

On the efficiency of nested GMRES preconditioners for 3D acoustic and elastodynamic \mathcal{H} -matrix accelerated Boundary Element Methods

Félix Kpadonou^a, Stéphanie Chaillat^{a,*}, Patrick Ciarlet Jr^a

^aLaboratoire POEMS (CNRS-INRIA-ENSTA Paris), ENSTA Paris, Institut Polytechnique de Paris, 828 Bd des Maréchaux, 91120 Palaiseau, FRANCE

Abstract

This article is concerned with the derivation of fast Boundary Element Methods for 3D acoustic and elastodynamic problems. In particular, we are interested in the acceleration of Hierarchical matrix (\mathcal{H} -matrix) based iterative solvers. While \mathcal{H} -matrix representations allow to reduce the storage requirements and the cost of a matrix-vector product, the number of iterations for an iterative solver, as the frequency or the problem size increases, remains an issue.

We consider an inner-outer preconditioning strategy, i.e., the preconditioner is applied through an iterative solver at the inner level. The preconditioner is defined as a \mathcal{H} -matrix representation of the system matrix with a given accuracy. We investigate the influence of various parameters of the preconditioner, i.e., the \mathcal{H} -matrix accuracy, the GMRES threshold and the maximum number of iterations of the inner solver. Different numerical results are presented to compare the efficiency of the preconditioner with respect to the unpreconditioned reference system. Finally, we propose a way to define the optimal setting for this preconditioner.

Keywords: Boundary Element Methods, Hierarchical Matrices, Wave propagation problems, Nested GMRES, Preconditionner

1. Introduction

The understanding of acoustic and elastodynamic wave propagation is important for a large range of real-life phenomena. They are involved, e.g, in the modelling and design of noise barriers for acoustic wave propagation, in nondestructive testing of material in nuclear area, in soil-structure interactions and site effects in seismic risk engineering to understand the elastic wave propagation in complex media.

Acoustic and elastodynamic wave propagation problems can be modelled by Boundary Integral Equations (BIEs) [1, 2]. This formulation is well-adapted to deal with unbounded domain problems, since the radiation conditions at infinity are exactly taken into account in the formulation with the Green's functions. The integral equations are commonly solved numerically using the Boundary Element Method (BEM) [3]. The main advantage is that only the domain boundary is meshed. Although this yields, at the discrete level, to a problem with a reduced size, the resulting system matrix is fully populated. Given N_{DOF} , the number of degrees of freedom (DOFs) on the boundary of the domain, the storage and matrix-vector product with the standard BEM are both of the order of $O(N_{\text{DOF}}^2)$. A direct solution, e.g, via a LU factorization is of the order of $O(N_{\text{DOF}}^3)$. For an iterative solver, the global solution complexity is $O(n_{\text{iter}} N_{\text{DOF}}^2)$; n_{iter} being the number of iterations. Hence, iterative solvers are more interesting than direct ones provided *a priori* that $n_{\text{iter}} \ll N_{\text{DOF}}$.

Whatever solver is used, due to computational and storage complexities, the BEM in its standard form is not usable in practice for problems with a large number of DOFs. Recently, BEM solvers have been speed

*Corresponding author

Email addresses: felix.kpadonou@ensta-paris.fr (Félix Kpadonou), stephanie.chaillat@ensta-paris.fr (Stéphanie Chaillat), patrick.ciarlet@ensta-paris.fr (Patrick Ciarlet Jr)

up with acceleration techniques yielding to fast BEMs. For smooth geometries, i.e., that are globally parameterised by spherical coordinates, a very efficient approach is to use high order spectral algorithms [4, 5]. For more general geometries, one well-known fast BEM is the Fast Multipole accelerated BEM (FM-BEM) [6, 7, 8]. The Fast Multipole Method (FMM) [9, 10] allows to compute efficiently the application of the integral operator to a given field. This method is exclusively designed for iterative solvers since it speeds up the matrix-vector product computation. Several versions of the FMM exist for Helmholtz type equations to reduce the complexity from $O(N_{\text{DOF}}^2)$ to $O(N_{\text{DOF}} \log(N_{\text{DOF}}))$ for the multi-level version (see, e.g., [11]). The FMM reduces the memory requirements by not assembling the system matrix. Since only the near contributions of the system matrix are stored, the bottleneck is the difficulty to define an efficient preconditioner for the iterative solver used with the FM-BEM. Nevertheless, several applications relative to electromagnetic or elastodynamic FM-BEMs use an incomplete LU factorization [12, 13], SParse Approximative Inverse [14, 15], multi-grid methods [16] as a preconditioning strategy. However their efficiency is limited and they do not lead to a drastic reduction of the number of iterations. Indeed, they may not contain enough information on the underlying physics.

Another approach to speed-up the BEM, that we consider, involves a hierarchical representation of the system matrix (\mathcal{H} -matrix) [17, 18] and will be referred as \mathcal{H} -BEM. The approach originally introduced in [19, 20] relies on a hierarchical partitioning of the system matrix. Through this partitioning, some blocks known *a priori* to be low-rank, thanks to an admissibility condition, are approximated using compression techniques such as the Adaptive Cross Approximation (ACA) [21, 22]. If \mathcal{H} -matrix representations allow to reduce the storage requirements and the cost of a matrix-vector product, the number of iterations for an iterative solver, as the frequency or the problem size increases, remains an issue. The advantage of this approach is however that the system matrix is available and one is not restricted in the exploration and definition of an efficient preconditioner for the iterative solver.

From a more general point of view, various algebraic preconditioning strategies which are not specifically related to the representation format of the system matrix have been proposed in the literature. To the authors best knowledge, there exist, to date, no satisfactory approach for all geometries in the context of fast BEMs for mid to high frequency oscillatory problems, i.e. when the mesh is adapted to the frequency of the problem. Algebraic preconditioners can mainly be classified into two kinds: *implicit* or *explicit* ones. For explicit preconditioners, the inverse of the preconditioner is explicitly computed and directly applied, while for implicit ones the application of the preconditioner requires the (iterative) solution of a linear system. In most existing algebraic approaches, a sparse approximation of the BEM matrix or of its inverse is proposed in order to reduce the computational cost (see [23] for a review). The sparsity pattern can be set using different strategies or heuristics [23, 26]. E.g., in [27, 24, 28] the Helmholtz BEM operator is split into a compact and a bounded operator. The sparse bounded operator is then used to precondition the system. In [25], an algebraic multi-grid method is used to construct a preconditioner for boundary element matrices arising from the Galerkin discretization of the single layer potential and the hypersingular boundary integral operators of Laplace problems. In the context of \mathcal{H} -matrices, most of existing approaches use the approximate inverse of the \mathcal{H} -matrix to compute the preconditioner [29, 30, 31]. But these approaches induce an important extra computational cost and implementation effort compared to the initial unpreconditioned system.

We are interested in this paper with the proposition of an efficient *implicit* preconditioner for iterative solvers (GMRES in this work [32]) in the context of fast BEMs (with a collocation discretization). The inverse of the preconditioner is computed with an iterative solver and thus leads to an inner-outer preconditioned solver which in practice can be solved through a nested-GMRES method. Contrary to the current trend in this field [33, 34, 35, 36, 37, 38], we do not resort to Calderón preconditioning. As a result, we do not have to compute an additional integral operator. We only make use of the information at hand in the system to be solved. This important feature makes the proposed approach very appealing especially in the mechanical engineering community.

The article is organized as follows. Section 2 is more introductory and recalls the different ingredients of BEM solvers for three-dimensional oscillatory kernels. The integral representation formulas are recalled, the discretization through collocation method is briefly sketched. The hierarchical representation for the BEM matrix is considered in Section 3. In Section 4, we present the iterative solver which involves an inner-outer solver and specify the preconditioner. It is defined as a rough approximation of the BEM matrix. The efficiency of the present framework is controlled thanks to the compression accuracy. The advantage with

the proposed approach is that no additional memory nor computational time [other than the ones already needed to build the fast BEM](#) are required. In Section 5 several numerical tests are performed to discuss the efficiency. We compare the preconditioned inner-outer (nested-GMRES) solver to the unpreconditioned one. Section 6 gives some conclusions and perspectives for future works.

2. Boundary Element Method for Acoustic and Elastodynamic wave propagation

We consider the propagation of time-harmonic acoustic and elastic waves in three-dimensional isotropic and homogeneous domains. We adopt the following notations: matrices are denoted in blackboard characters and vector quantities in boldface. Then, we denote by u and \mathbf{u} respectively the velocity and displacement fields of the acoustic and elastodynamic problems. $\Omega^- \in \mathbb{R}^3$ is the bounded domain representing the obstacle, with a closed Lipschitz boundary $\Gamma = \partial\overline{\Omega^-}$, Ω^+ is the exterior domain $\mathbb{R}^3 \setminus \overline{\Omega^-}$. \mathbf{n} represents the outward unit normal vector field on Γ . The acoustic and elastodynamic equations are respectively given by

$$\Delta u + \kappa^2 u = 0 \quad (1)$$

and

$$\operatorname{div}(\boldsymbol{\sigma}(\mathbf{u})) + \rho \omega^2 \mathbf{u} = 0. \quad (2)$$

The stress and strain tensors are respectively given by $\boldsymbol{\sigma}(\mathbf{u}) = \lambda(\operatorname{div} \mathbf{u})\mathbb{I}_3 + 2\mu\boldsymbol{\varepsilon}(\mathbf{u})$ and $\boldsymbol{\varepsilon}(\mathbf{u}) = \frac{1}{2}([\nabla \mathbf{u}] + [\nabla \mathbf{u}]^\top)$; where \mathbb{I}_3 is the 3-by-3 identity matrix and $[\nabla \mathbf{u}]$ is the 3-by-3 matrix whose β -th column is the gradient of the β -th component of \mathbf{u} , μ and λ being the Lamé parameters. κ represents the wavenumber of the acoustic wave. ω and ρ are respectively the circular frequency and the density for the elastic case. We denote by κ_p and κ_s the P and S wavenumbers defined as $\kappa_p^2 = \rho\omega^2(\lambda + 2\mu)^{-1}$ and $\kappa_s^2 = \rho\omega^2\mu^{-1}$. The Green's tensors for the case of an acoustic and elastic full-space are respectively given by

$$G(\mathbf{x}, \mathbf{y}; \kappa) = \frac{e^{i\kappa|\mathbf{x}-\mathbf{y}|}}{4\pi|\mathbf{x}-\mathbf{y}|} \quad (3)$$

and

$$\mathbf{G}(\mathbf{x}, \mathbf{y}; \omega) = \frac{1}{\rho\omega^2} (\mathbf{curl} \operatorname{curl}_{\mathbf{x}}[G(\mathbf{x}, \mathbf{y}; \kappa_s)\mathbb{I}_3] - \nabla_{\mathbf{x}} \operatorname{div}_{\mathbf{x}}[G(\mathbf{x}, \mathbf{y}; \kappa_p)\mathbb{I}_3]). \quad (4)$$

The index \mathbf{x} means that the differentiation is carried out with respect to \mathbf{x} and $\operatorname{div}_{\mathbf{x}} \mathbb{B}$ corresponds to the application of the divergence along each row of \mathbb{B} . We introduce the traction operator \mathbf{T} and the acoustic pressure operator denoted by \mathbb{T} for sake of genericity

$$\mathbf{T} = 2\mu \frac{\partial}{\partial \mathbf{n}} + \lambda \mathbf{n} \operatorname{div} + \mu \mathbf{n} \times \mathbf{curl} \quad \text{and} \quad \mathbb{T} = \frac{\partial}{\partial \mathbf{n}}. \quad (5)$$

The traction tensor (resp. the normal derivative) of the Green's tensor, obtained by applying the traction operators, are defined as follows

$$\mathbf{T}(\mathbf{x}, \mathbf{y}, \omega) = \mathbf{T}_{\mathbf{y}} \mathbf{G}(\mathbf{x}, \mathbf{y}; \omega) \quad \text{and} \quad T(\mathbf{x}, \mathbf{y}, \kappa) = \mathbb{T}_{\mathbf{y}} G(\mathbf{x}, \mathbf{y}; \kappa).$$

For the elastic case the operator \mathbf{T} applies to each column.

Boundary Integral Representations and Equations. We consider an exterior scattering problem by an obstacle Ω^- . We denote by $H^s(\Omega_-)$, $H^s(\Omega_+)$ and $H^s(\Gamma)$ the standard complex valued Hilbertian Sobolev spaces of order $s \in \mathbb{R}$ ($|s| \leq 1$ for $H^s(\Gamma)$), with the convention $H^0 = L^2$. The complex vector value Hermitian spaces are denoted in boldface, then $\mathbf{H}^s = [H^s]^3$.

We introduce the operators $\Delta^*u := \Delta u + \kappa^2 u$ and $\mathbf{\Delta}^*\mathbf{u} := \operatorname{div} \sigma(\mathbf{u}) + \rho w^2 \mathbf{u}$ and the energy operator spaces for acoustics

$$H_-^1(\Delta^*) := \{u \in H^1(\Omega_-) : \Delta^*u \in L^2(\Omega_-)\}$$

and

$$H_+^1(\Delta^*) := \{u \in H_{loc}^1(\overline{\Omega_+}) : \Delta^*u \in L_{loc}^2(\overline{\Omega_+})\};$$

where *loc* means local in the case of the exterior domain. u belongs to $H_{loc}^1(\Omega)$ if, and only if, $u|_K$ belongs to $H^1(\Omega)$ for every compact subset K of Ω , where given a subset S of \mathbb{R}^3 , 1_S denotes the indicator function of S . For elastodynamics, $\mathbf{H}_-^1(\mathbf{\Delta}^*)$ and $\mathbf{H}_+^1(\mathbf{\Delta}^*)$ are defined analogously.

The acoustic pressure and elastic traction traces are defined by $t|_\Gamma := \mathbf{T}u$ and $\mathbf{t}_\Gamma = \mathbf{T}\mathbf{u}$; the operators \mathbf{T} and \mathbf{T} being defined in (5). For any solution $u \in H_+^1(\Delta^*)$ of (1), the classical integral representation obtained using the Green's function and the normal trace is

$$u(\mathbf{x}) = \mathcal{D}u|_\Gamma(\mathbf{x}) - \mathcal{S}t|_\Gamma(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^3 \setminus \Gamma,$$

where $u|_\Gamma \in H^{\frac{1}{2}}(\Gamma)$ and $t|_\Gamma \in H^{-\frac{1}{2}}(\Gamma)$. Given $\varphi \in H^{-\frac{1}{2}}(\Gamma)$ and $\psi \in H^{\frac{1}{2}}(\Gamma)$, the single- and double-layer potentials are respectively defined by

$$\mathcal{S}\varphi(\mathbf{x}) = \int_\Gamma G(\mathbf{x}, \mathbf{y}; \kappa) \varphi(\mathbf{y}) ds(\mathbf{y}) \quad \text{and} \quad \mathcal{D}\psi(\mathbf{x}) = \int_\Gamma [T_{\mathbf{y}}G(\mathbf{x}, \mathbf{y}; \kappa)]^\top \psi(\mathbf{y}) ds(\mathbf{y}) \quad \text{for } \mathbf{x} \in \mathbb{R}^3 \setminus \Gamma.$$

For the analogous representation formula in elastodynamics, the single- and double-layer potentials are defined similarly: Given $\varphi \in \mathbf{H}^{-\frac{1}{2}}(\Gamma)$ and $\psi \in \mathbf{H}^{\frac{1}{2}}(\Gamma)$

$$\mathcal{S}\varphi(\mathbf{x}) = \int_\Gamma \mathbf{G}(\mathbf{x}, \mathbf{y}; \omega) \varphi(\mathbf{y}) ds(\mathbf{y}) \quad \text{and} \quad \mathcal{D}\psi(\mathbf{x}) = \int_\Gamma [\mathbf{T}_{\mathbf{y}}\mathbf{G}(\mathbf{x}, \mathbf{y}; \omega)]^\top \psi(\mathbf{y}) ds(\mathbf{y}).$$

The corresponding integral representation for a solution $\mathbf{u} \in \mathbf{H}_+^1(\mathbf{\Delta}^*)$ to the equation (2) satisfying the Kupradze radiation condition [1, 2] is

$$\mathbf{u}(\mathbf{x}) = \mathcal{D}\mathbf{u}|_\Gamma(\mathbf{x}) - \mathcal{S}\mathbf{t}|_\Gamma(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^3 \setminus \Gamma.$$

The single-layer potential \mathcal{S} (resp. \mathcal{S}) is continuous from $H^{-\frac{1}{2}}(\Gamma)$ to $H_-^1(\Delta^*) \cup H_+^1(\Delta^*)$ (resp. $\mathbf{H}^{-\frac{1}{2}}(\Gamma)$ to $\mathbf{H}_-^1(\mathbf{\Delta}^*) \cup \mathbf{H}_+^1(\mathbf{\Delta}^*)$). On the other hand, the double-layer potential \mathcal{D} (resp. \mathcal{D}) is continuous from $H^{\frac{1}{2}}(\Gamma)$ to $H_-^1(\Delta^*) \cup H_+^1(\Delta^*)$ (resp. $\mathbf{H}^{\frac{1}{2}}(\Gamma)$ to $\mathbf{H}_-^1(\mathbf{\Delta}^*) \cup \mathbf{H}_+^1(\mathbf{\Delta}^*)$). For any fields $\varphi \in H^{-\frac{1}{2}}(\Gamma)$ and $\psi \in H^{\frac{1}{2}}(\Gamma)$ (resp. $\varphi \in \mathbf{H}^{-\frac{1}{2}}(\Gamma)$ and $\psi \in \mathbf{H}^{\frac{1}{2}}(\Gamma)$), the potentials $\mathcal{S}\varphi$ and $\mathcal{D}\psi$ (resp. $\mathcal{S}\varphi$ and $\mathcal{D}\psi$) solve the Helmholtz (resp. elastodynamic) equation in Ω^+ and Ω^- . The exterior and interior Dirichlet γ_0^\pm (resp. γ_0^\pm) and traction γ_1^\pm (resp. γ_1^\pm) traces of \mathcal{S} and \mathcal{D} (resp. \mathcal{S} and \mathcal{D}) are given by

$$\gamma_0^\pm \mathcal{S} = \mathcal{S}, \quad \gamma_1^\pm \mathcal{S} = (\mp \frac{1}{2} + D'), \quad \gamma_0^\pm \mathcal{D} = (\pm \frac{1}{2} + D) \quad (6)$$

$$\gamma_0^\pm \mathcal{S} = \mathcal{S}, \quad \gamma_1^\pm \mathcal{S} = (\mp \frac{1}{2} \mathbb{I}_3 + D'), \quad \gamma_0^\pm \mathcal{D} = (\pm \frac{1}{2} \mathbb{I}_3 + D) \quad (7)$$

where the operators \mathcal{S} (resp. \mathcal{S}) and \mathcal{D} (resp. \mathcal{D}) are continuous from $H^{-\frac{1}{2}}(\Gamma)$ to $H^{\frac{1}{2}}(\Gamma)$ (resp. $\mathbf{H}^{-\frac{1}{2}}(\Gamma)$ to $\mathbf{H}^{\frac{1}{2}}(\Gamma)$) and from $H^{\frac{1}{2}}(\Gamma)$ to $H^{-\frac{1}{2}}(\Gamma)$ (resp. $\mathbf{H}^{\frac{1}{2}}(\Gamma)$ to $\mathbf{H}^{-\frac{1}{2}}(\Gamma)$) and are defined as follows

$$\mathcal{S}\varphi(\mathbf{x}) = \int_\Gamma G(\mathbf{x}, \mathbf{y}; \kappa) \varphi(\mathbf{y}) ds(\mathbf{y}) \quad \text{and} \quad \mathcal{D}\psi(\mathbf{x}) = \int_\Gamma [T_{\mathbf{y}}G(\mathbf{x}, \mathbf{y}; \kappa)]^\top \psi(\mathbf{y}) ds(\mathbf{y}) \quad \text{for } \mathbf{x} \in \Gamma;$$

\mathbf{S} and \mathbf{D} are defined analogously. We write integrals, that stand for duality products if the data is insufficiently smooth.

The scattering problems are then formulated as Boundary Integral Equations (BIEs): Given an incident wave u_{inc} which is assumed to solve the Helmholtz equation in the absence of obstacle, find u solution to (1) in Ω^+ which satisfies the Dirichlet boundary condition on Γ

$$u|_{\Gamma} + u_{inc} = 0.$$

The acoustic scattering problem is: Find $t|_{\Gamma} \in H^{-\frac{1}{2}}(\Gamma)$ such that

$$S(t|_{\Gamma} + t|_{\Gamma}^{inc})(\mathbf{x}) = u|_{\Gamma}^{inc}(\mathbf{x}), \quad \mathbf{x} \in \Gamma. \quad (8)$$

Similarly, for elastodynamic problems, it is: Find $\mathbf{t}|_{\Gamma} \in \mathbf{H}^{-\frac{1}{2}}(\Gamma)$ such that

$$\mathbf{S}(\mathbf{t}|_{\Gamma} + \mathbf{t}|_{\Gamma}^{inc})(\mathbf{x}) = \mathbf{u}|_{\Gamma}^{inc}(\mathbf{x}), \quad \mathbf{x} \in \Gamma. \quad (9)$$

Discretization of the BIE. At the discrete level, one deals with a linear system resulting from the discretization of the BIE. Several discretization techniques can be used in practice, typically the collocation and Galerkin methods. We consider the collocation technique which requires to satisfy the BIE at some arbitrarily chosen (collocation) points [1]. The domain boundary Γ is approximated by N_E triangular elements E_l . h is the size of mesh $\Gamma_h := \bigcup_{l=1:N_E} E_l$. We introduce the sets of points $X = (\mathbf{x}_i)_{i=1:N_c}$ and $Y = (\mathbf{y}_j)_{j=1:N_{DOF}}$, respectively the collocation and interpolation points that all belong to Γ_h . For the acoustic scattering problems, it reads

$$S_h(t|_{\Gamma} + t|_{\Gamma}^{inc})(\mathbf{x}_i) = u|_{\Gamma}^{inc}(\mathbf{x}_i), \quad \forall i \in \{1, \dots, N_c\};$$

S_h is the discretization of the single layer potential for Γ_h . We denote by $\mathbb{A} \in \mathbb{C}^{N_c \times N_{DOF}}$ and $\mathbf{b} \in \mathbb{C}^{N_c}$ respectively the system matrix (also called BEM matrix) and the right hand side associated to the incident wave. We are interested in the solution through an iterative solver of the linear system

$$\mathbb{A}\mathbf{x} = \mathbf{b}; \quad (10)$$

$\mathbf{x} \in \mathbb{C}^{N_{DOF}}$ being the vector of unknown DOFs. The system matrix \mathbb{A} is commonly non-hermitian and fully-populated, hence a prior and crucial point is its appropriate storage-friendly representation.

3. Hierarchical data-sparse representation

To speed-up the BEM, we use a data-sparse representation of the system matrix. It is denoted by $\mathbb{A}_{\eta,\varepsilon}$; the parameter η defines the “data-sparsity pattern” associated to a given partitioning of \mathbb{A} . The parameter $\varepsilon > 0$ is a given accuracy of the data-sparse representation such that, for a given norm $\|\cdot\|$

$$\|\mathbb{A} - \mathbb{A}_{\eta,\varepsilon}\| \leq \varepsilon \|\mathbb{A}\|. \quad (11)$$

Thus, instead of (10), we consider the problem

$$\mathbb{A}_{\eta,\varepsilon}\mathbf{x}_{\mathcal{H}} = \mathbf{b}. \quad (12)$$

It is worth noting that the fidelity of $\mathbf{x}_{\mathcal{H}}$ with respect to \mathbf{x} , solution of (10), is guaranteed [18]. Indeed, it is shown that, the error $\|\mathbb{A}\mathbf{x}_{\mathcal{H}} - \mathbf{b}\|$ is controlled by the sum of the solution error $\|\mathbb{A}_{\eta,\varepsilon}\mathbf{x}_{\mathcal{H}} - \mathbf{b}\|$ and an additional error taking into account the influence of the accuracy parameter ε (the quality of the approximation of $\mathbb{A}_{\eta,\varepsilon}$).

Hierarchical matrices or \mathcal{H} -matrices have been introduced by Hackbusch [19] 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 (i) to partition the matrix into blocks and (ii) to perform low-rank approximations of the blocks of the matrix which are known *a priori* (by using an admissibility condition)

to be accurately approximated by low-rank decompositions. Using low-rank representations, the memory requirements and costs of a matrix-vector product are reduced. We refer to [18] for a deep insight on \mathcal{H} -matrix representation of oscillatory kernels and [39] for general purpose on \mathcal{H} -matrix representation, while only the key points are recalled below.

Low-rank Admissibility (data-sparsity pattern). The essential idea of the \mathcal{H} -matrix representation of a given matrix resides in its hierarchical partitioning in order to exhibit some blocks which are accurately approximated by low-rank matrices. Let $\sigma \subset \{1, \dots, N_c\}$ and $\tau \subset \{1, \dots, N_{\text{DOF}}\}$ denote two sets of indices corresponding to the clusters of nodes $X_\sigma = (\mathbf{x}_i)_{i \in \sigma} \subset X$ and $Y_\tau = (\mathbf{y}_j)_{j \in \tau} \subset Y$. $\mathbb{A}_{\sigma \times \tau}$ is the block of \mathbb{A} restricted to the row and column indices corresponding to the interaction between the clusters of nodes X_σ and Y_τ . When $X = Y$, singularities mainly occur (for $\tau = \sigma$) such that the diagonal is composed of full rank blocks but the kernel function is smooth everywhere else. For the Laplace kernel this is transcribed by the asymptotically smooth property [40]. Thus, at the discrete level, some blocks $\mathbb{A}_{\sigma \times \tau}$, are known *a priori* to be low-rank using the admissibility condition. The condition depends on geometric characteristics such as the diameters of the clusters of points X_σ , Y_τ and the distance between them. The condition for admissible blocks for the Laplace (static) case is the η -admissibility condition and reads

$$\min(\text{diam}(X_\sigma), \text{diam}(Y_\tau)) \leq \eta \text{dist}(X_\sigma, Y_\tau); \quad (13)$$

where dist and diam respectively denote the Euclidean distance between two clusters and the diameter of a cluster.

For wave propagation problems, this condition should depend also on the wavenumber. However, in practice, the η -admissibility condition has been shown to be viable towards high frequency regime. In fact, it has been shown that there exists a pre-asymptotic regime [18] where the maximum numerical rank among all admissible blocks increases linearly.

Hierarchical representation. The key ingredient of hierarchical matrices is the recursive block subdivision of the geometry on which the matrix blockwise partitioning is based. The geometry nodes and corresponding indices in the matrix are permuted to reflect the physical distances. Indeed, consecutive indices of row and column should correspond respectively to collocation and interpolation points that interact at close range, i.e., belonging to the same cluster.

Clustering of the unknowns. For the sake of clarity, in this work \mathbb{A} is defined by the same set of indices $I = \{1, \dots, N\}$ for rows and columns. This corresponds to the typical case where the collocation nodes are chosen as the interpolation points. A binary tree \mathcal{T}_I is used to drive the clustering. 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, see Figure 1. A stopping criterion for the recursive subdivision of a node of indices $\sigma \subset I$ bears on the minimum number of indices in a block, i.e., it is defined such that the size of the block matrix satisfies $|\sigma| \leq n_{\text{leaf}}$, where n_{leaf} is given.

Subdivision of the matrix. Once 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 $\mathbb{A}_{\sigma \times \tau}$ of \mathbb{A} (see Figure 2).

In practice, 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 defined. Starting from the initial matrix \mathbb{A} , i.e., $\sigma = I$ and $\tau = I$, a given block matrix $\mathbb{A}_{\sigma \times \tau}$ is recursively subdivided into 2×2 subblocks $\mathbb{A}_{\sigma_\alpha \times \tau_\beta}$, $(\alpha, \beta) \in \{1, 2\}$ matrices with $\sigma = \sigma_1 \cup \sigma_2$, $\tau = \tau_1 \cup \tau_2$ until this block is either η -admissible or too small, i.e., $\min(|\sigma|, |\tau|) \leq n_{\text{leaf}}$.

The partition \mathcal{P} is thus subdivided into two subsets \mathcal{P}^{ad} and $\mathcal{P}^{\text{non-ad}}$ reflecting the possibility for a block $\tau \times \sigma$ to be either *admissible*, i.e., $\tau \times \sigma \in \mathcal{P}^{\text{ad}}$; or *non-admissible*, i.e., $\tau \times \sigma \in \mathcal{P}^{\text{non-ad}}$. $\mathcal{P} = \mathcal{P}^{\text{ad}} \cup \mathcal{P}^{\text{non-ad}}$. This hierarchical partitioning is preferable and optimal in regards to the uniform partitioning which relies only on the leaf nodes of the cluster tree \mathcal{T}_I , since admissible blocks can occur at higher levels of the block clustering tree.

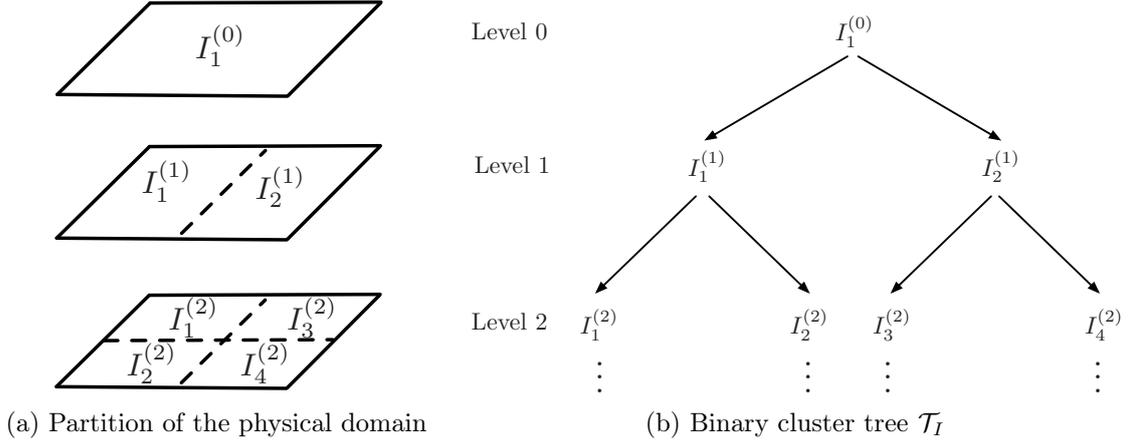


Figure 1: Illustration of the clustering of the degrees of freedom: (a) partition of the degrees of freedom in the domain and (b) corresponding binary tree.

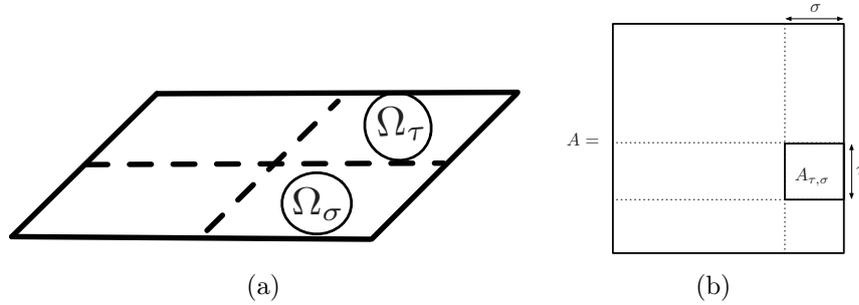


Figure 2: Illustration of the construction of the block cluster tree: (a) Clustering of the unknowns on the geometry and (b) corresponding block clustering in the matrix.

Low-rank (data-sparse) approximation. Once the *admissible* blocks are determined, an accurate rank-revealing algorithm is applied to determine the corresponding low-rank approximations. The idea with the low-rank representations, analogous to compression techniques, is to reveal and keep only the main significant information. The truncated SVD [41] gives the best low-rank approximation for unitary invariant norms. However, its computation is expensive and requires the complete storage of the matrix. The adaptive cross approximation (ACA) [17, 42, 22] offers an interesting alternative to the SVD.

The numerical rank obtained by the ACA being *sub-optimal*, in the sense that it is greater than the exact low-rank, it induces a slight additional memory consumption. Therefore, in practice a recompression of the ACA-based low-rank approximation is further made in order to obtain a more accurate numerical rank. Indeed, the ACA compression can be viewed as an intermediate step which allows to have at disposal a coarse low-rank approximation of the initial matrix which is not storage consuming, and on which the truncated SVD can be applied for a further more optimal low-rank computation. Reader may refer to [21, 43] for a deep insight on the matter and specifically to [18] for the extension of the low-rank approximation to vector-valued (elastodynamic) problem using the ACA.

4. Preconditioning a \mathcal{H} -BEM solver

Let \mathbb{M} and \mathbf{b} be respectively a generic matrix and a right hand side vector. The solution \mathbf{x} of the system $\mathbb{M}\mathbf{x} = \mathbf{b}$, through an iterative solver, is obtained by computing iteratively a sequence of vectors \mathbf{x}_k , approximating the exact solution \mathbf{x} . Starting from an initial guess \mathbf{x}_0 , the principle is to build the \mathbf{x}_k at

each iteration which minimizes the norm of the residual $\mathbf{r}_k = \mathbb{M}\mathbf{x}_k - \mathbf{b}$ over an appropriate subspace. Using a preconditioner, we are interested in solving an alternative system which is expected to have better spectral properties, i.e., better conditioning or better eigenvalues clustering in comparison to the original system. We denote by \mathbb{P} the preconditioner. For a right preconditioner, we solve the modified system

$$\mathbb{M}\mathbb{P}^{-1}\mathbb{P}\mathbf{x} = \mathbf{b} \Leftrightarrow \mathbb{M}\mathbb{P}^{-1}\mathbf{y} = \mathbf{b} \text{ with } \mathbb{P}\mathbf{x} = \mathbf{y}. \quad (14)$$

The application of the inverse of \mathbb{P} is required throughout the iterative solver. Instead of considering a ‘‘classical’’ approach, where a direct solver or explicit preconditioner is used, we consider that the preconditioning system is itself solved using an iterative solver. Hence we deal with a two-level iterative method. An outer solver that involves the original system matrix \mathbb{M} and the corresponding right hand side vector \mathbf{b} ; and an inner solver that involves the preconditioner \mathbb{P} of the outer problem, hereafter denoted \mathbb{P}_{out} . Of course, one could still assume that the inner solver is right preconditioned with an operator (that would be denoted in that case \mathbb{P}_{in}). The preconditioned system that we consider in the following thus reduces to

$$\mathbb{M}\mathbb{P}_{\text{out}}^{-1}\mathbf{y} = \mathbf{b} \text{ with } \mathbb{P}_{\text{out}}\mathbf{x} = \mathbf{y}. \quad (15)$$

With the \mathcal{H} -BEM, the complete system matrix is available conversely to the FM-BEM. Thus there are no constraints on how to define the preconditioner. In the sequel, we consider that the inner solver is not preconditioned. We denote by n_{out} the total number of outer iterations and n_{in} the cumulative, total number of iterations in the inner solver.

Theoretically, the best preconditioner is $\mathbb{P}_{\text{out}} = \mathbb{M}$. But, with this choice, the inner solver is equivalent to the outer solver. Hence the convergence is achieved after $n_{\text{out}} = 1$ outer iteration. The operations required, at each iteration, in the nested inner-outer iterative solver are the application of the system matrix \mathbb{M} and of the preconditioner \mathbb{P}_{out} to vectors. We look for a choice of \mathbb{P}_{out} such that the computational cost of $\mathbf{w} \leftarrow \mathbb{P}_{\text{out}}\mathbf{z}$ is low (with respect to $\mathbf{y} \leftarrow \mathbb{M}\mathbf{z}$). In the context of \mathcal{H} -BEM, we make the choice: $\mathbb{P}_{\text{out}} = \mathbb{A}_{\eta, \varepsilon'}$, i.e. we use a low-rank approximation of \mathbb{A} with a lower accuracy in the low-rank approximation of the preconditioner. Hence, we set $\varepsilon' \geq \varepsilon$, where ε is the accuracy of compression of the system matrix $\mathbb{M} = \mathbb{A}_{\eta, \varepsilon}$. The idea is to choose the preconditioner as a coarse approximation of \mathbb{M} (the theoretical best preconditioner) in order to achieve convergence through a minimum number of iterations while keeping the computational cost as low as possible. Indeed, for $\sigma \times \tau \in \mathcal{P}_{\text{ad}}$, the matrix-vector product cost is of the order of $O(|\sigma| + |\tau|) \times r(\varepsilon)$. Where $r(\varepsilon)$ is the numerical rank of a block matrix, and this rank decreases with the accuracy ε . The advantage with this choice of preconditioner is that there is no additional time for its computation nor additional storage requirements. It is already included in the computation of the system matrix $\mathbb{A}_{\eta, \varepsilon}$. Given an admissible block $\mathbb{A}_{\sigma \times \tau}$, $\sigma \times \tau \in \mathcal{P}^{\text{ad}}$ and $\varepsilon' \geq \varepsilon$, it is clear, from the construction of the low-rank approximation through successive rank-one matrix additions to the approximation, that $r(\varepsilon') \leq r(\varepsilon)$. Now, assume that the low-rank approximation of $\mathbb{A}_{\sigma \times \tau}$ for the accuracy ε is

$$\mathbb{B}_{\sigma \times \tau} = \mathbb{U}\mathbb{V}^*, \quad \mathbb{U} \in \mathbb{R}^{|\sigma| \times r(\varepsilon)} \text{ and } \mathbb{V} \in \mathbb{R}^{|\tau| \times r(\varepsilon)};$$

thus the low rank approximation for the accuracy ε' is

$$\mathbb{B}'_{\sigma \times \tau} = \mathbb{U}'\mathbb{V}'^*, \quad \mathbb{U}' \in \mathbb{R}^{|\sigma| \times r(\varepsilon')} \text{ and } \mathbb{V}' \in \mathbb{R}^{|\tau| \times r(\varepsilon')}.$$

Then the first $r(\varepsilon')$ columns of \mathbb{U} and \mathbb{V} form the columns of the matrices \mathbb{U}' and \mathbb{V}' .

Remark 4.1. In [30], the adopted preconditioner is a coarse (in regard to the accuracy of compression) LU factorization of the system \mathcal{H} -matrix with application to the Laplace problem with Galerkin method and to the magnetostatic problem with collocation method. Although, in practice, a low accuracy is sufficient for the computation of the LU factorization of the \mathcal{H} -matrix formatted preconditioner [30], an additional amount of storage memory is required for the storage of the lower and upper \mathcal{H} -matrices. The advantage with the proposed preconditioning strategy is that no extra memory is required in contrast to the LU factorization based one.

5. Numerical efficiency of the proposed strategy

To study the efficiency of the proposed preconditioning strategy, we consider the exterior domain scattering problem of a time harmonic acoustic plane wave in an homogeneous and isotropic medium by a sphere of radius r . We use the single-layer formulation (8). The discretization is performed with a collocation technique and a Lagrange P_1 interpolation. The parameter defining the hierarchical pattern and used for the low-rank admissibility condition is $\eta = 3$ and the minimum dimension of the block is $n_{\text{leaf}} = 100$. The accuracy of the \mathcal{H} -BEM matrix is set to $\varepsilon = 10^{-8}$. We assume that the number of points per wavelength is constant, i.e. a density $n_\lambda = 10$. The number of wavelengths along the characteristic dimension is denoted $d_\lambda = \frac{2r}{\lambda}$. The numerical tests are performed on a bi-processor Intel XEON E5-2637 machine where each processor is composed of 4 cores, with two threads per core and a RAM of 756 Gb. The solver (COFFEE¹ developed at POEMS) is implemented using a **shared** memory parallelization (OpenMP) **of the construction of the \mathcal{H} -matrix representation and of the matrix-vector product**. The different numerical tests are performed with 8 threads. **More details on the validation of this fast BEM solver are available in [18, 44].**

5.1. Preliminary tests: eigenvalue clustering

The efficiency of a preconditioner depends on the clustering of the eigenvalues of the preconditioned system, namely on the diameter of the cluster and the numerical range of the operator. Thus, we look for the influence of the accuracy ε' on the clustering of the eigenvalues of $\mathbb{A}_{\eta,\varepsilon} \mathbb{A}_{\eta,\varepsilon'}^{-1}$ and the condition number cond_2 in l_2 -norm.

The subsequent analysis is limited to a mesh with $N_{\text{DOF}} = 7680$. Indeed the computational cost of the complete eigenvalue decomposition is very prohibitive. The eigenvalues are computed using the library ARPACK [45] and the application of $\mathbb{A}_{\eta,\varepsilon'}^{-1}$ required by the routine is made using an iterative solver (a standard GMRES algorithm [46]) with a threshold equal to 10^{-6} .

Preconditioned inner-outer system. The diameters of the box surrounding the eigenvalues and the condition numbers corresponding to the different **accuracies used in the** preconditioners are reported on Table 1 (**for two frequencies and problem sizes**). We remark that the eigenvalues of the unpreconditioned system (i.e. with $\mathbb{P}_{\text{out}} = \mathbb{I}$) are less clustered and that the clustering, in regards to the values of the diameter of the box surrounding the eigenvalues improve when a preconditioner with higher accuracy is used. Finally the preconditioner with an accuracy $\varepsilon' = 10^{-1}$ appears to be useless in regard to the eigenvalues clustering without preconditioner.

Table 1: Information on the clustering of the eigenvalues and condition number cond_2 of the preconditioned system matrix for different values of ε' .

| $N_{\text{DOF}} = 7680, r\kappa = 14.32, d_\lambda = 4.6$ | | | $N_{\text{DOF}} = 6040, r\kappa = 12.71, d_\lambda = 4.05$ | |
|---|----------------------|-----------------|--|-----------------|
| ε' | box diam | cond_2 | box diam | cond_2 |
| 10^{-7} | $4.25 \cdot 10^{-6}$ | 1.00 | $4.10 \cdot 10^{-6}$ | 1.00 |
| 10^{-6} | $4.64 \cdot 10^{-6}$ | 1.00 | $4.22 \cdot 10^{-6}$ | 1.00 |
| 10^{-5} | $3.34 \cdot 10^{-5}$ | 1.00 | $3.23 \cdot 10^{-5}$ | 1.00 |
| 10^{-4} | $3.03 \cdot 10^{-4}$ | 1.00 | $3.02 \cdot 10^{-4}$ | 1.00 |
| 10^{-3} | $2.94 \cdot 10^{-3}$ | 1.00 | $2.56 \cdot 10^{-3}$ | 1.00 |
| 10^{-2} | $2.64 \cdot 10^{-2}$ | 1.03 | $2.32 \cdot 10^{-2}$ | 1.02 |
| 10^{-1} | 0.22 | 1.47 | 0.21 | 1.25 |
| \mathbb{I} | 0.18 | 44.24 | 0.20 | 33.92 |

¹<https://uma.ensta-paris.fr/soft/COFFEE/>

We have seen the influence of the parameter ε' on the outer solver. In the light of the previous test, we expect moderate numbers of outer iterations for ε' “sufficiently small”. As a consequence, the performances of the inner-outer solver will mostly depend on the behavior of the inner solver.

Conditioning of the inner solver. We report on Table 2 the values of the diameters and the condition numbers for the system matrices $\mathbb{A}_{\eta,\varepsilon'}$. For each non-dimensional wavenumber, we observe that, for the inner solver, the clustering of the eigenvalues are very similar and the values of the condition numbers are almost equal for the different accuracies. Thus, there is no a priori accuracy ε' to be favoured for the inner solver. The balance between the computational time per inner iteration and the total number of inner iterations will be essential for the performance of the inner-outer solver.

In the light of the previous tests on the eigenvalues, one can consider the inner-outer preconditioning for the accuracies ε' larger or equal to 10^{-2} . Subsequently, we will focus on the performances of the right preconditioned, inner-outer solver for different choices of parameters. These performances are compared to those obtained with the unpreconditioned reference solver.

Table 2: Information on the clustering of the eigenvalues and condition number of the unpreconditioned inner solver problem, for different values of ε' .

| $N_{\text{DOF}} = 7680, r\kappa = 14.32, d_\lambda = 4.6$ | | | $N_{\text{DOF}} = 6040, r\kappa = 12.71, d_\lambda = 4.05$ | |
|---|----------|-------------------|--|-------------------|
| ε' | box diam | cond ₂ | box diam | cond ₂ |
| 10^{-7} | 0.18 | 46.83 | 0.20 | 37.25 |
| 10^{-6} | 0.18 | 46.83 | 0.20 | 37.25 |
| 10^{-5} | 0.18 | 46.83 | 0.20 | 37.25 |
| 10^{-4} | 0.18 | 46.83 | 0.20 | 37.25 |
| 10^{-3} | 0.18 | 46.79 | 0.20 | 37.21 |
| 10^{-2} | 0.18 | 46.75 | 0.20 | 37.11 |
| 10^{-1} | 0.18 | 42.45 | 0.20 | 38.29 |

5.2. Reference unpreconditioned solver: performances

We first study the performances of the unpreconditioned reference solver. The following parameters are considered: for the outer GMRES solver, the maximum number of iterations is denoted by N_{out} and GMRES stopping criteria is denoted by $\varepsilon_{\text{tol}(\text{out})}$. The stopping criteria $\varepsilon_{\text{tol}(\text{out})}$ is set equal to 10^{-6} while the maximum number of iterations is 2000. The performances are checked on different meshes. Their corresponding numbers of degrees of freedom N_{DOF} , non-dimensional wavenumbers $r\kappa$ and numbers of wavelengths along the diameter d_λ are reported on Table 3.

Table 3: Number of DOFs, non-dimensional wavenumber and number of wavelengths along the characteristic dimension corresponding to each mesh.

| Mesh i | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|------------------|--------|--------|--------|---------|---------|---------|---------|
| N_{DOF} | 10 242 | 40 962 | 61 033 | 163 543 | 254 546 | 328 606 | 626 333 |
| $r\kappa_i$ | 16.64 | 33.27 | 40.91 | 66.61 | 83.60 | 92.79 | 127.84 |
| d_λ | 5.30 | 10.30 | 13.00 | 21.20 | 26.62 | 29.44 | 40.69 |

We denote by T_{ref} (resp. n_{ref}) the solution time (number of iterations) of the reference solver. The reference solver corresponds to the problem with the system matrix $\mathbb{A}_{\eta,\varepsilon}$ with $\varepsilon = 10^{-8}$.

The performances of the solver are reported on Table 4, where $T_{\mathcal{H}\text{-BEM}}$ is the time corresponding to the \mathcal{H} -matrix approximation of the BEM matrix while $\delta_{\mathcal{H}}$ is the compression rate, i.e., the number of entries

required for the storage of the \mathcal{H} -BEM matrix $\mathbb{A}_{\eta,\varepsilon}$ over the one of the standard fully dense BEM matrix \mathbb{A} . That is, $\delta_{\mathcal{H}} N_{\text{DOF}}^2$ is the total number of entries for the storage of the full and low-rank admissible blocks of $\mathbb{A}_{\eta,\varepsilon}$.

Table 4: Main characteristics of the reference solver.

| $r\kappa$ | n_{ref} | $T_{\text{ref}}(s)$ | $T_{\mathcal{H}\text{-BEM}}$ | $\delta_{\mathcal{H}}$ |
|-----------|------------------|---------------------|------------------------------|------------------------|
| 16.64 | 102 | 4.6 | 37.1 | 0.33 |
| 33.27 | 131 | 37.1 | 262.6 | 0.14 |
| 40.91 | 131 | 72.2 | 477.3 | 0.11 |
| 66.61 | 201 | 376.0 | 2046.8 | $5.9 \cdot 10^{-2}$ |
| 83.60 | 715 | 2316.0 | 4135.6 | $4.5 \cdot 10^{-2}$ |
| 92.79 | 929 | 3933.0 | 6055.2 | $3.8 \cdot 10^{-2}$ |
| 128 | 918 | 9322.0 | 17125.2 | $2.7 \cdot 10^{-2}$ |

For simplicity, we consider that the reference solver time is

$$T_{\text{ref}} = n_{\text{ref}} \times t_{\text{iter}}(\varepsilon_{\text{ACA}}),$$

where $t_{\text{iter}}(\varepsilon_{\text{ACA}})$ represents the computational time per iteration corresponding to a \mathcal{H} -matrix-vector product operation

$$t_{\text{iter}}(\varepsilon_{\text{ACA}}) = t_{\text{full}} + t_{\text{low}}(\varepsilon_{\text{ACA}});$$

t_{full} (resp. t_{low}) being the time associated to the matrix-vector product over the non-admissible (resp. admissible) blocks defined through the η -admissibility condition (resp. for an approximation accuracy ε_{ACA}). As expected, we observe that the number of iterations drastically increases with the non-dimensional frequency of the problem, motivating the need for an efficient preconditioner.

5.3. Preliminary experiments

We are looking for an inner-outer two-level GMRES solver which consistently outperforms the reference solver. There are many parameters involved: the maximum numbers of outer and inner iterations respectively denoted by N_{out} and N_{in} and the GMRES residual based stopping criteria of the outer and inner solvers respectively denoted by $\varepsilon_{\text{tol}(\text{out})}$ and $\varepsilon_{\text{tol}(\text{in})}$. The stopping criteria $\varepsilon_{\text{tol}(\text{out})}$ (resp. $\varepsilon_{\text{tol}(\text{in})}$) is set equal (resp. less than or equal) to 10^{-6} . The main concern is the choice of the outer preconditioner accuracy ε' and also the choices of the parameters $\varepsilon_{\text{tol}(\text{in})}$ and N_{in} used at the inner level to solve the preconditioning system. To drive this choice, the preconditioner must be cheap to compute and apply. Our preconditioner is available without any additional time. For a fixed parameter η , the unitary cost of application is $t_{\text{iter}}(\varepsilon')$. Let n_{out} and n_{in} respectively denote the cumulative numbers of outer and inner iterations before convergence of the inner-outer solver. Then, T_{prec} denotes the time spent on the \mathcal{H} -matrix-vector product operations in the inner-outer solver. It writes

$$T_{\text{prec}}(\varepsilon, \varepsilon') = T_{\text{out}}(\varepsilon) + T_{\text{in}}(\varepsilon');$$

where T_{out} and T_{in} are respectively the total times to perform matrix-vector products at the outer and inner levels of the preconditioned solver, with

$$T_{\text{out}}(\varepsilon) = n_{\text{out}} \times t_{\text{iter}}(\varepsilon) \text{ and } T_{\text{in}}(\varepsilon') = n_{\text{in}} \times t_{\text{iter}}(\varepsilon').$$

Since $t_{\text{iter}}(\varepsilon)$ and $t_{\text{iter}}(\varepsilon')$ depends on the choice of prescribed accuracy in the low-rank approximations, it is interesting to quantify this dependance.

Complexity of the computational time per iteration vs ε' . The time t_{iter} to perform a \mathcal{H} -matrix vector product is decomposed into two parts. The times t_{low} and t_{full} associated to the low-rank admissible and non-admissible blocks are respectively represented in Figure 3a and 3b for the six first meshes defined in

Table 3. At a fixed non-dimensional wavenumber, we observe that $t_{\text{low}}(\varepsilon_{\text{ACA}}) \sim O(|\log(\varepsilon_{\text{ACA}})|)$ while as expected t_{full} is constant for all values of ε_{ACA} . We assume this behavior in the following. We also observe that t_{low} dominates t_{full} .

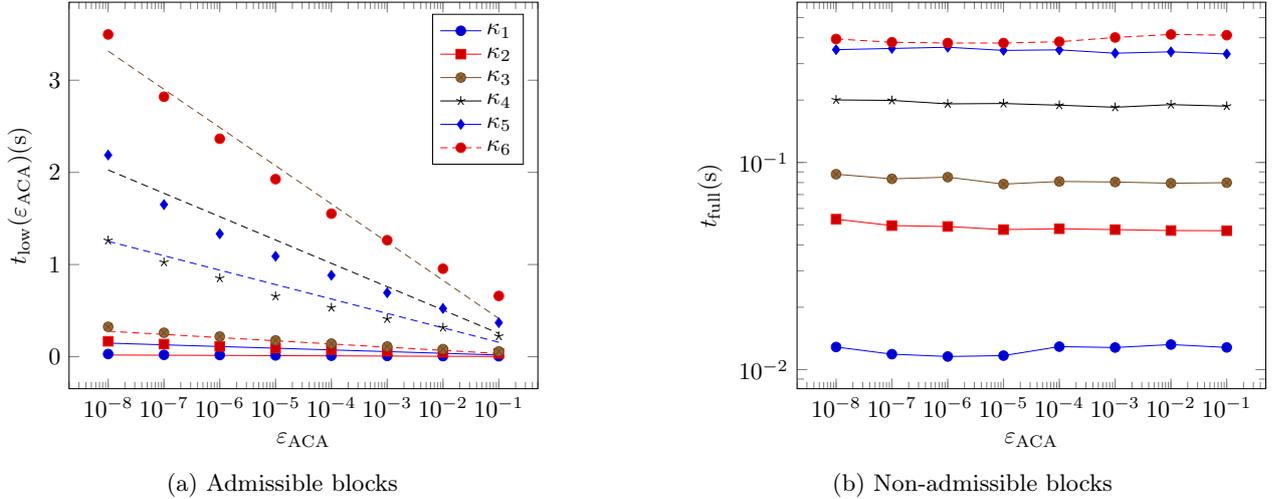


Figure 3: Illustration of the complexity of t_{low} and t_{full} (acoustic equation).

An insight on the expectable efficiency as ε' varies. We now have all the ingredients to study the effects of the different parameters of our preconditioner. Since we use \mathcal{H} -matrix representations with accuracies respectively ε' and ε , the times in the outer and inner levels T_{out} and T_{in} are given by

$$T_{\text{out}}(\varepsilon) = n_{\text{out}}(t_{\text{full}} + t_{\text{low}}(\varepsilon)) \text{ and } T_{\text{in}}(\varepsilon') = n_{\text{in}}(t_{\text{full}} + t_{\text{low}}(\varepsilon')). \quad (16)$$

By taking into account the fact that $t_{\text{low}}(\varepsilon_{\text{ACA}}) \simeq O(|\log(\varepsilon_{\text{ACA}})|)$, this yields

$$T_{\text{prec}} = n_{\text{prec}} t_{\text{full}} + t_{\text{low}}(\varepsilon)(n_{\text{out}} + c n_{\text{in}})$$

where $n_{\text{prec}} = n_{\text{out}} + n_{\text{in}}$ denotes the total number of iterations of the inner-outer solver and $c = \frac{\log(\varepsilon')}{\log(\varepsilon)} \in [0, 1]$. We introduce $\delta T = \frac{T_{\text{prec}}}{T_{\text{ref}}}$ the ratio of time of the preconditioned solver over the one of the reference solver which measures the gain obtained over the reference solver. It is given by

$$\delta T = \frac{n_{\text{prec}} t_{\text{full}} + t_{\text{low}}(\varepsilon)(n_{\text{out}} + c n_{\text{in}})}{n_{\text{ref}}(t_{\text{full}} + t_{\text{low}}(\varepsilon))} = \frac{n_{\text{prec}}}{n_{\text{ref}}} + \frac{t_{\text{low}}(\varepsilon)}{t_{\text{full}} + t_{\text{low}}(\varepsilon)} \frac{(c-1)n_{\text{in}}}{n_{\text{ref}}}.$$

Denoting by

$$\tau = \frac{t_{\text{low}}}{t_{\text{full}} + t_{\text{low}}}$$

the ratio of time devoted to the low-rank admissible blocks, we have

$$\delta T = \frac{n_{\text{prec}}}{n_{\text{ref}}} + \tau(c-1) \frac{n_{\text{in}}}{n_{\text{ref}}}. \quad (17)$$

For simplicity, we introduce the dimensionless numbers

$$\bar{n}_{\text{prec}} = \frac{n_{\text{prec}}}{n_{\text{ref}}} \text{ and } \bar{n}_{\text{in}} = \frac{n_{\text{in}}}{n_{\text{ref}}}. \quad (18)$$

The idea now is to have an insight on the bound of efficiency that can be obtained with this preconditioner. Thus we consider the following system of inequalities, with the dimensionless numbers

$$\begin{cases} \bar{n}_{\text{prec}} - \tau(1-c)\bar{n}_{\text{in}} \leq 1, & (19a) \\ -\bar{n}_{\text{prec}} + \bar{n}_{\text{in}} < 0, & (19b) \end{cases}$$

where the second inequality follows from the fact that $n_{\text{prec}} = n_{\text{out}} + n_{\text{in}}$. We also introduce \bar{N}_{in} and $\bar{N}_{\text{prec}} := \bar{N}_{\text{in}} + \bar{N}_{\text{out}}$ respectively the dimensionless maximum numbers of iterations of the inner and inner-outer, preconditioned solvers, defined as in (18).

We denote by \mathcal{S}_n the set of admissible values of $(\bar{n}_{\text{in}}, \bar{n}_{\text{prec}})$, filled in magenta in Figure 4, above the blue line

$$\mathcal{S}_n = \{(\bar{n}_{\text{in}}, \bar{n}_{\text{prec}}) \mid 0 < \bar{n}_{\text{in}} \leq \bar{N}_{\text{in}} \text{ and } \bar{n}_{\text{in}} < \bar{n}_{\text{prec}} \leq \bar{N}_{\text{prec}}\}.$$

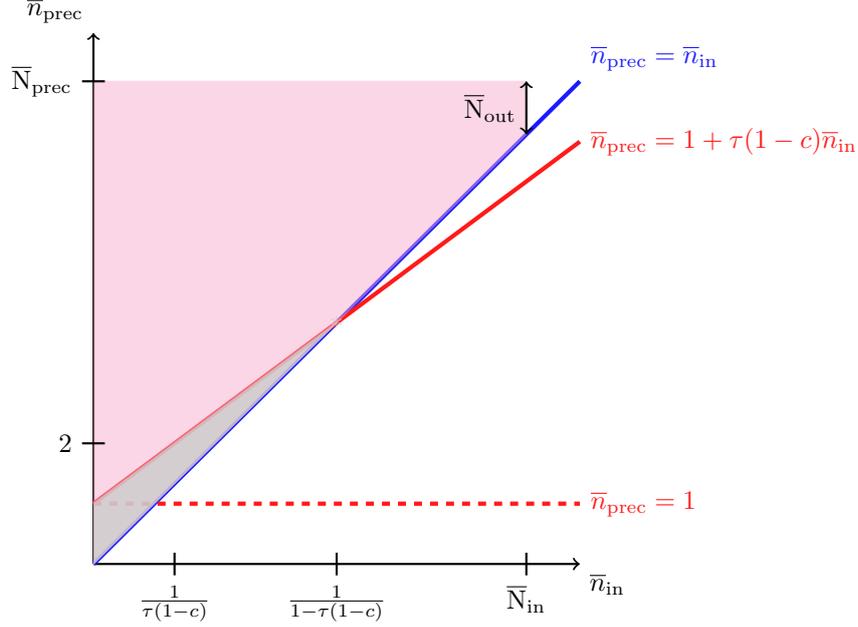


Figure 4: Geometrical illustration (filled in gray color) of the set of values of $(\bar{n}_{\text{in}}, \bar{n}_{\text{prec}})$ yielding to faster iterative solver in comparison to the reference solver performances (representation with $\tau(1-c) = 0.75$).

The corresponding area is

$$|\mathcal{S}_n| = \frac{1}{2}\bar{N}_{\text{in}}^2 + \bar{N}_{\text{in}}\bar{N}_{\text{out}}.$$

We denote by $\mathcal{S}_c \subset \mathcal{S}_n$ the subset of dimensionless couple $(\bar{n}_{\text{in}}, \bar{n}_{\text{prec}})$ leading to faster solver for the preconditioner defined with the parameter c

$$\mathcal{S}_c = \{(\bar{n}_{\text{in}}, \bar{n}_{\text{prec}}) \mid \bar{n}_{\text{prec}} - \tau(1-c)\bar{n}_{\text{in}} \leq 1\} \cap \mathcal{S}_n.$$

Its corresponding area is

$$|\mathcal{S}_c| = \min\left(\bar{N}_{\text{in}}, \frac{1}{1-\tau(1-c)}\right) - \frac{1}{2} \left[\min\left(\bar{N}_{\text{in}}, \frac{1}{1-\tau(1-c)}\right) \right]^2 (1-\tau(1-c)).$$

On Figure 4, $\bar{N}_{\text{in}} \geq \frac{1}{1-\tau(1-c)}$ such that the maximum area corresponding to the region $\mathcal{S}_c^{\text{max}}$, filled in gray is

$$|\mathcal{S}_c^{\text{max}}| := \frac{1}{2(1-\tau(1-c))}$$

while otherwise the corresponding area is

$$|\mathcal{S}_c| = \bar{N}_{\text{in}} - \frac{1}{2}\bar{N}_{\text{in}}^2(1-\tau(1-c)) < |\mathcal{S}_c^{\text{max}}|.$$

If \mathcal{S}_f denotes the global subset of parameters yielding to a faster solver comparatively to the reference solver and assuming that there exists c_{min} such that $c \in [c_{\text{min}}, 1]$, we have

$$\mathcal{S}_f := \bigcup_{c_{\text{min}} \leq c \leq 1} \mathcal{S}_c = \mathcal{S}_{c_{\text{min}}} \subset \mathcal{S}_0.$$

Indeed we have $\mathcal{S}_{c_1} \subseteq \mathcal{S}_{c_2}$ for $c_1 \geq c_2$.

Note that for some preconditioners such that $c \geq c_{\text{min}}$, the rate of values yielding to a faster solver among all the admissible values is equal to $\frac{|\mathcal{S}_{c_{\text{min}}}^{\text{max}}|}{|\mathcal{S}_n|} < \frac{|\mathcal{S}_0^{\text{max}}|}{|\mathcal{S}_n|}$. It is straightforward to see that

- at a fixed value of τ , i.e., for a given \mathcal{H} -matrix representation of the BEM problem we have

$$\partial_c |\mathcal{S}_c| < 0.$$

Indeed, the set of values \mathcal{S}_c yielding to a fast solver becomes smaller and smaller as the parameter c increases and tends to the following “limit” subset of values ($c = 1$)

$$\mathcal{S}_{\text{lim}} := \mathcal{S}_1.$$

For a very accurate preconditioner, whose evaluation is almost as expensive as the \mathcal{H} -BEM matrix, the efficiency occurs only when the number of iterations of the inner-outer solver is smaller or equal than n_{ref} , i.e., $\bar{n}_{\text{prec}} \leq 1$. As expected, requiring a high accuracy on the preconditioner is very constraining for efficiency.

- On the other hand, for a fixed preconditioner accuracy, i.e., for a fixed value of c , we have

$$\partial_\tau |\mathcal{S}_c| \geq 0.$$

Since τ increases with the non-dimensional wavenumber, it turns out that the rate of dimensionless numbers leading to a fast solver will increase with the non-dimensional wavenumber. On the contrary, when τ tends to 0 (i.e., a \mathcal{H} -BEM matrix tending to a standard full BEM matrix representation), the efficiency occurs only for the values of \mathcal{S}_{lim} . The proposed preconditioning strategy becomes inefficient when τ is small, i.e., either in the case of a \mathcal{H} -BEM with a very coarse accuracy ε or a \mathcal{H} -BEM with small number of low-rank admissible blocks.

Since the ratio of time of inner-outer solver over the time of the reference solver is given by (see (17), (18))

$$\delta T = \bar{n}_{\text{prec}} - \tau(1-c)\bar{n}_{\text{in}}$$

using inequality (19b), we obtain

$$\delta T > \bar{n}_{\text{prec}}(1-\tau(1-c)).$$

In other words, for an inner-outer solver of dimensionless number of iterations \bar{n}_{prec} , with a preconditioner of parameter $c = \frac{\log(\varepsilon')}{\log(\varepsilon)}$; *a priori*, the relative gain

$$\delta_r T = \frac{T_{\text{ref}} - T_{\text{prec}}}{T_{\text{ref}}}$$

is bounded above by the predicted gain

$$\hat{\delta}_r T = 1 - \bar{n}_{\text{prec}}(1 - \tau(1 - c)). \quad (20)$$

Subsequently, we consider two kinds of numerical tests to illustrate the effectiveness of the inner-outer preconditioning strategy. First, we assume that the residual based stopping criteria of the inner GMRES solver $\varepsilon_{\text{tol}(\text{in})}$ can vary: $\varepsilon_{\text{tol}(\text{in})} \in \{10^{-1}, \dots, 10^{-6}\}$. Then, we consider that the maximum number of iterations of the inner solver N_{in} can vary while the GMRES stopping residual is fixed to $\varepsilon_{\text{tol}(\text{in})} = 10^{-6}$. For both cases, the outer GMRES residual based stopping criteria is fixed to $\varepsilon_{\text{tol}(\text{out})} = 10^{-6}$ (similarly to what is prescribed in the reference solver).

5.4. Inner-outer solver: for different values of $\varepsilon_{\text{tol}(\text{in})}$

In this Section, the maximum number of inner iterations is fixed ($N_{\text{in}} = 2000$) per outer iteration and the parameter $\varepsilon_{\text{tol}(\text{in})}$ varies. For the different non-dimensional wavenumbers $r\kappa_i$, we report on the Table 5 the performances corresponding to the five most efficient preconditioners such that $\delta_r T > 0$. For each case, the number of couple of parameters $(\varepsilon', \varepsilon_{\text{tol}(\text{in})})$ tested is 36. We denote $\mathcal{P}_+ = \{(\varepsilon', \varepsilon_{\text{tol}(\text{in})}) \text{ s.t. } \delta T > 0\}$ and $|\mathcal{P}_+| = \text{Card}(\mathcal{P}_+)$ the number of efficient preconditioners.

Table 5: *The five most efficient inner-outer solvers for each non-dimensional wavenumber.*

| $r\kappa = 16.67$ and $\tau = 0.69$; $ \mathcal{P}_+ = 3$ | | | | | | |
|--|----------------|--------------------------------|------------------|-----------------|------------------------|------------------|
| Ranking | ε' | $\varepsilon_{\text{tol(in)}}$ | n_{out} | n_{in} | $\hat{\delta}_r T(\%)$ | $\delta_r T(\%)$ |
| 1 | 10^{-4} | 10^{-2} | 3 | 120 | 21.3 | 10.7 |
| 2 | 10^{-6} | 10^{-6} | 1 | 102 | 16.5 | 8.8 |
| 3 | 10^{-5} | 10^{-3} | 2 | 109 | 19.5 | 7.7 |
| $r\kappa = 33.27$ and $\tau = 0.76$; $ \mathcal{P}_+ = 13$ | | | | | | |
| 1 | 10^{-5} | 10^{-3} | 2 | 149 | 17.5 | 27.0 |
| 2 | 10^{-4} | 10^{-3} | 2 | 153 | 26.6 | 26.5 |
| 3 | 10^{-3} | 10^{-2} | 3 | 185 | 24.5 | 25.8 |
| 4 | 10^{-4} | 10^{-2} | 3 | 171 | 17.6 | 14.1 |
| 5 | 10^{-5} | 10^{-2} | 3 | 168 | 6.6 | 14.0 |
| $r\kappa = 40.91$ and $\tau = 0.79$; $ \mathcal{P}_+ = 20$ | | | | | | |
| 1 | 10^{-3} | 10^{-2} | 3 | 153 | 39.5 | 37.9 |
| 2 | 10^{-4} | 10^{-2} | 3 | 152 | 28.2 | 35.7 |
| 3 | 10^{-4} | 10^{-3} | 2 | 143 | 32.9 | 29.9 |
| 4 | 10^{-5} | 10^{-3} | 2 | 143 | 22.0 | 29.5 |
| 5 | 10^{-2} | 10^{-1} | 6 | 189 | 39.0 | 28.7 |
| $r\kappa = 66.61$ and $\tau = 0.86$; $ \mathcal{P}_+ = 19$ | | | | | | |
| 1 | 10^{-5} | 10^{-3} | 2 | 228 | 22.6 | 40.7 |
| 2 | 10^{-4} | 10^{-2} | 3 | 274 | 21.6 | 39.3 |
| 3 | 10^{-2} | 10^{-1} | 6 | 369 | 34.2 | 37.6 |
| 4 | 10^{-3} | 10^{-1} | 6 | 356 | 17.0 | 32.9 |
| 5 | 10^{-5} | 10^{-2} | 3 | 274 | 6.8 | 32.8 |
| $r\kappa = 83.60$ and $\tau = 0.86$; $ \mathcal{P}_+ = 14$ | | | | | | |
| 1 | 10^{-4} | 10^{-2} | 3 | 749 | 40.2 | 44.9 |
| 2 | 10^{-5} | 10^{-2} | 3 | 749 | 28.8 | 38.2 |
| 3 | 10^{-3} | 10^{-1} | 6 | 1137 | 26.3 | 36.4 |
| 4 | 10^{-6} | 10^{-2} | 3 | 749 | 17.5 | 25.2 |
| 5 | 10^{-4} | 10^{-1} | 6 | 1114 | 10.9 | 18.3 |
| $r\kappa = 92.79$ and $\tau = 0.88$; $ \mathcal{P}_+ = 16$ | | | | | | |
| 1 | 10^{-5} | 10^{-2} | 3 | 1096 | 20.7 | 38.4 |
| 2 | 10^{-3} | 10^{-1} | 6 | 1522 | 25.8 | 35.7 |
| 3 | 10^{-6} | 10^{-3} | 2 | 1046 | 12.0 | 30.0 |
| 4 | 10^{-6} | 10^{-2} | 3 | 1098 | 7.5 | 28.4 |
| 5 | 10^{-4} | 10^{-1} | 6 | 1475 | 10.6 | 28.2 |
| $r\kappa = 128$ and $\tau = 0.91$; $ \mathcal{P}_+ = 30$ | | | | | | |
| 1 | 10^{-4} | 10^{-2} | 3 | 1045 | 37.5 | 49.4 |
| 2 | 10^{-4} | 10^{-4} | 2 | 1126 | 32.8 | 45.5 |
| 3 | 10^{-3} | 10^{-1} | 6 | 1376 | 34.7 | 44.4 |
| 4 | 10^{-4} | 10^{-1} | 6 | 1276 | 23.6 | 42.2 |
| 5 | 10^{-4} | 10^{-5} | 2 | 1200 | 28.4 | 37.4 |

One observes that a small number of preconditioners are efficient for the smallest non-dimensional wavenumber, which has the lowest number of degrees of freedom and the smallest value of τ . Indeed, after recompression, the preconditioners may not contain enough information for this case. Therefore high

accuracies ε' are required to obtain an efficient preconditioner in this case. The total number of efficient preconditioners increases with the non-dimensional frequency with $|\mathcal{P}_+| = 30$ (among the 36 tested) observed for the largest non-dimensional wavenumber. The effectively observed gain $\delta_r T$ is greater than the maximum predicted gain $\hat{\delta}_r T$ except for $r\kappa = 16.67$. Indeed, our estimation does not take into account the time of the Arnoldi procedure which decreases considerably thanks to the balance of the total number of iterations at the outer and inner levels. The largest value of $\delta_r T = 50\%$ is achieved for $r\kappa = 128$. From this particular numerical example, a good choice of preconditioner can be the one with $\varepsilon' = 10^{-4}$ and $\varepsilon_{\text{tol(in)}} = 10^{-2}$. More generally, we remark that the residual of the inner solver of the efficient solver is such as $\varepsilon_{\text{tol(in)}} \geq \varepsilon'$. This corroborates with the fact that any iterative solver does not need to be as accurate as the \mathcal{H} -BEM matrix representation.

We represent on Figure 5 the efficiency $\delta_r T$ (as $\varepsilon_{\text{tol(in)}}$ varies) for preconditioners of \mathcal{P}_+ and their associated non-dimensional wavenumber $r\kappa_i$. Since the use of $\varepsilon' = 10^{-2}$ or 10^{-6} leads to a small number of efficient combinations, the results are put together on Figure 5a with the label κ_{i-n} meaning the case $r\kappa_i$ with $\varepsilon' = 10^{-n}$. For the sake of clarity we plot only the cases for which $\delta_r T > 0$.

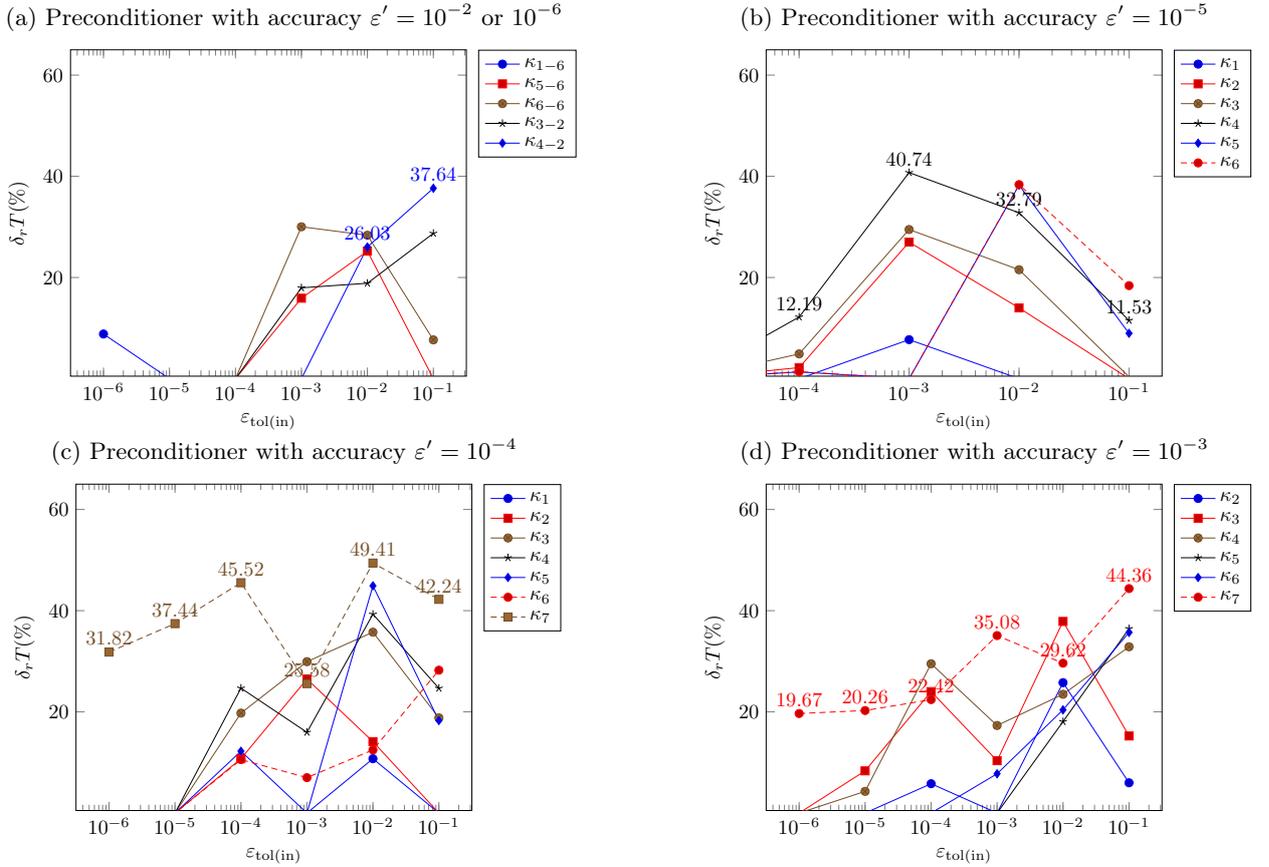


Figure 5: Relative gain with respect to the reference solver $\delta_r T$.

5.5. Inner-outer solver: for different N_{in}

Now, we consider a fixed inner solver tolerance $\varepsilon_{\text{tol(in)}} = 10^{-6}$ while the maximum number of inner iterations varies ($N_{\text{in}} \in \{10, 15, 30, 50, 60, 80, 100, 150\}$). Therefore, the number of iterations will mostly be the stopping criterion for the inner solver. For each case the number of couple of parameters $(\varepsilon', N_{\text{in}})$ is 56 since we have also tested the preconditioner of accuracy $\varepsilon' = 10^{-1}$. Again we denote $\mathcal{P}_+ = \{(\varepsilon', N_{\text{in}}) \text{ s.t. } \delta T > 0\}$ and $|\mathcal{P}_+| = \text{Card}(\mathcal{P}_+)$ the number of efficient preconditioners. We represent on Table 6 the five most efficient preconditioners for each non-dimensional wavenumber. In the above table, $\varepsilon_{\text{tol(in)}}^{\text{conv}}$ represents the

Table 6: The five most efficient inner-outer solvers for each non-dimensional wavenumber with varying N_{in} .

| $r\kappa = 16.67$ and $\tau = 0.69$; $ \mathcal{P}_+ = 9$ | | | | | | | |
|---|----------------|-----------------|------------------|-----------------|------------------------|------------------|--|
| Ranking | ε' | N_{in} | n_{out} | n_{in} | $\hat{\delta}_r T(\%)$ | $\delta_r T(\%)$ | $\varepsilon_{\text{tol(in)}}^{\text{conv}}$ |
| 1 | 10^{-4} | 60 | 2 | 120 | 21.9 | 25.5 | $1.87 \cdot 10^{-5}$ |
| 2 | 10^{-6} | 150 | 1 | 102 | 16.5 | 18.3 | $8.01 \cdot 10^{-7}$ |
| 3 | 10^{-5} | 60 | 2 | 120 | 11.5 | 15.0 | $1.56 \cdot 10^{-5}$ |
| 4 | 10^{-3} | 80 | 2 | 160 | 10.1 | 10.3 | $7.12 \cdot 10^{-6}$ |
| 5 | 10^{-7} | 150 | 1 | 102 | 17.8 | 9.7 | $8.02 \cdot 10^{-7}$ |
| $r\kappa = 33.27$ and $\tau = 0.76$; $ \mathcal{P}_+ = 29$ | | | | | | | |
| 1 | 10^{-4} | 100 | 2 | 200 | 4.3 | 37.1 | $8.49 \cdot 10^{-6}$ |
| 2 | 10^{-5} | 80 | 2 | 160 | 23.4 | 34.6 | $4.05 \cdot 10^{-5}$ |
| 3 | 10^{-4} | 30 | 7 | 210 | -2.8 | 31.2 | $1.10 \cdot 10^{-2}$ |
| 4 | 10^{-4} | 15 | 14 | 210 | -6.1 | 28.0 | $5.13 \cdot 10^{-2}$ |
| 5 | 10^{-4} | 60 | 4 | 240 | -15.6 | 26.6 | $8.02 \cdot 10^{-4}$ |
| $r\kappa = 40.91$ and $\tau = 0.79$; $ \mathcal{P}_+ = 42$ | | | | | | | |
| 1 | 10^{-2} | 50 | 4 | 200 | 36.2 | 47.8 | $2.62 \cdot 10^{-3}$ |
| 2 | 10^{-3} | 60 | 3 | 180 | 29.0 | 46.8 | $1.23 \cdot 10^{-3}$ |
| 3 | 10^{-3} | 30 | 6 | 180 | 27.9 | 44.3 | $1.24 \cdot 10^{-2}$ |
| 4 | 10^{-4} | 80 | 2 | 160 | 25.0 | 43.3 | $1.25 \cdot 10^{-4}$ |
| 5 | 10^{-2} | 15 | 13 | 195 | 35.0 | 42.9 | $4.32 \cdot 10^{-2}$ |
| $r\kappa = 66.608$ and $\tau = 0.86$; $ \mathcal{P}_+ = 47$ | | | | | | | |
| 1 | 10^{-3} | 15 | 19 | 285 | 30.3 | 47.6 | $5.59 \cdot 10^{-2}$ |
| 2 | 10^{-3} | 100 | 3 | 300 | 30.6 | 46.7 | $9.73 \cdot 10^{-4}$ |
| 3 | 10^{-4} | 100 | 3 | 300 | 14.3 | 45.7 | $9.33 \cdot 10^{-4}$ |
| 4 | 10^{-2} | 15 | 23 | 345 | 35.4 | 44.5 | $5.59 \cdot 10^{-2}$ |
| 5 | 10^{-4} | 15 | 18 | 270 | 18.5 | 44.1 | $5.59 \cdot 10^{-2}$ |
| $r\kappa = 83.602$ and $\tau = 0.86$; $ \mathcal{P}_+ = 30$ | | | | | | | |
| 1 | 10^{-3} | 50 | 22 | 1100 | 27.6 | 37.9 | $1.09 \cdot 10^{-2}$ |
| 2 | 10^{-2} | 10 | 105 | 1050 | 42.9 | 36.7 | $1.04 \cdot 10^{-1}$ |
| 3 | 10^{-3} | 30 | 36 | 1080 | 28.0 | 36.3 | $2.32 \cdot 10^{-2}$ |
| 4 | 10^{-3} | 60 | 19 | 1140 | 25.2 | 35.7 | $8.87 \cdot 10^{-3}$ |
| 5 | 10^{-3} | 80 | 15 | 1200 | 21.6 | 34.7 | $5.59 \cdot 10^{-2}$ |
| $r\kappa = 92.788$ and $\tau = 0.88$; $ \mathcal{P}_+ = 45$ | | | | | | | |
| 1 | 10^{-2} | 15 | 92 | 1380 | 46.0 | 59.9 | $6.81 \cdot 10^{-2}$ |
| 2 | 10^{-3} | 15 | 85 | 1275 | 34.0 | 59.2 | $6.81 \cdot 10^{-2}$ |
| 3 | 10^{-3} | 30 | 47 | 1410 | 29.3 | 58.7 | $2.54 \cdot 10^{-2}$ |
| 4 | 10^{-2} | 10 | 138 | 1380 | 44.3 | 56.2 | $1.10 \cdot 10^{-1}$ |
| 5 | 10^{-3} | 80 | 19 | 1520 | 25.3 | 55.3 | $7.38 \cdot 10^{-3}$ |
| $r\kappa = 128$ and $\tau = 0.91$; $ \mathcal{P}_+ = 55$ | | | | | | | |
| 1 | 10^{-2} | 50 | 39 | 1950 | 52.2 | 63.3 | $1.58 \cdot 10^{-2}$ |
| 2 | 10^{-2} | 30 | 65 | 1950 | 51.6 | 62.7 | $3.15 \cdot 10^{-2}$ |
| 3 | 10^{-2} | 60 | 34 | 2040 | 50.2 | 61.5 | $1.23 \cdot 10^{-2}$ |
| 4 | 10^{-2} | 15 | 121 | 1815 | 53.5 | 61.5 | $7.66 \cdot 10^{-2}$ |
| 5 | 10^{-3} | 50 | 38 | 1900 | 37.0 | 59.1 | $1.58 \cdot 10^{-2}$ |

minimum observed value of the residual in the inner systems after N_{in} iterations. In this case also, we remark that the efficiency $\delta_r T$ effectively observed is greater than the predicted one $\hat{\delta}_r T$ with some spurious negative values for $r\kappa = 33.27$. As for variable $\varepsilon_{\text{tol}(\text{in})}$, the number of efficient preconditioners $|\mathcal{P}_+|$ increases with the non-dimensional wavenumber and $|\mathcal{P}_+| = 55$ (among the 56 tested) for the largest non-dimensional wavenumber.

We observe that the number of outer iterations n_{out} has largely increased comparatively to the previous case. Indeed [it is more constraining to fix the parameter \$N_{\text{in}}\$](#) than to fix the GMRES residual to $\varepsilon_{\text{tol}(\text{in})} = 10^{-6}$. [It thus](#) leads to the use of more preconditioning Krylov subspace basis vectors for the outer solver. The value of $\delta_r T$ increases with the non-dimensional wavenumber. The largest value $\delta_r T = 60\%$ is achieved for $r\kappa_7$ and is greater than the highest one with $\varepsilon_{\text{tol}(\text{in})}$ varying. This increase in efficiency can be explained by a finer balance of the total number of iterations between the outer and inner levels in conjunction with a preconditioner of lower accuracy. The preconditioners with very high accuracy $\varepsilon' = 10^{-7}$ or 10^{-6} are among the top five ranked only for the smallest non-dimensional wavenumber. On the other hand, a preconditioner with a larger value of ε' can be used as the non-dimensional wavenumber increases. The same observation holds for the residual $\varepsilon_{\text{tol}(\text{in})}^{\text{conv}}$. We represent on Figure 6 the efficiency $\delta_r T$ (as N_{in} varies) for preconditioners of \mathcal{P}_+ and their associated non-dimensional wavenumber $r\kappa_i$. Figures 6(b-d) are restricted to only one ε' since a lot of combinations lead to an efficient preconditioner. On the contrary, since the use of $\varepsilon' = 10^{-5}$, 10^{-6} or 10^{-7} leads to a small number of efficient combinations, the results are put together on Figure 6(a) with the label κ_{i-n} meaning the case $r\kappa_i$ with $\varepsilon' = 10^{-n}$.

In this case, we have also tested the preconditioner of accuracy $\varepsilon' = 10^{-1}$. Surprisingly, despite the ‘‘poor spectral clustering’’ observed in the first part of the numerical tests, we observe in Figure 7 that it actually becomes efficient for the five largest non-dimensional wavenumbers, i.e when N_{DOF} becomes large yielding to a coarse approximation with a good amount of information. However its efficiency rapidly decreases with N_{in} and inefficiency occurs in general for large values of N_{in} . This is certainly due to unnecessary iterations performed with the intent to achieve the prescribed tolerance $\varepsilon_{\text{tol}(\text{in})} = 10^{-6}$ (too small compared to 10^{-1}). Finally, we note that the efficiency of the preconditioner with $\varepsilon' = 10^{-1}$ improves as the non-dimensional wavenumber increases.

Ultimately for problems with a high non-dimensional wavenumber $r\kappa$ one can consider a preconditioner with $\varepsilon' = 10^{-1}$ and $\varepsilon_{\text{tol}(\text{in})} \sim 10^{-1}$ or $\varepsilon' = 10^{-1}$ and a moderate maximum number of inner iterations, for instance $N_{\text{in}} \leq 60$.

Elastodynamic case

Finally we consider the \mathcal{H} -BEM iterative solver for an elastodynamic problem, which is a 3D vector-valued problem. We still consider the single-layer formulation, the spherical geometry and the same meshes as in the acoustic case. The density of points per S-wavelength is fixed to $n_{\lambda_S} = 10$, the corresponding S-non-dimensional wavenumber and N_{DOF} are reported in Table 7.

Table 7: Number of DOF, non-dimensional S-wavenumber and number of S-wavelengths along the characteristic dimension corresponding to each mesh (Elastodynamic case).

| Mesh i | 1 | 2 | 3 | 4 |
|------------------|--------|---------|---------|---------|
| N_{DOF} | 30 746 | 122 886 | 183 099 | 490 629 |
| $r\kappa_S$ | 16.64 | 33.27 | 40.91 | 66.61 |

The maximum number of iterations is 3000, i.e., $n_{\text{ref}} \leq 3000$. The performances of the reference solver are reported in Table 8.

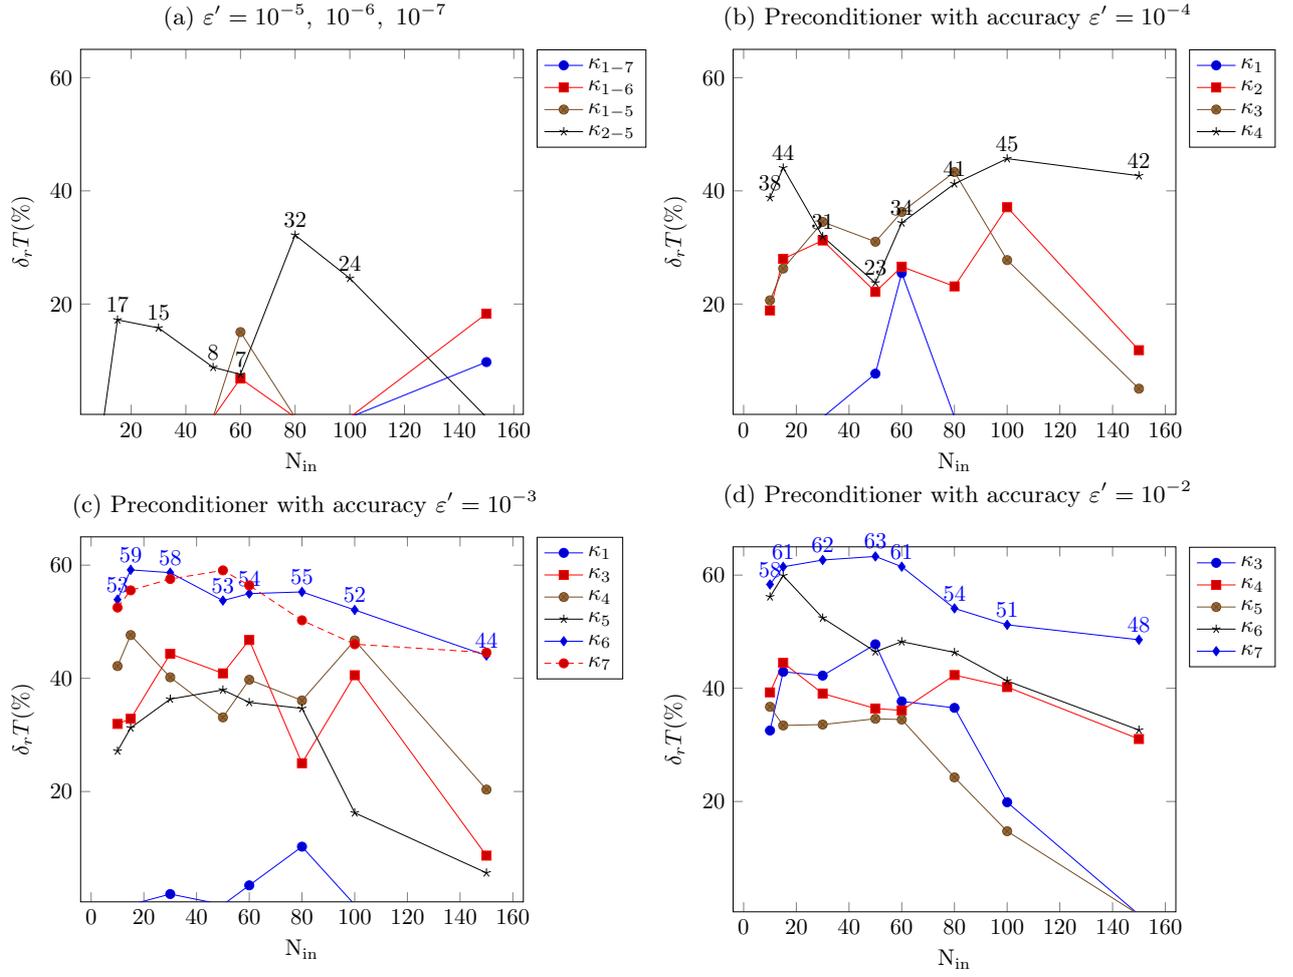


Figure 6: Relative gain with respect to the reference solver $\delta_r T$. For sake of readability we have rounded the value of $\delta_r T$.

Table 8: Performances obtained for the reference solver (Elastodynamic case).

| $r\kappa_S$ | n_{ref} | $T_{\text{ref}}(s)$ | $T_{\mathcal{H}\text{-BEM}(s)}$ | $\delta_{\mathcal{H}}$ |
|-------------|------------------|---------------------|---------------------------------|------------------------|
| 16.64 | 320 | 98.2 | 111.9 | 0.31 |
| 33.27 | 634 | 1210.4 | 885.9 | 0.13 |
| 40.91 | 437 | 1387.8 | 1634.2 | 0.10 |
| 66.61 | 2275 | 31 966.2 | 7177.0 | $5.47 \cdot 10^{-2}$ |

We consider that the GMRES threshold of the inner solver is fixed to $\varepsilon_{\text{tol}(\text{in})} = 10^{-6}$ and look for the efficient preconditioners yielding to better performances in comparison to the reference solver by varying the parameter N_{in} . The top five ranked preconditioners are reported in Table 9.

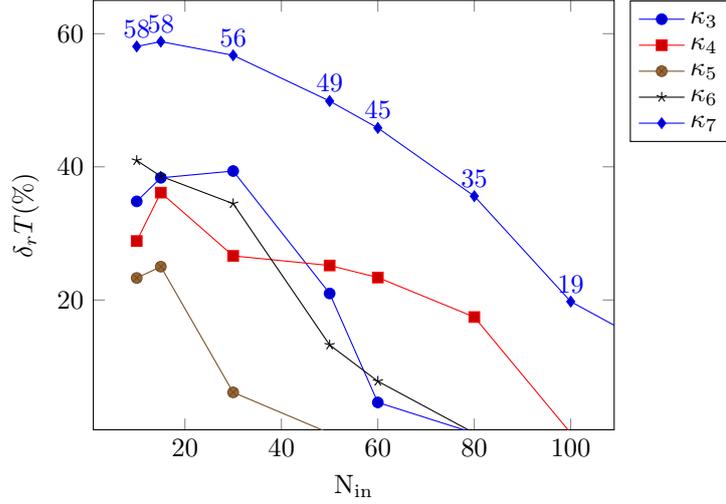


Figure 7: Performances of the inner-outer solver (for $\varepsilon' = 10^{-1}$) as the maximum number of iterations varies. For sake of readability we have rounded the value of $\delta_r T$.

Table 9: The five most efficient inner-outer solvers for each non-dimensional wavenumber (Elastodynamic case).

| $r\kappa_S = 16.67$ and $\tau = 0.66$; $ \mathcal{P}_+ = 21$ | | | | | | | |
|--|----------------|----------|-----------|----------|------------------------|------------------|--------------------------------|
| Ranking | ε' | N_{in} | n_{out} | n_{in} | $\hat{\delta}_r T(\%)$ | $\delta_r T(\%)$ | $\varepsilon_{tol(in)}^{conv}$ |
| 1 | 10^{-3} | 100 | 6 | 600 | -10.8 | 19.0 | $3.46 \cdot 10^{-3}$ |
| 2 | 10^{-3} | 30 | 20 | 102 | -13.3 | 16.5 | $5.26 \cdot 10^{-2}$ |
| 3 | 10^{-3} | 10 | 55 | 120 | -10.6 | 15.9 | $3.40 \cdot 10^{-1}$ |
| 4 | 10^{-4} | 30 | 19 | 160 | -22.9 | 14.2 | $5.25 \cdot 10^{-2}$ |
| 5 | 10^{-3} | 15 | 40 | 600 | -17.0 | 13.6 | $1.94 \cdot 10^{-1}$ |
| $r\kappa_S = 33.27$ and $\tau = 0.75$; $ \mathcal{P}_+ = 46$ | | | | | | | |
| 1 | 10^{-3} | 30 | 32 | 960 | 17.0 | 51.0 | $1.23 \cdot 10^{-1}$ |
| 2 | 10^{-3} | 15 | 62 | 930 | 17.0 | 48.5 | $2.69 \cdot 10^{-1}$ |
| 3 | 10^{-3} | 50 | 21 | 1050 | 10.4 | 47.9 | $7.74 \cdot 10^{-2}$ |
| 4 | 10^{-2} | 30 | 39 | 1170 | 16.8 | 47.8 | $1.23 \cdot 10^{-1}$ |
| 5 | 10^{-3} | 60 | 18 | 1080 | 8.2 | 45.6 | $6.69 \cdot 10^{-2}$ |
| $r\kappa_S = 40.91$ and $\tau = 0.78$; $ \mathcal{P}_+ = 51$ | | | | | | | |
| 1 | 10^{-2} | 60 | 12 | 720 | 30.9 | 56.7 | $6.11 \cdot 10^{-2}$ |
| 2 | 10^{-2} | 30 | 23 | 690 | 32.7 | 56.3 | $1.76 \cdot 10^{-1}$ |
| 3 | 10^{-3} | 150 | 4 | 600 | 29.4 | 55.1 | $5.54 \cdot 10^{-3}$ |
| 4 | 10^{-2} | 50 | 15 | 750 | 27.8 | 54.7 | $8.63 \cdot 10^{-2}$ |
| 5 | 10^{-2} | 150 | 5 | 750 | 28.7 | 54.3 | $5.75 \cdot 10^{-3}$ |
| $r\kappa_S = 66.61$ and $\tau = 0.84$; $ \mathcal{P}_+ = 56$ | | | | | | | |
| 1 | 10^{-1} | 50 | 75 | 3750 | 55.0 | 72.6 | $1.52 \cdot 10^{-1}$ |
| 2 | 10^{-1} | 60 | 65 | 3900 | 53.3 | 72.5 | $1.37 \cdot 10^{-1}$ |
| 3 | 10^{-1} | 30 | 118 | 3540 | 56.9 | 72.3 | $2.16 \cdot 10^{-1}$ |
| 4 | 10^{-2} | 60 | 58 | 3480 | 42.1 | 70.3 | $1.39 \cdot 10^{-1}$ |
| 5 | 10^{-2} | 50 | 67 | 3350 | 44.1 | 70.2 | $1.54 \cdot 10^{-1}$ |

As in acoustics, preconditioners with a coarse accuracy ε' become more efficient as the non-dimensional

wavenumber increases. The number of efficient preconditioners increases as the problem size increases. Full efficiency is even observed for the largest non-dimensional wavenumber, i.e. $|\mathcal{P}_+| = 56$. For the non-dimensional wavenumber $r\kappa_S = 66.61$, the GMRES residual is of the order of $\varepsilon_{\text{tol(in)}}^{\text{conv}} \sim 10^{-1}$. The value observed for $\delta_r T$ is still greater than the predicted maximum value $\hat{\delta}_r T$ (note that some negative spurious values are observed for $r\kappa_S = 16.67$). In Figure 8 we represent $\delta_r T$ for the accuracy ε' and non-dimensional wavenumber $r\kappa_S$ for which efficiency has occurred. The highest gains are observed for the highest non-dimensional wavenumbers with $\delta_r T = 72.6\%$. Generally, the preconditioning appears more efficient for elastodynamics in comparison to the acoustic case. Also, in this case, for problems with larger non-dimensional wavenumbers using the preconditioner of accuracy $\varepsilon' = 10^{-1}$ in conjunction to a moderate number of inner iterations or a GMRES threshold $\varepsilon_{\text{tol(in)}} \sim 10^{-1}$ remains a good compromise.

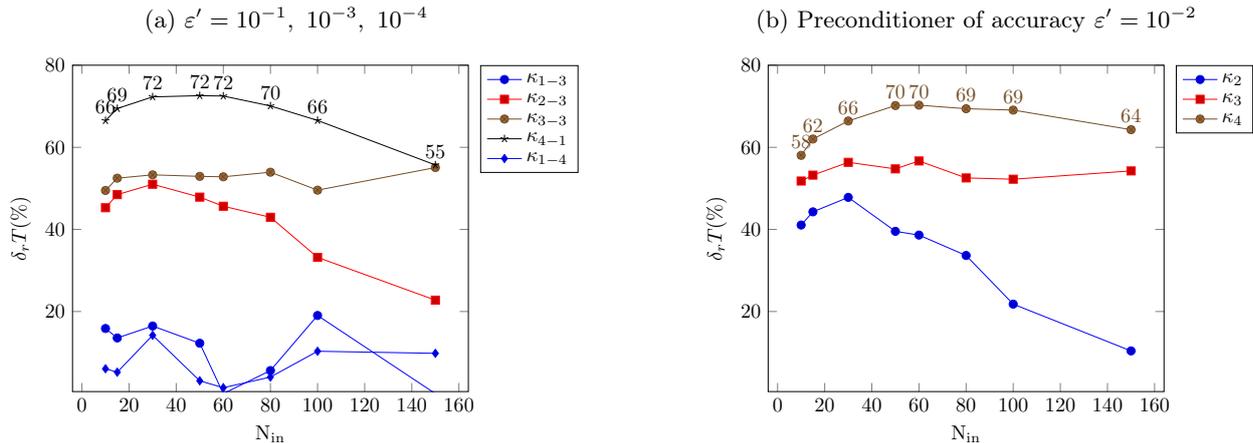


Figure 8: Relative gain with respect to the reference solver $\delta_r T$. Figure (a) reports the efficient preconditioners for $\varepsilon' = 10^{-1}, 10^{-3}$ or 10^{-4} while Figure (b) is restricted to efficient preconditioners for $\varepsilon' = 10^{-2}$ (since there are more efficient combinations in that case).

6. Conclusions

This article addresses the question of the efficient preconditioning of a \mathcal{H} -BEM iterative solver for wave propagation problems with oscillatory kernels. The adopted preconditioning strategy involves an inner-outer GMRES solver, which yields to a two-level iterative method.

The preconditioner is preferably set as a “poor”, coarse compression of the BEM system matrix. We have used the same data-sparse representation structure for both the \mathcal{H} -BEM and the preconditioner (i.e., a similar low-rank admissibility condition is used) while the accuracy of the low-rank blocks of the latter is worse. The advantage of doing so is that there are no additional memory requirements and no increased computational time since the low-rank approximation involves a cumulative rank-one matrix addition until the highest precision is reached.

Several numerical illustrations have been proposed. The efficiency of the preconditioner depends mainly on the maximum number of inner iterations and also on the inner GMRES threshold. We have considered numerical tests investigating the performances of the inner-outer solver as these parameters vary. We have derived an upper bound on the gain that can be computed with such a preconditioning strategy. We have sketched and illustrated throughout the numerical tests that for problems with a small non-dimensional wavenumber, the preconditioner has to be chosen with a high accuracy. While for a larger non-dimensional wavenumber (i.e., problems with a large number of degrees of freedom) a coarse approximation can be chosen. The inner GMRES threshold can be set of the order of 10^{-1} , or a moderate maximum number of iterations can be set.

The compression is only performed over the admissible blocks. Therefore the time of a matrix vector product for non admissible blocks is the same for both the original \mathcal{H} -matrix representation of the system

matrix and for the preconditioner, and may be a limiting factor for the gain. In this respect, one could think to set the preconditioner with a different sparsity pattern parameter (using a different η -admissibility condition). The drawback with such a choice is that it will require additional amount of storage. One can also consider the studies relative to the Krylov subspace based method by considering the question of the optimal setting of the initial solution of the inner solver problem or the recycling of a Krylov subspace previously computed [47]. Finally, one may consider the use of a flexible strategy [48], with a preconditioner varying throughout iterations [49]. In addition to the primordial aspect of preconditioning for iterative solver that we have been considered in this paper, one can also consider some High Performance Computing (HPC) implementations. For such an HPC implementation however, the issue of load balancing is crucial, and it requires very careful *a priori* estimates of block-vector multiplications. In particular, it must be known *a priori* whether each block is full or low-rank. This implies some precise work load distribution and is beyond the scope of our manuscript (see [50] and references therein).

Acknowledgment. This work was supported by a public grant as part of the Investissement d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH. The authors are highly grateful to the LMH (Labex Mathématiques Hadamard) for the funding of the post-doctoral research of D. Félix KPADONOU.

- [1] M. Bonnet, Boundary integral equation methods for solids and fluids, Wiley, 1999.
- [2] G. C. Hsiao, W. L. Wendland, Boundary Integral Equations, Springer Berlin Heidelberg, 2008. doi:10.1007/978-3-540-68545-6.
- [3] S. Kobayashi, N. Nishimura (Eds.), Boundary Element Methods, Springer Berlin Heidelberg, 1992. doi:10.1007/978-3-662-06153-4.
- [4] M. Ganesh, I. G. Graham, A high-order algorithm for obstacle scattering in three dimensions, Journal of Computational Physics 198 (1) (2004) 211–242.
- [5] F. Le Louër, A high order spectral algorithm for elastic obstacle scattering in three dimensions, Journal of Computational Physics 279 (2014) 1–17.
- [6] S. Chaillat, M. Bonnet, Recent advances on the fast multipole accelerated boundary element method for 3D time-harmonic elastodynamics, Wave Motion 50 (7) (2013) 1090–1104.
- [7] L. Greengard, J. Huang, V. Rokhlin, S. Wandzura, Accelerating fast multipole methods for the Helmholtz equation at low frequencies, IEEE Computational Science and Engineering 5 (3) (1998) 32–38.
- [8] 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 (5) (2012) 531–551.
- [9] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, Journal of Computational physics 73 (2) (1987) 325–348.
- [10] V. Rokhlin, Rapid solution of integral equations of classical potential theory, Journal of Computational Physics 60 (2) (1985) 187–207.
- [11] H. Cheng, W. Y. Crutchfield, Z. Gimbutas, L. F. Greengard, J. F. Ethridge, J. Huang, V. Rokhlin, N. Yarvin, J. Zhao, A wideband fast multipole method for the Helmholtz equation in three dimensions, Journal of Computational Physics 216 (1) (2006) 300–325. doi:10.1016/j.jcp.2005.12.001.
- [12] J. Lee, J. Zhang, C.-C. Lu, Incomplete LU preconditioning for large scale dense complex linear systems from electromagnetic wave scattering problems, Journal of Computational Physics 185 (1) (2003) 158–175. doi:10.1016/s0021-9991(02)00052-9.

- [13] K. Sertel, J. L. Volakis, Incomplete LU preconditioner for FMM implementation, *Microwave and Optical Technology Letters* 26 (4) (2000) 265–267. doi:10.1002/1098-2760(20000820)26:4<265::aid-mop18>3.0.co;2-o.
- [14] B. Carpentieri, I. S. Duff, L. Giraud, Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism, *Numerical Linear Algebra with Applications* 7 (7-8) (2000) 667–685. doi:10.1002/1099-1506(200010/12)7:7/8<667::aid-nla218>3.0.co;2-x.
- [15] B. Carpentieri, I. S. Duff, L. Giraud, M. Magolu monga Made, Sparse symmetric preconditioners for dense linear systems in electromagnetism, *Numerical Linear Algebra with Applications* 11 (89) (2004) 753–771. doi:10.1002/nla.345.
- [16] B. Carpentieri, A matrix-free two-grid preconditioner for solving boundary integral equations in electromagnetism, *Computing* 77 (3) (2006) 275–296. doi:10.1007/s00607-006-0161-7.
- [17] M. Bebendorf, Approximation of boundary element matrices, *Numerische Mathematik* 86 (4) (2000) 565–589.
- [18] S. Chaillat, L. Desiderio, 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) 165–186. doi:10.1016/j.jcp.2017.09.013.
- [19] W. Hackbusch, A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices, *Computing* 62 (2) (1999) 89–108.
- [20] W. Hackbusch, B. N. Khoromskij, A sparse matrix arithmetic. Part II: Application to multidimensional problems, *Computing* 64 (2000) 21–47.
- [21] M. Bebendorf, S. Kunis, Recompression techniques for adaptive cross approximation, *Journal of Integral Equations and Applications* 21 (2009) 331–357.
- [22] M. Bebendorf, S. Rjasanow, Adaptive low-rank approximation of collocation matrices, *Computing* 70 (1) (2003) 1–24.
- [23] K. Chen, On a class of preconditioning methods for dense linear systems from boundary elements, *SIAM Journal on Scientific Computing* 20 (2) (1998) 684–698. doi:10.1137/s1064827596304058.
- [24] P. J. Harris, K. Chen, On efficient preconditioners for iterative solution of a Galerkin boundary element equation for the three-dimensional exterior Helmholtz problem, *Journal of Computational and Applied Mathematics* 156 (2) (2003) 303–318. doi:10.1016/s0377-0427(02)00918-4.
- [25] U. Langer, D. Pusch, S. Reitzinger, Efficient preconditioners for boundary element matrices based on grey-box algebraic multigrid methods, *International Journal for Numerical Methods in Engineering* 58 (13) (2003) 1937–1953. doi:10.1002/nme.839.
- [26] S. A. Vavasis, Preconditioning for boundary integral equations, *SIAM Journal on Matrix Analysis and Applications* 13 (3) (1992) 905–925. doi:10.1137/0613055.
- [27] S. Amini, N. D. Maines, Preconditioned Krylov subspace methods for boundary element solution of the Helmholtz equation, *International Journal for Numerical Methods in Engineering* 41 (5) (1998) 875–898. doi:10.1002/(sici)1097-0207(19980315)41:5<875::aid-nme313>3.0.co;2-9.
- [28] L. Banjai, W. Hackbusch, Hierarchical matrix techniques for low-and high-frequency Helmholtz problems, *IMA Journal of Numerical Analysis* 28 (1) (2008) 46–79.
- [29] M. Bebendorf, Approximate inverse preconditioning of finite element discretizations of elliptic operators with nonsmooth coefficients, *SIAM Journal on Matrix Analysis and Applications* 27 (4) (2006) 909–929. doi:10.1137/s0895479803437621.

- [30] M. Bebendorf, Hierarchical LU decomposition-based preconditioners for BEM, *Computing* 74 (3) (2004) 225–247. doi:10.1007/s00607-004-0099-6.
- [31] I. Benedetti, M. Aliabadi, G. Davì, A fast 3D dual boundary element method based on hierarchical matrices, *International Journal of Solids and Structures* 45 (7-8) (2008) 2355–2376. doi:10.1016/j.ijsolstr.2007.11.018.
- [32] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on Scientific and Statistical Computing* 7 (3) (1986) 856–869. doi:10.1137/0907058.
- [33] P. Escapil-Inchauspé, C. Jerez-Hanckes, Fast calderón preconditioning for the electric field integral equation, *IEEE Transactions on Antennas and Propagation* 67 (4) (2019) 2555–2564.
- [34] I. Fierro, C. Jerez-Hanckes, Fast calderón preconditioning for helmholtz boundary integral equations, *Journal of Computational Physics* (2020) 109355.
- [35] R. Stevenson, R. van Venetië, Uniform preconditioners for problems of negative order, *Mathematics of Computation* 89 (322) (2020) 645–674.
- [36] R. Stevenson, R. van Venetië, Uniform preconditioners for problems of positive order, *Computers & Mathematics with Applications*.
- [37] S. B. Adrian, F. P. Andriulli, T. F. Eibert, On a refinement-free calderón multiplicative preconditioner for the electric field integral equation, *Journal of Computational Physics* 376 (2019) 1232–1252.
- [38] A. Kleanthous, T. Betcke, D. P. Hewett, M. W. Scroggs, A. J. Baran, Calderón preconditioning of pmchwt boundary integral equations for scattering by multiple absorbing dielectric particles, *Journal of Quantitative Spectroscopy and Radiative Transfer* 224 (2019) 383–395.
- [39] S. Börm, L. Grasedyck, W. Hackbusch, Introduction to hierarchical matrices with applications, *Engineering Analysis with Boundary Elements* 27 (5) (2003) 405–422.
- [40] M. Bebendorf, *Hierarchical matrices*, Springer, 2008.
- [41] G. Golub, C. Van Loan, *Matrix computations*, 3rd Edition, JHU Press, 2012.
- [42] M. Bebendorf, C. Kuske, R. Venn, Wideband nested cross approximation for Helmholtz problems, *Numerische Mathematik* 130 (1) (2015) 1–34.
- [43] L. Desiderio, \mathcal{H} -matrix based solvers for 3D elastodynamic boundary integral equations, Ph.D. thesis, Paris-Saclay University (2017).
- [44] S. Chaillat, S. P. Groth, A. Loseille, Metric-based anisotropic mesh adaptation for 3d acoustic boundary element methods, *Journal of Computational Physics* 372 (2018) 473–499.
- [45] R. B. Lehoucq, D. C. Sorensen, C. Yang, *ARPACK Users Guide: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*.
- [46] V. Frayassé, L. Giraud, S. Gratton, Langou, A set of gmres routines for real and complex arithmetics on high performance computers, Tech. Rep. TR/PA/03/03, CERFACS.
- [47] M. Parks, E. de Sturler, G. Mackey, D. Johnson, S. Maiti, Recycling Krylov Subspaces for Sequences of Linear Systems, *SIAM Journal on Scientific Computing* 28 (5) (2006) 1651–1674. arXiv:https://doi.org/10.1137/040607277, doi:10.1137/040607277. URL https://doi.org/10.1137/040607277
- [48] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, *SIAM Journal on Scientific Computing* 14 (2) (1993) 461–469. doi:10.1137/0914028.

- [49] L. Mariano Carvalho, S. Gratton, R. Lago, X. Vasseur, A flexible generalized conjugate residual method with inner orthogonalization and deflated restarting, *SIAM J. Matrix Analysis Applications* 32 (2011) 1212–1235. doi:10.1137/100786253.
- [50] H. Harbrecht, P. Zaspel, A scalable \mathcal{H} -matrix approach for the solution of boundary integral equations on multi-gpu clusters, arXiv preprint arXiv:1806.11558.