Inexact GMRES for singular linear systems

Xiuhong Du and Daniel B. Szyld

Report 07-09-18
September 2007

This report is available in the World Wide Web at
http://www.math.temple.edu/~szyld

INEXACT GMRES FOR SINGULAR LINEAR SYSTEMS*

XIUHONG DU† AND DANIEL B. SZYLD†

Abstract. Inexact Krylov subspace methods have been shown to be practical alternatives for the solution of certain linear systems of equations. In this paper, the solution of singular systems with inexact matrix-vector products is explored. Criteria are developed to prescribe how inexact the matrix-vector products can be, so that the computed residual remains close to the true residual, thus making the inexact method of practical applicability. Cases are identified for which the methods work well, and this is the case in particular for systems representing certain Markov chains. Numerical experiments illustrate the effectiveness of the inexact approach.

Key words. Singular linear systems. Inexact Krylov subspace methods. Inexact matrix-vector product. Solution of Markov chain problems.

1. Introduction. Consider the iterative solution of large sparse (symmetric or) nonsymmetric $n \times n$ linear systems of the form

$$Az = b, \tag{1.1}$$

with a Krylov subspace method such as GMRES [15]. In these iterative methods, at each iteration, a matrix-vector product with the matrix A is performed. We consider here inexact Krylov subspace methods, where instead of matrix-vector products with A, one uses $\mathbf{A} = A + E$, for some error matrix E, which usually changes from one iteration to the next; see section 2 for some further details. Simoncini and Szyld [16] studied these methods in the case of nonsingular A. They prescribed how to relax the exactness of the matrix-vector product and still achieve the desired convergence. In addition, similar criteria were given to guarantee that the computed residual with the inexact method is close to the true residual; see also [4], [21].

Our aim in this paper is to extend those results to the solution of (1.1) when the matrix A is singular or nearly singular, and the system (1.1) is consistent, i.e., $b \in \mathcal{R}(A)$, the range of A. For discussion and applications of the use of Krylov subspace methods for singular or nearly singular systems, see, e.g., [1], [5]–[8], [12], [13], [18], [22], [24]–[27].

In sections 3 and 4, we develop a theory on how to estimate a threshold η_k such that if at the kth iteration, $\|E\| \leq \eta_k \epsilon$, inexact GMRES produces a computed residual which does not differ from the true residual in norm by more than ϵ. Similarly to the nonsingular case [4], [16], the threshold η_k grows with k, i.e., as the iterations proceed, the matrix-vector product is allowed to be more inexact, and the norm of the matrix E can grow larger.

When the coefficient matrix A has index one (see below for a definition of index), and the null space of A is uni-dimensional, the threshold η_k which we prescribe, can be easily estimated. This is the case for many practical applications, such as the computation of stationary probability distributions of irreducible Markov chains; see, e.g., [19]. In section 5 we illustrate the use of this new criterion and of the use of inexact GMRES with several numerical experiments, and in particular for such Markov chains. When the coefficient matrix has index k_0 larger than one, both our theory and experiments indicate that the right hand side b needs to lie on (or close to) $\mathcal{R}(A^{k_0})$ for the methods to work.

*This version dated 18 September 2007

†Department of Mathematics, Temple University (038-16), 1805 N. Broad Street, Philadelphia, Pennsylvania 19122-6094, USA (dxhdxh@temple.edu, szyld@temple.edu). This research is supported in part by the U.S. National Science Foundation under grant CCF-0514889, and by the U.S. Department of Energy under grant DE-FG02-05ER25672.

We end this introduction by reviewing some definitions and notation, which we use throughout the paper. Let $\mathcal{N}(A)$ be the null space of A and let $\sigma_m(A)$ be the mth singular value of A. The index of a matrix A is the smallest integer $k_0 = \mathrm{ind}(A)$ such that $\mathcal{N}(A^k) = \mathcal{N}(A^{k+1})$ for all $k \geq k_0$. This is equivalent to saying that the largest Jordan block associated with the eigenvalue zero is of order k_0. By U^* we denote the conjugate transpose of U. Throughout the paper we use the 2-norm, simply denoted by $\|\cdot\|$.

2. Description of inexact Krylov subspace methods. For descriptions of GMRES [15], or other Krylov subspace methods, we refer the reader, e.g., to [14], [17]. Here we present some basic relations between matrices used in these methods, and mention how these relations differ when considering an inexact version. In our presentation, we assume exact arithmetic throughout. We note, however, that in Section 5 we study experimentally how the inexact methods behave when the right hand side (or the initial residual) is changed by perturbations (or round-off).

Let z_0 be an initial vector and $r_0 = b - Az_0$ the initial residual. Then, to find the solution of (1.1), it suffices to solve

$$Ax = r_0, \tag{2.1}$$

and then consider $z = z_0 + x$. Let x_m be the approximation to the solution of (2.1) at the mth iteration, so that $x_m \in \mathcal{K}_m(A, v_1)$, the Krylov subspace spanned by $\{v_1,\ Av_1,\ \ldots,\ A^{m-1}v_1\}$, with $v_1 = r_0/\|r_0\|$. Let $V_{m+1} = [v_1,\ v_2,\ \ldots,\ v_{m+1}]$ be a matrix whose columns are an orthonormal basis of $\mathcal{K}_{m+1}(A, v_1)$. It is usually obtained by the Arnoldi method, which consists of computing at the mth step, the product Av_m, which is then orthogonalized with respect to $v_1,\ v_2,\ \ldots,\ v_m$. One can then write the following Arnoldi relation

$$AV_m = V_{m+1}H_{m+1}, \tag{2.2}$$

where H_{m+1} is an $(m+1) \times m$ upper Hessenberg matrix.

In an inexact Krylov method, the product Av_m is not performed exactly. Instead, we have $(A + E_m)v_m$, where E_m is some error matrix. In this case, the usual Arnoldi relation (2.2) does not hold. In its place, we have the following

$$\begin{aligned} [(A + E_1)v_1\ (A + E_2)v_2\ \ldots\ (A + E_m)v_m] &= V_{m+1}H_{m+1}, \\ AV_m + [E_1v_1\ E_2v_2\ \ldots\ E_mv_m] &= V_{m+1}H_{m+1}. \end{aligned} \tag{2.3}$$

This relation is called the inexact Arnoldi relation. We note that while V_{m+1} has orthogonal columns (as in the exact case), its columns no longer span a Krylov subspace of the form $K_{m+1}(A, v_1)$.

The GMRES method computes the approximation $x_m = V_m y_m$, where y_m solves the least squares problem

$$\min_{y \in \mathbb{R}^m} \left\| H_{m+1}y - \|r_0\|e_1 \right\|, \tag{2.4}$$

and e_1 is the first canonical vector. In the inexact version of GMRES the approximation x_m is computed in the same manner using the matrices V_m and H_{m+1} satisfying the inexact Arnoldi relation (2.3). The corresponding residual is $r_m = r_0 - Ax_m$, and we call this the "true residual". Using the Arnoldi relation (2.2), it follows that

$$r_0 - Ax_m = r_0 - AV_m y_m = r_0 - V_{m+1}H_{m+1}y_m,$$

Thus, the "computed residual" $\tilde{r}_m = r_0 - V_{m+1}H_{m+1}y_m$ is built without a new matrix-vector product with the matrix A. The following result from [16], which was developed for nonsingular systems, provides conditions so that for a given $\epsilon > 0$, $\|r_m - \tilde{r}_m\| \leq \epsilon$. In the next section we explore similar conditions for singular systems.

THEOREM 2.1. *Let $\epsilon > 0$. Let $r_m = r_0 - Ax_m$, $\tilde{r}_m = r_0 - V_{m+1}H_{m+1}y_m$ be the true and computed residuals after m iterations of the inexact GMRES method, respectively, with y_m being the solution of* (2.4). *If for all $k = 1, ..., m$,*

$$\|E_k\| \leq \frac{\sigma_m(H_m)}{m} \frac{1}{\|\tilde{r}_{k-1}\|} \epsilon \,, \tag{2.5}$$

then $\|r_m - \tilde{r}_m\| \leq \epsilon$.

3. Inexact Krylov subspace method for singular systems. In the nonsingular case, it is customary to approximate the minimum singular value $\sigma_m(H_m)$ (which is usually not available *a priori*) in (2.5) with some estimate of the minimum singular value $\sigma_n(A)$; see, e.g., [16]. This is natural, since the first m rows of H_m represent the restriction of the operator A to $\mathcal{R}(V_m)$. (In the exact case $\mathcal{R}(V_m) = \mathcal{K}_m(A, v_1)$.) We should emphasize here that in many cases it is possible to estimate the value of $\sigma_n(A)$ from the problem at hand. We shall see later that this estimate does not have to be a necessarily good one for the inexact method to work.

Let us note that in the singular case, we have to consider two different situations. First, if for all $k = 1, \ldots, m$, we have that $\mathcal{R}(E_k) \subset \mathcal{R}(A^{k_0})$, with $k_0 = \mathrm{ind}(A)$, then we can think of the whole process taking place in the subspace $\mathcal{R}(A^{k_0})$. Since the restriction of A to this subspace, $A\,|_{\mathcal{R}(A^{k_0})}$, is nonsingular, we can apply the results from [16], and in particular Theorem 2.1, directly. Therefore, we can approximate $\sigma_m(H_m)$ in (2.5) with $\sigma_{\min}(A\,|_{\mathcal{R}(A^{k_0})}) = \sigma_{\min}(A)$, and use some estimate of the latter quantity.

On the other hand, if one cannot guarantee that $\mathcal{R}(E_k) \subset \mathcal{R}(A^{k_0})$, for all $k = 1, \ldots, m$, then, the results for nonsingular matrices cannot be used, and some new theory needs to be developed. This is what we do in the sequel, and it is one of the contributions of this paper.

A careful look at [16] reveals that Theorem 2.1 and its proof hold verbatim when A is singular. But in this case one has that $\sigma_n(A) = 0$. It follows then, that as m increases, $\sigma_m(H_m)$ tends to zero. Thus, for the singular case with arbitrary E_k, the use of $\sigma_n(A)$ as an estimate for $\sigma_m(H_m)$ is not practical. As an alternative, we develop a lower bound for $\sigma_m(H_m)$, using among other quantities, $\sigma_\ell(A)$, where ℓ is the rank of A. Of course, when A is nearly singular, we would use for ℓ the effective rank of A.

We mention here that the fact that $\sigma_m(H_m)$ may be small, indicates that the allowable errors in matrix-vector products may not be as large as in the nonsingular case. Nevertheless, as the iterations progress, i.e., as $\|\tilde{r}_{k-1}\|$ decreases, this allowable error may be allowed to grow; see the experiments in section 5 for illustrations of this.

We begin with a lemma which is a consequence of the Schur unitary triangularization theorem; see, e.g, [9, Theorem 2.3].

LEMMA 3.1. *Let A be a matrix with* $\mathrm{ind}(A) = 1$, *then A can be written as $A = UTU^*$ where U is a unitary matrix in $\mathbb{C}^{n\times n}$ and T is upper triangular with the last $n - \ell$ rows of zeros, where*

$\mathrm{rank}(A) = \mathrm{rank}(T) = \ell$, *i.e., T has the following form*

$$T = \begin{bmatrix} t_{11} & t_{12} & t_{13} & \cdots & \cdots & t_{1n} \\ 0 & t_{22} & t_{23} & \cdots & \cdots & t_{2n} \\ \vdots & & \ddots & \cdots & \cdots & \vdots \\ 0 & \cdots & 0 & t_{\ell\ell} & \cdots & t_{\ell n} \\ 0 & \ddots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & \cdots & \cdots & 0 \end{bmatrix}. \tag{3.1}$$

Proof. Consider a Schur factorization of A, i.e., $A = UTU^*$, with U unitary, and T upper triangular and let $\ell = \mathrm{rank}(A)$. One can choose the decomposition so that $t_{k,k} = 0$ for $k = \ell + 1, \ldots, n$. We want to show that all entries in these last $n - \ell$ rows are zeros. Since the last $n - \ell$ diagonal entries are zeros, we have that the matrix T^k has the last $n - \ell$ rows consisting of zeros for $k \geq n - \ell$, i.e., $\mathrm{rank}(T^{n-\ell}) = \ell$.

By hypothesis, $\mathrm{ind}(A) = 1$, i.e., $\mathcal{N}(A) = \mathcal{N}(A^k)$ for every integer $k \geq 1$. Let $X = [x_1,\ x_2,\ \ldots,\ x_{n-\ell}]$ be such that its columns span $\mathcal{N}(A)$, since $A = UTU^*$, then the columns of U^*X span $\mathcal{N}(T)$, and since U is a unitary matrix, $\mathrm{rank}(X) = \mathrm{rank}(U^*X)$. Therefore we have that $\dim\mathcal{N}(A) = \dim\mathcal{N}(T)$, and since $\mathcal{N}(A) = \mathcal{N}(A^k)$, we have that $\dim\mathcal{N}(T) = \dim\mathcal{N}(T^k)$ for $k \geq 1$, and in particular for $k = n - \ell$.

According to the rank-nullity theorem, $\dim\mathcal{N}(T^{n-\ell}) + \dim\mathcal{R}(T^{n-\ell}) = \dim\mathcal{N}(T) + \dim\mathcal{R}(T)$, we thus obtain $\dim\mathcal{R}(T^{n-\ell}) = \dim\mathcal{R}(T)$. Since the last $n - \ell$ diagonal entries of $T^{n-\ell}$ are zeros, we have that the triangular matrix T also has all zero entries in the last $n - \ell$ rows. Therefore T is of the form (3.1). Also, we conclude that $\mathrm{rank}(A) = \mathrm{rank}(T)$, and thus we have that $\ell = \mathrm{rank}(T) = \mathrm{rank}(A)$. ∎

The following lemma provides a matrix $\tilde{A}$ which is nonsingular and satisfies the same inexact Arnoldi relation (2.3) as does A, i.e., with the same matrices E_i and V_i. We will conceptually think of this nonsingular matrix $\tilde{A}$ and use it to estimate $\sigma_m(H_m)$; see (2.5).

LEMMA 3.2. *Let A be a singular matrix such that* $\mathrm{ind}(A) = 1$, *and let $r_0 \in \mathcal{R}(A)$. Then there exists a nonsingular matrix $\tilde{A}$ such that $A^i r_0 = \tilde{A}^i r_0$ for every positive integer i.*

Proof. Let $\ell = \mathrm{rank}(A)$, suppose that the Schur decomposition of A is $A = UTU^*$, and by Lemma 3.1 we further assume that T is of the form (3.1). Let $u_1, \ldots, u_n$ be the columns of the unitary matrix U. For every r_0 in $\mathcal{R}(A)$, $r_0 = Aw$ for some w, and we can write $r_0 = Aw = UTU^*w$. Since T has rows $\ell + 1, \ldots, n$ consisting of zeros, $a := TU^*w$ is a column vector with the last $n - \ell$ entries being zeros, i.e., of the form $a = [\alpha_1, \ldots, \alpha_\ell, 0, \ldots, 0]^T$, with α_j complex numbers. Thus, we have the following representation $r_0 = \sum_{j=1}^{\ell} \alpha_j u_j$. Let $\tilde{A} = U\tilde{T}U^*$, where $\tilde{T} = T + \lambda K$ with $\lambda \in \mathbb{C}$, $\lambda \neq 0$, arbitrary and $K = diag(0, \ldots, 0, 1, \ldots, 1)$, i.e., a diagonal matrix with the last $n - \ell$ entries being ones and the rest of them being zeros. That is, $\tilde{T}$ is of the following form

$$\tilde{T} = \begin{bmatrix} t_{11} & t_{12} & t_{13} & \cdots & \cdots & \cdots & t_{1n} \\ 0 & t_{22} & t_{23} & \cdots & \cdots & \cdots & t_{2n} \\ \vdots & & \ddots & \cdots & \cdots & \cdots & \vdots \\ 0 & \cdots & 0 & t_{\ell\ell} & \cdots & \cdots & t_{\ell n} \\ 0 & \cdots & 0 & 0 & \lambda & 0 & 0 \\ 0 & \cdots & 0 & \ddots & 0 & \ddots & 0 \\ 0 & \cdots & 0 & 0 & 0 & 0 & \lambda \end{bmatrix}. \tag{3.2}$$

Observe that $A^i r_0 = UT^i a$, and similarly $\tilde{A}^i r_0 = U\tilde{T}^i U^* r_0 = U\tilde{T}^i a$. Because of the structure of T, $\tilde{T}$, and α we have $T^i a = \tilde{T}^i a$, therefore $A^i r_0 = \tilde{A}^i r_0$, for every $i \geq 1$. ∎

We now obtain a similar result for A with $\mathrm{ind}(A) = k_0$, and $k_0 > 1$.

PROPOSITION 3.3. *Let A be a singular matrix such that* $\mathrm{ind}(A) = k_0$, $k_0 > 1$ *and* $r_0 \in \mathcal{R}(A^{k_0})$. *Then there exists a nonsingular matrix $\tilde{A}$ such that $A^i r_0 = \tilde{A}^i r_0$ for every positive integer i.*

Proof. Suppose $\mathrm{rank}(A^{k_0}) = m$, and $\mathrm{rank}(A) = \ell$, with m and ℓ not necessary equal to each other. Then $A = UT_1U^*$ with $T_1 = T + J$, where T is upper triangular with $t_{jj} \neq 0$ for $j = 1, \ldots, \ell$, and the last $n - \ell$ rows having all zeros, as in (3.1), and J is an upper bidiagonal matrix which has the first ℓ rows of zeros and Jordan blocks with 0 as the eigenvalue and at most k entries with value 1 on the superdiagonal consecutively in the last $n - \ell$ rows, i.e., T_1 is of the following form

$$T_1 = \begin{bmatrix} t_{11} & t_{12} & t_{13} & \cdots & \cdots & \cdots & \cdots & t_{1n} \\ 0 & t_{22} & t_{23} & \cdots & \cdots & \cdots & \cdots & t_{2n} \\ \vdots & & \ddots & & & & & \vdots \\ 0 & \cdots & 0 & t_{\ell\ell} & \cdots & \cdots & \cdots & t_{\ell n} \\ 0 & \cdots & 0 & 0 & 0 & 1 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & 0 & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & \cdots & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We can now use the same construction as in Lemma 2.1. Our goal is to show that $m = \ell$, i.e., that we can have the representation $r_0 = \sum_{j=1}^{\ell} \alpha_j u_j$, with u_j being the columns of U. Since U is a unitary matrix, we have that $A^{k_0} = UT_1^{k_0}U^*$. Since $\mathrm{ind}(A) = k_0$, we have $\mathrm{ind}(A^{k_0}) = 1$. Then since $\mathrm{rank}(A^{k_0}) = m$, by Lemma 3.1, we have that ${T_1}^{k_0}$ has the last $n - m$ rows of zeros. Therefore ${T_1}^{k_0}$ has m nonzero diagonal entries, but ${T_1}^{k_0}$ also has ℓ nonzero diagonals. Therefore $m = \ell$. Since $r_0 \in \mathcal{R}(A^{k_0})$, therefore we have $r_0 = \sum_{j=1}^{\ell} \alpha_j u_j$. Let $\tilde{A} = U\tilde{T}_1U^*$ with $\tilde{T}_1 = T_1 + \lambda K$, where $\lambda \in \mathbb{C}$, $\lambda \neq 0$, arbitrary, and $K = \mathrm{diag}(0, \ldots, 0, 1, \ldots, 1)$ is a diagonal matrix with the first ℓ entries being zeros and the last $n - \ell$ entries being ones, i.e., $\tilde{T}_1$ has the following form

$$\tilde{T}_1 = \begin{bmatrix} t_{11} & t_{12} & t_{13} & \cdots & \cdots & \cdots & t_{1n} \\ 0 & t_{22} & t_{23} & \cdots & \cdots & \cdots & t_{2n} \\ \vdots & & \ddots & & & & \vdots \\ 0 & \cdots & 0 & t_{\ell\ell} & \cdots & \cdots & t_{\ell n} \\ 0 & \ddots & 0 & 0 & \lambda & 1 & 0 \\ 0 & \cdots & 0 & 0 & 0 & \ddots & 1 \\ 0 & \cdots & 0 & 0 & 0 & 0 & \lambda \end{bmatrix}.$$

With this construction, we have that $T_1^i r_0 = \tilde{T}_1^i r_0$ for every integer i. Since $A = UT_1U^*$ and $\tilde{A} = U\tilde{T}_1U^*$, we have that $A^i r_0 = \tilde{A}^i r_0$ for every $i \geq 1$. ∎

We use the form of the nonsingular matrices $\tilde{A}$ constructed in Lemma 3.2 and Proposition 3.3 to show that as long as $r_0 \in \mathcal{R}(A^{k_0})$, where $k_0 = \mathrm{ind}(A)$, the solution of a system of the form $\tilde{A}x = r_0$ is also a solution of the original system $Ax = r_0$, defined by the singular matrix A.

PROPOSITION 3.4. *Let A be a singular matrix, and let* $\mathrm{rank}(A) = \ell$, $\mathrm{ind}(A) = k_0$ *($k_0 \geq 1$), and $r_0 \in \mathcal{R}(A^{k_0})$. Let $A = UTU^*$ be a Schur decomposition, with the last $n - \ell$ diagonal entries of T being zero. Let $\tilde{A} = U\tilde{T}U^*$, with $\tilde{T}$ constructed as in Lemma 3.2 or Proposition 3.3, depending on* $\mathrm{ind}(A)$. *If x is such that $\tilde{A}x = r_0$, then $Ax = r_0$.*

Proof. Since $\tilde{A}x = r_0$, and $r_0 = \sum_{j=1}^{\ell} \alpha_j u_j = Ua$, we have that $\tilde{T}U^*x = a$. Let $y = U^*x$, i.e., $\tilde{T}y = a$. Since a has the last $n - \ell$ entries being zeros, and $\tilde{T}$ is an upper triangular nonsingular matrix,

we have that y also has the last $n-\ell$ entries being zeros. Therefore we can write $Ty = \tilde{T}y = a$ and as a consequence, $Ax = UTU^*x = UTy = U\tilde{T}y = Ua = r_0$. ∎

REMARK 3.5. *We point out that the condition $r_0 \in \mathcal{R}(A^{k_0})$ is essential. In the case of index one ($k_0 = 1$), such as for irreducible Markov chains, and for $z_0 = 0$, i.e., $r_0 = b$, this simply means that the system needs to be consistent.*

4. Estimate of a threshold parameter for inexact matrix-vector products. For any singular matrix A, in the previous section we constructed a nonsingular matrix $\tilde{A}$, such that if $r_0 \in \mathcal{R}(A^{k_0})$, with $k_0 = \mathrm{ind}(A)$, the solution of $\tilde{A}x = r_0$, is also a solution of the the system $Ax = r_0$, which is the one we are interested in. With this in mind our approach consists then of estimating $\sigma_m(H_m)$ to be used in the criterion (2.5), by the value of $\sigma_n(\tilde{A})$. Since $\tilde{A}$ or $\tilde{T}$ are usually not available, we want in turn to estimate $\sigma_n(\tilde{A})$ using other quantities, especially by using $\sigma_\ell(A)$, the smallest positive singular value of A with $\ell = \mathrm{rank}(A)$. As mentioned in the introduction, in the case of nearly singular A, we would use for ℓ the effective rank of A.

Let $A = UTU^*$ be the Schur decomposition of A, then it easily follows that A, T, and T^* have the same singular values. Let $\tilde{A} = U\tilde{T}U^*$, where $\tilde{T}$ is composed using Lemma 3.2 or Proposition 3.3 according to $\mathrm{ind}(A)$. We conclude also that $\tilde{A}$, $\tilde{T}$, and $\tilde{T}^*$ have the same singular values.

LEMMA 4.1. *Let $\lambda \in \mathbb{C}$, $\lambda \neq 0$, be arbitrary. Let $A = UTU^*$ be a Schur decomposition of A, let $\mathrm{rank}(A) = n-1$ and $\mathrm{ind}(A) = 1$. Let $\hat{T}$ consist of the first $n-1$ rows of T, and let the row vector $v = \lambda e_n^*$, where e_n is the nth canonical vector. Furthermore, let $\|s_1\| = \min_{x\in\mathbb{R}^{n-1}} \|\hat{T}^*x - v^*\|$. Let $\tilde{T} = \begin{bmatrix} \hat{T} \\ v \end{bmatrix}$, an $n \times n$ matrix which has the form (3.2) (for $\ell = n-1$), and let $\tilde{A} = U\tilde{T}U^*$. Then,*

$$\sigma_n(\tilde{A}) \geq \gamma\ \sigma_{n-1}(A), \quad \text{where} \quad \gamma = \|s_1\|/(\sigma_1(A) + |\lambda|). \tag{4.1}$$

Proof. We have the following identity:

$$\tilde{T}\tilde{T}^* = \begin{bmatrix} \hat{T} \\ v \end{bmatrix} \cdot \begin{bmatrix} \hat{T}^* & v^* \end{bmatrix} = \begin{bmatrix} \hat{T}\hat{T}^* & \hat{T}v^* \\ v\hat{T}^* & vv^* \end{bmatrix}. \tag{4.2}$$

Because of (4.2), one has the following:

$$det(\tilde{T}\tilde{T}^*) = \prod_{i=1}^{n} \sigma_i^2(\tilde{T}),$$

$$det \begin{bmatrix} \hat{T}\hat{T}^* & \hat{T}v^* \\ v\hat{T}^* & vv^* \end{bmatrix} = \det(\hat{T}^*\hat{T})\beta_1,$$

where β_1 is the scalar Schur complement of vv^*, i.e.,

$$\beta_1 = vv^* - v\hat{T}^*(\hat{T}\hat{T}^*)^{-1}\hat{T}v^* = vv^* - v\hat{T}^*y = v(v^* - \hat{T}^*y) = vs_1 = \|s_1\|^2,$$

with $\|s_1\| = \min_{x\in\mathbb{R}^\ell} \|\hat{T}^*x - v^*\| = \|\hat{T}^*y - v^*\|$. We note that since $s_1 = v^* - \hat{T}^*y$ and y is the minimizer, with $\hat{T}$ upper triangular, v^* is decomposed into two components, s_1 and $\hat{T}^*y$, i.e., one component is in $\mathcal{R}(\hat{T}^*)$ the other component is in $\mathcal{N}(\hat{T}^*)$. It follows that $v \cdot s_1 = \|s_1\|^2$; cf. [3], [11]. Therefore we have the following identity,

$$\prod_{i=1}^{n} \sigma_i^2(\tilde{T}) = \prod_{i=1}^{n-1} \sigma_i^2(\hat{T})\|s_1\|^2. \tag{4.3}$$

From [23, sec. 47] applied to (4.2), it follows that the singular values of $\hat{T}$ interlace those of $\tilde{T}$ i.e., $\sigma_i(\tilde{T}) \geq \sigma_i(\hat{T}) \geq \sigma_{i+1}(\tilde{T})$, $i = 1, \ldots, n-1$. Therefore by (4.3), we have that

$$\frac{\sigma_n(\tilde{T})}{\sigma_{n-1}(\hat{T})} = \frac{\sigma_{n-2}(\hat{T})}{\sigma_{n-1}(\tilde{T})} \cdots \frac{\sigma_1(\hat{T})}{\sigma_2(\tilde{T})} \frac{\|s_1\|}{\sigma_1(\tilde{T})} \geq \frac{\|s_1\|}{\sigma_1(\tilde{T})} \ . \tag{4.4}$$

Replacing in (4.4) the following quantities $\sigma_n(\tilde{T}) = \sigma_n(\tilde{A})$, $\sigma_{n-1}(\hat{T}) = \sigma_{n-1}(T) = \sigma_{n-1}(A)$, and $\sigma_1(\tilde{A}) = \sigma_1(\tilde{T})$ we obtain

$$\frac{\sigma_n(\tilde{A})}{\sigma_{n-1}(A)} \geq \frac{\|s_1\|}{\sigma_1(\tilde{A})} \ .$$

Finally, since $\sigma_1(\tilde{A}) = \sigma_1(\tilde{T}) = \|\tilde{T}\| \leq \|T\| + |\lambda| = \sigma_1(T) + |\lambda| = \sigma_1(A) + |\lambda|$, we have that (4.1) holds. ∎

The bound (4.1) is the basis for our proposed threshold η_k when A has rank $n-1$ and index 1. We thus estimate $\sigma_m(H_m)$ by $\gamma\sigma_{n-1}(A)$, in (2.5) and use

$$\eta_k = \frac{\gamma\sigma_{n-1}(A)}{m^*} \frac{1}{\|\tilde{r}_{k-1}\|},$$

where m^* is the maximum number of iterations.

We can extend Lemma 4.1 to the case that $\operatorname{rank}(A) = \ell < n-1$, although in this case, the lower bound obtained is not so easily computable.

LEMMA 4.2. *Let $\lambda \in \mathbb{C}$, $\lambda \neq 0$, be arbitrary. Let $A = UTU^*$ be a Schur decomposition of A, let $\operatorname{rank}(A) = \ell$ and $\operatorname{ind}(A) = 1$. Assume further $\ell < n-1$. Let $\hat{T}$ consist of the first ℓ rows of matrix T, let*

$$\hat{T}_i = \begin{bmatrix} \hat{T} \\ \lambda K_i \end{bmatrix}, \tag{4.5}$$

where $K_i = [O \ I_i]$ and I_i is the $i \times i$ identity matrix. Let $\tilde{T} = \hat{T}_{n-\ell}$ and $\tilde{A} = U\tilde{T}U^*$. Let v_i be the ith row of λK_i and $\|s_i\| = \min_{x \in \mathbb{R}^{l+i-1}} \|\hat{T}_i^* x - v_i^*\|$, then we have

$$\sigma_n(\tilde{A}) \geq \gamma\sigma_\ell(A), \quad \text{where} \quad \gamma = \frac{\|s_1\|\|s_2\| \ldots \|s_{n-\ell}\|}{(\sigma_1(A) + (n-l)|\lambda|)^{n-\ell}}. \tag{4.6}$$

Proof. By applying the inequality (4.4) in the proof of Lemma 4.1 to the appropriate matrices and vectors, we have the following inequalities,

$$\frac{\sigma_{\ell+1}(\hat{T}_1)}{\sigma_\ell(T)} \geq \frac{\|s_1\|}{\sigma_1(\hat{T}_1)}, \ \frac{\sigma_{\ell+2}(\hat{T}_2)}{\sigma_{\ell+1}(\hat{T}_1)} \geq \frac{\|s_2\|}{\sigma_1(\hat{T}_2)}, \ \cdots, \ \frac{\sigma_n(\tilde{T})}{\sigma_{n-1}(\hat{T}_{n-1-\ell})} \geq \frac{\|s_{n-\ell}\|}{\sigma_1(\tilde{T})}. \tag{4.7}$$

Take the product of the inequalities in (4.7), where cancellation occurs, and use the fact that $\sigma_1(\tilde{T}) \geq \sigma_1(\hat{T}_{n-\ell-1}) \geq \cdots \geq \sigma_1(\hat{T}_\ell)$; it follows that

$$\frac{\sigma_n(\tilde{T})}{\sigma_\ell(T)} \geq \frac{\|s_1\| \ldots \|s_{n-\ell}\|}{\sigma_1(\tilde{T})^{n-\ell}}. \tag{4.8}$$

Finally, since $\sigma_1(\tilde{A}) = \sigma_1(\tilde{T}) = \|\tilde{T}\| \leq \|T\| + (n-\ell)|\lambda| = \sigma_1(T) + (n-\ell)|\lambda| = \sigma_1(A) + (n-\ell)|\lambda|$. Therefore we have (4.6). ∎

If $\operatorname{ind}(A) > 1$, we can obtain analogous results to Lemmas 4.1 and 4.2. The following result gives a lower bound for $\sigma_n(\tilde{A})$ in that general case. Since the proof of Theorem 4.3 is very similar to those of Lemmas 4.1 and 4.2, it is omitted.

THEOREM 4.3. *Let $\lambda \in \mathbb{C}$, $\lambda \neq 0$, be arbitrary. Let $A = UTU^*$ be a Schur decomposition of A, and* $\mathrm{rank}(A) = \ell$ *and* $\mathrm{ind}(A) = k$ *with $k > 1$. Let $\bar{T}_i$ be the first $\ell + i$ rows of matrix T, let*

$$\hat{T}_i = \bar{T}_i + \begin{bmatrix} O \\ \lambda K_i \end{bmatrix}, \tag{4.9}$$

where $K_i = [O\ I_i]$ and I_i is an $i \times i$ identity matrix, and let v_i be the ith row of λK_i. Let $\tilde{T}_1 = \hat{T}_{n-\ell}$ and $\tilde{A} = U\tilde{T}_1U^*$. Let $\|s_i\| = \min_{x\in\mathbb{R}^{\ell+i-1}} \|\hat{T}_i^*x - v_i^*\| = \|\hat{T}_i^*y - v_i^*\|$. Let

$$\gamma = \|s_1\|\|s_2\| \ldots \|s_{n-\ell}\|/(\sigma_1(A) + (n-\ell)\sqrt{\lambda^2+1})^{n-\ell},$$

then we have

$$\sigma_n(\tilde{A}) > \gamma\sigma_\ell(A). \tag{4.10}$$

The lower bounds (4.1), (4.6), and (4.10) include in the definition of γ the norm of one or more solutions of least squares problems using a portion of the triangular matrix T obtained from the Schur decomposition of A. While T may be easily available for some problems, we do not assume this. Instead, we simply estimate $\|s_1\| \approx 1$ in (4.1), for the case $\ell = n - 1$, and use this value in our calculations, reported in the next section. For $\ell << n$, we have not attempted to estimate these values.

5. Numerical Experiments. We begin with a simple example, which illustrates both the possible benefits as well as the limitations of the use of inexact matrix-vector products in the context of minimal residual methods.

Example 5.1. Consider the following bidiagonal matrix A of order $n = 100$ with $\mathrm{ind}(A) = 5$. Let $A = D + 0.1C$, where $D = \mathrm{diag}(0, 0, 0, 0, 0, 0.01, \ldots, 1)$ with the 95 nonzeros being equally spaced between 0.01 and 1, and C having only nonzeros in the first upper off-diagonal, consisting of all ones.

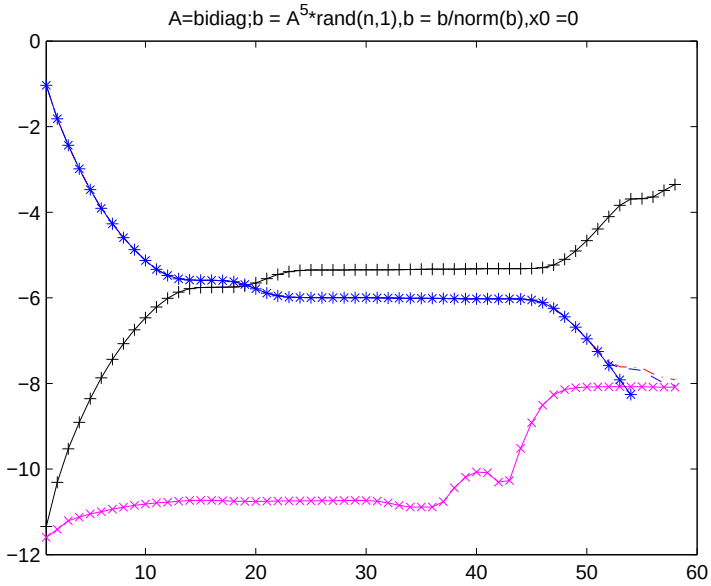

FIG. 5.1. *Convergence curves for GMRES and inexact GMRES. Matrix of Example 5.1. $b \in \mathcal{R}(A^5)$*

Due to the restrictions of our theory, we choose a right hand side $b \in \mathcal{R}(A^5)$ (and $x_0 = 0$). To that end we take a random vector c, we let $b_1 = A^5 c$, and normalize it to obtain $b = b_1/\|b_1\|$. We run GMRES with both exact and inexact matrix-vector products. For the inexact case, we used at iteration k, $A + E_k$, where

$$E_k = \frac{\sigma_{\min}(A)\epsilon}{m\|\tilde{r}_{k-1}\|} R_k, \tag{5.1}$$

where R_k is a random matrix with $\|R_k\| = 1$, cf. (2.5). Here we have chosen $m = 100$, and $\epsilon = 10^{-8}$. In Figure 5.1 we show three GMRES convergence curves (which are mostly on top of each other): the residual norm of exact GMRES (full lines with a symbol $*$), the true and computed residuals of inexact GMRES (dashed and dash-doted lines, respectively). We also report the norm $\|E_k\|$ used (with the symbol $+$), and the norm of the difference between the computed and the true residual (with the symbol $\times$). As it can be appreciated, the computed and true residuals remain close to each other, while the norm of the error in the matrix-vector product is allowed to grow (representing potentially less cost per iteration). Furthermore, the convergence curve of the inexact method mimics well the convergence curve of the exact one.

In Figure 5.2, we illustrate the effect of perturbing the right hand side with a random perturbation of the order of 10^{-10}. It can be appreciated that even if the right hand side is not exactly in $\mathcal{R}(A^{k_0})$, but close enough to it (possibly accounting for round-off errors), then, the inexact methods work well, with little deterioration as compared with the unperturbed right hand side.

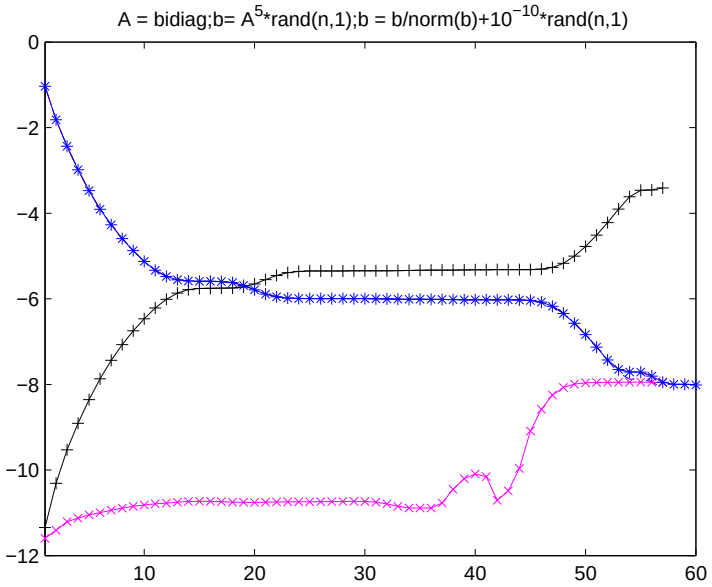

FIG. 5.2. *Convergence curves for GMRES and inexact GMRES. Matrix of Example 5.1. Right hand side with a perturbation of the order of* 10^{-10}

The plots in Figure 5.3 correspond to a normalized random right hand side, i.e., a vector which is not near $\mathcal{R}(A^5)$. As it can be appreciated, for this right hand side, the inexact method does not work, in the sense that the difference between the true and computed residual grows to unacceptable levels. This fact confirms the comments in Remark 3.5 that the condition $r_0 \in \mathcal{R}(A^{k_0})$ is essential.

In the next three examples we used sample matrices taken from other papers on GMRES for singular

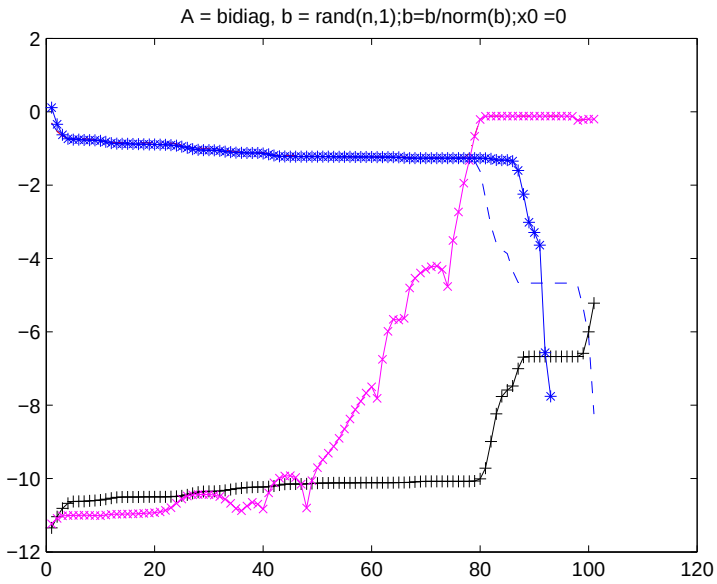

FIG. 5.3. *Convergence curves for GMRES and inexact GMRES. Matrix of Example 5.1. Right hand side is not near* $\mathcal{R}(A^5)$

systems, so that we can see how the inexact method works on them, and also with the purpose of analyzing how far are the solutions produced by the inexact method from those produced from (exact) GMRES.

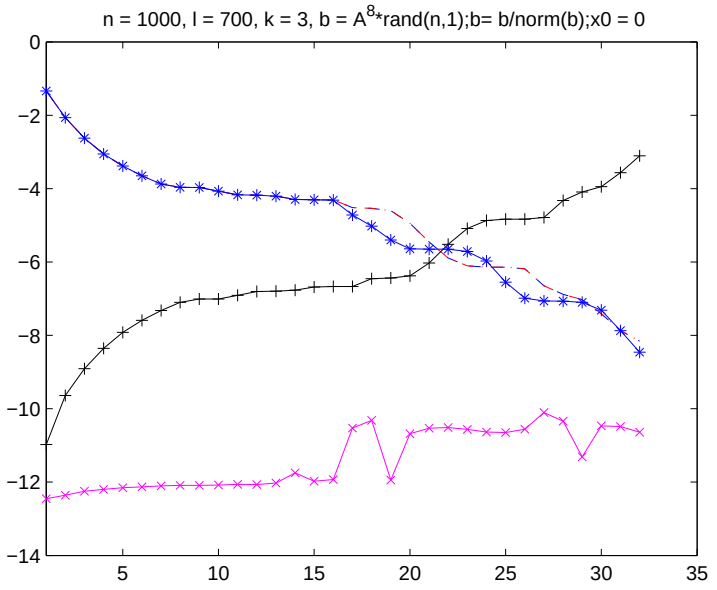

FIG. 5.4. *Convergence curves for GMRES and inexact GMRES. Matrix of example 5.2.* $b \in \mathcal{R}(A^8)$

Example 5.2. The matrix in this example comes from [13]. Consider an $n \times l$ rectangular matrix,

$l < n$, of the form

$$\tilde{A} = \begin{bmatrix} \tilde{A}_{11} & O \\ \tilde{A}_{21} & \tilde{A}_{22} \end{bmatrix},$$

where $\tilde{A}_{11} \in \mathbb{R}^{(n-k)\times(l-k)}$ and $\tilde{A}_{11} = L + 10I$, where L is a random lower triangular matrix, and $\tilde{A}_{21} \in \mathbb{R}^{k\times(l-k)}$ and $\tilde{A}_{22} \in \mathbb{R}^{k\times k}$ are random matrices generated by the MATLAB function `rand`. Finally, the matrix is completed with l columns of zeros as follows, assuring its singularity,

$$A = \begin{bmatrix} \tilde{A} & O \end{bmatrix}.$$

We choose $n = 1000$, $l = 700$ and $k = 3$. Since $\mathrm{ind}(A) = 8$, therefore we construct $b \in \mathcal{R}(A^8)$ as follows (and use $x_0 = 0$). We construct a random vector c, compute A^8c, and normalize it. The convergence curves, using the same conventions as the plots of the Example 5.1, are presented in Figure 5.4. In Figure 5.5 we present the result of perturbing the right hand side with a random perturbation of the order of 10^{-9}. As in the previous example, the effect of this large perturbation, while noticeable, does not prevent the inexact method from producing a solution with the computed residual close to the true residual.

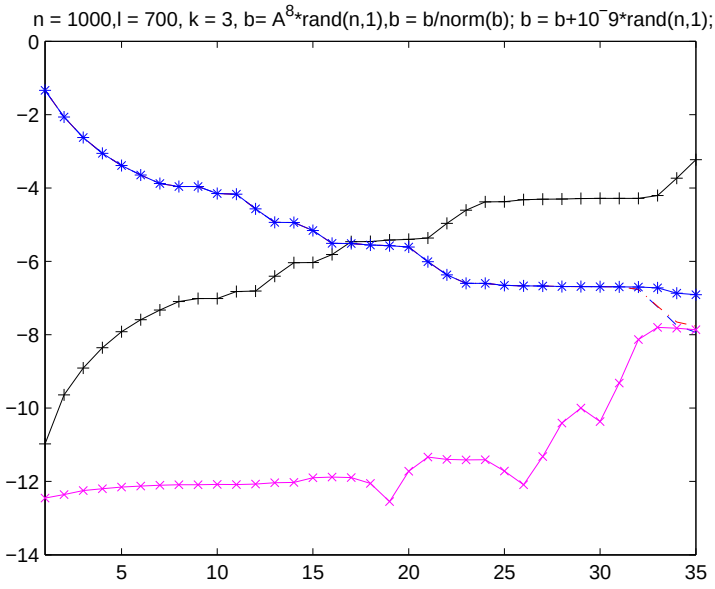

FIG. 5.5. *Convergence curves for GMRES and inexact GMRES. Matrix of Example 5.2. Right hand side with a perturbation of the order of* 10^{-9}

The next two examples involve discretizations of boundary value problems for partial differential equations of the following form

$$\Delta u + d\frac{\partial u}{\partial x_1} = f(x), \quad x = (x_1, x_2) \in \Omega = [0,1] \times [0,1], \tag{5.2}$$

where d is a constant and f is a given function; cf. [5]. In the experiments reported here, we discretized (5.2) with the usual second-order centered differences on a 50×50 mesh of equally spaced discretization points, so that the resulting linear system of order $n = 2500$.

Example 5.3. We first consider periodic boundary conditions. The matrix A is thus of the following form

$$A = \begin{bmatrix} T_m & I_m & & I_m \\ I_m & \ddots & \ddots & \\ & \ddots & \ddots & I_m \\ I_m & & I_m & T_m \end{bmatrix}, \quad \text{where } T_m = \begin{bmatrix} -4 & \alpha_+ & & & \alpha_- \\ \alpha_- & -4 & \alpha_+ & & \\ & \ddots & \ddots & \ddots & \\ & & \alpha_- & -4 & \alpha_+ \\ \alpha_+ & & & \alpha_- & -4 \end{bmatrix}, \tag{5.3}$$

$m = \sqrt{n} = 50$, $h = 1/m$, and $\alpha_\pm = 1 \pm dh/2$. Note that $\mathcal{N}(A) = span(1, 1, \ldots, 1)^T$ and $\text{ind}(A) = 1$. We choose $d = 1$ and as right hand side, the matrix A times a random vector x, so, on the one side we know the solution of the problem, and on the other, we guarantee that the system is consistent.

We performed two runs. One using GMRES, stopping when the residual norm is below 10^{-8}, and we call this solution x^G. We also ran inexact GMRES, using at each step $A + E_k$, where E_k is as in (5.1), with $m = 200$, stopping when the residual norm is below the same tolerance 10^{-8}. We call the latter approximation x^{IG}. The number of iterations for both runs were 200. In Table 5.1 we show the norm of the difference between the different solutions. One can see that the inexact method produces an approximation which is as close to the true solution as that with (exact) GMRES.

Example 5.4. In this experiment we imposed homogeneous Neumann boundary conditions, i.e., $\partial u(x)/\partial v = 0$ for $x \in \partial\Omega$. The matrix A is now given by

$$A = \begin{bmatrix} T_m & 2I_m & & & \\ I_m & T_m & I_m & & \\ & \ddots & \ddots & \ddots & \\ & & I_m & T_m & I_m \\ & & & 2I_m & T_m \end{bmatrix}, \quad \text{where } T_m = \begin{bmatrix} -4 & 2 & & & \\ \alpha_- & -4 & \alpha_+ & & \\ & \ddots & \ddots & \ddots & \\ & & \alpha_- & -4 & \alpha_+ \\ & & & 2 & -4 \end{bmatrix}. \tag{5.4}$$

As in the previous example, A has a one-dimensional null space, and has index one. As before, we choose $d = 1$, and fix a random vector x. We make the same runs as described in Example 5.3. The number of iterations for both runs were 147. In Table 5.1 we show the norm of the difference between the different solutions. As in the previous example, one can appreciate that the inexact method produces an approximation which is as close to the true solution as that with (exact) GMRES.

	Example 5.3	Example 5.4
$\|x - x^{IG}\|$	1.9486×10^{-6}	2.4886×10^{-7}
$\|x - x^{G}\|$	1.9481×10^{-6}	2.4886×10^{-7}
$\|x^G - x^{IG}\|$	4.8081×10^{-10}	3.7164×10^{-12}

TABLE 5.1
norm of the difference of inexact solution and true solution

In the rest of the paper, we report on several numerical experiments using Markov Chain matrices from the Marca collection [20], using an inexact preconditioner. The goal of these experiments is to present some real problems, where inexact GMRES can be used, and look at potential computational savings. All these matrices have $\text{ind}(A) = 1$.

For the solution of the linear system (2.1), we use a (right) preconditioner, i.e., we solve the system $AP^{-1}y = b$. We choose the block preconditioner P, in a way similar to Benzi and Uçar [2], as we explain in what follows.

We partition A by using the Metis software [10] into the following 2×2 block form:

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \tag{5.5}$$

where the order of A_{22} is much smaller than that of A_{11}, and where the $(1,1)$ block is block diagonal with p diagonal blocks, i.e., of the following form

$$A_{11} = \begin{bmatrix} A_1 & 0 & \cdots & \cdots \\ 0 & A_2 & 0 & \cdots \\ & & \ddots & \\ 0 & \cdots & 0 & A_p \end{bmatrix}. \tag{5.6}$$

We then choose the preconditioner as

$$P = \begin{bmatrix} A_{11} & A_{12} \\ 0 & \tilde{S} \end{bmatrix}, \tag{5.7}$$

where the matrix $\tilde{S}$ approximates the Schur complement $S = A_{22} - A_{21}A_{11}^{-1}A_{12}$. We use $\tilde{S} = A_{22} - A_{21}D_{11}^{-1}A_{12}$ with D_{11} diagonal, having the diagonal entries of A_{11}, following Benzi and Uçar [2], where it was shown that $\tilde{S}$ is nonsingular. At each step of the Krylov subspace method, we need to solve a linear system of the form $Pz = v$, and we do this inexactly as we explain below.

Consider v and z, partitioned into $p+1$ subvectors, conformally to (5.5)–(5.6), and denote these subvectors $v^1, v^2, \ldots, v^p, v^{p+1}$, and $z^1, z^2, \ldots, z^p, z^{p+1}$, respectively. Let also $(A_{12})_i$ denote the ith block of rows of A_{12} corresponding to A_i. In our illustrative experiments, since the size of the Schur complement is comparatively small, we solve $\tilde{S}z^{p+1} = v^{p+1}$ directly using an LU decomposition. The inexactness of the application of the matrix vector-product $AP^{-1}v$ at each iteration thus comes from the inexact solution of the p linear systems

$$A_i z^i = v^i - (A_{12})_i z_{p+1}, \quad i = 1, \ldots, p. \tag{5.8}$$

The following result adopted from [16] essentially translates Theorem 2.1 to the current situation, and provides the basis to obtain values of the tolerance $\epsilon_{i,k}$ for the solution of the systems (5.8) at the kth iteration; see further below.

PROPOSITION 5.1. *Let $\epsilon > 0$. Let $r_m = r_0 - Ax_m$, $\tilde{r}_m = b - V_{m+1}H_{m+1}y_m$ be the true and computed residuals after m iterations of the inexact GMRES method, respectively, with y_m being the solution of (2.4), and $x_m = V_m y_m$, and let P be a right preconditioner. Let us consider the inexact application of P^{-1} to v_k, i.e., we let $\tilde{z}_k \approx P^{-1}v_k$. If at each iteration $k \le m$ the inner residual*

$$p_k = v_k - P\tilde{z}_k \tag{5.9}$$

satisfies

$$\|p_k\| \le \frac{\sigma_m(H_m)}{m\|AP^{-1}\|} \frac{\epsilon}{\|\tilde{r}_{k-1}\|},$$

then $\|r_m - \tilde{r}_m\| \le \epsilon$.

Note that here our preconditioner P as in (5.7) is composed by $p+1$ blocks; see (5.6). We partition the residual accordingly and write $p_k = [\rho_{1,k}, \rho_{2,k}, \ldots, \rho_{p+1,k}]^T$. It follows from (5.9) and (5.7) that

$\rho_{i,k} = v_k^i - A_i\tilde{z}_k^i - (A_{12})_i\tilde{z}_k^{p+1}$, for $i = 1, 2, \ldots p$, and $\rho_{p+1,k} = \tilde{S}\tilde{z}_k^{p+1} - v_k^{p+1}$. Then, we have from Proposition 5.1 that if

$$\|\rho_{i,k}\| \leq \frac{\sigma_m(H_m)}{m(p+1)\|AP^{-1}\|}\frac{\epsilon}{\|\tilde{r}_{k-1}\|} := \epsilon_k (= \epsilon_{i,k}) \,, \tag{5.10}$$

for all $i = 1, 2, \ldots p+1$ and $k = 1, 2, \ldots m$, then $\|r_m - \tilde{r}_m\| \leq \epsilon$.

Following the development of the previous section, we approximate $\sigma_m(H_m)$ in (5.10) with $\tilde{\gamma}\sigma_1(AP^{-1})$, where

$$\tilde{\gamma} = 1/(\sigma_1(AP^{-1}) + 1); \tag{5.11}$$

cf. (4.1). In table 5.2 we compare the lower bound estimate $\tilde{\gamma}\sigma_{n-1}(AP^{-1})$ with $\sigma_\ell(H_\ell)$ for each of seven matrices taken from the Marca collection [20], where ℓ is the iteration number when inexact GMRES stops; see Table 5.3 (first run) below. For these small matrices, we computed P^{-1} using an LU factorization. Note that it is indeed a lower bound.

Matrix	n	$\tilde{\gamma}\sigma_{n-1}(AP^{-1})$	$\sigma_\ell(H_\ell)$
leaky530	530	0.0163	0.3423
mutex794	794	0.0245	0.3489
qnatm740	740	0.0540	0.4960
leaky1074	1074	0.0121	0.2783
leaky1618	1618	0.0125	0.3623
mutex3797	3797	0.0408	0.4573
mutex4095	4095	0.0193	0.3726

TABLE 5.2
Comparison of the lower bound $\tilde{\gamma}\sigma_{n-1}(AP^{-1})$ and $\sigma_\ell(H_\ell)$

Some remarks are in order. The fact that these lower bounds are about an order of magnitude smaller than the values of $\sigma_\ell(H_\ell)$ translates in more stringent tolerances for the solution of the systems (5.8) than strictly necessary. Nevertheless, this does not imply a big difference in computational times; cf. Table 5.3 below. As we show later, and in a way similar to what is done in the nonsingular case, instead of these lower bounds, some fixed small constant is used, with a very moderate increase in computational cost. In fact, since we have the theory for the existence of a bound for the tolerance needed for (5.8), one could simply try certain tolerances of the form $\delta\epsilon/\|\tilde{r}_{k-1}\|$, and find an appropriate value of δ.

We report on experimental results obtained with a Matlab 7.0.0 implementation on a 3GHz dual core Intel(R) Pentium(R) processor with 1GB main memory. In our first set of experiments, we present four different runs for each of the seven matrices of Table 5.3. We solve the homogeneous system $Ax = 0$, thus obtaining a (multiple of a) stationary probability distribution. The first run is using inexact GMRES with the following tolerance, which estimates the right hand side of (5.10),

$$\rho_{i,k} \leq \frac{\tilde{\gamma}\sigma_1(AP^{-1})}{m^*(p+1)\|AP^{-1}\|}\frac{\epsilon}{\|\tilde{r}_{k-1}\|} \,, \tag{5.12}$$

where $\tilde{\gamma}$ is as in (5.11), $m^* = 20$, $p = 1$ (i.e., A_{11} is a single block), and $\epsilon = 10^{-8}$. For the second run, we simply replace the left factor in (5.12) with 10^{-2} for all cases, i.e., we use

$$\rho_{i,k} \leq 10^{-2}\frac{\epsilon}{\|\tilde{r}_{k-1}\|} \,. \tag{5.13}$$

The third is taking a fixed value of $\epsilon_{i,k} = 10^{-10}$ for each of the systems (5.8), and finally solving them "exactly", for which we use a tolerance of 10^{-12}. We mention that in all cases, we solve (5.8) with GMRES preconditioned with $M_i = \text{tridiag}(A_i)$. This preconditioner works well here since in all cases the tridiagonal part has more than 50% of the weight of each row. We report in Table 5.3 the number of (outer GMRES) iterations and CPU times in seconds required to reach the overall (outer) stopping criterion of $\|r_k\|/\|r_0\| \leq 10^{-10}$. In the table n is the order of the matrix. We remark that the "leaky" matrices have null spaces with dimension larger than one, and this explains why the timings for these matrices are so much higher than for the others. For the matrices for which $\dim\mathcal{N}(A) = n - 1$, we compare the results of the first run, i.e., the one using (5.12), denoted by $\tilde{x}$ with GMRES with a tolerance of 10^{-14}, denoted x^{true}. In both cases we normalize the solution x so that $\sum_{i=1}^{n} x_i = 1$. We also performed these runs using left preconditioners, and the number of iterations and times are very similar to those in Table 5.3.

Matrix	n	it.	(5.12)	it.	$\frac{10^{-2}\epsilon}{\|\tilde{r}_{k-1}\|}$	it.	10^{-10}	it.	10^{-12}	$\|x^{true} - \tilde{x}\|$
leaky530	530	11	4.91	11	3.92	11	5.03	11	6.02	
mutex794	794	5	0.49	5	0.43	5	0.49	5	0.55	2.0704×10^{-10}
qnatm740	740	8	1.12	8	0.97	8	1.24	8	1.36	8.6973×10^{-11}
leaky1074	1074	13	19.22	13	14.62	13	20.58	13	24.63	
leaky1618	1618	13	20.62	13	16.06	13	21.84	13	25.65	
mutex3797	3797	6	5.73	6	4.85	6	5.44	6	5.87	5.6478×10^{-12}
mutex4095	4095	7	7.36	7	6.68	7	7.91	7	8.07	3.3872×10^{-12}

TABLE 5.3

Iterations and solution time for (exact) GMRES (next to last column) and inexact GMRES for different inexact tolerances

Observe that for each matrix, the four runs take essentially the same number of iterations. In other words, for these matrices, there is practically no delay in convergence when using inexact GMRES. One obtains almost the same convergence results, and almost up to convergence, the same convergence curves, with each step taking less computational work. Roughly speaking, in most cases reported here, the inexact method takes about 20% less time than the exact one, and about 10% less time than the runs with fixed inner tolerance.

Our second set of experiments with Markov chain matrices from [20] is with two larger matrices, and thus, we use $p > 1$, for several values of p. In table 5.4 we report on these experiments, where we solve (5.8) inexactly, this time using as preconditioner M_i a threshold ILU factorization of A_i, as in [2]. The threshold used are 0.01 for "ncd62196" and 0.2 for "twoD66177." For these inexact runs we use $\rho_{i,k} \leq 10^{-4}\frac{\epsilon}{\|\tilde{r}_{k-1}\|}$ with $\epsilon = 10^{-6}$. The other two runs are of the same type as described for Table 5.3, namely inexact GMRES with a fixed inner tolerance of 10^{-10}, or 10^{-12}. We consider the latter to be "exact", and we report the computational times relative to it. The runs stop when the relative residual is less than 10^{-10}. For these larger problems, we again see convergence in about the same number of iterations, and savings of the inexact version of up to 30%, although typical savings are smaller. We further illustrate the convergence behavior of inexact GMRES in Figure 5.6 for the matrix ncd62196 and $p = 16$. We use the same conventions as in the other figures.

6. Conclusion. In this paper we showed, theoretically and experimentally, that inexact GMRES can also be applied to singular systems, and in particular to those obtained from Markov Chain modeling.

Matrix	n	p	it.	inexact	it.	10^{-10}	it.	10^{-12}
ncd62196	62196	4	29	0.7040	20	0.9435	21	1.00
ncd62196	62196	16	17	0.9912	17	0.9794	17	1.00
ncd62196	62196	30	40	0.9584	40	0.9759	40	1.00
twoD66177	66177	4	12	0.9471	11	0.9730	11	1.00
twoD66177	66177	16	21	0.8644	21	0.9304	21	1.00
twoD66177	66177	32	27	0.8854	27	0.9268	27	1.00

TABLE 5.4
Iterations and relative convergence time for (exact) GMRES (last column) and inexact GMRES.

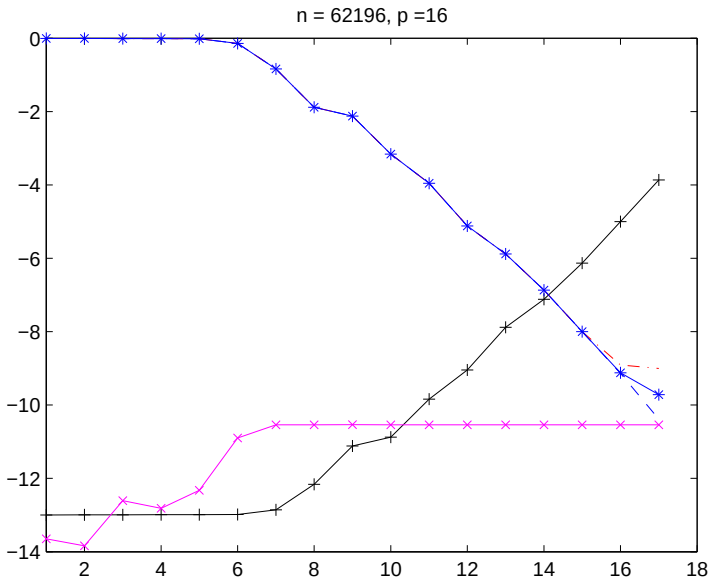

FIG. 5.6. *Convergence curves of inexact GMRES and GMRES, matrix ncd62196 with* $p = 16$.

In our numerical experiments, we compared inexact GMRES and (exact) GMRES for consistent singular linear systems obtained from Markov Chain matrices, and showed that one can indeed save computational effort using inexact matrix vector products. These savings, while not as dramatic as in the nonsingular case, nevertheless come with no deterioration of the overall convergence. In cases when the index k_0 of the matrix is larger than one, the right hand side needs to lie on (or close to) $\mathcal{R}(A^{k_0})$ for the methods to work well.

Acknowledgements. We thank Sébastien Loisel, Michele Benzi, Lothar Reichel, and Valeria Simoncini for their helpful comments.

REFERENCES

[1] Vincent A. Barker. Numerical solutions of sparse singular systems of equations arising from ergodic Markov chains. *Commun. Statist.-Stochastic Models*, 5:335–381, 1989.

[2] Michele Benzi and Bora Uçar. Block triangular preconditioners for M-matrices and Markov chains. *Electronic Transactions on Numerical Analysis*, 26:209–227, 2007.

[3] Åke Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.
[4] Amina Bouras and Valérie Frayssé. Inexact matrix-vector products in krylov methods for solving linear systems:a relaxation strategy. *SIAM Journal on Matrix Analysis and Applications*, 26:660–678, 2005.
[5] Peter Brown and Homer Walker. GMRES on (nearly) singular systems. *SIAM Journal of Matrix Analysis Applications*, 18:37–51, 1997.
[6] Daniela Calvetti, Brian W. Lewis, and Lothar Reichel. GMRES-type methods for inconsistent systems. *Linear Algebra and its Applications*, 316:157–169, 2000.
[7] Zhi hao Cao and Mei Wang. A note on Krylov subspace methods for singular systems. *Linear Algebra and Its Applications*, 350:285–288, 2002.
[8] Roland W. Freund and Marlis Hochbruck. On the use of two QMR algorithms for solving singular systems and applications in Markov chain modeling. *Numerical Linear Algebra with Applications*, 1:403–420, 1994.
[9] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, New York, NY, USA, 1999.
[10] George Karypis and Vipin Kumar. Metis: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices version 4.0. 1998.
[11] Christopher C. Paige and Zdeněk Strakŏs. Scaled total least squares fundamentals. *Numerische Mathematik*, 91:117–146, 2002.
[12] Bernard Philippe, Yousef Saad, and William J. Stewart. Numerical methods in Markov chain modeling. *Operations Research*, 40:1156–1179, 1992.
[13] Lothar Reichel and Qiang Ye. Breakdown-free GMRES for singular systems. *SIAM Journal of Matrix Analysis Applications*, 26:1001–1021, 2005.
[14] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. The PWS Publishing Company, Boston, U.S.A, 1996. Second edition, SIAM, Philadelphia, 2003.
[15] Yousef Saad and Martin H. Schultz. Gmres: a generalized minimal residual algorithm for solvingnonsymmetric linear systems. *Journal on Scientific and Statistical Computing*, 7:856–869, 1986.
[16] Valeria Simoncini and Daniel B. Szyld. Theory of inexact Krylov subspace mehods and applications to scientific computing. *SIAM Journal on Scientific Computing*, 25:454–477, 2003.
[17] Valeria Simoncini and Daniel B. Szyld. Recent computational developments in Krylov subspace methods for linear systems. *Numerical Linear Algebra with Applications*, 14:1–59, 2007.
[18] Laurent Smoch. Some results about GMRES in the singular case. *Numerical Algorithms*, 22:193–212, 1999.
[19] William J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, Princeton, NJ., 1994.
[20] William J. Stewart. Marca models: A collection of Markov chain models. Available at http://www.csc.ncsu.edu/faculty/stewart/MARCA Models/MARCA Models.html.
[21] Jasper van den Eshof and Gerard L. G. Sleijpen. Inexact Krylov subspace methods for linear systems. *SIAM Journal of Matrix Analysis Applications*, 26:125–153, 2004.
[22] Yimin Wei and Hebing Wu. Convergence properties of Krylov subspace methods for singular linear systems with arbitrary index. *Journal of Computational and Applied Mathematics*, 114:305–318, 2000.
[23] James H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, U.K., 1965.
[24] Naimin Zhang and Yimin Wei. A note on solving EP inconsistent linear systems. *Applied Mathematics and Computation*, 169:8–15, 2005.
[25] Naimin Zhang and Yimin Wei. Solving EP singular linear systems. *International Journal of Computer Mathematics*, 81:1395–1405, 2005.
[26] Jieyong Zhou and Yimin Wei. DFOM algorithm and error analysis for projection methods for solving singular linear system. *Applied Mathematics and Computation*, 157:313–329, 2004.
[27] Jieyong Zhou and Yimin Wei. Stagnation analysis of DGMRES. *Numerical Linear Algebra with Applications*, 151:27–39, 2004.