UNIVERSITAT POLITÈCNICA DE VALÈNCIA

Scalable Library for Eigenvalue Problem Computations

SLEPc

SLEPc Technical Report **STR-9**
Available at http://slepc.upv.es

Practical Implementation of Harmonic Krylov-Schur

J. E. Roman

Last update: July, 2009 (SLEPc 3.0.0)
Previous updates: –

About SLEPc Technical Reports: These reports are part of the documentation of SLEPc, the *Scalable Library for Eigenvalue Problem Computations*. They are intended to complement the Users Guide by providing technical details that normal users typically do not need to know but may be of interest for more advanced users.

Contents

1 Introduction

Krylov eigensolvers are among the most successful methods for approximating eigenvalues and eigenvectors of large, sparse matrices. It is well known that Krylov subspaces of increasing dimension contain increasingly better approximations of eigenspaces of a given matrix. These approximations can be retrieved by the Rayleigh-Ritz projection procedure. The overall process can be accomplished by means of the Lanczos and Arnoldi algorithms, for the Hermitian and non-Hermitian cases, respectively. The reader is referred to [Stewart, 2001b; van der Vorst, 2002] for details.

Since the convergence of Krylov subspaces to eigenspaces can be very slow, restarting is necessary, especially in the non-Hermitian case because Arnoldi needs to operate on all the available basis vectors. The key point is how to restart without throwing away the currently available spectral information that is relevant to the user. In this respect, two achievements have contributed to getting closer to a satisfactory solution: implicit restart and the Krylov-Schur method. Implicit restart was proposed by Sorensen [1992], and consists in combining the Arnoldi (or Lanczos) procedure with a few steps of the implicitly shifted QR iteration. The result is an algorithm that efficiently computes the Arnoldi reduction that would be obtained with a modified initial vector, on which a polynomial filter has been applied, but without computing this vector explicitly. In this way, at each restart it is possible to filter out the components in the direction of unwanted eigenvectors, thus favoring convergence to the wanted part of the spectrum. The Krylov-Schur method was introduced by Stewart [2001a, 2002] as an alternative to implicit restart that is mathematically equivalent, but easier to implement in a numerically robust manner. A symmetric variant of this process had already been proposed before [Wu and Simon, 2000].

Even with the filtering effect provided by implicit restart, some parts of the spectrum may be very difficult to compute. This is due to the fact that Krylov subspaces tend to converge first to eigenspaces associated to dominant eigenvalues or, in general, to those eigenvalues lying in the periphery of the spectrum. In applications where the sought-after eigenvalues are interior, restarted Krylov eigensolvers will likely fail. One of the reasons is that standard Rayleigh-Ritz may provide 'imposters' [Stewart, 2001b], that is, approximate eigenvalues whose associated vector is a linear combination of some irrelevant eigenvectors. If such vectors are chosen for the restart, then convergence will be damaged. Also, in cases where one is interested in computing righmost eigenvalues that are small relative to the dominant ones, difficulties arise as well because the filtering mechanism is not able to get rid of the components from largest magnitude eigenvalues. Smallest eigenvalues may suffer similar problems. In all these cases, a remedy is required.

One possibility for computing interior eigenvalues is to apply a spectral transformation such as shift-and-invert. In this case, the eigensolver builds a Krylov subspace associated to the shifted and inverted matrix, $(A - \tau I)^{-1}$, so that the eigenvalues closest to the parameter τ now become dominant and well separated. This approach has been used successfully for a long time, especially in the context of symmetric problems [Ericsson and Ruhe, 1980; Grimes *et al.*, 1994]. The main drawback of this technique is that inversion implies solving multiple linear systems within the Arnoldi procedure. The added cost may be prohibitive in some applications, especially for very large problems. Also, factorizing the matrix may be impractical in the context of parallel computation, or even impossible in applications where the matrix is not available explicitly. On the other hand, solving the linear systems with an iterative method is possible [Meerbergen and Roose, 1997], but in practice it may not be appropriate because these systems must be solved very accurately. If sufficient accuracy is not attained, then errors accumulate resulting in a computed subspace that differs from a Krylov subspace. In that case it may be better to resign from pursuing the Krylov character and build a completely different subspace, as in the case of preconditioned eigensolvers, see e.g. [Morgan, 2000b] and references therein.

A different alternative, on which we focus in this report, is harmonic extraction. The concept of harmonic Ritz values was used by Morgan [1991] as a means of achieving a similar effect as the shift-and-invert transformation, but without incurring the high cost associated with the inverted operator. Harmonic Ritz values of a matrix A are the reciprocals of Ritz values of A^{-1} with respect to the subspace $A\mathcal{V}$, where $\mathcal{V}$ is the space of approximants [Paige *et al.*, 1995; Sleijpen and van der Vorst, 1996]. Harmonic Ritz values can also be seen as the roots of quasi-kernel polynomials [Freund, 1992], or as directly related to Lehmann bounds [Paige *et al.*, 1995; Beattie, 1998]. For a general description of harmonic projection techniques, see [Stewart, 2001b; van der Vorst, 2002].

Harmonic Ritz values and vectors are relevant for the iterative solution of linear systems of equations [Morgan, 1995; Goossens and Roose, 1999; Morgan, 2000a, 2002]. Here, we are concerned with using them for computing approximations of eigenvalues and eigenvectors.

Although harmonic Ritz values seem to be very promising for computing smallest or internal eigenvalues, due to their relation to Ritz values of A^{-1}, it turns out that the performance of this technique is very poor compared to the shift-and-invert mapping in terms of convergence, although much cheaper. In fact, examples show that convergence rate is more or less equivalent to standard Ritz values [Paige *et al.*, 1995; Beattie, 1998]. However, when approximating internal eigenvalues, the convergence of harmonic Ritz values is monotonic as opposed to the irregular behaviour of Ritz values, and this can certainly have advantages. In the context of restarted eigensolvers, it can be beneficial to use harmonic Ritz vectors for restarting, as advocated by Morgan and Zeng [1998, 2006]. This is because harmonic Rayleigh-Ritz avoids the appearance of 'imposters' when computing internal eigenvalues. See also [Stewart, 2001b, §4.2]. In this report, we focus on the standard eigenproblem, although harmonic Rayleigh-Ritz can be extended to the generalized and polynomial eigenvalue problems [Stewart, 2001b; Hochstenbach and Sleijpen, 2008].

Another alternative for approximating eigenvalues is the so-called refined extraction, which is indeed compatible with harmonic extraction so that a combined method can be used [Jia, 2002, 2005; Chen and Jia, 2005]. It can be shown that there is a theoretical relation between refined and harmonic extraction [Sleijpen and van den Eshof, 2003].

As mentioned above, using a harmonic Rayleigh-Ritz procedure for obtaining accurate approximations of eigenpairs is especially important when these approximations are employed at the moment of restarting the method. Work by some authors has focussed on Krylov eigensolvers with either explicit restart [Chen and Jia, 2005] or implicit restart [Jia, 2002, 2005; Chen and Lin, 2008], or the Jacobi-Davidson method [Sleijpen and van der Vorst, 1996; Sleijpen *et al.*, 1998]. In this report, we consider the Krylov-Schur method. Our algorithm is equivalent to the one proposed by Morgan [2002] and Morgan and Zeng [2006], but here we consider a formulation entirely in terms of Krylov decompositions as defined by Stewart [2001a, 2002].

The definition of harmonic Ritz value has been generalized by Hochstenbach [2005] to a more general rational function of a matrix. This rational harmonic Ritz approach can be useful for computing rightmost eigenvalues or eigenvalues of structured matrices. The use of harmonic Ritz values has also been extended to the case of computing a partial singular value decomposition (SVD), where it can be useful for computing smallest or internal singular values [Kokiopoulou *et al.*, 2004; Hochstenbach, 2004; Baglama and Reichel, 2005, 2006]. This case will be mentioned also in this report.

In §1.1, we present the main theory of the Krylov-Schur method. Section 2.1 introduces the harmonic projection, and §2.2 discusses how it can be incorporated into the Krylov-Schur algorithm. In §2.3, an alternative path to the same result is presented. Section 2.4 provides some additional details concerning the symmetric case. Throughout this report $\|\cdot\|$ will denote the vector and matrix 2-norm. We will also use the colon notation to denote a range of indices that define a submatrix, for instance $Q_{1:\ell}$ denotes the matrix formed by the first ℓ columns of Q, and $S_{1:\ell,1:\ell}$ denotes the leading principal submatrix of order ℓ of S.

1.1 The Krylov-Schur Method

The Krylov-Schur method is based on so-called Krylov decompositions. Given an $n \times n$ matrix A, a Krylov decomposition of order m is a relation of the form

$$AU = UB + ub^*, \tag{1}$$

where B is a square matrix of order m and the columns of $[U, u]$ are linearly independent. The columns of $[U, u]$ span the space of the decomposition, that can be shown to be a Krylov subspace [Stewart, 2001a, Th. 2.2]. If these columns are orthonormal then the decomposition is said to be orthonormal.

The matrix B is called the Rayleigh quotient of the decomposition. To see why, consider the particular case of an orthonormal Krylov decomposition and premultiply (1) by U^*. Then,

$$U^* A U = B. \tag{2}$$

This fact allows the application of the Rayleigh-Ritz procedure, that is, if ζ is an eigenvalue of B and y is the corresponding eigenvector, then an approximate eigenpair of A can be obtained as $\tilde{\lambda} = \zeta$ (Ritz value) and $\tilde{x} = Uy$ (Ritz vector). Generalizations of this procedure will be considered later on.

If the spaces of two Krylov decompositions are the same, then the decompositions are said to be equivalent. Two classes of transformations maintain equivalence: similarity and translation. A similarity transformation is performed by postmultiplying the decomposition by the inverse of a non-singular matrix W,

$$A(UW^{-1}) = (UW^{-1})(WBW^{-1}) + u(b^* W^{-1}). \tag{3}$$

Note that the new Rayleigh quotient is similar to the original one, and the basis of the space has changed. In a translation transformation, the Rayleigh quotient is modified by a rank-one matrix,

$$AU = U(B + gb^*) + (u - Ug)b^*. \tag{4}$$

These two transformations can be used to build a sequence of equivalent Krylov decompositions with the aim of efficiently extracting the spectral information. In particular, they can be used for making the columns of $[U, u]$ orthogonal, for transforming an arbitrary Krylov decomposition to an equivalent Arnoldi decomposition, or for reducing the Rayleigh quotient to the Schur canonical form. Also, as pointed out by Stewart [2002], translations are relevant for computing harmonic Ritz values, and this will be discussed further in §2.2.

The Krylov-Schur method works with orthonormal Krylov decompositions and applies transformations in order to reduce them to a Krylov-Schur decomposition, in which B is in Schur form. The crux of the method is the observation that if B is in Schur form then the decomposition can be truncated at any point. This allows for a repetive scheme in which the decomposition is truncated and extended again, keeping the relevant spectral information and improving it at each cycle. Consider a Krylov-Schur decomposition written as

$$A \begin{bmatrix} U_1 & U_2 \end{bmatrix} = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} S_{11} & S_{12} \\ 0 & S_{22} \end{bmatrix} + u \begin{bmatrix} b_1^* & b_2^* \end{bmatrix}, \tag{5}$$

then equating the first block-column yields another Krylov-Schur decomposition,

$$AU_1 = U_1 S_{11} + u b_1^*. \tag{6}$$

This is possible because the bottom-left corner of the Rayleigh quotient is zero. An implication of this is that in the case of real arithmetic the Rayleigh quotient is quasi-triangular and truncation should not split any 2×2 diagonal blocks.

Algorithm 1 (Krylov-Schur Method)

1. *Start.* Given an initial vector, build an orthonormal Krylov decomposition of order m, for instance with the Arnoldi algorithm,
$$AU = UB + ub^*. \tag{7}$$

2. *Schur.* Compute a unitary matrix Q_1 that reduces B to Schur form, $S = Q_1^* B Q_1$, obtaining a Krylov-Schur decomposition,
$$AUQ_1 = UQ_1 S + ub^* Q_1. \tag{8}$$

3. *Sort.* Compute a unitary similarity transformation $\tilde{S} = Q_2^* S Q_2$ that sorts the diagonal blocks of the Rayleigh quotient in an appropriate order. Set $Q := Q_1 Q_2$, then
$$AUQ = UQ\tilde{S} + ub^* Q. \tag{9}$$

4. *Truncate.* Proceed as indicated in (5)–(6). Choose an appropriate dimension, ℓ, and explicitly compute $\hat{U} = UQ_{1:\ell}$ and $\hat{b} = Q_{1:\ell}^* b$. Set $\hat{S} := \tilde{S}_{1:\ell,1:\ell}$. The truncated Krylov-Schur decomposition is
$$A\hat{U} = \hat{U}\hat{S} + u\hat{b}^*. \tag{10}$$

5. *Lock.* Check the residual norm estimates and lock converged eigenpairs by setting the corresponding value of $\hat{b}$ to zero. If satisfied, stop.

6. *Extend.* Write (10) as
$$A\hat{U} = \left[\begin{array}{cc} \hat{U} & u \end{array}\right] \left[\begin{array}{c} \hat{S} \\ \hat{b}^* \end{array}\right], \tag{11}$$
then extend to a Krylov decomposition of order m by means of the Arnoldi algorithm taking u as initial vector. The computed quantities are appended as new columns to both sides of the equation to get a new U and B, as in (7). Go to step 2.

A few comments are in order to complement the description of Algorithm 1. In step 1, the Arnoldi algorithm is typically used (or a variant of Lanczos in the symmetric case). Remember that the Arnoldi method sequentially computes the columns of U and B, as follows. The $(k+1)$th column of U, u_{k+1}, is computed by orthogonalizing Au_k with respect to the previous k columns, and $b_{i,k}$ are set to the corresponding orthogonalization coefficients, for $i = 1, \ldots, k$. Finally, u_{k+1} is normalized and $b_{k+1,k}$ is set to the norm prior to normalization. After m steps of the algorithm, U has m columns and vector u in (7) represents the last computed vector. Vector b is a multiple of e_m, and this can be exploited to get minor savings.

The Arnoldi method is also used in step 6 of Algorithm 1, in this case taking u as the initial vector and considering also the columns of $[\hat{U}, u]$ in all subsequent orthogonalization operations. Note that u is orthogonal to the columns of $\hat{U}$ since $u^*\hat{U} = u^*UQ_{1:\ell}$ and $u^*U = 0$ by construction.

For the convergence test in step 5, it is necessary to compute all eigenvalues and eigenvectors of $\hat{S}$. For a given eigenpair $(\tilde{\lambda}, y)$, the corresponding Ritz pair is $(\tilde{\lambda}, \hat{U}y)$, and the associated residual norm estimate is

$$\|A\hat{U}y - \tilde{\lambda}\hat{U}y\| = \|(\hat{U}\hat{S} + u\hat{b}^*)y - \tilde{\lambda}\hat{U}y\| = \|u\hat{b}^*y\| = |\hat{b}^*y|. \tag{12}$$

In step 2, the diagonal blocks of the Schur Rayleigh quotient are sorted by moving them up and down by means of unitary similarities [Bai and Demmel, 1993]. If they are sorted in descending or ascending order, then the algorithm computes the largest or smallest Ritz values, respectively. In other words, step 2 moves the wanted part of the spectrum to the upper-left corner of $\tilde{S}$, so that truncation keeps these Ritz values together with the associated (Schur) Ritz vectors. For computing interior eigenvalues, a different sorting criterion could be used, such as the proximity to a given target value. However, this approach is not useful because Ritz vectors associated to interior eigenvalues are very poorly approximated and are not good for restart. As discussed in §1, harmonic Ritz vectors are preferred instead.

To conclude this section, we point out the simplifications that can be made in the case of a symmetric problem. If A is Hermitian, then B is real, symmetric and tridiagonal. In step 2, S is diagonal and therefore its diagonal entries are the Ritz values and the columns of Q_1 are the corresponding Ritz vectors. In step 3, Q_2 is simply a permutation matrix. Finally, the residual norm estimates are equal to the absolute values of the elements of $\hat{b}$.

2 Harmonic Extraction

In this section, we describe the idea of harmonic extraction and consider several ways of incorporating it into the Krylov-Schur eigensolver.

2.1 Harmonic Rayleigh-Ritz

As mentioned in the introduction, harmonic extraction intends to emulate the effect of the shift-and-invert spectral transformation. If we are interested in computing those eigenvalues of a matrix A that are closest to a given target value, τ, then the shift-and-invert approach consists in computing the largest eigenvalues of the eigenproblem

$$(A - \tau I)^{-1}x = \theta x. \tag{13}$$

The relation between the eigenvalues of A, λ, and the transformed eigenvalues, θ, is $\theta = (\lambda - \tau)^{-1}$. Note that the eigenvectors are not altered. Once the transformed eigenvalues have been computed, the original eigenvalues can be recovered by a simple operation,

$$\lambda = \theta^{-1} + \tau. \tag{14}$$

The effect of the spectral transformation is that eigenvalues closest to the target become dominant and with a better separation than in the original spectrum. Therefore, this approach is very effective in computing interior eigenpairs. However, the drawback is the high cost associated with the inverted operator $(A - \tau I)^{-1}$, that is usually handled by computing a triangular factorization and performing backsolves within the eigensolver.

The goal of harmonic Rayleigh-Ritz is to achieve a similar effect as shift-and-invert, but avoiding the inverted operator. For this, it is necessary to work with two bases, one that spans the subspace of approximants for matrix A, U, and another one that is defined as

$$V := (A - \tau I)U. \tag{15}$$

The trick is to try to cancel the inverted operator with the $(A - \tau I)$ factor.

Suppose we have built a Krylov decomposition associated with matrix A, (1). Then, the decomposition may be shifted as

$$(A - \tau I)U = U(B - \tau I) + ub^*. \tag{16}$$

Using the definition of V in (15), the above relation can be written as

$$V = U(B - \tau I) + ub^*. \tag{17}$$

Now, U can be isolated by postmultiplying the equation by the inverse of $(B - \tau I)$, provided it exists,

$$U = V(B - \tau I)^{-1} - ug^*, \tag{18}$$

with

$$g := (B - \tau I)^{-*}b. \tag{19}$$

Equation (18) can be seen as a Krylov decomposition of the shift-and-invert operator, since $U = (A - \tau I)^{-1}V$. Note that the basis of this decomposition, V, is not orthonormal, in contrast to the basis of the original decomposition, U. At this point, one may consider using a combination of similarity and translation transformations so that the columns of $[V, -u]$ become orthonormal. That would lead to a Rayleigh quotient whose eigenvalues would be related to the harmonic Ritz values of A with respect to τ. Since explicit orthonormalization may have an excessive cost, we consider alternative ways of manipulating this new decomposition to obtain harmonic Ritz values cheaply. This will be achieved by an appropriate projection.

In §1.1, we mentioned how the Rayleigh quotient in (2) can be obtained by premultiplying the Krylov decomposition by U^*. In terms of a projection method, this represents an orthogonal projection in which the approximate eigenvectors (Ritz vectors) are chosen from the subspace spanned by the columns of U, $\tilde{x} = Uy$, and the residual vectors corresponding to each Ritz pair $(\tilde{\lambda}, Uy)$ must be orthogonal to the same subspace, that is, they have to satisfy the Galerkin condition

$$U^*(AUy - \tilde{\lambda}Uy) = 0. \tag{20}$$

This yields the projected problem

$$U^*AUy = \tilde{\lambda}U^*Uy, \tag{21}$$

that simplifies to $By = \tilde{\lambda}y$, B as in (2), if U has orthonormal columns.

Now suppose that a Krylov decomposition (1) is available and we would like to apply the orthogonal projection method to matrix $(A-\tau I)^{-1}$ for Ritz vectors contained in the subspace spanned by the columns of V. In that case, if the Ritz pair is denoted as $(\tilde{\theta}, Vz)$, the Galerkin condition can be written as

$$V^*((A-\tau I)^{-1}Vz - \tilde{\theta}Vz) = 0, \tag{22}$$

or

$$V^*(A-\tau I)^{-1}Vz = \tilde{\theta}V^*Vz. \tag{23}$$

The next step is to write this equation in terms of the Krylov basis U by using the definition of V in (15),

$$U^*(A-\tau I)^*Uz = \tilde{\theta}U^*(A-\tau I)^*(A-\tau I)Uz. \tag{24}$$

From (16), this is equivalent to

$$\begin{aligned}(B-\tau I)^*z &= \tilde{\theta}(U(B-\tau I)+ub^*)^*(U(B-\tau I)+ub^*)z && (25)\\ &= \tilde{\theta}((B-\tau I)^*(B-\tau I)+bb^*)z, && (26)\end{aligned}$$

that is, the projected problem is a generalized eigenvalue problem in which the matrix on the right hand side is symmetric and positive definite. The values $\tilde{\theta}$ are Ritz values of $(A-\tau I)^{-1}$ with respect to V, or alternatively, the values $\tilde{\theta}^{-1}+\tau$ are harmonic Ritz values of A.

Equation (26) can also be formulated as a standard eigenvalue problem, as follows. First, pre-multiply the equation by the inverse of $(B-\tau I)^*$, to get

$$z = \tilde{\theta}((B-\tau I)+(B-\tau I)^{-*}bb^*)z, \tag{27}$$

or, equivalently, using the definition of g in (19),

$$z = \tilde{\theta}((B-\tau I)+gb^*)z. \tag{28}$$

Rearranging terms, the projected problem can be written as

$$(B+gb^*)z = (\tilde{\theta}^{-1}+\tau)z, \tag{29}$$

that is, the eigenvalues of matrix $B+gb^*$ are the harmonic Ritz values of A. The corresponding harmonic Ritz vectors are equal to Vz, or $(A-\tau I)Uz$. For practical purposes, the vector Uz can be taken instead.

The previous paragraphs show that harmonic Ritz pairs can be easily extracted from a Krylov decomposition of A. The key question is how this simple idea fits in the Krylov-Schur restart mechanism. But before addressing that issue, let us show that the same result can be obtained from a different perspective.

Instead of thinking in terms of an orthogonal projection, we could consider an oblique projection method applied to A. In an oblique projection, approximate eigenvectors are chosen to be in the range of U, as before, but in this case the residuals have to satisfy the Petrov-Galerkin condition,

$$V^*(AUy - \tilde{\lambda}Uy) = 0, \tag{30}$$

where we will choose the left subspace, V, to be the same as defined before in (15). The resulting projected eigenproblem is then

$$V^*AUy = \tilde{\lambda}V^*Uy. \tag{31}$$

Assuming we have a Krylov decomposition (1) available, then

$$V^*(UB + ub^*)y = \tilde{\lambda}V^*Uy, \tag{32}$$

or

$$(V^*UB + V^*ub^*)y = \tilde{\lambda}V^*Uy. \tag{33}$$

As before, we can reduce the problem to a standard eigenproblem, in this case pre-multiplying by the inverse of V^*U,

$$(B + (V^*U)^{-1}V^*ub^*)y = \tilde{\lambda}y. \tag{34}$$

At this point, using the definition of V in (15) gives the expression

$$(B + (U^*(A - \tau I)^*U)^{-1}U^*(A - \tau I)^*ub^*)y = \tilde{\lambda}y. \tag{35}$$

So it only remains to notice that $U^*(A - \tau I)^*U = (B - \tau I)^*$ and $U^*(A - \tau I)^*u = b$, resulting in the eigenproblem

$$(B + (B - \tau I)^{-*}bb^*)y = \tilde{\lambda}y, \tag{36}$$

in which the matrix is the same as in (29).

2.2 Harmonic Krylov-Schur

In the Krylov-Schur algorithm described in §1.1, the restarting mechanism is so effective because the truncation step, (5)–(6), satisfies two properties:

1. the Ritz values kept in the Rayleigh quotient of the truncated decomposition are those of interest, and

2. the space of the truncated decomposition contains the Ritz vectors associated to those particular Ritz values.

Analogously, if we use a Krylov decomposition of A to approximate harmonic Ritz values, the Krylov-Schur algorithm should be adapted in such a way that in the truncation step (i) the harmonic Ritz values of interest are kept in the Rayleigh quotient of the truncated decomposition, and (ii) the space of the truncated decomposition contains the wanted harmonic Ritz vectors.

Algorithm 2 describes the required modifications. The main idea is to apply a translation transformation (4) to the original Krylov decomposition, in order to express the Rayleigh quotient as required by (29), then perform the truncation on this Rayleigh quotient.

Algorithm 2 (Harmonic Krylov-Schur Method)

1. *Start.* Given an initial vector, build an orthonormal Krylov decomposition of order m, for instance with the Arnoldi algorithm,

$$AU = UB + ub^*. \tag{37}$$

2. *Translate.* Compute vector g as in (19) and compute a rank-one modification of B, as $\tilde{B} = B + gb^*$. This effects a translation on the original decomposition, resulting in

$$AU = U\tilde{B} + \tilde{u}b^*, \tag{38}$$

with $\tilde{u} = u - Ug$.

3. *Schur.* Compute a unitary matrix Q_1 that reduces $\tilde{B}$ to Schur form, $S = Q_1^*\tilde{B}Q_1$, obtaining a Krylov-Schur decomposition,

$$AUQ_1 = UQ_1S + \tilde{u}b^*Q_1. \tag{39}$$

4. *Sort.* Compute a unitary similarity transformation $\tilde{S} = Q_2^*SQ_2$ that sorts the diagonal blocks of the Rayleigh quotient in an appropriate order. Set $Q := Q_1Q_2$, then

$$AUQ = UQ\tilde{S} + \tilde{u}b^*Q. \tag{40}$$

5. *Truncate.* Proceed as indicated in (5)–(6). Choose an appropriate dimension, ℓ, and explicitly compute $\hat{U} = UQ_{1:\ell}$ and $\hat{b} = Q_{1:\ell}^*b$. Set $\hat{S} := \tilde{S}_{1:\ell,1:\ell}$. The truncated Krylov-Schur decomposition is

$$A\hat{U} = \hat{U}\hat{S} + \tilde{u}\hat{b}^*. \tag{41}$$

6. *Lock.* Check the residual norm estimates and lock converged eigenpairs by setting the corresponding value of $\hat{b}$ to zero. If satisfied, stop.

7. *Recover.* Recuperate orthonormality of the truncated decomposition by performing another translation, with

$$\gamma\hat{u} = \tilde{u} - \hat{U}\hat{g}, \qquad \hat{g} = -Q_{1:\ell}^*g, \tag{42}$$

$$\hat{B} = \hat{S} + \hat{g}\hat{b}^*. \tag{43}$$

8. *Extend.* Write the resulting Krylov decomposition as

$$A\hat{U} = \begin{bmatrix} \hat{U} & \hat{u} \end{bmatrix} \begin{bmatrix} \hat{B} \\ \gamma\hat{b}^* \end{bmatrix}, \tag{44}$$

then extend to a Krylov decomposition of order m by means of the Arnoldi algorithm taking $\hat{u}$ as initial vector. The computed quantities are appended as new columns to both sides of the equation to get a new U and B, as in (37). Go to step 2.

The most important difference of Algorithm 2 with respect to Algorithm 1 is the new step labelled *Translate*. Note that after this step the resulting decomposition is no longer orthonormal, because $U^*\tilde{u} \neq 0$. The aim of the other added step, labelled *Recover*, is to regain the orthonormal character. Note that vector $\hat{u}$ computed in (42) is orthogonal to the columns of $\hat{U}$, since

$$\hat{U}^*(\gamma\hat{u}) = \hat{U}^*(\tilde{u} - \hat{U}\hat{g}) = \hat{U}^*\tilde{u} - \hat{g} = Q_{1:\ell}^* U^*(u - Ug) + Q_{1:\ell}^* g = 0. \tag{45}$$

An important aspect from the perspective of a practical implementation is that some operations can be saved because vector $\tilde{u}$ is not really being used in steps 3 to 5. Thus, it is not necessary to compute it explicitly, if we reformulate (42) as

$$\gamma\hat{u} = (u - Ug) - \hat{U}\hat{g} = u - Ug + UQ_{1:\ell}Q_{1:\ell}^* g = u - U\tilde{g}, \tag{46}$$

where $\tilde{g} = (I - Q_{1:\ell}Q_{1:\ell}^*)g$, that is, the vector resulting from orthogonalizing g against the first ℓ columns of Q.

In step 6 (*Lock*), vector $\tilde{u}$ is implicitly used in the computation of the residual norm estimates (12), since in this case it has not been normalized. However, only its norm is necessary and this can be cheaply computed by noting the fact that $\tilde{u} = u - Ug$ with $u \perp U$, so $\tilde{u}$, u and $-Ug$ form a right-angled triangle and their Euclidean lengths are related by the Pythagorean theorem. Thus,

$$\|\tilde{u}\| = \sqrt{1 + \|g\|^2}. \tag{47}$$

Analogously, from (46) we get a similar expression for γ,

$$\gamma = \sqrt{1 + \|\tilde{g}\|^2}. \tag{48}$$

Furthermore, observe that the *Recover* step is in fact undoing the first translation, because (43) can be rewritten as

$$\hat{B} = Q_{1:\ell}^*\tilde{B}Q_{1:\ell} - Q_{1:\ell}^* g b^* Q_{1:\ell} = Q_{1:\ell}^*(\tilde{B} - gb^*)Q_{1:\ell} = Q_{1:\ell}^* B Q_{1:\ell}, \tag{49}$$

so $\hat{B}$ is a section of the original Rayleigh quotient B. This makes sense because the extension (step 8) should build a non-translated Krylov decomposition before returning to step 2 of the algorithm.

With the above observations, it is possible to write a cycle of Algorithm 2 in a more compact way.

Algorithm 3 (Harmonic Krylov-Schur cycle, compact form)

Input: the orthonormal Krylov decomposition $AU = UB + ub^*$ of (37)
Output: the orthonormal Krylov decomposition $A\hat{U} = \hat{U}\hat{B} + \hat{u}(\gamma\hat{b})^*$ of (44)

$g = (B - \tau I)^{-*}b$
$\tilde{B} = B + gb^*$
$[Q, \tilde{S}] = \mathtt{schur\_sorted}(\tilde{B}, \tau)$

$\hat{b} = Q_{1:\ell}^* b$
$\hat{g} = -Q_{1:\ell}^* g$
$\|\tilde{u}\| = \sqrt{1 + \|g\|^2}$
Test convergence, using $\|\tilde{u}\|$ and $\hat{b}$
$\hat{B} = \tilde{S}_{1:\ell,1:\ell} + \hat{g}\hat{b}^*$
$\tilde{g} = (I - Q_{1:\ell}Q_{1:\ell}^*)g$
$\hat{u} = u - U\tilde{g}$
$\gamma = \sqrt{1 + \|\tilde{g}\|^2}$
$\hat{U} = UQ_{1:\ell}$

In the algorithm, the operation `schur_sorted` represents the computation of the (real) Schur form, together with the appropriate ordering of the diagonal blocks [Bai and Demmel, 1993]. Since the diagonal blocks contain harmonic Ritz values, the sorting criterion must be the distance to the target value, τ.

Another observation is that if any of the eigenvalues of B is close to τ, then the computed vector g will have large errors, due to ill conditioning of $(B - \tau I)^{-*}$. However, these errors seem to be benign, as reported by experiments [Morgan and Zeng, 1998, 2006].

A further comment regarding errors is that harmonic Rayleigh-Ritz is better at providing good approximate eigenvectors for internal eigenvalues, but harmonic Ritz values themselves may be inaccurate. This is the reason why in the literature (e.g. [Morgan and Zeng, 1998]) it is recommended to compute the Rayleigh quotient of harmonic Ritz vectors to get the so-called ρ-values, which may be better approximations than harmonic Ritz values. These ρ-values can be computed cheaply as

$$\rho_i = \frac{z_i^* U^* A U z_i}{z_i^* U^* U z_i} = \tau + \frac{z_i^* U^* (A - \tau I) U z_i}{z_i^* z_i} = \tau + \frac{z_i^* (B - \tau I) z_i}{z_i^* z_i} = \frac{z_i^* B z_i}{z_i^* z_i}. \tag{50}$$

However, this approach cannot be directly adapted to Algorithm 2 since it uses (real) Schur forms instead of eigenvectors. On the other hand, the algorithm is based on keeping approximate eigenvectors of wanted eigenvalues, but without need to compute those eigenvalues accurately (eigenvalues are only used when sorting). This contrasts with other algorithms that make explicit use of approximate eigenvalues (e.g. as exact shifts in implicit restart [Jia, 2002]).

2.3 Alternative Derivation of Harmonic Krylov-Schur

It is possible to formulate the harmonic Krylov-Schur method in a slightly different way, by realizing that the two types of transformations, similarity (3) and translation (4), can be combined in a single transformation, as already pointed out in [Stewart, 2001a]. Suppose that a matrix

$$\hat{W} := \begin{bmatrix} W & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} I & g \\ 0 & \mu \end{bmatrix} = \begin{bmatrix} W & Wg \\ 0 & \mu \end{bmatrix} \tag{51}$$

is nonsingular, so its inverse is

$$\hat{W}^{-1} = \begin{bmatrix} W^{-1} & -\mu^{-1}g \\ 0 & \mu^{-1} \end{bmatrix}. \tag{52}$$

Write the Krylov decomposition (1) as

$$AU = \begin{bmatrix} U & u \end{bmatrix} \begin{bmatrix} B \\ b^* \end{bmatrix}, \tag{53}$$

then the combined transformation is effected by

$$AUW^{-1} = \begin{bmatrix} U & u \end{bmatrix} \hat{W}^{-1}\hat{W} \begin{bmatrix} B \\ b^* \end{bmatrix} W^{-1}, \tag{54}$$

$$AUW^{-1} = \begin{bmatrix} UW^{-1} & \mu^{-1}(u - Ug) \end{bmatrix} \begin{bmatrix} W(B + gb^*)W^{-1} \\ \mu b^* W^{-1} \end{bmatrix}. \tag{55}$$

Note that this transformation is equivalent to a translation followed by a similarity (the reverse order can be achieved by swapping the factors in the definition of $\hat{W}$). The value of μ is arbitrary and can be chosen so that vector $\mu^{-1}(u - Ug)$ has unit norm, if desired. If we assume that $[U, u]$ has orthonormal columns, this is equivalent to $\mu = \sqrt{1 + \|g\|^2}$.

Here, we are interested in a unitary transformation, so that the basis of the space has orthonormal columns after the transformation. Consider the matrix

$$\hat{Q} = \begin{bmatrix} Q^* & Q^*g \\ 0 & \mu \end{bmatrix} \tag{56}$$

where Q is unitary. The only possibility for $\hat{Q}$ to be unitary is $Q^*g = 0$ and $\mu = 1$. In this case, (55) turns into

$$AUQ = \begin{bmatrix} UQ & u \end{bmatrix} \begin{bmatrix} Q^*BQ \\ b^*Q \end{bmatrix}, \tag{57}$$

which is the simple similarity transformation used by the standard Krylov-Schur algorithm.

If we allow $\hat{Q}$ to be non-unitary, then the combined transformation (55) is

$$AUQ = \begin{bmatrix} UQ & \mu^{-1}(u - Ug) \end{bmatrix} \begin{bmatrix} Q^*(B + gb^*)Q \\ \mu b^* Q \end{bmatrix}, \tag{58}$$

where the resulting basis is not orthogonal in general because of the last vector, $\mu^{-1}(u - Ug)$. Note that in this case the condition $\mu = \sqrt{1 + \|g\|^2}$ is equivalent to the last column of $\hat{Q}$ having unit norm. The transformation could be reversed by applying another transformation with the inverse of $\hat{Q}$,

$$\hat{Q}^{-1} = \begin{bmatrix} Q & -\mu^{-1}g \\ 0 & \mu^{-1} \end{bmatrix}, \tag{59}$$

which can be thought of as a translation with $-Q^*g$ followed by a similarity with Q, or alternatively as a similarity with Q followed by a translation with $-g$. In the harmonic Krylov-Schur algorithm, only the translation is reversed, so the corresponding transformation would be

$$\bar{Q} = \begin{bmatrix} I & -\mu^{-1}Q^*g \\ 0 & \mu^{-1} \end{bmatrix}. \tag{60}$$

Now, we introduce the truncation step. Recall that in harmonic Krylov-Schur, we choose Q in such a way that a zero block is produced in the lower left corner of the projected matrix, i.e. $Q^*_{\ell+1:m}(B + gb^*)Q_{1:\ell} = 0$, so that we can eliminate the part of the decomposition corresponding to columns $\ell + 1 : m$ of the basis. But this can be done only *before* the last transformation. The accumulated transformation is the non-square matrix

$$\tilde{Q} = \begin{bmatrix} I_\ell & -\mu^{-1}Q^*_{1:\ell}g \\ 0 & \gamma\mu^{-1} \end{bmatrix} \begin{bmatrix} Q^*_{1:\ell} & Q^*_{1:\ell}g \\ 0 & \mu \end{bmatrix} = \begin{bmatrix} Q^*_{1:\ell} & 0 \\ 0 & \gamma \end{bmatrix}, \tag{61}$$

where the net effect is a truncated similarity. Note that an additional scaling factor γ has been introduced so that the last column of the resulting basis in normalized. The decomposition before the last translation is

$$AUQ_{1:\ell} = \begin{bmatrix} UQ_{1:\ell} & \mu^{-1}(u - Ug) \end{bmatrix} \begin{bmatrix} Q^*_{1:\ell}(B + gb^*)Q_{1:\ell} \\ \mu b^*Q_{1:\ell} \end{bmatrix}, \tag{62}$$

and, using the notation of Algorithm 3, the final one is

$$AUQ_{1:\ell} = \begin{bmatrix} UQ_{1:\ell} & \gamma^{-1}(u - U\tilde{g}) \end{bmatrix} \begin{bmatrix} Q^*_{1:\ell}BQ_{1:\ell} \\ \gamma\hat{b}^* \end{bmatrix}. \tag{63}$$

Now, starting from a shifted Krylov decomposition (16),

$$(A - \tau I)U = \begin{bmatrix} U & u \end{bmatrix} \begin{bmatrix} B - \tau I \\ b^* \end{bmatrix}, \tag{64}$$

and applying a transformation defined by the matrix

$$\hat{G} := \begin{bmatrix} (B - \tau I)^* & (B - \tau I)^*g \\ 0 & \mu \end{bmatrix} = \begin{bmatrix} (B - \tau I)^* & b \\ 0 & \mu \end{bmatrix} \tag{65}$$

with inverse

$$\hat{G}^{-1} = \begin{bmatrix} (B - \tau I)^{-*} & -\mu^{-1}g \\ 0 & \mu^{-1} \end{bmatrix}, \tag{66}$$

the resulting Krylov decomposition is

$$(A - \tau I)U(B - \tau I)^{-*} = \begin{bmatrix} U(B - \tau I)^{-*} & \mu^{-1}(u - Ug) \end{bmatrix} \begin{bmatrix} M \\ \mu b^*(B - \tau I)^{-*} \end{bmatrix}, \tag{67}$$

where the Rayleigh quotient is

$$\begin{aligned} M &:= (B-\tau I)^*(B-\tau I+gb^*)(B-\tau I)^{-*} && (68)\\ &= ((B-\tau I)^*(B-\tau I)+bb^*)(B-\tau I)^{-*}. && (69)\end{aligned}$$

Suppose $(\tilde{\theta}^{-1}, w)$ is an eigenpair of M, so that $Mw = \tilde{\theta}^{-1}w$ or

$$((B-\tau I)^*(B-\tau I)+bb^*)z = \tilde{\theta}^{-1}(B-\tau I)^*z, \tag{70}$$

with $z = (B-\tau I)^{-*}w$. Then $(\tilde{\theta}, z)$ satisfy the eigenvalue equation (26). Let Q_M be an orthonormal matrix such that $Q_M^* M Q_M = T_M$ is upper (quasi-)triangular with the eigenvalues sorted in the appropriate order (smallest $\tilde{\theta}^{-1}$ in the upper left corner), then a similarity transformation on decomposition (67) results in

$$(A-\tau I)UQ_Z = \left[\begin{array}{cc} UQ_Z & \mu^{-1}(u-Ug)\end{array}\right]\left[\begin{array}{c} T_M \\ \mu b^* Q_Z\end{array}\right], \tag{71}$$

with $Q_Z = (B-\tau I)^{-*}Q_M$. Unfortunately, Q_Z is not an orthonormal matrix in general, so this approach has no practical interest because an explicit orthonormalization of the new basis would be required.

2.4 The Symmetric Case and the SVD

Now we turn our attention to the case of a complex Hermitian (or real symmetric) problem matrix, $A = A^*$. Our goal is to determine whether the algorithm described in §2.2 can be tailored to exploit that feature. At the end of §1.1, we already pointed out some particularities of Krylov-Schur in the symmetric case, such as the special structure of the Rayleigh quotient B.

If $A = A^*$ and assuming that τ is a real scalar, it holds that $(B-\tau I)^* = B-\tau I$. Then, matrix M in (68) is the product of a symmetric positive definite matrix, $((B-\tau I)^*(B-\tau I)+bb^*)$, and a Hermitian matrix. Thus, M is a diagonalizable matrix, all of whose eigenvalues are real, [Horn and Johnson, 1985, Th. 7.6.3]. However, in general M is not unitarily diagonalizable. A practical procedure for computing the eigenpairs, $Mw = \tilde{\theta}^{-1}w$, could be to compute the following QR decomposition,

$$\left[\begin{array}{c} B-\tau I \\ b^*\end{array}\right] = QR, \tag{72}$$

where R is a square matrix of order m and Q is an $(m+1)\times m$ matrix with orthonormal columns. Then $M = R^*R\,(B-\tau I)^{-*}$, so it is possible to solve for

$$R^{-*}(B-\tau I)R^{-1}z = \tilde{\theta}z, \qquad z = R^{-*}w. \tag{73}$$

In this way, it would be possible to use a standard dense symmetric eigensolver for the projected eigenproblem. But again, this procedure has the shortcoming that the resulting transformation matrix is non-orthogonal. It seems that there is no way to exploit the symmetry of the Rayleigh quotient and maintain orthogonality when updating the basis U.

Therefore, one should desist from treating the projected problem as a symmetric one and stick to the general procedure of (29). Still, for a projected problem of large dimension, one may consider exploiting the rank-1 perturbation structure of $(B + gb^*)$, with an $O(m^2)$ complexity algorithm as proposed in [Eidelman *et al.*, 2008]. This will not be considered in this report.

In order to complement the discussion of symmetric problems, we now include a few comments regarding the computation of the singular value decomposition (SVD). In the context of Lanczos bidiagonalization for the computation of a partial SVD, it is convenient to use a thick-restart variant for enhanced convergence, as proposed by Baglama and Reichel [2005]; see also [Baglama and Reichel, 2006; Hernandez *et al.*, 2008]. Thick-restart can be seen as a customized version of the Krylov-Schur restarting mechanism tailored for symmetric problems. In [Baglama and Reichel, 2005], the authors already propose to augment the subspace with harmonic Ritz vectors when computing the smallest singular values. Thus, it would be rather straightforward to use the harmonic Krylov-Schur algorithm of §2.2 in this case.

3 The SLEPc Implementation

The harmonic Krylov-Schur method as expressed in Algorithm 3 has been added to SLEPc in version 3.0.0.

3.1 User Options

The user options are not specific of the Krylov-Schur solver. In order to use harmonic extraction, the value `EPS_HARMONIC` must be explicitly set with

```
EPSSetExtraction(EPS eps,EPSExtraction extr);
```

Also, the target value must be provided as well (default is $\tau = 0$), with

```
EPSSetTarget(EPS eps,PetscScalar target);
```

Note that a complex target can be used only if SLEPc has been configured with complex scalars.

A command line example would be:

```
ex5 -m 45 -eps_harmonic -eps_target 0.8 -eps_ncv 60
```

The example computes the eigenvalue closest to $\tau = 0.8$ of a real non-symmetric matrix of order 1035. Note that `ncv` has been set to a larger value than would be necessary for computing the largest magnitude eigenvalues. In general, users should expect a much slower convergence when computing interior eigenvalues compared to extreme eigenvalues. Increasing the value of `ncv` may help.

3.2 Known Issues and Applicability

The following comments have to be taken into account:

- The implementation does not exploit symmetry, that is, for a problem type `EPS_HEP` the solver will proceed as if the problem was non-Hermitian.
- In principle, harmonic extraction could be used for generalized eigenproblems, but there is no point in doing this because internally SLEPc's Krylov-Schur for generalized problems operates with matrix $B^{-1}A$, and therefore it would be more convenient to use a shift-and-invert `ST` with $\sigma = \tau$.
- Harmonic extraction works for both real and complex scalars.
- The `SVD` class does not provide support for harmonic extraction yet.

References

Baglama, J. and L. Reichel (2005). Augmented Implicitly Restarted Lanczos Bidiagonalization Methods. *SIAM J. Sci. Comput.*, 27(1):19–42.

Baglama, J. and L. Reichel (2006). Restarted Block Lanczos Bidiagonalization Methods. *Numer. Algorithms*, 43(3):251–272.

Bai, Z. and J. W. Demmel (1993). On Swapping Diagonal Blocks in Real Schur Form. *Linear Algebra Appl.*, 186:73–95.

Beattie, C. (1998). Harmonic Ritz and Lehmann Bounds. *Electron. Trans. Numer. Anal.*, 7:18–39.

Chen, G. and Z. Jia (2005). A Refined Harmonic Rayleigh-Ritz Procedure and an Explicitly Restarted Refined Harmonic Arnoldi Algorithm. *Math. Comput. Model.*, 41:615–627.

Chen, G. and J. Lin (2008). A New Shift Scheme for the Harmonic Arnoldi Method. *Math. Comput. Model.*, 48(11-12):1701–1707.

Eidelman, Y., L. Gemignani, and I. Gohberg (2008). Efficient Eigenvalue Computation for Quasiseparable Hermitian Matrices under Low Rank Perturbations. *Numer. Algorithms*, 47(3):253–273.

Ericsson, T. and A. Ruhe (1980). The Spectral Transformation Lanczos Method for the Numerical Solution of Large Sparse Generalized Symmetric Eigenvalue Problems. *Math. Comp.*, 35(152):1251–1268.

Freund, R. W. (1992). Quasi-kernel Polynomials and their Use in non-Hermitian Matrix Iterations. *J. Comput. Appl. Math.*, 43:135–158.

Goossens, S. and D. Roose (1999). Ritz and Harmonic Ritz Values and the Convergence of FOM and GMRES. *Numer. Linear Algebra Appl.*, 6(4):281–293.

Grimes, R. G., J. G. Lewis, and H. D. Simon (1994). A Shifted Block Lanczos Algorithm for Solving Sparse Symmetric Generalized Eigenproblems. *SIAM J. Matrix Anal. Appl.*, 15(1):228–272.

Hernandez, V., J. E. Roman, and A. Tomas (2008). A Robust and Efficient Parallel SVD Solver Based on Restarted Lanczos Bidiagonalization. *Electron. Trans. Numer. Anal.*, 31:68–85.

Hochstenbach, M. E. (2004). Harmonic and Refined Extraction Methods for the Singular Value Problem, with Applications in Least Squares Problems. *BIT*, 44(4):721–754.

Hochstenbach, M. E. (2005). Generalizations of Harmonic and Refined Rayleigh-Ritz. *Electron. Trans. Numer. Anal.*, 20:235–252.

Hochstenbach, M. E. and G. L. G. Sleijpen (2008). Harmonic and Refined Rayleigh-Ritz for the Polynomial Eigenvalue Problem. *Numer. Linear Algebra Appl.*, 15(1):35–54.

Horn, R. A. and C. R. Johnson (1985). *Matrix Analysis*. Cambridge University Press.

Jia, Z. (2002). The Refined Harmonic Arnoldi Method and an Implicitly Restarted Refined Algorithm for Computing Interior Eigenpairs of Large Matrices. *App. Numer. Math.*, 42(4):489–512.

Jia, Z. (2005). The Convergence of Harmonic Ritz Values, Harmonic Ritz Vectors and Refined Harmonic Ritz Vectors. *Math. Comp.*, 74(251):1441–1456.

Kokiopoulou, E., C. Bekas, and E. Gallopoulos (2004). Computing Smallest Singular Triplets with Implicitly Restarted Lanczos Bidiagonalization. *App. Numer. Math.*, 49(1):39–61.

Meerbergen, K. and D. Roose (1997). The Restarted Arnoldi Method Applied to Iterative Linear System Solvers for the Computation of Rightmost Eigenvalues. *SIAM J. Matrix Anal. Appl.*, 18(1):1–20.

Morgan, R. B. (1991). Computing Interior Eigenvalues of Large Matrices. *Linear Algebra Appl.*, 154–156:289–309.

Morgan, R. B. (1995). A Restarted GMRES Method Augmented with Eigenvectors. *SIAM J. Matrix Anal. Appl.*, 16(4):1154–1171.

Morgan, R. B. (2000a). Implicitly Restarted GMRES and Arnoldi Methods for Nonsymmetric Systems of Equations. *SIAM J. Matrix Anal. Appl.*, 21(4):1112–1135.

Morgan, R. B. (2000b). Preconditioning eigenvalues and some comparison of solvers. *J. Comput. Appl. Math.*, 123(1-2):101–115.

Morgan, R. B. (2002). GMRES with Deflated Restarting. *SIAM J. Sci. Comput.*, 24(1):20–37.

Morgan, R. B. and M. Zeng (1998). Harmonic Projection Methods for Large Non-Symmetric Eigenvalue Problems. *Numer. Linear Algebra Appl.*, 5(1):33–55.

Morgan, R. B. and M. Zeng (2006). A Harmonic Restarted Arnoldi Algorithm for Calculating Eigenvalues and Determining Multiplicity. *Linear Algebra Appl.*, 415(1):96–113.

Paige, C. C., B. N. Parlett, and H. A. van der Vorst (1995). Approximate Solutions and Eigenvalue Bounds from Krylov Subspaces. *Numer. Linear Algebra Appl.*, 2(2):115–133.

Sleijpen, G. L. G. and J. van den Eshof (2003). On the Use of Harmonic Ritz Pairs in Approximating Internal Eigenpairs. *Linear Algebra Appl.*, 358(1–3):115–137.

Sleijpen, G. L. G. and H. A. van der Vorst (1996). A Jacobi–Davidson Iteration Method for Linear Eigenvalue Problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425.

Sleijpen, G. L. G., H. A. van der Vorst, and E. Meijerink (1998). Efficient Expansion of Subspaces in the Jacobi–Davidson Method for Standard and Generalized Eigenproblems. *Electron. Trans. Numer. Anal.*, 7:75–89.

Sorensen, D. C. (1992). Implicit Application of Polynomial Filters in a k-Step Arnoldi Method. *SIAM J. Matrix Anal. Appl.*, 13:357–385.

Stewart, G. W. (2001a). A Krylov–Schur Algorithm for Large Eigenproblems. *SIAM J. Matrix Anal. Appl.*, 23(3):601–614.

Stewart, G. W. (2001b). *Matrix Algorithms. Volume II: Eigensystems*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

Stewart, G. W. (2002). Addendum to "A Krylov–Schur Algorithm for Large Eigenproblems". *SIAM J. Matrix Anal. Appl.*, 24(2):599–601.

van der Vorst, H. A. (2002). Computational Methods for Large Eigenvalue Problems. In *Handbook of Numerical Analysis* (edited by P. G. Ciarlet and J. L. Lions), volume VIII, pp. 3–179. Elsevier, Amsterdam.

Wu, K. and H. Simon (2000). Thick-Restart Lanczos Method for Large Symmetric Eigenvalue Problems. *SIAM J. Matrix Anal. Appl.*, 22(2):602–616.