

\mathcal{H} -MATRIX BASED SECOND MOMENT ANALYSIS FOR ROUGH RANDOM FIELDS AND FINITE ELEMENT DISCRETIZATIONS*

J. DÖLZ[†], H. HARBRECHT[†], AND M. D. PETERS[†]

Abstract. We consider the efficient solution of strongly elliptic partial differential equations with random load based on the finite element method. The solution’s two-point correlation can efficiently be approximated by means of an \mathcal{H} -matrix, in particular if the correlation length is rather short or the correlation kernel is nonsmooth. Since the inverses of the finite element matrices which correspond to the differential operator under consideration can likewise efficiently be approximated in the \mathcal{H} -matrix format, we can solve the correspondent \mathcal{H} -matrix equation in essentially linear time by using the \mathcal{H} -matrix arithmetic. Numerical experiments for three-dimensional finite element discretizations for several correlation lengths and different smoothness are provided. They validate the presented method and demonstrate that the computation times do not increase for nonsmooth or shortly correlated data.

Key words. finite element method, \mathcal{H} -matrix arithmetic, uncertainty quantification, nonlocal operators

AMS subject classifications. 60H35, 65C30, 65N30

DOI. 10.1137/16M1074813

1. Introduction. A lot of problems in science and engineering can be modeled in terms of strongly elliptic boundary value problems. While these problems are numerically well understood for input data which are given exactly, these input data are often only available up to a certain accuracy in practical applications, e.g., due to measurement errors or tolerances in manufacturing processes. In recent years, it has therefore become more and more important to take these inaccuracies in the input data into account and model them as random input parameters.

The Monte Carlo approach (see, e.g., [12] and the references therein) provides a straightforward approach to dealing with these random data, but it has a relatively slow convergence rate which is only in the sense of the root mean square error. This, in turn, means that a large number of samples have to be generated to obtain computational results with an acceptable accuracy, whereas the results still have a small probability of being too far away from the true solution. Therefore, in the past several years multiple deterministic approaches to overcoming this obstacle have been presented. For instance, random loads have been considered in [52, 57], random coefficients in [1, 2, 13, 18, 20, 42, 47, 49], and random domains in [36, 59].

For a domain $D \subset \mathbb{R}^d$ and a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, we consider the Dirichlet problem

$$(1) \quad \left. \begin{array}{ll} \mathcal{L}u(\omega, \mathbf{x}) = f(\omega, \mathbf{x}) & \text{for } \mathbf{x} \in D, \\ u(\omega, \mathbf{x}) = 0 & \text{for } \mathbf{x} \in \Gamma := \partial D \end{array} \right\} \mathbb{P}\text{-almost surely}$$

*Submitted to the journal’s Computational Methods in Science and Engineering section May 11, 2016; accepted for publication (in revised form) March 16, 2017; published electronically July 6, 2017.

<http://www.siam.org/journals/sisc/39-4/M107481.html>

Funding: This work was supported by the Swiss National Science Foundation (SNSF) through the project “ \mathcal{H} -matrix based first and second moment analysis.”

[†]Departement Mathematik und Informatik, Fachbereich Mathematik, Universität Basel, 4051 Basel, Switzerland (Juergen.Doelz@unibas.ch, Helmut.Harbrecht@unibas.ch, Michael.Peters@unibas.ch).

with random load $f(\omega, \mathbf{x})$ and a strongly elliptic partial differential operator \mathcal{L} of second order.

We can compute the solution’s mean

$$\mathbb{E}_u(\mathbf{x}) := \int_{\Omega} u(\omega, \mathbf{x}) \, d\mathbb{P}(\omega)$$

and also its two-point correlation

$$\text{Cor}_u(\mathbf{x}, \mathbf{y}) := \int_{\Omega} u(\omega, \mathbf{x})u(\omega, \mathbf{y}) \, d\mathbb{P}(\omega)$$

if the respective quantities of the input data are known. Namely, the mean \mathbb{E}_u satisfies

$$(2) \quad \mathcal{L}\mathbb{E}_u = \mathbb{E}_f \text{ in } D \quad \text{and} \quad \mathbb{E}_u = 0 \text{ on } \Gamma$$

due to the linearity of the expectation and the differential operator \mathcal{L} . Taking into account the multilinearity of the tensor product, one verifies by tensorizing (1) that

$$(3) \quad \begin{aligned} (\mathcal{L} \otimes \mathcal{L}) \text{Cor}_u &= \text{Cor}_f && \text{in } D \times D, \\ (\mathcal{L} \otimes \text{Id}) \text{Cor}_u &= 0 && \text{on } D \times \Gamma, \\ (\text{Id} \otimes \mathcal{L}) \text{Cor}_u &= 0 && \text{on } \Gamma \times D, \\ \text{Cor}_u &= 0 && \text{on } \Gamma \times \Gamma. \end{aligned}$$

From Cor_u , we can compute the variance \mathbb{V}_u of the solution due to

$$\mathbb{V}_u(\mathbf{x}) = \text{Cor}_u(\mathbf{x}, \mathbf{x}) - \mathbb{E}_u(\mathbf{x})^2.$$

If a low-rank factorization of Cor_f is available, (3) can easily be solved by standard finite element methods. The existence of an accurate low-rank approximation is directly related to the spectral decomposition of the associated Hilbert–Schmidt operator

$$(4) \quad (\mathcal{C}_f \psi)(\mathbf{x}) := \int_D \text{Cor}_f(\mathbf{x}, \mathbf{y})\psi(\mathbf{y}) \, d\mathbf{y}.$$

Let $\text{Cor}_f \in H^p(D) \otimes H^p(D)$; then, according to [26, 53], the eigenvalues of this Hilbert–Schmidt operator decay like

$$(5) \quad \lambda_m \lesssim m^{-2p/d} \text{ as } m \rightarrow \infty.$$

Unfortunately, the constant in this estimate behaves similarly to the $H^p(D) \otimes L^2(D)$ -norm of Cor_f . The following consideration shows that this can lead to large constants in the decay estimate if the correlation length is small. Let the correlation kernel $k(r)$ depend only on the distance $r = \|\mathbf{x} - \mathbf{y}\|$. Then, the derivatives $\partial_{\mathbf{x}}^{\alpha} \text{Cor}_f(\mathbf{x}, \mathbf{y})$ and $\partial_{\mathbf{y}}^{\alpha} \text{Cor}_f(\mathbf{x}, \mathbf{y})$ of the correlation

$$\text{Cor}_f(\mathbf{x}, \mathbf{y}) = k\left(\frac{\|\mathbf{x} - \mathbf{y}\|}{\ell}\right)$$

involve the factor $\ell^{-|\alpha|}$, leading to a constant ℓ^{-p} in the decay estimate of the eigenvalues (5). Thus, for a small correlation length ℓ , a low-rank approximation of Cor_f becomes prohibitively expensive to compute.

Different approaches to tackling the solution of (3) have been considered in several papers; what most of these approaches have in common is that they are in some sense based on a sparse tensor product discretization of the solution. For example, the computation of the second moment, i.e., Cor_u , has been considered for elliptic diffusion problems with random loads in [52] by means of a sparse tensor product finite element method. A sparse tensor product wavelet boundary element method was used in [36] to compute the solution's second moment for elliptic potential problems on random domains. In [32, 35], the computation of the second moment was done by multilevel finite element frames. Recently, this concept was simplified by using the combination technique; cf. [34]. Unfortunately, the sparse tensor product discretization needs to resolve the concentrated measure for short correlation lengths. This means that the number of hierarchies of the involved finite element spaces has to be doubled if the correlation length is halved to get the same accuracy.

The present paper discusses a different approach to approximating the full tensor product discretization without losing its resolution properties. In [14], it was demonstrated that the \mathcal{H} -matrix technique is a powerful tool to cope with Dirichlet data of low Sobolev smoothness if the problem is formulated as a boundary value problem. There, the similar behavior of two-point correlation kernels and boundary integral operators was exploited. In [15], \mathcal{H} -matrix compressibility of the solution was proven also in case of local operators on domains. In the present paper, we will combine this theoretical foundation with the \mathcal{H} -matrix technique used in [14] to efficiently solve (3) by the finite element method for a right-hand side Cor_f with small correlation length or low Sobolev smoothness.

The general concept of \mathcal{H} -matrices and the corresponding arithmetic were first introduced in [27, 29]. \mathcal{H} -matrices are feasible for the data-sparse representation of (block-) matrices which can be approximated blockwise with low rank and have originally been employed for the efficient treatment of boundary integral equations as they arise in the boundary element method.

The rest of this paper is organized as follows. In section 2, we introduce the Galerkin discretization of the problem under consideration. Section 3 discusses the compressibility of discretized correlation kernels and the efficient solution of general correlation equations. In section 4, we recall some specialities of \mathcal{H} -matrices in the context of finite elements. In particular, we restate a phenomenon, called “weak admissibility,” which produces a more data-sparse representation of the correlation matrices, and \mathcal{H} -matrix nested dissection techniques. In section 5, we present numerical examples to validate and quantify the proposed method. Finally, in section 6, we draw our conclusions from the theoretical findings and the numerical results.

2. Preliminaries. For the remainder of this paper, let $D \subset \mathbb{R}^d$ be a Lipschitz domain, $(\Omega, \mathcal{F}, \mathbb{P})$ a separable, complete probability space, and \mathcal{L} the linear differential operator of second order given by

$$(6) \quad (\mathcal{L}u)(\mathbf{x}) := -\text{div}(\mathbf{A}(\mathbf{x}) \cdot \nabla u(\mathbf{x})) + c(\mathbf{x})u(\mathbf{x}).$$

The differential operator shall be strongly elliptic in the sense that

$$\underline{\alpha}\|\boldsymbol{\xi}\|_2^2 \leq \boldsymbol{\xi}^\top \mathbf{A}(\mathbf{x})\boldsymbol{\xi} \leq \bar{\alpha}\|\boldsymbol{\xi}\|_2^2 \quad \text{for all } \boldsymbol{\xi} \in \mathbb{R}^d$$

with coefficients $\mathbf{A} \in W^{1,\infty}(D, \mathbb{R}^{d \times d})$, $0 \leq c \in L^\infty(D, \mathbb{R})$, and $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$.

Under these assumptions, for a given load $f \in L^2_{\mathbb{P}}(\Omega, H^{-1}(D))$, the Dirichlet

problem

$$\begin{aligned} \mathcal{L}u(\omega, \mathbf{x}) &= f(\omega, \mathbf{x}) && \text{for } \mathbf{x} \in D, \\ u(\omega, \mathbf{x}) &= 0 && \text{for } \mathbf{x} \in \Gamma := \partial D \end{aligned}$$

is known to have a unique solution $u(\omega, \cdot) \in H_0^1(D)$ for \mathbb{P} -almost every $\omega \in \Omega$; cf., e.g., [21]. As a result, the mean $\mathbb{E}_u \in H_0^1(D)$ and the correlation $\text{Cor}_u \in H_0^1(D) \otimes H_0^1(D)$ are well defined.

For the efficient numerical solution of (3), we use a finite element Galerkin scheme. To that end, we introduce a finite element space $V_N = \text{span}\{\varphi_1, \dots, \varphi_N\} \subset H_0^1(D)$. It is assumed that the mesh which underlies this finite element space is quasi-uniform. The basis functions $\{\varphi_i\}_i$ are supposed to be locally and isotropically supported such that $\text{diam}(\text{supp } \varphi_i) \sim N^{-1/d}$. In particular, we can assign to each degree of freedom $i \in \{1, \dots, N\}$ a suitable point $\mathbf{x}_i \in D$, e.g., the barycenter of the support of the corresponding basis function or the corresponding Lagrangian interpolation point if nodal finite element shape functions are considered.

The variational formulation of (3) is given as follows:

Find $\text{Cor}_u \in H_0^1(D) \otimes H_0^1(D)$ such that

$$((\mathcal{L} \otimes \mathcal{L}) \text{Cor}_u, v)_{L^2(D \times D)} = (\text{Cor}_f, v)_{L^2(D \times D)} \text{ for all } v \in H_0^1(D) \otimes H_0^1(D).$$

By replacing the energy space $H_0^1(D) \otimes H_0^1(D)$ in this variational formulation by the finite dimensional ansatz space $V_N \otimes V_N$, we arrive at

Find $\text{Cor}_{u,N} \in V_N \otimes V_N$ such that

$$(7) \quad ((\mathcal{L} \otimes \mathcal{L}) \text{Cor}_{u,N}, v)_{L^2(D \times D)} = (\text{Cor}_f, v)_{L^2(D \times D)} \text{ for all } v \in V_N \otimes V_N.$$

A basis in $V_N \otimes V_N$ is formed by the set of tensor product basis functions $\{\varphi_i \otimes \varphi_j\}_{i,j}$. Hence, representing $\text{Cor}_{u,N}$ by its basis expansion yields

$$\text{Cor}_{u,N} = \sum_{\ell, \ell'=1}^N u_{\ell, \ell'} (\varphi_\ell \otimes \varphi_{\ell'}).$$

Setting $\mathbf{C}_u := [u_{\ell, \ell'}]_{\ell, \ell'}$, we end up with the system of linear equations

$$(8) \quad (\mathbf{A} \otimes \mathbf{A}) \text{vec}(\mathbf{C}_u) = \text{vec}(\mathbf{C}_f),$$

where $\mathbf{C}_f := [(\text{Cor}_f, \varphi_\ell \otimes \varphi_{\ell'})_{L^2(D \times D)}]_{\ell, \ell'}$ is the discretized two-point correlation of the Dirichlet data f and $\mathbf{A} := [(\mathcal{L}\varphi_{\ell'}, \varphi_\ell)_{L^2(D)}]_{\ell, \ell'}$ is the system matrix of the second order differential operator (6). In (8), the tensor product, as usual in connection with matrices, has to be understood as the Kronecker product. Furthermore, for a matrix $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{m \times n}$, the operation $\text{vec}(\mathbf{B})$ is defined as

$$\text{vec}([\mathbf{b}_1, \dots, \mathbf{b}_n]) := \begin{bmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_n \end{bmatrix} \in \mathbb{R}^{mn}.$$

For matrices $\mathbf{B} \in \mathbb{R}^{k \times n}$, $\mathbf{C} \in \mathbb{R}^{\ell \times m}$, and $\mathbf{X} \in \mathbb{R}^{m \times n}$, there holds the relation

$$(\mathbf{B} \otimes \mathbf{C}) \text{vec}(\mathbf{X}) = \text{vec}(\mathbf{C}\mathbf{X}\mathbf{B}^\top).$$

Hence, we may rewrite (8) according to

$$(9) \quad \mathbf{A}\mathbf{C}_u\mathbf{A}^\top = \mathbf{C}_f.$$

An approach to dealing with nonhomogeneous boundary conditions was presented in [32].

3. \mathcal{H} -matrix approximation of correlation kernels. The matrix equation (9) has N^2 unknowns and is therefore not directly solvable if N is large due to memory and time consumption. Thus, an efficient compression scheme and a powerful arithmetic are needed to obtain its solution. In the following, we will restrict ourselves to *asymptotically smooth* correlation kernels Cor_f , i.e., correlation kernels satisfying the following definition.

DEFINITION 3.1. Let $k: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. The function k is called asymptotically smooth if for some constants $c_1, c_2 > 0$ and $q \in \mathbb{R}$ there holds

$$(10) \quad \left| \partial_{\mathbf{x}}^\alpha \partial_{\mathbf{y}}^\beta k(\mathbf{x}, \mathbf{y}) \right| \leq c_1 \frac{(|\alpha| + |\beta|)!}{c_2^{|\alpha| + |\beta|}} \|\mathbf{x} - \mathbf{y}\|_2^{-d-2q-|\alpha|-|\beta|}, \quad \mathbf{x} \neq \mathbf{y},$$

independently of α and β .

Examples for asymptotically smooth correlation kernels are the Matérn kernels, which include, in particular, the Gaussian kernel; cf. [46, 50] and the references therein. A main feature of such asymptotically smooth correlation kernels is that they exhibit a data-sparse representation by means of \mathcal{H} -matrices; cf., e.g., [6, 8, 28].

\mathcal{H} -matrices rely on local low-rank approximations of a given matrix $\mathbf{X} \in \mathbb{R}^{N \times N}$. For suitable nonempty index sets $\nu, \nu' \subset \{1, \dots, N\}$, a matrix block $\mathbf{X}_{|\nu \times \nu'}$ can be approximated by a rank- k matrix. This approximation can be represented in factorized form $\mathbf{X}_{|\nu \times \nu'} \approx \mathbf{Y}\mathbf{Z}^\top$ with factors $\mathbf{Y} \in \mathbb{R}^{\nu \times k}$ and $\mathbf{Z} \in \mathbb{R}^{\nu' \times k}$. Hence, if $k \ll \min\{\#\nu, \#\nu'\}$, the complexity for storing the block is considerably reduced. The construction of the index sets is based on the *cluster tree*.

3.1. Cluster tree. For a tree $\mathcal{T} = (V, E)$ with vertices V and edges E , we define its set of leaves by

$$\mathcal{L}(\mathcal{T}) := \{\sigma \in V : \sigma \text{ has no sons}\}.$$

Furthermore, we say that \mathcal{T} is a *cluster tree* for the set $\{1, \dots, N\}$ if the following conditions hold.

- $\{1, \dots, N\}$ is the *root* of \mathcal{T} .
- All $\sigma \in V \setminus \mathcal{L}(\mathcal{T})$ are the disjoint union of their sons.

The *level* of $\sigma \in \mathcal{T}$ is its distance to the root, i.e., the number of son relations that are required for traveling from $\{1, \dots, N\}$ to σ . We define the set of clusters on level j as

$$\mathcal{T}^{(j)} := \{\sigma \in \mathcal{T} : \sigma \text{ has level } j\}.$$

The construction of the cluster tree is based on the *support* of the clusters. The support Υ_σ of a cluster σ is defined as the union of the supports of the basis functions corresponding to their elements, that is,

$$\Upsilon_\sigma = \bigcup_{i \in \sigma} \Upsilon_i, \quad \text{where } \Upsilon_i := \text{supp } \varphi_i \text{ for all } i \in \{1, \dots, N\}.$$

For computing complexity bounds, the cluster tree should match the following additional requirements, uniformly as $N \rightarrow \infty$:

- The cluster tree is a balanced tree in the sense that the maximal level satisfies $J \sim \log_2 N$.
- The diameter of the support Υ_{σ_j} , $\sigma_j \in \mathcal{T}^{(j)}$, is local with respect to the level j , i.e., $\text{diam } \Upsilon_{\sigma_j} \sim 2^{-j/d}$. Moreover, the number $\#\sigma_j$ of indices contained in a cluster $\sigma_j \in \mathcal{T}^{(j)}$ scales approximately like 2^{J-j} , i.e., $\#\sigma_j \sim 2^{J-j}$.

Until further notice, a binary cluster tree \mathcal{T} with the indicated terms should be given for our further considerations. A common algorithm for its construction is based on a hierarchical subdivision of the point set which is associated with the basis functions; cf., e.g., [6, 8, 28]. We begin by embedding the point set $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ in a top-level bounding-box. This bounding-box is subsequently subdivided into two cuboids of the same size where the corresponding clusters are described by the points in each bounding-box. This process is iterated until a bounding-box encloses less than a predetermined number of points.

3.2. H-matrix approximation. H-matrices were originally invented in [27, 29] and are a generalization of cluster techniques for the rapid solution of boundary integral equations such as the fast multipole method [25], the mosaic skeleton approximation [56], and the adaptive cross approximation [3].

For the discretization of an asymptotically smooth correlation, we introduce a partition of its domain of definition which separates smooth and nonsmooth areas of the kernel function. It is based on the following.

DEFINITION 3.2. *Two clusters σ and σ' are called η -admissible if*

$$(11) \quad \max\{\text{diam}(\Upsilon_\sigma), \text{diam}(\Upsilon_{\sigma'})\} \leq \eta \text{ dist}(\Upsilon_\sigma, \Upsilon_{\sigma'})$$

holds for some fixed $\eta > 0$.

We can obtain the set of admissible blocks by means of a recursive algorithm: Starting with the root $(\sigma_{0,0}, \sigma_{0,0})$, the bounding-boxes of the current cluster pair are checked for admissibility. If they are admissible, the cluster pair is added to the set \mathcal{F} which corresponds to the correlation kernel's *farfield*. Otherwise, the admissibility check will be performed on all bounding-boxes of the possible pairs of son clusters of the two original clusters. When we arrive at a pair of leaf clusters with inadmissible bounding-boxes, the clusters are added to the set \mathcal{N} which corresponds to the correlation kernel's *nearfield*. The set $\mathcal{B} = \mathcal{F} \cup \mathcal{N}$ obviously inherits a tree structure from the recursive construction of \mathcal{F} and \mathcal{N} and is called the *block cluster tree*; see [6, 8, 28].

With the definition of the block cluster tree at hand, we are finally in the position to introduce H-matrices.

DEFINITION 3.3. *The set $\mathcal{H}(\mathcal{B}, k)$ of H-matrices of maximal block rank k is defined according to*

$$\mathcal{H}(\mathcal{B}, k) := \{\mathbf{X} \in \mathbb{R}^{N \times N} : \text{rank}(\mathbf{X}|_{\sigma \times \sigma'}) \leq k \text{ for all } (\sigma, \sigma') \in \mathcal{F}\}.$$

Note that all nearfield blocks $\mathbf{X}|_{\sigma \times \sigma'}$, $(\sigma, \sigma') \in \mathcal{N}$, are allowed to be full matrices.

In accordance with [6, 8, 28], the storage cost of an H-matrix $\mathbf{X} \in \mathcal{H}(\mathcal{B}, k)$ is $\mathcal{O}(kN \log N)$. Here, for asymptotically smooth correlation kernels, the rank k depends polylogarithmically on the desired approximation accuracy ε , which in turn usually depends on the degrees of freedom N . These remarks pertain to the approximation of an *explicitly* given, asymptotically smooth correlation kernel k , such as Cor_f in (3).

The compressibility of an *implicitly* given correlation kernel, such as Cor_u in (3), was studied in [15] for the case of smooth domains D . We restate the main theorem

for the setting of the present paper which employs that the Hilbert–Schmidt operator (4), related to the correlation kernel Cor_f , is in general a *pseudodifferential operator*; see, e.g., [38, 39, 40, 55] and the references therein.

THEOREM 3.4. *In the domain D with analytic boundary ∂D , assume that the correlation kernel Cor_f in (3) gives rise to an operator $\mathcal{C}_f \in OPS_{cl,s}^\theta$, i.e., to a classical pseudodifferential operator with symbol $a_f(\mathbf{x}, \boldsymbol{\xi})$ of order θ and of Gevrey class $s \geq 1$ in the sense of [10, Def. 1.1]. Assume further that the coefficients of the differential operator \mathcal{L} are smooth. Then, the correlation kernel Cor_u of (3) is the Schwartz kernel of an operator $\mathcal{C}_u \in OPS_{cl,s}^{\theta-4}$.*

Moreover, the kernel $\text{Cor}_u(\mathbf{x}, \mathbf{y})$ of the correlation operator \mathcal{C}_u is smooth in $D \times D$ outside of the diagonal $\Delta := \{(\mathbf{x}, \mathbf{y}) \in D \times D : \mathbf{x} = \mathbf{y}\}$, and there holds the pointwise estimate

$$(12) \quad |\partial_{\mathbf{x}}^\alpha \partial_{\mathbf{y}}^\beta \text{Cor}_u(\mathbf{x}, \mathbf{y})| \leq c \mathcal{A}^{|\alpha+\beta|} (|\boldsymbol{\alpha}|!)^s |\boldsymbol{\beta}|! \|\mathbf{x} - \mathbf{y}\|_2^{-\theta-d-|\alpha|-|\beta|+4}$$

for all $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{N}_0^d$, $(\mathbf{x}, \mathbf{y}) \in (D \times D) \setminus \Delta$ with some constants c and \mathcal{A} which depend only on D and a_f .

Obviously, for $s = 1$, estimate (12) directly implies condition (10) for the asymptotic smoothness of Cor_u , allowing us to approximate Cor_u by means of \mathcal{H} -matrices. In particular, [15] also provides some numerical evidence that this result could likely be extended to Lipschitz domains.

An example of correlation kernels for Cor_f satisfying the condition of this theorem for $s \geq 1$ is the Matérn class of kernels. We refer the reader to [15] for more details on how to verify the assumptions of the theorem for other correlation kernels.

3.3. \mathcal{H} -matrix arithmetic and iterative solution. An important feature of \mathcal{H} -matrices is that efficient algorithms for approximate matrix arithmetic operations are available. For two \mathcal{H} -matrices $\mathbf{H}_1, \mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$, the approximate *matrix-matrix addition* $\mathbf{H}_1 + \mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$ can be performed in $\mathcal{O}(k^2 N \log N)$ operations, while the approximate *matrix-matrix multiplication* $\mathbf{H}_1 * \mathbf{H}_2 \in \mathcal{H}(\mathcal{B}, k)$ can be performed in $\mathcal{O}(k^2 N \log^2 N)$ operations. Both of these operations are essentially block-matrix algorithms with successive recompression schemes. Moreover, employing the recursive block structure, the approximate inversion or the approximate computation of the LU-decomposition within $\mathcal{H}(\mathcal{B}, k)$ can also be performed in only $\mathcal{O}(k^2 N \log^2 N)$ operations. We refer the reader to [6, 8, 23, 28, 29] for further results and implementation details. In particular, the parallelization of the \mathcal{H} -matrix arithmetic and the \mathcal{H} -LU-decomposition was discussed in [43, 44].

In the context of correlation equations, this approximate \mathcal{H} -matrix arithmetic was successfully used in [14] to solve a problem, similar to (3), which has been discretized by the boundary element method. Then, the matrix \mathbf{A} in (9) corresponds to the stiffness matrix from the boundary element method which can naturally be approximated by the \mathcal{H} -matrix technique. The resulting matrix equation was solved using an iterative solver based on *iterative refinement* (cf. [22, 48, 58]), which we are also going to employ here. This method was originally introduced in [58] for the improvement of solutions to systems of linear equations computed by using an LU-factorization.

Having all matrices in (9) represented by \mathcal{H} -matrices, the solution can then be approximated as follows. Let $\mathbf{A} \approx \hat{\mathbf{L}} \hat{\mathbf{U}}$, where $\hat{\mathbf{L}}, \hat{\mathbf{U}} \in \mathcal{H}(\mathcal{B}, k)$ is an approximate LU-decomposition to \mathbf{A} , e.g., computed from \mathbf{A} by the \mathcal{H} -matrix arithmetic. Starting with the initial guess $\mathbf{C}_u^{(0)} = \hat{\mathbf{U}}^{-1} \hat{\mathbf{L}}^{-1} \mathbf{C}_f \hat{\mathbf{L}}^{-\top} \hat{\mathbf{U}}^{-\top}$, we iterate

$$(13) \quad \boldsymbol{\Theta}^{(i)} = \mathbf{C}_f - \mathbf{A} \mathbf{C}_u^{(i)} \mathbf{A}^\top, \quad \mathbf{C}_u^{(i+1)} = \mathbf{C}_u^{(i)} + \hat{\mathbf{U}}^{-1} \hat{\mathbf{L}}^{-1} \boldsymbol{\Theta}^{(i)} \hat{\mathbf{L}}^{-\top} \hat{\mathbf{U}}^{-\top}, \quad i = 0, 1, \dots$$

Note that we use, in contrast to [14], the LU-decomposition with forward and backward substitution algorithms, which avoids the expensive computation of an approximate inverse. Whether we use an approximate inverse as in [14] or an approximate LU-decomposition, the idea of iterative refinement stays the same: The residual $\Theta^{(i)}$ is computed with a higher precision than the correction $\hat{\mathbf{U}}^{-1}\hat{\mathbf{L}}^{-1}\Theta^{(i)}\hat{\mathbf{L}}^{-\top}\hat{\mathbf{U}}^{-\top}$. This yields an improved approximation to the solution in each step. Note that this algorithm also algebraically coincides with an undamped preconditioned Richardson iteration; see, e.g., [51].

In the following, we will elaborate on how this approach can be realized in the context of the finite element method. If \mathbf{A} is symmetric and positive definite, the LU-decomposition could also be replaced by a Cholesky decomposition. Nonetheless, we will see in the numerical experiments that the computation time of the decomposition is negligible compared to the overall computation time, and we prefer to stay in the more general, i.e., nonsymmetric, setting.

4. \mathcal{H} -matrices in the context of finite elements. Although a finite element matrix has a sparse structure, its inverse and both factors of its LU-decomposition are generally fully populated. Nevertheless, the inverse and the LU-decomposition exhibit a data-sparse structure in the sense that they are \mathcal{H} -matrix compressible. We recall the main concepts from the literature; see, e.g., [4, 7, 16, 17, 28].

4.1. General concepts. A rough argument for the \mathcal{H} -matrix compressibility of the inverse makes use of the Green's function G of \mathcal{L} . Let $\delta_{\mathbf{x}}$ denote the Dirac distribution at the point \mathbf{x} , and let $G: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy

$$\mathcal{L}_{\mathbf{y}}G(\mathbf{x}, \mathbf{y}) = \delta_{\mathbf{x}} \text{ and } G(\cdot, \mathbf{y})|_{\Gamma} = 0.$$

Then, the solution of

$$\begin{aligned} \mathcal{L}u(\mathbf{x}) &= f(\mathbf{x}) & \text{for } \mathbf{x} \in D, \\ u(\mathbf{x}) &= 0 & \text{for } \mathbf{x} \in \Gamma \end{aligned}$$

can be represented by

$$u(\mathbf{x}) = (\mathcal{L}^{-1}f)(\mathbf{x}) = \int_D G(\mathbf{x}, \mathbf{y})f(\mathbf{y}) \, d\mathbf{y}, \quad \mathbf{x} \in D.$$

If the Green's function is analytic away from the diagonal, e.g., in the case of constant coefficients of \mathcal{L} , we can approximate the Green's function away from the diagonal by local expansions of the kind

$$G(\mathbf{x}, \mathbf{y}) \approx \sum_{i=1}^k a(\mathbf{x})b(\mathbf{y}),$$

which is the theoretical basis for an \mathcal{H} -matrix approximation; see [6, 8, 28].

However, one of the advantages of the finite element method is that it can also be applied in the case of nonconstant coefficients. In [7], a proof was presented to guarantee the existence of an \mathcal{H} -matrix approximation to the inverse of the finite element stiffness matrix, even in the case of essentially bounded diffusion coefficients and the other coefficients set to zero. This result was then extended in [4] to allow all coefficients to be only essentially bounded, providing the theoretical foundation for an \mathcal{H} -matrix approximation to the inverse of the differential operator from (6). Having the \mathcal{H} -matrix approximability of the inverse to the finite element matrix available,

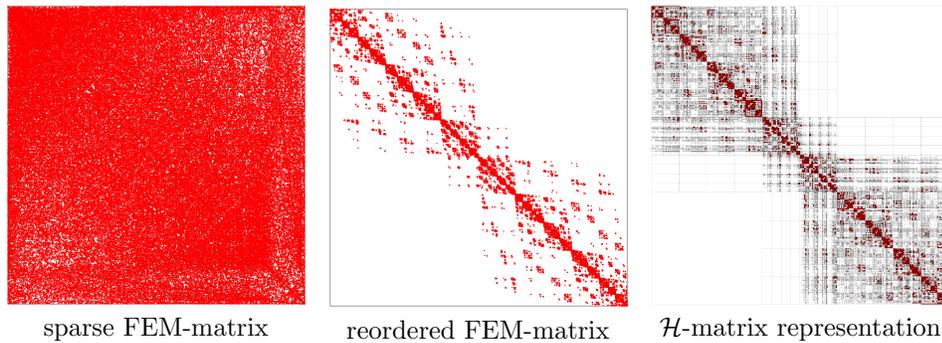


FIG. 1. Sparsity pattern of a 3D finite element matrix, its reordered finite element matrix, and the corresponding \mathcal{H} -matrix. Red blocks in the \mathcal{H} -matrix correspond to the nearfield, and white blocks correspond to the empty farfield.

the approximability of the LU-decomposition to the finite element matrix was then proven in [5].

While these first results hold up to the finite element discretization error, the results have recently been improved in [16, 17] to hold without additional error.

It remains to explain how to actually compute an \mathcal{H} -matrix approximation to the inverse or the LU-decomposition of a finite element stiffness matrix. To that end, note that a necessary condition for an entry \mathbf{A}_{ij} in the finite element matrix to be nonzero is that $\Upsilon_i \cap \Upsilon_j \neq \emptyset$; i.e., the intersection of the corresponding supports of the basis functions is nonempty. This yields together with the η -admissibility condition (11) that all entries of a finite element matrix have η -inadmissible supports; i.e., they are contained in the nearfield of an \mathcal{H} -matrix. A sparse finite element matrix can therefore be represented as an \mathcal{H} -matrix by reordering the index set corresponding to the clustering scheme introduced in section 3.1 and inserting the nonzero entries into the nearfield. An illustration of this procedure can be found in Figure 1.

Having the finite element matrix represented by an \mathcal{H} -matrix, the approximate inverse and the LU-decomposition can be computed by using the block algorithms of the \mathcal{H} -matrix arithmetic in $\mathcal{O}(k^2 N \log^2 N)$ operations; cf. [6, 8, 28]. We note especially that the computation of the LU-decomposition together with its forward and backward substitution algorithms still has an overall complexity of $\mathcal{O}(k^2 N \log^2 N)$, but with smaller constants than the computation and the application of an approximate inverse.

4.2. Weak admissibility. Approximate \mathcal{H} -matrix representations for the inverse or LU-factorizations of finite element matrices have been used to construct preconditioners for iterative solvers; see, e.g., [6] and the references therein. In [30], it was observed for the one-dimensional case that the computation of an approximate inverse can be considerably sped up by replacing the η -admissibility condition (11) by the following *weak admissibility condition*.

DEFINITION 4.1. *Two clusters σ and σ' are called weakly admissible if $\sigma \neq \sigma'$.*

We observe immediately that an η -admissible block cluster is also weakly admissible. Thus, by replacing the η -admissibility condition by weak admissibility, we obtain a much coarser partition of the \mathcal{H} -matrix. This leads to smaller constants in the storage and computational complexity; cf. [30]. Each row and each column of the finite element matrix has only $\mathcal{O}(1)$ entries. Thus, inserting the finite element matrix into a weakly admissible \mathcal{H} -matrix structure, the off-diagonal blocks of the \mathcal{H} -matrix

have low rank.

By partitioning the matrix according to the weak admissibility condition, we cannot ensure the exponential convergence of fast black box low-rank approximation techniques as used for boundary element matrices. For example, adaptive cross approximation relies on an admissibility condition similar to (11) to ensure exponential convergence; cf. [3]. Instead, the authors of [30] suggest assembling a weakly admissible matrix block according to the η -admissibility condition and transforming it on-the-fly to a low-rank matrix to obtain a good approximation.

The behavior of the ranks of the low-rank matrices in weakly admissible partitions compared to η -admissible partitions is not fully understood yet. Suppose that k_η is an upper bound for the ranks corresponding to an η -admissible partition, and suppose that k_w shall be an upper bound for the ranks to a weakly admissible partition. In [30], it is proven for one-dimensional finite element discretizations that one should generally choose

$$k_w = Lk_\eta$$

in order to obtain the same approximation accuracy in the weakly admissible case as in the η -admissible case. Here, L is a constant which depends on the depth of the block cluster tree and thus logarithmically on N . In the same paper, the authors remark in the numerical examples that this bound on k_w seems to be too pessimistic and one could possibly choose

$$(14) \quad k_w = c_{\eta \rightarrow w} k_\eta,$$

where $c_{\eta \rightarrow w} \in [2, 3.5]$.

Unfortunately, weak admissibility is not suitable for dimensions greater than one due to the fact that clusters can possibly intersect each other in $\mathcal{O}(N^\alpha)$ points, where $\alpha \geq 0$ depends on the spatial dimension. However, one can try to reduce this negative influence of the weak admissibility condition by mixing it with η -admissibility. In the software package HLib (cf. [9]), the authors use η -admissibility for all block clusters with block size larger than a given threshold and apply the weak admissibility condition for block clusters $\sigma \times \sigma'$ which are below that threshold provided that the condition

$$a_i^\sigma < \frac{a_i^{\sigma'} + b_i^{\sigma'}}{2} < b_i^\sigma \quad \text{or} \quad a_i^{\sigma'} < \frac{a_i^\sigma + b_i^\sigma}{2} < b_i^{\sigma'}$$

is satisfied for the corresponding bounding-boxes $\prod_{i=1}^3 [a_i^\mu, b_i^\mu]$, $\mu = \sigma, \sigma'$, in at most one coordinate direction. This condition restricts the application of weak admissibility to essentially one-dimensional cluster intersections with length below a certain threshold. The impact of this specific admissibility condition is illustrated in Figure 2.

4.3. Nested dissection. While weak admissibility takes the sparsity of the finite element matrix into account only during the construction of the \mathcal{H} -matrix, it is also possible to incorporate the sparsity during the construction of the cluster tree. One possibility was introduced in [24] and is based on nested dissection; cf. [11, 19, 37, 45] and the references therein. We briefly review the idea of nested dissection in the context of \mathcal{H} -matrices as discussed in [24] and refer the reader to [24] for more details.

The idea is to employ a recursive algorithm as follows.

1. Split degrees of freedom into three disjoint subsets $\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3$ according to the following conditions.
 - \mathcal{I}_1 and \mathcal{I}_2 should have comparable sizes.

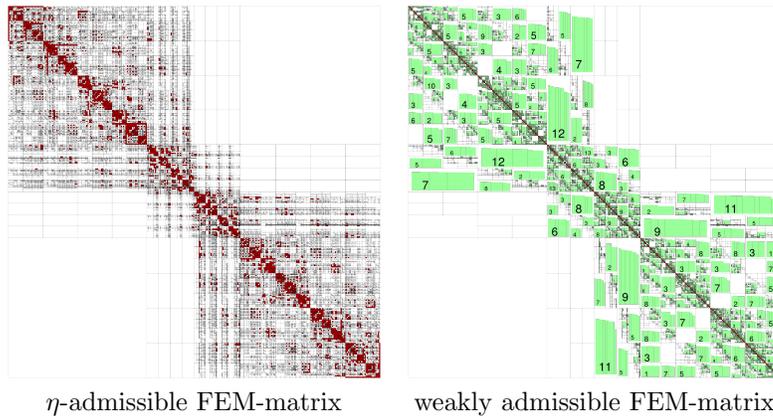


FIG. 2. Comparison of the partition for η -admissibility and for weak admissibility. Red blocks correspond to full matrices, green blocks correspond to low-rank matrices with inscribed rank, and white blocks are zero.

- \mathcal{I}_1 and \mathcal{I}_2 should not interact with each other; i.e., all entries $\mathbf{A}_{i,j}$, $i \in \mathcal{I}_1$, $j \in \mathcal{I}_2$ of the finite element matrix are zero.
 - \mathcal{I}_3 is the boundary layer between \mathcal{I}_1 and \mathcal{I}_2 .
2. Relabel indices in subsequent order: first \mathcal{I}_1 , then \mathcal{I}_2 , and then \mathcal{I}_3 .
 3. Proceed recursively with \mathcal{I}_1 and \mathcal{I}_2 (in the \mathcal{H} -matrix framework, \mathcal{I}_3 will also receive a recursive treatment).

Reordering the index sets of the finite element matrix in accordance with this procedure yields a sparsity pattern as illustrated in Figure 3. Due to the special construction, large parts of the matrix are zero and will remain zero in a subsequent LU-decomposition.

In the following, we take the approach of [24] to construct an \mathcal{H} -matrix which reorders the index set such that the pattern of the finite element matrix exposes a nested dissection ordering. We therefore recapitulate the construction of a cluster tree based on domain decomposition as proposed in [24]. For that purpose, the cluster algorithm distinguishes between *domain clusters* and *interface clusters*. The following algorithm is employed for the root $\{1, \dots, N\}$ and all domain clusters.

1. Given a cluster σ and its corresponding set of points $\{\mathbf{x}_1, \dots, \mathbf{x}_{\#\sigma}\}$, construct an axis-parallel bounding box Q_σ .
2. Cut Q_σ into two pieces Q_1 and Q_2 by halving the longest edge.
3. Define three disjoint sons of σ as follows:
 - $\sigma_1 = \{i: \mathbf{x}_i \in Q_1\}$,
 - $\sigma_2 = \{i: \Upsilon_i \cap \Upsilon_{\sigma_1} = \emptyset\}$,
 - $\sigma_3 = \sigma \setminus \{\sigma_1 \cup \sigma_2\}$.
4. Relabel indices of clusters in subsequent order: first σ_1 , then σ_2 , and then σ_3 .
5. Treat σ_1 and σ_2 as domain clusters and σ_3 as an interface cluster as indicated below.

Due to their special construction, the bounding-boxes of interface clusters are “flat” in one coordinate direction. The cluster algorithm therefore has to be adapted to fit the asymptotic requirements of section 3.1. Therefore, we define $\text{level}_{\text{int}}(\sigma)$ as the distance of σ to the nearest domain cluster in the cluster tree. We then employ the following cluster algorithm for the interface clusters and refer the reader to [24] for a detailed discussion.

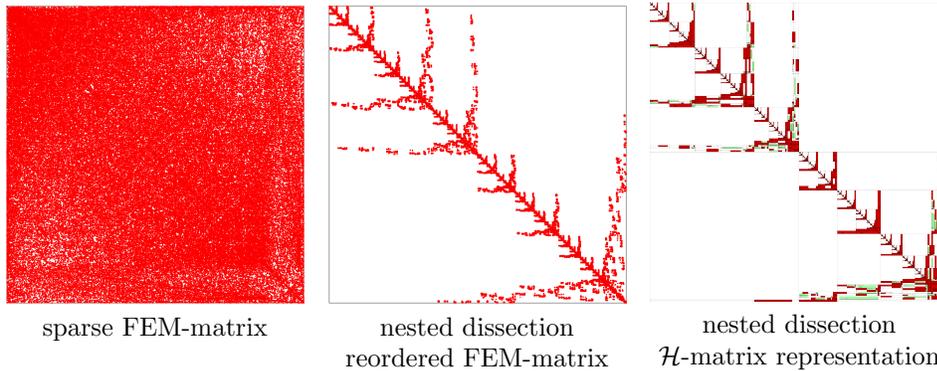


FIG. 3. Sparsity pattern of a 3D finite element matrix, its reordered finite element matrix according to nested dissection, and its corresponding \mathcal{H} -matrix.

1. Distinguish between two cases:

- If $\text{level}_{\text{int}}(\sigma) \equiv 0 \pmod{d}$, do not subdivide σ , and set $\sigma' = \sigma$ as its only son.
- If $\text{level}_{\text{int}}(\sigma) \not\equiv 0 \pmod{d}$, split the bounding box Q_σ axis parallel into two boxes Q_1 and Q_2 such that the “flat” direction is not modified. Then define the two son clusters $\sigma_1 = \{i: \mathbf{x}_i \in Q_1\}$ and $\sigma_2 = \{i: \mathbf{x}_i \in Q_2\}$.

2. Apply the cluster algorithm for interface clusters to all sons.

In order to translate the sparsity of the finite element matrix into the block structure of an \mathcal{H} -matrix, we can combine the η -admissibility condition from (11) and the weak admissibility condition from Definition 4.1 to a *nested dissection admissibility* condition.

DEFINITION 4.2. Two clusters σ and σ' are called nd-admissible if either

- $\sigma \neq \sigma'$ are both domain clusters or
- σ and σ' are η -admissible.

In fact, if $\sigma \neq \sigma'$ are both domain clusters, we can directly say that the corresponding \mathcal{H} -matrix block has rank zero. Figure 3 illustrates the sparsity pattern of the finite element matrix after the permutations determined by the cluster algorithms and how large parts of the corresponding \mathcal{H} -matrix have rank zero. The low-rank blocks in the representation are due to some internal checks of HLib which aim at replacing inadmissible blocks by low-rank matrices only if very few entries in the corresponding matrix block are nonzero.

The numerical experiments in the next section show that the sparse structure constructed here leads to smaller constants in the complexity of the solution algorithm.

5. Numerical results. Before we summarize the settings of the numerical experiments, we briefly recall that the algorithm of the presented method consists of the following three steps.

1. Compute the sparse finite element matrix \mathbf{A} in linear and the correlation \mathcal{H} -matrix \mathbf{C}_f in almost linear time.
2. Compute the approximate LU-decomposition of \mathbf{A} in \mathcal{H} -matrix format in almost linear time.
3. Solve the matrix equation (9) with iterative refinement (13) in almost linear time for each iteration.

The numerical experiments in this paper shall mainly focus on the third step and the

overall behavior of the method. We will see that only one iteration is required in the third step, which yields an almost linear overall complexity of the algorithm. To improve the computation time, we store and compute only the lower part of \mathbf{C}_f and \mathbf{C}_u .

All the computations in the following experiments have been carried out on a single core of a computing server with two Intel(R) Xeon(R) E5-2670 CPUs with a clock rate of 2.60GHz and a main memory of 256GB. For the \mathcal{H} -matrix computations, we use the software package HLib (cf. [9]), and for the finite element discretization we use the Partial Differential Equation Toolbox of MATLAB¹ which employs piecewise linear finite elements. The two libraries are coupled together in a single C-program (cf. [41]) using the MATLAB Engine interface. The meshes are generated by TETGEN (cf. [54]) and then imported into MATLAB.

5.1. Experimental setup. To obtain computational efficiency and to keep the ranks of the low-rank matrices under control, HLib imposes an upper threshold k_η for the ranks in the case of an η -admissible \mathcal{H} -matrix and a lower threshold n_{\min} for the minimal block size. For the application of the weak admissibility condition, we rely on the criterion of HLib, which considers the weak admissibility condition only if one of the index sets of the block cluster $\sigma \times \sigma'$ has a cardinality below 1,024, and the condition

$$a_i^\sigma < \frac{a_i^{\sigma'} + b_i^{\sigma'}}{2} < b_i^\sigma \quad \text{or} \quad a_i^{\sigma'} < \frac{a_i^\sigma + b_i^\sigma}{2} < b_i^{\sigma'}$$

is satisfied for the corresponding bounding boxes $\prod_{i=1}^3 [a_i^\mu, b_i^\mu]$, $\mu = \sigma, \sigma'$, in at most one coordinate direction. Otherwise, η -admissibility is used instead. In the case of a weakly admissible matrix block, HLib imposes an upper threshold of $k_w = 3k_\eta$, setting $c_{\eta \rightarrow w} = 3$ in (14). For our experiments, we choose $\eta = 2$, $k_\eta = 20$, and $n_{\min} = 50$ and employ either a geometric cluster strategy, i.e., the binary cluster strategy from the end of section 3.1, or the nested dissection cluster strategy as discussed in section 4.3. The iterative refinement is stopped if the absolute error of the residual in the Frobenius norm is smaller than 10^{-6} .

In the following examples, we want to study, in addition to other aspects, the influence of the weak admissibility condition and nested dissection clustering for the partitioning of the different \mathcal{H} -matrices. Namely, we successively want to replace η -admissibility by weak admissibility for a binary and a nested dissection cluster tree as described in Table 1 in order to lower the constants hidden in the complexity of the \mathcal{H} -matrix arithmetic and thus to improve the computation time. For the discretization of the correlation kernel Cor_f , we will always use adaptive cross approximation.

TABLE 1
The five combinations of admissibility conditions used for the partition of \mathcal{H} -matrices.

Case	Operator and admissibility		Cluster tree
	\mathcal{L} and \mathcal{L}^{-1}	Cor_f and Cor_u	
all- η	η -admissibility	η -admissibility	binary
weak-FEM	weak admissibility	η -admissibility	binary
all-weak	weak admissibility	weak admissibility	binary
nd- η	nd-admissibility	η -admissibility	nested dissection
nd-weak	nd-admissibility	weak-admissibility	nested dissection

The *all- η* case is the canonical case and was also investigated in the case of

¹Release 2015b, The MathWorks, Inc., Natick, MA.

the boundary element method in [14]. The *weak-FEM* case is a first relaxation to apply the weak admissibility condition. This is justified, since the stiffness matrix \mathbf{A} can exactly be represented as a weakly admissible \mathcal{H} -matrix and the iterative refinement only involves an approximate LU-decomposition. Hence, we expect at most an influence on the quality of the approximate LU-decomposition and thus on the number of iterations in the iterative refinement. We therefore have to investigate if possible additional iterations are compensated by the faster \mathcal{H} -matrix arithmetic.

The aforementioned cases have in common that they rely on the asymptotic smoothness of Cor_f and Cor_u and the η -admissibility which leads to exponential convergence of the \mathcal{H} -matrix approximation. In the case *all-weak*, we want to examine if there is some indication that weak admissibility could possibly also be considered for the partition of \mathcal{H} -matrices for Cor_f and Cor_u . To that end, we approximate Cor_f with adaptive cross approximation relative to the η -admissibility partition and convert it on-the-fly to the partition of the weak admissibility, as proposed in [30].

While the three aforementioned cases all rely on a binary cluster tree, the cases *nd- η* and *nd-weak* rely on a cluster tree which is constructed by nested dissection. In both cases, the finite element matrix \mathbf{A} is partitioned by nd-admissibility. For Cor_f and Cor_u , we use η -admissibility in the nd- η case and weak admissibility in the nd-weak case, where we assemble the matrix for Cor_f in the same way as in the all-weak case.

The following numerical examples are divided into two parts. In the first part, we demonstrate the convergence of the presented method by comparing it to a low-rank reference solution computed with the pivoted Cholesky decomposition; cf. [33]. In the second part, we will demonstrate that the presented method also works well in the case of correlation kernels with low Sobolev smoothness or small correlation length, where no low-rank approximations exist and sparse tensor product approximations fail to resolve the correlation length. Note that in both examples, due to the nonlocality of the correlation kernels and the Green's function, the computed system matrices are smaller than usual for the finite element method. In particular, the unknown in the system (9) of linear equations is a matrix with N^2 entries, whereas the corresponding mesh has only N degrees of freedom. The \mathcal{H} -matrix compression reduces the computational complexity for the assembly and the amount of required storage from N^2 to $\mathcal{O}(kN \log N)$, whereas the complexity of the solution algorithm decreases from $\mathcal{O}(N^3)$ to $\mathcal{O}(k^2 N \log^2 N)$.

5.2. Tests for the iterative solver. Due to the recompression schemes in the block matrix algorithms of the \mathcal{H} -matrix arithmetic, it is not immediately clear if the presented solver converges. Still, it can be shown that some iterative \mathcal{H} -matrix schemes converge up to a certain accuracy; cf. [31]. In the following, we want to demonstrate for a specific example that our iterative scheme provides indeed convergence.

On the dumbbell geometry pictured in Figure 4, we consider $\mathcal{L} = -\Delta$ in (3) and the Matérn-5/2 kernel as input correlation Cor_f , i.e., for $r = \|\mathbf{x} - \mathbf{y}\|_2$, we set

$$\text{Cor}_f(\mathbf{x}, \mathbf{y}) = \left(1 + \sqrt{5} \frac{r}{\ell} + \frac{5}{3} \frac{r^2}{\ell^2} \right) \exp \left(-\sqrt{5} \frac{r}{\ell} \right),$$

where $\ell \approx \text{diam}(D)$ denotes the correlation length. The conversion of the finite element matrix to an \mathcal{H} -matrix for the dumbbell geometry has already been illustrated in Figure 1. The difference between η -admissibility and weak admissibility is illustrated in Figure 2, and the effect of nested dissection ordering is illustrated in Figure 3.

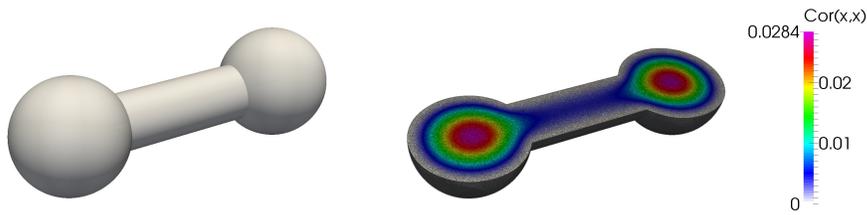


FIG. 4. The dumbbell geometry and its meshed cross section with the trace of the solution correlation kernel $\text{Cor}_u|_{\mathbf{x}=\mathbf{y}}$ for load data prescribed by the Matérn-5/2 kernel.

TABLE 2

Mesh points and degrees of freedom of the finite element mesh N and number of entries N^2 in the solution matrix for different levels of the dumbbell geometry.

Level	1	2	3	4	5	Reference mesh
Mesh points	238	1,498	6,958	34,112	175,562	1,033,382
N	4	201	1,742	13,341	98,177	756,626
N^2	16	40,401	3,034,564	$1.78 \cdot 10^8$	$9.64 \cdot 10^9$	$5.72 \cdot 10^{11}$

For determining a reference solution, we compute a low-rank approximation $\mathbf{C}_f \approx \mathbf{L}_f \mathbf{L}_f^T$ with the pivoted Cholesky decomposition as proposed in [33]. The numerical solution \mathbf{C}_u of (9) is then given by

$$\mathbf{C}_u \approx \mathbf{L}_u \mathbf{L}_u^T,$$

where \mathbf{L}_u solves $\mathbf{A} \mathbf{L}_u = \mathbf{L}_f$. To compute the error of the \mathcal{H} -matrix solution, we compare the correlation and the correlation's trace $\text{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ of the \mathcal{H} -matrix approximation with the correlation and the correlation's trace $\text{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ derived from the pivoted Cholesky decomposition on a finer reference mesh. We refer the reader to Table 2 for more details on the meshes under consideration.

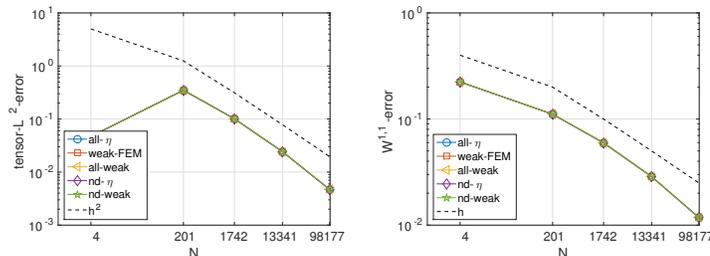


FIG. 5. L^2 -error of $\text{Cor}_{u,N}$ (left) and $W^{1,1}$ -error of $\text{Cor}_{u,N}|_{\mathbf{x}=\mathbf{y}}$ (right).

While the error of the correlation itself can be measured in the L^2 -norm on the tensor product domain, the appropriate norm for error measurements of its trace is the $W^{1,1}$ -norm. Due to the Poincaré–Friedrich and the Cauchy–Schwarz inequalities and

$$\|u^2 - u_h^2\|_{W^{1,1}(D)} \lesssim \|\nabla(u^2 - u_h^2)\|_{L^1(D)} \lesssim |u - u_h|_{H^1(D)} + \|u - u_h\|_{L^2(D)} \lesssim h,$$

we can expect a convergence rate in the $W^{1,1}$ -norm which is proportional to the mesh size h . A standard tensor product argument yields a convergence rate of order h^2 in

the tensor product L^2 -norm. Figure 5 shows that we indeed reach these rates for all five cases of admissibility which are considered in Table 1. In fact, the observed errors coincide in the first few digits.

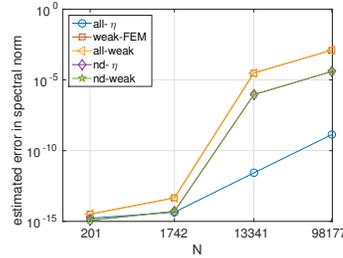


FIG. 6. The deviation $\hat{\mathbf{L}}\hat{\mathbf{U}} - \mathbf{A}$ in the estimated spectral norm for fixed rank.

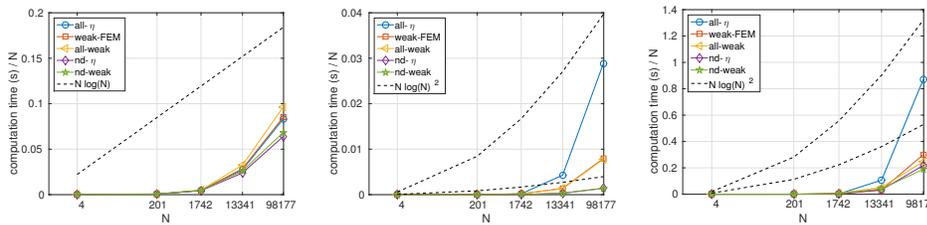


FIG. 7. Computation times in seconds for the computation of the data correlation \mathcal{H} -matrix \mathbf{C}_f (left), the approximate LU-decomposition of the system matrix \mathbf{A} (center), and for the iterative refinement (right) on the dumbbell geometry.

We are also interested in the quality of the approximate LU-decomposition $\mathbf{A} \approx \hat{\mathbf{L}}\hat{\mathbf{U}}$. We use a built-in function of HLib to estimate the deviation $\hat{\mathbf{L}}\hat{\mathbf{U}} - \mathbf{A}$ in the spectral norm by ten power iterations, which is a good indicator of the approximation quality of the LU-decomposition of the finite element matrix. The estimated errors are plotted in Figure 6. Note that the observed behavior is in contrast to the behavior typically observed for preconditioning (cf., e.g., [6]) since we do not increase the rank with the number of unknowns. We can see that the LU-decomposition is most accurate in the all- η case. Still, only one iteration is needed in the iterative refinement in all cases. When it comes to computation times, Figure 7 and Tables 3, 4, and 5 indicate that all cases of admissibility under consideration might yield essentially linear complexity, although the asymptotic regime seems not to be reached in the considered levels of refinement. Both the weak admissibility condition and the nested dissection approach lead to considerable speed-ups, where the combination of these approaches, the nd-weak case, seems to be the fastest approach. Figure 8 illustrates the required average and maximal ranks needed for the computations, whereas Figure 9 illustrates the amount of storage needed per mesh degree of freedom. For reasons of performance, HLib allocates the worst-case scenario for the ranks. Thus, in the latter case, only the different admissibilities for a single \mathcal{H} -matrix built from a binary and a nested dissection cluster tree have to be considered. In conclusion, nested dissection clustering consumes less computation time and less storage for the LU-decomposition.

Having verified the convergence of our solver, we now want to consider different correlation lengths and different classes of smoothness.

TABLE 3

Computation times in seconds to compute the data correlation \mathcal{H} -matrix \mathbf{C}_f on the dumbbell geometry.

Level	1	2	3	4	5
all- η	0.000954	0.123520	8.07820	367.071	8,158.12
weak-FEM	0.001476	0.109641	8.12224	380.533	8,370.79
all-weak	0.001113	0.113122	8.86575	434.289	9,464.19
nd- η	0.001569	0.151018	7.54221	319.773	6,276.75
nd-weak	0.000885	0.123694	8.07407	349.585	6,711.85

TABLE 4

Computation times in seconds to compute the approximate LU-decomposition of the finite element matrix on the dumbbell geometry.

Level	1	2	3	4	5
all- η	$2.2 \cdot 10^{-5}$	0.001274	0.339432	56.2521	2,806.5
weak-FEM	$2.9 \cdot 10^{-5}$	0.00143	0.315344	17.7348	743.624
all-weak	$2.2 \cdot 10^{-5}$	0.001831	0.316497	18.3358	746.364
nd- η	$2.5 \cdot 10^{-5}$	0.000513	0.048170	4.17870	135.778
nd-weak	$3.6 \cdot 10^{-5}$	0.000588	0.050896	4.05899	132.319

TABLE 5

Computation times in seconds for the iterative refinement on the dumbbell geometry.

Level	1	2	3	4	5
all- η	$5.2 \cdot 10^{-5}$	0.011115	4.68355	1,419.19	85,477.9
weak-FEM	0.000104	0.010592	2.64065	420.153	29,492.1
all-weak	$5.4 \cdot 10^{-5}$	0.042098	14.5121	691.129	24,225.6
nd- η	0.000102	0.039209	5.60769	443.310	21,390.0
nd-weak	$4.4 \cdot 10^{-5}$	0.061530	7.67542	544.080	18,570.5

5.3. Small correlation lengths. In the second part of the numerical experiments, we employ correlation kernels with smaller correlation lengths and lower regularity such that low-rank approximations would become prohibitively expensive and sparse tensor product approaches would fail to resolve the concentrated measure.

We consider the screw-nut geometry pictured in Figure 10 which is discretized by a mesh with 269,950 vertices, 197,480 mesh degrees of freedom, and a maximal element diameter of $h/\text{diam}(D) \approx 0.0225$, yielding a matrix equation with $3.90 \cdot 10^{10}$ unknowns. We choose $\mathcal{L} = -\Delta$ in (3), and as input correlation Cor_f we choose either the Gaussian kernel, i.e.,

$$\text{Cor}_f(\mathbf{x}, \mathbf{y}) = \frac{1}{\ell} \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|^2}{2\ell^2}\right),$$

or the exponential kernel, i.e.,

$$\text{Cor}_f(\mathbf{x}, \mathbf{y}) = \frac{1}{\ell} \exp\left(-\frac{\|\mathbf{x} - \mathbf{y}\|}{\ell}\right).$$

Herein, $\ell > 0$ denotes the correlation length.

In the following, we want to demonstrate that the presented method is well suited for small correlation lengths ℓ . We therefore choose the correlation lengths

$$\ell \in \left\{ \frac{\text{diam}(D)}{1}, \frac{\text{diam}(D)}{2}, \frac{\text{diam}(D)}{4}, \frac{\text{diam}(D)}{8}, \frac{\text{diam}(D)}{16}, \frac{\text{diam}(D)}{32} \right\}$$

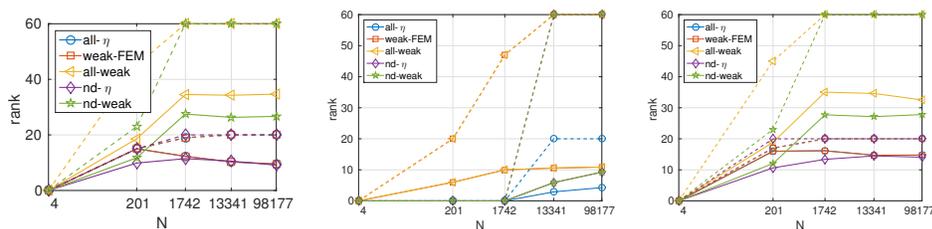


FIG. 8. Required ranks for the prescribed correlation C_f (left), the LU-decomposition of A (middle), and the solution correlation C_u (right). The straight lines indicate the average ranks, whereas the dashed lines illustrate the maximal rank attained.

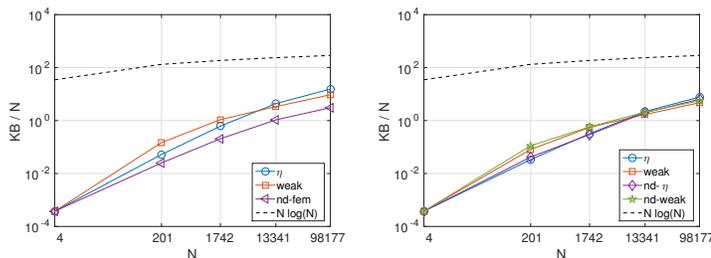


FIG. 9. Allocated storage per mesh degree of freedom for different admissibility conditions. Storage for nonsymmetric (left) and symmetric (right) matrices. The allocated storage is independent of the content of the matrix.

for both the Gaussian kernel and the exponential kernel and compute the corresponding correlation of the solution Cor_u of (3).

In our first experiment, we use the nd-weak case, as the previous section has shown that it is more memory efficient and has superior computation times. The computation time for the assembly of the prescribed correlation is around 20,000 seconds, and the computation time of the approximate LU-decomposition is around 400 seconds, whereas the computation times for the iterative refinement are contained in Table 6.

We do not tabulate the computation times for the Gaussian kernel for the correlation lengths $\ell/\text{diam}(D) = 1/16$ and $\ell/\text{diam}(D) = 1/32$ since the iterative refinement does not converge to the prescribed tolerance. In all other cases, the iterative refinement needs only one iteration.

Repeating the computations in the two problematic cases with increased k_η or in the nd- η instead of the nd-weak case also does not lead to convergence. However, repeating all computations in the all-weak case resolves the issue, as the computation times in Table 6 show. In the all-weak case, the computation time for the prescribed correlation is again around 20,000 seconds, and the computation time for the approximate LU-decomposition is around 1,700 seconds. The iterative refinement needs again one iteration in all tabulated cases.

The cross sections found in Figure 11 illustrate the different behavior of the correlation's trace $Cor_u|_{\mathbf{x}=\mathbf{y}}$ for the different correlation lengths in case of the exponential kernel. The related results for the Gaussian kernel are presented in Figure 12. It seems that a mass defect occurs in the correlation lengths $\ell/\text{diam}(D) = 1/16$ and $\ell/\text{diam}(D) = 1/32$. This could be due to the fact that the mesh size of the finite element method is not able to resolve the correlation length properly. Nevertheless,

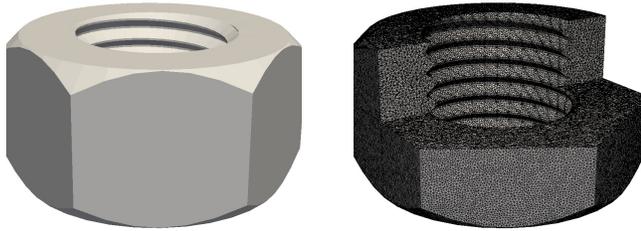


FIG. 10. The screw-nut geometry (left) and its meshed cross section (right).

TABLE 6

Computation times in seconds for the nd-weak case and for the all-weak case for the iterative refinement on the screw-nut geometry for the exponential kernel and the Gaussian kernel with different correlation lengths.

	$\ell/\text{diam}(D)$	1	1/2	1/4	1/8	1/16	1/32
Exponential	nd-weak	51,656.7	53,011.0	52,876.5	51,459.2	49,838.2	51,524.6
	all-weak	77,784.5	79,101.8	79,155.3	79,155.3	76,952.6	72,256.9
Gaussian	nd-weak	47,921.8	50,644.0	50,819.5	51,753.7	—	—
	all-weak	73,405.4	74,877.0	75,165.7	68,222.8	72,259.4	75,070.4

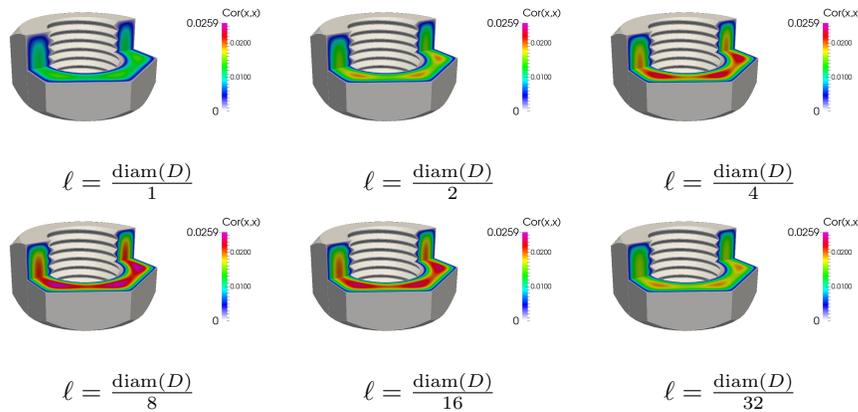


FIG. 11. Cross sections of the correlation of the solution through the screw-nut geometry for the exponential kernel with different correlation lengths ℓ .

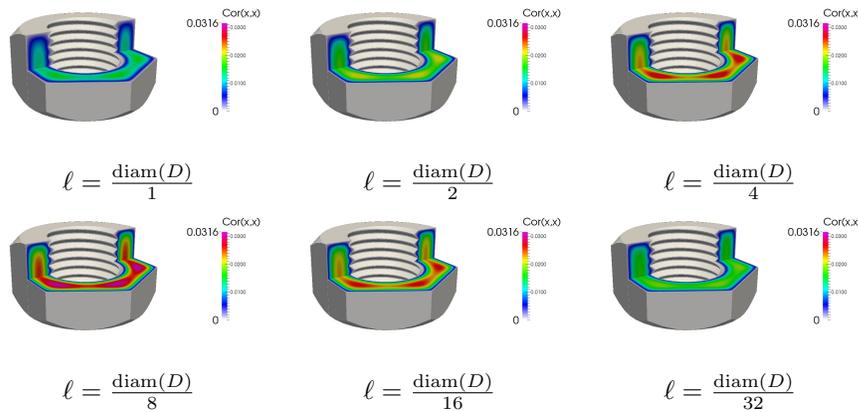


FIG. 12. Cross sections of the correlation of the solution through the screw-nut geometry for the Gaussian kernel with different correlation lengths ℓ .

the computation times are independent of ℓ , even if the underlying finite element method cannot resolve the correlation length. Moreover, the nested dissection clustering technique can lead to a speed-up, while the binary clustering technique seems to be more robust.

6. Conclusion. We considered the solution of strongly elliptic partial differential equations with random load by means of the finite element method. Approximating the full tensor approach by means of \mathcal{H} -matrices, we employed the \mathcal{H} -matrix technique to efficiently discretize the nonlocal correlation kernel of the data and approximate the LU-decomposition of the finite element stiffness matrix. The corresponding \mathcal{H} -matrix equation was then efficiently solved in essentially linear complexity by \mathcal{H} -matrix arithmetic.

Compared to sparse tensor product or low-rank approximations, the proposed method does not suffer in the case of shortly correlated data from large constants in the complexity estimates or lack of resolution of the roughness. This has been shown by numerical experiments on a nontrivial three-dimensional geometry. Indeed, neither the computation times nor the storage requirements increase for correlation kernels with short correlation length. It was moreover demonstrated that the use of the weak admissibility condition for the partition of the \mathcal{H} -matrix improves the constants in the computational complexity without having a significant impact on the solution accuracy. The use of a nested dissection clustering strategy can additionally lead to a speed-up of computation and save storage, whereas the binary clustering strategy seems to be the more robust approach.

Acknowledgments. The authors would like to thank Steffen Börm and the anonymous referees for valuable and helpful remarks on weak admissibility and the nested dissection \mathcal{H} -LU decomposition.

REFERENCES

- [1] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM J. Numer. Anal., 45 (2007), pp. 1005–1034, <https://doi.org/10.1137/050645142>.
- [2] I. BABUŠKA, R. TEMPONE, AND G. ZOURARIS, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM J. Numer. Anal., 42 (2004), pp. 800–825, <https://doi.org/10.1137/S0036142902418680>.
- [3] M. BEBENDORF, *Approximation of boundary element matrices*, Numer. Math., 86 (2000), pp. 565–589.
- [4] M. BEBENDORF, *Efficient inversion of the Galerkin matrix of general second-order elliptic operators with nonsmooth coefficients*, Math. Comp., 74 (2005), pp. 1179–1199.
- [5] M. BEBENDORF, *Why finite element discretizations can be factored by triangular hierarchical matrices*, SIAM J. Numer. Anal., 45 (2007), pp. 1472–1494, <https://doi.org/10.1137/060669747>.
- [6] M. BEBENDORF, *Hierarchical Matrices*, Lect. Notes Comput. Sci. Eng. 63, Springer, Berlin, Heidelberg, 2008.
- [7] M. BEBENDORF AND W. HACKBUSCH, *Existence of \mathcal{H} -matrix approximants to the inverse FE-matrix of elliptic operators with L^∞ -coefficients*, Numer. Math., 95 (2003), pp. 1–28.
- [8] S. BÖRM, *Efficient Numerical Methods for Non-local Operators*, EMS Tracts Math. 14, European Mathematical Society (EMS), Zürich, Switzerland, 2010.
- [9] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Hierarchical Matrices*, Tech. report 21, Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany, 2003.
- [10] L. BOUTET DE MONVEL AND P. KRÉE, *Pseudo-differential operators and Gevrey classes*, Ann. Inst. Fourier (Grenoble), 17 (1967), pp. 295–323.
- [11] I. BRAINMAN AND S. TOLEDO, *Nested-dissection orderings for sparse LU with partial pivoting*, SIAM J. Matrix Anal. Appl., 23 (2002), pp. 998–1012, <https://doi.org/10.1137/S0895479801385037>.

- [12] R. CAFLISCH, *Monte Carlo and quasi-Monte Carlo methods*, Acta Numer., 7 (1998), pp. 1–49.
- [13] M. DEB, I. BABUŠKA, AND J. ODEN, *Solution of stochastic partial differential equations using Galerkin finite element techniques*, Comput. Methods Appl. Mech. Engrg., 190 (2001), pp. 6359–6372.
- [14] J. DÖLZ, H. HARBRECHT, AND M. PETERS, *\mathcal{H} -matrix accelerated second moment analysis for potentials with rough correlation*, J. Sci. Comput., 65 (2015), pp. 387–410.
- [15] J. DÖLZ, H. HARBRECHT, AND C. SCHWAB, *Covariance regularity and \mathcal{H} -matrix approximation for rough random fields*, Numer. Math., 135 (2017), pp. 1045–1071.
- [16] M. FAUSTMANN, *Approximation inverser Finite Elemente- und Randelementmatrizen mittels hierarchischer Matrizen*, Ph.D. Thesis, Technische Universität Wien, Wien, Austria, 2015.
- [17] M. FAUSTMANN, J. M. MELENK, AND D. PRAETORIUS, *\mathcal{H} -matrix approximability of the inverses of FEM matrices*, Numer. Math., 131 (2015), pp. 615–642.
- [18] P. FRAUENFELDER, C. SCHWAB, AND R. TODOR, *Finite elements for elliptic problems with stochastic coefficients*, Comput. Methods Appl. Mech. Engrg., 194 (2005), pp. 205–228.
- [19] A. GEORGE, *Nested dissection of a regular finite element mesh*, SIAM J. Numer. Anal., 10 (1973), pp. 345–363, <https://doi.org/10.1137/0710032>.
- [20] R. GHANEM AND P. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer, New York, 1991.
- [21] D. GILBARG AND N. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, Springer, Berlin, Heidelberg, 1977.
- [22] G. GOLUB AND C. VAN LOAN, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, MD, 2013.
- [23] L. GRASEDYCK AND W. HACKBUSCH, *Construction and arithmetics of \mathcal{H} -matrices*, Computing, 70 (2003), pp. 295–334.
- [24] L. GRASEDYCK, R. KRIEMANN, AND S. LE BORNE, *Domain decomposition based \mathcal{H} -LU preconditioning*, Numer. Math., 112 (2009), pp. 565–600.
- [25] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.
- [26] M. GRIEBEL AND H. HARBRECHT, *Approximation of bi-variate functions: Singular value decomposition versus sparse grids*, IMA J. Numer. Anal., 34 (2014), pp. 28–54.
- [27] W. HACKBUSCH, *A sparse matrix arithmetic based on \mathcal{H} -matrices part I: Introduction to \mathcal{H} -matrices*, Computing, 62 (1999), pp. 89–108.
- [28] W. HACKBUSCH, *Hierarchical Matrices: Algorithms and Analysis*, Springer, Heidelberg, 2015.
- [29] W. HACKBUSCH AND B. KHOROMSKIJ, *A sparse \mathcal{H} -matrix arithmetic. General complexity estimates*, J. Comput. Appl. Math., 125 (2000), pp. 479–501.
- [30] W. HACKBUSCH, B. KHOROMSKIJ, AND R. KRIEMANN, *Hierarchical matrices based on a weak admissibility criterion*, Computing, 73 (2004), pp. 207–243.
- [31] W. HACKBUSCH, B. KHOROMSKIJ, AND E. TYRTYSHNIKOV, *Approximate iterations for structured matrices*, Numer. Math., 109 (2008), pp. 365–383.
- [32] H. HARBRECHT, *A finite element method for elliptic problems with stochastic input data*, Appl. Numer. Math., 60 (2010), pp. 227–244.
- [33] H. HARBRECHT, M. PETERS, AND R. SCHNEIDER, *On the low-rank approximation by the pivoted Cholesky decomposition*, Appl. Numer. Math., 62 (2012), pp. 28–440.
- [34] H. HARBRECHT, M. PETERS, AND M. SIEBENMORGEN, *Combination technique based k -th moment analysis of elliptic problems with random diffusion*, J. Comput. Phys., 252 (2013), pp. 128–141.
- [35] H. HARBRECHT, R. SCHNEIDER, AND C. SCHWAB, *Multilevel frames for sparse tensor product spaces*, Numer. Math., 110 (2008), pp. 199–220.
- [36] H. HARBRECHT, R. SCHNEIDER, AND C. SCHWAB, *Sparse second moment analysis for elliptic problems in stochastic domains*, Numer. Math., 109 (2008), pp. 385–414.
- [37] B. HENDRICKSON AND E. ROTHBERG, *Improving the run time and quality of nested dissection ordering*, SIAM J. Sci. Comput., 20 (1998), pp. 468–489, <https://doi.org/10.1137/S1064827596300656>.
- [38] L. HÖRMANDER, *The Analysis of Linear Partial Differential Operators. I: Distribution Theory and Fourier Analysis*, reprint of the 2nd (1990) ed., Classics in Mathematics, Springer, Berlin, 2003.
- [39] L. HÖRMANDER, *The Analysis of Linear Partial Differential Operators. III: Pseudo-Differential Operators*, reprint of the 1994 ed., Classics in Mathematics, Springer, Berlin, 2007.
- [40] G. HSIAO AND W. WENDLAND, *Boundary Integral Equations*, Appl. Math. Sci. 164, Springer, Berlin, 2008.
- [41] B. KERNIGHAN AND D. RITCHIE, *The C Programming Language*, Prentice-Hall, Upper Saddle River, NJ, 1988.

- [42] B. KHOROMSKIJ AND C. SCHWAB, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM J. Sci. Comput., 33 (2011), pp. 364–385, <https://doi.org/10.1137/100785715>.
- [43] R. KRIEMANN, *Parallel-matrix arithmetics on shared memory systems*, Computing, 74 (2005), pp. 273–297.
- [44] R. KRIEMANN, *\mathcal{H} -LU factorization on many-core systems*, Comput. Vis. Sci., 16 (2013), pp. 105–117.
- [45] R. J. LIPTON, D. J. ROSE, AND R. E. TARJAN, *Generalized nested dissection*, SIAM J. Numer. Anal., 16 (1979), pp. 346–358, <https://doi.org/10.1137/0716027>.
- [46] B. MATÉRN, *Spatial Variation*, Lecture Notes in Statistics 36, Springer, New York, 1986.
- [47] H. MATTHIES AND A. KEESE, *Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations*, Comput. Methods Appl. Mech. Engrg., 194 (2005), pp. 1295–1331.
- [48] C. MOLER, *Iterative refinement in floating point*, J. ACM, 14 (1967), pp. 316–321.
- [49] 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, <https://doi.org/10.1137/070680540>.
- [50] C. RASMUSSEN AND C. WILLIAMS, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press, Cambridge, MA, 2005.
- [51] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003, <https://doi.org/10.1137/1.9780898718003>.
- [52] C. SCHWAB AND R. TODOR, *Sparse finite elements for elliptic problems with stochastic loading*, Numer. Math., 95 (2003), pp. 707–734.
- [53] C. SCHWAB AND R. TODOR, *Karhunen-Loève approximation of random fields by generalized fast multipole methods*, J. Comput. Phys., 217 (2006), pp. 100–122.
- [54] H. SI, *Tetgen, a Delaunay-based quality tetrahedral mesh generator*, ACM Trans. Math. Softw., 41 (2015), 11.
- [55] M. TAYLOR, *Pseudodifferential Operators*, Princeton Math. Ser. 34, Princeton University Press, Princeton, NJ, 1981.
- [56] E. TYRTYSHNIKOV, *Mosaic-skeleton approximations*, Calcolo, 33 (1996), pp. 47–57.
- [57] T. VON PETERSDORFF AND C. SCHWAB, *Sparse finite element methods for operator equations with stochastic data*, Appl. Math., 51 (2006), pp. 145–180.
- [58] J. WILKINSON, *Rounding Errors in Algebraic Processes*, Prentice-Hall, Englewood Cliffs, NJ, 1963.
- [59] D. XIU AND D. TARTAKOVSKY, *Numerical methods for differential equations in random domains*, SIAM J. Sci. Comput., 28 (2006), pp. 1167–1185, <https://doi.org/10.1137/040613160>.