

Using spectral low rank preconditioners for large electromagnetic calculations

I. S. Duff^{1,2}, L. Giraud¹, J. Langou^{1,‡} and E. Martin^{1,*,†}

¹*CERFACS, 42 Avenue G. Coriolis, 31057 Toulouse Cedex, France*

²*RAL, Oxfordshire, England*

SUMMARY

For solving large dense complex linear systems that arise in electromagnetic calculations, we perform experiments using a general purpose spectral low rank update preconditioner in the context of the GMRES method preconditioned by an approximate inverse preconditioner. The goal of the spectral preconditioner is to improve the convergence properties by shifting by one the smallest eigenvalues of the original preconditioned system. Numerical experiments on parallel distributed memory computers are presented to illustrate the efficiency of this technique on large and challenging real-life industrial problems. Copyright © 2004 John Wiley & Sons, Ltd.

KEY WORDS: electromagnetic scattering problems; large dense complex linear systems; Frobenius-norm minimization preconditioner; spectral low rank update preconditioner

1. INTRODUCTION

In electromagnetic calculations, a classic problem is to compute the currents generated on the surface of an object illuminated by a given incident plane wave. Such calculations, relying on the Maxwell's equations, are required in the simulation of many industrial processes coming from antenna design, electromagnetic compatibility, computation of back-scattered fields, and so on. Recently, the boundary element method (BEM) has been successfully used in the numerical solution of this class of problems. The formulation considered in this paper is the electric field integral equation (EFIE) as it is the most general and does not require any assumption about the geometry of the object. The matrices associated with the resulting linear systems are large dense, non-Hermitian and complex. With the advent of parallel processing, this approach has become viable for large problems and the typical problem size in the electromagnetic industry is continually increasing. Nevertheless, nowadays, problems with a few hundred thousand variables

*Correspondence to: E. Martin, CERFACS, 42 Avenue G. Coriolis, 31057 Toulouse Cedex, France.

†E-mail: Emeric.martin@cerfacs.fr

‡Current address: Department of Computer Science, University of Tennessee, Knoxville, Tennessee, U.S.A.

Contract/grant sponsor: EADS, Corporate Research Centre, Toulouse

Received 5 November 2003

Revised 7 June 2004

Accepted 16 July 2004

can no longer be solved by direct solvers even if they are parallel and out-of-core because they require too much memory, CPU and disk resources. Iterative solvers appear as the only viable alternative since techniques based on multipole expansion [1, 2] have been developed to perform fast matrix–vector products without forming all the entries of the dense matrices. In particular, the fast multipole method (FMM) performs the matrix–vector product in $\mathcal{O}(n \log n)$ floating-point operations and can efficiently be implemented on parallel distributed platforms with some out-of-core techniques in order to tackle huge industrial problems [3]. The industrial problem we focus on in this paper is the monostatic radar cross-section calculation of an object. The procedure consists in considering a set of waves with the same wavelength but different incident angles that illuminate the object. For each of these waves we compute the electromagnetic field back-scattered in the direction of the incident wave. This requires the solution of one linear system per incident wave. For a complete radar cross-section calculation, from a few tens up to a few hundred waves have to be considered. We have then to solve a sequence of linear systems having the same coefficient matrix but different right-hand sides. The problem can be written as

$$A(x_1, \dots, x_p) = (b_1, \dots, b_p)$$

In this context it is particularly important to have a numerically efficient and easily parallelizable preconditioner. A preconditioner suitable for implementation in a multipole framework on parallel distributed platforms has been proposed in References [4, 5]; it is based on a sparse approximate inverse using a Frobenius norm minimization with an *a priori* sparsity pattern selection strategy. It has been shown [6, 7] that on medium size problems this technique gives rise to an effective preconditioner that outperforms more classical approaches like incomplete factorizations or other general purpose approximate inverses [8–10]. However, for large problems the preconditioner becomes sparser and sparser when the problem size increases and eventually performs poorly. Unfortunately, making it denser is not feasible due to memory and disk constraints [11]. In this paper, we investigate the use of a spectral low rank update [12] that attempts to shift by one the smallest eigenvalues of the preconditioned systems resulting in a faster convergence rate. The paper is organized as follows. In Section 2, we recall the main features of the spectral low rank update preconditioner and describe the real-life test problems we consider for the numerical experiments. Then, in Section 3, we illustrate the effect of the size of the update on the convergence rate of GMRES [13] for one right-hand side and show the gain for a complete radar cross-section in the following section. Since the spectral low rank update compensates for some possible weaknesses of the approximate inverse, we illustrate in Section 5 that a balance has to be found between the two components of the resulting preconditioner. In Section 6, we show that the gain induced by the low rank update becomes larger as the size of the restart in GMRES decreases. Finally, since many right-hand sides have to be solved, we illustrate in Section 7 the benefit of using the preconditioner in the context of seed GMRES that is a Krylov solver designed to deal with multiple right-hand sides. We conclude with some remarks and comments on future work in Section 8.

2. BACKGROUND AND TEST EXAMPLES

For the solution of large linear systems using fast multipole techniques on parallel distributed memory computers, we consider an approximate inverse preconditioner based on a Frobenius

norm minimization procedure and we denote it by M_{Frob} . Information from the underlying physical problem is exploited to prescribe in advance the sparsity pattern of the preconditioner [4, 6]. We note that its density in terms of number of non-zero entries per column can be adjusted through a simple parameter. For further details on its implementation in the fast multipole code as well as its behaviour on large problems we refer to References [3, 11]. In this paper, we investigate the behaviour of the low rank update preconditioner described in Reference [12] that attempts to improve the convergence rate of the Krylov solver by shifting by one the smallest eigenvalues that M_{Frob} leaves close to the origin. On these linear systems, GMRES has shown itself to be the most robust Krylov solver [3, 11]. We note that other related approaches that attempt to remove the effect of the smallest eigenvalues are available in the literature. In the context of adaptive preconditioners for restarted GMRES, we refer the reader to References [14–16].

Let us describe the spectral low rank update for the right preconditioner that we use later in our experiments. Given the right preconditioned linear system:

$$AM_{\text{Frob}}u = b \quad \text{with } x = M_{\text{Frob}}u$$

where the matrix A is $n \times n$ complex non-singular and M_{Frob} is the right preconditioner, the spectral low rank update is as follows. For the sake of simplicity, we assume that the preconditioned matrix AM_{Frob} is diagonalizable. That is,

$$AM_{\text{Frob}} = V\Lambda V^{-1}$$

with Λ the diagonal matrix formed by the eigenvalues $\{\lambda_i\}_{i \in \{1, n\}}$ ordered by increasing magnitude, and V the matrix whose columns correspond to the right eigenvectors. Let us consider the k smallest eigenvalues of AM_{Frob} denoted by λ_i with $i = 1, \dots, k$. We denote by V_k the set of the right eigenvectors associated with these k smallest eigenvalues. In Reference [12], it is shown that, if W is a n by k matrix such that $A_0 = W^H AM_{\text{Frob}} V_k$ has full rank, then setting

$$M_0 = M_{\text{Frob}} V_k A_0^{-1} W^H \quad \text{and} \quad M_{\text{SLRU}(k)} = M_{\text{Frob}} + M_0$$

we have that $AM_{\text{SLRU}(k)}$ is similar to a matrix whose eigenvalues are:

$$\begin{cases} \eta_i = \lambda_i & \text{if } i > k \\ \eta_i = 1 + \lambda_i & \text{if } i \leq k \end{cases} \quad (1)$$

The matrix M_0 is a rank- k correction of M_{Frob} , which ensures that the new system:

$$AM_{\text{SLRU}(k)}u = b \quad \text{with } x = M_{\text{SLRU}(k)}u$$

no longer has eigenvalues with magnitude smaller than $|\lambda_{k+1}|$. Note that we can choose other shifts in other contexts. In our case, as most of eigenvalues are already close to one, this shift makes sense. Furthermore, we set $W = V_k$ in our experiments. In Reference [12], some promising results in the context of BEM are shown on small examples with a few hundred unknowns; the spectral low rank update is combined with the GMRES method preconditioned by a Frobenius norm minimizer preconditioner. In this work, we investigate the same solver using the same combination of preconditioners but on large real industrial applications. The test geometries are shown in Figure 1. They consist of a wing with a hole referred to as




Figure 1. Various geometries used in the numerical experiments: (a) Cetaf; (b) civil aircraft from a European manufacturer; (c) Cobra; and (d) Almond.

Table I. Angular section of interest for each geometry; the discretization step is one degree.

Geometry	# unknowns	Density (%)	θ	φ	# rhs
Cetaf	5 391	3.3	(-90)-90	0	181
Aircraft	23 676	0.94	90	0-180	181
Cobra	60 695	0.24	0-90	0	91
Almond	104 793	0.19	90	0-180	181

Cetaf, an industrial civil aircraft from a European consortium producing a family of passenger aircrafts, an air intake referred to as Cobra, and finally an Almond. The Cetaf and Almond cases are classic test problems in the computational electromagnetic community; the other two have been kindly provided to us by EADS-CCR (Corporate Research Center of the European Aeronautic Defense and Space Company). Given these four geometries, the sets of angles of interest for the monostatic radar cross-section vary. In Table I we indicate, for each geometry, the number of unknowns associated with each linear system, the density of M_{Frob} , the angular




Figure 2. Spherical co-ordinates.

section considered (given in spherical co-ordinates, see Figure 2) as well as # rhs, the number of right-hand sides to be solved for the complete monostatic calculation.

The numerical experiments reported in this paper have been performed on a HP-Compaq Alpha (EV 68, 1.3 Gflops peak) server that is a cluster of symmetric multiprocessors. The stopping criterion consists in reducing the original residual by 10^{-3} that can then be related to a norm-wise backward error. Although this value may appear to be large, it is sufficient to obtain accurate radar cross-section results. Finally, the initial guess is set to the zero vector.

3. EFFICIENCY OF THE PRECONDITIONER WITH RESPECT TO THE RANK OF THE UPDATE

In Figure 3, we plot using the symbol '×' the spectrum of AM_{Frob} , the matrix preconditioned with the Frobenius preconditioner, for the Cetaf example. As can be observed, M_{Frob} succeeds in clustering most of the eigenvalues around (1.0, 0.0). Such a distribution is highly desirable to get fast convergence of Krylov solvers. Nevertheless the remaining eigenvalues nearest to zero can potentially slow down the convergence. Using the symbol 'o' we plot, in Figure 3, the spectrum of the matrix preconditioned with $M_{\text{SLRU}(20)}$. We observe that the 20 smallest eigenvalues of the matrix AM_{Frob} have been shifted close to one, in agreement with (1). Consequently, we expect the Krylov solver to perform better with $M_{\text{SLRU}(20)}$ than with M_{Frob} . In Figure 4, we plot the convergence histories obtained by varying the size of the low rank update. It can be observed that the larger the rank, the faster the convergence of full GMRES. However, in going from 15 to 20 the gain is negligible and going beyond 20 does not give further improvements.




Figure 3. Spectrum of AM_{Frob} denoted by \times and $AM_{\text{SLRU}(20)}$ depicted with \circ on the Cetaf test problem.




Figure 4. Convergence history when varying the rank of the update on the Cetaf example.




Figure 5. Number of full GMRES iterations when varying k in $M_{\text{SLRU}(k)}$ on the Cobra test problem.

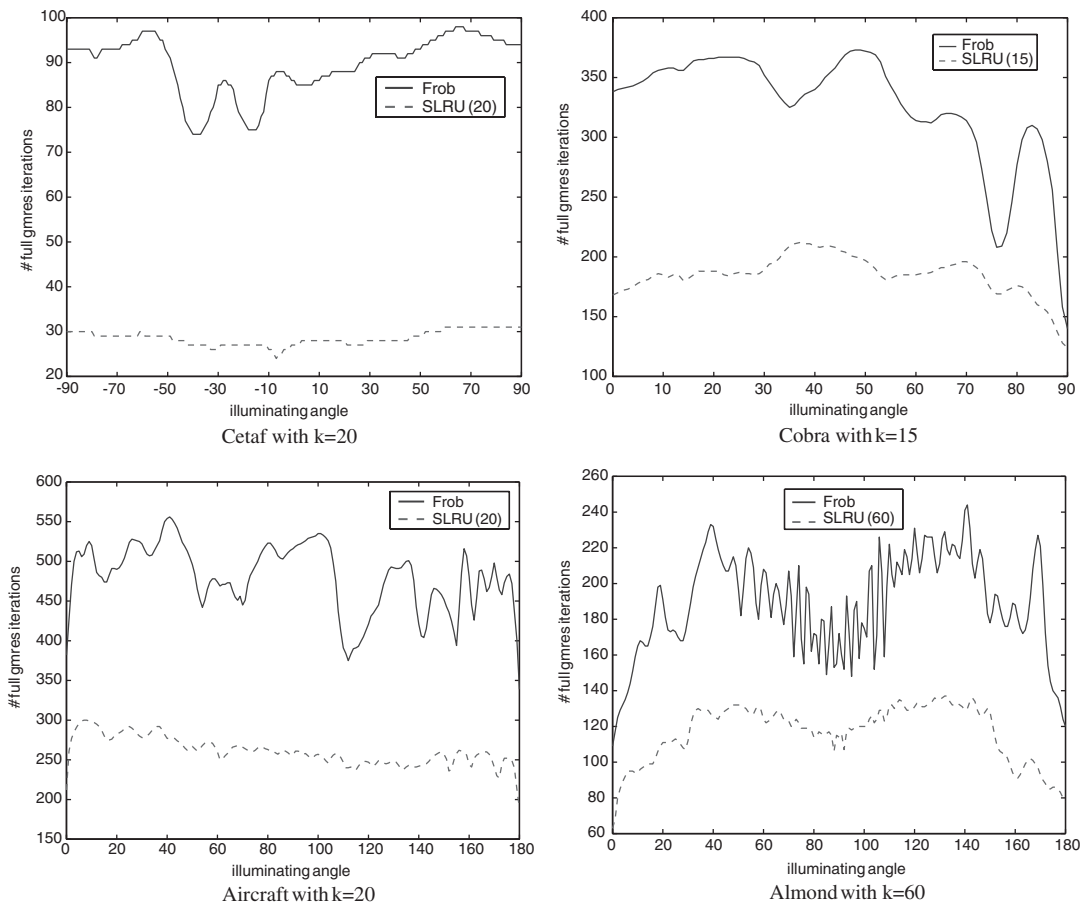
As we mentioned earlier, we have used this technique for monostatic calculation. As many linear systems with the same coefficient matrix but different right-hand sides have to be solved, some are easier to solve than others. In Figure 5, we illustrate an important feature of the low rank update by showing the number of iterations for convergence for a difficult right-hand side and an easy one. It can be observed that, when the number of shifted eigenvalues increases, the number of iterations to reach convergence decreases. Furthermore, there is not much difference between a difficult right-hand side and an easy one when the rank of the update increases. Later in this paper, we illustrate the advantage of this feature in the context of restarted GMRES and seed GMRES.

4. RESULTS ON A COMPLETE MONOSTATIC CALCULATION

The use of a preconditioner is often beneficial, however its usefulness depends not only on its effect on convergence but also its construction time and the time spent in applying it at each step. In Table II, we give the construction time for the FMM operator, for M_{Frob} and for $M_{\text{SLRU}(k)}$. In our experiments, the eigenvalue calculation is performed in a preprocessing phase using ARPACK [17], in a forward mode that represents the main part of the time required to setup $M_{\text{SLRU}(k)}$. In that table, we also display the average time spent in one FMM operation, one product by M_{Frob} , and one application of $M_{\text{SLRU}(k)}$. We expect to reduce the overall number

Table II. Average elapsed time per matrix–vector product.

Geometry	# proc.	k	Construction times			Application times		
			FMM (s)	M_{Frob}	$M_{\text{SLRU}(k)}$	FMM (s)	M_{Frob} (s)	$M_{\text{SLRU}(k)}$ (s)
Cetaf	8	20	13	25 s	45 s	0.21	0.03	0.04
Aircraft	32	20	27	51 s	19 min	0.83	0.12	0.15
Cobra	32	15	36	73 s	28 min	1.26	0.11	0.12
Almond	32	60	59	3 min	2 h	1.91	0.14	0.19

Figure 6. Number of full GMRES iterations with M_{Frob} and $M_{\text{SLRU}(k)}$ for the different incident angles for each geometry.

of iterations, but each new iteration is more expensive. One application of $M_{\text{SLRU}(k)}$ compared with one application of M_{Frob} introduces around $4.k.n$ additional flops, where k is the chosen number of eigenvalues and n the size of the problem. On our test examples the extra cost per

Table III. Cost for a complete monostatic calculation.

Geometry	# procs.	M_{Frob}		$M_{\text{SLRU}(k)}$	
		# iter	Times	# iter	Times
Cetaf	8	16 391	1 h 40 min	5 349	47 min
Aircraft	32	87 121	46 h	47 385	18 h 40 min
Cobra	32	29 777	21 h	16 921	8 h 30 min
Almond	32	34 375	25 h 30 min	21 273	14 h 40 min

Table IV. Cost of the eigencomputation preprocessing phase.

Geometry	# procs.	# mat-vec	Times	# rhs
Cetaf	8	170	45 s	3
Aircraft	32	1 000	19 min	6
Cobra	32	1 000	28 min	6
Almond	32	2 200	2 h	32

iteration in elapsed time ranges from 6% (Cobra case) to 35% (Almond case), but it remains small in comparison to the FMM application times.

In Figure 6, we show the number of full GMRES iterations for each right-hand side using either M_{Frob} (solid line) or $M_{\text{SLRU}(k)}$ (dashed line). For each geometry the value of k is given in the caption. While without the spectral preconditioner, the numbers of iterations from one right-hand side to another vary a lot, with the spectral preconditioner the number of iterations per right-hand side is more similar and almost constant for some geometries. This behaviour was already observed in Figure 5.

Table III summarizes the cumulated numbers of matrix-vector products and the total elapsed solution time for a complete radar cross-section calculation for each geometry using full GMRES. For all geometries except the Cetaf, 181 linear systems are solved; only 91 are considered for the Cetaf test problem. Depending on the geometry, the overall gain ranges from a factor of 1.6 to 3 for both CPU time and total number of GMRES iterations. It should be pointed out that this could be improved on some examples if more eigenvalues were shifted. Our purpose in these experiments is to illustrate the potential of $M_{\text{SLRU}(k)}$, and we did not try to find the best values of k for each geometry. For example, by shifting 10 more small eigenvalues for the Cobra case, we move from a factor 1.6 to a factor 3.

The extra cost for computing the eigenspace during the preprocessing phase in terms of matrix-vector products by AM_{Frob} as well as the corresponding elapsed time is displayed in Table IV. In that table, we also give the number of right-hand sides in the monostatic calculation, denoted # rhs, for which the gain introduced by $M_{\text{SLRU}(k)}$ compensates for the cost of computing the preconditioner. It can be seen that the preprocessing calculation is quickly amortized when a few right-hand sides need to be solved.

In addition, we should mention that the quality of the eigenvectors in terms of backward error does not need to be very accurate: using eigenvectors with a backward error of less than 10^{-3} does not significantly improve convergence if at all.

5. BALANCING THE TWO COMPONENTS OF THE PRECONDITIONER

Although we save a significant number of iterations using $M_{\text{SLRU}(k)}$, we might ask whether it could be more effective to use a better M_{Frob} , by allowing the preconditioner to have more non-zero entries, combined with a migration of only some smallest eigenvalues, or a worse M_{Frob} , but combined with a migration of many smallest eigenvalues. To deal with this, we first investigate the effect of the quality of M_{Frob} on the calculation time of the smallest eigenvalues of AM_{Frob} . In that respect, we vary the number of non-zeros per column leading to different densities of M_{Frob} . In Figure 7, we display the spectrum of AM_{Frob} on the Cetaf example for various values of the density. As we might expect, the denser M_{Frob} is, the better the clustering around one and the fewer eigenvalues close to zero; moreover these are better separated. A consequence for the eigensolver is that the few well separated eigenvalues that are near zero for the largest density of the preconditioner are more easily computed. When the density is relaxed, the eigenvalue cluster close to zero becomes wider and the eigensolver has more difficulty in computing those near zero. Indeed, it needs extra effort to identify the eigenvalues in a cluster.




Figure 7. Spectrum of AM_{Frob} for various density of M_{Frob} : (a) 0.7%; (b) 3.3%; and (c) 15.7%.

Table V. Number of matrix–vector products and elapsed times (one processor) required to compute the 60 smallest eigenvalues of AM_{Frob} for the Cetaf when the density of M_{Frob} is varied.

Density (%)	# mat–vec	Times
0.7	365	7 min 57 s
3.3	240	4 min 41 s
15.7	152	2 min 52 s

Table VI. Construction times when varying the M_{Frob} density on the Cetaf test problem.

M_{Frob} construction		Eigensolver		Total
Density (%)	Times (s)	k	Times	
0.7	41	54	7 min 49 s	8 min 30 s
3.3	143	20	2 min 41 s	5 min 04 s
15.7	1114	2	1 min 04 s	19 min 38 s

To illustrate this claim we show, in Table V, the number of matrix–vector products and the corresponding elapsed time required by ARPACK to find the eigenvectors associated with the 60 smallest eigenvalues as the density increases. As expected, we observe that the denser the preconditioner is, the easier it is for the eigensolver to find the smallest eigenvalues.

A question that seems natural to raise is: how many eigenvalues should be shifted for these three densities of M_{Frob} to get convergence in the same number of full GMRES iterations? On the Cetaf for a difficult angle and a density equal to 3.3%, full GMRES needs 98 iterations to converge. Using $M_{\text{SLRU}(20)}$ GMRES needs 31 iterations to converge. The number 31 is taken as a reference value to compare the performance with the three densities. Table VI shows the cost for ARPACK to compute the corresponding number of eigenvalues for each density and the total cost of computing the preconditioner.

To obtain the same number of full GMRES iterations, we need to shift more and more eigenvalues as we decrease the density. On the other hand, the construction cost of M_{Frob} with a low density is cheaper than with a higher density. There is a trade-off to be found between a cheap M_{Frob} that requires shifting many eigenvalues that might be difficult to compute, and a more expensive M_{Frob} where only a few eigenvalues need to be shifted to get a similar convergence behaviour. As Table VI shows, the medium density 3.3% offers the best trade-off among the three densities considered. The preconditioner already works well and only a few eigenvalues need to be shifted. In each case shown in Table VI, 31 iterations of preconditioned full GMRES are needed for convergence. The symbol ‘ k ’ is the number of eigenvalues computed by ARPACK.

Let us illustrate, on another example, the advantage of using $M_{\text{SLRU}(k)}$ rather than increasing the density of M_{Frob} . We consider now the Almond test problem, that is the biggest, using two densities for M_{Frob} . The targeted number of iterations of full GMRES is 157 that is obtained with a density of 0.19% and by shifting 30 eigenvalues using $M_{\text{SLRU}(30)}$. To get the same

Table VII. Comparison of a denser M_{Frob} with a sparser $M_{\text{SLRU}(30)}$ on eight processors on the Almond test example to obtain 157 iterations with full GMRES.

M_{Frob}		Setup			Solution		
Density (%)	Times	# Eigen.	Times	Total times	Application times		Total times
					M_{Frob}	$M_{\text{SLRU}(30)}$	
0.19	6 min	30	3 h	3 h 06 min	—	0.38 s	8 min 30 s
1.76	5 h 40 min	0	—	5 h 40 min	1.21 s	—	10 min 48 s

number of iterations without shifting eigenvalues, we need to increase the density of M_{Frob} to 1.76%. In Table VII, we show the computation cost for both preconditioners. It can be seen that the eigencalculation in the preprocessing phase of $M_{\text{SLRU}(30)}$ combined with the low cost of its M_{Frob} component is significantly less expensive than the denser M_{Frob} that exhibits the same convergence property. We compare the application times for these two approaches and the time to obtain solution. The second approach is four times as expensive needing 1.21 s per application as opposed to the first with 0.38 s. The use of M_{Frob} with a density of 0.19% in the first case would have cost 0.36 s; it means that $M_{\text{SLRU}(30)}$ yields an extra cost of only 0.02 s per application. Moreover, it gives a smaller solution time and a much cheaper setup time. Setup time appears to be too much important towards solution time, but these results are obtained on just one angle. For solutions on an angular section, we will pay this setup time for just one time.

6. SENSITIVITY OF THE RESTARTED GMRES

All the numerical experiments reported on so far have been obtained with full GMRES. In this section we will investigate the effect of restarted GMRES on the efficiency of the preconditioners. For each value of the restart, we show in Table VIII the number of GMRES iterations of M_{Frob} and $M_{\text{SLRU}(k)}$ on a easy and a hard right-hand side. The symbol ‘—’ means that convergence is not obtained within 5000 iterations. As can be seen, the smaller the restart is, the larger the improvement with $M_{\text{SLRU}(k)}$. With $M_{\text{SLRU}(20)}$ on the Cetaf example, we see that GMRES converges quickly and that GMRES(10) behaves very similarly to GMRES(∞). This observation is no longer true for the other geometries. It might depend on the number of eigenvalues that are computed: the choice is the best for the Cetaf example but not for the other geometries.

The $M_{\text{SLRU}(k)}$ preconditioner allows us to use a smaller restart than usual; this might be a significant asset on very large problems where using a large restart becomes a severe bottleneck because of the prohibitive memory requirements. There are some problems where M_{Frob} does not converge, for instance on the Aircraft case with a restart between 10 and 50, or on the Almond example with a restart of 10, while $M_{\text{SLRU}(k)}$ does converge in a reasonable number of iterations. On that later example, as we see in Figure 8, the convergence rate of GMRES increases with the restart whatever the chosen preconditioner. Furthermore, the slope of the convergence history for $M_{\text{SLRU}(k)}$ becomes quickly comparable to that for full GMRES, while for M_{Frob} this phenomenon takes more time to appear. Although not reported here, the same behaviour was observed on the other geometries.

Table VIII. Sensitivity to the GMRES restart parameter.

Restart	Cetaf		Aircraft		Cobra		Almond	
	M_{Frob}	$M_{\text{SLRU}(20)}$	M_{Frob}	$M_{\text{SLRU}(20)}$	M_{Frob}	$M_{\text{SLRU}(15)}$	M_{Frob}	$M_{\text{SLRU}(60)}$
<i>Easy case</i>								
10	271	30	—	639	514	378	492	95
30	128	27	—	439	280	196	149	67
50	100	27	—	390	264	188	118	68
∞	74	27	421	242	222	172	109	63
<i>Difficult case</i>								
10	669	37	—	1200	2624	481	—	1497
30	275	31	—	688	1031	232	429	163
50	197	31	—	608	760	207	334	144
∞	98	31	490	283	367	187	232	126




Figure 8. Convergence history varying the restart for a 'difficult' angle of the Almond test problem.

7. COMPLEMENTARITY OF $M_{\text{SLRU}(k)}$ AND THE SEED GMRES ALGORITHM

As mentioned already for a complete monostatic calculation, several linear systems with the same coefficient matrix but different right-hand sides have to be solved. In that framework, it

is crucial not only to use an efficient preconditioner but also a suitable Krylov solver. There are basically two classes of techniques designed for this purpose. These are the seed techniques and the block approaches [18–21]. In this section, we investigate the benefit of using $M_{\text{SLRU}(k)}$ in the context of seed GMRES.

For the sake of simplicity of the notation, we describe below the method for the solution of only two right-hand sides. The problem can be written as

$$\begin{cases} A(x_1, x_2) = (b_1, b_2) \text{ with} \\ (x_1^0, x_2^0) \text{ initial guesses} \end{cases}$$

The seed GMRES method starts by solving the first system $Ax_1 = b_1$ using the GMRES algorithm. Once that is converged, it computes an orthonormal basis of the Krylov space (V_{m+1}^1) and the upper-Hessenberg matrix (\bar{H}_m^1) such that $AV_m^1 = V_{m+1}^1 \bar{H}_m^1$. The idea consists in updating the initial guess x_2^0 , by minimizing the norm of the residual $r_2^0 = b_2 - Ax_2^0$ of the second system on the basis V_{m+1}^1 . We search a new initial guess \hat{x}_2^0 such as

$$\begin{cases} \hat{x}_2^0 = x_2^0 + V_m^1 y_2 \\ y_2 = \arg \min_y \|b_2 - A(x_2^0 + V_m^1 y)\|_2 \end{cases}$$

Using the fact that

$$\begin{aligned} \|b_2 - A(x_2^0 + V_m^1 y)\|_2^2 &= \|r_2^0 - AV_m^1 y\|_2^2 \\ &= \|(I_n - V_{m+1}^1 (V_{m+1}^1)^T) r_2^0 + V_{m+1}^1 (V_{m+1}^1)^T r_2^0 - AV_m^1 y\|_2^2 \\ &= \|(I_n - V_{m+1}^1 (V_{m+1}^1)^T) r_2^0 + V_{m+1}^1 ((V_{m+1}^1)^T r_2^0 - \bar{H}_m^1 y)\|_2^2 \\ &= \|(I_n - V_{m+1}^1 (V_{m+1}^1)^T) r_2^0\|_2^2 + \|V_{m+1}^1 ((V_{m+1}^1)^T r_2^0 - \bar{H}_m^1 y)\|_2^2 \\ &= \|(I_n - V_{m+1}^1 (V_{m+1}^1)^T) r_2^0\|_2^2 + \|(V_{m+1}^1)^T r_2^0 - \bar{H}_m^1 y\|_2^2 \end{aligned}$$

we finally should solve the least-squares problem: $\arg \min_y \|(V_{m+1}^1)^T r_2^0 - \bar{H}_m^1 y\|_2$, where we already have a QR factorization of \bar{H}_m^1 given by the GMRES algorithm. Once the new initial guess is computed, we simply run a new GMRES method for the second-right-hand side. When more than two right-hand sides are considered, the same algorithm can be applied but we have to solve a least-squares problem for each previous right-hand side after the first.

In order to illustrate the advantage of using $M_{\text{SLRU}(k)}$ in this context, we consider 11 right-hand sides: $\phi = 15^\circ : 1^\circ : 25^\circ$ for $\theta = 90^\circ$, involved in the radar cross-section calculation of the Aircraft. On these right-hand sides, seed GMRES combined with M_{Frob} behaves rather poorly. To illustrate this, we first compare the numerical behaviour of GMRES with three strategies for defining the initial guess: first using the zero vector, second using the solution of the previous linear system, and finally the initial guess computed by the seed technique. In Table IX, we display the number of iterations for each right-hand side for the three initial guess strategies.

Table IX. Number of iterations per right-hand side using three strategies for the initial guess on the Aircraft example with M_{Frob} preconditioner on eight processors.

(θ, φ)	M_{Frob} preconditioner				
	Zero guess # iter	Simple strat.		Seed GMRES	
		# iter	$\ r_0\ _2/\ b\ _2$	# iter	$\ r_0\ _2/\ b\ _2$
(90,15)	474	474	1	474	1
(90,16)	474	443	0.284	440	0.13
(90,17)	483	435	0.298	338	0.05
(90,18)	491	442	0.311	387	0.023
(90,19)	491	438	0.325	458	0.008
(90,20)	490	438	0.338	543	0.004
(90,21)	492	436	0.352	605	0.003
(90,22)	497	418	0.365	599	0.003
(90,23)	499	424	0.379	589	0.003
(90,24)	499	431	0.392	601	0.003
(90,25)	499	406	0.406	573	0.003
# iterations	5389	4785		5607	
elapsed time	4 h 50 min	4 h 17 min		8 h	

Table X. Number of iterations per right-hand side using three strategies for the initial guess on the Aircraft example with $M_{\text{SLRU}(20)}$ preconditioner on eight processors.

(θ, φ)	$M_{\text{SLRU}(20)}$ preconditioner				
	Zero guess # iter	Simple strat.		Seed GMRES	
		# iter	$\ r_0\ _2/\ b\ _2$	# iter	$\ r_0\ _2/\ b\ _2$
(90,15)	280	280	1	280	1
(90,16)	275	198	0.284	201	0.14
(90,17)	275	188	0.298	145	0.052
(90,18)	276	190	0.311	167	0.025
(90,19)	280	198	0.325	165	0.009
(90,20)	283	205	0.338	171	0.006
(90,21)	284	208	0.352	171	0.005
(90,22)	286	212	0.365	170	0.005
(90,23)	289	214	0.379	172	0.005
(90,24)	291	218	0.392	176	0.005
(90,25)	292	219	0.406	171	0.005
# iterations	3111	2330		1989	
elapsed time	2 h 19 min	1 h 44 min		1 h 32 min	

In the case of zero as initial guess, the initial backward error $\|r_0\|_2/\|b\|_2 = \|b - Ax_0\|_2/\|b\|_2$ is equal to one. In Table IX, we see that the seed GMRES method does a good job of decreasing the initial residual norm; it is always by far the smallest. Unfortunately, starting from the seed initial guess that is the closest (in the backward error sense) to the solution does not guarantee fast convergence. That is, from that good initial guess, GMRES performs rather poorly. It performs only slightly better than starting from zero and is outperformed by the approach




Figure 9. Convergence history of seed GMRES for the right-hand sides associated with $\phi = 15^\circ : 1^\circ : 25^\circ$ using M_{Frob} (\circ , in dashed line) and $M_{\text{SLRU}(20)}$ (\times , in solid line) on the Aircraft test problem.

that starts from the initial guess provided by the simple strategy of using the solution to the previous system. In other words and surprisingly enough, the approach that gives the smallest initial residual norm is not the method that gives the smallest number of iterations.

We have observed this behaviour of the seed GMRES method on some other difficult problems. Intuitively, it seems to us that an analogy exists between this behaviour and the observed stagnation of the classical restarted GMRES method. In the two cases, an initial guess is extracted from a Krylov space to generate a new Krylov space. As illustrated in the previous section, one possible remedy for the restarted GMRES method is to replace M_{Frob} by $M_{\text{SLRU}(k)}$.

In Table X, we investigate this possibility. We keep the same strategies as in Table IX but now use $M_{\text{SLRU}(20)}$ rather than M_{Frob} . The initial guess computed by the seed GMRES method becomes the best performing strategy. In Figure 9, for each right-hand side, $\phi = 15^\circ : 1^\circ : 25^\circ$, we plot the convergence history of the seed GMRES method with the two preconditioners. It can be seen that although the norm of the initial residuals are about the same for the two preconditioners, the rate of convergence is significantly improved by $M_{\text{SLRU}(20)}$. The seed GMRES method provides a small initial residual and $M_{\text{SLRU}(k)}$ ensures a fast rate of convergence of GMRES iterations: in a race starting close to the finish (seed strategy) and running fast ($M_{\text{SLRU}(k)}$ preconditioner) this ensures we finish first!




Figure 10. Influence of the multipole precision (prec) in ARPACK on the iterations on a difficult range for the Cobra test problems; the most accurate precision is prec 3.

8. CONCLUSIONS AND PROSPECTIVES

In this paper, we have shown that the $M_{\text{SLRU}(k)}$ preconditioner based on the M_{Frob} preconditioner is effective for solving large linear systems arising in challenging real-life electromagnetic applications. It is important to point out that, to be effective, the spectral low rank update should be built on top of a good enough preconditioner that already succeeds in clustering most of the eigenvalues close to a point far from the origin (one in our case). There are two main reasons that motivate this observation. Firstly, if only a few eigenvalues are left close to the origin, a small rank update will be sufficient to significantly improve the convergence. Secondly, these few isolated eigenvalues will be easily found by an eigensolver. Of course a trade-off between the two components of $M_{\text{SLRU}(k)}$ should be found as the low rank update might be unnecessary if M_{Frob} is very dense, or it might be too expensive to improve a poor M_{Frob} because too many eigenvalues, potentially difficult to compute, have to be shifted.

We observe that the convergence of GMRES using $M_{\text{SLRU}(k)}$ only weakly depends on the choice of the initial guess. This is particularly useful in the seed GMRES or restarted GMRES contexts.

When several right-hand sides have to be solved with the same coefficient matrix, the extra cost of the eigencalculation preprocessing phase is quickly amortized by the reduction in the number of iterations as the extra cost per iteration is negligible. However, the cost of the

preprocessing phase can be decreased in different ways. The first approach would consist in constructing the preconditioner as we solve different right-hand sides by extracting the spectral information from previous GMRES solutions. Another possibility, as the accuracy of the eigenvectors in terms of backward error does not need to be very good, would be to still compute them in a preprocessing phase but with a less accurate FMM. Implementing this idea for the Cobra example leads to the plot reported in Figure 10. In that figure, it can be seen that using either a high or the medium accuracy FMM leads to the same quality of the $M_{\text{SLRU}(k)}$ preconditioner, while using a low accuracy only deteriorates the efficiency by less than 20% in terms of iteration number. Consequently, using a low accurate FMM for the eigencalculation might also be a way of reducing the cost for the preprocessing phase.

A question that remains open is the *a priori* identification of the optimal number of eigenvalues to be shifted. We have seen that increasing this number is always beneficial, but the relative gain tends to vanish when this number becomes large. If the eigenvalue information was extracted at run-time, a possible strategy might be to increase the size of the rank from one right-hand side to the next as long as an improvement is observed. Further investigations deserve to be undertaken to study this possibility in the framework of the monostatic radar cross-section calculations.

ACKNOWLEDGEMENTS

We would like to thank Guillaume Alléon from EADS-CCR, for providing us some test examples considered in this paper. We wish to express our gratitude to Guillaume Sylvand from CERMICS-INRIA (CERMICS: mathematics, computer science and scientific computing teaching and research center; INRIA: national institute for research in computer science and control), for his friendly and permanent support and his advices in using the FMM components of the code. J. Langou was supported by EADS, Corporate Research Centre, Toulouse.

REFERENCES

1. Darve E. The fast multipole method (I): error analysis and asymptotic complexity. *SIAM Journal on Numerical Analysis* 2000; **38**(1):98–128.
2. Darve E. The fast multipole method: numerical implementation. *Journal on Computational Physics* 2000; **160**(1):195–240.
3. Sylvand G. Résolution Itérative de Formulation Intégrale pour Helmholtz 3D: Applications de la Méthode Multipôle à des Problèmes de Grande Taille. *Ph.D. Thesis*, Ecole Nationale des Ponts et Chaussées, 2002.
4. Alléon G, Benzi M, Giraud L. Sparse approximate inverse preconditioning for dense linear systems arising in computational electromagnetics. *Numerical Algorithms* 1997; **16**:1–15.
5. Carpentieri B. Sparse preconditioners for dense linear systems from electromagnetic applications. *Ph.D. Thesis*, CERFACS, Toulouse, France, 2002.
6. Carpentieri B, Duff IS, Giraud L. Sparse pattern selection strategies for robust Frobenius-norm minimization preconditioners in electromagnetism. *Numerical Linear Algebra with Applications* 2000; **7**(7–8):667–685.
7. Carpentieri B, Duff IS, Giraud L, Magolu monga Made M. Sparse symmetric preconditioners for dense linear systems in electromagnetism. *Technical Report TR/PA/01/35*, CERFACS, Toulouse, France, 2001 (Preliminary version of the paper to appear in *Numerical Linear Algebra with Applications*).
8. Benzi M, Meyer CD, Tūma M. A sparse approximate inverse preconditioner for the conjugate gradient method. *SIAM Journal on Scientific Computing* 1996; **17**:1135–1149.
9. Benzi M, Tūma M. A sparse approximate inverse preconditioner for nonsymmetric linear systems. *SIAM Journal on Scientific Computing* 1998; **19**:968–994.
10. Grote M, Huckle T. Parallel preconditionings with sparse approximate inverses. *SIAM Journal on Scientific Computing* 1997; **18**:838–853.

11. Carpentieri B, Duff IS, Giraud L, Sylvand G. Combining fast multipole techniques and an approximate inverse preconditioner for large parallel electromagnetics calculations. *Technical Report TR/PA/03/77*, CERFACS, Toulouse, France, 2003.
12. Carpentieri B, Duff IS, Giraud L. A class of spectral two-level preconditioner. *SIAM Journal on Scientific Computing* 2003; **25**(2):749–765.
13. Saad Y, Schultz MH. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing* 1986; **7**:856–869.
14. Baglama J, Calvetti D, Golub GH, Reichel L. Adaptively preconditioned GMRES algorithms. *SIAM Journal on Scientific Computing* 1999; **20**(1):243–269.
15. Erhel J, Burrage K, Pohl B. Restarted GMRES preconditioned by deflation. *Journal of Computational and Applied Mathematics* 1996; **69**:303–318.
16. Kharchenko SA, Yu Yeregin A. Eigenvalue translation based preconditioners for the GMRES(k) method. *Numerical Linear Algebra with Applications* 1995; **2**(1):51–77.
17. Lehoucq RB, Sorensen DC, Yang C. *ARPACK User's Guide: Solution of Large-Scale Problem with Implicitly Restarted Arnoldi Methods*. SIAM: Philadelphia, 1998.
18. Freund RW, Malhotra M. A block QMR algorithm for non-Hermitian linear systems with multiple right-hand sides. *Linear Algebra and its Applications* 1997; **254**(1–3):119–157.
19. Langou J. Iterative methods for solving linear systems with multiple right hand sides. *Ph.D. Thesis*, Institut National des Sciences Appliquées de Toulouse, 2003.
20. O'Leary DP. The block conjugate gradient algorithm and related methods. *Linear Algebra and its Applications* 1980; **29**:293–322.
21. Vital B. Etude de quelques méthodes de résolution de problèmes linéaires de grande taille sur multiprocesseur. *Ph.D. Dissertation*, Université de Rennes, 1990.