RAL-TR-1999-025 (Revised November 1999)

Dangers in changing poles in 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. The contribution of this paper is threefold. First, it is shown that the change of pole can be considered as a change of Lanczos basis. Second, moving the pole near a locked eigenvalue may prevent other eigenvalues to be computed accurately. Third, we show that a pole chosen close to an eigenvalue may lead to a loss of numerical stability. Usually, this will not lead to disaster but reduce the accuracy of the computed eigenvalues away from the pole. 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

November 25, 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. Unfortunately, this requires restarting the method from the same or 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 (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 (Arnoldi 1951, Saad 1992) (non-Hermitian problems). The rational Krylov method is not very often used for solving applications. One reason is that the method is not very well understood. It is derived in a different way as the Lanczos method and there is a large amount of freedom in setting the various parameters. In this paper, we look at the rational Krylov method from a different angle as Axel Ruhe. The only freedom in the method we propose is the pole. The contribution of this paper is threefold. First, it shows that the change of pole can be considered as a change of Lanczos basis. Second, moving the pole near a locked Ritz value may prevent other eigenvalues to be computed accurately. Third, we show that a pole chosen close to a Ritz value may lead to a loss of numerical stability. Usually, this will not lead to disaster but reduce the accuracy of the Ritz values a bit away from the pole. In this situation, we suggest a complete restart of the Lanczos method with the new pole.

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. This is not new work, but we show that changing the pole after locking Ritz pars may reduce the level of accuracy of the method. In §5, it is shown that the method may lose backward stability when the pole is chosen close to a Ritz value. In §6, we present a numerical result for an application coming from acoustics. We conclude the paper with the main conclusions in §7.

The dimension of the problem (1.1) is denoted by n. By e_j we denote the jth column of the identity matrix I. Throughout the paper, x^*y denotes the standard inner product, and $\|x\|_2$ 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 $C \in \boldsymbol{C}^{n\times n}$, denoted by $\|C\|_B$, is defined by $\max_{\|x\|_B=1} \|Cx\|_B$. This implies that $\|Cx\|_B \leq \|C\|_B\|x\|_B$ for any $x \in \boldsymbol{C}^n$ and any $C \in \boldsymbol{C}^{n\times n}$. We also define the B Frobenius norm of $C \in \boldsymbol{C}^{n\times k}$ as $\|C\|_{BF} = \left(\sum_{j=1}^{k} \|Ce_j\|_B^2\right)^{1/2}$. 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). We also use $\sigma_{\min}(C)$ and $\sigma_{\max}(C)$ to denote the minimum and maximum singular value of C respectively. The condition number is defined by $\kappa_2(C) = \sigma_{\max}(C)/\sigma_{\min}(C)$.

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 and $\mu \in \boldsymbol{R}$ is the pole. 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 (Ericsson and Ruhe 1980).

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 2.1 (Spectral transformation Lanczos)
1. Given $v_1 \in \boldsymbol{C}^n$ such that $\|v_1\|_B = 1$.
   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 : $\alpha_j = v_j^*Bw_j$.
   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 $V_{k+1} = [v_1, \ldots, v_{k+1}]$ 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 partial and external selective 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 2.1 leads to the three term recurrence relation

$$(A - \mu B)^{-1} B v_j = v_{j+1}\beta_j + v_j \alpha_j + v_{j-1}\beta_{j-1} \ . \tag{2.2}$$

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

$$(A - \mu B)^{-1} B V_k = V_{k+1} \underline{T}_k \ , \tag{2.3}$$

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 $(A - \mu B)^{-1} B x = \theta x$ 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 of $(A - \mu B)^{-1} B$ with residual

$$\begin{aligned} r &= (A - \mu B)^{-1} B y - \theta y \\ &= \beta_k v_{k+1} e_k^* z \end{aligned} \tag{2.4}$$

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

If $\beta_j = 0$ at Step 2.4 of Algorithm 2.1, then $\mathrm{span}\{v_1, \ldots, v_j\}$ forms an invariant subspace of (1.1). This means that the corresponding Ritz pairs are eigenpairs. 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.

In order to avoid confusion between Ritz pairs for $(A - \mu B)^{-1} B$ and $Ax = \lambda Bx$, we denote by θ a Ritz value for $(A - \mu B)^{-1} B$ and by η a Ritz value for $Ax = \lambda Bx$ defined as $\eta = \mu + \theta^{-1}$. The values η are also called harmonic Ritz values with target μ (Paige, Parlett and van der Vorst 1995, Sleijpen and van der Vorst 1996, Ruhe 1998). By λ we denote eigenvalues of $Ax = \lambda Bx$.

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 $\nu \in \boldsymbol{R}$ possible without restarting the method.

It is a special case of the rational Krylov method, which is extensively discussed by (Ruhe 1984, Ruhe 1994, Ruhe 1998).

By multiplying (2.3) by $A - \mu B$ and reorganising the terms, we have

$$AV_{k+1}\underline{T}_k = BV_{k+1}(\underline{I}_k + \mu\underline{T}_k) \ . \tag{3.1}$$

This is the rational Krylov recurrence relation of 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 $\nu \neq \mu$. Rewrite (3.1) as

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

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. Note that K_k is Hermitian since $K_k = W_k^* B(A - \nu B)^{-1}BW_k$. 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.

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 Lanczos three-term 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 3.1 (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 transformations P and $\tilde{P} = \begin{bmatrix} P & 0 \\ 0 & 1 \end{bmatrix}$ so that $\underline{K}_k := \tilde{P}^* \underline{K}_k P$ is tridiagonal.

4. Update $W_{k+1} = V_{k+1} Q \tilde{P}$.

We define the rational Lanczos method by a combination of Algorithms 2.1 and 3.1, i.e. after Step 2.5 in Algorithm 2.1, a change of pole is possible. An algorithm for a specific application is given in Algorithm 6.1. 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 computed by Algorithm 3.1, 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 $W_{k+1} = V_{k+1} Q \tilde{P}$, (3.4) follows. Equation (3.5) is shown as follows. Since $W_{k+1} = V_{k+1} Q \tilde{P}$ and $\underline{L}_k = Q \underline{R} = (Q \tilde{P})(\tilde{P}^* \underline{R})$, we have $W_{k+1} \tilde{P}^* \underline{R} = W_k P^* R = V_{k+1} \underline{L}_k$. Following the definition of $\underline{L}_k$ and (2.3), we have

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

The proof of (3.5) follows from the observation that $\text{Range}(W_k) = \text{Range}(V_{k+1} \underline{L}_k)$. □

The change of pole is very similar to an implicit restart in the Lanczos method (Sorensen 1992). In this case, a Lanczos process is implicitly applied to the new starting vector $w_1 = W_k e_1$, i.e. without explicitly computing $w_2, \ldots, w_{k+1}$, and the Krylov subspace is filtered by a matrix polynomial (Meerbergen and Spence 1997, Lehoucq 1996, Lehoucq 1999). Following (3.5), the eigenvalues near μ are enhanced in $\text{Range}(W_k)$ and those near ν are damped by the function $(\lambda - \nu)/(\lambda - \mu)$. (Note that Algorithm 3.1 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).)

When the pole is chosen equal to a harmonic Ritz value, the effect of a change of pole is given by the following theorem.

Theorem 3.2 *Let V_{k+1} and $\underline{T}_k$ be computed by the spectral transformation Lanczos method with pole μ. Let η_1 be a harmonic Ritz value. Choose the new pole $\nu = \eta_1$. Let W_{k+1} and $\underline{K}_k$ be computed by Algorithm 3.1, then after the change of pole, the harmonic Ritz value η_1 disappears and μ becomes a harmonic Ritz value with Ritz vector v_{k+1}. The other harmonic Ritz values and Ritz vectors remain unchanged.*

Proof First note that $\mathcal{R} := \text{Range}(W_k) = \text{Range}((A-\mu B)^{-1}(A-\nu B)X_k)$ where X_k is the matrix of Ritz vectors.

Let (θ_1, x_1) with $\theta_1 = (\eta_1 - \mu)^{-1}$ and $x_1 = V_k z_1$ be a Ritz pair and $r_1 = v_{k+1}\beta_k e_k^* z_1$ the corresponding residual. We shall show that $((\mu-\nu)^{-1}, r_1)$ is a Ritz pair for the pole ν. First, we prove that $r_1 \in \mathcal{R}$. Define $y = (A-\mu B)^{-1}(A-\nu B)x_1 \in \mathcal{R}$. Since $(A-\mu B)^{-1}Bx_1 = \theta_1 x_1 + r_1$, we have

$$\begin{aligned} y_1 &= x_1 + (\mu-\nu)\theta_1 x_1 + (\mu-\nu)r_1 \\ &= \frac{\eta_1 - \nu}{\eta_1 - \mu}x_1 + (\mu-\nu)r_1 = (\mu-\nu)r_1 \ . \end{aligned}$$

We now prove that $((\mu-\nu)^{-1}, r_1)$ is a Ritz pair of $(A-\mu B)^{-1}B$ after the change of pole. Therefore we show that

$$(A-\nu B)^{-1}Br_1 - \frac{1}{\mu-\nu}r_1 \perp_B \mathcal{R} \ .$$

This is equivalent to

$$\left((A-\mu B)^{-1}(A-\nu B)V_k\right)^* B \left((A-\nu B)^{-1}Br_1 - \frac{1}{\mu-\nu}r_1\right) = 0 \ .$$

Elaborating the left-hand side gives

$$V_k^* B(A-\mu B)^{-1}Br_1 - \frac{1}{\mu-\nu}V_k^* Br_1 - \frac{\mu-\nu}{\mu-\nu}V_k^* B(A-\mu B)^{-1}Br_1$$

which is zero since $V_k^* Br_1 = 0$.

Ley $(\eta \neq \eta_1, x)$ be another harmonic Ritz pair for pole μ. We now prove that (η, x) is a harmonic Ritz pair for pole ν. Let $y = (A-\mu B)^{-1}(A-\nu B)x \in \mathcal{R}$. Then

$$y = \frac{\eta-\nu}{\eta-\mu}x + (\mu-\nu)r \ .$$

Since $r = v_{k+1}\beta_k e_k^* z$ and y lie in $\mathcal{R}$, also x lies in $\mathcal{R}$. We still have to show that

$$(A-\nu B)^{-1}Bx - \frac{1}{\eta-\nu}x \perp_B \mathcal{R} \ .$$

This is equivalent to

$$\left((A-\mu B)^{-1}(A-\nu B)V_k\right)^* B \left((A-\nu B)^{-1}Bx - \frac{1}{\eta-\nu}x\right) = 0 \ .$$

Recalling that $(A-\mu B)^{-1}Bx = 1/(\eta-\mu)x + r$, elaboration of the terms in the left-hand side gives

$$\begin{aligned} V_k^* B(A-\mu B)^{-1}Bx - \frac{1}{\eta-\nu}V_k^* Bx - \frac{\mu-\nu}{\eta-\nu}V_k^* B(A-\mu B)^{-1}Bx \\ = \frac{\eta-\mu}{\eta-\nu}V_k^* B(A-\mu B)^{-1}Bx - \frac{1}{\eta-\mu}V_k^* Bx \\ = \frac{\eta-\mu}{\eta-\nu}\left(\frac{1}{\eta-\mu}V_k^* Bx + V_k^* Br\right) - \frac{1}{\eta-\nu}V_k^* Bx \ , \end{aligned}$$

which is zero since $V_k^* Br = 0$. □

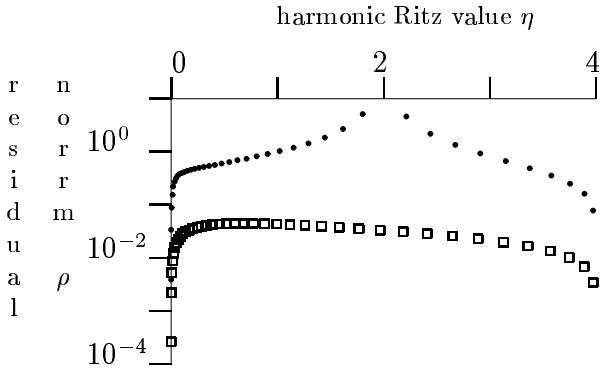

Figure 3.1: Residual norms versus harmonic Ritz values. The squares show the pairs (η, ρ) (see (2.4)) for 40 Lanczos steps with pole $\mu = -1$. The bullets show the pairs (η, ρ) after changing the pole into $\nu = 2$.

This is now illustrated by a numerical example.

Example 3.1 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 on a DEC Alpha. We performed $k = 40$ steps of the spectral transformation Lanczos method with $\mu = -1$ and an initial vector with equal elements.

The five smallest harmonic Ritz values are $\eta_1 \simeq 2.4494\,10^{-4}$, $\eta_2 \simeq 2.2460\,10^{-3}$, $\eta_3 \simeq 6.4159\,10^{-3}$, $\eta_4 \simeq 1.2947\,10^{-2}$, and $\eta_5 \simeq 2.1989\,10^{-2}$. The Ritz value corresponding to η_1 is $\theta_1 \simeq 0.9998$ and has residual norm $2.6\,10^{-4}$.

After altering the pole into $\nu = \eta_1$, the five smallest harmonic Ritz values are $\tilde{\eta}_1 \simeq -1.0000$, $\tilde{\eta}_2 \simeq 2.2460\,10^{-3}$, $\tilde{\eta}_3 \simeq 6.4159\,10^{-3}$, $\tilde{\eta}_4 \simeq 1.2947\,10^{-2}$, and $\tilde{\eta}_5 \simeq 2.1989\,10^{-2}$. The smallest value $\eta_1 = 2.4429\,10^{-4}$ has disappeared. Since ν is a Ritz value, the corresponding Ritz vector is filtered out of the Lanczos basis with the filter $(A-\mu B)^{-1}(A-\nu B)$. Following Theorem 3.2, the harmonic Ritz value η_1 is mapped onto $\mu = -1$.

After one additional Lanczos step with pole ν, the lost Ritz value is recovered, since one Lanczos steps adds w_{k+1} to the space in which the Ritz vectors are computed. The smallest harmonic Ritz value is again $2.4429\,10^{-4}$, and the other smallest harmonic Ritz values are now $2.1989\,10^{-3}$, $6.1193\,10^{-3}$, $1.2061\,10^{-2}$, and $2.0147\,10^{-2}$.

We did the same experiment but now changed the pole into $\nu = 2$. Figure 3.1 shows the harmonic Ritz values (η) with their residual norms (ρ) before (squares) and after (bullets) the change of pole. Since ν is not a harmonic Ritz value, Theorem 3.2 does not apply. Theorem 3.1 still applies, i.e. the subspace is filtered by the rational function $(\lambda - \nu)/(\lambda - \mu)$. This suggests that the Ritz values near ν are pushed away and that the corresponding residual norms grow more than the other residuals. Note that it is totally acceptable that residual norms grow with a factor $(\eta - \mu)/(\eta - \nu)$, since the Ritz values also grow with this factor.

In the implicitly restarted Lanczos method (Sorensen 1992, Calvetti, Reichel and Sorensen 1994) the new tridiagonal matrix can be computed by the action $Q^* T_k Q$ where Q is a unitary matrix. This is a symmetric formula. The new tridiagonal matrix is usually

computed with bulge-chasing (Golub and Van Loan 1996). Algorithm 3.1 presents a nonsymmetric form. A symmetric form can be obtained as follows. Rewrite (3.2) as

$$(A - \nu B)V_k T_k - BV_k L_k = -(A - \mu B)v_{k+1}\beta_k e_k^T \tag{3.6}$$

with $L_k = I + (\mu - \nu)T_k$. Factorize $L_k = QR$, then $R^{-1} = L_k^{-1}Q$. Also define $W_k = V_k Q$ and $K_k = Q^* T_k R^{-1}$, then

$$K_k = \frac{1}{\mu - \nu}(I - Q^* R^{-1}) = \frac{1}{\mu - \nu}(I - Q^* L_k^{-1} Q) \ ,$$

which is a symmetric formula. Unfortunately, it cannot be computed by a QR step (or bulge-chasing) on T_k. Moreover, we still have not yet computed $\underline{K}_k$. Similarly to the implicitly restarted Lanczos method (Sorensen 1992), this can be achieved by manipulating (3.6). We do not consider this approach any further, since there is not advantage compared to Algorithm 3.1.

4 Purging, implicit restarts and locking

When k is very large, the storage of the basis vectors becomes prohibitive and it may be necessary to restart the Lanczos method. An elegant restarting algorithm is the implicitly restarted Lanczos (or Arnoldi) method (Sorensen 1992, Calvetti et al. 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 significant 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).

We can use an implicit restart with exact shifts (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 $\Theta_k = \mathrm{diag}(\theta_1, \ldots, \theta_k)$ be the eigendecomposition of T_k. Multiplying (2.3) on the right by Z_k, we have

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

With

$$X_{k+1} = [X_k \ \ x_{k+1}] = V_{k+1} \begin{bmatrix} Z_k & \\ & 1 \end{bmatrix}$$

and

$$\underline{S}_k = \begin{bmatrix} Z_k & \\ & 1 \end{bmatrix}^* \underline{T}_k Z_k = \begin{bmatrix} \Theta_k \\ \beta_k e_k^* Z_k \end{bmatrix} \ ,$$

we get the relation

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

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 X_k are the B orthogonal set of Ritz vectors and the last column $x_{k+1} = v_{k+1}$ is the direction of the residual (2.4).

Suppose that the columns of X_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 of $\underline{S}_k$ into $\underline{S}_p$, we have

$$(A-\mu B)^{-1}BX_p = [X_p \;\; x_{k+1}]\underline{S}_p \quad \text{with} \quad \underline{S}_p = \begin{bmatrix} \theta_1 & & \\ & \ddots & \\ & & \theta_p \\ & \beta_k e_k^* Z_p & \end{bmatrix} . \tag{4.1}$$

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 X_p, yields a new recurrence relation

$$(A-\mu B)^{-1}B\tilde{V}_p - \tilde{V}_{p+1}\underline{\tilde{T}}_p = 0 \; . \tag{4.2}$$

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). From here, $k-p$ additional Lanczos iterations produce a Lanczos basis of dimension $k+1$.

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. By moving the corresponding residual terms to the right-hand side, (4.1) becomes

$$(A-\mu B)^{-1}BX_p - [X_p \;\; x_{k+1}] \begin{bmatrix} \Theta_q & 0 \\ 0 & \Theta_{p-q} \\ 0 & \beta_k e_k^* Z_{q-p} \end{bmatrix} = [x_{k+1}\beta_k e_k^* Z_q \;\; 0]$$

with $\Theta_q = \mathrm{diag}(\theta_1, \ldots, \theta_q)$ and $Z_p = [Z_q \;\; Z_{p-q}]$. Transformation to tridiagonal form gives us

$$(A-\mu B)^{-1}B \begin{bmatrix} X_q & \tilde{V}_{p-q} \end{bmatrix} - \begin{bmatrix} X_q & \tilde{V}_{p-q+1} \end{bmatrix} \begin{bmatrix} \Theta_q & 0 \\ 0 & \underline{\tilde{T}}_{p-q} \end{bmatrix} = [R_q \;\; 0] \tag{4.3}$$

with $\underline{\tilde{T}}_{p-q}$ a $p-q+1 \times p-q$ tridiagonal matrix and $R_q = x_{k+1}\beta_k e_k^* Z_q$. The Lanczos method can be continued from this point. Locking Ritz pairs means that they are considered as exact. We then assume the right-hand side of (4.3) zero. Locking introduces an error in the recurrence relation of the order $\|x_{k+1}\beta_k e_k^* Z_q\|_{BF} = |\beta_k| \|e_k^* Z_q\|_2$.

The new tridiagonal matrix $\underline{\tilde{T}}_{p-q}$ can be considered as been computed by the Lanczos method applied to

$$(I - X_q X_q^* B)(A-\mu B)^{-1}B(I - X_q X_q^* B) \; . \tag{4.4}$$

Note that 0 is an eigenpair with multiplicity q and the columns of X_q are eigenvectors. If $(\zeta \neq 0, y)$ is an eigenpair of (4.4), then $X_q^* B y = 0$, so $(I - X_q X_q^* B)(A-\mu B)^{-1}By = \zeta y$ and

$$\begin{aligned} (A-\mu B)^{-1}By - \zeta y &= X_q X_q^* B (A-\mu B)^{-1} By \\ &= X_q(\Theta_q X_q^* B y + R_q^* B y) \; . \end{aligned}$$

As a result,

$$\|(A-\mu B)^{-1}By - \zeta y\|_B = \|X_q R_q^* B y\|_B \leq \|R_q\|_{BF} \; .$$

Since the eigenpairs of $(A-\mu B)^{-1}B$ that still have to be computed are computed as eigenpairs of (4.4), the residual norm of the locked Ritz pairs limits the achievable accuracy.

When (θ, x) is a locked Ritz pair for $(A-\mu B)^{-1}B$, we consider (η, x) with $\eta = \mu + \theta^{-1}$ a locked harmonic Ritz pair for $Ax = \lambda Bx$. When the pole changes to ν, we consider $((\eta-\nu)^{-1}, x)$ a locked Ritz pair for $(A-\nu B)^{-1}B$. When the pole changes, there is a risk that the residual norm grows. Locking Ritz pairs with a relatively large residual norm prevents the accurate computation of other eigenpairs. In order to measure the error on the recurrence relation after locking of the Ritz vectors we can as well use the Rayleigh quotient $\zeta = x^*B(A-\nu B)^{-1}Bx$ instead of $1/(\eta-\nu)$, which gives the smallest residual norm for a given Ritz vector. We do not do this in practice, since this requires additional computational work.

Theorem 4.2 gives an upper bound on the residual norm after change of pole. The theorem uses the following lemma.

Lemma 4.1

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

Proof Let M be such that $B = M^*M$. First note that

$$\begin{aligned}
\|(A-\nu B)^{-1}B\|_B &= \max_{\|x\|_B=1} \|(A-\nu B)^{-1}Bx\|_B \\
&= \max_{\|Mx\|_2=1} \|(M(A-\nu B)^{-1}M^*)(Mx)\|_2 \\
&= \max_{\lambda\in\lambda(A,B)} \left|\frac{1}{\lambda-\nu}\right| \ .
\end{aligned}$$

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

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

which proves the lemma. □

Theorem 4.2 *Let $((\eta-\mu)^{-1}, x)$ be a Ritz pair for $(A-\mu B)^{-1}B$ with residual norm*

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

After the change of pole, the new residual norm satisfies

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

Proof By multiplying $(A-\mu B)^{-1}Bx - \theta x = r$ with $\theta = (\eta-\mu)^{-1}$ 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

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

The proof follows from

$$\|(A-\nu B)^{-1}(A-\mu B)r\|_B \leq \|(A-\nu B)^{-1}(A-\mu B)\|_B \|r\|_B \ .$$

□

Locking Ritz pairs generates an error in the recurrence relation of the order of the residual norms. When the pole changes, this norm changes as well and so does the error on the new recurrence relation. As long as the residual norm is much smaller than the norm of the tridiagonal matrix, there is no real problem. When the pole changes, the norm of the tridiagonal matrix is roughly speaking multiplied by $\|(A-\nu B)^{-1}(A-\mu B)\|_B$, and so it is acceptable that the residuals grow with this factor. However, the factor $|\eta-\mu|/|\eta-\nu|$ should not grow too large. This is not the case for the locked Ritz vectors far away from the new pole. As a conclusion, we better put the new pole away from the locked Ritz pairs. In practice picking the pole near locked Ritz vectors is not very interesting anyway since these have already been computed.

Example 4.1 Recall the problem from Example 3.1. The results of this example were generated using Matlab. We performed $k = 50$ steps of the spectral transformation Lanczos method with $\mu = -1$ and an initial vector with equal elements. The Ritz value with smallest residual norm is $\theta = 0.9976$ or $\eta = \mu + \theta^{-1} = 2.44\,10^{-4}$. The corresponding residual norm is $5.3\,10^{-7}$. Purging the corresponding Ritz pair in the recurrence relation creates an error of the order of $5\,10^{-7}$. Before purging, we have

$$\|E\|_F = \|(A-\mu B)^{-1}BV_k - V_{k+1}\underline{T}_k\|_2 = 0 \ .$$

The four first rows of Table 4.1 show the error when the pole changes after locking θ. The first pole $\nu = 4$ lies to the right of $\lambda(A,B)$. The second pole is the mean of the two smallest Ritz values.

In both cases, the relative error remains of the order of 10^{-6}. In the second case, the absolute error is much larger which means that the small eigenvalues of $(A-\nu B)^{-1}B$ cannot be found very accurately. Note however that the upper bound of the residual norm in Theorem 4.2 is rather pessimistic. For this problem the upper bound following Theorem 4.2 is 1.68, which is much bigger than $1\,10^{-3}$. The reason is that $\|(A-\nu B)^{-1}(A-\mu B)r\|_B$ is much smaller than $\|(A-\nu B)^{-1}(A-\mu B)\|_B\|r\|_B$. The residual r is poor in the dominant eigenvectors of $(A-\nu B)^{-1}(A-\mu B)$ which are well represented by the Krylov space. Also note that, in this case, the Rayleigh quotient ζ corresponds well to $1/(\eta-\nu)$. The reason is that the new residual is already fairly orthogonal to the Ritz vector.

Finally, we illustrate the difference between locking before and after the change of pole. The errors on the recurrence relation are given by Table 4.1. Recall that the locked

Table 4.1: Error in the recurrence relation after locking of a Ritz pair.

ν	4.0	$1.22\,10^{-3}$
Locking before change of pole:		
$\|E\|_F$	$5\,10^{-5}$	$1\,10^{-3}$
$\|E\|_F/\|\underline{K}_k\|_2$	$1\,10^{-6}$	$1\,10^{-6}$
Locking after change of pole:		
$\|E\|_F$	$2\,10^{-7}$	$5\,10^{-1}$
$\|E\|_F/\|\underline{K}_k\|_2$	$5\,10^{-9}$	$5\,10^{-4}$

harmonic Ritz value is $\eta = 2.44\,10^{-4}$. When the Ritz values are recomputed after the change of pole and then the smallest harmonic Ritz value is locked, the errors are smaller when the pole is far away ($\nu = 4$), and larger when the pole is near the locked Ritz value ($\nu = 1.22\,10^{-3}$). The reason is that the Krylov subspace is filtered with the matrix $(A-\mu B)^{-1}(A-\nu B)$ which improves Ritz values near μ and far away from ν but removes those close to ν.

5 The numerical stability of the rational Lanczos method

A remaining question 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 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.

We consider three sources of errors in the rational Lanczos method : the evaluation of the Lanczos recurrence relation ; the change of pole ; and the solution of linear systems with $A-\mu B$ or $A-\nu B$.

The Lanczos recurrence relation The first source of errors is in the evaluation of the Lanczos recurrence relation. Suppose that the spectral transformation Lanczos method produces an accurate recurrence relation, i.e.

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

where E has a small norm compared to the other terms in the equation. We say that the method is backward stable. We assume that when Gram-Schmidt orthogonalization (with reorthogonalization) is used for making $(A-\mu B)^{-1}Bv_j$ orthogonal versus v_{j-1} and

v_j in Algorithm 2.1, we have

$$\begin{aligned} \|E\|_{BF} &\sim \|\underline{T}_k\|_2 \mathbf{u} \\ &\leq \|(A-\mu B)^{-1}B\|_B \mathbf{u} \\ &= \max_{\lambda \in \lambda(A,B)} |\lambda - \mu|^{-1} \mathbf{u} \ . \end{aligned} \tag{5.1}$$

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. Obviously, the relative error, this is $\|E\|_{BF}/\|\underline{T}_k\|_2$ is always small, i.e. the large eigenvalues of $(A - \mu B)^{-1}B$ can always be computed relatively accurately. When partial reorthogonalization is used to maintain orthogonality of the Lanczos vectors, an additional error must be added to the right-hand side corresponding to terms which have been computed during the reorthogonalization. These are typically of the order of $\sqrt{\mathbf{u}}\|\underline{T}_k\|_2$.

Change of pole Suppose that Algorithm 3.1 is executed in exact arithmetic. We do assume, however, an error on Q and R. Following (Higham 1996, Chapter 18), we may assume that the backward error $\|\underline{\Lambda}_k\|_2$ in the QR factorization

$$Q\underline{R} = \underline{L}_k + \underline{\Lambda}_k$$

is bounded from above by $\|\underline{L}_k\|_2 \gamma \mathbf{u}$ where γ depends on k.

Then $Q^*\underline{L}_k R^{-1} = \underline{I}_k - Q^*\underline{\Lambda}_k R^{-1}$, so with $W_{k+1} = V_{k+1}Q$ and $\underline{K}_k = Q^*\underline{T}_k R^{-1}$, we have

$$\begin{aligned} (A-\nu B)^{-1}BW_{k+1}Q^*\underline{L}_k R^{-1} - W_{k+1}\underline{K}_k &= (A-\nu B)^{-1}(A-\mu B)ER^{-1} \\ (A-\nu B)^{-1}BW_k - W_{k+1}\underline{K}_k &= (A-\nu B)^{-1}(A-\mu B)ER^{-1} \\ &+ (A-\nu B)^{-1}BW_{k+1}Q^*\underline{\Lambda}_k R^{-1} \ . \end{aligned} \tag{5.2}$$

The norm of the first term in the right-hand side is roughly speaking bounded by

$$\|(A-\nu B)^{-1}(A-\mu B)\|_B \|R^{-1}\|_2 \|\underline{T}_k\|_2 \mathbf{u} \ , \tag{5.3}$$

where $\|R^{-1}\|_2 = 1/\sigma_{\min}(\underline{L}_k)$. The norm of the second term is roughly speaking bounded by

$$\|(A-\nu B)^{-1}B\|_B \kappa_2(\underline{L}_k)\mathbf{u} \tag{5.4}$$

where $\kappa_2(\underline{L}_k) = \sigma_{\max}(\underline{L}_k)/\sigma_{\min}(\underline{L}_k)$ is the condition number of $\underline{L}_k$. The factors $(A - \nu B)^{-1}(A - \mu B)$ and $(A - \nu B)^{-1}B$ usually have modest B norms unless ν is close to an eigenvalue. The singular values of $\underline{L}_k$ play a vital role.

Lemma 5.1 *Denote the eigenpairs of T_k by (θ_j, z_j) with $\|z_j\|_2 = 1$ for $j = 1, ..., k$ and $\theta_j = (\eta_j - \mu)^{-1}$. Define the residual norm $\rho_j = \|(A-\mu B)^{-1}BV_k z_j - \theta_j V_k z_j\|_B$. Then*

$$\begin{aligned} \sigma_{\min}(\underline{L}_k) &\leq \min_j \left(\left(\frac{\eta_j - \nu}{\eta_j - \mu} \right)^2 + (\mu - \nu)^2 \rho_j^2 \right)^{1/2} \ , \\ \sigma_{\max}(\underline{L}_k) &\geq \max_j \left(\left(\frac{\eta_j - \nu}{\eta_j - \mu} \right)^2 + (\mu - \nu)^2 \rho_j^2 \right)^{1/2} \ . \end{aligned}$$

Proof Following their definitions, we have $\sigma_{\min}(\underline{L}_k) \leq \|\underline{L}_k z\|_2 \leq \sigma_{\max}(\underline{L}_k)$ for any $z \in \boldsymbol{C}^k$ with $\|z\|_2 = 1$. From the assumptions of the lemma, we derive

$$\begin{aligned} \sigma^2_{\min}(\underline{L}_k) \leq \|\underline{L}_k z_j\|_2^2 &= \left\| \begin{bmatrix} (I + (\mu - \nu)T_k)z_j \\ (\mu - \nu)\beta_k e_k^* z_j \end{bmatrix} \right\|_2^2 \\ &= \left(\frac{\eta_j - \nu}{\eta_j - \mu}\right)^2 + (\mu - \nu)^2 \rho_j^2 \ , \end{aligned}$$

and we have a similar inequality for the largest singular value. □

When ν is close to η and the residual norm ρ is small, then $\sigma_{\min}(\underline{L}_k)$ is small. When there is a harmonic Ritz value near μ, $\sigma_{\max}(\underline{L}_k)$ is large. Roughly speaking, when the residual norms are small, the norm of the error in the recurrence relation after change of pole is bounded by

$$\left(\max_{\lambda \in \lambda(A,B)} |\lambda - \mu|^{-1}\right) \cdot \left(\max_{\lambda \in \lambda(A,B)} |\lambda - \nu|^{-1}\right)^2 \ .$$

The solution of linear systems The solution of the linear system $(A - \mu B)w_j = Bv_j$ has a residual f_j so that $f_j = Bv_j - (A - \mu B)w_j$. The backward error F is defined by $(A - \mu B + F)w_j = Bv_j$, so $Fw_j = f_j$ with $\|F\|_\infty \sim \mathbf{u}\|A - \mu B\|_\infty$ (Golub and Van Loan 1996). This implies that the Lanczos method is computing the eigenpairs of $(A - \mu B + F)^{-1}B$ instead of $(A - \mu B)^{-1}B$. When $A - \mu B$ is ill-conditioned, the difference between $(A - \mu B)^{-1}B$ and $(A - \mu B + F)^{-1}B$ can be quite high. We have that

$$(A - \mu B + F)^{-1}B - (A - \mu B)^{-1}B = -(A - \mu B)^{-1}F(A - \mu B + F)^{-1}B$$

which is normwise of the order of $\kappa_2(A - \mu B)\|(A - \mu B + F)^{-1}\|_2\mathbf{u}$. The Lanczos method produces $\underline{T}_k$ and V_{k+1} so that

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

with $\|E\|_{BF}$ small compared to $\|(A - \mu B + F)^{-1}B\|_B$. Assuming that no error is made with the change of pole, we have

$$\begin{aligned} (A - \nu B)^{-1}BW_k - W_{k+1}\underline{K}_k = \\ (A - \nu B)^{-1}(A - \mu B + F)ER^{-1} + (A - \nu B)^{-1}FV_{k+1}\underline{T}_k R^{-1} \ . \end{aligned}$$

The last term is also $(A - \nu B)^{-1}FW_{k+1}\underline{K}_k$ and is of the order of

$$\|(A - \nu B)^{-1}\|_2\|A - \mu B\|_2\|\underline{K}_k\|_2\mathbf{u} \ , \tag{5.5}$$

or roughly speaking,

$$\left(\max_{\lambda \in \lambda(A,B)} |\lambda - \nu|^{-1}\right)^2 \ .$$

Table 5.1: Illustration of growth of rounding errors for Example 5.1. We show bounds for the different errors in the previous analysis

		Bounds		Absolute	Relative
ν	(5.3)	(5.4)	(5.5)	error	error
0.0012	$9.3\,10^{-10}$	$7.5\,10^{-10}$	$1.4\,10^{-11}$	$1.8\,10^{-10}$	$1.7\,10^{-13}$
2.00	$1.6\,10^{-12}$	$1.0\,10^{-12}$	$8.6\,10^{-13}$	$4.6\,10^{-13}$	$3.8\,10^{-14}$

Numerical example The following example illustrates the theory in this section.

Example 5.1 Recall the eigenvalue problem from Example 3.1. The results of this example are generated using Matlab on a DEC Alpha. We performed $k = 50$ steps of the spectral transformation Lanczos method with $\mu = -1$ and an initial vector with equal elements. The explicitly computed error on the recurrence relation is

$$\|E\|_F = \|(A - \mu B)^{-1} B V_k - V_{k+1} \underline{T}_k\|_F \simeq 5.7\,10^{-16}$$

which satifies (5.1). The relative error $\|E\|_F / \|\underline{T}_k\|_2$ is of the same order since $\|\underline{T}_k\|_2 \simeq 1$. Table 5.1 shows the estimates (5.3), (5.4) and (5.5) for the error on the recurrence relation after the change of pole for two choices of ν. First, $\nu = 0.0012$ is the mean of the two smallest harmonic Ritz values, $\eta_1 = 2.4429\,10^{-4}$ and $\eta_2 = 2.1982\,10^{-3}$. Second, $\nu = 2$ lies in the middle of $\lambda(A, B)$. The relative error does not grow in the same way as the absolute error, since $\|\underline{K}_k\|_2$ is usually of the order of $\max_{\lambda(A,B)}(1/|\lambda - \nu|)$. We can see from Table 5.1 that finding the large eigenvalues of $(A - \nu B)^{-1} B$ is no problem whatsoever. The accurate computation of the smaller eigenvalues might be more difficult. It is important to note, though, that both the absolute and relative errors have grown.

The following experiment is an illustration of the danger of picking the pole near an eigenvalue.

Step A We ran $k = 50$ iterations of the Lanczos method with pole $\mu = -1$.

Step B We altered the pole into $\nu = \eta_1$ where η_1 is the smallest harmonic Ritz value. We performed 5 additional Lanczos iterations. We ended up with a Lanczos recurrence relation of order $k = 55$.

Step C We altered the pole again into $\mu = -1$ and performed 5 additional Lanczos steps so that $k = 60$.

Figure 5.1 shows the norms of the errors of the recurrence relations (2.2) for $j = 1, \ldots, k$ for the three steps. As we can see, we have after $k = 50$ Lanczos iterations

$$(A - \mu B)^{-1} B V_{50} - V_{51} \underline{T}_{50} = E_{50}$$

with $\|E_{50}\|_F \simeq 5.5\,10^{-16}$, which satisfies (5.1).

In Step B, the pole is altered into $\nu = \eta_1 \simeq 2.44\,10^{-4}$. We expect a significant growth of the absolute error in the recurrence relation since $\sigma_{\min}(\underline{L}_{50}) \simeq 5.\,10^{-7}$ and

$\kappa_2(\underline{L}_{50}) \simeq 1.6\,10^6$. After having changed the pole, the recurrence relation has an error of the order of $2\,10^2$ (see the peak at $j = 51$). The bound (5.3) is of this order since $\max_{\lambda\in\lambda(A,B)} |\lambda-\mu|/|\lambda-\nu| \simeq 1.5\,10^{12}$ and $\sigma_{\min}(\underline{L}_k) \simeq 5\,10^{-7}$. On the 51st iteration, the error drops to 10^{-4} which is of the order of $\max_{\lambda\in\lambda(A,B)} |\lambda-\nu|^{-1}\mathbf{u}$, and drops even further for $j > 51$.

After the second change of pole, the errors drop but not to the level of $\max_{\lambda\in\lambda(A,B)} |\lambda-\mu|^{-1}\mathbf{u}$. The 5 additional Lanczos iterations, however, produce recurrence relations of that order.

The following observation is also interesting. Suppose that we do not do 5 additional Lanczos iterations in Steps B and C. After Step B, we have (5.2). In Step C, we compute $\underline{L}_k^{(2)} = I + (\nu-\mu)\underline{K}_k = Z\underline{S}$ with Z unitary and $\underline{S}$ upper triangular. Since for Step B $\underline{K}_k = \frac{1}{\mu-\nu}\left(I - Q^* \begin{bmatrix} R^{-1} \\ 0 \end{bmatrix}\right)$, we find that $Z = Q^*$ and $\underline{S} = \begin{bmatrix} R^{-1} \\ 0 \end{bmatrix}$. So, after Step C, we get

$$(A-\mu B)^{-1}BW_k\underline{Z} - W_{k+1}ZZ^*\underline{K}_kS^{-1} = E + (A-\mu B)^{-1}BW_{k+1}Q^*\underline{\Lambda}_k \ .$$

Note that the two terms in the right-hand side have small norms of the order of $\mathbf{u}$. If Z and $\underline{S}$ are *computed* from $\underline{L}_k^{(2)}$, an error in the recurrence relation is introduced that can be of the order of $\mathbf{u}\kappa_2(\underline{L}_k^{(2)}) \simeq 1.6\,10^6\mathbf{u}$. Note that $\kappa_2(\underline{L}_{50}^{(2)}) = \kappa_2(\underline{L}_{50})$. In fact, in exact arithmetic $\underline{L}_{50}^{(2)}$ would be the generalized inverse of $\underline{L}_{50}$. In practice, however, we do not at all lose 6 digits but only 2 since

$$\|(A-\mu B)^{-1}BW_k\underline{Z} - W_{k+1}ZZ^*\underline{K}_kS^{-1}\|_F \simeq 2.5\,10^{-14} \ .$$

From this example, we conclude that picking a pole near a Ritz value may be dangerous. But picking a pole when the old pole is near a Ritz value may be dangerous as well.

We also conclude that in a practical algorithm it is sufficient to monitor $\sigma_{\min}(\underline{L}_k)$. If at the first change of pole, $\sigma_{\min}(\underline{L}_k^{(1)})$ is small, we may introduce a large absolute error in the recurrence relation. In the next change of pole this may lead to a large $\sigma_{\max}(\underline{L}_k^{(2)})$ and again large absolute errors may result.

6 Numerical example

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 fluid displacement components (weighted by the local porosity) (w) are selected as nodal variables. The finite-element

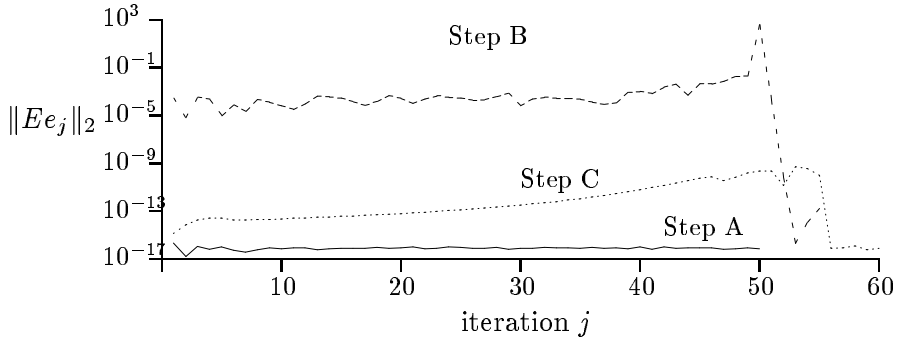

Figure 5.1: Norms of the errors of the recurrence relation (2.2). The solid line shows $\|Ee_j\|_2$ for $j = 1, \ldots, 50$ for the Lanczos method with pole $\mu = -1$ (Step A). The dashed line shows the errors after the change of pole for $j = 1, \ldots, 50$ and for 5 additional Lanczos iterations $j = 51, \ldots, 55$ (Step B). The dotted line shows the error on the recurrence relation when the pole is changed back into $\mu = -1$ and 5 additional Lanczos iterations are performed (Step C).

mesh has 324 nodes and 192 HEXA8 elements. The total number of degrees of freedom is $n = 1944$.

We use the Lanczos routine EA16 from the Harwell Subroutine Library (Release 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 6.1. 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 $\epsilon_{\rm tol} = \sqrt{\mathbf{u}}\|T_k\|_2$ with $\sqrt{\mathbf{u}} \simeq 1.05\,10^{-8}$, which is the square root of the machine precision on a DEC Alpha. The goal is to compute $s = 20$ eigenpairs on the right of 100. Ritz pairs are locked before the pole is changed. Recall the results of Example 4.1 where the accuracy of the recurrence relation after locking and changing the pole depends on whether the harmonic Ritz values are locked before or after the change of pole. In this code, no distinction is made between η_j's away from or close to ν : we always lock before the change of pole. The initial pole is $\mu = 100$. A new pole is chosen as the mean of two successive Ritz values, as indicated in Algorithm 6.1. We choose the new pole so that $|\eta_j - \nu| \geq |\eta_j - \mu|/50$ for $j = 1, \ldots, k$, which should avoid large $\kappa_2(\underline{L}_k)$ and $(A - \nu B)^{-1}B$.

Algorithm 6.1 (Rational Lanczos)
0. Given $v_1 \in \boldsymbol{R}^n$ such that $\|v_1\|_B = 1$.
   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 $\eta_j = \mu + \theta_j^{-1}$.

2.3. Sort the η_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 $\rho_j \leq \epsilon_{\text{tol}} \|T_k\|_2$ 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 $\nu = (\eta_j + \eta_{j+1})/2$ with $j \geq q$ and such that $|\eta_i - \nu| \geq |\eta_i - \mu|/50$ for $i = 1, \ldots, p$.

2.8. Change the pole using Algorithm 3.1. Factorize $A - \nu B = LDL^*$

end do

From (5.2), we can see that $\|(A - \nu B)^{-1}(A - \mu B)\|_B$ plays an important role in the error bounds. When $|\lambda - \nu| \geq |\lambda - \mu|/50$ for all eigenvalues λ of $Ax = \lambda Bx$, then

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

This limits the growth of the error in the recurrence relation after a change of pole. Lemma 5.1 is used for estimating $\sigma_{\min}(\underline{L}_p)$. This was illustrated by Example 5.1 and is confirmed again by the results for this example. Table 6.1 shows, after each implicit restart, the number of converged Ritz values and the new pole with the estimate for $\sigma_{\min}(\underline{L}_p)$. A picture of the spectrum is shown in Figure 6.1. 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.

The error on harmonic Ritz values η is bounded as follows.

Lemma 6.1 *Let η be a harmonic Ritz value corresponding to pole μ and let ρ be the residual norm for the Ritz pair of the spectral transformation. Then there is a $\lambda \in \lambda(A, B)$ so that*

$$|\lambda - \eta| \leq |\eta - \mu|^2 \rho \ .$$

Proof The proof is given in §3 of Ericsson and Ruhe (1980). □

Figure 6.2 shows the bounds on the errors on the harmonic Ritz values (Lemma 6.1) when the pole changes. We show the results for the unlocked Ritz values in the interval $[1000, 4500]$ for the first two changes of pole. In both cases, the bounds do not change very much with a change of pole, nor do the harmonic Ritz values. In addition, the harmonic Ritz values near 4000 seem to disappear after the change of pole.

After the last iteration, we found that, with full reorthogonalization at each iteration of the Lanczos process, $\max_{i,j \neq i} |w_i^* B w_j| \simeq 1\, 10^{-14}$ and

$$\|(A - \mu B)^{-1} B W_k - W_{k+1} \underline{K}_k\|_{BF} / \|\underline{K}_k\|_2 \simeq 4.9\, 10^{-8} \ ,$$

which is within the residual tolerance. With partial reorthogonalization, we have $\max_{i,j \neq i} |w_i^* B w_j| \simeq 3.5\, 10^{-11}$ and

$$\|(A - \mu B)^{-1} B W_k - W_{k+1} \underline{K}_k\|_{BF} / \|\underline{K}_k\|_2 \simeq 4.9\, 10^{-8} \ ,$$

Table 6.1: Number of locked Ritz values, the new pole and the estimate of $\sigma_{\min}(\underline{L}_p)$ following Lemma 5.1

restart	locked	pole	estimate $\sigma_{\min}(\underline{L}_p)$
1	5	1381.	$8.3\,10^{-2}$
2	12	1764.	$1.8\,10^{-1}$
3	17	1953.	$1.3\,10^{-1}$
4	25		

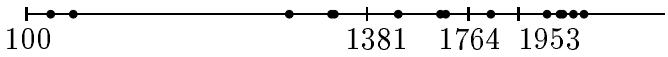

Figure 6.1: Ritz values computed via the rational Lanczos method. The vertical lines denote the positions of the poles.

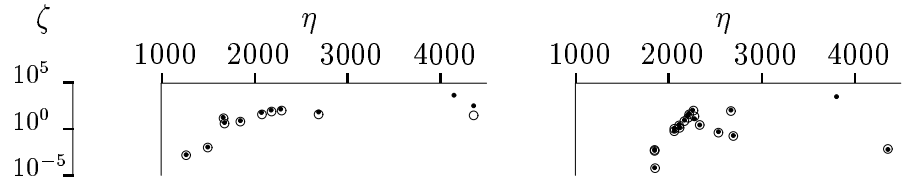

Figure 6.2: The pictures show the harmonic Ritz value η and the harmonic residual norms from Lemma 6.1 before (bullets) and after (circles) the change of pole. The left hand picture shows the transition for the change of $\mu = 100$ into $\nu = 1381$ and the right hand one for the change from $\mu = 1381$ into $\nu = 1764$.

We also performed the following experiment. We ran algorithm 6.1 but in Step 2.7, we select the pole $\nu = \eta_{q+1} \simeq 1264.112$. Before the change of pole, we have with partial reorthogonalization

$$\begin{aligned} \|(A-\mu B)^{-1}BV_k - V_{k+1}\underline{T}_k\|_{BF} &\simeq 1.2\,10^{-12} \\ \|(A-\mu B)^{-1}BV_k - V_{k+1}\underline{T}_k\|_{BF}/\|\underline{T}_k\|_F &\simeq 3.7\,10^{-11}\ . \end{aligned}$$

Note that $\sigma_{\min}(\underline{L}_k) \simeq 1.4\,10^{-6}$, $\kappa_2(\underline{L}_k) \simeq 9\,10^6$ and $\|(A-\nu B)^{-1}(A-\mu B)\|_B \geq \|I + (\nu - \mu)\underline{K}_k\|_2 \simeq 7.1\,10^5$. We expect the Lanczos recurrence relation to lose 12 digits after the change of pole. Since $\|\underline{K}_k\|_F$ is of the order of 10^3, we expect to keep only 1 significant digit. The computed error, however, is much larger since

$$\begin{aligned} \|(A-\nu B)^{-1}BW_k - W_{k+1}\underline{K}_k\|_{BF} &\simeq 1.2\,10^{3} \\ \|(A-\nu B)^{-1}BW_k - W_{k+1}\underline{K}_k\|_{BF}/\|\underline{K}_k\|_F &\simeq 2.0\,10^{0}\ . \end{aligned}$$

We keep no accuracy at all. For computing the term $(A-\nu B)^{-1}BW_k$, we used iterative refinement using GMRES in order to reduce the residual norm of the linear systems. (The reason why we used GMRES and not MINRES is that we had a code available for the former method. We used GMRES for the verification of the recurrence relation only, so performance is not an issue here.) The major problem here is that even with a small residual norm the linear systems are solved fairly inaccurately, since $\kappa_2(A-\nu B) \simeq 2.5\,10^{21}$. Before the change of pole, the Ritz value θ_{q+1} has a residual norm of the order of $1.35\,10^{-8}$. The Ritz value appears to be much closer to an eigenvalue than the residual norm suggests and $A-\nu B$ is almost singular. Continuing the Lanczos method with this pole is not a good idea since each connection with the original problem is lost.

If we do continue with this pole and use $\nu = (\eta_q + \eta_{q+1})/2 \simeq 1381$ as next pole, so that $\sigma_{\min}(\underline{L}_k) \simeq 2.0$, then we find

$$\begin{aligned} \|(A-\nu B)^{-1}BW_k - W_{k+1}\underline{K}_k\|_{BF} &\simeq 1.3\,10^{-3} \\ \|(A-\nu B)^{-1}BW_k - W_{k+1}\underline{K}_k\|_{BF}/\|\underline{K}_k\|_F &\simeq 6.5\,10^{-1}\ , \end{aligned}$$

which is very large. The connection with the original problem $Ax = \lambda Bx$ is completely lost. The only way to save the situation is restart the Lanczos method explicitly with the new pole $\nu \simeq 1381$.

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 for a Hermitian 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 reasons. 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. It is therefore advised to pick the pole away from the locked Ritz pairs.

Moving the pole near a Ritz value may not only lead to an increase of the error after the current change of pole, but may propagate to the next one. As Example 4.1 illustrates, the residual norm of a locked Ritz pair may decrease when the harmonic Ritz value is away from the new pole and may grow may it is close to the new pole. We have made no distinction between locked harmonic Ritz values near or away from the new pole. The major reason is to reduce the complexity of the software. In addition, since the poles are chosen away from the harmonic Ritz values, a significant decrease or growth of residual norm is rather unlikely.

One can pick poles near harmonic Ritz values, but we then suggest an explicit restart of the Lanczos method with the new pole after locking the converged Ritz pairs. Note that picking poles close to a harmonic Ritz value may speed up the convergence of one or a few Ritz pairs, but a larger number of poles may be required to compute the desired number of eigenpairs fairly accurately. From an efficiency point of view one should keep the number of pole changes (i.e. the number of sparse matrix factorizations) small.

Acknowledgements

The author is grateful to Jennifer Scott, Iain Duff, and Axel Ruhe for helpful suggestions that improved the quality of the paper. The author also wishes to thank the referees for their interesting comments that improved the quality of the paper.

References

Arnoldi, W. (1951), 'The principle of minimized iterations in the solution of the matrix eigenvalue problem', *Quart. Appl. Math.* **9**, 17–29.

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 Philadelphia.

HSL (1996), 'Harwell Subroutine Library. A Catalogue of Subroutines (Release 12)'. Information: Nick Brealey, AEA Technology, 477 Harwell, Didcot, Oxon, OX11 0RA, UK, E-mail: hsl@aeat.co.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. (1996), Restarting an Arnoldi reduction, Preprint MCS-P591-0496, Argonne National Laboratory, Argonne, IL 60439.

Lehoucq, R. (1999), On the convergence of an implicitly restarted Arnoldi method, Technical Report SAND99-1756J, Sandia National Laboratories, MS 1110, P.O. Box 5800, Albuquerque, NM 87185-1110.

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 Spence, A. (1997), 'Implicitly restarted Arnoldi and purification for the shift-invert transformation', *Math. Comp.* **66**, 667–689.

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.

Paige, C., Parlett, B. and van der Vorst, H. (1995), 'Approximate solutions and eigenvalue bounds from Krylov subspaces', *Num. Lin. Alg. Appl.* **2**, 115–133.

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 A.1

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. The Householder transformation is also applied on the front so that the $j+1$ st column is reduced to a column vector with $j-1$ zeros in the front.