



Institut für Numerische Simulation

Rheinische Friedrich-Wilhelms-Universität Bonn

Friedrich-Hirzebruch-Allee 7 • 53115 Bonn • Germany
phone +49 228 73-69828 • fax +49 228 73-69847
www.ins.uni-bonn.de

U. Seidler, M. Griebel

**A Dimension-adaptive Combination Technique
for Uncertainty Quantification**

INS Preprint No. 2205

April 2022

A Dimension-adaptive Combination Technique for Uncertainty Quantification

Uta Seidler ^{*} and Michael Griebel ^{*†}

Abstract. We present an adaptive algorithm for the computation of quantities of interest involving the solution of a stochastic elliptic PDE where the diffusion coefficient is parametrized by means of a Karhunen-Loève expansion. The approximation of the equivalent parametric problem requires a restriction of the countably infinite-dimensional parameter space to a finite-dimensional parameter set, a spatial discretization and an approximation in the parametric variables. We consider a sparse grid approach between these approximation directions in order to reduce the computational effort and propose a dimension-adaptive combination technique. In addition, a sparse grid quadrature for the high-dimensional parametric approximation is employed and simultaneously balanced with the spatial and stochastic approximation. Our adaptive algorithm constructs a sparse grid approximation based on the benefit-cost ratio such that the regularity and thus the decay of the Karhunen-Loève coefficients is not required beforehand. The decay is detected and exploited as the algorithm adjusts to the anisotropy in the parametric variables. We include numerical examples for the Darcy problem with a lognormal permeability field, which illustrate a good performance of the algorithm: For sufficiently smooth random fields, we essentially recover the rate of the spatial discretization as asymptotic convergence rate with respect to the computational cost.

Key words. parametric PDEs, high-dimensional approximation, uncertainty quantification, lognormal diffusion coefficient, truncated Karhunen-Loève expansion, anisotropic sparse approximation, adaptive sparse grids, combination technique

AMS subject classifications. 65N35, 65C20, 65D30, 65N30, 60H35

1. Introduction. Mathematical models and computational simulations are widely used to describe systems in the natural sciences and in engineering. Often the input parameters of the models are uncertain. In order to compute quantities of interest and to predict the behavior of a system, it is important to propagate such uncertainties through the model. This task can be high-dimensional and can thus be computationally challenging due to the curse of dimensionality.

In this paper we present an adaptive algorithm for the computation of quantities of interest for problems in uncertainty quantification which is based on the combination technique of the sparse grid method. To this end, as a model problem, we focus on the elliptic PDE

$$(1.1) \quad -\operatorname{div}(a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}) \quad \text{in } D \subset \mathbb{R}^n,$$

which holds for any realization ω in the sample space Ω . The spatial domain is denoted by $D \subset \mathbb{R}^n$, $n = 1, 2, 3$ and the derivatives are with respect to the spatial variable \mathbf{x} . The uncertainty in the model is introduced by the diffusion coefficient $a(\mathbf{x}, \omega)$ which is modeled by a random field.

^{*}Institute for Numerical Simulation, University of Bonn, Friedrich-Hirzebruch-Allee 7, 53115 Bonn, Germany
griebel@ins.uni-bonn.de, seidler@ins.uni-bonn.de

[†]Fraunhofer Institute for Algorithms and Scientific Computing SCAI, Schloss Birlinghoven, 53754 SanktAugustin, Germany

Given information about the distribution of $a(\mathbf{x}, \omega)$, our goal is to obtain statistical information of the solution to the PDE (1.1). In particular, we are interested in $\mathbb{E}[\mathcal{F}(u)]$ where \mathcal{F} either denotes a continuous linear functional $\mathcal{F}: H_0^1(D) \rightarrow \mathbb{R}$ or $\mathcal{F}(u) = u^p$ with $p = 1, 2$ for the computation of the first and second moment.

Although this elliptic PDE is a rather simple model, it is relevant in practice as it is in the core of many application problems ranging from thermal analysis in engineering over diffusion processes in biology and chemistry to subsurface and groundwater flow problems in hydrology. Moreover, for water management and in environmental and energy problems (see, e.g. [6]), it is important to understand the propagation of uncertainty through the flow model. There, the permeability of the medium under consideration is usually unknown and varies from point to point in the spatial domain D .

Equation (1.1) is often studied either with an affine (e.g. [21, 29, 37]) or a lognormal diffusion coefficient (e.g. [16, 28]). In this paper, we concentrate on the case of a lognormal random field, where its logarithm is a Gaussian field characterized by a mean and a covariance function and which captures a large range of values. Moreover, we consider a parametrized version of (1.1) where the stochastic dependence of the random field is expressed in terms of a sequence of random variables. This is obtained by a spectral expansion of the random field (see, e.g. [36]) such that the problem depends on a countable set of stochastic parameters $\mathbf{y} = (y_1, y_2, y_3, \dots)$, with $y_k \in \mathbb{R}$, for $k \in \mathbb{N}$. More precisely, we consider the logarithm of the diffusion coefficient represented in the Karhunen-Loève expansion (KL). Although this might not be always optimal as shown in [4], it provides a straightforward parametrization with independent random variables. The KL expansion is usually not known explicitly and needs to be truncated and further approximated. When moments are known, this requires solving an eigenvalue problem. Otherwise, the field needs to be obtained from discretized samples, compare also [22]. Either way, the exact knowledge of the eigenvalues of the KL expansion might not always be available.

In order to compute the quantities of interest numerically, several approximation steps need to be applied. First, the random diffusion coefficient a has to be replaced by a finite-dimensional approximation. For this, the Karhunen-Loève expansion is truncated after m terms where m can be large and $a(\mathbf{x}, \mathbf{y})$ and $u(\mathbf{x}, \mathbf{y})$ are replaced by the associated approximations $a_m(\mathbf{x}, \mathbf{y}_m)$ and $u_m(\mathbf{x}, \mathbf{y}_m)$, respectively, with $\mathbf{y}_m = (y_1, \dots, y_m)$. Furthermore, the expectation $\mathbb{E}[\mathcal{F}(u)]$ is replaced by an m -dimensional integral on $\mathcal{F}(u_m)$ with a proper m -dimensional measure. This high-dimensional integration problem needs then to be approximated by a quadrature rule. To this end, methods such as stochastic collocation or Monte Carlo methods rely on a combination of integrand evaluations on a set of deterministic points or random samples $\mathbf{y}_m^{(i)}$, $i = 1, \dots, N$. In either case, the PDE needs to be solved repeatedly for these different parameter values $\mathbf{y}_m^{(i)}$. This requires a third discretization in the spatial variable which approximates the solution of the corresponding deterministic PDEs for these fixed parameter values. For the spatial approximation, we consider in this work a finite element method with piecewise linear elements since the solution to (1.1) for fixed \mathbf{y}_m is in $H^2(D)$ for sufficiently smooth diffusion coefficients [13, 20] and a right hand side in $L^2(D)$. For the computation of the m -dimensional expectation, Monte Carlo or Quasi Monte Carlo methods are often applied [13, 20, 26, 38] as they exhibit convergence rates with respect to

the sample size that are independent or nearly independent of the dimension of the stochastic parameter space. Their convergence rates, however, are limited by a rate of $1/2$ and, depending on m , a rate of a bit less than 1, respectively, even if the solution is smooth with respect to the random parameters. In fact, it has been shown in [4, 26, 27] that the dependence of the solution $u(\mathbf{x}, \mathbf{y})$ on the parametric variables given by the Karhunen-Loève expansion is analytic in the lognormal case. Furthermore, the influence of each stochastic parameter decreases due to the decay of the eigenvalues in the KL expansion which enters the measure weights in the m -dimensional integration problem. However, these favorable regularity properties of the integrand cannot be exploited by conventional Monte Carlo or Quasi Monte Carlo quadratures.

In this work, we therefore choose for the parametric discretization a sparse grid type quadrature method based on univariate Gauss-Hermite quadrature rules such that exponential convergence can be obtained. Moreover, we exploit the prevailing anisotropy in the stochastic parameters and apply a decay-obeying, dimension-adaptive sparse grid quadrature which has already been studied in several works [3, 24, 31, 37]. This allows to harvest the favorable regularity properties of the integrand present in our specific problem.

Altogether, we propose in this paper an adaptive algorithm that combines these different approximation steps, i.e., truncation, dimension-adaptive sparse grid quadrature, and finite element discretization in a multilevel manner: To this end, in addition to the dimension-adaptive sparse grid quadrature, we also include the spatial approximation in our sparse grid approach and we do not fix the truncation level of the KL expansion beforehand. Similar to multilevel Monte Carlo methods [5, 21, 25, 28, 33, 37] which involve different discretization levels for the spatial approximation and the quadrature, this incorporation of the spatial approximation allows us to balance the finite element method with the parametric discretization and to benefit from possible mixed regularity between the spatial and parametric variables. Here, however, in contrast to Monte Carlo and Quasi Monte Carlo methods, we have more than just one index to describe the parametric approximation as we use a sparse grid quadrature. Moreover, the parametric dimensions depend on the truncation level which is not fixed. To overcome this issue, we embed the parameters into an infinite-dimensional space and view the truncation as fixing higher parameters to the constant value $y_{m+j} = 0$ for $j \geq 1$ which allows to easily modify the number of univariate quadrature rules. Overall, our final sparse grid approach balances simultaneously the error and cost of each univariate quadrature with the spatial approximation and adaptively adjusts the truncation level. To this end, we use an adaptive variant of the well-known sparse grid combination technique [23]. This has the advantage that we can employ straightforwardly standard finite element methods and quadratures. We modify the conventional dimension-adaptive combination technique [19] to match our spatial and active parametric quadrature approximation directions. The construction of the adaptive sparse grid structure relies on suitable reliable indicators which are based on the benefit-cost ratio, compare e.g. [9]. Therefore, different to an a-priori construction, knowledge of the regularity of the solution u and of explicit decay rates of the KL eigenfunctions and eigenvalues is not necessary. In fact, it suffices to be able to approximatively evaluate the KL truncation and to assume that the terms of the KL are sorted in decreasing order. This is similar to the well-known saturation assumption of adaptive finite element methods. Consequently, our algorithm can be applied to different types of diffusion coefficients without any

changes. It properly balances the associated truncation error, the quadrature error and the finite element discretization error in a cost-versus-accuracy effective multilevel manner and adjusts itself to the anisotropy in the parametric variables.

The remaining part of this paper is structured as follows: We start by stating the considered problem and its parametrized version more precisely in Section 2 and continue with a separate description of each approximation step in Section 3. Subsequently, in Section 4, we focus on the general construction of sparse grids based on hierarchical surpluses and the combination technique, as well as the difference operators for our approximation. The dimension-adaptive sparse grid combination algorithm is then described in Section 4.2. In Section 6, we present numerical examples which show the good convergence properties of our algorithm.

2. Problem setting. Let $(\Omega, \Sigma, \mathbb{P})$ be a complete separable probability space. As mentioned in the introduction, we want to compute the quantities of interest

$$\mathbb{E}[\mathcal{F}(u)] = \int_{\Omega} \mathcal{F}(u(\mathbf{x}, \omega)) d\mathbb{P}(\omega)$$

of the solution $u(\mathbf{x}, \omega)$ to the random PDE

$$-\operatorname{div}(a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}) \quad \text{for } \mathbf{x} \in D \quad \text{for a.e. } \omega \in \Omega$$

with $f \in L^2(D)$ and zero Dirichlet boundary conditions for simplicity of notation. We assume that $a(\mathbf{x}, \omega)$ is a lognormal random field, i.e.,

$$(2.1) \quad a(\mathbf{x}, \omega) = \exp(b(\mathbf{x}, \omega))$$

where $b(\mathbf{x}, \omega)$ is a Gaussian field expressed in form of a Karhunen-Loève expansion

$$(2.2) \quad b(\mathbf{x}, \omega) = \mathbb{E}[\log(a(\mathbf{x}))] + \sum_{k=1}^{\infty} \sqrt{\lambda_k} \psi_k(\mathbf{x}) y_k(\omega)$$

with $\lambda_1 \geq \lambda_2 \geq \dots$. The functions $\{\psi_k\}$ are pairwise orthonormal functions in $L^2(D)$ and y_k are independent standard normally distributed random variables, $y_k \sim \mathcal{N}(0, 1)$, for which we denote with ρ_k the standard Gaussian density function. In general, this expansion of b can be obtained by computing the eigenpairs of the operator

$$(Cv)(\mathbf{x}) = \int_D \operatorname{cov}_b(\mathbf{x}, \mathbf{x}') v(\mathbf{x}') d\mathbf{x}',$$

whereby $\operatorname{cov}_b(\mathbf{x}, \mathbf{x}')$ is the continuous covariance function of the Gaussian field. Alternatively, it must be appropriately approximated from measurements. In this paper we do not require explicit knowledge of the KL eigenpairs and their decay rates, but we just need to be able to evaluate truncated and approximated versions of (2.2).

Instead of the PDE (1.1) in the abstract probability space, we consider a parametrized version in terms of the variables y_k . We define the measure $\mathbb{P}_y = \mathbf{y}_{\#} \mathbb{P}$ as the pushforward of the probability measure on Ω through the map

$$\mathbf{y} : \Omega \rightarrow \mathbb{R}^{\mathbb{N}}, \quad \omega \mapsto (y_1(\omega), y_2(\omega), \dots)$$

and work on the probability space $(\mathbb{R}^{\mathbb{N}}, \mathcal{B}(\mathbb{R}^{\mathbb{N}}), \mu_\rho)$ with μ_ρ being the tensorized Gaussian probability measure with density $\rho(\mathbf{y}) = \prod_{k=1}^{\infty} \rho_k(y_k)$.

The dependence of $a(\mathbf{x}, \omega)$ on ω is substituted by a dependence on the parameters $\mathbf{y} = (y_1, y_2, \dots)$ and we write for the parametrized diffusion coefficient

$$(2.3) \quad a(\mathbf{x}, \mathbf{y}) = a_0(\mathbf{x}) \exp \left(\sum_{k=1}^{\infty} \sqrt{\lambda_k} \psi_k(\mathbf{x}) y_k \right).$$

The parametrized version of the PDE is then for a fixed parameter $\mathbf{y} \in \mathbb{R}^{\mathbb{N}}$ given by

$$(2.4) \quad -\operatorname{div}(a(\mathbf{x}, \mathbf{y}) \nabla u(\mathbf{x}, \mathbf{y})) = f(\mathbf{x}) \quad \mathbf{x} \in D.$$

For well-posedness of the problem, we need to ensure that the series in (2.3) converges in $L^\infty(D)$ almost surely and that the moments are well-defined, i.e., $u \in L_\rho^2(\mathbf{\Gamma}, H_0^1(D))$. To this end, we assume that there exists a measurable set $\mathbf{\Gamma} \subset \mathbb{N}^\infty$ such that $\mu_\rho(\mathbf{\Gamma}) = 1$ and $a(\mathbf{x}, \mathbf{y})$ converges in $L^\infty(D)$ for all $\mathbf{y} \in \mathbf{\Gamma}$. Furthermore, we define $a_{\min}(\mathbf{y}) = \operatorname{ess\,inf}_{\mathbf{x} \in D} a(\mathbf{x}, \mathbf{y})$ and $\operatorname{ess\,sup}_{\mathbf{x} \in D} a(\mathbf{x}, \mathbf{y}) = a_{\max}(\mathbf{y})$ and assume that $a_{\max} \in L_\rho^p(\mathbf{\Gamma})$ and $1/a_{\min} \in L_\rho^p(\mathbf{\Gamma})$ for any $p \in (0, \infty)$. By the Lax-Milgram lemma, (2.4) admits for $\mathbf{y} \in \mathbf{\Gamma}$ a unique solution $u(\mathbf{y}) \in H_0^1(D)$ and satisfies

$$\|u(\mathbf{y})\|_{H_0^1(D)} \leq \frac{C}{a_{\min}(\mathbf{y})} \|f\|_{L^2(D)}.$$

The expectation for the computation of the quantities of interest corresponds in the parametrized formulation then to the evaluation of the integral

$$(2.5) \quad \mathbb{E}[\mathcal{F}(u(\mathbf{y}(\omega)))] = \int_{\Omega} \mathcal{F}(u(\mathbf{y}(\omega))) d\mathbb{P}(\omega) = \mathbb{E}_{\mathbf{y}}[\mathcal{F}(u(\mathbf{y}))] = \int_{\mathbf{\Gamma}} \mathcal{F}(u(\mathbf{y})) \rho(\mathbf{y}) d\mathbf{y}$$

which is well-defined for $\mathcal{F}(u) = u^p$, $p = 1, 2$ or for continuous linear functionals $\mathcal{F} : H_0^1(D) \rightarrow \mathbb{R}$. In the case where the KL eigenfunctions and -values are known explicitly, these assumptions can be ensured if the sequence

$$(2.6) \quad \gamma_k = \sqrt{\lambda_k} \|\psi_k\|_{L^\infty(D)}$$

is summable, i.e., $\{\gamma_k\}_k \in \ell^1(\mathbb{N})$. In this case the set $\mathbf{\Gamma}$ is given by

$$\mathbf{\Gamma} = \left\{ \mathbf{y} \in \mathbb{R}^{\mathbb{N}} : \sum_{k=1}^{\infty} \gamma_k |y_k| < \infty \right\},$$

which is μ_ρ -measurable with $\mu_\rho(\mathbf{\Gamma}) = 1$ and $1/a_{\min}, a_{\max} \in L_\rho^p(\mathbf{\Gamma})$ for all $p \in (0, \infty)$ as shown in [34]. However, note that also weaker assumptions are possible to ensure the well-posedness of the parametric problem (see, e.g. [4, 11]) and we do not necessarily require summability of (2.6) in the following.

3. Numerical approximations. This section focuses on the methods needed to numerically compute an approximation of (2.5). In the following each approximation step is addressed separately. We give a basic idea of the convergence behavior that can be expected if sufficient regularity is provided, but refer to the literature for precise convergence rates and the regularity assumptions needed for a precise convergence order. The combination of the different methods will be discussed in the later sections.

3.1. Approximation by a finite-dimensional parameter set. First, since one can in practice not work with the infinite representation (2.2) of the random diffusion field a , we approximate a by its finite noise truncation, i.e., we cut off the series after m terms. This way, we obtain

$$a_m(\mathbf{x}, \mathbf{y}_m) = a_0(\mathbf{x}) \exp \left(\sum_{k=1}^m \sqrt{\lambda_k} \psi_k(\mathbf{x}) y_k \right),$$

which depends on a finite but possibly high-dimensional parameter set. Furthermore, we denote with u_m the solution of the PDE where the diffusion coefficient a is replaced by a_m , i.e., for any $\mathbf{y}_m \in \mathbb{R}^m$, the function $u_m(\mathbf{y}_m)$ solves

$$-\operatorname{div}(a_m(\mathbf{y}_m) \nabla u_m(\mathbf{y}_m)) = f.$$

For later purpose, we remark that the m -term truncation of the random diffusion coefficient a is equivalent to replacing the dependence of y_i for $i > m$ by the constant 0. Thus,

$$(3.1) \quad a_m(\mathbf{x}, \mathbf{y}_m) = a(\mathbf{x}, y_1, \dots, y_m, 0, 0, \dots) \quad \text{and} \quad u_m(\mathbf{x}, \mathbf{y}_m) = u(\mathbf{x}, y_1, \dots, y_m, 0, 0, \dots).$$

Note here that for $(y_1, \dots, y_m, 0, 0, \dots)$ the series in (2.3) converges as it consists only of finitely many terms and the solution of the associated PDE is well-defined.

Clearly, the error introduced by this truncation decreases with increasing m . The rate depends on the smoothness of the random field a which is reflected in the decay rates of the eigenvalues and eigenfunctions of its Karhunen-Loève expansion [35]. As shown in [11, 12, 13, 20] also the solution u_m converges to u and error estimates for both the strong and weak error are provided. These bounds depend on the decay properties of the coefficients in the KL expansion and require decay rates of λ_k and ψ_k in norms different to the $L^2(D)$ -norm which are however in general difficult to obtain. But note here that for our adaptive algorithm in Section 4.2 we do not require these bounds to be available beforehand. Instead we will observe that our algorithm adjusts to different decay rates of the KL coefficients and determines the required truncation level automatically. Hence, if the regularity of the random field a is a-priorily unknown, we can still apply our algorithm in an a-posteriori fashion for a single specific a in a similar way as adaptive finite element methods do.

3.2. Finite element method. Next, we address the spatial discretization. For any fixed parameter $\mathbf{y}_m^{(i)}, i = 1, \dots, N$, the solution $u_m(\cdot, \mathbf{y}_m^{(i)}) \in H_0^1(D)$ to the PDE (2.4) needs to be approximated. For this, we apply a linear finite element method and subdivide the domain D by a uniform and regular triangulation \mathcal{T}_h with diameter h . Note here that we deliberately stick to uniform triangulations and do not consider spatial adaptivity since we are in this

article merely interested in the interplay of truncation, plain discretization and quadrature. For each $\mathbf{y} \in \Gamma$, we denote with the index h the finite element solution $u_h(\cdot, \mathbf{y}) \in \mathcal{V}_h$ that satisfies

$$\int_D a(\mathbf{x}, \mathbf{y}) \nabla u_h(\mathbf{x}, \mathbf{y}) \nabla v_h(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x}$$

for all $v_h \in \mathcal{V}_h = \{v \in C(D) : v|_{\partial D} = 0 \text{ and } v|_K \in \mathcal{P}_1 \ \forall K \in \mathcal{T}_h\}$ where \mathcal{P}_1 is the space of linear functions and we denote as $u_{h,m}(\cdot, \mathbf{y}_m) \in \mathcal{V}_h$ the finite element solution of the truncated problem, i.e.,

$$\int_D a_m(\mathbf{x}, \mathbf{y}_m) \nabla u_{h,m}(\mathbf{x}, \mathbf{y}_m) \nabla v_h(\mathbf{x}) d\mathbf{x} = \int_D f(\mathbf{x}) v_h(\mathbf{x}) d\mathbf{x}.$$

For diffusion coefficients with sufficiently smooth realizations, the PDE solution satisfies $u(\cdot, \mathbf{y}) \in H^2(D)$ a.s. and the finite element approximation u_h exhibits an error rate of the form

$$\|u(\mathbf{x}, \mathbf{y}) - u_h(\mathbf{x}, \mathbf{y})\|_{H^1(D)} \leq Ch \sqrt{\frac{a_{\max}(\mathbf{y})}{a_{\min}(\mathbf{y})}} \|u(\mathbf{x}, \mathbf{y})\|_{H^2(D)}.$$

For any $q < \infty$, this implies $\|u - u_h\|_{L^q_\rho(\Gamma, H^1_0(D))} \leq Ch$ and it also implies $\|u - u_h\|_{L^q_\rho(\Gamma, L^2(D))} \leq Ch^2$ for the error in the spatial variable measured in the L^2 -norm. These bounds also hold for the error $u_m - u_{h,m}$ with constants independent of m . For the computation of the second moment, the spatial norms are replaced by the $W^{1,1}(D)$ -norm and the $L^1(D)$ -norm, respectively. Hence, depending on the quantity of interest and the error norm, the error bound is of the order $\mathcal{O}(h)$ or $\mathcal{O}(h^2)$. For less regular diffusion coefficient without continuously differentiable realizations, the rate reduces to $\mathcal{O}(h^s)$ and $\mathcal{O}(h^{2s})$, respectively, where $0 < s \leq 1$ depends on the regularity of the random field. We refer to [13, 20, 38] for more details on the spatial discretization error.

3.3. Quadrature method. As a third approximation step, we need to approximate the solution in the parametric variables and evaluate the integral

$$(3.2) \quad \mathbf{I}v = \int_{\mathbb{R}^m} v(\mathbf{y}_m) \rho(\mathbf{y}_m) d\mathbf{y}_m$$

where $v(\mathbf{y}_m) = \mathcal{F}(u_{h,m}(\mathbf{x}, \mathbf{y}_m))$. Note that, with the truncation of the Karhunen-Loève expansion after m terms, the integral is over the m -dimensional space \mathbb{R}^m . For this purpose, we consider quadrature rules of the type

$$\mathbf{I}v \approx \mathbf{Q}_N v = \sum_{i=1}^N w_i v(\mathbf{y}_m^{(i)})$$

which are based on the evaluation of the integrand at N quadrature points $\mathbf{y}_m^{(i)} \in \mathbb{R}^m$ that are taken into account with weight w_i , $i = 1, \dots, N$. As $u_{h,m}$ is also a function in \mathbf{x} , this

requires the FEM-discretization on mesh the \mathcal{T}_h and a solution of the resulting linear systems of equations for each of these N parameter values.

We apply a Gaussian quadrature rule that can exploit the regularity of the integrand with respect to the parameters y_1, y_2, \dots . For that, we combine m one-dimensional Gauss-Hermite quadrature rules as the underlying density function for each parameter y_k is the one-dimensional Gaussian. For each parameter y_k a univariate quadrature rule is considered which is based on the interpolation of the integrand among a set of deterministic points $\{y_k^{(i)}\}$ such that the highest degree of polynomial exactness is achieved. The quadrature points are chosen as the $N^{(k)}$ roots of the Hermite polynomials of degree $N^{(k)}$, which are orthogonal with respect to the inner product

$$(q, r)_{L^2_{\rho}(\mathbb{R})} = \int_{\mathbb{R}} q(y_k)r(y_k)\rho_k(y_k)dy_k,$$

and the quadrature weights are obtained by integration of the interpolating polynomial, i.e., $w_i = (L_i, 1)_{L^2_{\rho_k}(\mathbb{R})}$ where L_i denotes the i -th Lagrange polynomial with respect to the quadrature point set.

For the approximation of (3.2), we combine these m univariate Gauss-Hermite quadrature rules in a tensor product

$$\mathcal{Q}_{(N^{(1)}, N^{(2)}, \dots, N^{(m)})} v = Q_{N^{(1)}}^{(1)} \otimes Q_{N^{(2)}}^{(2)} \otimes \dots \otimes Q_{N^{(m)}}^{(m)} v$$

which requires the evaluation of v at the quadrature points $(y_1^{(i_1)}, y_2^{(i_2)}, \dots, y_m^{(i_m)})$ for $i_k = 1, \dots, N^{(k)}$ and $k = 1, \dots, m$. Here, $N = \prod_{k=1}^m N^{(k)}$.

If the integrand can be extended into a region in the complex plane, each univariate rule can achieve an exponential convergence rate. Hence, this choice of quadrature rule exploits the smooth dependence of the solution on the parameters y_k for $k = 1, \dots, m$, which is shown in [20, 26, 27]. However, the combination of the m quadrature rules in a full isotropic tensor product, i.e., with $N^{(1)} = N^{(2)} = \dots = N^{(m)}$, suffers from the curse of dimensionality as the number of quadrature points grows exponentially in the dimension m . Instead, an anisotropic version, as discussed in [3], can be used which combines the m rules with different approximation power and which can exploit the decreasing influence of the parameters y_k on the PDE solution u by choosing the numbers $N^{(k)}$ according to the anisotropy. If the decrease of influence of the stochastic parameters y_k is sufficiently fast, for example such that the errors of the univariate rules are balanced when $N^{(k)} = 2^{\alpha_k}$ with $\sum_{k=1}^{\infty} \alpha_k < \infty$, the curse of dimensionality can be broken. Moreover, we combine such an anisotropic approach with a sparse grid quadrature which leads to an anisotropic sparse grid quadrature similar to [24, 31]. The bounds in [3, 26, 27] on the derivatives of u with respect to the parameters \mathbf{y}_m justify that a sufficiently anisotropic sparse grid in the parametric variables is a sound approach. The construction of this anisotropic sparse grid quadrature is coupled with the KL-truncation and the spatial approximation which will become clearer in a moment.

4. Sparse grids. The adaptive algorithm that we propose is based on a sparse grid approach which balances the three different approximation methods. In the following, we give a short overview of the general construction of sparse grids in d dimensions with abstract operators before we address the adaptive combination technique for the computation of quantities

of interest. For more details on sparse grids we refer to [18] and [9]. The additional regularity between the spatial and parametric variables shown in [26, 28] for sufficiently smooth random fields justifies this approach.

For each discretization direction we consider a sequence of numerical approximation operators $\{P_l\}_l$ with increasing approximation power such that the limit leads to the true solution.

We define the difference operators between two approximations as

$$\Delta_l := \begin{cases} P_l - P_{l-1} & \text{if } l \geq 1, \\ P_0 & \text{if } l = 0, \end{cases}$$

such that P_L can be reconstructed from the difference operators up to L , i.e.,

$$P_L = \sum_{l=0}^L \Delta_l,$$

which is just a telescoping sum. For an approximation that involves d different approximation directions, we combine the d operators in a straightforward way and define the anisotropic full grid approximation as

$$P_{\mathbf{L}}u = \bigotimes_{i=1}^d P_{L_i}^{(i)}u = P_{L_1}^{(1)}(P_{L_2}^{(2)}(\dots P_{L_d}^{(d)}u)) = \sum_{\mathbf{l} \leq \mathbf{L}} \Delta_{\mathbf{l}}u$$

where $\mathbf{L} = (L_1, \dots, L_d)$ indicates the approximation levels in each direction. The d -dimensional difference operator is given by

$$\Delta_{\mathbf{l}} = \bigotimes_{i=1}^d \Delta_{l_i}^{(i)} \quad \text{for } \mathbf{l} = (l_1, \dots, l_d)$$

whereby the condition $\mathbf{l} \leq \mathbf{L}$ requires $l_i \leq L_i$ for all $i = 1, \dots, d$. Note at this point that the operators $P_{l_i}^{(i)}$ could in principle be of different type for different directions i , i.e., it could be a spatial energy projection in one direction and quadrature rules in the other directions, as it will be done later on in subsection 5.1.

A sparse grid approximation only includes differences for a subset $\mathcal{I} \subset \mathbb{N}^d$. For example, the regular sparse grid sums up the differences for $|\mathbf{l}|_1 \leq L + d - 1$, which yields a good approximation if the solution has some additional smoothness in form of bounded mixed derivatives. In the following, we apply a sparse grid with respect to a *general* index set \mathcal{I} and define the sparse grid approximation by

$$(4.1) \quad P_{\mathcal{I}}^{\text{SG}}u = \sum_{\mathbf{l} \in \mathcal{I}} \Delta_{\mathbf{l}}u.$$

The choice of the index set \mathcal{I} plays here a crucial role. The optimal choice clearly depends on the contribution of each increment $\Delta_{\mathbf{l}}u$ and its associated cost. In [9], it has been shown that, for a prescribed cost, the selection of the index set can be formulated as a binary knapsack problem to which a solution is obtained by including the indices with the highest

benefit-cost ratio, i.e., $\|\Delta_{\mathbf{l}}u\|$ in some norm divided by the associated computational cost. As a-priori estimates on the regularity of the solution and therefore on the size of $\Delta_{\mathbf{l}}u$ are not always available, we use the benefit-cost ratio to construct the index set in a greedy way in our adaptive algorithm later on (see Section 4.2).

4.1. Combination technique. For our algorithm we consider a somewhat different representation of the formula (4.1), the so-called combination technique, which was introduced in [23]. While efficient methods for the computation of the increments are necessary for (4.1), the combination technique relies on the combination of several anisotropic, full but low-order approximations.

The terms in (4.1) are rearranged such that full grid combinations are recovered and the combination technique approximation for (4.1) is defined as

$$(4.2) \quad P_{\mathcal{I}}^{\text{ct}}u = \sum_{\mathbf{l} \in \mathcal{I}} \alpha_{\mathbf{l}} P_{\mathbf{l}}u$$

with coefficients

$$\alpha_{\mathbf{l}} = \sum_{\mathbf{z}=\mathbf{0}}^{\mathbf{1}} (-1)^{|\mathbf{z}|_1} \chi^{\mathcal{I}}(\mathbf{l} + \mathbf{z}),$$

where $\mathbf{1} = (1, \dots, 1)$ and the characteristic function χ is given by

$$\chi^{\mathcal{I}}(\mathbf{l}) = \begin{cases} 1 & \text{if } \mathbf{l} \in \mathcal{I}, \\ 0 & \text{otherwise.} \end{cases}$$

In order to have an equivalence of the combination technique (4.2) and the sparse grid formulation based on the increments (4.1), we need to impose an admissibility condition on the index set \mathcal{I} as in [19]. We require that the index set \mathcal{I} is *downward closed*, meaning it satisfies the following property:

For any $\mathbf{l} \in \mathcal{I}$ it must hold that

$$(4.3) \quad \mathbf{l} - \mathbf{e}_j \in \mathcal{I} \quad \text{for } 1 \leq j \leq d \text{ with } l_j \geq 1$$

with \mathbf{e}_j being the j -th unit vector.

We will use the combination technique representation as it has the advantage that it relies on anisotropic, full grid approximations and does not require the computation of the increments $\Delta_{\mathbf{l}}$. Hence, no hierarchical structure needs to be assumed. Instead, the standard numerical methods discussed in Section 3 can be directly applied, such as standard finite element solvers and standard quadratures, and no nestedness of the quadrature methods is needed. Additionally, the structure of (4.2) is intrinsically parallel as the approximations $\{P_{\mathbf{l}}\}_{\mathbf{l}}$ can be computed independently. Note however that we will rely on the hierarchical operators $\Delta_{\mathbf{l}}$ for the adaptive construction of the sparse grid index set as explained later on.

4.2. Dimension-adaptive combination technique. Next, we describe the adaptive construction of the index set \mathcal{I} . In the following, we state the dimension-adaptive combination technique which was introduced for integration problems in [19] and which will be modified to fit our problem setting later on in Section 5.2. The basic algorithm with abstract operators P_l is stated in Algorithm 4.1. The dimension-adaptive algorithm heuristically constructs a sequence of admissible index sets $\mathcal{I}^{(1)} \subset \mathcal{I}^{(2)} \subset \dots \subset \mathcal{I}^{(t)}$ which indicate the increments that are included in the sparse grid formulation. Once such an index set has been found, the approximation is obtained using the generalized combination technique formula (4.2).

The algorithm starts off with one index in $\mathcal{I}^{(1)}$ and adds successively indices which do not destroy the admissibility of the index set. To ensure this, two index sets are introduced \mathcal{A} and \mathcal{O} with $\mathcal{I} = \mathcal{A} \cup \mathcal{O}$ and $\mathcal{A} \cap \mathcal{O} = \emptyset$. The set \mathcal{A} denotes the set of admissible neighbor indices and the set \mathcal{O} is referred to as the set of old indices. While the indices in \mathcal{O} have already been chosen to be considered for the sparse grid, the index set \mathcal{A} contains the indices in the neighborhood of \mathcal{O} , i.e., which result from a refinement along the axis, and which do not destroy the admissibility when added to \mathcal{O} .

In each iteration of the dimension-adaptive combination technique, the index with the highest contribution is selected among the indices in \mathcal{A} , removed from \mathcal{A} and added to the set \mathcal{O} . This selection process is based on the benefit-cost ratio [9] as mentioned before. Besides Δ_l , also the corresponding cost c_l is taken into account in the associated profit indicator η_l which will be explained in detail later on. As a result of this process, the neighborhood for the index set \mathcal{O} has changed such that the set \mathcal{A} needs to be updated by adding the indices to \mathcal{A} which still satisfy the admissibility condition. In a final step, the full grid solutions are combined according to (4.2). Since during the algorithm the solutions for $l \in \mathcal{A}$ were computed for the profit indicator, these indices are also included into the computation of the combination technique solution.

Algorithm 4.1 Dimension-adaptive combination technique [19]

Input: Tolerance $\varepsilon > 0$

Output: Index set \mathcal{I} and approximation $P_{\mathcal{I}}^{\text{ct}}u$

- 1: $l = (0, \dots, 0) \in \mathbb{N}^d$
- 2: $\mathcal{A} = \{l\}, \mathcal{O} = \emptyset$
- 3: compute local profit indicator η_l
- 4: $\eta = \eta_l$
- 5: **while** $\eta > \varepsilon$ **do**
- 6: Select the index $j \in \mathcal{A}$ with the largest profit η_j
- 7: $\mathcal{A} = \mathcal{A} \setminus \{j\}, \mathcal{O} = \mathcal{O} \cup \{j\}$
- 8: $\eta = \eta - \eta_j$
- 9: **for** $k = 1, \dots, d$ **do**
- 10: $l = j + e_k$
- 11: **if** $l - e_i \in \mathcal{O} \quad \forall 1 \leq i \leq d$ with $l_i > 0$ **then**
- 12: $\mathcal{A} = \mathcal{A} \cup \{l\}$
- 13: Compute full grid approximation $P_l u$
- 14: Compute local profit indicator η_l
- 15: $\eta = \eta + \eta_l$

16: **end if**
17: **end for**
18: **end while**
19: Combine approximations $\mathbf{P}_l u$ for $l \in \mathcal{I} = \mathcal{A} \cup \mathcal{O}$ according to (4.2) to obtain $\mathbf{P}_{\mathcal{I}}^{\text{ct}} u$
20: **return** $\mathcal{I} = \mathcal{A} \cup \mathcal{O}$ and $\mathbf{P}_{\mathcal{I}}^{\text{ct}} u$

5. Sparse Grids for Uncertainty Quantification. In the previous section, we had a fixed dimension d and described the basic so-called dimension-adaptive sparse grid approach for abstract approximation operators. Now, we generalize this approach to our uncertainty quantification problem and the computation of the quantities of interests (2.5). Here, the operators $P_{l_i}^{(i)}$ are of different type corresponding to the different approximation steps. The adaptation to uncertainty quantification (UQ) involves two main aspects: First, we employ a sparse grid for the quadrature of the parametric coordinates, but we also want to account for a possibly changing dimension to resolve a necessary dynamically adapted truncation length. This will be done via the embedding of the parametric coordinates and corresponding indices into infinite-dimensional spaces, as already indicated in (3.1). Second, we include the spatial discretization in the sparse grid formulation by adding a spatial index to also incorporate the treatment of the associated deterministic PDEs. In the following, we first define the difference operators for the UQ setting. We then present the generalization of the dimension-adaptive combination technique to this problem setting and in the last subsection we discuss the UQ-specific profit indicators with which we steer the overall adaptive construction.

5.1. Difference operators for UQ. For our setting, we have three numerical approximations that need to be properly taken into account. We now define the difference operators for these three approximation directions.

First, we consider a sequence of truncation operators $\{T_m\}_m$ which map the PDE solution to the solution of the truncated problem, i.e., $T_m : u \mapsto u_m$, and we define the difference between two levels by

$$\Delta_m^{(T)} u = \begin{cases} u_m - u_{m-1} & \text{if } m \geq 1, \\ u_0 & \text{if } m = 0. \end{cases}$$

Second, for the spatial discretization, we employ a sequence of uniform grids on D with width $h_{l_x} = \mathcal{O}(2^{-l_x})$ and consider the operator that maps the function u_m to the evaluation of \mathcal{F} of its finite element approximation $\mathcal{F}(u_{h_{l_x}, m})$. Due to the continuity of \mathcal{F} we have $\lim_{l_x \rightarrow \infty} \mathcal{F}(u_{h_{l_x}, m}) = \mathcal{F}(u_m)$. Here, in the spatial direction, we have not just a projection onto the finite element subspace (which itself involves the FEM-discretization and the solution of the arising linear system of equations), but also the corresponding evaluation required for the specific quantity of interest. Only in the case of computing the first moment, i.e., $\mathcal{F} = \text{Id}$, the operator corresponds to the projector onto the finite element space. Consequently, the difference between two spatial discretizations is defined as

$$(5.1) \quad \Delta_{l_x}^{(x)} u_m = \begin{cases} \mathcal{F}(u_{h_{l_x}, m}) - \mathcal{F}(u_{h_{l_x-1}, m}) & \text{if } l_x \geq 1, \\ \mathcal{F}(u_{h_0, m}) & \text{if } l_x = 0. \end{cases}$$

And finally, we approximate the expectation by a quadrature method for which we apply a sequence of quadrature rules that are themselves sparse grids. Here, given a truncation level m , we have m parametric variables. For each parameter y_k for $k = 1, \dots, m$, we consider a sequence of univariate Gauss-Hermite rules $\{Q_{l_k}^{(k)}\}_{l_k \geq 0}$ with increasing number of quadrature points $\{N_{l_k}^{(y_k)}\}_{l_k \geq 0}$ and associated differences operators

$$(5.2) \quad \Delta_{l_k}^{(Q_k)} = \begin{cases} Q_{l_k}^{(k)} - Q_{l_k-1}^{(k)} & \text{if } l_k \geq 1, \\ Q_0^{(k)} & \text{if } l_k = 0. \end{cases}$$

We combine the m one-dimensional difference operators and define the difference operator for the quadrature

$$(5.3) \quad \Delta_{\mathbf{l}}^{(Q)} = \bigotimes_{k=1}^m \Delta_{l_k}^{(Q_k)}$$

with \mathbf{l} being a m -dimensional array determining the level for each parametric variable. The summation of such differences over a set \mathcal{I}_Q of indices $\mathbf{l} = (l_1, \dots, l_m) \in \mathbb{N}^m$ then yields the sparse grid quadrature

$$\mathcal{Q}_{\mathcal{I}_Q}^{\text{SG}} v = \sum_{\mathbf{l} \in \mathcal{I}_Q} \Delta_{\mathbf{l}}^{(Q)} v.$$

In combination with the approximation of the random field $a(\mathbf{x}, \mathbf{y})$ and the spatial discretization, the integration is however not with respect to a fixed high-dimensional parameter space, but should rather adapt itself properly, which gives rise to a varying number of parametric dimensions. To this end, we exploit the fact that the m -term truncation of the KL expansion is equivalent to the constant evaluation at $y_k = 0$ for $k > m$ (c.f. (3.1)) and embed the parameters and quadrature rules into an infinite dimensional setting. A one-point Gauss-Hermite quadrature in y_k for $k > m$ performs the same approximation and thus replaces the integrand by just a constant function in y_k . We therefore assume $N_0^{(y_k)} = 1$, by choosing $N_{l_k}^{(y_k)} = 2^{l_k}$, and call the parameter dimension k inactive if $l_k = 0$. This lowest level of the univariate quadrature then corresponds to a constant approximation and disregards the dependence of the random field on this y_k . Conversely, the parameter dimension k is active when $l_k \geq 1$.

This allows us to represent the index indicating the quadrature levels by

$$(5.4) \quad \mathbf{l} = (l_1, l_2, \dots, l_m, 0, 0, \dots) \in \mathbb{N}^\infty,$$

which is in principle independent of the parameter dimension, possesses only finitely many non-zero entries, and altogether allows for straightforward adaptivity in the parametric dimensions, i.e., in the truncation length. The difference operator (5.3) is then equivalent to

$$(5.5) \quad \Delta_{\mathbf{l}}^{(Q)} v = \Delta_{(l_1, \dots, l_m, 0, 0, \dots)}^{(Q)} v = \bigotimes_{k=1}^{\infty} \Delta_{l_k}^{(Q_k)} v = \bigotimes_{k=1}^m \Delta_{l_k}^{(Q_k)} v$$

with the univariate differences as in (5.2). The equality holds because $Q_0^{(k)}v = v$ for $k > m$. Note also that the effective truncation level can be obtained from (5.4) and is given by $\max\{k \in \mathbb{N} : l_k \geq 1\}$.

In order to also include the spatial discretization and to combine it with the parametric discretization, we extend the index (5.4) by one entry which refers to the spatial discretization and write

$$(5.6) \quad \tilde{l} = (l_{\mathbf{x}}, l_{y_1}, \dots, l_{y_m}, 0, 0, \dots) \in \mathbb{N}^\infty.$$

Corresponding to this index, we set

$$\mathbf{P}_{\tilde{l}}u = Q_{l_{y_1}}^{(1)} \dots Q_{l_{y_m}}^{(m)} (\mathcal{F}(u_{h_{l_{\mathbf{x}}}, m})),$$

where $u_{h_{l_{\mathbf{x}}}, m}$ is the projection of u_m onto the FEM-space with triangulation $\mathcal{T}_{h_{l_{\mathbf{x}}}}$ (and involves thus discretization and solution of the resulting linear system of equations). Note here that the order of application of the quadrature, the application of \mathcal{F} , the FEM-projection (i.e., discretization and solution), and the truncation is not commutative and cannot be changed.

The generalized difference operator is then based on the difference operators of the different approximation directions

$$(5.7) \quad \Delta_{\tilde{l}}u = \Delta_l^{(Q)} \Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u,$$

where we use the notation $\Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u$ for

$$\Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u = \mathcal{F}(u_{h_{l_{\mathbf{x}}}, m}) - \mathcal{F}(u_{h_{l_{\mathbf{x}}-1}, m}) - \mathcal{F}(u_{h_{l_{\mathbf{x}}}, m-1}) + \mathcal{F}(u_{h_{l_{\mathbf{x}}-1}, m-1}).$$

As the quadrature and truncation levels are not completely independent, we only need to consider indices where the effective truncation level coincides with the truncation level, i.e., $m = \max\{k \in \mathbb{N} : l_{y_k} \geq 1\}$. In this case the differences can be written as

$$(5.8) \quad \begin{aligned} \Delta_{\tilde{l}}u &= \Delta_{(l_{y_1}, \dots, l_{y_m}, 0, 0, \dots)}^{(Q)} \Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u = \bigotimes_{k=1}^{\infty} \Delta_{l_{y_k}}^{(Q_k)} \left(\Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u \right) \\ &= \bigotimes_{k=1}^m \Delta_{l_{y_k}}^{(Q_k)} \left(\Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} \Delta_m^{(T)}u \right) \\ &= \bigotimes_{k=1}^{m-1} \Delta_{l_{y_k}}^{(Q_k)} \otimes \Delta_{l_m}^{(Q_m)} \left(\mathcal{F}(u_{h_{l_{\mathbf{x}}}, m}) - \mathcal{F}(u_{h_{l_{\mathbf{x}}-1}, m}) - \mathcal{F}(u_{h_{l_{\mathbf{x}}}, m-1}) + \mathcal{F}(u_{h_{l_{\mathbf{x}}-1}, m-1}) \right) \\ &= \bigotimes_{k=1}^{m-1} \Delta_{l_{y_k}}^{(Q_k)} \otimes \Delta_{l_{y_m}}^{(Q_m)} \left(\mathcal{F}(u_{h_{l_{\mathbf{x}}}, m}) - \mathcal{F}(u_{h_{l_{\mathbf{x}}-1}, m}) \right) \\ &= \bigotimes_{k=1}^m \Delta_{l_{y_k}}^{(Q_k)} \Delta_{l_{\mathbf{x}}}^{(\mathbf{x})} u_m, \end{aligned}$$

where we used the fact that

$$Q_{l_{y_m}}^{(m)} (\mathcal{F}(u_{h_{l_x}, m-1}) - \mathcal{F}(u_{h_{l_x-1}, m-1})) - Q_{l_{y_m}-1}^{(m)} (\mathcal{F}(u_{h_{l_x}, m-1}) - \mathcal{F}(u_{h_{l_x-1}, m-1})) = 0$$

as $l_{y_m} \geq 1$ and that constant functions are integrated exactly by the quadrature rule.

Hence, we obtained for (5.7) the natural simplified representation

$$(5.9) \quad \Delta_{\tilde{\mathcal{I}}} u = \bigotimes_{k=1}^m \Delta_{l_{y_k}}^{(Q_k)} \Delta_{l_x}^{(x)} u_m.$$

5.2. Dimension-adaptive algorithm for UQ. We are finally in the position to modify the procedure described in subsection 4.2 in order to match our approximation steps. To this end, we adapt the basic algorithm 4.1 to now construct a sequence of index sets $\tilde{\mathcal{I}}^{(1)} \subset \tilde{\mathcal{I}}^{(2)} \subset \dots \subset \tilde{\mathcal{I}}^{(t)} \subset \mathbb{N}^\infty$ which contain indices of the form (5.6). An index set $\tilde{\mathcal{I}}$ can then be used to obtain the sparse grid approximation

$$(5.10) \quad \mathbf{P}_{\tilde{\mathcal{I}}}^{\text{SG}} u = \sum_{\tilde{\mathcal{I}}} \Delta_{\tilde{\mathcal{I}}} u \quad \text{and} \quad \mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}} u = \sum_{\tilde{\mathcal{I}}} \alpha_{\tilde{\mathcal{I}}} \mathbf{P}_{\tilde{\mathcal{I}}} u$$

analogously to (4.1) and (4.2).

The construction of $\tilde{\mathcal{I}}$ is also based on two sets $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{O}}$, but we cannot store the infinite index arrays. However, since only a finite number of parametric dimensions are active at once, it suffices to store these active dimensions. To facilitate the implementation, we store all parameter dimensions which have been activated in course of the algorithm even if some of them are inactive at the current iteration. The number of parameter dimensions that have been activated is denoted by M , which will get larger in the course of the algorithm. In addition to the indices corresponding to these M activated quadratures dimensions, we store the spatial discretization level and a fixed number \widehat{M} of parameter dimensions that have not been activated yet to be able to increase the truncation level in future iterations of the algorithm. We therefore store indices of the form

$$(5.11) \quad \tilde{\mathcal{I}} = (l_x, l_{y_1}, \dots, l_{y_m}, 0, 0, \dots, 0) \in \mathbb{N}^{1+M+\widehat{M}}.$$

with $m \leq M + \widehat{M}$ and use a Boolean vector act to indicate whether a variable has been activated.

Our new adaptive method for UQ problems, which is presented in Algorithm 5.1, then proceeds similar to Algorithm 4.1. It starts with one index in $\tilde{\mathcal{A}}$ and adds successively indices with the highest profit indicator. Here, the set $\tilde{\mathcal{A}}$ includes indices which either refine the finite element discretization, refine the quadrature or raises the truncation level. The search for new admissible indices is separated into two parts: Lines 9-17 account for a refinement of the finite element discretization or the quadratures which have been activated before, whereby including $\tilde{\mathcal{I}} + \tilde{\mathcal{e}}_1$ in $\tilde{\mathcal{A}}$ increases the level index l_x and $\tilde{\mathcal{I}} + \tilde{\mathcal{e}}_k$ for $2 \leq k \leq 1 + M + \widehat{M}$ raises the level of the $k-1$ -th direction, i.e., the k -th quadrature.

The lines 18-27 become effective whenever $\tilde{\mathcal{I}}$ selected in line 6 activates a new variable which corresponds to a higher truncation level. In this case, the Boolean vector needs to be

adjusted and the number of stored parametric level is raised by one. This ensures that always \widehat{M} non-activated entries are included and the truncation level can be increased in subsequent iterations. Moreover, the index arrays which were added in previous steps of the algorithm need to be extended by one entry which is set to zero.

Note here that it would be sufficient to consider $\widehat{M} = 1$. By setting $\widehat{M} > 1$, the algorithm does not only consider raising the truncation level by one, but can consider the \widehat{M} forward neighbors as well. This avoids taking into account a too small number of stochastic variables as a parametric variable can be activated, even if not all parameters up to this parameter have been activated before. In our experiments it turned out that this approach with value $\widehat{M} = 5$ was more robust than just the choice $\widehat{M} = 1$. This lack of downward closedness in the truncation direction does not pose any problems for the combination technique due to the structure of the differences, compare also (5.8).

Algorithm 5.1 Adapted dimension-adaptive combination technique for UQ

Input: Tolerance $\varepsilon > 0$

Output: Index set $\tilde{\mathcal{I}}$ and approximation $\mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}}u$

- 1: $\tilde{\mathbf{l}} = (0, \dots, 0) \in \mathbb{R}^{1+\widehat{M}}$, $M = 0$, $\text{act} = (0, \dots, 0) \in \{0, 1\}^{\widehat{M}}$
- 2: $\tilde{\mathcal{A}} = \{\tilde{\mathbf{l}}\}$, $\tilde{\mathcal{O}} = \emptyset$
- 3: Compute local profit indicator $\eta_{\tilde{\mathbf{l}}}$
- 4: $\eta = \eta_{\tilde{\mathbf{l}}}$
- 5: **while** $\eta > \varepsilon$ **do**
- 6: Select index $\tilde{\mathbf{l}} \in \tilde{\mathcal{A}}$ with largest profit $\eta_{\tilde{\mathbf{l}}}$
- 7: $\tilde{\mathcal{A}} = \tilde{\mathcal{A}} \setminus \{\tilde{\mathbf{l}}\}$, $\tilde{\mathcal{O}} = \tilde{\mathcal{O}} \cup \{\tilde{\mathbf{l}}\}$
- 8: $\eta = \eta - \eta_{\tilde{\mathbf{l}}}$
- 9: **for** $k = 1, \dots, 1 + M + \widehat{M}$ **do**
- 10: $\tilde{\mathbf{j}} = \tilde{\mathbf{l}} + \tilde{\mathbf{e}}_k$
- 11: **if** $\tilde{\mathbf{j}} - \tilde{\mathbf{e}}_i \in \tilde{\mathcal{O}} \quad \forall 1 \leq i \leq 1 + M + \widehat{M}$ with $\tilde{j}_i > 0$ **then**
- 12: $\tilde{\mathcal{A}} = \tilde{\mathcal{A}} \cup \{\tilde{\mathbf{j}}\}$
- 13: Compute the approximation $\mathbf{P}_{\tilde{\mathbf{j}}}u$
- 14: Compute local profit indicator $\eta_{\tilde{\mathbf{j}}}$
- 15: $\eta = \eta + \eta_{\tilde{\mathbf{j}}}$
- 16: **end if**
- 17: **end for**
- 18: **if** it exists $n \in \{1, \dots, M + \widehat{M}\}$ s.t. $l_{y_n} > 0$ and $\text{act}_n = 0$ **then**
- 19: $\text{act}_n = 1$, extend act
- 20: $M = M + 1$
- 21: $\tilde{\mathbf{j}} = \tilde{\mathbf{e}}_{1+M+\widehat{M}}$
- 22: $\tilde{\mathcal{A}} = \tilde{\mathcal{A}} \cup \{\tilde{\mathbf{j}}\}$
- 23: Extend indices in $\tilde{\mathcal{A}}$, $\tilde{\mathcal{O}}$ by 0
- 24: Compute the approximation $\mathbf{P}_{\tilde{\mathbf{j}}}u$
- 25: Compute local profit indicator $\eta_{\tilde{\mathbf{j}}}$
- 26: $\eta = \eta + \eta_{\tilde{\mathbf{j}}}$
- 27: **end if**

28: **end while**

29: Combine approximations $\mathbf{P}_{\tilde{\mathbf{l}}}$ for $\tilde{\mathbf{l}} \in \tilde{\mathcal{A}} \cup \tilde{\mathcal{O}}$ according to (5.10) to obtain $\mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}}u$

30: **return** $\tilde{\mathcal{I}} = \tilde{\mathcal{A}} \cup \tilde{\mathcal{O}}$ and approximation $\mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}}u$

Different to the adaptive algorithms in [14, 15, 16, 30] which also construct sparse approximation schemes for the parametric approximation and increase the number of active variables, Algorithm 5.1 combines the quadrature method with the spatial discretization and the truncation level such that in each iteration either the finite element method, one Gauss-Hermite quadrature among the activated parametric dimensions or the truncation level is refined. This perspective is also reflected in the choice of our specific indicator which will be discussed in the next section.

5.3. Profit indicator. The crucial part of the adaptive algorithm is the refinement strategy which selects the index to be added to $\tilde{\mathcal{O}}$. This selection process is controlled by a local profit indicator $\eta_{\tilde{\mathbf{l}}}$ which should estimate the contribution of including the increment $\Delta_{\tilde{\mathbf{l}}}u$ in the sparse grid combination formula (5.10).

Since the global error can be bounded by

$$\|\mathbb{E}[\mathcal{F}(u)] - \mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}}u\|_{\mathcal{X}} \leq \sum_{\tilde{\mathbf{l}} \notin \tilde{\mathcal{I}}} \|\Delta_{\tilde{\mathbf{l}}}u\|_{\mathcal{X}},$$

we use

$$(5.12) \quad E_{\tilde{\mathbf{l}}} = \|\Delta_{\tilde{\mathbf{l}}}u\|_{\mathcal{X}}$$

as a local error estimator for adding $\tilde{\mathbf{l}}$ to $\tilde{\mathcal{I}}$. The increments are as discussed in Section 5.1 and the norm $\|\cdot\|_{\mathcal{X}}$ depends on the space in which the error is measured. Note at this point that, due to the upper bound by the triangle inequality, the local error indicator $E_{\tilde{\mathbf{l}}}$ is *reliable*. It is however not necessarily *efficient* in the terminology of error estimators for adaptive finite element methods (see, e.g. [1, 10]) as we have no lower bound $c \sum_{\tilde{\mathbf{l}} \notin \tilde{\mathcal{I}}} \|\Delta_{\tilde{\mathbf{l}}}u\|_{\mathcal{X}} \leq \|\mathbb{E}[\mathcal{F}(u)] - \mathbf{P}_{\tilde{\mathcal{I}}}^{\text{ct}}u\|_{\mathcal{X}}$, $0 < c \leq 1$.

Note here that the evaluation of $E_{\tilde{\mathbf{l}}}$ requires several anisotropic grid solutions. However, we only need to compute $\mathbf{P}_{\tilde{\mathbf{l}}}u$ and can reuse the other solutions that have been previously computed as the index set needs to be downward closed in the quadrature and finite element levels. Moreover, taking into account the computational cost, the local error estimator $E_{\tilde{\mathbf{l}}}$ might not be a well-suited refinement indicator, as including other indices might lead to a similar error reduction while having much lower computational cost. Therefore, we follow [9, 19] and use instead the benefit-cost ratio

$$\eta_{\tilde{\mathbf{l}}} = \frac{E_{\tilde{\mathbf{l}}}}{c_{\tilde{\mathbf{l}}}}$$

as refinement indicator where $c_{\tilde{\mathbf{l}}}$ denotes the cost of the evaluation of $\mathbf{P}_{\tilde{\mathbf{l}}}u$. With the notation for Algorithm 5.1, the cost is given by

$$c_{\tilde{\mathbf{l}}} = (N_{l_{\mathbf{x}}}^{(\mathbf{x})})^n \cdot \prod_{i=1}^m N_{l_{y_i}}^{(y_i)} \cdot m$$

where n denotes the spatial dimension and $m = \max\{k \in \mathbb{N} : l_{y_k} \geq 1\}$ denotes the actual truncation level. The first factor is the cost for the finite element method measured by the number of finite element nodes. Here we assume for reasons of simplicity an optimal solver like multigrid or a multilevel preconditioner, such that only a cost of $\mathcal{O}\left((N_{l_x}^{(\mathbf{x})})^n\right)$ is involved. The second factor is the number of quadrature points and the third factor accounts for the cost of evaluating the m -term truncation of the random field which needs to be included as well. This way, we balance the error with the required work and our algorithm selects the index from $\tilde{\mathcal{A}}$ with the best benefit-cost ratio and minimizes $\sum_{\tilde{\mathbf{l}} \in \tilde{\mathcal{I}}} \|\Delta_{\tilde{\mathbf{l}}} u\|_{\mathcal{X}}$ for a given cost. Here, the index with the highest contribution is only searched along the direct neighborhood of $\tilde{\mathcal{O}}$. Hence, we have to assume that the contributions of $\Delta_{\tilde{\mathbf{l}}} u$ decrease with the size of $\tilde{\mathbf{l}}$ to be able to ensure convergence. This is comparable to the saturation assumption for adaptive finite element methods. Having the coefficients in the Karhunen-Loève expansion sorted in decreasing order with respect to the $L^\infty(D)$ -norm implies already a decay for the truncation direction. Note that if this saturation assumption is not met, for example, when the diffusion field is rough and contains discontinuities or spikes on a very fine scale, also other methods such as multilevel Monte Carlo or Quasi Monte Carlo would run into similar problems as well.

At last, let us mention different stopping criteria. The criterion that is used in Algorithm 5.1 is based on a global profit indicator in the neighborhood set $\tilde{\mathcal{A}}$. This approach estimates how much profit is not covered yet. As the error contribution for indices in the complement of $\tilde{\mathcal{I}}$ is not known, the only available information are the contributions of the indices in $\tilde{\mathcal{A}}$. Therefore, an estimate of the global profit is

$$\eta = \sum_{\tilde{\mathbf{l}} \in \tilde{\mathcal{A}}} \eta_{\tilde{\mathbf{l}}}$$

and the algorithm stops when the global profit indicator is below a tolerance ε . It is also possible to replace in the stopping criterion the profit by the error contribution providing a global error indicator. Alternatively, the iteration could be stopped when a given upper bound on the total work $C = \sum_{\tilde{\mathbf{l}} \in \tilde{\mathcal{I}}} c_{\tilde{\mathbf{l}}}$ is exhausted. Another stopping criterion would look at the difference between the solutions of two iterations in order to estimate the error, i.e., $\|P_{\tilde{\mathcal{I}}^{(t+1)}}^{\text{ct}} u - P_{\tilde{\mathcal{I}}^{(t)}}^{\text{ct}} u\|_{\mathcal{X}} \leq \sum_{\tilde{\mathbf{l}} \in \tilde{\mathcal{I}}^{(t+1)} \setminus \tilde{\mathcal{I}}^{(t)}} \|\Delta_{\tilde{\mathbf{l}}} u\|_{\mathcal{X}}$. The algorithm stops if this difference is ζ -times in a row smaller than ε . Here, the parameter ζ is introduced to avoid stopping too early, when the solution differs only slightly from one iteration to another while the next solution would result in a larger difference.

6. Numerical examples. As indicated in the introduction, we focus on the model problem

$$(6.1) \quad -\operatorname{div}(a(\mathbf{x}, \omega) \nabla u(\mathbf{x}, \omega)) = f(\mathbf{x}) \quad \text{in } D,$$

where $a(\mathbf{x}, \omega) = \exp(b(\mathbf{x}, \omega))$, compare (2.1). Moreover, we set $n = 2$, $D = [0, 1]^2$, and apply homogeneous Dirichlet boundary conditions. For our numerical examples, we assume that the covariance function of the random field $b(\mathbf{x}, \omega)$ is in the class of Matérn covariance functions, see, e.g., [32]. Then, the regularity and hence the eigenvalue decay for the Karhunen-Loève expansion is explicitly known.

The Matérn covariance kernels are defined by a stationary covariance function. Let $r = |\mathbf{x} - \mathbf{x}'|$ denote the distance between two points and let $\xi, \sigma^2 > 0$ be the correlation length

and variance, respectively. Then the Matérn covariance function of order $\nu > 0$ is defined as

$$(6.2) \quad \text{cov}(r; \nu, \xi) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{r}{\xi} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{r}{\xi} \right),$$

where Γ is the Gamma function and K_ν is the modified Bessel function of the second kind of order ν . The formula (6.2) simplifies for $\nu = s + \frac{1}{2}$ with $s \in \mathbb{N}$ to

$$\text{cov} \left(r; s + \frac{1}{2}, \xi \right) = \sigma^2 \exp \left(-\frac{\sqrt{2s+1}r}{\xi} \right) \frac{s!}{(2s)!} \sum_{i=0}^s \frac{(s+i)!}{i!(s-i)!} \left(\frac{2\sqrt{2s+1}r}{\xi} \right)^{s-i}.$$

Furthermore, in the limit $\nu \rightarrow \infty$, we obtain a Gaussian covariance function

$$\text{cov}(r; \xi) = \sigma^2 \exp \left(-\frac{1}{2} \frac{r^2}{\xi^2} \right).$$

The parameter ν can be regarded as a smoothness parameter as it controls the regularity of the covariance function at the point $r = 0$ as well as the regularity in the spatial variable of realizations of the random field. Both parameters ξ and ν influence the decay of the Karhunen-Loève coefficients and therefore the stochastic regularity. In the case $\nu < \infty$ the asymptotic decay rates of the KL coefficients are known to be algebraic, while for $\nu = \infty$ the eigenvalues decrease exponentially. Moreover, their pre-asymptotic behavior is determined by the correlation length ξ . The decay of the eigenvalues for various parameters are shown in Figure 6.1 together with a realization of the random field $\log(a(x, \omega))$ for $\nu = 2.5$ and $\xi = 0.2$.

In the following we consider different modeling parameters ν and ξ leading to different decay rates in the Karhunen-Loève expansion. Even though the decay rates differ, we can apply the adaptive combination technique without any adjustments, which would be necessary for any a-priori construction of a sparse grid approximation.

We illustrate the performance of our new dimension-adaptive combination technique by computing the first moment $\mathbb{E}[u]$ for different permeability fields. For other quantities of interest such as the second moment and linear functionals of u , we observed similar results in further experiments. Here, we only present the results for the first moment in detail. We assume homogeneous Dirichlet boundary condition and set $f = 1$. The algorithm is applied with $\widehat{M} = 5$ variables and the error is measured in the $L^2(D)$ -norm with respect to a reference solution. Consequently, we use as a measure for the local contribution $\|\Delta_{\tilde{I}} u\|_{L^2(D)}$ throughout the algorithm. For solving the PDE for fixed parameter values on a uniform triangulation of D with mesh width $h_{l_x} = \mathcal{O}(2^{-l_x})$ and piecewise finite elements, the finite element software FEniCS [2] is used.

The behavior of the dimension-adaptive algorithm for the Matérn covariance function with parameter $\nu = 2.5$ and correlation length $\xi = 0.4$ is illustrated in Figure 6.2. We observe that the error decreases with increasing number of iterations. We observe also that the error decreases fast in the first iterations as the profits are high for these indices and that the global error estimators follow this behavior of the error.

Figure 6.2 (right) expresses the interplay of the three different numerical approximations and the evolution of the index set during the adaptive algorithm. It shows how the maximal level for each refinement direction increases throughout the process, whereby, for the




Figure 6.1. Left: Decay of Karhunen-Loève coefficients λ_k for different parameters ν and ξ . Right: Realization of random field $b(\mathbf{x}, \omega)$ for $\xi = 0.2$ and $\nu = 2.5$.




Figure 6.2. Left: convergence in terms of number of iterations. Right: maximal level in the index set for spatial refinement, quadrature and truncation level.

quadrature, the maximal level of all dimensions is considered. As expected, a much finer discretization of the spatial domain than for the stochastic parameters is needed, since the Gaussian quadrature exhibits exponential convergence and therefore not many levels are required for a small quadrature error. Moreover, the number of activated parameter dimensions increases successively such that, after 3000 iterations, our algorithm has activated 77 parameter dimensions which cover the variability of the random field to a large extent. Moreover, it balances the corresponding finite noise truncation error with the quadrature error and the spatial discretization error while also taking the associated costs into account.

The interplay of the truncation, the quadrature and the finite element method is also illustrated in Figure 6.3, where the index set $\tilde{\mathcal{O}}$ at iteration 2000 is depicted. Again, we give just the maximal level over all quadrature rules for the stochastic approximation direction. We observe the typical shape of a generalized sparse grid discretization in three dimensions:




Figure 6.3. Index set $\tilde{\mathcal{O}}$ at iteration 2000.




Figure 6.4. Convergence of dimensions-adaptive algorithm with respect to the computational cost.

A high resolution in one direction relates to a corresponding low resolution in the other two directions.

The convergence rate of the solution with respect to the computational cost is shown in Figure 6.4 for the example of random fields with the parameter values $\nu = 2.5$ and $\nu = \infty$ and the correlation lengths $\xi = 0.2$ and $\xi = 0.4$ which possesses an algebraic and an exponential decay of the coefficients, respectively. In all four cases, the decay of the coefficients is sufficiently fast such that the mixed regularity condition as in [26] for a sparse grid approximation between the spatial and the parametric variables is met. We can thus expect that the convergence rate is limited by the slowest approximation direction. As the realizations of the random fields are continuously differentiable, the PDE solution satisfies $u \in L^q(\Omega, H^2(D))$ for $q < \infty$ (see [13, 38]). Hence, for fixed \mathbf{y} , the finite element method achieves at most a convergence rate of order one with respect to the number of degrees of freedom (cf. Section 3.2), as we have $\mathcal{O}((N_{l_x}^{(\mathbf{x})})^2)$ degrees of freedom for a mesh with width $h_{l_x} = (N_{l_x}^{(\mathbf{x})})^{-1}$. The truncation error in terms of m decreases with a rate faster than one and the smooth dependence on the stochastic parameters yields also a faster convergence of the quadrature. Hence the convergence rate for smooth fields is limited by that of the finite element method. Indeed, we observe an asymptotic convergence rate that approaches $\mathcal{O}(N^{-1})$ with respect to the overall computational cost $N = \sum_{\tilde{i} \in \tilde{\mathcal{I}}} c_{\tilde{i}}$, which is consistent with the behavior that is expected from a-priori analysis. Compared to the often applied methods based on Monte Carlo quadrature with convergence rates limited to $1/2$, our adaptive approach with the Gaussian quadrature yields a better rate as it can exploit the regularity of the parametric variables.

To illustrate the adjustment of our algorithm to the specific problem under consideration, we plot the levels used for the parametric discretization in Figure 6.5. Here, the largest level for each one-dimensional quadrature level is shown. We observe that a higher refinement is applied for the first few parameter dimensions as their influence on the solution is larger, while for the later parameter dimensions only the lowest level quadrature is required. Our algorithm




Figure 6.5. Maximal quadrature level (black after iteration 100, dark gray after iteration 500 and light gray after iteration 2500) for different random fields. From top to bottom: $\nu = \infty$ and $\xi = 0.4$; $\nu = 2.5$ and $\xi = 0.4$; $\nu = \infty$ and $\xi = 0.2$; $\nu = 2.5$ and $\xi = 0.2$.

therefore detects the anisotropy in the parameters and adjusts the necessary quadrature levels accordingly. We also observe that fewer variables are activated for the smoother fields with $\nu = \infty$ than for the fields with $\nu = 2.5$. Due to the faster decay of the KL coefficients, it suffices to only consider a smaller number of parametric variables. The correlation length merely affects the size of the pre-asymptotic regime of the decay. For smaller correlation lengths the regime is larger, which is reflected in the quadrature levels by a larger number of parametric dimensions that are refined to a higher extent. Furthermore, slightly more variables are activated in the computation. However, the asymptotic convergence rate is not affected.

As a final example, we consider a problem with $\nu = 0.5$ and thus with an exponential covariance function. For $r = |\mathbf{x} - \mathbf{x}'|_1$ the Karhunen-Loève eigenvalues and eigenfunctions are explicitly known. In this case the almost sure convergence of $a(\mathbf{x}, \omega)$ in $L^\infty(D)$ is guaranteed [11], but the condition $\{\gamma_k\}_k \in \ell^1(\mathbb{N})$ (c.f. (2.6)) is not satisfied as the realizations of the diffusion coefficient are only Hölder continuous with exponent $1/2$. For this example, the assumptions typically made to obtain a-priori estimates on the size of the surpluses $\Delta_{\tilde{t}}u$ cannot be verified. Nevertheless, the set that includes the indices with the highest local benefit-cost ratio can be constructed with our adaptive algorithm in a straightforward way.

The rougher diffusion coefficient results in less regular solutions of the PDE which only satisfies $u(\mathbf{y}) \in H^{3/2}(D)$, see [13, 38]. Hence, we can only expect from the finite element method to achieve a convergence order of $1/2$. Indeed, we observe this slower decay of the $L^2(D)$ -error with our dimension-adaptive combination technique.

The associated slow decrease of the coefficients in the KL expansion results in a large




Figure 6.6. Convergence plot for $\nu = 0.5$

number of parametric variables that need to be taken into account. Our algorithm adjusts to this slower decay and activates 192 parametric variables for $\xi = 0.4$ at iteration 2500 (and 189 parametric variables for $\xi = 0.2$ at iteration 2500) which cover most of the variability of the random field. As the parametric dimension grows and more parameters have a strong influence on the solution, the sparse grid for the parametric discretization gets naturally more costly than for problems with only a few important parametric dimensions.

7. Concluding remarks. In this paper we presented a dimension-adaptive algorithm that computes quantities of interests for a given stochastic random diffusion coefficient and relies on a sparse grid approximation. The algorithm balances the spatial, stochastic and parametric approximation in a cost-efficient way and detects the necessary discretization levels to obtain a prescribed overall error. This method is especially useful when information about the decrease of the increments $\Delta_{\mathcal{I}}u$ is not available, as it does not require any a-priori knowledge on the regularity of the solution and on the decay of the KL coefficients.

Our adaptive algorithm is based on the benefit-cost ratio as reliable refinement indicator such that it minimizes the error for a prescribed computational cost. Further work may focus on the optimality of this refinement strategy and on the choice of efficient and reliable global error estimators.

Numerical examples showed that our obtained convergence order is comparable with the known a-priori convergence results for smooth random diffusion coefficients. Compared to other methods such as multilevel Monte Carlo and Quasi Monte Carlo, our algorithm can exploit the parametric regularity such that the convergence order in terms of the computational cost is limited by the finite element method for smooth fields. For problems with less smoothness, the algorithm also yields good results, but converges slower due to the lower regularity of the PDE solution.

As the algorithm is based on combinations of quadratures and finite element discretizations, the algorithm can be generalized using other spatial discretization techniques, such as

h-adaptive and h-p-adaptive finite element methods. By this, we expect to be able to benefit from the regularity in the case of smooth PDE solutions and to obtain a higher order of convergence than by a Quasi Monte Carlo method. Furthermore, our algorithm is not limited to computing first and second moment or linear functionals. In fact, it can easily be generalized to other quantities of interest involving a different function \mathcal{F} . This might however require an additional approximation step for the evaluation of \mathcal{F} which also needs to be balanced.

In this work, we concentrated on adaptivity with respect to the stochastic parameters and kept the spatial discretization uniform. Besides, there is also the possibility/necessity to apply local adaptive refinement in space, compare e.g. [7, 8, 17]. The study of adaptivity in the stochastic *and* the spatial domain is future work.

Acknowledgments. Michael Griebel was supported by the *Hausdorff Center for Mathematics* in Bonn, funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC-2047/1 - 390685813 and the CRC 1060 *The Mathematics of Emergent Effects* of the Deutsche Forschungsgemeinschaft.

REFERENCES

- [1] M. AINSWORTH AND J. ODEN, *A Posteriori Error Estimation in Finite Element Analysis*, Pure and Applied Mathematics, Wiley, New York, 2000.
- [2] M. ALNAES, J. BLECHTA, J. HAKE, A. JOHANSSON, B. KEHLET, A. LOGG, C. RICHARDSON, J. RING, M. ROGNES, AND G. WELLS, *The FEniCS Project Version 1.5*, Archive of Numerical Software, Vol 3 (2015).
- [3] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM Rev., 52 (2010), pp. 317–355.
- [4] M. BACHMAYR, A. COHEN, R. DEVORE, AND G. MIGLIORATI, *Sparse polynomial approximation of parametric elliptic PDEs. Part II: Lognormal coefficients*, ESAIM Math. Model. Numer. Anal., 51 (2017), pp. 341–363.
- [5] A. BARTH, C. SCHWAB, AND N. ZOLLINGER, *Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients*, Numer. Math., 119 (2011), pp. 123–161.
- [6] J. BEAR AND A. CHENG, *Modeling Groundwater Flow and Contaminant Transport*, Springer Netherlands, 2010.
- [7] A. BESPALOV, D. SILVESTER, AND F. XU, *Error estimation and adaptivity for stochastic collocation finite elements Part I: single-level approximation*, (2021), <https://arxiv.org/abs/2109.07320>.
- [8] A. BESPALOV AND D. J. SILVESTER, *Error estimation and adaptivity for stochastic collocation finite elements Part II: multilevel approximation*, (2022), <https://arxiv.org/abs/2202.08902>.
- [9] H.-J. BUNGARTZ AND M. GRIEBEL, *Sparse grids*, Acta Numer., 13 (2004), pp. 147–269.
- [10] C. CARSTENSEN, M. FEISCHL, M. PAGE, AND D. PRAETORIUS, *Axioms of adaptivity*, Comput. Math. Appl., 67 (2014), pp. 1195–1253.
- [11] J. CHARRIER, *Strong and weak error estimates for elliptic partial differential equations with random coefficients*, SIAM J. Numer. Anal., 50 (2012), pp. 216–246.
- [12] J. CHARRIER AND A. DEBUSSCHE, *Weak truncation error estimates for elliptic PDEs with lognormal coefficients*, Stoch. Partial Differ. Equ. Anal. Comput., 1 (2013), pp. 63–93.
- [13] J. CHARRIER, R. SCHEICHL, AND A. 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.
- [14] P. CHEN, *Sparse quadrature for high-dimensional integration with Gaussian measure*, ESAIM Math. Model. Numer. Anal., 52 (2018), pp. 631–657.
- [15] A. CHKIFA, A. COHEN, AND C. SCHWAB, *High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs*, Found. Comput. Math., 14 (2014), pp. 601–633.

- [16] O. ERNST, B. SPRUNGK, AND L. TAMELLINI, *Convergence of sparse collocation for functions of countably many Gaussian random variables (with application to elliptic PDEs)*, SIAM J. Numer. Anal., 56 (2018), pp. 877–905.
- [17] M. FEISCHL AND A. SCAGLIONI, *Convergence of adaptive stochastic collocation with finite elements*, Comput. Math. Appl., 98 (2021), pp. 139–156.
- [18] J. GARCKE, *Sparse grids in a nutshell*, in Sparse Grids and Applications, J. Garcke and M. Griebel, eds., vol. 88 of Lect. Notes Comput. Sci. Eng., Springer Berlin Heidelberg, 2012, pp. 57–80.
- [19] T. GERSTNER AND M. GRIEBEL, *Dimension adaptive tensor product quadrature*, Computing, 71 (2003), pp. 65–87.
- [20] I. GRAHAM, F. KUO, J. NICHOLS, R. SCHEICHL, C. SCHWAB, AND I. SLOAN, *Quasi-Monte Carlo finite element methods for elliptic PDEs with lognormal random coefficients*, Numer. Math., 131 (2015), pp. 329–368.
- [21] M. GRIEBEL, H. HARBRECHT, AND M. MULTERER, *Multilevel quadrature for elliptic parametric partial differential equations in case of polygonal approximations of curved domains*, SIAM J. Numer. Anal., 58 (2020), pp. 684–705.
- [22] M. GRIEBEL, G. LI, AND C. RIEGER, *On the numerical approximation of the Karhunen-Loève expansion for random fields with random discrete data*, (2021), <https://arxiv.org/abs/2112.02526>.
- [23] M. GRIEBEL, M. SCHNEIDER, AND C. ZENGER, *A combination technique for the solution of sparse grid problems*, in Iterative Methods in Linear Algebra, P. de Groen and R. Beauwens, eds., IMACS, Elsevier, 1992, pp. 263–281.
- [24] A.-L. HAJI-ALI, H. HARBRECHT, M. PETERS, AND M. SIEBENMORGEN, *Novel results for the anisotropic sparse grid quadrature*, J. Complexity, 47 (2018), pp. 62–85.
- [25] A.-L. HAJI-ALI, F. NOBILE, L. TAMELLINI, AND R. TEMPONE, *Multi-index stochastic collocation for random PDEs*, Comput. Methods Appl. Mech. Engrg. C, 306 (2016), pp. 95–122.
- [26] H. HARBRECHT, M. PETERS, AND M. SIEBENMORGEN, *Multilevel accelerated quadrature for PDEs with log-normally distributed diffusion coefficient*, SIAM/ASA J. Uncertain. Quantif., 4 (2016), pp. 520–551.
- [27] V. HOANG AND C. SCHWAB, *N-term Wiener chaos approximation rates for elliptic PDEs with lognormal Gaussian random inputs*, Math. Models Methods Appl. Sci., 24 (2014), pp. 797–826.
- [28] F. KUO, R. SCHEICHL, C. SCHWAB, I. SLOAN, AND E. ULLMANN, *Multilevel quasi-Monte Carlo methods for lognormal diffusion problems*, Math. Comp., 86 (2017), pp. 2827–2860.
- [29] F. KUO, C. SCHWAB, AND I. SLOAN, *Quasi-Monte Carlo finite element methods for a class of elliptic partial differential equations with random coefficients*, SIAM J. Numer. Anal., 50 (2012), pp. 3351–3374.
- [30] F. NOBILE, L. TAMELLINI, F. TESEI, AND R. TEMPONE, *An adaptive sparse grid algorithm for elliptic PDEs with lognormal diffusion coefficient*, in Sparse Grids and Applications - Stuttgart 2014, J. Garcke and D. Pflüger, eds., vol. 109, Springer International Publishing, 2016, pp. 191–220.
- [31] F. NOBILE, R. TEMPONE, AND C. WEBSTER, *An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM J. Numer. Anal., 46 (2008), pp. 2411–2442.
- [32] C. RASMUSSEN AND C. WILLIAMS, *Gaussian Processes for Machine Learning*, Adaptive Computation and Machine Learning, MIT Press, 2006.
- [33] P. ROBBE, D. NUYENS, AND S. VANDEWALLE, *A Multi-index quasi-Monte Carlo algorithm for lognormal diffusion problems*, SIAM J. Sci. Comput., 39 (2017), pp. S851–S872.
- [34] C. SCHWAB AND C. GITTELSON, *Sparse tensor discretizations of high-dimensional parametric and stochastic PDEs*, Acta Numer., 20 (2011), pp. 291–467.
- [35] C. SCHWAB AND R. TODOR, *Karhunen-Loève approximation of random fields by generalized fast multipole methods*, J. Comput. Phys., 217 (2006), pp. 100–122.
- [36] T. SULLIVAN, *Introduction to Uncertainty Quantification*, Springer International Publishing, 2015.
- [37] A. TECKENTRUP, P. JANTSCH, C. WEBSTER, AND M. GUNZBURGER, *A multilevel stochastic collocation method for partial differential equations with random input data*, SSIAM/ASA J. Uncertain. Quantif., 3 (2015), pp. 1046–1074.
- [38] A. TECKENTRUP, R. SCHEICHL, M. GILES, AND E. ULLMANN, *Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients*, Numer. Math., 125 (2013), pp. 569–600.