A class of spectral two-level preconditioners[1]

Bruno Carpentieri[2], Iain S. Duff[3] and Luc Giraud[4]

ABSTRACT

When solving the linear system $Ax = b$ with a Krylov method, the smallest eigenvalues of the matrix A often slow down the convergence. This is usually still the case even after the system has been preconditioned. Consequently if the smallest eigenvalues of A could be somehow "removed" the convergence would be improved. Several techniques have been proposed in the past few years that attempt to tackle this problem. The proposed approaches can be split into two main families depending on whether the scheme enlarges the generated Krylov space or adaptively updates the preconditioner. In this paper, we follow the second approach and propose a class of preconditioners both for unsymmetric and for symmetric linear systems that can also be adapted for symmetric positive definite problems. Our preconditioners are particularly suitable when there are only a few eigenvalues near the origin that are well separated. We show that our preconditioners shift these eigenvalues from close to the origin to near one. We illustrate the performance of our method through extensive numerical experiments on a set of general linear systems. Finally we show the advantages of the preconditioners for solving dense linear systems arising in electromagnetism applications that were the main motivation for this work.

Keywords: preconditioning iterative methods, convergence of iterative method, sparse complex systems, electromagnetics application, GMRES.

AMS(MOS) subject classifications: 65F05, 65F50.

---

[1]Current reports available by anonymous ftp to ftp.numerical.rl.ac.uk in directory pub/reports. This report is available in compressed postscript as file cadgRAL2002020.ps.gz or as the PDF file cadgRAL2002020.pdf. The report is also available through URL http://www.numerical.rl.ac.uk/reports/reports.html. Also published as Technical Report TR/PA/02/55 from CERFACS.

[2]b.carpentieri@dm.uniba.it, CERFACS, 42 Ave G. Coriolis, 31057 Toulouse Cedex, France. The work of this author was supported by I.N.D.A.M. (Rome, Italy)under a grant (Borsa di Studio per l'estero, Provvedimento del Presidente del 30 Aprile 1998).

[3]i.s.duff@rl.ac.uk, the work of this author was supported in part by the EPSRC Grant GR/R46441.

[4]giraud@cerfacs.fr, CERFACS, 42 Ave G. Coriolis, 31057 Toulouse Cedex, France.

Computational Science and Engineering Department
Atlas Centre
Rutherford Appleton Laboratory
Oxon OX11 0QX
June 5, 2002

Contents

1 Introduction

The starting point for this work was the iterative solution of linear systems that arise in electromagnetism applications. In the last few years we have studied preconditioning techniques based on sparse approximate inverses and have found them to be quite effective. These preconditioners are able to cluster most of the eigenvalues close to one but still leave a few close to the origin that are difficult to remove by tuning the parameter that controls our preconditioner (Alléon, Benzi and Giraud 1997, Carpentieri 2002, Carpentieri, Duff and Giraud 2000, Carpentieri, Duff, Giraud and Sylvand 2002). This is a fairly common situation for a wide range of problems and preconditioners. We address this in a more general context in this article even though we present some results for that particular application in a later section of this paper.

It is well known that, when solving the linear system $Ax = b$ with a Krylov method, the smallest eigenvalues of the matrix A often slow down the convergence. In the symmetric positive definite (SPD) case, this is clearly highlighted by the bound on the rate of convergence of the Conjugate Gradient method (CG) given by Golub and Loan (1996) viz.

$$||e^{(k)}||_A \leq 2 \cdot (\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1})^k ||e^{(0)}||_A,$$

where $e^{(n)} = x^* - x^{(n)}$ denotes the error associated with the iterate at step k and $\kappa(A) = \dfrac{\lambda_{max}}{\lambda_{min}}$ denotes the condition number. From this bound, it can be seen that increasing the size of the smallest eigenvalues will improve the convergence rate of CG. Consequently if the smallest eigenvalues of A could be somehow "removed" the convergence of CG will be improved (Kolotilina 1995, Mansfield 1988, Mansfield 1991, Nicolaides 1987). Similar arguments exist for unsymmetric systems to mitigate the bad effect of the smallest eigenvalues on the rate of convergence of the unsymmetric Krylov solver (Baglama, Calvetti, Golub and Reichel 1999, de Sturler 1996, Erhel, Burrage and Pohl 1996, Morgan 1995). The main argument is that the Krylov methods build a polynomial expansion that should be equal to one when the argument is zero and whose roots are the eigenvalues. To get fast convergence it is necessary to find a low order polynomial with these properties (for example, strategies have been developed to improve the convergence of GMRES (Saad and Schultz 1986)). Clearly the presence of eigenvalues close to the origin makes this difficult.

In exact arithmetic, the number of distinct eigenvalues determines the maximum dimension of the Krylov subspace. If the eigenvalues are not distinct but the diameters of the clusters are small enough, the eigenvalues within each cluster behave numerically like a single eigenvalue, and we expect a few iterations of a Krylov method to produce reasonably accurate approximations. Theoretical studies

have related superlinear convergence of GMRES to the convergence of the Ritz values (van der Vorst and Vuik 1993). Basically, convergence occurs if, at each iteration of GMRES, the next smallest eigenvalue in magnitude is removed from the system. As the restarting procedure destroys information about the Ritz values at each restart, the superlinear convergence may be lost. Thus removing the effect of small eigenvalues in the preconditioned matrix can have a beneficial effect on the convergence.

When GMRES is the Krylov solver there are essentially two different approaches for exploiting information related to the smallest eigenvalues. The first idea is to compute a few, k say, approximate eigenvectors of MA corresponding to the k smallest eigenvalues in magnitude, and augment the Krylov subspace with those directions. At each restart, let $u_1, u_2, ..., u_k$ be approximate eigenvectors corresponding to the approximate eigenvalues of MA closest to the origin. The updated solution of the linear system in the next cycle of GMRES is extracted from $Span\{r_0, Ar_0, A^2r_0, A^3r_0, ..., A^{m-k-1}r_0, u_1, u_2, ..., u_k\}$. This approach is referred to as the *augmented subspace approach* (see Morgan (2000), Morgan (1995), Saad (1993)). The approximate eigenvectors can be chosen to be Ritz vectors from the Arnoldi process. The standard implementation of the restarted GMRES algorithm is based on the Arnoldi process, and spectral information of MA might be recovered during the iterations. The second idea exploits spectral information gathered during the Arnoldi process to determine an approximation of an invariant subspace of A associated with the eigenvalues nearest the origin, and uses this information to construct a preconditioner or to update the preconditioner. The idea of using exact invariant subspaces to improve the eigenvalue distribution was proposed by Saad (1988). Information from the invariant subspace associated with the smallest eigenvalues and its orthogonal complement are used to construct a preconditioner in the approach proposed in Baglama et al. (1999). This information can be obtained from the Arnoldi decomposition of a matrix A of size n that has the form

$$AV_m = V_m H_m + f_m e_m^T$$

where $V_m \in R^{n\times m}, f_m \in R^n, V_m^T V_m = I_m, V_m^T f_m = 0$, and $H_m \in R^{m\times m}$ is an upper Hessenberg matrix. If the Arnoldi process is started from $V_m e_1 = r_0/\|r_0\|$, the columns of V_m span the Krylov subspace $K_m(A, r_0)$. Let the matrix $V_k \in R^{k\times n}$ consist of the first k columns $v_1, v_2, ..., v_k$ of V_m, and let the columns of the orthogonal matrix W_{n-k} span the orthogonal complement of $Span\{v_1, v_2, ..., v_k\}$. As $W_{n-k}^T W_{n-k} = I_{n-k}$, the columns of the matrix $[V_k\, W_{n-k}]$ form an orthogonal basis of R^n. Baglama et al. (1999) use the inverse of the matrix

$$M = V_k H_k V_k^T + W_{n-k} W_{n-k}^T$$

as a left preconditioner. It can be expressed as:

$$M^{-1} = V_k H_k^{-1} V_k^T + W_{n-k} W_{n-k}^T.$$

At each restart, the preconditioner is updated by extracting new eigenvalues which are the smallest in magnitude. The algorithm proposed uses the recursion formulae of the implicitly restarted Arnoldi (IRA) method described by Sorensen (1992), and the determination of the preconditioner does not require the evaluation of any matrix-vector products with the matrix A in addition to those needed for the Arnoldi process.

Another adaptive procedure to determine a preconditioner during GMRES iterations was introduced by Erhel et al. (1996). It is based on the same idea of estimating the invariant subspace corresponding to the smallest eigenvalues. The preconditioner is based on a deflation technique such that the linear system is solved exactly in an invariant subspace of dimension r corresponding to the smallest r eigenvalues of A.

Finally, a preconditioner for GMRES based on a sequence of rank-one updates that involve the left and right smallest eigenvectors is proposed by Kharchenko and Yeremin (1995). The method is based on the idea of translating isolated eigenvalues consecutively group by group into a vicinity of one using low-rank projections of the coefficient matrix of the form

$$\tilde{A} = A \cdot (I_n + u_1 v_1^H) \cdot \ldots \cdot (I_n + u_l v_l^H).$$

The vectors u_j and v_j, $j \in [1, l]$ are determined to ensure the numerical stability of consecutive translations of groups of isolated eigenvalues of $\tilde{A}$. After each restart of GMRES(m), approximations to the isolated eigenvalues to be translated are computed by the Arnoldi process. The isolated eigenvalues are translated towards one, and the next cycle of GMRES(m) is applied to the transformed matrix. The effectiveness of this method relies on the assumption that most of the eigenvalues of A are clustered close to one in the complex plane.

Most of these schemes are combined with the GMRES procedure as they derive information directly from its internal Arnoldi process. In our work, we consider an additional explicit eigencomputation that is used to update the selected preconditioner. This makes the preconditioner independent of the Krylov solver used for the actual solution of the linear system. This extra cost will be overcome if the same linear system with several right-hand sides has to be solved because the number of Krylov iterations can be significantly reduced.

The paper is organized as follows. In the following section, we describe the proposed preconditioners and prove their shifting capabilities on diagonalizable matrices. In Section 3, we illustrate the numerical efficiency of the proposed scheme on a set of unsymmetric and SPD linear systems from the Harwell-Boeing collection (Duff, Grimes and Lewis 1992). We devote Section 4 to a particular application in electromagnetism where the same linear system has to be solved with many different right-hand sides. This situation is particularly of interest for the preconditioners we propose as it enables us to amortize the extra eigencomputation required. Finally, we conclude with some remarks in Section 5.

2 Two-level preconditioner via low-rank update

Many of the preconditioners proposed in the literature succeed in clustering most of the eigenvalues of the preconditioned matrix MA (for left preconditioning) far from the origin. Such a distribution is highly desirable to get fast convergence of Krylov solvers. However, a few eigenvalues can be left close to zero and they potentially can significantly degrade the convergence. In order to tackle this difficulty we propose a refinement technique based on the introduction of low-rank corrections computed from spectral information associated with the smallest eigenvalues of MA. Roughly speaking, the proposed technique consists in solving exactly the preconditioned system in the low dimensional space spanned by the eigenvectors associated with the eigenvalues closest to the origin. This is then used to update the preconditioned residual. We first present our technique for unsymmetric linear systems and then derive a variant for symmetric and symmetric positive definite matrices. For simplicity, we first consider complex linear systems. We later indicate how it can be adapted for problems in real arithmetic.

We consider the solution of the linear system

$$Ax = b, \tag{2.1}$$

where A is a $n \times n$ unsymmetric complex nonsingular matrix, and x and b are vectors of size n. The linear system is solved using a preconditioned Krylov solver and we denote by M_1 the left preconditioner, meaning that we solve

$$M_1 Ax = M_1 b. \tag{2.2}$$

We assume that the preconditioned matrix $M_1 A$ is diagonalizable, that is:

$$M_1 A = V \Lambda V^{-1}, \tag{2.3}$$

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 $u_i^H v_i \neq 0, \forall i$ (Wilkinson 1965). 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.

Proposition 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 diagonalizable and we have $MA = V diag(\eta_i) V^{-1}$ with*

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

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

$$\begin{aligned} MAV_\varepsilon &= M_1 AV_\varepsilon + V_\varepsilon A_c^{-1} U_\varepsilon^H M_1 AV_\varepsilon \\ &= V_\varepsilon D_\varepsilon + V_\varepsilon I_k \\ &= V_\varepsilon (D_\varepsilon + I_k) \end{aligned}$$

where I_k denotes the $(k \times k)$ identity matrix, and

$$\begin{aligned} MAV_{\bar{\varepsilon}} &= M_1 AV_{\bar{\varepsilon}} + V_\varepsilon A_c^{-1} U_\varepsilon^H M_1 AV_{\bar{\varepsilon}} \\ &= V_{\bar{\varepsilon}} D_{\bar{\varepsilon}} + V_\varepsilon A_c^{-1} U_\varepsilon^H V_{\bar{\varepsilon}} D_{\bar{\varepsilon}} \\ &= V_{\bar{\varepsilon}} D_{\bar{\varepsilon}} \qquad \text{since } U_\varepsilon^H V_{\bar{\varepsilon}} = 0. \end{aligned}$$

We then have

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

■

Proposition 2 *Let W be such that $\tilde{A}_c = W^H AV_\varepsilon$ 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:

$$\begin{aligned} \tilde{M}AV_\varepsilon &= M_1 AV_\varepsilon + V_\varepsilon A_c^{-1} W^H AV_\varepsilon \\ &= V_\varepsilon D_\varepsilon + V_\varepsilon I_k \\ &= V_\varepsilon (D_\varepsilon + I_k) \\ \tilde{M}AV_{\bar{\varepsilon}} &= M_1 AV_{\bar{\varepsilon}} + V_\varepsilon A_c^{-1} W^H AV_{\bar{\varepsilon}} \\ &= V_{\bar{\varepsilon}} D_{\bar{\varepsilon}} + V_\varepsilon C \qquad \text{with } C = A_c^{-1} W^H AV_{\bar{\varepsilon}} \\ &= \begin{pmatrix} V_\varepsilon V_{\bar{\varepsilon}} \end{pmatrix} \begin{pmatrix} C \\ D_{\bar{\varepsilon}} \end{pmatrix} \end{aligned}$$

We then 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.

Proposition 3 *Let $A_c = U_\varepsilon^H AM_1 V_\varepsilon$, $M_c = M_1 V_\varepsilon A_c^{-1} U_\varepsilon^H$ and $M = M_1 + M_c$. Then AM is diagonalizable and we have $AM = V diag(\eta_i) V^{-1}$ with*

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

Proposition 4 *Let W be such that $\tilde{A}_c = W^H A M_1 V_\varepsilon$ has full rank, $\tilde{M}_c = M_1 V_\varepsilon \tilde{A}_c^{-1} W^H$ and $\tilde{M} = M_1 + \tilde{M}_c$. Then $A\tilde{M}$ is similar to a matrix whose eigenvalues are*

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

We should point out that, if the symmetry of the preconditioner has to be preserved, an obvious choice exists. For left preconditioning, we can set $W = V_\varepsilon$, but then $\tilde{A}_c$ may not have full rank. In the SPD case, these results extend as follows and lead to an expression that is similar to those proposed by Carvalho, Giraud and Tallec (2001) for two-level preconditioners in non-overlapping domain decomposition.

Proposition 5 *If A and M_1 are SPD, then $M_1 A$ is diagonalizable, and $\tilde{A}_c = V_\varepsilon^H A V_\varepsilon$ is SPD. The preconditioner defined by $\tilde{M} = M_1 + \tilde{M}_c$, with $\tilde{M}_c = V_\varepsilon \tilde{A}_c^{-1} V_\varepsilon^H$ is SPD and $\tilde{M}A$ 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

Because the matrix M_1 is SPD, there exists a unique SPD matrix $M_1^{\frac{1}{2}}$ that is the square root of M_1 (see for instance, Golub and Loan (1996)). Then, the matrix $M_1 A$ is similar to the matrix $M_1^{\frac{1}{2}} A M_1^{\frac{1}{2}}$ which is symmetric and consequently similar to a diagonal matrix. Therefore the matrix $M_1 A$ is diagonalizable.

By construction $\tilde{A}_c$ is symmetric, let us show that it is positive definite. V_ε is a $n \times k$ matrix. Let $z \in \mathbb{R}^k,\ z \neq 0$.

$$\begin{aligned} < \tilde{A}_c z, z > &= < V_\varepsilon^H A V_\varepsilon z, z > \\ &= < A V_\varepsilon z, V_\varepsilon z > \end{aligned}$$

$V_\varepsilon z \neq 0$ because V_ε has full rank. Then $< \tilde{A}_c z, z >$ is greater than 0 because A is SPD. Therefore $\tilde{A}_c$ is a SPD matrix and consequently has full rank.
Let $x \in \mathbb{R}^n,\ x \neq 0$.

$$\begin{aligned} < \tilde{M}_c x, x > &= < V_\varepsilon \tilde{A}_c^{-1} V_\varepsilon^H x, x > \\ &= < \tilde{A}_c^{-1} V_\varepsilon^H x, V_\varepsilon^H x > \\ &\geq 0 \text{ as } \tilde{A}_c \text{ is a SPD matrix.} \end{aligned}$$

Therefore $\tilde{M}_c$ is a positive semi-definite matrix and $\tilde{M} = M_1 + \tilde{M}_c$ is an SPD matrix because M_1 is SPD, and the results of Proposition 2 hold with $W = V_\varepsilon$.

■

For unsymmetric linear systems in real arithmetic some of the eigenvectors can be complex. If implemented as described so far, the preconditioner would be complex which is not desirable because all the calculations would have to be performed in complex arithmetic. If a complex eigenvector exists its conjugate is also an eigenvector. So the drawback just described can be overcome by considering not just the eigenvectors but a real basis of the plane spanned by those two conjugate eigenvectors; that are the two real vectors defined by the real part and the imaginary part of those vectors.

Finally we mention that we can use an additional scaling in the low-rank update so that the k smallest eigenvalues are not just shifted by one, but rather are all transformed to one with multiplicity equal to k. This feature is obtained by using $M_c = V_\varepsilon(I - D_\varepsilon)A_c^{-1}U_\varepsilon^H M_1$ in Proposition 1, and $\tilde{M}_c = V_\varepsilon(I - D_\varepsilon)\tilde{A}_c^{-1}W^H$ in Proposition 2. Similar transformations can be applied to get the same property for right preconditioning. This does not change the numerical behaviour of the method although it makes the expression of the preconditioner slightly more complicated. For those reasons we do not develop this variant further.

3 Numerical experiments

In order to illustrate the efficiency of the preconditioners, we first present numerical experiments on general linear systems that are either unsymmetric or symmetric indefinite. Then we consider SPD linear systems to assess the effectiveness of the preconditioners on those problems as well.

3.1 Non-Hermitian linear systems

In Table 3.1 we display the list of test problems from the Harwell-Boeing collection that we have considered for the experiments on general matrices. All the experiments have been performed in Matlab using $ILU(t)$ (Saad 1994) as the preconditioner M_1. The stopping criterion in all cases just consists in reducing the original residual by 10^{-6} that then can be related to a normwise backward error as we use the null vector as initial guess. In all the tables, the symbol "-" means that convergence is not obtained after 1000 iterations. The eigenvectors are computed using the Matlab function `eigs` that calls ARPACK (Lehoucq, Sorensen and Yang 1998).

Name	Size	Field	Characteristics
HOR131	434	Flow in networks	real unsymmetric
ORSIRR1	1030	Oil reservoir simulation	real unsymmetric
GRE1107	1107	Simulation studies in computer systems	real unsymmetric
YOUNG2C	841	Dynamic analysis in structural engineering	complex symmetric indefinite

Table 3.1: Set of non-Hermitian test matrices.

In Table 3.2, we show the number of iterations required by restarted GMRES and BiCGStab (van der Vorst 1992) varying the dimension of the low-rank correction in the range 1 to 10. The choice of the threshold for ILU has been set to illustrate the behaviour generally observed when the spectrum of the preconditioned system has only a few eigenvalues close to the origin. That is, when the preconditioner is already effective in clustering most of the eigenvalues. In Figure 3.1, we display the spectrum of the preconditioned matrices using only $ILU(t)$. For the experiments shown in Table 3.2, we use a left preconditioner and the formulation described in Proposition 1 that is $W^H = U_\varepsilon^H M_1$. Similar results are displayed in Table 3.3 using the formulation described in Proposition 2 that is with $W = V_\varepsilon$. In this latter case, the cost for the eigencomputation to setup the update is halved because only right eigenvectors need to be computed. As expected, it can be immediately seen that, for these two choices of W^H, the numerical trends are the same; that is, the larger the rank of the correction the faster the convergence. However, the decrease is not always monotonic with the dimension of the rank correction. This behaviour is observed both for GMRES and BiCGStab.

As can be seen, a correction of rank ten enables us to half the number of iterations in most of the cases. In general, a very small dimension correction (i.e. one or two) improves the convergence of the Krylov solver significantly. In some cases, it even enables convergence that was otherwise not obtained. This situation is illustrated in Figure 3.2 where we display the convergence history of GMRES(30) on the Grenoble test problem, GRE1107. Without correction the backward error stagnates (as well as with a correction of dimension one or two) and the convergence is obtained only for a correction of dimension larger than 3. Even though this aspect is discussed later, we can already observe the link that exists between the dimension of the update and the size of the restart of GMRES since GMRES(40) converges with a rank-one update (see Table 3.2) while GMRES(30) does not (see Figure 3.2). This linear system has a very small eigenvalue that is fairly isolated and that prevents the convergence of the solvers. Once this eigencomponent is removed by the rank-one update preconditioner both GMRES(40) and BiCGStab converge.

			Dimension of the small dimensional correction space										
Matrix	t	Solver	0	1	2	3	4	5	6	7	8	9	10
HOR131	$4 \cdot 10^{-2}$	GMRES(5)	106	76	66	60	61	58	54	43	41	39	39
		BiCGStab	19	14	12	13	12	10	10	9	9	8	8
ORSIRR1	$5 \cdot 10^{-2}$	GMRES(5)	95	88	77	71	68	65	62	59	55	54	50
		BiCGStab	28	29	24	26	22	21	18	18	18	16	16
GRE1107	$1 \cdot 10^{-2}$	GMRES(40)	-	159	120	75	67	39	37	35	34	32	30
		BiCGStab	-	80	80	61	63	51	58	44	33	29	27
YOUNG2C	$7 \cdot 10^{-2}$	GMRES(30)	-	535	494	478	330	316	299	322	299	298	293
		BiCGStab	76	60	57	58	46	46	47	47	47	46	46

Table 3.2: Number of iterations varying the dimension of the low-rank update with $W^H = U_\varepsilon^H M_1$.

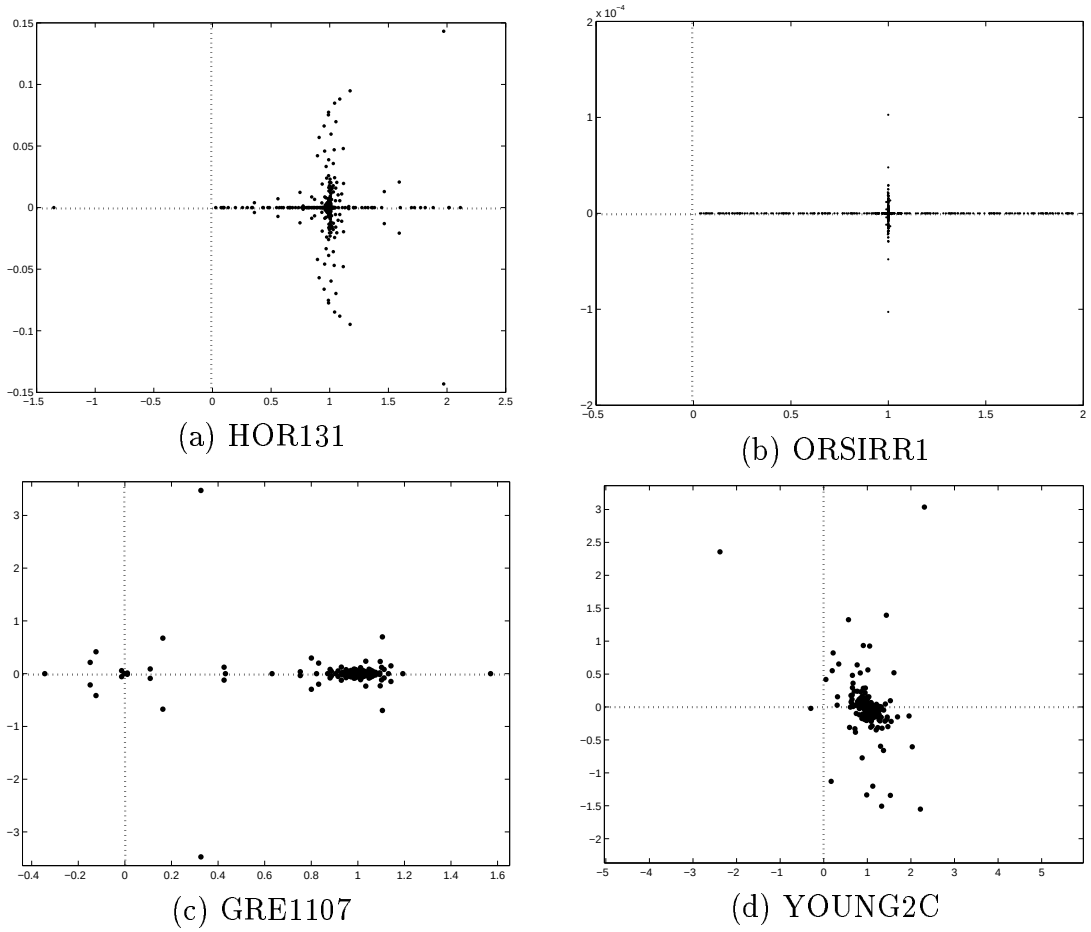

(a) HOR131 (b) ORSIRR1

(c) GRE1107 (d) YOUNG2C

Figure 3.1: The spectrum of the preconditioned matrices.

			Dimension of the small dimensional correction space										
Matrix	t	Solver	0	1	2	3	4	5	6	7	8	9	10
HOR131	$4 \cdot 10^{-2}$	GMRES(5)	106	71	62	64	59	60	53	44	46	44	41
		BiCGStab	19	14	13	13	11	10	10	9	9	9	9
ORSIRR1	$5 \cdot 10^{-2}$	GMRES(5)	95	88	77	71	68	65	62	59	55	54	50
		BiCGStab	28	28	24	24	21	22	20	18	18	17	16
GRE1107	$1 \cdot 10^{-2}$	GMRES(40)	-	160	88	76	62	39	60	36	58	57	57
		BiCGStab	-	79	87	74	84	82	56	47	65	67	40
YOUNG2C	$7 \cdot 10^{-2}$	GMRES(30)	-	562	497	474	358	297	297	299	301	297	296
		BiCGStab	76	59	55	52	46	44	46	45	47	46	46

Table 3.3: Number of iterations varying the dimension of the low-rank update with $W = V_\varepsilon$.

To illustrate that the proposed updates should be used to improve an already effective preconditioner, we report in Table 3.4 the number of iterations when the threshold of $ILU(t)$ is relaxed making the original preconditioner less and less

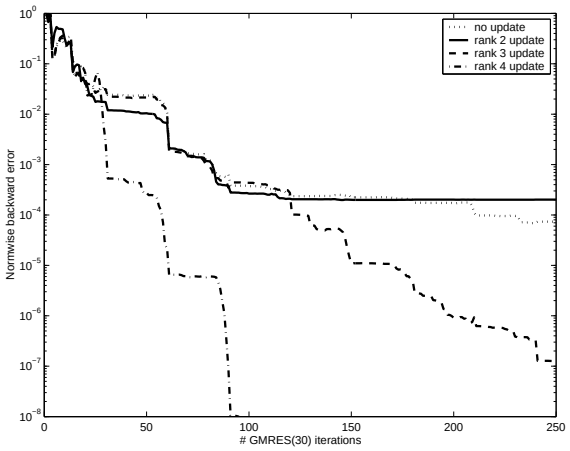

Figure 3.2: Convergence history varying the dimension of the low-rank correction for GRE1107 and $ILU(1 \cdot 10^{-2})$.

efficient. We see that, in that case, the update will only improve the convergence up to a certain level above which it does not have any effect. That corresponds to the situation where there are many eigenvalues close to zero and shifting a few of them does not further affect the convergence.

	# GMRES(5) iterations	
t	Without update	With update
$6 \cdot 10^{-2}$	150	150
$5 \cdot 10^{-2}$	95	93
$4 \cdot 10^{-2}$	106	58
$3 \cdot 10^{-2}$	85	55
$2 \cdot 10^{-2}$	52	32
$1 \cdot 10^{-2}$	31	18

Table 3.4: Number of GMRES(5) iterations varying the threshold for a low-rank update of dimension 5 for the matrix ORSIRR1.

As mentioned earlier, theoretical studies have related superlinear convergence of GMRES to the convergence of Ritz values (van der Vorst and Vuik 1993). Basically, convergence occurs if, at each iteration of GMRES, the next smallest eigenvalue in magnitude is removed from the system. As the restarting procedure destroys information about the Ritz values at each restart, the superlinear convergence may be lost. Thus removing the effect of small eigenvalues in the preconditioned matrix can have a beneficial effect on the convergence. This has been observed in the experiments that we have performed. However, this theoretical result can also be

read differently. That is, if the small eigenvalues are removed, the restart value for GMRES might no longer be critical and the convergence should not be much affected by the choice of the restart. We can see this in Figure 3.3 where we show, for different choices of the restart parameter, the number of GMRES iterations as a function of the dimension of the low-rank correction. It can be seen the number of iterations with all the restarts tend to behave as full-GMRES as the dimension of the update increases.

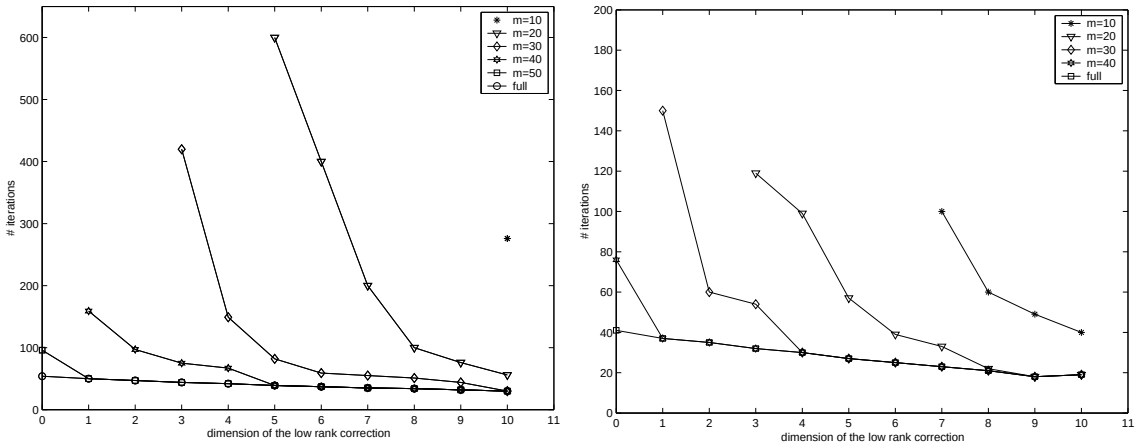

(a) $ILU(1 \cdot 10^{-2})$ (b) $ILU(7 \cdot 10^{-3})$

Figure 3.3: Sensitivity to the restart parameter of GMRES. GRE1107 test problem.

3.2 Symmetric positive definite linear systems

In this section we illustrate, on the set of SPD matrices listed in Table 3.5, the SPD variant of the update presented in Proposition 5. As a preconditioner we use $IC(t)$ (Meijerink and van der Vorst 1977).

Name	Size	Field
BCSSTK27	1224	Dynamic analyses in structural engineering - Buckling analysis
BCSSTK14	1806	Static analyses in structural engineering - Roof of the Omni Coliseum, Atlanta
BCSSTK16	4884	Static analyses in structural engineering - U.S. Army Corps of Engineers dam
S1RMQ4M1	5489	Structural mechanics - Cylindrical shell

Table 3.5: Set of SPD test matrices.

We observe a similar improvement for SPD linear systems to what was seen in the previous section. This is illustrated in Table 3.6 where we show the number of CG iterations as we vary the dimension of the positive semi-definite update. To show that the improvement of the update is not too closely related to the quality of the initial preconditioner we show, for BCSSTK27 and S1RMQ4M1, the number of iterations for two different thresholds for IC. For these examples, the relative gain

in terms of number of iterations does not depend much on the selected threshold although the absolute values do.

		Dimension of the small dimensional correction space										
Matrix	t	0	1	2	3	4	5	6	7	8	9	10
BCSSTK27	$5 \cdot 10^{-2}$	63	61	53	52	45	39	36	34	30	28	28
BCSSTK27	$1 \cdot 10^{-1}$	143	133	125	110	97	95	84	78	78	77	77
BCSSTK14	$5 \cdot 10^{-1}$	130	118	119	83	83	83	68	66	66	67	66
BCSSTK16	$5 \cdot 10^{-2}$	42	36	35	31	29	28	27	24	24	22	22
S1RMQ4M1	$5 \cdot 10^{-2}$	149	98	79	79	66	61	60	59	59	55	55
S1RMQ4M1	$1 \cdot 10^{-1}$	379	376	233	164	164	124	122	123	116	96	96

Table 3.6: Number of CG iterations varying the dimension of the low-rank update.

3.3 Sensitivity to the accuracy of the eigencomputation

As mentioned earlier, the eigenvalue calculation is performed in a pre-processing phase using ARPACK on the preconditioned matrix. When a set of isolated eigenvalues close to zero are computed, the backward error associated with the smallest ones is always the best. Even if we relax the stopping criterion, the smallest are still well computed. In order to investigate the sensitivity of the eigencomputation accuracy on the low-rank update improvement we would like to have a similar backward error on each eigenpair and to vary it. To do this, we compute the eigenpairs of a slightly perturbed matrix, $(M_1A + E)$, with $\frac{||E||}{||M_1A||} = \eta$, and we use these eigenvectors to build our preconditioners and compute the backward error of these eigenvectors as if they were eigenvectors of M_1A. By varying η, we can monitor the level of the backward error associated with each eigenvalue that then becomes comparable for each eigenvector.

In Table 3.7, we give the number of iterations of the Krylov solvers when varying the backward error of the computed eigenvectors. As we have one backward error per eigenvector, we give the average of them in the table. It can be seen that, in general, there is no need for very high accuracy in the computation of the eigenvectors. However, if some of the eigenvectors are ill-conditioned, even a small backward error might imply a large forward error and lead us to make a correction in the wrong space. Such a behaviour can be observed on the GRE1107 matrix.

4 A case study in electromagnetism applications

In recent years, there has been a significant amount of work on the simulation of electromagnetic wave propagation phenomena, addressing various topics ranging

GRE1107 - ILU($1 \cdot 10^{-2}$) - GMRES(40)											
Backward error	Dimension of the small dimensional correction space										
	0	1	2	3	4	5	6	7	8	9	10
$\approx 1 \cdot 10^{-15}$	-	80	78	37	35	32	30	28	26	24	23
$\approx 1 \cdot 10^{-5}$	-	155	97	77	40	39	38	37	35	33	31
$\approx 2 \cdot 10^{-5}$	-	134	134	79	79	75	61	40	39	38	36
$\approx 1 \cdot 10^{-4}$	-	-	-	-	-	440	160	116	114	95	80
HOR131 - ILU($5 \cdot 10^{-2}$) - GMRES(5)											
Backward error	Dimension of the small dimensional correction space										
	0	1	2	3	4	5	6	7	8	9	10
$\approx 1 \cdot 10^{-15}$	86	46	42	47	45	42	42	31	28	26	26
$\approx 3 \cdot 10^{-5}$	86	46	42	47	47	42	40	30	28	26	26
$\approx 5 \cdot 10^{-4}$	86	46	41	40	42	44	40	30	28	27	26
$\approx 3 \cdot 10^{-3}$	86	58	49	47	46	47	45	36	29	31	31
$\approx 3 \cdot 10^{-3}$	86	66	100	55	55	50	50	41	33	33	34
BCSSTK27 - IC($5 \cdot 10^{-2}$) - CG											
Backward error	Dimension of the small dimensional correction space										
	0	1	2	3	4	5	6	7	8	9	10
$\approx 1 \cdot 10^{-15}$	63	61	53	52	45	39	36	34	30	28	28
$\approx 2 \cdot 10^{-5}$	63	60	53	52	45	39	37	34	31	28	28
$\approx 1 \cdot 10^{-4}$	63	62	53	52	47	39	35	34	32	29	29
$\approx 2 \cdot 10^{-4}$	63	59	55	53	49	42	35	35	33	29	29
$\approx 1 \cdot 10^{-3}$	63	64	61	60	54	50	49	46	50	47	46
$\approx 2 \cdot 10^{-3}$	63	61	58	62	59	60	58	61	61	58	58

Table 3.7: Sensitivity of the preconditioner efficiency versus the accuracy of the eigencomputation.

from radar cross section to electromagnetic compatibility, to absorbing materials, and antenna design. To address these problems the Maxwell equations are often solved in the frequency domain leading to singular integral equations of the first kind. The discretization by the boundary element method (BEM) results in linear systems with dense complex matrices that are challenging to solve. The solution of these linear systems using iterative Krylov methods has recently become feasible thanks to a combination of the Fast Multipole Method (Sylvand 2002) and efficient preconditioners (Carpentieri 2002). In that framework, we have been working for the last few years on the design of approximate inverse preconditioners based on a Frobenius norm minimization with an *a priori* pattern selection strategy. We do not describe these preconditioners further but refer the reader to Alléon et al. (1997), Carpentieri (2002), and Carpentieri et al. (2000) for a detailed presentation.

The Frobenius-norm minimization preconditioner succeeds in clustering most of the eigenvalues far from the origin. This can be observed in Figure 4.1 where

we see a big cluster near one in the spectrum of the preconditioned matrix. This matrix is associated with a satellite that is particularly challenging to solve. The corresponding mesh, discretized with 1701 degrees of freedom, is displayed in Figure 4.2.

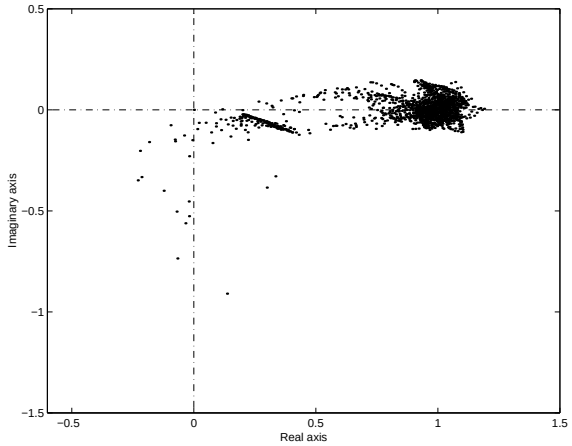

Figure 4.1: Eigenvalue distribution for the coefficient matrix preconditioned by the Frobenius-norm minimization method on the satellite problem.

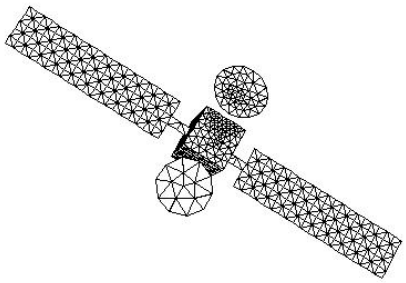

Figure 4.2: Mesh of a satellite with 1701 degrees of freedom.

The construction of the Frobenius-norm minimization preconditioner is inherently local. Each degree of freedom in the approximate inverse is coupled to only a very few neighbours and this compact support does not allow an exchange of global information. When the exact inverse is globally coupled, the lack of global information may have a severe impact on the quality of the preconditioner. Although the discrete Green's function in electromagnetic applications exhibits a

rapid decay, the exact inverse is dense and thus has global support. In that context, the use of the two-level spectral preconditioners seems appropriate. In addition, in such electromagnetism applications the same linear system has to be solved with many right-hand sides, when lighting an object with the same wave frequency but different incident angles. Depending on the object and the simulation, the number of right-hand sides can vary from a few to a few thousand. This situation is particularly suitable as the eigencomputation required to set up the correction can be compensated by the saving in iteration counts over the multiple right-hand sides.

All the numerical experiments are performed using a Fortran implementation in double precision complex arithmetic on a SGI Origin 2000. In these experiments, we consider low-rank updates of dimension up to 20, and different Krylov solvers. For this test problem, we perform experiments with two levels of accuracy in the GMRES solution to gain more insight into the robustness of our method. In that section, we show the qualitative numerical behaviour of our method on one test example that is representative of the general trend in electromagnetic applications (Carpentieri 2002). In Figure 4.3, we show the number of iterations required by GMRES(10) to reduce the normwise backward error to 10^{-8} and 10^{-5} for increasing dimension of the update. The numerical results show that the introduction of the low-rank updates can remarkably enhance the robustness of the approximate inverse. The plateau in Figure 4.3 that can be observed before a significant jump corresponds to a cluster of eigenvalues. When the eigenvalues within the cluster are shifted, a quick speedup of convergence is observed. By selecting up to 10 eigenpairs the number of iterations decreases by more than a factor of two on most of the experiments reported. The gain is more relevant in absolute value when high accuracy is required for the approximate solution, but remains almost constant in relative gain. As already observed on the other examples in the previous section, the preconditioning updates enable fast convergence of GMRES with a low restart within a tolerance of 10^{-8} whereas no convergence was obtained in 1500 iterations without updates. However, a substantial improvement in the convergence is observed also when low accuracy is required. In the most effective case, by selecting 10 corrections, the number of GMRES iterations needed to achieve convergence of 10^{-5} using low restarts reduces by more than a factor of two. If more eigenvectors are selected, generally no substantial improvement is observed.

Similarly to experiments reported in Section 3.1, we show, in Table 4.1, the number of iterations with two different choices for W. As expected, and already observed, with these two choices of W^H the numerical trends are the same; that is, the larger the rank of the correction the faster the convergence.

In Table 4.2, we show the number of matrix-vector products required by the ARPACK implementation of the IRA method to compute the smallest approximate eigenvalues and the associated approximate right eigenvectors (note that we do need the invert mode). We remark that the matrix-vector products do not include those required for the iterative solution. Although the computation can be expensive, the

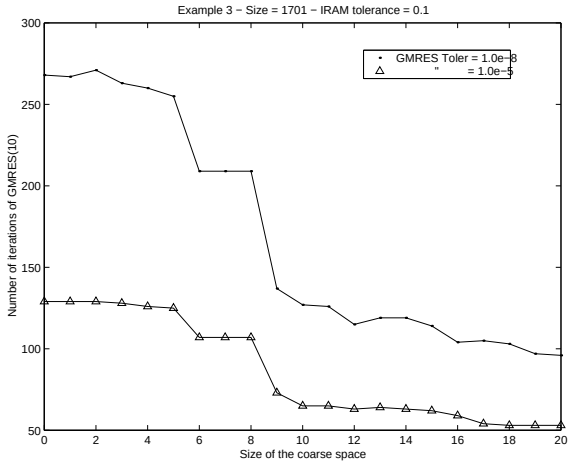

Figure 4.3: Number of iterations required by GMRES preconditioned by a Frobenius-norm minimization method updated with spectral corrections to reduce the normwise backward error by 10^{-8} and 10^{-5} for increasing number of corrections on the satellite.

cost can be amortized if the preconditioner is reused to solve linear systems with the same coefficient matrix and several right-hand sides. In the third column of this table we show the number of *amortization* vectors relative to GMRES(10) and a tolerance of 10^{-5}, that is the number of right-hand sides that have to be considered to amortize the extra cost for the eigencomputation. The localization of a few eigenvalues within a cluster may be more expensive than the computation of a full group of small eigenvalues. It can be seen that, for that example, the number of amortization vectors is reasonably small especially compared to real electromagnetic calculations where linear systems with the same coefficient matrix and up to thousands of right-hand sides are often solved.

In Figure 4.4 we display the number of iterations of SQMR (Freund and Nachtigal 1994) (QMR version for symmetric matrices with symmetric preconditioner). These experiments show that convergence of SQMR also benefits from the low-rank update. We also notice the remarkable robustness of this Krylov solver on electromagnetic applications; it clearly outperforms GMRES with large restart.

5 Concluding remarks

In this work, we consider a low-rank correction scheme that is particularly suited to improve a given preconditioner that leaves only few eigenvalues close to zero. The update of the preconditioner is beneficial to many Krylov solvers but requires an *a priori* eigencomputation that might be performed without too high an accuracy. Because the technique is used in combination with a first preconditioner, that already

Dimension of low-rank update	Choice for the operator W^H	
	$W^H = U_\epsilon^H M_1$	$W = V_\epsilon$
1	267	260
2	271	267
3	263	272
4	260	256
5	255	262
6	209	199
7	209	202
8	209	208
9	137	135
10	127	126
11	126	125
12	115	115
13	119	118
14	119	120
15	114	110
16	104	103
17	105	105
18	103	102
19	97	94
20	96	90

Table 4.1: Number of iterations required by GMRES(10) preconditioned by a Frobenius-norm minimization method updated with spectral corrections to reduce the normwise backward error by 10^{-8} for increasing number of corrections on the satellite. Different choices are considered for the operator W^H.

succeeds in clustering most of the eigenvalues close to one leaving only few isolated eigenvalues close to the origin, ARPACK in forward mode is an efficient approach to compute the associated eigenvectors. In that context, another advantage is that ARPACK computes with a better accuracy the smallest eigenpairs. These are the most important to compute accurately as they play a dominant role in the convergence of the Krylov solvers. This extra calculation can be amortized if several linear systems with the same coefficient matrix but different right-hand sides have to be solved. We indicate that, on real life problem arising in electromagnetism applications, this extra cost can be fairly quickly overcome. As an empiric criterion for the selection of the dimension of the low rank correction we can indicate that removing all the eigenvalues that are small and isolated (sometimes in a small cluster) is an effective approach. For instance in the electromagnetic application,

Dimension of low-rank update	# ARPACK Mat-vec	# Amortization rhs
1	120	-
2	336	-
3	290	290
4	250	84
5	192	48
6	183	9
7	175	8
8	165	8
9	154	3
10	169	3
11	157	3
12	219	4
13	224	4
14	212	4
15	223	4
16	202	3
17	226	4
18	264	4
19	264	4
20	300	4

Table 4.2: Number of matrix-vector products required by ARPACK to compute approximate eigenvalues nearest 0 and the corresponding right eigenvectors.

removing the few eigenvalues of magnitude less than 10^{-3} was enough to speedup the convergence on all our test examples (Carpentieri 2002). Finally, when the Krylov solver is GMRES, we suggest that the techniques described by Baglama et al. (1999) can be applied. This consists in recovering the eigenvectors from the Arnoldi process embedded in the GMRES iterations and then updates the preconditioner at each GMRES restart.

Acknowledgement

We would like to thank Serge Gratton for his fruitful and stimulating discussions during the development of this work.

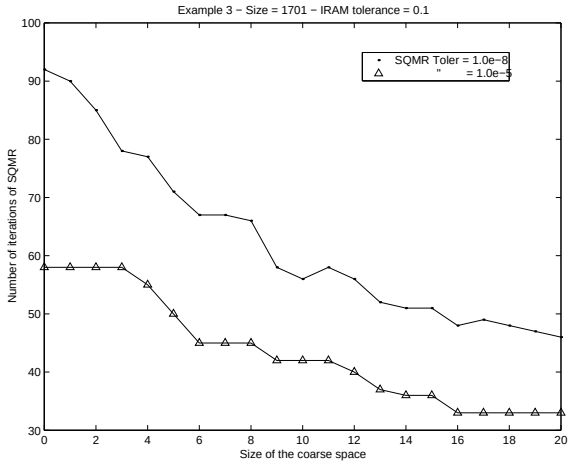

Figure 4.4: Number of iterations required by SQMR preconditioned by a Frobenius-norm minimization method updated with spectral corrections to reduce the normwise backward error by 10^{-5} for increasing number of corrections on the satellite. The symmetric formulation of Proposition 2 with the choice $W = V_\varepsilon$ is used for the low-rank updates.

References

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

J. Baglama, D. Calvetti, G. H. Golub, and L. Reichel. Adaptively preconditioned GMRES algorithms. *SIAM J. Scientific Computing*, **20(1)**, 243–269, 1999.

B. Carpentieri. *Sparse preconditioners for dense linear systems from electromagnetic applications*. PhD thesis, CERFACS, Toulouse, France, 2002.

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*, **7**(7-8), 667–685, 2000.

B. Carpentieri, I. S. Duff, L. Giraud, and G. Sylvand. Combining fast multipole techniques and an approximate inverse preconditioner for large parallel electromagnetics calculations. Technical Report in preparation, CERFACS, Toulouse, France, 2002.

L. M. Carvalho, L. Giraud, and P. Le Tallec. Algebraic two-level preconditioners for the Schur complement method. *SIAM J. Scientific Computing*, **22**(6), 1987–2005, 2001.

E. de Sturler. Inner-outer methods with deflation for linear systems with multiple right-hand sides. *in* 'Householder Symposium XIII, Proceedings of the Householder Symposium on Numerical Algebra, June 17 - 26, Pontresina, Switzerland', pp. 193–196, 1996.

I. S. Duff, R. G. Grimes, and J. G. Lewis. Users' guide for the Harwell-Boeing sparse matrix collection. Technical Report RAL 92-086, Rutherford Appleton Laboratory, Chilton, England, 1992. Also CERFACS Technical Report TR/PA/92/86.

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

R. W. Freund and N. M. Nachtigal. An implementation of the QMR method based on coupled two-term recurrences. *SIAM J. Scientific Computing*, **15(2)**, 313–337, 1994.

G. H. Golub and C. F. Van Loan. *Matrix computations.* Johns Hopkins Studies in the Mathematical Sciences. The Johns Hopkins University Press, Baltimore, MD, USA, third edn, 1996.

S. A. Kharchenko and A. Yu. Yeremin. Eigenvalue translation based preconditioners for the GMRES(k) method. *Numerical Linear Algebra with Applications*, **2(1)**, 51–77, 1995.

L. Yu. Kolotilina. Twofold deflation preconditioning of linear algebraic systems. I. Theory. Technical Report EM-RR 20/95, Elegant Mathematics, Inc., 1995. Available at: http://www.elegant-math.com/abs-emrr.htm.

R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK User's guide: Solution of large-scale problem with implicitly restarted Arnoldi methods.* SIAM, Philadelphia, 1998.

L. Mansfield. On the use of deflation to improve the convergence of conjugate gradient iteration. *Communs. Appl. Numer. Meth.*, **4**, 151–156, 1988.

L. Mansfield. Damped Jacobi preconditioning and coarse grid deflation for conjugate gradient iteration on parallel computers. *SIAM J. Scientific and Statistical Computing*, **12**, 1314–1323, 1991.

J. A. Meijerink and H. A. van der Vorst. An iterative solution method for linear systems of which the coeffcient matrix is a symmetric M-matrix. *Mathematics of Computation*, **31**, 148–162, 1977.

R. B. Morgan. A restarted GMRES method augmented with eigenvectors. *SIAM J. Matrix Analysis and Applications*, **16**, 1154–1171, 1995.

R. B. Morgan. Implicitely restarted GMRES and Arnoldi methods for nonsymmetric systems of equations. *SIAM J. Matrix Analysis and Applications*, **21(4)**, 1112–1135, 2000.

R. Nicolaides. Deflation of conjugate gradients with applications to boundary value problems. *SIAM J. Numerical Analysis*, **24**, 355–365, 1987.

Y. Saad. Projection and deflation methods for partial pole assignment in linear state feedback. *IEEE Trans. Automat. Contr.*, **33(3)**, 290–297, 1988.

Y. Saad. Analysis of augmented Krylov subspace techniques. *SIAM J. Scientific Computing*, **14**, 461–469, 1993.

Y. Saad. ILUT: A dual threshold incomplete LU factorization. *Numerical Linear Algebra with Applications*, **1**, 387–402, 1994.

Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Scientific and Statistical Computing*, **7**, 856–869, 1986.

D. C. Sorensen. Implicit application of polynomial filters in a k-step Arnoldi method. *SIAM J. Matrix Analysis and Applications*, **13**, 357–385, 1992.

G. Sylvand. *Résolution Itérative de Formulation Intégrale pour Helmholtz 3D : Applications de la Méthode Multipôle à des Problèmes de Grande Taille.* PhD thesis, Ecole Nationale des Ponts et Chaussées, 2002.

H. A. van der Vorst. Bi-CGSTAB: a fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Scientific and Statistical Computing*, **13**, 631–644, 1992.

H. A. van der Vorst and C. Vuik. The superlinear convergence behaviour of GMRES. *J. Comput. Appl. Math.*, **48**, 327–341, 1993.

J. H. Wilkinson. *The Algebraic Eigenvalue Problem.* Oxford University Press, Walton Street, Oxford OX2 6DP, UK, 1965.