Improved Convergence and Robustness of USM3D Solutions on Mixed Element Grids (Invited)

Mohagna J. Pandya[1]
NASA Langley Research Center, Hampton, Virginia 23681, USA

Boris Diskin[2]
National Institute of Aerospace, Hampton, Virginia 23666, USA

and

James L. Thomas[3], Neal T. Frink[4]
NASA Langley Research Center, Hampton, Virginia 23681, USA

Several improvements to the mixed-element USM3D discretization and defect-correction schemes have been made. A new methodology for nonlinear iterations, called the Hierarchical Adaptive Nonlinear Iteration Scheme (HANIS), has been developed and implemented. It provides two additional hierarchies around a simple and approximate preconditioner of USM3D. The hierarchies are a matrix-free linear solver for the exact linearization of Reynolds-averaged Navier Stokes (RANS) equations and a nonlinear control of the solution update. Two variants of the new methodology are assessed on four benchmark cases, namely, a zero-pressure gradient flat plate, a bump-in-channel configuration, the NACA 0012 airfoil, and a NASA Common Research Model configuration. The new methodology provides a convergence acceleration factor of 1.4 to 13 over the baseline solver technology.

I.   Introduction

USM3D flow solver is a component of the NASA Tetrahedral Unstructured Software System (TetrUSS) [1, 2] that was developed during the 1990s to provide a rapid aerodynamic analysis and design capability to applied aerodynamicists. The system is comprised of loosely integrated, user-friendly software that enables the application of advanced Euler and Navier-Stokes tetrahedral finite volume technology to complex aerodynamic problems. The system consists of component software for setting up geometric surface definitions (GridTool) [3], generating tetrahedral grids (VGRID) [4-6], computing Euler and Navier-Stokes flow solutions (USM3D) [7-12], and extracting meaningful information from analysis of results (SimpleView). The system is also coupled with the design tool CDISC [13]. Significant extensions to VGRID and USM3D capabilities are described in Ref. [2].

The USM3D tetrahedral grid flow solver has been widely used over the last two decades within NASA [14, 15], other U.S. government agencies [16] and industry [17] as a workhorse for aerodynamic analysis of complex configurations. A recent example is its role as a lead CFD flow solver within the Ares project of the NASA Constellation Program, furnishing over 7,000 solutions for the vehicles up through supersonic speeds [14,15]. USM3D was used in the aerodynamic development of the Crew Exploration Vehicle and its Launch Abort System configuration, providing over 1,000 solutions [2]. USM3D is also being used extensively in the NASA Aviation Safety Program to generate dynamic stability and control databases for civil transports for stall/post-stall flight regimes [18].

Some of the speed of USM3D computations is attributed to an analytical formulation of the spatial differencing stencil that exploits invariant features of tetrahedra [8,19], which renders it unnecessary to explicitly compute and

---

[1]Research Aerospace Engineer, Configuration Aerodynamics Branch, Mail Stop 499, Senior Member AIAA.
[2]NIA Research Fellow, 100 Exploration Way, Associate Fellow AIAA.
[3]Distinguished Research Associate, Computational Aerosciences Branch, Fellow AIAA.
[4]Senior Aerospace Engineer, Configuration Aerodynamics Branch, Mail Stop 499, Associate Fellow AIAA.

store the cell gradients. This advantage is complemented by the VGRID tetrahedral grid generator, which generates organized thin layers of tetrahedral cells in the near-wall boundary layer region using Advancing Layer Method (ALM) [5]. Unfortunately, the diagonal faces and skewed topology of the highly stretched tetrahedra result in a severely distorted spatial differencing stencil for computing fluxes, which may increase the numerical dissipation within the near-wall flow solution. Other cell topologies, such as hexahedra and prisms can be exploited to avoid this problem.

Since many recent USM3D applications have ventured into predominately separated flow regimes, attention is warranted on the sensitivity of smooth surface flow separation, such as a wing leading edge near stall, to near-wall cell topology. A similar case can be made for computing heat transfer on high-speed bodies. The USM3D simulations for such flows can be improved by using more flow-aligned anisotropic hexahedral or prismatic cells in the boundary layer and isotropic tetrahedral cells away from the surfaces.

Work to expand USM3D support for additional cell topologies, including hexahedra, prisms, and transitional pyramids, was presented in Ref. [20]. Extensions to the key elements of solution methodology, such as geometric quantities, cell- and face-centered gradients were described and a code verification study was presented. The study used three two-dimensional (2D) test cases available on the Turbulence Modeling Resource (TMR) website [21]: a zero-pressure-gradient flat plate, a planar shear case, and a bump-in-channel configuration.

The 2D bump-in-channel case in the verification study of the baseline mixed-element USM3D [20] showed that the solution convergence becomes sluggish, as the grid is refined. The convergence of forces and moment required as many as 500,000 nonlinear iterations on the finest hexahedral grid for this case. The principle reason for the slow convergence is that the baseline code advances the nonlinear solution in pseudo time using only a simple preconditioner. The preconditioner uses a low-fidelity defect-correction scheme and point-implicit Gauss-Seidel (G-S) iterations. A limiting aspect of this approach is that typically, a low Courant-Friedrichs-Lewy (CFL) number is needed to preserve solution robustness. A low CFL number adversely impacts iterative convergence and increases the time needed for converging a solution.

The objective of this study is to eliminate these bottlenecks. Several improvements to the mixed-element USM3D discretization and defect-correction schemes are made. The improvements include provisions to accommodate negative values for the turbulence variable in the Spalart-Allmaras model [22], a robust treatment of possible violations of realizability constraints, second-order approximations to the convective term of the SA turbulence model equation [23], and more accurate Jacobian approximations. A new methodology for nonlinear iterations, called the Hierarchical Adaptive Nonlinear Iteration Scheme (HANIS), has been implemented. It provides two additional hierarchies over the baseline USM3D preconditioner-alone (PA) solver. The hierarchies are an enhanced linear solver for the exact linearization of Reynolds-Averaged Navier Stokes (RANS) equations and a nonlinear control of the solution update. The linear solver uses matrix-free methods [24-26]. The nonlinear solution update strategy automatically adapts the under-relaxation parameter and pseudo-time step similar to recently reported approaches [27-29]. CFL adaptation is used as a comprehensive tool to address robustness, efficiency, and automation considerations. In this paper, key components of HANIS are described. Two different variants of HANIS methodology are assessed on four benchmark turbulent flow cases, namely, a 2D zero-pressure-gradient flat plate, a 2D bump-in-channel configuration, the 2D NACA 0012 airfoil, and a NASA Common Research Model (CRM) configuration used for the 4[th] AIAA Drag Prediction Workshop.

II. Baseline Mixed-Element USM3D

USM3D solves a system of nonlinear flow equations that can be formally represented as
$$\boldsymbol{R}(\boldsymbol{Q}) = 0. \tag{1}$$
The discrete nonlinear operator, $\boldsymbol{R}(\boldsymbol{Q})$, may, for example, represent a discretization of steady-state Reynolds-Averaged Navier Stokes (RANS) equations. The USM3D code extended for this study is a mixed-element cell-centered, finite volume CFD solver [20]. The term "cell-centered" means that the flow variables are solved at the centroid of each cell. The solution is advanced in pseudo time to a steady-state condition by an implicit backward-Euler scheme. Currently, three turbulence models are available: the Spalart-Allmaras (SA) one-equation model, the Menter Shear Stress Transport (SST) two-equation model, and the Langtry-Menter transition prediction model.

Inviscid fluxes are computed at each cell face using various upwind schemes, such as van Leer's Flux Vector Splitting (FVS), Roe's Flux Difference Splitting (FDS), Harten, Lax, van Leer, Einfeldt (HLLE), and Harten, Lax, van Leer – Contact (HLLC) schemes, among others. Spatial discretization is accomplished by a reconstruction process based on solution gradients computed within cells. The code has several options for the gradient computation within a cell, such as the Green-Gauss integration, unweighted or weighted least squares procedures, or the Mitchell stencil. Face gradients are needed to compute diffusive fluxes for the mean flow and turbulence model

equations. A face gradient can be evaluated from either Mitchell's stencil or an average of the gradients computed within two cells sharing the face. For the latter case, the averaged gradient is suitably augmented using a face-normal method, an edge-normal method, or a face-tangent method. Discrete residuals at any given iteration are computed using solution variables at the cells, nodes, boundary faces and ghost cells as well as cell and face gradients. Solution quantities at these locations are not synchronized. For example, solution variables at nodes are updated using the ghost cell values evaluated at the previous iteration.

The USM3D nonlinear iterations are based on a defect correction scheme

$$\frac{V}{\Delta \tau} \Delta \boldsymbol{Q} + \frac{\widehat{\partial \boldsymbol{R}}}{\partial \boldsymbol{Q}} \Delta \boldsymbol{Q} = -\boldsymbol{R}(\boldsymbol{Q}^n), \tag{2}$$

$$\boldsymbol{Q}^{n+1} = \boldsymbol{Q}^n + \Delta \boldsymbol{Q}. \tag{3}$$

Here, $\boldsymbol{Q}^n$ and $\boldsymbol{Q}^{n+1}$ are the solutions at iterations n and $n+1$, respectively; $\frac{\widehat{\partial \boldsymbol{R}}}{\partial \boldsymbol{Q}}$ is an approximation to Jacobian $\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}}$; V is a control volume; and $\Delta \tau$ is a pseudo-time step, which is set through a CFL specification. The mean flow and turbulence model equations are loosely coupled. The approximate Jacobian for the mean flow equations is formed using the linearization of the first-order FVS inviscid fluxes and a thin-layer approximation for the viscous fluxes. The approximate Jacobian for a turbulence-model equation includes the contributions from the advection, diffusion, and source terms. The advection term is linearized with a first-order approximation. A thin-layer approximation is used for the diffusion term. Only positive contributions from the source term linearization are added to the diagonal of the approximate Jacobian. The approximate Jacobians of the mean flow and turbulence model equations are updated at every nonlinear iteration. An option to use single precision for the approximate Jacobian off-diagonal terms and for the solution updates is available to reduce the memory footprint.

A few Gauss-Seidel (G-S) iterations are used to approximately solve Eq. (2). The number of G-S iterations is a user-defined input parameter; typically, 10 to 20 iterations are used. The resulting update, $\Delta \boldsymbol{Q}$, is applied to advance the nonlinear solution (Eq. (3)). The updates are fully applied for the mean flow equations, and under-relaxed for the turbulence model equations. The pseudo-time step is updated according to a pre-scheduled ramping of the CFL number starting from a very low initial value. The final (highest) CFL number is typically modest (50-150) for the sake of solution robustness. Robustness is improved further by initially using a first-order approximation to the nonlinear residuals for a certain number of nonlinear iterations. The target second-order approximation of the nonlinear residuals is used for the subsequent iterations. This start-up strategy may mitigate solution difficulties during initial transients.

III. New Developments in USM3D

A. Discretization Features

A fully-implicit formulation is implemented for the present study implying that the solution variables at the grid nodes and boundary faces as well as the cell gradients are computed solely from the current solution variables defined at the cell centers. Except for the symmetry boundary, ghost cells are eliminated for all boundaries. Present implementation precludes intersecting symmetry boundaries. Solution at a boundary node that is not on the symmetry plane is computed using extrapolation from the interior of the computational domain.

The cell-centered solution values are required to satisfy the realizability constraints, such as positive values of density and pressure. For first-order spatial accuracy, the solution reconstructed at a cell face is the same as the cell-centered solution. For second-order spatial accuracy, a realizability check is newly implemented for the reconstructed solution to avoid a catastrophic failure (an underflow condition associated with a square root of a negative number) in computing mean flow inviscid fluxes. If a realizability violation is detected at a face during any nonlinear iteration, then the face is temporarily designated as "first-order". The face realizability check is performed for the second-order reconstruction values for all faces including the "first-order" faces, however, the "first-order" faces employ the first-order reconstruction for flux computations. The "first-order" designation for a face is removed after the second-order reconstruction values have been realizable for a certain number (currently 20) of consecutive nonlinear iterations.

A negative formulation of SA model [22] has been implemented to improve numerical convergence on under-resolved grids. A face realizability check for the negative SA model is not needed.

B. Defect-Correction Solver Features

Linearization of the FDS scheme for the first-order mean flow inviscid fluxes has been implemented for the approximate Jacobian in Eq. (2). The FDS scheme is less dissipative than the FVS scheme and provides a better

approximation to the target second order accurate nonlinear operator that also uses the FDS scheme. Historically, the FVS scheme has been used for the approximate Jacobian in the USM3D solutions. The SA model linearization is somewhat modified. Similar to the baseline solver, the diagonal of the approximate Jacobian collects contributions from the advection, diffusion, and source terms. The new aspect is that the entire contribution from the linearized source term is added to the diagonal. Positivity check for the diagonal values is conducted before adding the pseudo-time term. Negative diagonal values are substituted by their absolute values.

Although not used in the present study, a provision for an intermittent Jacobian update is implemented for runtime efficiency. The Jacobian matrix can be frozen for several nonlinear iterations, subject to certain solution conditions and/or a user-defined schedule.

Residual reduction targets are used for an earlier termination of G-S iterations to improve the runtime efficiency. The G-S iterations are terminated if the rms norm of the residuals (Eq. (2)) has been reduced by one order of magnitude, any time after a minimum of 10 G-S iterations have been performed. The G-S iterations can terminate even before completing 10 G-S iterations if the rms norm of the residuals has been reduced by two orders of magnitude. The earlier termination option significantly improves the runtime efficiency when a user specifies a large maximum number of G-S iterations.

C. Preconditioner-Alone (PA) Nonlinear Iteration Scheme

PA is a nonlinear scheme similar to the scheme used in the baseline USM3D discussed in the Section II. It benefits from the algorithmic modifications described in sub-sections III A and III B. Figure 1a shows a flowchart for the PA scheme. The preconditoner and nonlinear update modules in Fig. 1a represent a linear solver associated with Eq. (2) and the solution update (Eq. (3)), respectively.

D. Hierarchical Adaptive Nonlinear Iteration Scheme (HANIS)

1. HANIS Overview

HANIS is an enhanced solver for Eq. (1) that seeks to improve (i) robustness of iterations by providing a mechanism to treat difficult (e.g., transient) solution states, (ii) efficiency by iterating at higher CFL numbers, and (iii) automation by reducing the number of ad hoc parameters. HANIS extends the PA scheme described in Section III C by providing two additional hierarchies around the preconditioner. The hierarchies are a matrix-free linear solver for the exact linearization of nonlinear equations and a nonlinear control of the solution updates. An economical version of the Generalized Conjugate Residual (GCR) method [25] is used as the matrix-free linear solver. Note that HANIS mainly uses GCR to control the stability of linear iterations, and therefore only a few search directions are used in GCR. The nonlinear control is used as a mechanism to optimize the reduction of nonlinear residuals.

The three HANIS hierarchies are interrelated. The preconditioner generates a new search direction for the GCR matrix-free linear solver. GCR suggests updates for the current nonlinear solution. The updates are checked for the solution realizability and passed on to the nonlinear control where they may be scaled.

HANIS specifies the maximum number of G-S iterations allowed in preconditioner and the maximum number of search directions allowed in GCR. Furthermore, HANIS also prescribes the residual reduction targets for the preconditioner, GCR, and nonlinear solution updates. The targets for GCR and nonlinear updates are quite soft.

HANIS uses CFL adaptation as a comprehensive tool to address robustness, efficiency, and automation considerations. HANIS tries to run nonlinear iterations at the highest possible CFL number; it increases the CFL number (currently by a factor 2) if all HANIS hierarchies have reported complete success. On the other hand, if any convergence difficulty has been encountered in the preconditioner, GCR, or during a nonlinear solution update, HANIS discards the suggested correction and aggressively reduces the CFL number (currently by a factor 10). The approach is based on the assumptions stated below.

Robustness Assumption. For any physically realizable initial approximation, there is a small enough CFL, with which the defect-correction iterations (Eqs. (2)-(3)) converge to the solution of Eq. (1).

Efficiency Assumption. In a vicinity of the converged solution, convergence rates of nonlinear iterations would benefit from using (infinitely) large CFL numbers.

The HANIS CFL adaptation methodology often results in large CFL oscillations. The ratio between CFL reduction and increase factors defines the frequency of failure. The current ratio implies that in normal operational conditions every third or fourth iteration is expected to fail and reduce the CFL number. An oscillatory CFL number is often interpreted as a sign of convergence difficulties. Some CFD codes that also use a CFL adaptation strategy put a limit on the maximum CFL number allowed. The CFL number is expected to settle at the specified maximum level. Generally, it is impossible to identify the right maximum level in advance. Any ad hoc cap is likely to be either useless, if the CFL operational range for a solution is well below the cap, or unduly restrictive, if the CFL

operational range is well above the cap. The HANIS methodology is different. CFL number oscillating at a higher level is preferable to a constant low-level CFL number.

The hierarchical structure of HANIS allows for a combination of tightly- and loosely-coupled iteration strategies. In USM3D, preconditioner iterations are always loosely coupled. This approach reduces memory footprint of approximate Jacobian, improves cache characteristics, and lowers the operation count of G-S iterations. The matrix-free GCR and nonlinear control methods can either loosely or tightly couple the mean flow and turbulence model equations, notwithstanding the formulation adopted for preconditioner.

Two HANIS versions, namely, HANIS-1 and HANIS-4, are considered in the present study. HANIS-1 uses just one GCR search direction, whereas, HANIS-4 can use a maximum of four search directions. A detailed parametric study to identify an optimal combination of HANIS parameters for minimizing solution time is needed in future.

The HANIS implementation in USM3D is such that any nonlinear iteration is always successful. This is guaranteed by using an infinite loop around the newly added HANIS hierarchies and updating the CFL number. The HANIS flowchart is shown in Fig. 1b. HANIS is implemented using six different modules, identified as, *Preconditioner*, *GCR Solver*, *Realizability Check*, *Nonlinear Control*, *Failure Handler*, and *Nonlinear Update*. The Preconditioner module provides a new search direction for the GCR Solver module. The GCR Solver module solves linearized equations and suggests an update, $\Delta \boldsymbol{Q}$, for the current nonlinear solution, $\boldsymbol{Q}^n$. The Realizability Check module assesses whether the updated solution violates any physical constraints. The Nonlinear Control module optimizes the suggested update. The Failure Handler module treats possible exceptions from the preceding modules. The Nonlinear Update module computes the updated solution. The infinite loop is terminated when the GCR Solver, Realizability Check, and Nonlinear Control modules report success. All HANIS modules are described below.

2. Preconditioner

Preconditioner approximately solves the system of linear equations for $\Delta \boldsymbol{Q}$:

$$\frac{V}{\Delta \tau} \Delta \boldsymbol{Q} + \frac{\widehat{\partial \boldsymbol{R}}}{\partial \boldsymbol{Q}} \Delta \boldsymbol{Q} = \boldsymbol{g}. \tag{4}$$

The mean flow and turbulence model equations are solved separately using point-implicit G-S iterations. The initial solution approximation is the identical zero. For the PA and HANIS-1 iterations, the right hand side of Eq. (4) is

$$\boldsymbol{g} = -\boldsymbol{R}(\boldsymbol{Q}^n). \tag{5}$$

For HANIS-4 iterations, the right hand side may also represent a current linear residual in GCR. The HANIS Preconditioner module introduces an additional condition for an earlier termination of G-S iterations. The preconditioner residuals are allowed to increase by a maximum of two orders. Termination of Preconditioner because of reaching the residual reduction target or performing maximum iterations is considered success. Termination because of exceeding the maximum residual increase is considered failure. In HANIS iterations, the update, $\Delta \boldsymbol{Q}$, and an indication of success or failure of the Preconditioner module are sent to the GCR Solver module to form a new search direction.

3. GCR Solver

GCR iterations approximately solve the linear equations

$$\frac{V}{\Delta \tau} \Delta \boldsymbol{Q} + \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}} \Delta \boldsymbol{Q} = -\boldsymbol{R}(\boldsymbol{Q}^n). \tag{6}$$

Here, $\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}}$ is the exact linearization of the nonlinear residual operator around the current solution, $\boldsymbol{Q}^n$. The presently implemented HANIS uses a tightly-coupled formulation in the GCR Solver module. This approach provides a better approximation to nonlinear residuals in the limit of the infinitely high CFL number. In the tightly-coupled GCR formulation, $\boldsymbol{Q}^n$, is a vector that combines the mean flow and the turbulence model variables. The nonlinear residual $\boldsymbol{R}(\boldsymbol{Q}^n)$ is a vector that combines the mean flow and turbulence model residuals. The exact linearization matrix is big and requires significant memory. Instead a matrix-free approach is followed. The matrix-vector product in Eq. (6) is approximated by a Frechet derivative

$$\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}} \Delta \boldsymbol{Q} \approx \frac{\boldsymbol{R}\left(\boldsymbol{Q}^n + \epsilon \frac{\Delta \boldsymbol{Q}}{|\Delta \boldsymbol{Q}|}\right) - \boldsymbol{R}(\boldsymbol{Q}^n)}{\epsilon} |\Delta \boldsymbol{Q}|. \tag{7}$$

Here ϵ is a small factor; the current implementation uses $\epsilon = \max(1, \|\boldsymbol{Q}^n\|) \, 10^{-7}$. Note that $|\cdot|$ and $\|\cdot\|$ denote the l_2 and rms norms of a vector, respectively. This approximation is accurate and insensitive to the size of $\Delta \boldsymbol{Q}$.

The HANIS-1 algorithm always uses one search direction; the HANIS-N algorithm allows for maximum of N search directions. The initial solution is set as $\Delta \boldsymbol{Q} = 0$, thus the initial GCR residual is

$$\boldsymbol{r}_0 = -\boldsymbol{R}(\boldsymbol{Q}^n). \tag{8}$$

For a k-th search direction ($k = 1, \dots, N$), an improved approximation, $\Delta \boldsymbol{Q}_k$, is computed by preconditioner,

$$\frac{V}{\Delta\tau}\Delta\boldsymbol{Q}_k + \frac{\widehat{\partial\boldsymbol{R}}}{\partial\boldsymbol{Q}}\,\Delta\boldsymbol{Q}_k = \boldsymbol{r}_{k-1}. \tag{9}$$

If the Preconditioner module fails, then the GCR Solver module immediately terminates with failure. If the Preconditioner module reports success, the new search direction, $\boldsymbol{b}_k$, is computed and normalized as

$$\boldsymbol{b}_k = \frac{V}{\Delta\tau}\Delta\boldsymbol{Q}_k + \frac{\boldsymbol{R}\left(\boldsymbol{Q}^n + \epsilon\frac{\Delta\boldsymbol{Q}_k}{|\Delta\boldsymbol{Q}_k|}\right) - \boldsymbol{R}(\boldsymbol{Q}^n)}{\epsilon}|\Delta\boldsymbol{Q}_k|; \quad \Delta\boldsymbol{Q}_k = \frac{\Delta\boldsymbol{Q}_k}{|\boldsymbol{b}_k|}, \quad \boldsymbol{b}_k = \frac{\boldsymbol{b}_k}{|\boldsymbol{b}_k|}. \tag{10}$$

Then the $\boldsymbol{b}_k$ vector is orthonormalized against previously stored search directions, $\boldsymbol{b}_j, 1 \le j < k$, simultaneously updating $\Delta\boldsymbol{Q}_k$. For each $\boldsymbol{b}_j$ ($j < k$),

$$\mu = \boldsymbol{b}_k^T\boldsymbol{b}_j, \qquad \boldsymbol{b}_k = \boldsymbol{b}_k - \mu\boldsymbol{b}_j, \quad \Delta\boldsymbol{Q}_k = \Delta\boldsymbol{Q}_k - \mu\Delta\boldsymbol{Q}_j, \qquad \Delta\boldsymbol{Q}_k = \frac{\Delta\boldsymbol{Q}_k}{|\boldsymbol{b}_k|}, \qquad \boldsymbol{b}_k = \frac{\boldsymbol{b}_k}{|\boldsymbol{b}_k|}. \tag{11}$$

After completing the orthonormalization procedure, the projection, γ_k, of the current residual on the k-th search direction is computed, and the correction and the linear residual are updated as

$$\gamma_k = \boldsymbol{b}_k^T\boldsymbol{r}_{k-1}, \quad \Delta\boldsymbol{Q} = \Delta\boldsymbol{Q} + \gamma_k\Delta\boldsymbol{Q}_k, \quad \boldsymbol{r}_k = \boldsymbol{r}_{k-1} - \gamma_k\boldsymbol{b}_k. \tag{12}$$

In rare situations, vectors $\boldsymbol{b}_k$ and $\boldsymbol{r}_{k-1}$ can be (nearly) orthogonal implying $\gamma_k \ll |\boldsymbol{r}_{k-1}|$, then (almost) no update is applied. These situations are called *GCR stall* and are treated as a GCR Solver module failure.

If the rms norm of the updated linear residual has been reduced more than the GCR residual reduction target, μ_{gcr}, (i.e., $\|\boldsymbol{r}_k\|/\|\boldsymbol{r}_0\| < \mu_{gcr}$), then the GCR Solver module terminates with success. Otherwise another search direction is generated. If the maximum number of search directions has been used, and the residual reduction target has not been met, the GCR Solver module terminates with failure. Note that by construction, the rms norm of the GCR residual cannot increase. To reduce the memory footprint, GCR search directions, $\boldsymbol{b}_k$, and updates, $\Delta\boldsymbol{Q}_k$, can be stored with single precision. In case of failure, the GCR Solver module transfers control to the Failure Handler module. In case of success, the suggested nonlinear solution update, $\Delta\boldsymbol{Q}$, is transferred to the Realizability Check module.

4. Realizability Check

The role of the Realizability Check module is to prevent solution updates that could lead to a non-physical solution. To check the realizability constraints, the update, $\Delta\boldsymbol{Q}$, is first added to the solution, $\boldsymbol{Q}^n$. If the updated solution at any cell center violates physical realizability conditions (positive pressure, density, etc.), the Realizability Check module restores solution $\boldsymbol{Q}^n$, declares failure, and transfers the control to the Failure Handler module. If the suggested update successfully passed the check, then $\boldsymbol{Q}^n$ is restored and the update, $\Delta\boldsymbol{Q}$, is sent to the Nonlinear Control module.

5. Nonlinear Control

The role of the Nonlinear Control module is to find an optimal under-relaxation parameter, w, for the suggested update, $\Delta\boldsymbol{Q}$. The Nonlinear Control module uses the tightly-coupled formulation. The initial value for the under-relaxation parameter is $w = 1$. The cumulative under-relaxation parameter, w_{cum}, is introduced for a subsequent use and is initialized as $w_{cum} = 1$. The updated solution,

$$\boldsymbol{Q} = \boldsymbol{Q}^n + w\Delta\boldsymbol{Q}, \tag{13}$$

is computed together with the pseudo-time nonlinear residual, $\boldsymbol{R}_\tau(\boldsymbol{Q})$,

$$\boldsymbol{R}_\tau(\boldsymbol{Q}) \equiv \frac{V}{\Delta\tau}(\boldsymbol{Q} - \boldsymbol{Q}^n) + \boldsymbol{R}(\boldsymbol{Q}). \tag{14}$$

The goal of the Nonlinear Control module is to reduce the rms norm of the residual (Eq. (14)) below the residual reduction target, μ_{nl}, to satisfy

$$\|\boldsymbol{R}_\tau(\boldsymbol{Q})\|/\|\boldsymbol{R}(\boldsymbol{Q}^n)\| \le \mu_{nl}. \tag{15}$$

In the current implementation,

$$\mu_{nl} = 0.5(1 + \mu_{gcr}). \tag{16}$$

If the nonlinear residual reduction target (Eq. (15)) is met with $w = 1$ and the update, $\Delta\boldsymbol{Q}$, is large enough to dominate the round-off error, then the Nonlinear Control module declares success and schedules the CFL increase for the next nonlinear iteration. Small updates are treated differently. If a CFL number adaptation strategy does not take into account the size of suggested corrections, on final stages of convergence, when round-off error plays a bigger role, the CFL number can react on this round-off error and frequently change for no reason. In the current implementation, the update with the rms norm less than a small fraction of the current solution norm is considered as dominated by the round-off error. Such an update is accepted as success, if it leads to any reduction of the nonlinear residual, and no change of CFL number is allowed.

If $\|\boldsymbol{R}_\tau(\boldsymbol{Q})\|/\|\boldsymbol{R}(\boldsymbol{Q}^n)\| > \mu_{nl}$, then an under-relaxation parameter is sought through a quadratic search. Two nonlinear scalar functions of w are defined,

$$f(w) \equiv \left\| \boldsymbol{R}(\boldsymbol{Q}^n + w\Delta\boldsymbol{Q}) + \frac{V}{\Delta\tau} w\Delta\boldsymbol{Q} \right\|, \tag{17}$$

$$h(w) \equiv \left\| \boldsymbol{R}(\boldsymbol{Q}^n) + \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}} w\Delta\boldsymbol{Q} + \frac{V}{\Delta\tau} w\Delta\boldsymbol{Q} \right\|, \tag{18}$$

describing rms norm variations with respect to w of the nonlinear (Eq. (14)) and linear (Eq. (6)) residuals, respectively.

The following properties are satisfied:

$$\|\boldsymbol{R}(\boldsymbol{Q}^n)\| = f(0) = h(0); \quad \|\boldsymbol{R}_\tau(\boldsymbol{Q}^n + \Delta\boldsymbol{Q})\| = f(1); \quad \left\| \boldsymbol{R}(\boldsymbol{Q}^n) + \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{Q}}\Delta\boldsymbol{Q} + \frac{V}{\Delta\tau}\Delta\boldsymbol{Q} \right\| = h(1). \tag{19}$$

Here, $f(0)$ and $h(0)$ are equal to the rms norm of the initial nonlinear residual, $f(1)$ is equal to the rms norm of the current nonlinear residual (Eq. (14)), and $h(1)$ is equal to the rms norm of the linear residual $\boldsymbol{r}_k$ (Eq. (12)) computed by the GCR Solver module. Assuming a quadratic variation for $f(w) \approx a\,w^2 + b\,w + c$ and a linear variation for $h(w) \approx b\,w + c$, the following relations can be derived $c \approx f(0), b \approx h(1) - c, a \approx f(1) - b - c$. The minimum of a quadratic function $a\,w^2 + b\,w + c$ is achieved at $w = -\frac{b}{2a}$. Thus, the under-relaxation parameter is chosen as

$$w = -\frac{h(1) - f(0)}{2\big(f(1) - h(1)\big)}. \tag{20}$$

The Nonlinear Control module restores the initial solution, $\boldsymbol{Q}^n$, updates $\Delta\boldsymbol{Q} = w\Delta\boldsymbol{Q}$, stores the cumulative under-relaxation parameter, $w_{cum} = w_{cum}w$, and repeats computations starting from Eq. (13). To avoid an infinite loop, the cumulative under-relaxation parameter is required to be greater than a specified minimum, $w_{cum} > w_{min}$. The choice of the nonlinear residual reduction target (Eq. (16)) ensures that if the linear solver meets the required linear residual reduction, μ_{gcr}, but the nonlinear residual does not achieve the target reduction, μ_{nl}, then the under-relaxation parameter (Eq. (20)) is bounded as $0 \le w \le 1$. The value of the minimum under-relaxation parameter is set as

$$w_{min} = \frac{1 - \mu_{nl}}{1 - \frac{\|\boldsymbol{r}_k\|}{\|\boldsymbol{r}_0\|}}, \tag{21}$$

where $\|\boldsymbol{r}_k\|/\|\boldsymbol{r}_0\|$ is the actual reduction of the rms norm of the linear residual achieved by the GCR method. This value of w_{min} is chosen from the consideration that with a cumulative under-relaxation parameter smaller than w_{min} ($w_{cum} < w_{min}$), the rms norm of the linear residual is reduced less than μ_{nl}, and thus, there is no reason to expect that the nonlinear residual reduction will meet the target.

If the reduction target (Eq. (15)) has not been met, and the correction is not dominated by the round-off error, then the Nonlinear Control module declares failure and transfers control to the Failure Handler module. If the residual reduction target is met with a cumulative w_{cum} satisfying $w_{min} \le w_{cum} < 1$, or the round-off dominated update provides any reduction for the rms norm of the nonlinear residual, then the Nonlinear Control module ends with success, but suggests no increase in the CFL number. In case of success, the Nonlinear Control module forms the nonlinear residual corresponding to the next HANIS iteration by subtracting the pseudo-time contribution from the current pseudo-time residual, restores the solution, $\boldsymbol{Q}^n$, terminates the infinite internal loop, and sends the scaled solution update and the suggested CFL update to the Nonlinear Update module. In a tightly-coupled formulation, the nonlinear residuals from the Nonlinear Control module can be used for the subsequent HANIS iteration, thus avoiding an additional residual evaluation.

6. Failure Handler

If any failure occurred within GCR Solver, Realizability Check, or Nonlinear Control modules, the Failure Handler module restores the nonlinear solution and residual to the initial states they had before the current nonlinear iteration. The module also reduces the CFL number, updates Jacobian to account for the CFL number reduction, and returns control to the beginning of the HANIS infinite loop. The mean flow Jacobian is updated through a fast procedure that replaces the pseudo-time contributions in the diagonal terms without recomputing flux linearization.

7. Nonlinear Update

The Nonlinear Update module performs the actual solution update, handles "first-order" designations of faces with violated face-realizability constraints, and prepares the CFL number for the next nonlinear iteration.

IV. Results and Discussion

Two variants of the HANIS scheme are investigated in the present study. The variants differ in the maximum number of search directions allowed for the GCR solver. The variant HANIS-1 uses one search direction only,

whereas, the variant HANIS-4 uses a maximum of four search directions. The PA iteration scheme serves as a benchmark that broadly represents the baseline solver technology.

The performance of HANIS, as implemented in the mixed element version of USM3D, is assessed using four viscous turbulent flow cases. The cases are a zero-pressure-gradient flat plate, a bump-in-channel configuration, the NACA 0012 airfoil, and a NASA Common Research Model (CRM) configuration. For each case, solutions are computed on several grids using three different methods, namely, PA, HANIS-1, and HANIS-4.

The first three cases are quasi 2D, having one hexahedral cell in the spanwise direction. For each of these cases, a series of nested hexahedral grids is available on the TMR website [21] and is used to compute solutions at low subsonic flow conditions. The fourth case, namely a NASA CRM configuration [30], is a three-dimensional (3D) case. For this case, four different grids are used to compute solutions at transonic flow conditions. Figures 2-5 provide a schematic representation for these cases including the configuration, computational domain, and boundary conditions.

The SA turbulence model is used for all solutions reported herein. An identical CFL number is used for both mean flow and SA model solutions. The start up CFL number is 1 for all solution methods. In a PA solution the CFL number is specified a priori for any iteration. The CFL number is ramped from a value of 1 to 150 over 150 iterations. The HANIS-1 and HANIS-4 methods update CFL number adaptively.

Some of the parameters used presently are chosen based on a few precursor investigations. In USM3D, both, the standard and negative versions of SA model are implemented. The negative model formulation allows the SA solution to become negative whereas, the standard model solution is typically restricted above a positive minimum value. USM3D uses a minimum value of 1.0×10^{-10}. In the first precursor study related to the SA model, the NACA 4412 airfoil ($M_\infty = 0.09$, $\alpha = 13.87°$, $Re_L = 1.52 \times 10^6$) PA solutions were computed using both variants of the SA model on a coarse 2x113x33 hexahedral grid. The SA model primary variable gravitated toward zero in these solutions. In both solutions the mean flow residual convergence was very similar. The SA model residual convergence stalled for the standard model. The residuals of the negative SA model converged satisfactorily (Fig. 6).

A parametric study for an optimal value of the maximum number of G-S iterations in the preconditioner was conducted on the NACA 0012 airfoil ($M_\infty = 0.15$, $\alpha = 10°$, $Re_L = 6 \times 10^6$). The PA and HANIS-1 methods were used to compute solutions on a 2x897x257 hexahedral grid. The preconditioner residual reduction target was set to 0.1. Four different values (15, 50, 200, and 500) for the maximum number of G-S iterations were considered. The study showed that a maximum of 500 G-S iterations significantly improves the performance of the HANIS-1 method. The PA solutions were relatively insensitive to the maximum-number variations (Fig. 7). The trend is logical considering that the PA method uses a lower CFL number thereby allowing the preconditioner to satisfy the linear residual reduction target with fewer than the prescribed maximum number of G-S iterations. A HANIS method generally operates at much higher CFL number that necessitates more G-S iterations for achieving the same linear residual reduction target. Figure 7 also highlights the importance of a good preconditoner in any HANIS algorithm.

Solution accuracy, robustness, and time-to-convergence are important metrics for any practical CFD solvers, as an applied aerodynamicist may need to compute hundreds to thousands of solutions to design and analyze complex configurations. High-fidelity solutions are typically computed using billions of degrees of freedom. The HANIS algorithm described in Section III has a significantly higher operation count as compared to the PA method for a given nonlinear iteration. Hence, for a significant improvement in convergence time, HANIS algorithm needs to operate at a much higher CFL number on an average. Keeping this in mind, in yet another precursor study for the NACA 0012 airfoil ($M_\infty = 0.15$, $\alpha = 10°$, $Re_L = 6 \times 10^6$) on a 2x897x257 hexahedral grid, a HANIS-4 solution was computed without limiting CFL number. The mean flow and SA model residual convergence stalled in this solution. After observing an operational range of CFL numbers in this stalled solution, another HANIS-4 solution was obtained in which CFL number was limited below 10,000. The latter solution exhibited better convergence (Fig. 8). It should be noted that the corresponding HANIS-1 solutions (not shown) behaved as expected; the solution using an unlimited CFL number converged faster than the one using a maximum CFL number of 10,000. The mean flow and SA model residuals were reduced to machine-zero level in both of the HANIS-1 solutions. It is not fully understood why the HANIS-4 solution stalled and a detailed investigation is pending. At this time, an unlimited CFL number for HANIS-1 and a maximum CFL number of 10,000 for HANIS-4 solutions is deemed a prudent strategy.

A detailed comparative study of PA, HANIS-1, and HANIS-4 performance follows for each of the four cases mentioned earlier. Unless specified otherwise, solutions are obtained using the same set of common input parameters. Some of these parameters are identified based on precursor studies and others are chosen based on past experience and judgment. A brief summary of these input parameters is provided below.

American Institute of Aeronautics and Astronautics

1. *Spatial accuracy for mean flow: second-order*
2. *Spatial accuracy for SA model: first-order*
3. *Cell gradients: Green-Gauss and nodal averaging scheme*
4. *Face gradients: Mitchell's method*
5. *Mean flow convective flux: Roe's FDS*
6. *Turbulence model: Negative SA model*
7. *Approximate Jacobian for the preconditioner: Roe's FDS*
8. *Maximum number of G-S iterations for preconditioner: 500*
9. *Linear residual reduction target for preconditioner: 0.1*
10. *Linear residual reduction target for GCR: 0.92*
11. *CFL number: Ramped from 1 to 150 over 150 iterations for PA; updated adaptively for HANIS with no limit for HANIS-1, a maximum of 10,000 for HANIS-4*

For each of the four cases discussed below, solution convergence is monitored by tracking the mean flow and SA model residuals as well as certain field and integrated surface quantities, such as maximum eddy viscosity, aerodynamic forces and pitching moment. Figures are included that show the convergence of these quantities with respect to run time. Variations of a composite solution residual, generated by combining the mean flow and SA model residuals, as well as the CFL number are plotted as well. The combined residuals are shown, as they are the main driver for the tightly-coupled HANIS formulation investigated here. Run time needed to converge various quantities to a specific tolerance is documented in the corresponding tables. The tables also specify the tolerance for each quantity.

A. 2D Zero Pressure Gradient Flat Plate

Four progressively finer nested structured hexahedral grids from the TMR "Turbulence Model Verification Cases and Grids" section [21] have been used for this case. The specific grids, denoted by the number of grid points in the spanwise, streamwise, and body-normal directions, are 2x69x49, 2x137x97, 2x273x193, and 2x545x385. Figure 2 presents a schematic view of the flat plate configuration, computational domain and the boundary conditions in the X-Z plane. Symmetry boundary condition is applied on the two spanwise planes. As noted earlier, the present USM3D version does not allow intersecting symmetry boundaries. Therefore in a departure from the suggested boundary conditions on the TMR website [21], a slip wall boundary condition is used on a Z=0 plane upstream of the flat plate (Fig. 2). This change may affect solutions reported herein. However, note that the current study is focused on assessing the new developments from the standpoint of iterative convergence.

Solutions have been computed for the freestream Mach number of 0.2 and Reynolds number of $5x10^6$ per unit length of the plate. The angle-of-attack is 0°. Two different sets of solutions are computed; in the first set the SA model convective term is first-order accurate, in the second set it is second-order accurate. Each set consists of the PA, HANIS-1, and HANIS-4 solutions on all four grids described earlier. For each set, solution convergence is demonstrated in a separate figure for each grid. A figure includes plots for the run time variations of six quantities, namely, (a) mean flow residuals, (b) SA model residuals, (c) combined mean flow and SA model residuals, (d) CFL number, (e) maximum eddy viscosity, and (f) viscous drag coefficient.

Figures 9-12 show the convergence of various quantities in the first solution set for the 2x69x49, 2x137x97, 2x273x193, and 2x545x385 grids, respectively. It is seen that HANIS solutions converge much faster than the PA solutions, and HANIS-4 operates at a higher CFL number and converges faster than HANIS-1. Only exception to this observed trend is somewhat superior convergence of maximum eddy viscosity and viscous drag in the HANIS-1 solution on the 2x69x49 grid (Fig. 9). Tables 1a and 1b below summarize the run time needed on various grids for the PA and HANIS methods to converge the coupled residuals and viscous drag to a specific level.

Table 1a. Residual error levels and convergence time for flat plate solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Residual level	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	**HANIS-1**	**HANIS-4**	**HANIS-1**	**HANIS-4**
2x69x49	1.0e-15	233	48	42	4.9	5.5
2x137x97	1.0e-14	1362	253	198	5.4	6.9
2x273x193	1.0e-15	13973	6981	3190	2.0	4.4
2x545x385	1.0e-14	118356	35481	15492	3.3	7.6

American Institute of Aeronautics and Astronautics

Table 1b. Viscous drag coefficient values and convergence time for flat plate solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Viscous drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x69x49	0.289279e-2	99	19	26	5.2	3.8
2x137x97	0.287039e-2	769	154	131	5.0	5.9
2x273x193	0.286291e-2	6937	1913	1117	3.6	6.2
2x545x385	0.286047e-2	70881	18157	10383	3.9	6.8

It can be seen from the Tables 1a and 1b that as compared to PA, HANIS-4 converges 3.8-7.6 faster, whereas, HANIS-1 converges between 2.0-5.4 times faster. The speed up factor of 2.0 is observed only in one computation on 2x273x193 grids and only for the residual convergence. All other speed up factors are greater than 3.3.

The convergence of the second set of solutions generated using a second-order approximation for the SA model convective term is presented in the Figs. 13-16 for the 2x69x49, 2x137x97, 2x273x193, and 2x545x385 grids, respectively. Figure 13 shows aberrant solutions on 2x69x49 grid. At certain locations in this grid the primary SA model variable became negative in solutions from all the three methods. In the PA solution, the primary variable in certain cells oscillated between negative and non-negative values throughout the run. In HANIS solutions these oscillations occurred for a shorter duration in the evolutionary stage and vanished later on. The PA solution exhibited limit-cycle behavior for the entire duration of the run. The HANIS-1 and HANIS-4 successfully exited from the stall behavior. On this grid, HANIS-1 solution operated at a lower CFL number on an average basis. Although, HANIS-1 and HANIS-4 solutions converged to machine-zero level, viscous drag coefficient from the two solutions differed slightly (0.288807e-2 for HANIS-1, 0.288800e-2 for HANIS-4) indicating non-unique solutions.

Figures 14-16 show that HANIS solutions substantially outperform PA solutions on other finer grids as well. HANIS-4 generally operates at higher CFL number and converges faster than HANIS-1. Table 1c and Table 1d summarize the time taken by each solution method to converge the coupled residuals and viscous drag to a specific level. It can be seen from the tables that as compared to PA, HANIS-4 speeds up convergence by a factor 2.2-8.3, and HANIS-1 speeds up convergence by a factor 2.2-4.8. The smallest speed up factor of 2.2 is observed only for the SA model (and combined) residual convergence and only on the 2x273x193 grid. All other speed up factors are greater than 3.0.

Table 1c. Residual error levels and convergence time for flat plate solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is second-order-accurate.

Grid	Residual level	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x69x49	1.0e-15	limit cycle	152	204	N/A	N/A
2x137x97	1.0e-14	1371	453	235	3.0	5.8
2x273x193	1.0e-15	14322	6462	6427	2.2	2.2
2x545x385	1.0e-14	120148	36586	14422	3.3	8.3

Table 1d. Viscous drag coefficient values and convergence time for flat plate solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is second-order-accurate.

Grid	Viscous drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x69x49	non-unique	limit cycle	95	72	N/A	N/A
2x137x97	0.286881e-2	711	147	131	4.8	5.4
2x273x193	0.286266e-2	6742	1803	1049	3.7	6.4
2x545x385	0.286057e-2	73863	17098	9808	4.3	7.5

B. 2D Bump-in-Channel

Five progressively finer nested hexahedral grids [21] have been used for this case. The specific grids, denoted by the number of grid points in the spanwise, streamwise, and body-normal directions, are 2x89x41, 2x177x81, 2x353x161, 2x705x321, and 2x1409x641. Figure 3 presents a schematic view of the bump configuration,

computational domain and the boundary conditions in the X-Z plane. Symmetry boundary condition is applied on the two spanwise planes. In a departure from the suggested boundary conditions on the TMR website [21], a slip wall boundary condition is applied on a Z=0 plane upstream and downstream of the bump (Fig. 3).

Solutions have been computed for the freestream Mach number of 0.2, Reynolds number of $3x10^6$ per unit grid length, and angle-of-attack of 0°. On each of the five grids, solutions are computed using three different methods, namely, the PA, HANIS-1, and HANIS-4. A first-order approximation is applied for the SA model convective term.

Iterative convergence is demonstrated in a separate figure for each grid. A figure includes plots for the run time variations of eight quantities, namely, (a) mean flow residuals, (b) SA model residuals, (c) combined mean flow and SA model residuals, (d) CFL number, (e) maximum eddy viscosity, (f) drag coefficient (g) lift coefficient, and (h) viscous drag coefficient. Figures 17-21 show the convergence of various quantities for the 2x89x41, 2x177x81, 2x353x161, 2x705x321, and 2x1409x641 grids, respectively. The PA solution on the 2x1409x641 grid is not fully converged yet, as seen from Fig. 21. It is evident that on all grids, HANIS solutions converge substantially faster than PA solutions.

HANIS-1 operates at higher levels of CFL number on all grids, as compared to HANIS-4. On the coarsest grid, CFL number in HANIS-1 solution responds to round-off errors and settles to a lower level in the post-convergence phase of the solution. It is noted that on all grids, HANIS-4 operates at the maximum-allowed CFL number of 10,000 almost throughout the solution cycle. HANIS solutions on finer grids exhibit highly oscillatory variations of aerodynamic forces prior to convergence. On the contrary, PA solutions show a monotonic (albeit slow) convergence of aerodynamic forces. This is expected because HANIS solutions, unlike PA solutions, use a much higher CFL number.

A more quantified evaluation of the PA, HANIS-1, and HANIS-4 solution convergence is presented in Tables 2a-2d for all grids. The tables summarize the time taken by each solution method to converge the combined residuals (Table 2a), drag coefficient (Table 2b), lift coefficient (Table 2c), and viscous drag coefficient (Table 2d) to a specific level. It can be readily seen from the tables that on first four grids, as compared to PA, HANIS-1 speeds up convergence by a factor 2.8-7.2, and HANIS-4 speeds up convergence by a factor 2.2-9.1. The largest speed up factor is observed on the coarsest grid, for either of the HANIS solutions. On the next three progressively finer grids, HANIS speed factors reduce gradually. Note that on the finest 2x1409x641 grid, HANIS-1 and HANIS-4 solutions have already converged, whereas, the PA solution is yet to converge. HANIS solutions are expected to indicate a speed up of about 3 over the converged PA solution on this grid.

Table 2a. Residual error levels and convergence time for bump-in-channel solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Residual level	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x89x41	1.0e-13	344	48	44	7.2	7.8
2x177x81	1.0e-13	2412	468	434	5.2	5.6
2x353x161	1.0e-13	23988	6183	7069	3.9	3.4
2x705x321	1.0e-13	343911	87431	103132	3.9	3.3
2x1409x641	1.0e-13	yet to converge	1247370	1250940	N/A	N/A

Table 2b. Drag coefficient values and convergence time for bump-in-channel solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x89x41	0.493627e-2	218	33	24	6.6	9.1
2x177x81	0.376135e-2	1240	283	252	4.4	4.9
2x353x161	0.360087e-2	13006	3578	4106	3.6	3.2
2x705x321	0.357887e-2	157939	47382	69754	3.3	2.3
2x1409x641	0.357386e-2	yet to converge	1008880	965300	N/A	N/A

Table 2c. Lift coefficient values and convergence time for bump-in-channel solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Lift coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x89x41	0.243922e-1	122	23	20	5.3	6.1
2x177x81	0.248145e-1	1018	243	217	4.2	4.7
2x353x161	0.248967e-1	10199	3461	4540	2.9	2.2
2x705x321	0.249228e-1	163454	48239	69492	3.4	2.4
2x1409x641	0.249456e-1	yet to converge	960810	902295	N/A	N/A

Table 2d. Viscous drag coefficient values and convergence time for bump-in-channel solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Viscous drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x89x41	0.327116e-2	157	22	21	7.1	7.5
2x177x81	0.322682e-2	1184	265	242	4.5	4.9
2x353x161	0.320438e-2	12455	3513	3160	3.5	3.9
2x705x321	0.319627e-2	152789	53755	50608	2.8	3.0
2x1409x641	0.319266e-2	yet to converge	806879	808462	N/A	N/A

C. 2D NACA 0012 Airfoil

Grids from the TMR "Cases and Grids for Turbulence Model Numerical Analysis" section [21] have been used for this case. Grids of three families, namely, Family I, Family II, and Family III are available. The families differ in the trailing edge streamwise spacing. The current study uses five of the seven nested structured grids of Family II. Family II grids have been used presently as they provide the highest resolution around the trailing edge. The specific grids, denoted by the number of grid points in the spanwise, streamwise, and body-normal directions, are 2x113x33, 2x225x65, 2x449x129, 2x897x257, and 2x1793x513. The current study does not use the finer two grids (2x3585x1025 and 2x7169x2049) in the seven grid series. Figure 4 presents a schematic view of the NACA 0012 airfoil, computational domain and the boundary conditions in the X-Z plane. Symmetry boundary condition is applied on the two spanwise planes.

Solutions have been computed for the freestream Mach number of 0.15, Reynolds number of $6x10^6$ per unit chord length, and angle-of-attack of 10°. On each of the five grids, solutions are computed using three different methods, namely, the PA, HANIS-1, and HANIS-4. A first-order approximation is applied for the SA model convective term.

Iterative convergence is demonstrated in a separate figure for each grid. A figure includes plots for the run time variations of eight quantities, namely, (a) mean flow residuals, (b) SA model residuals, (c) combined mean flow and SA model residuals, (d) CFL number, (e) maximum eddy viscosity, (f) drag coefficient (g) lift coefficient, and (h) pitching moment coefficient. Figures 22-26 present the convergence of various quantities for the 2x113x33, 2x225x65, 2x449x129, 2x897x257, and 2x1793x513 grids, respectively. It is evident that on all grids, HANIS solutions converge substantially faster than PA solutions.

Both HANIS-1 and HANIS-4 methods operate on an average at similar levels of CFL number for all grids, except the 2x225x65 grid (Fig. 23), where HANIS-1 solution shows higher levels of CFL number. It is noted that on all grids, HANIS-4 operates at the maximum-allowed CFL number of 10,000 almost throughout the solution cycle.

A more quantified evaluation of the PA, HANIS-1, and HANIS-4 convergence is presented in Tables 3a-3d for all grids. The tables summarize the time taken by each solution method to converge the combined residuals (Table 3a), drag coefficient (Table 3b), lift coefficient (Table 3c), and pitching moment coefficient (Table 3d) to a specific level. Note that pitching moment is computed with respect to the airfoil leading edge. It can be readily seen from the tables that as compared to PA, HANIS-1 speeds up convergence by a factor 1.8-11, and HANIS-4 speeds up convergence by a factor 3.6-13.0. The largest HANIS speed up is derived on the 2x225x65 grid, whereas, the smallest speed up is achieved on the 2x1793x513 grid. Disregarding the solution times on the coarsest 2x113x33 grid, it can be seen from the Tables 3a-3d that HANIS speed up gradually diminishes on the progressively finer grids. HANIS-4 significantly outperforms HANIS-1 on the finest two grids of the current study.

Table 3a. Residual error levels and convergence time for NACA 0012 solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Residual level	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x113x33	1.0e-12	99	14	13	7.1	7.6
2x225x65	1.0e-13	905	82	69	11.0	13.0
2x449x129	1.0e-13	8558	1317	1076	6.5	8.0
2x897x257	1.0e-13	94553	22187	16413	4.3	5.8
2x1793x513	1.0e-13	1585630	645271	309561	2.5	5.1

Table 3b. Drag coefficient values and convergence time for NACA 0012 solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x113x33	0.209708e-01	29	9	8	3.2	3.6
2x225x65	0.145285e-01	215	50	40	4.3	5.4
2x449x129	0.127287e-01	2693	732	610	3.7	4.4
2x897x257	0.123448e-01	39559	14268	10417	2.8	3.8
2x1793x513	0.122646e-01	762022	416972	211012	1.8	3.6

Table 3c. Lift coefficient values and convergence time for NACA 0012 solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Lift coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x113x33	0.101120e+01	46	8	8	5.8	5.8
2x225x65	0.107610e+01	198	38	29	5.2	6.8
2x449x129	0.108777e+01	2040	475	426	4.3	4.8
2x897x257	0.108947e+01	32286	6620	5710	4.9	5.7
2x1793x513	0.109001e+01	673057	318125	163592	2.1	4.1

Table 3d. Pitching moment values and convergence time for NACA 0012 solutions obtained using a sequence of hexahedral grids. Convective term in the SA model is first-order-accurate.

Grid	Pitching moment coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
2x113x33	-0.237382	28	9	7	3.1	4.0
2x225x65	-0.258021	326	43	38	7.6	8.6
2x449x129	-0.261365	2422	675	503	3.6	4.8
2x897x257	-0.261769	49115	13059	9179	3.8	5.4
2x1793x513	-0.261936	781751	397217	190948	2.0	4.1

D. NASA Common Research Model (CRM)

The NASA CRM wing/body/horizontal-tail configuration used presently has been the focus of the AIAA 4th Drag Prediction Workshop (DPW-4) [31]. The specific configuration that has tail incidence angle, $i_H = 0°$, is used in the current study. The DPW-4 organizing committee made several different families of the structured and unstructured grids available. One of the grid families, known as "VGRID LaRC" that consisted of only tetrahedral elements, was generated using VGRID [6]. The family consisted of two different sets; one set of grids was suitable for the cell-centered and another was suitable for the node-centered CFD solvers. The "VGRID LaRC (cell-centered)" grid set has four grids, namely, the "coarse", "medium", "fine", and "extra-fine" grids. These grids can be downloaded from a server (ftp://cmb24.larc.nasa.gov/outgoing/DPW4/unstructured_Larc/CellBase/) at NASA Langley Research Center.

Four different grids used for the present study are described next. A "coarse tetrahedral" grid is identical to the "coarse" "VGRID LaRC (cell-centered)" grid. The grid has 3,935,055 cells and 672,235 nodes. Other three grids are mixed-element grids. In the "coarse mixed" grid, boundary layer tetrahedral elements in the "coarse tetrahedral" grid have been merged into prismatic and pyramidal cells, using a utility code developed by Ed Parlette of ViGYAN, Inc. Similarly, "medium mixed", and "fine mixed" grids are derived from the corresponding "VGRID LaRC (cell-centered)" "medium" and "fine" grids, respectively. The "coarse mixed" grid has a total of 1,524,325 elements (1,226,266 prisms, 26,728 pyramids, and 271,331 tetrahedra) and 685,941 nodes. The "medium mixed" grid has a total of 4,025,920 elements (3,057,031 prisms, 47,028 pyramids, and 921,861 tetrahedra), and 1,736,808 nodes. The "fine mixed" grid has a total of 15,773,460 elements (9,630,012 prisms, 102,082 pyramids, and 6,041,366 tetrahedra), and 5,969,072 nodes.

A schematic representation of the NASA CRM configuration, computational domain and the boundary conditions used for the current study is provided in Fig. 5. Only half of the configuration is considered and a symmetry boundary condition is applied on the plane of symmetry.

Solutions have been computed for the freestream Mach number of 0.85, Reynolds number of $5x10^6$ per unit reference chord length, and angle-of-attack of $2°$. On each of the four grids, solutions are computed using three different methods, namely, the PA, HANIS-1, and HANIS-4. A first-order approximation is applied for the SA model convective term.

The objective of this study is to assess HANIS benefits for a practical configuration for which tetrahedral grid USM3D production version is typically used. Hence, best practices developed for the production USM3D version are adopted for PA solutions. For example, for the first 500 nonlinear iterations, solution is first-order accurate and then switched over to second-order accuracy. CFL number is scaled according to the deviation of cell aspect ratio from the ideal value of an isotropic tetrahedron [32]. Only 15 G-S iterations are used for the preconditioner. Furthermore, SA model solution updates are uniformly under-relaxed such that at any cell only 25% of the solution updates are applied. It is emphasized that for HANIS solutions the parameters summarized earlier in Section IV are used.

Solution convergence is demonstrated in a separate figure for each grid. A figure includes plots for the run time variations of eight quantities, namely, (a) mean flow residuals, (b) SA model residuals, (c) combined mean flow and SA model residuals, (d) CFL number, (e) maximum eddy viscosity, (f) drag coefficient (g) lift coefficient, and (h) pitching moment coefficient. Figures 27-30 present the convergence of various quantities for the "coarse tetrahedral", "coarse mixed", "medium mixed", and "fine mixed" grids, respectively. It is evident that on all grids, HANIS solutions converge asymptotically faster than PA solutions. Convergence acceleration from HANIS-1 and HANIS-4 is comparable in the "coarse tetrahedral" grid solutions. On mixed element grids, HANIS-1 significantly outperforms HANIS-4.

On "coarse tetrahedral" grid, HANIS methods operate at an average CFL number below 150, which is the specified global CFL value for PA solutions (Fig. 27). Inability of HANIS solutions to operate at higher levels of CFL number needs further investigation. For the three mixed-element grids, HANIS methods typically operate at significantly higher levels (Figs. 28-30).

A more quantified evaluation of the PA, HANIS-1, and HANIS-4 solution convergence is presented in two sets of tables. The sets use different tolerance settings for measuring solution convergence. Tight tolerance levels are applied in Tables 4a-4d, and somewhat loose tolerance levels, typically used for practical aerodynamic computations, are set in Tables 4e-4h. The tables summarize the time taken by each solution method to converge the combined residuals (Table 4a, Table 4e), drag coefficient (Table 4b, Table 4f), lift coefficient (Table 4c, Table 4g), and pitching moment coefficient (Table 4d, Table 4h) to a specific level.

It is noted that for tight tolerances, as compared to PA, HANIS-1 speeds up convergence by a factor 1.7-5.5, and HANIS-4 speeds up convergence by a factor 1.4-2.9 (Tables 4a-4d). For loose tolerances, as compared to PA, HANIS-1 speeds up convergence by a factor 1.4-4.7 and HANIS-4 speeds up convergence by a factor 0.8-3.0 (Tables 4e-4h). For all quantities, the smallest HANIS-1 speed up factors are observed on "coarse tetrahedral" grid. HANIS-1 performs comparatively better on mixed-element grids. HANIS-4 shows a much-degraded performance on this case and struggles on "coarse tetrahedral" and "fine mixed" grids. For loose tolerances, it converges slower than PA for certain quantities. For the present 3D case HANIS-1 is superior to HANIS-4. It can be surmised that additional search directions used in HANIS-4 merely consume more run time without achieving linear residual reduction targets in GCR.

Table 4a. Residual error levels and convergence time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Residual level	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	1.0e-8	778295	437008	445627	1.8	1.7
coarse mixed	1.0e-8	114359	20665	41867	5.5	2.7
medium mixed	1.0e-8	348447	68654	132401	5.1	2.6
fine mixed	1.0e-8	1397490	326735	809475	4.3	1.7

Table 4b. Drag coefficient values and convergence time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Drag coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	0.295854e-1	328597	172176	173135	1.9	1.9
coarse mixed	0.298837e-1	37657	8507	16417	4.4	2.3
medium mixed	0.266261e-1	126049	26556	68569	4.7	1.8
fine mixed	0.258262e-1	735403	209924	459099	3.5	1.6

Table 4c. Lift coefficient values and convergence time for NASA CRM solutions obtained using a tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Lift coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	0.445167	332666	191796	191588	1.7	1.7
coarse mixed	0.440453	38774	9232	17762	4.2	2.2
medium mixed	0.447966	143822	31908	72442	4.5	2.0
fine mixed	0.448297	714623	183635	458029	3.9	1.6

Table 4d. Pitching moment coefficient values and convergence time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Pitching moment coefficient	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	-0.406453e-1	506875	278983	278114	1.8	1.8
coarse mixed	-0.354578e-1	63064	13090	21846	4.8	2.9
medium mixed	-0.298685e-1	164865	40434	79762	4.1	2.1
fine mixed	-0.252248e-1	841354	246459	601875	3.4	1.4

Table 4e. Residual error reduction and convergence time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Residual reduction by order of magnitude	Time in seconds to converge to the specific level using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	5	142672	99939	99573	1.4	1.4
coarse mixed	5	20239	5030	7700	4.0	2.6
medium mixed	5	63037	16509	33090	3.8	1.9
fine mixed	5	269846	113624	250364	2.4	1.1

American Institute of Aeronautics and Astronautics

Table 4f. Drag coefficient convergence tolerance and time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Drag coefficient tolerance	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	±0.25%	122810	79517	78839	1.5	1.6
coarse mixed	±0.25%	12708	3985	6371	3.2	2.0
medium mixed	±0.25%	38998	13081	21543	3.0	1.8
fine mixed	±0.25%	169632	92264	216982	1.8	0.8

Table 4g. Lift coefficient convergence tolerance and time for NASA CRM solutions obtained using a tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Lift coefficient tolerance	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	±0.25%	108006	78150	77447	1.4	1.4
coarse mixed	±0.25%	12698	4263	7122	3.0	1.8
medium mixed	±0.25%	29122	13815	26865	2.1	1.1
fine mixed	±0.25%	172888	78312	185388	2.2	0.9

Table 4h. Pitching moment coefficient convergence tolerance and time for NASA CRM solutions obtained using tetrahedral and mixed element grids. Convective term in the SA model is first-order-accurate.

Grid	Pitching moment coefficient tolerance	Time in seconds to converge to the specific value using various methods			Speedup factor of HANIS over PA	
		PA	HANIS-1	HANIS-4	HANIS-1	HANIS-4
coarse tetrahedral	±0.25%	229114	133162	132419	1.7	1.7
coarse mixed	±0.25%	24840	5338	8380	4.7	3.0
medium mixed	±0.25%	88050	18647	43041	4.7	2.0
fine mixed	±0.25%	416245	129273	301482	3.2	1.4

V. Concluding Remarks

A new Hierarchical Adaptive Nonlinear Iteration Scheme (HANIS) methodology has been implemented to improve iterative convergence and robustness of the baseline mixed-element USM3D code [20]. The baseline code advances the nonlinear solution in pseudo time using only a simple preconditioner based on a defect-correction scheme and point-implicit Gauss-Seidel (G-S) iterations. A limiting aspect of this approach is that typically, a low Courant-Friedrichs-Lewy (CFL) number is needed to preserve solution robustness. A low CFL number adversely impacts iterative convergence and increases the time needed for converging a solution. The HANIS methodology provides two additional hierarchies over the USM3D preconditioner-alone (PA) solver. The hierarchies are an enhanced linear solver for the exact linearization of Reynolds-Averaged Navier Stokes (RANS) equations and a nonlinear control of the solution update. The preconditioner generates new search directions for a matrix-free Generalized Conjugate Residual (GCR) linear solver that uses Frechet derivatives to compute matrix-vector multiplications. The GCR method suggests an update for the current nonlinear solution. The update is checked for possible violations of realizability constraints and is scaled, if needed, to optimize the nonlinear residual reduction. CFL number adaptation is used as the main tool to provide robustness and improve efficiency. Depending on success reported by each of hierarchies, HANIS may reduce, increase, or maintain CFL number for the subsequent nonlinear iteration; analogously, the suggested update can be discarded entirely, or either fully or partially applied.

Two versions of HANIS, HANIS-1 and HANIS-4, have been assessed. HANIS-1 uses only one search direction in the GCR solver and does not limit operational CFL number; HANIS-4 uses the maximum of four search directions and sets a CFL maximum at 10,000. The performance of the HANIS methods is evaluated against a preconditioner-alone (PA) method representative of the current state-of-the-art solver technology. The PA solver uses a specified CFL ramping throughout several initial iterations and a constant CFL thereafter.

The PA, HANIS-1, and HANIS-4 methods use many advanced solver technology elements, recently implemented in USM3D, such as improved linearization schemes for the defect-correction preconditioner, adaptive discretization schemes for mean flow fluxes, a turbulence model formulation with improved numerical

characteristics, residual reduction targets for an earlier iteration termination, implementation strategies optimizing computer memory and operation count.

The PA and HANIS methods have been applied to four benchmark turbulent flow cases: a two-dimensional (2D) flat plate, a 2D bump in channel, the 2D NACA 0012 airfoil, and a three-dimensional (3D) NASA CRM wing-body-tail configuration used for the 4[th] AIAA Drag Prediction Workshop. The following observations have been made:

1. HANIS methods can significantly accelerate iterative convergence in terms of time-to-solution. In comparison to the PA solver, HANIS-1 always reduces the computational time required to converge either nonlinear residuals or an aerodynamic coefficient to a prescribed tolerance. For 3D computations and for 2D computations on finer grids the typical convergence acceleration factor ranges from 1.4 to 5. The factor is even higher on coarser 2D grids.

2. HANIS algorithms provide the largest acceleration when operating at high CFL numbers. On finer grids, HANIS-1 and HANIS-4 methods perform at a lower CFL level and produce diminished convergence acceleration.

3. HANIS-4 mostly outperformed HANIS-1 in 2D cases. However, HANIS-4 converged poorly on the fine grid for the 3D case. This degradation can be explained by the fact that the additional search directions do not significantly improve convergence of the GCR solver, but spend a significant amount of additional time per iteration.

VI. Future Directions

The present study exposed many avenues for further improvements in the new USM3D. For improved automation, it is important to reduce the number of user-prescribed parameters. Currently the HANIS performance is controlled by many parameters. For example, the preconditioner has a residual reduction target and the maximum number of G-S iterations as parameters. Performance of the current point-implicit G-S preconditioner strongly varies with the variations of these parameters. Softer (closer to 1) targets imply fewer G-S iterations, but less accurate solutions for the preconditioner equations; higher maximum number for G-S iterations improves probability to reach the target, but may increase the time spent in preconditioner. A line-implicit G-S iteration scheme is less sensitive to the variations of these parameters [29]. Eventually, a linear multigrid solver should completely eliminate these dependences, as a single multigrid cycle is expected to be sufficient to meet a reasonable residual reduction target in preconditioner computations on any grid.

A general GCR method may use several search directions and a linear residual reduction target. There is a trade-off between more search directions and a softer target. More search directions potentially provide a better accuracy for linear solutions and better chances to meet the GCR residual reduction target, but use more computational time and memory. Softer targets increase the operational CFL number range, but provide less accurate solution updates, which may potentially lead to a nonlinear convergence degradation. GCR with 1 search direction is a parameter free algorithm, which is a very attractive feature. If the GCR preconditioner met its own residual reduction target, but the corresponding search direction did not provide the sufficient reduction in the GCR residual, then the preconditioner solution is not a good approximation for the GCR solution. Chances are that asking for more search directions from a crude preconditioner may be expensive and may not bring any benefits (as 3D computations seem to indicate). A simpler and often more effective strategy is to declare a GCR failure and reduce CFL number, which automatically improves approximation property between the preconditioner and the GCR equations.

The nonlinear control of the solution update has a nonlinear residual reduction target and a minimum under-relaxation parameter. In the current implementation, the nonlinear residual reduction target is linked to the GCR linear residual reduction target, and it is not a free parameter. Similarly, the minimum under-relaxation parameter is linked to the nonlinear residual reduction target and to the actual reduction of the GCR residual, and it is not a free parameter either. With one search direction for the GCR method and a linear multigrid for the preconditioner, a single HANIS control parameter is left, which is a residual reduction target (either for GCR residual or for nonlinear residual).

In the present HANIS, preconditioner, GCR, and nonlinear residual reduction targets are 0.1, 0.92 and 0.96, respectively. The targets are the same on all grids. With fixed targets, HANIS performance diminishes on finer grids. HANIS solutions show a high rate of increase in the computation time needed to converge on progressively finer grids. For example, for NACA 0012 case, the convergence time for the HANIS solutions on the four finer grids increases by factors 1, 16, 270, and 7869. This degradation is expected because more G-S iterations are required to achieve the same preconditioner residual reduction on finer grids. Thus the HANIS solver with the same parameters typically operates with lower CFL numbers on finer grids. Multigrid for preconditioner will alleviate this problem, but will not solve it completely because convergence rates of defect-correction iterations may still be insufficient on

finer grids, leading to lower CFL operational range and less efficiency. The problem can be addressed either by developing an efficient linear multigrid solver for the GCR equations or by developing an efficient nonlinear multigrid solver. The recent positive experience with multigrid solutions reported in literature [23, 29, 33] suggests the latter route.

Acknowledgments

References

[1] Frink, N. T., Pirzadeh, S. Z., Parikh, P. C., Pandya, M. J., and Bhat, M. K., "The NASA tetrahedral unstructured software system (TetrUSS)," *The Aeronautical Journal*, Vol. 104, No. 1040, 2000, pp. 491-499.

[2] Pandya, M. J., Frink, N. T., Abdol-Hamid, K. S., Samareh, J. A., Parlette, E. B., and Taft, J. R., "Enhancements to TetrUSS for NASA Constellation Program," *Journal of Spacecraft and Rockets*, Vol. 49, No. 4, 2012, pp. 617-631.

[3] Samareh, J. A., "GridTool: A Surface Modeling and Grid Generation Tool," Proceedings of the Workshop on Surface Modeling, Grid Generation, and Related Issues in CFD Solutions, NASA CP–3291, 9–11 May, 1995.

[4] Lohner, R. and Parikh, P., "Three-dimensional grid generation by the advancing front method," *International Journal for Numerical Methods in Fluids*," Vol. 8, Issue 10, 1988, pp. 1135-1149.

[5] Pirzadeh, S. Z., "Three-dimensional unstructured viscous grids by the advancing layer method," *AIAA Journal*, Vol. 33, No. 1, 1996, pp. 43-49.

[6] Pirzadeh, S. Z., "Advanced unstructured grid generation for complex aerodynamic applications," *AIAA Journal*, Vol. 48, No. 5, 2010, pp. 904–915.

[7] Frink, N. T., "Upwind scheme for solving the Euler equations on unstructured tetrahedral meshes," *AIAA Journal*, Vol. 30, No. 1, 1992, pp. 70-77.

[8] Frink, N. T., "Tetrahedral unstructured Navier-Stokes method for turbulent flows," *AIAA Journal*, Vol. 36, No. 11, 1998, pp. 1975-1982.

[9] Pandya, M.J., Frink, N.T., Abdol-Hamid, K.S., and Chung, J.J., "Recent enhancements to USM3D unstructured flow solver for unsteady flows," AIAA Paper 2004-5201.

[10] Abdol-Hamid, K. S., Frink, N. T., Deere, K. A., and Pandya, M. J., "Propulsion Simulations Using Advanced Turbulence Models with the Unstructured-Grid CFD Tool, TetrUSS," AIAA Paper 2004-0714.

[11] Pandya, M. J., Abdol-Hamid, K. S., and Frink, N. T., "Enhancement of USM3D unstructured flow solver for high-speed high-temperature shear flows," AIAA Paper 2009-1329.

[12] Pandya, M.J., Frink, N.T., and Noack, R.W., "Overset-grid moving body capability in the USM3D unstructured flow solver," AIAA Paper 2005-5118.

[13] Campbell, R. L., "Efficient Viscous Design of Realistic Aircraft Configuration (Invited)," AIAA Paper 1998–2539.

[14] Bauer, S. X., Krist, S. E., and Compton, W. B. "Generation of the Ares I-X Flight Test Vehicle Aerodynamic Data Book and Comparison To Flight," AIAA Paper 2011–0011.

[15] Abdol-Hamid, K. S., Ghaffari, F., and Parlette, E. B., "Ares I Vehicle Computed Turbulent Ascent Aerodynamic Data Development and Analysis," *Journal of Spacecraft and Rockets*, Vol. 49, No. 4, 2012, pp. 596-608.

[16] Green, B. E., "Computational Prediction of Nose-Down Control for F/A-18E at High Alpha," *Journal of Aircraft*, Vol. 45, No. 5, 2008, pp. 1661-1668.

[17] Hooker, J. R., Wick, A., Zeune, C., Agelastos, A., "Over Wing Nacelle Installations for Improved Energy Efficiency", AIAA 2103-2920.

[18] Thompson, J. R., Frink, N. T., and Murphy, P. C., "Guidelines for Computing Longitudinal Dynamic Stability Characteristics of a Subsonic Transport," AIAA Paper 2010-4819.

[19] Frink, N. T., "Recent Progress Toward a Three-Dimensional Unstructured Navier-Stokes Flow Solver," AIAA Paper 1994-0061.

[20] Pandya, M. J., Frink, N. T., Ejiang Ding, and Edward Parlette, "Toward Verification of USM3D Extensions for Mixed Element Grids," AIAA Paper 2013-2541.

[21] Turbulence Modeling Resource website, URL http://turbmodels.larc.nasa.gov [cited 24 October 2014].

[22] Allmaras, S. R., Johnson, F. T., and Spalart, P. R., "Modifications and Clarifications for the Implementation of the Spalart-Allmaras Turbulence Model," Seventh International Conference on Computational Fluid Dynamics, Big Island, Hawaii, 2012.

[23] Diskin B. and Nishikawa H., "Evaluation of Multigrid Solutions for Turbulent Flows," AIAA-2014-0082.

[24] Knoll, D. A. and Keyes, D. E., "Jacobian-free Newton–Krylov Methods: A Survey of Approaches and Applications," Journal of Computational Physics 193 (2004) 357–397.

[25] Van der Vorst H. A. and Vuik C. "GMRESR: A Family of Nested GMRES Methods," Numerical Linear Algebra with Applications, Vol. 1, pp. 369-386, 1994.

[26] Lucas, P., van Zuijlen, A. H., and Bijl, H., "Fast Unsteady Flow Computations with a Jacobian-free Newton-Krylov Algorithm," Journal of Computational Physics, Vol. 229, No. 24, 2010, pp. 9201-9215.

[27] Ceze, M. and Fidkowski, K., "A Robust Adaptive Solution Strategy for High-Order Implicit CFD solvers," 20th AIAA Computational Fluid Dynamics Conference, AIAA Paper 2011-3696, Hawaii, 2011.

[28] Allmaras, S. R., Bussoletti, 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-Verlag, 2009, pp. 1-19.

[29] Nishikawa, H., Diskin, B., Thomas, J. L., and Hammond, D. P., "Recent Advances in Agglomerated Multigrid," AIAA Paper 2013-863.

[30] Vassberg, J. C., DeHann, M. A., Rivers, S. M., and Wahls, R. A., "Development of a Common Research Model for Applied CFD Validation Studies," AIAA Paper 2008-6919.

[31] 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," AIAA Paper 2010-4547.

[32] Frink, N. T., "Assessment of an Unstructured-Grid Method for Predicting 3-D Turbulent Viscous Flows," AIAA Paper 1996-0292.

[33] Langer, S., Schwoppe, and Kroll, N., "The DLR Flow Solver TAU – Status and Recent Algorithmic Developments," AIAA Paper 2014-0080.

This page is intentionally blank.

Begin a nonlinear iteration

nonlinear residuals
pseudo-time step
approximate Jacobian

Preconditioner

Nonlinear Update

(a) PA iterations

Begin a nonlinear iteration

nonlinear residuals
pseudo-time step
approximate Jacobian

Preconditioner

GCR Solver

GCR
failed?

Yes

No

Realizability Check

Realizability Check
failed?

Yes

No

Nonlinear Control

Nonlinear Control
failed?

Yes

No

Failure Handler

Nonlinear Update

(b) HANIS iterations

Figure 1. Flowchart of two nonlinear iteration schemes.
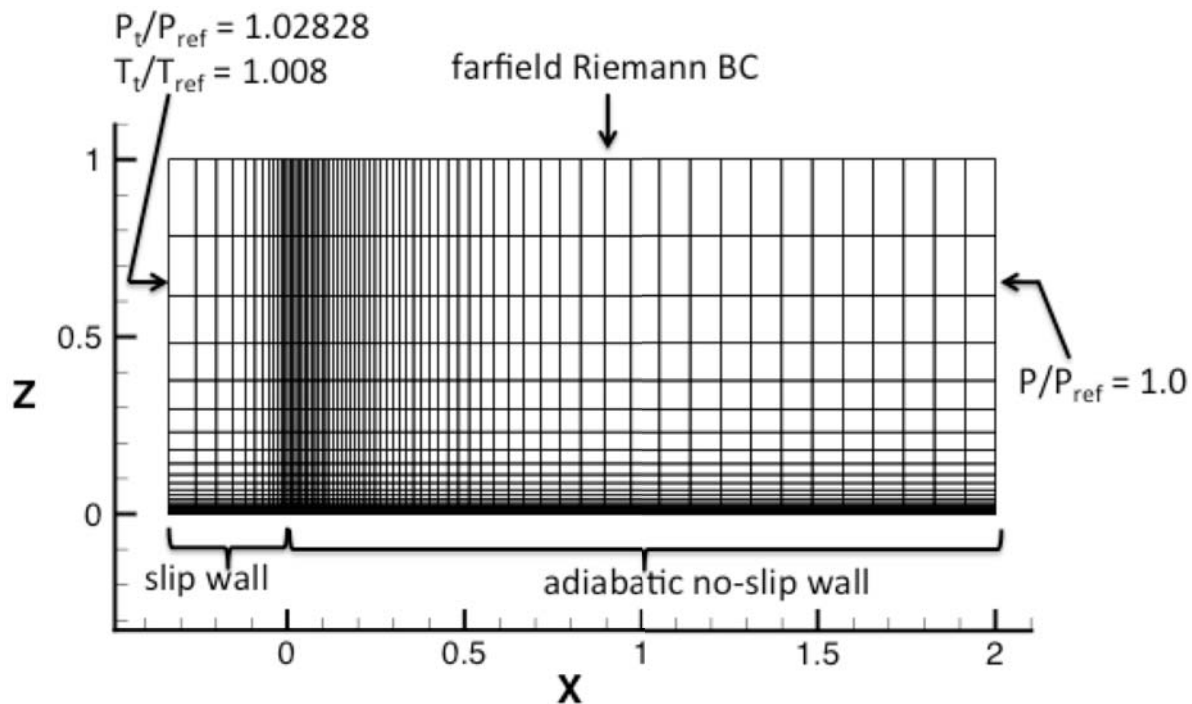
Figure 2. Schematic representation of the 2D zero pressure gradient flat plate configuration and its computational domain with boundary conditions (BC).
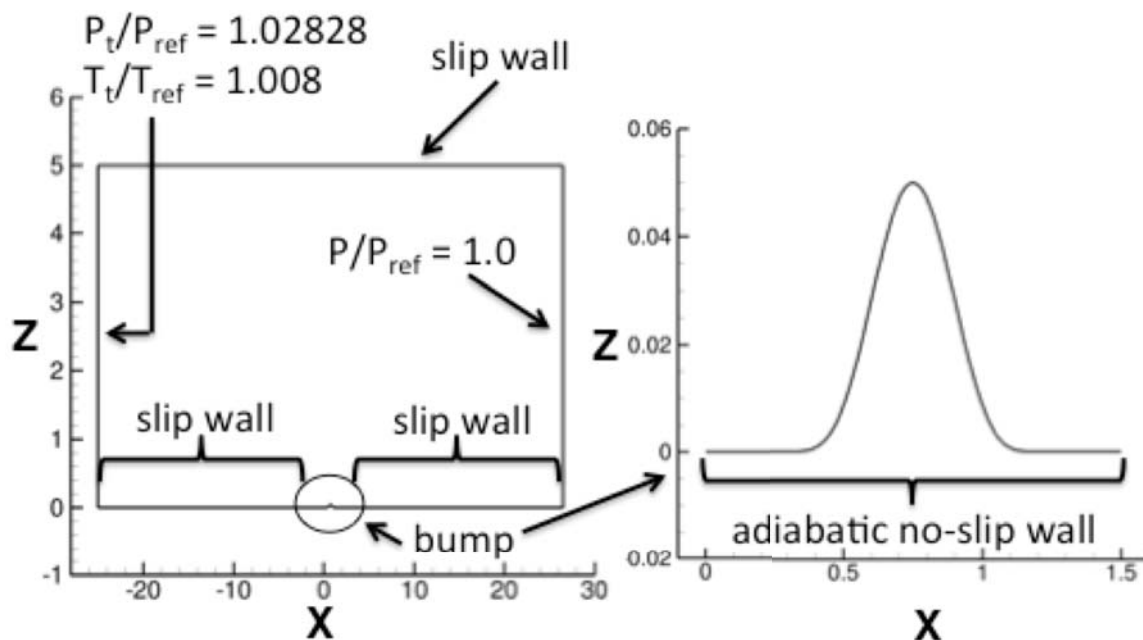
Figure 3. Schematic representation of the 2D bump-in-channel configuration and its computational domain with boundary conditions (BC).

American Institute of Aeronautics and Astronautics
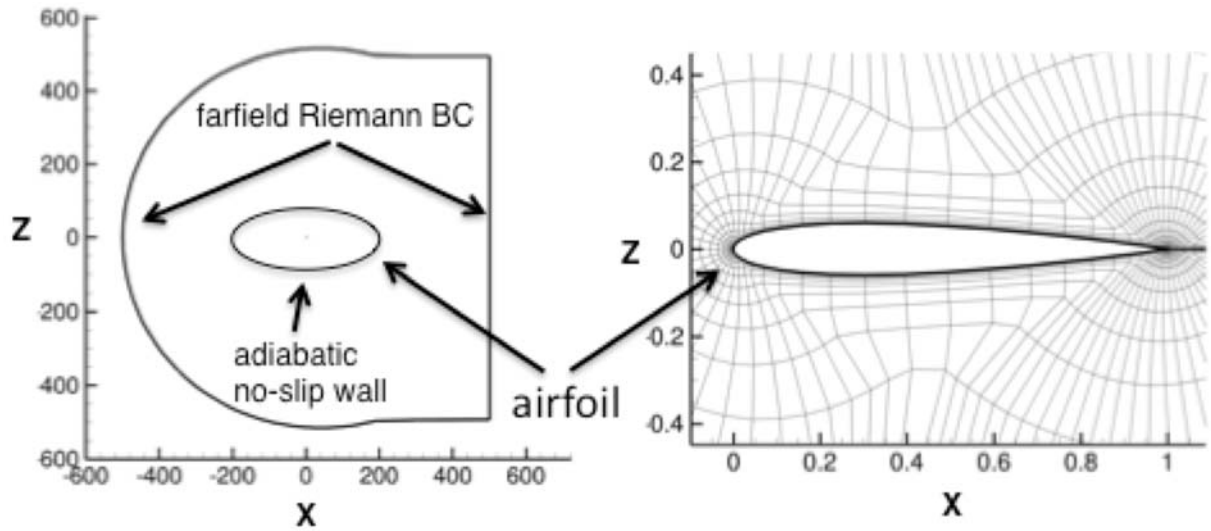
Figure 4. Schematic representation of the 2D NACA 0012 airfoil configuration and its computational domain with boundary conditions (BC).
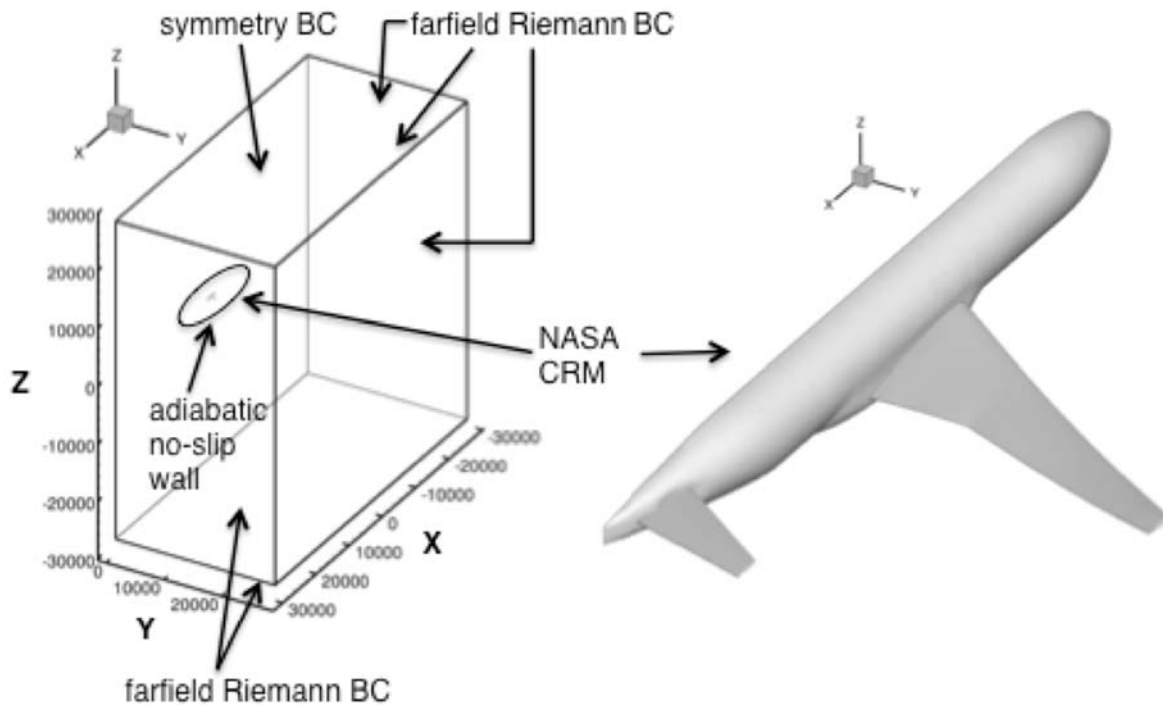
Figure 5. Schematic representation of the NASA Common Research Model (CRM) wing-body-tail configuration and its computational domain with boundary conditions (BC).
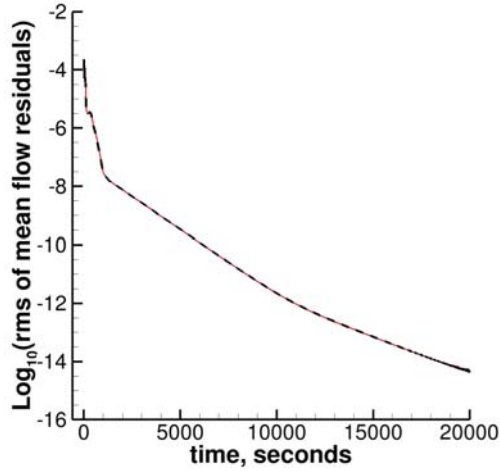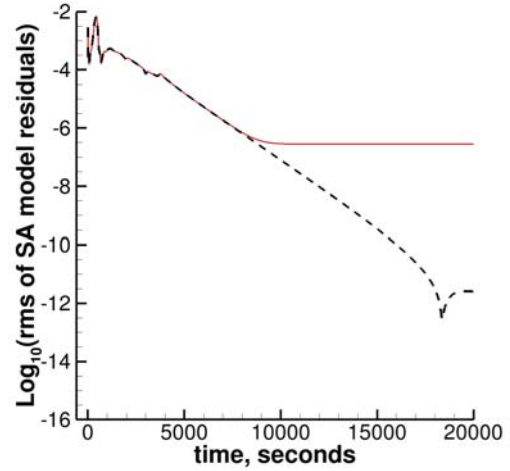
(a) rms of mean flow residuals (b) rms of SA model residuals

—————— **Standard SA model**

– – – – **Negative SA model**

Figure 6. NACA 4412 airfoil PA solution convergence history for 2x113x33 hexahedral grid using standard and negative variants of SA model. $M_\infty = 0.09$, $\alpha = 13.87°$, $Re_L = 1.52x10^6$. Convective term in SA model is first-order accurate.
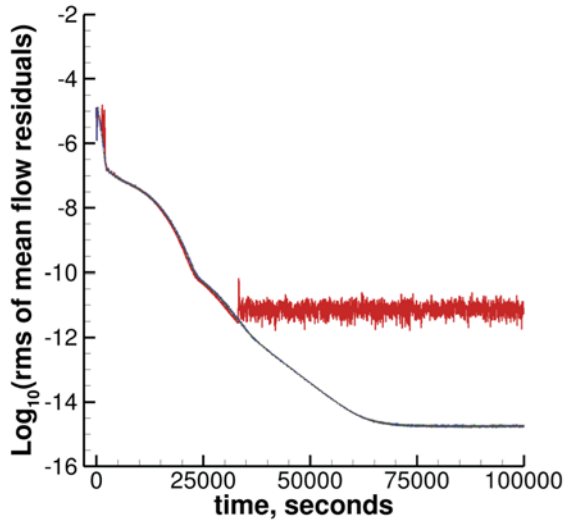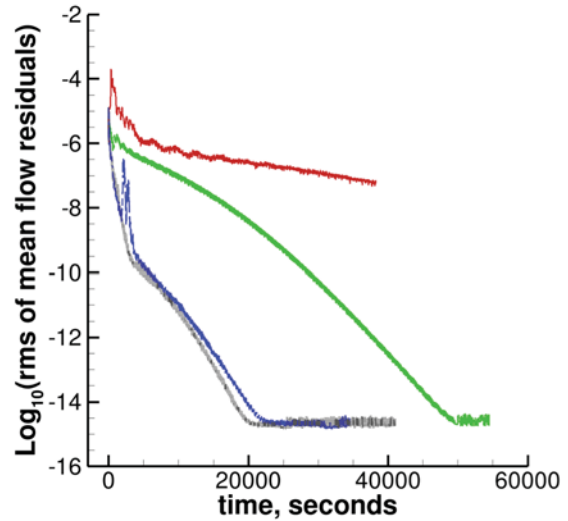
(a) PA (b) HANIS-1

—————— **15 iterations**

– – – – **50 iterations**

·–··–·· **200 iterations**

—————— **500 iterations**

Figure 7. NACA 0012 airfoil mean flow solution convergence history for 2x897x257 hexahedral grid using various number of G-S iterations in the preconditioner. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6x10^6$. Convective term in SA model is first-order accurate.

American Institute of Aeronautics and Astronautics

(a) SA model residual (b) CFL number

Maximum CFL unlimited
Maximum CFL 10000
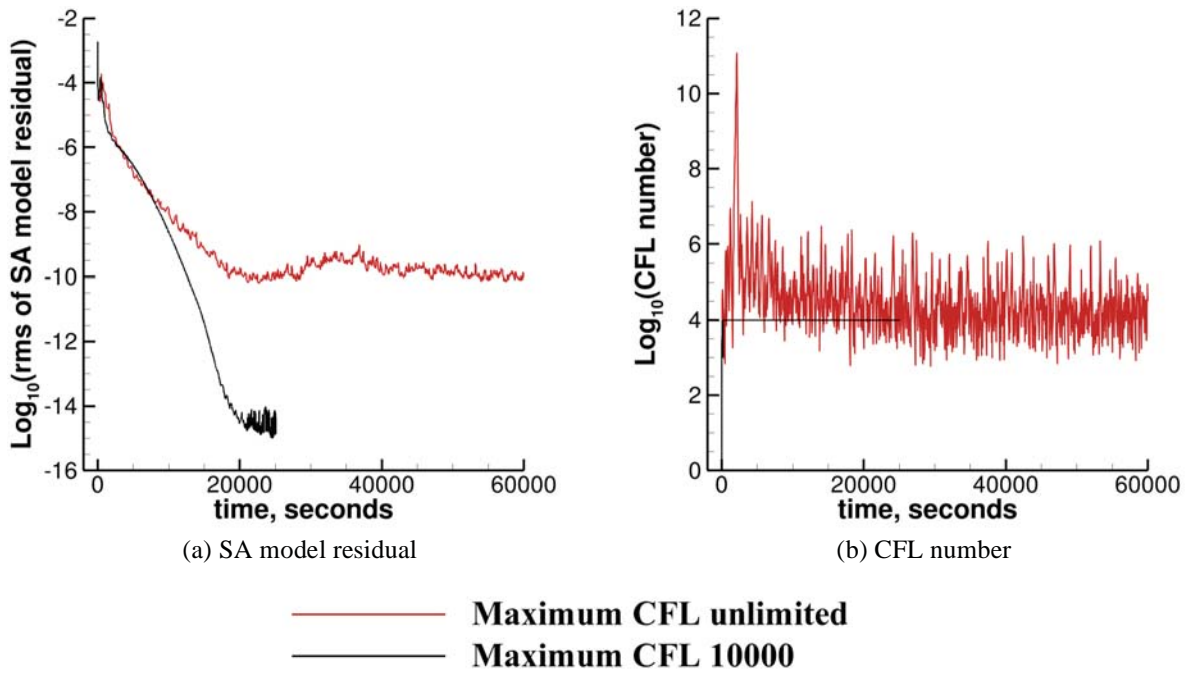
Figure 8. NACA 0012 airfoil HANIS-4 solution convergence history for 2x897x257 hexahedral grid using two different strategies for the maximum CFL number. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6x10^6$. Convective term in SA model is first-order accurate.
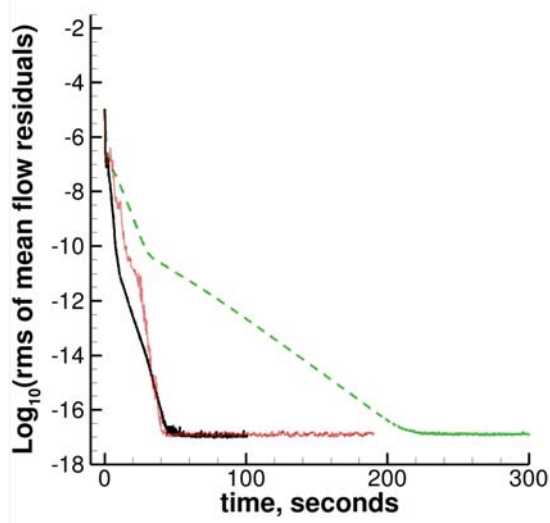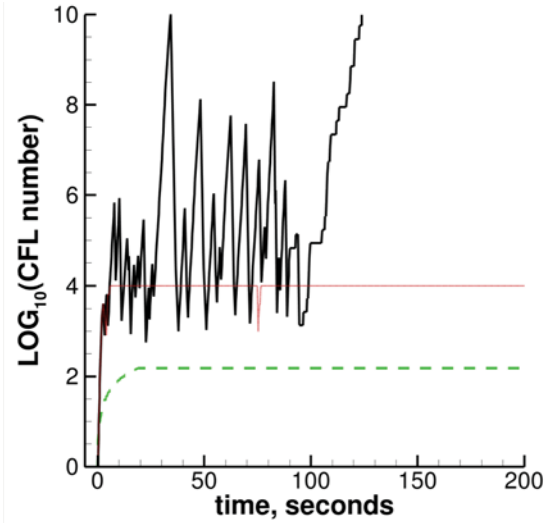
(a) rms of mean flow residuals
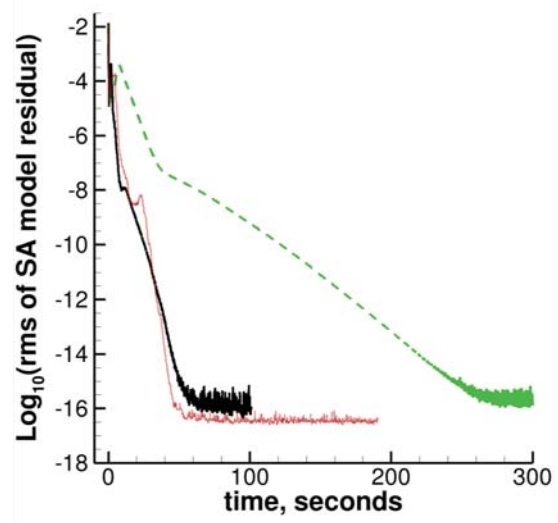
(b) rms of SA model residual
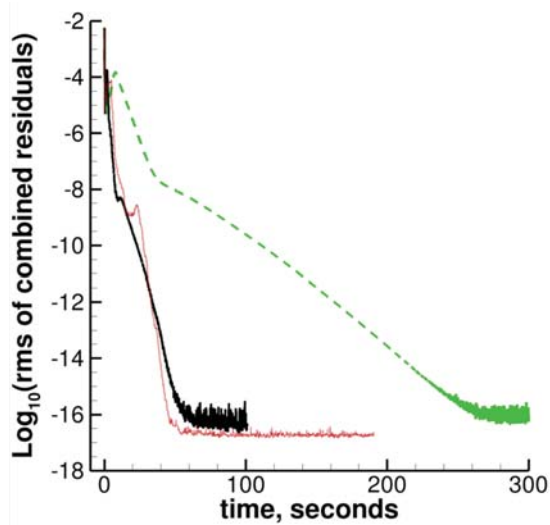
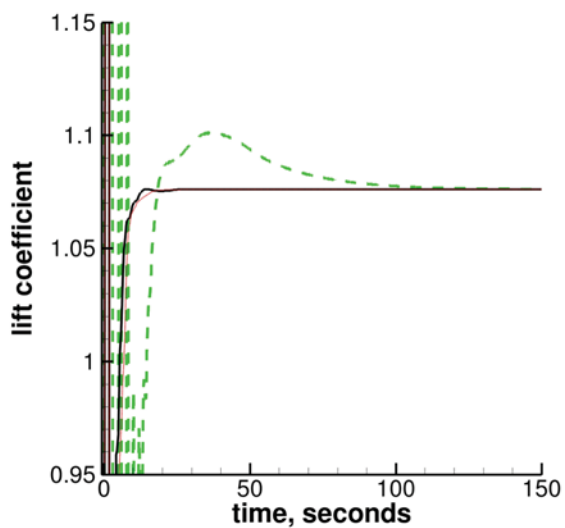(c) rms of combined mean flow and SA residuals
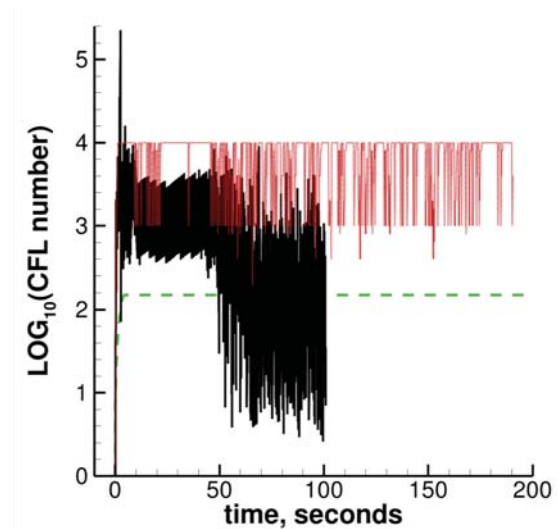
(d) CFL for mean flow and SA model

Figure 9. Flat plate solution convergence history using 2x69x49 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5x10^6$. Convective term in SA model is first-order accurate.

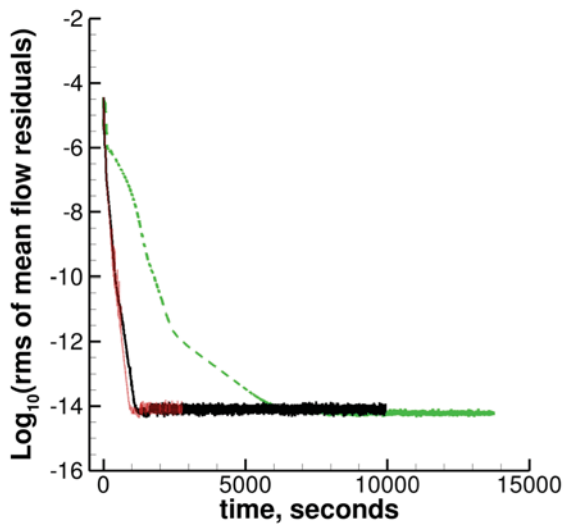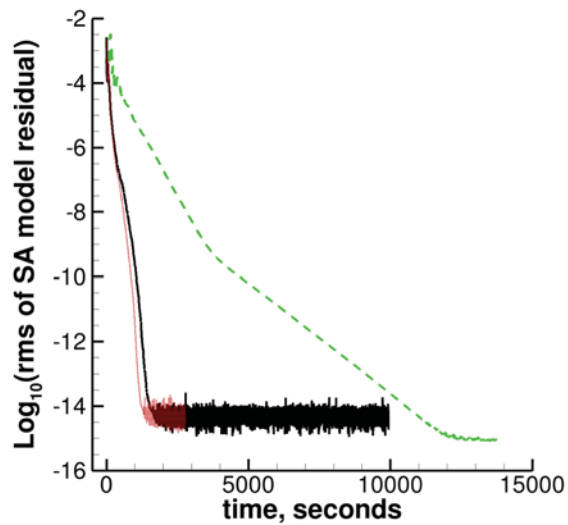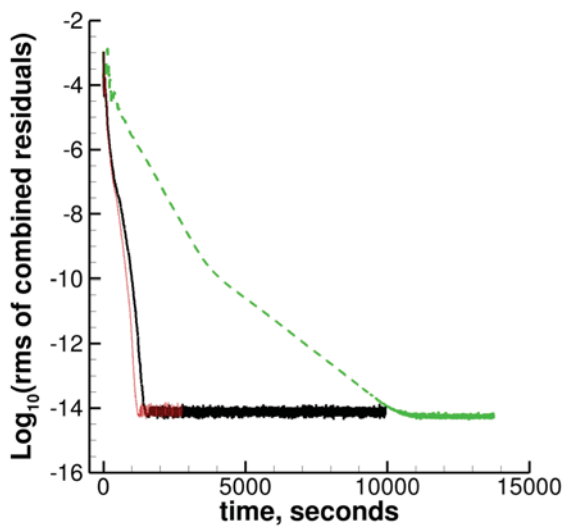(e) maximum normalized eddy viscosity   (f) viscous drag coefficient
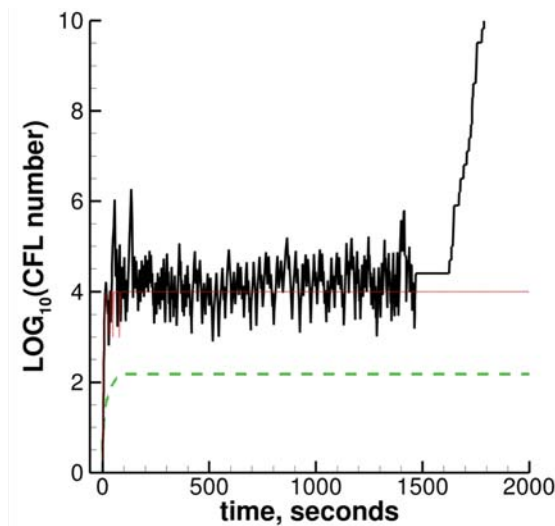
Figure 9. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 10. Flat plate solution convergence history using 2x137x97 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5x10^6$. Convective term in SA model is first-order accurate.

American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity   (f) viscous drag coefficient

Figure 10. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

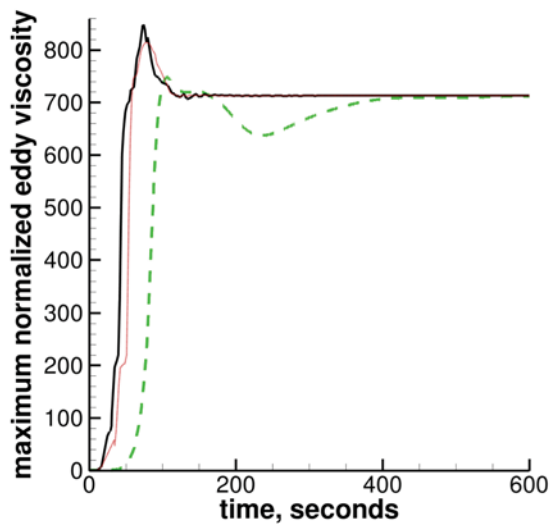
PA – – – – –

HANIS-1 ——————

HANIS-4 ——————

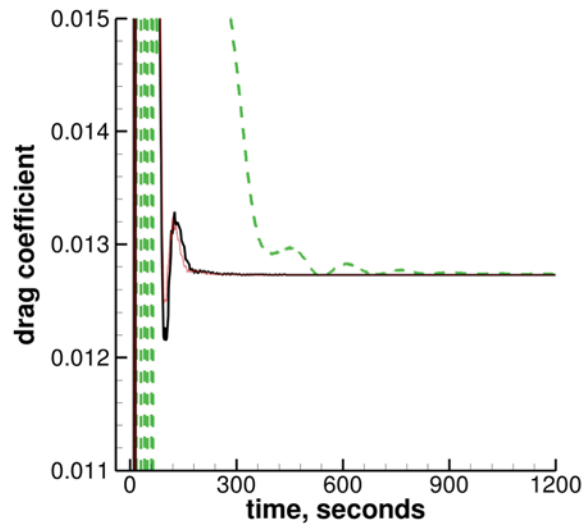(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 11. Flat plate solution convergence history using 2x273x193 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5x10^6$. Convective term in SA model is first-order accurate.

(e) maximum normalized eddy viscosity

(f) viscous drag coefficient

Figure 11. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 12. Flat plate solution convergence history using 2x545x385 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5 \times 10^6$. Convective term in SA model is first-order accurate.
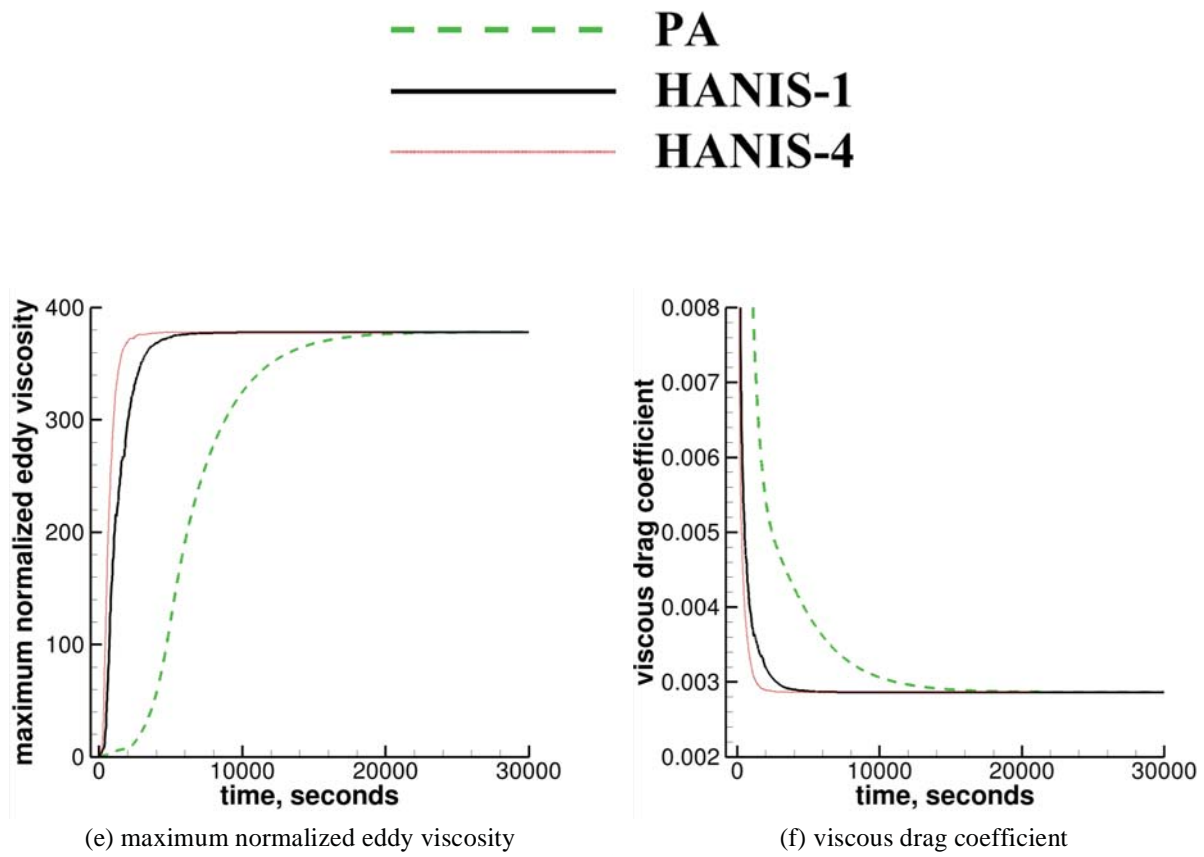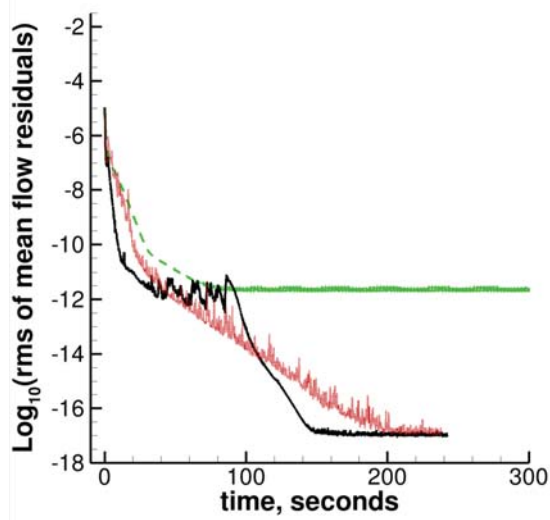
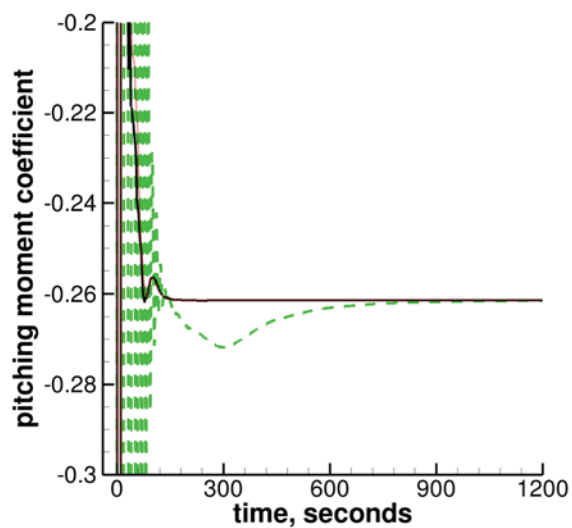(e) maximum normalized eddy viscosity   (f) viscous drag coefficient

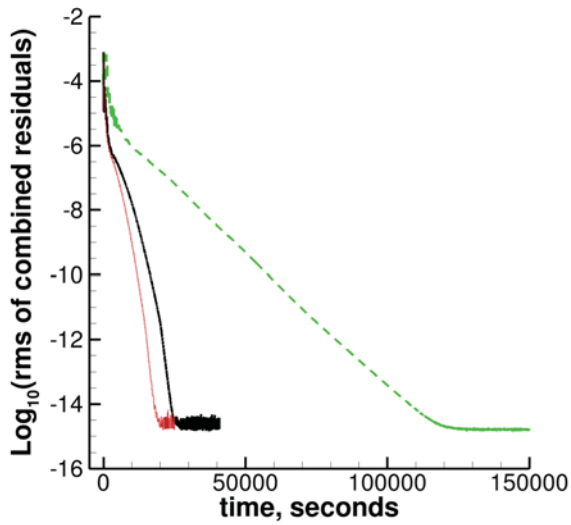**Figure 12. Concluded.**

American Institute of Aeronautics and Astronautics
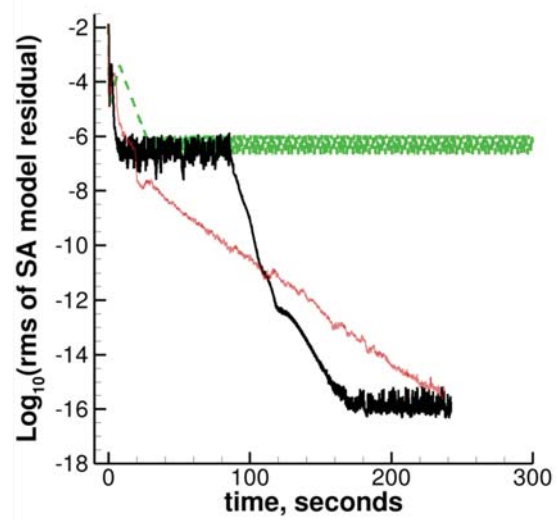
(a) rms of mean flow residuals

(b) rms of SA model residual
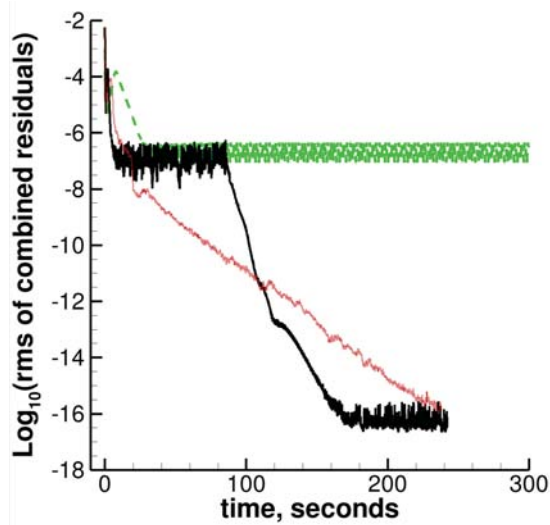
(c) rms of combined mean flow and SA residuals
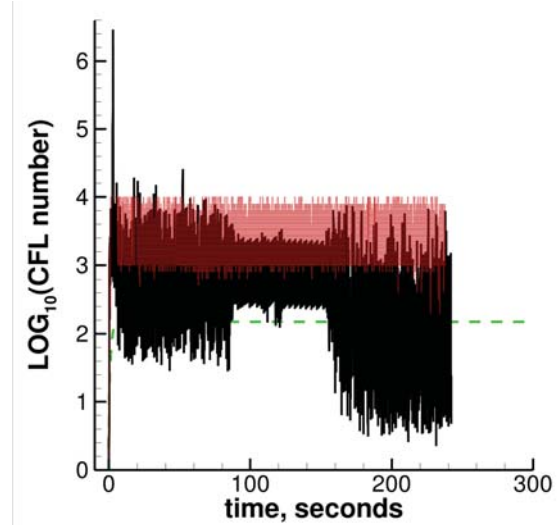
(d) CFL for mean flow and SA model

Figure 13. Flat plate solution convergence history using 2x69x49 hexahedral grid. M_∞ = 0.2, Re_L = 5x10^6. Convective term in SA model is second-order accurate.

(e) maximum normalized eddy viscosity   (f) viscous drag coefficient

Figure 13. Concluded.

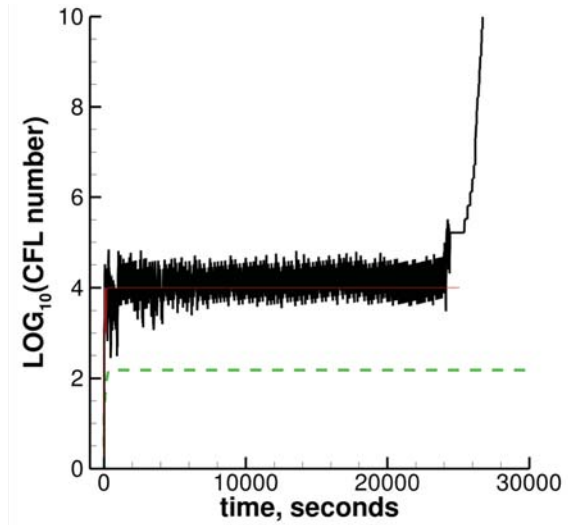American Institute of Aeronautics and Astronautics

(a) rms of mean flow residuals
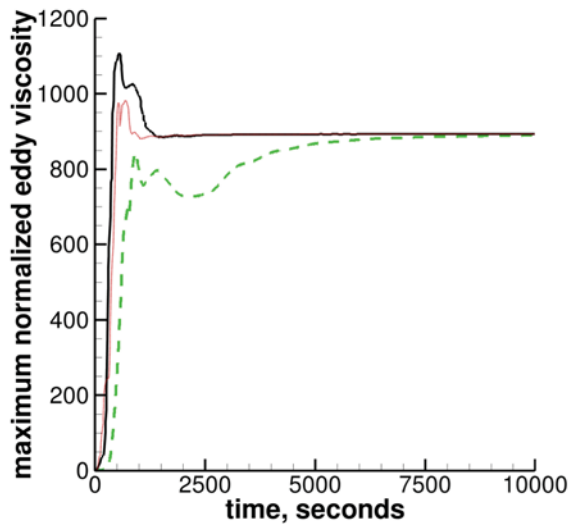
(b) rms of SA model residual

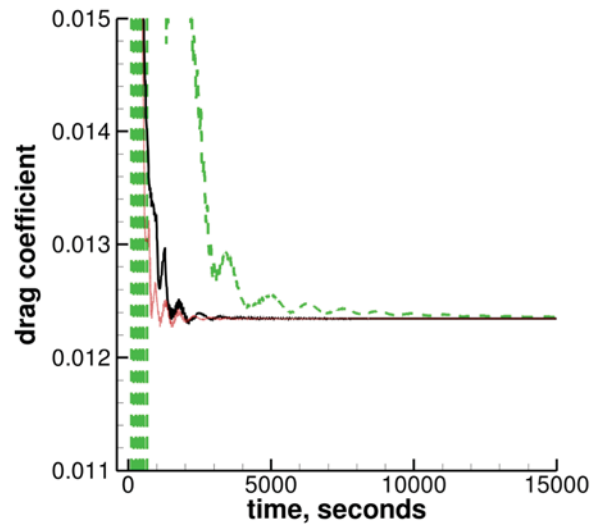(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 14. Flat plate solution convergence history using 2x137x97 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5x10^6$. Convective term in SA model is second-order accurate.

American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity   (f) viscous drag coefficient

Figure 14. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

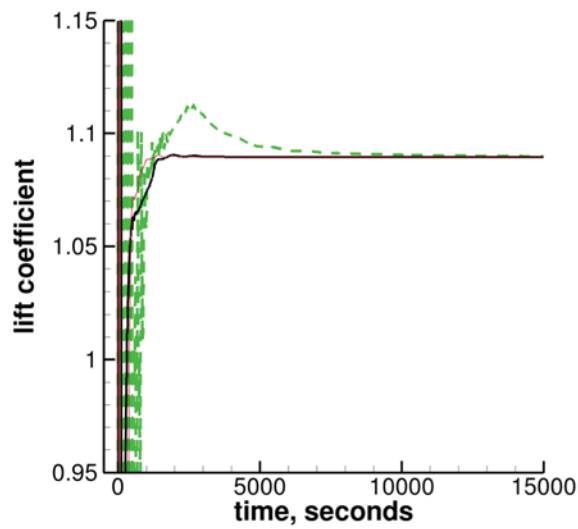(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 15. Flat plate solution convergence history using 2x273x193 hexahedral grid. M_∞ = 0.2, Re_L = 5x10^6. Convective term in SA model is second-order accurate.

38
American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity   (f) viscous drag coefficient
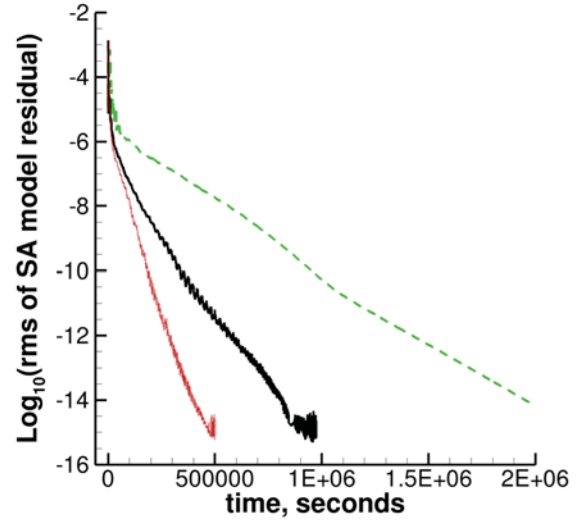
Figure 15. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

PA

HANIS-1

HANIS-4

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 16. Flat plate solution convergence history using 2x545x385 hexahedral grid. $M_\infty = 0.2$, $Re_L = 5x10^6$. Convective term in SA model is second-order accurate.

(e) maximum normalized eddy viscosity   (f) viscous drag coefficient

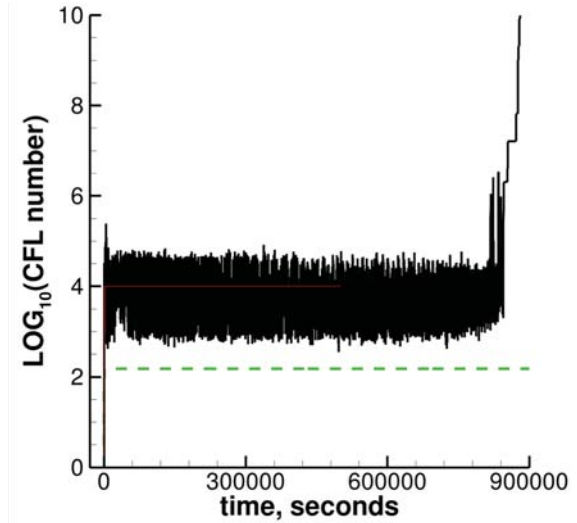**Figure 16. Concluded.**

(a) rms of mean flow residuals
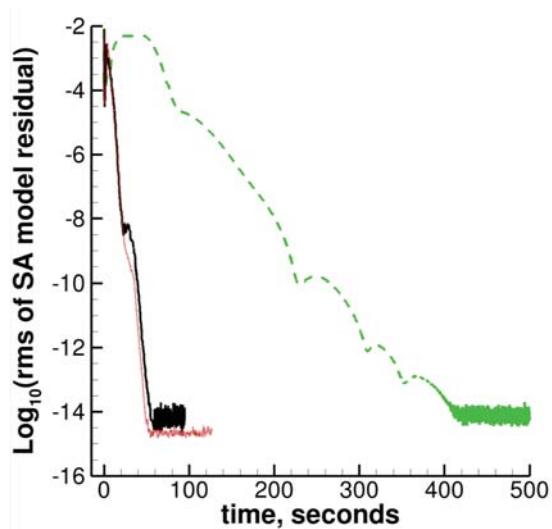
(b) rms of SA model residual
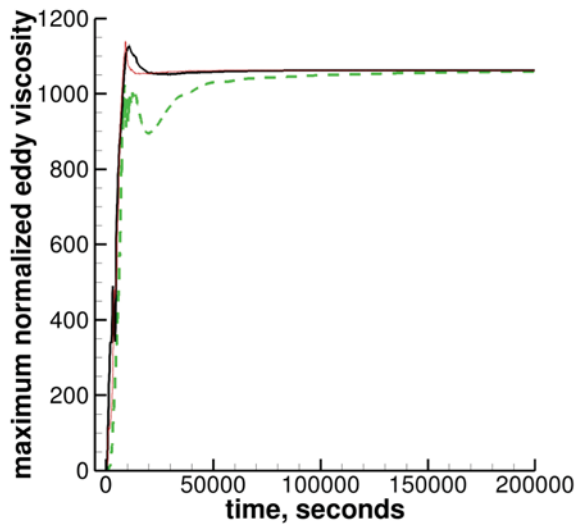

PA
HANIS-1
HANIS-4
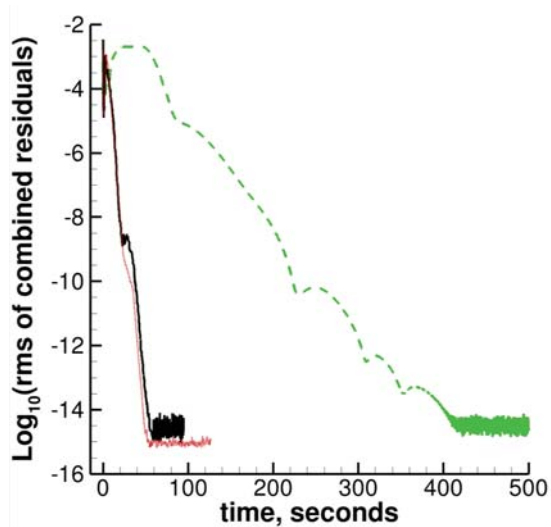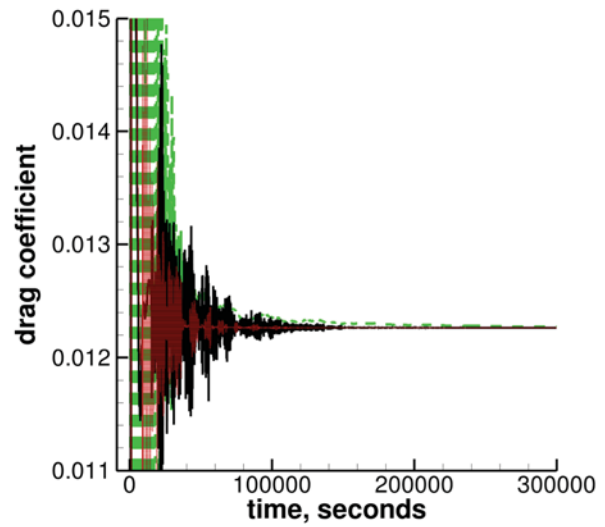

(c) rms of combined mean flow and SA residuals
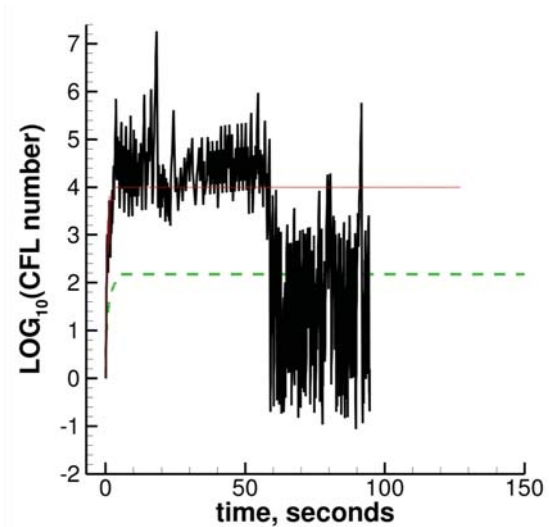
(d) CFL for mean flow and SA model

Figure 17. Bump-in-channel solution convergence history using 2x89x41 hexahedral grid. $M_\infty = 0.2$, $Re_L = 3x10^6$. Convective term in SA model is first-order accurate.

(e) maximum normalized eddy viscosity

(f) drag coefficient

Figure 17. Concluded.

PA
HANIS-1
HANIS-4

(g) lift coefficient

(h) viscous drag coefficient

Figure 17. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

PA

HANIS-1

HANIS-4

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model
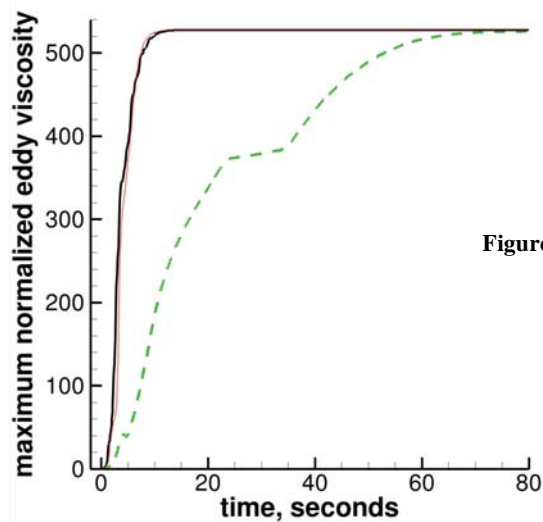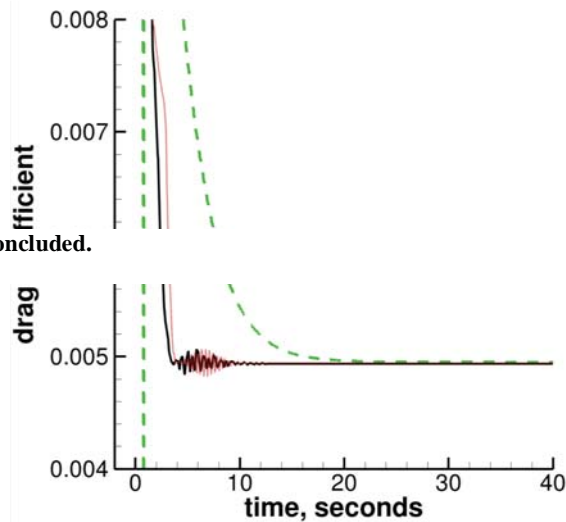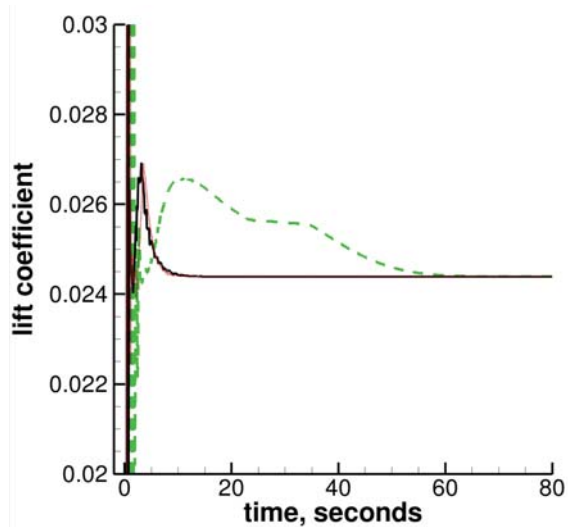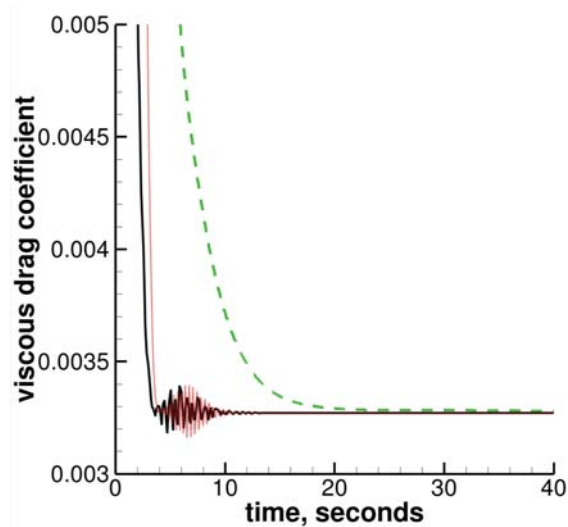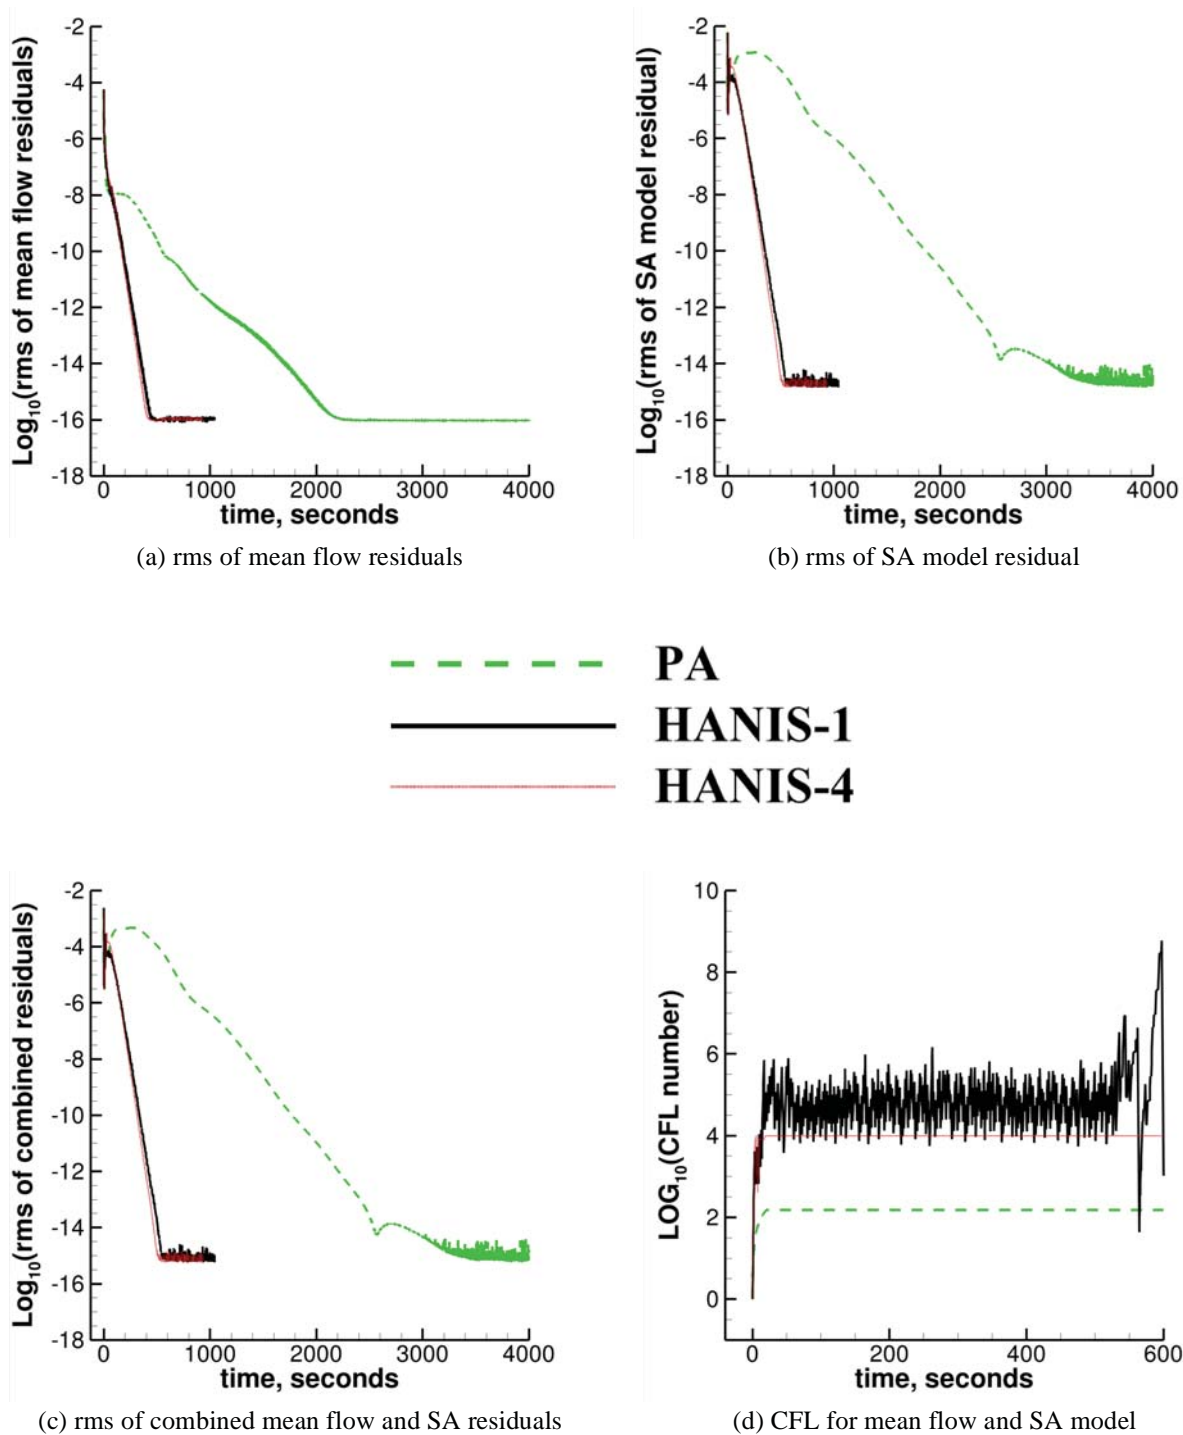
Figure 18. Bump-in-channel solution convergence history using 2x177x81 hexahedral grid. $M_\infty = 0.2$, $Re_L = 3x10^6$. Convective term in SA model is first-order accurate.
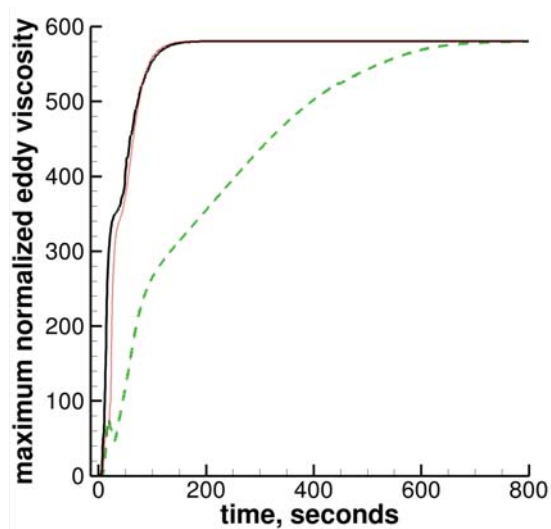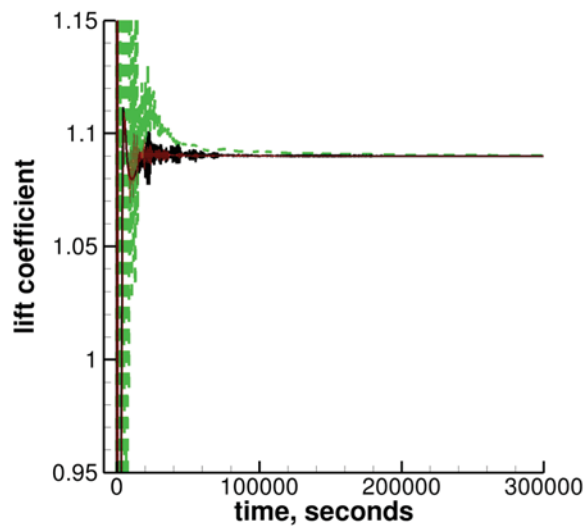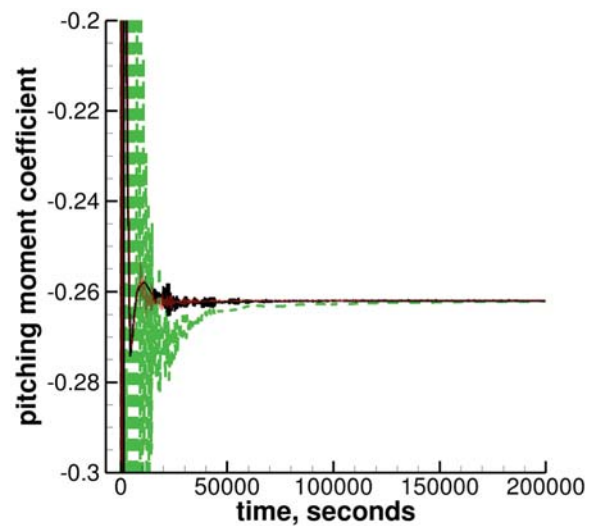
American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity

(f) drag coefficient

PA
HANIS-1
HANIS-4
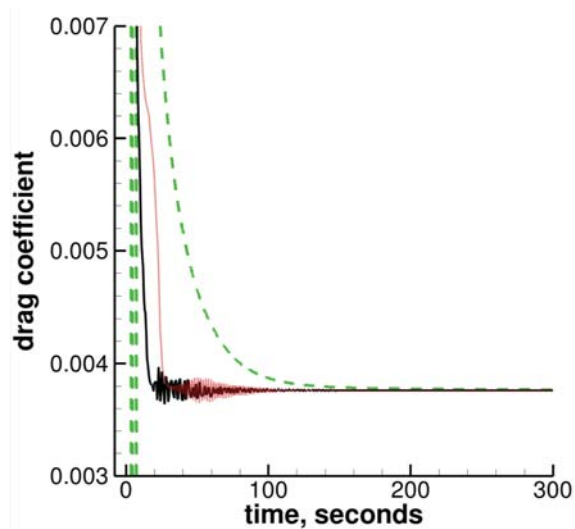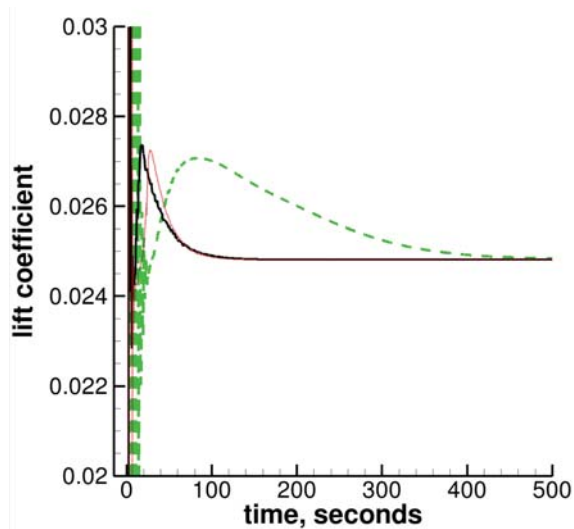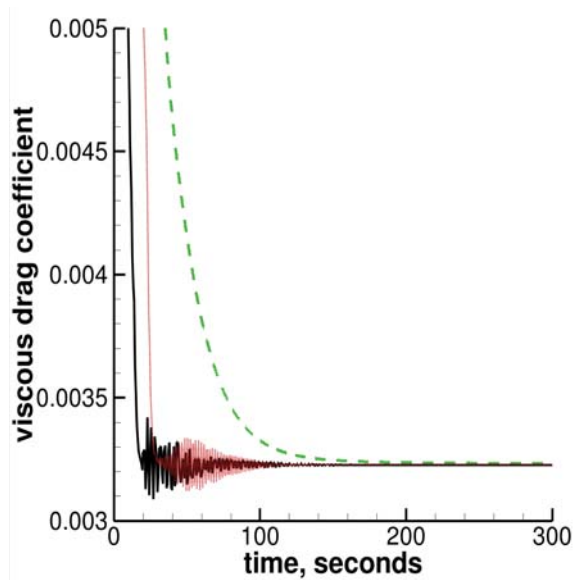
(g) lift coefficient

(h) viscous drag coefficient

Figure 18. Concluded.
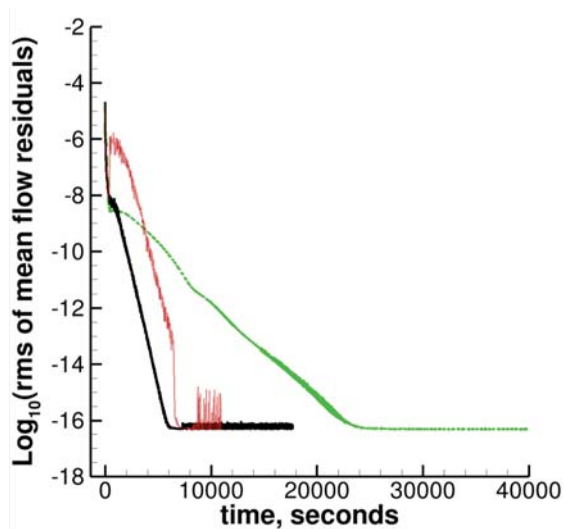
(a) rms of mean flow residuals
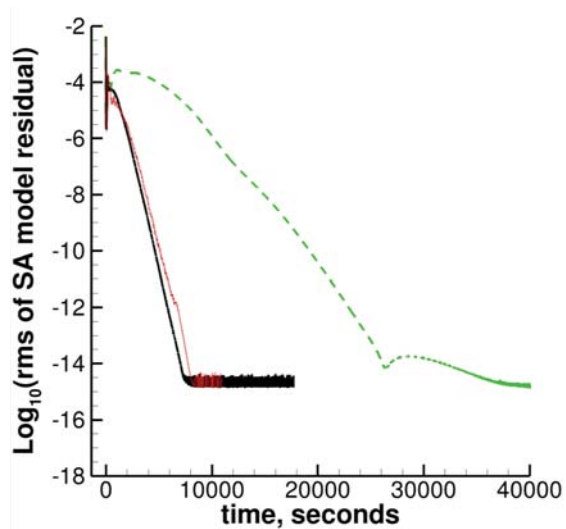
(b) rms of SA model residual

PA
HANIS-1
HANIS-4
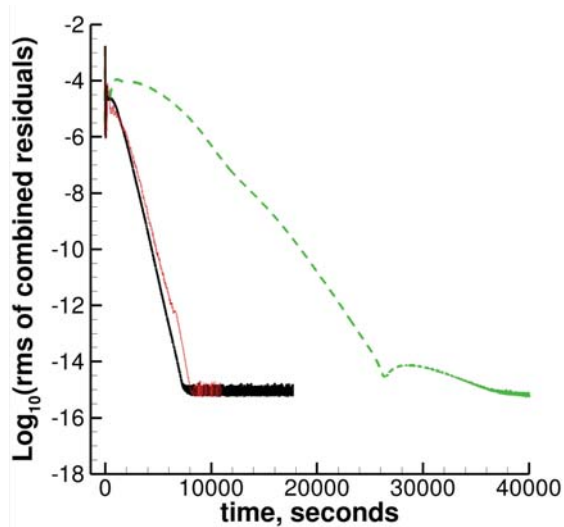
(c) rms of combined mean flow and SA residuals
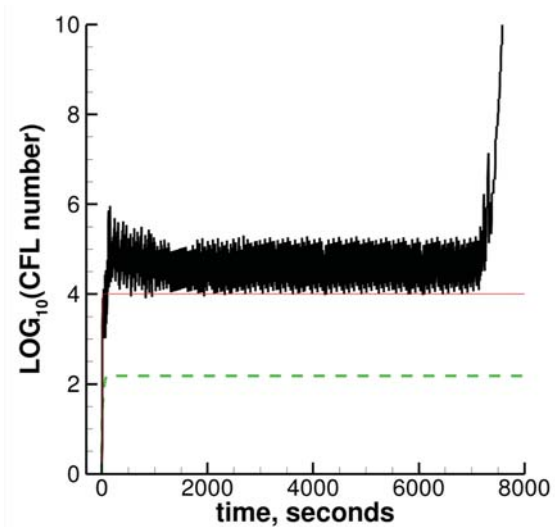
(d) CFL for mean flow and SA model

Figure 19. Bump-in-channel solution convergence history using 2x353x161 hexahedral grid. $M_\infty = 0.2$, $Re_L = 3x10^6$. Convective term in SA model is first-order accurate.
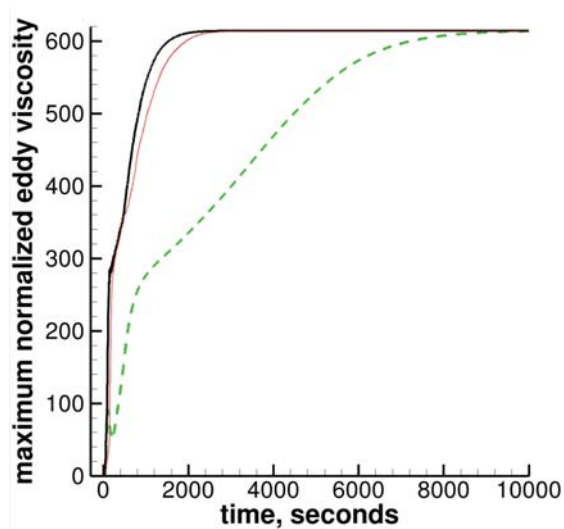
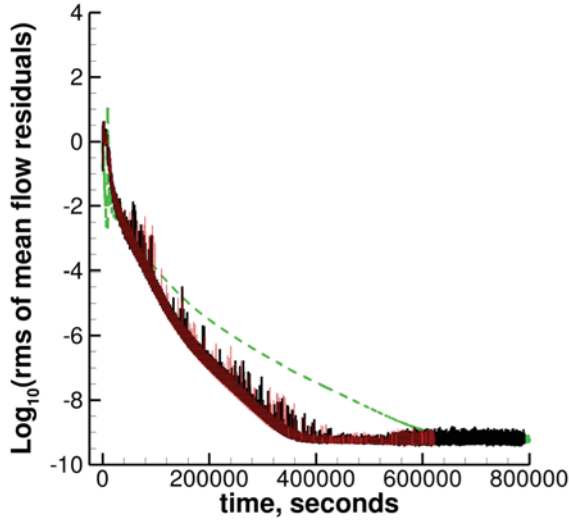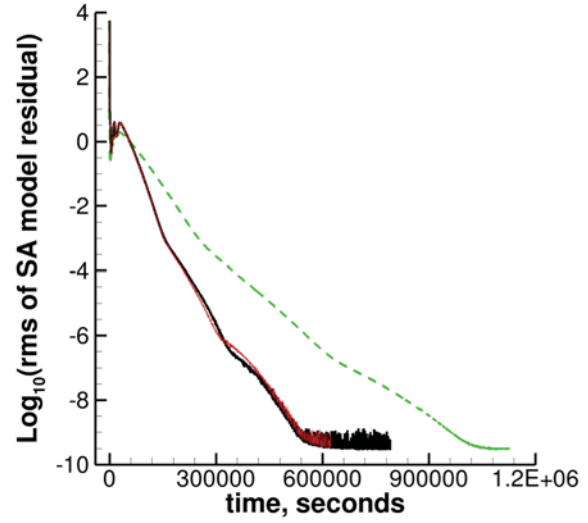(e) maximum normalized eddy viscosity
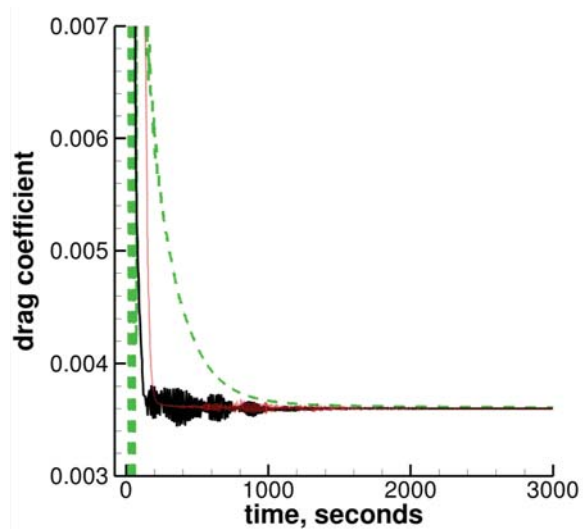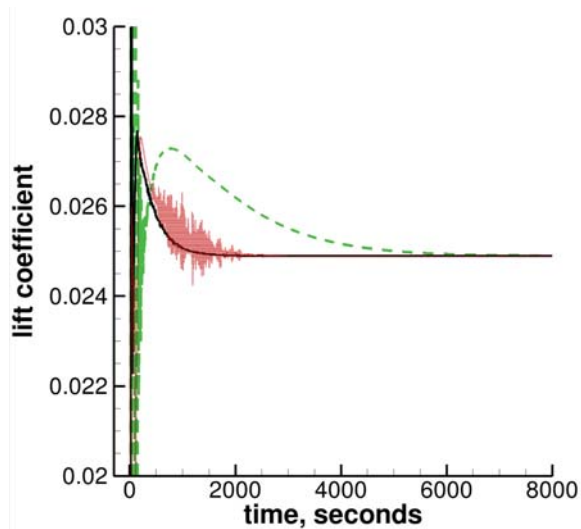
(f) drag coefficient
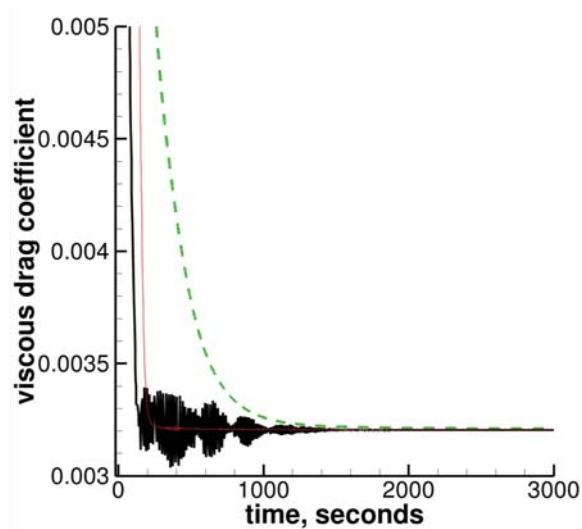
(g) lift coefficient

(h) viscous drag coefficient
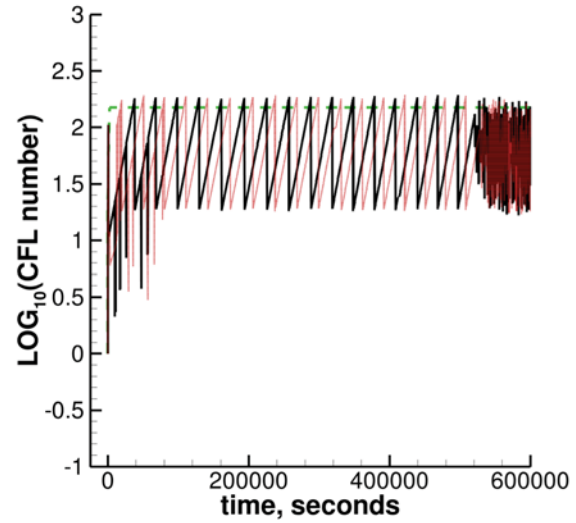
Figure 19. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

PA
HANIS-1
HANIS-4
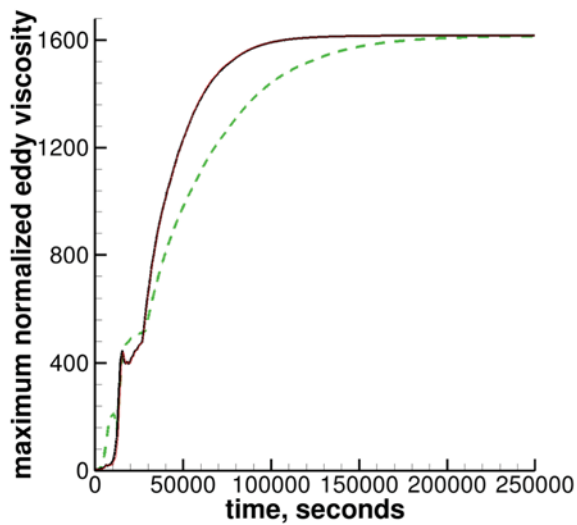
(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model
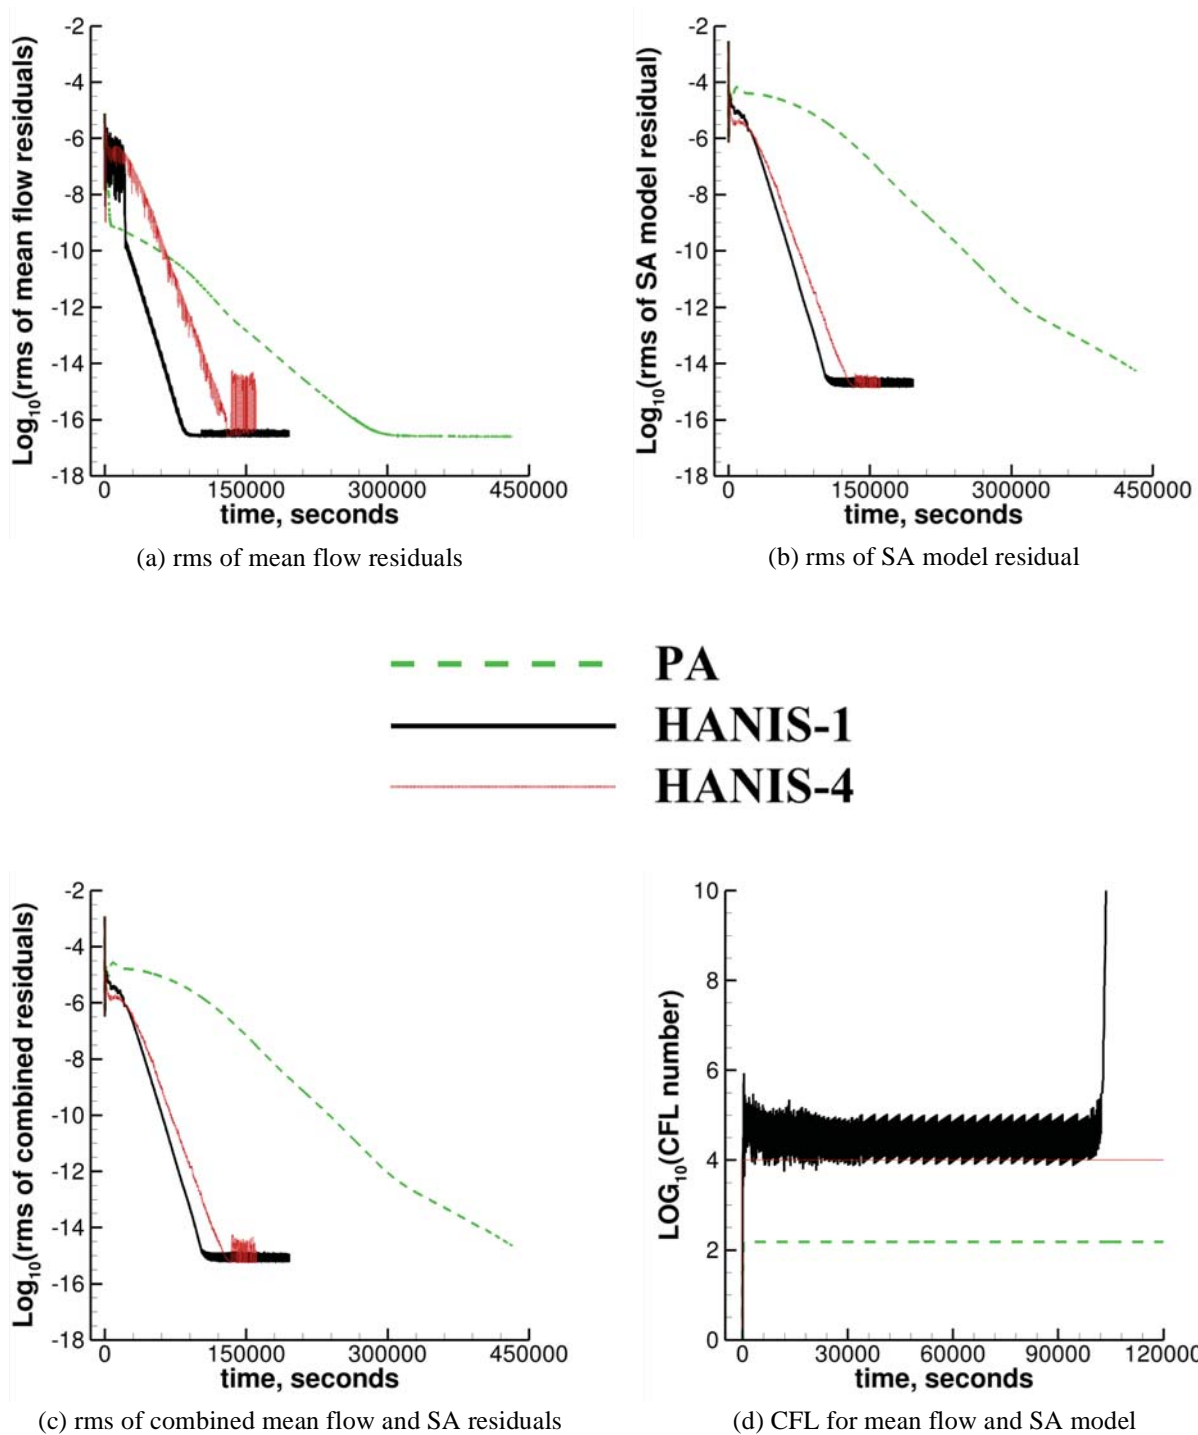
Figure 20. Bump-in-channel solution convergence history using 2x705x321 hexahedral grid. $M_\infty = 0.2$, $Re_L = 3x10^6$. Convective term in SA model is first-order accurate.
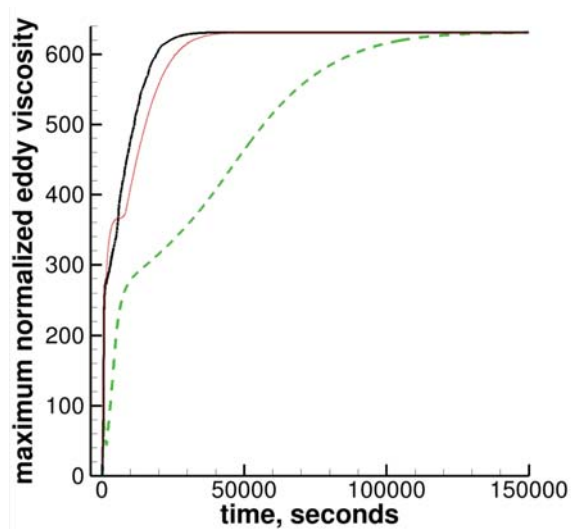
(e) maximum normalized eddy viscosity
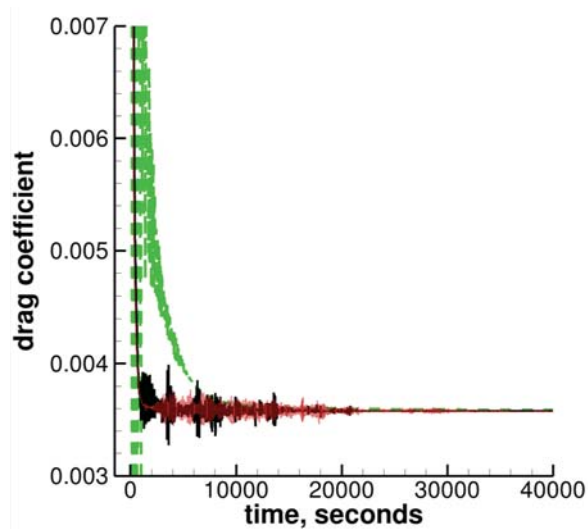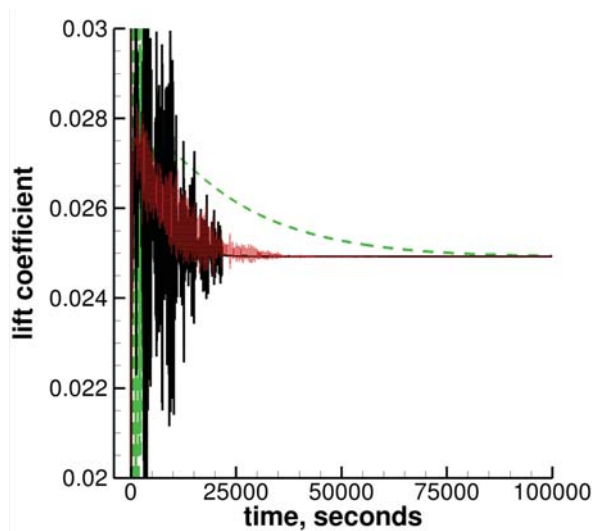
(f) drag coefficient
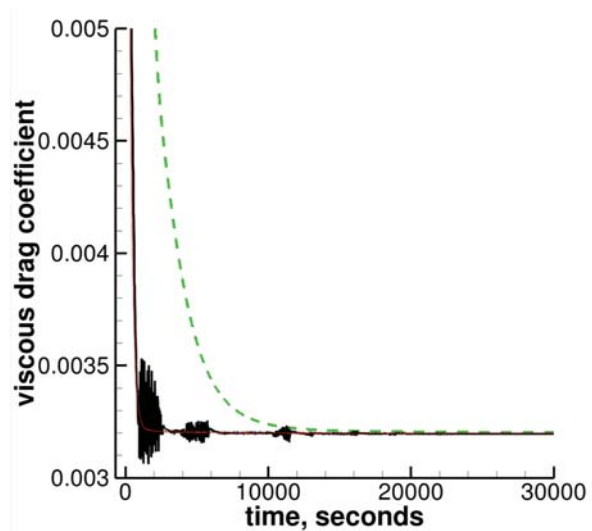
(g) lift coefficient

(h) viscous drag coefficient

Figure 20. Concluded.
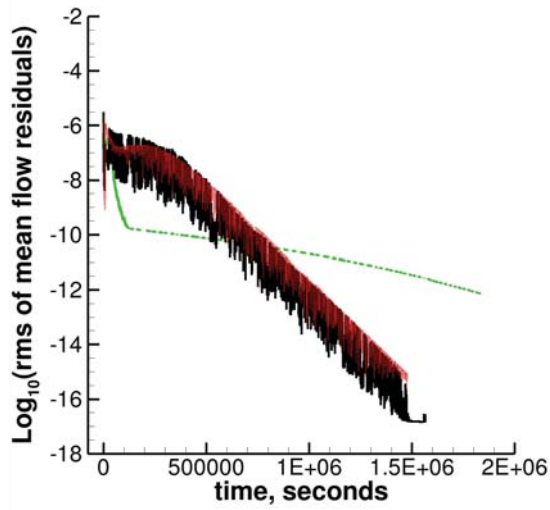
(a) rms of mean flow residuals
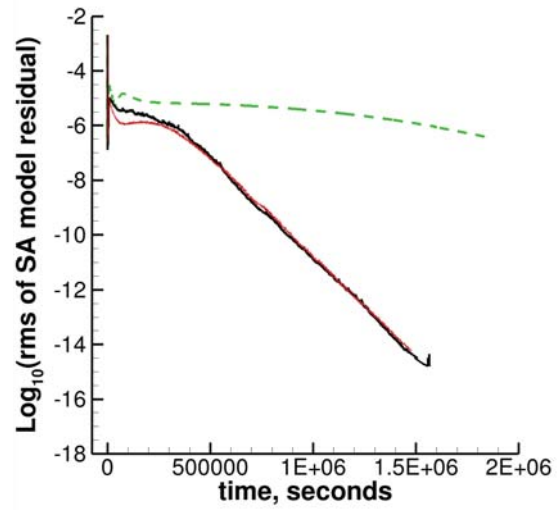
(b) rms of SA model residual

PA
HANIS-1
HANIS-4
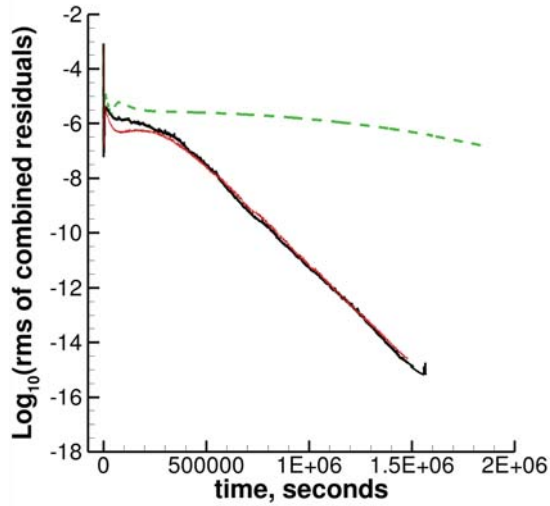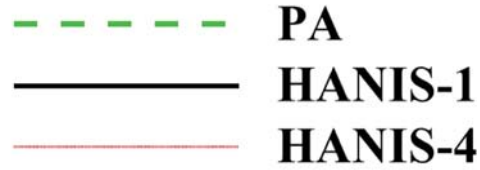
(c) rms of combined mean flow and SA residuals
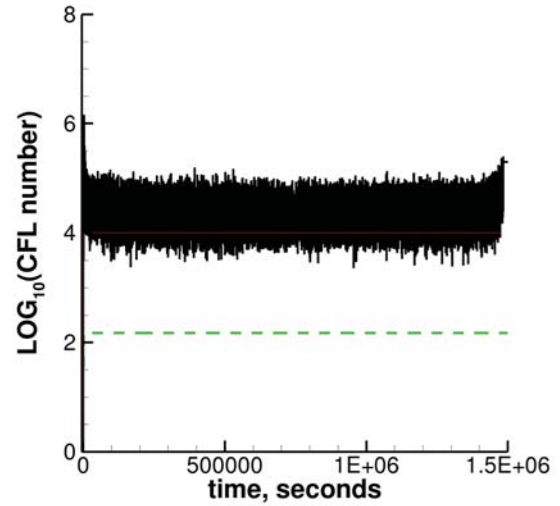
(d) CFL for mean flow and SA model

Figure 21. Bump-in-channel solution convergence history using 2x1409x641 hexahedral grid. $M_\infty = 0.2$, $Re_L = 3x10^6$. Convective term in SA model is first-order accurate.
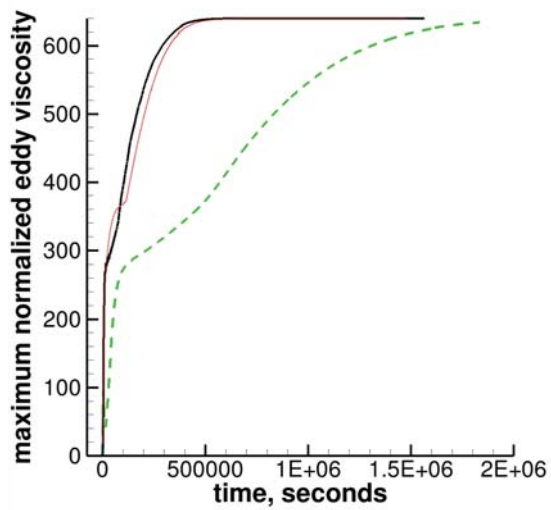
(e) maximum normalized eddy viscosity
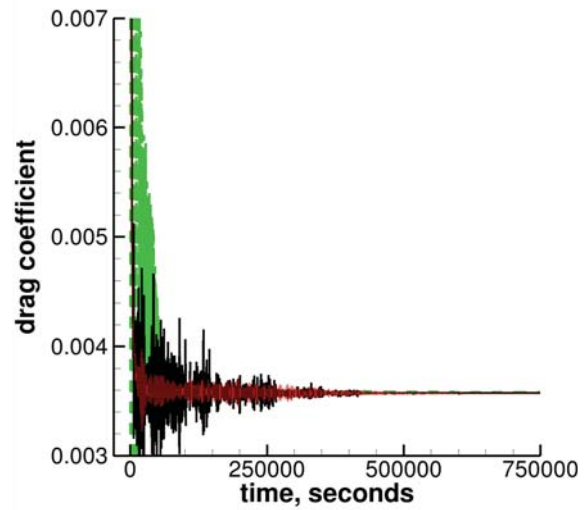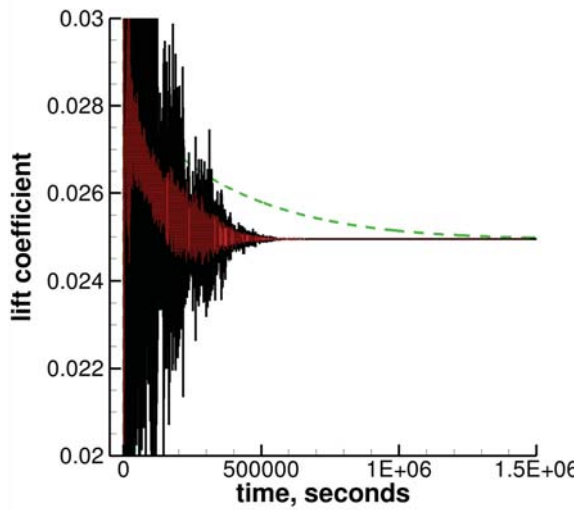
(f) drag coefficient
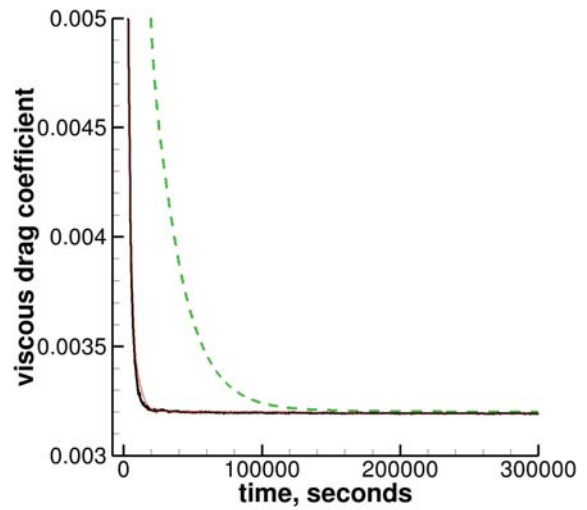
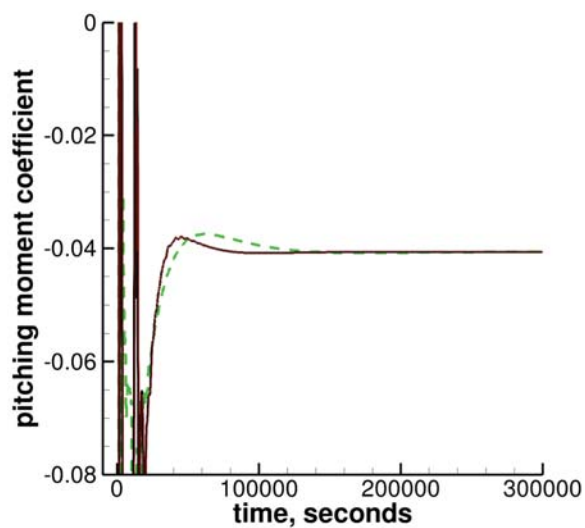

(g) lift coefficient

(h) viscous drag coefficient

Figure 21. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

PA

HANIS-1

HANIS-4

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model
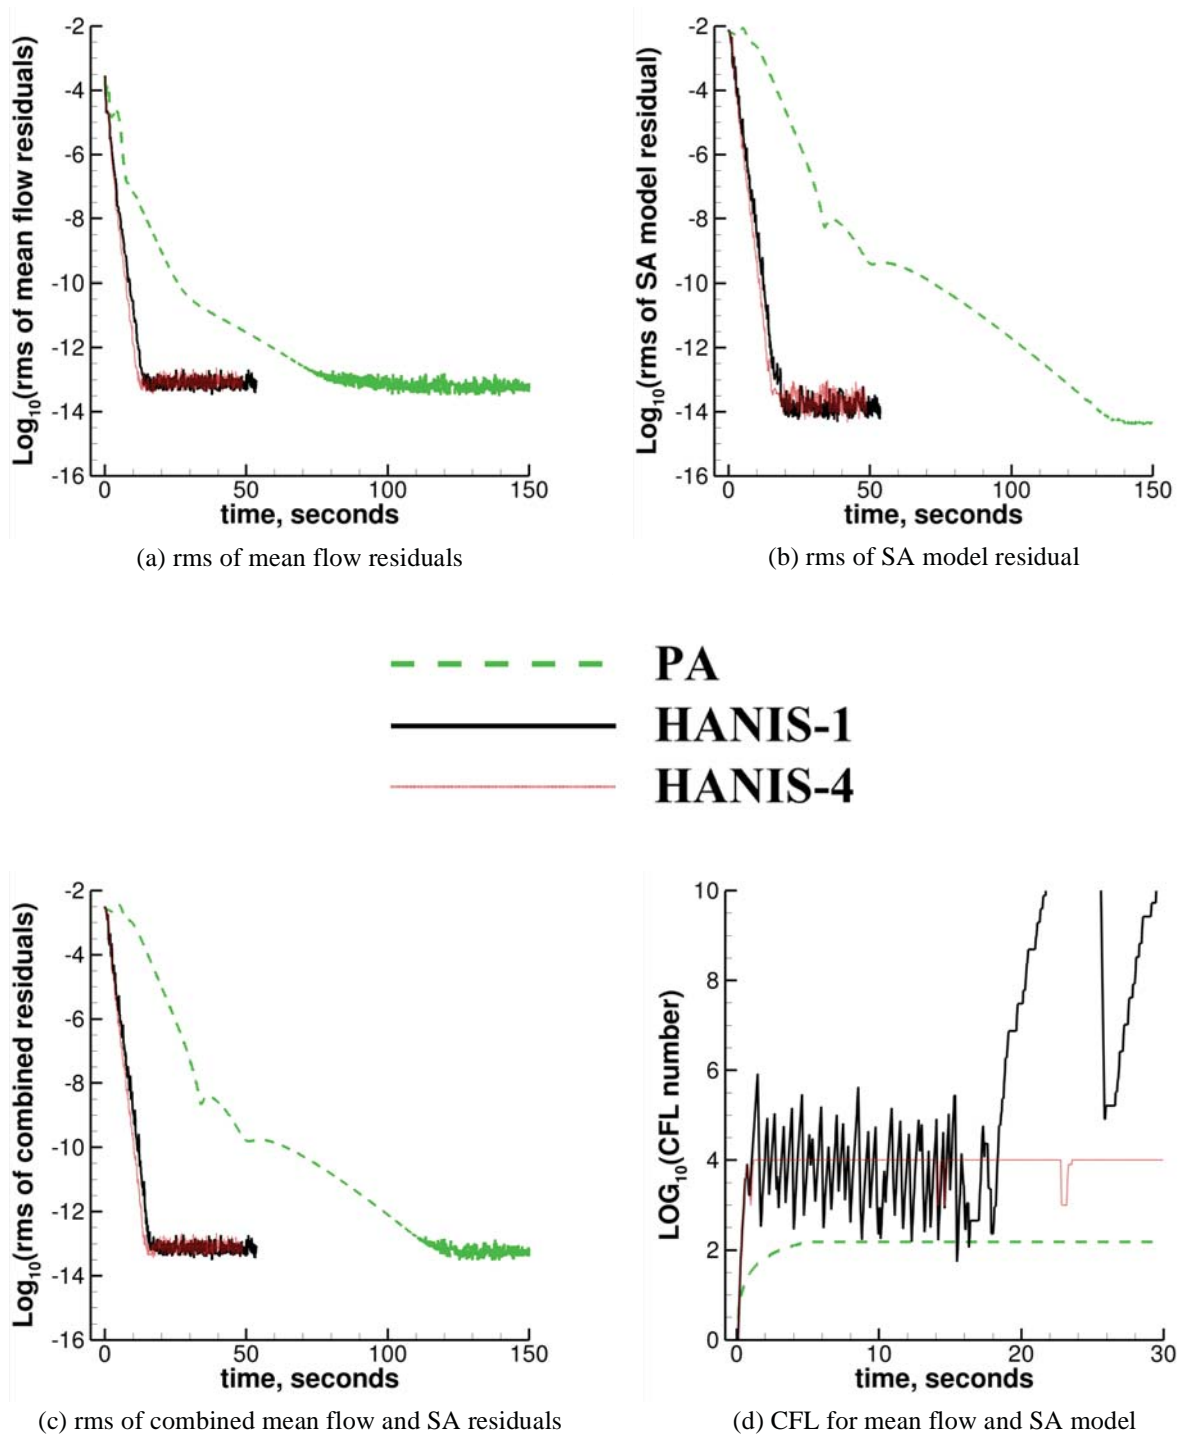
Figure 22. NACA 0012 solution convergence history using 2x113x33 hexahedral grid. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6 \times 10^6$. Convective term in SA model is first-order accurate.
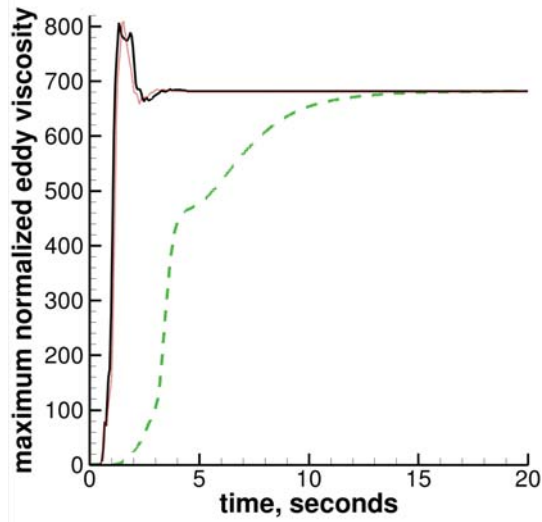
(e) maximum normalized eddy viscosity
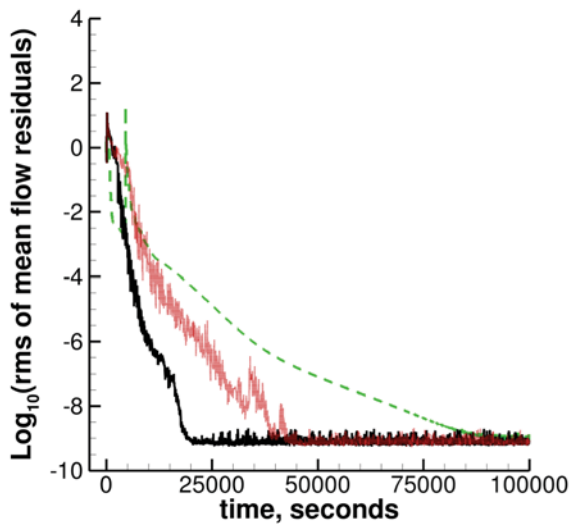
(f) drag coefficient

PA
HANIS-1
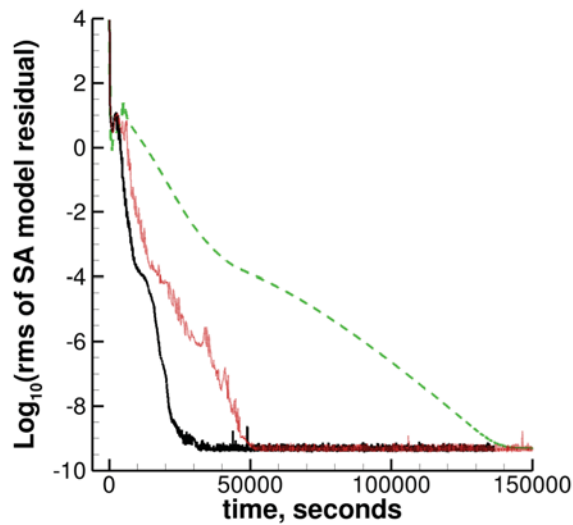HANIS-4
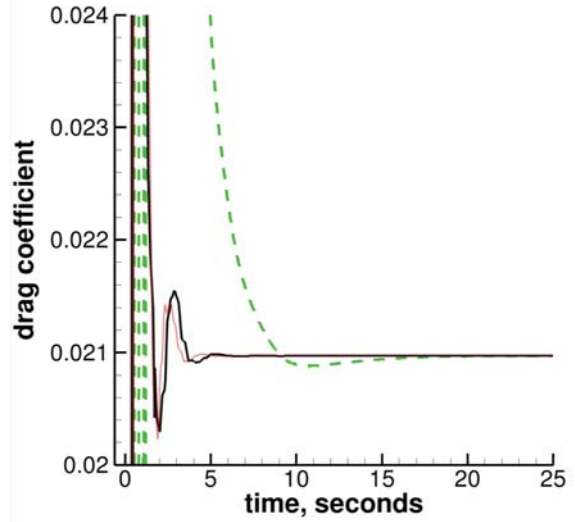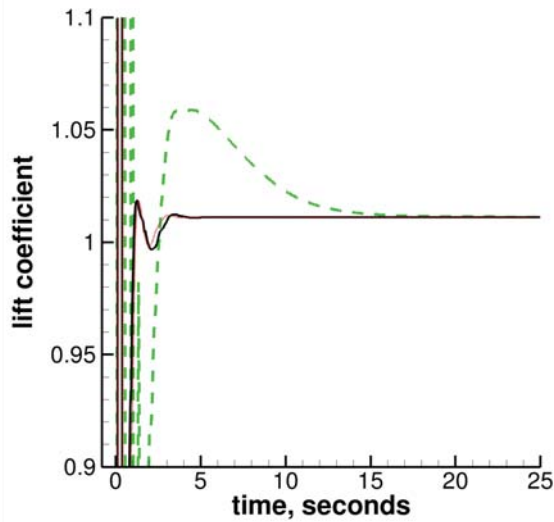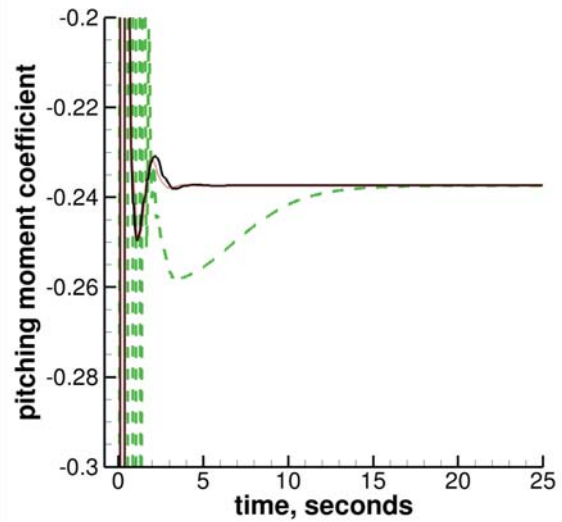
(g) lift coefficient

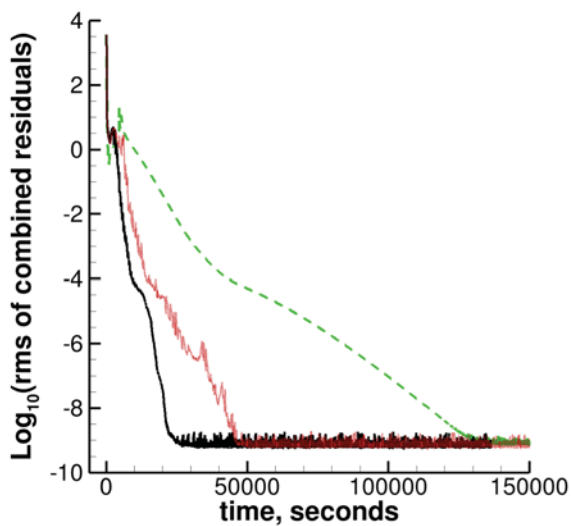(h) pitching moment coefficient
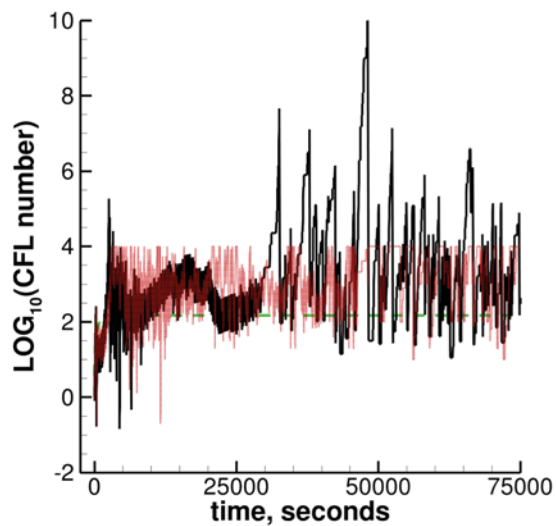
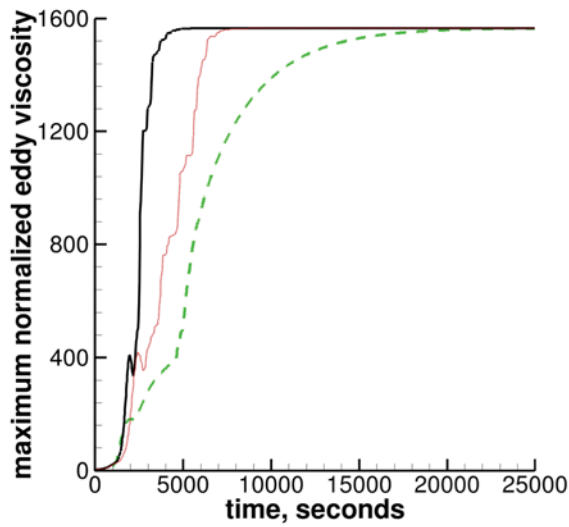**Figure 22. Concluded.**
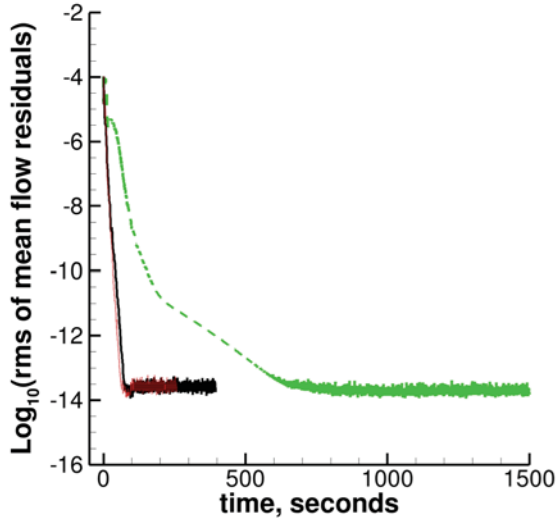
(a) rms of mean flow residuals
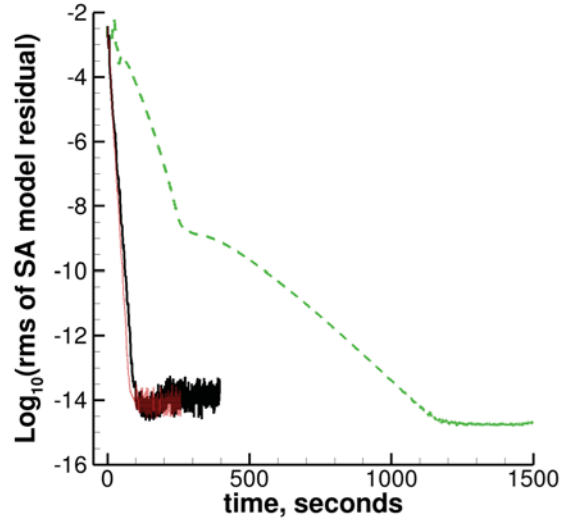
(b) rms of SA model residual
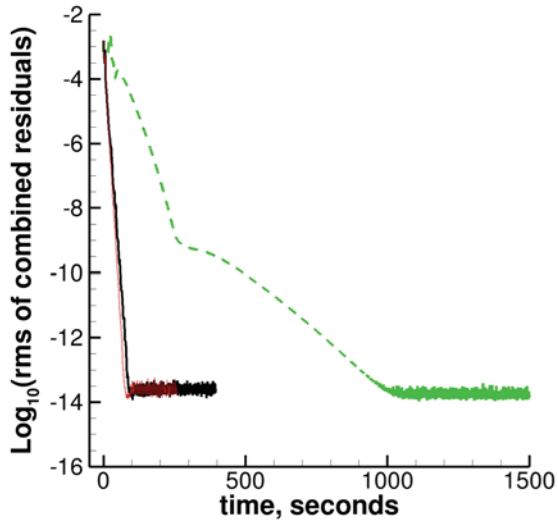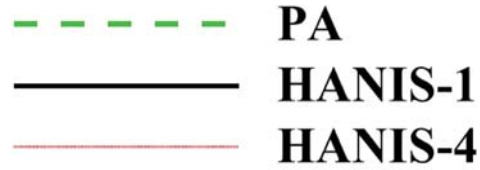
(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 23. NACA 0012 solution convergence history using 2x225x65 hexahedral grid. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6x10^6$. Convective term in SA model is first-order accurate.

(e) maximum normalized eddy viscosity

(f) drag coefficient

- - - - PA
——— HANIS-1
——— HANIS-4
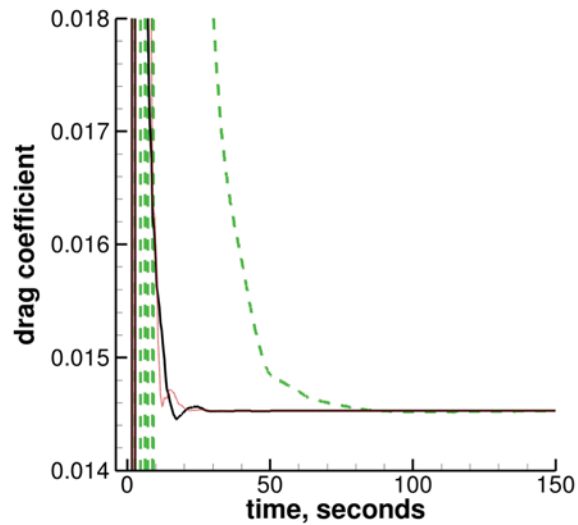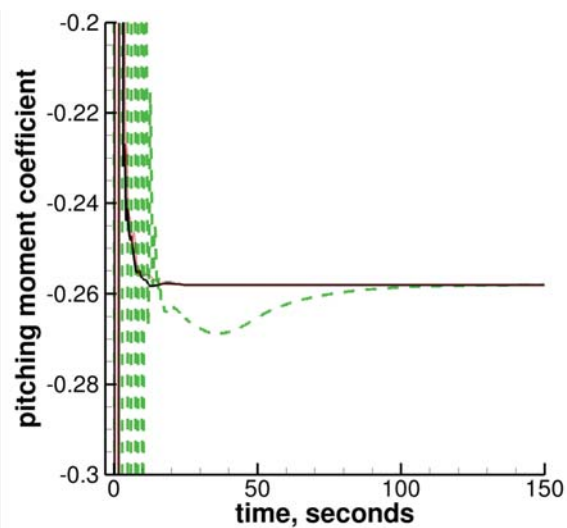
(g) lift coefficient

(h) pitching moment coefficient

Figure 23. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual
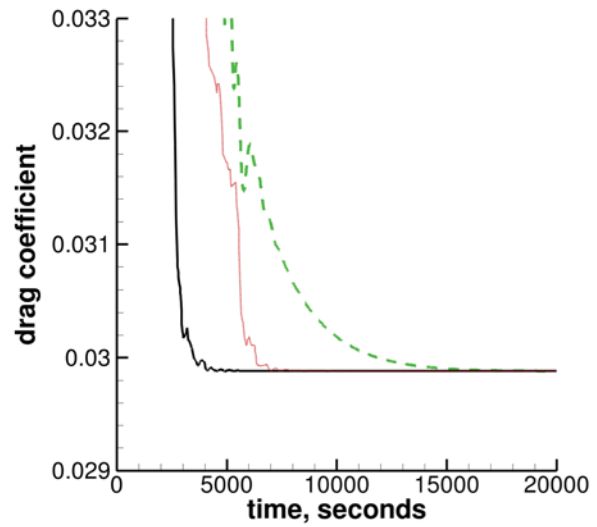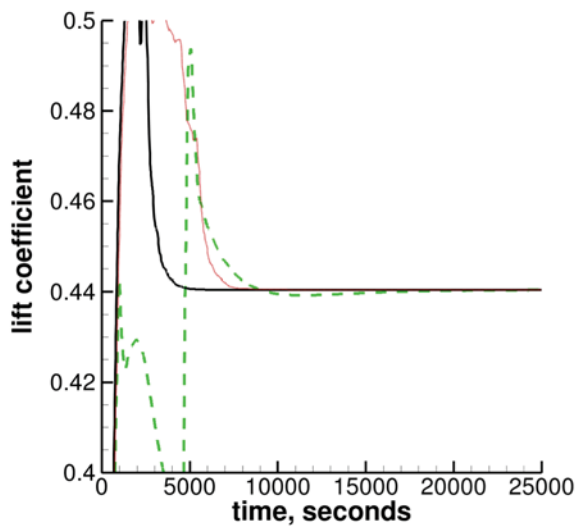
(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 24. NACA 0012 solution convergence history using 2x449x129 hexahedral grid. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6\text{x}10^6$. Convective term in SA model is first-order accurate.

American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity

(f) drag coefficient

(g) lift coefficient

(h) pitching moment coefficient

Figure 24. Concluded.
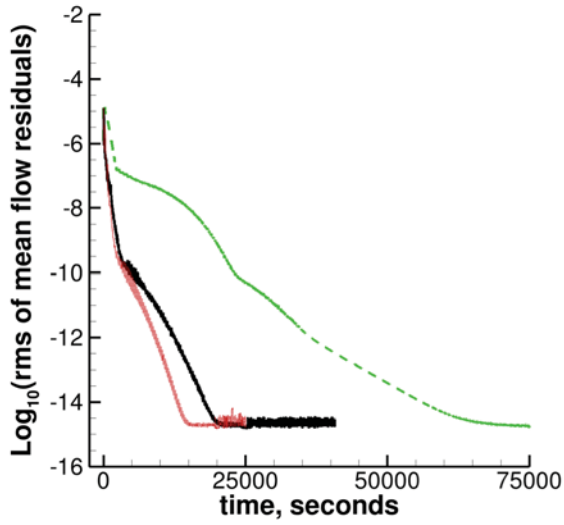
(a) rms of mean flow residuals
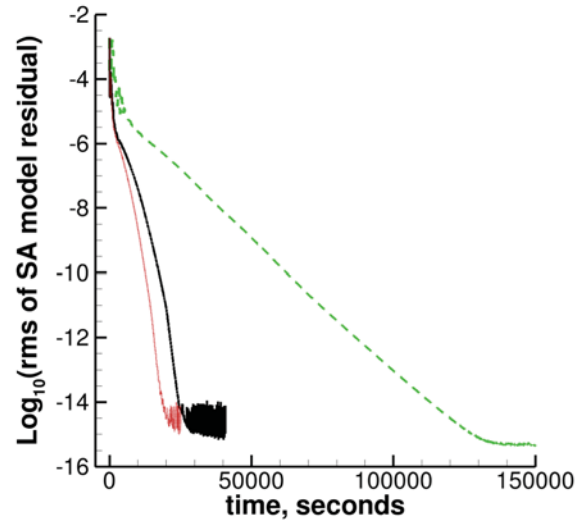
(b) rms of SA model residual

PA
HANIS-1
HANIS-4

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 25. NACA 0012 solution convergence history using 2x897x257 hexahedral grid. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6x10^6$. Convective term in SA model is first-order accurate.
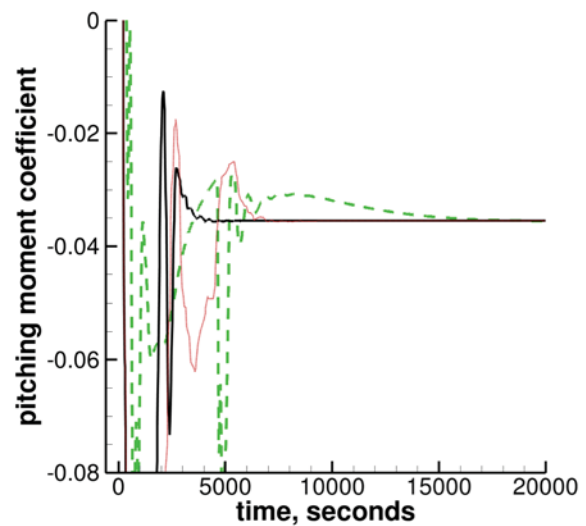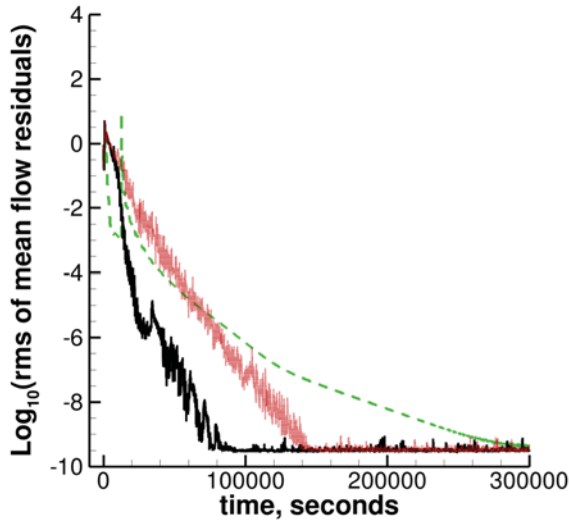
(e) maximum normalized eddy viscosity

(f) drag coefficient

PA
HANIS-1
HANIS-4

(g) lift coefficient
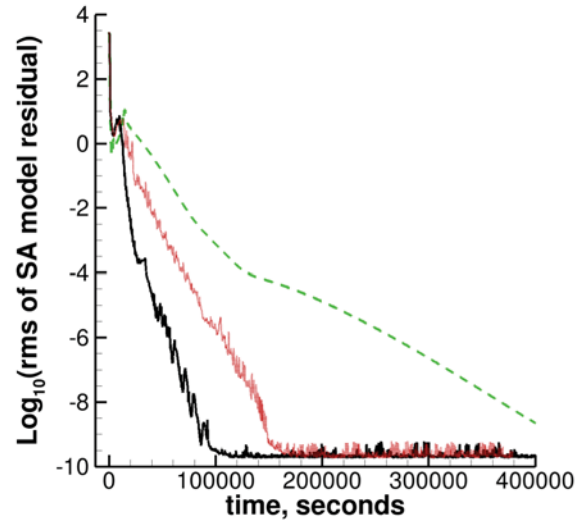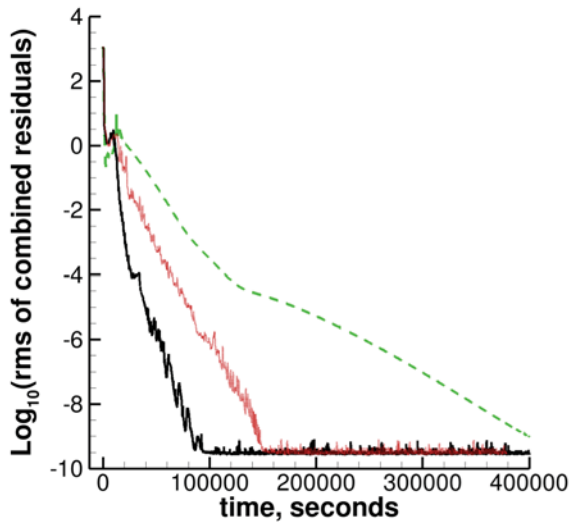
(h) pitching moment coefficient

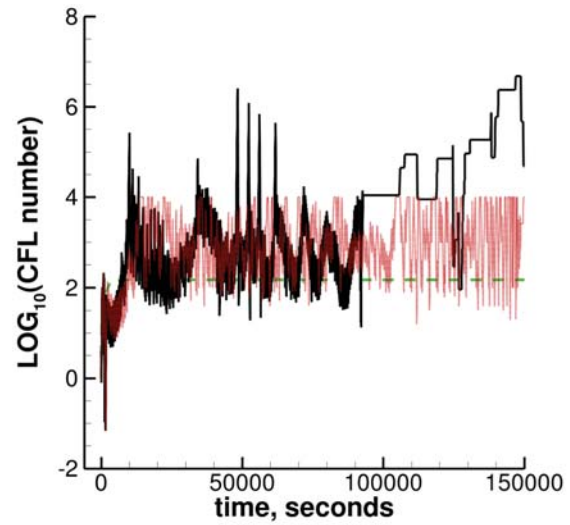**Figure 25. Concluded.**
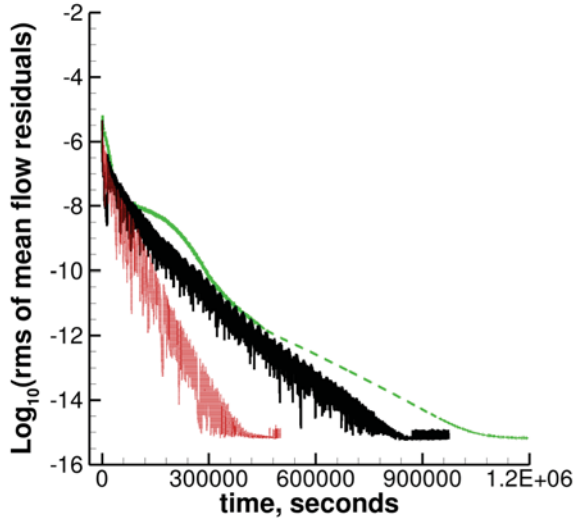
(a) rms of mean flow residuals

(b) rms of SA model residual

PA
HANIS-1
HANIS-4
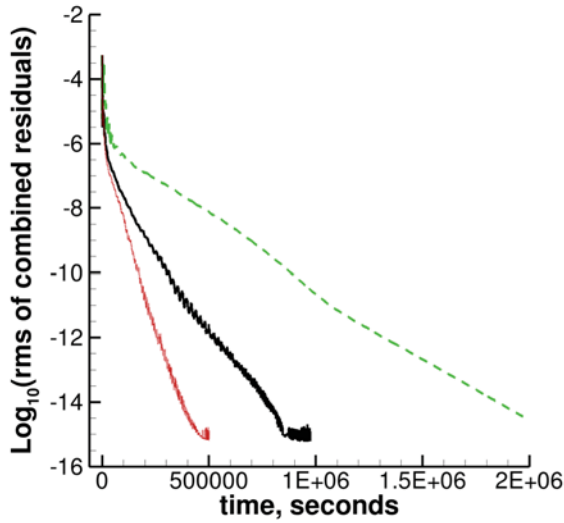
(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 26. NACA 0012 solution convergence history using 2x1793x513 hexahedral grid. $M_\infty = 0.15$, $\alpha = 10°$, $Re_c = 6x10^6$. Convective term in SA model is first-order accurate.

American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity

(f) drag coefficient

(g) lift coefficient

(h) pitching moment coefficient

Figure 26. Concluded.

(a) rms of mean flow residuals
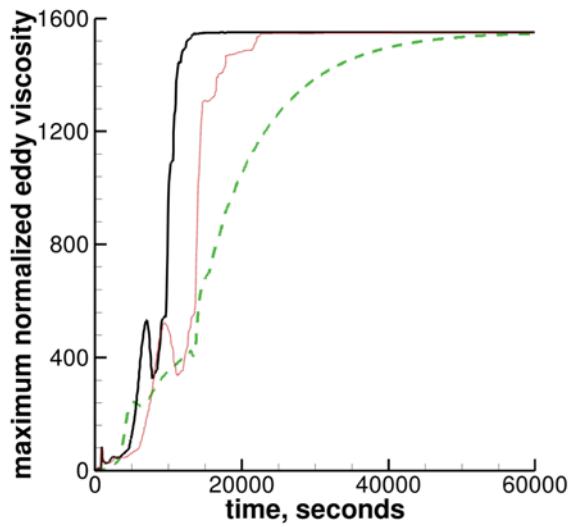
(b) rms of SA model residual
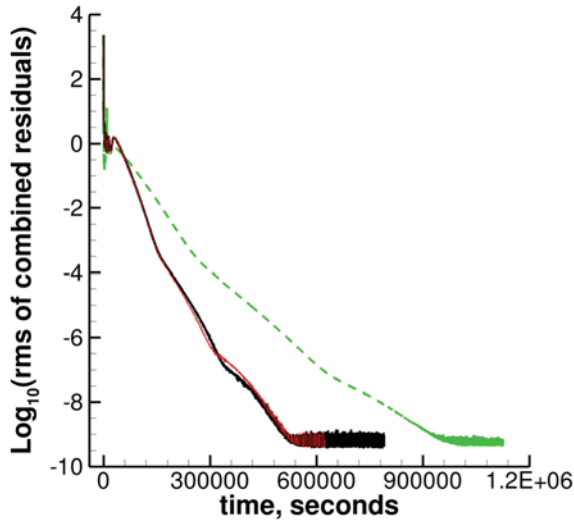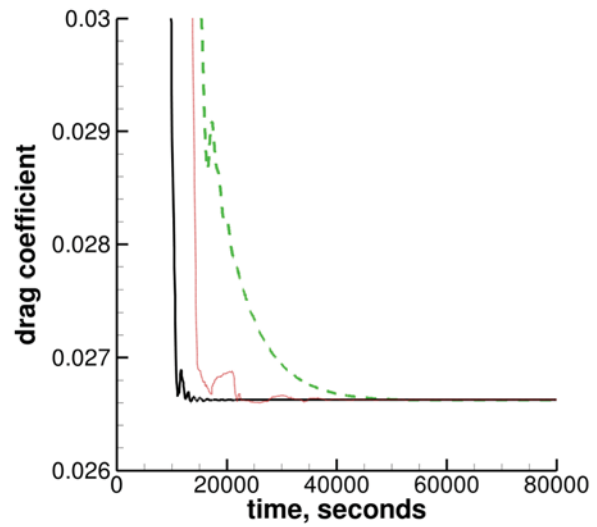
(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 27. NASA Common Research Model (CRM) solution convergence history using a coarse tetrahedral grid. $M_\infty = 0.85$, $\alpha = 2°$, $Re_c = 5x10^6$. Convective term in SA model is first-order accurate.
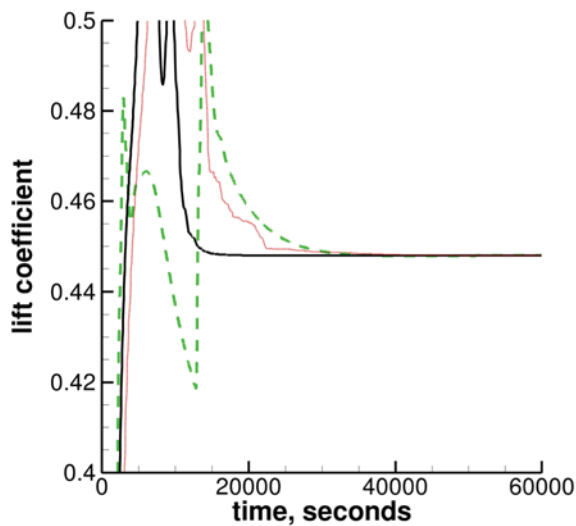
American Institute of Aeronautics and Astronautics
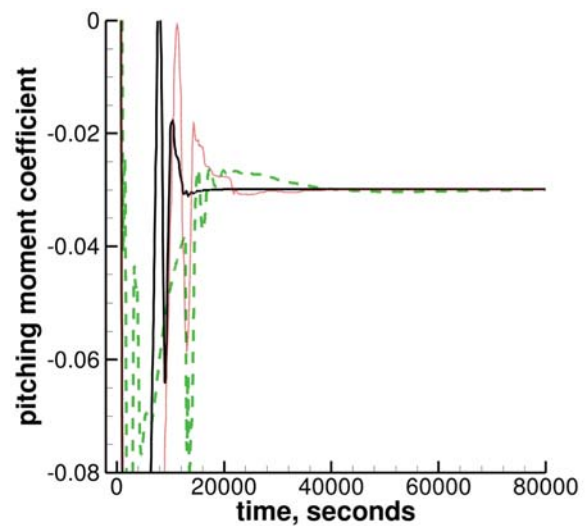
(e) maximum normalized eddy viscosity
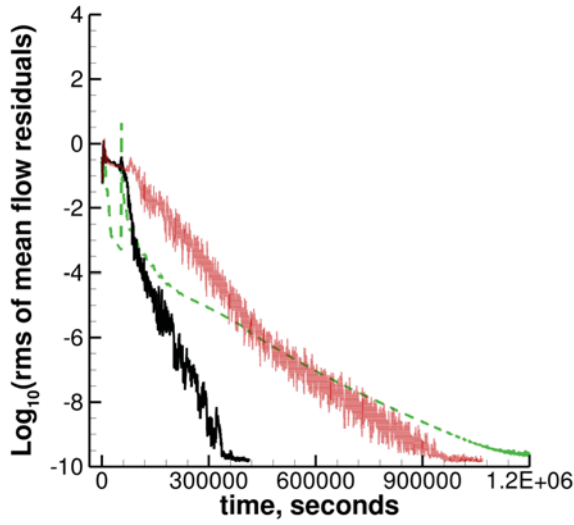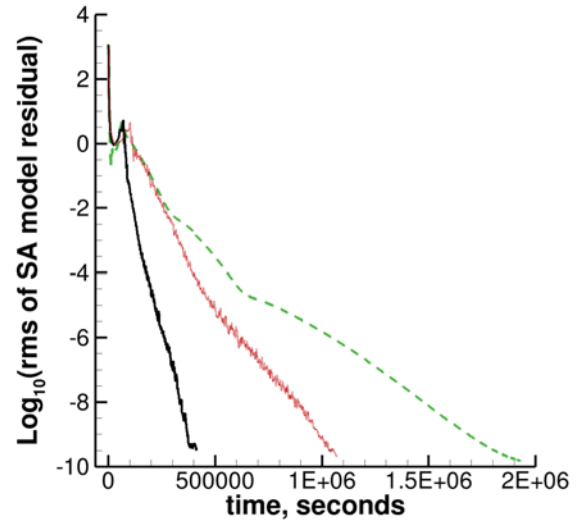
(f) drag coefficient

PA
HANIS-1
HANIS-4
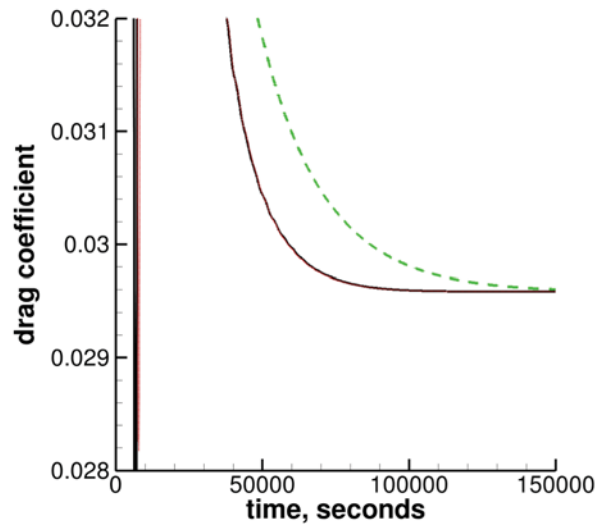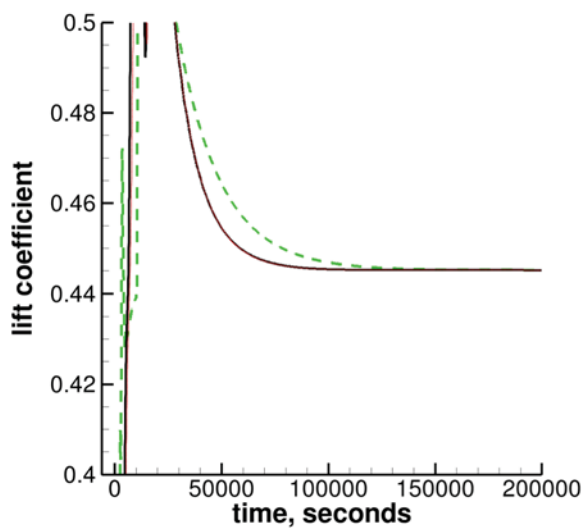
(g) lift coefficient

(h) pitching moment coefficient

Figure 27. Concluded.

(a) rms of mean flow residuals

(b) rms of SA model residual

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 28. NASA Common Research Model (CRM) solution convergence history using a coarse mixed element grid. $M_\infty = 0.85$, $\alpha = 2°$, $Re_c = 5x10^6$. Convective term in SA model is first-order accurate.

(e) maximum normalized eddy viscosity

(f) drag coefficient

PA
HANIS-1
HANIS-4

(g) lift coefficient

(h) pitching moment coefficient

Figure 28. Concluded.

American Institute of Aeronautics and Astronautics

(a) rms of mean flow residuals

(b) rms of SA model residuals

(c) rms of combined mean flow and SA residuals

(d) CFL for mean flow and SA model

Figure 29. NASA Common Research Model (CRM) solution convergence history using a medium mixed element grid. $M_\infty = 0.85$, $\alpha = 2°$, $Re_c = 5x10^6$. Convective term in SA model is first-order accurate.

American Institute of Aeronautics and Astronautics

(e) maximum normalized eddy viscosity

(f) drag coefficient

PA
HANIS-1
HANIS-4

(g) lift coefficient

(h) pitching moment coefficient

Figure 29. Concluded.

American Institute of Aeronautics and Astronautics

(a) rms of mean flow residuals

(b) rms of SA model residual
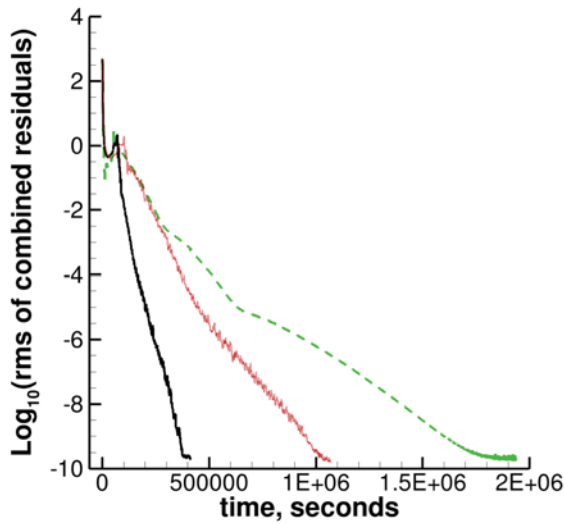
(c) rms of combined mean flow and SA residuals
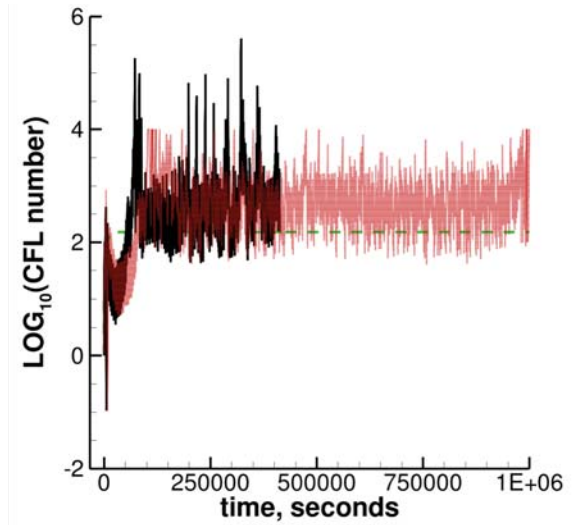
(d) CFL for mean flow and SA model

Figure 30. NASA Common Research Model (CRM) solution convergence history using a fine mixed element grid. $M_\infty = 0.85$, $\alpha = 2°$, $Re_c = 5\text{x}10^6$. Convective term in SA model is first-order accurate.
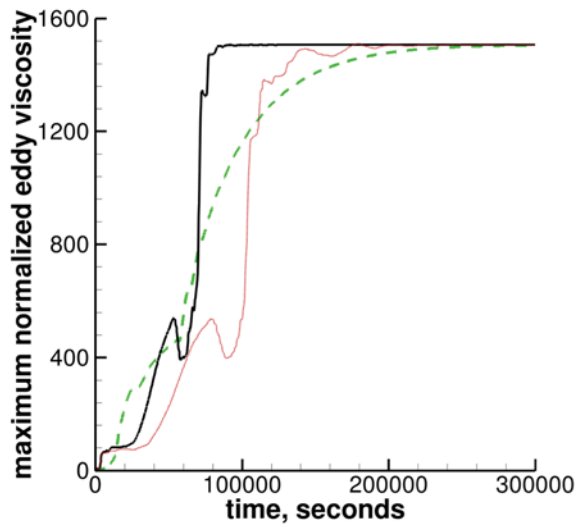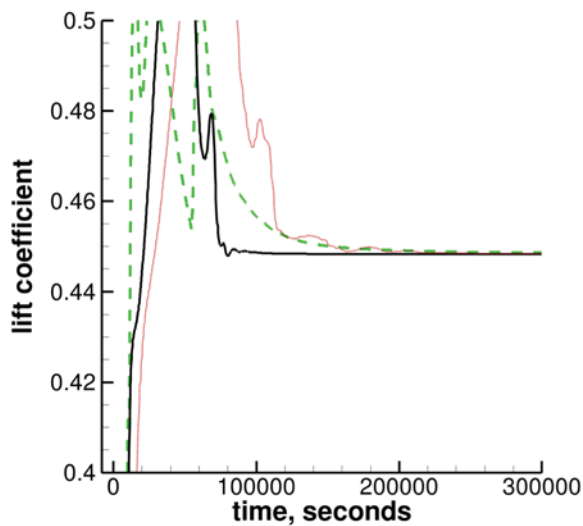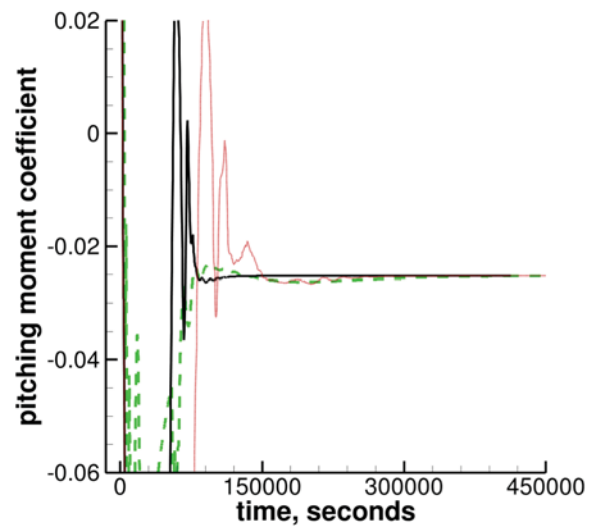
(e) maximum normalized eddy viscosity

(f) drag coefficient

PA
HANIS-1
HANIS-4
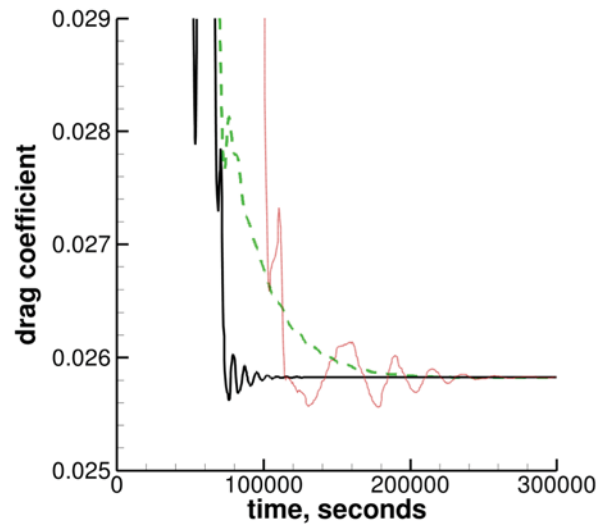
(g) lift coefficient

(h) pitching moment coefficient

Figure 30. Concluded.

69