CrossMark

A model and variance reduction method for computing statistical outputs of stochastic elliptic partial differential equations

F. Vidal-Codina [a],[*],[1],[2], N.C. Nguyen [a],[1], M.B. Giles [b], J. Peraire [a],[1]

[a] *Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[b] *Mathematical Institute, University of Oxford, Oxford, UK*

ARTICLE INFO

ABSTRACT

We present a model and variance reduction method for the fast and reliable computation of statistical outputs of stochastic elliptic partial differential equations. Our method consists of three main ingredients: (1) the hybridizable discontinuous Galerkin (HDG) discretization of elliptic partial differential equations (PDEs), which allows us to obtain high-order accurate solutions of the governing PDE; (2) the reduced basis method for a new HDG discretization of the underlying PDE to enable real-time solution of the parameterized PDE in the presence of stochastic parameters; and (3) a multilevel variance reduction method that exploits the statistical correlation among the different reduced basis approximations and the high-fidelity HDG discretization to accelerate the convergence of the Monte Carlo simulations. The multilevel variance reduction method provides efficient computation of the statistical outputs by shifting most of the computational burden from the high-fidelity HDG approximation to the reduced basis approximations. Furthermore, we develop *a posteriori* error estimates for our approximations of the statistical outputs. Based on these error estimates, we propose an algorithm for optimally choosing both the dimensions of the reduced basis approximations and the sizes of Monte Carlo samples to achieve a given error tolerance. We provide numerical examples to demonstrate the performance of the proposed method.

1. Introduction

The analysis of physical systems is often carried out by mathematical modeling and numerical simulation. For a given system, the corresponding mathematical model requires certain input data, such as material properties, forcing terms, boundary conditions and geometry information. For many problems of interest, input data are not known precisely. In such cases, one may need to consider input data as random variables and represent them in probabilistic terms. Mathematical models represented by partial differential equations with random input data are known as stochastic partial differential equations (SPDEs). Uncertainty in the input data may come from different sources. It can be that the physical system has some in-

---

trinsic variability, for example, uncertainty in the gust loads on an aircraft, or wind and seismic loading on civil structures. It is also possible that we are unable to effectively characterize the physical system with a mathematical model because, for instance, we may have errors in geometry, roughness of surfaces, or multiscale behavior that we are unable to capture. Therefore, there is a growing need to represent the uncertainty in the data and effectively propagate it through the mathematical model. The goal of this probabilistic approach resides in computing statistics of some observable outputs (quantities of interest), which are usually defined as functionals of the solution of the underlying SPDE.

There exist a number of different approaches to solve SPDEs and retrieve the statistics of the output. The most common approach is to use Monte Carlo (MC) methods [20,31]. Monte Carlo methods only need repeated evaluations of the output functional of the solution of the SPDEs for different instantiations of the random input. The main advantage of Monte Carlo methods is that their convergence rate is independent of the dimension of the stochastic space, namely, the number of random variables. The main caveat of these methods is their slow convergence rate, which demands a large amount of realizations to achieve accurate results. As a result, a number of techniques such as quasi Monte Carlo methods [8,48], Latin Hypercube Sampling [33,57], variance reduction methods [8] and multilevel Monte Carlo [27] have been proposed to alleviate the slow convergence rate of the standard Monte Carlo methods.

Another approach is stochastic Galerkin methods, first introduced by Ghanem et al. in [26], that generalize the theory of Wiener–Hermite polynomial chaos expansion [63] and combine it with a finite element method to model uncertainty in a SPDE. In this approach, the random variables are treated as additional dimensions of the problem and projected onto a stochastic space spanned by a set of orthogonal polynomials. The problem is then reduced to a system of deterministic equations, which couple the physical and stochastic dimensions. This methodology has proven to be very effective when solving SPDEs in a broad range of applications, such as diffusion problems and heat conduction [24,65,67], structural dynamics [25], transport in random media [23] and fluid dynamics [10,66]. The advantage of these methods is that they converge exponentially fast for a sufficiently regular solution field [2,3,18]. However, their main drawback is that their computational complexity grows combinatorially with the number of random variables and the number of expansion terms. As a consequence, they are not effective for solving problems with a large number of random variables.

A more recent approach is stochastic collocation methods (SCM), first introduced in [35] and further developed in [64]. The main idea is to compute deterministic solutions of the SPDE for certain instantiations of the random variables and then construct an interpolation function to approximate the response over the stochastic space. When the interpolation procedure is performed on tensor grids, these methods suffer from the exponential growth with the dimensionality of the stochastic space. To economize the interpolation process in large dimensions, sparse grids (Smolyak [56]) were introduced for elliptic problems [51,64], parabolic problems [49] and natural convection problems [21]. In [1], sparse grids were shown to achieve exponential convergence for problems with smooth solutions. However, like polynomial chaos expansions, sparse grids still suffer from the curse of dimensionality in the sense that the number of grid points grows rapidly with the dimension of the stochastic space. Recently, anisotropy and adaptivity on sparse grids [22,32] have been used in SCM [21,50] to mitigate the elevated cost in high dimensions.

In this paper, we present a model and variance reduction method for the fast and reliable computation of statistical outputs of stochastic elliptic partial differential equations. Our method consists of three main ingredients: (1) the hybridizable discontinuous Galerkin (HDG) discretization of elliptic partial differential equations (PDEs), which allows us to obtain high-order accurate solutions of the governing PDE; (2) a reduced basis method for the HDG discretization of the underlying PDE to enable real-time solution of the parameterized PDE in the presence of stochastic parameters; and (3) a multilevel variance reduction method that exploits the statistical correlation among the different reduced basis approximations and the high-fidelity HDG discretization to accelerate the convergence rate of the Monte Carlo simulations. The multilevel variance reduction method provides efficient computation of the statistical outputs by shifting most of the computational burden from the high-fidelity HDG approximation to the reduced basis approximations. Although the three ingredients of our approach exist in the literature, the main contribution of this paper is to put these methodologies into a unified framework that combines all of their strengths to tackle stochastic elliptic PDEs. Another important contribution of the paper is to develop *a posteriori* error bounds for the estimates of the statistical outputs and to introduce an algorithm for optimally choosing the dimensions of the reduced basis approximations and the sizes of MC samples to achieve a given error tolerance. Last but not least, we present a new HDG formulation that enables the efficient construction of reduced basis approximations for the HDG discretization of parameterized PDEs.

The HDG method was first introduced in [14] for elliptic problems, subsequently analyzed in [13,16,17], and later extended to a wide variety of PDEs [15,38–46,59]. The HDG method is particularly effective for solving elliptic PDEs because it possesses several unique features that distinguish it from other DG methods. First, it reduces the number of globally coupled unknowns to those required to represent the trace of the approximate solution on the element boundaries, thereby resulting in a smaller global system than other DG methods. Second, the method provides optimal convergence rates for both the solution and the flux. And third, its flux superconvergence properties can be exploited to devise a local postprocess that increases the convergence rate of the approximate solution by one order. These advantages are the main driver for the development of the Reduced Basis (RB) method for the HDG discretization of parameterized PDEs. While the RB method is well developed for the standard finite element discretization of parameterized PDEs [28,29,34,52,54,60,61], the RB method for the HDG approximation of parameterized PDEs has not been considered before. The HDG discretization has multiple field variables and various equivalent weak formulations, which make the application of the RB method non-straightforward.

Recently, the RB method has been applied to standard continuous Galerkin finite element solutions of stochastic elliptic PDEs [7,9,30]. In this approach, the stochastic PDE is first reformulated as a parametrized PDE over the coefficients of the Karhunen–Loève expansion of the random fields. The reduced basis approximation and associated *a posteriori* error estimation are then developed for the resulting parametrized PDE. Finally, the output statistics and their error estimates are computed with a MC simulation [7,30] or a stochastic collocation approach [9]. These approaches, which involve the RB method and its *a posteriori* error bounds to evaluate the output instead of the original finite element discretization, have been shown to outperform both standard MC and stochastic collocation. In this paper, we extend the previous work [7,30] in several important ways. We will use the HDG method to construct the RB approximation. We will adopt the multilevel Monte Carlo strategy [5,12,27,58] and demonstrate a significant computational gain relative to the standard MC approach. Moreover, we will provide *a posteriori* error estimates for our prediction of the statistical outputs without involving *a posteriori* error bounds for the RB approximation. This feature will broaden the applicability of our approach to a wide variety of stochastic PDEs for which *a posteriori* error bounds for the RB approximation are either not available or too expensive to compute.

According to the central limit theorem [19], the error in a Monte Carlo estimation of the expectation of an output is proportional to the square root of the ratio of the variance of the output and the number of samples. Therefore, in order to reduce the error one can increase the number of samples and/or decrease the variance of the output. Because increasing the number of samples leads to higher computational cost, various techniques such as the control variates method [6,8,31], the multilevel Monte Carlo method [5,12,27,58], and the multi-fidelity Monte Carlo method [36] have been proposed to reduce the variance of the output. The control variates method reduces the variance of the output by making use of the correlation between the output and a surrogate. The multi-fidelity Monte Carlo method makes use of the statistical correlation between the low-fidelity (surrogate) and high-fidelity outputs to reduce the number of high-fidelity evaluations needed to achieve a given accuracy of interest. The multilevel Monte Carlo method applies the principle of control variates to a sequence of lower fidelity outputs (multigrid approximations) to estimate the statistics of the high-fidelity output. Likewise, our method applies the principle of control variates to the HDG approximation and a sequence of reduced basis approximations, thereby shifting the computational burden from the high-fidelity HDG discretization to the lower fidelity RB approximations.

This article is organized as follows. In Section 2, we introduce a stochastic elliptic boundary value problem and describe a new weak HDG formulation particularly suited for the reduced basis method. In Section 3, we describe a reduced basis method for the HDG approximation of the stochastic elliptic boundary value problem. In Section 4, we develop a multilevel Monte Carlo method that incorporates the HDG approximation and its reduced basis models into a unified framework to provide rapid reliable computation of the statistical outputs. In Section 5, we present numerical results to demonstrate the performance of the proposed method. Finally, in Section 6, we discuss some directions for future research.

2. The hybridizable discontinuous Galerkin method

2.1. Problem statement

Let $\mathcal{D} \in \mathbb{R}^d$ be a regular domain with Lipschitz boundary $\partial \mathcal{D}$. We consider the following stochastic boundary value problem: find a function u such that,

$$-\nabla \cdot (\kappa \nabla u) + \varrho u = f, \qquad \forall \boldsymbol{x} \in \mathcal{D}, \tag{1a}$$

$$\kappa \nabla u \cdot \boldsymbol{n} + \nu u = g, \qquad \forall \boldsymbol{x} \in \partial \mathcal{D}, \tag{1b}$$

where f is the source term, κ is the diffusion coefficient, ϱ is the Helmholtz parameter, ν is the Robin coefficient, and g is the boundary data. In this problem, one or more than one of the quantities f, κ, ϱ, ν, g are stochastic functions. For simplicity of exposition we shall assume that κ is a real stochastic function and that f, ϱ, ν, g are deterministic. The generalization to the case where one or more of f, ϱ, ν, g are stochastic is straightforward. Note that since we allow f, ϱ, ν, g to be complex-valued functions, the solution u is in general a complex stochastic function.

We next introduce a probability space $(\Omega, \mathcal{F}, P)$, where Ω is the set of outcomes, $\mathcal{F}$ is the σ-algebra of the subsets of Ω, and P is the probability measure. If Z is a real random variable in $(\Omega, \mathcal{F}, P)$ and ω a probability event, we denote its expectation by $E[Z] = \int_\Omega Z(\omega) dP(\omega)$. We will consider random functions v in $L^2(\mathcal{D} \times \Omega)$ equipped with the following norm

$$\|v\|^2 = E\left[\int_{\mathcal{D}} |v(\boldsymbol{x}, \cdot)|^2 d\boldsymbol{x} \right] = \int_\Omega \int_{\mathcal{D}} |v(\boldsymbol{x}, \omega)|^2 d\boldsymbol{x} \, dP(\omega).$$

We will assume that $\kappa \in L^2(\mathcal{D} \times \Omega)$ and that $\kappa(\boldsymbol{x}, \omega)$ is bounded and strictly positive, i.e., there exist constants α_1 and α_2 such that

$$0 < \alpha_1 \leq \kappa(\boldsymbol{x}, \omega) \leq \alpha_2 < +\infty, \qquad \text{a.s. in } \mathcal{D} \times \Omega.$$

We next assume that the random function $\kappa(\boldsymbol{x}, \omega)$ can be written in the following form

$$\kappa(\boldsymbol{x}, \omega) = \bar{\kappa}(\boldsymbol{x}) + \sum_{q=1}^{Q} \psi_q(\boldsymbol{x}) y_q(\omega), \qquad (\boldsymbol{x}, \omega) \in \mathcal{D} \times \Omega,$$

where $\bar{\kappa}(\boldsymbol{x})$ is the expectation of κ, $\psi_q(\boldsymbol{x})$, $q = 1, \ldots, Q$ are uniformly bounded real functions, and $y_q(\omega)$ for $q = 1, \ldots, Q$ are mutually independent random variables with zero mean. In addition, we assume that each of the $y_q(\omega)$ is bounded in the interval $\Lambda_q = [-\gamma_q, \gamma_q]$ with a uniformly bounded probability density function $\rho_q : \Lambda_q \to \mathbb{R}^+$. It thus follows that, with a slight overloading of notation, we can write κ in the form

$$\kappa(\boldsymbol{x}, \boldsymbol{y}) = \bar{\kappa}(\boldsymbol{x}) + \sum_{q=1}^{Q} \psi_q(\boldsymbol{x}) y_q, \qquad (\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{D} \times \Lambda, \tag{2}$$

where $\boldsymbol{y} = (y_1, \ldots, y_Q)$ and $\Lambda = \prod_{q=1}^{Q} \Lambda_q$.

Therefore, the solution u of (1) can be written as a function of $\boldsymbol{y} \in \Lambda$, namely, $u(\boldsymbol{x}, \boldsymbol{y})$. Now let ℓ be a bounded linear functional. We introduce a random output s defined as

$$s(\boldsymbol{y}) = \ell(u(\cdot, \boldsymbol{y})).$$

We are interested in evaluating the expectation and variance of s as

$$E[s] = \int_{\Lambda} s(\boldsymbol{y}) \rho(\boldsymbol{y}) d\boldsymbol{y}, \qquad V[s] = \int_{\Lambda} (E[s] - s(\boldsymbol{y}))^2 \rho(\boldsymbol{y}) d\boldsymbol{y},$$

where $\rho(\boldsymbol{y}) = \prod_{q=1}^{Q} \rho_q(y_q)$. Below we describe the hybridizable discontinuous Galerkin method for solving the model problem (1) and the Monte Carlo simulation for computing estimates of $E[s]$ and $V[s]$.

2.2. HDG discretization

We begin by rewriting the governing equation (1) as a first-order system

$$\boldsymbol{q} - \nabla u = 0, \quad \text{in } \mathcal{D}, \tag{3a}$$

$$-\nabla \cdot \kappa \boldsymbol{q} + \varrho u = f, \quad \text{in } \mathcal{D}, \tag{3b}$$

$$\kappa \boldsymbol{q} \cdot \boldsymbol{n} + \nu u = g \quad \text{on } \partial \mathcal{D}. \tag{3c}$$

The physical domain $\mathcal{D}$ is triangulated into elements T forming a mesh $\mathcal{T}_h$ satisfying the standard finite element conditions [11]. Then, letting $\partial \mathcal{T}_h := \{\partial T : T \in \mathcal{T}_h\}$ and denoting by $\mathcal{F}_h$ the set of the faces F of the elements $T \in \mathcal{T}_h$, we seek a vector approximation $\boldsymbol{q}_h \in \boldsymbol{V}_h^p$ to $\boldsymbol{q}$, a scalar approximation $u_h \in W_h^p$ to u, and a scalar approximation $\widehat{u}_h \in M_h^p$ to the *trace* of u on element boundaries, where

$$\boldsymbol{V}_h^p = \{\boldsymbol{v} \in \boldsymbol{L}^2(\mathcal{D}) : \boldsymbol{v}|_T \in [P_p(T)]^d \ \forall T \in \mathcal{T}_h\},$$

$$W_h^p = \{w \in L^2(\mathcal{D}) : w|_T \in P_p(T) \ \forall T \in \mathcal{T}_h\},$$

$$M_h^p = \{\mu \in L^2(\mathcal{F}_h) : \mu|_F \in P_p(F) \ \forall F \in \mathcal{F}_h\},$$

and $P_p(D)$ is a space of complex-valued polynomials of degree at most p on D. Note that $\widehat{u}_h$ are defined only on the faces of the elements, hence they are single valued. We introduce the following inner products

$$(v, w)_{\mathcal{T}_h} := \sum_{T \in \mathcal{T}_h} (v, w)_T, \qquad \langle v, w \rangle_{\partial \mathcal{T}_h} := \sum_{T \in \mathcal{T}_h} \langle v, w \rangle_{\partial T},$$

where $(u, v)_D := \int_D u \bar{v} \, dx$ whenever D is a domain in $\mathbb{R}^d$, and $\langle u, v \rangle_D := \int_D u \bar{v} \, dx$ whenever D is a domain in $\mathbb{R}^{d-1}$. For vector-valued functions $\boldsymbol{v}$ and $\boldsymbol{w}$, the integrals are similarly defined with the integrand being the dot product $\boldsymbol{v} \cdot \overline{\boldsymbol{w}}$. Note that $\overline{w}$ denotes the complex conjugate of w.

The HDG approximations $(\boldsymbol{q}_h, u_h, \widehat{u}_h)$ in $\boldsymbol{V}_h^p \times W_h^p \times M_h^p$ are determined by requiring that

$$(\boldsymbol{q}_h, \boldsymbol{r})_{\mathcal{T}_h} + (u_h, \nabla \cdot \boldsymbol{r})_{\mathcal{T}_h} - \langle \widehat{u}_h, \boldsymbol{r} \cdot \boldsymbol{n} \rangle_{\partial \mathcal{T}_h} = 0, \tag{4a}$$

$$(\kappa \boldsymbol{q}_h, \nabla w)_{\mathcal{T}_h} - \langle \kappa \widehat{\boldsymbol{q}}_h \cdot \boldsymbol{n}, w \rangle_{\partial \mathcal{T}_h} + (\varrho u_h, w)_{\mathcal{T}_h} = (f, w)_{\mathcal{T}_h}, \tag{4b}$$

$$\langle \kappa \widehat{\boldsymbol{q}}_h \cdot \boldsymbol{n}, \mu \rangle_{\partial \mathcal{T}_h} + \langle \nu \widehat{u}_h, \mu \rangle_{\partial \mathcal{D}} = \langle g, \mu \rangle_{\partial \mathcal{D}}, \tag{4c}$$

hold for all $(\boldsymbol{r}, w, \mu)$ in $\boldsymbol{V}_h^p \times W_h^p \times M_h^p$, where the numerical flux $\widehat{\boldsymbol{q}}_h$ is defined as

$$\widehat{\boldsymbol{q}}_h = \boldsymbol{q}_h - \tau (u_h - \widehat{u}_h) \boldsymbol{n}, \qquad \text{on } \partial \mathcal{T}_h. \tag{5}$$

Here τ is the so-called *stabilization* parameter, a global constant with dimensions $\tau = 1/L$ where L is the characteristic lengthscale. We set $\tau = 1$ since we do not consider multiple physical scales in this work. Further discussions on τ may be found in [14,39]. By substituting (5) into (4) we obtain that $(\boldsymbol{q}_h, u_h, \widehat{u}_h) \in \boldsymbol{V}_h^p \times W_h^p \times M_h^p$ satisfies

$$(\boldsymbol{q}_h, \boldsymbol{r})_{\mathcal{T}_h} + (u_h, \nabla \cdot \boldsymbol{r})_{\mathcal{T}_h} - \langle \widehat{u}_h, \boldsymbol{r} \cdot \boldsymbol{n} \rangle_{\partial \mathcal{T}_h} = 0, \tag{6a}$$

$$(\kappa \boldsymbol{q}_h, \nabla w)_{\mathcal{T}_h} - \langle \kappa \boldsymbol{q}_h \cdot \boldsymbol{n} - \kappa \tau (u_h - \widehat{u}_h), w \rangle_{\partial \mathcal{T}_h} + (\varrho u_h, w)_{\mathcal{T}_h} = (f, w)_{\mathcal{T}_h}, \tag{6b}$$

$$\langle \kappa \boldsymbol{q}_h \cdot \boldsymbol{n} - \kappa \tau (u_h - \widehat{u}_h), \mu \rangle_{\partial \mathcal{T}_h} + \langle \nu \widehat{u}_h, \mu \rangle_{\partial \mathcal{D}} = \langle g, \mu \rangle_{\partial \mathcal{D}}, \tag{6c}$$

for all $(\boldsymbol{r}, w, \mu)$ in $\boldsymbol{V}_h^p \times W_h^p \times M_h^p$. This completes the definition of the HDG method.

The above weak formulation of the HDG method involves three field variables, namely, $\boldsymbol{q}_h$, u_h, and $\widehat{u}_h$. However, the first two equations (6a) and (6b) allow us to write both $\boldsymbol{q}_h$ and u_h in terms of $\widehat{u}_h$ at the element level due to the fact that our approximation spaces are discontinuous. Therefore, we can substitute both $\boldsymbol{q}_h$ and u_h from the first two equations into the last equation (6c) to obtain a weak formulation in terms of $\widehat{u}_h$ only: find $\widehat{u}_h \in M_h^p$ such that

$$\widehat{a}_h(\widehat{u}_h, \mu) = \widehat{b}_h(\mu), \qquad \forall \mu \in M_h^p. \tag{7}$$

Here we omit the derivation of the bilinear form $\widehat{a}_h$ and the linear functional $\widehat{b}_h$. Instead we refer the reader to [39] for a detailed discussion. The reduced weak formulation (7) gives rise to the following linear system

$$\widehat{\mathbf{A}}\widehat{\mathbf{u}} = \widehat{\mathbf{b}}, \tag{8}$$

where $\widehat{\mathbf{u}}$ is the vector containing the degrees of freedom of $\widehat{u}_h$. Because $\widehat{u}_h$ is single valued on the faces of the finite element mesh, it has significantly fewer degrees of freedom than u_h. As a result, the global matrix system (8) of the HDG method can be much smaller than that of other DG methods. This results in significant savings in terms of computational time and memory storage.

It turns out that although the formulation (7) results in the smallest possible system, it is not ideal to use it as the starting point for our reduced basis method. Substituting the first two equations (6a) and (6b) into the last equation (6c) results in the inverse of the material coefficients κ and ϱ, which renders the bilinear form $\widehat{a}_h$ *nonaffine* in the material coefficients. Although nonaffine parameter dependence can be treated by using the empirical interpolation method [4] or the best points interpolation method [37], such treatment incurs additional cost and is unnecessary. We are going to derive a new weak formulation of the HDG method, which is suited for the reduced basis method.

2.3. A new weak formulation of the HDG method

We begin by deriving a weak formulation of the HDG method upon which our reduced basis method is constructed. To this end, we introduce two lifting operators $\boldsymbol{l} : W_h^p \to \boldsymbol{V}_h^p$ and $\boldsymbol{m} : M_h^p \to \boldsymbol{V}_h^p$ defined as

$$(\boldsymbol{l}(w), \boldsymbol{r})_{\mathcal{T}_h} = -(w, \nabla \cdot \boldsymbol{r})_{\mathcal{T}_h}, \qquad \forall \boldsymbol{r} \in \boldsymbol{V}_h^p, \tag{9a}$$

$$(\boldsymbol{m}(\mu), \boldsymbol{r})_{\mathcal{T}_h} = \langle \mu, \boldsymbol{r} \cdot \boldsymbol{n} \rangle_{\partial \mathcal{T}_h}, \qquad \forall \boldsymbol{r} \in \boldsymbol{V}_h^p. \tag{9b}$$

It thus follows from (6a) and (9) that we can express $\boldsymbol{q}_h$ as a function of u_h and $\widehat{u}_h$ as

$$\boldsymbol{q}_h = \boldsymbol{l}(u_h) + \boldsymbol{m}(\widehat{u}_h). \tag{10}$$

By substituting (10) into (6b) and (6c) we arrive at the following weak formulation: find $(u_h, \widehat{u}_h) \in W_h^p \times M_h^p$ such that

$$\big(\kappa (\boldsymbol{l}(u_h) + \boldsymbol{m}(\widehat{u}_h)), \nabla w \big)_{\mathcal{T}_h} - \big\langle \kappa (\boldsymbol{l}(u_h) + \boldsymbol{m}(\widehat{u}_h)) \cdot \boldsymbol{n} - \kappa \tau (u_h - \widehat{u}_h), w \big\rangle_{\partial \mathcal{T}_h} + (\varrho u_h, w)_{\mathcal{T}_h} = (f, w)_{\mathcal{T}_h},$$

$$\big\langle \kappa (\boldsymbol{l}(u_h) + \boldsymbol{m}(\widehat{u}_h)) \cdot \boldsymbol{n} - \kappa \tau (u_h - \widehat{u}_h), \mu \big\rangle_{\partial \mathcal{T}_h} + \langle \nu \widehat{u}_h, \mu \rangle_{\partial \mathcal{D}} = \langle g, \mu \rangle_{\partial \mathcal{D}},$$

for all (w, μ) in $W_h^p \times M_h^p$. By setting the $\mathcal{N}$-dimensional approximation space to be $\boldsymbol{W}_h^p := W_h^p \times M_h^p$, $\boldsymbol{u}_h := (u_h, \widehat{u}_h)$, and $\boldsymbol{w} := (w, \mu)$ we obtain that $\boldsymbol{u}_h \in \boldsymbol{W}_h^p$, satisfies

$$a_h(\boldsymbol{u}_h, \boldsymbol{w}; (\kappa, \varrho, \nu)) = b_h(\boldsymbol{w}), \qquad \forall \boldsymbol{w} \in \boldsymbol{W}_h^p, \tag{11}$$

where the bilinear form a_h and the linear functional b_h are given by

$$a_h(\boldsymbol{v}, \boldsymbol{w}; (\kappa, \varrho, \nu)) = \big(\kappa (\boldsymbol{l}(v) + \boldsymbol{m}(\eta)), \nabla w \big)_{\mathcal{T}_h} - \big\langle \kappa (\boldsymbol{l}(v) + \boldsymbol{m}(\eta)) \cdot \boldsymbol{n} - \kappa \tau (v - \eta), w \big\rangle_{\partial \mathcal{T}_h}$$

$$+ (\varrho v, w)_{\mathcal{T}_h} + \big\langle \kappa (\boldsymbol{l}(v) + \boldsymbol{m}(\eta)) \cdot \boldsymbol{n} - \kappa \tau (v - \eta), \mu \big\rangle_{\partial \mathcal{T}_h} + \langle \nu \eta, \mu \rangle_{\partial \mathcal{D}}, \tag{12a}$$

$$b_h(\boldsymbol{w}) = (f, w)_{\mathcal{T}_h} + \langle g, \mu \rangle_{\partial \mathcal{D}}, \tag{12b}$$

for all $\boldsymbol{v} := (v, \eta) \in \boldsymbol{W}_h^p$ and $\boldsymbol{w} := (w, \mu) \in \boldsymbol{W}_h^p$. We note that the bilinear form (12a) is affine in $\boldsymbol{y} = (\kappa, \varrho, \nu)$. Furthermore, if we select $\boldsymbol{r} = \kappa (\boldsymbol{l}(v) + \boldsymbol{m}(\eta))$ in (9) and substitute into (12a), we recover a symmetric form, which is also coercive

provided $\kappa, \varrho, \nu > 0$. Henceforth, the choice $(\kappa, \varrho, \nu) = (1, 1, 1)$ allows us to equip the approximation space $\boldsymbol{W}_h^p$ with the inner product $(\boldsymbol{v}, \boldsymbol{w})_{\boldsymbol{W}} := a_h(\boldsymbol{v}, \boldsymbol{w}; (\kappa, \varrho, \nu) = (1, 1, 1))$ and the induced norm $\|\boldsymbol{w}\|_{\boldsymbol{W}} = \sqrt{(\boldsymbol{w}, \boldsymbol{w})_{\boldsymbol{W}}}$.

We now substitute the expression of κ from (2) into (12) to express a_h as

$$a_h\left(\boldsymbol{v}, \boldsymbol{w}; (\kappa, \varrho, \nu)\right) = a_h(\boldsymbol{v}, \boldsymbol{w}; \boldsymbol{y}) = a_h^0(\boldsymbol{v}, \boldsymbol{w}) + \sum_{q=1}^{Q} y_q a_h^q(\boldsymbol{v}, \boldsymbol{w}), \tag{13}$$

where the bilinear forms are given by $a_h^q(\boldsymbol{v}, \boldsymbol{w}) := a_h\left(\boldsymbol{v}, \boldsymbol{w}; (\psi_q, 0, 0)\right)$ for $1 \leq q \leq Q$ and $a_h^0(\boldsymbol{v}, \boldsymbol{w}) := a_h\left(\boldsymbol{v}, \boldsymbol{w}; (\bar{\kappa}, \varrho, \nu)\right)$. Therefore, we can write the weak formulation (11) as follows: for any $\boldsymbol{y} \in \Lambda$, $\boldsymbol{u}_h(\boldsymbol{y}) \in \boldsymbol{W}_h^p$ satisfies

$$a_h(\boldsymbol{u}_h, \boldsymbol{w}; \boldsymbol{y}) = b_h(\boldsymbol{w}), \qquad \forall \boldsymbol{w} \in \boldsymbol{W}_h^p. \tag{14}$$

Finally, we evaluate our realization output as

$$s_h(\boldsymbol{y}) = \ell_h(\boldsymbol{u}_h(\boldsymbol{y})),$$

where the linear functional ℓ_h is obtained from the HDG discretization of ℓ. The key point of the new HDG formulation (14) for an efficient performance of the reduced basis method is the affine representation (13). This aspect is of crucial importance, and the main reason we prefer (14) to the reduced weak formulation (7) for constructing the reduced basis approximation. Furthermore, the new formulation is optimal in terms of degrees of freedom, since we no longer account for the gradient $\boldsymbol{q}_h$. Finally, even though the parameter independent matrices arising from (13) are used for the reduced basis approximation, the solution $\boldsymbol{u}_h$ is never computed as the solution of the full system (14). Instead, we can invoke again discontinuity of the approximation spaces to write u_h in terms of $\widehat{u}_h$. This common strategy in HDG methods enables us to solve for the global degrees of freedom of $\widehat{u}_h$ only and then recover u_h efficiently.

2.4. Monte Carlo sampling with the HDG method

We are interested in evaluating statistics of the output $s_h(\boldsymbol{y})$ such as its expectation and variance. Let $Y_M = \{\boldsymbol{y}_m \in \Lambda, 1 \leq m \leq M\}$ be a set of random samples drawn in the parameter space Λ with the probability density function $\rho(\boldsymbol{y})$. We evaluate the following outputs

$$s_h(\boldsymbol{y}_m) = \ell_h(\boldsymbol{u}_h(\boldsymbol{y}_m)), \qquad m = 1, \ldots, M. \tag{15}$$

The Monte Carlo-HDG (MC-HDG) estimates of the expectation $E[s]$ and variance $V[s]$ can be computed, respectively, as

$$E_M[s_h] = \frac{1}{M} \sum_{m=1}^{M} s_h(\boldsymbol{y}_m), \qquad V_M[s_h] = \frac{1}{M-1} \sum_{m=1}^{M} \left(E_M[s_h] - s_h(\boldsymbol{y}_m)\right)^2. \tag{16}$$

We shall assume that $s_h(\boldsymbol{y})$ is indistinguishable from $s(\boldsymbol{y})$ for any $\boldsymbol{y} \in \Lambda$. Moreover, it is a known result that the estimators in (16) are unbiased and converge in distribution to

$$E[s_h] - E_M[s_h] \xrightarrow{d} N\left(0; \frac{V[s_h]}{M}\right), \qquad V[s_h] - V_M[s_h] \xrightarrow{d} N\left(0; \frac{V[(s_h - E[s_h])^2]}{M}\right).$$

Confidence intervals can be constructed employing the central limit theorem (CLT), that is for all $a > 0$ we have

$$\lim_{M \to \infty} \Pr\left(\left|E[s_h] - E_M[s_h]\right| \leq \Delta_{h,M}^E\right) = \operatorname{erf}\left(\frac{a}{\sqrt{2}}\right), \tag{17a}$$

$$\lim_{M \to \infty} \Pr\left(\left|V[s_h] - V_M[s_h]\right| \leq \Delta_{h,M}^V\right) = \operatorname{erf}\left(\frac{a}{\sqrt{2}}\right), \tag{17b}$$

where

$$\Delta_{h,M}^E = a \sqrt{\frac{V_M[s_h]}{M}}, \qquad \Delta_{h,M}^V = a \sqrt{\frac{V_M[(s_h - E_M[s_h])^2]}{M}}. \tag{18}$$

Therefore, in order to guarantee that $\left|E[s_h] - E_M[s_h]\right|$ is bounded by a specified error tolerance ϵ_{tol} with a high probability (say, greater than 0.95), we need to take $a \geq 1.96$ and $M \geq a^2 V_M[s_h]/\epsilon_{\text{tol}}^2$. As a result, M can be very large for a small error tolerance. Hence, the evaluations (15)–(16) can be very demanding.

The remaining goals of this paper are as follows. On one hand, we develop a reduced basis (RB) method for rapid reliable approximation of the stochastic HDG output $s_h(\boldsymbol{y})$ for any given parameter vector $\boldsymbol{y}$ in Λ. On the other hand, we develop a multilevel variance reduction method to accelerate the convergence of the Monte Carlo simulation by exploiting the exponentially fast convergence of the RB output to the high-fidelity HDG output as a function of the RB dimension. These two ingredients enable very fast reliable computation of the statistical outputs at a computational cost which is several orders of magnitude less expensive than that of the MC-HDG approach. We describe the reduced basis approach in Section 3 and the multilevel variance reduction method in Section 4.

3. Reduced basis method

We consider a "primal–dual" formulation [47,53] particularly well-suited to good approximation and error characterization of the output. To this end, we introduce the dual problem of (14): given $\boldsymbol{y} \in \Lambda$, the dual solution $\boldsymbol{\phi}_h(\boldsymbol{y}) \in \boldsymbol{W}_h^p$ satisfies

$$a_h(\boldsymbol{v}, \boldsymbol{\phi}_h; \boldsymbol{y}) = -\ell_h(\boldsymbol{v}), \qquad \forall \boldsymbol{v} \in \boldsymbol{W}_h^p.$$

The dual problem plays an important role in improving the convergence rate of both the RB output and associated error bound.

We next assume that we are given orthonormalized basis functions $\boldsymbol{\zeta}_n^{\mathrm{pr}}, \boldsymbol{\zeta}_n^{\mathrm{du}} \in \boldsymbol{W}_h^p$, $1 \leq n \leq N_{\max}$, such that $(\boldsymbol{\zeta}_m^{\mathrm{pr}}, \boldsymbol{\zeta}_n^{\mathrm{pr}})_{\boldsymbol{W}} = (\boldsymbol{\zeta}_m^{\mathrm{du}}, \boldsymbol{\zeta}_n^{\mathrm{du}})_{\boldsymbol{W}} = \delta_{mn}, 1 \leq m, n \leq N_{\max}$. We define the associated hierarchical RB spaces as

$$\boldsymbol{W}_N^{\mathrm{pr}} = \mathrm{span}\{\boldsymbol{\zeta}_n^{\mathrm{pr}}, 1 \leq n \leq N\}, \quad \boldsymbol{W}_N^{\mathrm{du}} = \mathrm{span}\{\boldsymbol{\zeta}_n^{\mathrm{du}}, 1 \leq n \leq N\}, \quad N = 1, \dots, N_{\max}.$$

In practice, the spaces $\boldsymbol{W}_N^{\mathrm{pr}}$ and $\boldsymbol{W}_N^{\mathrm{du}}$ consist of orthonormalized primal and dual solutions $\boldsymbol{\zeta}_n^{\mathrm{pr}}, \boldsymbol{\zeta}_n^{\mathrm{du}}$ at selected parameter values generated by a Greedy sampling procedure [28,55,61]. For our present purpose, however, $\boldsymbol{W}_N^{\mathrm{pr}}$ and $\boldsymbol{W}_N^{\mathrm{du}}$ can in fact represent any sequence of (low-dimensional) hierarchical approximation spaces. We then apply the Galerkin projection for both the primal and dual problems: Given $\boldsymbol{y} \in \Lambda$, we find a primal RB approximation $\boldsymbol{u}_N(\boldsymbol{y}) \in \boldsymbol{W}_N^{\mathrm{pr}}$ satisfying

$$a_h(\boldsymbol{u}_N(\boldsymbol{y}), \boldsymbol{w}; \boldsymbol{y}) = b_h(\boldsymbol{w}), \qquad \forall \boldsymbol{w} \in \boldsymbol{W}_N^{\mathrm{pr}}, \tag{19}$$

and a dual RB approximation $\boldsymbol{\phi}_N(\boldsymbol{y}) \in \boldsymbol{W}_N^{\mathrm{du}}$ satisfying

$$a_h(\boldsymbol{w}, \boldsymbol{\phi}_N(\boldsymbol{y}); \boldsymbol{y}) = -\ell_h(\boldsymbol{w}), \qquad \forall \boldsymbol{w} \in \boldsymbol{W}_N^{\mathrm{du}}.$$

We can now evaluate the RB realization output as

$$s_N(\boldsymbol{y}) = \ell_h(\boldsymbol{u}_N(\boldsymbol{y})) + a_h(\boldsymbol{u}_N(\boldsymbol{y}), \boldsymbol{\phi}_N(\boldsymbol{y}); \boldsymbol{y}) - b_h(\boldsymbol{\phi}_N(\boldsymbol{y})).$$

As discussed below, the online computational cost of evaluating the RB output depends only on N and Q. Hence, for small N and Q, the RB approximation can be significantly less expensive than the HDG approximation.

The RB output is then used as an approximation to the HDG output in the Monte Carlo simulation. The Monte Carlo-Reduced Basis (MC-RB) estimates of the expectation and variance of the output of interest are given by

$$E_M[s_N] = \frac{1}{M} \sum_{m=1}^{M} s_N(\boldsymbol{y}_m), \qquad V_M[s_N] = \frac{1}{M-1} \sum_{m=1}^{M} \left(E_M[s_N] - s_N(\boldsymbol{y}_m)\right)^2$$

for the same set of samples $Y_M = \{\boldsymbol{y}_m \in \Lambda, 1 \leq m \leq M\}$. Since the RB approximation is constructed upon the HDG approximation these quantities actually approximate the MC-HDG estimates. We next develop *a posteriori* error bounds for our MC-RB estimates relative to the MC-HDG estimates.

3.1. A posteriori *error estimation*

We note from (19) that the residuals $r_h^{\mathrm{pr}}(\boldsymbol{w}; \boldsymbol{y})$ and $r_h^{\mathrm{du}}(\boldsymbol{w}; \boldsymbol{y})$ associated with $\boldsymbol{u}_N(\boldsymbol{y})$ and $\boldsymbol{\phi}_N(\boldsymbol{y})$, respectively, are given by

$$r_h^{\mathrm{pr}}(\boldsymbol{w}; \boldsymbol{y}) = b_h(\boldsymbol{w}) - a_h(\boldsymbol{u}_N(\boldsymbol{y}), \boldsymbol{w}; \boldsymbol{y}), \quad r_h^{\mathrm{du}}(\boldsymbol{w}; \boldsymbol{y}) = -\ell_h(\boldsymbol{w}) - a_h(\boldsymbol{w}, \boldsymbol{\phi}_N(\boldsymbol{y}); \boldsymbol{y}),$$

for all $\boldsymbol{w} \in \boldsymbol{W}_h^p$. The dual norm of the primal residual and the dual norm of the dual residual are given by

$$\|r_h^{\mathrm{pr}}(\cdot; \boldsymbol{y})\|_{\boldsymbol{W}'} = \sup_{\boldsymbol{w} \in \boldsymbol{W}_h^p} \frac{r_h^{\mathrm{pr}}(\boldsymbol{w}; \boldsymbol{y})}{\|\boldsymbol{w}\|_{\boldsymbol{W}}}, \qquad \|r_h^{\mathrm{du}}(\cdot; \boldsymbol{y})\|_{\boldsymbol{W}'} = \sup_{\boldsymbol{w} \in \boldsymbol{W}_h^p} \frac{r_h^{\mathrm{du}}(\boldsymbol{w}; \boldsymbol{y})}{\|\boldsymbol{w}\|_{\boldsymbol{W}}}.$$

It is a standard result [47,55] that

$$\|\boldsymbol{u}_h(\boldsymbol{y}) - \boldsymbol{u}_N(\boldsymbol{y})\|_{\boldsymbol{W}} \leq \Delta_N^{\mathrm{pr}}(\boldsymbol{y}) \equiv \frac{\|r_h^{\mathrm{pr}}(\cdot; \boldsymbol{y})\|_{\boldsymbol{W}'}}{\tilde{\beta}(\boldsymbol{y})},$$

$$\|\boldsymbol{\phi}_h(\boldsymbol{y}) - \boldsymbol{\phi}_N(\boldsymbol{y})\|_{\boldsymbol{W}} \leq \Delta_N^{\mathrm{du}}(\boldsymbol{y}) \equiv \frac{\|r_h^{\mathrm{du}}(\cdot; \boldsymbol{y})\|_{\boldsymbol{W}'}}{\tilde{\beta}(\boldsymbol{y})},$$

$$|s_h(\boldsymbol{y}) - s_N(\boldsymbol{y})| \leq \Delta_N^s(\boldsymbol{y}) \equiv \tilde{\beta}(\boldsymbol{y}) \Delta_N^{\mathrm{pr}}(\boldsymbol{y}) \Delta_N^{\mathrm{du}}(\boldsymbol{y}),$$

where $\tilde{\beta}(\boldsymbol{y})$ is a positive lower bound for the Babuška "inf–sup" stability constant $\beta_h(\boldsymbol{y})$ defined as

$$0 < \beta_h(\boldsymbol{y}) \equiv \inf_{\boldsymbol{w} \in \boldsymbol{W}_h^p} \sup_{\boldsymbol{v} \in \boldsymbol{W}_h^p} \frac{a_h(\boldsymbol{w}, \boldsymbol{v}; \boldsymbol{y})}{\|\boldsymbol{w}\|_{\boldsymbol{W}} \|\boldsymbol{v}\|_{\boldsymbol{W}}},$$

that is, the minimum (generalized) singular value associated with the differential operator. It is critical to note that the output error (and output error bound) vanishes as the *product* of the primal and dual error (bounds), and hence much more rapidly than either the primal or dual error.

It thus follows that we can bound the errors in the MC-RB estimates relative to the MC-HDG estimates as

$$|E_M[s_h] - E_M[s_N]| \leq \frac{1}{M} \sum_{m=1}^{M} |s_h(\boldsymbol{y}_m) - s_N(\boldsymbol{y}_m)| \leq \frac{1}{M} \sum_{m=1}^{M} \Delta_N^s(\boldsymbol{y}_m) \equiv \Delta_{N,M}^E, \tag{20}$$

and

$$
\begin{aligned}
\left| V_M[s_h] - V_M[s_N] \right| &= \frac{1}{M-1} \left| \sum_{m=1}^{M} \left(E_M[s_h] - s_h(\boldsymbol{y}_m) \right)^2 - \left(E_M[s_N] - s_N(\boldsymbol{y}_m) \right)^2 \right| \\
&= \frac{1}{M-1} \left| \sum_{m=1}^{M} \left(s_h(\boldsymbol{y}_m) - s_N(\boldsymbol{y}_m) - E_M[s_h] + E_M[s_N] \right) \left(s_h(\boldsymbol{y}_m) + s_N(\boldsymbol{y}_m) \right) \right| \\
&\leq \frac{1}{M-1} \sum_{m=1}^{M} \left(\left| s_h(\boldsymbol{y}_m) - s_N(\boldsymbol{y}_m) \right| + \left| E_M[s_h] - E_M[s_N] \right| \right) \left| s_h(\boldsymbol{y}_m) + s_N(\boldsymbol{y}_m) \right| \\
&\leq \frac{1}{M-1} \sum_{m=1}^{M} \left(\Delta_N^s(\boldsymbol{y}_m) + \Delta_{N,M}^E \right) \left(\Delta_N^s(\boldsymbol{y}_m) + 2 \left| s_N(\boldsymbol{y}_m) \right| \right) \equiv \Delta_{N,M}^V.
\end{aligned}
\tag{21}
$$

It should be stated that this error bound is rather pessimistic, and that a more precise bound can be obtained by introducing suitable dual problems to recover a quadratically convergent bound for the variance, as reported in [30]. We can also bound the difference between the RB expected value and the true value. To this end, we note from the triangle inequality that

$$\left| E[s_h] - E_M[s_N] \right| \leq \left| E[s_h] - E_M[s_h] \right| + \left| E_M[s_h] - E_M[s_N] \right|. \tag{22}$$

Following from (17a), (18), (20), (21), and (22) we define the error bound

$$\widetilde{\Delta}_{N,M}^E = a \sqrt{\frac{(V_M[s_N] + \Delta_{N,M}^V)}{M}} + \Delta_{N,M}^E \tag{23}$$

such that

$$\lim_{M \to \infty} \text{Pr} \left(\left| E[s_h] - E_M[s_N] \right| \leq \widetilde{\Delta}_{N,M}^E \right) \geq \text{erf} \left(\frac{a}{\sqrt{2}} \right).$$

Clearly, the error bound (23) comprises two terms: the first term is due to the MC sampling, while the second term is due to the RB approximation.

3.2. Computational strategy

The linearity and parametric affinity of the problem allow for an efficient Offline–Online decomposition strategy. The Offline stage − parameter independent, computationally intensive but performed only once − comprises the greedy search for the selection of parameter values, the computation of snapshots ζ_n^{pr}, ζ_n^{du}, $1 \leq n \leq N_{\max}$ associated with the HDG approximation space at the selected parameter values and the formation and storage of several parameter-independent small matrices and vectors. The Online stage − parameter dependent, performed multiple times − evaluates $s_N(\boldsymbol{y})$, $\Delta_N^s(\boldsymbol{y})$ for any new $\boldsymbol{y}$ with complexity $\mathcal{O}\left(2N^3 + 2(Q+1)^2 N^2\right)$ independent of the dimension $\mathcal{N}$ of the HDG approximation space. The implications are twofold: first, if N and Q are indeed small, we shall achieve very fast output evaluation, usually several orders of magnitude faster than the HDG output; second, we may choose the HDG approximation very conservatively − to effectively eliminate the error between the exact output and HDG output − without adversely affecting the Online (marginal) cost. We refer the reader to [47,54] for a more thorough description of the Offline–Online procedure.

It is clear that the error in the RB expected value and its error bound depend on N and Q as well as on M. Typically, both the error and its error bound decrease very rapidly as a function of N, but very slowly as a function of the number of samples M. Hence, M should be chosen large, while N can be chosen to be much smaller. Indeed, the (Online) computational cost to evaluate the RB expected value $E_M[s_N]$ and its error bound $\widetilde{\Delta}_{N,M}^E$ scales as $\mathcal{O}\left(2M(N^3 + 2(Q+1)^2 N^2)\right)$. Since both Q and N are typically very small, the RB method can effect significant savings relative to the HDG method. Nevertheless, its performance can be affected by the accuracy of the RB outputs and the sharpness of the RB error bounds.

4. Model and variance reduction method

4.1. Control variates principle

We first review the essential idea of control variates, which will serve as a building block for our method. Let X be a random variable. We would like to estimate the expected value of X. Suppose that we have another random variable Y and that its expected value $E[Y]$ is either known or inexpensive to compute. We then introduce a new random variable

$$X^* = X + \gamma (E[Y] - Y),$$

where γ is a deterministic coefficient. It is obvious that $E[X^*] = E[X]$ for any choice of γ. However, the variance of X^* is different from that of X. Specifically, we have

$$V[X^*] = V[X] + \gamma^2 V[Y] - 2\gamma \, \text{Cov}(X, Y),$$

where $\text{Cov}(X, Y) = E[XY] - E[X]E[Y]$ is the covariance of X and Y. It can be easily shown that the following choice

$$\gamma = \frac{\text{Cov}(X, Y)}{V[Y]}$$

is optimal in the sense that it minimizes the variance of X^*. With this choice, we have

$$V[X^*] = V[X](1 - \rho^2(X, Y)),$$

where $\rho(X, Y) = \text{Cov}(X, Y)/\sqrt{V[X]V[Y]}$ is the correlation coefficient of X and Y. It is clear that if X and Y are highly correlated (i.e., $\rho(X, Y)$ is close to ± 1) then $V[X^*]$ is much smaller than $V[X]$. In that case, the MC simulation of $E[X^*]$ converges significantly faster than that of $E[X]$ according to the CLT.

In summary, control variate methods try to estimate $E[X]$ by using the "surrogate" expected value $E[Y]$ and sampling the reduced variance variable X^*. When the same principle is applied recursively to estimate $E[Y]$, the resulting method is called multilevel control variates.

4.2. Two-level Monte Carlo sampling

We now apply the above idea to compute an estimate of $E[s_h]$, where $s_h(\mathbf{y})$ is the stochastic output obtained by using the HDG method to solve the underlying stochastic PDE as described in Section 2. To achieve this goal, we introduce

$$s_h^*(\mathbf{y}) = s_h(\mathbf{y}) + \gamma (E[s_{N_1}] - s_{N_1}(\mathbf{y})),$$

where $s_{N_1}(\mathbf{y})$ is the RB output developed in Section 3 for some $N_1 \in [1, N_{\max}]$. Because $s_{N_1}(\mathbf{y})$ generally approximates $s_h(\mathbf{y})$ very well, the two outputs are highly correlated. Therefore, we choose $\gamma = 1$ as we expect that the optimal value of γ is close to 1. With this choice, we obtain

$$E[s_h] = E[s_h^*] = E[s_h - s_{N_1}] + E[s_{N_1}]. \tag{24}$$

The underlying premise here is that the two expectation terms on the right hand side can be computed efficiently by MC simulations owing to variance reduction and model reduction: the first term requires a small number of samples because its variance is generally very small, while the second term is less expensive to evaluate because it involves the RB output.

In particular, let $Y_{M_0}^0 = \{\mathbf{y}_m^0 \in \Lambda, 1 \le m \le M_0\}$ and $Y_{M_1}^1 = \{\mathbf{y}_m^1 \in \Lambda, 1 \le m \le M_1\}$ be two independent sets of random samples drawn in Λ with the probability density function $\rho(\mathbf{y})$. We calculate our Model and Variance Reduction (MVR) unbiased estimate of $E[s_h]$ as

$$E_{M_0, M_1}[s_h] = E_{M_0}[s_h - s_{N_1}] + E_{M_1}[s_{N_1}], \tag{25}$$

where

$$E_{M_0}[s_h - s_{N_1}] = \frac{1}{M_0} \sum_{m=1}^{M_0} \left(s_h(\mathbf{y}_m^0) - s_{N_1}(\mathbf{y}_m^0) \right), \qquad E_{M_1}[s_{N_1}] = \frac{1}{M_1} \sum_{m=1}^{M_1} s_{N_1}(\mathbf{y}_m^1) \tag{26}$$

We note that our approach computes an estimate of $E[s_h]$, while the MC-RB approach described in the previous section computes an estimate of $E[s_N]$.

Similarly, we exploit the control variates idea to compute an estimate of the true variance $V[s_h]$ given by

$$V_{M_0, M_1}[s_h] = E_{M_0}[\zeta_h - \zeta_{N_1}] + E_{M_1}[\zeta_{N_1}], \tag{27}$$

where $\zeta_h := \left(s_h - E_{M_0, M_1}[s_h] \right)^2$ and $\zeta_{N_1} := \left(s_{N_1} - E_{M_0, M_1}[s_h] \right)^2$ and the expectations in (27) are analogous to the expectations in (26). The variance estimate is negatively biased

$$E \left[V_{M_0, M_1}[s_h] - V[s_h] \right] = -\frac{V[s_h - s_{N_1}]}{M_0} - \frac{V[s_{N_1}]}{M_1},$$

as shown in Appendix A.

It remains to provide *a posteriori* estimates for the errors in the expectation and variance. Subtracting (25) from (24) we identify new random variables Z_0, Z_1 whose limiting distributions are normal, and since they are independent their sum is also normally distributed,

$$Z_0 = E[s_h - s_{N_1}] - E_{M_0}[s_h - s_{N_1}] \sim N\left(0; \frac{V[s_h - s_{N_1}]}{M_0}\right), \tag{28a}$$

$$Z_1 = E[s_{N_1}] - E_{M_1}[s_{N_1}] \sim N\left(0; \frac{V[s_{N_1}]}{M_1}\right), \tag{28b}$$

$$Z_0 + Z_1 = E[s_h] - E_{M_0,M_1}[s_h] \sim N\left(0; \frac{V[s_h - s_{N_1}]}{M_0} + \frac{V[s_{N_1}]}{M_1}\right). \tag{28c}$$

We invoke now the CLT to obtain an error estimate for the expectation error as

$$\lim_{M_0 \to \infty} \lim_{M_1 \to \infty} \mathrm{Pr}\left(\left|E[s_h] - E_{M_0,M_1}[s_h]\right| \le \Delta^E_{M_0,M_1}\right) = \mathrm{erf}\left(\frac{a}{\sqrt{2}}\right), \tag{29}$$

where

$$\Delta^E_{M_0,M_1} = a\sqrt{\frac{V_{M_0}[s_h - s_{N_1}]}{M_0} + \frac{V_{M_1}[s_{N_1}]}{M_1}}, \tag{30a}$$

$$V_{M_0}[s_h - s_{N_1}] = \frac{1}{M_0 - 1}\sum_{m=1}^{M_0}\left(E_{M_0}[s_h - s_{N_1}] - s_h(\boldsymbol{y}_m^0) + s_{N_1}(\boldsymbol{y}_m^0)\right)^2, \tag{30b}$$

$$V_{M_1}[s_{N_1}] = \frac{1}{M_1 - 1}\sum_{m=1}^{M_1}\left(E_{M_1}[s_{N_1}] - s_{N_1}(\boldsymbol{y}_m^1)\right)^2, \tag{30c}$$

and the variances in (28c) are estimated with their MC counterparts (30b)–(30c).

For the variance, we first define auxiliary variables $\widehat{\zeta_h} := (s_h - E[s_h])^2$ and $\widehat{\zeta}_{N_1} := (s_{N_1} - E[s_h])^2$ and the auxiliary variance

$$\widehat{V}_{M_0,M_1}[s_h] = E_{M_0}[\widehat{\zeta_h} - \widehat{\zeta}_{N_1}] + E_{M_1}[\widehat{\zeta}_{N_1}].$$

The MVR variance estimate in (27) can be rewritten as (see Appendix A)

$$V_{M_0,M_1}[s_h] = \widehat{V}_{M_0,M_1}[s_h] - \left(E[s_h] - E_{M_0,M_1}[s_h]\right)^2,$$

which implies

$$V_{M_0,M_1}[s_h] - V[s_h] = \left(\widehat{V}_{M_0,M_1}[s_h] - V[s_h]\right) - \left(E[s_h] - E_{M_0,M_1}[s_h]\right)^2. \tag{31}$$

Let us consider the two terms in the RHS in reverse order. Convergence in probability for the second term is guaranteed by (29), that is $E_{M_0,M_1}[s_h] - E[s_h] \xrightarrow{P} 0$. Furthermore, repeating the analysis in (28) for the first term leads to

$$\widehat{V}_{M_0,M_1}[s_h] - V[s_h] \sim N\left(0; \frac{V[\widehat{\zeta_h} - \widehat{\zeta}_{N_1}]}{M_0} + \frac{V[\widehat{\zeta}_{N_1}]}{M_1}\right). \tag{32}$$

Therefore the limiting distribution of $V_{M_0,M_1}[s_h] - V[s_h]$ is the same as the limiting distribution of $\widehat{V}_{M_0,M_1}[s_h] - V[s_h]$ (Slutzky's theorem), and the straightforward application of the CLT recovers

$$\lim_{M_0 \to \infty} \lim_{M_1 \to \infty} \mathrm{Pr}\left(\left|V[s_h] - V_{M_0,M_1}[s_h]\right| \le \Delta^V_{M_0,M_1}\right) = \mathrm{erf}\left(\frac{a}{\sqrt{2}}\right),$$

where

$$\Delta^V_{M_0,M_1} = a\sqrt{\frac{V_{M_0}[\zeta_h - \zeta_{N_1}]}{M_0} + \frac{V_{M_1}[\zeta_{N_1}]}{M_1}}, \tag{33a}$$

$$V_{M_0}[\zeta_h - \zeta_{N_1}] = \frac{1}{M_0 - 1}\sum_{m=1}^{M_0}\left(E_{M_0}[\zeta_h - \zeta_{N_1}] - \zeta_h(\boldsymbol{y}_m^0) + \zeta_{N_1}(\boldsymbol{y}_m^0)\right)^2, \tag{33b}$$

$$V_{M_1}[\zeta_{N_1}] = \frac{1}{M_1 - 1}\sum_{m=1}^{M_1}\left(E_{M_1}[\zeta_{N_1}] - \zeta_{N_1}(\boldsymbol{y}_m^1)\right)^2. \tag{33c}$$

The variances in (32) are again estimated with their MC simulations (33a), and $\widehat{\zeta}_h$, $\widehat{\zeta}_{N_1}$ are approximated with ζ_h and ζ_{N_1} by replacing $E[s_h]$ for its MVR estimate $E_{M_0,M_1}[s_h]$.

We would like to make two observations. First, the model and variance reduction approach described here requires M_0 realizations of the high-fidelity HDG output and M_1 realizations of the RB output, while the MC-RB approach described in the previous section requires M realizations of the RB output and its error bound. If we take $N_1 = N$ then it is reasonable to consider $M_1 \approx M$. Furthermore, we take M_0 such that the computational cost of M_0 realizations of the HDG output is commensurate with that of M realizations of the RB output bound. In this scenario, the two approaches have the same computational complexity. The advantage of the present approach is that it provides more accurate estimates than the MC-RB approach owing to the variance reduction. Second, unlike the MC-RB approach, the present approach does not require *a posteriori* error bounds for the RB output to obtain the error bounds for our estimates of the statistical outputs. As a result, our approach can be applied to problems for which *a posteriori* output bounds are either computationally expensive or theoretically difficult.

4.3. Multilevel Monte Carlo sampling

The method can be further generalized and improved by pursuing a multilevel control variate strategy. Given L different RB output models $s_{N_\ell}(\boldsymbol{y})$, $1 \leq \ell \leq L$, with $N_1 > N_2 > \ldots > N_L$,[3] we first express the expected value as

$$E[s_h] = E[s_h - s_{N_1}] + \sum_{\ell=1}^{L-1} E[s_{N_\ell} - s_{N_{\ell+1}}] + E[s_{N_L}].$$

We next introduce $L + 1$ independent sample sets $Y_{M_\ell}^\ell = \{\boldsymbol{y}_m^\ell \in \Lambda, 1 \leq m \leq M_\ell\}$, $0 \leq \ell \leq L$, which are drawn in Λ with probability density function $\rho(\boldsymbol{y})$. We then define our estimate of $E[s_h]$ as

$$E_{M_0,\ldots,M_L}[s_h] = E_{M_0}[s_h - s_{N_1}] + \sum_{\ell=1}^{L-1} E_{M_\ell}[s_{N_\ell} - s_{N_{\ell+1}}] + E_{M_L}[s_{N_L}].$$

Extending the analysis in (28) we apply the CLT to the multilevel case to obtain

$$\lim_{M_0 \to \infty} \cdots \lim_{M_L \to \infty} \Pr\left(\left| E[s_h] - E_{M_0,\ldots,M_L}[s_h] \right| \leq \Delta_{M_0,\ldots,M_L}^E \right) = \mathrm{erf}\left(\frac{a}{\sqrt{2}} \right),$$

$$\Delta_{M_0,\ldots,M_L}^E = a \sqrt{ \frac{V_{M_0}[s_h - s_{N_1}]}{M_0} + \sum_{\ell=1}^{L-1} \frac{V_{M_\ell}[s_{N_\ell} - s_{N_{\ell+1}}]}{M_\ell} + \frac{V_{M_L}[s_{N_L}]}{M_L} }.$$

Similarly, the estimate of the variance is defined as

$$V_{M_0,\ldots,M_L}[s_h] = E_{M_0}[\zeta_h - \zeta_{N_1}] + \sum_{\ell=1}^{L-1} E_{M_\ell}[\zeta_{N_\ell} - \zeta_{N_{\ell+1}}] + E_{M_L}[\zeta_{N_L}],$$

where the auxiliary variables are $\zeta_h := \left(s_h - E_{M_0,\ldots,M_L}[s_h] \right)^2$ and $\zeta_{N_\ell} := \left(s_{N_\ell} - E_{M_0,\ldots,M_L}[s_h] \right)^2$ for $\ell = 1, \ldots, L$. Combining the results in (31)–(32) with the CLT leads to the following error bound for the variance estimate

$$\lim_{M_0 \to \infty} \cdots \lim_{M_L \to \infty} \Pr\left(\left| V[s_h] - V_{M_0,\ldots,M_L}[s_h] \right| \leq \Delta_{M_0,\ldots,M_L}^V \right) = \mathrm{erf}\left(\frac{a}{\sqrt{2}} \right),$$

$$\Delta_{M_0,\ldots,M_L}^V = a \sqrt{ \frac{V_{M_0}[\zeta_h - \zeta_{N_1}]}{M_0} + \sum_{\ell=1}^{L-1} \frac{V_{M_\ell}[\zeta_{N_\ell} - \zeta_{N_{\ell+1}}]}{M_\ell} + \frac{V_{M_L}[\zeta_{N_L}]}{M_L} }.$$

Note that all expectations and variances are MC estimates through the sample sets $Y_{M_\ell}^\ell$ for $0 \leq \ell \leq L$.

We will refer to the general model and variance reduction method with a sequence of L reduced basis models as the L-MVR method. For clarity of notation, we shall identify $s_h - s_{N_1}$ as level 0, and the subsequent $s_{N_\ell} - s_{N_{\ell+1}}$ as level ℓ. The method allows us to transfer the computational burden from the higher-fidelity (expensive) outputs to the lower-fidelity (inexpensive) outputs. In particular, we can choose $N_1, N_2, \ldots, N_L$ so as to have $M_0 \ll M_1 \ll \ldots \ll M_L$. Hence, the number of evaluations of the higher-fidelity outputs are significantly smaller than those of the lower-fidelity outputs, thereby resulting in a significant reduction in the overall computational cost. Finally, we address the issue of how to determine the RB dimensions $N_1, N_2, \ldots, N_L$ and the number of samples $M_0, M_1, \ldots, M_L$ to achieve a specified error tolerance and minimize the computational cost.

---

[3] In our context, it is natural to number the levels from the finest RB approximation to the coarsest RB approximation because the finest RB level is closest to the HDG approximation.

4.4. Selection method

Let t_{N_ℓ} denote the (Online) wall time to compute the RB output $s_{N_\ell}(\mathbf{y})$ for $\ell \geq 1$, and t_h denote the wall time to compute the HDG output $s_h(\mathbf{y})$ for any given $\mathbf{y} \in \Lambda$. Note that t_{N_ℓ} depends on N_ℓ, while t_h depends on the finite element approximation spaces. The total (Online) wall time T_L of the L-MVR and the (Online) speedup π_L with respect to the MC-HDG method are given by

$$T_L = \left(t_h + t_{N_1}\right) M_0 + \sum_{\ell=1}^{L-1} M_\ell \left(t_{N_\ell} + t_{N_{\ell+1}}\right) + t_{N_L} M_L, \qquad \pi_L = \frac{t_h M}{T_L}. \tag{34}$$

We wish to find $(N_1, N_2, \ldots, N_L)$ and $(M_0, M_1, \ldots, M_L)$ so as to minimize T_L, while ensuring that $\Delta^E_{M_0,\ldots,M_L}$ is equal to a specified error tolerance ϵ_{tol}. This error condition is satisfied if we take

$$a^2 \frac{V_{M_L}[s_{N_L}]}{M_L} = w_L \epsilon_{\text{tol}}^2, \quad a^2 \frac{V_{M_0}[s_h - s_{N_1}]}{M_0} = w_0 \epsilon_{\text{tol}}^2, \quad a^2 \frac{V_{M_\ell}[s_{N_\ell} - s_{N_{\ell+1}}]}{M_\ell} = w_\ell \epsilon_{\text{tol}}^2, \quad \ell \geq 1, \tag{35}$$

for any given positive $w_\ell \in (0,1)$, $\ell = 0, \ldots, L$ such that $w_0 + w_1 + \ldots + w_L = 1$. The choice of the weights depends on how we would like to distribute the error among the levels. We combine expressions (34)–(35) to define the cost function

$$C_L = \frac{T_L \epsilon_{\text{tol}}^2}{a^2} = \frac{V_{M_0}[s_h - s_{N_1}]}{w_0}\left(t_h + t_{N_1}\right) + \sum_{\ell=1}^{L-1} \frac{V_{M_\ell}[s_{N_\ell} - s_{N_{\ell+1}}]}{w_\ell}\left(t_{N_\ell} + t_{N_{\ell+1}}\right) + t_{N_L} \frac{V_{M_L}[s_{N_L}]}{w_L}. \tag{36}$$

We need to determine $(M_0, M_1, \ldots, M_L)$ and $(N_1, N_2, \ldots, N_L)$ that minimize C_L. Unfortunately, this is a nonlinear integer optimization problem which is difficult to solve exactly. We thus solve an approximate problem as follows.

We first introduce a test sample set $Y_{\widehat{M}} = \{\widehat{\mathbf{y}}_m \in \Lambda, 1 \leq m \leq \widehat{M}\}$. We then precompute and store the HDG outputs $s_h(\widehat{\mathbf{y}}_m)$ for $m = 1, \ldots, \widehat{M}$ and the RB outputs $s_N(\widehat{\mathbf{y}}_m)$ for $m = 1, \ldots, \widehat{M}$ and $N = 1, \ldots, N_{\max}$. In addition, we also precompute and store t_h and t_N for $N = 1, \ldots, N_{\max}$. For any given strictly decreasing L-tuple $\mathbf{I} = (I_1, I_2, \ldots, I_L) \in [1, N_{\max}]^L$ and valid weights $\mathbf{w} = (w_0, \ldots, w_L)$, we can evaluate the equivalent cost function

$$\widehat{C}_L(\mathbf{I}, \mathbf{w}) = \sum_{\ell=0}^{L} \frac{\widehat{C}_L^\ell(\mathbf{I})}{w_\ell} = \frac{V_{\widehat{M}}[s_h - s_{I_1}]}{w_0}\left(t_h + t_{I_1}\right) + \sum_{\ell=1}^{L-1} \frac{V_{\widehat{M}}[s_{I_\ell} - s_{I_{\ell+1}}]}{w_\ell}\left(t_{I_\ell} + t_{I_{\ell+1}}\right) + t_{I_L} \frac{V_{\widehat{M}}[s_{I_L}]}{w_L}, \tag{37}$$

with $\mathcal{O}\left((L+1)\widehat{M}\right)$ operations count, where all the variances are computed using the test sample set $Y_{\widehat{M}}$. We now set

$$\mathbf{N} \equiv (N_1, N_2, \ldots, N_L) = \arg\min_{\mathbf{I}} \widehat{C}_L(\mathbf{I}, \mathbf{w}^{\mathbf{I}}),$$
$$\text{s.t.} \quad N_{\max} \geq I_1 > I_2 > \ldots > I_L \geq 1 \tag{38}$$

where the weights $\mathbf{w}^{\mathbf{I}}$ are the minimizers of the equivalent cost for any L-tuple $\mathbf{I}$, that is

$$\mathbf{w}^{\mathbf{I}} \equiv (w_0^{\mathbf{I}}, w_1^{\mathbf{I}}, \ldots, w_L^{\mathbf{I}}) = \arg\min_{\mathbf{w}} \widehat{C}_L(\mathbf{I}, \mathbf{w}),$$
$$\text{s.t.} \quad \sum_{\ell=0}^{L} w_\ell = 1, \quad w_\ell > 0. \tag{39}$$

The KKT conditions for (39) render the optimal weights for any L-tuple $\mathbf{I}$ as

$$w_\ell^{\mathbf{I}} = \frac{\sqrt{\widehat{C}_L^\ell(\mathbf{I})/\widehat{C}_L^0(\mathbf{I})}}{\displaystyle\sum_{\ell'=0}^{L} \sqrt{\widehat{C}_L^{\ell'}(\mathbf{I})/\widehat{C}_L^0(\mathbf{I})}}, \qquad \ell = 0, \ldots, L. \tag{40}$$

The minimization problem (38) can be approximately solved, for the weights defined in (40), by simply evaluating the cost function $\widehat{C}_L(\mathbf{I}, \mathbf{w}^{\mathbf{I}})$ for all feasible L-tuples $\mathbf{I}$ in $\mathcal{O}((L+1)\widehat{M}(N-L) \times \ldots \times (N-1)/L!)$ operations count. Even though we present here an optimal choice of the weights, any valid distribution can be employed.

Having determined the RB dimensions $\mathbf{N}$ and the weights $\mathbf{w}^{\mathbf{N}}$, we can now proceed with the MC simulations for all levels. We initially set $Y_{M_0}^0 = Y_{\widehat{M}}$, and thus reuse $s_h(\widehat{\mathbf{y}}_m)$, $m = 1, \ldots, \widehat{M}$. We then execute the MC processes for all the levels and enforce the error constraint $\Delta^E_{M_0,\ldots,M_L} = \epsilon_{\text{tol}}$ by adding new random parameters to the sample sets until the following inequalities

$$M_L \geq \frac{a^2 V_{M_L}[s_{N_L}]}{w_L^{\mathbf{N}} \epsilon_{\text{tol}}^2}, \quad M_0 \geq \frac{a^2 V_{M_0}[s_h - s_{N_1}]}{w_0^{\mathbf{N}} \epsilon_{\text{tol}}}, \quad M_\ell \geq \frac{a^2 V_{M_\ell}[s_{N_\ell} - s_{N_{\ell+1}}]}{w_\ell^{\mathbf{N}} \epsilon_{\text{tol}}}, \quad \ell = 1, \ldots, L-1,$$

(a) Exact solution $u(x, \boldsymbol{y})$ for several realizations of the diffusivity field.

(b) Average output error bound $\Delta^s_{N,\text{avg}}$ and average output error $\epsilon_{N,\text{avg}}$ for $\boldsymbol{y} \in Y_{\widehat{M}}$ vs. RB size N.
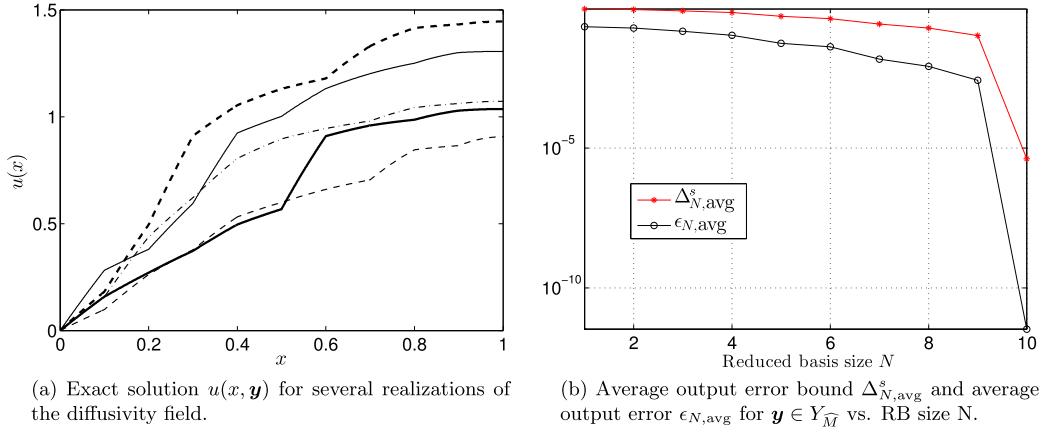
Fig. 1. Representative solutions and the RB convergence of the heat diffusion example.

are satisfied and the MC processes are terminated upon satisfaction of these conditions. Therefore, the sample sets $Y^{\ell}_{M_\ell}$, $\ell = 0, \ldots, L$ are continuously updated during the MC runs. Finally, to provide confidence in the application of the CLT we also need to enforce that M_ℓ are greater than a certain threshold, say 30.

Although we have assumed that the number of levels L is fixed, our approach also allows us to compare the computational costs for several values of L. Hence, we can determine not only the RB dimensions and the weights, but also the optimal number of levels, and it can be done efficiently evaluating expressions (37)–(39). This analysis provides inexpensive means to determine the optimal multilevel structure.

5. Numerical results

5.1. A coercive example: heat diffusion

In the first example, we consider the one dimensional steady-state heat equation in $\mathcal{D} = (0, 1)$:

$$-(\kappa u_x)_x = f(x), \quad \forall x \in \mathcal{D}, \tag{41a}$$

$$\kappa u_x = 0, \quad \text{on } x = 1, \tag{41b}$$

$$u = 0, \quad \text{on } x = 0, \tag{41c}$$

where $\kappa(x, \omega)$ is a piecewise constant function on a series of disjoint subdomains $\mathcal{D}_q = ((q-1)/Q, q/Q)$, $q = 1, \ldots, Q$, that is, $\kappa(x, \omega) = \sum_{q=1}^{Q} \kappa_q(\omega) \mathbf{1}_{\mathcal{D}_q}$, with $\kappa_q(\omega) \in [\gamma^-_q, \gamma^+_q]$ for all q. For this problem, we treat $\kappa_q(\omega)$ as i.i.d. uniform random variables; hence, we can write $\kappa(x, \omega) = \kappa(x, \boldsymbol{y}) = \sum_{q=1}^{Q} y_q \mathbf{1}_{\mathcal{D}_q}$, where y_q, $q = 1, \ldots, Q$ are i.i.d. random variables with uniform continuous distributions in the interval [0.1, 1]. The problem (40) has an analytic solution given by

$$u(x, \boldsymbol{y}) = \int_0^x \left(\frac{1}{\kappa(z, \boldsymbol{y})} \int_z^1 f(\xi) d\xi \right) dz. \tag{42}$$

The observable quantity is the average temperature on the domain, namely, $s(\boldsymbol{y}) = \int_0^1 u(x, \boldsymbol{y})$. The output s, its expectation $E[s]$ and variance $V[s]$ have closed analytic forms, thereby s will be here used instead of the HDG output s_h. Numerical results for a constant source term $f(x) = 1$ and $Q = 10$ are presented below.

We show in Fig. 1(a) different realizations of the exact solution in (42). Since the output is compliant, the dual problem coincides with the primal problem. We thus need to construct the reduced basis approximation for the primal problem only. Furthermore, since the bilinear form is coercive and the parameters are positive, we compute the stability constant $\beta_h(\boldsymbol{y})$ using a bound conditioner technique [62], which greatly simplifies the process. Our reduced basis is constructed with $N_{\text{max}} = 10$. We show in Fig. 1(b) the average output error $\epsilon_{N,\text{avg}}$ and the average output error bound $\Delta^s_{N,\text{avg}}$ as a function of N, where $\epsilon_{N,\text{avg}} = \sum_{\boldsymbol{y} \in Y_{\widehat{M}}} |s(\boldsymbol{y}) - s_N(\boldsymbol{y})| / \widehat{M}$ and $\Delta^s_{N,\text{avg}} = \sum_{\boldsymbol{y} \in Y_{\widehat{M}}} \Delta^s_N(\boldsymbol{y}) / \widehat{M}$, being $Y_{\widehat{M}}$ a test set of $\widehat{M} = 1000$ samples. We observe that the average output error and the average output error bound converge slowly up to $N = 9$ and drop rapidly at $N = 10$. This is because of the nature of the particular problem which requires $N = Q$ basis functions to capture all the possible solutions. When we use $N < Q$, we do not have enough basis functions to represent all the possible solutions, which in turn causes a slow convergence of the reduced basis approximation.

We now compare the performance of the MC-HDG method with a uniform mesh of $h = 1/10$, the MC-RB with $N = 9$, and the L-MVR with $L = 1$, $N_1 = 5$, $M_1 = M$ and $M_0 = M/10$ in estimating $E[s]$ and $V[s]$ as a function of M. For each M value

Table 1

The expectation error and its error bound for different values of M for the MC-HDG method, the MC-RB method and the 1-MVR method.

M	MC-HDG		MC-RB ($N = 9$)		1-MVR ($N_1 = 5$)	
	$\left|E[s] - E_M[s]\right|$	$\Delta_{h,M}^E$	$\left|E[s] - E_M[s_N]\right|$	$\widetilde{\Delta}_{N,M}^E$	$\left|E[s] - E_{M_0,M_1}[s]\right|$	Δ_{M_0,M_1}^E
10^2	2.10e−2	2.6e−2 a	2.19e−2	1.1e−1 + 7.0e−2 a	2.30e−2	2.9e−2 a
10^3	6.59e−2	8.3e−3 a	7.06e−3	1.1e−1 + 2.2e−2 a	7.30e−3	9.2e−3 a
10^4	2.08e−3	2.6e−3 a	3.21e−3	1.1e−1 + 7.0e−3 a	2.32e−3	2.9e−3 a
10^5	6.57e−4	8.3e−4 a	2.88e−3	1.1e−1 + 2.2e−3 a	7.29e−4	9.2e−4 a
10^6	2.11e−4	2.6e−4 a	2.90e−3	1.1e−1 + 7.0e−4 a	2.32e−4	2.9e−4 a

Table 2

The variance error and its error bound for different values of M for the MC-HDG method and the 1-MVR method.

M	MC-HDG		1-MVR ($N_1 = 5$)	
	$\left|V[s] - V_M[s]\right|$	$\Delta_{h,M}^V$	$\left|V[s] - V_{M_0,M_1}[s]\right|$	Δ_{M_0,M_1}^V
10^2	9.21e−3	1.1e−2 a	1.24e−2	1.5e−2 a
10^3	2.91e−3	3.6e−3 a	4.06e−3	5.0e−3 a
10^4	9.20e−4	1.2e−3 a	1.29e−3	1.6e−3 a
10^5	2.91e−4	3.6e−4 a	4.07e−4	5.1e−4 a
10^6	9.29e−5	1.1e−4 a	1.27e−4	1.6e−4 a

we repeat the simulations $H = 1000$ times, and present in Tables 1 and 2 the average values of the absolute errors and error bounds for the expectation and the variance respectively. We observe that the 1-MVR method significantly outperforms the MC-RB method. The improvement is noticeable when we increase M, since the MC-RB method stagnates around 2.9×10^{-3} whereas 1-MVR keeps reducing the error as the square root of the number of samples. The stagnation is caused by the inherent bias arising from the reduced basis method with $N = 9$, which provides outputs with a level of error of 3×10^{-3}, as seen in Fig. 1(b). With the MC-RB method we are unable to achieve more accurate estimators than the precision of the reduced basis output. Furthermore, the error bound of the MC-RB method, $\widetilde{\Delta}_{N,M}^E$ as defined in (23), is the sum of two terms: the first term $\Delta_{N,M}^E$ does not depend on the confidence level, whereas the second term $a\sqrt{(V_M[s_N] + \Delta_{N,M}^V)/M}$ does. We observe from Table 1 that increasing M does not improve the error bound of the MC-RB method since it is dominated by $\Delta_{N,M}^E$, which in many cases can be overly pessimistic. Variance estimations for the MC-RB are not included in Table 2, as they can only be worse than the expectation results.

The 1-MVR method does not suffer from this stagnation owing to the fact that it directly approximates $E[s]$ instead of $E[s_N]$. As a result, the expectation and variance error of the 1-MVR method can be made arbitrarily small. The same behavior is observed in the error bounds Δ_{M_0,M_1}^E and Δ_{M_0,M_1}^V defined in (30a) and (33a), which agree with the Monte Carlo dependence on the square root of the sample size. Even though the accuracy of the estimators and the sharpness of the bounds for 1-MVR is slightly worse than that of MC-HDG, the former performs ten times less full model evaluations than the latter. These numerical results show a considerable gain for model and variance reduction.

5.2. A noncoercive example: acoustic wave propagation

We consider a wave propagation problem as depicted in Fig. 2(a). A wave is excited by a Gaussian source term f centered at $\boldsymbol{x}_s$ and propagates through a heterogeneous medium $\kappa(\boldsymbol{x}, \boldsymbol{y})$. The governing equation for this model problem is given by

$$-\nabla \cdot (\kappa \nabla u) - k^2 u = f, \qquad \forall \boldsymbol{x} \in \mathcal{D},$$

$$\kappa \nabla u \cdot \boldsymbol{n} - iku = 0, \qquad \forall \boldsymbol{x} \in \partial \mathcal{D}_R,$$

$$\kappa \nabla u \cdot \boldsymbol{n} = 0, \qquad \forall \boldsymbol{x} \in \partial \mathcal{D}_N,$$

where $f = \dfrac{10}{\sqrt{2\pi}\sigma_s} \exp\left(-\dfrac{(x_1 - x_{s1})^2 + (x_2 - x_{s2})^2}{2\sigma_s^2}\right)$ for $\boldsymbol{x}_s = (-3, -16)$ and $\sigma_s = 0.25$ is the source term and $k = \sqrt{2}$ is the wavenumber. Here the physical domain is $\mathcal{D} = [-15, 15] \times [-20, 0]$. To describe the $\kappa(\boldsymbol{x}, \boldsymbol{y})$ field we use the example described in [9], namely

$$\kappa(\boldsymbol{x}, \boldsymbol{y}) = \bar{\kappa} + \sigma y_1 \sqrt{\frac{\lambda_0}{2}} + \sigma \sum_{n=1}^{8} \sqrt{\lambda_n} \left(\sin\left(n\pi \frac{x_1 + \ell/2}{\ell}\right) y_{2n} + \cos\left(n\pi \frac{x_1 + \ell/2}{\ell}\right) y_{2n+1}\right),$$

where

$$\sqrt{\lambda_n} = \left(\sqrt{\pi} L_c\right)^{1/2} \exp\left(-\frac{(n\pi L_c)^2}{8}\right), \quad n = 0, \ldots, 8,$$

(a) Geometry of the wave propagation problem. The source generates a wave that propagates through the medium.

(b) Triangular mesh of 1420 elements for the wave problem. Higher resolution is appreciated at the source and output locations.

(c) Random realization of the wave amplitude field, using elements of order $p = 4$ for a total of $\mathcal{N} = 32180$ degrees of freedom.

(d) Average output error bound $\Delta^s_{N,\text{avg}}$ and average output error $\epsilon_{N,\text{avg}}$ for $\boldsymbol{y} \in Y_{\widehat{M}}$ vs. RB size N.

Fig. 2. Problem specification, representative solution, and RB convergence of the wave propagation example.

for $\bar{\kappa} = 1, \sigma = 1/10$, and $L_c = 1/12$. Here the random variables y_n for $n = 1, \ldots, Q = 17$ are uncorrelated and uniformly distributed with zero mean and unit variance. Hence, we can write $\kappa(\boldsymbol{x}, \boldsymbol{y})$ in the form of the affine expansion (2) with $\Lambda = [-\sqrt{3}, \sqrt{3}]^Q$. We consider the following output

$$s(\boldsymbol{y}) = \frac{1}{\sqrt{2\pi}\sigma_0} \int_{\mathcal{D}} \Re(u(\boldsymbol{y})) \exp\left(-\frac{(x_1 - x_{O1})^2 + (x_2 - x_{O2})^2}{2\sigma_0^2}\right) d\boldsymbol{x},$$

for $\boldsymbol{x}_O = (5, -7)$ and $\sigma_0 = 0.25$, which corresponds to the real part of the amplitude at $\boldsymbol{x}_O$ regularized by a Gaussian field. The physical domain is discretized into a triangular mesh of 1420 elements as shown in Fig. 2(b) and polynomials of degree $p = 4$ are used to represent the numerical solution $u_h(\boldsymbol{y})$. Fig. 2(c) depicts a realization of the numerical solution obtained using the HDG method.

Since the exact values of $E[s_h]$ and $V[s_h]$ are not known, we approximate them using the MC-HDG method with a random sample set Y_{M^*} of $M^* = 6.5 \times 10^7$ and obtain $E_{M^*}[s_h] = 0.2576$ and $V_{M^*}[s_h] = 0.0596$ for a statistical error of 10^{-4} and 3.1×10^{-5} respectively, corresponding to 0.999 confidence level ($a = 3.3$). We are going to use $E[s_h] = E_{M^*}[s_h]$ and $V[s_h] = V_{M^*}[s_h]$ as the reference values to evaluate the performance of our method. We consider an error tolerance of 10^{-3} and a confidence level 0.95 for our estimators. To achieve this level of accuracy, the MC-HDG method requires a random sample set Y_M of size $M = 238\,447$ to compute the MC-HDG estimators $E_M[s_h], V_M[s_h]$.

We next pursue the RB method and show in Fig. 2(d) the average output error $\epsilon_{N,\text{avg}}$ and its error bound $\Delta^s_{N,\text{avg}}$ as a function of N. Here $\epsilon_{N,\text{avg}} = \sum_{\boldsymbol{y} \in Y_{\widehat{M}}} |s(\boldsymbol{y}) - s_N(\boldsymbol{y})|/\widehat{M}$ and $\Delta^s_{N,\text{avg}} = \sum_{\boldsymbol{y} \in Y_{\widehat{M}}} \Delta^s_N(\boldsymbol{y})/\widehat{M}$, where $Y_{\widehat{M}}$ is a test set of $\widehat{M} = 100$ samples. We observe that the RB error bound is about two orders of magnitude larger than the output error. The slow convergence rate of the RB error bound is expected because the problem is noncoercive and has (many) $Q = 17$ parameters. Since the RB error bounds are quite pessimistic, we will not consider the MC-RB method to compute the statistical outputs and their error bounds.

We now turn to the 1-MVR and enforce the tolerance $\Delta^E_{M_0,M_1} = \epsilon_{\text{tol}} = 10^{-3}$ and the confidence level of 0.95. We depict in Fig. 3(a) the computational speedup π_1 relative to the MC-HDG method, the original cost function C_1, and the equivalent cost function $\widehat{C}_1$ as a function of the RB dimension I_1 — results for each level size are averaged 8 times. The equivalent cost function $\widehat{C}_1$ approximates the original cost function C_1 reasonably well, despite being drastically less expensive to evaluate than the true cost function (and available *a priori*). The equivalent cost is minimized at $N_1 = 72$, requiring $(M_0, M_1) = (96, 278\,684)$ for a speedup $\pi_1 = 585$, whereas the true cost yields an optimal RB dimension $N_1 = 66$ and sample sizes

(a) Speedup π_1 and cost functions C_1, $\widehat{C}_1$.

(b) Variance of $s_h - s_{I_1}$ and relative time to solve RB output.

(c) Optimal weights.

(d) Expectation (left) and variance (right) estimators with 95% confidence interval.
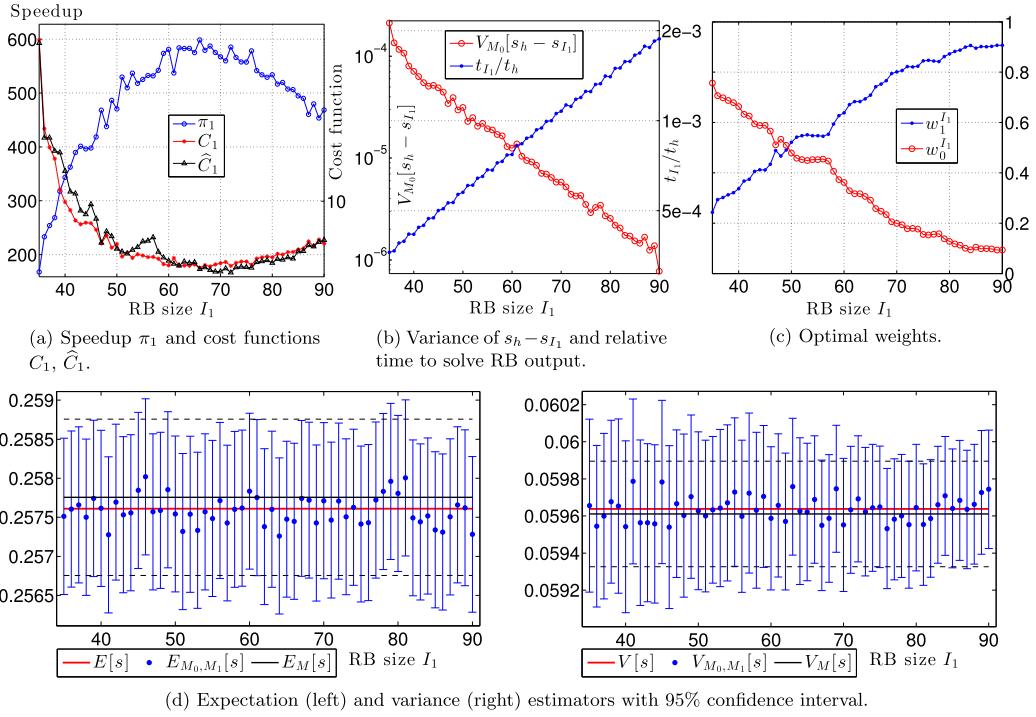
Fig. 3. Result for the 1-MVR method vs RB size I_1.

$(M_0, M_1) = (106, 308\,836)$ that achieve a speedup $\pi_1 = 599$. The model and variance reduction strategy is represented in Figs. 3(b)–3(c). For small reduced basis size more accuracy is demanded for level 1, that is $w_1 < w_0$, since the RB model is very inexpensive to evaluate and the variance of $s_h - s_{I_1}$ is large. Conversely, for increasing I_1 the RB model becomes more costly to compute, whereas the variance of $s_h - s_{I_1}$ rapidly decreases — therefore requiring very few full model evaluations. This change of behavior is detected by the level selection method by setting $w_0 < w_1$, that is requiring higher accuracy for level 0. The model and variance reduction method therefore seeks a balance between these two phenomena to achieve optimal efficiency. We next show in Fig. 3(d) the MC-HDG and 1-MVR expectation and variance as well as their error bars as a function of I_1. The expectation error bound is equal to the prescribed tolerance $\epsilon_{\text{tol}} = 10^{-3}$, while the error bound for the variance $\Delta^V_{M_0, M_1}$ decreases from about 4.5×10^{-4} for small I_1 to about 3.2×10^{-4} for larger reduced basis size.

We proceed analogously for the 2-MVR method computing the speedup with respect to the MC-HDG method for several values of level sizes I_1 and I_2. The computational gain is presented in Fig. 4(a), and the optimum speedup $\pi_2 = 1800$ is reached for $(N_1, N_2) = (122, 36)$, which needs $(M_0, M_1, M_2) = (30, 4102, 286\,152)$. The true cost and the equivalent cost present a very similar behavior, compare Figs. 4(b)–4(c), and the level selection method recovers as low fidelity models the bases $(N_1, N_2) = (139, 36)$, for sample sizes $(M_0, M_1, M_2) = (30, 3483, 281\,824)$ and a speedup $\pi_2 = 1778$. The accuracy of the estimators is shown in Fig. 4(d), and we observe that the true moments lie in all cases inside the confidence intervals (displayed as a surface) computed by the model and variance reduction method. We now analyze the performance of the L-MVR when an arbitrary number of levels are considered. The goal is to use the information on the test set to select not only the level sizes, but also the optimal multilevel model. For simplicity, we consider the same test set with $\widehat{M} = 100$ used for the previous cases, and obtain for each number of levels the optimal level sizes $\boldsymbol{N} = (N_1, \ldots, N_L)$, weights $\boldsymbol{w^N} = \left(w^{\boldsymbol{N}}_0, \ldots, w^{\boldsymbol{N}}_L\right)$ and the equivalent cost $\widehat{C}_L$ in (37). We then perform the L-MVR that provides the actual computational cost C_L in (36). Results corresponding to $L = 1, 2, 3, 4, 5$ are presented in Table 3, for a confidence level of 0.95 and averaged 8 times. The costs are normalized with respect to the minimum costs, attained with $L = 4$ for this problem. The actual speedup π_L in (34) with respect to MC-HDG is also shown.

The proposed selection method effectively predicts the *a priori* performance for each model, since $\widehat{C}_L/\widehat{C}_4$ replicates the behavior of C_L/C_4 quite well. The consideration of an arbitrary number of levels recovers even greater speedups, and the inexpensive *a priori* analysis enables the detection of the optimal model. Furthermore, the weights for each model are shown in Fig. 5(a), exhibiting a nonlinear behavior that truly motivates its selection in an optimal automated manner. The distribution of the weights enforces a larger error on the coarser levels, and requires a smaller error on the finer levels relying on the reduction of variance. The estimators for each number of levels also satisfy the required accuracy, as seen in Figs. 5(b)–5(c).

(a) Computational speedup π_2.   (b) Cost function C_2.   (c) Equivalent cost function $\widehat{C}_2$.

(d) Expectation (left) and variance (right) estimators with 95% confidence interval.
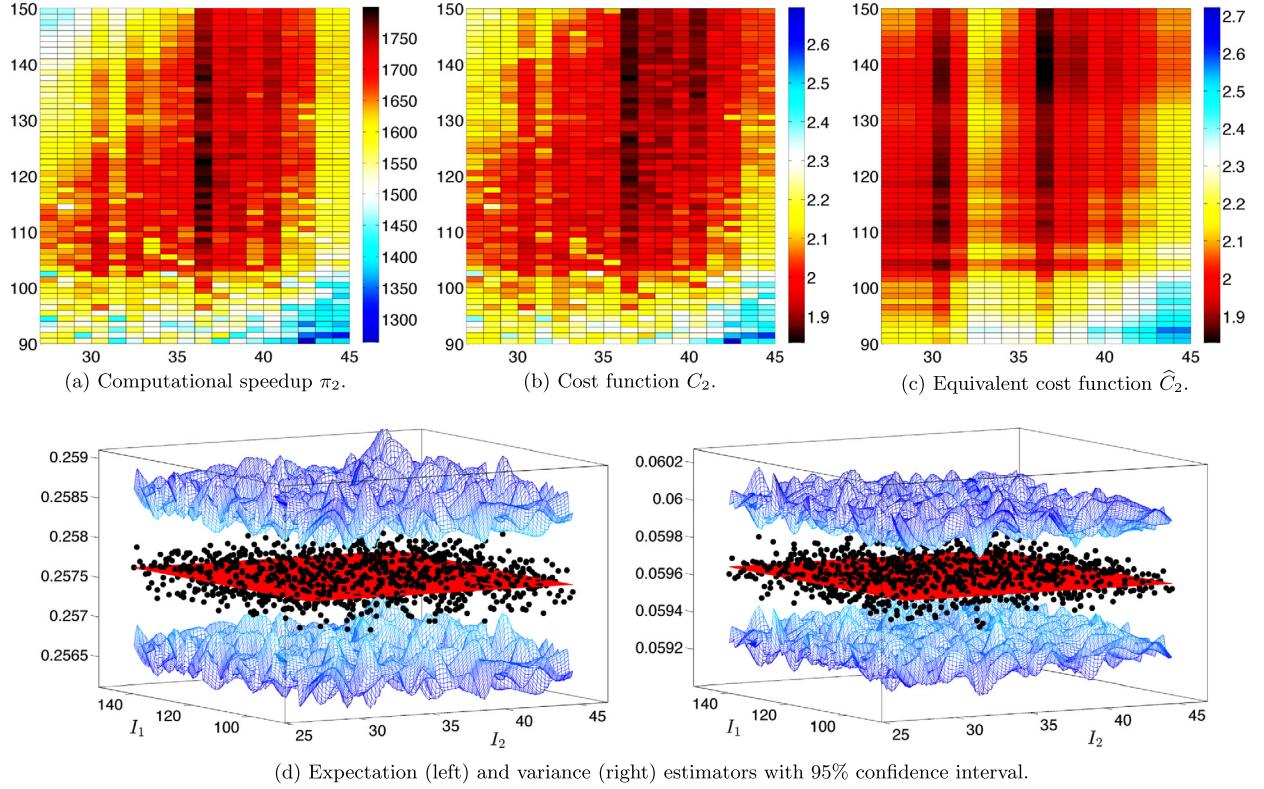
Fig. 4. Results for the 2-MVR method vs RB sizes I_2 and I_1.

Table 3
Predicted optimal level sizes, speedup π_L with respect to MC-HDG cost and relative real and equivalent multilevel cost.

	$L = 1$	$L = 2$	$L = 3$	$L = 4$	$L = 5$
$(N_1, \ldots, N_L)$	72	(139, 36)	(150, 52, 16)	(150, 84, 47, 16)	(150, 118, 84, 47, 16)
π_L	585	1778	2277	2363	2354
C_L/C_4	4.03	1.33	1.04	1	1.00
$\widehat{C}_L/\widehat{C}_4$	4.37	1.38	1.08	1	1.00

(a) Distribution of weights.   (b) Expectation estimator with 95% confidence interval.   (c) Variance estimator with 95% confidence interval.
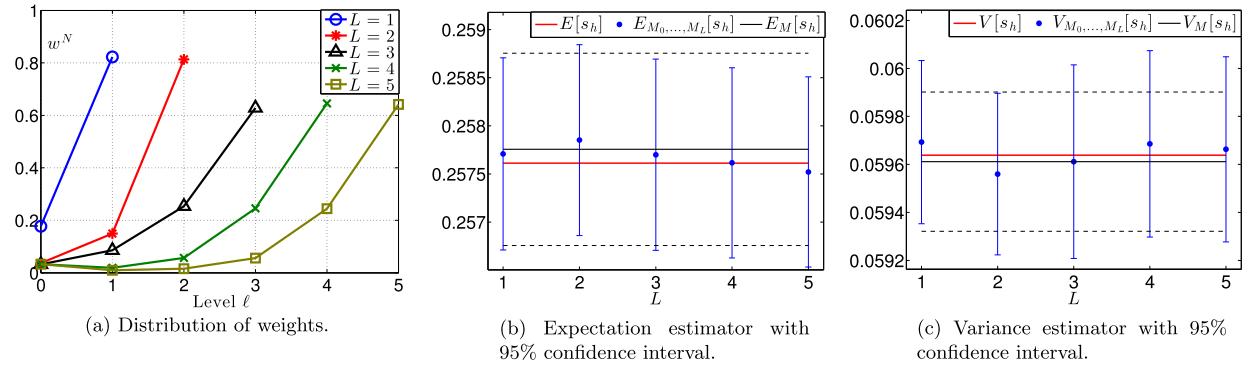
Fig. 5. Results for the L-MVR for arbitrary number of levels.

6. Conclusions

We have presented a model and variance reduction method for computing statistical outputs of stochastic elliptic PDEs. We first combined the reduced basis method with the hybridizable discontinuous Galerkin method by introducing a new HDG weak formulation that retains affine parametric dependence, hence providing rapid and accurate evaluation of the functional output of parametrized PDEs. We next incorporated them into the multilevel control variate framework to exploit

the statistical correlation between the RB approximation and the high-fidelity HDG discretization to accelerate the convergence rate of the Monte Carlo simulations by several orders of magnitude. We then introduced *a posteriori* error bounds for the estimates of the statistical outputs. Finally, we devised an algorithm to optimally select the RB dimensions and the number of levels L. We presented numerical results for both coercive and noncoercive elliptic problems. The results showed that the present method provides a significant speedup compared to both the MC-HDG method and the MC-RB method.

We conclude the paper by pointing out several possible extensions and directions for further research. Firstly, it would be interesting to address the computation of higher order moments, with the additional difficulties of determining the bias and the limiting distributions. Secondly, we would like to extend the proposed approach to nonlinear stochastic problems, which will broaden the application domain of our method. In this aspect, the main challenge remains the development of the RB method for the HDG discretization of nonlinear parametrized PDEs. We would also like to tackle stochastic optimization problems with stochastic PDE constraints, for which the rapid and reliable evaluation of statistical outputs and their derivatives are crucial to finding an optimal solution of any stochastic optimization problem. We would like to develop new methods that allow us to compute not only the statistical outputs but also their derivatives with respect to the decision variables.

Acknowledgements

Appendix A. Bias of estimators

To simplify the notation, we will use the following auxiliary variables

$$\zeta_h := \left(s_h - E_{M_0,\ldots,M_L}[s_h]\right)^2, \qquad \zeta_{N_\ell} := \left(s_{N_\ell} - E_{M_0,\ldots,M_L}[s_h]\right)^2, \ \ell = 1, \ldots, L,$$

$$\widehat{\zeta}_h := (s_h - E[s_h])^2, \qquad\qquad \widehat{\zeta}_{N_\ell} := \left(s_{N_\ell} - E[s_h]\right)^2, \ \ell = 1, \ldots, L,$$

$$\bar{s}_h = E_{M_0,\ldots,M_L}[s_h],$$

and the auxiliary (unbiased) variance

$$\widehat{V}_{M_0,\ldots,M_L}[s_h] = E_{M_0}[\widehat{\zeta}_h - \widehat{\zeta}_{N_1}] + \sum_{\ell=1}^{L-1} E_{M_\ell}[\widehat{\zeta}_{N_\ell} - \widehat{\zeta}_{N_{\ell+1}}] + E_{M_L}[\widehat{\zeta}_{N_L}]$$

which allows us to express the L-MVR variance estimate as

$$V_{M_0,\ldots,M_L}[s_h] = \widehat{V}_{M_0,\ldots,M_L}[s_h] - \left(E[s_h] - E_{M_0,\ldots,M_L}[s_h]\right)^2 \tag{A.1}$$

We first show the latter expression. We add and subtract $E[s_h]$ from every term within the expectations

$$V_{M_0,\ldots,M_L}[s_h] = E_{M_0}\left[\left(s_h - E[s_h] - \bar{s}_h + E[s_h]\right)^2 - \left(s_{N_1} - E[s_h] - \bar{s}_h + E[s_h]\right)^2\right]$$

$$+ \sum_{\ell=1}^{L-1} E_{M_\ell}\left[\left(s_{N_\ell} - E[s_h] - \bar{s}_h + E[s_h]\right)^2 - \left(s_{N_{\ell+1}} - E[s_h] - \bar{s}_h + E[s_h]\right)^2\right]$$

$$+ E_{M_L}\left[\left(s_{N_L} - E[s_h] - \bar{s}_h + E[s_h]\right)^2\right],$$

and now expanding the squares we arrive at

$$V_{M_0,\ldots,M_L}[s_h] = E_{M_0}\left[\widehat{\zeta}_h - \widehat{\zeta}_{N_1} - 2\left(\bar{s}_h - E[s_h]\right)\left(s_h - s_{N_1}\right)\right]$$

$$+ \sum_{\ell=1}^{L-1} E_{M_\ell}\left[\widehat{\zeta}_{N_\ell} - \widehat{\zeta}_{N_{\ell+1}} - 2\left(\bar{s}_h - E[s_h]\right)\left(s_{N_\ell} - s_{N_{\ell+1}}\right)\right]$$

$$+ E_{M_L}\left[\widehat{\zeta}_{N_L} - 2\left(\bar{s}_h - E[s_h]\right)\left(s_{N_L} - E[s_h]\right) + \left(\bar{s}_h - E[s_h]\right)^2\right].$$

Applying linearity of the MC expectation operator and grouping terms we arrive at (A.1). The bias of the L-MVR variance estimate is defined as

$$E\left[V_{M_0,\ldots,M_L}[s_h] - V[s_h]\right] = E\left[\widehat{V}_{M_0,\ldots,M_L}[s_h] - V[s_h]\right] - E\left[\left(E[s_h] - E_{M_0,\ldots,M_L}[s_h]\right)^2\right]$$

$$= E[s_h]^2 - E\left[E_{M_0,\ldots,M_L}^2[s_h]\right], \tag{A.2}$$

since $E_{M_0,\ldots,M_L}[s_h]$ and $\widehat{V}_{M_0,\ldots,M_L}[s_h]$ are unbiased. If we rename the RB output differences as

$$z_0 := s_h - s_{N_1}, \quad z_L := s_{N_L}, \quad z_\ell := s_{N_\ell} - s_{N_{\ell+1}}, \; \ell = 1, \ldots, L-1$$

the expression for expectation of the square of the L-MVR expectation estimate reads

$$E\left[E^2_{M_0,\ldots,M_L}[s_h]\right] = \sum_{\ell=0}^{L} \frac{1}{M_\ell^2} E\left[\left(\sum_{m=1}^{M_\ell} z_\ell(\boldsymbol{y}_m)\right)^2\right] + 2 \sum_{\substack{\ell<\ell'\\\ell=0}}^{L} E\left[\sum_{m=1}^{M_\ell} z_\ell(\boldsymbol{y}_m) \sum_{m'=1}^{M_{\ell'}} z_{\ell'}(\boldsymbol{y}_{m'})\right]$$

thanks to the linearity of the expectation operator. The latter expression can be further reduced with

$$E\left[\left(\sum_{m=1}^{M_\ell} z_\ell(\boldsymbol{y}_m)\right)^2\right] = M_\ell E[z_\ell^2] + (M_\ell^2 - M_\ell)E[z_\ell]^2,$$

$$E\left[\sum_{m=1}^{M_\ell} z_\ell(\boldsymbol{y}_m) \sum_{m'=1}^{M_{\ell'}} z_{\ell'}(\boldsymbol{y}_{m'})\right] = E[z_\ell]E[z_{\ell'}],$$

that hold because we consider independent samples within each level and independent samples among levels. We then have

$$E\left[E^2_{M_0,\ldots,M_L}[s_h]\right] = \sum_{\ell=0}^{L} \left(\frac{E[z_\ell^2] - E[z_\ell]^2}{M_\ell} + E[z_\ell]^2\right) + 2 \sum_{\substack{\ell<\ell'\\\ell=0}}^{L} E[z_\ell]E[z_{\ell'}],$$

and by induction on the number of levels, the cross-products can be reduced to

$$E\left[E^2_{M_0,\ldots,M_L}[s_h]\right] = E[s_h]^2 + \sum_{\ell=0}^{L} \frac{E[z_\ell^2] - E[z_\ell]^2}{M_\ell} = E[s_h]^2 + \sum_{\ell=0}^{L} \frac{V[z_\ell]}{M_\ell}. \tag{A.3}$$

Hence, if we combine equations (A.2)–(A.3) we obtain the bias of the L-MVR variance estimate

$$E\left[V_{M_0,\ldots,M_L}[s_h] - V[s_h]\right] = -\sum_{\ell=0}^{L} \frac{V[z_\ell]}{M_\ell}.$$

References

[1] I. Babuška, 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.
[2] I. Babuška, R. Tempone, G.E. Zouraris, Galerkin finite element approximations of stochastic elliptic partial differential equations, SIAM J. Numer. Anal. 42 (2) (2004) 800–825.
[3] I. Babuška, R. Tempone, G.E. Zouraris, Solving elliptic boundary value problems with uncertain coefficients by the finite element method: the stochastic formulation, Comput. Methods Appl. Mech. Eng. 194 (12) (2005) 1251–1294.
[4] M. Barrault, Y. Maday, N.C. Nguyen, A.T. Patera, An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations, C. R. Math. 339 (9) (2004) 667–672.
[5] 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.
[6] S. Boyaval, A fast Monte-Carlo method with a reduced basis of control variates applied to uncertainty propagation and Bayesian estimation, Comput. Methods Appl. Mech. Eng. (2012).
[7] S. Boyaval, C. Le Bris, T. Lelièvre, Y. Maday, N.C. Nguyen, A.T. Patera, Reduced basis techniques for stochastic problems, Arch. Comput. Methods Eng. 17 (4) (2010) 435–454.
[8] R. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta Numer. (1998) 1–49.
[9] P. Chen, A. Quarteroni, G. Rozza, Comparison between reduced basis and stochastic collocation methods for elliptic problems, J. Sci. Comput. 59 (1) (2014) 187–216.
[10] Q.-Y. Chen, D. Gottlieb, J. Hesthaven, Uncertainty analysis for the steady-state flows in a dual throat nozzle, J. Comput. Phys. 204 (2005) 387–398.
[11] P.G. Ciarlet, The Finite Element Method for Elliptic Problems, Elsevier, 1978.
[12] K.A. Cliffe, M.B. 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.
[13] B. Cockburn, B. Dong, J. Guzmán, A superconvergent LDG-hybridizable Galerkin method for second-order elliptic problems, Math. Comput. 77 (264) (2008) 1887–1916.
[14] B. Cockburn, J. Gopalakrishnan, R. Lazarov, Unified hybridization of discontinuous Galerkin, mixed, and continuous Galerkin methods for second order elliptic problems, SIAM J. Numer. Anal. 47 (2) (2009) 1319–1365.
[15] B. Cockburn, J. Gopalakrishnan, N.C. Nguyen, J. Peraire, F. Sayas, Analysis of HDG methods for Stokes flow, Math. Comput. 80 (274) (2011) 723–760.
[16] B. Cockburn, J. Gopalakrishnan, F. Sayas, A projection-based error analysis of HDG methods, Math. Comput. 79 (271) (2010) 1351–1367.

[17] B. Cockburn, J. Guzmán, H. Wang, Superconvergent discontinuous Galerkin methods for second-order elliptic problems, Math. Comput. 78 (265) (2009) 1–24.
[18] M.K. Deb, I.M. Babuška, J.T. Oden, Solution of stochastic partial differential equations using Galerkin finite element techniques, Comput. Methods Appl. Mech. Eng. 190 (48) (2001) 6359–6372.
[19] W. Feller, An Introduction to Probability Theory and its Applications, vol. 1, John Wiley & Sons, 1968.
[20] G. Fishman, Monte Carlo: Concepts, Algorithms, and Applications, Springer-Verlag, New York, 1996.
[21] B. Ganapathysubramanian, N. Zabaras, Sparse grid collocation schemes for stochastic natural convection problems, J. Comput. Phys. 225 (2007) 652–685.
[22] T. Gerstner, M. Griebel, Dimension adaptive tensor product quadrature, Computing 71 (2003) 65–87.
[23] R. Ghanem, Probabilistic characterization of transport in heterogeneous media, Comput. Methods Appl. Mech. Eng. 158 (3) (1998) 199–220.
[24] R. Ghanem, Higher-order sensitivity of heat conduction problems to random data using the spectral stochastic finite element method, Am. J. Heat Transf. 121 (1999) 290–298.
[25] R. Ghanem, A. Sarkar, Mid-frequency structural dynamics with parameter uncertainty, Comput. Methods Appl. Mech. Eng. 191 (47) (2002) 5499–5513.
[26] R. Ghanem, P.D. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer-Verlag, 1991.
[27] M.B. Giles, Multilevel Monte Carlo path simulation, Oper. Res. 56 (3) (2008) 607–617.
[28] M. Grepl, Reduced-basis approximations for time-dependent partial differential equations: application to optimal control, Ph.D. thesis, Massachusetts Institute of Technology, 2005.
[29] M.A. Grepl, A.T. Patera, A posteriori error bounds for reduced-basis approximations of parametrized parabolic partial differential equations, ESAIM Math. Model. Numer. Anal. 39 (1) (2005) 157–181.
[30] B. Haasdonk, K. Urban, B. Wieland, Reduced basis methods for parameterized partial differential equations with stochastic influences using the Karhunen–Loeve expansion, SIAM/ASA J. Uncertain. Quant. 1 (1) (2013) 79–105.
[31] J.M. Hammersley, D.C. Handscomb, G. Weiss, Monte Carlo methods, Phys. Today 18 (1965) 55.
[32] W.A. Klimke, Uncertainty modeling using fuzzy arithmetic and sparse grids, Citeseer, 2006.
[33] W. Loh, On latin hypercube sampling, Ann. Stat. 24 (1996) 2058–2080.
[34] L. Machiels, Y. Maday, I.B. Oliveira, A.T. Patera, D.V. Rovas, Output bounds for reduced-basis approximations of symmetric positive definite eigenvalue problems, C. R. Acad. Sci., Ser. 1 Math. 331 (2) (2000) 153–158.
[35] L. Mathelin, M.Y. Hussaini, A stochastic collocation algorithm for uncertainty analysis, Tech. rep., NASA/CR-2003-0212153, NASA Langley Research Center, 2003.
[36] L.W.T. Ng, D.B.P. Huynh, K. Willcox, Multifidelity uncertainty propagation for optimization under uncertainty, in: 12th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and 14th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, 2012.
[37] N.C. Nguyen, A.T. Patera, J. Peraire, A 'best points' interpolation method for efficient approximation of parametrized functions, Int. J. Numer. Methods Eng. 73 (4) (2008) 521–543.
[38] N.C. Nguyen, J. Peraire, Hybridizable discontinuous Galerkin methods for partial differential equations in continuum mechanics, J. Comput. Phys. 231 (18) (Jul. 2012) 5955–5988.
[39] N.C. Nguyen, J. Peraire, B. Cockburn, An implicit high-order hybridizable discontinuous Galerkin method for linear convection–diffusion equations, J. Comput. Phys. 228 (9) (2009) 3232–3254.
[40] N.C. Nguyen, J. Peraire, B. Cockburn, An implicit high-order hybridizable discontinuous Galerkin method for nonlinear convection–diffusion equations, J. Comput. Phys. 228 (23) (2009) 8841–8855.
[41] N.C. Nguyen, J. Peraire, B. Cockburn, A hybridizable discontinuous Galerkin method for Stokes flow, Comput. Methods Appl. Mech. Eng. 199 (9–12) (2010) 582–597.
[42] N.C. Nguyen, J. Peraire, B. Cockburn, A hybridizable discontinuous Galerkin method for the incompressible Navier–Stokes equations, in: Proceedings of the 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, Florida, 2010, AIAA Paper 2010-362.
[43] N.C. Nguyen, J. Peraire, B. Cockburn, A hybridizable discontinuous Galerkin method for the incompressible Navier–Stokes equations, in: Proceedings of the 48th AIAA Aerospace Sciences Meeting and Exhibit, Orlando, Florida, 2010, AIAA Paper 2010-362.
[44] N.C. Nguyen, J. Peraire, B. Cockburn, An implicit high-order hybridizable discontinuous Galerkin method for the incompressible Navier–Stokes equations, J. Comput. Phys. 230 (4) (2011) 1147–1170.
[45] N.C. Nguyen, J. Peraire, B. Cockburn, High-order implicit hybridizable discontinuous Galerkin methods for acoustics and elastodynamics, J. Comput. Phys. 230 (10) (2011) 3695–3718.
[46] N.C. Nguyen, J. Peraire, B. Cockburn, Hybridizable discontinuous Galerkin methods for the time-harmonic Maxwell's equations, J. Comput. Phys. 230 (19) (2011) 7151–7175.
[47] N.C. Nguyen, K. Veroy, A.T. Patera, Certified real-time solution of parametrized partial differential equations, in: Handbook of Materials Modeling, Springer, 2005, pp. 1529–1564.
[48] H. Niederreiter, Random Number Generation and Quasi-Monte Carlo Methods, SIAM, Philadelphia, 1992.
[49] F. Nobile, R. Tempone, Analysis and implementation issues for the numerical approximation of parabolic equations with random coefficients, Int. J. Numer. Methods Eng. 80 (6–7) (2009) 979–1006.
[50] F. Nobile, R. Tempone, C. Webster, An anisotropic sparse grid collocation method for elliptic partial differential equations with random input data, SIAM J. Numer. Anal. 46 (5) (2008) 2411–2442.
[51] F. Nobile, R. Tempone, C. Webster, A sparse grid collocation method for partial differential equations with random input data, SIAM J. Numer. Anal. 46 (5) (2008) 2309–2345.
[52] A.K. Noor, J.M. Peters, Reduced basis technique for nonlinear analysis of structures, AIAA J. 18 (4) (1980) 455–462.
[53] N.A. Pierce, M.B. Giles, Adjoint recovery of superconvergent functionals from PDE approximations, SIAM Rev. 42 (2) (2000) 247–264.
[54] C. Prud'homme, Y. Maday, A.T. Patera, G. Turinici, D.V. Rovas, K. Veroy, L. Machiels, Reliable real-time solution of parametrized partial differential equations: reduced-basis output bound methods, J. Fluids Eng. 124 (1) (2002) 70–80.
[55] G. Rozza, D.B.P. Huynh, A.T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, Arch. Comput. Methods Eng. 15 (3) (2008) 229–275.
[56] S. Smolyak, Quadrature and interpolation formulas for tensor product of certain classes of functions, Sov. Math. Dokl. 4 (1963) 240–243.
[57] M. Stein, Large sample properties of simulations using latin hypercube sampling, Technometrics 29 (1987) 143–151.
[58] A.L. Teckentrup, R. Scheichl, M.B. Giles, E. Ullmann, Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients, Numer. Math. 125 (3) (2013) 569–600.
[59] M.P. Ueckermann, P.F.J. Lermusiaux, High-order schemes for 2d unsteady biogeochemical ocean models, Ocean Dyn. 60 (6) (2010) 1415–1445.
[60] K. Veroy, A.T. Patera, Certified real-time solution of the parametrized steady incompressible Navier–Stokes equations: rigorous reduced-basis a posteriori error bounds, Int. J. Numer. Methods Fluids 47 (8–9) (2005) 773–788.
[61] K. Veroy, C. Prud'homme, D.V. Rovas, A.T. Patera, A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in: Proceedings of the 16th AIAA Computational Fluid Dynamics Conference, vol. 3847, 2003.
[62] K. Veroy, D.V. Rovas, A.T. Patera, A posteriori error estimation for reduced-basis approximation of parametrized elliptic coercive partial differential equations: "convex inverse" bound conditioners, ESAIM Control Optim. Calc. Var. 8 (2002) 1007–1028.

[63] N. Wiener, The homogeneous chaos, Am. J. Math. 60 (1938) 897–936.
[64] D. Xiu, J.S. Hesthaven, High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput. 27 (2005) 1118–1139.
[65] D. Xiu, G.E. Karniadakis, Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos, Comput. Methods Appl. Mech. Eng. 191 (43) (2002) 4927–4948.
[66] D. Xiu, G.E. Karniadakis, Modeling uncertainty in flow simulations via generalized polynomial chaos, J. Comput. Phys. 187 (1) (2003) 137–167.
[67] D. Xiu, G.E. Karniadakis, A new stochastic approach to transient heat conduction modeling with uncertainty, Int. J. Heat Mass Transf. 46 (24) (2003) 4681–4693.