An Adaptive Approximate Inverse-Based Preconditioner Combined with the Fast Multipole Method for Solving Dense Linear Systems in Electromagnetic Scattering

B. Carpentieri

Institute for Mathematics and Scientific Computing, Karl-Franzens University,
A-8010 Graz, Austria
bruno.carpentieri@uni-graz.at

Abstract – We discuss preconditioning strategies for solving large Electric Field Integral Equation systems. We consider several algebraic preconditioners for solving the dense linear system arising from the Galerkin discretization of the pertinent integral equation. We show that approximate inverse methods based on Frobenius-norm minimization techniques can be very effective to reduce the number of iterations of Krylov subspace solvers for this problem class. We describe the implementation of the preconditioner within the Fast Multipole Algorithm and we illustrate how to reduce the construction cost by using static pattern selection strategies. Finally, we present deflating techniques based on low-rank matrix updates to enhance the robustness of the approximate inverse on tough problems. Experiments are reported on the numerical behavior of the proposed method on a set of realistic industrial problems.

Keywords: spectral corrections, electromagnetic scattering applications, and Frobenius-norm minimization method.

I. INTRODUCTION

In this study we consider the scattering problem from a perfectly conducting object Ω with boundary Γ and assume that the domain Ω is illuminated by an incident plane wave $(\vec{E}_{inc}, \vec{H}_{inc})$ of angular frequency $\omega = ck = 2\pi c/\lambda$, where the constant c is the speed of light, k is the wavenumber and $\lambda = c/f$ is the wavelength (f is the frequency). We concentrate our attention on the Electric Field Integral Equation (EFIE) formulation that reads as: find the surface current $\vec{j}$ such that for all tangential test functions $\vec{j}^t$, we have,

$$\int_\Gamma \int_\Gamma G(|y-x|) \left(\vec{j}(x) \cdot \vec{j}^t(y) - \frac{1}{k^2} div_\Gamma \vec{j}(x) \cdot div_\Gamma \vec{j}^t(y) \right) dxdy = \frac{i}{kZ_0} \int_\Gamma \vec{E}_{inc}(x) \cdot \vec{j}^t(x) dx. \quad (1)$$

In equation (1) we denote by $G(|y-x|) = \dfrac{e^{ik|y-x|}}{4\pi|y-x|}$ the Green's function and by $Z_0 = \sqrt{\mu_0/\varepsilon_0}$ the characteristic impedance of vacuum (ϵ is the electric permittivity and μ the magnetic permeability). This formulation is the only one that can be used to model arbitrary geometries, including those with cavities, disconnected parts, breaks on the surface and is the most difficult to solve by iterative methods. However, the solution techniques described in this paper are applicable to other integral formulations as well, such as the Combined Field Integral Equation (CFIE) and the Magnetic Field Integral Equation (MFIE) [1]. The Galerkin discretization of equation (1) leads to dense and complex linear systems of equations,

$$Ax = b \quad (2)$$

whose coefficient matrix A is symmetric for EFIE, nonsymmetric for CFIE and MFIE. Each entry of the coefficient matrix is associated with the interaction of a pair of triangles in the mesh; the entries of the unknown vector x are associated with the vectorial flux across an edge in the mesh, and the right-hand side b depends on the frequency and the direction of the illuminating wave. Although efficient out-of-core direct solvers have been developed for this problem class [2, 3], the huge storage requirement remains the main bottleneck to the viability of integral equation methods for solving high-frequency scattering problems in electromagnetism. The use of iterative methods can solve the memory limits of direct solvers but their success depends much on the underlying integral formulation. The CFIE formulation gives rise to well conditioned systems, and the number of iterations of nonsymmetric Krylov solvers scale as $\mathcal{O}(n^{0.25})$. On EFIE, Krylov methods scale as $\mathcal{O}(n^{0.5})$, thus preconditioning is mandatory to use.

II. SOLUTION TECHNIQUES

The design of robust preconditioners for boundary integral equations can be challenging; many important research papers (see e.g. [4–7]) have addressed this issue in recent years. Simple preconditioners like the diagonal of A, diagonal blocks, or a band can be effective only when the coefficient matrix has some degree of diagonal

dominance depending on the integral formulation. Block diagonal preconditioners are generally more robust than their point-wise counterparts, but may require matrix permutations or renumbering of the grid points to cluster the large entries close to the diagonal.

In Table 1 we report on experiments with various algebraic preconditioners on a sphere of 1 meter length illuminated at 190 MHz, modeled using $EFIE$. The mesh is depicted in Fig. 1. Although the size is small, the problem is representative of realistic electromagnetic scattering calculations and is difficult to solve for many iterative solvers and preconditioners. Incomplete factorizations can be effective for solving nonsymmetric dense systems [5] and hybrid integral formulations [8], but on the EFIE the triangular factors can be very ill-conditioned due to the indefiniteness of A [9]. Although pivoting may help to circumvent numerical instabilities and improve the performance [10], the parallelization may require significant efforts. Approximate inverse methods are generally less prone to instabilities on indefinite systems and they are inherently parallel. Owing to the rapid decay of the discrete Green's function, the location of the large entries in the inverse matrix exhibit some structure, and only a very small number of its entries have large magnitude compared to the others that are much smaller. Several preconditioners of this type have been proposed in electromagnetism (see for instance [4, 7, 11–13]). The approximate inverse can be computed in factorized or unfactorized form, depending on the fact that the preconditioner is expressed as a single matrix or as the product of two (or more) matrices. For a small sphere, we display in Fig. 2 the sparsity pattern of A^{-1} (on the left) and L^{-1}, the inverse of its Cholesky factor (on the right), respectively, where all the entries smaller than 5.0×10^{-2} have been dropped after a symmetric scaling such that $\max_i |a_{ji}| = \max_i |\ell_{ji}| = 1$. The inverse factors L^{-1} can be totally unstructured (see Fig. 2(b)), while entries of A^{-1} decay very rapidly far from the diagonal (see Fig. 2(a)).

Table 1. Number of iterations required by Krylov solvers using various preconditioners to reduce the initial residual by six orders of magnitude on a model problem (see Fig. 1).

Precon	GMRES(50)	Bi-CGSTAB	UQMR	TFQMR
M_j	473	257	354	228
$SSOR$	245	185	281	266
$ILU(0)$	+500	385	394	439
$AINV$	+500	+500	+500	+500
$SPAI$	61	48	93	40

In this work, we describe an algebraic approximate inverse preconditioner based on Frobenius-norm minimization with a static pattern selection strategy for this problem class. The approximate inverse is computed as the matrix M that minimizes the Frobenius-norm of the error matrix $\|I - AM\|_F$, subject to certain sparsity constraints. The Frobenius norm allows the decoupling of the constrained minimization problems into n independent linear least-squares problems, one for each column (resp. row) of M when preconditioning from the right (resp. left). The independence of these least-squares problems follows immediately from the identity,

$$\|I - AM\|_F^2 = \sum_{j=1}^{n} \|e_j - Am_{\bullet j}\|_2^2 \tag{3}$$

where e_j is the jth canonical unit vector and $m_{\bullet j}$ is the column vector representing the jth column of M. Both the construction and the application of M are inherently parallel. The preconditioner is combined with the Multilevel Fast Multipole Algorithm (MLFMA) [14, 15] exploiting the box-wise partitioning of the object available in the MLFMA. MLFMA partitions the mesh of the object by recursive subdivision into disjoint aggregates of small size compared to the wavelength, each roughly formed by an equal number of separate triangles. The number of levels for the box-wise partitioning is determined so that the size of the smallest box is of the order of $\lambda/2$. We skip the description of the MLFMA as it is out of the scope of this paper; details of the parallel implementation we use are found in [16, 17]. The nonzero pattern of the approximate inverse is computed in advance using the sparsity structure of the near-field matrix. More precisely, the structure of the column of the preconditioner associated with a given edge in the mesh is defined by retaining all the edges within the box itself and one level of neighboring boxes [18]. Thus the preconditioner is constructed from a sparse approximation of the dense coefficient matrix and it has a sparse block structure; each block is a dense matrix associated with one box. Indeed the least-squares problems corresponding to edges within the same box are identical because they are defined using the same nonzero structure and the same set of entries of A. It

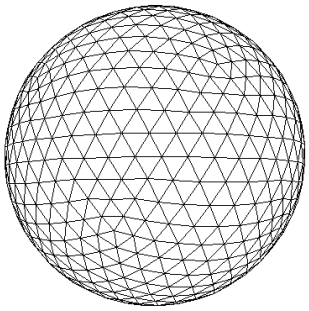

Fig. 1. Model problem, a sphere of 1 meter length, illuminated at 190 MHz. The mesh is discretized with 2430 edges.

means that we only have to compute one QR factorization per box. Blocking the columns enables us to reduce the algorithmic complexity of computing M to $\mathcal{O}(n)$. Parallelism can be exploited by assigning disjoint subsets of boxes to different processors and performing the least-squares solutions independently on each processor. We remark that the preconditioner computed by Frobenius-norm minimization is not guaranteed to be symmetric; we may enforce symmetry in M by reflecting at each step the computed entries with respect to the diagonal, and then solving a reduced least-squares problem to compute the remaining entries below the diagonal.

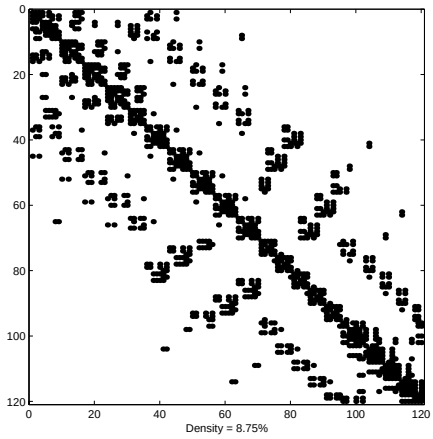

(a) Sparsity pattern of *sparsified*(A^{-1}).

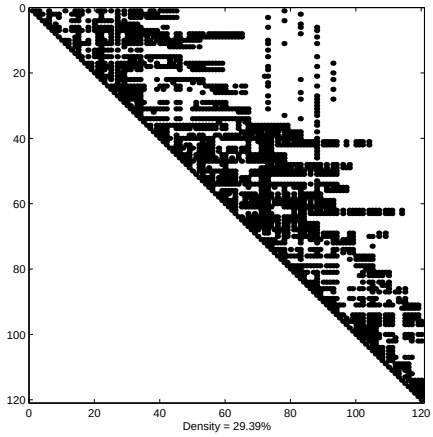

(b) Sparsity pattern of *sparsified*(L^{-1}).

Fig. 2. Sparsity patterns of the inverse of A (on the left) and of the inverse of its lower triangular factor (on the right), where all the entries whose relative magnitude is smaller than 5.0×10^{-2} are dropped. The test problem, representative of the general trend, is a small sphere.

In our numerical experiments, reported in the next section, we observe a lack of robustness of the approximate inverse on tough configurations due to the presence of small eigenvalues that cluster near zero in their natural trajectory towards point one of the spectrum of the preconditioned matrix under the action of the preconditioner. This consideration motivates us to introduce a stabilization step after computing M, which deflates a small group of eigenvalues close to zero in the spectrum of MA. Deflating techniques have proved to be useful to accelerate the convergence of iterative methods for general linear systems (e.g. [19–21]). We consider equation (2) and we denote by M_1 the left preconditioner, meaning that we solve,

$$M_1 A x = M_1 b. \tag{4}$$

We assume that the preconditioned matrix $M_1 A$ is diagonalisable, that is,

$$M_1 A = V \Lambda V^{-1} \tag{5}$$

with $\Lambda = diag(\lambda_i)$, where $|\lambda_1| \leq \ldots \leq |\lambda_n|$ are the eigenvalues and $V = (v_i)$ the associated right eigenvectors. We denote by $U = (u_i)$ the associated left eigenvectors; we then have $U^H V = diag(u_i^H v_i)$, with $u_i^H v_i \neq 0, \forall i$. Let V_ε be the set of right eigenvectors associated with the set of eigenvalues λ_i with $|\lambda_i| \leq \varepsilon$. Similarly, we define by U_ε the corresponding subset of left eigenvectors.

Theorem 1: Let

$$A_c = U_\varepsilon^H M_1 A V_\varepsilon,$$

$$M_c = V_\varepsilon A_c^{-1} U_\varepsilon^H M_1$$

and

$$M = M_1 + M_c.$$

Then MA is diagonalisable and we have $MA = V diag(\eta_i) V^{-1}$ with

$$\begin{cases} \eta_i = \lambda_i & \text{if} \quad |\lambda_i| > \varepsilon, \\ \eta_i = 1 + \lambda_i & \text{if} \quad |\lambda_i| \leq \varepsilon. \end{cases}$$

A_c represents the projection of the matrix $M_1 A$ on the coarse space defined by the approximate eigenvectors associated with its smallest eigenvalues. **Proof**

We first remark that $A_c = diag(\lambda_i u_i^H v_i)$ with $|\lambda_i| \leq \varepsilon$ and so A_c is nonsingular. A_c represents the projection of the matrix $M_1 A$ on the space spanned by the approximate eigenvectors associated with its smallest eigenvalues.
Let $V = (V_\varepsilon, V_{\bar{\varepsilon}})$, where $V_{\bar{\varepsilon}}$ is the set of $(n-k)$ right eigenvectors associated with eigenvalues $|\lambda_i| > \varepsilon$.
Let $D_\varepsilon = diag(\lambda_i)$ with $|\lambda_i| \leq \varepsilon$ and $D_{\bar{\varepsilon}} = diag(\lambda_j)$ with $|\lambda_j| > \varepsilon$.
The following relations hold: $MAV_\varepsilon = V_\varepsilon(D_\varepsilon + I_k)$ where I_k denotes the $(k \times k)$ identity matrix, and $MAV_{\bar{\varepsilon}} = V_{\bar{\varepsilon}} D_{\bar{\varepsilon}}$ since $U_\varepsilon^H V_{\bar{\varepsilon}} = 0$; then we have

$$MAV = V \begin{pmatrix} D_\varepsilon + I_k & 0 \\ 0 & D_{\bar{\varepsilon}} \end{pmatrix}.$$

Theorem 2: Let W be such that

$$\tilde{A}_c = W^H A V_\varepsilon \text{ has full rank,}$$
$$\tilde{M}_c = V_\varepsilon \tilde{A}_c^{-1} W^H$$

and

$$\tilde{M} = M_1 + \tilde{M}_c.$$

Then $\tilde{M}A$ is similar to a matrix whose eigenvalues are

$$\begin{cases} \eta_i = \lambda_i & \text{if} \quad |\lambda_i| > \varepsilon, \\ \eta_i = 1 + \lambda_i & \text{if} \quad |\lambda_i| \leq \varepsilon. \end{cases}$$

Proof

With the same notation as for Proposition 1 we have: $\tilde{M}AV_\varepsilon = V_\varepsilon(D_\varepsilon + I_k)$ and, $\tilde{M}AV_{\bar{\varepsilon}} = V_{\bar{\varepsilon}} D_{\bar{\varepsilon}} + V_\varepsilon C$ with $C = A_c^{-1} W^H A V_{\bar{\varepsilon}}$; then we have

$$\tilde{M}AV = V \begin{pmatrix} D_\varepsilon + I_k & C \\ 0 & D_{\bar{\varepsilon}} \end{pmatrix}.$$

For right preconditioning, that is $AM_1 y = b$, similar results hold. We should point out that in the nonsymmetric case a natural choice exists for the operator W, i.e. to select $W = V_\varepsilon$, that saves the computation of left eigenvectors. These formulations enable us to move to one any set of eigenvalues lying in any particular region of the spectrum; if for some particular applications some eigenvalues different from the smallest ones perturb the convergence they can be removed by the same technique. The application of the correction update at each iteration step costs $2nk + k^2$, where k is the size of the coarse space. The novelty of this study with respect to that conducted in [19] is to use MLFMA for computing approximations to the smallest eigenvalues and their corresponding approximate eigenvectors. We use the Implicitly Restarted Arnoldi Method implemented in the ARPACK package [22] that only requires matrix-vector products for the spectral computation. Thus the resulting preconditioner is nearly matrix-free.

III. PERFORMANCE ANALYSIS

In Table 2 we report on the results of the approximate inverse (referred to as $SPAI$) of an experiment on the Cobra problem (Fig. 3(a), $n = 60695$) and on the Almond problem (Fig. 3(b), $n = 104793$), two standard test cases in the electromagnetic community. For both geometries, the scattering problem is modeled using the EFIE and the integral equation is discretized by the Galerkin method. We use the Fast Multipole Method for computing approximate matrix-vector products; we refer to [17] for the numerical implementation of the multipole code. The experiment are run in single precision on eight processors of a Compaq Alpha server. The Compaq Alpha server is a cluster of symmetric multiprocessors. Each node consists of four DEC Alpha processors (EV 6, 1.3 GFlops peak) that share 512 MB of memory. We observe the favorable numerical scalability of the preconditioner for reasonably large value of the restart parameter and the $\mathcal{O}(n)$ complexity of computing M thanks to the blocking strategy.

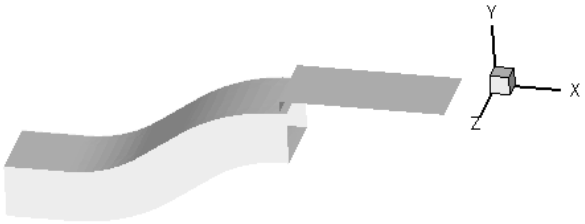

(a) The Cobra problem. It represents an air intake and has size 67.9 cm × 23.3 cm × 11 cm.

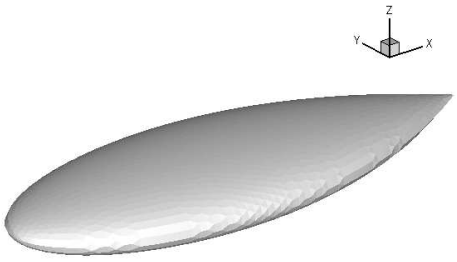

(b) The Almond problem. The size is 2.5 m.

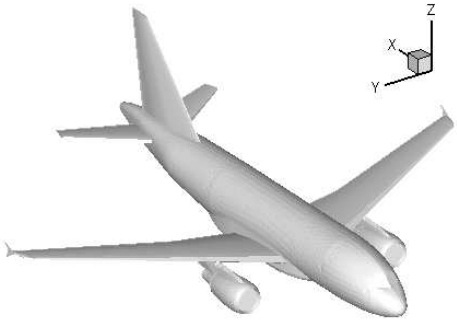

(c) The aircraft problem: an industrial civil aircraft from a European company. It represents a real-life model problem in an industrial context.

Fig. 3. Geometries considered for the numerical experiments. Courtesy of EADS-CCR Toulouse.

In Table 3, we show the parallel scalability of the implementation of the preconditioner in the FMM code on an industrial civil aircraft from a European company, a real-life model problem in an industrial context (the mesh is depicted in Fig. 3(c)). We solve systems of increasing size on a larger number of processors, keeping the number of unknowns per processor constant. It can be observed the very good parallel scalability of the construction and of the application of the preconditioner

typical of approximate inverse methods; for the matrix-vector product operation, the $nlogn$ factor appears in the results.

Table 2. Number of iterations and elapsed time required to reduce the initial residual by 10^{-3} on 8 processors of the Compaq machine, except those marked with $^{(k)}$, that were run on k processors.

Almond

Size	Density SPAI	Time SPAI	GMRES(∞)		GMRES(120)	
			Iter	Time	Iter	Time
104793	0.19	6m	234	20m	253	17m
419172	0.05	21m	413	2h 44m	571	2h 26m
943137	0.02	49m	454	3h 35m$^{(32)}$	589	5h 55m

Cobra

Size	Density SPAI	Time SPAI	GMRES(∞)		GMRES(120)	
			Iter	Time	Iter	Time
60695	0.24	2m	369	26m	516	23m
179460	0.09	7m	353	1h 11h	406	1h 2m

Table 3. Parallel scalability of the aircraft problem (see Figure 3(c)).

Problem size	Nb procs	Construction time (sec)	Elapsed time precond (sec)	Elapsed time mat-vec (sec)
112908	8	513	0.39	1.77
221952	16	497	0.43	2.15
451632	32	509	0.48	2.80
900912	64	514	0.60	3.80

Finally, in Tables 4 to 6 we analyze the effect of using low-rank deflation techniques on the robustness of the iterative method. We study initially the numerical behavior of the preconditioner on a set of small test problems that are representative of realistic scattering calculations in electromagnetism. The model problems are:

Example 1: a cylinder with a hollow inside (110 MHz, $n = 1080$);
Example 2: a cylinder with a break on the surface (60 MHz,$n = 1299$);
Example 3: a satellite (220 MHz, $n = 1701$);
Example 4: a parallelepiped (420 MHz, $n = 2016$); and
Example 5: a sphere (190 MHz, $n = 2430$).

For physical consistency, we have set the frequency of the incident wave so that there are about ten discretization points per wavelength. In each case, we take as initial guess $x_0 = 0$, and the right-hand side is such that the exact solution of the system is known. We consider the formulation described in Theorem 2 and we apply the spectral updates on top of the preconditioned system $AM_1y = b$. In Table 4 we show the number of iterations required by GMRES to obtain convergence for increasing size of the coarse space up to 20. The

Table 4. Number of iterations required by GMRES and SQMR preconditioned by a Frobenius-norm minimization method updated with spectral corrections to reduce the normwise backward error by 10^{-8} for increasing size of the coarse space.

Size of the coarse space	GMRES(m), Toler. 1e-8					SQMR
	m=10	m=30	m=50	m=80	m=110	
Example 1						
Unprec.	+1500	+1500	+1500	651	423	271
0	358	213	144	79	79	103
4	313	169	109	68	68	78
8	294	138	76	58	58	60
12	190	96	52	51	51	52
16	184	80	47	47	47	40
20	174	61	44	44	44	44
Example 2						
Unprec.	+1500	+1500	+1500	+1500	+1500	439
0	+1500	+1500	496	311	198	161
4	279	192	152	125	93	117
8	188	147	129	90	84	97
12	196	148	131	91	83	82
16	183	137	114	74	74	73
20	168	130	100	69	69	68
Example 3						
Unprec.	+1500	+1500	+1500	1404	1193	519
0	268	174	130	79	79	92
4	259	150	99	66	66	77
8	225	109	77	58	58	66
12	117	81	56	52	52	56
16	105	74	49	49	49	48
20	96	58	44	44	44	46
Example 4						
Unprec.	1100	566	434	309	262	185
0	145	113	90	71	71	61
4	125	97	74	61	61	54
8	101	78	58	56	56	49
12	86	70	52	51	51	42
16	81	64	49	49	49	41
20	77	62	47	47	47	39
Example 5						
Unprec.	1241	374	277	216	208	140
0	297	87	75	66	66	51
4	345	66	64	58	58	40
8	55	43	40	40	40	33
12	52	43	38	38	38	34
16	52	44	39	39	39	34
20	53	45	40	40	40	34

numerical experiments are performed in double precision complex arithmetic on a SGI Origin 2000 and the number of iterations are for right preconditioning. We observe that the linear systems are difficult to solve as GMRES does not converge or converges slowly with no preconditioner. We can see that the introduction of the low-rank updates can remarkably accelerate the iterative solution. By selecting up to 10 eigenpairs the number of iterations decreases by at least a factor of two on most of the reported experiments and convergence becomes nearly independent from the restart parameter. On Example 2,

the preconditioning updates enable fast convergence of GMRES with a low restart whereas no convergence was obtained in 1500 iterations without updates.

Table 5. Number of amortization vectors required by the IRAM algorithm to compute approximate eigenvalues nearest zero and the corresponding right eigenvectors. The computation of the amortization vectors is relative to GMRES(10) and a tolerance of 10^{-8}.

Size of the coarse space	Number of Amortization Vectors				
	Ex. 1	Ex. 2	Ex. 3	Ex. 4	Ex. 5
4	7	1	28	9	-
8	4	1	4	6	1
12	2	1	2	4	1
16	2	1	2	7	1
20	4	1	2	6	1

Table 6. Experiments using a sparse approximate inverse preconditioner and spectral deflation combined with MLFMA.

Cobra problem, $n = 60695$ - $frequency = 10.0$ GHz

	Dimension of the coarse space			
	0	5	10	15
GMRES(10)	2719 (1^h 10^m)	1458 (42^m)	594 (12^m)	517 (11^m)
GMRES(∞)	378 (18^m)	262 (9^m)	216 (7^m)	188 (6^m)

Almond problem, $n = 104793$ - $frequency = 2.6$ GHz

	Dimension of the coarse space			
	0	10	30	50
GMRES(50)	1524 (1^h 17^m)	883 (45^m)	368 (20^m)	284 (5^m)
GMRES(∞)	242 (14^m)	134 (9^m)	92 (6^m)	77 (6^m)

In Table 5 we show the number of *amortization* vectors relative to GMRES(10), that is the number of right-hand sides that have to be considered to amortize the extra cost for the eigencomputation. In bistatic radar cross section calculations, linear systems with the same coefficient matrix and up-to hundreds of different right-hand sides are solved, ranging over the complete set of directions between the transmitter and the receiver. In Table 6 the low-rank deflation technique is combined with the FMM. We report on an experiment on the Cobra problem (Fig. 3(a), $n = 60695$) and on the Almond problem (Fig. 3(b), $n = 104793$). We see on the Cobra problem that with a preconditioning update of only 10 eigenvectors and setting very low restart in GMRES, we are able to reduce the number of iterations by nearly a factor of six; a significant reduction of both number of iterations and solution time is also observed on the Almond problem as well owing to the fact that we use low-accurate MLFMA for computing the spectral information. On that computer, the temporary disk space that can be used by the out-of-core solver is around 189 GB. On that hardware the CPU time is also reduced by a factor of six.

IV. CONCLUDING REMARKS

We have presented experiments with an adaptive preconditioning method constructed on top of an approximate inverse preconditioner for solving dense linear systems of equations arising in electromagnetic scattering applications. The results show that the proposed method can be very effective to accelerate the convergence of iterative Krylov solvers. Sparse approximate inverses based on Frobenius-norm minimization are very good candidates for preconditioning this problem class for their inherent parallelism and their proved numerical stability on indefinite systems. We have shown that the construction cost can be controlled using static pattern selection strategies and the performance can be enhanced significantly using deflating techniques.

ACKNOWLEDGEMENTS

The author would like to thank the EMC-CERFACS Team (Toulouse, France) for providing with the MoM code and the EADS-CRC company (Toulouse, France) for providing with the MLFMA code and the large test examples for the numerical experiments. The author is grateful to the anonymous referees for their valuable comments and suggestions that greatly improved the presentation of this paper.

REFERENCES

[1] A.F. Peterson, S.L. Ray, and R. Mittra. *Computational Methods for Electromagnetics.* IEEE Press, 1997.

[2] G. Alléon, S. Amram, N. Durante, P. Homsi, D. Pogarieloff, and C. Farhat, "'Massively parallel processing boosts the solution of industrial electromagnetic problems: High performance out-of-core solution of complex dense systems,"' In M. Heath, V. Torczon, G. Astfalk, P. E. Bj?rstad, A. H. Karp, C. H. Koebel, V. Kumar, R. F. Lucas, L. T. Watson, and Editors D. E. Womble, editors, *Proceedings of the Eighth SIAM Conference on Parallel*. SIAM Book, Philadelphia, Minneapolis, Minnesota, USA, 1997.

[3] W. C. Chew and Y. M. Wang. A recursive T-matrix approach for the solution of electromagnetic scattering by many spheres. *IEEE Transactions on Antennas and Propagation*, vol. 41, no. 12, pp. 1633–1639, 1993.

[4] K. Chen. An analysis of sparse approximate inverse preconditioners for boundary integral equations. *SIAM J. Matrix Analysis and Applications*, vol. 22, no. 3, pp. 1058–1078, 2001.

[5] K. Sertel and J. L. Volakis. Incomplete LU preconditioner for FMM implementation. *Micro. Opt. Tech. Lett.*, vol. 26, no. 7, pp. 265–267, 2000.

[6] A. Grama, V. Kumar, and A. Sameh. On n-body simulations using message-passing parallel computers. In Sidney Karin, editor, *Proceedings of the 1995 SIAM Conference on Parallel Processing, San Francisco, CA, USA*, 1995.

[7] S. A. Vavasis. Preconditioning for boundary integral equations. *SIAM J. Matrix Analysis and Applications*, vol. 13, pp. 905–925, 1992.

[8] J. Lee, C.-C. Lu, and J. Zhang. Incomplete LU preconditioning for large scale dense complex linear systems from electromagnetic wave scattering problems. *J. Comp. Phys.*, vol. 185, pp. 158–175, 2003.

[9] B. Carpentieri, I. S. Duff, L. Giraud, and M. Magolu monga Made. Sparse symmetric preconditioners for dense linear systems in electromagnetism. Technical Report TR/PA/01/35, CERFACS, Toulouse, France, 2001. Also Technical Report RAL-TR-2002-016. Preliminary version of the article accepted for publication in *Numerical Linear Algebra with Applications*.

[10] T. Malas and L. Gürel. Incomplete LU preconditioning with multilevel fast multipole algorithm for electromagnetic scattering. *SIAM J. Scientific Computing*, 2007. Accepted for publication.

[11] G. Alléon, M. Benzi, and L. Giraud. Sparse approximate inverse preconditioning for dense linear systems arising in computational electromagnetics. *Numerical Algorithms*, vol. 16, pp. 1–15, 1997.

[12] A. R. Samant, E. Michielssen, and P. Saylor. Approximate inverse based preconditioners for 2D dense matrix problems. Technical Report CCEM-11-96, University of Illinois, 1996.

[13] J. Lee, C.-C. Lu, and J. Zhang. Sparse inverse preconditioning of multilevel fast multipole algorithm for hybrid integral equations in electromagnetics. Technical Report 363-02, Department of Computer Science, University of Kentucky, KY, 2002.

[14] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *Journal of Computational Physics*, vol. 73, pp. 325–348, 1987.

[15] J. Song, C.-C. Lu, and W. C. Chew. Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects. *IEEE Transactions on Antennas and Propagation*, vol. 45, no. 10, pp. 1488–1493, 1997.

[16] E. Darve. The fast multipole method: Numerical implementation. *J. Comp. Phys.*, vol. 160, no. 1, pp. 195–240, 2000.

[17] G. Sylvand. *La Méthode Multipôle Rapide en Electromagnétisme : Performances, Parallélisation, Applications*. PhD thesis, Ecole Nationale des Ponts et Chaussées, 2002.

[18] B. Carpentieri, I. S. Duff, and L. Giraud. Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism. *Numerical Linear Algebra with Applications*, vol. 7, no. 7-8, pp. 667–685, 2000.

[19] B. Carpentieri, I. S. Duff, and L. Giraud. A class of spectral two-level preconditioners. *SIAM J. Scientific Computing*, vol. 25, no. 2, pp. 749–765, 2003.

[20] J. Erhel, K. Burrage, and B. Pohl. Restarted GMRES preconditioned by deflation. *J. Comput. Appl. Math.*, vol. 69, pp. 303–318, 1996.

[21] Y. Saad. Projection and deflation methods for partial pole assignment in linear state feedback. *IEEE Trans. Automat. Contr.*, vol. 33, no. 3, pp. 290–297, 1988.

[22] R. Lehoucq, D. Sorensen, and P. Vu. *ARPACK User's Guide: Solution of LargeScale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. Society for Industrial and Applied Mathematics, Philadelphia, 1998.

Bruno Carpentieri received a Ph.D. degree in Computer Science from Institut National Polytechnique of Toulouse, France, in 2002. He is a Postdoctoral Fellow at the Institut of Mathematics and Scientific Computing of Karl-Franzens University in Graz, Austria. His research interests include numerical linear algebra, high performance computing, Computational Fluid Dynamics, Computational Electromagnetics.