Chapter VII

Arnoldi's Method

28 The computation of eigenpairs

As we have seen in Chapter V, the construction of a basis with Arnoldi's method for the Krylov subspace $\mathcal{K}^m(A; v_1)$ leads to an upper Hessenberg matrix that describes the relation between the basis vectors. Each new basis vector can be constructed from the existing set as

$$h_{i+1,i}v_{i+1} = Av_i - \sum_{j=1}^{i} h_{j,i}v_j, \tag{28.1}$$

where the $h_{j,i}$ follow from the orthogonality requirement $v_j^* v_{i+1} = 0$, and $h_{i+1,i}$ follows from the requirement that $\|v_{i+1}\|_2 = 1$. The above formulation, if carried out in numerical computation in this form, is a classical Gram-Schmidt formulation; in actual computation the orthogonalization is rather carried out with modified Gram-Schmidt in order to improve the orthogonality of the basis. In exact arithmetic it leads to the same coefficients and the expressions (28.1) can be formulated in matrix notation as

$$AV_m = V_{m+1}H_{m+1,m}, \tag{28.2}$$

where $H_{m+1,m}$ is an upper Hessenberg matrix.

In the previous chapter we have seen that, when A is Hermitian, interesting relations between eigenvalues of A and those of the often much smaller matrix $H_{m,m}$ could be derived. In principle, we can follow similar approaches for unsymmetric matrices, but the situation for non-Hermitian matrices is more complicated. For instance, there is no minimax characterization and that leads to less powerful tools.

Because of the orthogonality we have that

$$V_m^T AV_m = H_{m,m}. \tag{28.3}$$

The matrix $H_{m,m}$ is the projection of A onto $K^m(A; v_1)$.
If (θ, y) is an eigenpair of $H_{m,m}$ then

$$H_{m,m} y = \theta y$$

$$V_m^T A V_m y - \theta V_m^T V_m y = 0$$
$$V_m^T (A V_m y - \theta V_m y) = 0$$
$$V_m^T (As - \theta s) = 0,$$

where $s = V_m y$.
Hence, the residual for the approximate eigenpair (θ, s) is orthogonal to the current Krylov subspace. As before, the value θ is called a *Ritz value* of A with respect to $K^m(A; v_1)$, and s is the corresponding *Ritz vector*.

29 Numerical examples

We will now discuss aspects of the convergence behavior of the Arnoldi method, guided by observations from some numerical examples.

Example 1 Convergence with the Arnoldi method is much more complicated than with the Lanczos method, and in fact we have only the orthogonality relations to rely on. The following famous example, taken from SAAD and SCHULTZ [1986], shows the problem quite clearly.

Consider the matrix

$$A = [e_2|e_3|\cdots|e_n|e_1],$$

that is, the unity matrix in which the first column is removed and replaced as the last column. A is an orthonormal matrix with all its eigenvalues uniformly distributed over the unit circle in the complex plane.

However, when we form the Krylov subspace with initial vector $v_1 = e_1$, then clearly $Av_1 = e_2$, and we see that these unit vectors form an orthonormal basis for the Krylov subspace without any further orthogonalization. This implies that $h_{j,i} = 0$ for $j \leq i$ in relation (28.1).

As a consequence, the Hessenberg matrix associated with this Krylov subspace of dimension m is the leading minor of A. For all $m < n$ these matrices are largely filled with zeros, except for the ones on the sub-codiagonal. Obviously all Ritz values are equal to zero. Not only is there no separation property as in the symmetric case, we even have a severe case of stagnation. This is nothing similar as the near-stagnation that we have observed with

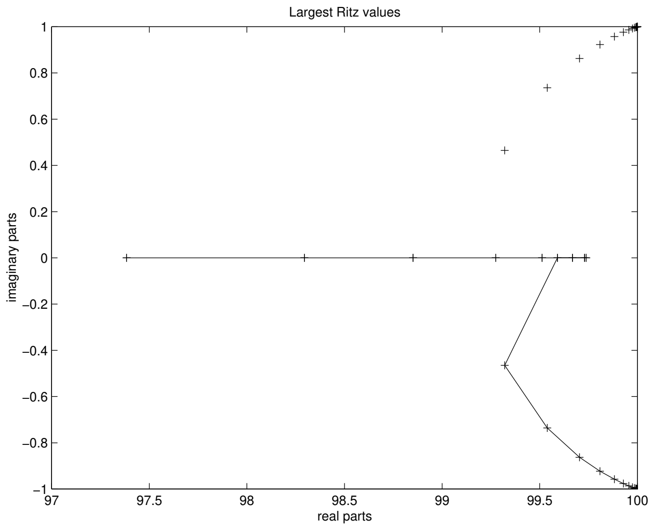

Figure 17: Largest Ritz values for $m = 5 : 30$

almost double eigenvalues in the symmetric case. All eigenvalues are nicely distributed over the unit circle and the condition of the eigenvector matrix is 1.

Note that because of the shift invariance of the Krylov subspace, the same may happen if the eigenvalues are located on a circle somewhere in the complex plane. Note also that stagnating convergence is not necessarily related to ill-conditioning of the eigenvector matrix of A. In principle, there is nothing to prevent such cases of misconvergence and the main trick to avoid this is to use the shift-and-invert approach with properly chosen shifts.

For some situations one can prove that this severe misconvergence cannot take place, for instance, when the exterior (in modulus) eigenvalues of A are well separated from the others and if the matrix is not too nonnormal. This is the same condition for which the Power Method converges.

Example 2 With the Lanczos method we observed cases of almost stagnating convergence in the presence of close eigenvalues. Likewise, one may expect similar stagnating effects if there are clusters of eigenvalues in the spectrum. In fact, as we have seen in the previous example, the eigenvalues need not even be close to each other.

This example shows that we may also see strange convergence effects when there are complex conjugate eigenvalues: $\lambda = a \pm b\imath$. This might be expected to cause some problems when $|b|$ is small, but even for relatively large b we may see strange convergence patterns of the Ritz values. In particular, when we use a real starting vector, the Ritz values may initially converge to a value on the real axis, before they split and set off to the members of the conjugate eigenpair.

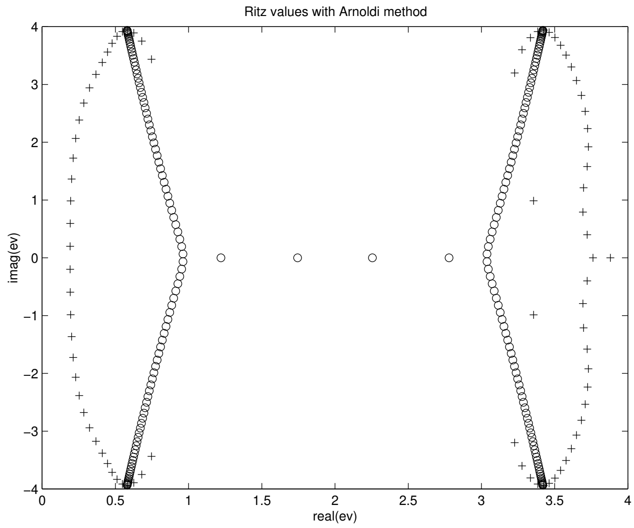

Figure 18: Ritz values for $m = 80$ (+), exact eigenvalues (o)

In Figure 17, we have plotted the largest Ritz value for $m = 5$ to $m = 30$, for a matrix with eigenvalues $1, 2, \ldots, 98$, and $100 + \imath$, and $100 - \imath$. If the largest Ritz value is complex, then we have also plotted the conjugate one. Furthermore, we have connected the largest Ritz values with negative imaginary part, in order to show the pattern of convergence.

For $m = 5$, the largest Ritz value has arrived at the left of the drawn line, and for increasing m, it moves to the right until it arrives, for $m = 11$ at the endpoint of this line on the x-axis. Then the largest Ritz value moves backwards and for $m \geq 13$ there are two complex conjugate Ritz values. We see, that from this phase on we get more or less regular convergence to the complex conjugate eigenvalues $100 \pm \imath$.

We can aggrevate this kind of convergence by making b smaller; this may lead to situations where it seems that the Ritz values are heading for some nearby point on the real axis, stagnate there for some values of m, and then suddenly set off for the complex conjugate eigenvalues. This type of convergence is quite similar to the situation for close (real) eigenvalues in the Hermitian case.

Example 3 We will describe some other effects for unsymmetric matrices, that shows that the convergence behavior for Krylov methods may be totally different from that what one experiences for symmetric matrices (with the Lanczos method). First we consider the matrix of order 200 of the following form:

$$A = \begin{bmatrix} 2 & 1 & -2 & & & \\ 1 & 2 & 1 & -2 & & \\ 2 & 1 & 2 & 1 & 2 & \\ & \ddots & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & & \\ & & & & & \end{bmatrix}. \qquad (29.1)$$

This matrix is nonnormal, the condition number of the eigenvector matrix is about 10^{10}. In Fig. 18 we see the exact eigenvalues, represented by small circles. The four eigenvalues in the middle of the spectrum are quite well separated from the others. If we carry out as many as 80 steps with Arnoldi, with a starting vector v with all ones, then we obtain the Ritz values represented by the $+$ signs.

We see that the interior isolated eigenvalues have not been discovered at all for $m = 80$ (for $m = n = 200$, they are well approximated). We also see some strange behavior for the Ritz values that are close to the two branches of eigenvalues at the right and left of the spectrum.

Example 4 We have aggrevated the situation slightly for the next example.

We now choose our matrix as

$$A = \begin{bmatrix} 2 & 1 & -.4 & & & & \\ 0 & 2 & 1 & -.4 & & & \\ 2 & 0 & 2 & 1 & -.4 & & \\ & \ddots & \ddots & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & & & \\ & & & & & & \end{bmatrix}.$$

The eigenvalues computed with Matlab for this 200 by 200 matrix have been indicated by small circles in Fig. 19.

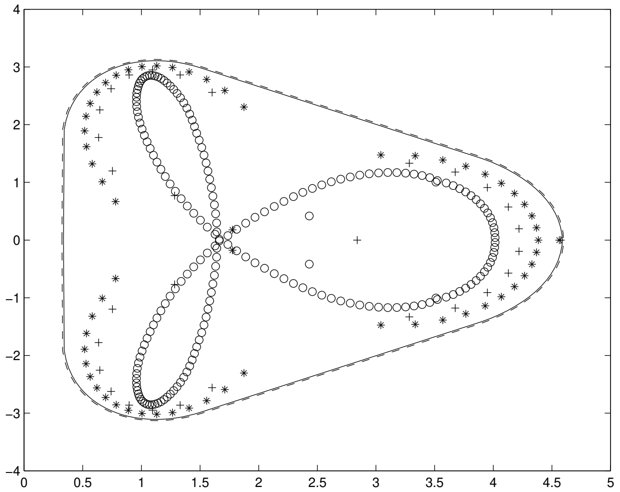

Figure 19: Ritz values for $m = 30$ $(+)$, $m = 60$ $(*)$, and exact eigenvalues (o).

Again, the starting vector for Arnoldi is the vector with all ones. We have indicated the Ritz values for $m = 30$ by a $+$, and for $m = 60$ by a $*$. We see that the Ritz values are located more or less on curves that follow the curves for the Matlab eigenvalues. If one would not have known these eigenvalues, then one might have expected from the trend in the Ritz values, when going from $m = 30$ to $m = 60$, that their limit values ar still located slightly more outside. Only for rather high values of m, relative to the order $n = 200$ of A, the Ritz values are located closer to the Matlab eigenvalues.

The *field of values* of a matrix is the set of all possible values of the Rayleigh quotient $(Ax, x)/(x, x)$, for all $x \in \mathbb{C}^n$. In Figure 19 we have indicated the field of values with the dashed curve. The field of values for $H_{30,30}$ is shown by the drawn curve. We see that the Arnoldi method leads to a fairly good estimate for the field of values; in this case much better than it does represent the eigenvalues. This phenomenon was observed by Toh and Trefethen [1996]. The field of values plays a role in some error bounds for the Arnoldi process, also in the context of solving linear systems, see Manteuffel [1977] and Embree [1999].

The seemingly strange convergence behavior of the Ritz values is not a side-effect of computational errors in the orthogonalization process for the basis vectors for the Krylov subspace. If we repeat this basis construction with iterated Gram-Schmidt, as in Algorithm 6, then we obtain in this case virtually the same results.

We have included this example as a warning for the interpretation of eigenvalue computations and the usage of 'nice' test examples. Note that we have avoided to use the notion of exact eigenvalues, instead we referred to the computed eigenvalues of A as the Matlab eigenvalues. These eigenvalues follow a very nice pattern and may seem very suspect for that reason for an unexperienced person.

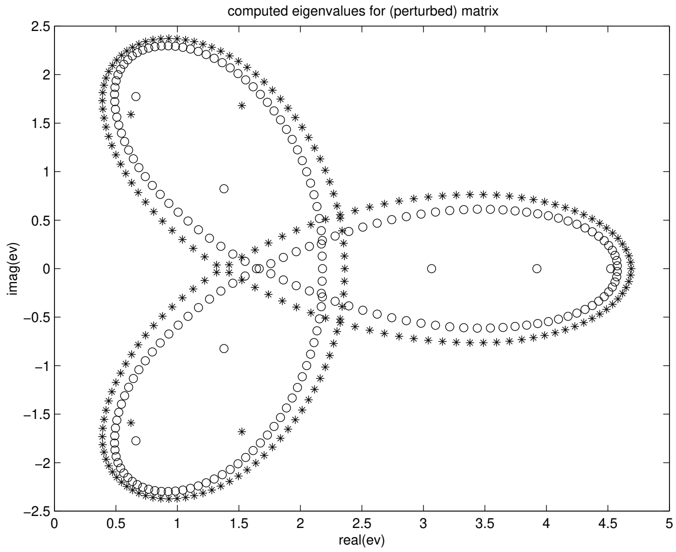

Figure 20: computed eigenvalues of A (o), and of a randomly perturbed A ($*$), with perturbations in the order of 10^{-12}

A warning might have been that the condition number of A's eigenvector matrix (computed with Matlab for smaller values of n and extrapolated to $n = 200$) is $\approx E46$. This indicates that small perturbations, even much smaller then in the order of rounding errors, may lead to noticeable perturbations to the computed eigenvalues. Indeed, this is the case here.

In Fig. 20 we have plotted the computed eigenvalues of A, indicated with o, and those of $A + B$, where $B = 10^{-12}\mathrm{rand}(200, 200)$ (rand is the Matlab random function). The computed eigenvalues of $A + B$ have been indicated by a $*$. It appears that the computed eigenvalues coincide more or less with

ϵ-level curves of the pseudospectra of A, that is the set of values z for which

$$\|(zI - A)^{-1}\|_2 = \epsilon.$$

The level curves in the pseudospectra reveal the sensitivity of the eigenvalues of A. In this particular case, they show that almost all eigenvalues are extremely sensitive for perturbations, except possible a few in the middle of the picture. See Chapter XII for further details on pseudospectra.

These examples illustrate that one should not easily trust the Ritz values for rather low values of m for nonnormal matrices. It might be helpful to compute condition numbers for the Ritz values of interest, but this requires also the computation of left eigenvectors.

30 Shift-and-invert

In practical situations subspace methods are often applied with $(A - \sigma I)^{-1}$, so that convergence towards eigenvalues close to σ is favored. This is called the *shift-and-invert* approach. This technique is attractive if one can compute vectors $(A - \sigma I)^{-1} z$ at relatively low costs. Note that with a direct solver this can be done with an LU decomposition of $A - \sigma I$, and this LU decomposition can be used for subsequent iteration steps.

Since the Krylov subspace methods work with vectors and the matrix is never needed explicitly, the procedure for shift-and-invert Arnoldi runs as follows. The eigenvalues μ_i of $(A - \sigma I)^{-1}$ are related to the eigenvalues λ_i of A as

$$\lambda_i = \frac{1}{\mu_i} + \sigma.$$

Suppose we need to compute a vector $y = (A - \sigma I)^{-1} z$ for the expansion of the Krylov subspace. Then we decompose $A - \sigma I = LU$. The vector y is computed in two steps. First solve w from $Lw = z$ and y is solved from $Uy = w$. Note that we have to compute the LU decomposition only once, when we work with a fixed shift σ. This means that the cost of an LU factorization can be amortized over the Arnoldi iterations.

We now repeat the Example 3, described in Section 29, with shift-and-invert. With a fixed shift $\sigma = 2.0$, and 15 Arnoldi iterations, we nicely recover two of the interior eigenvalues of the spectrum, as is shown in Fig. 21. This is in sharp contrast with what we have observed for much larger values of m with regular Arnoldi, see Figure 18. If one wishes to explore other parts of the spectrum, then one has to select proper shifts.

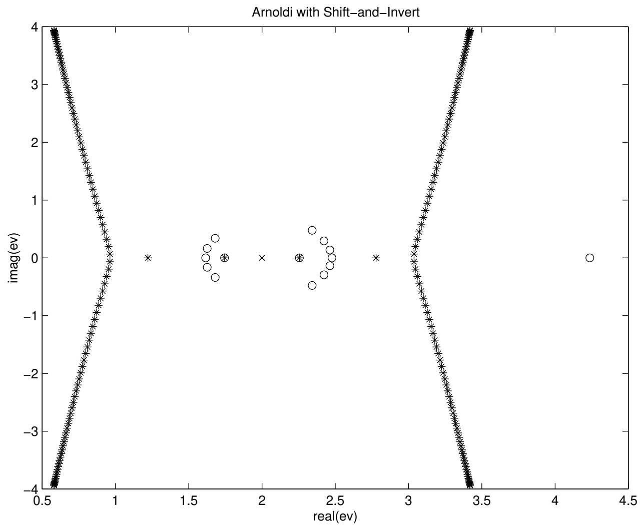

Figure 21: Ritz values for $m = 15$ (o), exact eigenvalues ($*$), and shift $s = 2$ ($\times$)

31 Implicitly Restarted Arnoldi

The main problem with Arnoldi's method is that it becomes increasingly expensive per iteration step. In order to restrict memory storage, as well as computational work, restarts are necessary. However, at a restart we do not want to throw away useful information. Instead of restarting with the most current appoximation for the approximated eigenvector (the Ritz vector), corresponding to the desired eigenvalue, it has been suggested to restart with a vector that is a mix of Ritz vectors, corresponding to relevant Ritz values. See Saad [1992, Chapter VII] and Petiton [1992] for such strategies. These strategies are related to an earlier approach suggested in Golub and Underwood [1977].

These strategies may have the advantage that the new starting vector also contains information for nearby eigenvalues. The main problem is that in this kind of restart we try to catch the information for an approximate subspace in one single vector, and apart from this, it is not easy to find the optimal mix.

The restart problem has been solved very elegantly by Sorensen [1992]. The idea behind his *Implicitly Restarted Arnoldi* (IRA) method is the following. Suppose we are at step $k+m$, and we want to shrink the current subspace to a subspace of dimension k again. Obviously we want to maintain the subspace spanned by the best k Ritz vectors, that is, the Ritz vectors corresponding to k Ritz values that we have selected, for instance the k rightmost Ritz values if we want to approximate the k rightmost eigenvalues of A.

We compute all $k+m$ Ritz values of H_{k+m} and then we carry out m shifted QR iterations with the m undesired Ritz values as shifts. This has the effect that the k wanted Ritz values are contained in the leading k by k part of the transformed matrix. Also the projection matrix V_{m+k} is transformed correspondingly to the transformations on H_{m+k}, and after the entire operation the first k columns of the transformed V_{m+k} span the subspace of the Ritz vectors for the wanted Ritz values.

We will show this in more detail. After $k+m$ Arnoldi steps, we have arrived at the following situation

$$\begin{aligned} AV_{k+m} &= V_{k+m+1}H_{k+m+1,k+m} \\ &= [V_{k+m}, v_{k+m+1}] \begin{bmatrix} H_{k+m,k+m} \\ h_{k+m+1,k+m}\, e_{k+m}^T \end{bmatrix}. \end{aligned} \tag{31.1}$$

We will write $r = h_{k+m+1,k+m}\, v_{k+m+1}$ for short.

With one QR step for the upper Hessenberg matrix $H_{k+m,k+m}$ with shift θ, we obtain

$$H_{k+m,k+m} - \theta I = Q_1 R_1, \tag{31.2}$$

where Q_1 is orthogonal and R_1 is upper triangular. See Chapter 10 for relevant details on the QR method.

Now we try to incorporate this step in the Arnoldi relations (31.1):

$$(A-\theta I)V_{k+m} - V_{k+m}(H_{k+m,k+m} - \theta I) = re_{k+m}^T \tag{31.3}$$

$$(A-\theta I)V_{k+m} - V_{k+m}Q_1R_1 = re_{k+m}^T. \tag{31.4}$$

As in the QR method, we apply these matrices again to Q_1, which gives

$$(A-\theta I)(V_{k+m}Q_1) - (V_{k+m}Q_1)R_1Q_1 = re_{k+m}^T Q_1 \tag{31.5}$$

$$A(V_{k+m}Q_1) - (V_{k+m}Q_1)(R_1Q_1 + \theta I) = re_{k+m}^T Q_1. \tag{31.6}$$

Precisely as in the QR iteration the matrix $R_1Q_1 + \theta I$ is again an upper Hessenberg matrix (remember that Q_1 is upper Hessenberg, see equation (31.2)).

Because the co-diagonal elements $h_{j+1,j}$ of $H_{k+m,k+m}$ are nonzero, we have, just as in the QR iteration, that the bottom diagonal element of R_1

equals 0 if θ is an eigenvalue of $H_{k+m,k+m}$. Likewise, if θ is close to an eigenvalue then this last element will be small. This plays a role in the strategy for the selection of the shifts. Of course, the shifts are carried out implicitly, by a bulge chasing algorithm, as explained in Section 11.

We will denote $H^{(1)}_{k+m,k+m} := R_1 Q_1 + \theta I = Q_1^T H_{k+m,k+m} Q_1$. Note that $V_{k+m} Q_1$ is orthogonal again and spans the same subspace. Equation (31.5) can be reformulated as

$$A(V_{k+m}Q_1) = (V_{k+m}Q_1, v_{k+m+1}) \begin{bmatrix} H^{(1)}_{k+m,k+m} \\ h_{k+m+1,k+m}\, e^T_{k+m} Q_1 \end{bmatrix}. \tag{31.7}$$

With another shift $\widetilde{\theta} \neq \theta$, we have $H^{(1)}_{k+m,k+m} - \widetilde{\theta} I = Q_2 R_2$ and define

$$H^{(2)}_{k+m,k+m} := R_2 Q_2 + \widetilde{\theta} I.$$

After multiplying (31.7) from the right with Q_2, we obtain

$$A\left(V_{k+m}Q_1Q_2\right) = \left(V_{k+m}Q_1Q_2, v_{k+m+1}\right) \begin{bmatrix} H^{(2)}_{k+m,k+m} \\ h_{k+m+1,k+m}\, e^T_{k+m} Q_1 Q_2 \end{bmatrix}.$$

Proceeding in the same manner with another $m-2$ well chosen shifts, we arrive at

$$A(V_{k+m}Q) = (V_{k+m}Q, v_{k+m+1}) \begin{bmatrix} H^{(m)}_{k+m,k+m} \\ h_{k+m+1,k+m}\, e^T_{k+m} Q \end{bmatrix}, \tag{31.8}$$

where $Q = Q_1 Q_2 \cdots Q_m$. Note that each of the matrices Q_i is $k+m$ by $k+m$ upper Hessenberg. This implies that the zero leading part of the vector e^T_{k+m} shrinks with one element after one multiplication with a Q_j. After m of these multiplications there are only $k-1$ zero leading elements left in de vector $e^T_{k+m}Q$, and we can write

$$h_{k+m+1,k+m}\, e^T_{k+m} Q = (\beta e^T_k, b^T), \tag{31.9}$$

where $b \in \mathbb{R}^m$.

Note that the first k columns of the matrix

$$\begin{bmatrix} H^{(m)}_{k+m,k+m} \\ h_{k+m+1,k+m}\, e^T_{k+m} Q \end{bmatrix}$$

still are in perfect upper Hessenberg form. We denote $\widetilde{V}_{k+m} = V_{k+m}Q$. Then we obtain the following equation for the first k columns of $\widetilde{V}_{k+m}$, with $\widetilde{v}_j$ the j-th column of $\widetilde{V}_{k+m}$:

$$\begin{aligned} A\widetilde{V}_k &= \widetilde{V}_k H^{(m)}_{k,k} + h^{(m)}_{k+1,k}\widetilde{v}_{k+1}e_k^T + \beta v_{k+m+1}e_k^T \\ &= \widetilde{V}_k H^{(m)}_{k,k} + \widetilde{h}_{k+1,k} s e_k^T, \end{aligned} \tag{31.10}$$

where s has been normalized: $\|s\|_2 = 1$.

We now define the upper Hessenberg matrix $\widetilde{H}_{k+1,k}$ as

$$\widetilde{H}_{k+1,k} = \left[\begin{array}{c} H^{(m)}_{k,k} \\ \widetilde{h}_{k+1,k}e_k^T \end{array} \right], \tag{31.11}$$

and we have arrived at

$$A\widetilde{V}_k = (\widetilde{V}_k, s)\widetilde{H}_{k+1,k}. \tag{31.12}$$

We see that this describes precisely the Arnoldi relations for the columns of $\widetilde{V}_k$, and expansion vector s (note that s is indeed orthogonal to the columns of $\widetilde{V}_k$). All these vectors are available after the m implicit QR steps.

What is the big advantage of this all? After $k+m$ Arnoldi steps, we typically arrive at a situation that the Krylov subspace contains some reasonable approximations for (wanted) eigenvalues and some poor approximations for other (unwanted) eigenvalues.

If we carry out m QR steps with the m poorest Ritz values as shifts, then the reduced system (31.12) contains all the information for the remaining k Ritz values. Essentially we have removed the m poorest Ritz vectors from the $k+m$ dimensional Krylov subspace, and the surprising result is that we have rearranged the remaining Ritz vectors so that they form again a Krylov subspace: $\mathcal{K}^k(A;\widetilde{v}_1)$. The vector $\widetilde{v}_1$ can be expressed in the form

$$\widetilde{v}_1 = \gamma(A-\theta_m)\cdots(A-\theta_1)v_1,$$

where γ is a normalization constant, and the θ_i represent the shifts.

Sorensen [1992] proposes to continue this new Arnoldi process with another m steps, and to repeat the sketched subspace shrinking procedure. The whole procedure can be viewed as a mechanism that rakes repeatedly new supplementary information with m vectors from the entire space, to filter the desired supplements for the k approximated Ritz values, and to keep the gradually improved k-dimensional subspace in stock.

The vector $\widetilde{h}_{k+1,k}s$ shows how close the subspace spanned by the columns of $\widetilde{V}_k$ is to a k dimensional invariant subspace, and its norm can be used for a stopping criterion.

The resulting so-called Implicitly Restarted Arnoldi Algorithm (*IRA*) can globally be described as in Figure 22.

> Select v_1, k, and m
> do k Arnoldi steps to get V_{k+1}, $H_{k+1,k}$
> $r = h_{k+1,k}v_{k+1}$
> **while** $\|r\|_2 > \varepsilon$ **do**
>   do m additional Arnoldi steps
>   Compute Ritz values of $H_{k+m,k+m}$
>   Select m shifts (for instance the
>   unwanted Ritz values)
>   Apply m QR iterations with implicit shifts:
> $\Rightarrow Q$, $\widetilde{V}_{k+m}$, $H^{(m)}_{k+m,k+m}$
> $t = h_{k+m+1,k+m}\,v_{k+m+1}e^T_{k+m}Qe_k$
> $V_k := \widetilde{V}_k$, $H_{k,k} = H^{(m)}_{k,k}$
> $r = h^{(m)}_{k+1,k}\widetilde{v}_{k+1} + t$
> $\gamma = \|r\|_2$, $v_{k+1} = r/\gamma$, $h_{k+1,k} = \gamma$

Figure 22: The Implicitly Restarted Arnoldi Algorithm

The attractive property of the IRA process is that we effectively explore subspaces of high-dimensional Krylov subspaces, without the high computational overhead that would go with the unrestarted process. The process has been further refined in LEHOUCQ [1995], who considers the many practical aspects involved in a careful implementation. In particular, an analysis of the restarted Arnoldi iteration, its numerical stability, and deflation rules that allow for computation of clustered and/or multiple eigenvalues, are examined. The IRA method is available as a collection of Fortran subroutines in the Package ARPACK (LEHOUCQ, SORENSEN and YANG [1998]). In LEHOUCQ and SCOTT [1996], this software is further evaluated and compared with other software.

The implicit-restart idea has been applied, in combination with the two-sided Lanczos method, for model reduction in electronical applications (polezero analysis, stability of CD-players, see GRIMME, SORENSEN and VAN DOOREN [1994]). Similar ideas have recently been exploited to improve the Jacobi-Davidson method as well (see Section 46).

32 Harmonic Ritz values for Arnoldi's process

It is well-known that the subspace methods lead to eigenvalue approximations that tend to converge towards exterior eigenvalues. It may happen that an eigenvalue approximation is close to an interior eigenvalue of A, but in the next iteration step this needs not to be the case. In that situation we say that the eigenvalue approximation was on its way towards some exterior eigenvalue.

It is easy to reverse the direction of convergence of the eigenvalue approximations to the interior eigenvalues closest to the origin, by working with A^{-1}, but this is expensive. It is also possible to obtain eigenvalue approximations that converge (slowly) to the eigenvalues of A closest to the origin, from the subspaces generated with A. This leads to the so-called *Harmonic Ritz* values. These harmonic Ritz values (and vectors) have been defined for a symmetric matrix A in Chapter 25. They can be defined and computed in a similar way in combination with the Arnoldi pocess, as we will explain now.

The Arnoldi proces leads to

$$AV_m = V_{m+1}H_{m+1,m},$$

with $H_{m+1,m}$ a upper Hessenberg matrix with $m+1$ rows and m columns.

With $H_{m+1,m}^T V_{m+1}^T = V_m^T A^T$ it follows that:

$$V_m^T A^T A V_m = H_{m+1,m}^T V_{m+1}^T V_{m+1} H_{m+1,m}$$

or

$$V_m^T A^T A V_m = H_{m+1,m}^T H_{m+1,m} := U_m^T U_m. \tag{32.1}$$

The matrix U_m could be taken as the Cholesky factor of $H_{m+1,m}^T H_{m+1,m}$, but, as we will see, there is no need to compute U_m explicitly. From the equation

$$U_m^{-T} V_m^T A^T A V_m U_m^{-1} = I, \tag{32.2}$$

it follows that the columns of $AV_m U_m^{-1}$ form a set of m orthonormal vectors.

Note that these vectors form an orthonormal basis for $A\mathcal{K}^m(A; v_1)$. The matrix A^{-1} maps $A\mathcal{K}^m(A; v_1)$ onto $\mathcal{K}^m(A; v_1)$, and we may try to find a suitable form for the projection of the operator A^{-1}, that is the orthogonal restriction of A^{-1} with respect to $A\mathcal{K}^m(A;v_1)$. To that end we use the orthonormal basis, and find that the projection can be given as:

$$\begin{aligned} U_m^{-T}(AV_m)^T A^{-1} AV_m U_m^{-1} &= U_m^{-T}(AV_m)^T V_m U_m^{-1} \\ &= U_m^{-T} H_{m+1,m}^T V_{m+1}^T V_m U_m^{-1} \\ &= U_m^{-T} H_{m,m}^T U_m^{-1}. \end{aligned} \tag{32.3}$$

Hence, approximations for the eigenvalues of A^{-1} follow from

$$U_m^{-T} H_{m,m}^T U_i^{-1} t = \theta t,$$

or, with $s = U_i^{-1} t$

$$(U_m^T U_m)^{-1} H_{m,m}^T s = \theta s\,. \tag{32.4}$$

The vector s can be represented in terms of the basisvectors for $A\mathcal{K}^m(A;v_1)$, and can then be viewed as an approximate eigenvector of A^{-1}:

$$y = AV_m s.$$

In fact, the vector $V_m s$ may be expected to be an even better approximation for an eigenvector, in view of the fact that we have applied the operator A^{-1} onto y (one step of the Power Method).

This leads to the approximate eigenpair $(\theta^{-1}, V_m s)$ for a small (in absolute value) eigenvalue λ of A.

We may expect that the exterior eigenvalues of the projected A^{-1} converge (slowly) to the exterior eigenvalues of A^{-1}. Note that these are the eigenvalues of A closest to the origin.

Because of the shift invariance of the Krylov subspace (see (14.1)), convergence of the Harmonic Ritz values can be directed towards eigenvalues close to a specified target τ. Define

$$M_{m,\tau} := (H_{m+1,m} - \tau I_{m+1,m})^T (H_{m+1,m} - \tau I_{m+1,m}), \tag{32.5}$$

then it may easily be verified that the solutions of

$$M_{m,\tau}^{-1} (H_{m,m} - \tau I_{m,m})^T s_\tau = \theta_\tau s_\tau, \tag{32.6}$$

lead to approximate eigenpairs $(\tau + \theta_\tau^{-1}, V_m s_\tau)$, that converge, for increasing m, towards some eigenpair (λ, x) of A, with λ close to τ.

In Paige, Parlett and van der Vorst [1995], these eigenvalue approximations, in connection with the related Lanczos process, were called

Harmonic Ritz values, and some nice relations for Harmonic Ritz values for symmetric indefinite matrices are given in that paper. Harmonic Ritz values had already been studied from a different viewpoint by other authors. FREUND [1992] studies them as the zeros of the GMRES and MINRES iteration polynomials. MORGAN [1991] has observed that the Harmonic Ritz values and vectors are very suitable for restarting purposes, if one wants to compute interior eigenvalues with subspaces of restricted dimension. This can be easily carried out in the Implicitly Restarted Arnoldi method. In SLEIJPEN and VAN DER VORST [1996] and FOKKEMA, SLEIJPEN and VAN DER VORST [1998], the Harmonic Ritz values are considered in connection with the Jacobi-Davidson process.

33 Numerical examples

In this section we will show the behavior of the Harmonic Ritz values for some examples. In this respect, the examples given in GOOSSENS and ROOSE [1999], in connection with the convergence behavior of GMRES, are also illustrative.

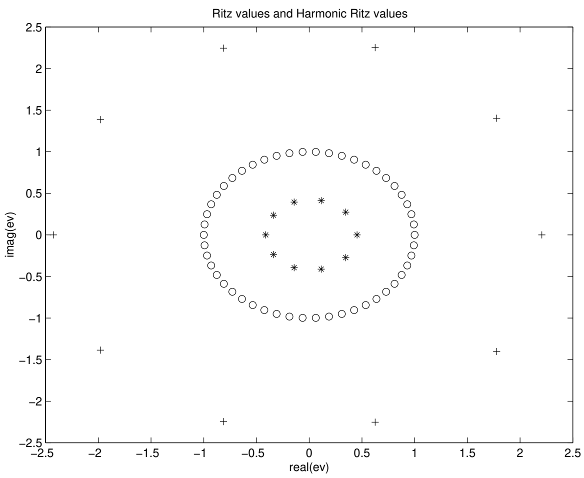

Figure 23: Ritz values (*), Harmonic Ritz values (+) and Eigenvalues (o), for Example 1

Example 1 First consider the Example 1 of Section 29 again. With a starting vector $v_1 = e_1$, the matrices $H_{m,m}$ have only eigenvalues zero for $m < n$, so that all Ritz values and Harmonic Ritz values are zero as well for these values of m. To make the situation slightly more interesting we have selected

$$v_1 = (0.0001, 0.0001, \ldots, 0.0001, 1)^T.$$

For the case $n = 50$, the exact eigenvalues (indicated by 'o') have been plotted in Figure 23. For $m = 10$ the Ritz values ('$*$') and the Harmonic Ritz values ('$+$') have been marked in the same figure as well. We see that the Ritz values approach the eigenvalues from inside, while the Harmonic Ritz values approach from outside.

The Harmonic Ritz values are not much better (maybe in a relative sense they are) than the Ritz values; together they indicate that there are convergence problems.

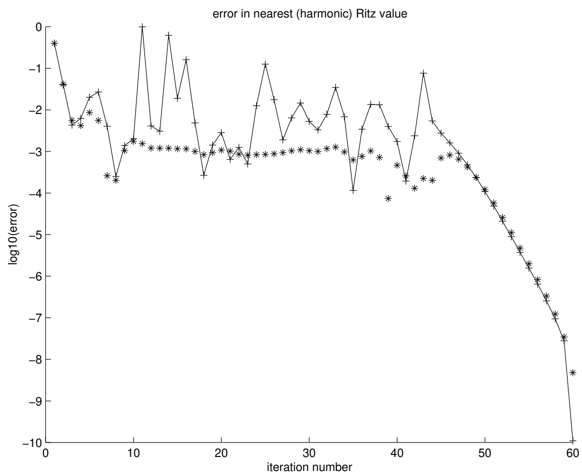

Figure 24: Errors for the Ritz values ($+$) and Rayleigh quotient of Harmonic Ritz vectors ($*$), closest to a selected eigenvalue for Example 2

Example 2 In this example we mimic the situation that A has an interior eigenvalue that is well separated from the other eigenvalues. Such a situation may be encountered for instance in acoustic problems, where the isolated interior eigenvalue may represent a resonance mode in a model.

Our example, however, is completely artificial with the purpose that it can easily be repeated. We construct a block diagonal matrix A, with

$$a(j,j) = j \quad \text{for} \quad j = 1, 2, 3, 10, 17, 18, \ldots, 100.$$

A has six 2 by 2 blocks on its diagonal:

$$\begin{bmatrix} a_{4,4} & a_{4,5} \\ a_{5,4} & a_{5,5} \end{bmatrix} = \begin{bmatrix} 4.5 & 1 \\ -1 & 4.5 \end{bmatrix} \quad , \quad \begin{bmatrix} a_{6,6} & a_{6,7} \\ a_{7,6} & a_{7,7} \end{bmatrix} = \begin{bmatrix} 6.5 & 3 \\ -3 & 6.5 \end{bmatrix},$$

$$\begin{bmatrix} a_{8,8} & a_{8,9} \\ a_{9,8} & a_{9,9} \end{bmatrix} = \begin{bmatrix} 8.5 & 5 \\ -5 & 8.5 \end{bmatrix} \quad , \quad \begin{bmatrix} a_{11,11} & a_{11,12} \\ a_{12,11} & a_{12,12} \end{bmatrix} = \begin{bmatrix} 11.5 & 5 \\ -5 & 11.5 \end{bmatrix},$$

$$\begin{bmatrix} a_{13,13} & a_{13,14} \\ a_{14,13} & a_{14,14} \end{bmatrix} = \begin{bmatrix} 13.5 & 3 \\ -3 & 13.5 \end{bmatrix}, \begin{bmatrix} a_{15,15} & a_{15,16} \\ a_{16,15} & a_{16,16} \end{bmatrix} = \begin{bmatrix} 15.5 & 1 \\ -1 & 15.5 \end{bmatrix}.$$

This means that A has eigenvalues

$$1,\ 2,\ 3,\ 4.5 \pm i,\ 6.5 \pm 3i,\ 8.5 \pm 5i,\ 10,$$

$$11.5 \pm 5i,\ 13.5 \pm 3i,\ 15.5 \pm i,\ 17,\ 18,\ \ldots,\ 100.$$

Clearly, the eigenvalue 10 is an interior eigenvalue that is relatively well-separated from the surrounding (complex) eigenvalues.

We start Arnoldi's method with a starting vector v of all ones, except $v(10) = 100$. This mimics the situation where we have some estimate for the wanted eigenvector (or in terms of our example: the resonance vector). By selecting this vector we avoid a phase of poor convergence of the Ritz and Harmonic Ritz approximations for $\lambda = 10$. For the Harmonic Ritz values we select a target value $\tau = 9.5$.

Then, in each iteration of Arnoldi's method we select the Ritz value closest to τ and we select the Harmonic Ritz value closest to τ. Since the Harmonic Ritz vector represents the good information for restart, we measure the value of the Harmonic Ritz vector u corresponding to the selected Harmonic Ritz value by computing the Raleigh quotient $\rho = u^* A u$.

In Fig. 24 we have plotted the $log10$ of the distances between the selected Ritz value (indicated by $+$; connected by the drawn line) to $\lambda = 10$, as well as the $log10$ of the distance of the Rayleigh quotient (for the selected Harmonic Ritz vector) to that eigenvalue (these distances are indicated with $*$).

We see that in average there is no advantage for the Harmonic Ritz values and vectors, but for many iterations the Ritz value (and corresponding Ritz vector) are significantly worse. They would have led to poor restart vectors, if the Arnoldi process were restarted at that iteration with the current best eigenvector approximation as the restart vector. In this sense, the example illustrates that the Harmonic Ritz approximations may have an advantage for restart purposes. If the process is not restarted then they have no advantage over the Ritz approximations.

CHAPTER VIII

The Two-sided Lanczos Method

34 The Petrov-Galerkin approach: Bi-Lanczos

If the matrix A is Hermitian then we can, at modest computational costs per iteration step, generate an orthonormal basis for the Krylov subspace $\mathcal{K}^m(A; v_1)$. It suffices to put each newly generated vector Av_j orthogonal to the two previous basis vectors v_j and v_{j-1} only. This leads to a reduced matrix that is symmetric tridiagonal. A theoretical result of FABER and MANTEUFFEL [1984] shows that this is not possible for general unsymmetric matrices.

In this chapter we will discuss a generalization of the Lanczos method for unsymmetric matrices. This method is known as the Two-sided Lanczos method, or Bi-Lanczos method.

For ease of notation, we will restrict ourselves to the real case. For an unsymmetric matrix $A \in \mathbb{R}^{n\times n}$ we will try to obtain a suitable non-orthogonal basis with a 3-term recurrence, by requiring that this basis is orthogonal with respect to some other basis.

This leads to the *Bi-Lanczos Method* (LANCZOS [950]), which is very similar to the Bi-CG method. Both methods lead to the same projected tridiagonal matrix. FREUND and NACHTIGAL [1990], FREUND, GUTKNECHT and NACHTIGAL [1993], FREUND and NACHTIGAL [1996], and CULLUM and WILLOUGHBY [1986] have published codes for the computation of eigenvalues using this approach.

We start by constructing a not necessarily orthogonal basis for the Krylov subspace:

$$h_{i+1,i}v_{i+1} = Av_i - \sum_{j=1}^{i} h_{j,i}v_j, \tag{34.1}$$

which can be rewritten in matrix notation as $AV_i = V_{i+1}H_{i+1,i}$.

The coefficients $h_{i+1,i}$ define a scaling for v_{i+1}, and a natural choice would be to select them such that $\|v_{i+1}\|_2 = 1$. In Bi-CG implementations, a popular choice is to select $h_{i+1,i}$ such that $\|v_{i+1}\|_2 = \|r_{i+1}\|_2$, where r_i denotes the residual of the linear system $r_i = b - Ax_i$.

Clearly, we cannot use V_i for the projection, but suppose we have a W_i for which $W_i^T V_i = D_i$ (an i by i diagonal matrix with diagonal entries d_i), and for which $W_i^T v_{i+1} = 0$.
Then

$$W_i^T A V_i = D_i H_{i,i}, \tag{34.2}$$

and now our goal is to find a W_i for which $H_{i,i}$ is tridiagonal. This means that $V_i^T A^T W_i$ should be tridiagonal too, and since this is a similar expression as in the left-hand side of (34.2) it suggests to generate the w_i with A^T.

We start with an arbitrary $w_1 \neq 0$, such that $w_1^* v_1 \neq 0$. Then we generate v_2 with (34.1), and orthogonalize it with respect to w_1, which means that

$$h_{1,1} = \frac{w_1^T A v_1}{w_1^T v_1}.$$

Since

$$w_1^T A v_1 = (A^T w_1)^T v_1 = v_1^T A^T w_1,$$

this implies that w_2, generated with

$$h_{2,1} w_2 = A^T w_1 - h_{1,1} w_1,$$

is also orthogonal to v_1.
This can be continued, and we see that we can create biorthogonal basis sets $\{v_j\}$, and $\{w_j\}$, by making the new v_i orthogonal to w_1 up to w_{i-1}, and then by generating w_i with the same recurrence coefficients[1], but with A^T instead of A.
Now we have that $W_i^T A V_i = D_i H_{i,i}$, and also that

$$V_i^T A^T W_i = D_i H_{i,i}. \tag{34.3}$$

This implies that $D_i H_{i,i}$ is a symmetric tridiagonal matrix, which gives us the desired 3-term recurrence relation for the v_j's, and the w_j's. Note that $v_1, \ldots, v_i$ form a basis for $K^i(A; v_1)$, and $w_1, \ldots, w_i$ form a basis for $K^i(A^T; w_1)$. In order to stress the tridiagonality of the reduced matrix $H_{i+1,i}$, we will further denote it as $T_{i+1,i}$ throughout this chapter.

In the unsymmetric Lanczos method (Lanczos [1950]), dual bases $\{v_j\}$ and $\{w_j\}$ are generated for the Krylov subspace $K^i(A; v_1)$ and its adjoint

[1] Note that in the complex case, the recursion coefficients for the w_i should be complex conjugated

$K^i(A^T; w_1)$. The v_j are generated with a three term recurrence relation, with A:

$$\gamma_j v_{j+1} = A v_j - \alpha_j v_j - \beta_{j-1} v_{j-1},$$

and the w_j with a similar recurrence for A^T:

$$\gamma_j w_{j+1} = A^T w_j - \alpha_j w_j - \beta_{j-1} w_{j-1}.$$

In Figure 25, we show schematically an algorithm for the two-sided Lanczos, suitable for execution with an unsymmetric matrix A.

Select a normalized pair v_1, w_1 (for instance, $w_1 = v_1$)
such that $w_1^T v_1 = \delta_1 \neq 0$
$\beta_0 = 0$, $w_0 = v_0 = 0$
for $i = 1, 2, ...$
  $p = A v_i - \beta_{i-1} v_{i-1}$
  $\alpha_i = w_i^T p / \delta_i$
  $p = p - \alpha_i v_i$
  $\gamma_{i+1} = \|p\|_2$
  $v_{i+1} = p / \gamma_{i+1}$
  $w_{i+1} = (A^T w_i - \beta_{i-1} w_{i-1} - \alpha_i w_i) / \gamma_{i+1}$
  $\delta_{i+1} = w_{i+1}^T v_{i+1}$
  $\beta_i = \gamma_{i+1} \delta_{i+1} / \delta_i$
end;

Figure 25: The Two-sided Lanczos algorithm

In this algorithm, we choose γ_i such that $\|v_i\|_2 = 1$, and use the same γ_i as a scaling for w_i.

We have tacitly assumed that the inner products $w_j^T v_j \neq 0$, since the value 0 would lead to an unsollicited breakdown of the method if v_j and w_j are not equal to zero. If v_i is equal to zero, then we have an invariant subspace, and in that case the eigenvalues of $T_{i,i}$ are eigenvalues of A. In Section 37 we will discuss possibilities for handling breakdown situations

Note that for symmetric matrices A, Bi-Lanczos generates the same solution as the Lanczos method, provided that $w_1 = v_1$. However, the biorthogonal variants does so at the cost of two matrix vector operations per iteration step.

There are alternative (and simpler) forms of deriving the two-sided Lanczos algorithm. For instance, WILKINSON [1965, Chapter 6.36] starts from two Arnoldi like recursions for the basis vectors of $K^m(A; v_1)$ and $K^m(A^T; w_1)$:

$$\begin{aligned} h_{j+1,j} v_{j+1} &= A v_j - \sum_{i=1}^{j} h_{i,j} v_i, \\ g_{j+1,j} w_{j+1} &= A^T w_j - \sum_{i=1}^{j} g_{i,j} w_i, \end{aligned}$$

and requires that v_{j+1} is orthogonal to all previous w_i and that w_{j+1} is orthogonal to all previous v_i. Clearly, this defines, apart from the constants $h_{j+1,j}$ and $g_{j+1,j}$, the vectors v_{j+1} and w_{j+1}, once the previous vectors are given. Then we have that

$$AV_j = V_{j+1} H_{j+1,j} \quad \text{and} \quad A^T W_j = W_{j+1} G_{j+1,j}.$$

Since each new v_{j+1} is only orthogonal with respect to the w_i, for $i < j$, and likewise for w_{j+1} with respect to the v_i, it follows that

$$W_j^T V_j = L_{j,j} \quad \text{and} \quad V_j^T W_j = K_{j,j},$$

where $L_{j,j}$ and $K_{j,j}$ are lower triangular. Clearly

$$K_{j,j}^T = L_{j,j},$$

so that both matrices are diagonal. Let us denote this matrix as D. Then, we have that

$$W_j^T A V_j = D H_{j,j}$$

and also

$$V_j^T A^T W_m = D G_{j,j}.$$

This shows that $H_{j,j}$ and $G_{j,j}$ must be tridiagonal.

Obviously, this derivation is clear and crisp, but we have chosen ours, because we believe that it brings in more naturally the necessity to work with A^T. CULLUM, KERNER and WILLOUGHBY [1989, p. 22] start the derivation of the two-sided Lanczos method with the assumption that there exists a nonsingular matrix V so that $V^{-1}AV$ is tridiagonal. This leads to an extremely simple derivation of the method, but the assumption may not come as a natural one.

There is freedom in the selection of the γ_i, which leads to different D_i and $T_{i,i}$. For instance, in Bi-CG the γ_i are chosen such that v_i are equal to the residuals in the iterative solution process for $Ax = b$. Wilkinson [1965, Chapter 6.36] views the choice $\gamma_i = 1$ as a natural choice for the purpose of analysis of the possible failure (that is when a $\delta_i = 0$). Saad [1992, Chapter 4.1] suggests to scale the vectors w_i and v_i by the same factor to the effect that $\delta_i = w_i^T v_i = 1$. This typically leads to a symmetric tridiagonal reduced matrix $T_{i,i}$. Saad also mentions that one can scale both w_i and v_i to unit norm length. Freund, Gutknecht and Nachtigal [1993] use this in their criteria for the look-ahead strategy. Cullum, Kerner and Willoughby [1989, p. 25] argue that the identical scaling of v_i and w_i, such that $\delta_i = 1$, symmetrizes the effects of rounding errors within the procedure. They suggest that this may be important, because it was seen as being important for the real symmetric case (Paige [1971]).

For the computation of approximate eigenpairs, we may proceed in a similar way as in the symmetric case:

$$AV_j = V_{j+1}T_{j+1,j}, \tag{34.4}$$

but here we use the matrix $W_j = [w_1, w_2, ..., w_j]$ for orthogonality conditions on an approximate vector. Let $V_j y$ be the wanted eigenvector approximation in V_j, with corresponding eigenvalue approximation, then we impose the *Petrov-Galerkin* condition

$$W_j^T(AV_j y - \theta V_j y) = 0,$$

or

$$W_j^T AV_j y - \theta W_j^T V_j y = 0.$$

This shows that y and θ form an eigenpair of the generalized symmetric eigenproblem

$$D_j T_{j,j} y = \theta D_j y,$$

or of the standard unsymmetric eigenproblem

$$T_{j,j} y = \theta y.$$

An interesting aspect of the two-sided Lanczos approach is that it also admits possibilities for the approximation of left eigenvectors. Let $p^* W_j^T$ denote an approximation for a left eigenvector corresponding to an approximate eigenvalue θ. Note that the eigenvectors of A and $T_{j,j}$ may be complex; this explains the $*$. Then the Petrov-Galerkin condition with respect to V_j leads to

$$p^* W_j^T AV_j - \theta p^* W_j^T V_j = 0, \tag{34.5}$$

which shows that the pair p, θ is a *left eigenpair* of the symmetric generalized eigenproblem

$$p^* D_j T_{j,j} = \theta p^* D_j.$$

Since $D_j T_{j,j}$ is real symmetric, the vector p is a right eigenvector of $T_{j,j}$. It is, becauses of the Petrov-Galerkin condition, natural to associate Petrov-Galerkin approximations for eigenvectors of A with the eigenpair y, θ of $T_{j,j}$. We will use the terminology of *Petrov value* for θ, *right Petrov vector* for $V_j y$, and *left Petrov vector* for $W_j y$.

This is a very convenient aspect of the Two-sided Lanczos process, because the left and right eigenvectors can be used to get an impression of the stability of the computed eigenvalue. As we have seen in Section 5, if x and z denote the right and left eigenvector corresponding to a single isolated eigenvalue λ of A, then

$$\kappa_\lambda = \frac{1}{z^* x}$$

is the condition number of λ. This condition number shows how perturbations to A are related to perturbations in λ (in first order). Note that this condition number can be conveniently estimated as

$$\kappa_\lambda \approx \kappa_\theta = \frac{1}{y^* W_j^* V_j y} = \frac{1}{y^* D_j y},$$

which can be computed right away from the right eigenpair of $T_{j,j}$.

Although we cannot say much about the actual convergence behavior of the Petrov values associated with the Two-sided Lanczos process, the typical behavior is very much the same as for Arnoldi's method. That is, the Petrov values tend to approximate isolated exterior eigenvalues in the spectrum of A first. If the wanted eigenvalues are not well separated from the unwanted ones, or when the wanted eigenvalues are in the interior of the spectrum, then one should consider to work with operations with the shift-and-inverted matrix $(A - \tau I)^{-1}$, for eigenvalues close to τ. This requires an efficient factorization of this matrix and if this is not well possible then working with the *Harmonic Petrov values* may be an alternative. In Section 36 we will discuss how we can define these harmonic Petrov values in the context of the Two-sided Lanczos process.

For unsymmetric problems the picture is still rather obscure. Although the Two-sided Lanczos process often leads to good approximations for some eigenvalues, its convergence behavior is not well understood. Even the convergence behavior of the Arnoldi method is not well understood, notwithstanding the fact that it can be viewed as an accelerated power iteration.

CULLUM [1996] shows that the convergence behavior of the Two-sided Lanczos process, for eigenvalues, can be mimicked by Arnoldi's process and vice versa. For the unsymmetric eigenproblem we also face the difficulty that there are no algorithms for computing the eigenvalues and eigenvectors of unsymmetric tridiagonal matrices, that are as efficient and robust as for the symmetric case.

In the complex case, most of the arithmetic can be kept real by writing the complex matrix A as $A_1 + \imath A_2$, and the eigenvector z as $z_1 + \imath z_2$, with $\imath = \sqrt{-1}$. The Two-sided Lanczos process can then be applied for the equivalent real system

$$\begin{bmatrix} A_1 & -A_2 \\ A_2 & A_1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} x \\ y \end{bmatrix}.$$

Despite the bad reputation of the standard unsymmetric Lanczos process, software presented in CULLUM and WILLOUGHBY [1986] has been quite successful. In their approach the inaccurate eigenapproximations, due to numerical instabilities, were identified from comparison of results for submatrices of $T_{j,j}$. The main idea is that the starting vector has nonzero components in directions of eigenvectors corresponding to eigenvectors that we are interested in. With this starting vector the Krylov subspaces are generated. When we compare the eigenvalues of $T_{j,j}$ with those of the tridiagonal matrix of order $j-1$ that is obtained if we skip the first row and column from $T_{j,j}$, then we compare Petrov values of the current Krylov subspace with Petrov values for a subspace from which the starting vector has been removed. Since this starting vector contains essential information of the desired eigenvector directions, no Petrov value can have converged *unless* all information for the corresponding eigenvector has entered through rounding errors. This kind of heuristics helps to identify the so-called *spurious* eigenvalues.

35 Numerical Example

We will present a simple example, that illustrates some of the effects that can be observed with the Two-sided Lanczos algorithm in Figure 25 . The matrix A has order 100 and has 2 by 2 blocks and 1 by 1 blocks along its diagonal (the other elements are equal to zero). The five 2 by 2 blocks are

$$\begin{bmatrix} 1 & 2 \\ -2 & 1 \end{bmatrix}, \begin{bmatrix} 3 & \frac{1}{2} \\ -\frac{1}{2} & 3 \end{bmatrix}, \begin{bmatrix} 5 & \frac{1}{4} \\ -\frac{1}{4} & 5 \end{bmatrix},$$

$$\begin{bmatrix} 95 & 5 \\ -5 & 95 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 100 & 10 \\ -10 & 100 \end{bmatrix}.$$

This means that the matrix has eigenvalues $1 \pm 2\imath$, $3 \pm \frac{1}{2}\imath$, $5 \pm \frac{1}{4}\imath$, $95 \pm 5\imath$, and $100 \pm 10\imath$. The other 90 eigenvalues are real and uniformly distributed in the interval $[20, 80]$. The starting vector has also been chosen real as a vector with all ones (and then normalized to length 1), $w_1 = v_1$.

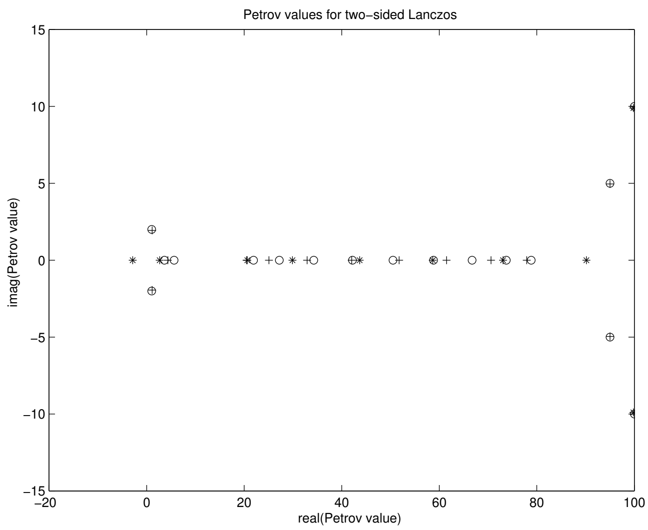

Figure 26: Petrov values for 10 ($*$), 15 ($+$), and 17 (o) iterations

In Fig. 26 we see the Petrov values for 10, 15, and 17 iterations with the Two-Sided Lanczos algorithm. We see that there is no monotonicity, as for the symmetric Lanczos process (nor can this be expected): after 10 iterations the smallest Petrov value is smaller than 1.

After some 17 iterations the largest Petrov values have more or less converged to the largest eigenvalues in absolute value of A ($100 \pm 10\imath$), which could have been expected, since this conjugate eigenpair is relatively well separated from the remainder of the spectrum. Also some of the other complex eigenpairs have been detected in some accuracy. After convergence of an eigenvalue to almost machine precision, the biorthogonality among the sets $\{w_i\}$ and $\{v_i\}$ gets lost. In the symmetric Lanczos process this had pretty harmless effects: it led to the occurrence of almost double Ritz values. This is not necessarily the case for the Two-sided Lanczos process. If an eigenvalue has converged, then rounding errors may lead to small values of the δ_i, and this may lead to Petrov values that can be located almost anywhere. These 'wild' Petrov values can be detected by a test on spurious Petrov values, but they may have a wider effect: if they are relatively large then they may also affect the accuracy in the 'good' Petrov values. Therefore, it is advised to apply some form of selective biorthogonalization with respect to the converged Petrov vectors, or to restart.

36 Harmonic Petrov values for Bi-Lanczos

The Lanczos proces for unsymmetric matrices leads to a set of basis vectors $v_1, v_2, \ldots, v_{i+1}$ for the Krylov subspace (for a general matrix A, and starting vector v_1), biorthogonal with respect to the set $w_1, w_2, \ldots, w_{i+1}$. For the matrices associated with these vectors we have that $W_i^T V_i = D_i$.

Then we have the following relations:

$$AV_i = V_{i+1}T_{i+1,i},$$

$$A^T W_i = W_{i+1}T_{i+1,i}.$$

From the biorthogonality relations it follows that:

$$W_i^T AAV_i = T_{i+1,1}^T W_{i+1}^T V_{i+1} T_{i+1,i}$$

or

$$W_i^T AAV_i = T_{i+1,i}^T D_{i+1} T_{i+1,i} = Q_i \Lambda Q_i^T,$$

with orthogonal Q_i and diagonal real Λ, since $T_{i+1,i}^T D_{i+1} T_{i+1,i}$ is a symmetric real matrix. From the equation

$$Q_i^T W_i^T AAV_i Q_i = \Lambda,$$

it follows that the columns of AV_iQ_i are biorthogonal with respect to those of $A^T W_i Q_i$.

Note that the columns of AV_iQ_i form a basis for $AK_i(A; v_1)$. The matrix A^{-1} maps $AK_i(A; v_1)$ onto $K_i(A; v_1)$, and we may try to find a suitable form for the projection of the operator A^{-1}, that is the biorthogonal restriction of A^{-1} with respect to $AK_i(A; v_1)$. To that end we use the biorthogonal bases, and we find that the Petrov-Galerkin conditions for an approximate eigenvector of A^{-1} in $AK_i(A; v_1)$ can be formulated as:

$$Q_i^T W_i^T A \left(A^{-1} AV_i Q_i y - \theta AV_i Q_i y\right) = 0,$$

because any vector in the subspace $AK_i(A; v_1)$ can be expressed in the form AV_iQ_iy.

This leads to

$$Q_i^T W_i^T AV_iQ_iy = \theta Q_i^T W_i^T AAV_iQ_iy,$$

so that the pair y, θ is an eigenpair of

$$Q_i^T D_iT_{i,i}Q_iy = \theta\Lambda y,$$

or, with $z = Q_i$,

$$D_iT_{i,i}z = \theta T_{i+1,i}^T D_{i+1}T_{i+1,i}z. \tag{36.1}$$

Note that this is again a symmetric generalized eigenproblem. The solution z, θ leads to a *Petrov vector* AV_iz with associated *Petrov value* θ for A^{-1}.

For a left eigenvector of A^{-1}, we may look in the subspace spanned by the vectors A^Tw_j. A left eigenvector of A^{-1} is a right eigenvector of A^{-T}. The candidates for good eigenvector approximations from the subspace may be expressed as $p = A^TW_iq$, and the Petrov-Galerkin conditions with respect to $AK_i(A; v_1)$, lead to

$$V_i^T A^T A^{-T} A^T W_iq - \theta V_iA^T A^T W_iq = 0,$$

or

$$V_i^T A^T W_iq - \theta V_iA^T A^T W_iq = 0,$$
$$D_iT_{i,i}q - \theta T_{i+1,i}^T D_{i+1}T_{i+1,i}q = 0,$$

so that the vector q is an eigenvector of the symmetric pencil

$$(D_iT_{i,i}, T_{i+1,i}^T D_{i+1}T_{i+1,i}).$$

This implies that from an eigenpair θ, y of this pencil, we can compute the *left Petrov vector* A^TW_iy, and the *right Petrov vector* AV_iy.

Note that since AV_iy is an approximate eigenvector of A^{-1}, we have that $r = A^{-1}AV_iy - \theta AV_iy$, so that V_iy is a right approximate eigenvector of A with residual $\frac{1}{\theta}r$, for the approximate eigenvalue θ^{-1}. Likewise, W_iy is an approximate left eigenvector of A, with approximate eigenvalue θ^{-1}. The vector V_iy is referred to as a *harmonic right Petrov vector*, the vector W_iy is called the em harmonic left Petrov vector, with associated *harmonic Petrov value* θ^{-1}.

Since the Petrov values converge to the dominant (that is exterior) eigenvalues of A^{-1}, we may expect that the inverses of the Harmonic Petrov values, namely the inverses of the Petrov values of A^{-1} converge towards the interior eigenvalues closest to the origin. We will illustrate this with an example.

Convergence towards a target τ The convergence of the Petrov values (and vectors) is shift independent, because the Krylov subspace is shift independent, that is $K^m(A;v) = K^m(A - \tau I; v)$. In relation with the Two-sided Lanczos process, the biorthogonalization procedure for $A - \tau I$ leads to exactly the same V_i, W_i, and D_i, as for A. For the shifted matrix we have that

$$(A - \tau I)V_i = V_{i+1}(T_{i+1,i} - \tau I_{i+1,i}),$$

and

$$(A^T - \tau I)W_i = W_{i+1}(T_{i+1,i} - \tau I_{i+1,i}),$$

with $T_{i+1,i}$ the same as for A. The harmonic Petrov values, however, are not independent of the shift. We can compute them from (36.1) with the shifted matrix

$$T^{(\tau)}_{i+1,i} := T_{i+1,i} - \tau I_{i+1,i},$$

from the generalized symmetric eigenproblem

$$D_i T^{(\tau)}_{i,i} z = \theta (T^{(\tau)}_{i+1,i})^T D_{i+1} T^{(\tau)}_{i+1,i} z. \tag{36.2}$$

Each τ leads to a different set of harmonic Petrov values and the harmonic Petrov values closest to τ represent the inverses of the approximated extremal eigenvalues of $(A - \tau I)^{-1}$. This suggests that the corresponding harmonic Petrov vectors $V_i z$ are good candidates for the construction of restart vectors, if one wants to compute interior eigenpairs.

It should be noted that although the convergence pattern for the harmonic Petrov values is oriented towards the eigenvalues closest to the shift τ, this does not imply that they are necessarily much better than the standard Petrov values. The reason is that the harmonic Petrov values are the Petrov values of $(A - \tau I)^{-1}$, *but with respect to the Krylov subspace generated with* A itself. This explains that the convergence of the Harmonic Petrov values should not be compared with the Petrov values that one obtains with a true shift-and-invert approach. The main niche for the harmonic Petrov values, in our opinion, is for situations where shift-and-invert is too expensive and where restarts are necessary.

Numerical example We will now describe an example that demonstrates that the Harmonic Petrov values may lead to better convergence behavior after restart than the standard Petrov values. Our example is rather small: the order of A is 100. Since the method does not take an advantage of a particular structure of the matrix, we have selected a block diagonal matrix

A. The first 83 diagonal elements are single elements: $-91, -90, \ldots, -9$. Then we have three 2 by 2 blocks:

$$\begin{bmatrix} -8 & \frac{1}{4} \\ -\frac{1}{4} & -7 \end{bmatrix}, \begin{bmatrix} -6 & \frac{1}{2} \\ -\frac{1}{2} & -5 \end{bmatrix}, \begin{bmatrix} -4 & 2 \\ -2 & -3 \end{bmatrix}.$$

Note that these blocks cause A to have three complex conjugate eigenpairs, e.g., the conjugate eigenpair nearest to the origin is $-3.5 \pm \imath\frac{1}{2}\sqrt{15}$.

The 90-th diagonal element is .25, followed by two 2 by 2 blocks

$$\begin{bmatrix} 4 & 10 \\ -10 & 5 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 6 & 5 \\ -5 & 7 \end{bmatrix},$$

and the 6 last diagonal elements are 8, 9, ..., 13. Clearly the eigenvalue closest to the origin is .25 and we pretend that we do not know this and want to identify it, as well as the associated left and right eigenvectors, with the Two-sided Lanczos method, as given in Figure 25. We take for $w_1 = v_1$ the normalized vector with all equal coordinates.

We carry out 8 cycles of 20 Lanczos iterations. After each cycle we determine the (harmonic) Petrov value closest to the target $\tau = 0$, and the corresponding (harmonic) Petrov vector $z = V_i y$. Then we restart with $v_1 = V_i y$, $w_1 = W_i y$. In order to improve the convergence we have approximated $A^{-1}V_i y$, with 30 steps of unpreconditioned GMRES, and taken this as the actual starting vector v_1 for every cycle of Lanczos iterations. We have done the same for w_1; this has been taken as the GMRES(30) approximation of $A^{-T}W_i y$. This is a very expensive process for such a small (trivial) matrix, but we expect similar convergence behavior for very large matrices that have a similar eigenvalue distribution. This small example indicates what we may expect for such larger problems. The 30 GMRES steps represent an approximation for A^{-1}, in actual problems one may take a suitable preconditioner instead of this.

In Figure 27 we have plotted the residual for the normalized right eigenvector approximation ($\|AV_i y - \theta V_i y\|_2 / \|V_i y\|_2$) after each cycle. The dotted curve represents the results for the standard Petrov values and vectors, the solid curve represents the results for the harmonic Petrov values.

37 Breakdown and Look-ahead

We have assumed that $\delta_i \neq 0$, that is $w_i^T v_i \neq 0$. The generation of the biorthogonal basis breaks down if for some i the value of $w_i^T v_i = 0$, which is

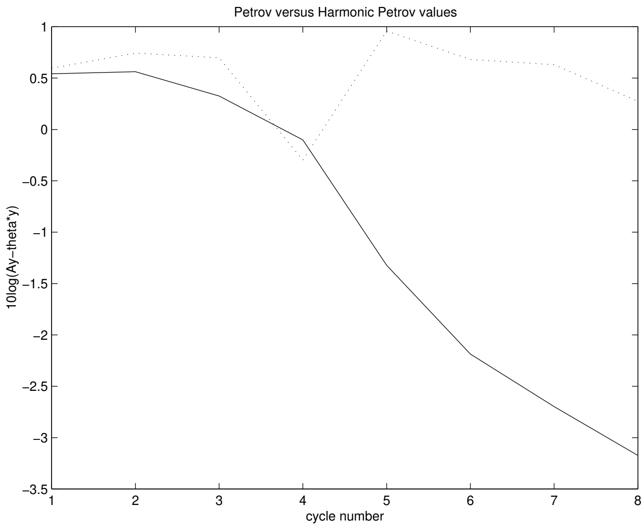

Figure 27: Restart with harmonic Petrov values (-line-) and standard Petrov values (-dotted-)

referred to in literature as a *serious breakdown*, if not at the same iteration $v_i = 0$ and $w_i = 0$. Likewise, when

$$\frac{w_i^T v_i}{\|w_i\|_2 \|v_i\|_2} \approx 0,$$

we have a near-breakdown. See Wilkinson [1965, Chapter 6.36] for a discussion on the failures in the Lanczos process.

A naive approach for overcoming a breakdown is the following strategy. We simply restart with the most recent eigenvector approximations as soon as a diagonal element gets small. Of course, this strategy is surprisingly simple indeed, but one should realise that at a restart the Krylov subspace, that has been built up so far, is thrown away, and this destroys the possibility of faster (i.e., superlinear) convergence.

A more mature way to get around a breakdown is the so-called *Look-ahead strategy*, which comes down to taking a number of successive basis vectors for the Krylov subspace together and to make them blockwise biorthogonal. This has been worked out in detail in Parlett, Taylor and Liu [1985], Freund, Gutknecht and Nachtigal [1993], Freund and Nachtigal [1990], and Freund and Nachtigal [1991].

In the early 1980s, TAYLOR [1982] and PARLETT, TAYLOR and LIU [1985], suggested to expand the Krylov subspace with sets of basis vectors that were block-dual. The idea is that certain dimensions, namely those for which breakdown occurs, cannot be used for projection. One has to postpone the inspection of the projected system until a block R_i of sufficient large dimension is discovered for which $S_i^T R_i$ is not too small. This process of delaying the actual projection step was called *Look-ahead*. In the period '85–'95 the look-ahead technique was generally accepted as a necessary element for the Lanczos method. GUTKNECHT [1992] gives a detailed theoretic basis for the look-ahead mechanism. FREUND and NACHTIGAL [1991] incorporated this technique in their QMR method, and it was shown later that their codes could also be used for effective solution of large sparse unsymmetric eigenproblems, see, for instance, FELDMAN and FREUND [1995] for an example of this. See also GUTKNECHT [1997] for a thorough overview of Bi-Lanczos techniques, including the treatment of breakdown.

BREZINSKI and REDIVO-ZAGLIA [1994], consider the breakdown problem from a polynomial point of view. The vectors in the Krylov subspaces can be seen as the results of matrix polynomials acting on the starting vectors, and likewise, the biorthogonality relations can be viewed as orthogonality relations for polynomials, with respect to a finite innerproduct (in which the weights are defined by the starting vectors). They showed that for some degrees the polynomials do not exist, and they also showed how higher degree orthogonal polynomials could be defined, by temporarily using polynomials of degree larger than 1 in the recursion formulas. This is equivalent with the block-wise expansion for the dual bases in the look-ahead approach.

We will now describe the look-ahead Lanczos algorithm in some more detail. For our description we follow closely the presentation given in FREUND and NACHTIGAL [1991, Section 2]. The main idea is that we give up single vector biorthogonality, since that leads sometimes to breakdowns, and instead of this we strive for block biorthogonality. This will give more freedom for the vectors that belong to one block. Let the generated vectors v_1, v_2, ..., v_m, and w_1, w_2, ..., w_m, that are generated by the look-ahead Two-sided Lanczos algorithm belong to k groups (blocks) of vectors. Let us suppose that in the process $k-1$ blocks have been formed already (we admit single vector blocks):

$$V_\ell = [v_{m_\ell}, v_{m_\ell+1}, \ldots, v_{m_{\ell+1}-1}], W_\ell = [w_{m_\ell}, w_{m_\ell+1}, \ldots, w_{m_{\ell+1}-1}],$$
$$\text{for} \quad \ell = 1, 2, \ldots, k-1.$$

The k-th block is still under construction and the already existing part is denoted by

$$V_k = [v_{m_k}, v_{m_k+1}, \ldots, v_m] \quad \text{and} \quad W_k = [w_{m_k}, w_{m_k+1}, \ldots, w_m].$$

The matrix D_k is now defined as $D_k = W_k^T V_k$. If all blocks have length 1, then we have the standard process and D_k is a 1 by 1 matrix. The subspaces cannot be expanded if D_k is singular. One may expect an unstable computational process if D_k is near singular. To check this, one can compute the smallest singular value of D_k. If D_k is not near-singular then one proceeds with a *regular step* for the expansion of the subspaces. At a regular step the previous block is supposed to be complete and the new regular step defines the start of a new block. For the case that all blocks have dimension 1 this would be a normal bi-Lanczos step, but otherwise we have to take into account that we have to biorthogonalize with respect to the two previous blocks. In general, a regular step for the blocked version looks like

$$\gamma_{m+1} v_{m+1} = A v_m - V_k D_k^{-1} W_k^T A v_m - V_{k-1} D_{k-1}^{-1} W_{k-1}^T A v_m,$$

and

$$\gamma_{m+1} w_{m+1} = A^T w_m - W_k D_k^{-1} V_k^T A^T w_m - W_{k-1} D_{k-1}^{-1} V_{k-1}^T A^T w_m.$$

The vectors v_{m+1} and w_{m+1} are the first columns of the new blocks V_{k+1} and W_{k+1}, respectively.

If a regular step is undesirable, because of near-singularity of the existing block, then we expand the block with something different. Now we have an option of choices: in principle, we can expand by any arbitrary three-term recursion (for instance, a Chebyshev recursion step), in order to maintain a (block) tridiagonal form. One could select new vectors in such a way that the smallest singular vector of the expanded block is as large as possible. In Freund and Nachtigal [1991], it is suggested that one could simply use Av_m and Aw_m and orthogonalize these vectors to the previous block (but not to the current block):

$$\gamma_{m+1} v_{m+1} = A v_m - V_{k-1} D_{k-1}^{-1} W_{k-1}^T A v_m,$$

and

$$\gamma_{m+1} w_{m+1} = A^T w_m - W_{k-1} D_{k-1}^{-1} V_{k-1}^T A^T w_m.$$

Such a step is called an *inner step*.

The look-ahead Lanczos algorithm leads to

$$A\,[V_1, V_2, \ldots, V_m] = [V_1, V_2, \ldots, V_{m+1}] \widetilde{T}_{m+1,m}$$

and

$$A^T\,[W_1, W_2, \ldots, W_m] = [W_1, W_2, \ldots, W_{m+1}]\widetilde{T}_{m+1,m},$$

with a block tridiagonal matrix $\widetilde{T}_{m+1,m}$ (with $m = 1$ block rows and m block columns).

For further details, and for implementation aspects we refer to Freund and Nachtigal [1991] and Freund, Gutknecht and Nachtigal [1993]. A Fortran implementation is available as part of QMRPACK (Freund and Nachtigal [1996]).