

Circulant embedding with QMC: analysis for elliptic PDE with lognormal coefficients

Ivan G. Graham¹ · Frances Y. Kuo² ·
Dirk Nuyens³ · Rob Scheichl¹ · Ian H. Sloan²

Received: 25 October 2017 / Revised: 15 March 2018 / Published online: 3 May 2018
© The Author(s) 2018

Abstract In a previous paper (Graham et al. in *J Comput Phys* 230:3668–3694, 2011), the authors proposed a new practical method for computing expected values of functionals of solutions for certain classes of elliptic partial differential equations with random coefficients. This method was based on combining quasi-Monte Carlo (QMC) methods for computing the expected values with circulant embedding methods for sampling the random field on a regular grid. It was found capable of handling fluid flow problems in random heterogeneous media with high stochastic dimension, but no convergence theory was provided. This paper provides a convergence analysis for the method in the case when the QMC method is a specially designed randomly shifted lattice rule. The convergence result depends on the eigenvalues of the underlying nested block circulant matrix and can be independent of the number of stochastic variables under certain assumptions. In fact the QMC analysis applies to general factorisations of the covariance matrix to sample the random field. The error analysis for

✉ Dirk Nuyens
dirk.nuyens@cs.kuleuven.be

Ivan G. Graham
I.G.Graham@bath.ac.uk

Frances Y. Kuo
f.kuo@unsw.edu.au

Rob Scheichl
R.Scheichl@bath.ac.uk

Ian H. Sloan
i.sloan@unsw.edu.au

¹ Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, UK

² School of Mathematics and Statistics, UNSW Sydney, Sydney, NSW 2052, Australia

³ Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001, Leuven, Belgium

the underlying fully discrete finite element method allows for locally refined meshes (via interpolation from a regular sampling grid of the random field). Numerical results on a non-regular domain with corner singularities in two spatial dimensions and on a regular domain in three spatial dimensions are included.

Mathematics Subject Classification 60G10 · 60G60 · 65C05 · 65C60 · 35Q86 · 65D32

1 Introduction

In the paper [12], the present authors proposed a new practical algorithm for solving a class of elliptic partial differential equations with coefficients given by statistically homogeneous lognormal random fields—and in particular for computing expected values of spatial functionals of such solutions. In this algorithm, the required expected value is written as a multidimensional integral of (possibly) high dimension, which is then approximated by a quasi-Monte Carlo (QMC) method. Each evaluation of the integrand is obtained by using a fully discrete finite element (FE) method to approximate the PDE. A key original feature of the method in [12] was the procedure for sampling the random field: instead of sampling the continuous random field by a truncated Karhunen–Loève (KL) expansion, the field was sampled discretely on a regular grid covering the domain and then interpolated at the (irregularly spaced) quadrature points. This completely eliminated the problem of truncation error from the KL expansion, but requires the factorisation of a dense matrix of dimension equal to the number of sample points. In [12] this was done using a circulant embedding technique. The method was found to be effective even for problems with high stochastic dimension, but [12] did not contain a convergence analysis of the algorithm.

The main purpose of the present paper is to provide an analysis for a method closely related to that of [12], with an error bound that is independent of stochastic dimension, and a convergence rate faster than that of a simple Monte Carlo method. The setting differs in two ways from [12]: first, the FE method considered here is the standard nodal FE method for elliptic problems, whereas in [12] the mixed FE method was used; and second, the QMC method considered here is a specially designed randomly shifted lattice rule (see (1.12) below), instead of using Sobol' points as in [12]. (We expect the present analysis can be extended to mixed FEs using results in [14], but do not attempt this here.)

Thus our PDE model (written initially in strong form) is

$$-\nabla \cdot (a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in D \subseteq [0, 1]^d, \quad \text{and almost all } \omega \in \Omega. \quad (1.1)$$

Given a functional \mathcal{G} of u with respect to the spatial variable \mathbf{x} , our aim here (as in [12]) is to compute efficiently and accurately $\mathbb{E}[\mathcal{G}(u)]$, the expected value of $\mathcal{G}(u(\cdot, \omega))$. The (spatial) domain $D \subset \mathbb{R}^d$ ($d = 1, 2, 3$) in (1.1) is assumed to be a bounded interval ($d = 1$), polygon ($d = 2$) or Lipschitz polyhedron ($d = 3$), while Ω is the set of events in a suitable probability space $(\Omega, \mathcal{A}, \mathbb{P})$. The solution u is required to satisfy the

homogeneous Dirichlet condition $u = 0$ on the boundary ∂D of D . The spatial domain D is allowed to be irregular but we assume for convenience that it can be embedded in the d -dimensional unit cube; this is always possible after a suitable affine scaling. (The length-scale of our random field is therefore always considered with respect to the unit cube.) The driving term f is for simplicity taken to be deterministic.

We consider the lognormal case where

$$a(\mathbf{x}, \omega) = \exp(Z(\mathbf{x}, \omega)), \quad (1.2)$$

with $Z(\mathbf{x}, \omega)$ a Gaussian random field with prescribed mean $\bar{Z}(\mathbf{x})$ and covariance

$$r_{\text{cov}}(\mathbf{x}, \mathbf{x}') := \mathbb{E}[(Z(\mathbf{x}, \cdot) - \bar{Z}(\mathbf{x}))(Z(\mathbf{x}', \cdot) - \bar{Z}(\mathbf{x}'))], \quad (1.3)$$

where the expectation is with respect to the Gaussian measure. Lognormal random fields are commonly used in applications, for example in hydrology (see, e.g., [18] and the references there). Throughout we will assume that Z is *stationary* (see, e.g., [1, p. 24]), i.e., its covariance function satisfies

$$r_{\text{cov}}(\mathbf{x}, \mathbf{x}') = \rho(\mathbf{x} - \mathbf{x}'). \quad (1.4)$$

Strictly speaking, ρ only needs to be defined on a sufficiently large ball $B(0, \text{diam}(D))$ for the prescription above, but as in many applications we assume that it is defined on all of \mathbb{R}^d . A particular case that will be discussed extensively, is the Matérn covariance, where ρ is isotropic, i.e., ρ depends on \mathbf{x} only through its Euclidean length $\|\mathbf{x}\|_2$.

In the present paper (1.1) is discretised by piecewise linear finite elements in space, using simplicial meshes with maximum diameter h , and a simple low order quadrature rule for suitable approximation of the associated stiffness matrix. In consequence, the only values of the stochastic coefficient $Z(\mathbf{x}, \omega)$ that enter the FE computation are its values at the quadrature points. However, the FE quadrature points will in general be irregularly distributed, and (for refined meshes) very large in number, typically rendering a direct evaluation of the field at the quadrature points prohibitively expensive. It is much more efficient, as explained below, to instead evaluate exactly the realisation of the field at a uniform grid of

$$M = (m_0 + 1)^d$$

points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M$ on the d -dimensional unit cube $[0, 1]^d$ (containing the domain D), with a fixed integer m_0 and with grid spacing $h_0 := 1/m_0$. We assume further that $h_0 \sim h$. (The extension to general tensor product grids with different mesh sizes in the different coordinate directions is straightforward and not discussed here.) We use *multilinear interpolation* to obtain a sufficiently good approximation of the field $Z(\mathbf{x}, \omega)$ at any other spatial point $\mathbf{x} \in D$, i.e., we use repeated linear interpolation in each coordinate direction with respect to the vertices of the surrounding grid cell.

At this stage of the algorithm, the output is the approximate FE solution $u_h(\mathbf{x}, \omega)$, which inherits randomness from the input data

$$\mathbf{Z}(\omega) := (Z(\mathbf{x}_1, \omega), \dots, Z(\mathbf{x}_M, \omega))^{\top}.$$

The M -vector $\mathbf{Z}(\omega)$ is a Gaussian random vector with a covariance structure inherited from the continuous field Z . Thus it has mean $\bar{\mathbf{Z}} := (\bar{Z}(\mathbf{x}_1), \dots, \bar{Z}(\mathbf{x}_M))^T$ and a positive definite covariance matrix

$$R = [\rho(\mathbf{x}_i - \mathbf{x}_j)]_{i,j=1}^M. \tag{1.5}$$

Because of its finite length, $\mathbf{Z}(\omega)$ can be expressed exactly (but not uniquely) as a linear combination of a finite number of i.i.d. standard normal random variables, i.e., as

$$\mathbf{Z}(\omega) = B\mathbf{Y}(\omega) + \bar{\mathbf{Z}}, \quad \text{where } \mathbf{Y} \sim \mathcal{N}(\mathbf{0}, I_{s \times s}). \tag{1.6}$$

for some real $M \times s$ matrix B with $s \geq M$ satisfying

$$R = BB^T. \tag{1.7}$$

To see this, note that (1.6) and (1.7) imply

$$\mathbb{E}[(\mathbf{Z} - \bar{\mathbf{Z}})(\mathbf{Z} - \bar{\mathbf{Z}})^T] = \mathbb{E}[B\mathbf{Y}\mathbf{Y}^T B^T] = B \mathbb{E}[\mathbf{Y}\mathbf{Y}^T] B^T = BB^T = R.$$

An efficient computation of a suitable factorisation (1.7) using the extension of R to a nested block circulant matrix and then diagonalisation using FFT (the ‘‘circulant embedding method’’) is described in detail in [5, 6, 9, 12]. It is essential for that approach that the random field is sampled on a uniform grid of points. In a related paper [13], we have analysed certain key properties of the circulant extension and its factorisation, which will be crucial for efficiency and for the dimension independence of the QMC convergence analysis in this paper. Other approaches, such as Cholesky factorisation or direct spectral decomposition, could also be used to find a factorisation of the form (1.7). These alternative approaches have the advantage that they do not require the sample grid to be uniform, but when M is large these approaches are likely to be prohibitively expensive. Some ideas of how to overcome this problem using a pivoted Cholesky factorisation or hierarchical matrices can be found in [15] or [4, 10], respectively.

From now on, realisations of the random vector $\mathbf{Y}(\omega)$ are denoted by \mathbf{y} . Thus, \mathbf{y} contains s independent realisations of $\mathcal{N}(0, 1)$. Hence, if $F : \mathbb{R}^s \rightarrow \mathbb{R}$ is any Lebesgue measurable function then the expected value of $F(\mathbf{Y}(\omega))$ may be written as

$$I_s(F) := \int_{\mathbb{R}^s} F(\mathbf{y}) \prod_{j=1}^s \phi(y_j) \, d\mathbf{y} = \int_{(0,1)^s} F(\Phi_s^{-1}(\mathbf{v})) \, d\mathbf{v}, \tag{1.8}$$

where ϕ is the one-dimensional standard normal probability density, and Φ_s^{-1} is the inverse of the cumulative normal distribution function applied componentwise on $(0, 1)^s$. Since $u_h(\cdot, \omega)$ is derived from $\mathbf{Y}(\omega)$, we make the notational convention

$$u_h(\mathbf{x}, \omega) = u_h(\mathbf{x}, \mathbf{y}), \tag{1.9}$$

and so our approximation to $\mathbb{E}[\mathcal{G}(u)]$ is

$$\mathbb{E}[\mathcal{G}(u_h)] = I_s(F), \quad \text{with } F(\mathbf{y}) = \mathcal{G}(u_h(\cdot, \mathbf{y})). \tag{1.10}$$

Because $s \geq M$, the integral (1.8) can have very high dimension. However, two important empirical findings in [12] were that (for all the applications considered) the accuracy of the QMC cubature rule based on Sobol’ points did not appear to be affected by the size of s and that it was always superior to classical Monte Carlo (MC) methods. Successful computations with $s \sim 4 \times 10^6$ were reported in [12]. One aim of the present work is to provide a rigorous proof of the independence of the QMC error on the dimension s (under appropriate conditions) for a specially designed randomly shifted lattice rule. Furthermore, we will also prove here the superior asymptotic convergence rate of QMC over MC in this setting.

The accuracy of the approximation to $\mathbb{E}[\mathcal{G}(u)]$ depends on the FE mesh diameter h (through the FE error and the interpolation error), as well as on the number n of QMC points. We analyse convergence with respect to both these parameters. Our first set of theoretical results concern the accuracy with respect to h . In particular, for bounded linear functionals \mathcal{G} (with respect to the spatial variable \mathbf{x}) and under suitable assumptions, one result, obtained in Sect. 2.2, is that

$$|\mathbb{E}[\mathcal{G}(u) - \mathcal{G}(u_h)]| \leq Ch^{2t}, \tag{1.11}$$

for some parameter $t \in (0, 1]$, determined by the smoothness of realisations of $a(\mathbf{x}, \omega)$, and with a constant C independent of h . This result differs from that of [21] through the use of interpolation to approximate the random vector \mathbf{Z} .

Then, a substantial part of the paper is concerned with the convergence of the QMC rules for (1.8). In particular, we consider randomly shifted lattice rules,

$$Q_{s,n}(\mathbf{\Delta}, F) := \frac{1}{n} \sum_{k=1}^n F\left(\Phi_s^{-1}(\mathbf{v}_k)\right), \quad \text{with } \mathbf{v}_k = \text{frac}\left(\frac{k\mathbf{z}}{n} + \mathbf{\Delta}\right), \tag{1.12}$$

where $\mathbf{z} \in \mathbb{N}^s$ is some suitably chosen generating vector, $\mathbf{\Delta} \in [0, 1]^s$ is a uniformly distributed random shift, and “frac” denotes taking the fractional part of every component in a vector. For the particular integrand $F(\mathbf{y}) := \mathcal{G}(u_h(\cdot, \mathbf{y}))$, we provide in Sect. 3 upper bounds on the error of approximating the integral (1.10) by (1.12). In particular, in Theorem 5, we give sufficient conditions on the matrix B appearing in (1.7) for the root-mean-square error to satisfy an estimate of the form:

$$\sqrt{\mathbb{E}_{\mathbf{\Delta}} [|I_s(F) - Q_{s,n}(\mathbf{\Delta}, F)|^2]} \leq C_{\delta} n^{-(1-\delta)}, \tag{1.13}$$

for arbitrary $\delta > 0$. Here, $\mathbb{E}_{\mathbf{\Delta}}$ denotes expectation with respect to the random shift $\mathbf{\Delta}$. Moreover, we also provide conditions under which the constant C_{δ} in (1.13) is independent of s . Our proof of Theorem 5 differs from the corresponding result in [11] because of the use of multilinear interpolation of the random field in space, and because

of the different meaning of the parameters y_j . The problem is no longer about truncating an infinite Karhunen–Loève expansion, but rather of dealing with a sequence of matrix problems with ever increasing size s .

Finally, combining (1.11) and (1.13), the overall error estimate is

$$\sqrt{\mathbb{E}_\Delta [|\mathbb{E}[\mathcal{G}(u)] - Q_{s,n}(\Delta, F)|^2]} \leq \sqrt{2} C h^{2t} + \sqrt{2} C_\delta n^{-(1-\delta)}.$$

Although the algorithm in [12] applies to both linear and nonlinear functionals, our theory at present is restricted to the linear case. In the numerical experiments, we will use the average of q different random shifts as our final QMC estimator, bringing the total number of integrand evaluations (i.e., PDE solves) to $N = qn$. We compare with a classical Monte Carlo (MC) method for which we use N i.i.d. random samples $w_k \sim U[0, 1)^s$, i.e.,

$$Q_{s,N}^{MC}(F) := \frac{1}{N} \sum_{k=1}^N F(\Phi_s^{-1}(w_k)), \quad \text{with } w_k \sim U[0, 1)^s. \tag{1.14}$$

The layout of the paper is as follows. The PDE with random coefficient and its FE approximation with quadrature are described in Sect. 2.1. The estimate (1.11) is proved in Sect. 2.2. The QMC theory is given in Sect. 3. In particular, one of the key results proved in Sect. 3.3 is the upper bound (1.13). A sufficient condition on B for this result in the case of circulant embedding is identified in Sect. 3.4. The circulant embedding algorithm is summarised briefly in Sect. 3.4 and we refer to [13] for its theoretical analysis. Numerical experiments are given in Sect. 4, illustrating the performance of the algorithm on PDE problems on an irregular domain with corners and holes in two space dimensions, as well as on the unit cube in three dimensions.

2 Finite element implementation and analysis

In Sect. 2.1, we first give the algorithmic details of our practical finite element method, before proving error estimates for this method in Sect. 2.2, in particular Theorem 1 and Corollary 1.

2.1 Model formulation and implementation

We start with (1.1) written pathwise in weak form: seek $u(\cdot, \omega) \in V := H_0^1(D)$ such that

$$\mathcal{A}(\omega; u(\cdot, \omega), v) = \langle f, v \rangle \quad \text{for all } v \in V \text{ and for almost all } \omega \in \Omega, \tag{2.1}$$

where

$$\mathcal{A}(\omega; w, v) := \int_D a(\mathbf{x}, \omega) \nabla w(\mathbf{x}) \cdot \nabla v(\mathbf{x}) \, d\mathbf{x}, \quad w, v \in V,$$

and $a(\mathbf{x}, \omega)$ is given by (1.2)–(1.4). The norm in V is $\|v\|_V := \|\nabla v\|_{L^2(D)}$. For simplicity we assume that $f \in L^2(D)$, so that $\langle f, v \rangle$ reduces to the $L^2(D)$ inner product in (2.1). In general, it denotes the duality pairing between V and its dual space $V' := H^{-1}(D)$.

To discretise (2.1) in the physical domain D , let $\{\mathcal{T}_h\}_{h>0}$ denote a family of conforming, simplicial meshes on D , parametrised by the maximum mesh diameter $h := \max_{\tau \in \mathcal{T}_h} \text{diam}(\tau)$ with $\text{diam}(\tau) := \max_{\mathbf{x}, \mathbf{x}' \in \tau} \|\mathbf{x} - \mathbf{x}'\|_2$. On this mesh, we let $V_h \subset V$ denote the usual finite element space of continuous piecewise linear functions that vanish on ∂D . We assume that $\dim(V_h) = \mathcal{O}(h^{-d})$. This includes many locally refined mesh families, including anisotropic refinement in 3D (e.g., [2,3]). Since any function in V_h has a piecewise constant gradient, we have

$$\mathcal{A}(\omega; w_h, v_h) = \sum_{\tau \in \mathcal{T}_h} a_\tau(\omega) (\nabla w_h(\mathbf{x}) \cdot \nabla v_h(\mathbf{x}))|_\tau \quad \text{for all } v_h, w_h \in V_h, \quad (2.2)$$

where

$$a_\tau(\omega) := \int_\tau a(\mathbf{x}, \omega) \, d\mathbf{x}, \quad \tau \in \mathcal{T}_h. \quad (2.3)$$

We approximate the required integrals (2.3) using a two-stage interpolation/quadrature process as follows. Recall the uniform grid on the cube containing D , with points \mathbf{x}_j for $j = 1, \dots, M$, and grid spacing $h_0 \sim h$, defined in Sect. 1. Let $\mathbf{x} \in \bar{D}$ and let $\{\mathbf{t}_{i,\mathbf{x}}\}_{i=1}^{2^d} \subset \{\mathbf{x}_1, \dots, \mathbf{x}_M\}$ be the vertices of the surrounding grid cell labelled in arbitrary order. Since multilinear interpolation is done by repeatedly applying linear interpolation in each coordinate direction, we can write the interpolated value of g at \mathbf{x} as a convex combination of the surrounding vertex values $\{\mathbf{t}_{i,\mathbf{x}}\}_{i=1}^{2^d}$, i.e.,

$$\mathcal{I}_{h_0} \left(g; \{\mathbf{x}_j\}_{j=1}^M \right) (\mathbf{x}) = \sum_{i=1}^{2^d} w_{i,\mathbf{x}} g(\mathbf{t}_{i,\mathbf{x}}), \quad \text{with} \quad \sum_{i=1}^{2^d} w_{i,\mathbf{x}} = 1 \quad \text{and} \quad 0 \leq w_{i,\mathbf{x}} \leq 1. \quad (2.4)$$

The operator $\mathcal{I}_{h_0} : C(\bar{D}) \rightarrow C(\bar{D})$ is linear and satisfies $\mathcal{I}_{h_0}(g; \{\mathbf{x}_j\}_{j=1}^M)(\mathbf{x}_i) = g(\mathbf{x}_i)$, for every point \mathbf{x}_i of the uniform grid.

Let us further define an r -point quadrature rule on each element τ , which is exact for constant functions, has positive weights $\mu_{\tau,k} \geq 0$ and only uses quadrature points $\mathbf{x}_{\tau,k} \in \tau$, i.e.,

$$Q_\tau(g) := \sum_{k=1}^r \mu_{\tau,k} g(\mathbf{x}_{\tau,k}), \quad \text{with} \quad \sum_{k=1}^r \mu_{\tau,k} = |\tau| \quad \text{and} \quad \mu_{\tau,k} \geq 0. \quad (2.5)$$

Here, $|\tau|$ denotes the volume of τ . The quadrature points $\mathbf{x}_{\tau,k}$ in (2.5) are unrelated to the uniform grid points \mathbf{x}_j in general. Examples of rules satisfying (2.5) are the centroid, nodal or face-centroid rules.

Using the rule (2.5) to approximate all the integrals (2.3) would require evaluating $a(\cdot, \omega)$ at the union of all the (in general irregularly distributed) quadrature points $\{\mathbf{x}_{\tau,k}\}$, which could be costly. We avoid that and compute the field only at the points

of the uniform grid. We then interpolate these values using $\mathcal{I}_{h_0}(a(\cdot, \omega); \{\mathbf{x}_j\}_{j=1}^M)$ and then approximate $a_\tau(\omega)$ using (2.5). In summary, we approximate the bilinear form in (2.2) by

$$\mathcal{A}_h(\omega; w_h, v_h) := \sum_{\tau \in \mathcal{T}_h} \widehat{a}_\tau(\omega) (\nabla w_h \cdot \nabla v_h)|_\tau, \tag{2.6}$$

where

$$\widehat{a}_\tau(\omega) := \mathcal{Q}_\tau(\mathcal{I}_{h_0}(a(\cdot, \omega); \{\mathbf{x}_j\}_{j=1}^M)) = (\mathcal{Q}_\tau \circ \mathcal{I}_{h_0})(a(\cdot, \omega)). \tag{2.7}$$

Proposition 1 *For all $\tau \in \mathcal{T}_h$, there is a sparse positive vector $\mathbf{p}_\tau = (p_{\tau,1}, \dots, p_{\tau,M}) \in \mathbb{R}^M$ such that*

$$\widehat{a}_\tau(\omega) = \sum_{j=1}^M p_{\tau,j} a(\mathbf{x}_j, \omega), \text{ and } \widehat{a}_\tau(\omega) \geq |\tau| a_{\min,M}(\omega), \text{ for all } \tau \in \mathcal{T}_h,$$

where $a_{\min,M}(\omega) := \min_{1 \leq j \leq M} a(\mathbf{x}_j, \omega)$.

Proof It follows from (2.7) together with (2.4) and (2.5) that

$$\widehat{a}_\tau(\omega) = \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} a(\mathbf{t}_{i,\mathbf{x}_{\tau,k}}, \omega). \tag{2.8}$$

The second result then follows from the definition of $a_{\min,M}(\omega)$ and the fact that the coefficients $\mu_{\tau,k}$ and $w_{i,\mathbf{x}_{\tau,k}}$ are all positive and their sum is $|\tau|$. □

Extending the notational convention (1.9), we may thus write our discrete finite element method for (2.1) as the problem of finding $u_h(\cdot, \mathbf{y})$ which satisfies

$$\mathcal{A}_h(\mathbf{y}; u_h(\cdot, \mathbf{y}), v_h) = \langle f, v_h \rangle, \text{ for all } v_h \in V_h \text{ } \mathbf{y} \in \mathbb{R}^S, \tag{2.9}$$

where $\mathcal{A}_h(\mathbf{y}; w_h, v_h)$ is identified with $\mathcal{A}_h(\omega; w_h, v_h)$.

2.2 Finite element error analysis

Let us first define some relevant function spaces. Let $C^1(\overline{D})$ denote the space of continuously differentiable functions on D with seminorm $|\phi|_{C^1(\overline{D})} := \sup_{\mathbf{x} \in \overline{D}} |\nabla \phi(\mathbf{x})|$. For $\beta \in (0, 1)$, let $C^\beta(\overline{D})$ denote the space of Hölder continuous functions on D with exponent β and let $|\phi|_{C^\beta(\overline{D})} := \sup_{\mathbf{x}_1, \mathbf{x}_2 \in \overline{D}: \mathbf{x}_1 \neq \mathbf{x}_2} |\phi(\mathbf{x}_1) - \phi(\mathbf{x}_2)| / \|\mathbf{x}_1 - \mathbf{x}_2\|_2^\beta < \infty$ denote the Hölder coefficient which is, in fact, a seminorm. Also let $L^p(\Omega, X)$ denote the space of all random fields in a Banach space X with bounded p th moments over Ω .

We assume throughout that $Z(\cdot, \omega) \in C^\beta(\overline{D})$, for some $\beta \in (0, 1]$, \mathbb{P} -almost surely. Since $Z(\mathbf{x}, \omega)$ is Gaussian, it follows from Fernique’s Theorem that $\|a\|_{L^p(\Omega, C^\beta(\overline{D}))}$ is finite, for all $p \in [1, \infty)$ (see [7]). Moreover, this implies that $\|a_{\max}\|_{L^p(\Omega)} < \infty$

and $\|1/a_{\min}\|_{L^p(\Omega)} < \infty$, for all $p \in [1, \infty)$, where $a_{\min}(\omega) := \min_{\mathbf{x} \in \overline{D}} a(\mathbf{x}, \omega)$ and $a_{\max}(\omega) := \max_{\mathbf{x} \in \overline{D}} a(\mathbf{x}, \omega)$.

Models where realisations of $a(\mathbf{x}, \omega)$ lack smoothness are often of interest in applications, and a class of coefficients of particular significance is given by the Matérn class with smoothness parameter $\nu \geq 1/2$, described in detail in Example 1. For $\nu \leq 1$, realisations are in $C^\beta(\overline{D})$ \mathbb{P} -almost surely, for all $0 < \beta < \nu$ (see, e.g., [11, 17]).

There are two factors that limit the convergence rate of the finite element error: (i) the regularity of the coefficient field $a(\cdot, \omega)$ and (ii) the shape of the domain D . Since $a(\cdot, \omega) \in C^\beta(\overline{D})$, then (if ∂D is smooth enough), we have $u(\cdot, \omega) \in H^{1+t}(D)$ for all $0 \leq t \leq \beta$. Here, when $\beta < 1$, the loss of H^2 regularity is global and the resulting reduction in the finite element convergence rate cannot be corrected by local mesh refinement. On the other hand, the influence of corner or edge singularities can typically be eliminated by suitable local mesh refinement near ∂D .

Using the notation in [21, Def. 2.1], let $\lambda_\Delta(D)$ be the order of the strongest singularity of the Dirichlet-Laplacian on D . Then $u(\cdot, \omega) \in H^{1+t}(D)$, for all $t \leq \lambda_\Delta(D)$ and $t < \beta$, and $\|u\|_{L^p(\Omega, H^{1+t}(D))}$ is bounded for all $p \in [1, \infty)$ (see [21, Lem. 5.2]). When $\lambda_\Delta(D) \geq \beta$ uniform mesh refinement leads to a best approximation error that satisfies

$$\inf_{v_h \in V_h} \|u(\cdot, \omega) - v_h\|_V \leq C_{FE}(\omega)h^t, \quad \text{for all } t < \beta, \tag{2.10}$$

with $C_{FE}(\omega) \sim \|u(\cdot, \omega)\|_{H^{1+t}(D)}$. When $\lambda_\Delta(D) < \beta$, (2.10) cannot be achieved by uniform refinement. However, it can be recovered by a suitable local refinement. For example, consider the 2D case where D is smooth except for a single reentrant corner with interior angle $\theta > \pi$ and where $W \subset D$ is a local neighbourhood of this corner. Then $\lambda_\Delta(D) = \pi/\theta$ and $u(\cdot, \omega) \in H^{1+t}(D \setminus W)$, for all $t < \beta$, but $u(\cdot, \omega) \notin H^{1+t}(W)$, for $\pi/\theta < t < \beta$. However, by considering the best approximation error over W and over $D \setminus W$ separately, we see that it suffices to grade the meshes such that the mesh size is $\mathcal{O}(h^{\beta\theta/\pi})$ near the reentrant corner and $\mathcal{O}(h)$ away from it. This is because

$$\begin{aligned} \inf_{v_h \in V_h} \|u(\cdot, \omega) - v_h\|_V &\leq C_1 \|u(\cdot, \omega)\|_{H^{1+t}(D \setminus W)} h^t \\ &\quad + C_2 \|u(\cdot, \omega)\|_{H^{1+\lambda_\Delta(D)}(D)} (h^{\beta\theta/\pi})^{\lambda_\Delta(D)} \end{aligned}$$

for all $0 < t < \beta$. Such a mesh grading can often be achieved while retaining the desired complexity estimate $\dim(V_h) \leq Ch^{-2}$ (e.g., [20]).

Thus, using similar techniques to those in the proof of [21, Lem. 5.2] it can be shown that (2.10) holds with $C_{FE}(\omega) \sim \|u(\cdot, \omega)\|_{H^{1+t}(D \setminus W)} + \|u(\cdot, \omega)\|_{H^{1+\lambda_\Delta(D)}(D)}$, for all $t < \beta$. The case of multiple reentrant corners can be treated in an identical fashion. Analogous but more complicated (anisotropic) refinement is needed in 3D, especially in the presence of edge-singularities (e.g., [2, 3]). In practice, such local refinements can be constructed adaptively. The important observation here is that the locally refined mesh needs to be constructed only once for $\exp(\overline{Z}(\cdot))$ (or for one sample of a), since the boundary singularities will be the same for all samples.

We start our analysis by estimating the error in approximating $a_\tau(\omega)$ by $\widehat{a}_\tau(\omega)$.

Lemma 1 Assume $a(\cdot, \omega) \in C^\beta(\overline{D})$ for some $\beta \in (0, 1]$. Furthermore, for $h_0 \sim h$ and for $\tau \subseteq D$ with $\text{diam}(\tau) \leq h$, let \mathcal{I}_{h_0} and \mathcal{Q}_τ be as defined in (2.4) and (2.5), respectively. Then, with $a_\tau(\omega)$ and $\widehat{a}_\tau(\omega)$ given by (2.3) and (2.7),

$$|a_\tau(\omega) - \widehat{a}_\tau(\omega)| \leq |\tau| h^\beta \gamma^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})},$$

with $\gamma = 1 + \sqrt{d} (h_0/h)$. If the quadrature points all lie on the regular grid and we do not need interpolation, we may take $\gamma = 1$.

Proof Using the fact that $a(\cdot, \omega)$ is continuous, the integral mean value theorem asserts the existence of an $\mathbf{x}_\tau^* \in \tau$ such that

$$a_\tau(\omega) = |\tau| a(\mathbf{x}_\tau^*, \omega) = \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} a(\mathbf{x}_\tau^*, \omega),$$

where we used (2.4) and (2.5). Then it follows from (2.8) that

$$\begin{aligned} |a_\tau(\omega) - \widehat{a}_\tau(\omega)| &= \left| \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} (a(\mathbf{x}_\tau^*, \omega) - a(\mathbf{t}_{i,\mathbf{x}_{\tau,k}}, \omega)) \right| \\ &\leq \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} |a(\mathbf{x}_\tau^*, \omega) - a(\mathbf{t}_{i,\mathbf{x}_{\tau,k}}, \omega)| \\ &\leq \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} \|\mathbf{x}_\tau^* - \mathbf{t}_{i,\mathbf{x}_{\tau,k}}\|_2^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})} \\ &\leq \sum_{k=1}^r \sum_{i=1}^{2^d} \mu_{\tau,k} w_{i,\mathbf{x}_{\tau,k}} (\|\mathbf{x}_\tau^* - \mathbf{x}_{\tau,k}\|_2 \\ &\quad + \|\mathbf{x}_{\tau,k} - \mathbf{t}_{i,\mathbf{x}_{\tau,k}}\|_2)^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})} \\ &\leq |\tau| (h + \sqrt{d}h_0)^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})}. \end{aligned}$$

In the last step we used the fact that the distance between a point in a cell of the regular grid and a vertex of that cell is at most $\sqrt{d} h_0$. If the quadrature points all lie on the regular grid then the second term can be omitted. This completes the proof. \square

Theorem 1 Suppose that $Z(\cdot, \omega) \in C^\beta(\overline{D})$ for some $\beta \in (0, 1)$ and suppose that there exists a family $\{\mathcal{I}_h\}_{h>0}$ of conforming, simplicial meshes on D such that (2.10) holds with $\text{dim}(V_h) \leq Ch^{-d}$. Let \mathcal{I}_{h_0} and \mathcal{Q}_τ be defined in (2.4) and (2.5), respectively. Then, we have \mathbb{P} -almost surely

$$\|u(\cdot, \omega) - u_h(\cdot, \omega)\|_V \leq C_{IQ}(\omega) h^t, \quad \text{for all } t < \beta,$$

with C_{IQ} a positive random variable that satisfies $\mathbb{E}[C_{IQ}^p] < \infty$, for all $p \in [1, \infty)$.

If $Z(\cdot, \omega) \in C^1(\overline{D})$ and (2.10) also holds for $t = 1$, then

$$\|u(\cdot, \omega) - u_h(\cdot, \omega)\|_V \leq C_{\mathcal{I}Q}(\omega) h.$$

Proof The proof follows that of [7, Prop. 3.13]. First, using Lemma 1 and the fact that $\nabla v_{h,\tau} := \nabla v_h|_\tau$ is constant, for all piecewise linear finite element functions $v_h \in V_h$, as well as applying the Cauchy–Schwarz inequality in the last step we obtain the estimate

$$\begin{aligned} |\mathcal{A}(\omega; w_h, v_h) - \mathcal{A}_h(\omega; w_h, v_h)| &= \sum_{\tau \in \mathcal{T}_h} |a_\tau(\omega) - \widehat{a}_\tau(\omega)| |(\nabla w_h \cdot \nabla v_h)|_\tau| \\ &\leq h^\beta \gamma^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})} \sum_{\tau \in \mathcal{T}_h} |\tau| |\nabla w_{h,\tau} \cdot \nabla v_{h,\tau}| \\ &\leq h^\beta \gamma^\beta |a(\cdot, \omega)|_{C^\beta(\overline{D})} \|v_h\|_V \|w_h\|_V. \end{aligned}$$

Now, using this bound in Strang’s First Lemma (cf. [7, Lem. 3.12]), we can write

$$\begin{aligned} &\|u(\cdot, \omega) - u_h(\cdot, \omega)\|_V \\ &\leq \inf_{v_h \in V_h} \left\{ \left(1 + \frac{a_{\max}(\omega)}{a_{\min}(\omega)} \right) \|u(\cdot, \omega) - v_h\|_V + h^\beta \gamma^\beta \frac{|a(\omega)|_{C^\beta(\overline{D})}}{a_{\min}(\omega)} \|v_h\|_V \right\}. \end{aligned} \tag{2.11}$$

Since

$$\|v_h\|_V \leq \|u(\cdot, \omega) - v_h\|_V + \|u(\cdot, \omega)\|_V \leq \|u(\cdot, \omega) - v_h\|_V + \frac{\|f\|_{L^2(D)}}{a_{\min}(\omega)},$$

we can combine (2.11) with (2.10) to establish the result.

The fact that the constant $C_{\mathcal{I}Q}(\omega)$ in the above bounds has bounded moments of any (finite) order is a consequence of our assumptions that $Z(\mathbf{x}, \omega)$ is Gaussian and that $Z(\cdot, \omega) \in C^\beta(\overline{D})$. As stated above, it can be proved as in [7] via Fernique’s Theorem. \square

An $\mathcal{O}(h^{2t})$ bound on the L^2 -norm of the error follows via the well-known Aubin–Nitsche trick (cf. [7, Cor. 3.10]). We omit this and finish the section with an error bound for linear functionals \mathcal{G} of u , which we have already stated in (1.11).

Corollary 1 *Let \mathcal{G} be a bounded linear functional on $L^2(D)$. Then, under the assumptions of Theorem 1, there exists a constant $C > 0$ independent of h and u such that*

$$\mathbb{E}[|\mathcal{G}(u) - \mathcal{G}(u_h)|] \leq Ch^{2t}, \quad \text{for all } t < \beta.$$

For $\beta = 1$, we get $\mathbb{E}[|\mathcal{G}(u) - \mathcal{G}(u_h)|] \leq Ch^2$.

Proof The proof follows, as in [21, Lem. 3.3], from Hölder’s inequality using the fact that Theorem 1 applies verbatim also to the FE error $\|z(\cdot, \omega) - z_h(\cdot, \omega)\|_V$ for the dual problem $\mathcal{A}(\omega; v, z(\cdot, \omega)) = \mathcal{G}(v)$, for all $v \in V$. \square

Using the techniques in [21, §3], this corollary can be extended in a straightforward way also to higher order moments of the error or to functionals \mathcal{G} of u that are random, nonlinear or bounded only on a subspace of $L^2(D)$. In summary, we have provided in this section a recipe for extending all results of [21, §3] to general meshes, with the random field being sampled on a regular grid and then interpolated onto the finite element grid.

3 QMC error analysis

The QMC theory for integrals of the form (1.8) is set in a special weighted Sobolev space. Provided the integrand lies in this space, we obtain an estimate for the root mean square error when a specially chosen, randomly shifted lattice rule (1.12) is used to approximate (1.8). The cost for explicitly constructing a good rule tailored to our analysis with n points in s dimensions grows log-linearly in n and quadratically in s (cf. Remark 1 below). However, applying the rule is essentially as cheap as obtaining samples from a random number generator, see, e.g., [16, §7]. Full details of the convergence theory are in other sources, e.g., [11, 16], so we will be brief here.

Later in this section we use this theory to estimate the error when the rule is applied to the particular F given in (1.10). We assume first that the random field Z is sampled by employing the quite general factorisation (1.7) of the covariance matrix R . Later, in Sect. 3.4 we will discuss the case when this is done by circulant embedding.

3.1 Abstract convergence result and proof strategy

The relevant weighted Sobolev norm is defined as:

$$\|F\|_{s,\gamma}^2 := \sum_{\mathbf{u} \subseteq \{1:s\}} \frac{J_{\mathbf{u}}(F)}{\gamma_{\mathbf{u}}}, \tag{3.1}$$

where

$$J_{\mathbf{u}}(F) := \int_{\mathbb{R}^{|\mathbf{u}|}} \left(\int_{\mathbb{R}^{s-|\mathbf{u}|}} \frac{\partial^{|\mathbf{u}|} F}{\partial \mathbf{y}_{\mathbf{u}}}(\mathbf{y}_{\mathbf{u}}; \mathbf{y}_{\{1:s\} \setminus \mathbf{u}}) \prod_{j \in \{1:s\} \setminus \mathbf{u}} \phi(y_j) \, d\mathbf{y}_{\{1:s\} \setminus \mathbf{u}} \right)^2 \times \prod_{j \in \mathbf{u}} \psi_j^2(y_j) \, d\mathbf{y}_{\mathbf{u}}.$$

Here, $\{1:s\}$ denotes the set $\{1, 2, \dots, s\}$, $\frac{\partial^{|\mathbf{u}|} F}{\partial \mathbf{y}_{\mathbf{u}}}$ denotes the mixed first order derivative with respect to the “active” variables y_j with $j \in \mathbf{u}$, $\mathbf{y}_{\{1:s\} \setminus \mathbf{u}}$ denotes the “inactive” variables y_j with $j \notin \mathbf{u}$, and ϕ is the univariate normal probability density (see (1.8)). The remaining ingredients in (3.1) are the *weight parameters* $\gamma_{\mathbf{u}}$ and the *weight functions* ψ_j , which are used, respectively, to moderate the relative importance of the derivatives of F with respect to $\mathbf{y}_{\mathbf{u}}$ and to control the behaviour of these derivatives asymptotically as $\|\mathbf{y}\|_{\infty} \rightarrow \infty$. As in [11], we shall restrict ourselves to the choice

$$\psi_j^2(y_j) = \exp(-2\alpha_j |y_j|), \quad \text{for some } \alpha_j > 0. \tag{3.2}$$

The following result is then essentially [11, Thm. 15] (see also [16]).

Theorem 2 *Suppose $\|F\|_{s,\gamma} < \infty$ and n is a power of a prime. Then, a generating vector $\mathbf{z} \in \mathbb{N}^s$ for a randomly shifted lattice rule (1.12) can be constructed so that the root mean square error in applying (1.12) to (1.8) satisfies*

$$\sqrt{\mathbb{E}_\Delta [|I_s(F) - Q_{s,n}(\Delta, F)|^2]} \leq \left(\frac{2}{n}\right)^{1/(2\kappa)} \tilde{C}_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa) \|F\|_{s,\gamma}, \tag{3.3}$$

for all $\kappa \in (1/2, 1]$, where

$$\tilde{C}_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa) = \left(\sum_{\emptyset \neq \mathbf{u} \subseteq \{1:s\}} \gamma_{\mathbf{u}}^\kappa \prod_{j \in \mathbf{u}} \varrho(\alpha_j, \kappa) \right)^{1/(2\kappa)}, \tag{3.4}$$

$$\varrho(\alpha, \kappa) = 2 \left(\frac{\sqrt{2\pi} \exp(\alpha^2/\eta(\kappa))}{\pi^{2-2\eta(\kappa)}(1 - \eta(\kappa))\eta(\kappa)} \right)^\kappa \zeta\left(\kappa + \frac{1}{2}\right), \quad \eta(\kappa) := \frac{2\kappa - 1}{4\kappa}, \tag{3.5}$$

and ζ is the Riemann zeta function.

A method for constructing the vector \mathbf{z} one component at a time is described in [19] for weights $\gamma_{\mathbf{u}}$ of a special form, see Remark 1 below.

In the remainder of this section we shall apply this theory to the function F given in (1.10). The main result is Theorem 5. It is obtained in the steps summarised as follows.

1. By differentiating the parametrised discrete weak form (2.9), we estimate the norms $\|(\partial^{|\mathbf{u}|} u_h / \partial \mathbf{y}_{\mathbf{u}})(\cdot, \mathbf{y})\|_V$ for any $\mathbf{u} \in \{1 : s\}$. The estimate (which uses an induction argument over the set of all partial derivatives of $u_h(\cdot, \mathbf{y})$) is given in Theorem 3 and involves the quantities

$$b_j := \|\mathbf{B}_j\|_\infty, \quad j = 1, \dots, s, \tag{3.6}$$

where \mathbf{B}_j is the j th column of the $M \times s$ matrix B introduced in (1.7). We let $\mathbf{b} \in \mathbb{R}^s$ be the vector $(b_1, \dots, b_s)^\top$.

2. Using the result from Step 1 and the linearity of \mathcal{G} , we estimate $\|F\|_{s,\gamma}$ in Theorem 4. The shape parameters α_j from (3.2) are constrained to be in the range $\alpha_j > b_j$. The precise values of α_j are arbitrary at this point, and so are the values of the weight parameters $\gamma_{\mathbf{u}}$.
3. We substitute the result from Step 2 into the right hand side of (3.3) to obtain an error bound for this particular F , and we choose the weight parameters $\gamma_{\mathbf{u}}$ and then the shape parameters α_j to minimise this bound. The end result, Theorem 5, is a convergence estimate with order $\mathcal{O}(n^{-1/(2\kappa)})$, valid for $\kappa \in (1/2, 1]$, and with the implied constant depending on the sum $\sum_{j=1}^s b_j^{2\kappa/(1+\kappa)}$.

The theory (in particular Theorem 5) is essentially independent of the choice of factorisation (1.7). In Sects. 3.4 and 4 we then focus on the circulant embedding approach.

3.2 Regularity of F

In this subsection it is helpful to introduce more general partial derivatives than those mixed first order derivatives which appear in (3.1). Thus for any multiindex $\mathbf{v} \in \mathbb{N}^s$ with order $|\mathbf{v}| = \sum v_j$, we let $\partial^{\mathbf{v}}$ denote the corresponding mixed derivative. For any multiindex $\mathbf{v} \in \mathbb{N}^s$ and any vector $\mathbf{c} \in \mathbb{R}^s$ we also write $\mathbf{c}^{\mathbf{v}} = \prod_{j=1}^s c_j^{v_j}$.

Theorem 3 *For any $\mathbf{y} \in \mathbb{R}^s$, any $f \in V'$, and for any multiindex $\mathbf{v} \in \mathbb{N}^s$, the solution $u_h(\cdot, \mathbf{y})$ of (2.9) satisfies*

$$\|\partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})\|_V \leq |\mathbf{v}|! \left(\frac{\mathbf{b}}{\log 2}\right)^{\mathbf{v}} \frac{1}{a_{\min, M}(\mathbf{y})} \|f\|_{V'}$$

where $\mathbf{b} = (b_1, \dots, b_s)^T$ is defined in (3.6) and $a_{\min, M}(\mathbf{y}) = \min_{1 \leq i \leq M} a(\mathbf{x}_i, \mathbf{y})$.

Proof The proof is similar to that of [11, Thm. 14], but there are some differences due to the fact that we are working with the FE discretisation (2.9) with quadrature and interpolation, and because of the finiteness of the ‘expansion’ (1.6). (In [11] an infinite KL expansion was used in the context of the continuous problem.)

To simplify the proof we introduce the \mathbf{y} -dependent discrete norm $\|\cdot\|$ on V_h :

$$\| \| v_h \| \|_{\mathbf{y}}^2 := \sum_{\tau \in \mathcal{T}_h} \hat{a}_{\tau}(\mathbf{y}) |\nabla v_h|^2|_{\tau}, \quad v_h \in V_h, \tag{3.7}$$

with $\hat{a}_{\tau}(\mathbf{y}) = \hat{a}_{\tau}(\omega)$ given by (2.7). Then we have $\| \| v_h \| \|_{\mathbf{y}}^2 = \mathcal{A}_h(\mathbf{y}; v_h, v_h)$, see (2.6). Since we consider piecewise linear finite elements, for $v_h \in V_h$ we have that ∇v_h is piecewise constant on each element τ .

We first prove by induction on $|\mathbf{v}|$ that the solution $u_h(\cdot, \mathbf{y})$ of (2.9) satisfies

$$\| \| \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y}) \| \|_{\mathbf{y}} \leq \Lambda_{|\mathbf{v}|} \mathbf{b}^{\mathbf{v}} \| \| u_h(\cdot, \mathbf{y}) \| \|_{\mathbf{y}}, \quad \mathbf{y} \in \mathbb{R}^s, \tag{3.8}$$

where the sequence $(\Lambda_n)_{n \geq 0}$ is defined recursively by

$$\Lambda_0 := 1 \quad \text{and} \quad \Lambda_n := \sum_{i=0}^{n-1} \binom{n}{i} \Lambda_i, \quad \text{for all } n \geq 1.$$

Clearly (3.8) holds for $|\mathbf{v}| = 0$. For $\mathbf{v} \neq \mathbf{0}$, we differentiate (2.9), using the multivariate Leibniz rule to obtain (since the right-hand side is independent of \mathbf{y}),

$$\sum_{m \leq \mathbf{v}} \binom{\mathbf{v}}{\mathbf{m}} \sum_{\tau \in \mathcal{T}_h} (\partial^{\mathbf{v}-\mathbf{m}} \hat{a}_{\tau}(\mathbf{y})) (\nabla \partial^{\mathbf{m}} u_h(\cdot, \mathbf{y}) \cdot \nabla v_h)|_{\tau} = 0, \quad \text{for all } v_h \in V_h. \tag{3.9}$$

Now inserting $v_h = \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})$ into (3.9), keeping the term with $\mathbf{m} = \mathbf{v}$ in the outer sum on the left-hand side and moving the remaining terms to the right-hand side, we have

$$\begin{aligned} \|\|\| \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}}^2 &= - \sum_{\substack{\mathbf{m} \leq \mathbf{v} \\ \mathbf{m} \neq \mathbf{v}}} \binom{\mathbf{v}}{\mathbf{m}} \sum_{\tau \in \mathcal{I}_h} (\partial^{\mathbf{v}-\mathbf{m}} \widehat{a}_{\tau}(\mathbf{y})) (\nabla \partial^{\mathbf{m}} u_h(\cdot, \mathbf{y}) \cdot \nabla \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})) \Big|_{\tau} \\ &\leq \sum_{\substack{\mathbf{m} \leq \mathbf{v} \\ \mathbf{m} \neq \mathbf{v}}} \binom{\mathbf{v}}{\mathbf{m}} \left(\max_{\tau \in \mathcal{I}_h} \left| \frac{\partial^{\mathbf{v}-\mathbf{m}} \widehat{a}_{\tau}(\mathbf{y})}{\widehat{a}_{\tau}(\mathbf{y})} \right| \right) \sum_{\tau \in \mathcal{I}_h} \widehat{a}_{\tau}(\mathbf{y}) |(\nabla \partial^{\mathbf{m}} u_h(\cdot, \mathbf{y}) \cdot \nabla \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})) \Big|_{\tau}|. \end{aligned}$$

Then, after an application of the Cauchy–Schwarz inequality and a cancellation, we obtain

$$\|\|\| \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}} \leq \sum_{\substack{\mathbf{m} \leq \mathbf{v} \\ \mathbf{m} \neq \mathbf{v}}} \binom{\mathbf{v}}{\mathbf{m}} \left(\max_{\tau \in \mathcal{I}_h} \left| \frac{\partial^{\mathbf{v}-\mathbf{m}} \widehat{a}_{\tau}(\mathbf{y})}{\widehat{a}_{\tau}(\mathbf{y})} \right| \right) \|\|\| \partial^{\mathbf{m}} u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}} . \tag{3.10}$$

To estimate (3.10), we have from Proposition 1 that $\widehat{a}_{\tau}(\mathbf{y}) = \sum_{i=1}^M p_{\tau,i} a(\mathbf{x}_i, \mathbf{y})$ with all $p_{\tau,i} \geq 0$, and we recall from (1.2) and (1.6) that $a(\mathbf{x}_i, \mathbf{y}) = \exp(\sum_{j=1}^s B_{i,j} y_j + \overline{Z}_i)$. Then, noting that $a(\mathbf{x}_i, \mathbf{y}) \geq 0$ it is easy to see that, for any multiindex \mathbf{v} ,

$$|\partial^{\mathbf{v}} a(\mathbf{x}_i, \mathbf{y})| = a(\mathbf{x}_i, \mathbf{y}) \prod_{j=1}^s |B_{i,j}^{v_j}| \leq a(\mathbf{x}_i, \mathbf{y}) \mathbf{b}^{\mathbf{v}} ,$$

which leads to $|\partial^{\mathbf{v}} \widehat{a}_{\tau}(\mathbf{y})| \leq \widehat{a}_{\tau}(\mathbf{y}) \mathbf{b}^{\mathbf{v}}$, and hence

$$\max_{\tau \in \mathcal{I}_h} \left| \frac{\partial^{\mathbf{v}-\mathbf{m}} \widehat{a}_{\tau}(\mathbf{y})}{\widehat{a}_{\tau}(\mathbf{y})} \right| \leq \mathbf{b}^{\mathbf{v}-\mathbf{m}} .$$

Inserting this into (3.10) we obtain

$$\|\|\| \partial^{\mathbf{v}} u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}} \leq \sum_{\substack{\mathbf{m} \leq \mathbf{v} \\ \mathbf{m} \neq \mathbf{v}}} \binom{\mathbf{v}}{\mathbf{m}} \mathbf{b}^{\mathbf{v}-\mathbf{m}} \|\|\| \partial^{\mathbf{m}} u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}} . \tag{3.11}$$

Using (3.11), the estimate (3.8) then follows by induction in exactly the same way as in [11, Thm. 14].

Now, using the definition of the discrete norm (3.7) and the fact that $u_h(\cdot, \mathbf{y})$ is the solution of (2.9), we have

$$\left(\min_{\tau \in \mathcal{I}_h} \frac{\widehat{a}_{\tau}(\mathbf{y})}{|\tau|} \right) \|u_h(\cdot, \mathbf{y})\|_{\mathbf{V}}^2 \leq \|\|\| u_h(\cdot, \mathbf{y}) \|\|\|_{\mathbf{y}}^2 = \langle f, u_h(\cdot, \mathbf{y}) \rangle \leq \|f\|_{\mathbf{V}'} \|u_h(\cdot, \mathbf{y})\|_{\mathbf{V}} . \tag{3.12}$$

Hence, using Proposition 1, we conclude that

$$\|u_h(\cdot, \mathbf{y})\|_V \leq \frac{\|f\|_{V'}}{a_{\min, M}(\mathbf{y})} \quad \text{and} \quad \| \|u_h(\cdot, \mathbf{y})\| \|_{\mathbf{y}} \leq \frac{\|f\|_{V'}}{\sqrt{a_{\min, M}(\mathbf{y})}}. \tag{3.13}$$

Using the same argument as for the lower bound in (3.12), together with Proposition 1 again, we obtain from (3.8) and (3.13) that

$$\sqrt{a_{\min, M}(\mathbf{y})} \|\partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})\|_V \leq \| \|\partial^{\mathbf{v}} u_h(\cdot, \mathbf{y})\| \|_{\mathbf{y}} \leq \Lambda_{|\mathbf{v}|} \mathbf{b}^{\mathbf{v}} \frac{\|f\|_{V'}}{\sqrt{a_{\min, M}(\mathbf{y})}}.$$

This together with the estimate $\Lambda_n \leq n! / (\log 2)^n$ (proved in [11, Thm. 14]) completes the proof. \square

We can now use this theorem to show that F lies in the weighted Sobolev space characterised by the norm (3.1). We make use of the s -dependent quantities

$$\|\mathbf{b}\|_{p,s} = \left(\sum_{j=1}^s |b_j|^p \right)^{1/p}, \quad p > 0, \quad \text{and} \quad \|\mathbf{b}\|_{\infty,s} = \max_{j \in \{1:s\}} |b_j|.$$

Theorem 4 *Suppose that $\|\mathbf{b}\|_{1,s}$ is uniformly bounded with respect to s . Suppose also that $\alpha_j > b_j$ for all j . Then, for any $f \in V'$ and for any linear functional $\mathcal{G} \in V'$, the integrand $F(\mathbf{y}) = \mathcal{G}(u_h(\cdot, \mathbf{y}))$ in (1.10) satisfies*

$$\|F\|_{s,\gamma} \leq C \left(\sum_{\mathbf{u} \subseteq \{1:s\}} \frac{1}{\gamma_{\mathbf{u}}} \left(\frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \right)^2 \prod_{j \in \mathbf{u}} \frac{\tilde{b}_j^2}{\alpha_j - b_j} \right)^{1/2},$$

where C is independent of s and

$$\tilde{b}_j = \frac{b_j}{2 \exp(b_j^2/2) \Phi(b_j)},$$

with Φ denoting the univariate standard cumulative normal distribution function.

Proof Using the linearity of \mathcal{G} , together with Theorem 3, but replacing the multiindex \mathbf{v} with any set $\mathbf{u} \subseteq \{1 : s\}$ (i.e., restricting to the case where all $v_j \leq 1$), we obtain the following estimate for the first order partial derivatives of F appearing in the norm (3.1),

$$\begin{aligned} \left| \frac{\partial^{|\mathbf{u}|} F}{\partial \mathbf{y}_{\mathbf{u}}}(\mathbf{y}) \right| &\leq \frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \left(\prod_{j \in \mathbf{u}} b_j \right) \frac{1}{a_{\min, M}(\mathbf{y})} \|f\|_{V'} \|\mathcal{G}\|_{V'} \\ &\leq \frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \left(\exp(\mathbf{b}^{\top} |\mathbf{y}|) \prod_{j \in \mathbf{u}} b_j \right) \left(\exp(\|\bar{\mathbf{Z}}\|_{\infty}) \|f\|_{V'} \|\mathcal{G}\|_{V'} \right), \end{aligned} \tag{3.14}$$

where we used the estimate $a_{\min, M}(\mathbf{y}) = \min_{1 \leq i \leq M} a(\mathbf{x}_i, \mathbf{y}) \geq \exp(-\|\bar{\mathbf{Z}}\|_\infty) \exp(-\mathbf{b}^\top |\mathbf{y}|)$.

Examining the right-hand side of (3.14), we see that the only factor which depends on \mathbf{y} is $\exp(\mathbf{b}^\top |\mathbf{y}|)$. An elementary calculation (see [11, Thm. 16]) shows that

$$\begin{aligned} & \int_{\mathbb{R}^{|\mathbf{u}|}} \left(\int_{\mathbb{R}^{s-|\mathbf{u}|}} \left(\exp(\mathbf{b}^\top |\mathbf{y}|) \prod_{j \in \mathbf{u}} b_j \right) \prod_{j \in \{1:s\} \setminus \mathbf{u}} \phi(y_j) \, d\mathbf{y}_{\{1:s\} \setminus \mathbf{u}} \right)^2 \prod_{j \in \mathbf{u}} \psi_j^2(y_j) \, d\mathbf{y}_{\mathbf{u}} \\ &= \left(\prod_{j \in \{1:s\} \setminus \mathbf{u}} \left(2\Phi(b_j) \exp(b_j^2/2) \right)^2 \right) \left(\prod_{j \in \mathbf{u}} \frac{b_j^2}{\alpha_j - b_j} \right) \\ &= \left(\prod_{j \in \{1:s\}} \left(2\Phi(b_j) \exp(b_j^2/2) \right)^2 \right) \left(\prod_{j \in \mathbf{u}} \frac{\tilde{b}_j^2}{\alpha_j - b_j} \right). \end{aligned}$$

Since $2\Phi(b_j) \leq 1 + 2b_j/\sqrt{2\pi} < \exp(b_j)$ for all j , we have

$$\prod_{j \in \{1:s\}} \left(2\Phi(b_j) \exp(b_j^2/2) \right)^2 \leq \exp(2\|\mathbf{b}\|_{1,s} + \|\mathbf{b}\|_{2,s}^2).$$

Thus, it follows from (3.14) and the definition of the norm (3.1) that

$$\begin{aligned} \frac{\|F\|_{s,\boldsymbol{\gamma}}}{\exp(\|\bar{\mathbf{Z}}\|_\infty) \|f\|_{V'} \|\mathcal{G}\|_{V'}} &\leq \exp(2\|\mathbf{b}\|_{1,s} + \|\mathbf{b}\|_{2,s}^2) \\ &\times \left(\sum_{\mathbf{u} \subseteq \{1:s\}} \frac{1}{\gamma_{\mathbf{u}}} \left(\frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \right)^2 \prod_{j \in \mathbf{u}} \frac{\tilde{b}_j^2}{\alpha_j - b_j} \right)^{1/2}. \end{aligned}$$

The final result, with the constant factor C being independent of s , is then a consequence of the assumption on \mathbf{b} . □

3.3 Error estimate

In order to obtain a dimension-independent estimate for the QMC method we need a stronger assumption on \mathbf{b} than that used in Theorem 4. The following theorem shows that under that stronger assumption there is a choice of $\boldsymbol{\gamma}$ and $\boldsymbol{\alpha}$ which ensures that the QMC error is bounded independently of s . The appropriate choice of $\boldsymbol{\gamma}$ is of ‘‘POD’’ type, which allows a good generating vector \mathbf{z} for the QMC rule to be efficiently computed by the ‘‘component-by-component’’ procedure, see Remark 1 below.

Theorem 5 *Under the assumptions of Theorem 4, let $\kappa \in (1/2, 1)$, set $p = 2\kappa/(1+\kappa)$ and assume in addition that $\|\mathbf{b}\|_{p,s}$ is uniformly bounded with respect to s . Then there exists a positive constant $C(\kappa)$ depending on κ (as well as on $\bar{\mathbf{Z}}, f, \mathcal{G}$) such that*

$$\sqrt{\mathbb{E}_{\boldsymbol{\Delta}} [|I_s(F) - Q_{s,n}(\boldsymbol{\Delta}, F)|^2]} \leq C(\kappa) n^{-1/(2\kappa)}. \tag{3.15}$$

Proof Some parts of the proof are similar to that of [11, Thm. 20], and for these parts we will be brief. We remind the readers that each vector $\mathbf{b} = (b_1, \dots, b_s)^\top$ depends fundamentally on s ; changing the value of s leads to completely different components b_j . This is different from the situation in [11] where there is just one infinite sequence \mathbf{b} that is truncated to s terms.

First, combining Theorems 2 and 4 we see that (3.15) holds with $C(\kappa)$ proportional to

$$C_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa) := \tilde{C}_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa) \left(\sum_{\mathbf{u} \subseteq \{1:s\}} \frac{1}{\gamma_{\mathbf{u}}} \left(\frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \right)^2 \prod_{j \in \mathbf{u}} \frac{\tilde{b}_j^2}{\alpha_j - b_j} \right)^{1/2},$$

with $\tilde{C}_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa)$ defined in (3.4). Now, we choose the weight parameters $\boldsymbol{\gamma}$ to minimise $C_s(\boldsymbol{\gamma}, \boldsymbol{\alpha}, \kappa)$. This minimisation problem was solved in [11, Lem. 18], yielding the solution:

$$\gamma_{\mathbf{u}} = \gamma_{\mathbf{u}}^* := \left(\left(\frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \right)^2 \prod_{j \in \mathbf{u}} \frac{\tilde{b}_j^2}{(\alpha_j - b_j) \varrho(\alpha_j, \kappa)} \right)^{1/(1+\kappa)}, \tag{3.16}$$

which is of ‘‘product and order dependent’’ (POD) form. With this choice, one can show that

$$C_s(\boldsymbol{\gamma}^*, \boldsymbol{\alpha}, \kappa) = S_s(\boldsymbol{\alpha}, \kappa)^{(\kappa+1)/(2\kappa)},$$

where

$$S_s(\boldsymbol{\alpha}, \kappa) = \sum_{\mathbf{u} \subseteq \{1:s\}} \left(\left(\frac{|\mathbf{u}|!}{(\log 2)^{|\mathbf{u}|}} \right)^2 \prod_{j \in \mathbf{u}} \frac{\varrho^{1/\kappa}(\alpha_j, \kappa) \tilde{b}_j^2}{\alpha_j - b_j} \right)^{\kappa/(1+\kappa)}. \tag{3.17}$$

It remains to estimate $S_s(\boldsymbol{\alpha}, \kappa)$. Apart from the constraint $\alpha_j > b_j$, the shape parameters α_j are still free at this stage and so we choose them to minimise the right-hand side of (3.17). This minimisation problem is also solved in [11, Cor. 21]); the solution is

$$\alpha_j := \frac{1}{2} \left(b_j + \sqrt{b_j^2 + 1 - \frac{1}{2\kappa}} \right). \tag{3.18}$$

Now, to estimate $S_s(\boldsymbol{\alpha}, \kappa)$, let b_{\max} be an upper bound on $\|\mathbf{b}\|_{\infty, s}$ for all s (guaranteed by assumption). Then $b_j \leq b_{\max}$ for all $j = 1, \dots, s$ and all s . For a given value of $\kappa \in (1/2, 1]$, let α_{\max} denote the value of (3.18) with b_j replaced by b_{\max} . Then we have $\alpha_j \leq \alpha_{\max}$ for all $j = 1, \dots, s$ and all s , and

$$\begin{aligned} \alpha_j - b_j &= \frac{1}{2} \frac{1 - 1/(2\kappa)}{\sqrt{b_j^2 + 1 - 1/(2\kappa)} + b_j} \geq \frac{1}{2} \frac{1 - 1/(2\kappa)}{\sqrt{b_{\max}^2 + 1 - 1/(2\kappa)} + b_{\max}} \\ &= \alpha_{\max} - b_{\max}. \end{aligned}$$

Note also that for ϱ defined in (3.5) we have $\varrho(\alpha_j, \kappa) \leq \varrho(\alpha_{\max}, \kappa)$ for all j and all s . Moreover, since $2 \exp(b_j^2/2)\Phi(b_j) \geq 1$, it follows that $\tilde{b}_j \leq b_j$, and so from (3.17) we can conclude that, with $p := 2\kappa/(1 + \kappa)$,

$$S_s(\alpha, \kappa) \leq \sum_{u \subseteq \{1:s\}} (|u|!)^p \prod_{j \in u} (\tau_\kappa b_j^2)^{p/2} = \sum_{\ell=0}^s (\ell!)^p \sum_{u \subseteq \{1:s\}, |u|=\ell} \prod_{j \in u} (\tau_\kappa b_j^2)^{p/2},$$

where

$$\tau_\kappa := \frac{\varrho^{1/\kappa}(\alpha_{\max}, \kappa)}{(\log 2)^2(\alpha_{\max} - b_{\max})}.$$

Now using the inequality

$$\ell! \sum_{\substack{u \subseteq \{1:s\} \\ |u|=\ell}} \prod_{j \in u} a_j \leq \left(\sum_{j=1}^s a_j \right)^\ell,$$

(which holds since for $|u| = \ell$ each term $\prod_{j \in u} a_j$ from the left-hand side appears in the expansion of the right-hand side exactly $\ell!$ times, but the right-hand side includes other terms) and with K denoting the assumed uniform bound on $\|b\|_{p,s}^p$, we obtain

$$S_s(\alpha, \kappa) \leq \sum_{\ell=0}^s (\ell!)^{p-1} \tau_\kappa^{p\ell/2} \left(\sum_{j=1}^s b_j^p \right)^\ell \leq \sum_{\ell=0}^\infty (\ell!)^{p-1} \tau_\kappa^{p\ell/2} K^\ell < \infty.$$

The finiteness of the right-hand side follows by the ratio test, on noting that $p < 1$. \square

Remark 1 A generating vector $z \in \mathbb{N}^s$ for a randomly shifted lattice rule with n points in s dimensions that achieves the desired error bound can be constructed using a component-by-component (CBC) algorithm, which goes as follows: (1) Set $z_1 = 1$. (2) For each $k = 2, 3, \dots, s$, choose z_k from the set $\{1 \leq z \leq n - 1 : \gcd(z, n) = 1\}$ to minimise

$$E_{s,n,k}^2(z_1, \dots, z_k) := \sum_{\emptyset \neq u \subseteq \{1:k\}} \frac{\gamma_u}{n} \sum_{i=1}^n \prod_{j \in u} \theta_j \left(\frac{iz_j}{n} \right),$$

where the function $\theta_j(x)$ is symmetric around $1/2$ for $x \in [0, 1]$ and can be computed for $x \in [0, 1/2]$ by

$$\theta_j(x) := \frac{x - \frac{1}{2} + \exp(2\alpha_j^2) [\Phi(2\alpha_j) - \Phi(2\alpha_j + \Phi^{-1}(x))]}{\alpha_j} - 2 \int_{-\infty}^0 \frac{\Phi(t)^2}{\psi_j(t)^2} dt.$$

The integral in the above formula for θ_j only needs to be calculated once, while the general formulation [19, Equation (50)] also has an integral for the first part,

which we evaluate explicitly here for our particular choice of ψ_j . Note that γ_u and α_j fundamentally depend on s through b_j . When (and only when) the algorithm reaches $k = s$, the expression $E_{s,n,s}^2(z_1, \dots, z_s)$ is the so-called *squared shift-averaged worst case error*. See [19] for the analysis of an efficient implementation of the algorithm for POD weights (3.16), so that the cost is $\mathcal{O}(sn \log n + s^2n)$ operations using FFT. We refer to the accompanying software of [16] for an implementation.

3.4 QMC convergence in the case of circulant embedding

The circulant embedding technique is a method of computing efficiently the factorisation (1.7), thus yielding a method of sampling the random vector \mathbf{Z} via (1.6). We describe the process briefly here before verifying the assumptions of Theorem 5. This section is a summary of our results in [13].

The $M = (m_0 + 1)^d$ points $\{x_i : i = 1, \dots, M\}$ are assumed to be uniformly spaced with spacing $h_0 := 1/m_0$ on a d -dimensional grid over the unit cube $[0, 1]^d$ enclosing the domain D . Using a vector notation, we may relabel the points x_1, \dots, x_M to be indexed by \mathbf{k} as

$$x_{\mathbf{k}} := h_0 \mathbf{k} \text{ for } \mathbf{k} = (k_1, \dots, k_d) \in \{0, \dots, m_0\}^d.$$

Then it is easy to see that (with analogous vector notation for the rows and columns) the $M \times M$ covariance matrix R defined in (1.5) can be written as

$$R_{\mathbf{k}, \mathbf{k}'} = \rho(h_0(\mathbf{k} - \mathbf{k}')), \quad \mathbf{k}, \mathbf{k}' \in \{0, \dots, m_0\}^d. \tag{3.19}$$

If the vectors \mathbf{k} are enumerated in lexicographical ordering, then we obtain a nested block Toeplitz matrix where the number of nested levels is the physical dimension d .

We extend R to a nested block circulant matrix R^{ext} . To do this, it is convenient to extend to the infinite grid:

$$x_{\mathbf{k}} := h_0 \mathbf{k} \text{ for } \mathbf{k} \in \mathbb{Z}^d.$$

Then, to define R^{ext} , we consider an enlarged cube $[0, \ell]^d$ of edge length $\ell := mh_0 \geq 1$ with integer $m \geq m_0$. We assume that m_0 (and hence h_0) is fixed and we enlarge m (or equivalently ℓ) as appropriate. We introduce a 2ℓ -periodic map on \mathbb{R} by specifying its action on $[0, 2\ell]$:

$$\varphi(x) := \begin{cases} x & \text{if } 0 \leq x \leq \ell, \\ 2\ell - x & \text{if } \ell \leq x < 2\ell. \end{cases}$$

Now we apply this map elementwise and define an extended version ρ^{ext} of ρ as follows:

$$\rho^{\text{ext}}(\mathbf{x}) := \rho(\varphi(x_1), \dots, \varphi(x_d)), \quad \mathbf{x} \in \mathbb{R}^d.$$

Note that ρ^{ext} is 2ℓ -periodic in each coordinate direction and $\rho^{\text{ext}}(\mathbf{x}) = \rho(\mathbf{x})$ when $\mathbf{x} \in [0, \ell]^d$. Then R^{ext} is defined to be the $s \times s$ symmetric nested block circulant matrix with $s = (2m)^d$, defined, analogously to (3.19), by

$$R_{\mathbf{k}, \mathbf{k}'}^{\text{ext}} = \rho^{\text{ext}}(h_0(\mathbf{k} - \mathbf{k}')), \quad \mathbf{k}, \mathbf{k}' \in \{0, \dots, 2m - 1\}^d. \tag{3.20}$$

It follows that R is the submatrix of R^{ext} in which the indices are constrained to lie in the range $\mathbf{k}, \mathbf{k}' \in \{0, \dots, m_0\}^d$. Since R^{ext} is nested block circulant, it is diagonalisable by FFT. The following theorem is taken from [12]:

Theorem 6 R^{ext} has the spectral decomposition:

$$R^{\text{ext}} = Q^{\text{ext}} \Lambda^{\text{ext}} Q^{\text{ext}},$$

where Λ^{ext} is the diagonal matrix containing the eigenvalues of R^{ext} , which can be obtained by \sqrt{s} times the Fourier transform on the first column of R^{ext} , and $Q^{\text{ext}} = \Re(\mathcal{F}) + \Im(\mathcal{F})$ is real symmetric, with

$$\mathcal{F}_{\mathbf{k}, \mathbf{k}'} = \frac{1}{\sqrt{s}} \exp\left(2\pi i \frac{\mathbf{k}' \cdot \mathbf{k}}{2m}\right)$$

denoting the d -dimensional Fourier matrix. If the eigenvalues of R^{ext} are all non-negative then the required B in (1.7) can be obtained by selecting M appropriate rows of

$$B^{\text{ext}} := Q^{\text{ext}} (\Lambda^{\text{ext}})^{1/2}.$$

The use of FFT allows fast computation of the matrix-vector product $B^{\text{ext}} \mathbf{y}$ for any vector \mathbf{y} , which then yields $B \mathbf{y}$ needed for sampling the random field in (1.6). Our algorithm from [13] for obtaining a minimal positive definite R^{ext} is given in Algorithm 1. Our algorithm from [13] for sampling an instance of the lognormal random field is given in Algorithm 2. Note that the normalisation used within the FFT routine differs among particular implementations. Here, we assume the Fourier transform to be unitary.

Algorithm 1 Input: d, m_0 , and covariance function ρ .

1. Set $m = m_0$.
2. Calculate \mathbf{r} , the first column of R^{ext} in (3.20).
3. Calculate \mathbf{v} , the vector of eigenvalues of R^{ext} , by d -dimensional FFT on \mathbf{r} .
4. If the smallest eigenvalue < 0 then increment m and go to Step 2.

Output: m, \mathbf{v} .

Algorithm 2 Input: d, m_0 , mean field \bar{Z} , and m and \mathbf{v} obtained by Algorithm 1.

1. With $s = (2m)^d$, sample an s -dimensional normal random vector \mathbf{y} .
2. Update \mathbf{y} by elementwise multiplication with $\sqrt{\mathbf{v}}$.
3. Set \mathbf{w} to be the d -dimensional FFT of \mathbf{y} .

4. Update \mathbf{w} by adding its real and imaginary parts.
5. Obtain \mathbf{z} by extracting the appropriate $M = (m_0 + 1)^d$ entries of \mathbf{w} .
6. Update \mathbf{z} by adding $\overline{\mathbf{z}}$.

Output: $\exp(\mathbf{z})$.

In the case of QMC sampling, the random sample \mathbf{y} in Step 1 of Algorithm 2 is replaced by a randomly shifted lattice point from $[0, 1]^s$, mapped to \mathbb{R}^s elementwise by the inverse of the cumulative normal distribution function (see (1.12)). The relative size of the quantities $b_j = \|\mathbf{B}_j\|_\infty$ (as defined in (3.6)) determines the ordering of the QMC variables in order to benefit from the good properties of lattice rules in earlier coordinate directions in the construction of the generating vector in Remark 1.

We prove in [13] (under mild conditions) that Algorithm 1 will always terminate. Moreover, in many cases the required m (equivalently ℓ) can be quite small. Theorem 7 below gives an explicit lower bound for the required value of ℓ .

Example 1 The Matérn family of covariances are defined by

$$\rho(\mathbf{x}) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{\lambda} \|\mathbf{x}\|_2 \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\lambda} \|\mathbf{x}\|_2 \right), \tag{3.21}$$

where Γ is the Gamma function and K_ν is the modified Bessel function of the second kind, σ^2 is the variance, λ is the correlation length and $\nu \geq 1/2$ is a smoothness parameter. The limiting cases $\nu \rightarrow 1/2$ and $\nu \rightarrow \infty$ correspond to the exponential and Gaussian covariances respectively, see, e.g., [11], however, using a slightly different scaling.

The following result, proved in [13, Thm. 2.10], shows that the growth of the size of ℓ with respect to the mesh size h_0 and with respect to the parameters in the Matérn family is moderate. In particular, for fixed $\nu < \infty$, it establishes a bound on ℓ that grows only logarithmically with λ/h_0 and gets smaller as λ decreases. Experiments illustrating the sharpness of this bound are given in [13].

Theorem 7 *Consider the Matérn covariance family (3.21) with $1/2 \leq \nu < \infty$ and $\lambda \leq 1$. Suppose $h_0/\lambda \leq e^{-1}$. Then there exist constants $C_1 > 0$ and $C_2 \geq 2\sqrt{2}$ which may depend on s but are independent of ℓ, h_0, λ, ν and σ^2 , such that R^{ext} is positive definite if*

$$\ell/\lambda \geq C_1 + C_2 \nu^{1/2} \log \left(\max\{\lambda/h_0, \nu^{1/2}\} \right).$$

In the case $\nu = \infty$, the bound on ℓ is of the form $\ell \geq 1 + \lambda \max\{\sqrt{2}\lambda/h_0, C_1\}$.

In order to verify the QMC convergence estimate given in Theorem 5 in the case of circulant embedding, we need to bound $\|\mathbf{b}\|_{p,s}$, where \mathbf{b} is defined in (3.6). Since every entry in $\Re(\mathcal{F}) + \Im(\mathcal{F})$ is bounded by $\sqrt{2/s}$, we have

$$b_j = \|\mathbf{B}_j\|_\infty \leq \sqrt{\frac{2}{s}} A_{s,j}^{\text{ext}}, \tag{3.22}$$

where $\Lambda_{s,j}^{\text{ext}}, j = 1, \dots, s$, are the eigenvalues of the nested block circulant matrix R^{ext} . Notice that we added ‘ s ’ explicitly to the notation to stress the dependence of these eigenvalues on s . A sufficient condition to ensure the uniform boundedness of $\|\mathbf{b}\|_{s,p}$ required in Theorem 5 is that there exists a constant $C > 0$, independent of s , such that

$$\sum_{j=1}^s \left(\frac{\Lambda_{s,j}^{\text{ext}}}{s} \right)^{p/2} \leq C. \tag{3.23}$$

It is thus important to investigate for what values of p this inequality holds. The smaller the value of p the faster the convergence will be in Theorem 5.

In [13, §3], we conjecture (with supporting mathematical arguments and empirical evidence) that the eigenvalues $\Lambda_{s,j}^{\text{ext}}$, when rearranged in non-increasing order, decay like $j^{-(1+2\nu/d)}$ in case of the Matérn covariance. This is the same as the decay rates of both the eigenvalues of the original nested block Toeplitz matrix R and of the KL eigenvalues of the underlying continuous field Z . Under this conjecture, it follows that the smallest value of p allowed for (3.23) to hold is just bigger than $2/(1 + 2\nu/d)$. In turn this yields a theoretical convergence rate of nearly

$$\mathcal{O}(n^{-\min(\nu/d,1)})$$

in Theorem 5 above, for any $\nu > d/2$, independently of s . To see this, recall that the convergence rate of $-1/(2\kappa)$ with respect to n in Theorem 5 is related to p via $p = 2\kappa/(1 + \kappa)$ with $\kappa \in (1/2, 1)$. These bounds on κ imply that, for the conjectured rate of decay of the eigenvalues $\Lambda_{s,j}^{\text{ext}}$, Theorem 5 is only applicable for $\nu > d/2$.

These conjectures will be investigated in detail in our numerical experiments in the next section. As we will see there, the theoretically predicted rates may be pessimistic. In the experiments here, we see nearly optimal QMC convergence, i.e., $\mathcal{O}(n^{-1})$, even when $\nu < d$, and at least as good convergence as for standard MC, i.e., $\mathcal{O}(n^{-1/2})$, even when $\nu < d/2$. All these findings are in line with the results we obtained in the case of KL expansions in [11], and they guarantee a dimension-independent optimal QMC convergence for sufficiently large smoothness parameter ν .

4 Numerical experiments

In this section we perform numerical experiments on problem (1.1) in 2D and 3D which illustrate the power of the proposed algorithm. Our quantity of interest will be the average value of the solution u

$$\mathcal{G}(u(\cdot, \mathbf{y})) = \frac{1}{|T|} \int_T u(\mathbf{x}, \mathbf{y}) \, d\mathbf{x}, \tag{4.1}$$

over some measurable $T \subseteq D$, with D being an L -shaped domain with a hole in 2D or the unit cube in 3D; all details to be specified below. In both cases, the domain D is contained in the unit cube $[0, 1]^d$, as assumed.

Table 1 The values of s needed to ensure positive definiteness for different combinations of parameters in 2D and 3D

$d = 2$	$\lambda = 0.2$			$\lambda = 0.5$		
	$\nu = 0.5$	$\nu = 2$	$\nu = 4$	$\nu = 0.5$	$\nu = 2$	$\nu = 4$
$m_0 = 12$ (nbe=424)	576 (0.07)	576 (0.03)	576 (0.04)	1,296 (0.11)	5,476 (0.22)	9,216 (0.15)
$m_0 = 24$ (nbe=1,255)	2,304 (0.04)	2,916 (0.04)	4,900 (0.05)	8,464 (0.09)	34,596 (0.26)	59,536 (0.39)
$m_0 = 48$ (nbe=5,559)	9,216 (0.03)	19,044 (0.05)	33,124 (0.07)	49,284 (0.15)	198,916 (0.32)	350,464 (0.44)
$m_0 = 96$ (nbe=23,202)	36,864 (0.02)	114,244 (0.06)	200,704 (0.11)	270,400 (0.14)	1,077,444 (0.40)	1,721,344 (0.49)

$d = 3$	$\lambda = 0.2$			$\lambda = 0.5$		
	$\nu = 0.5$	$\nu = 3$	$\nu = 4$	$\nu = 0.5$	$\nu = 3$	$\nu = 4$
$m_0 = 7$ (nbe=2,642)	2,744 (0.002)	2,744 (0.001)	2,744 (0.001)	64,000 (0.01)	97,336 (0.02)	125,000 (0.03)
$m_0 = 14$ (nbe=21,491)	27,000 (0.001)	39,304 (0.002)	39,304 (0.002)	1,061,208 (0.04)	2,000,376 (0.06)	2,406,104 (0.10)
$m_0 = 28$ (nbe=172,421)	438,976 (0.002)	778,688 (0.003)	941,192 (0.004)	15,625,000 (0.06)	30,371,328 (0.12)	37,933,056 (0.14)

The numbers in round brackets show the cost of random field generation as a fraction of the total computational cost per sample. These numbers increase with increasing s (from left to right) for a fixed FE mesh. For reference, we also provide the number of elements (nbe) for each of the FE meshes

Random field generation. In all experiments the random coefficient a is of the form (1.2) where Z is a Gaussian random field with the Matérn covariance (3.21). We take the mean \bar{Z} to be 0, the variance to be $\sigma^2 = 0.25$, and we consider two different values for the correlation length, namely $\lambda \in \{0.2, 0.5\}$, combined with three different values for the smoothness parameter $\nu \in \{0.5, 2, 4\}$ in 2D and $\nu \in \{0.5, 3, 4\}$ in 3D (thus illustrating the cases $\nu < d, \nu = d, \nu > d$ in each case). The forcing term is taken to be $f \equiv 1$.

For different values of m_0 , we first obtain values of the random field on a uniform grid with $(m_0 + 1)^d$ points on the unit cube $[0, 1]^d$ by circulant embedding as described in Sect. 3.4. We choose $m_0 \in \{12, 24, 48, 96\}$ in 2D and $m_0 \in \{7, 14, 28\}$ in 3D. The necessary length $\ell = m/m_0$ of the extended cube $[0, \ell]^d$ to ensure positive definiteness, where $m \geq m_0$, depends on the values of d, λ and ν , and affects the dimensionality $s = (2m)^d$ of \mathbf{y} . This dependence is investigated in detail in [13] (see Theorem 7). In Table 1, we summarise the values of s for the different combinations of parameters.

FE solution of the PDE. For each realisation of the random field (i.e., for each \mathbf{y}), we solve the PDE using a finite element (FE) method with piecewise linear elements on an unstructured grid produced with the help of the MATLAB PDE toolbox. The quadrature rule for the matrix assembly is based on the mid point rule, where the values of the random field at the centroids of the elements are obtained by multi-linear interpolation of the values on the uniform grid, computed with circulant embedding (see (2.6) and (2.7)). In order to balance the quadrature error and the FE discretisation error in light of Lemma 1 and Theorem 1, the maximum FE mesh diameter is chosen such that $h \approx \sqrt{d} h_0 = \sqrt{d}/m_0$. In particular, we choose $h \in \{0.12, 0.06, 0.03, 0.015\}$ in 2D and $h \in \{0.24, 0.12, 0.06\}$ in 3D, for each of the respective values of m_0 above. In 2D, the MATLAB function `adaptmesh` is used to build a family of adaptive meshes for the




Fig. 1 Adaptive FE mesh of an L -shaped domain with a hole. Left: $h = 0.12$ (424 elements). Middle: $h = 0.06$ (1255 elements). Right: $h = 0.03$ (5559 elements). Our fourth FE mesh not shown here: $h = 0.015$ (23,202 elements)




Fig. 2 A local view of the meshes from Fig. 1, showing the quadrature points at the centroids of the triangles (blue dots) and the uniform grid points where the random field is sampled (purple crosses). Left: $(m_0, h) = (12, 0.12)$. Middle: $(m_0, h) = (24, 0.06)$. Right: $(m_0, h) = (48, 0.03)$ (color figure online)

L -shaped domain with a hole (see Fig. 1). We use the same adaptive mesh, constructed with $a \equiv 1$, for all realisations. To find meshes with our desired maximum mesh diameters h , we gradually increase the `maxt` parameter of the MATLAB `adaptmesh` command. Figure 2 zooms in on the shaded region in the bottom left corner of each of the adaptive meshes to show the centroids of the triangles in relation to the uniform grids. The PDE is solved with the MATLAB function `assempde`. For the 3D problem, we use the MATLAB PDE toolbox to mesh and solve the PDE. The integral in (4.1) is approximated by applying the midpoint rule on each of the elements in T . In 2D, the resulting linear system is solved with the default sparse direct solver (“backslash”) in MATLAB. We believe that that is also the solver used in the MATLAB PDE toolbox for our 3D experiments, but we could not verify this.

As we can see from the fraction of time needed to construct the random field, which is shown in brackets in Table 1 for each case, the majority of time is spent on assembling and solving the FE systems. As expected this is even more pronounced in 3D, since sparse direct solvers are known to be significantly more expensive for 3D FE matrices with respect to the number of unknowns (both in terms of the order and in terms of the asymptotic constant), while the cost of the FFT factorisation grows log-linearly with the number of unknowns in 2D and in 3D. In any case, the random field generation is insignificant in the majority of cases and it takes less than 50% of the computational time in all cases.

Construction of lattice sequences. We approximate the expected value of (4.1) by randomly shifted lattice rules obtained using the fast CBC code from the QMC4PDE website <https://people.cs.kuleuven.be/~dirk.nuyens/qmc4pde/> (see also [16]). A typical call to the QMC4PDE construction script would be:

```
./lat-cbc.py -s=2000 -b_file=[f] -m=16 -d1=1 -d2=2 -a3=1
               -outdir=[o]
```

where $s=2000$ specifies an initial maximum number of dimensions and $[f]$ is a file containing the calculated values of b_j in (3.6) in nonincreasing order for a particular case of d , λ and ν . See [16] for an explanation of the other parameters.

In specifying parameters for the CBC construction, we follow the theory of [16] as closely as possible, but we make a couple of modifications for practical reasons. Firstly, the implementation follows [8] to construct “embedded” lattice sequences in base 2, so that in practice we can increase n gradually without throwing away existing function evaluations. At every power of 2, the points form a full lattice rule which complies with the theory from Sect. 3. Secondly, with the POD weights in (3.16) the CBC construction to find the generating vector \mathbf{z} has a cost of $\mathcal{O}(s^2n + sn \log n)$ operations which becomes quite high for the large s we are considering (see Table 1). Thus, we only carry out the CBC construction up to a certain dimension s^* and then randomly generate the remaining components of \mathbf{z} . In particular, we stop the CBC algorithm at the first component s^* where the generating vector has a repeated entry. Repeated components in \mathbf{z} yield bad two-dimensional projections of lattice points and randomly generated components are intuitively better in that situation. The highest dimensionality for the switch-over dimension for all cases in Table 1 is $s^* = 1811$. The only two cases where we did not need to add random components were $\nu = 2$ and $\nu = 4$, for $d = 2$, $m_0 = 12$ and $\lambda = 0.2$.

Estimation of (Q)MC error. For fixed h , we can compute the standard error on the QMC estimate of the expected value of (4.1) by using a number of random shifts. Specifically, for each case we took $q = 64$ independent random shifts of one n -point lattice rule, giving q independent approximations Q_1, \dots, Q_q to the expected value. We take their average $\bar{Q} = (Q_1 + \dots + Q_q)/q$ as our final approximation, and we estimate the standard error on \bar{Q} by $\sqrt{\frac{1}{q(q-1)} \sum_{i=1}^q (Q_i - \bar{Q})^2}$. The total number of function evaluations in this case is $N = qn$. According to our theory (see also [16]), the convergence rate for our randomised QMC method is of the order $q^{-1/2} n^{-r} = q^{r-1/2} N^{-r}$, with $r \approx \min(\nu/d, 1)$. Hence, for $r > 1/2$, the constant in any of the convergence graphs with respect to N depends on $q^{r-1/2}$. To provide less erratic curves, the number of random shifts is chosen to be fairly large here. In practice, e.g., $q = 16$ shifts would be sufficient. This would effectively push all convergence graphs down, leading to bigger gains for QMC.

We compare QMC with the simple Monte Carlo (MC) method (1.14) based on N random samples. Denoting the function values for these samples by Y_1, \dots, Y_N , then the MC approximation of the integral is $\bar{Y} = (Y_1 + \dots + Y_N)/N$. The standard error can be estimated by $\sqrt{\frac{1}{N(N-1)} \sum_{i=1}^N (Y_i - \bar{Y})^2}$. The expected MC convergence rate is $\mathcal{O}(N^{-1/2})$.

In our figures later, we plot the *relative standard error* obtained by dividing the estimated standard error by the estimated mean.

Computing environment. All our computations were run as serial jobs on reserved 8-core Intel Xeon E5-2650v2 2.60 GHz nodes on the computational cluster Katana at UNSW (or on almost identical hardware). Since they are embarrassingly parallel, both the MC and QMC simulations could easily be parallelised with roughly linear speedup. We chose to run different jobs in parallel instead of parallelising individual jobs, and to report the actual serial computation times for our test cases.

4.1 Results for an L -shaped domain in 2D

In this example the domain D is the complex 2D domain shown in Fig. 1: an L -shaped region with a hole. We consider five choices for the averaging domain $T \subseteq D$ in (4.1):

- T1** the full domain,
- T2** the bottom left corner with a circular segment cut out,
- T3** the lower right interior square, and
- T4** the upper left interior square in a symmetrical location to **T3**,
- T5** the L -shape near the reentrant corner.

Figure 1 shows the different averaging domains T , as well as some of the adaptive meshes that were used. Note that the circular sections of the boundary of D are approximated polygonally and the averaging domains **T2**, ..., **T5** are resolved on all meshes. The meshes are adapted to capture the loss of regularity near the reentrant corner, but nevertheless the number of elements grows roughly with $\mathcal{O}(h^{-2})$, the same as for a uniform family of meshes with mesh size h . We specify the domain in MATLAB by means of constructive solid geometry (CSG), i.e., the union and subtraction of the basic pieces. This ensures that all the averaging domains T are covered by complete elements.

Mesh errors. Before we compare the performance of our QMC method with the basic MC method, let us first estimate the discretisation errors for each of the adaptive meshes. In Table 2, we present results for **T1** and **T5** for the case $\lambda = 0.2$ and $\nu = 2$. The estimates of $\mathbb{E}[\mathcal{G}(u_h)]$, obtained using QMC, are stated together with the estimated standard error for each mesh. We use sufficiently many QMC cubature points, so that the significant figures of the estimates are not affected by QMC errors. The product of m_0 and h is kept fixed (approximately equal to \sqrt{d}), as discussed above. From the results we can clearly see that the convergence rate for the discretisation error is $\mathcal{O}(h^2)$ on the given meshes. This shows that the mesh refinement near the reentrant corner is working optimally. Using Richardson extrapolation, we can thus compute a higher order approximation of the limit of $\mathbb{E}[\mathcal{G}(u_h)]$ that is stated in the last row of Table 2. The relative error with respect to this extrapolated value is stated in the last column for both **T1** and **T5**.

The behaviour is similar for the other quantities of interest and for the other values of λ when $\nu = 2$. For $\nu = 0.5$, on the other hand, the solution is globally only in $H^{3/2}$, so that the local mesh refinement near the reentrant corner plays no role and the convergence of $\mathbb{E}[\mathcal{G}(u_h)]$ is only $\mathcal{O}(h)$. Thus, to achieve acceptable accuracy, comparable to the QMC errors we quote below, we would in practice need much finer

Table 2 Mesh convergence for the case $\lambda = 0.2$ and $\nu = 2$ for **T1** and **T5**

(m_0, h)	T1		T5	
	$\mathbb{E}[\mathcal{G}(u_h)]$	rel. h -error	$\mathbb{E}[\mathcal{G}(u_h)]$	rel. h -error
(12, 0.12)	$0.0114378 \pm 3.4e-08$	$4.6e-02$	$0.0108212 \pm 7.9e-08$	$7.9e-02$
(24, 0.06)	$0.0118095 \pm 2.8e-08$	$1.5e-02$	$0.0114086 \pm 6.6e-08$	$2.9e-02$
(48, 0.03)	$0.0119549 \pm 4.1e-08$	$3.3e-03$	$0.0116850 \pm 8.1e-08$	$5.9e-03$
(96, 0.015)	$0.0119846 \pm 4.9e-08$	$8.3e-04$	$0.0117372 \pm 1.1e-07$	$1.5e-03$
	0.0119945 (extrapolated)	$\sim h^2$	0.0117546 (extrapolated)	$\sim h^2$

The estimates of $\mathbb{E}[\mathcal{G}(u_h)]$ (stated together with one standard deviation) are computed with our randomised lattice rule with n sufficiently large, such that the standard error is significantly smaller than the discretisation error. Richardson extrapolation is used to compute a more accurate estimate of the limit of $\mathbb{E}[\mathcal{G}(u_h)]$, as $h \rightarrow 0$ (final row). The columns denoted “rel. h -error” give the relative error with respect to these extrapolated estimates

meshes for $\nu = 0.5$. However, since this would not affect the behaviour of the QMC cubature errors, we did not do that.

QMC convergence rates. For the remainder we will now demonstrate the dimension independence of our QMC method and its superiority over basic MC. In Fig. 3, we plot the relative standard error in (4.1) with $T = \mathbf{T1}$ against the total number of PDE solves (top) and against the total calculation time (bottom). We consider six combinations of the values of ν and λ and plot the graphs for four levels of mesh refinement. (The MC estimates were not computed on the finest mesh.) The convergence of the MC method is proportional to $N^{-0.5}$, as expected. The convergence of the QMC method ranges from $\mathcal{O}(N^{-0.72})$ up to $\mathcal{O}(N^{-0.89})$. For example, for $\nu = 0.5$ and $\lambda = 0.2$, to achieve the same relative standard error of 10^{-4} we need about 10^6 PDE solves with the MC method while the QMC method only needs about $3 \cdot 10^4$ PDE solves. Also in terms of computational time, all the results consistently show huge computational savings for the QMC method over the MC method, even with the relatively large number of $q = 64$ random shifts.

We note that the convergence graphs are meant to illustrate the convergence behaviour and in practice one would not try to achieve such high precision, especially not on the coarser meshes. One would rather aim to balance the QMC errors with the discretisation errors in Table 2. From the theory, we expect the smoothness ν of the random field to have an effect on the convergence rate of QMC. Specifically, the bound in Theorem 5 would suggest a rate of $\mathcal{O}(N^{-\min(\nu/2, 1)})$ in 2D, as discussed in Sect. 3.4. This effect is not immediately observed in the graphs. For $\nu = 2$ and 4 we expect an asymptotic convergence rate of order 1, while we see very good convergence rates on the graphs, we did not reach this asymptotic regime yet. On the other hand, we do observe excellent convergence behaviour for the case $\nu = 0.5$ for which our theory does not apply. From the graphs we also observe that a smaller correlation length λ corresponds to a better convergence rate, which is in full agreement with our findings in [12].

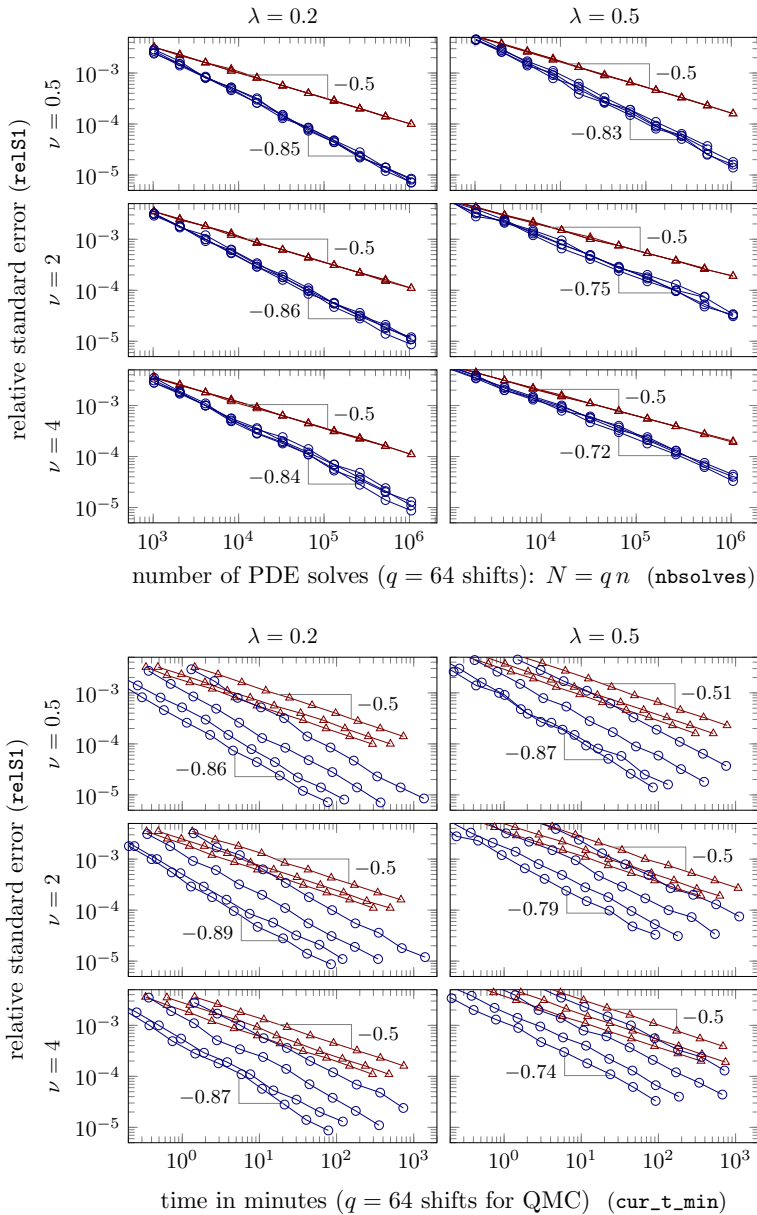


Fig. 3 Relative standard error for $T = \mathbf{T1}$ (the average solution over the entire domain) against total number of PDE solves (top) and execution time for Monte Carlo (red triangles) and QMC (blue circles). In the timing plot, results appear from bottom to top from the coarsest mesh with $(m_0, h) = (12, 0.12)$ to the finest mesh with $(m_0, h) = (96, 0.015)$ for QMC and $(m_0, h) = (48, 0.03)$ for MC, respectively (color figure online)

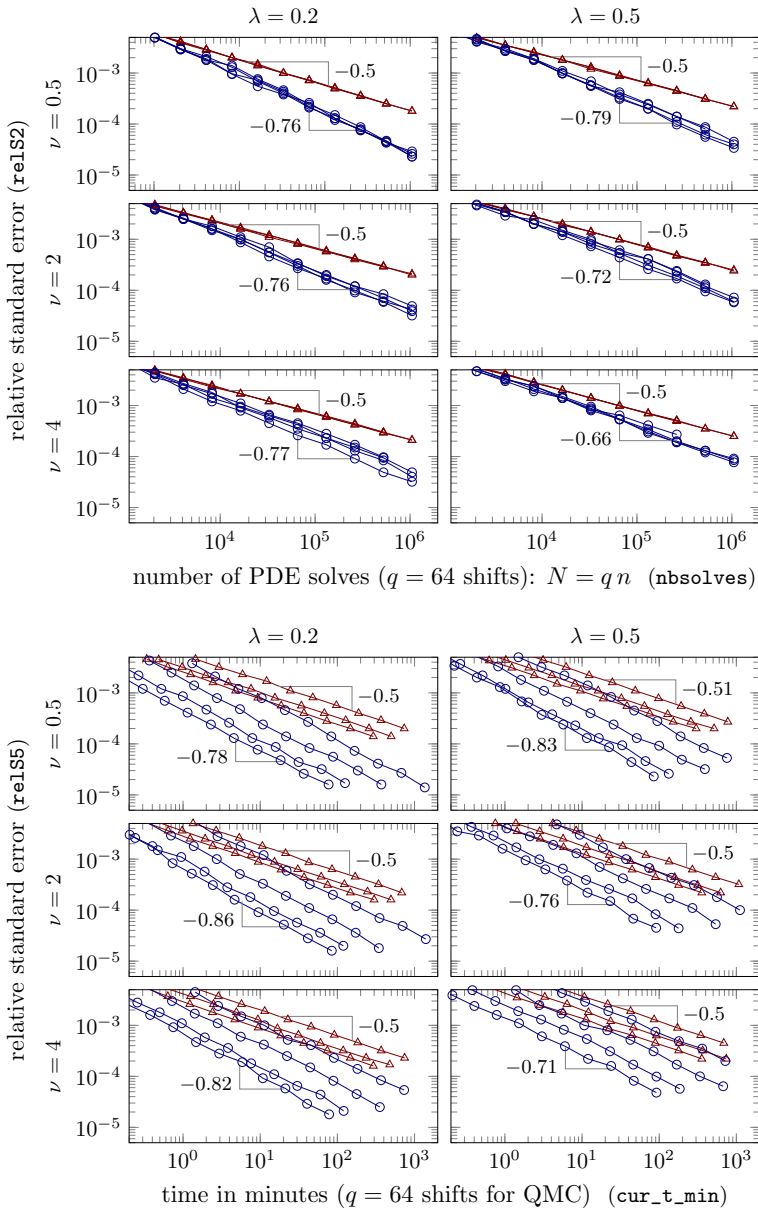


Fig. 4 Relative standard error for Monte Carlo (red triangles) and the lattice sequences (blue circles) for **T2** (the average over the bottom-left corner with circular cutout) versus number of PDE solves (top), as well as for **T5** (the average over the L-shaped region near the reentrant corner) versus execution time (bottom). In the bottom figure, the results appear from bottom to top from the coarsest mesh to the finest (color figure online)

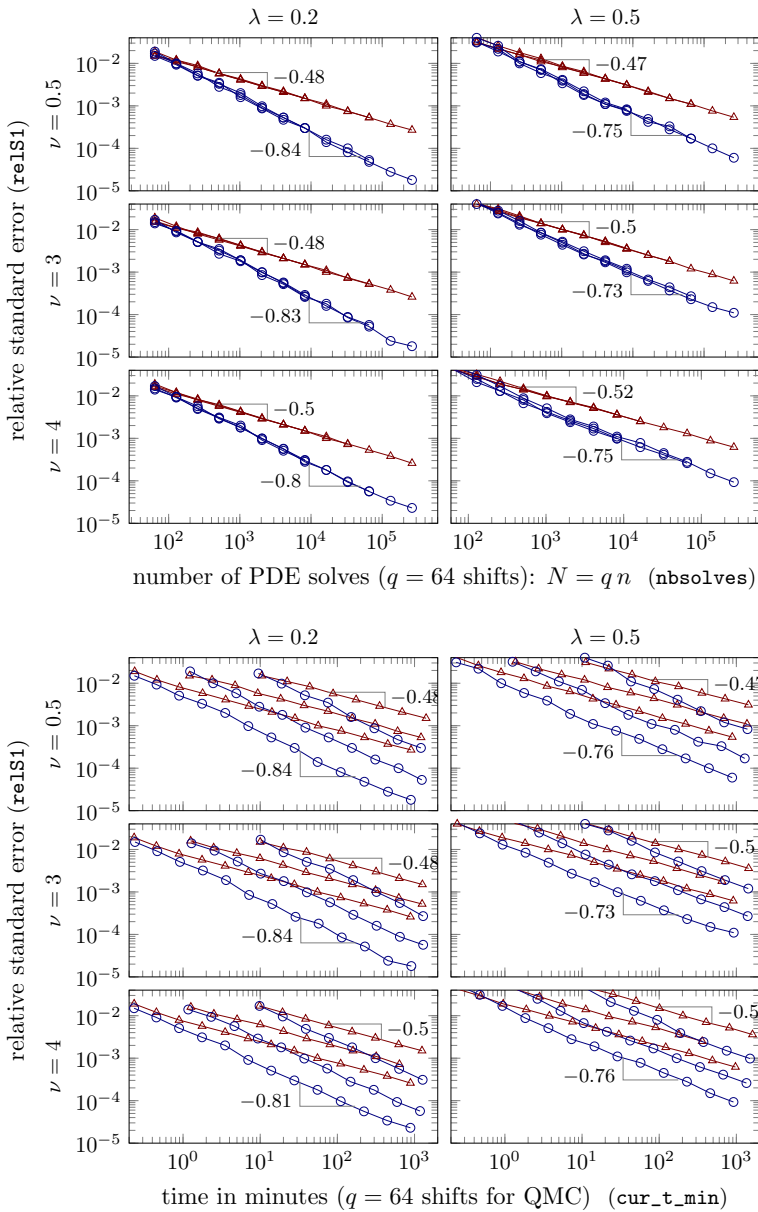


Fig. 5 Relative standard error of the average solution over the unit cube in 3D against the total number of PDE solves (top) and the execution time (bottom) for Monte Carlo (red triangles) and the lattice sequences (blue circles). As in Fig. 3, the results appear from bottom to top from the coarsest mesh with $(m_0, h) = (7, 0.24)$ to the finest mesh with $(m_0, h) = (28, 0.06)$ (color figure online)

A more important observation is the overlay of the convergence lines for the different meshes in the plots of relative standard error versus number of PDE solves in Fig. 3. For example, for the case $\lambda = 0.5$ and $\nu = 4$ the dimensionality s increases from about 9 thousand to 1 million as we increase m_0 from 12 to 96 (see Table 1), while the convergence rate and the asymptotic constant for the relative standard error are clearly independent of the increasing dimension.

In Fig. 4, we confirm that the superiority is independent of the quantity of interest, by presenting similar graphs as for **T1** also for **T2** and **T5**. We do not include the results for the two symmetrical squares **T3** and **T4**, which look very similar.

4.2 Results for the unit cube in 3D

Our second example is the unit cube in 3D and our quantity of interest is (4.1) with $T = D = [0, 1]^3$. We use the mesh generator and the solver from the MATLAB PDE toolbox to obtain three meshes with desired maximum mesh diameters h and to calculate the solution. To evaluate the random field Z at the centroids of the elements we used the function `interpolateSolution` from the toolbox.

In Fig. 5 we plot again the relative standard error for six combinations of parameters against the total number of PDE solves and against the calculation time. We observe convergence rates from $N^{-0.73}$ up to $N^{-0.84}$ for our QMC rules, and the expected $N^{-0.5}$ for the plain MC method. The lattice sequences were again constructed with the actual values of b_j from (3.6), except for the combination of $\lambda = 0.5$ and $m_0 = 28$ where we get the near astronomical dimensionalities ranging from about 15 to 37 million, as ν increases (see Table 1); for these cases we replaced b_j by the convenient upper bound given in (3.22).

Acknowledgements We gratefully acknowledge the financial support from the Australian Research Council (FT130100655 and DP150101770), the KU Leuven research fund (OT:3E130287 and C3:3E150478), the Taiwanese National Center for Theoretical Sciences (NCTS)—Mathematics Division, and the Statistical and Applied Mathematical Sciences Institute (SAMSI) 2017 Year-long Program on Quasi-Monte Carlo and High-Dimensional Sampling Methods for Applied Mathematics.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Adler, R.J.: The Geometry of Random Fields. Wiley, London (1981)
2. Apel, T.: Anisotropic Finite Elements: Local Estimates and Applications. Advances in Numerical Mathematics. Teubner, Stuttgart (1999)
3. Bacuta, C., Nistor, V., Zikatanov, L.: Improving the rate of convergence of high-order finite elements on polyhedra I: a priori estimates. Numer. Funct. Anal. Optim. **26**, 613–639 (2006)
4. Blanchard, P., Coulaud, O., Darve, E.: Fast hierarchical algorithms for generating Gaussian random fields. Preprint <https://hal.inria.fr/hal-01228519>, 18 Feb 2016
5. Chan, G., Wood, A.T.A.: Simulation of stationary Gaussian processes in $[0, 1]^d$. J. Comput. Graph. Stat. **3**, 409–432 (1994)

6. Chan, G., Wood, A.T.A.: Algorithm AS 312: an algorithm for simulating stationary Gaussian random fields. *Appl. Stat. J. R. Stat. Soc. Ser. C* **46**, 171–181 (1997)
7. Charrier, J., Scheichl, R., Teckentrup, A.: Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods. *SIAM J. Numer. Anal.* **51**, 322–352 (2013)
8. Cools, R., Kuo, F.Y., Nuyens, D.: Constructing embedded lattice rules for multivariate integration. *SIAM J. Sci. Comput.* **28**, 2162–2188 (2006)
9. Dietrich, C.R., Newsam, G.H.: Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput.* **18**, 1088–1107 (1997)
10. Feischl, M., Kuo, F.Y., Sloan, I.H.: Fast random field generation with H -matrices. *Numer. Math.* (to appear). [arXiv:1702.08637](https://arxiv.org/abs/1702.08637), 28 Feb 2017
11. Graham, I.G., Kuo, F.Y., Nichols, J.A., Scheichl, R., Schwab, Ch., Sloan, I.H.: Quasi-Monte Carlo finite element methods for elliptic PDEs with log-normal random coefficients. *Numer. Math.* **131**, 329–368 (2015)
12. Graham, I.G., Kuo, F.Y., Nuyens, D., Scheichl, R., Sloan, I.H.: Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications. *J. Comput. Phys.* **230**, 3668–3694 (2011)
13. Graham, I.G., Kuo, F.Y., Nuyens, D., Scheichl, R., Sloan, I.H.: Analysis of circulant embedding methods for sampling stationary random fields. *SIAM J. Numer. Anal.* (to appear). [arXiv:1710.00751](https://arxiv.org/abs/1710.00751), 28 Sept 2017
14. Graham, I.G., Scheichl, R., Ullmann, E.: Mixed finite element analysis of lognormal diffusion and multilevel Monte Carlo methods. *Stoch. PDE Anal. Comput.* **4**, 41–75 (2013)
15. Harbrecht, H., Peters, M., Schneider, R.: On the low-rank approximation by the pivoted Cholesky decomposition. *Appl. Numer. Math.* **62**(4), 428–440 (2012)
16. Kuo, F.Y., Nuyens, D.: Application of quasi-Monte Carlo methods to PDEs with random coefficients—analysis and implementation. *Found. Comput. Math.* **16**, 1631–1696 (2016)
17. Lord, G., Powell, C., Shardlow, T.: *An Introduction to Computational Stochastic PDEs*. Cambridge University Press, Cambridge (2014)
18. Naff, R.L., Haley, D.F., Sudicky, E.A.: High-resolution Monte Carlo simulation of flow and conservative transport in heterogeneous porous media I. Methodology and flow results. *Water Resour. Res.* **34**, 663–677 (1998)
19. Nichols, J.A., Kuo, F.Y.: Fast component-by-component construction of randomly shifted lattice rules achieving $\mathcal{O}(n^{-1+\delta})$ convergence for unbounded integrands in \mathbb{R}^s in weighted spaces with POD weights. *J. Complex.* **30**, 444–468 (2014)
20. Schatz, A.H., Wahlbin, L.B.: Maximum norm estimates in the finite element method on plane polygonal domains. Part 2: refinements. *Math. Comput.* **33**, 465–492 (1979)
21. Teckentrup, A.L., Scheichl, R., Giles, M.B., Ullmann, E.: Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficient. *Numer. Math.* **125**, 569–600 (2013)