

A continuation multilevel Monte Carlo algorithm

Nathan Collier · Abdul-Lateef Haji-Ali ·
Fabio Nobile · Erik von Schwerin · Raúl Tempone

Received: 10 February 2014 / Accepted: 29 June 2014 / Published online: 5 September 2014
© Springer Science+Business Media Dordrecht 2014

Abstract We propose a novel Continuation Multi Level Monte Carlo (CMLMC) algorithm for weak approximation of stochastic models. The CMLMC algorithm solves the given approximation problem for a sequence of decreasing tolerances, ending when the required error tolerance is satisfied. CMLMC assumes discretization hierarchies that are defined a priori for each level and are geometrically refined across levels. The actual choice of computational work across levels is based on parametric models for the average cost per sample and the corresponding variance and weak error. These parameters are calibrated using Bayesian estimation, taking particular notice of the deepest levels of the discretization hierarchy, where only few realizations are available to produce the estimates. The resulting CMLMC estimator exhibits a non-trivial split-

Communicated by Desmond Higham.

N. Collier
Environmental Sciences Division, Oak Ridge National Lab, Climate Change Science Institute (CCSI),
Oak Ridge, USA
e-mail: nathaniel.collier@gmail.com

A.-L. Haji-Ali (✉) · R. Tempone
Applied Mathematics and Computational Sciences, KAUST, Thuwal, Saudi Arabia
e-mail: abdullateef.hajiali@kaust.edu.sa

R. Tempone
e-mail: raul.tempone@kaust.edu.sa

F. Nobile
MATHICSE-CSQI, EPF de Lausanne, Lausanne, Switzerland
e-mail: fabio.nobile@epfl.ch

E. von Schwerin
Department of Mathematics, Kungliga Tekniska Högskolan, 100 44 Stockholm, Sweden
e-mail: Schwerin@kth.se

ting between bias and statistical contributions. We also show the asymptotic normality of the statistical error in the MLMC estimator and justify in this way our error estimate that allows prescribing both required accuracy and confidence in the final result. Numerical results substantiate the above results and illustrate the corresponding computational savings in examples that are described in terms of differential equations either driven by random measures or with random coefficients.

Keywords Multilevel Monte Carlo · Monte Carlo · Partial differential equations with random data · Stochastic differential equations · Bayesian inference

Mathematics Subject Classification 65C05 · 65N22

1 Introduction

Multilevel Monte Carlo Sampling was first introduced for applications in the context of parametric integration by Heinrich [18, 19]. Later, to consider weak approximation of stochastic differential equations (SDEs) in mathematical finance, Kebaier [24] introduced a two-level Monte Carlo technique in which a coarse grid numerical approximation of an SDE was used as a control variate to a fine grid numerical approximation, thus reducing the number of samples needed on the fine grid and decreasing the total computational burden. This idea was extended to a multilevel Monte Carlo (MLMC) method by Giles in [12], who introduced a full hierarchy of discretizations with geometrically decreasing grid sizes. By optimally choosing the number of samples on each level this MLMC method decreases the computational burden, not only by a constant factor as standard control variate techniques do, but even reducing the rate in the computational complexity to compute a solution with error tolerance $TOL > 0$ from $\mathcal{O}(TOL^{-3})$ of the standard Euler-Maruyama Monte Carlo method to $\mathcal{O}(\log(TOL)^2 TOL^{-2})$, assuming that the work to generate a single realization is $\mathcal{O}(TOL^{-1})$. For one-dimensional SDEs, the computational complexity of MLMC was further reduced to $\mathcal{O}(TOL^{-2})$ by using the Milstein Scheme [11]. Moreover, the same computational complexity can be achieved by using antithetic control variates with MLMC in multi-dimensional SDEs with smooth and piecewise smooth payoffs [16].

This standard MLMC method has since then been extended and applied in a wide variety of contexts, including jump diffusions [31] and Partial Differential Equations (PDEs) with random coefficients [5–7, 13, 30]. It is shown in [30, Theorem 2.3] that there is an optimal convergence rate that is similar to the previously mentioned complexity rates, but that depends on the relation between the rate of variance convergence of the discretization method of the underlying equation and the work complexity associated with generating a single sample of the quantity of interest. In fact, in certain cases, the computational complexity can be of the optimal rate, namely $\mathcal{O}(TOL^{-2})$.

To achieve the optimal MLMC complexity rate and to obtain an estimate of the statistical error, sufficiently accurate estimates of the variance on each level must be obtained. Moreover, finding the optimal number of levels requires a sufficiently accurate estimate of the bias. As such, an algorithm is needed to find these estimates without incurring a significant overhead to the estimation of the wanted quantity of

interest. In [12], Giles proposed an algorithm, henceforth referred to as *Standard MLMC* or *SMLMC*, that works by iteratively increasing the number of levels and using sample variance estimates across levels. Moreover, SMLMC uses an arbitrary fixed accuracy splitting between the bias and the statistical error contributions. Other works [7, 14, 15, 29] listed similar versions of this algorithm. We outline this algorithm in Sect. 3.

In Sect. 4, we propose a novel continuation type of MLMC algorithm that uses models for variance and weak convergence and for average computational work per sample. We refer to this algorithm as *Continuation MLMC* or *CMLMC*. The CMLMC algorithm solves the given problem for a sequence of decreasing tolerances, which plays the role of a continuation parameter, the algorithm ends when the required error tolerance is satisfied. Solving this sequence of problems allows CMLMC to find increasingly accurate estimates of the bias and variances on each level, in addition to the quantity of interest, which is the goal of the computation. In each case, given the current estimate of problem parameters, the optimal number of levels of the MLMC hierarchy is found. Moreover, we use a Bayesian inference approach to robustly estimate the various problem parameters. The CMLMC algorithm is able to relax the statistical error bound given the bias estimate, to achieve the optimal splitting between the two. These techniques improve the computational complexity of the CMLMC algorithm and decreases the variability of the running time of the algorithm.

The outline of this work is as follows: We start in Sect. 2 by recalling the MLMC method and the assumed models on work, and on weak and variance convergence. After introducing the algorithms in Sects. 3 and 4, Sect. 5 presents numerical examples, which include three-dimensional PDEs with random inputs and an Itô SDE. Finally, we finish by offering conclusions and suggesting directions for future work in Sect. 6.

2 Multilevel Monte Carlo

2.1 Problem setting

Let $g(u)$ denote a real valued functional of the solution, u , of an underlying stochastic model. We assume that g is either a bounded linear functional or Lipschitz with respect to u . Our goal is to approximate the expected value, $E[g(u)]$, to a given accuracy TOL and a given confidence level. We assume that individual outcomes of the underlying solution u and the evaluation of the functional $g(u)$ are approximated by a discretization-based numerical scheme characterized by a mesh size, h . The value of h will govern the weak and strong errors in the approximation of $g(u)$ as we will see below. To motivate this setting, we now give two examples and identify the numerical discretizations, the discretization parameter, h , and the corresponding rates of approximation. The first example is common in engineering applications like heat conduction and groundwater flow. The second example is a simple one-dimensional geometric Brownian motion with European call option.

Example 2.1 Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space and \mathcal{D} be a bounded convex polygonal domain in \mathbb{R}^d . Find $u : \mathcal{D} \times \Omega \rightarrow \mathbb{R}$ that solves almost surely (a.s.) the following equation:

$$\begin{aligned} -\nabla \cdot (a(\mathbf{x}; \omega) \nabla u(\mathbf{x}; \omega)) &= f(\mathbf{x}; \omega) & \text{for } \mathbf{x} \in \mathcal{D}, \\ u(\mathbf{x}; \omega) &= 0 & \text{for } \mathbf{x} \in \partial\mathcal{D}, \end{aligned}$$

where $\omega \in \Omega$ and the value of the diffusion coefficient and the forcing are represented by random fields, yielding a random solution. We wish to compute $E[g(u)]$ for some deterministic functional g which is globally Lipschitz satisfying $|g(u) - g(v)| \leq G\|u - v\|_{H^1(\mathcal{D})}$ for some constant $G > 0$ and all $u, v \in H^1(\mathcal{D})$. Following [30], we also make the following assumptions

- $a_{\min}(\omega) = \min_{\mathbf{x} \in \mathcal{D}} a(\mathbf{x}; \omega) > 0$ a.s. and $1/a_{\min} \in L^p_{\mathbb{P}}(\Omega)$, for all $p \in (0, \infty)$.
- $a \in L^p_{\mathbb{P}}(\Omega, C^1(\overline{\mathcal{D}}))$, for all $p \in (0, \infty)$.
- $f \in L^{p^*}_{\mathbb{P}}(\Omega, L^2(\mathcal{D}))$ for some $p^* > 2$.

Here, $L^p_{\mathbb{P}}(\Omega, \mathcal{B})$ is the space of \mathcal{B} -valued random fields with a finite p 'th moment of their \mathcal{B} -norm, where the p 'th moment is with respect to measure \mathbb{P} . On the other hand, $C^1(\overline{\mathcal{D}})$ is the space of continuously differentiable functions with the usual norm [6]. Note that with these assumptions and since \mathcal{D} is bounded, one can show that $\max_{\mathbf{x} \in \mathcal{D}} a(\mathbf{x}; \omega) < \infty$ a.s. A standard approach to approximate the solution of the previous problem is to use Finite Elements on regular triangulation. In such a setting, the parameter $h > 0$ refers to either the maximum element diameter or another characteristic length and the corresponding approximate solution is denoted by $u_h(\omega)$. For piecewise linear or piecewise d -multilinear continuous finite element approximations, and with the previous assumptions, it can be shown [30, Corollary 3.1] that asymptotically as $h \rightarrow 0$:

- $|E[g(u) - g(u_h)]| \lesssim Q_W h^2$ for a constant $Q_W > 0$.
- $\text{Var}[g(u) - g(u_h)] \lesssim Q_S h^4$ for a constant $Q_S > 0$.

Example 2.2 Here we study the weak approximation of Itô stochastic differential equations (SDEs),

$$du(t) = a(t, u(t))dt + b(t, u(t))dW(t), \quad 0 < t < T, \tag{2.1}$$

where $u(t; \omega)$ is a stochastic process in \mathbb{R}^d , with randomness generated by a k -dimensional Wiener process with independent components, $W(t; \omega)$, cf. [23,26], and $a(t, u) \in \mathbb{R}^d$ and $b(t, u) \in \mathbb{R}^{d \times k}$ are the drift and diffusion fluxes, respectively. For any given sufficiently well behaved function, $g : \mathbb{R}^d \rightarrow \mathbb{R}$, our goal is to approximate the expected value, $E[g(u(T))]$. A typical application is to compute option prices in mathematical finance, cf. [17,22], and other related models based on stochastic dynamics. When one uses a standard Euler Maruyama (Forward Euler) method based on uniform time steps of size h to approximate (2.1), the following rates of approximation hold: $|E[g(u(T)) - g(u_h(T))]| = Q_W h + o(h)$ and $E[(g(u(T)) - g(u_h(T)))^2] = Q_S h + o(h)$, for some constants, $0 < Q_W, Q_S < \infty$, different from the constants of the previous example. For suitable assumptions on the functions a, b and g , we refer to [25,28].

To avoid cluttering the notation, we omit the reference to the underlying solution from now on, simply denoting the quantity of interest by g . Following the standard

MLMC approach, we assume, for any given non-negative integer $L \in \mathbb{N}$, that we have a hierarchy of $L + 1$ meshes defined by a decreasing sequence of mesh sizes $\{h_\ell\}_{\ell=0}^L$ where $h_\ell = h_0\beta^{-\ell}$ for some $h_0 > 0$ and a constant integer $\beta > 1$. We denote the resulting approximation of g using mesh size h_ℓ by g_ℓ , or by $g_\ell(\omega)$ when we want to stress the dependence on an outcome of the underlying random model. Using the following notation:

$$G_\ell(\omega) = \begin{cases} g_0(\omega) & \text{if } \ell = 0, \\ g_\ell(\omega) - g_{\ell-1}(\omega) & \text{if } \ell > 0, \end{cases}$$

the expected value of the finest approximation, g_L , can be expressed as

$$E[g_L] = \sum_{\ell=0}^L E[G_\ell],$$

where the MLMC estimator is obtained by replacing the expected values in the telescoping sum by sample averages. We denote the sample averages by \tilde{G}_ℓ as

$$\tilde{G}_\ell = M_\ell^{-1} \sum_{m=1}^{M_\ell} G_\ell(\omega_{\ell,m}).$$

Each sample average, \tilde{G}_ℓ , is computed using $M_\ell \in \mathbb{Z}_+$ independent identically distributed (i.i.d.) outcomes, $\{\omega_{\ell,m}\}_{m=1}^{M_\ell}$, of the underlying, mesh-independent, stochastic model; i.e. $\omega_{\ell,m} \in \Omega$ for all ℓ and m . The MLMC estimator can then be written as

$$A = \sum_{\ell=0}^L \tilde{G}_\ell. \tag{2.2}$$

Note that the outcomes are also assumed to be independent among the different sample averages, $\{\tilde{G}_\ell\}_{\ell=0}^L$.

We use the following model for the expected value of the cost associated with generating one sample of G_ℓ , including generating all the underlying random variables:

$$W_\ell \propto h_\ell^{-\gamma} = h_0^{-\gamma} \beta^{\ell\gamma}$$

for a given γ . Note the cost of generating a sample of G_ℓ might differ for different realizations, for example due to different number of iterations in an iterative method or due to adaptivity of the used numerical method. The parameter γ depends on the number of dimensions of the underlying problem and the used numerical method. For example, $\gamma = 1$ for the one-dimensional SDE in Example 2.2. For the PDE in Example 2.1, if the number of dimensions is $d = 3$ then $\gamma = 3\tilde{\gamma}$, where $\tilde{\gamma}$ depends on the solver used to solve the resulting linear system. In that example, iterative methods may have a smaller value of $\tilde{\gamma}$ than direct methods. The theoretical best-case scenario

for iterative methods would be $\tilde{\gamma} = 1$ for multigrid methods. On the other hand, we would have $\tilde{\gamma} = 3$ if one used a direct method using a naive Gaussian elimination on dense matrices. The total work of the estimator (2.2) is

$$W = \sum_{\ell=0}^L M_{\ell} W_{\ell}.$$

We want our estimator to satisfy a tolerance with prescribed failure probability $0 < \alpha < 1$, i.e.,

$$P[|E[g] - \mathcal{A}| > \text{TOL}] \leq \alpha, \tag{2.3}$$

while minimizing the work, W . Here, we split the total error into bias and statistical error,

$$|E[g] - \mathcal{A}| \leq \underbrace{|E[g - \mathcal{A}]|}_{\text{Bias}} + \underbrace{|E[\mathcal{A}] - \mathcal{A}|}_{\text{Statistical error}},$$

and use a splitting parameter, $\theta \in (0, 1)$, such that

$$\text{TOL} = \underbrace{(1 - \theta)\text{TOL}}_{\text{Bias tolerance}} + \underbrace{\theta\text{TOL}}_{\text{Statistical error tolerance}}.$$

The MLMC algorithm should bound the bias, $B = |E[g - \mathcal{A}]|$, and the statistical error as follows:

$$B = |E[g - \mathcal{A}]| \leq (1 - \theta)\text{TOL}, \tag{2.4a}$$

$$|E[\mathcal{A}] - \mathcal{A}| \leq \theta\text{TOL}, \tag{2.4b}$$

where the latter bound should hold with probability $1 - \alpha$. Note that θ does not have to be a constant, indeed it can depend on TOL as we shall see in Sect. 4. In the literature, some authors (e.g. [12]) have controlled the mean square error (MSE),

$$\text{MSE} = |E[g - \mathcal{A}]|^2 + E[|E[\mathcal{A}] - \mathcal{A}|^2],$$

rather than working with (2.3). We prefer to work with (2.3) since it allows us to prescribe both the accuracy TOL and the confidence level, $1 - \alpha$, in our results. The bound (2.4b) leads us to require

$$\text{Var}[\mathcal{A}] \leq \left(\frac{\theta\text{TOL}}{C_{\alpha}}\right)^2, \tag{2.5}$$

for some given confidence parameter, C_{α} , such that $\Phi(C_{\alpha}) = 1 - \frac{\alpha}{2}$; here, Φ is the cumulative distribution function of a standard normal random variable. The bound

(2.5) is motivated by the Lindeberg Central Limit Theorem in the limit $\text{TOL} \rightarrow 0$, cf. Lemma 7 in the Appendix.

By construction of the MLMC estimator, $E[\mathcal{A}] = E[g_L]$, and denoting $V_\ell = \text{Var}[G_\ell]$, then by independence, we have $\text{Var}[\mathcal{A}] = \sum_{\ell=0}^L V_\ell M_\ell^{-1}$, and the total error estimate can be written as

$$\text{Total error estimate} = B + C_\alpha \sqrt{\text{Var}[\mathcal{A}]} \tag{2.6}$$

Given L and $0 < \theta < 1$ and minimizing W subject to the statistical constraint (2.5) for $\{M_\ell\}_{\ell=0}^L \in \mathbb{R}^{L+1}$ gives the following optimal number of samples per level ℓ :

$$M_\ell = \left(\frac{C_\alpha}{\theta \text{TOL}} \right)^2 \sqrt{\frac{V_\ell}{W_\ell}} \left(\sum_{\ell=0}^L \sqrt{V_\ell W_\ell} \right). \tag{2.7}$$

When substituting the optimal number of samples in all levels the optimal work can be written in terms of L as follows

$$W(\text{TOL}, L) = \left(\frac{C_\alpha}{\theta \text{TOL}} \right)^2 \left(\sum_{\ell=0}^L \sqrt{V_\ell W_\ell} \right)^2 \tag{2.8}$$

Of course, the number of samples on each level is a positive integer. To obtain an approximate value of the optimal integer number of samples, we take the ceiling of the real-valued optimal values in (2.7).

In this work, we assume the following models on the weak error and variance:

$$E[g - g_\ell] \approx Q_W h_\ell^{q_1}, \tag{2.9a}$$

$$\text{Var}[g_\ell - g_{\ell-1}] \approx Q_S h_{\ell-1}^{q_2}, \tag{2.9b}$$

for some constants $Q_W \neq 0, Q_S > 0, q_1 > 0$ and $0 < q_2 \leq 2q_1$. Note that the condition that is usually assumed for MLMC is $\min(q_2, d\gamma) \leq 2q_1$, cf. [30, Theorem 2.3], to ensure that the cost of MLMC is not dominated by the cost of a single sample on each level. We assume instead the slightly more restrictive condition $q_2 \leq 2q_1$. As an example, recall that the PDE in Example 2.1 has $q_2 = 2q_1$ and in Sect. 5, the PDE is solved using a finite element method with standard trilinear basis and it has $q_1 = 2$. On the other hand, for the SDE in Example 2.2 with Euler discretization, $q_1 = q_2 = 1$. Collectively, we refer to the parameters q_1, q_2, Q_S, Q_W and $\{V_\ell\}_{\ell=0}^L$ as problem parameters. Based on these models, we can write for $\ell > 0$

$$E[G_\ell] \approx Q_W h_0^{q_1} \beta^{-\ell q_1} (\beta^{q_1} - 1), \tag{2.10a}$$

$$\text{Var}[G_\ell] = V_\ell \approx Q_S h_0^{q_2} \beta^{-(\ell-1)q_2}. \tag{2.10b}$$

Specifically, as a consequence of (2.9a), the bias model is

$$B \approx |Q_W| h_0^{q_1} \beta^{-Lq_1}. \tag{2.11}$$

Finally, we note that the algorithms presented in this work are iterative. We therefore denote by $\bar{M}_\ell, \bar{G}_\ell$ and \bar{V}_ℓ the total number of samples of G_ℓ generated in all iterations and their sample average and sample variance, respectively. Explicitly, we write¹

$$\bar{G}_\ell = \frac{1}{\bar{M}_\ell} \sum_{m=1}^{\bar{M}_\ell} G_\ell(\omega_{\ell,m}), \tag{2.12a}$$

$$\bar{V}_\ell = \frac{1}{\bar{M}_\ell} \sum_{m=1}^{\bar{M}_\ell} (G_\ell(\omega_{\ell,m}) - \bar{G}_\ell)^2. \tag{2.12b}$$

3 Standard MLMC

3.1 Overview

While minor variations exist among MLMC algorithms listed in [12, 15, 16], we believe that there is sufficient commonality in them for us to outline here the overarching idea and refer to this collection of methods as the Standard MLMC algorithm or simply SMLMC. SMLMC solves the problem by iteratively increasing the number of levels of the MLMC hierarchy. In order to find the optimal number of samples of each level ℓ , an estimate of the variance V_ℓ is needed. If there were previously generated samples in previous iterations for a level ℓ , the sample variance \bar{V}_ℓ is used. Otherwise, an initial fixed number of samples, \tilde{M} , is generated. Moreover, in most works, the splitting between bias and statistical error, θ , is chosen to be 0.5.

After running the hierarchy, an estimate of the total error is computed. To this end, the work [12] approximates the absolute value of the constant, Q_W , using a similar expression to the following:

$$|Q_W| \approx \frac{\max(|\bar{G}_L|, |\bar{G}_{L-1}| \beta^{-q_1})}{h_0^{q_1} \beta^{-Lq_1} (\beta^{q_1} - 1)} := \tilde{Q}_W.$$

In other words, the absolute value of the constant Q_W is estimated using the samples generated on the last two levels. Thus, this estimate is only defined for $L \geq 2$. Next, the variance of the estimator, $\text{Var}[\mathcal{A}]$, is approximated by

$$\text{Var}[\mathcal{A}] \approx \sum_{\ell=0}^L \frac{\bar{V}_\ell}{\bar{M}_\ell} := \tilde{V}.$$

¹ For the variance estimator, one can also use the unbiased estimator; by dividing by $\bar{M}_\ell - 1$ instead of \bar{M}_ℓ . All discussions in this work still apply.

Finally, a total error estimate can be computed as outlined by (2.6)

$$\text{Total error estimate} = \tilde{Q}_W h_0^{q_1} \beta^{-Lq_1} + C_\alpha \sqrt{\tilde{V}}. \tag{3.1}$$

The complete algorithm is outlined in Algorithm 1.

Algorithm 1

```

1: function STANDARDMLMC(TOL,  $\tilde{M}$ ,  $\theta$ )
2:   Start with  $L = 0$ .
3:   loop
4:     Add new levels to  $\{h_\ell\}_{\ell=0}^L$ .
5:     Generate  $\tilde{M}$  samples for level  $L$  and estimate  $\bar{V}_L$ .
6:     Using sample variance estimates,  $\{\bar{V}_\ell\}_{\ell=0}^L$  from all iterations, and the constant  $\theta$ ,
       compute optimal number of samples,  $\{M_\ell\}_{\ell=0}^L$ , according to (2.7).
7:     Run the hierarchy using the optimal number of samples.
8:     If  $L \geq 2$  and the total estimate error (3.1) is less than TOL, then END.
9:     Otherwise, set  $L = L + 1$ .
10:  end loop
11: end function

```

Usually all samples from previous iterations are used in the algorithm to run the hierarchy in step 7 to calculate the required quantity of interest. However, the analysis of the bias and the statistical error of the resulting estimator is difficult and has not been done before, to the best of our knowledge.

3.2 Accuracy of the parameter estimates

In the standard algorithm, Q_W and the variances $\{V_\ell\}_{\ell=0}^L$ are needed and estimated. In this section, we look at the accuracy of the estimators for these problem parameters.

We examine the accuracy of the sample variance by computing its squared relative error for $\ell > 1$:

$$\begin{aligned} \frac{\text{Var}[\bar{V}_\ell]}{V_\ell^2} &= \frac{(\bar{M}_\ell - 1)^2}{\bar{M}_\ell^3 V_\ell^2} \left(\text{E}[(G_\ell - \text{E}[G_\ell])^4] - \frac{V_\ell^2(\bar{M}_\ell - 3)}{\bar{M}_\ell - 1} \right) \\ &= \frac{(\bar{M}_\ell - 1)^2}{\bar{M}_\ell^3} \left(\text{E}[(G_\ell - \text{E}[G_\ell])^4] V_\ell^{-2} - \frac{\bar{M}_\ell - 3}{\bar{M}_\ell - 1} \right) \\ &\approx \frac{(\bar{M}_\ell - 1)^2}{\bar{M}_\ell^3} \left(\text{E}[(G_\ell - \text{E}[G_\ell])^4] Q_S^{-2} h_\ell^{-2q_2} - \frac{\bar{M}_\ell - 3}{\bar{M}_\ell - 1} \right). \end{aligned}$$

Unless $\text{E}[(G_\ell - \text{E}[G_\ell])^4] \leq Ch_\ell^{2q_2}$, for some constant $C > 0$, or M_ℓ increases sufficiently fast, the relative error in the estimator \bar{V}_ℓ can become unbounded as

$\ell \rightarrow \infty$. Similarly, the relative error of the sample variance at level $\ell = 0$ can be shown to be bounded for instance by assuming that the second and fourth central moments of G_0 are bounded.

Next, for simplicity, we look at the squared relative error estimate of Q_W by assuming that it is estimated using samples on a single level, L , only.

$$\begin{aligned} \frac{\text{Var} \left[\left| \frac{\bar{G}_L}{h_0^{q_1} \beta^{-Lq_1} (\beta^{q_1} - 1)} \right| \right]}{Q_W^2} &= \frac{V_L}{Q_W^2 M_L h_0^{2q_1} \beta^{-2Lq_1} (\beta^{q_1} - 1)^2} \\ &= \frac{Q_S}{Q_W^2} \cdot \frac{h_0^{q_2} \beta^{-q_2 L}}{Q_W^2 M_L h_0^{2q_1} \beta^{-2Lq_1} (\beta^{q_1} - 1)^2} \\ &= \frac{Q_S h_0^{q_2 - 2q_1}}{Q_W^2 (\beta^{q_1} - 1)^2} \left(\frac{\beta^{L(2q_1 - q_2)}}{M_L} \right). \end{aligned}$$

Observe now that if $q_2 < 2q_1$ (as in Example 2.2), then, for the previous relative error estimate to be $o(1)$, we must have $M_L \propto \beta^{L(2q_1 - q_2)} \rightarrow \infty$ as $L \rightarrow \infty$. This analysis shows that in some cases, M_L will have to grow to provide an accurate estimate to Q_W , regardless of the optimal choice of the number of samples outlined in (2.7).

4 Continuation MLMC (CMLMC)

In this section we discuss the main contribution of this work, a continuation MLMC (CMLMC) algorithm that approximates the value $E[g(u)]$. We begin in the next subsection by giving an overview of the general idea of algorithm. Subsequent subsections discuss how to estimate all the required problem parameters that are necessary for running the algorithm. CMLMC is listed in Algorithm 2.

Algorithm 2

- 1: **function** CMLMC(Parameters summarized in Table 1)
 - 2: Compute with an initial hierarchy.
 - 3: Estimate problem parameters $\{V_\ell\}_{\ell=0}^L$, Q_S , Q_W , q_1 and q_2 according to Sect. 4.2.
 - 4: Set $i = 0$.
 - 5: **repeat**
 - 6: Find L according to (4.1).
 - 7: Add new levels to $\{h_\ell\}_{\ell=0}^L$.
 - 8: Using the variance estimates (4.5) and θ from (4.2), compute the optimal number of samples according to (2.7).
 - 9: Run the resulting hierarchy using the optimal number of samples.
 - 10: Estimate problem parameters, $\{V_\ell\}_{\ell=0}^L$, Q_S , Q_W , q_1 and q_2 , according to Sect. 4.2.
 - 11: Estimate the total error according to (2.6).
 - 12: Set $i = i + 1$
 - 13: **until** $i > i_E$ and the total error estimate is less than TOL
 - 14: **end function**
-

4.1 Overview

The main idea of CMLMC is to solve for $E[g(u)]$ with a sequence of decreasing tolerances. By doing this, CMLMC is able to increasingly improve estimates of several problem dependent parameters while solving relatively inexpensive problems corresponding to large tolerances. These parameters estimates are crucial to optimally distribute computational effort when solving for the last tolerance, which is the desired tolerance, TOL , or smaller. Moreover, the sequence is built such that the total work of the algorithm is close to the work of MLMC when solving for the desired tolerance, TOL , assuming all the necessary parameters are known a priori. To this end, we make the following choice for the sequence of decreasing tolerances TOL_i for $i = 0, 1, \dots$

$$TOL_i = \begin{cases} r_1^{i_E-i} r_2^{-1} TOL & i < i_E, \\ r_2^{i_E-i} r_2^{-1} TOL & i \geq i_E, \end{cases}$$

where $r_1 \geq r_2 > 1$. By imposing $TOL_0 = TOL_{\max}$ for some maximum tolerance, we have

$$i_E = \left\lceil \frac{-\log(TOL) + \log(r_2) + \log(TOL_{\max})}{\log(r_1)} \right\rceil.$$

Iterations for which $i \leq i_E$ are meant to obtain increasingly more accurate estimates of the problem parameters. The iteration i_E solves the problem for the tolerance $r_2^{-1}TOL$. Notice that the problem is solved for a slightly smaller tolerance than the required tolerance TOL . This tolerance reduction is to prevent extra unnecessary iterations due to slight variations in estimates of the problem parameters. This technique improves the overall average running time of the algorithm. Similarly, iterations $i > i_E$ have tolerances that are even smaller to account for cases in which estimates of the problem parameters are unstable. The parameters r_1 and r_2 are chosen such that the total work of the algorithm is not significantly more than the work of the final hierarchy that solves the problem with the required tolerance, TOL . For example, if the work of the MLMC estimator is $\mathcal{O}(TOL^{-2})$, we choose $r_1 = 2$ to ensure that the work of iteration i is roughly four times the work of iteration $i - 1$ for iterations for which $TOL_i \geq TOL$. The choice of $r_2 = 1.1$, on the other hand, ensures that for iterations for which $TOL_i < TOL$, the work of iterations of i is roughly 1.2 times the work of iteration $i - 1$.

Consider now the i -th iteration of CMLMC and assume that estimates for $\mathbf{Q} := \{q_1, q_2, Q_W, Q_S\}$ and $\{V_\ell\}_{\ell=0}^L$ are available from previous iterations; we will discuss how to obtain these estimate in Sect. 4.2. The i -th iteration begins by selecting the optimal number of levels $L[i]$ that solves the problem for the given tolerance, TOL_i , as follows

$$L[i] = \operatorname{argmin}_{L_{\min}[i] \leq L \leq L_{\max}[i]} W(TOL_i, L), \tag{4.1}$$

where $W(\text{TOL}_i, L)$ is defined by (2.8) and depends on all the parameters \mathbf{Q} and $\{V_\ell\}_{\ell=0}^L$ and $\theta = \theta(L)$ given by

$$\theta = 1 - \frac{|Q_W|h_L^{q_1}}{\text{TOL}_i} = 1 - \frac{|Q_W|h_0^{q_1}\beta^{-Lq_1}}{\text{TOL}_i}, \tag{4.2}$$

which comes from enforcing that the bias model (2.11) equals $(1 - \theta)\text{TOL}_i$. Moreover, L_{\min} should satisfy $Q_W h_{L_{\min}}^{q_1} = \text{TOL}_i$ or, since we have $h_\ell = h_0 \beta^{-\ell}$,

$$L_{\min}[i] = \max \left(L[i - 1], \frac{q_1 \log(h_0) - \log\left(\frac{\text{TOL}_i}{|Q_W|}\right)}{q_1 \log \beta} \right),$$

where $L[i - 1]$ is the number of levels from the previous iteration. This ensures that L does not decrease from one iteration to the next, which agrees with our intuition that L increases with $\log(\text{TOL}_i^{-1})$. On the other hand, L_{\max} is given by other considerations. For instance, it could be related to the minimum mesh size imposed by memory or computational restrictions. More practically, to ensure robustness, L_{\max} can be chosen to be $L_{\min} + L_{\text{inc}}$, for a given fixed integer L_{inc} , so that L has limited increments from one iteration to the next. Since only few values of L are considered in the optimization (4.1), it is easy to find the optimal L by exhaustive search. The choice (4.2) implies that the statistical constraint (2.5) is relaxed (or tightened) depending on the estimated bias of each hierarchy. The iteration then continues by running the resulting hierarchy with the optimal number of samples $\{M_\ell\}_{\ell=0}^L$ according to (2.7). Finally the iteration ends by improving the estimates of the problem parameters \mathbf{Q} and $\{V_\ell\}_{\ell=0}^L$ as well as the quantity of interest based on the newly available samples as described in Sect. 4.2.

To start CMLMC we compute with an initial, relatively inexpensive, hierarchy. The purpose of using this initial hierarchy is to obtain rough estimates of the problem parameters. Such a hierarchy cannot depend on estimates of problem parameters and should have at least three levels to allow estimating \mathbf{Q} ; these three levels are needed to be able to extrapolate (or interpolate) the weak error and variance estimates on all MLMC levels. The algorithm stops when the total error estimate is below the required tolerance TOL.

4.2 Parameters estimation

In this section, we discuss how to improve estimates of the parameters \mathbf{Q} as well as the variances V_ℓ based on the generated samples in all iterations and all levels. For easier presentation, we will also use the following notation

$$\begin{aligned} w_\ell(q_1) &= h_0^{q_1} \beta^{-\ell q_1} (\beta^{q_1} - 1), \\ s_\ell(q_2) &= h_0^{-q_2} \beta^{\ell q_2}. \end{aligned}$$

Thus, using the notation above, (2.10) becomes

$$E[G_\ell] \approx Q_W w_\ell(q_1), \tag{4.3a}$$

$$\text{Var}[G_\ell] = V_\ell \approx Q_S s_\ell^{-1}(q_2). \tag{4.3b}$$

4.2.1 Estimating variances V_ℓ

We first assume that we have estimates of q_1, q_2, Q_W and Q_S and discuss estimating the variances, $\{V_\ell\}_{\ell=0}^L$, and the total statistical error after computing with a given hierarchy. Estimating q_1, q_2, Q_W and Q_S is discussed in the next subsection.

Usually the variances $\{V_\ell\}_{\ell=0}^L$ are estimated by using the sample variance estimator (2.12b) to estimate the statistical error as well as the optimal number of samples $\{M_\ell\}_{\ell=0}^L$. However, sometimes there are too few samples in a given level to give a corresponding accurate variance estimate. This is specially acute on the deepest levels, and unlike the standard MLMC algorithm, we do not impose a minimum number of samples across levels to obtain a stable estimate of the sample variance. Recalling that we have the variance model (4.3b) at our disposal, we can use this model to estimate the variance at all levels $\ell > 0$. However, the model (4.3b) is only accurate asymptotically. We can use the generated samples on each level to locally improve the accuracy of the V_ℓ estimates. To this end, we use a Bayesian setting [27].

We assume that G_ℓ follows a normal distribution with mean μ_ℓ and precision λ_ℓ (precision is simply the inverse of the variance). To simplify the computation, we choose a normal-gamma prior on (μ_ℓ, λ_ℓ) – the conjugate prior of the normal likelihood. The resulting posterior probability density function (pdf) is also a normal-gamma distribution function. We choose the parameters $(\widehat{\mu}_\ell, \kappa_0, 0.5 + \widehat{\lambda}_\ell \kappa_1, \kappa_1)$ for the normal-gamma prior, such that it is maximized at $\widehat{\mu}_\ell$ and $\widehat{\lambda}_\ell$. The parameter $\widehat{\mu}_\ell$ and $\widehat{\lambda}_\ell$ serve as initial guesses for μ_ℓ and λ_ℓ , respectively. Moreover, κ_0 and κ_1 are positive constants that model our certainty in those respective guesses. We use the assumed models of the weak error and variance (4.3) to give the initial guesses

$$\widehat{\mu}_\ell = Q_W w_\ell(q_1), \tag{4.4a}$$

$$\widehat{\lambda}_\ell = Q_S^{-1} s_\ell(q_2). \tag{4.4b}$$

As mentioned, the posterior pdf is also a normal-gamma with parameters $(\Upsilon_{1,\ell}, \Upsilon_{2,\ell}, \Upsilon_{3,\ell}, \Upsilon_{4,\ell})$ and it is maximized at $(\Upsilon_{1,\ell}, \frac{\Upsilon_{3,\ell}-0.5}{\Upsilon_{4,\ell}})$. Specifically

$$\begin{aligned} \Upsilon_{3,\ell} &= 0.5 + \kappa_1 \widehat{\lambda}_\ell + \frac{\overline{M}_\ell}{2}, \\ \Upsilon_{4,\ell} &= \kappa_1 + \frac{1}{2} \left(\sum_{m=1}^{\overline{M}_\ell} (G_{\ell,m} - \overline{G}_\ell)^2 \right) + \frac{\kappa_0 \overline{M}_\ell (\overline{G}_\ell - \widehat{\mu}_\ell)^2}{2(\kappa_0 + \overline{M}_\ell)}. \end{aligned}$$

As such, we use the following estimate of the variance V_ℓ for $\ell > 0$

$$V_\ell \approx \frac{\Upsilon_{4,\ell}}{\Upsilon_{3,\ell} - 0.5}. \tag{4.5}$$

Estimating the variance at the coarsest mesh, V_0 , can be done using the sample variance. The number of samples on the coarsest level, M_0 , is usually large enough to produce a stable and accurate estimate. Using these estimates and the bias estimate (2.11), the total error can be estimated as (2.6).

4.2.2 Estimating \mathbf{Q}

To incorporate prior knowledge on q_1 and q_2 including initial guesses and the relation $q_2 \leq 2q_1$, we again follow a Bayesian setting to estimate these parameters and assume that G_ℓ follows a Gaussian distribution with mean $Q_W w_\ell(q_1)$ and variance $Q_S s_\ell^{-1}(q_2)$. In what follows, ℓ_0 is a non-negative integer. With these assumptions, the corresponding likelihood is

$$\begin{aligned} \mathcal{L} &= \left(\prod_{\ell=\ell_0}^L (2\pi Q_S s_\ell^{-1}(q_2))^{\frac{-\bar{M}_\ell}{2}} \right) \\ &\times \exp \left(-\frac{1}{2Q_S} \sum_{\ell=\ell_0}^L s_\ell(q_2) \sum_{m=1}^{\bar{M}_\ell} (G_{\ell,m} - Q_W w_\ell(q_1))^2 \right). \end{aligned} \tag{4.6}$$

Assuming a improper prior on Q_W and Q_S and maximizing the resulting posterior pdf with respect to Q_W and Q_S gives the following weighted least-squares solution:

$$Q_W^* = \left(\sum_{\ell=\ell_0}^L \bar{M}_\ell w_\ell^2(q_1) s_\ell(q_2) \right)^{-1} \sum_{\ell=\ell_0}^L w_\ell(q_1) s_\ell(q_2) \bar{M}_\ell \bar{G}_\ell, \tag{4.7a}$$

$$Q_S^* = \left(\sum_{\ell=\ell_0}^L \bar{M}_\ell \right)^{-1} \sum_{\ell=\ell_0}^L s_\ell(q_2) \sum_{m=1}^{\bar{M}_\ell} (G_{\ell,m} - Q_W w_\ell(q_1))^2. \tag{4.7b}$$

We can substitute the previous expressions for Q_W and Q_S in (4.6) to obtain a likelihood in terms of q_1 and q_2 . Denoting $\bar{M} = \sum_{\ell=\ell_0}^L \bar{M}_\ell$, we write

$$\begin{aligned} \mathcal{L}(q_1, q_2) &= \exp \left(-\frac{\bar{M}}{2} \right) \left(\sum_{\ell=\ell_0}^L \sum_{m=0}^{\bar{M}_\ell} s_\ell(q_2) G_{\ell,m}^2 \right. \\ &\quad \left. - \frac{\left(\sum_{\ell=\ell_0}^L s_\ell(q_2) w_\ell(q_1) \bar{M}_\ell \bar{G}_\ell \right)^2}{\sum_{\ell=\ell_0}^L \bar{M}_\ell w_\ell(q_1)^2 s_\ell(q_2)} \right)^{-\frac{\bar{M}}{2}}. \end{aligned}$$

Table 1 Summary of parameters in CMLMC

Parameter	Purpose
$\widehat{x}_0, \widehat{x}_1, \sigma_0$ and σ_1	Parameters to model the initial guess of q_1 and q_2 and the confidence in those estimates
κ_0 and κ_1	The confidence in the variance and weak convergence models, respectively
TOL _{max}	The maximum tolerance with which to start the algorithm
r_1 and r_2	Controls the computational burden to calibrate the problem parameters compared to the one taken to solve the problem
Initial hierarchy	The initial hierarchy to start the algorithm. Must be relatively inexpensive and has at least three levels
L_{inc}	Maximum number of values to consider when optimizing for L
\mathfrak{L}	Maximum number of levels used to compute estimates of Q_W and Q_S
C_α	Parameter related to the confidence in the statistical constraint

We can then assume a prior on q_1 and q_2 . However, remember that $q_2 \leq 2q_1$, and $q_1 > 0$. As such, we introduce the unconstrained parameters $x_0(q_1) = \log(q_1) \in \mathbb{R}$ and $x_1(q_1, q_2) = \log(2q_1 - q_2) \in \mathbb{R}$ and assume a Gaussian prior on them

$$\rho_{\text{prior}}(q_1, q_2) = \frac{1}{2\pi\sqrt{\sigma_0^2\sigma_1^2}} \exp\left(-\frac{(x_0(q_1) - \widehat{x}_0)^2}{2\sigma_0^2} - \frac{(x_1(q_1, q_2) - \widehat{x}_1)^2}{2\sigma_1^2}\right).$$

Here, \widehat{x}_0 and \widehat{x}_1 represent our initial guesses of x_0 and x_1 , respectively, which we can obtain from a rough analysis of the problem. Moreover, σ_1 and σ_2 model our confidence in those guesses. The more accurate our initial guesses are, the faster the algorithm converges. Finally, we numerically maximize the log of the posterior pdf with respect to $(x_0, x_1) \in \mathbb{R}^2$ using a suitable numerical optimization algorithm. For robustness, we choose $\ell_0 = 1$ to estimate q_1 and q_2 . In other words we include samples from all levels $\ell > 0$ for this estimation.

Given estimates of q_1 and q_2 , we can use the least-squares estimates Q_W^* and Q_S^* in (4.7) as estimates of Q_W and Q_S , respectively. However, usually not all levels follow the assumed asymptotic models (2.10) and as such special care must be taken to choose ℓ_0 in these estimates. The parameter Q_W must be accurate on deeper levels since it is used to compute the bias (2.11). Similarly, Q_S must be accurate on deeper levels where not many samples are available and the variance estimate (4.5) is mainly determined by the initial guess (4.4b). For these reasons, when computing Q_W^* and Q_S^* , we choose $\ell_0 = \max(1, L - \mathfrak{L})$ in (4.7) for some positive integer \mathfrak{L} that denotes the maximum number of levels use to compute the estimates. Finally, since Q_W has

an improper prior, its posterior is also a Gaussian with mean Q_W^* and variance

$$V_W := \sum_{\ell=\ell_0}^L \frac{Q_S}{\overline{M} w_\ell^2(q_2) s_\ell(q_1)}.$$

Motivated by the accuracy analysis of the Q_W estimate in Sect. 3.2, we use a worst-case estimate of Q_W instead of simply using the estimate Q_W^* in (4.7a). The worst-case estimate is produced by adding the maximum sampling error with $1 - \alpha$ confidence, namely $C_\alpha \sqrt{V_W}$, multiplied by the sign of Q_W^* . In other words, our estimate of Q_W is $Q_W^* + \text{sign}(Q_W^*) C_\alpha \sqrt{V_W}$.

4.3 Algorithm parameters

Table 1 summarizes the parameters that control the CMLMC algorithm. Some of these parameters need to be suitably chosen for the specific problem. However, while there might be optimal values for these parameters to minimize the average running time, it is our experience that reasonable values of these parameters are enough to get average running times that are near-optimal. In fact, similar results to those that we show Sect. 5 were obtained with variations of κ_1 and κ_2 ; namely $\kappa_1 = \kappa_2 \in \{0.05, 0.1, 0.2\}$.

5 Numerical tests

In this section, we first introduce the test problems. We then describe several implementation details and finish by presenting the actual numerical results.

5.1 Test problems

We look at three test problems: the first two are based on PDEs with random inputs and the last one is based on an Itô SDE.

5.1.1 Ex. 1

This problem is based on Example 2.1 in Sect. 2.1 with some particular choices that satisfy the assumptions therein. First, we choose $\mathcal{D} = [0, 1]^3$ and assume that the forcing is

$$f(\mathbf{x}; \omega) = f_0 + \hat{f} \sum_{i=0}^K \sum_{j=0}^K \sum_{k=0}^K \Phi_{ijk}(\mathbf{x}) Z_{ijk},$$

where

$$\Phi_{ijk}(\mathbf{x}) = \sqrt{\lambda_i \lambda_j \lambda_k} \phi_i(x_1) \phi_j(x_2) \phi_k(x_3),$$

and

$$\phi_i(x) = \begin{cases} \cos\left(\frac{5\Lambda i}{2}\pi x\right) & i \text{ is even,} \\ \sin\left(\frac{5\Lambda(i+1)}{2}\pi x\right) & i \text{ is odd,} \end{cases}$$

$$\lambda_i = (2\pi)^{\frac{7}{6}} \Lambda^{\frac{11}{6}} \begin{cases} \frac{1}{2} & i = 0, \\ \exp\left(-2\left(\frac{\Lambda i}{4}\pi\right)^2\right) & i \text{ is even,} \\ \exp\left(-2\left(\frac{\Lambda(i+1)}{4}\pi\right)^2\right) & i \text{ is odd,} \end{cases}$$

for given $\Lambda > 0$, and positive integer K and $\mathbf{Z} = \{Z_{ijk}\}$ a set of $(K + 1)^3$ i.i.d. standard normal random variables. Moreover, we choose the diffusion coefficient to be a function of two random variables as follows:

$$a(\mathbf{x}; \omega) = a_0 + \exp\left(4Y_1\Phi_{121}(\mathbf{x}) + 40Y_2\Phi_{877}(\mathbf{x})\right).$$

Here, $\mathbf{Y} = \{Y_1, Y_2\}$ is a set of i.i.d. normal Gaussian random variables, also independent of \mathbf{Z} . Finally we make the following choice for the quantity of interest, g :

$$g = (2\pi\sigma^2)^{-\frac{3}{2}} \int_{\mathcal{D}} \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_0\|_2^2}{2\sigma^2}\right) u(\mathbf{x}) d\mathbf{x},$$

and select the parameters $a_0 = 0.01$, $f_0 = 50$, $\hat{f} = 10$, $\Lambda = \frac{0.2}{\sqrt{2}}$, $K = 10$, $\sigma^2 = 0.02622863$ and $\mathbf{x}_0 = [0.5026695, 0.26042876, 0.62141498]$. Since the diffusion coefficient, a , is independent of the forcing, f , a reference solution can be calculated to sufficient accuracy by scaling and taking expectation of the weak form with respect to \mathbf{Z} to obtain a formula with constant forcing for the conditional expectation with respect to \mathbf{Y} . We then use stochastic collocation [3] with 11 Hermite quadrature points for each variable in \mathbf{Y} (thus totaling 121 points) and a Finite Difference method with centered differences and 128 equally spaced points in each dimension to produce the reference value $E[g]$. Using this method, the reference value 1.6026 was computed with an error estimate of 10^{-4} .

5.1.2 Ex.2

The second example is a slight variation of the first. First, we choose the following diffusion coefficient instead:

$$a(\mathbf{x}; \omega) = a_0 + \exp\left(Y_1\phi_{121}(\mathbf{x}) + Y_2\phi_{877}(\mathbf{x})\right).$$

Moreover, in this example \mathbf{Y} is a set of two i.i.d. uniform random variables in the range $[-1, 1]$, again independent of \mathbf{Z} . We also make the following choice for the quantity of interest g

$$g = 100 \left(2\pi\sigma^2\right)^{-\frac{3}{2}} \int_{\mathcal{D}} \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}_0\|_2^2}{2\sigma^2}\right) u(\mathbf{x}) d\mathbf{x},$$

and select the parameters $a_0 = 1$, $f_0 = 1$, $\hat{f} = 1$, $\Lambda = 0.2$, $K = 10$, $\sigma^2 = 0.01194691$ and $\mathbf{x}_0 = [0.62482261, 0.45530923, 0.49862328]$. We use the same method as in **Ex.1** to compute the reference solution, except that in this case, we use Legendre quadrature in the stochastic collocation method, instead of Hermite quadrature. The computed reference solution $E[g]$ in this case is 2.3627 with an error estimate of 10^{-4} .

5.1.3 Ex.3

The third example is a one-dimensional geometric Brownian motion based on Example 2.2. We make the following choices:

$$\begin{aligned} u(0) &= 1, T = 1, \\ a(t, u) &= 0.05u, \\ b(t, u) &= 0.2u, \\ g(u) &= 10 \max(u(1) - 1, 0). \end{aligned}$$

The exact solution can be computed using a standard change of variables and Itô's formula. For the selected parameters, the solution is $E[g] = 1.04505835721856$.

5.2 Implementation and Runs

All the algorithms mentioned in this work were implemented using the C programming language, with the goal that the software be as optimal as possible, while maintaining generality.

For implementing the solver for the PDE test problems (**Ex.1** and **Ex.2**), we use PetIGA [8,9]. While the primary intent of this framework is to provide high-performance B-spline-based finite element discretizations, it is also useful for applications where the domain is topologically square and subject to uniform refinements. As its name suggests, PetIGA is designed to tightly couple to PETSc [4]. The framework can be thought of as an extension of the PETSc library, which provides methods for assembling matrices and vectors related to the discretization of integral equations.

In our PDE numerical tests (**Ex.1** and **Ex.2**), we use a standard trilinear basis to discretize the weak form of the model problem, integrating with eight quadrature points. We also generate results for two linear solvers that PETSc provides an interface to. The first solver is an iterative GMRES solver that solves a linear system in almost linear time with respect to the number of degrees of freedom for the mesh sizes of interest; in other words $\tilde{\gamma} = 1$ in this case. The second solver we tried is a direct one, called MUMPS [1,2]. For the mesh sizes of interest, the running time of MUMPS varies from quadratic to linear in the total number of degrees of freedom. The best fit turns out to be $\tilde{\gamma} = 1.5$ in the case.

Table 2 Summary of problem parameters

	γ	q_1	q_2	s_1	s_2
Ex.1 and Ex.2 with GMRES solver	3	2	4	2	0
Ex.1 and Ex.2 with MUMPS solver	4.5	2	4	2.25	0
Ex.3	1	1	1	2	2

Table 3 Summary of parameters values to used in numerical tests

Parameter	Value for PDE examples (Ex.1 and Ex.2)	Value for SDE example (Ex.3)
h_0	1/4 for Ex.1 , 1/8 for Ex.2	1
β	2	2
κ_0 and κ_1	0.1 for both	0.1 for both
TOL_{max}	0.5	0.1
r_1 and r_2	2 and 1.1, respectively	2 and 1.1, respectively
Initial hierarchy	$L = 2$ and $h_\ell = \{4, 6, 8\}$ and $M_\ell = 10$ for all ℓ	$L = 2$ and $h_\ell = \{1, 2, 4\}$ and $M_\ell = 10$ for all ℓ
L_{inc}	2	2
\mathcal{L}	3	5
C_α	2	2

From [30, Theorem 2.3], the complexity rate for all the examples is expected to be $\mathcal{O}(TOL^{-s_1} \log(TOL)^{s_2})$, where s_1 and s_2 depend on q_1, q_2 and $\gamma = d\tilde{\gamma}$. These and other problem parameters are summarized in Table 2 for the different examples.

We run each algorithm 100 times for each tolerance and show in plots in the next section the medians with vertical bars spanning from the 5% percentile to the 95% percentile. Finally, all results were generated on the same machine with 52 gigabytes of memory to ensure that no overhead is introduced due to hard disk access during swapping that could occur when solving the three-dimensional PDEs with a fine mesh.

In order to compare CMLMC to SMLMC, and since the latter does not include a step to fit q_1 and q_2 , we assume that these parameters are both known as discussed in Examples 2.1 and 2.2. Moreover, we use the parameters listed in Table 3.

5.3 Results

Figure 1 shows that the running time of CMLMC follows the expected complexity rates $\mathcal{O}(TOL^{s_1} \log(TOL)^{s_2})$ as summarized in Table 2. Notice that the running time in this and all figures that we present in this work include the time necessary to sample the underlying stochastic solution and the time to do the necessary computation to

estimate the problem parameters. However, the computational complexity of calculating of problem parameters is largely dominated by the computational complexity of sampling the approximate solution to the differential equations at hand. Indeed, the computations described in Sect. 4.2 are inexpensive post-processing calculations of these samples. Moreover, our results show that the algorithm has the same complexity as the theoretical work (c.f. [30, Theorem 2.3]) of the last iteration where we effectively solve the problem with the required tolerance requirements. Next, Fig. 2 shows the number of levels, L , in the last iteration of CMLMC for different tolerances. As expected, even though L depends on the particular realization, it is well approximated by a linear function of $\log(\text{TOL}^{-1})$.

Next, Fig. 3 shows the computational errors of CMLMC that were computed using the reference solutions as listed in Sect. 5.1. This indicates that the imposed accuracy is achieved with the required confidence of 95 %—since $C_\alpha = 2$. Compare this figure to Fig. 4 which shows the computational errors of SMLMC. One can see that, in certain cases, SMLMC solves the problem for a smaller tolerance than the imposed TOL. This is because θ is fixed and the statistical error is not relaxed when the bias is small. This can be especially seen in Ex.2 where the choice $h_0 = 1/8$ produces a bias much smaller than 0.5TOL for the shown tolerances. On the other hand, Fig. 5 is a QQ-plot showing that the empirical cumulative distribution function (CDF) of the MLMC error estimates is well approximated by the standard normal CDF, even for finite tolerances.

Figure 6 shows a comparison of the running time of CMLMC and SMLMC. Notice that a good value of \tilde{M} in SMLMC is not known a priori and the computational time varies considerably for different values of \tilde{M} , especially for smaller tolerances in Ex.1 and Ex.2. Specifically, a larger \tilde{M} in SMLMC increases the computational time of the algorithm, but also decreases its variability. A smaller \tilde{M} gives a smaller computational time at the expense of increased variation. The variation of the running time is due to inaccurate estimates of V_ℓ due to the smaller number of initial samples. On the other hand, the running time of CMLMC is less varied, which is a reflection of the stability of the estimates of V_ℓ . The computational savings of CMLMC over SMLMC is an aggregate effect of the different improvements. This includes (1) a more stable variance and bias estimates as already discussed, (2) a better splitting of bias and statistical tolerances. This second point can be seen in Fig. 7, which shows the tolerance splitting parameter, θ , used in CMLMC as computed by (4.2). We can clearly see here that θ is not trivial and changes with the tolerance. Looking closely, one can notice sudden jumps in the values of θ due to changes in the discrete number of levels, L . Between jumps, θ changes continuously due to inaccuracies in the estimation of the weak error constant, Q_W . Specifically, notice that for $\text{TOL} \approx 0.015$ in Ex.1 when using the direct solver, the splitting parameter θ used in CMLMC is very close to 0.5 which explains why, for this case, the computational time of SMLMC is very close to the computational time of CMLMC as shown in Fig. 6.

Finally, the bias of the MLMC estimator when using samples generated in previous iterations to compute the quantity of interest is not well understood. Using CMLMC, generating new samples at each iteration, instead of using samples from previous iterations, does not add a significant overhead to the total running time of the

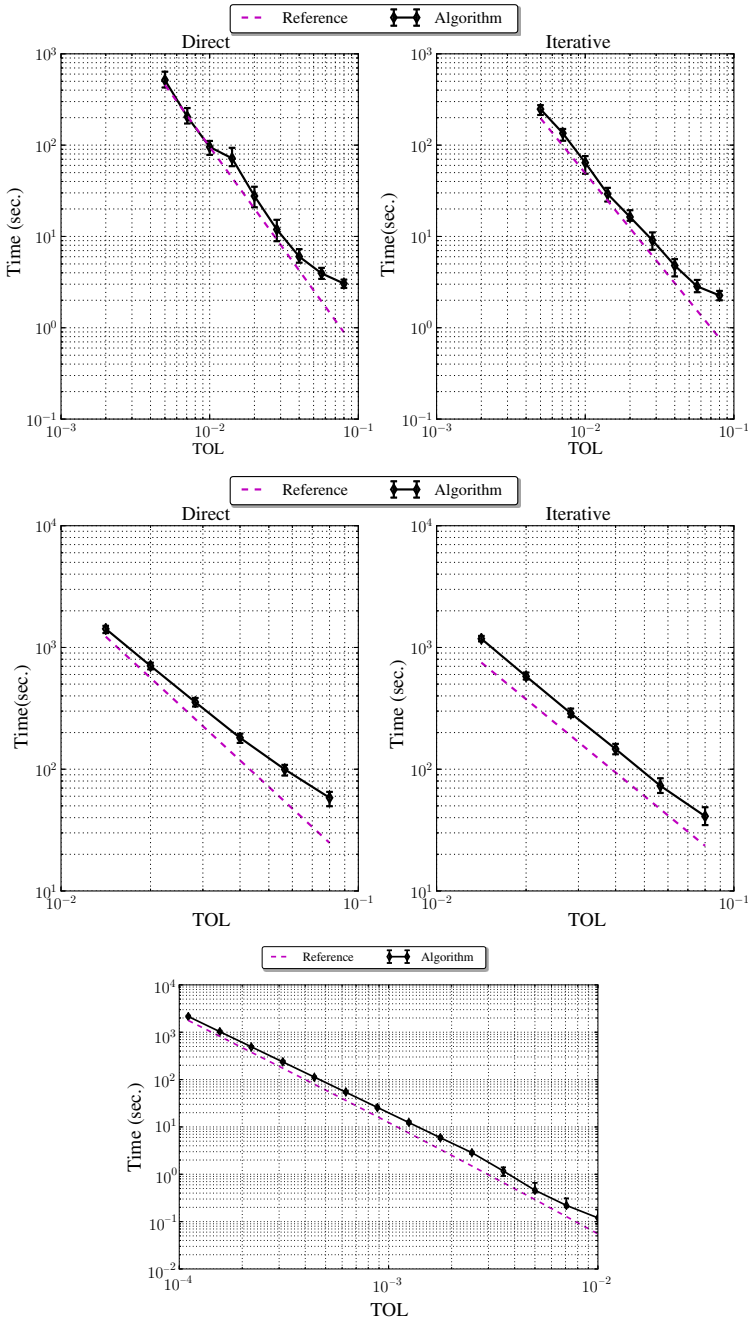


Fig. 1 From top: Ex.1, Ex.2, Ex.3. These plots show the total running time of CMLMC. The reference dashed lines are $\mathcal{O}(\text{TOL}^{-s_1} \log(\text{TOL})^{s_2})$ as summarized in Table 2. Notice that, asymptotically, the total running times seem to follow the expected rates. This shows that the algorithm, in our examples, has the same complexity as the theoretical work (c.f. [30, Theorem 2.3]) of the last iteration where we effectively solve the problem with the required tolerance requirements

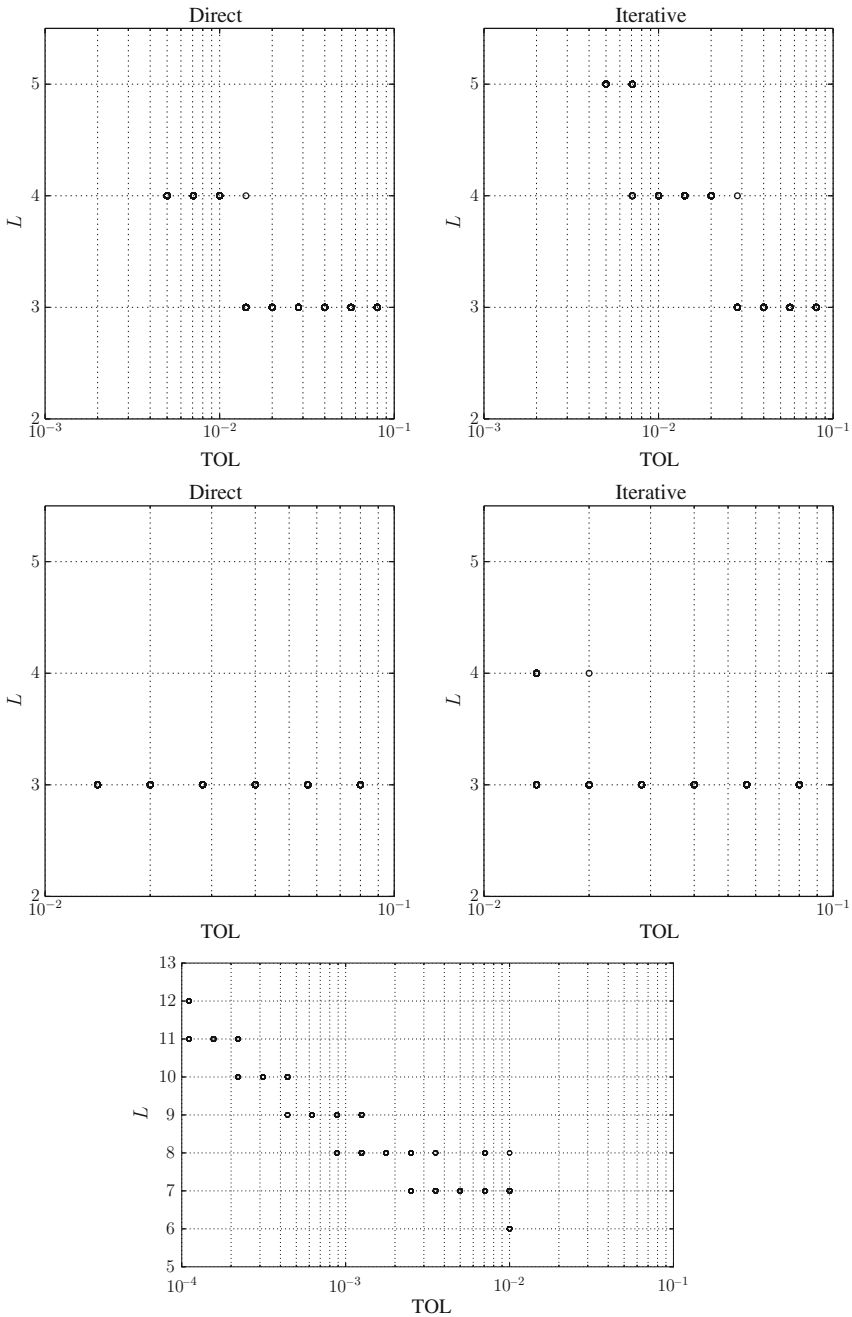


Fig. 2 From *top* Ex.1, Ex.2, Ex.3. These plots show the number of levels, L , for different tolerances, as produced in the last iteration of CMLMC. Here, it is clear that L depends on the particular realization. However, the relation between L and $\log(\text{TOL}^{-1})$ looks reasonably linear, as expected. Note that in Ex.2, L does not exhibit significant variations. This is because, for the tolerances considered, $L = 3$ already satisfies the bias constraint

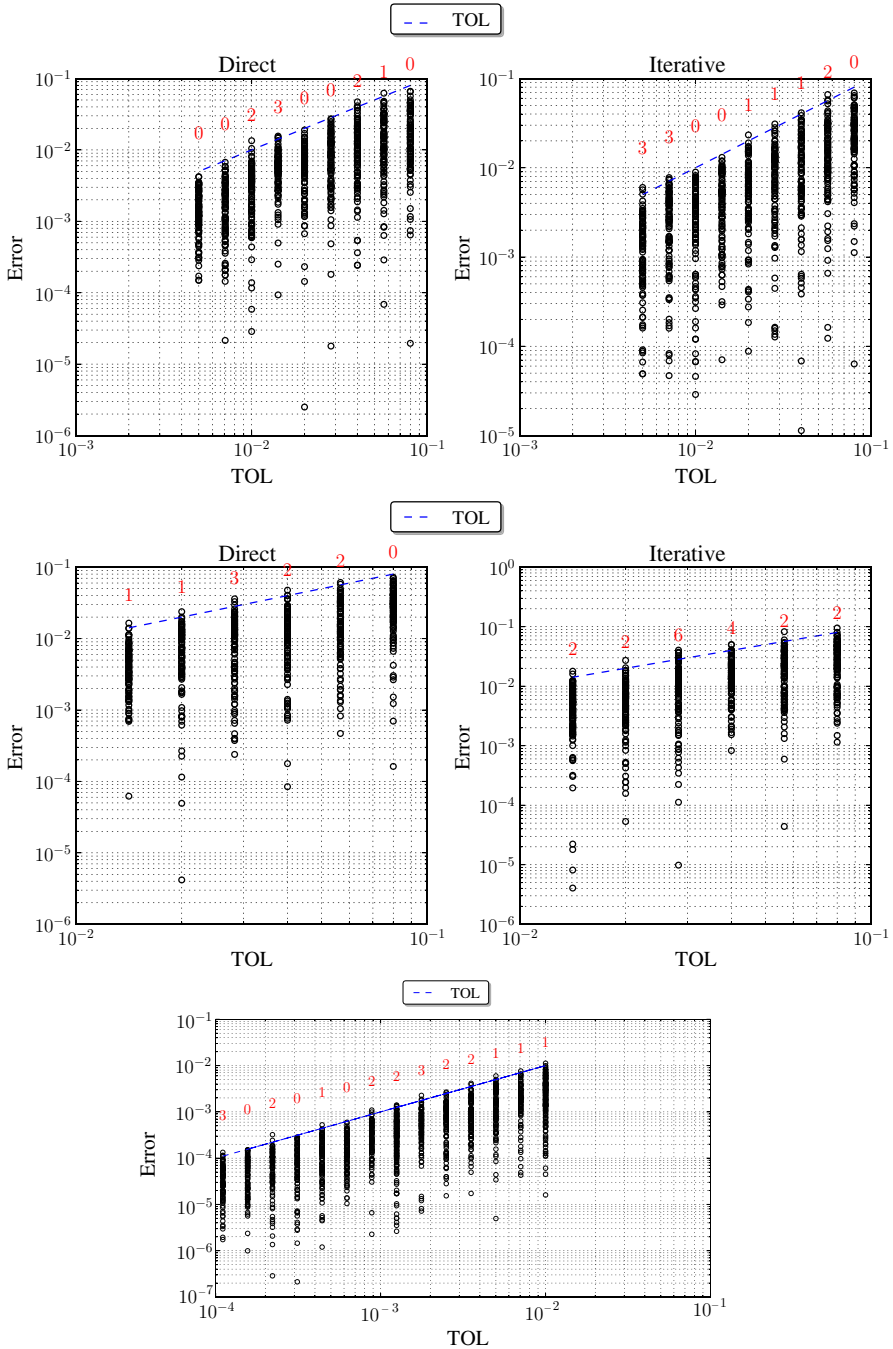


Fig. 3 From *top* Ex.1, Ex.2, Ex.3. Actual computational errors based on the reference solutions when using CMLMC. The numbers above the dashed line show the percentage of runs that had errors larger than the required tolerance. We observe that these results are consistent with the imposed error constraints with a 95 % confidence

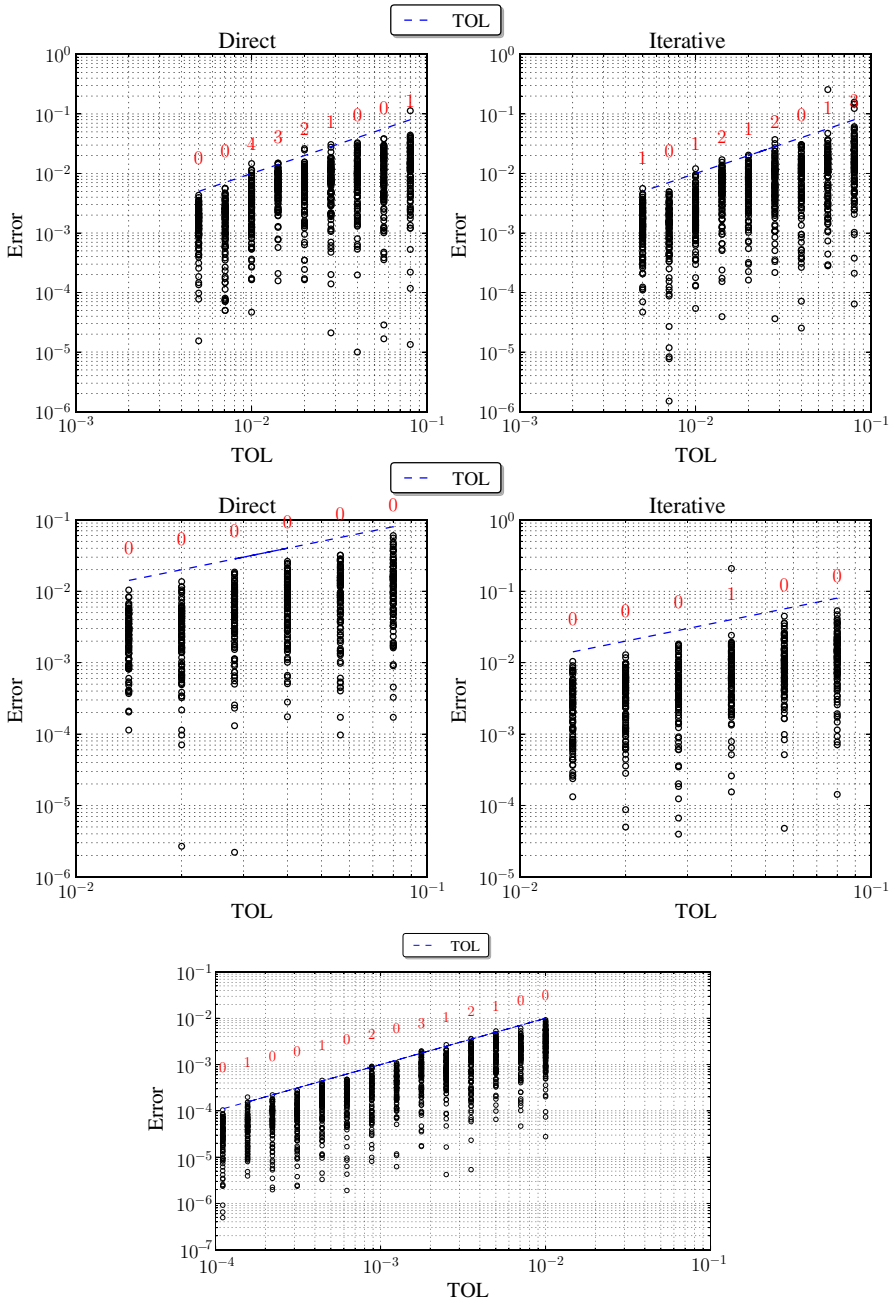


Fig. 4 From *top* Ex.1, Ex.2, Ex.3. Actual computational errors based on the reference solutions when using SMLMC. The *numbers* above the *dashed line* show the percentage of runs that had errors larger than the required tolerance. We observe that these results are consistent with the imposed error constraints with a 95 % confidence. However, for particular tolerances, the error is smaller than TOL because the statistical error is not relaxed when the bias is small since tolerance splitting parameter, θ , is kept fixed for all tolerances

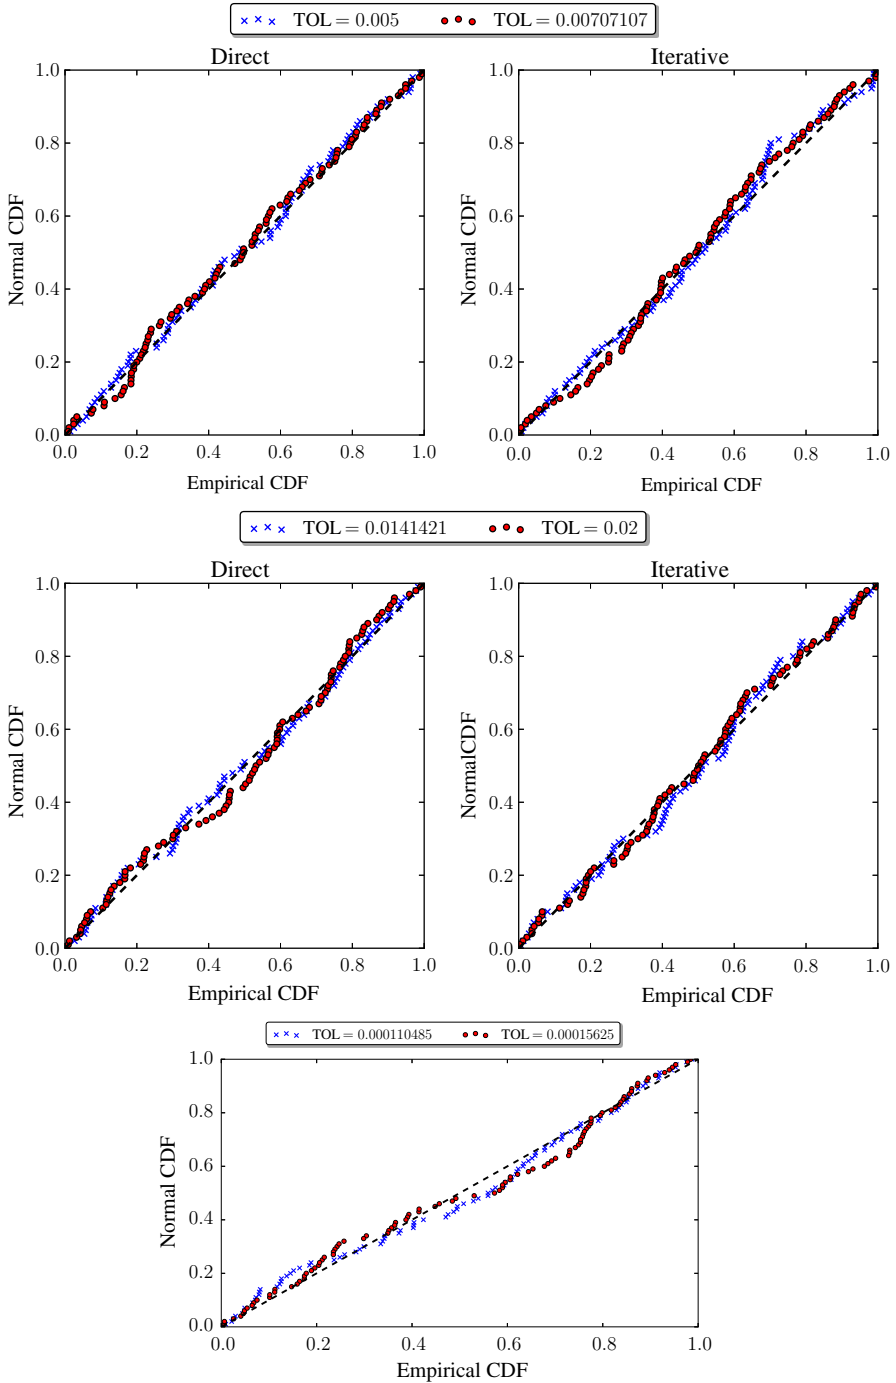


Fig. 5 From *top* Ex.1, Ex.2, Ex.3. Normalized empirical cumulative distribution function (CDF) of MLMC error estimates for different tolerances versus the standard normal CDF. Notice that, even for finite tolerances, the standard normal CDF is a good approximation of the CDF of the error estimates

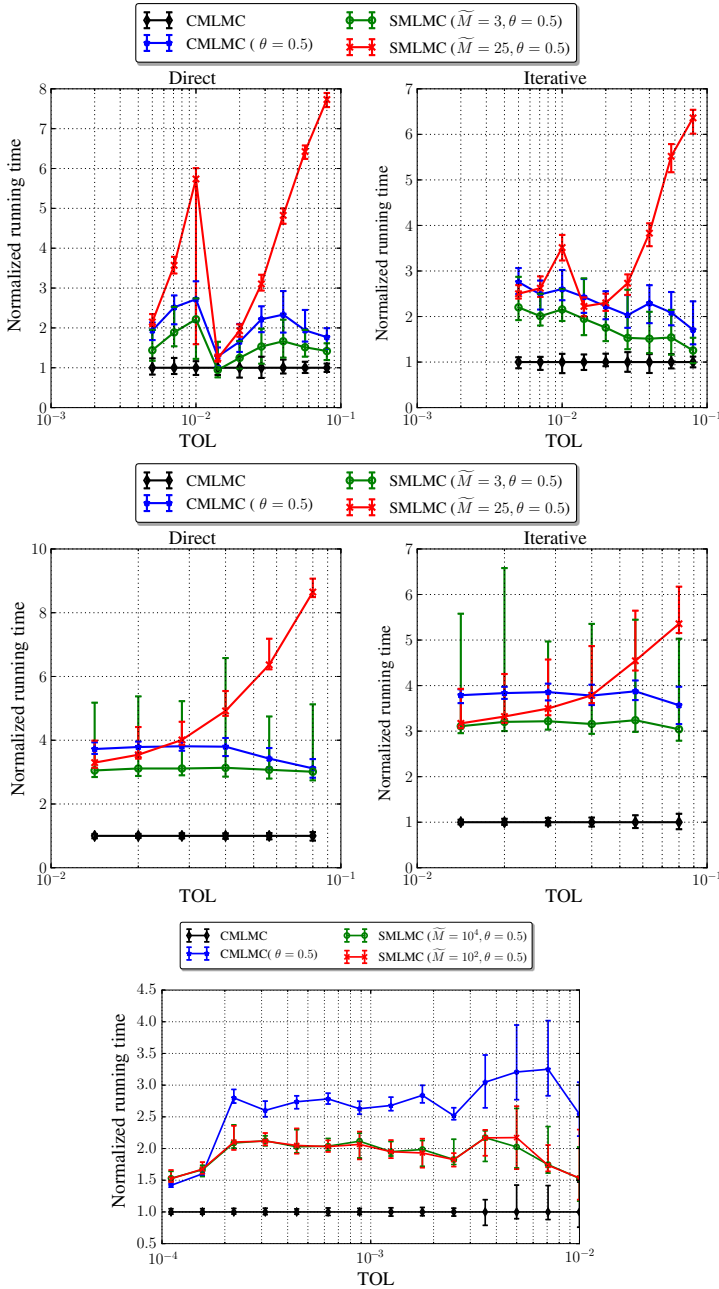


Fig. 6 From *top* Ex.1, Ex.2, Ex.3. The running time of CMLMC and SMLMC for different \tilde{M} and θ , normalized by the median running time of CMLMC. This plot shows that a larger \tilde{M} increases the median running time of the SMLMC but also decreases its variability. One sees that CMLMC outperforms SMLMC even for a small \tilde{M} in all numerical examples. Note that for Ex.1 using direct method and for TOL ≈ 0.015 , all algorithms perform similarly. This is because, for this example, the optimal error splitting parameter, θ , according to (4.2) is approximately 0.5

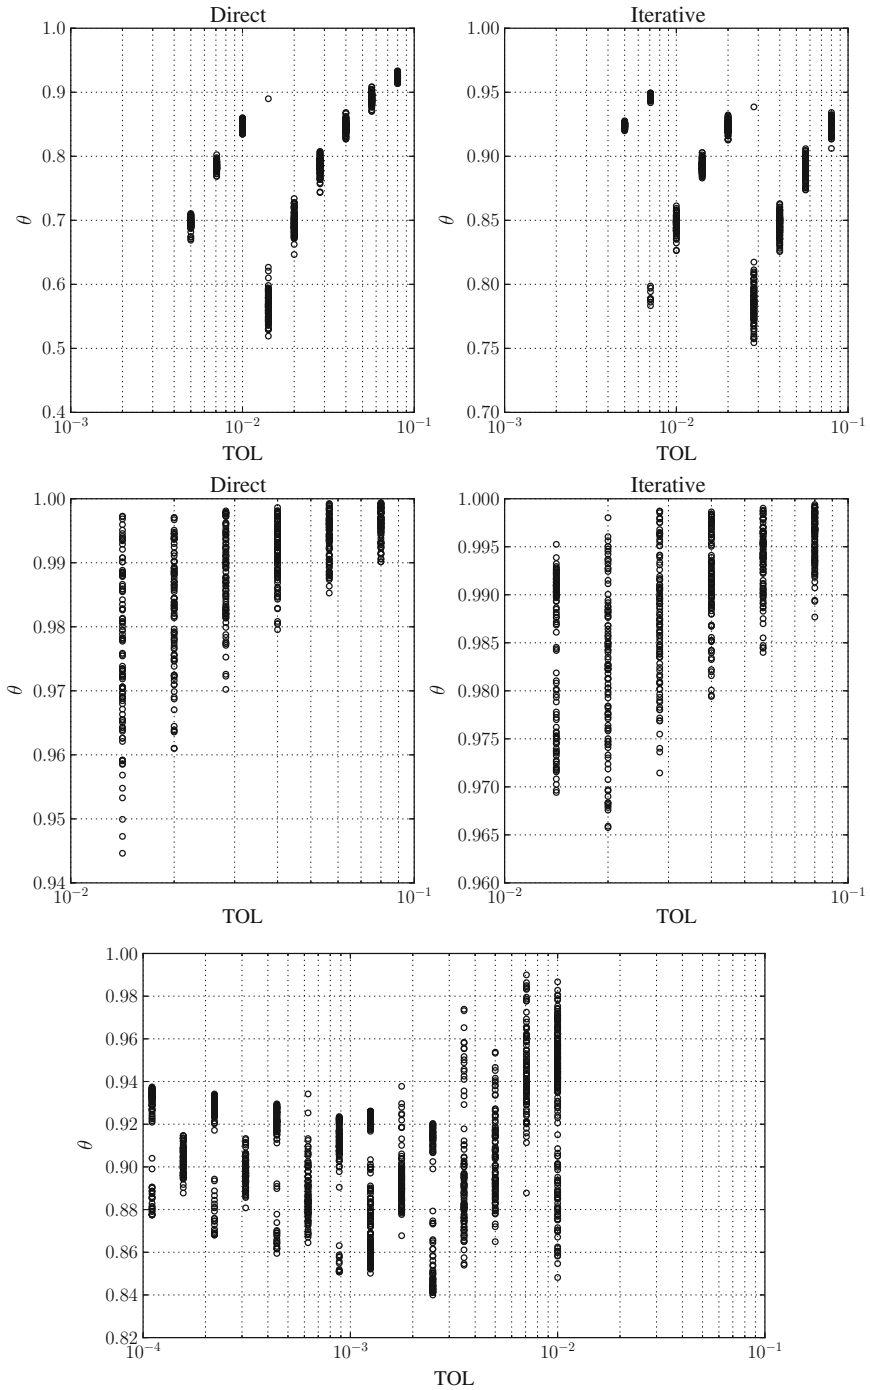


Fig. 7 From *top* Ex.1, Ex.2, Ex.3. The error splitting, θ , as computed in (4.2) and used in CMLMC, versus TOL. Observe the behavior of θ is non-trivial and can be far from $\frac{1}{2}$

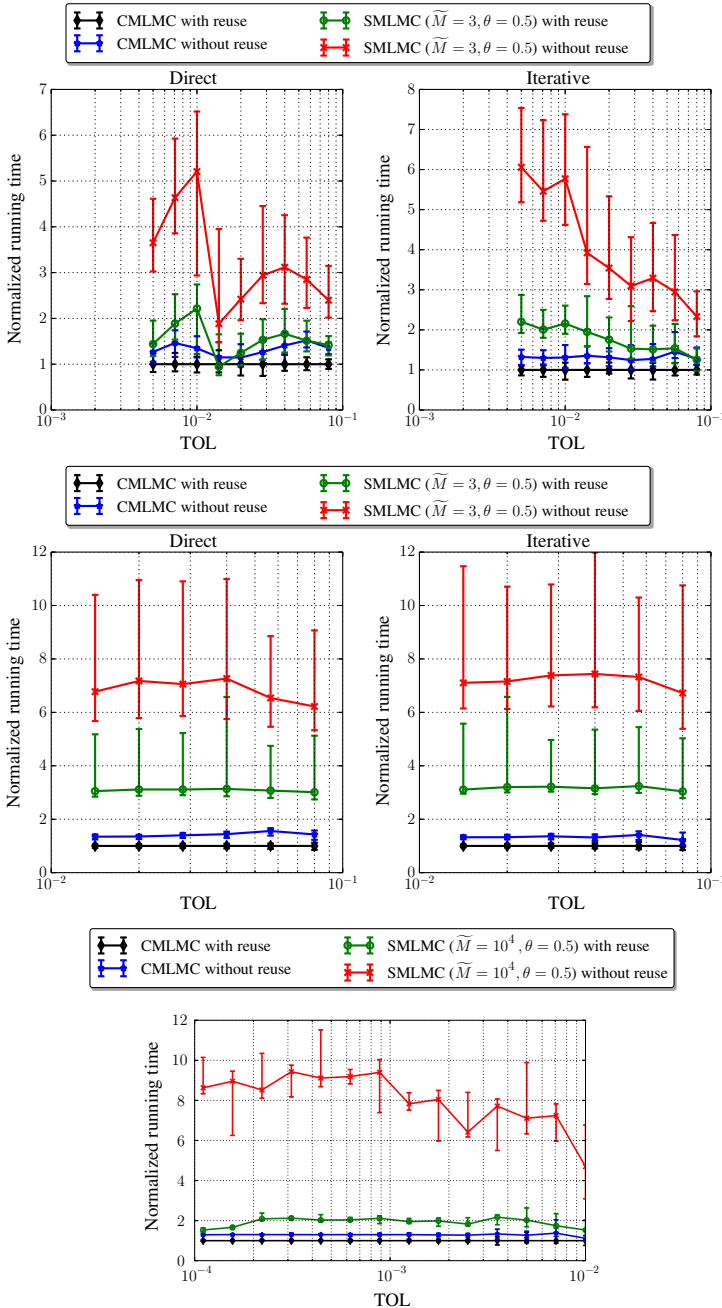


Fig. 8 From *top* Ex.1, Ex.2, Ex.3. Running time of CMLMC versus SMLMC when reusing samples for both. Also included, is CMLMC without reusing samples from previous iterations. All running times are normalized by the median of the running time of CMLMC when reusing samples. Notice that reusing samples in CMLMC does not add a significant advantage. Moreover, CMLMC still produces savings over SMLMC, even when reusing samples in the latter

algorithm. Figure 8 explains this point by comparing the running time of CMLMC for both cases for both CMLMC and SMLMC. This figure shows that computational savings of CMLMC over SMLMC whether we reuse samples or not in the former, mainly due to better splitting of the tolerance between bias and statistical errors. Moreover, it shows that reusing samples in CMLMC does not offer significant computational savings that justify the increased complexity in the analysis of the resulting estimator.

6 Conclusions

We have proposed a novel Continuation Multi Level Monte Carlo (CMLMC) algorithm for weak approximation of stochastic models. Our algorithm uses discretization hierarchies that are defined a priori for each level and are geometrically refined across levels. These hierarchies are either uniform at each level or obtained by regular subdivision of a non-uniform mesh.

The actual choice of computational work across levels uses the optimal number of samples per level given the variance and the work contribution from each level. Accurate computation of these relevant quantities is based on parametric models.

These parameters are calibrated using approximate samples, either produced before running the CMLMC and/or during the actual runs. We also propose a novel Bayesian estimation of the variance and weak convergence model parameters, taking particular notice of the deepest levels of the discretization hierarchy, where only a few realizations are available to produce the required estimates. The idea is to use results from coarser levels, where more samples are available, to stabilize the estimates in the deeper levels. The resulting MLMC estimator exhibits a non-trivial splitting between bias and statistical contributions. Indeed, the actual split depends on the given accuracy and other problem parameters. In fact, as the numerical examples show, there are cases where most of the accuracy budget is devoted to the statistical error. Finally, using the Lindeberg–Feller theorem, we also show the asymptotic normality of the statistical error in the MLMC estimator and justify in this way our error estimate that allows prescribing both required accuracy and confidence in the final result.

We presented three numerical examples to substantiate the above results, exhibiting the robustness of the new CMLMC Algorithm and to demonstrate its corresponding computational savings. The examples are described in terms of differential equations either driven by random measures or with random coefficients.

Other aspects of MLMC estimators can also be explored, such as the optimality of geometric hierarchies compared to non-geometric ones. This will be the subject of a forthcoming work, where extensions of the CMLMC to that setting will be considered.

Acknowledgments Raúl Tempone is a member of the Strategic Research Initiative on Uncertainty Quantification in Computational Science and Engineering at KAUST (SRI-UQ). The authors would like to recognize the support of King Abdullah University of Science and Technology (KAUST) AEA project “Predictability and Uncertainty Quantification for Models of Porous Media” and University of Texas at Austin AEA Round 3 “Uncertainty quantification for predictive modeling of the dissolution of porous and fractured media”. We would also like to acknowledge the use of the following open source software packages: PETSc [4], PetIGA [8], NumPy, matplotlib [21].

7 Appendix: Normality of MLMC estimator

Theorem 7.1 ([10, LindebergsFeller Theorem, p. 114]) *For each n , let $X_{n,m}$, for $1 \leq n \leq m$, be independent random variables (not necessarily identical). Denote*

$$\begin{aligned}
 a_n &= \sum_{m=1}^n X_{n,m}, \\
 Y_{n,m} &= X_{n,m} - E[X_{n,m}], \\
 s_n^2 &= \sum_{m=1}^n E[Y_{n,m}^2].
 \end{aligned}$$

Suppose the following Lindeberg condition is satisfied for all $\epsilon > 0$:

$$\lim_{n \rightarrow \infty} s_n^{-2} \sum_{m=1}^n E[Y_{n,m}^2 \mathbf{1}_{|Y_{n,m}| > \epsilon s_n}] = 0. \tag{7.1}$$

Then,

$$\lim_{n \rightarrow \infty} P\left[\frac{a_n - E[a_n]}{s_n} \leq z\right] = \Phi(z),$$

where $\Phi(z)$ is the normal cumulative density function of a standard normal random variable.

Lemma 7.1 *Consider the MLMC estimator \mathcal{A} given by*

$$\mathcal{A} = \sum_{\ell=0}^L \sum_{m=1}^{M_\ell} \frac{G_\ell(\omega_{\ell,m})}{M_\ell},$$

where $G_\ell(\omega_{\ell,m})$ denote as usual i.i.d. samples of the random variable G_ℓ . The family of random variables, $(G_\ell)_{\ell \geq 0}$, is also assumed independent. Denote $Y_\ell = |G_\ell - E[G_\ell]|$ and assume the following

$$C_1 \beta^{-q_3 \ell} \leq E[Y_\ell^2] \tag{7.2a} \quad \text{for all } \ell \geq 0,$$

$$E[Y_\ell^{2+\delta}] \leq C_2 \beta^{-\tau \ell} \tag{7.2b} \quad \text{for all } \ell \geq 0,$$

for some $\beta > 1$ and strictly positive constants C_1, C_2, q_3, δ and τ . Choose the number of samples on each level M_ℓ to satisfy, for $q_2 > 0$ and a strictly positive sequence $\{H_\ell\}_{\ell \geq 0}$

$$M_\ell \geq \beta^{-q_2 \ell} \text{TOL}^{-2} H_\ell^{-1} \left(\sum_{\ell=0}^L H_\ell \right) \quad \text{for all } \ell \geq 0. \tag{7.3}$$

Moreover, choose the number of levels L to satisfy

$$L \leq \max \left(0, \frac{c \log (\text{TOL}^{-1})}{\log \beta} + C \right) \tag{7.4}$$

for some constants C , and $c > 0$. Finally, denoting

$$p = (1 + \delta/2)q_3 + (\delta/2)q_2 - \tau,$$

if we have that either $p > 0$ or $c < \delta/p$, then

$$\lim_{\text{TOL} \rightarrow 0} P \left[\frac{\mathcal{A} - \mathbb{E}[\mathcal{A}]}{\sqrt{\text{Var}[\mathcal{A}]}} \leq z \right] = \Phi(z).$$

Proof We prove this lemma by ensuring that the Lindeberg condition (7.1) is satisfied. The condition becomes in this case

$$\lim_{\text{TOL} \rightarrow 0} \underbrace{\frac{1}{\text{Var}[\mathcal{A}]} \sum_{\ell=0}^L \sum_{m=1}^{M_\ell} \mathbb{E} \left[\frac{Y_\ell^2}{M_\ell^2} \mathbf{1}_{\frac{Y_\ell}{M_\ell} > \epsilon \sqrt{\text{Var}[\mathcal{A}]}} \right]}_{:=F} = 0,$$

for all $\epsilon > 0$. Below we make repeated use of the following identity for non-negative sequences $\{a_\ell\}$ and $\{b_\ell\}$ and $q \geq 0$.

$$\sum_{\ell} a_\ell^q b_\ell \leq \left(\sum_{\ell} a_\ell \right)^q \sum_{\ell} b_\ell. \tag{7.5}$$

First we use the Markov inequality to bound

$$\begin{aligned} F &= \frac{1}{\text{Var}[\mathcal{A}]} \sum_{\ell=0}^L \sum_{m=1}^{M_\ell} \mathbb{E} \left[\frac{Y_\ell^2}{M_\ell^2} \mathbf{1}_{Y_\ell > \epsilon \sqrt{\text{Var}[\mathcal{A}] M_\ell}} \right] \\ &\leq \frac{\epsilon^{-\delta}}{\text{Var}[\mathcal{A}]^{1+\delta/2}} \sum_{\ell=0}^L M_\ell^{-1-\delta} \mathbb{E} \left[Y_\ell^{2+\delta} \right]. \end{aligned}$$

Using (7.5) and substituting for the variance $\text{Var}[A]$ where we denote $\text{Var}[G_\ell] = \mathbb{E}[(G_\ell - \mathbb{E}[G_\ell])^2]$ by V_ℓ , we find

$$\begin{aligned}
 F &\leq \frac{\epsilon^{-\delta} \left(\sum_{\ell=0}^L M_\ell^{-1} V_\ell\right)^{1+\delta/2}}{\left(\sum_{\ell=0}^L V_\ell M_\ell^{-1}\right)^{1+\delta/2}} \sum_{\ell=0}^L V_\ell^{-1-\delta/2} M_\ell^{-\delta/2} \mathbb{E}\left[Y_\ell^{2+\delta}\right] \\
 &\leq \epsilon^{-\delta} \sum_{\ell=0}^L V_\ell^{-1-\delta/2} M_\ell^{-\delta/2} \mathbb{E}\left[Y_\ell^{2+\delta}\right].
 \end{aligned}$$

Using the lower bound on the number of samples M_ℓ (7.3) and (7.5) again yields

$$\begin{aligned}
 F &\leq \epsilon^{-\delta} \text{TOL}^\delta \left(\sum_{\ell=0}^L V_\ell^{-1-\delta/2} \beta^{\frac{\delta q_2 \ell}{2}} H_\ell^{\delta/2} \mathbb{E}\left[Y_\ell^{2+\delta}\right]\right) \left(\sum_{\ell=0}^L H_\ell\right)^{-\delta/2} \\
 &\leq \epsilon^{-\delta} \text{TOL}^\delta \left(\sum_{\ell=0}^L V_\ell^{-1-\delta/2} \beta^{(\delta/2)q_2 \ell} \mathbb{E}\left[Y_\ell^{2+\delta}\right]\right).
 \end{aligned}$$

Finally using the bounds (7.2a) and (7.2b)

$$\begin{aligned}
 F &\leq \epsilon^{-\delta} \text{TOL}^\delta \left(C_1^{-1-\delta/2} C_2 \sum_{\ell=0}^L \beta^{(1+\delta/2)q_3 \ell} \beta^{(\delta/2)q_2 \ell} \beta^{-\tau \ell}\right) \\
 &= \epsilon^{-\delta} \text{TOL}^\delta C_1^{-1-\delta/2} C_2 \frac{\beta^{(L+1)p} - 1}{\beta^p - 1},
 \end{aligned}$$

We distinguish two cases here, namely:

- If $p > 0$ is satisfied then $\lim_{\text{TOL} \rightarrow 0} F = 0$ for any $c > 0$.
- Otherwise, substituting (7.4) gives

$$F \leq \epsilon^{-\delta} \text{TOL}^\delta C_1^{-1-\delta/2} C_2 \frac{\text{TOL}^{-cp} \beta^{(C+1)p} - 1}{\beta^p - 1} = \mathcal{O}(\text{TOL}^{\delta-cp}),$$

and since in this case $cp < \delta$ then $\lim_{\text{TOL} \rightarrow 0} F = 0$. □

Remark 7.1 The choice (7.3) mirrors the choice (2.7), the latter being the optimal number of samples to bound the statistical error of the estimator by TOL. Specifically, $H_\ell \propto \sqrt{V_\ell W_\ell}$ where W_ℓ is the work per sample on level ℓ . Moreover, the choice (2.7) uses the variances $\{V_\ell\}_{\ell=0}^L$ or an estimate of it in the actual implementation. On the other hand, the choice (7.3) uses the upper bound of V_ℓ instead, if q_2 is the rate of variance convergence therein. Furthermore, if we assume the weak error model (2.9a) holds and $h_L = h_0 \beta^{-L}$ then we must have

$$Q_W h_L^{q_1} = Q_W h_0^{q_1} \beta^{-Lq_1} \leq (1 - \theta) \text{TOL},$$

which gives a lower bound on the number of levels L , namely

$$L \geq \frac{\log(\text{TOL}^{-1})}{q_1 \log(\beta)} + \frac{-\log(1 - \theta) + \log(Q_W) + q_1 \log(h_0)}{q_1 \log(\beta)},$$

to bound the bias by TOL.

Finally, in Example 2.1 the conditions (7.2) are satisfied for $q_3 = 2$ and, assuming $p^* > 3$, for $\delta = 1$ and $\tau = 6$. Similarly, Example 2.2 satisfies the conditions (7.2) are for $q_3 = 1$ and $\delta = 2$ and $\tau = 2$, cf. [20].

Remark 7.2 The assumption (7.2a) can be relaxed. For instance, one can assume instead that

$$\begin{aligned} V_{\ell+1} &\leq V_\ell \quad \text{for all } \ell \geq 1, \\ 0 &< \lim_{\ell \rightarrow \infty} \text{Var}[Y_\ell] \beta^{q_3 \ell} < \infty, \end{aligned}$$

and slightly different conditions on L .

References

1. Amestoy, P.R., Duff, I.S., L'Excellent, J.Y., Koster, J.: A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.* **23**, 15–41 (2001)
2. Amestoy, P.R., Guermouche, A., L'Excellent, J.Y., Pralet, S.: Hybrid scheduling for the parallel solution of linear systems. *Parallel Comput.* **32**(2), 136–156 (2006)
3. Babuška, I., Nobile, F., Tempone, R.: A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Rev.* **52**(2), 317–355 (2010)
4. Balay, S., Brown, J., Buschelman, K., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Smith, B.F., Zhang, H.: PETSc Web page (2013). <http://www.mcs.anl.gov/petsc>
5. Barth, A., Schwab, C., Zollinger, N.: Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients. *Numer. Math.* **119**(1), 123–161 (2011)
6. 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 J. Numer. Anal.* **51**(1), 322–352 (2013)
7. Cliffe, K., Giles, M., Scheichl, R., Teckentrup, A.: Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput. Vis. Sci.* **14**(1), 3–15 (2011)
8. Collier, N., Dalcin, L., Calo, V.: PetIGA: high-performance isogeometric analysis. (2013). <http://arxiv.org/abs/1305.4452>
9. Dalcin, L., Collier, N.: PetIGA: a framework for high performance isogeometric analysis (2013). <https://bitbucket.org/dalcin/petiga>
10. Durrett, R.: *Probability: Theory and Examples*, 2nd edn. Duxbury Press, Belmont (1996)
11. Giles, M.: Improved multilevel Monte Carlo convergence using the Milstein scheme. In: *Monte Carlo and quasi-Monte Carlo methods 2006*, pp. 343–358. Springer, Berlin (2008)
12. Giles, M.: Multilevel Monte Carlo path simulation. *Oper. Res.* **56**(3), 607–617 (2008)
13. Giles, M., Reisinger, C.: Stochastic finite differences and multilevel Monte Carlo for a class of SPDEs in finance. *SIAM J. Financ. Math.* **3**(1), 572–592 (2012)
14. Giles, M., Szpruch, L.: Antithetic multilevel Monte Carlo estimation for multidimensional SDEs. In: *Monte Carlo and Quasi-Monte Carlo Methods 2012*. Springer, Berlin (2013)
15. Giles, M., Szpruch, L.: Multilevel Monte Carlo methods for applications in finance, pp. 3–48. World Scientific, Singapore (2013)
16. Giles, M., Szpruch, L.: Antithetic multilevel Monte Carlo estimation for multi-dimensional SDEs without Lévy area simulation. *Ann. Appl. Probab.* (2014)

17. Glasserman, P.: Monte Carlo methods in financial engineering. Applications of Mathematics (New York), vol. 53. Stochastic Modelling and Applied Probability. Springer, New York (2004)
18. Heinrich, S.: Monte Carlo complexity of global solution of integral equations. *J. Complex.* **14**(2), 151–175 (1998)
19. Heinrich, S., Sindambiwe, E.: Monte Carlo complexity of parametric integration. *J. Complex.* **15**(3), 317–341 (1999)
20. Hoel, H., Schwerin, E.v., Szepessy, A., Tempone, R.: Adaptive multilevel Monte Carlo simulation. In: Engquist, B., Runborg, O., Tsai, Y.H. (eds.) Numerical Analysis of Multiscale Computations, no. 82 in Lecture Notes in Computational Science and Engineering, pp. 217–234. Springer, New York (2012)
21. Hunter, J.D.: Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**(3), 90–95 (2007)
22. Jouini, E., Cvitanic, J., Musiela, M. (eds.): Handbooks in Mathematical Finance. Option Pricing, Interest Rates and Risk Management. Cambridge University Press, Cambridge (2001)
23. Karatzas, I., Shreve, S.E.: Brownian Motion and Stochastic Calculus, Graduate Texts in Mathematics, 2nd edn, vol. 113. Springer, New York (1991)
24. Kebaier, A.: Statistical Romberg extrapolation: a new variance reduction method and applications to options pricing. *Ann. Appl. Probab.* **14**(4), 2681–2705 (2005)
25. Mordecki, E., Szepessy, A., Tempone, R., Zouraris, G.E.: Adaptive weak approximation of diffusions with jumps. *SIAM J. Numer. Anal.* **46**(4), 1732–1768 (2008)
26. Øksendal, B.: Stochastic differential equations. Universitext, 5th edn. Springer, Berlin (1998)
27. Sivia, D.S., Skilling, J.: Data Analysis.: A Bayesian Tutorial, 2nd edn. Oxford University Press, Oxford (2006)
28. Szepessy, A., Tempone, R., Zouraris, G.E.: Adaptive weak approximation of stochastic differential equations. *Commun. Pure Appl. Math.* **54**(10), 1169–1214 (2001)
29. Teckentrup, A.: Multilevel Monte Carlo methods and uncertainty quantification. PhD thesis, University of Bath (2013)
30. Teckentrup, A., Scheichl, R., Giles, M., Ullmann, E.: Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numer. Math.* **125**(3), 569–600 (2013)
31. Xia, Y., Giles, M.: Multilevel path simulation for jump-diffusion SDEs. In: Plaskota, L., Woźniakowski, H. (eds.) Monte Carlo and Quasi-Monte Carlo Methods 2010, pp. 695–708. Springer, Berlin (2012)