

THE INVERSE FAST MULTIPOLE METHOD: USING A FAST APPROXIMATE DIRECT SOLVER AS A PRECONDITIONER FOR DENSE LINEAR SYSTEMS

PIETER COULIER^{†‡}, HADI POURANSARI[‡], AND ERIC DARVE[‡]

Abstract. Although some preconditioners are available for solving dense linear systems, there are still many matrices for which preconditioners are lacking, in particular in cases where the size of the matrix N becomes very large. Examples of preconditioners include ILU preconditioners that sparsify the matrix based on some threshold, algebraic multigrid, and specialized preconditioners, e.g., Calderón and other analytical approximation methods when available. Despite these methods, there remains a great need to develop general purpose preconditioners whose cost scales well with the matrix size N . In this paper, we propose a preconditioner, with broad applicability and with cost $\mathcal{O}(N)$, for dense matrices, when the matrix is given by a smooth kernel. Extending the method using the same framework to general \mathcal{H}^2 -matrices (i.e., algebraic instead of defined in terms of an analytical kernel) is relatively straightforward, but this won't be discussed here. These preconditioners have a controlled accuracy (e.g., machine accuracy can be achieved if needed) and scale linearly with N . They are based on an approximate direct solve of the system. The linear scaling of the algorithm is achieved by means of two key ideas. First, the \mathcal{H}^2 -structure of the dense matrix is exploited to obtain an extended sparse system of equations. Second, fill-ins arising when performing the elimination of the latter are compressed as low-rank matrices if they correspond to well-separated interactions. This ensures that the sparsity pattern of the extended sparse matrix is preserved throughout the elimination, hence resulting in a very efficient algorithm with $\mathcal{O}(N \log(1/\varepsilon)^2)$ computational cost and $\mathcal{O}(N \log 1/\varepsilon)$ memory requirement, for an error tolerance $0 < \varepsilon < 1$. The solver is inexact, although the error can be controlled and made as small as needed. These solvers are related to ILU in the sense that the fill-in is controlled. However, in ILU, most of the fill-in (i.e., below a certain tolerance) is simply discarded, whereas here it is approximated using low-rank blocks, with a prescribed tolerance. Numerical examples are discussed to demonstrate the linear scaling of the method and to illustrate its effectiveness as a preconditioner.

Key words. Fast direct solver, preconditioner, \mathcal{H}^2 -matrices, extended sparsification, low-rank compression.

AMS subject classifications. 65F05, 65F08, 65Y20.

1. Introduction. Large linear systems of equations $\mathbf{Ax} = \mathbf{b}$ involving dense matrices \mathbf{A} arise in a variety of applications in science and engineering (e.g., in boundary element methods [6, 32], Kalman filtering [34], radial basis function interpolation [9, 23], etc.). Storing a dense matrix entrywise requires a quadratic amount of memory $\mathcal{O}(N^2)$ (where N indicates the matrix size), while solving the corresponding system of equations with a direct solver involves $\mathcal{O}(N^3)$ flops. In order to alleviate these stringent requirements, fast methods have been developed in the past decades, including the fast multipole method (FMM) [16, 21, 39, 43, 51], the panel clustering technique [29, 48], and methods based on hierarchical matrices [4, 8, 24, 26, 27]. These methods rely on the observation that many dense matrices can be represented in a hierarchical format, i.e., the original matrix can be subdivided into a hierarchy of smaller block matrices and most of these blocks are replaceable by low-rank approximations. Several categories of hierarchical matrix representations can be distinguished. In the simplest case, all off-diagonal blocks are assumed to be low-rank, which is referred to as a weak admissibility criterion [28]. If the matrix entries can be interpreted as

[†]KU Leuven, Department of Civil Engineering, Kasteelpark Arenberg 40, 3001 Leuven, Belgium (pieter.coulier@bwk.kuleuven.be).

[‡]Stanford University, Department of Mechanical Engineering, 496 Lomita Mall, 94305 Stanford, CA, USA ([pcoulier,hadip,darve@stanford.edu]).

interactions between points in a physical domain, this criterion implies that all the interactions of a certain cluster of points with any other cluster (except itself) are low-rank. This assumption leads to hierarchically off-diagonal low-rank (HODLR) matrices [3] and hierarchically semi-separable (HSS) matrices [12, 49]. In the latter, a nested approach is used (i.e., the low-rank basis at a certain hierarchical level is constructed using the low-rank basis at the child level), while this is not the case in the former. If a more general strong admissibility criterion is adopted, \mathcal{H} - and \mathcal{H}^2 -matrices are obtained for the non-nested and nested cases, respectively [4, 8]*. In these matrices, only interactions between non-neighboring clusters of points are assumed to be low-rank. The FMM can be considered as a subcategory of \mathcal{H}^2 -matrices [4]. The reader is referred to [4, 8] for a detailed overview of the various hierarchical structures. See also Appendix B for additional discussion of this classification.

The use of the aforementioned hierarchical matrix representations facilitates matrix-vector multiplications, hence paving the way for iterative solvers such as Krylov subspace methods for solving the linear system $\mathbf{Ax} = \mathbf{b}$ [45]. The precarious convergence behavior — and thus the number of iterations — is a major disadvantage of iterative solvers. This necessitates the incorporation of a preconditioning strategy to improve the latter, which is often a difficult and case-dependent problem. Typical preconditioners for dense matrices include the incomplete LU factorization (with thresholding) [37], methods based on the sparse approximate inverse [10], geometric and algebraic multigrid methods [36, 50], as well as analytical preconditioning techniques such as Calderón preconditioners [14, 15]. Although these techniques are usually well suited for the specific problem they have been designed for, their performance as robust and efficient black-box preconditioners for general problems remains rather limited. Another disadvantage of iterative solvers is that they do not lead to a factorization of \mathbf{A} that can easily be re-used, implying that the algorithm needs to be restarted if the right hand side \mathbf{b} of the system is modified. Direct solvers, on the other hand, are known to be more robust and are able to cope with multiple right hand sides, but at a high computational cost. Significant progress has been made in recent years regarding the development of fast direct solvers for hierarchical matrices, especially for the aforementioned category of HODLR and HSS matrices. For example, direct $\mathcal{O}(N)$ algorithms for HSS matrices have been presented by, among others, Martinsson and Rokhlin [38], Chandrasekaran et al. [12], and Gillman et al. [19], while fast solvers for HODLR matrices can be found in [3, 35]. These methods can be applied as high accuracy direct solvers but are often more effective if employed as preconditioners in an iterative solver. The framework used for HODLR and HSS matrices is appropriate for one-dimensional (1D) applications. For higher dimensional problems (2D or 3D), the assumption that all off-diagonal blocks are low-rank does not hold in that case, that is the rank typically grows with N . For 2D problems (points distributed on a 2D manifold or 2D plane), the rank grows like $\mathcal{O}(\sqrt{N})$ [1], hence jeopardizing the computational efficiency of the algorithms. In order to obtain a direct solver with a linear complexity $\mathcal{O}(N)$ for 2D and 3D problems, the more general framework of \mathcal{H} - and \mathcal{H}^2 -matrices should be adopted.

In this paper, the inverse fast multipole method (IFMM) is presented as a fast

*There is a confusing point of terminology associated with strong and weak, which cannot really be removed. If we use a strong admissibility criterion, we are considering that all well-separated blocks are low-rank. This therefore must correspond to weak hierarchical matrices or $w\mathcal{H}$ (e.g., FMM matrices). Similarly, if we use a weak admissibility criterion, we get strong hierarchical matrices or $s\mathcal{H}$ (e.g., HODLR and HSS). This terminology is imposed by the fact that we must have $s\mathcal{H} \subset w\mathcal{H}$.

direct solver for dense \mathcal{H}^2 -matrices, where “direct” refers to the use of an approximate LU factorization [2]. The solver is inexact, although the error can be controlled and made as small as needed. A low accuracy solver is hence applicable as a robust and efficient preconditioner for iterative schemes. The IFMM relies on two key ideas. First, the dense \mathcal{H}^2 -matrix is converted to an extended sparse matrix through the introduction of auxiliary variables. Second, fill-ins arising during the elimination of the latter are compressed as low-rank matrices if they correspond to well-separated (i.e., non-neighboring) interactions. As a result, the sparsity of the extended sparse matrix is maintained throughout the elimination. Specifically, there is an upper bound on the number of non-zero entries in the matrix independent of N , at all steps during the elimination. This results in a very efficient algorithm with linear complexity. To the authors’ knowledge, this is the first direct solver that achieves this complexity for the most general category of hierarchical matrices [2].

The proposed strategy shows similarities to multigrid methods as it uses a hierarchical decomposition of the domain, as well as to the incomplete LU (ILU) factorization and its variants as the sparsity pattern is preserved. Contrary to ILU, the sparsity is maintained using low-rank compression instead of thresholding.

The outline of this paper is as follows. Section 2 introduces the IFMM and focuses on the key ideas required to achieve linear complexity; special attention is paid to a graph representation of the method. Numerical examples are subsequently discussed in Section 3 to illustrate the computational efficiency of the methodology, both as a direct solver and as a preconditioner in an iterative scheme. The applicability of the method to a Stokes flow problem is finally demonstrated in Section 4, while summarizing conclusions are drawn in Section 5. In Appendix B, we included a discussion of theoretical connections between the method discussed in this paper and the algorithms developed by Hackbusch et al. [4, 8, 24, 26, 27].

2. The inverse fast multipole method: key ideas and algorithm.

2.1. Extended sparsification. It is the aim to obtain a fast solver for the dense linear system $\mathbf{Ax} = \mathbf{b}$, where \mathbf{A} is an \mathcal{H}^2 -matrix, i.e., a matrix constructed based on strong admissibility criteria (well-separated clusters are assumed to be low-rank only), and with a nested basis for the low-rank representations.

The first step in obtaining such a solver is to exploit the matrix’ \mathcal{H}^2 -structure in order to convert the original dense system into an extended sparse system, as solving the latter is more efficient from a computational point of view. The sparsification procedure is based on recursively decomposing a dense matrix into a sparse matrix (representing the “near field” interactions) and a low-rank approximation (representing the “far field” interactions), combined with the introduction of auxiliary variables. This is done in a multilevel fashion. Similar ideas have been considered for HSS matrices in [13, 31].

The sparsification concept is illustrated in the following paragraphs for a one-dimensional (1D) example, shown in Figure 1. A 1D domain Ω is hierarchically subdivided into clusters $\Omega_i^{(l)}$ (with super- and subscripts $^{(l)}$ and i referring to the hierarchical level and the cluster index at that level, respectively) until the leaf level $l = L$ is reached (with $L = 3$ in Figure 1). Blocks in the matrix \mathbf{A} corresponding to interactions between well-separated clusters $\Omega_i^{(l)}$ and $\Omega_j^{(l)}$ are replaced by low-rank approximations of rank r , i.e., $\mathbf{A}_{ij}^{(l)} \simeq \mathbf{U}_i^{(l)} \mathbf{K}_{ij}^{(l)} \mathbf{V}_j^{(l)\top}$ (see Figure 2b). For the example of Figure 1, the resulting matrix structure is shown in Figure 2a. Note that the nested structure of the \mathcal{H}^2 -matrix is not shown in this figure. For readers familiar with the

FMM, the matrices $\mathbf{U}_i^{(l)}$, $\mathbf{K}_{ij}^{(l)}$, and $\mathbf{V}_j^{(l)\top}$ can be identified as L2P- (or L2L-), M2L-, and P2M- (or M2M-) operators, respectively.




FIG. 1. One-dimensional domain Ω hierarchically subdivided into clusters $\Omega_i^{(l)}$.




FIG. 2. (a) Approximation of the dense matrix by an \mathcal{H}^2 -matrix. (b) Low rank approximation of an interaction between well-separated clusters: $\mathbf{A}_{ij}^{(l)} \simeq \mathbf{U}_i^{(l)} \mathbf{K}_{ij}^{(l)} \mathbf{V}_j^{(l)\top}$.

The dense matrix \mathbf{A} can consequently be decomposed at the leaf level as follows:

$$(1) \quad \mathbf{A} = \mathbf{S}^{(3)} + \mathbf{U}^{(3)} \mathbf{K}^{(3)} \mathbf{V}^{(3)\top}$$

The matrix $\mathbf{S}^{(3)}$ in Eq. (1) is a sparse matrix containing the interactions of leaf clusters with their neighbors (including the self-interactions). The full rank matrix blocks can be denoted as P2P-operators if the FMM terminology is employed. In 1D, $\mathbf{S}^{(3)}$ is a tri-diagonal matrix. However, for general point distributions, this matrix is a general sparse matrix, with no “particular” structure. The matrix $\mathbf{U}^{(3)} \mathbf{K}^{(3)} \mathbf{V}^{(3)\top}$ represents the dense fill-ins arising from far field interactions. The decomposition presented in Eq. (1) is pictorially depicted in Figure 3 for a 1D example and in Figure 4 for a 2D example.

The matrix $\mathbf{K}^{(3)}$ in Eq. (1) is representative for all far field interactions. It is still rather dense and is therefore further decomposed using a similar decomposition as for \mathbf{A} :

$$(2) \quad \mathbf{A} = \mathbf{S}^{(3)} + \mathbf{U}^{(3)} \left(\mathbf{S}^{(2)} + \mathbf{U}^{(2)} \mathbf{K}^{(2)} \mathbf{V}^{(2)\top} \right) \mathbf{V}^{(3)\top}$$

This decomposition at level $l = 2$ is illustrated in Figure 5 and 6 for a 1D and 2D example, respectively. The sparse matrix $\mathbf{S}^{(2)}$ in Eq. (2) represents near field interactions at level $l = 2$ (that have not been accounted for through $\mathbf{S}^{(3)}$), while $\mathbf{K}^{(2)}$




FIG. 3. Decomposition of the \mathcal{H}^2 -matrix \mathbf{A} at the leaf level L (with $L = 3$): $\mathbf{A} = \mathbf{S}^{(3)} + \mathbf{U}^{(3)}\mathbf{K}^{(3)}\mathbf{V}^{(3)\top}$.




FIG. 4. Quadtree decomposition of the matrix. This figure is showing the decomposition at level 3 for a 2D test case. This shows how the sketch in Figure 3 is generalized in dimensions higher than 1.

represents the remaining far field interactions. As no well-separated clusters exist at level $l = 1$, the matrix $\mathbf{K}^{(2)}$ can not be further decomposed.

Eq. (2) is still a dense system, but can now be converted into an extended sparse system by introducing auxiliary variables at each level that contains well-separated clusters (i.e., $l \geq 2$). More specifically, multipole coefficients $\mathbf{y}^{(l)}$ are defined as:

$$(3) \quad \mathbf{y}^{(3)} = \mathbf{V}^{(3)\top} \mathbf{x}$$

$$(4) \quad \mathbf{y}^{(2)} = \mathbf{V}^{(2)\top} \mathbf{y}^{(3)}$$

This allows rewriting the equation $\mathbf{A}\mathbf{x} = \mathbf{b}$ as:

$$(5) \quad \mathbf{S}^{(3)}\mathbf{x} + \mathbf{U}^{(3)} \left(\mathbf{S}^{(2)}\mathbf{y}^{(3)} + \mathbf{U}^{(2)}\mathbf{K}^{(2)}\mathbf{y}^{(2)} \right) = \mathbf{b}$$

Furthermore, local coefficients $\mathbf{z}^{(l)}$ are identified at each level:

$$(6) \quad \mathbf{z}^{(2)} = \mathbf{K}^{(2)}\mathbf{y}^{(2)}$$

$$(7) \quad \mathbf{z}^{(3)} = \mathbf{S}^{(2)}\mathbf{y}^{(3)} + \mathbf{U}^{(2)}\mathbf{z}^{(2)}$$

The original equation $\mathbf{A}\mathbf{x} = \mathbf{b}$ thus becomes:

$$(8) \quad \mathbf{S}^{(3)}\mathbf{x} + \mathbf{U}^{(3)}\mathbf{z}^{(3)} = \mathbf{b}$$




FIG. 5. Decomposition of the matrix $\mathbf{K}^{(3)}$ at level l (with $l = 2$): $\mathbf{K}^{(3)} = \mathbf{S}^{(2)} + \mathbf{U}^{(2)}\mathbf{K}^{(2)}\mathbf{V}^{(2)\top}$.




FIG. 6. Quadtree decomposition at level 2 for a 2D test case. Compare with Figure 5 for a 1D distribution of points.

Eq. (3)–(8) provide an extended sparse system of equations:

$$(9) \quad \begin{bmatrix} \mathbf{S}^{(3)} & \mathbf{U}^{(3)} \\ \mathbf{V}^{(3)\top} & -\mathbf{I} \\ & -\mathbf{I} & \mathbf{S}^{(2)} & \mathbf{U}^{(2)} \\ & & \mathbf{V}^{(2)\top} & -\mathbf{I} \\ & & & -\mathbf{I} & \mathbf{K}^{(2)} \end{bmatrix} \begin{Bmatrix} \mathbf{x} \\ \mathbf{z}^{(3)} \\ \mathbf{y}^{(3)} \\ \mathbf{z}^{(2)} \\ \mathbf{y}^{(2)} \end{Bmatrix} = \begin{Bmatrix} \mathbf{b} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{Bmatrix}$$

The extended sparse matrix in Eq. (9) (indicated as \mathbf{E} in the following) contains matrices from the decompositions introduced in Eq. (1) and (2) and exhibits a symmetric pattern in terms of fill-in, as is shown in Figure 7 and 8. In fact, for a symmetric matrix \mathbf{A} , the matrix \mathbf{E} is symmetric as well. The elimination and substitution processes are thus performed on this sparse system rather than on the original dense system.

This sparsification procedure is a very general and powerful technique that is not restricted to 1D or 2D problems. In general, any dense \mathcal{H}^2 -matrix can be recursively decomposed into sparse matrices and low-rank approximations to obtain an extended sparsified matrix.

In order to further illustrate the sparsification procedure, it is instructive to consider graph representations of the linear system. Figure 9a and 9b show the graphs of the original dense system and the extended sparse system corresponding to the 1D problem introduced in Figure 1, respectively. Vertices in a graph represent the variables, while the edges correspond to matrix operators; the same color convention as in Figure 7 is employed. The transformation from Figure 9a to 9b clearly demonstrates how one can take advantage of the matrix' \mathcal{H}^2 -structure to obtain a sparsified system. This graph representation will play a crucial role for understanding the next key step of the algorithm, as will be elucidated in the following subsection. Figure 9b also illustrates that the variables $\mathbf{y}^{(l)}$ at level l can be considered as unknowns $\tilde{\mathbf{x}}^{(l-1)}$




FIG. 7. A pictorial representation of the extended sparse system of Eq. (9) for a 1D distribution of points. Red blocks are P2P-operators, cyan blocks are L2P-, P2M-, L2L-, and M2M-operators, and orange blocks are M2L-operators. The black blocks represent negative identity matrices. The original unknowns \mathbf{x} are shown in blue, the locals $\mathbf{z}^{(l)}$ in light green, and the multipoles $\mathbf{y}^{(l)}$ in magenta. The right hand side vector \mathbf{b} is depicted in gray.




FIG. 8. Extended sparse matrix for a 2D distribution of points. Compare with the sketch shown in Figure 7 for a 1D distribution of points.

at the parent level $l - 1$, hence revealing the hierarchical structure of the \mathcal{H}^2 -matrix (with $\tilde{\mathbf{x}}^{(L)}$ corresponding to \mathbf{x}).

2.2. Compression and redirection of fill-ins. Transforming a dense \mathcal{H}^2 -matrix into an extended sparse matrix is a crucial but insufficient step to obtain a fast solver. Indeed, feeding a matrix such as the ones depicted in Figure 7 and Figure 8 to a conventional sparse solver leads to the creation of dense fill-ins throughout the elimination phase (or, equivalently, additional edges in the graph of Figure 9b), which is harmful for the scaling of the algorithm (see for example the thesis of Pals [40]). In order to preserve the initial sparsity pattern, a second key idea is therefore introduced:




FIG. 9. The graph of (a) the original dense system and (b) the extended sparse system, corresponding to the 1D problem of Figure 1. The same color convention as in Figure 7 is adopted.

fill-ins corresponding to interactions between well-separated clusters are compressed into a low-rank representation and are subsequently redirected through other operators/edges. This procedure is discussed step by step in the next paragraphs, focusing on a graph representation of the methodology. To this end, the one-dimensional example introduced in Subsection 2.1 is abandoned and a more general case is now considered. References to the original FMM-terminology (e.g., particles, multipoles, and locals) are omitted to stress the generality of the method.

The following simplifying notations are introduced (see Figure 10): nodes in a graph corresponding to the variables $\tilde{\mathbf{x}}_i^{(l)}$, $\mathbf{z}_i^{(l)}$, and $\mathbf{y}_i^{(l)}$ of a cluster $\Omega_i^{(l)}$ are referred to as nodes $\underline{\mathbf{i}}^{(l)}$, $\mathbf{i}^{(l)}$, and $\overline{\mathbf{i}}^{(l)}$, respectively (the position of the bar is a mnemonic corresponding to the location of the node in the tree, either below or above). The symbol $\mathbf{i}^\dagger^{(l)}$ represents the variable $\mathbf{z}_p^{(l-1)}$ of its parent cluster $\Omega_p^{(l-1)}$. The latter can also be denoted as $\mathbf{p}^{(l-1)}$. If no confusion is possible, the superscripts referring to the hierarchical level of a cluster will be omitted in the following for sake of readability.

With this notation in place, operators such as the inter- and ant interpolation[†] matrices \mathbf{U}_i and \mathbf{V}_i^T of a cluster Ω_i are written as $\mathbf{U}(\underline{\mathbf{i}}, \mathbf{i})$ and $\mathbf{V}^T(\mathbf{i}, \underline{\mathbf{i}})$, respectively.




FIG. 10. Alternative notation for the nodes related to cluster $\Omega_i^{(l)}$.

Figure 11a shows part of a graph of an extended sparse matrix at a certain stage of the elimination phase of level l , i.e., after elimination of levels $L, L-1, \dots, l+1$. This graph shows a “generic” description of the process of removing fill-in. The nodes shown are generic nodes in the graph of the matrix.

The variables $\tilde{\mathbf{x}}^{(l)}$ and $\mathbf{z}^{(l)}$ of some clusters at level l have already been eliminated at this point, e.g., $\tilde{\mathbf{x}}_h^{(l)}$ and $\mathbf{z}_h^{(l)}$ of cluster $\Omega_h^{(l)}$ (or, equivalently, the nodes $\underline{\mathbf{h}}$ and \mathbf{h}). This has created edges in the graph between neighboring clusters, e.g., connections $(\bar{\mathbf{h}}, \underline{\mathbf{i}})$ and $(\underline{\mathbf{i}}, \bar{\mathbf{h}})$ between $\mathbf{y}_h^{(l)}$ and $\tilde{\mathbf{x}}_i^{(l)}$ (and vice versa). Additional edges such as $(\bar{\mathbf{h}}, \bar{\mathbf{h}})$ (i.e., between $\mathbf{y}_h^{(l)}$ and itself) have been established as well. The variables $\tilde{\mathbf{x}}_i^{(l)}$ and $\mathbf{z}_i^{(l)}$ of cluster $\Omega_i^{(l)}$ that are to be eliminated next are highlighted in Figure 11a. Performing this elimination results in new interactions between all neighbors of $\Omega_i^{(l)}$ (including itself), as is illustrated in Figure 11b[‡]. Some of these fill-ins establish direct connections between clusters that did not exist prior to this elimination step, however. For example, fill-ins $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$ and $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ between $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$ arise (where \mathbf{E}' denotes a fill-in in the extended sparse matrix). As the aforementioned clusters are well-separated (i.e., their existing interactions are accounted for through the edges between $\bar{\mathbf{k}}$ and $\bar{\mathbf{j}}$ rather than between $\underline{\mathbf{k}}$ and $\underline{\mathbf{j}}$), these fill-ins are low-rank and can be approximated as follows:

$$(10) \quad \mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \simeq \mathbf{U}'(\underline{\mathbf{k}}, \underline{\mathbf{k}}) \mathbf{K}'(\bar{\mathbf{k}}, \bar{\mathbf{j}}) \mathbf{V}'^T(\underline{\mathbf{j}}, \underline{\mathbf{j}})$$

$$(11) \quad \mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \simeq \mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{K}'(\bar{\mathbf{j}}, \bar{\mathbf{k}}) \mathbf{V}'^T(\underline{\mathbf{k}}, \underline{\mathbf{k}})$$

It is assumed in the following that Eq. (10) and (11) represent partial singular value decompositions (SVD), implying that the matrices \mathbf{U}' and \mathbf{V}' are orthonormal and that \mathbf{K}' is a diagonal matrix containing the singular values (i.e., $\mathbf{K}'(\bar{\mathbf{k}}, \bar{\mathbf{j}}) = \boldsymbol{\Sigma}'(\bar{\mathbf{k}}, \bar{\mathbf{j}})$ and $\mathbf{K}'(\bar{\mathbf{j}}, \bar{\mathbf{k}}) = \boldsymbol{\Sigma}'(\bar{\mathbf{j}}, \bar{\mathbf{k}})$). The implications for the algorithm if this is not the case will be discussed in Subsection 2.4. It is also assumed that a criterion is available for

[†]Interpolation = from coarse to fine. Anterpolation = from fine to coarse.

[‡]Eliminating a node also implies that the right hand side of the system needs to be updated.

determining an appropriate rank for approximations (10) and (11) (see Subsection 2.4 for more details).

These low-rank approximations for the fill-in enable the calculation of new interpolation and antinterpolation operators for $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$, respectively. For example, a new interpolation operator $\widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ for $\Omega_j^{(l)}$ is obtained by concatenating the existing scaled interpolation operator $\widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}_{\mathbf{U}}$ (where $\boldsymbol{\Sigma}_{\mathbf{U}}$ contains the weights associated with each column of $\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$)[§] and the scaled fill-in operator $\mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}'(\overline{\mathbf{j}}, \overline{\mathbf{k}})$ and subsequently applying a low-rank recompression (thin SVD):

$$(12) \quad \left[\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}_{\mathbf{U}} \mid \mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}'(\overline{\mathbf{j}}, \overline{\mathbf{k}}) \right] \simeq \widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{U}} \boldsymbol{\Phi}_j$$

$\boldsymbol{\Phi}_j = [\phi_j \mid \phi'_j]$ represents a fat, row orthogonal, matrix. It is important to note that the recompression in Eq. (12) is applied to the scaled interpolation operators rather than to $\left[\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mid \mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \right]$, as this allows to account for the correct “weight” of each vector when obtaining a new basis.[¶] Eq. (12) allows expressing the existing and fill-in operators $\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ and $\mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ in terms of the new basis $\widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$:

$$(13) \quad \mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \simeq \widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{U}} \phi_j \boldsymbol{\Sigma}_{\mathbf{U}}^{-1} = \widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{r}_j$$

$$(14) \quad \mathbf{U}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \simeq \widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{U}} \phi'_j \boldsymbol{\Sigma}'(\overline{\mathbf{j}}, \overline{\mathbf{k}})^{-1} = \widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{r}'_j$$

The weights $\widehat{\boldsymbol{\Sigma}}_{\mathbf{U}}$ of the new interpolation basis $\widehat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ are stored, as they are needed in potential future recompressions. A new antinterpolation operator $\widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ for $\Omega_j^{(l)}$ is obtained in a similar way:

$$(15) \quad \left[\mathbf{V}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}_{\mathbf{V}} \mid \mathbf{V}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \boldsymbol{\Sigma}'(\overline{\mathbf{k}}, \overline{\mathbf{j}}) \right] \simeq \widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{V}} \boldsymbol{\Psi}_j$$

with $\boldsymbol{\Psi}_j = [\psi_j \mid \psi'_j]$, such that $\mathbf{V}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ and $\mathbf{V}'(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ can be written as:

$$(16) \quad \mathbf{V}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \simeq \widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{V}} \psi_j \boldsymbol{\Sigma}_{\mathbf{V}}^{-1} = \widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{t}_j$$

$$(17) \quad \mathbf{V}'(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \simeq \widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \widehat{\boldsymbol{\Sigma}}_{\mathbf{V}} \psi'_j \boldsymbol{\Sigma}'(\overline{\mathbf{k}}, \overline{\mathbf{j}})^{-1} = \widehat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{t}'_j$$

The same strategy as outlined in Eq. (12) to (17) is applicable for obtaining new operators $\widehat{\mathbf{U}}(\underline{\mathbf{k}}, \underline{\mathbf{k}})$ and $\widehat{\mathbf{V}}(\underline{\mathbf{k}}, \underline{\mathbf{k}})$ for $\Omega_k^{(l)}$.

The crucial benefit of compressing the dense fill-ins and updating the inter- and antinterpolation operators of the corresponding clusters is that the former can now be accounted for through the operators $\mathbf{K}(\overline{\mathbf{k}}, \overline{\mathbf{j}})$ and $\mathbf{K}(\overline{\mathbf{j}}, \overline{\mathbf{k}})$ between $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$. For example, the fill-in $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$ is redirected through the $\mathbf{K}(\overline{\mathbf{k}}, \overline{\mathbf{j}})$ -operator, which becomes:

$$(18) \quad \mathbf{K}(\overline{\mathbf{k}}, \overline{\mathbf{j}}) \leftarrow \mathbf{r}_k \mathbf{K}(\overline{\mathbf{k}}, \overline{\mathbf{j}}) \mathbf{t}_j^T + \mathbf{r}'_k \mathbf{K}'(\overline{\mathbf{k}}, \overline{\mathbf{j}}) \mathbf{t}'_j^T$$

[§]Initial weights $\boldsymbol{\Sigma}_{\mathbf{U}}$ for $\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ (i.e., upon initialization of the IFMM operators) can be obtained rigorously by applying an SVD to $[\dots \mid \mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}}) \mathbf{K}(\overline{\mathbf{j}}, \overline{\mathbf{q}}) \mathbf{V}^T(\underline{\mathbf{q}}, \underline{\mathbf{q}}) \mid \dots]$ ($\forall \Omega_q^{(l)} \in \mathcal{I}_j^{(l)}$, with $\mathcal{I}_j^{(l)}$ the interaction list of $\Omega_j^{(l)}$), or can be roughly estimated by sampling some matrix entries. The latter approach is less expensive in terms of computational cost.

[¶]The proof of this result, which is elementary, is omitted.

Updating the inter- and anteroplation operators also implies that all the other edges in the graph associated with $\Omega_j^{(l)}$ need to be adjusted. More specifically, the remaining operators $\mathbf{K}(\bar{\mathbf{q}}, \bar{\mathbf{j}})$ and $\mathbf{K}(\bar{\mathbf{j}}, \bar{\mathbf{q}})$ (with $\Omega_q^{(l)} \in \mathcal{N}_j^{(l)} \cup \mathcal{I}_j^{(l)} \setminus \Omega_k^{(l)}$) are replaced by:

$$(19) \quad \mathbf{K}(\bar{\mathbf{q}}, \bar{\mathbf{j}}) \leftarrow \mathbf{K}(\bar{\mathbf{q}}, \bar{\mathbf{j}}) \mathbf{t}_j^T$$

$$(20) \quad \mathbf{K}(\bar{\mathbf{j}}, \bar{\mathbf{q}}) \leftarrow \mathbf{r}_j \mathbf{K}(\bar{\mathbf{j}}, \bar{\mathbf{q}})$$

where $\mathcal{N}_j^{(l)}$ and $\mathcal{I}_j^{(l)}$ represent the neighbor list (including itself) and interaction list (i.e., well-separated clusters that are children of its parent's neighbors) of $\Omega_j^{(l)}$, respectively. Furthermore, the inter- and anteroplation operators of the parent cluster of $\Omega_j^{(l)}$ are updated as follows:

$$(21) \quad \mathbf{U}(\bar{\mathbf{j}}, \mathbf{j}^\dagger) \leftarrow \mathbf{r}_j \mathbf{U}(\bar{\mathbf{j}}, \mathbf{j}^\dagger)$$

$$(22) \quad \mathbf{V}(\mathbf{j}^\dagger, \bar{\mathbf{j}}) \leftarrow \mathbf{t}_j \mathbf{V}(\mathbf{j}^\dagger, \bar{\mathbf{j}})$$

Similar updates as in Eq. (18)–(22) are required for all edges associated with the cluster $\Omega_k^{(l)}$. Edges that need to be updated are highlighted in yellow in Figure 11c.

The technique outlined above leads to a significant reduction of the amount of fill-in and ensures that the sparsity pattern of the extended sparse matrix is preserved throughout the elimination phase. This leads to an algorithm with linear complexity $\mathcal{O}(N)$. This assumes that the rank remains bounded independent of the matrix size. This is expected to be the case, as only interactions between well-separated clusters are replaced by low-rank approximations.

This strategy of compressing and redirecting dense fill-ins is very similar to the initial sparsification approach introduced in Subsection 2.1, as low-rank interactions are represented at the parent level. Apart from fill-ins between the $\tilde{\mathbf{x}}$ -variables (e.g., between $\tilde{\mathbf{x}}_j^{(l)}$ and $\tilde{\mathbf{x}}_k^{(l)}$), fill-ins between the variables $\tilde{\mathbf{x}}^{(l)}$ and $\mathbf{y}^{(l)}$ are created as well throughout the elimination phase. If these occur between well-separated clusters (for example $\mathbf{E}'(\bar{\mathbf{h}}, \mathbf{j})$, $\mathbf{E}'(\mathbf{j}, \bar{\mathbf{h}})$ and $\mathbf{E}'(\bar{\mathbf{h}}, \underline{\mathbf{k}})$, $\mathbf{E}'(\underline{\mathbf{k}}, \bar{\mathbf{h}})$ in Figure 11b), the fill-ins are compressed and redirected through other operators at the parent level.

Once all unknowns $\tilde{\mathbf{x}}^{(l)}$ and variables $\mathbf{z}^{(l)}$ at level l have been eliminated, the next level is considered. From a conceptual point of view, variables $\mathbf{y}^{(l)}$ can be interpreted as unknowns $\tilde{\mathbf{x}}^{(l-1)}$ of the parent level $l-1$. More specifically, a set of variables $\mathbf{y}_i^{(l)}$ is combined to form $\tilde{\mathbf{x}}_p^{(l-1)}$, where the clusters $\Omega_i^{(l)}$ are children of the parent cluster $\Omega_p^{(l-1)}$:

$$(23) \quad \tilde{\mathbf{x}}_p^{(l-1)} = \left\{ \dots \mathbf{y}_i^{(l)} \dots \right\}^T \quad \forall \Omega_i^{(l)} \in \mathcal{C}_p^{(l-1)}$$

or, equivalently, the nodes $\bar{\mathbf{i}}^{(l)}$ are joined to form $\underline{\mathbf{p}}^{(l-1)}$:

$$(24) \quad \underline{\mathbf{p}}^{(l-1)} = \bigcup_{\Omega_i^{(l)} \in \mathcal{C}_p^{(l-1)}} \bar{\mathbf{i}}^{(l)}$$

$\mathcal{C}_p^{(l-1)}$ in Eq. (23) and (24) denotes the children list of $\Omega_p^{(l-1)}$. This merging procedure is illustrated in Figure 12. Furthermore, the corresponding operators are combined accordingly to form the inter- and anteroplation operators $\mathbf{U}(\underline{\mathbf{p}}^{(l-1)}, \mathbf{p}^{(l-1)})$ and

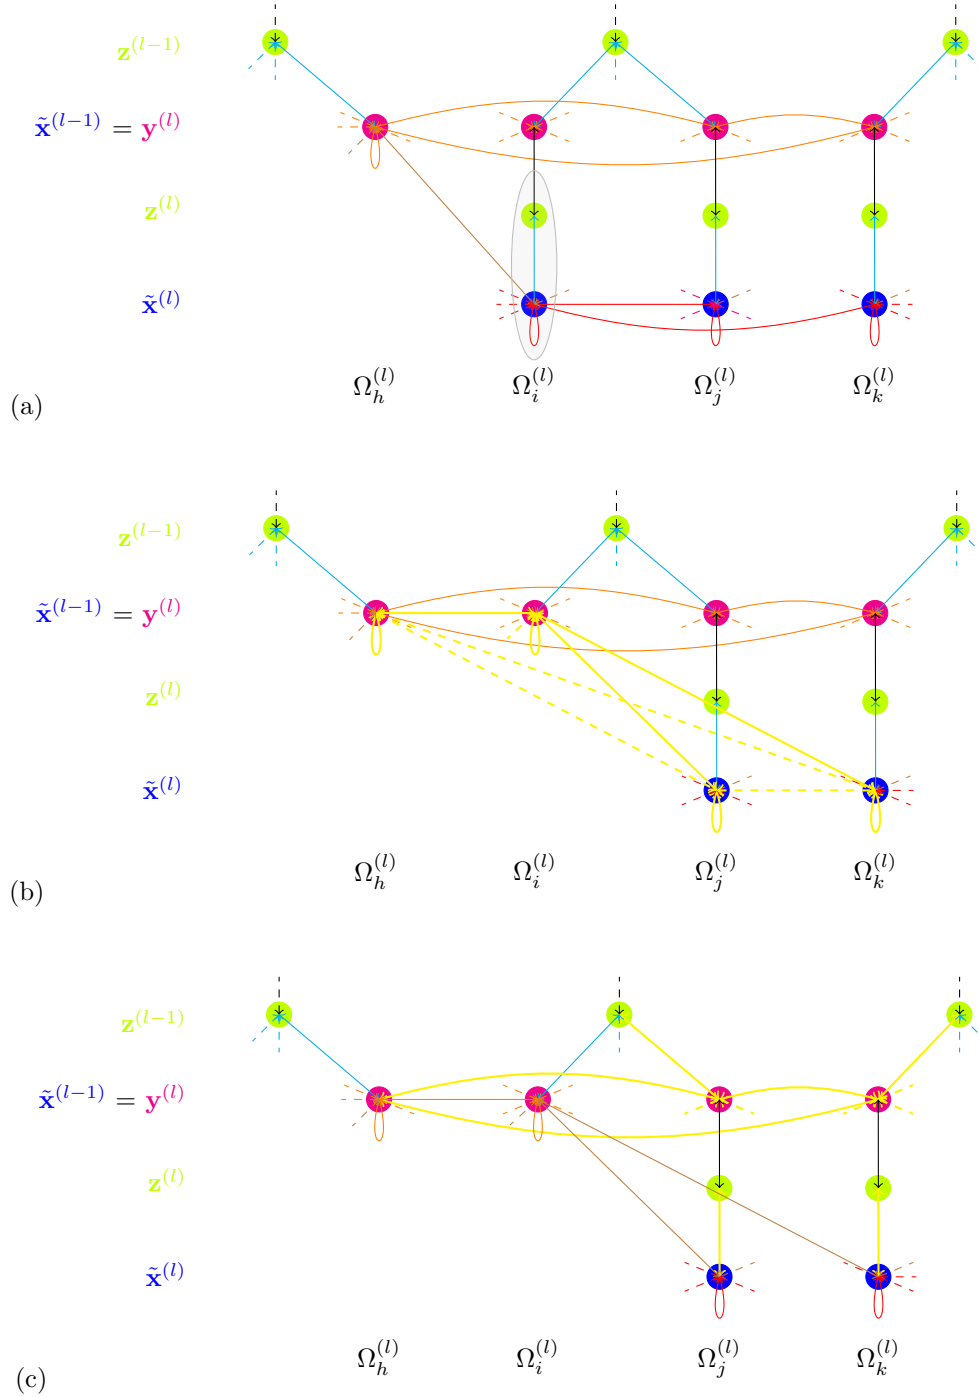


FIG. 11. Subgraph (a) before and (b) after the elimination of the variables $\tilde{\mathbf{x}}_i^{(l)}$ and $\mathbf{z}_i^{(l)}$. New and modified edges are highlighted in yellow in (b). Operators associated with edges between well-separated clusters (indicated by dashed yellow lines in (b)) are compressed and redirected, resulting in the subgraph shown in (c). The edges that are updated after compression are highlighted in yellow in (c). This strategy of compressing and redirecting dense fill-ins is similar to the initial sparsification approach introduced in Subsection 2.1, as low-rank interactions are represented at the parent level.

$\mathbf{V}(\underline{\mathbf{p}}^{(l-1)}, \underline{\mathbf{p}}^{(l-1)})$ of $\Omega_{\mathbf{p}}^{(l-1)}$:

$$(25) \quad \mathbf{U}(\underline{\mathbf{p}}^{(l-1)}, \underline{\mathbf{p}}^{(l-1)}) = \begin{bmatrix} \vdots \\ \mathbf{U}(\bar{\mathbf{i}}^{(l)}, \mathbf{i}^\dagger^{(l)}) \\ \vdots \end{bmatrix} \quad \forall \Omega_i^{(l)} \in \mathcal{C}_{\mathbf{p}}^{(l-1)}$$

$$(26) \quad \mathbf{V}(\underline{\mathbf{p}}^{(l-1)}, \underline{\mathbf{p}}^{(l-1)}) = \begin{bmatrix} \vdots \\ \mathbf{V}(\mathbf{i}^\dagger^{(l)}, \bar{\mathbf{i}}^{(l)}) \\ \vdots \end{bmatrix} \quad \forall \Omega_i^{(l)} \in \mathcal{C}_{\mathbf{p}}^{(l-1)}$$

A similar operation is performed to merge the interactions between variables $\mathbf{y}^{(l)}$ into interactions between unknowns $\tilde{\mathbf{x}}^{(l-1)}$. For example, the interaction between $\tilde{\mathbf{x}}_{\mathbf{p}}^{(l-1)}$ and $\tilde{\mathbf{x}}_{\mathbf{q}}^{(l-1)}$ of clusters $\Omega_{\mathbf{p}}^{(l-1)}$ and $\Omega_{\mathbf{q}}^{(l-1)}$ is assembled as follows:

$$(27) \quad \mathbf{K}(\underline{\mathbf{p}}^{(l-1)}, \underline{\mathbf{q}}^{(l-1)}) = \begin{bmatrix} \vdots & & \\ \dots & \mathbf{K}(\bar{\mathbf{i}}^{(l)}, \bar{\mathbf{j}}^{(l)}) & \dots \\ \vdots & & \end{bmatrix} \quad \forall \Omega_i^{(l)} \in \mathcal{C}_{\mathbf{p}}^{(l-1)}, \forall \Omega_j^{(l)} \in \mathcal{C}_{\mathbf{q}}^{(l-1)}$$

The resulting graph has exactly the same structure as the initial one, except for the fact that it consists of one level less. The elimination and merging procedure are repeated consecutively for each level $l \geq 2$ until only a small dense system for $\mathbf{y}^{(2)}$ (or, equivalently, $\tilde{\mathbf{x}}^{(1)}$) remains to be solved. Once this is done, the elimination phase is terminated and all the other variables can be retrieved through substitution, marching down the graph in the reverse order. In the end, the solution \mathbf{x} is obtained.

2.3. The algorithm. Based on the key ideas introduced in Subsection 2.1 and 2.2, the general IFMM algorithm can now be presented in a formal way. In the following, it is assumed that the \mathcal{H}^2 -structure of the matrix \mathbf{A} is based on a three-dimensional octree decomposition, implying that every cluster $\Omega_i^{(l)}$ can have up to $3^3 = 27$ clusters in its neighbor list $\mathcal{N}_i^{(l)}$ (including itself) and up to $6^3 - 3^3 = 189$ clusters in its interaction list $\mathcal{I}_i^{(l)}$.

Algorithm 1 presents the general outline of the IFMM. It consists of initializing the IFMM-operators, performing an upward pass to eliminate each level l (and forming the appropriate operators of the parent level $l-1$), solving for $\mathbf{y}^{(2)} (= \tilde{\mathbf{x}}^{(1)})$ at the top level, and substituting through a downward pass to finally retrieve the solution \mathbf{x} . The elimination procedure at each level l is detailed in Algorithm 2. Eliminating a cluster $\Omega_i^{(l)}$ leads to new interactions between its neighbors $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$ (with $\Omega_j^{(l)}, \Omega_k^{(l)} \in \mathcal{N}_i^{(l)}$): $\mathbf{E}'(\bar{\mathbf{k}}, \bar{\mathbf{j}})$ and $\mathbf{E}'(\bar{\mathbf{j}}, \bar{\mathbf{k}})$ if both $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$ have already been eliminated, $\mathbf{E}'(\underline{\mathbf{k}}, \bar{\mathbf{j}})$ and $\mathbf{E}'(\bar{\mathbf{j}}, \underline{\mathbf{k}})$ if $\Omega_j^{(l)}$ has been eliminated but $\Omega_k^{(l)}$ has not (and vice versa), and $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$ and $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ if neither $\Omega_j^{(l)}$ nor $\Omega_k^{(l)}$ has been eliminated yet. If $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$ are well-separated, these fill-ins are compressed into a low-rank approximation and subsequently redirected through other operators, as discussed in Subsection 2.2. This redirection procedure is described in Algorithm 3. If this is not the case, the additional interactions are added to the existing ones. The substitution




FIG. 12. Subgraph (a) before and (b) after merging the variables $\mathbf{y}^{(l)}$ of level l into unknowns $\tilde{\mathbf{x}}^{(l-1)}$ of the parent level $l-1$. The associated operators are combined accordingly.

phase proceeds level per level in the reverse elimination order and allows determining $\mathbf{z}^{(l)}$ and $\tilde{\mathbf{x}}^{(l)}$ based on $\mathbf{y}^{(l)}$. If l is not the leaf level, the variables $\tilde{\mathbf{x}}^{(l)}$ are nothing else than the variables $\mathbf{y}^{(l+1)}$ of the child level, cf., Eq. (23). At the leaf level, $\tilde{\mathbf{x}}^{(L)}$ corresponds to the desired solution \mathbf{x} . As the elimination phase is decoupled from the substitution phase (as in a standard direct solve), the algorithm is well suited to handle multiple right hand sides \mathbf{b} .

2.4. Some important remarks. Several techniques can be employed for obtaining a suitable low-rank approximation for the dense fill-ins \mathbf{E}' arising throughout the elimination phase (cf., Eq. (10) and (11)). The best possible approximation of rank k of a matrix (minimizing the error norm) is given by its k^{th} partial singular value decomposition (SVD) [20]. Unfortunately, the calculation of a complete SVD requires $\mathcal{O}(\min(mn^2, m^2n))$ arithmetical operations for a $(m \times n)$ matrix, making it rather unattractive from a computational point of view. The randomized SVD (rSVD) algorithm [30] provides a very efficient alternative for computing (approximations of) the first few singular values and vectors. This method exploits the power of randomization to identify a subspace that captures most of the action of a matrix [30]. It involves generating random matrices, performing matrix-matrix products, applying orthonormalizations, and computing the SVD of a much smaller matrix. Although it uses random matrices, the rSVD algorithm gives very accurate results with a high probability, provided that the singular values of the matrix decay sufficiently fast. Rather than specifying a fixed rank for the low-rank approximation of a dense fill-in, the rank can also be determined by prescribing a relative accuracy ε . The relative importance of a fill-in is hence evaluated by weighting its singular values (obtained, for example, with the rSVD algorithm) with respect to some reference value. In this paper, it is proposed to employ the largest singular value $\sigma_0(\mathbf{E})$ of the extended sparse

Algorithm 1 The inverse fast multipole method: overall algorithm

```

1: Construct an octree decomposition of the domain  $\Omega$ 
2: for  $l = L, L - 1, \dots, 2$  do
3:   Initialize the IFMM operators at level  $l$ 
4: end for
5: for  $l = L, L - 1, \dots, 2$  do
6:   Eliminate level  $l$  (Algorithm 2) and update the right hand side
7:   if  $l > 2$  then
8:     Form the appropriate operators of the parent level  $l - 1$ 
9:   end if
10: end for
11: Solve for  $\mathbf{y}^{(2)} (= \tilde{\mathbf{x}}^{(1)})$ 
12: for  $l = 2, \dots, L - 1, L$  do
13:   Substitute at level  $l$  for  $\mathbf{z}^{(l)}$  and  $\tilde{\mathbf{x}}^{(l)}$ 
14:   if  $l < L$  then
15:     Form  $\mathbf{y}^{(l+1)}$  based on  $\tilde{\mathbf{x}}^{(l)}$  (Eq. (23))
16:   else
17:     Set  $\mathbf{x} = \tilde{\mathbf{x}}^{(L)}$ 
18:   end if
19: end for

```

matrix \mathbf{E} for this purpose. The retained rank of a fill-in \mathbf{E}' is thus given by^{||}:

$$(28) \quad k(\varepsilon) = \min \{k \in \mathbb{N} : \sigma_{k+1}(\mathbf{E}') \leq \varepsilon \sigma_0(\mathbf{E})\}$$

$\sigma_0(\mathbf{E})$ can easily be estimated by application of the rSVD algorithm to the extended sparse matrix \mathbf{E} .

A low-rank recompression is also required for calculating new inter- and anteropola-tion operators of a cluster after compression of the fill-ins (cf., Eq. (12) and (15)). Once again, computing the SVD will provide the optimal recompression, but at a relatively high computational cost. Applying the rSVD algorithm will not be very efficient either in this case, as the rank of the new operators $\hat{\mathbf{U}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ and $\hat{\mathbf{V}}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ is at least as large as the rank of the original operators $\mathbf{U}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ and $\mathbf{V}(\underline{\mathbf{j}}, \underline{\mathbf{j}})$ (as the latter consist of orthogonal columns). This recompression hence differs from the compression of the fill-ins, where only the first few singular values are (expected to be) of importance. The fully pivoted adaptive cross approximation (ACA) algorithm [5, 42]**, followed by an SVD recompression (i.e., two skinny QR decompositions and an SVD performed on the product of the resulting triangular matrices), provides a more efficient way for performing the recompression and is therefore favored.

A final remark is related to the weighting of the basis vectors when computing new inter- and anteropola-tion operators, as presented in Eq. (12) and (15). The advantage of using rSVD and ACA+SVD is that the obtained singular values straightforwardly provide the weights for each vector in the basis. If the low-rank approximations in Eq. (10), (11), (12) and (15) do not correspond to singular value decompositions (e.g., when using a rank-revealing LU decomposition), the weighting needs to be chosen in a different manner. The pivots of the LU factorization could, for example, be employed for this purpose. Although this approach is less accurate than using the singular

^{||} k is also bounded by the size of \mathbf{E}' .

**This approach is the same as an LU factorization with full pivoting.

Algorithm 2 The inverse fast multipole method: elimination of level l . In order to simplify the notations as much as possible, we made the following choices. The letter \mathbf{E} denotes the sparse matrix on which we are doing the Gaussian elimination. The fill-in introduced is denoted by \mathbf{E}' . This fill-in is progressively compressed (if well-separated) and folded into existing edges (non-zero entries) of \mathbf{E} . This compression process is done using matrices denoted by \mathbf{U}' , \mathbf{V}' , and \mathbf{K}' . \mathbf{U}' is an interpolation operator ($\underline{\mathbf{i}}$ to $\underline{\mathbf{i}}$), while \mathbf{V}' is an antinterpolation operator ($\underline{\mathbf{i}}$ to \mathbf{i}). The entries denoted by \mathbf{K}' correspond to edges between nodes at the same level.

```

1: for all clusters  $\Omega_i^{(l)}$  at level  $l$  do
2:   Eliminate  $\underline{\mathbf{i}}$  and  $\mathbf{i}$  of  $\Omega_i^{(l)}$ 
3:   for all unique pairs  $\{\Omega_j^{(l)}, \Omega_k^{(l)}\}$  with  $\Omega_j^{(l)}, \Omega_k^{(l)} \in \mathcal{N}_i^{(l)}$  do
4:     if  $\underline{\mathbf{j}}$  and  $\mathbf{j}$  of  $\Omega_j^{(l)}$  have been eliminated then
5:       if  $\underline{\mathbf{k}}$  and  $\mathbf{k}$  of  $\Omega_k^{(l)}$  have been eliminated then
6:         Compute the fill-ins  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
7:         Add the fill-ins:  $\mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \leftarrow \mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) + \mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \leftarrow \mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) + \mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
8:       else
9:         Compute the fill-ins  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
10:      if  $\Omega_j^{(l)} \notin \mathcal{N}_k^{(l)}$  then
11:        Compress  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \simeq \mathbf{U}'(\underline{\mathbf{k}}, \mathbf{k}) \mathbf{K}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \simeq \mathbf{K}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \mathbf{V}'^T(\mathbf{k}, \underline{\mathbf{k}})$ 
12:        Update the edges associated with  $\Omega_k^{(l)}$  (Algorithm 3)
13:      else
14:        Add the fill-ins:  $\mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \leftarrow \mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) + \mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \leftarrow \mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) + \mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
15:      end if
16:    end if
17:  else
18:    if  $\underline{\mathbf{k}}$  and  $\mathbf{k}$  of  $\Omega_k^{(l)}$  have been eliminated then
19:      Compute the fill-ins:  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
20:    if  $\Omega_j^{(l)} \notin \mathcal{N}_k^{(l)}$  then
21:      Compress  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \simeq \mathbf{K}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \mathbf{V}'^T(\mathbf{j}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \simeq \mathbf{U}'(\underline{\mathbf{j}}, \mathbf{j}) \mathbf{K}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
22:      Update the edges associated with  $\Omega_j^{(l)}$  (Algorithm 3)
23:    else
24:      Add the fill-ins:  $\mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \leftarrow \mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) + \mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \leftarrow \mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) + \mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
25:    end if
26:  else
27:    Compute the fill-ins  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
28:  if  $\Omega_j^{(l)} \notin \mathcal{N}_k^{(l)}$  then
29:    Compress  $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \simeq \mathbf{U}'(\underline{\mathbf{k}}, \mathbf{k}) \mathbf{K}'(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \mathbf{V}'^T(\mathbf{j}, \underline{\mathbf{j}})$  and  $\mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \simeq \mathbf{U}'(\underline{\mathbf{j}}, \mathbf{j}) \mathbf{K}'(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \mathbf{V}'^T(\mathbf{k}, \underline{\mathbf{k}})$ 
30:    Update the edges associated with  $\Omega_j^{(l)}$  (Algorithm 3)
31:    Update the edges associated with  $\Omega_k^{(l)}$  (Algorithm 3)
32:  else
33:    Add the fill-ins:  $\mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) \leftarrow \mathbf{E}(\underline{\mathbf{k}}, \underline{\mathbf{j}}) + \mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$  and  $\mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) \leftarrow \mathbf{E}(\underline{\mathbf{j}}, \underline{\mathbf{k}}) + \mathbf{E}'(\underline{\mathbf{j}}, \underline{\mathbf{k}})$ 
34:  end if
35:  end if
36:  end if
37: end for
38: end for

```

Algorithm 3 The inverse fast multipole method: updating of the edges associated with $\Omega_j^{(l)}$ after compressing the interactions between $\Omega_j^{(l)}$ and $\Omega_k^{(l)}$

```

1:  $\mathbf{U}(\underline{j}, j) \leftarrow \widehat{\mathbf{U}}(\underline{j}, j), \Sigma_{\mathbf{U}} \leftarrow \widehat{\Sigma}_{\mathbf{U}}$ 
2:  $\mathbf{V}(j, \underline{j}) \leftarrow \widehat{\mathbf{V}}(j, \underline{j}), \Sigma_{\mathbf{V}} \leftarrow \widehat{\Sigma}_{\mathbf{V}}$ 
3:  $\mathbf{U}(\overline{j}, j^\dagger) \leftarrow \mathbf{r}_j \mathbf{U}(\overline{j}, j^\dagger)$ 
4:  $\mathbf{V}(j^\dagger, \overline{j}) \leftarrow \mathbf{t}_j \mathbf{V}(j^\dagger, \overline{j})$ 
5: for all  $\Omega_q^{(l)} \in \mathcal{N}_j^{(l)} \cup \mathcal{I}_j^{(l)} \setminus \Omega_k^{(l)}$  do
6:    $\mathbf{K}(\overline{q}, \overline{j}) \leftarrow \mathbf{K}(\overline{q}, \overline{j}) \mathbf{t}_j^T$ 
7:    $\mathbf{K}(\overline{j}, \overline{q}) \leftarrow \mathbf{r}_j \mathbf{K}(\overline{j}, \overline{q})$ 
8: end for
9: if  $\underline{k}$  and  $k$  of  $\Omega_k^{(l)}$  have been eliminated then
10:   $\mathbf{K}(\overline{k}, \overline{j}) \leftarrow \mathbf{K}(\overline{k}, \overline{j}) \mathbf{t}_j^T + \mathbf{K}'(\overline{k}, \overline{j}) \mathbf{t}_j'^T$ 
11:   $\mathbf{K}(\overline{j}, \overline{k}) \leftarrow \mathbf{r}_j \mathbf{K}(\overline{j}, \overline{k}) + \mathbf{r}_j' \mathbf{K}'(\overline{j}, \overline{k})$ 
12: else
13:   $\mathbf{K}(\overline{k}, \overline{j}) \leftarrow \mathbf{r}_k \mathbf{K}(\overline{k}, \overline{j}) \mathbf{t}_j^T + \mathbf{r}_k' \mathbf{K}'(\overline{k}, \overline{j}) \mathbf{t}_j'^T$ 
14: end if

```

values, it can still provide a reasonable measure of the scale of each basis vector when performing the recompression.

2.5. Computational cost. We will derive the computational cost in the case of a uniform tree decomposition of the matrix (as opposed to adaptive), although the proof extends to general adaptive trees.

Let's assume that we have a matrix \mathbf{A} in \mathcal{H}^2 -format. Given a tolerance ε , we assume that there is an upper bound r on the rank of all compressed blocks such that $r \in \mathcal{O}(\log 1/\varepsilon)$. One can prove that for the original matrix \mathbf{A} , in the case of smooth kernels (e.g., Green's functions), the rank r of the compressed blocks does satisfy this property [18]. However, in our case, as we proceed through the LU factorization, recompression is required. We have not proved that the rank does not increase in an unbounded fashion (for a fixed ε , as N varies) as a result of this process. However, we did observe experimentally that the rank is in fact bounded. See in particular Figure 16.

Under these assumptions, the cost of operating on a node at a given level is $\mathcal{O}(r^3)$. The cost of a leaf node elimination is $\mathcal{O}(N_0^3)$, where N_0 is an upper bound on the size of all leaf clusters. Choose $N_0 \in \Omega(r)$. This requirement sets the number of levels in the tree, and the number of leaf nodes. In that case, the cost of the elimination of all leaf nodes is $\mathcal{O}(Nr^2)$. The cost of compression of the far-away fill-in blocks, and computation of the updated parent-level matrix blocks is also $\mathcal{O}(Nr^2)$. Since at the next level, the number of clusters is divided at least by a fixed factor (e.g., 2 for a binary tree or 8 for an octree), the total computational cost of the method is $\mathcal{O}(Nr^2)$ or $\mathcal{O}(N \log(1/\varepsilon)^2)$.

3. Numerical examples. The IFMM algorithm has been implemented in C++, using the Eigen-library [22] for performing dense linear algebra operations. In the following subsections, numerical examples are considered to validate the implementation and to demonstrate the efficiency of the methodology. All computations have been performed on Intel[®] Xeon[®] E5-2650 v2 (2.60 GHz) CPUs. In each of these examples, the initial low-rank approximations in the \mathcal{H}^2 -matrix are created using interpolation

based on Chebyshev polynomials [18], followed by an additional SVD to reduce the rank. A uniform octree decomposition of the domain is employed; the number of levels is adjusted to ensure that the leaf clusters contain approximately 100 points. This value was found to provide a reasonable trade-off between the time spent at the leaf level and the time spent at higher levels in the tree.^{††} A relative accuracy $\varepsilon = 10^{-3}$ is used, both for compressing the fill-ins (with rSVD) and for updating the inter- and antepolation operators (with ACA).

3.1. The IFMM as a direct solver. Consider the linear system of equations $\mathbf{Ax} = \mathbf{b}$ with matrix entries $a_{ij} = K(\|\mathbf{r}_i - \mathbf{r}_j\|)$ depending on the kernel $K(r)$:

$$(29) \quad K(r) = \begin{cases} 1 & \text{if } r = 0, \\ \frac{r}{d} & \text{if } 0 < r < d, \\ \frac{d}{r} & \text{if } r \geq d. \end{cases}$$

with \mathbf{r}_i and \mathbf{r}_j indicating 3D position vectors and $r = \|\mathbf{r}_i - \mathbf{r}_j\|$. Figure 13 shows the kernel function $K(r)$ for three values of the parameter d in Eq. (29). This parameter strongly affects the condition number $\kappa(\mathbf{A})$ of the matrix \mathbf{A} : the larger d , the larger the condition number (for a fixed matrix size N). The condition number also increases with N for a fixed value of d , as illustrated in Figure 14. In the following, a known vector \mathbf{x} is chosen and the matrix-vector product \mathbf{Ax} is evaluated for obtaining \mathbf{b} . The IFMM is subsequently used as a direct solver, with the aforementioned vector \mathbf{b} as right hand side. The retrieved solution is denoted as $\hat{\mathbf{x}}$.




FIG. 13. Kernel function $K(r)$ for three values of the parameter d .

In a first example, N points \mathbf{r} are randomly distributed on the unit sphere. The parameter d is scaled according to $d = 10^{-3} (N/1000)^{-1/2}$ in order to ensure that the condition number of \mathbf{A} remains nearly constant for all N (see Figure 14a). Figure 15a shows the total time taken by the IFMM solver as a function of the number of points N , with the number of Chebyshev nodes n (in each dimension) varying from 1 to 4 (the total number of Chebyshev interpolation nodes employed for constructing the initial low-rank approximations hence equals n^3). These curves demonstrate that, as expected, a linear complexity $\mathcal{O}(N)$ is achieved. The computation time increases for larger values of n . The relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ for each of these cases is depicted in Figure 15b. The figure clearly shows a decrease of the error as n is increased. The errors are reasonably small for all problem sizes N .

^{††}Less points per leaf cluster makes the elimination of the leaves faster, but increases the number of levels and thus the time spent in the tree, and vice versa.




FIG. 14. Condition number $\kappa(\mathbf{A})$ as a function of the matrix size N , for points randomly distributed (a) on the unit sphere and (b) in the cube $[-1, 1]^3$.

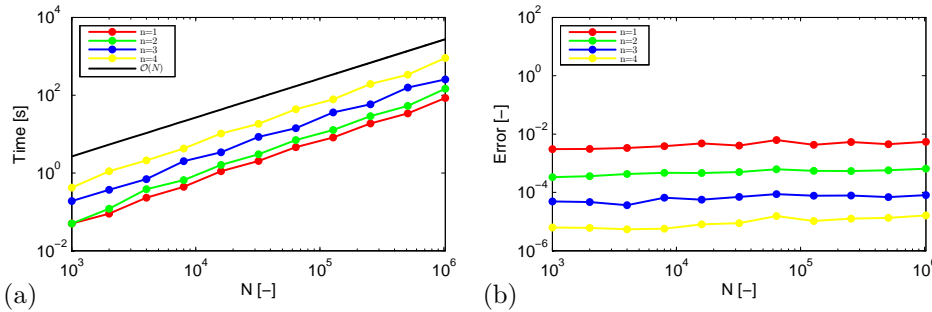


FIG. 15. (a) Total CPU time and (b) the relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ as a function of the matrix size N , for points randomly distributed on the unit sphere.

A crucial feature for achieving the linear complexity demonstrated in Figure 15a is the fact that the rank of the matrix blocks corresponding to far field interactions remains (almost) constant if the matrix size N is increased (due to the strong admissibility criterion employed for \mathcal{H}^2 -matrices). This is in contrast with HODLR and HSS matrices, where the rank grows like $\mathcal{O}(\sqrt{N})$ for 2D problems [1]. Figure 16 shows the maximum and average rank of the aforementioned matrix blocks as a function of N for the case of $n = 4$, confirming the previous statement (especially for $N > 10^4$). The kinks in the curve of the average rank are caused by varying the number of levels in the octree.

Next, a more general example of N points \mathbf{r} randomly distributed in the cube $[-1, 1]^3$ is investigated. The parameter d varies in this case as $d = 10^{-3} (N/1000)^{-1/3}$ to keep the condition number of \mathbf{A} nearly constant for all N (see Figure 14b). Figure 17a confirms that a linear scaling of the computation time is also obtained in this case (although only from $N \sim 10^4$ upwards), while Figure 17b demonstrates the convergence of the relative error as the number of Chebyshev nodes is increased. The computation times are slightly higher than in the case of points distributed on a sphere (Figure 15a). This is due to the fact that a uniform octree decomposition of a sphere leads to many empty clusters, implying that most of the clusters will have less than 27 neighbors, while it is likely that almost all clusters will possess the maximum number of neighbors in case the points are randomly distributed in a cube. The reader is referred to Appendix A for a detailed breakdown of the computational cost for this example.




FIG. 16. Maximum (black line) and average (grey line) rank of the matrix blocks corresponding to far field interactions as a function of the matrix size N , for points randomly distributed on the unit sphere (with $n = 4$).




FIG. 17. (a) Total CPU time and (b) the relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ as a function of the matrix size N , for points randomly distributed in the cube $[-1, 1]^3$.

Depending on the application of interest, the errors shown in Figure 15b and 17b might already be sufficiently small. The convergence of the IFMM is investigated in more detail in Figure 18, which shows the relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ as a function of the number of Chebyshev nodes for a fixed matrix size $N = 8 \times 10^3$. An increase of n leads to more accurate low-rank approximations of the far field interactions and consequently a smaller overall error in $\hat{\mathbf{x}}$. This figure demonstrates that the method is capable of serving as a highly accurate direct solver. The number of Chebyshev nodes needs to be increased significantly if 10 or more digits of accuracy are desired, however, which leads to a strong increase of the computational cost (even if the linear complexity is preserved). In most cases (depending on the number of right-hand sides), it is more efficient to apply a low accuracy IFMM solver as a preconditioner in an iterative method rather than using a highly accurate direct IFMM solver, as will be illustrated in the next subsection.

3.2. The IFMM as a preconditioner in an iterative scheme. In this subsection, the linear system $\mathbf{Ax} = \mathbf{b}$ is solved iteratively using a non-restarted GMRES-algorithm [47, 46] with a tolerance $\varepsilon = 10^{-14}$ for the relative residual $\|\mathbf{b} - \mathbf{Ax}\|_2 / \|\mathbf{b}\|_2$ and a maximum of 500 iterations. The same kernel $K(r)$ as in Eq. (29) is used. The black-box FMM [18] is employed for evaluating the matrix-vector products within GMRES, while a low accuracy IFMM solver is applied as preconditioner in each iterative step.

The first example involves $N = 10^5$ points randomly distributed in the cube




FIG. 18. The relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ as a function of the number of Chebyshev nodes n , for $N = 8000$ points randomly distributed in the cube $[-1, 1]^3$.

$[-1, 1]^3$, with $d = 10^{-3}$ (in Eq. (29)). Figure 19a shows the relative residual as a function of the number of GMRES-iterations for several preconditioning approaches. Without preconditioning, convergence is rather slow and GMRES seems to stall at certain steps of the iteration. A total of 187 iterations are required to achieve convergence. The application of the IFMM as preconditioner results in a much faster convergence. Even with a rank-1 approximation of the far field interactions (i.e., $n = 1$), a significant reduction of the number of iterations is obtained. An increase of the number of Chebyshev nodes leads to a more accurate preconditioner and hence a further reduction of the number of iterations; this comes at the price of a higher cost to construct the preconditioner, however. The computation times for the considered preconditioning strategies are presented in Figure 19b, where a decomposition is made into the time required to construct the preconditioner (i.e., the elimination phase in the IFMM) and the actual iteration time (i.e., the substitution phase in the IFMM, together with all GMRES-related computations). For the example under concern, the use of three Chebyshev nodes per dimension leads to the smallest overall computation time. The results are summarized in Table 1.

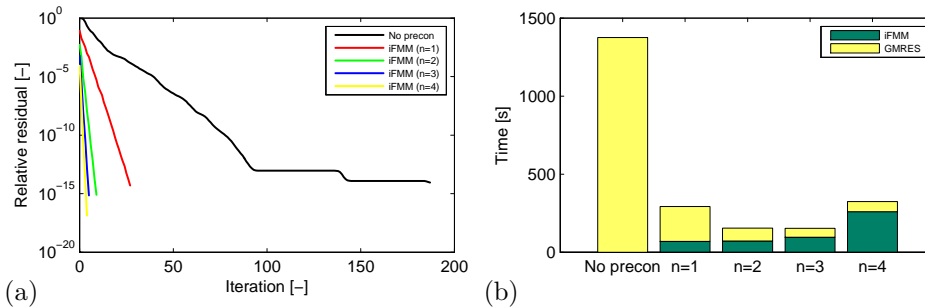


FIG. 19. (a) Relative residual $\|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|_2 / \|\mathbf{b}\|_2$ as a function of the iteration step and (b) total CPU time for $d = 10^{-3}$.

In the second example, the parameter d in Eq. (29) is increased from 10^{-3} to 10^{-2} , resulting in a more ill-conditioned matrix \mathbf{A} . Figure 20a demonstrates that the unpreconditioned GMRES-algorithm is now unable to converge to the specified residual within 500 iterations. Application of the IFMM as preconditioner ensures convergence, although a rank-1 approximation of the far field is not very efficient in this case, as the number of iterations remains rather large. As soon as two or more

TABLE 1
Overview of the results for various preconditioning strategies for $d = 10^{-3}$.

	# iterations	Total CPU time [s]	$\ \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\ _2 / \ \mathbf{b}\ _2$	$\ \mathbf{x} - \hat{\mathbf{x}}\ _2 / \ \mathbf{x}\ _2$
No precon	187	1376	9.0×10^{-15}	2.4×10^{-14}
IFMM ($n = 1$)	27	292	5.0×10^{-15}	3.4×10^{-15}
IFMM ($n = 2$)	9	154	8.0×10^{-16}	8.6×10^{-16}
IFMM ($n = 3$)	5	153	7.0×10^{-16}	8.6×10^{-16}
IFMM ($n = 4$)	4	325	1.3×10^{-17}	7.5×10^{-16}

Chebyshev nodes are employed, a remarkable reduction of the number of iterations is obtained and convergence is much faster. Figure 20b and Table 2 illustrate that the use of three Chebyshev nodes per dimension results once more in the smallest overall computation time. It is also important to note that although the relative residual $\|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|_2 / \|\mathbf{b}\|_2$ reaches the specified tolerance ε , the actual relative error $\|\mathbf{x} - \hat{\mathbf{x}}\|_2 / \|\mathbf{x}\|_2$ is a few order of magnitudes larger. This is due to the large condition number of \mathbf{A} .

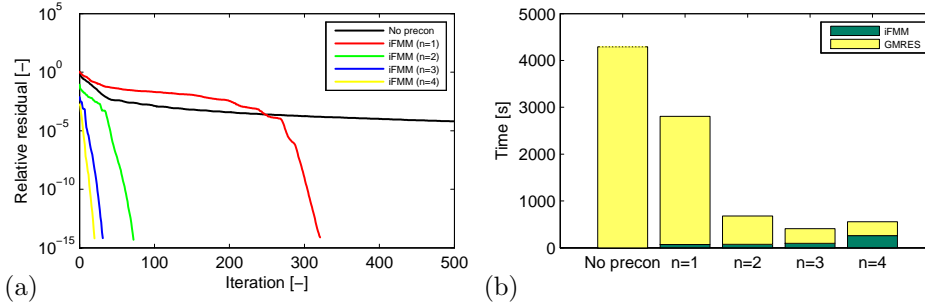


FIG. 20. (a) Relative residual $\|\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\|_2 / \|\mathbf{b}\|_2$ as a function of the iteration step and (b) total CPU time for $d = 10^{-2}$.

TABLE 2
Overview of the results for various preconditioning strategies for $d = 10^{-2}$. The results are surrounded with parentheses if no convergence was achieved.

	# iterations	Total CPU time [s]	$\ \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}\ _2 / \ \mathbf{b}\ _2$	$\ \mathbf{x} - \hat{\mathbf{x}}\ _2 / \ \mathbf{x}\ _2$
No precon	(500)	(4290)	(6.3×10^{-5})	(4.2×10^{-2})
IFMM ($n = 1$)	321	2808	7.8×10^{-15}	1.1×10^{-9}
IFMM ($n = 2$)	72	677	4.9×10^{-15}	5.4×10^{-10}
IFMM ($n = 3$)	31	408	6.6×10^{-15}	3.4×10^{-9}
IFMM ($n = 4$)	20	555	6.4×10^{-15}	7.6×10^{-11}

In order to provide more insight in the effect of the IFMM as preconditioner, the eigenvalue distributions of the original and preconditioned matrices \mathbf{A} and $\mathbf{P}^{-1}\mathbf{A}$ are investigated, where \mathbf{P}^{-1} represents the IFMM preconditioner. A smaller problem of size $N = 10^4$ is considered for this purpose. Figure 21a and 21b show the eigenvalues of \mathbf{A} and $\mathbf{P}^{-1}\mathbf{A}$ for $d = 10^{-3}$ and $d = 10^{-2}$, respectively. As \mathbf{A} is symmetric, all eigenvalues are real. It is clear from these figures that the eigenvalues of \mathbf{A} are not well clustered. This explains the slow convergence of GMRES if no preconditioning is employed. Application of the IFMM as preconditioner leads to a better clustering

of the eigenvalues. An increase of the number of Chebyshev nodes results in a more compact clustering of the eigenvalues near 1, and consequently a faster convergence of the iterative solver.

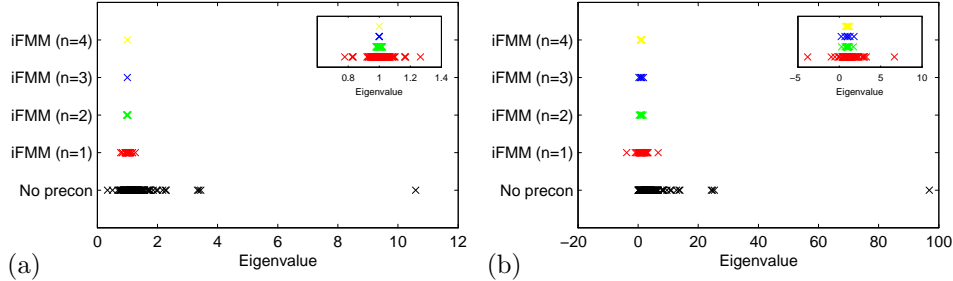


FIG. 21. Eigenvalue distribution of the original (black crosses) and the preconditioned (colored crosses) matrices of size $N = 10^4$, for (a) $d = 10^{-3}$ and (b) $d = 10^{-2}$.

The results presented in this subsection clearly illustrate the effectiveness of the IFMM as preconditioner for an iterative solver. There is a trade-off between the preconditioner's accuracy (and thus the number of iterations) and its construction cost, which has to be evaluated on a case-by-case basis.

4. Application: rigid bodies immersed in a Stokes flow. In order to illustrate the applicability of the IFMM to challenging science and engineering problems (as opposed to the synthetic examples discussed in Section 3), an application involving rigid bodies immersed in a Stokes flow is considered in this section. These rigid bodies represent, for example, colloidal particles suspended in a solvent [17]. Each rigid body is discretized with a number of points \mathbf{r}_i on the surface, and a force vector $\mathbf{F}_i \in \mathbb{R}^{3 \times 1}$ and velocity vector $\mathbf{U}_i \in \mathbb{R}^{3 \times 1}$ are associated with each point. The linear system of equations that needs to be solved reads as:

$$(30) \quad \mathcal{M}\mathbf{F} = \mathbf{U}$$

where $\mathbf{F} \in \mathbb{R}^{3N \times 1}$ and $\mathbf{U} \in \mathbb{R}^{3N \times 1}$ collect the forces and velocities of all N points, respectively. The mobility matrix $\mathcal{M} \in \mathbb{R}^{3N \times 3N}$ consists of $N \times N$ blocks of size 3×3 . A subblock $\mathbf{M}_{ij} \in \mathbb{R}^{3 \times 3}$ corresponding to the interaction between points \mathbf{r}_i and \mathbf{r}_j can be expressed as:

$$(31) \quad \mathbf{M}_{ij} = f(\|\mathbf{r}_{ij}\|) \mathbf{I} + g(\|\mathbf{r}_{ij}\|) \hat{\mathbf{r}}_{ij} \otimes \hat{\mathbf{r}}_{ij}$$

with $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$, while a hat in Eq. (31) denotes a unit vector. $f(r)$ and $g(r)$ are scalar functions of distance. In this paper, the commonly used Rotne-Prager-Yamakawa (RPY) kernel [44] is employed for $f(r)$ and $g(r)$, leading to a symmetric positive-semidefinite mobility matrix \mathcal{M} . The reader is referred to [33] for a detailed overview of the immersed boundary method that leads to Eq. (30).

The dense system presented in Eq. (30) is solved with GMRES; the matrix-vector products are evaluated with an FMM-scheme [33]. An arbitrary velocity vector \mathbf{U} is taken as right hand side. First, a regular cubic lattice consisting of $16 \times 16 \times 16$ spherical rigid bodies is considered. Each sphere is discretized with 42 particles, resulting in a system of equations with $3 \times 42 \times 4096 = 516096$ unknowns. Figure 22a shows the relative residual as a function of the iteration step in the GMRES-algorithm, and this for three different preconditioning approaches. If no preconditioning is applied,

convergence to the specified relative residual $\varepsilon = 10^{-8}$ is achieved after 231 iterations. The use of a simple block-diagonal preconditioner (where the diagonal subblocks of size 126×126 correspond to the mobility matrices of the individual spheres) already leads to a significant reduction of the number of iterations and hence the overall computation time (see Figure 22 and Table 3). This is a very cheap preconditioning strategy (i.e., the time needed to construct the preconditioner is negligible) that performs very well in this particular case, as the subblocks in the block-diagonal matrix represent the mobilities of disjoint rigid bodies. Application of the IFMM as preconditioner (with $n = 2$ Chebyshev nodes) results in a further reduction of the number of iterations. The overall speed-up remains rather limited, however, as the preconditioner’s construction cost is considerably higher. It is nevertheless remarkable that the IFMM still results in a speed-up, given the fact that the block-diagonal preconditioner already performs very well.

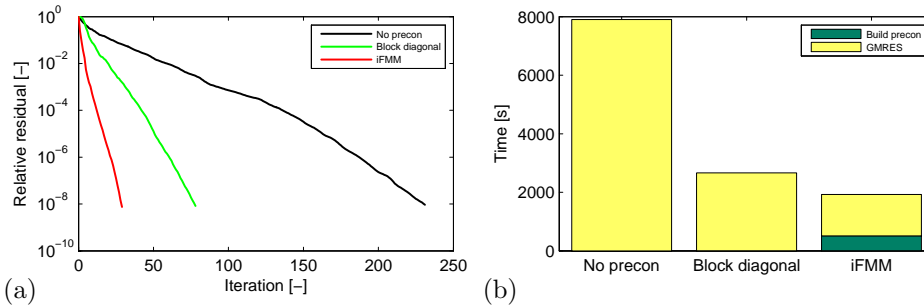


FIG. 22. (a) Relative residual $\|\mathbf{U} - \mathcal{M}\hat{\mathbf{F}}\|_2 / \|\mathbf{U}\|_2$ as a function of the iteration step and (b) total CPU time for a cubic lattice consisting of 4096 spheres, each discretized with 42 particles.

TABLE 3

Overview of the results for various preconditioning strategies for a cubic lattice consisting of 4096 spheres, each discretized with 42 particles.

	# iterations	Total CPU time [s]	$\ \mathbf{U} - \mathcal{M}\hat{\mathbf{F}}\ _2 / \ \mathbf{U}\ _2$
No preconditioner	231	7900	9.3×10^{-9}
Block diagonal	78	2662	8.3×10^{-9}
IFMM ($n = 2$)	29	1928	7.5×10^{-9}

The second example involves six concentric spherical shells with a discretization ranging from 42 particles for the smallest inner shell to 40962 particles for the largest outer shell; the total number of particles equals 54612. The mobility matrix \mathcal{M} is more ill-conditioned compared to the previous example. Once again, GMRES is employed for solving Eq. (30). Figure 23a and Table 4 indicate that 331 iterations are required if no preconditioner is incorporated. The use of a block-diagonal preconditioner (where diagonal subblocks of size 108×108 have been employed) does not result in a speed-up; it even deteriorates the convergence. While this was a successful preconditioning strategy in the previous example, such an approach is too simplistic to be effective in this test case, as the diagonal subblocks do not correspond any more to the mobility matrices of disjoint spheres. Application of the IFMM as preconditioner (with $n = 3$ Chebyshev nodes), on the other hand, leads once more to a very fast convergence of the relative residual and consequently a considerable reduction of the computation

time (see Figure 23b). Note that this is a relatively small problem compared to the previous example, as only $3 \times 54612 = 163836$ unknowns are considered. The relative speed-up of the IFMM is expected to be even more significant for larger problems.

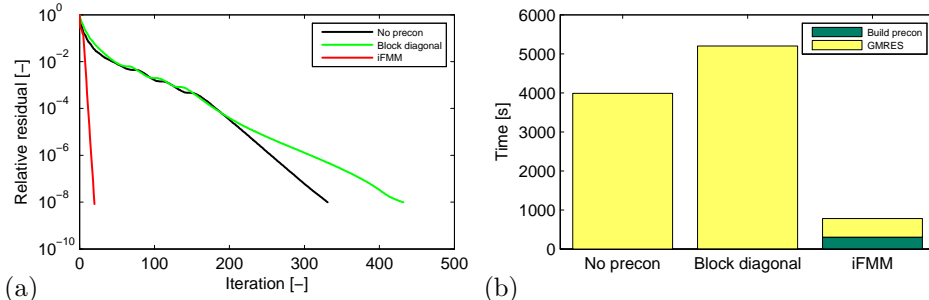


FIG. 23. (a) Relative residual $\|\mathbf{U} - \mathcal{M}\hat{\mathbf{F}}\|_2 / \|\mathbf{U}\|_2$ as a function of the iteration step and (b) total CPU time for six concentric spherical shells discretized with 54612 particles.

TABLE 4

Overview of the results for various preconditioning strategies for six concentric spherical shells discretized with 54612 particles.

	# iterations	Total CPU time [s]	$\ \mathbf{U} - \mathcal{M}\hat{\mathbf{F}}\ _2 / \ \mathbf{U}\ _2$
No precondition	331	3988	9.9×10^{-9}
Block diagonal	432	5202	9.8×10^{-9}
IFMM ($n = 3$)	20	778	8.3×10^{-9}

5. Conclusions and future work. In this paper, the inverse fast multipole method has been presented as a fast direct solver for linear systems of equations involving dense \mathcal{H}^2 -matrices. The method is inexact, although an arbitrary high accuracy can be achieved by tuning the parameters appropriately. Because the error can be related to the decay of the singular values for low-rank fill-in blocks, we expect the error to decay roughly geometrically in many applications as a function of the rank. The overall computational cost scales like $\mathcal{O}(r^2 N)$ where r is the average rank at the leaf level. In terms of the tolerance criterion ε , the cost is $\mathcal{O}(N \log(1/\varepsilon)^2)$.

The algorithm is based on two key ideas. First, the \mathcal{H}^2 -structure of the dense matrix is exploited to construct an extended sparse matrix. Second, fill-ins arising throughout the elimination of the latter are compressed and redirected if they correspond to well-separated clusters. As a result, the sparsity pattern is maintained (strictly speaking, there is a bound, independent of N , on the fill-in), leading to a fast solver that scales like $\mathcal{O}(N)$ for 3D problems. Numerical examples have been presented to validate the linear scaling of the algorithm. It has furthermore been demonstrated that a low accuracy IFMM solver performs extremely well as a preconditioner for an iterative solver; using the IFMM as a preconditioner is often favorable, compared to its use as an accurate direct solver. Furthermore, an application related to the interaction of rigid bodies immersed in a Stokes fluid has been discussed to illustrate the method's applicability to challenging science and engineering problems.

An important aspect to consider in future work is the parallelization of the proposed methodology. The algorithm is expected to exhibit good parallel scalability, as each cluster only interacts with a limited number of other clusters (i.e., those in

its neighbor and interaction list). Furthermore, alternative low-rank approximation schemes that might be more efficient than interpolation based on Chebyshev polynomials can be investigated.

Although the examples presented in this paper only involve uniform octrees, the method can easily be applied to adaptive trees with a varying number of levels [41, 51].

Appendix A. Additional information. This appendix provides some additional data related to the benchmarks presented in Subsection 3.1. More specifically, the computational cost for the IFMM as a direct solver, as depicted in Figure 17, is analyzed in detail in the following. Figure 24 shows a breakdown of the total CPU time into the time required for initializing the IFMM operators, estimating the largest singular value $\sigma_0(\mathbf{E})$ of the extended sparse matrix \mathbf{E} , and performing the elimination and substitution. The number of Chebyshev nodes (in each dimension) is varied from $n = 1$ to $n = 4$; an increase of n results in a higher overall computation time. It is clear from this figure that the elimination phase strongly dominates the total computation time, while the contribution of the other components, and especially of the substitution phase, is almost negligible.

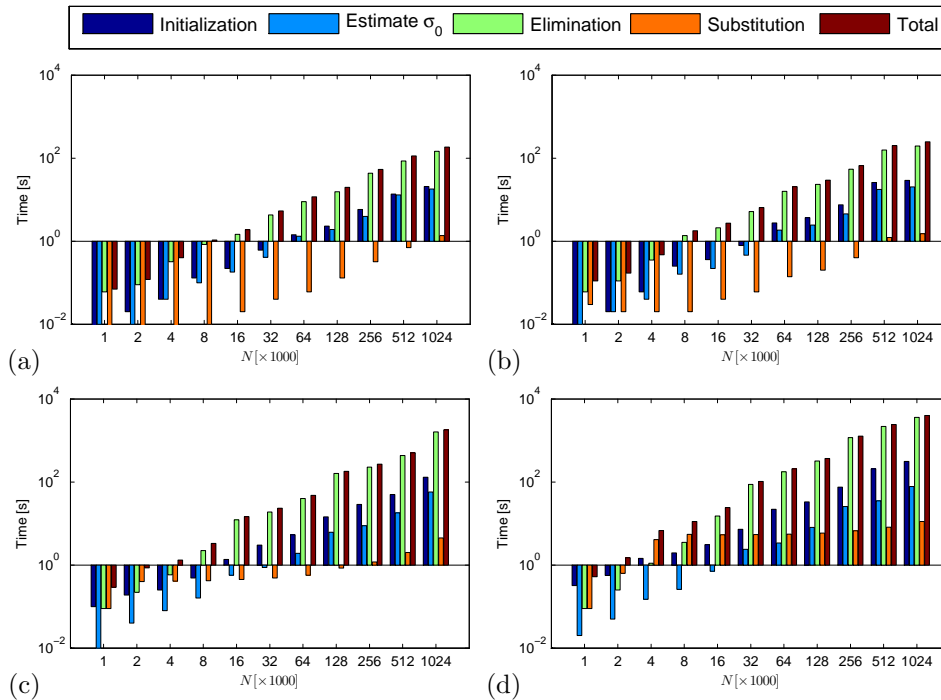


FIG. 24. Breakdown of the total CPU time for the IFMM as a direct solver in case of N points randomly distributed in the cube $[-1, 1]^3$. The number of Chebyshev nodes (in each dimension) used for the initial low-rank approximations corresponds to (a) $n = 1$, (b) $n = 2$, (c) $n = 3$, and (d) $n = 4$. From left to right, each bar represents: initialization, estimation of $\sigma_0(\mathbf{E})$, elimination, substitution, and total computation time (brown bars). The elimination step (light green bars) dominates the total time in most cases as expected.

A more detailed breakdown of the time spent on different types of operations throughout the elimination phase is provided in Figure 25. Following operations are distinguished: performing LU-factorizations and -solves, executing matrix-matrix multiplications, obtaining low-rank approximations through rSVD and ACA, and

transferring operators to the parent level (corresponding to Eq. (25)–(27)). The computation of low-rank approximations is dominant in case of $n = 1$ and $n = 2$, while matrix-matrix multiplications are the most time consuming operations for $n = 3$ and $n = 4$. In all cases, the contribution of transferring operators from one level to the parent level is significantly less important than the other components.

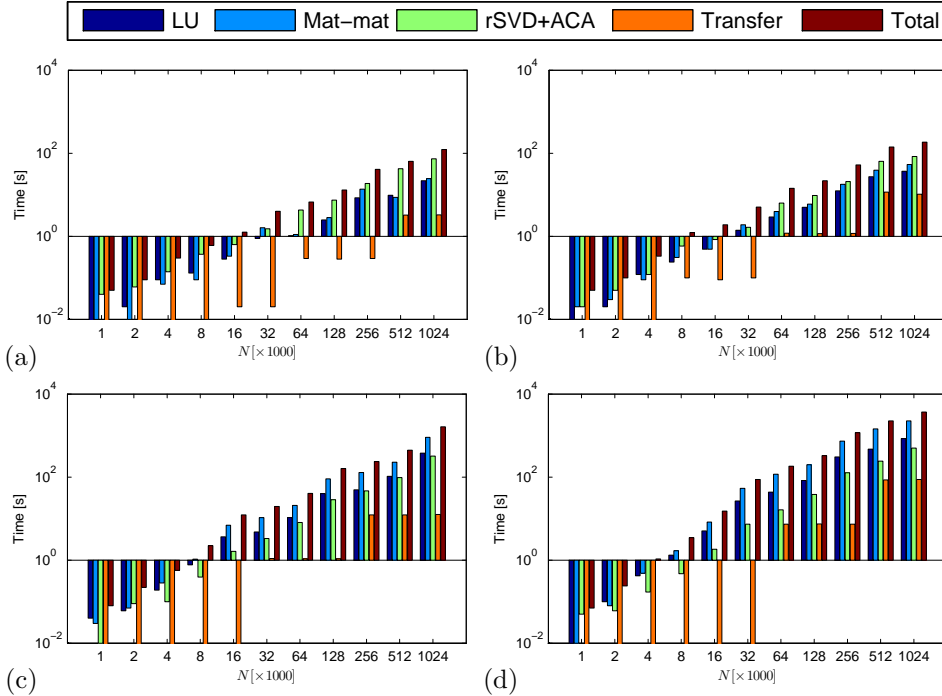


FIG. 25. Breakdown of the CPU time spent on different types of operations throughout the elimination phase of the IFMM algorithm. The number of Chebyshev nodes (in each dimension) used for the initial low-rank approximations corresponds to (a) $n = 1$, (b) $n = 2$, (c) $n = 3$, and (d) $n = 4$. From left to right, each bar represents: block LU factorization and triangular solves, matrix-matrix products (trailing matrix update in the block LU factorization), low-rank approximations, operator transfer (from child to parent), and total elimination time (brown bars). For small n , the low-rank approximations (light green bars) dominate, while for large n , the matrix-matrix products (light blue bars) dominate.

Finally, the different behavior of fill-ins arising between neighboring or well-separated clusters is illustrated. Figure 26 shows the singular values of some fill-ins $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$ that are encountered throughout the elimination phase, both for neighboring ($\Omega_j \in \mathcal{N}_k$) and well-separated ($\Omega_j \notin \mathcal{N}_k$) clusters Ω_j and Ω_k . This figure confirms that the decay of singular values is significantly faster in the case of well-separated clusters, indicating that $\mathbf{E}'(\underline{\mathbf{k}}, \underline{\mathbf{j}})$ can then indeed accurately be approximated using a low-rank representation.

Appendix B. Classification and relation between different methods. We briefly present the different \mathcal{H} -matrix formats and discuss connections between the algorithm presented in this paper and the fast solvers developed by Hackbusch et al. [4, 8, 24, 26, 27].

\mathcal{H} -matrices can be classified into four broad categories.

1) $w\mathcal{H}$ with non-nested basis. Off-diagonal blocks that correspond to clusters




FIG. 26. Singular values σ_k of fill-ins $\mathbf{E}'(\underline{k}, \underline{j})$ arising between neighboring (black line) and well-separated (grey line) clusters Ω_j and Ω_k .

that are well-separated are low-rank. The low-rank basis is non-nested, that is the low-rank basis of a node is unrelated to the basis of its children. This is the most general \mathcal{H} -format. The matrix is generally speaking expanded in the form:

$$\mathbf{A} = \mathbf{A}_s + \sum_{l=0}^{\ln N} \mathbf{U}_l \mathbf{\Sigma}_l \mathbf{V}_l^T$$

where \mathbf{A}_s is a block sparse matrix.

2) $w\mathcal{H}$ with nested basis. The low-rank basis of a node is a linear combination of the basis of its children nodes. In that case, the expansion can be written in recursive form:

$$\mathbf{A} = \mathbf{A}_s + \mathbf{U} \mathbf{A}_{\mathcal{H}} \mathbf{V}^T$$

where $\mathbf{A}_{\mathcal{H}}$ is a $w\mathcal{H}$ matrix of smaller size. FMM matrices [11, 18], and the matrices discussed in this paper are of this form. This also corresponds to \mathcal{H}^2 -matrices [7, 25, 26].

3) $s\mathcal{H}$ with non-nested basis. All off-diagonal blocks are assumed to be low-rank. The matrix decomposition is:

$$\mathbf{A} = \mathbf{A}_d + \sum_{l=0}^{\ln N} \mathbf{U}_l \mathbf{\Sigma}_l \mathbf{V}_l^T$$

where \mathbf{A}_d is a block diagonal matrix. This is the same as the HODLR matrix format [1].

4) $s\mathcal{H}$ with nested basis:

$$\mathbf{A} = \mathbf{A}_d + \mathbf{U} \mathbf{A}_{\mathcal{H}} \mathbf{V}^T$$

where $\mathbf{A}_{\mathcal{H}}$ is an $s\mathcal{H}$ matrix. This is the HSS format [12, 49].

We now discuss the connection between our algorithm and Hackbusch et al. [4, 8, 26, 24, 27]. Both methods use a tree decomposition of the matrix in order to identify low-rank blocks. Although the Hackbusch algorithm was written differently, this algorithm can be derived using the extended sparse matrix introduced in Subsection 2.1.

The key difference with this paper is that Hackbusch's algorithm performs the LU factorization using a depth-first in-order (symmetric) traversal. In contrast, our

method uses a breadth-first or level-by-level traversal, starting from the leaf level and moving up to the root.

This can be illustrated on a simple example for concreteness. Consider the simplest \mathcal{H} -matrix:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{00} & \mathbf{U}_0 \mathbf{V}_1^T \\ \mathbf{U}_1 \mathbf{V}_0^T & \mathbf{A}_{11} \end{pmatrix}$$

The extended sparse format is:

$$\mathbf{E} = \begin{pmatrix} \mathbf{A}_{00} & & \mathbf{U}_0 & \\ & \mathbf{A}_{11} & & \mathbf{U}_1 \\ \mathbf{V}_0^T & & & -\mathbf{I} \\ & \mathbf{V}_1^T & -\mathbf{I} & \end{pmatrix}$$

Hackbusch's algorithm is based on a block LU factorization. Such a factorization is obtained if, in \mathbf{E} , we eliminate first \mathbf{A}_{00} , then the auxiliary variables, then \mathbf{A}_{11} . Step-by-step, let's eliminate \mathbf{A}_{00} :

$$\begin{pmatrix} \mathbf{A}_{11} & & \mathbf{U}_1 \\ & -\mathbf{V}_0^T \mathbf{A}_{00}^{-1} \mathbf{U}_0 & -\mathbf{I} \\ \mathbf{V}_1^T & & -\mathbf{I} \end{pmatrix}$$

If we eliminate pivots 2 and 3 (before \mathbf{A}_{11}), we simply get:

$$\mathbf{A}_{11} - \mathbf{U}_1 \mathbf{V}_0^T \mathbf{A}_{00}^{-1} \mathbf{U}_0 \mathbf{V}_1^T$$

This is the expected Schur complement from a block LU factorization.

If we represent \mathbf{E} using a tree, for a 3-level \mathcal{H} -matrix, this elimination corresponds to the order shown in Figure 27.




FIG. 27. Tree node elimination order following Hackbusch's algorithm. The traversal is a depth-first in-order tree traversal.

In contrast, our algorithm follows a level-by-level elimination. Going back to \mathbf{E} , we first eliminate \mathbf{A}_{00} , then \mathbf{A}_{11} . This leads to a new linear system involving the auxiliary variables only:

$$\begin{pmatrix} -\mathbf{V}_0^T \mathbf{A}_{00}^{-1} \mathbf{U}_0 & -\mathbf{I} \\ -\mathbf{I} & -\mathbf{V}_1^T \mathbf{A}_{11}^{-1} \mathbf{U}_1 \end{pmatrix}$$

As a result, in terms of traversal ordering, this paper follows a level-by-level elimination. This is shown in Figure 28.




FIG. 28. Tree node elimination order following this paper’s algorithm. The traversal is a level-by-level (breadth-first) traversal.

Acknowledgments The authors would like to thank Prof. Aleksandar Donev and Dr. Floren Balboa (Courant Institute of Mathematical Sciences, New York University) for their assistance in obtaining the results presented in Section 4. The first author is a post-doctoral fellow of the Research Foundation Flanders (FWO) and a Francqui Foundation fellow of the Belgian American Educational Foundation (BAEF). The financial support is gratefully acknowledged. Funding from the “Army High Performance Computing Research Center” (AHPARC), sponsored by the U.S. Army Research Laboratory under contract No. W911NF-07-2-0027, at Stanford, supported in part this research.

REFERENCES

- [1] S. AMBIKASARAN AND E. DARVE, *An $\mathcal{O}(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices*, *Journal of Scientific Computing*, 57 (2013), pp. 477–501.
- [2] ———, *The Inverse Fast Multipole Method*, arXiv, (2014), p. 1407.1572.
- [3] A. AMINFAR, S. AMBIKASARAN, AND E. DARVE, *A fast block low-rank dense solver with applications to finite-element matrices*, arXiv, (2014), p. 1403.5337.
- [4] M. BEBENDORF, *Hierarchical Matrices: A Means to Efficiently Solve Elliptic Boundary Value Problems*, Springer Publishing Company, 1st ed., 2008.
- [5] M. BEBENDORF AND S. RJASANOW, *Adaptive low-rank approximation of collocation matrices*, *Computing*, 70 (2003), pp. 1–24.
- [6] M. BONNET, *Boundary integral equation methods for solids and fluids*, John Wiley & Sons, Chichester, United Kingdom, 1995.
- [7] S. BÖRM, *Data-sparse approximation of non-local operators by \mathcal{H}^2 -matrices*, *Linear Algebra and its Applications*, 422 (2007), pp. 380–403.
- [8] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Introduction to hierarchical matrices with applications*, *Engineering Analysis with Boundary Elements*, 27 (2003), pp. 405–422.
- [9] M. D. BUHMANN, *Radial basis functions: theory and implementations*, Cambridge University Press, Cambridge, UK, 2003.
- [10] B. CARPENTIERI, I. S. DUFF, L. GIRAUD, AND M. MAGOLU MONGA MADE, *Sparse symmetric preconditioners for dense linear systems in electromagnetism*, *Numerical Linear Algebra with Applications*, 11 (2004), pp. 753–771.
- [11] J. CARRIER, L. GREENGARD, AND V. ROKHLIN, *A fast adaptive multipole algorithm for particle simulations*, *SIAM Journal on Scientific and Statistical Computing*, 9 (1988), pp. 669–686.
- [12] M. CHANDRASEKARAN, S. GU AND T. PALS, *A fast ULV decomposition solver for hierarchically semi-separable representations*, *SIAM Journal on Matrix Analysis and Applications*, 28 (2006), pp. 603–622.
- [13] S. CHANDRASEKARAN, P. DEWILDE, M. GU, W. LYONS, AND T. PALS, *A fast solver for HSS representations via sparse matrices*, *SIAM Journal on Matrix Analysis and Applications*, 29 (2007), pp. 67–81.
- [14] S.H. CHRISTIANSEN AND J.-C. NÉDÉLEC, *A preconditioner for the electric field integral equation based on Calderon formulas*, *SIAM Journal on Numerical Analysis*, 40 (2002), pp. 1100–1135.

- [15] M. DARBAS, E. DARRIGRAND, AND Y. LAFRANCHE, *Combining analytic preconditioner and Fast Multipole Method for the 3-D Helmholtz equation*, Journal of Computational Physics, 236 (2013), pp. 289–316.
- [16] E. DARVE, *The Fast Multipole Method: Numerical Implementation*, Journal of Computational Physics, 160 (2000), pp. 195–240.
- [17] S. DELONG, F.B. USABIAGA, R. DELGADO-BUSCALIONI, B.E. GRIFFITH, AND A. DONEV, *Brownian dynamics without Green's functions*, The Journal of Chemical Physics, 140 (2014), p. 134110.
- [18] W. FONG AND E. DARVE, *The black-box fast multipole method*, Journal of Computational Physics, 228 (2009), pp. 8712–8725.
- [19] A. GILLMAN AND P.-G. MARTINSSON, *An $O(N)$ algorithm for constructing the solution operator to 2D elliptic boundary value problems in the absence of body loads*, Advances in Computational Mathematics, 40 (2014), pp. 773–796.
- [20] G.H. GOLUB AND C.F. VAN LOAN, *Matrix computations*, John Hopkins University Press, Baltimore, MD, 3d ed., 1996.
- [21] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, Journal of Computational Physics, 73 (1987), pp. 325–348.
- [22] G. GUENNEBAUD, B. JACOB, ET AL., *Eigen v3*. <http://eigen.tuxfamily.org>, 2010.
- [23] N.A. GUMEROV AND R. DURAIWAMI, *Fast radial basis function interpolation via preconditioned Krylov iteration*, SIAM Journal on Scientific Computing, 29 (2007), pp. 1876–1899.
- [24] W. HACKBUSCH, *A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices*, Computing, 62 (1999), pp. 89–108.
- [25] W. HACKBUSCH AND S. BÖRM, *Data-sparse approximation by adaptive \mathcal{H}^2 -matrices*, Computing, 69 (2002), pp. 1–35.
- [26] W. HACKBUSCH AND S. BÖRM, *\mathcal{H}^2 -matrix approximation of integral operators by interpolation*, Applied Numerical Mathematics, 43 (2002), pp. 129–143.
- [27] W. HACKBUSCH AND B.N. KHOROMSKIJ, *A sparse \mathcal{H} -matrix arithmetic: general complexity estimates*, Journal of Computational and Applied Mathematics, 125 (2000), pp. 479–501.
- [28] W. HACKBUSCH, B. N. KHOROMSKIJ, AND R. KRIEMANN, *Hierarchical matrices based on a weak admissibility criterion*, Computing, 73 (2004), pp. 207–243.
- [29] W. HACKBUSCH AND Z.P. NOWAK, *On the fast matrix multiplication in the boundary element method by panel clustering*, Numerische Mathematik, 54 (1989), pp. 463–491.
- [30] N. HALKO, P.-G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.
- [31] K.L. HO AND L. GREENGARD, *A fast direct solver for structured linear systems by recursive skeletonization*, SIAM Journal on Scientific Computing, 34 (2012), pp. 2507–2532.
- [32] G.C. HSIAO, *Boundary element methods – an overview*, Applied Numerical Mathematics, 56 (2006), pp. 1356–1369.
- [33] B. KALLEMOV, A. PAL SINGH BHALLA, B.E. GRIFFITH, AND A. DONEV, *An immersed boundary method for rigid bodies*, arXiv, (2015), p. 1505.07865.
- [34] R. KALMAN, *Mathematical description of linear dynamical systems*, SIAM Journal on Control and Optimization, 1 (1963), pp. 152–192.
- [35] W.Y. KONG, J. BREMER, AND V. ROKHLIN, *An adaptive fast direct solver for boundary integral equations in two dimensions*, Applied and Computational Harmonic Analysis, 31 (2011), pp. 346–369.
- [36] U. LANGER AND D. PUSCH, *Data-sparse algebraic multigrid methods for large scale boundary element equations*, Applied Numerical Mathematics, 54 (2005), pp. 406–424.
- [37] J. LEE, J. ZHANG, AND C. LU, *Incomplete LU preconditioning for large scale dense complex linear systems from electromagnetic wave scattering problems*, Journal of Computational Physics, 185 (2003), pp. 158–175.
- [38] P.-G. MARTINSSON AND V. ROKHLIN, *A fast direct solver for boundary integral equations in two dimensions*, Journal of Computational Physics, 205 (2005), pp. 1–23.
- [39] N. NISHIMURA, *Fast multipole accelerated boundary integral equation methods*, Applied Mechanics Reviews, 55 (2002), pp. 299–324.
- [40] T. PALS, *Multipole for Scattering Computations: Spectral Discretization, Stabilization, Fast Solvers*, PhD thesis, University of California Santa Barbara, 2004.
- [41] H. POURANSARI AND E. DARVE, *Optimizing the adaptive fast multipole method for fractal sets*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1040–A1066.
- [42] S. RJASANOW AND O. STEINBACH, *The Fast Solution of Boundary Integral Equations (Mathematical and Analytical Techniques with Applications to Engineering)*, Springer-Verlag New York, 2007.

- [43] V. ROKHLIN, *Rapid solution of integral equations of classical potential theory*, Journal of Computational Physics, 60 (1985), pp. 187–207.
- [44] J. ROTNE AND S. PRAGER, *Variational treatment of hydrodynamic interaction in polymers*, The Journal of Chemical Physics, 50 (1969), pp. 4831–4837.
- [45] Y. SAAD, *Krylov subspace methods for solving large unsymmetric linear systems*, Mathematics of Computation, 37 (1981), pp. 105–126.
- [46] ———, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM Journal on Scientific and Statistical Computing, 14 (1993), pp. 461–469.
- [47] Y. SAAD AND M.H. SCHULTZ, *GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [48] S. SAUTER, *The panel clustering method in 3-D BEM*, in Wave Propagation in Complex Media, G. Papanicolaou, ed., vol. 96 of The IMA Volumes in Mathematics and its Applications, Springer New York, 1998, pp. 199–224.
- [49] Z. SHENG, P. DEWILDE, AND S. CHANDRASEKARAN, *Algorithms to solve hierarchically semi-separable systems*, in System Theory, the Schur Algorithm and Multidimensional Analysis, Daniel Alpay and Victor Vinnikov, eds., vol. 176 of Operator Theory: Advances and Applications, Birkhäuser Basel, 2007, pp. 255–294.
- [50] U. TROTTEBERG, C.W. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Academic Press, 2000.
- [51] L. YING, G. BIROS, AND D. ZORIN, *A kernel-independent adaptive fast multipole algorithm in two and three dimensions*, Journal of Computational Physics, 196 (2004), pp. 591–626.