Flexible GMRES with deflated restarting

L. Giraud* S. Gratton† X. Pinel‡ X. Vasseur‡

CERFACS Technical Report TR/PA/08/128

Abstract

In many situations, it has been observed that significant convergence improvements can be achieved in preconditioned Krylov subspace methods by enriching them with some spectral information. On the other hand effective preconditioning strategies are often designed where the preconditioner varies from one step to the next so that a flexible Krylov solver is required. In this paper, we present a new numerical technique for non-symmetric problems that combines these two features. We illustrate the numerical behavior of the new solver both on a set of small academic test examples as well as on large industrial simulation arising in wave propagation simulations.

1 Introduction

The solution of large linear systems is a basic kernel in many large scale simulations and preconditioned Krylov subspace methods are among the most popular linear solvers. For non-symmetric problems the GMRES [22] method is often chosen because of its robustness [17, 18] and because the Euclidean norm of the system residual is decreasing along the iterations. In order to make the GMRES method affordable from a memory and operation-count point of view, a restarting process has to be implemented. In the classical restarted GMRES approach, the initial guess at restart is chosen to be the best (for the residual norm) known iterate, enforcing thereby the decrease of the residual norm even when a restart is performed. In such a situation the restart is performed with only one vector. In addition, it has been observed that reusing part of the current Krylov space (and not only one vector) for the construction of iterates in the next cycle of GMRES might significantly improve the convergence. In many approaches, some estimate of the invariant subspace is searched in the Krylov subspace and reused in the next restart either by augmenting the space [3, 14, 21], by deflating over the subspace [16] or by ensuring some orthogonality properties with respect to that space [19]. One of the most recent work in this field based on a deflation approach is GMRES-DR [16]. This method reduces to GMRES, when no deflation is applied, but may provide a much faster convergence than GMRES for well chosen deflation spaces as described in [16].

The methods mentioned above suppose that the preconditioner is a given matrix M that is not allowed to change along the iterations. However, there are situations where this is not true anymore, as e.g. in domain decomposition methods, when approximate solvers are considered for the interior problems (see references in [24, Sect. 4.4] or in [26, Sect. 4.3]). This approach is notably used when the size of the local subproblems is so large that obtaining an approximate solution using an iterative method is computationally more interesting than using a direct method. If the domain decomposition preconditioner is based on the use of approximate solvers, its application is not a linear operation in general, and *flexible methods*, such as the Flexible GMRES method (see the FGMRES method in [20]), are designed to handle this situation.

In this paper, we present a new approach that combines flexible iterations and a restarting strategy that exploits some spectral information. The paper is organized as follows: in Section 2

*INPT-IRIT, University of Toulouse and ENSEEIHT, 2 rue Camichel, BP 7122, F-31071 Toulouse cedex 7, France

†CNES, avenue Edouard Belin, F-31401 Toulouse cedex 4, France and CERFACS, 42, Avenue Gaspard Coriolis, F-31057 Toulouse Cedex 1, France

‡CERFACS, 42, Avenue Gaspard Coriolis, F-31057 Toulouse Cedex 1, France

we describe the FGMRES-DR method that combines the GMRES-DR [16] and the FGMRES [20] algorithms. Section 3 is devoted to numerical experiments where both academic and real life problems are considered to illustrate the numerical features of the new solver.

For the sake of generality we describe here the approaches for complex-valued linear systems. Note that everything also specializes to real arithmetic calculation. Let $A \in \mathbb{C}^{n\times n}$ be a square nonsingular $n \times n$ non-symmetric complex matrix, and $b \in \mathbb{C}^n$ be a complex vector of length n, that both define the linear system

$$Ax = b. \tag{1}$$

The solution techniques start from an initial guess $x_0 \in \mathbb{C}^n$ and they compute an approximation in a ℓ-dimensional affine space $x_0 + V_\ell\, y$ where $V_\ell \in \mathbb{C}^{n\times \ell}$. The basis of the ℓ-dimensional space spanned by the columns of V_ℓ is built using an Arnoldi process that enjoys the following equality:

$$AV_\ell = V_{\ell+1}\bar{H}_\ell$$

such that $r_0 = b - Ax_0$ belongs to the space V_ℓ. Here $\bar{H}_\ell$ denotes the $(\ell+1)\times\ell$ upper Hessenberg matrix defined in the Arnoldi process. Among the methods that are based on the construction of these spaces we consider the variant that searches for the approximation associated with a minimum norm residual property such as GMRES, where the considered norm $\|\cdot\|$ is the Euclidean norm.

2 The FGMRES-DR method

2.1 Derivation of the method

In many large-scale scientific and industrial applications one might not be able to consider a fixed preconditioner at each step of the GMRES method. This happens for instance when block preconditioners, including domain decomposition techniques, are considered where the blocks are too large to be handled by a direct solver. In such situations, an iterative solver has to be implemented to solve linear systems involving these blocks. Consequently the preconditioner varies from one step to the next and *flexible* Krylov solvers have been developed to manage this issue [20, 28].

Standard flexible iterative methods implement a scheme where, after a certain number of steps (denoted by m in this paper), the Krylov subspace is truncated and the method is restarted, in order to control the memory requirements and the cost of the orthogonalization scheme of the method. In the FGMRES method, the method is simply stopped and then restarted, taking as an initial guess the past iterate with the smallest residual norm. In the GMRES-DR algorithm, a more sophisticated scheme has been developed, but only in the case where a *fixed preconditioner* M is used; a full subspace of dimension k, $k < m$ (and not only the approximate solution with minimum residual norm) is retained in the restarting scheme and the success of this approach has been demonstrated on many academic examples [14]. We show in this section how GMRES-DR can be extended to handle the situation where the preconditioner is allowed to vary. Following the FGMRES description (proposed in [20]) we denote by M_i a matrix that represents the preconditioner at step i of the method. Algorithm 1 depicts the FGMRES-DR algorithm that we propose. As the Flexible GMRES algorithm, starting from an initial guess x_0, it generates, at each restart, the matrices $Z_m \in \mathbb{C}^{n\times m}$, $V_{m+1} \in \mathbb{C}^{n\times(m+1)}$ and $\bar{H}_m \in \mathbb{C}^{m+1\times m}$ such that $AZ_m = V_{m+1}\bar{H}_m$. An approximate solution $x_m \in \mathbb{C}^n$ is found by minimizing the residual norm $\|b - A(x_0 + V_m y)\|$ over the space $x_0 + \text{range}(V_m)$, the corresponding residual being $r_m = b - Ax_m \in \mathbb{C}^n$, with $r_m \in \text{range}(V_{m+1})$. The only step which remains unspecified is the restarting procedure, where V_{k+1}^{new}, Z_k^{new} and $\bar{H}_k^{new}$ are computed so that Equation (4) holds. Proposition 1 and Algorithm 2 show how this operation can be efficiently implemented. It relies on the use of harmonic Ritz vectors of a certain matrix with respect to a subspace defined by range(V_m). Let define a subspace $S \subset \mathbb{C}^n$ and a matrix $B \in \mathbb{C}^{n\times n}$. The pair $(\theta, y) \in \mathbb{C}\times S$ is referred to as a harmonic Ritz pair of B with respect to BS if and only if [23, Theorem. 5.1]:

$$S^H B^H (By - \theta y) = 0, \tag{2}$$

where y is called the harmonic Ritz vector associated with the harmonic Ritz value θ.

Before showing how Equation (4) is obtained, we start with some comments on the FGMRES-DR algorithm:

Algorithm 1 FGMRES with deflated restarting

1: *Initialization:* Choose $m > 0$, $k > 0$, $tol > 0$, $x_0 \in \mathbb{C}^n$ Let $r_0 = b - Ax_0$; $\beta = \|r_0\|$, $c = \beta e_1$, $v_1 = r_0/\beta$.

2: *First restart:* Apply the standard Arnoldi process with flexible preconditioner to construct V_{m+1}, Z_m and $\bar{H}_m$ so that:

$$AZ_m = V_{m+1}\bar{H}_m. \tag{3}$$

3: *Minimum norm solution:* Compute the minimal norm solution in the affine space $x_0 + \text{range}(Z_m)$; that is, $x_m = x_0 + Z_m y_m$ where $y_m = \text{argmin}_y \|c - \bar{H}_m y\|$. Set $x_0 = x_m$ and $r_0 = b - Ax_0$.

4: *Check convergence:* If $\|c - \bar{H}_m y_m\|/\|b\| \leq tol$, Exit

5: *Restarting procedure:* Perform a QR factorization of $[u_1, \dots, u_k]$ and store the Q-factor in V_k^{new}, where $u_i \in \mathbb{C}^n$, $i = 1, \dots, k$ are harmonic Ritz vectors of $AZ_m V_m^H$ with respect to $\text{range}(V_m)$ (see Algorithm 2). Use the Gram-Schmidt process to obtain v_{k+1}^{new}, such that $r_0 \in \text{range}(V_{k+1}^{new})$ holds, where $V_{k+1}^{new} = [V_k^{new}, v_{k+1}^{new}]$ is orthonormal. Compute Z_k^{new} and $\bar{H}_k^{new}$ so that

$$AZ_k^{new} = V_{k+1}^{new}\bar{H}_k^{new}. \tag{4}$$

6: Apply $(m-k)$ additional steps of the Flexible Arnoldi process on (4) so that:

$$AZ_m^{new} = V_{m+1}^{new}\bar{H}_m^{new}. \tag{5}$$

7: Restart: Set $c = (V_{k+1}^{new})^H r_0$, $Z_m = Z_m^{new}$, $V_{m+1} = V_{m+1}^{new}$, $\bar{H}_m = \bar{H}_m^{new}$, Goto 3

Algorithm 2 Computation of V_{k+1}^{new}, Z_k^{new} and $\bar{H}_k^{new}$

1: *Input:* A, Z_m and V_{m+1} such that $AZ_m = V_{m+1}\bar{H}_m$

2: *Compute k harmonic Ritz vectors.* Compute k independent eigenvectors g_i of the matrix $H_m + h_{m+1,m}^2 H_m^{-H} e_m e_m^T$ with $e_m^H = (0_{m-1}, 1)$ where 0_{m-1} is the $1 \times m-1$ zero row vector and H_m denotes the first m rows of $\bar{H}_m$. Set $G_k = [g_1, \dots, g_k] \in \mathbb{C}^{m \times k}$.

3: *Augmentation of G_k:*
Set

$$G_{k+1} = \left[\begin{bmatrix} G_k \\ 0_k \end{bmatrix}, c - \bar{H}_m y_m \right],$$

where 0_k is the $1 \times k$ zero row vector.

4: *Orthonormalization of the columns of G_{k+1}:* Make a QR-factorization of G_{k+1} as $G_{k+1} = P_{k+1}\Gamma_{k+1}$ and store P_{k+1}. Define, in Matlab notation, $P_k = P_{k+1}(1:k, 1:k)$.

5: Set $V_{k+1}^{new} = V_{m+1}P_{k+1}$, $Z_k^{new} = Z_m P_k$ and $\bar{H}_k^{new} = P_{k+1}^H \bar{H}_m P_k$.

1. **Breakdown of the algorithm**. Steps 2. and 6. of Algorithm 1 implement the non-restarted FGMRES algorithm, that may break down whereas the solution of the linear system has not been found (see [20]). Indeed, a breakdown occurs at step j when $AM_j v_j$ belongs to $\text{range}(V_j)$; this corresponds to a very particular property of M_j with respect to A and V_j, that, from our point of view, is not very likely to hold when M_j represents an iterative preconditioning step. Note that, from the GMRES theory, this situation cannot occur if the preconditioner is a nonsingular (constant) matrix M (see [20]) because in this case, when the algorithm breaks down, the solution is found. In addition, FGMRES-DR also relies on being able to compute k distinct eigenpairs of a matrix in Step 2. of Algorithm 2. This may not be possible if the matrix is not diagonalizable. Again, we consider as very unlikely in finite precision arithmetic that an eigenvalue solver like the QR method fails to find k linearly independent vectors of a given matrix. Strictly speaking, if a problem occurred in Steps 2. and 6. of Algorithm 1 or in Step 2. of Algorithm 2, the algorithm would effectively breakdown, but as for FGMRES and GMRES-DR, this is considered as very unlikely in practical computations.

2. **Use of the harmonic Ritz values associated with matrix $AZ_mV_m^H$ in Step 5. of Algorithm 1**. First we remind the situation for a constant preconditioner $M_i = M$ and define

$Z_m = MV_m$. The preconditioned GMRES-DR algorithm uses, in its restarting procedure, the harmonic Ritz vectors of AM w.r.t. V_m. From Equation (2), these vectors are readily the harmonic Ritz vectors of $AMV_mV_m^H = AZ_mV_m^H$ w.r.t. V_m, which is what we retained in FGMRES-DR. It follows that the restarting procedure of Algorithm 1 can be viewed as a generalization of the corresponding operation in the GMRES-DR algorithm to the flexible preconditioning framework.

Proposition 1. *Let P_k be the matrix defined in Algorithm 2. At each restart of FGMRES-DR the Arnoldi relation*

$$AZ_k^{new} = V_{k+1}^{new}\bar{H}_k^{new}$$

holds with

$$Z_k^{new} = Z_mP_k,$$
$$V_{k+1}^{new} = V_{m+1}P_{k+1},$$

and

$$\bar{H}_k^{new} = P_{k+1}^H\bar{H}_mP_k.$$

Proof. For a given Hessenberg matrix $\bar{H}_m$, our definition of the harmonic Ritz vectors g_i is the same as the definition used in [15]. Using [15, Theorem 5.5, p. 1121] we then get, for some $\alpha_i \in \mathbb{C}$,

$$\forall i \in \{1,\ldots,k\},\ \bar{H}_mg_i - \lambda_i\begin{pmatrix} g_i \\ 0 \end{pmatrix} = \alpha_i(c - \bar{H}_my_m) = \alpha_i\rho_m,$$

where $\rho_m = c - \bar{H}_my_m$. Multiplying the above equality by V_{m+1} on the left, we obtain

$$\forall i \in \{1,\ldots,k\},\ V_{m+1}\bar{H}_mg_i - \lambda_iV_{m+1}\begin{pmatrix} g_i \\ 0 \end{pmatrix} = \alpha_iV_{m+1}\rho_m,$$

which in turn yields, using the Arnoldi relation $AZ_m = V_{m+1}\bar{H}_m$

$$\forall i \in \{1,\ldots,k\},\ AZ_mg_i = V_{m+1}\left(\lambda_i\begin{pmatrix} g_i \\ 0 \end{pmatrix} + \alpha_i\rho_m\right). \tag{6}$$

Setting $G_k = [g_1,\cdots,g_k] \in \mathbb{C}^{m\times k}$ and $\alpha = [\alpha_1,\cdots,\alpha_k] \in \mathbb{C}^{1\times k}$, Equation (6) reads

$$AZ_mG_k = V_{m+1}\left[\begin{pmatrix} G_k \\ 0_k \end{pmatrix}, \rho_m\right]\begin{bmatrix} \operatorname{diag}(\lambda_1,\cdots,\lambda_k) \\ \alpha \end{bmatrix}. \tag{7}$$

Let us further denote by $G_k = P_k\Gamma_k$ the QR-factorization of G_k and orthogonalize the vector ρ_m against the columns of $\begin{bmatrix} P_k \\ 0_k \end{bmatrix}$ to obtain p_{k+1}. We have $ap_{k+1} = \rho_m - \begin{bmatrix} P_k \\ 0_k \end{bmatrix}u$ where $a = \|\rho_m - \begin{bmatrix} P_k \\ 0_k \end{bmatrix}u\|$ and $u_i = \begin{pmatrix} p_i \\ 0 \end{pmatrix}^H\rho_m$. The orthonormal matrix P_{k+1} writes $P_{k+1} = \left[\begin{pmatrix} P_k \\ 0_k \end{pmatrix}, p_{k+1}\right]$.

Regarding the Γ_{k+1} factor we have

$$\begin{aligned}\left[\begin{pmatrix} G_k \\ 0_k \end{pmatrix}, \rho_m\right] &= \left[\begin{pmatrix} P_k\Gamma_k \\ 0_k \end{pmatrix}, ap_{k+1} + \begin{pmatrix} P_k \\ 0_k \end{pmatrix}u\right] \\ &= \left[\begin{pmatrix} P_k \\ 0_k \end{pmatrix}, p_{k+1}\right]\begin{bmatrix} \Gamma_k & u \\ 0_k & a \end{bmatrix},\end{aligned}$$

from which follows that $\Gamma_{k+1} = \begin{bmatrix} \Gamma_k & u \\ 0_k & a \end{bmatrix}$.

Consequently, the following equality (7) now becomes

$$AZ_mP_k = V_{m+1}P_{k+1}\Gamma_{k+1}\begin{bmatrix} \operatorname{diag}(\lambda_1,\cdots,\lambda_k) \\ \alpha \end{bmatrix}\Gamma_k^{-1}.$$

Note that Γ_k is non singular if range(G_k) is equal to k, which is assumed, as explained above in the comment 1. on the breakdown of the algorithm.

Let us denote $Z_k^{new} = Z_m P_k$, $V_{k+1}^{new} = V_{m+1} P_{k+1}$ and

$$\bar{H}_k^{new} = \Gamma_{k+1} \begin{bmatrix} \mathrm{diag}(\lambda_1, \cdots, \lambda_k) \\ \alpha \end{bmatrix} \Gamma_k^{-1}.$$

This enables us to write the following Arnoldi relation

$$A Z_k^{new} = V_{k+1}^{new} \bar{H}_k^{new}.$$

We now derive the expression of $\bar{H}_k^{new}$ as the product of matrices of low dimension (basically of order k and m). The collinearity relation $\bar{H}_m g_i - \lambda_i \begin{pmatrix} g_i \\ 0 \end{pmatrix} = \alpha_i \rho_m$ reads in matrix form $\bar{H}_m G_k - \begin{bmatrix} G_k \\ 0_k \end{bmatrix} \mathrm{diag}(\lambda_1, \cdots, \lambda_k) = \rho_m \alpha$. Since $G_k = P_k \Gamma_k$, we have

$$\begin{aligned} \bar{H}_m P_k \Gamma_k &= P_{k+1} \Gamma_{k+1} \begin{bmatrix} \mathrm{diag}(\lambda_1, \cdots, \lambda_k) \\ \alpha \end{bmatrix}, \\ P_{k+1}^H \bar{H}_m P_k &= \Gamma_{k+1} \begin{bmatrix} \mathrm{diag}(\lambda_1, \cdots, \lambda_k) \\ \alpha \end{bmatrix} \Gamma_k^{-1}. \end{aligned}$$

Hence $\bar{H}_k^{new} = P_{k+1}^H \bar{H}_m P_k$ as it is defined in Algorithm 2.

□

We now consider the right-hand side c of the least-squares problem over y that has to be solved at each iteration:

$$y^* = \mathrm{argmin}_y \|c - \bar{H}_m y\|.$$

In the classical GMRES-DR algorithm the right-hand side c of the least-squares problems is computed as $c = V_{m+1}^H r_0$. However this computation might be expensive as it requires a matrix-vector product and $(m+1)$ dot products. This cost can be first reduced by observing that the residual r_0 belongs to the span of V_{k+1}, consequently only its first $(k+1)$ entries are non-zero and they can be computed as $V_{k+1}^H r_0$. By further exploiting this observation the calculation of c can be even more reduced as described in Proposition 2.

Proposition 2. *At restart the new residual r_0 lies in the span of the orthonormal basis V_{k+1}^{new}. Its components (that define the right-hand side of the new least-squares problem to be solved in the next restart) are given by the last column of the R factor of the QR factorization of the matrix*

$$\left[\begin{pmatrix} G_k \\ 0_k \end{pmatrix}, c - \bar{H}_m y_m \right]$$

Proof. Let $P_{k+1}\Gamma_{k+1}$ denote the QR factorization of $\left[\begin{pmatrix} G_k \\ 0_k \end{pmatrix}, c - \bar{H}_m y_m \right]$ with $\begin{pmatrix} u \\ a \end{pmatrix}$ be the last column of Γ_{k+1}. We have $c - \bar{H}_m y_m = P_{k+1} \begin{pmatrix} u \\ a \end{pmatrix}$. Consequently $r_0 = V_{m+1}(c - \bar{H}_m y_m) = V_{m+1} P_{k+1} \begin{pmatrix} u \\ a \end{pmatrix} = V_{k+1}^{new} \begin{pmatrix} u \\ a \end{pmatrix}$.

□

2.2 Practical implementation

The results shown in Propositions 1 and 2 enable the design of the final and effective implementation of the FGMRES-DR method described in Algorithm 3. Even though the algorithm has been described in Section 2.1 for complex arithmetic, it can be also used for real-valued matrices in real arithmetic by separating the real and imaginary parts of the harmonic Ritz vectors in Step 6. of Algorithm 3, as already suggested in the GMRES-DR algorithm [16].

Step 10. of Algorithm 3 offers two possible ways to evaluate the projected residual c, as suggested in Proposition 2. The expression involving γ_{k+1} is attractive because it avoids the matrix vector product $V_{k+1}^{new} r_0$. From our (somehow limited) experience either formula seem to produce the same convergence curves. Further work is therefore needed to fully understand which approach is the best in presence of round-off errors.

Concerning memory comparison, targeting for large scale problems, where n is usually much greater than m and k, we neglect all matrices and vectors, the dimension of which involves only m (and not n). With this convention FGMRES is twice as much expensive in memory as GMRES, and the same holds for FGMRES-DR and GMRES-DR. Indeed, in the flexible variant of the algorithms, memory is needed to store both the Krylov basis V and its preconditioned counterpart Z. An optimization of the memory is however still possible. It can indeed be seen that the variables Z^{new} and V^{new} can overwrite the variables Z and V. This can be accomplished by performing the matrix multiplications $V_{k+1} \leftarrow V_{m+1}P_{k+1}$, $Z_k \leftarrow Z_m P_k$ of Step 8. *in place*, i.e. within the arrays V and Z. The trick comes from the fact that multiplications involving triangular factors can be done in place. It is therefore feasible to perform an LU factorization with total pivoting of P_k (to get a very good approximation $P_k = LU$), and then, to perform successively the operations $X \leftarrow UX$ and $X \leftarrow LX$, for X being V and Z. This approach is clearly saving a lot of memory when k is close to m, but may introduce additional round-off errors that can hopefully be monitored by inspecting the ratio $\dfrac{\|P_k - LU\|}{\|P_k\|}$.

Algorithm 3 Practical implementation of FGMRES with deflated restarting: FGMRES-DR(m,k)

1: *Initialization.* Choose:
m: the maximum size of the subspace where the solution is built,
k: the desired number of approximate eigenvectors.
tol: the convergence threshold , x_0: an initial guess
Let $r_0 = b - Ax_0$; $\beta = \|r_0\|$, $c = \beta e_1$, $v_1 = r_0/\beta$.

2: *First iteration.* Apply standard FGMRES(m). It generates V_{m+1}, $\bar{H}_m$, Z_m.

3: *Minimum norm solution.* Solve $y_m = \text{argmin}_y ||c - \bar{H}_m y||$. Compute the approximate solution $x_m = x_0 + Z_m y_m$ and the residual $r_m = V_{m+1}^H (c - \bar{H}_m y_m)$. Set $x_0 = x_m$ and $r_0 = r_m$.

4: *Check convergence.* If $||r_0||/||b|| = ||c - \bar{H}_m y_m||/||b|| \le tol$, Exit.

5: *Compute harmonic Ritz vectors.* Compute k independent eigenvectors g_i of the matrix: $H_m + h_{m+1,m}^2 H_m^{-H} e_m e_m^T$. Then store the g_i, $i = 1 \dots k$, into the matrix G_k.

6: *Append the quasi-residual $\beta e_1 - \bar{H}_m y_m$ to G_k.*
$\star$ If A is complex: add a zero row to G_k and the column $c - \bar{H}_m y_m$.
$\star$ If A is real: if some of the g_i's are complex, split their real and imaginary parts (in the case, there would be more vectors than k, remove a real one). Then add a zero row to G_k and the column $c - \bar{H}_m y_m$. In the two cases, the new matrix is noted G_{k+1} and its dimensions are $(m+1) \times (k+1)$.

7: *Orthonormalization of the columns of G_{k+1}.* Make a QR-factorization of G_{k+1}: $G_{k+1} = P_{k+1}\Gamma_{k+1}$, store P_{k+1} and the last column of Γ_{k+1}, γ_{k+1}.

8: Let $V_{k+1}^{new} = V_{m+1}P_{k+1}$, $Z_k^{new} = Z_m P_k$ and $\bar{H}_k^{new} = P_{k+1}^H \bar{H}_m P_k$.

9: Apply $(m-k)$ additional steps of the Flexible Arnoldi process on $AZ_k^{new} = V_{k+1}^{new}\bar{H}_k^{new}$ so that $AZ_m^{new} = V_{m+1}^{new}\bar{H}_m^{new}$.

10: Restart: Set $c = (V_{k+1}^{new})^H r_0 = \begin{bmatrix} I_{k+1} \\ 0_{k+1} \end{bmatrix} \gamma_{k+1}$, $Z_m = Z_m^{new}$, $V_{m+1} = V_{m+1}^{new}$, $\bar{H}_m = \bar{H}_m^{new}$, Goto 3. In these expressions, 0_{k+1} is the $1 \times (k+1)$ zero row vector and I_{k+1} is the order $k+1$ identity matrix.

---

3 Numerical experiments

In this section we investigate the numerical behaviour of the FGMRES-DR(m,k) algorithm on both academic and realistic applications. We consider the case of both sparse or dense matrices in either real or complex arithmetic. All the examples include a detailed comparison with FGMRES(m). This allows us to show the effects of incorporating the deflation strategy in the flexible preconditioning framework.

In the following experiments, the right-hand sides are computed as $b = A\mathbf{1}$ where $\mathbf{1}$ is the vector of appropriate dimension with all components equal to one. A zero initial iterate x_0 is considered as an initial guess and the following stopping criterion is used:

$$\frac{\|b - A\,x_\ell\|}{\|b\|} \leq 10^{-12} \tag{8}$$

where ℓ represents the step when the iterations are stopped.

3.1 Harwell-Boeing and Matrix Market test problems

In order to illustrate the numerical behaviour of FGMRES-DR, we first consider a few test matrices from the Harwell-Boeing [11] and Matrix Market [2] libraries so that any reader could reproduce these experiments. The sparse matrices named Sherman4, Saylor4 and Young1c have been chosen. Sherman4 and Saylor4 are real matrices, whereas Young1c is a complex-valued one. They represent challenging sparse matrices coming from realistic applications (reservoir modelling, acoustics) that are often used to analyze the behaviour of numerical algorithms. For those experiments, the preconditioner consists in five steps of preconditioned full GMRES, where the preconditioner is based on an ILU(0) factorization. In the case of Sherman4 only, the inner solver corresponds to five steps of unpreconditioned full GMRES.

In Table 1, we depict the total number of matrix-vector products performed in the inner and outer parts of the solver (# Mv), the total number of dot products (# dot) for three combinations of flexible methods: FGMRES-DR(5,3), FGMRES(5) and FGMRES(7) respectively. In order to illustrate the possible benefit of using the implementation trick that alleviates the memory penalty due to the deflation at restart we effectively consider different restart parameters for FGMRES. Indeed the performance of FGMRES-DR(5,3) can be compared with FGMRES(5) if the trick is implemented or with FGMRES(7) if a straightforward implementation is considered. The total amount of floating-point operations including the cost of preconditioning has been computed for each solution method, excluding however the cost of the ILU(0) factorization that is the same for each proposed method. We choose FGMRES(7) as a reference solution method and report in Table 1 the ratio in terms of floating-point operations required by a given solution method compared with the reference solution method. In the three cases, it can be noticed that the deflation at restarting enables a faster convergence. It also results in a faster calculation since a significant amount of floating-point operations is saved.

	Sherman4			Saylor4			Young1c		
	# Mv	# dot	ratio	# Mv	# dot	ratio	# Mv	# dot	ratio
FGMRES-DR(5,3)	373	1288	0.42	115	384	0.44	1633	5698	0.35
FGMRES(5)	1273	3813	1.08	409	1221	1.36	6145	18430	1.07
FGMRES(7)	877	2771	1.00	295	931	1.00	5095	16126	1.00

Table 1: Performance of FGMRES and FGMRES-DR to satisfy the convergence threshold (8); # Mv is the total number of matrix vector products, # dot the total number of dot products and ratio the ratio of floating-point operations with respect to the FGMRES(7) method.

3.2 Two-dimensional Helmholtz problem

Our goal in this section is to illustrate the performance of FGMRES with deflated restarting on a simple two-dimensional partial differential equation model problem. A model wave propagation problem in a two-dimensional homogeneous medium is considered:

$$-\frac{\partial}{\partial x}(\frac{\partial u}{\partial x}) - \frac{\partial}{\partial y}(\frac{\partial u}{\partial y}) - \sigma^2\, u \quad = \quad f \qquad \text{in} \qquad \Omega =]0,1[^2 \tag{9}$$

with homogeneous Dirichlet boundary conditions $u = 0$ on the boundary $\delta\Omega$. The unknown u represents the pressure field in the frequency domain and σ the constant wavenumber. A second order finite difference discretization scheme of the Helmholtz equation (9) is used on a equidistant Cartesian grid of step size h with the following dispersion stability condition $\sigma\, h = 0.625$ [6] being satisfied. One V(1,1) cycle of a geometric multigrid method [27] is used as a preconditioner. This multigrid method uses a two-level hierarchy with a red-black Gauss-Seidel smoother, bilinear interpolation as prolongation and its adjoint as restriction operator. Galerkin coarse grid discretization is employed to build the coarse grid operator and a sparse direct solution method is used to solve the coarse grid systems. Numerical experiments with this two-grid preconditioner for FGMRES on two-dimensional wave propagation problems in geophysics with Robin boundary conditions have been reported in [10]. The discretization of the Helmholtz equation on a 64×64 grid with $\sigma = 40$ leads to a real-valued sparse symmetric indefinite matrix A, whose spectrum is shown in Figure 1 (left part). There are 117 negative eigenvalues for this choice of wavenumber and step size. The spectrum of the preconditioned operator AM is also shown in Figure 1 (right part). It exhibits both positive and negative isolated real eigenvalues and a cluster of eigenvalues around $(1, 0)$.

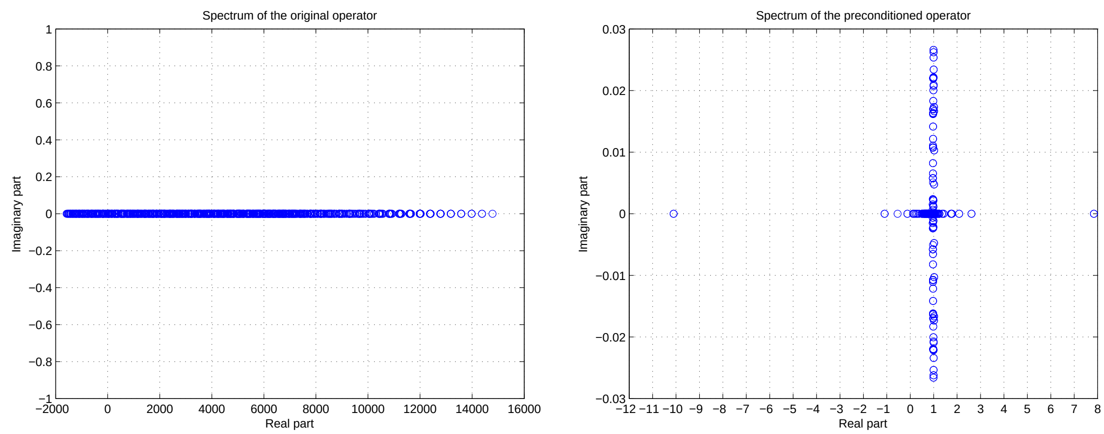

Figure 1: Spectrum of the original Helmholtz operator (left) and spectrum of the preconditioned operator (right), when a V(1,1) cycle of a multigrid method is used as a preconditioner. Case of $h = 1/64$ with $\sigma = 40$. Note the different scalings on both figures.

The harmonic Ritz vectors corresponding to the k eigenvalues λ_i of smallest magnitude of the matrix $H_m + h_{m+1,m}^2 H_m^{-H} e_m e_m^T$ have been chosen in Algorithm 3 (step 5.). In the sequel we call this strategy Smallest. However any combination of k harmonic vectors may be selected. Thus we have considered two other possibilities. The first one selects the k eigenpairs corresponding to the eigenvalues of largest magnitude. It is called Largest. The second deflation strategy retains the k eigenvectors associated with the eigenvalues such as $|1 - \lambda_i|$ is of largest magnitude. With this latter choice we aim at selecting eigenvalues located away from a cluster around the eigenvalue of the "ideal" preconditioned operator AM with $M^{-1} = A$. This eventually allows simultaneous deflation of eigenvalues of both smallest and largest magnitude. We call this strategy Cluster. We investigate the influence of the different deflation strategies (Smallest, Largest and Cluster respectively) and compare FGMRES-DR(m,k) with FGMRES(m) for different values of the restart parameter m. Table 2 gives the number of approximate eigenpairs k that led to the smallest number of iterations ℓ to satisfy the stopping criterion (8) for each deflation strategy. These results have been obtained by running FGMRES-DR(m,k) with $1 \leq k < m$ for each value of the restart parameter m. Numerical results show that the Cluster deflation strategy is almost the most efficient on this application (see bold values in Table 2) leading sometimes to a significant reduction in terms of iterations. The total amount of floating-point operations including the cost of preconditioning has been computed for each solution method. We choose FGMRES(m) as a reference solution method and report the following normalized quantity in Table 2:

$$ratio = \frac{flops_{Strat}}{flops_{FGMRES(m)}} \qquad (10)$$

where $Strat$ denotes the FGMRES-DR(m,k) solution method with a given deflation strategy among

SMALLEST, LARGEST and CLUSTER. Consequently values of *ratio* less than 1 indicate which solution methods are expected to be more efficient than FGMRES(m) in terms of computational work. In this table, we see that the SMALLEST deflation strategy yields the best performance with respect to floating-point operations on this application. The CLUSTER deflation strategy tends to favour values of k close to the restart parameter m to be most effective. Figure 2 shows a typical convergence history on this wave propagation problem for two different settings of (m,k).

	FGMRES		SMALLEST			LARGEST			CLUSTER		
m	ℓ	ratio	k	ℓ	ratio	k	ℓ	ratio	k	ℓ	ratio
10	492	1.00	6	**161**	0.55	4	355	0.98	7	175	0.71
12	194	1.00	2	97	0.55	4	181	1.25	9	**76**	0.87
14	148	1.00	5	75	0.66	3	138	1.11	10	**58**	0.80
16	152	1.00	8	61	0.60	6	127	1.16	11	**50**	0.62
18	124	1.00	6	54	0.55	2	103	0.90	16	**45**	1.25
20	101	1.00	7	52	0.65	4	85	0.98	16	**41**	0.90

Table 2: Wave propagation problem ($h = 1/64$, $\sigma = 40$). On each line is shown the iso-memory performance of FGMRES and FGMRES-DR; ℓ is the number of iterations required to satisfy the stopping criterion (8) and ratio the ratio of total floating-point operations v.s. FGMRES(m) (see Equation (10)). Best values of ℓ are marked in bold. Case of a constant two-grid preconditioner.

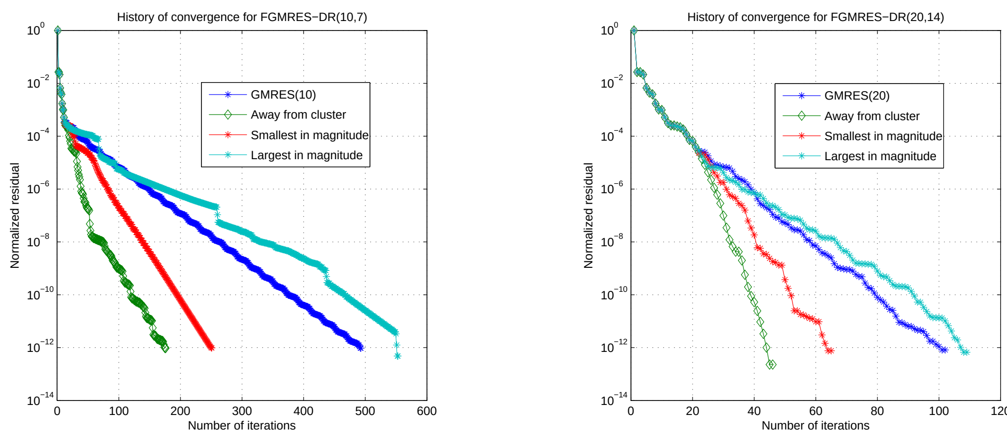

Figure 2: Convergence history of the scaled residual with respect to number of iterations on the wave propagation problem ($h = 1/64$, $\sigma = 40$). Experiments with FGMRES-DR(10,7) (left) and FGMRES-DR(20,14) (right). Case of a constant two-grid preconditioner.

The choice of the two-grid components has led to a fixed preconditioner. This allowed us to compute the spectrum of the preconditioned operator shown in Figure 1. The efficiency of FGMRES with deflated restarting has been shown on this simple model problem. This is of primary interest for three-dimensional wave propagation applications, where the coarse grid systems of the two-grid method can not be handled any more by a sparse direct solution method due to excessive memory requirements. Iterative methods are then required to solve the coarse grid systems only approximately. A non constant preconditioner is then obtained which requires the use of *flexible* Krylov subspace methods. The study of preconditioned FGMRES-DR for such three-dimensional wave propagation applications is beyond the scope of this paper and will be analyzed in the near future. Nevertheless we give an illustration of the potential benefits of FGMRES-DR on the two-dimensional model problem (9) when such inexact coarse grid solution method is used. As an example of approximate coarse grid solver, we consider now the use of an iterative method to solve the coarse grid system to a loose tolerance of 0.15 on the normalized residual. Table 3 reports the results for the two promising deflation strategies SMALLEST and CLUSTER in this setting. The same conclusions as in the constant preconditioner case hold: FGMRES with deflated restarting is efficient. This case study illustrates that there are possibly better choices than selecting the harmonic Ritz vectors corresponding to the harmonic Ritz values of smallest magnitude. If the goal is to minimize the number of matrix-vector products the CLUSTER policy is the most efficient on that problem.

	FGMRES		SMALLEST			CLUSTER		
m	ℓ	ratio	k	ℓ	ratio	k	ℓ	ratio
10	488	1.00	3	**115**	0.28	8	125	0.66
12	208	1.00	2	96	0.50	10	**87**	1.25
14	141	1.00	8	75	0.87	12	**74**	1.78
16	156	1.00	8	67	0.68	13	**54**	0.80
18	124	1.00	6	57	0.60	16	**47**	0.86
20	106	1.00	7	54	0.66	16	**46**	0.90

Table 3: Wave propagation problem ($h = 1/64$, $\sigma = 40$). On each line is shown the iso-memory performance of FGMRES and FGMRES-DR; ℓ is the number of iterations required to satisfy the stopping criterion (8) and ratio the ratio of total floating-point operations v.s. FGMRES(m) (see 10). Best values of ℓ are marked in bold. Case of a non constant two-grid preconditioner.

3.3 Three-dimensional Maxwell's equations in the frequency domain

The boundary element method has become a popular tool in computational electromagnetics for the solution of Maxwell's equations in the frequency domain. The analysis of those wave propagation phenomena has gained an increasing interest in recent years in the simulation of many challenging industrial processes, ranging from the prediction of the Radar Cross Section (RCS) of arbitrarily shaped 3D objects like aircraft, the study of electromagnetic compatibility of electrical devices with their environment, the design of antennae and absorbing materials, and many others. All these simulations are very demanding in terms of computer resources, and require fast and efficient numerical methods to compute an approximate solution of Maxwell's equations. Using the equivalence principle, Maxwell's equations can be recast in the form of integral equations. The discretization is performed on the surface of the object and gives rise to a linear system, where the matrix is dense and complex. Such a linear system can be solved without explicitly forming the matrix A thanks to the fast multipole (FMM) approximation [8, 9, 13, 25]. For perfectly conducting materials effective preconditioners based on sparse approximate inverse can be designed [1, 5]. In this framework, the features of the fast multipole techniques can be further exploited to design an inner-outer scheme. An accurate FMM is used within the outer solver as it governs the final accuracy of the computed solution. The inner solver that acts as a preconditioner consists in a few steps of full GMRES preconditioned by a sparse approximate inverse preconditioner and uses a less accurate FMM.

In this section, we extend the experiments reported in [4] and consider the solution of electromagnetic scattering problems relative to an impedance boundary condition on an obstacle of arbitrary shape in the frequency domain [7]. The geometry of the scattering object is displayed in Figure 3 and corresponds to an air intake. Such a cavity is known to be particularly challenging to solve. The dimension of the linear system is 16 950 for the frequency considered in that example.

In Figure 4 we depict the convergence history for both FGMRES and FGMRES-DR where the inner solver is one restart of GMRES(30) with a sparse approximate inverse preconditioner based on Frobenius norm minimization. The restart parameters of FGMRES and FGMRES-DR are chosen so that both solvers use the same amount of storage. For this implementation the trick based on the LU decomposition of P_k at restarting was not implemented in the prototype code. Based on a previous work [12], where a deflating preconditioning technique targeting the smallest eigenvalues in magnitude was very succesful for perfectly conducting materials, we target the same part of the spectrum for these experiments. The history is plotted at the iteration when the methods start generating different iterates; that is at the smallest restart considered for FGMRES-DR.

It can be seen that FGMRES-DR converges significantly faster than regular FGMRES, especially when the number of deflated directions is increased. As it could be expected if too many directions are deflated the performance deteriorates (see $k = 11$ v.s. $k = 13$ in the graph). The convergence remains worse than full FGMRES but FGMRES-DR is much less memory consuming. The gain becomes larger if accurate solutions are expected. On that large electromagnetics calculation, the extra $\mathcal{O}(k)$ operations are completely negligible and the saving in iteration count directly results in a computational time saving. For instance for a scaled residual lower than 10^{-11} on one processor of a Cray-XD1 computer, the CPU time is 5 hours 47 minutes with FGMRES(25) and only 3 hours 19 min for FGMRES-DR(19,11).

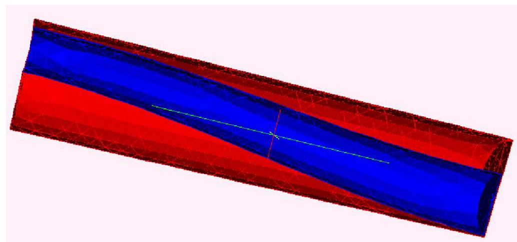

Figure 3: Cross-section of the air intake used in the numerical experiments.

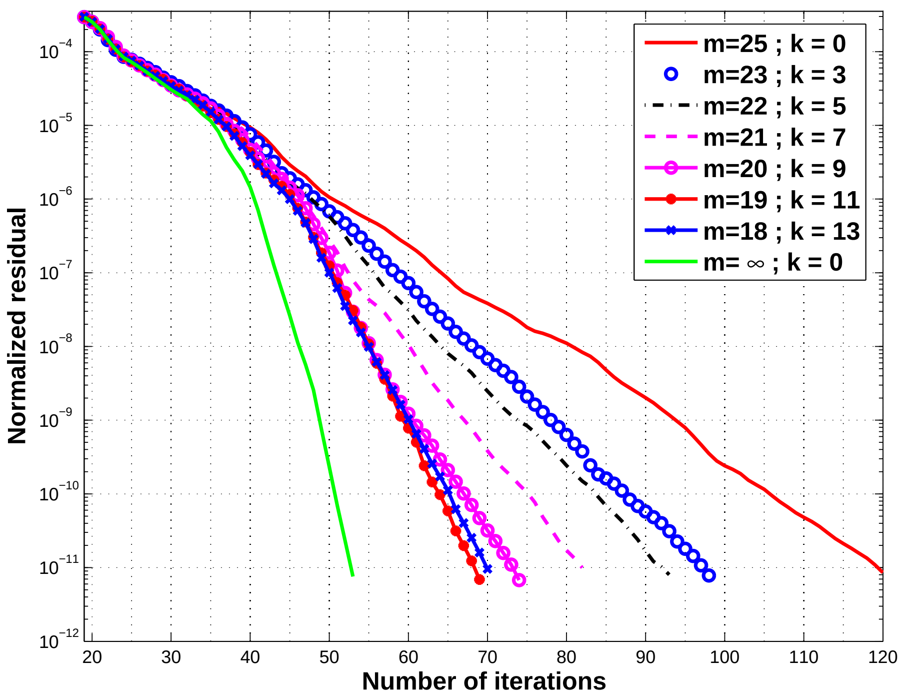

Figure 4: Convergence history of the scaled residual with respect to the iteration for the electromagnetics application.

4 Concluding remarks

There are many situations in scientific computing where variable preconditioners have to be considered for the iterative solution of a linear system. In that framework we have proposed a novel algorithm that attempts to combine the numerical features of GMRES-DR and the flexibility of

FGMRES. The new algorithm, referred to as FGMRES-DR, inherits from the attractive numerical properties of its two parents. We have shown, on a set of small test examples as well as on two real life applications in wave propagation that, after the first restart of the method, FGMRES-DR may outperform FGMRES; the benefit obtained is problem dependent. As for the GMRES-DR algorithm, the eigenvalues of smallest magnitude are often considered as good candidates for the restarting procedure. However, any other part of the spectrum can be considered; the best suited choice is again problem-dependent and can be based on the analysis of the effect of the preconditioner on the system matrix.

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. F. Boisvert, R. Pozo, K. Remington, R. Barrett, and J. J. Dongarra. The Matrix Market: A web resource for test matrix collections. In Ronald F. Boisvert, editor, *Quality of Numerical Software, Assessment and Enhancement*, pages 125–137, London, 1997. Chapman & Hall.

[3] C. Le Calvez and B. Molina. Implicitly restarted and deflated GMRES. *Numerical Algorithms*, 21:261–285, 1999.

[4] 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.

[5] K. Chen. An analysis of sparse approximate inverse preconditioners for boundary integral equations. *SIAM J. Matrix Analysis and Applications*, 22(3):1058–1078, 2001.

[6] G. Cohen. *Higher-order numerical methods for transient wave equations*. Springer-Verlag, 2002.

[7] F. Collino, F. Millot, and S. Pernet. Boundary-integral methods for iterative solution of scattering problems with variable impedance surface condition. *Progress In Electromagnetics Research*, 80:1–28, 2008.

[8] E. Darve. The fast multipole method (I) : Error analysis and asymptotic complexity. *SIAM J. Numerical Analysis*, 38(1):98–128, 2000.

[9] E. Darve. The fast multipole method: Numerical implementation. *J. Comp. Phys.*, 160(1):195–240, 2000.

[10] I. S. Duff, S. Gratton, X. Pinel, and X. Vasseur. Multigrid based preconditioners for the numerical solution of two-dimensional heterogeneous problems in geophysics. *International Journal of Computer Mathematics*, 84-88:1167–1181, 2007.

[11] I. S. Duff, R. G. Grimes, and J. G. Lewis. Users' guide for the Harwell-Boeing sparse matrix collection (Release I). Technical Report TR/PA/92/86, CERFACS, 1992.

[12] L. Giraud, S. Gratton, and E. Martin. Incremental spectral preconditioners for sequences of linear systems. *Applied Numerical Mathematics*, 57(11-12):1164–1180, 2007.

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

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

[15] R. B. Morgan. Implicitly restarted GMRES and Arnoldi methods for nonsymmetric ssystems of equations. *SIAM J. Matrix Analysis and Applications*, 21(4):1112–1135, 2000.

[16] R. B. Morgan. GMRES with deflated restarting. *SIAM J. Scientific Computing*, 24(1):20–37, 2002.

[17] C. C. Paige, M. Rozlozník, and Z. Strakos. Modified Gram-Schmidt (MGS), least squares, and backward stability of MGS-GMRES. *SIAM Journal on Matrix Analysis and Applications*, 28(1):264–284, 2006.

[18] C. C. Paige and Z. Strakos. Residual and backward error bounds in minimum residual Krylov subspace methods. *SIAM Journal on Scientific Computing*, 23(6):1898–1923, 2002.

[19] M. Parks, E. de Sturler, G. Mackey, D.D. Johnson, and S. Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM J. Scientific Computing*, 28(5):1651–1674, 2006.

[20] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Scientific and Statistical Computing*, 14:461–469, 1993.

[21] Y. Saad. Analysis of augmented Krylov subspace methods. *SIAM J. Matrix Analysis and Applications*, 18(2):435–449, 1997.

[22] 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.

[23] G. L. G. Sleijpen and H. A. Van der Vorst. A Jacobi–Davidson iteration method for linear eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 17(2):401–425, 1996.

[24] B. F. Smith, P. E. Bjørstad, and W. D. Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations.* Cambridge University Press, 1996.

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

[26] A. Toselli and O. Widlund. *Domain Decomposition methods - Algorithms and Theory.* Springer Series on Computational Mathematics, vol. 34, Springer., 2004.

[27] U. Trottenberg, C. W. Oosterlee, and A. Schüller. *Multigrid.* Academic Press Inc., 2001.

[28] H. A. van der Vorst and C. Vuik. GMRESR: A family of nested GMRES methods. *Numerical Linear Algebra with Applications*, 1:369–386, 1994.