

ON LOCAL FOURIER ANALYSIS OF MULTIGRID METHODS FOR PDEs WITH JUMPING AND RANDOM COEFFICIENTS*

PRASHANT KUMAR[†], CARMEN RODRIGO[‡], FRANCISCO J. GASPAR[†],
AND CORNELIS W. OOSTERLEE[§]

Abstract. In this paper, we propose a novel nonstandard local Fourier analysis (LFA) variant for accurately predicting the multigrid convergence of problems with random and jumping coefficients. This LFA method is based on a specific basis of the Fourier space rather than the commonly used Fourier modes. To show the utility of this analysis, we consider, as an example, a simple cell-centered multigrid method for solving a steady-state single phase flow problem in a random porous medium. We successfully demonstrate the predictive capability of the proposed LFA using a number of challenging benchmark problems. The information provided by this analysis could be used to estimate a priori the time needed for solving certain uncertainty quantification problems by means of a multigrid multilevel Monte Carlo method.

Key words. PDEs, random coefficients, multigrid, local Fourier analysis, multilevel Monte Carlo, uncertainty quantification

AMS subject classifications. 65F10, 65M22, 65M55

DOI. 10.1137/18M1173769

1. Introduction. A number of problems in science and engineering involve solving partial differential equations (PDEs) with random parameters or coefficients. The solution to such PDEs is of a stochastic nature, and the aim is to compute expected values and corresponding variances of a functional of the solution. For such uncertainty quantification (UQ) problems, Monte Carlo (MC) type methods are preferred due to their dimension-independent convergence. For any sampling-based approach, the availability of a highly efficient and robust (w.r.t. the random inputs) iterative solver becomes critical. In general, the samplewise computational cost can become highly heterogeneous, depending on the random inputs. Therefore, if the performance statistics of such solvers were known a priori, one could utilize this information to optimize and parallelize the MC simulations efficiently.

In this paper, we present a nonstandard local Fourier analysis (LFA) technique to predict the convergence rate of multigrid solvers for problems involving random and jumping coefficients. Standard LFA techniques are typically based on constant coefficient discretization stencils, whereas for stochastic PDEs we encounter varying coefficients throughout the computational domain, due to the randomness. One of the main contributions of this work is to generalize the LFA toward problems with

*Submitted to the journal's Methods and Algorithms for Scientific Computing section March 5, 2018; accepted for publication (in revised form) March 19, 2019; published electronically May 2, 2019.

<http://www.siam.org/journals/sisc/41-3/M117376.html>

Funding: The work of the second author was supported by the Spanish project FEDER/MCYT MTM2016-75139-R and the Diputación General de Aragón (Grupo de referencia APEDIF, ref. E24_17R). The work of the third author was supported by the European Union's Horizon 2020 research and innovation program under Marie Skłodowska-Curie grant agreement 705402, POROSOS.

[†]CWI, Centrum Wiskunde and Informatica, Amsterdam, The Netherlands (P.Kumar@cw.nl, gaspar@cw.nl, <http://www.unizar.es/pde/fjgaspar/>).

[‡]IUMA and Applied Mathematics Department, University of Zaragoza, Zaragoza, Spain (carmenr@unizar.es).

[§]CWI, Centrum Wiskunde and Informatica, Amsterdam, The Netherlands, and DIAM, Delft University of Technology, The Netherlands (c.w.oosterlee@cw.nl).

random and jumping coefficients, with the aim of predicting, a priori, the total time needed to solve UQ problems. Some efforts have already been made in [1] regarding the generalization of LFA for jumping coefficients. The novelty of our approach lies in the choice of basis functions. Here, we utilize a new basis from the Fourier space rather than the standard Fourier modes. We benchmark the prediction capability of the proposed LFA technique using a set of challenging jumping coefficient problems and a number of spatially correlated random fields with varying heterogeneity.

One of the results of this paper is the ability of the aforementioned LFA technique to predict accurately the multigrid convergence rates for elliptic PDEs with heterogeneous random coefficients, as encountered in the stochastic Darcy flow problem. In stochastic subsurface flow modeling, a lognormal random field often represents the permeability of an occurring heterogeneous porous medium [2, 3, 4]. Uncertainty may then be assessed by means of stochastic collocation approaches [5, 6], where the use of multigrid has been analyzed in [7]. In the current paper, we are interested in solving the PDE in the context of the multilevel Monte Carlo (MLMC) technique [8, 9, 10]. Within MC methods, samples of the stochastic random field are generated and the corresponding solution of the PDE with “frozen stochastic coefficients” is computed. Plain MC methods may require many thousands of random samples on a fine computational mesh, before the mean and variance of the numerical solution stabilize and converge. This may cost substantial CPU-time. MLMC methods are based on the decomposition of the expectation and variance operators into sequences of differences of these quantities on fine and coarse problem scales. The number of iterations required on each scale is different, where fewer samples (fewer random fields) are typically required on the finer scales, and larger numbers of samples are needed on the coarser, cheaper to compute, scales.

Combining MLMC methods with multigrid seems natural, where the multigrid method is employed for the numerical solution of a PDE which is based on a sample of the stochastic quantity on a certain (fine or coarse) scale. Due to the stochasticity, however, we deal with PDE problems with jumping coefficients, where different jump patterns are encountered each time a new random field is generated. The generalization of the LFA toward these PDEs will provide us insight into the average number of multigrid iterations and the spread of the convergence factors, among other things. This information helps to estimate the total CPU-time needed for the multiple multigrid computations in a MLMC setting. We consider it useful to construct a technique to assess the quality of the choice of the multigrid components in the context of the PDEs with random problems, before the actual multigrid computation has taken place.

In this work, we will employ a basic cell-centered multigrid (CCMG) algorithm for solving elliptic PDEs with a variable coefficients field. The components of this algorithm include a simple Gauss–Seidel iteration as the smoother, a piecewise constant prolongation operator and its adjoint as the restriction, and a direct discretization technique to define the discrete operators on the coarse grids. We show that for this special combination, the coarse-grid discretization operators are equivalent to the ones obtained from commonly used Galerkin operators [11]. We utilize this CCMG method to perform MLMC simulations with different permeability parameters. It may be surprising that such a basic algorithm converges well in the context of the generated random fields, where computation takes place for thousands of different samples. Although we restrict ourselves to this basic CCMG to demonstrate the accuracy of the predictions of the novel LFA technique, we emphasize that this approach can be used for a wider range of problems, discretizations, and multigrid methods. The proposed LFA technique allows us to deal with several challenging problems for

which it is not easily possible to apply the classical LFA.

The paper is organized as follows. In section 2 we introduce the context of PDEs with jumping and random coefficients, together with their discretization by a cell-centered finite volume scheme. A discussion on multigrid methods for this type of problem is also included, and the multigrid components that will be considered in this work are defined. Section 3 is devoted to the generalization of the LFA to deal with jumping coefficients and problems with random fields. In section 4, we present results obtained by this analysis for different benchmark problems with jumping coefficients. Section 5 presents the LFA results for problems with random coefficients, and in section 6 MLMC computations for PDEs with random coefficients are presented. Finally, in section 7 conclusions are drawn.

2. Jumping coefficients, random coefficients, multigrid. Robust and efficient iterative solution methods are very relevant for PDEs with variable coefficients. For PDEs with jumping coefficients, multigrid methods have already been shown to be this type of solver. When using MC methods, in the case of elliptic PDEs with random coefficient fields, many samples of the random field are generated and for each field the numerical solution should be computed. This can take substantial CPU-time if very many samples are required. For a fixed sample of the random field, we deal with an elliptic PDE with varying coefficients, due to the randomness. Multigrid comes in naturally as a highly efficient solution method for the resulting PDEs. In this section, we explain this setting and we briefly describe an efficient multigrid method based on a cell-centered grid and a finite volume discretization.

2.1. PDEs with jumping and with random coefficients. We start with classical PDE problems with jumping coefficients. In particular, we deal with the following two-dimensional diffusion equation on the square domain $D = (0, \ell)^2$:

$$\begin{aligned} (1) \quad & -\nabla \cdot (k(\mathbf{x}) \nabla u(\mathbf{x})) = f(\mathbf{x}), \quad \mathbf{x} \in D, \\ (2) \quad & u(\mathbf{x}) = g(\mathbf{x}), \quad \mathbf{x} \in \partial D, \end{aligned}$$

where $k(\mathbf{x})$ is a function which may be discontinuous across internal boundaries.

To discretize this problem, we use a *cell-centered finite volume method* based on the harmonic average of the diffusion coefficient $k(\mathbf{x})$. We consider a uniform grid D_h with the same step size $h = \ell/M$, $M \in \mathbb{N}$ in both directions,

$$(3) \quad D_h = \{(x_{i_1}, x_{i_2}); x_{i_\alpha} = (i_\alpha - 1/2)h, i_\alpha = 1, \dots, M, \alpha = 1, 2\}.$$

This gives, for each interior cell with center (x_{i_1}, x_{i_2}) , denoted by $D_h^{i_1, i_2}$, a five-point scheme

$$(4) \quad c_{i_1, i_2}^h u_{i_1, i_2} + w_{i_1, i_2}^h u_{i_1-1, i_2} + e_{i_1, i_2}^h u_{i_1+1, i_2} + s_{i_1, i_2}^h u_{i_1, i_2-1} + n_{i_1, i_2}^h u_{i_1, i_2+1} = f_{i_1, i_2}^h,$$

where

$$\begin{aligned} w_{i_1, i_2}^h &= -\frac{2}{h^2} \frac{k_{i_1, i_2} k_{i_1-1, i_2}}{k_{i_1, i_2} + k_{i_1-1, i_2}}, & e_{i_1, i_2}^h &= -\frac{2}{h^2} \frac{k_{i_1, i_2} k_{i_1+1, i_2}}{k_{i_1, i_2} + k_{i_1+1, i_2}}, \\ s_{i_1, i_2}^h &= -\frac{2}{h^2} \frac{k_{i_1, i_2} k_{i_1, i_2-1}}{k_{i_1, i_2} + k_{i_1, i_2-1}}, & n_{i_1, i_2}^h &= -\frac{2}{h^2} \frac{k_{i_1, i_2} k_{i_1, i_2+1}}{k_{i_1, i_2} + k_{i_1, i_2+1}}, \\ c_{i_1, i_2}^h &= -(w_{i_1, i_2}^h + e_{i_1, i_2}^h + n_{i_1, i_2}^h + s_{i_1, i_2}^h), \end{aligned}$$

with, for instance, k_{i_1, i_2} the diffusion coefficient associated with the cell $D_h^{i_1, i_2}$. By *interior cell* we mean a cell for which none of its edges lies at the boundary of the domain. This scheme is changed appropriately for the cells close to the boundary.

As mentioned, we also consider elliptic PDEs with random coefficient fields. The PDE of our interest describes the steady-state single-phase flow in a random porous medium. Denoting by ω an event in the probability space $(\Omega, \mathbb{F}, \mathbb{P})$, with sample space Ω , σ -algebra \mathbb{F} , and probability measure \mathbb{P} , the permeability in the porous medium is described by $k(\mathbf{x}, \omega) : \overline{D} \times \Omega \rightarrow \mathbb{R}_+$. The PDE is then given by

$$(5) \quad -\nabla \cdot (k(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}), \quad \mathbf{x} \in D, \omega \in \Omega,$$

with f as a source term.

The engineering interest in the solution of (5) is typically found in expected values of linear functionals of the solution u , denoted by $Q := Q(u)$.

To discretize these problems, we use the same cell-centered finite volume method based on the harmonic average of the random diffusion coefficient as previously described for problems with jumping coefficients. We make the common assumption that the permeability random field is constant over each cell of the grid.

2.2. Multigrid for PDEs with jumping and with random coefficients.

In this work, the multigrid components for the above cell-centered discrete problems are chosen as follows. We use a lexicographic Gauss-Seidel iteration as the multigrid smoother and we consider standard coarsening obtained by doubling the mesh size in both directions. Each coarse cell is the union of four fine cells, and, since the unknowns are located in the cell centers, this results in a nonnested hierarchy of grids. We consider a simple prolongation operator P_{2h}^h , that is, the piecewise constant interpolation operator. In stencil notation, it is given by

$$(6) \quad P_{2h}^h = \begin{bmatrix} 1 & & 1 \\ & \star & \\ 1 & & 1 \end{bmatrix}_{2h}^h,$$

where \star denotes the position of a coarse-grid unknown. The classical stencil notation shows the contribution of the coarse-grid node to the neighboring fine-grid nodes. The restriction operator R_h^{2h} is chosen as the scaled adjoint of the prolongation, given in stencil form by

$$(7) \quad R_h^{2h} = \frac{1}{4} \begin{bmatrix} 1 & & 1 \\ & \star & \\ 1 & & 1 \end{bmatrix}_h^{2h}.$$

The coarse-grid operators are constructed by direct discretization defining the diffusion coefficients at the edges of the coarse cells appropriately, which we will describe in more detail. We assume that the diffusion coefficient $k(\mathbf{x})$ is piecewise constant on the fine grid. The flux over an edge, dependent on the solution in the two adjacent cells, is calculated based on the harmonic average. The values of the diffusion coefficients at a coarse edge located between two coarse cells, however, are calculated as the arithmetic average of the corresponding fine-grid coefficients; see Figure 1 for a more detailed description. As pointed out in [12], this direct discretization procedure is equivalent to the often used Galerkin approach, i.e., $L_{2h} = \frac{1}{2} R_h^{2h} L_h P_{2h}^h$, but computationally more efficient. The factor $1/2$ in the previous expression is due to the lack of consistency of the operator $R_h^{2h} L_h P_{2h}^h$ with the differential operator [13]. In the next result, we prove that both discretizations are indeed equivalent.

PROPOSITION 1. *Let L_h be the fine grid operator based on the cell-centered finite volume discretization of problem (5) on a uniform grid of mesh size $h = \ell/M$ with*




FIG. 1. Schematic representation of permeability upscaling used in the multigrid hierarchy (h - $2h$ - $4h$). (Top) Permeability values generated at cell centers (blue dots). (Bottom left) Permeability values at face centers (red dots) obtained from the harmonic mean (HM) of permeabilities from two adjacent cell centers. (Bottom middle) Permeability at face centers (bigger red dots) of $2h$ -grid is the arithmetic mean (AM) of permeabilities from face centers of the h -grid. (Bottom right) Permeability at face centers (biggest red dots) of $4h$ -grid is the arithmetic mean of permeabilities from face centers of the $2h$ -grid.

M even. Let P_{2h}^h be the piecewise constant prolongation operator and R_h^{2h} its adjoint. Then, the Galerkin coarse-grid operator $L_{2h} = \frac{1}{2} R_h^{2h} L_h P_{2h}^h$ is equivalent to a direct discretization on the coarse-grid based on the arithmetic average of the corresponding fine-grid coefficients.

Proof. We prove the equivalence for a coarse-grid cell $D_{2h}^{i_1, i_2}$ such that none of its edges lies on the boundary of the domain. The equivalence for coarse cells close to boundaries with Dirichlet or Neumann boundary conditions can be proven similarly. By applying the restriction operator R_h^{2h} in (7), the equation associated with the cell $D_{2h}^{i_1, i_2}$ by using the Galerkin approach is given by

$$(8) \quad \frac{1}{2} (R_h^{2h} L_h P_{2h}^h u)_{i_1, i_2} = \frac{1}{8} ((L_h P_{2h}^h u)_{2i_1, 2i_2} + (L_h P_{2h}^h u)_{2i_1-1, 2i_2} + (L_h P_{2h}^h u)_{2i_1, 2i_2-1} + (L_h P_{2h}^h u)_{2i_1-1, 2i_2-1}).$$

Taking into account that the prolongation operator is piecewise constant, we obtain

$$(L_h P_{2h}^h u)_{2i_1, 2i_2} = e_{2i_1, 2i_2}^h u_{i_1+1, i_2} + w_{2i_1, 2i_2}^h u_{i_1, i_2} + n_{2i_1, 2i_2}^h u_{i_1, i_2+1} + s_{2i_1, 2i_2}^h u_{i_1, i_2} + c_{2i_1, 2i_2}^h u_{i_1, i_2}$$

and similar expressions for the other terms in (8). By substituting these expressions in (8), the following discretization for the coarse cell $D_{2h}^{i_1, i_2}$ is obtained:

$$(9) \quad \frac{1}{2} (R_h^{2h} L_h P_{2h}^h u)_{i_1, i_2} = c_{i_1, i_2}^{2h} u_{i_1, i_2} + w_{i_1, i_2}^{2h} u_{i_1-1, i_2} + e_{i_1, i_2}^{2h} u_{i_1+1, i_2} + s_{i_1, i_2}^{2h} u_{i_1, i_2-1} + n_{i_1, i_2}^{2h} u_{i_1, i_2+1},$$

where

$$\begin{aligned} w_{i_1, i_2}^{2h} &= \frac{1}{8} (w_{2i_1-1, 2i_2}^h + w_{2i_1-1, 2i_2-1}^h), & e_{i_1, i_2}^{2h} &= \frac{1}{8} (e_{2i_1, 2i_2}^h + e_{2i_1, 2i_2-1}^h), \\ s_{i_1, i_2}^{2h} &= \frac{1}{8} (s_{2i_1, 2i_2-1}^h + s_{2i_1-1, 2i_2-1}^h), & n_{i_1, i_2}^{2h} &= \frac{1}{8} (n_{2i_1-1, 2i_2}^h + n_{2i_1, 2i_2}^h), \\ c_{i_1, i_2}^{2h} &= -(w_{i_1, i_2}^{2h} + e_{i_1, i_2}^{2h} + n_{i_1, i_2}^{2h} + s_{i_1, i_2}^{2h}). \end{aligned}$$

We observe that this scheme is equivalent to a direct discretization on the coarse grid where the diffusion coefficients on the edges are the arithmetic averages of the corresponding fine-grid coefficients. \square

Remark. In order to achieve a mesh-independent multigrid convergence following the analysis from [14], the following condition must be satisfied:

$$(10) \quad m_p + m_r > M_{pde},$$

where the orders m_p and m_r are the highest degree plus one of the polynomials that are exactly interpolated by P_{2h}^h and R_h^{2h} , respectively, and M_{pde} is the order of the PDE to be solved. For PDE (5), we have $M_{pde} = 2$, and for the considered operators (6) and (7), we get $m_p = m_r = 1$, which does not satisfy the inequality (10). In [15] it is shown, however, that this condition is not needed to prove uniform convergence [16].


2.3. Discussion about other multigrid methods for jumping coefficients.

In the context of *algebraic multigrid* methods for the numerical solution of PDEs, basically two prevailing methods have proved their use for multiple engineering problems, i.e., algebraic multigrid and aggregation-based multigrid methods [17, 18, 19, 20, 21]. These methods converge remarkably well, for example, for scalar PDEs with jumping coefficients. It is not always easily understood why these methods, and particularly the aggregation-based method, converge so well.

The origin of these algebraic methods may be found in the early days of multigrid, where a black-box multigrid with operator-dependent transfer operators (restriction and prolongation) and Galerkin coarse-grid operators for structured vertex-centered Cartesian grids was proposed in [22, 23, 24]. This can be seen as a predecessor of classical AMG, where these components were essentially enhanced by a flexible coarsening strategy.

The aggregation-based multigrid methods, with their origin in the work by Mandel [21] (smoothed aggregation), may be related to the CCMG methods as proposed in [25, 26]. In [26], it was shown that constant, i.e., operator-*independent*, transfer operators, in combination with Galerkin coarse-grid discretization, provided highly efficient multigrid results for cell-centered discretizations of elliptic PDEs that included jumping coefficients. These CCMG components were augmented with robust smoothing, like incomplete lower-upper decomposition (ILU) relaxation. The individual contributions of the coarse-grid correction and the smoothing parts were difficult to distinguish. Also, a CCMG based on coarsening by a factor of three together with operator-dependent interpolations was explored in [27].

3. Local Fourier analysis for variable coefficients. In this section we describe LFA in a setting which allows us to estimate the multigrid convergence factors for problems with jumping coefficients and problems with random fields. A discrete linear operator with constant coefficients, which is formally defined on an infinite grid, is usually assumed for carrying out a standard LFA. As we will show, this assumption can be relaxed by considering a discrete operator with constant coefficients *in*

FIG. 2. Infinite grid D_h divided into the corresponding subgrids for $n = 2$.

appropriate infinite subgrids. This allows us to generalize the analysis to problems for which the discrete operator consists of different stencils. A key point in this improved analysis is to consider a specific basis of the Fourier space, rather than the standard basis which is based on the Fourier modes. The use of this new basis will simplify the analysis.

We start from a regular infinite grid D_h with grid size h in both directions. Such an infinite grid will be *split into* $n \times n$ subgrids in the following way. First of all, a window comprising $n \times n$ cells of the original grid is *adequately* chosen, and, subsequently, we consider its periodic extension. The choice of the size of the $n \times n$ window is made such that the variability of the discrete operator in the computational grid can be appropriately represented, as will be explained by means of examples of different nature. Once n is fixed, the infinite subgrids are defined as follows (see Figure 2 for an example with $n = 2$),

$$(11) \quad D_h^{kl} = \{(k, l)h + (nk_1, nk_2)h \mid k_1, k_2 \in \mathbb{Z}\}, \quad k, l = 0, \dots, n-1.$$

For each low frequency, $\theta^{00} \in \Theta_{nh} = (-\pi/nh, \pi/nh]^2$, we introduce the grid-functions:

$$(12) \quad \psi_h^{kl}(\theta^{00}, \mathbf{x}) = \varphi_h(\theta^{00}, \mathbf{x}) \chi_{D_h^{kl}}(\mathbf{x}), \quad k, l = 0, \dots, n-1, \quad \mathbf{x} \in D_h,$$

where $\varphi_h(\theta^{00}, \mathbf{x}) = e^{i\theta^{00} \cdot \mathbf{x}}$ is the standard Fourier mode on D_h corresponding to the frequency θ^{00} . It is easy to see that the subspace generated by these n^2 grid-functions,

$$(13) \quad \mathcal{F}_h^{n^2}(\theta^{00}) = \text{span}\{\psi_h^{kl}(\theta^{00}, \cdot), \quad k, l = 0, \dots, n-1\},$$

is the same as the one spanned by the n^2 Fourier modes $\varphi_h(\theta_{kl}^{00}, \cdot)$ associated with the frequencies:

$$(14) \quad \theta_{kl}^{00} = \theta^{00} + (k, l) \frac{2\pi}{nh}, \quad k, l = 0, \dots, n-1.$$

In the case $n = 2$, the basis $\{\psi_h^{00}(\theta^{00}, \cdot), \psi_h^{11}(\theta^{00}, \cdot), \psi_h^{10}(\theta^{00}, \cdot), \psi_h^{01}(\theta^{00}, \cdot)\}$ is related to the standard basis of Fourier modes $\{\varphi_h(\theta^{00}, \cdot), \varphi_h(\theta^{11}, \cdot), \varphi_h(\theta^{10}, \cdot), \varphi_h(\theta^{01}, \cdot)\}$ in the following way:

$$(15) \quad \begin{pmatrix} \psi_h^{00}(\boldsymbol{\theta}^{00}, \cdot) \\ \psi_h^{11}(\boldsymbol{\theta}^{00}, \cdot) \\ \psi_h^{10}(\boldsymbol{\theta}^{00}, \cdot) \\ \psi_h^{01}(\boldsymbol{\theta}^{00}, \cdot) \end{pmatrix} = \mathcal{M} \begin{pmatrix} \varphi_h(\boldsymbol{\theta}^{00}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{11}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{10}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{01}, \cdot) \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & -1 & 1 & -1 \end{pmatrix} \begin{pmatrix} \varphi_h(\boldsymbol{\theta}^{00}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{11}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{10}, \cdot) \\ \varphi_h(\boldsymbol{\theta}^{01}, \cdot) \end{pmatrix}.$$

It is well-known that Fourier modes are eigenfunctions of any constant coefficient linear discrete operator L_h , that is, $L_h \varphi_h(\boldsymbol{\theta}, \mathbf{x}) = \tilde{L}_h(\boldsymbol{\theta}) \varphi_h(\boldsymbol{\theta}, \mathbf{x})$. Therefore, the representation of L_h with respect to a basis of n^2 Fourier modes is a diagonal matrix with diagonal elements $\tilde{L}_h(\boldsymbol{\theta}^{kl})$ with $k, l = 0, \dots, n-1$. In general, the Fourier representation with respect to the basis of functions $\{\psi_h^{kl}\}_{k,l=0}^{n-1}$ is a dense matrix. We will denote it by $\hat{L}_h(\boldsymbol{\theta}^{00})$.

If we consider the five-point standard discretization of the Laplace operator on a uniform grid of mesh size h ,

$$(16) \quad \frac{1}{h^2} \begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix},$$

its Fourier symbol with respect to the standard basis of Fourier modes is a diagonal matrix with diagonal elements equal to

$$\frac{1}{h^2} (4 - 2 \cos(\theta_x^{kl}) - 2 \cos(\theta_y^{kl}))$$

(see [28], for instance), whereas the Fourier representation with respect to the new basis in the case $n = 2$ is given by

$$(17) \quad \hat{L}_h(\boldsymbol{\theta}^{00}) = \frac{2}{h^2} \begin{pmatrix} 2 & 0 & \cos(\theta_x^{00}) & \cos(\theta_y^{00}) \\ 0 & 2 & \cos(\theta_y^{00}) & \cos(\theta_x^{00}) \\ \cos(\theta_x^{00}) & \cos(\theta_y^{00}) & 2 & 0 \\ \cos(\theta_y^{00}) & \cos(\theta_x^{00}) & 0 & 2 \end{pmatrix} \text{ with } \boldsymbol{\theta}^{00} = (\theta_x^{00}, \theta_y^{00}).$$

Notice that, for example, the first row of the previous symbol is obtained by looking at the decomposition of the stencil (16) into the connections among the unknowns located at the different subgrids D_h^{kl} defined in (11). In particular, following the notation in Figure 2, the $\bullet-\bullet$, $\bullet-\square$ and $\bullet-\circ$ connections are given by the stencils

$$\frac{1}{h^2} [4], \quad \frac{1}{h^2} \begin{bmatrix} & -1 & \\ -1 & \bullet & -1 \end{bmatrix}, \quad \frac{1}{h^2} \begin{bmatrix} -1 \\ \bullet \\ -1 \end{bmatrix},$$

giving rise to the symbols

$$\frac{4}{h^2}, \quad \frac{2}{h^2} \cos(\theta_x^{00}), \quad \frac{2}{h^2} \cos(\theta_y^{00}),$$

which appear in the first row of (17), whereas there is no $\bullet-\times$ connection. The rest of the rows are analogously computed.

The procedure to obtain the Fourier symbol of a smoothing operator S_h , which is based on a splitting of the discrete operator $L_h = L_h^+ + L_h^-$, is analogous with the new basis. The smoothing iteration is given by

$$L_h^+ \bar{w}_h + L_h^- w_h = f_h$$

with w_h the approximation of the solution before the smoothing step and $\overline{w_h}$ the approximation after the smoothing step. By computing the symbols of L_h^+ and L_h^- as before, the Fourier symbol of the smoothing operator is given by

$$(18) \quad \widehat{S}_h(\boldsymbol{\theta}^{00}) = -(\widehat{L}_h^+)^{-1}(\boldsymbol{\theta}^{00})\widehat{L}_h^-(\boldsymbol{\theta}^{00}).$$

The Fourier symbol corresponding to a lexicographic Gauss–Seidel iteration for the five-point standard discretization of the Laplace operator on a uniform grid of mesh size h , in the case $n = 2$, is as in (18), where

$$(19) \quad \widehat{L}_h^+(\boldsymbol{\theta}^{00}) = \frac{2}{h^2} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 2 & \cos(\theta_y^{00}) & \cos(\theta_x^{00}) \\ \cos(\theta_x^{00}) & 0 & 2 & 0 \\ \cos(\theta_y^{00}) & 0 & 0 & 2 \end{pmatrix}, \quad \widehat{L}_h^-(\boldsymbol{\theta}^{00}) = \widehat{L}_h(\boldsymbol{\theta}^{00}) - \widehat{L}_h^+(\boldsymbol{\theta}^{00}).$$

Once the Fourier representation of the smoothing operator with respect to the new basis is obtained, we can define the smoothing factor by using the change of basis matrix. For example, for the case $n = 2$, the smoothing factor is obtained by

$$\mu(S_h) = \sup_{\boldsymbol{\theta}^{00} \in \Theta_{2h}} \rho(\mathcal{Q}_h \mathcal{M} \widehat{S}_h(\boldsymbol{\theta}^{00}) \mathcal{M}^{-1}),$$

where \mathcal{Q}_h is the projection operator onto the space of high-frequency components and \mathcal{M} is the change of basis matrix given in (15).

3.1. LFA formulations for cell-centered grids. In section 2, we didn't distinguish between cell- and vertex-centered grids. The generalized LFA indeed works well for both types of discretization. By introducing the coarse grids and their relation with the fine grids, we need to fix the approach of interest. Since here we will focus on cell-centered discretizations, from now on the description of the analysis will be given for this case, although it may be applied to the vertex-centered case in a similar way by defining appropriately the coarse meshes.

According to the location of the coarse-grid points in a regular cell-centered grid, we define for a fixed n the following infinite coarse subgrids of D_{2h} :

$$(20) \quad D_{2h}^{kl} = \{(h/2, h/2) + (k, l)2h + (nk_1, nk_2)h \mid k_1, k_2 \in \mathbb{Z}\}, \quad k, l = 0, \dots, n/2 - 1.$$

Due to the relation between the grid-functions of the new Fourier basis given in (12) and the standard Fourier modes, it can be shown that the coarse-grid correction operator $C_h = I_h - P_{2h}^h L_{2h}^{-1} R_{2h}^{2h} L_h$, where P_{2h}^h and R_{2h}^{2h} are the prolongation and restriction operators, L_h and L_{2h} are the fine- and coarse-grid operators, and I_h is the identity, satisfies the following invariance property:

$$C_h : \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00}) \rightarrow \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00}).$$

More concretely, for $\boldsymbol{\theta}^{00} \in \Theta_{2nh} = (-\pi/2nh, \pi/2nh]^2$, the following properties of the operators in C_h are fulfilled:

1. $L_h, I_h : \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00}) \rightarrow \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00})$,
2. $L_{2h} : \mathcal{F}_{2h}^{n^2/4}(2\boldsymbol{\theta}^{00}) \rightarrow \mathcal{F}_{2h}^{n^2/4}(2\boldsymbol{\theta}^{00})$,
3. $R_{2h}^{2h} : \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00}) \rightarrow \mathcal{F}_{2h}^{n^2/4}(2\boldsymbol{\theta}^{00})$,
4. $P_{2h}^h : \mathcal{F}_{2h}^{n^2/4}(2\boldsymbol{\theta}^{00}) \rightarrow \mathcal{F}_h^{n^2}(\boldsymbol{\theta}^{00})$.

From these invariance properties we can compute the Fourier representations of the corresponding operators. As an example, we will describe next the representation of R_h^{2h} with respect to the grid-functions $\{\psi_{2h}^{kl}(2\theta^{00})\}_{k,l=0}^{n/2-1}$ and $\{\psi_h^{kl}(\theta^{00})\}_{k,l=0}^{n-1}$, for the restriction operators considered in this work.

We first consider the basic restriction operator obtained as the adjoint of the piecewise constant prolongation operator with stencil form (7). Its Fourier representation with respect to the new Fourier basis is given by

$$\widehat{R}_h^{2h}(\theta^{00}) = \frac{1}{4} \begin{pmatrix} e^{-i(\theta_x^{00} + \theta_y^{00})/2} & e^{i(\theta_x^{00} + \theta_y^{00})/2} & e^{i(\theta_x^{00} - \theta_y^{00})/2} & e^{i(\theta_y^{00} - \theta_x^{00})/2} \end{pmatrix}.$$

In the case of the cell-centered restriction operator by Khalil and Wesseling [25], that is,

$$(21) \quad R_h^{2h} = \frac{1}{16} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 3 & 2 & 0 \\ 0 & 2 & 3 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}_h^{2h},$$

the Fourier representation is given by

$$\widehat{R}_h^{2h}(\theta^{00}) = \frac{1}{16} \begin{pmatrix} e^{-i(\theta_x^{00} + \theta_y^{00})/2} (2 + e^{2i\theta_y^{00}} + e^{2i\theta_x^{00}}) \\ e^{i(\theta_x^{00} + \theta_y^{00})/2} (2 + e^{-2i\theta_y^{00}} + e^{-2i\theta_x^{00}}) \\ e^{i(\theta_x^{00} - \theta_y^{00})/2} (3 + e^{2i(\theta_y^{00} - \theta_x^{00})}) \\ e^{i(\theta_y^{00} - \theta_x^{00})/2} (3 + e^{2i(\theta_x^{00} - \theta_y^{00})}) \end{pmatrix}^T.$$

As an immediate consequence of these invariance properties and the invariance property of the smoothing operator, also the two-grid operator $K_h^{2h} = C_h S_h^\nu$, where ν denotes the number of smoothing steps, leaves the subspaces $\mathcal{F}_h^{n^2}(\theta^{00})$ invariant. Its Fourier representation is given by

$$\widehat{K}_h^{2h}(\theta^{00}) = \widehat{C}_h(\theta^{00}) \widehat{S}_h^\nu(\theta^{00}) = (\widehat{I}_h(\theta^{00}) - \widehat{P}_{2h}^h(\theta^{00}) \widehat{L}_{2h}^{-1}(\theta^{00}) \widehat{R}_h^{2h}(\theta^{00}) \widehat{L}_h(\theta^{00})) \widehat{S}_h^\nu(\theta^{00}).$$

Finally, we can compute the asymptotic two-grid convergence factor as the supremum of the spectral radii of $(n^2 \times n^2)$ -matrices, as follows:

$$\rho(K_h^{2h}) = \sup_{\theta^{00} \in \widetilde{\Theta}_{2nh}} \rho(\widehat{K}_h^{2h}(\theta^{00})),$$

where $\widetilde{\Theta}_{2nh}$ is the subset of Θ_{2nh} in which we remove the frequencies θ^{00} such that the determinant of the Fourier symbol of L_h or L_{2h} vanishes.

4. LFA results for PDEs with jumping coefficients. In this section, we apply the proposed LFA to predict the two-grid convergence factors for a collection of benchmark problems with jumping coefficients taken from the literature [22, 29, 30]. The test cases cover a variety of possible inhomogeneities including jumps that are not aligned with the coarse grid. In all these problems, (1) is numerically solved in the domain $D = (0, 1)^2$ by using a mesh of grid-size $h = 1/128$. In particular, the following jumping coefficient benchmark problems, characterized by the distribution of the diffusion coefficient, are considered here:

1. Vertical jump. Function $k(x, y)$ is defined in the following way (see also Figure 3(a)):

$$k(x, y) = \begin{cases} 1 & \text{if } x < \frac{1}{2} + h, \\ 10^3 & \text{if } x \geq \frac{1}{2} + h. \end{cases}$$

2. Four corner problem. The domain is divided into four regions in which the diffusion coefficient is varying; see Figure 3(b). In particular,

$$k(x, y) = \begin{cases} 10^4 & \text{if } (x, y) \in (0, \frac{1}{2} + h)^2, \\ 1 & \text{if } (x, y) \in (0, \frac{1}{2} + h) \times (\frac{1}{2} + h, 1), \\ 10^{-2} & \text{if } (x, y) \in (\frac{1}{2} + h, 1) \times (\frac{1}{2} + h, 1), \\ 10^{-4} & \text{otherwise.} \end{cases}$$

3. Square inclusion. In this example we assume a square inhomogeneity in one cell within the square domain; see Figure 3(c). The diffusion coefficient is defined as

$$k(x, y) = \begin{cases} k_0 & \text{if } (x, y) \in (\frac{1}{2} - h, \frac{1}{2})^2, \\ 1 & \text{otherwise,} \end{cases}$$

where values $k_0 = 10^4$ and $k_0 = 10^{-4}$ are considered.

4. Periodic square inclusions. This test is taken from [31]. We consider a structured pattern of square inclusions of size $2h \times 2h$ as depicted in Figure 3(d). The diffusion parameter is $k(x, y) = 1$ inside the dark region and $k(x, y) = 1000$ inside the white region.
5. Periodic L-shaped inclusions. In the last test case, we consider a structured pattern of L-shaped inclusions as in Figure 3(e). The diffusion parameter is $k(x, y) = 10^4$ inside the white region and $k(x, y) = 1$ inside the dark region.

To perform the theoretical analysis, the periodic extension of a window of size 8×8 has been chosen, where the diffusion coefficient is prescribed in such a window according to its definition; see Figure 3 (right side). In all numerical tests a random initial guess is chosen, and the right-hand side and boundary conditions are set to zero to be able to determine asymptotic convergence factors. In this way, we avoid round-off errors, permitting us to perform as many iterations as needed. In practice, we have seen that 50 iterations are sufficient.

Next, we show the excellent correspondence between the theoretical analysis and the experimental results for these test cases. Two combinations of intergrid transfer operators are considered. The first combination, denoted here by (CP,CR), is based on the use of piecewise constant prolongation (6) and its adjoint as the restriction (7). In the second combination we change to a higher polynomial order interpolation operator which is the adjoint to the Wesseling–Khalil restriction (21). This choice is denoted by (WP,CR). Moreover, a standard damped Jacobi smoother (damping with $\omega = 0.8$) is considered as well as the proposed lexicographic Gauss–Seidel smoother.

In Table 1, for different numbers of smoothing steps, for two different smoothers, and for the two combinations of restriction and prolongation operators, we provide the two-grid convergence factors predicted by the novel LFA for each of the proposed numerical experiments. We also display in parentheses the average value after 50 iterations of the experimentally computed convergence factors by using two-grid multigrid cycles. For all these cases, we observe a very accurate match between the analysis results and the rates experimentally obtained. Regarding the size of the window to perform the LFA, we have observed that a window of size 8×8 is enough to

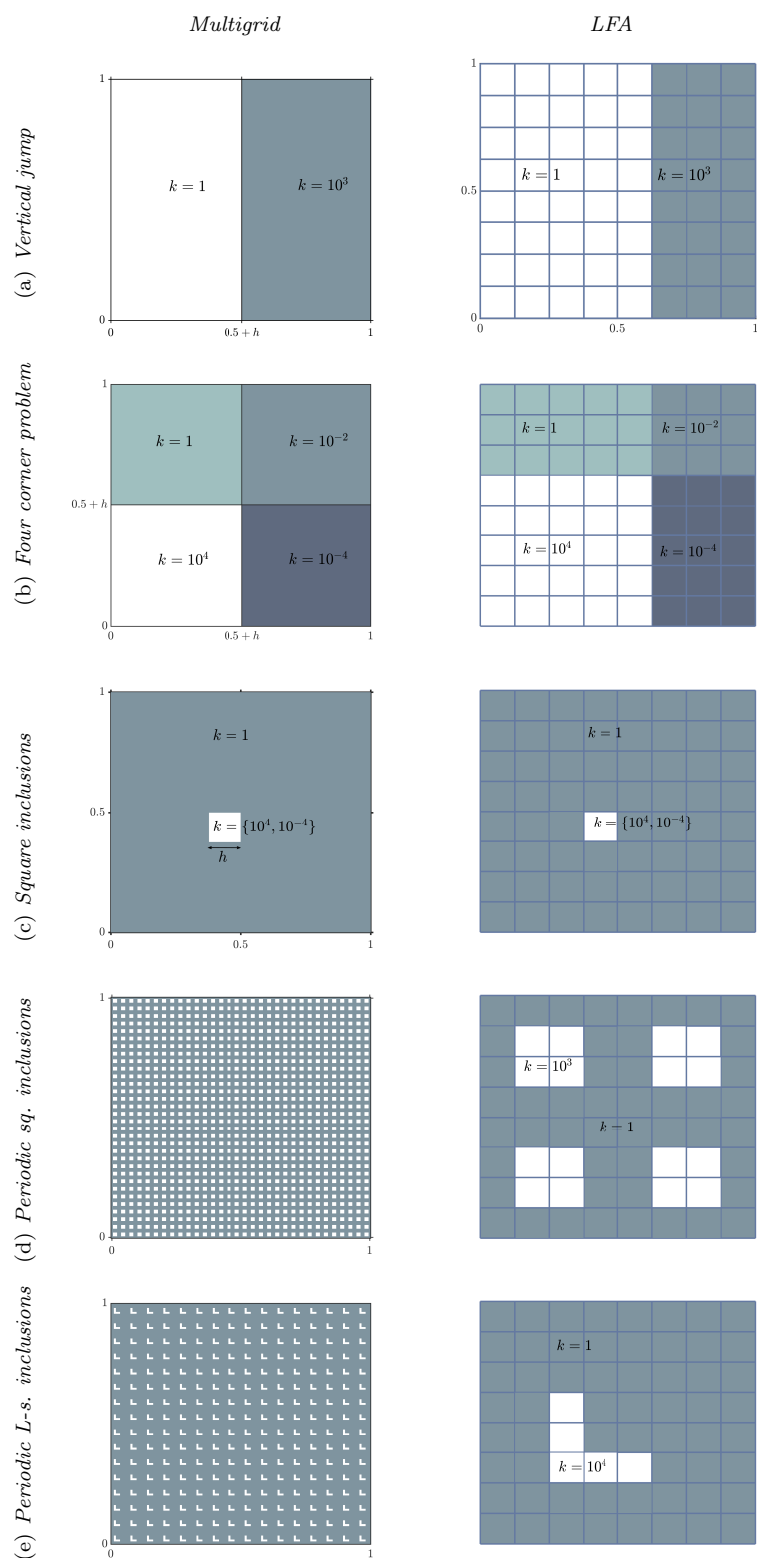


FIG. 3. Distribution of the diffusion coefficients for the five considered examples on a unit square domain and corresponding 8×8 window used in the LFA. L-s. = L-shaped.

TABLE 1

Asymptotic two-grid convergence factors predicted by LFA and the corresponding computed average multigrid convergence factors (in parentheses) using two-grid cycles and different pre- and postsmoothing steps for the five examples. *L-s.* = *L-shaped*; *GS* = *Gauss-Seidel*.

			Pre-, postsmoothing steps		
			(1, 0)	(1, 1)	(2, 2)
Vertical jump	GS	(CP,CR)	0.42(0.42)	0.18(0.19)	0.04(0.03)
		(WP,CR)	0.41(0.37)	0.19(0.16)	0.07(0.11)
	Jacobi	(CP,CR)	0.65(0.65)	0.43(0.42)	0.19(0.19)
		(WP,CR)	0.63(0.59)	0.40(0.35)	0.19(0.19)
Four corner problem	GS	(CP,CR)	0.42(0.37)	0.15(0.12)	0.04(0.03)
		(WP,CR)	0.40(0.39)	0.16(0.16)	0.09(0.09)
	Jacobi	(CP,CR)	0.63(0.61)	0.40(0.39)	0.16(0.16)
		(WP,CR)	0.62(0.62)	0.40(0.40)	0.19(0.18)
Sq. inclusion ($k = 10^4$)	GS	(CP,CR)	0.45(0.44)	0.21(0.19)	0.04(0.04)
		(WP,CR)	0.41(0.40)	0.18(0.17)	0.11(0.11)
	Jacobi	(CP,CR)	0.60(0.65)	0.36(0.42)	0.13(0.19)
		(WP,CR)	0.60(0.62)	0.38(0.37)	0.22(0.21)
Sq. inclusion ($k = 10^{-4}$)	GS	(CP,CR)	0.46(0.45)	0.21(0.20)	0.05(0.05)
		(WP,CR)	0.41(0.40)	0.19(0.19)	0.12(0.12)
	Jacobi	(CP,CR)	0.61(0.65)	0.38(0.42)	0.15(0.19)
		(WP,CR)	0.61(0.59)	0.39(0.39)	0.23(0.23)
Periodic Sq. inclusions	GS	(CP,CR)	0.64(0.61)	0.43(0.42)	0.41(0.41)
		(WP,CR)	0.62(0.61)	0.43(0.41)	0.41(0.40)
	Jacobi	(CP,CR)	0.81(0.78)	0.66(0.65)	0.44(0.46)
		(WP,CR)	0.81(0.80)	0.66(0.65)	0.44(0.43)
Periodic L-s. inclusions	GS	(CP,CR)	0.50(0.50)	0.32(0.26)	0.21(0.21)
		(WP,CR)	0.54(0.53)	0.40(0.40)	0.30(0.30)
	Jacobi	(CP,CR)	0.71(0.63)	0.54(0.48)	0.36(0.35)
		(WP,CR)	0.71(0.71)	0.56(0.56)	0.42(0.42)

achieve excellent predictions in all considered benchmark problems. For example, for the vertical jump test, the two-grid analysis considering four smoothing steps of the Gauss-Seidel smoother and the combination (CP,CR) of intergrid transfer operators provides a factor of 0.11 when a 2×2 window is used, a factor of 0.06 for a 4×4 window, and a factor of 0.04, which matches perfectly the real convergence, when the 8×8 window is considered.

We remark that with the current multigrid approach, the quality of the coarse-grid discretization may not be satisfactory, for example, for the chessboard or L-shaped inclusion example. For such cases, we recommend either using more *powerful* smoothers such as the ILU smoother or adapting the coarsening. Furthermore, for PDEs with strong local variations in the coefficient fields, *homogenization techniques* [22, 29, 30] may also be used to obtain the coarser representation of the fine-grid problem. Comparing the results of the two combinations of intergrid transfer operators, we observe a very similar performance for all five test cases studied here. In the rest of the paper, we therefore choose the strategy (CP,CR) because of its simplicity and low computational cost.

5. LFA results for PDEs with random coefficients. Here, we consider the stochastic PDE (5) defined on a unit square domain $D = (0, 1)^2$ with homogeneous Dirichlet boundary conditions. Two different types of diffusion coefficients based on random jumps and lognormal random fields are studied. The randomly jumping coefficient problem can be seen as a transition from the deterministic to a stochastic setting.




FIG. 4. An example of random realization of $U = \ln k$ with $m = 5$ on a unit square domain.

5.1. Randomly jumping coefficients. To simulate random jumps, the domain D is subdivided into square-blocks of size $[\frac{1}{8} \times \frac{1}{8}]$ and the value of the coefficients on each of the blocks is sampled as

$$(22) \quad k = e^U \quad \text{with} \quad U \sim \mathcal{U}\{-m, m\} \quad \text{and} \quad m \in \mathbb{Z}.$$

In other words, U is an independent identically distributed (i.i.d.) integer sampled from a discrete uniform distribution $\mathcal{U}\{-m, m\}$. Here, the integer m defines the order of magnitude of the jumps; an example for $m = 5$ is shown in Figure 4. Notice that for this choice of m , we may encounter interfaces with maximum jumps of magnitude equal to e^{10} .

For each random realization of the jumping coefficient field, we compare the LFA two-grid convergence factors with the computed asymptotic convergence factors of the multigrid method by using W-cycles. To perform the LFA, we again use a window of size 8×8 capturing the heterogeneity of the whole domain. Furthermore, the LFA equivalent of the randomly jumping coefficient problem is similar to the four corner problem in Figure 3(d) with the magnitude of each block given by (22) and the cross-point exactly at the center of the LFA block. Regarding the multigrid components, a lexicographic Gauss–Seidel iteration is employed as the smoother, and the simplest combination (CP,CR) of intergrid transfer operators is chosen. Also, we use a 4×4 grid as the coarsest in the multigrid hierarchy. To determine the asymptotic convergence factors of the multigrid method, the right-hand side is again set to zero.

The experimental convergence factor of the multigrid method for the i th realization of the random field is then computed, as follows:

$$(23) \quad \rho_i = \left\{ \frac{\|\text{res}^{k_i}\|_\infty}{\|\text{res}^0\|_\infty} \right\}^{1/k_i} \quad \text{for} \quad i = 1, 2, \dots, N_{MG},$$

where $\|\text{res}^0\|_\infty$ is the infinity norm of the residual obtained from an initial solution and $\|\text{res}^{k_i}\|_\infty$ is the residual after k_i iterations of the multigrid cycle. We use these quantities to calculate the average and the standard deviation of the asymptotic convergence factors:

$$(24) \quad \langle \rho \rangle_{MG} = \frac{1}{N_{MG}} \sum_{i=1}^{N_{MG}} \rho_i \quad \sigma_{MG} = \sqrt{\frac{1}{(N_{MG} - 1)} \sum_{i=1}^{N_{MG}} (\rho_i - \langle \rho \rangle_{MG})^2},$$




FIG. 5. Comparison of the mean and the standard deviation of the LFA (dashed line) and multigrid (solid line) convergence factors for different W -cycling strategies for randomly jumping coefficients with $m = 2$ (left) and $m = 5$ (right).

respectively. These averaged quantities are defined similarly for the LFA results (based on LFA two-grid factors) and are denoted, respectively, by $\langle \rho \rangle_{LFA}$ and σ_{LFA} .

In Figure 5, we show the comparison, the mean \pm standard deviation, of the LFA prediction and the multigrid convergence for jump parameter $m = 2$ (left) and $m = 5$ (right) computed using $N_{LFA} = N_{MG} = 100$. For clarity, the LFA predicted mean and standard deviation are plotted at an arbitrary location ($h = 1/200$) and the mean is extended throughout the x -axis. Overall, a good match between the LFA and multigrid convergence is seen up to one decimal place. We also observe that for this specific jumping coefficient problem, there is no further improvement with an increase in the number of smoothing steps after the $W(2, 2)$ -cycle.

5.2. Lognormal random fields. Next, we test the LFA prediction capability for a more realistic lognormal diffusion problem. Lognormality leads to positive permeability throughout the domain. The logarithm of the permeability field, $Z = \log k$, is modeled by a zero-mean Gaussian random field $Z : \Omega \times \overline{D} \rightarrow \mathbb{R}$. A simplification is the use of a homogeneous covariance function $C_\Phi : \mathbb{R} \rightarrow \mathbb{R}$, so that

$$(25) \quad \text{Cov}(Z(\mathbf{x}, \cdot), Z(\mathbf{y}, \cdot)) = C_\Phi(r) \text{ with } r = \|\mathbf{x} - \mathbf{y}\|_2.$$

By the so-called Matérn family of covariance functions, random coefficient fields with different degrees of smoothness can be generated. The Matérn covariance function [32] is characterized by a parameter set $\Phi = (\nu_c, \lambda_c, \sigma_c^2)$, as follows:

$$(26) \quad C_\Phi(r) = \sigma_c^2 \frac{2^{1-\nu_c}}{\Gamma(\nu_c)} \left(2\sqrt{\nu_c} \frac{r}{\lambda_c} \right)^{\nu_c} K_\nu \left(2\sqrt{\nu_c} \frac{r}{\lambda_c} \right).$$

Here, Γ is the gamma function and K_ν the modified Bessel function of the second kind. The different parameters have different roles regarding the field's randomness. Parameter ν_c defines the field's smoothness, σ_c^2 represents its variance, and λ_c is the correlation length of the covariance function. Moreover, parameters λ_c and σ_c^2 prescribe the number of peaks and the amplitude of the random field, respectively. When the smoothness parameter is $\nu_c = 1/2$, the Matérn function corresponds to an exponential model, whereas when $\nu_c \rightarrow \infty$ it represents a Gaussian model.

Realizations of the random field are “almost surely” Hölder continuous, $Z, k \in C^\eta(D)$, with $0 < \eta < \nu_c$ (see, e.g., [33, 34]). For (5) regularity results have been obtained by taking into account the regularity of the lognormal coefficient field; see [35, 36, 34].

Remark. The random fields are generated using the circulant embedding method [37, 38]. In Appendix A, we briefly describe this sampling algorithm. This method employs the fast Fourier transform (FFT) for the covariance matrix decomposition and, therefore, requires $\mathcal{O}(h^{-2} \log h^{-2})$ operations to generate one sample of the random field on a uniform grid mesh size h . Other techniques, e.g., the Karhunen–Loève (KL) expansion, can also be utilized to sample these fields. The benefit of the circulant embedding technique is that it yields an exact representation of the random field on the sampling mesh for all Φ , whereas the KL expansion gives a low-dimensional representation of the field introducing a bias (due to truncation after a finite number of eigenmodes). The number of terms in the KL expansion is dependent on Φ [39, 40] and can be large for stochastic processes with small correlation lengths.

To illustrate the performance of LFA for PDEs with random parameters, we consider four Matérn reference parameter sets Φ with increasing order of complexity, listed in Table 2. Random fields generated with these Matérn parameter sets, sampled on a uniform mesh, are presented in Figure 6.

TABLE 2

Different combinations of the Matérn reference parameters $\Phi = (\nu_c, \lambda_c, \sigma_c^2)$ with increasing complexity from left to right.

Φ_1	Φ_2	Φ_3	Φ_4
(1.5, 0.3, 1)	(0.5, 0.3, 1)	(1.5, 0.1, 3)	(0.5, 0.1, 3)




FIG. 6. *Logarithm of the permeability field, $\log_{10} k$, generated using four reference parameter sets $\Phi = (\nu_c, \lambda_c, \sigma_c^2)$.*

TABLE 3

Comparison of the average and the standard deviation of the LFA and multigrid convergence factors. Different numbers of smoothing steps, $h = 1/64$, and the four reference parameter sets in Table 2 are considered. The coarsest grid considered in the multigrid method is a 4×4 mesh.

		Φ_1		Φ_2		Φ_3		Φ_4	
		MG	LFA	MG	LFA	MG	LFA	MG	LFA
$W(1, 1)$	$\langle \rho \rangle$	0.20	0.20	0.20	0.20	0.19	0.20	0.23	0.21
	σ	0.004	0.002	0.004	0.005	0.005	0.01	0.04	0.02
$W(2, 2)$	$\langle \rho \rangle$	0.04	0.04	0.04	0.04	0.07	0.06	0.17	0.08
	σ	0.001	0.001	0.002	0.003	0.03	0.03	0.04	0.03
$W(3, 3)$	$\langle \rho \rangle$	0.01	0.02	0.02	0.02	0.05	0.03	0.13	0.06
	σ	0.001	0.002	0.002	0.004	0.02	0.02	0.03	0.03

For a fixed Φ and mesh size h , we generate N_{MG} realizations of the permeability field. For each of these random fields, we can compare the LFA two-grid convergence factors with the computed asymptotic convergence factors of the multigrid method by using the same components that were utilized for randomly jumping coefficient fields. The comparison is shown in Table 3, where different numbers of smoothing steps are considered and the comparison is done for the four reference parameter sets, described in Table 2. For all experiments, we set $h = 1/64$ and $N_{MG} = N_{LFA} = 100$. In general, we observe a good agreement between the experimental and the LFA quantities. For the first three sets of parameters excellent convergence factors are obtained already with two pre- and two postsmoothing steps. When the more difficult set of parameters is considered, however, more smoothing steps may be necessary to obtain a good convergence factor. It is also pointed out that the performance of $W(1,1)$ -cycle is well predicted by LFA for all Φ_i , $i = 1, 4$. A slight discrepancy is observed for the $W(2,2)$ - and $W(3,3)$ -cycles in the case of Φ_4 .

5.3. Mesh dependency. Next, we wish to study the influence of the size of the sampling mesh on the multigrid convergence. For this purpose, we choose two representative parameter sets describing a smooth and a highly oscillating random field that are generated by the parameter sets Φ_2 and Φ_4 , respectively. Figure 7 shows the average convergence factors for Φ_2 (left side) and Φ_4 (right side), predicted by LFA (top) and experimentally observed multigrid convergence (bottom), for different mesh sizes and different numbers of smoothing steps. For Φ_2 , the average reduction factor is roughly the same, independent of the size of the mesh. These predictions coincide well with the experimentally observed factors. For the results corresponding to parameter set Φ_4 , we observe robustness of the method when the mesh is sufficiently fine, which is also confirmed by the multigrid experiments. In the same figure, the standard deviation is presented, which decreases when $h \rightarrow 0$. The LFA predictions on the coarser grids are less reliable when compared to the multigrid convergence rates. We would like to mention that a three-grid LFA [41] is not helpful here. On coarse grids, boundary conditions have an impact on the method's convergence, but they are not taken into account in the analysis. Moreover, to analyze the robustness of the proposed multigrid method, in the next experiment we fix the field's smoothness parameter as $\nu_c = 0.5$ and vary the other parameters $\sigma_c^2 \in [0.5, 5]$ and $\lambda_c \in [0.05, 0.5]$. In Figure 8 we show the average LFA two-grid convergence factors when two smoothing steps are considered for $h = 1/32$ (left) and $h = 1/64$ (right). The multigrid convergence is very satisfactory for all combinations of the parameters and increases slightly when λ_c tends to be small and σ_c^2 becomes large. This case, however, represents a rather extreme situation in which the jumps in the permeability field are of more than

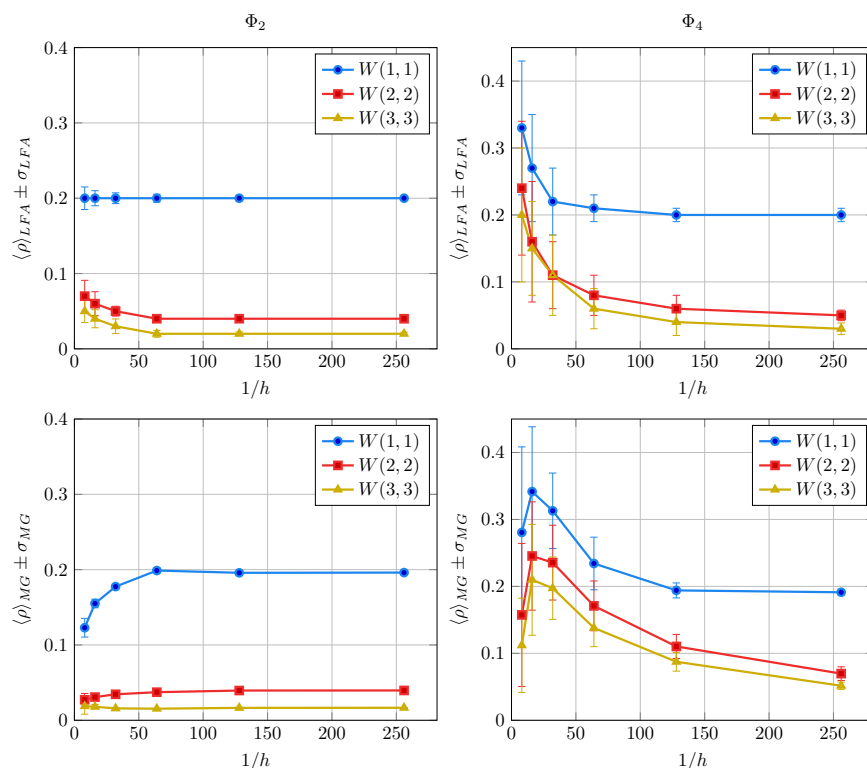


FIG. 7. Average reduction LFA two-grid factors (top) and asymptotic multigrid convergence (bottom) along with the standard deviation for two reference parameter sets Φ_2 (left column) and Φ_4 (right column); $N_{MG} = N_{LFA} = 100$.




FIG. 8. Contour of average LFA two-grid convergence factors, $\langle \rho \rangle_{LFA}$, with two smoothing steps for different λ_c and σ_c^2 , and for fixed $\nu_c = 0.5$ with $h = 1/32$ (left) and $h = 1/64$ (right).

15 orders of magnitude. Again, we see an improvement in the convergence rate with grid refinement. Note that with the considered range of parameters we cover all realistic cases.

6. Multilevel Monte Carlo computations. In the preceding section, we demonstrated that LFA performs well in predicting the mean convergence rate of the multigrid method for a given Matérn parameter set. In this section, we wish to utilize

LFA-based information to select the multigrid cycle for MLMC simulations. Within MLMC methods, a *hierarchy of grids*, $\{D_{h_\ell}\}_{\ell=0}^L$, characterized by $h_0 > \dots > h_L$ is defined to accurately estimate the quantity of interest Q , based on the linearity of the expectation operator, i.e.,

$$(27) \quad \mathbb{E}[Q_{h_L}] = \mathbb{E}[Q_{h_0}] + \sum_{\ell=1}^L \mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}].$$

The advantage of using the above telescopic decomposition is that on the coarsest grid, for $\ell = 0$, large numbers of samples to accurately determine $\mathbb{E}[Q_{h_0}]$ can be computed at a small computational effort. For larger values of ℓ the numerical solution is comparatively expensive; however, fewer samples are required on these finer levels, as the variance of the correction terms, $\mathbb{V}[Q_{h_\ell} - Q_{h_{\ell-1}}]$, is significantly smaller compared to the variance $\mathbb{V}[Q_{h_\ell}]$. A multilevel estimator for $\mathbb{E}[Q_{h_L}]$ can be defined in terms of the standard MC estimator as

$$(28) \quad \mathcal{E}_L^{ML}[Q_{h_L}] := \frac{1}{N_0} \sum_{i=1}^{N_0} Q_{h_0}(\omega_i) + \sum_{\ell=1}^L \left(\frac{1}{N_\ell} \sum_{i=1}^{N_\ell} (Q_{h_\ell}(\omega_i) - Q_{h_{\ell-1}}(\omega_i)) \right),$$

where the level dependent samples $N_\ell \in \mathbb{N}$ form a decreasing sequence for increasing ℓ -values. We will use the same geometric sequence of meshes in the MLMC estimator as employed for the multigrid method described in subsection 2.2, based on uniform coarsening such that $h_\ell = h_{\ell-1}/2 = 2^{-\ell}h_0$, with h_0 the coarsest mesh size. As each of the estimators in (28) is independent, the variance of the multilevel estimator is the sum of the variances of individual estimators, i.e.,

$$(29) \quad \mathbb{V}[\mathcal{E}_L^{ML}[Q_{h_L}]] = \sum_{\ell=0}^L \frac{\mathcal{V}_\ell}{N_\ell},$$

where the level dependent sample variance, $\mathcal{V}_\ell = \mathbb{V}[Q_{h_\ell} - Q_{h_{\ell-1}}] = \mathcal{O}(2^{-\beta\ell})$, $\beta > 0$, and $Q_{h_{-1}} = 0$. Assume that the mean finite volume error for mesh width h_ℓ satisfies

$$(30) \quad |\mathbb{E}[Q_{h_\ell} - Q]| = \mathcal{O}(h_\ell^\alpha)$$

with $\alpha > 0$. The mean square error of the MLMC estimator can be expressed as a sum of bias and sampling errors,

$$(31) \quad \begin{aligned} \mathbb{E}[(\mathcal{E}_L^{ML}[Q_{h_L}] - \mathbb{E}[Q])^2] &= (\mathbb{E}[Q_{h_L} - Q])^2 + \mathbb{V}[\mathcal{E}_L^{ML}[Q_{h_L}]], \\ &= c_0 h_L^{2\alpha} + \sum_{\ell=0}^L \frac{\mathcal{V}_\ell}{N_\ell}, \end{aligned}$$

with c_0 a constant independent of h_L and α . The level dependent sample size N_ℓ can be chosen such that the sampling errors of each of the MC estimators in (28) reduces to the size of the discretization error. This is achieved by taking N_ℓ such that $\mathcal{V}_\ell/N_\ell = \mathcal{O}(h_L^{2\alpha}) = \mathcal{O}(2^{-2\alpha L})$. Moreover, by fixing the number of the finest level samples N_L , the number of coarse-grid samples can be simply expressed as

$$(32) \quad N_\ell = \lceil N_L 2^{\beta(L-\ell)} \rceil.$$

In practice the value N_L is small $\sim \mathcal{O}(1)$ and can be chosen heuristically [42, 43]. The sampling error on the coarsest level does not depend on β , and therefore, by using (32), the balance $\mathcal{V}_0/N_0 = \mathcal{O}(h_L^{2\alpha})$ holds only asymptotically for large L . Alternatively, one can also solve an optimization problem, as in [8, 9, 10], which minimizes the total cost of the MLMC estimator for a given tolerance ε .

The computational cost of the MLMC estimator is given by

$$(33) \quad \mathcal{W}_L^{ML} = \sum_{\ell=0}^L N_\ell \cdot (\mathcal{W}_\ell + \mathcal{W}_{\ell-1}),$$

where \mathcal{W}_ℓ represents total number of arithmetic operations to obtain a sample of the quantity of interest Q on grid level ℓ . This typically includes the cost of generating the random permeability field, assembling and solving the linear system of equations, and postprocessing. Of these, the cost of solving the linear system is usually dominant, and, therefore, \mathcal{W}_ℓ will only represent the cost of the multigrid method. For the current combination of multigrid components, it is straightforward to show that the total amount of computational work for one multigrid cycle is proportional to the number of grid points on the finest multigrid level [28]. The cost of the multigrid solver can thus be expressed as

$$(34) \quad \mathcal{W}_\ell = \bar{c} \nu k_\ell h_\ell^{-2},$$

where \bar{c} is a constant, $k_\ell \in \mathbb{N}$ is the number of multigrid iterations, ν is the number of smoothing steps, and h_ℓ^{-2} represents the number of unknowns on grid level ℓ . The cost associated with the other components, such as computations of the defect, transfer to coarser grids, or interpolation of corrections, is accounted for by the constant \bar{c} , to facilitate the comparison of W-cycles with different numbers of smoothing steps.

The product $N_\ell \cdot (\mathcal{W}_\ell + \mathcal{W}_{\ell-1})$ in (33) gives the individual cost contributions at any level ℓ . For the proposed multigrid solver, we have $(\mathcal{W}_\ell + \mathcal{W}_{\ell-1}) = \mathcal{O}(h_\ell^{-2})$. For simplicity, let's assume that the cost of one sample on level ℓ is given by $\mathcal{O}(h_\ell^{-\gamma})$, $\gamma \geq d$, where d is the number of spatial dimensions. The computational cost (33) can be conveniently expressed as $\mathcal{W}_L^{ML} = \mathcal{O}(\sum_{\ell=0}^L 2^{(\gamma-\beta)\ell})$, leading to the following three cases. When the rate of decay of N_ℓ is faster than the growth in computational cost ($\beta > \gamma$), the dominating cost comes from the coarsest level. When the sample decay rate is similar to the growth in computational cost ($\beta = \gamma$), all levels equally contribute to the total cost. If the sample decay rate is slower than the growth in computational cost ($\beta < \gamma$), the dominant cost is on the finest level. This scenario is typically encountered in many practical settings involving three-dimensional problems with low regularity permeability models, resulting in a slow variance decay (or a small β).

Note that in many cases the rates α, β, γ are not available. A common practice is to use a few “warm-up samples and levels” to obtain an estimate for these quantities. Having a reliable LFA tool can be further helpful in efficiently optimizing the MLMC simulation, especially to assess the performance of a multigrid solver for different combinations of transfer operators and smoothers. This way, one can select the most suitable combination for the considered random input data. We can also use the LFA to obtain solver statistics on fine grids where computing warm-up samples is not feasible due to high computational cost.

6.1. MLMC numerical experiments. In this section we analyze the performance of the proposed multigrid MLMC method. We consider PDE (5) on domain $D \in (0, 1)^2$ with mixed Dirichlet–Neumann boundary conditions,

$$(35) \quad \begin{aligned} u(0, y, \omega) &= 1, & u(1, y, \omega) &= 0, & \text{and} \\ k(\mathbf{x}, \omega) \frac{\partial u}{\partial x} \Big|_{y=0} &= 0, & k(\mathbf{x}, \omega) \frac{\partial u}{\partial x} \Big|_{y=1} &= 0, \end{aligned}$$

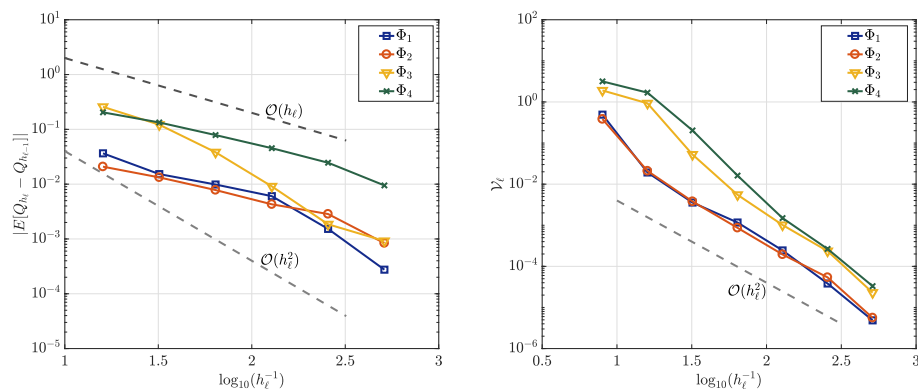


FIG. 9. (Left) Convergence of finite volume error. (Right) Decay of level dependent variance for the outflow, Q_{h_ℓ} , with mesh refinement for different Matérn parameters.

respectively. For all tests the source term is set to zero, i.e., $f = 0$. As the quantity of interest, the outflow through the boundary $x = 1, y = (0, 1)$, also referred to as *effective permeability*, is considered:

$$(36) \quad Q(u) = - \int_0^1 k(\mathbf{x}, \omega) \frac{\partial u}{\partial x}(\mathbf{x}, \omega) \Big|_{x=1} dy.$$

In Figure 9, we show the convergence of the FV bias, $|\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]|$, and the level dependent variance \mathcal{V}_ℓ for the considered quantity of interest (36). We observe that the rate of decay of the finite volume bias depends on the smoothness parameter ν_c of the random fields. In the case of parameter sets Φ_1 and Φ_3 , we see a second-order convergence, whereas a first-order convergence is observed for Φ_2 and Φ_4 . The correlation length λ_c and variance σ_c^2 only affect the proportionality constant. Also, for all four cases, the level dependent variance shows a quadratic decay.

To compute samples of $Q_{h_\ell}(\omega_i) - Q_{h_{\ell-1}}(\omega_i)$, we invoke the multigrid solver twice on level ℓ and $\ell - 1$ with the permeability field $k(\mathbf{x}_\ell, \omega_i)$ and an upscaled version of permeability $k(\mathbf{x}_{\ell-1}, \omega_i)$, respectively. We define $k(\mathbf{x}_{\ell-1}, \omega_i)$ using the *covariance upscaling* technique from [44] (see Appendix B). It is important to emphasize that homogenization techniques, e.g., [30, 31] for obtaining the coarse-scale solution, cannot be used to compute $Q_{h_{\ell-1}}(\omega_i)$. This is due to the fact that the homogenized permeabilities are typically obtained by averaging procedures that can modify the spatial covariance as well as the marginal distribution defined for the permeability field. This will lead to a violation of the telescopic identity (27). However, homogenization methods can still be used within a multigrid solver to obtain coarse-grid discretization operators.

Next, we conduct MLMC experiments for Matérn parameter sets Φ_2 and Φ_4 , using W(1,1), W(2,2), and W(3,3)-cycles. The multigrid components described in section 2.2 are used with stopping criterion $\varepsilon_{MG} = 10^{-10}$. We consider $\alpha = 1$ and $\beta = 2$ for both parameter sets Φ_2 and Φ_4 . Recall that the MLMC sampling strategy (32) requires the total number of MLMC levels, the number of finest level samples N_L , along with the rate β . We set $N_L = 64$ within MLMC for both parameter sets resulting in a sample sequence $N_\ell = N_L 4^{(L-\ell)}$. We also tested smaller values for N_L , where we still observed high statistical fluctuation in the estimates of $\mathbb{E}[Q_{h_\ell}]$, especially for parameter set Φ_4 . As mentioned earlier, the number of samples on all levels can also be determined by solving an optimization problem where the total cost of MLMC for a given tolerance is minimized [8, 9].

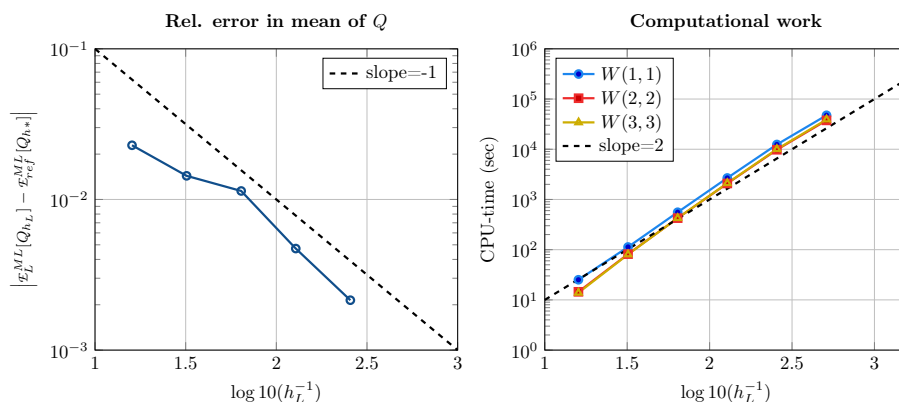


FIG. 10. (Left) Convergence of error in $\mathcal{E}_L^{ML}[Q_{h_L}]$ with increasing number of MLMC levels for parameter Φ_2 . The reference solution is based on mesh size $h^* = 1/512$. (Right) Mean CPU-times versus accuracy for different W-cycles. Computational cost proportional to $\mathcal{O}(h_L^{-2})$ is observed for all cycling strategies.

Remark. The experiments are performed on a parallel machine with a homogeneous set of cores. The MLMC algorithm allows for three degrees of parallelization: over samples, over MC levels, and within the multigrid solver. Recent advancements on parallelization can be found in [43, 45, 46]. We only consider parallelization over the MLMC samples and levels. For the considered MLMC sampling strategy with known number of samples and levels, a greedy algorithm based approach can be followed where cores are distributed between levels such that the sum of idle times for all cores is minimized.


6.1.1. MLMC results for Φ_2 . We first analyze the error convergence of the expected value and the variance of the quantity of interest. In Figure 10 (left), relative errors with respect to a reference solution computed using the MLMC estimator with the finest level $h^* = 1/512$ are shown. We see the error $|\mathcal{E}_L^{ML}[Q_{h_L}] - \mathcal{E}_{ref}^{ML}[Q_{h^*}]|$ converges with $\mathcal{O}(h_L)$. Next, we analyze the performance of the MLMC estimator for different multigrid cycling strategies. In Figure 10 (right), we present the CPU-times for the MLMC simulations with an increasing number of MLMC levels. These CPU-times are derived by summing up the run-times from the multigrid solves for all samples over all levels. Here, we do not include the cost of generation of the random field and the postprocessing costs as they are the same for different cycles. The cost scales as $\mathcal{O}(h_L^{-2})$ coinciding with the theoretical MLMC complexity when $\beta = \gamma$. No substantial difference in the CPU-times from the three cycles is observed. We also show the levelwise CPU-times in Table 4 for the three cycles with the same number of MLMC samples per level. In general, the computational cost improves with levels for all cycles. We observe that the W(2,2)-cycle slightly outperforms the other two variants.

We also show the distribution of the number of multigrid iterations for the W(2,2)-cycle on different grid levels along with the number of iterations predicted by the LFA, $\tilde{k}_\ell := \lceil \log(\varepsilon_{MG}) / \log(\langle \rho \rangle_{LFA, \ell}) \rceil$, in Figure 11. The cost is more heterogeneous on coarser levels with larger variance in the number of iterations and is homogeneous for $h_\ell = 1/32$ and onward. A high variance in the multigrid convergence rate was also predicted by the LFA experiments in the previous section. Notice that if the dominant cost of the MLMC simulation comes from the coarsest level, the solution time for the

TABLE 4

Comparison of the three W -cycles in terms of the levelwise CPU-times for parameter Φ_2 .

h_ℓ	N_ℓ	Levelwise CPU-time (sec.)		
		W(1,1)	W(2,2)	W(3,3)
1/8	262144	489	331	323
1/16	65536	528	391	390
1/32	16384	454	362	363
1/64	4096	407	336	348
1/128	1024	376	320	342
1/256	256	357	308	338
1/512	64	350	300	333

FIG. 11. Distribution of number of multigrid iterations of $W(2,2)$ -cycles for different mesh sizes for parameter Φ_2 to reach $\varepsilon_{MG} < 10^{-10}$ and the LFA predicted number of iteration \tilde{k}_ℓ .

MLMC technique will be varying significantly when the proposed multigrid solver is utilized on the coarse grids. In such cases, it is advised to use a sparse direct solver on these coarser levels.

In this work, the considered multigrid stopping criteria of 10^{-10} is quite conservative. For many engineering applications a residual reduction of 10^{-6} may already be sufficient to reach a converged solution, therefore reducing the computational cost roughly by a factor of two.

6.1.2. MLMC results for Φ_4 . Now we describe MLMC results for the challenging parameter set Φ_4 for the same quantity of interest. In Figure 12 (left), we display the convergence of the relative error in the expected value of the quantity of interest using a reference solution obtained with finest resolution $h^* = 1/1024$. Due to a large variance σ_c and a small correlation length λ_c of the random field, the error is larger, as compared to the results of parameter set Φ_2 for same mesh size h_L , and the *expected* convergence rates are visible on relatively fine grids. A similar trend is observed for the computational cost with scaling as $\mathcal{O}(h_L^{-2})$ in Figure 12 (right). Again, in Table 5, the levelwise CPU-times for the three cycles are listed. Similar to the results for set Φ_2 , the cost improves with levels but the gain from $W(2,2)$ -cycles is more prominent.

Last, the distribution of the number of iterations for $W(2,2)$ -cycles is depicted in Figure 13. High variability in the number of multigrid iterations persists until a relatively fine grid ($h = 1/128$). However, the average number of iterations predicted by LFA coincides very well with the mode of these distributions. Again, a direct solver can be applied on these coarser levels to get a more reliable estimate of the cost. Here, we point out that the $W(2,2)$ -cycle takes roughly seven iterations similar to parameter Φ_2 to reduce the residual by 10^{-10} , indicating robustness with respect to the Matérn parameter.

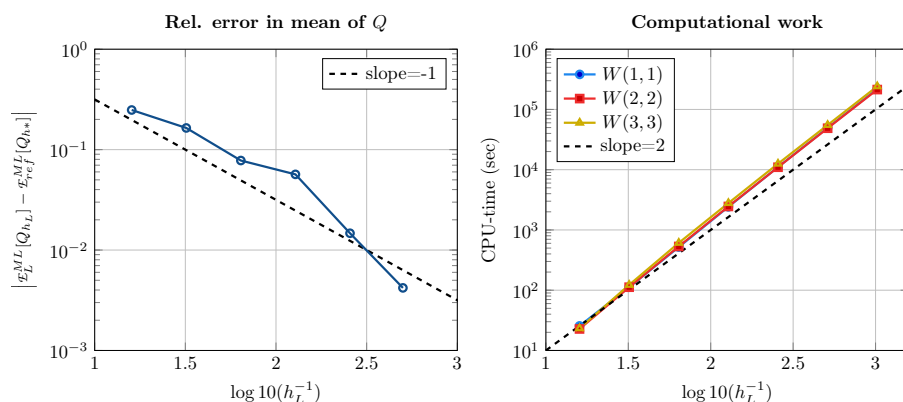


FIG. 12. (Left) Convergence of error in $x_L^{ML}[Q_{h_L}]$ with increasing number of MLMC levels for parameter Φ_4 . The reference solution is based on mesh size $h^* = 1/1024$. (Right) Mean CPU-times versus accuracy for different W -cycles. Computational cost proportional to $\mathcal{O}(h_L^{-2})$ is observed for all cycling strategies.

TABLE 5

Comparison of the three W -cycles in terms of the levelwise CPU-times for parameter Φ_4 .

h_ℓ	N_ℓ	Levelwise CPU-time (sec.)		
		W(1,1)	W(2,2)	W(3,3)
1/8	1048576	2.23[+3]	2.10[+3]	2.18[+2]
1/16	262144	2.42[+3]	2.52[+3]	2.80[+3]
1/32	65536	1.85[+3]	2.10[+3]	2.45[+3]
1/64	16384	1.56[+3]	1.59[+3]	1.98[+3]
1/128	4096	1.45[+3]	1.33[+3]	1.60[+3]
1/256	1024	1.40[+3]	1.23[+3]	1.40[+3]
1/512	256	1.38[+3]	1.23[+3]	1.34[+3]
1/1024	64	1.36[+3]	1.21[+3]	1.31[+3]

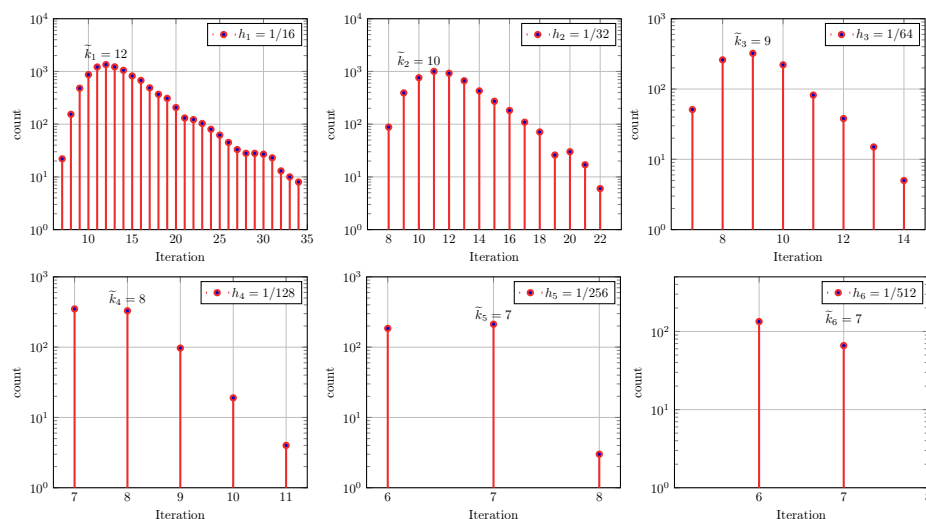


FIG. 13. Distribution of the number of multigrid iterations of $W(2,2)$ -cycles for different mesh sizes for parameter Φ_4 to reach $\varepsilon_{MG} < 10^{-10}$ and the LFA predicted number of iteration k_ℓ .

7. Conclusions. A novel, generalized local Fourier analysis which can be employed for quantitative assessment of multigrid methods for PDEs involving jumping and random coefficients has been presented. In particular, a cell-centered multigrid algorithm for solving a model problem based on a single-phase flow problem in a random porous medium has been used to show the accuracy of the proposed analysis. This technique, however, is appropriate for the prediction of the performance of a wider class of multigrid methods for solving PDEs with random fields. For instance, this approach can be extended to higher-order cell-centered and vertex-centered discretizations and corresponding multigrid methods, or even to multigrid methods for structured grid discretizations of coupled PDE systems, such as the Darcy–Stokes problem; see [47, 48]. Moreover, it should also be possible to apply it to triangular grids and finite element discretizations [49, 50]. Further, the novel local Fourier analysis could be used to estimate a priori the time needed for solving certain uncertainty quantification problems by using a multigrid multilevel Monte Carlo method.

Appendix A. Sampling Gaussian random fields. A fast algorithm for generating Gaussian random fields is critical for obtaining an efficient multigrid MLMC estimator. When dealing with stationary covariance functions, one can utilize spectral generators, e.g., in [37, 38, 51], that exploit the efficiency of the FFT algorithm to achieve fast sampling of these random fields on a uniform mesh. In this work, we use the fast Fourier transform moving average (FFT-MA) technique from [51] to decompose the covariance matrix $C_\Phi(r)$ (recall (26)). Although this sampling method is similar to the Cholesky factorization technique, the key idea is to make the computational domain periodic. Thus, the resulting covariance operator is also periodic, which can now be decomposed as a convolutional product. The samples of random fields are computed using a cheaper vector-vector product compared to the expensive matrix-vector operation required when using Cholesky factorization. As a periodic covariance function sampled on a uniform grid results in a *circulant* covariance matrix, these methods are sometimes also referred to as the circulant embedding technique. In the following, we provide a brief description of FFT-MA method from [51].

Samples of correlated Gaussian random vectors $\mathbf{z}_\ell(\omega)$ can be obtained using a Cholesky decomposition of the covariance matrix \mathbf{C}_Φ^ℓ on mesh D_ℓ as

$$(37) \quad \mathbf{C}_\Phi^\ell = \mathbf{L}_\ell \mathbf{L}_\ell^T \quad \text{and use} \quad \mathbf{z}_\ell = \mathbf{L}_\ell \mathbf{y}_\ell,$$

where \mathbf{y}_ℓ is a vector of i.i.d. samples from the standard normal distribution. The FFT-MA method is based on a decomposition of the covariance function $C_\Phi(r)$ as a convolutional product of some function $S_\Phi(r)$ and its transpose $S_\Phi'(r)$ ($S_\Phi'(r) = S_\Phi(-r)$). In a discrete setting, we can express this decomposition as

$$(38) \quad \mathbf{c}_\ell = \mathbf{s}_\ell * \mathbf{s}_\ell',$$

where $\mathbf{c}_\ell, \mathbf{s}_\ell$ are vectors obtained by evaluating $C_\Phi(r)$ and $S_\Phi(r)$, respectively, at grid points of the mesh D_ℓ . A correlated random vector \mathbf{z}_ℓ can now be synthesized by using the convolution product

$$(39) \quad \mathbf{z}_\ell = \mathbf{s}_\ell * \mathbf{y}_\ell.$$

The FFT-MA method performs the above computations in the frequency domain. First the vector \mathbf{c}_ℓ is transformed into a periodic signal, which is also real, positive, and symmetric. For more details on the practical aspects of this transformation, see [37, 38, 52, 44]. Note that the resulting vector \mathbf{s}_ℓ is also real, positive, and symmetric

and $\mathbf{s}_\ell = \mathbf{s}'_\ell$. As the convolution product in the spatial domain is equivalent to the componentwise product in the frequency domain, we can rewrite (38) as

$$(40) \quad \mathcal{F}(\mathbf{c}_\ell) = \mathcal{F}(\mathbf{s}_\ell) \cdot \mathcal{F}(\mathbf{s}_\ell) \implies \mathcal{F}(\mathbf{s}_\ell) = \sqrt{\mathcal{F}(\mathbf{c}_\ell)},$$

where \mathcal{F} denotes the discrete FFT and \cdot denotes componentwise multiplication. It is pointed out that the componentwise square-root operation does not pose any problems as the power spectrum $\mathcal{F}(\mathbf{c}_\ell)$ is real, positive. Next, we express the convolution product in (39) as a vector-vector product in the frequency domain as

$$(41) \quad \mathcal{F}(\mathbf{z}_\ell) = \mathcal{F}(\mathbf{s}_\ell * \mathbf{y}_\ell) = \mathcal{F}(\mathbf{s}_\ell) \cdot \mathcal{F}(\mathbf{y}_\ell).$$

In the final step, an inverse FFT is applied to obtain the samples for Gaussian random fields

$$(42) \quad \mathbf{z}_\ell = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{s}_\ell) \cdot \mathcal{F}(\mathbf{y}_\ell)).$$

Due to the periodicity in the covariance vector \mathbf{c}_ℓ , the resulting random field \mathbf{z}_ℓ is also periodic. Thus, we only retain the part of the vector that corresponds to the physical domain. We also remark that it takes two FFT evaluations to obtain one sample of \mathbf{z}_ℓ (ignoring the FFT operation in (40) that is performed just once). Therefore, in terms of the number of floating point operations, the sampling cost is significantly smaller compared to one multigrid solve for the mesh sizes considered.

Appendix B. Upscaling Gaussian random fields. While estimating the correction term $\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]$ in the telescopic sum (27), the approximations $Q_{h_\ell}(\omega_i)$ and $Q_{h_{\ell-1}}(\omega_i)$ need to be positively correlated such that the variance $\mathbb{V}[Q_{h_\ell} - Q_{h_{\ell-1}}]$ is small. This is typically achieved by first sampling the fine-grid permeability vector, $k(\mathbf{x}_\ell, \omega_i)$, to compute $Q_{h_\ell}(\omega_i)$ and using an upscaled version, $k(\mathbf{x}_{\ell-1}, \omega_i)$, for $Q_{h_{\ell-1}}(\omega_i)$. While performing such upscaling of random fields, it is important to ensure that the telescopic sum (27) is not violated. In other words, the expectation of the random variable Q_{h_ℓ} when estimating $\mathbb{E}[Q_{h_\ell} - Q_{h_{\ell-1}}]$ and $\mathbb{E}[Q_{h_{\ell+1}} - Q_{h_\ell}]$ should be the same, i.e.,

$$(43) \quad \mathbb{E}[Q_{h_\ell}]^{(coarse)} = \mathbb{E}[Q_{h_\ell}]^{(fine)} \quad \text{for } \ell = \{0, 1, \dots, L-1\}.$$

There are a number of ways to generate an upscaled version of a random field. For instance, the coarse-grid permeability can be computed by means of homogenization and variational coarsening with operator-dependent interpolation (regarding more general flow equations with a full tensor permeability and additional closure terms) [30, 31]. Upscaling algorithms based on homogenization techniques work well in the context of deterministic PDEs; however, these homogenization procedures may result in a modified covariance structure on the coarser levels for PDEs with random coefficients, violating (43). They may, however, be applicable for quantities of interest that do not depend on the permeability field. In general, the same covariance structure can be maintained using the *covariance upscaling* [44] that employs the spectral generator on two consecutive grids using the same normally distributed vector \mathbf{y}_ℓ . Furthermore, in the case of the FFT-MA algorithm, the vector \mathbf{y}_ℓ is associated with respective grid points, and coarser realizations of the fine-grid Gaussian random field \mathbf{z}_ℓ ($\log k(\mathbf{x}_\ell)$) can be obtained by using multidimensional averaging of the vector \mathbf{y}_ℓ . For instance, in two dimensions for the cell-centred mesh,

$$(44) \quad \mathbf{y}_{\ell-1}^{i,j} = \frac{1}{2} \left(\mathbf{y}_\ell^{2i-1, 2j-1} + \mathbf{y}_\ell^{2i-1, 2j} + \mathbf{y}_\ell^{2i, 2j-1} + \mathbf{y}_\ell^{2i, 2j} \right),$$




FIG. 14. Comparison of cross section of the pressure fields u at $x = 0.5$ (left) and $y = 0.5$ (right) obtained from a random permeability field (based on Φ_4) on a 256×256 grid and from the corresponding upscaled permeabilities at 128×128 and 64×64 grids.

where (i, j) is the cell index for the mesh $\mathcal{D}_{\ell-1}$. The scaling by a factor of 2 is needed to obtain a standard normal distribution for the averaged quantity $\mathbf{y}_{\ell-1}^{i,j}$. The coarse random field can now be simply assembled as

$$(45) \quad \mathbf{z}_{\ell-1} = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{s}_{\ell-1}) \cdot \mathcal{F}(\mathbf{y}_{\ell-1})).$$

Although this process can be recursively applied to generate upscaled random fields on all coarser scales, for the MLMC estimator, we just need this upscaling procedure for the next coarser grid. As the averaging in (44) smooths out high frequencies, the upscaled version $\mathbf{z}_{\ell-1}$ will also be smoother compared to \mathbf{z}_{ℓ} . In Figure 14, we compare the cross sections of pressure fields obtained from a permeability field sampled (using parameter set Φ_4) on a 256×256 grid and from the corresponding upscaled versions on 128×128 and 64×64 grids. The fine-scale properties are well preserved on the coarser levels.

REFERENCES

- [1] M. BOLTEN AND H. RITTICH, *Fourier analysis of periodic stencils in multigrid methods*, SIAM J. Sci. Comput., 40 (2018), pp. A1642–A1668, doi:10.1137/16M1073959.
- [2] J. P. DELHOMME, *Spatial variability and uncertainty in groundwater flow parameters: A geostatistical approach*, Water Resources Research, 15 (1979), pp. 269–280, doi:10.1029/WR015i002p00269.
- [3] R. A. FREEZE, *A stochastic-conceptual analysis of one-dimensional groundwater flow in nonuniform homogeneous media*, Water Resources Research, 11 (1975), pp. 725–741, doi:10.1029/WR011i005p00725.
- [4] R. J. HOEKSEMA AND P. K. KITANIDIS, *Analysis of the spatial structure of properties of selected aquifers*, Water Resources Research, 21 (1985), pp. 563–572, doi:10.1029/WR021i004p00563.
- [5] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., 45 (2007), pp. 1005–1034, doi:10.1137/050645142.
- [6] D. XIU, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, Princeton, NJ, 2010.
- [7] B. SEYNAEVE, E. ROSSEEL, B. NICOLA, AND S. VANDEWALLE, *Fourier mode analysis of multigrid methods for partial differential equations with random coefficients*, J. Comput. Phys., 224 (2007), pp. 132–149, doi:10.1016/j.jcp.2006.12.011.
- [8] K. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients*, Comput. Vis. Sci., 14 (2011), pp. 3–15, doi:10.1007/s00791-011-0160-x.

- [9] M. B. GILES, *Multilevel Monte Carlo path simulation*, Oper. Res., 56 (2008), pp. 981–986, doi:10.1287/opre.1070.0496.
- [10] M. B. GILES, *Multilevel Monte Carlo methods*, Acta Numer., 24 (2015), pp. 259–328, doi:10.1017/S096249291500001X.
- [11] A. BRANDT AND O. E. LIVNE, *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics*, Classics in Appl. Math. 67, SIAM, Philadelphia, 2011.
- [12] J. MOLENAAR, *A simple cell-centered multigrid method for 3D interface problems*, Comput. Math. Appl., 31 (1996), pp. 25–33, doi:10.1016/0898-1221(96)00039-9.
- [13] I. YAVNEH, *Coarse-grid correction for nonelliptic and singular perturbation problems*, SIAM J. Sci. Comput., 19 (1998), pp. 1682–1699, doi:10.1137/S1064827596310998.
- [14] W. HACKBUSCH, *Multigrid Methods and Applications*, Springer, New York, 1985.
- [15] J. H. BRAMBLE, R. E. EWING, J. E. PASCIAK, AND J. SHEN, *The analysis of multigrid algorithms for cell centered finite difference methods*, Adv. Comput. Math., 5 (1996), pp. 15–29, doi:10.1007/BF02124733.
- [16] J. H. BRAMBLE, J. E. PASCIAK, AND J. XU, *The analysis of multigrid algorithms with non-nested spaces or non-inherited quadratic forms*, Math. Comp., 56 (1991), pp. 1–34, doi:10.2307/2008527.
- [17] M. BREZINA, R. FALGOUT, S. MACLACHLAN, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Adaptive smoothed aggregation (α SA) multigrid*, SIAM Rev., 47 (2005), pp. 317–346, doi:10.1137/S1064827502418598.
- [18] K. STÜBEN, *Appendix A: An introduction to algebraic multigrid*, in Multigrid, U. Trottenberg, C. W. Oosterlee, and A. Schüller, eds., Academic Press, San Diego, CA, 2001, pp. 413–532.
- [19] P. VANĚK, M. BREZINA, AND J. MANDEL, *Convergence of algebraic multigrid based on smoothed aggregation*, Numer. Math., 88 (2001), pp. 559–579, doi:10.1007/s211-001-8015-y.
- [20] D. BRAESS, *Towards algebraic multigrid for elliptic problems of second order*, Computing, 55 (1995), pp. 379–393, doi:10.1007/BF02238488.
- [21] P. VANĚK, J. MANDEL, AND M. BREZINA, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196, doi:10.1007/BF02238511.
- [22] R. E. ALCOUFFE, A. BRANDT, J. E. DENDY, JR., AND J. W. PAINTER, *The multigrid methods for the diffusion equation with strongly discontinuous coefficients*, SIAM J. Sci. Statist. Comput., 2 (1981), pp. 430–454, doi:10.1137/0902035.
- [23] J. DENDY, *Black box multigrid*, J. Comput. Phys., 48 (1982), pp. 366–386, doi:10.1016/0021-9991(82)90057-2.
- [24] J. DENDY, *Black box multigrid for non-symmetric problems*, Appl. Math. Comput., 13 (1983), pp. 261–283, doi:10.1016/0096-3003(83)90016-4.
- [25] M. KHALIL AND P. WESSELING, *Vertex-centered and cell-centered multigrid for interface problems*, J. Comput. Phys., 98 (1992), pp. 1–10, doi:10.1016/0021-9991(92)90168-X.
- [26] P. WESSELING, *Cell-centered multigrid for interface problems*, J. Comput. Phys., 79 (1988), pp. 85–91, doi:10.1016/0021-9991(88)90005-8.
- [27] J. E. DENDY AND J. D. MOULTON, *Black box multigrid with coarsening by a factor of three*, Numer. Linear Algebra Appl., 17, pp. 577–598, doi:10.1002/nla.705.
- [28] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Academic Press, San Diego, CA, 2000.
- [29] S. KNAPEK, *Matrix-dependent multigrid homogenization for diffusion problems*, SIAM J. Sci. Comput., 20 (1998), pp. 515–533, doi:10.1137/S1064827596304848.
- [30] J. MOULTON, J. E. DENDY, AND J. M. HYMAN, *The black box multigrid numerical homogenization algorithm*, J. Comput. Phys., 142 (1998), pp. 80–108, doi:10.1006/jcph.1998.5911.
- [31] S. P. MACLACHLAN AND J. D. MOULTON, *Multilevel upscaling through variational coarsening*, Water Resources Research, 42 (2006), doi:10.1029/2005WR003940.
- [32] M. S. HANDCOCK AND J. R. WALLIS, *An approach to statistical spatial-temporal modeling of meteorological fields (with discussion)*, J. Amer. Statist. Assoc., 89 (1994), pp. 368–390, doi:10.2307/2290832.
- [33] R. J. ADLER, *The Geometry of Random Fields*, Classics in Math. Appl. 62, SIAM, Philadelphia, 2010, doi:10.1137/1.9780898718980.
- [34] F. NOBILE AND F. TESEI, *A multilevel Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients*, Stoch. Partial Differ. Equ. Anal. Comput., 3 (2015), pp. 398–444, doi:10.1007/s40072-015-0055-9.
- [35] J. CHARRIER, R. SCHEICHL, AND A. L. TECKENTRUP, *Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods*, SIAM J. Numer. Anal., 51 (2013), pp. 322–352, doi:10.1137/110853054.

- [36] J. GALVIS AND M. SARKIS, *Approximating infinity-dimensional stochastic Darcy's equations without uniform ellipticity*, SIAM J. Numer. Anal., 47 (2009), pp. 3624–3651, doi:10.1137/080717924.
- [37] C. R. DIETRICH AND G. N. NEWSAM, *Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix*, SIAM J. Sci. Comput., 18 (1997), pp. 1088–1107, doi:10.1137/S1064827592240555.
- [38] A. T. A. WOOD AND G. CHAN, *Simulation of stationary Gaussian processes in $[0, 1]^d$* , J. Comput. Graph. Statist., 3 (1994), pp. 409–432, doi:10.2307/1390903.
- [39] R. G. GHANEM AND P. D. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991, doi:10.1007/978-1-4612-3094-6.
- [40] CH. SCHWAB AND R. A. TODOR, *Karhunen-Loeve approximation of random fields by generalized fast multipole methods*, J. Comput. Phys., 217 (2006), pp. 100–122, doi:10.1016/j.jcp.2006.01.048.
- [41] R. WIENANDS AND C. W. OOSTERLEE, *On three-grid Fourier analysis for multigrid*, SIAM J. Sci. Comput., 23 (2001), pp. 651–671, doi:10.1137/S106482750037367X.
- [42] S. MISHRA AND CH. SCHWAB, *Sparse tensor multi-level Monte Carlo finite volume methods for hyperbolic conservation laws with random initial data*, Math. Comp., 81 (2012), pp. 1979–2018, doi:10.1090/S0025-5718-2012-02574-9.
- [43] S. MISHRA, CH. SCHWAB, AND J. ŠUKYS, *Multi-level Monte Carlo finite volume methods for nonlinear systems of conservation laws in multi-dimensions*, J. Comput. Phys., 231 (2012), pp. 3365–3388, doi:10.1016/j.jcp.2012.01.011.
- [44] S. MISHRA, CH. SCHWAB, AND J. ŠUKYS, *Multi-level Monte Carlo finite volume methods for uncertainty quantification of acoustic wave propagation in random heterogeneous layered medium*, J. Comput. Phys., 312 (2016), pp. 192–217, doi:10.1016/j.jcp.2016.02.014.
- [45] D. DRZISGA, B. GMEINER, U. RÜDE, R. SCHEICHL, AND B. WOHLMUTH, *Scheduling massively parallel multigrid for multilevel Monte Carlo methods*, SIAM J. Sci. Comput., 39 (2017), pp. S873–S897, doi:10.1137/16M1083591.
- [46] J. ŠUKYS, S. MISHRA, AND CH. SCHWAB, *Static load balancing for multilevel Monte Carlo finite volume solvers*, in *Parallel Processing and Applied Mathematics*, Springer, Berlin, 2012, pp. 245–254, doi:10.1007/978-3-642-31464-3_25.
- [47] P. LUO, C. RODRIGO, F. J. GASPAS, AND C. W. OOSTERLEE, *Uzawa smoother in multigrid for the coupled porous medium and Stokes flow system*, SIAM J. Sci. Comput., 39 (2017), pp. S633–S661, doi:10.1137/16M1076514.
- [48] P. KUMAR, P. LUO, F. J. GASPAS, AND C. W. OOSTERLEE, *A multigrid multilevel Monte Carlo method for transport in the Darcy-Stokes system*, J. Comput. Phys., 371 (2018), pp. 382–408, doi:10.1016/j.jcp.2018.05.046.
- [49] C. RODRIGO, P. SALINAS, F. J. GASPAS, AND F. J. LISBONA, *Local Fourier analysis for cell-centered multigrid methods on triangular grids*, J. Comput. Appl. Math., 259 (2014), pp. 35–47, doi:10.1016/j.cam.2013.03.040.
- [50] C. RODRIGO, F. J. GASPAS, X. HU, AND L. T. ZIKATANOV, *Stability and monotonicity for some discretizations of the Biot's consolidation model*, Comput. Methods Appl. Mech. Engrg., 298 (2016), pp. 183–204, doi:10.1016/j.cma.2015.09.019.
- [51] M. L. RAVALEC, B. NOETINGER, AND L. Y. HU, *The FFT Moving Average (FFT-MA) generator: An efficient numerical method for generating and conditioning Gaussian simulations*, Math. Geol., 32 (2000), pp. 701–723, doi:10.1023/A:1007542406333.
- [52] I. G. GRAHAM, F. Y. KUO, D. NUYENS, R. SCHEICHL, AND I. H. SLOAN, *Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications*, J. Comput. Phys., 230 (2011), pp. 3668–3694, doi:10.1016/j.jcp.2011.01.023.