

Combining Sparse Grids, Multilevel MC and QMC for Elliptic PDEs with Random Coefficients

Michael B. Giles, Frances Y. Kuo, and Ian H. Sloan

Abstract Building on previous research which generalized multilevel Monte Carlo methods using either sparse grids or Quasi-Monte Carlo methods, this paper considers the combination of all these ideas applied to elliptic PDEs with finite-dimensional uncertainty in the coefficients. It shows the potential for the computational cost to achieve an $O(\varepsilon)$ r.m.s. accuracy to be $O(\varepsilon^{-r})$ with $r < 2$, independently of the spatial dimension of the PDE.

1 Introduction

There has been considerable research in recent years into the estimation of the expected value of output functionals $P(u)$ arising from the solution of elliptic PDEs of the form

$$-\nabla \cdot (a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y})) = f(\mathbf{x}), \quad (1)$$

in the unit hypercube $[0, 1]^d$, with homogeneous Dirichlet boundary conditions. Here \mathbf{x} represents the d -dimensional spatial coordinates and the gradients are with respect to these, while \mathbf{y} represents the uncertainty. In this paper we will consider the simplest possible setting in which we have finite s -dimensional uncertainty where

Michael B. Giles (✉)

Mathematical Institute, University of Oxford, Woodstock Road, Oxford, OX2 6GG, United Kingdom
e-mail: mike.giles@maths.ox.ac.uk

Frances Y. Kuo

School of Mathematics and Statistics, University of New South Wales, Sydney NSW 2052, Australia
e-mail: f.kuo@unsw.edu.au

Ian H. Sloan

School of Mathematics and Statistics, University of New South Wales, Sydney NSW 2052, Australia
e-mail: i.sloan@unsw.edu.au

$$a(\mathbf{x}, \mathbf{y}) = a_0(\mathbf{x}) + \sum_{j=1}^s y_j a_j(\mathbf{x}),$$

with the y_j independently and uniformly distributed on the interval $[-\frac{1}{2}, \frac{1}{2}]$, with $0 < a_{\min} \leq a(\mathbf{x}, \mathbf{y}) \leq a_{\max} < \infty$ for all \mathbf{x} and \mathbf{y} . This is the so-called “uniform case”.

In this paper we consider several grid-based sampling methods, in all of which the PDE (1) is solved approximately by full or sparse grid-based methods with respect to \mathbf{x} , for selected values of \mathbf{y} . We will consider both multilevel and multi-index methods [10, 15], and compare Monte Carlo (MC) and Quasi-Monte Carlo (QMC) methods for computing expected values with respect to \mathbf{y} . We pay attention to the dependence of the computational cost on the spatial dimension d , and we assume throughout this paper that the stochastic dimension s is fixed, though possibly large, and we do not track the dependence of the cost on s .

As a general approach in a wide range of stochastic applications, the multilevel Monte Carlo (MLMC) approach [10] computes solutions with different levels of accuracy, using the coarser solutions as a control variate for finer solutions. If the spatial dimension d is not too large, this can lead to an r.m.s. accuracy of ε being achieved at a computational cost which is $O(\varepsilon^{-2})$, which is much better than when using the standard MC method.

The earliest multilevel research on this problem was on the use of the MLMC method for both this “uniform case” [1, 17] and the harder “lognormal case” [5, 6, 18, 25] in which $a(\mathbf{x}, \mathbf{y})$ has a log-normal distribution with a specified spatial covariance so that $\log a(\mathbf{x}, \mathbf{y})$ has a Karhunen-Loève expansion of the form $\log a(\mathbf{x}, \mathbf{y}) = \kappa_0(\mathbf{x}) + \sum_{j=1}^{\infty} y_j \sqrt{\lambda_j} \kappa_j(\mathbf{x})$, where the y_j are independent with a standard normal distribution, and λ_j and $\kappa_j(\mathbf{x})$ are the non-decreasing eigenvalues and orthonormal eigenfunctions of integral operator involving the covariance kernel. For simplicity we will restrict our discussions to the uniform case in this paper, but our results can be easily adapted for the lognormal case.

Subsequent research [7, 13, 19, 21, 22] combined the multilevel approach with the use of QMC points, to form multilevel Quasi-Monte Carlo (MLQMC). In the best cases, this can further reduce the computational cost to $O(\varepsilon^{-r})$ for $r < 2$.

The efficiency of both MLMC and MLQMC suffers when d is large, and the reason for this is easily understood. Suppose the numerical discretisation of the PDE has order of accuracy p , so that the error in the output functional is $O(h^p)$, where h is the grid spacing in each coordinate direction. To achieve an $O(\varepsilon)$ accuracy requires $h = O(\varepsilon^{1/p})$, but if this is the grid spacing in each direction then the total number of grid points is $O(\varepsilon^{-d/p})$. Hence, the computational cost of performing just one calculation on the finest level of resolution is $O(\varepsilon^{-d/p})$, and this then gives a lower bound on the cost of the MLMC and MLQMC methods.

This curse of dimensionality is well understood, and in the case of deterministic PDEs (i.e., without the uncertainty \mathbf{y}) it has been addressed through the development of sparse grid methods [4]. One variant of this, the sparse combination technique, was the inspiration for the development of the multi-index Monte Carlo (MIMC) method [15]. The latter is a generalization of MLMC which in the context of multi-

dimensional PDEs uses a variety of regular grids, with differing resolutions in each spatial direction.

In this paper we have two main contributions:

- we present alternative ways of combining MLMC with sparse grids, and discuss their relationship to the MIMC method;
- we extend these approaches by considering the use of randomised QMC points, and derive the resulting computational cost if certain conditions are met.

The paper begins by reviewing sparse grid, MLMC/MIMC and randomised QMC methods [4, 8, 11]. Next we consider the combination of MLMC with sparse grids, before adding randomised QMC to the combination. In doing so, we present meta-theorems on the resulting computational cost, based on key assumptions about the asymptotic behaviour of certain quantities.

2 Sparse Grid Methods

There are two main classes of sparse grid methods for deterministic PDEs: sparse finite elements and the sparse combination technique [4].

2.1 Sparse Finite Element Method

The sparse finite element method for elliptic PDEs uses a standard Galerkin finite element formulation but with a sparse finite element basis. One advantage of this approach is that most of the usual finite element numerical analysis remains valid; the accuracy of the method can be bounded by using bounds on the accuracy in interpolating the exact solution using the sparse finite element basis functions. The main disadvantage of the approach compared to the sparse combination technique (see the next subsection) is the difficulty of its implementation.

Following the very clear description of the method in [3], suppose that we are interested in approximating the solution of an elliptic PDE in d -dimensions. For a non-negative multi-index $\underline{\ell} = (\ell_1, \ell_2, \dots, \ell_d)$, let $\mathcal{V}_{\underline{\ell}}$ be the finite element space spanned by the usual d -linear hat functions on a grid with spacing $2^{-\ell_j}$ in dimension j for each $j = 1, \dots, d$. The difference space $\mathcal{W}_{\underline{\ell}}$ is defined by

$$\mathcal{W}_{\underline{\ell}} = \mathcal{V}_{\underline{\ell}} \ominus \left(\bigoplus_{j=1}^d \mathcal{V}_{\underline{\ell} - \underline{e}_j} \right)$$

where \underline{e}_j is the unit vector in direction j . Thus, $\mathcal{W}_{\underline{\ell}}$ has the minimal set of additional basis elements such that

$$\mathcal{V}_{\underline{\ell}} = \mathcal{W}_{\underline{\ell}} \oplus \left(\bigoplus_{j=1}^d \mathcal{V}_{\underline{\ell} - \mathbf{e}_j} \right).$$

A sparse finite element space is then defined by $\bigoplus_{\underline{\ell} \in \mathcal{L}} \mathcal{W}_{\underline{\ell}}$, for some index set \mathcal{L} . A simple and near-optimal choice for a given level of accuracy is the set $\mathcal{L} = \{\underline{\ell} : \|\underline{\ell}\|_1 \leq L\}$ for some integer L ; this is discussed in [3] (that paper also presents a slightly better choice). Having defined the finite element space used for both test and trial functions, the rest of the formulation is the standard Galerkin finite element method. In the following, the space \mathcal{H}_1 is the standard Sobolev space with mixed first derivatives in \mathbf{x} .

Theorem 1 (Sparse finite element method). *For fixed \mathbf{y} , if the PDE (1) is solved using the sparse finite element method with the index set specified by $\|\underline{\ell}\|_1 \leq L$, then the computational cost is $O(L^{d-1} 2^L)$. Moreover, the \mathcal{H}_1 solution accuracy is $O(2^{-L})$ if the solution u has sufficient mixed regularity, and the accuracy of simple output functionals P (such as smoothly weighted averages of the solution) is $O(2^{-2L})$. Hence, the cost to achieve a functional accuracy of ε is $O(\varepsilon^{-1/2} |\log \varepsilon|^{d-1})$.*

Proof. The cost and \mathcal{H}_1 solution accuracy are proved in [3, 14]. The super-convergence for output functionals is an immediate consequence of adjoint-based error analysis [12]. \square

2.2 Sparse Combination Method

The sparse combination method combines the results of separate calculations on simple tensor product grids with different resolutions in each coordinate direction [14]. For a given output functional P and multi-index $\underline{\ell} = (\ell_1, \dots, \ell_d)$, let $P_{\underline{\ell}}$ denote the approximate output functional obtained on a grid with spacing $2^{-\ell_j}$ in direction j for each $j = 1, \dots, d$. For convenience, we define $P_{\underline{\ell}} := 0$ if any of the indices in $\underline{\ell}$ is negative.

The backward difference in the j^{th} dimension is defined as $\Delta_j P_{\underline{\ell}} := P_{\underline{\ell}} - P_{\underline{\ell} - \mathbf{e}_j}$, and we define the d -dimensional mixed first difference as

$$\Delta P_{\underline{\ell}} := \left(\prod_{j=1}^d \Delta_j \right) P_{\underline{\ell}}.$$

For an arbitrary multi-index $\underline{\ell}'$, it can be shown that

$$P_{\underline{\ell}'} = \sum_{0 \leq \underline{\ell} \leq \underline{\ell}'} \Delta P_{\underline{\ell}}, \tag{2}$$

where the multi-index inequality $\underline{\ell} \leq \underline{\ell}'$ is applied element-wise (i.e. $\ell_j \leq \ell'_j, \forall j$). Taking the limit as $\underline{\ell}' \rightarrow \infty$ (i.e. $\ell'_j \rightarrow \infty, \forall j$) gives

$$P = \sum_{\underline{\ell} \geq 0} \Delta P_{\underline{\ell}}. \quad (3)$$

The sparse combination method truncates the summation to a finite index set, with a simple and near-optimal choice again being $\|\underline{\ell}\|_1 \leq \ell$. This gives the approximation

$$P_\ell := \sum_{\|\underline{\ell}\|_1 \leq \ell} \Delta P_{\underline{\ell}}, \quad (4)$$

where we are slightly abusing notation by distinguishing between the original $P_{\underline{\ell}}$ with a multi-index subscript (in bold type with a tilde underneath), and the new P_ℓ on the left-hand side of this equation with a scalar subscript (which is not in bold).

If we now define

$$S_\ell := \sum_{\|\underline{\ell}\|_1 = \ell} P_{\underline{\ell}} \quad (5)$$

and the backward difference $\Delta S_\ell := S_\ell - S_{\ell-1}$, then it can be shown [24] that

$$P_\ell = \Delta^{d-1} S_\ell = \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} S_{\ell-k}.$$

Hence, the computation of P_ℓ requires $O(\ell^{d-1})$ separate computations, each on a grid with $O(2^\ell)$ grid points. This leads to the following theorem.

Theorem 2 (Sparse combination method). *For fixed \mathbf{y} , if the PDE (1) is solved using the sparse combination method with the index set specified by $\|\underline{\ell}\|_1 \leq L$, then the computational cost is $O(L^{d-1} 2^L)$. Moreover, if the underlying PDE approximation has second order accuracy and the solution u has sufficient mixed regularity, then $\Delta P_{\underline{\ell}}$ has magnitude $O(2^{-2\|\underline{\ell}\|_1})$ so the error in P_L is $O(L^{d-1} 2^{-2L})$. Hence, the cost to achieve a functional accuracy of ε is $O(\varepsilon^{-1/2} |\log \varepsilon|^{3(d-1)/2})$.*

Proof. For the results on the cost and accuracy see [24]. The cost result is an immediate consequence. \square

3 MLMC and MIMC

3.1 MLMC

The multilevel Monte Carlo (MLMC) idea is very simple. As explained in a recent review article [11], given a sequence P_ℓ , $\ell=0, 1, \dots$ of approximations of an output functional P , with increasing accuracy and cost as ℓ increases, and defining $P_{-1} := 0$, we have the simple identity

$$\mathbb{E}[P] = \sum_{\ell=0}^{\infty} \mathbb{E}[\Delta P_{\ell}], \quad \Delta P_{\ell} := P_{\ell} - P_{\ell-1}.$$

The summation can be truncated to

$$\mathbb{E}[P] \approx \mathbb{E}[P_L] = \sum_{\ell=0}^L \mathbb{E}[\Delta P_{\ell}], \quad (6)$$

with L chosen to be sufficiently large to ensure that the weak error $\mathbb{E}[P - P_L]$ is acceptably small. Each of the expectations on the r.h.s. of (6) can be estimated independently using N_{ℓ} independent samples so that the MLMC estimator is

$$Y = \sum_{\ell=0}^L Y_{\ell}, \quad Y_{\ell} = \frac{1}{N_{\ell}} \sum_{i=1}^{N_{\ell}} \Delta P_{\ell}^{(i)}. \quad (7)$$

The computational savings comes from the fact that on the finer levels ΔP_{ℓ} is smaller and has a smaller variance, and therefore fewer samples N_{ℓ} are required to accurately estimate its expected value.

The optimal value for N_{ℓ} on level $\ell = 0, 1, \dots, L$ can be estimated by approximately minimising the cost for a given overall variance. This results in the following theorem which is a slight generalization of the original in [10].

Theorem 3 (MLMC). *Let P denote an output functional, and let P_{ℓ} denote the corresponding level ℓ numerical approximation. Suppose there exist independent estimators Y_{ℓ} of $\mathbb{E}[\Delta P_{\ell}]$ based on N_{ℓ} Monte Carlo samples and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3$, with $\alpha \geq \frac{1}{2} \min(\beta, \gamma)$, such that*

- i) $\left| \mathbb{E}[P_{\ell} - P] \right| \rightarrow 0$ as $\ell \rightarrow \infty$,
- ii) $|\mathbb{E}[\Delta P_{\ell}]| \leq c_1 2^{-\alpha \ell}$,
- iii) $\mathbb{E}[Y_{\ell}] = \mathbb{E}[\Delta P_{\ell}]$,
- iv) $\mathbb{V}[Y_{\ell}] \leq c_2 N_{\ell}^{-1} 2^{-\beta \ell}$,
- v) $\text{cost}(Y_{\ell}) \leq c_3 N_{\ell} 2^{\gamma \ell}$.

Then there exists a positive constant c_4 such that for any $\varepsilon < e^{-1}$ there are values L and N_{ℓ} for which the MLMC estimator (7) achieves the mean-square-error bound $\mathbb{E}[(Y - \mathbb{E}[P])^2] < \varepsilon^2$ with the computational cost bound

$$\text{cost}(Y) \leq \begin{cases} c_4 \varepsilon^{-2}, & \beta > \gamma, \\ c_4 \varepsilon^{-2} |\log \varepsilon|^2, & \beta = \gamma, \\ c_4 \varepsilon^{-2 - (\gamma - \beta)/\alpha}, & \beta < \gamma. \end{cases}$$

The proof of this theorem uses a constrained optimisation approach to optimise the number of samples N_{ℓ} on each level. This treats the N_{ℓ} as real variables, and then the optimal value is rounded up to the nearest integer. This rounding up improves

the variance slightly, so that we still achieve our target mean-square-error accuracy, but it also increases the cost by at most one sample per level. This additional cost is dominated by the cost of one sample on the finest level, which is $O(\varepsilon^{-\gamma/\alpha})$ since the weak convergence condition requires that the finest level satisfies $2^{-\alpha L} = O(\varepsilon)$. The condition in the theorem that $\alpha \geq \frac{1}{2} \min(\beta, \gamma)$ ensures that this additional cost is negligible compared to the main cost.

When applied to our model elliptic PDE, if one uses a tensor product grid with spacing $2^{-\ell}$ in each direction, then if the numerical discretisation has second order accuracy it gives $\alpha=2$ and $\beta=4$, while with an ideal multigrid solver the cost is at best proportional to the number of grid points which is $2^{d\ell}$ so $\gamma=d$. Hence, the cost is $O(\varepsilon^{-r})$ where $r = \max(2, d/2)$, except for $d=4$ for which $\beta=\gamma$ and hence there is an additional $|\log \varepsilon|^2$ factor. It is the dependence on d which will be addressed by incorporating sparse grid methods.

3.2 MIMC

The multi-index Monte Carlo (MIMC) method [15] is inspired by the sparse combination technique. Starting from (3), if each of the $\Delta P_{\underline{\ell}}$ is now a random variable due to the random coefficients in the PDE, we can take expectations of each side and truncate the sum to give

$$\mathbb{E}[P] \approx \mathbb{E}[P_L] = \sum_{\|\underline{\ell}\|_1 \leq L} \mathbb{E}[\Delta P_{\underline{\ell}}]. \quad (8)$$

This is now very similar to the telescoping sum (6) in MLMC, with the difference that the levels are now labelled by multi-indices, so allowing different discretizations in different directions. We can independently estimate each of the expectations on the r.h.s. of (8) using a number of independent samples $N_{\underline{\ell}}$ so that the MIMC estimator is

$$Y = \sum_{\|\underline{\ell}\|_1 \leq L} Y_{\underline{\ell}}, \quad Y_{\underline{\ell}} = \frac{1}{N_{\underline{\ell}}} \sum_{i=1}^{N_{\underline{\ell}}} \Delta P_{\underline{\ell}}^{(i)}. \quad (9)$$

The numbers $N_{\underline{\ell}}$ are optimised to minimise the cost of achieving a certain desired variance or mean-square-error.

The original paper [15] considers much more general circumstances: the different indices in $\underline{\ell}$ are not limited to the spatial discretizations in \mathbf{x} but can also involve quantities such as the number of particles in a system, or the number of terms in a Karhunen–Loève expansion (arising from dimension truncation in the stochastic variables \mathbf{y}). Here in the isotropic PDE case, in which the behaviour in each space dimension is similar, this leads to the following theorem.

Theorem 4 (MIMC). *Let P denote an output functional, and for each multi-index $\underline{\ell}$ let $P_{\underline{\ell}}$ denote the approximate output functional indexed by $\underline{\ell}$. Suppose for each*

multi-index $\underline{\ell}$ there exist independent estimators $Y_{\underline{\ell}}$ of $\mathbb{E}[\Delta P_{\underline{\ell}}]$ based on $N_{\underline{\ell}}$ Monte Carlo samples and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3$, with $\alpha \geq \frac{1}{2}\beta$, such that

- i) $\left| \mathbb{E}[P_{\underline{\ell}} - P] \right| \rightarrow 0$ as $\underline{\ell} \rightarrow \infty$ ($\ell_j \rightarrow \infty, \forall j$),
- ii) $\left| \mathbb{E}[\Delta P_{\underline{\ell}}] \right| \leq c_1 2^{-\alpha \|\underline{\ell}\|_1}$,
- iii) $\mathbb{E}[Y_{\underline{\ell}}] = \mathbb{E}[\Delta P_{\underline{\ell}}]$,
- iv) $\mathbb{V}[Y_{\underline{\ell}}] \leq c_2 N_{\underline{\ell}}^{-1} 2^{-\beta \|\underline{\ell}\|_1}$,
- v) $\text{cost}(Y_{\underline{\ell}}) \leq c_3 N_{\underline{\ell}} 2^{\gamma \|\underline{\ell}\|_1}$.

Then there exists a positive constant c_4 such that for any $\varepsilon < e^{-1}$ there are values L and $N_{\underline{\ell}}$ for which the MIMC estimator (9) achieves the mean-square-error bound $\mathbb{E}[(Y - \mathbb{E}[P])^2] < \varepsilon^2$ with the computational cost bound

$$\text{cost}(Y) \leq \begin{cases} c_4 \varepsilon^{-2}, & \beta > \gamma, \\ c_4 \varepsilon^{-2} |\log \varepsilon|^{e_1}, & \beta = \gamma, \\ c_4 \varepsilon^{-2-(\gamma-\beta)/\alpha} |\log \varepsilon|^{e_2}, & \beta < \gamma, \end{cases}$$

where

$$\begin{aligned} e_1 &= 2d, & e_2 &= (d-1)(2+(\gamma-\beta)/\alpha), \text{ if } \alpha > \frac{1}{2}\beta, \\ e_1 &= \max(2d, 3(d-1)), & e_2 &= (d-1)(1+\gamma/\alpha), \text{ if } \alpha = \frac{1}{2}\beta. \end{aligned}$$

Proof. This is a particular case of the more general analysis in [15, Theorem 2.2]. \square

In the case of MIMC, there are $O(L^{d-1})$ multi-indices on the finest level on which $\|\underline{\ell}\|_1 = L$. Hence the finest level is determined by the constraint $L^{d-1} 2^{-\alpha L} = O(\varepsilon)$, and the associated cost is $O(\varepsilon^{-\gamma/\alpha} |\log \varepsilon|^{(d-1)(1+\gamma/\alpha)})$. Given the assumption that $\alpha \geq \frac{1}{2}\beta$, this is not asymptotically bigger than the main cost except when $\alpha = \frac{1}{2}\beta$, in which case it is responsible for the e_2 and the $3(d-1)$ component in the maximum in e_1 .

When applied to our model elliptic PDE, if one uses a tensor product grid with spacing $2^{-\ell_j}$ in the j^{th} direction, and a numerical discretisation with second order accuracy, then we are likely to get $\alpha=2$ and $\beta=4$ if the solution has sufficient mixed regularity [24]. (Note that this is a much stronger statement than the $\alpha=2$, $\beta=4$ in the previous section; taking the case with $d=3$ as an example, with grid spacing h_1, h_2, h_3 in the three dimensions, Section 3.1 requires only that $\Delta P_{\underline{\ell}} = O(h^2)$ when all three spacings are equal to h , whereas in this section we require the product form $\Delta P_{\underline{\ell}} = O(h_1^2 h_2^2 h_3^2)$ which is much smaller when $h_1, h_2, h_3 \ll 1$.) With an ideal multigrid solver, the cost is proportional to $2^{\|\underline{\ell}\|_1}$, so $\gamma=1$. Since $\beta > \gamma$, the cost would then be $O(\varepsilon^{-2})$, regardless of the value of d .

4 Randomised QMC and MLQMC

4.1 Randomised QMC Sampling

A randomized QMC method with N deterministic points and R randomization steps approximates an s -dimensional integral over the unit cube $[-\frac{1}{2}, \frac{1}{2}]^s$ as follows

$$I := \int_{[-\frac{1}{2}, \frac{1}{2}]^s} g(\mathbf{y}) d\mathbf{y} \approx \bar{Q} := \frac{1}{R} \sum_{k=1}^R Q_k, \quad Q_k = \frac{1}{N} \sum_{i=1}^N g(\mathbf{y}^{(i,k)}).$$

For the purpose of this paper it suffices that we introduce briefly just a simple family of randomized QMC methods – randomly shifted lattice rules. We have

$$\mathbf{y}^{(i,k)} = \left\{ \frac{i\mathbf{z}}{N} + \boldsymbol{\Delta}^{(k)} \right\} - \frac{1}{2},$$

where $\mathbf{z} \in \mathbb{N}^s$ is known as the generating vector; $\boldsymbol{\Delta}^{(1)}, \dots, \boldsymbol{\Delta}^{(R)} \in (0, 1)^s$ are R independent random shifts; the braces indicate that we take the fractional part of each component in the vector; and finally we subtract $\frac{1}{2}$ from each component of the vector to bring it into $[-\frac{1}{2}, \frac{1}{2}]^s$.

Randomly shifted lattice rules provide unbiased estimators of the integral. Indeed, it is easy to verify that $\mathbb{E}_{\boldsymbol{\Delta}}[\bar{Q}] = \mathbb{E}_{\boldsymbol{\Delta}}[Q_k] = I$, where we introduced the subscript $\boldsymbol{\Delta}$ to indicate that the expectation is taken with respect to the random shifts. In some appropriate function space setting for the integrand function g , it is known (see e.g., [8]) that good generating vectors \mathbf{z} can be constructed so that the variance or mean-square-error satisfies $\mathbb{V}_{\boldsymbol{\Delta}}[\bar{Q}] = \mathbb{E}_{\boldsymbol{\Delta}}[(\bar{Q} - I)^2] \leq C_{\delta} R^{-1} N^{-2(1-\delta)}$, for some $\delta \in (0, 1/2]$ with C_{δ} independent of the dimension s . In practical computations, we can estimate the variance by $\mathbb{V}_{\boldsymbol{\Delta}}[\bar{Q}] \approx \sum_{k=1}^R (Q_k - \bar{Q})^2 / [R(R-1)]$. Typically we take a large value of N to benefit from the higher QMC convergence rate and use only a relatively small R (e.g., 20–50) for the purpose of estimating the variance.

There are other randomization strategies for QMC methods. For example, we can combine any digital net such as Sobol' sequences or interlaced polynomial lattice rules with digital shift or Owen scrambling, to get an unbiased estimator with variance close to $O(N^{-2})$ or $O(N^{-3})$. We can also apply randomization to a higher order digital net to achieve $O(N^{-p})$ for $p > 2$ in an appropriate function space setting for smooth integrands. For detailed reviews of these results see e.g., [8].

4.2 MLQMC

As a generalization of (7), the multilevel Quasi-Monte Carlo (MLQMC) estimator is

$$Y = \sum_{\ell=0}^L Y_{\ell}, \quad Y_{\ell} = \frac{1}{R_{\ell}} \sum_{k=1}^{R_{\ell}} \left(\frac{1}{N_{\ell}} \sum_{i=1}^{N_{\ell}} \Delta P_{\ell}^{(i,k)} \right). \quad (10)$$

Later in Theorem 5 we will state the corresponding generalization of Theorem 3.

The use of QMC instead of MC in a multilevel method was first considered in [13] where numerical experiments were carried out for a number of option pricing problems and showed convincingly that MLQMC improves upon MLMC. A meta-theorem similar to the MLMC theorem was proved in [9]. A slightly sharper version of the theorem, eliminating some $\log(\varepsilon)$ factors, will be stated and proved later in §6.

MLQMC methods have been combined with finite element discretizations for the PDE problems in [7, 21, 22]. The paper [22] studied the uniform case for the same elliptic PDE of this paper with randomly shifted lattice rules (which yield up to order 2 convergence in the variance); the paper [7] studied the uniform case for general operator equations with deterministic higher order digital nets; the paper [21] studied the lognormal case with randomly shifted lattice rules. A key analysis which is common among these papers is the required mixed regularity estimate of the solution involving both \mathbf{x} and \mathbf{y} , see [20] for a survey of the required analysis in a unified framework.

5 Combining Sparse Grids and MLMC

After this survey of the three component technologies, sparse grid methods, MLMC and MIMC, and randomised QMC samples, the first novel observation in this paper is very simple: MIMC is not the only way in which MLMC can be combined with sparse grid methods.

An alternative is to use the standard MLMC approach, but with samples which are computed using sparse grid methods. The advantage of this is that it can be used with either sparse finite elements or the sparse combination technique.

5.1 MLMC with Sparse Finite Element Samples

In Theorem 3, if P_{ℓ} is computed using sparse finite elements as described in Section 2.1 based on grids with index set $\|\mathbf{\ell}\|_1 \leq \ell$, and if the accuracy and cost are as given in Theorem 1, then we obtain $\alpha = 2 - \delta$, $\beta = 4 - \delta$, and $\gamma = 1 + \delta$ for any $0 < \delta \ll 1$. Here δ arises due to the effect of some additional powers of ℓ . So $\beta > \gamma$ and therefore the computational cost is $O(\varepsilon^{-2})$.

Recall that with the full tensor product grid we had $\alpha = 2$, $\beta = 4$, and $\gamma = d$. Hence the improvement here is in the removal of the dependence of the cost parameter γ on d .

5.2 MLMC with Sparse Combination Samples

The aim in this section is to show that the MIMC algorithm is very similar to MLMC using sparse combination samples.

Suppose we have an MIMC application which satisfies the conditions of Theorem 4. For the MLMC version, we use (4) to define the P_ℓ in Theorem 3. Since

$$\mathbb{E}[\Delta P_\ell] = \sum_{\|\underline{\ell}\|_1 = \ell} \mathbb{E}[\Delta P_{\underline{\ell}}], \quad (11)$$

the two algorithms have exactly the same expected value if the finest level for each is given by $\|\underline{\ell}\|_1 = L$ for the same value of L . The difference between the two algorithms is that MIMC independently estimates each of the expectations on the r.h.s. of (11), using a separate estimator $Y_{\underline{\ell}}$ for each $\mathbb{E}[\Delta P_{\underline{\ell}}]$ with independent samples of \mathbf{y} , whereas MLMC with sparse combination samples estimates the expectation on the l.h.s., using the combination

$$Y_\ell = \sum_{\|\underline{\ell}\|_1 = \ell} Y_{\underline{\ell}},$$

with the $Y_{\underline{\ell}}$ all based on the same set of N_ℓ random samples \mathbf{y} .

There are no more than $(\ell+1)^{d-1}$ terms in the summation in (11), so if the cost of $Y_{\underline{\ell}}$ for MIMC is $O(N_\ell 2^{\gamma_\ell})$ when $\|\underline{\ell}\|_1 = \ell$, then the cost of the sparse combination estimator Y_ℓ for MLMC is $O(N_\ell \ell^{d-1} 2^{\gamma_\ell}) = o(N_\ell 2^{(\gamma+\delta)\ell})$, for any $0 < \delta \ll 1$.

Likewise,

$$|\mathbb{E}[Y_\ell]| \leq \sum_{\|\underline{\ell}\|_1 = \ell} |\mathbb{E}[Y_{\underline{\ell}}]|,$$

so if $|\mathbb{E}[Y_{\underline{\ell}}]| = O(2^{-\alpha_\ell})$ when $\|\underline{\ell}\|_1 = \ell$, then $|\mathbb{E}[Y_\ell]| = o(2^{-(\alpha-\delta)\ell})$ for any $0 < \delta \ll 1$.

Furthermore, Jensen's inequality gives

$$\begin{aligned} \mathbb{V}[Y_\ell] &= \mathbb{E}[(Y_\ell - \mathbb{E}[Y_\ell])^2] = \mathbb{E}\left[\left(\sum_{\|\underline{\ell}\|_1 = \ell} (Y_{\underline{\ell}} - \mathbb{E}[Y_{\underline{\ell}}])\right)^2\right] \\ &\leq (\ell+1)^{d-1} \sum_{\|\underline{\ell}\|_1 = \ell} \mathbb{E}[(Y_{\underline{\ell}} - \mathbb{E}[Y_{\underline{\ell}}])^2] \\ &= (\ell+1)^{d-1} \sum_{\|\underline{\ell}\|_1 = \ell} \mathbb{V}[Y_{\underline{\ell}}], \end{aligned}$$

so if $\mathbb{V}[Y_{\underline{\ell}}] = O(N_\ell^{-1} 2^{-\beta_\ell})$, then $\mathbb{V}[Y_\ell] = o(N_\ell^{-1} 2^{-(\beta-\delta)\ell})$, for any $0 < \delta \ll 1$.

This shows that the α , β , γ values for the MLMC algorithm using the sparse combination samples are almost equal to the α , β , γ for the MIMC method, which leads to the following lemma.

Lemma 1. *If a numerical method satisfies the conditions for the MIMC Theorem 4, then the corresponding MLMC estimator with sparse combination samples will have a cost which is $O(\varepsilon^{-2})$, if $\beta > \gamma$, and $o(\varepsilon^{-2-(\gamma-\beta)/\alpha-\delta})$, $\forall 0 < \delta \ll 1$, if $\beta \leq \gamma$.*

As with MLMC with sparse finite element samples, the key thing here is that the level ℓ MLMC samples use a set of grids in which the number of grid points is $O(2^{\|\underline{\ell}\|_1}) = O(2^\ell)$. That is why the γ values for MIMC and MLMC are virtually identical.

If there is substantial cancellation in the summation, it is possible that $\mathbb{V}[Y_\ell]$ could be very much smaller than the $\mathbb{V}[Y_\ell]$ for each of the $\underline{\ell}$ for which $\|\underline{\ell}\|_1 = \ell$. However, we conjecture that this is very unlikely, and therefore we are not suggesting that the MLMC with sparse combination samples is likely to be better than MIMC. The point of this section is to show that it cannot be significantly worse. In addition, this idea of combining MLMC with sparse grid samples works for sparse finite elements for which there seems to be no natural MIMC extension.

5.3 Nested MLMC

Another alternative to MIMC is nested MLMC. To illustrate this in 2D, suppose we start by using a single level index ℓ_1 to construct a standard MLMC decomposition

$$\mathbb{E}[P] \approx \mathbb{E}[P_{L_1}] = \sum_{\ell_1=0}^{L_1} \mathbb{E}[\Delta P_{\ell_1}].$$

Now, for each particular index ℓ_1 we can take $\mathbb{E}[\Delta P_{\ell_1}]$ and perform a secondary MLMC expansion with respect to a second index ℓ_2 to give

$$\mathbb{E}[\Delta P_{\ell_1}] \approx \sum_{\ell_2=0}^{L_2} \mathbb{E}[Q_{\ell_1, \ell_2} - Q_{\ell_1, \ell_2-1}],$$

with $Q_{\ell_1, -1} := 0$. If we allow L_2 to possibly depend on the value of ℓ_1 , this results in an approximation which is very similar to the MIMC method,

$$\mathbb{E}[P] \approx \sum_{\underline{\ell} \in \mathcal{L}} \mathbb{E}[Q_{\ell_1, \ell_2} - Q_{\ell_1, \ell_2-1}],$$

with the summation over some finite set of indices \mathcal{L} . In contrast to the MIMC method, here $Q_{\ell_1, \ell_2} - Q_{\ell_1, \ell_2-1}$ is not necessarily expressible in the cross-difference form $\Delta P_{\underline{\ell}}$ used in MIMC. Thus, this method is a generalization of MIMC.

This approach is currently being used in two new research projects. In one project, the second expansion is with respect to the precision of floating point computations; i.e. half, single or double precision. This follows ideas presented in section 10.2 of [11] and also in [2]. In the other project [16], the second expansion uses Rhee &

Glynn's randomised multilevel Monte Carlo method [23] to provide an unbiased inner estimate in a financial nested expectation application.

6 MLQMC and MIQMC

The next natural step is to replace the Monte Carlo sampling with randomised QMC sampling to estimate $\mathbb{E}[\Delta P_\ell]$ or $\mathbb{E}[\underline{\Delta P}_\ell]$.

6.1 MLQMC (continued from §4.2)

In the best circumstances, using N_ℓ QMC deterministic points with $R_\ell = R$ randomisation steps to estimate $\mathbb{E}[\Delta P_\ell]$ gives a variance (with respect to the randomisation in the QMC points) which is $O(R^{-1}N_\ell^{-p}2^{-\beta\ell})$, with $p > 1$. This leads to the following theorem which generalizes Theorem 3.

Theorem 5 (MLQMC). *Let P denote an output functional, and let P_ℓ denote the corresponding level ℓ numerical approximation. Suppose there exist independent estimators Y_ℓ of $\mathbb{E}[\Delta P_\ell]$ based on N_ℓ deterministic QMC points and $R_\ell = R$ randomization steps, and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3, p$, with $p > 1$ and $\alpha \geq \frac{1}{2}\beta$, such that*

- i) $\left| \mathbb{E}[P_\ell - P] \right| \longrightarrow 0$ as $\ell \longrightarrow \infty$,
- ii) $|\mathbb{E}[\Delta P_\ell]| \leq c_1 2^{-\alpha\ell}$,
- iii) $\mathbb{E}_\Delta[Y_\ell] = \mathbb{E}[\Delta P_\ell]$,
- iv) $\mathbb{V}_\Delta[Y_\ell] \leq c_2 R^{-1} N_\ell^{-p} 2^{-\beta\ell}$,
- v) $\text{cost}(Y_\ell) \leq c_3 R N_\ell 2^{\gamma\ell}$.

Then there exists a positive constant c_4 such that for any $\varepsilon < e^{-1}$ there are values L and N_ℓ for which the MLQMC estimator (10) achieves the mean-square-error bound $\mathbb{E}_\Delta[(Y - \mathbb{E}[P])^2] < \varepsilon^2$ with the computational cost bound

$$\text{cost}(Y) \leq \begin{cases} c_4 \varepsilon^{-2/p}, & \beta > p\gamma, \\ c_4 \varepsilon^{-2/p} |\log \varepsilon|^{(p+1)/p}, & \beta = p\gamma, \\ c_4 \varepsilon^{-2/p - (p\gamma - \beta)/(p\alpha)}, & \beta < p\gamma. \end{cases}$$

Proof. We omit the proof here because the theorem can be interpreted as a special case of Theorem 6 below for which we will provide an outline of the proof. \square

6.2 MIQMC

As a generalization of (9), the MIQMC estimator is

$$Y = \sum_{\|\underline{\ell}\|_1 \leq L} Y_{\underline{\ell}}, \quad Y_{\underline{\ell}} = \frac{1}{R_{\underline{\ell}}} \sum_{k=1}^{R_{\underline{\ell}}} \left(\frac{1}{N_{\underline{\ell}}} \sum_{i=1}^{N_{\underline{\ell}}} \Delta P_{\underline{\ell}}^{(i,k)} \right), \quad (12)$$

where $Y_{\underline{\ell}}$ is an estimator for $\mathbb{E}[\Delta P_{\underline{\ell}}]$ based on $N_{\underline{\ell}}$ deterministic QMC points and $R_{\underline{\ell}}$ randomization steps.

Suppose that $Y_{\underline{\ell}}$ has variance and cost given by $\mathbb{V}_{\Delta}[Y_{\underline{\ell}}] = N_{\underline{\ell}}^{-p} v_{\underline{\ell}}$ and $\text{cost}(Y_{\underline{\ell}}) = N_{\underline{\ell}} c_{\underline{\ell}}$. The variance and total cost of the combined estimator Y are

$$\mathbb{V}_{\Delta}[Y] = \sum_{\|\underline{\ell}\|_1 \leq L} N_{\underline{\ell}}^{-p} v_{\underline{\ell}}, \quad \text{cost}(Y) = \sum_{\|\underline{\ell}\|_1 \leq L} N_{\underline{\ell}} c_{\underline{\ell}}.$$

Treating the $N_{\underline{\ell}}$ as real numbers, the cost can be minimised for a given total variance by introducing a Lagrange multiplier and minimising $\text{cost}(Y) + \lambda \mathbb{V}_{\Delta}[Y]$, which gives

$$N_{\underline{\ell}} = \left(\frac{\lambda p v_{\underline{\ell}}}{c_{\underline{\ell}}} \right)^{1/(p+1)}.$$

Requiring $\mathbb{V}_{\Delta}[Y] = \frac{1}{2} \varepsilon^2$ to achieve a target accuracy determines the value of λ and then the total cost is

$$\text{cost}(Y) = (2 \varepsilon^{-2})^{1/p} \left(\sum_{\|\underline{\ell}\|_1 \leq L} (c_{\underline{\ell}}^p v_{\underline{\ell}})^{1/(p+1)} \right)^{(p+1)/p}.$$

This outline analysis shows that the behaviour of the product $c_{\underline{\ell}}^p v_{\underline{\ell}}$ as $\underline{\ell} \rightarrow \infty$ is critical. If $c_{\underline{\ell}} = O(2^{\gamma \ell})$ and $v_{\underline{\ell}} = O(2^{-\beta \ell})$ where $\ell = \|\underline{\ell}\|_1$, then $c_{\underline{\ell}}^p v_{\underline{\ell}} = O(2^{(p\gamma - \beta)\ell})$.

If $\beta > p\gamma$, then the total cost is dominated by the contributions from the coarsest levels, and we get a total cost which is $O(\varepsilon^{-2/p})$.

If $\beta = p\gamma$, then all levels contribute to the total cost, and it is $O(L^{d(p+1)/p} \varepsilon^{-2/p})$.

If $\beta < p\gamma$, then the total cost is dominated by the contributions from the finest levels, and we get a total cost which is $O(L^{(d-1)(p+1)/p} \varepsilon^{-2/p} 2^{(p\gamma - \beta)L/p})$.

To complete this analysis, we need to know the value of L which is determined by the requirement that the square of the bias is no more than $\frac{1}{2} \varepsilon^2$. This can be satisfied by ensuring that

$$\text{bias}(Y) := \sum_{\|\underline{\ell}\|_1 > L} \left| \mathbb{E}[\Delta P_{\underline{\ell}}] \right| \leq \varepsilon / \sqrt{2}.$$

If $|\mathbb{E}[\Delta P_{\underline{\ell}}]| = O(2^{-\alpha\|\underline{\ell}\|_1})$, then the contributions to $\text{bias}(Y)$ come predominantly from the coarsest of the levels in the summation (i.e. $\|\underline{\ell}\|_1 = L+1$), and hence $\text{bias}(Y) = O(L^{d-1}2^{-\alpha L})$. The bias constraint then gives $L^{d-1}2^{-\alpha L} = O(\varepsilon)$ and hence $L = O(|\log \varepsilon|)$.

As discussed after the MLMC and MIMC theorems, the values for $N_{\underline{\ell}}$ need to be rounded up to the nearest integers, incurring an additional cost which is $O(\varepsilon^{-\gamma/\alpha} |\log \varepsilon|^{(d-1)(1+\gamma/\alpha)})$. If $\alpha > \frac{1}{2}\beta$ it is always negligible compared to the main cost, but it can become the dominant cost when $\alpha = \frac{1}{2}\beta$ and $\beta \leq p\gamma$. This corresponds to the generalization of Cases C and D in Theorem 2.2 in the MIMC analysis in [15].

This outline analysis leads to the following theorem in which we make various assumptions and then draw conclusions about the resulting cost.

Theorem 6 (MIQMC). *Let P denote an output functional, and for each multi-index $\underline{\ell}$ let $P_{\underline{\ell}}$ denote the approximate output functional indexed by $\underline{\ell}$. Suppose for each multi-index $\underline{\ell}$ there exist independent estimators $Y_{\underline{\ell}}$ of $\mathbb{E}[\Delta P_{\underline{\ell}}]$ based on $N_{\underline{\ell}}$ deterministic QMC samples and $R_{\underline{\ell}} = R$ randomization steps, and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3, p$, with $p > 1$ and $\alpha \geq \frac{1}{2}\beta$, such that*

$$i) \quad |\mathbb{E}[P_{\underline{\ell}} - P]| \longrightarrow 0 \text{ as } \underline{\ell} \longrightarrow \infty \text{ } (\ell_j \rightarrow \infty, \forall j),$$

$$ii) \quad |\mathbb{E}[\Delta P_{\underline{\ell}}]| \leq c_1 2^{-\alpha\|\underline{\ell}\|_1}$$

$$iii) \quad \mathbb{E}_{\Delta}[Y_{\underline{\ell}}] = \mathbb{E}[\Delta P_{\underline{\ell}}]$$

$$iv) \quad \mathbb{V}_{\Delta}[Y_{\underline{\ell}}] \leq c_2 R^{-1} N_{\underline{\ell}}^{-p} 2^{-\beta\|\underline{\ell}\|_1}$$

$$v) \quad \text{cost}(Y_{\underline{\ell}}) \leq c_3 R N_{\underline{\ell}} 2^{\gamma\|\underline{\ell}\|_1}.$$

Then there exists a positive constant c_4 such that for any $\varepsilon < e^{-1}$ there are values L and $N_{\underline{\ell}}$ for which the MIQMC estimator (12) achieves the mean-square-error bound $\mathbb{E}_{\Delta}[(Y - \mathbb{E}[P])^2] < \varepsilon^2$ with the computational cost bound

$$\text{cost}(Y) \leq \begin{cases} c_4 \varepsilon^{-2/p}, & \beta > p\gamma, \\ c_4 \varepsilon^{-2/p} |\log \varepsilon|^{e_1}, & \beta = p\gamma, \\ c_4 \varepsilon^{-2/p - (p\gamma - \beta)/p\alpha} |\log \varepsilon|^{e_2}, & \beta < p\gamma, \end{cases}$$

where

$$\begin{aligned} e_1 &= d(p+1)/p, \quad e_2 = (d-1)((p+1)/p + (p\gamma - \beta)/p\alpha), & \text{if } \alpha > \frac{1}{2}\beta, \\ e_1 &= \max(d(p+1)/p, (d-1)(1+\gamma/\alpha)), \quad e_2 = (d-1)(1+\gamma/\alpha), & \text{if } \alpha = \frac{1}{2}\beta. \end{aligned}$$

Proof. Based on the outline proof which indicates how to specify the near optimal number of QMC points for each level, the detailed proof follows the same lines as the main MIMC Theorem 2.2 in [15]. \square

The key observation here is that the dimension d does not appear in the exponent for ε in the cost bounds, so it is a significant improvement over the MLQMC result in which the cost is of the form ε^{-r} with $r = \max(2/p, d/2)$, which limits the multilevel benefits even for $d=3$ if $p > 4/3$.

It is interesting to compare the cost given by this theorem with that given by the MIMC Theorem 4. If $\beta > p\gamma$, then the use of QMC improves the cost from $O(\varepsilon^{-2})$ to $O(\varepsilon^{-2/p})$. This is because the dominant costs in this case are on the coarsest levels where many points have to be sampled, and therefore QMC will provide substantial benefits. On the other hand, if $\beta < \gamma$ then both approaches give a cost of approximately $O(\varepsilon^{-\gamma/\alpha})$ because in this case the dominant costs are on the finest levels, and on the finest levels the optimal number of QMC points is $O(1)$, which is why the additional cost of rounding up to the nearest integer often dominates the main cost. Hence the use of QMC points is almost irrelevant in this case. Fortunately, we expect that the favourable case $\beta > p\gamma$ is likely to be the more common one. It is clearly the case in our very simple elliptic model with $\beta = 4$ and $\gamma = 1$.

7 Concluding Remarks

In this paper we began by summarizing the meta-theorems for MLMC and MIMC in a common framework for elliptic PDEs with random coefficients, where we applied full or sparse grid methods with respect to the spatial variables \mathbf{x} and used MC sampling for computing expected values with respect to the stochastic variables \mathbf{y} .

Following this, our novel contributions were

- showing that, in this context, MIMC is almost equivalent to the use of MLMC with sparse combination samples;
- introducing the idea of a) MLMC with sparse finite element or sparse combination samples, and b) nested MLMC, as other alternatives to MIMC;
- deriving the corresponding meta-theorems for MLQMC and MIQMC in this context, concluding that the computational cost to achieve $O(\varepsilon)$ r.m.s. accuracy can be reduced to $O(\varepsilon^{-r})$ with $r < 2$ independent of the spatial dimension d .

Natural extensions of the results in this paper include allowing the different indices in $\underline{\ell}$ to cover also different levels of dimension truncation in the stochastic variables \mathbf{y} , as well as providing verifications of the precise parameters α , β , γ and p for specific PDE applications.

Acknowledgements The authors acknowledge the support of the Australian Research Council under the projects FT130100655 and DP150101770.

References

1. Barth, A., Schwab, Ch., Zollinger, N.: Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients. *Numerische Mathematik*, 119(1):123–161 (2011)
2. Brugger, C., de Schryver, C., Wehn, N., Omland, S., Hefter, M., Ritter, K., Kostiuk, A., Korn, R.: Mixed precision multilevel Monte Carlo on hybrid computing systems. Proceedings of the Conference on Computational Intelligence for Financial Engineering and Economics, IEEE (2014)
3. Bungartz, H.-J., Griebel, M.: A note on the complexity of solving Poisson’s equation for spaces of bounded mixed derivatives. *Journal of Complexity*, 15(2):167–199 (1999)
4. Bungartz, H.-J., Griebel, M.: Sparse grids. *Acta Numerica*, 13:1–123 (2004)
5. Charrier, J., Scheichl, R., Teckentrup, A.: Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods. *SIAM Journal on Numerical Analysis*, 51(1):322–352 (2013)
6. Cliffe, K.A., Giles, M.B., Scheichl, R., Teckentrup, A.: Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science*, 14(1):3–15 (2011)
7. Dick, J., Kuo, F.Y., Le Gia, Q.T., Schwab, Ch.: Multi-level higher order QMC Petrov–Galerkin discretization for affine parametric operator equations. *SIAM Journal on Numerical Analysis*, 54(4):2541–2568 (2016)
8. Dick, J., Kuo, F.Y., Sloan, I.H.: High-dimensional integration: the quasi-Monte Carlo way. *Acta Numerica*, 22:133–288 (2013)
9. Gerstner, T., Noll, M.: Randomized multilevel quasi-Monte Carlo path simulation. In *Recent Developments in Computational Finance*, pages 349–372. World Scientific (2013)
10. Giles, M.B.: Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617 (2008)
11. Giles, M.B.: Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328 (2015)
12. Giles, M.B., Süli, E.: Adjoint methods for PDEs: *a posteriori* error analysis and postprocessing by duality. *Acta Numerica*, 11:145–236 (2002)
13. Giles, M.B., Waterhouse, B.J.: Multilevel quasi-Monte Carlo path simulation. In *Advanced Financial Modelling*, Radon Series on Computational and Applied Mathematics, pages 165–181. De Gruyter (2009)
14. Griebel, M., Schneider, M., Zenger, C.: A combination technique for the solution of sparse grid problems. In P. de Groen and P. R. Beauwens, editors, *Iterative Methods in Linear Algebra*. IMACS (1992)
15. Haji-Ali, A.-L., Nobile, F., Tempone, R.: Multi Index Monte Carlo: when sparsity meets sampling. *Numerische Mathematik*, 132:767–806 (2016)
16. Haji-Ali, A.-L., Giles, M.B.: MLMC for Value-at-Risk. Presentation at *International Conference on Monte Carlo Methods and Applications*, 2017. (URL: <http://people.maths.ox.ac.uk/hajiali/assets/files/hajiali-mcm2017-var.pdf>)
17. Harbrecht, H., Peters, M., Siebenmorgen, M.: On multilevel quadrature for elliptic stochastic partial differential equations. In *Sparse Grids and Applications*, volume 88 of *Lecture Notes in Computational Science and Engineering*, pages 161–179. Springer (2013)
18. Harbrecht, H., Peters, M., Siebenmorgen, M.: Multilevel accelerated quadrature for PDEs with log-Normally distributed diffusion coefficient. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):520–551 (2016)
19. Harbrecht, H., Peters, M., Siebenmorgen, M.: On the quasi-Monte Carlo method with Halton points for elliptic PDEs with log-normal diffusion. *Mathematics of Computation*, 86:771–797 (2017)
20. Kuo, F.Y., Nuyens, D.: Application of quasi-Monte Carlo methods to elliptic PDEs with random diffusion coefficients – a survey of analysis and implementation. *Foundations of Computational Mathematics*, 16(6):1631–1696 (2016)
21. Kuo, F.Y., Scheichl, R., Schwab, Ch., Sloan, I.H., Ullmann, E.: Multilevel quasi-Monte Carlo methods for lognormal diffusion problems. *Mathematics of Computation*, 86:2827–2860 (2017)

22. Kuo, F.Y., Schwab, Ch., Sloan, I.H.: Multi-level quasi-Monte Carlo finite element methods for a class of elliptic partial differential equations with random coefficients. *Foundations of Computational Mathematics*, 15(2):411–449 (2015)
23. Rhee, C.-H., Glynn, P.W.: Unbiased estimation with square root convergence for SDE models. *Operations Research*, 63(5):1026–1043 (2015)
24. Reisinger, C.: Analysis of linear difference schemes in the sparse grid combination technique. *IMA Journal of Numerical Analysis*, 33(2):544–581 (2013)
25. Teckentrup, A., Scheichl, R., Giles, M.B., Ullmann, E.: Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numerische Mathematik*, 125(3):569–600 (2013)