

ACCELERATING KRYLOV SOLVERS WITH LOW RANK CORRECTION

L. Giraud* and S. Gratton†

*ENSEEIH-IRIT
2, rue Charles Camichel, BP 7122
F 31071 Toulouse Cedex 7
e-mail: giraud@enseeiht.fr

†CERFACS
42 Avenue G. Coriolis
F 31057 Toulouse Cedex 01
e-mail: gratton@cerfacs.fr

Key words: Preconditioning techniques, linear systems, sensitivity analysis, spectral correction, large parallel calculations

Abstract. *In this paper we present some aspects of recent works we have been developing on preconditioning techniques for accelerating Krylov solvers that are based on low rank corrections of a prescribed preconditioner. For SPD linear systems, we investigate the behaviour of some techniques based on spectral approaches when the eigenvalues are only known approximately. We use the first-order perturbation theory for eigenvalues and eigenvectors to investigate the behaviour of the spectrum of the preconditioned systems using first order approximation. For unsymmetric linear system, we present a similar technique suited for the solution of sequences of linear systems is described. This technique is a combination of a low rank update spectral preconditioner and a Krylov solver that computes on the fly approximations of the eigenvectors associated with the smallest eigenvalues. We illustrate the interest of this approach in large parallel calculations for electromagnetic simulations. In this latter context, the solution technique enables the reduction of the simulation times by a factor of up to eight; these simulation times previously exceeded several hours of computation on a modern high performance computer.*

1 INTRODUCTION

It is well known that the convergence of Krylov methods for solving the linear system often depends to a large extent on the eigenvalue distribution. In many cases, it is observed that “removing” the smallest eigenvalues can greatly improve the convergence. Several techniques have been proposed in the past few years that attempt to tackle this problem. In the first part of this paper we investigate the behaviour of some of these techniques when the eigenvalues are only known approximately. We use the first-order perturbation

theory for eigenvalues and eigenvectors to investigate the behaviour of the spectrum of the preconditioned systems using first order approximation. We illustrate the effect of the inexactness of the eigenlements on the behaviour of the resulting preconditioner when applied to accelerate the conjugate gradient method.

In the second part of this paper, we investigate a solution scheme suited for the solution of large linear systems involving the same unsymmetric matrix but different right-hand sides. Such a situation occurs in many numerical simulations in scientific and engineering applications. We illustrate the effectiveness of the approach on large real life problems arising from electromagnetics applications.

2 On the sensitivity of some spectral preconditioners

In many problems the convergence of Krylov solvers can be significantly slowed down by the presence of small eigenvalues in the spectrum of the matrices involved in the solution of the linear systems. This occurs for instance when the Conjugate Gradient (CG) method is implemented to solve linear systems arising from the discretization of second-order elliptic problems. For symmetric positive definite (SPD) linear systems it is well-known that the convergence of CG to solve $Ax = b$ depends to a large extent on the eigenvalue distribution of the coefficient matrix A . This can be illustrated by the bound on the rate of convergence of the CG method given by

$$\|x_k - x^*\|_A \leq 2\|x_0 - x^*\|_A \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^k, \quad (1)$$

where $\kappa = \lambda_{\max}(A)/\lambda_{\min}(A)$ is the spectral condition number of A , the A -norm of x is $\|x\|_A = \sqrt{x^T A x}$ and the exact solution is $x^* = A^{-1}b$. This analysis leads to the idea of improving the convergence of CG by using a preconditioner M such that the ratio $\lambda_{\max}(MA)/\lambda_{\min}(MA)$ is less than κ . In this paper, we are interested in spectral preconditioners (and some of their variants) that exploit some information on the eigenpairs of A . The underlying driving idea of these approaches is to capture in a low dimensional space the modes that do not quickly converge with a first level preconditioner. In order to be efficient and keep the dimension of the low dimensional space reasonably small, these techniques are generally used in combination with a first level preconditioner that does a good job of clustering most eigenvalues near to one with relatively few outliers near the origin [6, 9, 10, 24, 35]. These spectral preconditioners can be split into two main families depending on their effect on the spectrum. They are referred to as deflation preconditioners [9, 13, 26] if they attempt to move a subset of eigenvalues to a positive quantity σ ; they are referred to as coarse grid preconditioners [6, 14] if they only attempt to shift the subset close to σ . The name of those latter techniques comes from domain decomposition and was first introduced in [5]. For this reason $\sigma = 1$ is often considered in practice. An impressive example of the efficiency of a spectral preconditioner is provided by the atmosphere data assimilation area [13]. In this application, nonlinear least-squares problems with more than 10^7 unknowns are daily solved using a Gauss-Newton approach.

2.1 Spectral preconditioner variants

We first consider one representative of the deflation preconditioners and one of the coarse grid preconditioners. Let $V = [v_1, \dots, v_n] \in \mathbb{R}^{n \times n}$ be an eigen-basis of A and $\{\lambda_i\}_{i=1, \dots, n}$ be the set of corresponding eigenvalues sorted by increasing magnitude. In order to move $\{\lambda_i\}_{i=1, \dots, k}$ to σ , we define the following deflation preconditioner:

$$M^{def} = I + \sum_{i=1}^k \left(\frac{\sigma}{\lambda_i} - 1 \right) \frac{v_i v_i^T}{\|v_i\|^2}. \quad (2)$$

The columns of V also form an eigen-basis of $M^{def} A$. For $V_k = [v_1, \dots, v_k]$, this preconditioner is such that $M^{def} A V_k = \sigma V_k$ and $M^{def} A w = A w$ if $V_k^T w = 0$, which shows that M^{def} moves the eigenvalues $\{\lambda_i\}_{i=1, \dots, k}$ to σ and leaves the rest of the spectrum unchanged. This technique is expected to be especially efficient in the case where $\{\lambda_i\}_{i=k+1, \dots, n}$ are already in the neighbourhood of σ , in which case $\lambda_{\max}(MA)/\lambda_{\min}(MA)$ becomes close to one. Because (1) only provides an upper-bound for the convergence rate it might be argued that, if $\{\lambda_i\}_{i=k+1, \dots, n}$ are already close to σ , shifting $\{\lambda_i\}_{i=1, \dots, k}$ to any quantity close to σ (and not necessarily to σ exactly) does still improve the convergence of CG. To this end, if $\{\lambda_i\}_{i=1, \dots, k}$ are small we can use the coarse grid preconditioner

$$M^{coarse} = I + \sigma V_k (\text{diag}(\lambda_i))^{-1} V_k^T, \quad (3)$$

where $\text{diag}(\lambda_i)$ denotes the diagonal matrix with entries λ_i . The columns of V also form an eigen-basis of $M^{coarse} A$. This preconditioner is such that $M^{coarse} A v_i = (\sigma + \lambda_i) v_i$ and $M^{coarse} A w = A w$ if $V_k^T w = 0$. That is, the eigenvalues $\{\lambda_i\}_{i=1, \dots, k}$ are shifted to $\sigma + \lambda_i$, while the rest of the spectrum is unchanged. This latter technique is particularly suited when $\{\lambda_i\}_{i=1, \dots, k}$ are small.

In the previous paragraph the spectral transformations associated with M^{def} and M^{coarse} rely on the facts that (λ_i, v_i) are exact eigenpairs of a SPD matrix. This implies two intensively used properties that are $A v_i = \lambda_i v_i$ and $v_i^T v_j = 0$ if $i \neq j$. We assume now that we only have access to approximate spectral information, and denote by $(\tilde{\lambda}_i, \tilde{v}_i)$ the inexact eigenpairs such that the two latter properties do not necessarily hold. We only suppose that $\tilde{V}_k = [\tilde{v}_1, \dots, \tilde{v}_k]$ is such that $\tilde{V}_k^T \tilde{V}_k$ and $\tilde{V}_k^T A \tilde{V}_k$ are nonsingular. We can write the two above preconditioners in a form that does a weaker usage of the properties of the exact eigenpairs. The inexact ‘‘deflation’’ preconditioner then reads

$$M_1 = I + \tilde{V}_k \left(\sigma \left(\tilde{V}_k^T A \tilde{V}_k \right)^{-1} - \left(\tilde{V}_k^T \tilde{V}_k \right)^{-1} \right) \tilde{V}_k^T. \quad (4)$$

Noticing that $V_k^T A V_k$ is diagonal, we can be tempted to approximate it by the diagonal of the Rayleigh quotients. This gives

$$M_1^{rayl} = I + \tilde{V}_k \left(\sigma \left(\text{diag}(\tilde{v}_i^T A \tilde{v}_i) \right)^{-1} - \left(\tilde{V}_k^T \tilde{V}_k \right)^{-1} \right) \tilde{V}_k^T. \quad (5)$$

Furthermore, because the diagonal elements of M_1 are simply the eigenvalues, we can also use the following preconditioner formulation

$$M_1^{eig} = I + \tilde{V}_k \left(\sigma \left(\text{diag}(\tilde{\lambda}_i) \tilde{v}_i^T \tilde{v}_i \right)^{-1} - \left(\tilde{V}_k^T \tilde{V}_k \right)^{-1} \right) \tilde{V}_k^T. \quad (6)$$

Using the orthogonality property of the eigenvectors M_1 reduces to

$$M_{1,\perp}^{rayl} = I + \sum_{i=1}^k \left(\sigma \left(\frac{\tilde{v}_i^T A \tilde{v}_i}{\tilde{v}_i^T \tilde{v}_i} \right)^{-1} - 1 \right) \frac{\tilde{v}_i \tilde{v}_i^T}{\tilde{v}_i^T \tilde{v}_i}, \quad (7)$$

that could also lead to consider

$$M_{1,\perp}^{eig} = I + \sum_{i=1}^k \left(\frac{\sigma}{\tilde{\lambda}_i} - 1 \right) \frac{\tilde{v}_i \tilde{v}_i^T}{\tilde{v}_i^T \tilde{v}_i}. \quad (8)$$

Similarly, for the coarse grid preconditioners we consider the following variants

$$M_{+1} = I + \sigma \tilde{V}_k \left(\tilde{V}_k^T A \tilde{V}_k \right)^{-1} \tilde{V}_k^T, \quad (9)$$

$$M_{+1}^{rayl} = I + \sigma \tilde{V}_k \left(\text{diag}(\tilde{v}_i^T A \tilde{v}_i) \right)^{-1} \tilde{V}_k^T, \quad (10)$$

$$M_{+1}^{eig} = I + \sigma \tilde{V}_k \left(\text{diag}(\tilde{\lambda}_i \tilde{v}_i^T \tilde{v}_i) \right)^{-1} \tilde{V}_k^T. \quad (11)$$

To study the performance of the above preconditioners in the presence of inexact spectral information, we assume that the spectral information is not related to A but to a nearby matrix $A + tE$, where t is a real parameter and $\|E\| = 1$. Let denote $\lambda_i(t)$ and $v_i(t)$ the eigenvalues and eigenvectors of $A + tE$. If A has only simple eigenvalues, it can be shown [32] that the eigenvalues of $A + tE$ are differentiable functions of t in a neighbourhood \mathcal{V} of $t = 0$. If the eigenvectors are normalized using $v_i(t)^T v_i = 1$ the eigenvectors are also differentiable functions of t in a neighbourhood of $t = 0$. Note that none of the preconditioners assume that the eigenvectors have unit length. Indeed, the preconditioners are invariant by any nonzero scaling of the eigenvectors. Therefore the normalization $v_i(t)^T v_i = 1$ can be assumed for the analysis without loss of generality.

2.2 Sensitivity analysis

2.2.1 Notation

For any square matrix $X \in \mathbb{R}^{n \times n}$, let \underline{X}_i denote the $n \times (n - 1)$ matrix whose columns are those of X excepted for the i^{th} , that is $\underline{X}_i = [X(:, j)]_{j=1, \dots, n; j \neq i}$. For a $n \times n$ matrix X , $\{\lambda_1(X), \dots, \lambda_n(X)\}$ are the eigenvalues of X where multiple eigenvalues are repeated. We also assume that $|\lambda_1(X)| \leq \dots \leq |\lambda_n(X)|$. The i th eigenvalue of

A is denoted by λ_i when there is no possible confusion. Let $A \in \mathbb{R}^{n \times n}$ be a SPD matrix where

$$AV = VD \text{ with } V^T V = I \text{ and } D = \text{diag}(\lambda_i)_{i=1, \dots, n}$$

denotes its spectral decomposition. We assume that all the eigenvalues of A are simple.

For a vector x , $\|x\| = (\sum_{i=1}^n x_i^2)^{1/2}$ is the Euclidean vector norm, and $\|A\| = \max_{\|x\|=1} \|Ax\|$ is the spectral norm of the matrix A . The operator \circ denotes the Hadamard product: $A \circ B = [a_{ij}b_{ij}] \in \mathbb{C}^{m \times n}$, for A and $B \in \mathbb{C}^{m \times n}$. The spectral norm is submultiplicative with respect to the Hadamard product (see [2, p. 332]): $\|A \circ B\| \leq \|A\| \|B\|$.

Let \tilde{V} and \tilde{D} be defined by $\tilde{V} = [v_1(t), \dots, v_n(t)]$ and $\tilde{D} = \text{diag}(\tilde{\lambda}_i) = \text{diag}(\lambda_i(t))$. For sufficiently small $t \in \mathcal{V}$, we have

$$(A + tE)\tilde{V} = \tilde{V}\tilde{D}.$$

Note that $\tilde{V}(0) = V$. Using the Landau little “ o ” notation, the first order expansion of the eigenvalues and eigenvectors in the direction E [32], is $\tilde{V} = V + \delta V(t) + o(t)$ where the i^{th} column of $\delta V(t)$ is defined by

$$\delta v_i(t) = t \underline{V}_i (\lambda_i I - B_i)^{-1} \underline{V}_i^T E v_i,$$

and $\lambda_i(t) = \lambda_i + \delta \lambda_i(t) + o(t)$, where $\delta \lambda_i(t) = t v_i^T E v_i$ and the $(n-1) \times (n-1)$ diagonal matrix $B_i = \text{diag}(\lambda_j)_{j=1, \dots, n; j \neq i}$. The first order expansion of the eigenvalues of the preconditioned matrices will be expressed in terms of the following $k \times k$ matrices W and Y defined in function of the eigenvalues of A by their (ℓ, s) -entry :

$$W_{\ell, \ell} = 0; W_{\ell, s} = \frac{\sigma}{\lambda_s - \lambda_\ell} \sqrt{\frac{\lambda_\ell}{\lambda_s}} \text{ for } \ell \neq s, \quad (12)$$

$$Y_{\ell, \ell} = 0; Y_{\ell, s} = \frac{\sigma - \lambda_s}{\lambda_s - \lambda_\ell} \sqrt{\frac{\lambda_\ell}{\lambda_s}} \text{ for } \ell \neq s. \quad (13)$$

Similarly we also introduce the $k \times k$ matrix $\Delta = V_k^T E V_k$ and the diagonal matrix $J = \text{diag} \left(-\frac{\sigma}{\lambda_\ell} \right)_{\ell=1, \dots, k}$.

2.3 Sensitivity and backward errors

For all the preconditioners considered in this paper, the eigenvalues of the preconditioned matrices write $\mu_i(t) = \mu_i(0) + \lambda_i ((X_1 + X_2) \circ \Delta + X_1^T \circ \Delta^T) t + o(t)$, where the X_i are matrices depending on selected preconditioner and the targeted eigenvalues [16]. We summarize the various values of the matrices X_1 and X_2 for the different preconditioners in Table 1. We can therefore define a condition number κ_i for the eigenvalue μ_i in the direction of E [28] by

$$\kappa_i = \lim_{u \rightarrow 0} \sup_{0 < |t| < u} \frac{|\mu_i(t) - \mu_i(0)|}{|t|} = |\lambda_i ((X_1 + X_2) \circ \Delta + X_1^T \circ \Delta^T)|. \quad (14)$$

Notice that κ_i is not the usual condition number of the eigenvalue of a matrix [19, p. 323], but the condition number of the map $\mu_i(t)$. Taking norms and using the submultiplicativity of the spectral norm with respect to the Hadamard product yields [16]

$$\kappa_i \leq (2\|X_1\| + \|X_2\|), \tag{15}$$

where we have used that $\|E\| = 1$. Equations (14) and (15) show that if the entries of X_1 and X_2 are small, the condition number of the eigenvalues μ_i is small. By inspecting the equalities (12) and (13) follows that asymptotically for $t \rightarrow 0$,

- the preconditioners M_1 , M_{+1} and M_{+1}^{rayl} are stable (i.e. X_1, X_2 are the zero matrix),
- the preconditioners M_1^{rayl} , $M_{1,\perp}^{rayl}$, M_1^{eig} and $M_{1,\perp}^{eig}$ may present an instability if for some (s, ℓ) , the ratio $\frac{\sigma}{\lambda_s - \lambda_\ell} \sqrt{\frac{\lambda_\ell}{\lambda_s}}$ is large. In the above statement we have assumed that λ_s is far from σ , which seems to be a reasonable assumption as otherwise we would not have targeted this eigenvalue. This instability happens for instance if some eigenvalues are clustered or small and isolated.

In Table 1 we summarize the situation where a high sensitivity of the eigenvalues is expected. In the case where E is a real symmetric matrix, M_1^{eig} and $M_{1,\perp}^{eig}$ are the same

Prec	X_1	X_2	some cases of ill-conditioning
M_1	0	0	none
M_1^{rayl}	W	0	cluster, small
M_1^{eig}	W	J	cluster, small
$M_{1,\perp}^{rayl}$	Y	0	cluster, small
$M_{1,\perp}^{eig}$	Y	J	cluster, small
M_{+1}	0	0	none
M_{+1}^{rayl}	0	0	none
M_{+1}^{eig}	0	J	small

Table 1: Matrices X_1 and X_2 for the spectral preconditioners and some cases of ill-conditioning. The terms “cluster” and “small” refer respectively to the presence of cluster or of small isolated eigenvalues.

preconditioners, elementary calculations show that $X_1 \circ \Delta + X_1^T \circ \Delta^T$ are the same in both cases.

2.4 Use of spectral preconditioners in CG

In this section we illustrate the effect of using the approximate eigenpairs on the convergence behaviour of PCG. In that respect we consider the 685×685 BUS685 matrix,

denoted B_{685} , from the Harwell-Boeing collection. We compute an Incomplete Cholesky factorization (IC) CC^T with threshold $4 \cdot 10^{-1}$, which is our first level preconditioner and consequently $\sigma = 1$. We apply the various spectral preconditioners to the matrix $A = C^{-1}B_{685}C^{-T}$. As in the previous series of experiments, we use the eigenvectors of the perturbed matrix $A + tE$ to build the preconditioners. Consequently we use eigenelements that have a backward error of the order of t .

The right-hand side is chosen so that the solution of $Ax = b$ is the vector of all ones: $x = (1, \dots, 1)^T$, $b = Ax$. For the numerical experiments the initial guess is the zero vector and we decide to stop the PCG iterations when the normalized unpreconditioned residual is reduced by 10^{-9} , so that the stopping criterion is independent of the preconditioner. Even though this quantity might be a by-product of the PCG solver we explicitly compute the unpreconditioned residual to decide when to stop the iterations.

In Table 2 we report on the number of PCG iterations for the various deflation and coarse-grid preconditioners for $t = 10^{-12}$ and $t = 10^{-3}$. We vary from one to ten the number ne of eigenelements used to build the preconditioners. The eigenvalues of A obtained with `eig` are reported in Table 3. For the smallest perturbation, that is $t = 10^{-12}$, all the preconditioners behave exactly the same. Because the IC preconditioner has already clustered many of the eigenvalues close to one, moving the smallest eigenvalues (that vary from $5.67 \cdot 10^{-4}$ to $3.05 \cdot 10^{-2}$ see Table 3) exactly to one or shifting them by one leads to the same behaviour of PCG. However, when a perturbation is applied, that is when the eigenelements are less accurately computed, some differences appear. Both M_1 and M_{+1} perform similarly and outperform the others. Then the various variants that approximate the eigenvalues using Rayleigh quotients perform similarly. The worse behaviour is observed for the variants that make use of the approximate eigenvalues.

$t = 10^{-12}$										
ne	1	2	3	4	5	6	7	8	9	10
All Prec	161	147	132	118	106	95	88	81	81	79
$t = 10^{-3}$										
ne	1	2	3	4	5	6	7	8	9	10
M_1	168	154	142	129	122	114	107	104	101	98
M_1^{rayl}	168	155	152	142	144	139	133	131	126	121
M_1^{eig}	168	156	154	141	143	140	134	194	183	183
$M_{1,\perp}^{rayl}$	168	155	151	139	138	133	125	124	121	117
$M_{1,\perp}^{eig}$	168	156	152	140	139	133	125	184	173	171
M_{+1}	168	154	142	129	121	114	107	104	101	98
M_{+1}^{rayl}	168	155	152	139	138	133	125	124	121	117
M_{+1}^{eig}	168	156	152	140	139	133	125	184	174	171

Table 2: # iterations for PCG for perturbed eigenpairs on the BUS 685 matrix.

λ_1	λ_2	λ_3	λ_4	λ_5
5.67e-04	2.91e-03	3.58e-03	4.71e-03	6.19e-03
λ_6	λ_7	λ_8	λ_9	λ_{10}
8.35e-03	1.85e-02	2.05e-02	3.04e-02	3.05e-02

Table 3: The ten smallest eigenvalues targeted by the spectral preconditioners on the BUS 685 matrix.

To summarize we use the first-order perturbation theory for eigenvalues and eigenvectors to investigate the behaviour of deflation and coarse preconditioners for SPD linear systems. Our analysis shows a better stability of the preconditioners M_1 and M_{+1} compared to the other preconditioning variants that exploit some additional properties that are only true for exact eigenpairs. These results show that targeting small eigenvalues or small clusters may require a backward stable calculation of the eigenelements. An important result of this work is that the efficiency of a spectral preconditioner should not be assessed only using exact eigenpairs.

For a more detailed description of this work we refer to [16].

3 Incremental spectral preconditioners for sequences of linear systems

Many numerical simulations in scientific and engineering applications require the solution of set of large linear systems involving the same coefficient matrix but different right-hand sides. That is, the solution of

$$Ax^{(\ell)} = b^{(\ell)} \quad \text{for } \ell = 1, \dots, \quad (16)$$

where A is a nonsingular matrix in $\mathbb{C}^{n \times n}$, and $x^{(\ell)}$ and $b^{(\ell)}$ are vectors of \mathbb{C}^n . Such a situation occurs for instance in some parametric studies, in the calculation of eigenvalues using shift and invert techniques, in radar cross section for electromagnetic calculations. Several numerical techniques can be considered to attempt to reduce the cost of solving subsequent systems in the sequence. The approach to follow depends on the features of the problem at hand and might consist in adapting the Krylov solver or in improving the preconditioner. For simultaneous right-hand sides, block Krylov linear solvers [15, 25, 34] might be appropriate. For sequence of right-hand sides that do not vary much, a straightforward idea is to use the former solution as an initial guess for the next solve. A more sophisticated variant is the seed approach [31], that consists in choosing the initial guess vector of the current system such that it complies with an optimum norm or an orthogonality criterion over the Krylov spaces associated with the previous right-hand sides. Another alternative is GCRO-DR recently proposed in [27] that further exploits the deflating ideas presented in GMRES-E [22] or GMRES-DR [23].

Other possible complementary alternatives consist in improving a selected preconditioner. In most of the situations, the linear systems are solved using an application dependent preconditioner whose efficiency and cost are controlled by a few parameters.

Because the preconditioner is used for all the right-hand sides, some extra effort can be devoted to improve it. The extra work involved in its construction can sometimes be amortized as many systems have to be solved. For instance, if an incomplete factorization [4, 20, 30] is considered, the threshold parameter can be decreased to allow for more fill-in in the factors, giving rise to a more efficient preconditioner but more costly to build and to apply. Similarly, in the algebraic multigrid context we might decide to select techniques that leads to better preconditioners but that have a more expensive setup phase. Even though such an approach is certainly beneficial, the gain is often limited and other complementary techniques can be envisaged. One possibility is to perform a spectral update of the ad-hoc preconditioner as described for instance in [3, 7, 12].

In this section, we present a preconditioning technique that implements a spectral low rank increment of the preconditioner after the solution of each right-hand side. From one linear system to the next, an elementary low rank correction is added to the preconditioner and we therefore use the terminology “incremental preconditioner”. The spectral information required to perform the elementary update is recovered at very low computational cost from by-products of the selected Krylov linear solver. Among the preconditioners that can be considered to perform the elementary update [3, 7, 12], we focus in this paper on the variant described in [7]. Our solution technique is expected to be efficient if it is implemented to complement a prescribed ad-hoc and efficient preconditioner that clusters most of the eigenvalues near one and only leaves a few outliers near the origin. Such a situation appears in various applications and closely related spectral techniques have been recently applied with success in elastic wave propagation or thermal convection [35], porous media flow [24], semiconductor device modeling [29] or electromagnetic simulation [11].

3.1 Incremental spectral preconditioners formulation

Many preconditioners are able to cluster most of the eigenvalues close to one but still leave a few close to the origin. To move these eigenvalues close to one might be possible by tuning the parameters that control these preconditioners. However, this tuning is often difficult and might lead to very expensive preconditioners to setup and to apply. Furthermore, it is well known that the convergence of Krylov methods for solving the linear system often depends to a large extent on the eigenvalue distribution. In many cases, it is observed that “removing” the smallest eigenvalues can greatly improve the convergence. This can be illustrated in the symmetric positive definite case by the bound on the rate of convergence of the Conjugate Gradient method given by (1). Some arguments exist for unsymmetric systems to mitigate the bad effect of the smallest eigenvalues on the rate of convergence of the unsymmetric Krylov solver [3, 12, 21].

In the remainder of this section, we provide a heuristic motivation for our incremental preconditioning methods. In [7] a spectral low rank update (SLRU) preconditioning technique that exploits the knowledge of the eigenvectors associated with the smallest eigenvalues is described. The proposed preconditioners shift the corresponding eigenvalues

close to one and numerical examples show the relevance of this approach. Let us now briefly recall the formulations of these preconditioners as well as their spectral properties.

In the remainder of this section, we provide a heuristic motivation for our incremental preconditioning methods. In [7] a spectral low rank update (SLRU) preconditioning technique that exploits the knowledge of the eigenvectors associated with the smallest eigenvalues is described. The proposed preconditioners shift the corresponding eigenvalues close to one and numerical examples show the relevance of this approach. Let us now briefly recall the formulations of these preconditioners as well as their spectral properties. In particular, we derive a new proof on its spectral transformation.

Let us consider the solution of the first linear system

$$Ax^{(1)} = b^{(1)}, \quad (17)$$

where A is a $n \times n$ unsymmetric complex nonsingular matrix, and $x^{(1)}$ and $b^{(1)}$ are vectors of size n . The linear system is solved using a preconditioned Krylov solver and we denote by M the initial preconditioner. Let $\{\lambda_1, \dots, \lambda_n\}$ be the set of eigenvalues of MA where the multiple eigenvalues are repeated. Let the columns of U be the basis of a right invariant subspace of MA of dimension k . Suppose without loss of generality that $MAU = UJ_k$ where the eigenvalues of J_k are $\{\lambda_1, \dots, \lambda_k\}$. Using U we can design a preconditioner for the second linear system. The spectrum of the preconditioned matrix is given by the proposition below.

Proposition 1 *Let W be such that $A_c = WAU$ is nonsingular, and let define*

$$M_{ISLRU}^{(2)} = M + UA_c^{-1}W. \quad (18)$$

Then the eigenvalues of $M_{ISLRU}^{(2)}A$ (or $AM_{ISLRU}^{(2)}$) are

$$\begin{cases} \eta_i^{(2)} = \lambda_i & \text{if } i > k, \\ \eta_i^{(2)} = 1 + \lambda_i & \text{if } i \leq k. \end{cases}$$

For a sequence of right-hand sides, we consider a repeated use of the result above. Assume that after the solution of the first system we have $x^{(1)}$, $\{\lambda_1, \dots, \lambda_{k_1}\}$ and $U^{(1)}$ an associated basis for the right invariant subspace of MA . Consequently, $\{\lambda_1, \dots, \lambda_{k_1}\}$ will be shifted in $M_{ISLRU}^{(2)}A$. The next linear system to be solved is $M_{ISLRU}^{(2)}Ax^{(2)} = M_{ISLRU}^{(2)}b^{(2)}$. Similarly after the solution of this system, we suppose that we know $x^{(2)}$, $\{\lambda_{k_1+1}, \dots, \lambda_{k_2}\}$ and $U^{(2)}$ such that $\{\lambda_{k_1+1}, \dots, \lambda_{k_2}\}$ will be shifted in $M_{ISLRU}^{(3)}A$. Notice that following (18), $M_{ISLRU}^{(3)} = M_{ISLRU}^{(2)} + U^{(2)}(W^{(2)}AU^{(2)})^{-1}W^{(2)}$. Repeating this procedure until the ℓ -th linear system we can then update the preconditioner $M_{ISLRU}^{(\ell+1)}$ such that the eigenvalues of $M_{ISLRU}^{(\ell+1)}A$ are:

$$\begin{cases} \eta_i^{(\ell+1)} = \lambda_i & \text{if } i > k_\ell, \\ \eta_i^{(\ell+1)} = 1 + \lambda_i & \text{if } i \leq k_\ell. \end{cases} \quad (19)$$

The preconditioner for the $(\ell + 1)$ -th linear system can be written as:

$$M_{ISLRU}^{(\ell+1)} = M + \sum_{j=1}^{\ell} U^{(j)} (W^{(j)} A U^{(j)})^{-1} W^{(j)}. \quad (20)$$

Similar formulation and results can be established if right preconditioners are considered.

The sketch of the solution scheme for a sequence of right-hand sides in MATLAB like syntax is described in Algorithm 1, it is by no means well suited for practical implemen-

Algorithm 1 Basic scheme for a sequence of right-hand sides

- 1: $M_{ISLRU}^{(1)} = M$
 - 2: **for** $\ell = 1, 2, \dots$ **do**
 - 3: $[x^{(\ell)}] = \text{Solve}(b^{(\ell)}, M_{ISLRU}^{(\ell)}, A)$
 - 4: $[V^{(\ell)}] = \text{Right_Invariant_Space}(A M_{ISLRU}^{(\ell)})$
 - 5: $M_{ISLRU}^{(\ell+1)} = M_{ISLRU}^{(\ell)} \left(I + V^{(\ell)} \left(W^{(\ell)} A M_{ISLRU}^{(\ell)} V^{(\ell)} \right)^{-1} W^{(\ell)} \right)$
 - 6: **end for**
-

tation. The computation of an invariant subspace is generally far more expensive than the solution of a linear system. Furthermore, the low rank-update as performed by step 5 would generally fill the preconditioner which is unacceptable. The purpose of the next section is to show how this algorithm can be adapted so that it is suitable to practical implementation.

Finally, we assume that the eigenvalues $\{\lambda_1, \dots, \lambda_{k_\ell}\}$ are simple and that the columns of $V^{(\ell)}$ are the corresponding right eigenvectors. We select $W^{(\ell)} = (V^{(\ell)})^H$ as in [7, 11].

The spectral preconditioning techniques are likely to be particularly efficient if they are implemented to complement a prescribed ad-hoc and efficient preconditioner that only leaves a few outliers near the origin. Because our primary interest is to solve a sequence of linear systems we would like to recover the eigen-information almost for free. This information is either a by-product of the linear solver or can be computed at a low computational cost from information available in the linear solver. In that context, natural candidates among the Krylov linear solvers are those that rely on an Arnoldi procedure and belong to the variants of GMRES. In particular, because we are looking for the smallest eigenvalues and because for large scale computation a restart mechanism has to be implemented, GMRES-DR(m, k) [23] appears as a suited candidate.

The GMRES-DR(m, k) method is a variant of restarted GMRES(m) that retains an approximatively invariant subspace of dimension k between the restarts (i.e. after having built a Krylov space of dimension m). We do not go into further details on the GMRES-DR solution technique and refer to [23] for a detailed description.

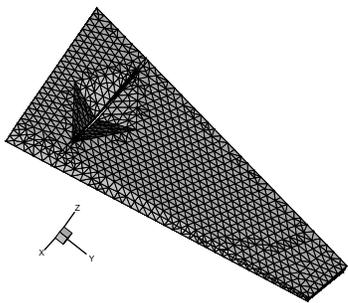
The GMRES-DR method exhibits nice capabilities to recover the spectral information we are targeting with our preconditioner. We explore in the rest of this paper the numerical behaviour of the combination of our incremental technique with this Krylov linear solver. The resulting implementation of the numerical method is obtained by replacing steps 3 and 4 of Algorithm 1 by a call to GMRES-DR where of course the preconditioner M_{ISLRU} is kept in implicit form. That is, it is never assembled and whenever a preconditioning operation is required we only have to perform matrix-vector products involving $V^{(\ell)}$ and to solve small linear systems involving $A_c^{(\ell)}$.

4 Implementation in large electromagnetism applications

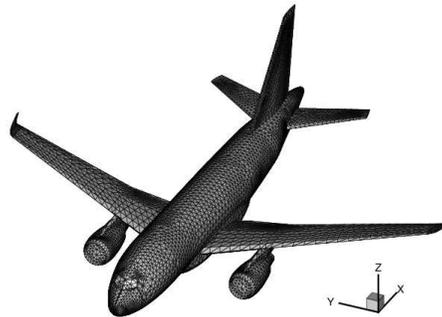
In electromagnetic calculations, a classic problem is to compute the currents generated on the surface of an object illuminated by a given incident plane wave. Such calculations, relying on the Maxwell's equations, are required in the simulation of many industrial processes coming from antenna design, 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 EFIE (Electric Field Integral Equation) as it applied to any object, without any assumption on its topological or geometrical properties. The matrices associated with the resulting linear systems are large dense, non-Hermitian and complex. Nowadays, problems with a few hundred thousand variables have to be solved and iterative solvers appear as the only viable alternative since techniques based on multipole expansion 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 $O(n \log n)$ 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 [33]. The industrial problem we focus on in this section 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 backscattered in the direction of the incident wave. This requires the solution of one linear system per incident wave. Therefore we have to solve a sequence of linear systems having the same coefficient matrix but different right-hand sides. 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 [1, 8]; it is based on a sparse approximate inverse using a Frobenius norm minimization with an a priori sparsity pattern selection strategy and is denoted by M_{Frob} . The parallel distributed code uses out-of-core data structure, in particular all the vectors involved in the linear algebra operations are out-of-core; for a detailed description of the other features we refer to [18, 33].

In the sequel, we investigate the behaviour of our incremental solution scheme on large real industrial applications. The test geometries are shown in Figure 1. They consist of a

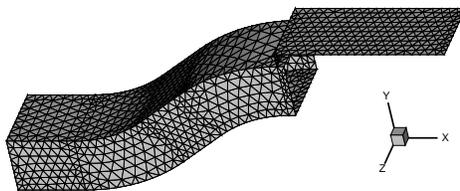
wing with a hole referred to as Cetaf, a civil aircraft and an air intake referred to as Cobra. The Cetaf is a classic test problem in the computational electromagnetics community; the other two have been kindly provided to us by EADS-CCR. In all the experiments, we con-



(a) Cetaf



(b) Civil aircraft



(c) Cobra

Figure 1: Various geometries used in the numerical experiments

sider a right preconditioner and the threshold for the stopping criterion is set to 10^{-3} on the scaled residual $\frac{\|r\|}{\|b\|}$. This tolerance is accurate for engineering purposes, as it enables the correct construction of the radar cross section of the object. The runs have been performed in single precision on a four-way SMP (Symmetric Multi-Processors) HP-Compaq Alpha cluster. Each node consists of four DEC Alpha processors (EV 6, 1.3 GFlops peak) that share 2 GB of memory. On that computer, the temporary disk space that can be used by the out-of-core solver is around 189 GB. For the experiments we consider a radar cross section that is a angular section of width 30° discretized every degree so that we end up with a sequence of 31 right-hand sides. In Figure 2, we display for the largest problems of each geometry, the number of iterations with and without the incremental preconditioner on the sequence of right-hand sides. It can be seen that the incremental

preconditioner enables us to significantly reduce the number of iterations. For the solution of the last right-hand side of the sequence, the reduction in GMRES-DR iterations is equal to about 33 for the Aircraft discretized with 94 704 degrees of freedom (dof), about 4.5 for the Cetaf with 264 159 dof and greater than 7.5 for the Cobra with 179 460 dof. In Table 4, we give more details on these numerical experiments. In that table,

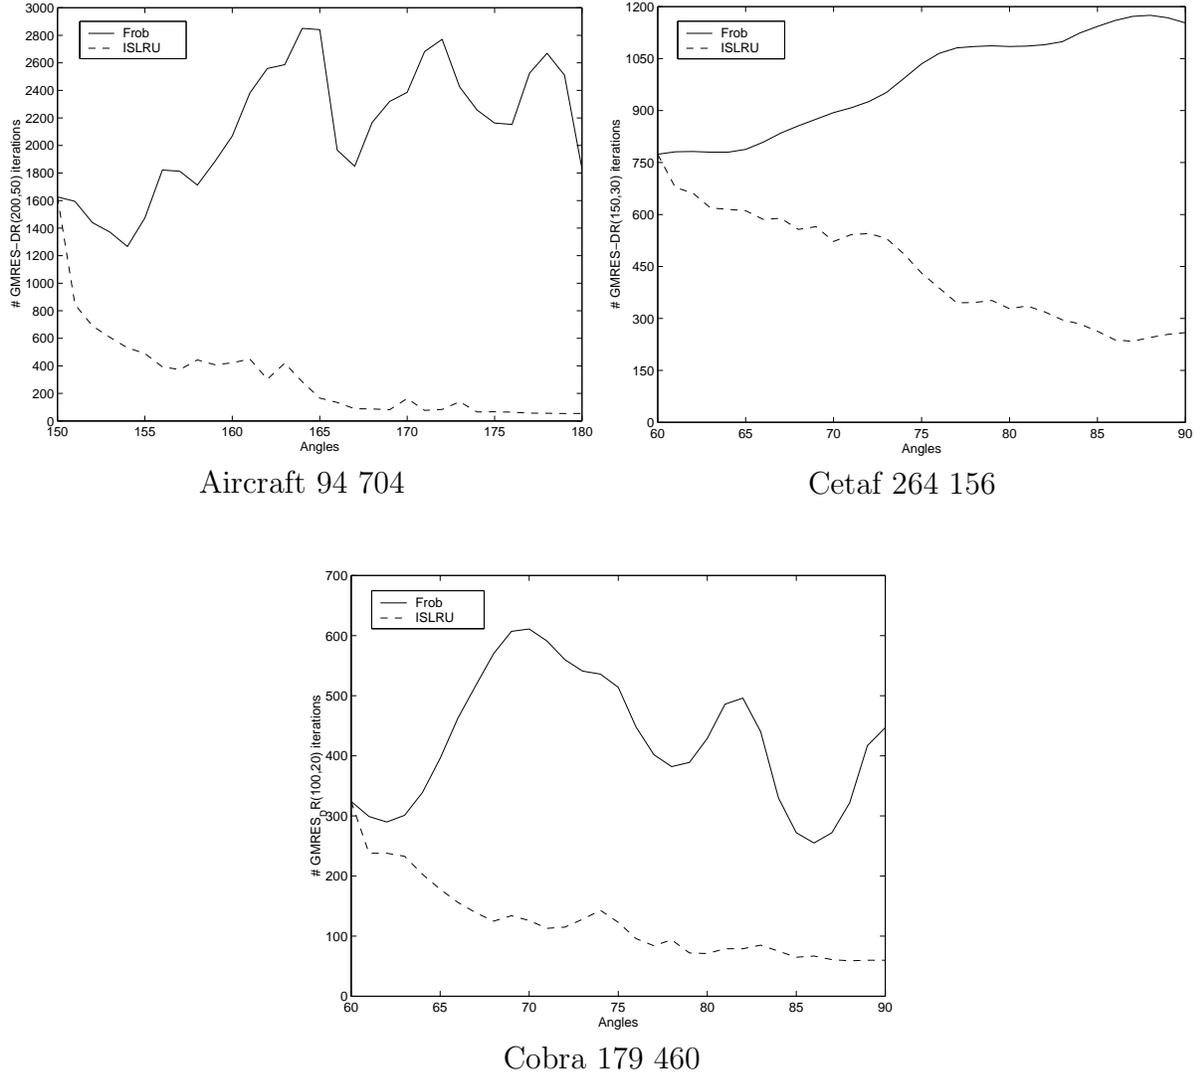


Figure 2: Number of GMRES-DR iterations with M_{Frob} and M_{ISLRU} for the different incident angles for each geometry. The sampling for the illuminating wave is one degree.

“# Proc” denotes the number of processors, “ (m, k) ” denote the restart and the number of harmonic Ritz vectors of the GMRES-DR solver, “# M.V” is the cumulated number of matrix-vector products and “Elap. time” is the parallel elapsed time to perform the

complete simulation. Finally, “ p ” denotes the total number of eigenvalues shifted by the incremental preconditioner. We mention that the selection of the number of processors used was mainly governed by the memory constraint and batch queue management implemented on our computing platform; it has not influence on the numerical behaviour of the algorithm. We can see that the incremental mechanism enables not only a significant reduction of the number of iterations but also a noticeable decrease of the elapsed time.

Geometry	Size	# Proc	(m, k)	M_{Frob}		M_{ISLRU}		
				# M.V	Elap. time	# M.V	p	Elap. time
Aircraft	23 676	8	(200,50)	20 613	12h 15m	2 688	246	1h 30m
Aircraft	94 704	31	(200,50)	66 411	2d 21h	9 801	686	10h
Cetaf	86 256	31	(150,30)	23 047	14h 30m	6 558	462	4h 30m
Cetaf	134 775	31	(150,30)	22 254	23h 20m	9 098	577	10h
Cetaf	264 159	31	(150,30)	30 804	2d 08h	13 921	770	1d 3h
Cobra	60 695	8	(100,20)	9 672	8h 30m	2 092	200	2h
Cobra	179 460	31	(100,20)	13 418	14h	3 876	365	4h

Table 4: Cost for monostatic calculations.

The gain in time ranges from two to eight depending on the problem and is almost proportional to the reduction in the total number of iterations. This can be explained by the fact that the cost of a matrix-vector product is quite high and the relative cost of our preconditioner remains low even for large p ; consequently any reduction in the number of iterations translates to a reduction in the computational time. We notice that this property is not necessarily true for sparse linear systems, where the cost of the incremental preconditioner might dominate even for small values of p so that the preconditioner might not be effective if it does not significantly reduce the number of iterations.

The proposed technique is a combination of a low rank update spectral preconditioner and a Krylov solver that recovers at run time approximations to the eigenvectors associated with the smallest eigenvalues, namely GMRES-DR in our case. It shows a clear interest in large parallel calculations for electromagnetic simulations. For more details on this work we refer to [17].

5 CONCLUSIONS

We use first-order perturbation theory for eigenvalues and eigenvectors to investigate the behaviour of deflation and coarse preconditioners for SPD linear systems. Our analysis shows a better stability of the preconditioners M_1 and M_{+1} compared to the other preconditioning variants that exploit some additional properties that are only true for exact eigenpairs. An important result of this work is that the efficiency of a spectral preconditioner should not be assessed only using exact eigenpairs. In practice these pre-

conditioners may be built using approximate information. Such a situation occurs for instance when a sequence of linear systems have to be solved. On unsymmetric problems we show how such a solution technique can be implemented and we illustrate its efficiency on large industrial electromagnetics examples.

Acknowledgments

We would like to thank Guillaume Alléon and Guillaume Sylvand from EADS-CCR, for providing us with some test examples. We also thanks E. Martin, who performed some of the numerical experiments reported here. Finally, we would like to acknowledge the computing support provided by the CINES (Centre Informatique National de l'Enseignement Supérieur) and the CEA (Commissariat à l'Énergie Atomique) that enabled us to access their computing facilities.

REFERENCES

- [1] 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.
- [2] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1995.
- [3] J. Baglama, D. Calvetti, G. H. Golub, and L. Reichel. Adaptively preconditioned GMRES algorithms. *SIAM J. Scientific Computing*, 20(1):243–269, 1999.
- [4] M. Benzi and M. Tuma. A sparse approximate inverse preconditioner for nonsymmetric linear systems. *SIAM J. Scientific Computing*, 19(3):968–994, 1998.
- [5] J. H. Bramble, J. E. Pasciak, and A. H. Schatz. The construction of preconditioners for elliptic problems by substructuring I. *Math. Comp.*, 47(175):103–134, 1986.
- [6] B. Carpentieri, I. S. Duff, and L. Giraud. A class of spectral two-level preconditioners. *SIAM J. Scientific Computing*, 25:749–765, 2003.
- [7] B. Carpentieri, I. S. Duff, and L. Giraud. A class of spectral two-level preconditioners. *SIAM J. Scientific Computing*, 25(2):749–765, 2003.
- [8] B. Carpentieri, I. S. Duff, L. Giraud, and G. Sylvand. Combining fast multipole techniques and an approximate inverse preconditioner for large electromagnetism calculations. *SIAM J. Scientific Computing*, 27(3):774–792, 2005.
- [9] B. Carpentieri, L. Giraud, and S. Gratton. Additive and multiplicative spectral two-level preconditioners for general linear systems. Technical Report TR/PA/04/38, CERFACS, Toulouse, France, 2004.

- [10] I. S. Duff, L. Giraud, J. Langou, and E. Martin. Using spectral low rank preconditioners for large electromagnetic calculations. *Int. Journal of Numerical Methods in Engineering*, 62(3):416–434, 2005.
- [11] I. S. Duff, L. Giraud, J. Langou, and E. Martin. Using spectral low rank preconditioners for large electromagnetic calculations. *Int. J. Numerical Methods in Engineering*, 62(3):416–434, 2005.
- [12] J. Erhel, K. Burrage, and B. Pohl. Restarted GMRES preconditioned by deflation. *J. Comput. Appl. Math.*, 69:303–318, 1996.
- [13] M. Fisher. Minimization algorithms for variational data assimilation. In *Proc. ECMWF seminar "Recent developments in numerical methods for atmospheric modelling*, pages 364–385, September 1998.
- [14] J. Frank and C. Vuik. On the construction of deflation-based preconditioners. *SIAM J. Scientific Computing*, 23:442–462, 2001.
- [15] R. W. Freund and M. Malhotra. A block QMR algorithm for non-Hermitian linear systems with multiple right-hand sides. *Linear Algebra and its Applications*, 254:119–157, 1997.
- [16] L. Giraud and S. Gratton. On the sensitivity of some spectral preconditioners. *SIAM Journal on Matrix Analysis and Applications*, 27(4):1089–1105, 2006.
- [17] L. Giraud, S. Gratton, and E. Martin. Incremental spectral preconditioners for sequences of linear systems. *Applied Numerical Mathematics*, Accepted.
- [18] L. Giraud, J. Langou, and G. Sylvand. On the parallel solution of large industrial wave propagation problems. *Journal of Computational Acoustics*, 14(1):83–111, 2006.
- [19] 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 edition, 1996.
- [20] M. Grote and T. Huckle. Parallel preconditionings with sparse approximate inverses. *SIAM J. Scientific Computing*, 18(3):838–853, 1997.
- [21] R. B. Morgan. A restarted GMRES method augmented with eigenvectors. *SIAM J. Matrix Analysis and Applications*, 16(4):1154–1171, 1995.
- [22] R. B. Morgan. Implicitly restarted GMRES and Arnoldi methods for nonsymmetric systems of equations. *SIAM J. Matrix Analysis and Applications*, 21(4):1112–1135, 2000.

- [23] R. B. Morgan. GMRES with deflated restarting. *SIAM J. Scientific Computing*, 24(1):20–37, 2002.
- [24] R. Nabben and C. Vuik. A comparison of deflation and coarse grid correction applied to porous media flow. *SIAM Journal on Numerical Analysis*, 42:1631–1647, 2004.
- [25] D. P. O’Leary. The block conjugate gradient algorithm and related methods. *Linear Algebra and its Applications*, 29:293–322, 1980.
- [26] A. Padiy, O. Axelsson, and B. Polman. Generalized augmented matrix preconditioning approach and its application to iterative solution of ill-conditioned algebraic systems. submitted to SIAM Journal of Matrix Anal. and Appl., to appear.
- [27] M. L. Parks, E. de Sturler, G. Mackey, D. D. Jhonson, and S. Maiti. Recycling Krylov subspaces for sequences of linear systems. Technical Report UIUCDCS-R-2004-2421 (CS), University of Illinois at Urbana-Champaign, 2004.
- [28] J. R. Rice. A theory of condition. *SIAM Journal on Numerical Analysis*, 3:287–310, 1966.
- [29] J.-C. Rioual. *Solving linear systems for semiconductor device simulations on parallel distributed computers*. PhD thesis, CERFACS, Toulouse, France, 2002.
- [30] Y. Saad. ILUT: A dual threshold incomplete LU factorization. *Numerical Linear Algebra with Applications*, 1:387–402, 1994.
- [31] V. Simoncini and E. Gallopoulos. An iterative method for nonsymmetric systems with multiple right-hand sides. *SIAM J. Scientific Computing*, 16(4):917–933, 1995.
- [32] G. W. Stewart. *Matrix algorithms. Volume II: Eigensystems*. Society for Industrial and Applied Mathematics, 2001.
- [33] G. Sylvand. *La méthode multipôle rapide en électromagnétisme : performances, parallélisation, applications*. PhD thesis, Ecole Nationale des Ponts et Chaussées, 2002.
- [34] B. Vital. *Étude de quelques méthodes de résolution de problèmes linéaires de grande taille sur multiprocesseur*. Ph.D. dissertation, Université de Rennes, November 1990.
- [35] H. Waisman, J. Fish, R. S. Tuminaro, and J. Shadid. The generalized global basis method. *Int J. Numerical Methods in Engineering*, 61:1243–1269, 2004.