



Institut für Numerische Simulation

Rheinische Friedrich-Wilhelms-Universität Bonn

Wegelerstraße 6 • 53115 Bonn • Germany
phone +49 228 73-3427 • fax +49 228 73-7527
www.ins.uni-bonn.de

F. Franzelin, P. Diehl, D. Pflüger

Non-intrusive Uncertainty Quantification with Sparse Grids for Multivariate Peridynamic Simulations

INS Preprint No. 1408

June 2014

Non-intrusive Uncertainty Quantification with Sparse Grids for Multivariate Peridynamic Simulations

Fabian Franzelin¹, Patrick Diehl², and Dirk Pflüger¹

Abstract

Peridynamics is an accepted method in engineering for modeling crack propagation on a macroscopic scale. However, the sensitivity of the method to two important model parameters – elasticity and the particle density – has not yet been studied. Motivated by [11, 25] we use Peridynamics to simulate a high-speed projectile impacting a plate and study the overall damage on the plate. We have extended the setting by the magnitude of the force of the indenter and selected the parameter range such that a sharp transition in the response function occurs.

We describe the simulation setting as an uncertainty quantification problem and use a non-intrusive stochastic collocation method based on spatially adaptive sparse grids to propagate the uncertainty. We show first convincing results of its successful application to Peridynamics and compare to Monte Carlo sampling. If the magnitude of the force is deterministic, a strong sensitivity of the damage in the plate with respect to the elasticity factor can be shown for the 2-dimensional setting. If it is non-deterministic, it dominates the simulation and explains most of the variance of the solution. The error of the expectation value estimation reaches an early saturation point for the studied collocation methods: We found parameter ranges where the quantity of interest oscillates. Moreover, faster convergence and higher robustness than for the Monte Carlo method can be observed.

Key words: uncertainty quantification, sparse grids, Peridynamics, sensitivity analysis, ANOVA

Universität Stuttgart, Universitätsstr. 38, 70569 Stuttgart
fabian.franzelin@ipvs.uni-stuttgart.de, dirk.pflueger@ipvs.uni-stuttgart.de
Universität Bonn, Wegelerstr. 6, 53115 Bonn diehl@ins.uni-bonn.de,
schweitzer@ins.uni-bonn.de

1 Introduction

Most simulations depend on certain parameters, such as material parameters or model parameters. Frequently, the exact values of parameters are not given, and the effect of uncertainty in the choice of parameters on certain output values are of interest. In plenty of applications, the influence of parameters to the simulation results are of high relevance. In this paper, we present a non-intrusive method for a d -dimensional parameter study for simulations, both to estimate the influence of parameters and to quantify uncertainties. To this end, we consider the peridynamic theory [24, 25], a non-local particle-based theory in continuum mechanics with a focus on discontinuous functions as they arise in fracture mechanics.

Silling and Askari used in their publication about Peridynamics [25] the impact of a spherical indenter on a cylindrical plate to show the method's capability to model complex fractures. Kidane et al. [11] employed a similar setting. They simulated the impact of steel projectiles with ballistics speed on aluminum plates, with regard to a data-on-demand uncertainty quantification (UQ) protocol.

Motivated by their work, this paper brings together these two aspects, Peridynamics and UQ, using stochastic collocation based on spatially adaptive sparse grids, a numerical scheme for higher-dimensional discretizations [4]. Our simulation setting is the impact of a high-speed projectile of steel on a ceramic plate. The aim of this paper is to quantify the sensitivity of the damage in the plate with regard to three relevant model parameters. First, we consider a scaling factor for the elasticity in the Prototype Brittle Material (PMB) model, which is used to model the ceramic material of the plate. The second one is the density of particles in the plate, which we model by the initial distance between two particles, and the third one is a scaling factor for the magnitude of the force of the indenter. We formulate the parameter study as a UQ problem and interpret these parameters as uncertain, i.e., as random variables.

Due to the uncertainty of the inputs, the model's outcome becomes uncertain as well, but neither statistical moments nor a probability distribution are known a priori. They depend on the uncertainty of the parameters, which have to be propagated through the simulation. We consider non-intrusive UQ, sample the stochastic space on so-called collocation nodes, compute the outcome of interest and analyze the effects. The most common method based on this principle is Monte Carlo sampling. It is independent of the dimensionality of the stochastic space, easy to implement, and it is guaranteed to converge to the true solution in the limit of infinitely many samples. Its main drawback is a rather slow convergence rate.

In a simulation context each collocation node requires a full-scale simulation and is often very costly. This motivated us to use the knowledge about the solution, namely the simulation results at the collocation nodes, and interpolate between them. A very common approach to construct such

an interpolant is the non-intrusive generalized polynomial chaos expansion (gPC), presented in [29, 30]. The expansion is based on the Askey Scheme, which assigns parameter distributions to optimized orthogonal bases. The corresponding coefficients are computed by spectral projection, which leads of course to a possibly high-dimensional quadrature problem. However, the basis functions have global support and suffer therefore Gibb’s phenomenon when sharp transitions occur [29]. Therefore, [28] introduced the multi-element gPC, which decomposes the stochastic space adaptively using an Analysis-of-Variance (ANOVA) based decomposition criterion. Oladyshkin extended the polynomial chaos expansion to arbitrary parameter distributions, which is known as the arbitrary polynomial chaos method [14].

However, we propose a different approach, which is called adaptive sparse grid collocation (ASGC) method in [7]. It is based on standard spatially adaptive sparse grids [4, 18, 19] and expands the stochastic space using a hierarchical basis. It can cope with the curse of dimensionality to a large extent per construction, which gPC methods can not. Moreover, the sparse grid method provides, due to its hierarchical approach, a direct estimate for the interpolation error for free, which allows one to adaptively refine to peculiarities of the underlying function. This grid-based approach scales linearly in the number of collocation nodes and then allows the efficient estimation of the statistical moments. The basis functions have local support and can thus resolve sharp transitions without suffering a lot from Gibb’s phenomenon. Moreover, we consider several strategies for adaptive refinement and are able to further reduce the number of samples. Sparse grid collocation is a highly active field of research in the context of UQ: Jakeman et al. worked on the detection of discontinuities in high-dimensional stochastic spaces [10], Archibald et al. extended it to hybrid parallel architectures [2], and Zhang et al. used it to improve Bayesian inference in UQ settings [31], to name but a few.

In this work, we employ the ASGC method to simulations with Peridynamics as a first, still low-dimensional proof of concept for its efficiency and robustness in a real-world example. To keep the number of simulation runs small we propose a new refinement criterion. Having the explicit function representation of the solution in hand, we estimate global sensitivity values by variance decomposition based on [26].

In the next section we introduce the UQ problem formally, and describe how we compute global sensitivity values. In Sect. 3 we present the adaptive sparse grid collocation method in detail. In Sect. 4, we introduce the peridynamic model. The simulation setting is described in Sect. 5. We present numerical results in Sect. 6 for a two- and three-dimensional setting that can explain properties of the underlying simulations, and conclude this work with some remarks in Sect. 7.

2 Problem Setup

The task we address in this paper is to estimate sensitivity values of peridynamic parameters with respect to some quantity of interest we compute out of the results of peridynamic simulations. We describe the variation of the input by probability density functions, which we propagate through the simulation. In this section we discuss first the forward propagation of uncertainty and second the sensitivity analysis based on ANOVA.

2.1 Forward Propagation of Uncertainty

Let the system be driven by some model $\mathcal{M} = (\boldsymbol{\theta}, p(\boldsymbol{\theta}), F(\boldsymbol{\theta}), u)$ which depends on a finite number of random parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_d) \in \Gamma$ with a given joint probability density function $p(\boldsymbol{\theta})$ and cumulative distribution function $F(\boldsymbol{\theta})$, and a response function $u: \Gamma \rightarrow \mathbb{R}$ for some quantity of interest. So we seek a functional representation of $u(\boldsymbol{\theta})$ under the model \mathcal{M} in order to extract its probabilistic characterization introduced by the probabilistic inputs $\boldsymbol{\theta}$. Knowledge of the probability laws of $\boldsymbol{\theta}$ leads to the probability laws of u .

However, the model can be arbitrarily complex and an analytical solution might not even exist. To overcome this problem we replace the real solution by some suitable approximation f such that

$$f(\boldsymbol{\theta}) \approx u(\boldsymbol{\theta}). \quad (1)$$

Once we have an approximation f we can give estimates of the probability density function (risk analysis), statistical moments (expectation value, variance), confidence intervals of u , and even do a sensitivity analysis.

However, computing global sensitivity values based on ANOVA includes computing multi-dimensional variances of f . Motivated by Lemma 1, we can neglect the probability density function $p(\boldsymbol{\theta})$ for the computation of these variances due to a suitable transformation of the stochastic space to the unit hypercube, as long as the random variables are independent.

Lemma 1. *Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_d) \in \Gamma$ be a multivariate random variable from which its components θ_i are arbitrarily distributed but independent. They have a continuous and differentiable cumulative function $F(\boldsymbol{\theta})$ and a continuous probability density function $p(\boldsymbol{\theta})$. Then it holds that $\boldsymbol{\xi} := F(\boldsymbol{\theta}) \in [0, 1]^d$ is uniformly distributed under F and the expectation value for some function $u: \Gamma \rightarrow \mathbb{R}$ is*

$$\mathbb{E}(u) = \int_{\Omega} u(\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{[0,1]^d} u(F^{-1}(\boldsymbol{\xi})) d\boldsymbol{\xi}. \quad (2)$$

Proof. Due to the independence of θ_i and θ_j for $i \neq j$ we can define ξ_i component-wise as

$$\xi_i = F(\theta_i) \quad \text{for } 1 \leq i \leq d. \quad (3)$$

Standard statistics secures that each ξ_i is uniformly distributed in $[0, 1]$. In order to proof Eq. (2) we just substitute $\boldsymbol{\theta}$ by $F^{-1}(\boldsymbol{\xi})$. According to the substitution theorem for integrals we need to transform the ranges of the integral and insert a volume correction factor to the integral, i.e.,

$$\frac{d\boldsymbol{\theta}}{d\boldsymbol{\xi}} = |\det J_{F^{-1}}(\boldsymbol{\xi})|, \quad (4)$$

where $J_{F^{-1}}(\boldsymbol{\xi})$ is the Jacobian matrix of F^{-1} . Moreover, we know that $F^{-1}(\boldsymbol{\xi})$ is bijective, Lebesgue-measurable and differentiable, and therefore the transformation theorem for probability densities [27] holds, and thus

$$p(\boldsymbol{\xi}) = |\det J_{F^{-1}}(\boldsymbol{\xi})|p(F^{-1}(\boldsymbol{\xi})) = |\det J_{F^{-1}}(\boldsymbol{\xi})|p(\boldsymbol{\theta}) = 1. \quad (5)$$

Inserting Eq. (4) and (5) in the definition of the expectation value leads to the general calculation rule for the expectation value under F

$$\begin{aligned} \mathbb{E}(u) &= \int_{\Gamma} u(\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta} \\ &= \int_{[0,1]^d} u(F^{-1}(\boldsymbol{\xi})) \underbrace{\frac{|\det J_{F^{-1}}(\boldsymbol{\xi})|}{|\det J_{F^{-1}}(\boldsymbol{\xi})|}}_{=1} p(\boldsymbol{\xi}) d\boldsymbol{\xi} \\ &= \int_{[0,1]^d} u(F^{-1}(\boldsymbol{\xi}))d\boldsymbol{\xi}. \end{aligned} \quad (6)$$

□

For the more general case where no independence assumption of the marginal distributions can be made, we want to refer the reader to other, more elaborate statistical transformations, for example the Rosenblatt-transformation [21] or the Nataf-transformation.

However, we approximate the variance of u by the variance of f using Steiner's translation theorem $\mathbb{V}(f) = \mathbb{E}(f^2) - \mathbb{E}(f)^2$ and estimate each term with Eq. (2).

With respect to Peridynamics, we consider a model $\mathcal{M}(\boldsymbol{\theta}, p(\boldsymbol{\theta}), F(\boldsymbol{\theta}), u)$ with the uncertain parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$. With θ_1 as a material parameter, θ_2 as a discretization parameter and θ_3 as the external force. The response function u maps the uncertain parameters via simulation to the damage of the material, our quantity of interest. The model is formally introduced in Sect. 4 and 5.

2.2 Global Sensitivity Analysis

For the global sensitivity analysis we use the unanchored ANOVA decomposition [26] of the surrogate f . In general, this would require $2^d - 1$ integrations. For large d this is often unfeasible. To overcome this, one can employ the anchored ANOVA approach where the integrals are replaced by function evaluations at some anchor point [8]. As we deal with rather low dimensionalities so far, the unanchored approach is sufficiently efficient for the purposes of this paper. To be specific, we compute the ANOVA representation of f as

$$f(\theta_1, \dots, \theta_d) = f_0 + \sum_{i=1}^d f_i(\theta_i) + \sum_{i=1}^d \sum_{i < j}^d f_{i,j}(\theta_i, \theta_j) + \dots + f_{1,2,\dots,d}(\theta_1, \dots, \theta_d). \quad (7)$$

This decomposition is unique if each member $f_{i_1, \dots, i_k} \neq f_0$ has zero mean, i.e.,

$$\int_{\Gamma_{i_1} \times \dots \times \Gamma_{i_k}} f_{i_1, \dots, i_k}(\theta_{i_1, \dots, i_k}) p(\theta_{i_1, \dots, i_k}) d\theta_{i_1, \dots, i_k} = 0 \quad (8)$$

for some uncorrelated and independent probability measure $p(\theta_{i_1, \dots, i_k}) = \prod_{j=1}^k p(\theta_{i_j})$ [8, 26]. In total we have 2^d members, which we identify in the following by $\mathcal{P}(\mathcal{D})$ being the power set over $\mathcal{D} := \{1, \dots, d\}$. On this set of sets we introduce a total order $\mathcal{J} \prec \mathcal{K} \Leftrightarrow |\mathcal{J}| < |\mathcal{K}|$, $\mathcal{J}, \mathcal{K} \in \mathcal{P}(\mathcal{D})$ and can now write Eq. (7) in a compact form (notation adapted from [9]) as

$$f(\theta_1, \dots, \theta_d) = f_\emptyset + \sum_{k=1}^d \sum_{\mathcal{J} \in \mathcal{P}(\mathcal{D})} f_{\mathcal{J}}(\theta_{\mathcal{J}}). \quad (9)$$

where $f_\emptyset = f_0$. To guarantee orthogonality, we obtain the constant term f_\emptyset from

$$f_\emptyset = \mathbb{E}(f) = \int_{\Gamma} f(\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (10)$$

and the higher-order terms from

$$\mathbb{E}(f_{\mathcal{J}}) = \int_{\Gamma_{\mathcal{D} \setminus \mathcal{J}}} f(\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}_{\mathcal{D} \setminus \mathcal{J}} = f_\emptyset + \sum_{\mathcal{K} \prec \mathcal{J}} f_{\mathcal{K}} + f_{\mathcal{J}}, \quad (11)$$

where we integrate with respect to all $\theta_k, k \in \mathcal{D} \setminus \mathcal{J}$. Note, that as f is a function, $\mathbb{E}(f_{\mathcal{J}})$ is as well a function that depends on $\theta_{\mathcal{J}}$.

If we insert the recursive definition of the lower order terms $f_{\mathcal{K}}$ we obtain $f_{\mathcal{J}}$ as a linear combination of the $f_{\mathcal{K}}$ with $\mathcal{K} \prec \mathcal{J}$,

$$f_{\mathcal{J}} = \mathbb{E}(f_{\mathcal{J}}) + \sum_{\mathcal{K} \prec \mathcal{J}} (-1)^{|\mathcal{J}| - |\mathcal{K}|} \mathbb{E}(f_{\mathcal{K}}). \quad (12)$$

Having the ANOVA decomposition at hand we can compute global sensitivity values for the input parameters. We want to know the contribution of the variance of each combination of parameters to the variance of the solution. We decompose the variance of the solution by writing it as a sum of individual variances using the ANOVA representation and obtain

$$\mathbb{V}(f) = \sum_{\mathcal{J} \in \mathcal{P}(\mathcal{D})} \mathbb{V}(f_{\mathcal{J}}) \quad (13)$$

if the random variables are independent. Due to the linearity of the variance operator and Eq. (12) we can write the variance of each member as

$$\mathbb{V}(f_{\mathcal{J}}) = \mathbb{E}(f_{\mathcal{J}}^2) + \sum_{\mathcal{K} \prec \mathcal{J}} (-1)^{|\mathcal{J}| - |\mathcal{K}|} \mathbb{E}(f_{\mathcal{K}}^2) \quad (14)$$

taking into account that each member has zero mean.

The ratio

$$S_{\mathcal{J}} = \frac{\mathbb{V}(f_{\mathcal{J}})}{\mathbb{V}(f)} \quad (15)$$

is now a global sensitivity measure called the *main effect*. Certainly, all $S_{\mathcal{J}}$ are non-negative, and they sum up to 1. The main effect describes how much of the total variance can be explained by one parameter combination. For the untruncated ANOVA decomposition one has 2^d combinations covering interactions up to d^{th} order. This might get confusing for higher-dimensional problems. Hence we want to introduce another sensitivity measure, which describes the overall contribution of one single parameter to the variance of the solution. It is called the *total effect* and is computed by summing up all the main effects where the parameter is involved, i.e.,

$$S_{\theta_k} = \sum_{\mathcal{J} \in \mathcal{D}: k \in \mathcal{J}} S_{\mathcal{J}}. \quad (16)$$

A total effect $S_{\theta_k} = 1$ means that the whole variance in the solution is caused by θ_k , a total effect of $S_{\theta_k} = 0$ that θ_k is irrelevant.

3 Adaptive Sparse Grid Collocation Method

Spatially adaptive sparse grids were first introduced to the setting of uncertainty quantification in [7] and have been employed to both interpolation and quadrature in various settings even before, see [4, 19] and the references therein. The main idea of ASGC is to use an adaptively refined sparse grid function as an approximation of u for the unknown functional dependencies of the random input parameters $\boldsymbol{\theta}$, as described in Sect. 2, and extract statistical quantities using the approximation. The approximation is constructed

by interpolation, which means that every sparse grid point is a collocation node in terms of UQ, and a parameter combination in terms of numerical simulations.

3.1 Sparse Grids

To briefly recall the most important properties and to clarify our notation, we describe the basic principles of sparse grids in the following; see, e.g., [4, 18–20] for further details. Sparse grids are based on a hierarchical (and thus inherently incremental and adaptive) formulation of a one-dimensional basis, which is then extended to the d -dimensional setting via a tensor product approach.

As discussed in Sect. 2, we scale the stochastic parameter space Γ to the unit-hypercube and restrict ourselves to piecewise d -linear functions $f_{\mathcal{I}} : [0, 1]^d \rightarrow \mathbb{R}$. They are defined on an equidistant mesh with N grid points and mesh-width $h_{l_k} := 2^{-l_k}$ in dimension k .

We denote \mathbf{l} and \mathbf{i} as multi-indices with level and index for each dimension, $\|\mathbf{l}\|_1$ as the classical l^1 norm for vectors, and comparison operators on multi-indices component-wise. We can define grid points $\boldsymbol{\theta}_{\mathbf{l}, \mathbf{i}} = (i_1 2^{l_1}, \dots, i_d 2^{l_d}) \in [0, 1]^d$ on a hierarchy of grids uniquely by the set $\mathcal{I}_{\mathbf{l}}$ of level-index tuples (\mathbf{l}, \mathbf{i}) with

$$\mathcal{I}_{\mathbf{l}} := \{(\mathbf{l}, \mathbf{i}) : 1 \leq i_k < 2^{l_k}, i_k \text{ odd}, k = 1, \dots, d\},$$

omitting even-indexed ones.

To obtain the space of piecewise d -linear functions, we first define one-dimensional basis functions $\phi_{l,i}$, depending on a level l and an index i , out of the reference hat function $\phi(\theta) := \max(1 - |\theta|, 0)$ via translation and scaling as $\phi_{l,i}(\theta) := \phi(2^l \theta - i)$, see Fig. 1 (left) for the basis functions up to level 3. We then obtain d -dimensional basis functions $\phi_{\mathbf{l}, \mathbf{i}}$ as a product of the respective one-dimensional ones,

$$\phi_{\mathbf{l}, \mathbf{i}}(\boldsymbol{\theta}) := \prod_{k=1}^d \phi_{l_k, i_k}(\theta_k),$$

which are centered at the grid points $\boldsymbol{\theta}_{\mathbf{l}, \mathbf{i}}$, see Fig. 1.

We then define hierarchical increment spaces $W_{\mathbf{l}}$ for which the grid points are the Cartesian product of the one-dimensional ones on the respective one-dimensional levels as $W_{\mathbf{l}} := \text{span}(\{\phi_{\mathbf{l}, \mathbf{i}} : (\mathbf{l}, \mathbf{i}) \in \mathcal{I}_{\mathbf{l}}\})$ and functions $f_{\mathcal{I}_{\mathbf{l}}}(\boldsymbol{\theta}) \in W_{\mathbf{l}}$ as a sum of basis functions weighted by hierarchical coefficients $v_{\mathbf{l}, \mathbf{i}}$ (so-called surpluses),

$$f_{\mathcal{I}_{\mathbf{l}}}(\boldsymbol{\theta}) = \sum_{(\mathbf{l}, \mathbf{i}) \in \mathcal{I}_{\mathbf{l}}} v_{\mathbf{l}, \mathbf{i}} \phi_{\mathbf{l}, \mathbf{i}}(\boldsymbol{\theta}). \quad (17)$$

Figure 1 (center) shows the grids of the two-dimensional hierarchical increments W_l up to level 3 in each dimension. Note that in each W_l , all basis functions have supports with piecewise disjoint interiors.

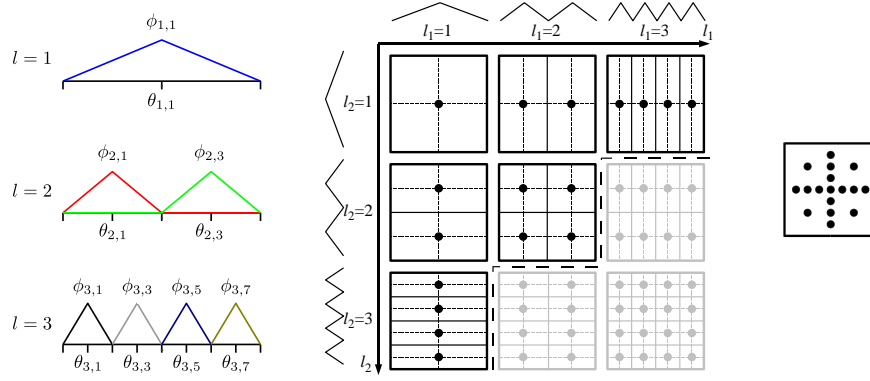


Fig. 1 One-dimensional basis functions up to level 3 (left), and tableau of hierarchical increments W_l up to level 3 in both dimensions (center). Leaving out the grayed-out W_l , we obtain the sparse grid of level 3 (right).

The hierarchical representation now allows one to select only those subspaces that contribute most to the overall solution. This can be done by an a priori selection (see [4] for details). We then obtain a sparse grid space such as

$$V_n^{(1)} := \bigoplus_{|l_1| \leq n+d-1} W_l,$$

which in this case is optimal with respect to both the L^2 -norm and the maximum-norm for suitably smooth functions. In the example in Fig. 1, we can neglect the gray W_l for $n = 3$, which leads to the regular (non adaptive) sparse grid in Fig. 1 (right). This reduces the number of collocation or grid points significantly from $O((2^n)^d)$ for full grids to $O(2^n n^{d-1})$ with only slightly deteriorated accuracy if the function f under consideration is sufficiently smooth, i.e., if the mixed second derivatives $|D^2 f| := \left| \frac{\partial^{2d}}{\partial \theta_1^2 \dots \partial \theta_d^2} f \right|$ are bounded.

In general, sparse grids are more beneficial for higher-dimensional approximation problems. But the hierarchical approach is also beneficial when every single collocation node is costly. Moreover, in their adaptive formulation, sparse grids are often able to reduce the number of grid points even more, especially when the function to be interpolated does not meet the smoothness requirements. We present spatially adaptive sparse grids in the next section.

3.2 Adaptive Refinement

We suggest to use adaptively refined grids to keep the number of collocation nodes and thus costly simulation runs as small as possible, and to be able to deal with sharp transitions in the quantity of interest. The main idea is to add new collocation nodes where the local error with respect to some metric is large. The optimal way would be to run the simulation for all new candidates and select the one which contributes most with respect to the quantity one wants to optimize. This is, of course, computationally not feasible in the context of UQ. Therefore we exploit the hierarchical structure of the method and use the hierarchical surpluses as an indicator for the local interpolation error. We thus select the most promising nodes of the current grid and create all their hierarchical descendants.

We propose three approaches to measure the contribution associated with the collocation points, from which the first is designed to minimize the interpolation error, the second refines the grid due to a local variance measure, and the third minimizes the error with respect to the expectation value. For this purpose, let $\mathcal{A} \subseteq \mathcal{I}$ be the admissible set of level index vectors (\mathbf{l}, \mathbf{i}) out of all collocation nodes \mathcal{I} , which are refinable, i.e. are leaf nodes in at least one direction. With \mathcal{I} we refer to an adaptively refined set of sparse grid collocation nodes and omit therefore the subscript used for regular (non-adaptive) sparse grids in the previous section. Furthermore, note that we applied the refinement criteria presented below individually and not sequentially.

1. *Absolute surplus refinement:* Let $f_{\mathcal{I}}$ be the sparse grid function which interpolates our solution. We use the coefficients $v_{\mathbf{l}, \mathbf{i}}$ as a local interpolation error measure [19] and refine the grid according to the largest one,

$$\max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |v_{\mathbf{l}, \mathbf{i}}|. \quad (18)$$

2. *Variance surplus refinement:* As proposed in [12], we interpret the hierarchical coefficients $w_{\mathbf{l}, \mathbf{i}}$ of the interpolant of the squared solution $f_{\mathcal{I}}^2 = \sum_{(\mathbf{l}, \mathbf{i}) \in \mathcal{I}} w_{\mathbf{l}, \mathbf{i}} \phi_{\mathbf{l}, \mathbf{i}}$ as a measure for the local variance. By refining the area where the local variance is large, we minimize the error in the global variance of the solution. Thus, we refine the collocation node (\mathbf{l}, \mathbf{i}) with the largest absolute coefficient

$$\max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |w_{\mathbf{l}, \mathbf{i}}|. \quad (19)$$

3. *Expectation value refinement:* The previous two are greedy criteria and may get stuck in local extrema as they just consider hierarchical coefficients for refinement. Moreover, since we want to assess a quadrature problem one should consider the contribution of each refinable node to its solution. Therefore, we propose a new refinement criterion, which refines the grid according to the largest contribution of a collocation node

to the expectation value, i.e.

$$\begin{aligned}
 \max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |\mathbb{E}(f_{\mathcal{I} \setminus \{(\mathbf{l}, \mathbf{i})\}}) - \mathbb{E}(f_{\mathcal{I}})| &= \max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} \int_{\Gamma} |f_{\mathcal{I} \setminus \{(\mathbf{l}, \mathbf{i})\}}(\boldsymbol{\theta}) - f_{\mathcal{I}}(\boldsymbol{\theta})| p(\boldsymbol{\theta}) d\boldsymbol{\theta} \\
 &= \max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} \int_{\Gamma} \left| \sum_{(\bar{\mathbf{l}}, \bar{\mathbf{i}}) \in \mathcal{I} \setminus \{(\mathbf{l}, \mathbf{i})\}} v_{\bar{\mathbf{l}}, \bar{\mathbf{i}}} \phi_{\bar{\mathbf{l}}, \bar{\mathbf{i}}}(\boldsymbol{\theta}) - \sum_{(\bar{\mathbf{l}}, \bar{\mathbf{i}}) \in \mathcal{I}} v_{\bar{\mathbf{l}}, \bar{\mathbf{i}}} \phi_{\bar{\mathbf{l}}, \bar{\mathbf{i}}}(\boldsymbol{\theta}) \right| p(\boldsymbol{\theta}) d\boldsymbol{\theta} \\
 &= \max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |v_{\mathbf{l}, \mathbf{i}}| \int_{\Gamma} \phi_{\mathbf{l}, \mathbf{i}}(\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (20)
 \end{aligned}$$

Observe that the integral $\int_{\Gamma} \phi_{\mathbf{l}, \mathbf{i}}(\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}$ reduces to $2^{-|\mathbf{l}|_1}$ if the random variables are independent and uniformly distributed on the unit hypercube. Thus, we refine the grid point (\mathbf{l}, \mathbf{i}) with maximum contribution

$$\max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |\mathbb{E}(f_{\mathcal{I} \setminus \{(\mathbf{l}, \mathbf{i})\}}) - \mathbb{E}(f_{\mathcal{I}})| = \max_{(\mathbf{l}, \mathbf{i}) \in \mathcal{A}} |v_{\mathbf{l}, \mathbf{i}}| 2^{-|\mathbf{l}|_1}. \quad (21)$$

This refinement criterion includes for each collocation node the volume of the basis function's support and describes therefore a local estimate for the quadrature error.

The sensitivity analysis is a quadrature problem on a sparse grid. Therefore we want to make two remarks here.

First, the sensitivity analysis described in Sect. 2.2 can be applied directly to sparse grids. Every ANOVA term but the constant one is a sparse grid function, which is obtained by dimension-wise integration. For all one-dimensional sub-grids, one just needs to sum over all hierarchical surpluses multiplied by the volume of their respective one-dimensional basis functions. The result is the coefficient of the lower-dimensional ANOVA term.

Second, Bungartz proposes in [3] to use balanced grids for quadrature. This makes sure that every collocation node has none or both children in every dimension, which balances the error cancellation. We therefore enforce balanced grids for all the criteria by applying a balancing step after refinement.

Finally, the complete ASGC pipeline for this task is shown in Fig. 2.

4 Peridynamics

The principle of this theory is that particles in a continuum interact with surrounding particles in a finite distance by exchanging forces. Some of these concepts are similar to concepts in molecular dynamics. We refer to the notation, used in [25], to describe the Peridynamics.

Fig. 3 shows the body \mathcal{R} of the continuum in a reference configuration with particles $P := \{p_i | i = 1, \dots, m\}$. Each particle p_i has an initial position

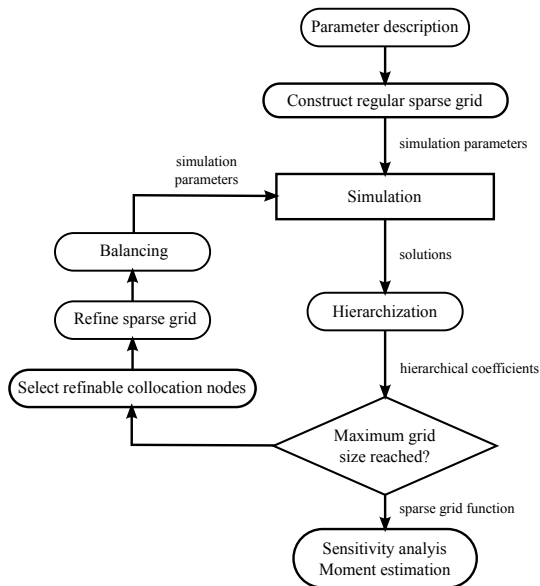


Fig. 2 The ASGC pipeline for a forward propagation problem in UQ.

x_i in the so-called reference configuration Ω and a neighborhood $B_\delta(x_i)$ with the radius δ , in which the particles interact by forces. Literally speaking, the particle does not “see” particles outside its own neighborhood and is not influenced by them. Within the neighborhood we consider the particles being connected by virtual bonds which can break during the simulation due to the displacement of the particles, i.e., when particle p_j leaves the fixed neighborhood $B_\delta(x_i)$. Computing the acceleration at position x in the reference configuration Ω at time t yields the integral equation

$$\varrho(x)\ddot{\mathbf{u}}(x,t) = \int_{B_\delta(x)} g(x',x,\mathbf{u}(\cdot,t))dx' + b(x,t). \quad (22)$$

Eq. (22) contains the mass density $\varrho(x)$ of the material at position x in the reference configuration Ω , \mathbf{u} is the displacement vector field, $b(x,t)$ denotes the external force at position x at time t , and g is the so-called pairwise force function.

Fig. 3 sketches the discretization in space by a collocation approach of the continuum by $P := \{p_i | i = 1, \dots, m\}$ particles on an equidistant lattice with mesh-widths Δx , Δy , and Δz . Each particle p_i has a surrounding volume $V_i = \Delta x \cdot \Delta y \cdot \Delta z$. To discretize the equation of motion (22), the set

$$\mathcal{F}_i = \{j \mid \|x_j - x_i\| \leq \delta, j \neq i\}, \quad (23)$$




Fig. 3 The body \mathcal{R} of the continuum with particles $P := \{p_i | i = 1, \dots, m\}$ at time $t = 0$ (the reference configuration), particle p_i at position x_i with the neighborhood $B_\delta(x_i)$, particle p_j at position x_j , the relative position $x_j - x_i$, and an exemplary surrounding volume $V = \Delta x \cdot \Delta y \cdot \Delta z$ on an equidistant lattice with mesh-widths Δx , Δy , and Δz .

is defined for each particle p_i . The set \mathcal{F}_i contains the indices of all particles, which are in the neighborhood $B_\delta(x_i)$ of particle p_i . This means that the distance between particle p_j and p_i is less or equal than δ with respect to the reference configuration Ω of the body.

We then obtain the semi-discrete Peridynamic equation,

$$\varrho(x_i) \ddot{\mathbf{u}}(x_i, t) = \sum_{j \in \mathcal{F}_i} g(x_j, x_i, \mathbf{u}(\cdot, t)) \tilde{V}_j + b(x_i, t), \quad (24)$$

with $\varrho(x_i)$ as the mass density function, g as the pairwise force function and b the external force. The volume \tilde{V}_j is the scaled volume V_j of particle p_j with the linear dimensionless scaling function

$$\nu(x_j, x_i) = \begin{cases} -\frac{1}{\Delta x} \|x_j - x_i\| + \left(\frac{\delta}{\Delta x} + \frac{1}{2}\right), & \delta - \frac{\Delta x}{2} \leq \|x_j - x_i\| \leq \delta \\ 1, & \|x_j - x_i\| \leq \delta - \frac{\Delta x}{2} \\ 0, & \text{otherwise.} \end{cases} \quad (25)$$

The scaling function ν is needed for particles p_j which are bound to particle p_i but are close to the neighborhood $B_\delta(x_i)$. A part of their volume is outside the neighborhood of particle p_i . The influence of such a particle differs from that of a particle, which lies completely within $B_\delta(x_i)$. If the distance $\|x_j - x_i\| = \delta$, then the volume V_j is simply scaled by 0.5, because nearly one half of the volume is inside and one half of the volume is outside $B_\delta(x_i)$.

4.1 Material Model

Eq. (22) describes the interaction between the particles, but gives no explicit information about the behavior of the simulated material. The modeling of the material is hidden in the kernel function g . To simulate brittle material, like ceramic or fused silica, the prototype microelastic brittle (PMB) material model is used. In the PMB material model the assumption is made that the pairwise force function for inner forces g depends only on the bond stretch s , which is defined by

$$s(x_j, x_i, \mathbf{u}(\cdot, t)) := \frac{\|\mathbf{u}(x_j, t) - \mathbf{u}(x_i, t) + x_j - x_i\| - \|x_j - x_i\|}{\|x_j - x_i\|}. \quad (26)$$

The easiest way to model failure is to let bonds break when they are stretched beyond a predefined constant value. The function s obviously attains positive values if the bond is under tension. An isotropic material has the property that it behaves identically in all directions.

Therefore the bond stretch is independent of the direction of $x_j - x_i$, and the pairwise force function g in a PMB material model is defined as

$$g(x_j, x_i, \mathbf{u}(\cdot, t)) = h(x_j, x_i, \mathbf{u}(\cdot, t)) \frac{\mathbf{u}(x_j, t) - \mathbf{u}(x_i, t) + x_j - x_i}{\|\mathbf{u}(x_j, t) - \mathbf{u}(x_i, t) + x_j - x_i\|}, \quad (27)$$

where h is a linear scalar-valued function which implements the behavior of the material and the decision if the bond is broken or “alive”. It is defined as

$$h(x_j, x_i, \mathbf{u}(\cdot, t)) = \begin{cases} c \cdot s(x_j, x_i, \mathbf{u}(\cdot, t)) \cdot \mu(x_j, x_i, \mathbf{u}(\cdot, t)), & \|x_j - x_i\| \leq \delta \\ 0, & \|x_j - x_i\| > \delta, \end{cases} \quad (28)$$

with c being the material dependent stiffness constant of the PMB material model and s the bond stretch (26). The function μ is an history dependent, scalar-valued function that models the inability of the material to “heal” broken bonds. It is defined as

$$\mu(x_j, x_i, \mathbf{u}(\cdot, t)) = \begin{cases} 1 & s(x_j, x_i, \mathbf{u}(\cdot, t)) < s_{00} - \alpha s_{\min}(t'), \forall 0 \leq t' \leq t, \\ 0 & \text{otherwise,} \end{cases} \quad (29)$$

where s_{00} is the critical stretch for bond failure in the material and $s_{\min}(t')$ is the minimal bond stretch up to t' , i.e. $s_{\min}(t') = \min_{x_j - x_i} s(x_j, x_i, \mathbf{u}(\cdot, t'))$. Hence the stiffness constant c and the critical stretch for bond failure s_{00} are material parameters in the PMB material model. However, the factor α is usually not considered as an material parameter, but fixed at 0.25 [25].

Nevertheless, α influences the elasticity of the material: increasing α leads to a higher critical stretch for the bonds and keeps them alive longer. This,

consequently, affects the local damage, which is defined at position x at time t as

$$\varphi(x, \mathbf{u}(\cdot, t)) = 1 - \frac{1}{V_{B_\delta(x)}} \int_{B_\delta(x)} \mu(x', x, u(\cdot, t)) dx'. \quad (30)$$

This yields to the total damage, our quantity of interest

$$c_t = \frac{1}{m} \sum_{i=1}^m \varphi(x_i, \mathbf{u}(\cdot, t)), \quad (31)$$

where $\varphi(x_i, \mathbf{u}(\cdot, t))$ is the relative number of bond failures for particle i at time t .

To study the influence of the uncertain parameters $\alpha =: \theta_1$, $\Delta x =: \theta_2$ and $K =: \theta_3$, the force the indenter exerts to the plate, to the total damage $c_t =: u$, our response function, we use a specific model $\mathcal{M}_{\text{Peridynamics}} = (\boldsymbol{\theta}, p(\boldsymbol{\theta}), F(\boldsymbol{\theta}), u)$.

5 Simulation

For the simulation we use a common peridynamic example which is described in [16, 17]. Figure 4 shows the CAD model of this experiment. The target is a homogeneous plate of fused silica with a diameter of 0.074 m and a height of 0.0025 m. The Peridynamics material parameters for fused silica are $s_{00} = 6.99 \cdot 10^{-6}$ and $c = 2.19 \cdot 10^{22}$. The indenter is a sphere with diameter 0.01 m and impacts the cylinder with a directed velocity of $v = 100$ m/s perpendicular to the surface of the target, see Fig. 5. The properties of the uncertain parameters we used for the simulations are shown in Tab. 1.

The particles are arranged in an equidistant grid with the same spacing Δx in all three dimension, a first parameter. Increasing Δx leads to a reduction of the particle density in the plate. A standard value for Δx in practice is 0.5 mm, which leads to 876,435 particles in the simulation and a total of 102,865,590 bonds. Decreasing Δx a lot gets very fast computationally unfeasible. Here, we consider $\Delta x \in [0.4, 0.6]$. Note, Δx defines as well the radius $\delta := 3\Delta x$ of the neighborball $B_{\delta(x_i)}$ of each particle. This means

Table 1 Range of the parameters that define the stochastic input space for the quantification of uncertainties.

Parameter	Min	Max	Unit	Distribution	Description
Δx	0.4	0.6	mm	$\mathcal{U}(0.4, 0.6)$	grid spacing
α	0	1	-	$\mathcal{U}(0, 1)$	model parameter for bond stretch failure
K	10^{12}	10^{20}	N/m ²	$\mathcal{U}(12, 20)$	magnitude of force on indenter

that by varying Δx the properties of the material of the plate vary as well. However, this interplay has yet not been quantified in peridynamic theory and is part of ongoing research.

The influence of the second parameter α has not yet been studied in detail. It is by definition restricted to the unit interval, but literature suggests to use $\alpha = 0.25$ as this leads to somehow realistic fracturing behavior. We want to especially study its influence on the peridynamic method in detail and set therefore the parameter range to $[0, 1]$.

For the external force K the *fix indent* command is used with a sphere that has a constant diameter of 0.01 m. The magnitude of the force K , which is exerted by the indenter, varies from 10^{12} up to 10^{20} on a logarithmic scale. The parameter range of it is chosen such that there is no damage in the plate for low K , and that the plate gets destroyed almost completely for large K , see Fig. 6.




Fig. 4 CAD model of the simulation configuration.




Fig. 5 Setting of the experiment.




Fig. 6 Simulation runs for $K = 10^{12}$ (left), $K = 10^{16}$ (center), and $K = 10^{20}$ (right) at the same time step $t = 300$, $\alpha = 0.25$, and $\Delta x = 0.0005$

Altogether, these three parameters define the stochastic space $\Gamma := [0.4, 0.6] \times [0, 1] \times [12, 20]$ for the UQ-setting. They are independent and uniformly distributed by definition, which reduces the transformation F according to Sect. 2 to be a linear scaling. As the peridynamic simulation framework we use LAMMPS [16], and SG^{++} [18] for sparse grids and UQ.

6 Numerical Results

We have studied our approach for both an analytical example and the peridynamic simulation setting. The analytical example is close to the peridynamic simulation with just the first two parameters, but does not exhibit the “real-world” issues as we will observe for the peridynamic simulations. It validates the ASGC method and shows its faster convergence with respect to the absolute error of the expectation value compared to vanilla Monte Carlo and quasi-Monte Carlo. For comparisons to polynomial chaos methods we refer to [12].

In the following numerical examples we use the mean squared error with respect to M uniformly drawn Monte Carlo samples as an error indicator for the interpolation quality of an interpolant $f_{\mathcal{I}}$ with $|\mathcal{I}| = N$ collocation nodes, i.e.

$$\epsilon_I := \frac{1}{M} \sum_{i=1}^M (f_{\mathcal{I}}(\boldsymbol{\theta}_i) - u(\boldsymbol{\theta}_i))^2. \quad (32)$$

The error in the expectation value is computed as

$$\epsilon_E := |\mathbb{E}(f_{\mathcal{I}}) - \mathbb{E}_r(u)|, \quad (33)$$

where the reference value $\mathbb{E}_r(u)$ is either the analytical solution of the quadrature problem or obtained by the sample mean over all Monte Carlo observations, i.e.

$$\mathbb{E}_r(u) = \frac{1}{M} \sum_{i=1}^M u(\boldsymbol{\theta}_i). \quad (34)$$

The respective reference value for the variance $\mathbb{V}_r(u)$ is computed accordingly.

We compared the convergence of the ASGC method with three other non-intrusive forward propagation methods: Monte Carlo with uniformly drawn samples, quasi-Monte Carlo with Sobol-sequences, and a full-grid interpolant with hierarchical basis where the collocation nodes are equidistantly distributed over the complete stochastic space. With respect to Monte Carlo, we generated a training set, from which we generated 20 Monte Carlo paths by permutation. We estimated different expectation values for each path by adding samples iteratively to it. By averaging over all paths we reduce the

variance of Monte Carlo and get one sufficiently smooth Monte Carlo path, which we used for comparison.

6.1 Analytical Example

We consider the function

$$u(\theta_1, \theta_2) = \arctan(50(\theta_1 - 0.35)) + \frac{\pi}{2} + 4\theta_2^3 + e^{\theta_1\theta_2-1} \quad (35)$$

with two stochastic parameters $\theta_1, \theta_2 \sim \mathcal{U}(0, 1)$. A graphical representation is given in Fig. 7. The function u is designed to represent the characteristics




Fig. 7 3-dimensional plot of $u(\theta_1, \theta_2)$.

of the peridynamic simulation just after the indenter has hit the plate. We can think of the parameters θ_1 and θ_2 as representatives for $1 - \alpha$ and Δx . Their functional dependencies are motivated as follows:

1. For high α the damage will be low, as the modeled material is rather hard compared to the softer material for lower α . In-between we expect a sharp transition at $\theta_1 = 0.35$, which we model with a shifted and steep arctan function depending on θ_1 .

2. The term depending just on θ_2 is motivated by the fact that a particle-based numerical simulation should converge to the true solution when one increases the number of particles in the plate. Increasing the number of particles means decreasing the initial distance Δx between them. We want to reach convergence with $\theta_2 \rightarrow 0$. We expect to overestimate the damage for too few particles and use therefore a cubic dependence on θ_2 .
3. The last term in u models an unknown dependency between the two parameters.

We first construct 4 interpolants of u : the regular sparse grid interpolant, and the adaptively refined sparse grid interpolants using the three refinement strategies of Sect. 3.2. Considering the refinement strategies, we just refined collocation nodes with a hierarchical surplus larger than $\epsilon := 10^{-10}$, per iteration we refined $\min\{10, N \frac{5}{100}\}$ of the collocation nodes with the largest contribution and stopped if either there were no more points to be refined or the upper limit of 3000 collocation nodes was reached. To estimate their interpolation quality we used 10^6 Monte Carlo test samples. The error in the expectation value is computed as the absolute difference to the up to 14 digits exact reference value $\mathbb{E}_r(u) = 3.51449126644638$. For the Monte Carlo method we computed 2^{14} training samples, and the same number of quasi-Monte Carlo samples.

With respect to the mean squared interpolation error ϵ_I , as shown in Fig. 8, we can observe that the full grid method converges as expected with $\mathcal{O}(N^{-2})$. The regular sparse grid method shows almost the same error as the full grid method on the same discretization level. And as it requires less grid points per level, it converges therefore faster with an upper bound of $\mathcal{O}(N^{-5} \log(N)^6)$ [4]. The adaptive methods converge with the same speed as the regular sparse grid method. However, the absolute interpolation error is three orders of magnitude lower compared to the regular sparse grid after the fourth refinement iteration with 20 collocation nodes. The surplus and the expectation value refinement approaches show similar behavior while the variance-based approach is one order of magnitude worse.

For the absolute error in the expectation value estimation, Fig. 9, we observe two important things: First, even though the interpolation error is monotonically decreasing, the error in the expectation value is not. Between the regular sparse grid of level 2 (25 collocation nodes) and level 3 (81) the error grows and cannot be fully recovered by the grid of level 4. The reason is cancellation of the integral around the steep transition in direction of θ_1 . A level 2 grid is too coarse to represent the transition well, but due to cancellation, the error in the expectation value is low. This cancellation effect is destroyed on the next level because of a better approximation of the steep transition on just one side, which is once more compensated on the next level. This effect is known [3, 19] and a major problem for adaptively refined grids. To overcome it to some extent we introduced a balancing step after refinement. But as we can see from the convergence plots it does not solve the problem completely. However, the expectation value refinement

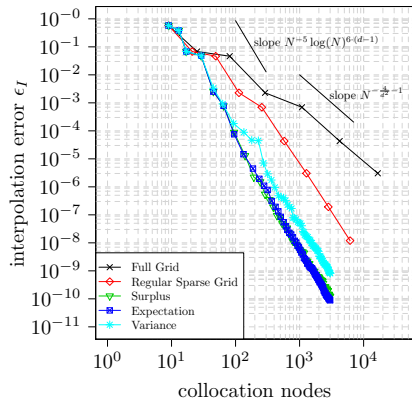


Fig. 8 Interpolation error measured by the mean squared error at 10^6 Monte Carlo test samples. The black lines in the right upper corner show the slopes for the regular sparse grid and the full grid [4].

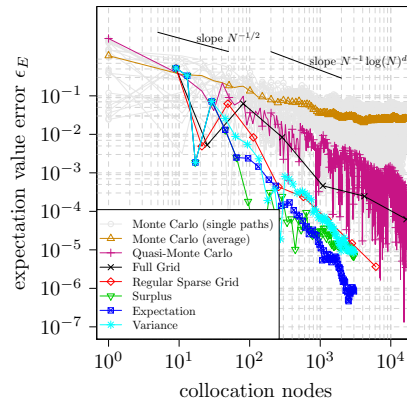


Fig. 9 Absolute error of the estimated expectation value for different methods. The black lines in the upper part of the figure show the slopes for Monte Carlo and quasi-Monte Carlo [5].

strategy reaches an accuracy of at least one order of magnitude higher than the other grid-based methods. The Monte Carlo methods show a significantly slower convergence. For 20 collocation nodes the accuracy of the grid-based methods is already one and for 2000 collocation nodes four order of magnitudes higher. The quasi-Monte Carlo method converges as expected with $\mathcal{O}(N^{-1} \log(N)^d)$ [5], which is still significantly slower than the sparse grid methods.

With respect to the global sensitivity values we computed main effects of $S_{\theta_1} = 0.577934314364$, $S_{\theta_2} = 0.421292846812$, and $S_{\theta_1, \theta_2} = 0.0007728388246$, which sum up to 1. Hence, most of the variance of u is explained by the arctan and the cubic term. The interactive term has almost no influence on the variance.

Starting with these promising results we assess the peridynamic example in the following section.

6.2 Peridynamic Example

For this example we use the simulation described in Sect. 5. We split the analysis in two parts. First we assess the 2-dimensional example where we vary α and Δx . After that we add the magnitude of force K to the setting and discuss the resulting 3-dimensional setting.

Both of the simulations are, of course, time dependent. Hence, we have for each time step a different sparse grid interpolant. For the 2-dimensional

setting we restrict ourselves to the study of the uncertainty at time step $t = 300$ of the simulation, at which the variance is largest, see Fig. 10. The




Fig. 10 The expectation value and the standard deviation over time, estimated using the ASGC method.

Monte Carlo and the quasi-Monte Carlo training sets for comparison contain 2500 samples each.

We determine the reference expectation value $\mathbb{E}_r(u)$ measured with $M = 6000$ Monte Carlo test samples as the exact value is unknown. The Monte Carlo samples have been drawn uniformly according to the distribution of the uncertainty in the parameters. However, the quality of this estimation clearly affects the convergence with respect to the expectation value of the compared methods. The evolution of the estimated reference expectation value $\mathbb{E}_r(u)$ of the Monte Carlo test set is shown in Fig. 11. We assume that the variance of u




Fig. 11 The Monte Carlo simulation confidence interval obtained with the central limit theorem with a confidence level of $\beta = 0.1$.

exists and apply the central limit theorem to estimate an upper bound for the error of $\mathbb{E}_r(u)$ [5] for the most common confidence levels $\beta \in \{0.1, 0.05, 0.01\}$ given the empirical variance $\mathbb{V}_r(u) = 0.06304$, i.e.,

$$|\mathbb{E}_r(u) - \mathbb{E}(u)| \leq \Phi_{1-\frac{\beta}{2}} \sqrt{\frac{\mathbb{V}_r(u)}{M}} \quad (36)$$

The resulting upper bounds are shown in Tab. 2. In these circumstances we can not expect that the accuracy of expectation value estimates of the competing methods to be reliable beyond $5.33 \cdot 10^{-3}$.

Table 2 Upper bound for the error of $\mathbb{E}_r(u)$ for different confidence levels, $M = 6000$ test samples, which have an empirical variance of $\mathbb{V}_r(u) = 0.06304$.

confidence level	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.01$
upper error bound	$5.33 \cdot 10^{-3}$	$6.35 \cdot 10^{-3}$	$8.35 \cdot 10^{-3}$

With this in mind, we focus now on the convergence of the expectation value at time step $t = 300$, shown in Fig. 12, and investigated the absolute error in the expectation value, see Eq. (33).

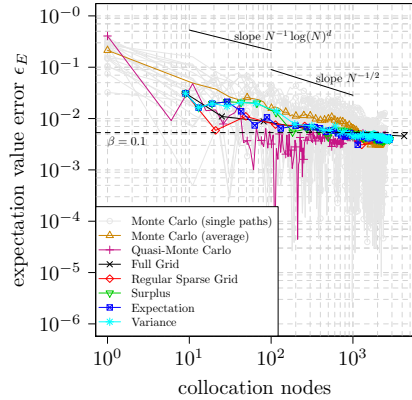


Fig. 12 Absolute error of the estimated expectation value for different methods. The black dashed line marks the upper bound for the error of the reference value at a confidence level of $\beta = 0.1$.




Fig. 13 Interpolation error measured at $t = 300$ by the mean squared difference at 6000 Monte Carlo test samples (uniformly distributed). The error has a local minimum between 300 and 1000 collocation nodes.

The error convergence of the expectation value estimation for most of the sparse grid methods is, in contrast to the analytical example, similar to that of the Monte Carlo results, while showing a lower error. The regular sparse

grid approach performs even a bit better than the adaptively refined ones for level 2, which is the same behavior as we have seen for the analytical example. The results for the adaptive sparse grids show that the criterion for adaptive refinement is very critical up to 200 collocation nodes. While the strategy based on the optimization of the expectation value performs best, its greedy counterparts, the absolute surplus and the variance surplus, cannot keep up and get stuck in local features.

However, the sparse grid methods, independent of the refinement criteria, reach a saturation point at about 110 collocation nodes. Runs with quasi-Monte Carlo samples using Sobol sequences show the same behavior, though reaching the trusted accuracy earlier at about 70 collocation nodes but oscillating more around the saturation value until 200 nodes. The Monte Carlo method reaches it with more than 1000 nodes, too. In contrast to the regular sparse grid and the ASGC methods, the Monte Carlo methods keep oscillating in a range of two order of magnitudes around the saturation point even for large numbers of collocation nodes without reducing the error. Furthermore, the hierarchical representation of sparse grids provides the means to easily detect the early saturation here, which is much more costly for the Monte Carlo approaches due to the highly oscillating behavior.

The early saturation point is also visible in the convergence plot of the interpolation error, see Fig. 13. We can observe a saturation for the error convergence, now for around 300 collocation points, which is much too early to be credited to numerical issues. Admittedly, we did not expect this behavior beforehand. The quantity of interest we study is the mean over the damage of each particle, which we expected to be smooth. Furthermore, the interpolation error should not increase spending more samples as it is the case for all grid-based methods. These observations suggest noise or at least high frequency oscillations in the quantity of interest. This would also explain why the Monte Carlo approaches stagnate with large oscillations, and why the error is smallest for grids, which cover the parameter range more regularly.

To support this suggestion we studied the hierarchical surpluses of the regular sparse grid and the full grid, see Fig. 14. For sufficiently smooth functions the hierarchical surpluses converge to zero with increasing level. This behavior can be exploited to find non-smooth dependencies in individual dimensions [3]. In the peridynamic simulation the interquartile range of the hierarchical surpluses decreases as expected up to level 6 for regular sparse grids and up to level 5 for the full grid. However, the hierarchical surpluses explode for higher levels and the interpolation error increases. These collocation nodes with the largest hierarchical surpluses lie mostly on the border of the domain and have a high level in at least one dimension. We investigated the collocation node with the highest hierarchical surplus 0.6397 added in the 8th iteration of the regular sparse grid. It is located at $\alpha = 1$ and $\Delta x = 0.76953125$. Its left and right neighbor in direction of Δx are spatially very close, i.e. 2^{-8} , due to the full grid resolution of the sparse grid on the




Fig. 14 Box plot of the hierarchical surpluses which are added when incrementing the level for the regular sparse grid (left) and the full-grid (right). The gray boxes mark the interquartile range, the black line within the boxes is the median, the so called “whiskers” mark the range where surpluses are located, which do not differ more than 1.5 times the interquartile range from the median. The surpluses, which lie outside of these ranges, are marked as circles. Note the exploding values of the hierarchical surpluses for increment iteration 7–8 for the regular sparse grid method, and on level 1 and 6 for the full grid method.

border. The quantity of interest, the average damage on the plate, for the left neighbor is approximately 0.2, for the right one is 0.21, while for the node itself is roughly 0.85. This explains the large hierarchical surplus and supports the statement of, at least, local high-frequency oscillations or, in other words, instabilities for extreme parameter combinations. The expectation value estimation is not influenced by these oscillations as the support of these high-level basis functions is rather small.

To investigate this behavior further, we show the adaptively refined sparse grid with 207 collocation nodes using the expectation value refinement strategy and the corresponding sparse grid function in Fig. 15 and 16, respectively. In Fig. 15 we can observe that most of the collocation nodes are spent in the regime $\alpha \in [0.25, 0.75]$. In between there is a transitional phase where the total damage decreases from 0.9 to 0.2. This is also the range in which most oscillations or noise occur, see the test samples with the highest errors in Fig. 17. The two steep slopes of the interpolant at about $\Delta x = 0.1$ and 0.5 have been successfully identified and resolved by the refinement strategy.

We furthermore investigated the interpolant and the error for three non-adaptive, regular sparse grids of level 1, 3 and 5, which are shown in Fig. 18. In the top left, the piecewise linear interpolant for a sparse grid of level 1 which consists of 9 collocation nodes is shown. As the real function varies much less in Δx compared to α , it can be much better approximated by a linear function in the Δx direction. This is directly reflected by the region in which most of the error occurs (bottom left).




Fig. 15 An adaptively refined sparse grid with expectation value optimization refinement strategy. Color and radius of the collocation nodes represent the hierarchical coefficients. The corner nodes have the highest contribution to the function while the coefficients of the inner nodes are oscillating around zero.




Fig. 16 Sparse grid function for the collocation nodes in Fig. 15. Note the steep slopes at $\Delta x \approx 0.1$ and 0.5 . These two regions are well resolved by the adaptivity criterion and are therefore not artifacts of the sparse grid function but the behavior of the true solution.




Fig. 17 The plot represents the largest local errors of the sparse grid interpolant represented in Fig. 16, using the 993 test samples, which explain 53% of the interpolation error ϵ_I . Plenty of samples which are close to each other have errors (colors) oscillating around zero, which indicates highly frequent oscillations or noise in that regime.

The transition in the α -direction is far from linear in the lower third of the parameter space. With a level 3 grid, however, the transition can now be better represented (center column). There remain mainly two vertical clusters, in the middle and on the left, and a horizontal region of large error. These are exactly the bends that can be seen for a higher sampling resolution on level 5 in the upper right plot. The remaining region of large error is spread out throughout vast parts of the domain.

With respect to the peridynamic theory and the sensitivities we want to study, several areas with different influence of Δx for fixed α can be roughly classified in Fig. 16. First, there are two regimes, $\Delta x \approx 0.5$ and $\Delta x \approx 0.1$, where the total damage changes rapidly with varying Δx . Increasing the particle density (reducing Δx) starting from $\Delta x = 1$, we observe similar behavior of the damage until about $\Delta x = 0.5$. One would expect that increasing the density leads to a convergence of the damage behavior (propagation of the displacement waves and cracks, amount of damage). This can be observed in simulations for $\Delta x < 0.5$, but fails for very high particle densities ($\Delta x \approx 0.1$), which is also reflected in simulations where the damage propagation changes significantly for very high particle densities. The latter effect is yet to be investigated further.

Considering α , there are two regions with less sensitivity on Δx . For $0.1 \leq \alpha \leq 0.35$ and $\alpha \geq 0.85$, the sensitivity on the particle density is very small, and the overall damage in the simulation during the penetration

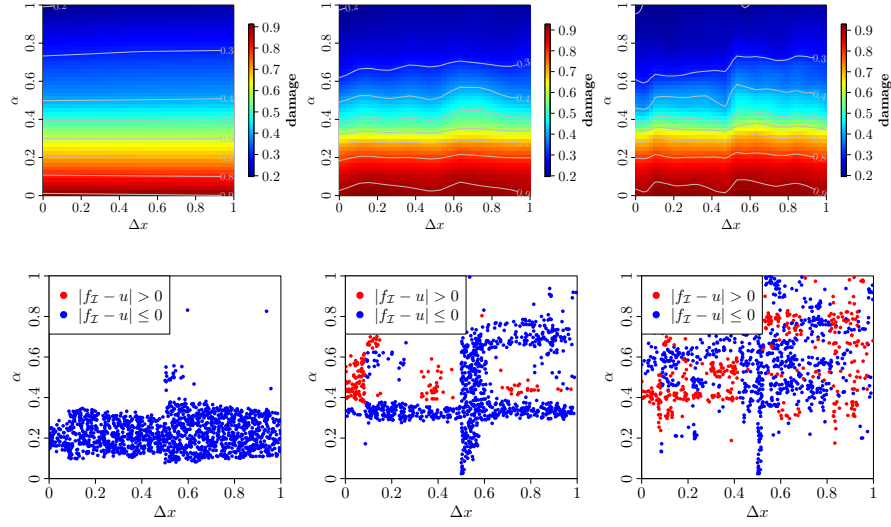


Fig. 18 Sparse grid functions for sparse grids with boundary points for levels 1, 3 and 5 (top), and the Monte Carlo test samples which cause at least 50% of the interpolation error (978, 743 and 784, bottom).

of the indenter varies very little. Silling and Askari proposed in [25] to choose $\alpha = 0.25$, for which the sensitivity on Δx is relatively small. In other regions of α , the influence of the particle density is significant, leading to variations of up to more than 15% of the total damage. The typical value $\alpha = 0.25$ delivers a relatively stable behavior of the damage with respect to Δx . Considering peridynamic simulations, this analysis enables us to derive some heuristics about appropriate ranges of the particle density for chosen values of α .

The standard deviation of the solution was already shown in Fig. 10. It reaches its maximum in the penetration phase of the projectile. In this phase ($0 < t < 400$) 96–99% of the total variance is explained by the variance of α , see the left part of Fig. 19. For $t \geq 400$ the second-order interactions between Δx and α become more important and explain up to 10% of the variance. The main effect of Δx has its maximum of 5% at time step 500. However, the sensitivity values do not show the oscillating behavior of the solution. This is not surprising, as they are comparatively small in the global context.

In the three-dimensional setting we observe similar behavior with respect to the convergence. Concerning the sensitivity analysis, the dominant parameter is the force parameter K as its total effect is $S_K = 0.995580573617$. The total effects $S_\alpha = 0.0135270393259$ and $S_{\Delta x} = 0.0021731614627$ are rather small in comparison. This is not surprising considering the choice of the parameters and their ranges. The force on the indenter must have the highest influence. However, the main effect of α has an impact of 12.9% at time step 100 when the indenter hits the plate. After the impact the sensitivity with

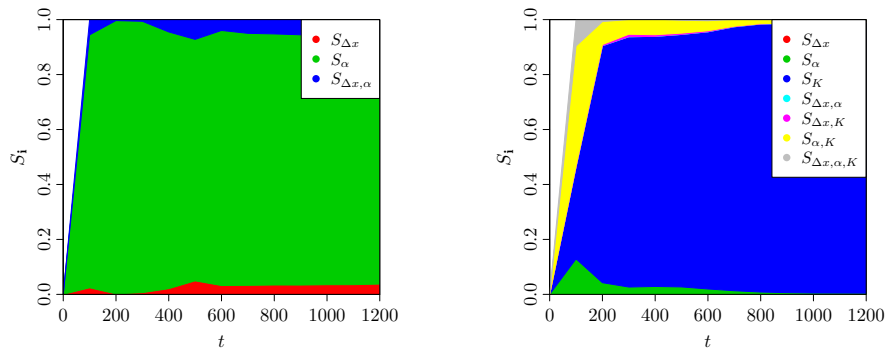


Fig. 19 Main effects and higher-order interactions of α , Δx and K for the 2-dimensional (left) and the three-dimensional (right) peridynamic UQ-setting.

respect to α reduces almost linearly until it vanishes after time step 800. Therewith the main effect of α directly depicts the behavior of the indenter itself. We want to recall that this values have been computed without additional samples and without introducing additional numerical errors. This is a great advantage over Monte Carlo since to obtain sensitivity values one would need $M(d + 2)$ number of samples [22]. Every sample is costly in Peridynamics and therefore it is unfeasible to do Monte Carlo based sensitivity analysis.

7 Conclusions

In this paper we have shown the first application of adaptive sparse grids for uncertainty quantification and sensitivity analysis in a peridynamic simulation setting. We have simulated the impact of a high-speed projectile on a ceramic plate using the PMB model of Peridynamics with two and three uncertain model parameters. The first two parameters were α , which describes the elasticity of the material, and Δx , which models the particle density in the plate. As a third parameter, we have considered the force K exerted by the projectile.

With regard to the adaptive sparse grid collocation method, we have introduced a new refinement criterion that is motivated by the UQ setting, and we have shown that it can be very effective compared to traditional and more greedy strategies. For the peridynamic setting, adaptivity did not pay off as much in terms of the number of collocation nodes as we expected, as it is known from other scenarios, and as the analytical example with simi-

lar behavior suggested. This was due to unexpected noise in the quantity of interest.

For both the 2- and the 3-dimensional simulation setting, the sparse grid methods outperformed vanilla Monte Carlo in terms of moment estimation by up to one order of magnitude for small numbers of collocation nodes. For the analytical example, this holds for quasi-Monte Carlo based on Sobol-sequences, too. For the peridynamic setting the accuracy was similar than that of quasi-Monte Carlo while oscillating much less. This is promising for UQ settings where the number of costly simulation runs has to be kept low.

In the simulation setting, we have observed an early saturation of the error convergence for the expectation value, which is due to noise or highly frequent oscillations in the transition phase between high and low damage on the plate with respect to changes in the particle density. Our approach allows us to identify parameter regions that show high sensitivity to the choice of parameters, which means unstable behavior from a simulation point of view. Furthermore, we have successfully computed reliable global sensitivity values for the uncertain parameters using the unanchored ANOVA decomposition method on sparse grids. We did this without increasing the number of samples and without introducing additional numerical errors through binning.

In contrast to stochastic approaches, sparse grids have several advantages. The hierarchical surpluses can be directly used to detect outliers or instabilities in the simulation, to detect an early saturation, and for adaptive refinement. Having constructed the sparse grid interpolant, it can be used to compute moments and sensitivity values without introducing further errors.

Considering adaptivity, the effects of the different refinement strategies have to be studied in more detail, especially for higher-dimensional settings. To overcome the problem of non-smooth dependencies in the UQ context as observed for peridynamic simulations, we propose to use a regression-based reconstruction of the response function instead of interpolation which has shown promising results in first tests. Another important point of future work will be to extend the current scenario to a higher-dimensional setting. We aim to add the bulk modulus and the critical stress intensity factor to study material properties of solids. Especially from an engineering point of view such a setting would be of major interest because the studied materials do not need to exist. Moreover, connecting the material properties with the initial mesh width would allow to study their relationship in detail.

References

1. M. Adams, A. Lashgari, B. Li, M. McKerns, J. Mihaly, M. Ortiz, H. Owhadi, A.J. Rosakis, M. Stalzer, and T.J. Sullivan, *Rigorous model-based uncertainty quantification with application to terminal ballistics Part II. Systems with uncontrollable inputs and large scatter*, *J. of the Mechanics and Physics of Solids* **60** (2012), no. 5, 1002–1019.

2. R. Archibald, R. Deiterding, and J. Jakeman, *Extending adaptive sparse grids for stochastic collocation to hybrid parallel architectures*, (2012), no. 4, 2–6.
3. H.-J. Bungartz and S. Dirnstorfer, *Multivariate quadrature on adaptive sparse grids*, *Computing* **71** (2003), no. 1, 89–114 (English).
4. H.-J. Bungartz and M. Griebel, *Sparse Grids*, *Acta Numerica* (2004), no. 13, 1–123.
5. Russel E Caffisch, *Monte carlo and quasi-monte carlo methods*, *Acta numerica* **7** (1998), 1–49.
6. R. I. Cukier, C. M. Fortuin, K. E. Shuler, A. G. Petschek, and J. H. Schaibly, *Study of the sensitivity of coupled reaction systems to uncertainties in rate coefficients. I Theory*, *Chem. Phys.* **59** (1973), no. 8, 3873–3878.
7. B. Ganapathysubramanian and N. Zabaras, *Sparse grid collocation schemes for stochastic natural convection problems*, *Comput. Phys.* **225** (2007), no. 1, 652–685.
8. M. Griebel, *Sparse grids and related approximation schemes for higher dimensional problems*, *Foundations of Computational Mathematics*, Cambridge University Press, 2005, pp. 106–161.
9. Michael Griebel and Markus Holtz, *Dimension-wise integration of high-dimensional functions with applications to finance*, 2009.
10. J. Jakeman, R. Archibald, and D. Xiu, *Characterization of discontinuities in high-dimensional stochastic problems on adaptive sparse grids*, *Comput. Phys.* **230** (2011), no. 10, 3977–3997.
11. A. Kidane, A. Lashgari, B. Li, M. McKerns, M. Ortiz, H. Owhadi, G. Ravichandran, M. Stalzer, and T.J. Sullivan, *Rigorous model-based uncertainty quantification with application to terminal ballistics, part I: Systems with controllable inputs and small scatter*, *J. of the Mechanics and Physics of Solids* **60** (2012), no. 5, 983–1001.
12. X. Ma and N. Zabaras, *An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations*, *Comput. Phys.* **228** (2009), no. 8, 3084–3113.
13. F. Nobile, R. Tempone, and C. G. Webster, *A sparse grid stochastic collocation method for partial differential equations with random input data*, *SIAM J. Numer. Anal.* **46** (2008), no. 5, 2309–2345.
14. S. Oladyshkin, H. Class, R. Helmig, and W. Nowak, *A concept for data-driven uncertainty quantification and its application to carbon dioxide storage in geological formations*, *Adv. Water Resour.* **34** (2011), no. 11, 1508–1518.
15. H. Owhadi, C. Scovel, T. J. Sullivan, M. McKerns, and M. Ortiz, *Optimal uncertainty quantification*, (2010).
16. M. L. Parks, R.B. Lehoucq, S. J. Plimpton, and S. Silling, *Implementing peridynamics within a molecular dynamics code*, *Comput. Phys. Commun.* **179** (2008), 777–783.
17. M.L. Parks, P. Seleson, S. J. Plimpton, S. Silling, and R.B. Lehoucq, *Peridynamics with LAMMPS: A User Guide v0.3 Beta*, SANDIA REPORT, Sandia National Laboratories, November 2011.
18. D. Pflüger, *Spatially adaptive sparse grids for high-dimensional problems*, Verlag Dr. Hut, München, August 2010.
19. ———, *Spatially adaptive refinement*, *Sparse Grids and Applications*, LNCSE, Springer, October 2012, pp. 243–262.
20. D. Pflüger, B. Peherstorfer, and H.-J. Bungartz, *Spatially adaptive sparse grids for high-dimensional data-driven problems*, *Journal of Complexity* **26** (2010), no. 5, 508–522.
21. M. Rosenblatt, *Remarks on a Multivariate Transformation*, *The Annals of Mathematical Statistics* **23** (1952), no. 1952, 470–472.
22. A. Saltelli, P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola, *Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index*, *Computer Physics Communications* **181** (2010), no. 2, 259–270.
23. Sandia National Laboratories, *LAMMPS Molecular Dynamics Simulator*.

24. S. Silling, *Reformulation of Elasticity Theory for Discontinuities and Long-Range Forces*, SANDIA REPORT SAND98-2176, Sandia National Laboratories, 1998 1998.
25. S. Silling and E. Askari, *A meshfree method based on the peridynamic model of solid mechanics*, *Computers & Structures* **83** (2005), no. 17-18, 1526–1535.
26. I.M. Sobol, *Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates*, *Math. Comput. Simulat.* **55** (2001), no. 1–3, 271–280.
27. D. Stirzaker, *Elementary probability*, Cambridge University Press, 2003.
28. X. Wan and E. Karniadakis, *Multi-Element Generalized Polynomial Chaos for Arbitrary Probability Measures*, *SIAM J. Sci. Comput.* (2006), 901–928.
29. D. Xiu and E. Karniadakis, *The Wiener–askey polynomial chaos for stochastic differential equations*, *SIAM J. Sci. Comput.* **24** (2002), no. 2, 619–644.
30. ———, *Modeling uncertainty in flow simulations via generalized polynomial chaos*, *Comput. Phys.* **187** (2003), no. 1, 137–167.
31. G. Zhang, D. Lu, M. Ye, M. Gunzburger, and C.G. Webster, *An adaptive sparse-grid high-order stochastic collocation method for Bayesian inference in groundwater reactive transport modeling*, *Adv. Water Resour.* **49** (2013), no. 10, 6871–6892.