A Low-rank Control Variate for Multilevel Monte Carlo Simulation of High-dimensional Uncertain Systems

Hillary R. Fairbanks^a, Alireza Doostan^{b,*}, Christian Ketelsen^c, Gianluca Iaccarino^d

^aApplied Mathematics and Statistics Department, University of Colorado, Boulder, CO 80309, USA

^bAerospace Engineering Sciences Department, University of Colorado, Boulder, CO 80309, USA

^cComputer Science Department, University of Colorado, Boulder, CO 80309, USA

^dMechanical Engineering Department, Stanford University, Stanford, CA 94305, USA

Abstract

Multilevel Monte Carlo (MLMC) is a recently proposed variation of Monte Carlo (MC) simulation that achieves variance reduction by simulating the governing equations on a series of spatial (or temporal) grids with increasing resolution. Instead of directly employing the fine grid solutions, MLMC estimates the expectation of the quantity of interest from the coarsest grid solutions as well as differences between each two consecutive grid solutions. When the differences corresponding to finer grids become smaller, hence less variable, fewer MC realizations of finer grid solutions are needed to compute the difference expectations, thus leading to a reduction in the overall work. This paper presents an extension of MLMC, referred to as *multilevel control variate (MLCV)*, where a low-rank approximation to the solution on each grid, obtained primarily based on coarser grid solutions, is used as a control variate for estimating the expectations involved in MLMC. Cost estimates as well as numerical examples are presented to demonstrate the advantage of this new MLCV approach over the standard MLMC when the solution of interest admits a low-rank approximation and the cost of simulating finer grids grows fast.

Keywords: Uncertainty Quantification, Stochastic PDEs, Multilevel Monte Carlo, Control Variate, Low-rank Approximation, Multifidelity, Interpolative Decomposition

1. Introduction

The use of uncertainty quantification as a tool to assess the prediction accuracy of simulation models of physical systems has been increasing at a rapid rate over the last decade. By accounting for the uncertainties of input data in models, such as initial conditions, boundary conditions, or other physical parameters, the objective is to establish the predictive capabilities of simulations by quantifying the uncertainty in the quantities of interest (QoI's). To this end and within the probabilistic framework, several methods, e.g., polynomial chaos expansions [1, 2, 3] and stochastic collocation [4, 5], have been developed and proven successful in various applications. However, it is known that the computational cost of these methods grows rapidly as a function of the number of random variables describing model uncertainties, a phenomenon referred to as *curse of dimensionality*.

An alternative class of techniques rely on the Monte Carlo (MC) simulation or its variants, where the statistics of the QoI are estimated using an ensemble of (random) realizations of the QoI. The cost of such estimations, while may be prohibitive, is *formally* independent of the number of input variables. In details, let $\xi = (\xi_1, \dots, \xi_d)$ denote the d -vector of random variables, with joint probability density function $\rho_\xi(\xi)$, representing the uncertainty in the inputs. Let $Q = Q(\xi)$ denote a scalar-valued QoI depending on ξ and Q_M its approximation obtained via simulation. The subscript M denotes the number of deterministic degrees of

*Corresponding Author: Alireza Doostan

Email address: alireza.doostan@colorado.edu (Alireza Doostan)

freedom, e.g., number of grid points in a finite element model, controlling the accuracy of Q_M relative to Q . The goal is to approximate the statistics of Q , e.g., the mean of Q , $\mathbb{E}[Q]$, using the realizations of Q_M . Given a set of N samples of inputs, each denoted by $\boldsymbol{\xi}^{(i)}$ and drawn according to $\rho_{\boldsymbol{\xi}}(\boldsymbol{\xi})$, and the corresponding realizations of Q_M , given by $Q_M^{(i)} = Q_M(\boldsymbol{\xi}^{(i)})$, the MC approximation of $\mathbb{E}[Q]$ is

$$\mathbb{E}[Q] \approx \mathbb{E}[Q_M] \approx \hat{Q}_{M,N}^{MC} = \frac{1}{N} \sum_{i=1}^N Q_M^{(i)}. \quad (1)$$

Following the notation in [6], $\hat{Q}_{M,N}^{MC}$ in (1) is the MC estimator of $\mathbb{E}[Q_M]$ using N samples of Q_M with the Mean Square Error (MSE)

$$\text{MSE}(\hat{Q}_{M,N}^{MC}, \mathbb{E}[Q]) = \frac{1}{N} \mathbb{V}[Q_M] + (\mathbb{E}[Q_M - Q])^2, \quad (2)$$

where $\mathbb{V}$ is the variance operator and $\text{MSE}(\hat{Q}_{M,N}^{MC}, \mathbb{E}[Q])$ denotes the MSE of $\hat{Q}_{M,N}^{MC}$ with respect to $\mathbb{E}[Q]$. We note that, in this paper, the hat operator indicates the MC estimator of the corresponding expectation. In (2), the MSE is decomposed into the *sampling error* $\frac{1}{N} \mathbb{V}[Q_M]$, controlled by the variance of Q_M and the number of samples, and the *discretization error* $(\mathbb{E}[Q_M - Q])^2$, which measures how closely the model simulates the true solution. As can be seen from (2), the sampling error decays slowly as a function of N , but with a rate that is independent of the dimension d , implying that the standard MC simulation does not formally suffer from the curse of dimensionality.

Aside from the necessary refinement in the model to reduce the discretization error, there are only two options to improve the MSE of a MC estimate: increasing the sample size N or using a variance reduction technique. Due to the cost incurred by the first option, it is more practical to consider the use of a variance reduction technique, such as importance sampling or control variates (CV) [7]. In particular, the CV approach considers a second quantity Z , such that it is correlated with Q_M , is cheaper than Q_M to evaluate, and whose expectation is either known or can be approximated with relatively small cost. Then a new variable,

$$W = Q_M - \theta(Z - \mathbb{E}[Z]),$$

is constructed that has the same mean as Q_M , i.e., $\mathbb{E}[W] = \mathbb{E}[Q_M]$, thus suggesting the use of MC estimate of $\mathbb{E}[W]$ as a proxy for $\mathbb{E}[Q_M]$. In doing so, the gain is that, depending on the choice of θ , the MC estimator of $\mathbb{E}[W]$ features a reduced MSE (or variance). Stated differently, a smaller number of W realizations, hence Q_M realizations, are needed for a comparable MSE when CV is applied. For scenarios when Z is poorly correlated with Q_M , a notable MSE reduction is not observed. If, in addition, the cost of estimating $\mathbb{E}[Z]$ is large, it is likely that this CV will not result in a cost improvement over standard MC.

Multilevel Monte Carlo (MLMC), proposed in [8, 9], is a generalization of CV, which constructs a sequence of control variates Z based on approximations of Q on a set of models that are cheaper to simulate than the one for Q_M , hence the term *multilevel*. A common example of a cheaper model is to approximate Q on coarser grids with number of degrees of freedom smaller than M . While the notion of levels can go beyond a grid-based construction, we limit the scope of this study to such an approach. For the interest of a simpler introduction, we delay the full presentation of MLMC to Section 2, and instead focus on the two-level formulation next.

Taking $\theta = 1$ and $Z = Q_m$, with $m < M$, to be the QoI approximated from a coarser grid than that of Q_M , the two-level MLMC variable is given by $W = Q_M - (Q_m - \mathbb{E}[Q_m]) = \mathbb{E}[Q_m] + (Q_M - Q_m)$, with expected value

$$\begin{aligned} \mathbb{E}[W] &= \mathbb{E}[Q_m] + \mathbb{E}[Q_M - Q_m] \\ &= \mathbb{E}[Q_M]. \end{aligned} \quad (3)$$

To approximate $\mathbb{E}[W]$, or equivalently $\mathbb{E}[Q_M]$, MC is applied independently to the two expectations in the right-hand-side of the first equation in (3),

$$\hat{W} = \frac{1}{N_m} \sum_{i=1}^{N_m} Q_m^{(i)} + \frac{1}{N_M} \sum_{i=1}^{N_M} (Q_M^{(i)} - Q_m^{(i)}). \quad (4)$$

As compared to the standard MC estimator of $\mathbb{E}[Q_M]$ given in (1), the estimation of $\mathbb{E}[Q_m]$ in (4) also involves drawing samples of Q_m , which are less expensive. More importantly, when Q_m is close to Q_M , estimating $\mathbb{E}[Q_M - Q_m]$ requires fewer samples of Q_M , as $(Q_M - Q_m)$ features a smaller variance. In practice, depending on the cost of simulating the two models as well as the variances of Q_m and $(Q_M - Q_m)$, the numbers of samples of Q_m , N_m , and Q_M , N_M , are selected such that the overall estimation cost, for a given accuracy, is minimal. MLMC expands upon this concept by including multiple levels, as delineated in Section 2.

Very recently the combination of MLMC and CV, referred to as *multilevel control variates* (MLCV), has been used with the aim of further reducing the variance of $(Q_M - Q_m)$ and thereby improving the computational cost [10]. As an approach to MLCV, in [10], the authors consider the solution to a linear diffusion problem, where the diffusion coefficient is modeled by a rough random field represented by a large number of independent random variables. An auxiliary diffusion problem with smoothed (low-dimensional) coefficient is developed, whose solution – computed via stochastic collocation – is employed as the control variate for the MLMC simulation of the original problem. This particular construction of the control variate relies on the accuracy of the smoothed problem in predicting the solution to the original one. Other related work [11, 12] establishes multilevel control variates within a MC simulation framework. For instance, in [12], a model reduction approach based on the reduced basis method, see, e.g., [13], is considered for the case of linear, stochastic elliptic PDEs. In particular, in an off-line stage, a set of basis functions, i.e., reduced basis, is identified from the realizations of the solution on a target mesh. The PDE solution is then approximated in an increasing number of reduced basis functions, via a discontinuous Galerkin formulation, which are then employed as control variates. The levels here are defined based on the size of the reduced basis used, as opposed to the number of grid points in the spatial discretization of the PDE.

1.1. Contributions of This Work

In this work we propose a different construction of a multilevel control variate that may be applicable to more general classes of problems. That is, our method may be applied to general nonlinear equations and is not limited to diffusion-type equations. In addition, we are able to use a standard PDE solver, rather than a different solver based on reduced basis, to generate the control variate and solution samples. The main idea behind this approach is to create a control variate Z for $(Q_M - Q_m)$ that is obtained based on a low-rank approximation of Q_M , using a *small* set of selected samples of fine grid solution identified from the realizations of the coarse grid solution. This low-rank approximation of the fine grid solution, based on ideas in [14, 15, 16] and following the algorithm described in [17], consists of three main steps. In the first step, a reduced basis, say of size $r \ll M$, is identified for the coarse grid solution, together with an interpolation rule that gives an arbitrary realization of the coarse grid solution in that basis. For this purpose, we borrow ideas from matrix interpolative decomposition (MID) as presented in [18]. The second step entails the identification of a reduced basis for the fine grid solution. For this, the fine counterpart to the coarse grid reduced basis is generated, which requires r realizations of the fine grid solution. In the third step, the same interpolation rule, as for the coarse grid solution, is applied to generate the low-rank approximation of its fine counterpart (and subsequently Q_M). Based on this construction, we are able to formulate our low-rank control variate Z for $(Q_M - Q_m)$, and then the control variate equation

$$W = (Q_M - Q_m) - \theta(Z - \mathbb{E}[Z])$$

is used to estimate $\mathbb{E}[Q_M - Q_m]$, which requires setting θ and estimating $\mathbb{E}[Z]$. For these, we consider work from [19], where θ is chosen such that the MSE of $\hat{W}$, with respect to $\mathbb{E}[Q_M - Q_m]$, is minimized. In addition, $\mathbb{E}[Z]$ is estimated primarily from an independent MC simulation of the coarse grid solution.

An important feature of this MLCV approach is worthwhile highlighting. Depending on the rank of the fine model solution and the cost of simulating Q_m , relative to Q_M , Z may be sampled with a cost that is considerably smaller than that of Q_M . This, together with the MSE reduction achieved by $\hat{W}$, will lead to a smaller number of samples of Q_M required for a similar accuracy as in standard MLMC.

This paper is organized into four sections. A background on MLMC is provided in Section 2, including the setup and improvements it provides over MC, as well as a discussion of MLMC theory. Section 3 focuses

on the formulation of the MLCV method of this paper, including the construction of the control variate and the MLCV estimator such that the MSE is minimized. This section also includes an algorithm to outline the process, as well as a brief discussion regarding considerations to take into account when taking this MLCV approach. In Section 4 the numerical results of this new method as well as comparison with MLMC are presented using two test problems. The first test considers an eigenvalue problem associated with a linear elasticity problem in an L-shaped domain, followed by the second case, a thermally driven flow problem in a square domain.

2. Background on Multilevel Monte Carlo

From [8, 9], we next describe the use of MLMC as a variance reduction technique for the MC simulation of differential equations with random data. Consider a sequence of spatial discretizations, i.e., grids, for the governing equations with increasing accuracy, where each discretization is indexed by a *level* parameter ℓ , $\ell = 0, \dots, L$. Let M_ℓ denote the number of degrees of freedom in the level ℓ discretization model, such that $M_0 < M_1 < \dots < M_L$. Unlike in MC, MLMC does not directly approximate the expectation of QoI on the finest (target) level L . Instead, it uses the telescoping sum

$$\mathbb{E}[Q] \approx \mathbb{E}[Q_L] = \mathbb{E}[Q_0] + \sum_{\ell=1}^L \mathbb{E}[Q_\ell - Q_{\ell-1}], \quad (5)$$

which is the expected value of the QoI on the coarse level plus the sum of expectations of correction terms. Here, Q_ℓ is the QoI determined from the level ℓ grid with M_ℓ degrees of freedom. Defining the multilevel correction variable

$$Y_\ell = Q_\ell - Q_{\ell-1}, \quad (6)$$

for $\ell = 1, \dots, L$, and $Y_0 = Q_0$, the sum in (5) can be rewritten as

$$\mathbb{E}[Q_L] = \sum_{\ell=0}^L \mathbb{E}[Y_\ell].$$

In MLMC, the expectation of each correction variable Y_ℓ is independently computed with N_ℓ MC samples of Y_ℓ ,

$$\mathbb{E}[Y_\ell] \approx \hat{Y}_\ell = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} Y_\ell^{(i)} = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (Q_\ell^{(i)} - Q_{\ell-1}^{(i)}), \quad (7)$$

where each $Q_\ell^{(i)}$ and $Q_{\ell-1}^{(i)}$ is generated by applying the same realization $\boldsymbol{\xi}^{(i)}$ to the level ℓ and $\ell - 1$ grids, respectively. Building on the notation once more, the MLMC estimator, $\hat{Q}_L^{ML}$, is defined to be

$$\hat{Q}_L^{ML} = \sum_{\ell=0}^L \hat{Y}_\ell.$$

When Q_ℓ converges to Q as a function of ℓ , the variance of Y_ℓ converges to zero. This allows for $\hat{Y}_\ell$ to converge with fewer samples of Q_ℓ for larger, and more expensive, levels ℓ . To determine how to pick N_ℓ such that the same level of accuracy as in MC is reached with a reduced computational cost, the MSE of $\hat{Q}_L^{ML}$ must first be considered.

2.1. MSE of MLMC

The MSE of the MLMC estimator, much like that of the MC estimator in (2), can be decomposed into the sampling and discretization errors, given by

$$\text{MSE}(\hat{Q}_L^{ML}, \mathbb{E}[Q]) = \sum_{\ell=0}^L \frac{1}{N_\ell} \mathbb{V}[Y_\ell] + (\mathbb{E}[Q_L - Q])^2, \quad (8)$$

where the sampling error of MLMC is a sum of MC sampling errors on each level. For a desired MSE bound of magnitude ε^2 , the sampling error (as well as the discretization error) is ideally required to be bounded by $\varepsilon^2/2$; that is,

$$\sum_{\ell=0}^L \frac{1}{N_\ell} \mathbb{V}[Y_\ell] \leq \frac{\varepsilon^2}{2}. \quad (9)$$

Provided an estimate of $\mathbb{V}[Y_\ell]$, (9) serves as a constraint to estimate the number of samples N_ℓ drawn from each level.

2.2. Number of Samples N_ℓ for $\hat{Y}_\ell$

In MLMC, N_ℓ is determined such that the total cost, given by

$$\mathcal{C}(\hat{Q}_L^{ML}) = \sum_{\ell=0}^L N_\ell \mathcal{C}_\ell, \quad (10)$$

is minimized, where $\mathcal{C}_\ell$ is the cost of generating a sample of Y_ℓ and is given by

$$\mathcal{C}_\ell = \mathcal{C}(Y_\ell) = \mathcal{C}(Q_\ell) + \mathcal{C}(Q_{\ell-1}), \quad (11)$$

for $\ell = 1, \dots, L$, and $\mathcal{C}_0 = \mathcal{C}(Q_0)$. In practice, the average CPU time for obtaining a sample of Y_ℓ can be used to determine $\mathcal{C}_\ell$. Alternatively, when the computational complexity of the adopted solvers with respect to M_ℓ is known, a relation of the form

$$\mathcal{C}_\ell \lesssim M_\ell^\gamma, \quad \gamma > 0,$$

may be considered, where $\lesssim$ denotes inequality up to a positive constant. By using the method of Lagrange multipliers to minimize the total cost in (10), with an $\varepsilon^2/2$ sampling error constraint in (9), the ideal number of samples N_ℓ from each level can be computed from

$$N_\ell \geq \frac{2}{\varepsilon^2} \left[\sum_{k=0}^L \sqrt{\mathbb{V}[Y_k] \mathcal{C}_k} \right] \sqrt{\frac{\mathbb{V}[Y_\ell]}{\mathcal{C}_\ell}}, \quad (12)$$

see, e.g., [20]. In order to calculate the values of N_ℓ , a *pilot run* of MLMC must first be implemented; that is, MLMC is completed on a relatively small sample size that may be constant across all levels. By doing so, the expectation, variance, and cost on each level can be estimated to determine the required number of levels and optimal N_ℓ from (12), such that a given MSE tolerance can be achieved. Subsequently, a second run of MLMC is completed, this time using only the calculated sample size N_ℓ for each level. To make the method as efficient as possible, simulations from the pilot run are incorporated into the second run of MLMC.

2.3. MLMC Convergence Guarantee

In order to guarantee an ε^2 bound on the MSE in (8), we turn to the MLMC theorem as presented in [6]. While the number of samples on each level can be determined such that the desired sampling error bound $\varepsilon^2/2$ is attained, it does not imply the same for the discretization error. The MLMC theorem provides the constraints in order to maintain the prescribed MSE bound. Assuming that there exists an integer $s > 1$ such that $M_\ell = sM_{\ell-1}$, for $\ell = 1, \dots, L$, and

1. $|\mathbb{E}[Q_\ell - Q]| \lesssim M_\ell^{-\alpha}$,
2. $\mathbb{V}[Y_\ell] \lesssim M_\ell^{-\beta}$,
3. $\mathcal{C}_\ell \lesssim M_\ell^\gamma$,

for some constants $\gamma, \beta > 0$ and $\alpha \geq \frac{1}{2} \min(\beta, \gamma)$, Theorem 1 in [6] states that there exists a positive integer L depending on ε and a sequence $\{N_\ell\}_{\ell=0}^L$, such that $\text{MSE}(\hat{Q}_L^{ML}, \mathbb{E}[Q]) < \varepsilon^2$. This relation indicates whether or not it is necessary to refine the model further to meet the desired discretization error. In Section 4, estimates of α and β will be provided as an illustration that the theory holds for the numerical experiments of this study.

3. Multilevel Control Variates

It has been shown that MLMC is more cost effective than MC for many stochastic differential equations [8, 9, 6, 21, 22]. Due to the success of this method, we consider continuing along the path of variance – more precisely MSE – reduction by applying control variates to each $\mathbb{E}[Y_\ell]$ estimate with the aim of reducing the number of samples N_ℓ required by MLMC or the required work to achieve a desired MSE ε^2 . In detail, given the multilevel correction variable Y_ℓ , as defined in (6), and a correlated variable Z_ℓ , a general MLCV correction variable is defined as

$$W_\ell = Y_\ell - \theta_\ell(Z_\ell - \mathbb{E}[Z_\ell]), \quad (13)$$

for $\ell = 1, \dots, L$, and $W_0 = Y_0$. If $\mathbb{E}[Z_\ell]$ is known, it follows that $\mathbb{E}[W_\ell] = \mathbb{E}[Y_\ell]$, and θ_ℓ can be set such that the MSE of the estimator of $\mathbb{E}[W_\ell]$, $\hat{W}_\ell$, is optimized. However, it is often the case that $\mathbb{E}[Z_\ell]$ is unknown and thus must be estimated. In the following subsections we will discuss our formulation of the control variate Z_ℓ and the approximation of $\mathbb{E}[Z_\ell]$. The discussion will include the decomposition of the MSE of $\hat{W}_\ell$ with respect to $\mathbb{E}[Y_\ell]$ and the optimal number of samples needed to reach a desired sampling error.

3.1. Z_ℓ Formulation

We consider a formulation of Z_ℓ based on a low-rank approximation of Y_ℓ in (6). Specifically,

$$Z_\ell = Q_\ell^{ID} - Q_{\ell-1}, \quad (14)$$

where $Q_{\ell-1}$ is as defined in (6), and Q_ℓ^{ID} is a low-rank approximation of Q_ℓ based on samples of $Q_{\ell-1}$ and a pre-determined basis for level ℓ solution described next. Let Q_ℓ depend on a solution-dependent, vector-valued quantity $\mathbf{q}_\ell(\boldsymbol{\xi}) \in \mathbb{R}^{m_\ell}$; that is,

$$Q_\ell = Q(\mathbf{q}_\ell(\boldsymbol{\xi})).$$

For instance, the (spatial) mean heat flux along a boundary, Q_ℓ , can be computed based on the values of the heat flux, $\mathbf{q}_\ell$, on the boundary grid nodes.

To construct Q_ℓ^{ID} , we assume the approximation to $\mathbf{q}$ on level ℓ , $\mathbf{q}_\ell$, admits an accurate representation in a reduced basis of (small) cardinality $r \ll m_\ell$. Inspired by [14, 15, 16, 17], such a low-rank approximation of $\mathbf{q}_\ell$ may be achieved by identifying a reduced basis for $\mathbf{q}_{\ell-1}$ along with an approximation rule representing an arbitrary sample of $\mathbf{q}_{\ell-1}$ in that basis. In details, some $N_\ell \geq m_{\ell-1}$ realizations of $\mathbf{q}_{\ell-1}$ corresponding to random samples $\{\boldsymbol{\xi}^{(i)}\}_{i=1}^{N_\ell}$, are generated and organized in an $m_{\ell-1} \times N_\ell$ *coarse grid data matrix*

$$\mathbf{U}_{\ell-1} := \begin{bmatrix} \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(1)}) & \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(2)}) & \cdots & \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(N_\ell)}) \end{bmatrix}.$$

We note that, here, N_ℓ is different from the sample size N_ℓ in MLMC and will be defined more precisely in Section 3.6.1. To find a reduced basis for $\mathbf{q}_{\ell-1}$, we consider a rank r factorization of $\mathbf{U}_{\ell-1}$ using a subset of its columns. While a number of tools are available for this purpose [23, 24, 25, 26], we employ the so-called interpolative decomposition (ID) of $\mathbf{U}_{\ell-1}$ [27, 18, 25],

$$\mathbf{U}_{\ell-1} \approx \mathbf{U}_{\ell-1}^c \mathbf{C}_{\ell-1}, \quad (15)$$

where the $m_{\ell-1} \times r$ *column skeleton matrix*

$$\mathbf{U}_{\ell-1}^c = \begin{bmatrix} \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(i_1)}) & \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(i_2)}) & \cdots & \mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(i_r)}) \end{bmatrix}$$

consists of r columns of $\mathbf{U}_{\ell-1}$ identified via pivoted rank-revealing QR factorization of $\mathbf{U}_{\ell-1}$, and $\mathbf{C}_{\ell-1}$ is an $r \times N_\ell$ *coefficient matrix* as specified in Appendix A. Stated differently, (15) gives a rank r approximation of $\mathbf{q}_{\ell-1}(\boldsymbol{\xi}^{(i)})$ in a *reduced basis* $\mathbf{U}_{\ell-1}^c$ consisting of r realizations of $\mathbf{q}_{\ell-1}$. For $\mathbf{q}_{\ell-1}(\boldsymbol{\xi})$ evaluated at an arbitrary $\boldsymbol{\xi}$, a rank r representation of the form

$$\mathbf{q}_{\ell-1}(\boldsymbol{\xi}) \approx \mathbf{U}_{\ell-1}^c \mathbf{c}_{\ell-1}(\boldsymbol{\xi}) \quad (16)$$

can be generated by computing the r -vector of coefficients $\mathbf{c}_{\ell-1}(\boldsymbol{\xi})$ via least squares approximation. Specifically,

$$\mathbf{c}_{\ell-1}(\boldsymbol{\xi}) = \arg \min_{\hat{\mathbf{c}} \in \mathbb{R}^r} \|\mathbf{U}_{\ell-1}^c \hat{\mathbf{c}} - \mathbf{q}_{\ell-1}(\boldsymbol{\xi})\|_2, \quad (17)$$

which can be computed by via SVD or QR decomposition. Further discussion related to the complexity of solving for the coefficients, $\mathbf{c}_{\ell-1}(\boldsymbol{\xi})$, can be found in Section 3.7.

Remark. It is worthwhile highlighting that, in practice, r is not known a priori and has to be chosen such that the approximation in (15) achieves a desired accuracy. Additionally, to reveal the rank r of $\mathbf{U}_{\ell-1}$ from (15), $m_{\ell-1}$ must satisfy the condition $m_{\ell-1} \geq r$, as $r \leq \min\{m_{\ell-1}, N_{\ell}\}$.

Remark. The rank r may vary from one discretization level to another; however, for the interest of a simpler notation, we suppress the dependence of r on ℓ .

From this process, we identify samples $\{\boldsymbol{\xi}^{(i_k)}\}_{k=1}^r$ using which the corresponding *fine grid* reduced basis and approximation can be determined. Specifically, following [14] and using the coarse grid coefficient vector $\mathbf{c}_{\ell-1}$ in (16), we now define the rank r approximation of $\mathbf{q}_{\ell}$ as

$$\mathbf{q}_{\ell}^{ID}(\boldsymbol{\xi}) = \mathbf{U}_{\ell}^c \mathbf{c}_{\ell-1}(\boldsymbol{\xi}), \quad (18)$$

where the reduced basis

$$\mathbf{U}_{\ell}^c = \begin{bmatrix} \mathbf{q}_{\ell}(\boldsymbol{\xi}^{(i_1)}) & \mathbf{q}_{\ell}(\boldsymbol{\xi}^{(i_2)}) & \cdots & \mathbf{q}_{\ell}(\boldsymbol{\xi}^{(i_r)}) \end{bmatrix}$$

is the $m_{\ell} \times N_{\ell}$ fine grid counterpart of $\mathbf{U}_{\ell-1}^c$. Finally, (14) is fully specified by setting

$$Q_{\ell}^{ID} = Q(\mathbf{q}_{\ell}^{ID}(\boldsymbol{\xi})).$$

For a discussion on the convergence of the coarse-fine approximation $\mathbf{q}_{\ell}^{ID}$ in (18), we refer the interested reader to Sections 3 and 4 of [14], where an identical construction is presented.

3.2. $\mathbb{E}[Z_{\ell}]$ Estimation

The numerical construction of Z_{ℓ} in (14) does not lead to an analytic value for $\mathbb{E}[Z_{\ell}]$, as desired in standard control variate approaches. However, approximations to $\mathbb{E}[Z_{\ell}]$, denoted hereafter by $\bar{Z}_{\ell}$, may be alternatively used as long as they are not expensive to generate [28, 29, 19]. While there are multiple techniques to estimate $\mathbb{E}[Z_{\ell}]$, an independent MC simulation is used here. In detail, let N'_{ℓ} be the number of MC samples collected from level $\ell - 1$ to estimate $\mathbb{E}[Z_{\ell}]$. We define the MC estimator of $\mathbb{E}[Z_{\ell}]$ to be

$$\bar{Z}_{\ell} = \frac{1}{N'_{\ell}} \sum_{i=1}^{N'_{\ell}} Z_{\ell}^{(i)} = \frac{1}{N'_{\ell}} \sum_{i=1}^{N'_{\ell}} \left(Q_{\ell}^{ID(i)} - Q_{\ell-1}^{(i)} \right), \quad (19)$$

where $Q_{\ell}^{ID(i)}$ is generated following the procedure of Section 3.1. An advantage to this MC method is that the estimate $\bar{Z}_{\ell}$ only relies on additional level $\ell - 1$ samples without any need to draw additional level ℓ samples. Thus, given the reduced bases $\mathbf{U}_{\ell-1}^c$ and $\mathbf{U}_{\ell}^c$, the primary cost of finding $\bar{Z}_{\ell}$ is running the level $\ell - 1$ model to obtain N'_{ℓ} coarse grid samples, which is given by

$$\mathcal{C}(\bar{Z}_{\ell}) = N'_{\ell} \mathcal{C}(Q_{\ell}^{ID}) + N'_{\ell} \mathcal{C}(Q_{\ell-1}) \approx N'_{\ell} \mathcal{C}(Q_{\ell-1}).$$

A more detailed quantification of $\mathcal{C}(Q_{\ell}^{ID})$ will be provided in Section 3.7. Next, we focus on the MSE reduction achieved with this method of control variates, and subsequently, the estimation of $\tilde{N}_{\ell}$ and N'_{ℓ} .

3.3. MSE Reduction via W_ℓ

Now that we have defined Z_ℓ and $\bar{Z}_\ell$ we can update (13) to get

$$W_\ell = Y_\ell - \theta_\ell(Z_\ell - \bar{Z}_\ell), \quad (20)$$

for $\ell = 1, \dots, L$, and $W_0 = Y_0$. Given a fixed (independently computed) $\bar{Z}_\ell$, $\hat{W}_\ell$ is a biased estimator of $\mathbb{E}[Y_\ell]$ with the bias $\theta_\ell \delta_\ell$, where $\delta_\ell = \bar{Z}_\ell - \mathbb{E}[Z_\ell]$. However, δ_ℓ may be thought of as a random variable with mean zero and variance $\mathbb{V}[Z_\ell]/N'_\ell$. Define $\tilde{N}_\ell$ to be the number of samples of (Y_ℓ, Z_ℓ) from (20) and recall N'_ℓ is the number of samples used to calculate $\bar{Z}_\ell$. Based on the work in [19], θ_ℓ can be selected such that the MSE of $\hat{W}_\ell$ – also averaged over possible realizations of δ_ℓ – is minimized.

By doing so, it follows that, see [19],

$$\text{MSE}(\hat{W}_\ell, \mathbb{E}[Y_\ell]) = \text{MSE}(\hat{Y}_\ell, \mathbb{E}[Y_\ell]) \left[1 - \rho_\ell^2 \left(\frac{1}{1 + \tilde{N}_\ell/N'_\ell} \right) \right], \quad (21)$$

where

$$\rho_\ell^2 = \frac{(\text{cov}(Y_\ell, Z_\ell))^2}{\mathbb{V}[Z_\ell]\mathbb{V}[Y_\ell]} \quad (22)$$

and

$$\theta_\ell^* = \frac{\text{cov}(Y_\ell, Z_\ell)}{\mathbb{V}[Z_\ell]} \left(\frac{1}{1 + \tilde{N}_\ell/N'_\ell} \right). \quad (23)$$

Since $0 \leq \rho_\ell^2 \leq 1$, (21) and (22) indicate that as the correlation between Y_ℓ and Z_ℓ increases, ρ_ℓ^2 approaches one, and thus the *MSE reduction factor* $\text{MSE}(\hat{W}_\ell, \mathbb{E}[Y_\ell])/\text{MSE}(\hat{Y}_\ell, \mathbb{E}[Y_\ell])$ goes to zero. Updating $\hat{W}_\ell$ to account for the optimal MSE, the MLCV correction estimator is given by

$$\hat{W}_\ell = \frac{1}{\tilde{N}_\ell} \sum_{i=1}^{\tilde{N}_\ell} \left(Y_\ell^{(i)} - \theta_\ell^*(Z_\ell^{(i)} - \bar{Z}_\ell) \right), \quad (24)$$

for $\ell = 1, \dots, L$, and $\hat{W}_0 = \hat{Y}_0$. Then the full MLCV estimator is defined as

$$\hat{Q}_L^{MLCV} = \sum_{\ell=0}^L \hat{W}_\ell.$$

In order to complete the task of calculating the number of samples $\tilde{N}_\ell$ per level to optimize the cost, we must find the MSE of this MLCV estimator first. It can be shown that

$$\text{MSE}(\hat{Q}_L^{MLCV}, \mathbb{E}[Q]) = \text{MSE}(\hat{Q}_L^{MLCV}, \mathbb{E}[Q_L]) + (\mathbb{E}[Q_L - Q])^2,$$

where the MSE's are also averaged over the realizations of δ_ℓ . Expanding the sampling error term, we find

$$\begin{aligned} \text{MSE}(\hat{Q}_L^{MLCV}, \mathbb{E}[Q_L]) &= \mathbb{E} \left[\left(\hat{Q}_L^{MLCV} - \mathbb{E}[Q_L] \right)^2 \right] \\ &= \mathbb{E} \left[\left(\hat{W}_0 - \mathbb{E}[Y_0] \right)^2 \right] + \dots + \mathbb{E} \left[\left(\hat{W}_L - \mathbb{E}[Y_L] \right)^2 \right] \\ &= \sum_{\ell=0}^L \text{MSE}(\hat{W}_\ell, \mathbb{E}[Y_\ell]) \\ &= \sum_{\ell=0}^L \text{MSE}(\hat{Y}_\ell, \mathbb{E}[Y_\ell]) \left[1 - \rho_\ell^2 \left(\frac{1}{1 + \tilde{N}_\ell/N'_\ell} \right) \right], \end{aligned}$$

where the last line is determined from (21). Thus the MSE of $\hat{Q}_L^{MLCV}$ is given by

$$\text{MSE}(\hat{Q}_L^{MLCV}, \mathbb{E}[Q]) = \sum_{\ell=0}^L \frac{1}{\tilde{N}_\ell} \mathbb{V}[Y_\ell] \left[1 - \rho_\ell^2 \left(\frac{1}{1 + \tilde{N}_\ell/N'_\ell} \right) \right] + (\mathbb{E}[Q_L - Q])^2, \quad (25)$$

where we used the relation $\text{MSE}(\hat{Y}_\ell, \mathbb{E}[Y_\ell]) = \mathbb{V}[Y_\ell]/\tilde{N}_\ell$. Since the discretization error term in (25) is the same as in (8), the MLMC convergence theory, as discussed in Section 2.3, can be applied to the simulations of W_ℓ in MLCV in order to guarantee a bound on the MSE.

3.4. Number of Samples $\tilde{N}_\ell$ for $\hat{W}_\ell$

In MLMC the number of samples needed for each level is estimated in order to attain a bound on the sampling error component of the MSE, see (12). In this formulation of MLCV, we must not only account for the sample size $\tilde{N}_\ell$, but also for N'_ℓ . The sample size $\tilde{N}_\ell$ is the required number of samples of (Y_ℓ, Z_ℓ) to achieve a sampling error bound of $\varepsilon^2/2$, while the value N'_ℓ is the number of samples used to find $\bar{Z}_\ell$.

In practice, we are able to determine the ratio $\tilde{N}_\ell/N'_\ell$ from a number of pilot runs, as discussed in Section 3.6, and update the MSE in (25) accordingly. Applying the same method to derive (12) and fixing the ratio $\tilde{N}_\ell/N'_\ell$, the optimal number of samples $\tilde{N}_\ell$ for $\hat{W}_\ell$ can be determined from (25) and is given by

$$\tilde{N}_\ell = \frac{2}{\varepsilon^2} \left[\sum_{k=0}^L \left(\mathbb{V}[Y_k] \left[1 - \rho_k^2 \left(\frac{1}{1 + \tilde{N}_k/N'_k} \right) \right] \mathcal{C}(W_k) \right)^{1/2} \right] \sqrt{\frac{\mathbb{V}[Y_\ell] \left[1 - \rho_\ell^2 \left(\frac{1}{1 + \tilde{N}_\ell/N'_\ell} \right) \right]}{\mathcal{C}(W_\ell)}}, \quad (26)$$

where ρ_k^2 is also estimated from the pilot simulations. We note that in (26), we take

$$\mathcal{C}(W_\ell) = \mathcal{C}_\ell.$$

This cost accounts for the samples of Y_ℓ and Z_ℓ . We consider the data acquisition for $\bar{Z}_\ell$ to be offline, thus the related cost is not included in the estimation of $\tilde{N}_\ell$; however, it is included in the total cost of implementing MLCV, as will be discussed in Section 3.7.

3.5. Number of Samples N'_ℓ for $\bar{Z}_\ell$

As can be observed from (25) or (26), the ratio $\tilde{N}_\ell/N'_\ell$ plays an important role in the MSE reduction achieved by $\hat{W}_\ell$ and the number of samples $\tilde{N}_\ell$ required by MLCV on each level. The smaller $\tilde{N}_\ell/N'_\ell$, the larger the MSE reduction and, hence, the smaller the estimated $\tilde{N}_\ell$ from (26). However, a small ratio $\tilde{N}_\ell/N'_\ell$ requires a large number N'_ℓ of independent samples to compute $\bar{Z}_\ell$ and, hence, a large overall computational cost. Therefore, it is necessary to determine an *optimal* value for $\tilde{N}_\ell/N'_\ell$.

In [19], the authors discuss the use of the so-called *generalized* MSE to find the optimal number of samples to minimize the product of the cost and $\text{MSE}(\hat{W}_\ell, \mathbb{E}[Y_\ell])$ from (21). Let $\zeta = \mathcal{C}(Q_{\ell-1})/\mathcal{C}_\ell$ denote the cost per sample of Z_ℓ (the cost of a level $\ell - 1$ sample) divided by the cost per sample of Y_ℓ , where $\mathcal{C}_\ell$ is given in (11). Following [19] and by minimizing the product of the total cost required on each level – given by $\zeta N'_\ell + (1 + \zeta)\tilde{N}_\ell$ – and $\text{MSE}(\hat{W}_\ell, \mathbb{E}[Y_\ell])$, the optimal ratio $\tilde{N}_\ell/N'_\ell$ is such that

$$N'_\ell = \max(0, s_1 - 1) \tilde{N}_\ell \quad (27)$$

with

$$s_1 = \left[\frac{\rho_\ell^2}{\zeta(1 - \rho_\ell^2)} \right]^{1/2}.$$

If $N'_\ell = 0$ this implies we do not implement MLCV on level ℓ , instead, we let $W_\ell = Y_\ell$.

The potential issue with using (27) is that for $\rho_\ell^2 \approx 1$, the value of N'_ℓ can get quite large. In some tests this may occur, and depending on the desired implementation, the number of samples may alternatively be determined by

$$N'_\ell = s_2 \tilde{N}_\ell, \quad (28)$$

where the integer $s_2 > 1$ is preselected. In the numerical experiments of this work, we choose the minimum N'_ℓ from (27) and (28), i.e.,

$$N'_\ell = \min(s_2, \max(0, s_1 - 1)) \tilde{N}_\ell, \quad (29)$$

with $s_2 = 10$. This selected value of $s_2 = 10$ will insure more than 90% of ρ_ℓ^2 is maintained in the MSE reduction factor (see (21)), when s_2 is used to evaluate (29). For values $s_2 > 10$, the benefit of maintaining a slightly larger percentage of ρ_ℓ^2 is more outweighed by the computational cost of generating additional solution realizations.

3.6. MLCV Implementation Details and Algorithms

In this section we discuss the implementation of MLCV as outlined in Algorithms 1 and 2. The MLCV algorithm can be organized into two different components: the pilot run and the full MLCV run. For both, we will describe the case that $\ell > 0$. When $\ell = 0$, we simply simulate the level $\ell = 0$ model and take a sample average to calculate $\hat{W}_0 = \hat{Y}_0$. Additionally, we assume that L is large enough to meet the desired discretization error. In practice, we check that convergence is achieved by estimating α and β , as introduced in Section 2.3.

3.6.1. Pilot Run – Reduced Basis and Sample Size Determination

As in the case of MLMC, the goal of the pilot run is to find estimates of $\mathbb{E}[Y_\ell]$ and $\mathbb{V}[Y_\ell]$, to ensure level L is fine enough such that the discretization error is bounded by $\varepsilon^2/2$, and to determine the sample sizes $\tilde{N}_\ell$ and N'_ℓ . However, determining $\tilde{N}_\ell$ and N'_ℓ requires the estimation of ρ_ℓ^2 and the derivation of the reduced basis $\mathbf{U}_\ell^c$ (along with the associated interpolation rules) to construct the control variates Z_ℓ . These are specific to this MLCV approach and will be additionally performed during the pilot run stage summarized in Algorithm 1.

In particular, a small set of $N_p \geq m_{L-1}$ samples of $\boldsymbol{\xi}$ is generated and applied to level $\ell - 1$ and ℓ models to obtain samples of Y_ℓ , as well as estimates of $\mathbb{E}[Y_\ell]$ and $\mathbb{V}[Y_\ell]$. We note that N_p may vary from one level to the other, provided that some prior knowledge of $\mathbb{V}[Y_\ell]$ is available. As discussed in Section 3.1, a reduced basis $\mathbf{U}_{\ell-1}^c$ and an associated interpolation rule for the level $\ell - 1$ solution vector $\mathbf{q}_{\ell-1}$ is found via the ID of the coarse grid data matrix $\mathbf{U}_{\ell-1}$. The level ℓ reduced basis $\mathbf{U}_\ell^c$ is generated by simulating the level ℓ model at the samples of $\boldsymbol{\xi}$ identified by the level $\ell - 1$ ID and used – in conjunction with the same interpolation rule – to acquire samples of $\mathbf{q}_\ell^{ID}$, and subsequently Q_ℓ^{ID} and Z_ℓ . These bases are stored and reused in the MLCV runs of Section 3.6.2. Following this, ρ_ℓ^2 and $\tilde{N}_\ell/N'_\ell$ are computed from (22) and (29), respectively. After each level the change in the sample means of Y_ℓ must be determined to verify that the intended discretization accuracy is met. If not, this process is repeated with larger values of L . If the value of L is sufficient, the pilot run ends and the values of $\tilde{N}_\ell$ and N'_ℓ are determined from (26) and (22), respectively.

Remark. While the variance and covariance are estimated from the pilot run, these values, in practice, may be updated during the main MLCV run described next.

3.6.2. Main MLCV Run

After the total number of samples $\tilde{N}_\ell$ and N'_ℓ for each level are estimated, the full MLCV run is completed following Algorithm 2. Much like the pilot run, $\tilde{N}_\ell$ realizations of $Y_\ell^{(i)}$ are collected from the level ℓ and $\ell - 1$ models. Using the $\tilde{N}_\ell$ samples of $\mathbf{q}_{\ell-1}$, and the level $\ell - 1$ reduced basis $\mathbf{U}_{\ell-1}^c$ from the pilot run, least squares (17) is performed on each sample to find the coefficient vectors $\mathbf{c}_{\ell-1}(\boldsymbol{\xi})$. Given this interpolation rule and the corresponding level ℓ basis $\mathbf{U}_\ell^c$, the samples of $\mathbf{q}_\ell^{ID}$ (hence Q_ℓ^{ID}) and subsequently $Z_\ell^{(i)}$ are generated using (18) and (14), respectively. Following these, $\bar{Z}_\ell$ is computed from (19) using N'_ℓ independent samples of Z_ℓ drawn as described above. After calculating θ_ℓ^* in (23), we have all the components of the MLCV estimator, and use (24) to find $\hat{W}_\ell$ and lastly $\hat{Q}_L^{MLCV}$.

Algorithm 1: MLCV Reduced Basis Identification and Sample Size Determination

```

for  $\ell = 0, \dots, L$  do
  if  $\ell = 0$  then
     $\{\xi^{(i)}\}_{i=1}^{N_p} \longrightarrow \{Q_\ell^{(i)}\}_{i=1}^{N_p} \longrightarrow \{Y_\ell^{(i)}\}_{i=1}^{N_p}$ 
  else
     $\{\xi^{(i)}\}_{i=1}^{N_p} \longrightarrow \{\mathbf{q}_{\ell-1}^{(i)}\}_{i=1}^{N_p}, \{\mathbf{q}_\ell^{(i)}\}_{i=1}^{N_p} \longrightarrow \{Q_{\ell-1}^{(i)}\}_{i=1}^{N_p}, \{Q_\ell^{(i)}\}_{i=1}^{N_p} \longrightarrow \{Y_\ell^{(i)}\}_{i=1}^{N_p}$ 

     $\mathbf{U}_{\ell-1} = [\mathbf{q}_{\ell-1}^{(1)} \cdots \mathbf{q}_{\ell-1}^{(N_p)}] \xrightarrow{\text{ID via (15)}} \mathbf{U}_{\ell-1}^c \mathbf{C}_{\ell-1}, \{\xi^{(i_k)}\}_{i=1}^r$ 

     $\{\xi^{(i_k)}\}_{i=1}^r \longrightarrow \mathbf{U}_\ell^c; \mathbf{U}_\ell^{ID} = \mathbf{U}_\ell^c \mathbf{C}_{\ell-1} \longrightarrow \{Q_\ell^{ID(i)}\}_{i=1}^{N_p} \longrightarrow \{Z_\ell^{(i)}\}_{i=1}^{N_p}$ 

    Calculate  $\rho_\ell^2$  and  $\tilde{N}_\ell/N'_\ell$ , respectively, via (22) and (29)
  end
  Estimate  $\mathbb{V}[Y_\ell]$ 
end
for  $\ell = 0, \dots, L$  do
  Calculate  $\tilde{N}_\ell$  using (26)
  Calculate  $N'_\ell$  using (29)
end

```

Algorithm 2: Main MLCV Run

```

for  $\ell = 0, \dots, L$  do
  if  $\ell = 0$  then
     $\{\xi^{(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow \{Q_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow \{Y_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow \hat{W}_\ell = \hat{Y}_\ell$ 
  else
    %  $\bar{Z}_\ell$  calculation:
     $\{\xi^{(i)}\}_{i=1}^{N'_\ell} \longrightarrow \{\mathbf{q}_{\ell-1}^{(i)}\}_{i=1}^{N'_\ell} \longrightarrow \{Q_{\ell-1}^{(i)}\}_{i=1}^{N'_\ell}$ 

     $\mathbf{c}_{\ell-1}^{(i)} = \min_{\hat{\mathbf{c}} \in \mathbb{R}^r} \|\mathbf{U}_{\ell-1}^c \hat{\mathbf{c}} - \mathbf{q}_{\ell-1}^{(i)}\|$  for  $i = 1, \dots, N'_\ell$ 

     $\mathbf{q}_\ell^{ID(i)} = \mathbf{U}_\ell^c \mathbf{c}_{\ell-1}^{(i)}$  for  $i = 1, \dots, N'_\ell \longrightarrow \{Q_\ell^{ID(i)}\}_{i=1}^{N'_\ell} \longrightarrow \{Z_\ell^{(i)}\}_{i=1}^{N'_\ell} \longrightarrow \bar{Z}_\ell$ 

    %  $Z_\ell$  sampling:
     $\{\xi^{(i)}\}_{i=1}^{\tilde{N}_\ell}$  (independent of  $\{\xi^{(i)}\}_{i=1}^{N'_\ell}$  above)  $\longrightarrow \{\mathbf{q}_{\ell-1}^{(i)}\}_{i=1}^{\tilde{N}_\ell}, \{\mathbf{q}_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow$ 
     $\{Q_{\ell-1}^{(i)}\}_{i=1}^{\tilde{N}_\ell}, \{Q_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow \{Y_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell}$ 

     $\mathbf{c}_{\ell-1}^{(i)} = \min_{\hat{\mathbf{c}} \in \mathbb{R}^r} \|\mathbf{U}_{\ell-1}^c \hat{\mathbf{c}} - \mathbf{q}_{\ell-1}^{(i)}\|$  for  $i = 1, \dots, \tilde{N}_\ell$ 

     $\mathbf{q}_\ell^{ID(i)} = \mathbf{U}_\ell^c \mathbf{c}_{\ell-1}^{(i)}$  for  $i = 1, \dots, \tilde{N}_\ell \longrightarrow \{Q_\ell^{ID(i)}\}_{i=1}^{\tilde{N}_\ell} \longrightarrow \{Z_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell}$ 

    %  $W_\ell$  sampling:
     $\{W_\ell^{(i)}\}_{i=1}^{\tilde{N}_\ell}$  from (24)  $\longrightarrow \hat{W}_\ell$ 
  end
end
Set  $\hat{Q}_L^{MLCV} = \sum_{\ell=0}^L \hat{W}_\ell$ 

```

3.7. MLCV Cost Breakdown

After determining the total number of samples required to optimize the cost with the MSE bound as a constraint, the total cost of implementing MLMC is given in (10). By modifying MLMC through the use of the control variate Z_ℓ , we are able to achieve further reduction in the MSE of the estimator $\hat{W}_\ell$, as long as Z_ℓ is sufficiently correlated with Y_ℓ . In doing so, the sample size for Y_ℓ and thus Z_ℓ needed to meet the sampling error bound, while optimizing the total cost, is reduced. However, sampling Z_ℓ and estimating its expectation, as required in (13), lead to additional computational cost that we quantify next.

Cost of pilot run. Recall that the construction of Z_ℓ is based on the identification of the reduced basis $\mathbf{U}_{\ell-1}^c$ (using which $\mathbf{U}_\ell^c$ is generated) and the associated least squares coefficients $\mathbf{c}_{\ell-1}$ computed from (17). Following Algorithm 1, the former is based on the ID of the coarse data matrix $\mathbf{U}_{\ell-1}$ consisting of N_p pilot samples of $\mathbf{q}_{\ell-1} \in \mathbb{R}^{m_{\ell-1}}$. From [30], the cost to implement a rank r ID on $\mathbf{U}_{\ell-1}$ is $\mathcal{O}(m_{\ell-1}\tilde{N}_\ell n_r)$, where $n_r < \tilde{N}_\ell$. This means that the cost per sample of Z_ℓ is $\mathcal{O}(m_{\ell-1}n_r)$. In many scenarios of practical interest, the QoI Q depends on a $\mathbf{q}$ consisting of a considerably smaller set of solution degrees of freedom, e.g., average outflow temperature or lift/drag on an airfoil, as opposed to those of the entire solution. In such cases, $m_{\ell-1} \ll M_{\ell-1}$ and we may ignore the cost of performing the ID of $\mathbf{U}_{\ell-1}$. Assuming that (i) the factorization to solve (17) is formed only once, e.g, using a reduced QR decomposition with cost $\mathcal{O}(r^2 m_{\ell-1})$, (ii) $r \ll m_\ell$, and (iii) $\tilde{N}_\ell \gg 1$, the cost of solving for the coefficients $\mathbf{c}_{\ell-1}$ for each sample and applying them to $\mathbf{U}_\ell^c$ to obtain $\mathbf{q}_\ell^{ID}$ is $\mathcal{O}(r m_\ell)$, which again may be ignored when $m_\ell \ll M_\ell$. From the N_p pilot samples, r of them form the reduced bases $\mathbf{U}_{\ell-1}^c$ and $\mathbf{U}_\ell^c$ the rest will be recycled in the main MLCV run, assuming $N_p \leq \tilde{N}_\ell$. Therefore, the net cost of performing the pilot run on each level is the same as the cost of generating r samples of level ℓ and $\ell - 1$ solutions to form the reduced bases $\mathbf{U}_{\ell-1}^c$ and $\mathbf{U}_\ell^c$. This gives rise to the total cost of $\sum_{\ell=1}^L r (\mathcal{C}(Q_{\ell-1}) + \mathcal{C}(Q_\ell))$ for the pilot run.

Cost of main MLCV run. Sampling Z_ℓ in (20) entails using the $\tilde{N}_\ell$ samples of $\mathbf{q}_{\ell-1}$ drawn for Y_ℓ and performing least squares regression to generate the approximation (16). As discussed previously, these least squares solves have negligible cost when $m_\ell \ll M_\ell$. Therefore, there is no additional cost for sampling Z_ℓ in (20) relative to its MLMC counterpart. With a similar argument, sampling Z_ℓ in (20) to compute $\tilde{Z}_\ell$ requires only generating N'_ℓ independent samples of $\mathbf{q}_{\ell-1}$. Therefore, the cost of the main run is $\tilde{N}_0 \mathcal{C}(Q_0) + \sum_{\ell=1}^L \tilde{N}_\ell (\mathcal{C}(Q_{\ell-1}) + \mathcal{C}(Q_\ell)) + \sum_{\ell=1}^L N'_\ell \mathcal{C}(Q_{\ell-1})$.

Total cost. Adding the cost of pilot and main MLCV runs, we arrive at the total cost of MLCV given by

$$\mathcal{C}(\hat{Q}_L^{MLCV}) = \tilde{N}_0 \mathcal{C}(Q_0) + \sum_{\ell=1}^L (\tilde{N}_\ell + r) (\mathcal{C}(Q_{\ell-1}) + \mathcal{C}(Q_\ell)) + \sum_{\ell=1}^L N'_\ell \mathcal{C}(Q_{\ell-1}). \quad (30)$$

3.8. Discussion on Use of MLCV

Before proceeding to the numerical experiments and results, we first discuss the types of problems that benefit from the application of MLCV. When deciding to apply a MC based method to solve a high-dimensional problem in uncertainty quantification, the practitioner must consider the QoI's variance – more specifically coefficient of variation (COV) – as well as the cost of simulating the model on a fine enough discretization. For problems with small COV and cheap-to-evaluate models, it can be the case that standard MC will perform as well, or may outperform MLMC. On the other hand, the practitioner will likely decide to use MLMC over MC if the QoI shows a large COV and the solver exhibits a fast cost growth between refined grids ($\mathcal{C}_\ell \gg M_\ell$). Since MLCV is a multilevel method, it benefits from the same traits as MLMC: a large COV and a fast cost growth. However, unlike MLMC, MLCV relies on a low-rank representation of the solution, in the sense defined in Section 3.1, thus making the approach most effective in outperforming MLMC on problems exhibiting small solution ranks on all the levels. When this occurs we find that Q_ℓ^{ID} approximates the fine grid data Q_ℓ accurately, and that the value for ρ_ℓ^2 is much closer to one, thus resulting in an improved MSE for $\hat{W}_\ell$.

A final comment must be made regarding the total cost of MLCV. In the case that significant MSE reduction for $\hat{W}_\ell$ does occur, we are not necessarily guaranteed cost reduction. In (30) the number of samples needed on levels $\ell > 0$ is always larger than r . In MLMC, this constraint is not present (see (10)). For scenarios where the data exhibits a small variance or small values of ε are provided as the desired tolerance, the resulting number of samples on the finer levels may be comparable to r . In such cases, MLCV may not outperform MLMC. This suggests MLCV is more beneficial for smaller selections of ε .

4. Numerical Results

In this section we consider two test cases to compare the resulting error and cost estimates for MLMC and MLCV. In the first case we consider the generalized eigenvalue problem for finding the frequency of the first natural mode of vibration of an L-shaped elastic structure. For the second case we consider a thermally driven cavity flow with stochastic boundary temperatures previously studied in [31, 32, 33, 34].

Following each case, we present several results when comparing MLCV and MLMC for both tests. In particular, we consider the convergence of ρ_ℓ^2 in (22), and subsequently the MSE at each level, which controls the sampling error of MLCV. For values of ρ_ℓ^2 close to one, a significant reduction in the MSE of $\hat{W}_\ell$ can be observed, and thus a reduction in the number $\tilde{N}_\ell$ of required samples per level. We further provide results on the total cost and relative error of MLCV and MLMC. The total cost, as a function of ε , is determined from (26), (29), and (30). For the eigenvalue problem, the cost is defined to be $\mathcal{C}_\ell = M_\ell$, while the thermally driven flow problem cost is approximated by an average CPU time. The final result is the relative error in estimating the mean of QoI – to be specified – on all discretization levels. The errors are generated based on a reference solution $\hat{Q}_L^{\text{ref}}$ obtained via sparse polynomial chaos expansion [3, 34] on the finest available discretization for each model. Specifically,

$$e_\ell^{\text{MLMC}} = \frac{\left| \sum_{k=0}^{\ell} \hat{Y}_k - \hat{Q}_L^{\text{ref}} \right|}{\left| \hat{Q}_L^{\text{ref}} \right|} \quad (31)$$

and

$$e_\ell^{\text{MLCV}} = \frac{\left| \sum_{k=0}^{\ell} \hat{W}_k - \hat{Q}_L^{\text{ref}} \right|}{\left| \hat{Q}_L^{\text{ref}} \right|}, \quad (32)$$

for $\ell = 1, \dots, L$.

4.1. Test 1: Natural Frequency of a 2D Linear Elasticity Problem

For the first test, we seek to find the expected value of the smallest vibration frequency, ω , of the linear elasticity problem on an L-shape domain $\mathcal{D}$ as shown in Figure 1 (a). We consider zero Dirichlet boundary conditions on the top and right sides of the domain along y and x directions, respectively. The Young's modulus of the medium, given by the log-normal random field

$$E(\mathbf{x}, \boldsymbol{\xi}) = \bar{E} + \exp(G(\mathbf{x}, \boldsymbol{\xi})), \quad \mathbf{x} \in \mathcal{D},$$

is the source of uncertainty in this test. Here, $G(\mathbf{x}, \boldsymbol{\xi})$ is a Gaussian random field represented by the Karhunen-Loève expansion

$$G(\mathbf{x}, \boldsymbol{\xi}) = \sum_{i=1}^d \sqrt{\lambda_i} \phi_i(\mathbf{x}) \xi_i,$$

where the random variables ξ_i are independent standard Gaussian and λ_i are the d largest eigenvalues with corresponding eigenfunctions $\phi_i(\mathbf{x})$ determined by the Gaussian covariance function

$$\mathcal{K}(\mathbf{x}_1, \mathbf{x}_2) = \sigma^2 \exp\left(-\frac{\|\mathbf{x}_1 - \mathbf{x}_2\|_2^2}{l}\right).$$

Figure 1: (a) Schematic of the L-shaped domain. (b)-(f) Meshes on the five different levels with increasing resolution.

In our experiments, we set $\bar{E} = 0.1$, $d = 36$, $\sigma^2 = 0.33$, and $l = 0.33$, and assume the medium has a unit density.

To compute ω , we consider the solution of the random eigenvalue problem

$$\mathbf{K}(\boldsymbol{\xi})\mathbf{q}(\boldsymbol{\xi}) = \omega^2(\boldsymbol{\xi})\mathbf{M}\mathbf{q}(\boldsymbol{\xi}), \quad (33)$$

where $\mathbf{K}(\boldsymbol{\xi})$ and $\mathbf{M}$ are, respectively, the stiffness and mass matrices associated with finite element discretization of the linear elasticity equations on uniform triangular meshes of varying size. Additionally, $\mathbf{q}(\boldsymbol{\xi})$ is the natural vibration mode (eigenvector) corresponding to the eigenvalue $\omega^2(\boldsymbol{\xi})$ given by

$$\omega^2(\boldsymbol{\xi}) = \frac{\mathbf{q}(\boldsymbol{\xi})^T \mathbf{K}(\boldsymbol{\xi}) \mathbf{q}(\boldsymbol{\xi})}{\mathbf{q}(\boldsymbol{\xi})^T \mathbf{M} \mathbf{q}(\boldsymbol{\xi})}. \quad (34)$$

Here, QoI is the smallest frequency ω satisfying (33) and is a functional of $\mathbf{q}(\boldsymbol{\xi})$, i.e. $Q = Q(\mathbf{q}(\boldsymbol{\xi}))$, through (34). With the given uncertainty, the COV of Q is approximately 15%.

The finite element discretization is done using FENiCS [35] with five different meshes of varying size, as displayed in Figure 1 (b)-(f). The number of degrees of freedom of the meshes are $M_0 = 36$, $M_1 = 106$, $M_2 = 212$, $M_3 = 856$, and $M_4 = 3296$. The resulting eigenvector $\mathbf{q}_\ell$, from each level ℓ simulation, has M_ℓ degrees of freedom. After eigenvectors $\mathbf{q}_\ell^{ID}$ are found, the corresponding (saved) stiffness and mass matrices, $\mathbf{K}$ and $\mathbf{M}$, are used in (34) to generate Q_ℓ^{ID} and the control variates Z_ℓ .

For the application of MLMC and MLCV we consider two cases to test the relative performance of these methods when altering the difference in the number of degrees of freedom between adjacent levels. In the first case, *A*, we apply MLMC and MLCV to all five levels. In the second case, *B*, we consider only using three of the five meshes from levels $\ell = 0, 2, 4$. This is worth testing, as both methods rely on the differences between QoIs on adjacent levels. The results between *A* and *B* will indicate which method can adjust to the larger gap in the number of degrees of freedom between levels for this problem.

To select the rank on each level, we consider values of the correlation ρ_ℓ^2 (see (22)) for different values of r , based on pilot run data. Figure 2 demonstrates how the values of ρ_ℓ^2 change with rank r for case *A* (left) and case *B* (right). In order to have $\rho_\ell^2 > 0.90$ we use $r = 10$ for all levels, with the exception of $r = 15$ on the finest level, $\ell = 4$, in case *A*.

Figure 2: Change of the correlation ρ_ℓ^2 in (22) as a function of the ID rank r for (left) case A and (right) case B.

4.1.1. Results of Test 1

For the results presented in this section, the “o” markers denote case A data and the “x” markers denote case B data. The first set of results are related to the MLMC convergence guarantee in Section 2.3. Figure 3 (left) shows the sample mean of Q_ℓ , Y_ℓ and W_ℓ as a function of the number of degrees of freedom on each level, M_ℓ . We note that the sample mean of Q_ℓ remains relatively constant, while those of Y_ℓ and W_ℓ decay at the same rate. The data displayed allows us to estimate the values of α . Based on the slope, we have $|\mathbb{E}[Q_\ell - Q_{\ell-1}]| \approx M_\ell^{-0.92}$ and $|\mathbb{E}[Q_\ell - Q_{\ell-1}]| \approx M_\ell^{-0.64}$ for case A and B, respectively. Thus it follows that $\alpha_A \approx 0.92$ and $\alpha_B \approx 0.64$. Figure 3 (right) shows the sample variance of Q_ℓ and Y_ℓ as a function of the degrees of freedom on each level, M_ℓ . The data displayed allows us to estimate the values of β . From Figure 3 (right) we have $\mathbb{V}[Q_\ell - Q_{\ell-1}] \approx M_\ell^{-1.7}$ and $\mathbb{V}[Q_\ell - Q_{\ell-1}] \approx M_\ell^{-1.3}$, for case A and B, respectively. This implies that $\beta_A \approx 1.7$ and $\beta_B \approx 1.3$. The cost of the FE solver, as well as the cost to compute ω_1^2 , indicate that $\gamma \approx 1$, i.e. $\mathcal{C}_\ell \lesssim M_\ell$, suggesting that for case A and B we may consider a tolerance of $\varepsilon \sim \mathcal{O}(0.0005)$. We remark that, while the value of α_B indicates the discretization error cannot meet the specified tolerance, we argue that this is an estimation of $\mathbb{E}[Q_L - Q]$, and that this approximation of α is better obtained from test A results.

Next we consider the MSE reduction in Figure 4. Figure 4 (left) shows the increasing values of ρ_ℓ^2 as a function of ℓ , as well as the decreasing values of the MSE reduction factor. For both cases A and B, $\rho_\ell^2 > 0.95$ for $\ell > 0$. This is ideal, as this reduction in the MSE means that fewer samples are required to obtain the MSE tolerance. Figure 4 (right) displays the number of samples required on each level for both methods when using $\varepsilon = 0.0005$. Note that the number of samples for MLCV includes those used to calculate $\bar{Z}_\ell$. For case A the number of samples required for MLCV is slightly smaller than that of MLMC. A more significant reduction in the number of samples needed for MLCV is observed for case B. This suggests that MLCV is more adaptable when large gaps are added between M_ℓ on adjacent levels.

To compare the costs for these two methods, we turn to Table 1, which displays the cost estimates for several cases as a function of prescribed ε based on (10) and (30), where $\mathcal{C}_\ell = M_\ell$. For comparison, the cost of MC on the finest level is also reported. For small values of ε , an improvement in cost is observed for MLCV. When requiring a tolerance of $\varepsilon = 0.0005$ we observe a cost ratio of 75% for case A and 58% for case B. Comparing the total costs of these two cases, we see that the MLCV method for case B outperforms

Figure 3: (Left) Convergence of the sample mean of Y_ℓ and W_ℓ , while that of Q_ℓ remains relatively constant. (Right) Convergence of the sample variance of Y_ℓ in comparison to that of Q_ℓ . Superscripts A and B refer to the case A and B of the L-shaped domain problem, respectively.

case A . For MLMC, it is the opposite. By adding larger gaps between levels, the performance of MLMC decays, while that of MLCV improves. When increasing the tolerance, there is a reduction in the cost gain of MLCV, as fewer samples are needed to maintain the MSE tolerance. The cost of the r samples for the basis is diminishing the success of the MLCV method. When $\varepsilon > 0.003$, the cost of performing MLCV is greater than the cost of performing MLMC.

Table 1: Cost of MC, MLMC, and MLCV for Test 1 problem. The difference in cost of MC for $\varepsilon = 0.0005$ is due to the small number of samples used.

ε	Levels	Cost MC	Cost MLMC	Cost MLCV	$\frac{\text{Cost MLCV}}{\text{Cost MLMC}}$
0.0005	0, 1, 2, 3, 4	8.7e7	2.8e6	2.1e6	0.75
0.0005	0, 2, 4	6.6e7	3.3e6	1.9e6	0.58
0.001	0, 1, 2, 3	2.2e7	6.5e5	5.0e5	0.77
0.003	0, 1, 2, 3	2.4e6	7.3e4	7.1e4	0.97
0.005	0, 1, 2, 3	8.7e5	2.7e4	3.7e4	1.40

These next results compare the estimated MSEs and relative errors of the methods. Figure 5 (left) displays the estimated MSE for both methods. For both case A and B , the MSE estimates for $\hat{Y}_\ell$ and $\hat{W}_\ell$ are on the same order, which is expected. The slight difference is due to the number of terms in the sampling error. Since case A has more levels, most of the sampling error terms should be smaller than those of case B , as the total sampling error for both test cases must attain the same tolerance of $\varepsilon^2/2$.

To see how this new MLCV method performs against MLMC in terms of relative error, we compare the values computed in (31) and (32). Figure 5 (right) displays the convergence of the relative error for both methods as the levels are refined. We note that, due to the cost of simulating the fine model, these results are based on an average of 10 runs of MLMC and MLCV, where the data in each run is not completely independent from data in other runs. This result indicates that, experimentally, both methods converge on the same order. However the two methods do not cost the same to derive this result. As stated earlier, using the cost and the total number of samples on each level, it is found that to achieve a sampling error of $\varepsilon = 0.0005$, MLCV only requires about 75% and 58% of the computational cost of MLMC for case A and

Figure 4: (Left) MSE reduction factor (MSERF) between $\hat{W}_\ell$ and $\hat{Y}_\ell$, as well as the value of ρ_ℓ^2 on each level. (Right) Total number of samples needed to achieve the sampling error $\varepsilon = 0.0005$ for MLMC and MLCV. MLCV sample size includes the number required to determine $\hat{Z}_\ell$. Superscripts A and B refer to the case A and B of the L-shaped domain problem, respectively.

B , respectively.

4.2. Test 2: Thermally Driven Cavity Flow

For the second implementation of MLCV, we consider a thermally driven flow problem in a square domain as described in [36, 37] and illustrated in Figure 6. The left vertical wall has a random temperature T_h with mean $\bar{T}_h$, while the right vertical wall, referred to as the cold wall, has a spatially varying stochastic temperature $T_c < T_h$ with constant mean $\bar{T}_c$. Both top and bottom walls are assumed to be adiabatic. The reference temperature and the reference temperature difference are defined as $T_{ref} = \bar{T}_c$ and $\Delta T_{ref} = T_h - \bar{T}_c$, respectively. Under small temperature difference assumption, i.e., Boussinesq approximation, the normalized governing equations are given by, [37],

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{\text{Pr}}{\sqrt{\text{Ra}}} \nabla^2 \mathbf{u} + \text{Pr} \Theta \mathbf{e}_y, \\ \nabla \cdot \mathbf{u} &= 0, \\ \frac{\partial \Theta}{\partial t} + \nabla \cdot (\mathbf{u} \Theta) &= \frac{1}{\sqrt{\text{Ra}}} \nabla^2 \Theta, \end{aligned} \tag{35}$$

where $\mathbf{e}_y$ is the unit vector $(0, 1)$, $\mathbf{u} = (u, v)$ is velocity vector field, $\Theta = (T - T_{ref})/\Delta T_{ref}$ is normalized temperature, p is pressure, and t is time. Zero velocity boundary conditions on all walls (in both directions) are assumed. For more details on the normalization of the variables in (35), we refer the interested reader to [36, 37]. Prandtl and Rayleigh numbers are defined, respectively, as $\text{Pr} = \nu/\alpha$ and $\text{Ra} = g\tau\Delta T_{ref}L^3/(\nu\alpha)$. Specifically, L is the length of the cavity, g is gravitational acceleration, ν is kinematic viscosity, α is thermal diffusivity, and the coefficient of thermal expansion is given by τ . In this example, we set $g = 10$, $L = 1$, $\tau = 0.5$, and $\text{Pr} = 0.71$. We use finite volume for the discretization of (35). The QoI Q is the spatial average of steady-state heat flux along the hot wall.

The uncertain quantities include the initial temperatures on the cold and hot wall (T_c and T_h , respectively), as well as the viscosity parameter ν . The cold wall temperature, T_c , is given by a Karhunen-Loève

Figure 5: (Left) MSE for $\hat{Y}_\ell$ and $\hat{W}_\ell$. (Right) Comparison of the relative errors (31) and (32) in estimating the mean QoI at each level for MLMC and MLCV. Superscripts A and B refer to the case A and B of the L-shaped domain problem, respectively.

Figure 6: Schematic of Test 2, the thermally driven flow problem with random temperature along the hot (west) wall and spatially varying stochastic temperature along the cold (east) wall.

Figure 7: Change of the correlation ρ_ℓ^2 in (22) as a function of the ID rank r for Test 2.

type expansion with 50 terms,

$$T_c(y, \boldsymbol{\xi}) = \bar{T}_c + \sigma \sum_{i=1}^{50} \sqrt{\lambda_i} \phi_i(y) \xi_i, \quad (36)$$

where λ_i and $\phi_i(y)$ are the eigenvalues and eigenfunctions of the exponential correlation function

$$\mathcal{K}(y_1, y_2) = \sigma^2 \exp\left(-\frac{|y_1 - y_2|}{l}\right), \quad (37)$$

with $\sigma = 2$ and $l = 0.15$. In (36), we set $\bar{T}_c = 100$ and assume the random variables $\xi_i \sim U[-1, 1]$, $i = 1, \dots, 50$, are independently and uniformly distributed over $[-1, 1]$. Additionally, we assume $T_h \sim U[105, 109]$ and $\nu \sim U[0.004, 0.01]$. This brings the total number of random inputs to $d = 52$. We note that the COV for the QoI of each simulation with this setup is about 25%.

For the numerical setup of this problem, we consider four nested uniform grids with the sizes of 16×16 , 32×32 , 64×64 , and 128×128 . This implies we have $M_\ell = s^\ell M_0$, with $s = 4$. The resulting vector of heat flux values along the hot wall, $\mathbf{q}$, is of length $m_\ell = \sqrt{M_\ell}$. As stated earlier, since $m_\ell \ll M_\ell$, the added cost of performing ID and least squares is discounted.

To select the rank on each level, we consider values of the correlation ρ_ℓ^2 (see (22)) for different values of r , based on pilot run data. Figure 7 demonstrates how the values of ρ_ℓ^2 change with rank r . In order to have $\rho_\ell^2 > 0.90$ we use $r = 10$ for all levels.

4.2.1. Results of Test 2

We first consider the results corresponding to the MLMC convergence guarantee in Section 2.3. Figure 8 (left) shows the sample mean of Q_ℓ , Y_ℓ and W_ℓ as a function of the degrees of freedom on each level, M_ℓ . We note that the sample mean of Q_ℓ remains relatively constant, while those of Y_ℓ and W_ℓ decay at the same rate. The data displayed allows us to estimate the values of α . Since the decay of data Y_ℓ , $\ell = 1, 2$, and 3 does not show a clear convergence, i.e., $\ell = 0$ discretization is *too coarse*, we use the data based on Y_2 and Y_3 to approximate α . By doing so, we estimate $\mathbb{E}[Q_\ell - Q_{\ell-1}] \approx M_\ell^{-1.0}$, indicating that $\alpha \approx 1.0$. We note this estimate was then confirmed by using data obtained from the 256×256 mesh. Figure 8 (right) shows the sample variance of Q_ℓ and Y_ℓ as a function of the degrees of freedom on each level, M_ℓ . The data displayed allows us to estimate the values of β . We estimate $\mathbb{V}[Q_\ell - Q_{\ell-1}] \approx M_\ell^{-1.8}$, indicating that

Figure 8: (Left) Convergence of the sample mean of Y_ℓ and W_ℓ , while that of Q_ℓ remains relatively constant. (Right) Convergence of the sample variance of Y_ℓ in comparison to that of Q_ℓ .

$\beta \approx 1.8$. By using the average CPU time to determine the cost of the solver, we find that $\gamma \approx 1.65$. This implies that the minimum MSE bound we can have is $\varepsilon \sim \mathcal{O}(0.00005)$ when using $M_L = 128^2$. Due to cost constraints, we select $\varepsilon = 0.001$ for the numerical experiments.

Next we consider the MSE reduction as well as the sample size reduction in Figure 9. Figure 9 (left) shows the values of ρ_ℓ^2 , where $\rho_\ell^2 > 0.85$ for levels 1, 2, and 3, as well as the decay of the MSE reduction factor. As the MSE reduction factor approaches zero, the number of samples MLCV requires on each level decays. Figure 9 (right) displays the number of samples required on each level for both methods when using $\varepsilon = 0.001$. It is clear that fewer samples are needed on each level for MLCV than for MLMC. And, as finer levels are approached, the gap between these two values increases. Because of this, we will see a significant cost payoff.

To compare the cost of performing MLMC and MLCV for this application, we consider the cost estimates displayed in Table 2, as determined by (10) and (30). Due to high simulation costs, we consider only four levels. We see that for $\varepsilon \leq 0.001$ MLCV halves the cost of MLMC. As the value of ε is increased, the cost difference of the two methods is smaller. For $\varepsilon \geq 0.005$, we can observe that MLCV has a larger computational cost than MLMC. The cost of the r samples for the basis is diminishing the success of the MLCV method.

The final set of results compare the estimated MSEs and relative errors of the two methods. Figure 10 (left) displays the estimated MSE for both methods. We see that the MSE estimates for $\hat{Y}_\ell$ and $\hat{W}_\ell$ are on the same order. To determine the relative accuracy of the two methods, we next compute the relative errors (31) and (32). Figure 10 (right) displays the convergence of the relative error for both methods as the levels are refined. We note that, due to the cost of simulating the fine model, these results are based on an average of 10 runs of MLMC and MLCV, where the data in each run is not completely independent from the other runs. As expected, we observe that as the mesh is refined for both methods, the relative errors improve. The notable result is that the relative errors for both methods remain essentially the same, while, in fact, MLCV only requires 48% of the computational cost of MLMC.

Figure 9: (Left) MSE reduction factor (MSERF) between $\hat{W}_\ell$ and $\hat{Y}_\ell$, as well as the value of ρ_ℓ^2 on each level. (Right) Total number of samples needed to attain the sampling error for MLMC and MLCV when $\varepsilon = 0.001$. MLCV sample size includes the number required to determine $\hat{Z}_\ell$.

Table 2: Cost of MC, MLMC, and MLCV for Test 2 problem.

ε	Levels	Cost MC	Cost MLMC	Cost MLCV	$\frac{\text{Cost MLCV}}{\text{Cost MLMC}}$
0.00005	0, 1, 2, 3	2.2e9	3.7e7	1.7e7	0.45
0.0005	0, 1, 2, 3	2.2e7	3.7e5	1.7e5	0.45
0.001	0, 1, 2, 3	5.6e6	9.3e4	4.4e4	0.48
0.003	0, 1, 2, 3	6.2e5	1.0e4	6.8e3	0.65
0.005	0, 1, 2, 3	2.2e5	3.8e3	4.0e3	1.1

4.3. Analysis

The preceding results show a cost improvement over MLMC when applying this MLCV method. In addition, other observations regarding these results are important to highlight. In Test 1, the difference between the number of degrees of freedom on adjacent levels has a greater negative impact on the MLMC results, than those of MLCV. When using MLCV, the process does not need to be broken up into as many levels with small changes in the degrees of freedom as in MLMC, as MLCV does not rely solely on the difference between QoIs on adjacent levels. This method is especially beneficial in cases where there is not sufficient flexibility in defining multiple coarse spatial discretizations, or when the solution converges slowly over the refinements.

The cost ratio between MLCV and MLMC is also significantly different when comparing Test 1 and Test 2. While the ρ_ℓ^2 results for both methods are favorable, the faster cost growth between levels in Test 2 results in greater cost savings. More precisely, in Test 1 we had $\gamma = 1$ and in Test 2 we had $\gamma = 1.65$. The payoff of placing more computational burden on the coarse level data than that of the fine level is more exaggerated for larger values of γ . This also suggests that MLCV will be even more advantageous over MLMC for three-dimensional (in physical space) problems, where M_ℓ grows faster.

Figure 10: (Left) MSE for $\hat{Y}_\ell$ and $\hat{W}_\ell$. (Right) Comparison of the relative errors (31) and (32) in estimating the mean of QoI at each level for MLMC and MLCV.

5. Conclusion

In this paper we have developed a non-intrusive, multilevel Monte Carlo method to approach uncertainty quantification problems with a large number of random inputs. The mathematical framework of this method relies heavily on the decomposition of the coarse level QoI data into a low-rank representation, from which a relatively accurate fine grid estimator can be calculated. Data with fast decaying singular values, combined with the use of control variates to minimize the MSE of the new variable results in this new MLCV method, with a total cost that is smaller than or comparable with that of MLMC.

Both MLCV and MLMC methods are applied to a generalized eigenvalue problem associated with a linear elasticity model as well as a thermally driven cavity flow problem in a square domain. While the data in both tests lends itself to approximations with different ranks, and costs to determine the QoIs, MLCV was found to outperform MLMC. To reach a fixed MSE tolerance, MLCV was determined to require fewer samples, and thus a cheaper implementation than that of MLMC.

Acknowledgements

This work was supported by the United States Department of Energy under the Predictive Science Academic Alliance Program 2 (PSAAP2) at Stanford University and NSF grant CMMI-145460. This material is based upon work supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0006402.

Appendix A

The following lemma from [30] outlines a number of key properties of the ID representation of $\mathbf{U}_{\ell-1}$, as seen in (15). For simplicity, we drop the subscripts of $\mathbf{U}$, $\mathbf{C}$, $\mathbf{U}^c$, m , and N .

Lemma 5.1. (Lemma 3.1 of [30].) *For any positive integer r with $r \leq \min\{m, N\}$, there exist a real $r \times N$ matrix $\mathbf{C}$, and a real $m \times r$ matrix $\mathbf{U}^c$ whose columns constitute a subset of the columns of $\mathbf{U}$, such that*

1. *some subset of the columns of $\mathbf{C}$ makes up the $r \times r$ identity matrix,*
2. *no entry of $\mathbf{C}$ has an absolute value greater than 1,*
3. $\|\mathbf{C}\| \leq \sqrt{r(N-r)+1},$

4. the least (that is, the r th greatest) singular value of $\mathbf{C}$ is at least 1,
5. $\mathbf{U} = \mathbf{U}^c \mathbf{C}$, when $r = m$ or $r = N$, and
6. $\|\mathbf{U} - \mathbf{U}^c \mathbf{C}\| \leq \sqrt{r(N-r)+1} \sigma_{r+1}$ when $r < \min\{m, N\}$, where σ_{r+1} is the $(r+1)$ st greatest singular value of $\mathbf{U}$.

The construction of matrix ID relies primarily on the rank-revealing QR factorization of $\mathbf{U}$ given by

$$\mathbf{U}\mathbf{P} \approx \mathbf{Q} [\mathbf{R}_{11} \mid \mathbf{R}_{12}], \quad (38)$$

where $\mathbf{P}$ is a $N \times N$ permutation matrix, $\mathbf{Q}$ an $m \times r$ matrix with orthonormal columns, $\mathbf{R}_{11}$ an $m \times r$ upper triangular matrix, and $\mathbf{R}_{12}$ an $r \times (N-r)$ matrix. In practice, the rank r is unknown and thus the pivoted Gram-Schmidt process involved in (38) is continued until $\|\mathbf{U} - \mathbf{U}^c \mathbf{C}\| \leq \epsilon$ is achieved for a predefined accuracy ϵ . Given (38), an $r \times (N-r)$ matrix $\mathbf{T}$ is sought for such that

$$\mathbf{R}_{11}\mathbf{T} = \mathbf{R}_{12}. \quad (39)$$

When $\mathbf{R}_{11}$ is ill-conditioned, [18] suggests a solution $\mathbf{T}$ with minimum $\|\mathbf{T}\|_F$. Using (39) in (38) we arrive at

$$\mathbf{U}\mathbf{P} \approx \mathbf{Q}\mathbf{R}_{11} [\mathbf{I} \mid \mathbf{T}], \quad (40)$$

or, equivalently,

$$\mathbf{U} \approx \mathbf{U}^c [\mathbf{I} \mid \mathbf{T}]\mathbf{P}^T = \mathbf{U}^c \mathbf{C}, \quad (41)$$

where $\mathbf{U}^c$ contains the first r columns of $\mathbf{U}\mathbf{P}$ and $\mathbf{C} = [\mathbf{I} \mid \mathbf{T}]\mathbf{P}^T$.

References

- [1] R. Ghanem, P. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Dover, 2002.
- [2] D. Xiu, G. Karniadakis, The Wiener-Askey polynomial chaos for stochastic differential equations, *SIAM Journal on Scientific Computing* 24 (2) (2002) 619–644.
- [3] A. Doostan, H. Owhadi, A non-adapted sparse approximation of PDEs with stochastic inputs, *Journal of Computational Physics* 230 (2011) 3015–3034.
- [4] D. Xiu, Y. Kevrekidis, R. Ghanem, An equation-free, multiscale approach to uncertainty quantification, *Computing in Science and Engineering* 7 (3) (2005) 16–23.
- [5] L. Mathelin, M. Hussaini, A stochastic collocation algorithm for uncertainty analysis, Tech. Rep. NAS 1.26:212153; NASA/CR-2003-212153, NASA Langley Research Center (2003).
- [6] K. Cliffe, M. Giles, R. Scheichl, A. Teckentrup, Multilevel monte carlo methods and applications to elliptic pdes with random coefficients, *Computing and Visualization in Science* 14 (1) (2011) 3–15.
- [7] S. Asmussen, P. Glynn, *Stochastic simulation: algorithms and analysis*, Vol. 57, Springer Science & Business Media, 2007.
- [8] S. Heinrich, Multilevel monte carlo methods, in: *Large-scale scientific computing*, Springer, 2001, pp. 58–67.
- [9] M. Giles, Multilevel monte carlo path simulation, *Operations Research* 56 (3) (2008) 607–617.
- [10] F. Nobile, F. Tesei, A multi level monte carlo method with control variate for elliptic pdes with log-normal coefficients, *Stochastic Partial Differential Equations: Analysis and Computations* 3 (3) (2015) 398–444.
- [11] A. Speight, A multilevel approach to control variates, *Journal of Computational Finance*, Forthcoming.
- [12] F. Vidal-Codina, N. Nguyen, M. Giles, J. Peraire, A model and variance reduction method for computing statistical outputs of stochastic elliptic partial differential equations, *Journal of Computational Physics* 297 (2015) 700–720.
- [13] G. Rozza, D. Huynh, A. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Archives of Computational Methods in Engineering* 15 (3) (2008) 229–275.
- [14] A. Narayan, C. Gittelsohn, D. Xiu, A stochastic collocation algorithm with multifidelity models, *SIAM Journal on Scientific Computing* 36 (2) (2014) A495–A521.
- [15] X. Zhu, A. Narayan, D. Xiu, Computational aspects of stochastic collocation with multifidelity models, *SIAM/ASA Journal on Uncertainty Quantification* 2 (1) (2014) 444–463.
- [16] A. Doostan, R. Ghanem, J. Red-Horse, Stochastic model reduction for chaos representations, *Computer Methods in Applied Mechanics and Engineering* 196 (37–40) (2007) 3951–3966.
- [17] A. Doostan, G. Geraci, G. Iaccarino, A bi-fidelity approach for uncertainty quantification of heat transfer in a rectangular ribbed channel, in: *ASME Turbo Expo 2016: Turbomachinery Technical Conference and Exposition*, American Society of Mechanical Engineers, 2016, pp. V02CT45A031–V02CT45A031.
- [18] H. Cheng, Z. Gimbutas, P. Martinsson, V. Rokhlin, On the compression of low rank matrices, *SIAM Journal on Scientific Computing* 26 (4) (2005) 1389–1404.
- [19] R. Pasupathy, B. Schmeiser, M. T. J. Wang, Control-variate estimation using estimated control means, *IIE Transactions* 44 (5) (2012) 381–385.
- [20] M. Giles, Multilevel monte carlo methods, in: *Monte Carlo and Quasi-Monte Carlo Methods 2012*, Springer, 2013, pp. 83–103.
- [21] A. Teckentrup, R. Scheichl, M. Giles, E. Ullmann, Further analysis of multilevel monte carlo methods for elliptic pdes with random coefficients, *Numerische Mathematik* 125 (3) (2013) 569–600.
- [22] R. Butler, T. Dodwell, T. Kim, S. Kynaston, R. Scheichl, R. Haftka, N. Kim, Uncertainty quantification of composite structures with defects using multilevel monte carlo simulations, in: *AIAA SciTech Conf*, 2015.
- [23] E. Elhamifar, R. Vidal, Sparse subspace clustering, in: *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, IEEE, 2009, pp. 2790–2797.
- [24] M. Mahoney, P. Drineas, Cur matrix decompositions for improved data analysis, *Proceedings of the National Academy of Sciences* 106 (2009) 697–702.
- [25] N. Halko, P.-G. Martinsson, J. Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review* 53 (2) (2011) 217–288.
- [26] E. Dyer, T. Goldstein, R. Patel, K. Kording, R. Baraniuk, Self-expressive decompositions for matrix approximation and clustering, *arXiv preprint arXiv:1505.00824*.
- [27] M. Gu, S. Eisenstat, Efficient algorithms for computing a strong rank-revealing qr factorization, *SIAM Journal on Scientific Computing* 17 (4) (1996) 848–869.
- [28] B. Schmeiser, M. Taaffe, J. Wang, Biased control-variate estimation, *IIE Transactions* 33 (3) (2001) 219–228.
- [29] M. Emsermann, B. Simon, Improving simulation efficiency with quasi control variates.
- [30] P. Martinsson, V. Rokhlin, M. Tygert, A randomized algorithm for the decomposition of matrices, *Applied and Computational Harmonic Analysis* 30 (1) (2011) 47–68.
- [31] O. L. Maitre, O. Knio, *Spectral Methods for Uncertainty Quantification with Applications to Computational Fluid Dynamics*, Springer, 2010.
- [32] O. Rubio, E. Bravo, J. Claeys, Thermally driven cavity flow with neumann condition for the pressure, *Applied numerical mathematics* 40 (1) (2002) 327–336.
- [33] J. Peng, J. Hampton, A. Doostan, A weighted ℓ_1 -minimization approach for sparse polynomial chaos expansions, *Journal of Computational Physics* 267 (2014) 92 – 111.
- [34] J. Hampton, A. Doostan, Compressive sampling of polynomial chaos expansions: Convergence analysis and sampling strategies, *Journal of Computational Physics* 280 (2015) 363 – 386.

- [35] A. Logg, K.-A. Mardal, G. N. Wells, et al., Automated Solution of Differential Equations by the Finite Element Method, Springer, 2012. [doi:10.1007/978-3-642-23099-8](https://doi.org/10.1007/978-3-642-23099-8).
- [36] P. L. Quéré, Accurate solutions to the square thermally driven cavity at high rayleigh number, Computers & Fluids 20 (1) (1991) 29–41.
- [37] O. LeMaitre, M. Reagan, H. Najm, R. Ghanem, O. Knio, A stochastic projection method for fluid flow. ii: Random process, J. Comp. Phys. 181 (2002) 9–44.