

EMBEDDED MULTILEVEL MONTE CARLO FOR UNCERTAINTY QUANTIFICATION IN RANDOM DOMAINS

SANTIAGO BADIA^{1,2}, JERRAD HAMPTON², AND JAVIER PRINCIPE^{2,3,*}

ABSTRACT. The multilevel Monte Carlo (MLMC) method has proven to be an effective variance-reduction statistical method for Uncertainty quantification in PDE models. It combines approximations at different levels of accuracy using a hierarchy of meshes in a similar way as multigrid. The generation of body-fitted mesh hierarchies is only possible for simple geometries. On top of that, MLMC for random domains involves the generation of a mesh for every sample. Instead, here we consider the use of embedded methods which make use of simple background meshes of an artificial domain (a bounding-box) for which it is easy to define a mesh hierarchy, thus eliminating the need of body-fitted unstructured meshes, but can produce ill-conditioned discrete problems. To avoid this complication, we consider the recent aggregated finite element method (AgFEM). In particular, we design an embedded MLMC framework for (geometrically and topologically) random domains implicitly defined through a random level-set function, which makes use of a set of hierarchical background meshes and the AgFEM. Performance predictions from existing theory are verified statistically in three numerical experiments, namely the solution of the Poisson equation on a circular domain of random radius, the solution of the Poisson equation on a topologically identical but more complex domain, and the solution of a heat-transfer problem in a domain that has geometric and topological uncertainties. Finally, the use of AgFE is statistically demonstrated to be crucial for complex and uncertain geometries in terms of robustness and computational cost.

Keywords: Partial Differential Equations, Finite Elements, Adaptive Mesh Refinement, Forest of Trees, Parallel algorithms, Scientific Software

1. INTRODUCTION

Uncertainty quantification (UQ) requires the solution of a stochastic partial differential equation (PDE) with random data. Several methods for solving stochastic PDEs, such as stochastic Galerkin [39, 26] or stochastic collocation [4], are based on a standard approximation in space like finite elements (FEs) or finite volumes, and different types of polynomial expansions in the stochastic space [68]. Although very powerful for some particular problems, most of these techniques suffer the so-called ‘curse of dimensionality’, i.e. a dramatic increase of computational cost with the number of stochastic variables. Besides, some of these methods are intrusive because given a code that can be used to solve a deterministic problem, it needs to be modified to solve a stochastic one.

In contrast, the classical Monte Carlo (MC) method does not present these drawbacks. On the one hand it is a sampling method, i.e., it only requires the repeated evaluation of a deterministic model, and it can be implemented in a non-intrusive way. On the other hand it is known to convergence to the exact statistics of the solution as the number of input samples tends to infinity, independently of the dimensionality of the stochastic space and mostly independently of the physics of the problem under consideration, as long as some moments of the quantity of interest (QoI) are bounded. However, the number of samples required to achieve statistical convergence combined

Date: November 28, 2019.

Key words and phrases. Multilevel Monte Carlo; Embedded Methods; Uncertainty Quantification; Topological Uncertainty; Geometric Uncertainty; Stochastic Partial Differential Equations; Random Geometry.

¹ School of Mathematics, Monash University, Clayton, Victoria, 3800, Australia.

² Centre Internacional de Mètodes Numèrics en Enginyeria, Esteve Terrades 5, E-08860 Castelldefels, Spain.

³ Universitat Politècnica de Catalunya, Campus Diagonal Besòs, Av. Eduard Maristany 16, Edifici A (EEBE), 08019, Barcelona, Spain

* Corresponding author.

E-mails: santiago.badia@monash.edu (SB), jhampton@cimne.upc.edu (JH), principe@cimne.upc.edu (JP).

This research was supported by the European Union’s Horizon 2020 research and innovation programme under the ExaQUTE project, with grant agreement No 800898. JH has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 712949 (TECNIOspring PLUS) and from the Agency for Business Competitiveness of the Government of Catalonia.

with the complexity of the computational model required to have enough spatial and/or temporal accuracy can make it prohibitively expensive.

To address this difficulty, the multilevel Monte Carlo (MLMC) method was introduced in [37] (with two levels) and in [27] (with several levels) and applied to solve stochastic PDEs with random input parameters in [28, 44, 43, 55]. The MLMC method exploits a hierarchy of discretizations of the underlying differential problem (using a hierarchy of meshes). The method computes expectations on the difference between solutions at two consecutive levels of accuracy, reducing the variance of the problem at hand and thus decreasing the number of samples required at finer discretization levels. By doing this, most of the computational cost is transferred to coarse discretization levels, where a large number of samples is used to control the statistical error of the MLMC estimator, whereas only few samples are used on the finest (and most costly) levels, to control the discretization error.

Depending on the application, different types of uncertain data can be considered, like material properties (PDE parameters), boundary and/or initial conditions, and the geometry/topology of the domain. The MLMC method can be used out-of-the-box in most cases but it requires a hierarchy of discretizations as input, which can be difficult to generate when dealing with complex geometries and unstructured body-fitted meshes. For this reason MLMC methods are barely used to perform UQ when complex geometries are considered. This is the first limitation we aim to address herein, utilizing embedded methods with MLMC in a method referred to as embedded MLMC (EMLMC).

On top of that, leaving simple academic constructions aside, geometry is hardly known exactly and surfaces are rarely smooth when looking at sufficiently small scale. Variation between different samples of natural materials, e.g. porous rocks, and cost limits of diminishing manufacturing tolerances of fabrication processes make geometry an epistemic uncertainty in general. Geometric uncertainties can have a strong impact on the solution of a PDE, e.g. they trigger boundary layer separation in fluids, and these effects are important in many fields, e.g. geology [23], tribology [36] and lubrication [31], nano-technology [47] and biology [52]. The first attempts to predict the impact of roughness was through deterministic parameterizations, like periodic indentations [65, 56], sinusoidal corrugation [25] and fractal representations, e.g. using the well-known Von Koch snowflake curve [14, 12]. After that, stochastic representations of surface roughness have been adopted frequently, which permit to incorporate the lack of information or measurement errors. Examples of these descriptions include random fields to represent the thickness in shell structures (see [62] and the references therein), random fractal representations [57] and random fields with a given spatial correlation structure to represent derivations from a nominal geometry [69, 64, 46, 70]. In engineering design, geometry is exactly defined using computational tools and in this case the uncertainty is rooted at manufacturing process, as mentioned. Uncertainty also arises when measuring, e.g. when the geometry is determined from digital images, sourced at image scanning and resolution. Using pixelated images for computing may result in significant errors, even in the limit of small resolution [2, 3]. However, it is possible to recover a definition of the geometry using level set methods [59] and, in fact, this can be done statistically, e.g. using polynomial chaos expansions, to define random level set representations [61]. Curiously, level set and phase field representations can also be obtained from noisy images solving stochastic PDEs [54].

There are three approaches to deal with geometric uncertainty. The first one is based on domain mappings, the second on perturbation methods and the third one on the use of fictitious domain methods, also known as embedded or immersed methods. The first approach started with the stochastic mapping concept that was introduced in [69] to transform a (deterministic or stochastic) problem in a random domain to a stochastic problem in a deterministic domain. This approach was applied to the stochastic analysis of roughness in flow problems in [64]. In [46] the main idea is to fix the mesh connectivity and change the coordinates of the nodes using deformation strategies that incorporate uncertainty through the solution of auxiliary PDE with random boundary conditions. In [16] a domain mapping approach is used to prove the convergence of a stochastic collocation method based on Smolyak grids. Because the construction of globally smooth mappings is difficult, a piecewise smooth domain mapping approach is proposed in [18]. The second approach was introduced in [57] where an analytical perturbation method was used to find solutions of the Laplace equation. A perturbation method developed using shape calculus (a Taylor expansion using shape derivatives and its linearization around a nominal domain) was proposed in [34, 33, 32]. Because shape derivatives are difficult to compute, a perturbation approach based on approximated stochastic boundary conditions (of Robin type) in [22, 21]. Although it is not based on a formal perturbation expansion, the stochastic smoothed profile method of [70] can be included in this category. This method transforms a random domain into a random force-term in a deterministic domain through the introduction of a smoothly spreading interface layer to represent rough boundaries and its error grows with the layer thickness [40]. By definition, these approaches are not applicable to complex geometrically and even topologically random domains.

The use of embedded methods, which rely on a fixed background grid and a discrete approximation of the geometry using a level set function, for the solution of PDEs in stochastic domains combined with polynomial chaos was proposed in [15] and applied to fluid dynamics problems in [53]. Among the several possible immersed methods, those in [15, 53] impose boundary conditions through Lagrange multipliers defined in a FE space built from a boundary discretization (independent of the domain discretization). This construction results in a saddle point linear system for which specific preconditioners have been recently developed [29]. Alternatively, the use of the so-called extended finite element method (XFEM) method [45, 63] for dealing with embedded stochastic geometries using polynomial chaos, was proposed in [51, 50, 49, 48] and named extended stochastic finite element method (XS-FEM). It was lately applied to heat transfer [38] and structural problems [58]. The lack of robustness of embedded FE methods is not addressed in these works.

The use of MLMC in complex random domains using body-fitted meshes would require to generate an unstructured mesh (and its corresponding mesh hierarchy) per sample, which is absolutely impractical. For this reason, and the previously described problems related to the generation of mesh hierarchies in general, we also consider embedded FE methods in this work. However, the naive use of methods exhibit some problems. The most salient one is the so-called small cut cell problem. The intersection of cells with the domain surface cannot be controlled and one can end up with cells for which the ratio between the cell volume and its portion in the domain interior is arbitrarily large. It is well-known that the condition number of the resulting linear system blows up with these ratios. As a result, plain embedded FE methods are not robust. Different remedies have been proposed so far in the literature. For conforming FE methods in which trace continuity must be enforced between cells, one can consider methods that introduce artificial dissipation (stabilization) in order to weakly enforce zero jump of derivatives (up to the FE order being used) of jumps on facets that belong to cut cells (see [13] for more details). Another approach is to create aggregated meshes such that every aggregate contains at least one interior cell. Whereas the generation of discontinuous FE spaces on aggregated meshes is straightforward (enforcing the local FE space to be the same polynomial space as in a standard cell), it is far more complicated for conforming FEs. Using such a naive approach for conforming methods would generate global constraints and too much *rigidization*, which would affect the convergence properties of the method, its implementation, and the sparsity of the resulting linear system; we note that this is in fact the *strong* version of the weakly enforced constraints in [13]. The aggregated FE method has been proposed in [9] to enable aggregation techniques for conforming methods. It has three key features: (a) the constraints are cell-local and their implementation in FE codes is easy; (b) it keeps the convergence order of standard body-fitted FE spaces; (c) it does not perturb the Galerkin formulation with any kind of artificial dissipation and does not involve the computation of (possibly) high-order derivatives on the element boundaries.

In this work, we propose the EMLMC framework. It combines the MLMC method on a hierarchy of background meshes (which can be simple structured meshes) and the aggregated finite element method (AgFEM) at every mesh level to capture every sample of the random domain. In particular, we consider random domains implicitly defined through random level-set functions and the marching tetrahedra scheme to approximate the random domain at every mesh level. The resulting method extends the applicability of MLMC to complex random domains and it is robust. We state the model problem in Section 2 and we describe the MLMC method in Section 3. The construction of the AgFEM discretization for each sample is then described in Section 4. Finally, a set of numerical experiments are described in Section 5 and we draw some concluding remarks in Section 6.

2. PROBLEM FORMULATION

As a model problem we consider the following elliptic stochastic problem. Given an oriented manifold $\mathcal{M}(\omega) \subset \mathbb{R}^d$ and its corresponding interior domain $\mathcal{D}(\omega)$, find $u(\mathbf{x}, \omega)$ such that

$$-\nabla \cdot (k \nabla u) = f \text{ in } \mathcal{D}(\omega), \quad u = u_0 \text{ on } \mathcal{M}(\omega), \quad (1)$$

almost surely (a.s.), i.e., for almost all $\omega \in \Omega$, which denotes the uncertainty, described by a complete probability space (Ω, \mathcal{F}, P) . Although stochastic coefficients, i.e. the diffusion $k = k(\mathbf{x}, \omega)$, forcing $f = f(\mathbf{x}, \omega)$, and boundary condition $u_0 = u_0(\mathbf{x}, \omega)$ could be considered random fields too, we assume they are deterministic, as it is usual in the literature when stochastic domains are considered [18, 46, 69, 34, 33, 32, 22, 21]. Let us assume that all realizations $\mathcal{M}(\omega)$ (and $\mathcal{D}(\omega)$) are bounded. We can define a bounded artificial domain \mathcal{B} that contains all possible realizations of $\mathcal{M}(\omega)$, i.e. $\mathcal{D}(\omega) \subset \mathcal{B}$, $\forall \omega \in \Omega$. We also assume that k and f are defined in \mathcal{B} , independently of $\omega \in \Omega$. With the random solution of this problem at hand we aim to compute $\mathbb{E}(Q(u))$ where Q is a deterministic QoI, e.g. an integral of u on a subregion or surface, and \mathbb{E} the expectation.

Although it is not strictly necessary (at least in deterministic domains [17, 66]), uniform ellipticity is often assumed to guarantee the well-posedness of Eq. (1) (see [10, 4, 18]). Therefore, we assume k uniformly bounded from below, i.e. there exists k_0 such that $k(\mathbf{x}) \geq k_0, \forall \mathbf{x} \in \mathcal{B}$ (and $\forall \omega \in \Omega$ if a stochastic diffusion $k(\mathbf{x}, \omega)$ is considered). Under these assumptions, the bilinear form associated to the weak form of (1) is bounded and coercive and the Lax-Milgram lemma guarantees a solution for any $\omega \in \Omega$ uniformly bounded by $\|f\|_{L^2(\mathcal{B})}$ [18].

3. MONTE CARLO AND ITS MULTILEVEL EXTENSION

3.1. The Monte Carlo method. The aim is to compute the expected value $\mathbb{E}(Q)$ of the QoI $Q(u)$. Using a discrete approximation in a mesh \mathcal{T}_h of size h having $M \propto h^{-d}$ degrees of freedom, we actually compute a discrete approximation to the QoI using u_h , the discrete FE solution of the continuous problem at hand. Let us denote $Q_h = Q(u_h)$. The standard MC algorithm computes the approximation of its expected value as the average

$$\mathbb{E}(Q) \approx \bar{Q}_h = \frac{1}{N} \sum_{i=1}^N Q_h^i,$$

where Q_h^i are obtained evaluating the QoI using the N realizations (samples) u_h^i of the stochastic solution $u_h(\mathbf{x}, \omega)$ on the given mesh \mathcal{T}_h . The mean square error of this approximation is

$$e^2(\bar{Q}_h) = \mathbb{E} \left[(\bar{Q}_h - \mathbb{E}(Q))^2 \right] = (\mathbb{E}(Q_h - Q))^2 + \mathbb{V}(\bar{Q}_h)$$

where \mathbb{V} is the variance. As it is well known, the first term is the (squared) discretization or bias error and its reduction requires refining the mesh, whereas the second one is the statistical error that decays as

$$\mathbb{V}(\bar{Q}_h) = N^{-1} \mathbb{V}(Q_h), \quad (2)$$

so its reduction requires increasing the number of samples. Denoting the complexity of evaluating Q_h^i as C_h , it is possible to estimate the total computational cost of the MC algorithm, C_{MC} , under the following assumptions.

Assumption 3.1. *There exist a and c_a such that*

$$|\mathbb{E}(Q_h - Q)| \leq c_a h^a.$$

Assumption 3.2. *There exist γ and c_γ such that*

$$C_h \leq c_\gamma h^{-\gamma}.$$

Register for free at <https://www.scipedia.com> to download the version without the watermark

By **Assumption (3.1)** the mesh size required to achieve a discretization error smaller than ε is $h \leq \varepsilon^{1/a}$ (hereafter $a \leq b$ means that there is a constant c , independent of h , such that $a \leq cb$). On the other hand, by Eq. (2), the number of samples required to achieve a statistical error $\left[\mathbb{V}(\bar{Q}_h) \right]^{1/2} \leq \varepsilon$ is $N \propto \varepsilon^{-2}$, where here $a \propto b$ means that there exist constants c and C such that $c\varepsilon^{-2} \leq N \leq C\varepsilon^{-2}$. Then

$$C_{MC} = NC_h \leq Nh^{-\gamma} \leq \varepsilon^{-(2+\gamma/a)} \quad (3)$$

Note that the factor $\varepsilon^{-\gamma/a}$ comes from the discretization error and represents the (asymptotic) cost of one sample. For the model problem of Section 2, **Assumption (3.1)** holds with $\alpha = 2$ [66]. Now the constant γ depends on d and the type of solver. For naive Gaussian elimination $\gamma = 3d$, whereas $\gamma = d$ for the optimal multigrid method. In between, for sparse linear solvers $\gamma = 3(d-1)$ (for $d = 2, 3$) whereas for the (unpreconditioned) conjugate gradient (CG) algorithm $\gamma = d + 1$. In the last case, for example, the computational cost scales as ε^{-4} when $d = 3$.

3.2. The Multilevel Monte Carlo method. The acMLMC algorithm exploits the linearity of the expectation using a hierarchy of $L + 1$ meshes $\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_L$ of sizes $h_0 > h_1 > \dots > h_L$. Although there are other options [30], we assume $h_l = h_0 s^{-l}$ (i.e. each mesh in the hierarchy is obtained by uniformly dividing each cell into s^d subcells). Performing $\{N_\ell\}_{\ell=0, \dots, L}$ simulations for different values of the random parameters ω on $\{\mathcal{T}_\ell\}_{\ell=0, \dots, L}$, the expectation on the coarse grid is corrected using the whole hierarchy as

$$\mathbb{E}(Q_L) = \mathbb{E}(Q_0) + \sum_{\ell=1}^L \mathbb{E}(Q_\ell - Q_{\ell-1}) \approx \sum_{\ell=0}^L \bar{Y}_\ell := \tilde{Q}_L \quad (4)$$

where $Q_\ell = Q(u_\ell)$ is the approximation of the QoI computed using the FE solution u_ℓ for the mesh \mathcal{T}_ℓ and $Y_\ell = Q_\ell - Q_{\ell-1}$ for $\ell = 1, \dots, L$ and $Y_0 = Q_0$. The total error of this approximation is now given by

$$e^2(\tilde{Q}_L) = \mathbb{E} \left[(\tilde{Q}_L - \mathbb{E}(Q))^2 \right] = (\mathbb{E}(Q_L - Q))^2 + \mathbb{V}(\tilde{Q}_L) \quad (5)$$

The first term, the (squared) discretization error, is the same as in the MC method and it requires mesh \mathcal{T}_L to be fine enough. For this term to be smaller than, e.g. $\varepsilon^2/2$, **Assumption (3.1)** requires $c_\alpha h_L \leq 2^{-1/2} \varepsilon^{1/\alpha}$, or, equivalently,

$$L = \lceil \frac{1}{\alpha} \log_s \left(\sqrt{2} c_\alpha \varepsilon^{-1} \right) \rceil$$

where, as usual, $\lceil x \rceil$ denote the unique integer n satisfying $x < n < x + 1$. The second term Eq. (5) is the statistical error, given by $\mathbb{V}(\tilde{Q}_L) = \sum_{\ell=0}^L N_\ell^{-1} \mathbb{V}(Y_\ell)$, is very different from the one in the MC method thanks to the following assumption.

Assumption 3.3. For any $\ell = 0, \dots, L$, there exist β and c_β such that

$$\mathbb{V}(Y_\ell) \leq c_\beta h_\ell^\beta.$$

Thanks to **Assumption (3.3)**, $\mathbb{V}(Y_\ell)$ tends to zero with h_ℓ and in this sense, the MLMC method can be understood as a variance reduction method. Under **Assumption (3.1)**, **Assumption (3.2)** and **Assumption (3.3)**, it is possible to explicitly bound the total computational cost and to minimize it to obtain the optimal number of samples to be taken on each level [27].

The number of samples on each level required to have a (squared) statistical error smaller than $\varepsilon^2/2$ is given by

$$N_\ell = \lceil 2\varepsilon^{-2} \sqrt{\frac{\mathbb{V}(Y_\ell)}{C_\ell}} \sum_{i=0}^L \sqrt{\mathbb{V}(Y_i) C_i} \rceil, \quad (6)$$

where C_ℓ is the complexity (computational cost) of evaluating Q_ℓ , i.e. taking one sample at level ℓ . Eq. (6) can also be written as

$$N_\ell = \lceil s^{\Gamma(L-\ell)} N_L \rceil, \quad (7)$$

where

$$\Gamma = \frac{1}{2}(\gamma + \beta).$$

The computational complexity of the MLMC algorithm is given by $C_{\text{MLMC}} = \sum_{\ell=0}^L N_\ell C_\ell$ and it can be bounded using **Assumption (3.2)** and **Assumption (3.3)** as

$$C_{\text{MLMC}} \leq \varepsilon^{-2 - \max(0, \frac{\gamma-\beta}{\alpha})}, \quad (8)$$

with an additional logarithmic factor when $\gamma = \beta$.

There are three different situations depending on whether the (computational cost required to reduce the) statistical error dominates, is of the same order of, or is dominated by the discretization error. In the first case, $\beta > \gamma$ and because the statistical error is eliminated using coarse grids, the complexity of the MLMC algorithm is independent of the fine mesh size. In the last one $\beta < \gamma$ and the complexity includes a factor $\varepsilon^{-\gamma/\alpha}$ that comes from the discretization error, which also appears in the MC complexity estimate (3). In applications of PDEs with random coefficients, it is typical to have $\beta = 2\alpha$ and in this case, $C_{\text{MLMC}} \leq \varepsilon^{-\gamma/\alpha}$ which is the complexity of one sample at the finest grid, that is, the MLMC algorithm does not *see* the sampling cost [55]. Thus, the relation between parameters β and γ determines whether the effort should be concentrated on coarse or fine meshes.

From the previous discussion the crucial role of the constants α , β , γ , c_α , c_β and c_γ is apparent. In general, these constants are unknown (they are problem and solver dependent) but their estimation is required to define the optimal number of levels and samples in MLMC. The so-called adaptive MLMC algorithms have been developed to face this challenge by estimating these constants with extrapolations and/or statistical inference [27, 19, 20]. For example, the variances in Eq. (6) can be estimated (a posteriori) using the sample variance

$$\mathbb{V}(Y_\ell) \approx V(Y_\ell) = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (Y_\ell^i - \bar{Y}_\ell)^2. \quad (9)$$

We do not follow this strategy herein, but we exploit the knowledge of the problem at hand. For the problem of Section 2 $\beta = 4$ [66] and using a CG algorithm $\gamma = d + 1$, so the statistical error dominates the computational cost

and the optimal relation between samples at the hierarchy are given by Eq. (7) with $\Gamma = 4$ for $d = 3$ and $\Gamma = 3.5$ for $d = 2$.

Observe that **Assumption (3.1)** and **Assumption (3.3)** (redefining c_α to include $h_0^{-\alpha}$ and c_β to include $h_0^{-\beta}$) can be written as

$$|\mathbb{E}(Q_\ell - Q)| \leq c_\alpha s^{-\alpha\ell}, \quad \mathbb{V}(Y_\ell) \leq c_\beta s^{-\beta\ell}. \quad (10)$$

The corresponding sample approximations $|Q - \tilde{Q}_\ell|$ (see Eq. (4)) and $V(Y_\ell)$ (see Eq. (9)) statistically converge to these values and thus, with a large enough number of samples, one can check spatial convergence (i.e., with respect to ℓ) for these quantities too. In the numerical examples of Section 5 $s = 2$ and therefore, the (logarithmic) convergence rate with respect to the level is $2 \log(2)$ and $4 \log(2)$ respectively.

4. THE AGGREGATED FE METHOD

The mesh hierarchy required by the MLMC algorithm is built from \mathcal{T}_0 , an initial shape-regular partition of \mathcal{B} , generating $\mathcal{T}_{\ell+1}$ by uniform refinement of \mathcal{T}_ℓ . In our implementation, we consider (a bounding box) $\mathcal{B} := \prod_{i=1}^d [x_m^i, x_M^i]$ and \mathcal{T}_0 a uniform Cartesian mesh. At every level ℓ , we consider a piecewise linear approximation $M_\ell(\omega)$ of the manifold $\mathcal{M}(\omega)$. When $\mathcal{M}(\omega)$ is already a polytopal surface mesh, e.g. an STL mesh, no approximation is required at this step. However, if the resolution of $\mathcal{M}(\omega)$ is very fine compared to the coarser mesh size, it would involve a very costly numerical integration at coarser levels. A case of practical interest that will be addressed in this work is when $\mathcal{M}(\omega)$ is represented implicitly with a level-set function $\phi(\mathbf{x})$. In this case, we construct a C^0 Lagrangian FE space of arbitrary order on the grid \mathcal{T}_ℓ , which is represented with $\tilde{\mathcal{V}}(\mathcal{T}_\ell)$. In this case, one can define $M_\ell(\omega)$ as the result of applying the marching tetrahedra algorithm [24] to the interpolation $\phi_\ell(\mathbf{x})$ of the level-set function $\phi(\mathbf{x})$ onto $\tilde{\mathcal{V}}(\mathcal{T}_\ell)$. When considering quad meshes, one can simply use the decomposition of an n-cube into n-tetrahedra before applying the marching tetrahedra algorithm.

Given the oriented manifold $M_\ell(\omega)$ (and its corresponding interior, the domain $\mathcal{D}_\ell(\omega)$) and the background mesh \mathcal{T}_ℓ , due to the piecewise linear nature of the manifold, it is possible to compute exactly the intersection of each cell $K \in \mathcal{T}_\ell$ with $M_\ell(\omega)$ and $\mathcal{D}_\ell(\omega)$. We represent with $\mathcal{T}_\ell^{\mathcal{D}}(\omega)$ (resp. $\mathcal{T}_\ell^{\mathcal{M}}(\omega)$) the set of active (resp., cut) cells in \mathcal{T}_ℓ that intersect $\mathcal{D}_\ell(\omega)$ (resp., $M_\ell(\omega)$); $\mathcal{T}_\ell^{\mathcal{D}}(\omega) \setminus \mathcal{T}_\ell^{\mathcal{M}}(\omega)$ is the set of interior cells in $\mathcal{D}_\ell(\omega)$. Numerical integration over cells $K \in \mathcal{T}_\ell^{\mathcal{M}}(\omega)$ can be carried out by generating a simplicial mesh $\mathcal{I}_K^{\mathcal{D}}$ of $K \cap \mathcal{D}_\ell(\omega)$, e.g. using a Delaunay triangulation. Interior cell volume integration is straightforward, and we can simply take $\mathcal{I}_K^{\mathcal{D}} = K$. We proceed analogously for the surface integration of $M_\ell(\omega)$ and the volume integration of the interior cells and the integration meshes for cut cells $\mathcal{I}_\ell^{\mathcal{D}}(\omega) := \bigcup_{K \in \mathcal{T}_\ell^{\mathcal{D}}(\omega)} \mathcal{I}_K^{\mathcal{D}}$; analogously for the surface mesh $\mathcal{I}_\ell^{\mathcal{M}}(\omega)$. These integration meshes always exist and are cell-wise local, i.e., no conformity is required among cells. We note that $\mathcal{I}_\ell^{\mathcal{D}}(\omega)$ is a body-fitted mesh of $\mathcal{D}_\ell(\omega)$ and $\mathcal{I}_\ell^{\mathcal{M}}(\omega)$ is a surface mesh of $M_\ell(\omega)$. However, the shape regularity of the meshes is not relevant since they are only used for integration purposes. For higher (than one) order approximations of $M(\omega)$, we can consider the same approach described above supplemented with an additional degree elevation of the linear cell-wise integration meshes. These techniques impose some constraints on the ratio between surface curvature and mesh resolution and are not being used in Section 5, since we restrict ourselves to first order FEs.

In this work, H^1 -conforming Lagrangian space of arbitrary order are considered for the discretization of the PDE problem at hand. We denote with $\tilde{\mathcal{V}}_\ell(\mathcal{T}^{\mathcal{D}}, \omega)$ the FE space on the grid $\mathcal{T}_\ell^{\mathcal{D}}(\omega)$. The space $\mathcal{T}_\ell^{\mathcal{D}}(\omega)$ is defined on $\tilde{\mathcal{D}}(\omega) \supseteq \mathcal{D}(\omega)$, where $\tilde{\mathcal{D}}(\omega)$ is the union of all cells in $\mathcal{T}^{\mathcal{D}}(\omega)$. Thus, it is not a body-fitted mesh of $\mathcal{D}(\omega)$ and embedded FE techniques must be used to properly capture the geometry. Let us consider the weak form of (1) in which the essential boundary conditions are imposed in a weak form using Nitsche's method. First, we define the cell-wise forms

$$\begin{aligned} \mathcal{A}_K^\ell(u, v) &:= \int_{K \cap \mathcal{D}_\ell(\omega)} \nabla u \cdot \nabla v dV + \int_{K \cap M_\ell(\omega)} (\tau_K uv - v(\mathbf{n} \cdot \nabla u) - u(\mathbf{n} \cdot \nabla v)) dS, \\ b_K^\ell(v) &:= \int_{K \cap \mathcal{D}_\ell(\omega)} f v dV + \int_{K \cap M_\ell(\omega)} (\tau_K u_0 - (\mathbf{n} \cdot \nabla v) u_0), \end{aligned}$$

where $\tau_K \approx O(h_\ell^{-2})$ is a positive parameter that must be large enough to ensure coercivity of the bilinear form. The global form $\mathcal{A}^\ell : H_0^1(\mathcal{D}_\ell(\omega)) \rightarrow H^{-1}(\mathcal{D}_\ell(\omega))$ and right-hand side term $b^\ell \in H^{-1}(\mathcal{D}_\ell(\omega))$ are stated as the sum of

the element contributions, i.e.,

$$\mathcal{A}^\ell(u, v) := \sum_{K \in \mathcal{T}_\ell^{\mathcal{D}}(\omega)} \mathcal{A}_K^\ell(u, v), \quad b^\ell(v) := \sum_{K \in \mathcal{T}_\ell^{\mathcal{D}}(\omega)} b_K^\ell(v).$$

The use of embedded FE methods has a dramatic impact on the condition number of the resulting linear systems, known as the *small cut cell* problem. Given that the mesh \mathcal{T}_ℓ is shape regular, one can define its characteristic cell size h_ℓ . It is well-known in body-fitted FE methods that the condition number of the linear system that results from the discretization of second order elliptic PDEs is $O(h_\ell^{-2})$. However, for unfitted FE methods, it also deteriorates with the maximum of $|K \cap \mathcal{D}_\ell(\omega)|^{-1}|K|$ among all cells in $K \in \mathcal{T}_\ell^{\mathcal{D}}$. The portion of the cell in the physical domain cannot be controlled and it can be arbitrarily close to zero. As a result, embedded FE methods can produce almost singular linear systems to be solved, and thus are not reliable.

To fix the small cut cell problem issue, we consider the AgFEM recently proposed in [9]. The idea is to start with the mesh $\mathcal{T}_\ell^{\mathcal{D}}$ and generate an aggregated mesh $\mathcal{G}_\ell(\omega)$ in which the cut cells with small volume ratio within the physical domain are aggregated to one interior cell (see [9, Alg. 2.1]). Such aggregation is next used to define constraints over the standard FE space $\tilde{\mathcal{V}}(\mathcal{T}_\ell^{\mathcal{D}}, \omega)$. We represent the resulting aggregated finite element (AgFE) space with $\mathcal{V}(\mathcal{G}_\ell, \omega) \subseteq \tilde{\mathcal{V}}(\mathcal{T}_\ell^{\mathcal{D}}, \omega)$. We refer the interested reader to [9] for a detailed exposition of the mesh aggregation algorithm, the computation of constraint, the implementation issues, and its numerical analysis, to [67] for its parallel implementation, and to [8] for a mixed AgFE space for the Stokes problem. AgFE spaces are not affected by the small cut cell problem and the condition number of the resulting matrix is $O(h_\ell^{-2})$ [9]. The numerical experiments in [67] show that a standard parallel algebraic multigrid solver is effective to solve AgFE linear systems using default settings. The AgFE space approximation of (1) at a given level ℓ reads as follows: find $u_\ell \in \mathcal{V}(\mathcal{G}_\ell, \omega)$ such that

$$\mathcal{A}^\ell(u_\ell, v_\ell) = b^\ell(v_\ell) \quad \text{for any } v_\ell \in \mathcal{V}(\mathcal{G}_\ell, \omega). \quad (11)$$

We note that all these terms can be computed using the cell-wise integration triangulations commented above. The condition number of the matrix resulting from the AgFEM discretization has been proven to be $O(h_\ell^{-2})$ as in body-fitted methods (see [9]). All the steps required to apply AgFEM in the frame of MLMC are listed in Alg. 1.

5. NUMERICAL EXAMPLES

Register for free at <https://www.scipedia.com> to download the version without the watermark

proposed approach. The first of them is a convergence test whose goal is to assess the actual convergence rates in the implementation. The second one is the solution of (1) in a complex random domain illustrating in a quantitative way the robustness of the proposed approach. Finally we solve (1) in a situation where the randomness may induce changes in the topology. These numerical experiments have been carried out using FEMPAR v1.0.0 [7, 6], an open source object-oriented scientific computing library for the simulation of complex multiphysics problems governed by PDEs at large scales that incorporates embedded FE machinery [67].

The implementation of EMLMC was made in FEMPAR by developing a new software layer which permits the concurrent execution of a model, understood as the actual implementation of an algorithm that permits the solution of a PDE. In our case, the model was developed using the numerical tools provided by FEMPAR to perform the automatic generation of background meshes \mathcal{T}_ℓ of arbitrary size, the level-set description of the geometry and its interpolation, the construction of AgFEM discretization, the numerical integration and the solution of the linear systems using domain decomposition preconditioners. The implementation of sampling methods made in this new module [5] exploits three levels of parallelism, across levels, across samples and in the evaluation of each sample, e.g. employing a domain decomposition method, although the latter feature has not been exploited in the examples below (each sample is computed using one processor).

5.1. A convergence test. In this section we consider an example with analytic solution proposed in [15]. The spatial domain $\mathcal{D}(R) \subset \mathbb{R}^2$ is a circle centered at $(0.5, 0.5)$ whose radius R is a truncated Gaussian random variable with mean $\mu = 0.3$, standard deviation $\sigma = 0.025$, and restricted to $[a, b]$ with $a = 0.2$ and $b = 0.4$, which we denote by $\mathcal{TN}(a, b, \mu, \sigma)$. In this random domain we solve problem Eq. (1) with $f = 4$, $k = 1$, and Dirichlet boundary conditions $u_0 = u_{\text{ex}}$ where

$$u_{\text{ex}}(\mathbf{x}) := R^2 - |\mathbf{x}_1 - 0.5|^2 - |\mathbf{x}_2 - 0.5|^2, \quad (12)$$

Algorithm 1: EMLMC

```

1 for  $\ell \in \{1, \dots, L\}$  do
2   Compute  $\mathcal{T}_\ell$  as uniform refinement of  $\mathcal{T}_{\ell-1}$ 
3   Generate the FE space  $\tilde{\mathcal{V}}(\mathcal{T}_\ell)$ 
4 end for
5  $\tilde{Q}_L = 0$ 
6 for  $\ell \in \{0, \dots, L\}$  do
7    $\tilde{Y}_\ell = 0$ 
8   for  $i \in \{1, \dots, N_\ell\}$  do
9     Sample the random variables and get  $\omega^i$ 
10    Set  $k = \ell$ 
11    Interpolate the level-set function  $\phi(x, \omega^i)$  in the FE space  $\tilde{\mathcal{V}}(\mathcal{T}_k, \omega^i)$  to obtain  $\phi_k(x, \omega^i)$ 
12    Run the marching tetrahedra algorithm for  $\phi_k(x, \omega^i)$  and compute  $\mathcal{M}_k(\omega^i)$  and  $\mathcal{D}_k(\omega^i)$ 
13    Compute the intersection of cells in  $\mathcal{T}_k(\omega^i)$  with  $\mathcal{M}_k(\omega^i)$  and  $\mathcal{D}_k(\omega^i)$  to obtain  $\mathcal{T}_k^{\mathcal{D}}(\omega^i)$  and  $\mathcal{T}_k^{\mathcal{M}}(\omega^i)$  and the cell-local
        integration meshes  $\mathcal{I}_k^{\mathcal{M}}(\omega^i)$  and  $\mathcal{I}_k^{\mathcal{D}}(\omega^i)$ 
14    Generate the FE space  $\tilde{\mathcal{V}}(\mathcal{T}_k^{\mathcal{D}}, \omega^*)$ 
15    Run the mesh aggregation algorithm in [9, Alg. 2.1] to compute the aggregated mesh  $\mathcal{G}_k(\omega^i)$ 
16    Generate the AgFE space  $\mathcal{V}(\mathcal{G}_k, \omega)$  as in [9, Sect. 3]
17    Compute the solution  $u_k(\omega^*)$  of (11)
18    Compute the  $Q_k = Q(u_k)$ 
19     $\tilde{Y}_\ell = \tilde{Y}_\ell + Q_k^i$ 
20    if  $\ell > 0$  then
21      Set  $k = \ell - 1$ 
22      Repeat lines 11 to 18.
23     $\tilde{Y}_\ell = \tilde{Y}_\ell - Q_k^i$ 
24    end if
25  end for
26   $\tilde{Q}_L = \tilde{Q}_L + \tilde{Y}_\ell/N_\ell$ 
27 end for

```

SCIPEDIA

is the exact solution. We consider two different (although similar) QoIs, the average of the solution over the whole

Register for free at <https://www.scipedia.com> to download the version without the watermark

that is,

$$Q_i(u) := \frac{1}{|\omega_i|} \int_{\omega_i} u \quad (13)$$

where $|\omega_i|$ denotes the area of ω_i . Using Eq. (12) we get $Q_1 = 0.5R^2$ and $Q_2 = R^2 - 2\delta^2/3$. The probability density function (PDF) of the \mathcal{TN} can be written in terms of the error function and the variance can be computed easily using, e.g., [1], which gives $\mathbb{E}(Q_1) = 0.04531216540324139$ and $\mathbb{E}(Q_2) = 0.08020766413981611$.

Because the MLMC estimation given by Eq. (4) is a random variable so is the squared error $(\tilde{Q}_L - \mathbb{E}(Q))^2$. Several realizations of the experiment will produce different results and its expectation error is determined by (5). Therefore, to evaluate the error decay in (10) we approximate it by a sample average, running each experiment K times, i.e., we compute the error as

$$\mathcal{E}_L^2 = \frac{1}{K} \sum_{k=1}^K (\tilde{Q}_L^k - \mathbb{E}(Q))^2, \quad (14)$$

where \tilde{Q}_L^k is the result of the k -th realization of the experiment. The numerical experiments have been performed with $K = 100$, whereas a value of $K = 30$ was found to be sufficient in the context of hyperbolic conservation laws [42]. Likewise we compute the averaged variances of the QoI differences between levels as

$$\mathcal{V}_\ell = \frac{1}{K} \sum_{k=1}^K V(Y_\ell^k) \quad (15)$$

where $V(Y_\ell^k)$ is given by Eq. (9) for the k -th realization of the experiment. In practice, each realized MLMC computation, here indexed by k , is computed with a different seed of the pseudo random number generator (PRNG).




FIGURE 1. Averaged error \mathcal{E}_L (given by Eq. (14)) for each level and sample size (as from Eq. (7)) evaluated with Q_1 (Full Domain) and Q_2 (Partial Domain) from Eq. (13) compared to the theoretical expected decay rate.

In the studies here we utilize the xoshiro256** algorithm [11] for generating pseudo random sequences, that has a 256 bit state. The algorithm is suited for this work as it is quick to advance the state, and statistically robust when seeded well. The detailed seeding procedure is not especially relevant to the results presented here, and those details are a bit beyond the scope of this paper, but the basic process can be given. The seeding process uses the recommended splitmix64 PRNG [60] to advance from an initial 64 bit seed. The outputs from this are used, together with physically generated random bits, to seed the state space of a Mersenne Twister [41], specifically MT19937-64, and the outputs of the Mersenne Twister are used, together with physically generated random bits, to seed the initial state of the xoshiro256**. The state space of the generator has a period of $2^{256} - 1$, meaning that after 2^{256} evaluations the state of the generator repeats, and there are no repetitions before that. The algorithm is adapted to parallelization by having easily computed a jump of 2^{128} states. Using this, jumps are performed to set the initial state for each sample. As a result, each process has an effective 2^{128} calls to the random number generator before refreshing state. This is realized by simple arithmetic, allowing for efficient and efficient state computation. The jumps are done consistently, so that the initial state depends only on the seed and sample number, to allow reproducibility of the stochastic inputs regardless of the scheduling of order of computations.

We consider $L = 5$ and the number of samples per level are defined according to Eq. (7) with $\Gamma = 7/2$. Each sample is computed taking $\mathcal{B} = [0, 1]^2$ and \mathcal{T}_0 a Cartesian mesh of 8×8 FEs and $\mathcal{T}_{\ell+1}$ is generated by dividing each element of \mathcal{T}_ℓ into four subelements (two per direction). Fig. (1) shows the average error \mathcal{E}_L as a function of the level for $N_5 = \{3, 6\}$. Even these low values of N_5 , shows a good convergence of the level estimates, in agreement with the expected decay rate in **Assumption (3.1)**.

Fig. (2) shows the sample variance estimate (15) as a function of the level. The variance of the QoI differences between levels decays with the level at the rate slightly higher than the one predicted by the theory in [66], verifying **Assumption (3.3)**. The determination of these variances is important in adaptive MLMC methods as they are required to obtain the optimal number of samples per level in Eq. (6).

The number of samples in Eq. (6) also depends on the complexity of sampling at each level C_ℓ which is shown in Fig. (3). This complexity is measured as the CPU time per sample averaged at each level, i.e., if t_ℓ^i is the elapsed CPU time in sample i of level ℓ , we estimate the complexity of sampling at each level as

$$C_\ell = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} t_\ell^i.$$

As it can be seen in Fig. (3), the scaling is in agreement with **Assumption (3.2)**.

After verifying that the EMLMC algorithm shows an error and variance decay and cost increase that is in agreement with **Assumption (3.1)** to **Assumption (3.3)** and because these assumptions permit to obtain the complexity bound Eq. (8) we expect the actual error to decrease with cost increase accordingly. The actual error





FIGURE 2. Averaged sample variance between levels (given by Eq. (15)) evaluated with Q_1 (Full Domain) and Q_2 (Partial Domain) from Eq. (13).



Register for free at <https://www.scipedia.com> to download the version without the watermark

FIGURE 3. Computational cost as a function of the level.

decrease with respect to the total complexity, measured as

$$C = \sum_{\ell=0}^L \sum_{i=1}^{N_{\ell}} t_{\ell}^i$$

is shown in Fig. (4) and the agreement with the scaling Eq. (8) is evident. The important gain of MLMC with respect to MC in deterministic domains with random coefficients is also observed for the EMLMC with this random geometry.

5.2. Dealing with complex random geometries. The aim of this example is to show the robustness of the proposed approach to perform UQ in complex random geometries. To this end, we consider a stochastic extension of a benchmark commonly used to test numerical methods designed to deal with interfaces described by a level-set [35, 13], the Poisson problem in a popcorn domain, a ball with spherical protuberances attached to its surface. Therefore, we aim to solve problem Eq. (1) with $k = 1$ and $f = -\nabla^2 u_{\text{ex}}$ and Dirichlet boundary conditions $u_0 = u_{\text{ex}}$, where

$$u_{\text{ex}} = \sin(k \|\mathbf{x} - \mathbf{x}_c\|)$$

is defined in $\mathcal{B} = [0, 1]^d$ with $\mathbf{x}_c = (1/2, 1/2, 1/2)$ and each sample is computed in a domain $\mathcal{D}(\omega)$ defined through a level set function, i.e. $\mathbf{x} \in \mathcal{D}(\omega)$ iff $\phi(\mathbf{x}, \omega) < 0$. The classical definition of the popcorn domain is

$$\phi(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_0\| - \rho_0 - \sum_{i=1}^n A \exp \|\mathbf{x} - \mathbf{x}_i\| / \sigma$$




FIGURE 4. Trade off between error and cost, evaluated with Q_1 (Full Domain) and Q_2 (Partial Domain) from Eq. (13).

with given parameters ρ_0 (the radius of the ball), \mathbf{x}_i (the location of spherical protuberances), A and σ (which control the size of the protuberances) (see e.g. [13]).

We consider a stochastic variation of this geometry defined by a random level set for $d = 2, 3$, although in this section we are primarily interested in the case $d = 3$. This function is build to represent a random central ellipse (labelled as 0) with smaller random ellipses attached to its boundary (labelled by j with $1 \leq j \leq n$) and is defined by

$$\phi(\mathbf{x}, \omega) = \min_{j \in \{0:n\}} (\|\mathbf{x} - \mathbf{x}_j\|_{A_j \mathbf{R}_j} - \rho_j),$$

where \mathbf{x}_j are random centers and ρ_j random radius. The deformation from circles to ellipses is achieved by the norm

$$\|\mathbf{y}\|_{A_j \mathbf{R}_j} = \sqrt{\langle (A_j \mathbf{R}_j)^{-1} \mathbf{y}, (A_j \mathbf{R}_j)^{-1} \mathbf{y} \rangle}.$$

which depends on stretching (A_j) and rotation (\mathbf{R}_j) matrices defined as follows. Distortion matrices A_j are square, diagonal matrices with diagonal entries $((A_j)_{ii})$ defined by parameters $A'_j(i)$ ($1 \leq i \leq d$)

$$(A_j)_{ii} := \frac{A'_j(i)}{\sqrt{\sum_{i=1}^d (A'_j(i))^2}}.$$

Rotations are defined in each (x_i, x_{i+1}) plane ($1 \leq i \leq d - 1$) through the rotation matrix $\mathbf{R}'(\Theta_j(i))$ defined from random parameters $\Theta_j(i)$ as

$$\mathbf{R}'(\Theta_j(i)) := \begin{pmatrix} \cos(\Theta_j(i)) & -\sin(\Theta_j(i)) \\ \sin(\Theta_j(i)) & \cos(\Theta_j(i)) \end{pmatrix},$$

Let $\mathbf{R}(\Theta_j(i))$, be the embedding of $\mathbf{R}'(\Theta_j(i))$ inside \mathbb{R}^d , acting as the identity on the other $d - 2$ coordinates of the vector. We then define the application of all rotations in \mathbb{R}^d by one matrix as \mathbf{R}_j defined by applying these rotations $\mathbf{R}(\Theta_j(1)), \dots, \mathbf{R}(\Theta_j(d - 1))$, in order, or as one single matrix denoted as,

$$\mathbf{R}_j := \prod_{i=1}^{d-1} \mathbf{R}(\Theta_j(i)) = \mathbf{R}(\Theta_j(d - 1)) \cdots \mathbf{R}(\Theta_j(1)).$$

The random parameters describing this geometry are distributed according to

- $n \in \mathcal{P}(11)$
- $A'_j \in \mathcal{U}(0.8, 1.3)^d$
- $\mathbf{x}_0 \in \mathcal{U}(0.4, 0.6)^d$
- $\rho_0 \in \mathcal{U}(0.1, 0.2)$
- $\Theta_j \in \mathcal{U}(0, 2\pi)^{d-1}$
- $\mathbf{y}_j \in \mathcal{S}^d$
- $\mathbf{x}_j = \mathbf{x}_0 + \rho_0 A_0 \mathbf{R}_0 \mathbf{y}_j$

- $\rho_j \in \mathcal{U}(0.03, 0.1)$

where \mathcal{P} and \mathcal{U} denote a Poisson and uniform distributions, respectively, and \mathcal{S}^d is a uniform distribution on the surface of a d -dimensional unit sphere. Note that using a level-set description permits to define complex domain through Boolean operations. In particular, given two level-set functions ϕ_1 and ϕ_2 describing Ω_1 and Ω_2 respectively, the function $\min(\phi_1, \phi_2)$ describes $\Omega_1 \cup \Omega_2$. A few samples of these stochastic domains are shown in Fig. (5), where the important variation of the morphology and size of the domains is apparent.

An important problem with sampling methods, as described in Section 1 is the robustness of the solver employed to compute samples and this is an important feature of the AgFEM employed herein and described in Section 4. We performed a quantitative evaluation of this feature by running the EMLMC method with and without the aggregation procedure using an iterative solver for the linear system arising from the FE discretization. Due to the symmetry of problem (1), we use the CG algorithm; as it is well-known, the number of iterations required to converge is proportional to the square root of the condition number of the system matrix. We compute statistics of the number of iterations (maximum, minimum and average) taking 1000 samples per level. The results obtained in a 2D version of the problem are shown in Fig. (6) whereas those obtained in the 3D setting are shown in Fig. (7).

5.3. Dealing with stochastic topologies. In this section we consider a problem in a stochastic domain where the randomness may induce changes in the topology. It consists of a square plate with two circular holes of uncertain position which may overlap to form a single one. We consider problem (1) with Dirichlet boundary conditions




FIGURE 5. Some realizations of the solution of Eq. (1) in a stochastic popcorn domain.

on the left and right sides and zero Neumann boundary on the top and bottom sides and on the boundary of the internal holes. Physically, this problem corresponds to the heat transfer in a plate where the left (low) and right (high) temperatures are prescribed and heat flows through the plate at a rate that depends on the location and size of the holes. The natural quantity of interest is therefore

$$Q(u) := \int_{x_1=0} \partial_{x_1} u, \quad (16)$$

The stochastic domain is defined, again, through a level set function $\phi(\mathbf{x}, \omega) = \phi_1(\mathbf{x}, \omega) + \phi_2(\mathbf{x}, \omega)$ where

$$\phi_i(\mathbf{x}, \omega) = \|\mathbf{x} - \mathbf{x}_i(\omega)\| - \rho_i$$

for $i = 1, 2$ represent the two interior circular holes (which do not conduct the heat) are distributed one along the upper half of the plate and one along the lower half of the plate. We study the influence of the radius on the average total heat flux through the plate for random position of the holes. We consider three cases $\rho_i = 0.18$, $\rho_i = 0.2$ or




FIGURE 6. Number of CG iterations required to converge to a tolerance of 10^{-8} in the relative norm the solution of each (2D) sample of the EMLMC as a function of the mesh size on each level, with and without aggregation.




FIGURE 7. Number of CG iterations required to converge to a tolerance of 10^{-8} in the relative norm the solution of each (3D) sample of the EMLMC as a function of the mesh size on each level, (a) with and (b) without aggregation (note the difference in the scales).

$\rho_i = 0.22$ for $i = 1, 2$. The first components of \mathbf{x}_i for $i = 1, 2$ (holes' centers) are drawn uniformly from $[0.23, 0.77]$. The upper circle has the x_2 coordinate of its center drawn uniformly from $[0.70, 0.76]$, and the lower circle has the x_2 coordinate of its center drawn uniformly from $[0.24, 0.30]$. The union of these circles form the non-conductive interior domain.

We consider $L = 5$ and the number of samples per level are defined according to Eq. (7) with $\Gamma = 7/2$. Each sample is computed by taking $\mathcal{B} = [0, 1]^2$ and \mathcal{T}_0 a Cartesian mesh of 8×8 FEs and $\mathcal{T}_{\ell+1}$ is generated by dividing each element of \mathcal{T}_ℓ into four subelements (two per direction). A few samples of the solution are shown in Fig. (9), where the topology changes in the domain can be observed. The proposed EMLMC method is able to automatically compute realizations using the same procedure followed in other cases, no special treatment of topology changes is required.

Fig. (8) shows \tilde{Q}_L for the three different radii. As it can be seen, and is expected, when the size of the holes is increased, the effective area available for heat conduction is decreased, thus resulting in a reduction of the heat flux through the plate.

6. CONCLUSIONS

In this article we propose the EMLMC method which consists of drawing samples in MLMC using a robust embedd method, the AgFEM, for the discretization of the problem. Once again, we emphasize that applying the standard MLMC on complex geometries can be hard if body-fitted methods are used for the problem discretization due to the difficulty of the mesh-hierarchy generation, which typically require human intervention. On top of that, when dealing with random geometries the construction of the mesh-hierarchy might even be impossible using body-fitted techniques or may simply fail for a particular sample. Using embedded discretization methods on a mesh-hierarchy generated on a bounding-box circumvents these problems.

The resulting algorithm is a powerful method to perform UQ on complex random domains thanks to the robustness provided by the AgFEM algorithm. The numerical examples presented herein show that the error and variance decays and the complexity are similar to those observed in the standard MLMC (with deterministic geometries and body-fitted discretizations) and are in line with theoretical expectations. Therefore, the EMLMC shows the same cost reduction with respect to MC than the standard MLMC making it an excellent method to perform UQ on complex random domains.

REFERENCES

- [1] <https://github.com/cossio/TruncatedNormal.jl>.
- [2] Ivo Babuska and Jan Chleboun. Effects of uncertainties in the Domain on the Solution of Neumann Boundary Value Problems. *Mathematics of Computation*, 71(240):1339–1370, 2001.
- [3] Ivo Babuska and Jan Chleboun. Effects of uncertainties in the domain on the solution of Dirichlet boundary value problems. *Numerische Mathematik*, 93(4):583–610, 2003.

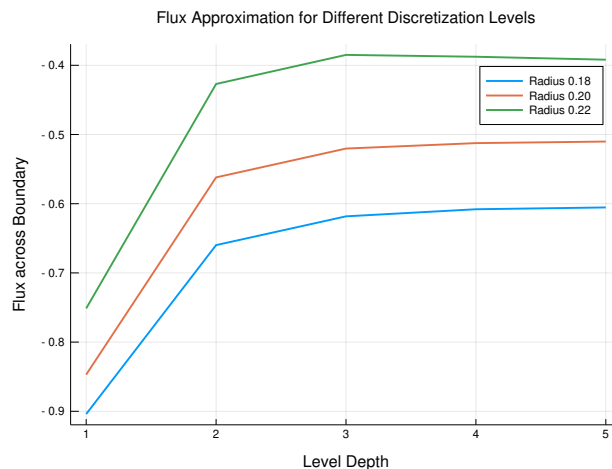


FIGURE 8. Average $Q(u)$ value as from Eq. (16) for three different radii at different levels of computation.

- [4] Ivo Babuska, Fabio Nobile, and Raúl Tempone. A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data. *SIAM Review*, 52(3):317–355, 2010.
- [5] Santiago Badia, Jerrad Hampton, and Javier Principe. A massively parallel implementation of Multilevel Monte Carlo methods. *In preparation*, 2019.
- [6] Santiago Badia and Alberto F. Martín. A tutorial-driven introduction to the parallel finite element library FEMPAR v1.0.0. *arXiv:1908.00891*, 2019.
- [7] Santiago Badia, Alberto F. Martín, and Javier Principe. FEMPAR: An Object-Oriented Parallel Finite Element Framework. *Archives of Computational Methods in Engineering*, 25(2):195–271, oct 2018.
- [8] Santiago Badia, Alberto F. Martin, and Francesc Verdugo. Mixed Aggregated Finite Element Methods for the Unfitted Discretization of the Stokes Problem. *SIAM Journal on Scientific Computing*, 40(6):B1541–B1576, jan 2018.
- [9] Santiago Badia, Francesc Verdugo, and Alberto F. Martín. The aggregated unfitted finite element method for elliptic problems. *Computer Methods in Applied Mechanics and Engineering*, 336:533–553, jul 2018.
- [10] Andrea Barth, Christoph Schwab, and Nathaniel Zollinger. Multi-level Monte Carlo Finite Element method for elliptic PDEs with stochastic coefficients. *Numerische Mathematik*, 119(1):123–161, sep 2011.
- [11] David Blackman and Sebastiano Vigna. Scrambled linear pseudorandom number generators. *arXiv preprint arXiv:1805.01407*, 2018.




FIGURE 9. Some realizations of the solution of Eq. (1) in a plate with two (stochastic) holes.

- [12] M G Blyth and C Pozrikidis. Heat conduction across irregular and fractal-like surfaces. *International Journal of Heat and Mass Transfer*, 46(8):1329–1339, 2003.
- [13] Erik Burman, Susanne Claus, P Hansbo, M G Larson, and André Massing. CutFEM: Discretizing Geometry and Partial Differential Equations. *International Journal for Numerical Methods in Engineering*, 104(7):472–501, 2015.
- [14] D O Cajueiro, V A de A Sampaio, C M C de Castilho, and R F S Andrade. Fractal properties of equipotentials close to a rough conducting surface. *Journal of Physics: Condensed Matter*, 11(26):4985, 1999.
- [15] Claudio Canuto and Tomas Kozubek. A fictitious domain approach to the numerical solution of PDEs in stochastic domains. *Numerische Mathematik*, 107(2):257–293, aug 2007.
- [16] Julio E Castrillón-Candás, Fabio Nobile, and Raúl F Tempone. Analytic regularity and collocation approximation for elliptic PDEs with random domain deformations. *Computers and Mathematics with Applications*, 71:1173–1197, 2016.
- [17] Julia Charrier, Robert Scheichl, and Aretha Leonore Teckentrup. Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods. *SIAM Journal on Numerical Analysis*, 51(1):322–352, 2013.
- [18] Jehanzeb H Chaudhry, Nathaniel Burch, and Donald Estep. Efficient Distribution Estimation and Uncertainty Quantification for Elliptic Problems on Domains with Stochastic Boundaries. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1127–1150, jan 2018.
- [19] K A Cliffe, M B Giles, R Scheichl, A L Teckentrup, C W Oosterlee, and A K Borzi A Cliffe. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput Visual Sci*, 14:3–15, 2011.
- [20] Nathan Collier, Abdul-Lateef Haji-Ali, Fabio Nobile, Erik von Schwerin, and Raúl Tempone. A continuation multilevel Monte Carlo algorithm. *BIT Numerical Mathematics*, 55(2):399–432, 2015.
- [21] M Dambrine, I Greff, H Harbrecht, and B Puig. Numerical Solution of the Poisson Equation on Domains with a Thin Layer of Random Thickness. *SIAM Journal on Numerical Analysis*, 54(2):921–941, jan 2016.
- [22] M Dambrine, I Greff, H Harbrecht, and B Puig. Numerical solution of the homogeneous Neumann boundary value problem on domains with a thin layer of random thickness. *Journal of Computational Physics*, 330:943–959, feb 2017.
- [23] K. S. Deffeyes, B. D. Ripley, and G. S. Watson. Stochastic geometry in petroleum geology. *Journal of the International Association for Mathematical Geology*, 14(5):419–432, oct 1982.
- [24] Akio Doi and Akio Koide. An Efficient Method of Triangulating Equi-Valued Surfaces by Using Tetrahedral Cells. *IEICE Transactions on Information and Systems*, (E74-D):214–224, 1991.
- [25] M M Fyrillas and C Pozrikidis. Conductive heat transport across rough surfaces and interfaces between two conforming media. *International Journal of Heat and Mass Transfer*, 44(9):1789–1801, 2001.
- [26] R Ghanem and P D Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag New York, 1991.
- [27] Michael B. Giles. Multilevel Monte Carlo Path Simulation. *Operations Research*, 56(3):607–617, jun 2008.
- [28] Michael B. Giles and Christoph Reisinger. Stochastic Finite Differences and Multilevel Monte Carlo for a Class of SPDEs in Finance. *SIAM Journal on Financial Mathematics*, 3(1):572–592, jan 2012.
- [29] Andrew Gordon and Catherine E Powell. A Preconditioner for Fictitious Domain Formulations of Elliptic PDEs on Uncertain Parameterized Domains. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):622–646, jan 2014.
- [30] Abdul-Lateef Haji-Ali, Fabio Nobile, Erik Von Schwerin, and Raúl Tempone. Optimization of mesh hierarchies in multilevel Monte Carlo samplers. 4:76–112, 2016.
- [31] Bernard J Hamrock, Steven R Schmid, and Bo O Jacobson. *Fundamentals of fluid film lubrication*. CRC press, 2004.
- [32] H Harbrecht, M. Peters, and M Siebenmorgen. Analysis of the domain mapping method for elliptic diffusion problems on random domains. *Numerische Mathematik*, 134(4):823–856, 2016.
- [33] Helmut Harbrecht and Jingzhi Li. First order second moment analysis for stochastic interface problems based on low-rank approximation. *ESAIM: Mathematical Modelling and Numerical Analysis*, 47(5):1533–1552, sep 2013.
- [34] Helmut Harbrecht, Reinhold Schneider, and Christoph Schwab. Sparse second moment analysis for elliptic problems in stochastic domains. *Numerische Mathematik*, 109(3):385–414, 2008.
- [35] Martin Hautefeuille, Chandrasekhar Annavarapu, and John E Dolbow. Robust imposition of {Dirichlet} boundary conditions on embedded surfaces. *International Journal for Numerical Methods in Engineering*, 90(1):40–64, apr 2012.
- [36] Ian Hutchings and Philip Shipway. *Tribology: friction and wear of engineering materials*. Butterworth-Heinemann, 2017.
- [37] Ahmed Kebaier. Statistical Romberg extrapolation: A new variance reduction method and applications to option pricing. *The Annals of Applied Probability*, 15(4):2681–2705, nov 2005.
- [38] Christopher Lang, Alireza Doostan, Kurt Maute, Christopher Lang, Alireza Doostan, and · K Maute. Extended stochastic FEM for diffusion problems with uncertain material interfaces. *Computational Mechanics*, 51(6):1031–1049, jun 2013.
- [39] Olivier P. Le Maitre and Omar M. Knio. *Spectral Methods for Uncertainty Quantification. With Applications to Computational Fluid Dynamics*. Springer, 2010.
- [40] Xian Luo, Martin R Maxey, and George Em Karniadakis. Smoothed profile method for particulate flows: Error analysis and simulations. *Journal of Computational Physics*, 228(5):1750–1769, 2009.
- [41] Makoto Matsumoto and Takuji Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation*, 8(1):3–30, 1998.
- [42] S Mishra and Ch. Schwab. Sparse tensor multi-level Monte Carlo finite volume methods for hyperbolic conservation laws with random initial data. *Mathematics of Computation*, 81(280):1979–2018, 2012.
- [43] S. Mishra, Ch. Schwab, and J. Sukys. Multi-level Monte Carlo finite volume methods for nonlinear systems of conservation laws in multi-dimensions. *Journal of Computational Physics*, 231(8):3365–3388, apr 2012.
- [44] S. Mishra, Ch. Schwab, J Sukys, and J. Sukys. Multilevel Monte Carlo Finite Volume Methods for Shallow Water Equations with Uncertain Topography in Multi-dimensions. *SIAM Journal on Scientific Computing*, 34(6):B761–B784, jan 2012.

- [45] Nicolas Moës, John Dolbow, and Ted Belytschko. A finite element method for crack growth without remeshing. *International Journal for Numerical Methods in Engineering*, 46(1):131–150, sep 1999.
- [46] P Surya Mohan, Prasanth B Nair, and Andy J Keane. Stochastic projection schemes for deterministic linear elliptic partial differential equations on random domains. *International Journal for Numerical Methods in Engineering*, 85(7):874–895, feb 2011.
- [47] Michael Nosonovsky and Bharat Bhushan. Biologically inspired surfaces: Broadening the scope of roughness. *Advanced Functional Materials*, 18(6):843–855, mar 2008.
- [48] A Nouy, M Chevreuil, and E Safatly. Fictitious domain method and separated representations for the solution of boundary value problems on uncertain parameterized domains. *Computer Methods in Applied Mechanics and Engineering*, 200(45-46):3066–3082, oct 2011.
- [49] A Nouy and A Clément. eXtended Stochastic Finite Element Method for the numerical simulation of heterogeneous materials with random material interfaces. *International Journal for Numerical Methods in Engineering*, 83(10):1312–1344, sep 2010.
- [50] A Nouy, A Clément, F Schoefs, and N Moës. An extended stochastic finite element method for solving stochastic partial differential equations on random domains. *Computer Methods in Applied Mechanics and Engineering*, 197(51-52):4663–4682, 2008.
- [51] Anthony Nouy, Franck Schoefs, and Nicolas Moës. X-SFEM, a computational technique based on X-FEM to deal with random shapes. *European Journal of Computational Mechanics*, 16(2):277–293, 2007.
- [52] Sang Woo Park, Marcos Intaglietta, and Daniel M Tartakovsky. Impact of endothelium roughness on blood flow. *Journal of Theoretical Biology*, 300:152–160, 2012.
- [53] Lucia Parussini, Valentino Pediroda, and Carlo Poloni. Prediction of geometric uncertainty effects on Fluid Dynamics by Polynomial Chaos and Fictitious Domain method. *Computers and Fluids*, 39(1):137–151, 2010.
- [54] Torben Pätz, Robert M Kirby, and Tobias Preusser. Ambrosio-Tortorelli Segmentation of Stochastic Images: Model Extensions, Theoretical Investigations and Numerical Methods. *International Journal of Computer Vision*, 103(2):190–212, jun 2013.
- [55] M. Piaroni, F. Nobile, and P. Leyland. A Continuation Multi Level Monte Carlo (C-MLMC) method for uncertainty quantification in compressible inviscid aerodynamics. *Computer Methods in Applied Mechanics and Engineering*, 326:20–50, 2017.
- [56] S. Richardson. A model for the boundary condition of a porous material. Part 2. *Journal of Fluid Mechanics*, 49(2):327–336, 1971.
- [57] Kausik Sarkar and Charles Meneveau. Gradients of potential fields on rough surfaces: Perturbative calculation of the singularity distribution function $\{f(\cdot)\}$ for small surface dimension. *Physical Review E*, 47(2):957–966, 1993.
- [58] Franck Schoefs, Mathilde Chevreuil, Olivier Pasqualini, and Mikaël Cazuguel. Partial safety factor calibration from stochastic finite element computation of welded joint with random geometries. 2016.
- [59] James Albert Sethian. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, volume 3. Cambridge university press, 1999.
- [60] Guy L Steele Jr, Doug Lea, and Christine H Flood. Fast splittable pseudorandom number generators. In *ACM SIGPLAN Notices*, volume 49, pages 453–472. ACM, 2014.
- [61] G. Stefanou, A. Nouy, and A. Clement. Identification of random shapes from images through polynomial chaos expansion of random level set functions. *International Journal for Numerical Methods in Engineering*, 79(2):127–155, jul 2009.
- [62] G Stefanou and M Papadarakakis. Stochastic finite element analysis of shells with combined random material and geometric properties. *Computer Methods in Applied Mechanics and Engineering*, 193(1-2):139–160, jan 2004.
- [63] N Sukumar, D L Chopp, N Moës, and T Belytschko. Modeling holes and inclusions by level sets in the extended finite-element method. *Computer Methods in Applied Mechanics and Engineering*, 190(46–47):6183–6200, 2001.
- [64] Daniel M Tartakovsky and Dongbin Xiu. Stochastic analysis of transport in tubes with rough walls. *Journal of Computational Physics*, 217(1):248–259, 2006.
- [65] G. I. Taylor. A model for the boundary condition of a porous material. Part 1. *Journal of Fluid Mechanics*, 49(02):319, sep 1971.
- [66] A. L. Teckentrup, R. Scheichl, M. B. Giles, and E. Ullmann. Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numer. Math*, 125:569–600, 2013.
- [67] Francesc Verdugo, Alberto F. Martín, and Santiago Badia. Distributed-memory parallelization of the aggregated unfitted finite element method. *Computer Methods in Applied Mechanics and Engineering*, 357:112583, dec 2019.
- [68] Dongbin Xiu and George Em Karniadakis. Modeling uncertainty in flow simulations via generalized polynomial chaos. *Journal of Computational Physics*, 187(1):137–167, 2003.
- [69] Dongbin Xiu and D.M. Tartakovsky. Numerical methods for differential equations in random domains. *SIAM Journal on Scientific Computing*, 28(3):1167–1185, jan 2006.
- [70] Mohsen Zayernouri, Sang-Woo Park, Daniel M Tartakovsky, and George Em Karniadakis. Stochastic smoothed profile method for modeling random roughness in flow problems. *Computer Methods in Applied Mechanics and Engineering*, 263:99–112, 2013.