ON THE ROLE OF ORTHOGONALITY IN THE GMRES METHOD

M. ROZLOŽNÍK, Z. STRAKOŠ, M. TŮMA * **

Institute of Computer Science, Academy of Sciences of the Czech Republic
Pod vodárenskou věží 2, 182 07 Praha 8, Czech Republic
e-mail: miro@uivt.cas.cz, strakos@uivt.cas.cz, tuma@uivt.cas.cz

Abstract. In the paper we deal with some computational aspects of the Generalized minimal residual method (GMRES) for solving systems of linear algebraic equations. The key question of the paper is the importance of the orthogonality of computed vectors and its influence on the rate of convergence, numerical stability and accuracy of different implementations of the method. Practical impact on the efficiency in the parallel computer environment is considered.

1 Introduction

Scientific and engineering research is becoming increasingly dependent upon development and implementation of efficient parallel algorithms on modern high-performance computers. Numerical linear algebra is an important part of such research and numerical linear algebra algorithms represent the most widely used computational tools in science and engineering. Matrix computations, including the solution of systems of linear equations, least squares problems, and algebraic eigenvalue problems, govern the performance of many applications on vector and parallel computers.

In recent years there has been a surge of interest for numerical solving of partial differential equations in numerous application fields. The major obstacle arising in solving the corresponding linear algebraic systems is the fact that these problems are usually very large and ill-conditioned. On the other hand, the resulting systems have often some special structure and they are usually sparse. These properties can be naturally exploited in the development of the efficient numerical algorithms for solving these problems.

This paper concentrates on the special important question in the solution of large sparse systems of linear algebraic equations. We consider the Generalized Minimal Residual Method (GMRES) [7], most widely known and used representative of the class of Krylov subspace methods, presented in recent years (for a deep survey we refer to [5]). In particular, we consider the preconditioned GMRES method and its possible implementations based on different orthogonalization techniques. We will examine the level of orthogonality among the

* Part of this work was performed while the second author visited Department of Mathematics and Computer Science, Emory University, Atlanta, USA

** This work was supported by the GA AS CR under grant 230401 and by the NSF grant Int 921824

computed basis vectors produced by these schemes in connection to the decrease of the norm of the GMRES residual. Presented results will lead to important conclusions about the parallel implementation and application of the method.

The organization of this paper is as follows. Section 2 briefly introduces the GMRES method, discusses some of its basic algorithmic aspects and its preconditioning. In Section 3 we give a description of practical application and preconditioning technique used in the numerical experiments presented in this paper. Section 4 concerns three important implementations of the GMRES method based on different orthogonalization schemes, namely the Householder orthogonalization, classical and modified variant of the Gram-Schmidt orthogonalization.

2 The Generalized Minimal Residual Method

Let $Ax = b$ be a system of linear algebraic equations, where A is a real nonsingular N by N matrix and b an N-dimensional real vector. The Generalized minimal residual method (GMRES) for solving this system starts with an initial guess x_0 and seeks the n-th approximate solution x_n of the form

$$x_n \in x_0 + K_n(A, r_0), \tag{2. 1}$$

where $r_0 = b - Ax_0$ is the initial residual and $K_n(A, r_0)$ is the n-th Krylov subspace generated by the matrix A and vector r_0,

$$K_n(A, r_0) = span\{r_0, Ar_0, \dots, A^{n-1}r_0\}. \tag{2. 2}$$

The approximate solution of the GMRES method minimizes the Euclidean norm of the residual, i.e.

$$\|b - Ax_n\| = \min_{u \in x_0 + K_n(A,r_0)} \|b - Au\|. \tag{2. 3}$$

The classical variant of the Generalized Minimal Residual Method consists of constructing the Arnoldi basis of the Krylov subspaces and then solving the transformed least squares problem at each individual iteration step.
The orthonormal basis $V_{n+1} = [v_1, \dots, v_{n+1}]$ of the Krylov subspace $K_n(A, r_0)$, called the Arnoldi basis, is computed via the recurrences

$$v_1 = r_0/\rho,\ \rho = \|r_0\|,\ \ AV_n = V_{n+1}H_{n+1,n},\ \ V_{n+1}^T V_{n+1} = I_{n+1}. \tag{2. 4}$$

The approximate solution x_n is then given by

$$x_n = x_0 + V_n y_n,$$

where y_n is chosen as the solution of the transformed least squares problem

$$\|\varrho e_1 - H_{n+1,n} y_n\| = \min_{y \in R^n} \|\varrho e_1 - H_{n+1,n} y\|. \tag{2. 5}$$

For more detailed description of the algorithm we refer to [7], [4].

In practice, preconditioning is the essential ingredient in the success of projection type methods like GMRES. It consists of transforming the original system in such a way that the resulting system would be easier to solve. For example, when the preconditioner is applied, we will solve the preconditioned linear system

$$AM^{-1}z = b, \text{ where } x = M^{-1}z, \tag{2. 6}$$

where the matrix M^{-1} is some approximation to the matrix A^{-1}. The relation (2. 6) represents so-called *explicit* preconditioning (see, e.g., [1]). Its application in the iteration cycle of GMRES reduces to matrix-vector products. On the contrary, a preconditioner is *implicit* if its application requires a solution of the special linear system within each step of the iterative method. It was realized that straightforward implementation of implicit preconditioning in conjugate gradient-like methods could lead to severe degradation of the performance on supercomputers. In particular, the sparse triangular solves involved in the ILU-type preconditioning were found to be a serial bottleneck (due to the recursive nature of the computation), thus limiting the effectiveness of this approach on vector and parallel computers. Perhaps the most important examples are provided by preconditioners based on an incomplete LU (ILU) decomposition. Here $M = \bar{L}\bar{U}$ where $\bar{L}$ and $\bar{U}$ are sparse triangular matrices which approximate the exact L and U factors of A. Applying the preconditioner requires the solution of two sparse triangular systems (the forward and backward solves). Other notable examples of preconditioners include the ILQ, SSOR and ADI preconditioners (an initial information can be found, e.g., in [1]).

3 Numerical experiment - description

We show the behaviour of the preconditioned GMRES method on a real-world problem. We have chosen a typical matrix which arises in the circuit simulation (in our case we used the matrix from the 32-bit adder design (see [3])). Its structural pattern is shown in Figure 1.

We have chosen explicit preconditioning technique based on a direct approximation of A^{-1}. We used a preconditioner AIBC (see [2]) with drop tolerance 0.6. Structural pattern of the preconditioner is shown in figure 2. It is worth to note how close is the structure of the preconditioner to the structure of the original nonsymmetric matrix although the complete inverse of A should in this case contain much more nonzero elements than shown on Figure 2. Numerical experiments were performed on the SGI workstation with the machine precision $\varepsilon = 2.2.10^{-16}$.

4 Implementations of GMRES with different orthogonalization schemes

Several algorithms for computing the Arnoldi basis exist. They differ in the way in which the orthogonalization in the Arnoldi process is carried out. Utilizing

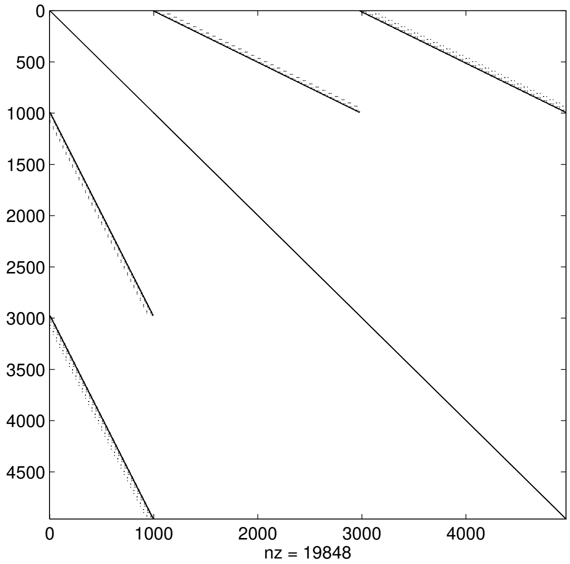

Fig. 1. Structural pattern of the matrix ADD1(4960). Dots (lines) represent positions of nonzero elements in the matrix.

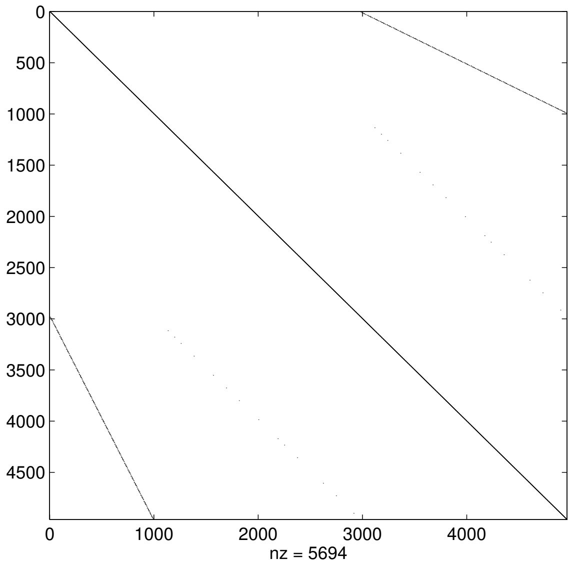

Fig. 2. Structural pattern of the approximate inverse preconditioner for the matrix ADD1(4960). Dots (lines) represent positions of nonzero elements in the matrix.

[7], [8], we will consider classical and modified Gram-Schmidt (MGS) and Householder (HH) orthogonalizations.

In exact arithmetic the Arnoldi vectors are orthogonal. However, in finite precision computation the orthogonality is lost, which may potentially affect both the convergence rate and the ultimate attainable accuracy of the computed approximate solution. In figure 3 we have plotted the loss of orthogonality of computed vectors of different implementations of the GMRES method. Solid line represents the loss of orthogonality $\|I - V_n^T V_n\|_F$ for the Householder implementation, dotted line and dashed lines, the loss of orthogonality for the modified and classical Gram-Schmidt implementations, respectively.

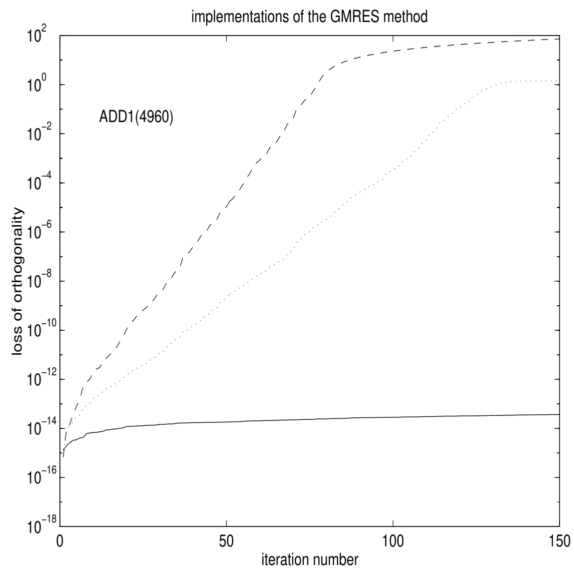

Fig. 3. The loss of orthogonality for different orthogonalization schemes. Solid line represents loss of orthogonality of the Householder orthogonalization. Dotted and dash-dotted lines represent loss of orthogonality of the MGS and CGS orthogonalization, respectively.

If we want to keep the computed orthogonality as close to the machine precision as possible, we need to use the proper orthogonal scheme, namely Householder orthogonalizations. It was shown in [4], that for the Householder implementation the loss of orthogonality of the computed vectors $\bar{V}_{n+1}$ is bounded independently of the system parameters,

$$\|I - \bar{V}_{n+1}^T \bar{V}_{n+1}\| \leq \zeta_1 n^{3/2} N \varepsilon. \tag{4. 7}$$

where ζ_1 is a moderate constant independent of n, N, ε and of other system parameters. Moreover, under obvious assumption on the numerical nonsingularity of the system matrix, the Householder implementation of GMRES is proved backward stable.

In Figure 4 we have plotted the norm of the relative true residual $\|b-A\bar{x}_n\|/\rho$ (solid line), norm of the relative Arnoldi residual, which is given as $\|\,\|r_0\|e_1 - \bar{H}_{n+1,n}\bar{y}_n\|/\rho$ (dashed line). As we can see these two quantities may differ in the finite precision arithmetic. The dash-dotted line represents the smallest singular value of the matrix $\bar{V}_n$ (here it is almost equal to 1 for any value of N) and the dotted line represents the loss of orthogonality of computed vectors.

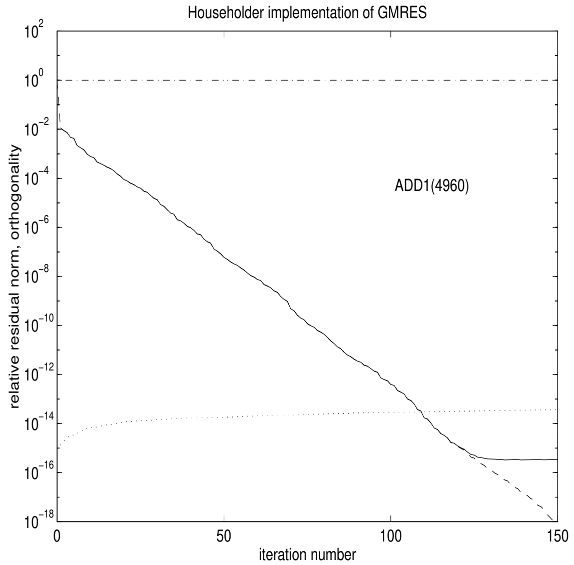

Fig. 4. Householder implementation of GMRES for matrix ADD1(4960). Solid line represents the norm of the relative true residual. Dashed line represents the norm of the relative Arnoldi residual. Dashed-dotted line represents the smallest singular value of the smallest singular value of the matrix $\bar{V}_n$.

The price for preserving of orthogonality near the machine precision is, unfortunately, too high for most of the applications. Moreover, the Householder implementation is less amenable to parallel computer architectures. The Gram-Schmidt process is a cheaper alternative and its modified version represents the most frequently used compromise. In the parallel computational environment, however, the classical Gram-Schmidt would be preferable, but it is not trusted. The loss of orthogonality in the classical Gram-Schmidt is much worse then in the modified version (Figure 3). The loss of orthogonality would be a crucial thing if the computed basis was a primary goal of our computation. But we are not interested in the basis itself – our goal is to solve the linear system. As we will explain, the emphasis on the orthogonality of the computed basis can lead to a completely wrong judgement about the numerical stability and applicability of the GMRES implementations.

In the GMRES context, there is a very important relation between the loss of orthogonality among the Arnoldi vectors and the decrease of the residual of the computed approximation. This relation has been described in [6]. It was shown that, the loss of orthogonality for the modified Gram-Schmidt implementation of GMRES is bounded by

$$\|I - \bar{V}_{n+1}^T \bar{V}_{n+1}\| \leq \zeta_2 n^{3/2} N \frac{\varepsilon \kappa(A)}{\|\hat{r}_n\|/\rho}, \tag{4. 8}$$

where $\kappa(A)$ is the condition number of the matrix A and the quantity $\|\hat{r}_n\|/\rho$ is close to the relative residual norm of the approximate solution in the n-th iteration step, ζ_2 is a moderate constant. Consequently, it can be shown that the Arnoldi vectors will loose their orthogonality completely *only after* the residual of the computed approximation is reduced close to its final level, which is proportional to the machine precision multiplied by the condition number of the system matrix. Until the orthogonality is completely lost, the modified Gram-Schmidt GMRES performs almost exactly as well as the Householder implementation. This suggests that unless the system matrix is extremely ill-conditioned, the use of the modified Gram-Schmidt GMRES is theoretically well justified. These results are illustrated on Figure 5. The proof of (4. 8) is far from trivial and its explanation is beyond the scope of this contribution.

In other words, despite the fact that the Arnoldi basis computed via the modified Gram-Schmidt is much worse than that one computed via the Householder reflections, in the GMRES context it does not really matter. This rises a question about the classical Gram-Schmidt GMRES. For the classical Gram-Schmidt GMRES implementation an analogous statement can be formulated (the final accuracy is, however, different). Although the theoretical analysis has not been finished yet and would need more work [6], we believe that it will lead to reinstating of the classical Gram-Schmidt as a suitable alternative for the parallel implementation of the preconditioned GMRES method.

References

1. R. Barret, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine and H. van der Vorst. Templates for the Solution of Linear Systems. SIAM, Philadelphia, 1994.
2. M. Benzi, M. Tůma: A sparse approximate inverse preconditioner for nonsymmetric linear systems. to appear in SIAM J. Sci. Comput..
3. T. Davis: Sparse matrix collection. *NA Digest*, Volume 94, Issue 42, October 1994.
4. J. Drkošová, A. Greenbaum, M. Rozložník, Z. Strakoš, *Numerical Stability of the GMRES Method*, BIT 3, pp. 309-330, 1995

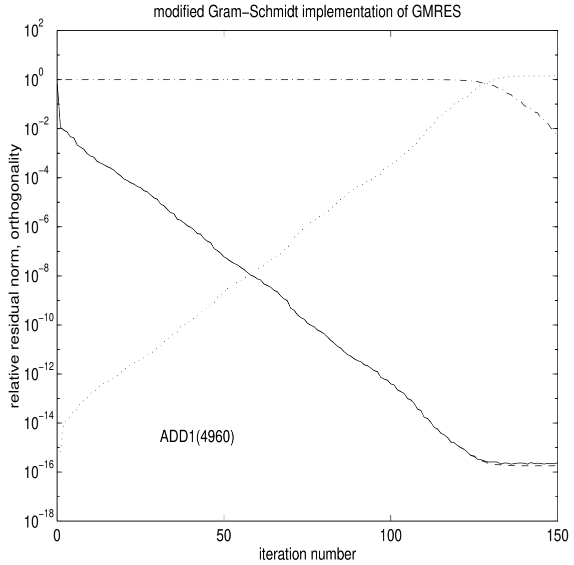

Fig. 5. modified Gram-Schmidt implementation of GMRES. Lines show the same quantities as on Figure 4.

5. R.W. Freund, G.H. Golub, N.M. Nachtigal, *Iterative Solution of Linear Systems*, Acta Numerica 1, pp. 1-44, 1992
6. A. Greenbaum, M. Rozložník, Z. Strakoš, *Numerical Behaviour of the Modified Gram Schmidt GMRES Method*, (submitted to BIT)
7. Y. Saad, M.H. Schultz, *GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*, SIAM J. Sci. Stat. Comput. 7, pp. 856-869, 1986
8. H. F. Walker, *Implementation of the GMRES method using Householder transformations*, SIAM J. Sci. Stat. Comput. 9, 1 (1988), pp. 152-163.

This article was processed using the LaTeX macro package with LLNCS style