

Multi-Element Probabilistic Collocation Method in High Dimensions

Jasmine Foo and George Em Karniadakis*

Division of Applied Mathematics, Brown University, Providence, RI 02912 USA

Abstract

We combine multi-element polynomial chaos with analysis of variance (ANOVA) functional decomposition to enhance the convergence rate of polynomial chaos in high dimensions and in problems with low stochastic regularity. Specifically, we employ the multi-element probabilistic collocation method MEPCM [1] and so we refer to the new method as MEPCM-A. We investigate the dependence of the convergence of MEPCM-A on two decomposition parameters, the polynomial order μ and the effective dimension ν , with $\nu \ll N$, and N the nominal dimension. Numerical tests for multidimensional integration and for stochastic elliptic problems suggest that $\nu \geq \mu$ for monotonic convergence of the method. We also employ MEPCM-A to obtain error bars for the piezometric head at the Hanford nuclear waste site under stochastic hydraulic conductivity conditions. Finally, we compare the cost of MEPCM-A against Monte Carlo in several hundred dimensions, and we find MEPCM-A to be more efficient for up to 600 dimensions for a specific multi-dimensional integration problem involving a discontinuous function.

1 Introduction

Partial differential equations with uncertain or stochastic parameters arise in many problems from engineering, biology, and various other fields. In many cases, using standard Monte Carlo techniques to calculate moment statistics of the solutions to these problems is prohibitively computationally expensive. In this paper we concentrate on a class of methods, called *stochastic spectral methods*, which are designed to efficiently calculate moments of solutions to systems with parametric uncertainty (see, e.g. [2–14] and references therein). These methods can provide considerable speed-up in computational

* Corresponding author.

Email address: `gk@dam.brown.edu` (George Em Karniadakis).

time when compared to Monte Carlo (MC) simulation for many problems of low to moderate random dimension. There are still, however, several important computational limitations associated with these techniques. Since the random inputs in many problems are stochastic processes approximated with truncated Karhunen-Loève expansions, the dimensionality of these inputs is dependent on the correlation length of these processes. For input processes with relatively low correlation length, the number of dimensions required for accurate representation can be extremely large. However, Karhunen-Loève inputs with moderately high dimensionality (e.g., greater than 10 dimensions) pose major computational challenges for stochastic spectral methods, which must utilize approximation techniques in the high-dimensional stochastic space. Low stochastic regularity of the solution poses another significant computational challenge. In this case, accurately capturing the solution requires a prohibitively high-order representation in the random space.

To deal with low stochastic regularity of solutions, domain decomposition in the random space for the stochastic Galerkin method was proposed in [2]. This method was then generalized in [15] to deal with arbitrary probability distributions with the numerical construction of generalized polynomial chaos bases on the fly. This approach, called the Multi-Element Generalized Polynomial Chaos (ME-gPC) method was found to effectively deal with problems exhibiting low regularity in parametric space as well as for long-time integration [15]. An analogous collocation-based method called the Multi-Element Probabilistic Collocation Method (ME-PCM) was developed and analyzed in [1]. In addition, Galerkin-based methods utilizing multi-resolution wavelet expansions [16,17] were developed to deal with solutions exhibiting discontinuities or sharp variation in the stochastic space. However, in high random dimensions, these methods remain too computationally expensive to be viable.

Collocation-based stochastic spectral methods were first introduced in [12], [18]. The use of quadrature formulae was introduced by Le Maitre and colleagues in [19] and later considered in [4]. The use of Lagrange interpolants in stochastic collocation was explored in [6] and [13], where convergence analysis for elliptic equations was performed. Since moment estimation in stochastic collocation essentially translates to numerical integration, this method benefits directly from advances in the field of high-dimensional integration and interpolation. In 2003, Matthies and Keese proposed the use of *sparse grid quadrature* for stochastic collocation [20]. Classical sparse grids, introduced by Smolyak [21] in 1963, are constructed from tensor products of one-dimensional quadrature formulas. Of all possible combinations of one-dimensional formulas, only the ones whose corresponding indices fall within the unit simplex are considered in the Smolyak construction. Sparse grids utilize the smoothness of the integrand to weaken the “curse of dimensionality” for certain classes of functions (e.g., functions with bounded mixed derivatives) in low to moderate dimensions. The errors and efficiency of sparse grid integration and in-

terpolation have been investigated extensively in [22–25]. In particular, in [22–24] the errors of integrating or interpolating functions with Sobolev regularity are analyzed for Smolyak constructions based on one-dimensional nested Clenshaw-Curtis rules. In [22] the degree of exactness of Smolyak quadrature using Clenshaw-Curtis and Gaussian one-dimensional rules is investigated. In [25] the integration error of sparse grids based on one-dimensional Kronrod-Patterson rules is bounded for analytic integrands. Some of these results have been used in recent works in analyzing the error bounds of sparse grid stochastic collocation.

In 2005, Xiu and Hesthaven [4] numerically demonstrated the efficiency of Clenshaw Curtis-based sparse grid stochastic collocation in comparison to other methods on an elliptic problem with low to moderate (< 50) dimensions. This method was extended to include adaptive refinement in stochastic dimensions and applied to stochastic natural convection problems in [14]. The method was also analyzed by Nobile et. al. [26] in 2006 for a linear elliptic PDE with random input. There, the sparse grid interpolation error bound developed in [23] was used to bound the stochastic discretization error and strong error estimates for the fully discrete solution were shown. Numerical experiments performed on the stochastic elliptic problem with 1 to 11 dimensions demonstrated the efficiency of this method in comparison to Monte Carlo and tensor product stochastic collocation.

Despite the considerable improvements in efficiency of the sparse grid collocation method over Monte Carlo and tensor product methods, the complexity estimate of sparse grids still depends heavily on the *dimension* and on the *regularity* of the integrand [24]. Thus, if the number of dimensions grows beyond moderate levels the convergence rate suffers increasingly. Similarly, solutions with low stochastic regularity will further hinder the convergence of the method. To push the dimensionality limitation higher, some variations of the sparse grid methods have been introduced.

One particularly interesting variation is the ‘dimension-adaptive tensor product quadrature’ method introduced by Gerstner and Griebel [27] in 2003. There it was recognized that in some applications the different dimensions of the integration problem are of *varying* importance. This is often the case for random inputs from Karhunen-Loève decompositions of random processes, or parametric sensitivity analysis problems. In [27] the authors utilized a generalized sparse grid formulation introduced in [28–30], where the space of possible combinations of one-dimensional formulas (index sets) was increased to include indices outside the unit simplex, which weights all dimensions equally. These generalized index sets can be chosen to provide greater resolution in selected dimensions only, according to chosen weights. However, it is not always known a priori which dimensions are of greater importance than others. The authors also developed a dimension-adaptive method to assess the importance of each

dimension by considering the selection of the index set as a binary knapsack optimization problem, and they demonstrated its effectiveness on integration problems with up to 256 dimensions.

The concept of dimensionally anisotropic sparse grids is certainly promising for stochastic collocation methods. However, the optimization problem used in the adaptive method of [27] involves calculating the integral for sparse grids resolved at one higher level in each dimension in order to assess the relative importance. This can be costly for stochastic collocation problems, where the calculation at each sample point involves the solution of a deterministic PDE. Thus, other methods of assessing the relative importance of different dimensions must be developed for use in conjunction with anisotropic sparse grids. The general sparse grid formulation described above is used by Nobile et. al. [31] for stochastic collocation of elliptic PDEs with random input data. There, the known regularity of the solution in each random dimension is used to provide the weighting of dimensions in the anisotropic grid, and estimates of the interpolation error are provided. However, we note that as stated in [27] these methods are not capable of accurately resolving solutions with large discontinuities.

For general problems where information regarding the dimension-wise stochastic regularity of the solution is unavailable, a variety of algorithms have been developed in numerical complexity theory to aid in designating important dimensions. This is particularly important for reducing computational effort in problems with high *nominal* random dimension. A 2005 report by Griebel [32] reviews a number of algorithms, beginning with methods stemming from Kolmogorov’s superposition theorem [33], where it is proven that continuous functions of several variables can be represented by superpositions of continuous functions of fewer variables. Of particular importance to stochastic collocation are analogous ideas found in statistics for regression problems and density estimation. The ANOVA (Analysis-of-Variance) decomposition, also described in [32], was introduced by Fisher in 1921 (see e.g. [34]) and utilized for studying U-statistics by Hoeffding in 1948 [35]. This approach involves splitting a multidimensional function into its contributions from different groups of sub-dimensions. This type of dimension-wise decomposition can effectively break the curse of dimensionality in certain approximation problems. The underlying idea involves the splitting of a one-dimensional function approximation space into the constant subspace and the remainder space. The associated splitting for multidimensional cases is formed via a product construction. In practice, one essentially truncates the ANOVA-type decomposition at a certain dimension d , thereby dealing with a series of low-dimensional ($\leq d$) approximation problems in lieu of one high-dimensional problem. This type of truncation can make high-dimensional approximation tractable for functions with high nominal dimension N but only low-order correlations amongst input variables (i.e. low effective dimension d). However, it is important to note that with

these types of decompositions, nothing is gained with respect to the curse of dimensionality when the *effective* dimension is close or equal to the nominal dimension of the problem. In other words, if the highest order terms in the ANOVA decomposition make large contributions, then the computational complexity of these methods is not smaller than that of the original high dimensional problem. This is not surprising as the same situation is well known also in quasi Monte Carlo (QMC) simulation studies [36,37]; QMC is fast only if the effective dimension is much smaller than the nominal dimension, i.e., $d \ll N$.

In [38], Rabitz and colleagues introduced two High Dimensional Model Reduction (HDMR) techniques to capture input-output relationships of physical systems with many input variables. These techniques (called ANOVA-HDMR and cut-HDMR) are based on ANOVA-type decompositions. ANOVA-HDMR utilizes the original ANOVA decomposition (which is also used statistics to analyze the variance of a random multivariate quantity). In this approach, multi-dimensional integrals must be evaluated to determine the component functions. The orthogonality condition imposed in this formulation ensures that the sum of variances of the component functions equals the full variance. Thus it can be used to provide information about the importance of different dimensions and of the correlations and interactions between groups of dimensions. In 2001, Sobol used this formulation to define *global sensitivity indices* [39] which reflect the relative variance contributions of different ANOVA terms, thus giving insight into the influence of individual parameters or subgroups of parameters and their interactions on the output. This same approach is explored by Griebel for use in numerical treatments of high-dimensional problems in [32]. However, a main drawback to this type of approach is the computational cost associated with evaluating the multi-dimensional integrals necessary for determining the component functions.

Rabitz et. al. [38] also introduced the cut-HDMR technique, in which the component functions are defined via hyperplane cuts through an anchor point in the center of the domain. This approach was shown to be computationally more efficient than the ANOVA-HDMR approach. However, in this formulation the sum of variances of the component functions is no longer equal the full variance. In 2006, Schwab and Todor [10] utilized the same approach to introduce a sparse polynomial chaos approximation for elliptic problems with stochastic coefficients, as an alternative to tensor product polynomial approximation bases. A stochastic collocation method based on this sparse basis was proposed and analyzed by Bieri and Schwab in 2008 [40], where the truncation dimension of the ANOVA-type series is chosen based on the size of the eigenvalues in the Karhunen-Loève random input expansion. They showed that for the stochastic elliptic model problem, random inputs of up to 80 dimensions were successfully addressed numerically with very low computational cost.

The Multi-Element Probabilistic Collocation Method (MEPCM) was shown in [1] to efficiently treat problems with discontinuous dependence on parameters and long time-integration. However, one main limitation of the MEPCM is the fact that elemental decomposition of high-dimensional random spaces is often prohibitively costly. Therefore, in this work we combine a cut-HDMR type ANOVA approach (analogous to the sparse chaos formulation of Schwab and colleagues) with the MEPCM and investigate its performance on high dimensional elliptic problems; we will refer to this method MEPCM-A. The use of the ANOVA-type decomposition reduces high-dimensional integration problems (e.g., stochastic collocation) to a series of low-dimensional problems. The MEPCM can then be applied to each of the low-dimensional problems where domain decomposition is *not* computationally prohibitive.

This paper is organized as follows. In Section 2 we introduce the general model problem and assumptions on the random inputs. In Section 4 the formulation of the MEPCM-A is given, and in Sections 5 and 6 we perform a variety of simple numerical studies to analyze the efficiency and accuracy of this method in comparison to other existing methods suitable for high-dimensional problems. To investigate convergence properties, we use moderate dimensional test integrands and a stochastic elliptic problem. We also investigate the convergence of the ANOVA-based single element collocation method (which we refer to as PCM-A) and its dependence on decomposition parameters. Next, we investigate the convergence of the MEPCM-A and compare the efficiency of MEPCM-A with and without a priori adapted meshes to sparse grid PCM as well as the PCM-A. For both PCM-A and MEPCM-A methods, we observe that the convergence is non-monotonic with respect to the computational costs, so the decomposition parameters must be chosen judiciously. We test the MEPCM-A method for integration of discontinuous functions in 100 to 500 dimensions using the GENZ package of testing functions for high dimensional integration. Lastly, in Section 7 we apply the MEPCM-A to aid in modeling subsurface contaminant transport through heterogeneous media at the U.S. Department of Energy’s Hanford Site in southeastern Washington state. We conclude that for problems with prohibitively high nominal dimension and possibly low stochastic regularity, the MEPCM-A can be a useful tool in solving otherwise computationally intractable problems, provided the solution has relatively *low effective dimension* or a decay in its interaction weights.

2 Model problem and assumptions

Let (Ω, \mathcal{A}, P) be a complete probability space, where Ω is the space of events, $\mathcal{A} \subset 2^\Omega$ is the σ -algebra of sets in Ω , and P is the probability measure. Also, define D to be a subset of $\mathbb{R}^d (d \in \{1, 2, 3\})$ with boundary ∂D . Let \mathcal{L} and \mathcal{R}

be operators on D and ∂D , respectively, where \mathcal{L} may depend upon $\omega \in \Omega$. In this work we consider the following problem: find $u : \overline{D} \times \Omega \rightarrow \mathbb{R}$ such that P -almost everywhere (a.e.) in Ω the following equation holds:

$$\begin{cases} \mathcal{L}(\mathbf{x}, \omega; u) = f(\mathbf{x}, \omega), & \mathbf{x} \in D, \\ \mathcal{R}(\mathbf{x}; u) = g(\mathbf{x}), & \mathbf{x} \in \partial D. \end{cases} \quad (1)$$

We assume that the boundary has sufficient regularity and that f and g are imposed so that the problem is well-posed P -a.e. We also assume that for P -a.e. $\omega \in \Omega$, the solution $u(\cdot, \omega)$ takes values in a Banach space, $\mathcal{W}(D)$, of functions over the physical domain taking values in \mathbb{R} . We assume additionally that the solution u is in $L^2(\Omega; \mathcal{W}(D))$. Note that the boundary operator \mathcal{R} may also have a random dependence, but for simplicity in this work we assume that \mathcal{R} is deterministic. In general, the various sources of randomness (e.g. random parameters in \mathcal{L} , initial/boundary conditions, driving terms) may arise from physically unrelated phenomena; thus, they can be independent and non-identically distributed.

In order to apply the methods that will be discussed later, the random dependence of operators \mathcal{L} and f must satisfy a few important properties. The first requirement, commonly known as a “finite dimensional noise assumption” [13,4], is that the random input can be represented with a finite-dimensional probability space. More specifically, the random input can be represented by a finite set of random variables $\{Y_1(\omega), Y_2(\omega), \dots, Y_N(\omega)\}$, with a known joint density function ρ . With this assumption on the random input, the problem (1) can be restated as follows. Find $u : \overline{D} \times \Omega \rightarrow \mathbb{R}$ such that

$$\mathcal{L}(\mathbf{x}, Y_1(\omega), Y_2(\omega), \dots, Y_N(\omega); u) = f(\mathbf{x}, Y_1(\omega), Y_2(\omega), \dots, Y_N(\omega))$$

holds $\forall \mathbf{x} \in D$ and for P -a.e. $\omega \in \Omega$, with corresponding boundary conditions. Using the Doob-Dynkin Lemma [41] we can assert that the solution $u(\mathbf{x}, \omega)$ can be written as $u(\mathbf{x}, \mathbf{Y}(\omega))$ with $\mathbf{Y} = (Y_1, Y_2, \dots, Y_N)$. Then, the problem may be recast from the space Ω into the target space of the N random variables. Let $\mathbf{y} = (y_1, y_2, \dots, y_N) \in \Gamma \equiv \prod_{j=1}^N \Gamma_j$, where Γ_j is the image of $Y_j(\Omega)$ for $j = 1, \dots, N$. Let $\rho(\mathbf{y})$ be the probability density function (PDF) of \mathbf{Y} . The problem can be restated: Find $u : \overline{D} \times \Gamma \rightarrow \mathbb{R}$ such that ρ -almost everywhere for $\mathbf{y} \in \Gamma$ the following equation holds:

$$\begin{cases} \mathcal{L}(\mathbf{x}, \mathbf{y}; u) = f(\mathbf{x}, \mathbf{y}), & \mathbf{x} \in D, \\ \mathcal{R}(\mathbf{x}, u) = g(\mathbf{x}), & \mathbf{x} \in \partial D. \end{cases} \quad (2)$$

Thus, the original problem (1) is recast as a fully deterministic problem in equation (2). It is sometimes useful to think of the solution u as a function

on Γ , taking values in $\mathcal{W}(D)$. In this case we would denote $u(\mathbf{y})$ to be the Banach-valued solution to the problem for a particular $\mathbf{y} \in \Gamma$.

3 ANOVA decomposition

The ANOVA decomposition dating back to Hoeffding in [35] states that an N -dimensional function f can be decomposed as follows:

$$\begin{aligned}
 f(x_1, x_2, \dots, x_N) = & f_0 + \sum_{j_1}^N f_{j_1}(x_{j_1}) + \sum_{j_1 < j_2}^N f_{j_1, j_2}(x_{j_1}, x_{j_2}) \\
 & + \sum_{j_1 < j_2 < j_3}^N f_{j_1, j_2, j_3}(x_{j_1}, x_{j_2}, x_{j_3}) + \dots + f_{j_1, \dots, j_N}(x_{j_1}, \dots, x_{j_N}),
 \end{aligned} \tag{3}$$

where f_0 is a constant, f_{j_1} is a one-dimensional function, and so on.

Based upon the decomposition idea of (3), a ‘sparse FEM’ method proposed in [40] deconstructs the N -dimensional polynomial chaos approximation problem into a series of approximation problems on ‘sparse’ polynomial chaos bases spanning different groups of subdimensions. If the high-dimensional problem has low *effective* dimension (i.e. it can be sufficiently approximated by a series of low-dimensional subproblems), this type of approach can greatly alleviate the computational burden. In the following we formulate the MEPCM-A which combines this subset decomposition approach with the MEPCM.

4 MEPCM-A Formulation

The formulation of MEPCM-A basically involves performing the standard MEPCM [1] on a series of low-dimensional problems using the aforementioned ANOVA decomposition. Therefore, we describe the MEPCM portion of the formulation briefly and refer the reader to [1] for more detail. The point we emphasize here is the subset decomposition of the high-dimensional integration problem into a series of low-dimensional integration problems.

Let $u : \overline{D} \times \Gamma \rightarrow \mathbb{R}$ be the exact solution to our general problem (2). Here $\Gamma = [a_i, b_i]^N$ is a rectangular hypercube in \mathbb{R}^N . For every subset $K \subset \{1, \dots, N\}$, let Γ_K be the tensor product of the slices of Γ restricted to each dimension in K . For example, if $K = \{1, 4\}$, $\Gamma_K = [a_1, b_1] \times [a_4, b_4]$. We define $\{A^{i,K}\}_{i=1}^{N_K}$ to be a nonoverlapping rectangular mesh of each Γ_K . Then, let $B^{i,K} = A^{i,K} \times \Gamma_{K^c}$.

where K' is the complement of K in $\{1, \dots, N\}$. The MEPCM-A stochastic semidiscrete solution is given by the following:

$$\mathcal{I}_{N,\mu,\nu}u(\mathbf{x}, \mathbf{y}) = \sum_{K \subset \{1, \dots, N\}, |K| \leq \nu} \gamma_{N,|K|,\nu} \left(\sum_{i=1}^{N_e^K} \mathbb{I}_{\mathbf{y} \in B^{i,K}} \mathcal{I}_{B^{i,K}}^\mu u(\mathbf{x}, \mathbf{y}) \right) \quad (4)$$

where

$$\mathcal{I}_{B^{i,K}}^\mu u(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{k} \in \mathbb{N}_0^N, k_i \leq \mu, \text{supp}(\mathbf{k})=K} u(\mathbf{x}, \mathbf{q}_{\mathbf{k}}^i) l_{\mathbf{k}}^i(\mathbf{y}) \quad (5)$$

and

$$\gamma_{N,j,\nu} = \sum_{r=j}^{\nu} (-1)^{r-j} \binom{N-j}{r-j}. \quad (6)$$

Here, μ represents the number of points used in the interpolation rule in each dimension and thus governs the order of polynomial interpolation (it does not need to be the same in each dimension but we assume this for simplicity). In addition, ν is the highest dimension of each subproblem, and N is the total number of dimensions in the problem. The interpolation abscissas $\mathbf{q}_{\mathbf{k}}^i = (q_{1,k_1}^i, \dots, q_{N,k_N}^i)$ are chosen to be the roots of tensorized orthogonal polynomials in $A^{i,K}$ of order μ . More specifically, for $n \in K$, the points $q_{n,m}^i$, $m = 1, \dots, \mu$ are the μ roots of the orthogonal polynomial of corresponding order in dimension n of $A^{i,K}$, and $q_{n,0}^i$ equals the midpoint of $[a_n, b_n]$ for any $n \in K'$. (Note here that we choose the midpoint as the anchor point for the ANOVA cuts. In future work we investigate techniques for choosing optimal anchor points to increase accuracy [?], preprint). The interpolating polynomials in each element are given by: $l_{\mathbf{k}}^i(\mathbf{y}) = \prod_{n=1}^N l_{n,k_n}^i(y_n)$. Here, if $n \in K$, then $l_{n,k_n}^i(y_n)$ is the Lagrange interpolation polynomial of degree μ through the point q_{n,k_n}^i restricted to $B^{i,K}$ satisfying:

$$l_{n,k_n}^i(q_{n,j_n}^i) = \delta_{k_n,j_n}, \quad j_n = 1, \dots, \mu.$$

Otherwise, if $n \in K'$, $l_{n,k_n}^i(y_n) = 1$. The $\gamma_{N,j,\nu}$ are the weights associated with each subproblem and derived from Proposition 5.1 in [10].

Remark: Some results regarding the error and convergence rate associated with this discretization in the single-element case ($N_e^K = 1$ for all subsets K) are presented in [40] for an elliptic boundary value problem with a stochastic diffusion coefficient possessing a piecewise analytic two-point correlation. Further analysis of the truncation error associated with ν in the full multi-element formulation is the subject of ongoing work ([?]).

Remark: We note that a Smolyak sparse grid interpolation operator can be used instead of full tensor product interpolation in (5), just as in the formulation of the MEPCM. In this case, the parameter μ governing the order

of polynomial interpolation in each dimension would be replaced by s , the *sparseness* parameter. As s increases, the sparsity of the grid decreases and thus the size of the point set increases. We refer the reader to [1,?] for a more detailed description of the use of Smolyak grids in stochastic collocation and the sparseness parameter.

The moments of u are then calculated as in the MEPCM using whichever quadrature rule is associated with the choice of points, either Gaussian quadrature or Smolyak sparse grid quadrature. The dimension of each MEPCM-A subproblem is bounded by ν , thus introducing an additional stochastic *truncation error* not dealt with in the standard MEPCM. The interplay of the two decomposition parameters ν and μ will be investigated in the following numerical examples.

5 PCM-A numerical studies

We begin by numerically demonstrating the ANOVA-type approach within the framework of the standard *one-element* probabilistic collocation in order to provide an understanding of the (ν, μ) relationship to integration errors.

5.1 Approximation of integrals

We first demonstrate this method on a few simple integrals to elucidate some basic behaviors. As a preliminary example we construct the 10-dimensional test function $u : [-1, 1]^{10} \rightarrow \mathbb{R}$.

$$u = \sum_{i=1}^{10} x_i^2 + \sum_{i=1}^9 x_i^2 x_{i+1}^2$$

and approximate the integrals

$$I : \int_{[-1,1]^{10}} u d\mathbf{x}$$

and

$$II : \int_{[-1,1]^{10}} u^2 d\mathbf{x}$$

The reference solutions (obtained using Mathematica) are $I : 4.43733(10)^3$ and $II : 2.16476(10)^4$. Varying parameters μ and ν we obtain relative errors for integral I . PCM-A is performed using a tensor product Gauss-Legendre collocation grid. The relative error is defined as $\epsilon = \frac{|I_{num} - I_{exact}|}{I_{exact}}$. Table 1 provides the MEPCM-A integration errors of function I for various μ and ν . We

Table 1

Relative errors for integral I , for various parameters ν, μ . Since the function is second-order polynomial and utilizes interactions of degree at most two, the choice $\nu, \mu = 2$ provides an exact solution.

ν	$\mu = 1$	$\mu = 2$	$\mu = 3$
1	1.308	0.231	0.231
2	3.385	$4.3(10)^{-15}$	$2.7(10)^{-15}$
3	3.385	$2.19(10)^{-14}$	$2.85(10)^{-14}$

Table 2

Relative errors for integral II . As expected, the minimum parameters for exactness are $\mu = 3$ and $\nu = 4$ since the squaring of the integrand has the effect of increasing the effective dimension of the function and Gauss quadrature with μ points is exact for polynomials of degree $2\mu - 1$.

ν	$\mu = 1$	$\mu = 2$	$\mu = 3$	$\mu = 4$
2	5.859	0.133	0.302	0.302
3	13.4	0.147	0.0327	0.0327
4	16.076	0.180	$4.369(10)^{-14}$	$1.160(10)^{-14}$

note that since the function is second-order and utilizes interactions of degree at most two, the choice $\nu, \mu = 2$ provides an exact solution. Table 2 provides the analogous errors for integral II . As expected, the minimum parameters for exactness are $\mu = 3$ and $\nu = 4$ since the squaring of the integrand has the effect of increasing the effective dimension of the function and Gauss quadrature with $\mu + 1$ points is exact for polynomials of degree $2(\mu + 1) - 1$.

Next we investigate the function $u : [-1, 1]^{10} \rightarrow \mathbb{R}$

$$u = x_1^2 + 3x_1^2x_2^2 + 3x_3^2x_4^2 + cx_5^2x_6^2x_7^2.$$

If we choose $c = 9$ the relative error using $\mu = 2, \nu = 2$ is 0.25, and for the choice $\mu = 2, \nu = 3$ we have relative error $O(10^{-16})$. Now, suppose we let $c = 0.1$ instead; then, the error for $\mu = 2, \nu = 2$ is just $3.7(10)^{-3}$ as we would expect, since this choice of c has an effect similar to lowering the effective dimension of the function.

5.2 Stochastic elliptic problem

Next we move on to the following classic example of a stochastic elliptic problem. Here we consider the following problem: find $u : \overline{D} \times \Omega \rightarrow \mathbb{R}$ such that

P -almost everywhere (a.e.) in Ω the following equation holds:

$$\begin{cases} -\nabla \cdot (a(\mathbf{x}; \omega) \nabla u(\mathbf{x}; \omega)) = f(\mathbf{x}) & \text{in } D, \\ u(\mathbf{x}; \omega) = g(\mathbf{x}) & \text{on } \partial D, \end{cases} \quad (7)$$

where $f(\mathbf{x})$ is assumed to be deterministic for simplicity and $a(\mathbf{x}; \omega)$ is a second-order random process satisfying the following assumption:

Assumption. Let $a(\mathbf{x}; \omega) \in L^\infty(\Omega; DW(D))$ be strictly positive with lower and upper bounds a_{min} and a_{max} , respectively,

$$0 < a_{min} < a_{max} \text{ and } P(a(\mathbf{x}; \omega) \in [a_{min}, a_{max}], \forall \mathbf{x} \in \overline{D}) = 1.$$

We assume that the random input $a(\mathbf{x}; \omega)$ can be expressed as a finite expansion of independent random variables. In the first example we choose $D = [0, 1]$, $u(0) = 0$ and $u(1) = 1$ with random input

$$a(x; \omega) = 20 + \sum_{i=1}^{10} \frac{10}{2^{i-1}} \sin(ix) Y_i(\omega), \quad (8)$$

where the Y_i are independent *uniform* random variables on $[-1, 1]$. Note that strong ellipticity is preserved $\forall \omega \in \Omega$.

The (single-element) PCM-A method is used with a full tensor product Gauss-Legendre grid of $(\mu + 1)$ points in each dimension for every subset. The spatial discretization error is negligible. The parameter $\mu = 3$ is kept constant while ν is increased, and the solution is compared with the Monte Carlo solution with 25 million samples. Figure 1 shows the spatial L^∞ errors in mean solution obtained. The L^∞ error of two numerically integrated functions $f(t_j)$ and $g(t_j)$, $j = 1 \dots n_t$ is defined to be $\epsilon_{L^\infty} = \max_{j=1 \dots n_t} \{|f(t_j) - g(t_j)|\}$. It can be seen in Figure 1 that the error is dominated by MC standard error for $\nu > 2$. In Figure 2 we plot the error between the PCM-A and a reference PCM solution with 1, 048, 576 Gauss-Legendre tensor product points. Here we observe convergence to the reference solution as ν is increased. To illustrate the computational cost of these calculations, Table 3 shows the total number of points used at each level of ν and also the number of subproblems required for the computation.

Next, we investigate the relationship between parameters μ and ν on the convergence of the PCM-A method for this problem. Again, a Gauss-Legendre tensor product grid is used in each subset. To test convergence, ν is increased from 1 to 4, and for each ν , μ is increased from 1 to 5. We define the L^2 error between two numerically integrated functions $g(t_j)$ and $f(t_j)$, $j = 1 \dots n_t$ as:




Fig. 1. Error in *mean* between PCM-A solutions of the stochastic elliptic problem with a 10-dimensional random input (8), at varying ν and a reference MC solution with 25 million samples. The Monte Carlo computed standard error is also plotted.




Fig. 2. Error between PCM-A solutions of stochastic elliptic problem with a 10-dimensional random input (8), at varying ν and a reference PCM solution with 1,048,576 points. The Monte Carlo computed standard error is also plotted for 25 million samples.

$$\epsilon_{L^2} = \frac{\frac{1}{n_t} \sqrt{\sum_{j=1}^{n_t} (g(t_j) - f(t_j))^2}}{\frac{1}{n_t} \sqrt{\sum_{j=1}^{n_t} f(t_j)^2}}, \quad (9)$$

where f is considered to be the reference solution. A reference solution is calculated using sparse grid PCM with sparseness parameter $s = 17$ and 652,065 nested Clenshaw-Curtis points.

Table 4 shows the number of points used for each parameter set as well as the L^∞ errors for the mean and variance of the solution. For a more visual representation of results, Figure 3(left) shows the L^2 errors for each (ν, μ)

ν	total # of points	# of subproblems
1	41	11
2	761	56
3	844	176
4	62201	386

Table 3

Total number of points and number of subproblems used for the PCM-A computations for the 10-dimensional random input (8) whose error is plotted in Figures 1 and 2.

parameter set. We observe that for each ν , there exists a lower bound on the error beyond which increasing μ does not improve the solution. This is therefore error the of truncating the stochastic interaction dimension of the problem to ν . We also note that the optimal path to convergence for this case is achieved by choosing $\mu = \nu + 1$ and increasing ν . However, the same convergence rate (but lower accuracy) is also achieved using the choice $\mu = \nu$.

- This suggests that μ and ν should be simultaneously increased to achieve monotonic convergence.

The errors are clearly the highest when μ is chosen to be 1 for any ν . This can be seen more clearly in Figure 3(right) where the errors are plotted as a function of computational work (total number of points sampled). There, it is clear that the relationship between computational effort and error is non-monotonic. Therefore, it is important to make a judicious choice for (ν, μ) when using this method.

We also consider the case where the order of the coefficients of $a(x, \omega)$ is reversed so that the largest coefficient multiplies the highest order sinusoidal function:

$$a(x; \omega) = 20 + \sum_{i=1}^{10} \frac{10}{2^{10-i}} \sin(ix) Y_i(\omega), \quad (10)$$

where the Y_i are independent *uniform* random variables on $[-1, 1]$. In Figure 4 we compare the PCM-A variance errors for the stochastic elliptic problem with exponentially *decaying* and exponentially *increasing* coefficients. Once again, Gauss-Legendre points are used for each subset and a sparse grid PCM solution is used as the reference. Choosing $\mu = \nu$ and $\mu = \nu + 1$ and increasing ν ensures a monotonic convergence. As expected, the errors for the exponentially increasing coefficients case are larger since the high-order spatial basis functions are given heavier weights.

Table 4

PCM-A is performed using Gauss-Legendre tensor product grid points in each subset for the stochastic elliptic problem (7) using a 10-dimensional random input with exponentially decaying coefficients (8). This table shows the computational costs associated with each (ν, μ) pair along with the L^∞ errors in mean and variance. The reference solution is calculated using sparse grid PCM with sparseness parameter $s = 17$ and 652,065 nested Clenshaw-Curtis points.

ν	μ	number of points	Mean error(L^∞)	Variance error (L^∞)
1	1	21	1.0625e-02	2.5084e-03
	2	31	1.4125e-04	7.5214e-05
	3	41	2.8307e-04	1.4376e-04
	4	51	2.8617e-04	1.4559e-04
	5	61	2.8625e-04	1.4564e-04
2	1	201	1.3488e-02	3.9159e-03
	2	436	2.1277e-04	1.0289e-04
	3	761	6.0328e-06	1.6686e-06
	4	1176	6.1993e-06	3.7767e-06
	5	1681	6.3043e-06	3.8624e-06
3	1	1161	1.3712e-02	4.0607e-03
	2	3676	2.1687e-04	1.0794e-04
	3	8441	6.6734e-06	3.9882e-06
	4	16176	2.4737e-07	1.0209e-07
	5	27601	7.7124e-08	8.2200e-08
4	1	4.521	1.3714e-02	4.0652e-03
	2	20686	2.1688e-04	1.0803e-04
	3	62201	6.6994e-06	4.0743e-06
	4	14743	2.7190e-07	1.8437e-07
	5	299761	1.2529e-08	8.9012e-09

6 MEPCM-A numerical studies

In this section we study how the convergence of MEPCM-A is affected by ν, μ and the adapted mesh discretization of the random space. We will also compare the errors of using several methods including MEPCM-A, sparse grid




Fig. 3. L^2 error of PCM-A solution of the stochastic elliptic problem (7) using a 10-dimensional random input (8) with Gauss-Legendre tensor product grids in each subset. Left: Varying colors indicate different levels of ν , and for each line, μ is increased from left to right in the range 1 to 5. The reference solution is calculated using sparse grid PCM with $q = 17$. Right: Errors are plotted versus the number of points used in each computation. The points corresponding to the cases $\mu = \nu$ and $\mu = \nu + 1$ are labeled.




Fig. 4. Comparison of PCM-A L^2 error in variance for the stochastic elliptic problem with exponentially decaying and exponentially increasing coefficients of a . A Gauss-Legendre collocation grid is used in each element. Errors are shown for increasing ν , for cases $\mu = \nu$ and $\mu = \nu + 1$.

PCM and PCM-A with the same computational costs. For these purposes, we will revisit the example one-dimensional stochastic elliptic problem in (7) with the stochastic input given in (8).

We are especially interested in investigating how adaptive mesh refinement of MEPCM affects the efficiency of the method. In the first example, we use an adapted mesh in which only the two dimensions with highest coefficients are discretized. In these top two directions, two elements per side are used, and in the remaining dimensions only one element is used. A Gauss-Legendre grid is used to prescribe collocation points in each element in each subset.

Table 5

MEPCM-A is performed using Gauss-Legendre tensor product grid points in each subset for the stochastic elliptic problem (7) using a 10-dimensional random input with exponentially decaying coefficients (8). The first two dimensions (with highest coefficients) are discretized in the MEPCM mesh. This table shows the computational costs associated with each (ν, μ) pair along with the L^∞ errors in mean and variance. The reference solution is calculated using sparse grid PCM with sparseness parameter $s = 17$ and 652065 nested Clenshaw-Curtis points.

ν	μ	number of points	Mean error(L^∞)	Variance error (L^∞)
1	1	25	2.9383e-04	1.4956e-04
	2	37	2.8631e-04	1.4568e-04
	3	49	2.8625e-04	1.4565e-04
	4	61	2.8625e-04	1.4565e-04
2	1	281	1.5423e-05	9.4854e-06
	2	613	6.3832e-06	3.9240e-06
	3	1073	6.3094e-06	3.8668e-06
	4	1661	6.3085e-06	3.8660e-06
3	1	1881	1.1368e-05	6.3268e-06
	2	6013	1.8814e-07	1.9281e-07
	3	13873	8.0918e-08	9.3400e-08
	4	26661	8.0219e-08	9.1476e-08
4	1	8377	1.1369e-05	6.2687e-06
	2	38899	1.5636e-07	1.0251e-07
	3	117809	3.2357e-09	2.6457e-09
	4	280411	5.2590e-10	7.4695e-10

Table 5 shows the computational cost associated with each choice of (ν, μ) as well as the L^∞ errors in mean and variance of the solution. Once again the

reference solution here is calculated using sparse grid PCM with sparseness parameter $s = 17$ and 652065 nested Clenshaw-Curtis points. In Figure 5(left) we plot the L^2 errors using MEPCM-A on this adapted mesh. We observe a similar type of convergence as ν and μ vary and again conclude that μ and ν should simultaneously be increased for fast convergence. Figure 5(right) shows the L^2 error in variance versus the amount of work performed for each computation. We observe similar convergence behavior as in the PCM-A case, wherein there exists a truncation dimension error associated with each ν and a point beyond which increasing μ has little improvement on the error. Other cases including refining only the first and first three dimensions were also considered and results will be shown in section 6.3.




Fig. 5. L^2 error of MEPCM-A solution of the stochastic elliptic problem (7) using a 10-dimensional random input (8) with Gauss-Legendre tensor product grids in each element of each subset. The first two dimensions (with highest coefficients) are discretized in the MEPCM-A mesh. Left: Varying colors indicate different levels of ν , and for each line, μ is increased from left to right in the range 1 to 4. Right: L^2 error of the variance vs. computational cost. The points corresponding to the cases $\mu = \nu$ are labeled along with selected other points. The reference solution is calculated using sparse grid PCM with $q = 17$.

6.2 Uniform mesh

Next we investigate the case where all dimensions are discretized equally with two elements on each side. In Figure 6 we compare the performance of the sparse grid MEPCM-A with two different types of grids in each element: Clenshaw-Curtis sparse grids versus full tensor product Gauss-Legendre grids. For sparse grids, the sparseness parameter s used in each subproblem is the same level above the dimension of each subset. We observe that the full tensor product grids are actually slightly more efficient than the nested sparse grids – and conclude that in the regime where MEPCM-A is most useful (when small ν is sufficient), sparse grids do not provide a significant advantage. Since Gaussian quadrature grids offer much more robust integration and the number of points between consecutive sparseness levels is large, we prefer to use this

choice for MEPCM-A.




Fig. 6. Error versus work for full and sparse grid MEPCM-A performed on the stochastic elliptic problem (7) using a 10-dimensional random input (8) with Gauss-Legendre tensor product grids in each element of each subset. Every dimension is discretized with two elements in the MEPCM-A mesh.

6.3 Comparison with other methods

In this section we compare the efficiency of MEPCM-A and PCM-A to the standard probabilistic collocation method (PCM) using nested sparse grids for the 10-dimensional elliptic problem. We use tensor product Gauss-Legendre grids for MEPCM-A and PCM-A, and sparse Clenshaw Curtis grids for the PCM. MEPCM-A is calculated using three meshes wherein the top one, two, and three dimensions (in importance) are discretized, and also using one mesh where all dimensions are discretized. We see in Figure 7 that we achieve convergence of the ANOVA-type methods by increasing μ and ν simultaneously, using either $\mu = \nu$ or $\mu = \nu + 1$. Convergence for PCM is achieved by increasing the sparseness parameter s , and a reference solution is calculated using sparse grid PCM with $s = 17$.

We observe in Figure 7 that while all of these methods converge to the correct solution, MEPCM-A provides an improvement over PCM-A for the adapted mesh cases and performs quite poorly for a uniform mesh, as expected. MEPCM-A with the top few dimensions discretized is competitive with sparse PCM computational efficiency for all error levels. Note that MEPCM-A provides the ability to deal with solutions with low stochastic regularity, a feature

that is not highlighted here since this example problem is smooth. These results suggest that with the coupling of an ANOVA-type decomposition to mesh-adaptive MEPCM, we may be able to address problems with low regularity in high dimensions, while still preserving the computational efficiency of single-element sparse grid methods which scale well with dimension.

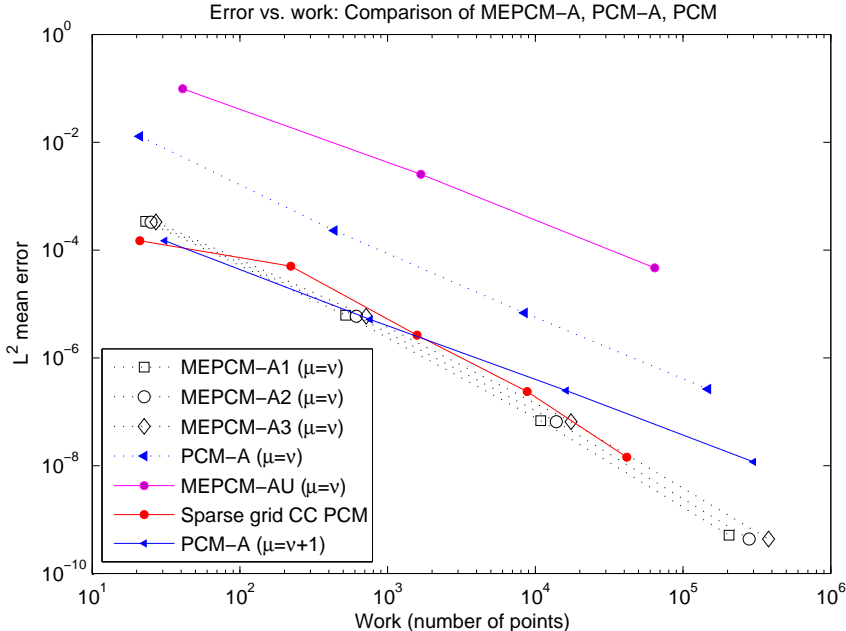


Fig. 7. L^2 error in mean versus computational cost for the stochastic elliptic problem with 10-dimensional random input. MEPCM-A1: Adapted mesh refined in top dimension, MEPCM-A2: Adapted mesh refined in top two dimensions, MEPCM-A3: Adapted mesh refined in top three dimensions, MEPCM-AU: Uniform mesh refined in all dimensions. All MEPCM-A and PCM-A computations use full tensor product Gauss-Legendre grids in each element of each subset. PCM is calculated using a sparse Clenshaw Curtis grid.

6.4 Random process inputs

We also consider a higher-dimensional case when a is defined to be a second-order random process with covariance kernel:

$$\text{Cov}(x, y) = \exp \left\{ -\frac{|x - y|}{b} \right\}. \quad (11)$$

For this covariance kernel, exact forms for the Karhunen-Loève (KL) eigenfunctions and eigenvalues can be found [42]. We utilize these eigenpairs to define a truncated expansion (with the same form as a KL expansion) in terms of independent uniform random variables on $[-1, 1]$. Note that this cannot technically be termed a KL expansion since the distributions of the random variables in a KL expansion are derived from the random process itself, and

here we specify only the covariance kernel of the process. Here we define the process a to be a truncated expansion in terms of uniform random variables having the desired covariance structure, by utilizing the KL expansion form and eigenpairs. We use a correlation length $b = 0.5$ and truncate the expansion after 50 terms. The random variables in the expansion are taken to be uniformly distributed on $[-1, 1]$. MEPCM-A is performed with an adapted mesh discretizing the first two dimensions and a Gauss-Legendre collocation grid is used in each element (12, 073 points are used in total). The variance of the solution is plotted in Figure 8. To calculate the error, sparse grid PCM with 171, 901 points was performed. The L^∞ error in mean and variance are $3.2(10)^{-8}$ and $3.2(10)^{-10}$, respectively, and the L^2 error in mean and variance are $3.9(10)^{-8}$ and $3.6(10)^{-6}$ respectively.




Fig. 8. Variance of MEPCM-A solution to the stochastic elliptic problem (7) with random input given by a 50-dimensional truncated Karhunen Loève expansion of a random process with exponential covariance kernel. The two dimensions with highest eigenvalues are discretized, $\nu = 2$, and $\mu = 2$. 12073 Gauss-Legendre points are used in total.

We also consider the *two-dimensional* stochastic elliptic problem on $D = [0, 1]^2$. We use a non-zero force term with homogeneous boundary conditions

$$f(\mathbf{x}) = \sin(x_1) \cos(x_2) \text{ and } \mathbb{E}[a](\mathbf{x}) = 1. \quad (12)$$

Assume that the random field $a(\mathbf{x}, \omega)$ satisfies the Gaussian correlation function:

$$K(\mathbf{x}_1, \mathbf{x}_2) = \delta^2 e^{\frac{-|\mathbf{x}_1 - \mathbf{x}_2|^2}{A}}$$

with $A = 1$ being the correlation length and $\delta = 0.3$ the standard deviation. Due to the analyticity of the Gaussian kernel, the eigenvalues decay exponentially [9]. The decay rate is determined by the value of the correlation length, where a larger A corresponds to a faster decay rate. Since the Gaussian kernel is analytic, high-order element methods for spatial discretization convergence

quickly, resulting in highly accurate numerical solutions for the eigenvalue problem in the Karhunen-Loève expansion. Once again, using the KL eigen-pairs calculated for the Gaussian kernel, the random input a is defined to be a truncated expansion in terms of uniform independent random variables on $[-1, 1]$.

Spectral/ hp discretization with 64 quadrilateral elements is used in the physical space [43]. In each element, an 8-th order Jacobi polynomial basis is used to construct the approximation space. Errors from spatial discretization are negligible relative to the stochastic discretization errors.

In Figure 9 we show the errors of the H_0^1 norm of the solution of MEPCM-A on two examples of this problem. This norm is defined as the standard Sobolev norm on the space D : $\|f\|_{H_0^1(D)} = \left(\int_D (|f|^2 + |\nabla f|^2) \right)^{1/2}$. In the left plot, the random input is taken to be the expansion truncated after $N = 10$ terms, and in the right plot 50 expansion terms are retained. The reference solution for each case is calculated via sparse grid PCM with sparseness level $s = N + 3$. MEPCM-A is performed using tensor product Gauss-Legendre grids in each element, and the two dimensions with highest eigenvalues are discretized. We observe that as expected, the errors decrease as both μ and ν are increased. We note that this behavior is not modified by greatly increasing the dimension of the problem.




Fig. 9. Errors of MEPCM-A solution to the two-dimensional stochastic elliptic problem (7) on $D = [0, 1]^2$ where a is a random process with a Gaussian covariance structure. The process a is defined using Karhunen-Loève eigen-pairs calculated for the Gaussian kernel as an expansion in terms of independent uniform random variables on $[-1, 1]$ (left: A 10-dimensional truncation of the expansion. right: a 50-dimensional truncated expansion of the same process). The two dimensions with highest KL eigenvalues are discretized.

6.5 Integration of Genz test functions

In this section we consider the some integrands in the multiple integration test suite of Genz [44]. We concentrate in particular on the integrands from this suite called ‘CONTINUOUS’ and ‘DISCONTINUOUS’ to compare the efficiency of the MEPCM-A to sparse grid PCM for problems with low stochastic regularity.

The Genz function f_6 (DISCONTINUOUS) is defined as:

$$f_6(\mathbf{x}) = \begin{cases} 0, & \text{if } x_1 > w_1 \text{ or } x_2 > w_2 \\ \exp \sum_{i=1}^N c_i x_i, & \text{otherwise} \end{cases} \quad (13)$$

where \mathbf{x} is an N -dimensional vector in $[0, 1]^N$, the w_i and c_i are constants between zero and 1.

Figure 10 shows the relative errors as a function of the computational effort for the MEPCM-A and sparse grid Clenshaw Curtis PCM methods. The random dimension N here is 10, the c_i are chosen to decay exponentially ($c_i = 10/2^i$) and $w_i = 0.5$. MEPCM-A is performed with Gauss-Legendre points on an adapted mesh where the first two dimensions are discretized so that the discontinuity falls on the boundary of the elements. The sparse grid PCM fails to converge with this example.




Fig. 10. Comparison of the relative errors between MEPCM-A and Clenshaw Curtis sparse grid PCM on the DISCONTINUOUS test function f_6 . The total number of random dimensions $N = 10$ and MEPCM-A is performed using Gauss-Legendre points on a mesh adapted to the discontinuity.

To contrast, we also consider the CONTINUOUS test function f_5 , using the

same coefficients c_i and w_i as for f_6 . This function is defined as:

$$f_5(\mathbf{x}) = \exp\left(-\sum_{i=1}^N c_i |x_i - w_i|^2\right)$$

In this case, the sparse grid PCM converges to the correct solution with very good efficiency. We perform MEPCM-A with the first three dimensions discretized and a Gauss-Legendre collocation grid in each element. In Figure 11 we observe the MEPCM-A still performs just as well as the highly efficient sparse grid PCM for this moderate dimensional problem. In the next section we consider high dimensional cases where $N \geq 100$.




Fig. 11. Comparison of the relative errors between MEPCM-A and Clenshaw Curtis sparse grid PCM on the CONTINUOUS test function f_5 . The total number of random dimensions $N = 10$ and MEPCM-A is performed using Gauss-Legendre points on a mesh discretized in the first 3 dimensions.

6.5.1 High dimensional integration

We next perform several high-dimensional integrations of f_6 where $N = 100, 200,$ and 300 using MEPCM-A. Once again an adapted mesh discretizing the first two dimensions is used, and the c_i are chosen to decay as $c_i = \exp^{-(35.0/(N-1))i}$ (so that the coefficients are bounded away from zero). Table 6 shows the computational costs and errors for various choices of (ν, μ) . We see that for the 100 dimensional case, error on the order of 10^{-2} is achieved with just 5 million points. For comparison, using either tensor product or sparse grid PCM in 100 dimensions with just one element, with a bare minimum of just 3 points in each dimension would be computationally intractable, requiring on the order of 10^{47} points to compute. As N increases, the number

of dimensions contributing to the integrand also increases; thus the error associated with addressing only subproblems of dimension less than or equal to ν increases as N increases. The number of points required for each (μ, ν) discretization increases with increasing N as well. This example suggests that MEPCM-A can be a powerful tool for problems with high nominal dimension N but low *effective* dimensionality.

Table 6

MEPCM-A errors and computational costs for 100 to 500-dimensional integration of test function f_6 (DISCONTINUOUS).

N	(ν, μ)	relative error	number of points
100	(1, 1)	$\mathcal{O}(1)$	103
	(2, 2)	0.0197	20801
	(3, 3)	0.0098	4677148
200	(1, 1)	$\mathcal{O}(1)$	203
	(2, 2)	0.067	81601
	(3, 3)	0.047	36714298
300	(1, 1)	$\mathcal{O}(1)$	303
	(2, 2)	0.12	182401
	(3, 3)	0.09	123111448
400	(1, 1)	$\mathcal{O}(1)$	403
	(2, 2)	0.22	323201
	(3, 3)	0.07	290868598
500	(1, 1)	$\mathcal{O}(1)$	503
	(2, 2)	0.43	504001
	(3, 3)	0.21	566985748

As another illustrative example, we now focus on the DISCONTINUOUS integrand f_6 given in (13). We choose: $c_i = \exp(-Ai)$, $A = 0.2$, and $w_i = 0.5$, and we set a relative error tolerance of $\epsilon = 10^{-4}$. As the dimension is increased, relative contributions from additional dimensions decrease in importance as would be expected in a Karhunen Loève input, for example. In Figure 12 we compare the minimum number of points required to achieve an error less than the tolerance for various random dimensions. The MC estimates are achieved by calculating the standard error of the estimator for the mean; we plot the number of samples required so that the MC sample mean is within ϵ relative error of the true solution with 95% confidence. The PCM and PCM-A methods both fail to converge for this problem due to the discontinuities. The MEPCM

is performed with an adapted mesh in which only the first two dimensions are discretized, with full Gauss-Legendre and sparse Clenshaw Curtis grids. Lastly, the MEPCM-A is computed with an adapted mesh in which only the first two dimensions are discretized, and full Gauss-Legendre grids are used in each element. We note that for the MEPCM-A, the choice $\nu = 4, \mu = 2$ approximates the function with sufficient accuracy at any dimension, due to the decaying nature of the coefficients. Beyond 100 dimensions, the exact reference solution is difficult to compute. Thus, the dashed line represents estimation of the continued MEPCM-A line by calculating the number of points required for the same ($\nu = 4, \mu = 2$) choice.

Remark: We note that this example illustrates the fact that boundary issues should be considered when making a choice of grids. In this case, the MEPCM with sparse grids is *less* efficient than with tensor product grids at low dimensions due to the fact that Clenshaw Curtis points fall on the boundary of the the elements, which meet at the discontinuities of f_6 . Thus, it is important to consider that any single grid point of a collocation grid is given a non-negligible weight, as opposed to a probability measure of zero, so care must be taken not to place points directly on discontinuities or singularities where the function is not indicative of surrounding points.

From this illustration we observe that the crossover point beyond which Monte Carlo is favorable to MEPCM-A (approximately 600 dimensions) is significantly higher than the crossover point for MEPCM full and sparse grid methods (less than 50 dimensions).

6.6 Hierarchical nature of the subset decomposition

Finally, we observe that all subsets used in the calculation of a solution with parameters $(\nu - 1, \mu)$ are also used in the calculation of our solution with parameters (ν, μ) . Thus, it is possible to retain the moments calculated in each subset and recompute the solution that would have been obtained with all lower values of ν . Then, the percentage difference between the (ν, μ) solution and solutions with lower values of ν can be calculated. This can be useful in determining when a particular level of ν is not sufficient for a given problem; if the percentage difference between the solutions at levels $\nu - 1$ and ν is large, it is likely that the *effective* dimension of the problem is higher than ν .

We demonstrate this on the function f_6 (DISCONTINUOUS) in 10 dimensions with exponentially decaying c_i in figure 13. For this problem the choice $\nu = 2$ is sufficient for a fairly good approximation of the integral, and subsequently higher levels of ν provide smaller contributions.




Fig. 12. Top: Minimum computational effort required to achieve error tolerance of $\epsilon = 10^{-4}$ for approximating the integral of DISCONTINUOUS (f_6) function with exponentially decaying c_i . Bottom: Same plot on log-log axes. For MC, we plot the number of samples required such that the MC sample mean is within ϵ relative error of the true solution with 95% confidence. The MEPCM is performed with an adapted mesh in which only the first two dimensions are discretized, with both full Gauss-Legendre and sparse Clenshaw Curtis points. The MEPCM-A is performed with an adapted mesh in which only the first two dimensions are discretized, with full Gauss-Legendre grids. We note that for the MEPCM-A, the choice $\nu = 4, \mu = 2$ approximates the function with sufficient accuracy at any dimension, due to the decaying nature of the coefficients. Beyond 100 dimensions, the exact reference solution is difficult to compute. Thus, the dashed line represents estimation of the continued MEPCM-A line by calculating the number of points required for the same (μ, ν) choice.




Fig. 13. Percentage error between MEPCM-A solutions with $\nu < 6$ and the MEPCM-A solution with $\nu = 6$. This is accomplished via re-postprocessing of the MEPCM-A solution at $\nu = 6$.

7 Modeling subsurface contaminant transport through heterogeneous media

To demonstrate the computational tractability of MEPCM-A on complex realistic problems, we apply this method to study subsurface contaminant transport in natural porous media. In such problems, uncertainty arises from a lack of information concerning physical features and events or a lack of understanding of the processes controlling groundwater flow and transport. It is important to understand these uncertainties and their potential impact on model results.

Until recently, quantifying uncertainty in the outcome of these problems was performed with Monte Carlo simulation. Unfortunately, the expense of MC simulation is too costly to provide reliable results. For simple probabilistic models of the hydraulic properties (with low random dimension), stochastic finite element techniques have been able to speed up the computations significantly. However, realistic representation of these features requires the capability of handling random inputs with high dimension. In this section we aim to apply the MEPCM-A to address this problem with high-dimensional random inputs, using results of simpler studies in the previous section to judiciously make parameter choices.

7.1 Background

We focus in particular on the US Department of Energy’s Hanford Site, which is located along the Columbia River in southeastern Washington. During the

Hanford site's 60-plus years of history, there have been more than 1000 individual sources of contaminants distributed over 200 square miles mostly along the Columbia River. A total of 450 billion gallons of contaminated liquid were spilled containing radioactive materials such as cesium, strontium, technetium, iodine, uranium, neptunium and plutonium and large quantities of organic contaminants such as carbon tetrachloride [45]. Stewardship efforts for this site require detailed modeling and prediction of the subsurface transport of these contaminants. However, the hydraulic properties of the subsurface are highly heterogeneous and difficult to fully characterize.

The ability to describe subsurface heterogeneities in quantitative terms and to incorporate this quantitative description into computer models is a key element in predicting transport of subsurface contaminants and the reduction of modeling uncertainties. While all subsurface environments are clearly deterministic, numerous experimental data sets suggest that the permeability varies significantly in space in a manner that cannot be described with certainty. Consequently, over the last decade it has become common to treat these quantities as spatially correlated random fields, thereby making velocity and saturation spatially correlated random fields as well. Hence, probabilistic descriptions using spatial stochastic processes may be employed to describe the variations in the geologic structure at the Hanford site. Figure 14 (from [46]) is a map showing the location and areas surrounding the Hanford site - Pasco Basin. The regional aquifer system extends from western Idaho through eastern Washington and northeastern Oregon. This problem has been studied extensively in stochastic hydrogeology (e.g. [47] and references therein).

7.2 *Mathematical model*

In this section we consider natural stochastic extensions of Darcy's law and mass conservation in order to address the effect of variability in the hydraulic conductivity, \mathbf{K} , on the piezometric head, h .

$$\mathbf{q} = -\mathbf{K}(\mathbf{x}; \omega) \nabla h(\mathbf{x}; \omega), \quad -\nabla \cdot \mathbf{q} + f = 0, \quad (14)$$

where we have assumed for simplicity that the source term is zero and that the problem is steady. Then, the system reduces to the stochastic elliptic problem

$$-\nabla \cdot (\mathbf{K}(\mathbf{x}; \omega) \nabla h(\mathbf{x}; \omega)) = 0. \quad (15)$$

The domain of this problem is the region within the model boundary detailed in Figure 14. Boundary conditions have been modeled by the authors of [46] via analysis of hydrologic properties including permeability of the bedrock, recharge volumes, runoff of precipitation, and many other factors. A schematic of the lateral boundary conditions is shown in Figure 15. There, various regions




Fig. 14. Extent of Regional and Local Groundwater Flow Systems Beneath the Hanford Site (from [46], section 4).

are marked where either the value of the piezometric head are specified, the Darcy flux is specified, or zero flow is imposed on the boundary.

7.3 Results

We perform the PCM-A on the problem in equation (15) with boundary conditions on the physical domain as given in Figure 15. The random hydraulic conductivity, \mathbf{K} , is modeled as a spatially correlated random process on the domain with a standard logarithmic transformation to ensure that it is bounded away from zero. The correlation length is roughly equal to one third the distance between the points furthest north and south on the domain. The mean of this process is $e^{0.5\sigma^2}$ with $\sigma = 0.05$, and the variance is equal to $e^{2\sigma^2} - e^{\sigma^2}$. Once again, using the KL eigenpairs calculated for the covariance kernel, the




Fig. 15. Groundwater flow model grid and lateral boundary conditions (from [46], section 4).

random input \mathbf{K} is defined to be a truncated expansion in terms of uniform independent random variables on $[-1, 1]$. The expansion is truncated after 100 terms. We demonstrate the capability of PCM-A to address high dimensional inputs on this problem.

The spatial discretization is performed using spectral elements on the approximately 15000-element mesh shown in Figure 15, with 3rd order Jacobi polynomials on each element. PCM-A is performed using Gauss-Legendre tensor product grids in each element. In Figure 16 we plot the first and second moments of the value of the piezometric head over the domain, for two choices of ANOVA discretization parameter $\nu = 1, 2$. We observe that the difference between solutions at $\nu = 1$ and $\nu = 2$ is small. To quantify the difference between these discretization levels we also calculated the mean of the H_0^1 norm of the solution for each case, shown in Table 7.

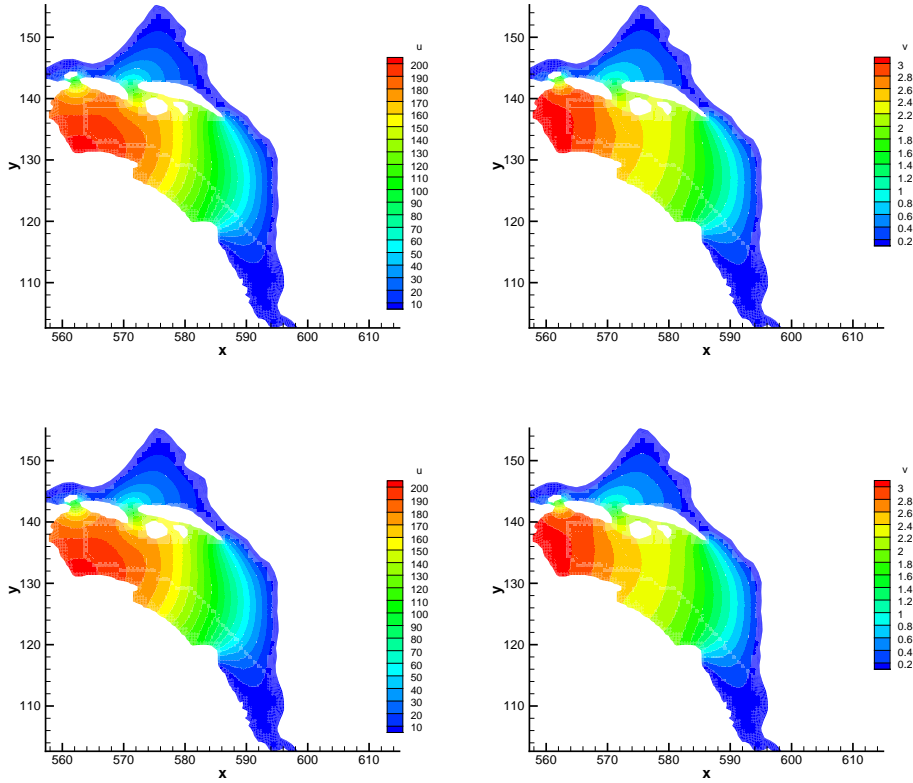


Fig. 16. PCM-A results for problem (15) quantifying the effect of uncertainty in hydraulic conductivity \mathbf{K} on the piezometric head h . The conductivity is modeled by a 100-term K-L expansion of a lognormal spatial random process. Top row: $\nu = 1$ results for mean (left) and standard deviation (right). Bottom row: $\nu = 2$ results for mean (left) and standard deviation (right).

Table 7

Number of points used for each ν and mean of the H_0^1 norm of the piezometric head.

ν	Number of points	$\mathbb{E}[\ h\ _{H_0^1(D)}]$
1	401	9.030922763071
2	44581	9.030966179659

8 Summary and Discussion

In this work a variant of the MEPCM method utilizing an ANOVA-type decomposition was proposed for dealing with problems of high random dimension. We first numerically investigated the dependence of the convergence of this method on the decomposition parameters μ and ν , concluding that in order to achieve monotonic convergence μ and ν must be increased simultaneously. In addition, it was shown that for the 10-dimensional stochastic elliptic problem, MEPCM-A with an adapted mesh performs better than PCM-A and

as well as the sparse grid PCM method, which can only be used for smooth problems. We also demonstrated the MEPCM-A for the integration of functions with low regularity in very high dimensions.

We also applied the PCM-A to obtain error bars for the piezometric head at the Hanford site under stochastic hydraulic conductivity conditions. For future uncertainty quantification at the Hanford site, data from drilling samples and knowledge of the subsurface composition will be incorporated into the stochastic models. It is highly likely that various regions of the domain will be modeled independently under different spatial random processes, depending on regional geologic variations. Within each region, the estimated correlation length of these processes may be relatively small. Thus, we expect that in order to realistically represent the variation in conductivity, very high-dimensional random inputs may be required. In future studies, the MEPCM-A may provide a useful tool for handling such high-dimensional inputs in this complex system. We conclude that for problems with prohibitively high nominal dimension and possibly low stochastic regularity, the MEPCM-A can be a useful tool in solving otherwise computationally intractable problems, provided the solution has relatively low effective dimension or a decay in its interaction weights.

We note that although we increase the computationally tractable number of dimensions by incorporating the ANOVA decomposition, there still exists a dimension limit beyond which Monte Carlo simulation is more efficient for moment calculation. This is due to the fact that the computational costs of MC simply scale more favorably with increasing dimension. We are then interested in asking the question: as we increase the number of stochastic dimensions in a problem, where is the crossover point beyond which Monte Carlo is more efficient than the MEPCM-A and other stochastic collocation variants? This crossover dimension is of course dependent on the problem and the particular error tolerance we seek, but addressing a simple example problem may shed light on relative differences between crossover dimensions of the various methods.

The results presented in this work suggest that the MEPCM-A can provide a more efficient alternative to Monte Carlo for some problems in higher-dimensional regimes, which may be intractable for other stochastic collocation methods. In general, the MEPCM-A is a promising tool for addressing problems of high dimension and low regularity; however, further investigation is needed to more clearly delineate the class of problems for which it provides the greatest benefit.

Specifically, it is of great practical importance that further effort is put into incorporating approaches from data mining, statistics and financial mathematics for estimating the *effective dimension*. Examples of such work include the report of Caffisch et. al. [48] on effective dimension in mortgage backed

securities, as well as the work of Wang and collaborators in [49,37] and references therein. With these techniques, certain classes of high dimensional PDE problems with lower effective dimension or a decay in the interaction weights of the solution may be detected and treated automatically [32]. It will be interesting to also extend MEPCM-A to non-elliptic PDEs; for *nonlinear evolution PDEs*, like the Navier-Stokes equations, computing the effective dimensionality is much more complex. (The more fundamental question on the existence of low-dimensional manifolds in turbulent flows will not be addressed here but physical and numerical evidence points to *coherent structures* and relatively low-dimensionality, see [50,51].) Finally, systematic work on obtaining the composite error bound due to error contributions from the MEPCM representation and the ANOVA decomposition is required.

Acknowledgements

We would like to thank Prof. Xiaoliang Wan for his help with the Hanford site simulations. In addition, we would like to thank Prof. Christoph Schwab and Dr. Marcel Beri at ETH, Zurich, for their helpful suggestions. We thank the anonymous referees for their comments and suggestions which greatly improved the manuscript. This work was partially supported by DOE, ONR and ONR/ESRDC Consortium, and computations were performed on the DoD/HPCM supercomputers. Jasmine Foo would like to acknowledge the support of DOE fellowship (CSGF) under grant DE-FG02-97ER25308 and the Krell institute and travel support by NSF grant OISE-0456114.

References

- [1] J. Foo, X. Wan, G. Karniadakis, The multi-element probabilistic collocation method: error analysis and simulation, *J. Comp. Phys.* 227 (2008) 9572–9595.
- [2] I. Babuška, R. Tempone, G. E. Zouraris, Galerkin finite element approximations of stochastic elliptic differential equations, *SIAM J. Numer. Analysis* 42 (2) (2004) 800–825.
- [3] D. Xiu, G. E. Karniadakis, The Wiener-Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.* 24 (2) (2002) 619–644.
- [4] D. Xiu, J. Hesthaven, High-order collocation methods for differential equations with random inputs, *SIAM J. Sci. Comput.* 27 (3) (2005) 1118–1139.
- [5] X. Wan, G. E. Karniadakis, Multi-element generalized polynomial chaos for arbitrary probability measures, *SIAM J. Sci. Comput.* 28 (3) (2006) 901–928.
- [6] L. Mathelin, M. Hussaini, A stochastic collocation algorithm for uncertainty analysis, NASA/CR-2003-212153.
- [7] M. K. Deb, I. Babuška, J. T. Oden, Solution of stochastic partial differential equations using Galerkin finite element techniques, *Comput. Methods Appl. Mech. Engrg.* 190 (2001) 6359–6372.
- [8] R. G. Ghanem, P. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag, New York, 1991.
- [9] P. Frauenfelder, C. Schwab, R. A. Todor, Finite elements for elliptic problems with stochastic coefficients, *Comput. Methods Appl. Mech. Engrg.* 194 (2005) 205–228.
- [10] R. A. Todor, C. Schwab, Convergence rates for sparse chaos approximations of elliptic problems with stochastic coefficients, *IMA J. Numer. Anal.* 27 (2) (2007) 232–261.
- [11] H. G. Matthies, A. Keese, Galerkin methods for linear and nonlinear stochastic partial differential equations, *Comput. Methods Appl. Mech. Engrg.* 194 (12-16) (2005) 1295–1331.
- [12] M. Tatang, G. McRae, Direct treatment of uncertainty in models of reaction and transport, MIT Technical Reports.
- [13] I. Babuška, F. Nobile, R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data, *SIAM J. Numer. Analysis* 45 (3) (2007) 1005–1034.
- [14] B. Ganapathysubramanian, N. Zabaras, Sparse grid collocation for stochastic natural convection problems, *J. Comput. Phys.* 225 (1) (2007) 652–685.
- [15] X. Wan, G. E. Karniadakis, An adaptive multi-element generalized polynomial chaos method for stochastic differential equations, *J. Comput. Phys.* 209 (2) (2005) 617–642.

- [16] O. P. L. Maitre, H. N. Najm, R. G. Ghanem, O. M. Knio, Uncertainty propagation using Wiener-Haar expansions, *J. Comput. Phys.* 197 (2004) 28–57.
- [17] O. P. L. Maitre, H. N. Najm, R. G. Ghanem, O. M. Knio, Multi-resolution analysis of Wiener-type uncertainty propagation schemes, *J. Comput. Phys.* 197 (2004) 502–531.
- [18] R. G. Ghanem, Hybrid stochastic finite elements: coupling of spectral expansions with monte carlo simulations, *ASME Journal of Applied Mechanics* 65 (1998) 1004–1009.
- [19] O. P. L. Matre, M. T. Reagan, H. N. Najm, R. G. Ghanem, O. M. Knio, A stochastic projection method for fluid flow: Ii. random process, *Journal of Computational Physics* 181 (1) (2002) 9–44.
- [20] A. Keese, H. Matthies, Numerical methods and Smolyak quadrature for nonlinear partial differential equations, *Informatikbericht 2003-5*, Technische Universitat Braunschweig.
- [21] S. Smolyak, Quadrature and interpolation formulas for tensor products of certain classes of functions, *Soviet Math. Dokl.* 4 (1963) 240–243.
- [22] E. Novak, K. Ritter, Simple cubature formulas with high polynomial exactness, *Constr. Approx.* 15 (1999) 499–522.
- [23] V. Barthelmann, E. Novak, K. Ritter, High dimensional polynomial interpolation on sparse grids, *Adv. Comput. Math.* 12 (2000) 273–288.
- [24] E. Novak, K. Ritter, High dimensional integration of smooth functions over cubes, *Numer. Math* 75 (1996) 79–97.
- [25] K. Petras, On the Smolyak cubature error for analytic functions, *Advances in Computational Mathematics* 12 (2000) 71–93.
- [26] F. Nobile, R. Tempone, C. Webster, A sparse grid collocation method for elliptic partial differential equations with random input data, *SIAM Journal on Numerical Analysis* 46 (5) (2008) 2309–2345.
- [27] T. Gerstner, M. Griebel, Dimension-adaptive tensor-product quadrature, *Computing* 71 (1) (2003) 65–87.
- [28] M. Hegland, Adaptive sparse grids, *Proceedings of CTAC, Brisbane* (2001) 16–18.
- [29] L. Plaskota, The discrepancy of sparse grids is at least 2.1933, *Adv. Comp. Math* 12 (2000) 3–24.
- [30] G. Wasilkowski, H. Wozniakowski, Weighted tensor product algorithms for linear multivariate problems, *J. Complexity* 15 (1999) 402–447.
- [31] F. Nobile, R. Tempone, C. Webster, An anisotropic sparse grid collocation method for elliptic partial differential equations with random input data, *SIAM Journal on Numerical Analysis* 46 (5) (2008) 2411–2442.

- [32] M. Griebel, Sparse grids and related approximation schemes for higher-dimensional problems, *Proc. Foundations of Computational Mathematics* 248.
- [33] A. Kolmogorov, On the representation of continuous functions of several variables by superpositions of continuous functions of several variables and addition, *Dokl. Akad. Nauk SSSR* 114 (1957) 953–956.
- [34] R. Fisher, *Statistical Methods for Research Workers*, Oliver and Boyd, 1925.
- [35] W. Hoeffding, A class of statistics with asymptotically normal distributions, *Annals of Math. Statist.* 19 (1948) 293–325.
- [36] I. Sloan, When are Quasi-Monte Carlo algorithms efficient for high-dimensional integrals?, *J. Complexity* 14 (1998) 1–33.
- [37] X. Wang, I. Sloan, Why are high-dimensional finance problems often of low effective dimension?, *SIAM J. Sci. Comp.* 27 (1) (2005) 159–183.
- [38] H. Rabitz, O. Alis, J. Shorter, K. Shim, Efficient input-output model representations, *Comput. Phys. Comm.* 117 (1999) 11–20.
- [39] I. Sobol, Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Mathematics and Computers in Simulation* 55 (2001) 271–280.
- [40] M. Bieri, C. Schwab, Sparse high order fem for elliptic spdes, *Comput. Methods Appl. Mech. Engrg.* 198 (2009) 1149–1170.
- [41] B. Oksendal, *Stochastic Differential Equations*, Springer-Verlag, 1998.
- [42] M. Jardak, C.-H. Su, G. E. Karniadakis, Spectral polynomial chaos solutions of the stochastic advection equation, *J. Sci. Comput.* 17 (2002) 319–338.
- [43] G. E. Karniadakis, S. J. Sherwin, *Spectral/hp Element Methods for CFD*, Oxford University Press, 2005.
- [44] A. Genz, A package for testing multiple integration subroutines, *Numerical integration: Recent developments, software and applications* (1987) 337–340.
- [45] U. S. Department of Energy, *Groundwater Remediation Project Fact Sheet*, June 2006.
- [46] P. D. Thorne, M. P. Bergeron, M. D. Williams, V. L. Freedman, *Groundwater data package for Hanford assessments*, Tech. Rep. PNNL-14753, Rev. 1, Pacific Northwest National Laboratory (2006).
- [47] B. Ye, D. Yang, D. L. Kane, Changes in lena river streamflow hydrology: Human impacts versus natural variations, *Water Resour. Res.* 39 (2003) 1200–1224.
- [48] R. Caflisch, W. Morokoff, A. Owen, Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension, *Finance* 1 (1) (1997) 27–46.
- [49] X. Wang, K.-T. Fang, The effective dimension and Quasi-Monte Carlo integration, *J. Complexity* 19 (2003) 101–124.

- [50] L. Keefe, P. Moin, J. Kim, The dimension of attractors underlying periodic turbulent Poiseuille flow, *J. Fluid Mech.* 242 (1992) 1–29.
- [51] K. Ball, L. Sirovich, L. Keefe, Dynamical eigenfunction decomposition of turbulent channel flow, *Int. J. Num. Meth. Fluids* 12 (1991) 585–604.