



ELSEVIER

Available online at www.sciencedirect.com



Comput. Methods Appl. Mech. Engrg. 192 (2003) 4723–4744

**Computer methods
in applied
mechanics and
engineering**

www.elsevier.com/locate/cma

A multigrid solver for two-dimensional stochastic diffusion equations

O.P. Le Maître ^{a,*}, O.M. Knio ^b, B.J. Debuschere ^c, H.N. Najm ^c,
R.G. Ghanem ^d

^a *Université d'Evry Val d'Essonne, Centre d'Etudes de Mécanique d'Ile de France, 40 Rue du Pelvoux CE 1455, 91 020 Evry cedex, France*

^b *Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218-2686, USA*

^c *Sandia National Laboratories, Combustion Research Facility, Livermore, CA 94550, USA*

^d *Department of Civil Engineering, Johns Hopkins University, Baltimore, MD 21218-2686, USA*

Received 4 October 2002; received in revised form 9 June 2003; accepted 11 July 2003

Abstract

Steady and unsteady diffusion equations, with stochastic diffusivity coefficient and forcing term, are modeled in two dimensions by means of stochastic spectral representations. Problem data and solution variables are expanded using the Polynomial Chaos system. The approach leads to a set of coupled problems for the stochastic modes. Spatial finite-difference discretization of these coupled problems results in a large system of equations, whose dimension necessitates the use of iterative approaches in order to obtain the solution within a reasonable computational time. To accelerate the convergence of the iterative technique, a multigrid method, based on spatial coarsening, is implemented. Numerical experiments show good scaling properties of the method, both with respect to the number of spatial grid points and the stochastic resolution level.

© 2003 Elsevier B.V. All rights reserved.

Keywords: Stochastic problem; Diffusion equation; Multigrid; Polynomial Chaos; Karhunen–Loève expansion; Random media

1. Introduction

Developments in the field of computational mechanics and physics are enabling the solution of increasingly more realistic engineering problems. These advances take advantage of (i) enhanced computational capabilities—including parallel platforms and parallel techniques; (ii) elaborate models to handle more physical effects with less approximation of the system dynamics; and (iii) the development of numerical methods to reduce computational time and/or improve accuracy. As with many other fields, the

* Corresponding author.

E-mail addresses: olm@cemif.univ-evry.fr (O.P. Le Maître), knio@jhu.edu (O.M. Knio), bjdebus@ca.sandia.gov (B.J. Debuschere), hnnajm@ca.sandia.gov (H.N. Najm), ghanem@jhu.edu (R.G. Ghanem).

field of computational stochastic mechanics has benefited from such recent developments, and reached a level of maturity that allows for simulations that provide meaningful predictions to problems involving uncertain data [36]. Computational stochastic mechanics is particularly attractive in engineering sciences and physics, where the system to be analyzed can rarely be characterized exactly, while numerical methods usually require deterministic inputs. The present work takes place in this context.

Uncertainties in simulations of mechanical systems can be related to an inexact knowledge of the system geometry (e.g. [4,32]), boundary and initial conditions (e.g. [23,24,43]), external forcing (e.g. [18]), physical properties, or model parameters. Uncertainties can sometimes be due to measurement difficulties or to the intrinsic randomness of the processes, as in the case of heterogeneous media [14,15,29]. In order to deal with these uncertainties, distinct computational strategies have emerged, including Monte Carlo simulations (MCS) [25] and the integration of stochastic differential equations (SDE) [17,20], or a blend of the two. In MCS, the response surface of the random process is estimated by computing the deterministic responses of the system for a (large) set of distinct conditions that appropriately sample the uncertainty domain. In contrast, integration of the SDE governing the system is usually more difficult than the corresponding deterministic problem. Usually, integration of the SDE is achieved approximately, through statistical linearizations, asymptotic expansions, perturbation methods or truncated spectral representations. In the present study, uncertainty is taken into account by means of spectral expansions along the stochastic dimensions using Polynomial Chaos (PC) representations [5,42]. This representation is used to obtain a full statistical characterization of the response, in contrast with the other cited approaches which are usually limited to the very first statistical moments. Moreover, the spectral representation is now well established. Over the last few years, this technique has been successfully applied in various settings, including both solid and fluid mechanics [14,18,19,21,23,24,26,28,29,33,36–38].

In this work, we focus on a “generic” diffusion problem for a quantity u , with a random, spatially varying isotropic diffusion coefficient λ , inside a two-dimensional domain \mathbf{D} . The general form of the governing equation for this problem can be expressed as:

$$\alpha \frac{\partial u(\mathbf{x}, t, \theta)}{\partial t} = \nabla \cdot [\lambda(\mathbf{x}, t, \theta) \nabla u(\mathbf{x}, t, \theta)] + s(\mathbf{x}, t, \theta), \quad (1)$$

where $\alpha = 0, 1$ in the steady and unsteady cases, respectively, s is a given stochastic source term, and $\theta(\omega)$ denotes the stochastic character of the solution. The formulation is completed by specifying boundary conditions for u (generally Neumann or Dirichlet conditions), as well as an initial condition in the unsteady case. The elliptic form of Eq. (1) has been thoroughly analyzed from the mathematical point of view; see for instance [3,7,20,40].

Our current interest in an efficient solution method for Eq. (1) comes from the simulation of reacting electrochemical microchannel flow [9], where the steady form of the equation governs the electric field. There, u is the electrostatic field potential, λ is the electrical conductivity, and s is the charge accumulation due to diffusion of dissociated species (ions). λ and s are expressed as:

$$\lambda = F^2 \sum z_i^2 \mu_i c_i, \quad s = F \sum z_i \nabla \cdot (D_i \nabla c_i), \quad (2)$$

where F is the Faraday constant, z_i , μ_i , D_i and c_i are the charge number, electrophoretic mobility, molar diffusivity and molar concentration of the species i respectively [35]. Eq. (1) becomes stochastic whenever λ , s or the initial/boundary conditions on u are uncertain. In addition to this specific example, Eq. (1) appears, by itself or as part of a larger system, in the formulation of many problems involving gradient diffusion processes [9,19–21,23,26,33,37], as well as a variety of problems such as 1D-linear elasticity problems [10] and electromagnetism [38]. Therefore, an abstract study in a broader context than electrochemical flow is well justified.

A major difficulty in the solution of Eq. (1) concerns the representation of the stochastic diffusivity and source fields, and of the solution itself. As stated before, we will make use of spectral representations [17]

for this purpose. For brevity, but without loss of generality, the diffusivity field will be represented in terms of the Karhunen–Loève (KL) expansion, which is assumed to be known. Following the discussion above, this enables us to avoid describing the auxiliary problem of explicitly modeling the uncertainty in λ , and consequently focus our attention on the solution for u , which is sought in terms of its PC representation [17,20,36]. While the source term s may also be given in terms of a KL representation or more generally by a PC expansion, we will restrict our attention in the numerical tests to the case where $s \equiv 0$, i.e. to the homogeneous form of Eq. (1). For the purpose of the present construction, this enables us to avoid unnecessary details associated with setting up a stochastic source field. Both of these restrictions, however, can be easily relaxed within the framework of the construction.

One potential drawback of the spectral approach is that the size of the system of equations that needs to be solved grows rapidly as the number of stochastic dimensions increases. Specifically, the size of the stochastic system scales with the number of spectral-expansion modes retained in the computation, which increases rapidly with the number of dimensions and with the order of the expansion. Furthermore, since the equations governing the uncertainty modes are generally coupled, the CPU time needed to solve the stochastic system can increase rapidly with system size. This poses a serious computational challenge, which requires the development of efficient solvers [30,31,34]. The present study specifically aims at this objective, in the context of the generic formulation given in Eq. (1). Specifically, we describe the adaptation of a (deterministic) multigrid (MG) technique [41] for the solution of the system of equations arising from the finite-difference discretization of the spectral representation of the stochastic diffusion equation.

This paper is organized as follows. In Section 2, we recall the basic concepts and properties of the PC expansion of a stochastic process. Using these concepts, the stochastic spectral formulation of Eq. (1) is derived in Section 3 and the difficulties inherent to the solution of the spectral equations are discussed. Next, the finite-difference discretization of the stochastic system is introduced (Section 4), and an iterative technique is proposed to solve the resulting set of equations. In Section 5, a multigrid technique, based on spatial coarsening, is developed to improve the convergence rate of the previous iterative method. The multigrid algorithm is applied to selected test problems in Section 6, and the tests are used to examine its efficiency and scalability properties. Major findings are summarized in Section 7.

2. Polynomial Chaos representation

2.1. Spectral representation

In this section, the spectral representation of the stochastic process $u(\mathbf{x}, t, \theta)$ by means of the PC system is introduced. We consider the case where the uncertainty is due to a set of N independent (uncorrelated) stochastic parameters. The problem is then said to possess N stochastic dimensions, denoted by ξ_1, \dots, ξ_N , which are considered as generators of new dimensions (in addition to the space and time dimensions) in the solution process. Thus we have $\theta \equiv \{\xi_1, \dots, \xi_N\}$. Noting that u is a non-linear functional of its stochastic argument θ , it is natural to look for an orthogonal expansion of u in terms of the random variables ξ_i , $i = 1, \dots, N$. This idea has led to the concepts of homogeneous chaos and of PC expansions [5,17,42], in the case of Gaussian variables. The case of non-Gaussian measures is discussed in [12,13], but will not be considered here.

The dependence of $u(\mathbf{x}, t, \theta)$ on its stochastic arguments is approximated with the following truncated expansion:

$$u(\mathbf{x}, t, \theta) = \sum_{k=0}^P u_k(\mathbf{x}, t) \Psi_k(\theta), \quad (3)$$

where u_k are deterministic coefficients and $\{\Psi_0, \dots, \Psi_P\}$ is a (truncated) orthogonal basis consisting of multidimensional Hermite polynomials in ξ_i . The truncation is such that the degree of the polynomials is at most equal to N_0 , the order of the expansion. The total number of modes, $P + 1$, depends on N and N_0 , according to [8]:

$$P + 1 = \frac{(N + N_0)!}{N!N_0!}. \quad (4)$$

The orthogonality of the spectral basis $\{\Psi_k, k = 0, \dots, P\}$ on which u is expanded is defined with respect to the inner product:

$$\langle \Psi_i, \Psi_j \rangle \equiv \int \dots \int \Psi_i(\theta) \Psi_j(\theta) g(\xi_1) \dots g(\xi_n) d\xi_1 \dots d\xi_n, \quad (5)$$

where

$$g(\xi) = \frac{e^{-\xi^2/2}}{\sqrt{2\pi}} \quad (6)$$

is the Gaussian measure. Since θ is a Gaussian vector, inner products (Eq. (5)) and higher moments can be efficiently computed using moment formulas [22], based on a straightforward generalization of Gauss quadrature in one spatial dimension [1].

2.2. PC expansions of field variables

In general, all field variables may exhibit a stochastic character and should therefore be represented with PC expansions. In particular, the diffusivity and source fields can be expressed as:

$$\lambda(\mathbf{x}, t, \theta) = \sum_{k=0}^P \lambda_k(\mathbf{x}, t) \Psi_k(\theta), \quad s(\mathbf{x}, t, \theta) = \sum_{k=0}^P s_k(\mathbf{x}, t) \Psi_k(\theta), \quad (7)$$

respectively. Clearly, if λ and s are deterministic, then all modes with index $k > 0$ vanish identically. When this is not the case, the solution process u immediately admits a stochastic character, even when the boundary and, if relevant, initial conditions are deterministic.

The formulation above is quite general, and enables us to accommodate situations where the initial and boundary conditions on u , the diffusivity field, λ , and source field, s , are all uncertain. While the general case may be of interest, its treatment would require a detailed analysis of the source of uncertainty, which would distract from the present objective. Thus, in order to limit the scope of the simulations, while at the same time provide a meaningful test to the solver below, we restrict our attention to the case of a random diffusivity, deterministic boundary conditions, and vanishing source field. The diffusivity λ is assumed to be given by a Gaussian process with an exponentially decaying covariance function:

$$C(\mathbf{x}, \mathbf{x}') = \sigma_\lambda^2 \exp - \frac{\|\mathbf{x} - \mathbf{x}'\|}{L_c}, \quad (8)$$

where L_c is the correlation length and σ_λ is the standard deviation. We then use KL expansion [16,27] to express $\lambda(\mathbf{x})$ as:

$$\lambda(\mathbf{x}) = \bar{\lambda} + \sum_{k=1}^{\infty} \sqrt{\beta_k} \lambda_k(\mathbf{x}) \zeta_k, \quad (9)$$

where $\bar{\lambda}$ is the mean value, ζ_k are uncorrelated Gaussian variables with zero mean and unit variance, while β_k and λ_k are, respectively, the eigenvalues and eigenfunctions appearing in the spectral representation of C :

$$C(\mathbf{x}, \mathbf{x}') = \sum_{k=1}^{\infty} \beta_k \lambda_k(\mathbf{x}) \lambda_k(\mathbf{x}'). \quad (10)$$

Note that, in general, $\bar{\lambda}$ may vary in space, but we shall restrict our attention in the computations below to the case of a uniform mean value.

In the computations below, the eigenvalues and eigenfunctions are obtained with the Galerkin procedure described in [14–16]. The eigenvalues, all positive, are arranged in decreasing magnitude, and the KL expansion is truncated after the first N terms. Also note that the first N Polynomial Chaos coincide with the normalized Gaussian variables ζ_k , i.e. $\Psi_k(\theta) = \zeta_k$ for $k = 1, \dots, N$. Thus, the KL representation of λ can be formally viewed as a special case of a PC representation (Eq. (7)) in which polynomials of degree larger than one have vanishing coefficient. Different cases are considered by varying σ_λ and L_c , and analyzing their effect on the performance of the solver.

3. Continuous formulation and time discretization

3.1. Stochastic spectral diffusion equation

Introducing the PC expansions of the diffusivity, source and solution fields into the diffusion equation (1), one gets:

$$\alpha \sum_{k=0}^P \frac{\partial u_k}{\partial t} \Psi_k = \sum_{l=0}^P \sum_{m=0}^P \Psi_l \Psi_m \nabla \cdot [\lambda_l(\mathbf{x}, t) \nabla u_m(\mathbf{x}, t)] + \sum_{k=0}^P s_k(\mathbf{x}, t) \Psi_k. \quad (11)$$

Then, multiplying this equation by Ψ_i , evaluating its expectation and taking into account the orthogonality of the PC basis, we obtain:

$$\alpha \frac{\partial u_i(\mathbf{x}, t)}{\partial t} = \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{ilm} \nabla \cdot [\lambda_l(\mathbf{x}, t) \nabla u_m(\mathbf{x}, t)] + s_i(\mathbf{x}, t) \quad \text{for } i = 0, \dots, P. \quad (12)$$

The *multiplication tensor*

$$\mathcal{M}_{ilm} \equiv \frac{\langle \Psi_i \Psi_l \Psi_m \rangle}{\langle \Psi_i \Psi_i \rangle} \quad (13)$$

is independent of the solution, and is therefore computed and stored during a pre-processing stage. A multi-dimensional Gauss-quadrature approach [22] is used for this purpose. The tensor \mathcal{M} is sparse with a structure that depends on the order of the PC expansion and on the number of stochastic dimensions. The sparse nature of \mathcal{M} comes from the fact that many of the triple products $\Psi_i \Psi_j \Psi_k$, have vanishing expectation. For instance, due to orthogonality $\langle \Psi_0 \Psi_i \Psi_j \rangle = \delta_{ij} \langle \Psi_i^2 \rangle$. Also note that the triple product $\Psi_i \Psi_j \Psi_k$ can be written as the product of 1D polynomials, according to:

$$\Psi_i \Psi_j \Psi_k(\zeta_1, \dots, \zeta_N) = \prod_{q=1}^N p_{ijk}^q(\zeta_q), \quad (14)$$

where $p_{ijk}^q(\zeta_q)$ is a triple product of 1D Hermite polynomials. Thus, in order for $\langle \Psi_i \Psi_j \Psi_k \rangle$ to vanish, it is sufficient that only one p_{ijk}^q has vanishing expectation. On the other hand, \mathcal{M} is generally *not* diagonal (except for $N_0 = 0$), which leads to coupling between the u_i modes. An immediate consequence of this

coupling is an increase in the size of the resulting system of equations to be solved, compared to the deterministic case.

For unsteady problems, the use of an explicit time integration scheme for Eq. (12), as proposed in [23] in the context of the Navier–Stokes equations, leads to a simple algorithm that requires direct evaluation of the coupling terms. The use of explicit time-schemes has shown its efficiency for transient computations, but explicit stability restrictions on the time step can prove prohibitive on fine grids. For steady-state problems, a pseudo-transient approach may also be conceived, but in this case as well stability restrictions may lead to poor computational efficiency. Consequently, the development of an efficient numerical solver for the coupled system of equations is needed. This approach is adopted in the development below.

3.2. Boundary and initial conditions

Boundary conditions, and, when relevant, initial conditions, are needed to solve Eq. (12). These are also implemented in a “weak sense”, i.e. the boundary conditions are also projected onto the PC basis, leading to explicit conditions for each of the u_i 's. As noted in the introduction, the boundary conditions in the present study can be either of the Neumann or Dirichlet type. For brevity, we assume here that the boundary conditions are deterministic; an example of the use of uncertain boundary conditions is given in [24]. Denoting by $\partial\mathbf{D}_D$ and $\partial\mathbf{D}_N$ the part of the boundaries of \mathbf{D} where Dirichlet and Neumann conditions apply, respectively, the boundary conditions for all modes are given by:

$$u_0(\mathbf{x}, t) = u^D(\mathbf{x}, t), \quad u_{i \in [1, P]}(\mathbf{x}, t) = 0 \quad \forall \mathbf{x} \in \partial\mathbf{D}_D, \quad (15)$$

$$\frac{\partial u_0}{\partial n}(\mathbf{x}, t) = g_0^N(\mathbf{x}, t), \quad \frac{\partial u_k}{\partial n}(\mathbf{x}, t) = 0, \quad k = 1, \dots, P \quad \forall \mathbf{x} \in \partial\mathbf{D}_N, \quad (16)$$

where n denotes the direction normal to the boundary.

Note that for steady problems involving only Neumann conditions, the modes of the source field must satisfy the integral constraints

$$\int_{\mathbf{D}} s_k(\mathbf{x}) \, d\mathbf{x} = \int_{\partial\mathbf{D}} \sum_{l=0}^P \mathcal{M}_{k|l0} \lambda_l g_0^N \, ds \quad \text{for } k = 1, \dots, P. \quad (17)$$

Here, ds is the surface element along $\partial\mathbf{D}$. For unsteady problems, an initial condition for u is required. This initial condition may be deterministic or uncertain. In the former case, we have

$$u_0(\mathbf{x}, t = 0) = u_0^0(\mathbf{x}), \quad u_k(\mathbf{x}, t = 0) = 0, \quad \text{for } k = 1, \dots, P. \quad (18)$$

On the other hand, when the initial condition is uncertain, initial conditions for all the modes u_k need to be specified.

3.3. Implicit time discretization

A simple, generic, example of an implicit time integration method is the Euler backward scheme, whose application to Eq. (12) results in the following semi-discrete form:

$$\frac{\alpha}{\Delta t} u_i^{n+1} - \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{ilm} \nabla \cdot [\lambda_i^{n+1}(\mathbf{x}) \nabla u_m^{n+1}] = s_i^{n+1}(\mathbf{x}) + \frac{\alpha}{\Delta t} u_i^n, \quad (19)$$

where Δt is the time step and the superscripts refer to the time level. In the following, the dependence of λ and s on time is dropped, since these fields are assumed to be given.

4. Finite-difference discretization

4.1. Spatial discretization

Let $\mathbf{D} \equiv [0, L] \times [0, H]$ be a rectangular domain discretized in a set of $N_x \times N_y$ non-overlapping cells with uniform size $\Delta x = L/N_x$ and $\Delta y = H/N_y$ in the x - and y -directions. We denote by $(\Phi)_{i,j}$, for $i = 1, \dots, N_x$ and $j = 1, \dots, N_y$ the cell-averaged value of Φ , so that

$$(\Phi)_{i,j} \equiv \frac{1}{\Delta x \Delta y} \int_{(i-1)\Delta x}^{i\Delta x} \int_{(j-1)\Delta y}^{j\Delta y} \Phi(\mathbf{x}) \, dx \, dy, \tag{20}$$

where Φ stands for any of the field variables u_k , λ_k and s_k . Using this convention, we rely on the following centered, second-order spatial discretization of Eq. (19):

$$\begin{aligned} \frac{\alpha}{\Delta t} (u_k^{n+1})_{i,j} - \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm} & \left[\frac{(\lambda_l)_{i+1,j} + (\lambda_l)_{i,j}}{2} \frac{(u_m^{n+1})_{i+1,j} - (u_m^{n+1})_{i,j}}{\Delta x^2} \right. \\ & - \frac{(\lambda_l)_{i,j} + (\lambda_l)_{i-1,j}}{2} \frac{(u_m^{n+1})_{i,j} - (u_m^{n+1})_{i-1,j}}{\Delta x^2} + \frac{(\lambda_l)_{i,j+1} + (\lambda_l)_{i,j}}{2} \frac{(u_m^{n+1})_{i,j+1} - (u_m^{n+1})_{i,j}}{\Delta y^2} \\ & \left. - \frac{(\lambda_l)_{i,j} + (\lambda_l)_{i,j-1}}{2} \frac{(u_m^{n+1})_{i,j} - (u_m^{n+1})_{i,j-1}}{\Delta y^2} \right] = (s_k^{n+1})_{i,j} + \frac{\alpha}{\Delta t} (u_k^n)_{i,j}, \quad k = 0, \dots, P. \end{aligned} \tag{21}$$

The above equation can be re-cast in the following generic form:

$$\begin{aligned} \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm} & \left[(W_l)_{i,j} (u_m^{n+1})_{i+1,j} + (E_l)_{i,j} (u_m^{n+1})_{i-1,j} + (N_l)_{i,j} (u_m^{n+1})_{i,j+1} \right. \\ & \left. + (S_l)_{i,j} (u_m^{n+1})_{i,j-1} + (C_l^k)_{i,j} (u_m^{n+1})_{i,j} \right] = (f_k^{n+1})_{i,j}, \quad k = 0, \dots, P, \end{aligned} \tag{22}$$

which shows that a linear system of $N_x \times N_y \times (P + 1)$ equations must be solved in order to advance the solution by one time step. Of course, in the steady case this system is solved only once, and the superscripts indicating the time level are no longer needed.

4.1.1. Treatment of boundary conditions

Both Dirichlet and Neumann conditions are implemented using ghost cell techniques. For the case of Dirichlet condition, a ghost cell is introduced at the mirror image with respect to the boundary point of the neighboring interior cell. The value of the solution at the ghost cell is then determined by linearly extrapolating the solution from the interior, leading to a linear combination of the known value at the boundary and the neighboring interior node. Using this relationship, the ghost variables are then eliminated from the equation system. A similar approach is used in the case of a Neumann condition, based on expressing the known value of the normal derivative in terms of a second-order centered difference formula involving the solution at the neighboring internal node and the corresponding ghost node. The resulting relationship is then substituted into the equation system in order to eliminate the ghost variable. This approach results in a modified system of the form

$$\begin{aligned} \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm} & \left[(\tilde{W}_l)_{i,j} (u_m^{n+1})_{i+1,j} + (\tilde{E}_l)_{i,j} (u_m^{n+1})_{i-1,j} + (\tilde{N}_l)_{i,j} (u_m^{n+1})_{i,j+1} \right. \\ & \left. + (\tilde{S}_l)_{i,j} (u_m^{n+1})_{i,j-1} + (\tilde{C}_l^k)_{i,j} (u_m^{n+1})_{i,j} \right] = (\tilde{f}_k^{n+1})_{i,j}, \end{aligned} \tag{23}$$

where the tildes are used to indicate the modified values after implementation of the boundary conditions.

4.2. Iterative method

Since the size of system (23) is large for most applications, iterative solution methods are preferred over direct schemes. In this work, Gauss–Seidel iterations are used [39].

4.2.1. Outer iterations

Let us denote by $(\tilde{u}_m)_{i,j}^{ou}$ the estimate of $(u_m^{n+1})_{i,j}$ after the ou -th Gauss–Seidel iteration. This estimate can be computed by applying the following algorithm, called *outer iterations*, in contrast with the *inner iterations* described later:

- Loop on ou (Gauss–Seidel index)
 - For $i = 1$ to N_x , do
 - For $j = 1$ to N_y , do
 - Find $(\tilde{u}_k)_{i,j}^{ou+1}$ such that:

$$\begin{aligned} \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm}(\tilde{C}_l)_{i,j}(\tilde{u}_m)_{i,j}^{ou+1} &= (\tilde{f}_k)_{i,j} - \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm} \left[(\tilde{W}_l)_{i,j}(\tilde{u}_m)_{i+1,j}^{ou} + (\tilde{E}_l)_{i,j}(\tilde{u}_m)_{i-1,j}^{ou+1} \right. \\ &\quad \left. + (\tilde{N}_l)_{i,j}(\tilde{u}_m)_{i,j+1}^{ou} + (\tilde{S}_l)_{i,j}(\tilde{u}_m)_{i,j-1}^{ou+1} \right] \\ &\equiv (Q_k)_{i,j}^{ou}, \quad k = 0, \dots, P \end{aligned} \tag{24}$$

End of loop on j

- End of loop on i

- End of loop on ou

Thus,

$$(R_k)_{i,j}^{ou} = (Q_k)_{i,j}^{ou} - \sum_{l=0}^P \sum_{m=0}^P \mathcal{M}_{klm}(\tilde{C}_l)_{i,j}(\tilde{u}_m)_{i,j}^{ou}$$

is the local residual of Eq. (24), for the k th mode, at the ou -th Gauss–Seidel iteration.

4.2.2. Inner iterations

For each point in space, Eq. (24) can be rewritten in vector form as:

$$\begin{bmatrix} \sum_{l=0}^P \mathcal{M}_{00l}(\tilde{C}_l) & \dots & \sum_{l=0}^P \mathcal{M}_{0Pl}(\tilde{C}_l) \\ \vdots & \ddots & \vdots \\ \sum_{l=0}^P \mathcal{M}_{P0l}(\tilde{C}_l) & \dots & \sum_{l=0}^P \mathcal{M}_{PPl}(\tilde{C}_l) \end{bmatrix} \cdot \begin{pmatrix} (\tilde{u}_0)^{ou+1} \\ \vdots \\ (\tilde{u}_P)^{ou+1} \end{pmatrix} = \begin{pmatrix} (Q_0)^{ou} \\ \vdots \\ (Q_P)^{ou} \end{pmatrix}, \tag{25}$$

where the grid-point indices have been dropped for clarity. Thus, at this stage, one has to solve a system of $P + 1$ equations to compute $(\tilde{u}_{k=0,\dots,P})_{i,j}^{ou+1}$ from Eq. (24). A standard relaxation method (SOR) [39] is employed for this purpose. Denoting by ω the over-relaxation parameter, and by $[A_{km}]$ the system matrix corresponding to (25), the iterations are performed according to:

- Loop over in (SOR index)
 - Do $k = 0, \dots, P$

Compute a new estimate of (\tilde{u}_k) solution of Eq. (25) using

$$(\tilde{u}_k)_{in+1} = (1 - \omega)(\tilde{u}_k)_{in} + \frac{\omega}{A_{kk}} \left((Q_k)_{in} - \sum_{m=0}^{k-1} A_{km}(\tilde{u}_m)_{in+1} - \sum_{m=k+1}^P A_{km}(\tilde{u}_m)_{in} \right) \tag{26}$$

- End of loop over k .
- End of loop over in .

Note that for convenience, the Gauss–Seidel index ou has been dropped in Eq. (26).

Remark. The above decomposition of the iterative scheme into outer and inner loops may appear artificial, since a global iteration on the three-dimensional system for $(u_k^{n+1})_{ij}$ could be constructed. However, in view of the implementation of the multigrid scheme, which is based on spatial coarsening, it is found more convenient to clearly distinguish the inner iterations—which locally update the spectral coefficients of the solution, from the outer iterations—which account for the spatial coupling. In addition, computational tests (not shown) indicate that the convergence of the outer GS iteration is greatly improved when a more accurate estimate of the exact solution of Eq. (25) is used.

4.3. Convergence of the iterative scheme

The efficiency of the overall iterative method proposed above is estimated through the convergence rate of $(\tilde{u})^{ou}$ towards $(u)^{n+1}$, as the number of iterations ou increases. This convergence rate depends on the spectral radius of the system (23). Since \mathcal{M}_{klm} does not depend on the solution variables or parameters, the spectral radius is only a function of the stochastic diffusivity field $\lambda(\mathbf{x}, \theta)$, of the time-step Δt (if relevant), and of Δx and Δy . For the deterministic problem ($P = 0$), it is known that $(\lambda)_{ij} \geq 0$ for all i, j is necessary to ensure convergence, and that the convergence rate deteriorates when $\alpha/\Delta t$ decreases. In the stochastic case, the positivity of $\lambda_{i,j}$ is not ensured for all possible realizations. On the other hand, the numerical experiences in Section 6 indicate that the iterations converge when the coefficient of variation (COV), $\sigma_\lambda/\bar{\lambda}$, is sufficient small. As COV increases, the convergence rate of the present iterative scheme deteriorates; convergence fails above a critical value. These experiences appear to be consistent with the theoretical results in [2], where uniformly distributed random variables were used to ensure positivity. In the latter case, it is shown [2] that when COV is small, the solution exists and may be obtained by successive approximation.

The measure of convergence is obtained through the L^2 -norm of the residual for a given mode l which is expressed as:

$$\mathcal{N}_l \equiv \left\{ \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} [(R_l)_{i,j}]^2 \Delta x \Delta y \right\}^{1/2} . \tag{27}$$

The convergence of the iterative method will be further analyzed in Section 6 by monitoring the evolution of the maximum (over all modes) normalized residual:

$$\mathcal{R}_p \equiv \frac{\max_l [\mathcal{N}_l(p)]}{\max_l [\mathcal{N}_l(p=0)]} \tag{28}$$

where the index p refers to the number of MG cycles.

5. Multigrid acceleration

It is known from the analysis of deterministic diffusion equations that the convergence rate is a function of spatial frequencies. Specifically, the longest wavelengths exhibit the lowest convergence rate, while short scales converge faster. To improve convergence, acceleration techniques based on spatial coarsening have been proposed in the literature [6,11,41]. In the present work, we develop a multigrid technique for the stochastic case.

The basic idea of the multigrid technique is to treat the modes with low spatial frequencies on coarser grids, since fine spatial resolution is not required for these modes. The gain of the method is due to the faster convergence of the long-wave modes on the coarser grids, as well as the lower CPU cost of the corresponding iterations. Since multigrid methods are widely used, we will just recall the main ingredients of the approach, namely (i) the definition of the grid levels, (ii) the projection step and (iii) the prolongation procedure.

5.1. Definition of grid levels

Thanks to the regular structure of the computational grid, the coarsening is made by merging a set of neighboring grid cells to give a single cell on the next (coarser) grid level. This leads to a hierarchical set of grids. In the current implementation, a coarsening step consists of merging four cells (two in each direction) with surface areas $\Delta x^k \times \Delta y^k$ each, to obtain a child cell with surface area $\Delta x^{k+1} \times \Delta y^{k+1} = 4\Delta x^k \times \Delta y^k$, the superscripts denoting the respective grid levels. Thus, starting from a grid level k , made of $Nx^k \times Ny^k$ cells, the next grid level contains $Nx^{k+1} \times Ny^{k+1} = (Nx^k \times Ny^k)/4$ cells. Clearly, this process can be repeated as long as Nx^k and Ny^k are even numbers. Whenever one of the number of cells in a direction is odd, the coarsening automatically switches to a one-dimensional coarsening procedure in which only two cells are merged to make a child cell. This procedure is illustrated in Fig. 1, where the successive grid levels are plotted. Clearly the procedure is optimal when Nx and Ny are powers of 2.

5.2. Projection and prolongation procedures

On the finest grid level, a small number N_{ou} of outer iterations is first performed. This provides approximate solutions $(\tilde{u}_m)_{i,j}^{N_{ou}}$ with residuals $(R_m)_{i,j}^{N_{ou}}$. These residuals are then projected onto the next coarser grid, where problem (23) is considered, with $(R_m)_{i,j}^{N_{ou}}$ as the right-hand side (in lieu of $(f_k)_{i,j}$), and with the same but *homogeneous* boundary conditions. (In other words, on the coarser grids, the *residual equation* is solved.) To do so, one has to provide an estimate of (λ) and (R_m) on consecutive grid levels. This is achieved by averaging their respective values over the *parent* cells as illustrated in the left scheme of Fig. 2a.

On the new grid level, a few outer iterations are performed, following the same methodology, to obtain an approximate solution and a residual. The projection process is repeated until the last grid level is reached. Then, from the coarser grid level, where an estimate solution for the residual equation has been obtained, it is first transferred to the previous grid level through a prolongation procedure and then used to correct the solution on that finer grid level. In the current implementation, this is achieved by summing the cell averaged solution at level k with the solution of its parent cell, as shown in Fig. 2b. When the solution has been prolonged onto level $k - 1$, a few outer iterations are performed (smoothing step) the process is repeated until the initial, fine, grid is reached.

5.3. Multigrid cycles

Starting from the original grid, the application of successive projections up to the coarsest grid level, followed by successive prolongations up to the starting grid, is referred to as a cycle. Different kinds of cycles may be used [41], according to the excursion path along the grid levels. For instance, the so-called *W*-cycles have been designed to improve the convergence rate of the multigrid method, and many other examples can be found in the literature. Since our objective here is simply to develop a multigrid methodology for stochastic diffusion equations, we have limited ourselves to the simplest case of the *V*-cycle as described above. Moreover, we use a constant number of Gauss–Seidel iterations, denoted N_{ou} , on every grid-level, after every projection or prolongation step.

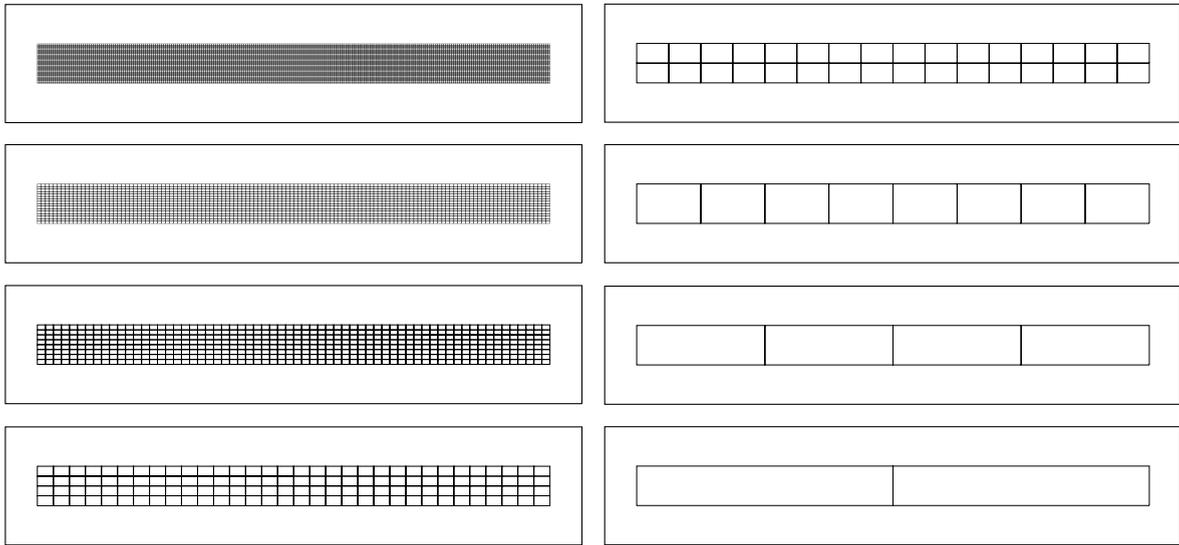


Fig. 1. Example of grid coarsening used for the multigrid method. The base grid consists of 32×256 cells (top left). The mesh is first coarsened by merging four (two in each direction) cells to form a coarser child cell. When the number of cells in one direction is odd, the coarsening process switches to 1D merging as in the last three grid levels.

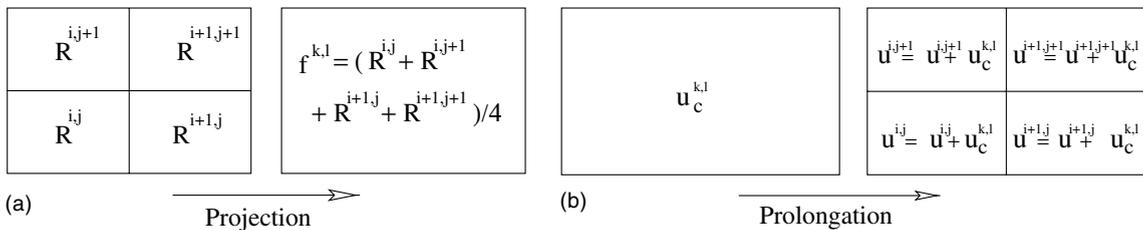


Fig. 2. Illustration of the (a) projection and (b) prolongation procedures to transfer data between two successive grid levels. In the projection step, the residual on a given grid is transferred to the next (coarser) grid level by spatial averaging. The same methodology is used to transfer the diffusivity data. For the prolongation of the solution from one grid level to the next (finer) one, simple addition is used.

5.4. Implementation of the multigrid scheme

Implementation of the MG scheme is now summarized as follows:

(1) **Initialization:**

- Determine spectral basis, compute and store the multiplication tensor \mathcal{M} .
- Compute the KL decomposition of the λ . (Alternatively the PC expansion of λ is imported (e.g. [8]) or set by the user.)
- Determine the system coefficients for all grid levels:

For $ig = 1, \dots, Ng$

- Determine the grid properties: $\Delta x^{ig} = 2^{ig-1} \Delta x$, $\Delta y^{ig} = 2^{ig-1} \Delta y$, $Nx^{ig} = Nx/2^{ig-1}$, $Ny = Ny/2^{ig-1}$.
- Compute the cell averaged diffusion field $(\lambda_i)_{i,j}^{ig}$, for $i = 1, \dots, Nx^{ig}$, $j = 1, \dots, Ny^{ig}$

- Using Eqs. (21) and (22), determine the system coefficients $(C_l^k)_{i,j}^{ig}$, $(E_l)_{i,j}^{ig}$, $(W_l)_{i,j}^{ig}$, $(N_l)_{i,j}^{ig}$, and $(S_l)_{i,j}^{ig}$.
- Compute and store the modified system coefficients accounting for the boundary conditions: $(\tilde{C}_l^k)_{i,j}^{ig}$, $(\tilde{E}_l)_{i,j}^{ig}$, $(\tilde{W}_l)_{i,j}^{ig}$, $(\tilde{N}_l)_{i,j}^{ig}$, $(\tilde{S}_l)_{i,j}^{ig}$.

End of loop over ig

- Initialize solution.

(2) **Loop over time index n :**

(a) Compute the right-hand side of system (23) on the first grid level: $(\tilde{f}_k^n)_{i,j}^{ig=1}$.

(b) Initialize solution on the first grid level: $(\tilde{u}_k)_{i,j}^{ig=1} = (\tilde{u}_k^n)_{i,j}$.

(c) **Beginning of V-cycle**

For $ig = 1, \dots, Ng$ (coarsening)

- If $ig > 1$ then initialize the solution $(\tilde{u})^{ig}$ to zero.

- **Outer loop:**

For $ou = 1, \dots, Nou$

- **Loop over spatial indices**

For $i = 1, \dots, N_x^{ig}$

For $j = 1, \dots, N_y^{ig}$

Using Eq. (24), compute the right-hand side of Eq. (25).

Inner loop

For $in = 1, \dots, Nin$

Loop over mode index

For $k = 0, \dots, P$

Apply Eq. (26) to $(\tilde{u}_k)_{i,j}^{ig}$

End of loop over k

End of loop over in

End of loop over i, j

End of loop over ou

- If $ig < Ng$, then

- Compute the local residual $(R_k)_{i,j}^{ig}$ of Eq. (23) on the current grid level.

- Project the local residuals to compute the right-hand side of Eq. (23) at the next grid level $ig + 1$, i.e. determine $(\tilde{f}_k)_{i,j}^{ig+1}$.

End of loop over ig

For $ig = Ng - 1, \dots, 1$ (refinement)

- Update solution $(\tilde{u}_k)^{ig}$ through the prolongation of $(\tilde{u}_k)^{ig+1}$.

- **Outer loop:**

For $ou = 1, \dots, Nou$

- **Loop over spatial indices**

For $i = 1, \dots, N_x^{ig}$

For $j = 1, \dots, N_y^{ig}$

Using Eq. (24), compute the right-hand side of Eq. (25).

Inner loop

For $in = 1, \dots, Nin$

Loop over mode index

For $k = 0, \dots, P$

Apply Eq. (26) to $(\tilde{u}_k)_{i,j}^{ig}$

End of loop over k

End of loop over in

End of loop over i, j

End of loop over ou

End of loop over ig

Compute local residual $(R_k)_{i,j}$ of Eq. (23) on the first grid level. If one of the norms \mathcal{N}_l from Eq. (27) is greater than the prescribed threshold, then a new V-cycle is performed starting from (c).

(d) Determine the solution: $(\tilde{u}_k^{n+1})_{i,j} = (\tilde{u}_k^{ig=1})_{i,j}$, for $i = 1, \dots, Nx$, $j = 1, \dots, Ny$ and $k = 1, \dots, P$.

(3) **End of time loop**

6. Results

We now present test results that show the behavior and convergence properties of the multigrid method. For the test cases below, we set $\alpha = 0$ and study the stochastic diffusion in a square domain, with unit edge-length and with no source term ($s \equiv 0$). Deterministic boundary conditions are used with Dirichlet conditions on $x = 0$ (where $u = 1$) and $x = 1$ (where $u = 0$), and homogeneous Neumann boundary conditions for the $y = 0$ and $y = 1$ edges. As stated previously, the spatially dependent diffusivity field is modeled using a truncated KL expansion involving N modes. It is characterized by its variance, coefficient of variation, and correlation length. To analyze the performance of the scheme, we monitor the evolution of the maximum (over all the modes) of the L^2 -norms of the normalized residuals, more specifically the decay of the peak residual as the number of multigrid cycles increases.

6.1. Multigrid acceleration

6.1.1. Dependence on grid size

We start by examining the dependence of the convergence rate on the number of points involved in the spatial discretization. To this end, the COV of the diffusivity field is set to 0.1, with a normalized correlation length $L_c = 5$. A KL expansion with 5 modes is used (i.e. $N = 5$), together with a second-order PC expansion ($N_0 = 2$). With these parameters, $P = 20$ and so the total number of modes equals 21. The multigrid parameters are selected as follows: $N_{ou} = N_{in} = 3$ and $\omega = 1.5$. The computations are performed for the spatial discretizations of $Nx = Ny = 16, 32, 64$ and 128; the corresponding number of grid levels are $Ng = 4, 5, 6$ and 7. For each, the maximum normalized residual is plotted against cycle number in Fig. 3.

These results clearly show the quasi-independence of the convergence rate with respect to the spatial discretization. There is a very weak improvement in the convergence rate at the lower values of Nx and Ny , which may be attributed to the lack of resolution in the representation of the KL modes on the coarser meshes. This claim is supported by the observation that the convergence rate tends to be grid-size independent when Nx and Ny increase. The weak dependence of the convergence rate on the grid size also highlights the excellent scalability of the method concerning the spatial discretization, as the CPU time scales roughly as $Nx \times Ny$. Note that the relaxation parameter, ω , and number of inner and outer iterations, N_{ou} and N_{in} , have been selected based on systematic tests (not shown) to determine their optimal value. While further refinement of these parameters may be possible, these values will be kept the same for the remaining cases below, unless explicitly stated.

6.1.2. Effect of grid levels on MG acceleration

Fig. 4 shows the evolution of the peak normalized residual with the number of cycles for a fixed grid with $Nx = Ny = 32$. Results obtained for different numbers of grid levels in the V-cycles are shown, namely $Ng = 1, 2, 3, 4$ and 5.

The results clearly show the effect of MG acceleration with increasing number of grid levels. The setting $Ng = 1$ corresponds to the Gauss–Seidel iteration, applied to the initial system of equations with no coarsening. Thus, after the first V-cycle (that is $2N_{ou} = 6$ GS iterations) the short-scales in the residual (mostly related to the Dirichlet boundary conditions) have been reduced and the convergence rate falls

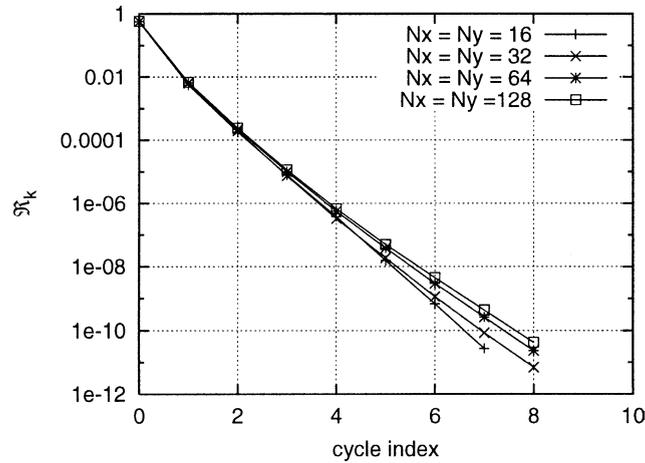


Fig. 3. Convergence of the iterative scheme for different spatial discretizations: $N_x = N_y = 16, 32, 64$ and 128 . The corresponding number of grid levels are $N_g = 4, 5, 6$ and 7 . $COV = 0.1, L_c = 5, N = 5, N_0 = 2 (P = 20), \omega = 1.5, N_{ou} = N_{in} = 3$.

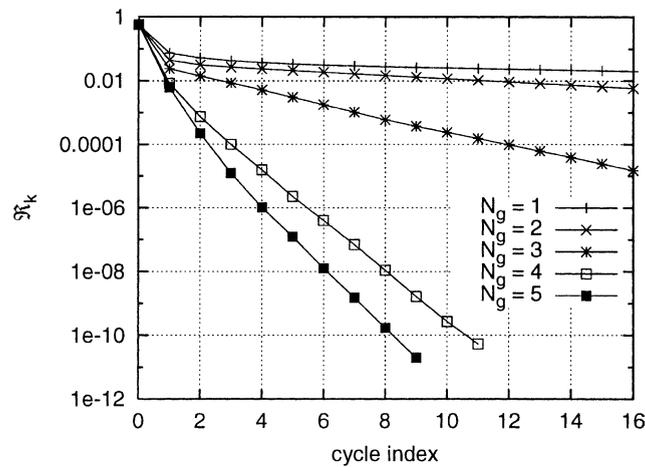


Fig. 4. Convergence of the iterative scheme for $N_x = N_y = 32$ and $N_g = 1, 2, 3, 4$ and 5 . $COV = 0.1, L_c = 1, N = 10, N_0 = 2 (P = 65), \omega = 1.5, N_{ou} = N_{in} = 3$.

dramatically. This clearly illustrates the lower convergence rate of the larger length scales. When the number of grid levels is increased to $N_g = 2$ the convergence rate is slightly improved, but the iterative method is still inefficient. In fact, the first significant improvement is reported for $N_g = 3$, where one observes a residual reduction factor per V-cycle of approximately 0.78. With $N_g = 4$, the convergence rate is much larger, as the residual reduction factor per V-cycle is approximately 0.2. As expected, the largest convergence rate is observed for $N_g = 5$, with a residual reduction factor close to 0.1. These tests show that the discretization parameters N_x and N_y should be selected, to the extent possible, so that the coarsest grid level has a minimum number of cells in each direction. Note, in particular, that the large improvement in convergence rate between $N_g = 3$ and $N_g = 5$ is achieved at a very low additional CPU cost, since the fourth and last grid levels only involve 16 and 4 cells respectively.

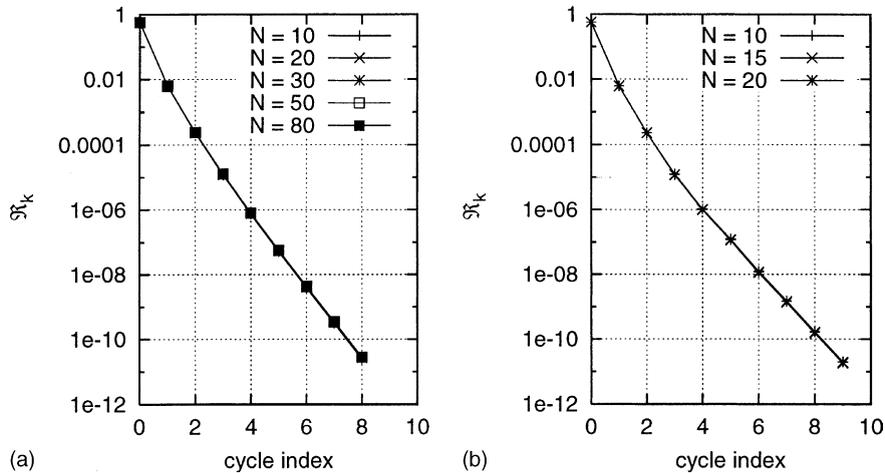


Fig. 5. Peak normalized residual versus cycle number: (a) first-order PC expansion with $N = 10, 20, 30, 50$ and 80 ; (b) second-order PC expansion with $N = 10$ ($P = 65$), 15 ($P = 135$) and 20 ($P = 230$). In both cases, $N_x = N_y = 32$, $L_c = 1$, $N_{in} = N_{ou} = 3$, and $\omega = 1.5$.

6.2. Influence of stochastic representation parameters

The behavior of the MG scheme is now investigated in terms of the stochastic representation parameters, namely the number, N , of KL modes used in the representation of the stochastic diffusivity field, and the order, N_0 , of the PC expansion. In the tests below, the spatial discretization parameters are held fixed, as are the over-relaxation parameter, $\omega = 1.5$, the number of grid levels, $N_g = 5$, and the number of iterations performed on each grid level, $N_{in} = N_{ou} = 3$. The coefficient of variation and correlation length are also held fixed, $\text{COV} = 0.1$ and $L_c = 1$, respectively.

6.2.1. Number of KL modes

In these tests, the impact of the number of modes retained in the KL expansion of the diffusivity field is investigated by varying N . Results are reported in Fig. 5, where the peak normalized residual is plotted against cycle number for (a) first-order PC expansion with N ranging from 10 to 80, and (b) second-order PC expansion with $N = 10, 15$ and 20 ($P = 65, 135$ and 230 , respectively). For both first- and second-order expansions, the evolution of the residual is independent of N , again showing the efficiency of the MG scheme. Note that in the present case one cannot infer from this behavior a linear relationship between N and the CPU time. The latter is in fact a strong function of the number of non-zero terms in \mathcal{M} , which depends on both N and N_0 . This contrasts with previous observation regarding scalability of the scheme with respect to the number of grid points.

6.2.2. Effect of PC expansion order on rate of convergence

The results of the previous section show a dependence of the convergence rate of the multigrid method on the order of the PC expansion. This dependence is further investigated by setting $N = 10$ and varying N_0 from 1 to 3; with $P = 10, 65$ and 285 , respectively. The convergence of the iterations for these cases is illustrated in Fig. 6, which depicts the behavior of the peak residual as the number of cycles increases.

The results indicate that the convergence rate decreases slightly as the order of the PC expansion increases. The residual reduction factor per cycle is about 0.09 for $N_0 = 1$ and approximately 0.2 for $N_0 = 3$. In light of the experiences above, it is evident that the present reduction in convergence rate is not due to the increase in number of modes P , but rather to the need for additional cycles in order to propagate the

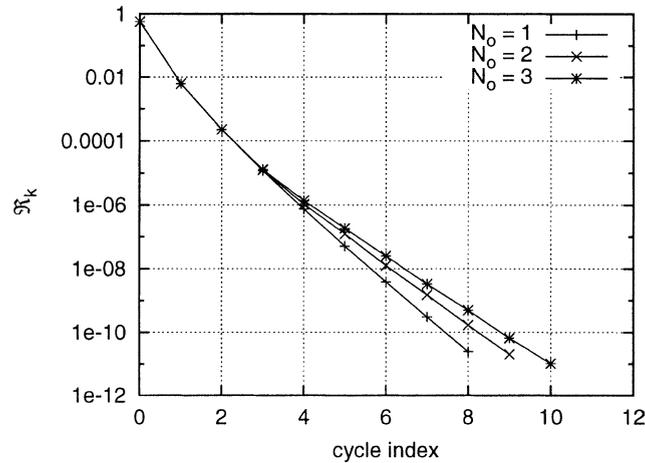


Fig. 6. Peak normalized residual versus cycle number for $N_0 = 1$ ($P = 10$), $N_0 = 2$ ($P = 65$), and $N_0 = 3$ ($P = 285$). In all cases, $N = 10$, $N_x = N_y = 32$, $L_c = 1$, $N_{in} = N_{ou} = 3$, and $\omega = 1.5$.

residual for coupled terms of different order. For the present examples, the convergence rate is still satisfactory for $N_0 = 3$. In situations requiring higher order expansions, however, further improvement may be required. This could be achieved for instance by blending the (spatial) MG concepts with a spectral (mode) coarsening procedure.

6.3. Effects of diffusivity field statistics

We now analyze the effects of the diffusivity field characteristics on the convergence rate, by varying its statistical parameters. We recall that the diffusivity is parametrized through its COV, which represents the normalized local statistical spread of the realizations about the expected value, and the correlation length, which accounts for the spatial variability of the process. The effects of these two parameters are analyzed separately below.

6.3.1. Effect of diffusivity variance

Tests on the effect of the variance of the diffusivity field are performed using $L_c = 1$, a KL expansion with $N = 25$, and PC expansion of first and second order ($P = 25$ and 350 , respectively). A 32×32 computational grid is used and the MG parameters are as follows: $\omega = 1.5$, $N_{ou} = N_{in} = 3$ and $Ng = 5$. Results with different values of σ_λ are reported in Fig. 7.

The results show that, as expected, the convergence rate is strongly dependent on the variance of the diffusivity field. Actually, the highest convergence rate is achieved for the lowest values of COV ($\sigma_\lambda = 0.025$) where the residual is reduced with each MG cycle, by a factor of about 0.05 for $N_0 = 1$ (Fig. 7a) and approximately 0.06 for $N_0 = 2$ (Fig. 7b). Consistent with previous findings, for the same values of σ_λ (and COV) the second-order PC expansion exhibits slower convergence rate than the first-order scheme. Moreover, when COV increases, the convergence rate decreases for both the first- and second-order expansions, but the reduction is more substantial in the latter case. As noted previously, the diffusivity field should be positive at any point inside the domain. However, the KL expansion does not guarantee this constraint, and the probability that this constraint is violated increases as COV increases. In fact, for $\text{COV} \geq 0.4$ and $N_0 = 2$, the solver did not converge. In contrast, with a first-order expansion the MG iterations did converge for $\text{COV} = 0.4$, but with a very low convergence rate (not shown). It should be

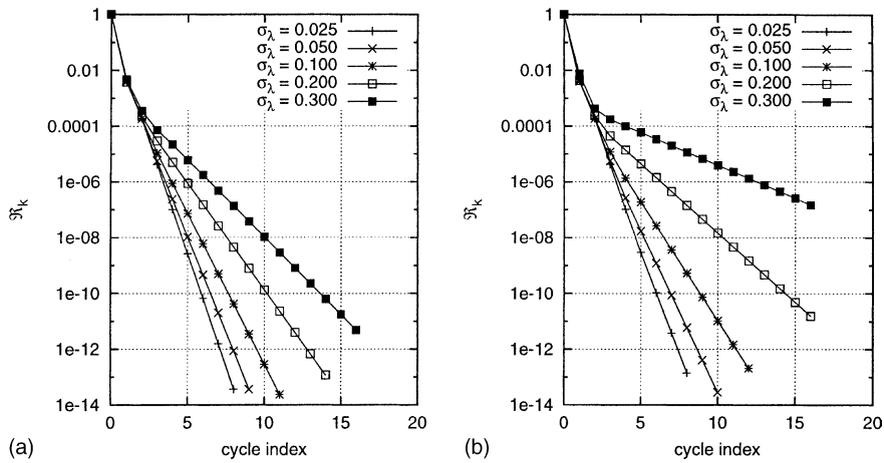


Fig. 7. Peak residual versus number of cycles for different values of σ_λ : (a) first-order PC expansion and (b) second-order PC expansion. In all cases, $L_c = 1$, $N = 25$, $\omega = 1.5$, $N_{ou} = N_{in} = 3$, $Ng = 5$ and $N_x = N_y = 32$.

emphasized, however, that the deterioration of convergence rate with increasing COV is due to the extreme behavior of the corresponding problem, and therefore is not inherent to the present multigrid scheme.

6.3.2. Influence of the correlation length

The effect of the correlation length is analyzed in this section by performing computations with a fixed variance ($COV = 0.1$) but varying L_c . As illustrated in Fig. 8, as L_c decreases, the spectrum of λ broadens with higher amplitudes in the small scales. Since the variance is fixed, however, the “energy” content of the spectrum remains constant.

Fig. 9 shows the convergence rate of the MG iterations for different correlation lengths, $L_c = 0.25, 0.5, 1, 2$ and 5 . Plotted are results obtained using both first- and second-order PC expansions. The results show a weak dependence of the convergence rate on L_c , indicating that the MG method effectively maintains its good convergence properties even as small-scale fluctuations in λ increase. The weak dependence of the convergence rate on spatial lengthscales of λ has also been observed in deterministic simulations (not

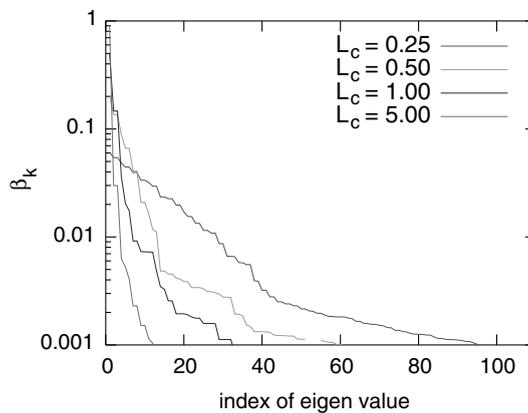


Fig. 8. Spectra of the eigenvalues, β_k , of the KL expansion for different correlation lengths $L_c = 0.25, 0.5, 1$ and 5 .

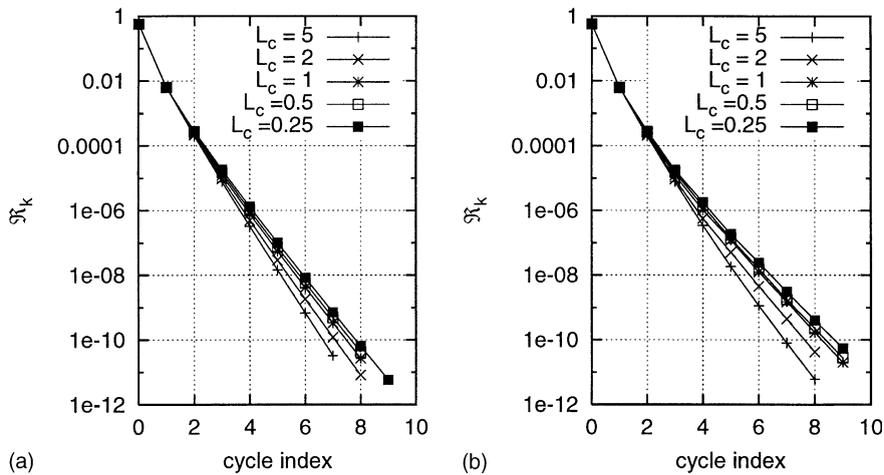


Fig. 9. Peak residual versus number of cycles for $N = 20$ KL modes and different values of L_c : (a) first-order PC expansion ($P = 20$); (b) second-order PC expansion ($P = 350$). In all cases, $\text{COV} = 0.1$, $N_{in} = N_{out} = 3$, $\omega = 1.5$ and $N_x = N_y = 32$.

shown). This shows that the extension of the MG scheme to stochastic problems does not adversely affect its effectiveness in dealing with spatially varying diffusivity. Closer analysis of the results in Fig. 9 also supports our previous observation that the convergence rate for first-order PC expansions is larger than for the second-order case, but differences are once again small.

As shown in Fig. 8, decreasing L_c results in a broader eigenvalue spectrum, which raises the question whether $N = 20$ KL modes is sufficient to capture all the relevant scales of λ . This question arises because truncation of the KL expansion removes the highest spatial frequencies, and leads to under-estimation of the variance. To verify that the near collapse of the curves in Fig. 9 with decreasing L_c is not due to such truncation, simulations were repeated using a first-order PC expansion and a higher number of KL modes, $N = 105$. The results (not shown) exhibit essentially the same convergence rate as with $N = 20$. This indicates that the truncation of the KL expansion does not affect the convergence rate of the multigrid solver.

6.4. Selection of multigrid parameters

The computational tests above were performed with fixed MG parameters, which enabled direct comparison between various cases and thus simplified the analysis. It is evident, however, that tuning these parameters can improve the efficiency of the method. For the test cases in the previous sections, selecting $\omega \in [1.2, 1.7]$ results in convergent iterations, but varying ω within this range affected the convergence rate. Specifically, for fixed tolerance on the peak residual (10^{-10}), the number of cycles needed to achieve this level varied between 3 and 5 cycles. Consequently, tests should generally be conducted in order to select the optimal ω value for the problem at hand. A similar optimization process should also be conducted for proper selection of the number of inner and outer iterations. Clearly, using a large number of outer iterations results in an inefficient method, since one does not want to perform a large number of outer iterations on the initial grid level. At the same time, a minimal number of outer iterations is required during prolongation in order to smooth the solution sufficiently before switching to the next grid level. Thus, N_{ou} should be carefully optimized. Meanwhile, N_{in} should be set to the minimum value above which the convergence rate starts to decrease.

Lastly, the efficiency of the multigrid procedure can also be drastically improved by designing cycles with a more complex structure than the simple V-cycle used in the present work. To illustrate the improvement

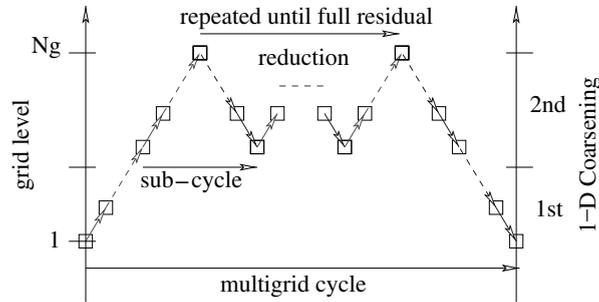


Fig. 10. Example of line-coarsening strategy for highly stretched grids. Instead of the usual V-cycle, the coarsening is first performed along the well resolved direction, and next in the second direction. Treatment (sub-cycles) of the second direction is repeated until the residual is reduced to the selected tolerance level.

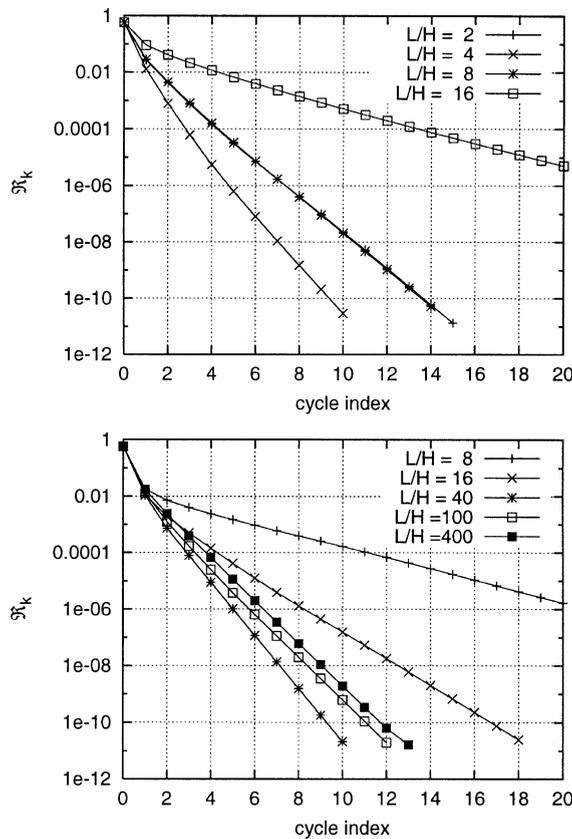


Fig. 11. Convergence rate of MG iterations: V-cycle approach (top) and line coarsening strategy (bottom). Curves are generated for solutions obtained in domains with different aspect ratio, L/H . In both sets of simulations, $N_x = 128$, $N_y = 32$, $N_{out} = 3$, $N_{in} = 2$, $\omega = 1.5$, $N_0 = 2$ and $N = 3$. The stochastic diffusivity field has $L_c = L$ and $COV = 0.1$.

that can be achieved by adaptation of cycle structure, a line-coarsening strategy, designed for highly stretched grids and/or domains with high aspect ratios, is briefly outlined below. Assuming that $\Delta x \gg \Delta y$,

the strategy consists of (i) performing a 1D-coarsening along the y -direction only, which eventually leads to a quasi-1D problem in x ; (ii) applying a 1D MG approach in the x -direction, which is iterated until the overall tolerance level is reached; and (iii) performing a prolongation in the y -direction only. This cycle, whose structure is schematically illustrated in Fig. 10, is repeated until the residual on the original fine grid drops below the desired tolerance level.

In Fig. 11 we contrast the convergence rates of the MG scheme using the V-cycles and of the adapted MG scheme outlined above. In both cases, the number of grid cells is fixed, $N_x = 128$ and $N_y = 32$, but the aspect ratio, L/H , of the domain is varied. Note that the case $L/H = 4$ corresponds to a grid with square cells, i.e. $\Delta x = \Delta y$. For the V-cycle iterations, the results indicate that the convergence rate deteriorates as the cell aspect ratio increases. Meanwhile, with the line-coarsening strategy, the convergence rate improves as L/H is increased from 8 to 40; for higher aspect ratios, up to $L/H = 400$, the convergence rate decreases slightly, but remains at a satisfactory level. In contrast, for such high aspect ratios, the regular V-cycle iterations are quite inefficient.

7. Conclusions

A multigrid scheme for the simulation of steady and unsteady stochastic diffusion equations was developed, and computational tests were conducted to analyze its behavior. In particular, these tests show that:

- The MG scheme exhibits a fast rate of convergence and good scalability with respect to spatial resolution in a fixed domain.
- The convergence rate is independent of the number of stochastic dimensions as long as the variance of the diffusivity field is held fixed.
- The convergence rate drops slowly as the order of the PC expansion increases.
- The convergence rate deteriorates substantially as the variance of the stochastic diffusivity field becomes large, but satisfactory convergence rates are still observed for COV up to 0.2. On the other hand, the convergence rate decreases slightly as the correlation length of the stochastic diffusivity field is decreased.
- Selection of MG parameters and cycle structure can drastically affect the efficiency of the iterations. These parameters should therefore be carefully optimized.

Future work will aim at enhancing the present MG approach, particularly for situations involving large variance and high order PC expansions. An attractive approach that is currently being explored is based on exploiting the structure and sparsity of the stochastic system. This structure suggests a hierarchical iterative strategy [34], which has been successfully exploited in the context of stochastic finite elements. In particular, it appears that incorporation of such an approach into the present MG framework could lead to a substantial performance enhancement.

Acknowledgements

This work was supported by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory, Air Force Material Command, USAF, under agreement number F30602-00-2-0612. The US government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Computations were performed at the National Center for Supercomputing Applications.

References

- [1] M. Abramowitz, L.A. Stegun, *Handbook of Mathematical Functions*, Dover, 1970.
- [2] I. Babuska, P. Chatzipantelidis, On solving elliptic stochastic partial differential equations, *Comput. Methods Appl. Mech. Engrg.* 191 (2002) 4093–4122.
- [3] I. Babuska, K.M. Liu, On solving stochastic initial-value differential equations, Technical report, TICAM Report 02-17, The University of Texas at Austin, 2002.
- [4] E. Bielewicz, J. Górski, Shells with random geometric imperfections simulation-based approach, *Int. J. Nonlinear Mech.* 37 (2002) 777–784.
- [5] R.H. Cameron, W.T. Martin, The orthogonal development of nonlinear functionals in series of fourier-hermite functionals, *Ann. Math.* 48 (1947) 385–392.
- [6] S.F. McCormick, *Multigrid Methods*, SIAM, 1987.
- [7] M.K. Deb, I.M. Babuska, J.T. Oden, Solution of stochastic partial differential equations using Galerkin finite element techniques, *Comput. Methods Appl. Mech. Engrg.* 190 (2001) 6359–6372.
- [8] B. Debusschere, H. Najm, A. Matta, O. Knio, R. Ghanem, O. Le Maître, Numerical simulation and quantitative uncertainty assessment of a reacting electrochemical microchannel, *Phys. Fluids* 15 (8) (2003) 2238–2250.
- [9] B. Debusschere, H. Najm, A. Matta, T. Shu, O. Knio, R. Ghanem, O. Le Maître, Uncertainty quantification in a reacting electrochemical microchannel flow model, in: 5th Int. Conf. on Mod. and Sim. of Microsystems, 2002.
- [10] O. Ditlevsen, N.J. Tarp-Johansen, Choice of input fields in stochastic finite-elements, *Probabilist. Engrg. Mech.* 14 (1999) 63–72.
- [11] C.A.J. Fletcher, *Computational Techniques for Fluid Dynamics*, vol. 1, Springer-Verlag, 1988.
- [12] R. Ghanem, Ingredients for a general purpose stochastic finite elements formulation, *Comput. Methods Appl. Mech. Engrg.* 168 (1999) 19–34.
- [13] R. Ghanem, The nonlinear gaussian spectrum of lognormal stochastic processes and variables, *ASME J. Appl. Mech.* 66 (1999) 964–973.
- [14] R. Ghanem, S. Dham, Stochastic finite element analysis for multiphase flow in heterogeneous porous media, *Transport Porous Med.* 32 (1998) 239–262.
- [15] R.G. Ghanem, Probabilistic characterization of transport in heterogeneous media, *Comput. Methods Appl. Mech. Engrg.* 158 (1998) 199–220.
- [16] R.G. Ghanem, P.D. Spanos, A spectral stochastic finite element formulation for reliability analysis, *J. Engrg. Mech. ASCE* 117 (1991) 2351–2372.
- [17] R.G. Ghanem, P.D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Springer Verlag, 1991.
- [18] H. Grudmann, H. Waubke, Non-linear stochastic dynamics of systems with random properties, a spectral approach combined with statistical linearization, *Int. J. Nonlinear Mech.* 31 (5) (1996) 619–630.
- [19] T.D. Hien, M. Kleiber, Stochastic finite element modelling in linear transient heat transfer, *Comput. Methods Appl. Mech. Engrg.* 144 (1997) 111–124.
- [20] H. Holden, B. Oksendal, J. Ubøe, T. Zhang, *Stochastic Partial Differential Equations*, Birkhauser, 1996.
- [21] M. Kaminski, T.D. Hien, Stochastic finite element modeling of trans-scient heat transfer in layered composites, *Int. Commun. Heat Mass* 26 (6) (1999) 801–810.
- [22] O.M. Knio, R.G. Ghanem, Polynomial Chaos product and moment formulas: A user utility, Technical report, Johns Hopkins University, 2001.
- [23] O.P. Le Maître, O.M. Knio, H.N. Najm, R.G. Ghanem, A stochastic projection method for fluid flow. I. Basic formulation, *J. Comput. Phys.* 173 (2001) 481–511.
- [24] O.P. Le Maître, M.T. Reagan, H.N. Najm, R.G. Ghanem, O.M. Knio, A stochastic projection method for fluid flow. II. Random process, *J. Comput. Phys.* 181 (2002) 9–44.
- [25] J.S. Liu, *Monte Carlo Strategies in Scientific Computing*, Springer, 2001.
- [26] N. Liu, B. Hu, Z.-W. Yu, Stochastic finite element method for random temperature in concrete structures, *Int. J. Solids Struct.* 38 (2001) 6965–6983.
- [27] M. Loève, *Probability Theory*, Springer, 1977.
- [28] H.G. Matthies, C.E. Brenner, C.G. Bucher, C.G. Scars, Uncertainties in probabilistic numerical analysis of structures and solids—stochastic finite-elements, *Struct. Saf.* 19 (3) (1997) 283–336.
- [29] H.G. Matthies, C.G. Bucher, Finite element for stochastic media problems, *Comput. Methods Appl. Mech. Engrg.* 168 (1999) 3–17.
- [30] H.G. Matthies, A. Keese, Multilevel solvers for the analysis of stochastic system, in: *First MIT Conference on Computational Fluid and Solid Mechanics*, Elsevier, 2001, pp. 1620–1622.
- [31] H.G. Matthies, A. Keese, Multilevel methods for stochastic systems, in: *Proceedings of the 2nd European Conference on Computational Mechanics*, Cracow, Poland, June 26, 2001.

- [32] J. Náprstek, Strongly non-linear stochastic response of a system with random initial imperfections, *Probabilist. Engrg. Mech.* 14 (1999) 141–148.
- [33] B.M. Nicola, B. Verlinden, A. Beuselincx, P. Jancsok, V. Quenon, N. Scheerlinck, P. Verbosen, J. de Baerdemaeker, Propagation of stochastic temperature fluctuations in refrigerated fruits, *Int. J. Refrig.* 22 (1999) 81–90.
- [34] M.F. Pellissetti, R.G. Ghanem, Iterative solution of systems of linear equations arising in the context of stochastic finite elements, *Adv. Engrg. Software* 31 (2000) 607–616.
- [35] R.F. Probstein, *Physicochemical Hydrodynamics*, Wiley, 1995.
- [36] G.I. Schuëller, Computational stochastic mechanics—recent advances, *Comput. Struct.* 79 (2001) 2225–2234.
- [37] A. Sluzalec, Random heat flow with phase change, *Int. J. Heat Mass Transfer* 43 (2000) 2303–2312.
- [38] A. Sluzalec, Stochastic finite element analysis of two-dimensional eddy current problems, *Appl. Math. Model.* 24 (2000) 401–406.
- [39] J. Stoer, R. Bulirsch, *Introduction to Numerical Analysis*, Springer, 1991.
- [40] T.G. Theting, Solving Wick-stochastic boundary value problems using a finite element method, *Stoch. Stoch. Rep.* 70 (2000) 241–270.
- [41] U. Trottenberg, C. Oosterlee, A. Schüller, *Multigrid*, Academic Press, 2001.
- [42] S. Wiener, The homogeneous chaos, *Amer. J. Math.* 60 (1938) 897–936.
- [43] D. Xiu, D. Lucor, G.E. Karniadakis, Modeling uncertainty in flow-structure interactions, in: *First MIT conference on Computational Fluid and Solid Mechanics*, Elsevier, 2001, pp. 1420–1423.