Hybrid Reynolds-Averaged / Large Eddy Simulation of the Flow in a Model SCRamjet Cavity Flameholder

R. A. Baurle [*]

NASA Langley Research Center, Hampton, Va 23681

Steady-state and scale-resolving simulations have been performed for flow in and around a model scramjet combustor flameholder. Experimental data available for this configuration include velocity statistics obtained from particle image velocimetry. Several turbulence models were used for the steady-state Reynolds-averaged simulations which included both linear and non-linear eddy viscosity models. The scale-resolving simulations used a hybrid Reynolds-averaged / large eddy simulation strategy that is designed to be a large eddy simulation everywhere except in the inner portion (log layer and below) of the boundary layer. Hence, this formulation can be regarded as a wall-modeled large eddy simulation. This effort was undertaken to not only assess the performance of the hybrid Reynolds-averaged / large eddy simulation modeling approach in a flowfield of interest to the scramjet research community, but to also begin to understand how this capability can best be used to augment standard Reynolds-averaged simulations. The numerical errors were quantified for the steady-state simulations, and at least qualitatively assessed for the scale-resolving simulations prior to making any claims of predictive accuracy relative to the measurements. The steady-state Reynolds-averaged results displayed a high degree of variability when comparing the flameholder fuel distributions obtained from each turbulence model. This prompted the consideration of applying the higher-fidelity scale-resolving simulations as a surrogate "truth" model to calibrate the Reynolds-averaged closures in a non-reacting setting prior to their use for the combusting simulations. In general, the Reynolds-averaged velocity profile predictions at the lowest fueling level matched the particle imaging measurements almost as well as was observed for the non-reacting condition. However, the velocity field predictions proved to be more sensitive to the flameholder fueling rate than was indicated in the measurements.

Nomenclature

Symbols

d	distance to the nearest solid surface
d^+	non-dimensional "law of the wall" coordinate
F	hybrid RAS/LES blending function
F_s	factor of safety used to assess grid convergence
f	generic functional used to assess grid convergence
k	turbulence kinetic energy
$\dot{m}$	mass flow rate
P	pressure
p	numerical scheme order of accuracy
Pr_t	turbulent Prandtl number
r	grid refinement factor
S_{ij}	mean strain rate tensor
S_{ij}^*	traceless mean strain rate tensor
Sc_t	turbulent Schmidt number
T	temperature
u, v, w	Cartesian velocity components
x, y, z	Cartesian coordinates
Y_m	Mass fraction of species "m"
δ	boundary layer thickness

---

[*]Aerospace Engineer, Hypersonic Airbreathing Propulsion Branch, Associate Fellow AIAA.

American Institute of Aeronautics and Astronautics

δ_{ij}	Kronecker delta
ϕ	fuel-air equivalence ratio
κ	MUSCL interpolation parameter
μ	molecular viscosity
μ_t	eddy viscosity
ρ	density
τ_{ij}	Reynolds stress tensor
τ_{ij}^*	traceless Reynolds stress tensor
ω	specific turbulence dissipation rate
∇	gradient operator
$\triangle$	SGS filter width or grid spacing

Subscripts

| i, j | coordinate indices |

Introduction

Reynolds-averaged Computational Fluid Dynamics (CFD) codes are the standard high-fidelity numerical tools utilized in the aerospace industry. These tools have revolutionized research and development practices, which as early as 15-20 years ago relied almost exclusively on extensive wind tunnel testing. Unfortunately, Reynolds-Averaged Simulations (RAS), which attempt to model all of the scales present in turbulent flows, have proven to be deficient in many challenging areas of interest to the aerospace community, *e.g.*:

- high lift devices (massive flow separation)

- combusting flows (particularly lean or rich flames near extinction)

- unsteady flows (rotorcraft, aeroacoustics, etc.)

- shock / boundary layer and shock / jet interactions

In many applications of engineering interest, the limitations associated with the turbulence closure models are the pacing items preventing the use of RAS in these (and other complex) settings as a true predictive tool.[1]

Large Eddy Simulation (LES) methods have the potential to reduce the modeling uncertainty inherent to RAS approaches, since the intent of LES is to resolve the large scale turbulent structures while modeling only the smallest scales. However, the computational expense of wall-resolved LES at Reynolds numbers relevant to engineering problems of interest is well beyond what can be deemed as practical. Hybrid RAS/LES approaches[2–5] have emerged to address this issue, and have provided a rational path forward towards extending LES into practical settings. These methodologies allow LES content to be resolved in areas that require a rigorous modeling approach, while maintaining a more cost effective RAS approach for benign regions of the flow (*e.g.* attached boundary layers).

The computational expense required for hybrid RAS/LES, while less than that of a full wall-resolved LES, is still formidable when compared with steady-state RAS. Moreover, the numerical algorithms required to partially resolve the turbulence scales present must have low numerical dissipation properties with minimal dispersive errors (particularly for high-speed flows where shock waves may be present). These observations demand an efficient high-order, low-dissipation numerical framework; a feature not typically required by pure RAS solvers. As a result, the extension of scale-resolving simulation approaches (such as hybrid RAS/LES) to engineering problems of interest will only become practical when substantial advancements to both the numerical and physical models have been realized. Towards this end, researchers in the Hypersonic Airbreathing Propulsion Branch (HAPB) and Computational AeroSciences Branch (CASB) of the NASA Langley Research Center have been extending the capabilities of the VULCAN-CFD flow solver[6,7] for scale-resolving simulations. Advancements made to the numerical framework[8] and the hybrid RAS/LES framework[9] have been documented for both low and high-speed benchmark flows.

The present effort is a continuation of the scramjet combustor flameholder simulations documented in Ref. 9 for a supersonic internal flow passing over a recessed cavity that is internally fueled with ethylene through an array of injection ports. Experimental data available for validation include two components of velocity (via particle image velocimetry) taken at the centerplane of the cavity flameholder.[10] Previous computational efforts[9] for this configuration considered only aerodynamic flowfield simulations that neglected the mixing processes associated with fuel injection. The present effort extends this work by considering simulations that account for fuel injection with and without combustion. The results obtained from non-reacting simulations are presented first to compare the steady-state RAS

data with the hybrid RAS/LES statistics. Experimental data was not available at non-reacting conditions, so this comparison allowed the higher fidelity hybrid RAS/LES solutions to serve as a surrogate "truth" model to assess the performance and tune the constants of the turbulent mixing sub-models used for the Reynolds-averaged simulations. These calibrated settings were used for all Reynolds-averaged combusting simulations performed for comparison with measurements.

Geometry Description and Flow Conditions

A schematic of the facility flowpath considered in this study is shown in Fig. 1. An asymmetric facility nozzle provides a continuous flow of Mach 2 (nominal) air to the constant area isolator section of the flowpath. The 7 inch long isolator starts at the $x = 0$ inch station and has a 2 inch height (y-direction) and a 6 inch width (z-direction). The 2.5° divergent portion of the lower wall initiates at the exit of the isolator ($x = 7$ inch station), and the cavity flameholder is located 3 inches downstream of this location. The cavity spans the entire width of the flowpath. The depth of the cavity is 0.65 inches and the length of the cavity floor is 1.815 inches. The aft wall of the cavity has a shallow angle (22.5° relative to the cavity floor) and houses an array of 11 evenly distributed fuel injection ports. The diameter of each injection port is 0.0078 inches, and the centerline of each port intersects the angled aft cavity surface at the x, y coordinate values of 12.0276 and 0.7703 inches, respectively. The cavity closeout location corresponds to the intersection of the cavity aft wall with the 2.5° facility lower wall surface.
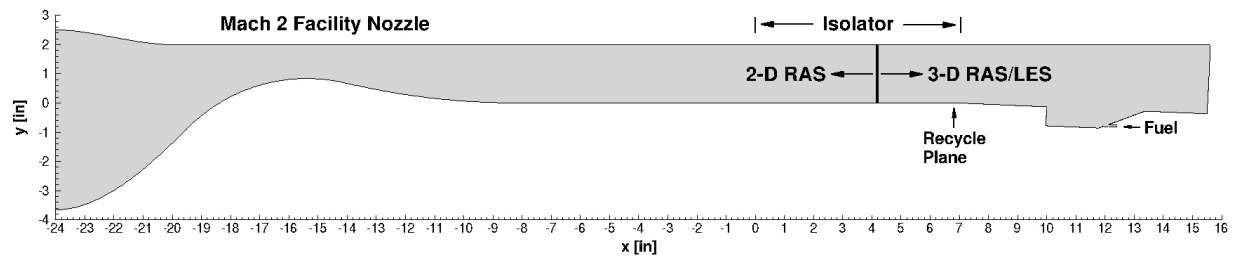

Figure 1: Schematic of the facility flowpath

Table 1: Facility Test Conditions

Nominal Conditions	Case 1	Case 2	Case 3
Air Mach Number	2.0^a	2.0^a	2.0^a
Air $T_\circ$ [K]	589.0	589.0	589.0
Air $P_\circ$ [kPa]	483.0	483.0	483.0
Fuel Flow Rate [SLPM]	0.0	56.0	99.0
Fuel $T_\circ$ [K]	N/A	310.0^b	310.0^b

^{a} facility nozzle design Mach number
^{b} estimated ethylene bottle temperature

The nominal flow conditions considered in the experiments conducted by Tuttle et al.[10] are given in Table 1. The fuel flow rates, reported in Standard Liters Per Minute (SLPM), are based on a reference state of 273 K and 1 atmosphere. The facility nozzle geometry was included in the simulated flowpath, so the Mach 2.0 value listed in Table 1 is a nominal value. All of the simulations performed in this work ignored any facility side-wall influences. This allowed the exit conditions from an *a priori* two-dimensional RAS of the facility nozzle (and 4.2 inches of the constant area isolator) to be used as the inflow condition for the three-dimensional simulations performed for the region of interest further downstream. The Mach number and pressure profiles extracted at this interface (see Fig. 1) are shown in Fig. 2. The simulation of the facility nozzle shows a weak shock structure propagating through the isolator as evidenced by the pressure variability along the profile. The boundary layer thickness at this station is 0.25 inches (1/4 of the duct half height) which is nearly 40% of the cavity depth.

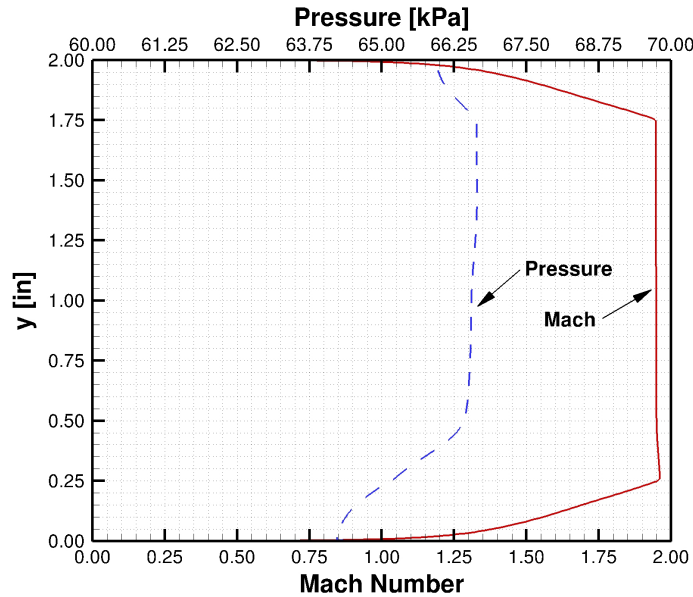American Institute of Aeronautics and Astronautics

Figure 2: Facility Mach number and pressure profiles at the $x = 4.2$ inch station used to prescribe the 3-D RAS and hybrid RAS/LES inflow conditions

Computational Grid Description

As mentioned previously, the simulations performed for this effort neglected any facility side-wall influences. This allowed the consideration of only a fraction of the facility width, providing for a more efficient means of examining numerical modeling sensitivities while retaining the salient flow features. The smallest spanwise width allowed by the symmetry present in the geometry is a 0.25 inch section that covers the region from the centerplane of an injection port to the centerplane between adjacent injector ports. This domain is appropriate for steady-state RAS, with symmetry conditions imposed at each bounding plane. The unsteady hybrid RAS/LES approach demands the use of periodic spanwise boundaries in lieu of symmetry conditions, so these simulations require (at a minimum) a domain that extends from the centerplane of one injector port to the centerplane of an adjacent port (or equivalently one that is bounded by the gap centerplane on either side of an injection port). However, as pointed out previously, the approach boundary layer has a thickness of 0.25 inches which will eventually transition to an even thicker free shear layer over the cavity flameholder. In order to properly accommodate the formation of resolved eddy structures, the largest of which are comparable to the thickness of the viscous shear layer, a domain width larger than 0.5 inches should be considered. More specifically stated, the width should be large enough to ensure that the largest eddy structures become decorrelated when separated by a distance larger than the half-width of the periodic domain. This prompted the consideration of a 1-inch domain width that encompasses two full injection ports. This domain matches that used in previous DES research efforts[11, 12] for this configuration.

A sequence of three progressively refined structured grids was created for the simulations. The refinement factor used was 1.5 in each coordinate direction. The grid attributes closely resembled those utilized for the DES studies performed in Refs. 11 and 12. However, the computational grids generated for this study were fully structured, while the grids in the aforementioned references made use of unstructured hexahedral cells in some regions of the flow. Each grid is comprised of nearly isotropic cells (except in the inner part of the boundary layers where RAS is utilized), and the grid is clustered to all surfaces such that d^+ is approximately unity. The isotropy attribute is driven by the desire to resolve eddy structures in the outer portion of the boundary layers and in free shear regions of the flow, while a d^+ wall spacing near unity is required for accurate RAS modeling of the boundary layers without resorting to the use of wall functions. The average grid spacing and wall distance (normalized by the incoming boundary layer thickness), average and maximum d^+, and the total number of cells associated with each grid generated for the RAS domain are listed in Table 2. The hybrid RAS/LES grids were obtained by duplicating and mirroring the RAS grids which results in a cell count that is 4 times larger than those listed in Table 2. Several images of the grid are shown in Fig. 3. Topological nesting was used in several areas to help maintain a nearly constant grid spacing as the cross-sectional area of the domain varied, and a non-trivial topology was implemented in the vicinity of the fuel injectors to provide additional

American Institute of Aeronautics and Astronautics

grid resolution for the smaller flow structures present when fuel is injected into the cavity. Also shown in this figure is the location of the recycling plane used to generate resolved turbulent content for the inflow of the scale-resolving simulations.

Table 2: Grid Attributes (RAS domain)

Attribute	Coarse	Medium	Fine
$\triangle / \delta$	0.063	0.042	0.028
d_1 / δ	0.0004	0.0004	0.0004
d^+ (ave, max)	0.50, 0.85	0.50, 0.92	0.50, 0.99
Total Cell Count	2,540,336	7,867,136	24,585,120

Reynolds-Averaged Simulation Settings

All of the Reynolds-averaged simulations were advanced in pseudo-time using an Incomplete LU(0) factorization scheme[13] with a Courant-Friedrichs-Lewy (CFL) number of 100. The inviscid fluxes were evaluated using the Low-Dissipation Flux Split Scheme of Edwards[14] with 3^{rd}-order cell interface variable reconstruction achieved via the κ=1/3 Monotone Upstream-centered Scheme for Conservation Laws (MUSCL). The van Leer flux limiter[15] was utilized to avoid spurious oscillations during this reconstruction process. The viscous fluxes were evaluated using 2^{nd}-order accurate central differences with the constituent viscosities and conductivities computed from the polynomial fits of McBride.[16,17] The RAS turbulence models considered were the Menter-SST k-ω model[18] and the Gatski-Rumsey Explicit Algebraic Stress (Gatski-EAS) model.[19] The Menter-SST model is a Linear Eddy Viscosity (LEV) model which assumes a direct proportionality between the turbulent stress tensor and the mean strain rate tensor (*i.e.* the Boussinesq approximation). The EAS model incorporates additional nonlinear functionals in the constitutive relationship for the stress tensor that depend on both the strain and rotation tensors, allowing the model to capture some aspects of the turbulent flow structure that standard LEV models can't (such as the effects of Reynolds stress anisotropies). The turbulent Prandtl (Pr_t) and Schmidt (Sc_t) numbers, which control the turbulent transport of energy and mass, were initially chosen as 0.9 and 0.5, respectively. However, as discussed below, the cavity fuel distribution was quite sensitive to the values chosen for turbulent Schmidt number. Hence, a sensitivity study was performed using Sc_t values that were half and twice the default value of 0.5.

Solution convergence of the steady-state RAS was monitored by assessing the L_2 norm of the equation set residual error, the mass flow error, and the surface friction force time histories. A sample convergence history is given in Fig. 4, which shows the normalized mass flow rate error and the L_2 norm of the steady-state residual error for the medium grid Gatski-EAS simulation. At a minimum, the following iterative convergence statements were satisfied for each non-combusting simulation:

- The L_2 norm of the steady-state residual error was 3 to 4 orders of magnitude smaller than the initial value.

- The surface friction force time history remained unchanged to 5 digits over the final 2500 iteration cycles.

- The relative mass flow rate error, $|\dot{m}_{out} - \dot{m}_{in}| / \dot{m}_{in}$, was on the order of $\times 10^{-7}$.

Hybrid RAS/LES Simulation Settings

The hybrid RAS/LES methodology used in this effort is based on the framework originally developed in Ref. 2, with subsequent variants described in Refs. 20 and 21. This framework is designed to enforce a RAS behavior near solid surfaces, and switch to an LES behavior in the outer portion of the boundary layer and free shear regions. Hence, this formulation can be thought of as a wall-modeled LES approach, where a RAS closure is used as the near-wall model. The basic idea is to blend the RAS eddy viscosity value with the LES SGS viscosity, along with any transport equation that involves a common RAS and SGS property. In this effort, the Menter-BSL k-ω RAS model[18] was blended with the one-equation SGS model of Yoshizawa.[22] The Yoshizawa model involves an evolution equation for the SGS turbulence kinetic energy, hence the blended expressions that are appropriate for this model combination are:

$$\text{Hybrid RAS/SGS viscosity} = (F)\left[\text{RAS viscosity}\right] + (1 - F)\left[\text{SGS viscosity}\right]$$
$$\text{Hybrid RAS/SGS k-equation} = (F)\left[\text{RAS k-equation}\right] + (1 - F)\left[\text{SGS k-equation}\right] \quad (1)$$
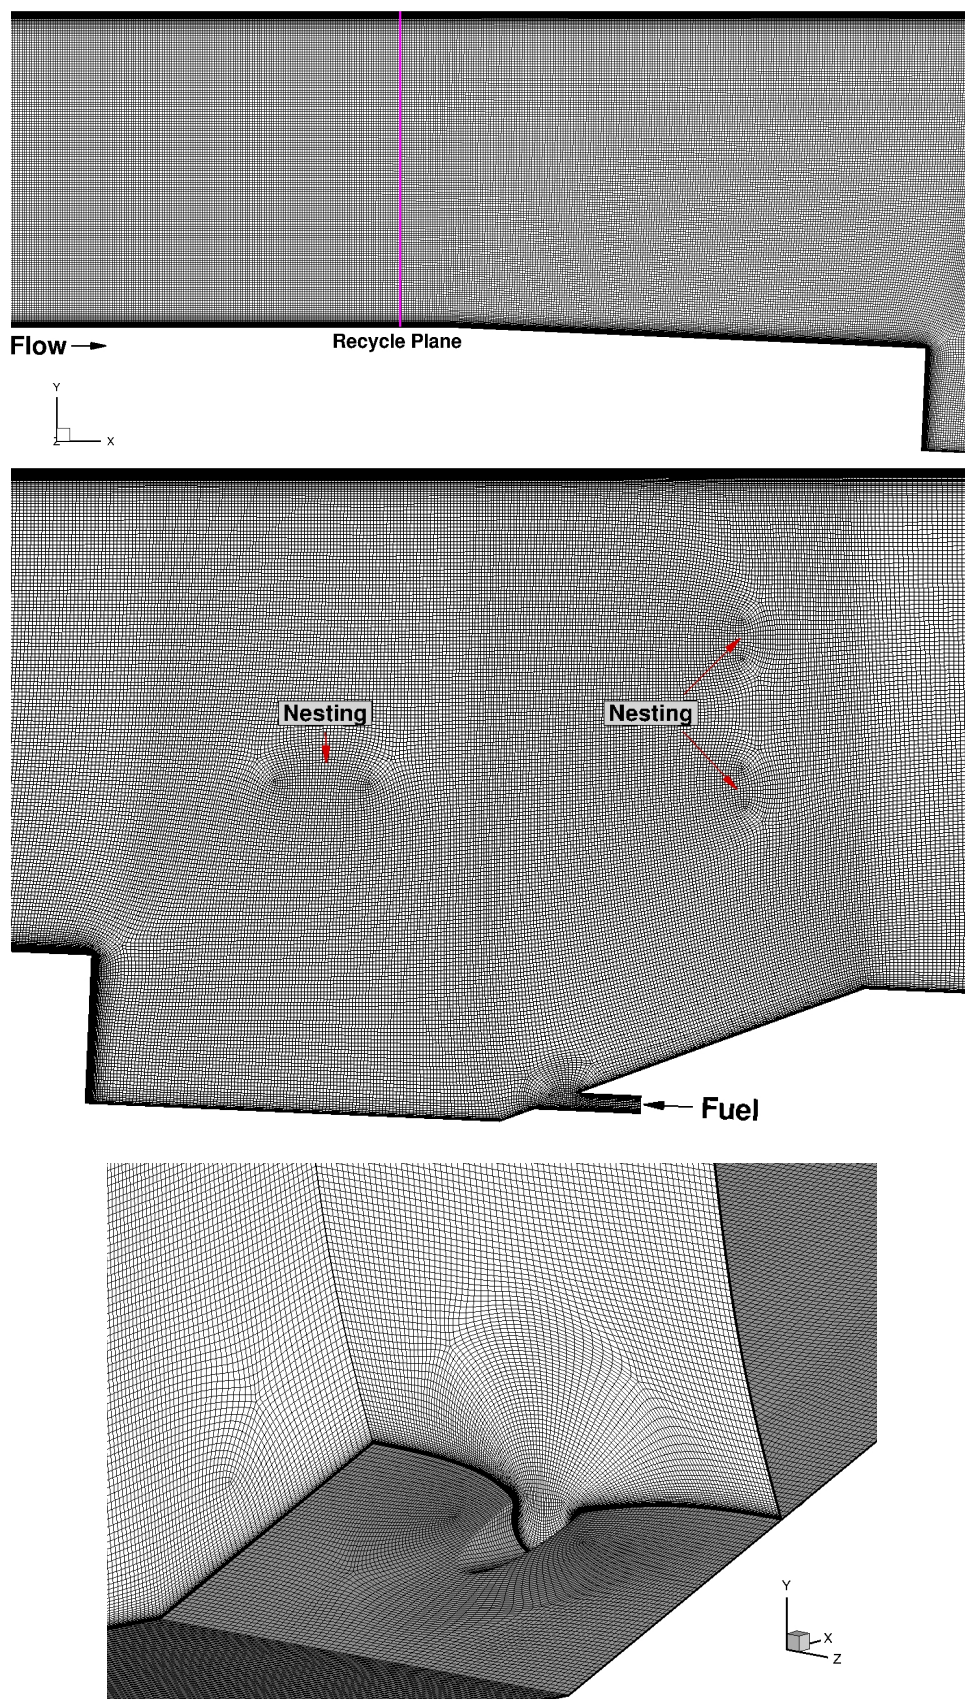
American Institute of Aeronautics and Astronautics

Figure 3: Grid images: overall view of the coarse grid (top 2 images), cavity fuel injection region (lower image)

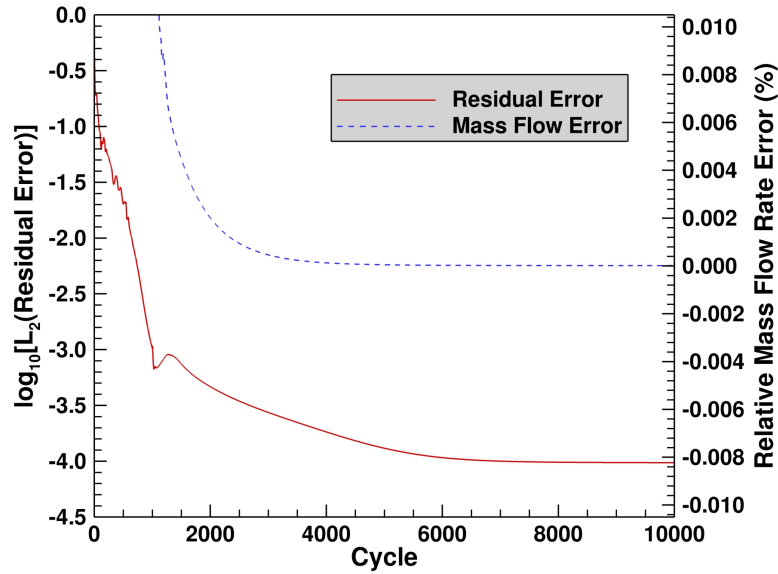American Institute of Aeronautics and Astronautics

Figure 4: Typical steady-state non-combusting RAS iterative convergence history

where F is a blending function[20] that varies between 0 and 1. Note that the transport equation for the RAS specific dissipation rate (ω) does not have an SGS counterpart. Hence, the blending is not applied to this equation, and the terms in this equation that involve the eddy viscosity, *i.e.* the turbulent diffusion terms, are evaluated based on the RAS relationships. Finally, the turbulent Prandtl (Pr_t) and Schmidt (Sc_t) numbers were both chosen as 0.9, resulting in a turbulent Lewis number of unity. Given that these parameters control only the small scale SGS transport (rather than the entire turbulent transport) of energy and mass, one would anticipate less variability in the scalar diffusion processes. Hence, the assumption of a unity Lewis number for SGS diffusion is expected to be a reasonable one.

The hybrid RAS/LES utilized the low-dissipation numerical framework described in Refs. 8 and 9 to encourage and maintain the development of resolved turbulent content. In this effort, the κ=1/3 MUSCL scheme was chosen as the dissipative reconstruction operator with the UNO limiter of Suresh and Huynh[23] used to suppress spurious oscillations during this reconstruction process. A simple 4^{th}-order symmetric reconstruction scheme was chosen for the non-dissipative operator. The final reconstructed states were obtained by blending the dissipative and non-dissipative states using a modified form of the Larsson sensor as described in Ref. 9. However, a small fraction (20%) of the dissipative reconstruction was required at all times for numerical stability when reconstructing the species composition variables. The viscous fluxes were evaluated in the same manner as that used for the Reynolds-averaged simulations described in the previous section. The time-accurate hybrid RAS/LES solutions were advanced in time using a dual time-stepping approach that combined a Diagonalized Approximate Factorization (DAF) scheme[24] for integration in pseudo-time, with a 3-point backwards finite difference approximation for integration in real-time. The values selected for the physical time-step and sub-iteration CFL constraint were 0.05 μs and 10.0, respectively. The time-step was chosen based on cell residence time considerations to ensure that turbulent structures would traverse less than one grid cell length per time-step. The sub-iteration process was carried out until the residual error dropped 2 orders of magnitude. This level of convergence typically required 7-10 sub-iterations for each physical time-step. The SGS filter width, $\triangle$, was defined as the maximum cell length, *i.e.* max($\triangle x, \triangle y, \triangle z$), of each computational cell. Finally, a recycling/rescaling procedure[25] was employed to provide an inflow condition with resolved turbulent content. Further details describing the effectiveness of these numerical settings can be found in previous simulation efforts for this flowpath.[9]

The initial state of the hybrid RAS/LES flowfield was defined using the converged RAS results together with artificial fluctuations added to the approach boundary layer flow to reduce the start-up time required to generate resolved turbulent content. The simulations were monitored as a function of time to assess the establishment of a statistically stationary state prior to gathering flowfield statistics (or analyzing any instantaneous flow properties). Particular emphasis was given to the mass flow error and the integrated friction force time histories. A sample time history is shown in Fig. 5. Based on these metrics, a statistically stationary state was deemed to be established after roughly 25,000 time-steps. This integration period corresponded to 10.6 flow-through times, defined as the time required for the freestream particles to traverse the length of the cavity. After this time period, flowfield statistics were

gathered for at least another 225,000 time-steps (95.4 flow-through times). The 3-fold increase in the integration time (compared to our previous aerodynamic simulations of this flowfield) was required to statistically converge the slow mass exchange process of the fuel air mixture within the recirculation zones of the cavity.
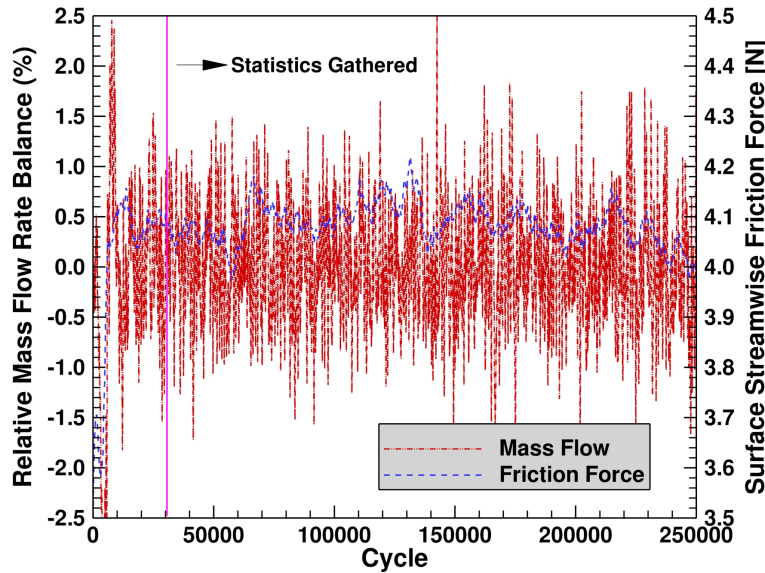

Figure 5: Typical non-combusting hybrid RAS/LES time history used to assess statistical stationarity

Non-Combusting Results

Figure 6 shows the time-averaged structure of the flow (via streamlines) inside the cavity flameholder as predicted by each class of turbulence model (LEV, EAS, and hybrid RAS/LES). Also shown are the streamlines extracted from time-averaged Particle Image Velocimetry (PIV) measurements taken by Tuttle *et al.*[10] The predicted flow structure is qualitatively similar to the measurements with a dominant clockwise rotating recirculation zone and a smaller counter-clockwise rotating recirculation zone near the front wall of the cavity. From a flameholding perspective, the primary recirculation zone provides the mass exchange between the core flow and the hot cavity combustion products. The secondary recirculation zone simply enhances the cavity flow residence time. The most notable difference between the model predictions is the size and shape of the secondary recirculation zone. The Gatski-EAS model and the hybrid RAS/LES model predict a larger and more elongated secondary bubble than that predicted by the Menter-SST model. This difference is likely due to Reynolds stress anisotropies that EAS models (or more generally, non-linear eddy viscosity models) and hybrid RAS/LES models are capable of capturing. The impact of capturing this feature will become evident in the subsequent discussions of profile properties extracted at various locations within the cavity. Figure 7 shows the specific locations where flowfield profiles have been extracted for detailed flow feature comparisons.

A quantitative solution verification procedure was performed to assess the numerical error for the outputs of interest extracted from each of the RAS data sets prior to any assessments being made to address the model-form uncertainties. The numerical errors in this effort were quantified using the Grid Convergence Index (GCI).[26] The GCI is a grid convergence estimator derived from the generalized Richardson extrapolation formula and can be written as follows:

$$GCI = F_s \frac{|f_1 - f_2|}{r^p - 1} \tag{2}$$

where f is some flow parameter of interest evaluated at two different grid resolutions (f_1 and f_2), r is the grid refinement ratio, and p is order of accuracy of the numerical scheme. Finally, F_s is a safety factor with recommended values taken to be 3 (as suggested by Roache[26]). It should be emphasized that the GCI, while based on Richardson Extrapolation, is not meant to be a "best estimate" of the numerical error. Instead, the intent is to provide a reasonable bound on the discretization error. The solution verification assessments were performed independently for each RAS turbulence model employed, however the GCI assessments in this study were not found to be strongly influenced by the choice
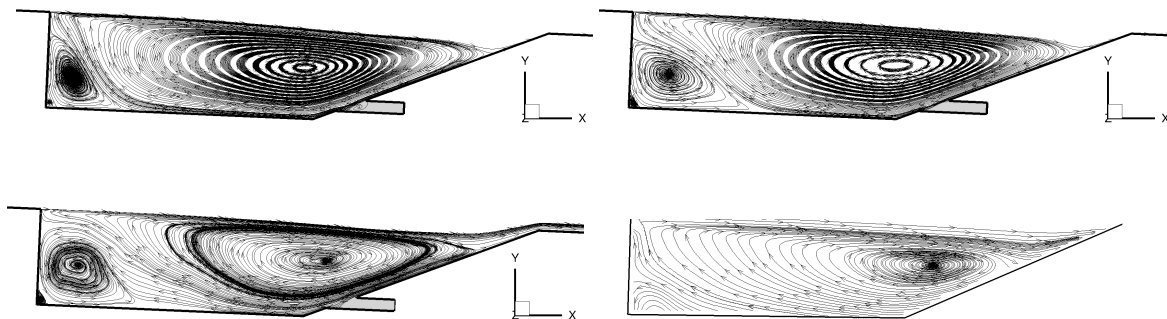
American Institute of Aeronautics and Astronautics

Figure 6: Streamlines extracted from measurements (bottom-right) and streamlines extracted from Case 1 simulations: Menter-SST (top left), Gatski-EAS (bottom-left), and hybrid RAS/LES (top right)

of turbulence model. Hence, for brevity, only the analysis for the Gatski-EAS model simulations are shown in Fig. 8. In this figure, the GCI measure of the numerical error is displayed as an error bar attached to the fine grid results, and should be interpreted as a bounding estimate of the error due to finite grid resolution. The error assessment for the fuel to air equivalence ratio (ϕ), which is an important parameter for flameholding effectiveness, was performed at every other profile station shown in Fig. 7. For this assessment, the equivalence ratio was spatially averaged in the spanwise (z) direction at each axial station to provide an overall integral value at each profile station. The analysis revealed the largest errors to appear in the secondary recirculation zone within the cavity, and near the injector exit (where the fuel plume is mixing rapidly with the surrounding cavity fluid). The observation of larger errors in the secondary recirculation zone is not all that surprising when one considers that the equivalence ratio in this region is governed by the mass exchange process between two recirculation zones. The numerical error is considerably smaller in the primary recirculation zone and in the shear layer that controls the mass exchange process with the supersonic air stream. Moreover, the numerical error is substantially smaller than the differences that appear due to modeling choices, which will become apparent in the results that follow.
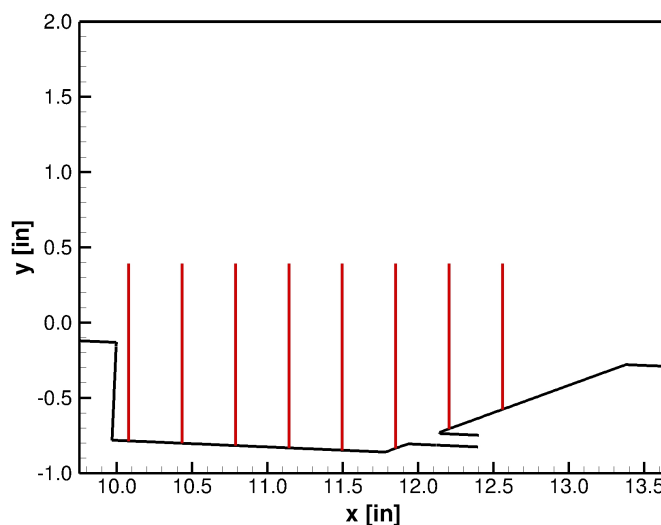

Figure 7: Locations within the cavity where profiles have been extracted

The equivalence ratio sensitivity to the choice of turbulence model and turbulent Schmidt number is shown in the contour images displayed in Fig. 9 for the highest fuel flow rate condition (Case 3). These images were extracted at the spanwise location that intersects the center of the cavity fuel injector port. In this figure, the images on the left were collected from the Menter-SST simulations (performed with three Sc_t values) and the images on the right

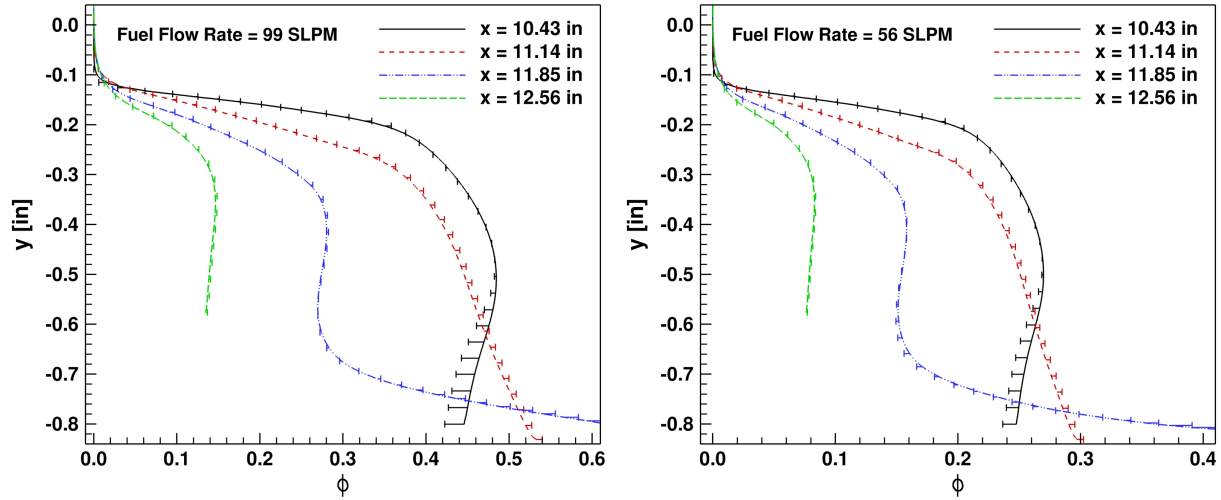American Institute of Aeronautics and Astronautics

Figure 8: Fuel-Air equivalence ratio profiles within the cavity with error bars depicting the GCI bounding estimate for the numerical error

are the corresponding Gatski-EAS results. Focusing first on the effect of Sc_t value, one immediately notices that this parameter has a significant effect on the fuel-air distribution within the cavity. The use of a low Sc_t value rapidly mixes the injected fuel and results in very uniform fuel-air distribution over most of the cavity area. This rate of mixing allows the mass exchange between the primary and secondary recirculation zones to equilibrate so that one can hardly recognize that a secondary recirculation region is even present. Simulations performed with an Sc_t value of 1.0 result in a more non-premixed fuel-air distribution within the cavity. The choice of turbulence model also has a significant effect on the fuel-air distribution within the flameholder. In general, the Gatski-EAS model predictions show a slower mass exchange (*i.e.* slower diffusion rates) between the cavity flameholder and the supersonic core flow (for a given Sc_t value) than that predicted by the Menter-SST model. This results in a higher overall equivalence ratio within the cavity flameholder for the Gatski-EAS model. This observation clearly demonstrates that the optimal choice for the turbulent Schmidt number will be dependent on the turbulence model employed.

A more quantitative assessment of the equivalence ratio sensitivity to the choice of turbulence model and Sc_t value is shown in the equivalence ratio profiles displayed in Fig. 10. These profiles were extracted at the spanwise location that corresponds to the midway point between adjacent fuel injector ports. The selected streamwise stations are identical to those used for the numerical error assessment. Again, these images show a significant level of variability to choices made to the physical sub-models, which was expected given that the region of interest is a large separated flowfield. In general, higher Sc_t values result in higher equivalence ratio values within the cavity (*i.e.* less mass exchange with the supersonic core flow). For the range of Sc_t values considered in this effort, the equivalence ratio at a given location in the flameholder can vary by as much as a factor of 2. A similar level of equivalence ratio variability is evident when the Sc_t value is held fixed and only the turbulence model is altered. This spread in the results is significantly larger than the numerical error that was quantified previously which indicates that the overall simulation uncertainty is dominated by the model-form uncertainty. Hence, some procedure is required to calibrate the RAS model prior to any attempt to simulate the combusting environment. Toward this end, the decision was made to use the higher-fidelity hybrid RAS/LES results to inform (or calibrate) the modeling decisions for the Reynolds-averaged simulations.

Before utilizing the hybrid RAS/LES simulations as a "truth" model for RAS model calibration, simulations were performed to assess the sensitivity of the hybrid RAS/LES results to the grid resolution and the chosen Sc_t value. Our previous work[9] showed that the coarse grid provided adequate resolution for LES based on a variety of measures that included a power spectral density analysis as well as comparisons of the fraction of SGS to resolved turbulence present in the time-averaged statistics. Nevertheless, the effect of grid resolution on the simulation statistics was not found to be negligible, so additional assessments have been performed for the present exercise. It should be noted that the variability in the statistics due to grid resolution is basically analogous to changing the form of the SGS turbulence model due to the fact that the filter width changes with grid resolution. The results of the sensitivity analysis are shown
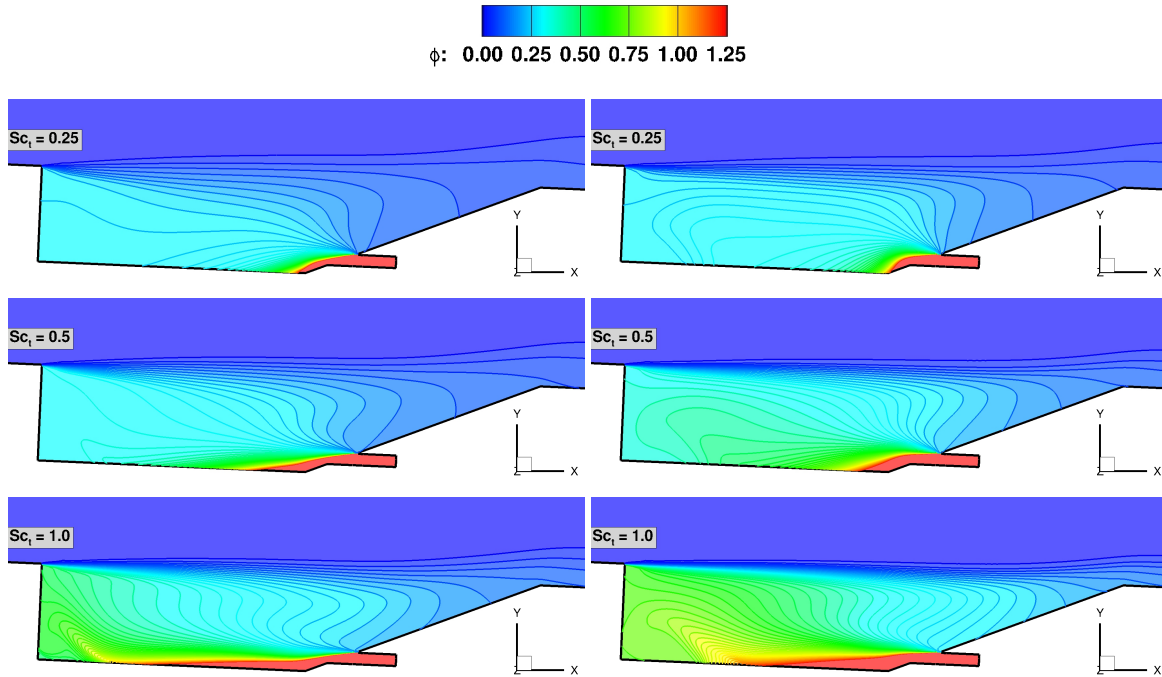
Figure 9: Fuel-Air equivalence ratio contours (injector centerplane) displaying the sensitivity to choice of turbulence model and Sc_t value for Case 3: Menter-SST results (left), Gatski-EAS results (right)

in Fig. 11, where the effect of grid resolution is displayed in the left image and the sensitivity to the Sc_t value is shown on the right. The error bars in these images represent the variability of the time-averaged equivalence ratio for the range of values used to vary the independent variable (grid resolution or Sc_t value). The grid sensitivity shows how the equivalence ratio deviated from the medium grid value when the grid was coarsened. A more conservative measure would have compared the variability from the fine grid to the coarse grid results. However, our previous work showed the variation of the velocity statistics between grid levels to diminish in a monotonic fashion (in the pure LES regions) as the grid was refined. Given the extreme cost of performing the fine grid simulations, only the variations from the medium to the coarse grid are considered here with the expectation that the variability from the medium to fine grid would be considerably smaller. The sensitivity displayed here is considerably smaller than what was observed with a change in turbulence model for the Reynolds-averaged simulations (see Fig. 10). The range of values used to assess the Sc_t sensitivity spanned the same range as that used for the RAS study (0.25 - 1.0). Moreover, the Sc_t sensitivity analysis was performed on the coarse grid, which should display the highest SGS response. Again, the sensitivity of the equivalence ratio to the Sc_t value is considerably smaller than that observed for RAS (see Fig. 10); a result that illustrates the fact that a large fraction of the turbulence is being directly resolved in the simulations. Overall, the variability of the equivalence ratio with grid resolution and Sc_t value was much smaller than what was observed for RAS, providing some degree of confidence towards using the hybrid RAS/LES data to calibrate the lower-fidelity RAS models.

Based on the results of the modeling sensitivity activities above and the grid resolution analysis of our previous work, the medium grid was chosen for the remaining hybrid simulations and the turbulent Lewis number was set to unity (Sc_t=0.9). Instantaneous images of the density gradient magnitude are shown in Fig. 12. This quantity is similar to a Schlieren image and allows both turbulent structures and inviscid structures (shock waves and expansions) to be visualized. A rich array of flow structures are resolved in the simulations, including eddy shocklets emanating from the movement of large coherent structures in the shear layer, a crisp shear layer re-attachment shock system near the back face of the flameholder, and the mixing dynamics of the fuel plume with the unsteady cavity flowfield. The mixing dynamics are detailed further in Fig. 13 which shows a snapshot of the equivalence ratio within the flameholder. The colormap for this image has been crafted to highlight the potential flammability region within the cavity.

An instantaneous snapshot and the time-average of the fuel-air equivalence ratio are shown in Fig. 14 for the highest fuel flow rate condition (Case 3). These images were extracted at the spanwise location that intersects the
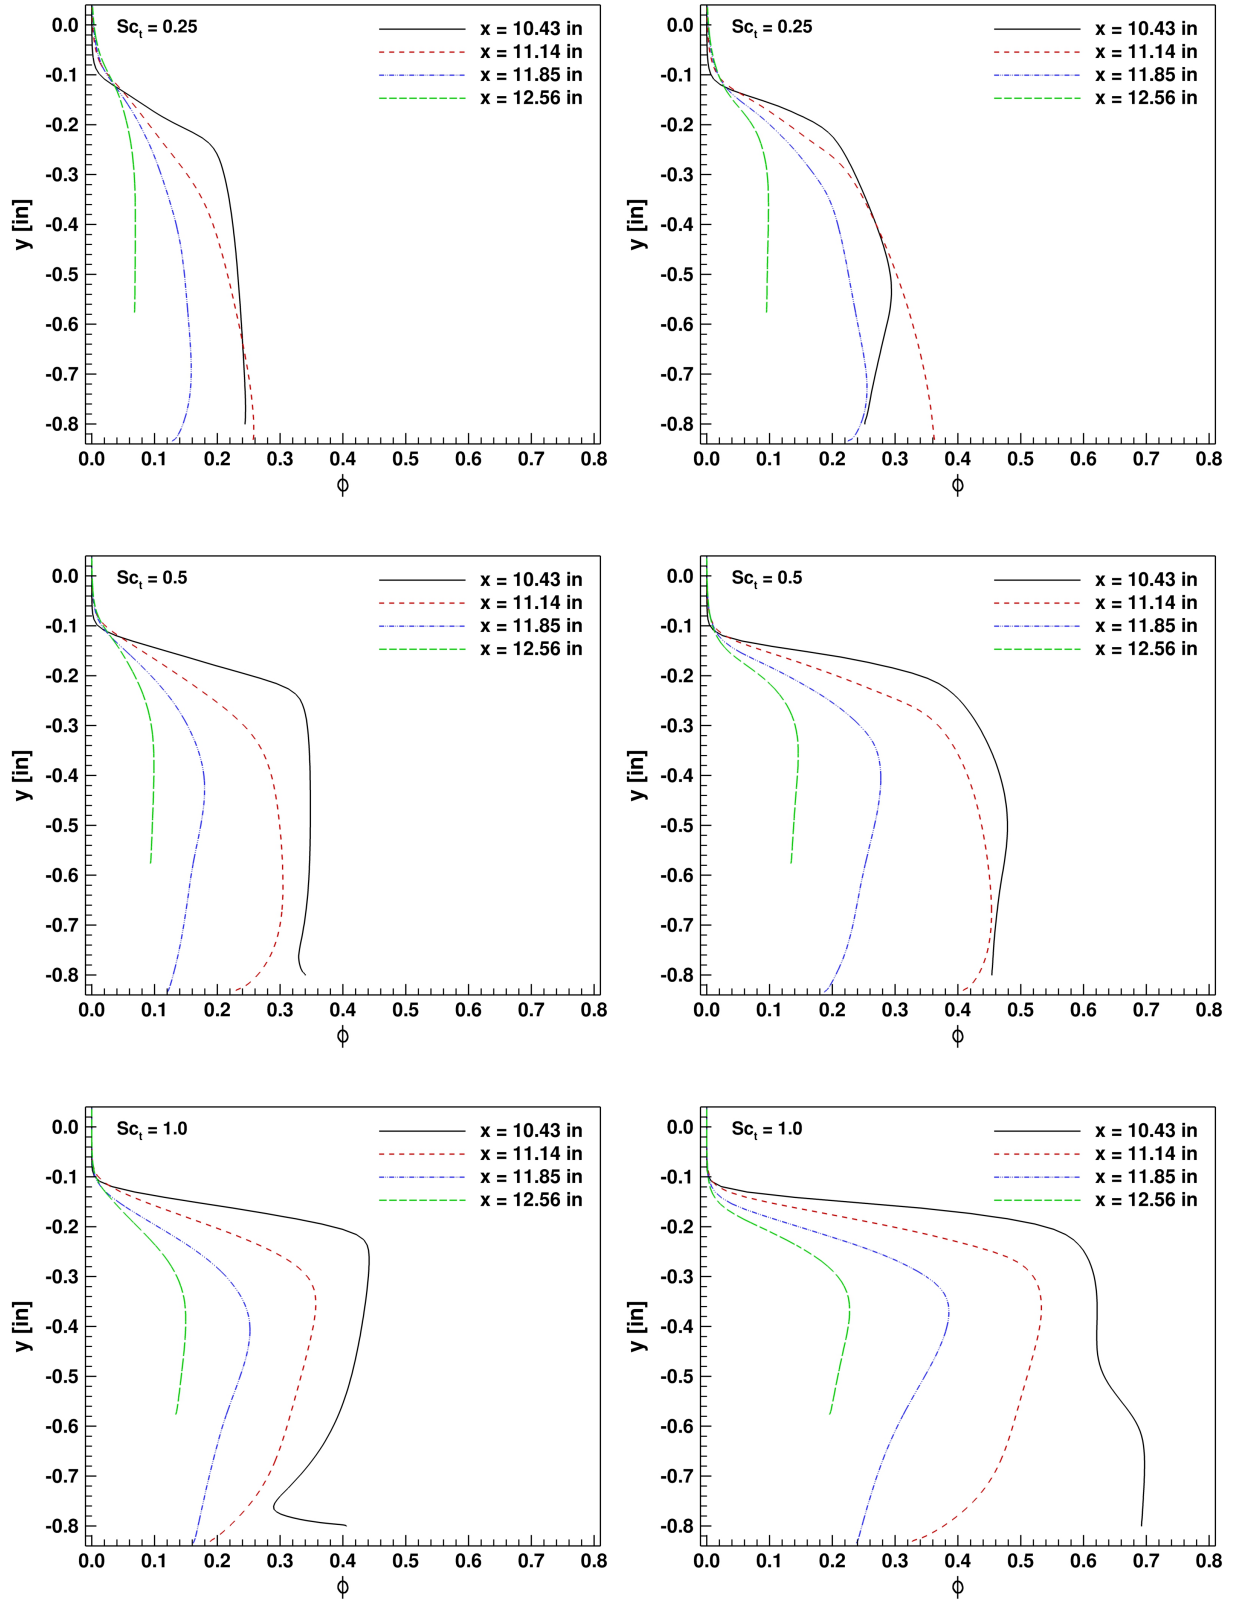
American Institute of Aeronautics and Astronautics

Figure 10: Fuel-Air equivalence ratio profiles (gap centerplane) displaying the sensitivity to choice of turbulence model and Sc_t value for Case 3: Menter-SST results (left), Gatski-EAS results (right)
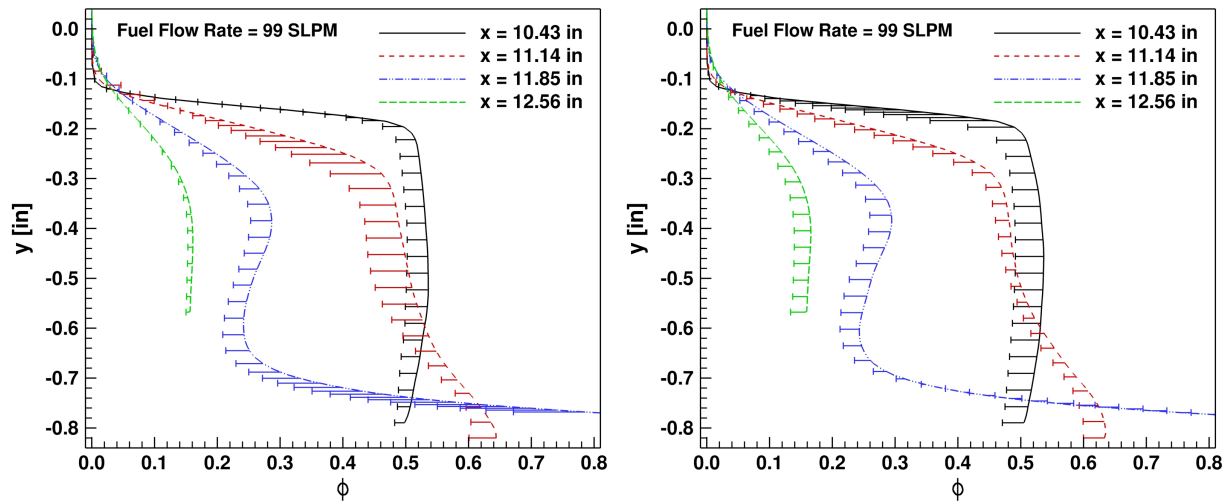
Figure 11: Sensitivity of the fuel-air equivalence ratio profiles (Case 3) within the cavity for hybrid RAS/LES: grid sensitivity (left), Sc_t sensitivity (right)
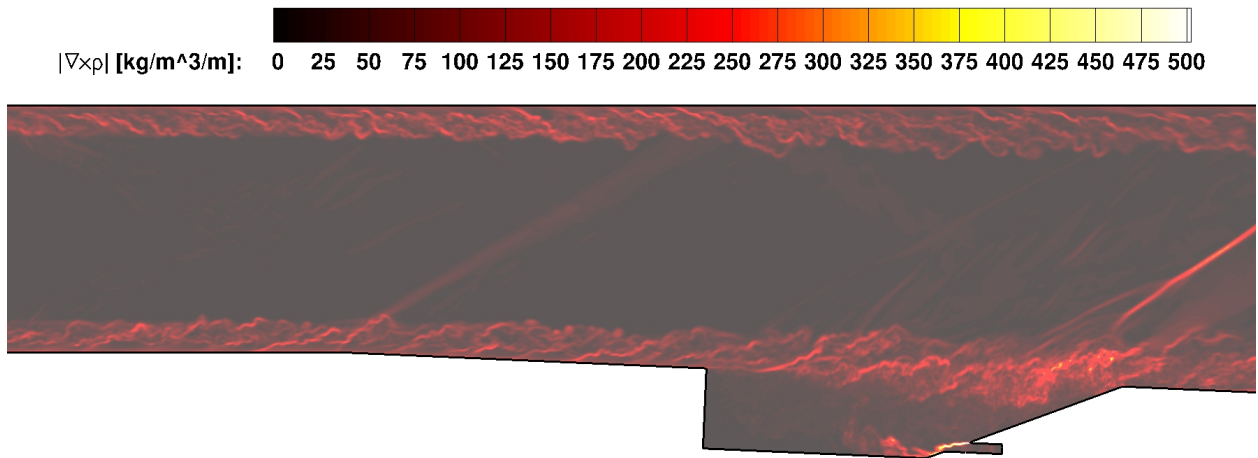

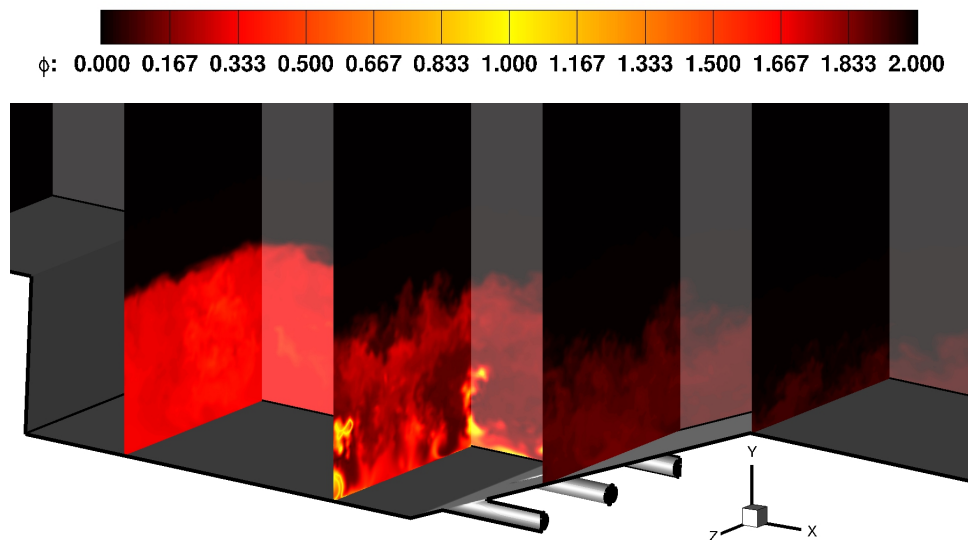Figure 12: Hybrid RAS/LES instantaneous images of density gradient magnitude for Case 2

Figure 13: Hybrid RAS/LES instantaneous images of equivalence ratio for Case 2

center of the cavity fuel injector port to allow direct comparisons with the RAS results shown in Fig. 9. A careful examination of the RAS and time-averaged hybrid RAS/LES results reveals that the cavity air equivalence ratio for this simulation bears many similarities with the Gatski-EAS result (especially at higher Sc_t values), while none of the Menter-SST results appear to replicate this fuel-air distribution. Hence, altering the rate of turbulent mixing by varying the Sc_t value was not sufficient to capture the hybrid RAS/LES behavior. In particular, the size and shape of the secondary recirculation zone alters the mass exchange process within the cavity, and this effect can't be captured by simply scaling the turbulent diffusion processes by a constant scaling factor. The ability of the Gatski-EAS model to more accurately capture this flow feature suggests this model to be the proper choice for this flowfield. As discussed in our previous work[9] the ability of the EAS model to account for Reynolds stress anisotropies is believed to be the mechanism that captures this flow feature. Although not shown, a comparison of the lower fuel flow rate simulations (Case 2) resulted in the same conclusions being drawn.
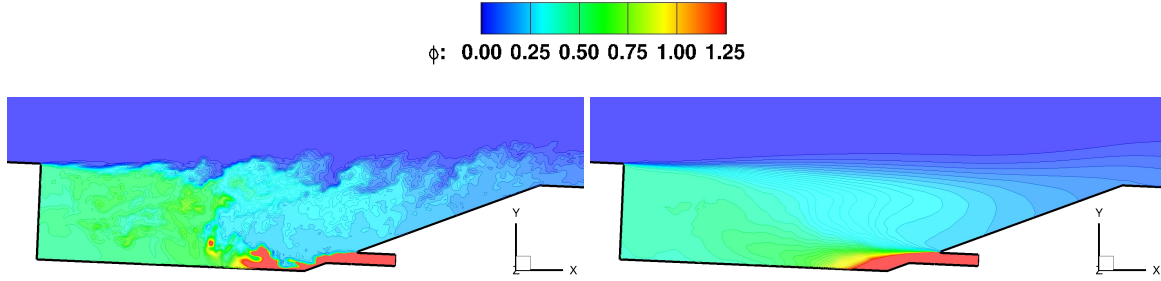

φ: 0.00 0.25 0.50 0.75 1.00 1.25

Figure 14: Fuel-Air equivalence ratio contours (injector centerplane) for Case 3: instantaneous snapshot (left), time-averaged results (right)

Given that the Gatski-EAS model proved to predict a similar cavity flow structure, the next step is to calibrate the Sc_t value. An examination of the time-averaged equivalence ratio from Fig. 14 with the Gatski-EAS results shown in Fig. 9 would suggest that a turbulent Schmidt number somewhere between a value of 0.5 and 1.0 to be an appropriate choice. In order to determine the "best fit" Sc_t value from the hybrid RAS/LES data, a scatter plot was generated to visualize how well the functional form of the RAS diffusion models correlate to the hybrid RAS/LES data. Standard RAS models assume that the Reynolds stress tensor is correlated to the mean strain rate, while the Reynolds mass flux vector is correlated to the mean mass fraction gradient, *i.e.*

$$\tau_{ij} = 2\mu_t \left(S_{ij} - \frac{1}{3}\delta_{ij}\frac{\partial \tilde{u}_k}{\partial x_k} \right) - \frac{2}{3}\delta_{ij}\bar{\rho}\tilde{k}, \qquad \bar{\rho}\widetilde{u_i''Y_m''} = -\frac{\mu_t}{Sc_t}\frac{\partial \tilde{Y}_m}{\partial x_i} \qquad (3)$$

A least-squares procedure can be used to determine the "best fit" μ_t value that correlates the Reynolds stress tensor to the mean strain rate. The same procedure can be applied to determine the "best fit" μ_t / Sc_t value that correlates the Reynolds mass flux vector to the mean mass fraction gradient to arrive at the following relationships:

$$\mu_t = \frac{\left(\tau_{ij} + \frac{2}{3}\delta_{ij}\bar{\rho}\tilde{k}\right)\left(S_{ij} - \frac{1}{3}\delta_{ij}\frac{\partial \tilde{u}_k}{\partial x_k}\right)}{2\left(S_{ij} - \frac{1}{3}\delta_{ij}\frac{\partial \tilde{u}_k}{\partial x_k}\right)\left(S_{ij} - \frac{1}{3}\delta_{ij}\frac{\partial \tilde{u}_k}{\partial x_k}\right)} \equiv \frac{\tau_{ij}^* S_{ij}^*}{2 S_{ij}^* S_{ij}^*}, \qquad \frac{\mu_t}{Sc_t} = \frac{-\bar{\rho}\widetilde{u_i''Y_m''}\frac{\partial \tilde{Y}_m}{\partial x_i}}{\frac{\partial \tilde{Y}_m}{\partial x_i}\frac{\partial \tilde{Y}_m}{\partial x_i}} \qquad (4)$$

The least-squares "best fit" Sc_t value is then obtained on a cell-by-cell basis using the above relationships. This information is displayed in the form of a scatter plot in Fig. 15. In general, there is not single Sc_t value that collapses all of the data. However, there is a clear clustering of points that correspond to a slope (Sc_t value) of 0.7, suggesting this to be the proper value given the constraint of a constant Sc_t value imposed by the standard gradient diffusion assumption.

Based on the information provided by the hybrid RAS/LES results, the Gatski-EAS model with Sc_t=0.7 was selected as the calibrated RAS model for predicting the mixing processes within the cavity flameholder. The equivalence ratio profiles obtained using the calibrated RAS model are compared with the hybrid RAS/LES values in Fig. 16. For this comparison, the equivalence ratio was spatially averaged in the spanwise (*z*) direction at each axial station to provide an overall integral value at each profile station. In general, the calibrated RAS model does a reasonable job of reproducing the time-averaged fuel-air distribution extracted from the hybrid RAS/LES data. The largest discrepancy is at the 11.85 inch station which is the closest station shown to the fuel injection plane. The fuel plume traverses along the lower wall of the cavity at this station and mixes with the surrounding cavity fluid. The mixing rate is high

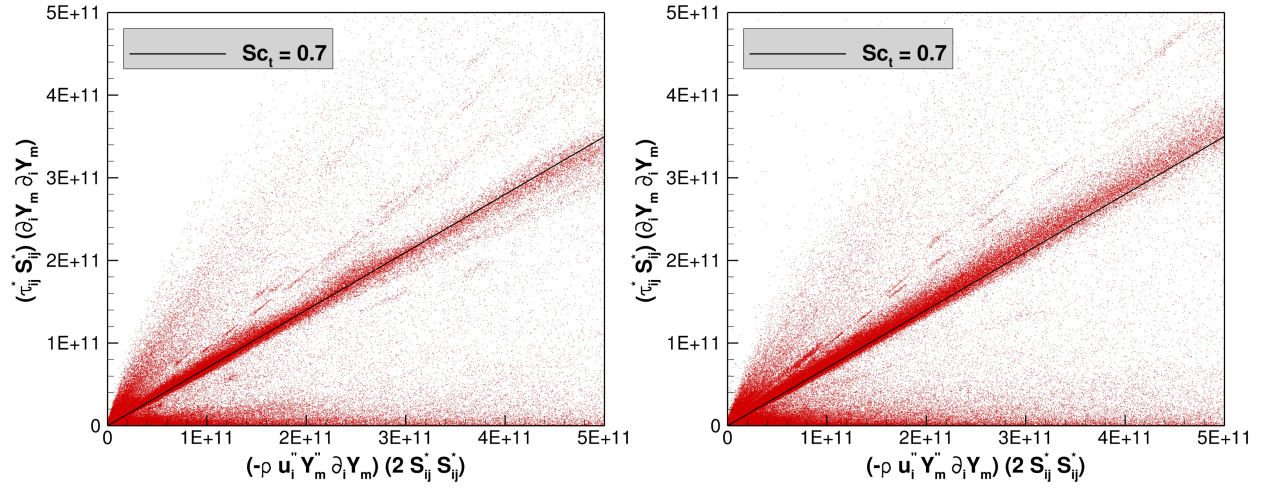American Institute of Aeronautics and Astronautics

Figure 15: Scatter plots of processed hybrid RAS/LES data to examine the correlation of the extracted Reynolds stress tensor and mass flux vector to their RAS modeled form: Case 2 (left), Case 3 (right)
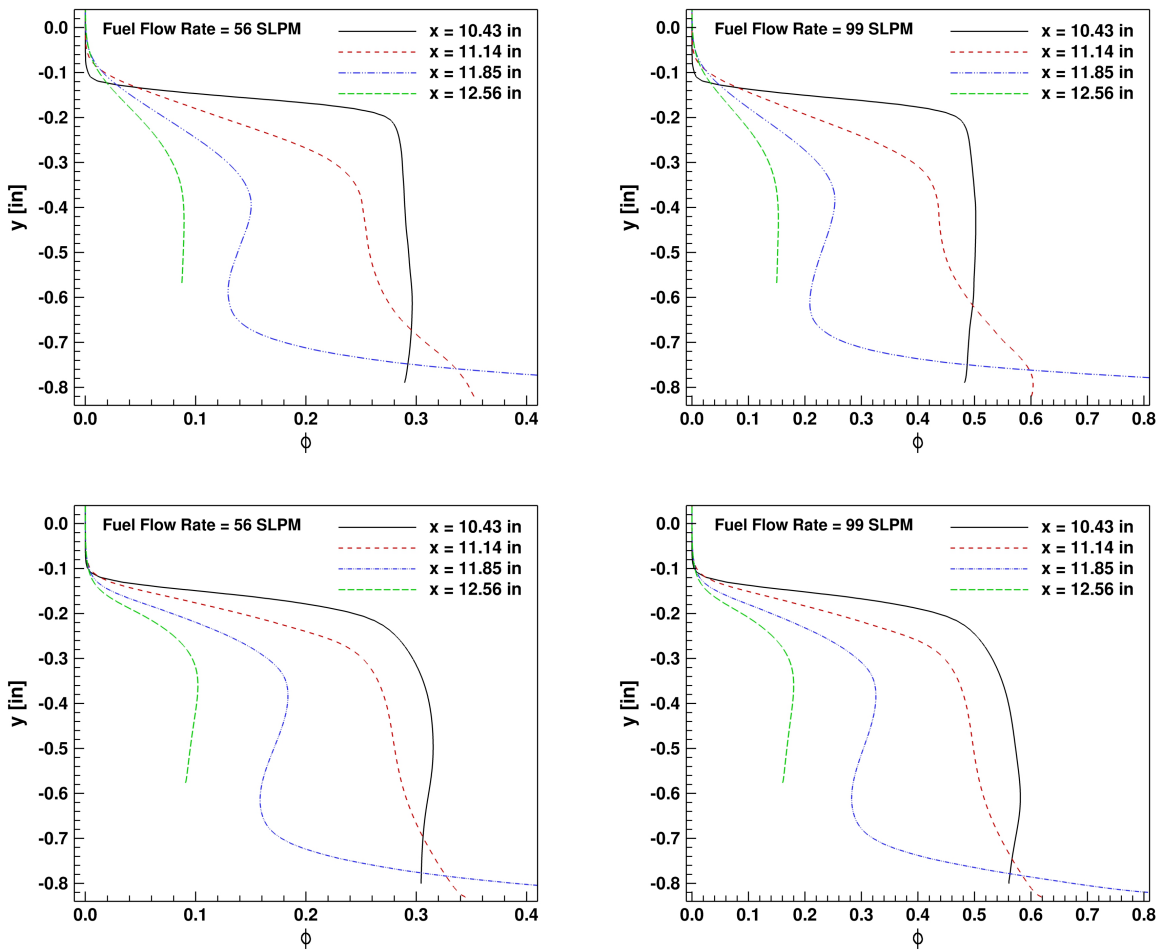

Figure 16: Hybrid RAS/LES (top) and calibrated RAS (bottom) equivalence ratio comparisons

American Institute of Aeronautics and Astronautics

in this region as evident by the large gradient in fuel mass fraction near the lower surface of the cavity. Overall, the calibrated RAS model is slightly overpredicting the mixing rate in this region. However, the overall equivalence ratio levels are well predicted at each profile station with an average percent difference on the order of 10%.

Before leaving this section, it is prudent to mention the relevance of the overall activity (or process) described above. Experimental testing at true hypersonic conditions at ground test facilities is not an easy task. The harsh thermodynamic flowfield environment associated with hypersonic flows severely limits what can be measured and the fidelity of these measurements. This is especially true for hypersonic air-breathing propulsion systems due to the environment, scale, and duration limitations of ground test facilities. As a result, many in the scramjet community see CFD as playing a major role in filling this void. Unfortunately, the current state-of-the-art for applying CFD to these devices relies heavily on the experience of the CFD practitioner to estimate the model-form uncertainty associated with their Reynolds-averaged simulations. This is typically accomplished through simple sensitivity studies and calibration efforts based on limited ground test data (*e.g.* time-averaged surface pressures). For example, the turbulent Schmidt number is often used as a calibration parameter to replicate the mean pressure distribution through the engine. However, the CFD practitioner has no idea if the turbulent mixing is truly the root cause of the discrepancy with measurements, and therefore has no knowledge if this is the proper calibration parameter. The introduction of a higher-fidelity computational capability can provide additional information in this regard (as illustrated above), which can influence the confidence one assigns to the Reynolds-averaged simulations that are performed. In other words, high-fidelity tools can play a role in the engineering analysis of real-world devices even if the computational cost prohibits their use as the primary analysis tool.

Combusting Results

The Reynolds-averaged simulations with combustion were performed using the calibrated settings described in the previous section (Gatski-EAS turbulence model, Pr_t=0.9, and Sc_t=0.7). All other simulation parameters were identical to that used for the non-combusting simulations except for the thermal wall treatment. The total temperature of the facility flow (589 K) was not too far removed from the ambient conditions present in the test facility allowing the use of adiabatic conditions for all flowpath surfaces in the non-reacting simulations. The 4140 steel uncooled facility hardware could not withstand the steady-state thermal environment within the cavity for the combusting tests, so a single thermocouple was located on the floor of the cavity approximately 1-inch from the front wall for health monitoring purposes. This thermocouple was mounted flush with the cavity surface and fully exposed to the flowfield. Once this thermocouple registered a reading of approximately 800 K, the experiment was shut down. Given the limited information available for setting an appropriate wall thermal condition, two options were considered. The first option defined the wall thermal condition to be the minimum of the adiabatic wall temperature and 800 K. This allowed the specification of the single measured surface temperature in the combusting regions of the flow while resorting to the adiabatic condition in regions where combustion is not present. The second option utilized a steady-state 1-D coupled fluid/solid heat transfer boundary condition.[27] This boundary condition requires the specification of the temperature at some thickness into the surface, allowing the temperature at the fluid/surface interface to float based on a 1-D steady-state balance of heat flow from the fluid and through the surface. As in the isothermal scenario, the resulting surface temperature was clipped to be no larger than the adiabatic wall temperature to recover the desired surface treatment in non-combusting regions. The finite-rate chemical-kinetic model chosen was a 22-species ethylene model,[28] which has been shown to replicate kinetic behavior of the parent full mechanism[29] that it was derived from quite accurately.

The thermal state of the cavity would not support auto-ignition of ethylene, so an ignition source was required to ignite the flow. This was accomplished by enforcing a minimum temperature of 1750 K for a region adjacent to the cavity floor that extended approximately 25% of the distance into the cavity. The ignition temperature was enforced as an isobaric process, and the cavity walls were treated as adiabatic surfaces during the ignition process to prevent heat loss through the walls. The ignition source was removed after 5000 iterations, and the adiabatic condition was replaced with one of the surface conditions described above after 10000 iterations. A typical convergence history that resulted from this simulation process is given in Fig. 17, which shows temporal history of the surface integrated heat load and the L_2 norm of the steady-state residual error for one of the cases. At a minimum, the following iterative convergence statements were satisfied for each combusting simulation:

- The L_2 norm of the steady-state residual error was at least 3 orders of magnitude smaller than the initial value.

- The surface heat load time history remained unchanged to 4 digits over the final 2500 iteration cycles.

- The relative mass flow rate error, $|\dot{m}_{out} - \dot{m}_{in}| / \dot{m}_{in}$, was on the order of $\times 10^{-7}$.
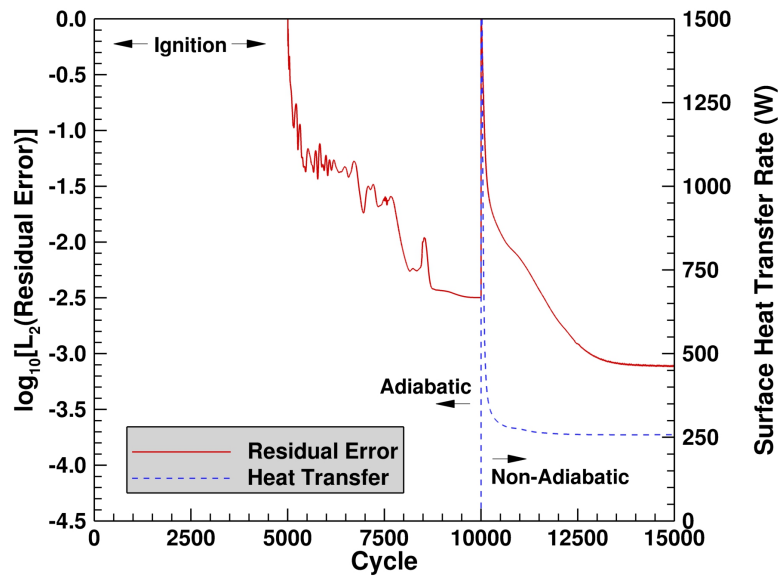
American Institute of Aeronautics and Astronautics

Figure 17: Typical steady-state combusting RAS iterative convergence history

Figure 18 compares the surface thermal environment for Case 3 that resulted from each of the surface boundary conditions described above. The internal temperature for the coupled fluid/solid heat transfer condition was specified as 765 K at a depth of 0.5 inches. This value recovered a surface temperature of approximately 800 K at the location of the thermocouple. The use of the adiabatic surface temperature as a minimum constraint allowed both boundary condition options to capture a lower wall temperature along the surface where unmixed fuel is preventing combustion from occurring (which is at least qualitatively expected). The coupled fluid/solid heat transfer condition also captured a higher surface temperature on the aft-wall of the cavity. The heat transfer is highest along this cavity surface due to the cavity shear layer reattachment and overall higher shear stress present along this surface. Hence, higher surface temperatures are to be expected in this region of the cavity. The use of the coupled fluid/solid heat transfer condition also resulted in an integrated heat loss that was approximately 10% lower than the isothermal condition. The impact of each surface treatment on several parameters of interest are shown in Fig. 19. This figure shows the cavity combustion processes are not overly sensitive to the treatment of the surface boundary condition. Hence, the coupled fluid/solid condition was retained for all remaining simulations since it qualitatively reproduces the expected surface temperature behavior.

The numerical error was assessed at several profile stations for both fueling rates in Fig. 20. For this assessment, the temperature and the equivalence ratio (elemental) was spatially averaged in the spanwise (z) direction at each axial station to provide an integral assessment at each profile station. In general, the numerical error for the combusting simulations was comparable to that observed for the non-combusting simulations, and is not a major contributor to the overall uncertainty of the simulations.

A variety of reacting flow properties are shown in Fig. 21. These contours were extracted from the centerplane between adjacent injector ports. The lower fuel flow rate (Case 2) results are shown on the left, while the higher fuel flow rate (Case 3) is shown on the right. The peak temperature for Case 3 is further aft in the cavity than what is predicted for Case 2. The reason behind this trend is revealed in the second row of images which displays the elemental equivalence ratio contours. The fuel-air ratio is at or near the stoichiometric value towards the aft end of the cavity for Case 3, while Case 2 shows the stoichiometric region to be closer to the center of the cavity. For both cases, the cavity is in a fuel rich state (for the most part), resulting in excess ethylene accumulating near the front wall of the cavity. This is in contrast to the non-combusting scenario which showed the cavity to be entirely in a fuel lean state. This difference is an artifact of the volumetric expansion due to heat release in the cavity as will be shown shortly. The distribution of combustion products (steam, carbon dioxide, and carbon monoxide) is given in the bottom three rows of images. The peak steam mass fraction values correlate well with the location of the stoichiometric regions in the cavity. As one might expect, the carbon dioxide mass fraction is a maximum in the fuel lean regions of the cavity, while the carbon monoxide mass fraction peaks in regions with larger fuel-air ratios.

An examination of the flowfield differences between the combusting and non-reacting cavity simulations revealed that the primary mechanism for the increased equivalence levels within the cavity under reacting conditions was due

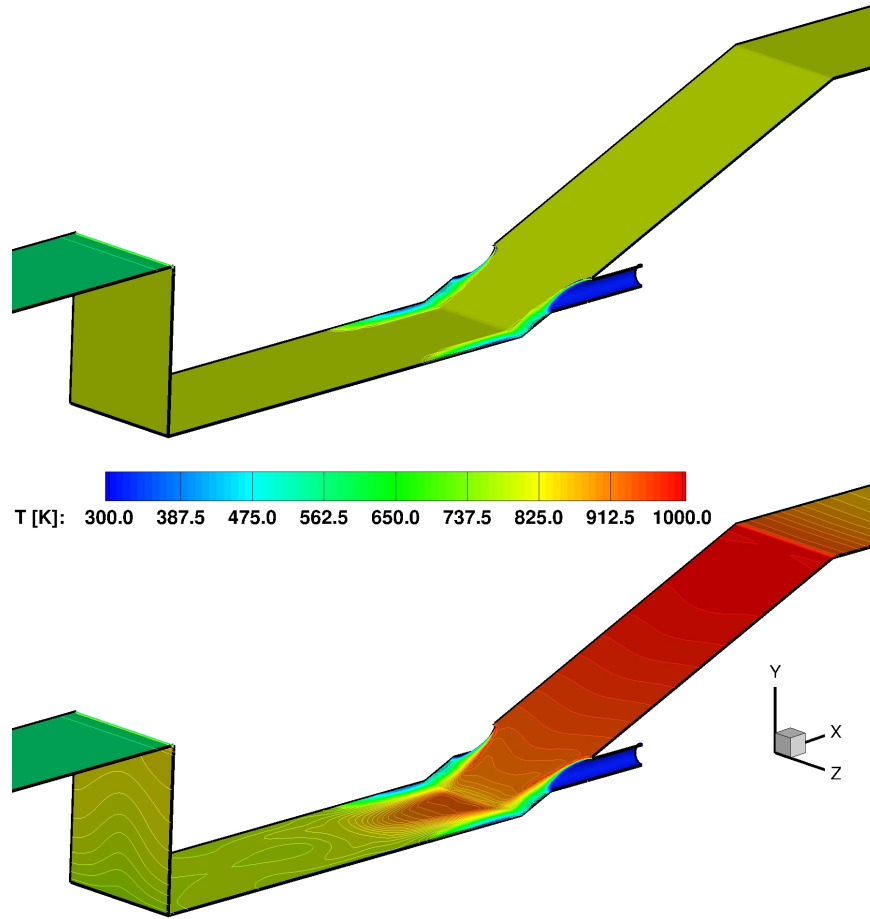American Institute of Aeronautics and Astronautics

Figure 18: Surface temperature comparison (Case 3): isothermal wall (top), 1-D coupled fluid/surface wall (bottom)
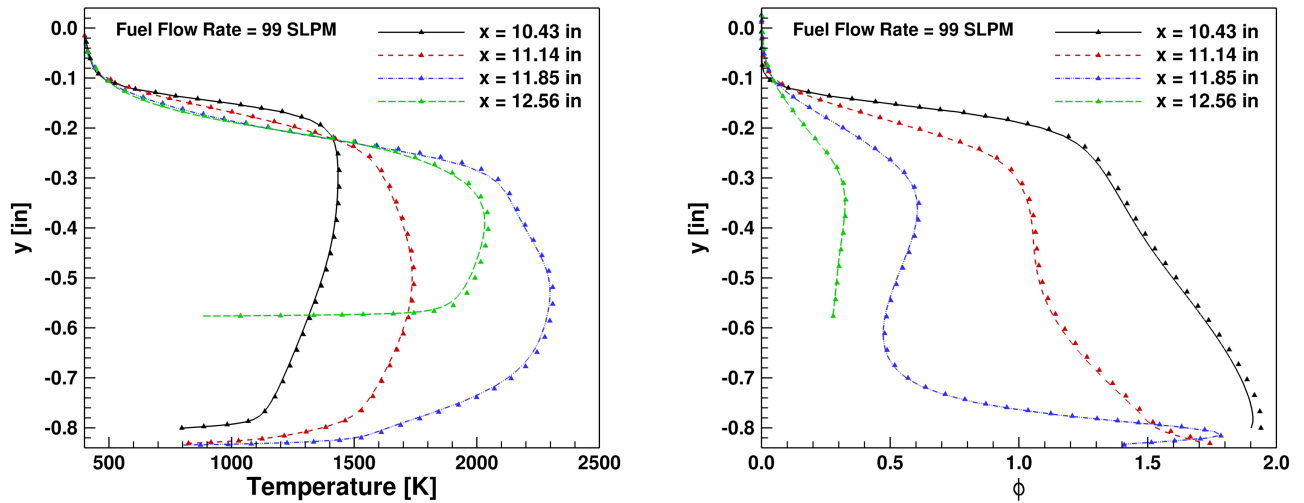

Figure 19: Temperature (left) and elemental fuel-air equivalence ratio (right) profiles displaying the sensitivity to choice of thermal surface treatment: isothermal (lines), 1-D coupled fluid/surface (symbols)
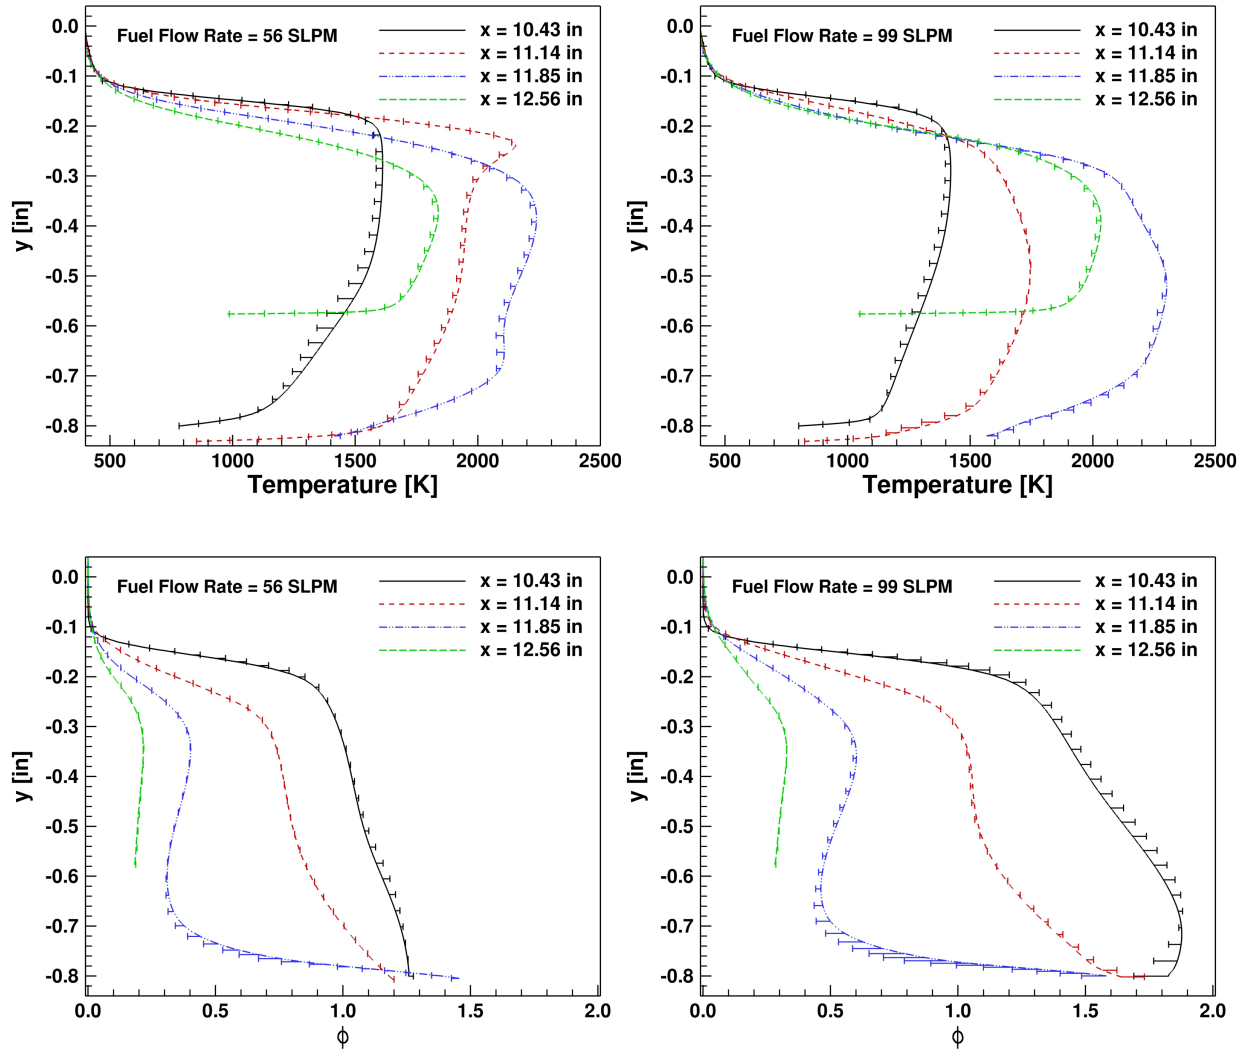
American Institute of Aeronautics and Astronautics

Figure 20: Temperature and elemental fuel-air equivalence ratio profiles within the cavity with error bars depicting the GCI bounding estimate for the numerical error
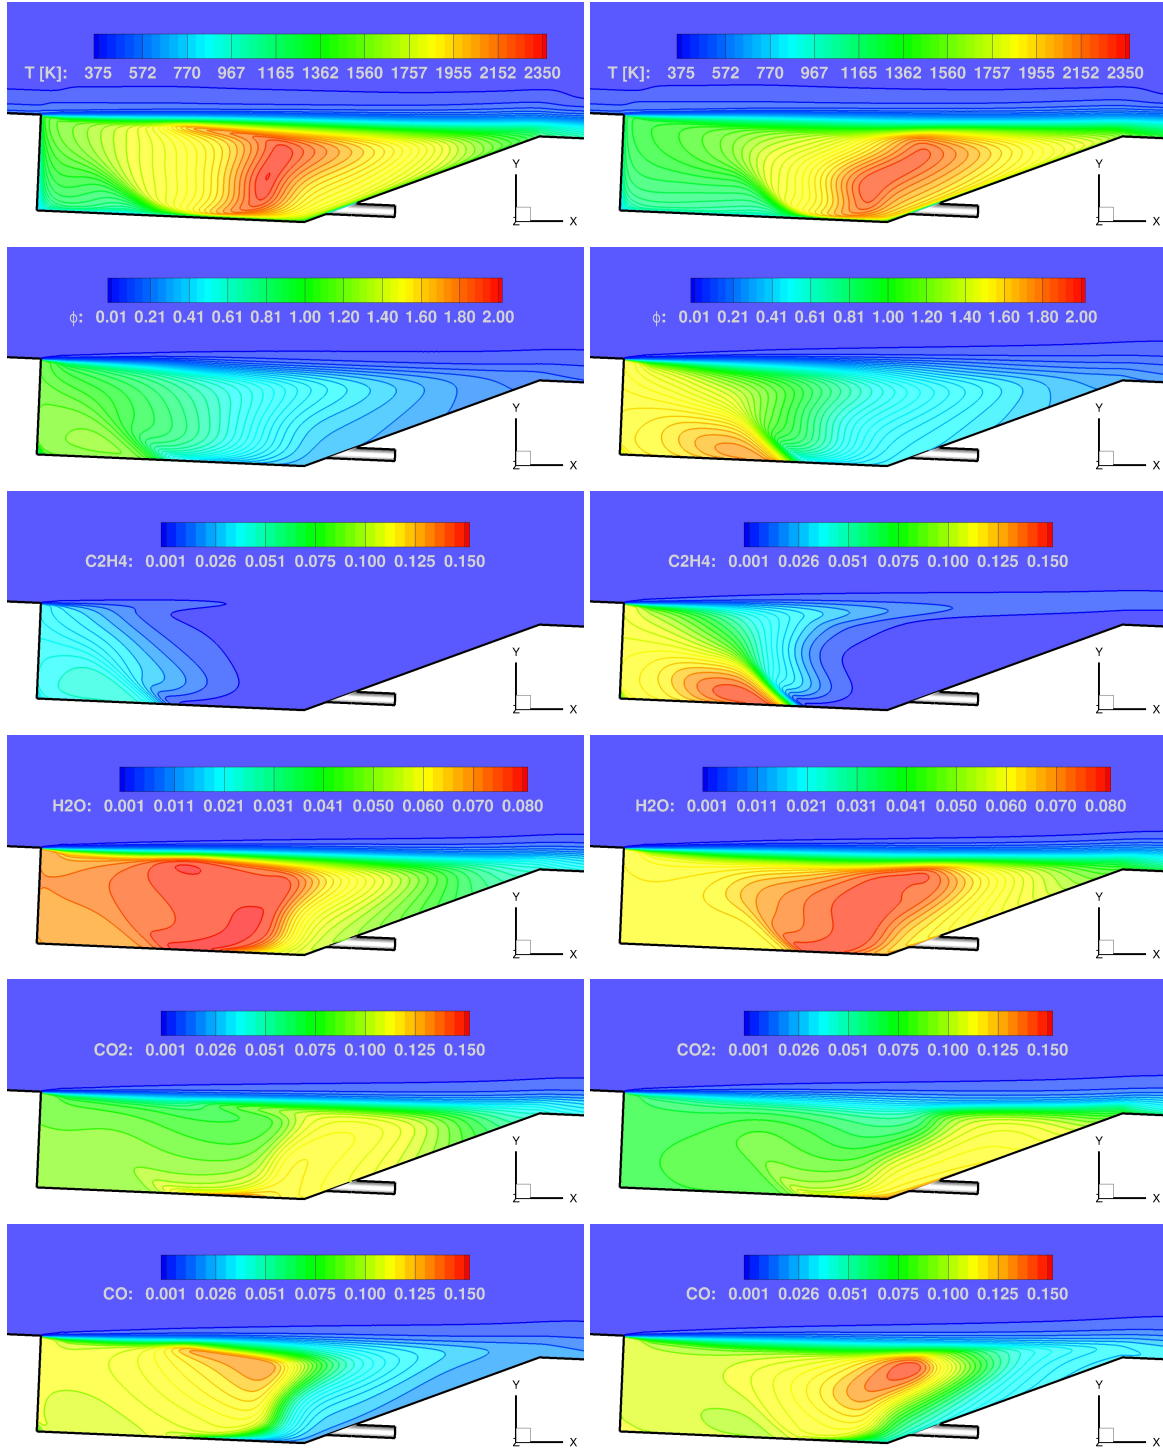
American Institute of Aeronautics and Astronautics

Figure 21: Cavity flameholder combusting flow properties (gap centerplane): Case 2 (left), Case 3 (right)

to volumetric expansion. The heat release significantly reduces the density in the cavity, resulting in nearly a 4-fold decrease in the eddy viscosity levels within the cavity recirculation zones. Fig. 22 compares the effective $(\mu + \mu_t)$ values under non-reacting (left image) and combusting (right image) conditions for Case 3. The effective viscosity in the primary recirculation zone is nearly 3 - 4 times smaller under combusting conditions. Given that the fuel is injected within the cavity at a rate that hardly perturbs the aerodynamically induced recirculation zones, the primary mechanism for fuel to escape the cavity is through turbulent transport, which is modeled in RAS as a diffusion process. Therefore, the reduced diffusion hinders the mass exchange with the core flow leaving more of the fuel inside of the cavity. This same trend was seen in the Sc_t sensitivity study where higher Sc_t values (*i.e.* lower diffusion rate) led to increased levels of fuel within the cavity (see Fig. 10). Note also that an increase in temperature tends to increase the molecular viscosity. Hence, the reduction in the effective viscosity is due to a reduction in the eddy viscosity in the flameholder under reacting conditions. Also shown in Fig. 22 is the effect that the heat release has on the velocity field within the flameholder. The center of the primary recirculation zone has shifted further upstream and the structure of the secondary recirculation zone is much more elongated and less pronounced than what was predicted in the non-reacting environment.
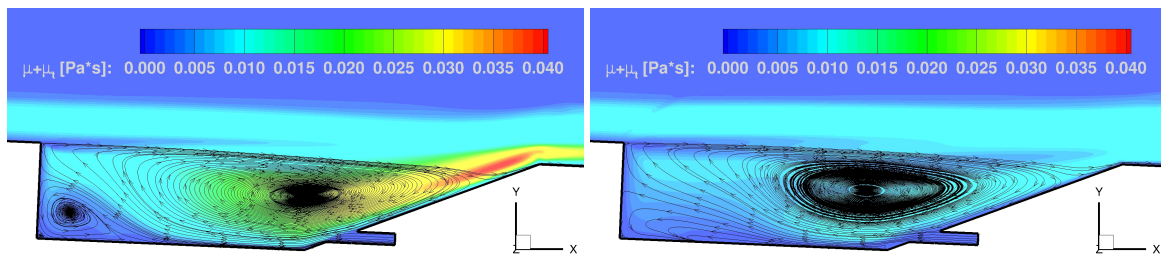

Figure 22: Comparison of non-combusting (left) and combusting (right) flow structure and mixture viscosity at the injector centerplane (Case 3)

A comparison of the predicted cavity velocity field with measurements is shown in Fig. 23 (Case 2) and Fig. 24 (Case 3). In general, the level of agreement for the streamwise velocity component for Case 2 is comparable to what was obtained in our previous simulations[9] without fuel injection (Case 1). The shear layer spreading rate is slightly underpredicted (as was the case for Case 1), and the reverse flow of the primary recirculation zone is slightly over-predicted. This latter feature was more accurately predicted in our previous Case 1 simulations. It should be noted that the window size used to post-process the raw PIV data was 0.07×0.07 inches (10 times coarser in each direction than the resolution of the fine grid simulations). This resolution is insufficient to fully resolve the steep gradients present at the early profile stations. Hence, part of the discrepancy at the early stages of the shear layer development is due to the implicit filter introduced by the finite PIV window size. Also, the core flow above the cavity shear layer was consistently overpredicted by all of the simulations, which may be a result of the flow being insufficiently seeded in this region. The flowfield was seeded in the boundary layer upstream of the cavity through an angled slot injector. The flow rate of the seeded air was intentionally kept small (10 SLPM) to minimize disturbances to the flow, so if the outer portion of the shear layer was insufficiently seeded, the net result would be a measurement that is biased towards lower velocity values. The transverse velocity is compared with the PIV measurements in the lower image of Fig. 23. The transverse velocity levels are very small throughout the cavity, so the differences present between the computations and measurements appear more magnified for this velocity component. However, for the most part the level of agreement visualized for this combusting simulation is arguably no worse than that obtained for Case 1.[9] The same discrepancies noted in the streamwise velocity component for Case 2 are also present for the higher fueling rate condition (Case 3). However, the overprediction of the reverse flow of the primary recirculation zone is more pronounced at this condition. The transverse velocity levels are also systematically larger (in magnitude) than the measurements indicate for this condition.

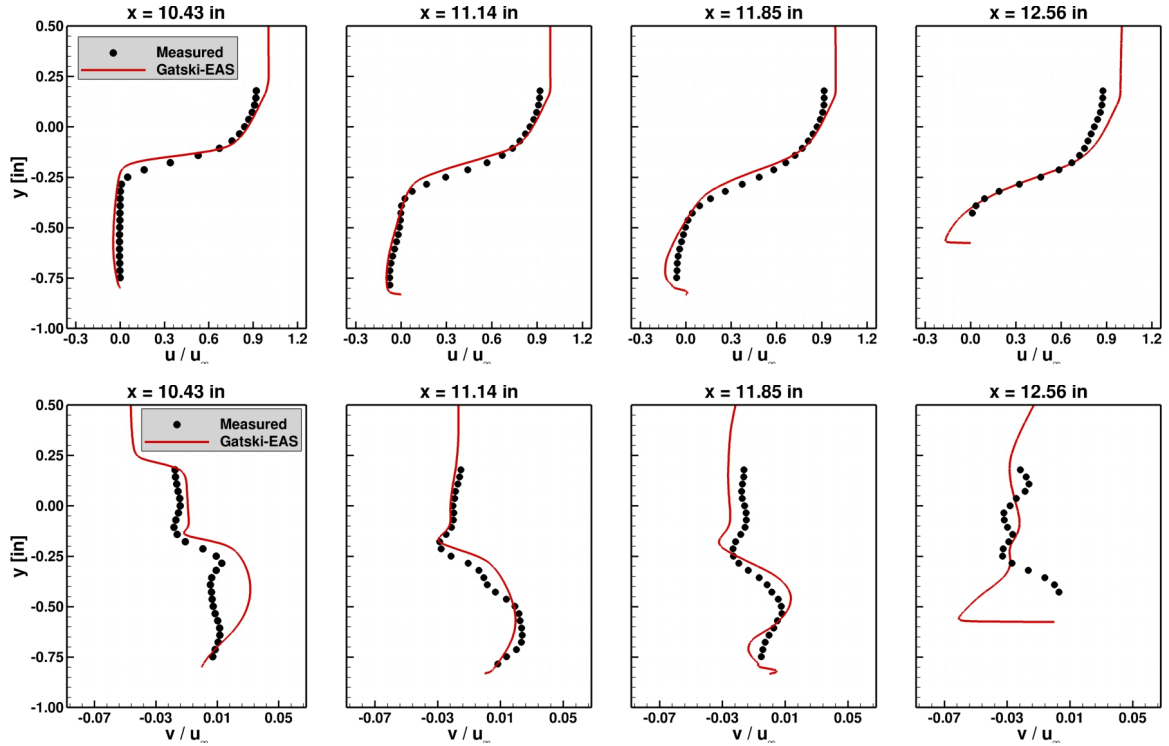American Institute of Aeronautics and Astronautics

Figure 23: Comparison of x-velocity and y-velocity components with the PIV measurements for Case 2

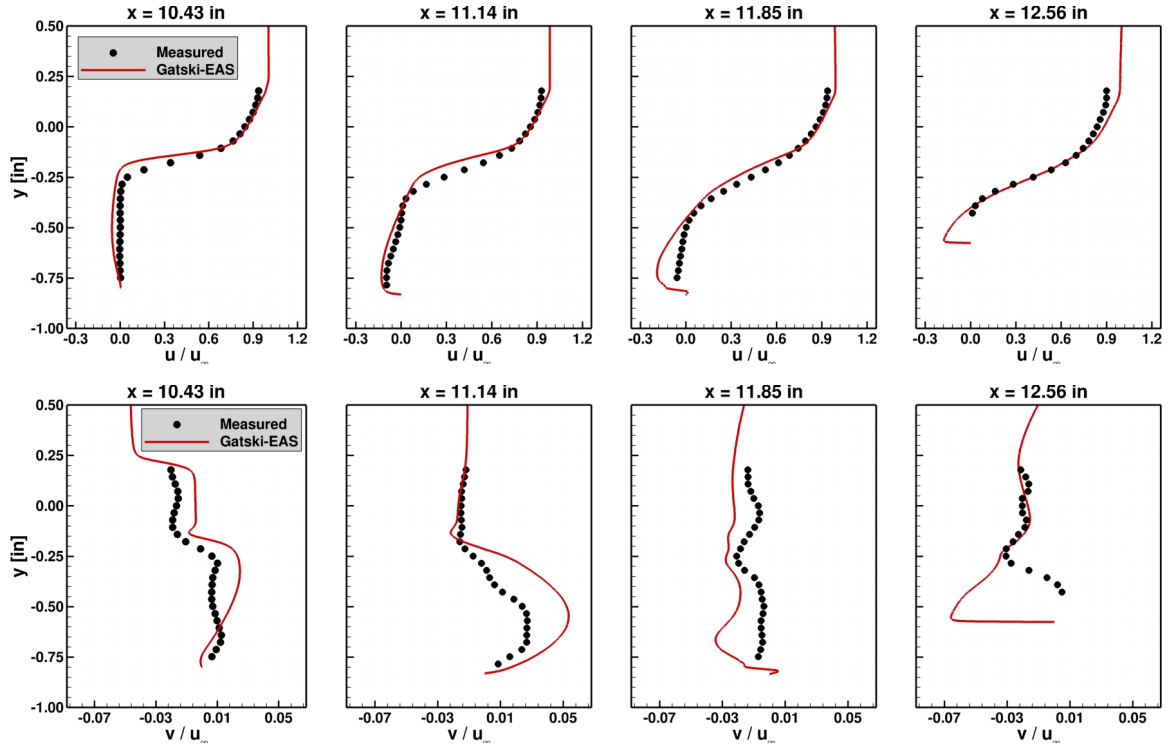Figure 24: Comparison of x-velocity and y-velocity components with the PIV measurements for Case 3

American Institute of Aeronautics and Astronautics

Summary and Future Work

Reynolds-averaged and hybrid Reynolds-averaged/large eddy simulations have been performed for a model scramjet cavity flow experiment. Particle Image Velocimetry data was taken at the centerplane of the cavity, allowing for detailed comparisons of the simulated cavity flow structure with measurements at both no-fuel and fueled reacting conditions. The purpose of the computational effort was to assess the state-of-the-art for both RAS and hybrid RAS/LES predictions for this flowfield, which is a relevant one to the scramjet research community. Moreover, this effort provided a real-world example of how high-fidelity scale-resolving simulations can be utilized to address the model-form uncertainty of Reynolds-averaged models.

The Reynolds-averaged non-reacting simulations of this flow were performed using a linear eddy viscosity model (Menter-SST k-ω model) and a non-linear model (k-ω based explicit algebraic stress model of Gatski and Rumsey). The grid convergence index was used to show that the numerical errors were negligible relative to the model-form variability associated with the choice of turbulence closure for the Reynolds stress tensor and the Reynolds mass flux vector. Experimental data was not available to assess this model form variability without bringing in the complexities of chemical reactions. Therefore, the decision was made to use the scale-resolving simulations as an intermediary "truth" model to address the RAS model-form uncertainties. The results of this assessment led to a calibrated RAS model (Gatski-EAS model with Sc_t=0.7) that was carried forward for the combusting simulations for comparison with measurements.

The heat release due to combustion had a pronounced effect on the fuel-air mixing within the cavity flameholder. The flameholder was operating under fuel lean conditions in a non-combusting environment, which switched to a fuel rich operation after ignition. The mechanism behind this transition was the volumetric expansion induced by the combustion process, which reduced the eddy viscosity by a factor of 4 (primarily through a reduction in the density). The lower eddy viscosity values are indicative of reduced turbulent transport processes relative to the non-reacting environment, which is the primary mass exchange mechanism between the supersonic core flow and the cavity flameholder. Hence, the cavity retained a larger fraction of the fuel that was directly injected into the cavity under combusting conditions. In general, the Reynolds-averaged velocity profile predictions at the lowest fueling level compared reasonably well with the particle imaging measurements (almost as well as the non-reacting condition without fuel injection). However, the velocity field predictions proved to be more sensitive to the flameholder fueling rate than was indicated in the measurements, resulting in less favorable comparisons with measurements when the fuel flow rate was increased.

As a final note, the computational cost of using a 22-species kinetic model with a scale-resolving simulation strategy proved to be too costly given the computational resources that were available for this effort. The hybrid RAS/LES simulations without finite-rate chemistry on the medium grid required 20 days of wall clock time on 768 cores. The cost of adding 20 additional species and activating the chemical kinetics would have increased this cost by at least a factor of 5 (for the same number of cores). This was unfortunate since flame-holding characteristics are known to be difficult to predict with RAS approaches, and one would expect a measurable improvement in the ability to predict these properties with a wall-modeled LES strategy. This clearly illustrates the need to find alternative methods to accurately parametrize the complex chemistry with only a limited number of scalars. Our group is currently examining a flamelet-based strategy[30,31] to address this issue.

Acknowledgments

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," NASA Contractor Report 218178, 2014.

[2] Baurle, R. A., Tam, C.-J., Edwards, J. R., and Hassan, H. A., "Hybrid Simulation Approach for Cavity Flows: Blending, Algorithm, and Boundary Treatment Issues," *AIAA Journal*, Vol. 41, No. 8, 2003, pp. 1463–1480.

[3] Menter, F. R. and Kuntz, M., "Adaptation of Eddy-Viscosity Turbulence Models to Unsteady Separated Flow Behind Vehicles," *The Aerodynamics of Heavy Vehicles: Trucks, Buses, and Trains*, edited by B. F. McCallen, R. and J. Ross, Springer Berlin Heidelberg, 2004, pp. 339–352.

[4] Girimaji, S. S., "Partially-Averaged Navier-Stokes Model for Turbulence: A Reynolds-Averaged Navier-Stokes to Direct Numerical Simulation Bridging Method," *Journal of Applied Mechanics*, Vol. 73, No. 3, 2006, pp. 413–421.

[5] Shur, M. L., Spalart, P. R., Strelets, M. K., and Travin, A. K., "A Hybrid RANS-LES Approach with Delayed-DES and Wall-Modelled LES Capabilities," *International Journal of Heat and Fluid Flow*, Vol. 29, No. 6, 2008, pp. 1638–1649.

[6] White, J. A. and Morisson, J. H., "Pseudo-Temporal Multi-Grid Relaxation Scheme for Solving the Parabolized Navier-Stokes Equations," AIAA Paper 99-3360, 1999.

[7] VULCAN, "http://vulcan-cfd.larc.nasa.gov/," 2015.

[8] White, J. A, Baurle, R. A., Fisher, T. C., Quinlan, J. R., and Black, W. S., "Low Dissipation Advection Schemes Designed for Large Eddy Simulations of Hypersonic Propulsion Systems," AIAA Paper 2012-4263, 2012.

[9] Baurle, R. A., "Hybrid Reynolds-Averaged / Large Eddy Simulations of a Cavity Flameholder; Assessment of Modeling Sensitivities," AIAA Paper 2015-0637, 2015.

[10] Tuttle, S., Carter, C., and Hsu, K., "Particle Image Velocimetry in an Isothermal and Exothermic High-Speed Cavity," AIAA Paper 2012-0330, 2012.

[11] Peterson, D. M., Hagenmaier, M. A., Carter, C. D., and Tuttle, S. G., "Hybrid Reynolds-Averaged and Large-Eddy Simulations of a Supersonic Cavity Flameholder," AIAA Paper 2013-2483, 2013.

[12] Peterson, D. M., Hassan, E. A., Tuttle, S. G., Hagenmaier, M. A., and Carter, C. D., "Numerical Investigation of a Supersonic Cavity Flameholder," AIAA Paper 2014-1158, 2014.

[13] Litton, D. K., Edwards, J. R., and White, J. A., "Algorithm Enhancements to the VULCAN Navier-Stokes Solver," AIAA Paper 2003-3979, 2003.

[14] Edwards, J. R., "A Low Diffusion Flux-Splitting Scheme for Navier-Stokes Calculations," *Computers & Fluids*, Vol. 26, No. 6, 1997, pp. 635–659.

[15] van Leer, B., "Towards the Ultimate Conservation Difference Scheme. II. Monotinicity and Conservation Combined in a Second Order Scheme," *Journal of Computational Physics*, Vol. 14, 1974, pp. 361–370.

[16] McBride, B. J. and Gordon, S., "Computer Program for Calculation of Complex Chemical Equilibrium Composition and Applications, I. Analysis," NASA Reference Publication 1311, 1994.

[17] McBride, B. J. and Gordon, S., "Computer Program for Calculation of Complex Chemical Equilibrium Composition and Applications, II. Users Manual and Program Description," NASA Reference Publication 1311, 1996.

[18] Menter, F. R., "Zonal Two Equation k-ω Models for Aerodynamic Flows," AIAA Paper 93-2906, 1993.

[19] Rumsey, C. L. and Gatski, T. B., "Summary of EASM Turbulence Models in CFL3D with Validation Test Cases," NASA Technical Report TM-2003-212431, 2003.

[20] Choi, J.-L., Edwards, J. R., and Baurle, R. A., "Compressible Boundary Layer Predictions at High Reynolds Number Using Hybrid LES/RANS Methods," *AIAA Journal*, Vol. 47, No. 9, 2009, pp. 2179–2193.

[21] Boles, J. A., Edwards, J. R., Choi, J.-L., and Baurle, R. A., "Simulations of High-Speed Internal Flows Using LES/RANS Models," AIAA Paper 2009-1324, 2009.

[22] Yoshizawa, A. and Horiuti, K., "A Statistically-Derived Subgrid Scale Kinetic Energy Model for Large-Eddy Simulation of Turbulent Flows," *Journal of the Physical Society of Japan*, Vol. 54, 1985, pp. 2834–2839.

[23] Suresh, A. and Huynh, H. T., "Numerical Experiments on a New Class of Nonoscillatory Schemes," AIAA Paper 92-0421, 1992.

[24] Krist, S. L., Biedron, R. T., and Rumsey, C. L., "CFL3D User's Manual (Version 5.0)," NASA Technical Report TM-1998-208444, 1998.

[25] Urbin, G., Knight, D., and Zheltovodov, A. A., "Large Eddy Simulation of a Supersonic Boundary Layer Using Unstructured Grids," *AIAA Journal*, Vol. 39, No. 7, 2001, pp. 1288–1295.

[26] Roache, P. J., *Verification and Validation in Computational Science and Engineering*, Hermosa Publishers, 1998.

[27] White, J. A, "A Modified Wall Matching Treatment to Account for Local Solid to Fluid Thermal Coupling," 2002 JANNAF CS/APS/PSHS/MSS Joint Meeting, 2002.

[28] Luo, Z., Yoo, C. S., Richardson, E. S., Chen, J. H., Law, C. K., and Lu, T. F., "Chemical Explosive Mode Analysis for a Turbulent Lifted Ethylene Jet Flame in Highly-Heated Coflow," *Combustion and Flame*, Vol. 159, No. 1, 2012, pp. 265–274.

[29] Wang, H., You, X., Joshi, A. V., Davis, S. G., Laskin, A., Egolfopoulos, F., and Law, C. K., "USC Mech Version II. High-Temperature Combustion Reaction Model of H2/CO/C1-C4 Compounds," http://ignis.usc.edu/USC_Mech_II.htm/, 2007.

[30] Quinlan, J., McDaniel, J. C., Drozda, T. G., Lacaze, G. and Oefelein, J. C., "A Priori Analysis of Flamelet-based Modeling for a Dual-Mode Scramjet Combustor," AIAA Paper 2014-3743, 2014.

[31] Quinlan, J., McDaniel, J. C., Drozda, T. G., Lacaze, G. and Oefelein, J. C., "A Priori Analysis of a Compressible Flamelet Model Using RANS Data for a Dual-Mode Scramjet Combustor," AIAA Paper 2015-3208, 2015.