Uncertainty Quantification of CFD Data Generated for a Model Scramjet Isolator Flowfield

R. A. Baurle and E. L. Axdahl
Hypersonic Airbreathing Propulsion Branch, NASA Langley Research Center
Hampton, Va 23681

ABSTRACT

Computational fluid dynamics is now considered to be an indispensable tool for the design and development of scramjet engine components. Unfortunately, the quantification of uncertainties is rarely addressed with anything other than sensitivity studies, so the degree of confidence associated with the numerical results remains exclusively with the subject matter expert that generated them. This practice must be replaced with a formal uncertainty quantification process for computational fluid dynamics to play an expanded role in the system design, development, and flight certification process. Given the limitations of current hypersonic ground test facilities, this expanded role is believed to be a requirement by some in the hypersonics community if scramjet engines are to be given serious consideration as a viable propulsion system. The present effort describes a simple, relatively low cost, nonintrusive approach to uncertainty quantification that includes the basic ingredients required to handle both aleatoric (random) and epistemic (lack of knowledge) sources of uncertainty. The nonintrusive nature of the approach allows the computational fluid dynamicist to perform the uncertainty quantification with the flow solver treated as a "black box". Moreover, a large fraction of the process can be automated, allowing the uncertainty assessment to be readily adapted into the engineering design and development workflow. In the present work, the approach is applied to a model scramjet isolator problem where the desire is to validate turbulence closure models in the presence of uncertainty. In this context, the relevant uncertainty sources are determined and accounted for to allow the analyst to delineate turbulence model-form errors from other sources of uncertainty associated with the simulation of the facility flow.

INTRODUCTION

A dual-mode scramjet engine (see Fig. 1) is composed of four main components: inlet, isolator, combustor, and exit nozzle. The inlet captures the air mass for the engine and conditions (compresses) the flow for subsequent combustion. The role of the isolator is to separate the combustion effects ($i.e.$, combustion-induced pressure rise) from the inlet, providing sufficient margin to prevent engine unstart. The combustor houses the fuel injection devices through which fuel is introduced and subsequently mixed with the incoming air allowing for combustion. Finally, the nozzle expands the resulting exhaust gases to produce thrust for the vehicle. At low hypersonic Mach numbers, the heat release associated with the combustion process can lead to (or at least approach) thermal choking of the engine flow, resulting in elevated pressure levels in the combustor. The adverse pressure gradient that forms under these conditions leads to the formation of a pre-combustion shock train within the isolator. This flow is characterized by regions of massive flow separation due to the adverse pressure gradient and the multiple shock/boundary layer interactions that result from

†Approved for public release; distribution is unlimited.

this shock system. The isolator length required to contain the pressure rise induced by the combustion processes is of critical importance to the engine designer, since this drives the size (and weight) of the isolator section.

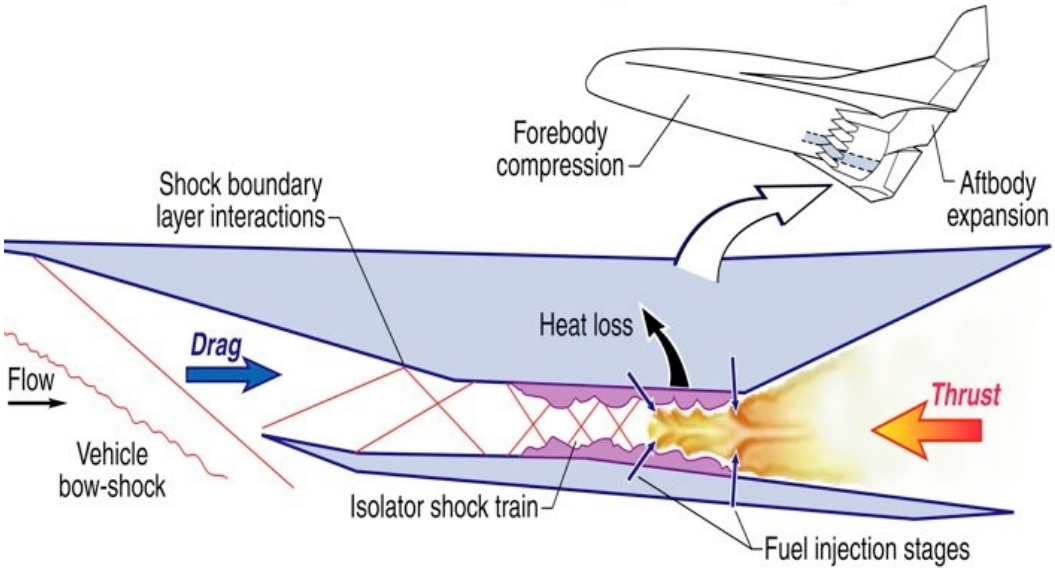

Figure 1: Dual-mode scramjet engine flowfield characteristics

Reynolds-Averaged Simulations (RAS) have been used almost exclusively as the high-fidelity tool of choice for the engineering design and development cycle of scramjet components due to the enormous computational resources required to perform scale-resolved calculations such as Large Eddy Simulation (LES). Unfortunately, the epistemic uncertainties (*i.e.*, systematic uncertainties) associated with imperfect physical models used within a RAS framework are often found to dominate the overall uncertainty in simulations of scramjet flowpaths.[1, 2] If Computational Fluid Dynamics (CFD) is to be used as a predictive tool in the design and analysis of scramjet components, these model-form uncertainties must be well quantified. The quantification of these uncertainties will require that the models in question be validated for their intended purposes (*i.e.*, for realistic geometries in representative flow environments). To date, while the use of CFD is prevalent, very few efforts have been undertaken that truly attempt to validate models for component level simulations. Instead, the current state-of-the-art relies heavily on the experience of the CFD practitioner to estimate the model-form uncertainty associated with their simulations through simple sensitivity studies. This practice will have to be replaced with a formal uncertainty quantification (UQ) process if CFD is to play an expanded role in the design and development research engineering community, development test and evaluation community, and ultimately certification for flight. This is especially true with hypersonic air-breathing propulsion systems due to the environment, scale, and duration limitations of ground test facilities.

The work performed here describes a process for uncertainty quantification that has been present in the literature for years,[3] but has perhaps not been readily adopted due to the perceived costs associated with a straight-forward application of the approach. The present effort uses reduced-order modeling as a cost savings tool to allow these concepts to be efficiently incorporated into the workflow process to illustrate how the chosen UQ process can affordably be incorporated into realistic research problems of interest to the hypersonics community. The model problem chosen involves a scramjet isolator, where the desire is to validate turbulence closure models in the presence of uncertainty. Accounting for uncertainties related to modeling of the facility environment, as well as those specific to the numerical simulation (but independent of the turbulence closure) allows the analyst to isolate the turbulence model deficiencies from other sources of uncertainty. In doing so, more informed decisions can be made concerning the level of model fidelity required for the application, while illuminating other error (or uncertainty) sources that may require resources

to address and/or reduce. This relatively simple (but relevant) model problem was selected over a more comprehensive task, $e.g.$, ground-to-flight mapping, in order to emphasize the UQ process without getting too bogged down into the details of a complex application of the process. However, a recipe that describes how this framework could be applied to more complex scenarios, in particular ground-to-flight mapping, is also provided.

UNCERTAINTY QUANTIFICATION FRAMEWORK

Given the prevalence of numerical simulations in the current research and engineering environment, where the trend has shown an ever increasing reliance on numerical data (with subsequent reduced reliance on test data), some sort of uncertainty quantification process should be expected as a standard practice. However, the trend in programmatic development (at least in the hypersonics community) has tended to utilize nearly all of the allotted computational resources ($i.e.$, time and people) in a deterministic fashion with the goal of maximizing the coverage of the design space, with less emphasis on quantifying the uncertainties associated with the analysis. Given the current balance of data provided from simulations relative to data provided by ground tests, this philosophy can no longer be justified. Hence, a change in culture is required; a prospect that as a general rule evolves slowly over time. As a result, even though there have been several recent government funded activities in uncertainty quantification methodology development,[4, 5] as well as the development of guides related to uncertainty quantification,[6] these ideas have not consistently found their way into the engineering development workflow. With this reality in mind, the framework described here is intended to be a minimalist approach (in terms of both complexity and time), that still retains the important basic features that those in the uncertainty quantification community tend to emphasize. In other words, the goal here is to document an approach that is within the reach of CFD practitioners that are not experts in uncertainty quantification, so that the benefits of a UQ-centric approach ($i.e.$, one that provides a means to communicate the confidence present in the CFD data) can be realized.

The uncertainty quantification framework chosen for this effort is that of Roy et al.[3] This framework was chosen because it is nonintrusive ($i.e.$, code modifications are not required to the CFD flow solver), and it properly handles both aleatoric and epistemic uncertainty sources. This framework has also been successfully adapted to other efforts aimed at quantifying uncertainty in complex engineering settings.[7, 8] The first crucial step in this UQ process is the determination and characterization of each source of uncertainty. The determination of potentially relevant uncertainty sources should not be taken lightly because the impact of any important uncertainty source missed here will not be accounted for, compromising the uncertainty quantification efforts. Hence, all members of the design and development team should be involved at some level during this critical initial stage of the process. Once the potential sources of uncertainty have been identified, the next step is the classification of each uncertainty as being either aleatoric or epistemic. Aleatoric uncertainty is any uncertainty source that can be considered as inherently random. These sources of uncertainty are typically not readily reducible, and are treated mathematically using probability distribution functions. Examples of aleatoric uncertainty sources for the present isolator simulations are the facility conditions (inflow stagnation conditions, outflow back pressure, etc.) and geometry variability due to manufacturing tolerances. Epistemic uncertainty sources can typically be regarded as some sort of bias error. These sources are commonly represented as being interval-bound with no belief that any single value within the interval is more likely to occur than any other value within the interval. Examples of epistemic uncertainties for the isolator simulations performed here are turbulence closure assumptions, onset of laminar to turbulent transition, and discretization error. These uncertainties, unlike aleatoric sources, can often be reduced as additional knowledge is gained. For example, turbulence closure uncertainties may be reduced by utilizing higher fidelity models, and discretization errors can be reduced by increasing the resolution of the computational grids employed in the simulations. The proper characterization of each uncertainty source can not be overemphasized. It may be tempting to employ what is referred to as a first-order uncertainty analysis where all sources of uncertainty are treated as aleatoric, however, treating epistemic uncertainties as aleatoric often leads to a drastic underprediction of their impact on the quantity of interest. The consequence of treating all uncertainty sources as aleatoric was illustrated in Ref. 9, where

the uncertainty in the prediction of radiative heat transfer of a reentry vehicle increased from 6% to 30% when the epistemic uncertainty sources were properly classified.

Once the uncertainties have been identified and classified, the next step is to perform a sensitivity analysis to estimate their influence on the system response (or output quantity of interest). The purpose of this step is to rank the relative importance of each uncertainty source toward the prediction of the system response. This information is used to potentially reduce the dimension of the uncertainty space by removing uncertainty sources that are deemed to have a negligible impact on the system response variability. Where this line is drawn is a function of the computational resources available to the CFD practitioner and the time that the "customer" is willing to accept for the additional information provided by an uncertainty analysis.

Given the set of relevant uncertainty sources being considered, the next step is to propagate them through the model to determine their impact on the system response. The manner in which this is done is dependent on the characterization of the uncertainty sources. As discussed previously, aleatoric parameters are random in nature and sampled based on an assigned probability. If multiple aleatoric uncertainty sources are considered, then each aleatoric parameter is independently sampled based on its Probability Density Function (PDF). These values are then passed to the computational model to obtain a single value for the quantity of interest. This process is repeated "N_a" times to obtain a random set of system response values from which a Cumulative Distribution Function (CDF) is formed. The CDF is a nondecreasing function that provides the probability that a random variable (in this case, the system response) will take on a value less than or equal to some chosen value. This function is formed by first sorting the system responses from its smallest to largest value forming the abscissa of Fig. 2. The CDF is then constructed by assigning a probability of $1/N_a$ to each response and adding them up as one traverses the abscissa from left to right. The end result is an "S" curve, as shown in the left image of Fig. 2. A steeper curve implies that the variability about the mean value (*i.e.*, the variance of the response) is relatively small and vice versa.

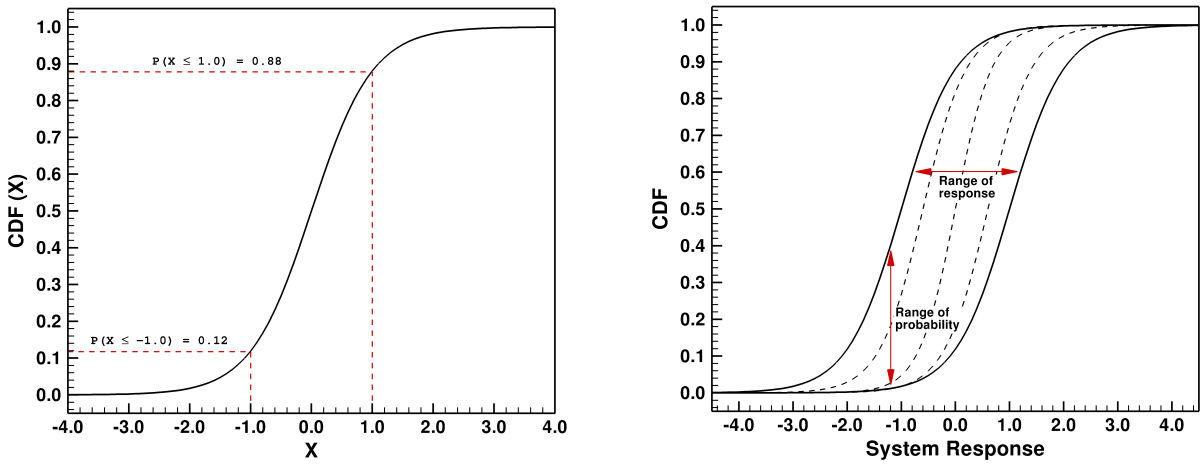

Figure 2: Representative CDF (left) and probability box (right) images

Epistemic parameters have no probability associated with them and instead one treats any value within the assumed interval as having an equal possibility of occurring. Therefore, propagating randomly sampled epistemic values (in the absence of any aleatoric uncertainties) simply results in a bounded interval for the system response, so that all system responses between the minimum and the maximum value obtained after propagation of "N_e" samples of the epistemic uncertainties are considered equally probable. The typical scenario in real-world problems is the situation where both aleatoric and epistemic uncertainty sources are present. In this case, one independently samples from each of the epistemic parameters prior to performing the uncertainty analysis for the aleatoric parameters as described in the previous paragraph to produce a CDF that represents a conditional probability. In other words, this CDF provides the probability of the system response "conditioned" on the specific epistemic samples that were chosen. Repeating the process for a different set of epistemic parameters would yield (in general) different CDFs. Therefore, when both

aleatoric and epistemic uncertainty sources are present, a second-order uncertainty analysis is required where a nested loop structure is formed with the epistemic parameters sampled in the outer loop, and the aleatoric parameters sampled in the inner loop. This process generates a family of conditional CDFs with the bounding curves representing a probability box, as shown in the image on the right of Fig. 2. The end result of this process is an interval-valued probability representation for the system response.

The computational expense of performing the process outlined above would be excessive if one attempted to propagate the uncertainties directly through the CFD software. A cost effective means of circumventing this difficulty is to construct a metamodel for the system response. Metamodeling is a technique by which a system response is approximated using a limited set of data over a defined design space (in this case, the aleatoric and epistemic parameters). This technique is also known as reduced-order modeling or emulation, but metamodeling is the preferred term here. The use of a metamodel is advantageous in situations where many calls to the computational model are required ($e.g.$, uncertainty propagation). Directly accessing the high-fidelity CFD solver in an on-demand manner during the uncertainty propagation process would simply not be practical due to the time required to obtain steady-state solutions from the discretized Navier-Stokes equations. Instead, a finite number of CFD simulations are performed at predefined locations within the design space. The metamodel is then formed by fitting the system responses extracted from the CFD solutions as some function of the design space.

FACILITY DESCRIPTION

The measured data used in this effort were obtained from experiments performed in the Isolator Dynamics Research Laboratory (IDRL) at the NASA Langley Research Center. A photograph and schematic of the facility hardware are given in Fig. 3. The flowpath is mounted vertically with facility flow directed from bottom to top. A unique aspect of this facility is the ability to translate and/or rotate the rig during facility operation. This capability was desired to simplify the application of advanced laser-based diagnostics for flowfield interrogation. The ability to move the test apparatus allows the diagnostic equipment to be set-up and aligned only once, so that the diagnostic components never have to be relocated. Air is supplied via 4 flexible hoses to a settling chamber that houses a honeycomb flow straightener to remove large-scale swirl and reduce freestream turbulence levels. The settling chamber connects to a two-dimensional Mach 2.5 nozzle that was designed using an inviscid method-of-characteristics code coupled with a two-dimensional full Navier-Stokes solver. The 24-inch isolator test section has a height (h) of 1-inch and a width (w) of 2-inches, and is assembled with interchangeable glass and metallic side walls. Hence, both wall measurements and laser-based in-stream measurements can be obtained simultaneously. The metal walls were installed for this effort, allowing for a total of 213 static pressure measurements in the arrangement shown in Fig. 4. The flow exits the isolator test section and enters a 15-inch long cylindrical expansion section that is machined with a 45° chamfer on the downstream end to mate with a back-pressure plug. The back-pressure plug is mounted in the spool section to a linear actuator that positions the plug with a precision of 1/1000 of an inch. The back-pressure plug simulates the thermal choking effect caused by combustion in a dual-mode scramjet engine, providing a means of creating the back-pressured isolator flowfield of interest. A detailed discussion of the facility is given in Ref. 10.

SIMULATION DESCRIPTION

The computational domain for this effort spans the region from the facility nozzle plenum to the end of the 24-inch isolator duct. The geometrical symmetry present in this flowpath, and the fact that only steady-state simulations are to be performed, permitted the use of only one quarter of the facility flowpath, as shown in Fig. 5. A structured grid with a total of 85,196,800 cells was created for the simulations with 18,636,800 cells placed in the nozzle block ($560 \times 160 \times 208$) and the remaining 66,560,000 cells ($2000 \times 160 \times 208$) in the isolator duct. Two coarser grids (10,649,600 and 1,331,200 cells) were obtained from this parent grid

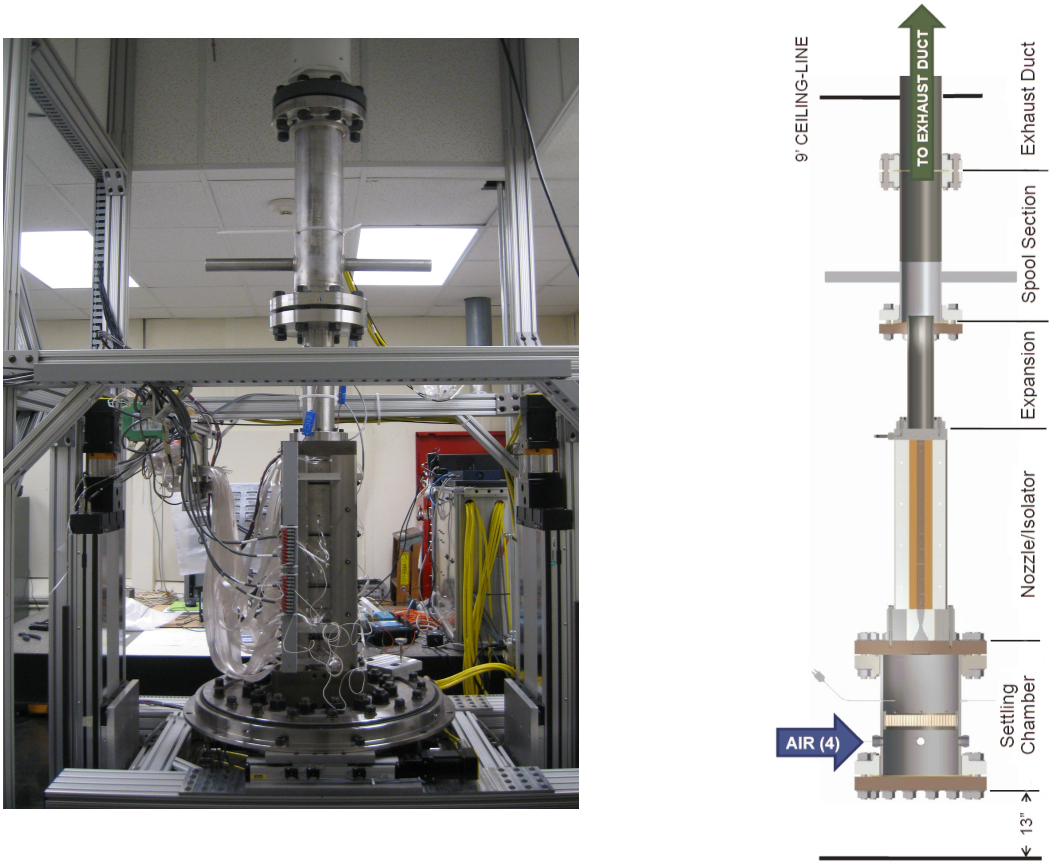

Figure 3: Photograph (left) and schematic (right) of the Isolator Dynamics Research Laboratory hardware

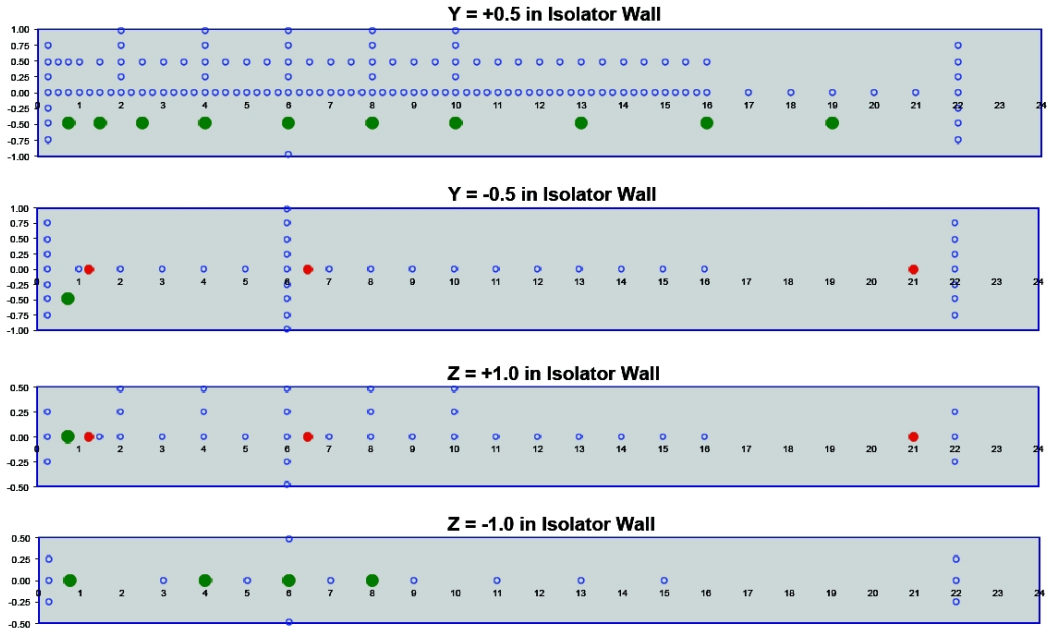

Figure 4: Instrumentation layout for the metallic isolator side walls (flow direction would be left to right): static pressure ports (blue), thermocouples (red), high frequency pressure ports (green)

via a factor of 2 coarsening in each coordinate direction to allow for an assessment of the numerical errors associated with the simulations. The grid points were clustered near the walls such that the maximum y^+ value at any cell center adjacent to the surface was no greater than 1 (approximately) for the coarsest grid utilized. The grid was also clustered in the streamwise direction at the nozzle throat (Δx = 0.0025 h) with the spacing gradually increased to 0.01 h at the nozzle exit. The streamwise spacing was held fixed at 0.01 h for the front half of the isolator to resolve the region where the shock train was expected to reside, and was stretched in the second half of the isolator to a value of 0.025 h at the isolator exit. The maximum grid spacing in the vertical (y) and spanwise (z) directions was 0.01 h (adjacent to the planes of symmetry). The entire grid was split up into a total of 1280 grid blocks, which yielded 95% or better load balance statistics for up to 438 processors.

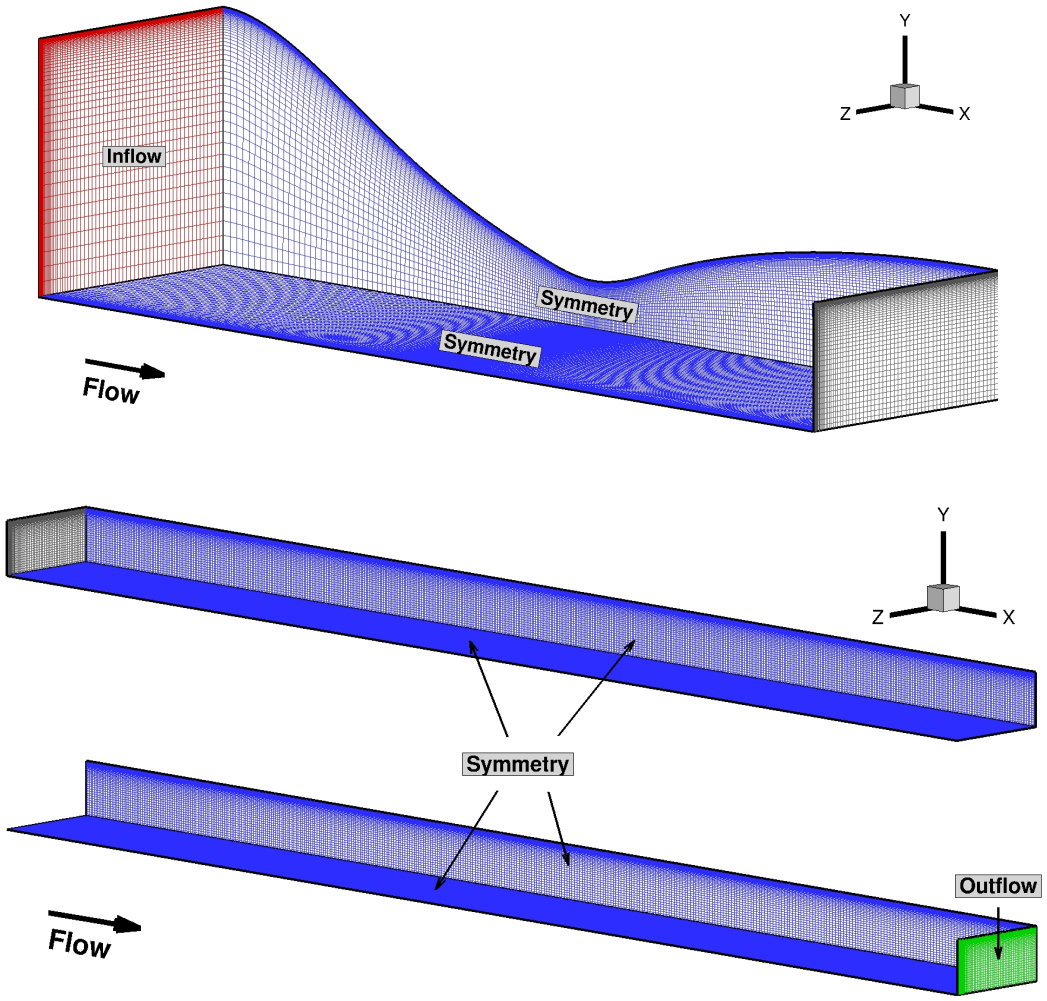

Figure 5: Grid created for the nozzle/isolator assembly (every other point removed for clarity): nozzle section (top), first half of isolator section (middle), second half of isolator section (bottom)

All of the CFD simulations were obtained using the VULCAN-CFD (Viscous Upwind aLgorithm for Complex flow ANalysis) Navier-Stokes code.[11] This CFD package is developed and maintained by researchers in the Hypersonic Airbreathing Propulsion Branch at the NASA Langley Research Center, with the development of this software being driven predominantly by the needs of the scramjet community. The CFD data obtained in this effort were acquired by integrating the Reynolds-Averaged Navier-Stokes equations until steady-state conditions were achieved. The solutions were advanced in pseudo-time via an incomplete LU factorization scheme (with planar relaxation) using a Courant-Friedrichs-Lewy (CFL) number of 100. The inviscid fluxes

were evaluated using the Low-Diffusion Flux Splitting Scheme of Edwards,[12] with cell interface variable reconstruction achieved via the κ=1/3 Monotone Upstream-centered Scheme for Conservation Laws. The Koren flux limiter[13] was utilized to avoid spurious oscillations during this reconstruction process. The viscous fluxes were evaluated using 2nd-order accurate central differences with the viscosity of air computed from the polynomial fit of McBride.[14, 15] The Menter baseline (BSL) closure for the Reynolds stress tensor was the turbulence model chosen for this validation study, and the molecular and turbulent Prandtl numbers were set to 0.72 and 0.9, respectively.

The boundary condition imposed along all solid surfaces was an adiabatic no-slip condition. The subsonic nozzle plenum inflow condition was specified by fixing the plenum pressure and temperature, with Mach number extrapolated from the interior. Finally, all properties were extrapolated at the subsonic exit of the isolator with the exception of the static pressure. The surface pressure at the isolator exit was defined by averaging the 20 surface pressure measurements taken around the periphery of the isolator duct at the x = 22 h station. This station was highly instrumented with pressure taps for the specific purpose of providing a well characterized outflow boundary condition. The static pressure profile applied at the exit of the isolator was then set according to the following relationship to account for the influence of the turbulence kinetic energy on the static pressure profile:

$$P = P_s - \frac{2}{3}\rho_i k_i \tag{1}$$

where P_s is the experimentally obtained surface pressure, and $\rho_i k_i$ is the product of density and turbulence kinetic energy extrapolated from the interior.

Solution Verification

Numerical errors are denoted here as spatial discretization errors and iterative convergence errors. Temporal discretization errors are of no consequence for this study since only steady-state simulations were performed in this effort. The most compelling proof of iterative convergence is the situation where the residual norms have been reduced to machine accuracy. Unfortunately, this situation is rarely realized in practice, particularly when simulating complex flows with shock waves and regions of flow separation. Alternatively, the L_2 norm of the residual is monitored until it levels out about some acceptably low value. In addition, the variability of integral properties of interest, such as mass flow rates and surfaces loads, are checked to offer further assurance that an adequate level of convergence has been achieved. In general, the following iterative convergence statements were satisfied for each simulation performed in this effort:

- The value of the residual L_2 norm was reduced by at least 5 orders of magnitude.
- The integrated surface load time histories were unchanged (to at least 5 significant digits) over the last 5000 iteration cycles.
- The mass flow rate was constant at every streamwise grid plane to at least 4 significant digits.

The Grid Convergence Index (GCI)[16] was used to quantify the spatial discretization (*i.e.*, grid convergence) error. The GCI is a grid convergence estimator derived from the generalized Richardson Extrapolation formula, and can be written as follows:

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

where f_2 - f_1 is the difference of some functional, f, evaluated using two different grid resolutions (with refinement ratio r); p is either the observed order of accuracy of the numerical scheme or taken to be the theoretical value (2 for the 2nd-order schemes used here); and F_s is a "factor of safety" with recommended values of 3 (if the observed order of accuracy is taken to be the theoretical value) or 1.25 (if the observed order of accuracy has been rigorously determined). The spatial discretization error estimates were found to be orders of magnitude larger than the iterative convergence errors, so the iterative convergence errors were omitted from the uncertainty assessment.

RESULTS AND DISCUSSION

A sample image of the simulated isolator flowfield is presented in Fig. 6. The global isolator shock train structure (visualized by the velocity divergence) consists of leading oblique shock waves followed by alternating expansions and quasi-normal shocks. The top image shows static pressure contours superimposed on an isosurface of the total pressure (P_o = 300 kPa). This image provides a visual of the core flow displacement caused by the interactions between the shock waves and the turbulent boundary layer that engulfs the isolator walls. The area blockage caused by the separated flow in the corners is a dominant feature that dictates (to a large extent) the shock train flow structure. The Mach number contours show a boundary layer structure on the vertical side walls that is notably thicker near the center of the side wall. This feature is a result of the crossflow pressure gradients imposed on the flat sidewalls of the facility nozzle, and makes this particular boundary layer more susceptible to large-scale flow separation. Mach disks are also seen in the center of the duct, and while it is difficult to visualize in this image, the interaction of the leading oblique shock wave and the Mach disk results in the formation of a shear layer. The flow in the isolator initially separates in the corners of the duct where the shear forces are weakest. This flow separation results in a conical shock that spreads diagonally across the isolator, diverting the flow toward the center of the duct. The high pressure that is produced in the corners also drives the boundary layer fluid away from the corners along the adjacent walls. The side walls (z = constant vertical surfaces) are half the width of the top and bottom walls (y = constant spanwise surfaces), so the effect of the adverse pressure gradient formed in the corners influences the center of the vertical walls more quickly than their spanwise counterparts. The additional width offered by the spanwise surfaces allows sufficient space for a second separation region to form at the z=0 centerplane of the duct. This separation bubble forms a strong oblique shock followed by a strong expansion as this separation zone reattaches to the surface. The resulting flow remains attached along the isolator centerplane even after further shock reflections downstream. The separation zones that originate in the corners, however, gradually extend across the side walls and persist for several duct heights further downstream. Moreover, the increased turbulence levels present in these separated flow zones do not dissipate as the flow progresses downstream. The eventual result of this unbalanced production of turbulence is a very large “effective” viscosity across the entire width of the duct that is 20,000 times larger than the local fluid viscosity. The net effect is a flowfield that exits the isolator with properties representative of a Reynolds number that is 4 orders of magnitude lower than the flow entering the isolator. These extreme values for the turbulent viscosity allow the Reynolds-averaged simulations to produce a steady-state solution for what would otherwise be a dynamic flow.

Turbulence Model Validation

The conditions of interest for the validation exercise are isolator conditions that result in a pressure ratio across the isolator shock train that approaches the operability bound of dual-mode scramjets. The IDRL isolator shock train length (L_{iso}), plotted as a percentage of the normal shock pressure ratio (based on the nominal isolator entrance Mach number of 2.5), is given in Fig. 7. The shock train length is defined as the distance from the location of the incipient shock-induced surface pressure rise to the exit of the isolator. This plot shows that the ability of an isolator to hold a given pressure ratio diminishes rapidly as the normal shock pressure ratio is approached. For these isolator conditions, a 9% increase in the back pressure from 82.25% to 89.82% of the normal shock pressure rise changes the shock train length by 8.6 x/h (or 35% of the available isolator length). The sensitivity of the shock position to the back pressure value in this regime is clearly a challenge for any simulation effort to accurately predict. This sensitivity also explains why dual-mode scramjets have historically been designed to limit the total combustor-induced pressure rise during low-speed operation to approximately 80% of the normal shock value based on the isolator approach flow Mach number.

The back-pressured facility isolator conditions chosen for the validation exercise are the experimental conditions provided in Table 1. The measured conditions in this table, and the measured surface pressures in the plots that follow, are an ensemble average of measurements taken over an approximate 5 second test in-

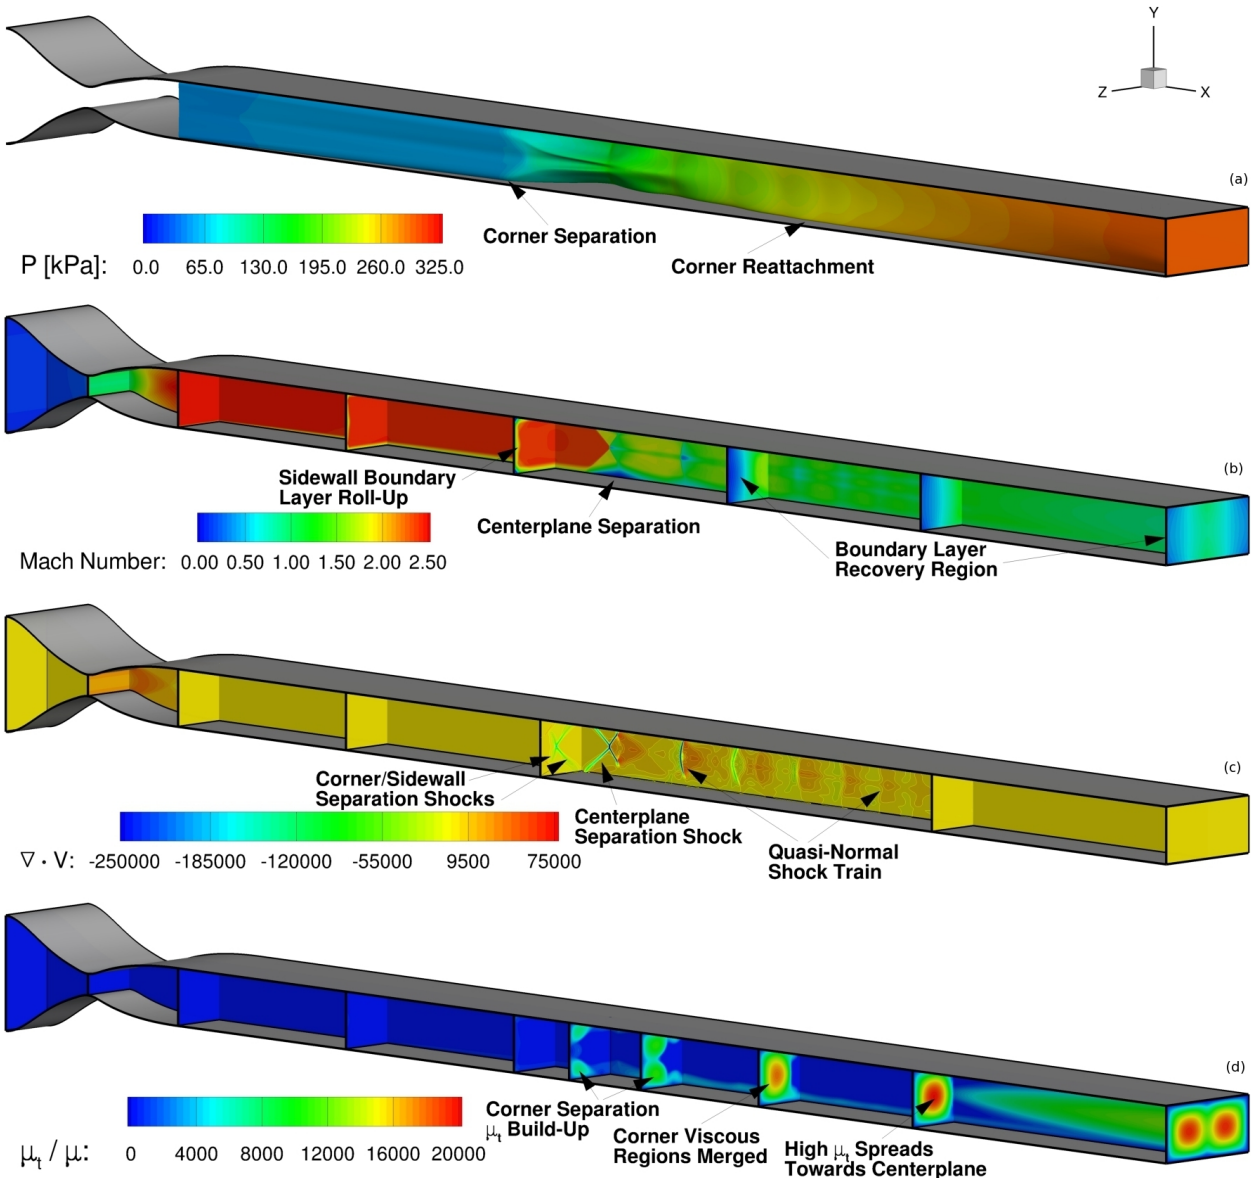

Figure 6: Isolator flowfield characteristics: (a) total pressure isosurface colored by static pressure contours, (b) Mach number contours (c) velocity divergence contours, (d) turbulent to molecular viscosity ratio contours

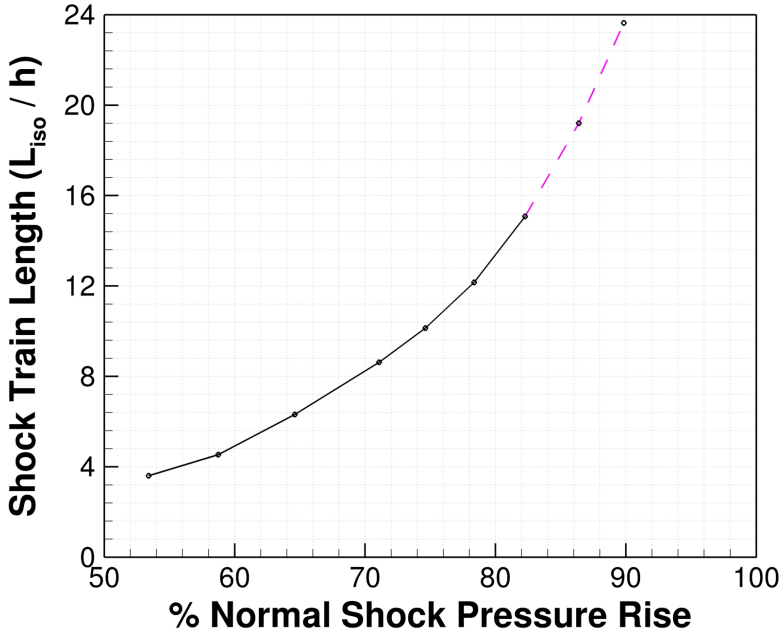

Figure 7: Isolator shock train length as a function of the normal shock pressure rise for a Mach 2.5 flow (dashed portion of the line represents an operability range beyond that of current designs)

terval. The values in parenthesis represent the RMS deviation in the measured values relative to the mean value obtained over the 5 second averaging interval. Note also that the isolator exit pressure measurements were taken at the x = 22 h station. The computational domain exit is located at the x = 24 h station, so the back pressure applied at the outflow boundary was increased slightly to 311.56 kPa to account for the additional pressure rise over the final 2 duct heights of the isolator.

Table 1: Nominal facility flow conditions selected for model validation

Facility Property	Operating Condition
Nozzle Plenum Pressure [kPa]	908.68* (0.027)
Nozzle Plenum Temperature [K]	292.91 (0.023)
Nozzle Exit Mach Number	2.5*
Nozzle Exit Pressure [kPa]	53.90
Isolator Exit Pressure [kPa]	304.51 (0.097)

* Value computed from isentropic relationships

The wall static pressures obtained from the CFD model at the nominal facility conditions described above are compared with the measurements in Fig. 8. The computed results shown are those obtained from a simulation that used the medium (10,649,600 cell) grid. The image on the left compares the pressure along the z=0 centerline of the top and bottom walls, while the image on the right compares the pressure along the y=0 centerline of the side walls. In general, the measured pressure distribution along the spanwise centerline (z=0) shows a monotonic rise, while the simulation does not; indicating a discrepancy in the separation zone structure along these surfaces. However, the computed side wall y=0 centerline pressure values display a monotonic rise, which is consistent with the measurements. The simulation has also reproduced the measured shock train length to within a fraction of a duct height. If one were to use this single deterministic simulation as the sole basis for assessing the adequacy of the chosen turbulence model, then one would be inclined to claim that the turbulence model has adequately predicted the shock train

length. However, a multitude of uncertainty sources exist that affect the ability of the simulation to replicate the measured isolator flowfield. As a result, the validation exercise should be performed in the presence of these uncertainty sources prior to making any definitive claims on the adequacy of the turbulence model for the chosen application.

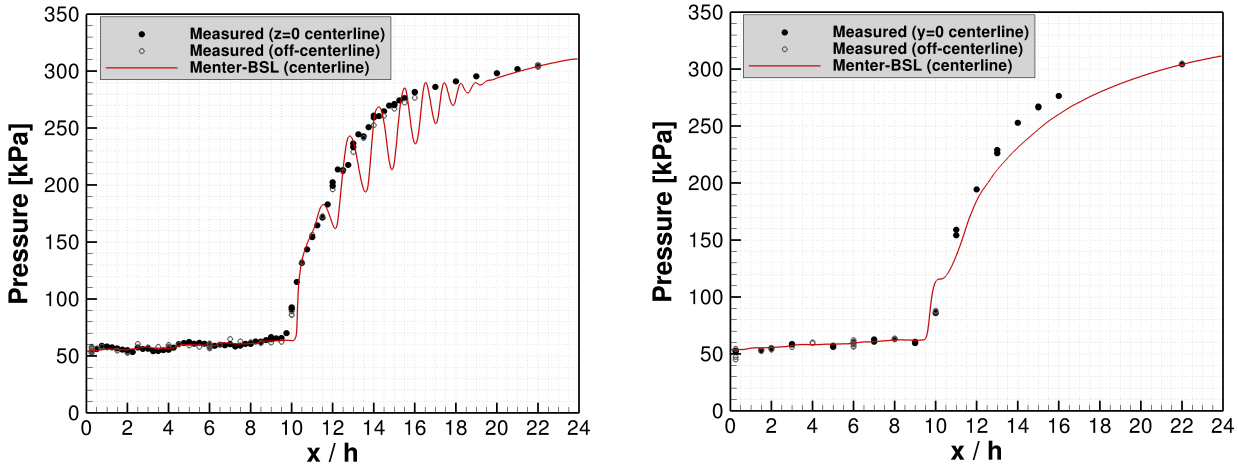

Figure 8: Comparison of measured wall pressure with computed results (medium grid): top and bottom walls (left), side walls (right)

As mentioned previously, the first step in the uncertainty quantification process is the determination and characterization of the uncertainties. Given knowledge of the isolator flow physics and the simulation models employed, the following potential sources of uncertainty were considered to examine their impact on the isolator effectiveness as defined by the shock train length (L_{iso}):

- spatial discretization error - EPISTEMIC
- onset of laminar to turbulence transition - EPISTEMIC
- facility supply pressure variability - ALEATORIC
- facility supply temperature variability - ALEATORIC
- facility supply turbulence intensity (0.01% to 5%) - EPISTEMIC
- facility wall temperature - EPISTEMIC
- isolator exit pressure variability - ALEATORIC

Given the isolated flow physics targeted by this validation exercise, the sensitivities in this effort were assessed by simply observing the variability in L_{iso} for a one-at-a-time variation of each uncertainty source. A more thorough process[7] would be required for more complicated scenarios, where the CFD practitioner has little $a\ priori$ knowledge of how the response might vary with each uncertainty parameter. Based on the sensitivity assessment, and given the emphasis of illustrating the UQ process (rather than enhancing the refinement of uncertainty margins), the uncertainty sources considered for this turbulence model validation exercise were limited to facility supply pressure, isolator exit pressure, transition onset, and numerical discretization error. The uncertainty space was reduced even further by recognizing correlations in the uncertainty sources. In particular, the ratio of the isolator exit to facility supply pressure was treated as a single uncertainty source, which was hypothesized through subject matter expertise and justified via sensitivity studies.

Having defined and characterized the relevant sources of uncertainty, the next step is to build the metamodel that will be used for the uncertainty propagation. The design space used to construct the metamodel must include the entire range of allowable values, requiring bounds to be defined for each uncertainty source considered. The facility pressure ratio uncertainty was assumed to span the range of values between the mean value and $\pm$ 3 σ, where σ is the standard deviation taken from the measured temporal histories. The onset of transition is an epistemic uncertainty that results from our "lack of knowledge" as to precisely where the flow transitions in the nozzle. In this effort, transition was assumed to take place somewhere between the nozzle entrance ($i.e.$, assumed turbulent everywhere) and the streamwise location of the nozzle inflection point downstream of the throat. If additional measurements were to be taken at the nozzle exit that provided information on the state of the boundary layer, then these uncertainty bounds could potentially be reduced. Finally, the spatial discretization error was accounted for as a categoric uncertainty using two grid levels. This information will later be transformed into a GCI measure to estimate the bounds of the discretization error. The design space for the metamodel developed for the uncertainty propagation is provided in Table 2. Note that the metamodel is set up to accept coded inputs ($i.e.$, the actual inputs are scaled to map between -1 and +1). By doing so, the coefficients of the model provide a more faithful estimate of the relative impacts of each input parameter.

Table 2: Inputs for the metamodel of isolator shock train length

Metamodel Input	**Classification**	**Type**	**Valid Settings**
Facility Pressure Ratio	Aleatoric	Continuous	[-1, +1]
Transition Location	Epistemic	Continuous	[-1, +1]
Grid Level	Epistemic	Categorical	L2 (-1), L1 (+1)

In the present study, the metamodel inputs were sampled using design of experiments methods. Design of experiments is a rigorous methodology for sampling a multidimensional design space in order to maximize the amount of information with a minimal number of samples. The benefits of using design of experiments become greater as the dimension of the design space increases relative to a full factorial method that utilizes regularly-spaced sampling. The strategy used in this effort was to first initialize the design space such that the corners and central locations were explored in an efficient manner. To do so, an I-Optimal design (which seeks to minimize the average prediction variance over the entire design space) was initialized with 9 samples in order to provide a perfect model fit with no degrees of freedom remaining for error characterization. Subsequently, four groups of five space filling samples were defined to gradually refine the population of samples within the design space to characterize the metamodel error. The space filling method used for this purpose was the Fast Flexible Filling approach.[17] At all times, the correlation between every design variable combination was kept less than 5% to ensure minimal metamodel correlation due to the sampling itself. Figure 9 shows the design space sample scatter for the facility pressure ratio and transition location for all grid level settings. Only the I-Optimal design ($i.e.$, the initial 9 samples) and the first generation of additional samples are shown because the additional generations were not used to refine the model in the present work.

Having defined the points in the uncertainty space to be sampled (14 total), the CFD simulations were performed to provide the 14 discrete responses (shock train length) to be fit by the metamodel. Without having a physics-based model form against which to fit the data, the classical approach to metamodeling of a system response is to use an empirical, 2nd-order model equation of the form:[18]

$$R(\mathbf{x}) = b_0 + \sum_{i}^{n} b_i x_i + \sum_{i}^{n} b_{ii} x_i^2 + \sum_{i}^{n-1} \sum_{j=i+1}^{n} b_{ij} x_i x_j \tag{3}$$

where n is the number of design variables. Higher-order terms may also be included to model additional nonlinear effects if desired. While the number of regression coefficients in a model is equal to the minimum number of samples required to produce a perfect fit with no characterization of error, best practice dictates

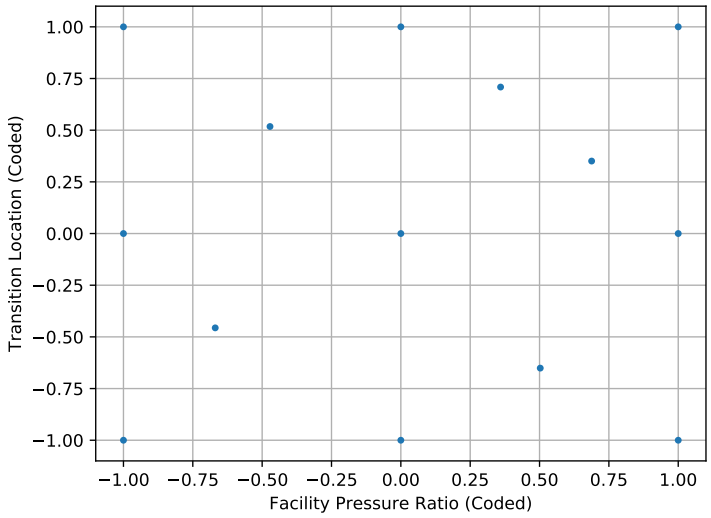

Figure 9: Design space points used to fit the metamodel shown for transition location versus facility pressure ratio for all grid levels

that a number of additional samples be used to produce a least squares fit representing the mean behavior of the response along with some modeling error. Thus, the final fitted form of the metamodel is

$$\hat{R}(\mathbf{X}) = [1, x_1, x_2, \ldots, x_1^2, x_2^2, \ldots, x_1 x_2, \ldots] \begin{bmatrix} \hat{b}_0 \\ \hat{b}_1 \\ \hat{b}_2 \\ \vdots \\ \hat{b}_{11} \\ \hat{b}_{22} \\ \vdots \\ \hat{b}_{12} \\ \vdots \end{bmatrix} + \epsilon = \mathbf{X^T}\,\hat{\mathbf{b}} + \epsilon \tag{4}$$

which includes an error term to capture scatter of the data about the mean response. The model parameters fit to this metamodel for shock train length are given in Table 3 along with a measure of the significance of each term. Smaller values in the 3rd column imply a larger significance. The 1st-order grid level term is the greatest contributor to the variability of the model (by an order of magnitude), followed by the 1st-order and 2nd-order transition location terms. The interaction between facility pressure ratio and transition location was of negligible importance, as was the interaction between the facility pressure ratio and the grid level. The interaction between the transition location and grid level and the 2nd-order facility pressure ratio term were also insignificant. While these could be removed from the model, thereby yielding additional available degrees of freedom for characterizing error, they were retained for the subsequent analysis. Pure categorical terms (*e.g.*, grid level) do not have 2nd-order terms, hence its omission from the table. Appendix A contains additional information on the processes used to assess the metamodel prediction error (*i.e.*, ϵ in Eq. 4) and the adequacy of the metamodel fit. The metamodel prediction error was included in the uncertainty quantification processes as an aleatoric uncertainty (normally distributed about the mean). The mean contribution of this error was 0.2% of the predicted response (L_{iso}) value.

The metamodel provides a simple low-cost functional relationship between the shock train length and the uncertainty (design variable) space, allowing for the efficient propagation of the uncertainty sources. This is

Table 3: Metamodel regression coefficients and statistical significance measures for each term

Metamodel Term	**Regression Coefficient**	**Probability** $>$ |**t**|
Intercept $[\hat{b}_0]$	13.5	$<$ 0.0001
Facility Pressure Ratio $[\hat{b}_1]$	0.0527	0.0012
Transition Location $[\hat{b}_2]$	-0.0805	0.0002
Grid Level $[\hat{b}_3]$	0.285	$<$ 0.0001
Facility Pressure Ratio * Facility Pressure Ratio $[\hat{b}_{11}]$	0.0159	0.281
Transition Location * Transition Location $[\hat{b}_{22}]$	-0.105	0.0005
Facility Pressure Ratio * Transition Location $[\hat{b}_{12}]$	-0.00271	0.791
Facility Pressure Ratio * Grid Level $[\hat{b}_{13}]$	-0.00567	0.5126
Transition Location * Grid Level $[\hat{b}_{23}]$	-0.00838	0.3422

the key cost saving step that enables the use of this UQ framework for real problems of engineering interest. In addition to the overall uncertainty provided by the formation of the probability box, the metamodel can also be used to construct conditional CDFs to examine the isolated impact of each epistemic uncertainty source. For example, the variability in the shock train length due to the uncertainty of where transition occurs in the facility nozzle is shown in the left image of Fig. 10 with the grid fixed at the L1 level (10,649,600 cells). This image shows that the shock train length varies by just over 0.2 duct heights depending on the location assumed for the transition onset. The right image of this figure shows the variability of the shock train length due to the use of either the L1 or L2 (1,331,200 cells) grid when transition is assumed to occur immediately (*i.e.*, fully turbulent). Since the grid level was a categoric uncertainty source, there are two discrete sets of CDFs with nothing in between. This image shows the shock position varies by approximately 0.6 duct heights depending on whether the L1 or L2 grid is used, suggesting that the shock train length uncertainty due to grid resolution is about 3 times larger than that due to the unknown location of laminar to turbulent transition. Note that conditional one-at-a-time comparisons are typically used for qualitative assessments only because the variability of the chosen epistemic parameter is computed assuming that all other epistemic sources are fixed. In reality, the amount of variability is contingent on the precise values assumed for the fixed epistemic uncertainty sources.

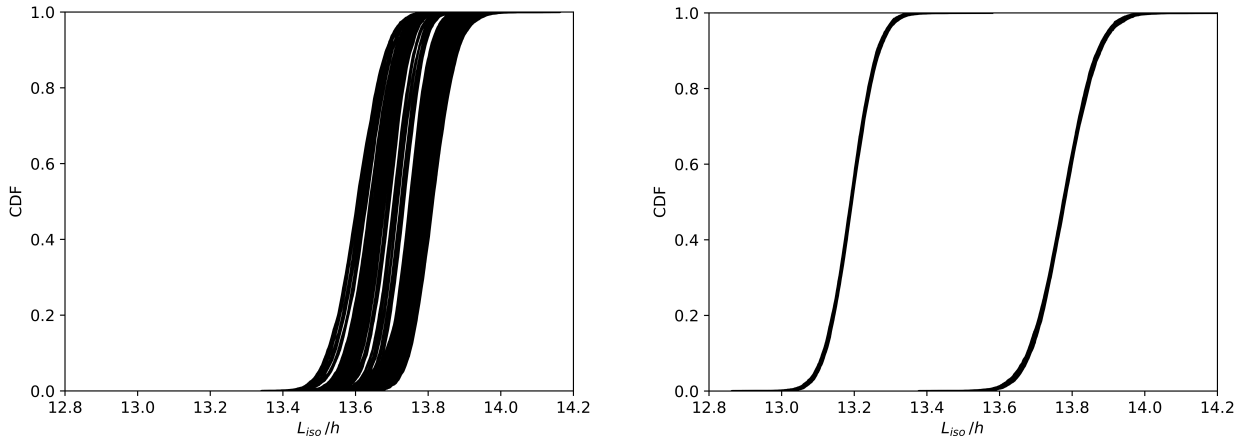

Figure 10: Conditional CDFs showing shock train length variability due to transition location (left), grid level (right)

The probability box that maps the design space uncertainties to a system response uncertainty is shown in the left image of Fig. 11. This plot provides the uncertainty estimate for the shock train length based on

uncertainties in facility pressure ratio, transition onset, metamodel error, and discretization error between the L1 and L2 grid levels. However, this latter effect only accounts for the numerical error between the two grid levels. What is really desired is the uncertainty in the shock train length due to discretization error. This information can be obtained by transforming the uncertainty between the L1 and L2 grid levels to a GCI measure (see Eq. 2), *i.e.*,

$$\text{CDF}(\text{L}^*) \;=\; \text{CDF}(\text{L1}) \;+\; F_s \,\frac{|\text{CDF}(\text{L1}) - \text{CDF}(\text{L2})|}{r^p - 1} \tag{5}$$

A rigorous solution verification exercise has previously been performed[19] for this isolator flowfield using the L1 and L2 grids along with the fine (L0) grid level. The additional grid level allowed tests to be performed to determine if the output quantities of interest were in the asymptotic convergence range, and the results of these tests suggested that this was indeed the case. This result allowed for the determination of the observed order of accuracy, which was performed for several flowfield metrics (including shock train length), and found to be 1 (*i.e.*, 1st-order) for this shock dominated flow. As a result of formally determining the observed order of accuracy, the factor of safety, F_s, that appears in Eq. 5 was taken to be 1.25. Moreover, the fact that this metric was found to be in the asymptotic convergence range allowed the treatment of discretization error as a true one-sided bias error. Hence, the absolute value of the difference between the L2 and L1 CDF values in Eq. 5 can be removed. From the perspective of the UQ framework, this additional knowledge allows this uncertainty source to be treated as a one-sided uncertainty, since it is known that further grid refinements will reduce the discretization error. The probability box that results from this transformation is shown in the right image of Fig. 11, where the red CDFs are the result of the GCI transformation. As a final note, the formal proof of CFD results being within the asymptotic convergence range is often an impractical task. Instead, one might be tempted to simply examine the CFD results on more than two grid levels to determine if the quantity of interest is monotonically changing as the grid is refined. While this provides an indication that the error due to finite grid resolution is continuously decreasing as the grid is refined, it is not certain. Hence, the CFD practitioner should consider this possibility when making decisions on how to handle this particular source of uncertainty if the order of accuracy is not formally assessed.

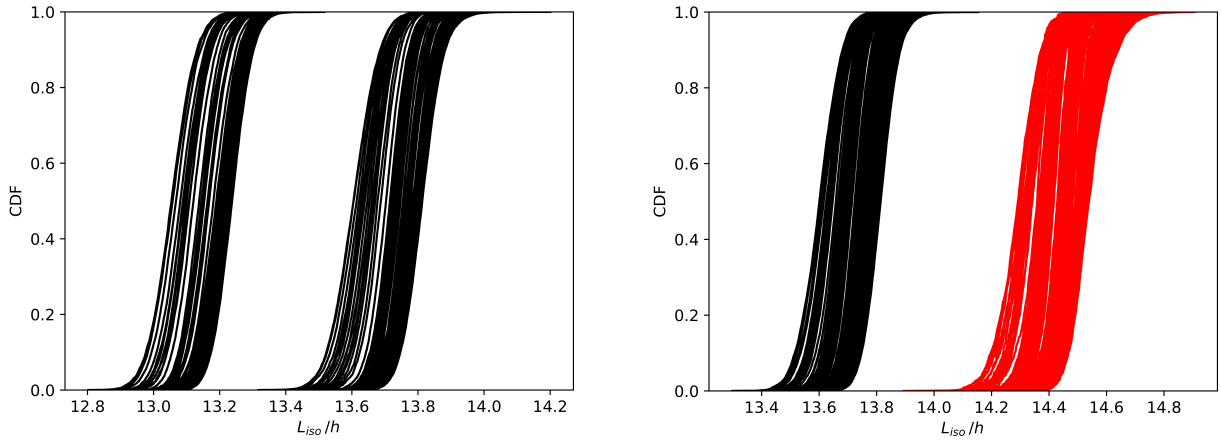

Figure 11: Probability box showing shock train length variability due to the raw uncertainty space (left) and GCI transformed uncertainty space (right)

The GCI transformed probability box, shown in the right image of Fig. 11, defines the uncertainty window of this validation exercise. Hence, if the measured shock train length is completely contained within the bounding CDFs in this image, then given our current state of knowledge, the turbulence model is validated with no measurable model-form error. An examination of this probability box shows that, given our current state of knowledge, the uncertainty in computing the shock train length is approximately 1 duct height. A significant portion of this uncertainty (nearly 50%) can be removed if the L0 and L1 grids were to be used

for the simulations. However, this grid pair utilizes 85,196,800 cells and 10,649,600 cells as opposed to the 10,649,600 cell L1 and 1,331,200 cell L2 grids, requiring an order of magnitude more cost for each CFD simulation used to build the metamodel. The extraction of information such as this is made possible by performing the validation within a UQ framework. The ability to compare costs and benefits is clearly a very valuable feature when making decisions on where to allocate resources, and what level of improvement to expect if given the additional resources.

The uncertainty information obtained from this UQ approach can also be displayed in a more conventional manner as a line plot with uncertainty bars. In this case, we know that the discretization error is a bias error, so the distance from the right-most CDF of the L1 grid family of curves to the right-most CDF of the GCI transformed family is solely the result of this bias error. The remaining uncertainty is contained within the width of the family of L1 CDF curves. This uncertainty space contains the portion of the discretization error that is mixed in with the rest of the epistemic parameters. Therefore, the uncertainty in this region is taken to be equally probable. The scenario just described is displayed pictorially in the left image of Fig. 12, and results in an asymmetric error bar for the CFD simulation performed at the nominal conditions of Table 1. The uncertainty results displayed in this manner, taken at the 50% aleatory uncertainty level, are shown in the right image of Fig. 12.

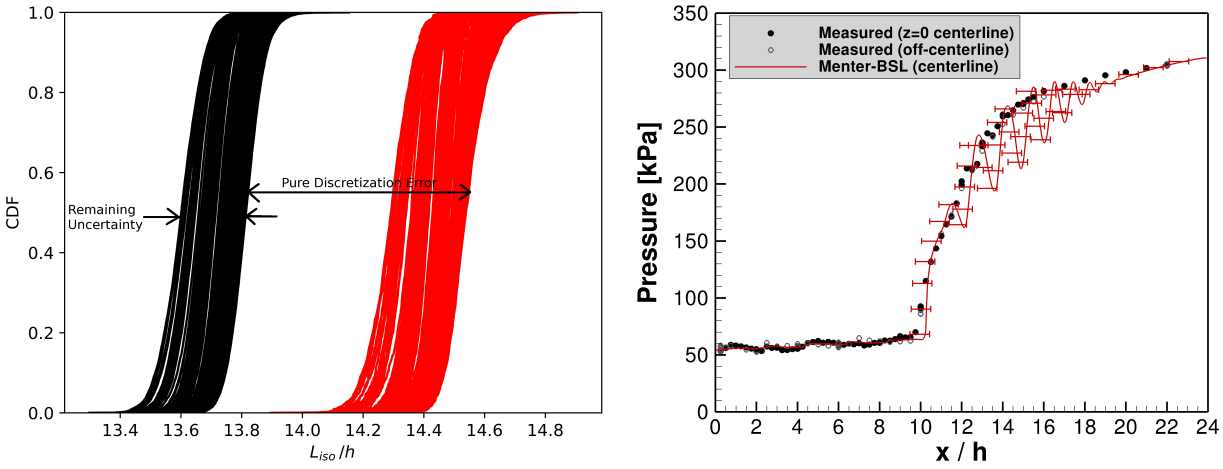

Figure 12: Comparison of measured wall pressure with computed results (L1 grid) with computational uncertainty shown

An examination of Fig. 12 shows that the measured shock train length is 14 h with an approximate uncertainty of 0.25 h (based on the proximity of adjacent pressure taps). This value lies completely within the probability box (left image of Fig. 12) of the CFD prediction, indicating that the turbulence closure model has been validated for this particular isolator condition with no evidence of model-form error. The question then arises as to how to handle the scenario when the measured results lie outside of the probability box of the prediction. Roy et al.[3] provides a detailed discussion on how this is dealt with, so the present discussion will be limited to the scenario where only a single instantiation of the measurements is available for validation purposes (which is commonly encountered in practice). This situation is shown schematically in the left image of Fig. 13, where the measured value is represented by the vertical red line, *i.e.*, a probabilistic representation of the measurement is not available, so the measured value is treated as a deterministic quantity. To assess the degree of disagreement, Roy et al.[3] suggests the use of an area validation metric defined as the area between the probability box and the measurement (shown as the shaded region and denoted as "A" in this image). This metric provides a measure of the level of disagreement between simulation and measurements in the presence of uncertainty, which can be used to to expand the probability box to account for model-form uncertainty (as shown in the right image of Fig. 13). This evaluation should be repeated as additional validation data is gathered to cover the range of conditions that the model is intended

to be validated for. Upon completion of the validation exercise, the largest area validation metric is used to pad the probability box of the simulations performed for any future predictive purposes. This represents the model-form uncertainty of the prediction based on the current state of knowledge provided by validation data.

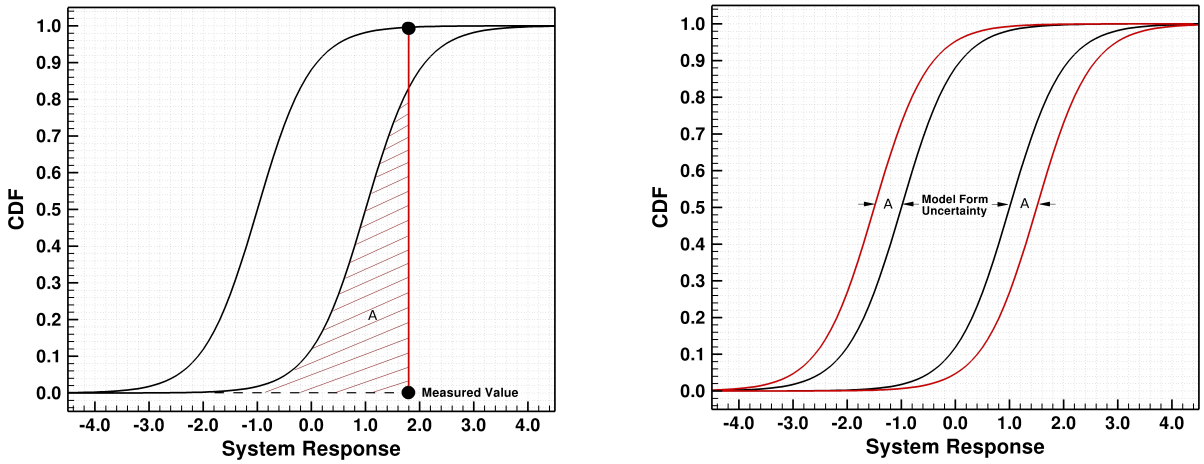

Figure 13: Schematic of the area validation metric (left) and a padded probability box that accounts for model-form uncertainty (right)

Ground-to-Flight Mapping

Although not a part of this effort, it is worthwhile stepping through how this same UQ framework might be applied to more complicated scenarios such as the ground-to-flight mapping of a scramjet propulsion system developed as part of a flight test program. This process represents the critical step of taking all of the knowledge gained from ground tests (measurements and simulation data) to estimate the performance and operability of the propulsion system in the true flight environment. The steps that follow are described in the context of a typical flight test development program where some combination of direct-connect, semi-freejet, and freejet tests (typically performed in this chronological order) are employed.

- Exercise the steps documented in this effort to validate the submodels used to simulate the direct-connect tests of the scramjet propulsion system. This process will inform the choices made for all of the relevant submodels required to predict the flow physics associated with the direct-connect tests. At the end of this campaign, a validated CFD model (in the presence of uncertainty) will have been developed with "best estimates" for the categorical choice of turbulence model, mixing model calibrations ($e.g.$, Sc_t and Pr_t values), categorical choice of chemical kinetic model (or perhaps calibration of a mixing controlled heat release model), etc. If the calibrated CFD model is unable to reproduce the measured quantities of interest, then model-form uncertainty increments should be added to the probability box and carried forward into the semi-freejet and freejet testing phases.
- Apply the CFD model validated in the direct-connect testing environment to the freejet (or semi-freejet) tests. The freejet tests introduce additional physical processes that were not considered in direct-connect tests (most notably inlet flow distortion), which may result in predictions that are outside of the model-form uncertainty estimates determined from the direct-connect tests. Moreover, the scale of the test apparatus may not match that of direct-connect testing ($e.g.$, 10X propulsion system development). The uncertainty environment associated with the new facility will also be different, potentially introducing new sources of uncertainty and altering the bounds of coexisting uncertainty sources. Given these new uncertainty inputs and bounds, one would assess (using the UQ framework) whether the CFD model developed in the direct-connect environment predicts the measured

quantities of interest within the uncertainty intervals defined by the freejet environment (including any model-form uncertainties determined during the direct-connect tests). If not, then either expand the model-form uncertainty increments, and/or recalibrate the submodel(s) using the more representative flight environment offered by the freejet tests. The freejet testing stage provides the last opportunity to add knowledge to the modeling process prior to flight.

- At this stage, one is ready to make preflight predictions and/or develop databases for flight. The epistemic uncertainty increments associated with the model-form are now fixed, so the CFD model is simply applied to the flight environment, which has its own set of epistemic and aleatoric uncertainty sources. These uncertainties are fed into the UQ framework when performing the flight simulations, so that the product delivered by the CFD practitioner includes a "best estimate" for the predicted quantities of interest along with uncertainty estimates for these predictions.

SUMMARY

In summary, a nonintrusive uncertainty quantification framework has been described and applied to the validation process of a turbulence closure model used in a highly back-pressured scramjet isolator flowfield. The uncertainty framework chosen for this exercise properly handles both aleatoric and epistemic uncertainty sources, providing the ability to perform a second-order uncertainty analysis. The framework encapsulates the uncertainty information in the form of a probability box involving an ensemble of cumulative distribution functions. Each cumulative distribution function captures the aleatoric uncertainty for a given sample of the epistemic uncertainty space. An enabling element in the application of this uncertainty quantification framework was the development of a metamodel for the propagation of uncertainties. This resulted in a major cost savings with regard to the computational time required to perform the uncertainty analysis. A total of only 14 high-fidelity CFD simulations were performed in this effort to build an adequate metamodel for the uncertainty propagation. Another unique element in the application of this framework was the treatment of discretization error. This error was treated as a categorical epistemic uncertainty source between two distinct grid resolution levels, and eventually transformed to a bounding estimate of spatial discretization error using the grid convergence index. A recipe for how this framework could be applied to more complex scenarios, in particular ground-to-flight mapping, was also provided.

ACKNOWLEDGMENTS

This effort was funded through the NASA Hypersonics Technology Project and carried out at the Hypersonic Airbreathing Propulsion Branch at the NASA Langley Research Center. Computational resources for this work were provided by the NASA Langley Research Center and the NASA Advanced Supercomputing (NAS) Division.

REFERENCES

[1] Baurle, R. A. and Eklund, D. R., ***Analysis of Dual-Mode Hydrocarbon Scramjet Operation at Mach 4–6.5***, Journal of Propulsion and Power, 18(5):990–1002 (Sept.-Oct. 2002).

[2] Hagenmaier, M. A., Gruber, M. R., Hsu, K.-Y., and Eklund, D. R., ***A Study of Shock Trains in Constant Area and Divergent Rectangular Ducts***, 2002 JANNAF CS/APS/PSHS/MSS Joint Meeting (April 2002).

[3] Roy, C. J. and Oberkampf, W. L., ***A Complete Framework for Verification, Validation, and Uncertainty Quantification in Scientific Computing***, AIAA Paper 2010-0124 (Jan. 2010).

[4] The Predictive Science Academic Alliance Program (PSAAP), ***http://www.lanl.gov/asci/university-partnerships.php*** (2017).

[5] Enabling Quantification of Uncertainty in Physical Systems (EQUiPS), ***https://www.darpa.mil/program/equips*** (2017).

[6] Mehta, U., Romero, V., Eklund, D., Pearce, J. and Keim, N., ***The JANNAF Simulation Credibility Guide on Verification Uncertainty Propagation and Quantification, and Validation***, AIAA Paper 2015-0502 (Jan. 2015).

[7] West, T. K., Hosder, S., and Johnston, C. O., ***Multistep Uncertainty Quantification Approach Applied to Hypersonic Reentry Flows***, Journal of Spacecraft and Rockets, 51(1):296–310 (2014).

[8] West, T. K., Reuter, B. W., Walker, E. L., Kleb, B., and Park, M. A., ***Uncertainty Quantification and Certification Prediction of Low Boom Supersonic Aircraft Configurations***, Journal of Aircraft, 54(1):40–53 (2017).

[9] Kleb, B. and Johnston, C. O., ***Uncertainty Analysis of Air Radiation for Lunar Return Shock Layers***, AIAA Paper 2008-6388 (Aug. 2008).

[10] Middleton, T. F., Balla, R. J., Baurle, R. A., Humphreys, W. M., and Wilson, L. G., ***The NASA Langley Isolator Dynamics Research Lab***, 2009 JANNAF CS/APS/PSHS Joint Meeting (Dec. 2009).

[11] VULCAN, ***http://vulcan-cfd.larc.nasa.gov/*** (2017).

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

[13] Koren, B., ***A Robust Upwind Discretization Method for Advection, Diffusion and Source Terms***, in C. B. Vreugdenhil and B. Koren, editors, *Numerical Methods for Advection-Diffusion Problems*, pages 117–138, Braunschweig: Vieweg Verlag (1993).

[14] McBride, B. J. and Gordon, S., ***Computer Program for Calculation of Complex Chemical Equilibrium Composition and Applications, I. Analysis***, NASA Reference Publication 1311 (Oct. 1994).

[15] 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 (June 1996).

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

[17] Lekivetz, R. and Jones, B., ***Fast Flexible Space-Filling Designs for Nonrectangular Regions***, Quality and Reliability Engineering International, 31(5):829–837 (July 2015).

[18] Meyers, R. H. and Montgomery, D. C., ***Response Surface Methodology: Process and Product Optimization Using Designed Experiments***, Wiley, 2nd edn. (2002).

[19] Baurle, R. A., Middleton, T. F., and Wilson, L. G., ***Reynolds-Averaged Turbulence Model Assessment for a Highly Back-Pressured Isolator Flowfield***, 2012 JANNAF CS/APS/PSHS Joint Meeting (Dec. 2012).

[20] Shapiro, S. S. and Wilk M. B., ***An Analysis of Variance Test for Normality (Complete Samples)***, Biometrika, 52(3/4):591–611 (Dec. 1965).

Appendix A. METAMODEL ERROR PREDICTION AND ADEQUACY OF FIT

For the purposes of including the prediction error of the metamodel itself in the uncertainty propagation process, the error was quantified mathematically assuming the underlying assumptions of the least squares methodology are not violated. The first step is to obtain the metamodel variance throughout the design space. This can be calculated after determining both the sensitivity of the model coefficients at a given point in the design space and the metamodel fit variance-covariance matrix, the latter of which is constant throughout the design space. Equation 1 demonstrates this calculation using the delta method, *i.e.*,

$$\sigma^2(\mathbf{X}) \;=\; \left[\frac{\partial \hat{R}\left(\hat{\mathbf{b}},\mathbf{X}\right)}{\partial \hat{b}}\right] \widehat{Cov}\left(\hat{\mathbf{b}}\right) \left[\frac{\partial \hat{R}\left(\hat{\mathbf{b}},\mathbf{X}\right)}{\partial \hat{b}}\right]^T \tag{1}$$

Once the metamodel variance is obtained, it can be transformed into a standard error for computing a confidence interval given the number of points (N) used to construct the metamodel, *i.e.*,

$$\epsilon_{std} \;=\; \frac{\sigma}{\sqrt{N}} \tag{2}$$

The metamodel prediction interval bounds can then be calculated throughout the design space by multiplying the standard error by the Student's t value and the square root of the number of samples used to build the metamodel. The resulting expression for the metamodel prediction error is then given by

$$\epsilon \;=\; \pm\, t\, \epsilon_{std} \sqrt{N} \tag{3}$$

A 95% prediction interval approaches a t value of 1.96 as the number of samples approaches infinity. The prediction interval was chosen over the confidence interval in this study in order to provide a larger error to capture the scatter of the data about the mean, rather than the uncertainty in the mean.

Metamodel adequacy checks allow the analyst to understand how suitable the metamodel is in representing the physics of the system as simulated by the source data. Quantitative and qualitative adequacy checks were used as indicators of the metamodel adequacy and include $R^2_{adjusted}$, $R^2_{predicted}$, actual versus predicted plots, residual versus predicted plots, and the distribution of studentized residuals. These quantities are defined as follows:

- $\mathbf{R^2_{adjusted}}$: The adjusted correlation coefficient is related to the correlation coefficient but also includes penalties to account for the number of modeling terms used, and will therefore always be less than the unadjusted correlation coefficient. Ideally it should be 0.99 or greater.
- $\mathbf{R^2_{predicted}}$: The predicted correlation coefficient is a cross-validation metric that removes one fit point from the metamodel at a time and refits the metamodel. It essentially validates the metamodel against points used to fit the model. Ideally it should be close to the adjusted correlation coefficient to provide confidence that the metamodel is relatively stable to removing points. This was the only cross-validation measure obtained for the metamodel, although additional CFD simulations could have been performed to augment this assessment.
- **Actual versus Predicted Plots**: The actual versus predicted plots should contain points that run closely to the 1:1 line with low dispersion about the perfect fit line and have no fanning or any other nonlinear trends.
- **Residual versus Predicted Plots**: The residual versus predicted plots should appear to have random scatter about the zero residual line with no observable trends that would indicate the metamodel is biased in any particular region of the design space. Ideally, the maximum residual magnitude should be less than 10% of the minimum response value.

- **Studentized Residuals**: The least squares methodology assumes that the error is normally distributed about the mean prediction. Plotting a histogram of studentized residuals of the metamodel relative to the data fit points is one way to validate this assumption. Ideally, the residuals should have a normal distribution with zero mean and a standard deviation of unity. The Shapiro-Wilk test[20] should be used to determine the probability that the data does not adhere to a normal distribution.

Metamodel adequacy checks are shown in Fig. A1 for the actual versus predicted and residual versus predicted metrics. The actual versus predicted trend stays close to the 1:1 line with no visible outliers. The residual versus predicted plot appears to have random scatter with no trends indicating model bias. The adjusted correlation coefficient is 0.99 (indicating good capture of the variability in the model), and the predicted correlation coefficient is 0.88 (which is a relatively high value, but suggests that the model is not fully converged). In these images, the 1's indicate sample points from the initial I-Optimal design, and the 2's indicate the second set of points used to augment the design. Finally, the studentized residual plot in Fig. A2 does not indicate the presence of outliers, and normality of the error is not rejected via the Shapiro-Wilk test.[20] This result provides sufficient confidence that the errors can be assumed to be normally distributed. The measure of adequacy provided by these metrics suggest that the model is suitable for the subsequent uncertainty propagation process.

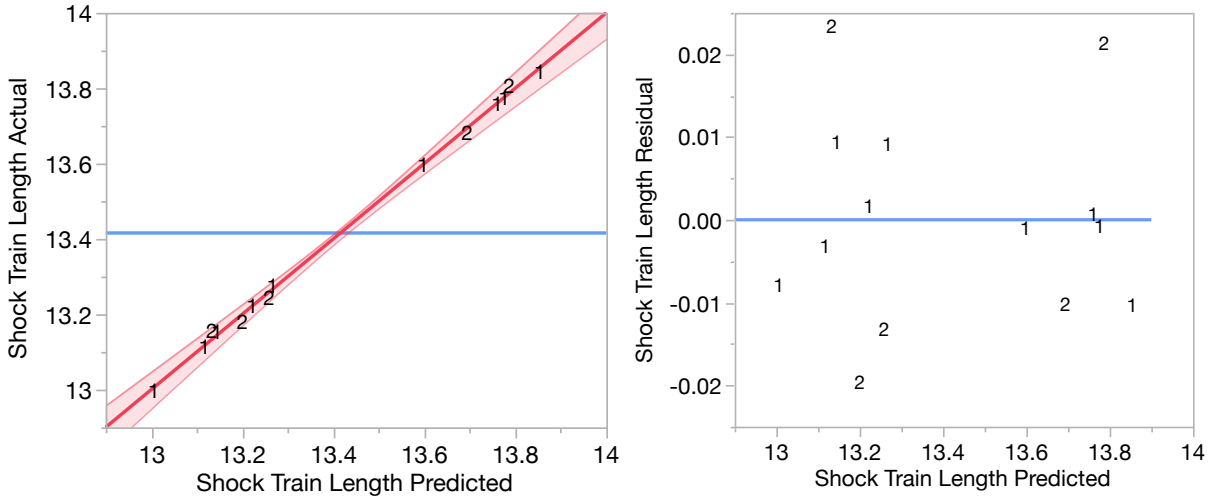

Figure A1: Actual versus predicted response with 95% metamodel significance confidence curves bounding the shaded region (left), Residual versus predicted response showing good, random dispersion in the results (right).

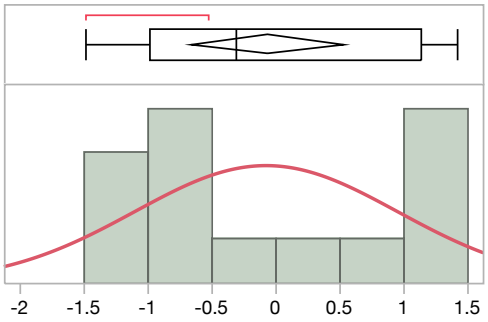

Figure A2: A histogram of the studentized residuals of the metamodel with normal distribution fit. An outlier box plot is also shown, including a confidence interval on the mean (diamond) and the most dense 50% of observations (red bracket).