On Convergence and Implementation of Minimal Residual Krylov Subspace Methods for Unsymmetric Linear Systems

J. Liesen[†], and M. Rozložník[‡] and Z. Strakoš[§]

[†]Center for Simulation of Advanced Rockets, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA, and Fakultät für Mathematik, Universität Bielefeld, 33501 Bielefeld, Germany (`liesen@mathematik.uni-bielefeld.de`). This work was partly supported by the SFB 343, Fakultät für Mathematik, Universität Bielefeld, Germany.

[‡]Seminar for Applied Mathematics, Swiss Federal Institute of Technology (ETH) Zürich, ETH-Zentrum, CH-8092 Zürich, Switzerland (`miro@sam.math.ethz.ch`). This work was supported by the Grant Agency of the Czech Republic under grants No. 201/98/P108 and No. 101/00/1035. Part of the work was performed while this author visited Emory University in May 1999.

[§]Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA and Institute of Computer Science, Academy of Sciences of the Czech Republic, Prague (`strakos@cs.cas.cz`).

On Convergence and Implementation of Minimal Residual Krylov Subspace Methods for Unsymmetric Linear Systems

J. Liesen[†], and M. Rozložník[‡] and Z. Strakoš[§]

Seminar für Angewandte Mathematik
Eidgenössische Technische Hochschule
CH-8092 Zürich
Switzerland

Abstract

Consider linear algebraic systems $Ax = b$ with a general unsymmetric nonsingular matrix A. We study Krylov subspace methods for solving such systems that minimize the norm of the residual at each step. Such methods are often formulated in terms of a sequence of least squares problems of increasing dimension. Therefore we begin with an overdetermined least squares problem $Bu \approx c$ and present several basic identities and bounds for the least squares residual $r = c - By$. Then we apply these results to minimal residual Krylov subspace methods, and formulate several theoretical consequences about their convergence. We consider possible implementations, in particular various forms of the GMRES method [26], and discuss their numerical properties. Finally, we illustrate our findings by numerical examples and draw conclusions.

Keywords: linear systems, least squares problems, Krylov subspace methods, minimal residual methods, GMRES, convergence, rounding errors

AMS Subject Classification: 65F10, 65F20, 65G50, 15A42

[†]Center for Simulation of Advanced Rockets, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA, and Fakultät für Mathematik, Universität Bielefeld, 33501 Bielefeld, Germany (`liesen@mathematik.uni-bielefeld.de`). This work was partly supported by the SFB 343, Fakultät für Mathematik, Universität Bielefeld, Germany.

[‡]Seminar for Applied Mathematics, Swiss Federal Institute of Technology (ETH) Zürich, ETH-Zentrum, CH-8092 Zürich, Switzerland (`miro@sam.math.ethz.ch`). This work was supported by the Grant Agency of the Czech Republic under grants No. 201/98/P108 and No. 101/00/1035. Part of the work was performed while this author visited Emory University in May 1999.

[§]Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA and Institute of Computer Science, Academy of Sciences of the Czech Republic, Prague (`strakos@cs.cas.cz`).

1 Introduction

Consider a linear algebraic system $Ax = b$, where $A \in \mathcal{R}^{N,N}$ is a nonsingular matrix (generally unsymmetric) and $b \in \mathcal{R}^N$. Krylov subspace methods for solving such systems start with an initial approximation x_0, compute the initial residual $r_0 = b - Ax_0$ and then determine a sequence of approximate solutions $x_1, \ldots, x_n$ such that x_n belongs to the linear manifold determined by x_0 and the n-th Krylov subspace

$$x_n \in x_0 + \mathcal{K}_n(A, r_0), \quad \mathcal{K}_n(A, r_0) \equiv \text{span}\{r_0, Ar_0, \ldots, A^{n-1}r_0\}. \quad (1)$$

The n-th residual then belongs to the manifold given by r_0 and the shifted Krylov subspace (also called Krylov residual subspace)

$$r_n \equiv b - Ax_n \in r_0 + A\mathcal{K}_n(A, r_0). \quad (2)$$

The choice of x_n is based on some particular additional condition. In this paper we focus on the *minimal residual principle*

$$\|r_n\| = \min_{u \in x_0 + \mathcal{K}_n(A, r_0)} \|b - Au\|, \quad (3)$$

which can be equivalently formulated as the *orthogonal projection principle*

$$r_n \perp A\mathcal{K}_n(A, r_0). \quad (4)$$

Since A is assumed to be nonsingular, both (3) and (4) determine the unique sequence of approximate solutions $x_1, \ldots, x_n$, see [26]. Mathematically (in exact arithmetic), there are several different methods and many of their algorithmic variants for generating this sequence. Computationally (in finite precision arithmetic), however, different algorithms for computing the same mathematical sequence may produce different results.

We will call the Krylov subspace methods (1) generating mathematically the sequence $x_1, \ldots, x_n$ uniquely determined by the minimal residual principle (3) (or by the equivalent orthogonal projection principle (4)) *minimal residual Krylov subspace methods* (MR methods).

The minimal residual principle (3) represents a least squares problem, and thus MR methods are often described as a sequence of least squares problems of increasing dimension [26]. Therefore we consider in Section 2 an overdetermined least squares problem $Bu \approx c$, and we present several basic identities and bounds for the residual $r = c - By$. In Section 3 we apply these results to MR methods for the problem $Ax = b$, and formulate several theoretical consequences about their convergence. Section 4 describes the main examples of the MR methods, in particular various forms of the GMRES method [26]. In Section 5 we discuss numerical properties of the important implementations, identify possible numerical weaknesses and illustrate our findings by numerical experiments.

We denote by $\sigma_i(\cdot)$ the i-th largest singular value and by $\sigma_{min}(\cdot)$ the smallest singular value of a given matrix. By $\kappa(\cdot)$ we denote the ratio of the largest to the

smallest singular value (condition number). We use $\|\cdot\|$ to denote the 2-norm, e_i to denote the i-th vector of the standard Euclidean basis, and I to denote the identity matrix.

2   Basic relations for the least squares residual

As the MR methods can be expressed as sequences of least squares problems, it will prove useful to collect some fundamental relations for the least squares residual. We will recall some known results, prove several new results and put the ones known previously in a new context.

Consider an overdetermined linear approximation problem

$$Bu \approx c, \quad B \in \mathcal{R}^{N,n},\ c \in \mathcal{R}^N,\ n < N,\ \mathrm{rank}(B) = n. \tag{5}$$

We denote by y the least squares (LS) solution of (5), and by $r = c - By$ the corresponding LS residual,

$$\|r\| \;=\; \|c - By\| \;=\; \min_u \|c - Bu\|\,. \tag{6}$$

For later convenience we introduce (similarly as in [22]) a real scaling parameter $\gamma > 0$ and consider a scaled version of (5),

$$Bz \approx c\gamma, \quad B \in \mathcal{R}^{N,n},\ c \in \mathcal{R}^N,\ n < N,\ \mathrm{rank}(B) = n. \tag{7}$$

Note that if the right hand side c is replaced in (5) and (6) by the scaled vector $c\gamma$, the LS solution and the LS residual scale trivially to $z = y\gamma$ and $r\gamma$.

We start with a general identity relating r to the matrix $[c\gamma, B]$.

Theorem 2.1. *Suppose that $[c, B] \in \mathcal{R}^{N,n+1}$ has full column rank and $r \neq 0$ is the residual of the least squares problem (5)–(6). Let $\gamma > 0$ be a real parameter. Then*

$$e_1^T\,[c\gamma, B]^\dagger = \frac{r^T}{\gamma\|r\|^2} \qquad and \qquad \gamma\|r\| = \frac{1}{\{e_1^T([c\gamma,B]^T[c\gamma,B])^{-1}e_1\}^{\frac{1}{2}}} \tag{8}$$

where $X^\dagger$ denotes the Moore-Penrose generalized inverse of a matrix X.

Proof. For any matrix X the Moore-Penrose pseudoinverse $X^\dagger$ satisfies $XX^\dagger X = X$ (see, e.g., [5]), which using the symmetry of $XX^\dagger$ gives $X^T = X^T X X^\dagger$. Substituting $X = [c\gamma, B]$, we receive the following simple identities:

$$\begin{aligned}
\gamma r^T &= [1, -\gamma y^T]\,[c\gamma, B]^T = [1, -\gamma y^T]\,[c\gamma, B]^T\,[c\gamma, B][c\gamma, B]^\dagger \\
&= \gamma r^T\,[c\gamma, B][c\gamma, B]^\dagger\,.
\end{aligned}$$

Since r is orthogonal to the columns of B, $\gamma r^T\,[c\gamma, B] = \gamma^2 (r^T c)\,e_1^T = \gamma^2\|r\|^2\,e_1^T$, which proves the first part of the theorem. The second part follows from the identity $\|e_1^T\,[c\gamma, B]^\dagger\|^2 = e_1^T([c\gamma,B]^T[c\gamma,B])^{-1}e_1$, which can be verified by a straightforward calculation. $\qquad\square$

The first equality in (8) was essentially proven (though neither the statement nor the proof were formulated explicitly in the form presented here) in [28, relations (2.5), (2.6), (3.7) and (3.8)]. Later it was presented (with $\gamma = 1$) in [16, Lemma 7.1] (see also other references given there).

It is important to notice that for an arbitrary nonsingular matrix $M \in \mathcal{R}^{n,n}$,

$$\|r\| \;=\; \|c - By\| \;=\; \|c - (BM)(M^{-1}y)\| \;=\; \min_u \|c - (BM)u\| \,.$$

As a consequence of this simple observation, (8) will hold when B is replaced by BM. A particularly useful choice is $M = R^{-1}$, where R is the upper triangular factor of a QR-factorization of B.

Theorem 2.1. *Using the assumptions and the notation of Theorem 2.1, and a QR-factorization $B = QR$ of the matrix B,*

$$e_1^T \, [c\gamma, Q]^\dagger = \frac{r^T}{\gamma \|r\|^2} \qquad and \qquad \gamma\|r\| = \frac{1}{\{e_1^T([c\gamma, Q]^T[c\gamma, Q])^{-1}e_1\}^{\frac{1}{2}}}\,. \qquad (9)$$

It may look a bit surprising that the first rows of the matrices $[c\gamma, B]^\dagger$ and $[c\gamma, Q]^\dagger$ are identical. A second look reveals that this fact is simple and natural:

Consider a full column rank matrix $X = [c\gamma, B] \in \mathcal{R}^{N,n+1}$. Then the rows of $X^\dagger$ are linear combinations of the rows of X^T (the transposed columns of X), and $X^\dagger X = I$. The last relation can be interpreted geometrically as an orthogonal relation between the rows of $X^\dagger$ and the columns of X. Denote by $s = e_1^T X^\dagger$ the first row of $X^\dagger$. Then s is orthogonal to all but the first column of X, i.e., it is orthogonal to the columns of the matrix B. Because s represents a linear combination of c^T and the transposed columns of B, it must be equal to a scalar multiple of the transposed residual $r^T = (c - By)^T$ for the least squares problem (5)–(6). The identity $(\zeta r^T)(c\gamma) = 1$ then immediately gives $\zeta = \gamma^{-1}\|r\|^{-2}$.

The orthogonality idea clearly applies with no change when B is replaced by any matrix BM, where $M \in \mathcal{R}^{n,n}$ is nonsingular. The geometrical interpretation of the generalized inverse is simple, but revealing.

The following theorem relates the norm of the LS residual (6) to the singular values of the matrices B, $[c\gamma, B]$ and $[c\gamma, Q]$.

Theorem 2.2. *Suppose that $[c, B] \in \mathcal{R}^{N,n+1}$ has full column rank and $r \neq 0$ is the residual of the least squares problem (5)–(6). Let $B = QR$ be a QR-factorization of the matrix B and $\gamma > 0$ be a real parameter. Then*

$$\|r\| \;=\; \frac{\sigma_{min}([c\gamma, B])}{\gamma} \prod_{j=1}^n \frac{\sigma_j([c\gamma, B])}{\sigma_j(B)} \qquad (10)$$

$$=\; \frac{1}{\gamma}\,\sigma_{min}([c\gamma, Q])\,\sigma_1([c\gamma, Q]). \qquad (11)$$

Furthermore,

$$\kappa([c\gamma, Q]) = \frac{\alpha + \left(\alpha^2 - 4\gamma^2 \|r\|^2\right)^{1/2}}{2\gamma \|r\|}, \quad \|r\| = \frac{\alpha}{\gamma} \frac{\kappa([c\gamma, Q])}{\kappa([c\gamma, Q])^2 + 1}, \quad (12)$$

where $\alpha \equiv 1 + \gamma^2 \|c\|^2$.

Proof. Using the orthogonality of the columns of the matrix Q, the right hand side c and the residual r are related by the identity

$$c = Qh + r, \quad h \equiv Q^T c, \quad \|c\|^2 = \|h\|^2 + \|r\|^2.$$

Now consider an orthogonal matrix $U \in \mathcal{R}^{n,n}, U^T U = I$, such that $Uh = \|h\|e_1$. Then

$$[c\gamma, Q]^T[c\gamma, Q] = \begin{bmatrix} 1 & 0 \\ 0 & U^T \end{bmatrix} \begin{bmatrix} \gamma^2\|c\|^2 & \gamma\|h\|e_1^T \\ \gamma\|h\|e_1 & I \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & U \end{bmatrix}, \quad (13)$$

$$[c\gamma, B]^T[c\gamma, B] = \begin{bmatrix} 1 & 0 \\ 0 & R^T \end{bmatrix} [c\gamma, Q]^T[c\gamma, Q] \begin{bmatrix} 1 & 0 \\ 0 & R \end{bmatrix}. \quad (14)$$

Identity (13) shows that all but two of the eigenvalues of $[c\gamma, Q]^T[c\gamma, Q]$ are equal to one. The two remaining eigenvalues are easily determined as the eigenvalues of the left principal two-by-two block,

$$\sigma_1^2([c\gamma, Q]) = \frac{\alpha + (\alpha^2 - 4\gamma^2\|r\|^2)^{1/2}}{2}, \quad \sigma_{min}^2([c\gamma, Q]) = \frac{\alpha - (\alpha^2 - 4\gamma^2\|r\|^2)^{1/2}}{2},$$

where $\alpha \equiv 1 + \gamma^2 \|c\|^2$ (notice that $\alpha^2 - 4\gamma^2\|r\|^2 \geq (1 - \gamma^2\|c\|)^2 \geq 0$). Using

$$\kappa([c\gamma, Q]) = \sigma_1([c\gamma, Q])/\sigma_{min}([c\gamma, Q]),$$

(12) is obtained by a simple algebraic manipulation.

Evaluating the determinants on both sides of (13) yields

$$\det([c\gamma, Q]^T[c\gamma, Q]) = \sigma_1^2([c\gamma, Q]) \, \sigma_{min}^2([c\gamma, Q]) = \gamma^2 \|r\|^2,$$

which shows (11). Similarly, transformation (14) yields

$$\begin{aligned} \det([c\gamma, B]^T[c\gamma, B]) &= \prod_{j=1}^{n+1} \sigma_j^2([c\gamma, B]) \\ &= \det([c\gamma, Q]^T[c\gamma, Q]) \det(R^T R) = \gamma^2 \|r\|^2 \prod_{j=1}^{n} \sigma_j^2(B), \end{aligned}$$

which proves (10). $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \square$

The relations (12) generalize results presented in [19, Section 5.5.2]. The identity (10) was (with $\gamma = 1$) first shown by Van Huffel and Vandewalle [29, Theorem 6.9], and it also appeared (with a different proof) in [20].

Van Huffel and Vandewalle [29, Theorem 6.10] gave the following lower and upper bounds for $\|r\|$ (with $\gamma = 1$) in terms of $\sigma_{min}([c\gamma, B])$ (see also [20]). Let

$$\delta(\gamma) \equiv \sigma_{min}([c\gamma, B])/\sigma_{min}(B). \qquad (15)$$

Then

$$\frac{\sigma_{min}([c\gamma, B])}{\gamma} \leq \|r\| \leq \frac{\sigma_{min}([c\gamma, B])}{\gamma} \left\{ 1 - \delta(\gamma)^2 + \frac{\gamma^2 \|c\|^2}{\sigma_{min}^2(B)} \right\}^{\frac{1}{2}}. \qquad (16)$$

Bounds for $\|r\|$ in terms of the minimal singular values of B and $[c\gamma, B]$ and as little additional information as possible were studied in detail in [20]. In particular, when B has full column rank and

$$c \not\perp \{\text{left singular vector subspace of } B \text{ corresponding to } \sigma_{min}(B)\}, \qquad (17)$$

then the following bounds were given in [20]:

$$\sigma_{min}([c\gamma, B]) \left\{ \gamma^{-2} + \|y\|^2 \right\}^{\frac{1}{2}} \leq \|r\|$$

$$\leq \sigma_{min}([c\gamma, B]) \left\{ \gamma^{-2} + \frac{\|y\|^2}{1 - \delta(\gamma)^2} \right\}^{\frac{1}{2}}. \qquad (18)$$

Though (16) can be derived from (18) (and not vice versa, see [20]), the upper bound in (18) is not always tighter than the upper bound in (16). When $\delta(\gamma) \approx 1$ and $\|r\| \approx \|c\|$, it is possible for the upper bound in (16) to be smaller than that in (18). But in this case the upper bound in (16) becomes trivial. For details see [20].

For $\delta(\gamma) = 1$ the upper bound in (18) does not exist. It was shown in [22] that if (17) holds, then $\delta(\gamma) < 1$ for all $\gamma > 0$. As explained in [22], the role of the assumption (17) is truly fundamental. If it does not hold, both theory and computation in errors-in-variables modeling is enormously complicated by the possible case $\delta(\gamma) = 1$. Fortunately nearly all practical problems will satisfy (17). Nevertheless, it is instructive to consider possible cases where (17) does not hold, so that $\delta(\gamma) = 1$ is possible.

The lower bound in (16) shows that we can make $\sigma_{min}([c\gamma, B])$ arbitrarily small by taking γ small, and thus ensure $\delta(\gamma) < 1$ in (15). How small must γ be to ensure this? The next theorem answers a variant of this question. Given $\sigma_{min}(B)$ and $\|c\|$, it shows that there is a γ_0 such that $\gamma < \gamma_0$ ensures $\delta(\gamma) < 1$ but $\gamma = \gamma_0$ does not.

Theorem 2.3. *Suppose that $[c, B] \in \mathcal{R}^{N,n+1}$ has full column rank, y is the solution and $r \neq 0$ is the residual of the least squares problem (5)–(6). Let*

$\gamma > 0$ *be a real parameter and* $\delta(\gamma) \equiv \sigma_{min}([c\gamma, B])/\sigma_{min}(B)$. *Define* $\gamma_0 \equiv \sigma_{min}(B)/\|c\|$. *Then*

$$\delta(\gamma) < 1 \;\; for\;all\;\; \gamma < \gamma_0. \tag{19}$$

Moreover,

$$y = 0 \;\; (r = c) \quad if\;and\;only\;if \quad \delta(\gamma_0) = 1. \tag{20}$$

Proof. Note that when $\gamma < \gamma_0$, then $\|c\gamma\| < \sigma_{min}(B)$. Therefore $\sigma_{min}([c\gamma, B]) < \sigma_{min}(B)$, i.e. $\delta(\gamma) < 1$.

Now assume that the least squares problem (5)–(6) has the trivial solution $y = 0$ $(r = c)$. Then $B^T c = 0$, which yields

$$[c\gamma, B]^T [c\gamma, B] = \begin{bmatrix} \|c\|^2 \gamma^2 & 0 \\ 0 & B^T B \end{bmatrix}.$$

Thus, $\sigma_{min}([c\gamma, B]) = \min\{\|c\|\gamma, \; \sigma_{min}(B)\}$, $\delta(\gamma) = \min\{\|c\|\gamma/\sigma_{min}(B), \; 1\}$, and $\delta(\gamma_0) = 1$. Conversely, (16) gives with $\gamma = \gamma_0$,

$$\delta(\gamma_0)\,\|c\| \;\le\; \|r\| \;\le\; \delta(\gamma_0)\,\|c\|\,\{2 - \delta(\gamma_0)^2\}^{1/2}, \tag{21}$$

which for $\delta(\gamma_0) = 1$ reduces to $\|c\| \le \|r\| \le \|c\|$, i.e. $\|r\| = \|c\|$, which completes the proof. □

We see that $\gamma_0 \equiv \sigma_{min}(B)/\|c\|$ represents an important number. For $\gamma < \gamma_0$ the value of $\delta(\gamma)$ is always strictly less than one, and $\delta(\gamma_0) = 1$ if and only if the least squares problem (5)–(6) has the trivial solution $y = 0$. Moreover, (21) shows that $\|r\|$ is significantly smaller than $\|c\|$ if and only if $\delta(\gamma_0)$ is significantly smaller than one. As an application, we will show in Section 3 how these results characterize stagnation or near-stagnation of MR methods.

One consequence of Theorem 2.3 can be stated as follows: Consider a rectangular matrix (here B) having full column rank, and an additional column (here $c\gamma$). If the norm of the additional column is smaller than the smallest singular value of the matrix (here if $\gamma < \gamma_0$), then appending the column necessarily decreases the smallest singular value. If the norm of the appended column is equal to the smallest singular value of the matrix (here if $\gamma = \gamma_0$), then appending the column to the matrix does not change the smallest singular value *if and only if* the appended column is orthogonal to all the columns of the original matrix. This is a somewhat specialized result because of the fixed norm of the added column. The general condition under which adding a column to a matrix does not alter the smallest singular value was given in [22]; Theorem 2.3 can also be derived from this.

Theorem 2.3 and the consequence stated above must be understood in their proper context. It was pointed out in [22], that nearly all practical problems will satisfy (17), that any problem $Bu \approx c$ can be reduced to a core problem satisfying (17), and that for many formulations it only makes sense to consider problems satisfying (17). Also if the problem satisfies (17) then $\delta(\gamma) < 1$ for all

$\gamma > 0$, and in this case (19) and (20) are irrelevant. On the other hand, (21) seems to be a generally useful result. Thus $\gamma_0 \equiv \sigma_{min}(B)/\|c\|$ is a significant quantity, as can be seen from the interesting but rarely practical properties (19) and (20), and the interesting and compact bounds (21). Note also that in many practical problems of interest, $\gamma_0 \equiv \sigma_{min}(B)/\|c\|$ will be a *very* small number. In particular, this suggests that for a general least squares problem the above 'column addition' result will be of minor practical use. It is, however, important theoretically because it offers a new insight into the stagnation or near stagnation of the MR methods.

Finally, consider (similarly as in Corollary 2.1) a QR-factorization $B = QR$. Replacing B by $BR^{-1} = Q$ and y by Ry (notice that $\|Ry\| = \|By\|$) gives the analogies of (16) and (18),

$$\frac{\sigma_{min}([c\gamma, Q])}{\gamma} \leq \|r\| \leq \frac{\sigma_{min}([c\gamma, Q])}{\gamma} \{1 - \sigma_{min}^2([c\gamma, Q]) + \gamma^2\|c\|^2\}^{\frac{1}{2}}, \quad (22)$$

$$
\begin{aligned}
\sigma_{min}([c\gamma, Q]) \{\gamma^{-2} + \|By\|^2\}^{\frac{1}{2}} &\leq \|r\| \\
&\leq \sigma_{min}([c\gamma, Q]) \{\gamma^{-2} + \frac{\|By\|^2}{1 - \sigma_{min}([c\gamma, Q])^2}\}^{\frac{1}{2}} \quad (23)
\end{aligned}
$$

Theorem 2.3 can be reformulated in a similar way. It is interesting to note that the bounds (22) do not give additional information. Indeed, since $\sigma_1([c\gamma, Q]) \geq 1$, the lower bound in (22) follows immediately from (11). And since $\{1 - \sigma_{min}^2([c\gamma, Q]) + \gamma^2\|c\|^2\}^{1/2} = \sigma_1([c\gamma, Q])$, the upper bound is a weak reformulation of (11) only.

In the following section we apply these results to the minimal residual Krylov subspace methods.

3 Characteristics of the basis and the size of the MR residual

Let $\rho_0 \equiv \|r_0\|$, $v_1 \equiv r_0/\rho_0$, $w_1 \equiv Av_1/\|Av_1\|$. Consider two sequences of orthonormal vectors, $v_1, v_2, \ldots$ and $w_1, w_2, \ldots$, such that for each iterative step n,

$$\mathcal{K}_n(A, r_0) = \text{span}\{v_1, \ldots, v_n\}, \quad V_n \equiv [v_1, \ldots, v_n], \quad V_n^T V_n = I, \quad (24)$$

$$A\mathcal{K}_n(A, r_0) = \text{span}\{w_1, \ldots, w_n\}, \quad W_n \equiv [w_1, \ldots, w_n], \quad W_n^T W_n = I. \quad (25)$$

Then the minimal residual principle (3) can be formulated as

$$\|r_n\| = \min_{u \in \mathcal{R}^n} \|r_0 - AV_n u\| \quad (26)$$

$$= \min_{u \in \mathcal{R}^n} \|r_0 - W_n u\|. \quad (27)$$

The MR residual at step n is therefore the LS residual for the LS problems $AV_n u \approx v_1 \rho_0$ and $W_n u \approx v_1 \rho_0$.

The application of the results presented in Section 2 to (26) and (27) is straightforward: For the n-th step of an MR method we consider $c \equiv r_0 = v_1 \rho_0$, $B \equiv AV_n$, $Q \equiv W_n$ and $r \equiv r_n$. The scaling parameter $\gamma > 0$ offers some flexibility. While it seems natural to use $\gamma \equiv \|r_0\|^{-1} = \rho_0^{-1}$, other values of γ also prove useful, cf. [21] and our discussion below.

With $\gamma \equiv \|r_0\|^{-1} = \rho_0^{-1}$, Theorem 2.2, relations (11) and (12), give the following identities for the relative residual norm $\|r_n\|/\rho_0$:

$$\begin{aligned} \|r_n\|/\rho_0 \quad &= \quad \sigma_{min}([v_1, W_n])\, \sigma_1([v_1, W_n]) & (28) \\ &= \quad \frac{2\,\kappa([v_1, W_n])}{\kappa([v_1, W_n])^2 + 1}. & (29) \end{aligned}$$

Identities (28) and (29) show that the conditioning of the basis $[v_1, W_n]$ of the Krylov subspace $\mathcal{K}_{n+1}(A, r_0)$ is fully determined (except for an unimportant multiplicative factor) by the convergence of the MR methods and vice versa. In other words,

$$\|r_n\| = \rho_0 \quad \text{if and only if} \quad \kappa([v_1, W_n]) = 1, \qquad (30)$$

and the relative residual norm $\|r_n\|/\rho_0$ is small if and only if $[v_1, W_n]$ is ill-conditioned.

The previous statement can also be mathematically expressed by inequalities. Dividing both the numerator and the denominator in (29) by $\kappa([v_1, W_n])$ gives in a simple way the bounds

$$\kappa([v_1, W_n])^{-1} \;\leq\; \|r_n\|/\rho_0 \;\leq\; 2\,\kappa([v_1, W_n])^{-1}. \qquad (31)$$

The upper bound in (31) was published by Walker and Zhou [34, Lemma 3.1].

Relations between the size of the residuals of the MR methods and the condition number of matrices $[v_1, W_n]$ and $[r_0, W_n]$ were studied in [19, Section 5.5.2]. We will generalize the result [19, (5.48)] and develop an elegant tool for quantification of the influence of the scaling parameter γ.

Theorem 3.1. *Let r_0, r_n and W_n be as in (27), $\rho_0 \equiv \|r_0\|, v_1 \equiv r_0/\rho_0$ and $\gamma > 0$. Then*

$$\kappa([r_0\gamma, W_n]) \;\geq\; \kappa([v_1, W_n]) + \frac{\gamma\,(\rho_0 - \gamma^{-1})^2}{2\|r_n\|}. \qquad (32)$$

Proof. Using (12) with the particular choices $c \equiv r_0$, $Q \equiv W_n$, $\gamma > 0$, and

9

$c \equiv r_0$, $Q \equiv W_n$, $\gamma \equiv \rho_0^{-1} = \|r_0\|^{-1}$ gives

$$\kappa([r_0\gamma, W_n]) \quad - \quad \kappa([v_1, W_n])$$

$$= \frac{1 + \gamma^2\rho_0^2 + [(1 + \gamma^2\rho_0^2)^2 - 4\gamma^2\|r_n\|^2]^{1/2}}{2\gamma\|r_n\|} - \frac{2 + [4 - 4\rho_0^{-2}\|r_n\|^2]^{1/2}}{2\rho_0^{-1}\|r_n\|}$$

$$= \frac{\gamma^{-1} + \gamma\rho_0^2 + [(\gamma^{-1} + \gamma\rho_0^2)^2 - 4\|r_n\|^2]^{1/2} - 2\rho_0 - [4\rho_0^2 - 4\|r_n\|^2]^{1/2}}{2\|r_n\|}$$

$$= \frac{\gamma(\rho_0 - \gamma^{-1})^2}{2\|r_n\|} + \frac{[(\gamma^{-1} + \gamma\rho_0^2)^2 - 4\|r_n\|^2]^{1/2} - [4\rho_0^2 - 4\|r_n\|^2]^{1/2}}{2\|r_n\|}$$

$$\geq \frac{\gamma(\rho_0 - \gamma^{-1})^2}{2\|r_n\|}.$$

$\square$

Clearly, $\kappa([r_0\gamma, W_n])$ is minimal for $\gamma = \rho_0^{-1}$, and the minimum is equal to $\kappa([v_1, W_n])$ (see also [8]) . If $\gamma \neq \rho_0^{-1}$, then with the residual norm $\|r_n\|$ decreasing towards zero the condition number $\kappa([r_0\gamma, W_n])$ grows much faster than $\kappa([v_1, W_n])$. The results considering the matrix $[r_0\gamma, W_n]$ will be particularly useful for our discussion of MR implementations based on the orthogonal projection principle (4) in Section 5.

With $c \equiv r_0$, $r \equiv r_n$, $y \equiv y_n$, and $B \equiv AV_n$, (18) gives the following bounds for the residual norm in terms of $\sigma_{min}([r_0\gamma, AV_n])$:

$$\sigma_{min}([r_0\gamma, AV_n])\{\gamma^{-2} + \|y_n\|^2\}^{\frac{1}{2}} \leq \|r_n\|$$
$$\leq \sigma_{min}([r_0\gamma, AV_n])\left\{\gamma^{-2} + \frac{\|y_n\|^2}{1 - \delta_n(\gamma)^2}\right\}^{\frac{1}{2}} \tag{33}$$

where $\delta_n(\gamma) \equiv \sigma_{min}([r_0\gamma, AV_n])/\sigma_{min}(AV_n)$. As mentioned in Section 2, the upper bound in (33) becomes for $\delta_n(\gamma) \approx 1$ intriguing, and for $\delta_n(\gamma) = 1$ it is not defined.

The convergence of the MR methods and the situation $\delta_n(\gamma) = 1$ or $\delta_n(\gamma) \approx 1$ are related by Theorem 2.3. Define $\gamma_0^{(n)} \equiv \sigma_{min}(AV_n)/\rho_0$. Then $\delta_n(\gamma) < 1$ for all $\gamma < \gamma_0^{(n)}$ and

$$\|r_n\| = \rho_0 \quad \text{if and only if} \quad \delta_n(\gamma_0^{(n)}) = 1. \tag{34}$$

Moreover, (21) gives

$$\delta_n(\gamma_0^{(n)}) \leq \|r_n\|/\rho_0 \leq \sqrt{2}\,\delta_n(\gamma_0^{(n)}), \tag{35}$$

which shows that the rate of convergence of the MR methods is determined by the size of $\delta_n(\gamma_0^{(n)})$. Summarizing, the MR methods stagnate in steps 1 through n if and only if $\delta_n(\gamma_0^{(n)}) = 1$, and they nearly stagnate in steps 1 through n

if and only if $\delta_n(\gamma_0^{(n)}) \approx 1$. However, this specific link between convergence of the MR methods and the value of $\delta_n(\gamma)$ can be made for $\gamma = \gamma_0^{(n)}$ only. In particular, when $\delta_n(\gamma_1) = 1$ for some $\gamma_1 > \gamma_0^{(n)}$, the MR methods do not necessarily stagnate or nearly stagnate. They may exhibit very fast convergence while $\delta_n(\gamma_1) \approx 1$ and very slow convergence while $\delta_n(\gamma_1) \ll 1$ (for more details see [21]).

Using the matrix $[r_0\gamma, AV_n]$ instead of $[r_0\gamma, W_n]$ may seem unwise because it necessarily brings into play the potentially ill-conditioned matrix AV_n (in comparison to W_n having orthonormal columns). However, as shown in [22, 21], bounds using the matrix $[r_0\gamma, AV_n]$ are very useful for the analysis of the modified Gram-Schmidt implementation of the GMRES method. Notice that the bounds (33) are not based on singular values only. Using $\|y_n\|$, the norm of the MR approximate solution, makes (33) amazingly tight [22]. The parameter γ offers flexibility required for the analysis of the GMRES method [21].

It is also possible to consider other bases of the Krylov subspaces or Krylov residual subspaces which lead to other matrices, identities and bounds. For example, Ipsen [16, 17] used the matrix $K_{n+1} = [r_0, Ar_0, \ldots, A^n r_0]$, got the identity

$$\|r_n\| = 1/\|e_1^T K_{n+1}^\dagger\| \tag{36}$$

(cf. Theorem 2.1) and developed the bound $\|r_n\|/\rho_0 \geq 1/(\|K_{n+1}\| \, \|K_{n+1}^\dagger\|)$. However, any bound based directly on the matrix K_{n+1} necessarily suffers from the potential ill-conditioning of the matrix $[Ar_0, \ldots, A^n r_0]$. Consider the QR-decomposition $[Ar_0, \ldots, A^n r_0] = W_n R_n$. In light of the results presented above, see, in particular, (9), (28) and (29), the upper triangular factor R_n containing *all* the potential ill-conditioning of the matrix $[Ar_0, \ldots, A^n r_0]$ is mathematically in no relation whatsoever to the residual r_n and to the convergence of any MR method measured by the residual norm. Except for some (rather special) examples, bounds based on the matrix K_{n+1} are therefore necessarily much weaker than the bounds based on the matrices $[r_0\gamma, W_n]$ and $[r_0\gamma, AV_n]$.

4   Implementations of the MR methods

Numerous residual norm minimizing Krylov subspace methods have been proposed in the last decades [18, 30, 35, 1, 11, 26]. Resulting from different approaches, they generate (under different assumptions) approximate solutions satisfying (3) and (4). Though they are, under some particular assumptions, mathematically equivalent, they differ in various algorithmic aspects, and, consequently, in their numerical behavior.

We will concentrate on two main approaches which explicitly compute the basis vectors $v_1, v_2, \ldots, v_n$ respectively $v_1, w_1, \ldots, w_{n-1}$ defined in (24) and (25). In the first approach, the approximate solution x_n is expressed as

$$x_n \;=\; x_0 \;+\; V_n \, y_n$$

for some y_n, and the residual norm is bounded in terms of $\sigma_{min}([v_1\rho_0\gamma, AV_n])$ via (33). In the second approach the approximate solution is expressed as

$$x_n \ = \ x_0 \ + \ [v_1, W_{n-1}]\, t_n$$

for some t_n, and for the residual norm we have the identities (28)–(29). At first sight the second approach seems more attractive because it gives a cleaner relation between the residual norm (which is minimized at every step) and the conditioning of the computed basis. Its implementation is also simpler. On the other hand, the fact that the approximate solution is in this approach determined via the basis vectors $v_1, w_1, \ldots, w_{n-1}$ which *are not mutually orthogonal* raises some suspicions about potential numerical problems. In this section we recall implementations of both approaches resulting in different variants of the GMRES algorithm. In Section 5 we will discuss their numerical properties.

A variety of MR methods that do not explicitly compute the vectors $v_1, v_2, \ldots, v_n$ or $v_1, w_1, \ldots, w_{n-1}$ have been proposed. For example, the method by Khabaza [18] uses the vectors $r_0, Ar_0, \ldots, A^{n-1}r_0$; Orthomin [30], Orthodir [35], Generalized Conjugate Gradient (GCG) [1, 2] and Generalized Conjugate Residual (GCR) [10, 11] compute an $A^T A$-orthogonal basis of $\mathcal{K}_n(A, r_0)$. These methods played an important role in the development of the field. In comparison to the approaches discussed in this paper they are, however, less numerically stable. Therefore we will not consider them below.

4.1 Minimal residual principle: Classical GMRES

Consider an initial approximation x_0 and the initial residual $r_0 = b - Ax_0$, $\rho_0 \equiv \|r_0\|$. In their classical paper [26], Saad and Schultz used the orthonormal basis (24) (Arnoldi basis). As noted in [33], this basis can be mathematically expressed as the Q-factor of a recursive columnwise QR-factorization

$$[r_0, AV_n] = V_{n+1}\, [e_1\rho_0, H_{n+1,n}], \ V_{n+1} \equiv [v_1, \ldots, v_{n+1}], \ V_{n+1}^T V_{n+1} = I_{n+1}. \tag{37}$$

Here $H_{n+1,n}$ is an $(n+1)$-by-n upper Hessenberg matrix with elements $h_{i,j}$, $h_{j+1,j} \neq 0, j = 1, 2, \ldots, n-1$. If at any stage $h_{n+1,n} = 0$, the algorithm would stop with $[r_0, AV_n] = V_n\, [e_1\rho_0, H_{n,n}]$. Using the substitution

$$x_n \ = \ x_0 \ + \ V_n\, y_n \tag{38}$$

and (37), the minimal residual principle (3) gives the least squares problem for the vector of coefficients y_n:

$$\|r_n\| \equiv \|b - Ax_n\| \ = \ \min_{y \in \mathcal{R}^n} \|r_0 - AV_n\, y\| \ = \ \min_{y \in \mathcal{R}^n} \|V_{n+1}\, (e_1\rho_0 - H_{n+1,n}\, y)\| \tag{39}$$

$$= \ \min_{y \in \mathcal{R}^n} \|e_1\rho_0 - H_{n+1,n}\, y\|. \tag{40}$$

To solve (39) we apply orthogonal rotations $J_1, J_2, \ldots, J_n$ sequentially to $H_{n+1,n}$ to bring it to the upper triangular form T_n:

$$J_n \cdots J_2 J_1\, H_{n+1,n} \ = \ \left[\begin{array}{c} T_n \\ 0 \end{array} \right].$$

The vectors y_n and $\|r_n\|$ then satisfy

$$\begin{bmatrix} T_n\,y_n \\ \|r_n\| \end{bmatrix} \;=\; J_1^T J_2^T \cdots J_n^T \, e_1 \rho_0. \tag{41}$$

The value of the (nonincreasing) residual norm is available without determining y_n and it can be easily updated at each iteration, while y_{n+1} and x_{n+1} will usually differ in every element from y_n and x_n respectively. We refer to this algorithm as GMRES or *classical* GMRES.

Several variants for computing the basis vectors $v_1, \ldots, v_n$ were proposed. Saad and Schultz [26] considered the modified Gram-Schmidt process. Walker [32, 33] presented classical GMRES based on Householder transformations. Iterated classical and iterated modified Gram-Schmidt versions were studied in [9].

A variety of parallel implementations [6, 3, 12, 23, 7, 27] use various techniques to increase the parallel efficiency of the basically sequential basis-generating process. Parallel aspects are out of the scope of this paper.

4.2 Orthogonal projection principle: Simpler GMRES

We now consider an implementation of GMRES derived from the orthogonal projection principle (4). The approach proposed by Walker and Zhou [34], called *simpler* GMRES, uses the orthonormal basis (25). This basis is computed by a recursive columnwise QR-factorization of the matrix $[Ar_0\gamma, AW_{n-1}]$. Based on Theorem 3.1 we set $\gamma = \rho_0^{-1}$, and we will use this value of the scaling parameter γ throughout the rest of this paper. Then

$$A[v_1, W_{n-1}] = [Av_1, AW_{n-1}] = W_n S_n, \;\; W_n \equiv [w_1, \ldots, w_n], \;\; W_n^T W_n = I_n, \tag{42}$$

where S_n is an n-by-n upper triangular matrix with elements $s_{i,j}$, $s_{j,j} \neq 0$. If at any stage $s_{n,n} = 0$, the algorithm would stop with $[Av_1, AW_{n-1}] = W_{n-1}[S_{n-1}, \hat{s}_n]$. Using the substitution

$$x_n \;=\; x_0 + [v_1, W_{n-1}]\, t_n\,, \tag{43}$$

the vector $t_n \in \mathcal{R}^n$ solves the least squares problem

$$\|r_n\| \equiv \|b - Ax_n\| \;=\; \min_{t \in \mathcal{R}^n} \|r_0 - A[v_1, W_{n-1}]\,t\| \tag{44}$$

$$=\; \min_{t \in \mathcal{R}^n} \|r_0 - W_n S_n\, t\|\,. \tag{45}$$

Solving the least squares problem (44)–(45) in a numerically stable way represents a more subtle task then solving (39)–(40). The main difference is in handling the right hand side vector r_0. In (39)–(40), r_0 is expressed in terms of the vectors $v_1, v_2, \ldots, v_n$ simply as $r_0 = v_1 \rho_0$. In finite precision arithmetic, until the linear independence of the vectors $v_1, v_2, \ldots, v_n$ is lost, this expression is maximally accurate. On the other hand, application of the orthogonal projection principle (4) directly to (44)–(45) gives the upper triangular system

$$S_n\, t_n \;=\; W_n^T r_0. \tag{46}$$

As demonstrated in [25], computing the vector of coefficients t_n from (46) leads to numerical difficulties. Numerically more stable implementations are described next.

Consider first the implementation of simpler GMRES using the modified Gram-Schmidt process for generating the basis vectors $w_1, \ldots, w_n$. A properly implemented algorithm for solving the least squares problem (44)–(45) applies the orthogonalization process also to the right hand side r_0, see [4, pp. 64-65]. Then, using the recursive columnwise modified Gram-Schmidt QR-factorization of the extended matrix $[Av_1, AW_{n-1}, r_0]$,

$$[Av_1, AW_{n-1}, r_0] \;=\; W_n\,[S_n, z_n] \,+\, [0, \frac{r_n}{\|r_n\|}] \left[\begin{array}{c} 0 \\ \|r_n\| \end{array} \right], \qquad (47)$$

the vector t_n solves the upper triangular system

$$S_n\, t_n \;=\; z_n. \qquad (48)$$

The j-th component of $z_n \equiv (\zeta_1, \ldots, \zeta_n)^T$ is determined by

$$\zeta_j \;=\; w_j^T\,(I - w_{j-1}w_{j-1}^T)\cdots(I - w_1 w_1^T)\,r_0 \;=\; w_j^T\, r_{j-1}\,, \qquad (49)$$

where we use

$$r_j \;=\; (I - w_j w_j^T)\cdots(I - w_1 w_1^T)\,r_0 \;=\; r_{j-1} \,-\, (w_j^T\, r_{j-1})\, w_{j-1}\,. \qquad (50)$$

A complete algorithm of the modified Gram–Schmidt implementation of simpler GMRES is given in the Appendix.

Now we consider the implementation of simpler GMRES based on Householder reflections. It transforms the matrix $[Av_1, AW_{n-1}]$ into upper triangular form,

$$P_n \cdots P_2 P_1\, [Av_1, AW_{n-1}] \;=\; \left[\begin{array}{c} S_n \\ 0 \end{array} \right], \qquad (51)$$

where P_j, $j = 1, \ldots, n$, are elementary Householder matrices (for details see [9, p. 312]). Then the transformed right hand side is determined as

$$z_n \;=\; [P_n \cdots P_1\, r_0]_{\{1:n\}}$$

where $[\cdot]_{\{1:n\}}$ denotes the first n elements of a vector. The vector of coefficients t_n is determined from (48). A complete algorithm of the Householder implementation of simpler GMRES is given in the Appendix.

Related to simpler GMRES are stabilized Orthodir [31], and the recent $A^T A$-variant of GMRES [25]. Both compute an $A^T A$-orthogonal basis of $\mathcal{K}_n(A, r_0)$, and thus each step of these methods requires about twice as much storage and also slightly more arithmetic operations than simpler GMRES. They are also numerically less stable than simpler GMRES. On the other hand, they allow easier parallel implementations, because they feature step by step updates of both the approximate solution and the residual vectors.

5 Numerical stability

In this section we analyze and compare the numerical stability of classical and simpler GMRES. As mentioned in Section 4, different orthogonalization techniques for computing the columns of V_n or W_n can be applied. Here we focus on implementations based on Householder transformations [32, 33] and on the modified Gram-Schmidt process [26].

For distinction, we denote quantities computed in finite precision arithmetic (with the machine precision ε) by bars. We assume the standard model of floating point arithmetic, see, e.g. [15, (2.4)]. In our bounds we present only those terms which are linear in ε and do not account for the terms proportional to higher powers of ε. By $p_k(n, m, N)$, $k = 1, 2, \ldots$, we denote low degree polynomials in the number of iteration steps n, the maximum number of nonzeros per row in the system matrix m and the dimension of the system N. They are introduced in a number of places in the text; some of them depend only on one or two variables. In all cases $p_k(n, m, N) \leq c_k N^{5/2}$, where $c_k > 0$ is a constant independent of n, m and N. This bound is, in general, very pessimistic; it accounts for the worst possible case. For details see [9, 14, 24].

5.1 Classical GMRES

In the classical GMRES implementation the computed approximate solution $\bar{x}_n$ satisfies

$$
\begin{aligned}
\bar{x}_n &= x_0 + \bar{V}_n \bar{y}_n + g_n, \\
\|g_n\| &\leq \varepsilon \|x_0\| + p_1(n) \varepsilon \|\bar{V}_n\| \|\bar{y}_n\|.
\end{aligned}
\tag{52}
$$

It was shown in [9, 14] that the computed matrix $\bar{V}_n = [\bar{v}_1, \bar{v}_2, \ldots, \bar{v}_n]$ satisfies the recurrence

$$
\begin{aligned}
[\bar{r}_0, A\bar{V}_n] &= \hat{V}_{n+1} \left[\bar{\rho}_0 e_1, \bar{H}_{n+1,n}\right] + [f_n, F_n], \\
\|f_n\| &\leq p_2(m, N) \varepsilon \|A\| \|x_0\| + p_3(N) \varepsilon \|b\|, \\
\|F_n\| &\leq p_4(n, m, N) \varepsilon \|A\| \|\bar{V}_n\|,
\end{aligned}
\tag{53}
$$

where the matrix $\hat{V}_{n+1}$ has exactly orthogonal columns ($\hat{V}_{n+1}^T \hat{V}_{n+1} = I_{n+1}$). The vector $\bar{y}_n$ is a computed solution of the finite precision analogue of the transformed least squares problem (40), and $\bar{r}_0$ satisfies

$$
\|\bar{r}_0 - (b - Ax_0)\| \leq p_5(m, N) \varepsilon \|A\| \|x_0\| + p_6(N) \varepsilon \|b\|.
\tag{54}
$$

For details we refer to [9] and also to [24, pp. 25–26].

Our goal is not to give a complete rounding error analysis of GMRES (for the Householder implementation of classical GMRES this was published in [9], and the modified Gram-Schmidt implementation of classical GMRES has been analyzed in [14, 24] and [21]). We wish to explain that there is a potential weakness of simpler GMRES which may negatively affect its computational behavior in comparison with classical GMRES. For this purpose we can simplify

our description of the GMRES convergence. This allows us to avoid tedious details which would make reading of this section difficult. We will describe the convergence of classical GMRES by the norm of the least squares residual associated with the matrix $A\bar{V}_n$ and the computed initial residual $\bar{r}_0$:

$$\|\hat{r}_n\| \equiv \|\bar{r}_0 - A\bar{V}_n\hat{y}_n\| = \min_y \|\bar{r}_0 - A\bar{V}_ny\|. \tag{55}$$

The analysis in [14, Section 3] as well as numerical experiments confirm that for classical GMRES $\|\hat{r}_n\|$ is close to the norm of the actually computed GMRES residual $\|b - A\bar{x}_n\|$.

It follows immediately from (16) that the residual norm (55) can be bounded in terms of the minimal singular values of matrices $[\bar{r}_0, A\bar{V}_n]$ and $A\bar{V}_n$ as

$$\sigma_{min}([\bar{r}_0, A\bar{V}_n]) \leq \|\hat{r}_n\| \leq \sigma_{min}([\bar{r}_0, A\bar{V}_n]) \left\{ 1 - \frac{\sigma_{min}^2([\bar{r}_0, A\bar{V}_n])}{\sigma_{min}^2(A\bar{V}_n)} + \frac{\|\bar{r}_0\|^2}{\sigma_{min}^2(A\bar{V}_n)} \right\}^{1/2}. \tag{56}$$

We see that convergence of the residual $\hat{r}_n$ is closely related to ill-conditioning of the matrix $[\bar{r}_0, A\bar{V}_n]$, i.e. decreasing $\|\hat{r}_n\|$ leads to ill-conditioning of $[\bar{r}_0, A\bar{V}_n]$. Moreover, it follows from (53) and from classical perturbation theory, see e.g. [13, p. 449], that the minimum singular values of the matrices $[\bar{r}_0, A\bar{V}_n]$ and $[\bar{\rho}_0 e_1, \bar{H}_{n+1,n}]$ are close to each other,

$$\left| \sigma_{min}([\bar{\rho}_0 e_1, \bar{H}_{n+1,n}]) - \sigma_{min}([\bar{r}_0, A\bar{V}_n]) \right| \leq \|[f_n, F_n]\|. \tag{57}$$

Consequently, decreasing $\|\hat{r}_n\|$ leads to ill-conditioning of the matrix $[\bar{\rho}_0 e_1, \bar{H}_{n+1,n}]$. The vector $\bar{y}_n$ from (52) is a computed solution of the least squares problem

$$\min_y \|e_1\bar{\rho}_0 - \bar{H}_{n+1,n}y\|. \tag{58}$$

Using (53), the extremal singular values of $\bar{H}_{n+1,n}$ can be bounded by

$$\|\bar{H}_{n+1,n}\| \leq \|A\bar{V}_n\| + \|F_n\| \leq \|A\|\|\bar{V}_n\| + \|F_n\|, \tag{59}$$

$$\sigma_{min}(\bar{H}_{n+1,n}) \geq \sigma_{min}(A\bar{V}_n) - \|F_n\| \geq \sigma_{min}(A)\,\sigma_{min}(\bar{V}_n) - \|F_n\|. \tag{60}$$

When $\|\hat{r}_n\|$ (and $\|b - A\bar{x}_n\|$) decreases, $\sigma_{min}([\bar{r}_0, A\bar{V}_n])$ and $\sigma_{min}([\bar{\rho}_0 e_1, \bar{H}_{n+1,n}])$ also decrease. However, while the columns of $\bar{V}_n$ (the Arnoldi vectors) keep their linear independence (while $\sigma_{min}(\bar{V}_n) \approx 1$), *the condition number of the computed upper Hessenberg matrix $\bar{H}_{n+1,n}$ is essentially bounded by the condition number of A*. Consequently, until the linear independence of the Arnoldi vectors begins to deteriorate, the solution $\bar{y}_n$ of the transformed least squares problem and the GMRES solution $\bar{x}_n$ are affected by rounding errors in a minimal possible way. This distinguishes classical GMRES from the other MR methods, in particular from simpler GMRES. Finite precision analysis of the QR-factorization of the matrix $\bar{H}_{n+1,n}$ via Givens rotations and of forming the GMRES solution can be found in [9] or [24, (4.6)–(4.12)].

It is important to note that not the orthogonality but the linear independence of the columns of $\bar{V}_n$ (measured by its extremal singular values) plays a decisive

role in the relations (59) and (60). If we use Householder reflections in the Arnoldi process, the loss of orthogonality among the computed columns of $\bar{V}_n$, and the extremal singular values of $\bar{V}_n$ are bounded independent of the system parameters,

$$1 - p_7(n, N)\,\varepsilon \leq \sigma_n(\bar{V}_n) \leq \|\bar{V}_n\| \leq 1 + p_7(n, N)\,\varepsilon. \tag{61}$$

Moreover, it was shown in [9] that the Householder implementation of classical GMRES is backward stable. Assuming that a conjecture similar to (61) holds, the same result can also be shown for the iterated classical or modified Gram-Schmidt implementations, see [9].

In practical computations cheaper orthogonalization techniques like the modified Gram-Schmidt algorithm are used. It is well known that the orthogonality among the columns of $\bar{V}_n$ computed via the modified Gram-Schmidt process will gradually deteriorate. However, from [14, (1.7) and Corollary 2.4] it follows that

$$\|\hat{V}_n - \bar{V}_n\| \leq p_8(n, m, N)\,\varepsilon\kappa([\bar{v}_1, A\bar{V}_{n-1}]) \tag{62}$$

and the minimal singular value and the norm of $\bar{V}_n$ are bounded by

$$1 - \frac{p_9(n, m, N)\,\varepsilon\kappa(A)}{\|\hat{r}_{n-1}\|/\bar{\rho}_0} \leq \sigma_n(\bar{V}_n) \leq \|\bar{V}_n\| \leq 1 + \frac{p_9(n, m, N)\,\varepsilon\kappa(A)}{\|\hat{r}_{n-1}\|/\bar{\rho}_0}. \tag{63}$$

The columns of $\bar{V}_n$ will thus begin to loose their linear independence *only after* the relative residual norm is reduced close to the level $\varepsilon\kappa(A)$. Up to that point the modified Gram-Schmidt implementation of classical GMRES behaves about as well as the Householder implementation.

It was shown in [20, 21] that there is a tight relation between the normwise backward error $\|b - Ax_n\|/(\|A\|\|x_n\| + \|b\|)$ associated with the approximate solution x_n and the condition number of the matrix $[r_0, AV_n]$. Finite precision analogy of this statement will prove normwise backward stability of the modified Gram-Schmidt implementation of classical GMRES. A formal proof will be given elsewhere.

5.2 Simpler GMRES

In simpler GMRES the approximate solution $\bar{x}_n$ computed in finite precision arithmetic satisfies

$$\begin{aligned} \bar{x}_n &= x_0 + [\bar{v}_1, \bar{W}_{n-1}]\,\bar{t}_n + g_n, \\ \|g_n\| &\leq \varepsilon\|x_0\| + p_1(n)\,\varepsilon\|[\bar{v}_1, \bar{W}_{n-1}]\|. \end{aligned} \tag{64}$$

Analogously to (53) for every iteration step n there exists a matrix $\hat{W}_n$ with exactly orthonormal columns ($\hat{W}_n^T \hat{W}_n = I$) such that

$$\begin{aligned} A[\bar{v}_1, \bar{W}_{n-1}] &= \hat{W}_n \bar{S}_n + F_n, \\ \|F_n\| &\leq p_4(n, m, N)\,\varepsilon\|A\|\|[\bar{v}_1, \bar{W}_{n-1}]\|. \end{aligned} \tag{65}$$

The vector of coefficients $\bar{t}_n$ is computed by solving the upper triangular system with the matrix $\bar{S}_n$. From (65) the extremal singular values of the matrix $\bar{S}_n$ are bounded by

$$
\begin{aligned}
\|\bar{S}_n\| &\leq \|A[\bar{v}_1, \bar{W}_{n-1}]\| + \|F_n\| \leq \|A\|\|[\bar{v}_1, \bar{W}_{n-1}]\| + \|F_n\|, \quad (66)\\
\sigma_{min}(\bar{S}_n) &\geq \sigma_{min}(A[\bar{v}_1, \bar{W}_{n-1}]) - \|F_n\|\\
&\geq \sigma_{min}(A)\,\sigma_{min}([\bar{v}_1, \bar{W}_{n-1}]) - \|F_n\|. \quad (67)
\end{aligned}
$$

The minimal singular value of the matrix $[\bar{v}_1, \bar{W}_{n-1}]$ can further be related to the minimal singular value of the matrix $[\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}]$, where $\hat{W}_{n-1}$ comes from the recurrence (65),

$$
\sigma_n([\bar{v}_1, \bar{W}_{n-1}]) \geq \sigma_n([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}]) - \|[\bar{v}_1 - \bar{r}_0/\|\bar{r}_0\|, \bar{W}_{n-1} - \hat{W}_{n-1}]\|. \quad (68)
$$

For the condition number $\kappa([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}])$ it follows from (29), that

$$
\kappa([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}]) = \frac{\|\bar{r}_0\| + (\|\bar{r}_0\|^2 - \|\hat{r}_{n-1}\|^2)^{1/2}}{\|\hat{r}_{n-1}\|}, \quad (69)
$$

where $\hat{r}_{n-1} \equiv (I - \hat{W}_{n-1}\hat{W}_{n-1}^T)\,\bar{r}_0$ is the least squares residual associated with the matrix $\hat{W}_{n-1}$, $\|\hat{r}_{n-1}\| = \min_y \|\bar{r}_0 - \hat{W}_{n-1}y\|$. The identity (69) proves that *convergence of the residual norm $\|\hat{r}_{n-1}\|$ and ill-conditioning of the matrix $[\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}]$ are closely related.*

Summarizing, small $\|\bar{W}_{n-1} - \hat{W}_{n-1}\|$ means $\kappa([\bar{v}_1, \bar{W}_{n-1}]) \approx \kappa([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}])$ (it can be shown that $\|\bar{v}_1 - \bar{r}_0/\|\bar{r}_0\|\| \leq (N+4)\varepsilon$, see [9]). Using (66) and (67), we conclude that *decreasing $\|\hat{r}_{n-1}\|$ may lead to ill-conditioning of the upper triangular matrix $\bar{S}_n$, and thus to a potentially large error in computing the vector $\bar{t}_n$, independent of the (well-)conditioning of the matrix A.* This important fact may negatively affect the numerical accuracy of the approximate solution $\bar{x}_n$ in simpler GMRES in comparison to classical GMRES.

Until $\bar{S}_n$ becomes pathologically ill-conditioned, $\|\hat{r}_n\|$ is (similarly to Subsection 5.1) close to $\|b - A\bar{x}_n\|$. After that the behavior of $\|\hat{r}_n\|$ and $\|b - A\bar{x}_n\|$ may be significantly different.

We have seen that the relation between the condition number of the matrix $\bar{S}_n$ and the condition number of the matrix $[\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}]$ (the decrease of $\|\hat{r}_n\|$) is strongly affected by the size of the term $\|\bar{W}_{n-1} - \hat{W}_{n-1}\|$. In the Householder implementation the computed matrix $\bar{W}_{n-1}$ is, up to a small multiple of the machine precision, close to the matrix $\hat{W}_{n-1}$ with exactly orthogonal columns,

$$
\|\bar{W}_{n-1} - \hat{W}_{n-1}\| \leq p_7(n, N)\,\varepsilon. \quad (70)
$$

It follows from (70) that the condition number $\kappa([\bar{v}_1, \bar{W}_{n-1}])$ is, up to terms proportional to the machine precision, equal to $\kappa([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}])$. In practice one frequently observes that after $\|b - A\bar{x}_n\|/\|\bar{r}_0\|$ reaches some particular point the norm of the computed vector $\bar{t}_n$ starts to increase dramatically (the computed results become irrelevant due to rounding errors) and the residual norm $\|b - A\bar{x}_n\|$ diverges.

For the modified Gram-Schmidt implementation we have

$$\|\bar{W}_{n-1} - \hat{W}_{n-1}\| \leq p_8(n, m, N)\,\varepsilon\kappa(A[\bar{v}_1, \bar{W}_{n-1}]). \tag{71}$$

Because $\kappa(A[\bar{v}_1, \hat{W}_{n-1}])$ is potentially much worse than $\kappa([\bar{v}_1, A\hat{V}_{n-1}])$, the linear independence of the columns of $\bar{W}_n$ often begins to deteriorate much sooner than the linear independence of the columns of $\bar{V}_n$ in classical GMRES. Until that point the modified Gram-Schmidt and Householder implementations of simpler GMRES behave similarly. In subsequent iterations, surprisingly, the behavior of the modified Gram-Schmidt implementation of simpler GMRES may be better than the behavior of the Householder implementation. For the Householder implementation of simpler GMRES the true residual $b - A\bar{x}_n$ often diverges. This has been linked to the tight relation between $\kappa([\bar{r}_0/\|\bar{r}_0\|, \hat{W}_{n-1}])$ and $\kappa([\bar{v}_1, \bar{W}_{n-1}])$, and, consequently, to the relation between the decrease of $\|\hat{r}_n\|$ and the simultaneous increase of $\kappa(\bar{S}_n)$. For the modified Gram-Schmidt implementation, after reaching a certain point no such relations hold. The norm of $\bar{t}_n$ does not diverge and the norm of the true residual remains (and often slightly oscillates) on or below the level corresponding to the turning point for the Householder implementation.

5.3 Numerical experiments

The different behavior of classical and simpler GMRES implementations is demonstrated by numerical examples with the matrix FS1836 from the Harwell-Boeing collection, $N = 183$, $\kappa(A) = 1.5 \times 10^{11}$, $\|A\| = 1.2 \times 10^9$. Experiments were performed using MATLAB 5.2, $\varepsilon = 1.1 \times 10^{-16}$. Householder and modified Gram-Schmidt orthogonalizations have been considered for both classical and simpler GMRES. In all experiments we used $x = (1, \ldots, 1)^T$, $b = Ax$ and $x_0 = 0$ ($\|\bar{r}_0\| = \|b\|$).

Figures 1 and 2 illustrate characteristics of the transformed least squares problem (58) for the Householder and the modified Gram-Schmidt implementations of classical GMRES. In both figures horizontal dotted lines represent $\|A\|$ and the minimal singular value $\sigma_{min}(A)$. The dashed lines show $\|\bar{H}_{n+1,n}\|$, the norm of the computed upper Hessenberg matrix (it almost coincides with $\|A\|$), and the minimal singular value $\sigma_{min}(\bar{H}_{n+1,n})$. The solid line stands for $\sigma_{min}(\bar{V}_n)$, the minimal singular value of the matrix of computed Arnoldi vectors, and the dots depict $\|\bar{y}_n\|$, the norm of the computed solution vector of (58). We see that until the linear independence of the columns of $\bar{V}_n$ in the modified Gram-Schmidt implementation begins to deteriorate, Figures 1 and 2 are almost identical. There is no substantial growth in $\|\bar{y}_n\|$ even after the linear independence of the computed Arnoldi vectors is completely lost (cf. Figure 2).

Similar quantities are illustrated in Figures 3 and 4 for the Householder and the modified Gram-Schmidt implementations of simpler GMRES. The dashed lines here represent $\|\bar{S}_n\|$, the norm of the computed upper triangular matrix, and its minimal singular value $\sigma_{min}(\bar{S}_n)$. The dots denote $\|\bar{t}_n\|$, the norm of the computed solution of the upper triangular system with the matrix $\bar{S}_n$.

Figure 1: Householder implementation of classical GMRES: $\|A\|$ and $\sigma_{min}(A)$ (dotted lines), $\|\bar{H}_{n+1,n}\|$ and $\sigma_{min}(\bar{H}_{n+1,n})$ (dashed lines), $\sigma_{min}(\bar{V}_n)$ (solid line) and $\|\bar{y}_n\|$ (dots).


Figure 2: Modified Gram-Schmidt implementation of classical GMRES: $\|A\|$ and $\sigma_{min}(A)$ (dotted lines), $\|\bar{H}_{n+1,n}\|$ and $\sigma_{min}(\bar{H}_{n+1,n})$ (dashed lines), $\sigma_{min}(\bar{V}_n)$ (solid line) and $\|\bar{y}_n\|$ (dots).

We see that the condition number of the matrix $\bar{H}_{n+1,n}$ is in Figure 1 (the Householder implementation of classical GMRES) approximately bounded by the condition number of A, and for Figure 2 (the modified Gram-Schmidt implementation of classical GMRES) the same is true until $\sigma_{min}(\bar{V}_n)$ begins to deteriorate. In contrast, in both implementations of simpler GMRES, the minimal singular value of $\bar{S}_n$ decreases very soon far below $\sigma_{min}(A)$. Consequently, the accuracy of the computed vector $\bar{t}_n$ deteriorates, and for the Householder implementation $\|\bar{t}_n\|$ diverges. Also note the difference between $\sigma_{min}(\bar{V}_n)$ and $\sigma_{min}(\bar{W}_n)$ in Figures 2 and 4.

In Figure 5 we compare the convergence characteristics for the Householder implementations of both classical GMRES ($\|b - A\bar{x}_n\|/\|b\|$ is represented by the solid line, $\|\hat{r}_n\|/\|b\|$ by dots) and simpler GMRES ($\|b - A\bar{x}_n\|/\|b\|$ is represented by the dashed line, $\|\hat{r}_n\|/\|b\|$ by the dotted line). Figure 5 illustrates our theoretical considerations and shows that the true residual norm of the Householder implementation of simpler GMRES may after some initial reduction diverge. Figure 6 uses similar notation for the illustration of the modified Gram-Schmidt implementations. The residual norm of simpler GMRES again exhibits worse behavior than the residual norm corresponding to classical GMRES.

6 Conclusions

Minimal residual Krylov subspace methods can be formulated and implemented using different bases and different orthogonalization processes. This paper shows that the choice of the basis is fundamental for getting revealing theoretical results about convergence of the methods. It is also important for getting a numerically stable implementation. The choice of the computed basis may strongly affect the numerical behavior of the implementation. It is explained that using the best orthogonalization technique in building the basis does not compensate for the possible loss of accuracy in a given method which is related to the choice of the basis. In particular, it is shown that the classical GMRES based on the Arnoldi process starting from the normalized initial residual (as proposed by Saad and Schultz) has numerical advantages over the simpler GMRES based on the shifted Arnoldi process.

Figure 3: Householder implementation of simpler GMRES: $\|A\|$ and $\sigma_{min}(A)$ (dotted lines), $\|\bar{S}_n\|$ and $\sigma_{min}(\bar{S}_n)$ (dashed lines), $\sigma_{min}(\bar{W}_n)$ (solid line) and $\|\bar{t}_n\|$ (dots).


Figure 4: Modified Gram-Schmidt implementation of simpler GMRES: $\|A\|$ and $\sigma_{min}(A)$ (dotted lines), $\|\bar{S}_n\|$ and $\sigma_{min}(\bar{S}_n)$ (dashed lines), $\sigma_{min}(\bar{W}_n)$ (solid line) and $\|\bar{t}_n\|$ (dots).

Figure 5: Householder implementation of classical and simpler GMRES: Normalized true residual norm $\|b - A\bar{x}_n\|/\|b\|$ (solid line – classical GMRES, dashed line – simpler GMRES), and $\|\hat{r}_n\|/\|b\|$ (dots – classical GMRES, dotted line – simpler GMRES).


Figure 6: Modified Gram-Schmidt implementation of classical and simpler GMRES: Normalized true residual norm $\|b - A\bar{x}_n\|/\|b\|$ (solid line – classical GMRES, dashed line – simpler GMRES), and $\|\hat{r}_n\|/\|b\|$ (dots – classical GMRES, dotted line – simpler GMRES).

7 Appendix

Here we present the implementations of simpler GMRES used throughout the paper.

Modified Gram-Schmidt implementation of simpler GMRES:

x_0, $r_0 = b - Ax_0$, $v_1 = r_0/\|r_0\|$, $w_0 = v_1$

$n = 1, 2, \ldots$

$\quad w_n = Aw_{n-1}$

$\quad j = 1, 2, \ldots, n-1$

$\qquad w_n \leftarrow w_n - \rho_{j,n}w_j, \;\; \rho_{j,n} = (w_n, w_j)$

$\quad w_n \leftarrow w_n/\rho_{n,n}, \;\; \rho_{n,n} = \|w_n\|$

$\quad S_n = \begin{pmatrix} S_{n-1} & \rho_{1,n} \\ & \vdots \\ 0 & \rho_{n,n} \end{pmatrix}, \;\; S_1 = (\rho_{1,1})$

$\quad r_n = r_{n-1} - \zeta_n w_n, \;\; \zeta_n = (r_{n-1}, w_n)$

$\quad$ Solve $S_n\, t_n = (\zeta_1, \ldots, \zeta_n)^T$

$\quad x_n = x_0 + [v_1, w_1, \ldots, w_{n-1}]\, t_n$

Householder implementation of simpler GMRES:

x_0, $r_0 = b - Ax_0$, $v_1 = r_0/\|r_0\|$, $(\zeta_1, \ldots, \zeta_N)^T = r_0$, $w_0 = v_1$

$n = 1, 2, \ldots$

Compute P_n such that $P_n\, Aw_{n-1} = (\rho_{1,n}, \ldots, \rho_{n,n}, 0, \ldots, 0)^T$

$\quad S_n = \begin{pmatrix} S_{n-1} & \rho_{1,n} \\ & \vdots \\ 0 & \rho_{n,n} \end{pmatrix}, \;\; S_1 = (\rho_{1,1})$

$\quad (\zeta_1, \ldots, \zeta_N)^T \leftarrow P_n\, (\zeta_1, \ldots, \zeta_N)$

$\quad$ Solve $\quad S_n\, t_n = (\zeta_1, \ldots, \zeta_n)^T$

$\quad r_n = r_{n-1} - \zeta_n w_n$

$\quad w_n = P_1 \ldots P_n\, e_n$

$\quad x_n = x_0 + [v_1, w_1, \ldots, w_{n-1}]\, t_n$

References

[1] O. Axelsson, *Conjugate gradient type methods for unsymmetric and inconsistent systems of linear equations*, Linear Algebra Appl., 29 (1980), pp. 1–16.

[2] O. Axelsson, *A generalized conjugate gradient, least square method*, Numer. Math., 51 (1987), pp. 209–227.

[3] Z. Bai, D. Hu, and L. Reichel, *A Newton basis GMRES implementation*, IMA J. Numer. Anal., 14 (1994), pp. 563–581.

[4] A. Björck, *Numerical methods for least squares problems*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1996.

[5] S. L. Campbell and C. D. Meyer, Jr., *Generalized inverses of linear transformations*, Pitman (Advanced Publishing Program), Boston, Mass., 1979.

[6] E. de Sturler, *A parallel variant of GMRES(m)*, in Proc. 13th World Congress on Computation and Applied Mathematics (Dublin IMACS'91), 1991, pp. 628–683.

[7] E. de Sturler and H. van der Vorst, *Reducing the effect of global communication in GMRES(m) and CG on parallel distributed memory computers*, Applied Numerical Mathematics (IMACS), 18 (1995), pp. 441–459.

[8] J. Demmel, *The condition number of equivalence transformations that block diagonalize matrix pencils*, SIAM J. Numer. Anal., 20 (1983), pp. 599–610.

[9] J. Drkošová, A. Greenbaum, M. Rozložník, and Z. Strakoš, *Numerical stability of GMRES*, BIT, 35 (1995), pp. 309–330.

[10] S. C. Eisenstat, H. C. Elman, and M. H. Schultz, *Variational iterative methods for nonsymmetric systems of linear equations*, SIAM J. Numer. Anal., 20 (1983), pp. 345–357.

[11] H. Elman, *Iterative methods for large sparse nonsymmetric systems of linear equations.*, PhD thesis, Yale University, New Haven, 1982.

[12] J. Erhel, *A parallel GMRES version for general sparse matrices*, Electron. Trans. Numer. Anal., 3 (1995), pp. 160–176 (electronic).

[13] G. H. Golub and C. F. van Loan, *Matrix computations*, The Johns Hopkins University Press, Baltimore, 3rd ed., 1996.

[14] A. Greenbaum, M. Rozložník, and Z. Strakoš, *Numerical behaviour of the modified Gram-Schmidt GMRES implementation*, BIT, 37 (1997), pp. 706–719.

[15] N. J. Higham, *Accuracy and stability of numerical algorithms*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1996.

[16] I. Ipsen, *A different approach to bounding the minimal residual norm in Krylov methods*, technical report, CRSC, Department of Mathematics, North Carolina State University, 1998.

[17] ——, *Expressions and bounds for the GMRES residual*, technical report, CRSC, Department of Mathematics, North Carolina State University, 1999.

[18] I. M. Khabaza, *An iterative least-square method suitable for solving large sparse matrices*, Comput. J., 6 (1963/1964), pp. 202–206.

[19] J. Liesen, *Construction and analysis of polynomial iterative methods for non-hermitian systems of linear equations*, PhD thesis, Fakultät für Mathematik, Universität Bielefeld, 1998. http://archiv.ub.uni-bielefeld.de/disshabi/mathe.htm.

[20] C. Paige and Z. Strakoš, *Bounds for the least squares distance via scaled total least squares problems*, submitted to Numerische Mathematik, (2000).

[21] ——, *Residual and backward error bounds in minimum residual Krylov subspace methods*, in preparation, (2000).

[22] ——, *Scaled total least squares fundamentals*, submitted to Numerische Mathematik, (2000).

[23] B. Philippe and R. Sidje, *Parallel algorithms for the Arnoldi procedure*, in Iterative Methods in Linear Algebra, II, vol. 3 of IMACS Series in Computational and Applied Mathematics, 1995, pp. 156–165.

[24] M. Rozložník, *Numerical stability of the GMRES method*, PhD thesis, Institute of Computer Science, Academy of Sciences of the Czech Republic, 1997. http://www.cs.cas.cz/~miro.

[25] M. Rozložník and Z. Strakoš, *Variants of the residual minimizing Krylov space methods*, in Proceedings of the XI-th Summer School on Software and Algorithms of Numerical Mathematics, I. Marek, ed., 1995, pp. 208–225.

[26] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.

[27] R. B. Sidje, *Alternatives for parallel Krylov subspace basis computation*, Numer. Linear Algebra Appl., 4 (1997), pp. 305–331.

[28] G. W. Stewart, *Collinearity and least squares regression*, Statist. Sci., 2 (1987), pp. 68–100. With discussion.

[29] S. Van Huffel and J. Vandewalle, *The total least squares problem*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1991.

[30] P. Vinsome, *Orthomin, an iterative method for solving sparse sets of simultaneous linear equations*, in Proceedings of the Fourth Symposium on Reservoir Simulation, Society of Petroleum Engineers of AIME, 1976, pp. 149–159.

[31] C. Vuik and H. A. van der Vorst, *A comparison of some GMRES-like methods*, Linear Algebra Appl., 160 (1992), pp. 131–162.

[32] H. F. Walker, *Implementation of the GMRES method using Householder transformations*, SIAM J. Sci. Stat. Comput., 9 (1988), pp. 152–163.

[33] H. F. Walker, *Implementations of the GMRES method*, Comput. Phys. Comm., 53 (1989), pp. 311–320. Practical iterative methods for large scale computations (Minneapolis, MN, 1988).

[34] H. F. Walker and L. Zhou, *A simpler GMRES*, Numer. Lin. Alg. Appl., 1 (1994), pp. 571–581.

[35] D. M. Young and K. C. Jea, *Generalized conjugate-gradient acceleration of nonsymmetrizable iterative methods*, Linear Algebra Appl., 34 (1980), pp. 159–194.

Research Reports

No.	Authors	Title
00-11	J. Liesen, M. Rozložník, Z. Strakoš	On Convergence and Implementation of Minimal Residual Krylov Subspace Methods for Unsymmetric Linear Systems
00-10	C. Schwab, O. Sterz	A scalar boundary integrodifferential equation for eddy current problems using an impedance boundary condition
00-09	M.H. Gutknecht, M. Rozložník	By how much can residual minimization accelerate the convergence of orthogonal error methods?
00-08	M. Rozložník, V. Simoncini	Short-term recurrences for indefinite preconditioning of saddle point problems
00-07	P. Houston, C. Schwab, E. Süli	Discontinuous hp-Finite Element Methods for Advection-Diffusion Problems
00-06	W.P. Petersen	Estimation of Weak Lensing Parameters by Stochastic Integration
00-05	M.H. Gutknecht	A Matrix Interpretation of the Extended Euclidean Algorithm
00-04	M.J. Grote	Nonreflecting Boundary Conditions for Time Dependent Wave Propagation
00-03	M.H. Gutknecht	On Lanczos-type methods for Wilson fermions
00-02	R. Sperb, R. Strebel	An alternative to Ewald sums. Part 3: Implementation and results
00-01	T. Werder, K. Gerdes, D. Schötzau, C. Schwab	hp Discontinuous Galerkin Time Stepping for Parabolic Problems
99-26	J. Waldvogel	Jost Bürgi and the Discovery of the Logarithms
99-25	H. Brunner, Q. Hu, Q. Lin	Geometric meshes in collocation methods for Volterra integral equations with proportional time delays
99-24	D. Schötzau, Schwab	An hp a-priori error analysis of the DG time-stepping method for initial value problems
99-23	R. Sperb	Optimal sub- or supersolutions in reaction-diffusion problems
99-22	M.H. Gutknecht, M. Rozložník	Residual smoothing techniques: do they improve the limiting accuracy of iterative solvers?
99-21	M.H. Gutknecht, Z. Strakoš	Accuracy of Two Three-term and Three Two-term Recurrences for Krylov Space Solvers
99-20	M.H. Gutknecht, K.J. Ressel	Look-Ahead Procedures for Lanczos-Type Product Methods Based on Three-Term Lanczos Recurrences