31st AIAA Applied Aerodynamics Conference, June 24-27, 2013, San Diego, CA

Results from the Fifth AIAA Drag Prediction Workshop obtained with a parallel Newton-Krylov-Schur flow solver discretized using summation-by-parts operators

Michal Osusky,* Pieter D. Boom,† and David W. Zingg‡

Institute for Aerospace Studies, University of Toronto, Toronto, Ontario, M3H 5T6, Canada

We present the solution of the test cases from the Fifth AIAA Drag Prediction Workshop computed with a novel Newton-Krylov-Schur parallel solution algorithm for the Reynolds-Averaged Navier-Stokes equations coupled with the Spalart-Allmaras one-equation turbulence model. The algorithm employs summation-by-parts operators on multi-block structured grids, while simultaneous approximation terms are used to enforce boundary conditions and coupling at block interfaces. Two-dimensional verification and validation cases highlight the correspondence of the current algorithm to established flow solvers as well as experimental data. The common grid study, using grids with up to 150 million nodes around the NASA Common Research Model wing-body configuration, demonstrates the parallel computation capabilities of the current algorithm, while the buffet study demonstrates the ability of the solver to compute flow with substantial recirculation regions and flow separation. The use of quadratic constitutive relations to modify the Boussinesq approximation is shown to have a significant impact on the recirculation patterns observed at higher angles of attack. The algorithm is capable of efficiently and accurately calculating complex three-dimensional flows over a range of flow conditions, with results consistent with those of well-established flow solvers using the same turbulence model.

I. Introduction

Recent advances in computer architecture and numerical methods have paved the way for massively parallel computations. Leveraging the ever-increasing access to supercomputer clusters with in excess of tens of thousands of processors, large-scale CFD simulations are becoming more suitable for dealing with problems of practical interest. However, accurate simulations of such flows necessitate the solution of very large problems, with the results from the AIAA Drag Prediction Workshop series[1] indicating that grids with over O(10^8) grid nodes are required for grid-converged lift and drag values for flows over a wing-body configuration. Following this trend, algorithms that scale well with thousands of processors are required to make efficient use of computational resources.

We have previously presented an efficient and robust Newton-Krylov-Schur flow solution algorithm, discretized using summation-by-parts operators and simultaneous approximation terms.[2] Extensive verification and validation tests were conducted to demonstrate the accuracy of the solution approach. In particular, the two-dimensional turbulence model verification and validation cases provided on the NASA Turbulence Modeling Resource website[3] have served as an invaluable tool in comparing the accuracy of the solver against established algorithms such as CFL3D[4] and FUN3D.[5] The excellent parallel scaling characteristics of the algorithm were demonstrated with over 6000 processors, making it well suited for the solution of the large

*PhD Candidate, AIAA Student Member

†PhD Candidate, AIAA Student Member

‡Professor and Director, Tier 1 Canada Research Chair in Computational Aerodynamics and Environmentally-Friendly Aircraft Design, J. Armand Bombardier Foundation Chair in Aerospace Flight, Associate Fellow AIAA

problems required for obtaining grid-converged force coefficients. The algorithm has also been applied to aerodynamic shape optimization by Osusky and Zingg,[6,7] demonstrating its effectiveness for three-dimensional aerodynamic shape optimization in the turbulent flow regime.

While simple test cases are an important part of the development process, a means of evaluating an algorithm for cases of practical interest is required. This opportunity is provided by the AIAA Drag Prediction Workshop (DPW) series, which invites participants from research institutes, industry, and academia to apply their algorithms to the solution of complex CFD problems. The first workshop, summarized by Levy *et al.*,[8] involved the solution of the DLR-F4 wing-body geometry; large discrepancies were observed between the solutions produced by the 18 participants. The second and third workshops, using the DLR-F6 geometry, resulted in similar observations. The large number of grids used by the various solution algorithms was thought to contribute to the scatter in results. The third workshop also introduced wing-alone configurations in an attempt to provide a simpler geometry with well-understood flow features. The fourth workshop made use of the NASA Common Research Model wing-body geometry,[9] designed to reduce the amount of side-of-body separation that was thought to be the cause of much of the scatter in the results for the DLR geometries. The results of this workshop are summarized by Vassberg *et al.*[10] Less scatter was observed in the results, but significant variance was still seen, especially between algorithms that make use of different turbulence models. The fifth workshop introduced a common grid family to be used by all participants,[11] in an effort to further reduce the variance in the solutions. The main focus of this paper is to present the results of our participation in the Fifth AIAA Drag Prediction Workshop (DPW5), as well as the algorithm modifications implemented based on DPW5 results. The cases demonstrate the efficiency with which the current algorithm can obtain accurate solutions for complex three-dimensional flows and the consistency of the results obtained with those of other solvers, especially those using the same turbulence model. The comparisons included provide excellent verification and validation of our solver.

The complexity of the cases studied as part of DPW5 provide an opportunity for the study of turbulence model applicability. In particular, the effect of the off-wall distance calculation and the assumptions made therein can have a substantial impact on the steady-state solution. This paper will demonstrate the use of a surface-node based calculation, as opposed to an off-wall distance calculation based on a full representation of the surface boundary. In particular, the accuracy of off-wall distance values is critical near the solid boundary, where the largest turbulence model source terms occur. Additionally, results from the fourth and fifth workshops have shown that the use of the Spalart-Allmaras one-equation turbulence model can result in large recirculation features at higher angles of attack, while more complex turbulence models do not produce this flow feature. The quadratic constitutive relations (QCR) presented by Spalart[12] provide an inexpensive modification to the use of turbulence models that may expand the capabilities of relatively simple models, especially for flows with substantial recirculation.

The paper is divided into the following sections. Section II presents a brief overview of the governing equations and the flow solution algorithm. Section III presents the cases that were part of DPW5, along with a summary of the geometry and computational grids used. Section IV contains results obtained with the current algorithm for steady flow solutions around the NASA Common Research Model (CRM) geometry. Conclusions are given in Section V. The Appendix contains results obtained for the optional turbulence model verification cases.

II. Flow Solution Algorithm

A. Governing Equations

1. The Navier-Stokes Equations

The three-dimensional Navier-Stokes equations, after the coordinate transformation $(x, y, z) \to (\xi, \eta, \zeta)$, are given by

$$\partial_t \hat{\mathbf{Q}} + \partial_\xi \hat{\mathbf{E}} + \partial_\eta \hat{\mathbf{F}} + \partial_\zeta \hat{\mathbf{G}} = \frac{1}{Re}\Big(\partial_\xi \hat{\mathbf{E}}_\mathbf{v} + \partial_\eta \hat{\mathbf{F}}_\mathbf{v} + \partial_\zeta \hat{\mathbf{G}}_\mathbf{v}\Big), \tag{1}$$

with

$$\hat{\mathbf{Q}} = J^{-1}\mathbf{Q},$$
$$\hat{\mathbf{E}} = J^{-1}\Big(\xi_x\mathbf{E} + \xi_y\mathbf{F} + \xi_z\mathbf{G}\Big), \quad \hat{\mathbf{E}}_\mathbf{v} = J^{-1}\Big(\xi_x\mathbf{E}_\mathbf{v} + \xi_y\mathbf{F}_\mathbf{v} + \xi_z\mathbf{G}_\mathbf{v}\Big),$$
$$\hat{\mathbf{F}} = J^{-1}\Big(\eta_x\mathbf{E} + \eta_y\mathbf{F} + \eta_z\mathbf{G}\Big), \quad \hat{\mathbf{F}}_\mathbf{v} = J^{-1}\Big(\eta_x\mathbf{E}_\mathbf{v} + \eta_y\mathbf{F}_\mathbf{v} + \eta_z\mathbf{G}_\mathbf{v}\Big),$$
$$\hat{\mathbf{G}} = J^{-1}\Big(\zeta_x\mathbf{E} + \zeta_y\mathbf{F} + \zeta_z\mathbf{G}\Big), \quad \hat{\mathbf{G}}_\mathbf{v} = J^{-1}\Big(\zeta_x\mathbf{E}_\mathbf{v} + \zeta_y\mathbf{F}_\mathbf{v} + \zeta_z\mathbf{G}_\mathbf{v}\Big),$$

where $Re = \frac{\rho_\infty a_\infty l}{\mu_\infty}$, ρ is the density, a the sound speed, l the chord length, μ the viscosity, and J the metric Jacobian that results from the coordinate transformation. The '∞' subscript denotes a free-stream value for the given quantity. The notation ξ_x, for example, is a shorthand form of $\partial_x\xi = \frac{\partial\xi}{\partial x}$. The conservative variables and inviscid and viscous fluxes are given by

$$\mathbf{Q} = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ e \end{bmatrix}, \quad \mathbf{E} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uw \\ u(e+p) \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho vw \\ v(e+p) \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} \rho v \\ \rho uw \\ \rho vw \\ \rho w^2 + p \\ w(e+p) \end{bmatrix},$$
$$\mathbf{E}_\mathbf{v} = \begin{bmatrix} 0 \\ \tau_{xx} \\ \tau_{xy} \\ \tau_{xz} \\ E_{v,5} \end{bmatrix}, \quad \mathbf{F}_\mathbf{v} = \begin{bmatrix} 0 \\ \tau_{yx} \\ \tau_{yy} \\ \tau_{yz} \\ F_{v,5} \end{bmatrix}, \quad \mathbf{G}_\mathbf{v} = \begin{bmatrix} 0 \\ \tau_{zx} \\ \tau_{zy} \\ \tau_{zz} \\ G_{v,5} \end{bmatrix},$$

where e is the energy, p the pressure, u, v, and w the Cartesian velocity components, and τ the Newtonian stress tensor. The preceding variables have been made dimensionless by the use of the free-stream values of density and sound speed, as well as chord length. Laminar viscosity is calculated using a dimensionless form of Sutherland's law:[13]

$$\mu = \frac{a^3(1 + S^*/T_\infty)}{a^2 + S^*/T_\infty}, \tag{2}$$

where $S^* = 198.6°\text{R}$, and T_∞ is typically set to 460°R.

2. Spalart-Allmaras One-Equation Turbulence Model

In order to solve turbulent flows, a turbulent, or eddy, viscosity, μ_t, can be added to the viscosity, μ, based on the Boussinesq approximation.[14] The Spalart-Allmaras one-equation turbulence model[15] is used to compute the turbulent viscosity. The model solves a transport equation for a turbulence variable, $\tilde{\nu}$, that is related to turbulent viscosity. The model itself is a sixth equation that is solved concurrently with the five Navier-Stokes equations in a fully-coupled implicit manner.

The version of the model used in this work, labeled "SA" on the NASA Turbulence Modeling Resource (TMR) website,[3] is given by

$$\begin{aligned} \frac{\partial\tilde{\nu}}{\partial t} + u_i\frac{\partial\tilde{\nu}}{\partial x_i} = \frac{c_{b1}}{Re}[1 - f_{t2}]\tilde{S}\tilde{\nu} + \frac{1 + c_{b2}}{\sigma_\mathrm{t} Re}\nabla\cdot[(\nu + \tilde{\nu})\nabla\tilde{\nu}] - \frac{c_{b2}}{\sigma_\mathrm{t} Re}(\nu + \tilde{\nu})\nabla^2\tilde{\nu} \\ - \frac{1}{Re}\left[c_{w1}f_w - \frac{c_{b1}}{\kappa^2}f_{t2}\right]\left(\frac{\tilde{\nu}}{d}\right)^2, \end{aligned} \tag{3}$$

where $\nu = \frac{\mu}{\rho}$.

The spatial derivatives on the left side of the equation represent advection. The first term on the right side represents production, while the second and third terms account for diffusion. The fourth term represents destruction. Reference 15 contains the values of all of the terms and constants that appear in (3). Importantly, precautions are taken to ensure that neither the vorticity, S, nor the vorticity-like term, $\tilde{S}$, approaches zero or becomes negative, which could lead to numerical problems. The off-wall distance at each computational node is denoted by d.

3. Quadratic Constitutive Relations

Originally motivated by the simulation of secondary flows in square ducts, Spalart[12] presented a modification to the Reynolds stress tensor that results from the Boussinesq approximation, termed as a quadratic constitutive relation (QCR). A nonlinear stress model is proposed as

$$\tau_{ij} = \bar{\tau}_{ij} - c_{nl1}\left(O_{ik}\bar{\tau}_{jk} + O_{jk}\bar{\tau}_{ik}\right), \tag{4}$$

where

$$\begin{aligned} O_{ik} &\equiv \frac{\partial_k u_i - \partial_i u_k}{\sqrt{\partial_n u_m \partial_n u_m}}, \\ \partial_n u_m \partial_n u_m &= (\partial_x u)^2 + (\partial_x v)^2 + (\partial_x w)^2 \\ &\quad + (\partial_y u)^2 + (\partial_y v)^2 + (\partial_y w)^2 \\ &\quad + (\partial_z u)^2 + (\partial_z v)^2 + (\partial_z w)^2, \end{aligned}$$

u_i are the velocity components (u, v, w), and $c_{nl1} = 0.3$. Additionally, $\bar{\tau}_{ij}$ is the Reynolds stress tensor obtained from the Boussinesq approximation. The implementation of the above modification does not add substantial development time and has negligible effect on the convergence characteristics of the algorithm,

The use of the above modification to the Reynolds stress tensor has seen limited adoption. However, some results presented during the fourth and fifth DPW have incorporated QCR,[10,16,17] with positive effects in the solutions presented. In particular, the use of QCR has a substantial impact on the size of the recirculation region at the wing-body junction once a critical angle of attack is reached.

B. Solution Methodology

The following sections present a brief overview of the solution algorithm developed for this work, called DIABLO, in terms of both the spatial discretization method used and the overall solution strategy. Further details of the algorithm can be found in references 18, 19 and 20, with a complete description of the RANS solution methodology of the current algorithm presented in reference 2.

1. Spatial Discretization

The spatial discretization of the Navier-Stokes equations and the turbulence model is obtained by the use of second-order Summation-By-Parts (SBP) operators, while inter-block coupling and boundary conditions are enforced weakly by the use of Simultaneous Approximation Terms (SATs). We decompose our domains into multiple blocks, resulting in multi-block structured grids. Not only does this type of blocking strategy work well with the SBP-SAT discretization, it also allows for relatively straightforward creation of meshes around complex geometries.

The SBP-SAT approach requires less information from adjoining blocks in order to obtain a discretization of the governing equations at block interfaces. This results in a reduced requirement for information sharing between blocks, which is especially advantageous for parallel algorithms, reducing the time spent in communication. Additionally, the fact that this discretization does not need to form any derivatives across interfaces reduces the continuity requirements for meshes at interfaces. In fact, only C^0 continuity is necessary, allowing the algorithm to provide accurate solutions even on grids with relatively high incidence angles for grid lines at interfaces, an example of which can be seen in Figure 1. This feature substantially reduces the burden placed on the grid generation process, especially for complex three-dimensional geometries, such as the NASA CRM.

Numerical dissipation is added using either the scalar dissipation model developed by Jameson *et al.*[21] and later refined by Pulliam,[22] or the matrix dissipation model of Swanson and Turkel.[23]

2. Iteration to Steady State

Applying the SBP-SAT discretization to the governing equations results in a large system of nonlinear equations:

$$\mathcal{R}(\mathcal{Q}) = 0, \tag{5}$$

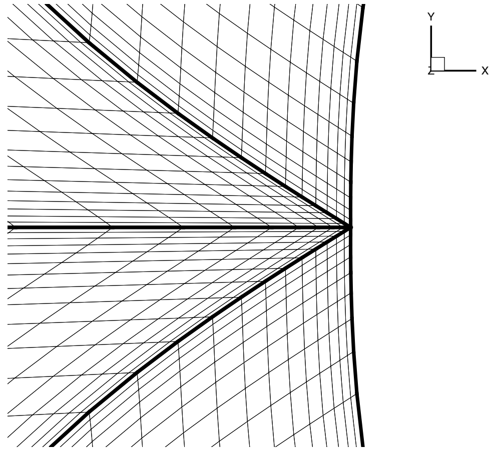

Figure 1. Blocks and grid-lines at leading edge of wing in an H-H topology grid with C^0 continuity at block interfaces

where $\mathcal{Q}$ represents the complete solution vector. When time-marched with the implicit Euler time-marching method and a local time linearization, this system results in a large set of linear equations of the form:[24]

$$\left(\frac{\mathcal{I}}{\Delta t} + \mathcal{A}^{(n)}\right) \Delta\mathcal{Q}^{(n)} = -\mathcal{R}^{(n)}, \tag{6}$$

where n is the outer (nonlinear) iteration index, Δt is the time step, $\mathcal{I}$ is the identity matrix, $\mathcal{R}^{(n)} = \mathcal{R}(\mathcal{Q}^{(n)})$, $\Delta\mathcal{Q}^{(n)} = \mathcal{Q}^{(n+1)} - \mathcal{Q}^{(n)}$, and

$$\mathcal{A}^{(n)} = \frac{\partial \mathcal{R}^{(n)}}{\partial \mathcal{Q}^{(n)}}$$

is the Jacobian.

In the infinite time step limit, the above describes Newton's method and will converge quadratically if a suitable initial iterate, $\mathcal{Q}^{(0)}$, is known. This initial iterate must be sufficiently close to the solution of (5). Since it is unlikely that any initial guess made for a steady-state solution will satisfy this requirement, the present algorithm makes use of a start-up phase whose purpose is to find a suitable initial iterate.

To solve the large set of discrete equations, we use a parallel Newton-Krylov algorithm using the flexible GMRES Krylov subspace iterative method with the approximate-Schur preconditioner.[25] Pseudo-transient continuation is used to march the solution from an initial guess to steady state, employing a spatially varying time step.

The general framework of the algorithm has also been leveraged to create an efficient time-accurate solution strategy.[26]

C. Off-Wall Distance Calculation

A key component affecting the accuracy of the Spalart-Allmaras turbulence model is the value of the off-wall distance, d, which represents the location of the closest solid surface to any node in the grid. This quantity has a direct impact on the source terms of the model, influencing the rate at which the turbulence quantity, $\tilde{\nu}$, is produced or destroyed, especially in the near-wall regions of the flow.

Several methods for calculating d exist, ranging from brute-force searches to the solution of differential equations.[27,28] Depending on the complexity of the geometry and the blocking strategy for the mesh, the calculation of d can present a substantial computational expense. Methods have also been developed to speed up the calculation of off-wall distances, relying on a priori knowledge of the grid topology; for example, some algorithms follow grid lines emanating from the surface and use the chord lengths along those lines in computing the off-wall distance of any nodes located on the grid lines. This approach has several disadvantages, foremost of which is the inability to easily generalize to multi-block and unstructured grids. While not hindered by assumptions about grid topology or geometry complexity, many differential approaches provide increasingly inaccurate values of d as one moves away from the surface. The extent to

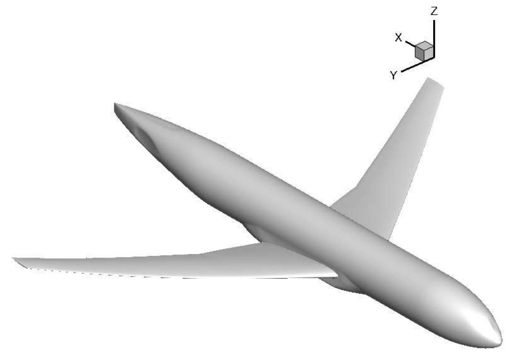

Figure 2. NASA CRM wing-body geometry used in DPW5

which this may affect the accuracy of the turbulence model depends on the specifics of the implementation and the grid. These approaches also require the solution of a global system of equations.

The current algorithm makes use of two search approaches. The first approach finds the distance from each internal node to the closest surface node (referred to as "method 1"), while the second takes into account distances to grid edges and surfaces on the solid boundary, in addition to the surface nodes ("method 2"). As will be demonstrated, using the closest surface node can have a significant impact on the computed flow solution. This is particularly noticeable on grids with non-orthogonal grid lines.

The calculation of d can become problematic for parallel algorithms, especially when the solid surface boundaries are located on blocks assigned to non-local processes. In the current implementation, all solid surface information is first collected into an array that is distributed to all processes, ensuring that all processes, regardless of whether or not they contain solid boundaries, are aware of the entire surface geometry. This framework lends itself easily to efficient parallel computations, as each process is responsible for calculating the the value of d only for local nodes. By ensuring that the number of nodes per processor remains approximately the same regardless of overall grid size, which is a necessary requirement for implicit algorithms due to memory requirements, the calculation of d scales with the total number of surface nodes, a value much smaller than the total number of nodes, resulting in fast off-wall distance computations for very large grids.

Both of the above approaches have been tested as part of the current algorithm. Their parallel nature results in fast calculation of off-wall distances for all nodes in the computational domain. For example, the recommended approach, taking into account solid boundary nodes, edges, and surfaces, can compute the value of d on a 35 million node grid in approximately 300 seconds using 128 processors. This is a small fraction of the overall flow solution time.

III. Test Cases for the Fifth Drag Prediction Workshop

This section presents the details of the test cases that are part of DPW5, along with the geometry specification and computational grid details. We focus on the mandatory cases involving the NASA CRM geometry.

A. Geometry

The focus of the required DPW5 cases was computation of flow around the NASA CRM wing-body geometry, which represents a state-of-the-art transonic aircraft configuration. For the purposes of the workshop, the model variant without a horizontal tail was considered, as shown in Figure 2. The configuration has a reference area, S_{ref}, of 594,720 in^2 and a reference chord, c_{ref}, of 275.80 in. The moment centre for the geometry is specified as ($x_{\mathrm{ref}} = 1325.90$ in, $z_{\mathrm{ref}} = 177.95$ in). All other model specifications can be found in the original reference.[9]

Table 1. **Grid characteristics of DPW5 common grid family (subdivided blocks)**

Grid level	Grid nodes, N	Blocks	Off-wall distance, ($\times c_{\mathrm{ref}}$)
T - tiny	746,200	88	2.4×10^{-6}
C - coarse	2,647,232	704	1.6×10^{-6}
M - medium	5,969,600	704	1.2×10^{-6}
F - fine	19,150,016	704	6.5×10^{-7}
X - extra-fine	44,239,552	704	3.9×10^{-7}
S - super-fine	146,284,992	960	1.9×10^{-7}

B. Computational Grids

In order to minimize the solution variability due to the use of customized grids in each solution algorithm, the organizers of DPW5 provided a set of common grids to all participants. Of note, the same grid nodes were used to create both structured and unstructured grids. The original grid family was constructed using a 2-to-3-to-4 node refinement strategy, as described by Vassberg.[11] All grids have been nondimensionalized using c_{ref}. The six grid levels effectively comprise two nested grid families, with the odd and even grid levels following the typical 1-to-2 node refinement strategy.

We use the O-O topology multi-block structured variant of the common grids provided for the workshop with minor modifications. These consisted of the subdivision of the original 5-block grids into larger numbers of blocks to facilitate the use of our parallel solution algorithm. A summary of the grid details can be seen in Table 1. It should be stressed that no points were removed from the original grids during the subdivision process; interfaces were simply introduced into select blocks until an adequate load distribution was achieved. The grids were not load balanced.

The off-wall spacing value range of 1.9×10^{-7} to 2.4×10^{-6} reference units corresponds to average y^+ values of 0.33 to 2.0.

C. Test Cases

As part of DPW5, participants are required to complete two mandatory cases. Additionally, an optional set of cases from the NASA TMR website is provided for the purpose of verification and validation of turbulence model implementations with simpler two-dimensional flows. We focus here on the mandatory cases involving the NASA CRM wing-body geometry. Results for the two-dimensional cases, including flow over a flat plate, flow over a bump-in-channel configuration, and flow over the NACA0012 airfoil, are provided in the Appendix.

Case 1 consists of a grid convergence study on the common family of grids. The flow conditions specify a Mach number of 0.85 and a lift coefficient, C_L, of 0.500($\pm$0.001). Additionally, a Reynolds number of 5 million (based on the reference length) is to be used, along with a reference temperature, T_{ref}, of 100°F. As part of this case, the grid convergence behaviour of force coefficients, as well as the angle of attack required to achieve the specified C_L, is reported.

Case 2 involves a buffet study at the Mach number and Reynolds number specified in Case 1. Only the medium grid level is used, and the angle of attack, α, is varied between 2.50° and 4.00° in 0.25° increments. Flow separation patterns and the angle of attack at which they occur are of interest.

All test cases are run assuming fully turbulent flow.

IV. Results

This section presents the results obtained for the DPW5 test cases using the flow solver DIABLO. Included are the common grid study (Case 1) and the buffet study (Case 2). For both cases, we present results comparing the use of the more accurate off-wall distance calculation against the original surface-node-based calculation method. The use of the QCR modification is also explored. Both additional data sets were motivated by discussions with participants at the workshop and the results presented in the workshop summary paper.[17]

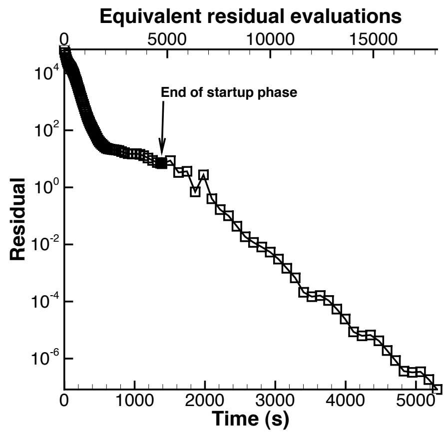

Figure 3. Convergence history for CRM flow on 19 million node grid (with 704 processors)

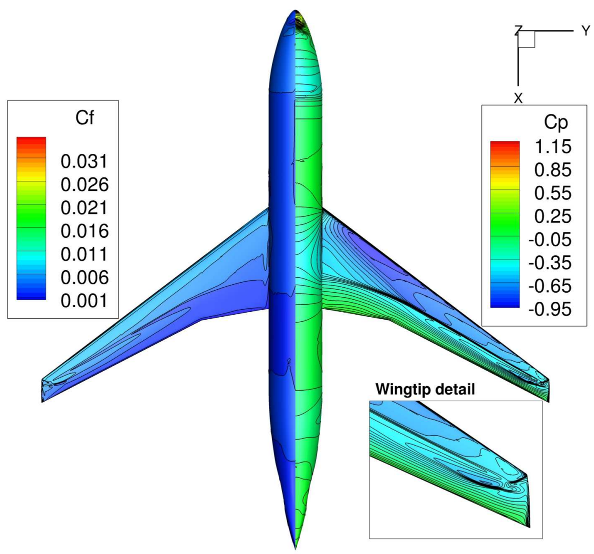

Figure 4. Surface coefficients for CRM flow

A. Case 1: Common Grid Study

Solutions were obtained on all six grid levels, converging the residual by 10 orders of magnitude on each grid. Both scalar and matrix dissipation are examined. Figure 3 presents the convergence history for the 19 million node fine ("F") grid level with matrix dissipation. This solution was computed with 704 processors, which were able to converge the solution in 70 minutes. A load balancing approach, as detailed by Apponsah and Zingg,[29] could be used to substantially improve the computational efficiency of the flow solves, either by reducing the time required to compute the solution using the same computational resources, or by obtaining the solution in the same amount of time using roughly 50% of the processors. Both stages of the solution algorithm, the globalization (approximate-Newton) phase and the inexact-Newton phase, are clearly visible in Figure 3. Figure 4 shows the contours of the coefficient of pressure, C_p, and coefficient of skin friction, C_f, on the top surface of the geometry. The C_p contours indicate the presence of a shock on the top surface of the wing, with a complex interplay of two shocks near the wingtip.

This case allows us to observe the grid convergence of the algorithm. In particular, the effect of the two off-wall distance calculation approaches and the use of QCR are of interest.

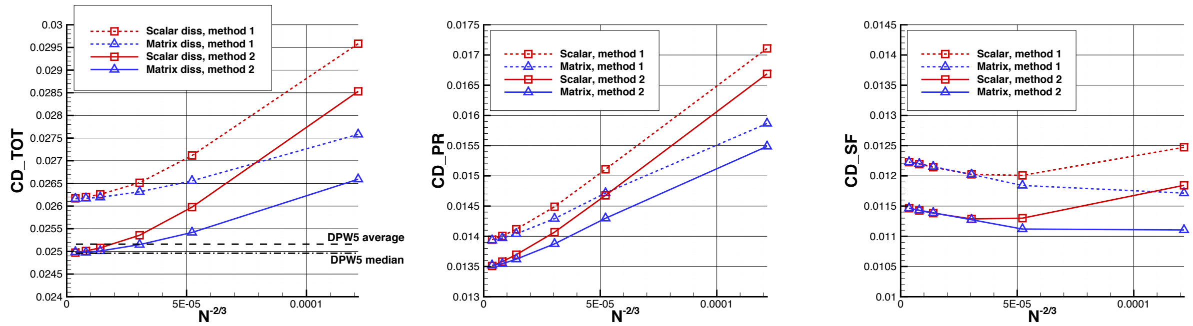

(a) C_D (b) C_{Dp} (c) C_{Df}

Figure 5. Grid convergence of C_D for Case 1 solutions (effect of off-wall distance)

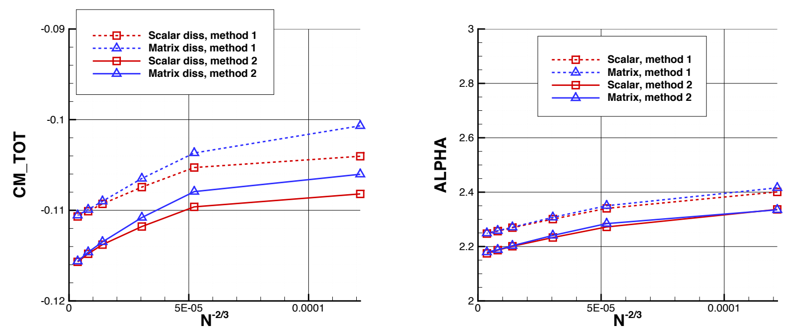

(a) C_M (b) α

Figure 6. Grid convergence of C_M and α for Case 1 solutions (effect of off-wall distance)

1. Effect of off-wall distance calculation method

Following participation in DPW5, it was observed that the grid-converged results for the coefficient of drag, C_D, obtained in DIABLO (0.02616) fell roughly 10 drag counts above the average and median values obtained by all participants, 0.02516 and 0.02496, respectively. From ensuing discussions, a possible source of this discrepancy was identified to be the off-wall distance calculation method. Instead of using the complete surface definition, the original calculation employed in DIABLO only made use of surface nodes (method 1). After implementing a more accurate off-wall distance calculation method (method 2), an updated set of results was obtained. As can be seen in Figure 5, using method 2 in the solution process has a substantial impact on the values of the force coefficients. While the grid convergence trends remain unaffected for both dissipation models, the more accurate distance calculation tends towards a grid-converged C_D of 0.02498, approximately 10 drag counts lower than the original result, and now corresponding to the median value for all workshop participants. The decrease in the drag coefficient appears to be grid independent. Both skin friction and pressure drag are affected. The moment coefficient, C_M, and the α required to achieve the specified lift on each grid level exhibit similar changes with the two wall-distance calculation methods, as seen in Figure 6. For both data sets, the matrix dissipation model produces a flatter curve for drag convergence. This signifies that, as expected, the model is more accurate on coarser grids. Additionally, a second-order algorithm will tend toward a straight line between three consecutive grid levels when force coefficients are plotted versus $N^{-2/3}$. This behavior can be observed for the finer grid levels. Using Richardson extrapolation[30] on the three finest grid levels, the order of convergence for C_D is calculated to be 1.60 and 3.32 for the scalar and matrix artificial dissipation models, respectively. Both dissipation models tend towards the same result as the grid is refined for each data set.

In order to understand why a change in the off-wall distance calculation method has such a significant impact on the steady-state force coefficients, we investigated the regions where the two calculation methods

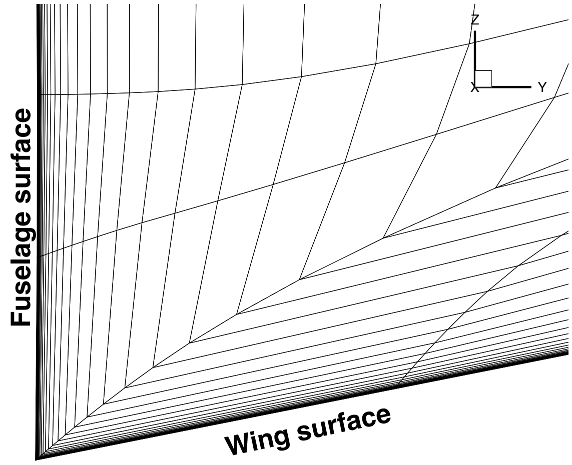

Figure 7. High grid incidence angles at wing-body junction (medium grid level)

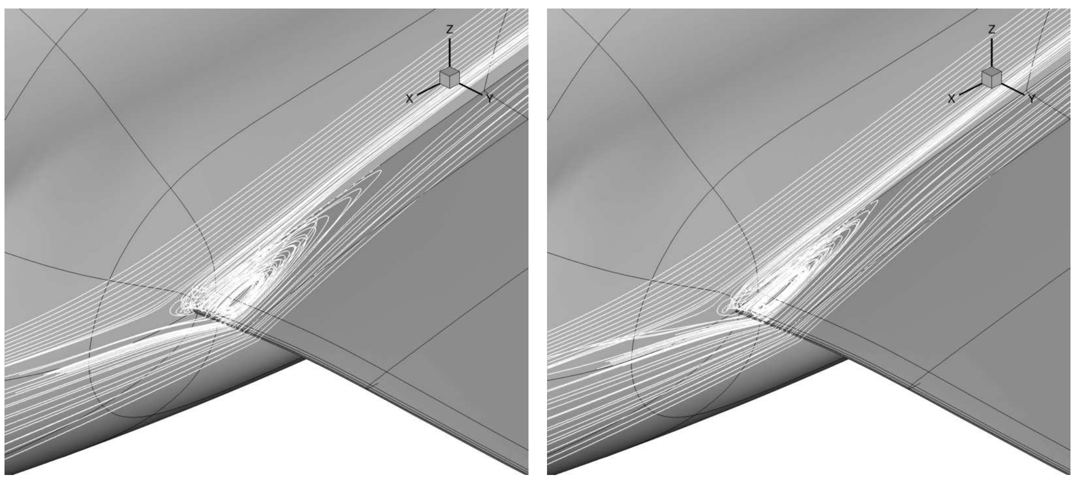

(a) Method 1

(b) Method 2

Figure 8. Recirculation bubble formed at trailing edge wing-body junction (effect of off-wall distance, scalar dissipation, medium grid level)

produce different values for d. By performing the off-wall distance calculation using both methods, and comparing the values of d near the surface of the body, where they have the most significant impact on the magnitudes of the turbulence model source terms, it was observed that the largest relative discrepancies occurred in the region where the trailing edge of the wing meets the fuselage. For example, on the medium grid level, method 1 identified the closest interior node to the solid surface to be at a distance of 5.26×10^{-6} reference units, while the more accurate method 2 calculated the distance to be 1.21×10^{-6} reference units, nearly five times smaller. The cause of this discrepancy can be seen in Figure 7; the high incidence angle (relative to the surface) of the grid lines emanating from the wing-body junction causes method 1, which relies on surface node coordinates only, to produce a larger than expected value of d for interior nodes closest to the surface boundary. The effect this discrepancy has on the values of force coefficients is exacerbated by the complex flow features found in the area of the wing-body junction. In particular, a recirculation region is present near the trailing edge of the wing, as shown in Figure 8. It can be seen that the two off-wall distance calculation methods produce recirculation bubbles of different sizes, even though all other algorithm parameters are identical.

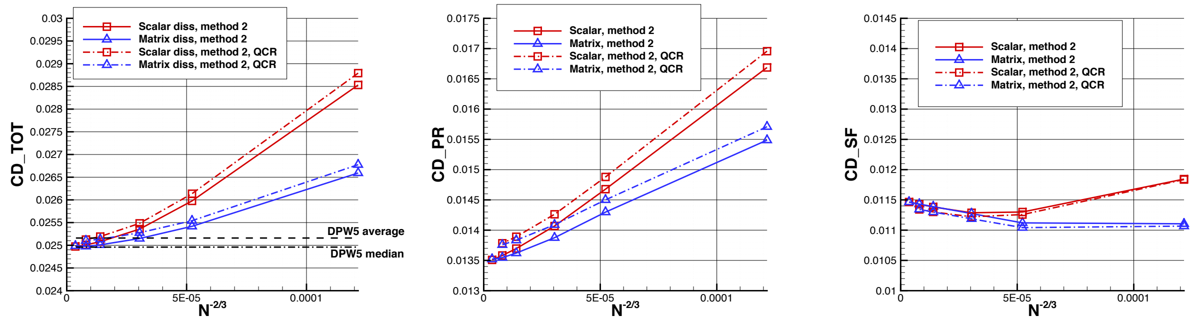

(a) C_D (b) $C_{D,p}$ (c) $C_{D,sf}$

Figure 9. Grid convergence of C_D for Case 1 solutions (effect of QCR)

Based on the above results, it is obvious that even small changes to the values of d in the computational domain can have a substantial impact on the steady-state solution obtained with the Spalart-Allmaras turbulence model. For Case 1, this is especially noticeable when complex flow features, such as the wing-body recirculation bubble, occur in regions of inaccurate d values. For this reason, the use of an accurate off-wall distance calculation method is critical for obtaining accurate force coefficient values on grids with non-orthogonal grid lines close to the surface.

2. Effect of QCR

Previous workshop results[16] have shown QCR to have a significant impact on the size of the recirculation region in the vicinity of the wing-body junction, especially at higher angles of attack. Since the recirculation bubble formed at the trailing edge of the wing-body junction, which has an impact on the overall solution, has been shown to vary with a simple change in the off-wall distance calculation method, the effect the use of QCR has on the flow solution is of interest.

Flow solutions are obtained on all but the finest grid level, with a visible difference in C_D. Figure 9 presents the values obtained with and without the use of QCR, using method 2 for the off-wall distance calculation. The grid-converged value of C_D increases by roughly 1 drag count when QCR is active ($C_D = 0.02509$), and still lies between the average and median values for all participants in DPW5. Interestingly, the pressure contribution to C_D increases by 2 drag counts with QCR active, while skin friction is decreased by 1 drag count.

As expected, Figure 10 shows that the use of QCR produces a visible change in the recirculation pattern found at the wing-body junction. With QCR active, the recirculation region is reduced on the medium grid level. Similar changes are observed for all grid levels in the common grid family.

B. Case 2: Buffet Study

The buffet study makes use of the medium grid level from the common grid family and is run with the scalar dissipation model. In addition to the required angles of attack, several additional angles were considered to obtain a more detailed understanding of the flow features present in certain regions of the α range. This section contains the results obtained using both off-wall distance calculation methods, as well as the use of QCR.

Figure 11 presents the lift, drag, and moment coefficients for all angles of attack. Of particular interest, a significant discontinuity can be seen for the force coefficients obtained using the algorithm without QCR, with a substantial drop in C_L. The angle of attack at which the discontinuity occurs is increased slightly with the use of the more accurate off-wall distance calculation method. Once the discontinuity is reached, the time-stepping parameters used in the algorithm have to be decreased in order to obtain the same residual reduction for all cases, adding significant computational expense. When QCR is used, the algorithm converges without difficulty even at higher angles of attack, and no discontinuity is present in the force coefficient data; in fact, the results obtained with QCR correspond extremely well with the majority of the DPW5 participants. Of the over 50 data sets submitted for Case 2, 21 exhibited a lift-break starting as early

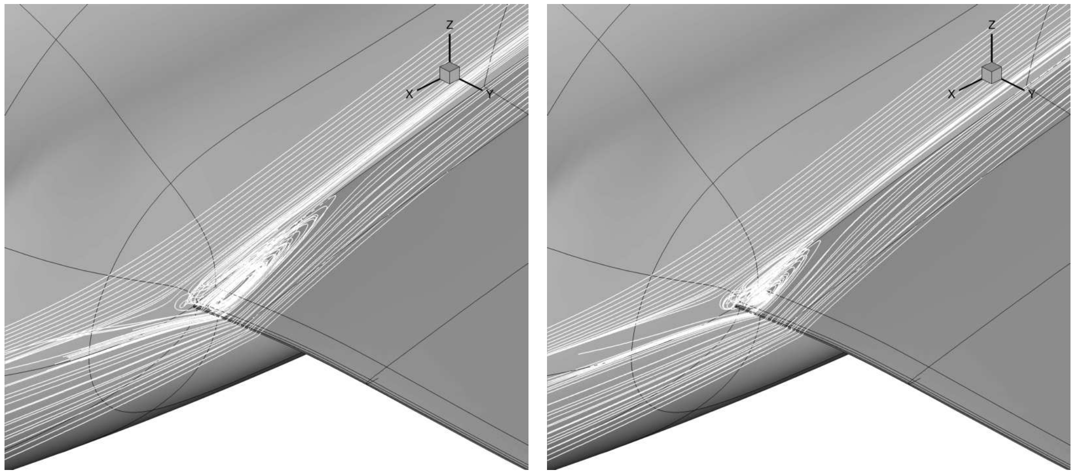

(a) Without QCR (b) With QCR

Figure 10. Recirculation bubble formed at trailing edge wing-body junction (effect of QCR, scalar dissipation, medium grid level)

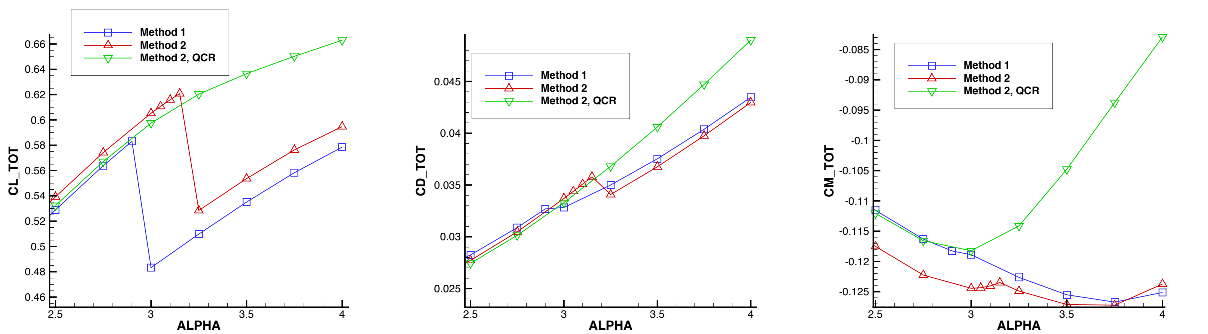

(a) C_L (b) C_D (c) C_M

Figure 11. CRM force coefficients for buffet study

as $\alpha = 3°$, spanning solutions with the Spalart-Allmaras, SST,[31] and Goldberg R_T[32] turbulence models. When present, the size and onset of the large recirculation region appears to be dependent on the algorithm used and the precise details of the implementation. This appears to signify that the large recirculation bubble, while representing a valid numerical steady-state solution, is caused by an interplay of turbulence model implementation and grid parameters in areas of interest, in this case near the trailing edge of the wing-body junction. The large recirculation region was not reported in any experimental results for the NASA CRM geometry, and no corresponding drop in lift was observed in force coefficient measurements. Figure 12 compares the residual convergence of the algorithm with and without QCR, showing the significant impact QCR has on the convergence for this case (shown for $\alpha = 3.75°$).

Examination of the flow features provides an explanation as to the cause of the discontinuity in the force coefficient curves and the increased difficulty of obtaining steady-state flow solutions once the drop in C_L occurs. Figure 13 shows the surface-adjacent streamlines above the wing for solutions at select angles of attack within the test range. Representative solutions are shown for all three data sets, spanning the use of the two different off-wall distance calculation methods, as well as the use of QCR. What becomes immediately obvious is that the discontinuity present in the results that do not use QCR is caused by the sudden appearance of a substantial recirculation region, spanning more than half of the chord on the wing near the fuselage. The presence of the large recirculation bubble alters the flow pattern on the upper surface of the wing, leading to a drop in C_L. When the more accurate off-wall distance calculation method is used, the precise angle at which the large recirculation bubble appears is delayed slightly, but not avoided. The

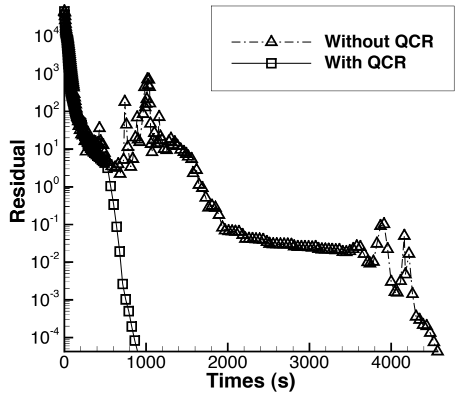

Figure 12. Convergence history for CRM flow at $\alpha = 3.75°$ (medium grid with 704 processors)

use of QCR does not produce the large recirculation region.

At all angles above 3.25°, a substantially different flow pattern is observed, as shown in Figure 13. The use of QCR, which alters the Reynolds stresses in areas of high recirculation, eliminates the large recirculation bubble completely, further pointing to subtle turbulence model interplay with grid features as the source of the erroneous flow feature. This case provides motivation for further research into turbulence modeling, especially with more and more complex cases being tackled by researchers with access to increasing computational resources.

V. Conclusions

The combination of a parallel Newton-Krylov-Schur algorithm with an SBP-SAT discretization was shown to provide accurate and efficient flow solutions to the three-dimensional Reynolds-averaged Navier-Stokes equations with the one-equation Spalart-Allmaras turbulence model. The algorithm was able to obtain accurate solutions for the mandatory cases of DPW5, with a grid-converged value of the drag coefficient in the common grid study within 1 drag count of the median obtained by all workshop participants.

Solutions of two-dimensional turbulence model test cases demonstrate the correspondence of the current algorithm with established flow solvers. The transonic flow computations over the NASA Common Research Model wing-body configuration demonstrate the efficiency and robustness of the algorithm, as well as its ability to accurately capture complex three-dimensional flows.

The importance of an accurate off-wall distance calculation which does not make assumptions about the surface geometry was clearly demonstrated with the grid convergence study. Using a surface-node-based calculation results in drag coefficient values that are more than 10 drag counts higher that the median obtained by all workshop participants. This discrepancy disappears once a more accurate off-wall distance calculation method is used. The buffet study provides motivation for further research into turbulence models and their range of applicability, along with further possible augmentations. The use of quadratic constitutive relations results in more consistent solutions at higher angles of attack, eliminating the large recirculation bubble that leads to a substantial drop in lift once the angle of attack was increased above 3.25° with the baseline Spalart-Allmaras model.

Appendix: Two-Dimensional Cases

The following sections present the results of the optional verification and validation cases from the Turbulence Modeling Resource website. As with the results presented for the mandatory cases of DPW5, the effect of the two off-wall distance calculation methods and the use of QCR is presented. It is important to ascertain that the use of QCR does not lead to inaccurate results for flow conditions that are accurately predicted using the baseline Spalart-Allmaras model.

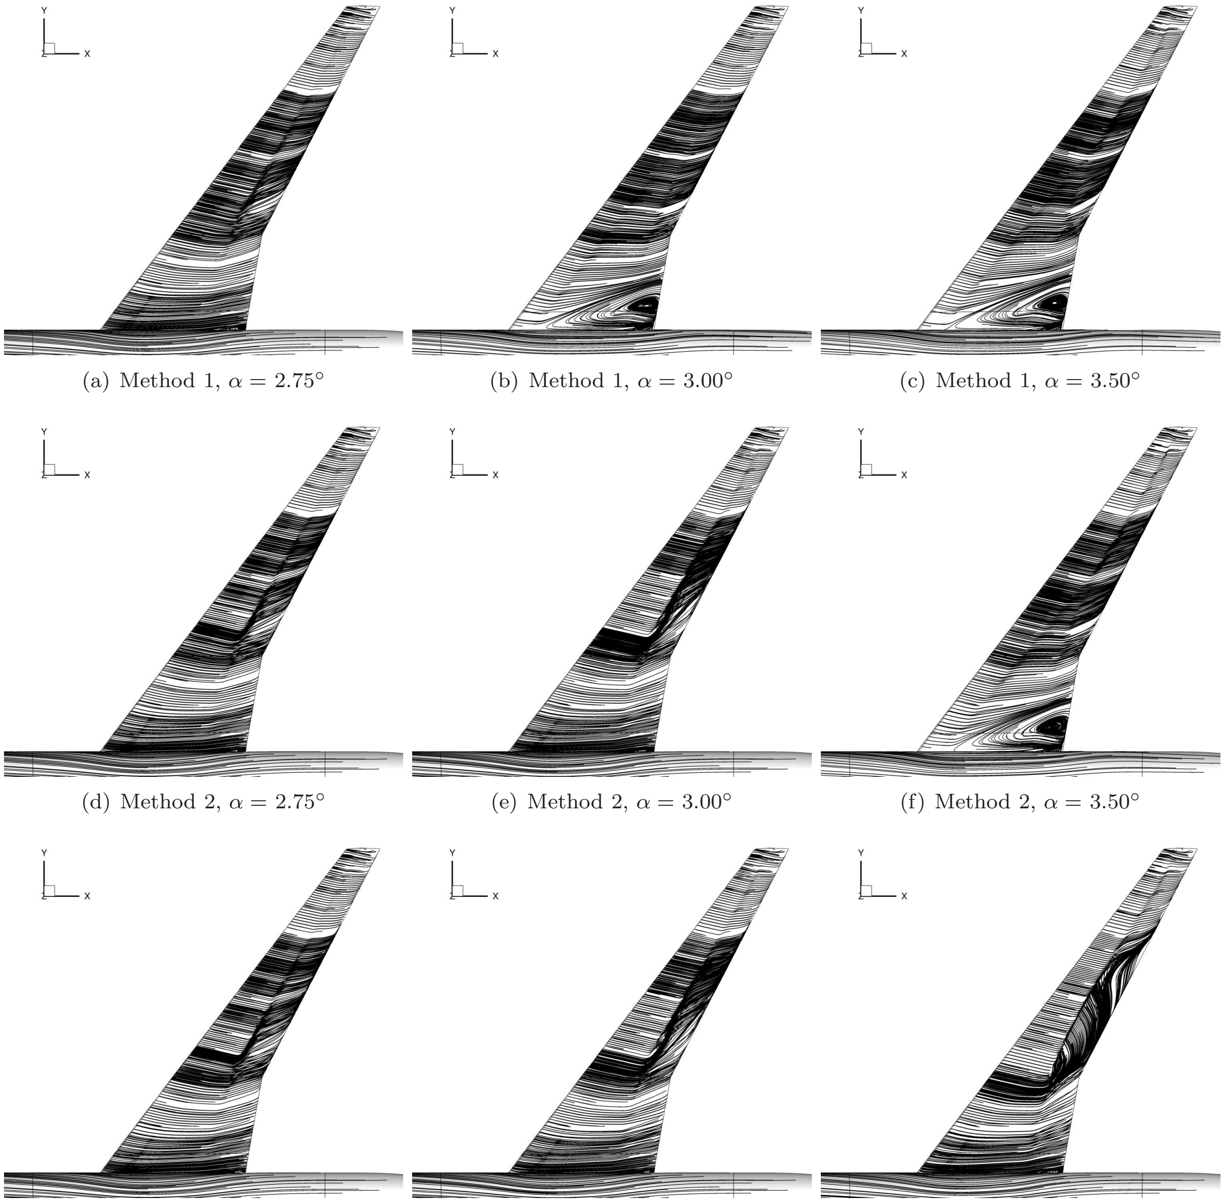

(a) Method 1, $\alpha = 2.75°$ (b) Method 1, $\alpha = 3.00°$ (c) Method 1, $\alpha = 3.50°$

(d) Method 2, $\alpha = 2.75°$ (e) Method 2, $\alpha = 3.00°$ (f) Method 2, $\alpha = 3.50°$

(g) Method 2, QCR, $\alpha = 2.75°$ (h) Method 2, QCR, $\alpha = 3.00°$ (i) Method 2, QCR, $\alpha = 3.50°$

Figure 13. Surface-adjacent streamlines at select angles of buffet study

2D zero-pressure-gradient flat plate

The first verification case considered is the two-dimensional flow over a flat plate. The flow conditions for this case are

$$M = 0.20,\ Re = 5 \times 10^6,\ T_{\text{ref}} = 540°\text{R}.$$

Three grid levels are considered, with node numbers ranging from 137×97 to 545×385. Successively finer off-wall spacing values are used, the finest of which is 5.0×10^{-7} chord units (the flat plat has a length of 2.0 chord units). The coarser meshes were created by successively removing every second node in each coordinate direction from their respective finer counterpart. The grids provide average y^+ values between 0.1 and 0.4, depending on the grid level. The data provided on the TMR website allow for a detailed comparison of the results obtained with the current algorithm with those obtained with CFL3D and FUN3D. Both scalar and matrix dissipation models were tested with this case.

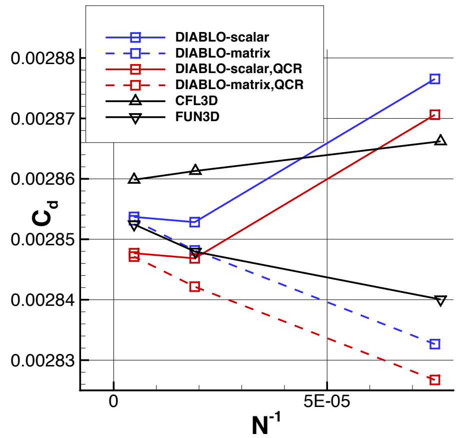

Figure 14. Grid convergence of drag for flow over a flat plate

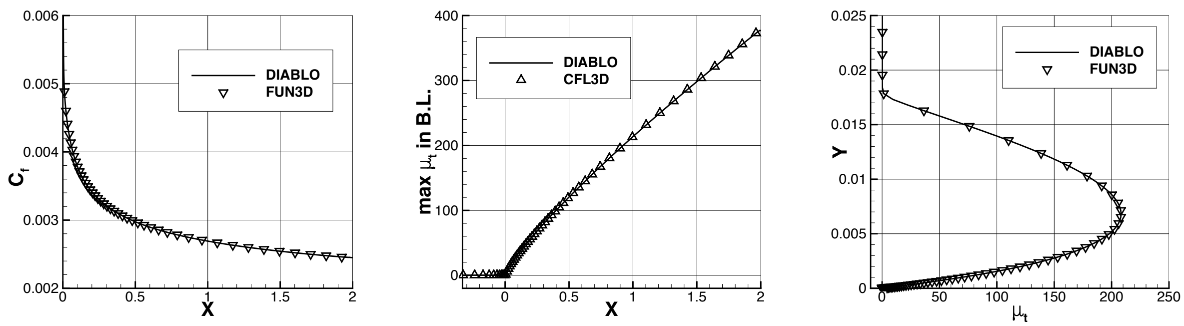

(a) Skin friction (b) Maximum μ_t in boundary layer (c) μ_t at $x = 0.97$

Figure 15. Flow solution comparisons for flow over flat plate (matrix dissipation)

Figure 14 shows the behavior of drag as the grid is refined, where N denotes the number of nodes in the grid. As a result of plotting versus N^{-1}, second-order behaviour should tend toward a straight line. Both dissipation models tend toward the same value of C_d as the grids are refined. However, the scalar dissipation model approaches this value in a non-monotonic manner from above, while the matrix dissipation model closely follows the trend of FUN3D, approaching from below. Both models are tending towards a grid-converged value of drag that lies between the trends of CFL3D and FUN3D. The flat plate geometry allows for the creation of perfectly orthogonal grid lines, resulting in identical solutions for both off-wall distance calculation methods. The use of QCR alters the values obtained with both dissipation models by less than 0.1 drag counts.

Comparisons of the coefficient of skin friction, C_f, maximum turbulent viscosity in the boundary layer, and the turbulent viscosity profile at a vertical section of the finest grid are presented in Figure 15. Since the scalar and matrix dissipation results on this grid level are nearly indistinguishable for all versions of the algorithm, only the matrix dissipation result is shown. Similarly, the results shown are obtained with the baseline Spalart-Allmaras turbulence model, as the QCR results are indistinguishable. Some data points are omitted for clarity. Each of the comparisons shows excellent correspondence between DIABLO and the other algorithms, with nearly identical distributions of all pertinent quantities.

2D bump-in-channel

In order to verify the algorithm in a more complex flow regime where pressure gradients are present, the second case considered was the two-dimensional bump-in-channel flow. Three grid levels are considered, with node numbers ranging from 353×161 to 1409×641, with successively finer off-wall spacing values, the

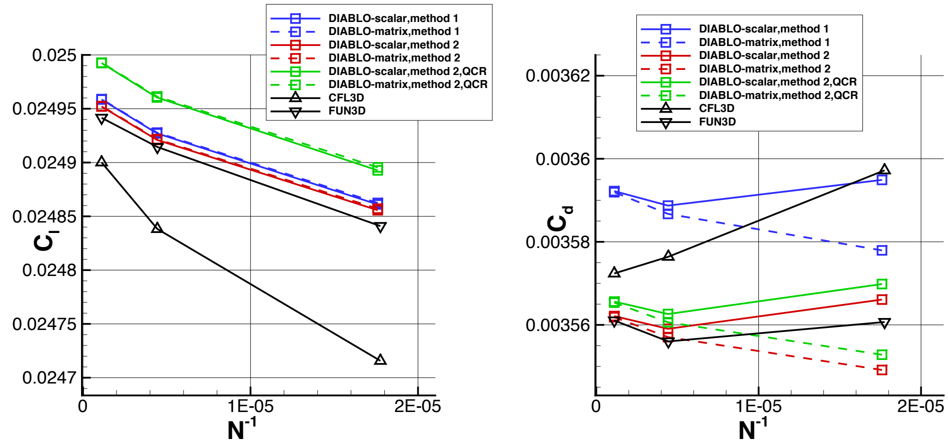

(a) Lift (b) Drag

Figure 16. Force coefficients for bump-in-channel flow

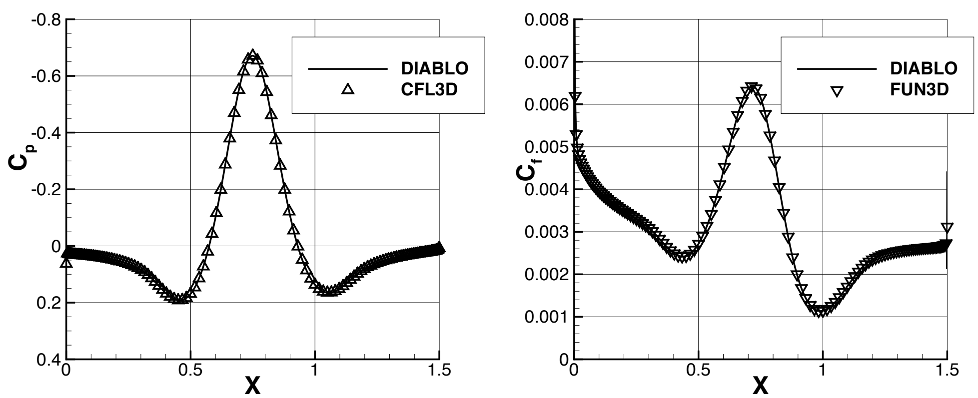

(a) Coefficient of pressure (b) Coefficient of friction

Figure 17. Surface coefficients for bump-in-channel flow (matrix dissipation)

finest of which is 5.0×10^{-7} chord units (the bump has a length of 1.5 chord units). As with the previous case, the coarser grid levels were created by removing every second node in each coordinate direction from the finer grid level. The grids provide average y^+ values between 0.06 and 0.23, depending on the grid level. The flow conditions for this case are

$$M = 0.20,\ Re = 3 \times 10^6,\ T_{\mathrm{ref}} = 540°\mathrm{R}.$$

Figure 16 provides an overview of the grid convergence behaviour of lift and drag for the scalar and matrix dissipation models. For both quantities, the results obtained with DIABLO lie very close to those of the other solvers, with slightly better correspondence to the results of FUN3D. The grid convergence trends of DIABLO are in line with those of CFL3D and FUN3D. Due to the curved nature of the geometry, the values of d produced by the two off-wall distance calculation methods are not identical, resulting in a decrease of roughly 1 drag count in the steady-state coefficient of drag, while lift remains unaffected. The use of QCR has a small impact on the force coefficients.

Additionally, Figures 17 and 18 highlight the excellent correspondence between the current algorithm and the established solvers. This is evident not only for the coefficients of pressure, C_p, and friction along the surface of the bump, but also for the value of μ_t in the boundary layer.

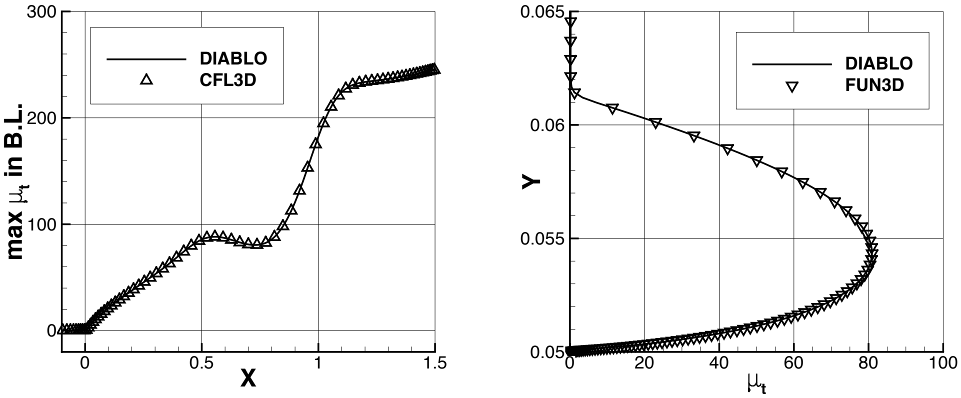

(a) Maximum μ_t in boundary layer

(b) μ_t at $x = 0.75$

Figure 18. μ_t **distribution comparisons for bump-in-channel flow (matrix dissipation)**

2D NACA0012 Airfoil

The final two-dimensional case considered is the flow over the NACA0012 airfoil. This case provides an opportunity not only to compare the current algorithm to CFL3D on a case of practical interest, but also to compare to the experimental data of Gregory and O'Reilly[33] (albeit at a lower Reynolds number of 3×10^6). The flow conditions are

$$M = 0.15,\ Re = 6 \times 10^6,\ T_{\text{ref}}{=}540°\text{R},\ \alpha = 0°,\ 10°,\ \text{and}\ 15°.$$

The grid consists of 1793×513 nodes, with an off-wall spacing of 4×10^{-7} chord units. This grid represents the finest grid level available for this test case on the TMR website, and provides an average y^+ of approximately 0.1. The focus of the results for this study is the distributions of C_p and C_f on the surface of the airfoil. Experimental data are provided for C_p, and the CFL3D C_f data provided is limited to the upper surface of the airfoil.

Figure 19 presents the comparisons for all three angles of attack for the scalar dissipation model, as the two models produced nearly indistinguishable results for this grid. Due to the size of the grid, data points are omitted from the CFL3D data for increased clarity. As can be seen from the figure, the results of DIABLO provide excellent correspondence to those of CFL3D, and line up well with the experimental results. This case provides verification through comparison with CFL3D for a range of flow conditions and a validation of the solver against experimental results, including high angles of attack where boundary-layer separation is present.

Finally, Table 2 presents the coefficients of lift and drag obtained with the three algorithm variants, along with the data available for CFL3D and FUN3D[a]. As expected, QCR has the largest impact at an angle of attack of 15°, where significant flow separation occurs.

These examples demonstrate that the use of QCR has little impact on the prediction of fully attached flows.

Acknowledgments

The authors gratefully acknowledge financial assistance from the Natural Sciences and Engineering Research Council (NSERC), MITACS, Bombardier, the Canada Research Chairs program, and the University of Toronto.

Computations were performed on the GPC supercomputer at the SciNet HPC Consortium and the Guillimin supercomputer of the CLUMEQ consortium, both part of Compute Canada. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

[a]CFL3D and FUN3D data available on 897×257 grid size only

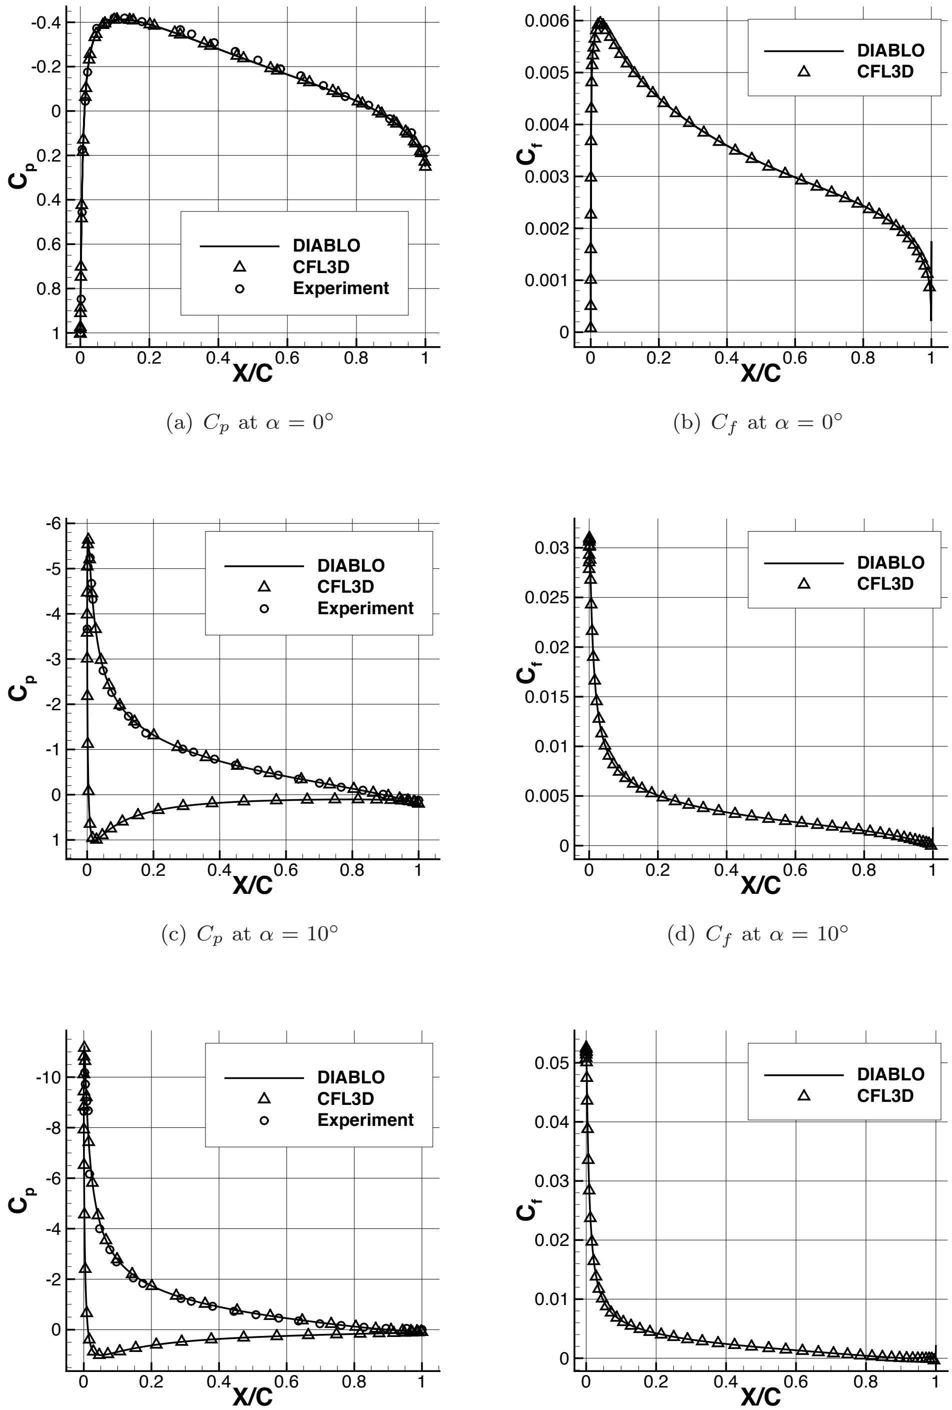

(a) C_p at $\alpha = 0°$
(b) C_f at $\alpha = 0°$
(c) C_p at $\alpha = 10°$
(d) C_f at $\alpha = 10°$
(e) C_p at $\alpha = 15°$
(f) C_f at $\alpha = 15°$

Figure 19. C_p and C_f comparison for NACA0012 flows

References

[1]Mavriplis, D. J., Vassberg, J. C., Tinoco, E. N., Mani, M., Brodersen, O. P., Eisfeld, B., Wahls, R. A., Morrison, J. H., Zickuhr, T., Levy, D., and Murayama, M., "Grid quality and resolution issues from the drag prediction workshop series," *46th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA–2008–0930, Reno, Nevada, Jan. 2008.

Table 2. **Force coefficients for NACA0012 validation case**

		$\alpha = 0°$	$\alpha = 10°$		$\alpha = 15°$	
Algorithm	Diss. model	C_D	C_L	C_D	C_L	C_D
DIABLO - Method 1	scalar	0.816×10^{-2}	1.095	0.124×10^{-1}	1.554	0.213×10^{-1}
	matrix	0.808×10^{-2}	1.095	0.123×10^{-1}	1.554	0.212×10^{-1}
DIABLO - Method 2	scalar	0.816×10^{-2}	1.095	0.124×10^{-1}	1.554	0.213×10^{-1}
	matrix	0.808×10^{-2}	1.095	0.123×10^{-1}	1.554	0.212×10^{-1}
DIABLO - M. 2, QCR	scalar	0.815×10^{-2}	1.091	0.124×10^{-1}	1.536	0.218×10^{-1}
	matrix	0.807×10^{-2}	1.090	0.123×10^{-1}	1.536	0.217×10^{-1}
CFL3D	–	0.819×10^{-2}	1.091	0.123×10^{-1}	1.546	0.212×10^{-1}
FUN3D	–	0.812×10^{-2}	1.098	0.124×10^{-1}	1.555	0.216×10^{-1}

[2]Osusky, M. and Zingg, D. W., "A parallel Newton-Krylov-Schur flow solver for the Navier-Stokes equations discretized using summation-by-parts operators," *AIAA Journal*, accepted May 2013.

[3]"NASA Turbulence Modeling Resource," http://turbmodels.larc.nasa.gov.

[4]Krist, S. L., , Biedron, R. T., and Rumsey, C. L., "CFL3D Users Manual (Version 5.0)," Tech. Rep. NASA TM 1998–208444, National Aeronautics and Space Administration, 1998.

[5]Nielsen, E. J., Walters, R. W., Anderson, W. K., and Keyes, D. E., "Application of Newton-Krylov methodology to a three-dimensional unstructured Euler code," *12th AIAA Computational Fluid Dynamics Conference*, AIAA–95–1733, San Diego, California, United States, June 1995.

[6]Osusky, L. and Zingg, D. W., "A Novel Aerodynamic Shape Optimization Approach for Three-Dimensional Turbulent Flows," *50th AIAA Aerospace Sciences Meeting and Aerospace Exposition*, AIAA–2012–0058, Nashville, Tennessee, United States, Jan. 2012.

[7]Osusky, L. and Zingg, D. W., "Application of an efficient Newton-Krylov aerodynamic shape optimization tool to Reynolds-Averaged Navier-Stokes design problems," *21st AIAA Computational Fluid Dynamics Conference*, San Diego, California, United States, June 2013.

[8]Levy, D. W., Zickuhr, T., Vassberg, J. C., Agarwal, S., Wahls, R. A., Pirzadeh, S., and Hemsch, M. J., "Summary of Data from the First AIAA CFD Drag Prediction Workshop," AIAA–2002–0841, 2002.

[9]Vassberg, J. C., DeHann, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," *26th AIAA Applied Aerodynamics Conference*, AIAA–2008–6919, Honolulu, Hawaii, United States, Aug. 2008.

[10]Vassberg, J. C., Tinoco, E. N., Mani, M., Rider, B., Zickuhr, T., Levy, D. W., Brodersen, O. P., Eisfeld, 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–2010–4547, Chicago, Illinois, United States, June 2010.

[11]Vassberg, J. C., "A Unified Baseline Grid about the Common Research Model Wing-Body for the Fifth AIAA CFD Drag Prediction Workshop," *29th AIAA Applied Aerodynamics Conference*, AIAA–2011–3508, Honolulu, Hawaii, United States, June 2011.

[12]Spalart, P. R., "Strategies for turbulence modeling and simulations," *International Journal of Heat and Fluid Flow*, Vol. 21, 2000, pp. 252–263.

[13]White, F. M., *Viscous Fluid Flow*, McGraw–Hill Book Company, New York, 1974.

[14]Wilcox, D. C., *Turbulence modeling for CFD*, Vol. 2, DCW industries La Canada, 1998.

[15]Spalart, P. R. and Allmaras, S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," *30th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA–92–0439, Reno, Nevada, United States, Jan. 1992.

[16]Yamamoto, K., Tanaka, K., , and Murayama, M., "Comparison study of drag prediction for the 4th CFD Drag Prediction Workshop using structured and unstructured mesh methods," *28th AIAA Applied Aerodynamic Conference*, AIAA–2010–4222, Chicago, Illinois, United States, 2010.

[17]Levy, D. W., Laflin, K. R., Tinoco, E. N., Vassberg, J. C., Mani, M., Rider, B., Rumsey, C. L., Wahls, R. A., Morrison, J. H., Brodersen, O. P., Crippa, S., Mavriplis, D. J., and Murayama, M., "Summary of Data from the Fifth AIAA CFD Drag Prediction Workshop," *51st AIAA Aerospace Sciences Meeting and Aerospace Exposition*, AIAA–2013–0046, Grapevine, Texas, United States, Jan. 2013.

[18]Hicken, J. E. and Zingg, D. W., "A parallel Newton-Krylov solver for the Euler equations discretized using simultaneous approximation terms," *AIAA Journal*, Vol. 46, No. 11, Nov. 2008, pp. 2773–2786.

[19]Osusky, M., Boom, P. D., Del Rey Fernández, D. C., and Zingg, D. W., "An efficient Newton-Krylov-Schur parallel solution algorithm for the steady and unsteady Navier-Stokes equations," *7th International Conference on Computational Fluid Dynamics*, ICCFD7–1801, Big Island, Hawaii, USA, July 2012.

[20]Osusky, M. and Zingg, D. W., "A parallel Newton-Krylov-Schur flow solver for the Reynolds-Averaged Navier-Stokes equations," *50th AIAA Aerospace Sciences Meeting and Aerospace Exposition*, AIAA–2012–0442, Nashville, Tennessee, United States, Jan. 2012.

[21]Jameson, A., Schmidt, W., and Turkel, E., "Numerical solution of the Euler equations by finite volume methods using

Runge-Kutta time-stepping schemes," *14th Fluid and Plasma Dynamics Conference*, AIAA–81–1259, Palo Alto, California, United States, 1981.

[22] Pulliam, T. H., "Efficient solution methods for the Navier-Stokes equations," Tech. rep., Lecture Notes for the von Kármán Inst. for Fluid Dynamics Lecture Series: Numerical Techniques for Viscous Flow Computation in Turbomachinery Bladings, Rhode-Saint-Genèse, Belgium, Jan. 1986.

[23] Swanson, R. C. and Turkel, E., "On central-difference and upwind schemes," *Journal of Computational Physics*, Vol. 101, No. 2, 1992, pp. 292–306.

[24] Lomax, H., Pulliam, T. H., and Zingg, D. W., *Fundamentals of Computational Fluid Dynamics*, Springer–Verlag, Berlin, Germany, 2001.

[25] Saad, Y. and Sosonkina, M., "Distributed Schur complement techniques for general sparse linear systems," *SIAM Journal of Scientific Computing*, Vol. 21, No. 4, 1999, pp. 1337–1357.

[26] Boom, P. D. and Zingg, D. W., "Time-Accurate Flow Simulations Using an Efficient Newton-Krylov-Schur Approach and High-Order Temporal and Spatial Discretization," *51st AIAA Aerospace Sciences Meeting and Aerospace Exposition*, AIAA–2013–0383, Grapevine, Texas, United States, Jan. 2013.

[27] Tucker, P. G., "Assessment of Geometric Multilevel Convergence Robustness and a Wall Distance Method for Flows with Multiple Internal Boundaries," *Applied Mathematical Modeling*, Vol. 22, 1998, pp. 293–311.

[28] Tucker, P. G., Rumsey, C. L., Spalart, P. R., Bartels, R. E., and Biedron, R. T., "Computations of Wall Distances Based on Differential Approaches," *AIAA Journal*, Vol. 43, No. 3, 2005, pp. 539–549.

[29] Apponsah, K. P. and Zingg, D. W., "A Load Balancing Tool for Structured Multi-Block Grid CFD Applications," *20th Annual Conference of the CFD Society of Canada*, Canmore, Alberta, Canada, May 2012.

[30] Baker, T. J., "Mesh generation: Art or science?" *Progress in Aerospace Sciences*, Vol. 41, No. 1, 2005, pp. 29–63.

[31] Menter, F. R., "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598–1605.

[32] Goldberg, U., Peroomian, O., Batten, P., and Chakravarthy, S., "The k-e-Rt Turbulence Closure," *Engineering Applications of Computational Fluid Dynamics*, Vol. 3, No. 2, 2009, pp. 175–183.

[33] Gregory, N. and O'Reilly, C. L., "Low-Speed Aerodynamic Characteristics of NACA 0012 Aerofoil Sections, including the Effects of Upper-Surface Roughness Simulation Hoar Frost," Tech. Rep. NASA R&M 3276, 1970.