Sparsification of BEM Matrices for Large-Scale Eddy Current Problems

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

[1]Department of Industrial Engineering, University of Padova, Padua 35131, Italy
[2]Consorzio RFX, Padua 35127, Italy
[3]Dipartimento di Ingegneria Elettrica Gestionale e Meccanica, University of Udine, Udine 33100, Italy

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. The method is first validated on a benchmark problem for which an analytical solution is available and then applied to a complex problem of engineering interest related to the computation of the currents induced in the stabilizing shell of a nuclear fusion device.

***Index Terms*— Adaptive cross approximation (ACA), eddy current problems, hierarchical matrix arithmetics, integral methods.**

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 discretization 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. In this case, the matrix size scales as n^2, where n is the number of unknowns, and its inversion has a computational cost of the order of n^3 for both direct and iterative solvers.

Therefore, to be able to solve realistic problems it is necessary to compress these matrices with suitable techniques. Maybe, among others, the fast multiple method (FMM) [1] is 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 integral formulations with the modest implementation effort. We refer to the formulation implemented in the electromagnetic code `CAFE` [3]. In contrast to other Boundary Element Method (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 the literature. The validity of the method is first demonstrated on a simple benchmark problem for which an analytical solution is available and then applied to a complex problem of engineering interest for which other sparsification approaches have been applied in the past [5]. Here, we focus on the computation of the current induced in the thin copper shell surrounding the plasma in a nuclear fusion device, RFX-mod (see Fig. 1).

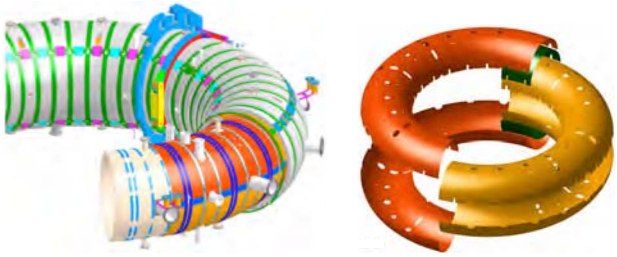

(a) (b)

Fig. 1. (a) RFX-mod, a toroidal magnetic confinement fusion device. (b) Exploded view of RFX-mod stabilizing shell (3 mm copper shell).

II. Integral Formulation

The mesh is formed by N nodes $\{N_i\}_{i=1}^{N}$, E edges $\{E_j\}_{j=1}^{E}$, and P polygons $\{P_k\}_{k=1}^{P}$. The elements of the primal cell complex $\mathcal{K}$ are defined as rectangular prisms $\{v_i\}_{i=1}^{P}$ constructed by considering the polygons with a thickness δ [see Fig. 2(a)]. Faces $\{f_i\}_{i=1}^{E}$ are defined as the lateral faces of the prisms (one-to-one with the mesh edges), and edges $\{e_i\}_{i=1}^{N}$ are those normal to the symmetry plane (one-to-one with the mesh nodes) [see Fig. 2(a)]. Then, the dual nodes $\tilde{n}$, dual edges $\tilde{e}$, and dual faces $\tilde{f}$ belonging to the dual complex $\tilde{\mathcal{K}}$ are constructed with the standard barycentric subdivision [see Fig. 2(b)]. We consider the incidence matrices $\mathbf{C}$ between pairs (f, e) and $\tilde{\mathbf{C}}$ between pairs $(\tilde{f}, \tilde{e})$, in regard to which $\tilde{\mathbf{C}} = \mathbf{C}^T$ holds.

The current 2-cochain $\mathbf{I}$ can be expressed by

$$\mathbf{I} = \mathbf{CT} + \mathbf{Hi} \tag{1}$$

Manuscript received July 3, 2015; revised September 1, 2015; accepted October 2, 2015. Date of publication October 8, 2015; date of current version February 17, 2016. Corresponding author: P. Bettini (e-mail: paolo.bettini@unipd.it).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TMAG.2015.2488699

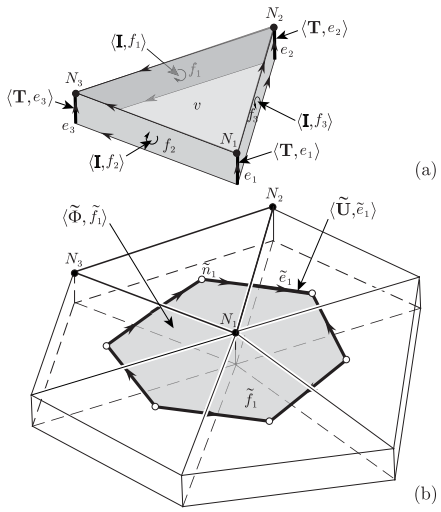

Fig. 2. Association of physical variables to geometric elements of (a) primal and (b) dual complexes.

where $\mathbf{T}$ is the current potential 1-cochain and $\mathbf{i}$ is an array of independent currents; the columns of $\mathbf{H}$ store the representatives of $H^2(\mathcal{K}-\partial\mathcal{K})$ generators [3] used to treat topologically non-trivial domains.

Then, the discrete Faraday's law is enforced

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

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

$$\tilde{\mathbf{U}} = \mathbf{RI} \text{ and } \tilde{\mathbf{A}} = \mathbf{MI} \tag{3}$$

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

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

$$\begin{bmatrix}\mathbf{C}^T\mathbf{KC} & \mathbf{C}^T\mathbf{KH}\\ \mathbf{H}^T\mathbf{KC} & \mathbf{H}^T\mathbf{KH}\end{bmatrix}\begin{bmatrix}\mathbf{T}\\ \mathbf{i}\end{bmatrix} = -i\omega\begin{bmatrix}\mathbf{C}^T\tilde{\mathbf{A}}_s\\ \mathbf{H}^T\tilde{\mathbf{A}}_s\end{bmatrix} \tag{4}$$

where $\mathbf{K} = \mathbf{R} + i\omega\mathbf{M}$. The matrix blocks $\mathbf{A_{ij}}$ are defined as

$$\mathbf{A} = \begin{bmatrix}\mathbf{A_{11}} & \mathbf{A_{12}}\\ \mathbf{A_{21}} & \mathbf{A_{22}}\end{bmatrix} = \begin{bmatrix}\mathbf{C}^T\mathbf{KC} & \mathbf{C}^T\mathbf{KH}\\ \mathbf{H}^T\mathbf{KC} & \mathbf{H}^T\mathbf{KH}\end{bmatrix}. \tag{5}$$

For problems of practical interest the square block $\mathbf{A_{11}}$ has a number of rows which is in the order of 10^3 times larger than the number of rows of the square block $\mathbf{A_{22}}$, with typical values in the order of 10^5 for the former and 10^2 for the latter. For this reason, and for the particular global meaning of the $\mathbf{i}$ unknown, only the square block $\mathbf{A_{11}}$ will be subject to sparsification.

III. Sparsification via Adaptive Cross Approximation and $\mathcal{H}$-Matrices

The most popular approach for the sparsification of matrices arising from the discretization of integral operators is probably FMM. In contrast to this method, in which the kernel is approximated by a sum of spherical multipole functions, the ACA generates low-rank approximations of far-field blocks directly from the entries of the original matrix without any need to modify the existing, robust, well-documented code. In this respect, the method is unobtrusive and extremely programmer friendly. 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 $\min\{\text{diam}(\sigma), \text{diam}(\tau)\} \leq \eta\text{dist}(\sigma, \tau)$, where $\text{diam}(\sigma)$ and $\text{diam}(\tau)$ are the cluster diameters, $\text{dist}(\sigma, \tau)$ is the 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.

In principle, if the singular value decomposition is applied to $\mathbf{W}$ in $\mathbb{R}^{m\times n}$, that represents a far-field interaction, only a few singular values are needed to represent the matrix, obtaining the low-rank approximation $\tilde{\mathbf{W}}_k$

$$\|\tilde{\mathbf{W}} - \tilde{\mathbf{W}}_k\|_F \leq \epsilon\|\tilde{\mathbf{W}}\|_F \tag{6}$$

where $k < m, n$ is the number of singular values used to represent $\tilde{\mathbf{M}}$, ϵ is a specified accuracy, and $\| \bullet \|_F$ is the Frobenius norm. The low-rank approximation can be obtained in a smarter way [6] without the construction of $\tilde{\mathbf{W}}$ choosing a subset of rows and columns, forming a cross, of the matrix such that

$$\tilde{\mathbf{W}}_k = \mathbf{UV}^T, \quad \mathbf{U} \in \mathbb{R}^{m\times k}, \quad \mathbf{V} \in \mathbb{R}^{n\times k}. \tag{7}$$

Since only a few entries of the original matrix are to 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 [7].

Fig. 3 shows the application of $\mathcal{H}$-matrix sparsification with ACA on one of the test cases chosen to demonstrate the usefulness of the method.

$\mathcal{H}$-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 incomplete LU decomposition to be used as preconditioner. The specific implementation of ACA within CAFE [3], aiming at the sparsification of the square block $\mathbf{A_{11}}$ in (4), has been achieved through the *HLIBpro* library [8]. A useful feature of this library is that it natively supports the assembly of a block $\mathcal{H}$-matrix, where the single blocks are $\mathcal{H}$-matrices computed by different routines like in the present case. Furthermore, the library is parallel for both the construction of the matrix and the solution of the resulting linear system of equations.

IV. Results

Two test cases are considered to validate the implementation in the frequency domain.

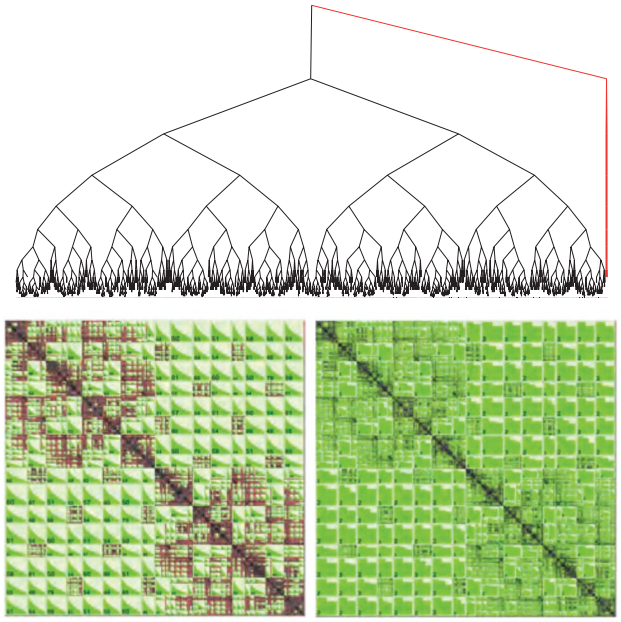

Fig. 3. Top: example of cluster-tree structure for a given matrix A (test case 2). Bottom: low-rank approximations obtained from the same block cluster tree with different values of accuracy. Left: $\epsilon = 10^{-12}$. Right: $\epsilon = 10^{-1}$.

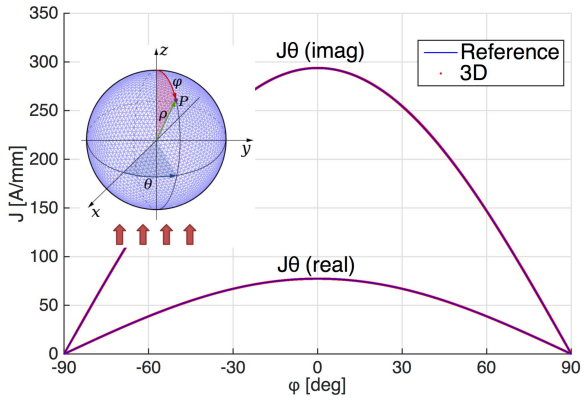

Fig. 4. Trivial case: 3-D numerical results versus reference (analytical) solution.

1) *Trivial Domain (Degrees of Freedoms (DoFs): $\boldsymbol{T}$):* A spherical shell of infinitesimal thickness (radius $a = 1$ m and resistivity $\rho = 1$ mΩ) discretized with 120 832 triangles, 181 248 edges, and 60 418 nodes, and subject to a uniform field [$B = 1$ T and $f = 100$ Hz (see Fig. 4)].
2) *Non-Trivial Domain (DoFs: $\boldsymbol{T}, \boldsymbol{i}$):* A thin toroidal copper shell (major radius $R = 1.995$ m, minor radius $a = 0.5125$ m, and thickness $\delta = 3$ mm) characterized by a poloidal gap, an inner equatorial gap, and 72 holes [circular or oblong port holes with typical area $A = 150$–200 cm^2 (see Fig. 7)], discretized with 181 572 triangles, 274 569 edges, and 92 924 nodes, and subject to a uniform field ($B = 1$ T and $f = 100$ Hz).

As far as the first test case is concerned, in which no generators are present (i.e., the DoFs are only associated with $\mathbf{T}$) and the whole matrix is sparsified, Fig. 4 shows a comparison between the analytical solution and the solution obtained with the presented formulation without sparsification. This was done in order to verify that the chosen discretization was not introducing significant errors. Fig. 5 shows the relative error maps for different values of the ϵ parameter. Table I shows the achievable memory compression ratios and the build time ratios for different values of the ϵ parameter. It can be noted that, in terms of both memory and time, savings of over 95% with respect to the full matrix can be achieved with negligible errors, shown in Fig. 6, for $\epsilon \leq 10^{-3}$.

The second test case, whose solution (3-D current density pattern) is shown in Fig. 7, is even more significant since it highlights the novelty of the presented approach. In this case, due to the holes in the structure, 74 generators are needed (i.e., 74 independent currents $\mathbf{i}$ are added to the

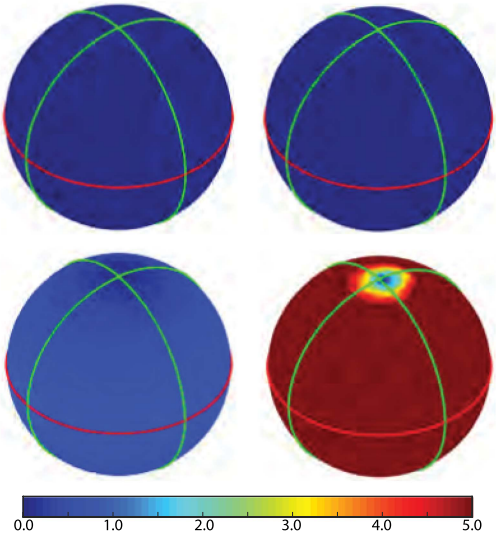

Fig. 5. Trivial case (hollow sphere): error for different values of the accuracy. Clockwise, from top-left: $\epsilon = [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$.

TABLE I
COMPRESSION AS A FUNCTION OF THRESHOLD ϵ FOR TRIVIAL CASE (HOLLOW SPHERE). FULL MATRIX SIZE: 56 Gb. BUILD TIME: 29221 s

ϵ	Build time	% Time ratio	Matrix size	% Memory ratio
1E-06	1098 s	3.76%	2.74 Gb	4.92%
1E-04	794 s	2.72%	1.70 Gb	3.05%
1E-03	643 s	2.2%	1.24 Gb	2.23%
1E-02	433 s	1.48%	0.89 Gb	1.60%
1E-01	317 s	1.09%	0.41 Gb	0.73%

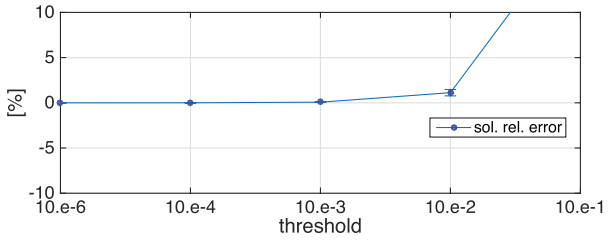

Fig. 6. Error (average and standard deviation) versus accuracy ϵ for test case 1.

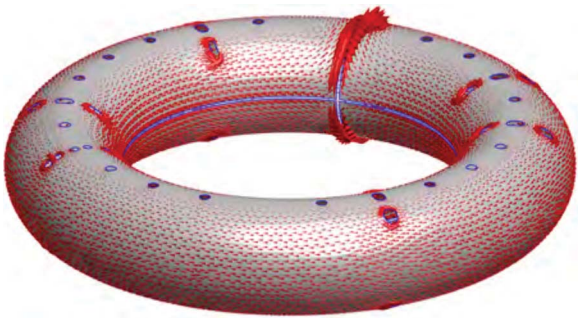

Fig. 7. Non-trivial case (toroidal structure with 2 gaps and 72 holes). 3-D current density pattern: bigger cones higher values.

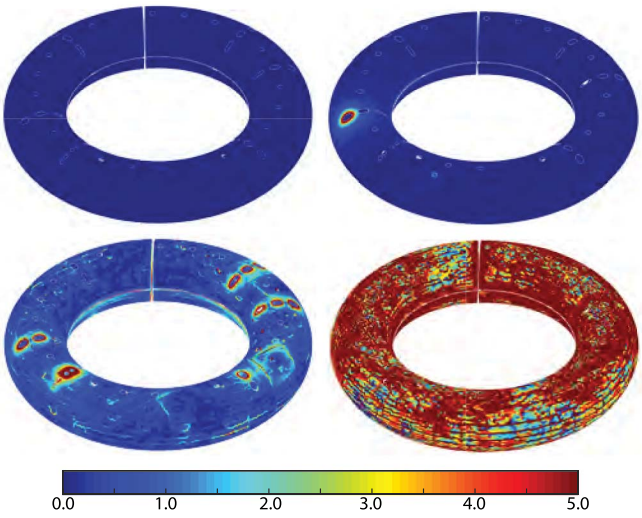

Fig. 8. Non-trivial case: error (current density) for different values of the accuracy. Clockwise, from top-left: $\epsilon = [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$.

TABLE II
COMPRESSION AS A FUNCTION OF THRESHOLD ϵ FOR NON-TRIVIAL CASE (TORUS). FULL MATRIX SIZE: 120 Gb. BUILD TIME: 55046 s

ϵ	Build time	% Time ratio	Matrix size	% Memory ratio
1E-06	2710 s	4.92%	6.43 Gb	5.37%
1E-04	1748 s	3.17%	3.86 Gb	3.22%
1E-03	1370 s	2.49%	2.82 Gb	2.36%
1E-02	1054 s	1.91%	2.02 Gb	1.69%
1E-01	768 s	1.40%	1.27 Gb	1.06%

DoFs associated with $\mathbf{T}$), but only the $\mathbf{A}_{11}$ block is subject to sparsification and the others ($\mathbf{A}_{12}$, $\mathbf{A}_{21}$, and $\mathbf{A}_{22}$) are kept full. Fig. 8 shows the relative error maps for different values of the ϵ parameter. Table II shows the achievable memory compression ratios and the build time ratios for different values of the ϵ parameter. It can be noted that also in this case in terms of both memory and time savings of over 95% with respect to the full matrix can be achieved with negligible errors, shown in Fig. 9, for $\epsilon \leq 10^{-4}$. Savings are slightly smaller than in the previous case partly due to the different structure and partly due to the presence of the unsparsified blocks.

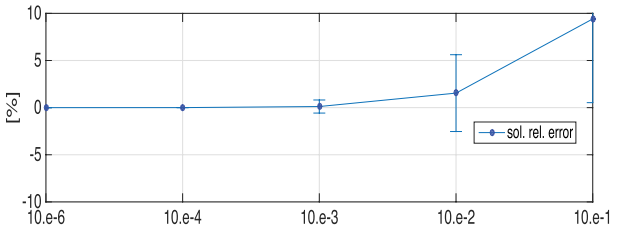

Fig. 9. Error (average and standard deviation) versus accuracy ϵ for test case 2.

V. CONCLUSION

This paper illustrates the usefulness of $\mathcal{H}$-matrices with ACA as an effective method to increase the size of the largest solvable algebraic systems arising from integral formulations. Furthermore, for the first time, the technique is applied to formulations which allow arbitrary topology which results in block matrices with unsparsifiable blocks. The method is first validated on a benchmark problem for which an analytical solution is available and then applied to a complex problem of engineering interest related to the stabilizing shell of a nuclear fusion device. In both cases, with full matrices of sizes in the order of 100 Gb, the savings both in terms of memory and time exceed 95% of the non-sparsified method with negligible loss in solution accuracy.

ACKNOWLEDGMENT

This work was supported in part by the Italian Ministry of Education, Universities and Research within the Projects of National Interest under Grant 2010SPS9B3 and in part by the University of Padua under Grant CPDA144817.

REFERENCES

[1] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.*, vol. 73, no. 2, pp. 325–348, Dec. 1987.
[2] W. Hackbusch, "A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices," *Computing*, vol. 62, no. 2, pp. 89–108, Apr. 1999.
[3] P. Bettini and R. Specogna, "A boundary integral method for computing eddy currents in thin conductors of arbitrary topology," *IEEE Trans. Magn.*, vol. 51, no. 3, Mar. 2015, Art. ID 7203904.
[4] A. Kameari, "Transient eddy current analysis on thin conductors with arbitrary connections and shapes," *J. Comput. Phys.*, vol. 42, no. 1, pp. 124–140, Jul. 1981.
[5] G. Rubinacci, S. Ventre, F. Villone, and Y. Liu, "A fast technique applied to the analysis of resistive wall modes with 3D conducting structures," *J. Comput. Phys.*, vol. 228, no. 5, pp. 1562–1572, Mar. 2009.
[6] S. Kurz, O. Rain, and S. Rjasanow, "The adaptive cross-approximation technique for the 3D boundary-element method," *IEEE Trans. Magn.*, vol. 38, no. 2, pp. 421–424, Mar. 2002.
[7] M. Bebendorf, *Hierarchical Matrices* (Lecture Notes in Computational Science and Engineering). New York, NY, USA: Springer-Verlag, 2008.
[8] R. Kriemann. *HLIBpro*. [Online]. Available: http://www.hlibpro.com, accessed Dec. 1, 2014.