Technical Report
RAL-TR-1999-025

The Rational Lanczos Method for the Hermitian Eigenvalue Problem

K Meerbergen

30th April 1999

COUNCIL FOR THE CENTRAL LABORATORY OF THE RESEARCH COUNCILS

The Central Laboratory of the Research Councils
Library and Information Services
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire
OX11 0QX
Tel: 01235 445384 Fax: 01235 446403
E-mail library@rl.ac.uk

ISSN 1358-6254

RAL-TR-1999-025

The rational Lanczos method for the Hermitian eigenvalue problem

Karl Meerbergen

ABSTRACT

Applications such as the modal analysis of structures and acoustic cavities require a number of eigenvalues and eigenvectors of large scale Hermitian eigenvalue problems. The most popular method is probably the spectral transformation Lanczos method. An important disadvantage of this method is that a change of pole requires a complete restart. In this paper, we investigate the use of the rational Krylov method for this application. This method does not require a complete restart after a change of pole. We prove that for a specific implementation, numerical instabilities can occur when the pole is not chosen carefully. Numerical examples illustrate the theory.

Keywords: Hermitian eigenvalue problem, Lanczos method, rational Krylov sequence, spectral transformation

AMS (MOS) subject classifications: 65F15

Computational Science and Engineering Department
Atlas Centre
Rutherford Appleton Laboratory
Oxon OX11 0QX

March 31, 1999

Contents

1 Introduction

The subject of this paper is the computation of a number of eigenvalues and the corresponding eigenvectors of the matrix pencil

$$Ax = \lambda Bx \quad , \quad x \neq 0 \; , \tag{1.1}$$

where A and B are large and sparse Hermitian matrices and B is positive (semi) definite. We call λ an eigenvalue and x a corresponding eigenvector ; (λ, x) is called an eigenpair. Applications of this form arise in structural engineering (Grimes, Lewis and Simon 1986) and acoustic modal analysis (Pierce 1981). The most commonly used methods are the spectral transformation Lanczos method (Ericsson and Ruhe 1980, Nour-Omid, Parlett, Ericsson and Jensen 1987) and its block version (Grimes, Lewis and Simon 1994). The Lanczos method builds a Krylov space for the spectral transformation $(A-\mu B)^{-1}B$ starting from a nonzero vector v_1. The method typically converges quickly for the eigenvalues of (1.1) near the pole μ. Very often, a relatively large number of eigenvalues is wanted, say 10, a 100 or more, and then it may be advantageous to change the pole to speed-up the convergence of the Lanczos method. Unfortunately, this requires restarting the method with a new v_1. The disadvantage is that the Krylov basis is completely thrown away. We propose using the rational Krylov method, developed by Ruhe (Ruhe 1984), which does not require a complete restart when μ changes.

In the Lanczos method, the approximate eigenpairs, called Ritz pairs, are computed from a Hermitian tridiagonal matrix and in the rational Krylov method from an upper Hessenberg matrix pair. In this paper, we use an alternative for the rational Krylov method, first formulated in (Ruhe 1998), so that the Ritz values are computed from a Hermitian matrix. For this reason, we call the method rational Lanczos. It differs from rational Krylov as Lanczos (Hermitian problems) differs from Arnoldi (non-Hermitian problems). The contribution of this paper is a stability analysis that shows that the pole should be selected carefully in order to prevent numerical instabilities.

The plan of the paper is as follows. First, we review the spectral transformation Lanczos method in §2. We introduce the rational Lanczos method in §3 and show some of its properties. In §4, we explain how converged Ritz pairs can be locked and how we can purge unwanted Ritz vectors for restarting purposes. This is not new work, but is used by a numerical example. In §5, the numerical stability of the method is studied. In §6, we present some numerical results. We conclude the paper with the main conclusions in §7.

Throughout the paper, x^*y denotes the standard inner product, and $\|x\|$ the induced two-norm, x^*By the B inner product and $\|x\|_B = \sqrt{x^*Bx}$ the induced B norm. The B norm of the matrix C, denoted by $\|C\|_B$, is defined by the two-norm $\|\sqrt{B}C\|$ with $(\sqrt{B})^2 = B$. By $\mathbf{u}$, we denote the machine precision defined as the difference between 1 and the next floating point number (Higham 1996). By $\lambda(A, B)$ we denote the set of eigenvalues of (1.1).

2 The spectral transformation Lanczos method

In this section, we describe the spectral transformation Lanczos method for Hermitian eigenvalue problems. For theoretical and implementation details, we shall refer to the literature.

The aim is to compute, for given v_1 and μ, a basis $v_1, \ldots, v_{k+1}$ for the Krylov space

$$\mathcal{K}_{k+1} = \text{span}\{v_1, (A-\mu B)^{-1}Bv_1, \ldots, ((A-\mu B)^{-1}B)^k v_1\} \ . \tag{2.1}$$

The matrix $(A-\mu B)^{-1}B$ is called the spectral transformation. The name spectral transformation comes from the fact that if (λ, x) satisfies (1.1) then $(\theta = (\lambda-\mu)^{-1}, x)$ is an eigenpair of $(A-\mu B)^{-1}B$. The spectral transformation Lanczos method computes eigenpairs (θ, x) of the spectral transformation. The relation $\lambda = \mu + \theta^{-1}$ allows us to compute the corresponding λ's.

The (Hermitian) Lanczos method (Lanczos 1950, Parlett 1980, Ericsson and Ruhe 1980, Nour-Omid et al. 1987) builds a Krylov space for Hermitian matrices, but because the spectral transformation is non-Hermitian, the method cannot be used for $(A-\mu B)^{-1}B$ in its standard form. Since $(A-\mu B)^{-1}B$ is self-adjoint with respect to the B inner product, the method can be used when B orthogonalization is employed. (See (Nour-Omid et al. 1987) for the details.) An algorithm is given below.

Algorithm 1 (Spectral transformation Lanczos)
1. Given $v_1 \in \boldsymbol{R}^n$ such that $v_1^* B v_1 = 1$.
   Let $\beta_0 = 0$ and $v_0 = 0$.
2. For $j = 1, ..., k$
   2.1. Transformation : $w_j = (A-\mu B)^{-1}Bv_j$.
   2.2. Computation of coefficient : $\alpha_j = v_j^* B w_j$.
   2.3. B orthogonalization : $s_j = w_j - \alpha_j v_j - \beta_{j-1} v_{j-1}$.
   2.4. Computation of coefficient : $\beta_j = \|s_j\|_B$.
   2.5. Normalization : $v_{j+1} = s_j/\beta_j$.
   Endfor

From Steps 2.1–2.5 it follows that $V_{k+1} = [v_1, ..., v_{k+1}]$ forms a B orthogonal basis of the Krylov space (2.1). In fact, Step 2.3 is a Gram-Schmidt orthogonalization step, where α_j and β_{j-1} are Gram-Schmidt coefficients. This orthogonalization step can lose stability and therefore one often employs modified Gram-Schmidt with partial reorthogonalization (Grimes et al. 1994) or even full reorthogonalization (Daniel, Gragg, Kaufman and Stewart 1976, Sorensen 1992). Elimination of w_j and s_j from Algorithm 1 leads to the three term recurrence relation

$$(A-\mu B)^{-1}Bv_j = v_{j+1}\beta_j + v_j\alpha_j + v_{j-1}\beta_{j-1} \ .$$

Collecting the three term recurrence relations for $j = 1, \ldots, k$ leads to

$$(A-\mu B)^{-1}BV_k = V_{k+1}\underline{T}_k \ , \tag{2.2}$$

where

$$\underline{T}_k = \begin{bmatrix} \alpha_1 & \beta_1 & & \\ \beta_1 & \ddots & \ddots & \\ & \ddots & \ddots & \beta_{k-1} \\ & & \beta_{k-1} & \alpha_k \\ & & & \beta_k \end{bmatrix}$$

is a $k+1 \times k$ tridiagonal matrix, formed by the coefficients α_j, β_j. Note that $\underline{T}_k = V_{k+1}^* B(A-\mu B)^{-1} B V_k$. Let T_k be the $k \times k$ leading submatrix of $\underline{T}_k$. Then another notation for the recurrence relation is

$$(A-\mu B)^{-1} B V_k = V_k T_k + v_{k+1}\beta_k e_k^* \ .$$

The B orthogonal projection of $(A-\mu B)^{-1}Bx = \theta x$ onto the Krylov subspace produces approximate eigenpairs, called Ritz pairs, and can be computed as follows. Let (θ, z) satisfy $T_k z = \theta z$, then (θ, y) with $y = V_k z$ is called a Ritz pair with residual

$$\begin{aligned} r &= (A-\mu B)^{-1} B y - \theta y \\ &= \beta_k v_{k+1} e_k^* z \end{aligned} \tag{2.3}$$

and residual norm $\rho = \|r\|_B = \sqrt{r^* B r} = \beta_k |e_k^* z|$.

If $\beta_j = 0$ at Step 2.4 of Algorithm 1, then $\mathrm{span}\{v_1, \ldots, v_j\}$ forms an invariant subspace of (1.1). This means that the corresponding Ritz pairs are exact. This is very unlikely to occur in practice and therefore, we assume that $\beta_1, \ldots, \beta_k > 0$. In this case, $\underline{T}_k$ is called unreduced (Golub and Van Loan 1996, page 346) and is of full rank.

3 The rational Lanczos method

The rational Lanczos method is an extension of the spectral transformation Lanczos method, that makes a change of pole μ to ν possible without restarting the method. It is a special case of the rational Krylov method, which is extensively discussed by Ruhe (Ruhe 1984, Ruhe 1994, Ruhe 1998).

By multiplying (2.2) by $A-\mu B$ and reorganising the terms, we have

$$A V_{k+1} \underline{T}_k = B V_{k+1} (\underline{I}_k + \mu \underline{T}_k) \ . \tag{3.1}$$

This is the rational Krylov recurrence relation of Ruhe (Ruhe 1984). (The columns of $\underline{T}_k$ follow from the orthogonalization process, while the columns of $\underline{I}_k$ are the 'continuation' vectors.) Suppose we want to change the pole to ν. Rewrite (3.1) as

$$(A-\nu B) V_{k+1} \underline{T}_k = B V_{k+1} (\underline{I}_k + (\mu-\nu)\underline{T}_k) \ . \tag{3.2}$$

Let $\underline{L}_k = \underline{I}_k + (\mu-\nu)\underline{T}_k$ and consider the QR decomposition

$$\underline{L}_k = Q\underline{R} = Q \begin{bmatrix} R \\ 0 \end{bmatrix}$$

where Q is a $k+1 \times k+1$ unitary matrix and R is upper triangular. It is important to note that $\underline{L}_k$ is unreduced tridiagonal and thus of full rank. This implies that R is invertible. Let $W_{k+1} = V_{k+1}Q$ and W_k denote the first k columns of W_{k+1}. Then by multiplying (3.2) on the left by $(A-\nu B)^{-1}$ and on the right by R^{-1}, we have

$$\begin{aligned} (A-\nu B)^{-1}BW_k &= W_{k+1}\underline{K}_k \\ \underline{K}_k &= Q^*\underline{T}_k R^{-1} \\ W_{k+1}^* B W_{k+1} &= I \\ W_{k+1}^* B(A-\nu B)^{-1}BW_k &= \underline{K}_k \ . \end{aligned} \tag{3.3}$$

These equations denote a relation similar to the Lanczos recurrence relation. The main difference is that $\underline{K}_k$ is not a tridiagonal matrix. The Krylov subspace can now be expanded by adding

$$(A-\nu B)^{-1}Bw_{k+1}, ((A-\nu B)^{-1}B)^2 w_{k+1}, ...$$

to Range(W_k), as for the Lanczos method.

Note that the tridiagonalization of $\underline{K}_k$ is straightforward by the application of orthogonal transformations on both sides of $\underline{K}_k$ and on the right of W_{k+1}. (See the discussion on the QR method in (Golub and Van Loan 1996) and Appendix A). In that case, (3.3) can really be considered as a three-term Lanczos recurrence relation.

We now have a procedure for changing the pole of the spectral transformation Lanczos method without a complete restart. Note that implicit restarts (Sorensen 1992), deflation and purging (Lehoucq and Sorensen 1996), developed for the Lanczos method, are now possible in combination with a change of pole. This approach is mathematically equivalent, but quite different in practice, from the implicitly filtered rational Krylov method (De Samblanx, Meerbergen and Bultheel 1997) that works immediately on the rational Krylov recurrence relation (3.1). An algorithm for the change of pole is given below.

Algorithm 2 (Change of Pole)

1. Form $\underline{L}_k = \underline{I}_k + (\mu-\nu)\underline{T}_k$.
2. Factorize $\underline{L}_k = Q \begin{bmatrix} R \\ 0 \end{bmatrix}$.
3. Compute $\underline{K}_k = Q^*\underline{T}_k R^{-1} = \frac{1}{\mu-\nu}\left(I - Q^* \begin{bmatrix} R^{-1} \\ 0 \end{bmatrix}\right)$.
4. Apply the orthogonal transformation P such that $\begin{bmatrix} P & 0 \\ 0 & 1 \end{bmatrix}^* \underline{K}_k P$ is tridiagonal.
4. Update $W_{k+1} = V_{k+1}Q \begin{bmatrix} P & 0 \\ 0 & 1 \end{bmatrix}$.

We define the rational Lanczos method by a combination of Algorithms 1 and 2, i.e. after Step 2.5 in Algorithm 1, a change of pole is possible at any time. An algorithm for a specific application is given in Algorithm 3. The rational Lanczos method is an extension of the Lanczos method, so it inherits many properties of the latter. The only difference lies in the change of pole. This has implications on the Krylov subspace :

Theorem 3.1 *Let V_{k+1} and $\underline{T}_k$ be computed by the spectral transformation Lanczos method with pole μ and let W_{k+1} and $\underline{K}_k$ be the result of the change of pole, then*

$$\begin{aligned} \text{Range}(W_{k+1}) &= \text{Range}(V_{k+1}) \ , & (3.4) \\ \text{Range}(W_k) &= \text{Range}((A-\mu B)^{-1}(A-\nu B)V_k) \ . & (3.5) \end{aligned}$$

Proof Since $W_{k+1} = V_{k+1}Q$, (3.4) follows. Equation (3.5) is shown as follows. Since $W_{k+1} = V_{k+1}Q$ and $\underline{L}_k = Q\underline{R}$, we have $W_{k+1}\underline{R} = W_k R = V_{k+1}\underline{L}_k$. Following the definition of $\underline{L}_k$ and (2.2), we have

$$\begin{aligned} V_{k+1}\underline{L}_k &= V_k + (\mu-\nu)V_{k+1}\underline{T}_k \\ &= V_k + (A-\mu B)^{-1}BV_k \\ &= (A-\mu B)^{-1}(A-\nu B)V_k \ . \end{aligned}$$

The proof follows from the observation that $\text{Range}(W_k) = \text{Range}(V_{k+1}\underline{L}_k)$. □

So, the eigenvalues near μ are enhanced in $\text{Range}(W_k)$ and those near ν are damped by the function $(\lambda-\nu)/(\lambda-\mu)$. This will be illustrated by a numerical example in §6. (Note that Algorithm 2 can be considered as a transition of harmonic Ritz values with target μ to harmonic Ritz values with target ν (Sleijpen and van der Vorst 1996).)

4 Purging, implicit restarts and locking

When k is very large, the storage of the basis vectors becomes prohibitive and it may be interesting to restart the Lanczos method. An elegant restarting algorithm is the implicitly restarted Lanczos (or Arnoldi) method (Sorensen 1992, Calvetti, Reichel and Sorensen 1994). This method compresses the Lanczos basis into one of smaller dimension by throwing away a part of the subspace that is unlikely to have a large contribution to the convergence of the wanted eigenvalues. After the compression, the Lanczos method can add new Lanczos vectors to the basis. A similar restart was developed for the rational Krylov method (Ruhe 1998, De Samblanx et al. 1997).

The following is not new, but we have to explain it in order to discuss the results of a numerical example in §6. We can use an implicit restart (Sorensen 1992) in the Lanczos method or equivalently, we can purge unwanted Ritz vectors (Lehoucq and Sorensen 1996, Morgan 1996). Let $T_k Z_k = Z_k \Theta_k$ with $\Theta_k = \text{diag}(\theta_1, \ldots, \theta_p)$ be the eigendecomposition of T_k. Multiplying (2.2) on the right by Z_k, we have

$$(A-\mu B)^{-1}BV_k Z_k = V_{k+1}\underline{T}_k Z_k \ .$$

With

$$Y_{k+1} = [Y_k \ \ y_{k+1}] = V_{k+1}\begin{bmatrix} Z_k & \\ & 1 \end{bmatrix} \quad \text{and} \quad \underline{S}_k = \begin{bmatrix} Z_k & \\ & 1 \end{bmatrix}^* \underline{T}_k Z_k = \begin{bmatrix} \Omega_k \\ \beta_k e_k^* Z_k \end{bmatrix} ,$$

we get the relation

$$(A-\mu B)^{-1}BY_k = Y_{k+1}\underline{S}_k \ .$$

Note that $\underline{S}_k$ is a matrix with $k+1$ rows and k columns that consists of a $k \times k$ diagonal matrix in the k first rows with the Ritz values $\theta_1, \ldots, \theta_k$ on the main diagonal and $\beta_k e_k^* Z_k$ in the $k+1$ st row. The last row contains the residual information for the Ritz pairs. The columns of Y_k are the B orthogonal set of Ritz vectors and the last column $y_{k+1} = v_{k+1}$ is the direction of the residual (5.2).

Suppose that the columns of Y_k and $\underline{S}_k$ are ordered so that the p Ritz values of interest are in the leading columns of $\underline{S}_k$ and the $k-p$ less interesting Ritz values are at the end. When we take the first p columns $\underline{S}_k$ into $\underline{S}_p$, we have

$$(A-\mu B)^{-1} B Y_p = [Y_p \;\; y_{k+1}] \underline{S}_p \quad \text{with} \quad \underline{S}_p = \begin{bmatrix} \theta_1 & & \\ & \ddots & \\ & & \theta_p \\ & \beta_k e_k^* Z_p & \end{bmatrix} .$$

Transforming the $p+1 \times p$ matrix $\underline{S}_p$ to tridiagonal form to $\tilde{T}_p$ by orthogonal transformations and applying these transformations to Y_p, yields a new recurrence relation

$$(A-\mu B)^{-1} B \tilde{V}_p - \tilde{V}_{p+1} \underline{\tilde{T}}_p = 0 \; . \tag{4.1}$$

From here, $k-p$ additional Lanczos iterations produce a Lanczos basis of dimension $k+1$. We call this truncation plus transformation to tridiagonal form purging. It is also achieved by $k-p$ implicit restarts with exact shifts (Sorensen 1992, Morgan 1996).

Ritz pairs with a residual norm $\beta_k |e_k^* z_j|$ smaller than a given tolerance are locked. They are considered as exact eigenpairs. We lock them in the recurrence relation by setting the elements $1, \ldots, q$ in the $p+1$ st row of $\underline{S}_p$ equal to zero. If we assume that the locked Ritz pairs are exact eigenpairs, (4.1) takes the form

$$(A-\mu B)^{-1} B \begin{bmatrix} Y_q & \tilde{V}_{p-q} \end{bmatrix} - \begin{bmatrix} Y_q & \tilde{V}_{p-q+1} \end{bmatrix} \begin{bmatrix} \Theta_p & 0 \\ 0 & \underline{\tilde{T}}_{p-q} \end{bmatrix} = 0 \tag{4.2}$$

with $\Theta_p = \mathrm{diag}(\theta_1, \ldots, \theta_q)$ and $\underline{\tilde{T}}_{p-q}$ a $p-q+1 \times p-q$ tridiagonal matrix. The Lanczos method can be continued from this point. Since the Ritz pairs are not exact eigenpairs, the right-hand side of (4.2) is nonzero : the error on the recurrence relation is of the order of the residual tolerance.

When (θ, x) is a locked Ritz pair for $(A-\mu B)^{-1} B$, we consider (λ, x) with $\lambda = \mu + \theta^{-1}$ a locked Ritz pair for $Ax = \lambda Bx$. When the pole changes to ν, we consider $((\lambda-\nu)^{-1}, x)$ a Ritz pair for $(A-\nu B)^{-1} B$. When a Ritz pair is locked and the pole changes, the residual norm also changes. There is a risk that the residual norm becomes much larger. Locking Ritz pairs with a relatively large residual norm prevents the accurate computation of other Ritz pairs.

The following theorem gives a new upper bound on the residual norm after change of pole.

Theorem 4.1 *Let $((\lambda-\mu)^{-1}, x)$ be a Ritz pair for $(A-\mu B)^{-1} B$ with residual norm*

$$\rho = \| (A-\mu B)^{-1} B x - (\lambda-\mu)^{-1} x \|_B \; .$$

After the change of pole, the new residual norm satisfies

$$\gamma = \|(A-\nu B)^{-1}Bx - (\lambda-\nu)^{-1}x\|_B \leq \left|\frac{\lambda-\mu}{\lambda-\nu}\right| \left(\max_{\lambda\in\lambda(A,B)} \left|\frac{\lambda-\mu}{\lambda-\nu}\right|\right) \|\rho\ .$$

Proof By multiplying $(A-\mu B)^{-1}Bx - \theta x = r$ with $\theta = (\lambda-\mu)^{-1}$ by $(A-\mu B)$, we have

$$Ax(-\theta) + Bx(1+\mu\theta) = (A-\mu B)r\ .$$

By reorganising the equation into

$$(A-\nu B)x(-\theta) + Bx(1+(\mu-\nu)\theta) = (A-\mu B)r$$

and by multiplying with $(A-\nu B)^{-1}$, we have

$$\begin{aligned}(A-\nu B)^{-1}Bx - \frac{\theta}{1+(\mu-\nu)\theta}x &= \frac{1}{1+(\mu-\nu)\theta}(A-\nu B)^{-1}(A-\mu B)r\ , \\ &= \frac{\lambda-\mu}{\lambda-\nu}(I+(\nu-\mu)(A-\nu B)^{-1}B)r\ .\end{aligned}$$

The proof follows from

$$\begin{aligned}\|(I+(\nu-\mu)(A-\nu B)^{-1}B)r\|_B &= \|(I+(\nu-\mu)\sqrt{B}(A-\nu B)^{-1}\sqrt{B})\sqrt{B}r\| \\ &\leq \left(\max_{\lambda\in\lambda(A,B)}\left|1+\frac{\nu-\mu}{\lambda-\nu}\right|\right)\|r\|_B\ .\end{aligned}$$

□

It is clear that when ν is close to an eigenvalue, $\|(A-\nu B)^{-1}(A-\mu B)\|$ can be very large, and so may the residual norms of locked vectors after the change of pole. This is a good reason for not choosing the pole close to a Ritz value. It is also clear that the factor $(\lambda-\mu)/(\lambda-\nu)$ does not change the residual norm much when λ is far away from μ and ν. Ritz values near μ get a smaller residual norm, however. Those near ν get a larger residual norm. As a conclusion, we better put the new pole away from the locked Ritz pairs. This is very natural, since it is useless picking the pole ν near locked Ritz values.

5 The numerical stability of the rational Lanczos method

A question that remains is the numerical stability of the rational Lanczos method. The numerical behaviour of the Lanczos method is well understood. Typically, the method behaves numerically well when it is carefully implemented : this includes modified Gram-Schmidt orthogonalization with partial reorthogonalization (Grimes et al. 1994) or full reorthogonalization (Daniel et al. 1976). An implicit restart may lose stability when the implicit shift is close to a Ritz value. Purging instead of implicit restarting is preferred in that situation (Lehoucq and Sorensen 1996). Also see (Meerbergen and Spence 1997) for another numerical stability result. The only difference between the Lanczos method and the rational Krylov method is the change of pole, so it is sufficient to study the stability of this step.

The following is the main conclusion of this paper.

The transition from pole μ to ν may make the recurrence relation lose accuracy, when ν lies close to an eigenvalue of $Ax = \lambda Bx$.

We now give a mathematical justification for this statement. Suppose that the spectral transformation Lanczos method produces an accurate recurrence relation, i.e.

$$(A - \mu B)^{-1} B V_k - V_{k+1} \underline{T}_k = E$$

with $\|E\|$ small compared to the other terms in the equation. We assume that when (modified) Gram-Schmidt orthogonalization (with reorthogonalization) is used for making $(A - \mu B)^{-1} B v_j$ orthogonal versus v_{j-1} and v_j in Algorithm 1, we have

$$\begin{aligned} \|E\|_B &\sim \|(A - \mu B)^{-1} B V_k\|_B \mathbf{u} \\ &\leq \|\sqrt{B}(A - \mu B)^{-1}\sqrt{B}\|\mathbf{u} \\ &= \max_{\lambda \in \lambda(A,B)} |\lambda - \mu|^{-1} \mathbf{u} \, . \end{aligned}$$

See (Björck 1996, §2.4.5) and (Higham 1996, Theorem 18.12) for modified Gram-Schmidt when $B = I$. The error E can become large when μ is close to an eigenvalue. Suppose that Algorithm 2 is executed in exact arithmetic. (This is a simplification of reality, but even under this assumption, loss of stability may occur.) Then

$$(A - \nu B)^{-1} B W_k - W_{k+1} \underline{K}_k = (A - \nu B)^{-1}(A - \mu B) E R^{-1} \, . \tag{5.1}$$

The factor $(A - \nu B)^{-1}(A - \mu B)$ usually has a modest norm unless ν is close to an eigenvalue, see the following lemma.

Lemma 5.1

$$\|(A - \nu B)^{-1}(A - \mu B)\|_B \leq \sqrt{\|B\|} \max_{\lambda \in \lambda(A,B)} \left| \frac{\lambda - \mu}{\lambda - \nu} \right| \, .$$

Proof First note that

$$\|\sqrt{B}(A - \nu B)^{-1}\sqrt{B}\| = \max_{\lambda \in \lambda(A,B)} \left| \frac{1}{\lambda - \nu} \right| \, .$$

Second, since $(A - \nu B)^{-1}(A - \mu B) = I + (\nu - \mu)(A - \nu B)^{-1} B$, we have

$$\begin{aligned} \|\sqrt{B}(A - \nu B)^{-1}(A - \mu B)\| &= \|(I + (\nu - \mu)\sqrt{B}(A - \nu B)^{-1}\sqrt{B})\sqrt{B}\| \\ &\leq \|I + (\nu - \mu)\sqrt{B}(A - \nu B)^{-1}\sqrt{B})\|\|\sqrt{B}\| \\ &= \sqrt{\|B\|} \max_{\lambda \in \lambda(A,B)} \left| 1 + \frac{\nu - \mu}{\lambda - \mu} \right| \, . \end{aligned}$$

The observation $\|(A - \nu B)^{-1}(A - \mu B)\|_B = \|\sqrt{B}(A - \nu B)^{-1}(A - \mu B)\|$ proves the lemma. □

The factor R^{-1} can play a vital role in (5.1). A lower bound for $\|R^{-1}\|$ can be found as follows.

Lemma 5.2 *Denote the eigenpairs of T_k by (θ_j, z_j) with $\|z_j\| = 2$ for $j = 1, ..., k$ with $\theta_j = (\lambda_j - \mu)^{-1}$. Define the residual norm $\rho_j = \|(A - \mu B)^{-1} B V_k z_j - \theta_j V_k z_j\|_B$. Then*

$$\|R^{-1}\| \geq 1/\min_j \left(\left| \frac{\lambda_j - \nu}{\lambda_j - \mu} \right|^2 + (\mu - \nu)^2 \rho_j^2 \right)^{1/2} .$$

Proof First note that $\|R^{-1}\| = 1/\sigma_{\min}(\underline{L}_k)$ where $\sigma_{\min}$ denotes the smallest singular value. Following the definition of the smallest singular value, we have that $\sigma_{\min}(\underline{L}_k) \leq \|\underline{L}_k z\|$ for any $z \in \boldsymbol{C}^k$ with $\|z\| = 1$. From the assumptions of the lemma, we derive

$$\begin{aligned} \sigma_{\min}^2(\underline{L}_k) \leq \|\underline{L}_k z_j\|^2 &= \left\| \begin{bmatrix} (I + (\mu - \nu)T_k) z_j \\ (\mu - \nu)\beta_k e_k^* z_j \end{bmatrix} \right\|^2 \\ &= \left(\frac{\lambda_j - \nu}{\lambda_j - \mu} \right)^2 + (\mu - \nu)^2 \rho_j^2 \end{aligned}$$

□

When ν is close to λ and the residual norm ρ is small, then $\|R^{-1}\|$ is large, which may imply a large error in the new recurrence relation. The following lemma allows us to make a more precise statement.

Lemma 5.3 *Let K_k denote the upper $k \times k$ submatrix of $\underline{K}_k$. Let $K_k u = \eta u$ with $\|u\| = 1$ and $x = W_k u$ be the corresponding Ritz vector. Then, with $\rho = |e_{k+1}^* \underline{K} u|$ and $\delta = 1 + 2|\mu - \nu| \max(\rho, |\eta|)$, we have*

$$\|(A - \nu B)^{-1} x - \eta x\|_B \leq \rho + \epsilon_1 \delta \tag{5.2}$$

$$|x^* B (A - \nu B)^{-1} B x - \eta| \leq \epsilon_2 \delta^2 . \tag{5.3}$$

with

$$\epsilon_1 = \|(A - \nu B)^{-1}(A - \mu B) E\|_B \leq \left(\max_{\lambda \in \lambda(A,B)} \left| \frac{\lambda - \nu}{\lambda - \mu} \right| \right) \|E\|_B \tag{5.4}$$

$$\epsilon_2 = \frac{2}{|\mu - \nu|} \|(A - \mu B) E\|_B + \|E^* (A - \mu B)(A - \nu B)^{-1} (A - \mu B) E\|_B \tag{5.5}$$

Proof First, with $\rho = |e_{k+1}^* \underline{K} u|$ and $K_k u = \eta u$, we have that

$$\underline{K}_k u = \frac{1}{\mu - \nu} \left(I - Q^* \begin{bmatrix} R^{-1} \\ 0 \end{bmatrix} u \right) = \begin{bmatrix} \eta u \\ \pm \rho \end{bmatrix} .$$

This implies

$$\begin{bmatrix} R^{-1} u \\ 0 \end{bmatrix} = Q(I - (\mu - \nu)\underline{K}_k) u$$

and

$$\|R^{-1} u\| \leq 1 + |\mu - \nu|(\rho + |\eta|) \leq \delta .$$

Second, from (5.1), we have that the exact residual is

$$(A-\nu B)^{-1}Bx-\eta x = r + (A-\nu B)^{-1}(A-\mu B)ER^{-1}u \tag{5.6}$$

with $r = w_{k+1}e_{k+1}^*\underline{K}_k u$ and $\|r\|_B = \rho \equiv |e_{k+1}^*\underline{K}_k u|$, which leads to

$$\|(A-\nu B)^{-1}x-\eta x\|_B \leq \rho + \|(A-\nu B)^{-1}(A-\mu B)ER^{-1}u\|_B$$

and proves (5.2).

In addition, with $f = (A-\mu B)ER^{-1}u$, we have, from (5.6),

$$x^*B(A-\nu B)^{-1}-\eta x^* = r^* + f^*(A-\nu B)^{-1} \ .$$

Using $x^*Bx = 1$ and $r^*Bx = 0$, multiplication on the right by Bx leads to

$$x^*B(A-\nu B)^{-1}Bx-\eta = x^*B(A-\nu B)^{-1}f \ .$$

Combining the last two equations leads to

$$x^*B(A-\nu B)^{-1}Bx-\eta = \eta x^*f + r^*f + f^*(A-\nu B)^{-1}f \ ,$$

from which (5.3) follows. This completes the proof. □

From this lemma, we see that (5.2) and (5.3) depend on δ, which is large when ρ or η are large. When ϵ_1 is modest and ρ is small, then the exact residual norm is still small compared to η, which guarantees a small relative error on the eigenvalue : the error on the eigenvalue is bounded from above by (5.2) following the Bauer-Fike theorem (Saad 1992, Theorem 3.6). As a conclusion, the factor R^{-1} is not dramatic in general. However, when ν is close to an eigenvalue, the factor $(A-\nu B)^{-1}$ may blow up the factor ϵ_1 and maybe ϵ_2.

6 Numerical examples

The first example illustrates the theory, and the second one solves a practical application.

Example 6.1 (Accuracy of the change of pole within rational Lanczos) This example illustrates the rounding error conclusion from §5.

The matrix A is real symmetric and tridiagonal with 2's on the main diagonal and -1's on the two off-diagonals and B is the identity matrix. Both have dimension 200×200. The results of this example are generated using Matlab. We performed $k = 20$ steps of the spectral transformation Lanczos method with $\mu = -1$ and an initial vector with equal elements. The explicitly computed error in the recurrence relation is

$$\begin{aligned} \|E\| = \|(A-\mu B)^{-1}BV_k - V_{k+1}\underline{T}_k\| &\approx 3.19\cdot 10^{-16} \ , \\ \|(A-\mu B)^{-1}BV_k - V_{k+1}\underline{T}_k\|/\|\underline{T}_k\| &\approx 3.19\cdot 10^{-16} \ . \end{aligned}$$

Relative and absolute errors match since $\|\underline{T}_k\| \approx 1$. The Ritz values and their residual norms are shown as bullets in Figure 6.1. Table 6.1 gives the error in the recurrence

Table 6.1: Illustration of Lemma 5.2 for Example 6.1

ν	$\sigma_{\min}(\underline{L}_k)$	absolute error	relative error
2.5	$6.4 \cdot 10^{-2}$	$1.54 \cdot 10^{-13}$	$3.50 \cdot 10^{-14}$
4.2	$7.1 \cdot 10^{-2}$	$1.31 \cdot 10^{-14}$	$4.53 \cdot 10^{-15}$
$3.160930 \cdot 10^{-4}$	$5.5 \cdot 10^{-4}$	$4.71 \cdot 10^{-9}$	$2.58 \cdot 10^{-12}$

relation after change of pole for three different values of ν. For the first and second pole, the Ritz values lie relatively far from the new poles (see the vertical dashed lines in Figure 6.1), so the singular values of $\underline{L}_k$ are not expected to be very small. The third pole is equal to a Ritz value. This implies that $\sigma_{\min}(\underline{L}_k)$ is bounded from above by the corresponding residual norm.

The eigenvalues and corresponding residual norms for the Lanczos process with pole $\mu = -1$, are shown in Figure 6.1. We changed the pole into $\nu_1 = 2.5$ and $\nu_2 = 4.2$. The Ritz values corresponding to the new Lanczos process are plotted as boxes for pole ν_1 and as circles for ν_2. This result corresponds to Theorem 3.1. The residual norms corresponding to the Ritz values close to μ are reduced, however, this is less pronounced when ν is closer to μ. The Ritz values close to ν disappear and/or get larger residual norms.

Table 6.2 illustrates Lemma 5.3 for $\nu = 3.1609 \cdot 10^{-4}$. More specifically, we show

$$\begin{aligned} \zeta_1 &= \|(A-\nu B)^{-1}Bx - \eta x - w_{k+1}e_{k+1}^* \underline{K} u\|_B \qquad (6.1) \\ &= \|(A-\nu B)^{-1}(A-\mu B)ER^{-1}u\|_B \\ \zeta_2 &= |x^*B(A-\nu B)^{-1}Bx - \eta| \qquad (6.2) \end{aligned}$$

and ζ_1/δ and ζ_2/δ^2. The third and fifth columns of Table 6.2 show that ζ_1 and ζ_2 are proportional to δ and δ^2 respectively. The results show that $\epsilon_2 \approx 5 \cdot 10^{-17}$ and $\epsilon_1 \approx 1.4 \cdot 10^{-12}$.

When we choose $\nu = 1.0001\lambda_1$ with λ_1 the smallest eigenvalue of A, then we have that ζ_1 varies between 10^{-9} and 10^{-5}, which is larger than ζ_1 in Table 6.2, and ζ_2 varies between 10^{-10} and 10^{-15}. We also have that $\epsilon_1 \approx 4 \cdot 10^{-9}$ and $\epsilon_2 \approx \cdot 10^{-17}$. Clearly, the factor $(A-\nu B)^{-1}$ blows up the error in the residual (ζ_1).

Example 6.2 (An application) This example is related to the acoustic simulation of a $0.4m \times 0.4m \times 0.06m$ sample made of a poro-elastic material. The material is modelled using a two-phase Biot model accounting for kinematic and mechanical interactions between the (elastic) skeleton and the pore (acoustic) fluid (Sandhu and Pister 1970, Simon, Wu, Zienkiewicz and Paul 1986). The following material properties have been selected : for the skeleton, the Young modulus is $140000N/m^2$, the Poisson ratio 0.35, and the density $1300kg/m^3$. The pore fluid has density $1.225kg/m^3$, the sound speed is $340m/s$, the porosity 0.95, the flow resistivity is 0, the Biot factor is 1, the fluid bulk modulus $141600N/m2$, and the tortuosity is 1.2. The discrete finite-element model relies on a u-w formulation (Simon et al. 1986) where skeleton displacement components (u) and relative

Table 6.2: Shows the errors on Ritz values and residual norms as the right-hand sides of (6.2) and (6.2)

η	ζ_2	ζ_2/δ^2	ζ_1	ζ_1/δ
180.5	$3.4 \cdot 10^{-10}$	$4.7 \cdot 10^{-17}$	$3.59 \cdot 10^{-09}$	$1.4 \cdot 10^{-12}$
47.57	$9.2 \cdot 10^{-11}$	$4.7 \cdot 10^{-17}$	$1.9 \cdot 10^{-09}$	$1.4 \cdot 10^{-12}$
21.10	$4.3 \cdot 10^{-11}$	$5.0 \cdot 10^{-17}$	$1.3 \cdot 10^{-09}$	$1.4 \cdot 10^{-12}$
11.68	$2.3 \cdot 10^{-11}$	$5.0 \cdot 10^{-17}$	$9.6 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
7.291	$1.4 \cdot 10^{-11}$	$4.9 \cdot 10^{-17}$	$7.5 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
4.905	$9.5 \cdot 10^{-12}$	$5.0 \cdot 10^{-17}$	$6.1 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
3.467	$6.4 \cdot 10^{-12}$	$4.8 \cdot 10^{-17}$	$5.0 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
2.536	$4.7 \cdot 10^{-12}$	$4.9 \cdot 10^{-17}$	$4.2 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
1.902	$3.4 \cdot 10^{-12}$	$5.0 \cdot 10^{-17}$	$3.6 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
1.453	$2.4 \cdot 10^{-12}$	$4.8 \cdot 10^{-17}$	$3.1 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
1.126	$1.8 \cdot 10^{-12}$	$4.9 \cdot 10^{-17}$	$2.6 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.8825	$1.2 \cdot 10^{-12}$	$4.7 \cdot 10^{-17}$	$2.2 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.6988	$9.2 \cdot 10^{-13}$	$5.0 \cdot 10^{-17}$	$1.9 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.5594	$6.1 \cdot 10^{-13}$	$4.7 \cdot 10^{-17}$	$1.5 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.4538	$4.1 \cdot 10^{-13}$	$4.9 \cdot 10^{-17}$	$1.3 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.3750	$2.5 \cdot 10^{-13}$	$4.8 \cdot 10^{-17}$	$1.0 \cdot 10^{-10}$	$1.4 \cdot 10^{-12}$
0.3181	$1.6 \cdot 10^{-13}$	$5.5 \cdot 10^{-17}$	$7.3 \cdot 10^{-11}$	$1.4 \cdot 10^{-12}$
0.2796	$5.1 \cdot 10^{-14}$	$4.1 \cdot 10^{-17}$	$4.5 \cdot 10^{-11}$	$1.3 \cdot 10^{-12}$
0.2573	$1.7 \cdot 10^{-14}$	$5.3 \cdot 10^{-17}$	$2.36 \cdot 10^{-11}$	$1.3 \cdot 10^{-12}$
-0.9982	$1.0 \cdot 10^{-10}$	$4.9 \cdot 10^{-17}$	$1.97 \cdot 10^{-09}$	$1.4 \cdot 10^{-12}$

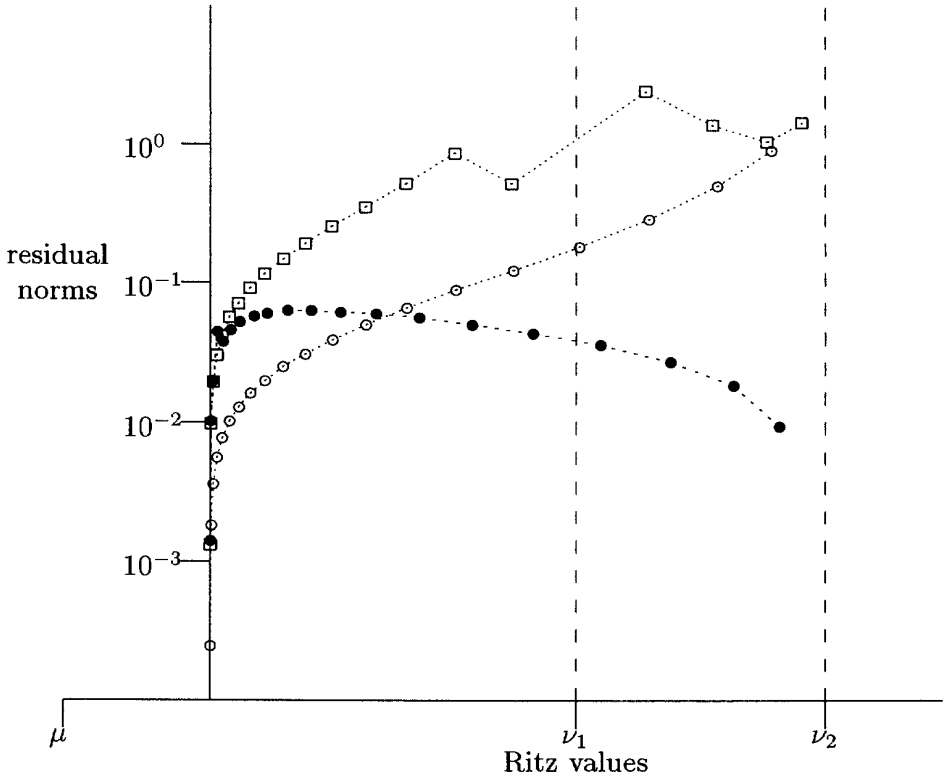

Figure 6.1: Residual norms as a function of the Ritz values. Bullets denote residual norms and Ritz values for Range(V_k). Boxes and circles denote residual norms and Ritz values for Range(W_k) for ν_1 and ν_2 respectively.

fluid displacement components (weighted by the local porosity) (w) are selected as nodal variables. The finite-element mesh has 324 nodes and 192 HEXA8 elements. The total number of degrees of freedom is $n = 1944$.

We use the Lanczos routine EA16 available with the next release of the Harwell Subroutine Library (HSL 2000). The implementation is based on the (block) Lanczos method with partial reorthogonalization (Grimes et al. 1994), purging of the unwanted Ritz pairs and locking of converged Ritz pairs (see §4), an implicit restart (Sorensen 1992), and change of pole (see §3). In this experiment, we use the Lanczos method with block size 1 in order to illustrate the theory in the paper. After purging, the Krylov subspace dimension is reduced from k to p. We used $k = 50$ and $p = 25$. The method is described in Algorithm 3. The sparse linear systems are solved by the Harwell Subroutine Library code MA27 (HSL 1996). For efficiency reasons, the matrix is factorized once for each pole and the Lanczos method uses backtransformations. The residual tolerance used for the eigenvalues was $\epsilon_{\text{tol}} = \sqrt{\mathbf{u}}||T_m||$ with $\sqrt{\mathbf{u}} \approx 1.05 \cdot 10^{-8}$, which is the square root of the machine precision on a SUN Ultra 1. The goal is to compute $s = 20$ eigenpairs on the right of 100.

A picture of the spectrum is shown in Figure 6.2. The intial pole is $\mu = 100$. A new pole is chosen as the mean of two successive Ritz values, as indicated in Algorithm 3. When the distance to a Ritz value is smaller than 50, we pick another pole. With this choice, we try to prevent that the pole is picked close to Ritz or eigenvalues.

Algorithm 3 (Rational Lanczos)

0. Given $v_1 \in \boldsymbol{R}^n$ such that $v_1^* B v_1 = 1$.
   Select an initial pole μ and factorize $A - \mu B = LDL^*$.
1. Perform p steps of the Lanczos method.
2. **do** until convergence :
   2.1. Expand the Krylov space from dimension p to k by the Lanczos method.
   2.2. Compute the eigenpairs (θ_j, z_j) $j = 1, ..., k$ of T_k and let $\lambda_j = \mu + \theta_j^{-1}$.
   2.3. Sort the λ_j's in ascending order and reorder the z_j accordingly.
   2.4. Compute residual norms $\rho_j = \beta_k |e_k^* z_j|$ for $j = 1, \ldots, k$.
   2.5. Stop if $\rho_j \leq \epsilon_{\text{tol}}||T_k||$ for $j = 1, \ldots, s$.
   2.6. Lock the converged Ritz pairs and truncate the Lanczos recurrence relation from order k to p.
   2.7. Compute the new pole as $\mu = (\lambda_j + \lambda_{j+1})/2$ with $j \geq q$ and such that $|\mu - \lambda_j| \geq 50$.
   2.8. Change the pole using Algorithm 2.
   end do

From Lemma 5.3 and (5.1), we can see that $||(A-\nu B)^{-1}(A-\mu B)||$ plays an important role in the error bounds. When $|\lambda - \nu| \geq 50$ for all eigenvalues λ of $Ax = \lambda Bx$, then

$$||(A - \nu B)^{-1}(A - \mu B)||_B \leq \sqrt{B} \max_{\lambda \in \lambda(A,B)} \left| \frac{\lambda - \mu}{\lambda - \nu} \right| \leq \sqrt{B}(1 + |\mu - \nu|/50) \ .$$

This limits the growth of the error in the recurrence relation after a change of pole. Lemma 5.2 is used for measuring the contribution from the factor R^{-1} in (5.1). This

Table 6.3: Number of locked Ritz values, the new pole and the estimate of $\sigma_{\min}(\underline{L}_p)$ following Lemma 5.2

restart	locked	pole	estimate $\sigma_{\min}(\underline{L}_p)$
1	5	1381.	$8.3 \cdot 10^{-2}$
2	12	1764.	$1.8 \cdot 10^{-1}$
3	18	2432.	$1.3 \cdot 10^{-1}$
4	25		

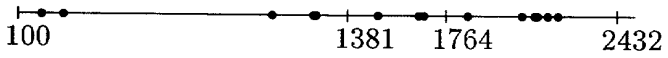

Figure 6.2: Ritz values computed via the rational Lanczos method. The vertical lines denote the positions of the poles.

was illustrated by Example 6.1 and is confirmed again by the results for this example. Table 6.3 shows, after each restart, the number of converged Ritz values and the new pole with the estimate for $\sigma_{\min}(\underline{L}_p)$. Using the inertia count (Grimes et al. 1994), we found that the number of eigenvalues between 100 and the poles 1380 and 1760 matches the number of locked Ritz values in the corresponding intervals. This shows the reliability of the algorithm for this example. After the last iteration, we found that, with full reorthogonalization at each iteration of the Lanczos process, $\max_{i,j \neq i} |w_i^* B w_j| \simeq 5 \cdot 10^{-15}$ and

$$\|(A - \mu B)^{-1} B W_k - W_{k+1} \underline{K}_k\|_F / \|(A - \mu B)^{-1} B W_k\|_F \simeq 2.4 \cdot 10^{-8} ,$$

which is within the residual tolerance. With partial reorthogonalization, we have $\max_{i,j \neq i} |w_i^* B w_j| \simeq 3 \cdot 10^{-8}$ and

$$\|(A - \mu B)^{-1} B W_k - W_{k+1} \underline{K}_k\|_F / \|(A - \mu B)^{-1} B W_k\|_F \simeq 2.5 \cdot 10^{-8} .$$

7 Conclusions

In this paper, we introduced the rational Lanczos method for the solution of Hermitian eigenvalue problems. It can be considered as a spectral transformation Lanczos method with implicit restart and change of pole, or as the rational Krylov method with a Hermitian projected problem. The major conclusion from this paper is that it may be dangerous to select a pole close to a Ritz value for the following reason. First, the linear systems may be difficult to solve, since they are nearly singular. Second, it may not only lead to large absolute and relative errors in the recurrence relation, but also to large residual norms for locked Ritz pairs after a change of pole. It is therefore advised to pick the pole away from the locked Ritz pairs.

The rational Lanczos method is suitable for computing the eigenvalues of structural eigenvalue problems in a frequency range (Grimes et al. 1994) or the right-most eigenvalue in the determination of the stability of steady state solutions of non-linear equations (Meerbergen and Roose 1996), where the pole is not selected near a Ritz value. It is not

advised for algorithms that focus on the convergence of a single eigenvalue by putting the pole as close as possible to that eigenvalue. In this case, numerical instabilities are likely to occur.

Acknowledgements

The author is grateful to Jennifer Scott, Iain Duff, and Axel Ruhe for helpful suggestions that improved the quality of the paper.

References

Björck, Å. (1996), *NUmerical methods for least squares problems*, SIAM, Philadelphia PA 19104-2688, USA.

Calvetti, D., Reichel, L. and Sorensen, D. (1994), 'An implicitly restarted Lanczos method for large symmetric eigenvalue problems', *ETNA* **2**, 1–21.

Daniel, J., Gragg, W., Kaufman, L. and Stewart, G. (1976), 'Reorthogonalization and stable algorithms for updating the Gram-Schmidt QR factorization', *Math. Comp.* **30**, 772–795.

De Samblanx, G., Meerbergen, K. and Bultheel, A. (1997), 'The implicit application of a rational filter in the RKS method', *BIT* **37**, 925–947.

Ericsson, T. and Ruhe, A. (1980), 'The spectral transformation Lanczos method for the numerical solution of large sparse generalized symmetric eigenvalue problems', *Math. Comp.* **35**, 1251–1268.

Golub, G. and Van Loan, C. (1996), *Matrix computations*, 3rd edn, The Johns Hopkins University Press.

Grimes, R., Lewis, J. and Simon, H. (1986), Eigenvalue problems and algorithms in structural engineering, *in* J. Cullum and R. Willoughby, eds, 'Large Scale Eigenvalue Problems', Elsevier Science Publishers B.V., pp. 81–93.

Grimes, R., Lewis, J. and Simon, H. (1994), 'A shifted block Lanczos algorithm for solving sparse symmetric generalized eigenproblems', *SIAM J. Matrix Anal. Applic.* **15**, 228–272.

Higham, N. (1996), *Accuracy and Stability of Numerical Algorithms*, SIAM philiadelphia.

HSL (1996), 'Harwell Subroutine Library. a Catalogue of Subroutines (Release 12)'. Information: Dr Scott Roberts, AEA Technology, 477 Harwell, Didcot, Oxon, OX11 0RA, UK.

Lanczos, C. (1950), 'An iteration method for the solution of the eigenvalue problem of linear differential and integral operators', *J. Res. Nat. Bur. Stand.* **45**, 255–282.

Lehoucq, R. and Sorensen, D. (1996), 'Deflation techniques within an implicitly restarted Arnoldi iteration', *SIAM J. Matrix Anal. Applic.* **17**, 789–821.

Meerbergen, K. and Roose, D. (1996), 'Matrix transformations for computing rightmost eigenvalues of real nonsymmetric matrices', *IMA J. Numer. Anal.* **16**, 297–346.

Meerbergen, K. and Spence, A. (1997), 'Implicitly restarted Arnoldi and purification for the shift-invert transformation', *Math. Comp.*

Morgan, R. (1996), 'On restarting the Arnoldi method for large nonsymmetric eigenvalue problems', *Math. Comp.* **65**, 1213–1230.

Nour-Omid, B., Parlett, B., Ericsson, T. and Jensen, P. (1987), 'How to implement the spectral transformation', *Math. Comp.* **48**, 663–673.

Parlett, B. (1980), *The symmetric eigenvalue problem*, Prentice Hall series in computational mathematics, Prentice-Hall.

Pierce, A. D. (1981), *Acoustics - An introduction to its physical principles and applications*, Mac Graw Hill, New York.

Ruhe, A. (1984), 'Rational Krylov sequence methods for eigenvalue computation', *Linear Alg. Appl.* **58**, 391–405.

Ruhe, A. (1994), 'Rational Krylov algorithms for nonsymmetric eigenvalue problems, II: Matrix pairs', *Linear Alg. Appl.* **197/198**, 283–296.

Ruhe, A. (1998), 'Rational Krylov: a practical algorithm for large sparse nonsymmetric matrix pencils', *SIAM J. Sci. Comput.* **19**(5), 1535–1551.

Saad, Y. (1992), *Numerical methods for large eigenvalue problems*, Algorithms and Architectures for Advanced Scientific Computing, Manchester University Press, Manchester, UK.

Sandhu, R. and Pister, K. (1970), 'A variational principle for linear coupled field problems in continuum mechanics', *Int. J. Eng. Sci.* **8**, pp. 989–999.

Simon, B., Wu, J., Zienkiewicz, O. and Paul, D. (1986), 'Evaluation of u-w and u-p finite element methods for the dynamic response of saturated porous media using one-dimensional models', *Int. J. Numer. Anal. Methods Geomech.* **10**, pp. 461–482.

Sleijpen, G. and van der Vorst, H. (1996), 'A Jacobi-Davidson iteration method for linear eigenvalue problems', *SIAM J. Matrix Anal. Applic.* **17**, 401–425.

Sorensen, D. (1992), 'Implicit application of polynomial filters in a k-step Arnoldi method', *SIAM J. Matrix Anal. Applic.* **13**, 357–385.

A Reduction of a rectangular matrix to tridiagonal form

Consider the $k+1 \times k$ matrix $\underline{T}^{(k)}$ with the upper $k \times k$ submatrix Hermitian. The following algorithm creates a unitary matrix P such that

$$T^{(1)} = \begin{bmatrix} P & 0 \\ 0 & 1 \end{bmatrix}^* \underline{T}^{(k)} P$$

is tridiagonal.

Algorithm 4

1. Let $P = I$
2. For $j = k$ to 2 step -1 do
   2.1. Let H_j be a Householder reflection such that $e_{j+1}^* \underline{T}^{(j)} H_j$ is a row vector with $j-1$ zeros in the front.
   2.2. Form
   $$\underline{T}^{(j-1)} = \begin{bmatrix} H_j & 0 \\ 0 & 1 \end{bmatrix}^* \underline{T}^{(j)} H_j$$
   2.3 Form $P = PH_j$

On iteration j the $j+1$ st row of $\underline{T}^{(j)}$ is reduced to a row with $j-1$ zeros in the front. Since the Householder transformation is also applied on the front without touching the $j+1$ st row, the $j+1$ st column is reduced to a column vector with $j-1$ zeros in the front.