Improvements in Iterative Convergence of FUN3D Solutions

Li Wang *
NASA Langley Research Center, Hampton, VA, USA
Boris Diskin†
National Institute of Aerospace, Hampton, VA, USA
Eric J. Nielsen ‡
NASA Langley Research Center, Hampton, VA, USA
Yi Liu §
National Institute of Aerospace, Hampton, VA, USA

This paper presents a hierarchical adaptive nonlinear iteration method (HANIM) implemented in the NASA computational fluid dynamics code, FUN3D, to improve robustness and computational efficiency. In contrast to the legacy FUN3D iterative solver that relies on an approximate Jacobian, a simple multicolor Gauss-Seidel point-implicit iteration scheme, and linear Courant-Friedrichs-Lewy number (CFL) ramping, HANIM is based upon a hierarchy of modules including preconditioner, generalized conjugate residual, realizability check, nonlinear control, and CFL adaption modules. HANIM performance is systematically compared with the performance of the legacy solver of FUN3D and a baseline solver based on a preconditioner alone. Iterative solutions are compared for three benchmark cases: a subsonic separated flow around a hemisphere cylinder, a supersonic flow through a long duct, and a subsonic flow over the NASA wing-fuselage juncture model. Two Reynolds-averaged Navier-Stokes turbulence models are used in these computations, namely, the negative variant of the linear one-equation Spalart-Allmaras model and its nonlinear extension based on quadratic constitutive relations.

I. Introduction

Computational fluid dynamics (CFD) is playing a significant and growing role in analysis and design of aerospace vehicles. Advanced simulation capabilities enable reduction in wind-tunnel and flight tests, provide physical insight, and reduce cost and risks of development and incorporation of new aerospace technologies. The effectiveness and impact of CFD on the design and analysis of aerospace systems is largely driven by the ability to reliably conduct many accurate simulations in a relatively short time. This task is complicated by a multitude of complex flow phenomena encountered in practical aerodynamics, such as boundary-layer separation, vortex shedding, and compressibility effects. Additional complications arise in the presence of increasingly irregular topology and highly stretched elements that are common in computational grids used for shock capturing and resolution in boundary layers. These challenges motivate the CFD community to advance solver technologies that improve efficiency, robustness, and scalability of CFD solutions.[1]

Newton iteration methods[2–4] are popular strategies for solving large systems of nonlinear equations. When the iterative solution arrives within a convergence basin, Newton iterations exhibit quadratic convergence rates and are often able to converge residuals of highly nonlinear equations to machine-zero levels. Although the machine-zero convergence criterion is not strictly required for most engineering problems, key aerodynamic quantities such as forces and particularly moments are not guaranteed to fully converge if the system residual has only been reduced by several orders of magnitude. Nonlinear iterations compute a solution update of a linear system that represents the exact linearization of the governing equations. The linearization is performed around the current nonlinear solution.

*Research Aerospace Engineer, AIAA Associate Fellow, li.wang@nasa.gov
†Sr. Research Fellow, AIAA Associate Fellow, bdiskin@nianet.org
‡Research Scientist, AIAA Associate Fellow, eric.j.nielsen@nasa.gov
§Sr. Research Engineer, AIAA member, yi.liu@nianet.org

Exact Newton steps are often computationally expensive, the exact linearization matrix may be ill-conditioned, and a good initial solution approximation is required to attain iterative convergence. In practice, various simplifications are introduced to reduce computational cost and dependence on the initial solution. A typical approach to incorporate Newton methods in CFD solvers is a pseudotransient continuation strategy,[5,6] in which a pseudotime term is added to the diagonal of the linearization matrix to enhance the diagonal dominance of the linear system. The pseudotime term is controlled by a pseudotime step and Courant-Friedrichs-Lewy number (CFL). Newton's method is recovered as the CFL number tends to infinity. The continuation strategy facilitates iterative convergence through transient stages, especially when highly anisotropic elements are present in boundary layers or wake regions. Various methods have been applied to solve the system of linear equations, including Jacobi and Gauss-Seidel (GS) point-iterative methods,[3,7] Krylov subspace constructions,[8–10] agglomeration[2] and algebraic[11,12] multigrid methods.

This paper reports on an implementation of a strong iterative nonlinear solver based on principles of Newton iterations in an operational CFD code, FUN3D, which is developed and supported by the NASA Langley Research Center (LaRC). FUN3D is a second-order, unstructured, node-centered, finite-volume CFD solver that is widely used for aerodynamic research and as a production tool by the U.S. Government[13–21] and aerospace industry.[22–25] Significant amount of research has been conducted in FUN3D in solver technology and high-performance computing for aerospace applications.[4,16,18,26–29] Recently, a novel hierarchical adaptive nonlinear iteration method (HANIM) has been implemented in FUN3D to improve efficiency and robustness for steady and unsteady Reynolds-averaged Navier-Stokes (RANS) simulations. The legacy FUN3D solver uses an approximate linearization and simple point-implicit multicolor GS iterations. In contrast, HANIM is based on a hierarchy of modules including a preconditioner, a generalized conjugate residual (GCR) linear solver, realizability check, nonlinear control, and CFL adaption modules. The specific HANIM implementation reported in this paper follows Refs. 30–32. While FUN3D HANIM has been successfully applied to several unsteady problems, including rotorcraft simulations on dynamic overset grids, this paper focuses on steady RANS solutions. Two one-equation eddy-viscosity turbulence models are employed. The first turbulence model is the negative variant of the linear Spalart-Allmaras (SA-neg) turbulence model;[33,34] the other turbulence model is a nonlinear extension of the SA-neg model that uses quadratic constitutive relations, SA-neg-QCR2000.[35]

In this paper, HANIM performance is studied for three benchmark cases: (1) a subsonic separated flow around a hemisphere-cylinder configuration, (2) a supersonic flow through a long duct with a square cross-section, and (3) a subsonic flow over a NASA wing-fuselage juncture configuration. The hemisphere-cylinder flow is simulated with the SA-neg turbulence model; the square-duct and juncture flows are simulated with the SA-neg-QCR2000 turbulence model. The hemisphere-cylinder and square-duct configurations are widely used verification cases posted in the section "Cases and Grids for Turbulence Model Numerical Analysis" of the NASA LaRC Turbulence Modeling Resource (TMR) website.[a] The square-duct and juncture flow cases are also SA-neg-QCR2000 verification cases selected for High-Fidelity CFD Workshop.[b]

The rest of the paper is organized as follows. Section II describes key elements of iterative solvers. Numerical results and discussions are provided in Section III, where the performance of HANIM is systematically compared with the performance of a preconditioner-alone (PA) solver and the legacy FUN3D solver. Section IV draws conclusions.

II. FUN3D Iterative Solvers

A. Discretization Scheme

FUN3D[36,37] solves the RANS equations discretized on unstructured mixed-element grids. The grids may contain tetrahedra, pyramids, prisms, and hexahedra. Median-dual control volumes are constructed around grid points. Inviscid fluxes are computed at edge medians by using an approximate Riemann solver. In the current study, Roe's flux difference splitting[38] is used. For second-order accuracy, density, pressure, and velocity at dual faces are reconstructed by a UMUSCL (Unstructured Monotonic Upstream-centered Scheme for Conservation Laws) scheme.[39,40] For mixed-element grids, the UMUSCL parameter is set as $\kappa = 0.5$ for the meanflow equations. The spatial discretization of the turbulence model uses a first-order accurate convection scheme. For discretization of viscous fluxes, the Green-Gauss theorem is used to compute cell-based gradients. In the case of tetrahedral meshes, this approach is equivalent to a Galerkin type approximation. For nontetrahedral meshes, cell-based Green-Gauss gradients are combined with edge-based gradients to improve stability of viscous operators and prevent odd-even decoupling. Both edge-normal and face-tangent augmented viscous gradients[41] have been implemented in FUN3D, and in this work, the face-tangent gradient construction scheme is used due to its robustness on highly-skewed grids. The diffusion term in the turbulence

[a] https://turbmodels.larc.nasa.gov, accessed 11/05/2020

[b] https://turbmodels.larc.nasa.gov/highfidelitycfd_workshop2022.html, accessed 11/05/2020

model is handled in the same fashion as the viscous terms in the meanflow equations. The vorticity-based source term for the turbulence model is computed using velocity gradients evaluated by the Green-Gauss method on dual volumes.

The boundary conditions involved in the present study include farfield Riemann-invariant, farfield Roe-based, symmetry, and viscous-wall boundary conditions.[42] Both the farfield Riemann-invariant and the Roe-based boundary conditions are weakly enforced by constructing an exterior state based on the freestream conditions to close the boundary fluxes. The symmetry boundary condition strongly enforces zero normal velocity. The tangential momentum, mass, and energy conservation at the symmetry plane are enforced through the conservation equations. The no-slip viscous-wall boundary condition explicitly sets zero velocities on the wall boundary; density is computed through fluxes satisfying the mass conservation equation. Either an isothermal or adiabatic wall boundary condition can be selected to close the energy equation. The isothermal boundary condition explicitly specifies temperature at the boundary nodes. Local pressure can be determined from temperature and density. The adiabatic condition is imposed either by setting an adiabatic-wall temperature derived from a flat-plate approximation or by specifying a true thermal insulation condition with zero numerical heat flux.

B. Legacy Iterative Solver

Discrete nonlinear equations describing conservation laws can be expressed as:

$$\mathbf{R}(\mathbf{Q}) = \mathbf{0}, \tag{1}$$

where $\mathbf{Q}$ denotes the flow solution and $\mathbf{R}$ denotes the nonlinear steady-state residual. This nonlinear system of equations is solved by a defect-correction method with a pseudotime term expressed as:

$$\left[\frac{V}{\Delta\tau}\mathbf{I} + \widetilde{\frac{\partial \mathbf{R}}{\partial \mathbf{Q}}}\right]\Delta\mathbf{Q} = -\mathbf{R}(\mathbf{Q}^n), \tag{2}$$

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + \omega\,\Delta\mathbf{Q}. \tag{3}$$

Here, V is a dual-cell volume; $\Delta\tau$ is the local pseudotime step; $\mathbf{I}$ is the identity matrix; the superscript n is the nonlinear iteration index, and ω is an underrelaxation parameter.

The linear iterations decouple the meanflow equations from the turbulence-model equation. For the meanflow equations, the approximate Jacobian matrix, $[\widetilde{\partial\mathbf{R}/\partial\mathbf{Q}}]$, is formed based on linearization of first-order inviscid fluxes and second-order viscous fluxes. The inviscid terms of the Jacobian are based on linearization of the vanLeer inviscid flux. The Jacobian of the turbulence model includes linearization of the convection, diffusion, and source terms. Positivity of the diagonal term in the linearization of the turbulence-model equation is enforced prior to adding the pseudotime term. Nonlinear iterations can reuse the Jacobian computed at a previous iteration; a "smart" Jacobian update is scheduled depending on residual convergence.

At each nonlinear iteration, the linear system represented by Eq. (2) is solved using a specified number of point-implicit GS sweeps with multicolor ordering; 15 and 10 sweeps are usually specified for the meanflow and turbulence-model equations, respectively. The resulting underrelaxed solution update, $\omega\Delta\mathbf{Q}$, is added to the nonlinear solution at the n-th iteration. The updated nonlinear solution is required to satisfy realizability constraints at each grid point. Realizability is satisfied if local values of density and pressure are greater than a small positive threshold. The solutions reconstructed at dual-cell faces should also satisfy realizability constraints. Reconstructing a negative pressure or density causes a so-called catastrophic failure. In the case of a catastrophic failure, the second-order solution reconstruction is automatically replaced by a zero-gradient (piecewise constant) reconstruction reducing accuracy to first order. The CFL number can be ramped linearly within a specified number of nonlinear iterations.

C. Preconditioner Alone Solver

In this work, a PA solver is also used for performance assessment. The PA solver closely resembles the legacy solver described in Section II.B. with two significant differences: (1) the approximate Jacobian is based on Roe's flux difference splitting scheme and (2) the Jacobian is updated in each nonlinear iteration. The PA solver uses the same number of point-implicit GS sweeps as the legacy solver for the linear system (Eq.(2)).

D. HANIM Algorithm

HANIM is based upon a hierarchy of modules illustrated in Fig. 1. Each HANIM iteration starts from a realizable solution. HANIM modules generate a solution update and report success or failure based on the level of residual

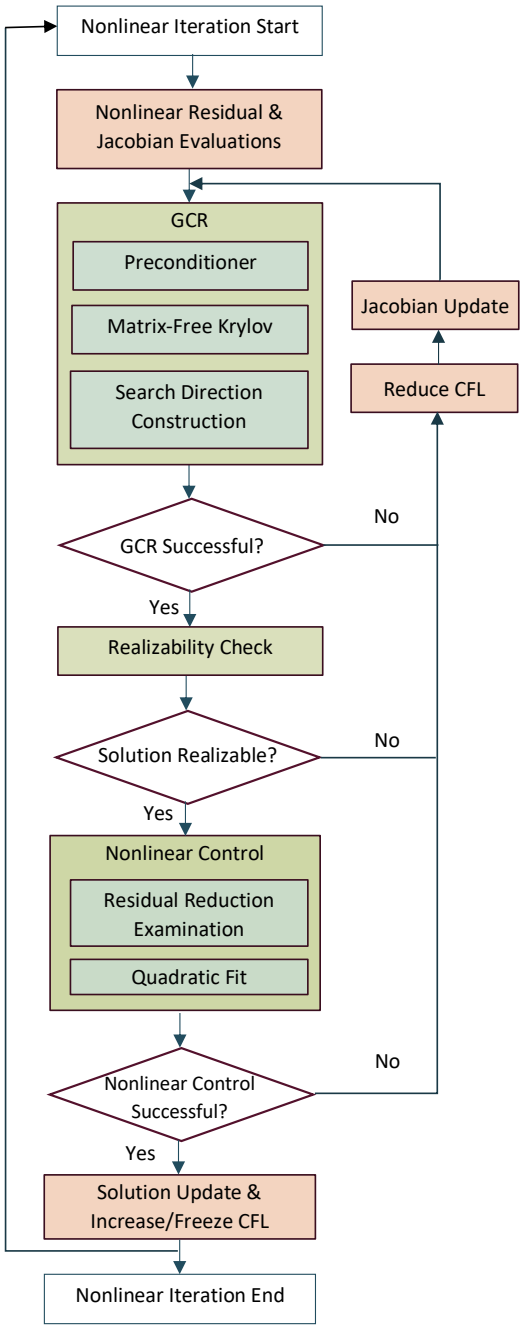


Figure 1. Flowchart illustrating key elements of HANIM algorithm.

reduction and realizability of the updated solution. If any of the modules fails, then the proposed update is discarded, the CFL number is reduced, and the nonlinear iteration repeats from the same solution. If all modules report success, the CFL number is increased, the solution is updated, and HANIM proceeds to the next nonlinear iteration. A similar HANIM algorithm has been previously implemented in a NASA cell-centered CFD code, USM3D.[30–32] The FUN3D implementation closely follows the USM3D HANIM formulation. Below, the key components of FUN3D HANIM are described in detail.

1. Preconditioner

Similar to the legacy iterative solver, the HANIM preconditioner uses a defect-correction scheme (Eq. (2)), a meanflow Jacobian based on a first-order approximation to inviscid fluxes, multicolor point-implicit GS iterations, and the same linearization of the turbulence-model equation. Unlike the legacy solver, the HANIM preconditioner always uses a Jacobian based on linearization of a consistent inviscid flux, i.e., the same approximate Riemann solver is used in evaluating the Jacobian and nonlinear residual. The preconditioner Jacobian is updated at the beginning of each nonlinear iteration using the most recent solution. The meanflow and turbulence-model equations are decoupled and solved in consecutive order. The initial solution update, $\Delta\mathbf{Q}$, is set to zero.

The preconditioner residual is defined as:

$$\mathbf{r} = -\mathbf{R}(\mathbf{Q}^n) - \left[\frac{V}{\Delta\tau}\mathbf{I} + \widetilde{\frac{\partial\mathbf{R}}{\partial\mathbf{Q}}}\right]\Delta\mathbf{Q}. \tag{4}$$

The multicolor point-implicit GS iterations (sweeps) continue until either the root-mean-square (RMS) norm of the preconditioner residual has attained the prescribed reduction level, the maximum number of sweeps have been performed, or preconditioner divergence has been detected. Divergence is detected when the RMS norm of the preconditioner residual has increased by more than two orders of magnitude. In the case of divergence, the preconditioner module reports failure and the solution process moves to the workflow handler that includes the CFL number and Jacobian updates and is discussed in Section II.D.5. If either the prescribed residual reduction or the maximum number of sweeps have been achieved, the preconditioner reports success. The preconditioner solution, $\Delta\mathbf{Q}$, is then used by the GCR solver to form a search direction.

2. GCR Solver

The GCR method solves a linear system of equations representing the exact linearization of the coupled meanflow and turbulence-model equations:

$$\left[\frac{V}{\Delta\tau}\mathbf{I} + \frac{\partial\mathbf{R}}{\partial\mathbf{Q}}\right]\Delta\mathbf{Q} = -\mathbf{R}(\mathbf{Q}^n). \tag{5}$$

Here, $[\partial\mathbf{R}/\partial\mathbf{Q}]$ is the exact linearization of the coupled nonlinear residual, Eq. (1). For second-order finite-volume discretizations, the exact linearization depends on solution values at neighbors of neighbor points and dramatically increases the computational cost and memory requirements. To avoid the increased computational cost and storage requirements, a Frechet derivative approach is used to evaluate the matrix-vector product shown in Eq. (5):

$$\frac{\partial\mathbf{R}}{\partial\mathbf{Q}}\Delta\mathbf{Q} \approx \frac{\mathbf{R}\left(\mathbf{Q} + \varepsilon\frac{\Delta\mathbf{Q}}{|\Delta\mathbf{Q}|}\right) - \mathbf{R}(\mathbf{Q})}{\varepsilon}|\Delta\mathbf{Q}|, \tag{6}$$

where $|\cdot|$ denotes an l_2 norm, $\Delta\mathbf{Q}$ is the preconditioner solution, and ε is a perturbation parameter, which is scaled with the solution magnitude:

$$\varepsilon = \max(1, \|\mathbf{Q}^n\|)\varepsilon_0. \tag{7}$$

Here, $\|\cdot\|$ denotes an RMS norm. Accuracy of the Frechet derivative is sensitive to round-off errors. The perturbation factor $\varepsilon_0 \approx 10^{-7}$ is suitable for turbulent flows with moderate or high speeds ($M_\infty > 0.1$). For low-speed turbulent or laminar flows, $\varepsilon_0 \approx 10^{-4}$ provides better accuracy of the Frechet derivative and enhances iterative convergence. Furthermore, to completely eliminate the round-off sensitivity associated with the Frechet derivative, the real-valued Frechet derivative can be replaced with a complex-valued Frechet derivative using an imaginary perturbation of magnitude $\varepsilon_1 = 10^{-50}$. In this case, the complex-valued Frechet derivative is evaluated as

$$\frac{\partial\mathbf{R}}{\partial\mathbf{Q}}\Delta\mathbf{Q} \approx \frac{imag\left[\mathbf{R}\left(\mathbf{Q} + i\varepsilon_1\frac{\Delta\mathbf{Q}}{|\Delta\mathbf{Q}|}\right)\right]}{\varepsilon_1}|\Delta\mathbf{Q}|. \tag{8}$$

This complex-variable derivative is practically exact, but carries a computational penalty as the complex-valued nonlinear residual evaluation is about three times slower than the corresponding real-valued approach. The real-valued Frechet derivative is used in all computations reported in this work.

The GCR module in the current HANIM implementation can use a maximum of K search directions to construct an approximate solution to Eq. (5); K is an input parameter. The corresponding solver is denoted as HANIM-K. Initially, the GCR correction is set to zero, $\Delta\mathbf{Q} = 0$. The initial GCR residual is equal to the nonlinear residual evaluated at the beginning of the n-th iteration,

$$\mathbf{r}_0 = -\mathbf{R}(\mathbf{Q}^n). \tag{9}$$

A new search direction is computed by a call to the preconditioner with the right-hand side set to the current GCR residual:

$$\left[\frac{V}{\Delta\tau}\mathbf{I} + \widetilde{\frac{\partial\mathbf{R}}{\partial\mathbf{Q}}}\right]\Delta\mathbf{Q}_k = \mathbf{r}_{k-1}, \quad k = 1,2,\cdots,K. \tag{10}$$

The preconditioner iterations are performed for the decoupled linear system. Here, $\Delta\mathbf{Q}_k$ is the preconditioner solution to be used in forming the k-th search direction; $\mathbf{r}_{k-1}$ represents the GCR residual computed after applying the $(k-1)$-th search direction. The new search direction, $\mathbf{b}_k$, is computed and normalized as follows:

$$\mathbf{b}_k = \frac{V}{\Delta\tau}\Delta\mathbf{Q}_k + \frac{\mathbf{R}\left(\mathbf{Q}^n + \varepsilon\dfrac{\Delta\mathbf{Q}_k}{|\Delta\mathbf{Q}_k|}\right) - \mathbf{R}(\mathbf{Q}^n)}{\varepsilon}\,|\Delta\mathbf{Q}_k|\,, \tag{11}$$

$$\mathbf{b}_k = \frac{\mathbf{b}_k}{|\mathbf{b}_k|}, \quad \Delta\mathbf{Q}_k = \frac{\Delta\mathbf{Q}_k}{|\mathbf{b}_k|}. \tag{12}$$

Next, an orthonormalization procedure is performed for the search direction $\mathbf{b}_k$ against previously stored search directions $\mathbf{b}_j$ $(1 \le j < k)$; solution updates are changed accordingly.

$$\mu_k = \mathbf{b}_k^T\mathbf{b}_j, \quad \mathbf{b}_k = \mathbf{b}_k - \mu\mathbf{b}_j, \quad \Delta\mathbf{Q}_k = \Delta\mathbf{Q}_k - \mu\Delta\mathbf{Q}_j,$$
$$\mathbf{b}_k = \frac{\mathbf{b}_k}{|\mathbf{b}_k|}, \quad \Delta\mathbf{Q}_k = \frac{\Delta\mathbf{Q}_k}{|\mathbf{b}_k|}. \tag{13}$$

The new orthonormal search direction $\mathbf{b}_k$ is stored for future search-direction orthonormalization. The residual projection on the k-th search direction is calculated, then the GCR correction and GCR residual are updated

$$\gamma_k = \mathbf{b}_k^T\mathbf{r}_{k-1}, \quad \Delta\mathbf{Q} = \Delta\mathbf{Q} + \gamma_k\Delta\mathbf{Q}_k, \quad \mathbf{r}_k = \mathbf{r}_{k-1} - \gamma_k\mathbf{b}_k. \tag{14}$$

The GCR module declares success if the RMS norm of GCR residual is reduced below a specified reduction target, μ_{GCR}:

$$\frac{\|\mathbf{r}_k\|}{\|\mathbf{r}_0\|} \le \mu_{GCR}. \tag{15}$$

The norms are computed for the coupled residual vectors. In the case of GCR success, HANIM proceeds to the realizability check module discussed in the next section. If the target reduction has not been met, GCR repeats the search-direction procedure, Eqs. (10)–(15). If the target GCR residual reduction is not satisfied after applying the K-th search direction, the GCR module reports failure. Failure reported either by the preconditioner or GCR modules results in interruption of the current nonlinear iteration, and HANIM proceeds to the workflow handler module.

3. Realizability Check

A successful GCR step offers a solution update, $\Delta\mathbf{Q}$, for the nonlinear solution. The updated solution should be realizable. The realizability check immediately follows the GCR module.

In the realizability check, the solution correction, $\Delta\mathbf{Q}$, is added to the nonlinear solution, $\mathbf{Q}^n$. If the updated density and pressure remain greater than the specified tolerance, HANIM proceeds to the nonlinear control module. If the updated density or pressure become smaller than the specified tolerance at any grid point, the realizability check reports failure, and HANIM proceeds to the workflow handler module.

4. Nonlinear Control

The nonlinear control module seeks a global underrelaxation parameter, $0 \leq \omega \leq 1$, to maximize the residual reduction for the updated solution, $\mathbf{Q}^n + \omega\Delta\mathbf{Q}$. The underrelaxation factor is first set to unity, and the pseudotime nonlinear residual, $\mathbf{R}_\tau$, is evaluated as follows:

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + \Delta\mathbf{Q}, \tag{16}$$

$$\mathbf{R}_\tau(\mathbf{Q}^{n+1}) = \frac{V}{\Delta\tau}\Delta\mathbf{Q} + \mathbf{R}(\mathbf{Q}^{n+1}). \tag{17}$$

The nonlinear control reports success if

$$\|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\| / \|\mathbf{R}(\mathbf{Q}^n)\| \leq \mu_{nl} \tag{18}$$

where μ_{nl} is the residual reduction target set as

$$\mu_{nl} = 0.5(1 + \mu_{GCR}). \tag{19}$$

If the target residual reduction is satisfied, HANIM proceeds to the workflow handler module.

In the case where $\|\mathbf{R}_\tau(\mathbf{Q})\| / \|\mathbf{R}(\mathbf{Q}^n)\| > \mu_{nl}$, the magnitude of the current correction, $\Delta\mathbf{Q}$, is compared to the magnitude of the round-off error. Round-off errors can play a noticeable role at the final stage of convergence. In the current work, corrections that satisfy $\|\Delta\mathbf{Q}\| / \|\mathbf{Q}^n\| \leq 10^{-12}$, are considered to be dominated by round-off errors. In such cases, if $\|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\|$ achieves any reduction relative to $\|\mathbf{R}(\mathbf{Q}^n)\|$, the nonlinear control module reports success, but the CFL number is frozen.

If $\|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\| / \|\mathbf{R}(\mathbf{Q}^n)\| > \mu_{nl}$ and $\|\Delta\mathbf{Q}\| / \|\mathbf{Q}^n\| > 10^{-12}$, then a quadratic search is performed to determine the underrelaxation factor in order to improve residual reduction. To determine the underrelaxation factor, two scalar functions are introduced to parameterize the variation of $\mathbf{R}_\tau$ and the GCR residual,

$$f(\omega) = \left\| \frac{V}{\Delta\tau}\omega\Delta\mathbf{Q} + \mathbf{R}(\mathbf{Q}^n + \omega\Delta\mathbf{Q}) \right\|, \tag{20}$$

$$g(\omega) = \left\| \frac{V}{\Delta\tau}\omega\Delta\mathbf{Q} + \frac{\partial\mathbf{R}}{\partial\mathbf{Q}}\omega\Delta\mathbf{Q} + \mathbf{R}(\mathbf{Q}^n) \right\|. \tag{21}$$

Assuming quadratic variation for $f(\omega) = a\omega^2 + b\omega + c$ and linear variation for $g(\omega) = b\omega + c$, the following relations hold: $c = f(0)$, $b = g(1) - c$, and $a = f(1) - b - c$. The quadratic function $f(\omega)$ yields a minimum at $\omega = -b/(2a)$. Substituting the expressions of a, b, and c yields:

$$\omega = -\frac{g(1) - f(0)}{2\,(f(1) - g(1))}. \tag{22}$$

Here, $f(0) = \|\mathbf{R}(\mathbf{Q}^n)\|$, $f(1) = \|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\|$, and $g(1) = \left\| \frac{V}{\Delta\tau}\Delta\mathbf{Q} + \frac{\partial\mathbf{R}}{\partial\mathbf{Q}}\Delta\mathbf{Q} + \mathbf{R}(\mathbf{Q}^n) \right\|$ is the RMS norm of the linear residual computed by the GCR module (c.f. Eqs. (5) and (6)). The underrelaxation factor determined by Eq. (22) is used to update the nonlinear solution as $\mathbf{Q}^{n+1} = \mathbf{Q}^n + \omega\Delta\mathbf{Q}$. If $\|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\|$ achieves the target residual reduction, Eq. (18), the nonlinear control module reports success, but the CFL number is frozen. If the target reduction of $\|\mathbf{R}_\tau(\mathbf{Q}^{n+1})\|$ is not attained and round-off error is not dominant, the nonlinear control module reports failure, and HANIM proceeds to the workflow handler module.

5. Workflow Handler

The workflow handler module processes the success/failure and CFL-update messages generated by other HANIM modules and prepares HANIM for the next nonlinear iteration.

In the case of success, the CFL number is either increased by a specified factor as $\mathrm{CFL} = \mathrm{CFL} \times \mathrm{CFL}_{incr}$ or remains frozen and the Jacobian matrix, $[\widetilde{\partial\mathbf{R}/\partial\mathbf{Q}}]$, is recomputed using the updated solution. In the case of failure, the solution is restored to the state, $\mathbf{Q}^n$, at the beginning of the n-th nonlinear iteration, and the CFL number is reduced by a specified decrease factor as $\mathrm{CFL} = \mathrm{CFL} \times \mathrm{CFL}_{decr}$. The Jacobian matrix is updated to accommodate the reduced CFL number. The diagonal terms of the Jacobian are updated without recomputing the linearized fluxes.

In the current work, the CFL adaptation parameters are set to $\text{CFL}_{incr} = 2$ and $\text{CFL}_{decr} = 0.1$. The CFL number is bounded by the maximum and minimum limits $\text{CFL}_{max} = 10^{10}$ and $\text{CFL}_{min} = 10^{-12}$. This range of 22 orders of magnitude is expected to be sufficient for many applications, although adjustments can be made if needed. If the CFL number reaches CFL_{min}, HANIM terminates the nonlinear iterations, as solution updates become too small to make any significant changes in the nonlinear solution.

6. Weighted Norms

The RMS and l_2 norms in HANIM can be weighted. In unstructured-grid CFD computations, large cells are often present near farfield boundaries. Small variations in a solution may lead to spurious residual fluctuations due to large grid volumes. Appropriate weights can be chosen to reduce the influence of large cells on the residual. The weighted norms are defined as

$$\text{RMS}_w = \sqrt{\frac{\sum_{i=1}^{N} (w_i r_i)^2}{N}}, \qquad l_{2,w} = \sqrt{\sum_{i=1}^{N} (w_i r_i)^2}, \tag{23}$$

where r_i is the i-th component of vector $\mathbf{r}$, w_i is the weight applied to this component, and N is the total number of equations solved on a given grid. Two weighting options have been implemented.. The first option sets the weight to be inversely proportional to the dual-cell volume with unity as the upper limit:

$$w_i = \min\left(1, \frac{1}{V_i}\right). \tag{24}$$

Here, V_i is the dual-cell volume. This option reduces the weight of larger cells while leaving the weight of smaller cells ($V_i < 1$) unchanged. The second option sets the weight to be inversely proportional to the characteristic cell size that is computed as the cube root of the dual-cell volume for a three-dimensional grid:

$$w_i = \frac{1}{\sqrt[3]{V_i}}. \tag{25}$$

This option reduces the weight of large cells and increases the weight of small cells, e.g., in a viscous boundary layer.

III. Numerical Results

In this section, HANIM performance on increasingly refined grids is compared with performance of the FUN3D legacy solver and the PA solver described in Sections II.B.–C. Iterative convergence is examined for three benchmark cases, namely, a subsonic separated flow around a hemisphere-cylinder configuration, a supersonic flow through a long duct with a square cross-section, and a subsonic flow over a NASA wing-fuselage juncture configuration. For all solutions shown in this work, the initialization uses freestream conditions, and the same number and type of processing cores are employed across all the solvers on any given grid. The dual-socket 20-core Intel Gold 6148 Skylake processors of the NASA LaRC K-Cluster and the dual-socket 14-core Intel Broadwell processors of the NASA Advanced Supercomputing (NAS) Pleiades system were used in the computations.

A. Subsonic Flow around Hemisphere Cylinder

The first test case is a subsonic flow around a three-dimensional hemisphere-cylinder configuration at a freestream Mach number of 0.6 and an angle of attack of 19 degrees. The geometry and flow conditions are presented in literature[26] and at the TMR website.[c] In the corresponding experiment,[43] the cylinder diameter (D) was 1 inch, the length was 10 inches, and the Reynolds number was 350,000 per inch. The reference temperature was 540 °Rankine.

Four computational grids containing prismatic and hexahedral elements are used in this study. More details on the grid metrics are provided in Table 1. The origin of the coordinate system is placed at the hemisphere apex. The positive x-direction coincides with the cylinder axis; the positive z-direction is up. A symmetry boundary condition is imposed at the $y = 0$ plane. A no-slip adiabatic-wall boundary condition is imposed at the hemisphere cylinder surface. Farfield freestream boundary conditions are specified at the outer quadrant of the sphere with a radius of $100D$ centered at $(x, y, z) = (10, 0, 0)$. The plane at $x = 10D$ is assigned a back-pressure outflow boundary condition[42] and the unit ratio

[c] `https://turbmodels.larc.nasa.gov/hc3dnumericspart2_val.html`, accessed 11/05/2020

Table 1. Grid and iterative convergence metrics for subsonic flow around the hemisphere cylinder.

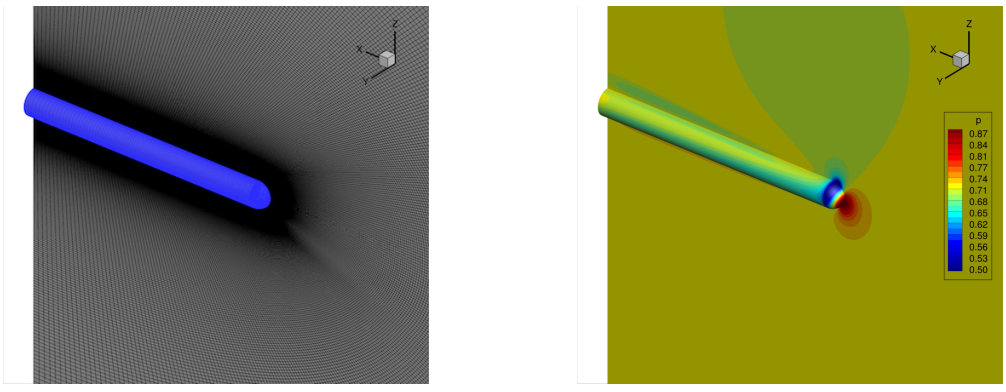
Grid	Nodes	Ncells	Iterations			Wall-clock seconds			Speedup over
			PA	Legacy	HANIM	PA	Legacy	HANIM	PA (Legacy)
1	71,368,353	78,643,200	10,907	11,398	1,874	107,044.0	100,840.5	51,646.9	2.1 (2.0)
2	8,995,153	9,830,400	8,976	9,809	675	39,195.6	31,099.3	15,002.6	2.6 (2.1)
3	1,143,081	1,228,800	2,984	3,693	359	3,018.5	2,829.0	1,190.3	2.5 (2.4)
4	147,637	153,600	1,626	2,039	389	1,126.3	942.5	428.3	2.6 (2.2)

of the static and reference pressure. The turbulent flow is simulated with the linear SA-neg turbulence model. The farfield Spalart turbulence variable is $\hat{\nu}_{farfield} = 3.0\hat{\nu}_\infty$, where $\hat{\nu}_\infty$ is the freestream laminar kinematic viscosity. The meanflow Prandtl number is 0.72, and the turbulent Prandtl number is 0.9. Figure 2(a) illustrates the geometry and the computational grid (grid 2) at the symmetry plane. Figure 2(b) displays pressure contours near the hemisphere-cylinder surface.

In HANIM, the maximum number of preconditioner sweeps is set as 200 for both the meanflow and turbulence equations. The target preconditioner residual reduction is 0.5. The GCR tolerance, μ_{GCR}, is set to 0.92 and only one search direction is used in each nonlinear iteration (HANIM-1). The weighting scheme corresponding to Eq. (25) is used.

The dual-socket 14-core Intel Broadwell processors of the NAS Pleiades system were used for the computations on the finest grid 1 with 1,400 processing cores. The dual-socket 20-core Intel Gold 6148 Skylake processors of the NASA LaRC K-Cluster were used for the computations on grids 2, 3, and 4 with 160, 80, and 8 processing cores, respectively. Grid convergence metrics including number of nonlinear iterations, wall-clock time in seconds, and HANIM speedup factors over the PA and legacy solvers on different grids are provided in Table 1. The wall-clock time corresponds to the elapsed real time for each individual solver to converge the unweighted RMS norm of the combined meanflow and turbulence-model residuals below 10^{-13}. HANIM provides a significant reduction of time to solution over the PA solver (a speedup factor range of 2.1–2.6 has been obtained for grids 1–4) and the legacy solver (a speedup factor range of 2.0–2.4 has been obtained for grids 1–4). Note that the PA solver uses fewer nonlinear iterations than the legacy solver; however, the legacy solver requires less time for convergence because of less frequent Jacobian evaluations.

Convergence histories of the PA, legacy, and HANIM solvers on grid 1 is shown in Figs. 3(a) and 3(b), where the unweighted RMS norm of the steady-state residual for the combined meanflow and turbulence equations is plotted versus the nonlinear iterations and computational time, respectively. Figure 3(c) shows the CFL histories for the HANIM, the baseline PA, and legacy solvers. To attain convergence, HANIM operates at much higher CFL numbers and uses an order of magnitude fewer nonlinear iterations than the PA and legacy solvers. Reference 44 reports a point-implicit preconditioned HANIM scheme implemented in the cell-centered finite-volume USM3D code and shows a similar speedup (1.7–2.2) over its baseline PA solver for the same hemisphere cylinder configuration but at a lower angle of attack (0 degrees). Reference 30 indicates that HANIM can provide much higher speedups for the same configuration with a line-implicit preconditioner. Line-implicit preconditioning will be considered for implementation

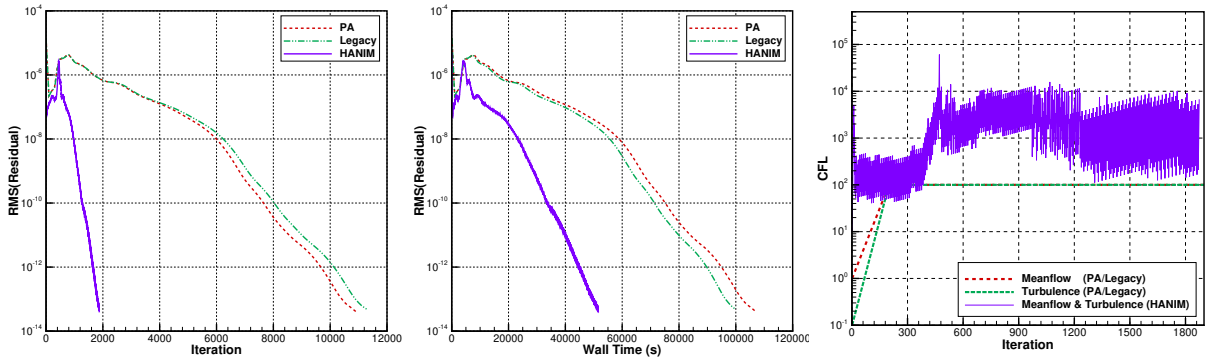


(a) Grid 2 (b) Pressure contours

Figure 2. Computational grid (grid 2) and pressure contours for subsonic flow around the hemisphere cylinder.

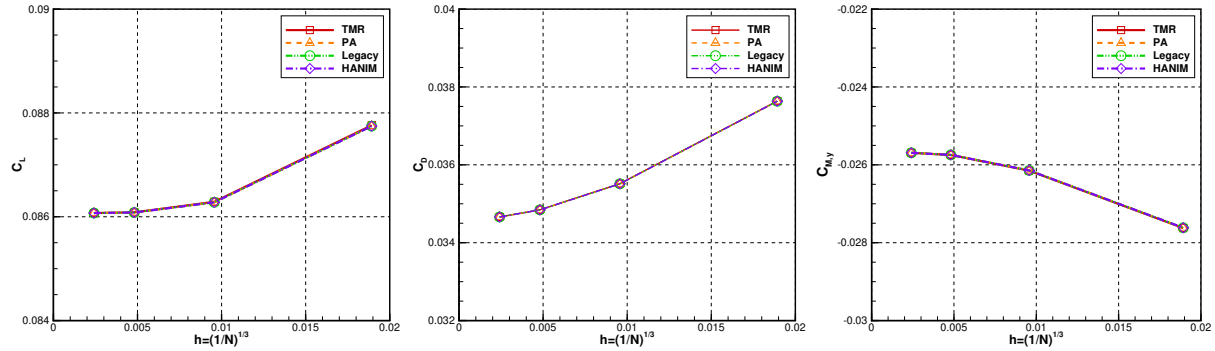
in the near future.

Figure 4 shows grid convergence of lift, drag, and pitching moment coefficients obtained by the PA, legacy, and HANIM solvers, as well as the reference convergence of these coefficients provided by the FUN3D solutions posted at the NASA TMR website. The force and moment coefficients obtained by the three solvers are identical to the reference TMR coefficients. As shown in Fig. 5, distributions of surface pressure, the x-component of skin friction along $y = 0$, and the surface pressure along $x = 5$ computed by the PA, legacy, and HANIM solvers on the finest grid (grid 1) match the TMR solution exactly and are also in good agreement with the experimental results.



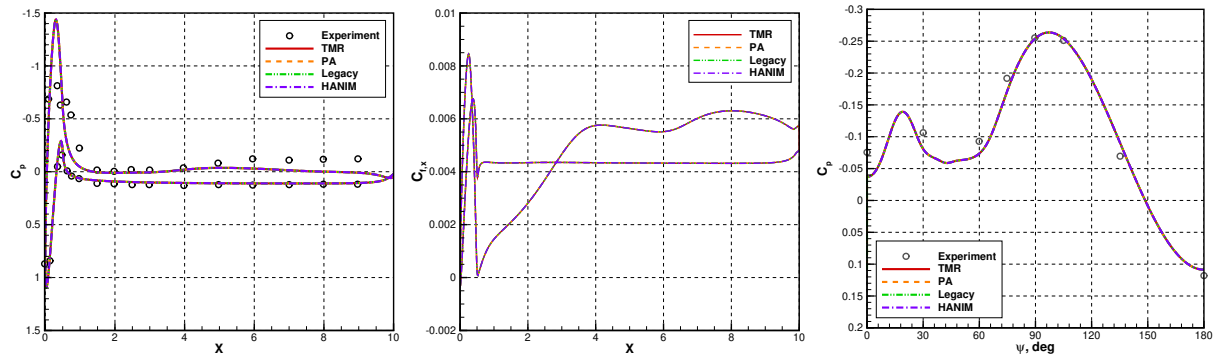
(a) Residual vs. iteration (b) Residual vs. time (c) CFL vs. iteration

Figure 3. Convergence of HANIM, PA, and legacy solvers for subsonic flow around the hemisphere cylinder on grid 1.



(a) Lift (b) Drag (c) Pitching moment

Figure 4. Grid convergence of lift, drag, and pitching moment coefficients in subsonic flow around the hemisphere cylinder.



(a) Surface pressure along $y = 0$ (b) x-component of skin friction along $y = 0$ (c) Surface pressure vs. azimuthal angle at $x = 5$

Figure 5. Comparisons of pressure and x-component of skin friction on grid 1 for subsonic flow around the hemisphere cylinder.

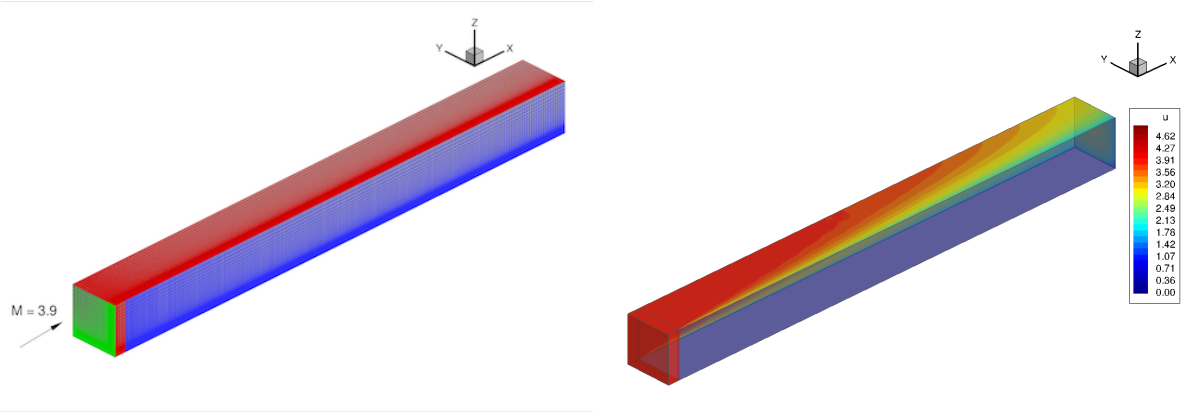
B. Supersonic Duct Flow

The second test case involves a supersonic turbulent flow in a long duct with a square cross-section. The reference Mach number is 3.9, the Reynolds number is 508,000 based on the duct height or width, $D = 25.4$mm, and the reference temperature is 520 °Rankine. The walls of the duct are modeled as adiabatic solid walls. Due to symmetry, only one quarter of the duct is computed; symmetry boundary conditions are applied at two boundaries.

The origin of the coordinate system is set at the lower-left corner of the duct inlet. The positive direction of the x-axis is the streamwise direction along the duct. The CFD grid is extended upstream of $x = 0$ to $x = -1.26829D$ to allow for a small run before the duct inlet using a symmetry boundary condition and the length of the duct is $52D$. Computational grids (grids 2, 3, and 4) with nested hexahedral elements are generated using a grid-generation tool available on the TMR website. These uniformly refined grids are employed to assess solver performance and solution accuracy. Figure 6(a) shows grid 4 that contains 3,170,833 nodes and 3,096,576 hexahedral elements; an enlarged scale for the y- and z-axes is used for clarity. The boundaries $y = 0.5D$ and $z = 0.5D$ (shown in red) are locations of symmetry boundary conditions. The boundaries $y = 0$ and $z = 0$ (shown in blue) are assigned a no-slip, adiabatic-wall boundary condition. The boundaries $x = -1.26829D$ and $x = 52D$ (shown in green) are the supersonic inflow and supersonic outflow boundary conditions, respectively. The simulations are performed with the SA-neg-QCR2000[35] turbulence model and without flux limiters or gradient (catastrophic) limiters. The farfield Spalart turbulence variable is $\hat{\nu}_{farfield} = 3.0\hat{\nu}_\infty$. The meanflow Prandtl number is 0.72, and the turbulent Prandtl number is 0.9. Figure 6(b) shows contours of streamwise velocity with an enlarged scale for the y- and z-axes.

In HANIM, the maximum number of multicolor sweeps in the preconditioner is 100 for both the meanflow and turbulence-model equations. The preconditioner residual reduction target is 0.2. The maximum number of search directions is 4 (HANIM-4) and the maximum CFL number is 1,000. The GCR residual reduction target is set to 0.92. The weighting scheme corresponding to Eq. (24) is used.

Table 2 lists the grid and convergence metrics corresponding to the PA, legacy, and HANIM solvers on grids 3 and 4. The legacy solver of FUN3D was not tried on grid 2 due to time constraints. The dual-socket 14-core Intel Broadwell processors of the NAS Pleiades system were used for the computations on grid 2 with 3,976 processing cores. The dual-socket 20-core Intel Gold 6148 Skylake processors of the NASA LaRC K-Cluster were used for the computations on grids 3 and 4 with 400 and 240 processing cores, respectively. Similar to the hemisphere-cylinder case presented in Section III.A, the PA solver uses fewer nonlinear iterations but more time than the legacy solver to



(a) Grid 4 (b) Streamwise velocity

Figure 6. Computational grid and contours of streamwise velocity (grid 4) for the supersonic square-duct flow. Scales on y and z axes are exaggerated ($x : y : z = 1 : 10 : 10$).

Table 2. Grid and convergence quantities for the supersonic square-duct flow.

Grid	Nodes	Ncells	Iterations			Wall-clock seconds			Speedup over
			PA	Legacy	HANIM	PA	Legacy	HANIM	PA (Legacy)
2	199,362,625	198,180,864	77,673	N/A	9,033	310,191.9	N/A	180,398.0	1.7 (N/A)
3	25,068,577	24,772,608	30,323	37,294	3,904	134,121.0	94,215.4	59,297.5	2.3 (1.6)
4	3,170,833	3,096,576	9,721	13,500	1,283	8,501.3	6,695.8	3,807.9	2.2 (1.8)

reach the target convergence criterion of 10^{-14}. The residual convergence target for grid 2 is set as 10^{-13}.

HANIM improves iterative convergence over the PA and legacy solvers in terms of nonlinear iterations and computational time, although the speedup factor is smaller compared to the hemisphere-cylinder case. HANIM attains the speedup factors of 2.2, 2.3, and 1.7 over the PA solver on grids 4, 3, and 2, respectively. HANIM shows less improvement on grid 2 than on coarser grid 3 because of more time spent on constructing GCR search directions. On grid 3, only four out of 3,904 nonlinear iterations used more than one GCR search direction. This HANIM behavior is a result of a relatively low maximum CFL limit imposed in this case. On grid 2, a total of 3,611 out of 9,033 nonlinear iterations used more than one GCR search direction; 333 nonlinear iterations used the maximum (four) search directions and did not meet the target GCR convergence criterion. Similarly, the HANIM speedup over the legacy solver appears to diminish on finer grids: a speedup factor of 1.8 is attained for grid 3 versus 1.6 for grid 2.

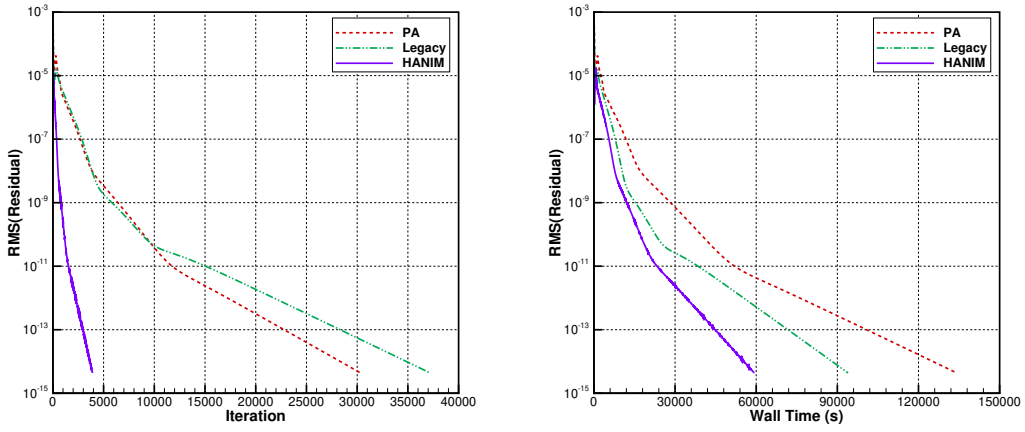
Convergence history of the unweighted RMS norm of combined meanflow and turbulence residuals on grid 3 is shown in Fig. 7 versus nonlinear iterations and computational time. To attain the residual level of 10^{-14}, the PA, legacy, and HANIM solvers use 30,323, 37,294, and 3,904 nonlinear iterations, respectively. After initial transients (after 30,000 seconds), HANIM residuals at any moment are at least an order of magnitude smaller than residuals of the PA and legacy solvers.

Figure 8 shows streamwise-velocity profiles along the vertical ($y = 0.5D$) and diagonal ($y = z$) cuts of the duct cross-section at two downstream locations, $x/D = 40$ and $x/D = 50$. The streamwise velocity is normalized by the centerline velocity, u_{CL}, obtained from the numerical solution at the specific x/D location. As expected, the PA, legacy, and HANIM solutions are identical; all of the solvers have converged to machine-zero residuals on the specific grid levels. The streamwise-velocity profiles are almost the same on all grids. These predicted velocity profiles are in good agreement with the experimental measurements.

Profiles of the skin friction coefficient along $y = 0$ at locations of $x = 40D$ and $x = 50D$ are plotted in Fig 9 and compared with the experimental measurements. The skin friction profiles correspond to PA and HANIM solutions on all of the three grid levels and legacy solutions on grids 3 and 4. The computed skin friction solutions from the PA, legacy, and HANIM solvers are identical. Below $2z/D < 0.7$, the variations of the computed skin friction solutions on different grids are negligibly small. In the region $0.7 \leq 2z/D \leq 1.0$, the level of skin friction decreases as the mesh resolution is increased. This behavior is more clearly observed in the $x/D = 50$ location (Fig. 9(b)). Moreover, the variation of skin friction between grids 2 and 3 significantly decreases compared to the variation between grids 3 and 4, approaching a grid-converged solution. The computed skin friction and the experiment are in good agreement overall; slight underprediction is seen in the vicinity of $2z/D = 1.0$.

C. Juncture Flow Model

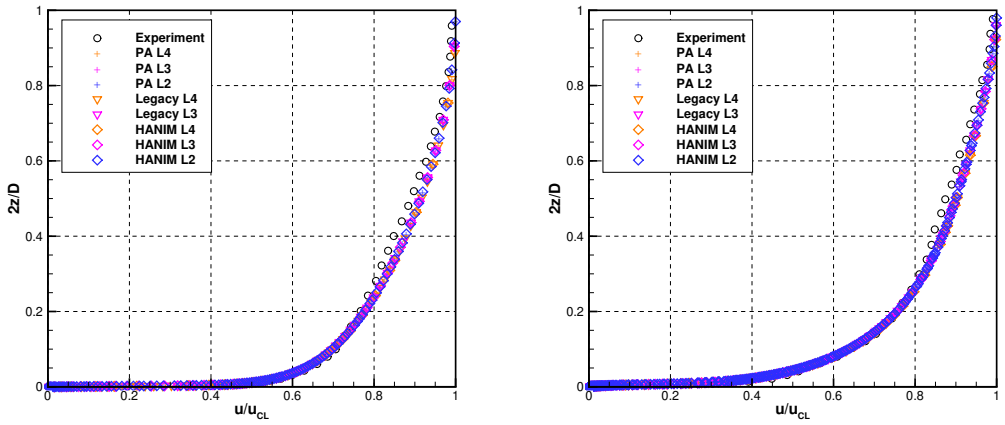
The next test case involves RANS simulations of flow around a NASA wing-fuselage juncture configuration. This flow has been extensively studied[13,45–48] to assess the capability of various turbulence models to predict juncture-flow separation. Reference 45 shows that nonlinear RANS models using quadratic constitutive relations provide improved separation predictions.



(a) Residual vs. iteration

(b) Residual vs. time

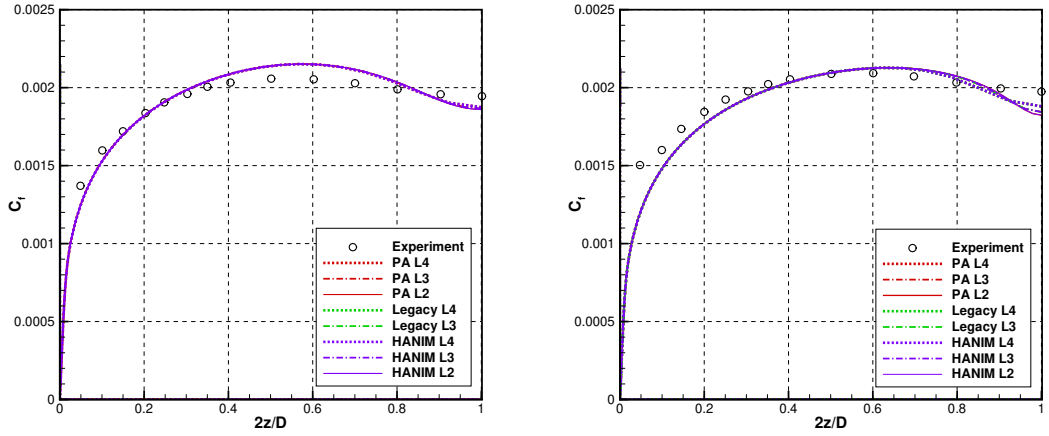
Figure 7. Convergence of PA, legacy, and HANIM solvers for the supersonic square-duct flow on grid 3.



(a) $x/D = 40$, vertical line (b) $x/D = 40$, diagonal line

(c) $x/D = 50$, vertical line (d) $x/D = 50$, diagonal line

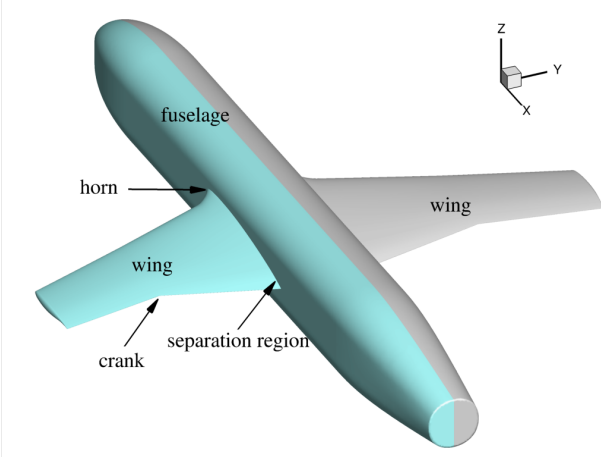
Figure 8. Comparison of nondimensional velocity profiles on vertical and diagonal cuts at locations of $x/D = 40$ and $x/D = 50$ for the supersonic square-duct flow.



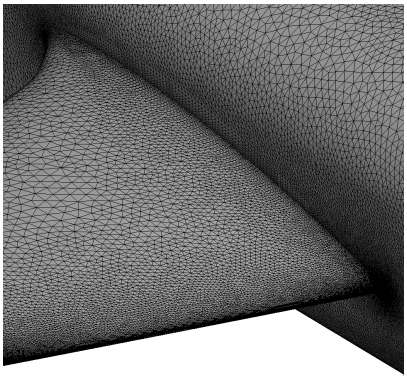
(a) $x/D = 40$ (b) $x/D = 50$

Figure 9. Comparison of the skin friction coefficients at $x/D = 40$ and $x/D = 50$ for the supersonic square-duct flow.

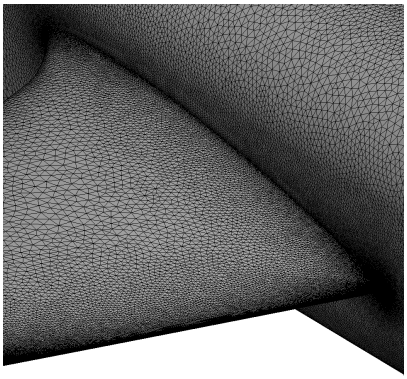
The flow conditions for this test example are a freestream Mach number of 0.189, a Reynolds number of 2,400,000 based on the crank chord of 557.17mm, a reference temperature of 288.84 Kelvin, and an angle of attack of 5 degrees. The full juncture-flow model is shown in Fig. 10(a), where the highlighted portion corresponds to the half model considered in the present study. A symmetry boundary condition is applied at the $y = 0$ plane. A farfield freestream



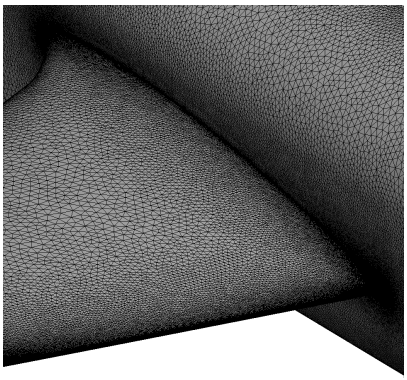
(a) Configuration



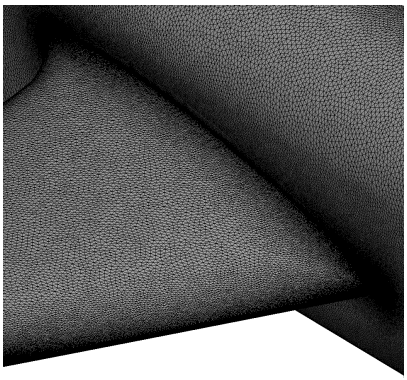
(b) Grid E (4.8 million nodes)



(c) Grid I (13 million nodes)



(d) Grid P (37 million nodes)



(e) Grid V (106 million nodes)

Figure 10. NASA juncture flow model configuration (half-span computational model highlighted in blue) and unstructured grids near wing-root junction region.

boundary condition is applied at the outer boundary. A no-slip adiabatic-wall boundary condition is applied at the fuselage and wing. A family of four unstructured, mixed-element grids have been generated. The coarsest grid in the family has 4.4 million grid points and 10.8 million cells. The finest grid in the family has about 106 million grid points and 420 million grid cells. Table 3 shows the grid and iterative convergence metrics corresponding to the PA, legacy, and HANIM solvers. The target residual level is set as 1.5×10^{-10}. Figures 10(b)–10(e) display closeup views of the surface grids in the vicinity of the wing-fuselage junction region of interest in which flow separation occurs. It is seen that the grid spacing is refined near the wing leading and trailing edges as well as the wing root areas.

The PA, legacy, and HANIM solvers are used to compute solutions on the four grids. The turbulent stresses are

computed based on SA-neg-QCR2000 turbulence model.[35] The HANIM setup includes the following parameters. The maximum number of multicolor sweeps in the preconditioner is set to 100 for the meanflow and turbulence equations. The preconditioner residual convergence target is 0.5. The maximum number of search directions is set to 4 (HANIM-4). The GCR residual reduction target is set to 0.92. The residual weighting based on estimated element sizes (Eq.(25)) is selected. The maximum CFL number is set to 1,000,000 for grids V, P, and I, while a much lower value of 100 is used for the coarsest grid E to prevent an occurrence of a limit cycle. The cause of the limit cycle observed with a higher CFL bound is not understood and needs to be investigated.

The computations on grids V and P were performed on the dual-socket 14-core Intel Broadwell processors of the NAS Pleiades system with 2,100 and 756 processing cores, respectively. The computations on grids I and E were performed on the dual-socket 20-core Intel Gold 6148 Skylake processors of the NASA LaRC K-Cluster with 320 and 160 processing cores, respectively. HANIM achieved the unweighted residual level of 1.5×10^{-10} or below on all of the grids. Table 3 shows the wall-clock time in seconds that was required by each solver to achieve this level of convergence. At the time of writing this paper, the PA and legacy solvers have not converged residuals to the HANIM level on the finest grid V. For these solutions, however, at least eight orders of residual reduction have been achieved. The computational times for the incomplete PA and legacy solver solutions are estimated through linear extrapolation. These projected values are marked with superscript asterisk (*) in Table 3 for distinction.

Overall, the performance of HANIM for this case is superior. Except for grid E where the maximum CFL number was limited, HANIM provides a speedup factor of 6.8–7.8 over the PA solver and a speedup factor of 4.1–5.6 over the FUN3D legacy solver.

Convergence of the unweighted RMS norm for the combined meanflow and turbulence-model residuals is shown in Fig. 11 versus nonlinear iterations and wall-clock time (seconds), as well as the CFL history. For grid E, all three solvers converge residuals to the target level. The HANIM CFL number increases to the maximum value in eight nonlinear iterations and remains at this level throughout the solution process (Fig. 11(c)). Because the HANIM CFL number is marginally higher than that of the PA and legacy solvers, only a moderate speedup is obtained: HANIM is about 36% faster than the PA solver and 17% faster than the legacy solver.

The legacy solver requires more nonlinear iterations but takes less time than the PA solver to reach the same residual level. The PA and legacy solvers show relatively fast initial convergence on grids P and I, but the convergence rate decreases substantially afterward. In contrast, HANIM achieves significantly greater residual reductions than the PA and legacy solvers. When HANIM achieves the target convergence level, its residual is much smaller than the residuals of the other solvers: five orders of magnitude smaller on grid I, three orders of magnitude smaller on grid P, and six orders of magnitude smaller on grid V. As shown in Figs. 11(f), 11(i), and 11(l), the operational CFL number in HANIM, which is in the range of 5,000 to 100,000, is 2–3 orders of magnitude greater than the CFL number used in the PA and legacy solvers.

Figure 12 shows streamline visualizations in the wing-fuselage junction region and streamwise skin-friction contours corresponding to the converged HANIM solutions on various grids. The size of the corner-flow separation near the wing trailing edge reduces consistently with increasing grid resolution. Table 4 provides quantitative characteristics of the separation zone, including the x-coordinate of the location of the separation onset, x-length of the separation vortex, and the predicted length of the separation bubble. The x-location is determined by detecting the first location where the streamwise skin friction becomes negative; the skin friction is observed in the wing-fuselage junction plane at $y = -236.1$mm. The x-length of the separation vortex is determined by the separation onset x-location subtracted from the wing-root trailing edge location, $x = 2961.9$mm. The bubble length at the fuselage along the wing is the x-length divided by $\cos(17.5^\circ)$ to account for the angle between the wing surface and the x-axis. The initial separation location appears to move downstream as the grid resolution increases.

Figure 13 shows convergence of lift and drag coefficients as a function of mesh resolution as well as conver-

Table 3. Grid and convergence metrics in the juncture flow model simulations for tolerance of 1.5×10^{-10}. Superscript * indicates extrapolated values.

Grid	Nodes	Ncells	Wall-clock seconds			Speedup over
			PA	Legacy	HANIM	PA (Legacy)
V	106,090,663	420,167,293	1,168,727.7 *	960,027.5*	172,487.7	6.8 (5.6)
P	37,131,844	113,428,664	1,525,525.6	932,174.0	195,346.3	7.8 (4.8)
I	13,036,210	37,126,819	453,359.3	274,738.9	66,795.5	6.8 (4.1)
E	4,784,640	13,480,180	142,648.0	109,066.4	91,006.8	1.6 (1.2)

Table 4. HANIM predictions of separation start x-location, bubble length in x-axis (x-length), and bubble length along the wing surface at $y = -236.1$mm. Experiment measurement of the start x-location based on oil flow visualization corresponds to 2851mm.

Grid	x-location (mm)	x-length (mm)	Bubble length (mm)
E	2794.23	167.67	175.81
I	2800.98	160.92	168.73
P	2801.35	160.55	168.34
V	2805.34	156.56	164.16

gence of their pressure- and viscous-force components. Monotonic convergence is observed for both forces and their components in grid refinement. The variations in the computed forces between the last two finer grids are seen to diminish. Although the integrated lift and drag coefficients are not available in the experiment, prior CFD simulations have shown a lift coefficient of 0.85–0.86 and a drag coefficient of 0.067–0.070,[47] which agrees with the present predictions. Moreover, the PA, legacy, and HANIM solutions are shown to be identical on each grid level.

Comparisons of surface pressure coefficient of the legacy and HANIM solutions and experimental measurements[49] are shown in Fig. 14 for various spanwise and streamwise stations on the wing and fuselage. The experimental results are plotted together with uncertainty bounds indicating the range of data from multiple experimental tests. The pressure distributions from the legacy and HANIM solutions are identical on each grid level; this is the same case as compared to the PA solutions (not shown). The present FUN3D solutions agree well with the experiment. As the mesh resolution increases, the agreement between the CFD solutions and experiments for surface pressure in the flow separation region is improved as indicated in Figs 14(b) and 14(c). Grid resolution also affects the surface pressure prediction at the wing tip near the trailing edge. Higher grid resolution improves the wing-tip vortex capturing, leading to a higher peak of wingtip loading (Fig. 14(h)).

Figure 15 depicts comparisons of the experimental LDV (Laser Doppler Velocimetry) measurements and computed velocities and Reynolds stresses along the y- and z-directions at various streamwise locations. The velocity profiles in the regions upstream of the separation bubble agree well with the experimental data. At these locations, an increase in the grid resolution results in relatively small variations in the velocity and Reynolds stresses. In contrast, much larger variations are observed inside the separation bubble as the grid is refined, see Figs. 15(g)–15(i) for the location at $x = 2852.6$mm (1mm away from the fuselage). Underpredictions and overpredictions of Reynolds stresses are observed with the exception of the computed shear stress $v'w'$, which has reasonable agreement with the experimental data. At the same streamwise location ($x = 2852.6$mm), 30mm away from the fuselage (Figs. 15(j)–15(l)), improved agreement of velocities and Reynolds stresses is observed between the computations and experiments. The legacy and HANIM solvers provide identical solutions on all grid levels: the velocity and Reynolds stress profiles obtained by the two solvers overplot each other. Although not shown, the computed velocity and Reynold stresses by the PA solver are also identical to those of the legacy and HANIM solutions.

IV. Conclusions

A hierarchical adaptive nonlinear iteration method (HANIM) has been implemented in a NASA computational fluid dynamics (CFD) solver, FUN3D. HANIM is based on a hierarchy of modules consisting of preconditioner, generalized conjugate residual, realizability check, nonlinear control, and CFL adaptation modules. This paper focuses on the HANIM performance for steady-state Reynolds-averaged Navier-Stokes equations. Two one-equation Spalart-Allmaras (SA) turbulence models are considered: one is the negative variant of the linear SA model; the other is a nonlinear SA model based on quadratic constitutive relations.

Numerical solutions have been computed for three benchmark flows including a subsonic separated flow around a hemisphere-cylinder configuration, a supersonic flow through a long square duct, and a subsonic flow over the NASA wing-fuselage juncture configuration. HANIM performance for various grid resolutions has been systematically compared with performance of a preconditioner-alone (PA) solver and the legacy iterative solver of FUN3D. The PA solver closely resembles the legacy solver but employs consistent inviscid-flux linearization and evaluates the Jacobian matrix in each nonlinear iteration.

In all of the test examples, HANIM improves computational efficiency of FUN3D solutions. The HANIM CFL number is adapted based on performance of individual modules and is often several orders of magnitude higher than that of the legacy FUN3D solver. For the hemisphere-cylinder and supersonic square-duct flow cases, HANIM provides a moderate speedup factor of 1.7–2.6 over the PA solver, and a speedup factor of 1.6–2.4 over the legacy solver.

In the NASA juncture flow model simulations, HANIM shows superior performance: a speedup factor of over 7.8 is shown over the PA solver and a speedup factor of 4.1–5.6 is achieved over the legacy solver of FUN3D. Various aerodynamic quantitates of interest computed from FUN3D solutions compare well with experimental measurements for all three benchmark examples.

Acknowledgments

The work is supported by the NASA Revolutionary Vertical Lift Technology Project, NASA Revolutionary Computational Aerosciences Project, and NASA Langley High Performance Computing Incubator Program. The support is greatly appreciated. Computations were performed on the NASA LaRC K-Cluster and NASA Advanced Supercomputing Pleiades system. The authors would also like to thank Dr. Steve L. Karman from Pointwise Inc. for generating the juncture-flow-model grids.

References

[1]Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., and Mavriplis, D., "CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences," Tech. Rep. NASA/CR-2014-218178, NASA Langley Research Center, Hampton, Virginia, 2014.

[2]Mavriplis, D. J., "An Assessment of Linear Versus Nonlinear Multigrid Methods for Unstructured Mesh Solvers," Tech. Rep. ICASE Report 2001-12, Institute for Computer Applications in Science and Engineering (ICASE), May 2001.

[3]Wang, L. and Mavriplis, D. J., "Implicit Solution of the Unsteady Euler Equations for High-order Accurate Discontinuous Galerkin Discretizations," *Journal of Computational Physics*, Vol. 225, No. 2, 2007, pp. 1994–2015, 10.1016/j.jcp.2007.03.002.

[4]Anderson, K., Newman, J., and Karman, S., "Stabilized Finite Elements in FUN3D," *Journal of Aircraft*, Vol. 55, No. 2, March–April 2018, pp. 696–714, 10.2514/1.C034482.

[5]Ceze, M. and Fidkowski, K. J., "A Robust Adaptive Solution Strategy for High-Order Implicit CFD Solvers," *Proceedings of 20th AIAA Computational Fluid Dynamics Conference*, June 27–30, 2012, AIAA Paper 2011–3696, Honolulu, Hawaii.

[6]Mavriplis, D. J., "Residual Smoothing Strategy for Accelerating Newton Method Continuation," *Computing Research Repository*, 2018, arXiv:1806.01437.

[7]Anderson, W. K., Rausch, R. D., and Bonhaus, D. L., "Implicit/Multigrid Algorithms for Incompressible Turbulent Flows on Unstructured Grids," *Journal of Computational Physics*, Vol. 128, No. 2, 1996, pp. 391–408, 10.1006/jcph.1996.0219.

[8]Saad, Y. and Schultz, M. H., "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1996, pp. 856–869, 10.1137/0907058.

[9]Knoll, D. A. and Keyes, D. E., "Jacobian-free Newton−Krylov Methods: A Survey of Approaches and Applications," *Journal of Computational Physics*, Vol. 193, No. 2, 2004, pp. 357–397, 10.1016/j.jcp.2003.08.010.

[10]Saad, Y., *Iterative Methods for Sparse Linear Systems*, 2nd ed., Society for Industrial and Applied Mathematics, 2003.

[11]Shapira, Y., *"Algebraic Multigrid". Matrix-Based Multigrid: Theory and Applications*, Springer, 2003, ISBN 978-1-4020-7485-1.

[12]Trottenberg, U., Oosterlee, C. W., and Schüller, A., *Multigrid*, Academic Press, 2001, ISBN 978-0-12-701070-0.

[13]Abdol-Hamid, K. S., Ahmad, N. N., Carlson, J.-R., and Biedron, R., "Juncture Flow Computations using kL-Based Turbulence Models," *Proceedings of the AIAA SciTech Forum*, January 6–10, 2020, Orlando, Florida, AIAA Paper 2020–1305, 10.2514/6.2020-1305.

[14]Wang, L., Diskin, B., Biedron, R. T., Nielsen, E. J., Sonneville, V., and Bauchau, O. A., "High-Fidelity Multidisciplinary Design Optimization Methodology with Application to Rotor Blades," *Journal of the American Helicopter Society*, Vol. 64, No. 3, July 2019, pp. 1–11, 10.4050/JAHS.64.032002.

[15]Wang, L., Diskin, B., Lopes, L. V., Nielsen, E. J., Lee-Rausch, E., and Biedron, R. T., "High-Fidelity Multidisciplinary Design Optimization of Low-Noise Rotorcraft," *Proceedings of VFS 75th Annual Forum*, May 13–16 2019, Philadelphia, PA.

[16]Walden, A., Nielsen, E., Diskin, B., and Zubair, M., "A Mixed Precision Multicolor Point-Implicit Solver for Unstructured Grids on GPUs," *2019 IEEE/ACM 9th Workshop on Irregular Applications: Architectures and Algorithms (IA3)*, 2019, pp. 23–30.

[17]Jain, R., "Sensitivity Study of High-Fidelity Hover Predictions on the Sikorsky S-76 Rotor,," *Journal of Aircraft*, Vol. 55, No. 1, January–February 2018, pp. 78–88, 10.2514/1.C034076.

[18]Nielsen, E. J. and Diskin, B., "High-Performance Aerodynamic Computations for Aerospace Applications," *Parallel Computing*, Vol. 64, 2017, pp. 20–32, 10.1016/j.parco.2017.02.004.

[19]Lee-Rausch, E. M., Rumsey, C. L., and Park, M. A., "Grid-Adapted FUN3D Computations for the Second High-Lift Prediction Workshop," *Journal of Aircraft*, Vol. 52, No. 4, 2015, pp. 1098–1111, 10.2514/1.C033192.

[20]Rumsey, C., Slotnick, J., Long, M., Stuever, R., and Wayman, T., "Summary of the First AIAA CFD High-Lift Prediction Workshop," *Journal of Aircraft*, Vol. 48, No. 6, May 2011, pp. 2068–2079, 10.2514/1.C031447.

[21]Biedron, R. T. and Lee-Rausch, E. M., "Rotor Airloads Prediction Using Unstructured Meshes and Loose CFD/CSD Coupling," *Proceedings of the 26th AIAA Applied Aerodynamic Conference, AIAA Paper 2008–7341*, August 18–21, 2008, Honolulu, Hawaii.

[22]Londenberg, W. K. and O. John Ohanian, I., "Application of CFD in the Design of Flow Control Concepts for a Ducted-Fan Configuration," *Proceedings of the International Powered Lift Conference*, October 5–7, 2010.

[23]Bartels, R., Scott, R., Funk, C., Allen, T., and Sexton, B., "Computed and Experimental Flutter/LCO Onset for the Boeing Truss-Braced Wing Wind-Tunnel Model," *Proceedings of 44th AIAA Fluid Dynamics Conference*, June 16–20 2014, AIAA Paper 2014-2446, Atlanta, Georgia.

[24]Bozeman, M., Tai, J., Robertson, B., and Mavris, D., "Prediction of Dynamic Stability Derivatives for Flexible Aircraft using FUN3D," *Proceedings of the AIAA SciTech Forum*, January 8–12 2018, AIAA Paper 2018-2063, Kissimmee, Florida.

[25]Edquist, K., Korzun, A., Bibb, K., Schauerhamer, G., Ma, E., McCloud, P., Palmer, G., and Monk, J., "Comparison of Navier-Stokes Flow Solvers to Falcon 9 Supersonic Retropropulsion Flight Data," *Proceedings of the AIAA SPACE Forum*, September 12–14 2017, AIAA Paper 2017-5296, Orlando, Florida.

[26]Diskin, B., Anderson, K., Pandya, M., Rumsey, C., Thomas, J., Liu, Y., and Nishikawa, H., "Grid Convergence for Three Dimensional Benchmark Turbulent Flows," *Proceedings of the AIAA SciTech Forum*, January 8–12 2018, AIAA Paper 2018-1102, Kissimmee, Florida.

[27]Diskin, B. and Nishikawa, H., "Evaluation of Multigrid Solutions for Turbulent Flows," *Proceedings of the AIAA SciTech Forum*, January 13–17 2014, AIAA Paper 2014-0082, National Harbor, Maryland.

[28]Diskin, B. and Thomas, J., "Comparison of Node-Centered and Cell-Centered Unstructured Finite Volume Discretizations: Inviscid Fluxes," *AIAA Journal*, Vol. 49, No. 4, April 2011, pp. 836–854, 10.2514/1.J050897.

[29]Diskin, B., Thomas, J., Nielsen, E., Nishikawa, H., and White, J., "Comparison of Node-Centered and Cell-Centered Unstructured Finite-Volume Discretizations Part I: Viscous Fluxes," *AIAA Journal*, Vol. 48, No. 7, July 2010, 10.2514/1.44940.

[30]Pandya, M. J., Jespersen, D. C., Diskin, B., Thomas, J. L., and Frink, N. T., "Accuracy, Scalability, and Efficiency of Mixed-Element USM3D for Benchmark Three-Dimensional Flows," *Proceedings of the AIAA SciTech 2019 Forum*, January 7–11, 2019, AIAA Paper 2019–2333, San Diego, California.

[31]Pandya, M., Diskin, B., Thomas, J., and Frink, N., "Assessment of USM3D Hierarchical Adaptive Nonlinear Method Preconditioners for Three-Dimensional Cases," *AIAA Journal*, Vol. 55, No. 10, October 2017, pp. 1–16, 10.2514/1.J055823.

[32]Pandya, M. J., Diskin, B., Thomas, J. L., and Frink, N. T., "Improved Convergence and Robustness of USM3D Solutions on Mix-Element Grids," *AIAA Journal*, Vol. 54, No. 9, 2016, pp. 2589–2596, 10.2514/1.J054545.

[33]Spalart, P. and Allmaras, S., "A One-Equation Turbulence Model for Aerodynamic Flows," *Le Recherche Aerospatiale*, Vol. 1, 1994, pp. 5–21, 10.2514/6.1992-439.

[34]Allmaras, S. R., Johnson, F. T., and Spalart, P. R., "Modifications and Clarifications for the Implementation of the Spalart-Allmaras Turbulence Model," *Proceedings of 7th International Conference on Computational Fluid Dynamics*, July 9–13 2012, ICCFD7–1902, BigIsland, Hawaii.

[35]Spalart, P., "Strategies for Turbulence Modelling and Simulation," *International Journal of Heat and Fluid Flow*, Vol. 21, No. 3, 2000, pp. 252–263, 10.1016/S0142-727X(00)00007-2.

[36]Anderson, W. K. and Bonhaus, D. L., "An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids," *Computers & Fluids*, Vol. 23, No. 1, 1994, pp. 1–21, 10.1016/0045-7930(94)90023-X.

[37]Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Gnoffo, P. A., Hammond, D. P., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Nielsen, E. J., Park, M. A., Rumsey, C. L., Thomas, J. L., and Wood, W. A., "FUN3D Manual: 13.6," NASA/TM-2019–220416, NASA Langley Research Center, Hampton, Virginia, 2019.

[38]Roe, P. L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," *Journal of Computational Physics*, Vol. 43, No. 2, 1981, pp. 357–372, 10.1016/0021-9991(81)90128-5.

[39]Burg, C. O. E., "Higher Order Variable Extrapolation For Unstructured Finite Volume RANS Flow Solvers," *Proceedings of the 17th AIAA Computational Fluid Dynamics Conference*, June 6–9 2005, AIAA Paper 2005–4999, Toronto, Ontario Canada.

[40]van Leer, B., "Towards the Ultimate Conservative Difference Scheme, V. A Second Order Sequel to Godunov's Method," *Journal of Computational Physics*, Vol. 32, No. 1, 1979, pp. 101–136, 10.1016/0021-9991(79)90145-1.

[41]Thomas, J. L., Diskin, B., and Nishikawa, H., "A Critical Study of Agglomerated Multigrid Methods for Diffusion on Highly-Stretched Grids," *Computers & Fluids*, Vol. 41, No. 1, February 2011, pp. 82–93, 10.1016/j.compfluid.2010.09.023.

[42]Carlson, J.-R., "Inflow/Outflow Boundary Conditions with Application to FUN3D," October 2011, NASA/TM–2011-217181, NASA Langley Research Center, Hampton, Virginia.

[43]Tsieh, T., "An Investigation of Separated Flow about A Hemisphere-Cylinder at 0- to 19-Deg Incidence In the Mach Number Range from 0.6 to 1.5," AEDC-TR-76-112, Propulsion Wind Tunnel Facility, Arnold Engineering Development Center, Air force Systems Command, Arnold Air Force Station, Tennessee, 1976.

[44]Pandya, M. J., Diskin, B., Thomas, J. L., and Frink, N. T., "Assessment of Preconditioner for a USM3D Hierarchical Adaptive Nonlinear Method (HANIM)," *Proceedings of the 54th AIAA Aerospace Sciences Meeting, AIAA Paper 2016–0860*, January 4–8, 2016, San Diego, California.

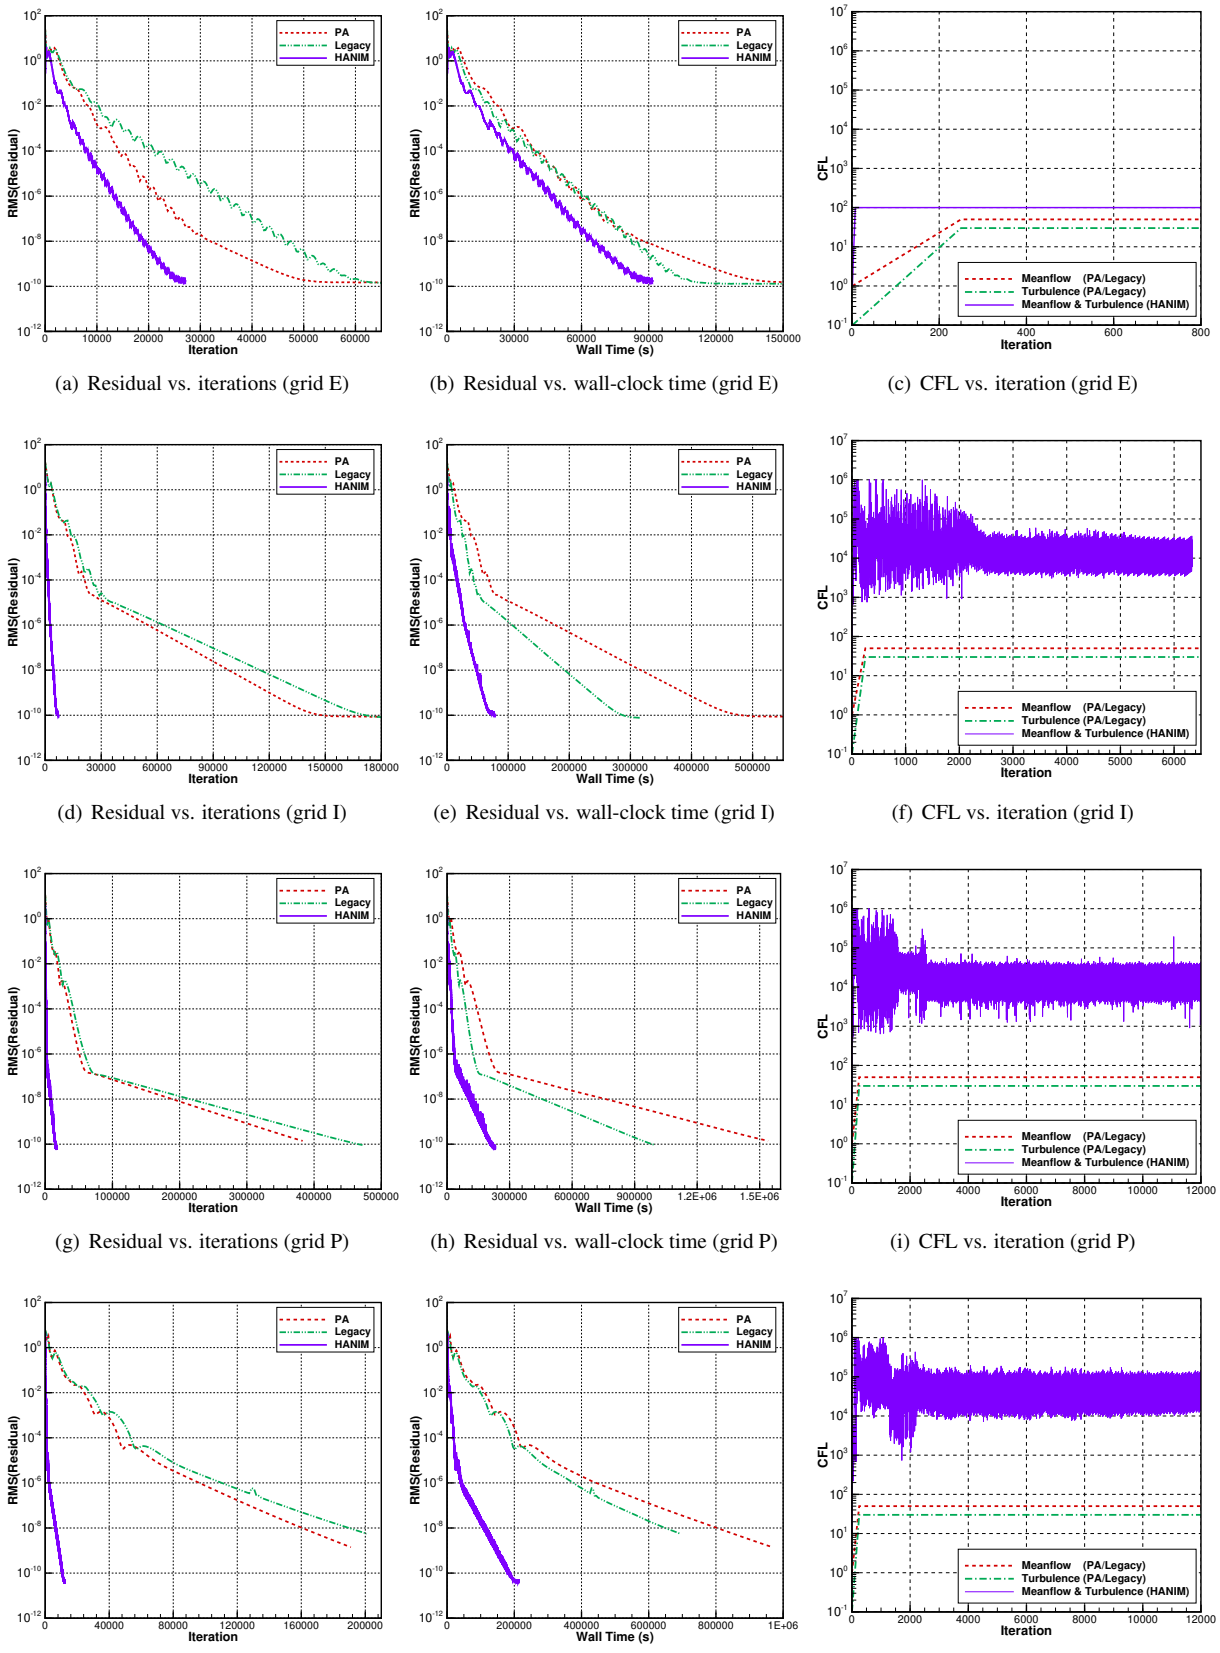
[45]Rumsey, C. L., Carlson, J.-R., and Ahmad, N. N., "FUN3D Juncture Flow Computations Compared with Experimental Data," *Proceedings of the AIAA SciTech Forum*, January 7–11 2019.

[46]Kegerise, M. A. and Neuhart, D. H., "An Experimental Investigation of a Wing-Fuselage Junction Model in the NASA Langley 14- by 22-Foot Subsonic Tunnel," NASA/TM-2019–220286, NASA Langley Research Center, June 2019, Hampton, Virginia.

[47]Rumsey, C. L., Lee, H. C., and Pulliam, T. H., "Reynolds-Averaged Navier-Stokes Computations of the NASA Juncture Flow Model Using FUN3D and OVERFLOW," *Proceedings of the AIAA SciTech Forum*, January 6–10, 2020, AIAA Paper 2020–1304, Orlando, Florida.

[48]Iyer, P. S. and Malik, M. R., "Wall-Modeled LES of the NASA Juncture Flow Experiment," *Proceedings of the AIAA SciTech Forum*, January 6–10, 2020, AIAA Paper 2020–1307, Orlando, Florida.

[49]Rumsey, C. L., Ahmad, N. N., Carlson, J.-R., Kegerise, M. A., Neuhart, D. H., Hannon, J. A., Jenkins, L. N., Yao, C.-S., Balakumar, P., Bartram, S. M., Pulliam, T. H., Olsen, M. E., and Spalart, P. R., "CFD Comparisons with Updated NASA Juncture Flow Data," *Proceedings of the AIAA SciTech Forum*, January 11-21, 2021.



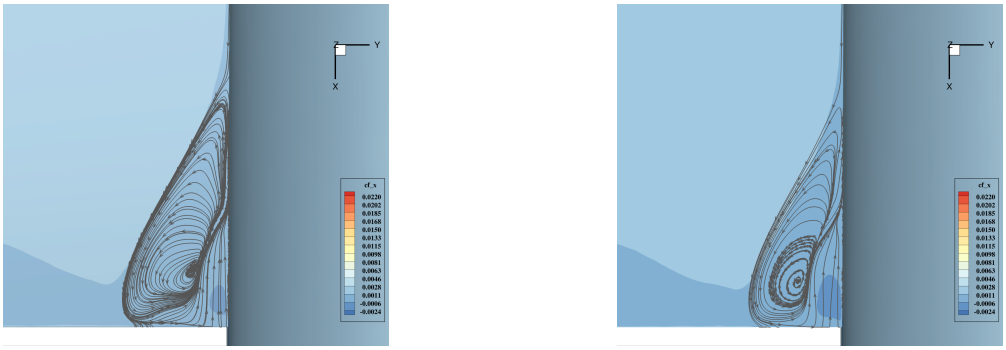
(a) Residual vs. iterations (grid E) (b) Residual vs. wall-clock time (grid E) (c) CFL vs. iteration (grid E)

(d) Residual vs. iterations (grid I) (e) Residual vs. wall-clock time (grid I) (f) CFL vs. iteration (grid I)

(g) Residual vs. iterations (grid P) (h) Residual vs. wall-clock time (grid P) (i) CFL vs. iteration (grid P)

(j) Residual vs. iterations (grid V) (k) Residual vs. wall-clock time (grid V) (l) CFL vs. iteration (grid V)

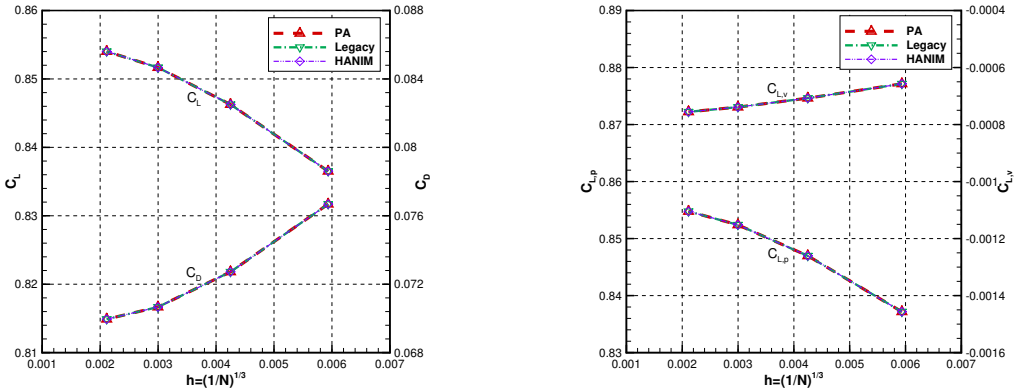
Figure 11. Convergence comparisons for PA, legacy, and HANIM solvers in the juncture-flow model simulations.



(a) Grid E (b) Grid I

(c) Grid P (d) Grid V

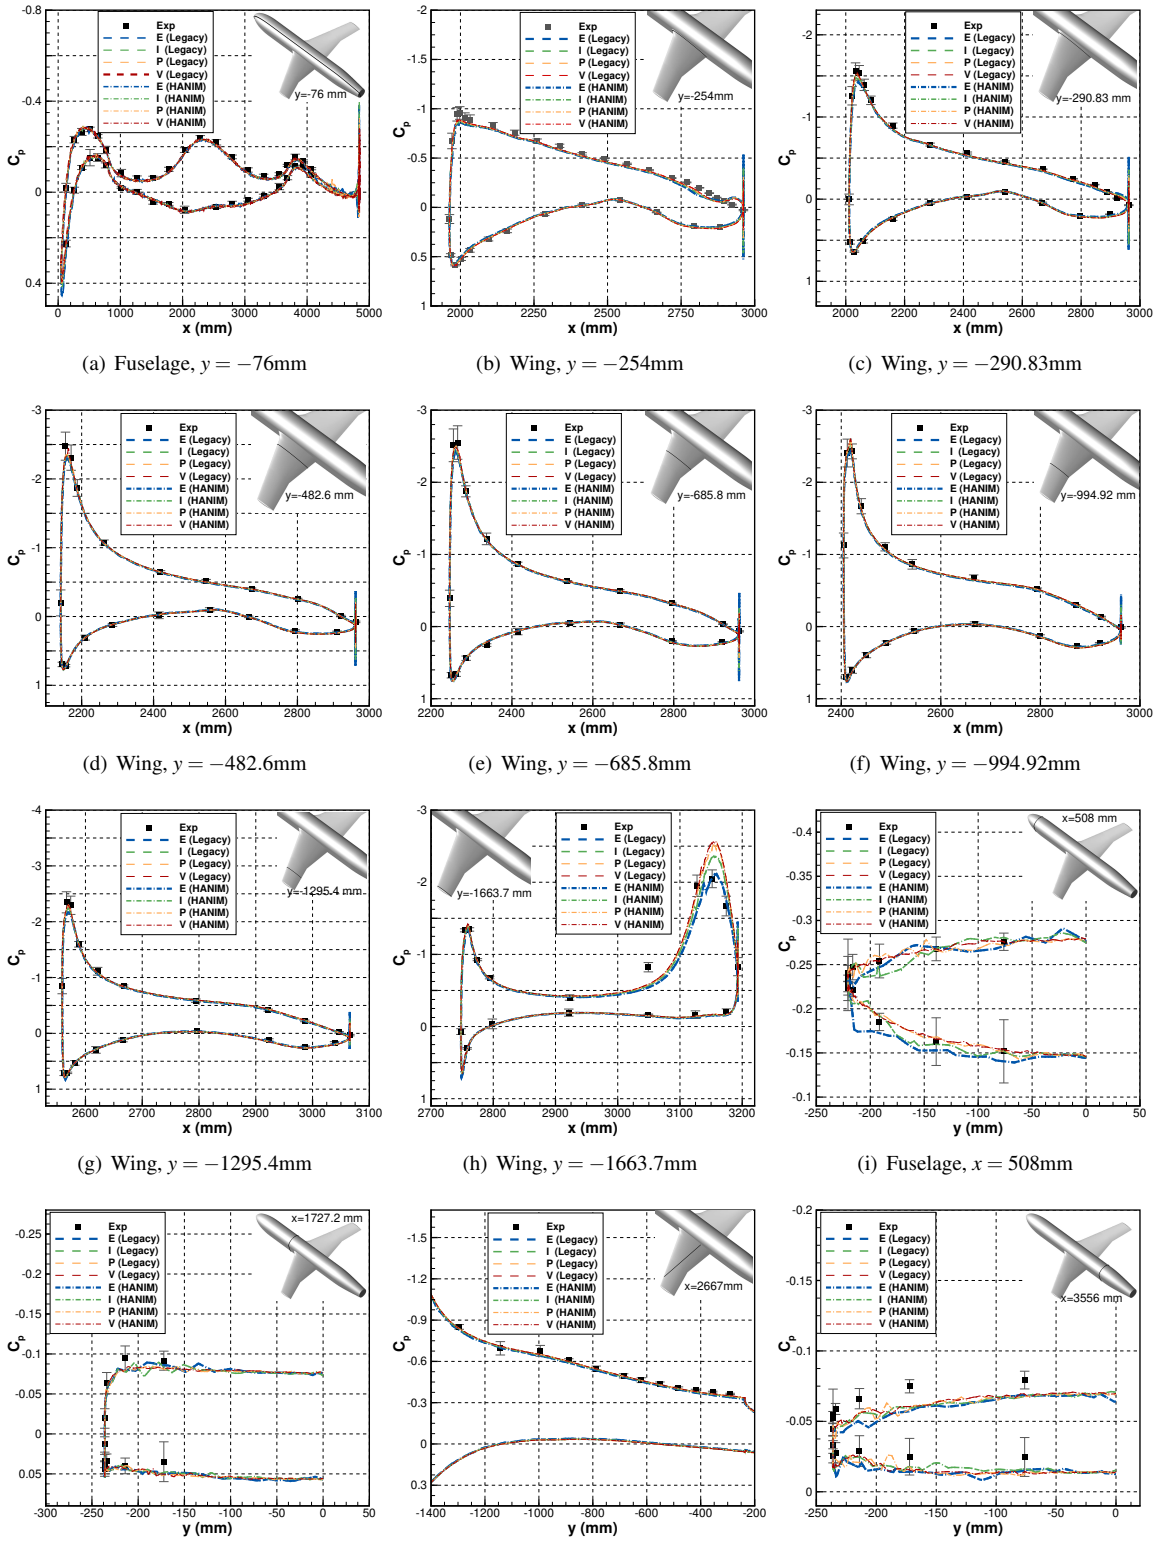
Figure 12. Streamlines near the corner-flow region colored by the x-component of skin friction. Converged solutions computed by HANIM.



(a) Lift and drag coefficients (b) Lift components

(c) Drag components

Figure 13. Convergence of forces and pressure and viscous components of forces.



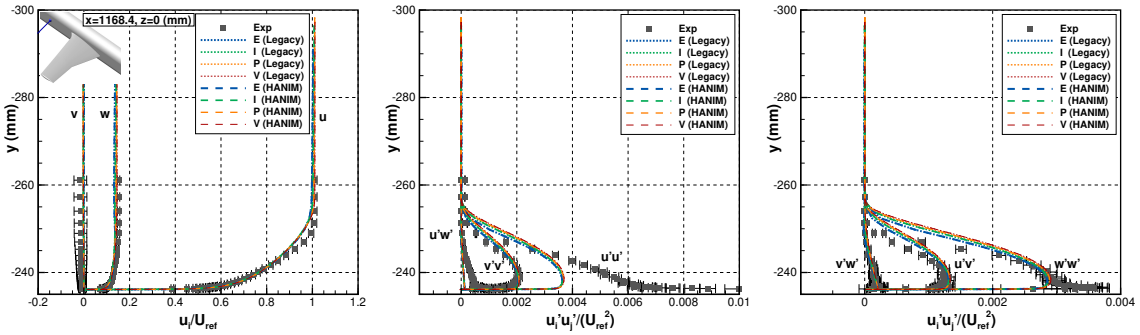
(a) Fuselage, $y = -76$mm (b) Wing, $y = -254$mm (c) Wing, $y = -290.83$mm

(d) Wing, $y = -482.6$mm (e) Wing, $y = -685.8$mm (f) Wing, $y = -994.92$mm

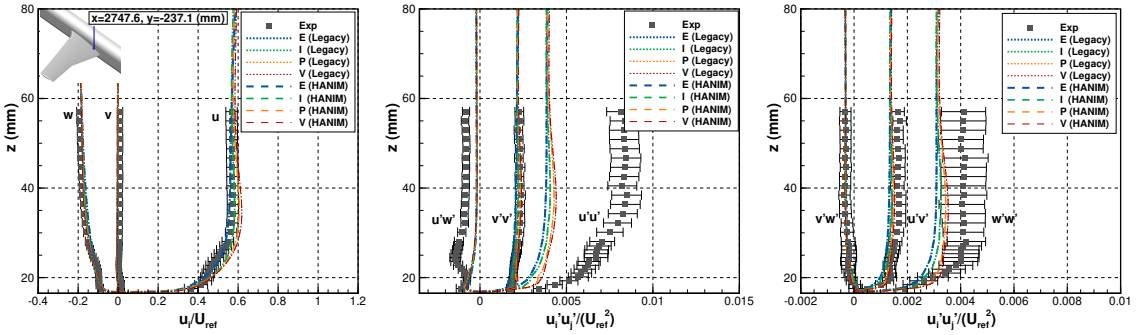
(g) Wing, $y = -1295.4$mm (h) Wing, $y = -1663.7$mm (i) Fuselage, $x = 508$mm

(j) Fuselage, $x = 1727.2$mm (k) Wing, $x = 2667$mm (l) Fuselage, $x = 3556$mm

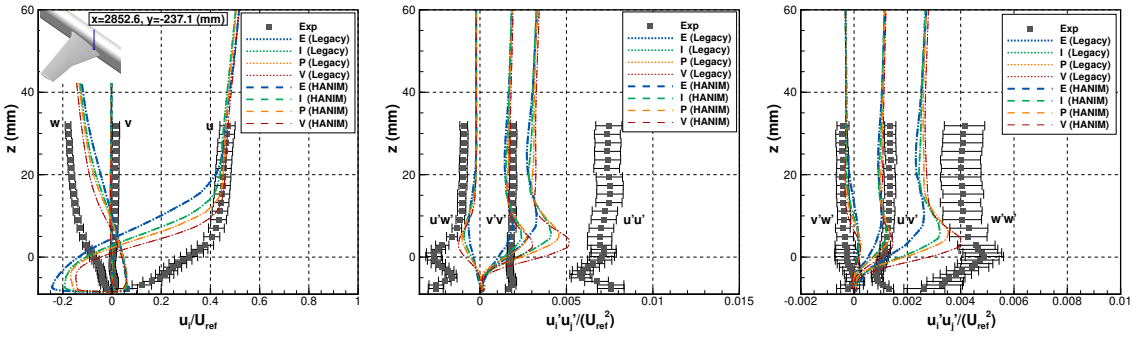
Figure 14. Comparisons of surface pressure at various stations on wing and fuselage of juncture flow model.



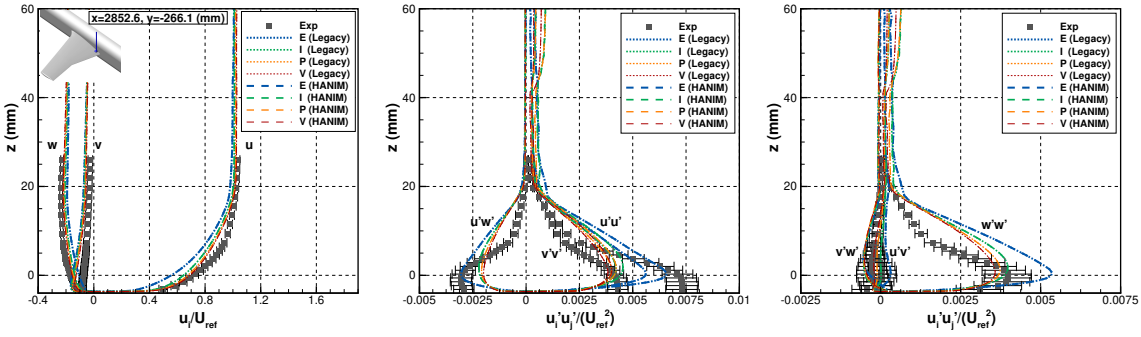
(a) u_i, $(x,z) = (1168.4, 0)$, mm (b) $u_i' u_j'$, $(x,z) = (1168.4, 0)$, mm (c) $u_i' u_j'$, $(x,z) = (1168.4, 0)$, mm



(d) u_i, $(x,y) = (2747.6, -237.1)$, mm (e) $u_i' u_j'$, $(x,y) = (2747.6, -237.1)$, mm (f) $u_i' u_j'$, $(x,y) = (2747.6, -237.1)$, mm



(g) u_i, $(x,y) = (2852.6, -237.1)$, mm (h) $u_i' u_j'$, $(x,y) = (2852.6, -237.1)$, mm (i) $u_i' u_j'$, $(x,y) = (2852.6, -237.1)$, mm



(j) u_i, $(x,y) = (2852.6, -266.1)$, mm (k) $u_i' u_j'$, $(x,y) = (2852.6, -266.1)$, mm (l) $u_i' u_j'$, $(x,y) = (2852.6, -266.1)$, mm

Figure 15. Comparisons of velocity and Reynolds stress of juncture flow model.