

MINRES-QLP: A KRYLOV SUBSPACE METHOD FOR INDEFINITE OR SINGULAR SYMMETRIC SYSTEMS*

SOU-CHENG T. CHOI[†], CHRISTOPHER C. PAIGE[‡], AND MICHAEL A. SAUNDERS[§]

Abstract. CG, SYMMLQ, and MINRES are Krylov subspace methods for solving symmetric systems of linear equations. When these methods are applied to an incompatible system (that is, a singular symmetric least-squares problem), CG could break down and SYMMLQ's solution could explode, while MINRES would give a least-squares solution but not necessarily the minimum-length (pseudoinverse) solution. This understanding motivates us to design a MINRES-like algorithm to compute minimum-length solutions to singular symmetric systems.

MINRES uses QR factors of the tridiagonal matrix from the Lanczos process (where R is upper-tridiagonal). MINRES-QLP uses a QLP decomposition (where rotations on the right reduce R to lower-tridiagonal form). On ill-conditioned systems (singular or not), MINRES-QLP can give more accurate solutions than MINRES. We derive preconditioned MINRES-QLP, new stopping rules, and better estimates of the solution and residual norms, the matrix norm, and the condition number.

Key words. MINRES, Krylov subspace method, Lanczos process, conjugate-gradient method, minimum-residual method, singular least-squares problem, sparse matrix

AMS subject classifications. 15A06, 65F10, 65F20, 65F22, 65F25, 65F35, 65F50, 93E24

DOI. xxx/xxxxxxxx

1. Introduction. We are concerned with iterative methods for solving a symmetric linear system $Ax = b$ or the related least-squares (LS) problem

$$\min \|x\|_2 \quad \text{s.t.} \quad x \in \arg \min_x \|Ax - b\|_2, \quad (1.1)$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and possibly singular, $b \in \mathbb{R}^n$, $A \neq 0$, and $b \neq 0$. Most of the results in our discussion are directly extendable to problems with complex Hermitian matrices A and complex vectors b .

The solution of (1.1), called the *minimum-length* or *pseudoinverse* solution [18], is formally given by $x^\dagger = (A^T A)^\dagger A^T b = (A^2)^\dagger A b = (A^\dagger)^2 A b$, where A^\dagger denotes the pseudoinverse of A . The pseudoinverse is continuous under perturbations E for which $\text{rank}(A + E) = \text{rank}(A)$ [49], and x^\dagger is continuous under the same condition. Problem (1.1) is then well-posed [19].

Let $A = U \Lambda U^T$ be an eigenvalue decomposition of A , with U orthogonal and $\Lambda \equiv \text{diag}(\lambda_1, \dots, \lambda_n)$. We define the condition number of A to be $\kappa(A) = \frac{\max |\lambda_i|}{\min_{\lambda_i \neq 0} |\lambda_i|}$, and we say that A is ill-conditioned if $\kappa(A) \gg 1$. Hence a singular matrix could be well-conditioned or ill-conditioned.

*Technical Report SOL 2010-3.
Revised March 31, 2011.

Submitted to SIAM J. Sci. Comput.

<http://www.siam.org/journals/sisc/>

[†]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305-4121 (schoi@stanford.edu). This author's research was partially supported by National Science Foundation grant CCR-0306662.

[‡]Computer Science, McGill University, Montreal, Quebec, Canada, H3A 2A7 (paige@cs.mcgill.ca). This author's research was partially supported by NSERC of Canada grant OGP0009236.

[§]Department of Management Science and Engineering, Stanford University, Stanford, CA 94305-4026 (saunders@stanford.edu). This author's research was partially supported by National Science Foundation grant CCR-0306662, Office of Naval Research grants N00014-02-1-0076 and N00014-08-1-0191, and AHPARC.

SYMMLQ and MINRES [39] are Krylov subspace methods for solving symmetric indefinite systems $Ax = b$. SYMMLQ is reliable on compatible systems even if A is ill-conditioned or singular, while on (singular) incompatible problems its iterates x_k diverge to a multiple of a nullvector of A [10, Proposition 2.15] and [10, Lemma 2.17]. MINRES seems more desirable to users because its residual norms are monotonically decreasing. On singular compatible systems, MINRES returns x^\dagger (see Theorem 3.1). On singular incompatible systems, MINRES is reliable if terminated with a suitable stopping rule involving $\|Ar_k\|$ (see Lemma 3.3), but the solution will not be x^\dagger .

Here we develop a new solver of this type named MINRES-QLP [10]. The aim is to deal reliably with compatible or incompatible systems and to return the *unique* solution of (1.1). We give theoretical reasons why MINRES-QLP improves the accuracy of MINRES on ill-conditioned systems, and illustrate with numerical examples.

Incompatible symmetric systems could arise from discretized semidefinite Neumann boundary value problems [27, section 4], and from any other singular systems involving measurement errors in b . Another potential application is large symmetric indefinite low-rank Toeplitz LS problems as described in [16, section 4.1].

1.1. Notation. The letters i, j, k denote integer indices, c and s cosine and sine of some angle θ , e_k the k th unit vector, e a vector of all ones, and other lower-case letters such as b, u , and x (possibly with integer subscripts) denote *column* vectors. Upper-case letters A, T_k, V_k, \dots denote matrices, and I_k is the identity matrix of order k . Lower-case Greek letters denote scalars; in particular, $\varepsilon \approx 10^{-16}$ denotes the floating-point precision. If a quantity δ_k is modified one or more times, we denote its values by $\delta_k, \delta_k^{(2)}, \delta_k^{(3)}, \dots$. The symbol $\|\cdot\|$ denotes the 2-norm of a vector or matrix. For an incompatible system, $Ax \approx b$ is shorthand for the LS problem $\min_x \|Ax - b\|$.

1.2. Overview. In sections 2–4 we briefly review the Lanczos process, MINRES, and QLP decomposition before introducing MINRES-QLP in section 5. We derive norm estimates in section 6 and preconditioned MINRES-QLP in section 7. Numerical experiments are described in section 8.

2. The Lanczos process. Given A and b , the Lanczos process [30] computes vectors v_k and tridiagonal matrices \underline{T}_k according to $v_0 \equiv 0$, $\beta_1 v_1 = b$, and then¹

$$p_k = Av_k, \quad \alpha_k = v_k^T p_k, \quad \beta_{k+1} v_{k+1} = p_k - \alpha_k v_k - \beta_k v_{k-1}$$

for $k = 1, 2, \dots, \ell$, where we choose $\beta_k > 0$ to give $\|v_k\| = 1$. In matrix form,

$$AV_k = V_{k+1} \underline{T}_k, \quad \underline{T}_k \equiv \begin{bmatrix} \alpha_1 & \beta_2 & & & \\ \beta_2 & \alpha_2 & \ddots & & \\ & \ddots & \ddots & \beta_k & \\ & & \beta_k & \alpha_k & \\ & & & \beta_{k+1} & \end{bmatrix} \equiv \begin{bmatrix} T_k \\ \beta_{k+1} e_k^T \end{bmatrix}, \quad V_k \equiv [v_1 \ \cdots \ v_k]. \quad (2.1)$$

In exact arithmetic, the columns of V_k are orthonormal and the process stops with $k = \ell$ and $\beta_{\ell+1} = 0$ for some $\ell \leq n$, and then $AV_\ell = V_\ell T_\ell$. For derivation purposes we assume that this happens, though in practice it is unlikely unless V_k is reorthogonalized for each k . In any case, (2.1) holds to machine precision and the computed vectors satisfy $\|V_k\|_1 \approx 1$ (even if $k \gg n$).

¹Numerically, $p_k = Av_k - \beta_k v_{k-1}$, $\alpha_k = v_k^T p_k$, $\beta_{k+1} v_{k+1} = p_k - \alpha_k v_k$ is slightly better [38].

2.1. Properties of the Lanczos process. The k th Krylov subspace generated by A and b is defined to be $\mathcal{K}_k(A, b) = \text{span}\{b, Ab, A^2b, \dots, A^{k-1}b\} = \text{span}(V_k)$. The following properties should be kept in mind:

1. If A is changed to $A - \sigma I$ for some scalar shift σ , T_k becomes $T_k - \sigma I$ and V_k is unaltered, showing that singular systems are commonplace. Shifted problems appear in inverse iteration or Rayleigh quotient iteration.
2. T_k has full column rank k for all $k < \ell$.
3. If A is indefinite, some T_k might be singular for $k < \ell$, but then by the Sturm sequence property (see [18]), T_k has exactly one zero eigenvalue and the strict interlacing property implies that $T_{k\pm 1}$ are nonsingular. Hence T_k cannot be singular twice in a row (whether A is singular or not).
4. T_ℓ is nonsingular if and only if $b \in \text{range}(A)$. (See appendix A.)

3. MINRES. Algorithm MINRES [39] is a natural way of using the Lanczos process to solve $Ax = b$ or $\min_x \|Ax - b\|$. For $k < \ell$, if $x_k = V_k y_k$ for some vector y_k , the associated residual is

$$r_k \equiv b - Ax_k = b - AV_k y_k = \beta_1 v_1 - V_{k+1} \underline{T}_k y_k = V_{k+1}(\beta_1 e_1 - \underline{T}_k y_k). \quad (3.1)$$

To make r_k small, it is clear that $\beta_1 e_1 - \underline{T}_k y_k$ should be small. At this iteration k , MINRES minimizes the residual subject to $x_k \in \mathcal{K}_k(A, b)$ by choosing

$$y_k = \arg \min_{y \in \mathbb{R}^k} \|\underline{T}_k y - \beta_1 e_1\|. \quad (3.2)$$

This subproblem is processed by the expanding QR factorization: $Q_0 \equiv 1$ and

$$Q_{k,k+1} \equiv \begin{bmatrix} I_{k-1} & & \\ & c_k & s_k \\ & s_k & -c_k \end{bmatrix}, \quad Q_k \equiv Q_{k,k+1} \begin{bmatrix} Q_{k-1} & \\ & 1 \end{bmatrix}, \quad Q_k [\underline{T}_k \quad \beta_1 e_1] = \begin{bmatrix} R_k & t_k \\ 0 & \phi_k \end{bmatrix}, \quad (3.3)$$

where c_k and s_k form the Householder reflector $Q_{k,k+1}$ that annihilates β_{k+1} in \underline{T}_k to give upper-tridiagonal R_k , with R_k and t_k being unaltered in later iterations.

When $k < \ell$, the unique solution of (3.2) satisfies $R_k y_k = t_k$. Instead of solving for y_k , MINRES solves $R_k^T D_k^T = V_k^T$ by forward substitution, obtaining the last column d_k of D_k at iteration k . At the same time, it updates x_k via $x_0 \equiv 0$ and

$$x_k = V_k y_k = D_k R_k y_k = D_k t_k = x_{k-1} + \tau_k d_k, \quad \tau_k \equiv e_k^T t_k. \quad (3.4)$$

When $k = \ell$, we can form T_ℓ but nothing else expands. In place of (3.1) and (3.3) we have $r_\ell = V_\ell(\beta_1 e_1 - T_\ell y_\ell)$ and $Q_{\ell-1} [T_\ell \quad \beta_1 e_1] = [R_\ell \quad t_\ell]$ and it is natural to choose y_ℓ from the subproblem

$$\min \|T_\ell y_\ell - \beta_1 e_1\| \quad \equiv \quad \min \|R_\ell y_\ell - t_\ell\|. \quad (3.5)$$

There are two cases to consider:

1. If T_ℓ is nonsingular, $R_\ell y_\ell = t_\ell$ has a unique solution. Since $AV_\ell y_\ell = V_\ell T_\ell y_\ell = b$, the problem is solved by $x_\ell = V_\ell y_\ell$ with residual $r_\ell = 0$ (the system is compatible, even if A is singular). Theorem 3.1 proves that $x_\ell = x^\dagger$.
2. If T_ℓ is singular, A and R_ℓ are singular ($R_{\ell\ell} = 0$) and both $Ax = b$ and $R_\ell y_\ell = t_\ell$ are incompatible. This case was not handled by MINRES in [39]. Theorem 3.2 proves that the MINRES point $x_{\ell-1}$ is a least-squares solution (but not necessarily x^\dagger). Theorem 5.1 proves that the MINRES-QLP point $x_\ell = V_\ell y_\ell^\dagger = x^\dagger$, where y_ℓ^\dagger is the min-length solution of (3.5).

3.1. Further details of MINRES. To describe MINRES-QLP thoroughly, we need further details of the MINRES QR factorization (3.3). For $1 \leq k < \ell$,

$$\begin{bmatrix} R_k \\ 0 \end{bmatrix} = \begin{bmatrix} \gamma_1 & \delta_2 & \epsilon_3 & & & \\ & \gamma_2^{(2)} & \delta_3^{(2)} & \ddots & & \\ & & \ddots & \ddots & & \\ & & & & \ddots & \\ & & & & & \delta_k^{(2)} \\ & & & & & \gamma_k^{(2)} \\ & & & & & 0 \end{bmatrix}, \quad \begin{bmatrix} t_k \\ \phi_k \end{bmatrix} \equiv \begin{bmatrix} \tau_1 \\ \tau_2 \\ \vdots \\ \tau_k \\ \phi_k \end{bmatrix} = \beta_1 \begin{bmatrix} c_1 \\ s_1 c_2 \\ \vdots \\ s_1 \cdots s_{k-1} c_k \\ s_1 \cdots s_{k-1} s_k \end{bmatrix} \quad (3.6)$$

(where the superscripts are defined in section 1.1). With $\phi_0 \equiv \beta_1 > 0$, the full action of $Q_{k,k+1}$ in (3.3), including its effect on later columns of T_j , $k < j \leq \ell$, is described by

$$\begin{bmatrix} c_k & s_k \\ s_k & -c_k \end{bmatrix} \begin{bmatrix} \gamma_k & \delta_{k+1} & 0 \\ \beta_{k+1} & \alpha_{k+1} & \beta_{k+2} \end{bmatrix} \begin{array}{c} \phi_{k-1} \\ 0 \end{array} = \begin{bmatrix} \gamma_k^{(2)} & \delta_{k+1}^{(2)} & \epsilon_{k+2} \\ 0 & \gamma_{k+1} & \delta_{k+2} \end{bmatrix} \begin{array}{c} \tau_k \\ \phi_k \end{array}, \quad (3.7)$$

where $s_k = \beta_{k+1} / \|\gamma_k \ \beta_{k+1}\| > 0$, giving $\gamma_1, \gamma_k^{(2)} > 0$ with R_j nonsingular for each $j \leq k < \ell$. Thus the d_j in (3.4) can be found from

$$R_k^T D_k^T = V_k^T : \begin{cases} d_1 = v_1 / \gamma_1, & d_2 = (v_2 - \delta_2 d_1) / \gamma_2^{(2)}, \\ d_j = (v_j - \delta_j^{(2)} d_{j-1} - \epsilon_j d_{j-2}) / \gamma_j^{(2)}, & j = 3, \dots, k. \end{cases} \quad (3.8)$$

Also, $\tau_k = \phi_{k-1} c_k$ and $\phi_k = \phi_{k-1} s_k > 0$. Hence from (3.1)–(3.3),

$$\|r_k\| = \|\underline{T}_k y_k - \beta_1 e_1\| = \phi_k \quad \Rightarrow \quad \|r_k\| = \|r_{k-1}\| s_k, \quad (3.9)$$

which is nonincreasing and tending to zero if $Ax = b$ is compatible.

REMARK 3.1. If $k < \ell$ and T_k is singular, we have $\gamma_k = 0$, $s_k = 1$, and $\|r_k\| = \|r_{k-1}\|$ (not a strict decrease), but this cannot happen twice in a row (cf. section 2.1).

REMARK 3.2. If T_ℓ is singular, MINRES sets the last element of y_ℓ to be zero. The final point and residual stay as $x_{\ell-1}$ and $r_{\ell-1}$ with $\|r_{\ell-1}\| = \phi_{\ell-1} = \beta_1 s_1 \cdots s_{\ell-1} > 0$.

3.2. Compatible systems. The following theorem assures us that MINRES is a useful solver for compatible linear systems even if A is singular.

THEOREM 3.1 ([10, Theorem 2.25]). If $b \in \text{range}(A)$, the final MINRES point x_ℓ is the minimum-length solution of $Ax = b$ (and $r_\ell = b - Ax_\ell = 0$).

Proof. If $b \in \text{range}(A)$, the Lanczos process gives $AV_\ell = V_\ell T_\ell$ with nonsingular T_ℓ , and MINRES terminates with $Ax_\ell = b$ and $x_\ell = V_\ell y_\ell = Aq$, where $q = V_\ell T_\ell^{-1} y_\ell$. If some other point \hat{x} satisfies $A\hat{x} = b$, let $p = \hat{x} - x_\ell$. We have $Ap = 0$ and $x_\ell^T p = q^T Ap = 0$. Hence $\|\hat{x}\|^2 = \|x_\ell + p\|^2 = \|x_\ell\|^2 + 2x_\ell^T p + \|p\|^2 \geq \|x_\ell\|^2$. \square

3.3. Incompatible systems. For a singular LS problem $Ax \approx b$, the optimal residual vector \hat{r} is unique, but infinitely many solutions x give that residual. In the following example, MINRES finds a least-squares solution (with optimal residual) but not the minimum-length solution.

EXAMPLE 3.1. Let $A = \text{diag}(1, 1, 0)$ and $b = e$. The minimum-length solution to $Ax \approx b$ is $x^\dagger = [1 \ 1 \ 0]^T$ with residual $\hat{r} = b - Ax^\dagger = e_3$ and $A\hat{r} = 0$. MINRES returns the solution $x^\# = e$ (with residual $r^\# = b - Ax^\# = e_3 = \hat{r}$ and $Ar^\# = 0$).

THEOREM 3.2 ([10, Theorem 2.27]). If $b \notin \text{range}(A)$, then $\|Ar_{\ell-1}\| = 0$ and the MINRES $x_{\ell-1}$ is an LS solution (but not necessarily x^\dagger).

Proof. Since $b \notin \text{range}(A)$, T_ℓ is singular and $R_{\ell\ell} = \gamma_\ell = 0$. By Lemma 3.3 below, $A(Ax_{\ell-1} - b) = -Ar_{\ell-1} = -\|r_{\ell-1}\| \gamma_\ell v_\ell = 0$. Thus $x_{\ell-1}$ is an LS solution. \square

3.4. Norm estimates in MINRES. For incompatible systems, r_k (3.1) will never be zero. However, all LS solutions satisfy $A^2x = Ab$, so that $Ar = 0$. We therefore need a new stopping condition based on the size of $\|Ar_k\|$. In applications requiring nullvectors, $\|Ax_k\|$ is also useful. We present efficient recurrence relations for $\|Ar_k\|$ and $\|Ax_k\|$ in the following Lemma, which was not considered in the framework of MINRES when it was originally designed for nonsingular systems [39].

LEMMA 3.3 (Ar_k , $\|Ar_k\|$, $\|Ax_k\|$ for MINRES). *If $k < \ell$,*

$$\begin{aligned} Ar_k &= \|r_k\| (\gamma_{k+1}v_{k+1} + \delta_{k+2}v_{k+2}) && (\text{where } \delta_{k+2}v_{k+2} = 0 \text{ if } k = \ell - 1), \\ \psi_k^2 &\equiv \|Ar_k\|^2 = \|r_k\|^2 ([\gamma_{k+1}]^2 + [\delta_{k+2}]^2) && (\text{where } \delta_{k+2} = 0 \text{ if } k = \ell - 1), \\ \omega_k^2 &\equiv \|Ax_k\|^2 = \omega_{k-1}^2 + \tau_k^2, \quad \omega_0 \equiv 0. \end{aligned}$$

Proof. For $k < \ell$, R_k is nonsingular. From (3.1)–(3.4) with $R_k y_k = t_k$ we have

$$r_k = V_{k+1} Q_k^T \left(\begin{bmatrix} t_k \\ \phi_k \end{bmatrix} - \begin{bmatrix} R_k \\ 0 \end{bmatrix} y_k \right) = \phi_k V_{k+1} Q_k^T e_{k+1}, \quad (3.10)$$

$$Ar_k = \phi_k V_{k+2} \underline{T}_{k+1} Q_k^T e_{k+1},$$

$$Q_k \underline{T}_{k+1}^T = Q_k [T_{k+1} \quad \beta_{k+2} e_{k+1}] = Q_k \begin{bmatrix} T_k & \beta_{k+1} e_k & 0 \\ \beta_{k+1} e_k^T & \alpha_{k+1} & \beta_{k+2} \end{bmatrix},$$

$$e_{k+1}^T Q_k \underline{T}_{k+1}^T = [0 \quad \gamma_{k+1} \quad \delta_{k+2}],$$

see (3.7), where $AV_{k+1} = V_{k+1}T_{k+1}$ and we take $\delta_{k+2} = 0$ if $k = \ell - 1$, so

$$\begin{aligned} Ar_k &= \phi_k V_{k+2} [0 \quad \gamma_{k+1} \quad \delta_{k+2}]^T = \phi_k (\gamma_{k+1}v_{k+1} + \delta_{k+2}v_{k+2}), \\ \psi_k^2 &\equiv \|Ar_k\|^2 = \|r_k\|^2 ([\gamma_{k+1}]^2 + [\delta_{k+2}]^2). \end{aligned}$$

For the recurrence relations of Ax_k and its norm, we have

$$\begin{aligned} Ax_k &= AV_k y_k = V_{k+1} \underline{T}_k y_k = V_{k+1} Q_k^T \begin{bmatrix} R_k \\ 0 \end{bmatrix} y_k = V_{k+1} Q_k^T \begin{bmatrix} t_k \\ 0 \end{bmatrix}, \\ \omega_k^2 &\equiv \|Ax_k\|^2 = \|t_k\|^2 = \|t_{k-1}\|^2 + \tau_k^2 = \omega_{k-1}^2 + \tau_k^2. \end{aligned}$$

□

Note that even using finite precision the expression for ψ_k^2 is extremely accurate for the versions of the Lanczos algorithm given in section 2, since (taking $\|v_j\| = 1$ with negligible error), $\|Ar_k\|^2 = \phi_k^2 ([\gamma_{k+1}]^2 + 2\gamma_{k+1}\delta_{k+2}v_{k+1}^T v_{k+2} + [\delta_{k+2}]^2)$, where from (3.7) $|\delta_{k+2}| \leq \beta_{k+2}$, while from [38, (18)] $|\beta_{k+2} v_{k+1}^T v_{k+2}| \leq O(\varepsilon)\|A\|$, and with $|\gamma_{k+1}| \leq \|A\|$, see [38, (19)], we see that $|\gamma_{k+1}\delta_{k+2}v_{k+1}^T v_{k+2}| \leq O(\varepsilon)\|A\|^2$.

Typically $\|Ar_k\|$ is not monotonic, while clearly $\|r_k\|$ and $\|Ax_k\|$ are monotonic. In the eigensystem $A = U\Lambda U^T$, let $U = [U_1 \ U_2]$, where the eigenvectors U_1 correspond to nonzero eigenvalues. Then $P_A \equiv U_1 U_1^T$ and $P_A^\perp \equiv U_2 U_2^T$ are orthogonal projectors [53] onto the range and nullspace of A . For general linear LS problems, Chang et al. [7] characterize the dynamics of $\|r_k\|$ and $\|A^T r_k\|$ in three phases defined in terms of the ratios among $\|r_k\|$, $\|P_A r_k\|$, and $\|P_A^\perp r_k\|$, and propose two new stopping criteria for iterative solvers. The expositions [1, 26] show that these estimates are cheaply computable in CGLS and LSQR [40, 41].

3.5. Effects of rounding errors in MINRES. MINRES should stop if R_k is singular (which theoretically implies $k = \ell$ and A is singular). Singularity was not discussed by Paige and Saunders [39], but they did raise the question: Is MINRES stable when R_k is ill-conditioned? Their concern was that $\|D_k\|$ could be large in (3.8), and there could then be cancellation in forming $x_{k-1} + \tau_k d_k$ in (3.4).

Sleijpen, Van der Vorst, and Modersitzki [47] analyzed the effects of rounding errors in MINRES and reported examples of apparent failure with a matrix of the form $A = QDQ^T$, where D is an ill-conditioned diagonal matrix and Q involves a single plane rotation. We were unable to reproduce MINRES's performance on the two examples defined in Figure 4 of their paper, but we modified the examples by using an $n \times n$ Householder transformation for Q , and then observed similar difficulties with MINRES—see Figure 8.2. The recurred residual norm ϕ_k^M is a good approximation to the directly computed $\|r_k^M\|$ until the last few iterations. The recurred norms ϕ_k^M then keep decreasing but the directly computed norms $\|r_k^M\|$ become stagnant or even increase (see the lower subplots in Figure 8.2).

REMARK 3.3. *Note that we want ϕ_k to keep decreasing on compatible systems, so that the test $\phi_k \leq \text{tol}(\|A\|\|x_k\| + \|b\|)$ with $\text{tol} \geq \varepsilon$ will eventually be satisfied even if the computed $\|r_k\|$ is no longer as small as ϕ_k .*

The analysis in [47] focuses on the rounding errors involved in the n lower triangular solves $R_k^T D_k^T = V_k^T$ (one solve for each row of D_k), compared to the single upper triangular solve $R_k y_k = t_k$ (followed by $x_k = V_k y_k$) that would be possible at the final k if all of V_k were stored as in GMRES [44]. We shall see that a key feature of MINRES-QLP is that a single lower triangular solve suffices with no need to store V_k , much the same as in SYMMLQ.

4. Orthogonal decompositions for singular matrices. In 1999 Stewart proposed the *pivoted QLP decomposition* [51], which is equivalent to two consecutive QR factorizations with column interchanges, first on A , then on R^T :

$$Q_R A \Pi_R = \begin{bmatrix} R & S \\ 0 & 0 \end{bmatrix}, \quad Q_L \begin{bmatrix} R^T & 0 \\ S^T & 0 \end{bmatrix} \Pi_L = \begin{bmatrix} \hat{R} & 0 \\ 0 & 0 \end{bmatrix}, \quad (4.1)$$

giving nonnegative diagonal elements, where Π_R and Π_L are permutations chosen to maximize the next diagonal element of R and \hat{R} at each stage. This gives $A = QLP$, where

$$Q = Q_R^T \Pi_L, \quad L = \begin{bmatrix} \hat{R}^T & 0 \\ 0 & 0 \end{bmatrix}, \quad P = Q_L \Pi_R^T,$$

with Q and P orthogonal. Stewart demonstrated that the diagonals of L (the *L-values*) give better singular-value estimates than the diagonals of R (the *R-values*), and the accuracy is particularly good for the extreme singular values σ_1 and σ_n :

$$R_{ii} \approx \sigma_i, \quad L_{ii} \approx \sigma_i, \quad \sigma_1 \geq \max_i L_{ii} \geq \max_i R_{ii}, \quad \min_i R_{ii} \geq \min_i L_{ii} \geq \sigma_n. \quad (4.2)$$

The first permutation Π_R in pivoted QLP is important. The main purpose of the second permutation Π_L is to ensure that the L -values present themselves in decreasing order, which is not always necessary. If $\Pi_R = \Pi_L = I$, it is simply called the *QLP decomposition*.

5. MINRES-QLP. We now develop MINRES-QLP for solving ill-conditioned or singular symmetric systems $Ax \approx b$. The Lanczos framework is the same as in MINRES, but we handle T_ℓ in (3.5) with extra care when it is rank-deficient. In this case, the normal approach to solving (3.5) is via a QLP decomposition of T_ℓ to obtain the (unique) minimum-length solution y_ℓ [51, 18]. Thus in MINRES-QLP we use a QLP decomposition of T_k in subproblem (3.2) for all $k \leq \ell$. This is the MINRES QR (3.3) followed by an LQ factorization of R_k :

$$Q_k \underline{T}_k = \begin{bmatrix} R_k \\ 0 \end{bmatrix}, \quad R_k P_k = L_k, \quad \text{so that} \quad Q_k \underline{T}_k P_k = \begin{bmatrix} L_k \\ 0 \end{bmatrix}, \quad (5.1)$$

where Q_k and P_k are orthogonal, R_k is upper tridiagonal and L_k is lower tridiagonal. When $k < \ell$, R_k and L_k are nonsingular. MINRES-QLP obtains the same solution as MINRES, but by a different process (and with different rounding errors). Defining u by $y = P_k u$, we see from (3.3) that

$$Q_k (\underline{T}_k y - \beta_1 e_1) = \begin{bmatrix} L_k \\ 0 \end{bmatrix} u - \begin{bmatrix} t_k \\ \phi_k \end{bmatrix},$$

and (3.2) is solved by $L_k u_k = t_k$ and $y_k = P_k u_k$. The MINRES-QLP estimate of x is therefore $x_k = V_k y_k = V_k P_k u_k = W_k u_k$, with theoretically orthonormal $W_k \equiv V_k P_k$.

We will see that only the last three columns of V_k are needed to update x_k .

5.1. The QLP factorization of T_k . The QLP decomposition of each T_k must be without permutations in order to ensure inexpensive updating of the factors as k increases. Our experience is that the desired rank-revealing properties (4.2) tend to be retained, perhaps because of the tridiagonal structure of T_k and the convergence properties of the underlying Lanczos process.

For $k < \ell$, the QLP decomposition of T_k (5.1) gives nonsingular tridiagonal R_k and L_k . As in MINRES, Q_k is a product of Householder reflectors, see (3.3) and (3.7), while P_k involves a product of *pairs* of essentially 2×2 reflectors:

$$Q_k = Q_{k,k+1} \cdots Q_{3,4} \quad Q_{2,3} \quad Q_{1,2}, \quad P_k = P_{1,2} \quad P_{1,3} P_{2,3} \cdots P_{k-2,k} P_{k-1,k}.$$

For MINRES-QLP to be efficient, in the k th iteration ($k \geq 3$) the application of the left reflector $Q_{k,k+1}$ is followed immediately by the right reflectors $P_{k-2,k}, P_{k-1,k}$, so that only the last 2×2 principal submatrix of the transformed T_k will be changed in future iterations. These ideas can be understood more easily from Figure 5.1 and the following compact form, which represents the actions of right reflectors on T_k (additional to $Q_{k,k+1}$ (3.7)):

$$\begin{aligned} & \begin{bmatrix} \gamma_{k-2}^{(5)} & & \epsilon_k \\ \vartheta_{k-1} & \gamma_{k-1}^{(4)} & \delta_k^{(2)} \\ & & \gamma_k^{(2)} \end{bmatrix} \begin{bmatrix} c_{k2} & s_{k2} \\ & 1 \\ s_{k2} & -c_{k2} \end{bmatrix} \begin{bmatrix} 1 & & \\ & c_{k3} & s_{k3} \\ & s_{k3} & -c_{k3} \end{bmatrix} \\ &= \begin{bmatrix} \gamma_{k-2}^{(6)} & & \\ \vartheta_{k-1}^{(2)} & \gamma_{k-1}^{(4)} & \delta_k^{(3)} \\ \eta_k & & \gamma_k^{(3)} \end{bmatrix} \begin{bmatrix} 1 & & \\ & c_{k3} & s_{k3} \\ & s_{k3} & -c_{k3} \end{bmatrix} = \begin{bmatrix} \gamma_{k-2}^{(6)} & & \\ \vartheta_{k-1}^{(2)} & \gamma_{k-1}^{(5)} & \\ \eta_k & \vartheta_k & \gamma_k^{(4)} \end{bmatrix}. \quad (5.2) \end{aligned}$$

5.2. The diagonals of L_k . Figure 5.2 shows the relation between the singular values of A and the diagonal elements of R_k and L_k with $k = 19$. This illustrates (4.2) for matrix ID 1177 from [54] with $n = 25$.




FIG. 5.1. QLP with left and right reflectors interleaved on T_5 . This figure can be reproduced with the help of `QLPfig5.m`.

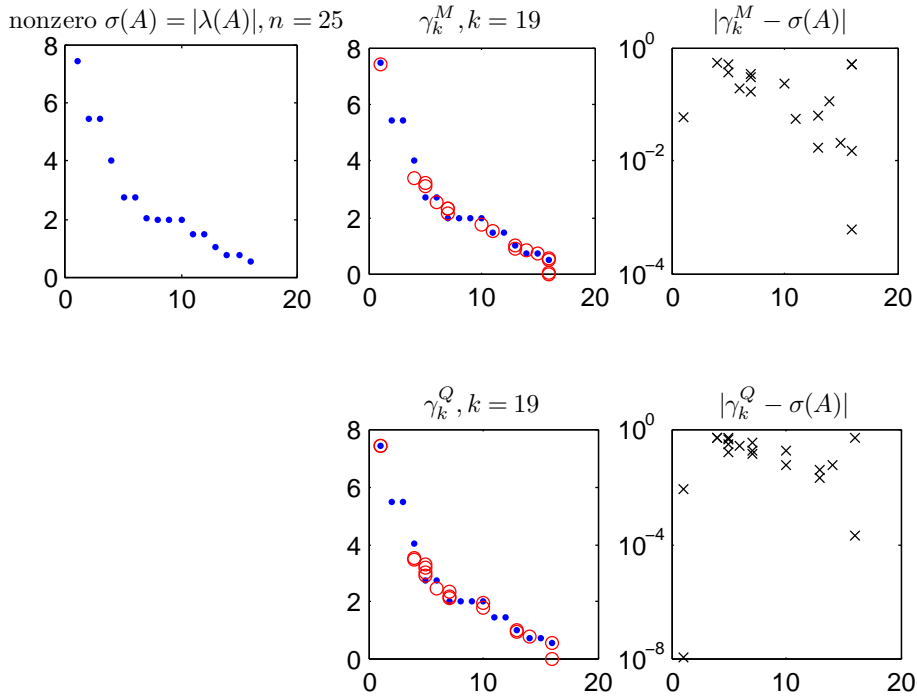


FIG. 5.2. **Upper left:** Nonzero singular values of A sorted in decreasing order. **Upper middle and right:** The diagonals γ_k^M of R_k (red circles) from MINRES are plotted as red circles above or below the nearest singular value of A . They approximate the extreme nonzero singular values of A well. **Lower:** The diagonals γ_k^Q of L_k (red circles) from MINRES-QLP approximate the extreme nonzero singular values of A even better. An implication is that the ratio of the largest and smallest diagonals of L_k provides a good estimate of $\kappa(A)$. To reproduce this figure, run `test_minresqlp_fig3(2)`.

solution of $\|Ax - b\|$ in $\mathcal{K}_\ell(A, b)$, i.e., $x_\ell = \arg \min\{\|x\| \mid A^2x = Ab, x \in \mathcal{K}_\ell(A, b)\}$. Likewise $y_\ell = P_\ell u_\ell$ is the min-length LS solution of $\|T_\ell y - \beta_1 e_1\|$ and so $y_\ell \in \text{range}(T_\ell)$, i.e. $y_\ell = T_\ell z$ for some z . Thus $x_\ell = V_\ell y_\ell = V_\ell T_\ell z = AV_\ell z \in \text{range}(A)$. We know that $x^\dagger = \arg \min\{\|x\| \mid A^2x = Ab, x \in \mathbb{R}^n\}$ is unique and $x^\dagger \in \text{range}(A)$. Since $x_\ell \in \text{range}(A)$, we must have $x_\ell = x^\dagger$. \square

5.5. Transfer from MINRES to MINRES-QLP. On well-conditioned systems, MINRES and MINRES-QLP behave very similarly. However, MINRES-QLP requires one more vector of storage, and each iteration needs 4 more axpy's ($y \leftarrow \alpha x + y$) and 3 more vector scalings ($x \leftarrow \alpha x$). Thus it would be a desirable feature to invoke MINRES-QLP from MINRES only if A is ill-conditioned or singular. The key idea is to transfer to MINRES-QLP at an iteration where \underline{T}_k is not yet too ill-conditioned. The MINRES and MINRES-QLP solution estimates are the same, so from (3.4), (5.9), and (5.3): $x_k^M = x_k \iff D_k t_k = W_k u_k = W_k L_k^{-1} t_k$. Now from (3.8), (5.1), and (5.6),

$$D_k L_k = (V_k R_k^{-1})(R_k P_k) = V_k P_k = W_k, \quad (5.10)$$

and the last three columns of W_k can be obtained from the last three columns of D_k and L_k . (Thus, we transfer the three MINRES basis vectors d_{k-2}, d_{k-1}, d_k to w_{k-2}, w_{k-1}, w_k .) In addition, we need to generate $x_{k-2}^{(2)}$ using (5.8):

$$x_{k-2}^{(2)} = x_k^M - w_{k-1}^{(3)} \mu_{k-1}^{(2)} - w_k^{(2)} \mu_k.$$

It is clear from (5.10) that we still need to do the right transformation $R_k P_k = L_k$ in the MINRES phase and keep the last 3×3 principal submatrix of L_k for each k so that we are ready to transfer to MINRES-QLP when necessary. We then obtain a short recurrence for $\|x_k\|$ (see section 6.5) and for this computation we save flops relative to the original MINRES algorithm, where $\|x_k\|$ is computed directly.

In the implementation, the MINRES iterates transfer to MINRES-QLP iterates when an estimate of the condition number of T_k (see (6.3)) exceeds an input parameter trancond . Thus, $\text{trancond} > 1/\varepsilon$ leads to MINRES iterates throughout, while $\text{trancond} = 1$ generates MINRES-QLP iterates from the start.

5.6. Comparison of Lanczos-based solvers. We compare MINRES-QLP with CG, SYMMLQ, and MINRES in Tables 5.1–5.2 in terms of subproblem definitions, basis, solution estimates, flops and memory. A careful implementation of SYMMLQ provides a point in $\mathcal{K}_{k+1}(A, b)$ as shown. All solvers need storage for v_k, v_{k+1}, x_k , and a product $p_k = Av_k$ each iteration. Some additional work-vectors are needed for each method (e.g., d_{k-1} and d_k for MINRES, giving 7 work-vectors in total).

6. Stopping conditions and norm estimates. This section derives several norm estimates that are computed in MINRES-QLP. As before, we assume exact arithmetic throughout, so that V_k and Q_k are orthonormal. Table 6.1 summarizes how the norm estimates are used to formulate three groups of stopping conditions. The second NRBE test $\|Ar_k\| \leq \|A\| \|r_k\| \text{tol}$ is from Stewart [50] with symmetric A .

6.1. Residual and residual norm. First we derive recurrence relations for r_k and its norm $\phi_k \equiv \|r_k\|$.

LEMMA 6.1 (r_k and $\|r_k\|$ for MINRES-QLP and monotonicity of $\|r_k\|$).

- If $k < \ell$, then $\text{rank}(L_k) = k$, $r_k = s_k^2 r_{k-1} - \phi_k c_k v_{k+1}$, and $\phi_k = \phi_{k-1} s_k > 0$.
- If $\text{rank}(L_\ell) = \ell$, then $r_\ell = 0$.
- If $\text{rank}(L_\ell) = \ell - 1$, then $r_\ell = r_{\ell-1} \neq 0$, and $\|r_\ell\| = \phi_{\ell-1} > 0$.

TABLE 5.1

Subproblems defining x_k for CG, SYMMLQ, MINRES, and MINRES-QLP.

Method	Subproblem	Factorization	Estimate of x_k
cgLanczos [24, 39, 48]	$T_k y_k = \beta_1 e_1$	Cholesky: $T_k = L_k D_k L_k^T$	$x_k^C = V_k y_k$ $\in \mathcal{K}_k(A, b)$
SYMMLQ [39, 45]	$y_{k+1} = \arg \min_{y \in \mathbb{R}^{k+1}} \ y\ $ s.t. $\underline{T}_k^T y = \beta_1 e_1$	LQ: $\underline{T}_k^T Q_k^T = [L_k \ 0]$	$x_k^L = V_{k+1} y_{k+1}$ $\in \mathcal{K}_{k+1}(A, b)$
MINRES [39]	$y_k = \arg \min_{y \in \mathbb{R}^k} \ \underline{T}_k y - \beta_1 e_1\ $	QR: $Q_k \underline{T}_k = \begin{bmatrix} R_k \\ 0 \end{bmatrix}$	$x_k^M = V_k y_k$ $\in \mathcal{K}_k(A, b)$
MINRES-QLP [10]	$y_k = \arg \min_{y \in \mathbb{R}^k} \ y\ $ s.t. $y \in \arg \min \ \underline{T}_k y - \beta_1 e_1\ $	QLP: $Q_k \underline{T}_k P_k = \begin{bmatrix} L_k \\ 0 \end{bmatrix}$	$x_k^Q = V_k y_k$ $\in \mathcal{K}_k(A, b)$

TABLE 5.2

Bases, subproblem solutions, storage, and work for each method.

Method	New basis	z_k, t_k, u_k	x_k estimate	vecs	flops
cgLanczos	$W_k \equiv V_k L_k^{-T}$	$L_k D_k z_k = \beta_1 e_1$	$x_k^C = W_k z_k$	5	$8n$
SYMMLQ	$W_k \equiv V_{k+1} Q_k^T \begin{bmatrix} I_k \\ 0 \end{bmatrix}$	$L_k z_k = \beta_1 e_1$	$x_k^L = W_k z_k$	6	$9n$
MINRES	$D_k \equiv V_k R_k^{-1}$	$t_k = \beta_1 [I_k \ 0] Q_k e_1$	$x_k^M = D_k t_k$	7	$9n$
MINRES-QLP	$W_k \equiv V_k P_k$	$L_k u_k = \beta_1 [I_k \ 0] Q_k e_1$	$x_k^Q = W_k u_k$	8	$14n$

Proof. If $k < \ell$, the residual is the same as for MINRES. We have $\|r_k\| = \phi_k = \phi_{k-1} s_k > 0$; see (3.6)–(3.9). Also from $r_k = \phi_k V_{k+1} Q_k^T e_{k+1}$ (3.10) we have

$$\begin{aligned}
r_k &= \phi_k \begin{bmatrix} V_k & v_{k+1} \end{bmatrix} \begin{bmatrix} Q_{k-1}^T & \\ & 1 \end{bmatrix} \begin{bmatrix} I_{k-1} & & \\ & c_k & s_k \\ & s_k & -c_k \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad \text{by (3.3),} \\
&= \phi_k \begin{bmatrix} V_k & v_{k+1} \end{bmatrix} \begin{bmatrix} Q_{k-1}^T & \\ & 1 \end{bmatrix} \begin{bmatrix} s_k e_k \\ -c_k \end{bmatrix} = \phi_k \begin{bmatrix} V_k & v_{k+1} \end{bmatrix} \begin{bmatrix} s_k Q_{k-1}^T e_k \\ -c_k \end{bmatrix} \\
&= \phi_k s_k V_k Q_{k-1}^T e_k - \phi_k c_k v_{k+1} = \phi_{k-1} s_k^2 V_k Q_{k-1}^T e_k - \phi_k c_k v_{k+1} \\
&= s_k^2 r_{k-1} - \phi_k c_k v_{k+1} \quad \text{by (3.10).}
\end{aligned}$$

If T_ℓ is nonsingular, $r_\ell = 0$. Otherwise $Q_{\ell-1, \ell}$ has made the last row of R_ℓ zero, so the last row and column of L_ℓ are zero; see (5.5). Thus $r_\ell = r_{\ell-1} \neq 0$; see Remark 3.2. \square

6.2. Norm of Ar_k . Next we derive recurrence relations for Ar_k and its norm $\psi_k \equiv \|Ar_k\|$, and we show that Ar_k is orthogonal to $\mathcal{K}_k(A, b)$.

LEMMA 6.2 (Ar_k and $\psi_k \equiv \|Ar_k\|$ for MINRES-QLP).

- If $k < \ell$, then $\text{rank}(L_k) = k$, $Ar_k = \|r_k\|(\gamma_{k+1} v_{k+1} + \delta_{k+2} v_{k+2})$ and $\psi_k = \|r_k\| \|\begin{bmatrix} \gamma_{k+1} & \delta_{k+2} \end{bmatrix}\|$, where $\delta_{k+2} = 0$ if $k = \ell - 1$.
- If $\text{rank}(L_\ell) = \ell$, then $Ar_\ell = 0$ and $\psi_\ell = 0$.
- If $\text{rank}(L_\ell) = \ell - 1$, then $Ar_\ell = Ar_{\ell-1} = 0$, and $\|\psi_\ell\| = \psi_{\ell-1} = 0$.

Proof. For the first case, the proof is essentially the same as the proof of Lemma 3.3. For the other two cases, the results follow directly from Lemma 6.1. \square

TABLE 6.1

Stopping conditions in MINRES-QLP. NRBE means normwise relative backward error, and tol , $maxit$, $maxcond$, and $maxnorm$ are input parameters. All norms and $\kappa(A)$ are estimated by MINRES-QLP.

Lanczos	NRBE	Regularization attempts
$\beta_{k+1} \leq n\ A\ \varepsilon$	$\ r_k\ / (\ A\ \ x_k\ + \ b\) \leq tol$	$\kappa(A) \geq maxcond$
$k = maxit$	$\ Ar_k\ / (\ A\ \ r_k\) \leq tol$	$\ x_k\ \geq maxnorm$

6.3. Matrix norms. For Lanczos-based algorithms, $\|A\| \geq \|V_{k+1}^T AV_k\| = \|\underline{T}_k\|$. Define

$$\mathcal{A}^{(0)} \equiv 0, \quad \mathcal{A}^{(k)} \equiv \max_{j=1,\dots,k} \left\{ \|\underline{T}_j e_j\| \right\} = \max \left\{ \mathcal{A}^{(k-1)}, \|\underline{T}_k e_k\| \right\} \text{ for } k \geq 1. \quad (6.1)$$

Then $\|A\| \geq \|\underline{T}_k\| \geq \mathcal{A}^{(k)}$. Clearly, $\mathcal{A}^{(k)}$ is monotonically increasing and is thus an improving estimate for $\|A\|$ as k increases. By the property of QLP decomposition in (4.2) and (5.2), we could easily extend (6.1) to include the largest diagonal of L_k :

$$\mathcal{A}^{(0)} \equiv 0, \quad \mathcal{A}^{(k)} \equiv \max \left\{ \mathcal{A}^{(k-1)}, \|\underline{T}_k e_k\|, \gamma_{k-2}^{(6)}, \gamma_{k-1}^{(5)}, |\gamma_k^{(4)}| \right\} \text{ for } k \geq 1. \quad (6.2)$$

Some other schemes inspired by Larsen [31, section A.6.1], Higham [25], and Chen and Demmel [8] follow. For the latter scheme, we use an implementation by Kaustuv [28] for estimating the norms of the rows of A .

1. [31] $\|T_k\|_1 \geq \|\underline{T}_k\|$
2. [31] $\sqrt{\|T_k^T \underline{T}_k\|_1} \geq \|T_k\|$
3. [31] $\|T_j\| \leq \|T_k\|$ for small $j = 5$ or 20
4. [25] MATLAB function NORMEST(A), which is based on the power method
5. [8] $\max_i \|h_i\| / \sqrt{m}$, where h_i^T is the i th row of AZ , each column of $Z \in \mathbb{R}^{n \times m}$ is a random vector of ± 1 's, and m is a small integer (e.g., $m = 10$).

Figure 6.1 plots estimates of $\|A\|$ for 12 matrices from the Florida sparse matrix collection [54] whose sizes n vary from 25 to 3002. In particular, scheme 3 above with $j = 20$ gives significantly more accurate estimates than other schemes for the 12 matrices we tried. However, the choice of j is not always clear and the scheme adds a little to the cost of MINRES-QLP. Hence we propose incorporating it into MINRES-QLP (or other Lanczos-based iterative methods) if very accurate $\|A\|$ is needed. Otherwise (6.2) uses quantities readily available from MINRES-QLP and gives us satisfactory estimates for the order of $\|A\|$.

6.4. Matrix condition numbers. We again apply the property of the QLP decomposition in (4.2) and (5.2) to estimate $\kappa(\underline{T}_k)$, which is a lower bound for $\kappa(A)$:

$$\begin{aligned} \gamma_{\min} &\leftarrow \min\{\gamma_1, \gamma_2^{(2)}\}, & \gamma_{\min} &\leftarrow \min\{\gamma_{\min}, \gamma_{k-2}^{(6)}, \gamma_{k-1}^{(5)}, |\gamma_k^{(4)}|\} \text{ for } k \geq 3, \\ \kappa^{(0)} &\equiv 1, & \kappa^{(k)} &\equiv \max \left\{ \kappa^{(k-1)}, \frac{\mathcal{A}^{(k)}}{\gamma_{\min}} \right\} \text{ for } k \geq 1. \end{aligned} \quad (6.3)$$

6.5. Solution norms. We derive a recurrence relation for $\|x_k\|$ whose cost is as low as computing the norm of a 3- or 4- vector.

Since $\|x_k\| = \|V_k P_k u_k\| = \|u_k\|$, we can estimate $\|x_k\|$ by computing $\chi_k \equiv \|u_k\|$. However, the last two elements of u_k change in u_{k+1} (and a new element μ_{k+1} is




FIG. 6.1. Relative errors in different estimates of $\|A\|$. This figure can be reproduced by `testminresQLPNormA8`.

added). We therefore maintain χ_{k-2} by updating it and then using it according to

$$\chi_{k-2}^{(2)} = \|\chi_{k-3}^{(2)} \mu_{k-2}^{(3)}\|, \quad \chi_k = \|\chi_{k-2}^{(2)} \mu_{k-1}^{(2)} \mu_k\| \quad \text{cf. (5.8) and (5.9)}.$$

Thus $\chi_{k-2}^{(2)}$ increases monotonically but we cannot guarantee that $\|x_k\|$ and its recurred estimate χ_k are increasing, and indeed they are not in some examples (see Figure 8.1).

6.6. Projection norms. Sometimes the projection of the right-hand side vector b onto $\mathcal{K}_k(A, b)$ is required (for example, see [46]). A simple recurrence relation is $\omega_k^2 \equiv \|Ax_k\|^2 = \omega_{k-1}^2 + \tau_k^2$ and we can derive it in the same way as shown in Lemma 3.3. With $\omega_0 \equiv 0$ we have $\omega_k \equiv \|Ax_k\| = \|\omega_{k-1} \tau_k\|$.

7. Preconditioned MINRES and MINRES-QLP. It is often asked: How can we construct a preconditioner for a linear system solver so that the same problem is solved with fewer iterations? Previous work on preconditioning the symmetric solvers CG, SYMMLQ, or MINRES includes [43, 37, 17, 12, 14, 35, 42, 34, 20, 2, 52].

We have the same question for singular symmetric systems $Ax \approx b$. Two-sided preconditioning is needed to preserve symmetry. We can still solve compatible systems, but we will no longer obtain the minimum-length solution. For incompatible systems, preconditioning alters the “least squares” norm. To avoid this difficulty we must work with larger equivalent systems that are compatible. We consider each case in turn, using a positive-definite preconditioner $M = CC^T$ with MINRES and MINRES-QLP to solve symmetric compatible systems $Ax = b$. Implicitly, we are solving equivalent symmetric systems $C^{-1}AC^{-T}y = C^{-1}b$, where $C^T x = y$. As usual, it is possible to work with M itself, so without loss of generality we can assume $C = M^{\frac{1}{2}}$.

7.1. Derivation. We derive preconditioned MINRES for compatible $Ax = b$ by applying MINRES to the equivalent problem $\bar{A}\bar{x} = \bar{b}$, where $\bar{A} \equiv M^{-\frac{1}{2}}AM^{-\frac{1}{2}}$, $\bar{b} \equiv M^{-\frac{1}{2}}b$, and $x = M^{-\frac{1}{2}}\bar{x}$.

7.1.1. Preconditioned Lanczos process. Let v_k denote the Lanczos vectors of $\mathcal{K}(\bar{A}, \bar{b})$. With $v_0 = 0$ and $\beta_1 v_1 = \bar{b}$, for $k = 1, 2, \dots$ we define

$$z_k = \beta_k M^{\frac{1}{2}} v_k, \quad q_k = \beta_k M^{-\frac{1}{2}} v_k, \quad \text{so that } Mq_k = z_k. \quad (7.1)$$

Then $\beta_k = \|\beta_k v_k\| = \|M^{-\frac{1}{2}} z_k\| = \|z_k\|_{M^{-1}} = \|q_k\|_M = \sqrt{q_k^T z_k}$, where the square root is well defined because M is positive definite, and the Lanczos iteration is

$$\begin{aligned} p_k &= \bar{A}v_k = M^{-\frac{1}{2}}AM^{-\frac{1}{2}}v_k = M^{-\frac{1}{2}}Aq_k/\beta_k, \\ \alpha_k &= v_k^T p_k = q_k^T Aq_k/\beta_k^2, \\ \beta_{k+1}v_{k+1} &= M^{-\frac{1}{2}}AM^{-\frac{1}{2}}v_k - \alpha_k v_k - \beta_k v_{k-1}. \end{aligned}$$

Multiplying the last equation by $M^{\frac{1}{2}}$ we get

$$\begin{aligned} z_{k+1} &= \beta_{k+1}M^{\frac{1}{2}}v_{k+1} = AM^{-\frac{1}{2}}v_k - \alpha_k M^{\frac{1}{2}}v_k - \beta_k M^{\frac{1}{2}}v_{k-1} \\ &= \frac{1}{\beta_k}Aq_k - \frac{\alpha_k}{\beta_k}z_k - \frac{\beta_k}{\beta_{k-1}}z_{k-1}. \end{aligned}$$

The last expression involving consecutive z_j 's replaces the three-term recurrence in v_j 's. In addition, we need to solve a linear system $Mq_k = z_k$ (7.1) each iteration.

7.1.2. Preconditioned MINRES. From (3.4) and (3.8) we have the following recurrence for the k th column of $D_k = V_k R_k^{-1}$ and \bar{x}_k :

$$d_k = (v_k - \delta_k^{(2)}d_{k-1} - \epsilon_k d_{k-2})/\gamma_k^{(2)}, \quad \bar{x}_k = \bar{x}_{k-1} + \tau_k d_k.$$

Multiplying the above two equations by $M^{-\frac{1}{2}}$ on the left and defining $\bar{d}_k = M^{-\frac{1}{2}}d_k$, we can update the solution of our original problem by

$$\bar{d}_k = \left(\frac{1}{\beta_k}q_k - \delta_k^{(2)}\bar{d}_{k-1} - \epsilon_k \bar{d}_{k-2} \right) / \gamma_k^{(2)}, \quad x_k = M^{-\frac{1}{2}}\bar{x}_k = x_{k-1} + \tau_k \bar{d}_k.$$

We list the algorithm in [10, Table 3.4].

7.1.3. Preconditioned MINRES-QLP. A preconditioned MINRES-QLP can be derived very similarly. The additional work is to apply right reflectors P_k to R_k , and the new subproblem bases are $W_k \equiv V_k P_k$, with $\bar{x}_k = W_k u_k$. Multiplying the new basis and solution estimate by $M^{-\frac{1}{2}}$ on the left, we obtain

$$\begin{aligned} \bar{W}_k &\equiv M^{-\frac{1}{2}}W_k = M^{-\frac{1}{2}}V_k P_k, \\ x_k &= M^{-\frac{1}{2}}\bar{x}_k = M^{-\frac{1}{2}}W_k u_k = \bar{W}_k u_k = x_{k-2}^{(2)} + \mu_{k-1}^{(2)}\bar{w}_{k-1}^{(3)} + \mu_k \bar{w}_k^{(2)}. \end{aligned}$$

Algorithm 1 lists all steps. Note that \bar{w}_k is written as w_k for all relevant k . Also, the output x solves $Ax \approx b$ but the other outputs are associated with $\bar{A}\bar{x} \approx \bar{b}$.

Remark. The requirement of positive-definite preconditioners M in MINRES and MINRES-QLP may seem unnatural for a problem with indefinite A because we cannot achieve $M^{-\frac{1}{2}}AM^{-\frac{1}{2}} \approx I$. However, as shown in [17], we can achieve $M^{-\frac{1}{2}}AM^{-\frac{1}{2}} \approx \begin{bmatrix} I & \\ & -I \end{bmatrix}$ using an approximate block-LDL^T factorization $A \approx LDL^T$ to get $M = L|D|L^T$, where D is indefinite with blocks of order 1 and 2, and $|D|$ has the same eigensystem as D except negative eigenvalues are changed in sign.

SQMR [15] without preconditioning is analytically equivalent to MINRES. Unlike MINRES, SQMR can work directly with an indefinite preconditioner (such as block-LDL^T). However, in finite precision, SQMR needs “look-ahead” to prevent numerical breakdown.

Algorithm 1: Preconditioned MINRES-QLP to solve $(A - \sigma I)x \approx b$.

input: A, b, σ, M

1 $z_0 = 0, \quad z_1 = b, \quad \text{Solve } Mq_1 = z_1, \quad \beta_1 = \sqrt{b^T q_1} \quad [\text{Initialize}]$

2 $w_0 = w_{-1} = 0, \quad x_{-2} = x_{-1} = x_0 = 0$

3 $c_{0,1} = c_{0,2} = c_{0,3} = -1, \quad s_{0,1} = s_{0,2} = s_{0,3} = 0, \quad \phi_0 = \beta_1, \quad \tau_0 = \omega_0 = \chi_{-2} = \chi_{-1} = \chi_0 = 0$

4 $\delta_1 = \gamma_{-1} = \gamma_0 = \eta_{-1} = \eta_0 = \eta_1 = \vartheta_{-1} = \vartheta_0 = \vartheta_1 = \mu_{-1} = \mu_0 = 0, \quad \mathcal{A} = 0, \quad \kappa = 1$

5 $k = 0$

6 **while** no stopping condition is satisfied **do**

7 $k \leftarrow k + 1$

8 $p_k = Aq_k - \sigma q_k, \quad \alpha_k = \frac{1}{\beta_k^2} q_k^T p_k \quad [\text{Preconditioned Lanczos}]$

9 $z_{k+1} = \frac{1}{\beta_k} p_k - \frac{\alpha_k}{\beta_k} z_k - \frac{\beta_k}{\beta_{k-1}} z_{k-1}$

10 **Solve** $Mq_{k+1} = z_{k+1}, \quad \beta_{k+1} = \sqrt{q_{k+1}^T z_{k+1}}$

11 **if** $k = 1$ **then** $\rho_k = \|[\alpha_k \ \beta_{k+1}]\|$ **else** $\rho_k = \|[\beta_k \ \alpha_k \ \beta_{k+1}]\|$

12 $\delta_k^{(2)} = c_{k-1,1} \delta_k + s_{k-1,1} \alpha_k \quad [\text{Previous left reflection...}]$

13 $\gamma_k = s_{k-1,1} \delta_k - c_{k-1,1} \alpha_k \quad [\text{on middle two entries of } \underline{T}_k e_k \dots]$

14 $\epsilon_{k+1} = s_{k-1,1} \beta_{k+1} \quad [\text{produces first two entries in } \underline{T}_{k+1} e_{k+1}]$

15 $\delta_{k+1} = -c_{k-1,1} \beta_{k+1}$

16 $c_{k1}, s_{k1}, \gamma_k^{(2)} \leftarrow \text{SymOrtho}(\gamma_k, \beta_{k+1}) \quad [\text{Current left reflection}]$

17 $c_{k2}, s_{k2}, \gamma_{k-2}^{(6)} \leftarrow \text{SymOrtho}(\gamma_{k-2}^{(5)}, \epsilon_k) \quad [\text{First right reflection}]$

18 $\delta_k^{(3)} = s_{k2} \vartheta_{k-1} - c_{k2} \delta_k^{(2)}, \quad \gamma_k^{(3)} = -c_{k2} \gamma_k^{(2)}, \quad \eta_k = s_{k2} \gamma_k^{(2)}$

19 $\vartheta_{k-1}^{(2)} = c_{k2} \vartheta_{k-1} + s_{k2} \delta_k^{(2)}$

20 $c_{k3}, s_{k3}, \gamma_{k-1}^{(5)} \leftarrow \text{SymOrtho}(\gamma_{k-1}^{(4)}, \delta_k^{(3)}) \quad [\text{Second right reflection...}]$

21 $\vartheta_k = s_{k3} \gamma_k^{(3)}, \quad \gamma_k^{(4)} = -c_{k3} \gamma_k^{(3)} \quad [\text{to zero out } \delta_k^{(3)}]$

22 $\tau_k = c_{k1} \phi_{k-1} \quad [\text{Last element of } t_k]$

23 $\phi_k = s_{k1} \phi_{k-1}, \quad \psi_{k-1} = \phi_{k-1} \|[\gamma_k \ \delta_{k+1}]\| \quad [\text{Update } \|r_k\|, \|Ar_{k-1}\|]$

24 **if** $k = 1$ **then** $\gamma_{\min} = \gamma_1$ **else** $\gamma_{\min} \leftarrow \min\{\gamma_{\min}, \gamma_{k-2}^{(6)}, \gamma_{k-1}^{(5)}, |\gamma_k^{(4)}|\}$

25 $\mathcal{A}^{(k)} = \max\{\mathcal{A}^{(k-1)}, \rho_k, \gamma_{k-2}^{(6)}, \gamma_{k-1}^{(5)}, |\gamma_k^{(4)}|\} \quad [\text{Update } \|A\|]$

26 $\omega_k = \|[\omega_{k-1} \ \tau_k]\|, \quad \kappa \leftarrow \mathcal{A}^{(k)} / \gamma_{\min} \quad [\text{Update } \|Ax_k\|, \kappa(A)]$

27 $w_k = -(c_{k2} / \beta_k) q_k + s_{k2} w_{k-2}^{(3)} \quad [\text{Update } w_{k-2}, w_{k-1}, w_k]$

28 $w_{k-2}^{(4)} = (s_{k2} / \beta_k) q_k + c_{k2} w_{k-2}^{(3)}$

29 **if** $k > 2$ **then** $w_k^{(2)} = s_{k3} w_{k-1}^{(2)} - c_{k3} w_k, \quad w_{k-1}^{(3)} = c_{k3} w_{k-1}^{(2)} + s_{k3} w_k$

30 **if** $k > 2$ **then** $\mu_{k-2}^{(3)} = (\tau_{k-2} - \eta_{k-2} \mu_{k-4}^{(4)} - \vartheta_{k-2} \mu_{k-3}^{(3)}) / \gamma_{k-2}^{(6)} \quad [\text{Update } \mu_{k-2}]$

31 **if** $k > 1$ **then** $\mu_{k-1}^{(2)} = (\tau_{k-1} - \eta_{k-1} \mu_{k-3}^{(3)} - \vartheta_{k-1} \mu_{k-2}^{(3)}) / \gamma_{k-1}^{(5)} \quad [\text{Update } \mu_{k-1}]$

32 **if** $\gamma_k^{(4)} \neq 0$ **then** $\mu_k = (\tau_k - \eta_k \mu_{k-2}^{(3)} - \vartheta_k \mu_{k-1}^{(2)}) / \gamma_k^{(4)}$ **else** $\mu_k = 0$ **[Compute** $\mu_k]$

33 $x_{k-2}^{(2)} = x_{k-3}^{(2)} + \mu_{k-2}^{(3)} w_{k-2}^{(3)} \quad [\text{Update } x_{k-2}]$

34 $x_k = x_{k-2}^{(2)} + \mu_{k-1}^{(2)} w_{k-1}^{(3)} + \mu_k w_k^{(2)} \quad [\text{Compute } x_k]$

35 $\chi_{k-2}^{(2)} = \|[\chi_{k-3}^{(2)} \ \mu_{k-2}^{(3)}]\| \quad [\text{Update } \|x_{k-2}\|]$

36 $\chi_k = \|[\chi_{k-2}^{(2)} \ \mu_{k-1}^{(2)} \ \mu_k]\| \quad [\text{Compute } \|x_k\|]$

37 $x = x_k, \quad \phi = \phi_k, \quad \psi = \phi_k \|[\gamma_{k+1} \ \delta_{k+2}]\|, \quad \chi = \chi_k, \quad \mathcal{A} = \mathcal{A}^{(k)}, \quad \omega = \omega_k$

output: $x, \phi, \psi, \chi, \mathcal{A}, \kappa, \omega$

38

$[c, s \leftarrow \text{SymOrtho}(a, b)]$ is a stable form for computing $r = \sqrt{a^2 + b^2}$, $c = \frac{a}{r}$, $s = \frac{b}{r}$

7.2. Preconditioning singular $Ax = b$. For singular compatible systems, MINRES and MINRES-QLP find the minimum-length solution (see Theorems 3.1 and 5.1). If M is nonsingular, the preconditioned system is also compatible and the solvers return its minimum-length solution. The unpreconditioned solution solves $Ax \approx b$, but is not necessarily a minimum-length solution.

EXAMPLE 7.1. Let $A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$ and $b = \begin{bmatrix} 6 \\ 9 \\ 6 \\ 3 \end{bmatrix}$. Then $\text{rank}(A) = 3$ and $Ax = b$ is a singular compatible system. The minimum-length solution is $x^\dagger = [2 \ 4 \ 3 \ 2]^T$. By binormalization [33] we construct the matrix $D = \text{diag}([0.84201 \ 0.81228 \ 0.30957 \ 3.2303])$. The minimum-length solution of the diagonally preconditioned problem $DADy = Db$ is $y^\dagger = [3.5739 \ 3.6819 \ 9.6909 \ 0.93156]^T$. Then $x = Dy^\dagger = [3.0092 \ 2.9908 \ 3.0000 \ 3.0092]^T$ is a solution of $Ax = b$, but $x \neq x^\dagger$.

7.3. Preconditioning singular $Ax \approx b$. We propose the following techniques for obtaining minimum-residual solutions of singular incompatible problems. In each case we use an equivalent but larger *compatible* system to which MINRES may be applied. Even if the larger system is singular, Theorem 3.1 shows that the minimum-length solution of the larger system will be obtained. The required x will be part of this solution. Preconditioning still gives a minimum-residual solution of $Ax \approx b$, and in *some* cases x will be x^\dagger . If the systems are ill-conditioned, it will be safer and more efficient to apply MINRES-QLP to the original incompatible system. However, preconditioning will give an x that is “minimum length” in a different norm.

7.3.1. Augmented system. When A is singular, so is the augmented system

$$\begin{bmatrix} I & A \\ A & \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}, \quad (7.2)$$

but it is always compatible. Preconditioning with symmetric positive-definite M gives us a solution $\begin{bmatrix} r \\ x \end{bmatrix}$ in which r is unique, but x may not be x^\dagger .

7.3.2. A giant KKT system. Problem (1.1) is equivalent to $\min_{r,x} x^T x$ subject to (7.2), which is an equality-constrained convex quadratic program. The corresponding KKT system [36, section 16.1] is both symmetric and compatible:

$$\begin{bmatrix} & & I & A \\ & -I & A & \\ I & A & & \\ A & & & \end{bmatrix} \begin{bmatrix} r \\ x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ b \\ 0 \end{bmatrix}. \quad (7.3)$$

Although this is still a singular system, the upper-left 3×3 block-submatrix is nonsingular and therefore r , x , and y are unique and a preconditioner applied to the KKT system would give x as the minimum-length solution of our original problem.

7.3.3. Regularization. If the rank of a given matrix A is ill-determined, we may want to solve the *regularized* problem [13, 22] with parameter $\delta > 0$:

$$\min_x \left\| \begin{bmatrix} A \\ \delta I \end{bmatrix} x - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|^2. \quad (7.4)$$

The matrix $\begin{bmatrix} A \\ \delta I \end{bmatrix}$ has full rank and is always better conditioned than A . LSQR [40, 41] may be applied, and its iterates x_k will reduce $\|r_k\|^2 + \delta^2 \|x_k\|^2$ monotonically. Alternatively, we could transform (7.4) into the following symmetric compatible systems and apply MINRES or MINRES-QLP. They tend to reduce $\|Ar_k - \delta^2 x_k\|$ monotonically.

Normal equation:

$$(A^2 + \delta^2 I)x = Ab. \quad (7.5)$$

Augmented system:

$$\begin{bmatrix} I & A \\ A & -\delta^2 I \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

A two-layered problem: If we eliminate v from the system

$$\begin{bmatrix} I & A^2 \\ A^2 & -\delta^2 A^2 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} = \begin{bmatrix} 0 \\ Ab \end{bmatrix}. \quad (7.6)$$

we obtain (7.5). Thus x is also a solution of our regularized problem (7.4). This is equivalent to the two-layered formulation (4.3) in Bobrovnikova and Vavasis [5] (with $A_1 = A$, $A_2 = D_1 = D_2 = I$, $b_1 = b$, $b_2 = 0$, $\delta_1 = 1$, $\delta_2 = \delta^2$). A key property is that $x \rightarrow x^\dagger$ as $\delta \rightarrow 0$.

A KKT-like system: If we define $y = -Av$ and $r = b - Ax - \delta^2 y$, then we can show (by eliminating r and y from the following system) that x in

$$\begin{bmatrix} I & A \\ -I & A \\ I & A & \delta^2 I \\ A \end{bmatrix} \begin{bmatrix} r \\ x \\ y \\ v \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ b \\ 0 \end{bmatrix} \quad (7.7)$$

is also a solution of (7.6) and thus of (7.4). The upper-left 3×3 block-submatrix of (7.7) is nonsingular, and the correct limiting behavior occurs: $x \rightarrow x^\dagger$ as $\delta \rightarrow 0$. In fact, (7.7) reduces to (7.3).

7.4. General preconditioners. The construction of preconditioners is usually problem-dependent. If not much is known about the structure of A , we can only consider general methods such as diagonal preconditioning and incomplete Cholesky factorization. These methods require access to the nonzero elements of A . (They are not applicable if A exists only as an operator for returning the product Ax .)

For a comprehensive survey of preconditioning techniques, see Benzi [3]. We discuss a few methods for symmetric A that also require access to the nonzero A_{ij} .

7.4.1. Diagonal preconditioning. If A has entries that are very different in magnitude, diagonal scaling might improve its condition. When A is diagonally dominant and nonsingular, we can define $D = \text{diag}(d_1, \dots, d_n)$ with $d_j = 1/|A_{jj}|^{1/2}$. Instead of solving $Ax = b$, we solve $DADy = Db$, where DAD is still diagonally dominant and nonsingular with all entries ≤ 1 in magnitude, and $x = Dy$.

More generally, if A is not diagonally dominant and possibly singular, we can safeguard division-by-zero errors by choosing a parameter $\delta > 0$ and defining

$$d_j(\delta) = 1/\max\{\delta, \sqrt{|A_{jj}|}, \max_{i \neq j} |A_{ij}|\}, \quad j = 1, \dots, n. \quad (7.8)$$

EXAMPLE 7.2.

1. If $A = \begin{bmatrix} -1 & 10^{-8} \\ 10^{-8} & 1 & 10^4 \\ & 10^4 & 0 \end{bmatrix}$, then $\kappa(A) \approx 10^4$. Let $\delta = 1$, $D = \begin{bmatrix} 1 & & \\ & 10^{-2} & \\ & & 10^{-2} & 1 \end{bmatrix}$
in (7.8). Then $DAD = \begin{bmatrix} -1 & 10^{-10} \\ 10^{-10} & 10^{-4} & 1 \\ & 1 & 0 \end{bmatrix}$ and $\kappa(DAD) \approx 1$.

2. $A = \begin{bmatrix} 10^{-4} & 10^{-8} & & & \\ 10^{-8} & 10^{-4} & 10^{-8} & & \\ & 10^{-8} & & 10^{-8} & \\ & & & & 0 \\ & & & & & 0 \end{bmatrix}$ contains mostly very small entries, and $\kappa(A) \approx 10^{10}$.
 Let $\delta = 10^{-8}$ and $D = \begin{bmatrix} 10^2 & & & & \\ & 10^2 & & & \\ & & 10^8 & & \\ & & & 10^8 & \\ & & & & 10^8 \end{bmatrix}$. Then $DAD = \begin{bmatrix} 1 & 10^{-4} & & & \\ 10^{-4} & 1 & 10^2 & & \\ & 10^2 & 1 & 10^2 & \\ & & & 1 & 0 \\ & & & & 0 \end{bmatrix}$
 and $\kappa(DAD) \approx 10^2$. (The choice of δ makes a critical difference in this case: with $\delta = 1$, we have $D = I$.)

7.4.2. Binormalization (BIN). Livne and Golub [33] scale a symmetric matrix by a series of k diagonal matrices on both sides until all rows and columns of the scaled matrix have unit 2-norm: $DAD = D_k \cdots D_1 A D_1 \cdots D_k$. See also Bradley [6].

EXAMPLE 7.3. If $A = \begin{bmatrix} 10^{-8} & & & & \\ 1 & 10^{-8} & 10^4 & & \\ & 10^4 & & 10^4 & \\ & & & & 0 \end{bmatrix}$, then $\kappa(A) \approx 10^{12}$. With just one sweep of BIN, we obtain $D = \text{diag}(8.1e-3, 6.6e-5, 1.5)$, $DAD \approx \begin{bmatrix} 6.5e-1 & 5.3e-1 & 0 \\ 5.3e-1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$ and $\kappa(DAD) \approx 2.6$ even though the rows and columns have not converged to one in the two-norm. In contrast, diagonal scaling (7.8) defined by $\delta = 1$ and $D = \text{diag}(1, 10^{-4}, 10^{-4})$ reduces the condition number to approximately 10^4 .

7.4.3. Incomplete Cholesky factorization. For a sparse symmetric positive definite matrix A , we could compute a preconditioner by the incomplete Cholesky factorization that preserves the sparsity pattern of A . This is known as IC0 in the literature. Often there exists a permutation P such that the IC0 factor of PAP^T is more sparse than that of A .

When A is semidefinite or indefinite, IC0 may not exist, but a simple variant that may work is the incomplete Cholesky-infinity factorization [55, section 5].

8. Numerical experiments. We compare the computed results of MINRES-QLP and various other Krylov subspace methods to solutions computed directly by the eigenvalue decomposition (EVD) and the truncated eigenvalue decompositions (TEVD) of A . For TEVD we have

$$x_t \equiv \sum_{|\lambda_i| > t \|A\|_\varepsilon} \frac{1}{\lambda_i} u_i u_i^T b, \quad \|A\| = \max |\lambda_i|, \quad \kappa_t(A) = \frac{\max |\lambda_i|}{\min_{|\lambda_i| > t \|A\|_\varepsilon} |\lambda_i|}$$

with parameter $t > 0$. Often t is set to 1, and sometimes to a moderate number such as 10 or 100; it defines a cut-off point relative to the largest eigenvalue of A . For example, if most eigenvalues are of order 1 in magnitude and the rest are of order $\|A\|_\varepsilon \approx 10^{-16}$, we expect TEVD to work better when the small eigenvalues are excluded, while EVD (with $t = 0$) could return an exploding solution.

In the tables of results, MATLAB MINRES and MATLAB SYMMLQ are MATLAB's implementation of MINRES and SYMMLQ respectively. They incorporate *local reorthogonalization* of the Lanczos vector v_2 , which could enhance the accuracy of the computations if b is close to an eigenvector of A [32]:

$$\text{Second Lanczos iteration: } \beta_1 v_1 = b, \text{ and } q_2 \equiv \beta_2 v_2 = Av_1 - \alpha_1 v_1$$

$$\text{Local reorthogonalization: } q_2 \leftarrow q_2 - (v_1^T q_2) v_1.$$

Lacking the correct stopping condition for singular problems, MATLAB SYMMLQ works more than necessary and then selects the smallest residual norm from all computed iterates; it would sometimes report that the method did not converge although the selected estimate appeared to be reasonably accurate.

MINRES SOL and SYMMLQ SOL are implementations based on [39]. MINRES⁺ and SYMMLQ⁺ are the same but with additional stopping conditions for singular incompatible systems (see Lemma 3.3 and [10, Proposition 2.12]).

The computations in this section were performed on a Windows XP machine with a 3.2GHz Intel Pentium D Processor 940 and 3GB RAM ($\varepsilon \approx 10^{-16}$). Tests were performed with each solver on five types of problem:

1. symmetric, nonsingular linear systems
2. symmetric, singular linear systems
3. mildly incompatible symmetric (and singular) systems (meaning $\|r\|$ is rather small with respect to $\|b\|$)
4. symmetric (and singular) LS problems
5. Hermitian systems.

We present a few examples that illustrate the key features of MINRES-QLP. For a larger set of tests and results, such as applying MINRES-QLP and other Krylov methods to Hermitian systems with preconditioning, we refer to [10, Chapter 4].

For a compatible system, we generate a random vector b that is in the range of the test matrix ($b \equiv Ay$, $y_i \sim i.i.d. U(0, 1)$, i.e., y_1, \dots, y_n are independent and identically distributed random variables, whose values are drawn from the standard uniform distribution with support $[0, 1]$). For an LS problem, we generate a random b that is *not* in the range of the test matrix ($b_i \sim i.i.d. U(0, 1)$ often suffices).

If A is Hermitian, then $v^H Av$ is real for all complex vectors v . Numerically (in double precision), $\alpha_k = v_k^H Av_k$ is likely to have small imaginary parts in the first few Lanczos iterations and snowball to have large imaginary parts in later iterations. This would result in a poor estimation of $\|T_k\|_F$ or $\|A\|_F$, and unnecessary errors in the Lanczos vectors. Thus we made sure to *typecast* $\alpha_k = \text{real}(v_k^H Av_k)$ in MINRES-QLP and MINRES SOL.

We could say from the results that the Lanczos-based methods have built-in regularization features [29], often matching the TEVD solutions very well.

8.1. A Laplacian system $Ax \approx b$ (almost compatible). Our first example involves a singular indefinite Laplacian matrix A of order $n = 400$. It is block-tridiagonal with each block being a tridiagonal matrix T of order $N = 20$ with all nonzeros equal to 1:

$$A = \begin{bmatrix} T & T & & & \\ T & T & \ddots & & \\ & \ddots & \ddots & T & \\ & & T & T & \end{bmatrix}_{n \times n}, \quad T = \begin{bmatrix} 1 & 1 & & & \\ 1 & 1 & \ddots & & \\ & \ddots & \ddots & 1 & \\ & & & 1 & 1 \end{bmatrix}_{N \times N}. \quad (8.1)$$

MATLAB's `eig(A)` reports the following data: 205 positive eigenvalues in the interval $[6.1e-2, 8.87]$, 39 almost-zero eigenvalues in $[-2.18e-15, 3.71e-15]$, 156 negative eigenvalues in $[-2.91, -6.65e-2]$, numerical rank = 361.

We used a right-hand side with a small incompatible component: $b = Ay + 10^{-8}z$ with y_i and $z_i \sim i.i.d. U(0, 1)$. Results are summarized in Table 8.1. In the column labeled "C?", the value "Y" denotes that the associated algorithm in the row has converged to the desired NRBE tolerances within *maxit* iterations (cf. Table 6.1); otherwise, we have values "N" and "N?", where "N?" indicates that the algorithm could have converged if more relaxed stopping conditions were used. The column "Av" shows the total number of matrix-vector products, and column "x(1)" lists the

TABLE 8.1

Finite element problem $Ax \approx b$ with b almost compatible. Laplacian on a 20×20 grid, $n = 400$, $maxit = 1200$, $shift = 0$, $tol = 1.0e-15$, $maxnorm = 100$, $maxcond = 1e15$, $\|b\| = 87$. To reproduce this example, run `test_minresqlp_eg7.1(24)`.

Method	C?	Av	$x(1)$	$\ x\ $	$\ e\ $	$\ r\ $	$\ Ar\ $	$\ A\ $	$\kappa(A)$
EVD	-	-	-7.39e5	4.12e7	4.1e7	1.7e-7	7.8e-7	8.9e0	1.1e17
TEVD	-	-	3.89e-1	1.15e1	0.0e0	1.7e-8	1.4e-12	8.9e0	1.5e2
MATLAB SYMMLQ	N?	371	3.89e-1	1.15e1	1.4e-7	1.8e-7	5.8e-7	-	-
SYMMLQ SOL	N	447	-3.08e0	9.63e1	9.5e1	1.4e2	4.4e2	9.6e1	1.3e1
SYMMLQ ⁺	N	447	2.94e6	4.27e8	4.3e8	1.8e2	6.5e2	8.6e0	1.3e1
MATLAB MINRES	N	1200	-7.50e5	2.10e7	2.1e7	1.5e7	9.1e7	-	-
MINRES SOL	N	1200	9.89e5	6.10e7	6.1e7	2.3e7	1.5e8	1.8e2	1.5e1
MINRES ⁺	N	611	1.02e0	9.28e1	9.2e1	1.7e-8	2.5e-11	8.6e0	6.9e13
MINRES-QLP	Y	612	3.89e-1	1.15e1	3.7e-11	1.7e-8	9.3e-11	8.7e0	4.3e13
MATLAB LSQR	Y	1462	3.89e-1	1.15e1	2.3e-13	1.7e-8	3.3e-13	-	-
LSQR SOL	Y	1464	3.89e-1	1.15e1	2.4e-13	1.7e-8	3.9e-13	1.5e2	6.4e3
MATLAB GMRES(30)	N?	1200	3.90e-1	1.15e1	5.2e-2	3.4e-3	9.4e-4	-	-
SQMR	N	1200	-2.58e8	3.74e10	3.7e10	4.6e3	2.3e4	-	-
MATLAB QMR	N?	798	3.89e-1	1.15e1	5.2e-7	1.9e-8	2.6e-8	-	-
MATLAB BICG	N?	790	3.89e-1	1.15e1	4.7e-7	3.9e-8	1.9e-7	-	-
MATLAB BICGSTAB	N?	2035	3.89e-1	1.15e1	4.2e-7	1.7e-8	4.3e-13	-	-

first element of the final solution estimate x for each algorithm. For GMRES, the integer in parentheses is the value of the restart parameter.

MINRES SOL gives a larger solution than MINRES-QLP. This example has a residual norm of about 1.7×10^{-8} , so it is not clear whether to classify it as a linear system or an LS problem. To the credit of MATLAB SYMMLQ, it thinks the system is linear and returns a good solution. For MINRES-QLP, the first 410 iterations are in standard “MINRES mode”, with a transfer to “MINRES-QLP mode” for the last 202 iterations. LSQR [40, 41] converges to the minimum-length solution but with more than twice the number of iterations of MINRES-QLP. The other solvers fall short in some way.

8.2. A Laplacian LS problem $\min \|Ax - b\|$. This example uses the same Laplacian matrix A (8.1) but with a clearly incompatible $b = 10 \times \text{rand}(n, 1)$, i.e., $b_i \sim i.i.d. U(0, 10)$. The residual norm is about 17. Results are summarized in Table 8.2. MINRES gives an LS solution, while MINRES-QLP is the only solver that matches the solution of TEVD. The other solvers do not perform satisfactorily.

8.3. Regularizing effect of MINRES-QLP. This example illustrates the regularizing effect of MINRES-QLP with the stopping condition $\chi_k \leq maxnorm$. For $k \geq 18$ in Figure 8.1, we observe the following values:

$$\begin{aligned} \chi_{18} &= \|[2.51 \quad 3.87e-11 \quad 1.38 \times 10^2]\| = 1.38 \times 10^2, \\ \chi_{19} &= \|[2.51 \quad -8.00e-10 \quad -1.52 \times 10^2]\| = 1.52 \times 10^2, \\ \chi_{20} &= \|[2.51 \quad 1.62e-10 \quad -1.62 \times 10^6]\| = 1.62 \times 10^6 > maxnorm \equiv 10^4. \end{aligned}$$

Because the last value exceeds $maxnorm$, MINRES-QLP regards the last diagonal element of L_k as a singular value to be ignored (in the spirit of truncated SVD solutions). It discards the last element of u_{20} and updates

$$\chi_{20} \leftarrow \|[2.51 \quad 1.62e-10 \quad 0]\| = 2.51.$$

TABLE 8.2

Finite element problem $\min \|Ax - b\|$. Laplacian on a 20×20 grid, $n = 400$, $\text{maxit} = 500$, $\text{shift} = 0$, $\text{tol} = 1.0e-14$, $\text{maxnorm} = 1e4$, $\text{maxcond} = 1e14$, $\|b\| = 120$. To reproduce this example, run `test_minresqlp_eg7.1(25)`.

Method	C?	Av	$x(1)$	$\ x\ $	$\ e\ $	$\ r\ $	$\ Ar\ $	$\ A\ $	$\kappa(A)$
EVD	-	-	-7.39e14	4.12e16	4.1e16	1.8e2	7.9e2	8.9e0	1.1e17
TEVD	-	-	-8.75e0	1.43e2	0.0e0	1.7e1	4.1e-12	8.9e0	1.5e2
MATLAB SYMMLQ	N	1	2.74e-1	1.52e1	1.4e2	6.0e1	2.9e2	-	-
SYMMLQ SOL	N	228	-7.70e2	9.93e3	9.9e3	7.0e3	3.4e4	6.8e1	9.7e0
SYMMLQ ⁺	N	228	-7.70e2	9.93e3	9.9e3	7.0e3	3.4e4	7.6e0	9.7e0
MATLAB MINRES	N	500	2.80e14	4.07e16	4.1e16	2.3e2	1.4e3	-	-
MINRES SOL	N	500	-1.46e14	2.11e16	2.1e16	1.1e2	6.6e2	1.5e2	1.4e1
MINRES ⁺	N	381	3.88e1	6.90e3	6.9e3	1.7e1	1.2e-5	7.9e0	1.6e10
MINRES-QLP	Y	382	-8.75e0	1.43e2	1.7e-6	1.7e1	1.7e-5	8.6e0	3.5e10
MATLAB LSQR	Y	1000	-8.75e0	1.43e2	2.0e-5	1.7e1	1.4e-5	-	-
LSQR SOL	Y	1000	-8.75e0	1.43e2	2.3e-5	1.7e1	1.1e-5	1.2e2	4.4e3
MATLAB GMRES(30)	N	500	-8.84e0	1.25e2	4.8e1	1.7e1	8.2e-1	-	-
SQMR	N	500	9.58e15	1.39e18	1.4e18	1.2e11	6.7e11	-	-
MATLAB QMR	N	556	-7.30e0	1.92e2	1.4e2	1.7e1	1.2e1	-	-
MATLAB BICG	N	2	1.40e0	1.71e1	1.4e2	6.0e1	2.6e2	-	-
MATLAB BICGSTAB	N	104	-1.12e1	1.40e2	9.6e1	2.6e1	1.8e1	-	-

The full truncation strategy used in the implementation is justified by the fact that $x_k = W_k u_k$ with W_k orthogonal. When $\|x_k\|$ becomes large, the last element of u_k is treated as zero. If $\|x_k\|$ is still large, the second-to-last element of u_k is treated as zero. If $\|x_k\|$ is *still* large, the third-to-last element of u_k is treated as zero.

8.4. Effects of rounding errors in MINRES-QLP. The recurred residual norms ϕ_k^M in MINRES usually approximate the directly computed ones $\|r_k^M\|$ very well until $\|r_k^M\|$ becomes small. We observe that ϕ_k^M continues to decrease in the last few iterations, even though $\|r_k^M\|$ has become stagnant. This is desirable in the sense that the stopping rule will cause termination, although the final solution is not as accurate as predicted.

We present similar plots of MINRES-QLP in the following examples, with the corresponding quantities as ϕ_k^Q and $\|r_k^Q\|$. We observe that except in very ill-conditioned LS problems, ϕ_k^Q approximates $\|r_k^Q\|$ very closely.

Figure 8.2 illustrates four singular compatible linear systems.

Figure 8.3 illustrates four singular LS problems.

9. Conclusion. MINRES constructs its k th solution estimate from the recursion $x_k = D_k t_k = x_{k-1} + \tau_k d_k$ (3.4), where n separate triangular systems $R_k^T D_k^T = V_k^T$ are solved to obtain the n elements of each direction d_1, \dots, d_k . (Only d_k is obtained during iteration k , but it has n elements.)

In contrast, MINRES-QLP constructs its k th solution estimate using orthogonal steps: $x_k^Q = (V_k P_k) u_k$; see (5.3)–(5.9). Only one triangular system $L_k u_k = Q_k(\beta_1 e_1)$ is involved for each k .

Thus MINRES-QLP overcomes the potential instability predicted by the MINRES authors [39] and analyzed by Sleijpen et al. [47]. The additional work and storage are moderate, and maximum efficiency is retained by transferring from MINRES to the MINRES-QLP iterates only when the estimated condition number of A exceeds a specified value.

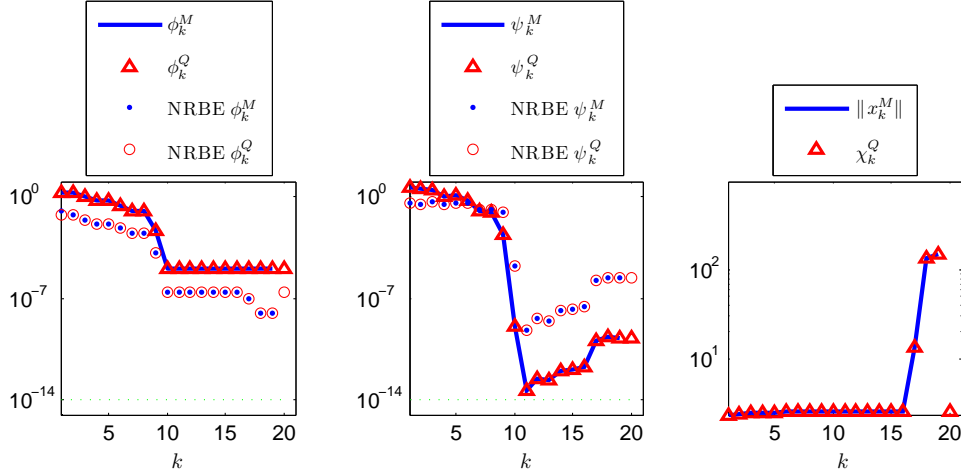


FIG. 8.1. Recurred $\phi_k \approx \|r_k\|$, $\psi_k \approx \|Ar_k\|$, and $\|x_k\|$ for MINRES and MINRES-QLP. The matrix A (ID 1177 from [54]) is positive semidefinite, $n = 25$, and b is random with $\|b\| \simeq 1.7$. Both solvers could have achieved essentially the TEVD solution of $Ax \simeq b$ at iteration 11. However, the stringent $\text{tol} = 10^{-14}$ on the recurred normwise relative backward errors (NRBE in Table 6.1) prevents them from stopping “in time”. MINRES ends with an exploding solution, yet MINRES-QLP brings it back to the TEVD solution at iteration 20. **Left:** ϕ_k^M and ϕ_k^Q (recurred $\|r_k\|$ of MINRES and MINRES-QLP) and their NRBE. **Middle:** ψ_k^M and ψ_k^Q (recurred $\|Ar_k\|$) and their NRBE. **Right:** $\|x_k^M\|$ (norms of solution estimates from MINRES) and χ_k^Q (recurred $\|x_k\|$ from MINRES-QLP) with $\text{maxnorm} = 10^4$. This figure can be reproduced by `test_minresqlp_fig7.1(2)`.

MINRES and MINRES-QLP are readily applicable to Hermitian matrices, once α_k is typecast as a real scalar in finite-precision arithmetic. For both algorithms, we derived recurrence relations for $\|Ar_k\|$ and $\|Ax_k\|$ and used them to formulate new stopping conditions for singular problems.

TEVD or TSVD are commonly known to use rank- k approximations to A to find approximate solutions to $\min \|Ax - b\|$ that serve as a form of *regularization*. Krylov subspace methods also have regularization properties [23, 21, 29]. Since MINRES-QLP monitors more carefully the rank of T_k , which could be k or $k-1$, we may say that regularization is a stronger feature in MINRES-QLP, as we have shown in our numerical examples.

It is important to develop robust techniques for estimating an a priori bound for the solution norm since the MINRES-QLP approximations are not monotonic as is the case in CG and LSQR. Ideally, we would also like to determine a practical threshold associated with the stopping condition $\gamma_k^{(4)} = 0$ in order to handle cases when $\gamma_k^{(4)}$ is numerically small but not exactly zero. These are topics for future research.

10. Software and reproducible research. MATLAB 7.6, Fortran 77, and Fortran 90 implementations of MINRES with new stopping conditions $\|Ar_k\| \leq \text{tol}\|A\|\|r_k\|$ and $\|Ax_k\| \leq \text{tol}\|A\|\|x_k\|$, and a MATLAB 7.6 implementation of MINRES-QLP are available from SOL [48].

Following the philosophy of reproducible computational research as advocated in [11, 9], for each figure and example in this paper we mention either the source or the specific MATLAB command. Our MATLAB scripts are available at SOL [48].

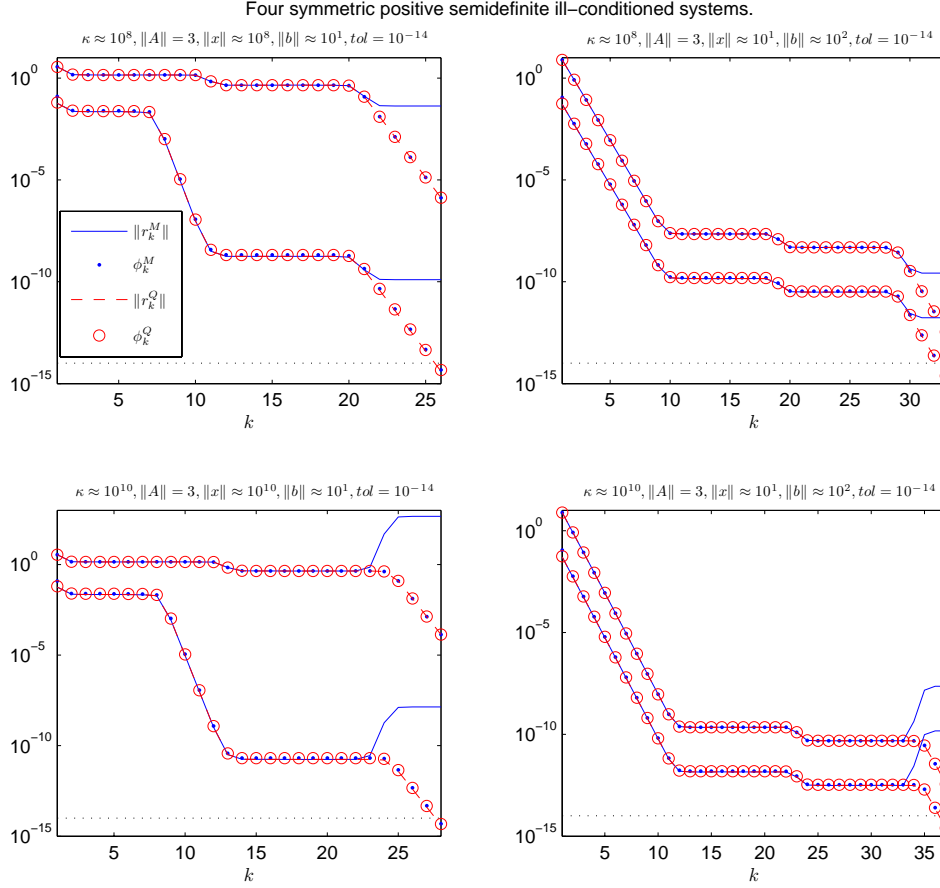


FIG. 8.2. Solving $Ax = b$ with semidefinite A similar to an example of Sleijpen et al. [47]. $A = Q \text{diag}([0_5, \eta, 2\eta, 2 : \frac{1}{789} : 3])Q$ of dimension $n = 797$, nullity 5, and norm $\|A\| = 3$, where $Q = I - (2/n)ww^T$ is a Householder matrix generated by $v = [0_5, 1, \dots, 1]^T$, $w = v/\|v\|$. These plots illustrate and compare the effect of rounding errors in MINRES and MINRES-QLP.

The upper part of each plot shows the computed and recurred residual norms, and the lower part shows the computed and recurred normwise relative backward errors (NRBE, defined in Table 6.1). MINRES and MINRES-QLP terminate when the recurred NRBE is less than the given $\text{tol} = 10^{-14}$.

Upper left: $\eta = 10^{-8}$ and thus $\kappa(A) \approx 10^8$. Also $b = e$ and therefore $\|x\| \gg \|b\|$. The graphs of MINRES's directly computed residual norms $\|r_k^M\|$ and recurrently computed residual norms ϕ_k^M start to differ at the level of 10^{-1} starting at iteration 21, while the values $\phi_k^Q \approx \|r_k^Q\|$ from MINRES-QLP decrease monotonically and stop near 10^{-6} at iteration 26.

Upper right: Again $\eta = 10^{-8}$ but $b = Ae$. Thus $\|x\| = \|e\| = O(\|b\|)$. The MINRES graphs of $\|r_k^M\|$ and ϕ_k^M start to differ when they reach a much smaller level of 10^{-10} at iteration 30. The MINRES-QLP ϕ_k^Q 's are excellent approximations of ϕ_k^Q , with both reaching 10^{-13} at iteration 33.

Lower left: $\eta = 10^{-10}$ and thus A is even more ill-conditioned than the matrix in the upper plots. Here $b = e$ and $\|x\|$ is again exploding. MINRES ends with $\|r_k^M\| \approx 10^2$, which means no convergence, while MINRES-QLP reaches a residual norm of 10^{-4} .

Lower right: $\eta = 10^{-10}$ and $b = Ae$. The final MINRES residual norm $\|r_k^M\| \approx 10^{-8}$, which is satisfactory but not as accurate as ϕ_k^M claims at 10^{-13} . MINRES-QLP again has $\phi_k^Q \approx \|r_k^Q\| \approx 10^{-13}$ at iteration 37.

This figure can be reproduced by `DptestSing7.m`.




FIG. 8.3. Solving $Ax = b$ with semidefinite A similar to an example of Sleijpen et al. [47]. $A = Q \text{diag}([0_5, \eta, 2\eta, 2 : \frac{1}{789} : 3])Q$ of dimension $n = 797$ with $\|A\| = 3$, where $Q = I - (2/n)ee^T$ is a Householder matrix generated by $e = [1, \dots, 1]^T$. (We are not plotting the NRBE quantities because $\|A\| \|r_k\| \approx 6$ throughout the iterations in this example.)

Upper left: $\eta = 10^{-2}$ and thus $\text{cond}(A) \approx 10^2$. Also $b = e$ and therefore $\|x\| \gg \|b\|$. The graphs of MINRES's directly computed $\|Ar_k^M\|$ and recurrently computed ψ_k^M , and also $\psi_k^Q \approx \|Ar_k^Q\|$ from MINRES-QLP, match very well throughout the iterations.

Upper right: Here, $\eta = 10^{-4}$ and A is more ill-conditioned than the last example (upper left). The final MINRES residual norm $\psi_k^M \approx \|Ar_k^M\|$ is slightly larger than the final MINRES-QLP residual norm $\psi_k^Q \approx \|Ar_k^Q\|$. The MINRES-QLP ψ_k^Q are excellent approximations of $\|Ar_k^Q\|$.

Lower left: $\eta = 10^{-6}$ and $\text{cond}(A) \approx 10^6$. MINRES's ψ_k^M and $\|Ar_k^M\|$ differ starting at iteration 21. Eventually, $\|Ar_k^M\| \approx 3$, which means no convergence. MINRES-QLP reaches a residual norm of $\psi_k^Q = \|Ar_k^Q\| = 10^{-2}$.

Lower right: $\eta = 10^{-8}$. MINRES performs even worse than in the last example (lower left). MINRES-QLP reaches a minimum $\|Ar_k^Q\| \approx 10^{-7}$ but $\text{tol} = 10^{-8}$ does not shut it down soon enough. The final $\psi_k^Q = \|Ar_k^Q\| = 10^{-2}$. The values of ψ_k^Q and $\|Ar_k^Q\|$ differ only at iterations 27–28.

This figure can be reproduced by `DptestLSSing5.m`.

Acknowledgements. We thank Jan Modersitzki, Gerard Sleijpen, Henk Van der Vorst, and also Kaustuv for providing us with their MATLAB scripts, which have aided us in producing parts of Figures 6.1, 8.2, and 8.3. We also thank Michael Friedlander, Rasmus Larsen, and Lek-Heng Lim for their finest examples of work, discussion, and support. We are most grateful to our anonymous reviewers for their insightful suggestions for improving this manuscript. Last but not least, we dedicate this paper to the memory of our colleague and friend, Gene Golub.

Appendix A. Proof that T_ℓ is nonsingular iff $b \in \text{range}(A)$ (section 2.1).

If T_ℓ is nonsingular, we can solve $T_\ell y = \beta_1 e_1$ and then $AV_\ell y = V_\ell T_\ell y = b$. Conversely, if $b \in \text{range}(A)$, then $\text{range}(V_\ell) \subseteq \text{range}(A)$ and $\text{null}(A) \cap \text{range}(V_\ell) = \{0\}$. We also know that $\text{rank}(V_\ell) = \ell$ and $\text{rank}(T_\ell) = \text{rank}(V_\ell T_\ell) = \text{rank}(AV_\ell) = \text{rank}(V_\ell) - \dim[\text{null}(A) \cap \text{range}(V_\ell)]$; see [4, Fact 2.10.4 ii]. Thus $\text{rank}(T_\ell) = \ell$ and so T_ℓ is nonsingular. \square

REFERENCES

- [1] M. Arioli and S. Gratton. Least-squares problems, normal equations, and stopping criteria for the conjugate gradient method. Technical Report RAL-TR-2008-008, Rutherford Appleton Laboratory, Oxfordshire, UK, 2008.
- [2] Z.-Z. Bai and G.-Q. Li. Restrictively preconditioned conjugate gradient methods for systems of linear equations. *IMA J. Numer. Anal.*, 23(4):561–580, 2003.
- [3] M. Benzi. Preconditioning techniques for large linear systems: a survey. *J. Comput. Phys.*, 182(2):418–477, 2002.
- [4] D. S. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas*. Princeton University Press, Princeton, NJ, 2nd edition, 2009.
- [5] E. Y. Bobrovnikova and S. A. Vavasis. Accurate solution of weighted least squares by iterative methods. *SIAM J. Matrix Anal. Appl.*, 22(4):1153–1174, 2001.
- [6] A. M. Bradley. *Algorithms for the Equilibration of Matrices and Their Application to Limited-Memory Quasi-Newton Methods*. PhD thesis, ICME, Stanford University, 2010.
- [7] X.-W. Chang, C. C. Paige, and D. Tittley-Péloquin. Stopping criteria for the iterative solution of linear least squares problems. *SIAM J. Matrix Anal. Appl.*, 31(2):831–852, 2009.
- [8] T.-Y. Chen and J. W. Demmel. Balancing sparse matrices for computing eigenvalues. *Linear Algebra Appl.*, 309(Issues 1-3):261–287, 2000.
- [9] S.-C. Choi, D. L. Donoho, A. G. Flesia, X. Huo, O. Levi, and D. Shi. About Beamlab—a toolbox for new multiscale methodologies. <http://www-stat.stanford.edu/~beamlab/>, 2002.
- [10] S.-C. T. Choi. *Iterative Methods for Singular Linear Equations and Least-Squares Problems*. PhD thesis, ICME, Stanford University, 2006.
- [11] J. Claerbout. Hypertext documents about reproducible research. <http://sepwww.stanford.edu/doku.php?id=sep:research:reproducible>.
- [12] F. A. Dul. MINRES and MINERR are better than SYMMLQ in eigenpair computations. *SIAM J. Sci. Comput.*, 19(6):1767–1782, 1998.
- [13] L. Eldén. Algorithms for the regularization of ill-conditioned least squares problems. *Nordisk Tidskr. Informationsbehandling (BIT)*, 17(2):134–145, 1977.
- [14] B. Fischer, A. Ramage, D. J. Silvester, and A. J. Wathen. Minimum residual methods for augmented systems. *BIT*, 38(3):527–543, 1998.
- [15] R. W. Freund and N. M. Nachtigal. A new Krylov-subspace method for symmetric indefinite linear systems. In W. F. Ames, editor, *Proceedings of the 14th IMACS World Congress on Computational and Applied Mathematics*, pages 1253–1256. IMACS, 1994.
- [16] K. A. Gallivan, S. Thirumalai, P. Van Dooren, and V. Vermaut. High performance algorithms for Toeplitz and block Toeplitz matrices. In *Proc. Fourth Conference of the International Linear Algebra Society (Rotterdam, 1994)*, volume 241/243, pages 343–388, 1996.
- [17] P. E. Gill, W. Murray, D. B. Ponceleón, and M. A. Saunders. Preconditioners for indefinite systems arising in optimization. *SIAM J. Matrix Anal. Appl.*, 13(1):292–311, 1992.
- [18] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, 3rd edition, 1996.
- [19] J. Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton University Bulletin*, XIII(4):49–52, 1902.
- [20] W. W. Hager. Iterative methods for nearly singular linear systems. *SIAM J. Sci. Comput.*, 22(2):747–766, 2000.
- [21] M. Hanke and J. G. Nagy. Restoration of atmospherically blurred images by symmetric indefinite conjugate gradient techniques. *Inverse Problems*, 12(2):157–173, 1996.
- [22] P. C. Hansen. Truncated singular value decomposition solutions to discrete ill-posed problems with ill-determined numerical rank. *SIAM J. Sci. Statist. Comput.*, 11(3):503–518, 1990.
- [23] P. C. Hansen and D. P. O’Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM J. Sci. Comput.*, 14(6):1487–1503, 1993.
- [24] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Research Nat. Bur. Standards*, 49:409–436, 1952.

- [25] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, PA, 2nd edition, 2002.
- [26] P. Jiránek and D. Tittley-Péloquin. Estimating the backward error in LSQR. *SIAM J. Matrix Anal. Appl.*, 31(4):2055–2074, 2010.
- [27] E. F. Kaasschieter. Preconditioned conjugate gradients for solving singular systems. *J. Comput. Appl. Math.*, 24(1-2):265–275, 1988.
- [28] Kaustuv. *IPSOL: An Interior Point Solver for Nonconvex Optimization Problems*. PhD thesis, SCCM Program, Stanford University, 2008.
- [29] M. Kilmer and G. W. Stewart. Iterative regularization and MINRES. *SIAM J. Matrix Anal. Appl.*, 21(2):613–628, 1999.
- [30] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Research Nat. Bur. Standards*, 45:255–282, 1950.
- [31] R. M. Larsen. *Efficient Algorithms for Helioseismic Inversion*. PhD thesis, Dept of Computer Science, University of Aarhus, 1998.
- [32] J. G. Lewis. *Algorithms for Sparse Matrix Eigenvalue Problems*. PhD thesis, Dept of Computer Science, Stanford University, 1976.
- [33] O. E. Livne and G. H. Golub. Scaling by binormalization. *Numer. Algorithms*, 35(1):97–120, 2004.
- [34] M. F. Murphy, G. H. Golub, and A. J. Wathen. A note on preconditioning for indefinite linear systems. *SIAM J. Sci. Comput.*, 21(6):1969–1972, 2000.
- [35] M. G. Neytcheva and P. S. Vassilevski. Preconditioning of indefinite and almost singular finite element elliptic equations. *SIAM J. Sci. Comput.*, 19(5):1471–1485, 1998.
- [36] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [37] Y. Notay. Solving positive (semi)definite linear systems by preconditioned iterative methods. In *Preconditioned Conjugate Gradient Methods (Nijmegen, 1989)*, volume 1457 of *Lecture Notes in Math.*, pages 105–125. Springer, Berlin, 1990.
- [38] C. C. Paige. Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix. *J. Inst. Math. Appl.*, 18(3):341–349, 1976.
- [39] C. C. Paige and M. A. Saunders. Solution of sparse indefinite systems of linear equations. *SIAM J. Numer. Anal.*, 12(4):617–629, 1975.
- [40] C. C. Paige and M. A. Saunders. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software*, 8(1):43–71, 1982.
- [41] C. C. Paige and M. A. Saunders. Algorithm 583; LSQR: Sparse linear equations and least-squares problems. *ACM Trans. Math. Software*, 8(2):195–209, 1982.
- [42] W.-Q. Ren and J.-X. Zhao. Iterative methods with preconditioners for indefinite systems. *J. Comput. Math.*, 17(1):89–96, 1999.
- [43] M. Rozložník and V. Simoncini. Krylov subspace methods for saddle point problems with indefinite preconditioning. *SIAM J. Matrix Anal. Appl.*, 24(2):368–391 (electronic), 2002.
- [44] Y. Saad and M. H. Schultz. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.*, 7(3):856–869, 1986.
- [45] M. A. Saunders. Solution of sparse rectangular systems using LSQR and Craig. *BIT*, 35(4):588–604, 1995.
- [46] M. A. Saunders. Computing projections with LSQR. *BIT*, 37(1):96–104, 1997.
- [47] G. L. G. Sleijpen, H. A. Van der Vorst, and J. Modersitzki. Differences in the effects of rounding errors in Krylov solvers for symmetric indefinite linear systems. *SIAM J. Matrix Anal. Appl.*, 22(3):726–751, 2000.
- [48] Systems Optimization Laboratory (SOL), Stanford University, downloadable software. <http://www.stanford.edu/group/SOL/software.html>.
- [49] G. W. Stewart. On the continuity of the generalized inverse. *SIAM J. Appl. Math.*, 17:33–45, 1969.
- [50] G. W. Stewart. Research, development and LINPACK. In J. R. Rice, editor, *Mathematical Software III*, pages 1–14. Academic Press, New York, 1977.
- [51] G. W. Stewart. The QLP approximation to the singular value decomposition. *SIAM J. Sci. Comput.*, 20(4):1336–1348, 1999.
- [52] K.-C. Toh, K.-K. Phoon, and S.-H. Chan. Block preconditioners for symmetric indefinite linear systems. *Internat. J. Numer. Methods Engrg.*, 60(8):1361–1381, 2004.
- [53] L. N. Trefethen and D. Bau, III. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, 1997.
- [54] University of Florida sparse matrix collection. <http://www.cise.ufl.edu/research/sparse/matrices/>.
- [55] Y. Zhang. Solving large-scale linear programs by interior-point methods under the MATLAB environment. *Optim. Methods Softw.*, 10(1):1–31, 1998.