Scalable Environment for Quantification of Uncertainty and Optimization in Industrial Applications (SEQUOIA)

Juan J. Alonso[*1], Richard W. Fenrich[†1], Victorien Menier[‡1], Gianluca Iaccarino[§1], Aashwin A. Mishra[¶1], Michael S. Eldred[‖2], John Jakemann[**2], Paul Constantine[††3] and Karthik Duraisamy[‡‡4]

[1]Department of Aeronautics & Astronautics, Stanford University
[2]Sandia National Laboratories
[3]Department of Applied Mathematics & Statistics, Colorado School of Mines
[4]Department of Aerospace Engineering, University of Michigan

As part of the DARPA EQUiPS (Enabling Quantification of Uncertainty in Physical Systems) program, the SEQUOIA project provides an integrated plan for performing uncertainty quantification (UQ) and design under uncertainty (DUU) that aggressively pursues new frontiers in scale and complexity. In particular, the coordinated investments in this effort will create advancements in scalable forward and inverse UQ algorithms and the rigorous quantification of model inadequacy, providing the primary foundation for the development of generalized stochastic design approaches that address the robustness and reliability of complex multi-disciplinary systems. This project will demonstrate new UQ methods on high-performance aircraft nozzle analysis and design problems that are simultaneously designed for aerodynamic performance, thermal and pressure loads, and fatigue, while subject to geometric constraints to be fully integrated with complex vehicle shapes. The ability to handle all kinds of uncertainty at very large scale will enable the design of future components and vehicles that can have a substantial impact on DARPAs mission. This paper describes the efforts of our team thus far, and the accomplishments we have completed in the pursuit of the overall goals of the project.

I. Introduction

THE The rapid advancement of both computer hardware and physical simulation capabilities has revolutionized science and engineering, placing computational simulation on an equal footing with theoretical analysis and physical experimentation. This rapidly increasing reliance on the predictive capabilities of computational models has created the need for rigorous quantification of the effect that all types of uncertainties have on these predictions.

In this project and paper, we present an integrated plan for performing uncertainty quantification (UQ) and design under uncertainty (DUU) that aggressively pursues new frontiers in scale and complexity. In particular, the coordinated investments in this proposal are intended to create advancements in scalable

---

[*]Professor, AIAA Associate Fellow
[†]Graduate student, AIAA Member
[‡]Post-doctoral Associate, AIAA Member
[§]Associate Professor, AIAA Associate Fellow
[¶]Post-doctoral Associate, AIAA Member
[‖]Research Staff, AIAA Associate Fellow
[**]Research Staff, AIAA Member
[††]Assistant Professor, AIAA Member
[‡‡]Assistant Professor, AIAA Member

American Institute of Aeronautics and Astronautics

forward and inverse UQ algorithms and the rigorous quantification of model inadequacy, providing the primary foundation for the development of generalized stochastic design approaches that address the robustness and reliability of complex multi-disciplinary systems. This project will demonstrate new UQ methods on high-performance aircraft nozzle problems that are simultaneously designed for aerodynamic performance, thermal and pressure loads, and fatigue, while subject to geometric constraints to be fully integrated with complex vehicle shapes. The ability to handle all kinds of uncertainty at very large scale will enable the design of future components and vehicles that can have a substantial impact on DARPA's mission.

We have assembled a team of world-class engineers, scientists, and applied mathematicians with a demonstrated track record of tackling major challenges and arriving at solutions that are widely recognized as seminal in their fields. In addition, our team has pioneered the use of open-source software (e. g. Dakota and SU2) to leverage investments in fundamental research across a broad spectrum of users. The team includes researchers from Stanford University, Sandia National Laboratories, the University of Michigan, and the Colorado School of Mines.

Our effort is divided into research elements in three separate Thrust Areas (TAs): Thrust Area 1 - Scalable Algorithms, Thrust Area 2 - Model-Form Uncertainty, and Thrust Area 3 - Design and Decision Making Under Uncertainty.

Within Thrust Area 1, we will combine scalable polynomial chaos expansions with explicit dimension reduction in order to accurately resolve the most influential subspace within large-scale parameter domains. These approaches will be tailored and applied within both forward and inverse UQ contexts, as we seek to discover low-dimensional structure within high-dimensional domains, and will enable the use of very large numbers of uncertain parameters in our UQ and DUU efforts.

In Thrust Area 2, we will address the critical challenge of model inadequacies and develop methodologies for determining the source and extent of model-form discrepancies in simulation models by using observational data as well as data from high-fidelity reference simulations. These methods enable accurate predictions with quantified uncertainty, despite the presence of imperfect sub-model empiricisms and closures. A novel formulation for embedded discrepancy is proposed, combining scalable inversion techniques in Thrust Area 1 along with large-scale machine learning techniques. Multifidelity algorithms for forward and inverse UQ will leverage the predictive capability of less expensive models through sparse, adaptive recovery of model discrepancy.

In Thrust Area 3, we build on the foundation provided by the previous thrust areas and incorporate rigorous, scalable UQ within the context of specific design problems, optimizing for robustness and reliability in the presence of variability, lack of knowledge, and model inadequacy. We employ a trust-region model management framework for DUU that exploits synergies within a hierarchy of statistical fidelity, provably achieving high-fidelity results at much lower cost.

Fundamental research breakthroughs in the efficacy of UQ and DUU methods for complex multiphysics systems will enable the deployment of UQ algorithms to higher-fidelity computational simulations operating in more complex high-dimensional uncertain environments, and paving the way for greater robustness and reliability in advanced vehicle designs of interest to DARPA and extending the predictive capabilities of computational models into new regimes.

II.   Details of Individual Efforts in SEQUOIA Project

In a large and integrated effort like SEQUOIA that involves multiple institutions, the work is usually divided into a number of separate Tasks that contribute some of the foundational capabilities and that are then combined into significant integrated milestones and accomplishments. In this paper, we briefly highlight each of the major tasks in our work and what has been accomplished thus far. The Tasks belong to different Thrust Areas (as defined above). A more detailed description of research directions in each of the Thrust Areas is provided below, prior to the technical work in each of the tasks.

Within Thrust Area 1, we will combine scalable polynomial chaos expansions with explicit dimension reduction in order to accurately resolve the most influential subspace within large-scale parameter domains. These approaches will be tailored and applied within both forward and inverse UQ contexts, as we seek to discover low-dimensional structure within high-dimensional domains, and will enable the use of very large numbers of uncertain parameters in our UQ and DUU efforts.

In Thrust Area 2, we will address the critical challenge of model inadequacies and develop methodologies for determining the source and extent of model-form discrepancies in simulation models by using i)

American Institute of Aeronautics and Astronautics

observational/high-fidelity simulation data, and ii) Physics constraints. These methods enable accurate predictions with quantified uncertainty, despite the presence of imperfect sub-model empiricisms and closures and are constrained by the underlying physics. A novel formulation for embedded discrepancy is pursued, combining scalable inversion techniques in Thrust Area 1 along with large-scale machine learning techniques. Multifidelity algorithms for forward and inverse UQ will leverage the predictive capability of less expensive models through sparse, adaptive recovery of model discrepancy.

In Thrust Area 3, we build on the foundation provided by the previous thrust areas and incorporate rigorous, scalable UQ within the context of specific design problems, optimizing for robustness and reliability in the presence of variability, lack of knowledge, and model inadequacy. We employ a trust-region model management framework (TRMM) for DUU that exploits synergies within a hierarchy of statistical fidelity, provably achieving high-fidelity results at much lower computational cost.

Moreover, we also focus on the application of advances in these three research areas to the problem of stochastic design of advanced aircraft nozzle systems. For this purpose we construct a hierarchy of simulation capabilities of increasing fidelity and cost, so that research contributions can be tested early (using low-fidelity models) in relevant environments, and so that the different simulation capabilities can be combined into multi-fidelity techniques using common parameterization and stochasticity.

Fundamental research breakthroughs in the efficacy of UQ and DUU methods for complex multiphysics systems will enable the deployment of UQ algorithms to higher-fidelity computational simulations operating in more complex high-dimensional uncertain environments, and paving the way for greater robustness and reliability in advanced vehicle designs of interest to DARPA and extending the predictive capabilities of computational models into new regimes.

In each of the research task descriptions below, the Thrust Area that the effort belongs to and the lead institution in the team (SU - Stanford University, SNL - Sandia National Labs, CSM - Colorado School of Mines, and UM - University of Michigan) are listed next to the task Title.

SCALABLE TECHNIQUES FOR FORWARD UQ (TA1, SNL):   In this task, we have have completed the development of minimal cubature rules for multivariate integration. Work focused on extending the quadrature method so that it can now be used to compute quadrature rules on the zonotopes of active subspaces. The development of minimal cubature rules in combination with the previously reported methods for building approximations of high-dimensional function has provided us with an essential and strong foundation for combining scalable forward UQ methods with dimension reduction . The aforementioned work will also be extremely useful for developing efficient bayesian inference methods using surrogates and dimension reduction, as part of this project.

SCALABLE FORWARD UQ USING DIMENSION REDUCTION (TA1, CSM):   In this effort, we have completed applying an initial set of tests to identify low-dimensional structure on the combined design / uncertain parameter space in quantities of interest from the nozzle model. The results of these tests will greatly reduce the cost of the full-scale design under uncertainty problem through low-dimensional surrogates and quadrature for both objectives and constraints. To build on this work we have also demonstrated the efficacy of using compressed sensing and cross validation to respectively build PCE over active subspaces and to select the 'best' dimension of the active subspace given accuracy requirements of the DUU of the low-fidelity nozzle. We have also deployed our minimal cubature technique (Task 1) to estimate moments of important QoI, of the low-fidelity nozzle, which appear in our DUU formulation. These activities culminate in the capability to combine formal dimension reduction methods with scalable UQ algorithms, enabling us to tackle large-scale UQ problems with greatly reduced expense. A more detailed discussion of these capabilities can be found later on in this paper.

EFFICIENT GRADIENT SAMPLING FOR ACTIVE SUBSPACES (TA1, CSM):   As part of this effort, we have completed a very useful proof. Active subspaces are defined by the eigenspaces of a matrix of integrals related to the function's gradient. We have proved that the integrands in these particular integrals have effective dimension 2—*no matter the dimension of the input space.* Quasi Monte Carlo (QMC) methods are well known to perform well on functions that have low effective dimension. This result provides the theoretical foundation for superior performance of QMC over simple Monte Carlo for estimating active subspaces. This impetus to study this question came from the recent Monte Carlo Quasi Monte Carlo (MCQMC) conference at Stanford. These activities culminate in a solid theoretical foundation for using QMC to estimate active

American Institute of Aeronautics and Astronautics

subspaces. When this result is written up and complemented with numerical experiment, it will represent a significant advance.

GRADIENT-FREE METHODS FOR ACTIVE SUBSPACES (TA1, CSM):   We have completed analysis showing conditions when statistical methods for *sufficient dimension reduction* (SDR), such as sliced inverse regression (SIR) and sliced average variance estimation (SAVE), identify active subspaces. In the process, we laid the theoretical framework for applying and interpreting SDR methods with a set of samples from deterministic computer simulations. These activities culminate in a set of statistical methods—none of which require gradients—for estimating active subspaces for complex simulation models. We have applied the simplest of these approaches—the ordinary least squares method—to the nozzle application to reveal simplified relationships and low intrinsic dimensionality in several quantities of interest.

ERROR ESTIMATORS FOR ROMS OF NOZZLE APPLICATION (TA1, SU):   Over the last year, we have refined our residual-based error indicator for *nonlinear* dynamical systems such as those associated with nonlinear structural dynamics problems. We have also established that this indicator is reliable and computationally efficient. We have successfully used it for accelerating an adaptive parameter sampling strategy for the construction of both global and local reduced-order bases. We have assessed its performance for various solid mechanics and structural vibrations problems, including all nozzle problems discussed in this report. Specifically, we have exploited it to construct effective nonlinear reduced-order and hyper reduced-order structural dynamics models.

STOCHASTIC ROM CONSTRUCTION APPROACHES (TA1, SU):   We have extended our nonparametric probabilistic approach for modeling uncertainties associated with nonlinear, projection-based, model order reduction methods to the case of hyper reduced-order models in order to achieve real-time performance for uncertainty quantification. Specifically, we have developed two mathematical approaches that combine our Energy Conserving Sampling and Weighting hyper reduction methodology for nonlinear computational models with our nonparametric probabilistic approach for quantifying uncertainties to enable the construction of parametric stochastic hyper reduced-order models (SHROMs) that can perform in real-time. The first approach can be described as a zero-level formulation where the sampled mesh elements and associated weights constituting the reduced mesh are treated as fixed quantities during the stochastic analysis, so that they can be rapidly determined. The resulting SHROM successfully accounts for the interpolation-induced errors, but the approach increases a priori the width of the confidence region of the random QoI for a given level of probability. The second approach fixes the latter issue by recomputing the weights in the stochastic framework for each determined set of hyper parameters until convergence of the coupled weight-hyper-parameter problem is attained. Both approaches have been mathematically analyzed and are being programmed for application to nozzle problems.

INFERENCE FOR EMBEDDED MODEL DISCREPANCY (TA2, UM):   In this effort, along with other efforts described later on we have focused on developing a data-driven embedded UQ strategy to address model-from discrepancy. Data from higher-fidelity simulations (and experiments, if available) is used in a statistical inference framework to infer the spatial distribution of modeling discrepancies.

In previous work, demonstrations were shown on 1D channel flow problems. In the current status of this work, we have extended the framework to infer the spatial form of turbulence model discrepancies in the main application - *i.e.* axisymmetric nozzle flow. Specifically, we used reference skin-friction data from a full Reynolds stress closure to infer the field discrepancy in an eddy viscosity (Spalart-Allmaras) model. The inverse problem consisting of 30,000 dimensions was accelerated by a) dimensionality-reduction using the active subspace of the Maximum a posteriori estimate, and b) stratified MCMC sampling. This inversion was performed over a large number of nozzle configurations encountered in the medium-fidelity DUU computations.

PHYSICS-INSPIRED MODEL ERROR ESTIMATES (TA2, SU):   Over the past year, we have formulated and validated a comprehensive methodology for uncertainty estimation of general RANS-based turbulence closures. This involves injection of perturbations directly into the eigenvalues and eigenvectors of the modeled Reynolds stress tensor. To circumvent any additional assumptions, this methodology conforms to the single-point modeling framework. In this vein, the perturbations of the Reynolds stress eigenvalues are governed by sampling from the extreme states of the turbulence componentiality. Similarly, the perturbations to

American Institute of Aeronautics and Astronautics

the Reynolds stress eigenvectors are guided by the maximal states of the production mechanism. Jointly, these provide for a schema wherein the envelope of model-form uncertainty can be estimated within two simulations of the RANS model, minimizing the computational overhead.

We have applied this methodology to a series of canonical turbulent flows, with and without separation, while contrasting the results against numerical and experimental data. In each case, we were able to obtain *prudent* uncertainty bounds. For instance, the results in the turbulent flow in an asymmetric diffuser are outlined in Figure 1. As can be seen, uncertainty bounds on the k-ω-SST model can be obtained with just two RANS simulations, minimizing the computational overhead, for both the skin friction and the flow streamlines. Similarly in Fig. 2 we demonstrate envelopes for one of the axisymmetric nozzles under investigation in the project demonstrating the ability of this approach to envelope both wall quantities – skin friction on the nozzle internal surface – and flow field quantities, i.e. the Mach number on the nozzle axis.

In addition we have continued to investigate the model uncertainties in RANS closures as described solely in terms of the Reynolds stresses. For a wide variety of flows, we studied the prediction intervals engendered. For instance, we found that while strain-dominated flows were highly sensitive to this inadequate representation of the turbulent flow field, rotation dominated flows were largely unaffected. This study provides us with valuable insight into the physics underlying one of the major sources of single-point uncertainty in RANS models.
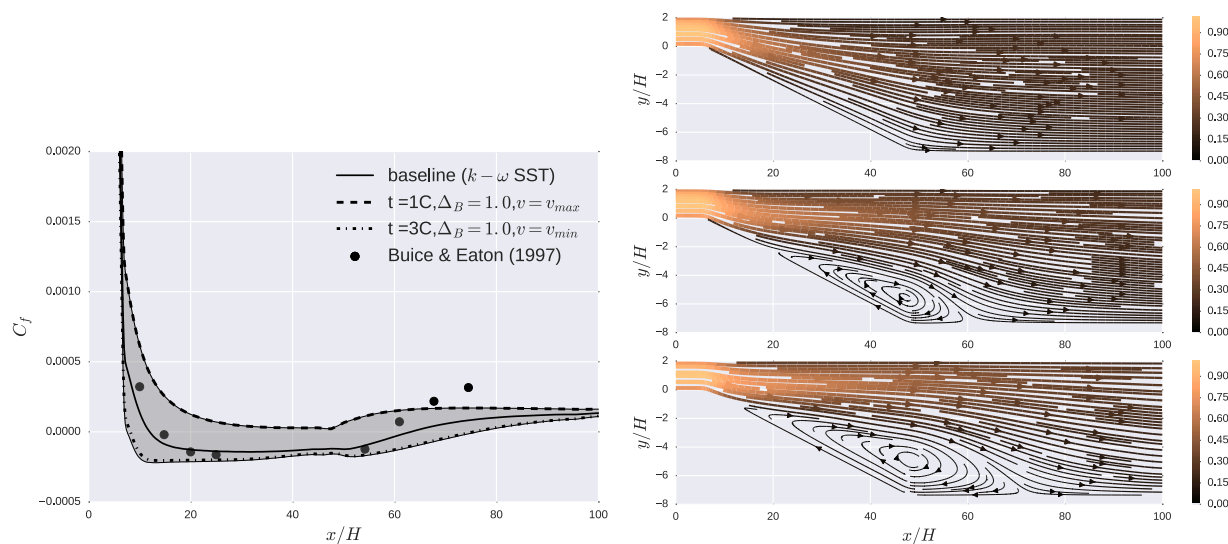

Figure 1. (Left) Uncertainty bounds on the coefficient of friction across the bottom wall of the asymmetric diffuser, (Right) Flow streamlines for the bounding and the baseline cases.
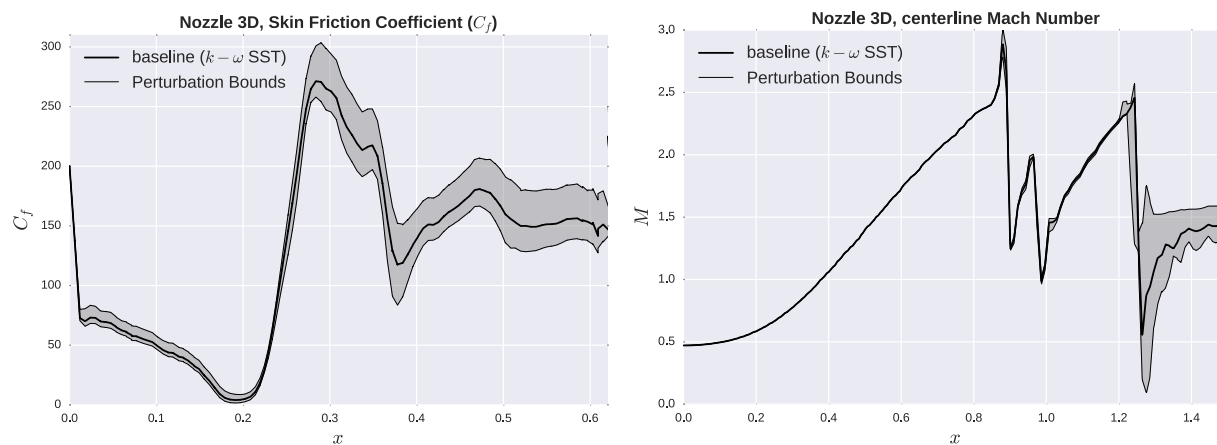

Figure 2. Uncertainty bounds on the coefficient of friction on the nozzle internal surface (left) and the Mach number on the nozzle axis (right).

American Institute of Aeronautics and Astronautics

MULTIFIDELITY METHODS FOR FORWARD UQ (TA2, SNL): This effort focuses on multifidelity methods that exploit multiple model forms, multilevel methods that exploit multiple discretization levels for a particular model form, or multilevel-multifidelity methods that exploit an ensemble containing both model forms and discretization levels. Significant progress has been made in developing new multilevel-multifidelity approaches and deploying to an ensemble of nozzle models comprised of the low-fidelity model and three discretizations of the medium-fidelity model.

The primary multilevel-multifidelity emphasis has been on refining our formulation for multilevel-control variate Monte Carlo. After observing small correlations between low-fidelity (LF) and medium-fidelity (MF) model discrepancies on model problems, we have modified our formulation to replace control variate estimation based on discrepancy-based correlations with a more targeted control variate estimation based on correlations among the individual QoI that comprise the discrepancies. This latest approach has been deployed to the nozzle application, and its performance is compared to multilevel Monte Carlo and standard Monte Carlo in Figure 3.
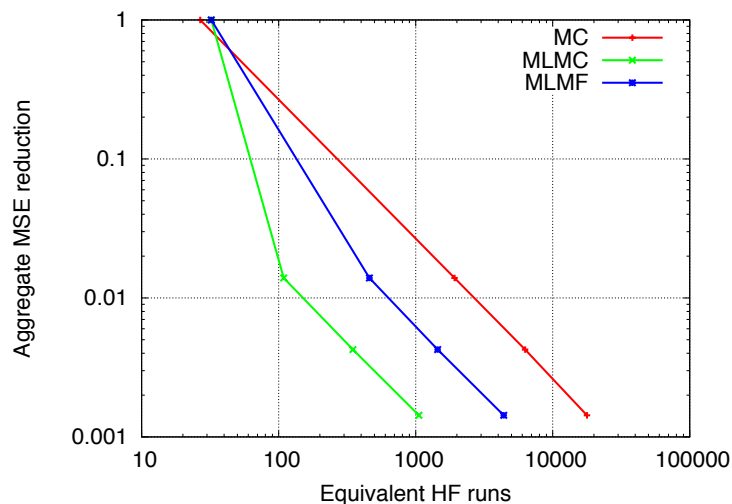

Figure 3. Convergence comparison for multilevel and multilevel-multifidelity methods applied to forward UQ using LF and MF nozzle models.

The corresponding results are reported in Table 1 where the sample profiles for both the LF and MF models are reported as a function of the target accuracy.

Target accuracy	LF	MF		
	Coarse	Coarse	Medium	Fine
0.01	21143	1757	20	20
0.003	69580	5775	36	20
0.001	212828	17715	109	34

Table 1. MLMF sample profile for LF and MF models as function of the mean square error reduction.

Focusing on the control variate portion across model fidelities, we first note that there is only one LF discretization whereas there are three MF model discretizations. We elect to apply the LF model as a control variate in a ragged configuration, i.e., only for the coarse level, since this is typically where the bulk of the variance is present. The correlation between the LF and MF models is reported in Table 2, whereas the variance reduction attained on each level is reported in Table 3.

Turning our attention to the multilevel component, the variance breakdown for the different QoIs is reported in Table 4.

From these results, it is evident that the algorithm is able to obtain a large variance reduction for the thrust, but the much lower correlation for mechanical and thermal stresses does not allow for reduction in

American Institute of Aeronautics and Astronautics

QoI	correlation
Thrust	9.9865e-01
Mech. stresses	4.8106e-03
Thermal stresses	3.7389e-01

Table 2. Pearson correlations between LF and MF models for the different QoIs.

Target accuracy	QoI		
	Thrust	Mech. stresses	Therm. stresses
0.01	9.14e+01	4.56e-02	1.43e+01
0.003	9.14e+01	2.06e-02	1.27e+01
0.001	9.14e+01	2.12e-03	1.28e+01

Table 3. Variance reduction (coarse level) as function of the target accuracy.

Level	QoI		
	Thrust	Mech. stresses	Therm. stresses
Coarse	1.9231e+07	4.2124e+14	2.0025e+14
Medium	3.1952e+04	4.9853e+10	3.5519e+10
Fine	2.3231e+05	6.0606e+07	3.3891e+10

Table 4. Variance as function of the discretization level for the different QoI.

these variances. In aggregate, the variance reduction across the set of QoI is not sufficient to outperform multilevel Monte Carlo for this instantiation of the low fidelity model. This leads to several refinements: *(i)* evaluation of alternatives for a more predictive LF structural model, *(ii)* evaluation of control variate alternatives for ragged model ensembles, and *(iii)* hardening of the MLCV algorithm to adaptively manage the model hierarchy when low correlations are encountered (abandon poor performing control variates).

The family of multilevel Monte Carlo approaches has primary strengths in insensitivity to nonsmoothness and to dimension; thus, they are robust methods that can be very effective for large-scale problems. In this nozzle application deployment, we considered a set of 15 random variables. For this moderate dimension, we can also evaluate methods that exploit smoothness in the QoIs by using our multilevel polynomial chaos expansion (ML PCE) approach. To support random sample sets on unstructured grids (as well as legacy data, fault tolerance, etc.), a compressed sensing strategy can be used to efficiently recover a sparse set of polynomial coefficients. From the hybridization of multilevel sampling and polynomial chaos with compressed sensing, we can replace the Monte Carlo estimator with known estimator variance with an improved estimator with parameterized estimator variance (more details are available in Eldred *et al.*, SIAM UQ 2016). In Figure 4, the performance of this ML PCE is compared with the previous MC, MLMC and MLMF results. The ML PCE employs only realizations from the MF model on the same coarse, medium and fine discretizations of the previous MLMF results; no LF evaluations are used for the current algorithm (prior to planned development of a MLMF PCE approach). For the ML PCE estimator, the actual number of realizations per-level is reported in Table 5 for different target tolerances.

Target accuracy	MF		
1E-1	11	10	10
1E-3	118	10	10
1E-4	374	10	10
1E-5	1182	35	11
1E-6	4048	132	70

Table 5. Samples profile for ML PCE as function of the target normalized MSE.

It is evident that a significant improvement in convergence rate is obtained, reducing evaluation requirements by orders of magnitude for a given accuracy target. Overall, this set of multilevel PCE and multilevel MC approaches provide a very useful tandem for exploiting smoothness in addressing smooth problems in

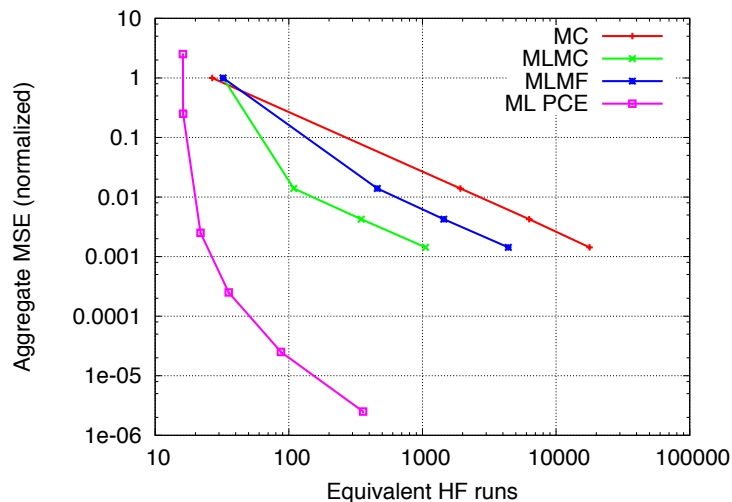American Institute of Aeronautics and Astronautics

Figure 4. Convergence comparison for sparse regression multilevel polynomial chaos expansion (ML PCE) against MC, MLMF and MLMC methods. MSE values are normalized with respect to the accuracy obtained by a set of 20 pilot samples per level.

moderate dimension and robustly addressing nonsmooth responses in higher dimension.

These activities over the first year of the project culminate in an extensive set of algorithms for exploiting a model ensemble in order to accelerate forward UQ. Algorithms include multifidelity methods (multifidelity PCE and control variate MC/LHS), multilevel methods (multilevel MC/LHS and multilevel PCE), and hybrid methods (multilevel control variate MC/LHS). In deploying these algorithms to model problems (elliptic and parabolic PDEs), we have demonstrated: (1) the variance reductions from multilevel and control variate approaches can be effectively combined, resulting in largely additive contributions to the variance reduction, and (2) replacing basic Monte Carlo ensemble-average estimators with more sophisticated sparse recovery estimators extends both the applicability of the multilevel algorithmic machinery and the convergence performance of multilevel methods. In deploying these algorithms to the nozzle application, we have discovered low correlations between low and medium fidelity predictions for some QoI. This hampers the control variate step, indicating the need to either refine the low fidelity predictions (currently under discussion with the modeling team) or adaptively discard poorly correlated models within the overall model ensemble.

MATHEMATICAL FOUNDATIONS FOR DUU FRAMEWORK (TA3, SNL): This task integrates foundational work in a variety of areas in order to facilitate affordable and effective design under uncertainty using high-fidelity computational models. In particular, our DUU processes currently integrate with scalable forward UQ, dimension reduction based on active subspaces, reduced-order modeling, and multilevel-multifidelity approaches that exploit model hierarchies, and integration with embedded model form uncertainty is planned. Multiple threads of DUU effort this quarter are highlighted below.

These activities over the first year of the project culminate in an extensive array of DUU approaches, all based on effective management of approximations of various forms. Looking forward, effective exploitation of multiple acceleration technologies and roll up of a comprehensive set of uncertainties will render DUU approaches tractable and meaningful for real world design of DOD systems using high-fidelity modeling and simulation.

COMPUTATIONAL MODELS FOR LOW- AND MEDIUM-FIDELITY APPLICATIONS (App, SU): We have completed, fine-tuned, and exercised the one-and-a-half-way coupling of the medium-fidelity thermal and structural models with the medium-fidelity fluid model. The coupled medium-fidelity model has now been successfully deployed for the analysis of thousands of different nozzle geometries. In addition, the nozzle has been upgraded from a single-layer single-material design to a bi-layer bi-material design after discussion with

American Institute of Aeronautics and Astronautics

experts in the application area. The upgraded nozzle design implements an inner thermal protection layer made from a ceramic matrix composite material, and an outer structural layer made from a composite panel structure using a titanium honeycomb core. Models for the nozzle thrust and volume reflect the changes to the nozzle model.

Lastly, the low- and medium-fidelity models have been translated from Matlab to a unified Python multifidelity model framework for ease-of-access. Robustness of the convergence of the low- and medium-fidelity models remains the same as before, but now team members are more easily able to run studies and optimizations on computer clusters using the models and switch between fidelity levels.

The activities over the first year of the project have culminated in the construction of two primary probabilistic nozzle models, a low- and a medium-fidelity model, both of which have a variety of parameters that can be varied to obtain various versions of *fidelity* (for example, convergence tolerances, mesh resolution, etc.). Both models estimate primary quantities of interest for the nozzle design problem such as weight (or volume), thrust, mechanical stress, and thermal stress given a set of deterministic or stochastic inputs such as the shape of the nozzle, the material properties, the inlet conditions of the nozzle, and the atmospheric conditions outside of the nozzle. The nozzle geometry shared between all fidelity levels is parameterized using a B-spline which can contain an arbitrary number of knots. The thickness of both layers of the nozzle wall is parameterized using piecewise-linear functions comprised of an arbitrary number of segments.

The low-fidelity model runs in approximately five seconds and solves the quasi-1D Navier-Stokes equation to obtain flow characteristics within the nozzle. A conjugate heat transfer problem is solved using a thermal resistance representation of the nozzle walls and heat transfer to the surrounding aircraft structure and atmosphere, thereby obtaining temperatures in the nozzle wall. Lastly, hoop stresses and thermal stresses are calculated in the nozzle wall. The model is quick, efficient, and captures the salient trends of the nozzle design problem. Agreement with the medium-fidelity model is fairly good, with the exception of the stresses which are grossly underestimated by the low-fidelity model due to complex Poisson and thermal effects taking place in the nozzle structure. For these reasons, we are planning an update to the low-fidelity model that will use the simplified full-structure finite element model, with the medium-fidelity model incorporating a more sophisticated structural model that the currently implemented one.

The medium-fidelity model runs in approximately thirty seconds and uses the SU2 software to solve the Euler and RANS equations and obtain flow characteristics within the nozzle. Robustness issues in the computational fluid dynamics calculation have been largely addressed, and adequate convergence is obtained for nearly all nozzle geometries. An adaptive mesh refinement capability has also been developed and is available as an additional level of fidelity in our model analysis hierarchy. A one-and-a-half-way coupled heat transfer and structural finite element problem is then solved using the AERO-S software with nozzle wall temperatures transferred from the CFD solution. A variety of mesh resolution levels are also available in the thermal-structural analysis for multifidelity studies.

SCALABLE TECHNIQUES FOR INVERSE UQ (TA1, SNL): Over the first year of the project, we completed our work on observation-consistent Bayesian inference. A paper was submitted to SISC for review. We also developed a procedure based on our inference approach for choosing optimal experimental designs. This work was documented in Sandia's Center for Computing Research (CCR) Summer Proceedings. These activities, along with the others reported over the first year of this project, have enhanced our ability to tackle the challenges of performing inverse UQ on large scale models, which is particularly relevant for estimating model-form uncertainty (see Tasks 7 and 8).

SCALABLE INVERSION USING DIMENSION REDUCTION (TA1, CSM): This task, which is currently being started, will aim to discover and exploit active subspaces in the misfit function defining the inverse problem. Given the low-dimensional, data-informed active subspace, we will efficiently compute posterior statistics of model parameters and unmeasurable quantities of interest using either Markov chain Monte Carlo or appropriate Gaussian quadrature on the coordinates of the active subspace.

MACHINE LEARNING FOR EMBEDDED MODEL DISCREPANCY (TA2, UM): In work presented earlier, statistical inference was used to extract model discrepancies from data. This was repeated over a wide range of nozzle configurations and the spatial form of the discrepancies was extracted for each configuration. Following the field inversion over multiple nozzle computations, the functional form of the discrepancy (and its distribution) was reconstructed using Machine Learning techniques. Specifically, we used AdaBoost (from

the openly available SKLearn library) and Neural Networks. Based on different training sets, we built ensembles of machine learning maps which were embedded in predictive computations to obtain uncertainty envelopes. Figure 5 shows the output of the machine learning (AdaBoost) on a sample nozzle data set.
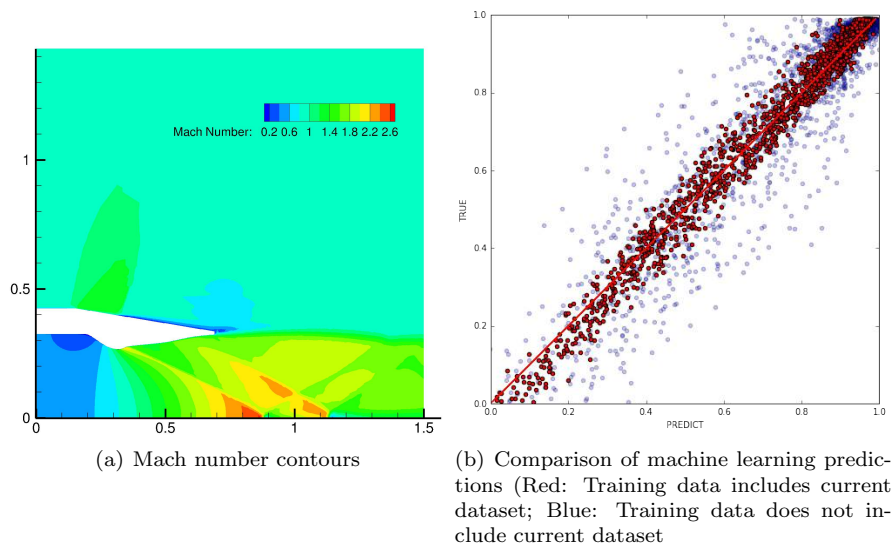

(a) Mach number contours

(b) Comparison of machine learning predictions (Red: Training data includes current dataset; Blue: Training data does not include current dataset

Figure 5. Machine learning prediction of model discrepancy for a sample nozzle data set

Reduced-order models for DUU (TA3, SU):   Although this effort is formally initiating at the time of writing, some preliminary work was undertaken to consider theoretical implications of parametric variations involving changes to the model geometry on the reduced-order modeling framework, and to upgrade the reduced-order modeling capability of the AERO-S code to support this class of parametric ROMs. Previous applications of the reduced-order modeling framework that will be utilized in this task initially focused on accurately reproducing high-fidelity simulation results in near-real time without any parametric variation, and more recently on parametric ROMs for parameters associated with source terms that did not involve any changes in the model geometry. In the case of geometric parameterization, it was not immediately obvious that solution snapshots collected by sampling the parameter space could be combined and compressed to generate a single global basis. It has now been established that the framework naturally extends to parametric ROMs in which the model geometry is parameterized provided that certain aspects are treated consistently, and a strategy for training the ROM has been developed and implemented in AERO-S. Using this new capability, a ROM was built for finite element structural model of a rocket nozzle with two geometric parameters representing a scaling of the length and radius. The dimension of the ROM was varied to study its effect on accuracy, a reassuring degree of convergence was observed and modest errors were measured for a relatively wide range of parameter variations (see Figure 6). Although these results were obtained without hyper reduction and thus no appreciable speed-ups were attained, preliminary results suggest that a similar accuracy accompanied by substantial speed-ups around two orders of magnitude will be demonstrated in the very near future by incorporating hyper reduction.

Integration of emerging capabilities within DUU framework (TA3, SNL):   We are designing our trust-region model management approaches with dimension reduction in mind. In particular, we are developing infrastructure to form active subspace approximations for a set of designed QoI within a trust region over the design variables, and then subsequently to adapt this trust region as part of our surrogate-based optimization processes. Later in this document we describe an initial demonstration of active subspace-based design under uncertainty, prior to the inclusion of trust region model management.

III.   Major Milestones Completed and Upcoming

This section describes the major milestones that the project has completed after 12 month of execution. These milestones are part of a planned effort that spans a total of 3 years and that culminates with significant
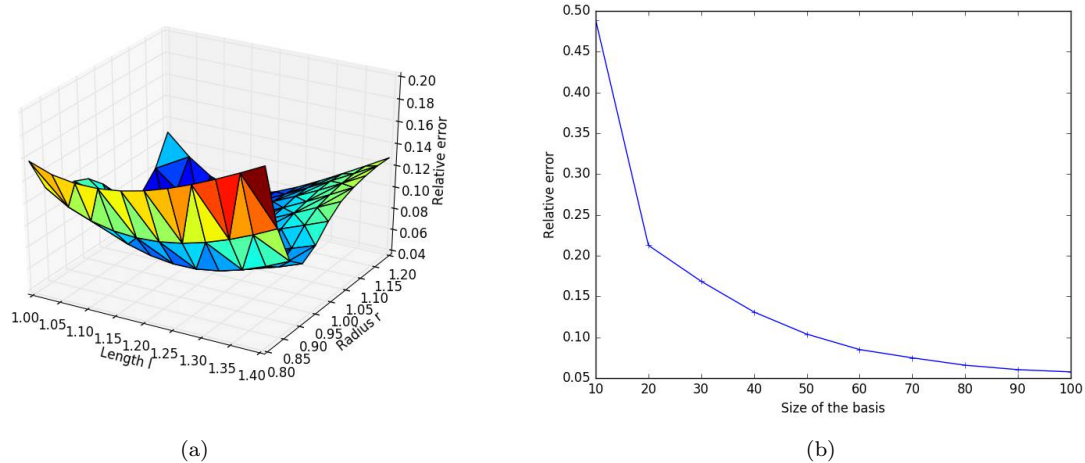
American Institute of Aeronautics and Astronautics

Figure 6. Relative error of ROM without hyper reduction shown (a) for fixed dimension while varying the parameters, and (b) for fixed parameters while varying the dimension of the reduced-order basis.

demonstrations in all Thurst Areas, as the apply to the aero-thermal-structural nozzle analysis and design problem.

Milestone SF1: Forward UQ integration of scalable methods including dimensionality reduction We have completed theory and computational implementations—complemented with numerical experiments—for exploiting low-dimensional structure in a given high-dimensional input/output map to propagate uncertainties from inputs to outputs. The preprint by Constantine, Eftekhari, and Ward (2016) develops an efficient procedure for *ridge approximation*, which approximates a function of several variables by some function of a few linear combinations of the variables. The ridge function is fitted with an alternating minimization heuristic over the Grassmann manifold of subspaces. Given a data set of input/output pairs $\{(\mathbf{x}_i, f(\mathbf{x}_i))\}$ and an initial subspace $\mathbf{U} \in \mathbb{R}^{m \times n}$,

1. fit a model $p(\mathbf{y})$ (e.g., a sparse polynomial chaos expansion) to the pairs $\{(\mathbf{U}^T\mathbf{x}_i, f(\mathbf{x}_i))\}$, which occurs on the lower n-dimensional active subspace;

2. solve the Grassmann-constrained minimization

$$\mathbf{V}_* = \underset{\mathbf{V}}{\text{minimize}} \quad \sum_i \left(f(\mathbf{x}_i) - p(\mathbf{V}^T\mathbf{x}_i) \right)^2,$$
$$\text{subject to} \quad \mathbf{V} \in \mathbb{G}(n, m),$$

where $\mathbb{G}(n, m)$ is the Grassmann manifold of dimension n in $\mathbb{R}^m$;

3. set $\mathbf{U} = \mathbf{V}_*$ and repeat.

We propose the model's active subspace as an initial subspace for the minimization. The ridge approximation can be used as an efficient response surface model for uncertainty propagation.

In addition to these fundamental developments, We deployed a combination of dimension reduction and scalable forward UQ methods to our nozzle application. Specifically we combined the use of sparse PCE and active subspaces to build approximations of the nozzle model over important combinations of 15 uncertain model parameters. The goal of building these approximations is to efficiently evaluate the objectives and constraints of the robust design under uncertainty of the nozzle. The accuracy of the approximations is related to the requirements of the optimization algorithm, specifically the final accuracy requested and the accuracy of each optimization step. The accuracy required at each step typically can be set quite low at the beginning of the optimization and increased as the optimization approaches the optimal design. Consequently we need a means to choose the accuracy of the surrogate depending on the situation. Figure 7 shows how the accuracy of a PCE approximation, of volume and thrust, can be tailored by tuning the number of model evaluations and the dimensionality of the subspace. Choosing the dimensionality of the volume based active

American Institute of Aeronautics and Astronautics

subspace to $d = 4$ provides the most efficient means of achieving an accuracy of $O(10^{-4})$, however if an accuracy of $O(10^{-6})$ is required the best subspace dimension is $d = 5$. Similar results hold for thrust and (although not shown) for the other QOI as well.
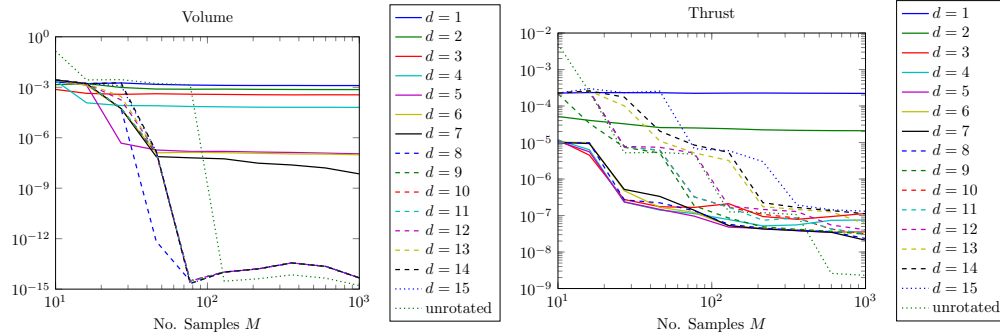

Figure 7. The relative ℓ^2-error in PCE approximatons built over active subspaces, derived from the nozzle, for varying sample sizes and active subspace dimensionality.

The PCE approximations described above can be used to compute moments of the QoI, needed by DUU, by random sampling the surrogate. However in some cases it may be more prudent to compute moments directly with cubature rules, without building a PCE approximation. Such an approach has been challenging in the past because the density on the active subspace is not a standard distribution, like a uniform distribution, even if the distribution of the full dimensional model-parameters is a standard distribution. Moreover for bounded variables, the domain of the active subspace is defined multivariate polygon (zonotope) which is unlikely to be a hypercube, and so most cubature techniques are inappropriate. In Figure 8 we apply our minimal cubature rule to compute moments of high-dimensional functions active subspaces of varying dimension. It can be seen that we can indeed compute moments using active subspaces and again the accuracy obtained can be tuned to the requirements of the DUU by changing the order of the cubature rule (number of samples) and the dimensionality of the subspace.
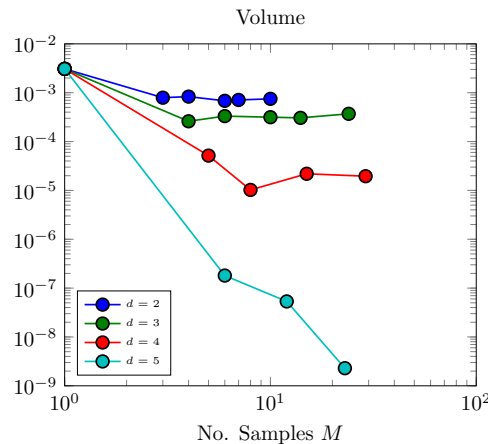

Figure 8. The relative error in the expected volume computed using active subspaces, derived from the nozzle, for varying sample sizes (degree of qudrature rule) and active subspace dimensionality.

It is worth noting that in meeting this milestone we were able to identify strong connections with the minimal cubature rules used here, DUU and the multi-fidelity quadrature being developed by Task 9. Specifically robust design under uncertainty requires estimation of moments (mean and variance) to compute a design that is insensitive to variation due to uncertainty. To date in this project we have developed multi-fidelity quadrature in parallel to minimal cubature as separate means, with differing strengths, with the intention of using these approaches to compute moments for our DUU for the nozzle. However in certain circumstances it has become apparent that these approaches could be combined to produce an even more powerful approach for computing moments. This will be the topic of future work.

Milestone MI1 (month 12): Inference methods for embedded model discrepancy The objective of MI1 is to develop and demonstrate approaches to address model form uncertainties introduced by turbulence closure. Two approaches were pursued:

A physics-based approach (PBA) which relies on realizability properties, symmetries and bounding principles to construct Reynolds stresses envelopes; underlying assumptions in classical models are eliminated and no data is required. The resulting enveloping models are embedded in simulation codes and lead to envelopes on predictions that provide *prudent* estimates of the potential impact of model-form uncertainties. A set of 6 computations are required to produce the envelopes, although in many applications only two are found to be sufficient.

A data-based approach (DBA) which relies on field inversion to extract the spatial form of the discrepancy and machine learning to convert this into a functional representation. The posterior from the machine learning is then sampled to obtain realizations of the prediction, an ensemble of which is used to generate bounds of predictive quantities.

MI1 was completed by extending both approaches to RANS computations of current axisymmetric nozzle application. As a demonstrative example, Figure 9 shows predictive computations augmented by machine learning models (trained on inverse solutions of benchmark data (symbols)). The base model (which is the eddy viscosity-based Spalart Allmaras model) results are shown as blue lines. Using high-fidelity data from a full Reynolds stress model, field inference and machine learning was used to reconstruct the form of the model discrepancy, introduced as a correction to the production of eddy viscosity. Predictive computations were then augmented with this machine learned ensembles and the envelope of the predictions is shown in the shaded red region.

In other work, analogous envelopes were shown for the same nozzle using the PBA.

Note that , if the PBA envelopes do *include* the truth (experimental observations or high-fidelity data), then one-point closures are sufficient to predict the flow and DBA can lead to improvements of the baseline predictions. In the future, we will combine both approaches by using the PBA envelopes to define the prior for DBA inference strategy.

Both of these techniques for model-form uncertainty quantification will be embedded into the predictive part of the DUU cycle.
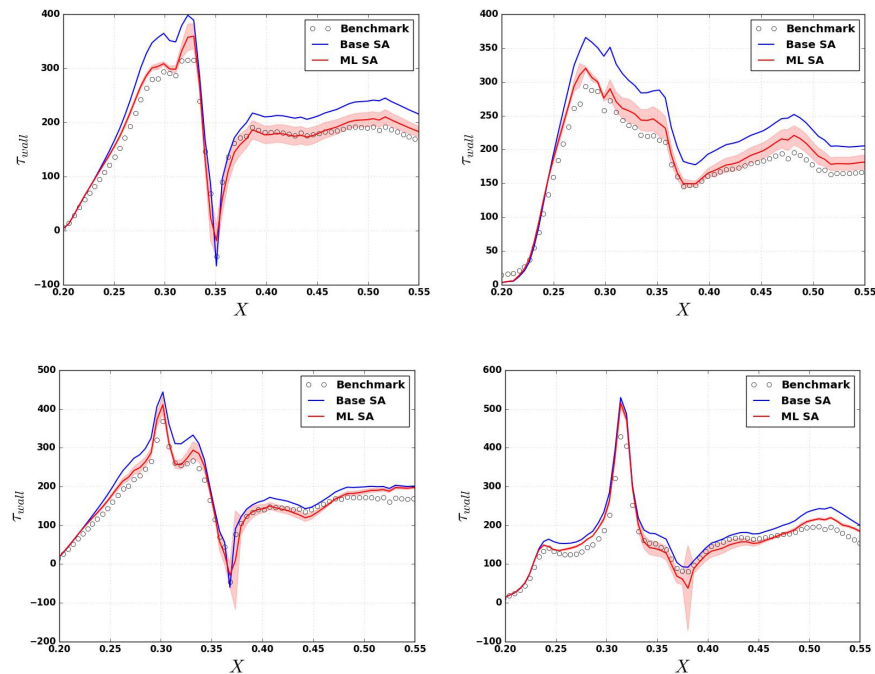

Figure 9. Prediction (with data-based embedded model-form UQ) of skin friction over the nozzle wall for four different nozzle flow conditions

American Institute of Aeronautics and Astronautics

Milestone DUU1: Mathematical foundations for DUU framework We are moving towards a shared formulation for robust optimization of the nozzle design that will minimize the variance of thrust subject to constraints on expected values of other metrics:

$$
\begin{aligned}
\underset{\mathbf{x}}{\text{minimize}} \quad & \mathbb{V}[\ T\ ] \\
\text{subject to} \quad & \mathbb{E}[\ W\ ] \leq \overline{W} \\
& \mathbb{E}[\ T\ ] \geq \overline{T} \\
& \mathbb{E}[\ ||T_w||\ ] \leq \overline{T}_w \\
& \mathbb{E}[\ ||\sigma||\ ] \leq \overline{\sigma}
\end{aligned}
\tag{1}
$$

The current problem formulation is described by a total of 36 parameters – 29 are design and 7 are uncertain parameters. The design parameters consist of 21 B-spline coefficients that describe the shape of the nozzle and 8 parameters controlling the nozzle wall thickness. The 7 uncertain parameters relate to the flight conditions (inlet stagnation temperature and pressure, and freestream temperature and pressure) and material properties (elastic modulus, thermal conductivity, and thermal diffusivity) of the nozzle.

In the following sections, we describe three investigations that target milestone DUU1. The first involves the extension of model management frameworks to handle multilevel-multifidelity hierarchies such as those present with the low and medium fidelity nozzle models. The second describes a trust region model management approach based on polynomial chaos expansions of varying statistical fidelity. And the third describes the exploitation of low-dimensional active subspaces within DUU. This is a forward looking effort book in terms of integration of emerging capabilities and in terms of the use of chance constraints.

MULTILEVEL-MULTIFIDELITY DUU: We have continued to invest in multilevel-multifidelity methods for forward UQ. We can leverage and extend this same model management infrastructure within multilevel-multifidelity approaches for DUU. In particular, we employ a generalized model management framework that flexibly accommodates multiple simulation fidelities each with multiple discretization levels (deterministic level), and augments with the additional dimension of multiple UQ algorithm forms and multiple levels of statistical approximation (stochastic level). To exploit this special structure, we focus this effort on optimization strategies that effectively manage hierarchical modeling.

We target two algorithmic components for developing provably-convergent capabilities for multilevel-multifidelity optimization, that have been investigated and reported previously:

- Multigrid optimization (MG/Opt) performs optimization with a hierarchy of discretization levels, effectively applying a multigrid V cycle to a hierarchy of optimization solves, including prolongation and restriction operations when needed to manage discretization-dependent design parameterizations.

- Trust region model management (TRMM) can handle the more general case of optimizing a hierarchy of model forms, lacking any special structure other than a presumed correlation between accuracy and cost. The actual numerical optimization is only performed at the lowest fidelity (as opposed to every level in the case of MG/Opt). In the case of more than two model forms, multiple trust regions manage model corrections over the full range of the model hierarchy.

We first explored multigrid optimization (MG/Opt) solvers on model problems, and demonstrated their efficacy with gradient-based optimizers. We then explored two prototype algorithms for multilevel-multifidelity optimization that leverage MG/Opt and TRMM. Formulations and initial computational experiments using this hybrid multilevel-multifidelity approach have been accepted for publication in an accompanying paper at AIAA SciTech 2017. We have also focused on promoting our algorithmic prototypes into Dakota, adding multigrid optimization (MG/Opt) for model discretization hierarchies, generalizing trust-region model management (TRMM) for model form hierarchies of arbitrary depth, and hybrid MG/Opt + TRMM for multilevel-multifidelity hierarchies.

Multifidelity DUU has been performed to solve the robust optimization problem in Eq. 1, initially using a bi-fidelity formulation. The high-fidelity statistics were obtained by forming a level 1 sparse grid PCE over an Euler CFD simulation with a medium discretization. The low-fidelity statistics used a level 1 sparse grid PCE over the same Euler CFD simulation with a coarse discretization. The optimization used the TRMM approach in Dakota with first-order consistency enforcement, providing provable convergence to the optimum of the HF statistical model. Design gradients were evaluated numerically by finite differencing

at the simulation level, enabling analytic gradients of moments at the design level. The TRMM algorithm performed 5 iterations and all were accepted using a filter method that accepts non-dominated objective-constraint violation pairs. Figure 10 compares the geometry and pressure contours of the baseline nozzle to the optimized nozzle. Tuning of this process to improve consistency among the finite-difference step size and CFD solver tolerances is in progress, although the deployment of adjoint sensitivities in future model realizations will eliminate this need.

(a) Baseline
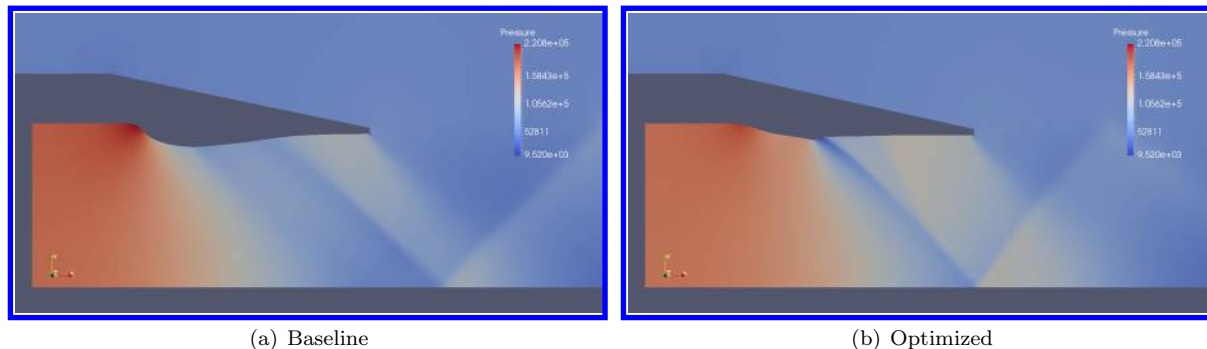
(b) Optimized

Figure 10. Nozzle geometry and pressure contours resulting from the robust DUU problem, a) baseline configuration b) optimized nozzle geometry.

TRMM WITH MULTIFIDELITY PCE SURROGATES: In addition, a trust region model management (TRMM) framework was recently implemented successfully to obtain an optimal nozzle shape for a reliable DUU problem using the low-fidelity model. The problem minimized weight and required thrust to be greater than 25,000 Newtons 99% of the time. Figure 11 compares the optimal nozzle design with the baseline design, deterministic optimal design, and a previous control variate optimal design. The CDF of thrust for the optimal nozzle design is also shown. Note that the chance constraint for thrust has been satisfied well. From Figure 11 it can be seen that the TRMM optimization finds a design very similar to the deterministic optimization, but that is more reliable.
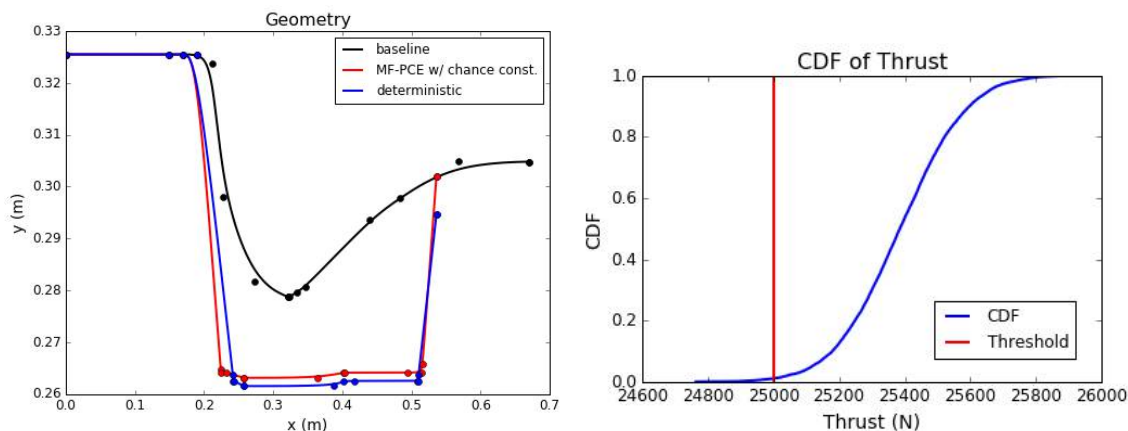

Figure 11. Comparison of optimal designs for a TRMM optimization using PCEs with a chance constraint on thrust, a similar deterministic optimization, and the baseline nozzle design. The CDF for the TRMM optimization is shown on the right with the accompanying thrust threshold of 25000 Newtons.

Although the dimension of the problem was low (21 deterministic design variables relating to the B-spline coefficients of the nozzle wall shape and 4 random variables relating to the inlet and atmospheric pressures and temperatures), the approach is easily scalable to higher-dimensions and paves the way for the solution of such problems within a reliable design setting. The use of the TRMM framework coupled with the type of higher- and lower-fidelity surrogates used for the thrust function enabled a large savings in computational cost. The optimization converged in 12 iterations with approximately 7,000 total function calls to the truth function for thrust (including finite difference derivatives). Comparatively, evaluating the thrust function once using simple Monte Carlo requires approximately 10,000 function calls for similar accuracy.

American Institute of Aeronautics and Astronautics

In the optimization above, two different fidelity surrogates of the truth model for thrust were constructed. A higher-fidelity surrogate of thrust was built using a polynomial chaos expansion (PCE) and a lower-fidelity surrogate of thrust was built first using an anchored decomposition of the truth thrust function, followed by a PCE of the stochastic components of the decomposition. The anchored decomposition effectively separates the multivariate thrust function into the sum of several functions of a smaller number of variables, thereby making the use of a PCE tractable. The anchor was chosen to be the center of the trust region (for the deterministic variables) and the expectation of the random variables.

DUU WITH ACTIVE SUBSPACES: We have also investigated the use of other multifidelity approaches for DUU including the control variate method and a technique based on anchored decomposition embedded within a TRMM framework. Several minimum weight optimization problems were formed and solved using the control variate method applied to two sub-fidelities within the low-fidelity model. The primary constraint of interest was a nonlinear thrust constraint which took several forms including the expectation or the expectation minus three standard deviations. Optimization results led to a minimum weight design, but were either too relaxed or conservative in terms of robustness with respect to generation of adequate thrust.

For more robust design of a minimum weight nozzle with a thrust constraint, a new approach was taken which focused on implementing chance constraints of the the form $P[thrust > 25,000 \text{ N}] \geq 0.99$. The approach again used two sub-fidelities within the low-fidelity model, but this time implemented polynomial chaos expansions for calculation of the constraints and an anchored decomposition of the original function for the lower-fidelity. By correcting the lower-fidelity model and embedding both fidelities within a TRMM framework, the optimization quickly converged to a minimum weight design that robustly generates thrust over the range of provided uncertainties.

For exploiting active subspaces within design, we have tested the DUU quantities of interest in the medium fidelity model for low-dimensional structure with a gradient-free linear model technique related to active subspaces. We independently tested (i) the quantities of interest as a function of the uncertain variables at several design points and (ii) statistics of the quantities of interest as functions of the design variables. All quantities currently considered admit low-dimensional relationships with the uncertain variables; these relationships can be exploited when computing statistics inside the DUU method. Figure 12 shows the results of this test. The response plots show samples of the responses at different designs as a function of one linear combination of the uncertain inputs; the weights of these combinations are in the corresponding weight plots. The labels on the horizontal axis in the weight plots correspond to the uncertain parameters. Given the relationship revealed in the response plots, the weights also measure sensitivity of response to each input; a zero weight indicates that the response does not depend on the corresponding parameter. This dimension reduction leads to increased efficiency in estimating mean, variance, and probability of failure—statistics present in DUU constraints.

Moreover, some statistics arising in the constraints admitted low-dimensional, monotonic relationships with the design variables through their one-dimensional active subspace—again, revealed with a gradient-free linear model test. Figure 13 shows the results of this test applied to the *volume* objective and probabilities of success / failure. The weight plots correspond to 29 design varibles that parameterize the nozzle shape and wall thicknesses. These revealed relationships imply that the expensive-to-evaluate statistical constraint can be replaced by a few linear inequality constraints on the design variables to create a surrogate DUU problem whose solution estimates the desired DUU solution. At the least, the solution to the surrogate DUU problem provides a superior starting point for local, trust-region based DUU approaches.

We can use the relationships in Figures 12 and 13 to construct a linear inequality constrained linear program that is a surrogate for the following DUU problem:

$$
\begin{aligned}
\underset{\mathbf{x}}{\text{minimize}} \quad & \text{Volume}(\mathbf{x}) \\
\text{subject to} \quad & \mathbb{P}[\text{ Thrust}(\mathbf{x}, \xi) > 42\,\text{kN}] > 0.99 \\
& \mathbb{P}[\text{ Stress}_{\text{mech}}(\mathbf{x}, \xi) > 322\,\text{GPa}] < 0.01 \\
& \mathbb{P}[\text{ Stress}_{\text{therm}}(\mathbf{x}, \xi) > 240\,\text{GPa}] < 0.01 \\
& \mathbf{x} \in \mathcal{D}, \ \ \xi \in \mathcal{U}
\end{aligned}
\tag{2}
$$

Solving the surrogate linear program produces a point in the design space. We can exploit the linear relationships from Figure 12 to perform full forward UQ propagation with only 10 runs and a linear surrogate model. Figure 14 shows the estimated cumulative distribution functions for the nozzle model quantities of
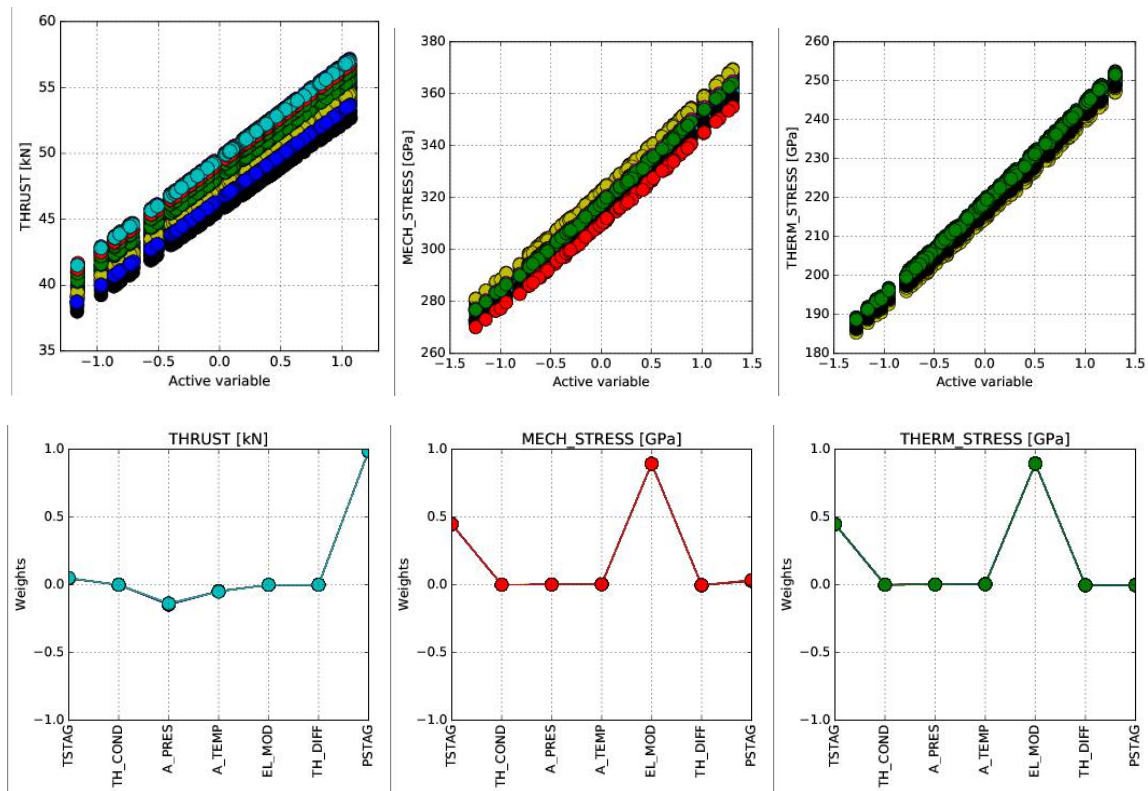
Figure 12. The nozzle model's output quantities that define DUU constraints show evidence of univariate linear structure as a function of the uncertain variables—consistently across several design. This structure was revealed through a gradient-free linear model graphical test related to discovering active subspaces. The structure can be exploited to compute statistics efficiently over design variables in a DUU framework.
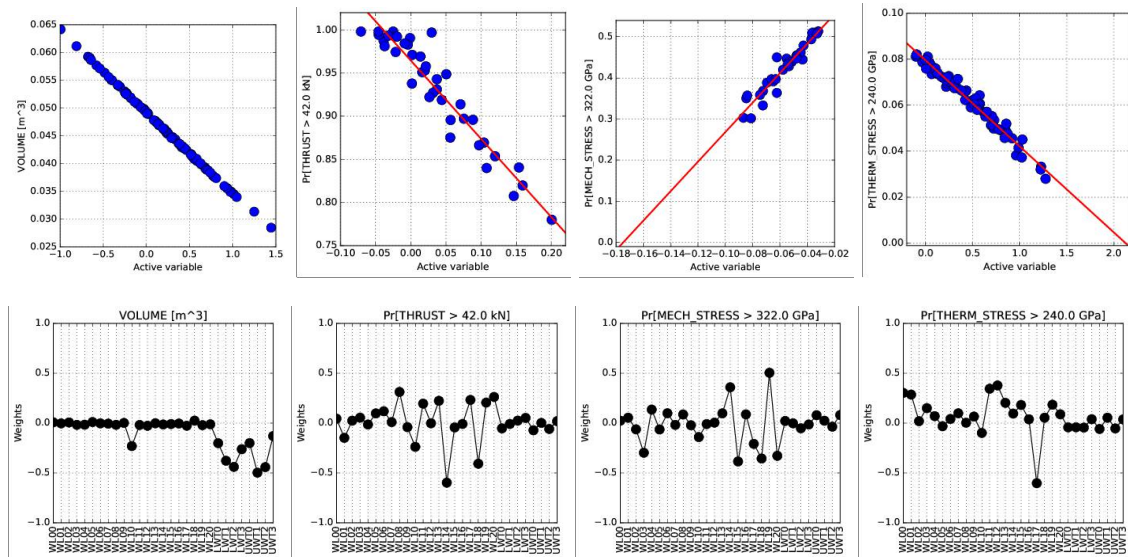

Figure 13. The DUU constraints show evidence of univariate linear structure as a function of the design variables, though the relationship is not a clear as in the case of the uncertain variables. This structure was revealed through a gradient-free linear model graphical test related to discovering active subspaces. The structure can be exploited to create a surrogate DUU problem whose solution estimates the full DUU solution.

interest for both the baseline model and the design from the DUU surrogate linear program. The green vertical lines show the threshold values. The surrogate DUU produced a design with more favorable statistical performance characteristics than the baseline across all quantities of interest. Moreover, the *volume* objective

decreased from 0.034 m^3 in the baseline design to 0.008 m^3 in the surrogate DUU design.
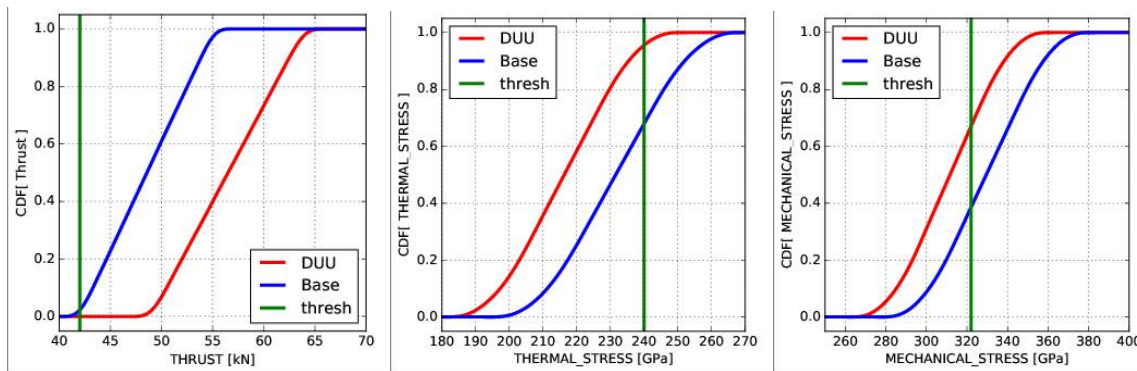

Figure 14. Estimated cumulative distribution functions for nozzle model quantities of interest. The blue lines show the baseline design, and the red lines show the design from the solution to the surrogate DUU problem. In all cases, the surrogate DUU solution has better statistical performance characteristics than the baseline design.

The value of the linear program surrogate DUU is two-fold. First, the design from the DUU surrogate solution can be used as an excellent starting design for a trust-region-based DUU method. Moreover, the gradient-free test for one-dimensional active subspaces revealed exploitable low-dimensional structure that can increase the efficiency of the full DUU problem, e.g., to estimate statistics over uncertain variables within the DUU method.

Milestone APP1: Computational model for low- and medium-fidelity applications Our model development processes have culminated in the construction of a multifidelity modeling tool MULTI-F for the aero-thermal-structural analysis of an entire nozzle that shares a consistent parameterization and QoI definitions across fidelity and sub-fidelity levels. In addition, MULTI-F is easily integrated with Sandia's DAKOTA software, easily accessing a powerful array of uncertainty quantification techniques and optimization methods. Both low- and medium-fidelity models have been successfully implemented in MULTI-F and run on computer clusters for a variety of design and analysis problems, and the code structure allows easy addition of further fidelity levels. MULTI-F is written in Python, shared openly and used widely by the team through GitHub. MULTI-F has now been deployed across our entire team and, as evidenced in this report, is being used to demonstrate algorithmic advances across all Thrust Areas. It is a focus of our program to ensure that methodology and algorithmic improvements are demonstrated in the full aero-thermal-structural nozzle problem and not just in canonical problems that hide some of the realities observed when dealing with the full problem.

The low-fidelity model incorporates the solution of the quasi-1D Navier Stokes equation with a conjugate heat transfer problem and also contains simplified stress calculations. The medium-fidelity model uses SU2 to solve the Euler/RANS equations and AERO-S to solve one-and-a-half-way coupled heat transfer and structural problems.

The heat transfer problem is discretized with 8-node hexahedra elements. Both the inner insulating layer and the outer load-bearing structural layer are modeled, with a prescribed temperature boundary condition (obtained from the SU2 solution) on the inside surface of the inner layer and a convection boundary condition on the outside surface of the outer layer. The structural problem is discretized with 3-node shell elements. Only the outer layer is modeled in this case. A fixed displacement boundary condition is used at the inlet and the external load is generated by a non-uniform surface pressure and temperature (obtained from the SU2 and AERO-S solutions, respectively). Typical solutions obtained from the thermal and structural analyses are shown in Figure 15. The thermal and structural meshes are fully parameterized and automatically generated using the transfinite algorithm of the Gmsh library. Furthermore, the structural mesh generator allows both longitudinal and circumferential stiffeners to be included, although this feature is not yet supported by the python user interface. Although many improvements to the medium-fidelity model have been made during the quarter, a few details remain to be included such as allowance for material anisotropy, variation in the ambient temperature along the length of the nozzle, and modeling of the external support structure.
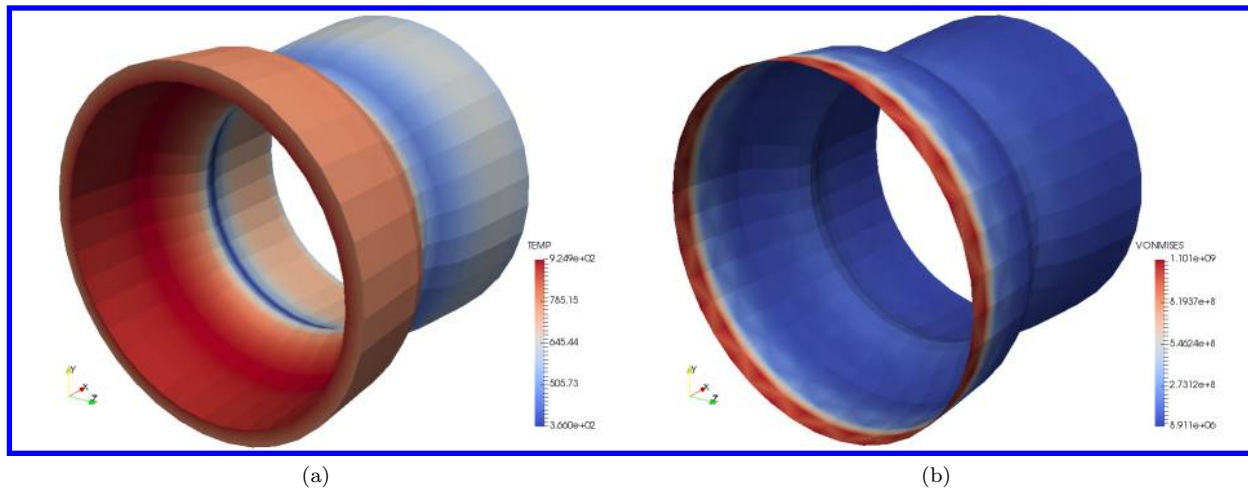
American Institute of Aeronautics and Astronautics

Figure 15. Representative medium-fidelity finite element model discretization and solution (a) for thermal analysis with contours of temperature, and (b) for structural analysis with contours of Von Mises stress.

IV.   Conclusion

This paper has presented a quick overview of the SEQUOIA project, one of the projects that participate in the DARPA EQUiPS program. Progress in efforts in various thrust areas have been described, and milestone accomplishments thus far have been detailed. Work will continue over the next two years to realize the vision of design under uncertainty for nozzle systems, while considering a significant number of aleatoric uncertain parameters, and the quantification of important model-form uncertainties present in our simulations.

V.   Acknowledgments

This article has been cited by:

1. Vishal Srivastava, Karthik Duraisamy. Aerodynamic Design of Aircraft Engine Nozzles with Consideration of Model Form Uncertainties . [Citation] [PDF] [PDF Plus]

2. Michael S. Eldred, Gianluca Geraci, Alex Gorodetsky, John Jakeman. Multilevel-Multifidelity Approaches for Forward UQ in the DARPA SEQUOIA project . [Citation] [PDF] [PDF Plus]

3. Zach Grey, Paul Constantine. Characterizing Subspaces of Engineering Shapes using Differential Geometry . [Citation] [PDF] [PDF Plus]

4. Martijn Roelofs, Roelof Vos. Uncertainty-Based Design Optimization and Technology Evaluation: A Review . [Citation] [PDF] [PDF Plus]

5. Gianluca Geraci, Michael S. Eldred, Gianluca Iaccarino. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications . [Citation] [PDF] [PDF Plus]