The Role of the Inner Product in Stopping Criteria for Conjugate Gradient Iterations*

S. F. Ashby† M. J. Holst‡ T. A. Manteuffel§
P. E. Saylor¶

May 14, 1999

Abstract.

Two natural and efficient stopping criteria are derived for conjugate gradient (CG) methods, based on iteration parameters. The derivation makes use of the inner product matrix B defining the CG method. In particular, the relationship between the eigenvalues and B-norm of a matrix is investigated, and it is shown that the ratio of largest to smallest eigenvalues defines the B-condition number of the matrix. Upper and lower bounds on various measures of the error are also given. The compound stopping criterion presented here is an obvious "default" in software packages because it does not require any additional norm computations. *Key words:* stopping criteria, conjugate gradient methods, B-normal matrices

1 Introduction

Unlike a direct method, an iterative method for solving a linear system $Ax = b$ must be told when to stop. The choice of stopping criterion is important. If it is too stringent, the iteration will expend unnecessary effort computing (or attempting to compute) an excessively accurate approximation to the solution. On the other hand, if the stopping criterion is too lax, a poor approximation may be returned.

*This work was supported in part by Department of Energy grants DE-FG02-90ER25086, DE-FG02-91ER25099, and DE-FG03-93ER25217; National Science Foundation grants DMS-9015259, DMS-8911410, and DMS-8704169; the Schlumberger Foundation; and the Applied Mathematical Sciences subprogram of the Office of Scientific Computing, Department of Energy, by Lawrence Livermore National Laboratory under contract W-7405-ENG-48. Effort sponsored in part by the Air Force Office of Scientific Research, Air Force Materials Command, USAF, under grant number 1-5-20644 at UIUC. The US Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the US Government. This work also supported in part by NASA CAN NCC S5-153.

†Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, P.O. Box 808, L-561, Livermore, CA 94551 (sfashby@llnl.gov).

‡Department of Mathematics, UC San Diego, La Jolla, CA 92093 (mholst@math.ucsd.edu).

§Department of Applied Mathematics, Campus Box 526, University of Colorado at Boulder, Boulder, CO 80309-0526 (tmanteuf@colorado.edu).

¶Department of Computer Science, University of Illinois, Urbana, IL 61801 (saylor@cs.uiuc.edu).

In most implementations of an iterative method, the iteration is halted when some norm of either the residual or relative residual is reduced to a user-specified tolerance. This stopping criterion is often convenient: The residual is readily available and its norm is easily computed. However, in many conjugate gradient (CG) iterations, this norm computation represents an additional expense. That is, the *norm* of the residual is not needed for the iteration, but only for the stopping criterion. For example, in the preconditioned conjugate gradient (PCG) method [8], the inner product $\langle r, s\rangle$ is available, but not $\|r\|_2$ or $\|s\|_2$. (Here r is the residual and $s = Cr$ is the preconditioned residual, defined in §2.) It would be more efficient to employ a stopping criterion that makes use of only previously computed iteration quantities. Moreover, stopping criteria based on the residual are not necessarily the most desirable. One might prefer to stop when the error is reduced (in some norm) to some tolerance.

In this paper we derive several bounds that give rise to stopping criteria based on the error. In each case, an estimate of a condition number of a matrix (or a related quantity) is required. Although this may seem prohibitive, modern computer codes are capable of computing such estimates dynamically, especially for conjugate gradient methods because of their relationship to the Lanczos procedure. In particular, we will introduce an inexpensive stopping criterion based on the same norm used to define the CG optimization property. An algorithm using this stopping criterion would halt after the quantity being minimized was brought below some threshold. Such an inexpensive stopping criterion based on the relative error would be a good default criterion in a software package of iterative methods.

1.1 Preliminaries

In this section we establish some notation and terminology. We denote the Euclidean inner product and spectral condition number by $\langle\cdot,\cdot\rangle$ and $\kappa(\cdot)$, respectively. In addition, $\|\cdot\|$ will denote both the Euclidean vector norm and the corresponding spectral matrix norm, depending on the context. If B is an Hermitian positive definite (hpd) matrix, it may be used to define a new inner product, norm, and condition number, namely, $\langle B\cdot,\cdot\rangle$, $\|\cdot\|_B \equiv \langle B\cdot,\cdot\rangle^{1/2}$, and $\kappa_B(G) = \|G\|_B\|G^{-1}\|_B$. In comparing different condition numbers, we will sometimes refer to $\kappa(\cdot)$ as the I-condition number to emphasize the underlying inner product matrix.

Throughout this paper we make use of the theoretical framework for CG methods that was developed in [6]. In particular, we assume that the reader is familiar with the concepts of B-normality and B-definiteness. The former is simply a generalization of the usual definition of a normal matrix to the B-inner product. The matrix G is *B-normal* (for some B) if and only if it has a complete set of B-orthogonal eigenvectors. If, in addition, the matrix has collinear eigenvalues, it is said to be *B-normal(1)*. The matrix G is *B-self-adjoint* (i.e., self-adjoint with respect to the B-inner product) if and only if BG is Hermitian, a property that is easily verified. Thus, G is B-self adjoint if $\langle Gu, v\rangle_B = \langle u, Gv\rangle_B$, i.e., $\langle Gu, Bv\rangle = \langle u, BGv\rangle$ or $G^*B = BG$. We also remark that the matrix G is B-normal(1) if and only if it has the form $e^{i\theta}(irI + H)$, where r is real and H is B-self-adjoint. (See the paper of Faber and Manteuffel for details [14].)

Another important concept is that of definiteness. The matrix G is *B-definite* if and only if BG is definite, that is, $0 \notin F_B(G)$, where $F_B(G) = \{\langle BGv, v\rangle/\langle Bv, v\rangle : v \neq 0\}$ is the *B-field of values*. Also, A is *B-indefinite* if BG is indefinite. The *field of values* of G is the I-field of values of G, and G is *definite* if it is I-definite. Also observe as an incidental fact that G is not necessarily B-definite if it is definite, and conversely: For

$G = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$, $B = \begin{bmatrix} 1 & 0 \\ 0 & 1/4 \end{bmatrix}$, $v = \begin{bmatrix} 1 \\ -2 \end{bmatrix}$, G is definite, whereas $\langle v, BGv \rangle = 0$. A useful relationship is that if a matrix G is B-normal, then $F_B(G) = H(G)$, where $H(G)$ denotes the *convex hull of the eigenvalues of* G (usually shortened to *convex hull*. See [6, 20, 21] for a detailed discussion of B-normality and B-definiteness.

Here is an overview of the paper. In the next section we briefly review the theory of conjugate gradient methods, paying particular attention to the importance of the inner product. In §3 we review several well-known bounds and introduce a few variants. We explore the relationship between the B- and Euclidean inner products in §4, and establish several results which relate the norms, spectral radii, and condition numbers defined by the two inner products. In §5, these relationships are used to derive a general result that is the basis for the "natural" stopping criterion we propose. An alternative criterion is presented in §6 and condition number estimation is discussed in §7. Other approaches are surveyed in §8; we see that our estimates are generalizations of those in the literature. Our stopping criteria are instantiated for twelve CG methods in §9. Results of numerical experiments are given in §10, where we see that a practical criterion based on eigenvalue estimates is not a strict upper bound in the early stages of an iteration. In §11, we summarize our results and propose future work concerning improvements in the upper bound on which our (primary) criterion is based.

2 CG methods

It was shown in [6] that any conjugate gradient method is characterized by just three matrices: an hpd *inner product matrix* B, a *left preconditioning matrix* C, and the original system matrix A. The resulting method, denoted by CG(B, C, A), minimizes the B-norm of the error over a *translated Krylov subspace*. To be precise, at step i it minimizes

$$\|e_i\|_B \equiv \langle Be_i, e_i \rangle^{1/2} \tag{2.1}$$

over

$$x_0 + K_i(Cr_0, CA)$$

where

$$K_i(Cr_0, CA) \equiv \text{sp}\{Cr_0, CACr_0, (CA)^2Cr_0, \ldots, (CA)^{i-1}Cr_0\}. \tag{2.2}$$

Note that the matrix B is used to define a norm in which minimization is sought.

To implement this method, one must employ an *algorithm*. The optimality property is usually realized by enforcing an equivalent orthogonality condition. In the most commonly used algorithms, Orthodir and Orthomin [31], this requires a Gram-Schmidt-like procedure for generating a B-orthogonal basis for the Krylov subspace. For general matrices, this would entail storage of all past basis vectors, which is prohibitive when A is large. However, if the *preconditioned matrix* CA is B-normal(1), one may use an efficient three-term recursion [14]. For such matrices, Orthodir and Orthomin simplify, yielding algorithms we call Odir and Omin (Fig. 2.1). The former is the most faithful implementation of a three-term CG method: it converges whenever CA is B-normal(1). The latter is cheaper, but is more restrictive in its application: it converges whenever CA is B-normal(1) and BCA is definite. A particular CG method is obtained by specifying B and C, which must be chosen so that the resulting method is *computable*, i.e., expressible in terms of known quantities. For example, if A is hpd, then $B = A$ and $C = I$ yields the original CG method of Hestenes and Stiefel [19], which we call CGHS. For A Hermitian, the choice $B = A^2$ and $C = I$ gives the conjugate residual (CR) method. It is computable because $\langle Be_i, p_i \rangle = \langle r_i, p_i \rangle$. Several other choices are

Odir	Omin
$p_0 = Cr_0$	$\widehat{p}_0 = Cr_0$
$\alpha_i = \langle Be_i, p_i\rangle / \langle Bp_i, p_i\rangle$	$\widehat{\alpha}_i = \langle Be_i, \widehat{p}_i\rangle / \langle B\widehat{p}_i, \widehat{p}_i\rangle$
$x_{i+1} = x_i + \alpha_i p_i$	$x_{i+1} = x_i + \widehat{\alpha}_i \widehat{p}_i$
$r_{i+1} = r_i - \alpha_i A p_i$	$r_{i+1} = r_i - \widehat{\alpha}_i A\widehat{p}_i$
$s_{i+1} = Cr_{i+1}$	$s_{i+1} = Cr_{i+1}$
$\gamma_i = \langle BCAp_i, p_i\rangle / \langle Bp_i, p_i\rangle$	
$\sigma_i = \langle BCAp_i, p_{i-1}\rangle / \langle Bp_{i-1}, p_{i-1}\rangle$	$\beta_i = -\langle BCr_{i+1}, \widehat{p}_i\rangle / \langle B\widehat{p}_i, \widehat{p}_i\rangle$
$p_{i+1} = CAp_i - \gamma_i p_i - \sigma_i p_{i-1}$	$\widehat{p}_{i+1} = s_{i+1} + \beta_i \widehat{p}_i$

Figure 2.1: Odir and Omin implementations of CG(B, C, A)

shown in Table 2.1. For each method, B must be hpd, which imposes a restriction on the class of matrices for which the Odir algorithm is applicable.[1] We remark that the algorithms given in Fig. 2.1 may be rewritten to avoid unnecessary inner products and matrix-vector products. The most efficient implementation depends on the particular method, that is, the choice of B, C, and A. See [6] for details.

For some methods, additional restrictions may be required to guarantee the convergence of the more efficient Omin algorithm, as in the conjugate residual (CR) method. The methods PCG and PCR are preconditioned variants of CGHS and CR, respectively. The methods PCGCA and PPCG are polynomial preconditioned CG methods. The former is obtained by applying CGHS to $C(A)A$ and the latter is simply PCG with $C = C(A)$. As we will see, we can exploit the commutativity of $C(A)$ and A to obtain efficient stopping criteria. The last six methods are all variants of the (preconditioned) normal equations. In Table 1 and elsewhere, the matrix M stands for some approximation to A whereas C is the total preconditioner. See [6] for a thorough discussion of each of these methods and their underlying patterns.

3 Error bounds and stopping criteria

The simplest and most common stopping criteria are based on the absolute and relative residual:

$$\|r_i\| \leq \epsilon \tag{3.1}$$

and

$$\frac{\|r_i\|}{\|b\|} \leq \epsilon \tag{3.2}$$

where $\epsilon > 0$ is a user-provided error tolerance. Since one usually prefers a relative stopping criterion in practice, we will focus on (3.2). Observe that the first criterion can be obtained from the second by an appropriate scaling of ϵ by $\|b\|$.

Despite its simplicity, this stopping criterion has two disadvantages. First, it is biased toward the large eigenvalues of the matrix A. This subtle fact, which is often

[1] For simplicity we list sufficient conditions in the Table 2.1. See [6] for necessary and sufficient conditions.

Table 2.1: Twelve CG methods.

Method	B	CA	Odir Restrictions	Omin Restrictions
CGHS	A	A	A hpd	A hpd
CR	A^2	A	A herm	A hpd
PCG	A	CA	A hpd, C herm	A hpd, C hpd
PCR	ACA	CA	A herm, C hpd	A hpd, C hpd
PCGCA	$C(A)A$	$C(A)A$	$C(A)A$ hpd	$C(A)A$ hpd
PPCG	A	$C(A)A$	A hpd	A hpd, $C(A)$ hpd
CGNR	A^*A	A^*A	none	none
CGNE	I	A^*A	none	none
PCGNS	$(M^{-1}A)^*M^{-1}A$	$(M^{-1}A)^*M^{-1}A$	none	none
PCGNE	I	$(M^{-1}A)^*M^{-1}A$	none	none
PCGNR	A^*A	$(M^*M)^{-1}A^*A$	none	none
PCGNM	M^*M	$(M^*M)^{-1}A^*A$	none	none

overlooked, can skew one's conclusions about the results of a numerical investigation. Second, although the residual is usually available, its norm may not be, as in the PCG method. To implement (3.2) would require an additional norm computation, which might add significantly to the overall cost of the iteration.

To overcome the first deficiency, we prefer stopping criteria based on the relative *error*. Unfortunately, this quantity is generally not known, and so we must base our stopping criterion on an error *bound* that is easily computed from available iteration parameters. To elucidate this point, let S be an arbitrary hpd matrix (not necessarily related to the CG inner product matrix). Then the following bounds on the relative error in the S-norm are easily derived:

$$\kappa_S^{-1}(A)\frac{\|r_i\|_S}{\|b\|_S} \;\le\; \frac{\|e_i\|_S}{\|x\|_S} \;\le\; \kappa_S(A)\frac{\|r_i\|_S}{\|b\|_S} \tag{3.3}$$

and

$$\kappa_S^{-1}(CA)\frac{\|s_i\|_S}{\|Cb\|_S} \;\le\; \frac{\|e_i\|_S}{\|x\|_S} \;\le\; \kappa_S(CA)\frac{\|s_i\|_S}{\|Cb\|_S}. \tag{3.4}$$

The former is based on the original matrix A and residual r_i; the latter is based on their preconditioned counterparts, CA and s_i. To obtain a stopping criterion, one halts when the right hand side of (3.3) or (3.4) is less than a tolerance ϵ. Of course, such a stopping criterion requires an estimate for $\kappa_S(A)$ or $\kappa_S(CA)$. As we will show in §7, estimates for $\kappa_S(CA)$ can be obtained from the CG iteration parameters when $S = B$ by exploiting the relationship between CG and Lanczos.

An obvious choice is $S = I$, in which case we obtain upper bounds on the relative

error in the Euclidean norm,

$$\frac{\|e_i\|}{\|x\|} \leq \kappa(A)\frac{\|r_i\|}{\|b\|} \tag{3.5}$$

and

$$\frac{\|e_i\|}{\|x\|} \leq \kappa(CA)\frac{\|s_i\|}{\|Cb\|}. \tag{3.6}$$

Unfortunately, to implement a stopping criterion based on either bound, one might need to compute an inner product (e.g., a norm) that would not otherwise be needed to execute the iteration. For example, in the PCG method, the inner product $\langle r_i, s_i\rangle$ is available, but not $\|r_i\|^2 = \langle r_i, r_i\rangle$ or $\|s_i\|^2 = \langle s_i, s_i\rangle$. Of course, one could explicitly compute either norm, but that would represent a 20% increase in the amount of vector-vector work in a typical CG algorithm (assuming the stopping criterion were checked at every step). More important, it is difficult to estimate either $\kappa(A)$ or $\kappa(CA)$. Contrary to popular belief, the CG–Lanczos relationship cannot be exploited. The reason is this: although one can estimate easily the extreme eigenvalues of CA (but not those of A, in general), their ratio does *not* necessarily yield the *spectral* condition number of CA. Instead, this ratio defines the B-condition number, as we will explain below. Thus, neither (3.5) nor (3.6) is an attractive stopping criterion. Instead, we wish to devise a stopping criterion that makes use of quantities already computed during the course of the iteration.

4 Relationships between the B- and Euclidean inner products

In the context of conjugate gradient methods, it is natural to consider $S = B$, where B is the inner product matrix defining the method. (Recall that a CG method minimizes the B-norm of the error at each step of the iteration over a certain subspace.) Since B varies with the CG method, the stopping criterion will depend on the method. Although this makes it difficult to compare different methods, it does enable an inexpensive and natural stopping criterion that would be a reasonable default in software packages.

Since CG(B, C, A) is defined with respect to the B-inner product, we will need to generalize several familiar results to this new inner product. Recall that if the matrix G is normal (with respect to the Euclidean inner product), then its spectral matrix norm is given by its spectral radius, $\rho(G)$. The following lemma generalizes this result to the B-inner product.

LEMMA 4.1. *If G is B-normal, then $\|G\|_B = \rho(G)$.*

PROOF. Since B is hpd, we have

$$\|G\|_B = \|B^{1/2}GB^{-1/2}\| = \sigma_{\max}(B^{1/2}GB^{-1/2}) = \rho^{1/2}(B^{-1/2}G^*BGB^{-1/2})$$

where $\sigma_{\max}$ denotes the largest singular value. Since G is B-normal, $G = QJQ^{-1}$ for some diagonal J and B-orthogonal Q (i.e., $Q^*BQ = I$). Thus,

$$B^{-1/2}G^*BGB^{-1/2} = B^{-1/2}Q^{-*}J^*JQ^{-1}B^{-1/2} = B^{1/2}QJ^*JQ^{-1}B^{-1/2}$$

which is similar to J^*J. Therefore, $\|G\|_B = \rho(J) = \rho(G)$. □

In what follows, let $\sigma(G)$ denote the spectrum of the matrix G.

COROLLARY 4.2. *If G is B-normal, then $\kappa_B(G) = \max_{\lambda\in\sigma(G)}|\lambda| / \min_{\lambda\in\sigma(G)}|\lambda|$.*

PROOF. The result follows from Lemma 4.1 applied to G and G^{-1}. □

In other words, if a matrix G is B-normal, the ratio of its largest to smallest eigenvalue is the B-condition number of G. It is important to emphasize that this ratio

is, in general, different from the spectral condition number of G—unless the matrix happens also to be I-normal.

The distinction between the B- and I-condition numbers is an important one and is crucial to the choice of stopping criteria. Despite this, it has been largely ignored in the literature. One reason is this: they are the same for some of the most common CG methods. To see this, we will first prove a few results, and then show in §9 that their hypotheses are satisfied by several well-known CG methods.

LEMMA 4.3. *Let B be hpd. If either $B^{1/2}G = G^*B^{1/2}$ or $B^{1/2}G = GB^{1/2}$, then $\|G\|_B = \|G\|$. Consequently, $\kappa_B(G) = \kappa(G)$.*

PROOF. Recall from the proof of Lemma 4.1 that $\|G\|_B = \|B^{1/2}GB^{-1/2}\|$. If either hypothesis holds, the result follows immediately. □

Remark: The first condition means that G is $B^{1/2}$-self-adjoint. However, note that if G is merely B-self-adjoint, then in general $\|G\|_B \neq \|G\|$. For example, consider

$$G = \begin{pmatrix} 1 & 5 \\ 0 & 2 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} 1 & 0 \\ -5 & 1 \end{pmatrix} \begin{pmatrix} 1 & -5 \\ 0 & 1 \end{pmatrix}.$$

The matrix G is B-self-adjoint, but $2 = \|G\|_B \neq \|G\| \approx 5.5$. Also note that if the second hypothesis of Lemma 4.3 is true, then G is B-self-adjoint if and only if it is $B^{1/2}$-self-adjoint.

The next result, which follows from the preceding lemma, provides a simple and useful sufficient condition under which a B-normal matrix is also I-normal.

COROLLARY 4.4. *If either $G = q(B)$ or $B = q^2(G)$, where q is a polynomial with real coefficients, then $\kappa_B(G) = \kappa(G)$.*

PROOF. If either hypothesis holds, then $B^{1/2}G = GB^{1/2}$. The result then follows from Lemma 4.3. □

Although the hypotheses of these results may seem opaque, one or all are satisfied for some of the familiar conjugate gradient methods. We will revisit this point when we discuss stopping criteria for particular CG methods.

Finally, note that while the Euclidean and B-norms of the relative error are not, in general, equal, they are equivalent. Precise bounds are given in the following lemma.

LEMMA 4.5. *Let B be hpd. Then the relative errors in the B- and Euclidean norms satisfy:*

$$\kappa^{-1/2}(B)\frac{\|e_i\|_B}{\|x\|_B} \leq \frac{\|e_i\|}{\|x\|} \leq \kappa^{1/2}(B)\frac{\|e_i\|_B}{\|x\|_B}. \tag{4.1}$$

PROOF. For any y, $\|y\|_B = \langle By, y\rangle^{1/2} \leq \|B\|^{1/2}\|y\|$. Similarly,

$$\|y\| = \langle y, y\rangle^{1/2} = \langle B^{-1}B^{1/2}y, B^{1/2}y\rangle^{1/2} \leq \|B^{-1}\|^{1/2}\|y\|_B.$$

Combining the two inequalities appropriately for $y = x$ and $y = e_i$ yields (4.1). □

Note. For many important problems, $1 \ll \kappa(B)$; for example, if A results from the solution of a PDE and $B = A^*A$, then $\kappa(B) = O\left(1/h^4\right), \quad h \ll 1$. If this is the case, then a CG algorithm based on B may not be the best choice of an iterative method when the user wants to make $\|e_i\|/\|b_i\|$ small.

5 A natural stopping criterion when BCA is definite

We will now derive a stopping criterion that uses only previously computed iteration parameters. Here we assume that BCA is definite (i.e., CA is B-definite), in which

case the Omin algorithm (Fig. 2.1) is applicable. The key quantity is the numerator of $\widehat{\alpha}_i$, the current steplength. This is given by

$$\langle Be_i, \widehat{p}_i \rangle \tag{5.1}$$

where $\widehat{p}_i$ is the current direction vector. (The $\widehat{p}_j$, $j = 0, \ldots, i$, form a B-orthogonal basis for K_{i+1}.) Since the error e_i is unknown, one must be careful to choose B and C so that this quantity is computable, that is, expressible in terms of known quantities. We regard r_i as the basic known quantity and consider A and b to be the basic irreducible entities. One may show that (5.1) is computable whenever C^*Be_i is computable [6]. For instance, it is a standard CG equality [6, Cor. 4.1] that

$$\langle Be_i, \widehat{p}_i \rangle = \langle C^*Be_i, r_i \rangle. \tag{5.2}$$

This quantity, which is computed during the course of the Omin iteration, is used in the following theorem to establish bounds that are the basis for our "natural" CG stopping criterion.

THEOREM 5.1. *If CA is B-normal and B-definite then*[2]

$$\left(\phi_B^{-1}(CA) \left| \frac{\langle C^*Be_i, r_i \rangle}{\langle C^*Bx, b \rangle} \right| \right)^{1/2} \;\leq\; \frac{\|e_i\|_B}{\|x\|_B} \;\leq\; \left(\phi_B(CA) \left| \frac{\langle C^*Be_i, r_i \rangle}{\langle C^*Bx, b \rangle} \right| \right)^{1/2}, \tag{5.3}$$

where

$$\phi_B(CA) = \frac{\rho(CA)}{\psi_B(CA)}, \qquad \textit{and} \qquad \psi_B(CA) = \min_{\lambda \in F_B(CA)} |\lambda|.$$

PROOF. Although the proof follows directly from the definition of $F_B(CA)$, we give the details because they are important in the sequel. Since CA is B-definite (i.e., $\langle CAv, Bv \rangle \neq 0$ if $v \neq 0$), we have for any nonzero z,

$$0 < \min_{\lambda \in F_B(CA)} |\lambda| \leq \left| \frac{\langle BCAz, z \rangle}{\langle Bz, z \rangle} \right| \leq \max_{\lambda \in F_B(CA)} |\lambda|. \tag{5.4}$$

(Unless CA is B-definite, the lower bound is zero.) Since CA is B-normal, we have $F_B(CA) = H(CA)$, where again $H(CA)$ denotes the convex hull of CA. Thus, $\max_{\lambda \in F_B(CA)} |\lambda| = \rho(CA)$. Then, from the definition of $\psi_B(CA)$ and the fact $|\langle BCAz, z \rangle| = |\langle Bz, CAz \rangle|$, we have

$$\frac{|\langle Be_i, CAe_i \rangle|}{\rho(CA)} \leq \langle Be_i, e_i \rangle \leq \frac{|\langle Be_i, CAe_i \rangle|}{\psi_B(CA)}$$

and

$$\frac{|\langle Bx, CAx \rangle|}{\rho(CA)} \leq \langle Bx, x \rangle \leq \frac{|\langle Bx, CAx \rangle|}{\psi_B(CA)}$$

for $z = e_i$ and $z = x$, respectively. Combining these two inequalities gives

$$\frac{\psi_B(CA)}{\rho(CA)} \left| \frac{\langle C^*Be_i, r_i \rangle}{\langle C^*Bx, b \rangle} \right| \leq \frac{\langle Be_i, e_i \rangle}{\langle Bx, x \rangle} \leq \frac{\rho(CA)}{\psi_B(CA)} \left| \frac{\langle C^*Be_i, r_i \rangle}{\langle C^*Bx, b \rangle} \right|.$$

The result now follows from the definition of $\phi_B(CA)$. □

[2] Although we are usually interested in B-normal(1) matrices—for which three-term CG methods are possible—many of our results hold for the more general B-normal case.

(a) (b)

Figure 5.1: In (a), Corollary 5.2 fails: For K a real Hermitian indefinite matrix, the convex hull of eigenvalues of $CA = I - iK$ is a vertical line segment through $z = 1$ in the complex plane. The point on this line of smallest absolute value is neither extreme eigenvalue but $z = 1$. Thus $\psi_B(CA) = \min_{\lambda \in F_B(CA)} |\lambda|$ is attained at an interior point. In (a), the Corollary holds if 1 is an eigenvalue. In (b), the Corollary holds: The upper horizontal line segment is the convex hull of eigenvalues of a matrix of the form $CA = iI + H$ where H is hpd. The point of smallest absolute value is closest to the origin. Also the Corollary holds for $e^{i\theta}CA$. The convex hull of $e^{i\theta}CA$ is shown for $\theta = -\pi/4$.

In the important special case of a B-self-adjoint matrix, the last theorem simplifies.

COROLLARY 5.2. *If CA is B-self-adjoint and B-definite then*

$$(5.5)\quad \left(\kappa_B^{-1}(CA)\left|\frac{\langle C^*Be_i, r_i\rangle}{\langle C^*Bx, b\rangle}\right|\right)^{1/2} \;\le\; \frac{\|e_i\|_B}{\|x\|_B} \;\le\; \left(\kappa_B(CA)\left|\frac{\langle C^*Be_i, r_i\rangle}{\langle C^*Bx, b\rangle}\right|\right)^{1/2}.$$

PROOF. The assumptions imply that the eigenvalues of CA are real and either all positive or all negative. (Recall that CA is B-self-adjoint if and only if BCA is Hermitian.) Therefore, $\psi_B(CA) = \min_{\lambda \in F_B(CA)} |\lambda| = \min_{\lambda \in \sigma(CA)} |\lambda|$, where $\sigma(CA)$ denotes the spectrum of CA. This implies that $\phi_B(CA) = \max_{\lambda \in \sigma(CA)} |\lambda| / \min_{\lambda \in \sigma(CA)} |\lambda|$, and the result follows from Corollary 4.2. □

Remark: Bound (5.5) for the case $B = A$ appears in [23]. The Corollary does not hold for B-definite B-normal(1) matrices in general. Although such matrices have collinear eigenvalues, the line segment connecting the extreme eigenvalues (i.e., the convex hull) is such that $\psi_B(CA)$ may be attained at an interior point. Here is a simple example: $CA = I - iK$, where K is a real Hermitian indefinite matrix. However, the Corollary does hold for matrices of the form $CA = e^{i\theta}(irI + H)$, where H is hpd; see Fig. 5.1. One final observation: It might fortuitously happen that the minimum occurs at an interior point. If so, then this interior point is also an eigenvalue, i.e., $|\lambda_i| = \psi_B(CA)$, where λ_i is an eigenvalue. Also in this case, the Corollary holds. For example, if in (a) of Fig. (5.1), 1 is an eigenvalue, then the Corollary holds.

The bound (5.5) involves the B-condition number of CA rather than the usual spectral condition number. Nonetheless, this is the "correct" quantity to use in at least two respects. First, it is defined by the ratio of the largest to smallest eigenvalue of CA, which is easily estimated via the CG iteration parameters (§7). Second, since $\mathrm{CG}(B, C, A)$ minimizes the B-norm of the error, it is appropriate to have a bound based on this norm. Nevertheless, if one prefers the I-norm, the following bound from

(4.1) may be used if an estimate for $\kappa(B)$ is available:

$$\frac{\|e_i\|}{\|x\|} \;\le\; \kappa^{1/2}(B)\frac{\|e_i\|_B}{\|x\|_B} \le \left(\kappa(B)\kappa_B(CA)\left|\frac{\langle C^*Be_i, r_i\rangle}{\langle C^*Bx, b\rangle}\right|\right)^{1/2}. \tag{5.6}$$

This bound is different from (5.5) by a factor of $\kappa^{1/2}(B)$. The advantage of this bound over (3.5) is that it does not require the norm of the residual, which may not be readily available. On the other hand, it does require an estimate for $\kappa(B)$. We emphasize that the bounds (5.3) and (5.5) require that BCA be definite, under which condition the Omin algorithm will not break down. We also remark that the bounds (5.3) and (5.5) are sharp for any B, C, and A. As explained earlier, we obtain our stopping criterion by halting the iteration when the upper bound is less than the user-specified tolerance.

Finally, note that C^*Bx in (5.3) or (5.5) is computable if and only if C^*Be is computable, and so the bound is always computable. In addition, the inner product $\langle C^*Bx, b\rangle$ needs to be computed just once at the beginning of the iteration. In some methods, this computation may be redundant when $x_0 = 0$ (i.e., $b = r_0$).

6 An alternative stopping criterion

In this section we present an inexpensive stopping criterion that is more general than the one given in the preceding section, but perhaps somewhat less robust, which we discuss in the next paragraph. Consider (3.4) with $S = B$:

$$\kappa_B^{-1}(CA)\frac{\|s_i\|_B}{\|Cb\|_B} \;\le\; \frac{\|e_i\|_B}{\|x\|_B} \;\le\; \kappa_B(CA)\frac{\|s_i\|_B}{\|Cb\|_B}. \tag{6.1}$$

The upper bound yields a stopping criterion for any matrix CA. As with our other criteria, it requires an estimate for $\kappa_B(CA)$, but it also requires $\|s_i\|_B$. At first glance this seems to require an additional inner product, which is the very expense we are trying to avoid. However, it so happens that we can compute this quantity from previously computed iteration parameters, as we will now show.

Suppose, first, that CA is B-normal(1) and we are using the Omin algorithm (Fig. 2.1), and that it does not breakdown (e.g., because CA is B-definite). In general, Omin is subject to breakdown and is not as robust as Odir. We have

$$\widehat{p}_i = s_i + \beta_{i-1}\widehat{p}_{i-1}.$$

Since the direction vectors are mutually B-orthogonal, we obtain

$$\|s_i\|_B^2 = \|\widehat{p}_i\|_B^2 + |\beta_{i-1}|^2\|\widehat{p}_{i-1}\|_B^2 \tag{6.2}$$

and so the B-norm of the preconditioned residual can be calculated from previously computed iteration parameters. (Recall that $\|\widehat{p}_i\|_B^2 = \langle B\widehat{p}_i, \widehat{p}_i\rangle$ is the denominator of $\widehat{\alpha}_i$.)

Next, suppose we are using the Odir algorithm because BCA is indefinite. Here the calculation of $\|s_i\|_B$ requires greater effort. One may use induction or the matrix formulation developed in [3, 4] to show

$$s_i = -\alpha_{i-1}p_i + \xi_{i-1}p_{i-1}$$

and therefore

$$\|s_i\|_B^2 = |\alpha_{i-1}|^2\|p_i\|_B^2 + |\xi_{i-1}|^2\|p_{i-1}\|_B^2 \tag{6.3}$$

where

$$\xi_{i-1} = \frac{\langle Bs_i, p_{i-1}\rangle}{\langle Bp_{i-1}, p_{i-1}\rangle} = e^{2i\theta}\frac{\langle Be_i, p_i\rangle}{\langle Bp_{i-1}, p_{i-1}\rangle}$$

which again is in terms of previously computed quantities. Recall that a B-normal(1) matrix has the form $e^{i\theta}(irI + H)$.

Finally, observe that the bounds in (6.1) are weaker than those in (5.5). The former can vary by as much as $\kappa_B^2(CA)$, whereas the latter vary by no more than $\kappa_B(CA)$.

7 Estimating $\phi_B(CA)$ and $\kappa_B(CA)$

The bounds derived in the previous sections require estimates either for $\phi_B(CA)$ or $\kappa_B(CA)$. It is possible to estimate these quantities from the CG iteration parameters by exploiting the "equivalence" of the CG and Lanczos methods. We will consider three cases: when CA is B-self-adjoint and B-definite, when CA is B-normal and B-definite, and when CA is B-normal but B-indefinite.

When CA is B-self-adjoint and B-definite, Corollary 5.2 is applicable, and we wish to estimate $\phi_B(CA) = \kappa_B(CA)$. Since $\kappa_B(CA) = \lambda_{\max}/\lambda_{\min}$, we can estimate this condition number by estimating the extreme eigenvalues of CA. We may obtain such estimates by computing the extreme eigenvalues of a small tridiagonal matrix, H_k, whose entries consist of the CG iteration parameters. For example, in the case of the Omin algorithm, we have [6]

$$H_k = \text{tridiag}\left(-\left(\frac{\widehat{\delta}_{j-1}}{\widehat{\delta}_{j-2}}\right)^{1/2}\frac{1}{\widehat{\alpha}_{j-2}},\ \text{Re}\left(\frac{1+\beta_{j-1}}{\widehat{\alpha}_{j-1}}\right),\ -\left(\frac{\widehat{\delta}_j}{\widehat{\delta}_{j-1}}\right)^{1/2}\frac{1}{\widehat{\alpha}_{j-1}}\right), \tag{7.1}$$

where $\widehat{\delta}_j = \langle B\widehat{p}_j, \widehat{p}_j\rangle$. We shall summarize some results from [25]. The eigenvalues of this matrix, η_j, are the roots of the B-orthogonal polynomials that generate the CG direction vectors. A second set of estimates, ν_j, can be obtained from the roots of the underlying residual polynomials. They are the eigenvalues of

$$G_k = \text{tridiag}\left(\frac{1-\rho_{j-1}}{\mu_{j-1}\rho_{j-1}}, \frac{1}{\mu_{j-1}}, -\frac{1}{\rho_{j-1}\mu_{j-1}}\right), \tag{7.2}$$

where

$$\mu_j = \frac{\widehat{\alpha}_j\widehat{\alpha}_{j-1}}{\widehat{\alpha}_{j-}\beta_{j-1}} \quad \text{and} \quad \rho_j = 1 + \frac{\widehat{\alpha}_j\beta_{j-1}}{\widehat{\alpha}_{j-1}}. \tag{7.3}$$

In both (7.1) and (7.2), the tridiagonal elements are from row j. Both sets of eigenvalues are guaranteed to lie in the convex hull of CA when CA is B-self-adjoint. Moreover, in either case, successive sets of eigenvalues interlace one another, and so the condition number estimates are monotonically nondecreasing. See [4, 6, 25] for details. We will use the approximation

$$\kappa_B(CA) \geq \frac{\max\{|\eta_j|, |\nu_j|\}}{\min\{|\eta_j|, |\nu_j|\}}. \tag{7.4}$$

Although the estimated condition number is less than the true value, it can still be used effectively in a stopping criterion as will be seen in §10.

When CA is B-normal(1) and B-definite (which is not needed until the end of this paragraph) but not necessarily B-self adjoint, the eigenvalues of CA lie on a line segment in the complex plane, which we denote by its endpoints: $[\lambda_1, \lambda_n]$. The roots of the B-orthogonal polynomials of degree k will be the eigenvalues of the matrix H_k in (7.1) and will also lie in this line segment (see [25]), say $[\eta_1, \eta_k]$. The roots of

the residual polynomials do not in general lie on this line segment. We approximate $\kappa_B(CA)$ by the same bound as given in (7.4), using both sets of roots. However, when estimating $\phi_B(CA)$, we used the line segment $[\eta_1, \eta_k]$, i.e.,

$$\phi_B(CA) \approx \frac{\max\{|\eta_1|, |\eta_k|\}}{\min_{\alpha\in[0,1]}\{|\alpha\eta_1 + (1-\alpha)\eta_k|\}}.$$

The denominator is the distance from $[\eta_1, \eta_k]$ to the origin, which is nonzero since CA is B-definite.

The third case, when CA is B-normal(1) but B-indefinite, is the most difficult. Since Theorem 5.1 is not applicable, we will use the bound (6.1) as the basis for our stopping criterion. (Note that the quantity $\phi_B(CA)$ is undefined because its denominator is zero.) As described earlier, one can compute $\|s_i\|_B$ without expending any unnecessary inner products, but we still must estimate $\kappa_B(CA)$. This can be done using the roots of the orthogonal and residual polynomials, in much the same fashion as before. Specifically, we will use the approximation

$$\kappa_B(CA) \geq \frac{\max\{|\eta_j|\}}{\min\{|\nu_j|\}}. \tag{7.5}$$

It is shown in [25] that in this case an η_i may be zero. No ν_j may ever equal zero because the ν_j's are the roots of a residual polynomial: A residual polynomial is, by definition, a polynomial that is 1 at the origin, i.e., 0 is not a root; also again see [25]. When CA is B-normal(1) and B-indefinite, $H(CA)$ is a line segment passing through the origin, but we can still use (7.5) to estimate $\kappa_B(CA)$.

Remark: In each of the cases above, we underestimate $\kappa_B(CA)$. Nonetheless, it is our numerical experience that these estimates are fairly accurate, and so we recommend their use in practical stopping criteria. However, if the tolerance ϵ is relatively large, the iteration may be halted too soon because it will not have had sufficient time to find reasonable eigenvalue estimates. This is a consequence of the CG–Lanczos connection. We shall discuss this further in §10.

7.1 When to update the estimates

The simplest course is to update the estimate for $\phi_B(CA)$ or $\kappa_B(CA)$ at every step, but this is unnecessary. It should be updated when the stopping criterion using the current approximation is satisfied. For example, if (5.5) is used, the approximation for $\kappa_B(CA)$ need be updated only when the right-hand side of (5.5) is less than the tolerance ϵ. If this inequality is still satisfied using the new approximation for the B-condition number, then the iteration may be stopped. (Recall that the interlacing property of the Lanczos estimates guarantees that successive condition number estimates will be nondecreasing.)

It is also possible to have a *compound* stopping criterion. For example, if CA is B-normal(1) and definite, then either (5.3) or (6.1) can be used to cheaply bound the relative error in the B-norm. Instead of choosing one over the other, both can checked. If either is satisfied, we may consider stopping. Specifically, we would update the estimate for $\phi_B(CA)$ and recheck the criterion. If it is still satisfied, the iteration may be halted. The advantage of this approach is that it may permit us to halt sooner than would a simple stopping criterion, thereby saving the cost of unwarranted iterations.

Finally, we wish to emphasize that the extreme eigenvalues of the small tridiagonal matrices yield an estimate for $\phi_B(CA)$, but *not* necessarily for its Euclidean counterpart. Since the B-norm is less than intuitive in many cases, one would prefer to

interpret the results in some equivalent norm that is more readily understood. The following results are helpful in this respect and will be used in the next section.

LEMMA 7.1. *If G is both B_1-normal and B_2-normal, then $\kappa_{B_1}(G) = \kappa_{B_2}(G)$.*

PROOF. The result follows immediately from Lemma 4.1. □

LEMMA 7.2. *Let F be hpd. Then G is G^*FG-normal if and only if G is F-normal.*

PROOF. Recall that G is G^*FG-normal if and only if it has a complete set of G^*FG-orthogonal eigenvectors, $\{v_i\}$. That is, $\langle G^*FGv_i, v_j\rangle = 0$ for $i \neq j$. Since

$$\langle G^*FGv_i, v_j\rangle = \langle FGv_i, Gv_j\rangle = \lambda_i|\lambda_j|\langle Fv_i, v_j\rangle$$

G also has a complete set of F-orthogonal eigenvectors, and so G is F-normal. The reverse direction is similarly proved. □

8 General Remarks and Other Approaches

In our approach to stopping criteria, we emphasize the B-norm since CG minimizes the error with respect to this, but there are many other norms (see (3.3, 3.4)) and criteria. We shall survey some of these.

Many criteria are based on aspects of an engineering or scientific problem. As an example, one scientific group may stop an iteration when $\|D^{-1}Ae_k\|_2$ is small, where D is the diagonal of A because for certain applications it has more physical meaning than when the residual is used. Another group may use the ℓ_1 norm, normalized by dividing by the number of unknowns because it is thought to work better when maintaining material balances. Often several criteria are watched, with the max norm (ℓ_∞), normalized ℓ_1 and ℓ_2 norms applied to the residual vector. With large engineering budgets dependent on accuracy, a numerical approximation is accepted only when all criteria are satisfied, an example of a compound stopping criterion.

There is some preference for the max norm on the principle that it is the value of at least one component whereas the ℓ_2 norm in averaging over all components blurs individual component values. The view that individual components are important finds expression in the work of Arioli, Duff and Ruiz [2] which starts from the backward error analysis of Oettli and Prager [26], also Skeel [27, 28], on componentwise backward error analysis. Arioli, Duff and Ruiz recommend stopping when

$$\|r_k\|_\infty/(\|A\|_\infty\|x_k\|_1 + \|b\|_\infty) \leq \epsilon.$$

Their views are strongly stated: "[We recommend] ... this stopping criteria be used as the standard one in the context of the current efforts to standardize the interface and environment for iterative solvers." They also recommend terminating the iteration when this bound "increases or oscillates significantly." Unfortunately, the max norm may be difficult to estimate either if the system is preconditioned or if an explicit matrix is not available.

In their treatise on iterative methods, Hageman and Young use B-norms ($B = I$ in (8.1) below) in their discussions of stopping criteria [18]. They recommend a criterion equivalent to stopping when

$$\langle Be_i, CAe_i\rangle/\lambda_1 \leq \epsilon\|x_i\| \tag{8.1}$$

where λ_1 is the smallest eigenvalue of CA and ϵ is the user tolerance [18, p. 149]. (We have simplified things: The authors use a symmetrizing matrix W to transform one system to another, equivalent to a preconditioning transformation. The results are

extended in [13] to a skew-symmetric matrix Z.) In order to estimate λ_1, they use a tridiagonal matrix in the same way we use H_k in §7.

To compare this criterion to ours, we shall modify an inequality of Hageman and Young [18, (7-5.2) p. 148] to get

$$\frac{\|e_i\|_B}{\|x\|_B} \leq \frac{(\langle Be_i, CAe_i\rangle/\lambda_1)^{1/2}}{\|x_i\|_B}. \tag{8.2}$$

(Also see [24, 32].) For the convenience of the reader, we state (a part of) (5.5) here for a side-by-side comparison with (8.2):

$$\frac{\|e_i\|_B}{\|x\|_B} \;\leq\; \left(\kappa_B(CA)\frac{|\langle C^*Be_i, r_i\rangle|}{|\langle C^*Bx, b\rangle|}\right)^{1/2}. \tag{8.3}$$

Inequality (8.2) improves (8.3) through using only λ_1 instead of the larger value $\kappa_B(CA) = \lambda_N/\lambda_1$. Estimating $\|x\|_B$ by $\|x_i\|_B$ in the denominator of (8.3) is the key to dispensing with the larger eigenvalue λ_N. Although this does not in general preserve the inequality, the approximation is accurate in the pivotal decision period when x_i is close to x. Moreover, if $x_0 = 0$, the inequality is in fact preserved for the B-norm since $\|x_i\|_B \leq \|x\|_B$ in this case. If $x_0 \neq 0$, then b can be replaced by $b - Ax_0$ to obtain an equivalent system for which the initial guess is zero, at the cost of an extra matrix vector multiplication. Note that in the case of the B-norm as with other norms, approximating λ_1 with a larger quantity, which occurs with the Lanczos algorithm, still affects the inequality.

Experiments show (later below in §11 we report on this) that the bound in (8.2) floats well below the bound given by (the exact value of) $\kappa_B(CA)$ in (8.3). Thus, there is an improvement but at a cost, this being the computation of the B-norm of the current iterate. In general the gain in accuracy does not justify a criterion based on the right side of (8.3), if $\|x_i\|_B$ is used at every step: the savings in stopping early is less than the extra cost. However, if the matrix is large and dense (cases that do arise), then the extra cost would be negligible compared to the savings in stopping early.

The difficulties due to the possibility that $\|x_i\|_B > \|x\|_B$ are handled in an elegant way in a paper of Kaaschieter [22]. (The analysis is also valid for norms not induced by an inner product, such as the max norm.) He looks at the effect not only of approximating $\|x\|_B$ but also of approximating λ_1. To summarize, we shall consider only the effect of approximating $\|x\|_B$. Kaaschieter's analysis yields

$$\frac{\|e_i\|_B}{\|x\|_B} \leq \frac{\delta}{1-\delta}$$

where

$$\delta = \frac{(\langle Be_i, CAe_i\rangle/\lambda_1)^{1/2}}{\|x_i\|_B}.$$

This result shows that the factor $1/(1-\delta)$ is just the factor by which the right side of (8.2) must be increased for the inequality to hold.

Next we come to a different approach to estimating the error, as proposed in [10, 11, 15]. The common idea in these papers is to use the orthogonal polynomials generated by the conjugate gradient method in a Gaussian quadrature formula to approximate the inner product $\langle e_k, Be_k\rangle$, which may be viewed as a quadrature with respect to a weight function defined by the initial error. A difficulty with this approach is in estimating the initial error; see [15] for an interesting treatment based on continued

fractions and their relation to orthogonal polynomials. For another approach based on continued fractions see [30]. Many interesting ideas may also be found in [16, 17].

Finally on the topic of upper bounds we mention that inequality (3.3) appears also in a paper of Vuik on GMRES applied to the Navier-Stokes equation [29]. The author recommends this inequality with an estimate for $\kappa_M(A)$.

Although we mainly discuss upper bounds, (5.5) includes a lower bound. One valuable property of a good lower bound is of course in showing how accurate the upper bound is by providing an interval in which the (relative) error lies. There is a useful suggestion due to Auchmuty [7] that improves (5.5). It is

$$\|r_k\|_B^2/\|A^*r_k\|_B \le \|e_k\|_B \tag{8.4}$$

where A^* is the Hermitian transpose of A. The inequality results from

$$\langle r_k, Ae_k\rangle_B = \langle A^*r_k, e_k\rangle_B \le \|A^*r_k\|_B \ \|e_k\|_B$$

which follows by Schwarz's inequality. Experiments confirm that Auchmuty's lower bound is an improvement albeit at significant extra cost.

9 Stopping criteria for twelve CG methods

In this section we instantiate the results of §5 for the twelve CG methods given in Table 2.1. We assume that the Odir restrictions on C and A listed in the table are satisfied. It is easy to verify that, under these conditions, CA is B-self-adjoint for each of the methods. Initially, we will assume also that the stricter Omin restrictions are met, in which case BCA is hpd and Corollary 5.2 is applicable. (We defer the BCA indefinite case to the end of the section.)

9.1 The BCA definite case

The primary stopping criterion for each method is then

$$\frac{\|e_i\|_B}{\|x\|_B} \ \le \ \left(\kappa_B(CA)\frac{|\langle C^*Be_i, r_i\rangle|}{|\langle C^*Bx, b\rangle|}\right)^{1/2} \ \le \ \epsilon \tag{9.1}$$

from (5.5) of Corollary 5.2. This criterion, which bounds the B-norm of the relative error, is always inexpensive to implement because it uses a previously computed inner product, namely, the numerator of $\widehat{\alpha}_i$. (Recall that C^*Bx is computable whenever C^*Be_i is computable.) Since B varies with the method, so will the stopping criterion. We will point out where the B-condition number of the matrix (necessary for the stopping criterion) coincides with the more familiar I-condition number. Whenever feasible, we will also give the stopping criterion (5.6), which is based on the Euclidean norm of the relative error. Of course, in some cases, this criterion will coincide with (3.5) or its preconditioned counterpart. We emphasize that the quantities appearing as the bounding expressions in the following inequalities are inexpensively computed as stopping criteria, despite their complex appearance. Specifically, in each of the methods below, the numerator of $\widehat{\alpha}_i$ is the key quantity in the upper bound that comprises the stopping criterion. We also wish to point out that many of these results generalize to the more general B-normal case when full recursions (i.e., Orthomin) are used. Finally, we remind the reader that the alternative criterion (6.1) is also available, but since it is easily interpreted, there is no need to instantiate it for each of the methods.

CGHS. Here $B = A$ and $C = I$ and so BCA is hpd. By Corollary 4.4, we have $\kappa_B(A) = \kappa(A)$, and so (9.1) becomes

$$\frac{\|e_i\|_A}{\|x\|_A} \;\leq\; \kappa^{1/2}(A)\frac{\|r_i\|}{\|b\|} \;\leq\; \epsilon. \tag{9.2}$$

If one prefers the Euclidean norm, (3.5) may be used.

PCG. As with CGHS, $B = A$, but now $C \neq I$, and so in general, $\kappa_B(CA) \neq \kappa(CA)$. (See PPCG for the special case of polynomial preconditioned CG.) If C is hpd, as is usually the case, BCA is hpd, and (9.1) becomes

$$\frac{\|e_i\|_A}{\|x\|_A} \;\leq\; \kappa_A^{1/2}(CA)\frac{\|r_i\|_C}{\|b\|_C} \;\leq\; \epsilon. \tag{9.3}$$

Unfortunately, since $\kappa(A)$ is not generally available, we cannot use (5.6) or (3.5) to obtain an inexpensive bound on the relative error in the Euclidean norm. Note that it is an option to compute the A-norm of the residual occasionally but that the estimate is a way to avoid this expensive step.

PPCG. In this variant of PCG ($B = A$), we employ a polynomial preconditioner, $C = C(A)$, where $C(\lambda)$ is a polynomial with real coefficients. In this special case, $\kappa_B(C(A)A) = \kappa(C(A)A)$, from Corollary 4.4. Moreover, it is possible to estimate both $\kappa(A)$ and $\kappa(C(A)A)$. If the preconditioning polynomial is chosen appropriately, then $C(A)$ is hpd—even when the polynomial is dynamically determined [5]. Thus, BCA is hpd and the stopping criteria

$$\frac{\|e_i\|_A}{\|x\|_A} \;\leq\; \kappa^{1/2}(C(A)A)\frac{\|r_i\|_C}{\|b\|_C} \;\leq\; \epsilon \tag{9.4}$$

and

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa^{1/2}(A)\kappa^{1/2}(C(A)A)\frac{\|r_i\|_C}{\|b\|_C} \;\leq\; \epsilon \tag{9.5}$$

are both cheap to implement.

PCGCA. In this polynomial preconditioned CG method we have $B = C(A)A$ and $C = C(A)$, and so again the B- and I-condition numbers of $C(A)A$ are the same. The stopping criteria are

$$\frac{\|e_i\|_{C(A)A}}{\|x\|_{C(A)A}} \;\leq\; \kappa^{1/2}(C(A)A)\frac{\|C(A)r_i\|}{\|C(A)b\|} \;\leq\; \epsilon \tag{9.6}$$

and

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa(C(A)A)\frac{\|C(A)r_i\|}{\|C(A)b\|} \;\leq\; \epsilon. \tag{9.7}$$

Observe that if $C(A)$ is a good preconditioner, then $C(A)A \approx I$, and the two criteria are nearly the same.

CR. Here $B = A^*A$ and $C = I$. If A is hpd, $BCA = A^3$ is hpd. Once again $\kappa_B(A) = \kappa(A)$, and so (9.1) becomes

$$\frac{\|r_i\|}{\|b\|} \;\leq\; \kappa^{1/2}(A)\frac{\|r_i\|_A}{\|b\|_A} \;\leq\; \epsilon. \tag{9.8}$$

Observe that the A^2-norm of the relative error is the same as the I-norm of the relative residual for Hermitian A. Here we see that unlike the previous cases, the B-norm of the error is computable. However, even if possible to compute, it is still an extra

computation, as we pointed out in our discussion of (3.2). If a bound on the relative error in the Euclidean norm is needed, one may use the stopping criterion

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa^{3/2}(A)\frac{\|r_i\|_A}{\|b\|_A} \;\leq\; \epsilon \tag{9.9}$$

derived from (9.1) and (5.6). Alternatively, one may use (3.5), but this requires an extra inner product per stopping criterion evaluation to compute the Euclidean norm of the residual.

PCR. In this preconditioned variant of CR we have $B = ACA$, which is hpd if and only if the preconditioner C is hpd. If, in addition, A is definite, then BCA is definite (i.e., CA is B-definite) and (9.1) is applicable. It gives the stopping criterion

$$\frac{\|r_i\|_C}{\|b\|_C} \;\leq\; \kappa_{ACA}^{1/2}(CA)\frac{\|Cr_i\|_A}{\|Cb\|_A} \;\leq\; \epsilon. \tag{9.10}$$

Since $\kappa(A)$ and $\kappa(ACA)$ are not generally available, we cannot use (3.5) or (5.6) to obtain an inexpensive bound on the relative error in the Euclidean norm. Note the similarity with PCG in this respect: the use of a preconditioner precludes the use of these simple stopping criteria.

CGNR. In this variant of the normal equations, $B = A^*A$ and $C = A^*$, and so $\kappa_B(A^*A) = \kappa(A^*A) = \kappa^2(A)$. The inequalities (9.1) and (5.6) become

$$\frac{\|r_i\|}{\|b\|} \;\leq\; \kappa(A)\frac{\|A^*r_i\|}{\|A^*b\|} \;\leq\; \epsilon \tag{9.11}$$

and

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa^2(A)\frac{\|A^*r_i\|}{\|A^*b\|} \;\leq\; \epsilon \tag{9.12}$$

respectively.

CGNE. This version of the normal equations, which is also known as Craig's method [9], minimizes the Euclidean norm of the error at each step of the iteration. That is, $B = I$ and $C = A^*$. Here the inequalities (9.1) and (5.6) coincide:

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa(A)\frac{\|r_i\|}{\|b\|} \;\leq\; \epsilon. \tag{9.13}$$

PCGNR. Here $B = A^*A$ and $C = (M^*M)^{-1}A^*$. In general, the B- and I-condition numbers of CA are not the same. The inequality (9.1) gives

$$\frac{\|r_i\|}{\|b\|} \;\leq\; \kappa_{A^*A}^{1/2}((M^*M)^{-1}A^*A)\frac{\|(AM^{-1})^*r_i\|}{\|(AM^{-1})^*b\|} \;\leq\; \epsilon. \tag{9.14}$$

The use of inequality (5.6) is practicable when the condition number of A^*A is known.

PCGNM. Here $B = M^*M$ and $C = (M^*M)^{-1}A^*$. In general, $\kappa_B(CA) \neq \kappa(CA)$. The inequality (9.1) gives

$$\frac{\|Me_i\|}{\|Mx\|} \;\leq\; \kappa_{M^*M}^{1/2}((M^*M)^{-1}A^*A)\frac{\|r_i\|}{\|b\|} \;\leq\; \epsilon. \tag{9.15}$$

If $M \approx A$, this stopping criterion halts when the approximate relative residual is small enough. The use of inequality (5.6) is practicable when $\kappa(M^*M)$ is unknown.

PCGNE. Here $B = I$ and $C = (M^*M)^{-1}A$, and so (9.1) yields

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa(A^*M^{-1})\frac{\|M^{-*}r_i\|}{\|M^{-*}b\|} \;\leq\; \epsilon. \tag{9.16}$$

PCGNS. Here $B = (M^{-1}A)^*M^{-1}A$ and $C = (M^{-1}A)^*M^{-1}$. By Corollary 4.4, we have $\kappa_B(CA) = \kappa(CA)$. Thus, (9.1) gives

$$\frac{\|M^{-1}r_i\|}{\|M^{-1}b\|} \;\leq\; \kappa(M^{-1}A)\frac{\|(M^{-1}A)^*M^{-1}r_i\|}{\|(M^{-1}A)^*M^{-1}b\|} \;\leq\; \epsilon. \tag{9.17}$$

Since $\kappa(B) = \kappa^2(M^{-1}A)$, we also have

$$\frac{\|e_i\|}{\|x\|} \;\leq\; \kappa^2(M^{-1}A)\frac{\|(M^{-1}A)^*M^{-1}r_i\|}{\|(M^{-1}A)^*M^{-1}b\|} \;\leq\; \epsilon. \tag{9.18}$$

Observe that if $M \approx A$, then $M^{-1}A \approx I$ and $M^{-1}r_i \approx e_i$.

9.2 The BCA indefinite case

When BCA is indefinite, one must use the Odir or an Odir/Omin hybrid algorithm for a robust implementation of a CG method. An indefinite BCA is allowed in four of the twelve methods in Table 2.1: PCG, PPCG, CR, and PCR. However, Corollary 5.2 no longer holds because $\psi_B(CA) = 0$. Consequently, the stopping criterion (9.1) is unavailable and one must use the alternative (6.1) or a more expensive stopping criterion. An indefinite BCA is rare in the case of PCG and PPCG (because one usually has an hpd preconditioner), but is encountered in CR and PCR when the matrix A is Hermitian indefinite.

To use (6.1), one must be able to compute $\|s_i\|_B$ and to estimate $\kappa_B(CA)$. In §6, we showed how one can use formula (6.3) to compute $\|s_i\|_B$ without expending additional inner products. In §7, we described how one can estimate $\kappa_B(CA)$ by computing the roots of the orthogonal and residual polynomials; see equation (7.5). Recall that these roots are the eigenvalues of two small tridiagonal matrices, the entries of which are derived from the CG iteration parameters.

10 Numerical Experiments

All experiments were run using MATLAB 4.2a. The initial vector was the zero vector. The true solution was computed using the MATLAB command `rand`, such that each usage begins with the same seed, and the right hand side computed from a matrix vector multiplication. The initial guess is taken to be zero. We shall describe our test matrices and experiments next.

In our experiments two sets of matrices begin either with a matrix generated from the centered difference approximation to

$$-(e^{-xy}u_x)_x - (e^{xy}u_y)_y + \delta[(x+y)u_y + ((x+y)u)_y] + [1/(1+x+y)]u$$

or with a complex symmetric matrix resulting from the finite element approximation to Maxwell's equation in 3D. The latter is a 125×125 matrix. The first case is an example due to Elman [12], and the second arises in the solution of a scattering problem. Also we used a diagonal matrix obtained as follows. We start with an $n \times n$ diagonal matrix D for which the diagonal elements are the first n positive integers. Then we let $A = D^{2.5}$ and use A for a test matrix. (The exponent disperses the uniform spectrum of D.)

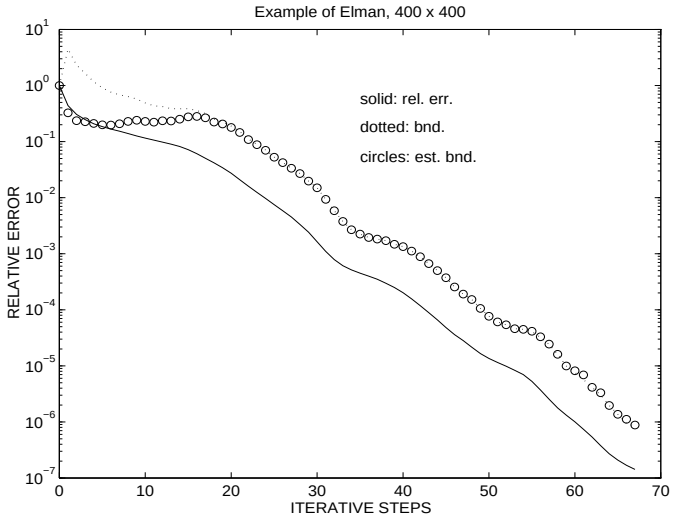

Figure 10.1: Estimated upper bound for the relative error in the Elman example. The estimated bound converges to the true (upper) bound.

We shall use the classical conjugate gradient method (CGHS) to solve an hpd system. The matrices in either Elman's example or the scattering problem are not hpd. In both cases, we arrive at an hpd matrix by the practical means of solving a system for which the matrix is the Hermitian part of the system matrix. (If M is a matrix, the *Hermitian part* of M is defined to be $(M + M^*)/2$.) In both cases, the Hermitian part is positive definite. The Hermitian part could then be used as a preconditioner for the non-Hermitian system. Finally, in the case of the scattering problem, we obtain a third test matrix by forming the normal equations, $A^*Ax = A^*b$.

Our stopping criterion is to halt when the right inequality in (9.1) is satisfied. In the case of CGHS, when $C = I, B = A$ and $Be_i = r_i$, this criterion reduces to (stopping when):

$$\kappa(A)^{1/2} \frac{\|r_i\|}{\|b\|} \leq \epsilon. \tag{10.1}$$

We shall refer to the left side of (10.1) as the *(relative error) true (upper) bound.*

We estimate the condition number using the Lanczos algorithm, which yields an estimate $\widetilde{\kappa}(A)$ that is less that the true value of the condition number: $\widetilde{\kappa}(A) < \kappa(A)$. In turn the approximate value yields the criterion (to stop when)

$$\widetilde{\kappa}(A)^{1/2} \frac{\|\|r_i\|}{\|b\|} \leq \epsilon \tag{10.2}$$

We shall refer to the left side of (10.2) as the *(relative error) estimated (upper) bound.*

One concern is whether the relative error estimated bound is in fact an upper bound, i.e., whether

$$\frac{\|e_i\|_B}{\|x\|_B} \leq \widetilde{\kappa}(A)^{1/2} \frac{\|\|r_i\|}{\|b\|} \quad (\leq \epsilon) \tag{10.3}$$

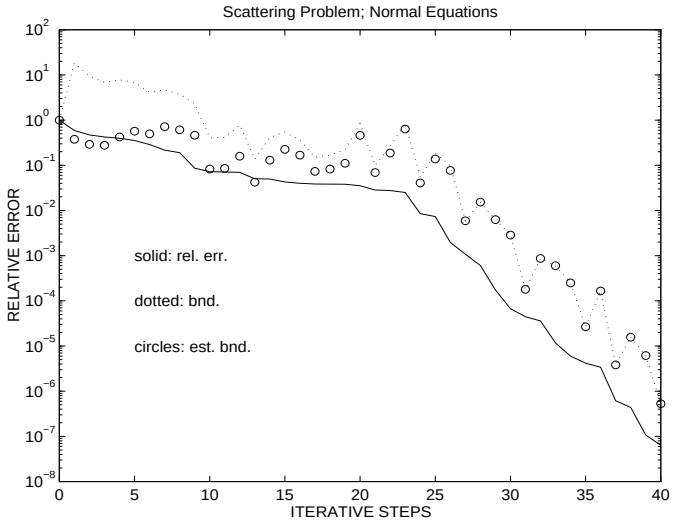

Figure 10.2: Convergence of the estimated upper bound to the true upper bound. The system is the normal equations for the solution of a system with the matrix from a scattering problem.

is true. We shall look at this question as well as the general question of whether the estimated upper bound is too pessimistic.

In Fig. 10.1, with the Elman matrix, we see that when the iteration starts, the early Lanczos approximation causes the estimated bound to be slightly less than the relative error. In the later stages of the iteration, when the approximation may be satisfactory and a stopping criteria ready to apply, the Lanczos approximation to $\kappa^{1/2}(A)$ is so accurate that the estimated bound is the same as the true bound. In experiments for which the matrix arises from taking the Hermitian part of the scattering problem, the same results hold as for the Elman matrix.

In Fig. 10.2, the system matrix is A^*A where A results from the scattering problem. We see the same behavior of the estimated bound converging to the true bound in a region where a stopping criterion would be applied. However, in the earlier stages, the estimated upper bound fails to be an upper bound. In our judgment, despite this, the estimated (upper) bound yields a practical stopping criterion. This situation, here and with the Elman matrix, has been commented upon elsewhere. Kaaschieter observed (with respect to a condition guaranteeing that (8.2) holds for CGHS) "... it can be concluded that the termination criterion ... is very reliable, provided that the conjugate gradient process is not stopped in a too early phase ..." [22, p. 321].

11 Summary and Future Work

We have presented two efficient stopping criteria for conjugate gradient methods. These criteria can be evaluated inexpensively using parameters already computed during the course of the iteration. Moreover, they are based on the B-norm of the relative error, which is the same norm used to define the CG optimization property. Conse-

quently, a CG algorithm employing either criterion would halt when the quantity being minimized was brought below the specified tolerance. Such a criterion would be a good default in software featuring iterative methods.

We began with a review of twelve well-known conjugate gradient methods (Table 2.1), each of which can be characterized by just three matrices: an hpd inner product matrix B, a left preconditioning matrix C, and the system matrix A. After discussing stopping criteria, we analyzed the relationship between the B- and Euclidean inner products. Some intermediate results were obtained, including relationships between the B-norm of a matrix and its spectral radius, between the B-condition number of a matrix and its extreme eigenvalues, and between the B- and Euclidean norms of the relative error. Our analysis bears on a common misconception: the ratio of the extreme eigenvalues of CA do *not* necessarily yield the Euclidean condition number of CA, but rather the B-condition number of CA.

Following this analysis, the main results were given in Theorem 5.1 and Corollary 5.2, which provide bounds on the B-norm of the relative error. The upper bound is the basis for the first proposed stopping criterion. When CA is B-normal and definite, the criterion requires an estimate for $\phi_B(CA)$, the ratio of the largest to the smallest point in the B-field of values. When CA is B-self-adjoint and B-definite, this quantity reduces to $\kappa_B(CA)$, the B-condition number of CA. Since neither quantity is usually known a priori, we described how they may be estimated dynamically from the CG iteration parameters. All other quantities needed to evaluate the bounds are already available from the CG iteration. We surveyed other approaches and theories before giving a detailed discussion on implementing our criteria. The particular form that this stopping criterion takes for each of the twelve CG methods of Table 2.1 also was discussed. For each method, the resulting bound was derived and found to be inexpensive to compute, requiring no additional vector work. In a set of numerical experiments we found that our criterion is satisfactory when Lanczos eigenvalue estimates are used.

We also presented an alternative upper bound on the B-norm of the relative error, one based on the B-norm of the preconditioned residual. At first glance this seems to require an additional norm computation, but we showed how one could express this quantity in terms of previously computed iteration parameters. We recommend that the two criteria be combined into a *compound* stopping criterion. This natural and efficient criterion would be more reliable than either of its constituent criteria taken alone.

This concludes our summary. As a final topic, we report on an observation proposed for further investigation. In principle the estimated (upper) bound for the relative error is less than the true bound but in our experience the estimated bound often converges so quickly to the true bound that the two coincide. It would be an advantage if they did not since the true upper bound is often an order of magnitude larger than the relative error. One approach to reducing the estimated upper bound is to reduce the estimate of $\kappa_B(CA)$. Recall that we estimated $\kappa_B(CA)$ by taking the maximum and minimum of the extreme eigenvalues of two tridiagonal matrices (7.1) and (7.2). We modify this estimate by extracting a lower right principal submatrix from each and use the maximum and minimum of the eigenvalues of the principal submatrices to estimate the extreme eigenvalues. The result is an estimate bounded above by the estimate obtained from the full tridiagonal matrices. (We could also use estimates from only one tridiagonal matrix; there are many possibilities.)

In Fig. 11.1, for the Elman matrix, we see that this approach yields a better estimate. The results are identical for the Hermitian matrix resulting from the scattering problem but are not shown. However, the estimate may not be greater than the true relative

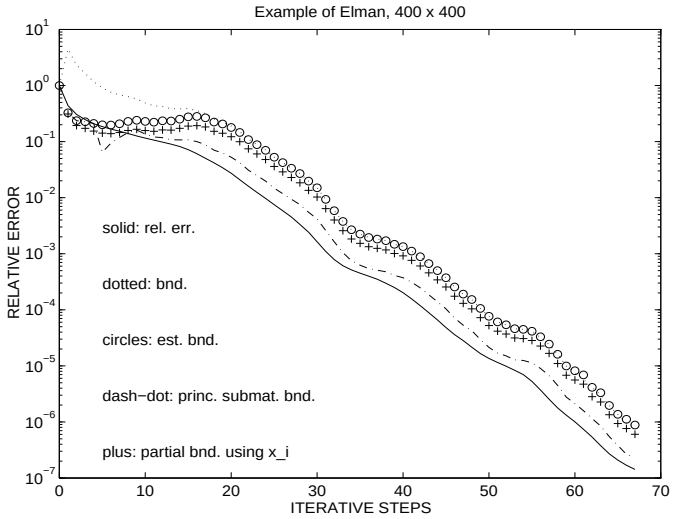

Figure 11.1: The dash-dot curve plots the estimates obtained by using 5×5 principal submatrices extracted from the lower right corner of tridiagonal matrices. The "+" 's mark the estimates due to using x_i's in (8.2).

error: In Fig. 11.2, the estimate from principal submatrices crosses below the relative error at an early stage as does the estimated upper bound. In both cases we have included the bound (8.2) for comparison. In this experiment, the Lanczos estimate is used for the value of λ_1 in (8.2). (It is labeled as a "partial" bound in the figure since the upper bound is obtained using only one extreme eigenvalue.)

Practical necessity requires giving up strict inequality in (10.3) at least in the early stages of the iteration. The even greater risk that we would lose strict inequality in the use of bounds from principal submatrices may not be acceptable to conservative users, already unhappy with the violation of inequality in (10.3). There is a heuristic remedy: We have a monotonically increasing sequence of estimates $\widetilde{\kappa}_B^{(i)}(CA)$ converging to $\kappa_B(CA)$ as the iteration number increases, with convergence of the estimates attained before a user would want to terminate the iteration. The heuristic is simply to test convergence of the condition number estimates by some means. For example, compute the first and second differences of the sequence $\{\widetilde{\kappa}_B^{(i)}(CA) : 0 \leq i\}$, then accept (10.2) only if magnitudes of these differences are sufficiently small to show convergence of the $\widetilde{\kappa}_B^{(i)}(CA)$'s.

Acknowledgments

Thanks to Joe Grcar for pointing out reference [7]. We appreciate helpful comments that Beresford Parlett made to one of the authors (Saylor). We are indebted to Anne Greenbaumn for her constructive remarks.

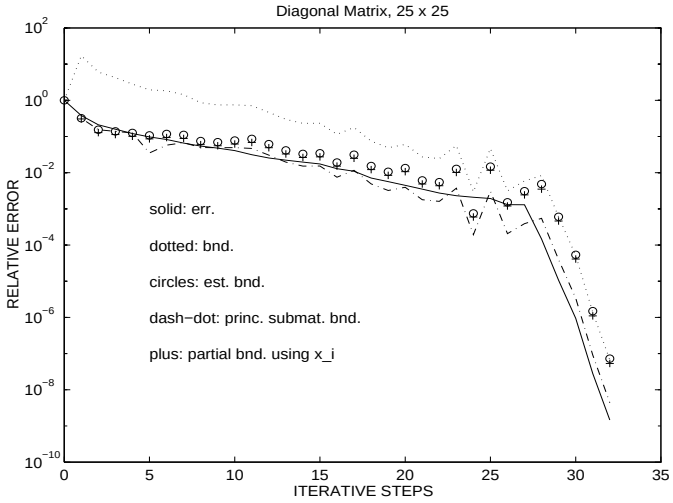

Figure 11.2: The use of principal submatrices to estimate the condition number may yield an estimate of the relative error that is less than the true relative error.

REFERENCES

1. M. Arioli, I. S. Duff, and D. Ruiz. Stopping criteria for iterative solvers. Technical Report RAL-91-057, Rutherford Appleton Laboratory, August 1991.
2. Mario Arioli, Ian Duff, and Daniel Ruiz. Stopping criteria for iterative solvers. *SIAM J. Matrix Anal. Appl.*, 13(41):138–144, 1992.
3. S. F. Ashby and M. H. Gutknecht. A matrix analysis of conjugate gradient algorithms. In M. Natori and T. Nodera, editors, *Proc. Ninth Symposium on Preconditioned Conjugate Gradients (Parallel Processing for Scientific Computing)*, pages 32–47. Keio University, 1993. Held at Keio University, Yokohama, Japan, March 2, 1993.
4. S. F. Ashby and M. H. Gutknecht. A matrix analysis of conjugate gradient algorithms. Technical Report UCRL-JC-113560, Lawrence Livermore National Laboratory, 1994.
5. S. F. Ashby, T. A. Manteuffel, and J. S. Otto. A comparison of adaptive Chebyshev and least squares polynomial preconditioning for Hermitian positive definite linear systems. *SIAM J. Sci. Statist. Comput.*, 13(1):1–29, 1992.
6. Steven F. Ashby, Thomas A. Manteuffel, and Paul. E. Saylor. A taxonomy for conjugate gradient methods. *SIAM J. Numer. Anal.*, 27(6):1542–1568, 1990.
7. Giles Auchmuty. A posteriori error-estimates for linear equations. *Numer. Math.*, 61(1):1–6, 1992.
8. P. Concus, Gene H. Golub, and D. P. O'Leary. A generalized conjugate gradient method for the numerical solution of elliptic partial differential equations. In J. R. Bunch and D. J. Rose, editors, *Sparse Matrix Computations*, pages 309–332. Academic Press, New York, 1976.

9. E. J. Craig. The N-step iteration procedures. *J. Math. Physics*, 34(1):64–73, 1955.
10. Germund Dahlquist, Stanley C. Eisenstat, and Gene H. Golub. Bounds for the error of linear systems of equations using the theory of moments. *J. Math. Anal. Applic.*, 37:151–166, 1972.
11. Germund Dahlquist, Gene H. Golub, and S. Nash. Bounds for the error in linear systems. In R. Hettich, editor, *Proceedings of the workshop on semi-infinite programming*, pages 154–172, New York, 1978. Springer.
12. H. C. Elman. *Iterative Methods for Large, Sparse Nonsymmetric Systems of Linear Equations*. PhD thesis, Dept. of Computer Science, Yale University, 1982. Available as Tech. Report 229.
13. D. J. Evans and C. Li. New stopping criteria for some iterative methods for a class of unsymmetric linear systems. *SINUM*, 12(3):549–551, 1991.
14. V. Faber and T. A. Manteuffel. Necessary and sufficient conditions for the existence of a conjugate gradient method. *SIAM J. Numer. Anal.*, 21(2):352–362, 1984.
15. Bernd Fischer and Gene H. Golub. On the error computation for polynomial based iteration methods. In *Recent Advances in Iterative Methods*, pages 59–69, New York, 1993. Springer.
16. Roy G. Gordon. Error bounds in equilibrium statistical mechanics. *Journal of Mathematical Physics*, 9(5):655–663, 1967.
17. Roy G. Gordon. Error bounds in spectroscopy and nonequilibrium statistical mechanics. *Journal of Mathematical Physics*, 9(7):1087–1092, 1967.
18. L. A. Hageman and D. M. Young. *Applied Iterative Methods*. Academic Press, New York, 1981.
19. M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Nat. Bur. Standards*, 49:409–435, 1952.
20. R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, New York, 1991.
21. W. D. Joubert and D. M. Young. Necessary and sufficient conditions for the simplification of generalized conjugate gradient algorithms. *Linear Algebra Appl.*, 88/89:449–485, 1987.
22. E. F. Kaaschieter. A practical termination criterion for the conjugate gradient method. *BIT*, 28:308–322, 1988.
23. Carl Timothy (C. T.) Kelley. *Iterative Methods for Linear and Nonlinear Equations*, volume 16 of *SIAM Frontiers Series*. SIAM, Philadelphia, 1995.
24. D. R. Kincaid, J. R. Respess, D. M. Young, and R. G. Grimes. ITPACK 2C: A FORTRAN package for solving large sparse linear systems by adaptive accelerated iterative methods. *ACM TOMS*, 8(3):302–322, 1982.
25. T. Manteuffel and J. Otto. On the roots of the orthogonal polynomials and residual polynomials associated with a conjugate gradient method. *J. Num. Lin. Alg. Appl.*, pages 449–475, 1994.
26. W. Oettli and W. Prager. Compatibility of approximate solution of linear equations with given error bounds for coefficients and right-hand sides. *NM*, 58:343–361, 1984.
27. Robert D. Skeel. Scaling for numerical stability in Gaussian elimination. *SIAM J. Sci. Statist. Comput.*, 26:494–526, 1979.
28. Robert D. Skeel. Iterative refinement implies numerical stability for gaussian elimination. *Math. Comp.*, 35:817–832, 1980.

29. C. Vuik. Solution of the discretized incompressible Navier-Stokes equations with the GMRES method. *J. Assoc. Comp. Mach.*, 16:507–523, 1983.
30. Karl F. Warnick. Continued fraction error bound for conjugate gradient method. Technical Report ECEN Department Report No. TR-L100-98.3, Dept. of Electrical and Computer Engineering, Brigham Young University, Provo, UT 84602, 9 June 1997.
31. D. M. Young and K. C. Jea. Generalized conjugate gradient acceleration of non-symmetrizable iterative methods. *Linear Algebra Appl.*, 34:159–194, 1980.
32. D. M. Young and T. Mai. ITPACK 3A user's guide (preliminary version). Technical Report CNA-197, Center for Numerical Analysis, University of Texas, Austin, 1984.