

HIFiRE Flight 2 Unstart Reliability Analysis

Thomas K. West IV*

NASA Langley Research Center, Hampton, VA, 23681

Michael D. Bynum†

Analytical Mechanics Associates, Inc., Hampton, VA, 23681

Nomenclature

CR	Confidence Ratio	q_∞	Freestream Dynamic Pressure
D	Deterministic Variables	R	Stochastic Response
d_L	Left-Sided Validation Metric	Sc_t	Turbulent Schmidt Number
d_R	Right-Sided Validation Metric	U	Uncertainty
M	Margin	α	Polynomial Chaos Expansion Coefficient
N_t	Number of Terms in a Total-Order Polynomial Chaos Expansion	ϵ	Surrogate Model Error
n	Number of Random Dimensions	κ	von Karman's Constant
P	Pressure	Ψ	Random Basis Function
p	Polynomial Order	σ	Diffusion Coefficient
		ξ	Standard Input Random Variable

I. Introduction

The Hypersonic International Flight Research Experimentation (HIFiRE) program [1] is a collaboration between the Air Force Research Laboratory (AFRL) and the Australian Defence Science and Technology Organisation (DSTO). The goals of the HIFiRE program are to further the understanding of basic hypersonic phenomena and to enable research and exploration using flight experiments. Investments are targeted at phenomena and configurations that are difficult to model using CFD and challenging to test in existing ground test facilities. The primary objective of HIFiRE Flight 2 (HF2) was to explore mode transition and combustion of a hydrocarbon fueled scramjet engine between Mach 6 and 8.

The HF2 payload was flown on May 1, 2012, from the Pacific Missile Range Facility in Hawaii. During the flight, the three-stage sounding rocket successfully accelerated the payload from Mach 6 to Mach 8+ at a nearly constant dynamic pressure. The forebody/inlet shroud was deployed at Mach 5.2, then the fuel schedule was initiated and ignition of the combustor was achieved. As the vehicle continued to accelerate, the combustor transitioned from dual-mode to scram-mode operation at a nearly constant fuel–air equivalence ratio (near unity). Data were acquired from more than 600 sensors for the entire duration of the flight test up to and beyond the end of the experiment.[2]

Prior to the HF2 experiment, a full-scale direct-connect version of the HF2 flowpath, the HIFiRE Direct-Connect Rig (HDCR), was thoroughly characterized in NASA's Arc-Heated Scramjet Test Facility.[3] The ground test rig geometry was identical to the flight experiment from the isolator entrance to the combustor exit, including scale, with only minor differences in instrumentation. At the time of the ground test campaign, the minimum allowable isolator margin to be used in the flight experiment was set at 50% of the isolator length. This conservative isolator margin was assumed to provide ample margin for preventing an unstart condition during the flight experiment. Robust operation of the isolator and combustor was confirmed by repeated and thorough testing at and near the expected flight-fueling conditions. A full summary of the HDCR ground test campaign can be found in Refs. [4–6]

To quantify the isolator margin in the HF2 experiment, an end-to-end uncertainty quantification (UQ) process is required that captures deficiencies in both modeling and the measurements used to inform said modeling obtained from HDCR testing. Storch and Bynum [7] conducted such an analysis that resulted in a ground-to-flight prediction of isolator margin. Though this previous work did not rigorously establish uncertainty impacts, it did establish that the

* Aerospace Engineer, Vehicle Analysis Branch, Systems Analysis and Concepts Directorate, Member AIAA.

† Aerospace Engineer.

flight risk associated with combustor-induced unstart were low, a conclusion substantiated by the HF2-flight test. Note that this study seeks to expand upon this prior analysis, by performing a more systematic quantification of uncertainty impacts on isolator-margin metrics.

The objective of this work was to revisit the HF2 isolator unstart analysis and conduct a comprehensive reliability assessment of the HF2 experiment. The goal was to roll the uncertainty from the ground testing and subsequent model validation exercises up to a prediction of combustion-induced unstart reliability/isolator margin at the Mach 6.5 condition that is inclusive of known flight uncertainties. To conduct the model validation and compute the model-form uncertainty of the computational fluid dynamics (CFD) model, the uncertainty in the HDCR measurements and the CFD were both quantified. The determined model-form uncertainty was then fed forward as an additional source of uncertainty in the CFD model of the flight vehicle.

The overall framework for assessing the reliability of the HF2 experiment was based on the quantification of margins and uncertainties (QMU) [8]. This framework offers a flexible and practical approach to determining confidence of a complex system, like that of an air-breathing propulsion vehicle. Forward propagation of input uncertainties was conducted using non-intrusive polynomial chaos for efficiency due to the computational cost of the CFD models. Model-form uncertainty of the CFD model was computed using the area metric approach described by Brune et al. [9]. All of these components rolled into an end-to-end UQ and reliability framework constitute a unique and widely applicable approach to determining system level performance that includes uncertainty.

The next section describes the computational modeling of the HF2 and HDCR geometries. Section III details the UQ and model validation approaches used in this study followed by Section IV, which describes the framework used for assessing reliability. Section V contains the results and discussion of the current study. Lastly, Section VI summarizes key findings of this work.

II. Description of the CFD Model, Grids and Spatial Errors

In this section, the HF2 geometry and computational modeling approach are presented. The combustion-induced shock location, which is the quantity of interest for measuring isolator unstart reliability, is then described. In particular, the approaches for obtaining this quantity from both the HDCR measurements and the CFD are discussed in detail. Furthermore, the unstart threshold is also discussed, which is later used as the operating limit in the QMU framework. Lastly, solution verification errors in the CFD model are determined, which will be included in the validation and reliability assessments.

A. Flowpath Geometry

The flowpath geometry used for the CFD simulations can be divided into three main sections: the forebody and inlet, the isolator, and the combustor. The forebody and inlet geometry [10] was only used in the HF2 flight test and corresponding CFD results. Figure 1 shows a three-dimensional view of the HF2 flowpath integrated into the flight test article. Inflow for the HDCR ground test came from the direct-connect facility nozzle, which is connected to the isolator. The forebody and inlet were not part of the direct connect test configuration.




Fig. 1 HIFiRE Flight 2 forebody and inlet geometry integrated with the flight vehicle.

Figure 2 shows a two-dimensional schematic of the isolator and combustor flowpath. The coordinate frame origin for the flight test article is located in the center of the flowpath and the forebody virtual leading edge. The isolator entrance is located 40.438 in. downstream of the forebody virtual leading edge. The flowpath is symmetric about the $z = 0$ plane.

0 and $y = 0$ planes. Port and starboard isolator and combustor surfaces are vertical planes separated at a constant width of 4 in. Table 1 lists the x and y values for several key geometric features on the cowl wall.




Fig. 2 Two-dimensional profile of the HIFiRE Flight 2 isolator and combustor geometry.

Table 1 HIFiRE Flight 2 cowl wall flowpath geometry specifications.

Feature	x [in]	y [in]
Isolator Start	40.438	0.5
Isolator End	48.438	0.5
Combustor Start	48.438	0.5
P1 Injectors	50.034	0.536
Cavity Step Start	52.034	0.582
Cavity Step End	52.019	1.256
Cavity Closeout Start	54.588	1.315
Cavity Closeout End	56.232	0.677
S1 Injectors	56.938	0.693
Combustor End	68.441	0.954

The injectors used in both the ground and flight experiments were the P1 and S1 injector stations (Fig.2). The P1 injector station, located upstream of the cavity, consists of four injectors equally spaced apart at 1 in. and at 1/2 in. from the side walls with diameters of 1/8 in. The S1 injector station, located downstream of the cavity, consists of four injectors equally spaced apart at 1 in. and at 1/2 in. from the side walls with diameters of 0.094 in. P1 injectors are inclined at 15° to the flowpath surface, and S1 injectors are normal to the flowpath surface.

B. Computational Modeling

The computational modeling for this effort was performed using VULCAN-CFD 7.1 [11]. VULCAN-CFD is a compressible, turbulent, Navier-Stokes flow solver with allowances for finite-rate chemistry and/or thermodynamic non-equilibrium. Multi-block structured or mixed element (tetrahedral, prismatic, pyramid, and hexahedral) unstructured grids can be used for discretization of the computational domain, and solutions can be obtained using explicit or implicit time advancement schemes. Wall matching functions and multigrid methods for elliptic and space marching schemes are available to reduce computational cost. Flow solver capabilities include models for compressibility, Reynolds stress anisotropies, turbulent diffusivity, finite-rate chemistry, and turbulence-chemistry interaction effects. Available turbulence models include several one- and two-equation models, as well as Large Eddy Simulation (LES) and hybrid Reynolds-Averaged Navier-Stokes (RANS)/LES options.

For the current work, computational analysis implemented thermally perfect and steady state assumptions. The turbulence model used was the blended 1988 Wilcox $k - \omega$ and the Jones Launder $k - \epsilon$ model of Menter [12] in

conjunction with the wall matching functions of Wilcox [13]. Wall matching functions reduce the overall number of grid points required and reduce the aspect ratio of the cells near the walls. Inviscid fluxes were evaluated using the low-diffusion flux splitting scheme of Edwards [14]. To ensure the solution satisfied the total variation diminishing (TVD) property, the van Leer flux limiter [15] was used. The monotone upstream centered scheme for conservation laws (MUSCL) [15] interpolation coefficient, κ , was chosen to be one-third.

The wall boundary condition used was a no-slip wall with a one-dimensional, wall-normal heat transfer model. This model requires a wall thickness over conductivity specification and a back side wall temperature. Thickness to conductivity values matched the vehicle wall hardware. Back side wall temperatures for the isolator and combustor were set to values obtained in the ground and flight test, respectively. The turbulent Prandtl number was set at 0.90.

The base combustion model used was the eddy dissipation concept (EDC) of Magnussen and Hjertager [16]. This model allowed for significantly reduced computational requirements and is valid for mixing-limited combustor operation such as in the current work. The base Magnussen model was modified according to the methods prescribed by Norris [17] to obtain a modified Magnussen model. This modified Magnussen model attempts to incorporate the predictive capabilities of a finite-rate kinetics model, but with the computational cost of the Magnussen model. The chemical composition of the fuel used was a 64%/36% molar mixture of ethylene and methane.

A three-dimensional, structured, quarter-symmetry computational grid was created for the HF2 and HDCR as-designed geometries. A block-wise representation of the grids created is shown in Fig.3. The HDCR and HF2 grids shared the common isolator and combustor blocks and only differed by the upstream components. Each grid contained 8.7 million cells and 19.1 million cells for HDCR and HF2, respectively. The grids were clustered near the wall to maintain an average y^+ of 10 throughout the flowpath with a maximum y^+ of 50. Simulation convergence was determined based on integrated quantities of interest and residual reduction. Convergence was declared when the integrated forces on the surfaces in the x, y, and z directions, the integrated wall heat flux, and the mass flow conservation were nominally invariant with iteration history (unsteady variations of no more than 0.5 percent); and when the L2-norm of the residual error had dropped by at least 2.5 orders of magnitude from the initial residual.




Fig. 3 Block-wise representation of the grids used for the HF2 and HDCR CFD computations.

C. Combustion Induced Shock Location

The methodology used to determine the location of the combustion-induced shock in the isolator for the ground test data differed between the original ground test campaign and the current work. In the original HDCR ground test campaign, the combustion-induced shock location was determined by the location of the first disturbed pressure tap compared to the tare data [7]. An example of this methodology can be seen in Fig.4(a). The methodology chosen to determine the combustion-induced shock location for the current work utilized a method presented by Hutzet et al. [18]. This approach can be described as the location at which the combustion data is greater than 1.5 times the tare data. More specifically the combustion-induced shock is said to be at the intersection of the two linearly interpolated lines formed by the combustion data and 1.5 times the tare data. An example of this methodology can be seen in Fig.4(b).




Fig. 4 HDCR combustion-induced shock location.

Both of the previously mentioned methodologies for determining the combustion-induced shock location rely on the presence of tare data for each combusting test point. This requirement is often easy to meet for ground test data; however, it is rarely true of CFD results as uncertainty may necessitate a tare solution be generated for every combusting solution. Creating a CFD tare solution for every combusting case analyzed would constitute a large expense compared to the traditional approach of having a single or significantly fewer tare solutions for the various combusting solutions. In the absence of a tare solution for each combusting solution, a methodology to determine the shock-induced pressure rise must rely solely on the combusting data itself. Also, because the spacial resolution of the CFD data is orders of magnitude higher than the test data, more continuous methods of determining the shock location can be used.

The method chosen for determining the combustion-induced shock location for the CFD data was based on the slope of the centerline wall pressure. The shock location was chosen to be the location where the axial pressure gradient, dP/dx , on the wall was greater than 37.9 psi/in. An example of this methodology can be seen in Fig.5. The magnitude of the dP/dx threshold was determined for this particular problem by finding the maximum dP/dx of the baseline tare CFD solution in the isolator, which, for the flight geometry, has oblique shocks caused by the inlet, and then providing ample margin to avoid false positives for the location of the combustion-induced shock.

D. Unstart Limit

To determine the unstart limit, the current work attempts to quantify the most-upstream possible, stable shock location in the isolator. The Langley Isolator Model (LIM) [19] was used to determine what upstream portion of the isolator was not able to hold a combustion-induced pressure rise. The LIM provides a formulation of the centerline pressure derivative with the streamwise direction in the isolator.

This formulation can be used to estimate the maximum back pressure possible for a given shock location in the isolator. It can also be used to predict the locations in the isolator where the boundary layer is too weak to support a given back pressure causing the shock to move discontinuously across this weak boundary layer region. A typical example of this type of analysis can be seen in Fig.6. The centerline skin friction coefficient, which is a measure of boundary layer strength, is plotted versus a portion of the isolator length. Local minima in the centerline skin friction coefficient are caused by the oblique shock wave pattern in the isolator. These regions of low skin friction correspond to the locations where the LIM predicts discontinuous movement in the shock location. The most upstream local minimum of the skin friction coefficient is the unusable portion of the isolator for this particular example, with approximately 0.6 inches of the isolator being unusable. This analysis must be performed for different inflow conditions and modeling approach as the shock pattern in the isolator may change.




Fig. 5 Current CFD combustion-induced shock location methodology.




Fig. 6 Minimum predicted stable shock location in the isolator using the Langley Isolator Model.

E. Grid Convergence and Spatial Resolution Error

An estimate of the spatial discretization error present in the computations was formulated using the Grid Convergence Index (GCI) method of Roache [20] with the combustion-induced shock location being the metric of interest for the GCI study. The first-order formulation of the GCI method was used with a safety factor of 1.25. This formulation provides additional margin as compared to the second-order formulation of the GCI. A factor of two grid coarsening was performed in each grid spatial direction for the HDCR and HF2 grids. The results of this coarser grid and the baseline grid were used to obtain the spatial error estimation.

Because only two grid levels were used for the spatial error estimation, the spatial error change with increasing grid resolution could not be determined to be monotonic for the baseline grid. As a result, instead of assigning the spatial error estimation as a one-sided error estimate, the spatial error was taken as a two-sided error estimate. The GCI of the combustion-induced shock location for the HF2 grid was determined to be ± 0.051 inches. The GCI of the combustion-induced shock location for the HDCR grid was determined to be ± 0.049 inches.

III. Uncertainty Quantification and Model Validation

To make accurate predictions about flight performance, there is a series of steps that must be completed to incorporate uncertainty into the predictions. First, sources of uncertainty in the computational model must be identified and propagated through the model. Then, verification errors are appended to the model outputs, such as spatial grid resolution error. Lastly, the model-form uncertainty, determined through a validation process, is appended to the model outputs. Note that model calibration can be used to reduce this model-form uncertainty, but it may not be eliminated for complex models. All together, the input uncertainty, verification errors, and model-form uncertainty are combined to provide an estimation of the model output variation. This process is illustrated in Fig.7.




Fig. 7 Computational model uncertainty quantification.

Each of the steps above and the approaches used to efficiently and accurately execute them are described in the remainder of this section. Types of uncertainty in computational models are first described, followed by a brief discussion of how uncertainty is propagated through a model. Next, the polynomial chaos approach is outlined, which is used as a means to efficiently propagate uncertainty through a computationally expensive model. Then, the model verification, validation, and calibration approaches used in the present work are described.

A. Types of Uncertainty in Numerical Modeling

A critical step in any uncertainty analysis is the identification and classification of uncertain parameters. These parameters may be mathematically represented differently based on the nature of their uncertainty. Incorrect classification and/or treatment of uncertain parameters can result in widely varying output uncertainty.

Two main types of uncertainty exist in numerical modeling: aleatory uncertainty and epistemic uncertainty [21]. Aleatory uncertainty is the inherent variation of a physical system. Such variation is due to the random nature of input data and can be mathematically represented by a probability density function if substantial experimental data is available for estimating the distribution. An example of this type of uncertainty could be the fluctuations in freestream quantities. These random variations are not controllable and are sometimes referred to as irreducible uncertainties.

Epistemic uncertainty in a stochastic problem comes from several potential sources. These include a lack of knowledge or incomplete information about the behavior of a particular variable. Also, ignorance or negligence with regards to accurate treatment of model parameters is a source of epistemic uncertainty. In the context of a mathematical model, assumptions and inadequate knowledge of a physical process leads to what is termed model-form uncertainty. Contrary to aleatory uncertainty, epistemic uncertainty is sometimes referred to as reducible uncertainty. An increase in knowledge regarding the physics of a problem, along with accurate modeling, may reduce this type of uncertainty, but typically comes at some cost. Epistemic uncertainty is often modeled using intervals because the use of probabilistic distributions (even a uniform distribution) can lead to inaccurate predictions in the amount of uncertainty in a system. Upper and lower bounds of these intervals can be drawn from limited experimental data or from expert predictions and judgment.

An additional, special case of epistemic uncertainty is computational error. This uncertainty is common in numerical modeling and is defined as a recognizable deficiency in any phase or activity of modeling and simulations that is not due to lack of knowledge of the physical system. In CFD, an example of this type of uncertainty would be the discretization error in both the temporal and spatial domains that comes from the numerical solution of the partial differential equations that govern the system. This uncertainty can be well understood and controlled through code verification and grid convergence studies.

B. Uncertainty Propagation

Many stochastic problems may contain both epistemic and aleatory uncertainty. The desired approach is to consider the contribution of both types of uncertainty simultaneously by propagating this mixed uncertainty through the stochastic model. This can be done using a procedure known as second-order probability (SOP)[22]. SOP is a type of double loop sampling, shown in Fig.8. In the outer loop, a vector of values for the epistemic variables is passed into the inner loop where the stochastic model is sampled for the single epistemic vector and every aleatory vector. Following the repeated process for all epistemic vectors, the total number of samples is the product of the number of epistemic samples and the number of aleatory samples.




Fig. 8 Schematic of second-order probability.

Each iteration of the outer loop generates a cumulative distribution function (CDF) based on the aleatory uncertainty analysis in the inner loop. After completion of the process, what remains is a family of CDFs, which gives intervals of the output variable from the model at different probability levels (i.e., a probability box or “p-box” representation of mixed uncertainty output). Important information can be taken from p-boxes, including confidence intervals of desired moments. For the case of mixed uncertainty, one approach to obtaining the 95% confidence interval, for example, is to take the upper 97.5% probability level and the lower 2.5% probability level as the interval. The p-box and confidence interval measurement are illustrated in Fig.9. Note that the boundaries of the p-box typically hold the greatest interest during analysis and design, as will be shown in later sections.




Fig. 9 P-box representation of mixed uncertainty output and confidence interval measurement.

C. Point-Collocation, Non-Intrusive Polynomial Chaos

Polynomial chaos is a surrogate modeling technique based on a spectral representation of the uncertainty. An important aspect of spectral representations is the decomposition of a response value or random function R into a linear combination of separable deterministic and stochastic components, as shown in Eq. (1) for a polynomial of order p and

dimension n .

$$R(\mathbf{D}, \boldsymbol{\xi}) \approx \sum_{i=0}^{N_t-1} \alpha_i(\mathbf{D}) \Psi_i(\boldsymbol{\xi}) + \epsilon \quad (1)$$

Here, α_i is the deterministic component and Ψ_i is the random variable basis functions corresponding to the i^{th} mode. The basis functions, Ψ_i , of each random variable are determined using the Askey key [23] and are dependent on the distribution of each random variable. The response, R , is a function of independent, deterministic variables, $\mathbf{D}$, and n independent, standard random variables, $\boldsymbol{\xi}$. Note that this series is, by definition, an infinite series; however, in practice, it is truncated, and a sum is taken over a discrete number of output modes, leaving a finite truncation error, ϵ . To form a complete basis or a total order expansion, N_t terms are required, which can be computed from Eq. (2) for a polynomial chaos expansion (PCE) of order p and a number of random dimensions or variables, n .

$$N_t = \frac{(n+p)!}{n!p!} \quad (2)$$

Special consideration is needed when epistemic variables are present. For epistemic variables with a bounded domain (i.e., intervals), the Legendre polynomials can serve as an appropriate basis, as they too have a bounded domain. Then, the second-order probability approach described in the previous section can be applied to propagate mixed uncertainty through the PCE and obtain intervals at select probability levels of interest. Further details on polynomial chaos theory is given by Ghanem and Spanos [24]. A description of more recent developments, including the analysis of mixed uncertainty with stochastic expansions is given by Eldred [25].

The objective with any PCE method is to determine the expansion coefficients, α_i . To do this, polynomial chaos methods can be implemented using an intrusive or a non-intrusive approach. While an intrusive method may appear straightforward in theory, for complex problems this process may be computationally expensive and difficult to implement as changing the deterministic model is required [26]. In contrast, the non-intrusive approach can be easily implemented to construct a surrogate model that represents a complex computational simulation because no modification to the deterministic model is required. The non-intrusive methods require only the response (or sensitivity [27–29]) values at selected sample points to approximate the stochastic response surface.

Several methods have been developed for non-intrusive polynomial chaos (NIPC). Of these, the point-collocation NIPC method has been used extensively in many aerospace simulation and CFD problems [27, 30–32] for improved computational efficiency and tractability for high-dimension problems over other spectral-projection-based approaches. The point-collocation method starts with replacing a stochastic response or random function with its PCE by using Eq. (1). Then, N_t sample vectors are chosen in random space, and the deterministic code is evaluated at these points, which is the left hand side of Eq. (1). Following this, a linear system of N_t equations can be formulated and solved for the expansion coefficients of the PCE. This system is shown in Eq. (3).

$$\begin{pmatrix} R(\mathbf{D}, \boldsymbol{\xi}_1) \\ R(\mathbf{D}, \boldsymbol{\xi}_2) \\ \vdots \\ R(\mathbf{D}, \boldsymbol{\xi}_{N_s}) \end{pmatrix} = \begin{pmatrix} \Psi_0(\boldsymbol{\xi}_1) & \Psi_1(\boldsymbol{\xi}_1) & \cdots & \Psi_{N_t-1}(\boldsymbol{\xi}_1) \\ \Psi_0(\boldsymbol{\xi}_2) & \Psi_1(\boldsymbol{\xi}_2) & \cdots & \Psi_{N_t-1}(\boldsymbol{\xi}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_0(\boldsymbol{\xi}_{N_s}) & \Psi_1(\boldsymbol{\xi}_{N_s}) & \cdots & \Psi_{N_t-1}(\boldsymbol{\xi}_{N_s}) \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_{N_t-1} \end{pmatrix} \quad (3)$$

Note that for this linear system, N_t is equal to the minimum number of deterministic samples, N_s , required to obtain a direct solution. If more samples are available and are linearly independent from other samples, the system is considered overdetermined and can be solved using a least squares approach. The number of samples over the required minimum is represented by the use of an oversampling ratio (OSR), defined as the ratio of number of actual samples to the minimum number required (i.e., N_t). In general, the number of collocation points can be determined by multiplying Eq. (2) by an OSR. Hosder et al. [33] determined that the PCE is dependent on the number of collocation points and an effective OSR of two was sufficient for the stochastic model problems studied. In the present work, collocation points are chosen using a combination of a D-optimal design and Latin Hypercube sampling for pseudo random space filling. This ensures coverage of the distributions for each random variable. More advanced sampling approaches, such as adaptive sampling, are the subject of future work.

D. Surrogate Model Verification

When using a surrogate model, such as PCE, to propagate input uncertainty through a model, verification errors arise from two sources: the training data and the model fit. In the present study, the verification error in the training data

is the result of solution verification errors. Spatial grid convergence error estimation was discussed in the previous section and represents the total solution verification error. Note that iterative convergence errors were determined to be negligible in this study.

For the surrogate model, verification is a necessary step to ensure the PCE is an accurate representation of the deterministic model on the random input domains. Before training any PCE, decisions must be made regarding polynomial basis order and, for the point-collocation approach, the number of collocation points used to train the PCE. Many approaches exist for verifying the accuracy of PCEs. According to the theorem by Cameron and Martin [34], a PCE expansion should converge in the L_2 sense. Slower convergence may be an indication of a highly nonlinear response on the input domain. For computationally expensive problems, however, raising the polynomial order high enough to observe this level of convergence may be prohibitive. One alternative approach is to use additional test points, not used in the development of the model. Another is to use cross validation by removing one or more points from the training set to determine the dependency on specific collocation points. These verification tests address the choice of PCE order and can be used to quantify truncation and regression error of the PCE. In the present study, Leave-One-Out cross validation is used to determine the verification error in the PCE.

E. Model Validation and Total Model Uncertainty

One approach to quantifying the discrepancy between a model prediction and experimental data is the area metric proposed by Ferson et al. [35]. In the present study, a modified version proposed by Brune et al. [9] is used to correctly account for mixed uncertainty. This metric is given by Eq. (4).

$$d_L, d_R = \int_{-\infty}^{\infty} g_L(x) dx, \int_{-\infty}^{\infty} g_R(x) dx \quad (4)$$

where,

$$g_L(x) = \begin{cases} 0 & \text{for } x \text{ such that } F_L(x) \geq S_L(x) \\ S_L(x) - F_L(x) & \text{for } x \text{ such that } F_L(x) < S_L(x) \end{cases}$$

and

$$g_R(x) = \begin{cases} 0 & \text{for } x \text{ such that } F_R(x) \leq S_R(x) \\ F_R(x) - S_R(x) & \text{for } x \text{ such that } F_R(x) > S_R(x) \end{cases}$$

Here, S is the experimental value and F is the model prediction. The resulting metrics d_L and d_R quantify the evidence for disagreement between both the left and right bounding probability distributions, while preserving the range of possible distributions from both the prediction and measurements. Note that in this study, there is no assumption of symmetry with regards to validation errors. The disagreement between a model and measurements may be one-sided, and Eq. (4) accurately captures this scenario.

The total model prediction uncertainty is generated by first propagating input uncertainty through the PCE model using the SOP approach. This generates a p-box to which the total verification error and validation metrics from Eq. (4) are appended. An example of the total model prediction uncertainty is illustrated in Fig.10.

IV. Reliability Assessment using Quantification of Margins and Uncertainties

Traditional reliability (or conversely, probability of failure) is based on estimating the probability a system or component will fail given uncertain operating conditions and modeling assumptions. As will be shown later, the HF2 experiment was predicted to have a 0% chance of failure (or 100% reliability) given known sources of uncertainty, in this traditional sense. However, in general, hypersonic systems are highly sensitive to any number of changes or variations, many of which are either not being modeled or are not well understood. That is to say, there is an unknown amount of epistemic uncertainty present in the flight predictions.

Instead of artificially appending unsubstantiated factors of safety or knockdown factors to predictions, an alternative approach to assessing reliability is based on the quantification of margins and uncertainties (QMU) framework [8, 36]. For a select probability level, the QMU approach communicates the relationship between a quantity of interest, uncertainty in that quantity, uncertainty in the performance limit, and the margin between them. Specific to inlet-isolator operability of a dual-mode ramjet system, the key measure of reliability is the amount of remaining usable isolator. As backpressure increases, the combustion-induced shock location advances forward, towards the inlet, through the isolator. If the backpressure is too large, the combustion-induced shock advances beyond the isolator causing the inlet to unstart.




Fig. 10 Total model prediction uncertainty build-up.

Reliability is then a function of usable isolator remaining between the shock location and the last stable position near the inlet throat. This is illustrated in Fig.11 for an inlet-isolator system.




Fig. 11 Schematic of inlet-isolator operability.

Notice here that both the combustion-induced shock and last stable shock locations are not deterministic points in this figure. Instead, these locations are ranges of possibilities, which are a function of both the uncertainty in predictability with modeling and simulations, as well as measurement capability during ground testing. One way of communicating the reliability of a system with uncertainty in the quantity of interest (e.g., shock location) and a performance limit (e.g., last stable shock location) is through a confidence ratio, CR , given by Eq. (5).

$$CR = \frac{M}{U} \quad (5)$$

Here, U is an aggregate of the uncertainty in both the quantity of interest and performance limit, and M is the margin between the two, which is the operating margin. While CR is not a direct measure of reliability, it does give an indication as to how uncertain a systems performance may be versus the amount of margin built into the system. For example, in the case of the inlet-isolator system, if the separation location prediction has very large uncertainty, yet maintains positive margin, a $CR < 1$ may indicate a low level of confidence, even if no possibility of failure is observed after quantifying all known sources of uncertainty. On the contrary, a $CR > 1$ may suggest that there is excess margin relative to the amount of uncertainty. The CR may be particularly insightful for complex systems with a known lack of understanding (i.e., epistemic uncertainty) in modeling and testing. Note that, as described by West et al. [36], the margin and uncertainty are determined at a select probability level, thereby making CR specific to a desired level of uncertainty.

Dual-mode ramjets pose a particular challenge when determining operating margin. As mentioned previously, the quantity of interest when assessing the reliability of a dual-mode ramjet system is the combustion-induced separation

location within the isolator. The operating margin, M , is then a measurement between the performance prediction and operating limit with uncertainty included. That is to say a measure between the combustion-induced shock location and the last stable shock location. Historically, this margin was a direct length between these two spatial locations in the isolator.

Note that the combustion-induced shock location in an isolator containing oblique shock waves caused by the flow being processed through an inlet system often moves discontinuously through an isolator. This occurs even if the back pressure in the combustor is varied in a continuous manner. The primary reason for the discontinuous movement of the combustion-induced shock location is due to regions of weakened boundary layer caused by shock-boundary layer interactions in the isolator. A result of this discontinuous movement of the combustion-induced shock location is that the linear distance from a combustion-induced shock location in the isolator to the isolator entrance may not be an adequate representation of the available operating margin. In other words, a portion of the isolator is unusable, as illustrated in Fig.11. A more accurate measure of margin would exclude regions of the isolator where weakened boundary layer cannot support the combustion-induced shock.

However, the HF2 is a unique case. As a result of the inlet design [37], the oblique shocks within the isolator are weak enough not to render the boundary layer completely unable to support a combustion-induced shock. This is evident by the LIM output in Fig.12 showing the backpressure as a function of shock location at nominal flight conditions. Notice that there are no discontinuous regions where the combustion-induced shock could exist at two locations in the isolator for a given backpressure. Only a small of portion of the isolator near 43 in. may be unusable. While not a significant issue for the HF2 system, future analyses may require an alternative approach to margin measurement to account for unusable portions of the isolator.




Fig. 12 Predicted separation location as a function of combustor backpressure.

V. Results and Discussion

This section presents a reliability assessment of the HIFiRE Flight 2. First, the sources of uncertainty considered in this study are described. Then, the validation approach is applied to a CFD model of the HDCR experiments. Comparison of these two enable a determination of model-form uncertainty that is then carried forward to flight reliability predictions. Finally, an analysis of the operating limit is conducted, followed by an application of the QMU framework to quantify the HF2 system confidence ratio.

A. Sources of Uncertainty

In the present study, there were a total of eight uncertain parameters in the computational model identified for the HIFiRE Flight 2 experiment. These parameters are listed in Table 2. Epistemic uncertainty sources were chosen to

be distribution free on intervals. Additionally, all of the subsequent results include verification error from the CFD grid convergence study (Section II.E) and any errors associated with PCE model using a leave-one-out cross-validation approach (Section III.D). These errors are appended to the input uncertainty, as described in Section III.E.

Table 2 HIFiRE Flight 2 uncertainty sources before calibration to HDCR data.

Parameter	Distribution	Mean/ Min	Std. Dev./ Max
Mach	Gaussian	6.5	0.1
q_∞ (psf)	Gaussian	2000.0	40.0
$\sigma_{\omega 1}$	Epistemic (interval)	0.3	0.7
κ	Epistemic (interval)	0.38	0.42
Sc_t	Epistemic (interval)	0.5	0.9
EDC_t (1/s)	Epistemic (interval)	30000.0	300000.0
ϕ factor	Epistemic (interval)	0.9	1.1
EDC_1	Epistemic (interval)	1.5	4.0

Uncertainty in the freestream Mach number and dynamic pressure, q_∞ , are representative of potential initial state variation during flight of the experiment. Turbulence model parameter uncertainties were taken from previous studies [38–40], which stem from possible ranges of the turbulence model closure coefficients of the Menter baseline model. Note that only those coefficients that were found to contribute significantly to uncertainty in various output quantities studied in previously performed sensitivity studies are considered here.

Within the combustion model, the EDC_1 constant controls the rate of heat release while EDC_t was added to the modified Magnussen model [17] that limits the turbulent mixing time scale, which functions to set an upper limit for the reaction rate in regions of high turbulence. The turbulent Schmidt number, Sc_t , effectively controls mixing in the combustor. EDC_t , EDC_1 , and Sc_t uncertainties are taken from expert opinion regarding the possible ranges of these combustion model parameters in the absence of calibration data. The impact of calibration on the combustion model to the HDCR data is discussed later. The fuel-air equivalence ratio uncertainty, ϕ , was the possible variation in the combustor equivalence ratio. The equivalence ratio being used as an uncertainty factor combines the fuel mass flow rate and the inflow air mass flow rate uncertainties into a single parameter. Uncertainty for this parameter was taken from the pre-flight HF2 values [7], which were obtained from a combination of the mechanical fuel system uncertainties and expert opinion.

Another possible source of uncertainty was the angle of attack. Including angle of attack as a factor in a CFD run matrix is costly. Angle of attack variations necessitates doubling the number of cells from a quarter-symmetry grid to a half-symmetry grid. Additionally, replacing a symmetry boundary condition plane with a solved domain to accommodate angle of attack changes can often result in solution unsteadiness for massively separated isolator flows.

To address whether angle of attack should be included in the UQ effort, a one-factor-at-a-time study was performed using three different angles of attack: 0.0 degrees, 0.4 degrees, and 0.8 degrees. The angle of attack range was taken from expected variation during the flight test. All three angles of attack were run on the half symmetry grid, which allowed the 0.0 degree case to provide some insight into the potential asymmetry of the combustor solution despite symmetric inflow boundary conditions. For all three cases, the combustion-induced shock location exhibited unsteady behavior. To get a single solution for these cases, the time integration scheme was changed from a local time step to a global time step. The solutions were then ensemble averaged over two cycles of the combustion-induced shock location variation. An example of the combustion-induced shock location variation magnitude as a function of time, as well as a representative time-averaging window for one case, can be seen in Fig.13.

The ensemble averaged centerline pressure results of the angle of attack study are presented in Fig.14. A total variation in the combustion-induced shock location for the range of angle of attack values investigated was 0.22 inches on the body-side isolator wall and 0.05 inches on the cowl-side isolator wall. Given the large computational cost of including angle of attack as a source of input uncertainty to the CFD model, the uncertainty in the results introduced by eliminating angle of attack as a variable in this study was deemed acceptably small.




Fig. 13 Unsteady variation in the combustion-induced pressure rise as a function of time.




Fig. 14 Centerline wall pressure results from the angle of attack study.

B. CFD Model-Form Uncertainty

The only source of information available to perform a validation study of the CFD model, pre-flight, was the HDCR data. A compilation of combustion-induced shock location as a function of time for different ground test runs at the Mach 6.5 case are shown in Fig.15. The time variation is irrelevant for this study as each time instance for each run is assumed to be a single data point. In this figure, the different colors and symbols represent different runs at the same tunnel conditions, which is also not important for the current study.

All together, this data set of about 150 points is a measure of ground test repeatability and is the only source of quantifiable uncertainty. Measurement uncertainty was investigated, but was determined to be very small relative to the observed repeatability. The horizontal lines on Fig.15 are the 95% coverage interval based on the median absolute deviation. As described by Hemsch [41], the presence of outliers or data points that do not belong to the population can off-set the mean and inflate the sample standard deviation. Therefore, using the median is a more robust approach to measuring the statistics of this dataset.

From Fig.15, there are multiple outlier data points. This suggested that further investigation of these points was warranted to determine if these points should be excluded from the dataset. Unfortunately, no evidence in the HDCR ground test logs could be found that would justify removing these points. Several theories were formulated, but nothing




Fig. 15 Scatter of HDCR combustion-induced shock location measurement.

definitive panned out. Therefore, the outliers were retained in the dataset for validation purposes. This may be typical due to complexities associated with ground testing air-breathing propulsion systems, which highlights the importance of capturing the uncertainty when making performance predictions.

CFD of the HDCR test setup was performed using the modeling approach described in Section II.B. Additionally, turbulence and combustion model input uncertainty sources were propagated through the CFD model using a second order PCE. Based on Eq. (2), at least 21 CFD model evaluations are required for five uncertain parameters. In total, 33 were used in this study to improve the least-squares regression accuracy. A comparison of the CFD p-box and the HDCR shock location uncertainty are shown in Fig.16.




Fig. 16 P-boxes of the HDCR and CFD predicted combustion-induced shock locations.

From this result, a left validation metric, d_L , was found to be 1.84 inches, while there was no evidence for disagreement on the right (i.e., $d_R = 0$). Note here that variation in the HDCR measurements are treated as random, hence the single discrete CDF of the measurements. This 1.84 inches is then the model-form uncertainty that was appended to the flight prediction of the shock location in the isolator, which is discussed later.

A note regarding this comparison should be made. The flow coming into the isolator was assumed to be uniform and free of substantial dispersion in the CFD. Additionally, the smooth wall, as-designed geometry was assumed in the CFD, which based on evaluations of the HDCR rig during the course of testing, proved not to be the case [4, 6]. However, this is the only data available to quantify model-form uncertainty. Because of the direct connect test nature and how the CFD was run, this model-form uncertainty estimate is expected to be conservative.

C. Inlet Unstart Threshold

To determine the uncertainty in the last stable shock location in the isolator (or the operating limit), the skin friction, momentum thickness, and dynamic pressure distributions along the isolator, upstream of the combustion-induced shock were passed to the LIM for each CFD run. A second order PCE of the last stable shock location was constructed, and the uncertainty in Table 2 was then propagated through the PCE by using the second-order probability approach. With eight sources of uncertainty, Eq. (2) dictates that at least 45 CFD model evaluations are needed. In this study, 70 are used to improve the accuracy of the least-squares regression.

After propagating the uncertainty through the PCE, a 95% uncertainty interval of the last stable shock location or operating limit was found to be [40.46, 41.03] in. This result includes the verification errors associated with grid error and surrogate model fit error, both of which were relatively small. As expected, the combustion model parameters (EDC_t , EDC_1 , and the $\dot{m}_f$ factor) have no impact on the variation of the last stable shock location. However, about 50% of the variation came from Mach number uncertainty. The remainder comes from the turbulence model parameter and dynamic pressure uncertainties, with each contributing about the same amount to the total variance.

D. Uncalibrated Flight Reliability

After determining the uncertainty in the last stable shock location, the next step is to quantify the uncertainty in the combustion-induced shock location. The same 70 CFD model evaluations were used, and a new PCE model was constructed for the shock location. The p-box of the combustion-induced shock location, along with the p-box of the operating limit are given in Fig.17. From this result, the 95% probability length margin was found to be 3.68 in., which is measured from the 97.5% probability level of the last stable shock location and the lower 2.5% probability level of the shock location. The 95% probability level confidence ratio, CR , was then calculated to be 0.91. In terms of sensitivities, about 45% of the variation in the shock location comes from the turbulent Schmidt number uncertainty, 30% from freestream Mach number variations, and about 15% from equivalence ratio uncertainty. The remainder of the parameters contribute relatively small amounts of the total shock location variation.

While this measure of reliability is encouraging in terms of the system reliability, there is one key component missing from the uncertainty build-up. While this reliability estimate is inclusive of input uncertainty sources and verification errors in the CFD model, what is missing is a measure of the predictive capability of the model, or rather, the amount of model-form uncertainty. This estimate is determined by comparing model predictions to test data and, in this work, by using the validation approach outlined in Section III.E.

Appending the model-form uncertainty estimate to the previously determined input uncertainty and verification errors leads to the p-boxes shown in Fig.17. From this result, the length based margin was found to be 1.84 in. So, by definition, this system was predicted to have positive margin, albeit a small amount.

With the model-form uncertainty included, the resulting CR was found to be about 0.31. This means that the amount of margin in the system is low relative to the amount of uncertainty. While the system may have positive margin, it may be considered to have low reliability. This is an important result because the reliability of the HF2 system, in a statistically traditional sense, was found to be much greater than 99%. This can be misleading for such a complex system that is known to be sensitive to even small changes in operating conditions (geometry, boundary conditions, etc.). The impact of the CR result of the traditional reliability measure is that the CR communicates that there is not much room for variability in system operation outside of what was modeled here. Any unknown variation or unmodeled effect (such as angle of attack uncertainty already discussed and excluded due to computational cost) may result in the system having negative margin.

Note that this result is a function of the full uncertainty ranges used for the combustion model parameters, namely EDC_t , EDC_1 , and Sc_t . One of the primary purposes of the direct connect test was to calibrate the heat release model, which would in turn reduce the variability in the combustion-induced shock location in the isolator. However, because of the upstream-bias nature of the model-form uncertainty, reducing the range of predicted shock locations through calibration of the combustion model will only increase model-form uncertainty (assuming that the current parameter uncertainties bound all possible solutions). Due to the noted issues with the current comparison and model-form uncertainty estimate, a more representative test, such as a free-jet test, or a more controlled validation quality experiment is needed to better improve the model-form uncertainty estimate.

VI. Conclusions

Overall, this work describes an end-to-end process for determining the unstart reliability of a dual-mode ramjet system. A framework based on the quantification of margins and uncertainties was employed to communicate the




Fig. 17 P-boxes of the last stable and combustion-induced shock locations.

confidence of a system in terms of operating margin and uncertainty. The process of performing model validation to determine model-form uncertainty using test data and then propagating that forward, along with other sources of uncertainty, to determine flight reliability was illustrated. The HIFiRE Flight 2 experiment and related ground testing was used as the problem of interest. Results of the reliability analysis suggested that the system may have very high reliability in a traditional sense, but the confidence in the system may be low as a result of the amount of uncertainty relative to the margin. However, designing a system with a high confidence ratio may be impractical. If a complete accounting of uncertainties, including testing, model-form, and inputs, is performed, then engineering judgment may be used to determine if the system is considered reliable. Comprehensive and detailed UQ is the key to designing reliable and robust hypersonic air-breathing propulsion systems.

Acknowledgments

This work was funded by NASA's NASA Hypersonic Technology Project. Michael Bynum was supported through the 80LARC17C0003 (Technology Engineering and Aerospace Mission Support 3 – TEAMS3) contract between NASA Langley Research Center and AMA, Inc.

References

- [1] Dolvin, D., "Hypersonic International Flight Research and Experimentation Technology Development and Flight Certification Strategy," AIAA Paper 2009-7228, 2009. doi:10.2514/6.2009-7228.
- [2] Jackson, K., Gruber, M., and Buccellato, S., "HIFiRE Flight 2 Project Overview and Status Update 2011," AIAA Paper 2011-2202, 2011. doi:10.2514/6.2011-2202.
- [3] Guy, R., Rogers, R., Puster, R., Rock, K., and Diskin, G., "The NASA Langley Scramjet Test Complex," AIAA Paper 96-3243, 1996. doi:10.2514/6.1996-3243.
- [4] Hass, N., Cabell, K., and Storch, A., "HIFiRE Direct-Connect Rig (HDCR) Phase I Ground Test Results from the NASA Langley Arc-Heated Scramjet Test Facility," JANNAF 31st Airbreathing Propulsion Subcommittee Meeting, 2009.
- [5] Hass, N., Cabell, K., Storch, A., and Gruber, M., "HIFiRE Direct-Connect Rig (HDCR) Phase I Scramjet Test Results from the NASA Langley Arc-Heated Scramjet Test Facility," AIAA Paper 2011-2248, 2011. doi:10.2514/6.2011-2248.
- [6] Hass, N., and Cabell, K., "Results from HDCR Phase 2 Testing at the NASA LaRC AHSTF," JANNAF 45th Joint CS/APS/EPSS/PSHS Meeting, 2012.
- [7] Storch, A. M., and Bynum, M. D., "HIFiRE Flight 2 Ground-to-Flight Dual-Mode Scramjet Operability Prediction Methodology," *JANNAF Journal of Propulsion and Energetics*, Vol. 6, No. 2, 2014, pp. 107–122.

- [8] Sharp, D. H., and Wood-Schultz, M. M., “QMU and Nuclear Weapons Certification. What’s Under the Hood,” *Los Alamos Science*, Vol. 1, No. 28, 2003, pp. 47–53.
- [9] Brune, A. J., West, T. K., and White, L. M., “Calibration Probe Uncertainty and Validation for the Hypersonic Material Environmental Test System,” *Journal of Thermophysics and Heat Transfer*, Vol. 34, No. 2, 2020, pp. 404–420. doi: 10.2514/1.T5839.
- [10] Ferlemann, P., and McDaniel, K., “HIFiRE Flight 2 Forebody, Inlet, and Nozzle Design Validation Using Flight Data,” *JANNAF Journal of Propulsion and Energetics*, Vol. 6, No. 2, 2014, pp. 123–138.
- [11] R. A. Baurle, T. G. D., J. A. White, and Norris, A. T., “VULCAN-CFD User Manual: Ver. 7.1.0,” Tech. Rep. TM-2020-5000767, NASA, 2020.
- [12] Menter, F. R., “Two-equation eddy-viscosity turbulence models for engineering applications,” *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598–1605. doi:10.2514/3.12149.
- [13] Wilcox, D., “Wall matching, a rational alternative to wall functions,” AIAA Paper 1989-611, 1989. doi:10.2514/6.1989-611.
- [14] Edwards, J. R., “A low-diffusion flux-splitting scheme for Navier-Stokes calculations,” *Computers and Fluids*, Vol. 26, No. 6, 1997, pp. 635 – 659. doi:10.1016/S0045-7930(97)00014-5.
- [15] Hirsch, C., *Numerical Computation of Internal and External Flows, Volume 2: Computational Methods for Inviscid and Viscous Flows*, John Wiley and Sons, Inc., New York, 1990.
- [16] Magnussen, B., and Hjertager, B., “On mathematical modeling of turbulent combustion with special emphasis on soot formation and combustion,” *Symposium (International) on Combustion*, Vol. 16, No. 1, 1977, pp. 719 – 729. doi:10.1016/S0082-0784(77)80366-4.
- [17] Norris, A. T., “Tuning Process for the Modified Magnussen Combustion Model,” AIAA Paper 2018-4739, 2018. doi: 10.2514/6.2018-4739.
- [18] Hutzler, J., Decker, D., Cobb, R., King, P., Veth, M., and Donbar, J., “Scramjet Isolator Shock Train Location Techniques,” AIAA Paper 2011-402, 2011. doi:10.2514/6.2011-402.
- [19] Auslender, A. H., “An Overview of the Langley-Isolator-Model Analysis Methodology,” JANNAF 35st Airbreathing Propulsion Subcommittee Meeting, 2016.
- [20] Roache, P. J., “QUANTIFICATION OF UNCERTAINTY IN COMPUTATIONAL FLUID DYNAMICS,” *Annual Review of Fluid Mechanics*, Vol. 29, No. 1, 1997, pp. 123–160. doi:10.1146/annurev.fluid.29.1.123.
- [21] Oberkampf, W., Helton, J., and Sentz, K., “Mathematical representation of uncertainty,” AIAA Paper 2001-1645, 2001. doi:10.2514/6.2001-1645.
- [22] Swiler, L. P., and Eldred, M. S., “Efficient algorithms for mixed aleatory-epistemic uncertainty quantification with application to radiation-hardened electronics. Part I, algorithms and benchmark results.” , September 2009. doi:10.2172/972887.
- [23] Xiu, D., and Karniadakis, G., “The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations,” *SIAM Journal on Scientific Computing*, Vol. 24, No. 2, 2002, pp. 619–644. doi:10.1137/S1064827501387826.
- [24] Ghanem, R. G., and Spanos, P. D., *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, 1991. doi:10.1007/978-1-4612-3094-6.
- [25] Eldred, M. S., “Recent Advances in Non-Intrusive Polynomial Chaos and Stochastic Collocation Methods for Uncertainty Analysis and Design,” AIAA Paper 2009-2274, 2009. doi:10.2514/6.2009-2274.
- [26] Hosder, S., and Bettis, B., “Uncertainty and Sensitivity Analysis for Reentry Flows with Inherent and Model-Form Uncertainties,” *Journal of Spacecraft and Rockets*, Vol. 49, No. 2, 2012, pp. 193–206. doi:10.2514/1.A32102.
- [27] West, T. K., Hosder, S., and Johnston, C. O., “Multistep Uncertainty Quantification Approach Applied to Hypersonic Reentry Flows,” *Journal of Spacecraft and Rockets*, Vol. 51, No. 1, 2014, pp. 296–310. doi:10.2514/1.A32592.
- [28] Lockwood, B., and Mavriplis, D., “Gradient-based methods for uncertainty quantification in hypersonic flows,” *Computers and Fluids*, Vol. 85, 2013, pp. 27 – 38. doi:10.1016/j.compfluid.2012.09.003.

- [29] Roderick, O., Anitescu, M., and Fischer, P., "Polynomial Regression Approaches Using Derivative Information for Uncertainty Quantification," *Nuclear Science and Engineering*, Vol. 164, No. 2, 2010, pp. 122–139. doi:10.13182/NSE08-79.
- [30] Bettis, B., Hosder, S., and Winter, T., "Efficient Uncertainty Quantification in Multidisciplinary Analysis of a Reusable Launch Vehicle," AIAA Paper 2011-2393, 2011. doi:10.2514/6.2011-2393.
- [31] Hosder, S., Walters, R. W., and Balch, M., "Point-Collocation Nonintrusive Polynomial Chaos Method for Stochastic Computational Fluid Dynamics," *AIAA Journal*, Vol. 48, No. 12, 2010, pp. 2721–2730. doi:10.2514/1.39389.
- [32] Han, D., and Hosder, S., "Inherent and Model-Form Uncertainty Analysis for CFD Simulation of Synthetic Jet Actuators," AIAA Paper 2012-0082, 2012. doi:10.2514/6.2012-82.
- [33] Hosder, S., Walters, R. W., and Balch, M., "Efficient Sampling for Non-Intrusive Polynomial Chaos Applications with Multiple Uncertain Input Variables," AIAA Paper 2007-0125, 2007. doi:10.2514/6.2007-1939.
- [34] Cameron, R. H., and Martin, W. T., "The Orthogonal Development of Non-Linear Functionals in Series of Fourier-Hermite Functionals," *Annals of Mathematics*, Vol. 48, No. 2, 1947, pp. 385–392. doi:10.2307/1969178.
- [35] Ferson, S., Oberkampf, W. L., and Ginzburg, L., "Model validation and predictive capability for the thermal challenge problem," *Computer Methods in Applied Mechanics and Engineering*, Vol. 197, No. 29, 2008, pp. 2408–2430. doi:10.1016/j.cma.2007.07.030.
- [36] West, T. K., Hosder, S., and Winter, T., "Quantification of Margins and Uncertainties for Integrated Spacecraft Systems Models," *Journal of Spacecraft and Rockets*, Vol. 52, No. 2, 2015, pp. 450–461. doi:10.2514/1.A33067.
- [37] Hass, N., and Cabell, K., "Forebody and Inlet Design for the HIFiRE 2 Flight Test," 55th JANNAF Propulsion Meeting, 2008.
- [38] Erb, A., and Hosder, S., "Uncertainty Analysis of Turbulence Model Closure Coefficients for Wall-Bounded Attached and Separated Flows," AIAA Paper 2017-1952, 2017. doi:10.2514/6.2017-1952.
- [39] Erb, A., and Hosder, S., "Uncertainty Analysis of Turbulence Model Closure Coefficients for Shock Wave-Boundary Layer Interaction Simulations," AIAA Paper 2018-2077, 2018. doi:10.2514/6.2018-2077.
- [40] Schaefer, J., Hosder, S., West, T., Rumsey, C., Carlson, J.-R., and Kleb, W., "Uncertainty Quantification of Turbulence Model Closure Coefficients for Transonic Wall-Bounded Flows," *AIAA Journal*, Vol. 55, No. 1, 2017, pp. 195–213. doi:10.2514/1.J054902.
- [41] Hemsch, M. J., "Statistical Analysis of Computational Fluid Dynamics Solutions from the Drag Prediction Workshop," *Journal of Aircraft*, Vol. 41, No. 1, 2004, pp. 95–103. doi:10.2514/1.1796.