



HAL
open science

An efficient preconditioner for adaptive Fast Multipole accelerated Boundary Element Methods to model time-harmonic 3D wave propagation

Faisal Amlani, Stéphanie Chaillat, Adrien Loseille

► **To cite this version:**

Faisal Amlani, Stéphanie Chaillat, Adrien Loseille. An efficient preconditioner for adaptive Fast Multipole accelerated Boundary Element Methods to model time-harmonic 3D wave propagation. *Computer Methods in Applied Mechanics and Engineering*, 2019, 352 (1), pp.189-210. 10.1016/j.cma.2019.04.026 . hal-02113702

HAL Id: hal-02113702

<https://hal.science/hal-02113702v1>

Submitted on 29 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An efficient preconditioner for adaptive Fast Multipole accelerated Boundary Element Methods to model time-harmonic 3D wave propagation

Faisal Amlani^a, Stéphanie Chaillat^a, Adrien Loseille^b

^aLaboratoire POEMS (CNRS-ENSTA-INRIA),

ENSTA-UMA, 828 Bd des Maréchaux, 91120 Palaiseau Cedex, FRANCE

^bGamma 3 Team, INRIA Saclay-Ile de France, 1 Rue Honoré d'Estienne d'Orves, 91120 Palaiseau, FRANCE

Abstract

This paper presents an efficient algebraic preconditioner to speed up the convergence of Fast Multipole accelerated Boundary Element Methods (FM-BEMs) in the context of time-harmonic 3D wave propagation problems and in particular the case of highly non-uniform discretizations. Such configurations are produced by a recently-developed anisotropic mesh adaptation procedure that is independent of partial differential equation and integral equation. The new preconditioning methodology exploits a complement between fast BEMs by using two nested GMRES algorithms and rapid matrix-vector calculations. The fast inner iterations are evaluated by a coarse hierarchical matrix (\mathcal{H} -matrix) representation of the BEM system. These inner iterations produce a preconditioner for FM-BEM solvers. It drastically reduces the number of outer GMRES iterations. Numerical experiments demonstrate significant speedups over non-preconditioned solvers for complex geometries and meshes specifically adapted to capture anisotropic features of a solution, including discontinuities arising from corners and edges.

Keywords: Boundary Element Method, Fast Multipole Method, Anisotropic meshes, Preconditioning, Hierarchical matrices

1. Introduction

This paper is concerned with the efficient numerical solution of 3D wave propagation problems in large-scale unbounded domains. The accurate numerical modeling of acoustic, elastodynamic or electromagnetic problems is a challenging task. It is also a timely research field due to the variety of possible applications, e.g., seismology, non-destructive testing, noise reduction, radar stealth, etc. We focus on an exterior Helmholtz problem for clarity, but the extension to other wave propagation problems (e.g., elastodynamics and electromagnetism) is straightforward. Let $\Omega \subset \mathbb{R}^3$ be a closed bounded domain with boundary Γ and let $\Omega^+ := \mathbb{R}^3 \setminus \bar{\Omega}$ denote the exterior scattering domain. Given an incident plane wave $u^i(\mathbf{x}) = e^{ik\mathbf{d}\cdot\mathbf{x}}$, $\mathbf{x} \in \mathbb{R}^3$ of wavelength λ (wavenumber $k := 2\pi/\lambda$) and direction \mathbf{d} , the total field $u = u^i + u^s$ (where u^s is the scattered field) in the exterior domain satisfies the time-harmonic acoustic wave equation given by

$$\nabla^2 u + k^2 u = 0 \quad \text{in } \Omega^+. \quad (1a)$$

We impose a sound-soft boundary condition as

$$u = 0 \quad \text{on } \Gamma, \quad (1b)$$

together with a Sommerfeld radiation condition on the scattered field $u^s = u - u^i$ given as

$$\lim_{r \rightarrow \infty} \left(\frac{\partial u^s}{\partial r} - ik u^s \right) = 0, \quad r = |\mathbf{x}|. \quad (1c)$$

Numerical methodologies to solve this system fall into two main categories: those obtained by domain methods requiring discretization of the entire physical domain of interest (e.g., finite difference [36], finite element [23, 41, 43, 32], finite volume and spectral methods [45]), and boundary methods [46, 49, 33]

requiring discretization of a reduced dimension such as the obstacle surface for an exterior scattering problem. Domain methods require a discretization of the exterior domain Ω^+ that leads to large but sparse algebraic systems. These systems are not well-suited to unbounded domains since the computational domain must be artificially truncated, leading to additional computational difficulties [22, 8, 53, 55]. In addition, the radiation condition (1c) must be approximated. A more suitable alternative is to recast the system (1) as a boundary integral equation (BIE). This subsequently leads to discretizations only on Γ which can be provided by a boundary element method (BEM). Where applicable, these methods are very efficient due to exact representations of both the boundary as well as radiation conditions at infinity. For a comparison between boundary element and finite element methods see, e.g., [40].

There exist various boundary integral equations to model acoustic wave scattering, with different advantages and drawbacks. We consider in this work the simplest approach. We will show that our methodology is nevertheless independent of the BIE considered. The solution u for an obstacle surface Γ can be sought as a boundary integral representation in the form of a single-layer potential operator \mathcal{S} given by [24]

$$u(\mathbf{x}) = \mathcal{S}q(\mathbf{x}) = \int_{\Gamma} \Phi_k(\mathbf{x}, \mathbf{y})q(\mathbf{y})ds(\mathbf{y}), \quad \mathbf{x} \notin \Gamma, \quad \mathbf{x} \in \Omega^+, \quad (2a)$$

where the fundamental solution of the 3D Helmholtz equation is given by

$$\Phi_k(\mathbf{x}, \mathbf{y}) = \frac{e^{ik|\mathbf{x}-\mathbf{y}|}}{4\pi|\mathbf{x}-\mathbf{y}|}, \quad \mathbf{x} \neq \mathbf{y}, \quad k \in \mathbb{R}^+.$$

Expressing the unknown density q as $q = -\partial u / \partial \mathbf{n}$ (or, for geometries of no thickness, the jump $[\partial u / \partial \mathbf{n}]_{\Gamma}$ across Γ) in the direction of the outer unit normal vector \mathbf{n} to Γ , one can derive [24, 61] a corresponding BIE for (2a) as

$$-u^i = \int_{\Gamma} \Phi_k(\mathbf{x}, \mathbf{y})q(\mathbf{y})ds(\mathbf{y}), \quad \mathbf{x} \in \Gamma. \quad (2b)$$

Clearly many realistic applications in acoustics (and other wave propagation problems) may also require Neumann or Robin boundary conditions. The methodology presented in this work can be straightforwardly extended to these cases. For the sake of simplicity, and without loss of generality, we restrict our simulations to the Dirichlet case.

A number of factors determine the effectiveness of a BEM applied for the solution of a general boundary integral equation. These include the choice of integral equation, their subsequent numerical resolution and the adopted discretization strategy [57]. Standard implementations of BEM lead to dense and (possibly) non-symmetric linear systems whose solutions become prohibitively expensive for large-scale problems. This is particularly true for obstacles that are large relative to the wavelength and for problems employing highly non-uniform discretizations. The storage of such systems is of the order $\mathcal{O}(N^2)$, where N is the number of degrees of freedom corresponding to the discretization of the obstacle boundary (e.g., the number of nodes in the mesh). The computational complexity of solving such a dense system using a direct method like a Gaussian elimination is $\mathcal{O}(N^3)$, whereas resolution via an iterative method like GMRES is $\mathcal{O}(N_{\text{iter}}N^2)$ if N_{iter} is the number of iterations. Many approaches have been proposed to speedup the iterative resolution of these dense systems [31, 30, 13, 12], among which Fast Multipole Methods (FMMs) have enjoyed considerable success in mechanical engineering problems [18, 62, 14, 48] by enabling a fast evaluation of the matrix-vector product required by the iterative solver. Initially developed for N-body simulations, the FMM has since been extended to oscillatory kernels [27, 37, 54] and thus expanded its efficacy to many applications. An algebraic alternative designed for dense systems is based on the concept of hierarchical matrices (\mathcal{H} -matrices) [7]. The principle of \mathcal{H} -matrices is to partition the initial dense linear system and reduce it to a data-sparse one by finding sub-blocks in the matrix that can be accurately approximated by low-rank matrices. The efficiency of hierarchical matrices relies on the possibility to approximate, under certain conditions, the underlying kernel function by low-rank matrices. The approach has been shown to be very efficient for asymptotically smooth kernels (e.g., the Laplace kernel) and relatively efficient in a large frequency range for oscillatory kernels such as Helmholtz or elastodynamic kernels [25, 19, 39]. These two approaches have advantages and drawbacks in the context of the acceleration of BEMs. On the one hand, \mathcal{H} -matrix based BEMs (\mathcal{H} -BEM) are very easy to implement and have the advantage to be an algebraic approach. Most of the computational time of

\mathcal{H} -BEMs is spent to evaluate the data-sparse approximation of the BEM matrix, the cost of a matrix-vector is then very cheap. On the other hand, Fast Multipole accelerated BEMs (FM-BEMs) have been shown to have an optimal complexity in terms of computational times and memory requirements for wave propagation problems. The counterpart is a more involved implementation effort. In the following, our fast BEM solver is based on the FM-BEM for wave propagation problems.

Appropriate mesh adaptations can further improve the computational efficiency of the FM-BEM (particularly for complex geometries) by accommodating anisotropic features of a solution (e.g., some elastic materials) as well as discontinuities near geometric singularities (e.g., corners and edges). These behaviors are difficult to capture and ultimately diminish the regularity of the boundary solution and subsequent performance of a BEM. The underlying principle of a refinement algorithm relies on a transformation of an initial mesh into an improved one according to error estimates calculated at each step. The ultimate goal is a reduction in the number of degrees of freedom required to resolve a solution within a desired level of accuracy. Adaptation is particularly important for scattering obstacles that contain geometric singularities leading to rapid variations of surface solutions. If an extensive literature exists for volume-based methods [1, 34], relatively little research attention has been paid to error estimators and mesh adaptation techniques in the BEM community. In general, most adaptation techniques for BEMs have centered on isotropic techniques [29, 5, 9, 44] that improve performance but do not always recover optimal convergence rates for 3D problems with anisotropic features [4]. In this work, we employ a recently introduced *anisotropic* mesh adaptation (AMA) strategy using a metric-based error estimator whose effectiveness in recovering convergence order was first demonstrated for volumetric (finite element) methods [51, 52] and only recently for first-order BEMs [20]. The methodology is independent of discretization technique as well as the choice of partial differential equation (PDE) and integral equation formulation, iteratively constructing meshes refined in size, shape and orientation according to an “optimal” metric relying on a reconstructed Hessian of the boundary solution. The resulting adaptation is anisotropic in nature and accommodates geometric complexities that include engineering detail.

Since we are interested in large-scale scattering problems requiring iterative solutions, it is important to note that anisotropically (and even isotropically) adapted meshes worsen the conditioning of the overall FM-BEM system [35]. This leads to an increase in the number of iterations and thus overall computational time required to resolve the system. Even with the accelerations provided by the FMM, global solution times can still be prohibitively expensive, necessitating the use of preconditioning techniques to improve convergence of an applied iterative method. This is already an active area of research for uniform meshes [21, 63]. We are interested here primarily by algebraic preconditioners which can be generally applied to a wide variety of problems as opposed to analytic preconditioners, which tend to be very efficient but problem specific [47, 2, 26, 56]. This is motivated by the fact that FM-BEMs have been successfully applied to a wide variety of wave propagation problems, and the fact that the anisotropic mesh adaptation procedure used herein is independent of PDE and integral representation. Traditional algebraic preconditioning approaches such as incomplete LU, Sparse Approximative Inverse [17, 16], multi-grid methods [15] and nested GMRES algorithms [21] have been applied to acoustic, electromagnetic or elastodynamic FM-BEMs. Since to reduce memory requirements, the FM-BEM does not provide access to the full system but only close interactions, these algebraic preconditioners have shown moderate efficiency. They probably do not contain enough information on the physics of the underlying continuous operator.

Our objective in this paper is to maintain the optimal complexity of the FM-BEM in terms of memory requirements and computational costs while complementing it with an algebraic preconditioner constructed from the complete BEM system. We are particularly interested in non-uniform meshes constructed by an anisotropic mesh adaptation procedure that is independent of the PDE and integral representation. In this context, our approach effectively utilizes the complementarities of the FM-BEM and the \mathcal{H} -BEM to treat geometries modeled by non-uniform meshes. It is based on a nested sequence of GMRES iterations. Our methodology is presented as follows: in Section 2 we summarize the Boundary Element Method and its multi-level Fast Multipole-acceleration; in Section 3 we present the nested GMRES implementation using hierarchical matrices and consider some verification examples; and, finally, we demonstrate in Section 4 the efficacy of the iterative solver applied to anisotropic meshes through a variety of numerical experiments on an infinitely thin planar screen and a more complex F-15 aircraft.

2. Adaptive fast multipole-accelerated boundary element methods

We begin by presenting an adaptive FM-BEM for 3D wave propagation problems. Sections 2.1 and 2.2 summarize the general-order triangular Boundary Element formulation and subsequent multi-level Fast Multipole acceleration. Section 2.3 summarizes the mesh adaptation algorithm utilized to generate anisotropic meshes tailored to solutions of scattering problems on complex domains.

2.1. Boundary element methods

The numerical solution of the boundary integral equation (2b), known as the Boundary Element Method [10], is obtained by discretization of the surface Γ into N_E finite (boundary) elements. Importantly in the BEM, each element E_e must be analytically described. Each physical element E_e on the approximate boundary is mapped onto a reference element Δ_e in the (ξ_1, ξ_2) -plane via

$$\boldsymbol{\xi} \in \Delta_e \rightarrow \mathbf{y}(\boldsymbol{\xi}) \in E_e, \quad 1 \leq e \leq N_E.$$

The boundary of the geometry is parametrically represented by using a number of d nodes and shape functions v_k^e , per element. We use as nodes \mathbf{y}^k all the vertices for a \mathbb{P}_0 or \mathbb{P}_1 mesh ($d = 3$) and all vertices plus midpoints for a \mathbb{P}_2 mesh ($d = 6$) such that the geometry is approximated by

$$\mathbf{y}(\boldsymbol{\xi}) \approx \sum_{k=1}^d \mathbf{y}^k v_k^e(\boldsymbol{\xi}).$$

The shape functions are related to the canonical basis $(\hat{v}_k)_{1 \leq k \leq d}$ defined on the reference element Δ_e by $v_k^e(\mathbf{y}(\boldsymbol{\xi})) = \hat{v}_k(\boldsymbol{\xi})$.

Then, each unknown field q is also approximated on the element E_e with interpolation functions $(w_i(\mathbf{y}))_{1 \leq i \leq N}$: $w_i(\mathbf{y}_j) = \delta_{ij}$ for $1 \leq i, j \leq N$ by

$$q(\mathbf{y}) \approx \sum_{k=1}^{d'} q^k w_k^e(\mathbf{y}),$$

where q^k denotes the approximation of the nodal value of $q(\mathbf{y}_k)$. For isoparametric elements such as \mathbb{P}_1 and \mathbb{P}_2 , $d = d'$, the shape functions serve as interpolation functions as well, and the N interpolation points coincide with the nodes previously defined. For a \mathbb{P}_0 mesh, $d' = 1$, the N interpolation points correspond to the centers of elements and we use piecewise constant interpolation functions.

To discretize the boundary integral equation (2b), we employ the collocation approach. It consists in enforcing the equation at a finite number of collocation points \mathbf{x} . To have a solvable discrete problem, one has to choose N collocation points. The N interpolation points thus defined also serve as collocation points, i.e., $(\mathbf{x}_i)_{1 \leq i \leq N} = (\mathbf{y}_j)_{1 \leq j \leq N}$. In addition, a standard Gaussian quadrature formula is used to evaluate the integral in (2b). This discretization process transforms (2b) into a square complex-valued linear system of size N of the form

$$\mathbb{A}\mathbf{q} = \mathbf{b}, \tag{3}$$

where the (N) -vector \mathbf{q} collects the degrees of freedom (DOFs) while the (N) -vector \mathbf{b} arises from the imposed incident wave field. Assembling the full dense matrix \mathbb{A} requires the computation of all element integrals for each collocation point, thus requiring an $\mathcal{O}(N^2)$ computational time and memory. If an iterative method is used, each GMRES iteration requires one evaluation of $\mathbb{A}\mathbf{q}$ for a given \mathbf{q} , a task requiring a computing time of order $\mathcal{O}(N^2)$. To lower this $\mathcal{O}(N^2)$ complexity, prohibitive for large BEM models, fast BEM solution techniques such as the Fast Multipole Method (FMM) must be employed.

2.2. Fast multipole-accelerated boundary element methods for wave propagation

The goal of the FMM is to accelerate the evaluation of the matrix-vector product $\mathbb{A}\mathbf{q}$ for a given \mathbf{q} , required at each iteration of an iterative solver applied to the BEM-discretized equations. Moreover, the governing BEM matrix is never explicitly formed, which leads to a storage requirement well below $\mathcal{O}(N^2)$. Substantial savings in both computational time and memory requirements are thus achieved. In general terms, the FMM

exploits a reformulation of the fundamental solution in terms of products of functions of \mathbf{x} and of \mathbf{y} , so that (unlike in the traditional BEM) integrations with respect to \mathbf{y} can be reused when the collocation point \mathbf{x} is changed. On decomposing the position vector $\mathbf{r} = \mathbf{y} - \mathbf{x}$ into $\mathbf{r} = (\mathbf{y} - \mathbf{y}_0) + \mathbf{r}_0 - (\mathbf{x} - \mathbf{x}_0)$, where \mathbf{x}_0 and \mathbf{y}_0 are two poles and $\mathbf{r}_0 = \mathbf{y}_0 - \mathbf{x}_0$, and invoking a well-known plane wave decomposition, the Helmholtz fundamental solution is written as [27]:

$$\Phi_k(\mathbf{x}, \mathbf{y}) = \lim_{L \rightarrow +\infty} \int_{\hat{\mathbf{s}} \in S} e^{ik\hat{\mathbf{s}} \cdot (\mathbf{y} - \mathbf{y}_0)} \mathcal{G}_L(\hat{\mathbf{s}}; \mathbf{r}_0; k) e^{-ik\hat{\mathbf{s}} \cdot (\mathbf{x} - \mathbf{x}_0)} d\hat{\mathbf{s}}, \quad (4)$$

where S is the unit sphere of \mathbb{R}^3 and the transfer function $\mathcal{G}_L(\hat{\mathbf{s}}; \mathbf{r}_0; k)$ is defined in terms of Legendre polynomials P_p and spherical Hankel functions of the first kind $h_p^{(1)}$ by:

$$\mathcal{G}_L(\hat{\mathbf{s}}; \mathbf{r}_0; k) = \frac{ik}{16\pi^2} \sum_{0 \leq p \leq L} (2p+1) i^p h_p^{(1)}(k|\mathbf{r}_0|) P_p(\cos(\hat{\mathbf{s}}, \mathbf{r}_0)). \quad (5)$$

It can be shown that expression (4) is valid only for well-separated sets of collocation and integration points clustered around poles \mathbf{x}_0 and \mathbf{y}_0 [27].

In the single-level FMM, a 3D cubic grid of linear spacing d embedding the whole boundary Γ is then introduced to drive the computation. The FMM basically consists in using decomposition (4), with the poles \mathbf{x}_0 and \mathbf{y}_0 being chosen as the cell centres of the cells \mathcal{C}_x and \mathcal{C}_y , whenever \mathbf{x} and \mathbf{y} belong to non-adjacent cubic cells. The treatment of such Fast Multipole (FM) contributions exploits the plane wave decomposition (4) of the fundamental solution, truncated at a finite L and in a manner suggested by its multiplicative form. When \mathbf{x} and \mathbf{y} belong to adjacent cells, traditional BEM evaluation methods are employed. To improve further the computational efficiency of the FM-BEM, standard (i.e., non-FMM) calculations must be confined to the smallest possible spatial regions while retaining the advantage of clustering the computation of influence terms into non-adjacent large groups whenever possible. This is achieved by recursively subdividing cubic cells into eight smaller cubic cells. New pairs of non-adjacent smaller cells, to which plane wave expansions are applicable, are thus obtained from the subdivision of pairs of adjacent cells. The cell-subdivision approach is systematized by means of an octree structure of cells. At each level ℓ , the linear cell size is denoted d^ℓ . The level $\ell = 0$, composed of only one cubic cell containing the whole surface Γ , is the tree root. The subdivision process is further repeated until the finest level $\ell = \bar{\ell}$, implicitly defined by a preset subdivision-stopping criterion ($d^{\bar{\ell}} \geq d^{\min}$), is reached. Level- $\bar{\ell}$ cells are usually termed leaf cells. This is the essence of the multi-level FMM, whose theoretical complexity is $\mathcal{O}(N \log N)$ per GMRES iteration both for computational time and memory requirements. More details on the implementation of the FM-BEM for waves can be found, e.g., in [18] where the method has been successfully employed in elastodynamics.

2.3. Anisotropic mesh adaptation for fast multipole-accelerated boundary element methods

The influence of the mesh discretization on the accuracy of a numerical solution still poses a challenge to the BEM community [9, 5, 20]. Anisotropic features of a solution (e.g., some elastic materials) as well as discontinuities near geometric singularities (e.g., corners and edges) are difficult to capture and ultimately diminish the regularity of the boundary solution and subsequent performance of a BEM. This is particularly true when uniform meshes are employed as singularities will degrade the convergence order. To this end, iterative mesh refinement schemes have been proposed to transform an initial mesh into an improved one according to error estimates calculated at each step, with the goal of reducing the number of degrees of freedom required to resolve a solution within a desired level of accuracy. Most work to date for BEMs has been confined to isotropic techniques based on *a posteriori* error analysis from which error indicators can be derived, see e.g., [5, 9] for wave propagation problems. These indicators steer refinement by systematically marking and subdividing only elements where the error is above a specified threshold—a process known as Dörfler marking. Not only is this process unable to produce anisotropic discretizations, it is also heavily reliant on discretization strategy and model [28] (e.g., often requiring the more computationally expensive Galerkin approach). Finally, this strategy is not optimal for high-order discretizations [9].

We employ here a truly *anisotropic* metric-based boundary mesh adaptation recently introduced [20]. It does not use a Dörfler marking but rather generates a sequence of non-nested meshes. This is a non-linear process that is solved iteratively. At each iteration, it seeks a new mesh that minimizes the L^2 linear interpolation error on the exact solution u .

A convenient framework to generate anisotropic meshes is that of Riemannian metric spaces. A Riemannian metric space is defined by a metric tensor $\mathcal{M}(\mathbf{x})$, i.e., is a smoothly varying function of the physical variable \mathbf{x} . A surface mesh, associated with a Riemannian metric space $\mathbf{M} = (\mathcal{M}(\mathbf{x}))_{\mathbf{x} \in \Gamma}$, is a triangulation of the surface. Each mesh vertex \mathbf{x} has an assigned value $\mathcal{M}(\mathbf{x})$ which dictates the size and orientation of adjacent elements. By generating a unit mesh in the corresponding Riemannian metric space, we obtain an anisotropic mesh refined in Euclidean space according to the metric \mathcal{M} . This is the fundamental idea of metric-based mesh adaptation. Since the existence of a conforming unit mesh in which each triangle is perfectly unit with respect to a given Riemannian metric space is not guaranteed in general, the constraint must be relaxed somewhat. In practice, we seek a quasi-unit mesh with respect to \mathcal{M} which is a mesh composed of quasi-unit triangles. A triangle, which is defined by its edges $\{\mathbf{e}_i\}_{i=1,3}$, is quasi-unit with respect to \mathcal{M} if

$$\frac{1}{\sqrt{2}} \|\mathbf{e}_i\|_{\mathcal{M}}, \quad 1 \leq i \leq 3 \quad \text{and} \quad |K|_{\mathcal{M}} = \frac{\sqrt{3}}{4}.$$

The principle of our adaptive mesh strategy is then to minimize the error coming from the linear interpolation of the (unknown) exact solution. The discrete optimization problem reads:

$$\text{Find the surface mesh } \mathcal{T}_N^{opt} \text{ with } N \text{ nodes, minimizing the discrete interpolation error } \|u - \Pi_h u\|_{L^p(\Gamma_h)}, \quad (6)$$

where Γ_h is the discrete support representing Γ and $\Pi_h u$ is the linear interpolant of u on the mesh \mathcal{T}_N^{opt} . This problem is generally intractable practically and would require some simplifications. The originality of our approach is to replace the discrete interpolation error with its continuous equivalent. Using the mathematical tools defined on the continuous mesh space instead of the discrete one, the optimization problem is simplified. In the continuous mesh setting [52], we need to introduce the complexity \mathcal{N} of a mesh. It is the continuous counterpart of the number of vertices and it governs the size of the meshes, i.e.,

$$\mathcal{N} := \int_{\Gamma} d(\mathbf{x}) d\mathbf{x},$$

where $d = (h_1 h_2)^{-1}$ is the density and $(h_i)_{i=1,2}$ are the local sizes along the principal directions of the metric. The global optimization problem becomes to find the optimal continuous mesh \mathbf{M}^* minimizing the continuous interpolation error $e_{R_s(\mathcal{M})}$ in the L^p norm:

$$\text{Given a complexity } \mathcal{N}, \text{ find } \mathbf{M}^* = \operatorname{argmin}_{\mathbf{M}} \left(\int_{\Gamma} (e_{R_s(\mathcal{M})}(\mathbf{x}))^p d\mathbf{x} \right)^{1/p}. \quad (7)$$

We have shown [20] that the solution to (7) is given by

$$\mathcal{M} = N \left(\int_{\Gamma} \det(|R_s(H)|)^{\frac{1}{3}} \right)^{-1} \det(|R_s(H)|)^{-\frac{1}{6}} |R_s(H)|, \quad (8)$$

where R_s is a restriction operator for surfaces and H is the Hessian of the solution. The mesh generated by this metric minimizes the L^2 -interpolation error.

The above results are valid for the linear interpolant $\Pi_h u$ of the exact solution u . However, neither $\Pi_h u$ or u are known. Only the approximation of u obtained by the FM-BEM is at hand. In order to relate the interpolation error to the approximation error, we consider a reconstruction operator \mathcal{R}_h which is applied to the numerical approximation u^h . This operator can be a recovery process, a hierarchical basis, or an operator connected to an a posteriori estimate. There are numerous techniques in the literature to approximate or “recover” the Hessian (see, e.g., [58, 64]). From a practical point of view, we only need to recover a pointwise 3D Hessian at the nodes of the current surface mesh. The method we employ here is based on a local quadratic recovery and uses the boundary solution only, see [20] for more details.

With the minimizing mesh in hand, one repeats the process for larger complexities until some ad-hoc stopping criterion is achieved. More precisely, given a user-prescribed sequence of complexity $\mathbf{N} = [\mathcal{N}_1, \dots, \mathcal{N}_k]$, we seek the sequence of corresponding optimal meshes $\mathcal{T} = [\mathcal{T}_1, \dots, \mathcal{T}_k]$. This process is non-linear, i.e., both the mesh and the solution have to be converged. The following iterative algorithm is used to generate \mathcal{T}_i :

Start from the mesh $\mathcal{T}_i^0 = \mathcal{T}_{i-1}^0$ (or from the initial uniform mesh \mathcal{T}_0 at iteration 0).

for all $k = 0 : n_{\text{iter}}$ **do**

 Compute the approximation u^h on the mesh \mathcal{T}_i^k with the iterative FM-BEM solver (software COFFEE).

 Compute the recovered Hessian from u^h and deduce Metric (8), with complexity \mathcal{N}_i .

 From (8), a new quasi-unit mesh \mathcal{T}_i^{k+1} of complexity \mathcal{N}_i is generated with Fello.a [50].

end for

The proposed strategy constructs adapted meshes that have been shown to recover optimal convergence rates for domains with corners and edges [20]. The main advantage of this strategy is to provide manageable system sizes for a desired accuracy (to achieve the same with a uniform mesh, one may require a discretization exceeding memory capabilities) to the FM-BEM solver. Importantly, the computation of the metric and subsequent generation of the mesh itself is a rather inexpensive calculation, negligible relative to the corresponding FM-BEM solution time (usually less than 5% of the overall computational cost). It is additionally independent of the PDE and integral equation formulation and hence can be applied to a wide variety of wave propagation problems. Figure 1 illustrates a corner portion of the three iterations of meshes generated for the planar screen geometry considered in Section 4.1. One can observe clear refinement at the edges. Unfortunately, such heavily non-uniform discretizations (note the stretched elements next to regular elements) lead to poorer conditioning of the FM-BEM system [35], further motivating the need for an efficient iterative solver.




Figure 1: An example of the first, third and fifth iterations (in the mesh adaptation process) employed in the anisotropic screen meshes considered in Section 4.1. The complexities input for each case were $\mathcal{N} = 350, 1400$ and 5600 , respectively.

3. A nested GMRES-based preconditioner for adapted FM-BEM

Section 3.1 briefly discusses the theory, construction and deployment of an \mathcal{H} -matrix representation of the BEM system for wave problems. It is subsequently incorporated into a nested GMRES iterative solver described in Section 3.2. Section 3.3 provides several validating examples of the implementation of both the FM-BEM solver and the preconditioner constructed by \mathcal{H} -matrix representations.

3.1. Principles of the \mathcal{H} -matrix representation of the BEM system for wave propagation

Hierarchical matrices or \mathcal{H} -matrices have been introduced by Hackbusch [38] to compute a data-sparse representation of some special dense matrices (e.g., matrices resulting from the discretization of non-local operators). The principle of \mathcal{H} -matrices is (1) to partition the matrix into blocks and (2) to perform low-rank approximations of the blocks of the matrix which are known *a priori* (via an admissibility condition) to be accurately approximated by low-rank decompositions. With these two ingredients it is possible to define fast iterative and direct solvers for matrices having a hierarchical representation. Using low-rank representations, the memory requirements and costs of a matrix-vector product are reduced.

The key component of hierarchical matrices is the recursive block subdivision. The first step prior to the partition of the matrix is a partitioning based on the geometry of the set of row and column indices of the system \mathbb{A} . The objective is to permute the indices in the matrix in order to reflect the physical distance

and hence interaction between degrees of freedom. Consecutive indices should correspond to DOFs that interact at close range. A binary tree \mathcal{T}_I is used to drive the clustering where each node of the tree defines a subset of indices $\sigma \subset I$ and each subset corresponds to a part in the partition of the domain. For each node we determine the box enclosing all the points in the cloud and subdivide it into 2 boxes, along the largest dimension. The subdivision is stopped when a minimum number of DOFs per box is reached (N_{leaf}).

After the clustering of the unknowns is performed, a block cluster representation $\mathcal{T}_{I \times I}$ of the matrix \mathbb{A} is defined by going through the cluster tree \mathcal{T}_I . Each node of $\mathcal{T}_{I \times I}$ contains a pair (σ, τ) of indices of \mathcal{T}_I and defines a block of \mathbb{A} . This uniform partition defines a block structure of the matrix with a full pattern but is not optimal. As a matter of fact, some parts of the matrix \mathbb{A} can be accurately approximated by a low-rank matrix at a higher level (i.e. for larger clusters). Such blocks are said to be *admissible*. A hierarchical representation $\mathcal{P} \subset \mathcal{T}_{I \times I}$ that uses the cluster tree \mathcal{T}_I and the existence of *admissible* blocks is more appropriate. Starting from the initial matrix, each block is recursively subdivided until it is either *admissible* or the leaf level is achieved. For complex 3D geometries, an admissibility condition based on the geometry and the interaction distance between points is used to determine *a priori* the *admissible* blocks. In summary, the blocks of the partition can be of three types: at the leaf level a block can be either an *admissible* block or a *non-admissible* block, at a non-leaf level a block can be either an *admissible* block or an \mathcal{H} -matrix (i.e. a block that will be subsequently hierarchically subdivided).

\mathcal{H} -matrix representations have been derived for some specific problems and will not result in efficient algorithms for all equations or matrices. The crucial point is to know a priori (1) if savings will be obtained when trying to approximate admissible blocks with a sum of separated variable functions and (2) which blocks are admissible since the explicit computation of the rank of all the blocks would be too expensive. In the case of asymptotically smooth kernels $G(\mathbf{x}, \mathbf{y})$, it is proved that under some *a priori* condition on the distance between \mathbf{x} and \mathbf{y} , the kernel is a degenerate function. More precisely, the admissibility condition is given by

$$\max(\text{diam } X, \text{diam } Y) \leq \eta \text{ dist}(X, Y), \quad (9)$$

where $\text{diam } X$ (resp. $\text{diam } Y$) is the diameter of a block domain X (resp. Y), $\text{dist}(X, Y) = \inf\{\|\mathbf{x} - \mathbf{y}\|, \mathbf{x} \in X, \mathbf{y} \in Y\}$ is the Euclidean distance between the domains X and Y and $\eta > 0$ is the admissibility constant of the method (fixed to $\eta = 3$ in the following sections). After discretization, this property is reformulated as the efficient approximation of blocks of the matrix by low-rank matrices. The Laplace Green's function is an example of an asymptotically smooth kernel for which \mathcal{H} -matrix representations have been shown to be very efficient. The Helmholtz Green's functions are not asymptotically smooth but oscillatory. It has been shown in [11] that even though the method is less optimal in that case, substantial memory savings are still achieved with a complexity level lower than $\mathcal{O}(N^{3/2} \log N)$ in practice.

Once the \mathcal{H} -matrix representation of a matrix is computed, it is easy to derive an \mathcal{H} -matrix based iterative solver. The only operation required is an efficient matrix-vector product. It is performed hierarchically by going through the block cluster tree. At the leaf level, there are two possibilities. If the block does not admit a low-rank approximation (*non-admissible* block), then the standard matrix-vector product is used. Otherwise, the block is marked as admissible such that a low-rank approximation (via adaptive cross approximation with stopping tolerance ε_{ACA} corresponding to the accuracy of the low-rank representation) has been computed. The cost of this part of the matrix-vector product is then reduced. The interested reader is referred to [19] for full details concerning the \mathcal{H} -matrix construction for oscillatory Helmholtz kernels.

3.2. A preconditioned adapted FM-BEM solver

The objective of this work is to use the complementary aspects of the fast multipole method and \mathcal{H} -matrix representations of the system. Our overall aim is to reduce total computational cost of the adaptive FM-BEM. Since the FM-BEM solvers have been shown to have an optimal complexity for wave propagation problems, our goal is to supplement the FM-BEM with a preconditioning provided by the \mathcal{H} -matrix representation of the BEM system. The goal is to solve the right preconditioned FM-BEM system (3), i.e., find \mathbb{M} such that

$$\mathbb{A}\mathbb{M}^{-1}\mathbb{M}\mathbf{q} = \mathbf{b}.$$

By doing so, we will solve

$$\mathbb{A}\mathbb{M}^{-1}\mathbf{y} = \mathbf{b} \quad \text{and} \quad \mathbb{M}\mathbf{q} = \mathbf{y}$$

with the goal of improving the conditioning of \mathbb{A} . The best preconditioner is of course $\mathbb{M} = \mathbb{A}$ but the cost of evaluating \mathbb{M}^{-1} would be equivalent to solving the initial system. Instead, we want to use an approximation of \mathbb{A} . In this work, we study the efficiency of the use of an \mathcal{H} -matrix approximation of the BEM system to precondition the FM-BEM system in the context of anisotropic meshes. Since we only need a coarse approximation of \mathbb{M}^{-1} , we will use a two level GMRES. The inner GMRES is used to calculate the preconditioner while the outer GMRES solves the corresponding preconditioned system.

A flexible GMRES [59] allows the preconditioner to change at every iteration, hence permitting the use of nested “inner-outer iterations” which can invoke an iterative method itself as the preconditioner (e.g., another GMRES). Since the subsequent matrix-vector product is very cheap, an \mathcal{H} -matrix based calculation is an ideal candidate for these inner iterations: it is algebraically formed and thus general, all the while respecting the integrity of the BEM system. Algorithm 1 summarizes the nested GMRES implementation employing such a preconditioner applied to system (3). The principle difference from the standard GMRES is found in line 3. Instead of the application of a fixed preconditioner \mathbb{M} (e.g., right-preconditioning $\mathbf{z}_j = \mathbb{M}^{-1}\mathbf{v}_j$), we solve $\mathbb{A}\mathbf{q}_j = \mathbf{v}_j$ approximately by a given iterative method—in this case, GMRES using (coarse) \mathcal{H} -matrix representations. This inner solve step preconditions the outer solve whose matrix-vector products are facilitated by the FMM.

Algorithm 1 Nested GMRES for preconditioning fast BEMs

```

1:  $\mathbf{r}_0 = \mathbf{b} - (\mathbb{A}\mathbf{q}_0)_{\text{FMM}}, \beta = \|\mathbf{r}_0\|_2 \implies \mathbf{v}_1 = \mathbf{r}_0/\beta$  // evaluate the initial residual with FM-BEM
2: for  $j = 1, 2, \dots$  until convergence do // outer GMRES iterations with FM-BEM system
3:   Solve  $\mathbb{A}_{\mathcal{H}\text{mat}}\mathbf{z}_j := \mathbf{v}_j$  // inner GMRES with  $\mathcal{H}$ -mat to define preconditioner
4:    $\mathbf{w} := (\mathbb{A}\mathbf{z}_j)_{\text{FMM}}$  // FMM approximation
5:   for  $i = 1$  to  $j$  do // Construction of the Krylov subspace
6:      $h_{i,j} := \mathbf{z}_i \cdot \mathbf{w}$ 
7:      $\mathbf{w} := \mathbf{w} - h_{i,j}\mathbf{z}_i$ 
8:   end for
9:    $h_{j+1,j} = \|\mathbf{w}\|_2$  and  $\mathbf{v}_{j+1} = \mathbf{w}/h_{j+1,j}$ 
10:   $\mathbf{y}_j = \arg \min_{\mathbf{y}} \|\beta\mathbf{e}_1 - \mathbb{H}(1:j+1, 1:j)\mathbf{y}\|_2$  //  $\mathbb{H}_{ij} = h_{ij}$ , solve minimization pb. to find the update
11:   $\mathbf{q}_j = \mathbf{q}_0 + [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_j]\mathbf{y}_j$  // updated solution
12: end for

```

Figure 2: The nested GMRES employing fast-multipole approximations preconditioned by a GMRES sequence of coarse \mathcal{H} -matrix approximations.

The brief diagram of Figure 3 summarizes the full solver. Initial coarse meshes are constructed to have a density of approximately $d_\lambda = 5$ collocation points per wavelength, leading to a small initial calculation of the boundary solution. The adaptive sequence computes the Hessian and corresponding metric, ultimately generating a quasi-unit mesh within this space. An FM-BEM is then applied to this new mesh to rapidly produce a solution for the next iteration. The metric calculation and mesh generation are very cheap—the complexity of our solver is dominated by the FM-BEM. Hence any improvement in computational cost is felt by every iteration of the adaptive mesh procedure.

3.3. Implementation details and validations of the preconditioner strategy

All the aforementioned algorithms have been implemented in the FORTRAN language and within the in-house accelerated boundary element software for elastodynamics and acoustics known as COFFEE. Numerical experiments were conducted on a high-performance computer equipped with four multicore Intel Xeon E7-4820 processors for a total of 32 cores. All anisotropic meshes in Section 4 are constructed by the metric-based adaptation procedure (Section 2.3) using Metrix and Feflo.a software packages from the Gamma3 team at INRIA Saclay [50, 51, 52].

As a verification of the FM-BEM code (whose FMM-speedups for piecewise linear elements have already been demonstrated in [20, 18]), we consider diffraction by a sphere of radius R of an incident field with




Figure 3: Summary of the adaptive FM-BEM solver with preconditioning.

direction $\mathbf{d} = (-1, 0, 0)^T$ and a non-dimensional wavenumber $kR = 10$ for a sequence of uniformly-refined discretizations employing $\mathbb{P}_0, \mathbb{P}_1$ and curved \mathbb{P}_2 elements. The FM-BEM system is solved by using the simplest preconditioner, i.e., the right-diagonally preconditioned GMRES [60] algorithm which scales all the diagonal entries of the system matrix to 1. We hence solve the system (3) as

$$\mathbb{A}\mathbb{M}^{-1}\mathbf{y} = \mathbf{b}, \quad \mathbf{q} = \mathbb{M}^{-1}\mathbf{y} \quad (10)$$

for a preconditioner \mathbb{M} defined as

$$\mathbb{M}_{ij} = \begin{cases} \mathbb{A}_{ij}, & \text{if } i = j, \\ 0, & \text{if } i \neq j. \end{cases} \quad (11)$$

Figure 4 reports the L^2 -norm of the error, between the approximated surface density q_N and the analytical density q for the integral equation on the boundary of a sphere for a fixed number of elements N_E using $\mathbb{P}_0, \mathbb{P}_1$ and \mathbb{P}_2 interpolations and given by

$$\frac{(\int_{\partial\Omega} (q - q_N)^2 d\mathbf{x})^{1/2}}{(\int_{\partial\Omega} q^2 d\mathbf{x})^{1/2}}.$$

There is little error analysis for 3D collocation BEM solvers but there exist some empirical studies, see, e.g. [3]. As this is a case of a smooth geometry, one readily appreciates the recovery of the expected behaviour of the L^2 convergence order in all three cases.

We fix the GMRES tolerance to 10^{-4} , which is a common choice in the BEM community. It has also been shown to produce the expected convergence rates in terms of errors between the discrete and exact physical solution, in problems of interest in this work [20]. Since (11) is a simple preconditioning technique commonly used for single-layer operators, we henceforth refer to solutions using this diagonal-preconditioner as the original (“regular”) “FM-BEM” and to solutions using the preconditioning nested GMRES—implemented with \mathcal{H} -matrix representations—as “preconditioned FM-BEM” (denoted by “prec




Figure 4: *Diffraction by a uniformly discretized sphere*: convergence results of the computed density on the boundary for $kR = 10$ for discretizations with \mathbb{P}_0 , \mathbb{P}_1 and curved \mathbb{P}_2 elements.

FM-BEM” in plots). To verify the preconditioner implementation, we again consider the unit sphere with an incident field of direction $\mathbf{d} = (-1, 0, 0)^T$, but for non-dimensional wavenumbers $kR = 17, 24, 30$ and 43 . We uniformly discretize the surface for each wavenumber using \mathbb{P}_1 elements (with a density d_λ of approximately 15 collocation points per wavelength) to obtain a corresponding series of meshes ranging in size from $N \approx 30\,000$ to $N \approx 180\,000$ vertices. Figure 5 presents an illustrative discretization as well as the solution (real part and complex amplitude) for the case of $kR = 43$. For each discretization, \mathcal{H} -matrix representations are calculated with a coarse resolution of $\varepsilon_{\text{ACA}} = 10^{-1}$, admissibility condition (9) with $\eta = 3$ and a minimum block size stopping criterion $N_{\text{leaf}} = 100$. These parameters are chosen empirically through a series of numerical experiments, a version of which is presented in this section for uniformly discretized planar screens and later for anisotropically discretized planar screens. Each case is iteratively resolved either by a (right) diagonal-preconditioned FM-BEM (11) (“FM-BEM”) or the nested \mathcal{H} -matrix preconditioned FM-BEM (“prec FM-BEM”), with the (outer) GMRES tolerance fixed to 10^{-4} (500 maximum iterations) and the inner tolerance of the nested GMRES fixed to 10^{-1} (10 maximum inner iterations).




Figure 5: *Diffraction by a uniformly discretized sphere*: a portion of the mesh; the real part of the boundary solution; and the complex amplitude of the boundary solution for $kR = 43$ ($N \approx 178\,000$).

Defining the computational setup time for the FM-BEM system and \mathcal{H} -matrix representations (for an approximation tolerance $\varepsilon = \varepsilon_{\text{ACA}}$) as $t_{\text{setup}}^{\text{FMM}}$ and $t_{\text{setup}}^{\mathcal{H}_{\text{mat}}, \varepsilon}$, respectively, the corresponding setup (precompu-

tation) times are given for FM-BEM and preconditioned FM-BEM respectively as

$$\begin{aligned} t_{\text{setup}} &= t_{\text{setup}}^{\text{FMM}}, \\ t_{\text{setup}}^{\text{prec}} &= t_{\text{setup}}^{\text{FMM}} + t_{\text{setup}}^{\mathcal{H}\text{mat},\varepsilon}. \end{aligned} \quad (12)$$

The left plot of Figure 6 presents the corresponding times for each discretization and non-dimensional wavenumber of the sphere example. The addition of an \mathcal{H} -matrix construction for the nested preconditioner results in a slightly longer precomputation time. Although the complexity of constructing the \mathcal{H} -matrix representation is generally higher than the FMM complexity, it can be shown [7, 19] that the growth in complexity is bounded as $\mathcal{O}(N^{3/2} \log N)$. This property is also verified in the figure where we have additionally plotted the setup time $t_{\text{setup}}^{\mathcal{H}\text{mat},\varepsilon}$. The right plot gives corresponding costs of a single (outer) iteration of FM-BEM and preconditioned FM-BEM defined as

$$\begin{aligned} t_{\text{iter}} &= t_{\text{iter}}^{\text{FMM}}, \\ t_{\text{iter}}^{\text{prec}} &= t_{\text{iter}}^{\text{FMM}} + N_{\text{iter}}^{\mathcal{H}\text{mat}} t_{\text{iter}}^{\mathcal{H}\text{mat}}, \end{aligned} \quad (13)$$

where $N_{\text{iter}}^{\mathcal{H}\text{mat}}$ is the number of iterations of the inner sequence of GMRES facilitated by \mathcal{H} -matrix representations. One similarly observes that the single iteration times are longer for the nested GMRES since they include FMM matrix-vector computations as well as inner \mathcal{H} -matrix computations (which are themselves rather cheap and are applied a maximum of ten times an outer iteration). The maximum numerical rank among all admissible blocks in the \mathcal{H} -matrix representation is provided in parentheses and grows with complexity as expected.




Figure 6: *Diffraction by a uniformly discretized sphere*: setup times ($t_{\text{setup}}, t_{\text{setup}}^{\text{prec}}$) and single iteration times ($t_{\text{iter}}, t_{\text{iter}}^{\text{prec}}$) applied to meshes of fixed density $d_\lambda = 15$. The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

We then report in Table 1 the global time-to-solution given for each case by

$$\begin{aligned} t_{\text{global}} &= N_{\text{iter}} t_{\text{iter}}^{\text{FMM}} + t_{\text{setup}}^{\text{FMM}}, \\ t_{\text{global}}^{\text{prec}} &= N_{\text{iter}}^{\text{prec}} (t_{\text{iter}}^{\text{FMM}} + N_{\text{iter}}^{\mathcal{H}\text{mat}} t_{\text{iter}}^{\mathcal{H}\text{mat}}) + t_{\text{setup}}^{\text{FMM}} + t_{\text{setup}}^{\mathcal{H}\text{mat},\varepsilon}. \end{aligned} \quad (14)$$

In spite of these marginal increases in setup and single iteration times, it is significantly less for the preconditioned FM-BEM owing to the significant reduction in the number of overall (outer GMRES) iterations $N_{\text{iter}}^{\text{prec}}$ when resolving the FM-BEM system. Defining the corresponding speedup S of preconditioned FM-BEM versus FM-BEM,

$$S = \frac{t_{\text{global}} - t_{\text{global}}^{\text{prec}}}{t_{\text{global}}} \times 100,$$

one already observes during the verification the larger than 50% speedup in computation time when employing the preconditioned solver for uniformly discretized spheres. The computational gains are demonstrated in Section 4 to be even more significant for the anisotropically-discretized meshes and higher-order numerical quadratures that are of interest.

N	kR	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
27 870	17	84	9m56s	23	4m14s	57%
54 117	24	89	25m32s	22	10m05s	61%
87 615	30	115	59m00s	32	26m36s	55%
178 472	43	209	188m09s	58	84m09s	55%

Table 1: *Diffraction by a uniformly discretized sphere*: number of solver iterations and global solution times (including setup) for meshes of fixed density $d_\lambda = 15$.

Geometric singularities and the subsequent discontinuous behavior of the corresponding boundary solution can lead to a significant performance loss in a BEM. An infinitely thin planar screen presents a challenging example in this context: it has sharp and thin geometric features that encourage, like many realistic configurations, a mesh adaptivity suitable to capture the behavior of the wave propagation around these geometric singularities. This is a case of particular interest in Section 4.1 and so in this spirit we consider a series of verifying numerical examples for scattering of plane waves of direction $\mathbf{d} = (1, 0, 0)^T$ by a uniformly discretized screen of characteristic length, i.e., edge length L .




Figure 7: *Diffraction by a uniformly discretized screen*: a portion of the mesh; the real part of the boundary solution; and the complex amplitude of the boundary solution for $kL = 43$ ($N \approx 30\,000$).

As in for the sphere, we consider the screen uniformly discretized by piecewise-linear elements as seen in Figure 7, which presents an illustrative discretization as well as a solution (real part and complex amplitude) for the case of $kL = 43$. A density d_λ of approximately 15 collocation points per wavelength is chosen to obtain a corresponding series of meshes ranging in size from $N \approx 300\,000$ to $N \approx 500\,000$ vertices. The corresponding setup and single-iteration times are presented in Figure 8. Again we can observe that the addition of an \mathcal{H} -matrix construction for the nested preconditioning results in a slightly longer precomputation time due to the \mathcal{H} -matrix representation, but is coherent with theoretical complexity. Similarly again, the iteration times are longer as well since they include an FMM matrix-vector computation as well as inner \mathcal{H} -matrix computations (the numerical rank of the largest approximated admissible block is included in parentheses). The global time-to-solution given in Table 2, however, is significantly less for the preconditioned FM-BEM solver owing to the significant reduction in the number of overall (outer) GMRES iterations $N_{\text{iter}}^{\text{prec}}$ when resolving the FM-BEM system, enabling more than 50% reduction in the global computation time.

We now consider the efficacy of the nested GMRES-based preconditioner with respect to the parameters most important to its additional computational time, i.e., that of the construction of the \mathcal{H} -matrix representation. In order to optimize performance, a balance must be struck between precomputational cost and

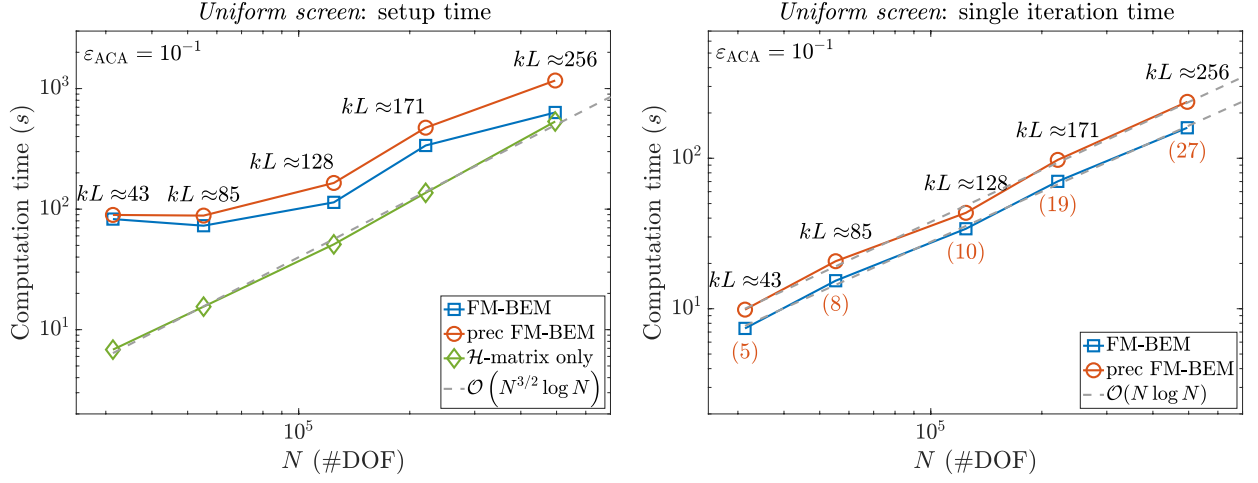


Figure 8: *Diffraction by a uniformly discretized screen*: setup times ($t_{\text{setup}}, t_{\text{setup}}^{\text{prec}}$) and single iteration times ($t_{\text{iter}}, t_{\text{iter}}^{\text{prec}}$) applied to meshes of fixed density $d_\lambda = 15$. The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

N	kL	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
31365	43	45	6m56s	15	3m54s	45%
55224	85	45	12m44s	13	5m55s	54%
124497	128	48	29m06s	13	12m01s	59%
220888	171	51	65m20s	13	28m48s	56%
495774	256	56	159m30s	13	70m16s	56%

Table 2: *Diffraction by a uniformly discretized screen*: number of solver iterations and global solution times (including setup) for meshes of fixed density $d_\lambda = 15$.

the number of iterations during which the corresponding matrix-vector product is applied. Similar studies are presented for anisotropic discretizations of the screen in Section 4.1. For each uniform discretization we solve the resulting system via FM-BEM and the preconditioned FM-BEM system for various non-dimensional wavenumbers and various accuracies for the low-rank approximations in the \mathcal{H} -matrix representation. For a fixed non-dimensional wavenumber $kL = 50$ and low-rank approximation tolerances of

$$\varepsilon_{ACA} = [10^{-1}, 2 \cdot 10^{-1}, 5 \cdot 10^{-1}],$$

Figure 9 presents setup and single iteration times for a series of meshes ranging in size from $N \approx 30\,000$ to $N \approx 200\,000$ (these are no longer of fixed density). These parameters are of interest because they are very coarse approximations and are relatively quick to precompute. One can observe similar performance for each tolerance ε_{ACA} , with the finest tolerance ε_{ACA} having the shortest single iteration time of the resulting nested GMRES iterative solver.

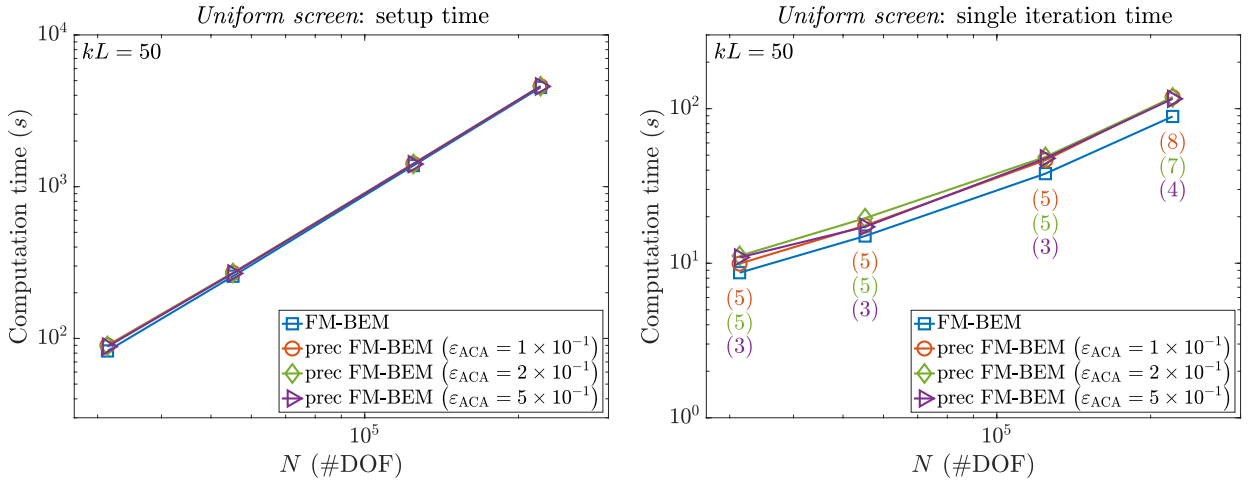


Figure 9: *Diffraction by a uniformly discretized screen*: setup times ($t_{\text{setup}}, t_{\text{setup}}^{\text{prec}}$) and single iteration times ($t_{\text{iter}}, t_{\text{iter}}^{\text{prec}}$) for non-dimensional wavenumber $kL = 50$ and different \mathcal{H} -matrix approximation tolerances ε_{ACA} . The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

Fixing $\varepsilon_{ACA} = 10^{-1}$, we consider construction times and subsequent speedups for varying non-dimensional wavenumbers

$$kL = [25, 50, 100].$$

It is known that the \mathcal{H} -matrix representations in their current form are not optimized for very high non-dimensional wavenumbers [19]. Table 3, which presents speedups over unpreconditioned FM-BEM, indicates however that this possibly is not the case when employing a coarse approximation only as a preconditioner; indeed, one can appreciate the significant speedups in all three cases, including for increasing wave numbers.

N	$kL = 25$		$kL = 50$		$kL = 100$	
	$N_{\text{iter}}^{\text{prec}}$	S	$N_{\text{iter}}^{\text{prec}}$	S	$N_{\text{iter}}^{\text{prec}}$	S
31 365	15	36%	16	48%	15	49%
55 224	13	39%	14	51%	14	51%
124 497	14	33%	13	44%	13	45%
220 888	14	31%	13	39%	13	39%

Table 3: *Diffraction by a uniformly discretized screen*: number of solver iterations and global solution times (including setup) for non-dimensional wavenumbers $kL = 25, 50$ and 100 .

Remark 3.1. The experiments above have employed modest numbers of Gaussian quadrature points to perform the numerical integrations of triangles required to construct the BEM system. For adapted meshes containing highly stretched, flat and anisotropic elements, this can be cumbersome as there may not be an adequate distribution of points, ultimately leading to poorer convergence of the iterative system. A larger number (13 points) of quadrature points may need to be used and are, in fact, required to maintain the properties of using higher-order element discretizations in a BEM. The rest of the experiments in this paper, unless otherwise noted, utilize a higher order of Gaussian numerical quadrature. This may increase certain computational costs but will reduce the number of iterations. Furthermore, as demonstrated in Table 4 for the above sphere and screen examples, the speedup obtained by the preconditioned FM-BEM may be even greater for more accurate quadrature routines.

N	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
27 870	55	17m52s	11	6m34s	63%
54 117	58	54m22s	12	19m45s	64%
87 615	76	129m18s	18	50m29s	61%
178 472	137	355m58s	35	130m47s	63%
N	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
31 365	30	17m15s	3	7m12s	58%
55 224	30	28m00s	3	9m16s	67%
124 497	32	50m30s	3	16m53s	67%
220 888	33	128m16s	3	40m15s	69%

Table 4: Number of solver iterations and global solution times (including setup) for discretizations of fixed density $d_\lambda = 15$ for the sphere (top) and screen (bottom) using 13 Gaussian quadrature points in integrations over triangles.

4. Efficiency of the preconditioner in the presence of an anisotropic mesh

The sound-soft boundary condition (1b), when compared to other kinds of boundary conditions, gives rise to more severe singularities in the boundary solution when considering scattering by obstacles with corners or edges. This suggests the use of an anisotropic mesh adaptation procedure to account for the loss of regularity of the corresponding boundary solutions and hence order of convergence. Unfortunately, both isotropically and anisotropically adapted meshes worsen the conditioning of the overall BEM system, leading to an increase in the number of iterations and thus overall computational time required to iteratively resolve the solution [35]. This naturally motivates the use of an effective preconditioner to accommodate the subsequent increase in computational cost. Furthermore, our anisotropic mesh adaptation procedure is itself iterative and requires several FM-BEM solves; hence any decrease in computational time is felt multifold. In what follows, we consider the efficacy of the \mathcal{H} -matrix based nested GMRES preconditioning on such adapted meshes constructed for problems of scattering by a planar screen and a more complex F-15 aircraft.

4.1. Diffraction by a screen

For the thin planar screen example of Section 3.3, the usage of an adaptation algorithm and, in particular an anisotropic one, is prudent: it is well known (see, e.g., [42]) that at the edge of a screen, the jump behaves as

$$\left[\frac{\partial u}{\partial \mathbf{n}} \right] \sim (kr)^{-0.5} \quad \text{as } kr \rightarrow 0, \quad (15)$$

where r is the distance from the edge. At the corner of the screen, this singularity is more severe, and has been shown [6] to take the form

$$\left[\frac{\partial u}{\partial \mathbf{n}} \right] \sim (kr)^{-0.704} \quad \text{as } kr \rightarrow 0.$$

Owing to the singular behavior, the approximated solution on a sequence of uniformly refined meshes should yield a sub-optimal convergence rate. Hence planar screens are an excellent example of the capabilities of an (iterative) adapted mesh strategy, an example of which is given in Figure 1 for a corner of the screens used in the subsequent experiments. One observes in these images highly stretched and anisotropic elements near the corners and edges. These ultimately degrade the performance of an iterative solver, and hence it is of great interest to study preconditioning in this context.




Figure 10: *Diffraction by an anisotropically discretized screen*: setup times ($t_{\text{setup}}, t_{\text{setup}}^{\text{prec}}$) and single iteration times ($t_{\text{iter}}, t_{\text{iter}}^{\text{prec}}$) for non-dimensional wavenumber $kL = 10$ and different \mathcal{H} -matrix approximation tolerances ε_{ACA} . The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

Figure 10 presents the setup time and single iteration times for three levels of accuracy in \mathcal{H} -matrix representations, i.e.

$$\varepsilon_{\text{ACA}} = [10^{-1}, 2 \cdot 10^{-1}, 5 \cdot 10^{-1}],$$

for a sequence of anisotropic meshes constructed with the mesh adaptation procedure for an incident plane wave of non-dimensional wavenumber $kL = 10$. There are expected slight increases in both setup time and iteration time for each choice of ε_{ACA} . We observe each of these cases is almost identical in cost (since they are very similar in the maximum low rank approximation), and so we adopt henceforth $\varepsilon_{\text{ACA}} = 10^{-1}$ in order to remain consistent with parameters found for uniform discretizations.

In addition, we report in Table 5 the characteristics of the anisotropic meshes, i.e., maximum and minimum length size together with anisotropic ratios. If we define γ_K to be the ratio of the maximum edge length to minimum edge length in a triangle K , then we define the anisotropic ratio γ of a mesh \mathcal{T}_N as

$$\gamma_{\max} = \max_{K \in \mathcal{T}_N} \gamma_K, \quad \gamma_{\text{mean}} = \frac{1}{\#\mathcal{T}_N} \sum_{K \in \mathcal{T}_N} \gamma_K,$$

where one can observe that the more anisotropic the mesh, the larger the value of γ_K . It should be noted that these experiments are conducted on several meshes of a single adaptation sequence (i.e., as discretization increases so does the level of anisotropy).

N	Min length	Max length	γ_{mean}	γ_{max}	$\varepsilon = 10^{-1}$		$\varepsilon = 2 \cdot 10^{-1}$		$\varepsilon = 5 \cdot 10^{-1}$	
					$N_{\text{iter}}^{\text{prec}}$	S	$N_{\text{iter}}^{\text{prec}}$	S	$N_{\text{iter}}^{\text{prec}}$	S
107	$8.63 \cdot 10^{-2}$	$1.64 \cdot 10^{-1}$	$1.80 \cdot 10^0$	$1.90 \cdot 10^0$	3	60%	3	58%	3	62%
705	$9.53 \cdot 10^{-3}$	$1.82 \cdot 10^{-1}$	$1.29 \cdot 10^1$	$1.91 \cdot 10^1$	3	58%	3	58%	3	58%
1535	$2.14 \cdot 10^{-3}$	$1.71 \cdot 10^{-1}$	$5.35 \cdot 10^1$	$7.98 \cdot 10^1$	3	57%	3	56%	3	57%
3370	$5.01 \cdot 10^{-4}$	$1.52 \cdot 10^{-1}$	$1.88 \cdot 10^2$	$3.04 \cdot 10^2$	3	45%	3	45%	4	40%

Table 5: *Diffraction by an anisotropically discretized screen*: number of solver iterations and speedups for non-dimensional wavenumber $kL = 10$ and varying accuracies ε_{ACA} .

Figure 11 indicates setup times and single iteration times for the cases of $kL = 50, 100$ and 200 , respectively. Each case carries a different series of meshes which are refined according to the mesh adaptation procedure for their respective wavelengths. Again, we observe slight increases in setup time, but the iteration time increases with level of anisotropy and size of the grid. Table 6 shows that, although with the more marginal increases in computation time, the method achieves even more significant speedups than that of a uniform discretization: on the order of 70% in most cases. This demonstrates the effectiveness of the preconditioned FM-BEM solver. One additionally observes a smaller speedup for finer (and more anisotropic discretizations). This is due to the non-optimality of the \mathcal{H} -matrix representation for non-uniform discretizations. For example, the stopping criterion N_{leaf} (minimal box size in the subdivision) must be tweaked to accommodate highly stretched elements.

$kL = 50$									
N	Min length	Max length	γ_{mean}	γ_{max}	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
2221	$1.57 \cdot 10^{-2}$	$3.52 \cdot 10^{-2}$	$2.16 \cdot 10^0$	$2.24 \cdot 10^0$	19	0m59s	3	0m18s	69%
8156	$1.81 \cdot 10^{-3}$	$5.56 \cdot 10^{-2}$	$2.13 \cdot 10^1$	$3.08 \cdot 10^1$	28	4m15s	3	1m14s	71%
18840	$2.18 \cdot 10^{-4}$	$6.33 \cdot 10^{-2}$	$1.54 \cdot 10^2$	$2.91 \cdot 10^2$	34	12m21s	3	4m40s	62%
$kL = 100$									
N	Min length	Max length	γ_{mean}	γ_{max}	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
8549	$7.42 \cdot 10^{-3}$	$1.84 \cdot 10^{-2}$	$2.38 \cdot 10^0$	$2.48 \cdot 10^0$	22	5m15s	3	1m16s	76%
33087	$6.85 \cdot 10^{-4}$	$2.73 \cdot 10^{-2}$	$2.17 \cdot 10^1$	$3.98 \cdot 10^1$	31	19m56s	4	5m53s	71%
76666	$3.86 \cdot 10^{-5}$	$2.75 \cdot 10^{-2}$	$2.12 \cdot 10^2$	$7.13 \cdot 10^2$	37	57m46s	3	23m05s	60%
$kL = 200$									
N	Min length	Max length	γ_{mean}	γ_{max}	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
34241	$3.33 \cdot 10^{-3}$	$9.31 \cdot 10^{-3}$	$2.56 \cdot 10^0$	$2.80 \cdot 10^0$	25	29m37s	3	6m34s	78%
134701	$2.69 \cdot 10^{-4}$	$1.54 \cdot 10^{-2}$	$2.78 \cdot 10^1$	$5.75 \cdot 10^1$	34	99m01s	3	26m48s	73%
312687	$1.72 \cdot 10^{-5}$	$1.21 \cdot 10^{-2}$	$1.82 \cdot 10^2$	$7.05 \cdot 10^2$	39	266m27s	4	121m40s	54%

Table 6: *Diffraction by an anisotropically discretized screen*: number of solver iterations and global solution times (including setup) for non-dimensional wavenumbers $kL = 50, 100$ and 200 .

Finally, it is important to recall that we have used in all examples a simple diagonal-preconditioner even in the solutions denoted “FM-BEM”. The solutions corresponding to the nested GMRES implemented with H-matrix representations are denoted “prec FM-BEM”. In Table 7 we report also the number of iterations and solution times for the diffraction by an anisotropically discretized screen with $kL = 50$. These results illustrate the already important improvement of the conditioning due to the diagonal preconditioner, leading to an even better behaviour of our nested GMRES preconditioner.




Figure 11: Diffraction by an anisotropically discretized screen: setup times (t_{setup} , $t_{\text{setup}}^{\text{prec}}$) and single iteration times (t_{iter} , $t_{\text{iter}}^{\text{prec}}$) for non-dimensional wavenumbers $kL = 50, 100$ and 200 . The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

N	N_{iter} (no diag prec)	t_{global} (no diag prec)	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$
2 221	26	1m23s	19	0m59s	3	0m18s
8 156	62	8m57s	28	4m15s	3	1m14s
18 840	124	38m37s	34	12m21s	3	4m40s

Table 7: *Diffraction by an anisotropically discretized screen*: number of solver iterations and global solution times (including setup) for non-dimensional wavenumbers $kL = 50$ utilizing no preconditioner, a diagonal-preconditioner and the nested GMRES preconditioner.

4.2. Diffraction by an F-15 aircraft

We revisit the geometry of the unarmed F-15 aircraft (of characteristic length L in the longest direction) previously considered by the mesh adaptation algorithm in [20] for an incident plane wave with direction $\mathbf{d} = (1/2, 1, 0)^T$ and non-dimensional wavenumber $kL = 114$. A series of adapted meshes are considered from a sequence of 15 that are generated (specifically, the third, sixth and ninth). Each mesh is an increasingly anisotropic discretization by piecewise-linear elements as seen in Figure 12, which presents an illustrative discretization of the anisotropic adaptation on a portion of the wing as well as the boundary solution (real part and complex amplitude). We see the high directionality of the diffracted waves with different patterns for the front and back of the aircraft, and most of the resulting anisotropy in the meshes are reached in areas where the diffracted waves have low regularity.




Figure 12: *Diffraction by an F-15 aircraft*: a portion of the adapted mesh where a wing meets the body; the real part of the boundary solution; and the complex amplitude of the boundary solution for $kL = 114$ ($N \approx 25\,000$).

The corresponding setup and single-iteration times are presented in Figure 13. Again we can observe that the addition of an \mathcal{H} -matrix construction for the nested preconditioning results in a slightly longer precomputation time due to the \mathcal{H} -matrix representation. Similarly again, the single iteration times are longer as well since they include an FMM matrix-vector computation as well as inner \mathcal{H} -matrix computations. The largest approximated admissible low-rank block is included in parentheses. The global time-to-solution given in Table 8, however, is significantly less for the preconditioned FM-BEM solver owing to the significant reduction in the number of overall (outer) GMRES iterations when resolving the FM-BEM system, enabling up to more than 80% reduction in global computation time, significantly better performance than the uniform discretizations treated in Section 3.3.

Remark 4.1. As discussed in Remark 3.1, a number of 13 points was used to perform numerical integration for the experiments in this section. For the sake of consistency, we can also consider the parameters of [20] where three points are used for non-singular integration over a triangle. The resulting linear system requires significantly more iterations to resolve as seen in Table 9 due to the stretching of elements. Indeed, even in this case, speedups up to 80% are achieved when employing the preconditioned FM-BEM.




Figure 13: *Diffraction by an anisotropically discretized F-15 aircraft*: setup times ($t_{\text{setup}}, t_{\text{setup}}^{\text{prec}}$) and single iteration times ($t_{\text{iter}}, t_{\text{iter}}^{\text{prec}}$) for non-dimensional wavenumber $kL = 114$. The maximum numerical rank along all admissible blocks in the \mathcal{H} -matrix representation is given in parentheses.

N	Min length	Max length	γ_{mean}	γ_{max}	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
10 562	$2.96 \cdot 10^{-3}$	$9.51 \cdot 10^{-1}$	$2.03 \cdot 10^2$	$3.21 \cdot 10^2$	210	59m06s	18	10m19s	83%
15 687	$1.97 \cdot 10^{-3}$	$8.23 \cdot 10^{-1}$	$3.22 \cdot 10^2$	$4.18 \cdot 10^2$	180	66m48s	21	18m23s	73%
25 428	$9.75 \cdot 10^{-4}$	$1.14 \cdot 10^0$	$7.44 \cdot 10^2$	$1.17 \cdot 10^3$	180	97m41s	25	38m22s	61%

Table 8: *Diffraction by an F-15 aircraft*: number of solver iterations and global solution times (including setup) for non-dimensional wavenumber $kL = 114$.

5. Conclusions and outlook

We have presented a nested GMRES-based adaptive fast BEM to model 3D wave propagation problems. It relies on an inner GMRES sequence of coarse and very fast \mathcal{H} -matrix approximated matrix-vector products to precondition the Fast Multipole accelerated Boundary Element system used in the outer GMRES. We have demonstrated a reduction in the (outer) GMRES iterations and hence the subsequent invocations of the more expensive FM-based matrix-vector product, resulting in significant global time-to-solution speedups versus conventionally-preconditioned GMRES methods. In particular, we have studied the preconditioned solver applied to the problem of diffraction by an infinitely thin planar screen and an F-15 aircraft, both for uniform and highly non-uniform discretizations. The algorithm has been shown to provide significant speedups for most configurations including up to 80% for some anisotropic meshes.

Efficient construction and application of the \mathcal{H} -matrix representations in general is still very much a work in progress for wave propagation. The current implementation is promising and is already useful to improve the cost of a mesh adaptation procedure. Further work will involve adapting the \mathcal{H} -matrix construction specifically for highly anisotropic meshes including those with curved elements. An obvious first step would be to alter the distance function for partitioning based on the anisotropy of the mesh—this information can be provided by a metric that is derived from the adaptation procedure. And, finally, the methodology herein can be extended for double- and combined-layer formulations of the boundary integral equations, as well as extended for more complex physics such as elastodynamics problems; both the \mathcal{H} -matrix approximation and the anisotropic mesh procedure are independent of choice of PDE and integral equation.

Acknowledgements

This work has been funded by the Agence Nationale de la Recherche (RAFFINE project, ANR-12-MONU-0021). The research of F. Amlani has been partially funded by DGA, Direction Générale de l’Armement

N	Min length	Max length	γ_{mean}	γ_{max}	N_{iter}	t_{global}	$N_{\text{iter}}^{\text{prec}}$	$t_{\text{global}}^{\text{prec}}$	S
10 562	$2.96 \cdot 10^{-3}$	$9.51 \cdot 10^{-1}$	$2.03 \cdot 10^2$	$3.21 \cdot 10^2$	395	59m54s	26	10m32s	82%
15 687	$1.97 \cdot 10^{-3}$	$8.23 \cdot 10^{-1}$	$3.22 \cdot 10^2$	$4.18 \cdot 10^2$	527	93m27s	32	19m26s	79%
25 428	$9.75 \cdot 10^{-4}$	$1.14 \cdot 10^{-0}$	$7.44 \cdot 10^2$	$1.17 \cdot 10^3$	726	169m49s	76	79m27s	53%

Table 9: *Diffraction by an F-15 aircraft*: number of solver iterations and global solution times (including setup) for $kL = 114$ using 3 Gaussian quadrature points in regular integrations over triangles. The number of iterations reported here is slightly larger than those reported in [20]; this is due to the latter use of an initial guess which is not employed in this experiment.

(Grant number 2015 60 0006).

References

- [1] M. AINSWORTH AND J. ODEN, *A posteriori error estimation in finite element analysis*, vol. 37, John Wiley & Sons, 2011.
- [2] X. ANTOINE, A. BENDALI, AND M. DARBAS, *Analytic preconditioners for the boundary integral solution of the scattering of acoustic waves by open surfaces*, Journal of Computational Acoustics, 13 (2005), pp. 477–498.
- [3] K. ATKINSON, *The numerical solution of integral equations of the second kind.*, Cambridge Monographs on Applied and Computational Mathematics, (1997).
- [4] M. AURADA, S. FERRAZ-LEITE, AND D. PRAETORIUS, *Estimator reduction and convergence of adaptive BEM*, Applied Numerical Mathematics, 62 (2012), pp. 787 – 801.
- [5] M. BAKRY, S. PERNET, AND F. COLLINO, *A new accurate residual-based a posteriori error indicator for the BEM in 2D-acoustics*, Computers & Mathematics with Applications, 73 (2017), pp. 2501–2514.
- [6] Z. P. BAŽANT, *Three-dimensional harmonic functions near termination or intersection of gradient singularity lines: a general numerical method*, International Journal of Engineering Science, 12 (1974), pp. 221–243.
- [7] M. BEBENDORF, *Hierarchical Matrices: A Means to Efficiently Solve Elliptic Boundary Value Problems*, Springer, 2008.
- [8] J. BERENGER, *A perfectly matched layer for the absorption of electromagnetic waves*, Journal of computational physics, 114 (1994), pp. 185–200.
- [9] A. BESPALOV, T. BETCKE, A. HABERL, AND D. PRAETORIUS, *Adaptive BEM with optimal convergence rates for the Helmholtz equation*, Computer Methods in Applied Mechanics and Engineering, 346 (2019), pp. 260–287.
- [10] M. BONNET, *Boundary integral equations methods in solids and fluids*, John Wiley and sons, 1999.
- [11] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Hierarchical matrices*, tech. rep., Max Planck Institute for Mathematics in the Sciences, 01 2003.
- [12] O. BRUNO, C. GEUZAIN, J. MONRO, AND F. REITICH, *Prescribed error tolerances within fixed computational times for scattering problems of arbitrarily high frequency: the convex case*, Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 362 (2004), pp. 629–645.
- [13] O. BRUNO AND L. KUNYANSKY, *A fast, high-order algorithm for the solution of surface scattering problems: basic implementation, tests, and applications*, Journal of Computational Physics, 169 (2001), pp. 80–110.

- [14] Y. CAO, L. WEN, J. XIAO, AND Y. LIU, *A fast directional BEM for large-scale acoustic problems based on the Burton–Miller formulation*, *Engineering Analysis with Boundary Elements*, 50 (2015), pp. 47–58.
- [15] B. CARPENTIERI, *A matrix-free two-grid preconditioner for solving boundary integral equations in electromagnetism*, *Computing*, 77 (2006), pp. 275–296.
- [16] B. CARPENTIERI, I. DUFF, L. GIRAUD, AND G. SYLVAND, *Combining fast multipole techniques and an approximate inverse preconditioner for large electromagnetism calculations*, *SIAM Journal on Scientific Computing*, 27 (2005), pp. 774–792.
- [17] B. CARPENTIERI, I. S. DUFF, AND L. GIRAUD, *Sparse pattern selection strategies for robust frobenius-norm minimization preconditioners in electromagnetism*, *Numerical Linear Algebra with Applications*, 7 (2000), pp. 667–685.
- [18] S. CHAILLAT, M. BONNET, AND J.-F. SEMBLAT, *A multi-level fast multipole BEM for 3-D elastodynamics in the frequency domain*, *Computer Methods in Applied Mechanics and Engineering*, 197 (2008), pp. 4233 – 4249.
- [19] S. CHAILLAT, L. DESIDERIO, AND P. CIARLET, *Theory and implementation of \mathcal{H} -matrix based iterative and direct solvers for Helmholtz and elastodynamic oscillatory kernels*, *Journal of Computational Physics*, 351 (2017), pp. 165–186.
- [20] S. CHAILLAT, S. P. GROTH, AND A. LOSEILLE, *Metric-based anisotropic mesh adaptation for 3D acoustic boundary element methods*, *Journal of Computational Physics*, 372 (2018), pp. 473 – 499.
- [21] S. CHAILLAT, J.-F. SEMBLAT, AND M. BONNET, *A preconditioned 3-D multi-region fast multipole solver for seismic wave propagation in complex geometries*, *Communications in Computational Physics*, 11 (2012), pp. 594–609.
- [22] R. CLAYTON AND B. ENGQUIST, *Absorbing boundary conditions for acoustic and elastic wave equations*, *Bulletin of the seismological society of America*, 67 (1977), pp. 1529–1540.
- [23] G. COHEN AND S. PERNET, *Finite element and discontinuous Galerkin methods for transient wave equations*, Springer, 2017.
- [24] D. COLTON AND R. KRESS, *Integral Equation Methods in Scattering Theory*, Wiley, 1983.
- [25] P. COULIER, S. FRANÇOIS, G. LOMBAERT, AND G. DEGRANDE, *Application of hierarchical matrices to boundary element methods for elastodynamics based on Green’s functions for a horizontally layered halfspace*, *Engineering Analysis with Boundary Elements*, 37 (2013), pp. 1745–1758.
- [26] 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.
- [27] E. DARVE, *The fast multipole method: Numerical implementation*, *Journal of Computational Physics*, 160 (2000), pp. 195 – 240.
- [28] C. ERATH, S. FUNKEN, P. GOLDENITS, AND D. PRAETORIUS, *Simple error estimators for the galerkin BEM for some hypersingular integral equation in 2D*, *Applicable Analysis*, 92 (2013), pp. 1194–1216.
- [29] M. FEISCHL, M. KARKULIK, J. MELENK, AND D. PRAETORIUS, *Quasi-optimal convergence rate for an adaptive boundary element method*, *SIAM Journal on Numerical Analysis*, 51 (2013), pp. 1327–1348.
- [30] M. GANESH AND S. HAWKINS, *A fast high order algorithm for multiple scattering from large sound-hard three dimensional configurations*, *Journal of Computational and Applied Mathematics*, (doi:10.1016/j.cam.2018.10.053).
- [31] M. GANESH AND S. C. HAWKINS, *A spectrally accurate algorithm for electromagnetic scattering in three dimensions*, *Numerical Algorithms*, 43 (2006), pp. 25–60.

- [32] M. GANESH AND C. MORGENSTERN, *High-order FEM domain decomposition models for high-frequency wave propagation in heterogeneous media*, Computers & Mathematics with Applications, 75 (2018), pp. 1961–1972.
- [33] L. GAUL, M. KÖGL, F. MOSER, AND M. SCHANZ, *Boundary Element Methods for the Dynamic Analysis of Elastic, Viscoelastic, and Piezoelectric Solids*, Encyclopedia of Computational Mechanics Second Edition, (2017), pp. 1–21.
- [34] O. GORYNINA, A. LOZINSKI, AND M. PICASSO, *An easily computable error estimator in space and time for the wave equation*, arXiv preprint arXiv:1710.08410, (2017).
- [35] I. GRAHAM AND W. MCLEAN, *Anisotropic mesh refinement: the conditioning of Galerkin boundary element matrices and simple preconditioners*, SIAM journal on numerical analysis, 44 (2006), pp. 1487–1513.
- [36] R. GRAVES, *Simulating seismic wave propagation in 3D elastic media using staggered-grid finite differences*, Bulletin of the Seismological Society of America, 86 (1996), pp. 1091–1106.
- [37] L. GREENGARD, J. HUANG, V. ROKHLIN, AND S. WANDZURA, *Accelerating fast multipole methods for the helmholtz equation at low frequencies*, IEEE Computational Science and Engineering, 5 (1998), pp. 32–38.
- [38] W. HACKBUSCH, *A Sparse Matrix Arithmetic Based on H-Matrices. Part I: Introduction to H-Matrices.*, Computing, 62 (1999), pp. 89–108.
- [39] A. HAIDER AND M. SCHANZ, *Adaptive Cross Approximation for Bem in Elasticity*, Journal of Theoretical and Computational Acoustics, (doi:10.1142/S2591728518500603).
- [40] I. HARARI AND T. HUGHES, *A cost comparison of boundary element and finite element methods for problems of time-harmonic acoustics*, Computer Methods in Applied Mechanics and Engineering, 97 (1992), pp. 77–102.
- [41] J. JIN, *The finite element method in electromagnetics*, John Wiley & Sons, 2015.
- [42] P. JUHL, *A note on the convergence of the direct collocation boundary element method*, Journal of Sound and Vibration, 212 (1998), pp. 703–719.
- [43] T. KHAJAH, X. ANTOINE, AND S. BORDAS, *Isogeometric finite element analysis of time-harmonic exterior acoustic scattering problems*, arXiv preprint arXiv:1610.01694, (2016).
- [44] E. KITA AND N. KAMIYA, *Error estimation and adaptive mesh refinement in boundary element method, an overview*, Engineering Analysis with Boundary Elements, 25 (2001), pp. 479–495.
- [45] D. KOMATITSCH AND J. TROMP, *Introduction to the spectral element method for three-dimensional seismic wave propagation*, Geophysical journal international, 139 (1999), pp. 806–822.
- [46] R. KRESS, *Boundary integral equations in time-harmonic acoustic scattering*, Mathematical and Computer Modelling, 15 (1991), pp. 229–243.
- [47] D. LEVADOUX, F. MILLOT, AND S. PERNET, *A well-conditioned boundary integral equation for transmission problems of electromagnetism*, Journal of Integral Equations and Applications, 27 (2015), pp. 431–454.
- [48] Y. LIU, *Fast multipole boundary element method: theory and applications in engineering*, Cambridge university press, 2009.
- [49] Y. LIU, S. MUKHERJEE, N. NISHIMURA, M. SCHANZ, W. YE, A. SUTRADHAR, E. PAN, N. DUMONT, A. FRANGI, AND A. SAEZ, *Recent advances and emerging applications of the boundary element method*, Applied Mechanics Reviews, 64 (2011), p. 030802.

- [50] A. LOSEILLE, *Unstructured Mesh Generation and Adaptation*, in Handbook of Numerical Methods for Hyperbolic Problems - Applied and Modern Issues, R. Abgrall and C.-W. Shu, eds., Elsevier, 2017, pp. 263 – 302.
- [51] A. LOSEILLE AND F. ALAUZET, *Continuous mesh framework part I: well-posed continuous interpolation error*, SIAM Journal on Numerical Analysis, 49 (2011), pp. 38–60.
- [52] ———, *Continuous mesh framework part II: validations and applications*, SIAM Journal on Numerical Analysis, 49 (2011), pp. 61–86.
- [53] V. MATTESI, M. DARBAS, AND C. GEUZAINÉ, *A high-order absorbing boundary condition for 2D time-harmonic elastodynamic scattering problems*, Computers & Mathematics with Applications, (doi:10.1016/j.camwa.2018.05.014).
- [54] M. MESSNER, M. SCHANZ, AND E. DARVE, *Fast directional multilevel summation for oscillatory kernels based on Chebyshev interpolation*, Journal of Computational Physics, 231 (2012), pp. 1175–1196.
- [55] A. MODAVE, C. GEUZAINÉ, AND X. ANTOINE, *Corner treatment for high-order local absorbing boundary conditions in high-frequency acoustic scattering*. Under review, Nov. 2018.
- [56] K. NIINO AND N. NISHIMURA, *Preconditioning based on Calderon’s formulae for periodic fast multipole methods for Helmholtz’ equation*, Journal of Computational Physics, 231 (2012), pp. 66–81.
- [57] C. PÉREZ-ARANCIBIA AND O. BRUNO, *High-order integral equation methods for problems of scattering by bumps and cavities on half-planes*, JOSA A, 31 (2014), pp. 1738–1746.
- [58] M. PICASSO, F. ALAUZET, H. BOROUCHAKI, AND P. GEORGE, *A numerical study of some Hessian recovery techniques on isotropic and anisotropic meshes*, SIAM Journal on Scientific Computing, 33 (2011), pp. 1058–1076.
- [59] Y. SAAD, *A flexible inner-outer preconditioned gmres algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469.
- [60] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd ed., 2003.
- [61] E. P. STEPHAN, *Boundary integral equations for screen problems in \mathbb{R}^3* , Integral Equations and Operator Theory, 10 (1987), pp. 236–257.
- [62] T. TAKAHASHI, *A wideband fast multipole accelerated boundary integral equation method for time-harmonic elastodynamics in two dimensions*, International Journal for Numerical Methods in Engineering, 91 (2012), pp. 531–551.
- [63] T. TAKAHASHI, P. COULIER, AND 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), pp. 406–428.
- [64] M. VALLET, C. MANOLE, J. DOMPIERRE, S. DUFOUR, AND F. GUIBAULT, *Numerical comparison of some Hessian recovery techniques*, International Journal for Numerical Methods in Engineering, 72 (2007), pp. 987–1007.