolerance
η_{κ^H}	=	elemental adaptive indicator
κ^H	=	element of computational mesh
κ_{SA}	=	scaling factor for discrete Spalart–Allmaras equation
μ_∞	=	freestream laminar dynamic viscosity

ρ	=	density
Ψ	=	discrete adjoint state
$\boldsymbol{\psi}$	=	adjoint state vector
$\boldsymbol{\psi}^{H,p}$	=	elementwise polynomial representation of the adjoint

I. Introduction

THE use of computational fluid dynamics (CFD) tools in engineering analysis and design has steadily increased in the past several decades. With the evolution of algorithms and the substantial enhancement of computational power, CFD tools now provide the ability to explore new configurations and test flow conditions that may be otherwise difficult to produce experimentally. As the range of applications becomes wider and the number of simulations increases, requirements of high accuracy and robustness present challenges for the CFD development community [1].

One application of CFD that demands high accuracy is drag prediction for a large transport aircraft: seemingly small variations in drag can significantly impact the aircraft’s payload [2,3]. Aerodynamic flow over an aircraft, as for many other cases, exhibits features with unknown spatial distribution, and the range of these features’ length scales can easily span six orders of magnitude. Furthermore, flows can exhibit singularities that pose additional challenges for the prediction of drag and other outputs. The trivial solution to these problems is to globally refine the mesh. However, this strategy is generally inefficient due to its refinement of unnecessary regions and the resulting very large grid sizes.

AIAA organizes drag and lift prediction workshops (DPW and HLPW, respectively) with the purpose of assessing the capability of state-of-the-art computational methods and turbulence modeling for predicting forces and moments on relevant geometries in the aeronautical industry. In these workshops, starting meshes are generated based on industry’s best practices, and mesh independence is generally sought via uniform-refinement studies. Nevertheless, the spread of results can be significant [4–7].

Solution-based adaptive methods present an attractive opportunity for accurate calculations on affordable grid sizes. These methods rely on the definition of an adaptive indicator, which localizes the regions of the computational domain that need mesh modification through refinement, coarsening, or node movement. An effective indicator is obtained through adjoint-based error estimation methods, which have already been demonstrated for many complex problems, including those in aerospace applications [8]. The goal of these methods is to provide confidence measures in the form of error bars for scalar outputs of engineering interest. In addition, one can use the error contributions of different elements or volumes of the computational mesh as an adaptive indicator that specifically targets errors in the outputs of interest [9–14].

Adaptive mechanics can include both mesh size, h refinement, and approximation order, p increment. Choosing the correct mechanics is

Presented as Paper 2013-0051 at the 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Grapevine, TX, 7–10 January 2013; received 12 September 2013; revision received 24 November 2013; accepted for publication 16 December 2013; published online 27 January 2014. Copyright © 2013 by Marco Ceze and Krzysztof J. Fidkowski. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-3868/14 and \$10.00 in correspondence with the CCC.

*Postdoctoral Research Fellow, Department of Aerospace Engineering, Member AIAA.

†Assistant Professor, Department of Aerospace Engineering, Senior Member AIAA.

important for efficient prediction of an output, and we employ a systematic cost/benefit approach to make this decision [15]. Increasing p requires high-order capability, and one candidate discretization that has gained popularity in aerodynamic applications is the discontinuous Galerkin (DG) method. Its popularity is due in part to its suitability for high-order discretizations of convection-dominated problems on unstructured meshes. In addition, DG's finite-element formulation naturally supports output error estimation and handles meshes with hanging nodes, both of which are important for this work.

The structure of this paper is as follows. Section II describes the flow solver and the discretization method. Section III outlines the output error estimation process that drives the mesh adaptation described in Sec. IV. Results are presented in Sec. V, and conclusions and future work are discussed in Sec. VI.

II. Solver Description

For our simulations, we use the XFlow code, a high-order DG finite-element solver for general equation sets with Message Passing Interface-based, distributed-memory parallel capabilities. XFlow serves as a platform for development in research areas such as error estimation, mesh adaptation, and solver algorithms [15–17]. In the case of this work, the equations are Reynolds-averaged Navier–Stokes (RANS) with a modified version of the Spalart–Allmaras (SA) [18,19] turbulence model.

The RANS–SA equations are written in their compact, conservative form as

$$\partial_t \mathbf{u}_s + \partial_i \mathcal{C}_{is}(\mathbf{u}) - \partial_i \mathcal{D}_{is}(\mathbf{u}, \nabla \mathbf{u}) = \mathcal{S}_s(\mathbf{u}) \quad (1)$$

where $i \in [1, \dots, \text{dim}]$ indexes the spatial dimensions, and s indexes the equations of conservation of mass, momentum, energy, and turbulent viscosity. Accordingly, the state vector is denoted by $\mathbf{u} = [\rho, \rho v_i, \rho E, \rho \tilde{\nu}]^T$.

The DG spatial discretization of the flow equations approximates the solution in a space $\mathcal{V}^{H,p}$ of piecewise polynomials of degree p with local support on each element $\kappa^H \in T^H$, where T^H is the set of elements resulting from a subdivision of the spatial domain. The resulting weak form reads

$$(\partial_t \mathbf{u}^{H,p}, \mathbf{w}^{H,p}) + \mathbb{R}(\mathbf{u}^{H,p}, \mathbf{w}^{H,p}) = 0 \quad \forall \mathbf{w}^{H,p} \in \mathcal{V}^{H,p} \quad (2)$$

where (\cdot, \cdot) denotes an inner product, and the semilinear form $\mathbb{R}(\mathbf{u}^{H,p}, \mathbf{w}^{H,p})$ includes source, convective, and diffusive terms.

In a DG approximation, the state can be discontinuous between elements, just like in a finite-volume method. The Riemann flux

involved in the convective term at element interfaces is approximated with Roe's [20] solver in which the SA working variable $\tilde{\nu}$ is transported as a conserved scalar $\rho \tilde{\nu}$. The diffusion term is discretized using the second form of Bassi and Rebay [21], and the SA source term is discretized according to Oliver's [19] modifications to the original SA model [18]. These modifications ensure stability of the model at negative $\tilde{\nu}$ and they are specifically suited for discontinuous Galerkin discretizations.

For problems with shocks, we use a modified version of Persson and Peraire's [22] shock-capturing scheme that uses a switch based on a regularity estimate of the density approximation; this involves the current p th-order solution and its projection onto $\mathcal{V}^{H,p-1}$. These modifications are described in [23].

The discrete system is obtained by expanding the components the state and the test functions in terms of basis functions $\phi^{H,p}(x)$, where $\mathcal{V}^{H,p} = \text{span}\{\phi^{H,p}(x)\}$. The resulting discrete system reads

$$M \frac{dU}{dt} = -R(U) \quad (3)$$

The mass matrix is block diagonal, and it consists of volume integrals of basis function products on each element in the mesh.


We use the constrained pseudotransient continuation [23] method for marching Eq. (3) in time. This method incorporates physical realizability constraints in the solution path and thus improves the robustness of the DG solver. At each pseudotransient continuation step, a linear system is solved to yield a state update direction ΔU . The coefficient matrix in this linear system is the residual Jacobian regularized by a matrix with the same footprint as the mass matrix in Eq. (3). In this work, we solve the linear systems using the generalized minimal residual method [24] (GMRES) with an element-line Jacobi preconditioner [3].

Most practical cases in the aeronautical industry are in the Reynolds number regime of $10^6 \rightarrow 10^7$. In this regime, the SA working variable $\tilde{\nu}$ typically spans several orders of magnitude. Therefore, it is also desirable to choose an appropriate scale for $\rho \tilde{\nu}$. The scale used in this work is


$$(\rho \tilde{\nu})' = \frac{\rho \tilde{\nu}}{\kappa_{SA} \mu_\infty} \quad (4)$$

where $(\rho \tilde{\nu})'$ is the scaled conserved variable that is stored and evolved by the solver. Essentially, we are nondimensionalizing $\rho \tilde{\nu}$ by a factor larger than the laminar viscosity, and the motivation for this scaling is to make the conserved variable $(\rho \tilde{\nu})'$ on the order of unity, on par with the other variables, in relevant regions of the computational domain.

To exemplify the effect of κ_{SA} , we show in Fig. 1 the residual history for two flows at $Re = 6.5 \times 10^6$, one subsonic and one



a) $M_\infty = 0.3, Re = 6.5 \times 10^6, \alpha = 2.31 \text{ deg}$



b) $M_\infty = 0.734, Re = 6.5 \times 10^6, \alpha = 2.79 \text{ deg}$

Fig. 1 Residual convergence for an RAE 2822 airfoil, using $p = 1$ for different $\tilde{\nu}$ scaling factors (κ_{SA}).

Table 1 Results for case in Fig. 1a solution costs normalized with $\kappa_{SA} = 1$

Quantity	$\kappa_{SA} = 1$	$\kappa_{SA} = 100$	$\kappa_{SA} = 1000$
C_D	1.2238×10^{-2}	1.2238×10^{-2}	1.2238×10^{-2}
C_L	4.5079×10^{-1}	4.5079×10^{-1}	4.5079×10^{-1}
$(\rho v_x)_{\max}$	1.25222	1.25222	1.25222
$(\rho \tilde{v})'_{\max}$	9.64975×10^2	9.64975	9.64974×10^{-1}
GMRES iterations per nonlinear iteration	1.0 (13.5)	1.023	1.1054
Total GMRES iterations	1.0 (648 iterations)	0.9799	0.8982
Total CPU time	1.0 (197 s)	0.9728	0.8804

Table 2 Results for case in Fig. 1b: solution costs normalized with $\kappa_{SA} = 1$

Quantity	$\kappa_{SA} = 1$	$\kappa_{SA} = 100$	$\kappa_{SA} = 1000$
C_D	1.9730×10^{-2}	1.9755×10^{-2}	1.9755×10^{-2}
C_L	7.3054×10^{-1}	7.3164×10^{-1}	7.3164×10^{-1}
$(\rho v_x)_{\max}$	1.11406	1.11178	1.11178
$(\rho \tilde{v})'_{\max}$	1.66004×10^3	1.63789×10^1	1.63789
GMRES iterations per nonlinear iteration	1.0 (23.9)	1.035	1.089
Total GMRES iterations	1.0 (3253 iterations)	0.6394	0.7925
Total CPU time	1.0 (745 s)	0.6481	0.7103

transonic. In the transonic case, we use the modified version of Persson and Peraire's [22] shock-capturing scheme mentioned in Sec. II. We solve the linear system in each nonlinear iteration to relative tolerances of 10^{-3} and 10^{-4} , respectively. For each case, we use three scaling factors, $\kappa_{SA} = 1, 100,$ and 1000 , and we see that κ_{SA} significantly affects the convergence history. Note that the initial residual norm for $\kappa_{SA} = 1$ is higher than for $\kappa_{SA} = 100$ and 1000 . Furthermore, when the residual norm for $\kappa_{SA} = 1$ drops to a similar order of magnitude as the initial residuals for $\kappa_{SA} = 100$ and 1000 , a secondary transient starts and the residual climbs again. The larger values of κ_{SA} ameliorate this secondary transient, which is also discussed by Burgess and Mavriplis [25].

By choosing an appropriate scale for the SA equation, we make the discrete residual for the SA equation similar in magnitude to the individual discrete residuals corresponding to the other conservation equations. Therefore, when we solve the linear systems for the state update direction to a finite tolerance, we evolve the solution to each of individual conservation equations together. This is desirable for implicit methods that make use of line searches, which is the case of pseudotransient continuation.

Tables 1 and 2 show the force coefficients, the maximum values for representative conserved quantities, and solution cost metrics. First, we analyze the effect of κ_{SA} on solution cost. The average cost of the linear solves, measured by the number of GMRES iterations per nonlinear step, is not significantly affected by the SA scaling. The reason for this cost being approximately constant across κ_{SA} is because the linear solves are not exact at each nonlinear step. Given that this cost is roughly constant, the overall solution cost is mainly dictated by the number of nonlinear steps. Note that the total number of nonlinear iterations in Fig. 1 correlates with the total number of GMRES iterations and with the total CPU time.

We now analyze the effect of SA scaling on the solution. In the fully subsonic case, the scaling has virtually no effect on the results, while in the transonic case, κ_{SA} has a slight effect on the force coefficients. This effect is due to the highly nonlinear nature of the shock-capturing scheme, which makes the output computations acutely sensitive to residuals. The difference in drag and lift in this case between $\kappa_{SA} = 100$ and $\kappa_{SA} = 1000$ is on the order of the residual tolerance.

We see from these two results that $\kappa_{SA} > 1$ can be beneficial for the solver, and we have found that $\kappa_{SA} \approx 100 \rightarrow 1000$ works well for Reynolds numbers in the range $10^6 \rightarrow 10^7$. Determining a general guideline for setting κ_{SA} at other Reynolds numbers is a subject of ongoing work.

III. Output Error Estimation

Output-based error estimation techniques identify all areas of the domain that are important for the accurate prediction of an engineering output. The resulting estimates properly account for error propagation effects that are inherent to hyperbolic problems, and they can be used to ascribe confidence levels to outputs or to drive adaptation. A key component of output error estimation is the solution of an adjoint equation for the output of interest. In a continuous setting, an adjoint, $\boldsymbol{\psi} \in \mathcal{V}$, is a Green's function that relates residual source perturbations to a scalar output of interest, $J(\boldsymbol{u})$, where \mathcal{V} is an appropriate function space. Specifically, given a variational formulation of a partial differential equation, determine \boldsymbol{u} such that

$$\mathbb{R}(\boldsymbol{u}, \boldsymbol{w}) = 0, \quad \forall \boldsymbol{w} \in \mathcal{V} \quad (5)$$

the adjoint $\boldsymbol{\psi} \in \mathcal{V}$ is the sensitivity of J to an infinitesimal source term added to the left-hand side of the original PDE. $\boldsymbol{\psi}$ satisfies a linear equation:

$$\mathbb{R}'[\boldsymbol{u}](\boldsymbol{w}, \boldsymbol{\psi}) + J'[\boldsymbol{u}](\boldsymbol{w}) = 0, \quad \forall \boldsymbol{w} \in \mathcal{V} \quad (6)$$

where the primes denote Fréchet linearization with respect to the arguments in square brackets. Details on the derivation of the adjoint equation can be found in many sources, including the review in [8]. Specifically, in the present work, we employ the discrete adjoint method, in which the system is derived systematically from the discretized primal system [26,27].

An adjoint solution can be used to estimate the numerical error in the corresponding output of interest. The resulting adjoint-weighted residual method is based on the observation that a solution $\boldsymbol{u}^{H,p}$ in a finite-dimensional approximation space will generally not satisfy the original PDE. The adjoint $\boldsymbol{\psi} \in \mathcal{V}$ translates the residual perturbation to an output perturbation via

$$\delta J = J(\boldsymbol{u}^{H,p}) - J(\boldsymbol{u}) \approx -\mathbb{R}(\boldsymbol{u}^{H,p}, \boldsymbol{\psi}) \quad (7)$$

This expression is based on a linear analysis, and hence for nonlinear problems and finite-size perturbations, the result is approximate.

Although the continuous solution \boldsymbol{u} is not required directly, the continuous adjoint $\boldsymbol{\psi}$ must be approximated to make the error

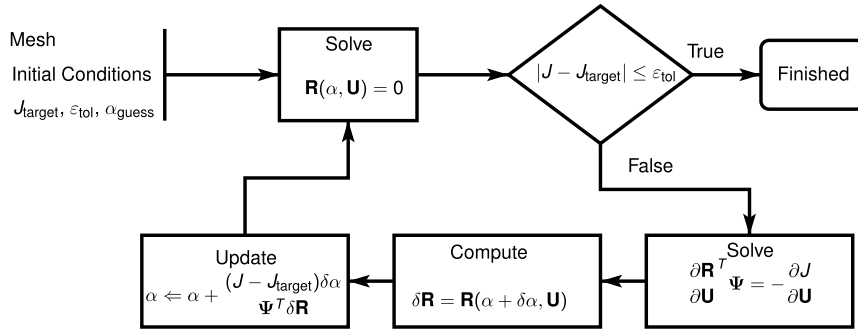


Fig. 2 Adjoint-based boundary-condition parameter correction, here shown for the angle of attack α .

estimate in Eq. (7) computable. In practice, ψ^{h,p^+} is solved approximately or exactly on a finer finite-dimensional space $\mathcal{V}^{h,p^+} \mathcal{V}^{H,p}$ [28–30]. This finer space can be obtained through mesh subdivision and/or approximation order increase [19,31,32], denoted here by changes in the superscript H and p , respectively.

The adjoint-weighted residual evaluation in Eq. (7) can be localized to yield an adaptive indicator consisting of the relative contribution of each element to the total output error. In this work, the finer space is obtained by approximation order increment, $\mathcal{V}^{h,p^+} \mathcal{V}^{H,p}$, and ψ^{h,p^+} is approximated by injecting $\psi^{H,p}$ into \mathcal{V}^{h,p^+} and applying element block Jacobi smoothing iterations. Via experimentation with a variety of flow problems, we found that five block Jacobi iterations are generally enough to provide good error estimates while maintaining the computational cost affordable.

The output perturbation in Eq. (7) is approximated as

$$\delta J \approx - \sum_{\kappa^H \in T^H} \mathbb{R}_{\kappa^H} (I_{H,p}^{H,p+1}(\mathbf{u}^{H,p}), \tilde{\psi}^{H,p+1} - I_{H,p}^{H,p+1}(\psi^{H,p})) \quad (8)$$

where $I_{H,p}^{H,p+1}(\cdot)$ is an injection operator from p to $p + 1$ in the coarse mesh T^H ; $\tilde{\psi}^{H,p+1}$ is the approximated fine-space adjoint; and \mathbb{R}_{κ^H} corresponds to the elemental residual as defined in Eq. (2). Note, the difference between the coarse-space and fine-space adjoints is not strictly necessary due to Galerkin orthogonality [8]. However, when the primal residual is not fully converged to machine precision levels, the use of the adjoint perturbation gives better error estimates. Equation (8) expresses the output error in terms of contributions from each coarse element. A common approach for obtaining an adaptive indicator is to take the absolute value of the elemental contribution in Eq. (8) [14,29,33–36],

$$\eta_{\kappa^H} = |\mathbb{R}_{\kappa^H} (I_{H,p}^{H,p+1}(\mathbf{u}^{H,p}), \tilde{\psi}^{H,p+1} - I_{H,p}^{H,p+1}(\psi^{H,p}))| \quad (9)$$

With systems of equations, indicators are computed separately for each equation and summed together. Because of the absolute values, the sum of the indicators, $\sum_{\kappa^H} \eta_{\kappa^H}$, is greater or equal to the original output error estimate. However, it is not a bound on the actual error because of the approximations made in the derivation.

Frequently in the aeronautical industry, CFD simulations are conducted under trimmed conditions, meaning under fixed, user-defined values of certain outputs, typically lift or pitching moment. This means that certain boundary condition parameters (e.g., angle of attack) depend on outputs computed from the flow solution. Thus, a feedback loop must be used to correct those input parameters.

The feedback loop used in this work is illustrated in Fig. 2, where J_{target} is the target value of the output for which the parameter α is trimmed. The cycle starts by solving the flow equations using an initial guess for α . Then, J is computed and checked against J_{target} under a trimming tolerance. Until this tolerance is met, α is corrected using Newton’s method for which the sensitivity of J with respect to α is needed. This sensitivity is computed via an inner product between an adjoint for J and a residual perturbation δR resultant from a perturbation $\delta\alpha$. This residual perturbation is computed by evaluating

the residual with the boundary condition perturbed by a small, user-defined $\delta\alpha$. In this work, $\delta\alpha$ is one one-thousandth of a radian. When the output depends directly on the input parameter, $\partial J / \partial \alpha$ is added to the output sensitivity, and this is also computed by finite differences.

In cases where the target value for the output is not achievable or the initial guess is bad, the cycle in Fig. 2 may not converge. In those cases, a contingency plan is needed (e.g., a maximum number of iterations is assigned or the cycle is restarted with a better initial guess). In the output-based adaptation framework presented in this work, the boundary conditions are only trimmed if the error estimate for J is smaller than its trimming tolerance, ϵ_{tol} .

IV. Mesh Adaptation Mechanics

The elemental adaptive indicator η_{κ^H} drives a fixed-fraction hanging-node adaptation strategy. In this strategy, which was chosen for simplicity and predictability of the adaptive algorithm, a certain fraction f^{adapt} of the elements with the largest values of η_{κ^H} is marked for refinement. Marked elements are refined according to discrete options that correspond to subdividing the element in different directions or increasing the approximation order. For quadrilaterals, the discrete options are: x -refinement, y -refinement, and xy -refinement, as depicted in Fig. 3. Although the option of modifying the local polynomial approximation order is possible in this framework [15], here we consider only h -adaptation. The directions x and y refer to reference-space coordinates of elements that can be arbitrarily oriented and curved in physical space. Also, the subelements created through refinement inherit the approximation order from the original element. In three dimensions, a hexahedron can be refined in seven ways: three single-plane cuts, three double-plane cuts, and isotropic refinement.

In this work, h -refinement is performed in an element’s reference space by employing the coarse element’s reference-to-global coordinate mapping in calculating the refined element’s geometry node coordinates. The refined elements inherit the same geometry approximation order and quadrature rules as the parent coarse element. As a result, there is no loss of element quality when a nonlinear mapping is used to fit the element to a curved geometry. Therefore, curved elements near a boundary can be efficiently refined to capture boundary layers in viscous flow. For simplicity of implementation, the initial mesh is assumed to capture the geometry




Fig. 3 Quadrilateral h -refinement options. The dashed lines indicate the neighbors of the refined element.




Fig. 4 Example of hanging-node refinement: shaded element is marked for refinement, and the dashed lines indicated edges created to keep one level of refinement difference between neighbors.

sufficiently well, through a high enough order of geometry interpolation on curved boundaries, such that no additional geometry information is used throughout the refinements. That is, refinement of elements on the geometry boundary does not change the geometry. We note that, for highly anisotropic meshes, curved elements may be required away from the boundary, and for simplicity, we use meshes with curved elements throughout the domain.

Elements created in a hanging-node refinement can be marked for h -refinement again in subsequent adaptation iterations. In this case, neighbors will be cut to keep one level of refinement difference between adjacent cells. This is illustrated in Fig. 4.

The choice of a particular refinement option is made locally in each element flagged for refinement. This choice is made by defining a merit function that ranks each available refinement option i . This function is defined as

$$m(i) = \frac{b(i)}{c(i)} \quad (10)$$

The benefit and cost measures depend on the method used for solving the flow equations and they should be tailored for each specific solver.

During calculation of the merit function, local mesh and data structures are created, one for each element, that include the flagged element and its first-level neighbors along with the corresponding primal and adjoint states. In these local structures, the central element is refined in turn according to each of the discrete options. On the refined local mesh, the merit function is computed, and the refinement option with the largest value of $m(i)$ is chosen.

Because we are seeking the most efficient way of locally refining an element amongst the available options, it is important that the cost and benefit measures in Eq. (10) are accurate but tractable representations of the computational expense and gain in accuracy, respectively. In [15,23], we discuss these aspects at length.

In an output-based mesh adaptation cycle, the steady-state residual is driven to zero at each adaptive step. Therefore, mesh modification on the element level can be interpreted as uncovering local residual perturbations. Because an adjoint solution represents the sensitivity of an output with respect to a residual perturbation, we define our benefit function as

$$b(i) = \sum_{\kappa^h \in \kappa^H} |R_{\kappa^h}(U_k T_{ki}(i))_j| |\Psi_k T_{kj}(i)| \quad (11)$$

where $R_{\kappa^h}(\cdot)_j$ is a discrete residual component in the embedded element, $T(i)$ is a matrix that transfers the discrete primal and adjoint states to the local meshes for each refinement i , and Ψ is the discrete coarse-space adjoint solution. Note that the adjoint variables act as positive weights for each of the perturbations.

In this work, most of the computational time is spent in the GMRES algorithm used to solve the linear systems that arise at each step of the pseudotransient continuation method. These systems are sparse; hence, we approximate the number of floating point operations in applying GMRES by the number of nonzero entries in the residual Jacobian matrix. Based on this observation, we define the cost measure as

$$c(i) = \sum_{\kappa^h \in \kappa^H} \left\{ (p+1)^{2 \cdot \dim} + \sum_{N_{\text{iface}}(i)} (p+1)^{2 \cdot \dim} \right\} \quad (12)$$




Fig. 5 NACA 0012, $M_\infty = 0.15$, $Re = 6 \times 10^6$, drag polar: initial mesh for $\alpha = 10$ deg (720 quartic elements).

where p denotes the polynomial approximation order, and $N_{\text{iface}}(i)$ is the number of internal faces associated with refinement option i . The first term in Eq. (12) accounts for the self-blocks of the residual Jacobian matrix corresponding to each of the subelements. The second term accounts for the off-diagonal blocks (i.e., the dependence of the subelements' residuals on the neighboring states).

V. Results

A. NACA 0012, $M_\infty = 0.15$, $Re = 6 \times 10^6$, Drag Polar

This case is one of the NASA's Turbulence Modeling Resource cases [37]. The purpose of this case is to validate the modifications made to the SA model. As suggested by NASA's Turbulence Modeling Resource, the domain's outer boundary is located 500 chord lengths away from the airfoil. We consider eight angles of attack in the drag polar: $\alpha = 0, 2, 4, 6, 8, 10, 12$, and 15 deg. For each angle of attack, an initial quartic mesh is generated by agglomerating 16 quadrilaterals from a linear mesh. We modify the original NACA 0012 geometry by closing the trailing edge according to [37]. The linear meshes are generated so that the cells downstream from the airfoil are approximately aligned with the wake. Figure 5 shows an example of an initial quartic mesh. Because of the geometrical simplicity of this case, we can have the first grid spacing off the wall on the initial mesh such that $y_{\text{max}}^+ \approx 0.06$ for $\alpha = 10$ and 15 deg and $y_{\text{max}}^+ \approx 0.03$ for $\alpha = 0$ deg.

The scheme's polynomial approximation order is $p = 2$, and the discretized SA equation is scaled by $\kappa_{\text{SA}} = 1000$. The adaptation is driven by drag error with $f^{\text{adapt}} = 10\%$, and the residual norm is reduced by eight orders of magnitude at each adaptive step. To simplify our analyses, we limit the number of the adaptive steps to six for all the angles and measure the error level of the final result. Figure 6 shows the drag convergence with degrees of freedom for three representative angles of attack. The largest final error estimate over all the angles of attack is approximately three drag counts ($\sim 3\%$) in the $\alpha = 15$ deg case.

Figure 7 compares our adaptive results with Ladson's experimental data [38] and with results computed with CFL3D [39] on a fine, 897×257 element, structured grid [37]. The experimental data consist of three sets of wind tunnel runs with varying roughness of carborundum strips to force transition to turbulence at the 5% position along the chord. This reduces transition effects and allows for a more adequate comparison with fully turbulent simulations.

In spite of the adaptation being driven by drag error, the lift values in Fig. 7a are in close agreement with the experimental data. Our computed drag values are within 3% difference with respect to CFL3D's results. This difference is within the spread of 4% in the CFD results reported for the SA model in [37]. With respect to the experimental values in Fig. 7b, the simulations show slightly larger drag values. We attribute these differences to the turbulence model and possibly to experimental measurement precision as the adjoint-based error estimation and adaptation only targets, and provides an error estimate for, the discretization error.

The adjoint solution offers insight on regions of the computational domain where residual errors affect the output of interest. Figure 8b shows the x -momentum drag-adjoint solution for the $\alpha = 10$ deg case. The most notable feature of this adjoint solution is the

stagnation streamline, which, in the inviscid limit, is a weak inverse-square-root singularity [40]. This sharp variation of the adjoint is reflected in the adapted mesh in Fig. 8a.

The outer edge of the boundary layer, where $\rho\tilde{v}$ exhibits strong variation, is also heavily targeted for refinement, as observed in [15]. We note that, because the initial solutions across all angles of attack had small values of y^+ , the first layer of cells off the wall is scarcely marked for refinement. Yet, between the initial and final solutions, there are large variations in force coefficients. This emphasizes a key aspect of adjoint-based error localization and adaptation, which is the reduced reliance on meshing guidelines for obtaining accurate results.

Other features that are important for accurate prediction of drag are the upper surface acceleration region, the trailing edge, and the wake. These regions are also frequently targeted for refinement because they present large magnitudes and variations of the adjoint variables.

B. Common Research Model: Wing–Body Geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$

This case consists of transonic, turbulent flow over NASA’s Common Research Model (CRM) [41]. This wing–body geometry mimics a modern passenger aircraft, and its purpose is to establish a reference for testing computational tools for simulation and design. Recently, this case has been added to the International Workshop on High-Order CFD methods [42], and it is considered a difficult case in the high-order CFD community because it challenges the robustness of curved mesh generation, nonlinear solution strategies, and

adaptation methods. Here, we consider only $p = 1$ solution approximation order due to the large cost of higher p -orders combined with the difficulty in making high-quality, curved, coarse meshes around the CRM geometry. The results serve as a proof of the concept of adaptive discontinuous finite-element methods applied to a problem with industrial relevance.

The cubic mesh used in this case was generated via agglomeration of linear cells. We do not consider mesh coarsening in our adaptation mechanics. Hence, the initial linear mesh was generated with the tradeoff of being coarse to use in our adaptation routine but fine enough to represent the geometry adequately. Figure 9 shows the linear and the agglomerated meshes. The off-wall spacing in the agglomerated mesh is such that $y^+ \approx 1$, computed from the initial solution, for most of the fuselage and the wing.

The discretized SA equation is scaled by $\kappa_{SA} = 1000$, and we use the modified version of Persson and Peraire’s [22] shock-capturing scheme mentioned in Sec. II. The convergence criterion is a residual-norm reduction of eight orders of magnitude from its initial value.

We consider anisotropic h -adaptation at fixed $p = 1$ with $f^{adapt} = 10\%$. Converging the initial solution for this problem is difficult. The physics-constrained solver with line search and the monotonic Residual Difference Method CFL strategy [23] is used for the first primal solve. In addition, one step of mesh adaptation based on the physics constraints [43] is taken to help the solver to converge. In subsequent solves, converging the residual is significantly easier.

The output used for adaptation is the total drag at a fixed lift. That is, at each primal solve, the angle of attack is trimmed so that the

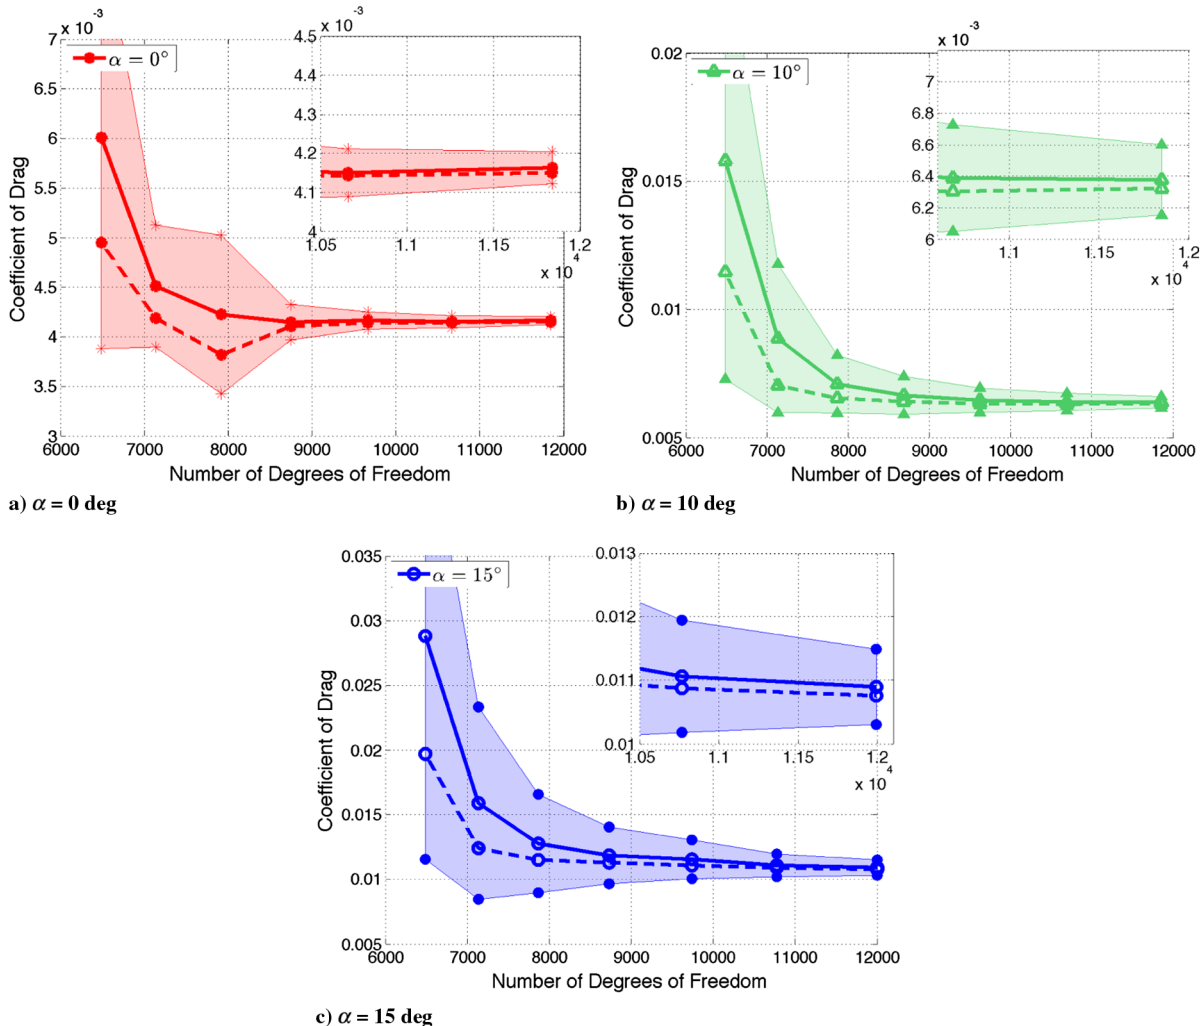
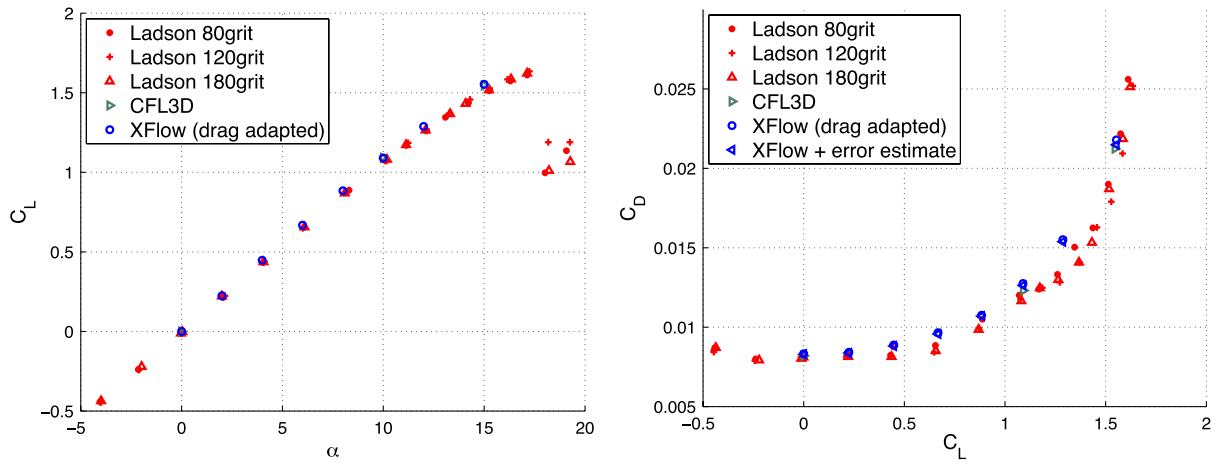


Fig. 6 NACA 0012, $M_\infty = 0.15$, $Re = 6 \times 10^6$, drag polar: for three angles of attack: solid lines: drag values; dashed lines: drag corrected by its error estimate; shading: sum of error indicators.

coefficient of lift is $C_{L_{target}} = 0.5 \pm 0.001$. The method for trimming α is described in Sec. V.A.

Because of a lack of spatial resolution in the initial mesh in the streamwise direction, the flow separates (Fig. 10a) before the lift


requirement is achieved. The solution in the initial mesh is nearly unsteady, which makes the adjoint problem very ill conditioned and, consequently, causes the error estimates to be very large, as shown in Figs. 11a and 11b. In this situation, the lift requirement is relaxed, and



a) Lift coefficient versus angle of attack

b) Drag polar

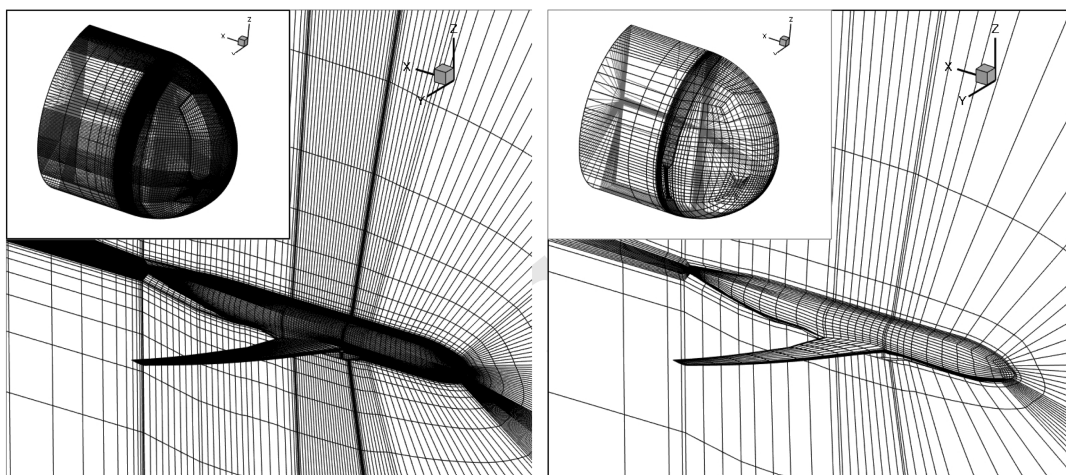
Fig. 7 NACA 0012, $M_\infty = 0.15$, $Re = 6 \times 10^6$, drag polar: comparison with experimental data.



a) 6th drag-adapted mesh and SA-working variable contours

b) 6th drag-adapted x-momentum adjoint solution for drag

Fig. 8 NACA 0012, $M_\infty = 0.15$, $Re = 6 \times 10^6$, drag polar: final mesh, $\rho\tilde{v}$ contours, and drag adjoint for $\alpha = 10$ deg.



a) Linear mesh used for agglomeration (1218375 elements)

b) Cubic mesh generated via agglomeration (45125 elements)

Fig. 9 CRM, wing-body geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$: linear and agglomerated cubic meshes.

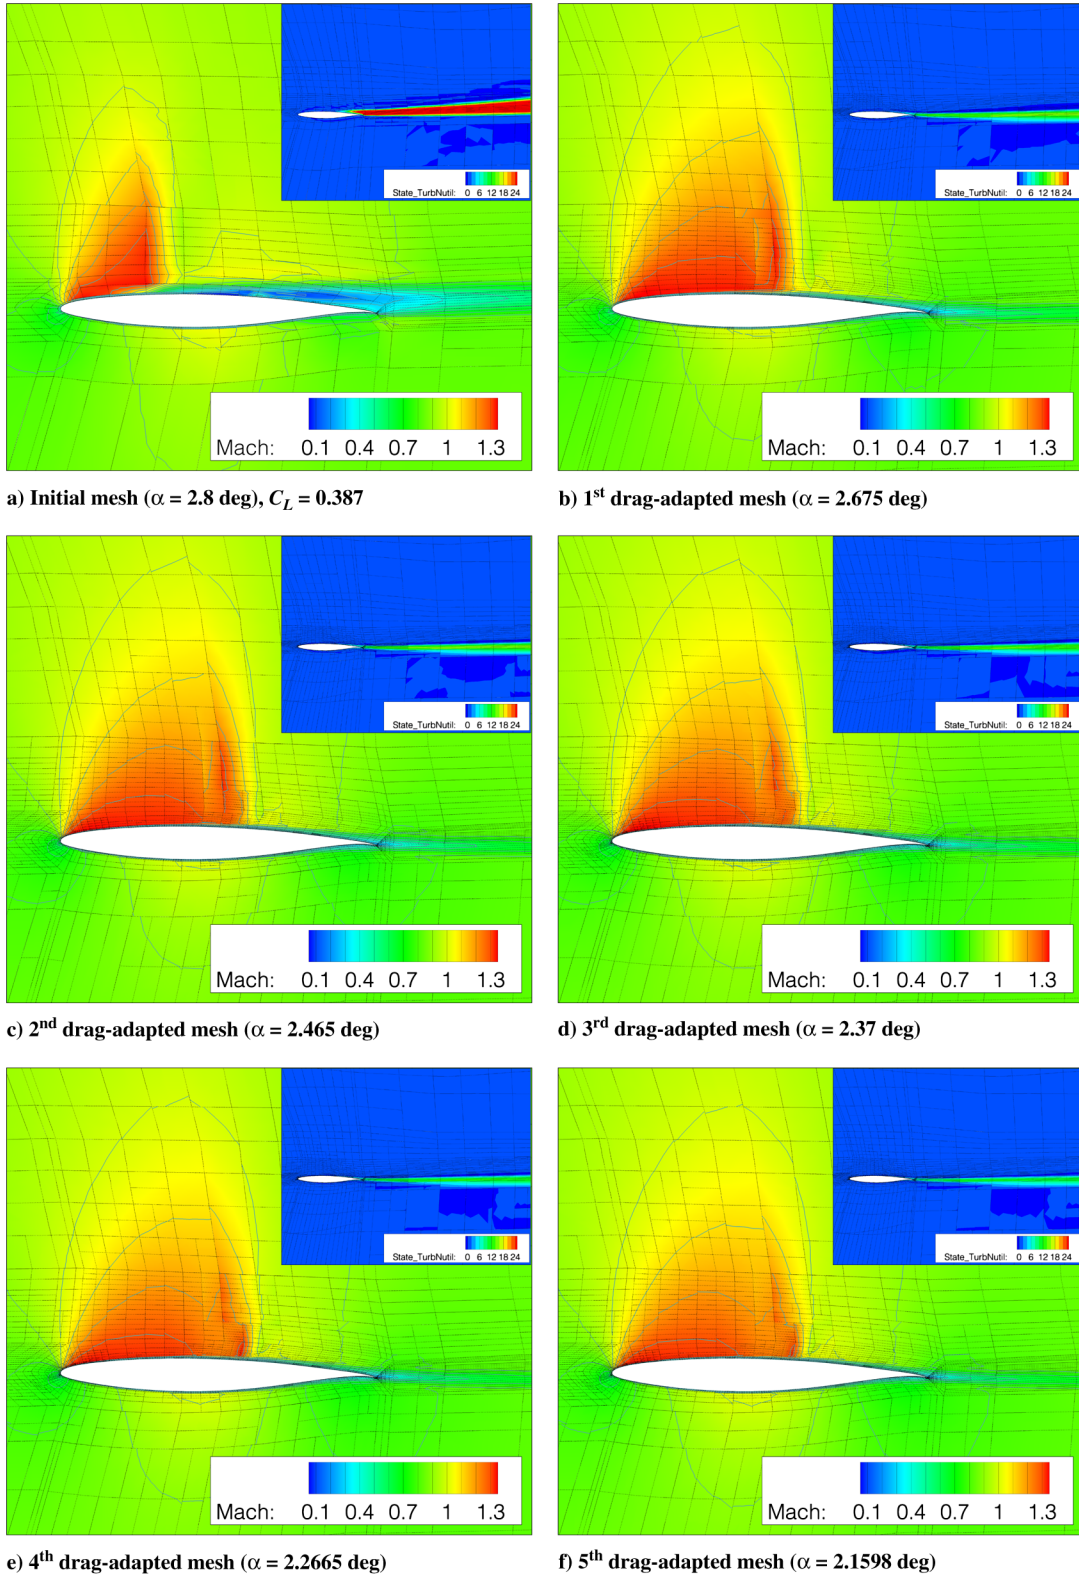


Fig. 10 CRM, wing-body geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$; slice at 37% of the span (428 in.). Note: on the initial mesh, the flow separates at $C_L = 0.387$.

the adaptive process proceeds. This decision is not yet automated and is one of the aspects of this problem that could benefit from further research.

After the first drag-based adaptation step, the flowfield is significantly different (Fig. 10b). The supersonic region is larger, and no visible flow separation is present. The lift requirement is now satisfied, and the error estimates for lift and drag are significantly smaller (Fig. 11).

The Mach number contours shown in Fig. 10 do not present large differences after the second adaptation step. Also, the areas targeted for adaptation are similar to the regions observed in the DPW III - W1 case presented in [15]. These regions are the stagnation streamline, the sonic transition, the shock-boundary-layer interaction, the wake, and edge of the boundary layer where $\rho\tilde{v}$ transitions from zero its maximum value. The final off-wall spacing is such that $y^+ \approx 0.2$ for most of the fuselage and wing surfaces, and the elements immediately

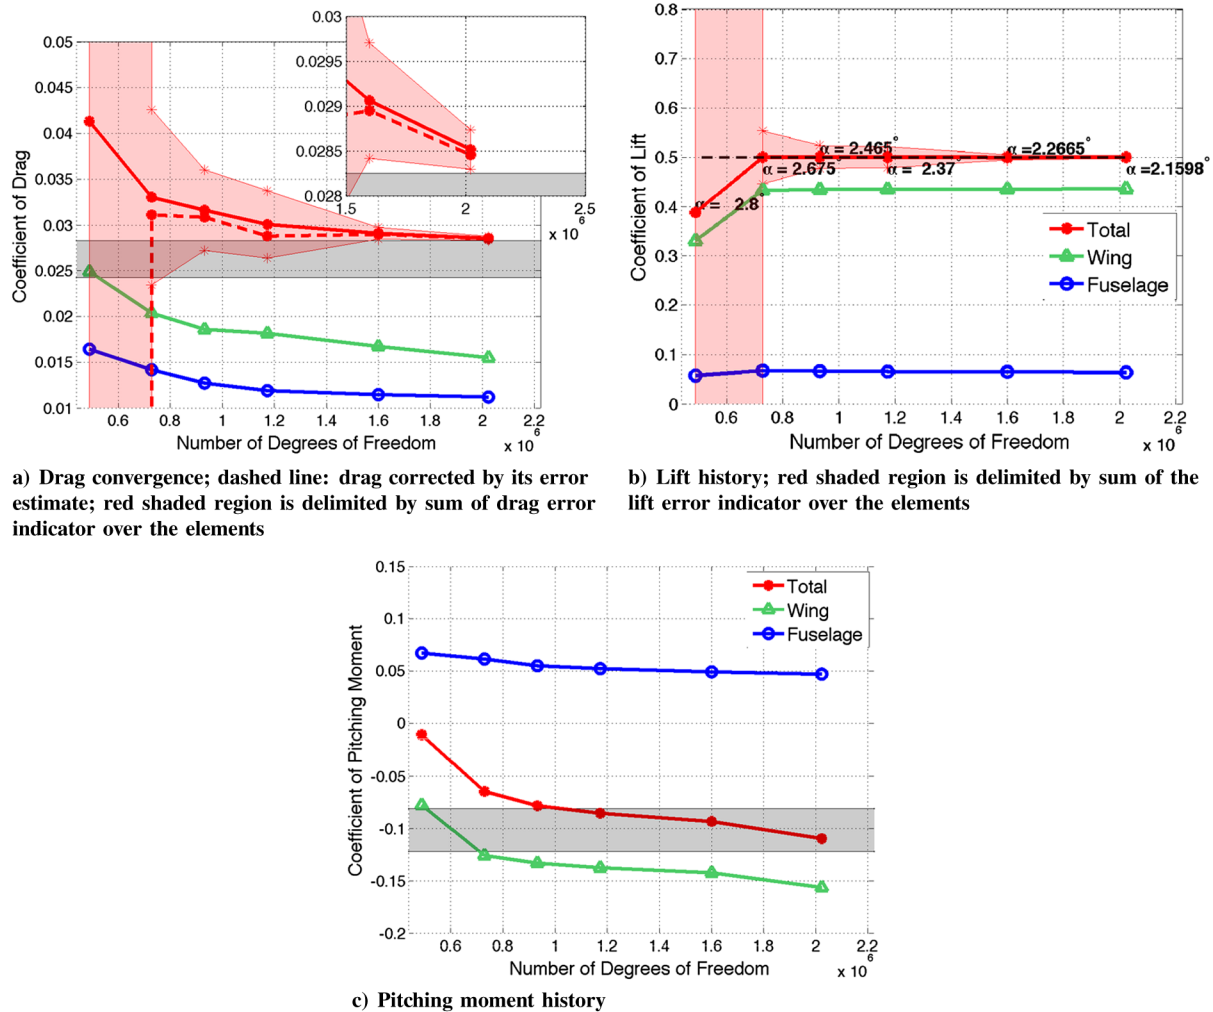


Fig. 11 CRM, wing-body geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$, sequence of adapted meshes; gray shaded region: range of data submitted to DPW-V. Note: on the initial mesh, the flow separates at $C_L = 0.387$.

attached to the surface are marked for refinement more often than in the two-dimensional case presented here.

Figure 12 compares the pressure coefficient at two span locations with the corresponding experimental data [44].[‡] Note that the initial result is very far from the experiments. However, after one adaptation step, the pressure distribution is much closer to the experimental data, and as the adaptation progresses, the shock profile becomes sharper, and the changes in pressure distribution become smaller.

Figure 11 shows the convergence history for drag, lift, and pitching moment. Note that our results for pitching moment are within the range of data submitted to the workshop, while the drag values are above the range of results from the workshop. However, it is worth emphasizing that the finest solution presented here has a factor of five to 10 fewer degrees of freedom than the midrange meshes used in the uniform refinement studies in DPW-V.

Note that the drag error correction for this case is not as effective as in the two-dimensional results. Here, two aspects are affecting the quality of the drag error estimate. One is the robustness of the fine-space approximation of the adjoint solution, which is affected by the under-resolution of the mesh. The second aspect is the angle of attack changing from one mesh to the next due to lift changes caused by the mesh refinement. This suggests that variations in the freestream boundary conditions due to lift error should be

incorporated in the drag error estimate, and this is the subject of current work.

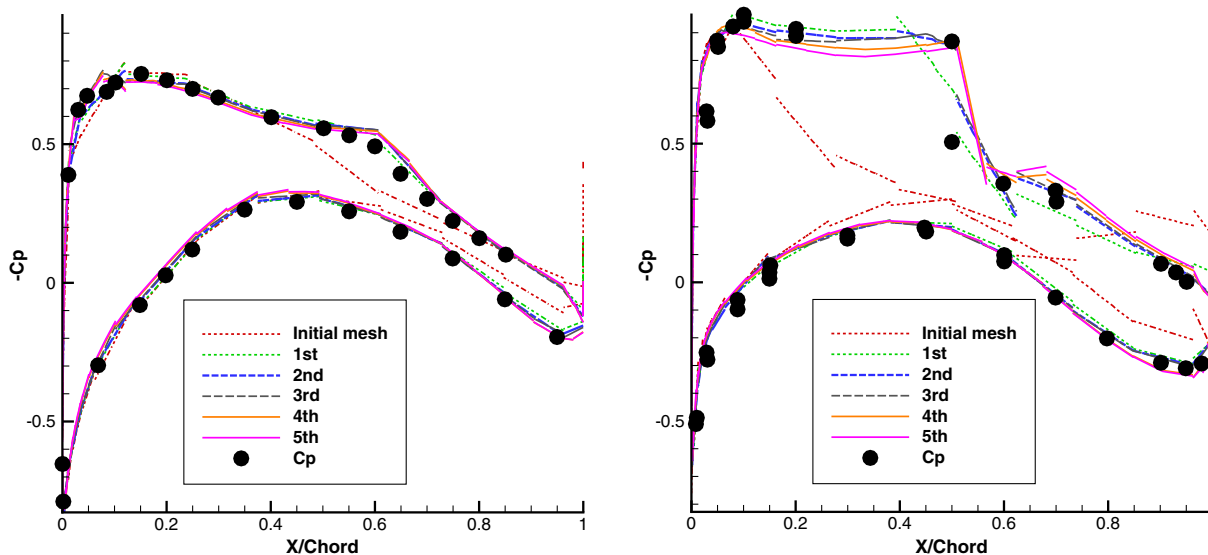
The mesh agglomeration algorithm used in this work receives as input a linear, multiblock mesh, where each block has a number of nodes that satisfies the following law:

$$N_{\text{node}} = (q \cdot N_i + 1) \cdot (q \cdot N_j + 1) \cdot (q \cdot N_k + 1) \quad (13)$$

where q is the desired polynomial order for the geometry representation, and $[N_i, N_j, N_k]$ are the number of q th order elements in the $[i, j, k]$ directions. Note that, with the previous rule for the number of nodes, geometric irregularities of order $g \leq q$ on the agglomerated mesh can only exist at element borders. The relevance of this observation is that the block boundaries in the linear mesh should coincide with the geometric irregularities of order $g \leq q$ in the source geometry. A mitigation to this problem is to generate locally finer meshes or to put block boundaries close to geometric irregularities. However, this is not always possible or easy to achieve.

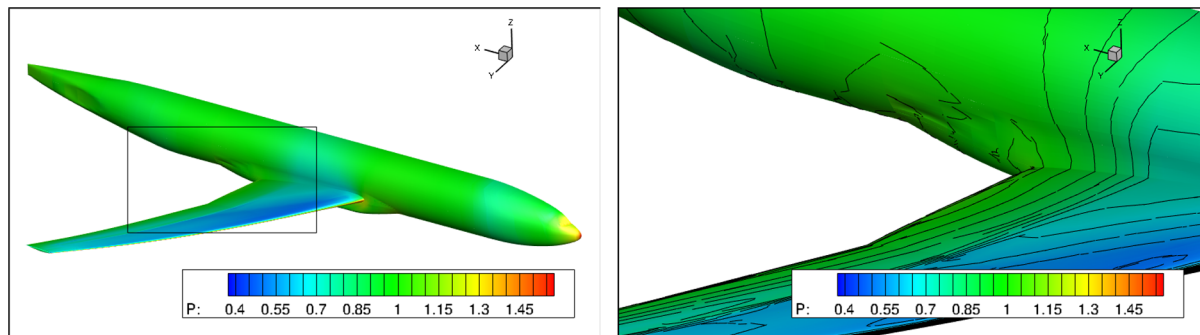
Another challenge is managing geometric irregularities as oscillations may occur in the high-order geometry representation. A way to address this problem is to adapt the surface elements of the initial mesh based on the integrated distance between the high-order elements' solid boundaries and the source geometry and to reproject the new nodes created in this adaptation onto the source geometry. This allows for explicit control over the geometry representation error with the caveat that we need to ensure volume positivity of the curved elements at the boundary.

[‡]Experimental data were digitized from the 5th Drag Prediction Workshop summary presentation.



a) Pressure coefficient at 13.06% of the reference span b) Pressure coefficient at 50.24% of the reference span

Fig. 12 CRM, wing-body geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$: comparison of pressure coefficient with experimental data. Note: the solution for the initial mesh is at $C_L = 0.387$.



a) Pressure contours on final adapted mesh b) Waviness on fairing-body junction.

Fig. 13 CRM, wing-body geometry, $M_\infty = 0.85$, $C_L = 0.5$, $Re_{MAC} = 5 \times 10^6$: waviness on fairing-body junction.

The irregularities identified on the CRM geometry are wing trailing edge, wing-fairing junction, cockpit-nose-cone junction, and fairing-body junction. In the case of our block topology for the initial mesh, a block boundary is not aligned with the aft portion of the fairing-body junction (Fig. 13). This causes oscillations in the geometry of the agglomerated mesh that affect the computed drag and the robustness of the error estimates. The slightly wavy geometry where $g \leq q$ is a possible source of larger drag in our results (Fig. 11a), as the geometry is held fixed throughout the adaptation.

VI. Conclusions

The use of an adaptive discontinuous finite-element method to predict drag was demonstrated. The two-dimensional results show that adjoint-based adaptation reduces the discretization error to acceptable levels and hence allows for quantification of the remaining error due to turbulence modeling and measurement errors. Furthermore, the increase in the number of degrees of freedom from the initial meshes to the adapted meshes is roughly 85%, which is much less than the increase due to one level of uniform refinement, which would quadruple the mesh size.

The drag prediction results for the CRM geometry show that adjoint-based mesh adaptation can significantly save degrees of freedom in comparison with the uniform refinement studies performed the DPW-V participants. The variational formulation of the discontinuous Galerkin method allows for straightforward output error estimation, but our current error estimation framework does not account for the effect of discretization error on the angle of attack that

yields the target lift. This suggests that the adaptation should be driven by a combination of drag and lift errors.

Finally, three-dimensional RANS simulations using discontinuous finite elements are still very challenging both in terms of robustness and in terms of computational expense. The results, however, are very promising, and further research in these topics will certainly be beneficial to future codes for aerodynamic performance prediction.

Acknowledgments

The authors acknowledge the support given by the University of Michigan to the development of this work and appreciate the computational resources provided by NASA, which were essential to obtaining the results presented here. Funding for this research was provided by the U.S. Air Force Office of Scientific Research under contract number FA9550-10-C-0040.

References

- [1] Buttazzo, G., Frediani, A., Allmaras, S. R., Bussolletti, J. E., Hilmes, C. L., Johnson, F. T., Melvin, R. G., Tinoco, E. N., Venkatakrishnan, V., Wigton, L. B., and Young, D. P., "Algorithm Issues and Challenges Associated with the Development of Robust CFD Codes," *Variational Analysis and Aerospace Engineering*, Vol. 33, Springer Optimization and Its Applications, Springer, New York, 2009, pp. 1–19.
- [2] Vassberg, J. C., DeHaan, M. A., and Sclafani, T. J., "Grid Generation Requirements for Accurate Drag Predictions Based on OVERFLOW Calculations," *16th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 2003-4124, June 2003.

- [3] Fidkowski, K. J., "A High-Order Discontinuous Galerkin Multigrid Solver for Aerodynamic Applications," M.S. Thesis, Massachusetts Inst. of Technology, Dept. of Aeronautics and Astronautics, June 2004.
- [4] Levy, D. W., Zickuhr, T., Vassberg, J., Agrawal, S., Wahls, R. A., Pirzadeh, S., and Hemsch, M. J., "Data Summary from the First AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 40, No. 5, 2003, pp. 875–882. doi:10.2514/2.6877
- [5] Laffin, K. R., Klausmeyer, S. M., Zickuhr, T., Vassberg, J. C., Wahls, R. A., Morrison, J. H., Brodersen, O. P., Rakowitz, M. E., Tinoco, E. N., and Godard, J.-L., "Data Summary from the Second AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 42, No. 5, 2005, pp. 1165–1178. doi:10.2514/1.10771
- [6] Vassberg, J. C., Tinoco, E. N., Mani, M., Brodersen, O. P., Einfeld, B., Wahls, R. A., Morrison, J. H., Zickuhr, T., Laffin, K. R., and Mavriplis, D. J., "Abridged Summary of the Third AIAA Computational Fluid Dynamics Drag Prediction Workshop," *Journal of Aircraft*, Vol. 45, No. 3, May 2008, pp. 781–798.
- [7] Vassberg, J. C., Tinoco, E. N., Mani, M., Rider, B., Zickuhr, T., Levy, D. W., Brodersen, O. P., Einfeld, B., Crippa, S., Wahls, R. A., Morrison, J. H., Mavriplis, D. J., and Murayama, M., "Summary of the Fourth AIAA CFD Drag Prediction Workshop," *28th AIAA Applied Aerodynamics Conference*, AIAA Paper 2010-4547, June–July 2010.
- [8] Fidkowski, K. J., and Darmofal, D. L., "Review of Output-Based Error Estimation and Mesh Adaptation in Computational Fluid Dynamics," *AIAA Journal*, Vol. 49, No. 4, 2011, pp. 673–694. doi:10.2514/1.J050073
- [9] 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. doi:10.1016/S0021-9991(03)00074-3
- [10] Park, M. A., "Adjoint-Based, Three-Dimensional Error Prediction and Grid Adaptation," *AIAA Journal*, Vol. 42, No. 9, 2004, pp. 1854–1862. doi:10.2514/1.10051
- [11] Hartmann, R., and Houston, P., "Goal-Oriented a Posteriori Error Estimation for Multiple Target Functionals," *Hyperbolic Problems: Theory, Numerics, Applications*, edited by Hou, T., and Tadmor, E., Springer-Verlag, Berlin, 2003, pp. 579–588.
- [12] Fidkowski, K. J., and Darmofal, D. L., "An Adaptive Simplex Cut-Cell Method for Discontinuous Galerkin Discretizations of the Navier–Stokes Equations," *18th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 2007-3941, June 2007.
- [13] Nemeč, M., Aftosmis, M. J., and Wintzer, M., "Adjoint-Based Adaptive Mesh Refinement for Complex Geometries," *46th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2008-0725, Jan. 2008.
- [14] Ceze, M., and Fidkowski, K. J., "Output-Driven Anisotropic Mesh Adaptation for Viscous Flows Using Discrete Choice Optimization," *48th AIAA Aerospace Sciences Meeting*, AIAA Paper 2010-0170, Jan. 2010.
- [15] Ceze, M., and Fidkowski, K. J., "Anisotropic hp -Adaptation Framework for Functional Prediction," *AIAA Journal*, Vol. 51, No. 2, Feb. 2012, pp. 492–509.
- [16] Fidkowski, K. J., and Roe, P. L., "An Entropy Adjoint Approach to Mesh Refinement," *SIAM Journal on Scientific Computing*, Vol. 32, No. 3, 2010, pp. 1261–1287. doi:10.1137/090759057
- [17] Fidkowski, K. J., and Luo, Y., "Output-Based Space-Time Mesh Adaptation for the Compressible Navier–Stokes Equations," *Journal of Computational Physics*, Vol. 230, No. 14, 2011, pp. 5753–5773. doi:10.1016/j.jcp.2011.03.059
- [18] Spalart, P. R., and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," *30th Aerospace Sciences Meeting and Exhibit*, AIAA Paper 1992-0439, Jan. 1992.
- [19] Oliver, T. A., "A High-Order, Adaptive, Discontinuous Galerkin Finite Element Method for the Reynolds-Averaged Navier–Stokes Equations," Ph.D. Dissertation, Massachusetts Inst. of Technology, Cambridge, MA, 2008.
- [20] Roe, P. L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372. doi:10.1016/0021-9991(81)90128-5
- [21] Bassi, F., and Rebay, S., "GMRES Discontinuous Galerkin Solution of the Compressible Navier–Stokes Equations," *Discontinuous Galerkin Methods: Theory, Computation and Applications*, edited by Cockburn, B., Karniadakis, G. E., and Shu, C.-W., Springer, Berlin, 2000, pp. 197–208.
- [22] Persson, P.-O., and Peraire, J., "Sub-Cell Shock Capturing for Discontinuous Galerkin Methods," *44th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2006-0112, Jan. 2006.
- [23] Ceze, M. D., "A Robust hp -Adaptation Method for Discontinuous Galerkin Discretizations Applied to Aerodynamic Flows," Ph.D. Thesis, Univ. of Michigan, Ann Arbor, MI, 2013.
- [24] Saad, Y., "A Flexible Inner-Outer Preconditioned GMRES Algorithm," *SIAM Journal on Scientific Computing*, Vol. 14, No. 2, 1993, pp. 461–469. doi:10.1137/0914028
- [25] Burgess, N. K., and Mavriplis, D. J., "Robust Computation of Turbulent Flows using a Discontinuous Galerkin Method," *50th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2012-0457, Jan. 2012.
- [26] Giles, M. B., Duta, M. C., Müller, J.-D., and Pierce, N. A., "Algorithm Developments for Discrete Adjoint Methods," *AIAA Journal*, Vol. 41, No. 2, 2003, pp. 198–205. doi:10.2514/2.1961
- [27] Mader, C. A., Martins, J. R., Alonso, J. J., and van der Weide, E., "ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers," *AIAA Journal*, Vol. 46, No. 4, 2008, pp. 863–873. doi:10.2514/1.29123
- [28] Rannacher, R., "Adaptive Galerkin Finite Element Methods for Partial Differential Equations," *Journal of Computational and Applied Mathematics*, Vol. 128, Nos. 1–2, 2001, pp. 205–233. doi:10.1016/S0377-0427(00)00513-6
- [29] Barth, T., and Larson, M., "A Posteriori Error Estimates for Higher Order Godunov Finite Volume Methods on Unstructured Meshes," *Finite Volumes for Complex Applications III*, edited by Herban, R., and Kröner, D., Hermes Penton, London, 2002, pp. 41–63.
- [30] Solin, P., and Demkowicz, L., "Goal-Oriented hp -Adaptivity for Elliptic Problems," *Computer Methods in Applied Mechanics and Engineering*, Vol. 193, Nos. 6–8, 2004, pp. 449–468. doi:10.1016/j.cma.2003.09.015
- [31] Lu, J., "An A Posteriori Error Control Framework for Adaptive Precision Optimization Using Discontinuous Galerkin Finite Element Method," Ph.D. Thesis, Massachusetts Inst. of Technology, Cambridge, MA, 2005.
- [32] Hartmann, R., "Adjoint Consistency Analysis of Discontinuous Galerkin Discretizations," *SIAM Journal on Numerical Analysis*, Vol. 45, No. 6, 2007, pp. 2671–2696. doi:10.1137/060665117
- [33] Venditti, 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. doi:10.1006/jcph.2001.6967
- [34] Becker, R., and Rannacher, R., "An Optimal Control Approach to A Posteriori Error Estimation in Finite Element Methods," *Acta Numerica*, edited by Iserles, A., Cambridge Univ. Press, Cambridge, England, U.K., 2001, pp. 1–102.
- [35] Hartmann, R., and Houston, P., "Adaptive Discontinuous Galerkin Finite Element Methods for the Compressible Euler Equations," *Journal of Computational Physics*, Vol. 183, No. 2, 2002, pp. 508–532. doi:10.1006/jcph.2002.7206
- [36] Giles, M. B., and Stili, E., "Adjoint Methods for PDEs: A Posteriori Error Analysis and Postprocessing by Duality," *Acta Numerica*, Vol. 11, 2002, pp. 145–236.
- [37] "Turbulence Modeling Resource," NASA, 2012, <http://turbmodels.larc.nasa.gov> [retrieved 21 May 2012].
- [38] Ladson, C. L., "Effects of Independent Variation of Mach and Reynolds Numbers on the Low-Speed Aerodynamic Characteristics of the NACA 0012 Airfoil Section," NASA TM-4074, 1988.
- [39] "CFL3D Home Page," NASA, 2012, <http://cfl3d.larc.nasa.gov> [retrieved 21 May 2012].
- [40] Giles, M. B., and Pierce, N. A., "Adjoint Equations in CFD: Duality, Boundary Conditions and Solution Behaviour," *13th Computational Fluid Dynamics Conference*, AIAA Paper 1997-1850, June–July 1997.
- [41] Vassberg, J. C., DeHaan, M. A., Rivers, M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," *26th AIAA Applied Aerodynamics Conference*, AIAA Paper 2008-6919, Aug. 2008.
- [42] Wang, Z., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, R., Hillewaert, K., Huynh, H., Kroll, N., May, G., Persson, P.-O., van Leer, B., and Visbal, M., "High-Order CFD Methods: Current Status and Perspective," *International Journal for Numerical Methods in Fluids*, Vol. 72, No. 8, 2013, pp. 811–845. doi:10.1002/ld.v72.8
- [43] Ceze, M., and Fidkowski, K. J., "A Robust Adaptive Solution Strategy for High-Order Implicit CFD Solvers," *20th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 2011-3696, June 2011.
- [44] Tinoco, E. N., Levy, D., and Brodersen, O. P., "DPW 5 Summary of Participant Data," AIAA, Reston, VA, June 2012, <http://aac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw> [retrieved 29 Oct. 2012].