Journal of Computational Physics 341 (2017) 386–396

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Multi-fidelity stochastic collocation method for computation of statistical moments ☆

Xueyu Zhu [a], Erin M. Linebarger [b], Dongbin Xiu [c,*]

[a] Department of Mathematics, University of Iowa, Iowa City, IA 52242, USA
[b] Department of Mathematics, University of Utah, Salt Lake City, UT 84112, USA
[c] Department of Mathematics, The Ohio State University, Columbus, OH 43210, USA

ARTICLE INFO

Article history:
Received 14 June 2016
Received in revised form 6 April 2017
Accepted 7 April 2017
Available online 14 April 2017

Keywords:
Uncertainty quantification
Stochastic collocation
Multi-fidelity

ABSTRACT

We present an efficient numerical algorithm to approximate the statistical moments of stochastic problems, in the presence of models with different fidelities. The method extends the multi-fidelity approximation method developed in [18,26]. By combining the efficiency of low-fidelity models and the accuracy of high-fidelity models, our method exhibits fast convergence with a limited number of high-fidelity simulations. We establish an error bound of the method and present several numerical examples to demonstrate the efficiency and applicability of the multi-fidelity algorithm.

1. Introduction

In recent years, many efforts have been devoted to the development of efficient numerical methods for uncertainty quantification (UQ). In practical computations, the most widely used method is stochastic collocation, as it is nonintrusive sampling based and allows one to use existing deterministic codes. Unfortunately, this number of deterministic simulations required by accurate stochastic collocation methods grows very rapidly for high dimensional random inputs — the curse of dimensionality. For large scale simulations, the computational cost can become prohibitive, as the computation of each individual deterministic sample is highly costly. Many options have been investigated to tackle this challenge. For example, methods that explore more efficient sampling strategies using sparse grids, adaptivity, smoothness or sparsity of the solutions, cf. [25,2,1,11,5,8,9,12,16,17,19,20,23]. There is also a recent surge of interest in multilevel Monte Carlo method, which uses the hierarchy models by physical space refinement to achieve variance reduction in random space, cf., [14,4,3,7,21]. Other approaches to achieve variance reduction have also been presented, cf., [6,22].

In this paper we focus primarily on the computation of solution statistics using models with different fidelities. In particular, we focus on the case with one high-fidelity model and one low-fidelity model. Here, the high-fidelity model is able to produce high resolution solution to the underlying physical problem. The simulation cost is high, thus preventing us from using the standard sampling strategy (Monte Carlo, sparse grids, etc.). The low-fidelity model, on the other hand, is not highly accurate but can capture the essential behavior of the underlying problem. It is computationally cheap and can be sampled a large number of times. Typically, the low-fidelity models are constructed using simplified physics and/or

☆ This work is partially supported by AFOSR FA95501410022, DARPA N660011524053, and NSF DMS 1418771.
* Corresponding author.
E-mail addresses: xueyu-zhu@uiowa.edu (X. Zhu), aerinline@sci.utah.edu (E.M. Linebarger), xiu.16@osu.edu (D. Xiu).

http://dx.doi.org/10.1016/j.jcp.2017.04.022
0021-9991/

much coarser discretization. Examples are abundant in many problems, for example, the fine-scale versus the coarse grained models in multi-scale problems.

We present an efficient stochastic collocation algorithm for computing solution statistics using a high-fidelity model and a low-fidelity model. A distinct feature of our method is that it "seperate" the low-fidelity solutions and high-fidelity solutions. It uses the low-fidelity solutions, which consist of a large number of samples, to construct a best approximation of the target solution statistics (mean, variance, etc.), and then apply the best approximation to the high-fidelity samples. Our method is essentially a "learning" algorithm, where the low-fidelity samples are used to "train" the best approximation. It is different from most of the existing methods, which usually achieve improved performance via variance reduction or by exploring the hierarchical structure (if available) of the models. The current method is an extension of the method developed in [18,26], where the same training idea was first proposed and used to predict the solutions at arbitrary sample locations. A straightforward way to use the technique of [18,26] to compute solution statistics is to compute the bi-fidelity solutions at every sample points and then compute the statistics. Although computing each bi-fidelity solution is efficient, computing such solutions at a large number of samples becomes expensive. The major contribution of this paper is to present a mathematically equivalent algorithm that directly computes the solution statistics and bypasses the step of a large number of bi-fidelity computations. Here we show that the method can be highly efficient in approximating the statistics (mean, variance, etc.) of the underlying stochastic problem. We establish an error bound of the method and use extensive numerical examples to demonstrate its performance. In the examples with varying multiple dimensions, accuracy solutions can be obtained by $O(10)$ number of high-fidelity simulation samples.

2. Problem setup

Let w be the solution of a system of governing equations in a bounded spatial domain $D \subset \mathbb{R}^\ell$, $\ell = 1, 2, 3$, and a random parameter domain $I_Z \subseteq \mathbb{R}^d$, $d \geq 1$. For general discussion we do not assume any specific form of the governing equations. We are interested in a quantity-of-interest (QoI), which is a function of the solution w, i.e.,

$$v = q(w) : \bar{D} \times I_Z \to \mathbb{R}. \tag{2.1}$$

Hereafter we denote $x = (x_1, \ldots, x_\ell)$ the spatial variable and $z = (z_1, \ldots, z_d)$ the random variable. Let $\rho : I_z \to \mathbb{R}^+$ be the probability distribution function of z. We are interested in evaluating the statistical average of the QoI, $v : \bar{D} \to \mathbb{R}$,

$$v(\cdot) = \mathbb{E}[v] = \int v(\cdot, z)\rho(z)dz. \tag{2.2}$$

For example, when $v = w^k$, $k \geq 1$, it stands for the k-th moment of the solution.

2.1. Numerical approximations

For numerical approximation, we seek an approximate solution u in a linear subspace V for any fixed random variables,

$$u : I_z \to V. \tag{2.3}$$

Obviously, the choice of the linear subspace V depends on the chosen numerical method. We assume that the numerical method is deterministic and satisfies

$$u(\cdot, z) \approx v(\cdot, z), \qquad \forall z \in I_Z,$$

in a proper norm in the physical space.

Since the solution dependence in the random space can also be complex, the mean operator $\mathbb{E}$ in (2.2) also needs to be approximated. In this paper we focus on linear sampling based approximation, which is the predominant approach in practice. Let $\Theta = \{z_1, \ldots, z_m\} \subset I_Z$ be a set of samples, then for any integrable function $f : I_Z \to \mathbb{R}$ we define

$$\widetilde{\mathbb{E}}[f; \Theta] := \sum_{i=1}^{m} w_i f(z_i) \approx \mathbb{E}[f], \tag{2.4}$$

where w_i is the weight at the sample z_i, for $i = 1, \ldots, m$. For example, the standard Monte Carlo method has an uniform weight $w_i \equiv 1/m$, whereas for most cubature rules the weights are non-uniform. Hereafter we assume the weights satisfy

$$\sum_{i=1}^{m} w_i = 1, \qquad \|w\|_{\ell_2} < \infty, \tag{2.5}$$

where $\|w\|_{\ell_2}$ is the 2-norm of $w = (w_1, \ldots, w_m)$. Although it is highly desirable to have $w_i > 0$, this is not the case for many cubature rules.

With the approximations in both the physical space and the random space, we have

$$\mu(\cdot) = \widetilde{\mathbb{E}}[u; \Theta] \approx v(\cdot) \tag{2.6}$$

as an approximation to the true statistical average (2.2).

2.2. High and low fidelity approximations

We assume there is a high-fidelity deterministic numerical approximation

$$u^H : I_Z \to V^H, \tag{2.7}$$

in a high-fidelity Hilbert subspace V^H, equipped with inner product $\langle \cdot, \cdot \rangle^H$ and its induced norm $\| \cdot \|^H$. Similarly, we assume there is a low-fidelity approximation

$$u^L : I_Z \to V^L, \tag{2.8}$$

equipped with inner product $\langle \cdot, \cdot \rangle^L$ and norm $\| \cdot \|^L$. We assume that u^H is much more accurate than u^L, but also much more time consuming to obtain.

Let $\Gamma_N = \{z_1, \cdots, z_N\} \subset I_Z$ be a set of dense samples, from which an accurate approximation of the expectation $\widetilde{\mathbb{E}}$ (2.4) can be obtained. We assume $N \gg 1$. We then readily have the low-fidelity estimation of the mean

$$\mu^L(\cdot) = \widetilde{\mathbb{E}}[u^L; \Gamma_N] = \sum_{i=1}^{N} w_i u^L(\cdot, z_i), \tag{2.9}$$

which is not highly accurate due to the limited accuracy of u^L at each sample points. The high-fidelity approximation $\widetilde{\mathbb{E}}[u^H]$

$$\mu^H(\cdot) = \widetilde{\mathbb{E}}[u^H; \Gamma_N] = \sum_{i=1}^{N} w_i u^H(\cdot, z_i) \tag{2.10}$$

should be accurate, but it is not computable due to the exceedingly large simulation cost of $N \gg 1$. Our goal is to construct a bi-fidelity approximation $\mu^B \approx \mu^H$, by using the low-fidelity mean (2.9) and a limited number of high-fidelity samples.

We remark that the method developed in [18,26] allows one to compute bi-fidelity solutions at any given sample point, i.e., $u^B(z) \approx u^H(z), z \in I_Z$. It is then natural to approximate the high-fidelity mean by using the bi-fidelity solutions, i.e.,

$$\widetilde{\mu}^B(\cdot) = \widetilde{\mathbb{E}}[u^B; \Gamma_N] = \sum_{i=1}^{N} w_i u^B(\cdot, z_i).$$

This straightforward approach requires the bi-fidelity approximations of the solution at the dense sample set Γ_N and can become computationally expensive. In the following, we shall present a mathematically equivalent method that directly constructs the bi-fidelity mean and avoids the bi-fidelity approximation at the dense set Γ_N. The new method is thus much more efficient.

3. Bi-fidelity method for expectation

In this section we describe the bi-fidelity algorithm for computing the expectation (2.4). The method is an extension of the work of [18,26], which developed a multi-fidelity approximation of the solution at any sample location z.

3.1. The main algorithm

Again, let $\Gamma_N = \{z_1, \ldots, z_N\} \subset I_Z$ with $N \gg 1$ be a set of dense samples, from which an accurate approximation of the expectation $\widetilde{\mathbb{E}}$ (2.4) can be obtained. Let

$$\gamma_n = \{z_{i_1}, \ldots, z_{i_n}\} \subset \Gamma_N \tag{3.1}$$

be a subset of samples. The general procedure of our bi-fidelity algorithm is as follows.

(1) Conduct the low-fidelity sampling at the dense sample set Γ_N and obtain the low-fidelity approximation $\mu^L(\cdot)$ via (2.9).
(2) Select a subset of samples $\gamma_n \subset \Gamma_N$, where $n \ll N$.
(3) Conduct the high-fidelity computations at the subset γ_n and obtain the high-fidelity solution samples, $u^H(\gamma_n) = \{u^H(\cdot, z_{i_1}), \ldots, u^H(\cdots, z_{i_n})\}$.
(4) Construct a bi-fidelity approximation $\mu^B = \mu^B(\mu^L, u^H(\gamma_n))$.

The key issue lies in Step (2), the selection of the subset γ_n, and Step (4), the construction of the bi-fidelity approximation μ^B. We now present the detail of these two steps.

3.1.1. Selection of the subset γ_n

For any set of samples $\delta_k = \{z_1, \ldots, z_k\}$, we denote

$$u^L(\delta_k) = \left\{u^L(z_1), \ldots, u^L(z_k)\right\}, \tag{3.2}$$

and define the corresponding space

$$U^L(\delta_k) = \text{span}(u^L(\delta_k)) = \text{span}\{u^L(z_1), \cdots, u^L(z_k)\}. \tag{3.3}$$

Similarly, for the high-fidelity model, we define

$$u^H(\delta_k) = \left\{u^H(z_1), \ldots, u^H(z_k)\right\}, \qquad U^H(\delta_k) = \text{span}(u^H(\delta_k)). \tag{3.4}$$

Our choice of the n-point subset γ_n follows the procedure proposed in [18,26]. It is a greedy algorithm that seeks to find the next sample whose corresponding low-fidelity solution is furthest to the space spanned by the existing low-fidelity solution set. Starting from a trivial initial choice $\gamma_0 = \{\}$, we let $\gamma_k = \{z_{i_1}, \ldots, z_{i_k}\} \subset \Gamma_N$ be the k-point existing subset in Γ_N. We then find the $(k+1)$-th point by

$$z_{i_{k+1}} = \underset{z \in \Gamma_N}{\text{argmax}}\, \text{dist}(u^L(z), U^L(\gamma_k)), \qquad \gamma_{k+1} = \gamma_k \cup z_{i_{k+1}}, \tag{3.5}$$

where the distance function $\text{dist}(g, G)$ between the function $g \in u^L(\Gamma_N)$ and the space $G \subset U^L(\Gamma_N)$ follows the standard definition. This greedy algorithm can be readily implemented via simple operations of numerical linear algebra.

- Let $\mathbf{W}$ be the Gramian matrix of the low-fidelity solution $u^L(\Gamma_N)$, i.e.,

$$\mathbf{W} = (w_{ij})_{1 \le i,j \le N}, \qquad w_{ij} = \left\langle u^L(z_i), u^L(z_j) \right\rangle^L. \tag{3.6}$$

- Apply the pivoted Cholesky decomposition to the matrix $\mathbf{W}$,

$$\mathbf{W} = \mathbf{P}^T \mathbf{L}\mathbf{L}^T \mathbf{P}, \tag{3.7}$$

  where $\mathbf{L}$ is lower-triangular and $\mathbf{P}$ is a permutation matrix due to pivoting. This will produce an ordered permutation vector $P = (i_1, \ldots, i_N)$, from which we choose the first n points to define $\gamma_n = \{z_{i_1}, \ldots, z_{i_n}\}$.

More details and properties of the algorithm can be found in [18,26].

3.1.2. Bi-fidelity approximation

Once the steps (1)–(3) are finished in the main algorithm in Section 3.1, we possess the following

- Low-fidelity solutions at the dense sample set Γ_N, $u^L(\Gamma_N)$, and the corresponding low-fidelity average, μ^L (2.9).
- The subset samples $\gamma_n \subset \Gamma_N$.
- High-fidelity solutions in the subset γ_n, $u^H(\gamma_n)$.

To construct the bi-fidelity approximation of the average, we first construct the best approximation of μ^L in the low-fidelity space $U^L(\gamma_n)$, and then apply the same construction to the high-fidelity space $U^H(\gamma_n)$. To accomplish this, we first construct the orthogonal projection operator that maps the low-fidelity solution space $U^L(\gamma_n)$ onto the mean μ^L. That is, we define

$$\mathbb{P}[U^L(\gamma_n); \mu^L] := \sum_{k=1}^{n} c_k u^L(\cdot, z_{i_k}) \approx \mu^L(\cdot), \tag{3.8}$$

where the coefficients are determined by solving, for each sample $z_{i_j} \in \gamma_n$, $j = 1, \ldots, n$,

$$\sum_{k=1}^{n} \langle u^L(z_{i_k}), u^L(z_{i_j}) \rangle^L c_k = \langle \mu^L, u^L(z_{i_j}) \rangle^L, \qquad j = 1, \ldots, n, \tag{3.9}$$

which is a linear system of equations. The notation $\mathbb{P}[U^L(\gamma_n); \mu^L]$ implies this operator is constructed by using μ^L and U^L constrained on the set γ_n.

Our bi-fidelity approximation of the average is then defined by

$$\mu^B(\cdot) := \mathbb{P}[U^H(\gamma_n); \mu^L] = \sum_{k=1}^{n} c_k u^H(\cdot, z_{i_k}), \tag{3.10}$$

where the coefficients c_k are computed in (3.9). This implies that the bi-fidelity approximation is the output of the same operator $\mathbb{P}$, which is constructed using μ^L, with $U^H(\gamma_n)$ as input.

3.2. Error analysis

The algorithm is a straightforward extension of the bi-fidelity method in [18,26], which seeks to approximate the solution at arbitrary parameter location of $z \in I_Z$. (Note that the method in [18,26] does not require probability.) For any $z \in I_Z$, we denote

$$\epsilon(z) = \left\| u^H(z) - \mathbb{P}[U^H(\gamma_n); u^L(z)] \right\|^H < \infty \tag{3.11}$$

the error bound of the bi-fidelity approximation. The bound was established in Theorem 4.4 of [18]. Its detailed expression is difficult to summarize without invoking too many notations unnecessary for this paper. Interested reader should consult Theorem 4.4 of [18].

Based upon this result, we derive the following error estimate for the bi-fidelity average $\mu^B(\cdot)$, compared to the high-fidelity average $\widetilde{\mathbb{E}}[u^H; \Gamma_N]$ (which is not feasible to compute).

Theorem 3.1. *Let μ^B be the bi-fidelity mean (3.10) and μ^H be the high-fidelity mean (2.10). If the assumptions in Theorem 4.4 of [18] are satisfied, then*

$$\left\| \mu^B(\cdot) - \mu^H(\cdot) \right\|^H \le C_w \|\epsilon(z)\|_{\ell^2(\Gamma_N)}, \tag{3.12}$$

where $\epsilon(z)$ is defined in (3.11) and $C_w = \|w\|_{\ell_2}$ is the 2-norm of the weight vector $w = (w_1, \ldots, w_N)$. Furthermore, if all weights are positive, i.e., $w_i > 0$, $i = 1, \ldots, N$, then

$$\left\| \mu^B(\cdot) - \mu^H(\cdot) \right\|^H \le \max_{z \in \Gamma_N} \epsilon(z). \tag{3.13}$$

Proof. The linearity of the operator $\mathbb{P}$ immediately gives us

$$\mu^B(\cdot) = \mathbb{P}[U^H(\gamma_n); \mu^L] = \mathbb{P}\left[U^H(\gamma_n,); \widetilde{\mathbb{E}}[u^L; \Gamma_N]\right] = \widetilde{\mathbb{E}}\left[\mathbb{P}\left[U^H(\gamma_n); u^L\right]; \Gamma_N\right].$$

Then,

$$\|\mu^B - \mu^H\|^H = \left\| \widetilde{\mathbb{E}}\left[\mathbb{P}\left[U^H(\gamma_n); u^L\right] - u^H; \Gamma_N\right] \right\|^H = \left\| \sum_{i=1}^N w_i \left(\mathbb{P}[U^H(\gamma_n); u^L(z_i)] - u^H(z_i) \right) \right\|^H.$$

A straightforward use of the Cauchy–Schwarz inequality completes the proof. □

4. Numerical examples

In this section, we present several numerical examples to illustrate the effectiveness and efficiency of our method. For benchmarking purpose, all examples have relatively small computational cost so that we can run the high-fidelity model many times to compute the reference solution μ^H, from which we compute the numerical errors of the bi-fidelity solutions. The errors are reported as the standard L^2 norm in physical space. Without loss of generality, we employ uniformly distributed random variables in all the examples.

4.1. Function example

We first consider a simple example using a function with known analytical form,

$$u(x, z) = g(x, z + \epsilon z^2) = \cos(x(z + \epsilon z^2)), \qquad (x, z) \in [-1, 1] \times [0, 10\pi]. \tag{4.1}$$

To approximate this function we employ polynomials and choose the linear space V to be Π_K, the linear space of polynomial of degrees up to K. The orthogonal projection of g has the following analytical form,

$$g_K(x, z) := \mathbb{P}_{\Pi_K} g = \sum_{k=0}^K \hat{g}_k(z) \tilde{L}_k(x), \tag{4.2}$$

where $\tilde{L}_k$ is the normalized Legendre polynomial,

$$\hat{g}_k = c_k \sqrt{\frac{\pi(2k+1)}{|z|}} J_{k+1/2}(|z|), \qquad z \neq 0, \qquad c_k = \Re[e^{(\mathrm{isgn}z)} i^k],$$

and J_k is the k-th orde Bessel function of the first kind.

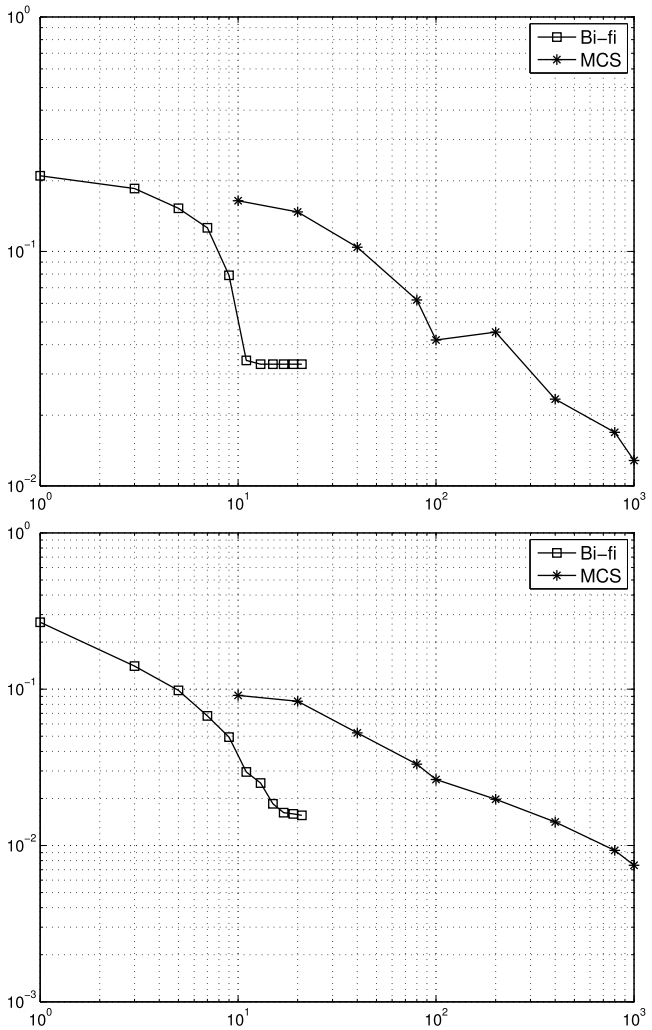

Fig. 4.1. Function example: Decay of errors in the mean (top) and the second moment (bottom) with respect to the number of high-fidelity simulations by the bi-fidelity algorithm and the standard Monte Carlo method.

We choose the low-fidelity model to be g_{35} and high-fidelity model to be g_{100}. The approximation space in physical space V are the piecewise linear polynomials with 500 uniform grids. That is, $V^L = V^H$. The dense sampling set $\Gamma_N = 1,000$ Monte Carlo points.

In Fig. 4.1 the numerical errors in the first moment (mean) and second moment by the bi-fidelity algorithm are shown, with respect to the number of high-fidelity samples. We clearly observe the very fast error decay with only about 10 samples. For reference the error convergence of the high-fidelity Monte Carlo method (MCS) is also shown, which has the standard $N^{-1/2}$ convergence rate. The bi-fidelity method exhibits drastic improvement in performance over the standard MCS. Note that the error of the bi-fidelity method will saturate at about $O(10)$ high-fidelity samples. This is expected as the bi-fidelity method relies on the mapping between $U^L(\gamma_n)$ and μ^L, as shown in (3.8), whose accuracy is limited by the quality of the low-fidelity model.

4.2. One dimensional stochastic elliptic equation

We now consider an elliptic equation with random diffusivity, a standard benchmark problem in stochastic computing.

$$\begin{cases} -(a(x,z)u_x(x,z))_x = 1, & (x,z) \in (0,1) \times [-1,1]^d, \\ u(0,z) = 0, & u(1,z) = 0, \end{cases} \tag{4.3}$$

where the diffusivity field is modeled as

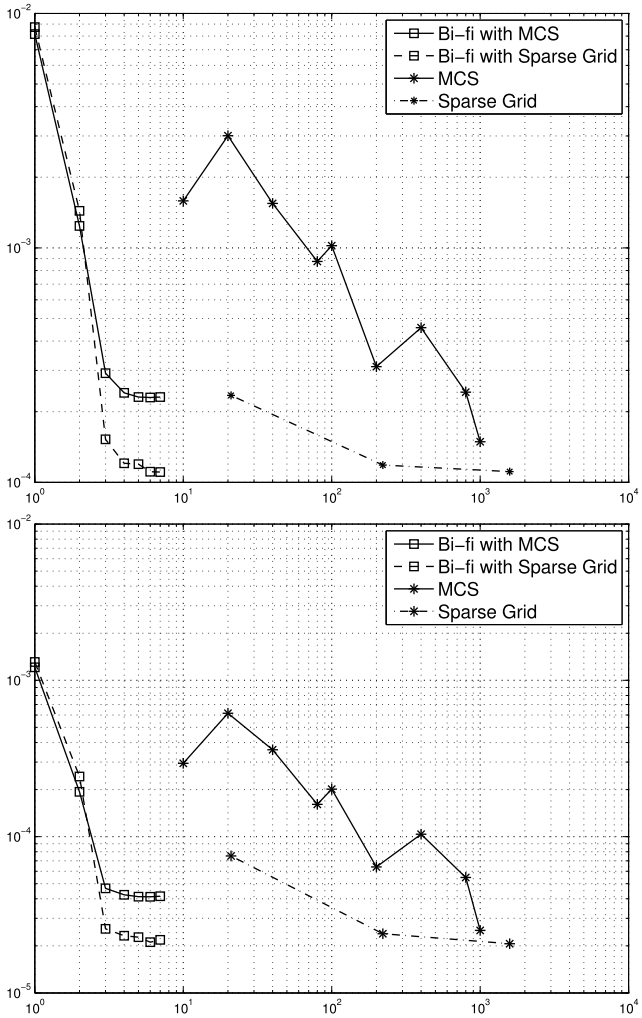

Fig. 4.2. One-dimensional diffusion problem in $d = 10$: Numerical errors in the mean (top) and second moment (bottom) with respect to the number of high-fidelity simulations by the standard Monte Carlo method, the sparse grids method, bi-fidelity method based on Monte Carlo, and bi-fidelity method based on sparse grids.

$$a(x,z) = 1 + \sigma \sum_{k=1}^{d} \frac{1}{(k\pi)^2} \cos(2\pi kx) z_k, \qquad d > 1. \tag{4.4}$$

Here we let $d = 10$ and $\sigma = 4$. (The positivity of a is strictly enforced for the choice of the parameters.)

The Chebyshev collocation method is employed to solve the problem in physical space. The low-fidelity model are the solution using 8-point Chebyshev collocation method; whereas the high-fidelity model is based on 128-point Chebyshev collocation method. Consequently, the low-fidelity approximation space V^L is different from the high-fidelity approximation space V^H. Two dense sampling sets are employed to compute the accurate mean. One is the random Monte Carlo sample set with $N = 10,000$ points, and the other one is the level 3 Clenshaw–Curtis sparse grids set ([25]). At $d = 10$, the sparse grids set has $N = 1,581$ points, which have non-uniform and even some negative weights. This is to illustrate the fact that the bi-fidelity method can be applied to any low-fidelity model using the weighted sum form (2.9) for computing the mean. Having these two ways to compute the statistical averages over the dense sets, we thus have two versions of bi-fidelity algorithms.

The convergence of the errors in mean and second moment is shown in Fig. 4.2 with respect to increasing number of high-fidelity simulations, for both the bi-fidelity algorithm based on Monte Carlo simulation and the bi-fidelity method based on sparse grids. For comparison we also show the convergence of the Monte Carlo method and the sparse grids method. Again we observe very fast error decay by the bi-fidelity algorithms. The errors become exceedingly small after only less than 10 high-fidelity samples. This is a drastic improvement over the standard MCS and sparse grids.

4.3. Two dimensional stochastic elliptic equation

We now consider the following 2D stochastic elliptic equation:

$$\begin{cases} -\nabla \cdot (a(x, y, z)\nabla u) = 0, & (x, y) \in (-1, 1)^2, \\ u(-1, y, z) = -1, \quad u(1, y, z) = 1, & u_y(x, -1, z) = 0, \quad u_y(x, 1, z) = 0, \end{cases} \tag{4.5}$$

where the diffusivity field is modeled via the following Karhunen–Loeve expansion,

$$a(x, y, z) = 1 + \sum_{k=1}^{d} \sqrt{\lambda_k}\psi_k(x, y)z_k. \tag{4.6}$$

We let $z \in (-1, 1)^d$ be uniformly distributed random variables and let $(\lambda_k, \psi_k)_{k=1}^d$ be the eigen-pairs in the Karhunen–Loeve expansion of a random process with the following covariance function

$$C(x_1, y_1, x_2, y_2) = \sigma \exp(-|x_1 - x_2| - |y_1 - y_2|),$$

with $\sigma = 0.3$. The KL expansion has explicitly computable eigenvalue-eigenfunction pairs based on its one-dimensional counterpart (cf., [13,24]). We truncate the expansion at $d = 17$, which keeps 91.4% of the total spectral energy.

We construct the following models:

- *Model 1*. This is a one-dimensional approximation to the two-dimensional problem (4.5). Due to the boundary conditions in (4.5), the problem exhibits certain "near" symmetry. Consequently, we define

$$\begin{cases} -(\hat{a}(x, z)u_x(x, z))_x = 0, & x \in (-1, 1), \\ u(-1, z) = -1, & u(1, z) = 1, \end{cases} \tag{4.7}$$

  where $\hat{a}(x, z) = a(x, 0, z) = 1 + \sum_{k=1}^{d} \sqrt{\lambda_k}\psi_k(x, 0)z_k$ is the one-dimensional version (at $y = 0$) of (4.6). We then employ the P_1-finite element method with 80 uniform elements to solve (4.7).
- *Model 2*. This is the original two-dimensional problem (4.5), where we employ the P_1 finite elements with a very coarse mesh − 128 uniform triangular elements with size $h = 1/4$.
- *High-fidelity model*. Here we employ the P_1 finite element method with a uniform triangular mesh of 12,800 elements with size $h = 1/40$. This is our high-fidelity model and can well resolve the solution.

We then examine the following bi-fidelity cases.

- *Bi-fidelity 1*. Here we let *Model 1* be the low-fidelity model and construct bi-fidelity approximation in conjunction with the high-fidelity model.
- *Bi-fidelity 2*. Here we let *Model 2* be the low-fidelity model and conduct the bi-fidelity approximation with the high-fidelity model.

In both cases, the dense point set Γ_N has $N = 1,000$ i.i.d samples.

The numerical errors of mean and second moment of "bi-fidelity 1" and "bi-fidelity 2" approximations are plotted in Fig. 4.3. We first notice the fast error decay of the "bi-fidelity 2" case, whose errors decay nearly exponentially fast and reach very low level with only a few high-fidelity simulation runs. This indicates that *Model 2*, despite its coarse mesh ($h = 1/4$), is able to capture the solution variation in the random space reasonably well. On the contrary, the errors in "bi-fidelity 1" saturate at a higher level after a few high-fidelity runs. This is not surprising, as the one dimensional low-fidelity model used in the *Bi-fidelity 1* algorithm is rather over-simplified and has lower accuracy. Again, we observe both bi-fidelity approximations show better convergence behavior than the standard Monte Carlo approach.

4.4. Acoustic horn problem

We now present results for the two-dimensional Helmholtz equation. It is an acoustic horn problem from [26]. The acoustic field is described by the following time-harmonic Helmholtz equation,

$$\begin{cases} \triangle u + 4u = 0, & \\ (2i + 1/25)u + \frac{\partial u}{\partial n} = 0, & \partial D_{out}, \\ 2iu + \frac{\partial u}{\partial n} = 4i, & \partial D_{in}, \\ \frac{\partial u}{\partial n} = 0, & \partial D_j,\ j = 3, 8, \\ i\kappa_j u + \frac{\partial u}{\partial n} = 0, & \text{on other boundaries}, \end{cases} \tag{4.8}$$

where $i^2 = -1$ and $\mu = (\mu_1, \mu_2, \mu_4, \mu_5, \mu_6, \mu_7, \mu_9, \mu_{10}) \in [0, 1]^8$ is modeled as a uniformly distributed random parameter governing the material properties. The geometry is shown in Fig. 4.4, where $R = 12.5$, $a = 0.5$, $c = 0.1$, $d = 5$, $l = 5$. The

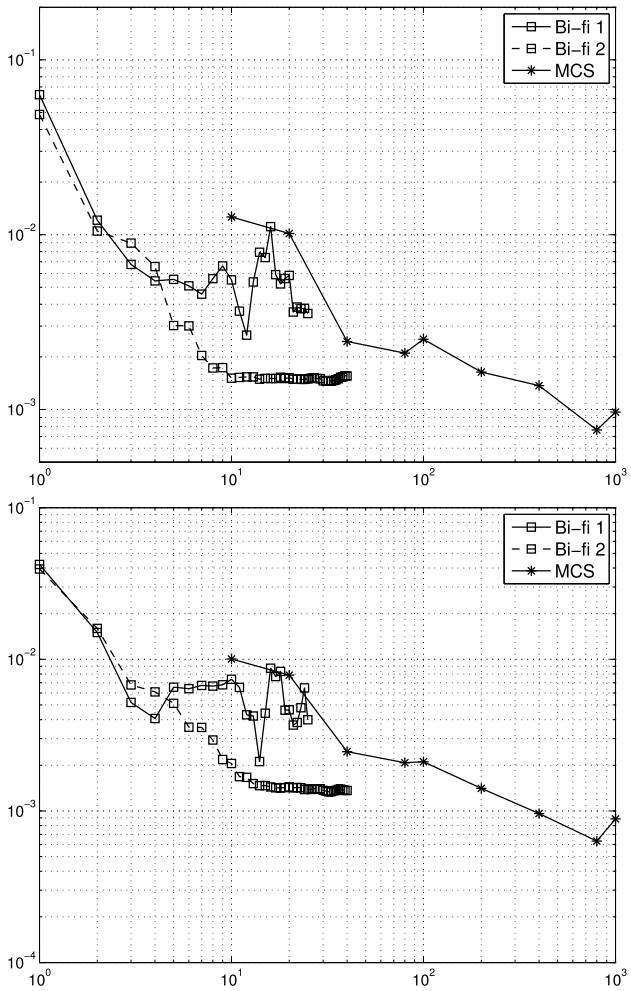

Fig. 4.3. Two dimensional diffusion problem (4.7): Numerical errors of the mean (top) and second moment (bottom) with respect to the number of high-fidelity simulations by the standard Monte Carlo method, the "bi-fidelity 1" and "bi-fidelity 2" method.

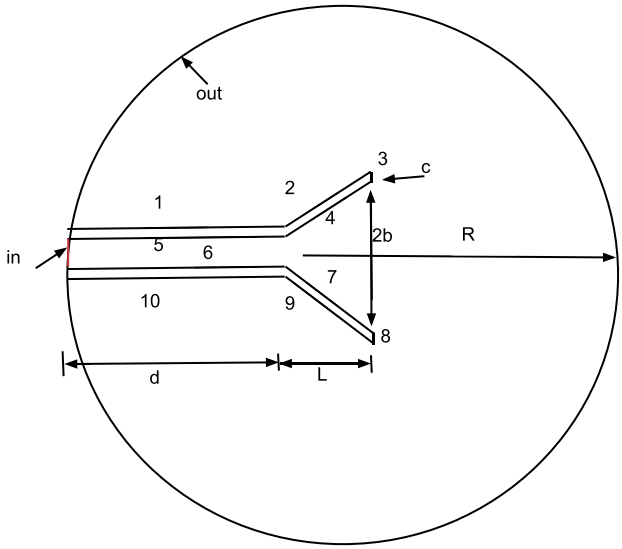

Fig. 4.4. The domain of the acoustic horn problem.

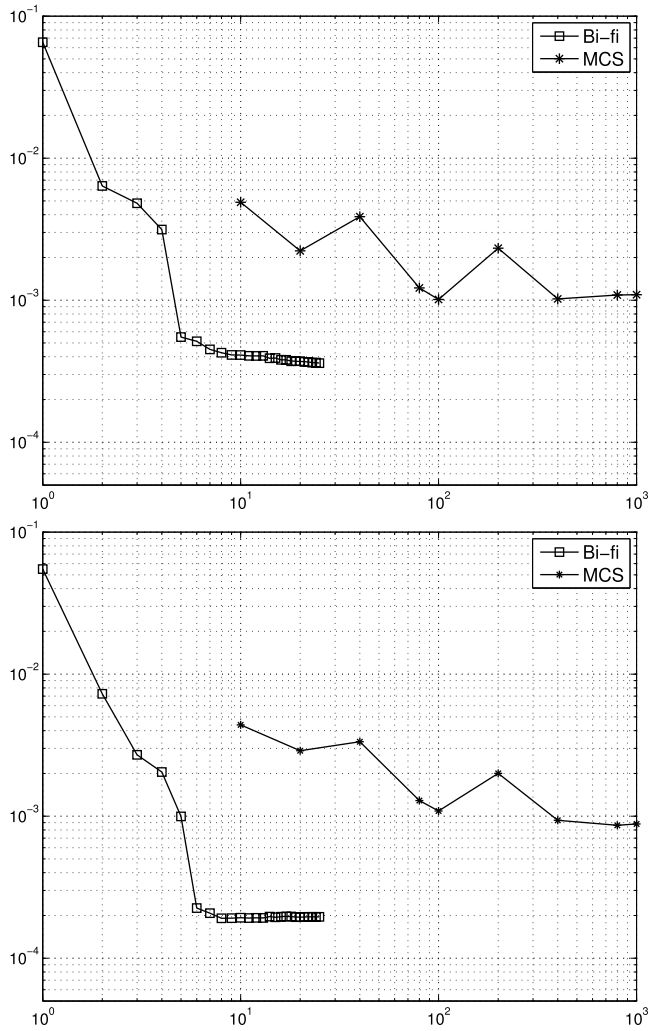

Fig. 4.5. Acoustic horn problem: Numerical errors in mean (top) and second moment (bottom) with respect to the number of high-fidelity simulations by standard Monte Carlo method and bi-fidelity method.

boundary condition at ∂D_{out} is the lowest order Enquist–Majda absorbing boundary condition [10], to reduce the reflection effect due to the artificial boundary. The boundaries $\partial D_{3,8}$ consist of sound-hard materials and the rest of ∂D_i mimic sound-soft materials determined by the parameters κ_i.

We solve this problem using P_2 bubble finite elements using the finite element package Freefem++ [15]. The *low-fidelity model* uses a mesh of $2,061$ elements; and the *high-fidelity model* uses a fine mesh of $22,810$ elements. The dense point set Γ_N has $N = 5,000$ i.i.d. random points.

The convergence results are shown in Fig. 4.5. Again, we observe very fast error convergence in the bi-fidelity simulations. With less than $n = 10$ high-fidelity simulations, the errors reduce to the level of the spatial discretization errors. The improvement over the standard Monte Carlo method is obvious.

5. Summary

In this paper we present a bi-fidelity algorithm for approximating the statistical moments of stochastic problems. The method extends the earlier work by [18,26], which approximates solutions at given sample locations. The current bi-fidelity method takes advantage of the low cost of the low-fidelity model and the high accuracy of the high-fidelity model. It uses the low-fidelity average and a limited number of high-fidelity simulations to construct a highly accurate approximation of the high-fidelity average. Error bound is established to ensure the method is well behaved. Numerous examples are presented to demonstrate the effectiveness of the bi-fidelity method. In most cases, the current bi-fidelity method is able to produce highly accurate results with only $O(10)$ high-fidelity samples.

References

[1] N. Agarwal, N. Aluru, A domain adaptive stochastic collocation approach for analysis of MEMS under uncertainties, J. Comput. Phys. 228 (20) (2009) 7662–7688.
[2] I. Babuska, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, SIAM J. Numer. Anal. 45 (3) (2007) 1005–1034.
[3] A. Barth, A. Lang, C. Schwab, Multilevel Monte Carlo method for parabolic stochastic partial differential equations, BIT Numer. Math. 53 (1) (2012) 3–27.
[4] A. Barth, C. Schwab, N. Zollinger, Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients, Numer. Math. 119 (1) (2011) 123–161.
[5] M. Bieri, C. Schwab, Sparse high order FEM for elliptic sPDEs, Comput. Methods Appl. Mech. Eng. 198 (2009) 1149–1170.
[6] P. Chen, A. Quarteroni, G. Rozza, Reduced order methods for uncertainty quantification problems, 2015.
[7] K. Cliffe, M. Giles, R. Scheichl, A.L. Teckentrup, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, Comput. Vis. Sci. 14 (1) (2011) 3–15.
[8] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, J. Comput. Phys. 230 (8) (2011) 3015–3034.
[9] M. Eldred, Recent advances in non-intrusive polynomial chaos and stochastic collocation methods for uncertainty analysis and design, in: 50th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2009, AIAA-2009-2249.
[10] B. Engquist, A. Majda, Absorbing boundary conditions for numerical simulation of waves, Proc. Natl. Acad. Sci. 74 (5) (1977) 1765–1766.
[11] J. Foo, X. Wan, G. Karniadakis, The multi-element probabilistic collocation method (ME-PCM): error analysis and applications, J. Comput. Phys. 227 (22) (2008) 9572–9595.
[12] B. Ganapathysubramanian, N. Zabaras, Sparse grid collocation methods for stochastic natural convection problems, J. Comput. Phys. 225 (1) (2007) 652–685.
[13] R. Ghanem, P. Spanos, Stochastic Finite Elements: a Spectral Approach, Springer-Verlag, 1991.
[14] M. Giles, Multilevel Monte Carlo path simulation, Oper. Res. 56 (3) (2008) 607–671.
[15] F. Hecht, New development in Freefem++, J. Numer. Math. 20 (3–4) (Jan. 2012).
[16] J. Jakeman, S. Roberts, Stochastic Galerkin and collocation methods for quantifying uncertainty in differential equations: a review, in: G. Mercer, A. Roberts (Eds.), Proceedings of the 14th Biennial Computational Techniques and Applications Conference, CTAC-2008, in: ANZIAM J., vol. 50, Feb. 2009, pp. C815–C830.
[17] X. Ma, N. Zabaras, An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations, J. Comput. Phys. 228 (2009) 3084–3113.
[18] A. Narayan, C. Gittelson, D. Xiu, A stochastic collocation algorithm with multifidelity models, SIAM J. Sci. Comput. 36 (2) (2014) A495–A521.
[19] A. Narayan, D. Xiu, Stochastic collocation methods on unstructured grids in high dimensions via interpolation, SIAM J. Sci. Comput. 34 (3) (2012) A1729–A1752.
[20] F. Nobile, R. Tempone, C. Webster, A sparse grid stochastic collocation method for partial differential equations with random input data, SIAM J. Numer. Anal. 46 (5) (2008) 2309–2345.
[21] A. Teckentrup, R. Scheichl, M. Giles, E. Ullmann, Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients, Numer. Math. 125 (3) (2013) 569–600.
[22] F. Vidal-Codina, N.-C. Nguyen, M. Giles, J. Peraire, A model and variance reduction method for computing statistical outputs of stochastic elliptic partial differential equations, preprint, arXiv:1409.1526, 2014.
[23] D. Xiu, Efficient collocational approach for parametric uncertainty analysis, Commun. Comput. Phys. 2 (2) (2007) 293–309.
[24] D. Xiu, Numerical Methods for Stochastic Computations, Princeton University Press, Princeton, New Jersey, 2010.
[25] D. Xiu, J. Hesthaven, High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput. 27 (3) (2005) 1118–1139.
[26] X. Zhu, A. Narayan, D. Xiu, Computational aspects of stochastic collocation with multifidelity models, SIAM/ASA J. Uncertain. Quantificat. 2 (1) (2014) 444–463.