

Algebraic Inverse Fast Multipole Method: A fast direct solver that is better than HODLR based fast direct solver

Vaishnavi Gujjula^a, Sivaram Ambikasaran^a

^a*Department of Mathematics, Indian Institute of Technology Madras, Chennai, 600036, Tamil Nadu, India*

Abstract

This article presents a fast direct solver, termed Algebraic Inverse Fast Multipole Method (from now on abbreviated as AIFMM), for linear systems arising out of N -body problems. AIFMM relies on the following three main ideas: (i) Certain sub-blocks in the matrix corresponding to N -body problems can be efficiently represented as low-rank matrices; (ii) The low-rank sub-blocks in the above matrix are leveraged to construct an extended sparse linear system; (iii) While solving the extended sparse linear system, certain fill-ins that arise in the elimination phase are represented as low-rank matrices and are "redirected" through other variables maintaining zero fill-in sparsity. The main highlights of this article are the following: (i) Our method is completely algebraic (as opposed to the existing Inverse Fast Multipole Method [1, 2, 3], from now on abbreviated as IFMM). We rely on our new Nested Cross Approximation [4] (from now on abbreviated as NNCA) to represent the matrix arising out of N -body problems. (ii) A significant contribution is that the algorithm presented in this article is more efficient than the existing IFMMs. In the existing IFMMs, the fill-ins are compressed and redirected as and when they are created. Whereas in this article, we update the fill-ins first without affecting the computational complexity. We then compress and redirect them only once. (iii) Another noteworthy contribution of this article is that we provide a comparison of AIFMM with Hierarchical Off-Diagonal Low-Rank (from now on abbreviated as HODLR) based fast direct solver and NNCA powered GMRES based fast iterative solver. (iv) Additionally, AIFMM is also demonstrated as a preconditioner.

Keywords: Fast direct solver, Extended sparsification, Fast Multipole Method, Hierarchical matrices, Low-Rank matrices, Nested Cross Approximation, Preconditioner
2000 MSC: 65F05, 65F08, 65Y20

1. Introduction

This article focuses on solving linear systems that arise out of N -body problems. Such N -body problems arise frequently in many applications such as electrostatics, integral equation solvers, radial basis function interpolation, inverse problems, Gaussian process regression, wave scattering, etc. A linear system can be solved using a direct solver or an iterative solver. Both have their advantages over the other and it is highly problem specific to choose the solver that is better.

An important and widely used iterative technique is the Krylov subspace technique, which involves matrix-vector products. To speed up these matrix-vector products, fast summation techniques such as Fast Multipole Method (FMM) [5], Barnes-Hut [6], FFT, etc. are used. Further for fast convergence in problems with high condition numbers, an iterative solver is coupled with a preconditioner.

On the other hand, direct solvers involve a factorization step followed by a solve step. The factorization step comprises of an LU factorization or QR factorization, etc., which generally is computationally more expensive than an iterative technique. But direct solvers are more robust and accurate than iterative solvers. Further, direct solvers are advantageous when one is interested in multiple right-hand sides. A naive direct solver costs $O(N^3)$, which is prohibitively large for large system sizes. To reduce the computational complexity, fast methods are used. Many dense matrices arising out of N -body problems possess a hierarchical low-rank structure. This low-rank structure is exploited to construct hierarchical matrices and hierarchical matrices based fast direct solvers [7, 8, 9, 10].

While constructing a hierarchical matrix, the low-rank bases of the sub-blocks that are compressed can be obtained in a nested or a non-nested approach. In the nested approach, the low-rank bases at a parent level in the hierarchy are

constructed from the bases at the child level. The class of hierarchical matrices that follow the nested approach are called \mathcal{H}^2 matrices.

The matrix sub-blocks that are low-rank approximated in a hierarchical matrix are identified based on an admissibility condition. The most widely used admissibility conditions are the weak admissibility and strong or standard admissibility conditions.

Hierarchically Off-Diagonal Low-Rank (HODLR) and Hierarchically Semi-Separable (HSS) matrices [11, 12, 13, 14, 15, 16, 17] are sub-classes of Hierarchical matrices that follow the weak admissibility condition, wherein all the off-diagonal sub-blocks are approximated by low-rank matrices. The former follows a non-nested approach in the construction of the bases and the latter follows the nested approach. The major drawback with these classes of Hierarchical matrices is that the ranks of the compressed sub-matrices are not "truly" low-rank. For instance in 2D, the ranks of the compressed sub-matrices grow as $O(\sqrt{N} \log^2(N))$ [18], in 3D it is $O(N^{\frac{2}{3}} \log^3(N))$, where N is the size of the compressed sub-matrix. In a d -dimensional setting, the rank of the compressed sub-matrices grow as $O(N^{\frac{d-1}{d}} \log^d(N))$ [19]. So, the direct solvers developed for HODLR and HSS classes of matrices are not linear in complexity.

Another sub-class of Hierarchical matrices is the \mathcal{H}^2 matrices with strong admissibility condition [9, 20, 21], wherein interactions between neighboring clusters of particles are not compressed and interactions between well-separated clusters of particles are approximated by low-rank matrices. This sub-class of Hierarchical matrices can be considered to be the algebraic generalization of FMM, so they are also referred as FMM matrices. The strong admissibility condition guarantees that the interactions between well-separated clusters of particles (when the underlying matrix is generated from singular kernels) [19] do not scale with the number of particles in the cluster.

There exists a vast literature on \mathcal{H}^2 matrices and solvers for linear systems involving \mathcal{H}^2 matrices in almost linear complexity [22, 23, 24, 25, 26]. The pre-factors in the scaling term of these methods tend to be large. Inverse fast multipole method (from now on abbreviated as IFMM), a fast direct solver for FMM matrices with linear complexity, was introduced [1], whose pre-factor in the scaling term is not that large. Another related work on fast direct solvers for FMM matrices is the strong skeletonization based factorization method [27, 28].

The work developed in this article is a variant of the IFMM developed in [1, 2, 3]. One of the key ideas based on which the IFMM is developed is the extended sparsification technique that was earlier used in [11, 29, 30]. In IFMM, an extended sparse system of size $O(N)$ is developed by introducing auxiliary variables - which are the locals and multipoles of nodes at various levels of the FMM tree. The advantage of sparsification is that the computational complexity of the solver gets improved, provided the fill-ins are minimal. In IFMM, the fill-ins corresponding to well-separated clusters of particles are compressed and redirected via the existing non-zero entries, which contributes to its linear scaling.

With the auxiliary variables in IFMM being the locals and multipoles of nodes at various levels of the FMM tree, the extended sparse matrix is an assembly of the FMM operators, which are the L2P/L2L (Local-To-Particle/Local-To-Local), M2L (Multipole-To-Local), P2M/M2M (Particle-To-Multipole/Multipole-To-Multipole), and P2P (Particle-To-Particle) operators.

In this article, we develop an Algebraic Inverse Fast Multipole Method (from now on abbreviated as AIFMM), wherein we employ a **new** Nested Cross Approximation (NNCA) [4], an algebraic technique, to obtain the L2P/L2L, M2L and P2M/M2M operators. Nested Cross Approximation (from now on abbreviated as NCA) [31, 32, 4] is a nested version of Adaptive Cross Approximation (from now on abbreviated as ACA), that forms low-rank bases in a nested fashion. NNCA [4] differs from the NCAs described in [31, 32] in the technique of choosing pivots, a key step of the approximation. The search space for far-field pivots of a hypercube (a node belonging to the 2^d tree) is considered to be its interaction list region in the former and is considered to be its entire far-field region in the latter. So the time to build the former approximation is lower than that of the latter, for no significant difference in accuracy.

There have been a few articles that have presented IFMM for various applications. Following are the differences between this article and the earlier articles [1, 2, 3]:

1. In this present article, the FMM operators which are used to form the extended sparse matrix, are obtained using NNCA [4], a purely algebraic technique. While in [2], it is done using Chebyshev interpolation, an analytic technique. And in [3], the extended sparse matrix is assembled using the Low Frequency FMM (LFFMM) operators, which is also an analytic technique. The advantages of an algebraic technique are that i) the method can be used in a black box fashion irrespective of the application ii) the ranks obtained are usually lower than

that of the analytic techniques as the bases obtained through an algebraic method are problem and domain specific.

2. In this article, the fill-in compression and redirection is performed using rank revealing QR (RRQR). While in [2, 3] it is done using randomized SVD.
3. In this article, a more efficient elimination algorithm than the one stated in the existing IFMMs [1, 2, 3] is presented. In the existing IFMMs, fill-ins are compressed and redirected as and when they are created, which could happen multiple times in the elimination process. Whereas in this article, we do not compress and redirect a fill-in as and when created. We update the fill-ins without affecting the computational complexity. We then compress and redirect only once.
4. In this article, we demonstrate AIFMM as a preconditioner in the high frequency scattering problem. While in [2, 3], IFMM is studied as a preconditioner in a Stokes flow problem and a 3D Helmholtz BEM at low frequencies respectively.

Below are the highlights of the AIFMM presented in this article:

1. It is a completely algebraic method, i.e., it does not use any analytic techniques such as the interpolation techniques [2], multipole expansions [3], etc, to obtain the low-rank factorizations.
2. AIFMM is demonstrated as a direct solver for linear systems involving non-oscillatory Green's functions and the 2D Helmholtz function at low frequency.
3. To the best of our knowledge, this work is one of the first to provide a comparison of the performance of AIFMM with that of i) a HODLR based fast direct solver [14] ii) GMRES, an iterative solver. It is observed that AIFMM is faster than HODLR, and GMRES is faster than AIFMM. But when one is interested in solving for multiple right hand sides, AIFMM is faster than GMRES.
4. AIFMM is also demonstrated as a preconditioner in an iterative scheme for high frequency scattering problem. It is observed that AIFMM as a preconditioner is better than the block-diagonal preconditioner, but not as good as the HODLR preconditioner.

The rest of the article is organized as follows: Section 2 describes the preliminaries to develop AIFMM, which are the construction of FMM tree, identification of the low-rank sub-blocks, and assembly of the various FMM operators using NNCA. Section 3 describes AIFMM, which includes the construction of the extended sparse system, elimination phase, and back substitution phase. Section 4, illustrates various numerical benchmarks of AIFMM in comparison to those of GMRES and HODLR.

2. Preliminaries

Let $u \in \mathbb{R}^{N \times d}$ be the coordinates of N targets in d dimensions (we will be referring them as target points), $v \in \mathbb{R}^{N \times d}$ be the coordinates of N sources in d dimensions (we will be referring them as source points). Let $A \in \mathbb{C}^{N \times N}$ be the matrix that captures the pair-wise interaction between these points, i.e., A_{ij} is the interaction between the source and target located at v_j and u_i respectively. Such interaction matrices arise in many applications; to name a few integral equation solvers, particle simulations, covariance matrices, electrostatics, scattering, etc. Electrostatic problems are studied extensively in the literature and the naming conventions in most of the research articles are based on it. Hence, in this article, we follow the nomenclature of Electrostatics.

We assume unknown charges of strength $x \in \mathbb{C}^{N \times 1}$ are located at source points v and the potential $b \in \mathbb{C}^{N \times 1}$ at the target points u is known. We are interested in finding the unknown charges x , given the potential b or in other words solve the system of equations,

$$Ax = b. \tag{1}$$

A key idea of the inverse fast multipole method is to introduce auxiliary variables and then create an extended sparse system of size $O(N)$. The advantage of sparsification is that it reduces the complexity of the problem as some of the fill-ins that get created in the elimination phase are compressed and redirected through the existing non-zero entries, resulting in a linear complexity algorithm.

The extended sparse matrix is created by constructing the FMM matrix representation of A . The multipoles and locals that are formed at various levels of the FMM tree are considered to be the auxiliary variables.

The steps involved in constructing the extended sparse matrix are i) construction of FMM tree ii) identification of the low-rank matrix sub-blocks, and iii) assembly of the various FMM operators. We now describe each of these steps below.

2.1. Construction of FMM tree

We consider a smallest hypercube that contains the support of the particles to be the domain $\Omega \in \mathbb{R}^d$. A 2^d uniform tree is constructed over Ω . The hypercube at level 0, is the domain Ω itself. A hypercube at level l is subdivided into 2^d hypercubes, which are considered to be at level $l+1$ of the tree. The former is considered to be the parent of the latter and the latter are considered to be the children of the former. And this subdivision is carried on hierarchically until a level L is reached where the hypercubes contain no more than n_{max} particles. We will be referring to the hierarchical tree as \mathcal{T}^L . The notations associated with a hypercube i at level l are described in Table 1. For $d = 2$, the construction of quad-tree and the numbering of the nodes till level 2 is illustrated in Figure 1.

$i^{(l)}$	Hypercube (also referred to as node or box) i at level l of the tree
$\mathcal{P}(i^{(l)})$	Parent of $i^{(l)}$
$\mathcal{C}(i^{(l)})$	$\{j^{(l+1)} : j^{(l+1)} \text{ is a child of } i^{(l)}\}$

Table 1: Notations associated with hypercube i at level l




Figure 1: The numbering convention followed at levels 0, 1, and 2 of a quad-tree.

2.2. Identification of the low-rank matrix sub-blocks

Let t^X and s^X be the index sets that store indices of points u and v respectively that lie in hypercube $X^{(l)}$.

$$t^X = \{m : u_m \in X^{(l)}\} \quad (2)$$

$$s^X = \{n : v_n \in X^{(l)}\} \quad (3)$$

In this section and the upcoming sections, we omit the superscript that represents the level at some places, to improve the readability of notations, in the hope that the level can be understood from the context. We follow the strong admissibility condition to identify the sub-blocks of the matrix that can be efficiently approximated by a low-rank matrix, i.e., the interaction between the clusters of particles located in hypercubes $X^{(l)}$ and $Y^{(l)}$, $A_{t^X s^Y}$, is approximated by a low-rank matrix, only if

$$\max\{\text{diam}(X^{(l)}), \text{diam}(Y^{(l)})\} \leq \eta \text{dist}(X^{(l)}, Y^{(l)}), \text{ where} \quad (4)$$

$$\begin{aligned} \text{diam}(X^{(l)}) &= \sup\{\|x - y\|_2 : x, y \in X^{(l)}\}, \\ \text{dist}(X^{(l)}, Y^{(l)}) &= \inf\{\|x - y\|_2 : x \in X^{(l)}, y \in Y^{(l)}\}. \end{aligned}$$

If $X^{(l)}$ and $Y^{(l)}$ satisfy the above stated strong admissibility criterion, then $X^{(l)}$ and $Y^{(l)}$ are said to be well-separated and the interaction matrix $A_{t^X s^Y}$ is said to be admissible. Further $A_{t^X s^Y}$ is considered to be a far-field interaction. If $X^{(l)}$ and $Y^{(l)}$ do not agree with the strong admissibility criterion, then $A_{t^X s^Y}$ is said to be non-admissible and is considered to be a near-field interaction. In this article, we consider $\eta = \sqrt{d}$.

2.2.1. FMM matrix structure

For each node i at level l , we introduce the neighbors and interaction list, described in Table 2. We illustrate the same for a node in 2D in Figure 2.

$\mathcal{N}(i^{(l)})$	Neighbors of $i^{(l)}$ that consists of hypercubes at level l , that do not satisfy the admissibility condition for low-rank.
$\mathcal{IL}(i^{(l)})$	Interaction list of hypercube $i^{(l)}$ that consists of children of $i^{(l)}$'s parent's neighbors that are not its neighbors.

Table 2: Neighbors and interaction list of hypercube i at level l




Figure 2: Illustration of neighbors and interaction list at different levels in 2D.

Let $K^{(l)}$ be the FMM matrix representation of A at level l . The sub-matrix notation in 2D is shown in Equation (5), where the ordering of boxes as shown in Figure 1 is followed. The notation $K_{ab}^{(l)}$ represents the interaction between the source points and target points of nodes b and a of level l respectively. Consider the sub-matrix $K_{01}^{(1)}$. It gets approximated at level 2, as in Equation (7), where only those interactions between boxes which are in each other's interaction list are approximated by a low-rank matrix. In this way, at each level, all the interactions between boxes which are in each other's interaction list are approximated by a low-rank matrix. The resulting low-rank structure of the matrix A at levels 2 and 3 arising in 2D problems is shown in Figure 3.

$$A = K^{(0)} = K^{(1)} = \begin{bmatrix} K_{00}^{(1)} & K_{01}^{(1)} & K_{02}^{(1)} & K_{03}^{(1)} \\ K_{10}^{(1)} & K_{11}^{(1)} & K_{12}^{(1)} & K_{13}^{(1)} \\ K_{20}^{(1)} & K_{21}^{(1)} & K_{22}^{(1)} & K_{23}^{(1)} \\ K_{30}^{(1)} & K_{31}^{(1)} & K_{32}^{(1)} & K_{33}^{(1)} \end{bmatrix} \quad (5)$$

$$K_{01}^{(1)} = \begin{bmatrix} K_{04}^{(2)} & K_{05}^{(2)} & K_{06}^{(2)} & K_{07}^{(2)} \\ K_{14}^{(2)} & K_{15}^{(2)} & K_{16}^{(2)} & K_{17}^{(2)} \\ K_{24}^{(2)} & K_{25}^{(2)} & K_{26}^{(2)} & K_{27}^{(2)} \\ K_{34}^{(2)} & K_{35}^{(2)} & K_{36}^{(2)} & K_{37}^{(2)} \end{bmatrix} \quad (6)$$

$$\approx \begin{bmatrix} U_0^{(2)} A_{04}^{(2)} V_4^{(2)*} & U_0^{(2)} A_{05}^{(2)} V_5^{(2)*} & U_0^{(2)} A_{06}^{(2)} V_6^{(2)*} & U_0^{(2)} A_{07}^{(2)} V_7^{(2)*} \\ K_{14}^{(2)} & U_1^{(2)} K_{15}^{(2)} V_5^{(2)*} & U_1^{(2)} K_{16}^{(2)} V_6^{(2)*} & K_{17}^{(2)} \\ K_{24}^{(2)} & U_2^{(2)} K_{25}^{(2)} V_5^{(2)*} & U_2^{(2)} K_{26}^{(2)} V_6^{(2)*} & K_{27}^{(2)} \\ U_3^{(2)} K_{34}^{(2)} V_4^{(2)*} & U_3^{(2)} K_{35}^{(2)} V_5^{(2)*} & U_3^{(2)} K_{36}^{(2)} V_6^{(2)*} & U_3^{(2)} K_{37}^{(2)} V_7^{(2)*} \end{bmatrix} \quad (7)$$

2.3. Assembly of FMM operators

To assemble the various FMM operators we use a new Nested Cross Approximation (NNCA) [4], a nested version of Adaptive Cross Approximation (ACA), which produces nested bases. The low-rank approximation of an admissible sub-block $A_{t^x s^y}$, with a controlled error of $\mathcal{O}(\epsilon_A)$, using NNCA takes the form

$$A_{t^x s^y} \approx A_{t^x s^x,i} (A_{t^x,i s^x,i})^{-1} A_{t^x,i s^y,o} (A_{t^y,o s^y,o})^{-1} A_{t^y,o s^y} \quad (8)$$




Figure 3: Low-Rank structure of the matrix A at levels 2 and 3 arising in 2D problems

where $t^{X,i}$, $s^{X,i}$, $t^{Y,o}$, and $s^{Y,o}$ are termed pivots, and $t^{X,i} \subset t^X$, $s^{X,i} \subset \mathcal{F}^{X,i}$, $t^{Y,o} \subset \mathcal{F}^{Y,o}$ and $s^{Y,o} \subset s^Y$. $\mathcal{F}^{X,i}$ and $\mathcal{F}^{Y,o}$ are defined as

$$\mathcal{F}^{X,i} = \{s^{X'} : X' \in \mathcal{IL}(X^{(l)})\} \text{ and} \quad (9)$$

$$\mathcal{F}^{Y,o} = \{t^{Y'} : Y' \in \mathcal{IL}(Y^{(l)})\}. \quad (10)$$

For more details on the construction of NNCA, error bounds, and the technique we use in identifying the pivots, we direct the readers to [4, 31]. Here we summarise the various FMM operators that are constructed using NNCA.

$$V_B^{(L)*} = (A_{t^{B,o} s^{B,o}})^{-1} A_{t^{B,o} s^B} \text{ of } B^{(L)} \quad (11)$$

$$V_B^{\dagger(l)*} = (A_{t^{P,o} s^{P,o}})^{-1} A_{t^{P,o} s^{B,o}} \text{ of } B^{(l)} \text{ where } P^{(l)} = \mathcal{P}(B^{(l)}) \quad (12)$$

$$A_{BD}^{(l)} = A_{t^{B,i} s^{D,o}} \text{ of } B^{(l)} \text{ where } D \in \mathcal{IL}(B^{(l)}) \quad (13)$$

$$U_B^{\dagger(l)} = A_{t^{B,i} s^{P,i}} (A_{t^{P,i} s^{P,i}})^{-1} \text{ of } B^{(l)} \quad (14)$$

$$U_B^{(L)} = A_{t^B s^{B,i}} (A_{t^{B,i} s^{B,i}})^{-1} \text{ of } B^{(L)} \quad (15)$$

$$K_{BX}^{(l)} = A_{t^B s^X} \text{ of } B^{(l)} \text{ where } X \in \mathcal{N}(B^{(l)}) \quad (16)$$

where $l \in \{0, 1, 2, \dots, L\}$. We describe each of these operators in Table 3.

$V_B^{(L)*}$	P2M (Particle To Multipole) operator of hypercube $B^{(L)}$ that translates the sources of hypercube $B^{(L)}$ to its multipoles
$V_B^{\dagger(l)*}$	M2M (Multipole To Multipole) operator of hypercube $B^{(l)}$ that translates the multipoles of hypercube $B^{(l)}$ to its parent's multipoles
$A_{BD}^{(l)}$	M2L (Multipole To Local) operator between hypercubes $B^{(l)}$ and $D^{(l)}$ that finds the locals (local potential) of hypercube $B^{(l)}$ due to the multipoles of hypercube $D^{(l)}$
$U_B^{\dagger(l)}$	L2L (Local To Local) operator of hypercube $B^{(l)}$ that translates the locals of its parent to its locals
$U_B^{(L)}$	L2P (Local To Particle) operator of hypercube $B^{(L)}$ that translates its locals to its potential
$K_{BX}^{(l)}$	P2P (Particle To Particle) operator between hypercubes $B^{(l)}$ and $X^{(l)}$ that finds the potential in hypercube $B^{(l)}$ due to the sources in hypercube $X^{(l)}$

Table 3: Various FMM operators

3. The algebraic inverse fast multipole method (AIFMM)

AIFMM has three main steps. The first step is to construct the extended sparse matrix from the given matrix using NNCA. The second step is to perform elimination. The third step is to find the unknowns using back substitution. We now describe each of these in the following subsections.

3.1. Construction of the extended sparse system

The construction of the extended sparse matrix representation of the dense matrix A involves

1. the construction of a 2^d hierarchical tree of depth L , $\mathcal{T}^{(L)}$, as described in Subsection 2.1.
2. the identification of neighbors and interaction list of each hypercube at all levels of the tree, as described in Subsection 2.2.
3. the introduction of auxiliary variables:
 - (a) multipoles at levels $2 \leq l \leq L$, i.e., $\{y^{(l)}\}_{l=2}^L$ where $y^{(l)} = [y_{i_1}^{(l)}; y_{i_2}^{(l)}; \dots; y_{i_{2^d}}^{(l)}]$ and $\{y_{i_c}^{(l)}\}$ indicates the multipoles of hypercube $i_c^{(l)}$.
 - (b) locals at levels $2 \leq l \leq L$, i.e., $\{z^{(l)}\}_{l=2}^L$ where $z^{(l)} = [z_{i_1}^{(l)}; z_{i_2}^{(l)}; \dots; z_{i_{2^d}}^{(l)}]$ and $\{z_{i_c}^{(l)}\}$ indicates the locals of hypercube $i_c^{(l)}$.

Here we followed MATLAB notation to represent the column vectors $y^{(l)}$ and $z^{(l)}$.

Remark 1. For any hypercube at levels 0 and 1, its interaction list is empty, so only the multipoles and locals of hypercubes at levels $2 \leq l \leq L$ are considered to be the auxiliary variables.

3.1.1. Unknown variables of the extended sparse system

For a hypercube $i^{(L)}$, its particles $x_i^{(L)}$ are defined as

$$x_i^{(L)} = x(i^i) \quad (17)$$

For all non-leaf levels, the multipoles at the child level are interpreted to be the particles at the parent level, i.e., we define the particles of a hypercube $B^{(l)}$ where $l < L$, to be the multipoles of its children, as defined in Equation (18).

$$x_B^{(l)} = [y_{B_1}^{(l+1)}; y_{B_2}^{(l+1)}; \dots; y_{B_{2^d}}^{(l+1)}], \quad B_c^{(l+1)} \in C(B^{(l)}) \quad \forall c \in \{1, 2, \dots, 2^d\} \quad (18)$$

Here MATLAB notation is followed to represent the column vector $x_B^{(l)}$. Accordingly, the child M2M operators get combined to form the parent's P2M operator, and similarly, the child L2L operators get combined to form the parent's L2P operator. This has been written in detail in Table 4.

$V_B^{(l)*}$	P2M (Particle To Multipole) operator of hypercube $B^{(l)}$ that translates the particles of hypercube $B^{(l)}$ to its multipoles. $V_B^{(l)*} = [V_{B_1}^{\dagger(l+1)*} \ V_{B_2}^{\dagger(l+1)*} \ \dots \ V_{B_{2^d}}^{\dagger(l+1)*}]$ where $l \in \{0, 1, 2, \dots, L-1\}$, $B_c^{(l+1)} \in C(B^{(l)})$, and $c \in \{1, 2, 3, \dots, 2^d\}$
$U_B^{(l)}$	L2P (Local To Particle) operator of hypercube $B^{(l)}$ that translates the locals of hypercube $B^{(l)}$ to the potential of its children. $U_B^{(l)} = [U_{B_1}^{\dagger(l+1)} \ U_{B_2}^{\dagger(l+1)} \ \dots \ U_{B_{2^d}}^{\dagger(l+1)}]$ where $l \in \{0, 1, 2, \dots, L-1\}$, $B_c^{(l+1)} \in C(B^{(l)})$, and $c \in \{1, 2, 3, \dots, 2^d\}$

Table 4: P2M and L2P operators at non-leaf level

Each hypercube $i^{(l)}$ at level $2 \leq l \leq L$ is therefore associated with the unknown variables described in Table 5.

$x_i^{(l)}$	particles of hypercube $i^{(l)}$
$y_i^{(l)}$	multipoles of hypercube $i^{(l)}$
$z_i^{(l)}$	locals of hypercube $i^{(l)}$

Table 5: Unknown variables of the extended sparse system

3.1.2. Governing equations of the extended sparse system

The equations governing the multipoles and the potential are given below.

At the leaf level,

$$y_i^{(L)} = V_i^{(L)*} x_i^{(L)} \quad \text{denoted as Equation } Y_i^{(L)} \quad (19)$$

$$b_i^{(L)} = U_i^{(L)} z_i^{(L)} + \sum_{j^{(L)} \in \mathcal{N}(i^{(L)})} K_{ij}^{(L)} x_j^{(L)} \quad \text{denoted as Equation } X_i^{(L)} \quad (20)$$

where $b_i^{(L)} = [b_{c_1}, b_{c_2}, \dots, b_{c_{m_i}}]$ and $\{c_e\}_{e=1}^{m_i} = i^i$.

For $l \in \{0, \dots, L-1\}$,

$$y_i^{(l)} = \sum_{i'^{(l)} \in \mathcal{C}(i^{(l)})} V_{i'}^{(l)*} y_{i'}^{(l+1)} \quad (21a)$$

$$= V_i^{(l)*} x_i^{(l)} \quad \text{denoted as Equation } Y_i^{(l)} \quad (21b)$$

wherein we have combined the multipoles at a child level to form the particles at the parent level as in Equation (18).

For $l \in \{0, 1, \dots, L\}$, the equation governing the locals is given by

$$z_i^{(l)} = U_i^{(l)} z_i^{(l-1)} + \sum_{j^{(l)} \in \mathcal{IL}(i^{(l)})} A_{ij}^{(l)} y_j^{(l)} \quad (22)$$

where $i'^{(l-1)} = \mathcal{P}(i^{(l)})$.

For $l \in \{1, 2, \dots, L-1\}$, Equation (22) takes the form of Equation (20), written out in Equation (23), when the multipoles at level l are combined to form the particles at level $l-1$ as in Equation (18).

$$b_i^{(l)} = U_i^{(l)} z_i^{(l)} + \sum_{j^{(l)} \in \mathcal{N}(i^{(l)})} K_{ij}^{(l)} x_j^{(l)} \quad \text{denoted as Equation } X_i^{(l)} \quad (23)$$

Here $z_i^{(1)} = \vec{0}$, $U_i^{(l)}$ is defined in Table 3 and

$$K_{ij}^{(l)} = \begin{bmatrix} A_{i_1 j_1}^{(l+1)} & A_{i_1 j_2}^{(l+1)} & \dots & A_{i_1 j_{2d}}^{(l+1)} \\ A_{i_2 j_1}^{(l+1)} & A_{i_2 j_2}^{(l+1)} & \dots & A_{i_2 j_{2d}}^{(l+1)} \\ \vdots & \ddots & \ddots & \vdots \\ A_{i_{2d} j_1}^{(l+1)} & A_{i_{2d} j_2}^{(l+1)} & \dots & A_{i_{2d} j_{2d}}^{(l+1)} \end{bmatrix}, \quad b_i^{(l)} = \begin{bmatrix} z_{i_1}^{(l+1)} \\ z_{i_2}^{(l+1)} \\ \vdots \\ z_{i_{2d}}^{(l+1)} \end{bmatrix}. \quad (24)$$

Summary of equations $\forall l \in \{1, 2, \dots, L\}$

$$y_i^{(l)} = V_i^{(l)*} x_i^{(l)} \quad \text{denoted as Equation } Y_i^{(l)} \quad (25)$$

$$b_i^{(l)} = U_i^{(l)} z_i^{(l)} + \sum_{j^{(l)} \in \mathcal{N}(i^{(l)})} K_{ij}^{(l)} x_j^{(l)} \quad \text{denoted as Equation } X_i^{(l)} \quad (26)$$

The various FMM operators: L2L/L2P, M2M/P2M, and M2L are obtained using NNCA as described in Subsection 2.3.

The system of equations with the unknowns x and the auxiliary variables is given in Equation (27).

$$\tilde{A} \begin{bmatrix} x \\ z^{(L)} \\ y^{(L)} \\ z^{(L-1)} \\ y^{(L-1)} \\ \vdots \\ z^{(l)} \\ y^{(l)} \\ \vdots \\ z^{(2)} \\ y^{(2)} \end{bmatrix} = \begin{bmatrix} b \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix} \quad (27)$$

The ordering of equations or rows is given by

$$\begin{aligned} &\{X_{i_1}^{(L)}, Y_{i_1}^{(L)}, X_{i_2}^{(L)}, Y_{i_2}^{(L)}, \dots, X_{i_{2dL}}^{(L)}, Y_{i_{2dL}}^{(L)}, \\ &X_{i_1}^{(L-1)}, Y_{i_1}^{(L-1)}, X_{i_2}^{(L-1)}, Y_{i_2}^{(L-1)}, \dots, X_{i_{2d(L-1)}}^{(L-1)}, Y_{i_{2d(L-1)}}^{(L-1)}, \\ &\dots, \\ &X_{i_1}^{(2)}, Y_{i_1}^{(2)}, X_{i_2}^{(2)}, Y_{i_2}^{(2)}, \dots, X_{i_{2d}}^{(2)}, Y_{i_{2d}}^{(2)}, \\ &X_{i_1}^{(1)}, X_{i_2}^{(1)}, \dots, X_{i_{2d}}^{(1)}\}. \end{aligned}$$

A reordering of unknowns or columns of Equation (27) is performed, such that $z^{(l)}$ is interleaved in between $x^{(l)}$ for all levels l from L to 2 as follows

$$\begin{aligned} &[x_{i_1}^{(L)*} z_{i_1}^{(L)*} x_{i_2}^{(L)*} z_{i_2}^{(L)*} \dots x_{i_{2dL}}^{(L)*} z_{i_{2dL}}^{(L)*} \\ &x_{i_1}^{(L-1)*} z_{i_1}^{(L-1)*} x_{i_2}^{(L-1)*} z_{i_2}^{(L-1)*} \dots x_{i_{2d(L-1)}}^{(L-1)*} z_{i_{2d(L-1)}}^{(L-1)*} \\ &\dots \\ &x_{i_1}^{(2)*} z_{i_1}^{(2)*} x_{i_2}^{(2)*} z_{i_2}^{(2)*} \dots x_{i_{2d}}^{(2)*} z_{i_{2d}}^{(2)*} \\ &x_{i_1}^{(1)*} x_{i_2}^{(1)*} \dots x_{i_{2d}}^{(1)*}]^*. \end{aligned}$$

It ensures that, when an elimination in standard ordering is performed, the fill-ins occur symmetrically. Let the new system after reordering be

$$\tilde{A}\tilde{x} = \tilde{b} \quad (28)$$

The structures of the matrices \tilde{A} constructed out of a 2D problem at levels 2 and 3 are illustrated in Figures 4 and 5 respectively.

3.2. Elimination or the Factorization Phase

Given the extended sparse system of Equations (28), the next task is to solve for the unknowns \tilde{x} . It is solved using Gaussian Elimination followed by Back Substitution. The elimination phase is not the naive Gaussian Elimination but the elimination process is interleaved with the compression and redirection of fill-ins corresponding to well-separated hypercubes through non-zero entries. In this subsection, we describe this process.

We show in Figures 6 and 8, the graphical representation of the extended sparse matrix, constructed out of a 2D problem at levels 2 and 3 respectively. For better clarity, we show in Figure 7 a partial graph of the graph shown in Figure 8. The nodes correspond to the variables and the incoming edges to a node constitute an equation. The same color notation followed in Figures 4 and 5 is followed in Figures 6, 7 and 8.




Figure 4: Structure of the extended sparse matrix at level 2 arising in 2D problems




Figure 5: Structure of the extended sparse matrix at level 3 arising in 2D problems




Figure 6: Graph of the Extended Sparse system at level 2




Figure 7: Partial graph of the Extended Sparse system at level 3, showing few nodes and edges for better readability




Figure 8: Graph of the Extended Sparse system at level 3

We eliminate the variables in standard ordering, i.e., the order in which the variables are arranged. When variables $x_i^{(l)}$ and $z_i^{(l)}$ get eliminated, which we term as the hypercube getting eliminated, it results in an update of the graph that involves nullification of some edges, updation of some edges and creation of new edges. The new edges that get created are termed fill-ins. The various fill-ins that get created are described in Table 6.

$P2P_{ij}^{(l)}$	The P2P fill-in that represents the potential of hypercube $i^{(l)}$ due to the particles of hypercube $j^{(l)}$
$M2P_{ij}^{(l)}$	The M2P fill-in that represents the potential of hypercube $i^{(l)}$ due to the multipoles of hypercube $j^{(l)}$
$P2L_{ij}^{(l)}$	The P2L fill-in that represents the local potential of hypercube $i^{(l)}$ due to the particles of hypercube $j^{(l)}$
$M2L_{ij}^{(l)}$	The M2L fill-in that represents the local potential of hypercube $i^{(l)}$ due to the multipoles of hypercube $j^{(l)}$

Table 6: Fill-in interactions

Upon elimination of hypercube $i^{(l)}$, fill-ins among its neighboring hypercubes get created as described below. For $p^{(l)}, q^{(l)} \in \mathcal{N}(i^{(l)})$,

1. if $p^{(l)}$ and $q^{(l)}$ have been eliminated, $M2L_{pq}^{(l)}$ and $M2L_{qp}^{(l)}$ get created.
2. if $p^{(l)}$ has been eliminated and $q^{(l)}$ has not been eliminated, fill-ins $P2L_{pq}^{(l)}$ and $M2P_{qp}^{(l)}$ get created.
3. if $q^{(l)}$ has been eliminated and $p^{(l)}$ has not been eliminated, fill-ins $P2L_{qp}^{(l)}$ and $M2P_{pq}^{(l)}$ get created.
4. if $p^{(l)}$ and $q^{(l)}$ have not been eliminated, $P2P_{pq}^{(l)}$ and $P2P_{qp}^{(l)}$ get created.

For a more detailed understanding of the fill-in creation, we refer the readers to the graphs in [1].

Theorem 1. Consider a hypercube $i^{(L)}$ and hypercubes $\{j^{(L)}, k^{(L)}\} \in \mathcal{N}(i^{(L)})$ such that $j^{(L)} \in \mathcal{I}\mathcal{L}(k^{(L)})$. Let N target points and N source points be distributed uniformly in each of the hypercubes. If the hypercubes $j^{(L)}$ and $k^{(L)}$ are not eliminated from the extended sparse system, then the P2P fill-in $P2P_{jk}^{(L)}$ that gets created upon elimination of hypercube $i^{(L)}$ is rank deficient.

Proof. Upon elimination of hypercube $i^{(L)}$ from the extended sparse system, the P2P fill-in $P2P_{jk}^{(L)}$ gets created as $P2P_{jk}^{(L)} := -K_{ji}^{(L)} K_{ii}^{(L)} K_{ik}^{(L)}$. The rank of $K_{ii}^{(L)}$ is N , as it is a self interaction. From [19], the rank of interaction between particles of hypercubes that (i) share a vertex scales as $O(\log(N) \log^d(\log(N)))$; (ii) share a hypersurface of dim d' scales as $O(N^{\frac{d'}{d}} \log^d(N))$, $d' \in \{1, 2, \dots, d-1\}$.

- If at least one of the hypercubes $j^{(L)}, k^{(L)}$ shares a vertex with hypercube $i^{(L)}$ then

$$\text{rank}(P2P_{jk}^{(L)}) \leq \min\{\text{rank}(K_{ji}^{(L)}), \text{rank}(K_{ii}^{(L)}), \text{rank}(K_{ik}^{(L)})\} \quad (29)$$

$$\leq O(\log(N) \log^d(\log(N))). \quad (30)$$

- If $j^{(L)}$ and $k^{(L)}$ share a hypersurface of dim $d_j \in \{1, 2, \dots, d-1\}$ and $d_k \in \{1, 2, \dots, d-1\}$ with hypercube $i^{(L)}$ respectively, then

$$\text{rank}(P2P_{jk}^{(L)}) \leq \min\{\text{rank}(K_{ji}^{(L)}), \text{rank}(K_{ii}^{(L)}), \text{rank}(K_{ik}^{(L)})\} \quad (31)$$

$$\leq O(N^{\frac{d'}{d}} \log^d(N)) \quad (32)$$

where $d' = \min\{d_j, d_k\}$ and $d' \in \{1, 2, \dots, d-1\}$.

□

We show in Theorem 1, under the assumption that the particles are uniformly distributed, that the P2P fill-ins corresponding to well-separated hypercubes at leaf level are rank deficient. We assume that this is true at higher levels

as well and also when the particles are distributed non-uniformly. Further, the bounds obtained in Theorem 1 are very conservative, as the numerical illustrations in [1] show that the ranks are almost constant.

The ranks of P2L and M2P fill-ins do not scale with N , as they are equal to the number of locals and the number of multiples respectively.

In conclusion, a fill-in corresponding to an interaction between well-separated hypercubes is low-rank and therefore can be efficiently approximated by a low-rank matrix. Further, the compression is redirected through existing operators as described later in the section.

In the process of elimination, due to the creation of fill-ins, and due to the compression and redirection of fill-ins corresponding to well-separated hypercubes, Equation (26) gets modified as

$$b_i^{(l)} = U_i^{(l)} z_i^{(l)} + \sum_{j^{(l)} \in \mathcal{N}(i^{(l)})} \left((1 - E_j^{(l)}) K_{ij}^{(l)} x_j^{(l)} + E_j^{(l)} M2P_{ij}^{(l)} y_j^{(l)} \right) \quad (33)$$

where $E_j^{(l)}$ takes values 0 or 1. It being 1, indicates that node $j^{(l)}$ is eliminated and 0, indicates that node $j^{(l)}$ is not eliminated. We continue the process of elimination until when the multipoles at level 2 are the only variables left. This entire elimination process is described in Algorithm 1.

3.2.1. Compression and redirection of P2P fill-in

Consider a P2P fill-in $P2P_{ij}$ where hypercubes i and j at level l are well-separated. $P2P_{ij}$ can be efficiently approximated by a low-rank matrix and this interaction can be redirected through an already existing interaction via the path $x_j \rightarrow z_j \rightarrow y_j \rightarrow y_i \rightarrow z_i \rightarrow x_i$ as shown in Figure 9. This redirection results in an update of (i) P2M V_j^* ; (ii) M2L A_{ij} ; (iii) L2P U_i ; (iv) M2M $V_j^{\dagger*}$; (v) L2L U_i^\dagger ; (vi) Other M2Ls $\{A_{ic} : c \in \{\mathcal{IL}(i) \setminus j\} \cup \mathcal{N}(i)\}$ and $\{A_{dj} : d \in \{\mathcal{IL}(j)\} \cup \mathcal{N}(j)\}$. We now describe how each of these updates is done.

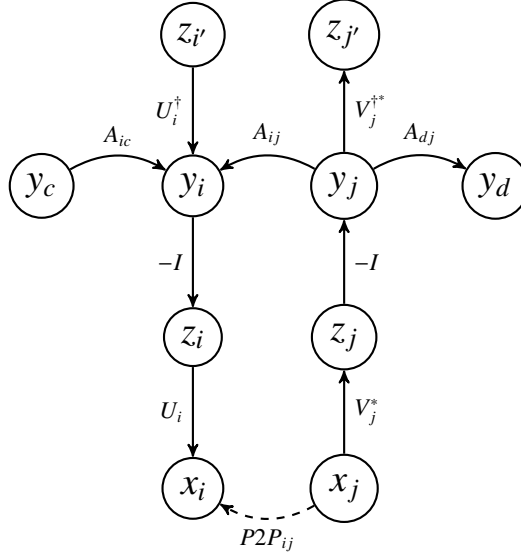


Figure 9: Illustration of $P2P_{ij}$ compression and its redirection. The fill-in is shown through a dashed edge. Here $i' = \mathcal{P}(i)$ and $j' = \mathcal{P}(j)$. (It is to be noted that a part of the graph, with only a few connections that get updated as a result of the fill-in redirection, is shown here.)

P2M, M2L, L2P update. First, we find the new orthogonal column basis, that spans the existing column basis of i at level l , i.e., U_i , and the columns of the fill-in $P2P_{ij}$ by finding the RRQR decomposition of the augmented matrix

$$[U_i | P2P_{ij}] = \tilde{U}_i L_{ij}. \quad (34)$$

A threshold ϵ_A is used as an input to the RRQR routine, such that the relative residual in the decomposition is equal

Algorithm 1 Elimination algorithm

```

1: procedure ELIMINATION( $n_{\max}, \epsilon_A$ )
2:                                      $\triangleright n_{\max}$  is the maximum number of particles at leaf level
3:   Form  $\mathcal{T}^L$ , where  $L = \min \{l : |x_i^{(l)}| < n_{\max}; \forall \text{ hypercubes } i \text{ at level } l\}$ 
4:   Perform NNCA with tolerance  $\epsilon_A$  to find the  $L2L/L2P, M2L, M2M/P2M$  operators of all hypercubes at all
   levels
5:   for  $l = L : 2$  do
6:     for  $i = e_1^{(l)}, e_2^{(l)}, \dots, e_{2^{dl}}^{(l)}$  do
7:       Eliminate  $x_i^{(l)}$  and  $z_i^{(l)}$  from the extended sparse system using Equations (33) and (25)
8:        $E_i^{(l)} := 1$ 
9:       for  $(p^{(l)}, q^{(l)})$  in  $\{(r^{(l)}, s^{(l)}) : r^{(l)} \in \mathcal{N}(i^{(l)}), s \in \mathcal{N}(i^{(l)})\}$  do
10:        if  $p^{(l)}$  is eliminated then
11:          if  $q^{(l)}$  is eliminated then
12:            Results in the update of  $M2L_{pq}^{(l)}$  and  $M2L_{qp}^{(l)}$ 
13:          else
14:            Results in fill-ins  $P2L_{pq}^{(l)}$  and  $M2P_{qp}^{(l)}$ 
15:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
16:              Compress  $P2L_{pq}^{(l)}$  and  $M2P_{qp}^{(l)}$  and update the relevant operators as in Subsections 3.2.2
              and 3.2.3.
17:            end if
18:          end if
19:        else
20:          if  $q^{(l)}$  is eliminated then
21:            Results in fill-ins  $P2L_{qp}^{(l)}$  and  $M2P_{pq}^{(l)}$ 
22:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
23:              Compress  $P2L_{qp}^{(l)}$  and  $M2P_{pq}^{(l)}$  and update the relevant operators as in Subsections 3.2.2
              and 3.2.3.
24:            end if
25:          else
26:            Results in the update of  $P2P_{pq}^{(l)}$  and  $P2P_{qp}^{(l)}$ 
27:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
28:              Compress  $P2P_{pq}^{(l)}$  and  $P2P_{qp}^{(l)}$  and update the relevant operators as in Subsection 3.2.1
29:            end if
30:          end if
31:        end if
32:      end for
33:    end for
34:  end for
35: end procedure

```

to $O(\epsilon_A)$. Considering the matrix L_{ij} to be an augmented matrix of the form $L_{ij} \equiv [L_i | \tilde{L}_{ij}]$, results in

$$U_i = \tilde{U}_i L_i \quad (35a)$$

$$P2P_{ij} = \tilde{U}_i \tilde{L}_{ij}, \quad (35b)$$

Next, we find the orthogonal row basis, that spans the existing row basis of j at level l , i.e., V_j^* , and a row basis of the fill-in $P2P_{ij}$, \tilde{L}_{ij} , by finding the RRQR decomposition of the augmented matrix

$$[V_j | \tilde{L}_{ij}^*] = \tilde{V}_j R_{ij} \quad (36)$$

By expressing R_{ij} as the augmented matrix $[R_j | \tilde{R}_{ij}]$, we have

$$V_j = \tilde{V}_j R_j, \quad (37a)$$

$$\tilde{L}_{ij}^* = \tilde{V}_j \tilde{R}_{ij}, \quad (37b)$$

Using Equations (35b) and (37b), we have,

$$P2P_{ij} = \tilde{U}_i \tilde{L}_{ij} = \tilde{U}_i \tilde{R}_{ij}^* \tilde{V}_j^*. \quad (38)$$

Using Equations (38), (35a), and (37a), $U_i A_{ij} V_j^* + P2P_{ij}$ can be expressed as

$$U_i A_{ij} V_j^* + P2P_{ij} = \tilde{U}_i L_i A_{ij} R_j^* \tilde{V}_j^* + \tilde{U}_i \tilde{R}_{ij}^* \tilde{V}_j^* \quad (39)$$

$$= \tilde{U}_i (L_i A_{ij} R_j^* + \tilde{R}_{ij}^*) \tilde{V}_j^*. \quad (40)$$

We then make the following assignments, which update the old operators with the new ones.

$$A_{ij} := L_i A_{ij} R_j^* + \tilde{R}_{ij}^* \quad (41)$$

$$U_i := \tilde{U}_i \quad (42)$$

$$V_j := \tilde{V}_j \quad (43)$$

Other M2L updates. For $c \in \{\mathcal{IL}(i) \setminus j\} \cup \mathcal{N}(i)$, the value of the old potential due to y_c at particles x_i should be equal to the value of the new potential due to y_c at particles x_i , as in equation 44, because the potential due to y_c at x_i is not dependent on the fill-in between hypercubes i and j at level l .

$$U_i A_{ic} y_c = \tilde{U}_i \tilde{A}_{ic} y_c \quad (44)$$

As Equation (44) holds true $\forall y_c \in \mathbb{C}^{k \times 1}$, it can be equivalently written as

$$U_i A_{ic} = \tilde{U}_i \tilde{A}_{ic} \quad (45)$$

Further, since $\tilde{U}_i^* \tilde{U}_i = I$, Equation (45), can be written as

$$\tilde{A}_{ic} = \tilde{U}_i^* U_i A_{ic} = L_i A_{ic}. \quad (46)$$

We then make the following assignment, which updates the old operator with the new one.

$$A_{ic} := \tilde{A}_{ic} \quad (47)$$

Similarly, for $d \in \{\mathcal{IL}(j)\} \cup \mathcal{N}(j)$, the value of the old locals due to x_j should be equal to the value of the new locals due to x_j , because the locals of d due to particles of j is not dependent on the fill-in between i and j .

$$A_{dj} V_j^* x_j = \tilde{A}_{dj} \tilde{V}_j^* x_j \quad (48)$$

As Equation (48) holds true $\forall x_j \in \mathbb{C}^{k \times 1}$, it can be equivalently written as

$$A_{dj} V_j^* = \tilde{A}_{dj} \tilde{V}_j^* \quad (49)$$

Further, since $\tilde{V}_j^* \tilde{V}_j = I$, Equation (49), can be written as

$$\tilde{A}_{dj} = A_{dj} V_j^* \tilde{V}_j = A_{dj} R_j^*. \quad (50)$$

We then make the following assignment, which updates the old operator with the new one.

$$A_{dj} := \tilde{A}_{dj} \quad (51)$$

M2M update. The old and new multipoles of hypercube j at level l are given by

$$y_j = V_j^* x_j, \quad (52)$$

$$\tilde{y}_j = \tilde{V}_j^* x_j \quad (53)$$

respectively. The fill-in $P2P_{ij}$ has no influence on the multipoles $y_{j'}$. So, the old and the new contribution of the multipoles of j at the multipoles of its parent j' must be equal and hence it follows that

$$V_j^{\dagger*} y_j = \tilde{V}_j^{\dagger*} \tilde{y}_j \quad (54)$$

Using Equations (52), (53), and, (54)

$$V_j^{\dagger*} V_j^* x_j = \tilde{V}_j^{\dagger*} \tilde{V}_j^* x_j \quad (55)$$

As Equation (55) holds true $\forall x_j \in \mathbb{C}^{k \times 1}$, it can be equivalently written as

$$V_j^{\dagger*} V_j^* = \tilde{V}_j^{\dagger*} \tilde{V}_j^* \quad (56)$$

Further, since $\tilde{V}_j^* \tilde{V}_j = I$, Equation (56) can be written as

$$\tilde{V}_j^{\dagger*} = V_j^{\dagger*} V_j^* \tilde{V}_j = V_j^{\dagger*} R_j^*. \quad (57)$$

We then make the following assignment, which updates the old operator with the new one.

$$V_j^{\dagger*} := \tilde{V}_j^{\dagger*} \quad (58)$$

L2L update. A similar analysis as done in updating the M2M on the L2L operator results in its update as follows:

$$U_i^\dagger := \tilde{U}_i^\dagger = \tilde{U}_i^* U_i U_i^\dagger = L_i U_i^\dagger. \quad (59)$$

3.2.2. Compression and redirection of P2L fill-in

Consider a P2L fill-in $P2L_{ij}$ where hypercubes i and j at level l are well-separated. Then $P2L_{ij}$ can be efficiently approximated by a low-rank matrix and this interaction can be redirected through an already existing interaction via the path $x_j \rightarrow z_j \rightarrow y_j \rightarrow y_i$, as shown in the Figure 10. This redirection results in an update of (i) P2M V_j^* ; (ii) M2L A_{ij} ; (iii) M2M $V_j^{\dagger*}$; (iv) Other M2Ls $\{A_{dj} : d \in \{\mathcal{IL}(j) \setminus i\} \cup \mathcal{N}(j)\}$. We now describe how each of these updates is done.




Figure 10: Illustration of $P2L_{ij}$ compression and its redirection. The fill-in is shown through a dashed edge. Here $j' = \mathcal{P}(j)$. It is to be noted that a part of the graph, with only a few connections that get updated as a result of the fill-in redirection, is shown here.

P2M, M2L update. We find the new orthogonal row basis, that spans the existing row basis of j at level l , i.e., V_j^* , and the rows of the fill-in $P2L_{ij}$, by finding the RRQR decomposition of the augmented matrix

$$[V_j | P2L_{ij}^*] = \tilde{V}_j R_{ij} \quad (60)$$

By expressing R_{ij} as the augmented matrix $[R_j | \tilde{R}_{ij}]$, we have

$$V_j = \tilde{V}_j R_j, \quad (61a)$$

$$P2L_{ij}^* = \tilde{V}_j \tilde{R}_{ij}, \quad (61b)$$

Using Equations (61a) and, (61b), $A_{ij}V_j^* + P2L_{ij}$ can be expressed as

$$A_{ij}V_j^* + P2L_{ij} = A_{ij}R_j^* \tilde{V}_j^* + \tilde{R}_{ij}^* \tilde{V}_j^* \quad (62)$$

$$= (A_{ij}R_j^* + \tilde{R}_{ij}^*) \tilde{V}_j^*. \quad (63)$$

We then make the following assignments, which update the old operators with the new ones.

$$A_{ij} := A_{ij}R_j^* + \tilde{R}_{ij}^* \quad (64)$$

$$V_j := \tilde{V}_j \quad (65)$$

Other M2L updates. As a result of the redirection of the fill-in $P2L_{ij}$, the M2Ls A_{dj} where $d \in \{\mathcal{IL}(j) \setminus i\} \cup \mathcal{N}(j)\}$ get updated. The updates follow the same lines described in Subsubsection 3.2.1.

M2M update. As a result of the redirection of the fill-in $P2L_{ij}$, M2M $V_j^{\dagger*}$ gets updated. The updates follow the same lines described in Subsubsection 3.2.1.

3.2.3. Compression and redirection of M2P fill-in

Consider a M2P fill-in $M2P_{ij}$ where hypercubes i and j at level l are well-separated. Then $M2P_{ij}$ can be efficiently approximated by a low-rank matrix and this interaction can be redirected through an already existing interaction via the path $y_j \rightarrow y_i \rightarrow z_i \rightarrow x_i$, as shown in the Figure 11. This redirection results in an update of (i) L2P U_i ; (ii) M2L A_{ij} ; (iii) L2L U_i^\dagger ; (iv) Other M2Ls $\{A_{ic} : c \in \{\mathcal{IL}(i) \setminus j\} \cup \mathcal{N}(i)\}$. We now describe how each of these updates is done.




Figure 11: Illustration of $M2P_{ij}$ compression and its redirection. The fill-in is shown through a dashed edge. Here $i' = \mathcal{P}(i)$. It is to be noted that a part of the graph, with only a few connections that get updated as a result of the fill-in redirection, is shown here.

L2P, M2L update. We find the new orthogonal column basis, that spans the existing column basis of i at level l i.e., U_i , and the columns of the fill-in $M2P_{ij}$, by finding the RRQR decomposition of the augmented matrix

$$[U_i | M2P_{ij}] = \tilde{U}_i L_{ij} \quad (66)$$

By expressing L_{ij} as the augmented matrix, $[L_i | \tilde{L}_{ij}]$, we have

$$U_i = \tilde{U}_i L_i, \quad (67a)$$

$$M2P_{ij} = \tilde{U}_i \tilde{L}_{ij}, \quad (67b)$$

Using Equations (67a) and (67b), $U_i A_{ij} + M2P_{ij}$ can be expressed as

$$U_i A_{ij} + M2P_{ij} = \tilde{U}_i L_i A_{ij} + \tilde{U}_i \tilde{L}_{ij} \quad (68)$$

$$= \tilde{U}_i (L_i A_{ij} + \tilde{L}_{ij}). \quad (69)$$

We then make the following assignments, which update the old operators with the new ones.

$$A_{ij} := L_i A_{ij} + \tilde{L}_{ij} \quad (70)$$

$$U_i := \tilde{U}_i \quad (71)$$

Other M2L updates. As a result of the redirection of the fill-in $M2P_{ij}$, the M2Ls A_{ic} where $c \in \{\mathcal{I}\mathcal{L}(i) \setminus j\} \cup \mathcal{N}(i)$ get updated. The updates follow the same lines as described in Subsubsection 3.2.1.

L2L update. As a result of the redirection of the fill-in $M2P_{ij}$, the L2L U_i^\dagger gets updated. The updates follow the same lines described in Subsubsection 3.2.1.

3.3. A more efficient elimination algorithm

For $d > 1$, a fill-in corresponding to well-separated hypercubes, say $P2P_{pq}$, could get created or updated multiple times during the elimination process. It is because there could be many hypercubes $i^{(l)}$ such that hypercubes $p^{(l)}, q^{(l)} \in \mathcal{N}(i^{(l)})$. To avoid the compression and redirection multiple times, we choose not to compress and redirect as and when a fill-in gets created as in Algorithm 1, but to update the fill-ins multiple times and compress and redirect only once, just before either $p^{(l)}$ or $q^{(l)}$ gets eliminated as in Algorithm 2.

In Algorithm 2 vectors $vP2P$ and $vP2L$ are used to keep track of the fill-ins corresponding to well-separated hypercubes. For the fill-ins $P2P_{pq}^{(l)}$ and $P2P_{qp}^{(l)}$ where $p^{(l)} \in \mathcal{IL}(q^{(l)})$, only one ordered pair $(p^{(l)}, q^{(l)})$ is stored in $vP2P$ as they always occur in a pair. Similarly for the fill-ins $P2L_{pq}^{(l)}$ and $M2P_{qp}^{(l)}$ where $p^{(l)} \in \mathcal{IL}(q^{(l)})$, only one ordered pair $(p^{(l)}, q^{(l)})$ is stored in $vP2L$. Before a node $i^{(l)}$ gets eliminated, the vectors $vP2P$ and $vP2L$ are searched for an ordered pair with $i^{(l)}$ as one of its elements. If it exists then the associated fill-ins are compressed and redirected.

3.4. Back Substitution or Solve phase

The third step of AIFMM is the back substitution or solve phase, wherein we solve for the multipoles at level 2 and then find the unknowns by back substitution. The pseudo-code is described in Algorithm 3.

Remark 2. *The elimination process, similar to the factorize phase in a direct solver, can be decoupled from the right-hand side. So the elimination phase can be considered as the factorize phase and the back substitution phase can be considered as the solve phase.*

4. Numerical Results

We perform a total of five experiments to demonstrate the performance of AIFMM as a direct solver and as a preconditioner.

In Experiment 1, the validation, convergence and various benchmarks of AIFMM are presented. In Experiments 2 to 4, AIFMM is compared with HODLR [14, 34], a direct solver, and with GMRES [35, 36], an iterative solver. HODLR solver hierarchically partitions the matrix and constructs low-rank approximations of the off-diagonal blocks to a user-specified tolerance ϵ_H .

In Experiment 5, AIFMM is demonstrated as a preconditioner. GMRES with AIFMM as preconditioner is compared with i) GMRES with no preconditioner ii) GMRES with HODLR as preconditioner iii) block-diagonal preconditioner. HODLR and AIFMM are used as preconditioners by constructing low-accuracy direct solvers, i.e., a high value of ϵ_H and ϵ_A are used respectively.

GMRES involves the computation of a matrix-vector product in each of its iterations. In Experiments 2 to 4, this computation is accelerated using NNCA-based fast \mathcal{H}^2 matrix-vector product, described in [4]. While in Experiment 5, where we solve the high frequency scattering problem, we use NNCA-based Directional Algebraic Fast Multipole Method (DAFMM), described in [33]. Let the compression tolerance of these fast summation techniques be denoted by ϵ_G .

All experiments were carried out on an Intel Xeon 2.5GHz processor. In Experiments 1 to 4, we solve for x , in $Ax = b$, where

- b is considered to be a random vector and
- the particles $\{u_i\}_{i=1}^N$ and $\{v_i\}_{i=1}^N$ are considered to be same and are distributed uniformly in the domain $[-1, 1]^2$.

Before presenting the experiments, we describe some notations that are used in this section in Table 7.

4.1. Experiment 1: Validation and convergence of AIFMM

Here we consider the 2D Helmholtz function with the wavenumber set to 1. To have a well-conditioned matrix, we consider the entries of the matrix to be

$$A_{i,j} = \begin{cases} \sqrt{1000N} & \text{if } i = j \\ \frac{\epsilon}{4} H_0^{(1)}(\|x_i - x_j\|_2) & \text{else} \end{cases}. \quad (72)$$

We plot r_m , assembly time, factorization time, solve time, and relative error versus N in Figure 12 for various values of ϵ_A . The following inferences are to be noticed from the figure.

1. The relative error for a given ϵ_A is almost constant as N increases.
2. The relative error decreases as ϵ_A decreases, which validates the convergence of AIFMM.
3. Assembly time, solve time, and factorization time scale linearly with N .

Algorithm 2 Efficient Elimination algorithm

```
1: procedure EFFICIENT_ELIMINATION( $n_{\max}, \epsilon_A$ )
2:   Form  $\mathcal{T}^L$ , where  $L = \min \{l : |x_i^{(l)}| < n_{\max}, \forall \text{ hypercubes } i \text{ at level } l\}$ 
3:   Perform NNCA with tolerance  $\epsilon_A$  to find the  $L2L/L2P, M2L, M2M/P2M$  operators of all hypercubes at all
   levels
4:   Declare sets  $vP2P$  and  $vP2L$  that holds integer ordered pairs
5:   for  $l = L : 2$  do
6:     for  $i = e_1^{(l)}, e_2^{(l)}, \dots, e_{2^{dl}}^{(l)}$  do
7:       for  $(r^{(l)}, s^{(l)})$  in  $vP2P$  do
8:         if  $(i == r^{(l)} \parallel i == s^{(l)})$  then
9:           Compress  $P2P_{rs}^{(l)}$  and  $P2P_{sr}^{(l)}$  and update the relevant operators as in Subsection 3.2.1.
10:          Erase  $(r^{(l)}, s^{(l)})$  in  $vP2P$ 
11:        end if
12:      end for
13:      for  $(r^{(l)}, s^{(l)})$  in  $vP2L$  do
14:        if  $(i == s^{(l)})$  then
15:          Compress  $P2L_{rs}^{(l)}$  and  $M2P_{sr}^{(l)}$  and update the relevant operators as in Subsections 3.2.2 and 3.2.3
16:          respectively.
17:          Erase  $(r^{(l)}, s^{(l)})$  in  $vP2L$ 
18:        end if
19:      end for
20:      Eliminate  $x_i^{(l)}$  and  $z_i^{(l)}$  from the extended sparse system using Equations (33) and (25)
21:       $E_i^{(l)} := 1$ 
22:      for  $(p^{(l)}, q^{(l)})$  in  $\{(r^{(l)}, s^{(l)}) : r^{(l)} \in \mathcal{N}(i^{(l)}), s \in \mathcal{N}(i^{(l)})\}$  do
23:        if  $p^{(l)}$  is eliminated then
24:          if  $q^{(l)}$  is eliminated then
25:            Results in the update of  $M2L_{pq}^{(l)}$  and  $M2L_{qp}^{(l)}$ 
26:          else
27:            Results in the update of  $P2L_{pq}^{(l)}$  and  $M2P_{qp}^{(l)}$ 
28:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
29:               $vP2L.\text{push\_back}((p^{(l)}, q^{(l)}))$ 
30:            end if
31:          end if
32:        else
33:          if  $q^{(l)}$  is eliminated then
34:            Results in the update of  $P2L_{qp}^{(l)}$  and  $M2P_{pq}^{(l)}$ 
35:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
36:               $vP2L.\text{push\_back}((q^{(l)}, p^{(l)}))$ 
37:            end if
38:          else
39:            Results in the update of  $P2P_{pq}^{(l)}$  and  $P2P_{qp}^{(l)}$ 
40:            if  $p^{(l)}$  and  $q^{(l)}$  are well-separated then
41:               $vP2P.\text{push\_back}((p^{(l)}, q^{(l)}))$ 
42:            end if
43:          end if
44:        end for
45:      end for
46:    end for
47: end procedure
```

Algorithm 3 Back Substitution Algorithm

```

1: procedure BACK_SUBSTITUTION
2:   Solve for the multipoles at level 2,  $y^{(2)}$ , directly
3:   for  $l = 2 : L$  do
4:     for  $i = e_{2^{dl}}^{(l)}, e_{2^{dl}-1}^{(l)}, \dots, e_1^{(l)}$ , do
5:       Find  $x_i^{(l)}$  and  $z_i^{(l)}$  by back substitution using Equations (33) and (25).
6:        $E_i^{(l)} := 0$ 
7:       Find  $\{y_{i_c}^{(l+1)}\}_{c=1}^{2^d}$  from  $x_i^{(l)}$  using Equation (18).
8:     end for
9:   end for
10: end procedure

```

N	System size that denotes the number of particles in the domain.
ϵ_A	Tolerance set for NNCA and RRQR, of AIFMM.
r_m	Maximum rank of the compressed blocks, which includes the interactions and the fill-ins corresponding to well-separated hypercubes.
T_{Aa}	Time taken to construct the extended sparse matrix using NNCA.
T_{Af}	Time taken by the elimination phase of AIFMM excluding the time taken to perform the Schur complement operations on the rhs.
T_{As}	Sum of the time taken by the back substitution phase of AIFMM and the time taken to perform the Schur complement operations on the rhs, i.e., the respective operations that are to be performed on the rhs during the elimination phase.
E_A	Relative forward error of AIFMM measured using $\ \cdot\ _2$.
ϵ_{GMRES}	The relative residual $\frac{\ A\hat{x}-b\ _2}{\ b\ _2}$, that is used as the stopping criterion for GMRES, where \hat{x} is the solution computed using GMRES.
T_{Ga}	For problems involving non-oscillatory Green's functions and the Helmholtz function at low frequency it is the time taken to construct the \mathcal{H}^2 matrix representation. For problems involving high frequency Helmholtz function it is the time taken to construct the DAFMM matrix [33].
T_{Gs}	Time taken to solve the system using GMRES
I_G	Number of iterations it takes for convergence by GMRES with no preconditioner.
E_G	Relative forward error of GMRES measured using $\ \cdot\ _2$.
T_{Ha}	Time taken to assemble the matrix in HODLR form.
T_{Hf}	Time taken to factorize using HODLR.
T_{Hs}	Time taken to solve using HODLR.
E_H	Relative forward error of HODLR measured using $\ \cdot\ _2$.
I_{pA}	Number of iterations it takes for convergence by GMRES with AIFMM as a preconditioner.
I_{pH}	Number of iterations it takes for convergence by GMRES with HODLR preconditioner.
I_{BD}	Number of iterations it takes for convergence by GMRES with block-diagonal preconditioner.
T_{BD}	Time taken to solve by GMRES with block-diagonal preconditioner.
T_{pAs}	Time taken to solve by GMRES with AIFMM preconditioner.
T_{pHs}	Time taken to solve by GMRES with HODLR preconditioner.
relative error	relative forward error in the solution measured using $\ \cdot\ _2$.

Table 7: List of notations followed in this section

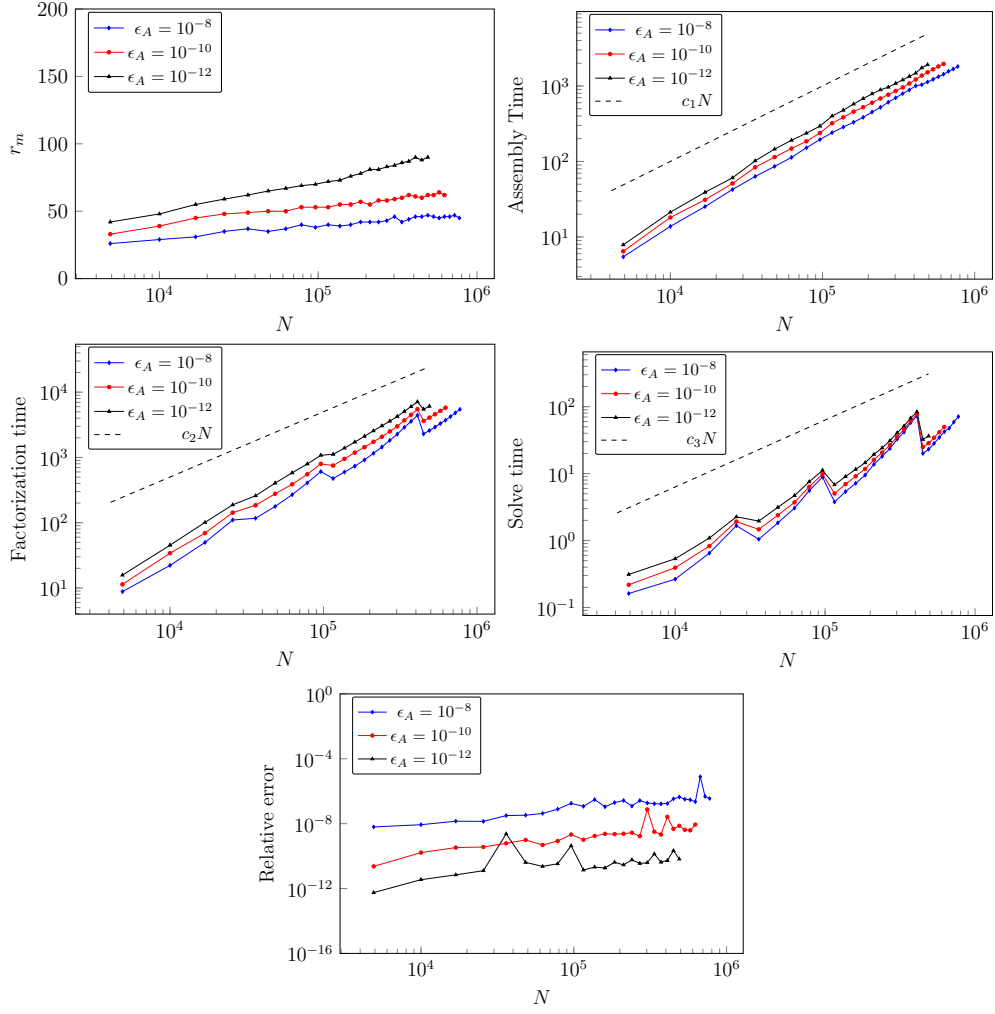


Figure 12: Results obtained with Experiment 1; Various benchmarks of AIFMM plotted for different values of ϵ_A

4.2. Experiment 2: Comparison of AIFMM with HODLR and GMRES for the 2D Helmholtz kernel

Here we consider the same matrix as considered in Subsection 4.1. ϵ_A , ϵ_G , and ϵ_H are set to 10^{-10} . ϵ_{GMRES} is also set to 10^{-10} . We tabulate the various CPU times and the relative errors of the three solvers AIFMM, GMRES, and HODLR in Table 8. Some of these benchmarks are also illustrated in Figure 13.

N	Assembly				Factorization			Solve					Error		
	T_{Ga}	T_{Ha}	T_{Aa}	$\frac{T_{Ha}}{T_{Aa}}$	T_{Hf}	T_{Af}	$\frac{T_{Hf}}{T_{Af}}$	T_{Gs}	T_{Hs}	T_{As}	$\frac{T_{Gs}}{T_{As}}$	$\frac{T_{Hs}}{T_{As}}$	E_G	E_H	E_A
4900	4.3	20.1	6.4	3.2	9.1	11.4	0.8	19.6	0.1	0.2	89.7	0.3	4e-11	9e-11	1e-11
16900	21.5	191.8	30.9	6.2	103.5	69.3	1.5	92.4	0.4	0.8	112.1	0.5	4e-10	3e-11	3e-10
36100	65.0	793.8	83.8	9.5	439.7	184.3	2.4	140.8	1.2	1.4	97.9	0.8	5e-10	3e-11	5e-10
62500	100.7	2231.1	143.3	15.6	1283.3	385.8	3.3	440.7	2.9	3.7	119.5	0.8	5e-10	1e-09	5e-10
115600	254.2	7211.8	318.1	22.7	4199.6	750.9	5.6	504.6	7.1	5.0	100.1	1.4	1e-09	3e-10	1e-09
160000	347.2	13347.6	450.1	29.7	7658.6	1178.5	6.5	1025.4	10.9	9.1	113.0	1.2	2e-09	6e-10	2e-09
240100	459.7	-	657.4	-	-	2060.6	-	2288.1	-	20.5	111.8	-	3e-09	-	3e-09
336400	617.2	-	976.4	-	-	3708.5	-	4552.1	-	45.8	99.4	-	3e-09	-	3e-09
448900	1099.2	-	1376.8	-	-	3619.1	-	2636.3	-	25.0	105.4	-	5e-09	-	5e-09
577600	1393.8	-	1809.0	-	-	5166.7	-	4373.4	-	41.5	105.3	-	4e-09	-	4e-09
672400	1586.7	-	2120.3	-	-	6429.6	-	6201.4	-	56.6	109.7	-	7e-09	-	8e-09

Table 8: Results obtained with experiment 2; CPU times and relative errors of the three solvers




Figure 13: Results obtained with Experiment 2; Plots of r_m , assembly time, factorisation time, and solve time versus N of AIFMM in comparison to those of HODLR and GMRES

4.3. Experiment 3: Comparison of AIFMM with HODLR and GMRES for the 2D Laplace kernel

Here we consider the 2D Laplace kernel. Again to ensure a well-conditioned matrix, we consider the entries of the matrix to be

$$A_{i,j} = \begin{cases} \sqrt{1000N} & \text{if } i == j \\ \frac{1}{\|x_i - x_j\|_2} & \text{else} \end{cases}. \quad (73)$$

ϵ_A , ϵ_G and ϵ_H are set to 10^{-10} , 10^{-8} , and 10^{-10} respectively. We used different compression tolerances to ensure that the relative errors of the three solvers AIFMM, GMRES, and HODLR are of the same order so that the CPU times of the solvers can be compared and an inference can be drawn on which solver performs better. ϵ_{GMRES} is set to 10^{-10} . We tabulate the various CPU times and the relative errors of the three solvers AIFMM, GMRES and HODLR in Table 9. Some of these benchmarks are also illustrated in Figure 14.

N	Assembly				Factorization			Solve					Error		
	T_{Ga}	T_{Ha}	T_{Aa}	$\frac{T_{Ha}}{T_{Aa}}$	T_{Hf}	T_{Af}	$\frac{T_{Hf}}{T_{Af}}$	T_{Gs}	T_{Hs}	T_{As}	$\frac{T_{Gs}}{T_{As}}$	$\frac{T_{Hs}}{T_{As}}$	E_G	E_H	E_A
4900	0.7	6.1	0.7	9.0	6.8	4.5	1.5	0.8	0.1	0.1	5.4	0.4	2e-08	5e-11	2e-08
16900	3.7	86.7	4.0	21.7	81.2	37.4	2.2	4.6	0.4	0.5	8.4	0.6	5e-08	5e-09	5e-08
36100	11.1	393.9	11.9	33.0	333.6	125.5	2.7	13.7	1.3	1.3	10.8	1.1	2e-07	5e-09	2e-07
62500	21.7	1147.4	22.7	50.6	966.1	275.7	3.5	38.3	2.7	2.4	15.7	1.1	2e-07	2e-06	2e-07
96100	34.2	2559.5	34.5	74.2	2030.0	526.0	3.9	204.1	5.1	4.6	44.1	1.1	8e-07	2e-07	7e-07
136900	69.4	5809.5	69.0	84.2	4187.1	844.9	5.0	118.9	8.4	6.1	19.4	1.4	3e-07	3e-07	3e-07
184900	91.9	10780.8	94.9	113.7	7265.3	1283.2	5.7	224.6	14.5	8.9	25.3	1.6	4e-07	4e-07	4e-07
240100	128.5	-	125.8	-	-	1819.3	-	351.6	-	12.7	27.7	-	4e-07	-	4e-07
336400	179.2	-	183.8	-	-	2914.8	-	648.5	-	21.4	30.3	-	6e-07	-	6e-07
490000	317.4	-	301.8	-	-	4612.9	-	796.2	-	27.5	28.9	-	7e-07	-	7e-07

Table 9: Results obtained with experiment 3; CPU times and relative errors of the three solvers




Figure 14: Results obtained with Experiment 3; Plots of r_m , assembly time, factorisation time, and solve time versus N of AIFMM in comparison to those of HODLR and GMRES

4.4. Experiment 4: Comparison of AIFMM with HODLR and GMRES in solving an integral equation

Consider the Fredholm integral equation of the second kind as defined in Equation (74),

$$\sigma_x + \int_{\Omega} K(x, y) \sigma_y dy = f(x) \quad (74)$$

where $K(x, y) = \log(\|x - y\|_2)$ and $\Omega = [-1, 1]^2$. We discretize Equation (74) using the Nystrom discretization on a uniform grid, which yields a linear system of the form $Ax = b$. Here we consider b to be a random vector. ϵ_A , ϵ_G , and ϵ_H are set to 10^{-10} , 10^{-8} , and 10^{-8} respectively. ϵ_{GMRES} is set to 10^{-10} . We tabulate the various CPU times and the relative errors of the three solvers AIFMM, GMRES, and HODLR in Table 10. Some of these benchmarks are also illustrated in Figure 15.

N	Assembly				Factorization			Solve					Error		
	T_{Ga}	T_{Ha}	T_{Aa}	$\frac{T_{Ha}}{T_{Aa}}$	T_{Hf}	T_{Af}	$\frac{T_{Hf}}{T_{Af}}$	T_{Gs}	T_{Hs}	T_{As}	$\frac{T_{Gs}}{T_{As}}$	$\frac{T_{Hs}}{T_{As}}$	E_G	E_H	E_A
4900	0.6	0.7	0.8	0.9	0.5	1.8	0.3	1.7	0.0	0.0	38.5	0.2	7e-11	4e-10	2e-10
16900	2.9	5.9	3.9	1.5	4.7	10.4	0.4	7.0	0.1	0.2	46.3	0.4	6e-10	2e-10	5e-10
36100	9.2	23.5	11.4	2.1	18.0	29.5	0.6	10.1	0.2	0.3	35.6	0.6	4e-09	6e-10	4e-09
62500	14.2	65.3	19.1	3.4	49.6	55.2	0.9	27.6	0.5	0.7	41.3	0.7	2e-09	3e-09	2e-09
96100	20.0	148.3	29.8	5.0	109.1	105.2	1.0	61.7	0.9	1.6	39.0	0.6	6e-09	2e-09	6e-09
160000	50.4	378.6	64.1	5.9	277.0	168.7	1.6	57.8	2.0	1.7	33.2	1.1	1e-08	1e-09	1e-08
240100	68.2	806.8	89.2	9.0	571.2	275.2	2.1	123.8	3.5	3.4	36.4	1.0	4e-09	3e-09	4e-09
409600	107.6	-	156.4	-	-	686.3	-	343.9	-	11.8	29.2	-	6e-09	-	6e-09
672400	253.9	-	324.1	-	-	905.1	-	307.6	-	10.4	29.5	-	6e-09	-	6e-09
883600	309.8	-	401.4	-	-	1322.3	-	511.9	-	17.6	29.0	-	5e-09	-	5e-09
1000000	333.3	-	477.5	-	-	1614.3	-	649.3	-	22.6	28.8	-	9e-09	-	9e-09

Table 10: Results obtained with experiment 4; CPU times and relative errors of the three solvers

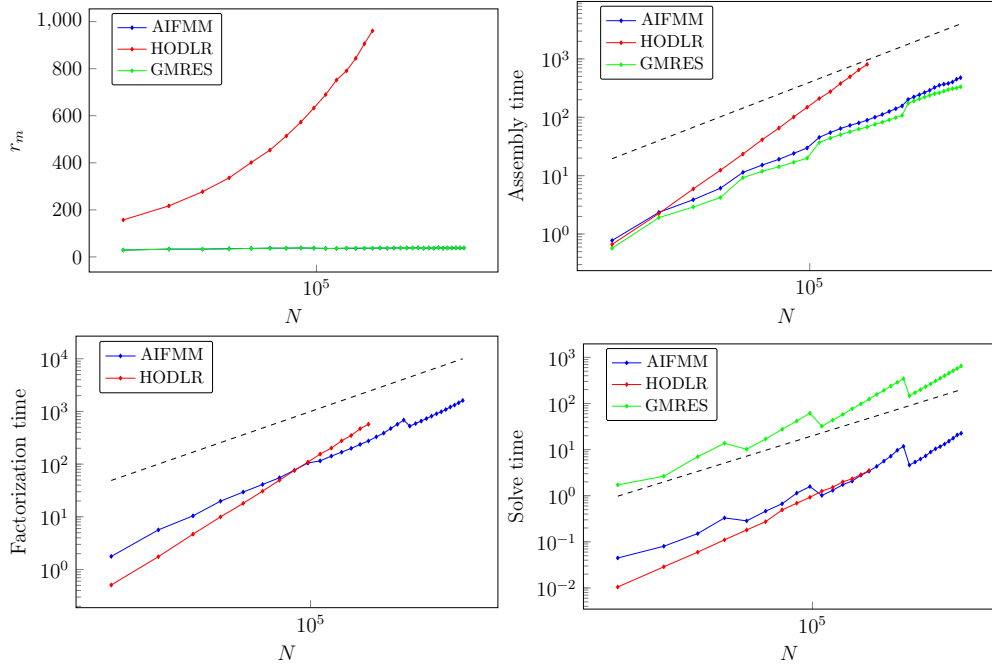


Figure 15: Results obtained with Experiment 4; Plots of r_m , assembly time, factorisation time, and solve time versus N of AIFMM in comparison to those of HODLR and GMRES

4.5. Wave scattering in 2D

We now demonstrate AIFMM on the matrix arising in solving acoustic or electromagnetic scattering from penetrable media using the Lippmann-Schwinger equation. It arises in many applications such as medical imaging, sonar, radar, geophysics, remote sensing, etc.

4.5.1. Formulation of the Lippmann-Schwinger equation in 2D

We now brief the formulation of the Lippmann-Schwinger equation in 2D. For a detailed description of the formulation of the Lippmann-Schwinger equation, we refer the readers to [33, 37].

Let $q(x)$, having compact support in the domain Ω , be the contrast function (or the susceptibility) of the penetrable medium. Let $u^{inc}(x)$ be the incident field and $u^{scat}(x)$ be the scattered field. Assume κ to be the wavenumber of the incident field. The total field $u(x)$ satisfies the time-harmonic Helmholtz equation

$$\nabla^2 u(x) + \kappa^2(1 + q(x))u(x) = 0, \quad x \in \mathbb{R}^2. \quad (75)$$

Assuming the incident field, $u^{inc}(x)$, satisfies the homogeneous Helmholtz equation, the scattered field $u^{scat}(x)$ satisfies

$$\nabla^2 u^{scat}(x) + \kappa^2(1 + q(x))u^{scat}(x) = f(x) \quad (76)$$

where $f(x) = -\kappa^2 q(x)u^{inc}(x)$. Expressing $u^{scat}(x)$ as the volume potential in terms of an unknown density function $\psi(x)$ and the Green's function, as in Equation (77)

$$u^{scat}(x) = V[\psi](x) = \int_{\Omega} G_{\kappa}(x, y)\psi(y)dy, \quad (77)$$

results in the Lippmann-Schwinger equation

$$\psi(x) + \kappa^2 q(x)V[\psi](x) = f(x). \quad (78)$$

To discretize the Lippmann-Schwinger equation, a balanced quad-tree is constructed. The tensor product Chebyshev nodes of size p^2 of each leaf box are considered to be the grid points. A local polynomial approximation of $\psi(x)$ in each leaf box B is built as following

$$\psi(x) \approx \psi^B(x) = \sum_{l=1}^{N_p} c_l^B b_l \left(\frac{\zeta_1 - \alpha_1^B}{\beta^B}, \frac{\zeta_2 - \alpha_2^B}{\beta^B} \right) \quad \forall x = (\zeta_1, \zeta_2) \in B. \quad (79)$$

where N_p is chosen to be $p(p+1)/2$, (α_1^B, α_2^B) are the coordinates of the center of box B , β^B is the half side length of B , and $\{b_l\}_{l=1}^{N_p}$ are the polynomial basis functions. Let $\{x_i^B\}_{i=1}^{p^2}$ be the gridpoints of box B , at which the unknown $\psi(x)$ is evaluated. Vector $\vec{\psi}^B = [\psi(x_1^B), \psi(x_2^B), \dots, \psi(x_{p^2}^B)]^T$ is expressed in terms of vector $\vec{c}^B = [c_1^B, c_2^B, \dots, c_{N_p}^B]^T$ as

$$\vec{\psi}^B = Q \vec{c}^B \quad (80)$$

where $Q \in \mathbb{R}^{p^2 \times N_p}$ is the interpolation matrix, whose entries are given by

$$Q_{il} = b_l \left(\frac{\zeta_{i,1}^B - \alpha_1^B}{\beta^B}, \frac{\zeta_{i,2}^B - \alpha_2^B}{\beta^B} \right), \quad x_i^B = (\zeta_{i,1}^B, \zeta_{i,2}^B). \quad (81)$$

By taking the pseudo-inverse of Q , we obtain \vec{c}^B in terms of $\vec{\psi}^B$

$$\vec{c}^B = Q^\dagger \vec{\psi}^B. \quad (82)$$

Using Equations (82) and (79), we have

$$\psi^B(x) = \sum_{l=1}^{N_p} Q^\dagger(l, :) \vec{\psi}^B b_l \left(\frac{\zeta_1 - \alpha_1^B}{\beta^B}, \frac{\zeta_2 - \alpha_2^B}{\beta^B} \right). \quad (83)$$

Using Equation (83) we build an approximate of $V[\psi](x)$ as

$$V[\psi](x) \approx \sum_{\mathcal{L}} \int_B G_{\kappa}(x, y) \sum_{l=1}^{N_p} Q^\dagger(l, :) \vec{\psi}^B b_l \left(\frac{\zeta_1 - \alpha_1^B}{\beta^B}, \frac{\zeta_2 - \alpha_2^B}{\beta^B} \right) dy. \quad (84)$$

where \mathcal{L} is the set of all leaf boxes. Using Equations (84) and (78), the discretized version of the Lippmann-Schwinger equation is obtained,

$$\psi(x) + \kappa^2 q(x) \sum_{\mathcal{L}} \int_B G_\kappa(x, y) \sum_{l=1}^{N_p} Q^\dagger(l, :) \vec{\psi}^B b_l \left(\frac{\zeta_1 - \alpha_1^B}{\beta^B}, \frac{\zeta_2 - \alpha_2^B}{\beta^B} \right) dy \approx f(x). \quad (85)$$

By enforcing Equation (78) at the grid points of all the leaf nodes of the tree, and using Equation (84), we obtain the linear system

$$A \vec{\psi} = \vec{f} \quad (86)$$

where $\vec{\psi}$ is a vector that contains function values of $\psi(x)$ evaluated at the grid points of leaf boxes of the quad-tree. The $(i, j)^{th}$ entry of A , that represents the contribution of the j^{th} grid point at the i^{th} grid point is given by

$$A_{ij} = \delta_{ij} + \kappa^2 q(x_i) \int_{\Omega \in B} G_\kappa(x_i, y) \sum_{l=1}^{N_p} Q_{l,j}^\dagger b_l \left(\frac{y_1 - \alpha_1^B}{\beta^B}, \frac{y_2 - \alpha_2^B}{\beta^B} \right) dy_1 dy_2 \quad (87)$$

where $j' = 1 + (j - 1) \bmod p^2$ and B is a leaf box that contains the support of the grid point x_j . The entries f_j of the rhs vector \vec{f} are given by $f_j = f(x_j)$.

We solve for $\vec{\psi}$ and then use it to find $u^{scat}(x)$. $u^{scat}(x)$ is obtained by discretizing Equation (77) (in the same way that the Lippmann-Schwinger equation is discretized) and performing fast directional summation using the Directional Algebraic Fast Multipole Method (DAFMM) [33].

We find the error in the solution, using function $E(x)$, defined as the residual of Equation (85) normalized with κ^2 .

$$E(x) = \left| \frac{\psi(x)}{\kappa^2} + q(x) \sum_{\mathcal{L}} \int_B G_\kappa(x, y) \sum_{l=1}^{N_p} Q^\dagger(l, :) \vec{\psi}^B b_l \left(\frac{y_1 - \alpha_1^B}{\beta^B}, \frac{y_2 - \alpha_2^B}{\beta^B} \right) dy - \frac{f(x)}{\kappa^2} \right|, \quad x \in \Omega. \quad (88)$$

We define vector \vec{E} , where the entries take the function values of $E(x)$ at the grid points of all the leaf nodes. We report $\|\vec{E}\|_2$ as well as illustrate $E(x)$ pictorially. We further define a few notations to represent this error for various solvers in Table 11.

E_G^*	Error $\ \vec{E}\ _2$ of GMRES
E_{pH}^*	Error $\ \vec{E}\ _2$ of GMRES with HODLR as preconditioner
E_{BD}^*	Error $\ \vec{E}\ _2$ of GMRES with block-diagonal preconditioner
E_{pA}^*	Error $\ \vec{E}\ _2$ of GMRES with AIFMM as preconditioner

Table 11: Errors of various solvers for the Lippmann-Schwinger equation

4.5.2. Experiment 5: AIFMM as a preconditioner in an iterative solver for Lippmann-Schwinger equation at high frequency

In this experiment, we demonstrate AIFMM as a preconditioner in solving the Lippmann-Schwinger equation using GMRES. We consider Gaussian contrast defined as

$$q(x) = 1.5 \exp(-160(x_1^2 + x_2^2)). \quad (89)$$

κ is set to 300. The depth of the uniform quad tree is set to 6. The leaf size p^2 is varied to get different system sizes as shown in Table 12. The same tree that is used for discretization of the Lippmann-Schwinger equation is used for building AIFMM, HODLR, and DAFMM routines. ϵ_G and ϵ_{GMRES} are set to 10^{-10} and 10^{-8} respectively.

We compared AIFMM as a preconditioner to a block-diagonal preconditioner and HODLR preconditioner. The block diagonal preconditioner is constructed by choosing the block size to be equal to the leaf size. In Table 13, we illustrate the CPU times and the errors of the four iterative solvers: GMRES solver with no preconditioner, GMRES

p^2	36	64	100
N	147456	262144	409600

Table 12: Different leaf sizes and the corresponding problem sizes

with block-diagonal preconditioner, GMRES with HODLR as a preconditioner, and GMRES with AIFMM as a preconditioner. In Figures 16a and 16b, for $N = 147456$, the real part of the scattered field, and the log plot of the error function obtained using GMRES with AIFMM as preconditioner with ϵ_A set to 10^{-5} are plotted. In Figures 16c and 16d, relative residual versus iteration count for different values of ϵ_A and time taken by the different iterative solvers are plotted.

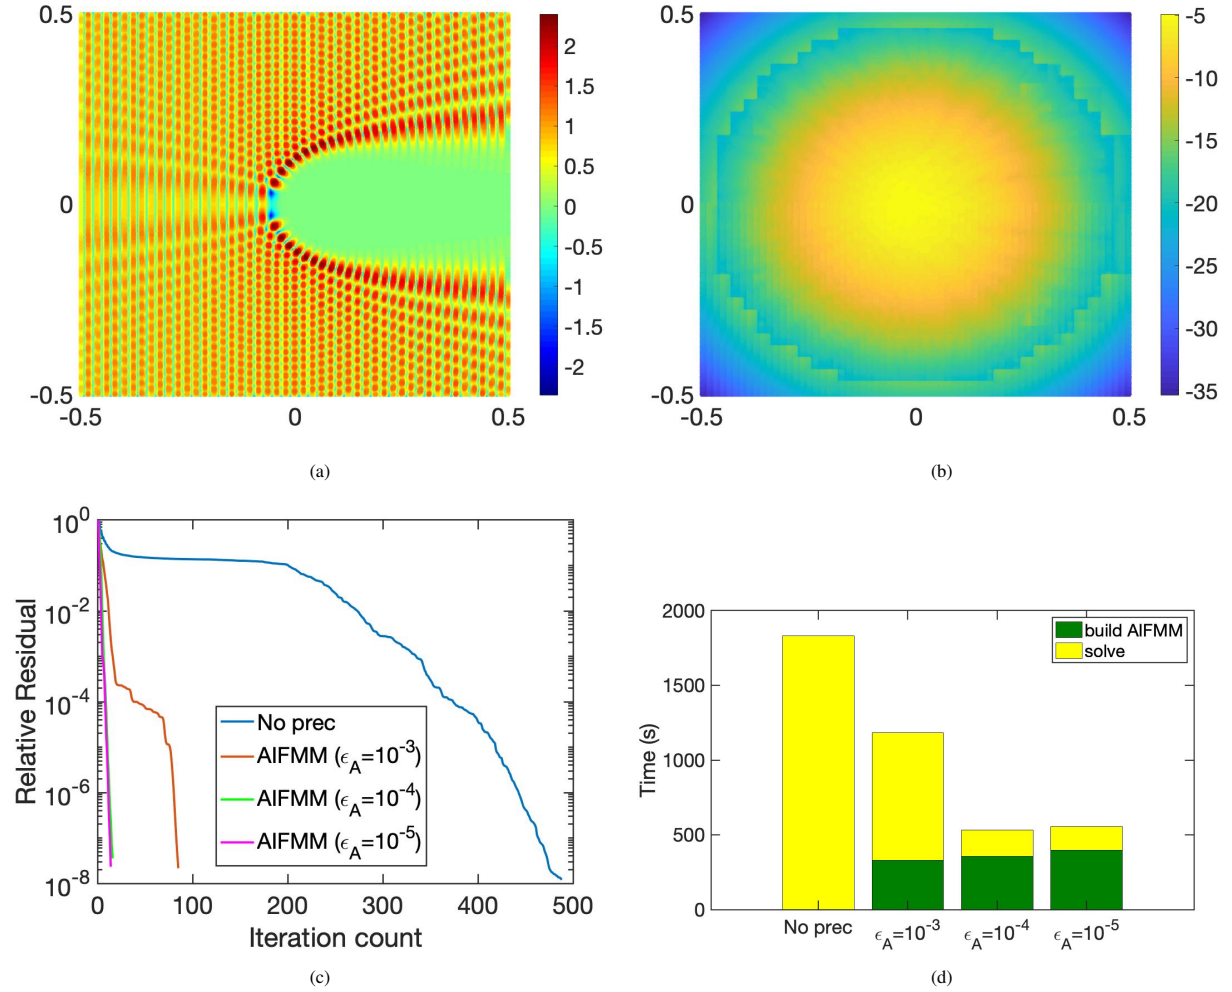


Figure 16: Results obtained with Experiment 5 with $N = 147456$; a) real part of the scattered field b) log plot of the error function $E(x)$ c) relative residual versus iteration count for various values of ϵ_A d) Time taken by the iterative solver, where the green bar indicates the time taken to assemble and factorize the AIFMM preconditioner and the yellow bar indicates the taken by the GMRES solver that includes the time to apply the preconditioner

4.6. Inferences

The following inferences are to be noticed from Figures 14 to 15 and Tables 9 to 10.

	N	147456	262144	409600
	T_{Ga}	856.04	699.70	1222.23
GMRES with no preconditioner	T_{Gs}	1822.34	3862.73	10353.70
	I_G	488	365	366
	E_G^*	1.8×10^{-4}	-	-
Block Diagonal preconditioner	T_{BD}	1170	3347.6	8306.9
	I_{BD}	294	295	293
	E_{BD}^*	1.8×10^{-4}	-	-
GMRES with HODLR ($\epsilon_H = 10^{-5}$) preconditioner	T_{Ha}	133.85	242.76	442.34
	T_{Hf}	33.78	65.24	116.52
	T_{pHs}	62.75	192.55	457.66
	I_{pH}	16	17	15
	E_{pH}^*	1.8×10^{-4}	-	-
GMRES with AIFMM ($\epsilon_A = 10^{-5}$) preconditioner	T_{Aa}	43.10	95.50	167.49
	T_{Af}	360.69	1711.02	5843.42
	T_{pAs}	151.77	535.00	1644.67
	I_{pA}	14	12	12
	E_{pA}^*	1.8×10^{-4}	-	-
	$\frac{T_{Gs}}{T_{Aa}+T_{Af}+T_{pAs}}$	3.28	1.65	1.35
	$\frac{T_{BD}}{T_{Aa}+T_{Af}+T_{pAs}}$	2.11	1.43	1.08

Table 13: Results obtained with Experiment 5; CPU times, relative error, and number of iterations it takes using GMRES with no preconditioner and GMRES with AIFMM, HODLR, and block diagonal preconditioners

1. The maximum rank of HODLR is proportional to N , whereas that of AIFMM does not scale with N .
2. Assembly time, solve time, and factorization time scale linearly with N for AIFMM and GMRES, whereas those of HODLR do not scale linearly.
3. AIFMM is faster than HODLR for the examples considered. The speedup can be observed from Tables 9 to 10.
4. The assembly time of GMRES and AIFMM are nearly equal and the solve time of GMRES is higher than that of AIFMM. When the total CPU time is considered, $T_{Ga} + T_{Gs}$ for GMRES and $T_{Aa} + T_{Af} + T_{As}$ for AIFMM, GMRES is faster than AIFMM for the examples considered. But when one is interested in multiple right-hand sides, it is advantageous to use AIFMM over GMRES, as the solve time of AIFMM is lower than that of GMRES.

From Table 13, it can be noticed that AIFMM performs well as a preconditioner for the high frequency scattering problem and is better than the block diagonal preconditioner, but not as good as the HODLR preconditioner.

In summary, for the problems considered we observed that

- The time complexity of AIFMM scales linearly with N .
- In problems involving the low frequency Helmholtz function and non-oscillatory Green's functions, AIFMM performs better than HODLR as a direct solver. And AIFMM performs better than GMRES when one considers multiple right hand sides.
- In the high frequency scattering problem, HODLR as a preconditioner performs better than AIFMM as a preconditioner. And AIFMM as a preconditioner performs better than the block diagonal preconditioner.

5. Conclusions

A completely algebraic, linear complexity, direct solver for FMM matrices is presented. The various FMM operators are obtained using NNCA, that algebraically obtains nested bases. The advantages of an algebraic technique are (i) the ranks obtained are lower because the method is domain and problem specific; (ii) it can be used in black box fashion, independent of the application. The key ideas of the AIFMM are i) to construct an extended sparse system; (ii) and then perform elimination and substitution, wherein in the elimination phase, the fill-ins corresponding to well-separated hypercubes are compressed and redirected using RRQR. Various numerical experiments were presented to demonstrate the scaling and accuracy of AIFMM as a direct solver. It is shown that AIFMM is faster than HODLR, a direct solver. Further, when multiple right hand sides are to be solved for, then AIFMM is better than GMRES. It is also shown that for the high frequency scattering problem, it can be used as a preconditioner and it performs better than the block-diagonal preconditioner.

Acknowledgments

The authors acknowledge HPCE, IIT Madras for providing access to the AQUA cluster. Vaishnavi Gujjula acknowledges the support of Women Leading IITM (India) 2022 in Mathematics (SB22230053MAIITM008880). Sivaram Ambikasaran acknowledges the support of Young Scientist Research Award from Board of Research in Nuclear Sciences, Department of Atomic Energy, India (No.34/20/03/2017-BRNS/34278) and MATRICS grant from Science and Engineering Research Board, India (Sanction number: MTR/2019/001241).

References

- [1] S. Ambikasaran, E. Darve, The inverse fast multipole method, arXiv preprint arXiv:1407.1572 (2014).
- [2] P. Coulier, H. Pouransari, E. Darve, The inverse fast multipole method: Using a fast approximate direct solver as a preconditioner for dense linear systems, SIAM Journal on Scientific Computing 39 (3) (2017) A761–A796.
- [3] T. Takahashi, P. Coulier, E. Darve, Application of the inverse fast multipole method as a preconditioner in a 3d helmholtz boundary element method, Journal of Computational Physics 341 (2017) 406–428.
- [4] V. Gujjula, S. Ambikasaran, A new nested cross approximation, arXiv preprint arXiv:2203.14832 (2022).
- [5] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, Journal of computational physics 73 (2) (1987) 325–348.
- [6] J. Barnes, P. Hut, A hierarchical $O(n \log n)$ force-calculation algorithm, nature 324 (6096) (1986) 446–449.
- [7] W. Hackbusch, A sparse matrix arithmetic based on \mathcal{H} -matrices. part i: Introduction to \mathcal{H} -matrices, Computing 62 (2) (1999) 89–108.
- [8] L. Grasedyck, W. Hackbusch, Construction and arithmetics of h-matrices, Computing 70 (4) (2003) 295–334.

- [9] S. Börm, L. Grasedyck, W. Hackbusch, Introduction to hierarchical matrices with applications, *Engineering analysis with boundary elements* 27 (5) (2003) 405–422.
- [10] M. Bebendorf, Hierarchical matrices: a means to efficiently solve elliptic boundary value problems, *Universität Leipzig*, 2007.
- [11] K. L. Ho, L. Greengard, A fast direct solver for structured linear systems by recursive skeletonization, *SIAM Journal on Scientific Computing* 34 (5) (2012) A2507–A2532.
- [12] W. Y. Kong, J. Bremer, V. Rokhlin, An adaptive fast direct solver for boundary integral equations in two dimensions, *Applied and Computational Harmonic Analysis* 31 (3) (2011) 346–369.
- [13] P.-G. Martinsson, V. Rokhlin, A fast direct solver for boundary integral equations in two dimensions, *Journal of Computational Physics* 205 (1) (2005) 1–23.
- [14] S. Ambikasaran, E. Darve, An $O(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices, *Journal of Scientific Computing* 57 (3) (2013) 477–501.
- [15] S. Chandrasekaran, P. Dewilde, M. Gu, T. Pals, A.-J. van der Veen, Fast stable solver for sequentially semi-separable linear systems of equations, in: *International Conference on High-Performance Computing*, Springer, 2002, pp. 545–554.
- [16] S. Chandrasekaran, M. Gu, T. Pals, A fast ulv decomposition solver for hierarchically semiseparable representations, *SIAM Journal on Matrix Analysis and Applications* 28 (3) (2006) 603–622.
- [17] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, T. Pals, A fast solver for hss representations via sparse matrices, *SIAM Journal on Matrix Analysis and Applications* 29 (1) (2007) 67–81.
- [18] V. Kandappan, V. Gujjula, S. Ambikasaran, Hodlr2d: A new class of hierarchical matrices, *arXiv preprint arXiv:2204.05536* (2022).
- [19] R. Khan, V. Kandappan, S. Ambikasaran, Numerical rank of singular kernel functions, *arXiv preprint arXiv:2209.05819* (2022).
- [20] M. Bebendorf, *Hierarchical matrices*, Springer, 2008.
- [21] W. Hackbusch, *Hierarchical matrices: algorithms and analysis*, Vol. 49, Springer, 2015.
- [22] W. Hackbusch, S. Börm, H2-matrix approximation of integral operators by interpolation, *Applied numerical mathematics* 43 (1-2) (2002) 129–143.
- [23] M. Bebendorf, Hierarchical lu decomposition-based preconditioners for bem, *Computing* 74 (3) (2005) 225–247.
- [24] S. Börm, \mathcal{H}^2 -matrix arithmetics in linear complexity, *Computing* 77 (1) (2006) 1–28.
- [25] S. Börm, Efficient numerical methods for non-local operators: H2-matrix compression, algorithms and analysis, Vol. 14, *European Mathematical Society*, 2010.
- [26] S. Börm, K. Reimer, Efficient arithmetic operations for rank-structured matrices based on hierarchical low-rank updates, *Computing and Visualization in Science* 16 (6) (2013) 247–258.
- [27] V. Minden, K. L. Ho, A. Damle, L. Ying, A recursive skeletonization factorization based on strong admissibility, *Multiscale Modeling & Simulation* 15 (2) (2017) 768–796.
- [28] D. Sushnikova, L. Greengard, M. O’Neil, M. Rachh, Fmm-lu: A fast direct solver for multiscale boundary integral equations in three dimensions, *arXiv preprint arXiv:2201.07325* (2022).
- [29] T. P. Pals, S. Chandrasekaran, H. Lee, Multipole for scattering computations: Spectral discretization, stabilization, fast solvers, Ph.D. thesis, *University of California Santa Barbara, USA*, aAI3159308 (2004).
- [30] L. Greengard, D. Gueyffier, P.-G. Martinsson, V. Rokhlin, Fast direct solvers for integral equations in complex three-dimensional domains, *Acta Numerica* 18 (2009) 243–275.
- [31] M. Bebendorf, R. Venn, Constructing nested bases approximations from the entries of non-local operators, *Numerische Mathematik* 121 (4) (2012) 609–635.
- [32] Y. Zhao, D. Jiao, J. Mao, Fast nested cross approximation algorithm for solving large-scale electromagnetic problems, *IEEE Transactions on Microwave Theory and Techniques* 67 (8) (2019) 3271–3283.
- [33] V. Gujjula, S. Ambikasaran, A new directional algebraic fast multipole method based iterative solver for the lippmann-schwinger equation accelerated with hodlr preconditioner, *Communications in Computational Physics* 32 (4) (2022) 1061–1093.
- [34] S. Ambikasaran, K. R. Singh, S. S. Sankaran, Hodlrlib: a library for hierarchical matrices, *Journal of Open Source Software* 4 (34) (2019) 1167.
- [35] Y. SAAD, A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput* 7 (1985) 417–424.
- [36] Y. Saad, *Iterative methods for sparse linear systems*, SIAM, 2003.
- [37] S. Ambikasaran, C. Borges, L.-M. Imbert-Gerard, L. Greengard, Fast, adaptive, high-order accurate discretization of the lippmann–schwinger equation in two dimensions, *SIAM Journal on Scientific Computing* 38 (3) (2016) A1770–A1787.