Sparsification of BEM matrices for large-scale eddy current problems

Piergiorgio Alotto[1], Paolo Bettini[1,2], and Ruben Specogna[3]

[1]University of Padova, Department of Industrial Engineering, 35131 Padova, Italy, piergiorgio.alotto@unipd.it
[2]Consorzio RFX, 35127 Padova, Italy, paolo.bettini@unipd.it
[3]University of Udine, DIEGM, 33100 Udine, Italy, ruben.specogna@uniud.it

Integral formulations can be convenient for computing eddy currents in complicated electromagnetic systems. However, large-scale problems may quickly exceed the memory capacity even of very large machines since the matrices are fully populated. We aim at illustrating how $\mathcal{H}$-matrices with Adaptive Cross Approximation can provide an effective method to increase the size of the largest solvable problems by means of Boundary Element Methods based on *stream functions* with modest implementation effort.

***Index Terms*—Boundary Element Method, Adaptive Cross Approximation, Hierarchical matrix arithmetics, Eddy current problems**

I. Introduction

INTEGRAL formulations can be convenient for computing eddy currents in complicated electromagnetic systems, consisting of many interconnected parts or components, since they do not require the discretisation of non-conducting subdomains. However, the inclusion of fine geometrical details may easily lead to impractical memory and computational time requirements if the problem is not carefully addressed, since integral formulations require the storage of fully populated matrices[1].

Therefore, to be able to solve realistic problems it is necessary to compress these matrices with suitable techniques. Among others, the Fast Multiple Method (FMM) [1] is maybe the most popular one for both low and high frequency problems. A more recent approach for integral operators with asymptotically smooth kernels is based on the Adaptive Cross Approximation (ACA) coupled with hierarchical matrix ($\mathcal{H}$-matrix) arithmetics [2].

In this paper, we aim at illustrating how $\mathcal{H}$-matrices with ACA can provide an effective method to increase the size of the largest solvable problems by means of Boundary Element Methods (BEM) with modest implementation effort. We refer to the formulation implemented in the electromagnetic code `CAFE` [3]. In contrast to other BEM formulations, it is based on *stream functions* [4] and features also non-local equations whose influence on the ACA sparsification, as far as we know, is not yet documented in literature.

II. BEM formulation

The numerical domain is discretised with a polygonal mesh whose incidences are stored in the cell complex $\mathcal{K}$ [5]. The current per unit of thickness 1-cochain $\mathbf{I}$ can be expressed by

$$\mathbf{I} = \mathbf{G}\boldsymbol{\Psi} + \mathbf{H}\mathbf{i}, \tag{1}$$

where $\boldsymbol{\Psi}$ is the 0-cochain whose coefficients are the values of the *stream function* sampled on mesh nodes and $\mathbf{i}$ is an array of *independent currents*; matrix $\mathbf{G}$ stores the edge-node incidences and the columns of $\mathbf{H}$ store the representatives of $H^1(\mathcal{K} - \partial\mathcal{K})$ generators [3] used to treat non-trivial domains.

Then, the discrete Faraday's law is enforced

$$\mathbf{G}^T\tilde{\mathbf{U}} + i\omega\tilde{\boldsymbol{\Phi}} = -i\omega\mathbf{G}^T\tilde{\mathbf{A}}_s, \tag{2}$$

where $\tilde{\mathbf{U}}$ is the electromotive force (e.m.f.) on dual edges, $\tilde{\boldsymbol{\Phi}}$ is the magnetic flux produced only by the eddy currents on dual faces and $\tilde{\mathbf{A}}_s$ is the circulation of the magnetic vector potential due to the source currents on dual edges. The two constitutive laws are expressed in the discrete setting as

$$\tilde{\mathbf{U}} = \mathbf{R}\mathbf{I} \ \text{ and } \ \tilde{\mathbf{A}} = \mathbf{M}\mathbf{I}, \tag{3}$$

where $\mathbf{R}$ and $\mathbf{M}$ are the classical resistance mass matrix and the magnetic matrix [4], respectively.

By substituting (1), (3) and $\tilde{\boldsymbol{\Phi}} = \mathbf{G}^T\tilde{\mathbf{A}}$ inside (2) and taking into account also non-local Faraday's laws, enforced on $H_1(\tilde{\mathcal{K}}) \simeq H^1(\mathcal{K} - \partial\mathcal{K})$ generators, one gets

$$\begin{bmatrix} \mathbf{G}^T\mathbf{K}\mathbf{G} & \mathbf{G}^T\mathbf{K}\mathbf{H} \\ \mathbf{H}^T\mathbf{K}\mathbf{G} & \mathbf{H}^T\mathbf{K}\mathbf{H} \end{bmatrix} \begin{bmatrix} \boldsymbol{\Psi} \\ \mathbf{i} \end{bmatrix} = -i\omega \begin{bmatrix} \mathbf{G}^T\tilde{\mathbf{A}}_s \\ \mathbf{H}^T\tilde{\mathbf{A}}_s \end{bmatrix}. \tag{4}$$

where $\mathbf{K} = \mathbf{R} + i\omega\mathbf{M}$.

III. Sparsification via adaptive cross approximation and $\mathcal{H}$-matrices

ACA can directly use the computational routines of the existing electromagnetic code without any major change. The specific implementation within `CAFE` [3], aiming at the sparsification of the LHS in (4), has been achieved through the *hlibpro* library [6].

In contrast to the FMM, where the kernel is approximated by a sum of spherical multipole functions, ACA generates low-rank approximations of far-field blocks from the entries of the original matrix. The method works as follows. In a first step the degrees of freedom are partitioned and clustered according to a geometrical criterion. Then each cluster pair σ, τ, corresponding to the sub matrix $A_{(\sigma,\tau)}$ is tested against the admissibility criterion $\min\{\mathrm{diam}(\sigma), \mathrm{diam}(\tau)\} \leq \eta\,\mathrm{dist}(\sigma,\tau)$ where $\mathrm{diam}(\sigma)$, $\mathrm{diam}(\tau)$ are the cluster diameters, $\mathrm{dist}(\sigma,\tau)$ is the

[1]The matrix size scales as n^2, n being the number of unknowns, and its inversion has a computational cost of the order of n^3 if a direct solver is used.

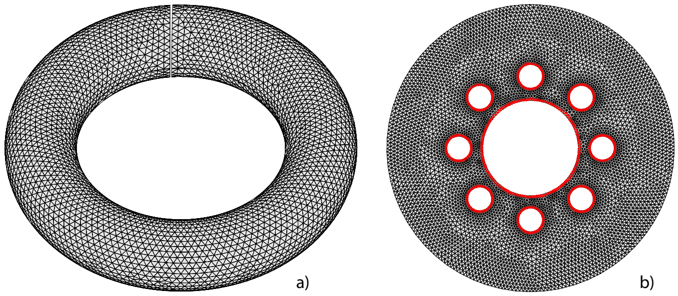

Fig. 1. a) Geometry of the conducting structure of test case 1. b) Geometry of the conducting structure of test case 2. Example of $H^1(\mathcal{K} - \partial\mathcal{K})$ generators arising from holes [3] are shown in red.

distance between the clusters and η the admissibility parameter. If the cluster pair satisfies the criterion, the corresponding matrix block belongs to the far-field, otherwise the clusters are halved and the procedure is applied recursively until the number of elements is larger than a specified threshold. The matrix blocks are then stored with a hierarchical $\mathcal{H}$-matrix structure. The near-field submatrices are calculated exactly, whereas the far-field interactions are approximated with the ACA technique. Consider matrix block $\mathbf{W} \in \mathbb{R}^{m \times n}$ that represent a far-field interaction. In principle, if the singular value decomposition is applied to $\mathbf{W}$, only a few singular values are needed to represent the matrix, obtaining the low rank approximation $\tilde{\mathbf{W}}_k$:

$$\|\tilde{\mathbf{W}} - \tilde{\mathbf{W}}_k\|_F \leq \epsilon \|\tilde{\mathbf{W}}\|_F \tag{5}$$

where $k < m, n$ is the number of singular values used to represent $\tilde{\mathbf{M}}$, ϵ a specified accuracy and $\| \bullet \|_F$ the Frobenius norm. The low rank approximation can be obtained in a smarter way [7] without the construction of $\tilde{\mathbf{W}}$ choosing a subset of rows and columns, forming a *cross*, of the matrix such that

$$\tilde{\mathbf{W}}_k = \mathbf{U}\mathbf{V}^{\mathrm{T}}, \quad \mathbf{U} \in \mathbb{R}^{m \times k}, \mathbf{V} \in \mathbb{R}^{n \times k} \tag{6}$$

Since only a few entries of the original matrix must be computed, it can be proved that the computational cost, as well as the memory, of the matrix partitioning and the ACA approximation have linear-logarithmic complexity [8]. $\mathcal{H}$-matrix arithmetics can be used to define the inversion operator or the LU decomposition of the $\mathcal{H}-$matrix. When Krylov-type solver are used, it is possible to choose a suitable $k' < k$ to perform an approximate LU decomposition to be used as preconditioner.

IV. Results

Two test cases are considered to validate the implementation.

1) Trivial domain: a metallic torus with two gaps (Fig. 1a), discretized with 5457 triangles, 8290 edges, 2834 nodes, subject to a uniform vertical magnetic field imposed by proper BCs (sinusoidal $\tilde{\mathbf{A}}_s$ at 100Hz).
2) Non-trivial domain: a thin 3D metallic shield with 9 holes (Fig. 1b), discretized with 11170 triangles, 17113 edges, 5935 nodes, subject to the field produced by a circular loop (AC current, 50Hz) placed before it.

Tables I-II show the level of compression that can be achieved as a function of the threshold ϵ which is used both in the sparsification of the matrix as well as in the construction of the LU preconditioner. The full matrix size would have been 105.14MBytes and 417.05MBytes, respectively. $Iter$ indicates the number of iterations of the preconditioned GMRES solver.

Note that the system matrices can be compressed to about 1/4 of the original size without significant errors in the solution. Further reductions may be possible but require different settings for the thresholds used in the sparsification of the system matrix and the preconditioner.

Similar trends are also obtained for the large-scale eddy current problems which will be presented in the extended version of paper.

The pseudo-code of the algorithm that performs the sparsification of the LHS in (4) will be documented in the full paper for the most general case (i.e. including non-trivial topologies). A detailed analysis of the error of the solution (approximated vs uncompressed) as a function of the threshold ϵ will also be presented with different metrics for local and global quantities. Another topic which will be addressed in the extended paper will be the influence of minimal and non-minimal cohomology generators on the obtainable sparsification. Finally, the issue of introducing some sparsification also in the off-diagonal blocks of (4) will also be addressed.

TABLE I
Compression as a function of threshold ϵ for test case 1

ϵ	Size [MB]	% Compr. ratio	Iter
1E-06	58.07	55.23%	1
1E-05	50.33	47.86%	2
1E-04	42.19	40.13%	2
1E-03	32.72	31.12%	3
1E-02	23.89	22.72%	4
1E-01	16.25	15.46%	5

TABLE II
Compression as a function of threshold ϵ for test case 2

ϵ	Size [MB]	% Compr. ratio	Iter
1E-06	408.43	97.93%	2
1E-05	400.74	96.09%	2
1E-04	383.39	91.93%	3
1E-03	332.05	79.62%	4
1E-02	217.76	52.21%	4
1E-01	114.68	27.50%	5

References

[1] L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *J. of Comp. Pys.*, vol. 73, no. 1, pp. 325–348, 1987.
[2] W. Hackbusch, “A sparse matrix arithmetic based on $\mathcal{H}$-matrices. part i: Introduction to $\mathcal{H}$-matrices,” *Computing*, vol. 62, pp. 89–108, 1999.
[3] P. Bettini, R. Specogna, “A boundary integral method for computing eddy currents in thin conductors of arbitrary topology,” *IEEE Trans. Magn.*, 10.1109/TMAG.2014.2347894
[4] A. Kameari, “Transient eddy current analysis on thin conductors with arbitrary connections and shapes,” J. Comp. Phys., Vol. 42, pp. 124–140, 1981
[5] K. Itō, *Encyclopedic Dictionary of Mathematics*, 2nd ed., MIT Press, Cambridge, MA, 1987.
[6] http://www.hlibpro.com
[7] S. Kurz, O. Rain, and S. Rjasanow, “The adaptive cross-approximation technique for the 3-D boundary-element method,” *IEEE Trans. Magn.*, vol. 38, no. 2, pp. 421–424, 2002.
[8] M. Bebendorf, *Hierarchical Matrices*, ser. Lecture Notes in Computational Science and Engineering, Springer Verlag, 2008.