SIAM J. MATRIX ANAL. APPL.
Vol. 18, No. 1, pp. 37–51, January 1997

004

GMRES ON (NEARLY) SINGULAR SYSTEMS*

PETER N. BROWN† AND HOMER F. WALKER‡

Abstract. We consider the behavior of the GMRES method for solving a linear system $Ax = b$ when A is singular or nearly so, i.e., ill conditioned. The (near) singularity of A may or may not affect the performance of GMRES, depending on the nature of the system and the initial approximate solution. For singular A, we give conditions under which the GMRES iterates converge safely to a least-squares solution or to the pseudoinverse solution. These results also apply to any residual minimizing Krylov subspace method that is mathematically equivalent to GMRES. A practical procedure is outlined for efficiently and reliably detecting singularity or ill conditioning when it becomes a threat to the performance of GMRES.

Key words. GMRES method, residual minimizing methods, Krylov subspace methods, iterative linear algebra methods, singular or ill-conditioned linear systems

AMS subject classification. 65F10

PII. S0895479894262339

1. Introduction.

The generalized minimal residual (GMRES) method of Saad and Schultz [16] is widely used for solving a general linear system

$$Ax = b, \qquad A \in \mathbb{R}^{n\times n}, \tag{1.1}$$

and its behavior is well understood when A is nonsingular. Our purpose here is to examine the behavior of GMRES when A is singular or nearly so, i.e., ill conditioned, and to formulate practically effective ways of recognizing singularity or ill conditioning when it might significantly affect the performance of the method.

Abstractly, GMRES begins with an initial approximate solution x_0 and initial residual $r_0 \equiv b - Ax_0$ and characterizes the kth approximate solution as $x_k = x_0 + z_k$, where z_k solves

$$\min_{z\in\mathcal{K}_k} \|b - A(x_0 + z)\|_2 = \min_{z\in\mathcal{K}_k} \|r_0 - Az\|_2. \tag{1.2}$$

Here, $\mathcal{K}_k$ is the kth Krylov subspace determined by A and r_0, defined by

$$\mathcal{K}_k \equiv \text{span}\{r_0, Ar_0, \ldots, A^{k-1}r_0\}.$$

There are a number of ways of implementing GMRES, but in each one generates a basis of $\mathcal{K}_k$ and then replaces (1.2) by an unconstrained k-dimensional least-squares

* Received by the editors January 31, 1994; accepted for publication (in revised form) by R. Freund February 17, 1996.

http://www.siam.org/journals/simax/18-1/26233.html

† Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94550 (pnbrown@llnl.gov). The research of this author was supported in part by the Applied Mathematical Sciences subprogram of the Office of Scientific Computing, U.S. Dept. of Energy, by Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

‡ Department of Mathematics and Statistics, Utah State University, Logan, UT 84322-3900 (walker@math.usu.edu). The work of this author was supported in part by U.S. Air Force Office of Scientific Research grant AFOSR-91-0294, U.S. Dept. of Energy grants DE-FG02-92ER25136 and DE-FG03-94ER25221, and National Science Foundation grant DMS-9400217, all with Utah State University. It was done in part during visits to the Computing and Mathematics Research Division, Lawrence Livermore National Laboratory, and the Center for Research on Parallel Computation, Rice University.

problem. We shall not be more specific about the basis generating process at this point, except to assume that it successfully generates a basis if and only if $\dim \mathcal{K}_k = k$, where "dim" denotes dimension.

Note that, trivially, $\dim A(\mathcal{K}_k) \leq \dim \mathcal{K}_k \leq k$ for each k. We shall say that GMRES *does not break down* at the kth step if $\dim A(\mathcal{K}_k) = k$. In this case, $\dim A(\mathcal{K}_k) = \dim \mathcal{K}_k$ and, hence, (1.2) has a unique solution. Furthermore, since $\dim \mathcal{K}_k = k$, a basis of $\mathcal{K}_k$ is successfully generated and the k-dimensional least-squares problem also has a unique solution. This definition addresses two distinct kinds of breakdown: *rank deficiency of the least-squares problem* (1.2), which occurs when $\dim A(\mathcal{K}_k) < \dim \mathcal{K}_k$, and *degeneracy of* $\mathcal{K}_k$, which occurs when $\dim \mathcal{K}_k < k$. The definition is intended to focus on essential breakdown of the method, as opposed to breakdown associated with any particular implementation or ancillary algorithm used in it. Note that if $\dim A(\mathcal{K}_k) < k$ for some k, then $\mathcal{K}_j = \mathcal{K}_k$ for all $j \geq k$ and no further improvement is possible, even if subsequent $z_j \in \mathcal{K}_j$ are well defined in some way.

For perspective, we recall that Proposition 2, p. 865, of [16] ensures that, if A is nonsingular, then GMRES does not break down until the solution of (1.1) has been found. Breakdown in [16, Prop. 2, p. 865] is associated specifically with breakdown of the Arnoldi process used in the GMRES implementation in [16], but the statement remains true with our definition (see section 2 below).

In contrast to the nonsingular case, anything may happen when A is singular. Example 1.1 below shows that GMRES may break down before getting anywhere at all, even when the system has a solution, or it may determine a least-squares solution[1] or the pseudoinverse solution[2] without breaking down. Example 1.2 shows that even if a least-squares solution or the pseudoinverse solution is reached, this may not be evident from the behavior of GMRES; indeed, GMRES may continue for a number of additional steps without breakdown (or further progress).

Example 1.1. Suppose that

$$A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \qquad b = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad x_0 = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

Then $r_0 = (1,0)^T$ and $Ar_0 = (0,0)^T$, and GMRES breaks down at the first step. Note that x_0 is not a (least-squares) solution. If A is changed to

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix},$$

then, for the same b and x_0, we have $r_0 = (1,0)^T = Ar_0$, and GMRES determines without breakdown $x_1 = (1,0)^T$, which is a least-squares solution but not the pseudoinverse solution. If we also change b to $b = (1,1)^T$, then, for the same x_0, we have $r_0 = (1,1)^T$ and $Ar_0 = (2,0)^T$, and GMRES determines without breakdown $x_1 = (1/2, 1/2)^T$, which is the pseudoinverse solution. Note that $\dim A(\mathcal{K}_2) = 1$ in these last two cases, so GMRES breaks down at the step after the least-squares or pseudoinverse solution has been found.

Example 1.2. For arbitrary n, let A be the "shift" operator with ones on the first subdiagonal and zeros elsewhere. Then for $b = (1,0,\ldots,0)^T$ and $x_0 = (0,\ldots,0)^T$,

[1] An $x \in \mathbb{R}^n$ for which $\|b - Ax\|_2$ is minimal.
[2] The least-squares solution x such that $\|x\|_2$ is minimal.

x_0 itself is the pseudoinverse solution, but GMRES proceeds without breakdown (or progress) until the nth step, at which point it breaks down with $\dim A(\mathcal{K}_n) = n - 1$.

In section 2 below, we explore the theoretical behavior of GMRES when A is singular and, in particular, determine circumstances in which the GMRES iterates converge without breakdown to a least-squares solution or the pseudoinverse solution of (1.1). We also discuss the conditioning of the least-squares problem (1.2) prior to breakdown, since this is crucial to the practical performance of the method. The results in section 2 apply not only to GMRES but also to any mathematically equivalent method, i.e., any method that takes steps characterized by the residual minimizing property (1.2). (See [8, sect. 2.4] for a discussion of mathematically equivalent methods.) Thus in section 2, one can think of GMRES as a *generic* minimal residual method that characterizes corrections by (1.2). In section 3, we discuss further how singularity or ill conditioning can appear in GMRES and affect its practical performance. We outline an efficient and reliable way of detecting singularity or ill conditioning when it threatens to cause breakdown or otherwise degrade the performance of the method. In section 4, we discuss several illustrative numerical experiments.

Others have considered GMRES and related methods on singular or ill-conditioned systems. It is noted in [3] and [15] that GMRES can be used to solve singular homogeneous systems that arise in Markov chain modeling. In [9], conditions are given for the convergence of general Krylov subspace methods on singular systems, and particular results are derived for the QMR [10] and TFQMR [7] methods (see section 2 below), with applications to Markov chain modeling. Deflation-like modifications of GMRES based on truncated singular value decomposition solutions have recently been considered in [12]; see also [13] and the references in [12] and [13] for more on deflation techniques for nearly singular systems. In [14], extensions of GMRES are considered in which Krylov subspaces are augmented with approximate eigenvectors generated during previous iterations. These extensions appear to be most effective when there are a few relatively small eigenvalues.

In the following, we denote the null space and range of A by $\mathcal{N}(A)$ and $\mathcal{R}(A)$, respectively, and say that (1.1) is *consistent* if $b \in \mathcal{R}(A)$. We set $r_k \equiv b - Ax_k$ for each k and denote the restriction of A to a subspace $\mathcal{S} \subseteq \mathbb{R}^n$ by $A|_{\mathcal{S}}$. As a convention, we always regard x_0 as determined without breakdown at the "0th" step and define $\mathcal{K}_0 \equiv \{0\}$. Also, we assume that GMRES terminates immediately upon breakdown.

2. Theoretical discussion. Although our interest is primarily in (1.1) when A is singular, the results in this section also apply, as appropriate, when A is nonsingular. The questions of interest are the following:

- Will GMRES determine a least-squares solution without breakdown?
- When has a least-squares solution been reached by GMRES?
- When is a least-squares solution determined by GMRES the pseudoinverse solution?
- How ill conditioned can the GMRES least-squares problem (1.2) become?

We begin with several general results.

LEMMA 2.1. *Apply GMRES to* (1.1) *and suppose that* $\dim \mathcal{K}_k = k$ *for some* $k \geq 0$. *Then exactly one of the following holds:*

(i) $\dim A(\mathcal{K}_k) = k - 1$ *and* $A(x_0 + z) \neq b$ *for every* $z \in \mathcal{K}_k$;
(ii) $\dim A(\mathcal{K}_k) = k$, $\dim \mathcal{K}_{k+1} = k$, x_k *is uniquely defined, and* $Ax_k = b$;
(iii) $\dim A(\mathcal{K}_k) = k$, $\dim \mathcal{K}_{k+1} = k + 1$, x_k *is uniquely defined, and* $Ax_k \neq b$.

Proof. First, note that if $\dim \mathcal{K}_k = k$ and $k > 0$, then $\dim A(\mathcal{K}_{k-1}) = k - 1$. Indeed, in this case r_0, Ar_0, ..., $A^{k-1}r_0$ constitute a basis of $\mathcal{K}_k$ and, therefore,

$Ar_0, \ldots, A^{k-1}r_0$ constitute a basis of $A(\mathcal{K}_{k-1})$. With this observation and the fact that $A(\mathcal{K}_{k-1}) \subseteq A(\mathcal{K}_k)$ for $k > 0$, it is clear that the assumption $\dim \mathcal{K}_k = k$ implies $k - 1 \leq \dim A(\mathcal{K}_k) \leq k$ for all $k \geq 0$. Note also that $r_0 \notin A(\mathcal{K}_{k-1})$ if $k > 0$.

If $\dim A(\mathcal{K}_k) = k - 1$, then conclusions (ii) and (iii) cannot hold. Furthermore, $k > 0$ and $A(\mathcal{K}_{k-1}) = A(\mathcal{K}_k)$ in this case, and, since $r_0 \notin A(\mathcal{K}_{k-1})$, it follows that $r_0 \notin A(\mathcal{K}_k)$. Then $A(x_0 + z) \neq b$ for every $z \in \mathcal{K}_k$, and (only) conclusion (i) holds.

Suppose that $\dim A(\mathcal{K}_k) = k$. Then x_k is uniquely defined; furthermore, since $A(\mathcal{K}_k) \subseteq \mathcal{K}_{k+1}$, we have $k = \dim A(\mathcal{K}_k) \leq \dim \mathcal{K}_{k+1} \leq k + 1$. If $\dim \mathcal{K}_{k+1} = k$, then we must have $A(\mathcal{K}_k) = \mathcal{K}_{k+1}$ and, hence, $r_0 \in A(\mathcal{K}_k)$. It follows from (1.2) that $r_k = 0$ and $Ax_k = b$; thus (only) conclusion (ii) holds. If $\dim \mathcal{K}_{k+1} = k + 1$, then $r_0 \notin A(\mathcal{K}_k)$, $r_k \neq 0$, $Ax_k \neq b$, and (only) conclusion (iii) holds. ∎

This lemma implies the following result.

THEOREM 2.2. *Apply GMRES to* (1.1). *Then, at some step, either*

(a) *GMRES breaks down through rank deficiency of the least-squares problem* (1.2) *without determining a solution or*

(b) *GMRES determines a solution without breakdown and then breaks down at the next step through degeneracy of the Krylov subspace.*

Proof. We have $\dim \mathcal{K}_0 = 0$. Assume that for some $k \geq 0$ GMRES has proceeded to the kth step with $\dim \mathcal{K}_k = k$. Then exactly one of the three conclusions of Lemma 2.1 must hold. If conclusion (i) holds, then we have (a) above. If conclusion (ii) holds, then we have (b). If conclusion (iii) holds, then $Ax_k \neq b$ and the iteration continues to the next step. The theorem follows by induction. ∎

The alternatives of this theorem give useful insights into the eventual outcome of applying GMRES to (1.1). For example, if (1.1) is not consistent, then breakdown through rank deficiency of (1.2) will eventually occur; in practice, this may be preceded by dangerous ill conditioning, as discussed further below. Conversely, breakdown through degeneracy of the Krylov subspace occurs if and only if (1.1) is consistent and the solution has been found. Also, these results imply the result in [16, Prop. 2, p. 865] cited earlier: if A is nonsingular, then GMRES does not break down until the solution of (1.1) has been found. Indeed, if A is nonsingular, then GMRES cannot break down through rank deficiency of (1.2), and the second alternative must hold. However, the reader is cautioned to make inferences carefully; e.g., Example 1.1 above shows that there can be breakdown through rank deficiency in the consistent case before a solution is found.

The next result characterizes circumstances in which a least-squares solution has been reached.

LEMMA 2.3. *At the kth step, GMRES determines a least-squares solution of* (1.1) *without breakdown if and only if*

$$\dim A^T(\mathcal{K}_{k+1}) = \dim A(\mathcal{K}_k) = k. \tag{2.1}$$

Proof. By definition, GMRES does not break down at the kth step if and only if $\dim A(\mathcal{K}_k) = k$. Thus we need only show that x_k is a least-squares solution of (1.1) if and only if $\dim A^T(\mathcal{K}_{k+1}) = \dim A(\mathcal{K}_k)$.

From (1.2), we have that x_k is a least-squares solution of (1.1) if and only if it is possible to reach a least-squares solution of (1.1) through *some* correction in $\mathcal{K}_k$, i.e., if and only if there is some $z \in \mathcal{K}_k$ such that

$$0 = A^T [b - A(x_0 + z)] = A^T(r_0 - Az). \tag{2.2}$$

But (2.2) holds for some $z \in \mathcal{K}_k$ if and only if $A^T r_0 \in A^T A(\mathcal{K}_k)$, which is equivalent to $A^T(\mathcal{K}_{k+1}) = A^T A(\mathcal{K}_k)$. To complete the proof, we note that $\dim A^T A(\mathcal{K}_k) =$

$\dim A(\mathcal{K}_k)$. Indeed, we clearly have $\dim A^T A(\mathcal{K}_k) \leq \dim A(\mathcal{K}_k)$. If $\dim A^T A(\mathcal{K}_k) < \dim A(\mathcal{K}_k)$, then there is a $w \in \mathcal{K}_k$ such that $Aw \neq 0$ and $A^T Aw = 0$. But then $0 = w^T A^T Aw = \|Aw\|_2^2$, which is a contradiction. ∎

With Lemma 2.1, one can easily extend Lemma 2.3 to conclude additionally that if (2.1) holds, then (1.1) is consistent if and only if $\dim \mathcal{K}_{k+1} = k$; i.e., GMRES breaks down at step $k+1$ through degeneracy of the Krylov subspace.

We use Lemma 2.3 to characterize the property of A that yields the most satisfactory answers to the questions posed at the beginning of this section. This property is $\mathcal{N}(A) = \mathcal{N}(A^T)$, equivalently, $\mathcal{N}(A) = \mathcal{R}(A)^\perp$, which holds when A is normal, e.g., when it is symmetric or skew symmetric. It also clearly holds when A is nonsingular. In general, this property holds if and only if $\mathcal{N}(A)^\perp$ is an invariant subspace of A. Also, it holds only if all eigenvectors of A associated with nonzero eigenvalues are orthogonal to $\mathcal{N}(A)$. Note that it does *not* hold for the matrices of Example 1.1; indeed, it holds for $A \in \mathbb{R}^{2\times 2}$ if and only if A is either nonsingular or symmetric. Neither does it hold for the "shift" operator of Example 1.2.

THEOREM 2.4. *GMRES determines a least-squares solution of* (1.1) *without breakdown for all b and x_0 if and only if $\mathcal{N}(A) = \mathcal{N}(A^T)$. If $\mathcal{N}(A) = \mathcal{N}(A^T)$ and a least-squares solution is reached at some step, then GMRES breaks down at the next step, with breakdown through degeneracy of the Krylov subspace if* (1.1) *is consistent and through rank deficiency of the least-squares problem* (1.2) *otherwise. Furthermore, if* (1.1) *is consistent and $x_0 \in \mathcal{R}(A)$, then the solution reached is the pseudoinverse solution.*

Proof. First, suppose that $\mathcal{N}(A) \neq \mathcal{N}(A^T)$. One can choose b and x_0 such that $r_0 \in \mathcal{N}(A)$ and $A^T r_0 \neq 0$. Then x_0 is not a least-squares solution. Furthermore, $\dim A(\mathcal{K}_1) = 0$, so GMRES breaks down at the first step before reaching a least-squares solution.

Now assume $\mathcal{N}(A) = \mathcal{N}(A^T)$. Then for each k, we have $\dim A^T(\mathcal{K}_{k+1}) = \dim A(\mathcal{K}_{k+1})$, and (2.1) becomes

$$\dim A(\mathcal{K}_{k+1}) = \dim A(\mathcal{K}_k) = k.$$

This condition must hold for some k, $0 \leq k \leq n$, and it follows from Lemma 2.3 that GMRES determines a least-squares solution x_k without breakdown at the kth step. Furthermore, since $\dim A(\mathcal{K}_{k+1}) = k$, GMRES breaks down at step $k+1$. One concludes from Theorem 2.2 that breakdown is through degeneracy of the Krylov subspace if (1.1) is consistent and through rank deficiency of the least-squares problem (1.2) otherwise. If (1.1) is consistent, then x_k is a solution and, furthermore, $\mathcal{K}_k \subseteq \mathcal{R}(A)$. If in addition $x_0 \in \mathcal{R}(A)$, then $x_k = x_0 + z_k \in x_0 + \mathcal{K}_k \subseteq \mathcal{R}(A) = \mathcal{N}(A)^\perp$. Since a (least-squares) solution of (1.1) is the pseudoinverse solution if and only if it lies in $\mathcal{N}(A)^\perp$, it follows that x_k is the pseudoinverse solution. ∎

If it is known that $\mathcal{N}(A) = \mathcal{N}(A^T)$, then Theorem 2.4 provides theoretical assurance not only that GMRES will determine a least-squares solution of (1.1) without breakdown but also that reaching it will be indicated by breakdown at the next step. If (1.1) is consistent as well, then choosing $x_0 \in \mathcal{R}(A)$, e.g., $x_0 = 0$, will yield the pseudoinverse solution without breakdown, and reaching it will be indicated by zero residual norm.

If $\mathcal{N}(A) = \mathcal{N}(A^T)$ and (1.1) is consistent, then the least-squares problem (1.2) will remain as well conditioned as the nature of A will allow until a solution of (1.1) is reached. Indeed, if we denote

$$A_k \equiv A|_{\mathcal{K}_k},$$

then the appropriate condition number for (1.2) is $\kappa_2(A_k)$, which satisfies

$$\kappa_2(A_k) \equiv \frac{\max\limits_{z\in\mathcal{K}_k,\ z\neq 0} \|Az\|_2 \big/ \|z\|_2}{\min\limits_{z\in\mathcal{K}_k,\ z\neq 0} \|Az\|_2 \big/ \|z\|_2} \leq \frac{\max\limits_{z\in\mathcal{R}(A),\ z\neq 0} \|Az\|_2 \big/ \|z\|_2}{\min\limits_{z\in\mathcal{R}(A),\ z\neq 0} \|Az\|_2 \big/ \|z\|_2} \equiv \kappa_2(A|_{\mathcal{R}(A)}) \tag{2.3}$$

since $\mathcal{K}_k \subseteq \mathcal{R}(A)$ in the consistent case. Note that, since $\mathcal{R}(A) = \mathcal{N}(A^T)^\perp = \mathcal{N}(A)^\perp$, $\kappa_2(A|_{\mathcal{R}(A)})$ is just the ratio of the largest singular value of A to the smallest positive one. Also, recall from above that, in the consistent case, if a solution is reached at some step, then breakdown of GMRES at the next step occurs because of degeneracy of the Krylov subspace and not because of rank deficiency of the least-squares problem (1.2). These reassuring results are to be expected, for if $\mathcal{N}(A) = \mathcal{N}(A^T)$ and (1.1) is consistent, then everything reduces to the nonsingular case on $\mathcal{R}(A) = \mathcal{N}(A)^\perp$.

If $\mathcal{N}(A) = \mathcal{N}(A^T)$ but (1.1) is not consistent, then, despite the theoretical guarantee of Theorem 2.4 that GMRES will not break down, the least-squares problem (1.2) may necessarily become dangerously ill conditioned before a least-squares solution of (1.1) is reached, regardless of the conditioning of $A|_{\mathcal{R}(A)}$. This is shown by Theorem 2.5 below. It is, perhaps, not surprising, because if a least-squares solution is reached at some step, then, in the inconsistent case, breakdown at the next step occurs because of rank deficiency of the least-squares problem (1.2), rather than degeneracy of the Krylov subspace.

THEOREM 2.5. *Suppose that $\mathcal{N}(A) = \mathcal{N}(A^T)$, and denote the least-squares residual for* (1.1) *by r_*. If $r_{k-1} \neq r_*$ for some k, then*

$$\kappa_2(A_k) \geq \frac{\|A_k\|_2}{\|\bar{A}_k\|_2} \cdot \frac{\|r_{k-1}\|_2}{\sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}}, \tag{2.4}$$

where $A_k \equiv A|_{\mathcal{K}_k}$ and $\bar{A}_k \equiv A|_{\mathcal{K}_k + \operatorname{span}\{r_*\}}$.

Proof. Note that $r_* \in \mathcal{R}(A)^\perp = \mathcal{N}(A)$ and $r_{k-1} - r_* \in \mathcal{R}(A) = \mathcal{N}(A)^\perp$. Then, since $r_{k-1} - r_* \in \mathcal{K}_k + \operatorname{span}\{r_*\}$, we have

$$\begin{aligned} \|Ar_{k-1}\|_2 &= \|A(r_{k-1} - r_* + r_*)\|_2 = \|A(r_{k-1} - r_*)\|_2 \\ &\leq \|\bar{A}_k\|_2 \cdot \|r_{k-1} - r_*\|_2 = \|\bar{A}_k\|_2 \cdot \sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}, \end{aligned}$$

whence

$$\frac{\|Ar_{k-1}\|_2}{\|r_{k-1}\|_2} \leq \|\bar{A}_k\|_2 \cdot \frac{\sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}}{\|r_{k-1}\|_2}. \tag{2.5}$$

Since $r_{k-1} \in \mathcal{K}_k$, (2.4) follows from (2.5) and the definition of $\kappa_2(A_k)$ (see (2.3)). ∎

If (1.1) is consistent, then $r_* = 0$ and $\bar{A}_k = A_k$. It follows that (2.4) is just the trivial bound $\kappa_2(A_k) \geq 1$ in this case. In general, we have $1 \geq \|A_k\|_2/\|\bar{A}_k\|_2 \geq \|A_k\|_2/\|A\|_2$, and (2.4) yields

$$\kappa_2(A_k) \geq \frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}}, \tag{2.6}$$

which may be more easily applied in the inconsistent case.

If A is singular and $\mathcal{N}(A) = \mathcal{N}(A^T)$, then it is evident from (2.6) that, for an unfortunate choice of b and x_0, the least-squares problem (1.2) will become so ill

conditioned before breakdown that little or no accuracy can be expected in a solution computed in finite-precision arithmetic. Indeed, in view of (2.6), one would expect that, in many cases, the residual for the computed solution will first decrease in norm for a number of iterations and then lose accuracy and perhaps increase as a least-squares solution is approached and accuracy is degraded by increasing ill conditioning. (This is seen in Experiment 4.2 below.) In such cases, it would clearly be desirable to terminate the iterations when approximately optimal accuracy has been reached. Note that the usual termination criteria based on the size of the residual norm are unlikely to be of any use in this case; some alternative criterion is needed.

We show how (2.6) can be used to derive a heuristic guideline for terminating the iterations at an approximately optimal point in finite-precision arithmetic. We make two assumptions that are reasonable but by no means the only possible assumptions; our main purpose is to demonstrate the method of derivation. (The guideline resulting from these assumptions is borne out well in Experiment 4.2 below.) The first assumption is that $\kappa_2(A_k)$ is about as small as possible, given the lower bound (2.6), i.e., that

$$\kappa_2(A_k) \approx \frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}}.$$

The second assumption is that the computed value of r_k, denoted by $\hat{r}_k$, satisfies

$$\frac{\|\hat{r}_k - r_k\|_2}{\|r_0\|_2} \approx \mathbf{u}\kappa_2(A_k),$$

where $\mathbf{u}$ is unit rounding error. A rigorous worst-case bound on $\|\hat{r}_k - r_k\|_2/\|r_0\|_2$ would require $\mathbf{u}\kappa_2(A_k)$ multiplied by a polynomial of low degree in n and k (see [11, Chap. 5]), but this is not necessary here. With these assumptions, we have

$$\begin{aligned}
\frac{\|\hat{r}_k - r_*\|_2}{\|r_0\|_2} &\leq \frac{\|\hat{r}_k - r_k\|_2}{\|r_0\|_2} + \frac{\|r_k - r_*\|_2}{\|r_0\|_2} \\
&\approx \mathbf{u}\kappa_2(A_k) + \frac{\sqrt{\|r_k\|_2^2 - \|r_*\|_2^2}}{\|r_0\|_2} \\
&\leq \mathbf{u}\kappa_2(A_k) + \frac{\sqrt{\|r_{k-1}\|_2^2 - \|r_*\|_2^2}}{\|r_0\|_2} \\
&\approx \mathbf{u}\kappa_2(A_k) + \frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\|r_0\|_2} \cdot \frac{1}{\kappa_2(A_k)} \\
&= B(\kappa_2(A_k)),
\end{aligned} \tag{2.7}$$

where

$$B(\kappa) \equiv \mathbf{u}\kappa + \frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\|r_0\|_2} \cdot \frac{1}{\kappa}.$$

It is easily seen that B is minimized when

$$\kappa = \kappa_{\min} \equiv \sqrt{\frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\|r_0\|_2} \cdot \frac{1}{\mathbf{u}}}, \tag{2.8}$$

which suggests a heuristic guideline as follows: If the iterations are terminated with $\kappa_2(A_k) \approx \kappa_{\min}$ given by (2.8), then (2.7) gives an approximate minimal bound

$$\frac{\|\hat{r}_k - r_*\|_2}{\|r_0\|_2} \leq B(\kappa_{\min}) = 2\sqrt{\frac{\|A_k\|_2}{\|A\|_2} \cdot \frac{\|r_{k-1}\|_2}{\|r_0\|_2} \cdot \mathbf{u}}\,. \tag{2.9}$$

This can be simplified for practical purposes by assuming that $\|A_k\|_2/\|A\|_2 \approx 1$ and $\|r_{k-1}\|_2 \approx \|\hat{r}_{k-1}\|_2$. We discuss how to monitor $\kappa_2(A_k)$ efficiently in practice in section 3.

If $\mathcal{N}(A) \neq \mathcal{N}(A^T)$, then it follows from Theorem 2.4 that, for *some* b and x_0, GMRES will break down before determining a least-squares solution of (1.1). However, there is an important special case in which GMRES still reliably determines a least-squares solution, viz., that in which $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$ and (1.1) are consistent. This occurs, e.g., in Experiment 4.3 below.

THEOREM 2.6. *Suppose that $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$. If (1.1) is consistent, then GMRES determines a solution without breakdown at some step and breaks down at the next step through degeneracy of the Krylov subspace.*

Proof. Since (1.1) is consistent, $r_0 \in \mathcal{R}(A)$ and $\mathcal{K}_k \subseteq \mathcal{R}(A)$ for each k. Since $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$, this implies that $\dim A(\mathcal{K}_k) = \dim \mathcal{K}_k$ for each k. Then there cannot be breakdown through rank deficiency of the least-squares problem (1.2), and the theorem follows from Theorem 2.2. ∎

Conditions that are essentially equivalent to those in Theorem 2.6 appear in [9]. The *index* of A, denoted index(A), is defined to be the smallest integer q such that A^q and A^{q+1} have the same rank. It is easily seen that index$(A) = 1$ if and only if A is singular and $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$. For a consistent system (1.1) with index$(A) = 1$, general conditions are given in [9] under which a Krylov subspace method is convergent. It is further shown in [9] that the QMR and TFQMR methods are convergent for such a system.

If $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$ and (1.1) is consistent, then $\kappa_2(A_k)$ satisfies (2.3). However, note that if $\mathcal{N}(A) \neq \mathcal{N}(A^T)$, then $\min_{z \in \mathcal{R}(A),\ z \neq 0} \|Az\|_2/\|z\|_2$ may be smaller than the smallest positive singular value of A, and so $\kappa_2(A|_{\mathcal{R}(A)})$ may be larger than the ratio of the largest singular value of A to the smallest positive one. Still, the least-squares problem (1.2) is as well conditioned as the nature of A will allow and cannot become arbitrarily ill conditioned before a solution is determined by GMRES through an unfortunate choice of b and x_0. This is not surprising, since GMRES breakdown occurs because of degeneracy of the Krylov subspace, rather than rank deficiency of the least-squares problem (1.2). As when (1.1) is consistent and $\mathcal{N}(A) = \mathcal{N}(A^T)$, the setting reduces to the nonsingular case on $\mathcal{R}(A)$, although now $\mathcal{R}(A)$ may not be $\mathcal{N}(A)^\perp$. When (1.1) is not consistent, breakdown must occur because of rank deficiency of (1.2), and in general we cannot expect (1.2) to remain well conditioned, whether or not a least-squares solution is reached.

We conclude this section by noting that, in some applications, one can easily project b onto $\mathcal{R}(A)$. For example, in each of Experiments 4.2 and 4.3 below, $\mathcal{N}(A^T)$ is one dimensional, and it is not difficult to determine a unit vector in $\mathcal{N}(A^T)$ and then to project b onto $\mathcal{N}(A^T)^\perp = \mathcal{R}(A)$. In such an application, if GMRES can be expected to behave well on a consistent system, e.g., if $\mathcal{N}(A) = \mathcal{N}(A^T)$ or $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$, then it is clearly desirable to project b onto $\mathcal{R}(A)$ before starting GMRES. By doing this, one can determine a least-squares solution for the original b without risking the dangerous ill conditioning that may precede GMRES breakdown with rank deficiency

of (1.2). In addition, if $\mathcal{N}(A) = \mathcal{N}(A^T)$, then one can determine the pseudoinverse solution by taking $x_0 \in \mathcal{R}(A)$, e.g., $x_0 = 0$.

3. Practical handling of (near) singularity. In section 2, we considered the conditioning of the least-squares problem (1.2) and how it might be affected by A and perhaps b and x_0. In this section, we look further into how singularity or ill conditioning can arise in GMRES and discuss how conditioning can be monitored efficiently in practice.

Recall from section 1 that, prior to breakdown, an implementation of GMRES generates a basis of $\mathcal{K}_k$ for each k. We denote the matrix having the basis vectors as columns by $B_k \in \mathbb{R}^{n\times k}$. The kth GMRES correction z_k, which is the solution of (1.2), is not computed for each k, but when desired, it is determined by first finding y_k that solves

$$\min_{y\in\mathbb{R}^k} \|r_0 - AB_k y\|_2 \tag{3.1}$$

and then forming $z_k = B_k y_k$. Thus ill conditioning or singularity is a concern in GMRES only if it becomes manifested in ill conditioning or rank deficiency of AB_k or B_k.

Sound GMRES implementations are designed so that, as much as possible, each B_k is well conditioned regardless of the conditioning of A. For example, the standard implementation of [16] and Householder variants in [18] determine ideally conditioned B_k such that $B_k^T B_k = I_k$ (in exact arithmetic). Other implementations in [2] and [19] generate B_k that are usually well conditioned, if not ideally conditioned. In any event, in well-constructed GMRES implementations, the conditioning of B_k does not suffer directly from ill conditioning of A; furthermore, any ill conditioning of B_k seems likely to be reflected in ill conditioning of AB_k. Therefore, we focus on the conditioning of AB_k here.

In practice, a reasonable course is to monitor the conditioning of AB_k and terminate the GMRES iterations if excessive ill conditioning or rank deficiency appears. Typically, the solution of (3.1) is computed using a factorization $AB_k = Q_k R_k$, where $Q_k \in \mathbb{R}^{n\times k}$ has orthonormal columns and $R_k \in \mathbb{R}^{k\times k}$ is upper triangular. It is reasonable to assume that this factorization is determined using one or more stable factorization techniques. For example, the implementations of [16] and [18] first use modified Gram–Schmidt or, respectively, Householder transformations to produce $AB_k = B_{k+1}H_k$, where $H_k \in \mathbb{R}^{(k+1)\times k}$ is upper Hessenberg, and then use plane rotations $J_1, \ldots, J_k$ to obtain $A_k B_k = Q_k R_k$ with $Q_k = B_{k+1}J_1^T \ldots J_k^T (I_k, 0)^T$ and $R_k = (I_k, 0)J_k \ldots J_1 H_k$. In general, each Q_k may be only implicitly specified, as in the implementations of [16] and [18], but each R_k is always produced explicitly. Then, since the conditioning of AB_k is determined by that of R_k, it suffices to monitor the conditioning of R_k and terminate the iterations if excessive ill conditioning or singularity appears.

In the important case in which $B_k^T B_k = I_k$, as in the implementations of [16] and [18], we have $\kappa_2(R_k) = \kappa_2(AB_k) = \kappa_2(A_k) \le \kappa_2(A)$, where $A_k = A|_{\mathcal{K}_k}$ as above. This inequality need not be strict; for example, if A is nonsingular and GMRES proceeds for n steps without breakdown, then $A_n = A$ and $\kappa_2(R_n) = \kappa_2(A_n) = \kappa_2(A)$. Thus R_k can become fully as ill conditioned as A. However, if r_0 lies in an invariant proper subspace, then $\kappa_2(R_k)$ may remain much less than $\kappa_2(A)$. The following example illustrates extreme behavior.

Example 3.1. Assume that $B_k^T B_k = I_k$ for each k. Suppose that we have $\sigma_1 \ge$

$\cdots \geq \sigma_{n-1} = \sigma_n > 0$, and define

$$A \equiv \begin{pmatrix} 0 & \cdots & 0 & \sigma_{n-1} & 0 \\ \sigma_1 & \ddots & \vdots & 0 & 0 \\ 0 & \ddots & 0 & \vdots & \vdots \\ \vdots & \ddots & \sigma_{n-2} & 0 & 0 \\ 0 & \cdots & 0 & 0 & \sigma_n \end{pmatrix}.$$

Clearly, $\sigma_1, \ldots, \sigma_n$ are the singular values of A, and $\kappa_2(A) = \sigma_1/\sigma_n$. For $i = 1, \ldots, n$, let e_i denote the ith column of I_n. If $r_0 = e_1$, then we have $\mathcal{K}_k = \text{span}\{e_1, \ldots, e_k\}$ and $\kappa_2(R_k) = \kappa_2(A_k) = \sigma_1/\sigma_k$ for $k = 1, \ldots, n-1$. In particular, the solution is reached at the $(n-1)$st step with $\kappa_2(R_{n-1}) = \sigma_1/\sigma_{n-1} = \sigma_1/\sigma_n = \kappa_2(A)$. However, if $r_0 = e_n$, then the solution is reached at the first step with $\kappa_2(R_1) = \sigma_n/\sigma_n = 1$.

A very efficient means of monitoring the conditioning of R_k is provided by *incremental condition estimation* (ICE) [4], [5]. This determines estimates of the largest and smallest singular values of each R_k in $O(k)$ arithmetic operations, given estimates of the largest and smallest singular values of R_{k-1}. Thus one can begin with $k = 1$ and use ICE to estimate incrementally the condition number of each successive R_k as k increases. Over a cycle of m GMRES steps, the total cost of estimating the condition number of each R_k, $1 \leq k \leq m$, is $O(m^2)$ arithmetic operations, which is negligible in most applications. A well-developed Fortran implementation of ICE is provided by auxiliary routine xLAIC1 of LAPACK [1], where x = S for single precision or x = D for double precision. This implementation was used in all of the numerical experiments reported in section 4.

4. Numerical experiments. In this section, we discuss several numerical experiments that illustrate the theoretical and practical points brought out above. A standard modified Gram–Schmidt GMRES implementation, as originally outlined in [16], was used in all experiments. Recall that with this implementation, the basis matrix B_k is ideally conditioned, with $B_k^T B_k = I_k$. This implementation was augmented with routine DLAIC1 of LAPACK for monitoring conditioning of the triangular factor of AB_k as discussed above. In all experiments, we took the zero vector to be the initial approximate solution and specified a stopping tolerance tol so that the GMRES iterations would terminate when $\|r_k\|_2 \leq \text{tol}\|b\|_2$. Of course, there was no expectation of stopping on the basis of such a test in cases in which (1.1) was not consistent; in these cases, termination was based on other criteria noted below. All computing was done in double precision Fortran on Sun Microsystems Sparc architectures.

Experiment 4.1. This experiment, which involves a contrived problem, points out the danger of not monitoring the conditioning of AB_k and terminating when excessive ill conditioning appears. The matrix A is from the example in [6, sect. 6],

$$A = \begin{pmatrix} 0 & 1 & & \\ -1 & \ddots & \ddots & \\ & \ddots & \ddots & 1 \\ & & -1 & 0 \end{pmatrix}.$$

We assume that n is odd, in which case A is singular with

$$\mathcal{N}(A) = \text{span}\{(1, 0, 1, 0, \ldots, 0, 1)^T\}.$$

Since A is skew symmetric, the conclusions of Theorem 2.4 hold, at least in exact arithmetic, and GMRES should find a least-squares solution of (1.1) without breakdown and then exhibit breakdown at the next step. In floating point arithmetic, however, GMRES produced misleading results.

We took $n = 49$, tol $= 10^{-6}$ and first ran GMRES with

$$b = (1/\sqrt{2}, 0, \ldots, 0, -1/\sqrt{2})^T,$$

for which (1.1) is consistent. GMRES safely terminated with a computed residual norm of 1.57×10^{-16} when the pseudoinverse solution was reached at the 24th step; the largest observed condition number estimate was 12.7. We then ran GMRES with $b = (1/\sqrt{2}, 0, \ldots, 0, 1/\sqrt{2})^T$, for which (1.1) is not consistent; the least-squares residual norm is $\sqrt{2}/5$. In exact arithmetic, a least-squares solution would have been obtained at the 24th step, and this would have been indicated by breakdown at the 25th step in the form of rank deficiency in the least-squares problems (1.2) and (3.1). Because of rounding error, exact breakdown did not occur, nor were any arithmetic exceptions such as overflow observed. However, the condition number estimate went from 12.7 at the 24th step to 1.47×10^{16} and 1.79×10^{30} at the 25th and 26th steps, respectively. We allowed GMRES to continue, restarting every 49 steps, until it declared *successful* termination at the 185th step with a computed residual norm of 6.68×10^{-7}. Of course, this was the value of the residual norm maintained recursively by GMRES and not the true residual norm, which was 9.14×10^{12} on termination!

We also note that the GMRES implementation used in these experiments did not re-evaluate the residual and its norm directly at each restart; i.e., it did not multiply the current approximate solution by A and subtract the result from b. Instead, it updated the residual at each restart by forming a certain linear combination of the Arnoldi basis vectors generated in the previous cycle of steps. Such updating saves an A-product at each restart and is usually a safe thing to do, unless extreme residual norm reduction is desired. In this example, however, it was not safe, and re-evaluating the residual directly at restarts would have indicated that GMRES had gone astray.

The next two experiments involve discretizations of boundary value problems for the partial differential equation

$$\Delta u + d\frac{\partial u}{\partial x_1} = f(x), \quad x = (x_1, x_2) \in \Omega \equiv [0,1] \times [0,1], \tag{4.1}$$

where d is a constant and f is a given function. In the experiments reported here, we discretized (4.1) with the usual second-order centered differences on a 100×100 mesh of equally spaced discretization points, so that the resulting linear systems were of dimension $10,000$. We took $d = 10$ and preconditioned the discretized problems on the right with a fast Poisson solver from FISHPACK [17]. This preconditioner is very effective for this fairly small value of d. We took tol $= 10^{-10}$ in order to see how GMRES behaved with a tight stopping tolerance. We also stopped the iterations when the condition number estimate became greater than $1/(50\mathbf{u}) \approx 10^{14}$. In the trials outlined below, there was no need to restart GMRES; in each case, there was termination because of either sufficient residual norm reduction or excessive ill conditioning before the maximum allowable number of iterations (50) had been reached.

In each of these two experiments, it is possible to give a simple characterization of $\mathcal{N}(A^T)$. In each, then, we first consider a b for which (1.1) is not consistent and then project it onto $\mathcal{R}(A)$ to obtain a consistent system that is effectively solved by

TABLE 4.1
GMRES iterations 9–19 on problem (4.1) with periodic boundary conditions.

Iteration no.	GMRES recursive residual norm	Computed residual norm	Condition no. estimate
9	99.000000080681	99.000000080680	7.80×10^3
10	99.000000005202	99.000000005201	4.17×10^4
11	99.000000000146	99.000000000145	1.65×10^5
12	99.000000000008	99.000000000007	9.97×10^5
13	99.000000000002	99.000000000000	4.71×10^6
14	99.000000000002	99.000000000000	3.20×10^7
15	99.000000000001	99.000000000001	1.76×10^8
16	98.999999999935	99.000000000068	1.33×10^9
17	98.999999997323	99.000000002679	8.41×10^9
18	98.999999811806	99.000000188196	7.05×10^{10}
19	98.999990468226	99.000009534599	5.02×10^{11}

GMRES. The result is both an approximate solution of the consistent system and an approximate least-squares solution of the original inconsistent system.

Experiment 4.2. In this experiment, we imposed periodic boundary conditions: $u(x_1, 0) = u(x_1, 1)$ and $u(0, x_2) = u(1, x_2)$ for $0 \le x_1, x_2 \le 1$. The matrix A is given as follows:

$$A = \frac{1}{h^2}\begin{pmatrix} T_m & I_m & & I_m \\ I_m & \ddots & \ddots & \\ & \ddots & \ddots & I_m \\ I_m & & I_m & T_m \end{pmatrix}, \quad T_m = \begin{pmatrix} -4 & \alpha_+ & & \alpha_- \\ \alpha_- & \ddots & \ddots & \\ & \ddots & \ddots & \alpha_+ \\ \alpha_+ & & \alpha_- & -4 \end{pmatrix} \in \mathbb{R}^{m\times m},$$

and $m = \sqrt{n} = 100$, $h = 1/m$, and $\alpha_\pm = 1 \pm dh/2$. It is easy to verify that A is normal and that

$$\mathcal{N}(A) = \mathcal{N}(A^T) = \text{span}\{(1, 1, \ldots, 1)^T\}; \tag{4.2}$$

then Theorems 2.4 and 2.5 are applicable.

We first took b to be a discretization of $f(x) = x_1 + x_2$. For this b, (1.1) is not consistent; the least-squares residual norm is 99. GMRES began with an initial residual norm of 107.1 and terminated after 21 iterations with a condition number estimate greater than the termination value $1/(50\mathbf{u}) \approx 10^{14}$. A subset of the iterations is shown in Table 4.1, which gives to 14-digit accuracy both the residual norm values maintained recursively by GMRES and the directly computed residual norms, as well as the condition number estimates. Note that the two norm values agree well and decrease toward the least-squares residual norm through iteration 15, but then the computed norms begin to increase while the recursive norm values continue erroneously to decrease below the least-squares residual norm. Since $\mathbf{u} \approx 2.2 \times 10^{-16}$ here, the heuristic guideline developed in section 2 would have called for termination when the condition number estimate was about 10^8. Table 4.1 shows that this would have been a very good point at which to terminate: the computed residual norm would have been near its minimum value, and the recursive residual norm value would have still been accurate. Note the pessimism of the bound (2.9) in this case.

Using the characterization of $\mathcal{N}(A^T)$ in (4.2), we next projected the above b onto $\mathcal{R}(A)$ to obtain a consistent system. The initial residual norm was 40.82. After 17 iterations, GMRES successfully met the termination test based on tol $= 10^{-10}$ and

terminated with a residual norm of 2.441×10^{-9}. No major inaccuracy was observed; the recursive residual norm value agreed with the directly computed residual norm to five significant digits. Since $\mathcal{N}(A) = \mathcal{N}(A^T)$ and the initial guess was zero, the final iterate was an approximate pseudoinverse solution of not only the consistent system but also the inconsistent system with the original b.

Experiment 4.3. In this experiment, we imposed Neumann boundary conditions: $\partial u(x)/\partial \nu = 0$ for $x \in \partial\Omega$. The matrix A is now given by

$$A = \frac{1}{h^2}\begin{pmatrix} T_m & 2I_m & & & \\ I_m & T_m & I_m & & \\ & \ddots & \ddots & \ddots & \\ & & I_m & T_m & I_m \\ & & & 2I_m & T_m \end{pmatrix}, \quad T_m = \begin{pmatrix} -4 & 2 & & & \\ \alpha_- & -4 & \alpha_+ & & \\ & \ddots & \ddots & \ddots & \\ & & \alpha_- & -4 & \alpha_+ \\ & & & 2 & -4 \end{pmatrix} \in \mathbb{R}^{m\times m},$$

and m, h, and $\alpha_\pm$ are as in Experiment 4.2. We have $\mathcal{N}(A) = \text{span}\{(1, 1, \cdots, 1)^T\}$ as before, but now $\mathcal{N}(A^T) \neq \mathcal{N}(A)$. Indeed, we determine $\mathcal{N}(A^T)$ as follows: Set

$$D_m \equiv \text{diag}\,(1, 2/\alpha_-, 2\alpha_+/\alpha_-^2, \ldots, 2\alpha_+^{m-3}/\alpha_-^{m-2}, \alpha_+^{m-2}/\alpha_-^{m-2}) \in \mathbb{R}^{m\times m},$$

and define a block-diagonal matrix $D = \text{diag}\,(D_m, 2D_m, \ldots, 2D_m, D_m) \in \mathbb{R}^{n\times n}$. Then one can verify that DA is symmetric, and it therefore follows that $\mathcal{N}(A^T) = \text{span}\{D(1, 1, \ldots, 1)^T\}$. With this characterization of $\mathcal{N}(A^T)$, one sees that $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$; then Theorem 2.6 applies when (1.1) is consistent.

The procedures and observations in this experiment were much like those in Experiment 4.2. We first took b to be a discretization of $f(x) = x_1 + x_2 + \sin 10x_1 \cos 10x_2 + e^{10x_1x_2}$. This gave somewhat more dramatic results than the choice of f in Experiment 4.2. For this b, (1.1) is not consistent; the least-squares residual is 5.302×10^4. GMRES began with an initial residual norm of 1.232×10^5 and terminated after 30 iterations with a condition number estimate greater than $1/(50\mathbf{u}) \approx 10^{14}$. The final computed residual norm was 6.305×10^4, which suggests that the GMRES iterates were not converging to a least-squares solution (at least not in any practical sense, given the very large condition number). We next used the characterization $\mathcal{N}(A^T) = \text{span}\{D(1, 1, \ldots, 1)^T\}$ to project this b onto $\mathcal{R}(A)$ and to obtain a consistent system. The initial residual norm was 1.112×10^5. After 23 iterations, GMRES successfully met the termination test based on $\text{tol} = 10^{-10}$ and terminated with a residual norm of 8.716×10^{-6}. No major inaccuracy was observed; the recursive residual norm agreed with the directly computed residual norm to three significant digits. In this case, the final iterate was not a pseudoinverse solution of either the consistent system or the inconsistent system with the original b.

5. Summary discussion. We have addressed the performance of GMRES on a linear system $Ax = b$ when A is singular or ill conditioned. Theoretical results are given that are of interest primarily in the singular case; these hold not only for GMRES but also for any mathematically equivalent method. In general, at some step, GMRES will either (a) break down through rank deficiency of the GMRES least-squares problem without determining a solution or (b) determine a solution without breakdown and then break down at the next step through degeneracy of the Krylov subspace.

More extensive results hold when $\mathcal{N}(A) = \mathcal{N}(A^T)$. This condition is necessary and sufficient for GMRES to determine a least-squares solution without breakdown for *all* b and x_0. If $\mathcal{N}(A) = \mathcal{N}(A^T)$ and the system is consistent, then the condition

number of the GMRES least-squares problem remains bounded by $\kappa_2(A|_{\mathcal{R}(A)})$, which, in this case, is the ratio of the largest singular value of A to the smallest positive one. If $x_0 \in \mathcal{R}(A)$ as well, then the solution determined by GMRES is the pseudoinverse solution. If $\mathcal{N}(A) = \mathcal{N}(A^T)$ and the system is not consistent, then, for some b and x_0, the GMRES least-squares problem will necessarily become dangerously ill conditioned before a least-squares solution is reached, despite the theoretical guarantee of no breakdown. However, one may be able to use the condition number of the GMRES least-squares problem to determine when to terminate with nearly the best obtainable accuracy.

If $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$ and the system is consistent, then GMRES will produce a solution without breakdown, even if $\mathcal{N}(A) \neq \mathcal{N}(A^T)$. In this case, the condition number of the GMRES least-squares problem again remains bounded by $\kappa_2(A|_{\mathcal{R}(A)})$, but this may be larger than the ratio of the largest singular value of A to the smallest positive one. Still, this condition number cannot become arbitrarily large through an unfortunate choice of b and x_0.

In some applications in which the system is not consistent, it may be possible to project b onto $\mathcal{R}(A)$. If GMRES can be expected to solve consistent systems reliably, e.g., if $\mathcal{N}(A) = \mathcal{N}(A^T)$ or $\mathcal{N}(A) \cap \mathcal{R}(A) = \{0\}$, then applying GMRES to the consistent system with the projected b will safely yield a least-squares solution of the original inconsistent system.

In practice, the kth GMRES step is obtained by reducing the GMRES least-squares problem to an unconstrained k-dimensional least-squares problem, which is solved through QR factorization. In numerically sound GMRES implementations, singularity or ill conditioning of A is a concern only if it becomes manifested in singularity or ill conditioning of the upper-triangular factors, which may or may not occur before a solution is found. The condition numbers of these factors can be estimated very efficiently using incremental condition estimation (ICE) [4], [5].

Acknowledgment. The authors are grateful to the referees and the editor for their helpful comments and suggestions, which significantly improved the paper.

REFERENCES

[1] E. Anderson, et al., *LAPACK User's Guide*, SIAM, Philadelphia, PA, 1992.
[2] Z. Bai, D. Hu, and L. Reichel, *A Newton Basis GMRES Implementation*, Tech. report 91-03, Department of Mathematics, University of Kentucky, Lexington, KY, April 1991.
[3] V. A. Barker, *Numerical solution of sparse singular systems of equations arising from ergodic Markov chain modeling*, Comm. Statist. Stochastic Models, 5 (1989), pp. 335–381.
[4] C. H. Bischof, *Incremental condition estimation*, SIAM J. Matrix Anal. Appl., 11 (1990), pp. 312–322.
[5] C. H. Bischof and P. T. P. Tang, *Robust Incremental Condition Estimation*, Tech. report CS-91-133, LAPACK Working Note 33, Computer Science Department, University of Tennessee, Knoxville, TN, May 1991.
[6] P. N. Brown, *A theoretical comparison of the Arnoldi and GMRES algorithms*, SIAM J. Sci. Statist. Comput., 12 (1991), pp. 58–78.
[7] R. W. Freund, *A transpose-free quasi-minimal residual algorithm for non-Hermitian linear systems*, SIAM J. Sci. Comput., 14 (1993), pp. 425–448.
[8] R. W. Freund, G. H. Golub, and N. M. Nachtigal, *Iterative solution of linear systems*, Acta Numerica, 1 (1992), pp. 57–100.
[9] R. W. Freund and M. Hochbruck, *On the use of two QMR algorithms for solving singular systems and applications in Markov chain modeling*, Numer. Linear Algebra Appl., 1 (1994), pp. 403–420.
[10] 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.

[11] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 2nd ed., The Johns Hopkins University Press, Baltimore, MD, 1989.
[12] J. C. Meza, *A Modification to the GMRES Method for Ill-Conditioned Linear Systems*, Tech. report SAND95-8220, Sandia National Laboratories, Livermore, CA, April 1995.
[13] J. C. Meza and W. W. Symes, *Deflated Krylov subspace methods for nearly singular linear systems*, J. Optim. Theory Appl., 72 (1992), pp. 441–457.
[14] R. B. Morgan, *A restarted GMRES method augmented with eigenvectors*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 1154–1171.
[15] B. Philippe, Y. Saad, and W. J. Stewart, *Numerical methods in Markov chain modeling*, Oper. Res., 40 (1992), pp. 1156–1179.
[16] Y. Saad and M. H. Schultz, *GMRES: a generalized minimal residual method for solving nonsymmetric linear systems*, SIAM J. Sci. Statist. Comput., 7 (1986), pp. 856–869.
[17] P. N. Swarztrauber and R. A. Sweet, *Efficient fortran subprograms for the solution of elliptic partial differential equations*, ACM Trans. Math. Software, 5 (1979), pp. 352–364.
[18] H. F. Walker, *Implementation of the GMRES method using Householder transformations*, SIAM J. Sci. Statist. Comput., 9 (1988), pp. 152–163.
[19] H. F. Walker and L. Zhou, *A simpler GMRES*, Numer. Linear Algebra Appl., 1 (1994), pp. 571–581.