CSMA 2019
14ème Colloque National en Calcul des Structures
13-17 Mai 2019, Presqu'île de Giens (Var)

An efficient preconditioner for $\mathcal{H}$-matrix accelerated Boundary Element Methods for 3D wave propagation problems

D.F. Kpadonou[1], S. Chaillat[1], P. Ciarlet[1]

[1] *Laboratoire POEMS (CNRS-INRIA-ENSTA), UMA-ENSTA*
{felix.kpadonou,stephanie.chaillat,patrick.ciarlet}@ensta-paristech.fr

Abstract — Hierarchical matrix or $\mathcal{H}$-matrix provides a data-sparse representation of system matrices resulting from the discretization of non-local operators. We are concerned with $\mathcal{H}$-matrix based iterative solver in the scope of fast boundary elements methods. We focus specifically on the definition of efficient preconditioners in order to speed up the solution time of such iterative solvers and also to reduce the number of iterations as the problem complexity increases. We consider a nested outer-inner solver preconditioning strategy. The efficiency of the preconditioning will be shown with some numerical tests.
Mots clés — wave propagation, $\mathcal{H}$-BEM, Preconditioning, nested GMRES, outer-inner solver.

1 Introduction

Acoustic and Elastodynamic wave propagation problems can be modeled by Boundary Integral Equations (BIEs) [6]. 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 using the Boundary Element Method (BEM), [15]. The main advantage is that only the domain boundary (a surface) mesh is required. Although this yields, at a discrete level, to a problem with reduced size, the resulting system matrix is fully populated. As a result, the complexities of the computational times and storage costs are prohibitive. Given N, 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 $O(N^2)$. A direct solution, for example, with the LU factorization is of the order of $O(N^3)$. For an iterative solver, the global solution complexity is $O(n_{\text{iter}}\, N^2)$; n_{iter} being the number of iterations. Hence, iterative solvers are more interesting than direct ones provided *a priori* that $n_{\text{iter}} << N$.
However, due to computational and storage complexities, the BEM in its standard form is not usable in practice for problem with large number of degrees of freedom. Recently, BEM solvers have been speeded up with some acceleration techniques yielding to fast BEMs. One well known fast BEM is the Fast Multipole accelerated BEM (FM-BEM) [9], [12]. The FM-BEM [17] allows to compute efficiently the application of the integral operator to any given field. This method is exclusively designed for iterative solvers since it speeds up the matrix-vector product computation. This approach helps to lower 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 [16], SParse Approximative Inverse [8], multi-grid methods [7] as 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 of the underlying physics. Another approach to speed-up the BEM, that we consider, involves a hierarchical representation of the system matrix ($\mathcal{H}$-matrix) [1], [10] and will be referred as $\mathcal{H}$-BEM. The approach originally introduced in [13, 14] relies on an 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) [3, 5]. The advantage of this approach is that the system matrix is available and one is not restricted in the exploration and definition of an efficient preconditioner for the iterative solver. Various preconditioning strategies exist in the litterature, namely the implicit and explicit ones. For the explicit

preconditioner the inverse of the preconditioning operator is explicitely computed and directly applied, while for the implicit one the application of the preconditioner requires the (iterative) solution of a linear system. A key and common requirement of both classes of preconditioners is that the evaluation of the preconditioner has to be computationally cheap. In this regard, the explicit preconditioner are usually approximated using several techniques such as the Sparse Approximate Inverser (SPAI), incomplete LU factorization, etc.
We are thus interested, in this contribution, with the proposition of efficient implicit preconditioners for the iterative solver. The linear system at hand results from a BEM with collocation technique and then the system matrix is typically non hermitian. Therefore we consider iterative solver based on GMRES (General Minimal RESidual) algorithm [19]. We avoid the computation of the inverse of the preconditioning operator and consider implicit preconditioners in conjunction with the Flexible variant of the GMRES algorithm which involves an outer-inner preconditioning (iterative) solver [18]. The first two parts of this submission recalls the boundary integral equations and several aspects of the $\mathcal{H}$-matrix data-sparse representation of BEM for wave propagation. Further, in section 4, we give the general fashion of the outer-inner iterative solver and propose a preconditioner. Finally, in section 5 we illustrate numerically the efficiency of the approach.

2 Boundary Integral represention and equations

We are concerned with 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 $\boldsymbol{u}$ respectively the velocity and displacement fields of the acoustic and elastodynamic problems. Let us denote by $\Omega^- \in \mathbb{R}^3$ the bounded domain representing the obstacle, with a closed Lipschtiz boundary $\Gamma = \partial\overline{\Omega^-}$ and Ω^+ the exterior domain $\mathbb{R}^3\backslash\overline{\Omega^-}$. $\mathbf{n}$ represents the outward unit normal vector field on Γ. The acoustic and elastodynamic equations are respectively given by

$$\Delta u + \kappa^2 u = 0 \tag{1}$$

and

$$\operatorname{div}(\sigma(\boldsymbol{u})) + \rho w^2 \boldsymbol{u} = 0. \tag{2}$$

The stress and strain tensors are respectively given by $\sigma(\boldsymbol{u}) = \lambda(\operatorname{div}\boldsymbol{u})\mathbb{I}_3 + 2\mu\varepsilon(\boldsymbol{u})$ and $\varepsilon(\boldsymbol{u}) = \frac{1}{2}([\nabla\boldsymbol{u}] + [\nabla\boldsymbol{u}]^\intercal)$; where $\mathbb{I}_3$ is the 3-by-3 identity matrix and $[\nabla\boldsymbol{u}]$ is the 3-by-3 matrix whose β-th column is the gradient of the β-th component of $\boldsymbol{u}$, 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 or elastic full-space are respectively given by

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

and

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

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

$$\boldsymbol{T} = 2\mu\frac{\partial}{\partial\boldsymbol{n}} + \lambda\boldsymbol{n}\operatorname{div} + \mu\boldsymbol{n}\times\mathbf{curl} \text{ and } T = \frac{\partial}{\partial\boldsymbol{n}}. \tag{5}$$

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

$$\boldsymbol{T}(\boldsymbol{x},\boldsymbol{y},\omega) = \boldsymbol{T}_{\boldsymbol{y}}\boldsymbol{G}(\boldsymbol{x},\boldsymbol{y};\omega) \text{ and } T(\boldsymbol{x},\boldsymbol{y},\kappa) = T_{\boldsymbol{y}}G(\boldsymbol{x},\boldsymbol{y};\kappa).$$

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

Boundary Integral Representation. The classical integral representation is given by

$$u(\boldsymbol{x}) = \mathcal{D}u_{|\Gamma}(\boldsymbol{x}) - \mathcal{S}t_{|\Gamma}(\boldsymbol{x}),\ \boldsymbol{x} \in \mathbb{R}^3 \backslash \Gamma,$$

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

$$\mathcal{S}\varphi(\boldsymbol{x}) = \int_\Gamma G(\boldsymbol{x},\boldsymbol{y};\kappa)\varphi(\boldsymbol{y})ds(\boldsymbol{y}) \text{ and } \mathcal{D}\psi(\boldsymbol{x}) = \int_\Gamma [T_{\boldsymbol{y}} G(\boldsymbol{x},\boldsymbol{y};\kappa)]^\intercal \psi(\boldsymbol{y})ds(\boldsymbol{y}).$$

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

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

Boundary Integral Equations. Scattering problems can be formulated as BIEs as follows: 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

$$\mathcal{S}(t_{|\Gamma} + t^{inc}_{|\Gamma})(\boldsymbol{x}) = u^{inc}_{|\Gamma}(\boldsymbol{x}),\ \boldsymbol{x} \in \Gamma. \tag{6}$$

Similarly, for the elastodynamics problem, it is: Find $\boldsymbol{t}_{|\Gamma} \in \boldsymbol{H}^{-\frac{1}{2}}(\Gamma)$ such that

$$\boldsymbol{\mathcal{S}}(\boldsymbol{t}_{|\Gamma} + \boldsymbol{t}^{inc}_{|\Gamma})(\boldsymbol{x}) = \boldsymbol{u}^{inc}_{|\Gamma}(\boldsymbol{x}),\ \boldsymbol{x} \in \Gamma. \tag{7}$$

Discretization of the BIE. At a discrete level, one deals with a linear system resulting from the discretization of the BIE with the BEM. 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 [6]. Let then introduce the sets of points $\mathrm{X} = (\boldsymbol{x}_i)_{i=1:N_c}$ and $\mathrm{Y} = (\boldsymbol{y}_j)_{j=1:N}$, respectively the collocation points and the points associated to the degrees of freedom of the mesh $\Gamma_h \subset \Gamma$ of the domain boundary. For the acoustic scattering problem it reads

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

h is the mesh parameter and S_h is the discretization of the single layer potential; for $\Gamma_h := \bigcup_{l=1:N_E} E_l$

$$S_h\varphi(x) = \sum_{l=1}^{N_E} \int_{E_l} (\cdots) ds(\boldsymbol{y}).$$

We denote by $\mathbb{A} \in \mathbb{R}^{N_c \times N}$ and $\mathbf{b} \in \mathbb{R}^{N_c}$ respectively the system matrix (also called BEM matrix) and the right hand side associated to the incident wave. We are interested by the solution through an iterative solver of the linear system

$$\mathbb{A}\mathbf{x} = \mathbf{b}; \tag{8}$$

$\mathbf{x} \in \mathbb{R}^N$ being the vector of unknown degrees of freedom. 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

We are interested with an alternative (data-sparse) representation of the system matrix, which will be 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 $\|\cdot\|$

$$\|\mathbb{A} - \mathbb{A}_{\eta,\varepsilon}\| \leq \varepsilon \|\mathbb{A}\|.$$

Thus, instead of (8), one consider the problem (9)

$$\mathbb{A}_{\eta,\varepsilon}\mathbf{x}_{\mathcal{H}} = \mathbf{b}. \tag{9}$$

It is worth noting that the fidelity of $\mathbf{x}_{\mathcal{H}}$ with respect to $\mathbf{x}$, solution of (8), is guarantee [10]. Indeed, it is shown that, the error $\|\mathbb{A}\mathbf{x}_{\mathcal{H}} - \mathbf{b}\|$ is controlled by the sum of the resolution or stability error $\|\mathbb{A}_{\eta,\varepsilon}\mathbf{x}_{\mathcal{H}} - \mathbf{b}\|$ and an additional error taking into account for the influence of the accuracy parameter ε (the quality of approximation of $\mathbb{A}_{\eta,\varepsilon}$).

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 exhibits some blocks which are low-rank. Let us introduce $\sigma \subset \{1,\cdots,N_c\}$ and $\tau \subset \{1,\cdots,N\}$ two sets of indexes corresponding to the clusters of nodes $\mathrm{X}_\sigma = (\boldsymbol{x}_i)_{i\in\sigma} \subset \mathrm{X}$ and $\mathrm{Y}_\tau = (\boldsymbol{y}_j)_{j\in\tau} \subset \mathrm{Y}$. We denote by $\mathbb{A}_{\sigma\times\tau}$ the block of $\mathbb{A}$ restricted to the row and column indexes corresponding to the interaction between the clusters of nodes X_σ and Y_τ. Basically, when $\mathrm{X} = \mathrm{Y}$, singularities (therefore full rank blocks) mainly occur at the diagonal and the kernel function is typically smooth everywhere else. For the Laplace kernel this is transcribed by the asymptotically smoothness property [2]. Thus, at the discrete level, some blocks $\mathbb{A}_{\sigma\times\tau}$, can be known and *a priori* identified as of low-rank using the admissibility condition, i.e. $\sigma\times\tau \in \mathcal{P}^{\mathrm{ad}}$ a partition of $\{1,\cdots,N_c\}\times\{1,\cdots,N\}$. 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

$$\min(\mathrm{diam}(\mathrm{X}_\sigma), \mathrm{diam}(\mathrm{Y}_\tau)) \leq \eta\, \mathrm{dist}(\mathrm{X}_\sigma, \mathrm{Y}_\tau);$$

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

For the 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 there exists a pre-asymptotic regime [10] where the maximum numerical rank among the admissible blocks increases linearly with the frequency.

The η-admissibility condition is used to define the "data-sparsity pattern" corresponding to a given partitioning of $\mathbb{A}$. The partitioning of the matrix is based on the clustering of the nodes of X and Y. Indeed the rows or columns of a given block must correspond to the indexes of nodes or degrees of freedom interacting at close range. . As stopping criterion: a given block matrix is recursively subdivised in a 2×2 subblock matrices until this block is either η-*admissible* or $\min(|\sigma|,|\tau|) \leq n_{\mathrm{leaf}}$ a prescribed (minimum) number of points of the clusters.

Low-rank approximation (data-sparse representation). Once the *admissible* blocks are determined, an accurate rank-revealing algorithm is applied to determine the corresponding low-rank approximations. The truncated SVD [11] gives the best low-rank approximation for unitary invariant norms, its computation is expensive and requires the complete storage of the matrix. The adaptive cross approximation (ACA) [1,4,5] offers an interesting alternative to the SVD technique. Reader may refer to [3] for a deep insight on the matter and specifically to [10] for the extension of the low-rank approximation to vector value problem using the ACA.

4 Nested GMRES based preconditioned solver

Let $\mathbb{M}$ and $\mathbf{b}$ respectively be a generic matrix and right hand side vector. The solution $\mathbf{x}$ of a system $\mathbb{M}\mathbf{x} = \mathbf{b}$, through an iterative method, 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 $\mathbf{x}_k$

which minimize the norm of the residual $\mathbf{r}_k = \mathbb{M}\mathbf{x}_k - \mathbf{b}$ at each iteration over an appropriate subspace. Using a preconditioner, one is generally interested by solving an alternative system which is expected to have some better spectral properties, i.e. better conditioning or eigenvalues clustering in comparison to the original system.
Let us denote by $\mathbb{P}$ the preconditioning operator. For a right preconditioning, one solves instead

Preconditioning strategy $\mathbb{M}\mathbb{P}^{-1}\mathbb{P}\mathbf{x} = \mathbf{b} \Leftrightarrow \mathbb{M}\mathbb{P}^{-1}\mathbf{y} = \mathbf{b}$ with $\mathbb{P}\mathbf{x} = \mathbf{y}$. (10)

As can be seen in (10), the application of the inverse of $\mathbb{P}$ is required throughout the iterative solver. Instead of consider a "classical" approach, where a direct solver or explicit preconditioner is used, we consider that the preconditioning system is solved using an iterative solver. This results in a two-level iterative method. The solver at the first level is called the outer solver. It involves the original system matrix $\mathbb{M}$ and the corresponding right hand side vector $\mathbf{b}$. The second solver is called the inner solver and it involves the preconditioner $\mathbb{P}$ of the outer problem, hereafter denoted $\mathbb{P}_{\text{out}}$. Of course, one can still assume that the inner solver is right preconditioned with an operator $\mathbb{P}_{\text{in}}$. Finally, one will consider the procedure described in Figure 1:

Two-level iterative solver (right preconditioning)
Input arguments: $(\mathbb{M}, \mathbf{b}, \mathbb{P}_{\text{out}}, \mathbb{P}_{\text{in}})$
Outer (main) solver: linear system $\mathbb{M}\mathbf{x} = \mathbf{b}$

⋮

– **Inner solver: preconditioning system** $\mathbb{P}_{\text{out}}\mathbf{z}_{k_{\text{out}}} = \mathbf{v}_{k_{\text{out}}}$

⋮

· ***Preconditioning system*** $\mathbb{P}_{\text{in}}\mathbf{z}_{k_{\text{in}}} = \mathbf{v}_{k_{\text{in}}}$ *(e.g. with* ***direct method****)*

⋮

⋮

⋮

Figure 1: Nested outer-inner iterative solver.

In the above procedure, k_{out} is the number of outer iteration and k_{in} is the number of inner iteration at the $k_{\text{out}}th$ outer iteration. A direct method can be used for the preconditioning of the inner solver, i.e the operator $\mathbb{P}_{\text{in}}^{-1}$ explicitly computed.

Preconditioners tailored for an $\mathcal{H}$-matrix based solver

In our scope of $\mathcal{H}$-matrix based BEM, with an $\mathcal{H}$-matrix representation of the system matrix; the complete system matrix is available, thus one is not restricted for the definition of the preconditioners, contrarily to a FM-BEM. We are interested with the definition of the outer preconditioners $\mathbb{P}_{\text{out}}$. We consider that the inner solver is not preconditioned, i.e. $\mathbb{P}_{\text{in}} = \mathbb{I}$. We denote by n_{out} the total number of outer iterations and n_{in} the total number of iterations in the inner solver.
On the one hand, 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 convergence is achieved after $n_{\text{out}} = 1$ outer iteration. On the other hand, the operations required, at each iteration, in the nested outer-inner iterative solver are the application of the system matrix $\mathbb{M}$ and the preconditioning matrix $\mathbb{P}_{\text{out}}$ to vectors. One will therefore be interested with a choice of $\mathbb{P}_{\text{out}}$ such that the computational cost of the action $\mathbf{w} \leftarrow \mathbb{P}_{\text{out}}\mathbf{z}$ is low (comparatively to $\mathbf{y} \leftarrow \mathbb{M}\mathbf{z}$). In the context of $\mathcal{H}$-BEM, we make the choice: $\mathbb{P}_{\text{out}} = \mathbb{A}_{\eta,\varepsilon_{\text{out}}}$, with an accuracy of compression of the low-rank blocks of the preconditioner $\varepsilon_{\text{out}} \geq \varepsilon$, where ε is the accuracy of compression of the system matrix $\mathbb{M} = \mathbb{A}_{\eta,\varepsilon}$. The idea is to set 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 $O(|\sigma|+|\tau|)\times r(\mu)$; where $r(\mu)$ is the numerical rank of a block matrix. This rank increases with the accuracy μ. The advantage with this choice of preconditioner is that there is no additional time for its computation nor additional storage requirements. Indeed they are already included in the computation of the system matrix $\mathbb{A}_{\eta,\varepsilon}$.

5 Illustration of the preconditioning strategy

We consider the exterior domain scattering problem of a time harmonic acoustic plane wave by a sphere of unit radius in an homogeneous and isotropic medium. We consider the single-layer formulation (6). The discretization is performed using a collocation technique and a Lagrange P_1 interpolation. The accuracy of the $\mathcal{H}$-matrix is set to $\varepsilon = 10^{-8}$. We assume that the number of points per wavelength is constant, i.e. a density $n_\lambda = 10$.

Eigenvalues clustering: influence of the parameter ε_{out}

The efficiency of a preconditioner typically depends on the clustering of the eigenvalues of the preconditioned system. Thus, we look for the influence of the accuracy ε_{out} on the clustering of the eigenvalues of $\mathbb{A}_{\eta,\varepsilon}\mathbb{A}^{-1}_{\eta,\varepsilon_{\text{out}}}$. The subsequent analysis is limited to a mesh with $N = 6040$ degrees of freedom (wavenumber is $\kappa \simeq 13$). Indeed the computational cost of the complete eigenvalue decomposition is very prohibitive. The diameters of the boxes surrounding the eigenvalues are reported on Table 1. We remark that, as expected, the eigenvalues are close to the point $(1,0)$ and the diameter of the box decreases as the accuracy increases. The quality of the preconditioning, in view of the clustering, is better and better for the higher accuracies.

Table 1: *Information on the clustering of the eigenvalues of the preconditioned system matrix: the diameter of the boxes surrounding the eigenvalues and the corresponding center coordinates as the accuracy ε_{out} of the preconditioner decreases.*

ε_{out}	10^{-7}	10^{-5}	10^{-3}	10^{-1}
box center	$(1.00, 1.45\times10^{-7})$	$(1.00, -7.33\times10^{-7})$	$(1.00, -1.25\times10^{-5})$	$(1.01, -4.77\times10^{-3})$
diameter	4.01×10^{-6}	3.30×10^{-5}	2.63×10^{-3}	2.11×10^{-1}

Performances of the different preconditioners

We are now interested by the performances of the right preconditioned outer-inner solver for different choices of preconditioner. In the light of the previous test on the eigenvalues, we consider the performances for the accuracies ε_{out} higher or equal to 10^{-3}, which we compare to those obtained with a GMRES solver. The number of points per wavelength is still 10, and the wavenumber corresponding to the different meshes ranges from 17 to 93. The maximum numbers of outer and inner iterations are set to 2000 while the convergence thresholds of both outer and inner solvers are equal to $\varepsilon_{\text{tol(out)}} = \varepsilon_{\text{tol(in)}} = 10^{-6}$. The total number of iterations of the outer-inner solver is denoted $n_{\text{iter}} := n_{\text{out}}(n_{\text{in}})$.

We report on Table 2 the performances for the different preconditioners and those obtained with the GMRES solver. We remark that at fixed frequency, the number of iterations of the solver decreases as the accuracy increases. Generally, the solver with the accuracy $\varepsilon_{\text{out}} = 10^{-7}$ yields to the best performances. Moreover, convergence is achieved after one outer iteration. This behaviour is similar to the one expected with the best theoretical preconditioner, i.e $\mathbb{P}_{\text{out}} = \mathbb{A}_{\eta,\varepsilon}$.

Table 2: *The number of outer-inner iterations and the solution times of the different solvers for different wavenumbers.*

# dofs	κ	GMRES		10^{-3}		10^{-5}		10^{-7}	
10242	16.64	102	4.59	2(290)	9.56	2(270)	10.15	1(102)	4.78
40962	33.27	131	37.11	2(398)	52.39	2(356)	57.09	1(131)	30.98
61033	40.91	131	72.17	2(311)	62.70	2(313)	81.91	1(131)	50.03
163543	66.61	201	375.95	2(521)	413.74	2(503)	519.41	1(201)	263.98
328606	92.79	929	3933.02	3(4381)	11902.75	2(2451)	7718.07	1(929)	3268.97

For the highest frequency ($\kappa \simeq 93$), the performances of the outer-inner solver increase with the accuracy. But, for the lower frequencies, the solver with $\varepsilon_{\text{out}} = 10^{-3}$ performs better than the one with the accuracy $\varepsilon_{\text{out}} = 10^{-5}$. This behavior may be explained by the fact that, at fixed frequency, the numerical rank and thus CPU time per inner iteration increases with the accuracy - see Figure 2-left (with *semilogy* representation). Indeed, the $\mathcal{H}$-matrix-vector product complexity is of the order of $O(\max(r_{\max}(\varepsilon_{\text{out}}), n_{\text{leaf}}) N \log_2 N)$; $r_{\max}(\varepsilon_{\text{out}})$ being the maximum value of numerical rank among the admissible blocks, for the accuracy ε_{out}. Therefore at comparable number of iterations, the preconditioner of coarsest accuracy yields to the best performances. The efficiency of the preconditioner of highest accuracy comes out when the number of iterations diminishes significantly to overcome the increase of the $\mathcal{H}$-matrix-vector product cost.
The plot drawn on Figure 2-right represents the CPU time per outer iteration as the wavenumber increases, with a *semilogy* representation. Generally, the CPU time per outer iteration is larger than the one of the GMRES solver. Indeed, for the outer-inner solver it corresponds to the CPU time for one preconditioning system solution while for the GMRES solver it corresponds essentially to one $\mathcal{H}$-matrix $\mathbb{A}_{\eta,\varepsilon}$-vector product. On Figure 2-right, we remark that the time per outer iteration of the GMRES solver matches perfectly the theoretical complexity $O(\kappa^{2.5}\log_2(\kappa))$ of the $\mathcal{H}$-matrix-vector product operation. Although the CPU time per outer iteration of the different outer-inner solvers is larger than the one of the GMRES solver, the better performances for the preconditioner of accuracy $\varepsilon_{\text{out}} = 10^{-7}$ steams from the fact that the number of outer iterations is reduced.

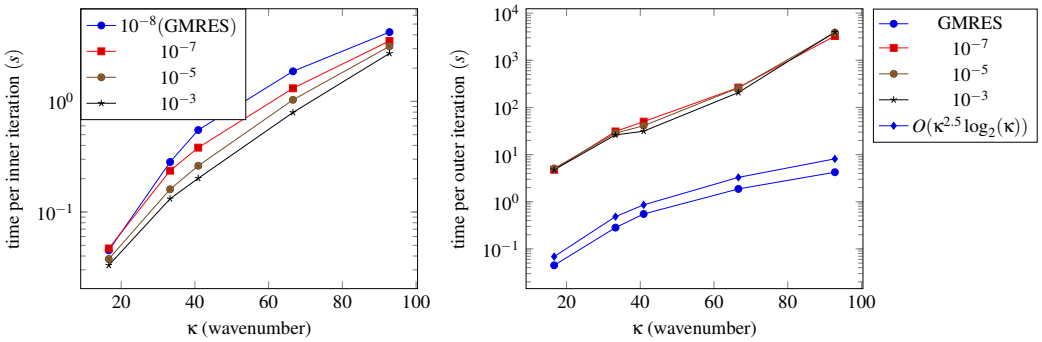

Figure 2: *From left to right, the CPU times for per inner iteration (the matrix vector product) and per outer iteration vs the wavenumber.*

The numerical tests considered for the illustration of the preconditioning strategy are very disadvantageous in the sense that the convergence threshold required in the inner solver is equal to the one of the outer solver (10^{-6}). This certainly explains the poor performances of the preconditioners of coarse accuracies and the fact that the good performances are obtained with the preconditioner of highest accuracy. Indeed, requiring an inner convergence threshold $\varepsilon_{\text{tol(in)}} > \varepsilon_{\text{out}}$ may leads to useless iterations. It is worth noting that the efficiency of the preconditioner is not only defined by the accuracy ε_{out}. Its application also depends on the parameters $\varepsilon_{\text{tol(in)}}$ and the maximum number of inner iterations. For instance, although considering $\mathbb{P}_{\text{out}} = \mathbb{A}_{\eta,\varepsilon}$ (the best theoretical preconditioner) convergence will not certainly be archived throughout one outer iteration if the inner convergence threshold were chosen as follows $\varepsilon_{\text{tol(in)}} = c \times \varepsilon_{\text{tol(out)}}$, $c > 10$.
Several perspectives can be defined for this preliminary work: for instance, the use of a flexible strategy [18] with preconditioners varying throughout iterations and also the investigation of the choice of the "optimal" parameters (convergence thresholds, maximum numbers of iterations) of the outer and inner solvers. The other perspective can be the choice of two different hierarchical structures for the preconditioners and the system matrix (i.e. $\mathbb{P}_{\text{out}} = \mathbb{A}_{\eta_{\text{out}},\eta_{\text{out}}}$, with $\eta_{\text{out}} \neq \eta$), at the price of an additional computational costs and storage requirements for the preconditioners.

6 Conclusion

We propose in this contribution a preconditioner for the $\mathcal{H}$-BEM iterative solvers. The preconditioner is implicit, i.e. the explicit computation of its inverse is not required, since the iterative solver adopted

involves an outer-inner preconditioning strategy. By doing so, the application of the inverse of the preconditioner is made through an (iterative) inner solver. An interesting point is that the preconditioner is naturally available, once the $\mathcal{H}$-matrix representation of the system is computed for an appropriate accuracy. As a result there is no additional time nor memory requirements for their setting up. The emerging trend, from the illustrative example, is that the preconditioner with the highest accuracy yields to the best performances.

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.

References

[1] M. Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86(4):565–589, 2000.

[2] M. Bebendorf. *Hierarchical matrices*. Springer, 2008.

[3] M. Bebendorf and S. Kunis. Recompression techniques for adaptive cross approximation. 2009.

[4] M. Bebendorf, C. Kuske, and R. Venn. Wideband nested cross approximation for Helmholtz problems. *Numerische Mathematik*, 130(1):1–34, 2015.

[5] M. Bebendorf and S. Rjasanow. Adaptive low-rank approximation of collocation matrices. *Computing*, 70(1):1–24, 2003.

[6] M. Bonnet. *Boundary integral equation methods for solids and fluids*. Wiley, 1999.

[7] B. Carpentieri. A matrix-free two-grid preconditioner for solving boundary integral equations in electromagnetism. *Computing*, 77(3):275–296, 2006.

[8] B. Carpentieri, I. S. Duff, L. Giraud, and M. Magolu monga Made. Sparse symmetric preconditioners for dense linear systems in electromagnetism. *Numerical Linear Algebra with Applications*, 11(89):753–771, 2004.

[9] S. Chaillat and M. Bonnet. Recent advances on the fast multipole accelerated boundary element method for 3d time-harmonic elastodynamics. *Wave Motion*, 50(7):1090–1104, 2013.

[10] S. Chaillat, L. Desiderio, and P. Ciarlet. Theory and implementation of h-matrix based iterative and direct solvers for helmholtz and elastodynamic oscillatory kernels. *Journal of Computational Physics*, 351:165–186, 2017.

[11] G.H. Golub and C.F. Van Loan. *Matrix computations*. JHU Press, 3rd edition, 2012.

[12] 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(3):32–38, 1998.

[13] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62(2):89–108, 1999.

[14] W. Hackbusch and B. N. Khoromskij. A sparse matrix arithmetic. Part II: Application to multidimensional problems. *Computing*, 64:21–47, 2000.

[15] S. Kobayashi and N. Nishimura, editors. *Boundary Element Methods*. Springer Berlin Heidelberg, 1992.

[16] J. Lee, J. Zhang, and C.-C. Lu. Incomplete LU preconditioning for large scale dense complex linear systems from electromagnetic wave scattering problems. *Journal of Computational Physics*, 185(1):158–175, feb 2003.

[17] V. Rokhlin. Rapid solution of integral equations of classical potential theory. *Journal of Computational Physics*, 60(2):187–207, 1985.

[18] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing*, 14(2):461–469, mar 1993.

[19] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869.