

Smoothed Circulant Embedding with Applications to Multilevel Monte Carlo Methods for PDEs with Random Coefficients

Anastasia Istratuca*, Aretha Teckentrup†

Abstract

We consider the computational efficiency of Monte Carlo (MC) and Multilevel Monte Carlo (MLMC) methods applied to partial differential equations with random coefficients. These arise, for example, in groundwater flow modelling, where a commonly used model for the unknown parameter is a random field. We make use of the circulant embedding procedure for sampling from the aforementioned coefficient. To improve the computational complexity of the MLMC estimator in the case of highly oscillatory random fields, we devise and implement a smoothing technique integrated into the circulant embedding method. This allows to choose the coarsest mesh on the first level of MLMC independently of the correlation length of the covariance function of the random field, leading to considerable savings in computational cost. We illustrate this with numerical experiments, where we see a saving of factor 5-10 in computational cost for accuracies of practical interest.

Keywords: PDEs with random coefficients, multilevel Monte Carlo, log-normal random fields, circulant embedding, smooth periodisation

1 Introduction

Uncertainty quantification (UQ) is a contemporary research area concerned with identifying, assessing and reducing uncertainties related to physical models, numerical algorithms, experiments and their predicted outcomes or quantities of interest. Methodical computations of uncertainties and errors in simulations, and careful exploration of how they propagate through a model and impact its expected outcome, are vital in many applications for principled risk assessments and decision making.

One of the most common approaches to modelling physical phenomena consists of partial differential equations (PDEs), which allow for computer simulations through the use of modern numerical solvers. Such models arise, for instance, in electrodynamics [Tovmasyan et al., 1994], and fluid [Howe, 1980] and quantum mechanics [Herron and Foster, 2008]. In this case, the uncertainty arises from unknown parameters whose estimation from physical or experimental data is impractical.

The specific UQ path we follow in this paper rests on modelling the unknown PDE coefficients as random parameters, so that their uncertainty can be defined in terms of statistical quantities, such as probability density functions or covariance functions. Further, this uncertainty can be tracked through the model to the output by resolving statistical estimates of various functionals related to the solution of the PDE model. These estimates are called quantities of interests, and they can take, for example, the form of expected values or variances.

*School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, King's Buildings, Edinburgh EH9 3FD, UK. Email: A.Istratuca@sms.ed.ac.uk

†School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, King's Buildings, Edinburgh EH9 3FD, UK. Email: A.Teckentrup@ed.ac.uk

An example of where this problem arises in practice is the modelling of groundwater flow. Specifically, water resources, generally comprising of ground and surface water, must be preserved free of pollution. Thus, efficient methods for modelling and forecasting the movement of impurities through aquifers, which are used as supplies for potable water, are necessary. Such impurities can contaminate the groundwater flowing beneath earth’s surface in various ways, such as carbon capture and underground storage, fracking, accidental spills or spent nuclear fuel repositories. A mathematical model for simulating groundwater flow rests on Darcy’s Law [De Marsily, 1986], where the main parameter is the hydraulic conductivity, that is, the ease with which a fluid can move through porous media or fractures under a given pressure gradient. In practice, this can only be measured at a finite, usually small, number of geographical points. However, for numerical simulations, the value of this parameter is usually required at all the points in the computational domain, which constitutes the main source of uncertainty for this problem.

There is already an abundant literature investigating PDEs with random coefficients using a variety of methods. Possible approaches include the stochastic Galerkin and stochastic collocation approaches, or sparse tensor discretisations, see e.g. the surveys [Schwab and Gittelson, 2011; Gunzburger et al., 2014; Cohen and DeVore, 2015] and the references therein. However, the computational cost of these techniques soars up when treating random field models with low regularity and short correlation length, which are of interest in this paper.

An effective approach for tackling these problems remains the Monte Carlo (MC) method. This entails generating realisations of parameter values and subsequently solving the governing equation for many such samples to approximate a specific quantity of interest. The benefits of MC are its ease of implementation, as well as the fact that the associated computational complexity does not grow with dimension [Kuo and Sloan, 2005]. Of course, these benefits are severely weakened if one takes into account that, in this context, one or more PDEs must be solved for each MC sample, the cost of which is naturally dimension-dependent. In addition, MC suffers from an extremely slow convergence rate, rendering this approach intractable.

One alternative which aims to alleviate this issue is the Multilevel Monte Carlo (MLMC) method [Heinrich, 2001; Giles, 2008]. This involves defining multiple levels of approximation that differ in computational cost. In the context of solving PDEs with random coefficients, the MLMC levels can be defined by different grid resolutions to solve the governing equation on. Thereafter, we cluster much of the computational effort in cheap estimations on a coarse mesh with low accuracy. To increase the quality of the approximation, we compute “correction terms” on finer levels with progressively higher precision. Since most of the uncertainty can be captured on the coarsest grid, comparatively fewer samples are required on the subsequent finer levels to obtain a sufficiently good estimate of the quantity of interest. This architecture of levels, which divides the work required to achieve a certain accuracy, leads to a significant reduction in the overall computational complexity of MLMC compared to standard MC.

One technique which is widely used for sampling from a random field is the Karhunen-Loève (KL) expansion [Kosambi, 1943]. This has been extensively studied both theoretically and numerically for PDEs with log-normal random input with MLMC [Cliffe et al., 2011; Teckentrup, 2012; Charrier, 2012; Teckentrup et al., 2013; Gittelson et al., 2013; Charrier et al., 2013; Giles, 2013; Schwab and Stein, 2023], as well as with Quasi-Monte Carlo methods (QMC) and Multilevel Quasi-Monte Carlo methods (MLQMC), see, for example, the survey [Kuo and Nuyens, 2018] and the references therein. Further alterations of this algorithm to accommodate the difficulties arising from random field models with short correlation length have been proposed in [Teckentrup et al., 2013; Gittelson et al., 2013; Schwab and Stein, 2023]. Other sampling methods have also been studied in the context of MLMC or MLQMC for PDEs with random coefficients, such as wavelet reconstruction approach [Bachmayr et al., 2018] or the white noise sampling method [Crocì et al., 2021].

In this paper, we focus on the circulant embedding technique for sampling from the random

parameter. While this offers a discrete representation of a random field on a given mesh, it is both exact at these grid points, and computationally efficient. This is in comparison with the KL-expansion, which, once truncated, offers a continuous approximation whose rate of convergence is intrinsically linked to the decay of the eigenvalues of the covariance function’s kernel, which has a direct impact on the computational cost. The circulant embedding method has been integrated with MLMC and MLMQMC in e.g. [Graham et al., 2011; Kuo et al., 2017].

The main novelty in this paper is adjusting the circulant embedding method to produce smooth approximations on coarse meshes of samples from extremely oscillatory random fields, i.e. with small correlation length, high variance and low regularity. For the application in groundwater flow outlined above, this problem usually arises when we consider large computational domains where the correlation length is considerably smaller compared to the size of the domain. This is of particular interest in MLMC for PDEs with random coefficients, as the computational gains prompted by MLMC rest on using coarse meshes, which is usually not feasible for the problem at hand. A similar algorithm has been proposed in [Sawko and Zimoń, 2022], albeit for random fields with large-scale fluctuations which still require a fine grid representation. In particular Sawko et al. focus on deriving low rank approximations of the covariance matrix of the random field, and do not provide a theoretical analysis of the corresponding error.

The novel contributions of this work can be summarised as follows:

- We adapt the circulant embedding method to sample from smoothed versions of a Gaussian field with short correlation length, and prove a bound on the introduced smoothing error. This method is not specific to applications in MLMC, but is a general method for obtaining smoothed realisations of highly oscillatory Gaussian fields which can accurately be represented on coarse grids.
- We integrate the smoothed circulant embedding method into MLMC, derive an optimal coupling between the smoothing and grid sizes on each level, and prove a complexity theorem for the resulting MLMC estimator.
- We demonstrate the gains in computational cost prompted by introducing smoothing into MLMC in a simple test problem in groundwater flow, where we observe a saving of factor 5-10 for accuracies of practical interest.

The remainder of this manuscript is organised as follows. Section 2 offers an overview of circulant embedding (CE) methods, as well as the proposed smoothing technique (CES), together with the corresponding error analysis. Further, Section 3 outlines key aspects of the Monte Carlo and Multilevel Monte Carlo approaches to uncertainty quantification, with a focus on their application to PDEs with random input and their integration with CE and CES. The paper concludes with Section 4, which summarises numerical experiments aimed at illustrating the potential of the smoothing procedure, particularly when combined with MLMC. We provide a short discussion and conclusion in Section 5.

2 Circulant Embedding with Smoothing

In this section, we begin with a brief review of the standard circulant embedding method for sampling from a stationary Gaussian random field. This relies on embedding the covariance matrix of the random field in a larger matrix which has a block circulant structure [Dietrich and Newsam, 1993], and is computationally very cheap because it exploits the Fast Fourier Transform (FFT) algorithm [Frigo and Johnson, 2005] for factorising the embedding matrix. We then briefly outline the smooth periodisation approach for ensuring the positive definiteness of the embedding matrix [Bachmayr et al., 2018]. Finally, we describe the smoothing technique for highly oscillatory random fields.

For simplicity, we shall write $a \lesssim b$ if a/b is uniformly bounded independent of any parameters for two positive quantities a and b . In addition, we write $a \simeq b$ if $a \lesssim b$ and $b \lesssim a$.

2.1 Circulant Embedding

Suppose we are interested in sampling from a stationary Gaussian field Z on a domain $D \subset \mathbb{R}^d$ and probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with continuous covariance function $r(\mathbf{x}, \mathbf{y})$, so that:

$$\mathbb{E}[Z(\mathbf{x}, \cdot)] = 0 \quad \text{and} \quad \mathbb{E}[Z(\mathbf{x}, \cdot)Z(\mathbf{y}, \cdot)] = r(\mathbf{x}, \mathbf{y}) = C(\mathbf{x} - \mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y} \in D, \quad (2.1)$$

where $C : \mathbb{R}^d \rightarrow \mathbb{R}$ is a suitably chosen function. For simplicity, we set the mean of Z to 0 throughout this work, but we note that everything can be easily extended to non-zero mean $m(\mathbf{x}) : D \rightarrow \mathbb{R}$.

Further, given a zero-mean Gaussian vector $\mathbf{Z} \in \mathbb{R}^M$ representing the field Z at M locations in D , with positive definite covariance matrix $R \in \mathbb{R}^{M \times M}$ so that $R = \mathbb{E}[\mathbf{Z}\mathbf{Z}^T]$, a sample from this discrete representation of the random field Z can be generated from any factorisation of its covariance matrix R of the form:

$$R = \Theta\Theta^T, \quad \Theta \in \mathbb{R}^{M \times M}.$$

In fact, for any vector $\boldsymbol{\xi} := (\xi_1(\omega), \xi_2(\omega), \dots, \xi_M(\omega))^T$ of independent standard Gaussian random variables, we can simply take:

$$\mathbf{Z} := \Theta\boldsymbol{\xi}. \quad (2.2)$$

This defines a suitable realisation from the discrete representation of the random field Z , since we can easily check that Eq. (2.2) implies $\mathbb{E}[\mathbf{Z}\mathbf{Z}^T] = R$, as required.

There are multiple techniques to factorising the covariance matrix R of the random field Z , such as the Cholesky decomposition. However, such commonly used algorithms have computational complexity $\mathcal{O}(n^3)$ [Higham, 2002]. On the other hand, the circulant embedding method provides a fast approach to factorising R by first embedding it in a larger matrix with circulant structure, which can then be factorised using the FFT, with complexity $\mathcal{O}(n \log_2 n)$ [Davis, 1979]. To describe this procedure, let us consider the one- and two-dimensional cases individually. Higher dimensional cases follow analogously.

2.1.1 One-Dimensional Case

Suppose now that Z is a stationary line Gaussian process, with zero mean and covariance function $r(x, y) := C(x - y)$, as described in section 2.1. Further, suppose that the domain is $D = [0, 1]$, which is discretised using a grid \mathcal{T} consisting of $m + 1$ equispaced sampling points:

$$\mathcal{T} = \{x_k = k\Delta x\}_{k=0}^m = \left\{x_k = \frac{k}{m}\right\}_{k=0}^m,$$

where $\Delta x = \frac{1}{m}$ is the mesh size. The advantage of considering equidistant points is that it allows us to characterise the entries in the covariance matrix R as follows:

$$R_{j,k} = C(x_j - x_k) = C(|j - k|\Delta x) = C\left(\frac{|j - k|}{m}\right) =: C_{|j-k|},$$

as $C((j - k)\Delta x) = C((k - j)\Delta x) = C(|j - k|\Delta x)$, $\forall j, k = 0, \dots, m$. Thus, the covariance matrix R has the following form:

$$R = \begin{bmatrix} C_0 & C_1 & C_2 & \dots & C_m \\ C_1 & C_0 & C_1 & \dots & C_{m-1} \\ C_2 & C_1 & C_0 & \dots & C_{m-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ C_m & C_{m-1} & C_{m-2} & \dots & C_0 \end{bmatrix}.$$

In other words, R is a *symmetric Toeplitz* matrix, which we denote by $\text{Toeplitz}(C_0, \dots, C_m)$. We then embed R into a larger circulant matrix S by mirroring its components, i.e. $S = \text{Toeplitz}(C_0, \dots, C_m, C_{m-1}, \dots, C_1)$, so that $S \in \mathbb{R}^{2m \times 2m}$.

Under the assumption that the symmetric matrix S is positive definite, we can diagonalise it using the Fourier transform, yielding the following eigenvalue decomposition [Barnett, 1990]:

$$S = F\Lambda F^* = F^*\Lambda F. \quad (2.3)$$

Here, $F \in \mathbb{C}^{2m \times 2m}$ is the one-dimensional unitary discrete Fourier matrix, and $F^* \in \mathbb{C}^{2m \times 2m}$ is its complex conjugate transpose, whose entries are given by:

$$F_{p,q} = \frac{1}{\sqrt{2m}} \exp\left(\frac{2\pi i(p-1)(q-1)}{2m}\right), \quad p, q = 1, \dots, 2m. \quad (2.4)$$

Further, Λ is the diagonal matrix of eigenvalues of S , which are given by:

$$\lambda = \sqrt{2m}Fs,$$

where $\mathbf{s} \in \mathbb{R}^{2m}$ is the first row of the matrix S .

Returning to the problem at hand, recall that we can now factorise the circulant matrix S as in Eq. (2.3). In some applications, such as quasi-Monte Carlo methods for PDEs with random coefficients [Graham et al., 2011], we require a real factorisation of S . Thus, to avoid complex values, we convert the Fourier factorisation into a Hartley transform [Hartley, 1942] by combining the real and imaginary parts of F , denoted by $\Re(F)$ and $\Im(F)$, respectively. This yields $S = G\Lambda G^T$, where $G \in \mathbb{R}^{2m \times 2m}$ is given by $G = \Re(F) + \Im(F)$, as proven in [Graham et al., 2011, Lemma 4].

Further, once we obtain a factorisation of the form $S = G\Lambda G^T$, we can take $\boldsymbol{\xi} \sim \mathcal{N}(0, I)$ to obtain $\mathbf{Z} = (G\Lambda^{\frac{1}{2}}\boldsymbol{\xi})_R$ a sample from the random field [Graham et al., 2011, Corollary 3]. Here, the notation $(G\Lambda^{\frac{1}{2}}\boldsymbol{\xi})_R$ represents the entries in this vector that correspond to the location of the matrix R in the embedding matrix S . In this case, it corresponds to the first $m+1$ entries in $G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}$.

2.1.2 Two-Dimensional Case

In this section, we let Z be a stationary Gaussian process, with mean 0 and covariance function $r(\mathbf{x}, \mathbf{y}) := C(\mathbf{x} - \mathbf{y})$, as characterised in Section 2.1. In two dimensions, the circulant embedding method can be extended to a rectangular domain D discretised uniformly. For simplicity, let us assume that the domain is the unit square $D = [0, 1]^2$, and that its corresponding two-dimensional grid \mathcal{T} is formed by $(m_1 + 1) \times (m_2 + 1)$ points of the form:

$$\mathbf{x}_{p_1, p_2} = (x_{p_1}, y_{p_2}),$$

where $x_{p_1} = p_1\Delta x$ and $y_{p_2} = p_2\Delta y$ with $p_i = 0, \dots, m_i$, $i = 1, 2$. Here, $\Delta x = \frac{1}{m_1}$ and $\Delta y = \frac{1}{m_2}$ denote the horizontal and vertical mesh sizes used in discretising D , respectively.

Suppose now that the points are ordered lexicographically first in the x direction, and then in the y direction. Stationarity and equispaticity ensure that the entries in the covariance matrix R can be expressed as follows:

$$R_{(p_1, p_2), (q_1, q_2)} = C\left(\begin{bmatrix} x_{p_1} \\ y_{p_2} \end{bmatrix} - \begin{bmatrix} x_{q_1} \\ y_{q_2} \end{bmatrix}\right) = C\left(\begin{bmatrix} \frac{|p_1 - q_1|}{m_1} \\ \frac{|p_2 - q_2|}{m_2} \end{bmatrix}\right) =: C_{|p_1 - q_1|, |p_2 - q_2|},$$

using the symmetry of C as in the one-dimensional case, for $p_i, q_i = 0, \dots, m_i$, $i = 1, 2$. Thus, the covariance matrix is a *symmetric block Toeplitz* matrix. Specifically, it has the following form:

$$R = \text{Toeplitz}(C_{0,0}, C_{1,0}, \dots, C_{m_1,0}, C_{0,1}, C_{1,1}, \dots, C_{m_1,1}, C_{0,2} \dots C_{m_2, m_2}).$$

Hence, the covariance matrix is actually block Toeplitz and block symmetric, with blocks which are, in turn, Toeplitz and symmetric. In other words, the covariance matrix is given by:

$$R = \text{Toeplitz}(R_0, R_1, \dots, R_{m_2}),$$

where each block R_j , $j = 0, \dots, m_2$, is a Toeplitz matrix of dimension $(m_1 + 1) \times (m_1 + 1)$:

$$R_j = \text{Toeplitz}(C_{0,j}, C_{1,j}, \dots, C_{m_1,j}).$$

We can now follow a similar procedure to the method introduced in Section 2.1.1 for embedding R into a circulant matrix. Accordingly, we first embed each block R_j , $j = 0, \dots, m_2$, into a circulant matrix S_j of dimension $2m_1 \times 2m_1$. Subsequently, the resulting blocks are embedded into a large circulant matrix S of dimension $4m_1m_2 \times 4m_1m_2$, so that S becomes a block circulant matrix with circulant blocks. The embedding matrix S can then be characterised as:

$$S = \text{Toeplitz}(S_0, S_1, \dots, S_{m_2}, S_{m_2-1}, \dots, S_1),$$

where each block S_j is itself circulant:

$$S_j = \text{Toeplitz}(C_{0,j}, C_{1,j}, \dots, C_{m_1,j}, C_{m_1-1,j}, \dots, C_{1,j}), \quad j = 0, \dots, m_2.$$

As S is a circulant matrix, it may be diagonalised via the two-dimensional discrete Fourier transform [Barnett, 1990]. This is obtained by applying the one-dimensional Fourier transform, as described in the previous section, to each of the two coordinate directions in turn. Thus, we have:

$$S = F\Lambda F^* = F^*\Lambda F,$$

where F is the unitary two-dimensional Fourier matrix, whose entries are given by:

$$F_{(p_1, p_2), (q_1, q_2)} = \frac{1}{\sqrt{4m_1m_2}} \exp\left(\frac{2\pi i (p_1 - 1)(q_1 - 1)}{2m_1}\right) \exp\left(\frac{2\pi i (p_2 - 1)(q_2 - 1)}{2m_2}\right),$$

with $p_i, q_i = 1, \dots, 2m_i$, $i = 1, 2$. Here, F^* and Λ are as in Section 2.1.1. Since S is circulant, it is again uniquely defined by its first row (or column), and so its eigenvalues are specified by:

$$\boldsymbol{\lambda} = \sqrt{4m_1m_2} F \mathbf{s},$$

where \mathbf{s} is the first row of the matrix S .

In an analogous fashion to the previous section, we can modify the obtained Fourier factorisation into a Hartley transform. The corresponding argument is given in [Graham et al., 2011, Lemma 5]. Similar to the one-dimensional case, $\mathbf{Z} = (G\Lambda^{\frac{1}{2}}\boldsymbol{\xi})_R$, for $\boldsymbol{\xi} \sim \mathcal{N}(0, I)$, gives a sample from the Gaussian random field. Here, $(G\Lambda^{\frac{1}{2}}\boldsymbol{\xi})_R$ corresponds to selecting the first $m_1 + 1$ row entries in the first $m_2 + 1$ columns of $G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}$.

2.2 Smooth Periodisation

There are instances when the circulant embedding matrix S is not positive definite, depending generally on the covariance function used. As such, the Fourier diagonalisation cannot be applied directly. To ensure the positive definiteness of S , Bachmayr et al. suggest in [Bachmayr et al., 2020] periodically extending the covariance function of the Gaussian field and padding the resulting covariance matrix before embedding it to obtain a circulant structure. Here, the analysis is developed for the case when the function C is the Matérn covariance, given by:

$$C(\mathbf{t}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}\|\mathbf{t}\|_2}{\lambda}\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}\|\mathbf{t}\|_2}{\lambda}\right), \quad (2.5)$$

where $\nu > 0$ is the smoothness exponent, $\lambda > 0$ denotes the length scale, $\sigma^2 > 0$ is the marginal variance and K_ν stands for the modified Bessel function of the second kind. Without loss of generality, we assume that the sampling domain D is contained in $[0, 1]^d$, so that $C : [-1, 1]^d \rightarrow \mathbb{R}$.

Specifically, Bachmayr et al. develop two types of periodisation involving a non-smooth and smooth truncation of the covariance function, respectively. The former intrinsically depends on the choice of discretisation parameters m_i , $i = 1, \dots, d$, and length scale λ , so that the analysis only holds for $\frac{1}{m_i} \ll \lambda$, $i = 1, \dots, d$, which is not always satisfied in practice. On the other hand, the smooth periodisation does not involve such an assumption, and, for this reason, this is the method we use in the analysis and numerical examples below.

Typically, there are two elements required in the periodisation approach: the padding of the covariance matrix and the choice of periodic covariance function C^{ext} . The padding technique was first introduced by Graham et al. in [Graham et al., 2011] to ensure the positive definiteness of the embedding matrix S . The procedure involves extending the covariance matrix, i.e. “padding” it, before mirroring its components. For example, in the two-dimensional case, the extended covariance matrix has the following form:

$$R = \text{Toeplitz}(R_0, R_1, \dots, R_{m_2}, R_{m_2+1}, \dots, R_{m_2+J_2}),$$

where each R_j , $j = 0, \dots, m_2 + J_2$, is a Toeplitz matrix of dimension $(m_1 + J_1 + 1) \times (m_1 + J_1 + 1)$:

$$R_j = \text{Toeplitz}(C_{0,j}, C_{1,j}, \dots, C_{m_1,j}, C_{m_1+1,j}, \dots, C_{m_1+J_1,j}),$$

where $J_1, J_2 \in \mathbb{N}_0$ denote the padding parameters in the x and y directions, respectively. Upon mirroring, this yields the following structure on the embedding matrix:

$$S = \text{Toeplitz}(S_0, S_1, \dots, S_{m_2}, S_{m_2+1}, \dots, S_{m_2+J_2}, S_{m_2+J_2-1}, \dots, S_1),$$

where each block S_j is itself circulant:

$$S_j = \text{Toeplitz}(C_{0,j}, C_{1,j}, \dots, C_{m_1,j}, C_{m_1+1,j}, \dots, C_{m_1+J_1,j}, C_{m_1+J_1-1,j}, \dots, C_{1,j}).$$

Note that with slight abuse of notation, we still denote the extended covariance and embedded covariance matrices by R and S , respectively. This is to unify and simplify notation.

Next, the padding values are obtained by evaluating the periodic covariance function C^{ext} at the points (i, j) , with $i = m_1 + 1, \dots, m_1 + J_1, \dots, m_1 + 1$ and $j = m_2 + 1, \dots, m_2 + J_2, \dots, m_2 + 1$. As explained in [Bachmayr et al., 2020], for the periodisation technique, a 2ℓ -periodic extension of C can be constructed by first choosing a cutoff function φ satisfying:

$$\varphi = 1 \quad \text{on} \quad [-1, 1] \quad \text{and} \quad \varphi = 0 \quad \text{on} \quad \mathbb{R} \setminus [-\kappa, \kappa], \quad \kappa := 2\ell - 1 \geq \ell,$$

where $\text{supp}(C) = [-1, 1]^d$ and $\ell = 1 + J/m \geq 1$ with J padding and m discretisation parameters. Here, we assume, for simplicity, that $J \equiv J_i$ and $m \equiv m_i$, for $i = 1, \dots, d$. Then, the periodic extension C^{ext} is given by:

$$C^{\text{ext}}(\mathbf{x}) = \sum_{\mathbf{n} \in \mathbb{Z}^d} (C\varphi_\kappa)(\mathbf{x} + 2\ell\mathbf{n}), \quad \mathbf{x} \in \mathbb{R}^d, \quad (2.6)$$

where $\varphi_\kappa(\mathbf{x}) := \varphi(\|\mathbf{x}\|_\infty)$. For the smooth periodisation case, φ must also satisfy the following:

$$\ell > 1 \quad \text{and} \quad \varphi \in C_0^\infty(\mathbb{R}), \quad \text{with} \quad \text{supp}(\varphi) = [-\kappa, \kappa].$$

An example of such a cutoff function which is easy to implement in practice is given by:

$$\varphi(t) = \frac{\eta\left(\frac{\kappa-|t|}{\kappa-1}\right)}{\eta\left(\frac{\kappa-|t|}{\kappa-1}\right) + \eta\left(\frac{|t|-1}{\kappa-1}\right)},$$

where:

$$\eta(x) = \begin{cases} \exp(-x^{-1}), & x > 0, \\ 0, & x \leq 0. \end{cases}$$

In addition, to ensure that $C^{\text{ext}} = C$ on $[-1, 1]^d$ we require [Bachmayr et al., 2020]:

$$\ell \geq \frac{\kappa + \sqrt{d}}{2}.$$

Further, Theorem 10 in [Bachmayr et al., 2020] gives the following lower bound on κ such that the embedding matrix is symmetric positive definite:

$$\frac{\kappa}{\lambda} \geq C_1 + C_2 \max \left\{ \nu^{\frac{1}{2}} (1 + |\ln(\nu)|), \nu^{-\frac{1}{2}} \right\}. \quad (2.7)$$

2.3 Eigenvalue Decay

A key factor in the study of the proposed smoothing technique is the decay rate of the circulant embedding eigenvalues $\{\lambda_j\}_{j=1}^s$ in the factorisation $S = GAG^T$. Accordingly, we exploit the remark in [Bachmayr et al., 2018] which establishes that the aforementioned eigenvalues are intricately linked to the exact Karhunen–Loève (KL) eigenvalues a_j of the covariance operator with kernel C defined on $L^2(D)$. These emerge in the KL expansion of the continuous random field Z , given by:

$$Z(\mathbf{x}, \omega) = \sum_{j=1}^{\infty} \sqrt{a_j} b_j(\mathbf{x}) \xi_j(\omega), \quad (2.8)$$

where $\{\xi_j\}_{j \in \mathbb{N}} \sim \mathcal{N}(0, 1)$ i.i.d., and $\{b_j\}_{j \in \mathbb{N}}$ are the $L^2(D)$ -orthonormal eigenfunctions of the covariance operator.

In particular, Bachmayr et al. explain in [Bachmayr et al., 2018] that the ordered eigenvalues of the (extended) embedding matrix S in the smooth periodisation case decay at the same rate as the exact KL eigenvalues. This is proven for a general class of covariance functions in [Bachmayr et al., 2018, Theorem 3.1], however here we focus on the Matérn case.

Lemma 1. *Let C be the Matérn covariance (2.5), and suppose ν, λ and κ satisfy Eq. (2.7). Let $\{\lambda_j\}_{j=1}^s$ denote the eigenvalues of the (extended) embedding matrix $S \in \mathbb{R}^{s \times s}$, with $s = \prod_{i=1}^d 2(m_i + J_i)$. Then for $j = 1, \dots, s$:*

$$\lambda_j \lesssim j^{-1 - \frac{2\nu}{d}}. \quad (2.9)$$

In our numerical experiments below, we also consider the separable exponential covariance, given by:

$$C(\mathbf{t}) := \sigma^2 \exp\left(\frac{-\|\mathbf{t}\|_p}{\lambda}\right), \quad (2.10)$$

with $p = 1$. Here, $\|\cdot\|_p$ denotes the ℓ_p norm on \mathbb{R}^d , and the parameters σ^2 and λ stand for the marginal variance and the correlation length of C , respectively. In this case, the eigenvalues used in the circulant embedding procedure are, in fact, real and positive. What is more, the decay rate of the aforementioned eigenvalues can be derived from the Matérn case outlined above, as shown in the following lemma:

Lemma 2. *Let C be the separable exponential covariance (2.10), and let $\{\lambda_j\}_{j=1}^s$ denote the eigenvalues of the embedding matrix $S \in \mathbb{R}^{s \times s}$, with $s = 2^d \prod_{i=1}^d m_i$. Then for $j = 1, \dots, s$:*

$$\lambda_j \in \mathbb{R} \quad \text{and} \quad \lambda_j > 0, \quad (2.11)$$

and

$$\lambda_j \lesssim j^{-2}. \quad (2.12)$$

Proof. First, Eq. (2.11) is a direct application of Theorem 2 in [Dietrich and Newsam, 1997]. Further, this implies that no padding is necessary, as the embedding matrix is symmetric positive definite.

Second, for Eq. (2.12), note that in the one-dimensional case, the exponential covariance is equivalent with the Matérn covariance with $\nu = 0.5$. Hence, we can apply Eq. (2.9) with no periodisation to infer that the corresponding eigenvalues satisfy:

$$\lambda_j \lesssim j^{-2}.$$

To extend this to higher dimensions, we exploit the fact that the exponential covariance with norm $p = 1$ is separable, so that:

$$\mathbf{s}^{nD} = \frac{1}{\sigma^n} \mathbf{s}^{1D} \otimes \dots \otimes \mathbf{s}^{1D},$$

where \mathbf{s}^{nD} is the first row of the n -dimensional embedding matrix S , and \otimes denotes the Kronecker product. Further, the discrete Fourier matrix satisfies [Cooley et al., 1969]:

$$F^{nD} = F^{1D} \otimes \dots \otimes F^{1D},$$

so that:

$$\boldsymbol{\lambda}^{nD} = \boldsymbol{\lambda}^{1D} \otimes \dots \otimes \boldsymbol{\lambda}^{1D}.$$

Hence, the eigenvalues corresponding to the n -dimensional case decay at the same rate:

$$\lambda_j^{nD} \lesssim j^{-2}$$

as the ones in the one-dimensional case (see e.g. [Charrier, 2012]). \square

2.4 Smoothing

Here, we focus on the case when the random field Z we aim to sample from is extremely oscillatory, so that its fluctuations can only be resolved on a fine mesh, which can render associated simulations computationally expensive. To this end, we smooth the generated realisation by setting the smallest eigenvalues in the factorisation of the circulant embedding matrix, which correspond to highly oscillatory terms, to zero. This approach is similar in nature to the level-dependent truncation of the KL-expansion proposed in [Teckentrup et al., 2013].

With this in mind, recall that given a zero-mean stationary Gaussian random field Z , we now have $\mathbf{Z} = (G\Lambda^{\frac{1}{2}}\boldsymbol{\xi})_R$, for $\boldsymbol{\xi} \sim \mathcal{N}(0, I)$, a sample from the discrete representation of Z . Here, Λ is a diagonal matrix whose entries are the eigenvalues $\{\lambda_j\}_{j=1}^s$ corresponding to the embedding matrix $S \in \mathbb{R}^{s \times s}$. Let k denote the number of eigenvalues we discard, so that $\tilde{\Lambda}$ is a diagonal matrix whose entries are the truncated eigenvalues $\{\lambda_j\}_{j=1}^{s-k}$ and $\mathbf{0} \in \mathbb{R}^k$. Further, let $\tilde{\mathbf{Z}} = (G\tilde{\Lambda}^{\frac{1}{2}}\boldsymbol{\xi})_R$ represent the corresponding smoothed realisation.

Then, the underlying smoothing procedure is as follows:

1. sort eigenvalues $\{\lambda_j\}_{j=1}^s$ in non-increasing order;
2. choose truncation index k such that $\mathbb{E} \left[\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_{\infty} \right]$ is less than a given threshold ε ;
3. compute smoothed realisation $\tilde{\mathbf{Z}}$ using the same $\boldsymbol{\xi}$ as in the definition of \mathbf{Z} .

For example, Figure 1 depicts a Gaussian field sample with covariance function (2.10), where we take $D = [0, 1]^2$, norm $p = 1$ and parameters $\sigma^2 = 1$ and $\lambda = 0.1$. In addition, in Figures 1b and 1d we drop $k = \frac{7s}{8}$ eigenvalues, yielding a noticeably smoother approximation of the sample in Figures 1a and 1c.

To be able to implement this algorithm in practice, we require a theoretical bound on the error between the original and the smoothed samples \mathbf{Z} and $\tilde{\mathbf{Z}}$, respectively. This is given in the following theorem:

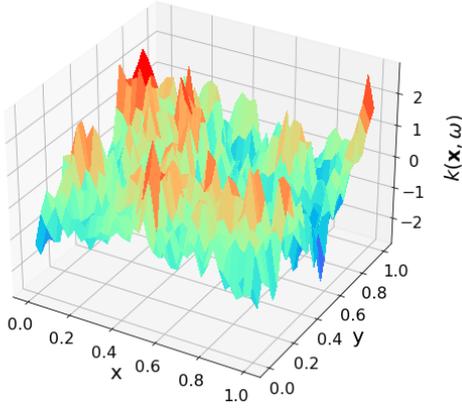
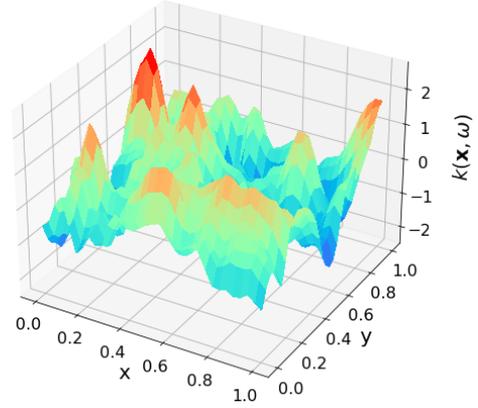
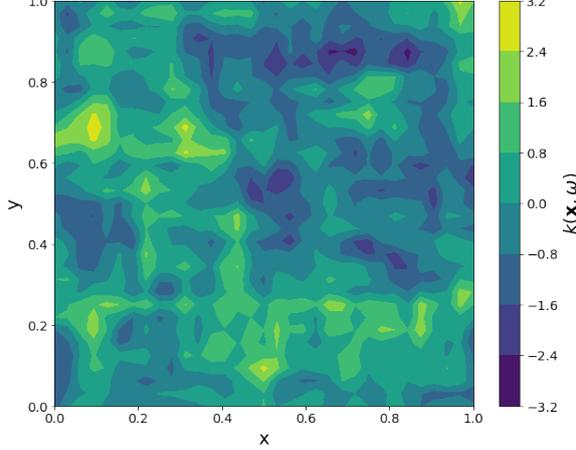
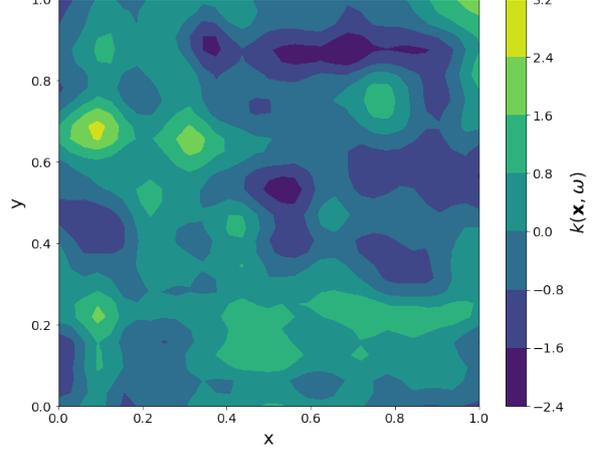
Gaussian Field sample for $\rho = 0.1$ and $\sigma = 1$, no smoothing(a) Surface plot of $\mathbf{Z}(\mathbf{x}, \omega)$.Gaussian Field sample for $\rho = 0.1$ and $\sigma = 1$, smoothing(b) Surface plot of $\tilde{\mathbf{Z}}(\mathbf{x}, \omega)$.Gaussian Field sample for $\rho = 0.1$ and $\sigma = 1$, no smoothing(c) Contour plot of $\mathbf{Z}(\mathbf{x}, \omega)$.Gaussian Field sample for $\rho = 0.1$ and $\sigma = 1$, smoothing(d) Contour plot of $\tilde{\mathbf{Z}}(\mathbf{x}, \omega)$.

Figure 1: Plots of Gaussian field samples $\mathbf{Z}(\mathbf{x}, \omega)$ and $\tilde{\mathbf{Z}}(\mathbf{x}, \omega)$ obtained using circulant embedding for the exponential covariance (2.10) with $p = 1$, $\sigma^2 = 1$ and $\lambda = 0.1$.

Theorem 1. Let Z be a d -dimensional zero-mean stationary Gaussian random field, and let \mathbf{Z} be the discrete representation of Z obtained using the circulant embedding method on a uniform grid \mathcal{T} with mesh sizes m_i , $i = 1, \dots, d$, of the domain $D \subset \mathbb{R}^d$. Suppose the embedding matrix $S \in \mathbb{R}^{s \times s}$ is symmetric positive definite upon smooth periodisation with padding parameters $\{J_i\}_{i=1}^d \in \mathbb{N}_0^d$. Let k be the truncation index and $\tilde{\mathbf{Z}}$ be the resulting smoothed sample. Then, for any $p \in [1, \infty)$:

$$\mathbb{E} \left[\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_\infty^p \right] \lesssim s^{-\frac{p}{2}} \left(\max_{j=s-k+1, \dots, s} \sqrt{\lambda_j} \right)^p k^p,$$

where $s = \prod_{i=1}^d 2(m_i + J_i)$.

Proof. Using the definition of \mathbf{Z} and $\tilde{\mathbf{Z}}$, note that:

$$\begin{aligned}
\mathbb{E} \left[\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_\infty^p \right] &= \mathbb{E} \left[\left\| \left(G\Lambda^{\frac{1}{2}} \boldsymbol{\xi} \right)_R - \left(G\tilde{\Lambda}^{\frac{1}{2}} \boldsymbol{\xi} \right)_R \right\|_\infty^p \right] \\
&\lesssim \mathbb{E} \left[\left\| G \left(\Lambda^{\frac{1}{2}} - \tilde{\Lambda}^{\frac{1}{2}} \right) \boldsymbol{\xi} \right\|_\infty^p \right] \\
&= \mathbb{E} \left[\max_{i=1, \dots, s} \left| \sum_{j=s-k+1}^s g_{ij} \sqrt{\lambda_j} \xi_j \right|^p \right] \\
&\lesssim \mathbb{E} \left[\max_{i=1, \dots, s} \left(\max_{j=s-k+1, \dots, s} |g_{ij}| \max_{j=s-k+1, \dots, s} \sqrt{\lambda_j} \sum_{j=s-k+1}^s |\xi_j| \right)^p \right] \\
&= \max_{i=1, \dots, s} \max_{j=s-k+1, \dots, s} |g_{ij}|^p \left(\max_{j=s-k+1, \dots, s} \sqrt{\lambda_j} \right)^p \mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right].
\end{aligned}$$

Now, $g_{ij} = \Re(f_{ij}) + \Im(f_{ij})$, where f_{ij} is an entry in the d -dimensional unitary Discrete Fourier Transform (DFT). This is a tensor product of d one-dimensional unitary DFT's [Cooley et al., 1969], whose entries are given in Eq. (2.4), so that the following holds:

$$|f_{ij}| \leq \frac{1}{\sqrt{s}}. \quad (2.13)$$

Further, applying the identity $|\Re(z) + \Im(z)| \leq \sqrt{2}|z|$, for any $z \in \mathbb{C}$, to $g_{ij} = \Re(f_{ij}) + \Im(f_{ij})$, together with Eq. (2.13) we derive the following:

$$\max_{i=1, \dots, s} \max_{j=s-k+1, \dots, s} |g_{ij}| \leq \sqrt{\frac{2}{s}},$$

which yields:

$$\mathbb{E} \left[\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_\infty^p \right] \lesssim s^{-\frac{p}{2}} \left(\max_{j=s-k+1, \dots, s} \sqrt{\lambda_j} \right)^p \mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right].$$

Next note that, using the linearity of expectation, we obtain the following:

$$\mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right] = \sum_{i_1=s-k+1}^s \dots \sum_{i_p=s-k+1}^s \mathbb{E} \left[\prod_{m=1}^p |\xi_{i_m}| \right].$$

Since some elements in the sum above are correlated, we observe that every term is, in fact, an element of the following set:

$$\Xi^p = \left\{ \prod_{i=1}^m \mathbb{E}[|\xi^{p_i}|] : p_i \in \{1, 2, \dots, p\}, i = 1, \dots, m, \sum_{i=1}^m p_i = p \right\},$$

where $\xi \sim \mathcal{N}(0, 1)$. Hence, we can write:

$$\begin{aligned}
\mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right] &\leq k^p \max_{m=1, \dots, p} \max_{i_m=s-k+1, \dots, s} \mathbb{E} \left[\prod_{m=1}^p |\xi_{i_m}| \right] \\
&= k^p \max_{q_i \in \Xi^p} q_i \\
&= k^p q^* \\
&= k^p \prod_{i=1}^{m^*} \mathbb{E}[|\xi|^{p_i^*}],
\end{aligned}$$

with $\sum_{i=1}^{m^*} p_i^* = p$ and $p_i^* \in \{1, 2, \dots, p\}, i = 1, \dots, m^*$. In the second step, we used that $\mathbb{E}[\prod_{m=1}^p |\xi_{i_m}|] \in \Xi^p$ by definition of the set Ξ^p . Since Ξ^p is finite, its maximum element, which we denote by q^* , will also be an element of Ξ^p . In the final step, we simply replace q^* with its form as given in Ξ^p . Further, using Jensen's inequality for $\varphi(x) = |x|^{\frac{p}{p_i}}$, we obtain:

$$\mathbb{E}[|\xi|^{p_i}] \leq \mathbb{E}[|\xi|^p]^{\frac{p_i}{p}},$$

for $i = 1, \dots, m$, so that:

$$\prod_{i=1}^m \mathbb{E}[|\xi|^{p_i}] \leq \mathbb{E}[|\xi|^p]^{\frac{\sum_{i=1}^m p_i}{p}} = \mathbb{E}[|\xi|^p].$$

Putting everything together, we determine:

$$\mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right] \leq k^p \mathbb{E}[|\xi|^p], \quad p \in \mathbb{N}.$$

From [Elandt, 1961], we know that if $\xi \sim \mathcal{N}(\mu, \sigma^2)$, then $|\xi|$ is a folded normal random variable. For the special case when $\mu = 0$, $|\xi|$ is a half-normal random variable, whose higher moments are:

$$\begin{aligned} \mathbb{E}[|\xi|^{2n}] &= \frac{(2n)!}{n! 2^n}, \\ \mathbb{E}[|\xi|^{2n+1}] &= \sqrt{\frac{2}{\pi}} 2^n n!, \end{aligned}$$

for $n \in \mathbb{N} \cup \{0\}$ and $\sigma = 1$. Using these formulae, we can infer the following:

$$\mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right] \leq \begin{cases} \frac{(2n)!}{n! 2^n} k^{2n}, & \text{if } p = 2n, \\ \sqrt{\frac{2}{\pi}} 2^n n! k^{2n+1}, & \text{if } p = 2n + 1. \end{cases}$$

Re-arranging, this reduces to:

$$\mathbb{E} \left[\left(\sum_{j=s-k+1}^s |\xi_j| \right)^p \right] \lesssim k^p,$$

with constant dependent on p , but not on k .

Combining everything gives us the claimed result for $p \in \mathbb{N}$, but this can be extended to $p \in [1, \infty)$ using Jensen's inequality. The claim of the Theorem then follows. \square

In practice, we observe that the error $\mathbb{E}[\|\mathbf{Z} - \tilde{\mathbf{Z}}\|_\infty]$ decays much faster as $k \rightarrow 0$, that is, for small values of k . This is to be expected, as the case $k = 0$ yields $\mathbf{Z} - \tilde{\mathbf{Z}} = \mathbf{0}$. As the number of eigenvalues we discard increases, we recover the decay rate that we theoretically derive. What is more, this bound is directly linked with the largest eigenvalue we keep, so that the faster the eigenvalues of the embedding matrix decline, the faster the smoothing error decreases.

3 Multilevel Monte Carlo Methods

In this section, we focus on an individual application of the smoothing technique, namely Multilevel Monte Carlo methods for PDEs with random coefficients. We first introduce the PDE model of interest and outline key features of the MC and MLMC estimators, where we make use of circulant embedding methods for sampling from the aforementioned random parameter. Finally, we describe how to integrate the smoothing approach in these methods, and give estimates on the error introduced by smoothing samples from the random field and the complexity of the resulting estimators.

3.1 PDE with random coefficient model

A mathematical model for simulating single-phase, stationary groundwater flow rests on Darcy's Law combined with the *Law of conservation of mass* [De Marsily, 1986], which leads to a linear relationship between the pressure difference of the groundwater and its flow rate, as follows:

$$-\nabla \cdot (k(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}), \quad \mathbf{x} \in D. \quad (3.1)$$

For simplicity, we let $D = (0, 1)^2$, and take the boundary conditions to be:

$$\begin{aligned} u|_{x_1=0} &= 1, & u|_{x_1=1} &= 0, \\ \frac{\partial u}{\partial \mathbf{n}} \Big|_{x_2=0} &= 0, & \frac{\partial u}{\partial \mathbf{n}} \Big|_{x_2=1} &= 0. \end{aligned}$$

Here, u denotes the pressure of the fluid, and $f := -\nabla \cdot \mathbf{g}$, where \mathbf{g} represents the source terms. Further, k stands for the hydraulic conductivity, that is, the ease with which a fluid can move through porous media or fractures under a given pressure gradient. Moreover, we apply the flow cell boundary conditions, which comprise of Dirichlet boundary conditions, ensuring that the groundwater stream is from $x_1 = 0$ to $x_1 = 1$, and Neumann boundary conditions, forcing the flow to remain vertically inside the domain.

To account for uncertainties in the coefficient k , we assume that it is a log-normal random field, i.e. $k(\mathbf{x}, \omega) = \exp(Z(\mathbf{x}, \omega))$. Here, Z is a zero-mean stationary Gaussian field, so that it has the structure outlined in Eq. (2.1). One example of covariance function $C : [-1, 1]^2 \rightarrow \mathbb{R}$ associated with the coefficient k was suggested by Hoeksema and Kitanidis in [Hoeksema and Kitanidis, 1985], and it has the form given in Eq. (2.10). In this instance, samples $Z(\cdot, \omega)$ are Hölder continuous with respect to \mathbf{x} , with Hölder exponent $\alpha < 1/2$, so that:

$$|Z(\mathbf{x}, \cdot) - Z(\mathbf{y}, \cdot)| \lesssim \|\mathbf{x} - \mathbf{y}\|^\alpha, \quad \forall \mathbf{x}, \mathbf{y} \in D.$$

Another example of C commonly used in this application is the Matérn covariance function given in Eq. (2.5).

Generally, we are interested in computing quantities of interest $\mathbb{E}[Q]$ in the form of expected values of functionals Q of the PDE solution u . An easily computed example of physical interest is the pressure of the water at a given point $\mathbf{x}^* \in D$ or the L_2 norm of the solution $\|u\|_{L_2(D)}$.

3.2 Finite Elements discretisation

To compute quantities of interest Q of the PDE solution u , we adopt the finite element method to obtain numerical approximations u_h of u and, subsequently Q_h of Q . We do not dwell on the practical implementation, as we exploit the Python FEniCS software [Alnaes et al., 2015; Logg et al., 2012] for this. Rather, we focus on theoretical bounds on the error $\|u - u_h\|_{H^1(D)}$, as this is crucial in our simulations. For simplicity, we consider zero Dirichlet boundary conditions $u = 0$ on ∂D for the analysis, but we expect similar results to hold with the mixed boundary conditions presented above.

To this end, consider the variational formulation of the PDE (3.1) for a given sample $k(\mathbf{x}, \omega)$:

$$\int_D k(\cdot, \omega) \nabla u \cdot \nabla v \, d\mathbf{x} = \int_D f v \, d\mathbf{x}, \quad \forall v \in \mathcal{V},$$

where $\mathcal{V} = H_0^1(D)$ is the space of test functions. Upon discretisation, we seek solutions to this equation in a subspace $\mathcal{V}_h \subset \mathcal{V}$ spanned by piece-wise linear polynomials. This is because the low regularity of the coefficient k limits the accuracy we can achieve with a higher order polynomial basis. For a more detailed implementation of finite elements for this problem see e.g. [Cliffe et al., 2000] or [Graham et al., 2011], and for a standard cell-centred finite volume approach see [Cliffe et al., 2011].

The error $\|u - u_h\|_{H^1(D)}$ is quantified in [Charrier et al., 2013, Proposition 3.13], where u_h is the finite element approximation obtained by applying either the midpoint rule or the trapezoidal rule when assembling the resulting linear system. In particular, this takes into account both the finite element error and the quadrature error, and u_h depends on k only through its values at the grid points. The result is summarised below.

Proposition 1. *Let $u(\cdot, \omega), u_h(\cdot, \omega) \in H^1(D)$ be as above. Then, for $0 < s < t \leq 1$ and $0 < h < 1$:*

$$\|u - u_h\|_{L^p(\Omega, H^1(D))} \lesssim h^s,$$

for any $p \in [1, \infty)$ and $s \neq \frac{1}{2}$, under the assumptions that:

- (A) the domain $D \subset \mathbb{R}^d$, $d = 1, 2$, is a bounded, convex, Lipschitz polygonal domain.
- (B) $k_{\min}(\omega) := \min_{\mathbf{x} \in \bar{D}} k(\mathbf{x}, \omega) > 0$ almost surely and $1/k_{\min} \in L^p(\Omega)$, for all $p \in [1, \infty)$.
- (C) $k \in L^p(\Omega, C^t(\bar{D}))$, for some $0 < t \leq 1$ and for all $p \in [1, \infty)$.
- (D) $f \in H^{t-1}(D)$.

Note that Assumption A was replaced by \mathcal{C}^2 bounded domains $D \subset \mathbb{R}^d$ in [Charrier et al., 2013], but the result extends to polygonal domains using the results on the regularity of u in [Teckentrup et al., 2013]. Assumptions B and C are satisfied for log-normal random fields as described in section 3.1. We can extend Proposition 1 to the error $\|Q - Q_h\|_{L^p(\Omega)}$, where $Q = \mathcal{G}(u(\cdot, \omega))$ and $Q_h = \mathcal{G}(u_h(\cdot, \omega))$, for u and u_h as above, and for some bounded functional $\mathcal{G} : H^1(D) \rightarrow \mathbb{R}$. This is stated in [Teckentrup et al., 2013, Lemma 3.2]:

Lemma 3. *Let $u(\cdot, \omega), u_h(\cdot, \omega) \in H^1(D)$ be as above. Suppose $\mathcal{G} : H^1(D) \rightarrow \mathbb{R}$ is a bounded functional satisfying the following assumption:*

- (E) \mathcal{G} is continuously Fréchet differentiable, and there exists $C_F \in L^q(\Omega)$, for all $q \in [1, \infty)$, such that:

$$|\overline{D_v \mathcal{G}}(u, u_h)| \lesssim C_F(\omega) \|v\|_{H^1(D)}, \quad \text{for all } v \in H_0^1(D) \quad \text{and for almost all } \omega \in \Omega,$$

where:

$$\overline{D_v \mathcal{G}}(u, u_h) := \int_0^1 D_v \mathcal{G}(u + \theta(u_h - u)) d\theta.$$

Here, the Gateaux derivative of \mathcal{G} at \tilde{v} and in the direction v is defined as:

$$D_v \mathcal{G}(\tilde{v}) := \lim_{\varepsilon \rightarrow 0} \frac{\mathcal{G}(\tilde{v} + \varepsilon v) - \mathcal{G}(\tilde{v})}{\varepsilon}, \quad \forall v, \tilde{v} \in H^1(D).$$

Then, for all $p \in [1, \infty)$:

$$\|Q - Q_h\|_{L^p(\Omega)} \lesssim \|u - u_h\|_{L^p(\Omega, H^1(D))}.$$

Throughout the remainder of this paper, we will use the notation u_h to denote the FE solution which uses any quadrature rule based on values of the coefficient k at the grid points, such as the trapezoidal rule, for the linear system assembly, and Q_h the resulting functional.

3.3 Monte Carlo estimator

To be able to define an estimator of the quantity of interest $\mathbb{E}[Q]$, let us assume that, for a finite element approximation Q_h of Q with mesh size h , the following holds:

$$\mathbb{E}[Q_h] \rightarrow \mathbb{E}[Q] \text{ as } h \rightarrow 0,$$

with mean order of convergence α , so that:

$$\mathbb{E}[Q_h - Q] \leq C_\alpha h^\alpha, \quad \alpha > 0.$$

A non-intrusive method for estimating $\mathbb{E}[Q]$ is Monte Carlo, which, for the application at hand, is given by the following equation:

$$\widehat{Q}_{h,N}^{\text{MC}} := \frac{1}{N} \sum_{i=1}^N Q_h^{(i)}, \quad (3.2)$$

where $\{Q_h^{(i)}\}_{i=1}^N$ are independently sampled from the distribution of Q_h . To achieve this, we first require a sample from the Gaussian field Z , which then yields a realisation from the coefficient k . For each such sample, we compute a finite element approximation u_h of the solution u , which then allows us to compute an estimate Q_h of the quantity of interest Q .

Further, the mean squared error (MSE) between the MC estimator in Eq. (3.2) and the quantity of interest $\mathbb{E}[Q]$ can be expanded as follows:

$$e\left(\widehat{Q}_{h,N}^{\text{MC}}\right)^2 := \mathbb{E}\left[\left(\widehat{Q}_{h,N}^{\text{MC}} - \mathbb{E}[Q]\right)^2\right] = \frac{1}{N}\mathbb{V}[Q_h] + (\mathbb{E}[Q_h - Q])^2. \quad (3.3)$$

This highlights the two contributions to the MSE: the variance of the MC estimator, and the bias introduced by replacing the functional Q with its numerical approximation Q_h .

Consequently, Eq. (3.3) suggests a strategy for identifying the necessary number of MC samples N and the grid size h such that $e\left(\widehat{Q}_{h,N}^{\text{MC}}\right) < \varepsilon$, for a given accuracy ε . Specifically, if we allow equal contributions of $\varepsilon^2/2$ to the two terms in Eq. (3.3), we can choose:

$$N > 2\mathbb{V}[Q_h]\varepsilon^{-2} \quad \text{and} \quad h < \left(\frac{\varepsilon}{\sqrt{2}C_\alpha}\right)^{1/\alpha}. \quad (3.4)$$

Further, assume that the cost of computing one sample of Q_h satisfies:

$$\mathcal{C}\left(Q_h^{(i)}\right) \leq C_\gamma h^{-\gamma}, \quad \gamma > 0. \quad (3.5)$$

Typically, the two main contributions to the cost $\mathcal{C}\left(Q_h^{(i)}\right)$ are sampling from the parameter $k(\cdot, \cdot)$, and numerically solving the PDE (3.1). For an appropriate FFT implementation, such as FFTW [Frigo and Johnson, 2005], and an optimal iterative linear solver, such as Generalized minimal residual method with incomplete LU factorisation as a preconditioner in FEniCS, a suitable choice is $\gamma \approx d$. In addition, C_γ might depend on the parameters of the covariance function of the random field, such as σ^2 , λ , and ν , but it is independent of h .

Let $\mathcal{C}_\varepsilon(\widehat{Q}_{h,N}^{\text{MC}})$ be the computational ε -cost, quantified by the number of floating point operations (FLOPS) required in order to satisfy $e(\widehat{Q}_{h,N}^{\text{MC}}) < \varepsilon$. Using Eq. (3.5) and under the assumption that $\mathbb{V}[Q_h]$ is approximately constant independent of h , the ε -cost $\mathcal{C}_\varepsilon\left(\widehat{Q}_{h,N}^{\text{MC}}\right)$ associated with the estimator (3.2) is given by:

$$\mathcal{C}_\varepsilon\left(\widehat{Q}_{h,N}^{\text{MC}}\right) \lesssim \varepsilon^{-2-\gamma/\alpha}. \quad (3.6)$$

Therefore, Eq. (3.6) emphasises that the ε -cost associated with the MC estimator is directly affected by the parameters γ and α . In particular, the latter is related to the regularity of the functional under consideration, so that the more irregular the functional is, the smaller α is, and, consequently, the larger the cost, rendering this approach computationally expensive.

3.4 Multilevel Monte Carlo simulations

The fundamental principle at the basis of Multilevel Monte Carlo simulations rests on a multilevel variance reduction technique for the standard MC method [Giles, 2008]. This is achieved by sampling from several approximations Q_h of Q , rather than just one. To this end, let $\{h_\ell : \ell = 0, \dots, L\}$ be a decreasing sequence in \mathbb{Q} defining the MLMC levels, so that $h_0 > \dots > h_L := h$. Further, suppose for simplicity that there exists an integer $r \in \mathbb{N} \setminus \{1\}$ satisfying:

$$h_{\ell-1} = rh_\ell, \quad \forall \ell = 1, \dots, L.$$

In other words, the MLMC levels are given, in this case, by different grid resolutions used in approximating the solution to the PDE (3.1). For example, we can take $h_\ell = 2^{-\ell}$, so that $r = 2$.

Next, letting $Y_\ell := Q_{h_\ell} - Q_{h_{\ell-1}}$, $\ell = 1, \dots, L$, with $Y_0 := Q_{h_0}$, the idea is to express the approximation $\mathbb{E}[Q_h]$ of the quantity of interest $\mathbb{E}[Q]$ as:

$$\mathbb{E}[Q_h] = \mathbb{E}[Q_{h_0}] + \sum_{\ell=1}^L \mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}] = \sum_{\ell=0}^L \mathbb{E}[Y_\ell],$$

via the linearity of expectation. Thus, we avoid estimating $\mathbb{E}[Q_{h_\ell}]$ on level ℓ directly. Rather, we approximate the correction term $\mathbb{E}[Y_\ell]$ with respect to the subsequent lower level. Then, the MLMC estimator is defined as [Giles, 2008]:

$$\widehat{Q}_L^{\text{MLMC}} = \widehat{Q}_{h_0, N_0}^{\text{MC}} + \sum_{\ell=1}^L \widehat{Y}_{\ell, N_\ell}^{\text{MC}}, \quad (3.7)$$

where $\widehat{Y}_{\ell, N_\ell}^{\text{MC}}$ is the standard MC approximation of $\mathbb{E}[Y_\ell]$, $\ell = 0, \dots, L$ with N_ℓ samples, namely:

$$\widehat{Y}_{\ell, N_\ell}^{\text{MC}} := \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \left(Q_{h_\ell}^{(i)} - Q_{h_{\ell-1}}^{(i)} \right),$$

the number of samples per level satisfying $N_0 > N_1 > \dots > N_L$. Note that it is essential to use the same underlying sample $k(\mathbf{x}, \omega^{(i)})$ when computing the difference $Q_{h_\ell}^{(i)} - Q_{h_{\ell-1}}^{(i)}$.

In order to determine the number of levels L and of samples N_ℓ , $\ell = 0, \dots, L$ required to achieve a given accuracy ε , similar to the derivation in (3.3), we obtain:

$$e \left(\widehat{Q}_L^{\text{MLMC}} \right)^2 := \mathbb{E} \left[\left(\widehat{Q}_L^{\text{MLMC}} - \mathbb{E}[Q] \right)^2 \right] = \sum_{\ell=0}^L \frac{1}{N_\ell} \mathbb{V}[Y_\ell] + (\mathbb{E}[Q_h - Q])^2. \quad (3.8)$$

Hence, we find the same two terms contributing to the MSE between the quantity of interest $\mathbb{E}[Q]$ and the MLMC estimator (3.7): its sample variance and bias. Consequently, we can apply the same strategy to identifying h and N_ℓ , $\ell = 0, \dots, L$, namely that each term in Eq. (3.8) is less than $\varepsilon^2/2$. In particular, since the bias is identical to the approximation error in Eq. (3.3), we can simply take $h = h_L \lesssim \varepsilon^{1/\alpha}$, as in Eq. (3.4).

Moreover, if we consider each N_ℓ to be a continuous variable for $\ell = 0, \dots, L$, then the optimal number of samples N_ℓ for each MLMC level can be established by solving the following optimisation problem:

$$\begin{aligned} \min_{N_\ell} \quad & \sum_{\ell=0}^L N_\ell \mathcal{C}_\ell \\ \text{s.t.} \quad & \sum_{\ell=0}^L \frac{1}{N_\ell} V_\ell = \frac{\varepsilon^2}{2}. \end{aligned}$$

Here, $\mathcal{C}_\ell := \mathcal{C}(Y_\ell^{(i)})$ represents the cost of computing one sample $Y_\ell^{(i)}$, while $V_\ell := \mathbb{V}[Y_\ell]$ denotes the variance of one such sample. In other words, we minimise the cost of the MLMC estimator required to achieve a corresponding variance of $\varepsilon^2/2$. This gives:

$$N_\ell = 2\varepsilon^{-2} \left(\sum_{\ell=0}^L \sqrt{\mathcal{C}_\ell V_\ell} \right) \sqrt{\frac{V_\ell}{\mathcal{C}_\ell}}, \quad \ell = 0, \dots, L, \quad (3.9)$$

which suggests that the number of samples N_ℓ satisfies:

$$N_\ell \rightarrow 1 \quad \text{as} \quad \ell \rightarrow \infty,$$

under the assumption that Q_{h_ℓ} converges to Q in mean square, that is:

$$\mathbb{E}[|Q_{h_\ell} - Q|^2] \rightarrow 0 \quad \text{as} \quad \ell \rightarrow \infty.$$

In addition, Eq. (3.9) implies that the ε -cost of the MLMC estimator in (3.7) is given by:

$$\mathcal{C}_\varepsilon(\widehat{Q}_L^{\text{MLMC}}) = 2\varepsilon^{-2} \left(\sum_{\ell=0}^L \sqrt{\mathcal{C}_\ell V_\ell} \right)^2. \quad (3.10)$$

In other words, the ε -cost of MLMC is intrinsically conditional on whether the sum in Eq. (3.10) is uniformly bounded, which, in turn, hinges on whether the product $\mathcal{C}_\ell V_\ell$ increases or decreases with level ℓ , or, equivalently, on whether the cost \mathcal{C}_ℓ grows more quickly than the variance V_ℓ decreases. Hence, we distinguish the following scenarios:

- If the cost \mathcal{C}_ℓ increases more slowly than the variance V_ℓ declines, the term which prompts the largest contribution to the cost $\mathcal{C}_\varepsilon(\widehat{Q}_L^{\text{MLMC}})$ is $\mathcal{C}_0 V_0$, so that $\mathcal{C}_\varepsilon(\widehat{Q}_L^{\text{MLMC}}) \approx \varepsilon^{-2} \mathcal{C}_0 V_0$. On the other hand, using Eq. (3.3), the ε -cost of the MC estimator satisfies $\mathcal{C}_\varepsilon(\widehat{Q}_{h,N}^{\text{MC}}) \approx \varepsilon^{-2} \mathcal{C}_L V_0$. Thus, the cost savings are approximately $\mathcal{C}_0/\mathcal{C}_L \simeq \varepsilon^{\gamma/\alpha}$.
- If the cost \mathcal{C}_ℓ grows faster than the variance V_ℓ decays, the main contribution to the cost is given by the approximation on the finest level L , so that $\mathcal{C}_\varepsilon(\widehat{Q}_L^{\text{MLMC}}) \approx \varepsilon^{-2} \mathcal{C}_L V_L$. Similarly as above, the computational savings are roughly V_L/V_0 .
- If the product $\mathcal{C}_\ell V_\ell$ does not vary with level, the cost is $\varepsilon^{-2} L^2 V_0 \mathcal{C}_0 = \varepsilon^{-2} L^2 V_L \mathcal{C}_L$ [Giles, 2013].

The above comments are formally developed in [Cliffe et al., 2011, Theorem 1], which states the following:

Theorem 2. *Suppose that there exist positive constants $\alpha, \beta, \gamma > 0$ such that $\alpha \geq \frac{1}{2} \min(\beta, \gamma)$, and the following hold:*

- (i) $|\mathbb{E}[Q_{h_\ell} - Q]| \lesssim h_\ell^\alpha$,
- (ii) $\mathbb{V}[Y_\ell] \lesssim h_\ell^\beta$,
- (iii) $\mathcal{C}_\ell \lesssim h_\ell^{-\gamma}$.

Then, for any $\varepsilon < e^{-1}$, there exists a value L (and correspondingly $h \equiv h_L$) and a sequence $\{N_\ell\}_{\ell=0}^L$ such that:

$$e(\widehat{Q}_L^{\text{MLMC}})^2 = \mathbb{E} \left[\left(\widehat{Q}_L^{\text{MLMC}} - \mathbb{E}[Q] \right)^2 \right] < \varepsilon^2,$$

and

$$\mathcal{C}_\varepsilon(\widehat{Q}_L^{\text{MLMC}}) \lesssim \begin{cases} \varepsilon^{-2}, & \text{if } \beta > \gamma, \\ \varepsilon^{-2} (\log \varepsilon)^2, & \text{if } \beta = \gamma, \\ \varepsilon^{-2 - (\gamma - \beta)/\alpha}, & \text{if } \beta < \gamma. \end{cases}$$

A possible implementation of the MLMC estimator is proposed by Giles et al. in [Giles, 2008] and Cliffe et al. in [Cliffe et al., 2011], where they adopt an adaptive approach to computing the number of levels.

3.5 MC and MLMC with Circulant Embedding

Using the observations in Section 2, we can write the MC and MLMC estimators $\widehat{Q}_h(\mathbf{Z})$, where we use the $\widehat{Q}_h(\mathbf{Z})$ notation to highlight the dependency on the Gaussian field \mathbf{Z} , as follows:

$$\widehat{Q}_{h,N}^{\text{MC}}(\mathbf{Z}) = \frac{1}{N} \sum_{i=1}^N Q_h((G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}^{(i)})_R),$$

$$\widehat{Q}_L^{\text{MLMC}} = \frac{1}{N_0} \sum_{i=1}^{N_0} Q_{h_0}((G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}^{(i,0)})_R) + \sum_{\ell=1}^L \left(\frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \left(Q_{h_\ell}((G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}^{(i,\ell)})_R) - Q_{h_{\ell-1}}((G\Lambda^{\frac{1}{2}}\boldsymbol{\xi}^{(i,\ell)})_R) \right) \right),$$

where $\{\boldsymbol{\xi}^{(i,\ell)}\}$ is a collection of IID standard Gaussian random vectors. Here, we use the same discretisation mesh \mathcal{T} of the domain D with mesh size $h \equiv \frac{1}{m_i}$, $i = 1, \dots, d$, for sampling using the circulant embedding scheme and for computing the finite element approximation Q_h .

Thus, the following algorithm is adapted from [Graham et al., 2011, Section 5.3] and it outlines the steps required for estimating $\widehat{Y}_{\ell, N_\ell}^{\text{MC}}$ in the MLMC method on level ℓ , where samples of $k(\mathbf{x}, \cdot)$ are obtained using circulant embedding:

Algorithm 1. *Computes an MC estimate of $\mathbb{E}[Y_\ell]$ on a level ℓ using the circulant embedding method to sample from $k(\cdot, \cdot)$ given the number of samples N_ℓ .*

1. Construct covariance matrix R from discretisation mesh \mathcal{T} and covariance function C .
2. Embed R into a $4mn \times 4mn$ block circulant with circulant blocks matrix S .
3. Compute $\boldsymbol{\lambda} = \sqrt{4mn}F\mathbf{s}$, where \mathbf{s} is the first row of S , and F is the 2D Fourier matrix.
 - (a) Check whether the eigenvalues real and positive. If yes, go to (4). If not, go to (3b).
 - (b) Pad the embedding matrix S to obtain a symmetric positive definite matrix using the smooth periodisation technique and the bound (2.7), and compute updated $\boldsymbol{\lambda}$.
4. For each MC iteration on level ℓ :
 - (a) Generate $\boldsymbol{\xi} \sim \mathcal{N}(0, I)$.
 - (b) Evaluate $\mathbf{w} := \boldsymbol{\lambda} \odot \boldsymbol{\xi}$.
 - (c) Compute $\mathbf{v} := F\mathbf{w}$.
 - (d) Take $\mathbf{u} := \Re(\mathbf{v}) + \Im(\mathbf{v})$.
 - (e) Set $\mathbf{Z}_\ell := (\mathbf{u})_{R_\ell}$ and $\mathbf{Z}_{\ell-1} := (\mathbf{u})_{R_{\ell-1}}$.
 - (f) Let $\mathbf{k}_\ell := \exp(\mathbf{Z}_\ell)$ and $\mathbf{k}_{\ell-1} := \exp(\mathbf{Z}_{\ell-1})$.
 - (g) Solve Eq. (3.1) in turn for \mathbf{k}_ℓ and $\mathbf{k}_{\ell-1}$.

Here, \odot denotes the Hadamard product, so that, for $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$, $\{\mathbf{a} \odot \mathbf{b}\}_j = \mathbf{a}_j \mathbf{b}_j$, $j = 1, \dots, d$.

It is worth mentioning that the 2D discrete Fourier transform in Step 3 and Step 4(c) can be computed using an FFT implementation, such as FFTW [Frigo and Johnson, 2005], significantly reducing the computational complexity. In addition, S is a circulant matrix and, therefore, uniquely defined by its first row or column. As a consequence, Step 1 and Step 2 can be altered to only construct the first row or column of S , rather than the full matrix, diminishing the memory requirements.

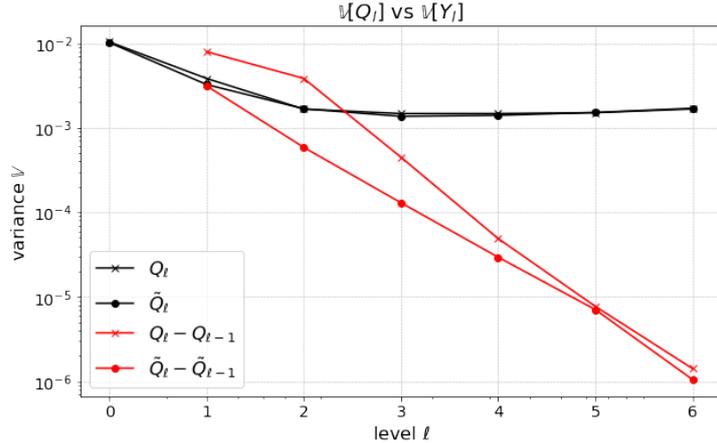


Figure 2: Plot of variance decay for the quantity of interest $Q = p(\mathbf{x}^*)$ using the Matérn covariance function for the random coefficient with $\sigma^2 = 1$, $\nu = 1.5$ and $\lambda = 0.03$.

3.6 MC and MLMC with Smoothed Circulant Embedding

One argument which supports the computational efficiency of the MLMC estimator (3.7) over the standard MC approximation (3.2) is that the coarsest level is invariant for any imposed accuracy ε , and, consequently, the cost \mathcal{C}_0 does not increase as $\varepsilon \rightarrow 0$. Generally, a minimum value h_0 is necessary to provide a basic level of detail on the problem. This is selected depending on the regularity of the solution to the PDE (3.1). While this can render the MLMC estimator more costly for large accuracies ε , it still leads to significant reductions in cost as $\varepsilon \rightarrow 0$.

In some applications, the choice of coarsest mesh is limited by the length scale λ and the smoothness parameter ν of the random field $k(\cdot, \cdot)$. This is illustrated in Figure 2 for the functional $Q = u\left(\frac{7}{15}, \frac{7}{15}\right)$ of the solution to the PDE (3.1), where we use the Matérn covariance for the random parameter $k(\cdot, \cdot)$ with $\sigma^2 = 1$, $\nu = 1.5$ and $\lambda = 0.03$. In particular, the plot highlights that on the first levels, which correspond to coarse meshes, the variance $\mathbb{V}[Y_\ell]$ of the difference $Q_{h_\ell} - Q_{h_{\ell-1}}$ is larger than then the variance $\mathbb{V}[Q_{h_\ell}]$ of the quantity of interest. Hence, for these levels, the contributions to the cost $\mathcal{C}_\varepsilon\left(\hat{Q}_L^{\text{MLMC}}\right)$ will actually be more significant than those using standard MC, rendering the MLMC approach more expensive.

For this reason, Cliffe et al. [Cliffe et al., 2011] make the following observation at the end of Section 4.1:

$$h_0 \leq \lambda, \quad (3.11)$$

for the covariance function in Eq. (2.10). For the Matérn case, we use the heuristically derived relationship between the correlation length ρ , smoothness parameter ν and length scale λ [Lindgren et al., 2011]:

$$\rho = \sqrt{8\nu\lambda},$$

so that a suitable estimate for the coarsest mesh in the MLMC estimator is:

$$h_0 \leq \sqrt{8\nu\lambda}. \quad (3.12)$$

This corresponds to the point where the two lines of $\mathbb{V}[Y_\ell]$ and $\mathbb{V}[Q_{h_\ell}]$ intersect in Figure 2, and limits severely the computational gains prompted by MLMC for small λ and ν , as it restricts the number of (computationally cheap) levels we can include in the MLMC estimator.

One solution to this problem was proposed in [Teckentrup et al., 2013] where the KL-expansion is used to sample from the random coefficient. To mitigate the aforementioned issue, Teckentrup et al. suggest a *level-dependent* truncation of the KL-expansion, where the most oscillatory terms in the sum are neglected to obtain a more accurate approximation of the

random field on coarser levels. This then allows to choose h_0 independent of λ , leading to noteworthy reductions in the cost. Similar ideas on level-dependent truncation have been used in [Gittelsohn et al., 2013; Schwab and Stein, 2023].

The goal is to replicate this idea in the context of circulant embedding methods. It is worth highlighting that the truncated KL-expansion is an approximation of the random field, and so the accuracy of the estimate is directly affected by where the sum is truncated. On the other hand, circulant embedding methods are exact on the discretisation mesh, and so we can no longer distinguish different levels of approximation of the random field. Rather, we limit the extent of fluctuations in $k(\mathbf{x}, \omega)$ on the mesh considered on each level of the MLMC estimator by discarding highly oscillatory terms corresponding to the smallest eigenvalues, as outlined in Section 2.4. In doing so, consecutive differences in the definition of $\widehat{Q}_L^{\text{MLMC}}$ have smaller variance. This leads to a choice of h_0 independent of the parameters λ and ν .

To justify using the CES approach in the MLMC estimator for the quantity of interest $\mathbb{E}[Q]$ from the PDE model (3.1), the following theorem provides an estimate of the error introduced by smoothing the random field samples on a given level ℓ :

Theorem 3. *Let $Z_{\mathcal{T}}$ and $\tilde{Z}_{\mathcal{T}}$ be piece-wise linear interpolants of the non-smoothed and smoothed discrete representations \mathbf{Z} and $\tilde{\mathbf{Z}}$, respectively, on a given grid \mathcal{T} of the domain D . Let $\tilde{k} = \exp(\tilde{Z}_{\mathcal{T}})$ and denote by \tilde{u}_h the corresponding finite element solution. Let $\mathcal{G} : H^1(D) \rightarrow \mathbb{R}$ be a bounded functional which satisfies assumption (E) in Lemma 3 with u_h and \tilde{u}_h instead of u and u_h . Then:*

$$\mathbb{E} \left[|Q_h - \tilde{Q}_h|^p \right] \lesssim \mathbb{E} \left[\|Z_{\mathcal{T}} - \tilde{Z}_{\mathcal{T}}\|_{C^0(\bar{D})}^{pq} \right]^{\frac{1}{q}},$$

for any $p, q \in [1, \infty)$.

Proof. From Lemma 3, we have:

$$\mathbb{E} \left[|Q_h(\cdot, \omega) - \tilde{Q}_h(\cdot, \omega)|^p \right] \lesssim \mathbb{E} \left[\|u_h(\cdot, \omega) - \tilde{u}_h(\cdot, \omega)\|_{H^1(D)}^p \right].$$

Then, from [Dodwell et al., 2015, Lemma 4.8], and since the finite element solution computed with a quadrature rule only depends on the values of the coefficient on the grid \mathcal{T} , we know that:

$$\mathbb{E} \left[\|u_h(\cdot, \omega) - \tilde{u}_h(\cdot, \omega)\|_{H^1(D)}^p \right] \lesssim \mathbb{E} \left[\|k(\cdot, \omega) - \tilde{k}(\cdot, \omega)\|_{C^0(\bar{D})}^p \right],$$

where $k = \exp(Z_{\mathcal{T}})$. Using the inequality $|e^x - e^y| \leq (e^x + e^y)|x - y|, \forall x, y \in \mathbb{R}$ and Hölder's inequality, we can derive:

$$\begin{aligned} \mathbb{E} \left[\|\exp(Z_{\mathcal{T}}(\cdot, \omega)) - \exp(\tilde{Z}_{\mathcal{T}}(\cdot, \omega))\|_{C^0(\bar{D})}^p \right] &\leq \\ \mathbb{E} \left[\|\exp(Z_{\mathcal{T}}(\cdot, \omega)) + \exp(\tilde{Z}_{\mathcal{T}}(\cdot, \omega))\|_{C^0(\bar{D})}^{pr} \right]^{\frac{1}{r}} \mathbb{E} \left[\|Z_{\mathcal{T}}(\cdot, \omega) - \tilde{Z}_{\mathcal{T}}(\cdot, \omega)\|_{C^0(\bar{D})}^{pq} \right]^{\frac{1}{q}}, \end{aligned}$$

where $q, r \in [1, \infty]$ are Hölder exponents. Applying Fernique's Theorem [Fernique et al., 1975], we obtain that the first term on the right-hand side is bounded for $r \in [1, \infty)$, which concludes the proof. \square

To obtain a bound for the error between Q_h and \tilde{Q}_h in terms of the number of dropped eigenvalues k , we combine Theorems 1 and 3 to infer the following Corollary:

Corollary 1. *Let the assumptions of Theorems 1 and 3 hold. Then, for any $p \in [1, \infty)$:*

$$\mathbb{E} \left[|Q_h - \tilde{Q}_h|^p \right] \lesssim s^{-\frac{p}{2}} \left(\max_{j=s-k+1, \dots, s} \sqrt{\lambda_j} \right)^p k^p.$$

Proof. This follows directly from Theorems 1 and 3, together with the observation that since $Z_{\mathcal{T}}$ and $\tilde{Z}_{\mathcal{T}}$ are piece-wise linear interpolants of \mathbf{Z} and $\tilde{\mathbf{Z}}$, respectively, we have that:

$$\|Z_{\mathcal{T}}(\cdot, \omega) - \tilde{Z}_{\mathcal{T}}(\cdot, \omega)\|_{C^0(\bar{D})} = \|\mathbf{Z}(\omega) - \tilde{\mathbf{Z}}(\omega)\|_{\infty}.$$

□

This immediately leads to the following Corollary:

Corollary 2. *Let Q_h and \tilde{Q}_h be as in Corollary 1, and let C be the separable exponential covariance (2.10). Then, for any $p \in [1, \infty)$:*

$$\mathbb{E} \left[|Q_h - \tilde{Q}_h|^p \right] \lesssim s^{-\frac{p}{2}} (s - k)^{-p} k^p.$$

Proof. This follows immediately from Corollary 1 and Lemma 2. □

For the Matérn covariance (2.5) where we use smooth periodisation to ensure the positive-definiteness of the (extended) embedding matrix S , we can directly use Eq. (2.9) to infer:

Corollary 3. *Let Q_h and \tilde{Q}_h be as in Corollary 1, and let C be the Matérn covariance (2.5) satisfying the assumptions in Lemma 1. Then, for any $p \in [1, \infty)$:*

$$\mathbb{E} \left[|Q_h - \tilde{Q}_h|^p \right] \lesssim s^{-\frac{p}{2}} (s - k)^{-\frac{p}{2}(1 + \frac{2\nu}{d})} k^p.$$

Proof. This follows immediately from Corollary 1 and Lemma 1. □

The results in Corollaries 2 and 3 quantify the error introduced by smoothing the random field Z for the application at hand. In addition, they also provide a rule for choosing the number of eigenvalues k_ℓ to drop on each level in the MLMC estimator. More specifically, observe that:

$$\begin{aligned} \mathbb{E}[|Q_{h_\ell} - \tilde{Q}_{h_{\ell-1}}|] &\leq \mathbb{E}[|Q_{h_\ell} - Q_{h_{\ell-1}}|] + \mathbb{E}[|Q_{h_{\ell-1}} - \tilde{Q}_{h_{\ell-1}}|] \\ &\leq C_\alpha h_\ell^\alpha + \mathbb{E}[|Q_{h_{\ell-1}} - \tilde{Q}_{h_{\ell-1}}|]. \end{aligned}$$

For the exponential covariance with $p = 1$, we obtain:

$$\mathbb{E}[|Q_{h_\ell} - \tilde{Q}_{h_{\ell-1}}|] \leq C_\alpha h_\ell^\alpha + C_s s_\ell^{-\frac{1}{2}} (s_\ell - k_\ell)^{-1} k_\ell.$$

Equating the decay rate of the discretisation error with the smoothing error, and using that, in this case, no padding is needed, so that $s_\ell = 4h_\ell^{-2}$, we find the following expression for k_ℓ :

$$k_\ell = \frac{s_\ell^{-\frac{\alpha-3}{2}}}{\frac{C_s}{2C_\alpha} + s_\ell^{-\frac{\alpha-1}{2}}}. \quad (3.13)$$

Analogously, for the Matérn covariance function, we recover:

$$\mathbb{E}[|Q_{h_\ell} - \tilde{Q}_{h_{\ell-1}}|] \leq C_\alpha h_\ell^\alpha + C_s s_\ell^{-\frac{1}{2}} (s_\ell - k_\ell)^{-\frac{1}{2}(1+\nu)} k_\ell.$$

Note that if the smooth periodisation extension in Section 2.2 is applied to ensure the positive definiteness of the embedding matrix, we have $s_\ell = 4(h_\ell^{-1} + J_\ell)^2$, where J_ℓ is the smoothing parameter. In this case, we determine the following rule for choosing k_ℓ :

$$(4(h_\ell^{-1} + J_\ell)^2 - k_\ell)^{-\frac{1}{2}(1+\nu)} k_\ell = \frac{C_\alpha}{2C_s} h_\ell^\alpha (h_\ell^{-1} + J_\ell).$$

In practice, we solve this equation numerically for k_ℓ , with negligible computational cost.

In the context of the MLMC estimator, it is necessary to highlight that in order to obtain an accurate approximation of the solution, only the random field samples on the levels preceding the finest level L should be smoothed. This is because the approximation on level L in the telescopic sum must always use a complete representation of the random field, in order to avoid introducing an additional error in the final estimate. In other words, Eq. (3.7) becomes:

$$\widehat{Q}_L^{\text{MLMC-CES}} = \sum_{\ell=0}^{L-1} \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \tilde{Y}_\ell^{(i)} + \frac{1}{N_L} \sum_{i=1}^{N_L} Q_{h_L}^{(i)} - \tilde{Q}_{h_{L-1}}^{(i)}, \quad (3.14)$$

where $\tilde{Y}_\ell^{(i)} = \tilde{Q}_{h_\ell}^{(i)} - \tilde{Q}_{h_{\ell-1}}^{(i)}$, and $\tilde{Y}_0^{(i)} = \tilde{Q}_{h_0}^{(i)}$. An alternative implementation involves first finding the level $\ell^* < L$ up to which smoothing is necessary using Eq. (3.11) or (3.12), and subsequently applying the smoothing technique only for $\ell = 0, \dots, \ell^*$. However, in our numerical experiments below, this did not appear to be favourable.

Hence, we adapt Algorithm 1 to estimate $\tilde{Y}_{\ell, N_\ell}^{\text{MC}}$ in the MLMC method (3.14) by altering step 4(b) as follows:

Algorithm 2. *Computes an MC estimate of $\mathbb{E}[Y_\ell]$ on a level ℓ using the smoothed circulant embedding method to sample from $\tilde{k}(\cdot, \cdot)$ given the number of samples N_ℓ and the number of eigenvalues to drop k_ℓ .*

(b) *Find permutation π such that $\{\lambda_{\pi(j)}\}_{j=1}^{s_\ell}$ are in non-increasing order. Set $\{\lambda_{\pi(j)}\}_{j=s_\ell-k_\ell+1}^{s_\ell} = 0$ and let $\tilde{\lambda}$ denote the resulting vector. Evaluate $\mathbf{w} := \tilde{\lambda} \odot \boldsymbol{\xi}$.*

4 Numerical Results

In this section, we propose a series of numerical experiments which are intended to establish the performance of the smoothed circulant embedding method. We focus on MLMC applied to the PDE (3.1) with a log-normal random field $k(\mathbf{x}, \omega)$ on $D = (0, 1)^2$ and for $f \equiv 1$. We also offer a brief comparison with the level-dependent truncation of the KL-expansion in [Teckentrup et al., 2013].

We consider two different quantities of interest, namely $Q = u\left(\frac{7}{15}, \frac{7}{15}\right)$ and $Q = \|u(\cdot, \omega)\|_{L^2(D)}$. In fact, the point evaluation of the solution to the PDE (3.1) is not a bounded functional on the Sobolev space $H_0^1(D)$ in domains with more than one dimension. Nonetheless, convergence rates of the form in Proposition 1 can still be obtained as shown in [Teckentrup, 2012].

In the absence of an analytical solution to the two-dimensional PDE in Eq. (3.1), we must perform a set of additional steps in order to obtain an estimate for the MSEs $e\left(\widehat{Q}_{h, N}^{\text{MC}}\right)$ and $e\left(\widehat{Q}_L^{\text{MLMC}}\right)$. In particular, a term which arises in both derivations (3.3) and (3.8) is the bias $(\mathbb{E}[Q_h - Q])^2$, which can be computed by first assuming that the decay in $|\mathbb{E}[Q_h - Q]|$ is monotonic when $h \leq h'$, for some $h' \in \mathbb{Q}$, and that:

$$|\mathbb{E}[Q_h - Q]| = C_\alpha h^\alpha, \quad \alpha > 0.$$

Then, using Richardson's extrapolation [Richardson, 1911] with extrapolant Q_{2h} , we obtain:

$$|\mathbb{E}[Q_h - Q]| = \left| \frac{\mathbb{E}[Q_h - Q_{2h}]}{1 - 2^\alpha} \right|,$$

which we can now use to compute the bias numerically. However, it is important to distinguish that, when smoothing is introduced, the extrapolant used in Richardson's extrapolation is \tilde{Q}_{2h} . Under the following assumption:

$$|\mathbb{E}[\tilde{Q}_h - Q]| = \tilde{C}_\alpha h^\alpha, \quad \alpha > 0,$$

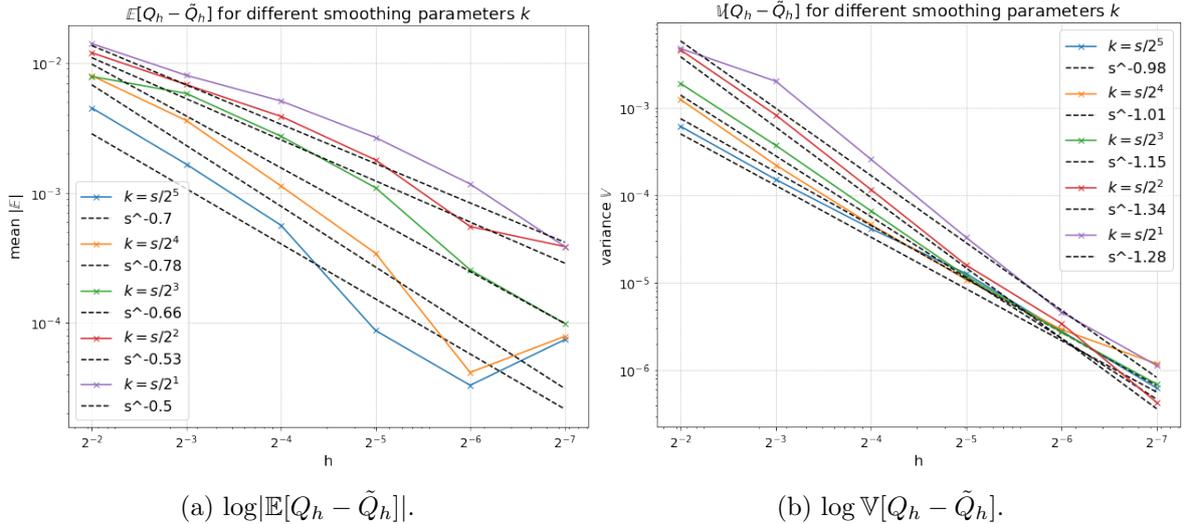


Figure 3: Decay rate of $|\mathbb{E}[Q_h - \tilde{Q}_h]|$ and $\mathbb{V}[Q_h - \tilde{Q}_h]$ with h on a log scale, for $\lambda = 0.1$ and $\sigma^2 = 1$, using different values of k . The quantity of interest is $Q = \|u(\cdot, \omega)\|_{L^2(D)}$.

we can alter the above derivation accordingly to obtain:

$$|\mathbb{E}[Q_h - Q]| = \left| \frac{\mathbb{E}[Q_h - \tilde{Q}_{2h}]}{1 - \frac{\tilde{C}_\alpha}{C_\alpha} 2^\alpha} \right|.$$

While [Teckentrup et al., 2013] gives theoretical values for α and β for the functionals under consideration, here we determine these numerically by estimating $|\mathbb{E}[Q_h - Q_{2h}]|$ on consecutively finer meshes, and computing the best linear fit through these approximations on a log scale. In particular, the expectation of the difference above is replaced by its MC estimate in order to approximate it:

$$|\mathbb{E}[Q_h - Q_{2h}]| = \frac{1}{N} \left| \sum_{i=1}^N Q_h^{(i)} - Q_{2h}^{(i)} \right|,$$

where N is chosen large enough to obtain a suitable estimate, and $Q_h^{(i)}$ and $Q_{2h}^{(i)}$ are computed with the same underlying sample of the random field $k(\mathbf{x}, \omega^{(i)})$. In the case of smoothing, $\tilde{Q}_{2h}^{(i)}$ is calculated with the smoothed sampled used in the approximation of $Q_h^{(i)}$. An analogous approach can be followed to obtain estimates for γ and β . These are necessary for establishing the theoretical bound for the cost of the MLMC algorithm using Theorem 2.

Further, throughout all our numerics, we use $h_\ell = 2^{-\ell} h_0$ as discretisation levels in the MLMC estimator. We use the same mesh size on level ℓ for generating the random field samples.

Finally, we use CPU times to quantify the computational cost of the studied algorithms. These were obtained by running our Python implementation remotely on the Cirrus High Performance Computing system. In addition, we exploit Python's `time` library to measure the CPU time, and quantify the error in terms of the root mean squared error (RMSE), given by $e(\hat{Q}_{h,N}^{\text{MC}})$ and $e(\hat{Q}_L^{\text{MLMC}})$, i.e. the square root of the MSE.

4.1 Smoothing error

Here, we explore a numerical example aimed at demonstrating the theoretical bounds in Theorem 3. We choose the exponential covariance (2.10) with $p = 1$, $\sigma^2 = 1$, and $\lambda = 0.1$, so that we are in the setting of Corollary 2. For the functional $Q = \|u(\cdot, \omega)\|_{L^2(D)}$, we expect $\alpha = 1$

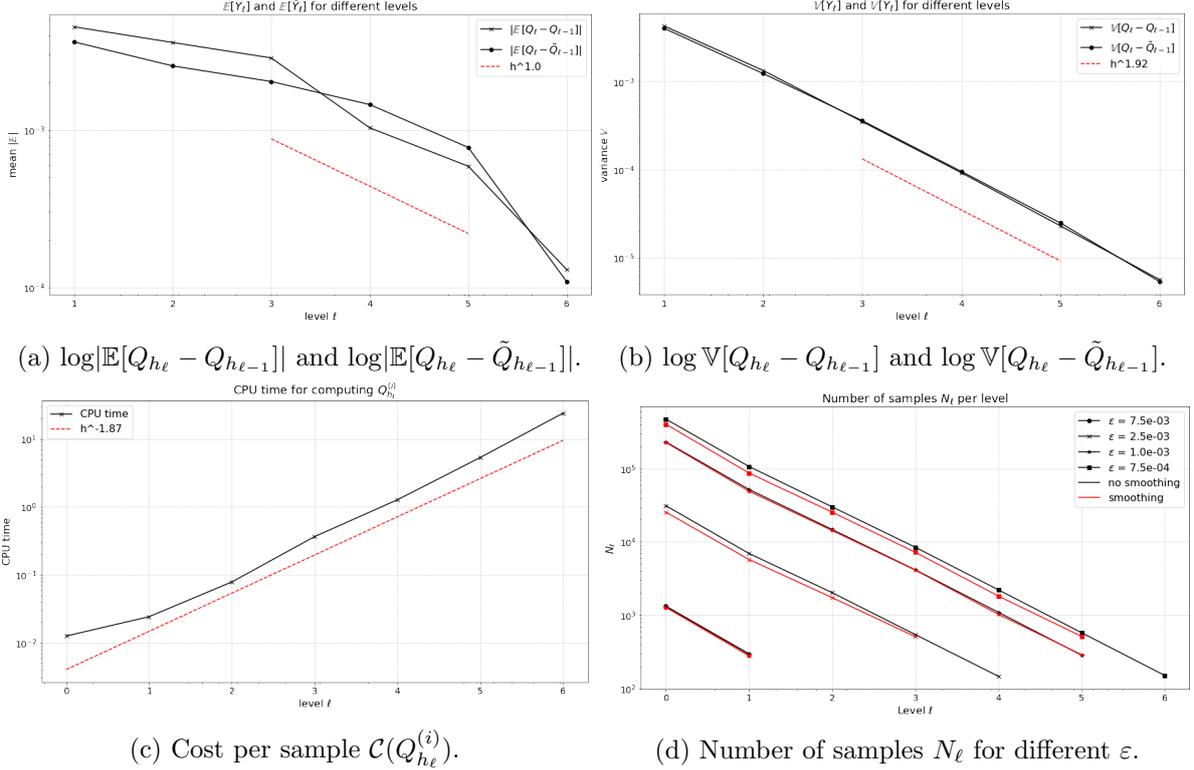


Figure 4: Plots of α , β , and γ and their best linear fit, and of the number of samples N_ℓ on a log scale computed with and without smoothing for $\lambda = 0.3$ and $\sigma^2 = 1$ using MLMC with $\ell = 0, \dots, 6$. The quantity of interest is the mean value of the pressure at $\mathbf{x}^* = (\frac{7}{15}, \frac{7}{15})$.

from [Teckentrup et al., 2013]. Hence, letting $\alpha = 1$ in Eq. (3.13), we obtain:

$$k = \frac{s}{\frac{C_s}{2C_\alpha} + 1}.$$

From Corollary 2, this choice of k yields the following bounds:

$$\mathbb{E}[|Q_h - \tilde{Q}_h|] \lesssim s^{-\frac{1}{2}} \quad \text{and} \quad \mathbb{E}[|Q_h - \tilde{Q}_h|^2] \lesssim s^{-1}.$$

To illustrate the role of the constants C_s and C_α in the decay rate above, we use $k = \frac{s}{a}$ for different values of a . Then, we numerically compute $\mathbb{E}[|Q_h - \tilde{Q}_h|]$ and $\mathbb{V}[Q_h - \tilde{Q}_h] \lesssim \mathbb{E}[|Q_h - \tilde{Q}_h|^2]$ for different values of h , and, implicitly, s .

Thus, in Figure 3, we show the decline in $\mathbb{E}[|Q_h - \tilde{Q}_h|]$ and $\mathbb{V}[Q_h - \tilde{Q}_h]$ with h on a log scale for different values of a . In particular, Figure 3a highlights that for small values of k , such as $k = \frac{s}{25}$, the decay rate of the mean is slightly faster than the expected $s^{-\frac{1}{2}}$. As k becomes larger, for example $k = \frac{s}{2}$, we recover the predicted decay rate. In addition, Figure 3b reproduces an analogous behaviour for the decay rate of the variance $\mathbb{V}[Q_h - \tilde{Q}_h]$. Specifically, for large values of k , we retrieve s^{-1} , as expected, while for small values, we observe a slightly faster decay.

4.2 Random field with exponential covariance

In this section, we analyse the computational complexity of MC and MLMC using circulant embedding with and without smoothing. In particular, we focus on the exponential covariance function (2.10) with norm $p = 1$ and correlation lengths $\lambda = 0.3$ and $\lambda = 0.1$.

Accordingly, Figure 4 illustrates the values obtained for α and β for the MC and MLMC algorithms used to compute the quantity of interest $\mathbb{E}[Q]$. In this case, we model the random

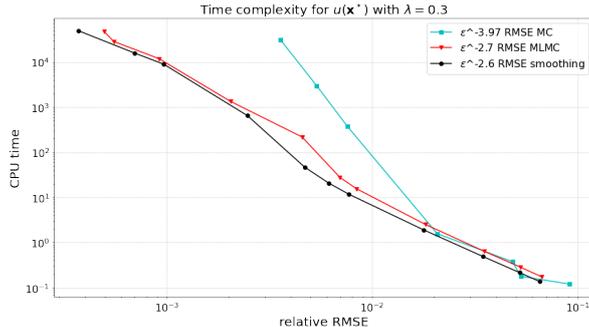


Figure 5: Performance plot on a log scale for $\lambda = 0.3$ and $\sigma^2 = 1$ using Monte Carlo and Multilevel Monte Carlo, with and without smoothing. The quantity of interest is the mean value of the pressure at $\mathbf{x}^* = \left(\frac{7}{15}, \frac{7}{15}\right)$.

coefficient in Eq. (3.1) as a random field whose covariance function has correlation length $\lambda = 0.3$ and variance $\sigma = 1$. As noted in (3.11), the optimal choice for the coarsest MLMC level is such that h_0 is slightly smaller than λ , which yields $h_0 = 2^{-2}$.

Specifically, Figure 4a shows the decay in $\log |\mathbb{E}[Q_h - Q_{2h}]|$, which, due to the definition of the MLMC levels, is equal to $\log |\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]|$, for $\ell = 1, \dots, 6$. The line through these points has a slope of around 1, which implies $\alpha = 1$. Further, we use $k_\ell = \frac{s_\ell}{2}$ for smoothing the samples on level ℓ , and the same coarsest mesh $h_0 = 2^{-2}$ for MLMC with smoothed random field samples. As expected, this gives the same slope α , with a slightly smaller constant \tilde{C}_α .

In addition, Figure 4b portrays the decrease in $\log \mathbb{V}[Y_\ell]$ with level ℓ , $\ell = 1, \dots, 6$. Following the same procedure as above, this gives $\beta \approx 1.92$, which confirms that $\beta \approx 2\alpha$, as expected from the theory. Similar to the α plot, if we introduce smoothing in the MLMC algorithm with $k_\ell = \frac{s_\ell}{2}$, the $\mathbb{V}[Y_\ell]$ values are unaltered, so that β remains unchanged.

Further, Figure 4c depicts the growth in the cost of computing one sample $Q_{h_\ell}^{(i)}$ with level ℓ . On a log scale, the resulting line has a slope of approximately 1.87, so that $\gamma \approx 1.87$. This does not change when smoothing is introduced for the following reasons:

- on the one hand, the main contribution to the cost arises from numerically solving the PDE (3.1) which is invariable to the type of samples involved;
- the complexity of the FFT routine of $\mathcal{O}(s \log s)$ does not depend on the entries in the vector of eigenvalues $\boldsymbol{\lambda} \in \mathbb{R}^s$, only on its dimension.

Moreover, Figure 4d presents the number of samples N_ℓ required per level $\ell = 0, \dots, L$ for different accuracies ε , where N_ℓ is computed using Eq. (3.9). This is shown for the MLMC estimator where the random field samples are smoothed with $k_\ell = \frac{s_\ell}{2}$, as before. For the imposed RMSE, we adopt equally spaced values between $\varepsilon = 10^{-1}$ and $\varepsilon = 5 \cdot 10^{-4}$. This figure also emphasises the discrepancy between the number of samples needed on the coarsest level h_0 and the finest level h_L . For example, to achieve an accuracy of 10^{-3} , five levels are necessary for the MLMC estimator which integrates smoothing technique. Particularly, $N_0 \approx 10^5$ samples are computed on the coarsest level, while $N_5 \approx 10^2$ samples are estimated on the finest level.

Lastly, we expect $\mathcal{C}_\varepsilon \left(\widehat{Q}_{h,N}^{\text{MC}} \right) \lesssim \varepsilon^{-2-\gamma/\alpha} \approx \varepsilon^{-3.87}$ for the MC estimator, as explained at the end of Section 3.3. We also anticipate from Theorem 2 that $\mathcal{C}_\varepsilon \left(\widehat{Q}_L^{\text{MLMC}} \right) \lesssim \varepsilon^{-2}(\log \varepsilon)^2$ for the MLMC estimator both with and without smoothing, as $\beta \approx \gamma$ in both cases. Accordingly, Figure 5 displays, on a log scale, the CPU time taken to achieve a relative RMSE less than ε with each of the MC, MLMC with and without smoothing estimators. For the MLMC simulations we use the same values of ε as in the N_ℓ figure, while for the MC simulations we only consider equally spaced values between $\varepsilon = 10^{-1}$ and $\varepsilon = 5 \cdot 10^{-3}$, due to the long CPU times involved.

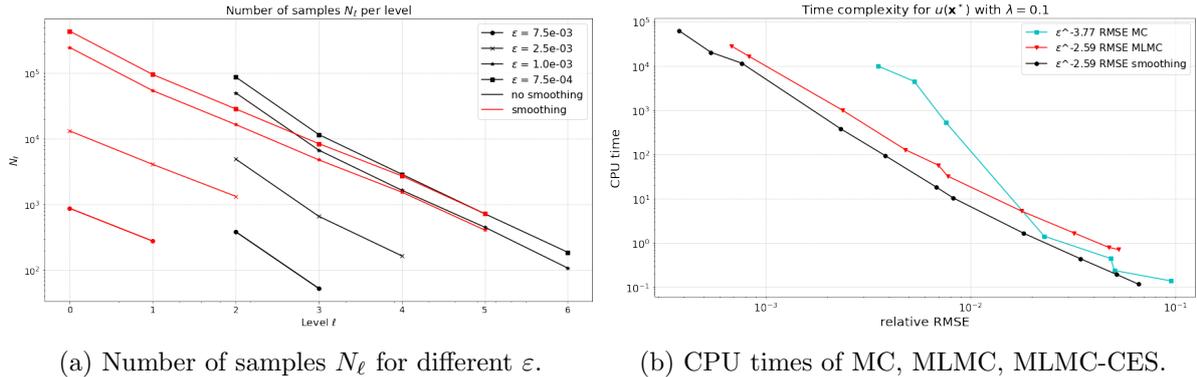


Figure 6: Plot of the number of samples N_ℓ on a log scale and relative RMSE for $\lambda = 0.1$ and $\sigma^2 = 1$ using Monte Carlo and Multilevel Monte Carlo with and without smoothing with $\ell = 0, \dots, 6$. The quantity of interest is the mean value of the pressure at $\mathbf{x}^* = (\frac{7}{15}, \frac{7}{15})$.

Indeed, this graph illustrates that the best performing algorithm is the MLMC with smoothing. In particular, we observe that the MC cost has a slope of roughly -3.97 , as expected, while both versions of the MLMC estimator achieve a computational complexity of $C_\varepsilon \varepsilon^{-2.6}$. The main contribution that the smoothing technique brings is to decrease the constant C_ε . Specifically, if no smoothing is introduced, $C_\varepsilon \approx 7 \cdot 10^{-5}$, while with smoothing we have $C_\varepsilon \approx 5.5 \cdot 10^{-5}$.

Note that integrating the smoothing technique in the MLMC algorithm does not yield major improvements in the computational cost in this case. This is expected because we use $h_0 = 2^{-2}$ for the coarsest grid in both the smoothing and no-smoothing examples. Let us now investigate the case when the covariance function of $k(\mathbf{x}, \omega)$ has correlation length $\lambda = 0.1$. Here, the MLMC method is severely limited in the choice of coarsest mesh, namely $h_0 = 2^{-4}$, so that smoothing the random field samples plays an essential role, since it allows $h_0 = 2^{-2}$ as before.

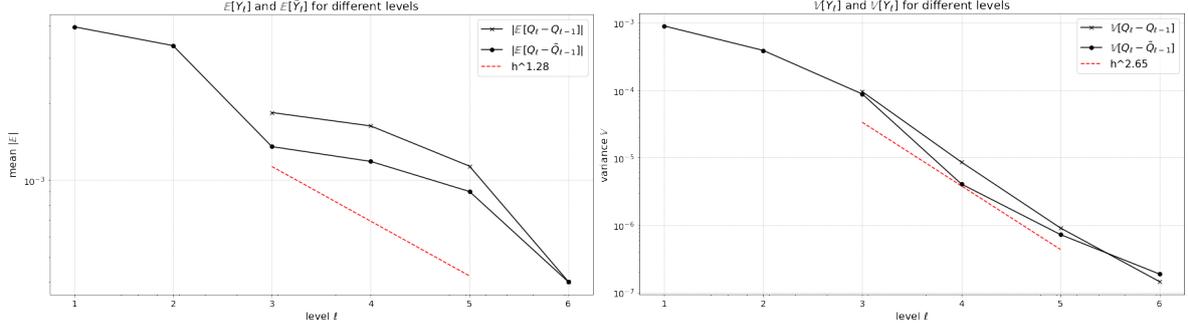
In this case, the plots for α , β , and γ look analogous, as only the constants C_α , C_β and C_γ are different, and so we do not present these here. Rather, we focus on how the number of samples per level N_ℓ varies, and how the overall computational complexities of MC and MLMC with and without smoothing change.

Specifically, Figure 6a illustrates the number of samples N_ℓ per level for different accuracies ε . It is important to highlight that, in this case, level 0 for MLMC without smoothing has a mesh size $h_0 = 2^{-4}$, which corresponds to level 2 for the MLMC with smoothing estimator. Hence, we observe that, for an accuracy of $\varepsilon = 10^{-3}$ for instance, standard MLMC uses five levels, while MLMC-CES uses six levels. However, in the former, the mesh sizes vary from $h_0 = 2^{-4}$ to $h_4 = 2^{-8}$, while in the latter the grid resolutions range from $h_0 = 2^{-2}$ to $h_5 = 2^{-7}$. This yields significant computational savings, as it allows for cheap approximations to be exploited.

The argument above is supported by Figure 6b, which displays the computational complexity of the three surveyed algorithms. In particular, we expect $C_\varepsilon \left(\widehat{Q}_{h,N}^{\text{MC}} \right) \lesssim \varepsilon^{-3.87}$ and $C_\varepsilon \left(\widehat{Q}_L^{\text{MLMC}} \right) = C_\varepsilon \varepsilon^{-2} (\log \varepsilon)^2$ for the MLMC estimator both with and without smoothing, as α , β and γ are as in the $\lambda = 0.3$ case. Indeed, these are very closely recovered in our experiments, as highlighted in this figure. However, one crucial aspect here is the fact that the MC estimator outperforms MLMC for large accuracies, behaviour alleviated by introducing the smoothing. In addition, the constant C_ε decreases from roughly $1.6 \cdot 10^{-4}$ in the no smoothing case to around $6.5 \cdot 10^{-5}$ for the estimator with smoothing, giving rise to the computational savings observed.

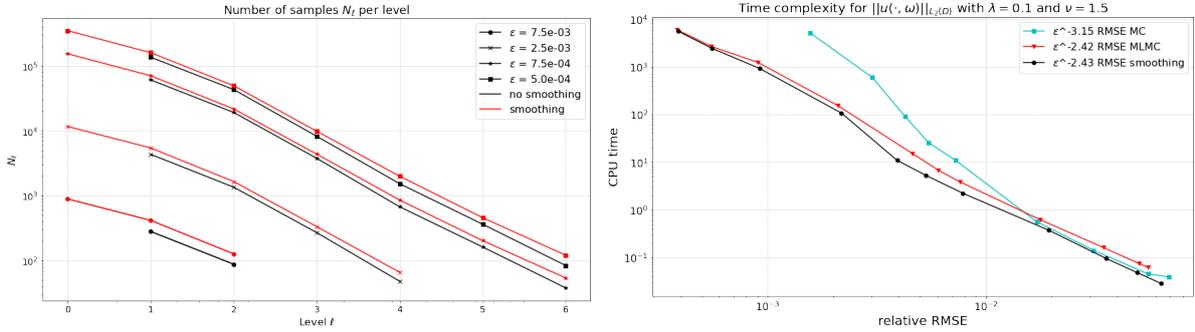
4.3 Random field with Matérn covariance

Throughout this section, we concentrate on the Matérn covariance in Eq. (2.5). We fix the variance $\sigma^2 = 1$ and the smoothness parameter $\nu = 1.5$, and vary the length scale λ . In partic-



(a) $\log|\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]|$ and $\log|\mathbb{E}[Q_{h_\ell} - \tilde{Q}_{\ell-1}]|$. (b) $\log \mathbb{V}[Q_{h_\ell} - Q_{h_{\ell-1}}]$ and $\log \mathbb{V}[Q_{h_\ell} - \tilde{Q}_{\ell-1}]$.

Figure 7: Plots of α and β on a log scale for $\lambda = 0.03$, $\sigma^2 = 1$ and $\nu = 1.5$ using Multilevel Monte Carlo with and without smoothing with $\ell = 0, \dots, 6$. The quantity of interest is the mean value of the L2 norm $\|u(\cdot, \omega)\|_{L^2(D)}$.



(a) Number of samples N_ℓ for different ϵ . (b) CPU times of MC, MLMC, MLMC-CES.

Figure 8: Plot of the number of samples N_ℓ on a log scale and relative RMSE for $\lambda = 0.1$, $\sigma^2 = 1$ and $\nu = 1.5$ using Monte Carlo and Multilevel Monte Carlo with and without smoothing with $\ell = 0, \dots, 6$. The quantity of interest is the mean value of the L2 norm $\|u(\cdot, \omega)\|_{L^2(D)}$.

ular, we explore the computational complexity of MC, MLMC and MLMC-CES for estimating the quantity of interest $Q = \|u(\cdot, \omega)\|_{L^2(D)}$ using $\lambda = 0.1$ and $\lambda = 0.03$.

Note that changing the covariance function does not alter the computational cost of computing one sample $\mathcal{C}(Q_{h_\ell}^{(i)})$, and so we keep $\gamma \approx 1.87$ throughout this section.

Now, Figure 7 portrays the values obtained for α and β for the MC and MLMC algorithms used to compute the quantity of interest $\mathbb{E}[Q]$ when the random field coefficient in (3.1) has length scale $\lambda = 0.03$. Since the two cases $\lambda = 0.1$ and $\lambda = 0.03$ give similar results for α and β , we only focus on the latter here, similar to the previous section. In this case, Eq. (3.12) suggests $h_0 = 2^{-4}$, which is the coarsest mesh we use in estimating α and β without smoothing.

In particular, Figure 7a illustrates the mean order of convergence α . This is achieved by computing $\log|\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]|$ and $\log|\mathbb{E}[Q_{h_\ell} - \tilde{Q}_{\ell-1}]|$ on different levels. Both lines have a slope of $\alpha \approx 1.28$ on average, but with two different constants C_α and \tilde{C}_α .

Moreover, Figure 7b depicts the decay in $\log \mathbb{V}[Q_{h_\ell} - Q_{h_{\ell-1}}]$ and $\log \mathbb{V}[Q_{h_\ell} - \tilde{Q}_{\ell-1}]$ with level ℓ . In particular, this graph indicates the MLMC estimators with and without smoothing have approximately the same $\beta \approx 2.6$, so that we have, again, $\beta \approx 2\alpha$.

In addition, Figure 8a portrays the number of samples N_ℓ needed per level $\ell = 0, \dots, L$ for various accuracies ϵ . This figure demonstrates that the drop from N_L to N_0 is significant, as in the previous example. For instance, to achieve the smallest error considered here, namely $\epsilon = 7.5 \cdot 10^{-4}$, roughly 10^5 samples are necessary on the coarsest level of the MLMC estimator with smoothing with $h_0 = 2^{-1}$. Likewise, around $2 \cdot 10^2$ estimates are computed on level five,

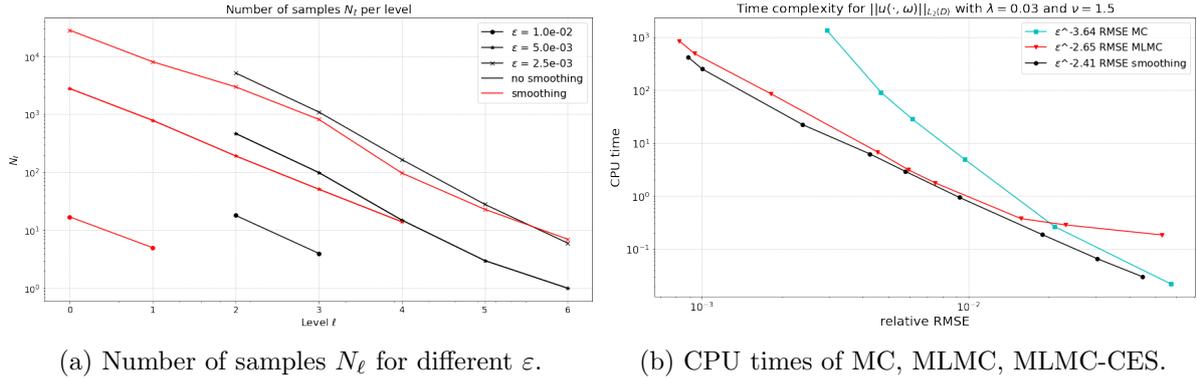


Figure 9: Plot of the number of samples N_ℓ on a log scale and relative RMSE for $\lambda = 0.03$, $\sigma^2 = 1$ and $\nu = 1.5$ using Monte Carlo and Multilevel Monte Carlo with and without smoothing with $\ell = 0, \dots, 6$. The quantity of interest is the mean value of the L2 norm $\|u(\cdot, \omega)\|_{L^2(D)}$.

where $h_5 = 2^{-5}$. The dip is similar for the number of samples in the standard MLMC estimator, even though we use $h_0 = 2^{-2}$ in this case.

Figure 5 displays, on a log scale, the computational efficiency of the three surveyed algorithms to achieve a relative RMSE less than ε , for given ε . We anticipate $\mathcal{C}_\varepsilon(\hat{Q}_{h,N}^{\text{MC}}) \lesssim \varepsilon^{-2-\gamma/\alpha} \approx \varepsilon^{-3.32}$ for the MC estimator, which is comparable to $\varepsilon^{-3.15}$ observed in our numerics. While expect from Theorem 2 that $\mathcal{C}_\varepsilon(\hat{Q}_L^{\text{MLMC}}) \lesssim \varepsilon^{-2}$ as $\gamma < \beta$ for the MLMC estimators with and without smoothing, we still observe $\mathcal{C}_\varepsilon(\hat{Q}_L^{\text{MLMC}}) \lesssim \varepsilon^{-2}(\log \varepsilon)^2$. This could be due to the initial drop in the $\mathbb{V}[Y_\ell]$, which is significantly faster than the decay corresponding to the later levels. Specifically, if we use the last three levels to estimate β , we obtain, in fact, $\beta \approx 2$.

Certainly, this plot establishes, again, that the most efficient method is the MLMC estimator with integrated smoothing. First, the regression line through the points $(\varepsilon, \mathcal{C}_\varepsilon(\hat{Q}_{h,N}^{\text{MC}}))$ has, on a log scale, a slope of -3.15 . Second, both versions of the MLMC approach achieve a cost of roughly $C_\varepsilon \varepsilon^{-2.6}$. The difference rests on the value of C_ε . Specifically, if no smoothing is used, we have $C_\varepsilon \approx 4 \cdot 10^{-5}$, while with smoothing we achieve $C_\varepsilon \approx 2.5 \cdot 10^{-5}$.

Similar to the $\lambda = 0.1$ case in the above section, the computational savings prompted by the smoothing technique are not significant. This is because the coarsest mesh we can use in the standard MLMC estimator does not yield meaningful limitations in terms of number of levels we can exploit. As already mentioned, if we use $\nu = 1.5$ and $\lambda = 0.03$ in the Matérn covariance, we obtain $h_0 = 2^{-4}$, in which case we expect considerable savings in computational cost.

To support this, Figure 9a illustrates the number of samples N_ℓ in the case when $\lambda = 0.3$. Since level 0 for MLMC without smoothing has a mesh size $h_0 = 2^{-4}$, this corresponds to level 2 for the MLMC-CES estimator. While we observe that $N_0 \approx 2 \cdot 10^4$ are necessary for an accuracy of $2.5 \cdot 10^{-3}$ when smoothing is employed, these samples are computed on a grid with resolution 2^{-2} , which are extremely cheap to estimate. In comparison, less than 10 samples are needed on the finest level. On the other hand, around $5 \cdot 10^3$ samples are sufficient for standard MLMC, but these are computed on a mesh of size 2^{-4} .

This discussion is precisely echoed in Figure 9b, which showcases the computational complexity of MC, MLMC and MLMC-CES. In particular, note that the estimator with smoothing outperforms the other two algorithms. While it is substantially better than MC, whose cost satisfies $\mathcal{C}_\varepsilon(\hat{Q}_{h,N}^{\text{MC}}) \lesssim \varepsilon^{-3.64}$, the main gain over the MLMC cost hinges on the constant C_ε in $\mathcal{C}_\varepsilon(\hat{Q}_L^{\text{MLMC}}) = C_\varepsilon \varepsilon^{-2}(\log \varepsilon)^2$. Specifically, for standard MLMC we recover $C_\varepsilon \approx 10^{-5}$, while for MLMC with smoothing we estimate $C_\varepsilon \approx 5 \cdot 10^{-6}$.

This feature is especially visible for larger accuracies ε . In fact, note that for $\varepsilon \geq 2.5 \cdot 10^{-2}$,

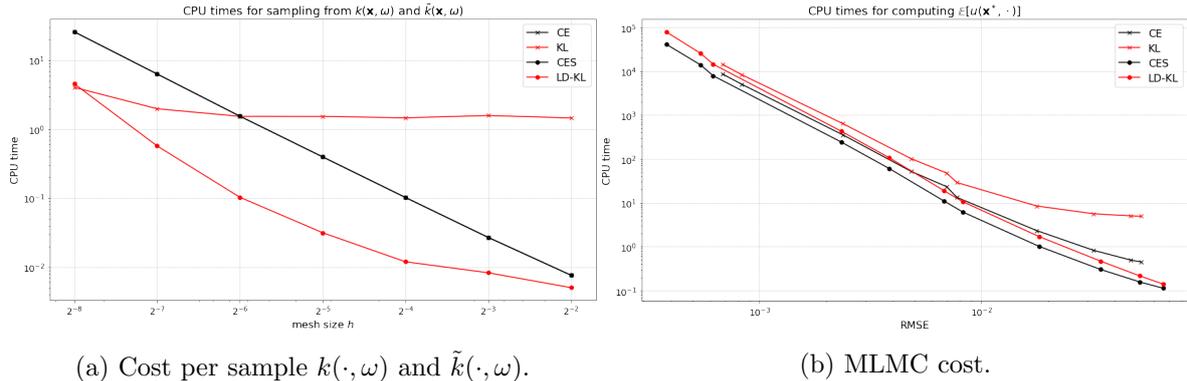


Figure 10: Comparison plots for circulant embedding with and without smoothing, uniform and level-dependent truncation of KL-expansion using covariance (2.10) with parameters $\lambda = 0.1$ and $\sigma^2 = 1$. The quantity of interest is the mean value of the pressure at $\mathbf{x}^* = (\frac{7}{15}, \frac{7}{15})$.

the MLMC estimator with no smoothing is more costly to compute than the standard MC one. On the other hand, if we smooth the random field samples, the resulting procedure is faster than both. In particular, it is 10 times faster than the standard MLMC method for large accuracies. As we decrease the accuracy and more levels are included in the approximation, the performance of the two MLMC methods becomes comparable. Specifically, for $\varepsilon = 10^{-3}$, MLMC-CES is roughly twice as fast than the standard MLMC.

Certainly, the benefits prompted by integrating the smoothing technique in the MLMC algorithm would be more compelling for even smaller correlation lengths. Guided by Eq. (3.11) and (3.12), for $\nu = 0.5$ and $\lambda = 0.01$, we require $h_0 = 2^{-7}$ in the exponential covariance case and $h_0 = 2^{-6}$ in the Matérn case, which limits severely the subsequent grid resolutions we can use. Hence, the circulant embedding with smoothing method would be essential in these cases.

4.4 Comparison with KL-expansion

An alternative solution to the problem at hand is the level-dependent truncation of the KL-expansion of the random field $k(\mathbf{x}, \omega)$, proposed in [Teckentrup et al., 2013; Gittelsohn et al., 2013; Schwab and Stein, 2023]. The aim of this section is to provide a brief comparison of the computational cost of this method and the circulant embedding with smoothing approach in the context of MLMC applied to the PDE (3.1).

To this end, we focus on the exponential covariance in Eq. (2.10) with norm $p = 1$, variance $\sigma^2 = 1$, and correlation length $\lambda = 0.1$, and consider the cost of estimating the mean value of the point evaluation $u(\mathbf{x}^*, \cdot)$, for $\mathbf{x}^* = (\frac{7}{15}, \frac{7}{15})$.

Thus, Figure 10 offers performance plots for the circulant embedding method with and without smoothing and the KL-expansion with uniform and level-dependent truncation. Specifically, Figure 10a illustrates the computational cost achieved for computing one sample of $k(\mathbf{x}, \omega)$ and $\tilde{k}(\mathbf{x}, \omega)$, respectively, for different grid resolutions $h \in \{2^{-2}, \dots, 2^{-8}\}$. As explained in Section 4.2, there is practically no difference in the CE and CES costs. Further, for the KL-expansion, we use $m_{\text{KL}} = 2500$ modes, and for the level-dependent truncation we use the rule in [Teckentrup et al., 2013] to choose the number of modes to include on each mesh. In the first case, the main contribution to the cost stems from computing the eigenvalues and eigenfunctions corresponding to the expansion (2.8), so that the discretisation parameter h is negligible. This explains the plateau up to $h = 2^{-7}$, where the small mesh size induces a more significant cost. In the case of the level-dependent truncation, as the number of modes is considerably smaller on each grid, the computational cost also scales accordingly. In particular, we observe that for one realisation $\tilde{k}(\mathbf{x}, \omega)$, the level-dependent truncation is more efficient than the smoothed circulant embedding.

On the other hand, Figure 10b demonstrates that, once we sample repeatedly from the random field and also take into account the cost of solving the PDE, the MLMC with circulant embedding and its smoothing variant are more cost-effective than MLMC with the uniform and level-dependent truncation of the KL-expansion by roughly a factor of two. This is because, in order to integrate the KL-expansion in the finite element method, in addition to computing the eigenvalues and eigenfunctions in (2.8), for each PDE solve we must perform a matrix-vector multiplication of the form $A\xi$, where $A \in \mathbb{R}^{n \times m_{\text{KL}}}$ and $\xi \in \mathbb{R}^{m_{\text{KL}}}$. Here, m_{KL} is the number of modes we include in the expansion, and n is the number of discretisation points of the mesh we are solving the PDE on. On the other hand, the circulant embedding method is inherently discrete, and so for each PDE solve we need only compute the FFT of $\lambda \odot \xi$, where $\lambda, \xi \in \mathbb{R}^s$, with $s = 4n^2$ generally, which is considerably faster. This makes the computational cost diminish tenfold when using MLMC with standard circulant embedding compared to using a uniformly truncated KL-expansion, particularly for large accuracies.

5 Conclusions and Outlook

Circulant Embedding methods are computationally efficient methods for sampling from random fields, as they exploit the Fast Fourier Transform algorithm to achieve speed-up. In the context of uncertainty quantification using Multilevel Monte Carlo for PDEs with random coefficients, this can be applied to generate realisations from the random parameter.

In some applications, such as groundwater flow, we are interested in the case when the covariance function associated with the random parameter has a short correlation length relative to the size of the computational domain, so that realisations exhibit small-scale fluctuations. In such a case, it is not possible to use very coarse meshes in the MLMC estimator, which limits the subsequent number of levels we can exploit. The circulant embedding with smoothing technique developed here alleviates this issue by producing approximate samples from the random field, which allows to freely choose the mesh size used in the coarsest level approximation of MLMC. This paper looks not only at the theoretical analysis of the corresponding approximation error, but also showcases numerical results which illustrate the computational savings that the smoothing technique can achieve in practical examples.

A straightforward extension is to integrate the proposed approach with MLQMC. As QMC estimators are known to outperform standard MC, we expect that this would yield a more efficient method. In addition, the smoothing technique is similar in nature to the level-dependent truncation of the KL-expansion. Analogous extensions could be developed for the spectral representation method [Rice, 1944; Shinozuka and Jan, 1972], the wavelet reconstruction approach [Zeldin and Spanos, 1996] or the white noise sampling scheme [Lindgren et al., 2011].

Acknowledgements

AI was supported by the EPSRC Centre for Doctoral Training in Mathematical Modelling, Analysis and Computation (MAC-MIGS) funded by the UK Engineering and Physical Sciences Research Council (grant EP/S023291/1), Heriot-Watt University and the University of Edinburgh. This work used the Cirrus UK National Tier-2 HPC Service at EPCC (<http://www.cirrus.ac.uk>) funded by the University of Edinburgh and EPSRC (EP/P020267/1). AI and ALT would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme *Mathematical and statistical foundation of future data-driven engineering* where work on this paper was undertaken. This work was supported by EPSRC grant no EP/R014604/1.

References

- M. S. Alnaes, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, and G. N. Wells. The FEniCS Project Version 1.5. *Archive of Numerical Software*, 3, 2015.
- M. Bachmayr, A. Cohen, and G. Migliorati. Representations of Gaussian Random Fields and Approximation of Elliptic PDEs with Lognormal Coefficients. *J. Fourier Anal. Appl.*, 24: 621–649, 2018.
- M. Bachmayr, I. G. Graham, V. K. Nguyen, and R. Scheichl. Unified Analysis of Periodization-Based Sampling Methods for Matérn Covariances. *SIAM J. Numer. Anal.*, 58:2953–2980, 2020.
- S. Barnett. *Matrices : methods and applications*. Oxford University Press, Oxford, 1990.
- J. Charrier. Strong and Weak Error Estimates for Elliptic Partial Differential Equations with Random Coefficients. *SIAM J. Numer. Anal.*, 50:216–246, 2012.
- J. Charrier, R. Scheichl, and A. L. Teckentrup. Finite Element Error Analysis of Elliptic PDEs with Random Coefficients and Its Application to Multilevel Monte Carlo Methods. *SIAM J. Numer. Anal.*, 51:322–352, 2013.
- K. Cliffe, I. Graham, R. Scheichl, and L. Stals. Parallel Computation of Flow in Heterogeneous Media Modelled by Mixed Finite Elements. *J. Comput. Phys.*, 164:258–282, 2000.
- K. A. Cliffe, M. B. Giles, R. Scheichl, and A. L. Teckentrup. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput. Visual Sci.*, 14:3–15, 2011.
- A. Cohen and R. DeVore. Approximation of high-dimensional parametric PDEs. *Acta Numer.*, 24:1–159, 2015.
- J. Cooley, P. Lewis, and P. Welch. The finite Fourier transform. *IEEE Trans. Audio Electroacoust.*, 17:77–85, 1969.
- M. Croci, M. Giles, and P. E. Farrell. Multilevel Quasi Monte Carlo Methods for Elliptic PDEs with Random Field Coefficients via Fast White Noise Sampling. *SIAM J. Sci. Comput.*, 43: A2840–A2868, Jan. 2021.
- P. J. Davis. *Circulant Matrices*. John Wiley & Sons Inc, New York, 1979.
- G. De Marsily. *Quantitative hydrogeology*. Academic Press, London, 1986.
- C. R. Dietrich and G. N. Newsam. A fast and exact method for multidimensional gaussian stochastic simulations. *Water Resour. Res.*, 29:2861–2869, 1993.
- C. R. Dietrich and G. N. Newsam. Fast and Exact Simulation of Stationary Gaussian Processes through Circulant Embedding of the Covariance Matrix. *SIAM J. Sci. Comput.*, 18:1088–1107, 1997.
- T. J. Dodwell, C. Ketelsen, R. Scheichl, and A. L. Teckentrup. A Hierarchical Multilevel Markov Chain Monte Carlo Algorithm with Applications to Uncertainty Quantification in Subsurface Flow. *SIAM/ASA J. Uncertain. Quantif.*, 3:1075–1108, 2015.
- R. C. Elandt. The Folded Normal Distribution: Two Methods of Estimating Parameters from Moments. *Technometrics*, 3:551–562, 1961.

- X. M. Fernique, J. P. Conze, and J. Gani. *Ecole d'Eté de Probabilités de Saint-Flour IV—1974*, volume 480 of *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1975.
- M. Frigo and S. G. Johnson. The Design and Implementation of FFTW3. *Proceedings of the IEEE*, 93:216–231, 2005.
- M. B. Giles. Multilevel Monte Carlo Path Simulation. *Oper. Res.*, 56:607–617, 2008.
- M. B. Giles. Multilevel Monte Carlo Methods. In J. Dick, F. Y. Kuo, G. W. Peters, and I. H. Sloan, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2012*, pages 83–103, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-41095-6.
- C. J. Gittelsohn, J. Könnö, C. Schwab, and R. Stenberg. The multi-level Monte Carlo finite element method for a stochastic Brinkman Problem. *Numer. Math.*, 125:347–386, 2013.
- I. G. Graham, F. Y. Kuo, D. Nuyens, R. Scheichl, and I. H. Sloan. Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications. *J. Comput. Phys.*, 230:3668–3694, 2011.
- M. D. Gunzburger, C. G. Webster, and G. Zhang. Stochastic finite element methods for partial differential equations with random input data. *Acta Numer.*, 23:521–650, 2014.
- R. Hartley. A More Symmetrical Fourier Analysis Applied to Transmission Problems. *Proc. IRE*, 30:144–150, 1942.
- S. Heinrich. Multilevel monte carlo methods. In *Large-Scale Scientific Computing: Third International Conference, LSSC 2001 Sozopol, Bulgaria, June 6–10, 2001 Revised Papers 3*, pages 58–67. Springer, 2001.
- I. H. Herron and M. R. Foster. *Partial Differential Equations in Fluid Dynamics*. Cambridge University Press, Cambridge, 2008.
- N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, second edition, 2002.
- R. J. Hoeksema and P. K. Kitanidis. Analysis of the Spatial Structure of Properties of Selected Aquifers. *Water Resour. Res.*, 21:563–572, 1985.
- R. Howe. Quantum mechanics and partial differential equations. *J. Funct. Anal.*, 38:188–254, 1980.
- D. D. Kosambi. Statistics in function space. *J. Indian Math. Soc. (N.S.)*, 7:76–88, 1943.
- F. Y. Kuo and D. Nuyens. Application of Quasi-Monte Carlo Methods to PDEs with Random Coefficients – An Overview and Tutorial. In A. B. Owen and P. W. Glynn, editors, *Monte Carlo and Quasi-Monte Carlo Methods*, pages 53–71, Cham, 2018. Springer International Publishing.
- F. Y. Kuo and I. H. Sloan. Lifting the curse of dimensionality. *Notices of the AMS*, 52:1320–1328, 2005.
- F. Y. Kuo, R. Scheichl, C. Schwab, I. H. Sloan, and E. Ullmann. Multilevel Quasi-Monte Carlo methods for lognormal diffusion problems. *Math. Comp.*, 86:2827–2860, 2017.
- F. Lindgren, H. Rue, and J. Lindström. An Explicit Link between Gaussian Fields and Gaussian Markov Random Fields: The Stochastic Partial Differential Equation Approach. *J. R. Stat. Soc., B: Stat. Methodol.*, 73:423–498, 2011.

- A. Logg, K.-A. Mardal, G. N. Wells, and others. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012.
- S. O. Rice. Mathematical analysis of random noise. *Bell Syst. Tech. J.*, 23:282–332, 1944.
- L. F. Richardson. The Approximate Arithmetical Solution by Finite Differences of Physical Problems Involving Differential Equations, with an Application to the Stresses in a Masonry Dam. *Philos. Trans. Royal Soc. A*, 210:307–357, 1911.
- R. Sawko and M. J. Zimoń. Effective Generation of Compressed Stationary Gaussian Fields. *SIAM/ASA J. Uncertain. Quantif.*, 10:439–452, 2022.
- C. Schwab and C. J. Gittelsohn. Sparse tensor discretizations of high-dimensional parametric and stochastic PDEs. *Acta Numer.*, 20:291–467, 2011.
- C. Schwab and A. Stein. Multilevel Monte Carlo FEM for Elliptic PDEs with Besov Random Tree Priors, Feb. 2023. arXiv:2302.00522.
- M. Shinozuka and C.-M. Jan. Digital simulation of random processes and its applications. *J. Sound Vib.*, 25:111–128, 1972.
- A. L. Teckentrup. Multilevel Monte Carlo methods for highly heterogeneous media. In *Proceedings Title: Proceedings of the 2012 Winter Simulation Conference (WSC)*, pages 1–12, Berlin, Germany, 2012. IEEE.
- A. L. Teckentrup, R. Scheichl, M. B. Giles, and E. Ullmann. Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numer. Math.*, 125:569–600, 2013.
- N. E. Tovmasyan, G. Zakaryan, and G. L.Z. *Boundary value problems for partial differential equations and applications in electrodynamics*. World Scientific Publishing, Singapore, 1994.
- B. A. Zeldin and P. D. Spanos. Random Field Representation and Synthesis Using Wavelet Bases. *J. Appl. Mech.*, 63:946–952, 1996.