

Restarted GMRES with Inexact Matrix–Vector Products

Gerard L.G. Sleijpen[†], Jasper van den Eshof[‡], and Martin B. van Gijzen[§]

[†]Department of Mathematics, Utrecht University, P.O. Box 80.010, 3508 TA Utrecht, The Netherlands. sleijpen@math.uu.nl <http://www.math.uu.nl/people/sleijpen/>

[‡]Department of Mathematics, University of Düsseldorf, Universitätsstr. 1, 40224, Düsseldorf, Germany. eshof@am.uni-duesseldorf.de
<http://www.am.uni-duesseldorf.de/~eshof/>

[§]CERFACS, 42, avenue Gaspard Coriolis, 31057 Toulouse cedex 1, France.
gijzen@cerfacs.fr <http://www.cerfacs.fr/algor>

CERFACS Report TR/PA/04/75

Abstract. This paper discusses how to control the accuracy of inexact matrix-vector products in restarted GMRES. We will show that the GMRES iterations can be performed with relatively low accuracy. Furthermore, we will study how to compute the residual at restart and propose suitable strategies to control the accuracy of the matrix-vector products in this computation.

1 Introduction

Iterative Krylov subspace solvers are widely used for solving large systems of linear equations. In recent years Krylov subspace methods have been used more and more for solving linear systems and eigenvalue problems in applications with full coefficient matrix where the matrix-vector products can be approximated reasonably effectively with some (iteration) method. Examples include simulations in quantum chromodynamics [9], electromagnetic applications [5], the solution of generalized eigenvalue problems [6], and of Schur complement systems [3, 8]. Relaxation strategies for controlling the accuracy of inexact matrix-vector products within Krylov methods have attracted considerable attention in the past years, see [1–3, 8, 10, 11]. These strategies allow the error in the matrix-vector product to grow as the Krylov method converges, without affecting the final accuracy too much. Relaxation strategies have been proposed for a range of different Krylov methods, and have shown to be surprisingly effective for the applications mentioned above.

In this paper we will discuss techniques for controlling the error in the matrix-vector product in the context of restarted GMRES. We will argue that the GMRES iterations at the inner level can be performed with relatively low accuracy for the matrix-vector products, since only a residual reduction has to be achieved. Moreover, this low accuracy can be further reduced by applying a relaxation strategy. At restart, a suitable strategy to control the error in the matrix-vector

product for the computation of the residual depends on the method by which the residual is updated. We will discuss two different methods and propose corresponding strategies to control the accuracy of the matrix-vector products in this computation. We will illustrate our ideas with two different types of Schur complement systems.

2 Restarted GMRES with inexact matrix-vector products

2.1 Preliminaries

The central problem is to find a vector \mathbf{x}' that approximately satisfies the equation

$$\mathbf{Ax} = \mathbf{b} \quad \text{such that} \quad \|\mathbf{b} - \mathbf{Ax}'\|_2 < \epsilon, \quad (1)$$

for some user specified, predefined value of ϵ . For ease of presentation we will assume that the problem is scaled such that $\|\mathbf{b}\| = 1$. We will solve the above problem with the restarted GMRES-method but assume that the matrix \mathbf{A} is not explicitly available. Instead, vectors of the form \mathbf{Av} are replaced with an approximation $\mathcal{A}_\eta(\mathbf{v})$ that has a precision η :

$$\mathcal{A}_\eta(\mathbf{v}) = \mathbf{Av} + \mathbf{f} \quad \text{with} \quad \|\mathbf{f}\|_2 \leq \eta \|\mathbf{A}\|_2 \|\mathbf{v}\|_2. \quad (2)$$

2.2 Outline of the algorithm

In restarted GMRES, in every (restart) step a correction to the current approximation is computed by at most m iterations of the GMRES method. Restarting the method is very attractive if inexact matrix-vector products are used, since in the inner loop (the actual GMRES iterations) only a residual reduction has to be achieved that can be much smaller than the final norm of the residual $\|\mathbf{b} - \mathbf{Ax}'\|_2 < \epsilon$ we aim for. As a consequence, lower tolerances can be used in the inner loop for the inexact matrix-vector products [11].

In Figure 1 we have summarized the general structure of the restarted GMRES method. The GMRES iterations are terminated after m steps, or if a residual reduction of ϵ_{inner} has been achieved. In the $j + 1$ -th GMRES iteration of the k -th step of the method we have to prescribe the tolerance $\eta_{j,k}$ for the precision of the matrix-vector product. Since we allow at most a residual reduction of ϵ_{inner} between two restarts, we can apply the relaxation strategy that was proposed in [2] to further reduce the cost of the inexact matrix-vector products. Translated in our setting, this strategy reads

$$\eta_{j,k} = \frac{\|\mathbf{r}_k\|_2}{\|\mathbf{r}_{j,k}\|_2} \epsilon_{inner}. \quad (3)$$

Here $\|\mathbf{r}_{j,k}\|_2$ is the norm of the residual computed in the $j + 1$ -th GMRES iteration of the k -th step. Notice that the limitation on the residual reduction between restarts to ϵ_{inner} is convenient in our context since it allows us to tune the accuracy of the matrix-vector multiplications within the GMRES steps using (3). In the next sections we focus on the computation of the residual at restart.

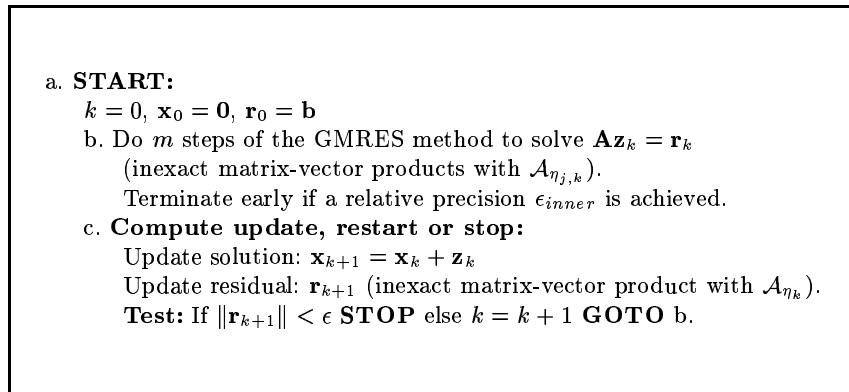


Fig. 1. Restarted GMRES with inexact matrix-vector products

2.3 Restarting: directly computed residuals

At every restart we need to compute the residual corresponding to the newly formed iterate \mathbf{x}_{k+1} . The usual way to compute it is directly from

$$\mathbf{r}_{k+1} = \mathbf{b} - \mathcal{A}_{\eta_k}(\mathbf{x}_{k+1}). \quad (4)$$

The question is how to determine a suitable strategy for choosing the precisions η_k . As was noted in [2] the usual strategy, where the precision is chosen high at the first restart and then decreased at subsequent restarts, does not work in this case. A suitable strategy for choosing η_k can be derived by exploiting that after the GMRES iterations we have for the true residual norm

$$\begin{aligned} \|\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1}\|_2 &= \|\mathbf{b} - \mathbf{A}(\mathbf{x}_k + \mathbf{z}_k)\|_2 \\ &\leq \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2 + \|\mathbf{r}_k - (\mathbf{b} - \mathbf{A}\mathbf{x}_k)\|_2. \end{aligned} \quad (5)$$

The first term in the last expression results from the error that remains after the last GMRES iterations. The second term is the result of the error that we have made in the computation of the residual vector by (4) in the previous restart. Using this expression we find that

$$\|\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1}\|_2 \leq \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2 + \eta_{k-1} \|\mathbf{A}\|_2 \|\mathbf{x}_k\|_2.$$

This bound suggests to choose $\eta_{k-1} = \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2$ for the tolerances. This choice assures us that the second term is of the same order as the first term:

$$\|\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1}\|_2 \leq (1 + \|\mathbf{A}\|_2 \|\mathbf{x}_k\|_2) \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2.$$

As soon as $\|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2$ drops below the precision ϵ we expect to have a solution that has a backward error of at most 2ϵ .

Unfortunately, $\|\mathbf{r}_{k+1} - \mathbf{A}\mathbf{z}_{k+1}\|_2$ is not known in advance, but has to be estimated. Hereto we make two realistic assumptions. The first assumption is

that the residual reduction in iteration $k + 1$ is at most ϵ_{inner} , which implies that

$$\epsilon_{inner} \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2 \leq \|\mathbf{r}_{k+1} - \mathbf{A}\mathbf{z}_{k+1}\|_2.$$

Secondly, we assume that the residual norm as computed by the inexact GMRES process is approximately equal to the true residual norm. This means that at the end of the k -th cycle of GMRES iterations we have that

$$\|\mathbf{r}_{m,k}\|_2 \approx \|\mathbf{r}_k - \mathbf{A}\mathbf{z}_k\|_2.$$

With these assumptions it is easy to see that a suitable choice for η_k is given by $\eta_k = \epsilon_{inner} \|\mathbf{r}_{m,k}\|_2$. The analysis of this choice gives rise to complicated formulas which we do not give here.

Notice that the accuracy with which the residual is computed at restarts is *increased* when the process comes closer to the solution. An advantage of the use of (4) is that the precision ϵ does not have to be decided a priori. Furthermore, there is no accumulation of errors (so the number of restarts does not appear in the expression for the final residual). For this reason, computation of the residuals at restart using (4) can be necessary in some applications. See e.g. [7] where the authors discuss the approximate solution of infinite dimensional systems.

Some problems require the solution of a linear system of equations for each computation of a matrix-vector products. This is the case, for example, for the Schur complement problems that we consider in Section 3. We notice that, if the restarted GMRES method converges, we have that $\mathbf{x}_k \approx \mathbf{x}_{k+1}$. We can exploit this by using the solution of the system that has to be solved for the computation of the matrix-vector product in step k as starting vector for the system that has to be solved in step $k + 1$. If, in addition, the GMRES method is restarted in every iteration, then the resulting method is related to the much used Uzawa iteration method, see [4]. Restarting less frequently can be interpreted as an accelerated Uzawa type method.

2.4 Restarting: recursively updated residuals

As an alternative for the computation of the residual by means of (4), we can compute the residual at restart by exploiting that $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{z}_k$:

$$\mathbf{r}_{k+1} = \mathbf{r}_k - \mathcal{A}_{\eta_k}(\mathbf{z}_k). \quad (6)$$

In this case we find with $\mathbf{r}_0 = \mathbf{b}$ and $\mathbf{x}_{k+1} = \sum_{j=0}^k \mathbf{z}_j$ that

$$\|\mathbf{r}_{k+1} - (\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1})\|_2 \leq \|\mathbf{A}\|_2 \sum_{j=0}^k \eta_j \|\mathbf{z}_j\|_2 \quad (7)$$

Furthermore, with the estimate $\|\mathbf{z}_j\|_2 \leq \|\mathbf{A}^{-1}\|_2 (\|\mathbf{r}_j\|_2 + \|\mathbf{r}_j - \mathbf{A}\mathbf{z}_j\|_2)$, we find

$$\|\mathbf{r}_{k+1} - (\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1})\|_2 \leq 2\|\mathbf{A}\|_2 \|\mathbf{A}^{-1}\|_2 \sum_{j=0}^k \eta_j \|\mathbf{r}_j\|_2. \quad (8)$$

Here, we have assumed that $\|\mathbf{r}_j - \mathbf{A}\mathbf{z}_j\|_2 \leq \|\mathbf{r}_j\|_2$ which can be shown to be true (up to a small factor). Given the relation (5) we want to achieve that the size of $\|\mathbf{r}_{k+1} - (\mathbf{b} - \mathbf{A}\mathbf{x}_{k+1})\|_2$ is of the order of $\|\mathbf{r}_{k+1} - \mathbf{A}\mathbf{z}_{k+1}\|_2$, as in the previous section. Therefore we choose the tolerance η_k equal to $\epsilon/\|\mathbf{r}_k\|_2$. This is the *relaxation strategy* proposed in [2].

We have seen that when the residuals at restart are computed recursively using (6), the precision of the matrix-vector products is *decreased* during the process, as opposed to the previous section. The advantage here is that no a priori knowledge is required about the expected residual reduction between two restarts. A disadvantage is that at termination, when $\|\mathbf{r}_{k+1}\|_2 \leq \epsilon$, the upper bound on the norm of the true residual contains the number of restarts and furthermore depends on the condition number of the matrix \mathbf{A} instead of on the norm of the computed solution. This means that one must be cautious for situations where the norm of the solution is much smaller than the inverse of the smallest singular value of the matrix.

3 Numerical experiments

3.1 Description of the test problem

As a model problem we consider the following set of partial differential equations on the unit square Ω :

$$-\nabla^2\psi - \frac{\partial\psi}{\partial x} - \alpha\nabla^2\zeta = f, \quad \nabla^2\psi + \zeta = 0 \quad \text{in } \Omega \quad (9)$$

plus boundary conditions $\psi = \frac{\partial\psi}{\partial n} = 0$ on Γ , the edge of the domain.

The above model problem is a simplified version of the example from oceanography that we considered in [11]. Discretisation yields the following block system

$$\begin{pmatrix} \mathbf{K} & \alpha\mathbf{L} \\ -\mathbf{L}^H & \mathbf{M} \end{pmatrix} \begin{pmatrix} \psi \\ \zeta \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{0} \end{pmatrix} \quad (10)$$

The size of the test problem we use in our experiments is 16642. From (10) we can eliminate either ζ , which yields the upper Schur complement system

$$(\mathbf{K} + \alpha\mathbf{L}\mathbf{M}^{-1}\mathbf{L}^H)\psi = \mathbf{f}, \quad (11)$$

or ψ , which yields the lower Schur complement system

$$(\mathbf{M} + \alpha\mathbf{L}^H\mathbf{K}^{-1}\mathbf{L})\zeta = \mathbf{L}^H\mathbf{K}^{-1}\mathbf{f}. \quad (12)$$

These two systems have very different numerical characteristics. The upper Schur complement is a fourth order bi-harmonic operator which becomes rapidly ill conditioned if the mesh size is decreased. The lower Schur complement has the characteristics of a second order operator and hence is, for fine enough mesh size, better conditioned than the biharmonic operator. To illustrate this, the MATLAB-routine `condst` gives 500 for the condition number of the lower Schur complement, and a condition number of 10^6 for the upper Schur complement. Systems with \mathbf{M} , on the other hand, are easier to solve (since the mass matrix is a discretised unit operator) than systems with the convection-diffusion operator.

3.2 Solution methods

For each multiplication with one of the Schur complements a linear system has to be solved, with the matrix \mathbf{K} for the upper Schur complement, and with the matrix \mathbf{M} for the lower Schur complement. In our experiments we solve these systems with preconditioned Bi-CGstab, with as preconditioner ILU of \mathbf{K} resp. of \mathbf{M}), using a drop tolerance of 10^{-2} . These systems are solved up to a precision (residual reduction) η . Note that this does not mean that the matrix-vector multiplication with the Schur complements is performed with accuracy η (see also [8]), in theory an extra constant has to be taken into account.

The parameter α is taken rather small in the experiments: $\alpha = 10^{-3}$. For this reason we have applied ILU with drop tolerance 10^{-2} of \mathbf{M} (of \mathbf{K}) as right preconditioner for the lower (resp. upper) Schur complement systems.

We consider four different methods to control the accuracy of the matrix-vector products:

- The systems to evaluate the matrix-vector products with the Schur complement are all solved to fixed accuracy $\epsilon = 10^{-8}$.
- Within GMRES the system are solved with reduced accuracy $\epsilon_{inner} = 10^{-3}$. At restart the accuracy $\epsilon = 10^{-8}$ is used.
- Within GMRES, relaxation is applied by $\eta_{j,k} = 10^{-3} \cdot (\|\mathbf{r}_k\|_2 / \|\mathbf{r}_{j,k}\|_2)$. The residuals at restart are computed directly from \mathbf{x}_{k+1} , using precision $\eta_k = \max(10^{-3} \cdot \|\mathbf{r}_{m,k}\|_2, 10^{-8})$.
- The above relaxation strategy is used within GMRES. The residuals at restart are computed recursively using precision $\eta_k = 10^{-8} / \|\mathbf{r}_k\|_2$.

3.3 The upper Schur complement

Table 1 shows the numerical results for the upper Schur complement. The first column gives the method to control the accuracy. The second column gives the number of Bi-CGstab iterations. Bi-CGstab allows for two tests of the residual norm and hence may terminate halfway an iteration, which explains the fractions of two. The number of Bi-CGstab iterations gives a measure for the work in the inexact matrix-vector products. The third column gives the number of GMRES-iterations. In our experiments we restart (at least) every 20 iterations. The fourth column gives the number of restarts. The true residual norm at the end of the iterative process, which is computed using an exact matrix-vector product, is given in the fifth column. In practice the true residual norm is not available; one has at its disposal only $\|\mathbf{r}_k\|_2$ that is computed using inexact matrix-vector products. This value is given in the sixth column. Note that this value is used in the convergence test. The results tabulated in Table 1 show that the most important saving is obtained by using a lower accuracy for the GMRES iterations. No extra saving is obtained by applying the error-control strategy for directly computed residuals. The savings that are achieved in the initial iterations are lost in the extra few (costly) iterations that are needed due to the extra perturbations that are introduced in the process. The relaxation strategy at restart

Method	Iterations	Iterations	Restarts	$\ \mathbf{b} - \mathbf{A}\mathbf{x}_k\ _2$	$\ \mathbf{r}_k\ _2$
	Bi-CGstab	GMRES			
Full precision	4462.5	1700	85	$3.3 \cdot 10^{-8}$	$8.9 \cdot 10^{-9}$
Low accuracy GMRES	2317.5	2060	103	$3.3 \cdot 10^{-8}$	$8.9 \cdot 10^{-9}$
Directly computed residuals at restart	2345	2120	106	$3.3 \cdot 10^{-8}$	$9.1 \cdot 10^{-9}$
Recursively computed residuals at restart	1903.5	1800	90	$2.1 \cdot 10^{-6}$	$9.0 \cdot 10^{-9}$

Table 1. Solution of the upper Schur system: four different ways to control the accuracy of the matrix-vector products and their effect on the efficiency and on the final accuracy.

with recursively updated residuals, on the other hand, yields a small but significant extra saving in computational cost. Note that on average the tolerances used in the matrix-vector product are lower than for the strategy with directly computed residuals. Moreover, the number of restarts (and GMRES iterations) is less for this example. An important disadvantage of this strategy, however, is that the true residual norm stagnates at around 10^{-6} . This can be explained by the fact the Schur complement is ill conditioned, and relatively many restarts are required to solve the system. Note that both the condition number of the Schur complement and the number of restarts negatively influence the bound (8). If residuals are calculated directly, however, the true residual norm decreases close to the target 10^{-8} .

3.4 The lower Schur complement

The numerical results for the experiments with the lower Schur complement are tabulated in Table 2. The number of GMRES iterations and restarts are

Method	Iterations	Iterations	Restarts	$\ \mathbf{b} - \mathbf{A}\mathbf{x}_k\ _2$	$\ \mathbf{r}_k\ _2$
	Bi-CGstab	GMRES			
Full precision	2955	160	8	$1.0 \cdot 10^{-8}$	$9.9 \cdot 10^{-9}$
Low accuracy GMRES	489.5	180	9	$7.2 \cdot 10^{-9}$	$5.7 \cdot 10^{-9}$
Directly computed residuals at restart	464.5	200	10	$1.2 \cdot 10^{-9}$	$4.7 \cdot 10^{-9}$
Recursively computed residuals at restart	372	200	10	$2.3 \cdot 10^{-8}$	$4.4 \cdot 10^{-9}$

Table 2. Solution of the lower Schur system: four different ways to control the accuracy of the matrix-vector products and their effect on the efficiency and on the final accuracy.

considerably less than for the upper Schur complement system. This is because this system is better conditioned. This fact is also reflected in the accuracy that is achieved using recursively updated residuals plus relaxation: the norm of the true residual stagnates around the target value 10^{-8} . The savings that are obtained for this example are very significant.

4 Conclusions

A considerable saving in computational cost can be obtained for restarted GMRES with inexact matrix-vector products by using a low accuracy for the GMRES iterations, in combination with a high accuracy to compute the residuals at restart. The accuracy of the matrix-vector products at restart can be reduced by a strategy that depends on the way the residual is calculated. If the residual is calculated directly from the latest iterate then the precision of the matrix-vector product has to be increased as the method comes closer to the solution, whereas if the residuals are calculated recursively a relaxation strategy can be applied. This latter strategy, however, has the disadvantage that the achieved accuracy may be well above the target accuracy for ill-conditioned problems that require many restarts.

Acknowledgments. We thank Daniel Loghin for providing us with the test matrices and Françoise Chaitin-Chatelin for valuable comments.

References

1. A. Bouras and V. Frayssé, *A relaxation strategy for Arnoldi method in eigenproblems*, Technical Report TR/PA/00/16, CERFACS, France, 2000.
2. ———, *A relaxation strategy for inexact matrix-vector products for Krylov methods*, Technical Report TR/PA/00/15, CERFACS, France, 2000.
3. A. Bouras, V. Frayssé, and L. Giraud, *A relaxation strategy for inner-outer linear solvers in domain decomposition methods*, Technical Report TR/PA/00/17, CERFACS, France, 2000.
4. James H. Bramble, Joseph E. Pasciak, and Apostol T. Vassilev, *Analysis of the inexact Uzawa algorithm for saddle point problems*, SIAM J. Numer. Anal. **34** (1997), no. 3, 1072–1092.
5. B. Carpentieri, *Sparse preconditioners for dense complex linear systems in electromagnetic applications*, Ph.D. dissertation, INPT, April 2002.
6. F. Chaitin-Chatelin and T. Meškauskas, *Inner-outer iterations for mode solvers in structural mechanics: application to the Code Aster*, Contract Report TR/PA/01/85, CERFACS, Toulouse, France, 2001.
7. P. Favati, G. Lotti, O. Menchi, and F. Romani, *Solution of infinite linear systems by automatic adaptive iterations*, Linear Algebra Appl. **318** (2000), no. 1-3, 209–225.
8. V. Simoncini and D. B. Szyld, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput. **25** (2003), no. 2, 454–477.
9. J. van den Eshof, A. Frommer, Th. Lippert, K. Schilling, and H.A. van der Vorst, *Numerical methods for the QCD overlap operator: I. sign-function and error bounds*, Comput. Phys. Comm. **146** (2002), 203–224.
10. J. van den Eshof and G. L. G. Sleijpen, *Inexact Krylov subspace methods for linear systems*, Preprint 1224, Dep. Math., University Utrecht, Utrecht, the Netherlands, February 2002, To appear in SIAM J. Matrix Anal. Appl. (SIMAX).
11. J. van den Eshof, G. L. G. Sleijpen, and M. B. van Gijzen, *Relaxation strategies for nested Krylov methods*, Preprint 1268, Dep. Math., University Utrecht, Utrecht, the Netherlands, March 2003.