Pretest Simulations of a Supersonic Mixing and Combustion Validation Experiment to Assess Sensitivities

Robert A. Baurle
NASA Langley Research Center, Hampton, VA

Brian Shi
North Carolina State University, Raleigh, NC

Joseph Johnson
University of Florida, Gainesville, FL

David M. Peterson
Air Force Research Laboratory, Wright-Patterson AFB, OH

Cannon J. DeBardelaben
Innovative Scientific Solutions Inc., Dayton, OH

ABSTRACT

The reliance on CFD simulations to develop, design, and optimize scramjet systems (or components) has become commonplace. This reliance inevitably hinges on the ability of the computational analyst to quantify the level of confidence in their computational results. Unfortunately, nearly all measured data available for this assessment comes from antiquated experimental datasets, or from tests that focused on the extraction of scramjet system (or component) performance. The objective of a CFD validation experiment is to quantify the predictive accuracy of one or more of the CFD physics submodels, implying that other uncertainties related to replicating the facility flow environment (e.g., knowledge of boundary conditions) must be minimized to the extent possible. This inevitably places stringent requirements on the quality and quantity of measurements taken to accurately specify inflow, outflow, and surface conditions for the CFD simulations, in addition to the measurements taken for the validation of physics submodels. This places additional demands on the experimental process above and beyond those for test article performance assessment. A recent high-speed code credibility workshop series sponsored by AFRL identified a gap in existing validation data for fundamental assessments of turbulent mixing and combustion CFD closure models at scramjet engine relevant conditions. To address this gap, engineers at AFRL have designed a coaxial jet flame configuration that will be tested at two facilities (Research Cell 19 at the Air Force Research Lab, and at Purdue University). The effort described here documents pretest simulations of this validation experiment with the goal of fleshing out the extent of the facility flowpath that must be included to adequately reproduce the facility test section flow environment. The findings indicate that the flow around the support structure for the fuel injection centerbody upstream of the facility nozzle generates disturbances that persist throughout the nozzle expansion process, corrupting the azimuthal symmetry that was desired in the fuel/air mixing region of the test section. Simulations without this support structure maintained a high degree of azimuthal symmetry up until the fuel injection plane. However, even in this scenario, the azimuthal symmetry was not maintained once the centerbody boundary layer transitioned to a wake flow downstream of the fuel injection plane.

INTRODUCTION

The accurate prediction of scramjet engine combustor performance hinges on the ability of CFD tools to predict the following flow phenomena:

(1) Fuel/Air Mixing: The prediction of the rate of mixing can accept little margin for error given the limited flow residence time within the scramjet combustor.

(2) Flameholding: Maintaining a robust flame given the limited flow residence time requires subcomponents for flameholding, which must be implemented in a manner that avoids excessive drag penalties.

(3) Engine Ignition: Propagating the flame from flameholding regions (where ignition is initiated) to the flammable portion of the supersonic core flow is a challenge.

(4) Heat Release Distribution: At low hypersonic flight conditions, the addition of heat must be controlled spatially to ensure that the thermally throated engine does not result in engine unstart due to isolator over-pressurization.

(5) Heat Loads: All the surfaces of a scramjet engine will experience high levels of heat transfer. However, the combustor component is particularly susceptible to localized hot spots near surfaces adjacent to combusting flow as well as local stagnation regions in the vicinity of shear layer reattachments (e.g., aft face of cavity flameholders).

(6) Shock/Boundary Layer Interactions: Dual-mode engine operation results in shock boundary layer interactions within the isolator, which by design result in massive flow separation.

(7) Turbulence-Chemistry Interactions: The influence of turbulence on the chemical kinetic processes results in “unmixedness” that typically reduces the effectiveness of the combustion process. The impact of these interactions can affect the operability of the engine, particularly near the edges of engine operability.

All the above phenomena are difficult to predict with a high degree of confidence, and validation data are sorely needed to quantify the predictive accuracy of the physics models used to predict them. The validation experiment that is the focus here will provide data for physics model assessments related to items 1, 2, and 7 above.

Seleznev et al.[1] provide a comprehensive overview of existing experimental databases and corresponding simulations relevant to scramjet engine flowpaths. In general, the existing databases can be broken down into three categories: full engine flowpaths, simplified model geometries or individual components, and canonical geometries. Full engine flowpaths contain multiple engine components: inlet, isolator, combustor, and exhaust nozzle. The combustor is also likely to contain one or more flameholding devices and multiple fuel injectors. Hence, full engine databases inevitably contain complex and often interacting flow physics. Such databases are important, as accurately predicting the entire engine flowfield is the ultimate test of a simulation approach used for the design or evaluation of operational scramjet engines. However, using these databases for the validation of a simulation approach comes with many challenges. Given the complexity of the full engine flowfield with all the interacting flow physics, it is generally not possible to isolate individual effects and opens up the possibility of getting “the right answer for the wrong reasons.” For example, it is possible to obtain a good match to the measured pressure distribution by underpredicting mixing, but overpredicting heat release. Moreover, full engine configurations usually have limited quantitative measurements (exacerbating the difficulty with isolating cause and effect) and limited information for boundary condition specification. As a result, these databases are of little use for the purpose of validating specific physics submodels present in CFD simulation software.

The second category of existing scramjet engine databases, simplified model geometries or individual components, typically offers a richer set of measurements for model validation, while also reducing the extent of flow physics interactions as compared to full engine flowpath databases. Facilities used for these measurements also tend to be designed to provide more access for additional measurements. In addition to mean wall pressure measurements, nonintrusive measurements of the flowfield using laser-based diagnostics, such as Planar Laser-Induced Fluorescence (PLIF) and Coherent Anti-Stokes Raman spectroscopy (CARS), are often possible. This database class often provides more highly resolved data as well. There are quite a few examples of existing scramjet engine datasets that fall into this category.[2–4] The simplified geometry of these configurations, along with the availability of nonintrusive instream flow measurements, make these databases very useful for model validation. However, these databases tend to also suffer from poor boundary condition characterization, and the flow physics can still be complex and difficult to isolate.

The final database category is that of simple canonical experiments. These configurations usually contain minimal geometrical complexity, and may bear little resemblance to the components found in a full engine flowpath. Such simple canonical experiments are intended to isolate a small subclass of flow physics, which for a scramjet combustor might simply involve supersonic jet or shear layer flows. These types of experiments can more readily be performed in a variety of facilities, include ample optical accessibility for advanced diagnostics, and are more likely to have multiple measurements available to thoroughly assess the adequacy of specific physics submodels. As such, these types of experiments are preferred for submodel validation purposes. Unfortunately, databases in this category have not received adequate attention given the recent emphasis on flight demonstration systems[5,6] where hardware performance testing is the overarching objective. Historical databases exist,[7,8] but these measurements were either taken using intrusive diagnostics or, if modern nonintrusive diagnostics were used, the experiments were likely not carried out using accepted guidelines[9] for obtaining validation quality data.

The purpose of this computational effort was to perform pretest simulations of a coaxial jet experiment[10] that was designed and developed based on a gap analysis from the High Speed Code Credibility Workshop series sponsored by the Air Force Research Laboratory. One of the goals of this workshop was to identify key gaps in validation quality data for airbreathing and tactical boost glide hypersonic systems, and fundamental fuel/air mixing data in a relevant scramjet engine flow environment was one of the gaps identified. This experimental campaign is intended to provide high quality data for validating turbulent mixing and combustion submodels selected for scramjet engine simulations using modern guidelines for experimentation targeted for model validation (rather than hardware testing or physics discovery). The specific objective of this effort was to assess the extent of the facility flowpath required in the computational domain chosen for this experiment to adequately replicate the flowfield entering the test section of the facility. This paper documents the work performed for this computational effort and is organized as follows:

- The computational methodology that uses anisotropic tetrahedral grid adaptation is described.
- The efficacy of this computational approach is then demonstrated on a supersonic jet flow problem.
- The coaxial jet validation experiment geometry and the flow conditions chosen for this activity are described next along with the various facility geometrical simplifications considered.
- Finally, the results of this pretest activity are discussed followed by a summary of the findings.

COMPUTATIONAL METHODOLOGY

All results were obtained using the VULCAN-CFD[11] unstructured grid solver. This solver allows for mixed element unstructured meshes that contain any combination of tetrahedral, prismatic, pyramidal, and hexahedral cell elements, but the anisotropic adaptive grid methodology used in this effort is currently limited to tetrahedral elements. The simulations were all performed using the LDFSS scheme[12] to construct the inviscid fluxes, with the higher-order cell face variable reconstruction limited using the van Leer limiter function.[13] The viscous fluxes were evaluated by averaging adjacent cell properties to obtain cell face values with the face-tangent augmentation method[14] used to stabilize the face gradients. Gradients were computed using an unweighted least squares approach for the higher-order variable reconstructions required to form the inviscid fluxes, while an inverse-distance weighted least squares approach was used for the gradients required to form the viscous and source terms. A node-neighbor stencil was used to form all gradients, and a point implicit Symmetric Gauss-Seidel (SGS) scheme was used to advance the solution to steady-state. The turbulence model chosen for this effort was the Menter-SST[15] model with the turbulence production terms evaluated using vorticity magnitude rather than strain rate magnitude, and the turbulent Prandtl number was set to 0.9.

The primary objective of this effort was to assess the impact of facility geometrical artifacts required for mounting of (and fuel delivery to) the centerbody fuel injection system for the AFRL coaxial jet experiment. Given that geometry is the primary parametric control variable, it was advantageous to use the *refine*[16] anisotropic tetrahedral adaptive gridding workflow available within VULCAN-CFD. This feature eliminates the human-in-the-loop grid generation process, and adaptively refines the grid as part of the automated workflow. Hence, this workflow is ideal for introducing high-fidelity CFD earlier into the design process, and the adaptive refinement tends to produce meshes that are substantially smaller than manually crafted grids to achieve a level of grid convergence adequate for engineering decision making. Moreover, the adaptation workflow naturally allows for grid convergence assessments as part of the automated workflow.

In the *refine* workflow, the user must supply a watertight geometry for the computational domain boundaries. From there, an initial mesh is generated directly onto this geometry representation. During the initial mesh generation phase, an association is made between the discrete CFD surface mesh and the underlying CAD representation, stored as an EGADS (Electronic Geometry Aircraft Desigh System[17]) file, allowing later adaptations of the CFD mesh to be projected directly onto the underlying CAD. Compared to the traditional CFD workflow, the initial CFD mesh is typically very coarse, where refinement is only applied to regions of high geometric curvature. The CFD solver is used to then obtain an initial steady-state solution on this coarse mesh. An error estimate is made on the initial solution based on the Hessian matrix of a chosen scalar sensor field (e.g., Mach number). This error estimate is converted into a 3D tensor field describing the desired spacing and orientation of a new mesh to reduce the estimated error. The next mesh is then adapted until it conforms to this 3D tensor field. The previous CFD solution is then interpolated to this adapted mesh for use as an initial condition for the next CFD solve, and the process continues yielding increasingly accurate CFD simulations until the CFD analyst elects to halt the process. The process just described is shown schematically in Fig. 1.

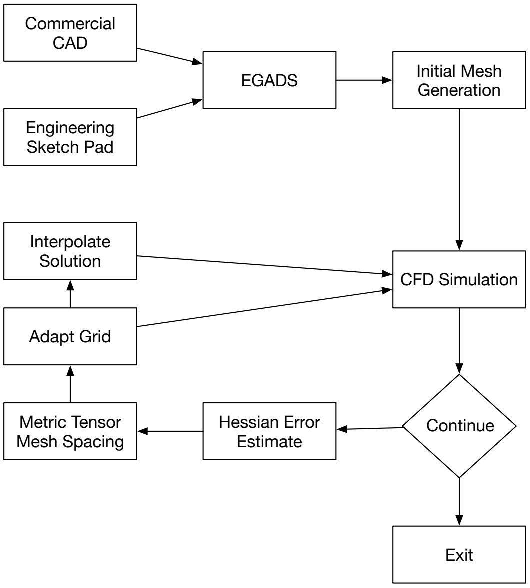

Figure 1: Unstructured grid adaptation workflow.

VERIFICATION OF THE ANISOTROPIC TETRAHEDRAL ADAPTATION PROCESS

Prior to exercising the *refine* process for the coaxial jet validation experiment, a set of simulations was performed on a jet flow configuration that was the focus of the 6^{th} AIAA Propulsion Aerodynamic Workshop (PAW6). This PAW6 jet configuration was taken from experiments[18,19] performed in the Small Hot Jet Acoustic Rig (SHJAR) located within the Aeroacoustic Propulsion Laboratory at the NASA Glenn Research Center. In addition to the measured data, there are simulation results available from a variety of CFD codes. Moreover, this axisymmetric jet configuration allowed grid convergence to be achieved with modest computational resources using a sequence of uniformly refined structured grids. Hence, the existing structured grid simulation data can be leveraged to ascertain whether the same grid converged result can be attained using solution adapted triangular meshes (or tetrahedral meshes in 3D). Recent work[20] suggests that grid convergence is not easily achieved for jet flows given a sequence of manually refined tetrahedral grids.

The computational domain for the PAW6 axisymmetric jet experiment is shown in Fig. 2. The Mach 1.63 jet conditions are obtained by expanding the air flow through the converging/diverging nozzle shown on the lower left corner of this image. The nozzle plenum inflow condition was defined using the measured stagnation conditions and assumed values for the turbulence intensity and eddy viscosity to molecular viscosity ratio (see Table 1). Mach number was extrapolated from the adjacent interior cells to fully characterize the subsonic plenum inflow state. Characteristic boundary conditions were imposed at the ambient inflow and farfield boundaries given the external state conditions provided in Table 1. A specified pressure condition was enforced at the outflow plane, which was placed 40 jet diameters (D) downstream of the nozzle exit. The internal and external nozzle walls were treated as adiabatic surfaces using the Wilcox[21] wall matching procedure, with the exception of the horizontal portion of the external nozzle surface where a slip wall condition was used. A slip condition was used for this segment to avoid a boundary condition conflict with the characteristic inflow condition, which assumes an isentropic process across the inflow boundary. Note that the structured grid simulations used a wall-resolved adiabatic condition along the nozzle surfaces. Wall matching versions of these conditions were required for the adaptive grids, since the near wall y^+ values were significantly greater than unity for much of the grid adaptation process. Local time stepping was used with a CFL value of 20 to advance the simulations in pseudotime early in the grid adaptation process. The time stepping strategy was switched to a constant time step approach ($\Delta t = 1.0 \times 10^{-5}$ s) after the 10^{th} adaptation step to achieve deeper iterative convergence on the finer grids. A similar strategy was used by most workshop participants with the provided structured grids.

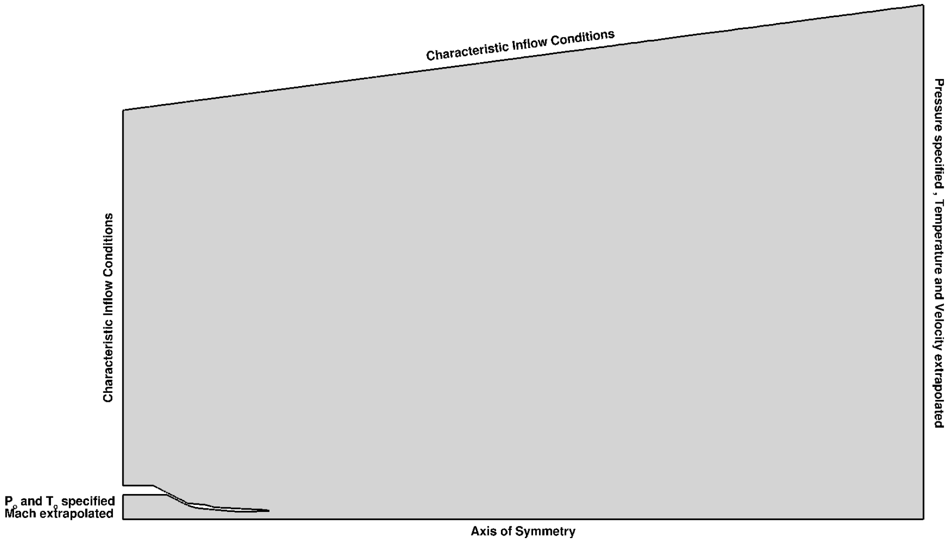

Figure 2: Computational domain used to simulate the PAW6 jet flow experiment.

Table 1: PAW6 jet flow conditions.

Facility condition	Plenum Inflow	Ambient Inflow	Outflow
Pressure [kPa]	440.8032*	99.28	99.28
Temperature [K]	799.4*	288.9	N/A
Turbulence Intensity	0.01	0.001	N/A
μ_t/μ	0.10	0.010	N/A

* stagnation condition

All boundaries of this computational domain involve straight line segments with the exception of the internal nozzle contour. As a result, this geometry was readily defined directly by an OpenCSM[22] script using straight line segment specifications and splining of the internal nozzle coordinates that were provided. The *refine* process is typically more robust when the geometry is prescribed natively in this fashion, so this approach is preferred over importing CAD into the OpenCSM script directly when it is feasible to do so. A single face body (or solid body for the 3D sector mesh) representation of the complete geometry specification was then dumped to an Engineering Geometry Aerospace Design System (EGADS)[17] file as required by the *refine* process. With the geometry and mesh association to this geometry defined by the EGADS file, the fluid domain is then filled with triangles (or tetrahedra for the 3D sector mesh) using the publicly available TetGen[23] software, which is then adapted by *refine* to better capture geometric features.

The *refine* adaptation process in VULCAN-CFD is controlled by an adaptation control file. Required adaptation control parameters include: initial complexity, complexity multiplier, complexity stride, and choice of sensor variable. The grid complexity is a measure of the size of the grid and is approximately equal to half of the number of grid nodes. The initial complexity chosen for the axisymmetric jet simulations was 2000. The complexity multiplier specifies the relative increase in the complexity when required during the adaptation process, and the complexity stride specifies the number of adaptation steps to perform before the next complexity increase. For adaptation steps between strides, the size of the grid is nominally fixed, implying that the grid nodes will only be redistributed. The simulations performed here used a complexity multiplier of 1.25 and a complexity stride of 1, so the grid size increased by 25% in each step of the adaptation process. Mach number was chosen as the sensor variable for the adaptation process. Mach number is often a good choice since it is sensitive to each compressible flow phenomena (expansions, compressions, contact surfaces) as well as regions of shear (e.g., boundary and free shear layers). Additional adaptation control variables are also available, but the only one used in this effort was the allowable cell aspect ratio limit. A value of 4 was chosen for this parameter. Smaller values for this parameter tend to provide higher quality grids, but at an increased computational cost since it constrains the movement of grid nodes during the adaptation process. A value this low is only practical given the axisymmetric nature of this jet geometry.

The initial axisymmetric grid and the subsequent solution adapted grids after 5, 10, and 15 adaptation steps are shown in Fig. 3. These images only highlight the jet nozzle and nearfield exhaust region to better elucidate how the grid is evolving to the flow features. Images of the flow structure that led to the evolution of the grid adaptation process are provided in Fig. 4. In these images, the flowfield has been reflected across the axis of symmetry (x-axis) to provide a better visual of the jet flow structure. The initial grid is shown in the top image of Fig. 3, where it is clear that refinement has only been applied to geometric features. After 5 adaptation cycles, the number of grid nodes has increased by a factor of approximately 3 (1.25^5), and the adapted grid is clearly tracking the shear layer between the ambient air and the jet potential core. The adaptation has also detected the formation of the boundary layer along the internal nozzle surface, and the wave structures in the diverging section of the nozzle are better resolved. However, the error indicator has yet to signal a need to resolve the boundary layer region along the no-slip portion of the external nozzle surface due to the low (nearly ambient) flow velocity there. In fact, the error indicator has redistributed a large portion of the initial grid nodes in this portion of the domain to other areas of the flow to the point of almost losing the flow separation after the first turn of the exterior nozzle surface. After another threefold increase in grid size (10th adaptation step), the adaptation process has been sensitized to the external surface boundary layer to some degree. Further refinement of the exhaust shear region, internal nozzle surface boundary layer, and supersonic nozzle wave structure are also evident. By adaptation step 15 (another threefold increase in grid size) essentially all the expected flow structures have been resolved to some extent. In particular, the weak wave reflections in the jet potential core are now more clearly discernible. These features are weak and were not targeted for resolution early in the adaptation process because the jet exit condition was pressure matched (by design) to the ambient environment. Improved resolution of the separation zone along the external nozzle surface is also evident.

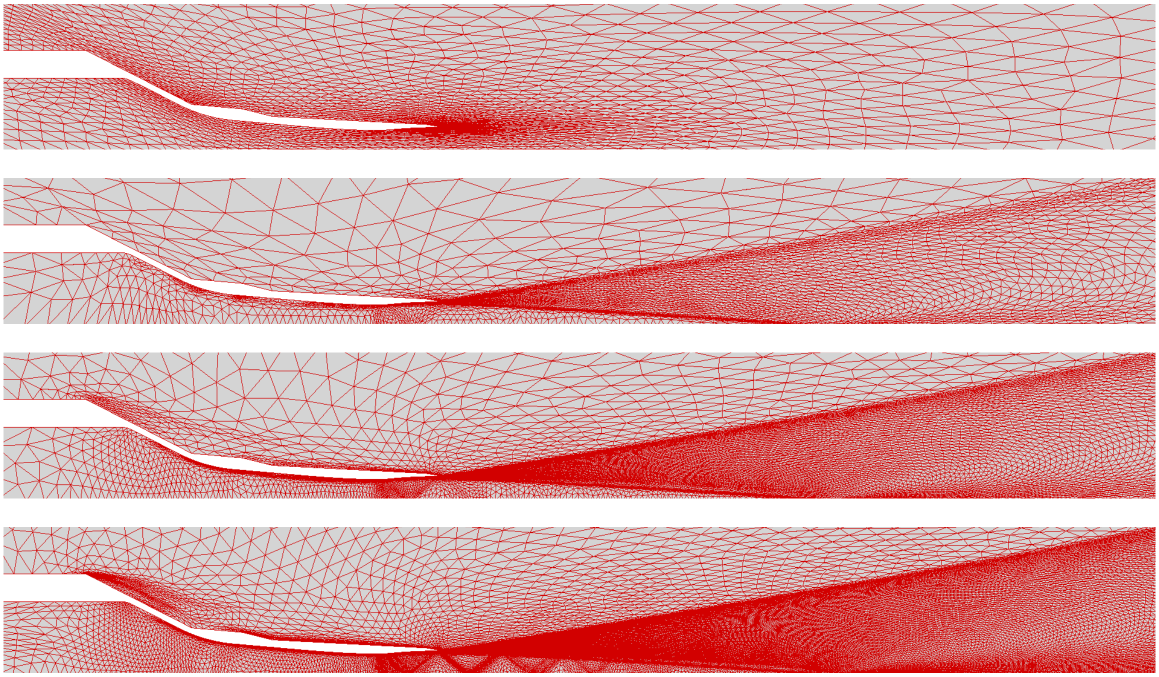

Figure 3: Sequence of adapted grids: initial grid, followed by adaptation steps 5, 10, and 15.

Figure 5 compares predictions of the streamwise velocity along the axis of symmetry (top image) and lip line (bottom image). The lip line is defined here as the horizontal line initiated from where the internal nozzle expansion surface ends (i.e., $y/D = 0.5$). The adapted simulation results were extracted from adaptation steps 10, 15, and 20 (indicated by the red, blue, and green lines, respectively). Also shown are structured grid results on the finest grid (604,800 cells) provided to the PAW6 participants. Simulation results from this grid are shown for three separate CFD solvers (VULCAN-CFD, FUN3D, and WIND-US), each with identical physics model settings (to the extent possible). Data obtained from flowfield measurements has been included as well. This comparison clearly emphasizes the importance of model validation. Even for this canonical jet flow problem, Reynolds-averaged turbulence closure models struggle with reproducing the measurements. In fact, nearly all widely used closure models (without compressibility corrections) overpredict the spreading rate of high-speed jet flow problems. Compressibility corrections[24,25] have been developed to address this deficiency, and these models often improve the comparisons with available measurements. However, for this particular jet condition, these model extensions tend to overcompensate and underpredict the measured spreading rate.

For the purposes here, we are concerned with assessing whether the use of adapted anisotropic grids can replicate the results obtained from “grid converged” hexahedral (e.g., structured) grids, particularly in light of the work of Dippold[20] that illustrated the difficulties with achieving grid independence when tetrahedral meshes are used for jet problems. In other words, the focus here is a solution verification activity, rather than a turbulence model validation activity. From a solution verification viewpoint, the results here are extremely encouraging. Even the results from the coarsest adapted grid shown in this image closely matched the VULCAN-CFD results obtained from the finest structured grid provided to the PAW6 participants. Moreover, this level of agreement was achieved with only 6% as many grid cells. In fact, the differences between all 3 adaptation cycles shown (spanning a factor of 9 in total number of grid cells) and the VULCAN-CFD structured grid result is less than the code-to-code variability using the same structured grid.

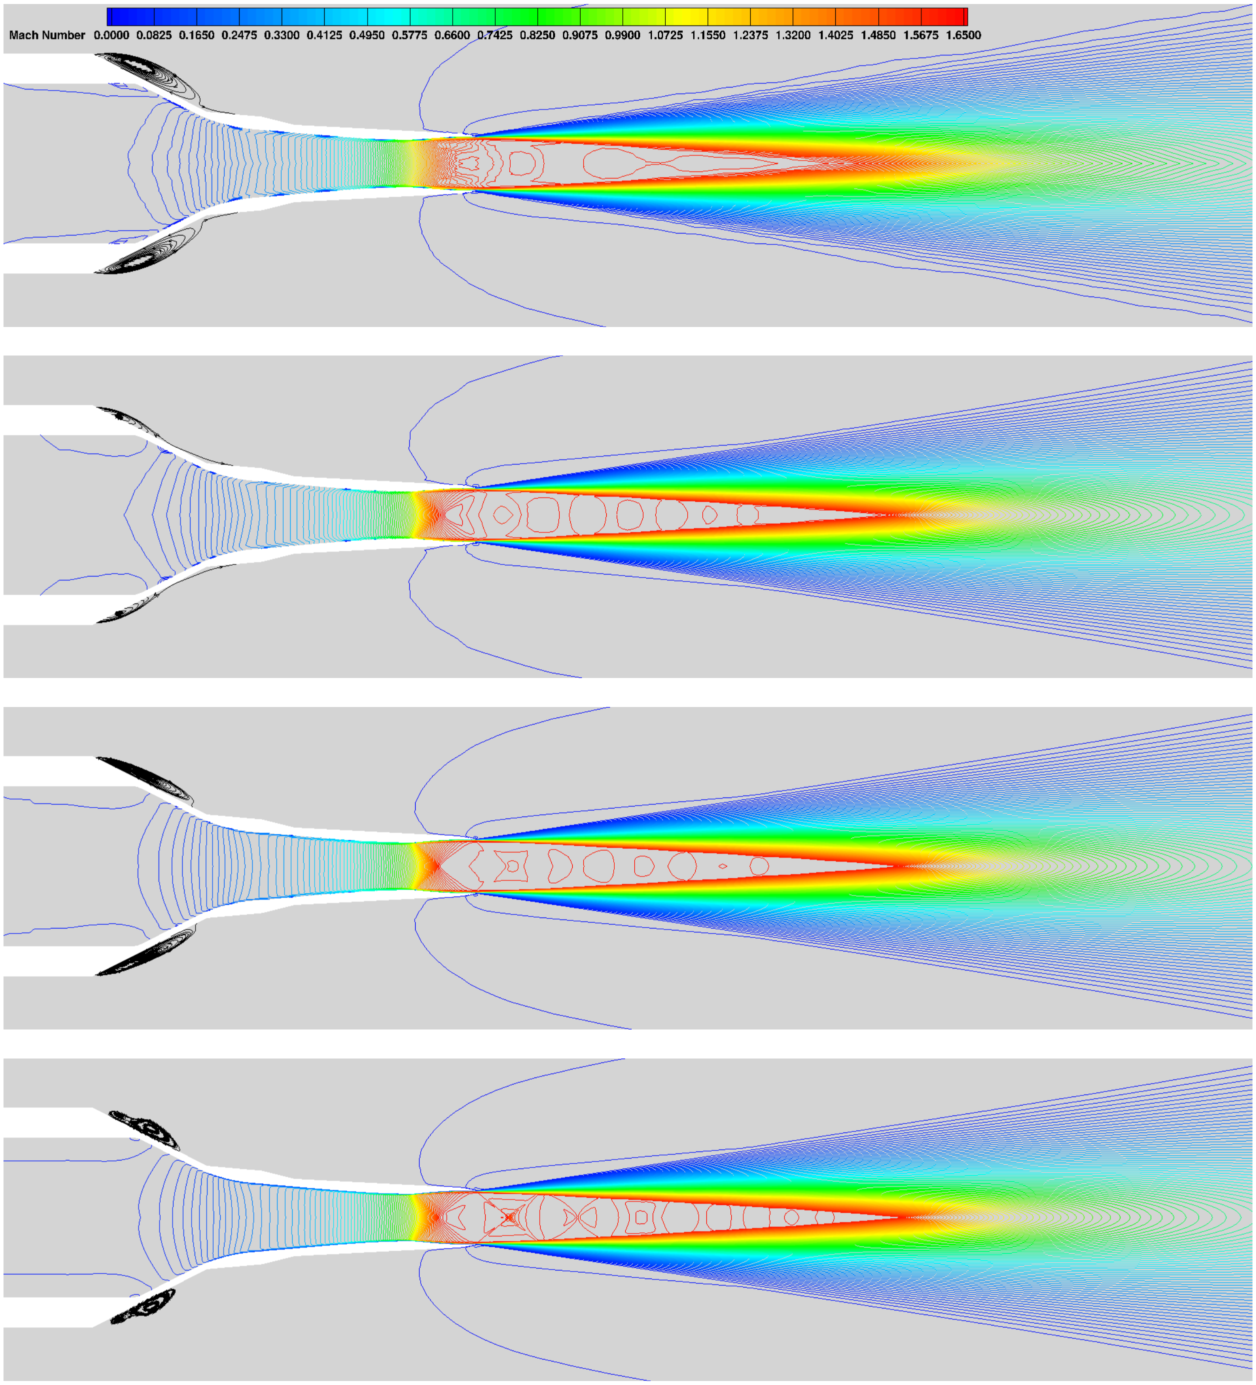

Figure 4: Mach contours (jet nearfield): initial grid, followed by adaptation steps 5, 10, and 15.

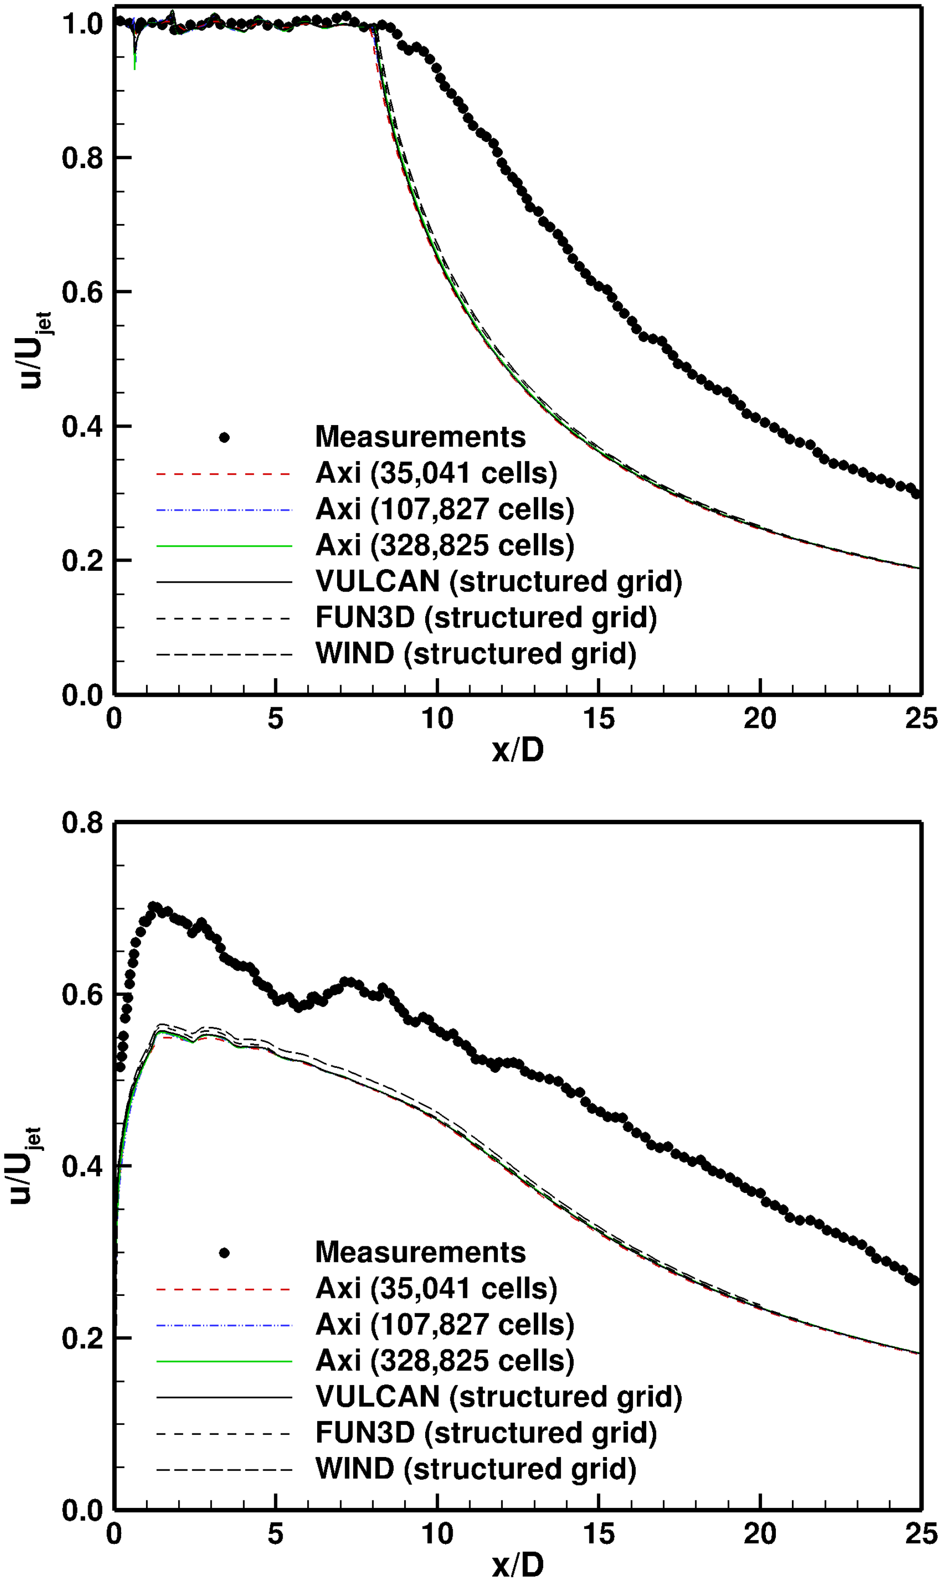

Figure 5: Streamwise velocity along the axis of symmetry (top) and lip line (bottom).

While the axisymmetric adapted grid simulation results were encouraging, these meshes were created by rotating a 2D triangular grid about the axis of symmetry to form a 3D grid composed of a single cell in the azimuthal direction. The cell types in this grid are all prismatic elements except for those adjacent to the centerline that contain alternating tetrahedral and pyramidal elements. Given this observation, additional adaptive simulations were performed on a 3D domain that spans a 30° sector. This provides for a more stringent verification on the efficacy of using tetrahedral dominant meshes for jet flow problems. The adaptation process for this 3D domain was identical to that used for the axisymmetric simulations, except the initial complexity was increased by a factor of 30 to loosely account for the additional cells required to fill a 3D volume as compared to a 2D plane.

Figure 6 compares predictions of the streamwise velocity and turbulent kinetic energy along the axis of symmetry, and Figure 7 compares the same variables along the lip line. Structured grid results (on the finest PAW6 grid) from the 3 different codes are also provided along with the measured data. Once again, the adapted unstructured grid results spanning a factor of 9 increase in cell count are practically indistinguishable from each other and are nearly a perfect match to the VULCAN-CFD structured grid result. This was the case even for the turbulent kinetic energy that is governed by a transport equation with a source term proportional to the square of the velocity gradient. As a final comparison, radial (y-direction) profiles of streamwise velocity and Reynolds shear stress extracted at 3 axial stations are shown in Fig. 8. Again, the anisotropic adaptive tetrahedral grid results are practically identical to the VULCAN-CFD structured grid result. This level of agreement (particularly for the Reynolds shear stresses that are functions of flowfield gradients) far exceeded expectations, and offers a high level of confidence that this numerical approach can be trusted for the pretest simulations of the AFRL validation experiment.

COAXIAL EXPERIMENT GEOMETRY AND TEST CONDITIONS

The full computational domain for the AFRL coaxial jet experiment is shown in Fig. 9. This configuration consists of a centerbody fuel injection system anchored in the center of the facility flowpath, and is sized to provide a significant base flow recirculation zone for flameholding support at the fuel injection plane under combusting conditions. The centerbody extends upstream of the converging-diverging facility nozzle, and is anchored to the facility walls within the nozzle plenum via 2 support struts to minimize flow distortion. Fuel is supplied to the centerbody through one (not both) of the support struts, which dictates that 1/2 (rather than 1/4) of the facility flowpath be included in the computational domain. The converging-diverging facility nozzle provides Mach 1.5 flow to the test section. The constant area test section has a rectangular cross section with allowances for ample optical access from all sides. The fuel injector exit port section can be swapped out to accommodate other port designs, but only a constant area 2.5 mm diameter port was used for this activity. The reader is referred to DeBardelaben et al.[10] for further details on the design of this apparatus and the diagnostics planned for the experiments.

The facility plenum inflow condition was set using the nominal stagnation conditions shown in Table 2 and assumed values for the turbulence intensity and eddy viscosity to molecular viscosity ratio of 0.001 and 0.01, respectively. Mach number was extrapolated from the adjacent interior cells to fully characterize the subsonic facility plenum inflow conditions. The mass flux and total temperature were specified at the fuel inflow plane, and pressure was extrapolated from the interior to complete the fuel flow specification. The outflow plane was placed just downstream of the rectangular to circular expansion section (rather than the end of the test section) to increase the likelihood of fully supersonic flow exiting the domain, permitting the use of a simple extrapolation boundary condition there. The fuel piping surfaces were assumed to be sufficiently isolated from the heated facility flow to permit the use of an adiabatic surface condition. The thermal boundary condition for the remaining facility surfaces were set by balancing the heat flux at the wall with the wall-normal heat flux through the steel structure given an assumed back side temperature of 325 K at a depth of 0.5 inches. The surface temperature is expected to be far from uniform under combusting conditions, so specifying the wall conditions in this way allows the surface temperature to vary spatially based on local flow conditions. Finally, wall matching versions of all no-slip wall boundary conditions were

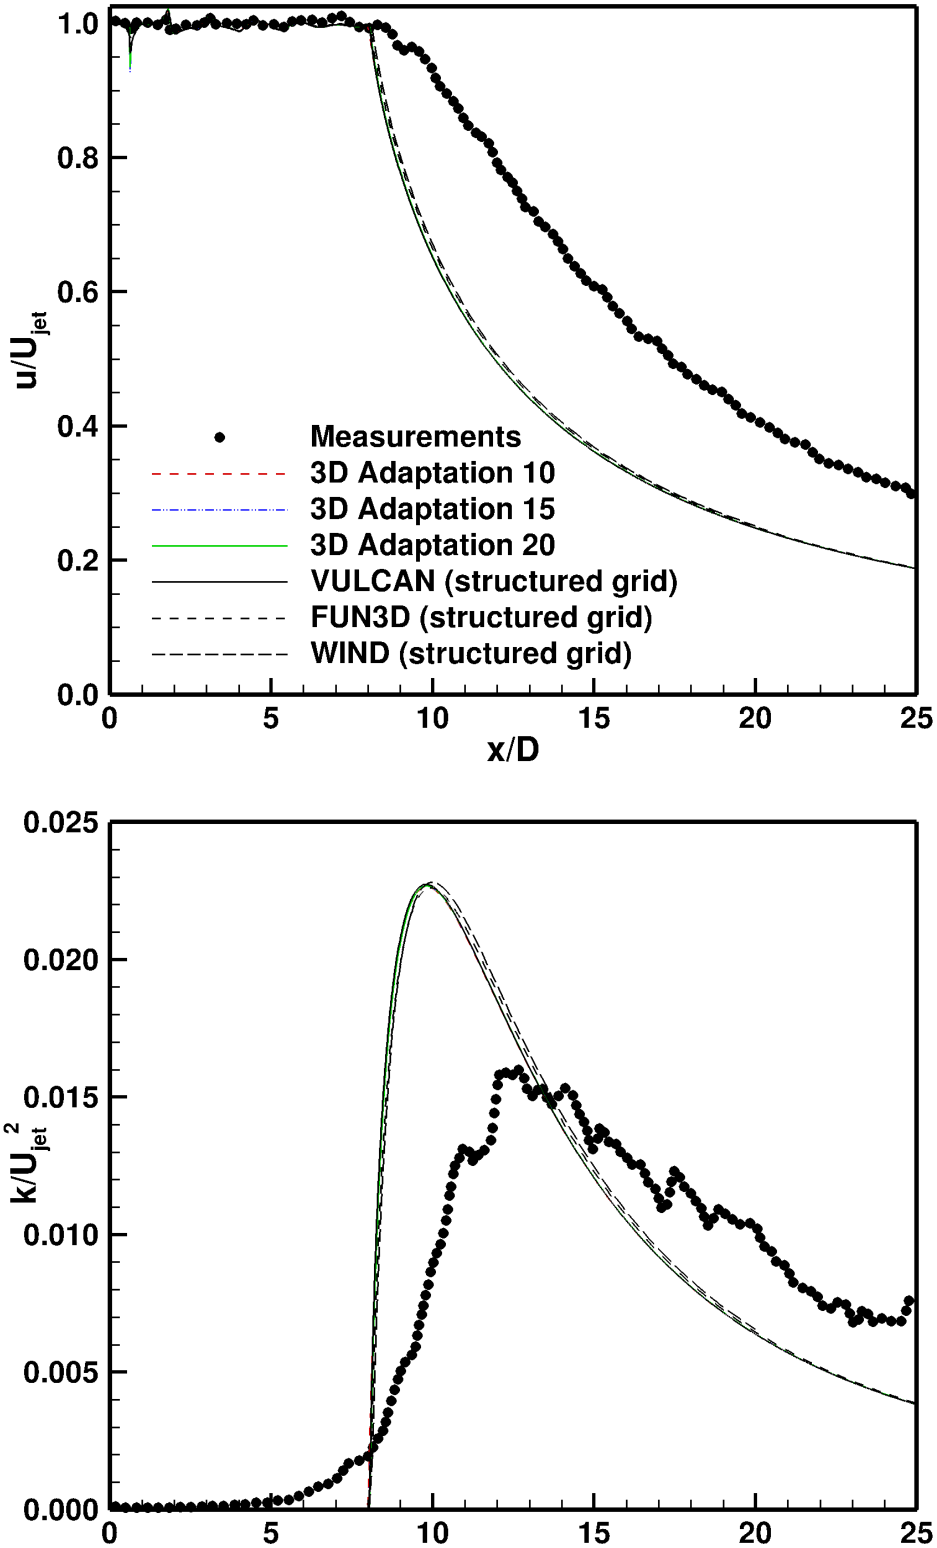

Figure 6: Streamwise velocity (top) and turbulent kinetic energy (bottom) along the axis of symmetry.

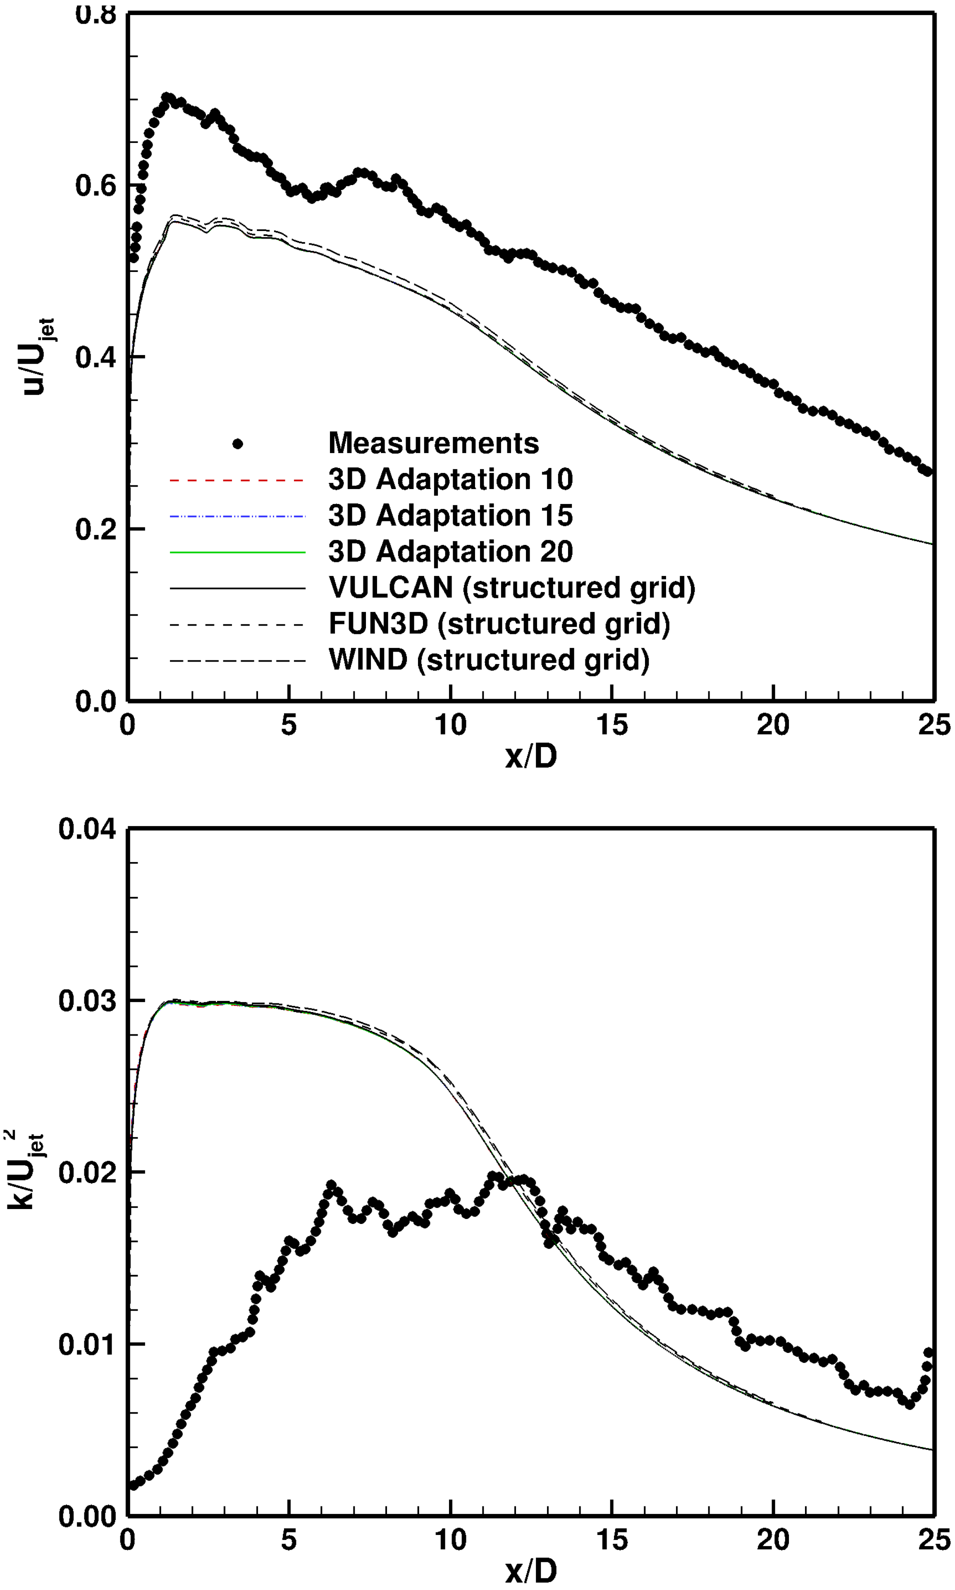

Figure 7: Streamwise velocity (top) and turbulent kinetic energy (bottom) along the lip line.

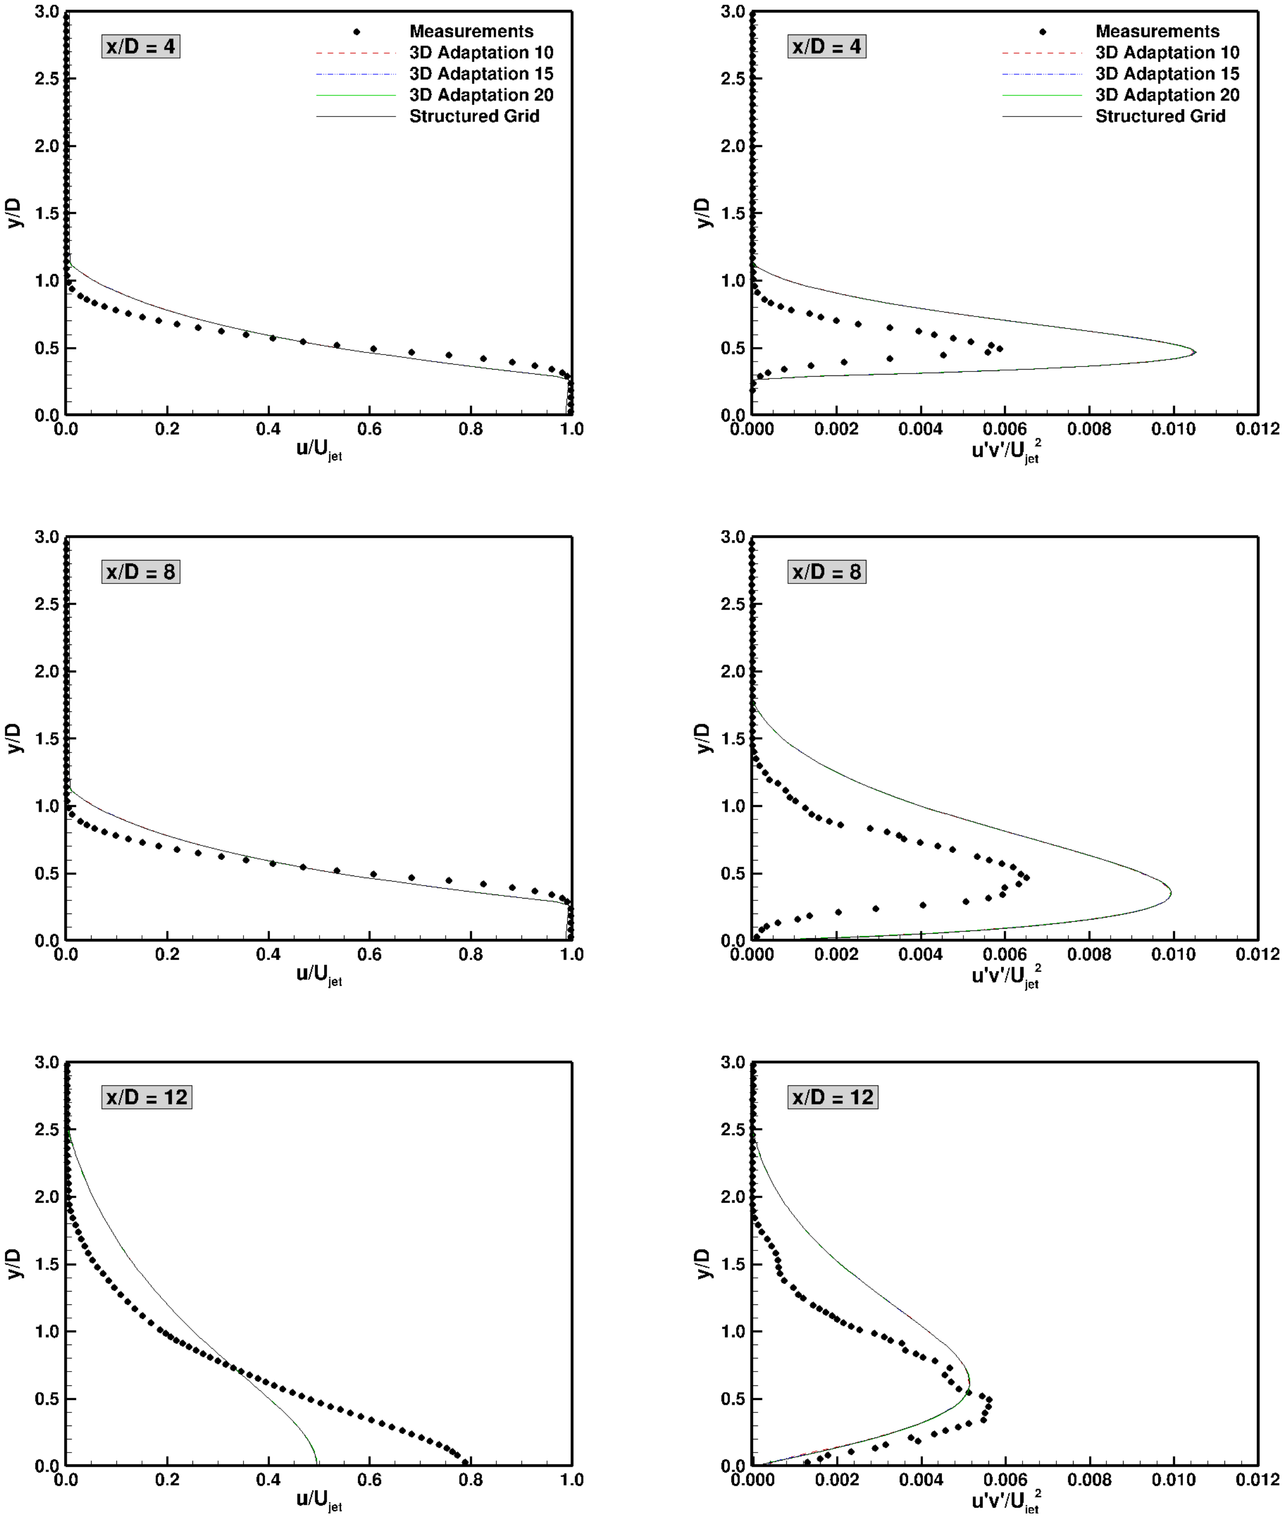

Figure 8: Streamwise velocity (left) and Reynolds shear stress (right) profiles.

used, since the wall adjacent y^+ values were significantly greater than unity for much of the grid adaptation process.

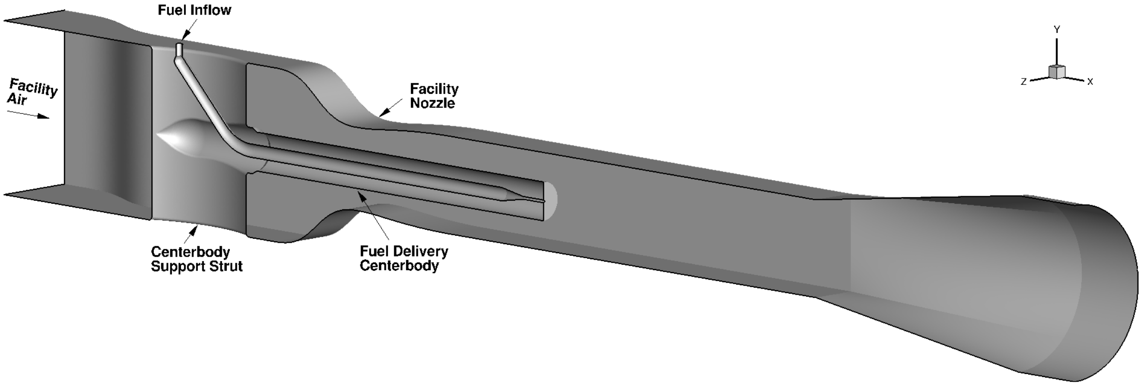

Figure 9: Full computational domain used for the AFRL coaxial jet experiment.

Table 2: AFRL coaxial jet flow conditions.

Facility condition	Facility Inflow	Ethylene Inflow
Total Pressure [kPa]	827.371	N/A
Mass Flow Rate [kg/s]	N/A	0.01
Total Temperature [K]	699.778	300.0
Turbulence Intensity	0.001	0.001
μ_t/μ	0.01	0.010

As mentioned previously, the main objective of this effort was to determine the extent of the facility flowpath geometry required to adequately replicate the facility test section flow environment. Toward that end, simplified alternative computational domains were investigated as shown in Fig. 10. The 1^{st} simplification (top image) simply truncated most of the upstream fuel plumbing. This simplification allows 1/4 of the flowpath to be simulated, reducing the computational costs by at least a factor of 2. The 2^{nd} simplification (bottom image) removes the upstream portion of the nozzle plenum containing the centerbody support structure in addition to truncating the upstream fuel plumbing. Note that the images shown here contain the same 1/2 facility flowpath required by the full configuration, but all simulations of these truncated configurations took advantage of both top-to-bottom and side-to-side symmetry.

RESULTS AND DISCUSSION

The numerical approach used for this configuration closely matched that used for the PAW6 simulations, but there were a few differences. The turbulence model chosen for this configuration was the 2003 version of the Menter-SST[26] model (again with vorticity used in lieu of strain rate in the turbulence production terms). The 1994 version of the model was used for the PAW6 verification activity to match what was used by the other flow solvers for the code-to-code comparisons. The only notable difference between the 2 versions is that the 2003 variant uses the strain rate magnitude to limit the shear stress rather than the vorticity magnitude. This difference sensitizes the stress limiter to shock waves, which mitigates the excessive amplification of transported turbulence properties across strong shocks that often plagues two-equation closures (particularly when the grid is highly resolved). The turbulent Prandtl number value chosen was 0.9 (same as the PAW6 simulations). The presence of multiple species necessitates the specification of a turbulent Schmidt number for these simulations, and a value of 0.7 was chosen based on calibration activities performed under similar flow conditions.[27,28] These calibrations used scale-resolving simulations to determine the Sc_t value that best correlated the resolved Reynolds mass flux vector to the gradient

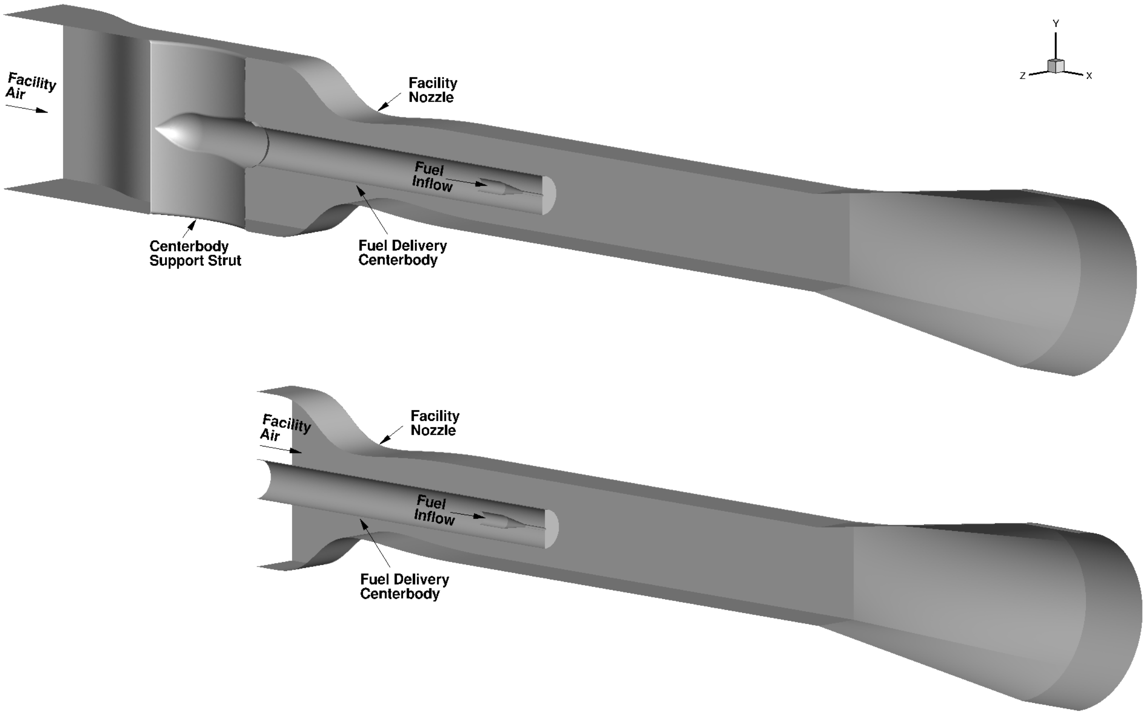

Figure 10: Alternative truncated computational domains considered for the AFRL coaxial jet experiment.

diffusion functional form used to model it for Reynolds-averaged simulations. Local time stepping was used with a CFL value of 10 to advance the simulations in pseudotime for these reacting flow simulations, and the chemistry model chosen was the 3-step, 6-species ethylene model calibrated by Hassan et al.[29]

Images of the combusting flowfield extracted from the last adaptation cycle are provided in Fig. 11. Mach number contours shown in the top image highlight the aerodynamic features of this flowfield. The facility nozzle is seen to provide nearly distortion free Mach 1.5 flow to the test section, and the evolution of the boundary layer growth along the facility surfaces is clearly visible. The combusting wake structure downstream of the centerbody remains subsonic until the aft expansion section is reached, highlighting the importance of including some portion of the facility expansion section. An expansion fan forms at the start of the expanding rectangular to circular shape transition section, which is followed by the formation of a shock wave at the end of the expansion section. Both of these features have been well resolved by the adaptation process. A zoomed-in view of the injected fuel jet flow (denoted by the purple box) is shown in the upper right corner. This image highlights the level of resolution obtained by the adaptation process for the underexpanded fuel jet flow, which consists of a well-developed initial Mach disk followed by a shock diamond pattern formed by multiple shock and expansion waves as the fuel jet adjusts to the surrounding flow conditions. The contours of temperature show where combustion is occurring within this configuration. The base of the centerbody has provided adequate flameholding support as evidenced by the elevated temperatures adjacent to the base. The flame (aided by this flameholding support) initiates shortly downstream of the centerbody base and spreads through the wake region as the flow convects downstream. A high level of symmetry is also noted in both images despite the fact that the fuel delivery system is not top-to-bottom symmetric.

The adaptation settings chosen for this configuration matched those used for the PAW6 simulations except for the initial complexity and allowable aspect ratio limit. The initial complexity was set to 120,000 (a larger value was used here given the larger computational domain), and the allowable aspect ratio was increased to 100 to allow *refine* more freedom with grid point redistribution. Without the change, *refine* was particularly slow with resolving the underexpanded fuel jet structure. The initial grid and the subsequent solution adapted grids after 5, 10, 15, and 20 adaptation steps are shown in Fig. 12. Corresponding images of the

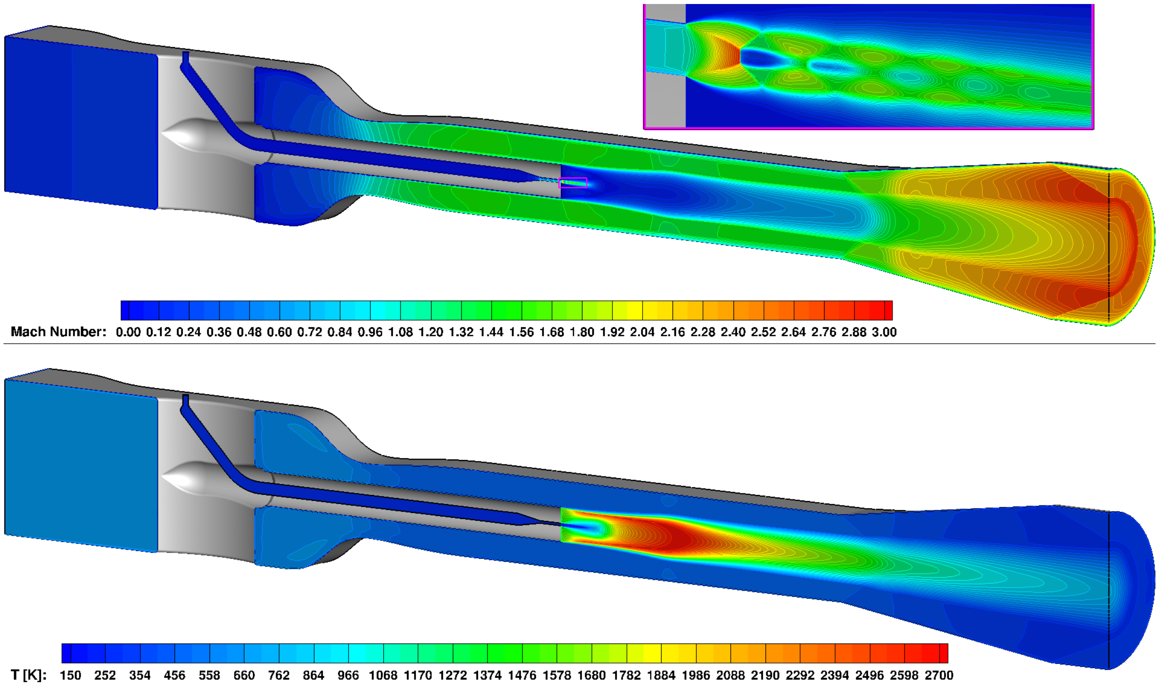

Figure 11: Mach number (top) and temperature (bottom) contours for the AFRL coaxial jet experiment.

flow structure that led to the evolution of the grid adaptation process are provided in Fig. 13. These images only highlight the underexpanded fuel jet flow region, which proved to be the portion of the domain that the adaptation process was slowest to respond to. The initial grid and Mach number field are shown in the top image of Figs. 12 and 13, respectively, with each subsequent image corresponding to adaptation steps with a threefold increase (1.25^5) in the number of cells. The first 10 adaptation steps showed only minor progress toward resolving the underexpanded jet flow structure. Instead, the error estimator prioritized other portions of the computational domain for additional resolution, e.g., the facility wall boundary layers and the downstream expansion fan (and shock wave) in the facility expansion section. Evidence of the full jet flow shock diamond structure did not appear until roughly the 15^{th} adaptation step (29.2 million cells). By the 20^{th} adaptation step (87.9 million cells), the error estimator has clearly placed more emphasis toward resolving the underexpanded jet flow structure. However, a significant change in jet flow structure occurred between the 15^{th} and 20^{th} adaptation steps, so further adaptation cycles would be required to confirm whether the grid converged jet structure has been obtained. For the purposes of this pretest effort, where only sensitivities related to facility geometry simplifications are being assessed, the adaptation process was halted at this point. It is worth noting the benefit of the systematic grid refinement inherent to the *refine* workflow, which allows the analyst to make informed engineering judgment calls such as this to the CFD process.

Figures 14, 15, 16, and 17 compare Mach number and H_2O mass fraction contours at several test section streamwise stations. In all these figures, the right half of each image contains the results extracted from the full geometry, while the left half corresponds to results extracted from the modified geometries shown in Fig. 10. The simulations for the modified geometries took advantage of top-to-bottom symmetry, so the results were mirrored to provide a one-to-one correspondence with the full configuration. Figures 14 and 15 show the impact of neglecting most of the fuel piping within the centerbody. This simplification is commonly applied when defining the computational domain for scramjet fuel injectors. However, given the extremely long (in terms of L/D) piping length for this configuration, and the fact that fuel is supplied from only one side of the centerbody (rather than in a symmetric fashion from both sides), tests seemed warranted to verify this approximation. A careful comparison of these images show very little differences, and the level of symmetry captured from the full computational domain deviates very little from the perfect symmetry present in the truncated fuel plumbing geometry results. A more quantitative assessment will be made later, but

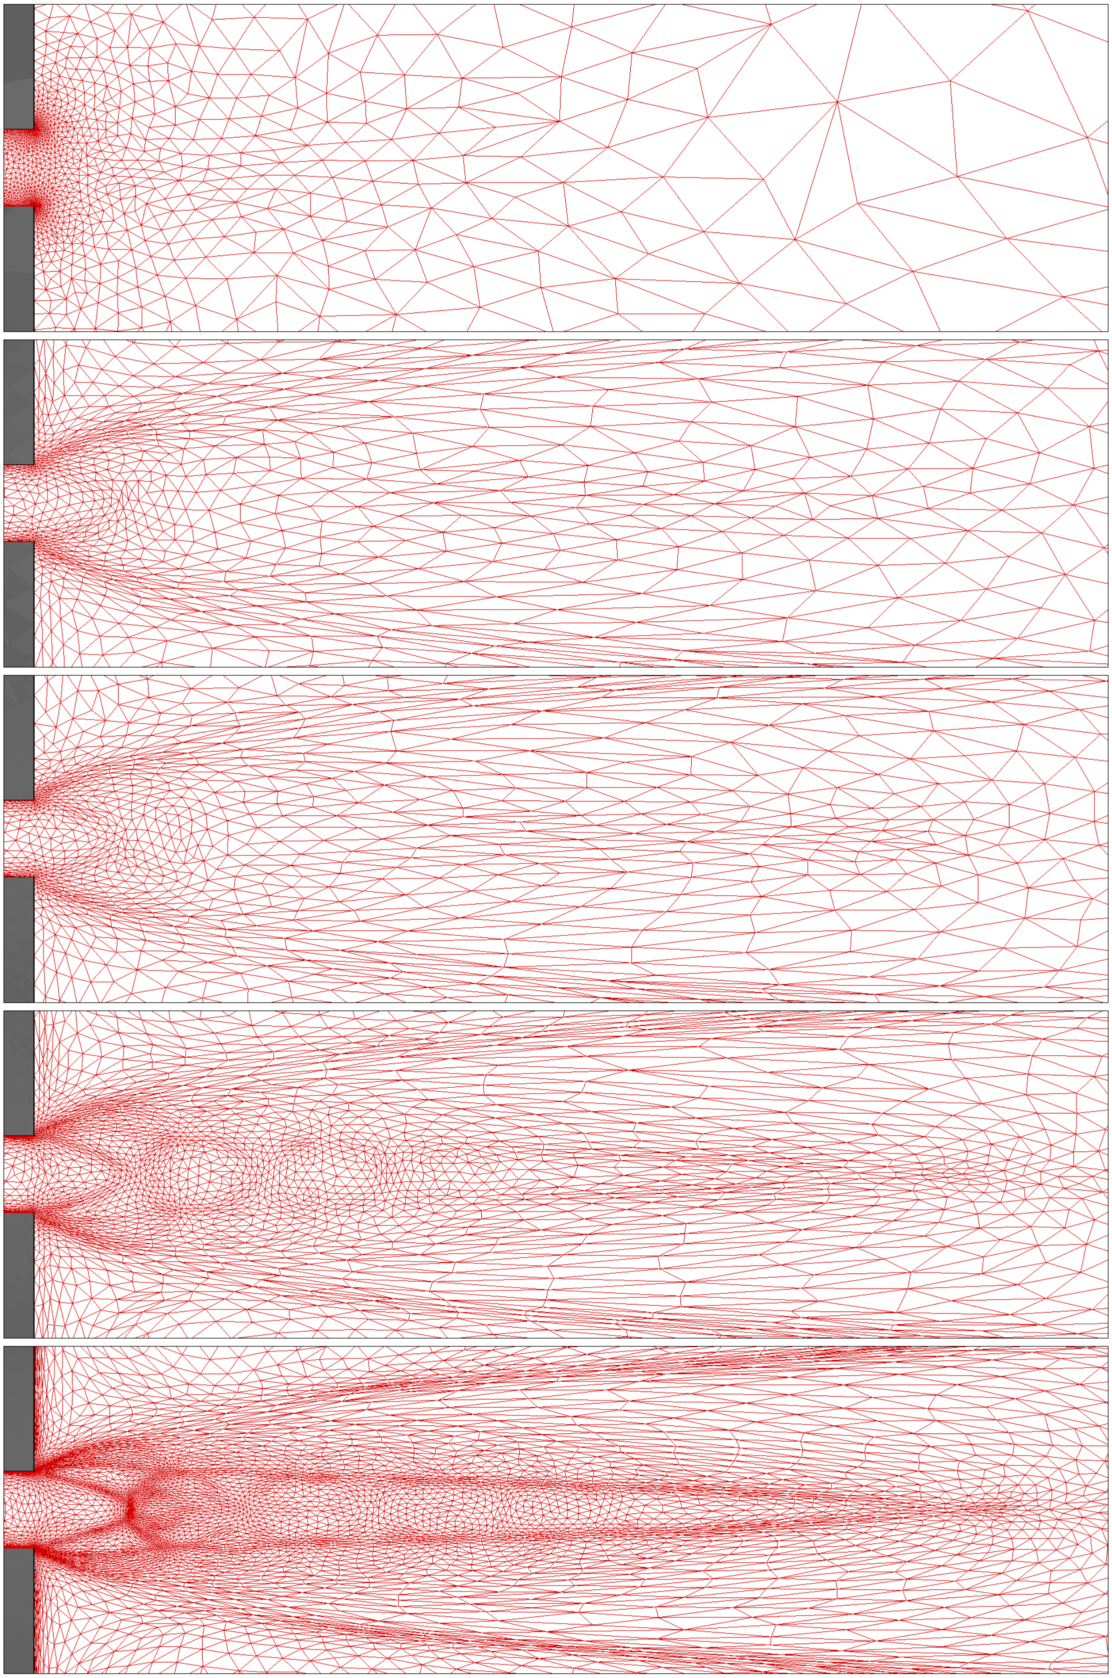

Figure 12: Sequence of adapted grids: initial grid, followed by adaptation steps 5, 10, 15, and 20.

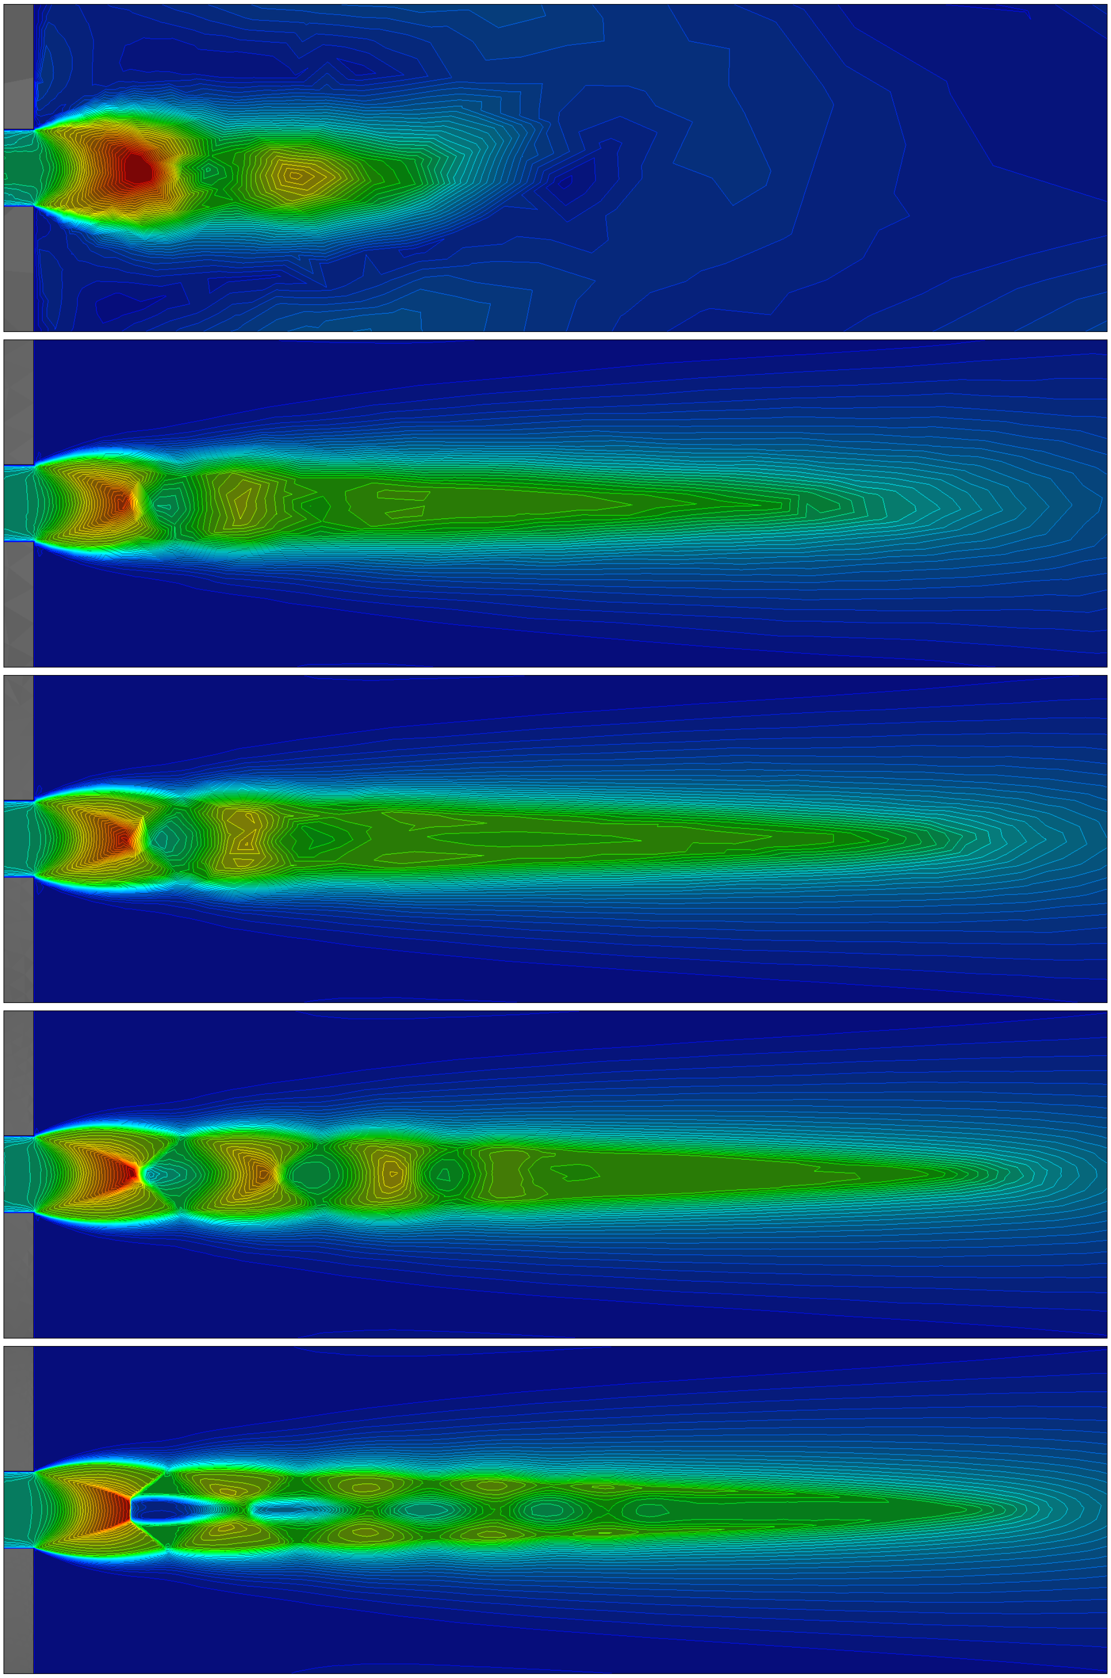

Figure 13: Mach contours (jet nearfield): initial grid, followed by adaptation steps 5, 10, 15, and 20.

the level of agreement shown here would seem to justify the use of this modified computational domain (which reduces the cost of the simulation by a factor of 2). However, the contours from both simulations show some deviation from azimuthal symmetry around the centerbody. In particular, there is a slight build up of boundary layer fluid on the $y = 0$ centerline (i.e., the top-to-bottom facility centerplane) seen at the fuel injection station. By the 2-inch station, the circumferential shape of the wake flow more closely resembles a hexagon than a circle, so the flowfield is not as azimuthally symmetric as was originally expected.

Figures 16 and 17 show the impact of also neglecting the upstream nozzle plenum section that housed the rounded leading edge of the centerbody and the centerbody support struts. By keeping these features in the low speed plenum region, it was hoped that their impact on the downstream flow would be minimal. Unfortunately that was not the case. When these artifacts are neglected, the resulting wake flow region appears considerably more azimuthally symmetric. The noticeable fluid dynamic impact of the upstream centerbody support structure is not an ideal situation, but a substantial effort is planned to thoroughly characterize the flow entering the test section to define (or at least verify) the CFD boundary condition specification. This emphasis is one of the primary differences between experimentation performed for the purpose of CFD validation versus experimentation performed for hardware testing or physics discovery. The specific aspect of the upstream flow that caused the hexagonal wake flow cross-sectional shape will be revealed later.

Figure 18 provides a more quantitative assessment of the degree of symmetry maintained through the wake flow downstream of the centerbody. These images show radial H_2O mass fraction profiles extracted at the side-to-side facility centerplane ($z = 0$) and the top-to-bottom facility centerplane ($y = 0$) at 3 streamwise stations. The left image compares the radial extractions for the full configuration (black lines) with those from the 1^{st} domain modification that simply neglected most of the fuel piping. In these images, the radial distance has been normalized by the centerbody radius, R_c. At the base of the centerbody ($x = 0$), all the radial profiles agree well, indicating a high level of azimuthal symmetry for both configurations. However, a significant deviation is evident between the $y = 0$ and $z = 0$ radial extractions in the wake region further downstream. The comparison of radial profiles between the 2 configurations is almost a perfect match though, validating the qualitative assessment made earlier. The right image shows the corresponding comparison between the full configuration and the 2^{nd} domain modification that neglects the upstream portion of the centerbody and support struts. This comparison shows a similar discrepancy in symmetry between the $y = 0$ and $z = 0$ radial extractions at the 2-inch and 4-inch wake flow stations. Hence, the contours shown for the 2^{nd} alternative computational domain in Figs. 16 and 17 only appear to be azimuthally symmetric. Moreover, given that the asymmetry in the $y = 0$ and $z = 0$ radial extractions is present even without the upstream centerbody and support strut hardware, these artifacts are not the root cause of this asymmetry. Some level of asymmetry is to be expected when an axisymmetric body is placed within a constrained rectangular flowpath, but this concession was necessary to maximize optical access for the nonintrusive measurement techniques planned for the experimental campaign. There was a slight difference in each radial profile at the downstream stations between the 2 configurations (particularly at larger radial values). This difference is an artifact of flow structures formed by the upstream hardware within the nozzle plenum.

The root cause of the slight difference noted in each radial profile between the full configuration and the configuration without the upstream hardware can be traced back to vortical structures that formed along the centerbody. These structures are shown in Fig. 19 in the form of a fence plot of the turbulence kinetic energy. This image shows several streamwise planes starting just upstream of the leading edge of the centerbody and ending near the facility nozzle throat. The immersion of the centerbody and the support struts in the flowpath introduces a change in cross-sectional flow area. The first 4 planes correspond to an area decrease with downstream distance, which implies that a favorable streamwise pressure gradient is present. The 5^{th} plane shows a stark increase in the turbulence kinetic energy near the corners and at both ends of the centerbody. This plane is on the backside of the strut where the flow area is starting to increase with streamwise distance. This induces a slight pressure increase away from the centerbody. The radial extent of the centerbody is still increasing, so the pressure in this vicinity is slightly lower. This causes the boundary layer fluid to migrate toward the middle of the centerbody from both sides, forming a counterrotating vortex pair. As the flow expands through the converging portion of the facility nozzle, this vortex extends further into the core flow. The remnants of this vortex persists through the rest of the nozzle expansion process and leads to the build-up of boundary layer fluid noted at the $y = 0$ centerplane in the test section crossflow images. Note that this vortical structure is formed by spanwise pressure gradients rather than Reynolds stress anisotropies, so the linear eddy viscosity model used here can properly capture this features. Although not shown, an additional simulation was performed with the Spalart Quadratic Constitutive Relationship[30] to confirm this assertion.

SUMMARY

Pretest simulations have been performed to assess geometric sensitivities of a recently designed coaxial reacting jet experiment for CFD model validation. The concept for the experiment was derived from a validation data gap analysis performed in the High Speed Code Credibility Workshop series sponsored by the Air Force Research Laboratory. The approach taken for the CFD simulations took advantage of the *refine* computational workflow, which eliminates the human-in-the-loop grid generation process and adaptively refines the grid based on an error metric derived from each previous steady-state solution performed. A verification of the *refine* process was performed using a supersonic jet flow configuration that was the focus of the 6^{th} AIAA Propulsion Aerodynamic Workshop. The results from this verification exercise definitively demonstrated that solutions obtained on adapted anisotropic tetrahedral grids could replicate the “grid converged” jet flow results produced from handcrafted structured grids with significantly fewer grid cells. Having verified the *refine* workflow for an existing coaxial jet problem, attention was then focused on the AFRL reacting jet configuration. In particular, the impact of removing 2 geometrically complex hardware features was investigated. The 1^{st} modification involved removing most of the upstream fuel piping from the computational domain. The upstream fuel piping was the only geometrical feature preventing symmetry between the top and bottom halves of the facility flowpath, so if proved unimportant, the steady-state simulation costs are reduced by a factor of 2. The 2^{nd} geometrical simplification involved removing the leading portion of the coaxial fuel jet centerbody and all the strut support structure associated with it. All this hardware was placed upstream of the converging-diverging facility nozzle in an attempt to minimize any impact it may have on the test section flow structure. Neglecting the upstream fuel piping had an inconsequential effect on the coaxial jet flowfield, but the same could not be said when the upstream centerbody hardware artifacts were removed. Weak pressure gradients formed as the nozzle plenum flow traversed around the strut/centerbody assembly, producing counterrotating vortex pairs on each side of the centerbody. These vortex structures persisted through the nozzle expansion process, and led to a slight build-up of boundary layer fluid along the centerbody at the facility centerplane locations. Regardless of whether the upstream support structure was included or not, the wake region downstream of the centerbody did not maintain azimuthal symmetry. In general, the radial extent of the wake structure was smaller along the $y = 0$ facility centerplane as compared to the $z = 0$ centerplane. Some level of asymmetry is inevitable when an axisymmetric body is placed within a rectangular facility flow, and this concession was necessary to prevent problems with optical distortion for the nonintrusive measurements planned.

ACKNOWLEDGMENTS

This effort was funded through the NASA Hypersonics Technology Project and carried out in the Hypersonic Airbreathing Propulsion Branch at the NASA Langley Research Center. Computational resources for this work were provided by the NASA Langley Research Center (Hypersonic Airbreathing Propulsion Branch beowulf cluster, Hypersonics Technology Project steppenwulf cluster, and the Langley K cluster). A special thanks is extended to Nicholas Georgiadis and Vance Dippold for providing their PAW6 computational results obtained using the WIND-US and FUN3D flow solvers.

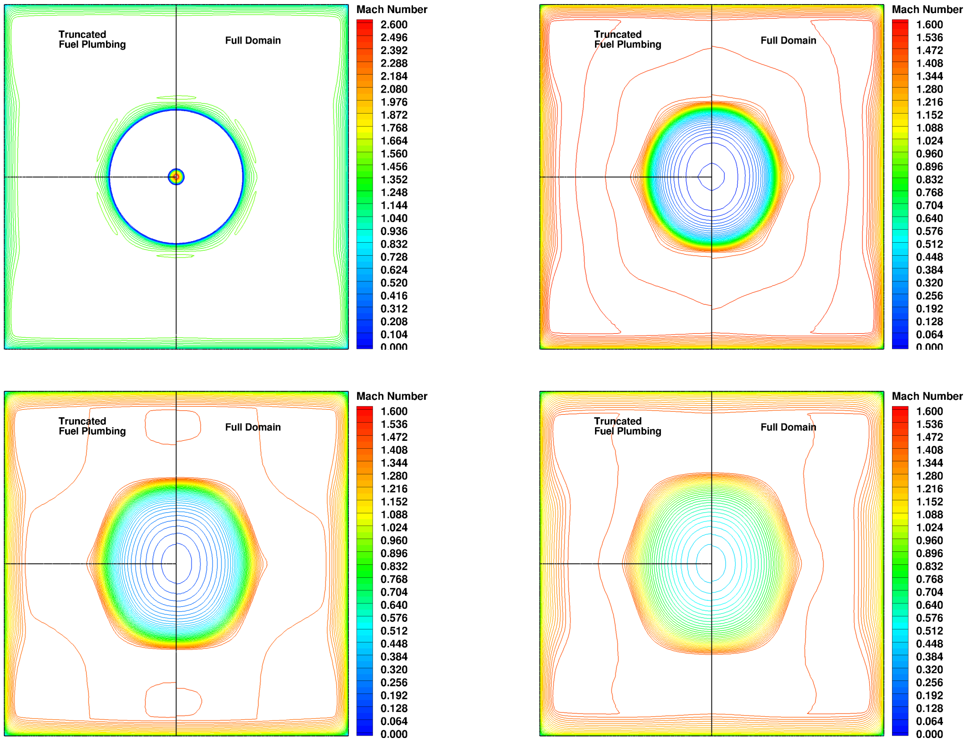

Figure 14: Mach number comparison for 1^{st} alternative geometry: fuel injection plane (top left), 2 inch station (top right), 4 inch station (bottom left), 8 inch station (bottom right).

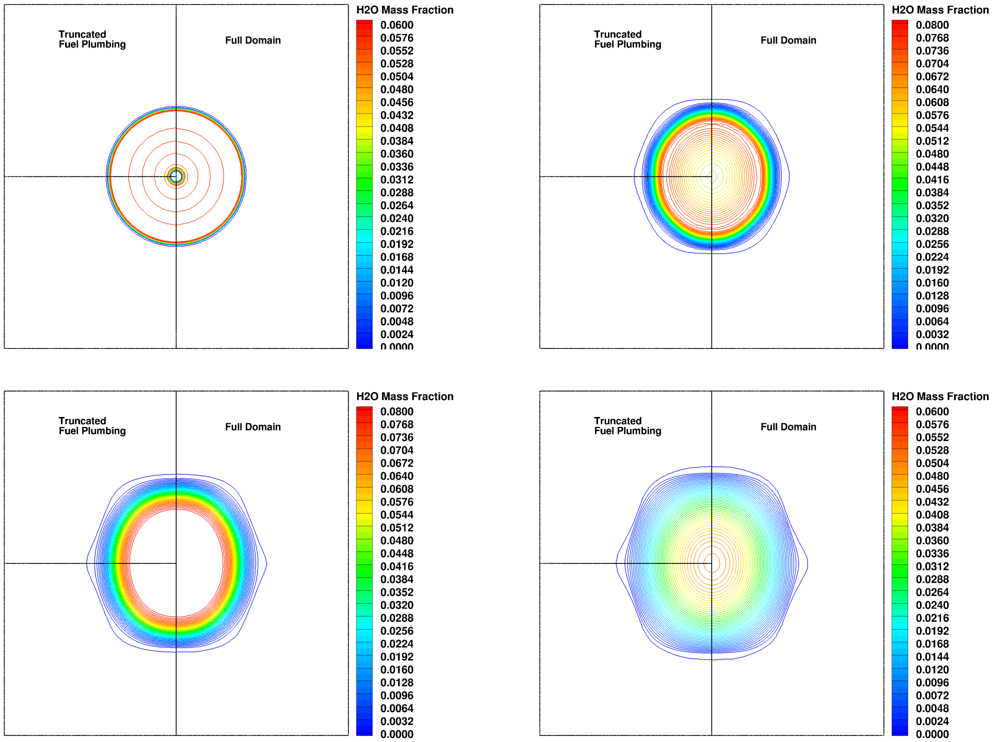

Figure 15: H_2O mass fraction comparison for 1^{st} alternative geometry: fuel injection plane (top left), 2 inch station (top right), 4 inch station (bottom left), 8 inch station (bottom right).

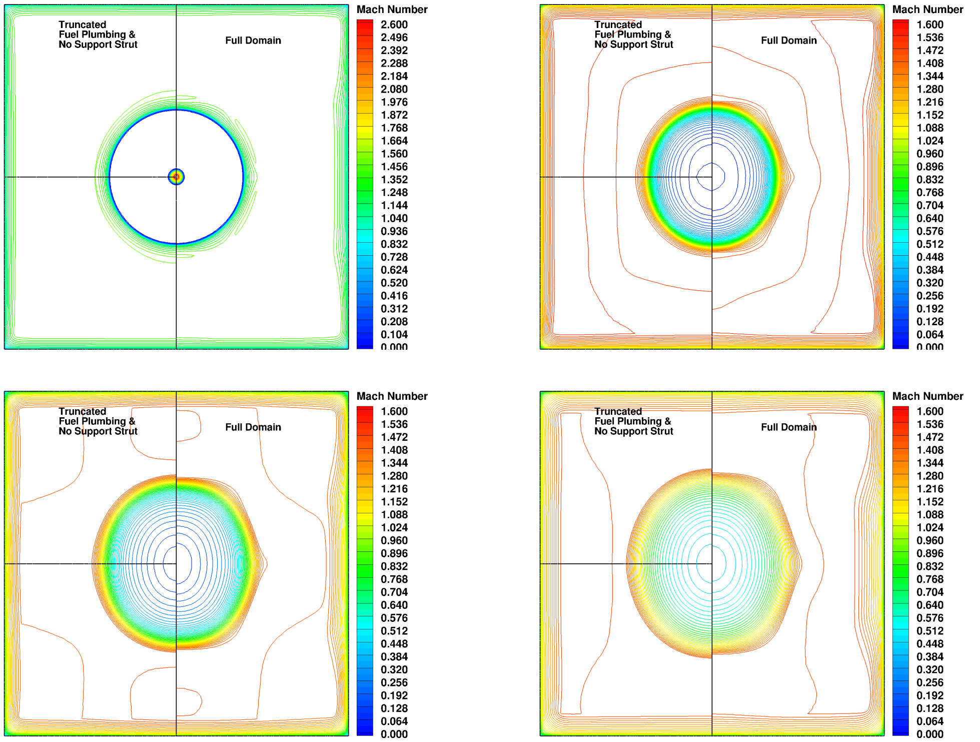

Figure 16: Mach number comparison for 2^{nd} alternative geometry: fuel injection plane (top left), 2 inch station (top right), 4 inch station (bottom left), 8 inch station (bottom right).

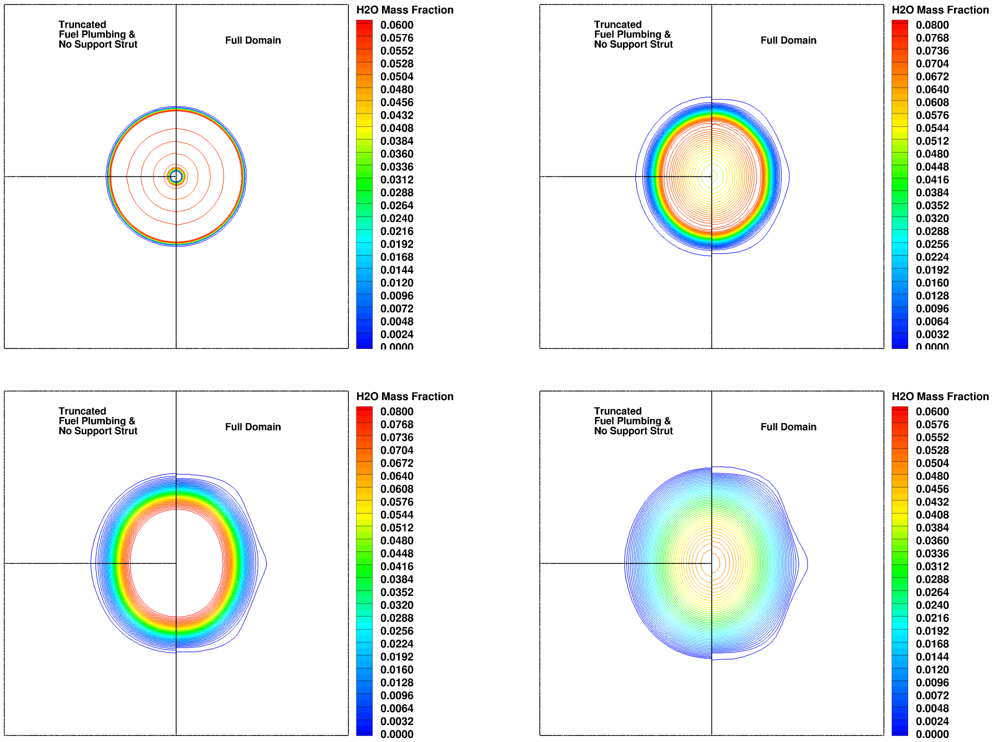

Figure 17: H_2O mass fraction comparison for 2^{nd} alternative geometry: fuel injection plane (top left), 2 inch station (top right), 4 inch station (bottom left), 8 inch station (bottom right).

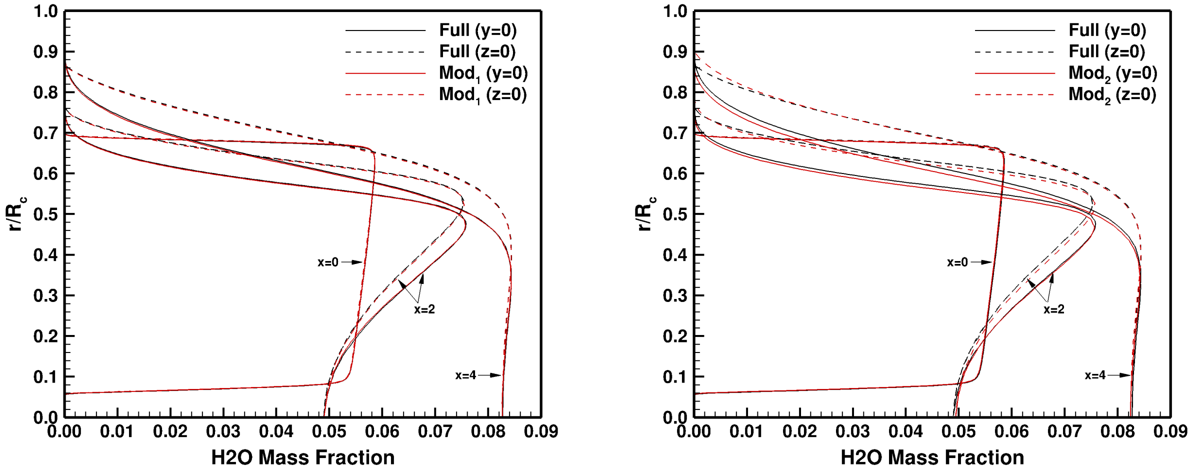

Figure 18: Radial H_2O mass fraction distributions: 1^{st} alternative geometry (left), 2^{nd} alternative geometry (right).

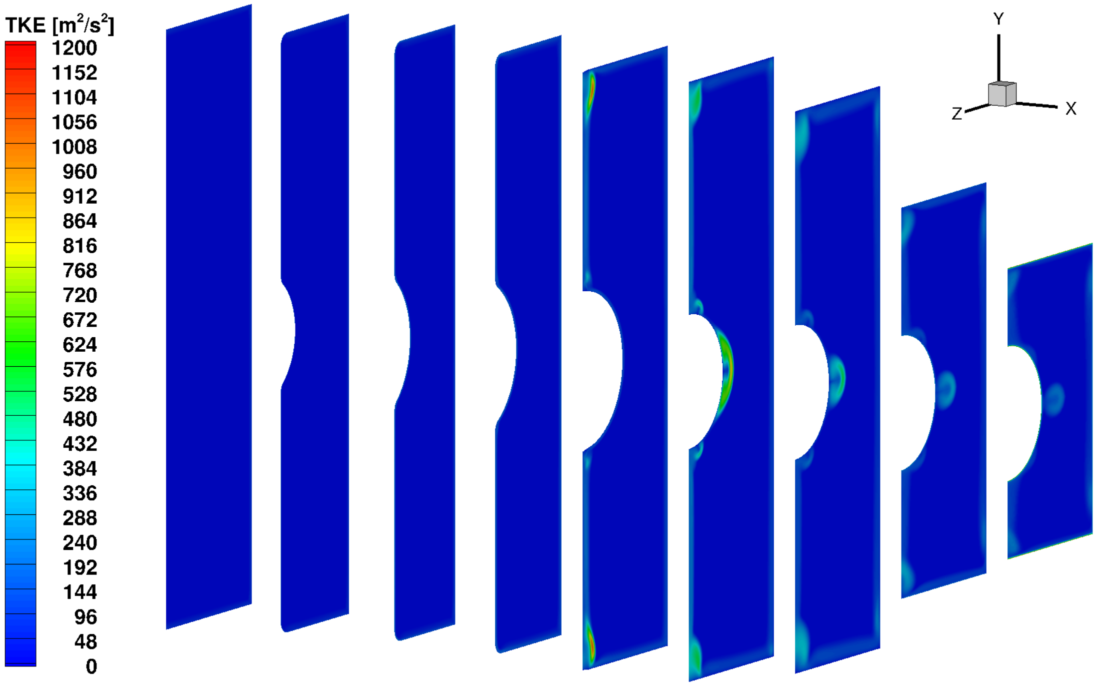

Figure 19: Evolution of turbulence kinetic energy over the centerbody.

References

[1] Selesnev, R. K., Surkzhikov, S. T., and Shang, J. S., ***A Review of the Scramjet Experimental Database***, Progress in Aerospace Sciences, 106:43–70 (2019).

[2] Cutler, A. D., Danehy, P. M., Springer, R. R., O'Byrne, S., Capriotti, D. P., and Deloach, R., ***Coherent Anti-Stokes Raman Spectroscopic Thermometry in a Supersonic Combustor***, AIAA Journal, 41(12):2451–2459 (Dec. 2003).

[3] O'Byrne, S., Danehy, P. M., and Cutler, A. D., ***Dual-Pump CARS Thermometry and Species Measurements in a Supersonic Combustor***, AIAA Paper 2004-0710 (Jan. 2004).

[4] Tuttle, S., Hsu, K.-Y., and Carter, C. D., ***Particle Image Velocimetry in a Nonreacting and Reacting High-Speed Cavity***, Journal of Propulsion and Power, 30(3):576–591 (May-June 2014).

[5] Voland, R. T., Huebner, L. D., and McClinton, C. R., ***X-43A Hypersonic Vehicle Technology Development***, in *56th International Astronautical Congress of the International Astronautical Federation, the International Academy of Astronautics, and the International Institute of Space Law* (2012).

[6] Hank, J. M., Murphy, J. S., and Mutzman, R. C., ***The X-51A Scramjet Engine Flight Demonstration Program***, AIAA Paper 2008-2540, AIAA 7th International Spaceplanes and Hypersonics Systems and Technology Conference (2008).

[7] Burrows, M. C. and Kurkov, A. P., ***Analytic and Experimental Study of Supersonic Combustion of Hydrogen in a Vitiated Airstream***, NASA Technical Report TM-X-2828 (1973).

[8] Cheng, T. S., Wehrmeyer, J. A., Pitz, R. W., Jarrett, O. Jr., and Northam, G. B., ***Raman-Measurement of Mixing and Finite-rate Chemistry in a Supersonic Hydrogen-Air Diffusion Flame***, Combustion and Flame, 99:157–173 (1994).

[9] Oberkampf, W. L. and Smith, B., ***Assessment Criteria for Computational Fluid Dynamics Validation Benchmark Experiments***, AIAA Paper 2014-0205 (Jan. 2014).

[10] DeBardelaben, C. J., Helm, C., Ochs, B. A., Hammack, S. D., and Peterson, D. M., ***Design of a Supersonic Centerbody Jet-Flame Apparatus for Code Validation***, 2017 JANNAF JPM/CS/APS/EPSS/PSHS Joint Meeting (Dec. 2017).

[11] ***VULCAN***, vulcan-cfd.larc.nasa.gov (2023).

[12] Edwards, J. R., ***A Low Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations***, Computers & Fluids, 26(6):635–659 (1997).

[13] van Leer, B., ***Towards the Ultimate Conservation Difference Scheme. II. Monotonicity and Conservation Combined in a Second Order Scheme***, Journal of Computational Physics, 14:361–370 (1974).

[14] Hasselbacher, A. and Blazek, J., ***On the Accurate and Efficient Discretization of the Navier-Stokes Equations on Mixed Grids***, AIAA Paper 98-0612 (Jan. 1998).

[15] Menter, F. R., ***Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications***, AIAA Journal, 32(8):1598–1605 (1994).

[16] Park, M. A., Kleb, B., Jones, W. T., Krakos, J. A., Michal, T., Loseille, A., Haimes, R., and Dannenhoffer, J. F. III, ***Geometry Modeling for Unstructured Mesh Adaptation***, AIAA Paper 2019-2946 (June 2019).

[17] Haimes, R. and Drela M., ***On the Construction of Aircraft Conceptual Geometry for High-Fidelity Analysis and Design***, AIAA Paper 2012-0683 (Jan. 2012).

[18] Wernet, M. P., Georgiadis, N. J., and Locke, R. J., ***Velocity, Temperature and Density Measurements in Supersonic Jets***, AIAA Paper 2021-0596 (Jan. 2021).

[19] Georgiadis, N. J., Wernet, M. P., Locke, R. J., and Eck, D. G., ***Mach Number and Heating Effects in Supersonic Jets***, AIAA Paper 2021-2834 (Aug. 2021).

[20] Dippold, V. F., ***Generating a Grid for Unstructured RANS Simulations of Jet Flows***, AIAA Paper 2018-3223 (June 2018).

[21] Wilcox, D. C., ***Wall Matching, a Rational Alternative to Wall Functions***, AIAA Paper 89-0611 (Jan. 1989).

[22] Dannenhoffer, III, J. F., ***OpenCSM: An Open-Source Constructive Solid-Modeler for MDAO***, AIAA Paper 2013-0701 (Jan. 2013).

[23] ***TetGen***, http://tetgen.org (2023).

[24] Sarkar, S., Erlebacher, G., Hussaini, M. Y., and Kreiss, H. O., ***The Analysis and Modeling of Dilatational Terms in Compressible Turbulence***, ICASE Report 89-79 (1989).

[25] Wilcox, D. C., ***Turbulence Modeling for CFD***, DCW Industries, Inc., 3rd edn. (2006).

[26] Menter, F. R., Kuntz, M., and Langtry, R., ***Ten Years of Industrial Experience with the SST Turbulence Model***, in *Proceedings of the Fourth International Symposium on Turbulence, Heat and Mass Transfer*, pages 625–632, Begell House, Inc. (2003).

[27] Zilberter, I. A. and Edwards, J. R., ***Large-Eddy Simulation / Reynolds-Averaged Navier-Stokes Simulations of High-Speed Mixing Processes***, AIAA Journal, 52(7):1486–1501 (2014).

[28] Baurle, R. A., ***Hybrid Reynolds-Averaged / Large Eddy Simulation of the Flow in a Model Scramjet Cavity Flameholder***, AIAA Journal, 48(3):544–560 (2016).

[29] Hassan, E., Peterson, D. M., Lieu, J., and Hagenmaier, M. A., ***Reacting Hybrid Reynolds-Averaged Navier-Stokes/Large-Eddy Simulation of a Supersonic Cavity Flameholder***, AIAA Paper 2016-4566 (2016).

[30] Spalart, P. R., ***Strategies for Turbulence Modelling and Simulation***, International Journal of Heat and Mass Transfer, 21:252–263 (2000).