PROBABILISTIC BOUNDS FOR THE MATRIX CONDITION NUMBER WITH EXTENDED LANCZOS BIDIAGONALIZATION

SARAH W. GAAF* AND MICHIEL E. HOCHSTENBACH*

Abstract. Reliable estimates for the condition number of a large (sparse) matrix A are important in many applications. To get an upper bound for the condition number $\kappa(A)$, a lower bound for the smallest singular value is needed. Krylov subspaces are usually unsuitable for finding a good approximation to the smallest singular value. Therefore, we study extended Krylov subspaces which turn out to be ideal for the simultaneous approximation of both the smallest and largest singular value of a matrix. First, we develop a new extended Lanczos bidiagonalization method. With this method we obtain a guaranteed lower bound for the condition number. Moreover, the method also yields a probabilistic upper bound for $\kappa(A)$. This probabilistic upper bound holds with a user-chosen probability.

Key words. Extended Lanczos bidiagonalization, extended Krylov method, matrix condition number, lower bound, probabilistic upper bound.

AMS subject classifications. 65F15, 65F35, 65F50, 65F10.

1. Introduction.

Let $A \in \mathbb{R}^{n\times n}$ be a large, nonsingular matrix. Let $A = X\Sigma Y^T$ be the singular value decomposition of A, where X and Y are $n \times n$ matrices with orthonormal columns containing the left and right singular vectors of A, respectively. Furthermore, Σ is an $n \times n$ diagonal matrix with positive real entries containing the singular values of A that are numbered in decreasing order,

$$\sigma_1 \geq \cdots \geq \sigma_n > 0.$$

We are interested in the important problem of approximating the condition number of A

$$\kappa(A) = \|A\|\,\|A^{-1}\| = \frac{\sigma_1}{\sigma_n},$$

where $\|\cdot\|$ stands for the 2-norm. The (Golub–Kahan–)Lanczos bidiagonalization method [2] gives an approximation, a lower bound, for the maximum singular value σ_1 of A. In addition, an upper bound for the minimum singular value is obtained, but this is usually a rather poor bound (see, for example, the experiments in Section 5). To approximate the condition number, good approximations to σ_n are needed.

This paper has three contributions. First, we develop a new extended Lanczos bidiagonalization method. The method generates a basis for the extended Krylov subspace:

$$\mathcal{K}^{k+1,k+1}(A^TA, \mathbf{v}) = \mathrm{span}\{(A^TA)^{-k}\mathbf{v}, \ldots, (A^TA)^{-1}\mathbf{v}, \mathbf{v}, A^TA\mathbf{v}, \ldots, (A^TA)^k\mathbf{v}\}.$$

Extended Krylov subspace methods have been studied recently by various authors [1, 6, 7, 8, 12]. The second contribution of this paper is that we obtain a guaranteed lower bound for σ_1 and a guaranteed upper bound for σ_n (that is much smaller than the one obtained with the standard Lanczos bidiagonalization), which leads to

*Version January 10, 2014. Department of Mathematics and Computer Science, TU Eindhoven, PO Box 513, 5600 MB, The Netherlands, `s.w.gaaf@tue.nl`, `www.win.tue.nl/~hochsten`. These authors were supported by an NWO Vidi research grant.

a guaranteed lower bound of good quality for the condition number of A. Third, we obtain a probabilistic upper bound for the condition number. Recently, probabilistic techniques have become popular, see for instance [3, 4, 13]. An important feature of the Lanczos bidiagonalization procedure is that the starting vector can be (and often is) chosen randomly. The probability that this vector has a very small component in the direction of the singular vector we are interested in is small. Another characteristic of the procedure is that during the bidiagonalization process polynomials implicitly arise. These two properties are exploited in [4] to obtain a probabilistic upper bound for σ_1. The bound holds with a user-chosen probability. In the extended Lanczos bidiagonalization we will expand the techniques from [4] to obtain both a probabilistic lower bound for σ_n and a probabilistic upper bound for σ_1, leading to a probabilistic upper bound for $\kappa(A)$.

2. Extended Lanczos bidiagonalization. The algorithm we develop starts with a random vector $\mathbf{v}_0$ with norm one. We express $\mathbf{v}_0$ as linear combination of the right singular vectors $\mathbf{y}_i$

$$\mathbf{v}_0 = \sum_{i=1}^{n} \gamma_i \, \mathbf{y}_i. \tag{2.1}$$

Notice that both the $\mathbf{y}_i$ and γ_i are not known. The extended Lanczos bidiagonalization method repeatedly applies the matrices A, A^T, A^{-T}, and A^{-1}. In every step a generated vector is orthogonalized with respect to the previous constructed vectors, and then normalized. This procedure can be visualized as a string of operations working on vectors:

$$\mathbf{v}_0 \xrightarrow{A} \mathbf{u}_0 \xrightarrow{A^T} \mathbf{v}_1 \xrightarrow{A^{-T}} \mathbf{u}_{-1} \xrightarrow{A^{-1}} \mathbf{v}_{-1} \xrightarrow{A} \mathbf{u}_1 \xrightarrow{A^T} \ldots$$

Note that in this visualization the orthogonalization and orthonormalization of the vectors are not shown. In this scheme, applying the operation A^{-T} after A^T (and A after A^{-1}) may seem contradictory, but since the vectors are orthogonalized in between this truly yields new vectors. Another way to represent this procedure is the table below:

Step	Action	Generated	Action	Generated	Action	Generated	Action	Generated
0	$A\mathbf{v}_0$	$\mathbf{u}_0$	$A^T\mathbf{u}_0$	$\mathbf{v}_1$	$A^{-T}\mathbf{v}_1$	$\mathbf{u}_{-1}$	$A^{-1}\mathbf{u}_{-1}$	$\mathbf{v}_{-1}$
1	$A\mathbf{v}_{-1}$	$\mathbf{u}_1$	$A^T\mathbf{u}_1$	$\mathbf{v}_2$	$A^{-T}\mathbf{v}_2$	$\mathbf{u}_{-2}$	$A^{-1}\mathbf{u}_{-2}$	$\mathbf{v}_{-2}$
$\vdots$								
$k-1$	$A\mathbf{v}_{-k+1}$	$\mathbf{u}_{k-1}$	$A^T\mathbf{u}_{k-1}$	$\mathbf{v}_k$	$A^{-T}\mathbf{v}_k$	$\mathbf{u}_{-k}$	$A^{-1}\mathbf{u}_{-k}$	$\mathbf{v}_{-k}$

During the procedure, the generated vectors $\mathbf{v}_j$ are normalized after being orthogonalized with respect to all previous generated $\mathbf{v}_i$, i.e., for $k \geq 1$

$$\mathbf{v}_k \perp \{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_{k-1}, \mathbf{v}_{-k+1}\}, \qquad \mathbf{v}_{-k} \perp \{\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_{-k+1}, \mathbf{v}_k\}.$$

Similarly, all generated vectors $\mathbf{u}_j$ have unit norm and

$$\mathbf{u}_{k-1} \perp \{\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{k-2}, \mathbf{u}_{-k+1}\}, \qquad \mathbf{u}_{-k} \perp \{\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{-k+1}, \mathbf{u}_{k-1}\}.$$

Define the matrices

$$\begin{aligned} V_{2k} &= [\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_{-1}, \ldots, \mathbf{v}_k], \\ V_{2k+1} &= [V_{2k}, \mathbf{v}_{-k}], \\ U_{2k-1} &= [\mathbf{u}_0, \mathbf{u}_{-1}, \mathbf{u}_1, \ldots, \mathbf{u}_{k-1}], \\ U_{2k} &= [U_{2k-1}, \mathbf{u}_{-k}]. \end{aligned}$$

The columns of these matrices are orthonormal and they span the corresponding subspaces $\mathcal{V}_{2k}$, $\mathcal{V}_{2k+1}$, $\mathcal{U}_{2k-1}$, and $\mathcal{U}_{2k}$, respectively. After k steps the algorithm gives rise to the following expressions:

$$\begin{aligned} AV_{2k+1} &= U_{2k+1}T_{2k+1,2k+1}, \\ A^T U_{2k-1} &= V_{2k}(T_{2k-1,2k})^T &&= V_{2k-1}(T_{2k-1,2k-1})^T + \beta_{k-1}\,\mathbf{v}_k\,\mathbf{e}_{2k-1}^T, \\ A^{-T}V_{2k} &= U_{2k}H_{2k,2k}^T, \\ A^{-1}U_{2k} &= V_{2k+1}H_{2k+1,2k} &&= V_{2k}H_{2k,2k} + \delta_k\,\mathbf{v}_{-k}\,\mathbf{e}_{2k}^T. \end{aligned} \tag{2.2}$$

Here, $\mathbf{e}_{2k}$ is the $2k$th unit vector and the coefficients β_{k-1} and δ_k are entries of the matrices T and H and will be specified below. The matrices T and H turn out to be tridiagonal matrices with a special structure as we will show in the next proposition. The leading submatrices $T_{2k-1} \in \mathbb{R}^{(2k-1)\times(2k-1)}$ and $H_{2k} \in \mathbb{R}^{(2k)\times(2k)}$ are given by

$$T_{2k-1} = U_{2k-1}^T A V_{2k-1}, \qquad H_{2k} = V_{2k}^T A^{-1} U_{2k}.$$

PROPOSITION 2.1.

(a) The matrix T_{2k-1} is tridiagonal and of the form

$$\begin{bmatrix} \alpha_0 & \beta_0 & & & & & \\ & \alpha_1 & & & & & \\ & \beta_{-1} & \alpha_{-1} & \beta_1 & & & \\ & & & \alpha_2 & & & \\ & & & \beta_{-2} & \alpha_{-2} & \beta_2 & \\ & & & & & \alpha_3 & \\ & & & & & & \ddots \end{bmatrix},$$

where its entries satisfy

$$\begin{aligned} t_{2j,2j} &= \alpha_j &&= \|A^{-T}\mathbf{v}_j\|^{-1} = \|A^T\mathbf{u}_{-j}\|, \\ t_{2j+1,2j} &= \beta_{-j} &&= \mathbf{u}_j^T A\mathbf{v}_j, \\ t_{2j+1,2j+1} &= \alpha_{-j} &&= \mathbf{u}_j^T A\mathbf{v}_{-j}, \\ t_{2j+1,2j+2} &= \beta_j &&= \|A^T\mathbf{u}_j - (\mathbf{u}_j^T A\mathbf{v}_j)\mathbf{v}_j - (\mathbf{u}_j^T A\mathbf{v}_{-j})\mathbf{v}_{-j}\|. \end{aligned}$$

(b) The matrix H_{2k} is tridiagonal and of the form

$$\begin{bmatrix} \alpha_0^{-1} & \delta_0 & & & & & \\ & \alpha_1^{-1} & & & & & \\ & \delta_1 & \alpha_{-1}^{-1} & \delta_{-1} & & & \\ & & & \alpha_2^{-1} & & & \\ & & & \delta_2 & \alpha_{-2}^{-1} & \delta_{-2} & \\ & & & & & \alpha_3^{-1} & \\ & & & & & & \ddots \end{bmatrix},$$

where its entries satisfy

$$\begin{aligned} h_{2j,2j} &= \alpha_j &&= \|A^{-T}\mathbf{v}_j\|^{-1} = \|A^T\mathbf{u}_{-j}\|, \\ h_{2j+1,2j} &= \delta_j &&= \|A\mathbf{v}_{-j}\|^{-1}, \\ h_{2j+1,2j+1} &= \alpha_{-j} &&= \mathbf{u}_j^T A\mathbf{v}_{-j}, \\ h_{2j+1,2j+2} &= \delta_{-j} &&= \|A^T\mathbf{u}_j - (\mathbf{u}_j^T A\mathbf{v}_j)\mathbf{v}_j - (\mathbf{u}_j^T A\mathbf{v}_{-j})\mathbf{v}_{-j}\|. \end{aligned}$$

Let $\theta_1^{(2k-1)} \geq \cdots \geq \theta_{2k-1}^{(2k-1)}$ be the singular values of T_{2k-1}, and $\xi_1^{(2k)} \geq \cdots \geq \xi_{2k}^{(2k)}$ be the singular values of H_{2k}. These values are approximations of the singular values of A and A^{-1}, respectively. We will avoid the use of superscripts if this is clear from the context. Further, let $\mathbf{c}_j$ and $\mathbf{d}_j$ indicate the corresponding right singular vectors of T_{2k-1} and H_{2k}, respectively. We will now study the behavior of these values θ_j and ξ_j to obtain guaranteed bounds for the extreme singular values of A.

PROPOSITION 2.2.

(a) *For* $1 \leq j \leq 2k-1$ *the singular values of* T_{2k-1} *converge monotonically to the largest singular values of* A*:*

$$\theta_j^{(2k-1)} \leq \theta_j^{(2k)} \leq \sigma_j(A).$$

(b) *For* $1 \leq j \leq 2k-1$ *the inverse singular values of* H_{2k} *converge monotonically to the smallest singular values of* A*:*

$$\sigma_{n-j+1}(A) = \big(\sigma_j(A^{-1})\big)^{-1} \leq \big(\xi_j^{(2k)}\big)^{-1} \leq \big(\xi_j^{(2k-1)}\big)^{-1}.$$

Proof. The matrix T_{2k-1} can be seen as the matrix T_{2k} from which the $2k$th row and column have been deleted. The same can be said of the matrices H_{2k-1} and H_{2k}. Now we apply [5, Cor. 3.1.3]. □

Proposition 2.2 shows in particular that the largest singular value of the matrices T_{2k-1} converges monotonically to σ_1, and the inverse of the largest singular value of the matrices H_{2k} converges monotonically to σ_n. After the kth step of the procedure, we obtain the value $\theta_1^{(2k-1)}$, a guaranteed lower bound for σ_1, and the value $(\xi_1^{(2k)})^{-1}$, a guaranteed upper bound for σ_n.

PROPOSITION 2.3. *After the* k*th step of the extended Lanczos bidiagonalization we obtain a guaranteed lower bound for the condition number of* A*:*

$$\kappa_{\text{low}}(A) = \theta_1\xi_1 \leq \frac{\sigma_1}{\sigma_n} = \kappa(A). \tag{2.3}$$

The experiments in Section 5 show for different matrices that the guaranteed lower bound achieved by the extended Lanczos bidiagonalization may often be very good.

We can reformulate the expressions in (2.2) to see the similarities with the extended Lanczos method, see, e.g., [7], with starting vector $\mathbf{v}_0$ and matrix A^TA:

$$\begin{aligned} A^TAV_{2k-1} &= A^TU_{2k-1}T_{2k-1,2k-1} \\ &= V_{2k-1}(T_{2k-1,2k-1})^T T_{2k-1,2k-1} + \beta_{k-1}\,\mathbf{v}_k\,\mathbf{e}_{2k-1}^T T_{2k-1,2k-1} \\ (A^TA)^{-1}V_{2k} &= A^{-1}U_{2k}H_{2k,2k}^T \\ &= V_{2k}H_{2k,2k}H_{2k,2k}^T + \delta_k\,\mathbf{v}_{-k}\,\mathbf{e}_{2k}^T H_{2k,2k}^T. \end{aligned} \tag{2.4}$$

In this extended Lanczos process the symmetric matrices $(T_{2k-1,2k-1})^T T_{2k-1,2k-1}$ and $H_{2k,2k}H_{2k,2k}^T$ are generated. They are both the formation of two tridiagonal matrices with a special structure, namely the matrices obtained from the extended Lanczos bidiagonalization, and they appear to be pentadiagonal. Furthermore, this way of representing the procedure will be convenient in the next section where we will introduce Laurent polynomials.

3. Polynomials arising in extended Lanczos bidiagonalization. In every step of the extended Lanczos bidiagonalization procedure four different vectors are generated (assuming there is no breakdown). For brevity we will focus on the vectors $\mathbf{v}_k$ and $\mathbf{v}_{-k}$, but analogous statements can be made about the vectors $\mathbf{u}_k$ and $\mathbf{u}_{-k}$. Since these vectors lie in an extended Krylov subspace, they can be expressed using polynomials:

$$\begin{aligned}\mathbf{v}_k &= p_k(A^TA)\mathbf{v}_0 &&\in \mathcal{K}^{k,k+1}(A^TA,\mathbf{v}_0)\\ \mathbf{v}_{-k} &= p_{-k}(A^TA)\mathbf{v}_0 &&\in \mathcal{K}^{k+1,k+1}(A^TA,\mathbf{v}_0).\end{aligned}$$

The polynomials p_k and p_{-k} are Laurent polynomials of the form

$$p_k(t) = \sum_{j=-k+1}^{k} a_j^{(k)}\, t^j, \qquad p_{-k}(t) = \sum_{j=-k}^{k} a_j^{(-k)}\, t^j.$$

Recall that $\theta_1^{(2k-1)} \geq \cdots \geq \theta_{2k-1}^{(2k-1)}$ are the singular values of T_{2k-1}, the singular values of H_{2k} are $\xi_1^{(2k)} \geq \cdots \geq \xi_{2k}^{(2k)}$, and $\mathbf{c}_j$ and $\mathbf{d}_j$ indicate the corresponding right singular vectors of T_{2k-1} and H_{2k}, respectively.

PROPOSITION 3.1. *The zeros of the polynomial p_k are exactly $\theta_1^2,\ldots,\theta_{2k-1}^2$, and the zeros of the polynomial p_{-k} are exactly $\xi_1^{-2},\ldots,\xi_{2k}^{-2}$.*

Proof. Let $j \in \{1,\ldots,2k-1\}$. Using (2.4) it can be easily seen that the Galerkin condition holds for the pair $(\theta_j^2, V_{2k-1}\mathbf{c}_j)$:

$$A^TAV_{2k-1}\mathbf{c}_j - \theta_j^2 V_{2k-1}\mathbf{c}_j \perp \mathcal{V}_{2k-1}.$$

Further, since $V_{2k-1}\mathbf{c}_j \in \mathcal{V}_{2k-1}$ it follows that

$$(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j \in \operatorname{span}\{(A^TA)^{-k+1}\mathbf{v}_0,\ldots,(A^TA)^k\mathbf{v}_0\}.$$

For all $j = 1,\ldots,2k-1$ we have that $(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j \in \mathcal{V}_{2k}$ but is orthogonal to $\mathcal{V}_{2k-1}$. This means that for all $j = 1,\ldots,2k-1$ the vector $(A^TA - \theta_j^2 I)V_{2k-1}\mathbf{c}_j$ is a nonzero multiple of $\mathbf{v}_k = p_k(A^TA)\mathbf{v}_0$. Hence $p_k(t)$ should contain all factors $t - \theta_j^2$. We know that $p_k(t) = \sum_{j=-k+1}^{k} a_j^{(k)} t^j$, and thus p_k has to be a nonzero multiple of

$$\nu(t) = t^{-k+1}\cdot(t-\theta_1^2)\cdots(t-\xi_{2k-1}^2).$$

Similarly, let $i \in \{1,\ldots,2k\}$. Using (2.4) it can be easily seen that the Galerkin condition holds for the pair $(\xi_i^2, V_{2k}\mathbf{d}_i)$:

$$(A^TA)^{-1}V_{2k}\mathbf{d}_i - \xi_i^2 V_{2k}\mathbf{d}_i \perp \mathcal{V}_{2k}.$$

Further, since $V_{2k}\mathbf{d}_i \in \mathcal{V}_{2k}$ it follows that

$$((A^TA)^{-1} - \xi_i^2 I)V_{2k}\mathbf{d}_i \in \text{span}\{(A^TA)^{-k}\mathbf{v}_0, \dots, (A^TA)^k\mathbf{v}_0\}.$$

For all $i = 1, \dots, 2k$ we have that $((A^TA)^{-1} - \xi_i^2 I)V_{2k}\mathbf{d}_i \in \mathcal{V}_{2k+1}$, but also orthogonal to $\mathcal{V}_{2k}$. This means that for all $i = 1, \dots, 2k-1$ the vector $((A^TA)^{-1} - \xi_i^2 I)V_{2k}\mathbf{d}_i$ is a nonzero multiple of $\mathbf{v}_{-k} = p_{-k}(A^TA)\mathbf{v}_0$. Hence $p_{-k}(t)$ should contain all factors $t^{-1} - \xi_i^2$. We know that $p_{-k}(t) = \sum_{i=-k}^{k} a_i^{(-k)} t^i$, and thus p_{-k} has to be a nonzero multiple of

$$\mu(t) = t^k \cdot (t^{-1} - \xi_1^2) \cdots (t^{-1} - \xi_{2k}^2).$$

Note that the proofs in [4, p. 467] and [11, p. 266–267] follow the same line of reasoning. □

We recall from Proposition 2.2 that for increasing k the largest singular value of T_{2k-1} converges monotonically to σ_1, and the inverse of the largest singular value of H_{2k} converges monotonically to σ_n. This implies that the largest zero of polynomial p_k increases monotonically to σ_1^2. Likewise, the smallest zero of polynomial p_{-k} decreases monotonically to σ_n^2. These polynomials are used in the next section to obtain probabilistic bounds for both the largest and smallest singular value of A.

4. Probabilistic bounds for the condition number. After step k, the extended Lanczos bidiagonalization implicitly provides Laurent polynomials p_k and p_{-k}. In the previous section we have seen what the zeros of p_k and p_{-k} are (Proposition 3.1). Moreover, the polynomials $|p_k|$ and $|p_{-k}|$ are strictly increasing to the right of their largest zero and also to the left of their smallest zero, for $t \to 0$. These properties will lead to the derivation of a probabilistic upper bound for $\kappa(A)$. Therefore, we first observe the two equalities

$$\begin{aligned} 1 = \|\mathbf{v}_k\|^2 &= \|p_k(A^TA)\mathbf{v}_0\|^2 = \|\sum_{i=1}^n p_k(A^TA)\gamma_i\mathbf{y}_i\|^2 = \sum_{i=1}^n \gamma_i^2 p_k(\sigma_i^2)^2, \\ 1 = \|\mathbf{v}_{-k}\|^2 &= \|p_{-k}(A^TA)\mathbf{v}_0\|^2 = \|\sum_{i=1}^n p_{-k}(A^TA)\gamma_i\mathbf{y}_i\|^2 = \sum_{i=1}^n \gamma_i^2 p_{-k}(\sigma_i^2)^2. \end{aligned}$$

Here we used, in view of (2.1), that $A^TA\mathbf{y}_i = \sigma_i^2\mathbf{y}_i$ and the fact that the right singular vectors $\mathbf{y}_i$ are orthonormal. Since the obtained sums only consist of nonnegative terms, we conclude that

$$|p_k(\sigma_1^2)| \le \frac{1}{|\gamma_1|}, \qquad \text{and} \qquad |p_{-k}(\sigma_n^2)| \le \frac{1}{|\gamma_n|}. \tag{4.1}$$

If γ_1 would be known, the first estimate in (4.1) would provide an upper bound for $\|A\|^2 = \sigma_1^2$, namely the largest zero of the function

$$f_1(t) = |p_k(t)| - \frac{1}{|\gamma_1|}.$$

Similarly, if γ_n would be known, the second estimate in (4.1) would provide a lower bound for $\|A^{-1}\|^{-2} = \sigma_n^2$, namely the smallest zero of the function

$$f_2(t) = |p_{-k}(t)| - \frac{1}{|\gamma_n|}.$$

However, both γ_1 and γ_n are unknown. Therefore we will compute a value δ that will be a lower bound for $|\gamma_1|$ and $|\gamma_n|$ with a user-chosen probability. Suppose that $|\gamma_1| < \delta$. Then the largest zero of $f_1^\delta(t) = |p_k(t)| - \frac{1}{\delta}$ is smaller than the largest zero of $f_1^{\gamma_1}(t) = |p_k(t)| - \frac{1}{|\gamma_1|}$ and thus may be less then σ_1^2. This means that δ may not give an upper bound for σ_1. We now compute the value δ such that the probability that $|\gamma_1| < \delta$ (or $|\gamma_n| < \delta$) is small, namely ε. Let S^{n-1} be the unit sphere in $\mathbb{R}^n$. We choose the starting vector $\mathbf{v}_0$ randomly from a uniform distribution on S^{n-1} (Matlab code: `v1=randn(n,1); v1=v1/norm(v1)`), see, e.g., [9, p. 1116], which (by an orthogonal transformation) implies that $(\gamma_1, \ldots, \gamma_n)$ is also random with respect to the uniform distribution on S^{n-1}.

Lemma 4.1. *Assume that the starting vector* $\mathbf{v}_0$ *has been chosen randomly with respect to the uniform distribution over the unit sphere* S^{n-1} *and let* $\delta \in [0, 1]$. *Then*

$$P(|\gamma_1| \le \delta) = 2\,B(\tfrac{n-1}{2}, \tfrac{1}{2})^{-1} \cdot \int_0^{\arcsin(\delta)} \cos^{n-2}(t)\,dt,$$

where B *denotes Euler's Beta function:* $B(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt$, *and* P *stands for probability.*

Proof. See [13, Lemma 3.1]. □

The user selects the probability $\varepsilon = P(|\gamma_1| \le \delta)$, i.e., the probability that the computed bound may not be an upper bound for the singular value σ_1. Given this user-chosen ε we have to determine the δ for which

$$\int_0^{\arcsin(\delta)} \cos^{n-2}(t)\,dt = \tfrac{1}{2}\,\varepsilon\,B(\tfrac{n-1}{2}, \tfrac{1}{2}). \tag{4.2}$$

The δ can be computed for instance by using Newton's method. With this δ we can compute two probabilistic bounds, namely the square root of the largest zero of the function f_1^δ and square root of the smallest zero of the function f_2^δ. Computing these values can be done with Newton's method or with bisection. We thus acquire a probabilistic upper bound for σ_1 and a probabilistic lower bound for σ_n:

$$\sigma_1 < \sigma_{\text{up}}^{\text{prob}} \qquad \text{and} \qquad \sigma_n > \sigma_{\text{low}}^{\text{prob}}.$$

Both inequalities are true with probability at least $1 - \varepsilon$. Since the coefficients γ_1 and γ_n are chosen independently, the probability that both inequalities hold is at least $1 - 2\varepsilon$. This proves the following theorem and corollary.

Theorem 4.2. *Assume that the starting vector* $\mathbf{v}_0$ *has been chosen randomly with respect to the uniform distribution over* S^{n-1}. *Let* $\varepsilon \in (0, 1)$ *and let* δ *be given by* (4.2). *Then* $\sigma_{\text{up}}^{\text{prob}}$, *the square root of the largest zero of the polynomial*

$$f_1^\delta(t) = |p_k(t)| - \frac{1}{\delta} \tag{4.3}$$

is an upper bound for σ_1 *with probability at least* $1-\varepsilon$. *Also,* $\sigma_{\text{low}}^{\text{prob}}$, *the square root of the smallest zero of the polynomial*

$$f_2^\delta(t) = |p_{-k}(t)| - \frac{1}{\delta} \tag{4.4}$$

is a lower bound for σ_n *with probability at least* $1-\varepsilon$.

Combining these two bounds leads to a probabilistic upper bound for the condition number of A.

COROLLARY 4.3. *The inequality*

$$\kappa(A) = \frac{\sigma_1}{\sigma_n} \leq \frac{\sigma_{\text{up}}^{\text{prob}}}{\sigma_{\text{low}}^{\text{prob}}} = \kappa_{\text{up}}(A). \tag{4.5}$$

holds with probability at least $1-2\varepsilon$.

5. Numerical experiments. We present the pseudocode for the extended Lanczos bidiagonalization method including the computation of a guaranteed lower bound and a probabilistic upper bound for the condition number. We test the method to quickly estimate the condition number for some large matrices. The matrices we choose are real and nonsymmetric (except for the second) and the number contained in their name indicates the size of the square matrix (`memplus` is an 17758×17758 matrix and `sherman5` is 3312×3312). Most of these matrices can be found in the Matrix Market [10]. The extended Lanczos bidiagonalization method constructs a basis for the Krylov subspace $\mathcal{K}^{k+1,k+1}(A^TA, \mathbf{v}_0)$ with dimension $2k+1$. Therefore, we compare the method to standard Lanczos bidiagonalization that builds a Krylov subspace $\mathcal{K}^{2k+1}(A^TA, \mathbf{v}_0)$ of dimension $2k+1$. Standard Lanczos bidiagonalization builds a bidiagonal matrix $B_{2k+1,2k+2}$ (see, for instance, [2]), and its condition number gives a lower bound for $\kappa(A)$. The starting vector $\mathbf{v}_0$ is randomly chosen from a uniform distribution on S^{n-1} as explained in Section 4. For these experiments we choose $\varepsilon = 0.01$ which corresponds to a reliability of at least 98% for the bounds for the condition number to be true (see Section 4). We use bisection to compute δ and to compute the largest and smallest zero of f_1^δ and f_2^δ, respectively (see (4.3) and (4.4)). Extended Lanczos bidiagonalization is computationally more expensive compared to standard Lanczos bidiagonalization, but since the bounds are so superior to those of the standard Lanczos bidiagonalization this extra cost is fully justified. Finally, we compute the ratios $\kappa(A)/\kappa_{\text{low}}(A)$ and $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$ (see (2.3) and (4.5)). In our experiments we used a fixed k. The next step would be to adaptively choose k given a user-selected ε and desired ratio $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$.

6. Discussion and conclusions. We have proposed a new extended Lanczos bidiagonalization method. This method leads to tridiagonal matrices with a special structure. Further, the method provides a guaranteed lower bound for $\kappa(A)$ of good quality. Also, it yields a tight probabilistic upper bound for $\kappa(A)$, already after a modest number of steps. Extended Lanczos bidiagonalization is more expensive than the standard Lanczos bidiagonalization, but it outperforms the standard method giving much tighter guaranteed lower bounds for the condition number and also probabilistic upper bounds. In our experiments we used a fixed k. Currently, we are developing a code that chooses k given a user-selected ε and desired ratio $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$.

Algorithm: Extended Lanczos bidiagonalization method with lower and probabilistic upper bounds

Input: Nonsingular matrix A, random starting vector $\mathbf{v}_0$, probability level ε, extended Krylov dimension $2k+1$ (we build a basis for $\mathcal{K}^{k+1,k+1}(A^TA, \mathbf{v}_0)$).

Output: A guaranteed lower bound $\kappa_{\text{low}}(A)$ and a probabilistic upper bound $\kappa_{\text{up}}(A)$ for the condition number $\kappa(A)$. The probability that $\kappa(A) \le \kappa_{\text{up}}(A)$ holds is at least $1-2\varepsilon$.

Determine δ from n and ε, see (4.2)
for $j = 0, \dots, k-1$
  $\mathbf{u} = A\mathbf{v}_{-j}$
  $\alpha_{-j} = \|\mathbf{u}\|$
  $\mathbf{u}_j = \mathbf{u} \,/\, \alpha_{-j}$
  $\mathbf{v} = A^T\mathbf{u}_j$
  if $j > 0$
    $\beta_{-j} = \mathbf{v}_j^T\mathbf{v}$
    $\mathbf{v} = \mathbf{v} - \beta_{-j}\mathbf{v}_j$
  end
  $\mathbf{v} = \mathbf{v} - \alpha_{-j}\mathbf{v}_{-j}$
  $\beta_j = \|\mathbf{v}\|$
  $\mathbf{v}_{j+1} = \mathbf{v} \,/\, \beta_j$
  $\mathbf{w} = A^{-T}\mathbf{v}_{j+1}$
  $\alpha_{j+1} = \|\mathbf{w}\|^{-1}$
  $\mathbf{u}_{-(j+1)} = \alpha_{j+1}\mathbf{w}$
  $\mathbf{z} = A^{-1}\mathbf{u}_{-(j+1)}$
  $\delta_{-j} = \mathbf{v}_{-j}^T\mathbf{z}$
  $\mathbf{z} = \mathbf{z} - \delta_{-j}\mathbf{v}_{-j} - \alpha_{j+1}^{-1}\mathbf{v}_{j+1}$
  $\delta_{j+1} = \|\mathbf{z}\|$
  $\mathbf{v}_{-(j+1)} = \mathbf{z} \,/\, \delta_{j+1}$
end
Determine the largest singular values θ_1 of T_{2k-1} and ξ_1 of H_{2k}
Compute guaranteed lower bound $\kappa_{\text{low}}(A) = \theta_1\xi_1$ for $\kappa(A)$ (2.3)
Determine probabilistic upper bound $\sigma_{\text{up}}^{\text{prob}}$ for σ_1 with probability $\ge 1-\varepsilon$ using f_1^δ (4.3)
Determine probabilistic lower bound $\sigma_{\text{low}}^{\text{prob}}$ for σ_n with probability $\ge 1-\varepsilon$ using f_2^δ (4.4)
Compute probabilistic upper bound $\kappa_{\text{up}}(A) = \sigma_{\text{up}}^{\text{prob}} \,/\, \sigma_{\text{low}}^{\text{prob}}$ for $\kappa(A)$ (4.5)

Table 5.1: The approximations of the condition number of different matrices using a standard Lanczos bidiagonalization method, and using extended Lanczos bidiagonalization. Both methods give a guaranteed lower bound for the condition number of the matrix. Extended Lanczos bidiagonalization also gives a probabilistic upper bound for the condition number that holds with probability at least 98% ($\varepsilon = 0.01$). We also give the ratios $\kappa(A)/\kappa_{\text{low}}(A)$ and $\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$.

		Lan. BD	**Ext. Lan. BD**		**Ratios**	
Matrix A	$\kappa(A)$	Guar. l.b.	$\kappa_{\text{low}}(A)$	$\kappa_{\text{up}}(A)$	$\kappa(A)/\kappa_{\text{low}}(A)$	$\kappa_{\text{up}}(A)/\kappa_{\text{low}}(A)$
`af23560`	$1.98 \cdot 10^4$	$1.80 \cdot 10^1$	$1.81 \cdot 10^4$	$2.05 \cdot 10^4$	1.10	1.13
`diag(1:10000)`	$1.00 \cdot 10^4$	$1.37 \cdot 10^1$	$9.72 \cdot 10^3$	$1.23 \cdot 10^4$	1.03	1.27
`dw2048`	$2.09 \cdot 10^3$	$1.31 \cdot 10^1$	$1.99 \cdot 10^3$	$2.52 \cdot 10^3$	1.05	1.26
`dw8192`	$3.81 \cdot 10^6$	$4.64 \cdot 10^1$	$3.78 \cdot 10^6$	$4.32 \cdot 10^6$	1.01	1.14
`grcar10000`	$3.63 \cdot 10^0$	$3.59 \cdot 10^0$	$3.57 \cdot 10^0$	$4.75 \cdot 10^0$	1.02	1.33
`memplus`	$1.29 \cdot 10^5$	$7.62 \cdot 10^1$	$1.28 \cdot 10^5$	$1.42 \cdot 10^5$	1.01	1.11
`olm5000`	$3.71 \cdot 10^7$	$1.49 \cdot 10^1$	$3.67 \cdot 10^7$	$4.55 \cdot 10^7$	1.01	1.24
`sherman5`	$1.88 \cdot 10^5$	$6.02 \cdot 10^1$	$1.88 \cdot 10^5$	$1.88 \cdot 10^5$	1.00	1.00
`tols2000`	$6.00 \cdot 10^6$	$1.50 \cdot 10^1$	$5.81 \cdot 10^6$	$6.97 \cdot 10^6$	1.03	1.20
`tols4000`	$2.36 \cdot 10^7$	$1.60 \cdot 10^1$	$2.31 \cdot 10^7$	$2.75 \cdot 10^7$	1.02	1.19
`utm5940`	$4.35 \cdot 10^8$	$1.60 \cdot 10^1$	$4.07 \cdot 10^8$	$5.04 \cdot 10^8$	1.07	1.24

REFERENCES

[1] V. L. Druskin and L. A. Knizhnerman, *Extended Krylov subspaces: approximation of the matrix square root and related functions*, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 755–771.

[2] G. H. Golub and C. F. Van Loan, *Matrix Computations*, The John Hopkins University Press, Baltimore, London, 3rd ed., 1996.

[3] N. Halko, P. G. Martinsson, and J. A. Tropp, *Finding structures with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.

[4] M. E. Hochstenbach, *Probabilistic upper bounds for the matrix two-norm*, J. Sci. Comput., 57 (2013), pp. 464–476.

[5] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.

[6] C. Jagels and L. Reichel, *The extended Krylov subspace method and orthogonal Laurent polynomials*, Linear Algebra Appl., 431 (2009), pp. 441–458.

[7] ———, *Recursion relations for the extended Krylov subspace method*, Linear Algebra Appl., 434 (2011), pp. 1716–1732.

[8] L. Knizhnerman and V. Simoncini, *A new investigation of the extended Krylov subspace method for matrix function evaluations*, Numer. Linear Algebra Appl., 17 (2010), pp. 615–638.

[9] J. Kuczyński and H. Woźniakowski, *Estimating the largest eigenvalue by the power and Lanczos algorithms with a random start*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 1094–1122.

[10] *The Matrix Market*. `http://math.nist.gov/MatrixMarket`, a repository for test matrices.

[11] B. N. Parlett, *The Symmetric Eigenvalue Problem*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.

[12] V. Simoncini, *A new iterative method for solving large-scale Lyapunov matrix equations*, SIAM J. Sci. Comput., 29 (2007), pp. 1268–1288.

[13] J. L. M. van Dorsselaer, M. E. Hochstenbach, and H. A. van der Vorst, *Computing probabilistic bounds for extreme eigenvalues of symmetric matrices with the Lanczos method*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 837–852.