

Computing Interior Eigenvalues of Large Matrices*

Ronald B. Morgan

Department of Mathematics

The University of Missouri

Columbia, Missouri 65211

Dedicated to Gene Golub, Richard Varga, and David Young

Submitted by Robert J. Plemmons

ABSTRACT

Computing eigenvalues from the interior of the spectrum of a large matrix is a difficult problem. The Rayleigh-Ritz procedure is a standard way of reducing it to a smaller problem, but it is not optimal for interior eigenvalues. Here a method is given that does a better job. In contrast with standard Rayleigh-Ritz, *a priori* bounds can be given for the accuracy of interior eigenvalue and eigenvector approximations. When applied to the Lanczos algorithm, this method yields better approximations at early stages. Applied to preconditioning methods, the convergence rate is improved.

1. INTRODUCTION

We wish to compute a few eigenvalues from the interior of the spectrum of a large matrix. The accompanying eigenvectors may also be desired. The eigenvalues of interest are far enough in the middle of the spectrum that it is impractical to develop approximations to all of the eigenvalues to either side.

This problem occurs in a number of different applications. Only certain eigenvalues in the middle of the range are wanted when studying tidal motion [2]. Interior eigenvalues are desired in adaptive polynomial preconditioning for indefinite linear systems [1]. In some weather forecasting models, only eigenvalues with small real parts are wanted, while there are other eigenvalues with large positive and with large negative real parts [16].

*This research was supported by the National Science Foundation under contract CCR-8801605.

Sometimes molecular chemists are only interested in particular vibrational levels. And only a certain range of vibrational frequencies of buildings will be affected by earthquakes.

We assume that factorization is impractical due to the size and sparsity structure of the matrix. This makes the problem quite difficult. We will present a new approach that generally gives better results than the standard method. The matrix can be either symmetric or nonsymmetric, but this paper discusses only the symmetric case.

The rest of this section discusses the standard method. Section 2 presents the new approach, including two variations, and gives *a priori* bounds on the effectiveness. Section 3 gives an interior version of the Lanczos algorithm and discusses how long it will take to develop good approximations. Roundoff error effects are analyzed. Section 4 has an interior version for preconditioning methods. These methods can give much better results than the Lanczos algorithm for interior eigenvalues.

The Rayleigh-Ritz procedure [13] is the standard way to extract approximate eigenvectors from a subspace. Let the eigenvalue problem be

$$Az = \lambda z,$$

where A is an n by n symmetric matrix. Suppose a j dimensional subspace of \mathbf{R}^n has been chosen. The Rayleigh-Ritz procedure reduces to a smaller problem. Let Q be an n by j orthonormal matrix whose columns span the subspace. Then the reduced problem

$$Q^T A Q g = \theta g \tag{1}$$

is solved. The approximate eigenvalues, called Ritz values, are the θ 's, and the approximate eigenvectors or Ritz vectors are $y = Qg$. For Q not orthonormal, the generalized Rayleigh-Ritz procedure can be used, and the reduced problem is

$$Q^T A Q g = \theta Q^T Q g. \tag{2}$$

But the Rayleigh-Ritz procedure is not optimal for interior eigenvalues. Rayleigh-Ritz is optimal in a global sense and is optimal for exterior eigenvalues [13, p. 215]. However, there are no guarantees for interior eigenvalues. Also, it is difficult to tell if a Ritz value is meaningful without computing the accompanying Ritz vector and testing its accuracy. Spurious Ritz values or ghost values [15] can occur.

EXAMPLE 1. Consider the 3 by 3 diagonal matrix A with entries $-1, 0$ and 1 . Let Q have columns $q_1 = [-0.1, 0.99, 0.1]^T$ and $q_2 = [1/\sqrt{2}, 0, 1/\sqrt{2}]^T$. Then the subspace is good if you want to approximate the eigenpair at 0 , because q_1 is close to the eigenvector e_2 . However, the Rayleigh-Ritz procedure gives Ritz values -0.14 and 0.14 with Ritz vectors $[-0.57, 0.70, -0.43]^T$ and $[0.43, 0.70, 0.57]^T$. While both Ritz values appear to be approximating the eigenvalue at 0 , actually both are ghost values. Their associated vectors are not accurate at all. Contributions less than and greater than zero cancel out in the Ritz value.

For interior eigenvalues, Rayleigh-Ritz is best when used with an inverted operator [15], say $(A - \sigma I)^{-1}$. The eigenvalues near σ are mapped to the exterior of the spectrum of the inverted matrix. However, we have assumed that the factorization needed to implement this is impractical. Another approach is needed.

2. A MODIFIED RAYLEIGH-RITZ PROCEDURE

It is possible to use the inverted operator implicitly. Let the columns of P span the given subspace. Then the inverted operator can be used with the subspace spanned by $Q = (A - \sigma I)P$. While Q does not represent the desired subspace, once the form of the approximate eigenvectors from Q have been determined, equivalent approximate eigenvectors from P can be calculated. These are more accurate, at least for eigenvalues near σ .

Here the method is derived. Start with the inverted problem

$$(A - \sigma I)^{-1}z = \frac{1}{\lambda - \sigma}z, \quad (3)$$

and apply Equation (2), the generalized Rayleigh-Ritz method:

$$Q^T(A - \sigma I)^{-1}Qg = \frac{1}{\theta - \sigma}Q^TQg. \quad (4)$$

Ideally we would like to solve Equation (4) with Q representing the desired subspace. Since that is not possible, let $Q = (A - \sigma I)P$, with P representing the desired subspace. Then Equation (4) becomes

$$P^T(A - \sigma I)Pg = \frac{1}{\theta - \sigma}P^T(A - \sigma I)^2Pg. \quad (5)$$

The reduced problem given by Equation (5) yields approximate eigenpairs $(1/(\theta - \sigma), Qg)$ for the inverted problem. These approximations are optimally chosen from the subspace spanned by Q , since the Rayleigh-Ritz procedure is done implicitly with respect to the inverted operator. For the matrix A , (θ, Qg) is an approximate eigenpair. But actually (ρ, Pg) is better, where ρ is the Rayleigh quotient of Pg with respect to A . Pg is equivalent to Qg with one step of inverse iteration applied. Q is not needed in the calculations, but it is remarkable that in deriving the method, moving to the poorer subspace corresponding to Q improves the resolution.

By starting with

$$(A - \sigma I)^{-2} z = \frac{1}{(\lambda - \sigma)^2} z$$

instead of Equation (3), another method is derived with reduced problem

$$P^T P g = \frac{1}{(\theta - \sigma)^2} P^T (A - \sigma I)^2 P g. \quad (6)$$

This can also be derived by simply applying the generalized Rayleigh-Ritz procedure to $(A - \sigma I)^2$.

So we have two methods, which we will call method 1 and method 2. They are effective at finding eigenvalues near σ . The problem of identifying ghost values [15] is eliminated.

EXAMPLE 1 (Continued). If these methods are applied to the problem in Example 1, much better results are obtained. For method 2, with $\sigma = 0$, the approximate eigenvector $[-0.10, 0.99, 0.10]^T$ is obtained. This is the best possible. And with $\sigma = 0.1$, $[-0.08, 0.99, 0.12]^T$ is about as good. For method 1, with $\sigma = 0$, we get $[-0.20, 0.98, 0.00]^T$. With $\sigma = 0.1$, the result is about the same. This is less accurate than method 2, but far better than standard Rayleigh-Ritz.

Next the methods are outlined.

METHOD 1.

1. Choose a subspace of dimension j . Compute P , an n by j matrix with columns spanning the subspace.
2. Calculate $H_1 = P^T (A - \sigma I) P$ and $H_2 = P^T (A - \sigma I)^2 P$.
3. Compute the desired number of the largest eigenpairs of $H_1 g_i = \alpha_i H_2 g_i$.

4. Compute approximate eigenpairs (ρ_i, y_i) of A , where $y_i = Pg_i$, and where if g_i is normalized, $\rho_i = \sigma + g_i^T H_1 g_i$. Also compute, as desired, the norms of the residual vectors, $r_i = Ay_i - \rho_i y_i$, to check the accuracy.

For numerical considerations, it is preferable if P is orthonormal. See [13, pp. 69, 222] for the significance of the residual norms.

METHOD 2. Same as method 1, except for step 3:

3. Compute the desired number of the smallest eigenpairs of $H_2 g_i = \alpha_i g_i$.

If H_1 is not calculated, then the ρ_i 's will have to be computed in another way.

Next the optimality properties are examined and some *a priori* bounds are given. First a lemma gives some standard results that will be useful. For notation, let $\varrho(s; A) = s^T A s / s^T s$, the Rayleigh quotient of s with respect to A . Let (λ_1, z_1) be the eigenpair of A with the smallest eigenvalue, and (θ_1, y_1) be the smallest Ritz pair from standard Rayleigh-Ritz.

LEMMA. For standard Rayleigh-Ritz with the matrix A and subspace \mathcal{S} ,

$$\lambda_1 \leq \theta_1 \leq \varrho(s; A) \quad \text{for all } s \in \mathcal{S}, \tag{7}$$

and

$$\sin^2 \angle(y_1, z_1) \leq \frac{\theta_1 - \lambda_1}{\lambda_2 - \lambda_1}. \tag{8}$$

Proof. Let N be an orthogonal matrix whose first column is a multiple of s and whose first j columns span \mathcal{S} . Then Equation (7) follows from the Cauchy interlacing theorem applied to $N^T A N$. For (8), let $\phi = \angle(y_1, z_1)$, and let $y_1 = z_1 \cos \phi + u \sin \phi$, where $u \perp z_1$ and all vectors are of unit length. Then

$$\theta_1 - \lambda_1 = y_1^T (A - \lambda_1 I) y_1 = \sin^2 \phi u^T (A - \lambda_1 I) u.$$

So

$$\sin^2 \phi = \frac{\theta_1 - \lambda_1}{u^T (A - \lambda_1 I) u} \leq \frac{\theta_1 - \lambda_1}{\lambda_2 - \lambda_1}. \quad \blacksquare$$

The lemma says that the vector in the subspace with the most negative Rayleigh quotient is extracted as a Ritz vector. The same is true for the vector with the most positive Rayleigh quotient. Bounds can also be given for Ritz approximations to interior eigenvalues in terms of the accuracy of all the approximations that are exterior to them [13, p. 227].

For results on the modified versions of Rayleigh-Ritz, instead of the usual eigenvalue ordering, we will number them according to their distance from σ . Let λ_1 be the eigenvalue nearest σ , λ_2 be the next nearest, and λ_n the furthest from σ . For method 1, let (θ_i, g_i) be the solution of Equation (5). For method 2, (θ_i, g_i) is the solution of (6). The θ_i 's are also ordered according to their distance from σ . Let $w_i \equiv Qg_i$ and $y_i \equiv Pg_i$. So y_i is the better approximate eigenvector. Let \mathcal{S} be the desired subspace, which is spanned by the columns of P . Denote by $(A - \sigma I)\mathcal{S}$ the subspace spanned by the columns of Q .

THEOREM 1. *In method 1, y_1 minimizes the magnitude of the quantity*

$$\frac{x^T(A - \sigma I)^2 x}{x^T(A - \sigma I)x} \quad (9)$$

over all $x \in \mathcal{S}$.

Proof. Assume first that $\sigma \neq \lambda_1$. Method 1 is equivalent to standard Rayleigh-Ritz with matrix $(A - \sigma I)^{-1}$ and with subspace $(A - \sigma I)\mathcal{S}$. From Equation (7) and the corresponding result for the largest Ritz value, it follows that $w_1 = (A - \sigma I)y_1$ has the Rayleigh quotient with respect to $(A - \sigma I)^{-1}$ with maximum magnitude. If

$$\frac{x^T(A - \sigma I)(A - \sigma I)^{-1}(A - \sigma I)x}{x^T(A - \sigma I)(A - \sigma I)x} = \frac{x^T(A - \sigma I)x}{x^T(A - \sigma I)^2 x}$$

is maximized over all $x \in \mathcal{S}$, then the quantity (9) is minimized.

As long as $z_1 \notin \mathcal{S}$, the matrix $P^T(A - \sigma I)^2 Pg$ stays away from being singular, and the eigenvalues of (5) are continuous with respect to σ . For $\sigma \neq \lambda_1$ the minimum of (9) is just $\theta_1 - \sigma$, the reciprocal of an eigenvalue of (5). Because of the continuity, the minimum at $\sigma = \lambda_1$ is from a solution of (5).

If $\sigma = \lambda_1$ and $z_1 \in \mathcal{S}$, then there is no minimum. ■

Note in the ideal case where $z_1 \in \mathcal{S}$, z_1 minimizes (9), so y_1 is chosen to be z_1 . This is also true for method 2, as the next theorem shows.

THEOREM 2. *In method 2, y_1 minimizes the quantity*

$$\frac{x^T(A - \sigma I)^2 x}{x^T x} \tag{10}$$

over all $x \in \mathcal{S}$. So y_1 minimizes the norm of this modified residual vector:

$$\frac{(A - \sigma I)x}{\|x\|}.$$

Proof. Method 2 is equivalent to standard Rayleigh-Ritz applied to $(A - \sigma I)^2$. By the lemma, y_1 gives the algebraically smallest Rayleigh quotient with respect to $(A - \sigma I)^2$. This Rayleigh quotient is nonnegative, so it is also the minimum in magnitude. ■

Some *a priori* bounds will now be given for the approximation to the eigenpair (λ_1, z_1) . Bounds can also be given for approximations to other eigenpairs, based on the accuracy of previous approximations.

THEOREM 3. *For convenience, assume that $\lambda_1 < \sigma$. For method 1,*

$$\frac{1}{\lambda_1 - \sigma} \leq \frac{1}{\theta_1 - \sigma} \leq \varrho(v; (A - \sigma I)^{-1}) \quad \text{for all } v \in (A - \sigma I)\mathcal{S}, \tag{11}$$

$$\sin^2 \angle(w_1, z_1) \leq \frac{1/(\theta_1 - \sigma) - 1/(\lambda_1 - \sigma)}{1/(\lambda_2 - \sigma) - 1/(\lambda_1 - \sigma)}, \tag{12}$$

$$\sin^2 \angle(y_1, z_1) \leq \left(\frac{\lambda_1 - \sigma}{\lambda_2 - \sigma} \right)^2 \tan^2 \angle(w_1, z_1), \tag{13}$$

$$|\rho_1 - \lambda_1| \leq \sin^2 \angle(y_1, z_1) |\lambda_n - \lambda_1|. \tag{14}$$

Proof. Equations (11) and (12) follow from the lemma, since $(1/(\theta_1 - \sigma), w_1)$ is a Ritz pair for the operator $(A - \sigma I)^{-1}$ with subspace $(A - \sigma I)\mathcal{S}$.

For Equation (13), using $y_1 = (A - \sigma I)^{-1}w_1$, one can derive that

$$\begin{aligned} \sin^2 \angle(y_1, z_1) &\leq \left(\frac{\lambda_1 - \sigma}{\lambda_2 - \sigma}\right)^2 \frac{\sin^2 \angle(w_1, z_1)}{\cos^2 \angle(w_1, z_1) + \left(\frac{\lambda_1 - \sigma}{\lambda_2 - \sigma}\right)^2 \sin^2 \angle(w_1, z_1)} \\ &\leq \left(\frac{\lambda_1 - \sigma}{\lambda_2 - \sigma}\right)^2 \tan^2 \angle(w_1, z_1). \end{aligned}$$

For (14), let $\phi = \angle(y_1, z_1)$, and let $y_1 = z_1 \cos \phi + u \sin \phi$, where $u \perp z_1$ and all vectors are of unit length. Then

$$\rho_1 - \lambda_1 = y_1^T (A - \lambda_1 I) y_1 = \sin^2 \phi u^T (A - \lambda_1 I) u,$$

and

$$|\rho_1 - \lambda_1| \leq \sin^2 \phi |\lambda_n - \lambda_1|. \quad \blacksquare$$

THEOREM 4. For method 2,

$$(\lambda_1 - \sigma)^2 \leq (\theta_1 - \sigma)^2 \leq \varrho(x; (A - \sigma I)^2) \quad \text{for all } x \in \mathcal{S}, \quad (15)$$

$$\sin^2 \angle(y_1, z_1) \leq \frac{(\theta_1 - \sigma)^2 - (\lambda_1 - \sigma)^2}{(\lambda_2 - \sigma)^2 - (\lambda_1 - \sigma)^2}, \quad (16)$$

$$|\rho_1 - \lambda_1| \leq \sin^2 \angle(y_1, z_1) |\lambda_1 - \lambda_n|. \quad (17)$$

Proof. Equations (15) and (16) follow from the lemma, since $((\theta_1 - \sigma)^2, y_1)$ is a Ritz pair for the operator $(A - \sigma I)^2$ with subspace \mathcal{S} . Equation (17) is the same as (14). \blacksquare

The bounds in Theorem 3 are not defined when σ equals λ_1 , and they may not be good in nearby cases because the subspace $(A - \sigma I)\mathcal{S}$ is very poor. It appears that method 1 may not work as well as method 2 when σ is near λ_1 . For method 2, the bounds on (ρ_1, y_1) decay when the two closest eigenvalues to σ are nearly the same distance away. This happens even if λ_1 and λ_2 are well separated, on opposite sides of σ . The eigenvalues λ_1 and λ_2 are mapped together by $(A - \sigma I)^2$, and the eigenvectors blur together. So

method 1 may perform better in this situation. These two cases will be discussed further with examples in the next section.

Because of these problem cases and because of the indirectness of the bounds on the approximate eigenvalues ρ_i , the bounds are not as appealing as those with standard Rayleigh-Ritz for exterior eigenvalues. However, it is nice that with our methods some bounds can be given for approximations to interior eigenvalues. If there is information about the subspace, more specific bounds can be given by choosing the vector x or v . This will be done in the next section with Krylov subspaces.

3. WITH LANCZOS

For large symmetric matrices, the symmetric Lanczos algorithm [13] is a standard method for computing eigenvalues. The Lanczos algorithm is the Rayleigh-Ritz procedure applied to a Krylov subspace [13]. The orthonormal matrix spanning this subspace is generated iteratively using an efficient three-term recurrence. The reduced matrix in Equation (1) is tridiagonal.

We now apply the modified versions of Rayleigh-Ritz to Lanczos. Let P be the orthonormal matrix generated by the Lanczos iteration. For method 1, Equation (5) becomes the j by j generalized eigenvalue problem

$$Tg = \frac{1}{\theta - \sigma} Vg, \quad (18)$$

where $T = P^T(A - \sigma I)Pg$ is tridiagonal and $V = P^T(A - \sigma I)^2Pg$ is pentadiagonal. T is the standard tridiagonal matrix from the Lanczos method except for shifting of the diagonal elements by σ . For method 2, Equation (6) becomes an eigenvalue problem with V .

The θ values have been used before in the case $\sigma = 0$ for applications such as orthogonal polynomials [5, 7]. However, we are interested in the more accurate ρ values and in solving eigenvalue problems.

The subspace that the approximate solutions are extracted from is the same as the subspace for the standard Lanczos algorithm. Once the subspace is good enough, standard Lanczos will usually give good approximations even to interior eigenvalues. So the new interior versions of Lanczos will not give faster convergence rates. But they will give better approximations at early stages. If eigenvectors are desired and storage is limited, then restarting is necessary. In that case, the best possible approximations are important. So the interior versions are useful if rough approximations to the eigenvalues are desired quickly or if restarting is used.

A key point is that in terms of n , the expense of these methods is the same as for standard Lanczos. V need not be formed; instead let $T = Q\bar{R}$ be an orthogonal QR-type factorization of T . Then use the Lanczos recurrence $(A - \sigma I)P = PT + \beta_j p_{j+1} e_j^T$, where p_{j+1} will be the next orthonormal column of P . Then

$$\begin{aligned}
 V &= P^T (A - \sigma I)^2 P \\
 &= (PT + \beta_j p_{j+1} e_j^T)^T (PT + \beta_j p_{j+1} e_j^T) \\
 &= T^T T + \beta_j^2 e_j e_j^T \\
 &= \bar{R}^T \bar{R} + \beta_j^2 e_j e_j^T \\
 &= R^T R,
 \end{aligned} \tag{19}$$

where R is upper triangular and differs from \bar{R} only in the jj element. For that element, $r_{jj}^2 = \bar{r}_{jj}^2 + \beta_j^2$. A similar formula was derived by Paige and Saunders [12] with their MINRES routine for solving indefinite systems of linear equations.

The approximate eigenvalues and their residual norms can be computed fairly cheaply. For g_i of unit length,

$$\rho_i = (Pg_i)^T APg_i = \sigma + g_i^T Tg_i. \tag{20}$$

A formula for the residual norm in both method 1 and method 2 is

$$\|r_i\|^2 = \|Tg_i + (\sigma - \rho_i)g_i\|^2 + \beta_{ji}^2, \tag{21}$$

where β_{ji} is β_j times the last component of g_i . For method 1 we also have

$$\|r_i\|^2 = (\theta_i - \rho_i)(\rho_i - \sigma), \tag{22}$$

and for method 2,

$$\|r_i\|^2 = (\theta_i - \sigma)^2 - (\rho_i - \sigma)^2. \tag{23}$$

So convergence can be monitored as the iteration progresses.

The Lanczos recurrence is independent of σ . Only the main diagonal of T varies. So σ can be easily changed, and there is no additional expense in terms of n . It is only necessary to do a QR factorization of the shifted T and solve a different reduced eigenvalue problem (18). It is possible to compute approximate eigenvalues over a range by letting σ vary.

EXAMPLE 2. For a test matrix, let A be diagonal and of dimension 500. The diagonal entries are 240 values equally spaced from 0.0 to 9.0, then an entry at 10.0, and then 259 values equally spaced from 11.0 to 20.0. Table 1 gives the residual norms every 25 steps for the best approximations to the eigenvalues at 10.0 and 11.0. The computations were done in double precision (about 16 decimal digits) on an IBM 4381-R14. The starting vector has all entries the same. Method 1 and method 2 use $\sigma = 10.1$. For computing the eigenvalue at 10.0, they give about the same results, and both are better than standard Lanczos. For example, at $j = 100$, the interior methods give residual norms that are more than an order of magnitude better than standard Lanczos. On the other hand, while the convergence of standard Lanczos is not as steady, it is at about the same rate as for the interior methods. It is interesting that at steps where Lanczos lags behind the most, there is a ghost Ritz value between 9 and 11. The interior methods also give better results for calculating the eigenvalue at 11.0. But here method 1 is better than method 2. For example at $j = 100$, Lanczos has residual norm 0.22, method 1 has $0.62E-1$, and method 2 has 0.12. The eigenvalue at 9.0 is not given in the table, but method 2 has even greater problems there. At $j = 100$, the residual norms are 0.18 for Lanczos, $0.67E-1$ for method 1, and

TABLE 1
 $\sigma = 10.1$

j	Residual norm					
	Ritz value near 10.0			Ritz value near 11.0		
	Lanczos	Method 1	Method 2	Lanczos	Method 1	Method 2
25	2.51	0.52	1.07	1.10	1.00	3.64
50	1.98	0.25	0.24	0.66	0.38	0.51
75	$0.33E-1$	$0.19E-1$	$0.19E-1$	0.26	0.10	0.17
100	$0.34E-1$	$0.17E-2$	$0.17E-2$	0.22	$0.62E-1$	0.12
125	$0.20E-3$	$0.12E-3$	$0.12E-3$	$0.99E-1$	$0.39E-1$	$0.62E-1$
150	$0.81E-4$	$0.96E-5$	$0.97E-5$	0.11	$0.26E-1$	$0.47E-1$
175	$0.11E-5$	$0.60E-6$	$0.60E-6$	$0.35E-1$	$0.15E-1$	$0.23E-1$
200	$0.19E-6$	$0.45E-7$	$0.45E-7$	$0.62E-1$	$0.90E-2$	$0.15E-1$

0.35 for method 2. Method 2 has trouble because of the nearness of $|9.0 - \sigma|$ to $|\lambda_i - \sigma|$ for several λ_i near 11.2.

EXAMPLE 3. Method 1 with the case where σ is equal or near to an eigenvalue is now examined. First the same test as in Example 1 is performed, except with $\sigma = 10.0$. Instead of just one approximation to the eigenpair at 10.0, there are many of them. At $j = 100$, the most accurate approximation has residual norm $0.67\text{E}-2$, compared to $0.17\text{E}-2$ in the previous example with $\sigma = 10.1$. Approximations to other eigenvalues do not appear. This is because the eigenvalue problem

$$(A - \sigma I)z = \frac{1}{\lambda - \sigma}(A - \sigma I)^2 z$$

that is approximated in Equation (5) has a continuous spectrum. The symmetry around 10.0 is not the reason. Results are similar with eigenvalues between 12 and 21 instead of between 11 and 20. Also, for $\sigma = 11.0$, extra approximations to the eigenvalue 11.0 show up, though not until the subspace is larger (around $j = 150$). In several other tests we let σ move away from λ_1 and fix j . Table 2 has residual norms for the eigenvalues at 10.0 and 11.0, for $j = 100$ and 150. For $j = 100$, by $\sigma = 10.00001$, the residual norm of the approximation to 10.0 improves to $0.17\text{E}-2$. Approximations to other eigenvalues appear soon after that, and at $\sigma = 10.01$ they attain almost the same accuracy as for $\sigma = 10.1$. The situation for σ near but not equal to the

TABLE 2
 σ CHANGING

σ	Residual norm			
	$j = 100$		$j = 150$	
	Ritz value 10	11	10	11
10.0	$0.67\text{E}-2$	—	$0.45\text{E}-4$	—
10.000000001	$0.67\text{E}-2$	—	$0.96\text{E}-5$	—
10.000001	$0.60\text{E}-2$	—	$0.96\text{E}-5$.25
10.00001	$0.17\text{E}-2$	—	$0.96\text{E}-5$	$0.35\text{E}-1$
10.0001	$0.17\text{E}-2$	0.51	$0.96\text{E}-5$	$0.27\text{E}-1$
10.01	$0.17\text{E}-2$	$0.65\text{E}-1$	$0.96\text{E}-5$	$0.26\text{E}-1$
10.1	$0.17\text{E}-2$	$0.62\text{E}-1$	$0.96\text{E}-5$	$0.26\text{E}-1$

eigenvalue should improve as the iteration progresses, because as the subspace \mathcal{S} improves, the subspace $(A - \sigma I)\mathcal{S}$ also improves. And then the ability of Rayleigh-Ritz with $(A - \sigma I)^{-1}$ to extract the correct vectors from that subspace improves. For $j = 150$, the interval of difficulty around 10.0 shrinks. The approximation to λ_1 is as accurate for $\sigma = 10.000000001$ as for $\sigma = 10.1$. The approximation to the eigenpair at 11.0 is about as accurate for $\sigma = 10.0001$ as for $\sigma = 10.1$.

We conclude that method 1 is useful even for σ near an eigenvalue, and that any ill effects will diminish as j increases. If there is some reason to be particularly concerned, it is possible to monitor the behavior of the method and, when needed, either change σ or switch to method 2 or standard Lanczos. All of these can be done with the same Lanczos iteration.

The problems of method 2 when eigenvalues are about the same distance away from σ but on opposite sides were shown in Example 1 with the second and third eigenvalues. Another case where this happens is for σ near 10.5. At $\sigma = 10.5$ and $j = 100$, the approximation to λ_1 has residual norm $0.52E-2$ instead of $0.17E-2$. This does not seem too bad, but at $\sigma = 10.51$, the result is $0.33E-1$. The problem case for method 2 does not seem as easy to detect as that for method 1. Generally, we favor method 1 over method 2.

A priori bounds similar to the Kaniel-Saad bounds [9, 13, 14] for standard Lanczos can be given. We will just give the bound for θ_1 from method 1. Bounds for w_1 , y_1 , and ρ_1 then follow from Equations (12) through (14), and bounds for method 2 can be similarly derived.

THEOREM 5. *For convenience, assume that $\lambda_1 < \sigma$ and $\lambda_2 > \sigma$. Let f be the starting vector. For method 1,*

$$0 \leq \lambda_1 - \theta_1 \leq (\sigma - \theta_1) \left(\frac{\sigma - \lambda_1}{\lambda_2 - \sigma} + 1 \right) \times \left[\tan \angle((A - \sigma I)f, z_1) \max_{i \neq 1} p(\lambda_i) \right]^2, \tag{24}$$

where p is any polynomial of degree $j - 1$ or less such that $p(\lambda_1) = 1$.

Proof. Let $h = (A - \sigma I)f$, $\phi = \angle(h, z_1)$, and $h = z_1 \cos \phi + u \sin \phi$, where u is orthogonal to z_1 and h , u , and z_1 are all unit vectors. \mathcal{S} is a

Krylov subspace. So all $v \in (A - \sigma I)\mathcal{S}$ can be written as $p(A)h$, for p a polynomial of degree $j - 1$ or less. Now

$$\begin{aligned} & \varrho \left(p(A)h; (A - \sigma I)^{-1} - \frac{1}{\lambda_1 - \sigma} I \right) \\ &= \frac{1}{\|p(A)h\|^2} \left[z_1^T \left((A - \sigma I)^{-1} - \frac{1}{\lambda_1 - \sigma} I \right) p^2(A) z_1 \cos^2 \phi \right. \\ &\quad \left. + u^T \left((A - \sigma I)^{-1} - \frac{1}{\lambda_1 - \sigma} I \right) p^2(A) u \sin^2 \phi \right] \\ &\leq \frac{1}{\|p(A)h\|^2} \left[0 + \left(\frac{1}{\lambda_2 - \sigma} - \frac{1}{\lambda_1 - \sigma} \right) \|p(A)u\|^2 \sin^2 \phi \right] \\ &\leq \left(\frac{1}{\lambda_2 - \sigma} - \frac{1}{\lambda_1 - \sigma} \right) \left[\tan \phi \max_{i \neq 1} |p(\lambda_i)| \right]^2, \end{aligned}$$

using that $\|p(A)u\| \leq \max_{i \neq 1} |p(\lambda_i)|$ and $\|p(A)h\| \geq p(\lambda_1) \cos \phi = \cos \phi$. For Equation (24), apply this to Equation (11) and use a little algebra. ■

For a good result, the polynomial p in Theorem 5 must be much larger at λ_1 than at the other eigenvalues. In the Kaniel-Saad bounds, the polynomial is chosen to be a shifted Chebyshev polynomial. The steep climb of the Chebyshev polynomial outside of the specified interval explains why convergence is rapid for a well-separated smallest eigenvalue. However, we must have a polynomial that is small on the two intervals containing λ_2 through λ_n and is large in the middle at λ_1 . To demonstrate the difficulty of this task, we will look at a special case. Let $\lambda_1 = 0$, and assume that the two intervals containing the other eigenvalues are symmetric about 0. For results about optimum polynomials over more complicated pairs of intervals, see [4] and its references.

THEOREM 6. *The polynomial p of degree $2k$ satisfying $p(0) = 1$ that minimizes*

$$\max_{x \in [\alpha, \beta] \text{ or } [-\beta, -\alpha]} p(x)$$

is

$$p(x) = \frac{T_k(cx^2 + d)}{T_k(d)},$$

where T_k is the degree k Chebyshev polynomial on $[-1, 1]$, and where $c = -2/(\beta^2 - \alpha^2)$ and $d = (\beta^2 + \alpha^2)/(\beta^2 - \alpha^2)$. The maximum of p on the two intervals is $1/T_k(d)$.

Proof. The proof is similar to that for standard Chebyshev results. The polynomial p attains its maximum or minimum $k + 1$ times in both of the two intervals (this follows from properties of T_k). If there is a polynomial q that satisfies the same conditions and has a lower maximum over the intervals, then $p - q$ has $2k + 1$ zeros. Since $p - q$ is of degree $2k$, it must be identically zero. So p is the minimal polynomial. ■

To compare this with a corresponding result for exterior eigenvalues, let the smallest eigenvalue of A be 0 and all of the other eigenvalues be in the interval $[\alpha, \beta]$. For $p(x)$ of degree k such that $p(0) = 1$, the minimum of $\max_{x \in [\alpha, \beta]} p(x)$ is

$$\left[T_k \left(\frac{\beta + \alpha}{\beta - \alpha} \right) \right]^{-1}.$$

Comparing $(\beta + \alpha)/(\beta - \alpha)$ with $(\beta^2 + \alpha^2)/(\beta^2 - \alpha^2)$ shows that convergence will be much slower for an interior eigenvalue. The convergence factors are roughly

$$1 - 2\sqrt{\frac{\alpha}{\beta - \alpha}}$$

for the exterior problem and

$$1 - \frac{\alpha}{\beta - \alpha}$$

for the interior problem. Approximately

$$2\sqrt{\frac{\beta - \alpha}{\alpha}}$$

times as many iterations are required for the interior problem as for the exterior problem.

So Krylov subspaces are more appropriate for exterior eigenvalues. The best way to find interior eigenvalues is to use the shift-and-invert Lanczos method [6]. The subspace is generated by an inverted operator and is much better. So not only are inverted operators better with the Rayleigh-Ritz extraction from the subspace, they also produce a better subspace. But as before, we are assuming that implementation of the inverted operator is impractical.

EXAMPLE 4. To demonstrate the difficulty of computing interior eigenvalues, let $A = \text{diag}\{0.1, 0.2, \dots, 29.9, 30.0\}$. Let $\sigma = 15.02$. To simulate the situation where the eigenvector is needed and the storage is limited, let the maximum size subspace be of dimension 100. The method is then restarted. Table 3 gives the results for 10 such runs to $j = 100$. The interior methods not only give better results, but they also show steady improvement, unlike standard Lanczos. Perhaps more significant is the fact that convergence is extremely slow for all methods. This problem is addressed in the next section.

It is well known that roundoff error causes the Lanczos recurrence to behave very differently in practice than it does in theory. Nevertheless, the standard Lanczos method produces accurate eigenvalue estimates. In the examples that have been given, the interior version of Lanczos also performed well in the presence of roundoff error. We conclude the section with an explanation of this. However, the explanation relies on theory developed

TABLE 3
DIFFICULT EXAMPLE

Run	Residual norm for Ritz value near 15.0		
	Lanczos	Method 1	Method 2
1	1.02	0.23	0.23
2	1.06	0.12	0.12
3	0.81	0.83E-1	0.86E-1
4	1.02	0.67E-1	0.70E-1
5	1.21	0.58E-1	0.60E-1
6	0.90	0.49E-1	0.51E-1
7	0.87	0.40E-1	0.43E-1
8	0.50	0.32E-1	0.36E-1
9	0.54	0.26E-1	0.30E-1
10	0.35	0.21E-1	0.24E-1

by Greenbaum [8], which in turn uses work by Paige [11]. The details will be left out here.

The basic idea is that while T differs from what it would be in exact arithmetic, it can be theoretically extended to a larger tridiagonal matrix whose eigenvalues all are near those of A . Now the pentadiagonal matrix V can also be extended, and then Equation (18) will have the same eigensolution as does the extended T . So while the solution of (18) is not the same as it would be in exact arithmetic, it is an intermediate step of an accurate answer.

LEMMA (Greenbaum). *Let m be the number of converged Ritz vectors at step j of a perturbed Lanczos recurrence. The tridiagonal matrix T is the same as one generated by an exact Lanczos recurrence applied to a matrix with eigenvalues close to those of A that has dimension $n + m$ or less (for convenience, assume $n + m$). The recurrence can be theoretically extended, and β_{n+m} will be zero. T_{n+m} then will have eigenvalues close to those of A .*

THEOREM 7. *Equation (18) generated by a perturbed Lanczos recurrence is the same as that generated by an exact Lanczos recurrence applied to a matrix whose eigenvalues are all close to those of A . The Lanczos recurrence can be theoretically extended so that the extended Equation (18) gives only ρ_i 's close to the eigenvalues of A .*

Proof. Extend the Lanczos recurrence as is done in the lemma. Then note that $V = T^2 + \beta_j^2 e_j^T e_j$. Thus $V_{n+m} = T_{n+m}^2$, since $\beta_{n+m} = 0$. So Equation (18) gives the same eigensolution as does T_{n+m} [the (ρ_i, g_i) 's are eigenpairs of T_{n+m}]. And by the lemma, the eigenvalues of T_{n+m} are all near those of A . ■

4. WITH PRECONDITIONING METHODS

Krylov subspaces are usually not very good for computing interior eigenvalues. Convergence can be very slow, as in Example 4. This section looks at methods that build a subspace to target specific eigenvalues.

Here we apply interior Rayleigh-Ritz method 1 to Davidson's method [3] and the GD method [10]. These methods use preconditioning to improve the convergence. Davidson's method has diagonal preconditioning. In GD, the preconditioning techniques developed for linear equations can be applied to eigenvalue problems.

Davidson's method and GD develop their subspace iteratively, using information at one step to find a new vector for the next step. The new trial vectors for the subspace are

$$v = (M - \alpha I)^{-1}(A - \rho I)y,$$

where (ρ, y) is the current Ritz pair approximating the desired eigenpair, where M is an approximation to A , and where α is usually chosen to be ρ or σ . The matrix $M - \alpha I$ is a preconditioner for $A - \rho I$. As ρ approaches λ_1 , $(M - \alpha I)^{-1}(A - \rho I)$ has an eigenvector approaching z_1 , and the eigenvalue is generally better separated in the spectrum of this preconditioned matrix than λ_1 is in the spectrum of A .

These methods are more expensive than the Lanczos algorithm, because there is no three-term recurrence. Also, the cost of an approximate factorization for the preconditioner may be significant. But this cost is less than for the complete factorization necessary for the shift-and-invert Lanczos method.

When the interior Rayleigh-Ritz method is used, there are better approximate eigenpairs developed at early stages. These are used to generate a better subspace. So the method can actually converge faster.

Full Gram-Schmidt orthogonalization is needed. Both P and $(A - \sigma I)P$ are saved to keep the number of matrix-vector products to one per iteration. There are three ways to perform the orthogonalization of P : with respect to the standard inner product, with respect to the $(A - \sigma I)^2$ inner product, and with respect to the indefinite $A - \sigma I$ inner product. The first way is probably best, because it is straightforward and it allows easy changing of σ . The expense is approximately $6jn$ multiplications per step for all length n operations except the matrix-vector product and the preconditioning. The regular Davidson's method and GD require $5jn$. Equation (5) becomes a j by j generalized eigenvalue problem. The second way yields a standard eigenvalue for Equation (5). The cost in terms of n is the same, but the orthogonalization is a little more complicated. The third way of orthogonalizing is less expensive at approximately $5jn$ multiplications per step, but stability is not guaranteed. It gives a generalized eigenvalue problem.

EXAMPLE 5. Here A is the tridiagonal matrix with $0.2, 0.4, \dots, 59.8, 60.0$ on the main diagonal and 1 's on the super- and subdiagonal. The distribution of eigenvalues makes computation of interior eigenvalues as difficult as in Example 4. Let $M = \text{diag}(A)$, $\sigma = 27.05$, and $\alpha = \sigma$. The convergence criterion is that the residual norm be less than 10^{-6} . First a poor starting vector with all entries the same is used. Table 4 gives residual norms for computing the eigenvalue 27.0 with two versions of Davidson's method. Davidson's

TABLE 4
PRECONDITIONING METHODS

<i>j</i>	Residual norm for Ritz value near 27.0			
	Poor starting vector		Good starting vector	
	Standard	Interior	Standard	Interior
1	17.3	17.3	0.67E-2	0.67E-2
5	2.38	0.56	0.45E-3	0.12E-3
10	1.02	0.31	0.63E-3	0.43E-4
15	0.86	0.23	0.12E-2	0.15E-4
20	1.55	0.40E-1	0.16E-2	0.45E-5
25	0.49	0.12E-1	0.18E-3	0.14E-5
30	1.98	0.88E-3	0.30E-2	0.13E-6
35	0.29	0.35E-5	0.13E-6	0.88E-8
37	2.73	0.45E-6	0.12E-6	0.74E-9

method with the interior Rayleigh-Ritz method converges by step 37. This is quite rapid compared to the convergence of Lanczos methods. Standard Lanczos gives a residual norm of only 1.6 at step 150, while interior Lanczos has 0.30 at $j = 150$. If interior Davidson is allowed to continue, three eigenvalues converge by step 53. Now for the standard Davidson method, the approximate eigenpair nearest to but less than 27.05 is chosen to be (ρ, y) . But the method flounders because this choice does not give the proper approximate eigenpair. Ghost Ritz values interfere.

If there is an accurate starting vector, the situation is different. The selection of the proper Ritz value is not difficult in the standard Davidson method. Pick the solution g_i of the reduced problem (1) with the largest first component, so that the approximate eigenvector y_i will have the largest component in the direction of the starting vector. Table 4 also compares the standard and interior Davidson methods with an accurate starting vector. Standard Davidson finds the eigenvalue 27.0 in 33 steps, but the convergence is uneven. Nearby ghost values decrease the accuracy of the best approximation. Interior Davidson converges in 26 iterations. It appears that the interior method helps not only in the selection of the proper Ritz pair, but also in the development of a good subspace.

5. CONCLUSION

For finding interior eigenvalues, the Raleigh-Ritz procedure is best with an inverted operator. But often implementation of the inverted operator is

too expensive or requires too much storage. Given a subspace spanned by the columns of P , it is possible to implicitly use the inverted operator or the inverted operator squared on the subspace spanned by $Q = (A - \sigma I)P$. And approximations from the subspace spanned by P can be formed. The results are generally better than for the standard Rayleigh-Ritz procedure, and the additional costs are small.

The two variations of interior Rayleigh-Ritz are fairly comparable. Both have problems if σ is unfortunately located. Method 1 may compute approximations to an eigenvalue problem with a continuous spectrum. Method 2 may alias distinct eigenvalues. But even then, results may be better than for standard Rayleigh-Ritz. It is possible to adjust σ dynamically to avoid the problem areas. We prefer method 1 because the problem case seems less likely and easier to determine.

The interior version of the Lanczos algorithm uses the same subspace as does standard Lanczos, but extracts the solution better. Interior Lanczos is worthwhile if rough approximations are desired quickly or if restarting is used.

With preconditioning methods, the interior Rayleigh-Ritz method is particularly important because the development of the subspace depends on having accurate approximate eigenpairs along the way. It is also worth noting that the preconditioning methods can be much more effective than Lanczos for finding interior eigenvalues. The approximate factorization generally used for a preconditioner can provide a compromise between the often slow convergence of Lanczos with A and the great factorization expense of shift-and-invert Lanczos.

The author wishes to thank the referees for their insightful suggestions that improved the presentation.

REFERENCES

- 1 S. F. Ashby, T. A. Manteuffel, and P. E. Saylor, Adaptive polynomial preconditioning for Hermitian indefinite linear systems, *BIT* 29:583–609 (1989).
- 2 A. K. Cline, G. H. Golub, and G. W. Platzman, Calculation of normal modes of oceans using a Lanczos method, in *Sparse Matrix Computations* (J. R. Bunch and D. J. Rose, Eds.), Academic, New York, 1976, pp. 409–426.
- 3 E. R. Davidson, The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices, *J. Comput. Phys.* 17:87–94 (1975).
- 4 C. de Boor and J. R. Rice, Extremal polynomials with application to Richardson iteration for indefinite linear systems, *SIAM J. Sci. Statist. Comput.* 3:47–57 (1982).

- 5 A. Draux, *Polynômes Orthogonaux Formels-Applications*, Lecture Notes in Math. 974, (Springer-Verlag, 1983).
- 6 T. Ericsson and A. Ruhe, The spectral transformation Lanczos method for the numerical solution of large sparse generalized symmetric eigenvalue problems, *Math. Comp.* 35:1251–1268 (1980).
- 7 R. Freund, Private communication, 1989.
- 8 A. Greenbaum, Behavior of slightly perturbed Lanczos and conjugate-gradient recurrences, *Linear Algebra Appl.* 113:7–63 (1989).
- 9 S. Kaniel, Estimates for some computational techniques in linear algebra, *Math. Comp.* 20:369–378 (1966).
- 10 R. B. Morgan and D. S. Scott, Generalizations of Davidson's method for computing eigenvalues of sparse symmetric matrices, *SIAM J. Sci. Statist. Comput.* 7:817–825 (1986).
- 11 C. C. Paige, The Computation of Eigenvalues and Eigenvectors of Very Large Sparse Matrices, Ph.D. Dissertation, Univ. of London, 1971.
- 12 C. C. Paige and M. A. Saunders, Solution of sparse indefinite systems of linear equations, *SIAM J. Numer. Anal.* 12:617–629 (1975).
- 13 B. N. Parlett, *The Symmetric Eigenvalue Problem*, Prentice-Hall, Englewood Cliffs, N.J., 1980.
- 14 Y. Saad, On rates of convergence of the Lanczos and the block-Lanczos methods, *SIAM J. Numer. Anal.* 17:687–706 (1980).
- 15 D. S. Scott, The advantages of inverted operators in Rayleigh-Ritz approximations, *SIAM J. Sci. Statist. Comput.* 3:68–75 (1982).
- 16 H. Tanaka, Global energetics analysis by expansion into three-dimensional normal mode functions during FGGE winter, *J. Meteorol. Soc. Japan* 63:180–200 (1985).

Received 15 January 1990; final manuscript accepted 16 August 1990