53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference
20th AI
23 - 26 April 2012, Honolulu, Hawaii

AIAA 2012-1852

Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation

L. W. T. Ng*

Massachusetts Institute of Technology, Cambridge, MA 02139

M. S. Eldred†

Sandia National Laboratories,‡Albuquerque, NM 87185

This paper explores the extension of multifidelity modeling concepts to the field of uncertainty quantification. Motivated by local correction functions that enable the provable convergence of a multifidelity optimization approach to an optimal high-fidelity point solution, we extend these ideas to global discrepancy modeling within a stochastic domain and seek convergence of a multifidelity uncertainty quantification process to globally integrated high-fidelity statistics. For constructing stochastic models of both the low fidelity model and the model discrepancy, we employ stochastic expansion methods (nonintrusive polynomial chaos and stochastic collocation) computed from sparse grids, where we seek to employ a coarsely resolved grid for the discrepancy in combination with a more finely resolved grid for the low fidelity model. The resolutions of these grids may be statically defined or determined through uniform and adaptive refinement processes. Adaptive refinement is particularly attractive, as it has the ability to preferentially target stochastic regions where the model discrepancy becomes more complex; i.e., where the predictive capabilities of the low-fidelity model start to break down and greater reliance on the high fidelity model (via the discrepancy) is necessary. These adaptive refinement processes can either be performed separately for the different sparse grids or within a unified multifidelity algorithm. In particular, we propose an adaptive greedy multifidelity approach in which we extend the generalized sparse grid concept to consider candidate index set refinements drawn from multiple sparse grids. We demonstrate that the multifidelity UQ process converges more rapidly than a single-fidelity UQ in cases where the variance of the discrepancy is reduced relative to the variance of the high fidelity model (resulting in reductions in initial stochastic error) and/or where the spectrum of the expansion coefficients of the model discrepancy decays more rapidly than that of the high-fidelity model (resulting in accelerated convergence rates).

I. Introduction

In engineering design, it is desirable to quantify the effect of input uncertainties on the system responses. If sufficient information is available to characterize the input uncertainties, then probabilistic methods can be used to propagate the uncertainties through the system. The results are statistics of the system responses that can be used in the engineering decision process. For example, the variance of a performance metric and the probability of exceeding a failure threshold can be used to assess the robustness and the reliability of the design, respectively. The prevalence of uncertainties in practical engineering problems suggests the

*Graduate Research Assistant, Department of Aeronautics and Astronautics.

†Distinguished Member of Technical Staff, Optimization and Uncertainty Estimation Department, AIAA Associate Fellow.

‡Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

need for efficient uncertainty propagation methods that can minimize the number of evaluations of expensive high-fidelity models.

The issue of model form is prevalent in the field of uncertainty quantification. The computational simulation of a particular physical phenomenon often has multiple discrete model selection possibilities. Here, we will broadly characterize this issue into two classes: a model hierarchy and a model ensemble. In the former case of a hierarchy, a clear preference structure exists among the models such that "high-fidelity" and "low-fidelity" judgements are readily assigned. Here, the goal is to manage the trade-off between accuracy and expense among the different model fidelities in order to achieve high quality statistical results at lower cost. In the latter case, a clear preference structure is lacking and there is additional uncertainty created by the lack of a clear "truth model." In this case, the goal becomes one of management and propagation of this model form uncertainty, or in the presence of experimental data for performing inference, model form selection.

In this paper, we address the model hierarchy case where the system responses can be obtained accurately by evaluating an expensive high-fidelity model or less accurately by evaluating a cheap low-fidelity model. The low-fidelity model may be based on simplified physics, coarser discretization of the high-fidelity model, or reduced-order models. We investigate a multifidelity approach to compute the statistics of the high-fidelity model without the expense of relying exclusively on high-fidelity model evaluations. Such multifidelity approaches have been developed for the optimization of expensive high-fidelity models. In the multifidelity trust-region model-management approach, the optimization is performed on a corrected low-fidelity model. The correction function can be additive, multiplicative, or a combination of the two and is updated occasionally by high-fidelity model evaluations.[1,2] First-order or second-order polynomials are used to enforce local first-order or second-order consistency, respectively, at high-fidelity model evaluation points.[1,2] Other variations have employed global correction functions (typically enforcing zeroth-order consistency) based on the interpolation of the discrepancy between the high-fidelity model evaluations and the low-fidelity model evaluations.[3–6] The idea is that a surrogate based on a physics-based low-fidelity model and an interpolant of the discrepancy may provide a more cost-effective approximation of the high-fidelity model than a surrogate based only on interpolating the high-fidelity model. We carry this idea over to uncertainty propagation and construct a surrogate of the high-fidelity model based on the low-fidelity model and a correction function using global polynomials in terms of the stochastic parameters. Two types of surrogates are considered, non-intrusive polynomial chaos using orthogonal polynomials and stochastic collocation using interpolation polynomials, and we are interested in discrepancy models that can reproduce the high-fidelity results at each of the high-fidelity collocation points to zeroth- and first-order (i.e., that exactly interpolate discrepancy values and first derivatives).

The polynomial chaos method expands the system response as a truncated series of polynomials that are orthogonal with respect to the probability density functions of the stochastic parameters[7,8] and exponential convergence in integrated statistical quantities (e.g., mean, variance) can be achieved for smooth functions with finite variance. The chaos coefficients are obtained by projecting the system onto each basis. In the non-intrusive case, the projection can be approximated by a multi-dimensional numerical integration and is sometimes known as the pseudo-spectral method. Stochastic collocation is a related stochastic expansion method which constructs multidimensional interpolation polynomials over the system responses evaluated at a structured set of collocation points.[9,10] If the collocation points are selected to be the Gaussian quadrature nodes associated with the same orthogonal polynomials as the polynomial chaos expansion, then the same exponential convergence properties can be observed. We apply the stochastic expansion methods to the low-fidelity model and to the discrepancy between the high-fidelity model and the low-fidelity model. The two expansions are then combined to create a surrogate stochastic expansion of the high-fidelity model from which the desired statistics are obtained. If the low-fidelity model is sufficiently predictive, a lower-order expansion of the discrepancy can be used, reducing the number of high-fidelity model evaluations necessary to obtain the response statistics at the desired accuracy.

In the following, we first review the stochastic expansion methods in Section II and their multi-dimensional construction via sparse grids. In Section III, we present the extension to the multifidelity case and provide an algorithmic framework. Finally, we demonstrate our approach in Section IV with computational experiments and provide concluding remarks in Section V.

Table 1: Some standard continuous probability distributions and their corresponding Askey polynomial bases. $B(\alpha,\beta)$ is the Beta function and $\Gamma(\alpha)$ is the Gamma function.

Distribution	Density Function	Polynomial Basis	Orthogonality Weight	Support
Normal	$\frac{1}{\sqrt{2\pi}}e^{\frac{-x^2}{2}}$	Hermite $He_n(x)$	$e^{\frac{-x^2}{2}}$	$[-\infty,\infty]$
Uniform	$\frac{1}{2}$	Legendre $P_n(x)$	1	$[-1,1]$
Beta	$\frac{(1-x)^\alpha(1+x)^\beta}{2^{\alpha+\beta+1}B(\alpha+1,\beta+1)}$	Jacobi $P_n^{(\alpha,\beta)}(x)$	$(1-x)^\alpha(1+x)^\beta$	$[-1,1]$
Exponential	e^{-x}	Laguerre $L_n(x)$	e^{-x}	$[0,\infty]$
Gamma	$\frac{x^\alpha e^{-x}}{\Gamma(\alpha+1)}$	Gen. Laguerre $L_n^{(\alpha)}(x)$	$x^\alpha e^{-x}$	$[0,\infty]$

II. Stochastic Expansions

In this section, we briefly review the non-intrusive polynomial chaos expansion method (PCE) and the stochastic collocation (SC) method. Let $R(\boldsymbol{\xi})$ be a "black box" that takes d stochastic parameters $\boldsymbol{\xi} = (\xi_1,\ldots,\xi_d)$ as inputs and return the system response as the output. Both methods construct a global polynomial approximation to the system response and have been shown to be point-wise equivalent if non-nested Gaussian quadrature nodes are used.[11] However, one form of the polynomial expansion may be preferred over the other, depending on the needs of the application (e.g., support for unstructured grids, fault tolerance, or local error estimation).

II.A. Non-Intrusive Polynomial Chaos

In polynomial chaos, one must estimate the chaos coefficients for a set of basis functions. The basis functions are obtained from the Askey family of hypergeometric orthogonal polynomials.[12] To reduce the nonlinearity of the expansion and improve convergence, the polynomial bases are chosen such that their orthogonality weighting functions match the probability density functions of the stochastic parameters up to a constant factor. Table 1 lists the appropriate polynomial bases for some commonly used continuous probability distributions. If the stochastic parameters do not follow these standard probability distributions, then the polynomial bases may be generated numerically.[13–15] Alternatively or if correlations are present, variable transformations may be used.[16]

The system response is approximated by the expansion

$$R(\boldsymbol{\xi}) \approx \sum_{\mathbf{i}\in\mathcal{I}_\mathbf{p}} a_\mathbf{i}\Psi_\mathbf{i}(\boldsymbol{\xi}),$$

where the basis functions $\Psi_\mathbf{i}(\boldsymbol{\xi})$ with multi-index $\mathbf{i} = (i_1,\ldots,i_d)$, $i_k = 0,1,2,\ldots$ are the product of the appropriate one-dimensional orthogonal polynomial basis of order i_k in each dimension $k = 1,\ldots,d$. The series is typically truncated in one of two ways. For the total-order expansion of order p, the index set is defined as

$$\mathcal{I}_p = \{\mathbf{i} : |\mathbf{i}| \le p\},$$

where $|\mathbf{i}| = i_1 + \ldots + i_d$, while for the tensor-product expansion of order $\mathbf{p} = (p_1,\ldots,p_d)$, the index set is defined as

$$\mathcal{I}_\mathbf{p} = \{\mathbf{i} : i_k \le p_k,\ k = 1,\ldots,d\}.$$

The number of terms required in each case is[16]

$$N_{\text{total-order}} = \frac{(d+p)!}{d!\,p!} = \binom{d+p}{d},$$

and

$$N_{\text{tensor-product}} = \prod_{k=1}^{d}(p_k+1),$$

respectively.

One approach to calculate the chaos coefficients $a_{\mathbf{i}}$ is the spectral projection method that takes advantage of the orthogonality of the bases. This results in

$$a_{\mathbf{i}} = \frac{\langle R(\boldsymbol{\xi}), \Psi_{\mathbf{i}}(\boldsymbol{\xi})\rangle}{\langle \Psi_{\mathbf{i}}^2(\boldsymbol{\xi})\rangle} = \frac{1}{\langle \Psi_{\mathbf{i}}^2(\boldsymbol{\xi})\rangle}\int_{\Omega} R(\boldsymbol{\xi})\Psi_{\mathbf{i}}(\boldsymbol{\xi})\rho(\boldsymbol{\xi})\,\mathrm{d}\boldsymbol{\xi},$$

where $\rho(\boldsymbol{\xi}) = \prod_{k=1}^{d}\rho_k(\xi_k)$ is the joint probability density of the stochastic parameters over the support $\Omega = \Omega_1 \times \ldots \times \Omega_d$. Thus, the bulk of the work is in evaluating the multi-dimensional integral in the numerator. Tensor product quadrature or, if d is moderately large, sparse grid quadrature may be employed. Gaussian quadrature rules can be chosen to match the orthogonal polynomial bases so that the weight function of the quadrature matches the probability density in the integral. Furthermore, if sparse grid quadrature is used to estimate the coefficients, the polynomial chaos expansion terms should be constructed from a sum of tensor expansions in order to avoid numerical noise in the coefficients of higher order terms.[11]

Once the chaos coefficients are known, the statistics of the system response can be estimated directly and inexpensively from the expansion. For example, the mean and the variance can be obtained analytically as

$$\mu_R = \langle R(\boldsymbol{\xi})\rangle \approx \sum_{\mathbf{i}\in\mathcal{I}_{\mathbf{p}}} a_{\mathbf{i}}\langle\Psi_{\mathbf{i}}(\boldsymbol{\xi})\rangle = a_{\mathbf{0}} \tag{1}$$

and

$$\sigma_R^2 = \left\langle R(\boldsymbol{\xi})^2\right\rangle - \mu_R^2 \approx \sum_{\mathbf{i}\in\mathcal{I}_{\mathbf{p}}}\sum_{\mathbf{j}\in\mathcal{I}_{\mathbf{p}}} a_{\mathbf{i}}a_{\mathbf{j}}\langle\Psi_{\mathbf{i}}(\boldsymbol{\xi})\Psi_{\mathbf{j}}(\boldsymbol{\xi})\rangle - a_{\mathbf{0}}^2 = \sum_{\mathbf{i}\in\mathcal{I}_{\mathbf{p}}\setminus\mathbf{0}} a_{\mathbf{i}}^2\left\langle\Psi_{\mathbf{i}}^2(\boldsymbol{\xi})\right\rangle \tag{2}$$

where $\mathbf{0}$ is a vector of zeros (i.e., the first multi-index). Other statistics such as probabilities can be estimated by sampling the polynomial expansion.

II.B. Stochastic Collocation

In stochastic collocation, a multivariate polynomial interpolant is formed over the system responses evaluated at a set of collocation points. In one-dimension, a degree n expansion (identified by the index i) has the form

$$R(\xi) \approx \mathcal{U}_i(R) \stackrel{\text{def}}{=} \sum_{j=0}^{n} R\left(\xi^{(j)}\right)\ell^{(j)}(\xi)$$

where $R\left(\xi^{(j)}\right)$ is the system response evaluated at collocation points $\xi^{(j)}$, $j = 0, \ldots, n$ and the basis $\ell^{(j)}(\xi)$ is the Lagrange polynomial

$$\ell^{(j)}(\xi) = \prod_{k=0,\ k\neq j}^{n}\frac{\xi - \xi^{(k)}}{\xi^{(j)} - \xi^{(k)}}.$$

The collocation points are chosen to be the nodes of the Gaussian quadrature rules associated with the same orthogonal polynomials as the polynomial chaos expansion. In the multivariate case, a tensor product formulation with the multi-index $\mathbf{i} = (i_1, \ldots, i_d)$, $i_k = 0, 1, 2, \ldots$ is applied:

$$R(\boldsymbol{\xi}) \approx \boldsymbol{\mathcal{U}}_{\mathbf{i}}(R) \stackrel{\text{def}}{=} (\mathcal{U}_{i_1}\otimes\ldots\otimes\mathcal{U}_{i_d})(R) = \sum_{j_1=0}^{n_1}\cdots\sum_{j_d=0}^{n_d} R\left(\xi_1^{(j_1)},\ldots,\xi_d^{(j_d)}\right)\ell_1^{(j_1)}(\xi_1)\ldots\ell_d^{(j_d)}(\xi_d).$$

Statistics of the system response such as the mean and the variance of a tensor product expansion can be obtained analytically as

$$\begin{aligned}\mu_R = \langle R(\boldsymbol{\xi})\rangle &\approx \sum_{j_1=0}^{n_1}\cdots\sum_{j_d=0}^{n_d} R\left(\xi_1^{(j_1)},\ldots,\xi_d^{(j_d)}\right)\left\langle\ell_1^{(j_1)}(\xi_1)\ldots\ell_d^{(j_d)}(\xi_d)\right\rangle \\ &= \sum_{j_1=0}^{n_1}\cdots\sum_{j_d=0}^{n_d} R\left(\xi_1^{(j_1)},\ldots,\xi_d^{(j_d)}\right)w_1^{(j_1)}\ldots w_d^{(j_d)} \\ &\stackrel{\text{def}}{=} \boldsymbol{\mathcal{Q}}_{\mathbf{i}}(R)\end{aligned}$$

and

$$\begin{aligned}\sigma_R^2 &= \left\langle R(\boldsymbol{\xi})^2 \right\rangle - \mu_R^2 \\ &\approx \sum_{j_1=0}^{n_1} \cdots \sum_{j_d=0}^{n_d} \sum_{k_1=0}^{n_1} \cdots \sum_{k_d=0}^{n_d} R\left(\xi_1^{(j_1)}, \ldots, \xi_d^{(j_d)}\right) R\left(\xi_1^{(k_1)}, \ldots, \xi_d^{(k_d)}\right) \left\langle \ell_1^{(j_1)}(\xi_1) \ldots \ell_d^{(j_d)}(\xi_d) \ell_1^{(k_1)}(\xi_1) \ldots \ell_d^{(k_d)}(\xi_d) \right\rangle - \mu_R^2 \\ &= \sum_{j_1=0}^{n_1} \cdots \sum_{j_d=0}^{n_d} R^2\left(\xi_1^{(j_1)}, \ldots, \xi_d^{(j_d)}\right) w_1^{(j_1)} \ldots w_d^{(j_d)} - \mu_R^2 \\ &\stackrel{\text{def}}{=} \boldsymbol{\mathcal{Q}}_{\mathbf{i}}\left(R^2\right) - \mu_R^2,\end{aligned}$$

where the expectation integrals are evaluated using the tensor product Gaussian quadrature corresponding to the set of collocation points. Using the property that $\ell^{(s)}\left(\xi^{(t)}\right) = \delta_{s,t}$, the numerical quadrature of the expectation integral leaves only the quadrature weights w. Higher moments can be obtained analytically in a similar manner and other statistics such as probabilities can be estimated by sampling the expansion.

If d is moderately large, then a sparse grid construction may be used to alleviate the exponential increase in the number of collocation points with respect to d. The sparse grid expansion is formed by a linear combination of tensor products $\boldsymbol{\mathcal{U}}_{\mathbf{i}}(R)$ at different multi-indices $\mathbf{i}$ in such a way that the exponential increase in the number of collocation points is alleviated up to a logarithmic factor. Note that while the tensor product expansion interpolates the system responses at the collocation points, the sparse grid expansion may not interpolate unless a nested set of collocation points is used.[17]

The relationship between the index i_k and the number of collocation points n_k in each dimension $k = 1, \ldots, d$ is called the growth rule and is important during the sparse grid construction. If the collocation points are chosen based on a fully nested quadrature rule, then a nonlinear growth rule that approximately doubles n_k with every increment in i_k should be used to reuse model evaluations. If the collocation points are based on a weakly-nested or non-nested quadrature rule, then a linear growth rule may alternatively be used to provide finer granularity in the degree of the interpolant.

The formulas $\mu_R \approx \mathcal{Q}_{\mathbf{i}}(R)$ and $\sigma_R^2 \approx \boldsymbol{\mathcal{Q}}_{\mathbf{i}}\left(R^2\right) - \mu_R^2$ can be extended to estimate the mean and variance of the sparse grid stochastic collocation using a linear combination of $\boldsymbol{\mathcal{Q}}_{\mathbf{i}}(R)$ or $\boldsymbol{\mathcal{Q}}_{\mathbf{i}}\left(R^2\right)$, respectively, at multi-indices $\mathbf{i}$ corresponding to the sparse grid construction. In the case of weakly-nested or non-nested quadrature rules, $\sum_{j=0}^{N} R_j w_j$ and $\sum_{j=0}^{N} R_j^2 w_j$ may still be used for $\langle R \rangle$ and $\left\langle R^2 \right\rangle$, but these raw moments are no longer equivalent to the expectations of the expansion and the square of the expansion due to the loss of the interpolation property in non-nested sparse grids (the numerical quadrature does not simplify as in the case of the tensor product expansion). Rather, they instead correspond to the sparse numerical integration of R and R^2.

II.C. Sparse Grid Construction

Since our multifidelity approach focuses on sparse grids, we briefly review isotropic and generalized sparse grid constructions here. Sparse grids, as employed in numerical integration[18,19] and interpolation,[17,20] are contructed from a linear combination of tensor product grids with relatively small numbers of grid points in such a way that preserves a high level of accuracy. The isotropic sparse grid at level q where $q = 0, 1, 2, \ldots$ is defined as

$$\mathcal{A}_{q,d}(R) = \sum_{q-d+1 \le |\mathbf{i}| \le q} (-1)^{q-|\mathbf{i}|} \binom{d-1}{q-|\mathbf{i}|} \boldsymbol{\mathcal{U}}_{\mathbf{i}}(R), \tag{3}$$

where the tensor product interpolation formulas $\boldsymbol{\mathcal{U}}_{\mathbf{i}}(R)$, $\mathbf{i} = (i_1, \ldots, i_d)$, $i_k = 0, 1, 2, \ldots$ can be replaced by tensor product quadrature formulas $\boldsymbol{\mathcal{Q}}_{\mathbf{i}}(R)$ for the case of sparse grid integration. Alternatively, Equation 3 may be expressed in terms of the difference formulas:

$$\mathcal{A}_{q,d}(R) = \sum_{|\mathbf{i}| \le q} \boldsymbol{\Delta}_{\mathbf{i}}(R), \tag{4}$$

where $\boldsymbol{\Delta}_{\mathbf{i}}(R) = (\Delta_{i_1} \otimes \ldots \otimes \Delta_{i_d})(R)$ and $\Delta_{i_k} = \mathcal{U}_{i_k} - \mathcal{U}_{i_k - 1}$ with $\mathcal{U}_{-1} = 0$. Thus, the tensor product grids used in the isotropic sparse grid construction are those whose multi-indices lie within a simplex defined by the sparse grid level q.

The generalized sparse grid construction relaxes the simplex constraint on the multi-indices in Equation 3 to provide flexibility for adaptive refinement. In the relaxed constraint, the set of multi-indices $\mathcal{J}$ is admissible if $\mathbf{i} - \mathbf{e}_k \in \mathcal{J}$ for all $\mathbf{i} \in \mathcal{J}$, $i_k \geq 1$, $k = 1, \ldots, d$, where $\mathbf{e}_k$ is the k^{th} unit vector.[19] Thus, admissible multi-indices can be added one by one starting with $\mathbf{i} = \mathbf{0}$ based on a refinement metric and the generalized sparse grid is then

$$\mathcal{A}_{q,d}(R) = \sum_{\mathbf{i} \in \mathcal{J}} \boldsymbol{\Delta}_{\mathbf{i}}(R).$$

III. Multifidelity Extensions

Let $R_{\text{high}}(\boldsymbol{\xi})$ be the system response obtained by evaluating the expensive high-fidelity model and $R_{\text{low}}(\boldsymbol{\xi})$ be the system response obtained by evaluating the cheap low-fidelity model. In a multifidelity stochastic expansion, the low-fidelity model values are corrected to match the high-fidelity model values (and potentially their derivatives) at the high-fidelity collocation points.

III.A. Corrected Low-Fidelity Model

We investigate additive correction, multiplicative correction, and combined additive and multiplicative correction for the low-fidelity model. Defining the additive correction function and multiplicative correction function as

$$C_\alpha(\boldsymbol{\xi}) = R_{\text{high}}(\boldsymbol{\xi}) - R_{\text{low}}(\boldsymbol{\xi})$$

and

$$C_\beta(\boldsymbol{\xi}) = \frac{R_{\text{high}}(\boldsymbol{\xi})}{R_{\text{low}}(\boldsymbol{\xi})},$$

respectively, then

$$R_{\text{high}}(\boldsymbol{\xi}) = R_{\text{low}}(\boldsymbol{\xi}) + C_\alpha(\boldsymbol{\xi})$$

or

$$R_{\text{high}}(\boldsymbol{\xi}) = R_{\text{low}}(\boldsymbol{\xi}) C_\beta(\boldsymbol{\xi}).$$

In the combined case,

$$R_{\text{high}}(\boldsymbol{\xi}) = \gamma(R_{\text{low}}(\boldsymbol{\xi}) + C_\alpha(\boldsymbol{\xi})) + (1-\gamma) R_{\text{low}}(\boldsymbol{\xi}) C_\beta(\boldsymbol{\xi}),$$

where $\gamma \in [0, 1]$ defines a convex combination that determines the proportion of additive correction or multiplicative correction employed in the combined correction.

Let $S_{q,d}[R]$ be the stochastic expansion (non-intrusive polynomial chaos or stochastic collocation) of $R(\boldsymbol{\xi})$ at sparse grid level q with dimension d. We add the superscript "pc" or "sc" when we refer specifically to non-intrusive polynomial chaos or stochastic collocation, respectively. Also, let $N_{q,d}$ be the number of model evaluations required to construct $S_{q,d}[R]$. Thus, $R_{\text{high}}(\boldsymbol{\xi}) \approx S_{q,d}[R_{\text{high}}](\boldsymbol{\xi})$. We further approximate the stochastic expansion of the high-fidelity model with $S_{q,d}[R_{\text{high}}](\boldsymbol{\xi}) \approx \tilde{R}_{\text{high}}(\boldsymbol{\xi})$, where $\tilde{R}_{\text{high}}(\boldsymbol{\xi})$ is the multifidelity stochastic expansion based on the additive, multiplicative, or combined correction of the low-fidelity model

$$\tilde{R}_{\text{high}} = S_{q,d}[R_{\text{low}}] + S_{q-r,d}[C_\alpha],$$

$$\tilde{R}_{\text{high}} = S_{q,d}[R_{\text{low}}] S_{q-r,d}[C_\beta],$$

$$\tilde{R}_{\text{high}} = \gamma(S_{q,d}[R_{\text{low}}] + S_{q-r,d}[C_\alpha]) + (1-\gamma) S_{q,d}[R_{\text{low}}] S_{q-r,d}[C_\beta],$$

respectively, where $r < q$ is a sparse grid level offset between the stochastic expansion of the low-fidelity model and the stochastic expansion of the correction function. Thus, the multifidelity stochastic expansion at sparse grid level q can be constructed with $N_{q-r,d}$ instead of $N_{q,d}$ high-fidelity model evaluations plus $N_{q,d}$ low-fidelity model evaluations. If the low-fidelity model is much cheaper to evaluate than the high-fidelity model, then significant computational savings will be obtained. Furthermore, if the low-fidelity model is sufficiently predictive, then accuracy comparable to the single fidelity expansion with $N_{q,d}$ high-fidelity model evaluations can be achieved. This notion of a sparse grid level offset is used in the case of

predetermined grid levels and enforces computational savings explicitly, with less regard to managing the accuracy in $\tilde{R}_{\text{high}}(\boldsymbol{\xi})$. In the case of adaptive refinement, we instead manage the accuracy explicitly by investing resources where they are most needed for resolving statistics of $\tilde{R}_{\text{high}}(\boldsymbol{\xi})$, and the computational savings that result are achieved implicitly.

In the combined correction, there is flexibility in the choice of γ. One option is to compute γ based on a regularization of the combined correction function to prevent overfitting. We propose a simple expression that can be computed analytically by minimizing the magnitude of the additive and multiplicative correction in the mean-square sense:

$$\min_{\gamma\in[0,1]} \Big\langle \gamma^2 C_\alpha^2(\boldsymbol{\xi}) + (1-\gamma)^2 C_\beta^2(\boldsymbol{\xi}) \Big\rangle,$$

which gives

$$\gamma = \frac{\Big\langle C_\beta^2(\boldsymbol{\xi}) \Big\rangle}{\langle C_\alpha^2(\boldsymbol{\xi})\rangle + \Big\langle C_\beta^2(\boldsymbol{\xi}) \Big\rangle},$$

where the second moments of the $C_\alpha(\boldsymbol{\xi})$ and $C_\beta(\boldsymbol{\xi})$ can be estimated analytically from their stochastic expansions as described in Sections II.A and II.B. This choice of γ balances the additive correction and the multiplicative correction such that neither becomes too "large".

III.B. Analytic Moments

In order to compute the moments of the multifidelity stochastic expansion analytically, we collapse the sum or product of the expansion of the low-fidelity model and the expansion of the correction function into a single expansion and then employ the standard moment calculation techniques from Sections II.A and II.B. This is most straightforward for non-intrusive polynomial chaos expansions with additive correction, and we will start with this case using a statically determined sparse grid offset. Let $\mathcal{J}_{q,d}$ be the set of multi-indices of the d-dimensional polynomial chaos expansion bases at sparse grid level q. Then,

$$S^{\text{pc}}_{q,d}[R_{\text{low}}](\boldsymbol{\xi}) = \sum_{\mathbf{i}\in\mathcal{J}_{q,d}} a_{\text{low}\,\mathbf{i}}\Psi_{\mathbf{i}}(\boldsymbol{\xi})$$

and

$$S^{\text{pc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi}) = \sum_{\mathbf{i}\in\mathcal{J}_{q-r,d}} a_{C_\alpha\,\mathbf{i}}\Psi_{\mathbf{i}}(\boldsymbol{\xi}).$$

Since the $\mathcal{J}_{q-r,d} \subset \mathcal{J}_{q,d}$, the bases $\Psi_{\mathbf{i}}(\boldsymbol{\xi})$, $\mathbf{i} \in \mathcal{J}_{q-r,d}$ are common between $S_{q,d}[R_{\text{low}}]$ and $S_{q-r,d}[C_\alpha]$. Therefore, the chaos coefficients of those bases can be added to produce a single non-intrusive polynomial chaos expansion

$$S^{\text{pc}}_{q,d}[R_{\text{low}}](\boldsymbol{\xi}) + S^{\text{pc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi}) = \sum_{\mathbf{i}\in\mathcal{J}_{q-r,d}} (a_{\text{low}\,\mathbf{i}} + a_{C_\alpha\,\mathbf{i}})\Psi_{\mathbf{i}}(\boldsymbol{\xi}) + \sum_{\mathbf{i}\in\mathcal{J}_{q,d}\setminus\mathcal{J}_{q-r,d}} a_{\text{low}\,\mathbf{i}}\Psi_{\mathbf{i}}(\boldsymbol{\xi}).$$

As shown in Equations 1 and 2, the mean and the variance can be computed directly from the multifidelity non-intrusive polynomial chaos as

$$\mu_R \approx a_{\text{low}\,\mathbf{0}} + a_{C_\alpha\,\mathbf{0}}$$

$$\sigma_R^2 \approx \sum_{\mathbf{i}\in\mathcal{J}_{q-r,d}\setminus\mathbf{0}} (a_{\text{low}\,\mathbf{i}} + a_{C_\alpha\,\mathbf{i}})\big\langle\Psi_{\mathbf{i}}^2(\boldsymbol{\xi})\big\rangle + \sum_{\mathbf{i}\in\mathcal{J}_{q,d}\setminus\mathcal{J}_{q-r,d}} a_{\text{low}\,\mathbf{i}}\big\langle\Psi_{\mathbf{i}}^2(\boldsymbol{\xi})\big\rangle.$$

In the multiplicative correction case for non-intrusive polynomial chaos, we again combine the low fidelity and discrepancy expansions and then compute the moments from the aggregated expansion. Multiplication of chaos expansions is a kernel operation within stochastic Galerkin methods. The coefficients of a product expansion are computed as follows (shown generically for $z = xy$ where x, y, and z are each expansions of

arbitrary form):

$$\begin{aligned}\sum_{k=0}^{P_z} z_k \Psi_k(\boldsymbol{\xi}) &= \sum_{i=0}^{P_x}\sum_{j=0}^{P_y} x_i y_j \Psi_i(\boldsymbol{\xi})\Psi_j(\boldsymbol{\xi}) \\ z_k &= \frac{\sum_{i=0}^{P_x}\sum_{j=0}^{P_y} x_i y_j \langle \Psi_i \Psi_j \Psi_k \rangle}{\langle \Psi_k^2 \rangle}\end{aligned}$$

where three-dimensional tensors of one-dimensional basis triple inner products $\langle \psi_i \psi_j \psi_k \rangle$ are typically sparse and can be efficiently precomputed using one dimensional quadrature for fast lookup within the multidimensional basis triple inner products $\langle \Psi_i \Psi_j \Psi_k \rangle$. The form of the high-fidelity expansion must first be defined to include all polynomial orders indicated by the products of each of the basis polynomials in the low fidelity and discrepancy expansions. These are readily estimated from total-order, tensor, or sum of tensor expansions since they involve simple polynomial order additions for each tensor or total-order expansion product.

Evaluating the moments for stochastic collocation with either additive or multiplicative correction involves forming a new interpolant on the more refined (low fidelity) grid. Therefore, we perform an additional step to create a single stochastic expansion

$$\begin{aligned}S^{\text{sc}}_{q,d}\big\{S^{\text{sc}}_{q,d}[R_{\text{low}}] + S^{\text{sc}}_{q-r,d}[C_\alpha]\big\} \\ S^{\text{sc}}_{q,d}\big\{S^{\text{sc}}_{q,d}[R_{\text{low}}] S^{\text{sc}}_{q-r,d}[C_\beta]\big\}\end{aligned}$$

from which the variance and higher moments can be obtained analytically. This requires evaluating $S^{\text{sc}}_{q,d}[R_{\text{low}}](\boldsymbol{\xi})$ and either $S^{\text{sc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi})$ or $S^{\text{sc}}_{q-r,d}[C_\beta](\boldsymbol{\xi})$ at the collocation points associated with the multi-indices in $\mathcal{J}_{q,d}$. For the former, the low-fidelity model values at all of the collocation points are readily available and can be used directly. For the latter, the discrepancy expansion $S^{\text{sc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi})$ or $S^{\text{sc}}_{q-r,d}[C_\beta](\boldsymbol{\xi})$ must be evaluated. Since the correction function values are available at collocation points associated with the multi-indices of $\mathcal{J}_{q-r,d}$, it may be tempting to only evaluate $S^{\text{sc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi})$ or $S^{\text{sc}}_{q-r,d}[C_\beta](\boldsymbol{\xi})$ at collocation points associated with the multi-indices in $\mathcal{J}_{q,d} \setminus \mathcal{J}_{q-r,d}$. However, because sparse grid stochastic collocation does not interpolate unless the set of collocation points are nested,[17] the function values of $C_\alpha(\boldsymbol{\xi})$ and $C_\beta(\boldsymbol{\xi})$ are not consistent with those from $S^{\text{sc}}_{q-r,d}[C_\alpha](\boldsymbol{\xi})$ and $S^{\text{sc}}_{q-r,d}[C_\beta](\boldsymbol{\xi})$ and should not be mixed together.

III.C. Adaptive Refinement

Conceptually speaking, the goal of an adaptive refinement procedure applied to multifidelity modeling should be to preferentially refine where the model discrepancy has the greatest complexity. This corresponds to regions of the stochastic domain where the low fidelity model becomes less predictive. It is often the case in real-world applications that low fidelity models may be predictive for significant portions of a parameter space, but, in other portions of the space, the simplifying assumptions break down and a higher-fidelity model must be relied upon. By selectively targeting these regions, we rely on the low fidelity model where it is effective and more faithfully resolve the discrepancy where it is not. Thus, adaptive refinement procedures can extend the utility of multifidelity uncertainty quantification approaches in cases where the predictive capability of low fidelity models is strongly parameter-dependent.

For adaptive refinement in a multifidelity context, we will employ a greedy adaptation based on the generalized sparse grid procedure described in Section II.C. One option is to separately adapt the low fidelity and discrepancy models; however, an optimal approach for controlling the relative levels of refinement (e.g., enforcing different convergence tolerances for the low fidelity and discrepancy adaptations) is not obvious in this case. Moreover, the effect of the individual candidate refinements on their individual statistical QOIs (e.g., low fidelity or discrepancy variance) are not viewed in the aggregated multifidelity context (e.g., high fidelity variance or failure probability). Thus, we instead propose a further generalization to generalized sparse grids in which we consider candidate index sets from multiple sparse grids simulaneously, and measure their effects within the aggregated context using appropriate cost normalization. The algorithmic steps can be summarized as:

1. *Initialization:* Starting from an initial reference sparse grid for the lowest fidelity model and each level of discrepancy within a multifidelity hierarchy, accept these reference index sets as the old set and define active index sets using the admissible forward neighbors of all old index sets.

2. *Trial set evaluation:* For each trial active index set within each of the sparse grids, perform the tensor grid evaluations of either the low fidelity or discrepancy functions, form the tensor polynomial chaos expansion or tensor interpolant corresponding to the grid, combine the trial expansion with the reference expansion for the particular level in the expansion hierarchy to which it corresponds (update $S_{q_{low},d}[R_{\text{low}}]$, $S_{q_\alpha,d}[C_\alpha]$, or $S_{q_\beta,d}[C_\beta]$), and then combine each of the levels to generate a high fidelity expansion ($\tilde{R}_{\text{high}}$). Note that index sets associated with discrepancy expansions require evaluation of two levels of fidelity, so caching and reuse of the lowest and all intermediate fidelity evaluations should be performed among the different sparse grids. Bookkeeping should also be performed to allow efficient restoration of previously evaluated tensor expansions, as they will remain active until either selected or processed in the finalization step.

3. *Trial set selection:* From among all of the candidates, select the trial index set that induces the largest change in the high fidelity statistical QOI, normalized by the cost of evaluating the trial index set (as indicated by the number of new collocation points and the average model run time per point). Initial estimates of relative simulation cost among the different fidelities are thus required to appropriately bias the adaptation, and discrepancy evaluations must incur the cost of two fidelities. To achieve greater parallelism, the best n index sets may be selected simultaneously, resulting in additional trial sets on the following cycle.

4. *Update sets:* If the largest change induced by the active trial sets exceeds a specified convergence tolerance, then promote the selected trial set(s) from the active set to the old set and update the active set with new admissible forward neighbors; return to step 2 and evaluate all active trial sets with respect to the new reference grid. If the convergence tolerance is satisfied, advance to step 5. In the common case of two levels of fidelity, we can choose to always promote selected discrepancy index sets into the old sets for both the low fidelity and discrepancy grids, such that the discrepancy index sets are always a subset of the low fidelity index sets. While there can be a small penalty for doing this (additional low fidelity trial sets that are generated may not be in perfect overlap with the additional discrepancy trial sets, due to different active set states), this backfill is consistent with the intent of more fully resolving the low-fidelity expansion and the overhead in any additional low-fidelity evaluations should be small. For multiple levels of fidelity, a full recursion in the promotion process can be similarly enforced so that all index sets in higher level grids are subsets of those in the lower levels; however, the additional degrees of freedom in this case admit other strategies as well (e.g., a minimal enforcement of overlap in the reference grids would ensure that the lowest fidelity grid contains all old index sets in the discrepancy level immediately above and all mid-level discrepancies contain all old index sets included within *both* of its nearest neighbors in the hierarchy).

5. *Finalization:* Promote all remaining active sets to the old set, update all expansions within the hierarchy, and perform a final combination of the low fidelity and discrepancy expansions to arrive at the final result for the high fidelity statistical QOI.

In the limiting case where the low fidelity model provides no useful information, then this algorithm will prefer refinements to the model discrepancy and will closely mirror the single-fidelity case, with the penalty of carrying along the low fidelity evaluations needed to resolve the discrepancy. This suggests an additional adaptive control, in which one drops low (and intermediate) fidelity models from the hierarchy that are adding expense but not adding value (as measured by their frequency of selection in step 3). In addition, this general framework can be extended to include pointwise local refinement using hierarchical surpluses[21] as well as adjoint-enhanced approaches.[22,23]

IV. Computational Results

We compare the performance of multifidelity stochastic expansion and single-fidelity stochastic expansion for several algebraic and PDE models. We demonstrate cases for which the multifidelity stochastic expansion converges more quickly than the single-fidelity stochastic expansion as well as cases for which the multifidelity stochastic expansion offers no efficiency gain.

IV.A. Simple One-Dimensional Example

First, we present a simple example to motivate the approach and demonstrate the efficiency improvements that are possible when an accurate low-fidelity model is available. The system responses of the high-fidelity model and the low-fidelity model are, respectively,

$$R_{\text{high}}(\xi) = e^{-0.05\xi^2}\cos 0.5\xi - 0.5e^{-0.02(\xi-5)^2}$$

$$R_{\text{low}}(\xi) = e^{-0.05\xi^2}\cos 0.5\xi,$$

where $\xi \sim \text{Uniform}[-8, 12]$. An additive correction $C_\alpha(\xi) = R_{\text{high}}(\xi) - R_{\text{low}}(\xi)$ is used, which is just the second term of $R_{\text{high}}(\xi)$. In Figure 1, we compare the convergence in mean and standard deviation of the single (high) fidelity PCE with the convergence of the multifidelity PCE. The multifidelity PCE is constructed from a PCE of the correction function at order 1 to 20 combined with a PCE of the low-fidelity model at order 60 (for which the low-fidelity statistics are converged to machine precision). This corresponds to the simpler case where low-fidelity expense can be assumed to be negligible, and by eliminating any issues related to low fidelity accuracy, we can focus more directly on comparing the convergence of the correction function with convergence of the high-fidelity model (in the short column example to follow, we will employ grid level offsets that accommodate nontrivial low-fidelity expense). The error is plotted against the number of high-fidelity model evaluations and is measured with respect to an overkill single-fidelity PCE solution at order 60. The multifidelity stochastic expansion converged much more rapidly because the additive correction function in the example has lower complexity than the high-fidelity model. This can be seen from comparison of the decay of the normalized spectral coefficients as plotted in Figure 2, which shows that the discrepancy expansion converges more rapidly allowing the moments of the discrepancy expansion to achieve a given accuracy using fewer PCE terms than that required by the high-fidelity model.

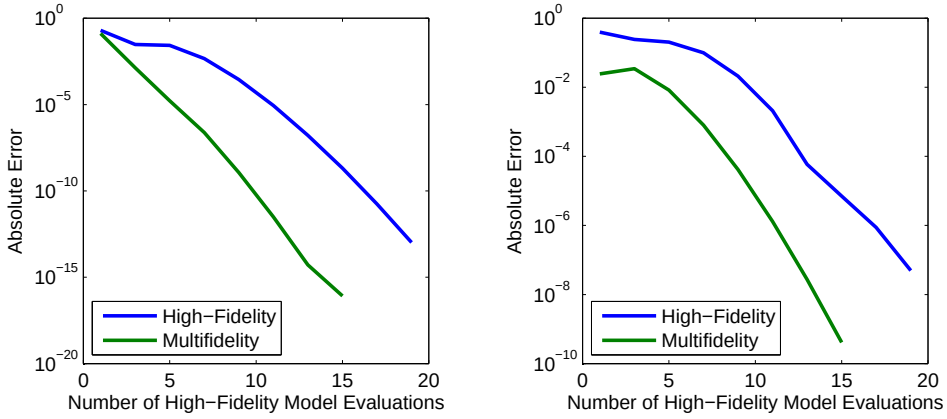

(a) Error in mean

(b) Error in standard deviation

Figure 1: Convergence of single-fidelity PCE and multifidelity PCE with additive correction for the one-dimensional example.

IV.B. Short Column Example

The five-dimensional short column example[24] serves to demonstrate a more realistic, albeit still algebraic, problem involving multifidelity models. For this example, we will employ relative grid refinement levels between low-fidelity and high-fidelity, which are more appropriate for cases where the low-fidelity model expense in non-negligible. Let the system response of the high-fidelity model be

$$R_{\text{high}}(\boldsymbol{\xi}) = 1 - \frac{4M}{bh^2Y} - \left(\frac{P}{bhY}\right)^2, \tag{5}$$

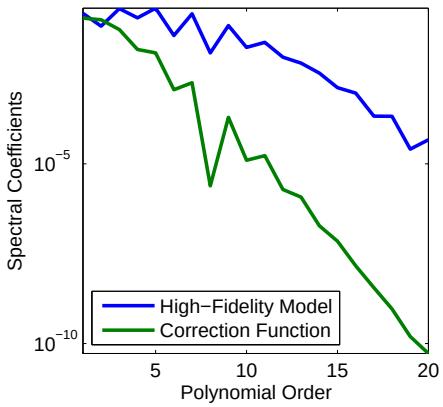

Figure 2: Normalized spectral coefficients of the high-fidelity model and the additive correction function for the one-dimensional example.

where $\boldsymbol{\xi} = (b, h, P, M, Y)$, $b \sim \text{Uniform}[5, 15]$, $h \sim \text{Uniform}[15, 25]$, $P \sim \text{Normal}[500, 100]$, $M \sim \text{Normal}[2000, 400]$, and $Y \sim \text{Lognormal}[5, 0.5]$ and we neglect the traditional correlation between P and M for simplicity. We consider three artificially constructed low-fidelity models of various predictive quality:

$$R_{\text{low}1}(\boldsymbol{\xi}) = 1 - \frac{4P}{bh^2Y} - \left(\frac{P}{bhY}\right)^2, \qquad C_{\alpha 1}(\boldsymbol{\xi}) = \frac{4(P-M)}{bh^2Y},$$

$$R_{\text{low}2}(\boldsymbol{\xi}) = 1 - \frac{4M}{bh^2Y} - \left(\frac{M}{bhY}\right)^2, \qquad C_{\alpha 2}(\boldsymbol{\xi}) = \frac{M^2 - P^2}{(bhY)^2},$$

$$R_{\text{low}3}(\boldsymbol{\xi}) = 1 - \frac{4M}{bh^2Y} - \left(\frac{P}{bhY}\right)^2 - \frac{4(P-M)}{bhY}, \qquad C_{\alpha 3}(\boldsymbol{\xi}) = \frac{4(P-M)}{bhY}.$$

In Figure 3, PCE with isotropic sparse grids is used and the offset r in the sparse grid level between the low-fidelity model and the correction function is fixed at one. It can be seen that the multifidelity case using $R_{\text{low}1}(\boldsymbol{\xi})$ results in a reduction in the number of high-fidelity model evaluations required for a given error compared to the single-fidelity case using $R_{\text{high}}(\boldsymbol{\xi})$. For example, at 10^{-5} error, the number of high-fidelity model evaluations is reduced from about 500 to about 100. While still a rational function with broad spectral content, the correction function, $C_{\alpha 1}(\boldsymbol{\xi})$, is similar to the middle term in Eq. 5 and has eliminated the final term possessing the greatest nonlinearity. Conversely, the correction function for the second low-fidelity model, $C_{\alpha 2}(\boldsymbol{\xi})$, is similar to the final term in Eq. 5 and has expanded it to include an additional dimension, resulting in a larger number of high-fidelity model evaluations for a given error compared to the single-fidelity case. For the third low-fidelity model, the convergence rate is faster than the single-fidelity case but the starting error is also larger, resulting in a break-even point at about 11 high-fidelity model evaluations. This suggests that while a smoother and/or less complex correction function provides a faster convergence rate, it is also important to consider the magnitude of the correction.

We modify $R_{\text{low}3}(\boldsymbol{\xi})$ by changing the scalar in the last term from 4 to 0.4 and label it $R_{\text{low}4}(\boldsymbol{\xi})$. Similarly, we also change the scalar in the last term from 4 to 40 and label it $R_{\text{low}5}(\boldsymbol{\xi})$. Thus, the correction functions $C_{\alpha 3}(\boldsymbol{\xi})$, $C_{\alpha 4}(\boldsymbol{\xi})$, and $C_{\alpha 5}(\boldsymbol{\xi})$ have the same smoothness and spectral content, but the magnitude of the correction is an order of magnitude smaller for $C_{\alpha 4}(\boldsymbol{\xi})$ and an order of magnitude larger for $C_{\alpha 5}(\boldsymbol{\xi})$. As plotted in Figure 4, the means have similar convergence rates, but a smaller correction results in lower error.

Figure 5 is the same as Figure 3 but with the sparse grid level offset, r, increased from one to two, resulting in greater resolution in the low fidelity expansion for a particular discrepancy expansion resolution. The results from Figure 3 are included as solid lines, and the new results with $r = 2$ are shown as dashed lines. A small improvement can be seen in the mean convergence for multifidelity using $R_{\text{low}1}(\boldsymbol{\xi})$ and $R_{\text{low}3}(\boldsymbol{\xi})$ and for standard deviation convergence using $R_{\text{low}1}(\boldsymbol{\xi})$, but results are mixed and it is unclear whether the benefit of increasing the offset is worth the additional low-fidelity evaluations, especially in the

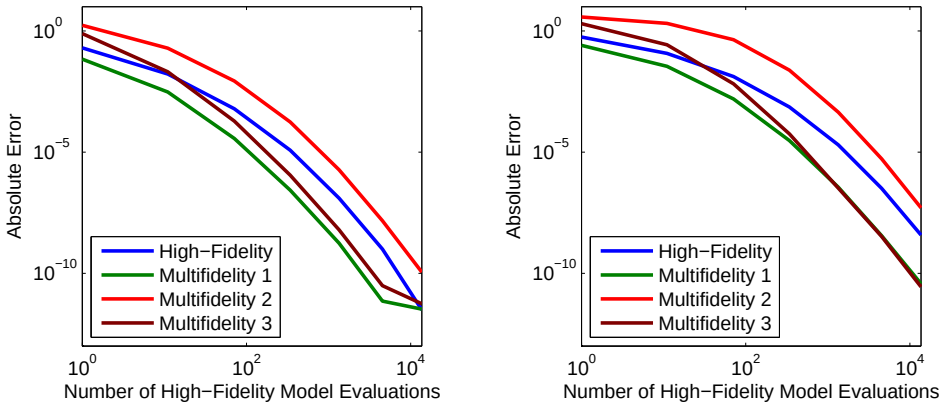

(a) Error in mean (b) Error in standard deviation

Figure 3: Convergence of single-fidelity PCE and multifidelity PCE with additive correction using isotropic sparse grids for the short column example. The multifidelity sparse grid level offset is set to one.

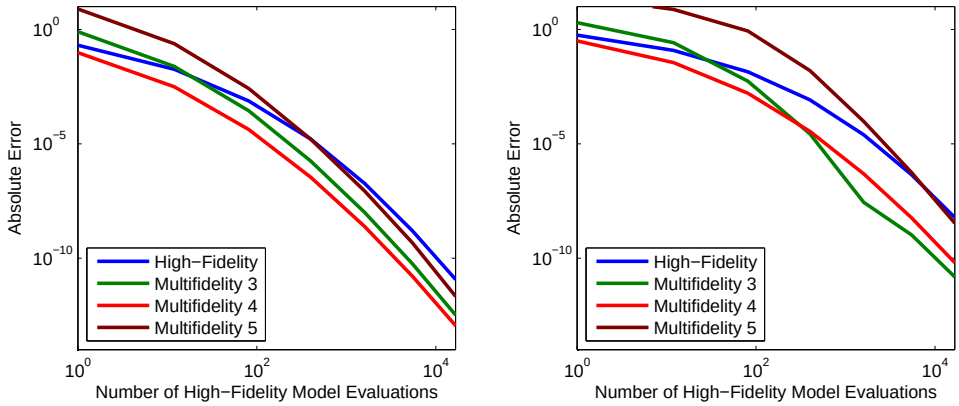

(a) Error in mean (b) Error in standard deviation

Figure 4: Convergence of single-fidelity PCE and multifidelity PCE with additive correction using isotropic sparse grids for the short column example. The multifidelity sparse grid level offset is set to one.

case where their expense is nontrivial. Thus, it is evident that an automated procedure will be needed to optimize these offsets accounting for relative cost. This motivates the adaptive refinement strategy described in Section III.C.

IV.C. Elliptic PDE Example

Next, we consider the stochastic PDE in one spatial dimension

$$-\frac{d}{dx}\left[\kappa(x,\omega)\frac{du(x,\omega)}{dx}\right] = 1, \qquad x \in (0,1), \qquad u(0,\omega) = u(1,\omega) = 0.$$

The coefficient is described by the following 10-dimensional Karhunen-Loève expansion

$$\kappa(x,\omega) = 0.1 + 0.03\sum_{k=1}^{10}\sqrt{\lambda_k}\phi_k(x)Y_k(\omega), \qquad Y_k \sim \text{Uniform}[-1,1]$$

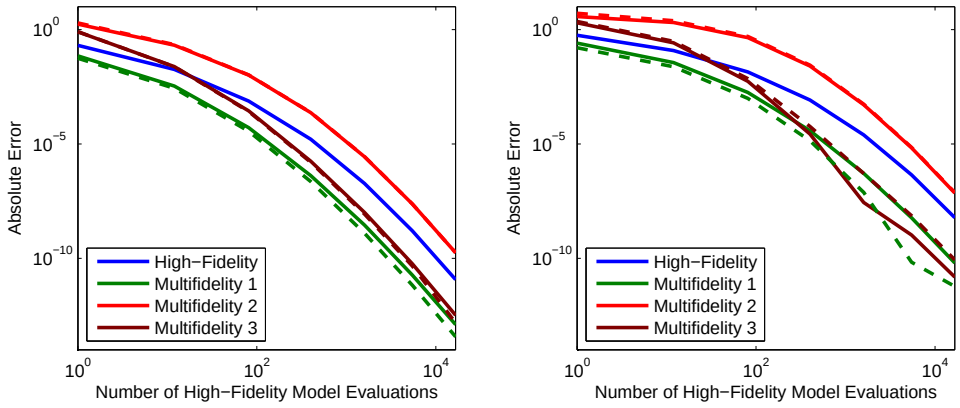

(a) Error in mean (b) Error in standard deviation

Figure 5: Convergence of single-fidelity PCE and multifidelity PCE with additive correction using isotropic sparse grids for the short column example. The multifidelity sparse grid level offset is compared using $r = 1$ from Figure 3 (solid lines) and $r = 2$ (dashed lines).

for the Gaussian covariance kernel

$$C_{\kappa\kappa}(x, x') = \exp\left[-\left(\frac{x - x'}{0.2}\right)^2\right].$$

The PDE is solved by finite elements and the output of interest is $u(0.5, \omega)$. We use a fine spatial grid with 500 states for the high-fidelity model and a coarse spatial grid with 50 states for the low-fidelity model. The ratio of average run time between the high-fidelity model and the low-fidelity model is $r_{\text{work}} = 40$.

We compute the mean and standard deviation using multifidelity PCE with sparse grid level 4 applied to the low-fidelity model and sparse grid level 3 applied to the additive correction function (i.e., sparse grid level offset $r = 1$). Table 2 compares the relative error with single-fidelity PCE of the high-fidelity model at sparse grid level 4 and at sparse grid level 3. It can be seen that the multifidelity PCE is able to achieve lower relative error than the single-fidelity PCE at sparse grid level 3 while using the same number of high-fidelity evaluations. The cost of low-fidelity evaluations is equivalent to about 325 additional high-fidelity evaluations, resulting in greater than an 80% reduction in total cost for comparable accuracy.

Table 2: Comparison of the relative error and the number of model evaluations for the elliptic PDE example.

	Relative Error in Mean	Relative Error in Std Deviation	High-Fidelity Evaluations	Low-Fidelity Evaluations
Single-Fidelity ($q = 3$)	5.3×10^{-6}	2.7×10^{-4}	1981	–
Single-Fidelity ($q = 4$)	4.1×10^{-7}	2.3×10^{-5}	12,981	–
Multifidelity ($q = 4$, $r = 1$)	4.7×10^{-7}	2.6×10^{-5}	1981	12,981

Figure 6 shows the convergence for the single-fidelity and multifidelity PCE with additive correction function based on adaptive refinement using generalized sparse grids. The single-fidelity case uses the standard generalized sparse grid procedure,[19] whereas the multifidelity case uses the algorithm described in Section III.C. The initial grid for both the low-fidelity model and the correction function is a level one sparse grid (requiring 11 model evaluations). We use the equivalent number of high-fidelity model evaluations, defined as $N_{\text{eqv}} = N_{\text{high}} + N_{\text{low}}/r_{\text{work}}$, to include the additional cost of low-fidelity model evaluations in the comparison with the single-fidelity case. By considering the potential error reduction per unit cost of refining the sparse grid of the correction function versus that of refining the sparse grid of the low-fidelity model, the multifidelity adaptive algorithm is able to achieve a faster convergence than the single-fidelity

adaptive generalized sparse grid. Relative to the non-adaptive results from Table 2, the adaptive multifidelity algorithm reduces the equivalent number of high-fidelity evaluations by 33% for the mean and 62% for the standard deviation.

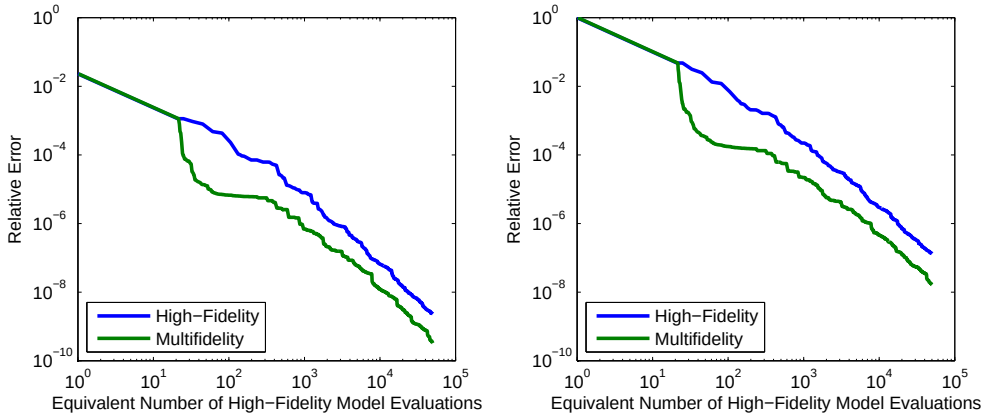

(a) Error in mean (b) Error in standard deviation

Figure 6: Convergence of single-fidelity PCE and multifidelity PCE with additive correction using adaptive generalized sparse grids for the elliptic PDE example.

IV.D. Horn Acoustics Example

We model the propagation of acoustic waves through a two-dimensional horn with the non-dimensional Helmholtz equation $\nabla^2 u + k^2 u = 0$. The incoming wave of wave number k enters the waveguide and exits the flare of the horn into the exterior domain with a truncated absorbing boundary.[25] The horn geometry is illustrated in Figure 7. The stochastic parameters are the wave number $k \sim \text{Uniform}[1.3, 1.5]$, upper horn wall impedance $z_u \sim \text{Normal}[50, 9]$, and lower horn wall impedance $z_l \sim \text{Normal}[50, 9]$, where the latter two represent imperfections in the horn wall. We compute the mean and standard deviation of the reflection coefficient, where a low reflection coefficient is desired for an efficient horn.

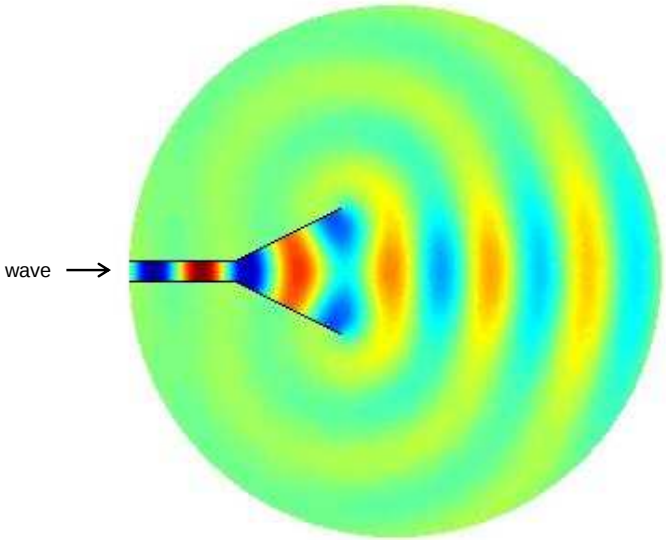

Figure 7: 2-D horn geometry and the propagation of acoustic waves.

The high-fidelity model solves the Helmholtz equation by finite elements using 35895 states and the low-fidelity model is a reduced basis model constructed from the finite element discretization[26] using 50 bases.

The ratio of average run time between the high-fidelity model and the low-fidelity model is $r_{\text{work}} = 40$.

A comparison of the convergence between single-fidelity PCE and the multifidelity PCE with an additive correction function based on adaptive refinement with generalized sparse grids is shown in Figure 8. For this problem, the multifidelity approach offers no efficiency gain over the single-fidelity PCE. This is because the reduced basis model (i.e., the low-fidelity model) interpolates the finite element model at the 50 snapshots used to generated the bases. Despite the accuracy of the reduced basis model (the maximum discrepancy between the reduced basis model and the finite element model is about 2%), its interpolatory nature results in oscillations that require a higher order PCE expansion to resolve. This example demonstrates that the performance of the multifidelity approach depends crucially on the choice of the low-fidelity model and that the ideal low-fidelity model would capture the true high order components of the high-fidelity model (resulting in low order model discrepancy) without introducing new high order components that are not present in the high-fidelity model (resulting in high order model discrepancy, with order possibly higher than the original high-fidelity model). It is worth noting that the multifidelity performance is not worse than the single-fidelity approach, despite the introduction of erroneous oscillations in the low-fidelity model. In this case, the low-fidelity and discrepancy expansions will be adapted to similar levels and the effects of the oscillations will be cancelled in the resolution of $\tilde{R}_{\text{high}}$.

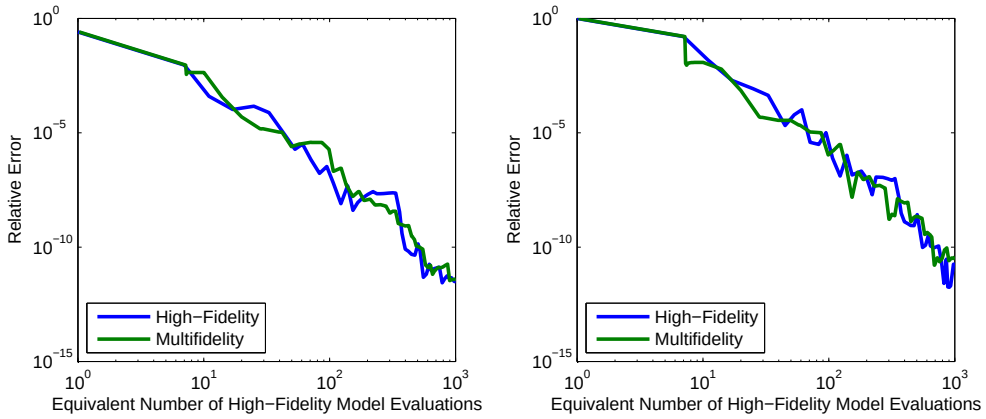

(a) Error in mean (b) Error in standard deviation

Figure 8: Convergence of single-fidelity PCE and multifidelity PCE with additive correction using adaptive generalized sparse grids for the acoustic horn example.

V. Conclusions

In this paper, we have presented a general framework for constructing stochastic expansions (non-intrusive polynomial chaos and stochastic collocation) in a multifidelity setting using a hierarchical approximation approach in which we resolve expansions for the low fidelity model and one or more levels of model discrepancy. Compared to the approach of directly estimating the statistics of the system response from a limited number of expensive high-fidelity model evaluations, greater computational efficiency can be obtained by incorporating information about the system response from a less expensive low-fidelity model at these and other locations in the stochastic space.

An adaptive multifidelity algorithm has been presented which extends the generalized sparse grid approach to consider candidate index sets from multiple sparse grids. Using normalization by relative cost of the different model fidelities, this adaptive procedure can select the refinements that provide the greatest benefit per unit cost in resolving the high fidelity statistics.

The ideal low-fidelity model for multifidelity UQ is one in which the discrepancy between the low- and high-fidelity models is a less complex function than the high-fidelity model (i.e., the spectrum of coefficients of the discrepancy expansion decays more rapidly than the high-fidelity expansion) and/or has lower variance than the high-fidelity model. Examples with good low-fidelity models (short column $R_{\text{low}1}$, elliptic PDE) have

demonstrated approximately an 80% reduction in high fidelity evaluations. In non-ideal cases where the low-fidelity model is non-informative or actually introduces erroneous high order information (horn acoustics), the multifidelity approach does not appear to degrade significantly from the single-fidelity performance, as the algorithms can fall back to reliance on resolving the original high-fidelity trends.

Acknowledgments

This collaborative work between the Sandia National Laboratories and the Massachusetts Institute of Technology is supported by the Sandia Computer Science Research Institute (CSRI) and the MURI/AFOSR Grant FA9550-09-0613.

References

[1]Alexandrov, N. M., Lewis, R. M., Gumbert, C. R., Green, L. L., and Newman, P. A., "Approximation and Model Management in Aerodynamic Optimization with Variable Fidelity Models," *AIAA Journal of Aircraft*, Vol. 38, No. 6, 2001, pp. 1093–1101.

[2]Eldred, M. S., Giunta, A. A., and Collis, S. S., "Second-Order Corrections for Surrogate-Based Optimization with Model Hierarchies," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, No. AIAA 2004-4457, Albany, NY, August 2004.

[3]Gano, S. E., Renaud, J. E., and Sanders, B., "Hybrid Variable Fidelity Optimization by Using a Kriging-Based Scaling Function," *AIAA Journal*, Vol. 43, No. 11, 2005, pp. 2422–2430.

[4]Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A., "Sequential Kriging Optimization Using Multiple-Fidelity Evaluations," *Structural and Multidisciplinary Optimization*, Vol. 32, No. 5, 2006, pp. 369–382.

[5]March, A. and Willcox, K., "Convergent Multifidelity Optimization using Bayesian Model Calibration," *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, No. AIAA 2010-9198, Fort Worth, TX, September 2010.

[6]Rajnarayan, D., Haas, A., and Kroo, I., "A Multifidelity Gradient-Free Optimization Method and Application to Aerodynamic Design," *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, No. AIAA 2008-6020, Victoria, BC, Canada, September 2008.

[7]Ghanem, R. G. and Spanos, P. D., *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, NY, 1991.

[8]Xiu, D. and Karniadakis, G. E., "The Wiener-Askey Polynomial Chaos for Stochastic Differential Equations," *SIAM Journal on Scientific Computing*, Vol. 24, No. 2, 2002, pp. 619–644.

[9]Babuška, I., Nobile, F., and Tempone, R., "A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data," *SIAM Journal on Numerical Analysis*, Vol. 45, No. 3, 2007, pp. 1005–1034.

[10]Xiu, D. and Hesthaven, J. S., "High-Order Collocation Methods for Differential Equations with Random Inputs," *SIAM Journal on Scientific Computing*, Vol. 27, No. 3, 2005, pp. 1118–1139.

[11]Constantine, P. G., Eldred, M. S., and Phipps, E. T., "Sparse Pseudospectral Approximation Method," *Computer Methods in Applied Mechanics and Engineering*, 2012, in press.

[12]Askey, R. and Wilson, J., *Some Basic Hypergeometric Orthogonal Polynomials that Generalize Jacobi Polynomials*, No. 319 in Memoirs of the American Mathematical Society, AMS, Providence, RI, 1985.

[13]Gautschi, W., *Orthogonal Polynomials: Computation and Approximation*, Oxford University Press, New York, 2004.

[14]Golub, G. H. and Welsch, J. H., "Calculation of Gauss Quadrature Rules," *Mathematics of Computation*, Vol. 23, No. 106, 1969, pp. 221–230.

[15]Witteveen, J. A. S. and Bijl, H., "Modeling Arbitrary Uncertainties Using Gram-Schmidt Polynomial Chaos," *44th AIAA Aerospace Sciences Meeting and Exhibit*, No. AIAA 2006-896, Reno, NV, January 2006.

[16]Eldred, M. S., "Design Under Uncertainty Employing Stochastic Expansion Methods," *International Journal for Uncertainty Quantification*, Vol. 1, No. 2, 2011, pp. 119–146.

[17]Barthelmann, V., Novak, E., and Ritter, K., "High Dimensional Polynomial Interpolation on Sparse Grids," *Advances in Computational Mathematics*, Vol. 12, No. 4, 2000, pp. 273–288.

[18]Gerstner, T. and Griebel, M., "Numerical Integration Using Sparse Grids," *Numerical Algorithms*, Vol. 18, No. 3, 1998, pp. 209–232.

[19]Gerstner, T. and Griebel, M., "Dimension-Adaptive Tensor-Product Quadrature," *Computing*, Vol. 71, No. 1, 2003, pp. 65–87.

[20]Bungartz, H. J. and Griebel, M., "Sparse Grids," *Acta Numerica*, Vol. 13, 2004, pp. 147–269.

[21]Jakeman, J. D. and Roberts, S. G., "Local and Dimension Adaptive Stochastic Collocation for Uncertainty Quantification," *Proceedings of the Workshop on Sparse Grids and Applications*, Bonn, Germany, May 16–20 2011.

[22]Eldred, M. S., Phipps, E. T., and Dalbey, K. R., "Adjoint enhancement within global stochastic methods," *Proceedings of the SIAM Conference on Uncertainty Quantification*, Raleigh, NC, April 2–5 2012.

[23]Adams, B. M., Bohnhoff, W. J., Dalbey, K. R., Eddy, J. P., Eldred, M. S., Hough, P. D., Lefantzi, S., Swiler, L. P., and Vigil, D. M., "DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation,

Uncertainty Quantification, and Sensitivity Analysis: Version 5.2 Theory Manual," Tech. Rep. SAND2011-9106, Sandia National Laboratories, Albuquerque, NM, 2011.

[24]Kuschel, N. and Rackwitz, R., "Two Basic Problems in Reliability-Based Structural Optimization," *Math. Method Oper. Res.*, Vol. 46, 1997, pp. 309–333.

[25]Eftang, J. L., Huynh, D. B. P., Knezevic, D. J., and Patera, A. T., "A Two-Step Certified Reduced Basis Method," *Journal of Scientific Computing*, 2012, in press.

[26]Rozza, G., Huynh, D. B. P., and Patera, A. T., "Reduced Basis Approximation and A Posteriori Error Estimation for Affinely Parametrized Elliptic Coercive Partial Differential Equations," *Archives of Computational Methods in Engineering*, Vol. 15, No. 3, 2008, pp. 229–275.