

Accuracy of Two Three-term and Three Two-term Recurrences for Krylov Space Solvers*

M.H. Gutknecht and Z. Strakoš[†]

Research Report No. 99-21

October 1999

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

*To appear in SIAM J. Matrix Anal. Appl. (accepted September 22, 1999)

[†]Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (on leave from Institute of Computer Science, Academy of Sciences of the Czech Republic, Prague). This author's work was supported by the ASCR Grant A2030706 and by the GA CR Grant 205/96/0921.

Accuracy of Two Three-term and Three Two-term Recurrences for Krylov Space Solvers*

M.H. Gutknecht and Z. Strakoš[†]

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

Research Report No. 99-21

October 1999

Abstract

It has been widely observed that Krylov space solvers based on two three-term recurrences can give significantly less accurate residuals than mathematically equivalent solvers implemented with three two-term recurrences. In this paper we attempt to clarify and justify this difference theoretically by analyzing the gaps between recursively and explicitly computed residuals. It is shown that, in contrast with the two-term recurrences analyzed by Sleijpen, van der Vorst, and Fokkema (*Numer. Algor.*, 7 (1994), pp. 75–109) and Greenbaum (*SIAM J. Matrix Anal. Appl.*, 18 (1997), pp. 535–551), in the two three-term recurrences the contributions of the local roundoff errors to the analyzed gaps may be dramatically amplified while propagating through the algorithm. This result explains, for example, the well-known behavior of three-term-based versions of the biconjugate gradient method, where large gaps between recursively and explicitly computed residuals are not uncommon. For the conjugate gradient method, however, such a devastating behavior — although possible — is not observed frequently in practical computations, and the difference between two-term and three-term implementations is usually moderate or small. This can also be explained by our results.

Keywords: system of linear algebraic equations, iterative method, Krylov space method, conjugate gradient method, three-term recurrence, accuracy, roundoff

AMS Subject Classification: 65F10, 65G05

*To appear in *SIAM J. Matrix Anal. Appl.* (accepted September 22, 1999)

[†]Department of Mathematics and Computer Science, Emory University, Atlanta, GA 30322, USA (on leave from Institute of Computer Science, Academy of Sciences of the Czech Republic, Prague). This author's work was supported by the ASCR Grant A2030706 and by the GA CR Grant 205/96/0921.

1 Introduction

Among the Krylov space solvers for linear systems $\mathbf{Ax} = \mathbf{b}$ (with \mathbf{A} an $N \times N$ -matrix and \mathbf{b} an N -vector) there are quite a few that are based on three-term recurrences for both the *residuals* \mathbf{r}_n and the *iterates* \mathbf{x}_n . Given an initial approximation \mathbf{x}_0 , we let $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}_0$, $\mathbf{r}_{-1} = \mathbf{o}$, $\mathbf{x}_{-1} = \mathbf{o}$, $\beta_{-1} = 0$ and consider for $n \geq 0$, while $\gamma_n \neq 0$,

$$\begin{aligned}\mathbf{r}_{n+1} &= (\mathbf{Ar}_n - \mathbf{r}_n\alpha_n - \mathbf{r}_{n-1}\beta_{n-1})/\gamma_n, \\ \mathbf{x}_{n+1} &= -(\mathbf{r}_n + \mathbf{x}_n\alpha_n + \mathbf{x}_{n-1}\beta_{n-1})/\gamma_n.\end{aligned}\tag{1}$$

In order that the recurrences (1) be consistent with the residual definition $\mathbf{r}_n \equiv \mathbf{b} - \mathbf{Ax}_n$, the scaling coefficients γ_n need to be chosen according to

$$\gamma_n = -(\alpha_n + \beta_{n-1}),\tag{2}$$

which means that the tridiagonal matrix with coefficients β_{n-1} , α_n , and γ_n in its $(n+1)$ st column has column sums zero; see, for example, § 4.3 of [14].

The list of algorithms based on (1) and (2) includes the Chebyshev iteration [24, 21, 19], the second-order Richardson iteration [21] (which is the stationary form of the Chebyshev iteration), the three-term versions (ORES) of the conjugate-gradient (CG) and the conjugate residual (CR) methods [24, 15], and the three-term version (BIORES) of the unsymmetric or two-sided Lanczos method [18, 14] (which is a variation of the biconjugate gradient or BICG method); see also [2, 15]. On the other hand, for example, neither the version of CG suggested by Rutishauser [21] (based on recurrences for the increments in \mathbf{x} and \mathbf{r}) nor the MINRES algorithm of Paige and Saunders [20], which implements the CR method for symmetric indefinite matrices, nor their SYMMLQ algorithm are covered by our assumptions. An interesting contribution to the rounding error analysis of MINRES and SYMMLQ can be found in [23].

CG, CR, and BICG have better known versions (OMIN and BIOMIN) that are instead based on three two-term recurrences involving in addition to the iterates and their residuals also direction vectors \mathbf{p}_n : for $n \geq 0$,

$$\begin{aligned}\mathbf{p}_n &= \mathbf{r}_n + \mathbf{p}_{n-1}\psi_{n-1}, \\ \mathbf{r}_{n+1} &= \mathbf{r}_n - \mathbf{Ap}_n\omega_n, \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \mathbf{p}_n\omega_n,\end{aligned}\tag{3}$$

with $\mathbf{p}_0 = \mathbf{r}_0$. Other methods like ORTHOMIN [28] use the last two of these recurrences, but have a more complex update formula for the direction vectors. In principle, the version (3) can be obtained from the three-term version (1)-(2) by an LU decomposition of the tridiagonal matrix of recurrence coefficients; see [1, 5, 14, 20]. The folklore is that implementations based on the two-term recurrences (3) are less affected by roundoff than those based on the three-term recurrences (1)-(2). It should be pointed out that the meaning of the phrase *less affected by roundoff* should be carefully specified, otherwise the previous statement is imprecise and can be misleading.

Recent work of Greenbaum [10, 11] shows that under the sole assumption that the last two recurrences (3) hold, there is a limitation on the accuracy of the iterates computed in finite precision arithmetic; the corresponding residuals $\mathbf{b} - \mathbf{A}\mathbf{x}_n$ cannot be expected to decrease below a certain level. (A similar, but somewhat weaker result was given by Sleijpen, van der Vorst, and Fokkema [22].) This level depends primarily on the largest norm of an approximate solution \mathbf{x}_n generated during the iteration, but it does not explicitly depend on how the coefficients ω_n and ψ_n are determined. Since, for example, the BICG method may produce very large intermediate iterates and residuals, this result is of great importance in practice. In contrast, related work on GMRES showed that the size of intermediate iterates does not play a role [4, 12].

In this paper we investigate and answer the question when and why algorithms based on two three-term recurrences of the form (1)–(2) usually do not produce as small residuals as mathematically equivalent algorithms based on three two-term recurrences (3). Similarly to [10, 11, 22, 4], we investigate the gap $\mathbf{f}_n \equiv (\mathbf{b} - \mathbf{A}\mathbf{x}_n) - \mathbf{r}_n$ between the explicitly computed residuals $\mathbf{b} - \mathbf{A}\mathbf{x}_n$ and the recursively updated residuals \mathbf{r}_n . We will refer to the former as *true* residuals and to the latter as *updated* residuals. We show that for computations based on (1)–(2), the gap \mathbf{f}_n satisfies a nonhomogeneous second order difference equation. By writing n steps of this difference equation as the superposition of $n + 1$ homogeneous difference equations (in a different context, this idea has been used by Grcar [8]), we receive an explicit formula for \mathbf{f}_n in terms of the local roundoff errors. The resulting formula contains, in addition to the sum of local errors (which is the analog of the sum that represents the gap \mathbf{f}_n in the case of two-term recurrences analyzed by Greenbaum), each local error multiplied by a set of potentially large multipliers. Moreover, the local errors may become for the two three-term recurrences much larger than for two-term recurrences.

Assume that—in any application for which they are suitable—the methods based on the recurrences (1)–(2) or (3) will eventually produce small updated residuals (whose norm will decrease to the level of roundoff occurring in the finite precision computation of the residual $\mathbf{b} - \mathbf{A}\mathbf{x}$ for the exact solution \mathbf{x}). Then the size of the gap \mathbf{f}_n determines the ultimate attainable accuracy measured by the size of the true residual; a large gap will eventually mean a poor residual $\mathbf{b} - \mathbf{A}\mathbf{x}_n$. The methods based on (1)–(2) are proven to be *in this sense* potentially much less accurate than those based on (3). In this sense, the folklore statement mentioned above is correct.

Our theoretical conclusions are well supported by numerical experiments.

It should be mentioned that the question of the ultimate attainable accuracy of iterative methods was studied by several other authors in addition to those mentioned above; see, for example, [3, 17, 25, 26, 27]. For a more detailed discussion we refer to [11]. However, up to our knowledge, the problem of numerical differences between the recurrences (1)–(2) and (3) was not analyzed in these papers.

2 Local roundoff and the basic recurrence for the gap

In finite precision arithmetic, recurrences (1) have to be replaced by

$$\begin{aligned}\mathbf{r}_{n+1} &= (\mathbf{A}\mathbf{r}_n - \mathbf{r}_n\alpha_n - \mathbf{r}_{n-1}\beta_{n-1} + \mathbf{g}_n)/\gamma_n, \\ \mathbf{x}_{n+1} &= -(\mathbf{r}_n + \mathbf{x}_n\alpha_n + \mathbf{x}_{n-1}\beta_{n-1} - \mathbf{h}_n)/\gamma_n,\end{aligned}\tag{4}$$

where \mathbf{g}_n and \mathbf{h}_n contain all the local rounding errors produced at the step $n+1$, and \mathbf{r}_n , \mathbf{x}_n , etc., denote the actually computed quantities.

The first step of the analysis consists in estimating these local errors. We make the usual assumption that the floating-point arithmetic with roundoff unit ϵ satisfies

$$\text{fl}(a \pm b) = a(1 + \epsilon_1) \pm b(1 + \epsilon_2), \quad |\epsilon_1|, |\epsilon_2| \leq \epsilon, \tag{5}$$

$$\text{fl}(a \text{ op } b) = (a \text{ op } b)(1 + \epsilon_3), \quad |\epsilon_3| \leq \epsilon, \quad \text{op} = *, /.\tag{6}$$

Then the roundoff in the matrix-vector multiplication (computed in a standard way) is bounded according to

$$|\text{fl}(\mathbf{A}\mathbf{p}) - \mathbf{A}\mathbf{p}| \leq m \epsilon |\mathbf{A}| |\mathbf{p}| + \mathcal{O}(\epsilon^2), \tag{7}$$

where $|\mathbf{A}|$ and $|\mathbf{p}|$ denote the elementwise absolute values of \mathbf{A} and \mathbf{p} , and m is the maximal number of nonzeros in any row of \mathbf{A} . Assuming that the first and the third term in (1) are summed up first, we get by applying these rules

$$|\mathbf{g}_n| \leq ((m+3)|\mathbf{A}||\mathbf{r}_n| + 3|\mathbf{r}_n\alpha_n| + 4|\mathbf{r}_{n-1}\beta_{n-1}|) \epsilon + \mathcal{O}(\epsilon^2), \tag{8}$$

$$|\mathbf{h}_n| \leq (3|\mathbf{r}_n| + 3|\mathbf{x}_n\alpha_n| + 4|\mathbf{x}_{n-1}\beta_{n-1}|) \epsilon + \mathcal{O}(\epsilon^2).\tag{9}$$

Both \mathbf{g}_n and \mathbf{h}_n are bounded by a quantity proportional to ϵ , but the behavior of their bounds close to convergence is different. While the updated residual will become eventually small in reasonable computations, and the bound for $|\mathbf{g}_n|$ will decrease correspondingly, the bound for $|\mathbf{h}_n|$ will not. Note that we could consider a norm of \mathbf{g}_n and \mathbf{h}_n here, but there is no real need for this.

In the following estimates we assume that the computed coefficients α_n , β_{n-1} , and γ_n satisfy, in analogy to (2),

$$\gamma_0 = -\alpha_0, \quad \gamma_n = -(\alpha_n + \beta_{n-1}) + \varepsilon_n \quad (n > 0) \tag{10}$$

with error terms ε_n (note that this is another symbol than ϵ) that are bounded by

$$|\varepsilon_n| \leq (|\alpha_n| + |\beta_{n-1}|) \nu \epsilon \quad (n > 0), \tag{11}$$

where ν is a suitable small constant. Note that $\nu = 1$ when γ_n is computed using (2). For later convenience we set $\varepsilon_0 = 0$.

We want to estimate the norm of the difference (or, gap) between updated and true residuals, hence, of

$$\mathbf{f}_n \equiv \mathbf{b} - \mathbf{A}\mathbf{x}_n - \mathbf{r}_n.$$

For $n = 0$, the gap \mathbf{f}_0 is the roundoff in computing \mathbf{r}_0 from \mathbf{A} , \mathbf{x}_0 , and \mathbf{b} ; that is, $\mathbf{f}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0 - \text{fl}(\mathbf{b} - \mathbf{A}\mathbf{x}_0)$, and this is bounded by

$$|\mathbf{f}_0| \leq ((m+1)|\mathbf{A}|\mathbf{x}_0| + |\mathbf{b}|) \epsilon + \mathcal{O}(\epsilon^2). \quad (12)$$

Inserting the recursions (4) and the equality (10) we have

$$\begin{aligned} \mathbf{f}_{n+1} &= \mathbf{b} + (\mathbf{A}\mathbf{r}_n + \mathbf{A}\mathbf{x}_n\alpha_n + \mathbf{A}\mathbf{x}_{n-1}\beta_{n-1} - \mathbf{A}\mathbf{h}_n) \frac{1}{\gamma_n} \\ &\quad - (\mathbf{A}\mathbf{r}_n - \mathbf{r}_n\alpha_n - \mathbf{r}_{n-1}\beta_{n-1} + \mathbf{g}_n) \frac{1}{\gamma_n} \\ &= - [(\mathbf{b} - \mathbf{A}\mathbf{x}_n - \mathbf{r}_n)\alpha_n + (\mathbf{b} - \mathbf{A}\mathbf{x}_{n-1} - \mathbf{r}_{n-1})\beta_{n-1} - \mathbf{b}\varepsilon_n + \mathbf{A}\mathbf{h}_n + \mathbf{g}_n] \frac{1}{\gamma_n} \\ &= - [\mathbf{f}_n\alpha_n + \mathbf{f}_{n-1}\beta_{n-1} - \mathbf{b}\varepsilon_n + \mathbf{A}\mathbf{h}_n + \mathbf{g}_n] \frac{1}{\gamma_n}. \end{aligned} \quad (13)$$

Let us gather the last three terms, the local errors, in

$$\mathbf{l}_n \equiv (-\mathbf{b}\varepsilon_n + \mathbf{A}\mathbf{h}_n + \mathbf{g}_n) \frac{1}{\gamma_n}.$$

By inserting the estimates (8), (9), and (11) we see that

$$\begin{aligned} |\mathbf{l}_n| &\leq [|\mathbf{b}|(|\alpha_n| + |\beta_{n-1}|)\nu + (m+6)|\mathbf{A}|\mathbf{r}_n| + 3(|\mathbf{A}|\mathbf{x}_n| + |\mathbf{r}_n|)|\alpha_n| \\ &\quad + 4(|\mathbf{A}|\mathbf{x}_{n-1}| + |\mathbf{r}_{n-1}|)|\beta_{n-1}|] \frac{\epsilon}{|\gamma_n|} + \mathcal{O}(\epsilon^2). \end{aligned}$$

For $n = 0$, we have $\gamma_0 = -\alpha_0$, $\varepsilon_0 = 0$, and thus

$$\mathbf{l}_0 = (\mathbf{A}\mathbf{h}_0 + \mathbf{g}_0) \frac{1}{\gamma_0}, \quad \mathbf{f}_1 = \mathbf{f}_0 - \mathbf{l}_0.$$

In summary, (13) yields for the gaps \mathbf{f}_n the linear second order difference equation

$$\mathbf{f}_1 = \mathbf{f}_0 - \mathbf{l}_0, \quad \mathbf{f}_{n+1} = - \left(\mathbf{f}_n \frac{\alpha_n}{\gamma_n} + \mathbf{f}_{n-1} \frac{\beta_{n-1}}{\gamma_n} + \mathbf{l}_n \right) \quad (n \geq 1), \quad (14)$$

or, equivalently, the pair of first order difference equations

$$\begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_{n+1} \end{bmatrix} = \begin{bmatrix} \mathbf{O} & \mathbf{I} \\ -\frac{\beta_{n-1}}{\gamma_n}\mathbf{I} & -\frac{\alpha_n}{\gamma_n}\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{f}_{n-1} \\ \mathbf{f}_n \end{bmatrix} - \begin{bmatrix} \mathbf{o} \\ \mathbf{l}_n \end{bmatrix} \quad (n \geq 1) \quad (15)$$

with $\mathbf{f}_1 = \mathbf{f}_0 - \mathbf{l}_0$. These recurrences describe the propagation of the local rounding errors \mathbf{l}_k , $k = 0, \dots, n$. We see that the gap \mathbf{f}_n between the updated and the true residuals after n steps is determined by a nonhomogeneous second order difference equation. This is in sharp contrast to the error behavior of the coupled two-term recurrences, where the gap after n steps is just a simple sum of local errors; see [11]. Consequently, as we will see in the next section, the two three-term recurrences may suffer from a strong amplification of the local errors.

3 Formula for the gap between true and updated residuals

For the moment, assume that the term ε_n in (10) vanishes, that is,

$$-\frac{\alpha_n}{\gamma_n} - \frac{\beta_{n-1}}{\gamma_n} = 1 \quad (16)$$

holds even in finite precision arithmetic. Denote by $\mathbf{z}_{n+1} = \mathcal{D}(\mathbf{z}_{n-m+1}, \mathbf{z}_{n-m}; m)$ the result of m steps of the recurrence

$$\mathbf{z}_{k+1} = -\mathbf{z}_k \frac{\alpha_k}{\gamma_k} - \mathbf{z}_{k-1} \frac{\beta_{k-1}}{\gamma_k}, \quad k = n - m + 1, \dots, n, \quad (17)$$

started at the step $n-m$. Note that due to (16), $\mathbf{z}_{n-m+k+1} = \mathcal{D}(\mathbf{z}_{n-m+1}, \mathbf{z}_{n-m}; k) = \mathbf{z}_{n-m}$ for all k whenever $\mathbf{z}_{n-m+1} = \mathbf{z}_{n-m}$. Our discussion will rely heavily on this fact.

First, we derive how the gap \mathbf{f}_{n+1} is affected by \mathbf{f}_0 . Clearly, the part of this gap that depends on \mathbf{f}_0 is given by

$$\mathcal{D}(\mathbf{f}_0, \mathbf{f}_0; n) = \mathbf{f}_0,$$

that is, \mathbf{f}_0 is not amplified in the process. Next we have to analyze the dependence of \mathbf{f}_{n+1} on the elementary rounding errors \mathbf{l}_0 born in the first step of the algorithm. Clearly, considering (14) for $n = 1$, subtracting and adding $\mathbf{l}_0 \frac{\beta_0}{\gamma_1}$, the contribution of \mathbf{l}_0 to the gap \mathbf{f}_{n+1} can be decomposed into two parts: the part which propagates through the recurrence without any change,

$$\mathcal{D}(-\mathbf{l}_0, -\mathbf{l}_0; n) = -\mathbf{l}_0,$$

and the part depending on the modified local error of the first step,

$$\tilde{\mathbf{l}}_1 \equiv \mathbf{l}_0 \frac{\beta_0}{\gamma_1} + \mathbf{l}_1,$$

which has yet to be analyzed. Repeating the same idea for the steps 2 through n , we can conclude that the gap \mathbf{f}_{n+1} can be written as the following superposition of effects of local errors:

$$\begin{aligned} \mathbf{f}_{n+1} = \mathbf{f}_0 & - \mathbf{l}_0 \\ & - \mathbf{l}_0 \frac{\beta_0}{\gamma_1} - \mathbf{l}_1 \\ & - \mathbf{l}_0 \frac{\beta_0 \beta_1}{\gamma_1 \gamma_2} - \mathbf{l}_1 \frac{\beta_1}{\gamma_2} - \mathbf{l}_2 \\ & \vdots \\ & - \mathbf{l}_0 \frac{\beta_0 \beta_1 \cdots \beta_{n-1}}{\gamma_1 \gamma_2 \cdots \gamma_n} - \dots - \mathbf{l}_{n-1} \frac{\beta_{n-1}}{\gamma_n} - \mathbf{l}_n \end{aligned} \quad (18)$$

Let us give another derivation of this fundamental result. From (15) we see that, in view of $\mathbf{f}_1 = \mathbf{f}_0 - \mathbf{l}_0$,

$$\begin{aligned} \begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_{n+1} \end{bmatrix} &= \prod_{k=1}^n \begin{bmatrix} \mathbf{O} & \mathbf{I} \\ -\frac{\beta_{k-1}}{\gamma_k} \mathbf{I} & -\frac{\alpha_k}{\gamma_k} \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{f}_0 \\ \mathbf{f}_0 \end{bmatrix} \\ &\quad - \sum_{j=0}^n \prod_{k=j+1}^n \begin{bmatrix} \mathbf{O} & \mathbf{I} \\ -\frac{\beta_{k-1}}{\gamma_k} \mathbf{I} & -\frac{\alpha_k}{\gamma_k} \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{o} \\ \mathbf{l}_j \end{bmatrix}. \end{aligned} \quad (19)$$

Here, due to (16), the matrices in the first product leave $[\mathbf{f}_0^\top \ \mathbf{f}_0^\top]^\top$ invariant. In the product that appears after the sum, we split off the last matrix (the one where $k = j + 1$) and apply it to $[\mathbf{o}^\top \ \mathbf{l}_j^\top]^\top$ to get

$$\begin{bmatrix} \mathbf{l}_j \\ -\mathbf{l}_j \frac{\alpha_{j+1}}{\gamma_{j+1}} \end{bmatrix} = \begin{bmatrix} \mathbf{l}_j \\ \mathbf{l}_j \end{bmatrix} + \begin{bmatrix} \mathbf{o} \\ \mathbf{l}_j \frac{\beta_j}{\gamma_{j+1}} \end{bmatrix}.$$

Now we have again a first term that is left invariant by the matrices it is multiplied with and a second term of the form $[\mathbf{o}^\top \ \star]^\top$ that can be treated in the same way that $[\mathbf{o}^\top \ \mathbf{l}_j^\top]^\top$ was treated before. Repeating this trick we finally obtain

$$\begin{aligned} \begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_{n+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{f}_0 \\ \mathbf{f}_0 \end{bmatrix} - \sum_{j=0}^{n-1} \begin{bmatrix} \mathbf{l}_j \\ \mathbf{l}_j \end{bmatrix} \left(1 + \frac{\beta_j}{\gamma_{j+1}} + \dots + \frac{\beta_j \cdots \beta_{n-2}}{\gamma_{j+1} \cdots \gamma_{n-1}} \right) \\ &\quad - \sum_{j=0}^n \begin{bmatrix} \mathbf{o} \\ \mathbf{l}_j \end{bmatrix} \frac{\beta_j \cdots \beta_{n-1}}{\gamma_{j+1} \cdots \gamma_n}, \end{aligned} \quad (20)$$

which is the same as formula (18), written for both \mathbf{f}_n and \mathbf{f}_{n+1} .

Now we describe how the picture changes when the coefficients α_n , β_{n-1} , and γ_n are computed imprecisely, that is, when (16) is replaced by (10). We can follow the analysis described above with the only difference that we should add the effect of the quantity $\mathbf{f}_0 \varepsilon_1 / \gamma_1$ propagating through $n - 1$ steps of the recurrence (17) with $\mathbf{z}_1 := \mathbf{o}$, the effect of $\tilde{\mathbf{l}}_1 \varepsilon_2 / \gamma_2$ propagating through $n - 2$ steps of (17) with $\mathbf{z}_2 := \mathbf{o}$, and so on. As long as the constant ν is small and ε_n is close to the machine precision ϵ , these modifications will only cause effects proportional to $\mathcal{O}(\epsilon^2)$. In (18) we should therefore add terms $\mathcal{O}(\epsilon^2)$ to individual terms of the sum. However, once the size of these terms is considered, the new $\mathcal{O}(\epsilon^2)$ contributions can be thought of as being incorporated in the $\mathcal{O}(\epsilon^2)$ terms already present in the bounds for \mathbf{f}_0 , \mathbf{l}_0 , \dots , \mathbf{l}_n . Therefore we can use (18) in the further analysis with no change and no limitation.

We summarize our main result in the following theorem.

Theorem 3.1 *Up to a term $\mathcal{O}(\epsilon^2)$, the gap \mathbf{f}_{n+1} between true and updated residuals is given by the formula*

$$\begin{aligned}
\mathbf{f}_{n+1} = \mathbf{f}_0 & - \sum_{j=0}^n \mathbf{l}_j \\
& - \mathbf{l}_0 \left(\frac{\beta_0}{\gamma_1} + \frac{\beta_0\beta_1}{\gamma_1\gamma_2} + \dots + \frac{\beta_0\cdots\beta_{n-1}}{\gamma_1\cdots\gamma_n} \right) \\
& - \mathbf{l}_1 \left(\frac{\beta_1}{\gamma_2} + \dots + \frac{\beta_1\cdots\beta_{n-1}}{\gamma_2\cdots\gamma_n} \right) \\
& \vdots \\
& - \mathbf{l}_{n-1} \frac{\beta_{n-1}}{\gamma_n}.
\end{aligned} \tag{21}$$

It is tempting to estimate $\|\mathbf{f}_n\|$ directly on the basis of (19), using an appropriate norm for the 2×2 block matrices. However, the resulting estimate is too generous as it does not take into account the fundamental special properties of these block matrices.

4 Comparison with three coupled two-term recurrences

In our notation, Greenbaum's gap [11] for the coupled two-term recurrences (3) is

$$\mathbf{f}_{n+1}^G = \mathbf{f}_0 - \sum_{j=0}^n \mathbf{l}_j^G, \quad \text{where } \mathbf{l}_j^G \equiv \mathbf{A}\mathbf{h}_j^G + \mathbf{g}_j^G, \tag{22}$$

with \mathbf{g}_n^G and \mathbf{h}_n^G denoting the local rounding errors in the computation of the first two recurrences of (3), analogously to \mathbf{g}_n and \mathbf{h}_n in (4). A comparison of (22) with (21) is instructive.

We point out that the size of the local rounding errors may be considerably larger in the two three-term recurrences than in the three two-term recurrences; the size of the local error \mathbf{l}_j^G in the step n is essentially bounded by $\mathcal{O}(\epsilon)\|\mathbf{A}\| \max_{1 \leq j \leq n} \|\mathbf{x}_j\|$, see [11], where $\|\mathbf{A}\|$ denotes the spectral norm of \mathbf{A} . In our case, a similar term in the bound for $\|\mathbf{l}_n\|$ would be multiplied by the factor $(3|\alpha_n| + 4|\beta_{n-1}|)/|\gamma_n|$, which can be substantially larger than 1; see Section 5 below for the specific case of the CG method. Nevertheless, as documented by our numerical experiments in Section 6, the difference between the implementations based on the two three-term recurrences (1)–(2) and those using the three two-term recurrences (3) cannot be explained by the size of the local rounding error terms only. The amplification of the local errors due to possibly large multipliers plays a substantial if not decisive role: the additional terms in (21) can be of similar size as or even dominate the sum of local rounding errors. If the multipliers become very large, then the two three-term recurrences (1)–(2) are likely to exhibit a dramatically wider gap than the two-term recurrences (3).

Assuming, as in [11], that the updated residuals become eventually negligible, the relations (21) and (22) determine the ultimate attainable accuracy of

the methods based on (1)–(2) and (3), respectively, measured by the norm of the true residuals.

5 Example: conjugate gradient method

For the following discussion of the size of the multiplicative factors

$$\prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j}, \quad (1 \leq i \leq k)$$

we restrict ourselves to symmetric positive definite matrices \mathbf{A} and to the method of conjugate gradients (CG). First, for the simplicity of our exposition, we assume exact arithmetic.

The coefficients in the two-term recurrences (3) are for CG given by [16]

$$\omega_n = \frac{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}{\langle \mathbf{p}_n, \mathbf{A}\mathbf{p}_n \rangle}, \quad \psi_n = \frac{\langle \mathbf{r}_{n+1}, \mathbf{r}_{n+1} \rangle}{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}. \quad (23)$$

Both ω_n and ψ_n are positive. Without specific knowledge about \mathbf{A} and \mathbf{r}_0 we cannot say anything more about their values. More precisely, given any two sequences of positive numbers, $\omega_0, \dots, \omega_{N-1}$ and $\psi_0, \dots, \psi_{N-2}$, there is a symmetric positive definite matrix \mathbf{A} and a vector \mathbf{r}_0 such that the classical OMIN form (Hestenes-Stiefel implementation, denoted further as HS) of the CG method applied to \mathbf{A} with the initial residual \mathbf{r}_0 generates the given coefficients; see Theorem 18:3 of Hestenes and Stiefel [16]. This result allows us to construct examples having any given set of multipliers, and thus to find some with very large gaps. On the other hand, if the matrix \mathbf{A} is reasonably well conditioned and if the CG method converges well, then the bounds derived for the multipliers will show that no substantial amplification of the local rounding errors will occur.

It is well known [20, 5, 1], that by eliminating the direction vectors \mathbf{p}_n in (3) we obtain the three-term (ORES) variant of the CG method with recurrences of the form (1)–(2). From the orthogonality of the residuals we receive

$$\alpha_n = \frac{\langle \mathbf{r}_n, \mathbf{A}\mathbf{r}_n \rangle}{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}, \quad \beta_{n-1} = \gamma_{n-1} \frac{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}{\langle \mathbf{r}_{n-1}, \mathbf{r}_{n-1} \rangle}. \quad (24)$$

Using (23) and $\gamma_n = -(\alpha_n + \beta_{n-1})$, we see that the coefficients of the two implementations are related by

$$\gamma_n = -\frac{1}{\omega_n} < 0, \quad \frac{\beta_{n-1}}{\gamma_n} = \frac{\psi_{n-1}\omega_n}{\omega_{n-1}} \geq 0, \quad \frac{\alpha_n}{\gamma_n} = -1 - \frac{\psi_{n-1}\omega_n}{\omega_{n-1}} \leq -1, \quad (25)$$

where $\psi_{-1} = 0$, $\omega_{-1} = 1$. The equality is attained in the last two formulas only if $\mathbf{x}_n = \mathbf{x}$, that is, if we have reached the solution. We conclude that the

multiplicative factors in (18) have the form

$$\prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} = \frac{\omega_k}{\omega_{i-1}} \prod_{j=i}^k \psi_{j-1}, \quad (26)$$

and therefore, they may exhibit, in general, an arbitrary behavior.

For a given matrix \mathbf{A} and an initial residual \mathbf{r}_0 , it is possible to relate the size of the multipliers to the condition number of \mathbf{A} and the convergence of the CG process measured by the norm of the residuals. First, according to Theorem 5:5 in [16],

$$\frac{\langle \mathbf{p}_n, \mathbf{A} \mathbf{p}_n \rangle}{\langle \mathbf{p}_n, \mathbf{p}_n \rangle} < \frac{1}{\omega_n} = |\gamma_n| < \frac{\langle \mathbf{r}_n, \mathbf{A} \mathbf{r}_n \rangle}{\langle \mathbf{r}_n, \mathbf{r}_n \rangle},$$

which yields, with the spectral norm,

$$\frac{1}{\|\mathbf{A}^{-1}\|} = \frac{1}{\sigma_{\min}(\mathbf{A})} < \frac{1}{\omega_n} = |\gamma_n| < \|\mathbf{A}\|. \quad (27)$$

Rewriting the multipliers in the form

$$\prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} = \frac{\omega_k}{\omega_{i-1}} \frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2},$$

we receive the following bounds:

$$\frac{1}{\kappa(\mathbf{A})} \frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2} \leq \prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} \leq \kappa(\mathbf{A}) \frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2}, \quad (28)$$

where $\kappa(\mathbf{A})$ is the spectral condition number of the matrix \mathbf{A} . Note that

$$\frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2} = \frac{\|\mathbf{A}^{1/2} \mathbf{A}^{1/2} (\mathbf{x} - \mathbf{x}_k)\|^2}{\|\mathbf{A}^{1/2} \mathbf{A}^{1/2} (\mathbf{x} - \mathbf{x}_{i-1})\|^2} \leq \frac{\|\mathbf{A}\|}{\sigma_{\min}(\mathbf{A})} \frac{\|\mathbf{x} - \mathbf{x}_k\|_{\mathbf{A}}^2}{\|\mathbf{x} - \mathbf{x}_{i-1}\|_{\mathbf{A}}^2} \leq \kappa(\mathbf{A})$$

due to the monotonicity of the \mathbf{A} -norm of the error. Consequently,

$$\prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} \leq \kappa^2(\mathbf{A}).$$

As mentioned in Section 2, the bound for the size of the local rounding errors \mathbf{l}_n in the two three-term recurrences (1)–(2) contains the factors $|\alpha_n/\gamma_n|$ and $|\beta_{n-1}/\gamma_n|$. In view of (24) and (27) we have $0 \leq \alpha_n \leq \|\mathbf{A}\|$ and $|\gamma_n|^{-1} \leq \|\mathbf{A}^{-1}\|$. Using (25), we obtain the estimate

$$0 \leq \frac{\beta_{n-1}}{\gamma_n} \leq \left| \frac{\alpha_n}{\gamma_n} \right| \leq \kappa(\mathbf{A}). \quad (29)$$

Surprisingly, to establish that the developed bounds remain relevant in the case of finite precision computation we do not need any extra work: the results

of [9] and [13] imply that in finite precision arithmetic the following slightly relaxed bounds hold:

$$(1 - \vartheta) \frac{1}{\kappa(\mathbf{A})} \frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2} \leq \prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} \leq (1 + \vartheta) \kappa(\mathbf{A}) \frac{\|\mathbf{r}_k\|^2}{\|\mathbf{r}_{i-1}\|^2}, \quad (30)$$

$$\prod_{j=i}^k \frac{\beta_{j-1}}{\gamma_j} \leq (1 + \vartheta) \kappa^2(\mathbf{A}), \quad (31)$$

$$\frac{\beta_{n-1}}{\gamma_n} \leq \left| \frac{\alpha_n}{\gamma_n} \right| \leq (1 + \vartheta) \kappa(\mathbf{A}), \quad (32)$$

where $0 \leq \vartheta \ll 1$ (here, we make the usual assumption about the numerical nonsingularity of the matrix \mathbf{A} ; for details see the references mentioned above). Note, however, that the conclusion we just made is far from trivial. The values of the actually computed recurrence coefficients and of the residual norms may be completely different from their theoretical counterparts. But still, essentially the same bounds hold!

The large size of the upper bounds for ill-conditioned \mathbf{A} suggest that though the size of the local errors may contribute to a possibly large gap between true and updated residuals, the further amplification of the local errors due to large multipliers may have a much stronger effect.

6 Numerical experiments with the conjugate gradient method

The construction of our numerical experiments follows ideas from [16].

Example 1. We consider $N = 48$ and aim at the following values of the coefficients (23) for the classical HS form of the CG method:

$$\begin{aligned} \omega_0 &= \omega_1 = \dots = \omega_{47} = 1, \\ \psi_0 &= 10, \quad \psi_1 = \psi_3 = \dots = \psi_{43} = 0.01, \quad \psi_2 = \dots = \psi_{42} = 100, \\ \psi_{44} &= 10^{-2}, \quad \psi_{45} = 10^{-3}, \quad \psi_{46} = 10^{-4}. \end{aligned} \quad (33)$$

Using the well-known formulas [9]

$$\begin{aligned} \mathbf{T}_{0,0} &= \frac{1}{\omega_0}, \\ \mathbf{T}_{i,i} &= \frac{1}{\omega_i} + \frac{\psi_{i-1}}{\omega_{i-1}}, \\ \mathbf{T}_{i,i-1} &= \mathbf{T}_{i-1,i} = \frac{\sqrt{\psi_{i-1}}}{\omega_{i-1}}, \quad i = 1, \dots, N-1, \end{aligned} \quad (34)$$

we construct an N by N symmetric positive definite tridiagonal matrix \mathbf{T} with spectral norm $\|\mathbf{T}\| = 102$ and condition number $\kappa(\mathbf{T}) \approx 2 \times 10^6$ (for $N = 48$).

For any unitary N by N matrix \mathbf{V} , the CG method (3), (23) applied to the system $\mathbf{Ax} = \mathbf{b}$ with $\mathbf{A} = \mathbf{VTV}^*$ and $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}_0 = \mathbf{Ve}_1$ then generates in steps 1 to N the prescribed coefficients $\omega_j, \psi_j, j = 0, \dots, N-1$, and the residual norms

$$\begin{aligned}\|r_j\| &= 10^{1/2} & \text{for } j = 1, 3, \dots, 43, \\ \|r_j\| &= 10^{-1/2} & \text{for } j = 2, 4, \dots, 44,\end{aligned}$$

with $\|r_j\|$ sharply decreasing in the steps 45 through 48. For an initial residual different from \mathbf{Ve}_1 the behavior of the residual norms will be different, but we may still expect some oscillations and, consequently, some large multipliers.

We have used the construction described above, choosing \mathbf{V} as the unitary matrix resulting from the QR decomposition of a randomly generated N by N matrix; in MATLAB notation $[\mathbf{V}, \mathbf{R}] = qr(\text{randn}(N, N))$. Furthermore, we have chosen $\mathbf{x} = (1, \dots, 1)^\top$, $\mathbf{b} = \mathbf{Ax}$, $\mathbf{x}_0 = \mathbf{o}$, $\mathbf{r}_0 = \mathbf{b}$. Hence, $\mathbf{r}_0 \neq \mathbf{Ve}_1$. Experiments were performed on an Sun Ultra 10 Workstation with $\epsilon \approx 1.11 \times 10^{-16}$ using MATLAB 5.0.

Three implementations of the CG method have been compared: except for Figure 9, solid lines always represent results of the classical OMIN or Hestenes-Stiefel version HS given by (3) and (23), dots those of the Rutishauser variant (R) described in [21], and dashed lines those of the ORES implementation of the form (1)–(2) presented, for example, in [15], p. 143, and denoted here as HY. In Rutishauser's variant the recurrences are, for $n \geq 0$, of the form

$$\begin{aligned}\Delta \mathbf{r}_n &= (-\mathbf{Ar}_n + \Delta \mathbf{r}_{n-1} \eta_{n-1}) \tau_n^{-1}, & \mathbf{r}_{n+1} &= \mathbf{r}_n + \Delta \mathbf{r}_n, \\ \Delta \mathbf{x}_n &= (\mathbf{r}_n + \Delta \mathbf{x}_{n-1} \eta_{n-1}) \tau_n^{-1}, & \mathbf{x}_{n+1} &= \mathbf{x}_n + \Delta \mathbf{x}_n,\end{aligned}\quad (35)$$

and they are started with $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}_0$, $\Delta \mathbf{r}_{-1} = \mathbf{o}$, $\Delta \mathbf{x}_{-1} = \mathbf{o}$, and $\eta_{-1} = 0$. The coefficients are computed according to

$$\tau_n = \frac{\langle \mathbf{r}_n, \mathbf{Ar}_n \rangle}{\langle \mathbf{r}_n, \mathbf{r}_n \rangle} - \eta_{n-1}, \quad \eta_n = \tau_n \frac{\langle \mathbf{r}_{n+1}, \mathbf{r}_{n+1} \rangle}{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}.$$
 (36)

In the HY variant, the following recurrences are used for $n \geq 0$:

$$\begin{aligned}\mathbf{r}_{n+1} &= \theta_{n+1}(-\mu_{n+1} \mathbf{Ar}_n + \mathbf{r}_n) + (1 - \theta_{n+1}) \mathbf{r}_{n-1}, \\ \mathbf{x}_{n+1} &= \theta_{n+1}(\mu_{n+1} \mathbf{r}_n + \mathbf{x}_n) + (1 - \theta_{n+1}) \mathbf{x}_{n-1}.\end{aligned}\quad (37)$$

They are started with $\mathbf{r}_0 = \mathbf{b} - \mathbf{Ax}_0$, $\theta_1 = 1$, $\mathbf{x}_{-1} = \mathbf{o}$, and $\mathbf{r}_{-1} = \mathbf{o}$, and the coefficients are computed according to

$$\mu_n = \frac{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}{\langle \mathbf{r}_n, \mathbf{Ar}_n \rangle}, \quad \theta_{n+1} = \left(1 - \frac{\mu_{n+1}}{\mu_n} \frac{\langle \mathbf{r}_n, \mathbf{r}_n \rangle}{\langle \mathbf{r}_{n-1}, \mathbf{r}_{n-1} \rangle} \frac{1}{\theta_n} \right)^{-1}.$$
 (38)

Clearly, the finite precision equivalent of (37) can be written in the form (4). Consequently, Theorem 3.1 applies, although the bounds for the size of the local errors derived in Section 2 have to be modified slightly.




Figure 1: Example 1: Norms of the updated residuals for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

Norms of the updated residuals are compared in Figure 1. We can see the oscillations followed by the fast convergence for n around 70. Of course, theoretically the method should converge in 48 steps, but, as can be explained by the analysis in [9, 13], the convergence is delayed due to roundoff effects. Norms of the true residuals $\|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$ are shown in Figure 2. Clearly, residual norms of the HY variant stagnate at a significantly worse level than those of the HS variant, as predicted by our analysis.

In Figure 3 the norms of the gaps \mathbf{f}_n we investigated, that is, of the differences between true and updated residuals, are displayed. Note that for the HY variant the gap starts to grow soon, much earlier than one can detect from the two previous figures. Figure 4 shows the behavior of the error norms $\|\mathbf{x} - \mathbf{x}_n\|$. Surprisingly, the differences in the error norms are much less pronounced than those in the true residuals.

Example 2. The second example makes use of the same construction, but now, again for $N = 48$, we aim at

$$\begin{aligned} \omega_0 &= \omega_1 = \dots = \omega_{47} = 1, \\ \psi_0 &= \psi_1 = \dots = \psi_{39} = \sqrt{2}, \quad \psi_{40} = \dots = \psi_{46} = 2^{-7}, \end{aligned} \quad (39)$$

which gives $\|\mathbf{T}\| \approx 4.8$ and $\kappa(\mathbf{T}) \approx 6 \times 10^7$. Again, we consider the system $\mathbf{A}\mathbf{x} = \mathbf{b}$, $\mathbf{A} = \mathbf{V}\mathbf{T}\mathbf{V}^*$, where \mathbf{V} is determined as in Example 1, $\mathbf{x} = (1, \dots, 1)^\top$,




Figure 2: Example 1: Norms of the true residuals computed as $\|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$ for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

$\mathbf{b} = \mathbf{A}\mathbf{x}$. If we chose \mathbf{x}_0 so that $\mathbf{r}_0 = \mathbf{V}\mathbf{e}_1$, we would find residuals with

$$\|r_n\| = (\sqrt{2})^n \quad \text{for } n = 1, 2, \dots, 40,$$

and a sharply decreasing norm in the subsequent steps. However, we have again chosen \mathbf{x}_0 differently, namely $\mathbf{x}_0 = \mathbf{o}$, so that $\mathbf{r}_0 = \mathbf{b}$. Then we do not find an initially increasing, but a quickly decreasing residual norm, both for the updated (Figure 5) and the true residual (Figure 6); note the significant oscillation around $n = 45$. The norm of the true residuals of the HY variant stagnates again at a significantly worse level than in the HS variant. Figure 7 shows the norm of the gaps \mathbf{f}_n . The differences in the norms of the errors, displayed in Figure 8, are again less pronounced.

To illustrate the contribution of the size of local rounding errors to the gap \mathbf{f}_n , we plotted in Figure 9 the size of the coefficients $|\alpha_n/\gamma_n|$, $|\beta_n/\gamma_n|$ and $|1/\gamma_n|$. Clearly, while the gap exhibits a loss of accuracy of about six orders of magnitude, the anticipated contribution of the local errors to this gap is not greater than about two orders of magnitude. The disastrous difference between updated and true residuals must therefore be caused by an amplification of the local rounding errors due to large multipliers. In the analogous figure (not shown) for Example 1 the same behavior is slightly less pronounced.

A detailed explanation of the performance of the Rutishauser variant and of the behavior of the error in all variants requires further work.




Figure 3: Example 1: Norms of the differences (gaps) \mathbf{f}_n between the true and updated residuals for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

7 Conclusions

We have explained why the ultimate attainable accuracy measured by the norm of the true residual $\mathbf{b} - \mathbf{A}\mathbf{x}_n$ can be much worse for implementations of Krylov space methods based on the two three-term recurrences (1)–(2) than for the corresponding implementations based on two-term recurrences of the form (3). For example, in the three-term (ORES) version of the CG method, the gap between true and updated residuals is affected not only by the maximum size of the intermediate iterates $\|\mathbf{x}_k\|$ as in the coupled two-term (OMIN) version, but also by oscillations of the squared norms of the residuals, that is, the quantities $\|\mathbf{r}_k\|^2/\|\mathbf{r}_{i-1}\|^2$, $1 \leq i \leq k$.

Many well-known algorithms like MINRES and SYMMLQ [20], or the three-term and the coupled two-term versions of the QMR method [6, 7], as well as the Rutishauser variant of the CG method are not of the form (1)–(2) or (3). Hence, the results presented in this paper do not apply to them.

Chris Paige suggested another derivation of the results presented in this paper, based entirely on matrix formulations of the algorithms. His approach brings some additional insight into the problem and has potential for further generalization of the results. We hope to report about the results of the joint subsequent work in the near future.

Acknowledgment. The authors would like to thank Anne Greenbaum, Ger-




Figure 4: Example 1: Norms of the errors $\|\mathbf{x} - \mathbf{x}_n\|$ for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

ard Meurant, Chris Paige, Lisa Perrone and Miro Rozložník for their helpful comments.

References

- [1] S. F. ASHBY AND M. H. GUTKNECHT, *A matrix analysis of conjugate gradient algorithms*, in *Advances in Numerical Methods for Large Sparse Sets of Linear Systems*, M. Natori and T. Nodera, eds., no. 9 in *Parallel Processing for Scientific Computing*, Keio University, Yokohama, Japan, 1993, pp. 32–47.
- [2] S. F. ASHBY, T. A. MANTEUFFEL, AND P. E. SAYLOR, *A taxonomy for conjugate gradient methods*, *SIAM J. Numer. Anal.*, 27 (1990), pp. 1542–1568.
- [3] J. A. M. BOLLEN, *Numerical stability of descent methods for solving linear equations*, *Numer. Math.*, 43 (1984), pp. 361–377.
- [4] J. DRKOŠOVÁ, A. GREENBAUM, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Numerical stability of the GMRES method*, *BIT*, 35 (1995), pp. 308–330. .

- [5] R. FLETCHER, *Conjugate gradient methods for indefinite systems*, in Numerical Analysis, Dundee, 1975, G. A. Watson, ed., vol. 506 of Lecture Notes in Mathematics, Springer, Berlin, 1976, pp. 73–89.
- [6] R. W. FREUND AND N. M. NACHTIGAL, *QMR: a quasi-minimal residual method for non-Hermitian linear systems*, Numer. Math., 60 (1991), pp. 315–339. .
- [7] ———, *An implementation of the QMR method based on coupled two-term recurrences*, SIAM J. Sci. Comput., 15 (1994), pp. 313–337.
- [8] J. F. GRGAR, *Analyses of the Lanczos algorithm and of the approximation problem in Richardson’s method*, PhD thesis, University of Illinois at Urbana-Champaign, 1981. Report No. UIUCDCS-R-81-1074.
- [9] A. GREENBAUM, *Predicting the behavior of finite precision Lanczos and conjugate gradient computations*, Linear Algebra Appl., 113 (1989), pp. 7–63.
- [10] ———, *Accuracy of computed solutions from conjugate-gradient-like methods*, in Advances in Numerical Methods for Large Sparse Sets of Linear Systems, M. Natori and T. Nodera, eds., no. 10 in Parallel Processing for Scientific Computing, Keio University, Yokohama, Japan, 1994, pp. 126–138.
- [11] ———, *Estimating the attainable accuracy of recursively computed residual methods*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 535–551. .
- [12] A. GREENBAUM, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Numerical behaviour of the modified Gram-Schmidt GMRES implementation*, BIT, 37 (1997), pp. 706–719.
- [13] A. GREENBAUM AND Z. STRAKOŠ, *Predicting the behavior of finite precision Lanczos and conjugate gradient computations*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 121–137.
- [14] M. H. GUTKNECHT, *Lanczos-type solvers for nonsymmetric linear systems of equations*, Acta Numerica, 6 (1997), pp. 271–397.
- [15] L. HAGEMAN AND D. YOUNG, *Applied Iterative Methods*, Academic Press, Orlando, 1981.
- [16] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, J. Res. Nat. Bureau Standards, 49 (1952), pp. 409–435.
- [17] N. J. HIGHAM AND P. A. KNIGHT, *Componentwise error analysis for stationary iterative methods*, in Linear Algebra, Markov Chains, and Queueing Models, C. D. Meyer and R. J. Plemmons, eds., vol. 48 of IMA Volumes in Mathematics and its Applications, Springer-Verlag, New York, 1993, pp. 29–46.

- [18] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Nat. Bureau Standards, 45 (1950), pp. 255–281.
- [19] T. A. MANTEUFFEL, *The Tchebyshev iteration for nonsymmetric linear systems*, Numer. Math., 28 (1977), pp. 307–327.
- [20] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [21] H. RUTISHAUSER, *Theory of gradient methods*, in Refined Iterative Methods for Computation of the Solution and the Eigenvalues of Self-Adjoint Boundary Value Problems, Mitt. Inst. angew. Math. ETH Zürich, Nr. 8, Birkhäuser, Basel, 1959, pp. 24–49.
- [22] G. L. G. SLEIJPEN, H. A. VAN DER VORST, AND D. R. FOKKEMA, *BiCGstab(l) and other hybrid Bi-CG methods*, Numerical Algorithms, 7 (1994), pp. 75–109. .
- [23] G. L. G. SLEIJPEN, H. A. VAN DER VORST, AND J. MODERSITZKI, *The main effects of rounding errors in Krylov solvers for symmetric linear systems*, SIAM J. Matrix Anal. Appl. Submitted.
- [24] E. STIEFEL, *Relaxationsmethoden bester Strategie zur Lösung linearer Gleichungssysteme*, Comm. Math. Helv., 29 (1955), pp. 157–179.
- [25] H. WOŹNIAKOWSKI, *Numerical stability of the Chebyshev method for the solution of large linear systems*, Numer. Math., 28 (1977), pp. 191–209.
- [26] ———, *Round-off error analysis of iterations for large linear systems*, Numer. Math., 30 (1978), pp. 301–314.
- [27] ———, *Round-off error analysis of a new class of conjugate-gradient algorithms*, Linear Algebra Appl., 29 (1980), pp. 507–529.
- [28] D. M. YOUNG AND K. C. JEA, *Generalized conjugate-gradient acceleration of nonsymmetrizable iterative methods*, Linear Algebra Appl., 34 (1980), pp. 159–194.

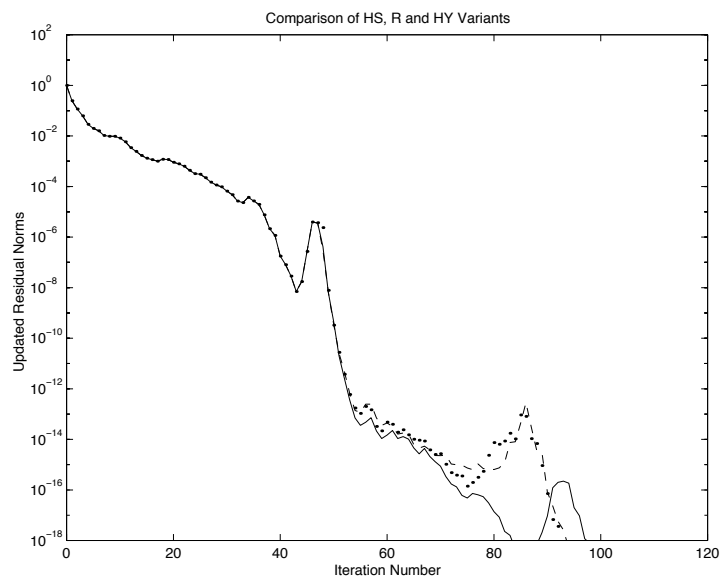


Figure 5: Example 2: Norms of the updated residuals for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.




Figure 6: Example 2: Norms of the true residuals computed as $\|\mathbf{b} - \mathbf{A}\mathbf{x}_n\|$ for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

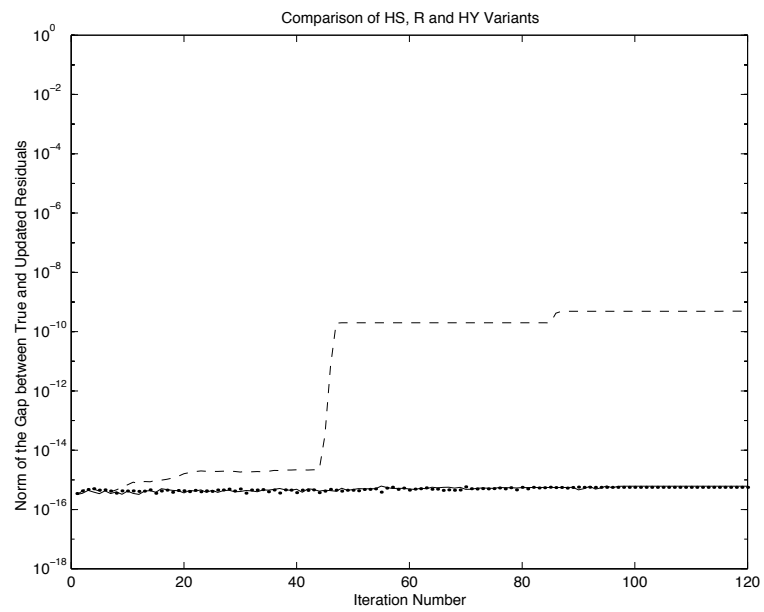


Figure 7: Example 2: Norms of the differences (gaps) \mathbf{f}_n between the true and updated residuals for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.




Figure 8: Example 2: Norms of the errors $\|\mathbf{x} - \mathbf{x}_n\|$ for the two-term (HS; solid line), three-term (HY; dashed line), and Rutishauser (R; dots) variants of the CG method.

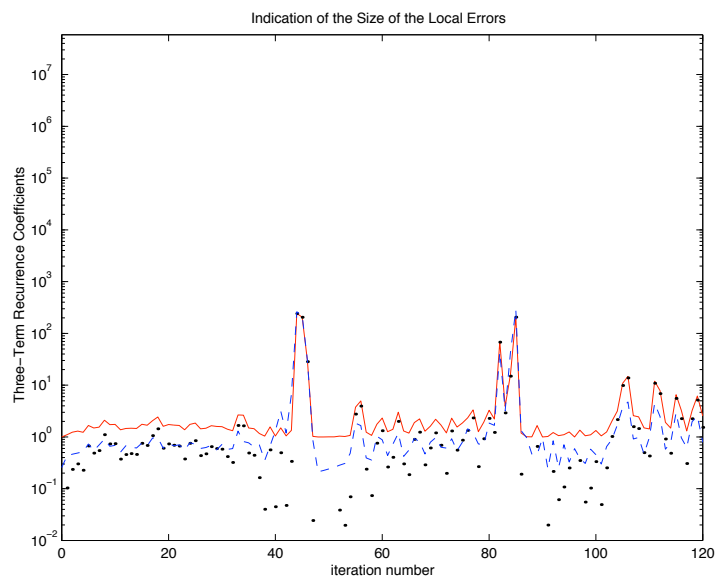


Figure 9: Example 2: Size of the three-term recurrence coefficients $|\alpha_n/\gamma_n|$ (solid line), β_n/γ_n (dots) and $|1/\gamma_n|$ (dashed line) for the HY variant of the CG method

Research Reports

No.	Authors	Title
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
99-19	M. Grote	Nonreflecting Boundary Conditions For Elastodynamic Scattering
99-18	J. Pitkäranta, A.-M. Matache, C. Schwab	Fourier mode analysis of layers in shallow shell deformations
99-17	K. Gerdes, J.M. Melenk, D. Schötzau, C. Schwab	The <i>hp</i> -Version of the Streamline Diffusion Finite Element Method in Two Space Dimensions
99-16	R. Klees, M. van Gelderen, C. Lage, C. Schwab	Fast numerical solution of the linearized Molodensky problem
99-15	J.M. Melenk, K. Gerdes, C. Schwab	Fully Discrete <i>hp</i> -Finite Elements: Fast Quadrature
99-14	E. Süli, P. Houston, C. Schwab	<i>hp</i> -Finite Element Methods for Hyperbolic Problems
99-13	E. Süli, C. Schwab, P. Houston	<i>hp</i> -DGFEM for Partial Differential Equations with Nonnegative Characteristic Form
99-12	K. Nipp	Numerical integration of differential algebraic systems and invariant manifolds
99-11	C. Lage, C. Schwab	Advanced boundary element algorithms
99-10	D. Schötzau, C. Schwab	Exponential Convergence in a Galerkin Least Squares <i>hp</i> -FEM for Stokes Flow
99-09	A.M. Matache, C. Schwab	Homogenization via <i>p</i> -FEM for Problems with Microstructure
99-08	D. Braess, C. Schwab	Approximation on Simplices with respect to Weighted Sobolev Norms
99-07	M. Feistauer, C. Schwab	Coupled Problems for Viscous Incompressible Flow in Exterior Domains
99-06	J. Maurer, M. Fey	A Scale-Residual Model for Large-Eddy Simulation
99-05	M.J. Grote	Am Rande des Unendlichen: Numerische Verfahren für unbegrenzte Gebiete
99-04	D. Schötzau, C. Schwab	Time Discretization of Parabolic Problems by the <i>hp</i> -Version of the Discontinuous Galerkin Finite Element Method
99-03	S.A. Zimmermann	The Method of Transport for the Euler Equations Written as a Kinetic Scheme