Electronic Transactions on Numerical Analysis.
Volume 20, pp. 235-252, 2005.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

GENERALIZATIONS OF HARMONIC AND REFINED RAYLEIGH–RITZ*

MICHIEL E. HOCHSTENBACH†

Abstract. We investigate several generalizations of the harmonic and refined Rayleigh–Ritz method. These may be practical when one is interested in eigenvalues close to one of two targets (for instance, when the eigenproblem has Hamiltonian structure such that eigenvalues come in pairs or quadruples), or in rightmost eigenvalues close to (for instance) the imaginary axis. Our goal is to develop new methods to extract promising approximate eigenpairs from a search space, for instance one generated by the Arnoldi or Jacobi–Davidson method. We give theoretical as well as numerical results of the methods, and recommendations for their use.

AMS subject classifications. 65F15, 65F50

Key words. Rational harmonic Rayleigh–Ritz, rightmost eigenvalue, structured eigenproblem, Hamiltonian matrix, Rayleigh–Ritz, harmonic Rayleigh–Ritz, refined Rayleigh–Ritz, subspace method, subspace extraction, Jacobi–Davidson

1. Introduction.

Let A be a (real or complex) large sparse $n \times n$ matrix. Suppose that p is a polynomial, and that we are looking for one or more normalized eigenvectors x such that the Euclidean norm $\|p(A)x\|$ is small. Since for an eigenvector we have

$$\|p(A)x\| = \|p(\lambda)x\| = |p(\lambda)|, \tag{1.1}$$

we see that these eigenvectors correspond to eigenvalues λ that have a small $|p(\lambda)|$. An example is the situation where we are interested in (interior) eigenpairs (λ, x) of A of which the eigenvalue λ close to a target τ. Then a natural choice is

$$p(z) = z - \tau; \tag{1.2}$$

indeed, if x is an eigenvector with a small $\|(A - \tau I)x\|$, where I is the identity matrix, its corresponding eigenvalue must be close to τ.

Subspace methods are often used for the computation of eigenpairs of large sparse matrices. An important part of these methods is the subspace extraction, where the task is to identify promising approximate eigenpairs $(\theta, u) \approx (\lambda, x)$, of which the approximate eigenvector u is in the search space. A good extraction process is important, especially at the time of a restart, when we reduce the dimension of the search space to save both computer time and memory. With a poor extraction at that time, we may discard the most relevant part of the search space, leading to very slow convergence or no convergence at all.

It is well known (see, for instance, [17, Section 4.4.1]) that the standard Rayleigh–Ritz extraction often yields good extraction results for exterior eigenpairs, but disappointing results for interior eigenpairs; simply selecting the Ritz pair (θ, u) where θ is the Ritz value closest to τ generally yields a poor approximate eigenvector u (see, for instance [17, p. 282]). Generally, the harmonic [12] or refined [4] Rayleigh–Ritz approach are preferred for interior eigenpairs.

We briefly review the harmonic Rayleigh–Ritz approach for the standard eigenvalue problem $Ax = \lambda x$, see also [17, Section 4.4.4], [3]. The idea behind this method is to apply a transformation to A such that the interior eigenvalues close to the target τ are mapped to

*Received September 15, 2004. Accepted for publication October 23, 2005. Recommended by R. Lehoucq.
†Department of Mathematics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106-7058, USA (hochsten@case.edu, www.case.edu/artsci/math/hochstenbach). The research of this author was supported in part by NSF grant DMS-0405387.

the exterior of the spectrum. For exterior eigenvalues, a Galerkin condition generally works favorably. Therefore, we start with the rational transformation

$$(A-\tau I)^{-1}x = (\lambda-\tau)^{-1}x.$$

If we have a k-dimensional search space $\mathcal{U}$ (where typically $k \ll n$), then we look for an approximate eigenpair $(\theta, u) \approx (\lambda, x)$ with $u \in \mathcal{U}$. As the exact equality $(A-\tau I)^{-1}u = (\theta-\tau)^{-1}u$ will not be reachable if $\mathcal{U}$ does not contain an eigenvector, we instead impose a Galerkin condition

$$(A-\tau I)^{-1}u - (\theta-\tau)^{-1}u \perp \mathcal{V} \tag{1.3}$$

for a certain *test space* $\mathcal{V}$. To avoid working with the inverse of a large sparse matrix, we want to choose the test space $\mathcal{V}$ in a suitable way. Taking $\mathcal{V} = (A-\tau I)^*\mathcal{U}$ leads to the standard Rayleigh–Ritz extraction $(A-\theta I)\,u \perp \mathcal{U}$. Therefore, we choose $\mathcal{V} = (A-\tau I)^*(A-\tau I)\,\mathcal{U}$. This is equivalent to requiring

$$(A-\tau I)^{-1}u - (\theta-\tau)^{-1}u \perp_{(A-\tau I)^*(A-\tau I)} \mathcal{U},$$

which means orthogonality with respect to the $(A-\tau I)^*(A-\tau I)$ inner product:

$$[x,\, y]_{(A-\tau I)^*(A-\tau I)} := y^*(A-\tau I)^*(A-\tau I)x.$$

(This is an inner product as long as τ is not an eigenvalue of A.) With this choice for $\mathcal{V}$, the characterizing Petrov–Galerkin constraint on the residual $(A-\theta I)\,u$ for harmonic Rayleigh–Ritz is

$$(A-\theta I)\,u \perp (A-\tau I)\,\mathcal{U}, \tag{1.4}$$

where θ is the *harmonic Ritz value* and u the *harmonic Ritz vector*. Let U be an $n \times k$-matrix with columns that form an orthonormal basis for $\mathcal{U}$. Then we can write $u = Uc$, for a low-dimensional k-vector c with unit norm. Using this, we get the projected generalized eigenvalue problem

$$U^*(A-\tau I)^*(A-\tau I)\,Uc = (\theta-\tau)\,U^*(A-\tau I)^*Uc. \tag{1.5}$$

The harmonic Ritz value θ satisfies

$$\theta = \tau + \frac{\|(A-\tau I)u\|^2}{u^*(A-\tau I)^*u} = \tau + \frac{\|(A-\tau I)u\|}{\cos((A-\tau I)\,u, u)}. \tag{1.6}$$

Let ξ assume the role of $\theta-\tau$. Any pair (ξ, c) solving the generalized eigenvalue problem

$$U^*(A-\tau I)^*(A-\tau I)\,Uc = \xi\, U^*(A-\tau I)^*Uc \tag{1.7}$$

satisfies (left-multiply (1.7) by c^* and use Cauchy–Schwarz, see [17, p. 293])

$$\|(A-\tau I)u\| \le |\xi|. \tag{1.8}$$

Hence, if $|\xi|$ is small, or equivalently if θ is close to τ, the residual norm of the approximate eigenpair (τ, u) (or, to lesser extent, of (θ, u)) is small. This not only forms a justification of the harmonic Rayleigh–Ritz method, but also suggests which approximate eigenvector should be taken from the search space: the vector Uc, where c is the eigenvector of the generalized eigenproblem (1.7) corresponding to the eigenvalue ξ with minimal $|\xi|$. In the

ETNA
Kent State University
etna@mcs.kent.edu

standard Rayleigh–Ritz method for interior eigenvalues, it is often far from clear which Ritz vector is the best approximation, see for example [17, p. 282]. For interesting other results see for instance [6, 1].

There are at least two reasons to investigate generalizations of the harmonic Rayleigh–Ritz approach. First, one might be interested in eigenvalues that are close to a set of complex numbers that contains more than one point τ. An important example of this situation is when one looks for the rightmost eigenvalue (that is, with maximal real part) of a matrix of which we know that the eigenvalues (almost) satisfy $\mathrm{Re}(\lambda) \le 0$. In this case, we may take the imaginary axis as the set of interest.

Second, if A has a certain structure then (1.5) loses this structure in general; we might say that *harmonic Rayleigh–Ritz does not respect the structure of the eigenproblem*. For instance, suppose A is (complex) Hamiltonian, then its eigenpairs come in pairs λ and $-\overline{\lambda}$. However, the generalized eigenvalue problem (1.5) generally has no Hamiltonian properties. As a consequence, the harmonic Ritz pairs, solutions to (1.5), do not come in θ and $-\overline{\theta}$ pairs, and the combination τ and $-\overline{\tau}$ is not dealt with in an equal manner. The wish to have a responsible extraction process for interior eigenvalues without sacrificing the structure of the problem too much leads us to consider *structure-preserving variants* of harmonic and refined Rayleigh–Ritz.

We will now study alternative extraction processes, in particular for the two situations mentioned above. In Section 2 we will study various generalizations of harmonic Rayleigh–Ritz and in Section 3 we will propose a generalization of refined Rayleigh–Ritz. Section 4 treats the situation where one is interested in the eigenvalues with maximum real part (the "rightmost eigenvalues"). Sections 5 applies the new approaches to the case of a matrix where the eigenvalues come in pairs with unrelated eigenvectors; we will take a Hamiltonian matrix as a role model. After dealing with various practical issues in Section 6, we finish with numerical experiments in Section 7 and a conclusion in Section 8.

2. Generalizations of harmonic Rayleigh–Ritz. We now generalize the harmonic Rayleigh–Ritz approach, which uses polynomial (1.2), by considering methods using rational polynomials of the form

$$r(z) = p(z)/q(z).$$

As we will see below, this corresponds to the situation in which we are interested in eigenvectors x for which $\|p(A)x\|$ is small compared to $\|q(A)x\|$. Since

$$\frac{\|p(A)x\|}{\|q(A)x\|} = \left|\frac{p(\lambda)}{q(\lambda)}\right|, \tag{2.1}$$

these correspond to eigenvalues λ for which is $|p(\lambda)|$ is relatively small compared to $|q(\lambda)|$. We get back the standard harmonic Rayleigh–Ritz approach of the previous section by taking $p(z) = z - \tau$ and $q \equiv 1$.

As in Section 1, we would like to map the eigenvalues λ for which $r(\lambda)$ is small to the exterior of the spectrum. Assume that $r(A)$ exists (that is, $q(A)$ is invertible) and that $r(A)$ (or, equivalently, $p(A)$) is invertible (otherwise, one of the roots of p or q is an eigenvalue). Since $r(A)x = r(\lambda)x$ and $r(A)^{-1}x = r(\lambda)^{-1}x$, a generalization of (1.3) is formed by the Galerkin condition

$$p(A)^{-1}q(A)\,u - r(\theta)^{-1}u \perp \mathcal{V}.$$

If we take $\mathcal{V} = p(A)^*\mathcal{U}$, we get a generalization of standard Rayleigh–Ritz for the problem $p(A)\,x = r(\lambda)\,q(A)\,x$:

$$p(A)\,u - r(\theta)\,q(A)\,u \perp \mathcal{U}.$$

However, similar to the standard Rayleigh–Ritz approach, if we select any pair (η, Uc) solving

$$U^*p(A)\,Uc = \eta\, U^*q(A)\,Uc,$$

for instance the pair with minimal $|\eta|$, we have no guarantee that $\|p(A)u\|$ will be small compared to $\|q(A)u\|$, as desired.

Therefore, we take $\mathcal{V} = p(A)^*p(A)\,\mathcal{U}$. Remembering that $u = Uc$, the following are equivalent:

(i) $p(A)^{-1}q(A)\,u - r(\theta)^{-1}u \quad \perp p(A)^*p(A)\,\mathcal{U}$,

(ii) $p(A)^{-1}q(A)\,u - r(\theta)^{-1}u \quad \perp_{p(A)^*p(A)}\,\mathcal{U}$,

(iii) $p(A)\,u - r(\theta)\,q(A)\,u \quad \perp p(A)\,\mathcal{U}$,

(iv) $U^*p(A)^*p(A)\,Uc \quad = r(\theta)\,U^*p(A)^*q(A)\,Uc$.

We call this method the *rational harmonic Ritz approach*. We call a pair (θ, u) satisfying any of the above equivalent requirements a *rational harmonic Ritz pair*, consisting of a *rational harmonic Ritz value* θ and a *rational harmonic Ritz vector* u.

We can interpret characterization (iii) as follows. For a generalized eigenvalue problem $Fx = \lambda Gx$, the harmonic Ritz approach with target $\tau = 0$ is (e.g., [17, p. 296])

$$U^*F^*FUc = \widetilde{\theta}\, U^*F^*GUc,$$

where $(\widetilde{\theta}, Uc)$ is a harmonic Ritz pair. Apparently, the rational harmonic extraction can be seen as a (standard) harmonic extraction applied to the spectrally transformed problem $p(A)\,x = r(\lambda)\,q(A)\,x$.

A first justification of this approach lies in the fact that any pair (ξ, c) solving the generalized eigenvalue problem

$$U^*p(A)^*p(A)\,Uc = \xi\, U^*p(A)^*q(A)\,Uc \tag{2.2}$$

satisfies (left-multiply by c^* and use Cauchy–Schwarz)

$$\|p(A)u\| \le |\xi|\; \|q(A)u\| \le |\xi|\; \|q(A)U\|. \tag{2.3}$$

Hence, if there is an eigenpair (ξ, c) of (2.2) with a small $|\xi|$, we know that $\|p(A)u\|$ is relatively small compared to $\|q(A)u\|$. Note that ξ satisfies (cf. (1.6))

$$\xi = \frac{\|p(A)u\|^2}{u^*p(A)^*q(A)\,u} = \frac{\|p(A)u\|}{\|q(A)u\|} \cdot \cos^{-1}(p(A)\,u,\, q(A)\,u). \tag{2.4}$$

We call ξ the *(rational harmonic) factor* of the rational harmonic Ritz vector. There are k factors $\xi_1, \dots, \xi_k$, which are eigenvalues of (2.2), belonging to k rational harmonic Ritz vectors; we are particularly interested in the smallest factor(s) in absolute value sense with corresponding rational harmonic vector(s).

Given a rational harmonic Ritz vector u, there does not need to be a unique *rational harmonic Ritz value* θ. However, only the factor ξ and not the value θ will be of importance in the rest of this paper.

A second justification of this method is given by the following proposition, which states that an eigenpair is also a rational harmonic Ritz pair. The proof is by direct verification, for instance using item (iii).

PROPOSITION 2.1. *If (λ, x) is an eigenpair then for all polynomials p and q it is also a rational harmonic Ritz pair with respect to the search space* $\mathrm{span}(x)$.

ETNA
Kent State University
etna@mcs.kent.edu

We note that if x is part of a larger subspace $\mathcal{U}$ the rational harmonic approach may still have difficulties selecting the right vector if there are (nearly) equal rational harmonic Ritz values; see [5] for the standard harmonic Ritz extraction, and Theorem 2.3 below. This, however, rarely forms a problem in practical processes, since we can just continue the subspace method by expanding the search space and performing a new extraction process.

Suppose we have a rational harmonic Ritz vector with a small ξ, what does this say about the existence of an eigenpair of interest in the neighborhood? The following theorem sheds some more light on this situation. Part (a) is a generalization of Bauer–Fike (see, for instance, [13, Th. 3.6]). For (b), which is a generalization of [17, p. 288], we use the fact that if λ is a simple eigenvalue of A with corresponding eigenvector x, $\|x\| = 1$, and if $[x\ Y_1]$ is an orthonormal basis, then there exist y and X_1 such that $[y\ Y_1]^*[x\ X_1] = I$ and

$$\begin{bmatrix} y^* \\ Y_1^* \end{bmatrix} A\,[x\ X_1] = \begin{bmatrix} \lambda & 0 \\ 0 & L \end{bmatrix}, \tag{2.5}$$

see, e.g., [17, p. 244]; in this case A is said to have the spectral representation $A = \lambda x y^* + X_1 \Lambda Y_1^*$. For a matrix Z, let $\sigma_{\min}(Z)$, $\sigma_{\max}(Z)$, and $\kappa(Z) := \|Z\|\,\|Z^{-1}\|$ denote its smallest singular value, largest singular value, and condition number, respectively.

THEOREM 2.2. *Let u be a rational harmonic Ritz vector with factor ξ.*
(a) If $A = X\Lambda X^{-1}$ is diagonalizable, there exists an eigenvalue λ such that

$$|p(\lambda)| \le |\xi|\,\kappa(X)\,\|q(A)u\|,$$

(b) Let λ be a simple eigenvalue such that A has the spectral decomposition $A = \lambda x y^* + X_1 L Y_1^*$ as in (2.5). Then

$$\sin(u,x) \le \frac{\|p(A)u\|}{\sigma_{\min}(p(L))} \le |\xi|\,\frac{\|q(A)u\|}{\sigma_{\min}(p(L))},$$

where $\sigma_{\min}(p(L)) > 0$ if the eigenvalues of A other than λ are no zeros of p.

Proof. We have

$$\|p(A)u\| \ge \sigma_{\min}(p(A)) \ge \sigma_{\min}(X)\,\sigma_{\min}(p(\Lambda))\,\sigma_{\min}(X^{-1}) = \kappa(X)^{-1}\,\min_j |p(\lambda_j)|$$

from which, together with (2.3), we conclude (a). For (b), we note that from $Y_1^*A = LY_1^*$ it follows that $Y_1^* p(A) = p(L) Y_1^*$. So

$$\sin(u,x) = \|Y_1^* u\| = \|p(L)^{-1} Y_1^* p(A)\,u\| \le \|p(L)^{-1}\|\,\|p(A)u\|,$$

from which the result follows. Note that the eigenvalues of L are the eigenvalues λ_j of A with the exception of λ; the eigenvalues of $p(L)$ are $p(\lambda_j)$, $\lambda_j \ne \lambda$. ∎

Three very heuristic remarks about this theorem are in order. First, decomposing $u = \cos(u,x)\,x + \sin(u,x)\,w$, where $w \perp x$, we get

$$\|q(A)u\| \le \frac{1}{\sqrt{1-\sin^2(u,x)}}\,|q(\lambda)| + \sin(u,x)\,\|q(A)\|. \tag{2.6}$$

If $|\xi|$ is small, then from (b) $\sin(u,x) = \mathcal{O}(|\xi|)$, from (2.6) $\|q(A)u\| \lesssim |q(\lambda)| + \mathcal{O}(|\xi|)$, and from (a) $|p(\lambda)| \lesssim |\xi|\,\kappa(X)\,|q(\lambda)|$, giving a relative estimate

$$|p(\lambda)|\,/\,|q(\lambda)| \lesssim |\xi|\,\kappa(X). \tag{2.7}$$

ETNA
Kent State University
etna@mcs.kent.edu

Second, we note that both expressions (a) and (2.7) may be a crude overestimation in practice. We may get a more realistic, but difficult to assess, first order estimate as follows.

If $|\xi|$ is small then $\sigma_{\min}(p(A)) \le \|p(A)u\| \le |\xi|\,\|q(A)u\| \lesssim |\xi|\,|q(\lambda)|$ is small. First we have to determine the smallest $\varepsilon_0 \ge 0$ such that there exists an E with $\|E\| = 1$ and $p(A + \varepsilon_0 E)$ singular. This looks like a hard problem for general p; it seems reasonable to assume that ε_0 is related to $|\xi|\,|q(\lambda)|$.

With this ε_0 and if $p(A + \varepsilon_0 E)$ is diagonalizable, there must be an eigenvalue $\widetilde{\lambda}$ of $A + \varepsilon_0 E$ such that $p(\widetilde{\lambda}) = 0$. Considering only the first order perturbation, there must be an eigenvalue λ of A with $|\lambda - \widetilde{\lambda}| \lesssim \varepsilon_0\kappa(\lambda)$, where $\kappa(\lambda) = |y^*x|^{-1}$ is the usual condition number of λ. For $p(\lambda)$ we then have for small enough $\varepsilon_0\kappa(\lambda)$

$$|p(\lambda)| \approx |p(\widetilde{\lambda}) + (\lambda - \widetilde{\lambda})p'(\lambda)| \lesssim \varepsilon_0\,\kappa(\lambda)\,|p'(\lambda)|.$$

If ε_0 is related to $|\xi|\,|q(\lambda)|$, this gives a heuristic first order bound on $|p(\lambda)|\,/\,|q(\lambda)|$.

Third, we shall see cases and numerical examples in the rest of the paper in which $|\xi| = \mathcal{O}(1)$; then Theorem 2.2 has little or nothing to say. However, as we will see, the approaches can still be successful for extracting promising approximate eigendata from a search space.

The following result is a generalization of a result of Chen and Jia [2] for harmonic Ritz vectors, which is on its turn inspired by a well-known theorem for Ritz vectors due to Saad [13, Th. 4.6]. It gives an upper bound of the angle of the best rational harmonic Ritz vector in terms of the quality of the search space.

Let (ξ, u) satisfy (2.2) (so u is a rational harmonic Ritz vector with factor ξ), and let $[u\ V\ W]$ be a unitary matrix such that $\mathrm{span}([u\ V]) = \mathcal{U}$ and $\mathrm{span}(W) = \mathcal{U}^\perp$. (By "span", we here mean the span of the columns of a matrix.) We first write

$$C = [u\ V]^* p(A)^* p(A) [u\ V] \quad \text{and} \quad B = [u\ V]^* p(A)^* q(A) [u\ V],$$

and we will assume that B is invertible. Denote by e_1 the first column of the identity matrix. Since (cf. (2.2))

$$p(A)^* p(A)\, u - \xi\, p(A)^* q(A)\, u \perp \mathcal{U},$$

we have $B^{-1}Ce_1 = \xi e_1$, so $B^{-1}C$ is of the form

$$B^{-1}C = \begin{bmatrix} \xi & g^* \\ 0 & G \end{bmatrix}. \tag{2.8}$$

Since the eigenvalues of $B^{-1}C$ are the rational harmonic factors, the eigenvalues of G are the rational harmonic factors other than ξ.

THEOREM 2.3. *Let u be a rational harmonic Ritz vector with factor ξ and B be invertible. Then*

$$\sin(u, x) \le \sin(\mathcal{U}, x) \sqrt{1 + \frac{\gamma^2\,\|B^{-1}\|^2}{\delta^2}},$$

where

$$\gamma = \|P_{\mathcal{U}}\, p(A)^* (p(A) - r(\lambda) q(A))(I - P_{\mathcal{U}})\|,$$
$$\delta = \mathrm{sep}(r(\lambda),\, G) := \sigma_{\min}(G - r(\lambda) I) \le \min_{\xi_j \ne \xi} |\xi_j - r(\lambda)|,$$

where $P_{\mathcal{U}}$ is the orthogonal projection on $\mathcal{U}$, and G is defined as in (2.8).

ETNA
Kent State University
etna@mcs.kent.edu

Proof. Introduce a new variable $z = [z_1^T\ z_2^T\ z_3^T]^T = [u\ V\ W]^* x$. From $p(A)x = r(\lambda)q(A)x$ we get $p(A)^*p(A)x = r(\lambda)p(A)^*q(A)x$ and

$$\begin{bmatrix} u^* \\ V^* \\ W^* \end{bmatrix} p(A)^*p(A)[u\ V\ W] \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix} = r(\lambda) \begin{bmatrix} u^* \\ V^* \\ W^* \end{bmatrix} p(A)^*q(A)[u\ V\ W] \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix},$$

which we can write as

$$\begin{bmatrix} C & C_1 \\ C_2 & C_3 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix} = r(\lambda) \begin{bmatrix} B & B_1 \\ B_2 & B_3 \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix},$$

where $B_1 = [u\ V]^*p(A)^*q(A)W$ and $C_1 = [u\ V]^*p(A)^*p(A)W$. Therefore we get

$$C \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + C_1 z_3 = r(\lambda)B \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} + r(\lambda)B_1 z_3$$

or

$$(B^{-1}C - r(\lambda)I) \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = B^{-1}(r(\lambda)B_1 - C_1)z_3. \tag{2.9}$$

For the right-hand side of (2.9) we have

$$\begin{aligned} \|B^{-1}(r(\lambda)B_1 - C_1)z_3\| &\le \|B^{-1}\|\ \|r(\lambda)B_1 - C_1\|\ \|z_3\| \\ &= \|B^{-1}\|\ \|[u\ V]^*p(A)^*(r(\lambda)q(A) - p(A))W\|\ \|z_3\| \end{aligned}$$

while the left-hand side of (2.9) can be bounded from below by

$$\left\| \begin{bmatrix} \xi - r(\lambda) & g^* \\ 0 & G - r(\lambda)I \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} \right\| \ge \|(G - r(\lambda)I)z_2\| \ge \mathrm{sep}(r(\lambda), G)\ \|z_2\|.$$

Combining these two bounds, we have

$$\|z_2\| \le \frac{\gamma\,\|B^{-1}\|}{\delta}\ \|z_3\|$$

Since $\sin^2(\mathcal{U}, x) = \|W^*x\|^2 = \|z_3\|^2$ and $\sin^2(x, u) = \|[V\ W]^*x\|^2 = \|z_2\|^2 + \|z_3\|^2$, the result now follows. □

We see that, besides the quantities γ and δ similar to those that are also present in Saad's original theorem, we have an additional factor $\|B^{-1}\|$ in this case. The quantity δ can be interpreted as follows: since $r(\lambda)$ is the factor corresponding to the eigenvector x (see (2.4)), δ is bounded by above by the minimum distance between the factor of x and the factors of the rational harmonic Ritz vectors not equal to u. In particular, $\delta > 0$ if ξ is a simple factor.

Apart from the rational harmonic Ritz value θ not being unique in general, it may form a poor approximate eigenvalue (see [16] for comments on the standard harmonic extraction). When we have a rational harmonic vector u, we can get a sensible approximate eigenvalue by taking the Rayleigh quotient $\rho(u) = u^*Au$; this ρ minimizes the residual norm $\|Au - \theta u\|$ over all values of θ. Taking an extra Rayleigh quotient is also advised in the context of (standard) harmonic extraction, see [16].

3. A generalization of refined Rayleigh–Ritz. Refined Rayleigh–Ritz is an alternative to harmonic Rayleigh–Ritz to ensure a small residual norm. For a target τ and a search space $\mathcal{U}$, one minimizes

$$\|A\widehat{u} - \tau\widehat{u}\| \qquad \text{over} \qquad \widehat{u} \in \mathcal{U},\ \|\widehat{u}\| = 1.$$

Generalizing this idea, *polynomial refined Rayleigh–Ritz* minimizes

$$\|p(A)\widehat{u}\| \qquad \text{over} \qquad \widehat{u} \in \mathcal{U},\ \|\widehat{u}\| = 1; \tag{3.1}$$

we will denote the minimum by $\widehat{\xi}$. In the context of the previous section, we here take $q \equiv 1$, which means we are interested in eigenvectors x for which $|p(\lambda)| = \|p(A)x\|$ is small. We call the minimizing vector $\widehat{u} \in \mathcal{U}$ a *polynomial refined Ritz vector*. Note that it is given by $\widehat{u} = U\widehat{c}$, where $\widehat{c}$ is the smallest (right) singular vector of the $n \times k$ matrix $p(A)\,U$. The next result follows immediately using the same arguments as in Theorem 2.2.

COROLLARY 3.1. *Let $\widehat{\xi}$ and $\widehat{u}$ as in (3.1). With the same assumptions and notations as in Theorem 2.2, there is a λ such that*

$$|p(\lambda)| \le \widehat{\xi}\,\kappa(X).$$

Moreover,

$$\sin(\widehat{u}, x) \le \frac{\widehat{\xi}}{\sigma_{\min}(p(L))}.$$

The next theorem, a generalization of a part of [17, p. 290], gives a bound of the residual norm in terms of the quality of the search space; compare this result with (1.1).

THEOREM 3.2. *Let $\widehat{\xi}$ and $\widehat{u}$ as in (3.1) and let x be an eigenvector. For the residual of the polynomial refined Ritz vector $\widehat{u}$ we have*

$$\widehat{\xi} = \|p(A)\widehat{u}\| \le \frac{|p(\lambda)| + \sin(\mathcal{U}, x)\,\|p(A)\|}{\sqrt{1 - \sin^2(\mathcal{U}, x)}}.$$

Proof. Decompose $x = \gamma_U x_U + \sigma_U e_U$, where $x_U := UU^*x/\|UU^*x\|$ is the orthogonal projection of x onto $\mathcal{U}$, $\|x_U\| = \|e_U\| = 1$, $\gamma_U = \cos(\mathcal{U}, x)$, and $\sigma_U = \sin(\mathcal{U}, x)$. Since $p(A)x_U = (p(A)x - \sigma_U p(A)e_U)/\gamma_U$, we have by the definition of a polynomial refined Ritz vector

$$\|p(A)\widehat{u}\| \le \|p(A)x_U\| \le (|p(\lambda)| + \sigma_U\,\|p(A)\|)/\gamma_U.$$

□

Having the polynomial refined Ritz vector $\widehat{u}$, it is advisable—as for the rational harmonic approach—to take an extra Rayleigh quotient of $\widehat{u}$ to determine a sensible approximate eigenvalue.

4. Rightmost eigenvalues. We are going to apply the techniques of the previous two sections to two different problems in the next two sections. This section focusses on the computation of the rightmost eigenvalue. Suppose we know that all eigenvalues of A are (almost) located in the left-half plane, that is, $\mathrm{Re}(\lambda) \lesssim 0$ for all eigenvalues. For many

ETNA
Kent State University
etna@mcs.kent.edu

applications, for instance stability considerations, we are interested in the rightmost eigenvalues, i.e., the ones with the largest real part. The problem of finding rightmost eigenvalues has already been studied by Meerbergen and his collaborators, for instance using generalized Cayley transforms [10, 7, 8] or the matrix exponential [9] in the subspace expansion phase. However, here we use generalized Cayley transforms and several novel generalizations *in the subspace extraction phase* of the process. This has the advantage that the subspace expansion may take place with well-established methods such as Arnoldi or Jacobi–Davidson [15]; this expansion may also be (much) cheaper and more practical than expansion using (shift-and-invert or generalized Cayley) transformations of the matrix.

We can use a generalization of harmonic Rayleigh–Ritz introduced in Section 2 in the following way. Let τ_1 and τ_2 be complex numbers with $\tau_2 = -\overline{\tau_1}$ and $\mathrm{Re}(\tau_2) < \mathrm{Re}(\tau_1)$. The generalized Cayley transform

$$r(z) = \frac{z-\tau_1}{z-\tau_2}$$

has the following mapping properties:

$$\begin{array}{lcl} \{z \mid \mathrm{Re}(z) = 0\} \text{ (imaginary axis)} & \stackrel{r}{\mapsto} & \text{unit circle } \{z \mid |z| = 1\}, \\ \{z \mid \mathrm{Re}(z) < 0\} \text{ (left-half plane)} & \stackrel{r}{\mapsto} & \text{exterior of the unit circle,} \\ \{z \mid \mathrm{Re}(z) > 0\} \text{ (right-half plane)} & \stackrel{r}{\mapsto} & \text{interior of the unit circle.} \end{array}$$

As an example, we take $\tau_1 = 1$, $\tau_2 = -1$ and plot a contour plot of the function $r = |z-1|/|z+1|$ in Figure 4.1(a). We can see that the line $\mathrm{Re}(z) = 0$ is one of the contours, namely the 1-level curve. This means that when we look for the eigenvalue that is closest to the imaginary axis, or the only eigenvalue that is in the right-half plane, we can try to find the eigenvalue with the smallest r-value. This means we can use the rational harmonic method of Section 2 with $p(z) = z-1$ and $q(z) = z+1$.

Returning to more general targets τ_1 and τ_2, the corresponding Galerkin condition

$$U^*(A-\tau_1 I)^{-1}(A-\tau_2 I)u - \xi^{-1}u \perp \mathcal{V}$$

leads, with the choice $\mathcal{V} = (A-\tau_1 I)^*(A-\tau_1 I)\mathcal{U}$, to

$$U^*(A-\tau_1 I)^*(A-\tau_2 I)Uc = \xi^{-1}U^*(A-\tau_1 I)^*(A-\tau_1 I)Uc.$$

For a rational harmonic Ritz vector satisfying this equation we get

$$\|(A-\tau_1 I)\,u\| \le |\xi|\,\|(A-\tau_2 I)\,u\|.$$

However, in view of Figure 4.1(a), if there are for instance eigenvalues $-0.5+2i$ and -0.3, the extraction process may select the first since it may have a lower r-value than the second eigenvalue. Therefore, it may select an incorrect eigenvalue as being the rightmost one. We are interested in level curves that run more vertically. There are at least two approaches to try to improve this. The first is to take the targets somewhat remote from the likely interesting part of the spectrum, see the numerical experiments in Section 7. The second is to add two extra targets.

Let τ_1, τ_2, τ_3, and τ_4 be complex numbers with $\tau_2 = -\overline{\tau_1}$ and $\tau_4 := -\overline{\tau_3}$. Then the map

$$r(z) = \frac{(z-\tau_1)(z-\tau_3)}{(z-\tau_2)(z-\tau_4)},$$

the product of two generalized Cayley transforms, also has the imaginary axis as the 1-level curve, but favorably, the neighboring curves are (at least locally) more vertical then in the

ETNA
Kent State University
etna@mcs.kent.edu

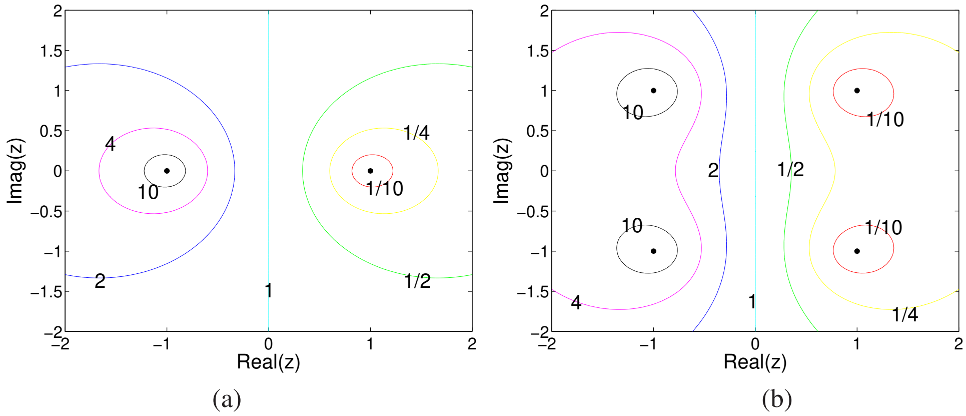

FIG. 4.1. *Contour plot of* $\left|\frac{z-1}{z+1}\right|$ *(a) and* $\left|\frac{(z-1-i)(z-1+i)}{(z+1-i)(z+1+i)}\right|$ *(b).*

case with just two targets. In Figure 4.1(b), we take $\tau_1 = 1-i$, $\tau_2 = -1-i$, $\tau_3 = 1+i$, and $\tau_4 = -1+i$ as an example. Therefore, we can also seek for rightmost eigenvalues by a rational harmonic approach with polynomials of the form $p(z) = (z-\tau_1)(z-\tau_3)$ and $q(z) = (z-\tau_2)(z-\tau_4)$.

Of course, we can scale the axes of Figure 4.1(b) by changing the real and/or imaginary components of the τ_j. Also, instead of the line $\mathrm{Re}(z) = 0$ we can easily take another vertical line, for instance in the case of two targets by selecting different τ_1 and τ_2 with $\mathrm{Im}(\tau_1) = \mathrm{Im}(\tau_2)$; in this case, the 1-level curve is the line $\mathrm{Re}(z) = \mathrm{Re}((\tau_1+\tau_2)/2)$.

In Figure 4.2(a), we show another alternative to find the most rightmost eigenvalue with only three target points: $\tau_1 = 1-i$, $\tau_2 = 1+i$, and $\tau_3 = -1$. So here p is of the form $(z-\tau_1)(z-\tau_2)$ and q is of the form $z-\tau_3$. Here only the p-polynomial is of degree two instead of both p and q. However, the line $\mathrm{Re}(z) = 0$ is no contour line and looks far from horizontal for $|\,\mathrm{Im}(z)| > 1$. This implies that for this choice (degree(p) = 2, degree(q)=1), it is important to have a good suspicion where the spectrum of A is located.

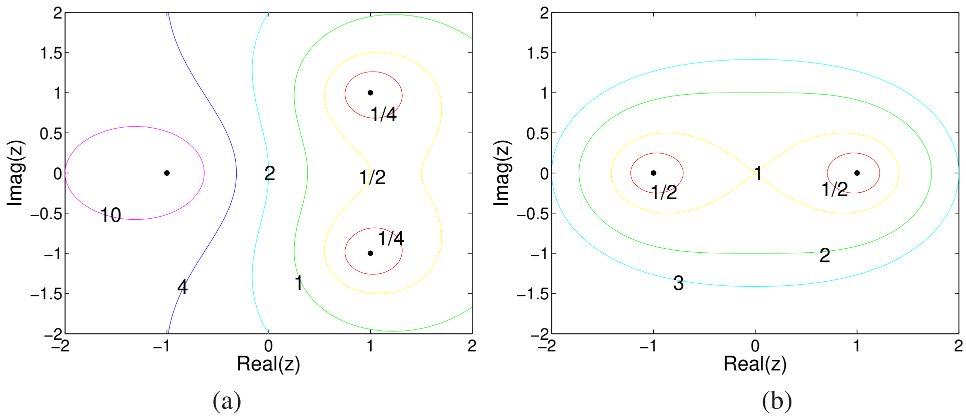

FIG. 4.2. *Contour plot of* $\left|\frac{(z-1-i)(z-1+i)}{z+1}\right|$ *(a) and* $|(z+1)(z-1)|$ *(b).*

We will experiment with these approaches in Section 7.

ETNA
Kent State University
etna@mcs.kent.edu

5. Eigenvalues close to one of two targets for structured eigenproblems. Now we consider the problem where we are interested in eigenvalues that are close to one of two targets τ_1, τ_2. This situation may occur when one studies an eigenvalue problem with structure. As a role model, we take the example of a (complex) Hamiltonian matrix.

Suppose that A with size $2m \times 2m$ is Hamiltonian, that is, $(JA)^* = JA$ for $J = \left[\begin{smallmatrix} 0 & I_m \\ -I_m & 0 \end{smallmatrix}\right]$. Then if λ is an eigenvalue, $-\overline{\lambda}$ is also an eigenvalue, while the corresponding (right) eigenvectors are not related in general. (However, the left eigenvector corresponding to $-\overline{\lambda}$ is closely related to the right eigenvector corresponding to λ.)

If we would like to preserve some structure of the Hamiltonian eigenproblem in the subspace extraction, we may wish to look for eigenvalues close to τ and $-\overline{\tau}$ *simultaneously*. Other examples of structured eigenvalue problem where eigenvalues occur in pairs are

- complex skew-Hamiltonian matrices: $(JA)^* = -(JA)$. If λ is an eigenvalue, $\overline{\lambda}$ is also an eigenvalue. We may take $\tau_2 = \overline{\tau_1}$.
- complex conjugate-symplectic matrices: $A^*JA = J$. If λ is an eigenvalue, $1/\overline{\lambda}$ is also an eigenvalue. We may take $\tau_2 = 1/\overline{\tau_1}$.

Assume we have a structured eigenproblem where we are interested in the targets τ_1 and τ_2. The simplest function r that is small for both targets is given by $r(z) = (z - \tau_1)(z - \tau_2)$; a contour plot of the example $r(z) = (z - 1)(z + 1)$ is given in Figure 4.2(b). In this case, $p(z) = r(z)$ and $q \equiv 1$. The Galerkin corresponding condition

$$(A - \tau_1 I)^{-1}(A - \tau_2 I)^{-1}u - \xi^{-1}u \perp \mathcal{V}$$

leads, with the choice $\mathcal{V} = (A - \tau_1 I)^*(A - \tau_2 I)^*(A - \tau_2 I)(A - \tau_1 I)\mathcal{U}$, to

$$U^*(A - \tau_1 I)^*(A - \tau_2 I)^*Uc = \xi^{-1}U^*(A - \tau_1 I)^*(A - \tau_2 I)^*(A - \tau_2 I)(A - \tau_1 I)Uc.$$

This may seem computationally unattractive, but we will deal with this in Section 6.2. For a rational harmonic Ritz vector satisfying this equation we get

$$\|(A - \tau_1 I)(A - \tau_2 I)\, u\| \le |\xi|.$$

A disadvantage of this could be that the r-function is not very distinctive: every z in (say) the unit circle has a modest r-value (see Figure 4.2(b)).

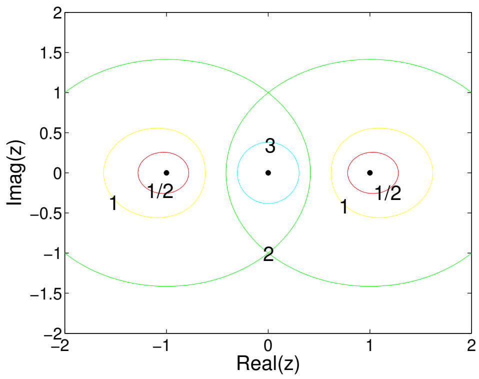

(a)

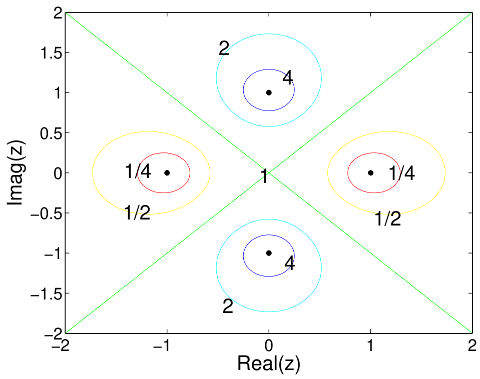

(b)

FIG. 5.1. *Contour plot of* $\left|\dfrac{(z-1)(z+1)}{z}\right|$ *(a) and* $\left|\dfrac{(z-1)(z+1)}{(z-i)(z+i)}\right|$ *(b).*

In Figure 5.1, we give two alternative q-functions (instead of $q \equiv 1$) such that the regions where $|r(z)|$ is small are smaller and more separated. In Figure 5.1(a), we squeeze a third target between the two targets (example: $q(z) = z$), while in Figure 5.1(b) we add two targets (example: $q(z) = (z-i)(z+i)$). As can be seen in the latter case, this approach could be suitable to find eigenvalues in a cone.

For some problems, eigenpairs even come in quadruples, for instance in the following cases:

- real Hamiltonian matrices: the eigenvalues come in quadruples λ, $\overline{\lambda}$, $-\lambda$, $-\overline{\lambda}$.
- real symplectic matrices: the eigenvalues come in quadruples λ, $\overline{\lambda}$, $1/\lambda$, $1/\overline{\lambda}$.

For the first case, Mehrmann and Watkins [11] suggested to work with the rational transformation

$$(A-\tau I)^{-1}\,(A+\tau I)^{-1}\,(A-\overline{\tau} I)^{-1}\,(A+\overline{\tau} I)^{-1} \tag{5.1}$$

for finding the eigenvalues nearest the quadruple $(\tau, \overline{\tau}, -\tau, -\overline{\tau})$, or, in case of either a real or purely imaginary target, with

$$(A-\tau I)^{-1}\,(A+\tau I)^{-1}. \tag{5.2}$$

For large sparse matrices A, these approaches may often not be feasible because of the prohibitively expensive matrix inverses. However, with a p-polynomial of the form $p(z) = (z-\tau)(z+\tau)$, our rational harmonic approach actually *employs a transformation of the form (5.2) in the subspace extraction in an implicit and practical way*.

The form (5.1) has additional disadvantages. In the first place the corresponding p-polynomial

$$p(z) = (z-\tau)(z+\tau)(z-\overline{\tau})(z+\overline{\tau})$$

is of degree four and hence $p(A)$ may have an enormous condition number, even for a modestly ill-conditioned matrix A. This makes the computation of the generalized harmonic pairs (2.2) unattractive and possibly unstable (see also Section 6.2). In addition, in a subspace method as Arnoldi or Jacobi–Davidson, this approach costs three matrix-vector products (MVs) extra per iteration step, see also Section 6.2.

Moreover, in the examples mentioned above, the eigenvectors are related: eigenpairs come in pairs (λ, x) and $(\overline{\lambda}, \overline{x})$. Thus, an approximate vector $u \approx x$ naturally leads to an approximate eigenvector $\overline{u} \approx \overline{x}$, such that there is hardly a need to simultaneously look for approximate eigenvectors corresponding to λ and $\overline{\lambda}$. Concluding, *focussing simultaneously on four targets may not only be unattractive, but also unnecessary*.

Therefore, we propose to use a transformation of the form

$$(A-\tau I)^{-1}(A+\overline{\tau} I)^{-1}$$

for both real and non-real targets τ. Strictly speaking, this is not a structure-preserving transformation: the operator is neither Hamiltonian nor skew-Hamiltonian. But this approach does take into account the targets τ and $-\overline{\tau}$ simultaneously, and via the conjugate of the search space, the targets $\overline{\tau}$ and $-\tau$.

As mentioned in [11], another natural rational transformation that is sometimes applied to Hamiltonian matrices is the product of two generalized Cayley transforms. Given a Hamiltonian A and a target τ, one then works with the transformation

$$(A-\tau I)^{-1}\,(A+\tau I)\,(A-\overline{\tau} I)^{-1}\,(A+\overline{\tau} I).$$

ETNA
Kent State University
etna@mcs.kent.edu

If τ is real, one may use the simpler generalized Cayley transform

$$(A - \tau I)^{-1}\,(A + \tau I).$$

In [11], this is done because the resulting operator is symplectic; in this paper, we proposed transformations of these two forms for finding rightmost eigenvalues in the previous section (cf. Figure 4.1).

6. Various issues.

6.1. Relation with standard harmonic Ritz in the linear case. In case both p and q are polynomials of degree one, the rational harmonic Ritz approach is related to the standard harmonic Ritz method as follows. Suppose $p(z) = z - \alpha$ and $q(z) = z - \beta$. (If exactly one of p and q has degree zero, one may check that we get the (standard) Ritz or (standard) harmonic extraction.) Then

$$(A - \alpha I)\,u - \xi\,(A - \beta I)\,u \perp (A - \alpha I)\,\mathcal{U}$$

is equivalent to

$$\Big(A - \frac{\beta\xi - \alpha}{\xi - 1}\,I\Big)u \perp (A - \alpha I)\,\mathcal{U}.$$

Comparing this with the standard harmonic Rayleigh–Ritz approach (1.4), we have the one-one correspondence

$$\theta = \frac{\beta\xi - \alpha}{\xi - 1}, \qquad \xi = \frac{\theta - \alpha}{\theta - \beta}.$$

This means that if p and q have degree one, then the rational harmonic Ritz vectors are equal to the harmonic Ritz vectors, with an ordering

$$\frac{|\theta_1 - \alpha|}{|\theta_1 - \beta|} \le \cdots \le \frac{|\theta_k - \alpha|}{|\theta_k - \beta|}$$

on the harmonic Ritz values. For instance, in the case that $\alpha = 1$, $\beta = -1$, the rational harmonic approach will extract the harmonic Ritz pair(s) (θ, u) with the smallest corresponding $\frac{|\theta-1|}{|\theta+1|}$. Since these roughly correspond to the harmonic Ritz values closest to the imaginary axis, one could also try to extract the rightmost harmonic Ritz value each iteration. However, the imaginary axis is no level curve of the function $r(z) = z - \tau$, the relevant function for the standard harmonic extraction with target τ, which suggests that this idea is not likely to work in practice. This suspicion was confirmed by numerical experiments (not reported here).

6.2. The practical computation of rational harmonic and polynomial refined Rayleigh–Ritz. Let us consider the practical computation of the rational harmonic Rayleigh–Ritz method. In view of (2.2), we might fear that we have to cope with $[\kappa(p(A))]^2$. Fortunately, this can be avoided with a technique similar to the standard harmonic Rayleigh–Ritz approach. With the QR-decomposition $p(A)\,U = QR$, (2.2) becomes

$$Rc = \xi\, Q^* q(A)\, Uc. \tag{6.1}$$

The advantages of this approach are twofold. In the first place, it makes the computations more efficient: both the QR-decomposition and $q(A)\,U$ can be computed incrementally (column per column) and hence a new step only requires a few inner products and the solution of

ETNA
Kent State University
etna@mcs.kent.edu

a small generalized eigenvalue problem. Moreover, we avoid squaring the condition number of $p(A)\,U$. As the last detail, we do not compute the pair(s) (ξ, c) of (6.1) with the smallest $|\xi|$, but the pair(s) (ξ, c) of

$$Q^* q(A)\, U c = \xi^{-1} R c$$

with the largest $|\xi|$, since this is generally more stable, cf. [17, p. 294].

For the polynomial refined Rayleigh–Ritz, we may use the same incremental QR decomposition. In step k, we then have to determine

$$\min_{\|c\|=1} \|p(A)Uc\| = \min_{\|c\|=1} \|QRc\| = \min_{\|c\|=1} \|Rc\|,$$

the minimal right singular vector of a $k \times k$ upper triangular matrix R.

Compared to Jacobi–Davidson with standard, harmonic, or refined Rayleigh–Ritz in the extraction phase, the new approaches take $\max(\deg(p), \deg(q)) - 1$ extra matrix-vector products per outer iteration step. Therefore, linear polynomials p and q do not involve extra MVs, and since often the subspace expansion (solving the correction equation) will take some MVs as well, the costs of approaches with quadratic p and/or q may be relatively modest.

6.3. An overview of the generalizations. In Table 6.1, we summarize the different generalizations of the harmonic and refined Rayleigh–Ritz method with their use. We name the harmonic methods after the degree of their p and q polynomial. The refined method with two targets is called "double refined".

TABLE 6.1
Overview of the different generalizations and their uses.

Method	Deg(p)	Deg(q)	Use	See Fig.
Standard harmonic	1	0	one target	—
(1,1)-harmonic	1	1	rightmost	4.1(a)
(2,1)-harmonic	2	1	rightmost	4.2(a)
(2,2)-harmonic	2	2	rightmost	4.1(b)
(2,0)-harmonic	2	0	two targets	4.2(b)
(2,1)-harmonic	2	1	two targets	5.1(a)
(2,2)-harmonic	2	2	two targets	5.1(b)
Standard refined	1	–	one target	–
Double refined	2	–	two targets	4.2(b)

7. Numerical experiments. EXPERIMENT 7.1. First, we test various approaches to find the rightmost eigenvalue. We take three challenging test matrices:

(1) $\texttt{diag}(-\texttt{rand}(1000,1) + i(\texttt{rand}(1000,1) - \frac{1}{2}))$, a diagonal matrix with eigenvalues with real part between -1 and 0, and imaginary part between -0.5 and 0.5;
(2) the same matrix, but then with the value 0.1 on the superdiagonal to make it nonnormal;
(3) $\texttt{diag}(-398, \ldots, -2, -1, 52i, -52i)$; this 400×400 matrix forms a similar example as is used in [10]. This spectrum is typical for a certain double diffusive convection problem.

To find the rightmost eigenvalue, we compare the Jacobi–Davidson (JD) method with the following subspace extraction methods:

- extracting the rightmost Ritz pair in every step (the "standard JD" approach);
- the best approximate pair according to the (1,1)-harmonic approach, so $p(z) = z - \tau$, $q(z) = z + \overline{\tau}$, with different choices for the target τ;

ETNA
Kent State University
etna@mcs.kent.edu

- the best approximate pair according to the (2,2)-harmonic approach, which means $p(z) = (z-\tau_1)(z-\tau_2)$, $q(z) = (z+\overline{\tau_1})(z+\overline{\tau_2})$, with different choices for the targets $\tau_{1,2}$.

The Jacobi–Davidson parameters: the correction equations are solved by 10 steps of unpreconditioned GMRES, the dimension of the search space is kept between 10 and 20 by restarts, and the tolerance of the method is 10^{-6}. In Figure 7.1(a) we display part of the eigenvalues of the first two test matrices, with arrows indicating the first, second, and third rightmost eigenvalue: $\lambda_1 \approx -0.00047 - 0.17i$, $\lambda_2 \approx -0.00055 + 0.41i$, $\lambda_3 \approx -0.00077 - 0.20i$ (while $\lambda_4 \approx -0.00570 - 0.40i$).

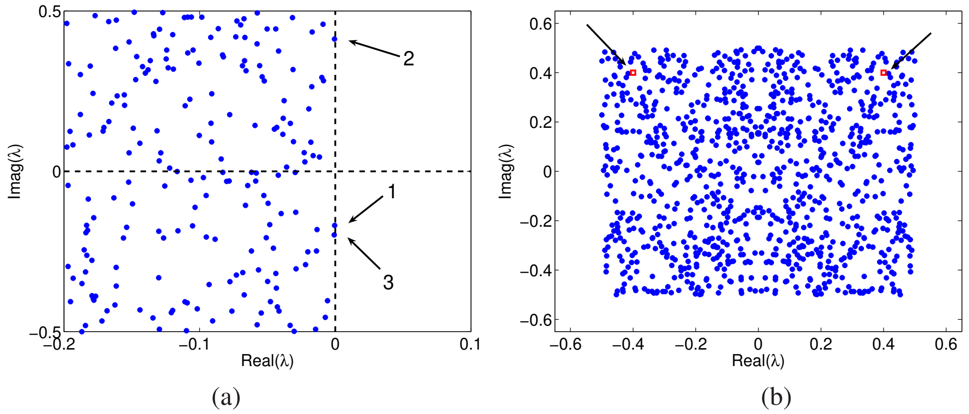

FIG. 7.1. *(a) Part of the spectrum of a* 1000×1000 *matrix with eigenvalues distributed in the square with corners* $(-\frac{1}{2}i, \frac{1}{2}i, -1+\frac{1}{2}i, -1-\frac{1}{2}i)$. *The first three rightmost eigenvalues are indicated by arrows. (b) The spectrum and targets (indicated by arrows) of a random* 1000×1000 *(diagonal) Hamiltonian matrix.*

In Table 7.1, we summarize the percentages of the cases in which Jacobi–Davidson with the three different types of extraction method converged to the rightmost eigenvalues. All methods used the same set of 100 random initial vectors.

TABLE 7.1
Percentage of the cases in which convergence to the (first, second, third) rightmost eigenvalues occurred for the 1000×1000 *diagonal and bidiagonal matrix with random eigenvalues in* $[-1, 0] \times [-i/2, i/2]$, *and the* 400×400 *matrix* $\mathrm{diag}(-398, \ldots, -1, \pm 52i)$ *for standard JD, the (1,1)-harmonic approach (with different targets), and the (2,2)-harmonic approach (with different target sets), each with 100 different random initial vectors.*

Method	τ	diag	bidiag	"$\pm 52i$"
Standard JD	–	(39, 57, 1)%	(15, 41, 0)%	3%
(1,1)-harmonic	1	(36, 39, 3)%	(7, 34, 0)%	0%
(1,1)-harmonic	2	(33, 44, 1)%	(13, 42, 0)%	0%
(1,1)-harmonic	$1+3i$	(46, 44, 0)%	(19, 38, 1)%	0%
(1,1)-harmonic	$1+10i$	(38, 56, 1)%	(17, 44, 0)%	0%
(1,1)-harmonic	$1+20i$	(41, 54, 1)%	(17, 42, 0)%	0%
(2,2)-harmonic	$0.1 \pm i$	(36, 23, 5)%	(10, 46, 0)%	100%
(2,2)-harmonic	$0.1 \pm 3i$	(38, 50, 0)%	(15, 45, 1)%	100%
(2,2)-harmonic	$1 \pm i$	(26, 30, 5)%	(10, 37, 0)%	100%
(2,2)-harmonic	$1 \pm 3i$	(35, 49, 2)%	(13, 34, 0)%	99%

The second column contains the zero(s) of the p-polynomial in the (1,1)-harmonic and (2,2)-harmonic approaches. The third and fourth column contain the number of instances where the methods converged to the first, second, and third rightmost eigenvalue. The fifth column describes the number of cases where the methods detected either one of $\pm 52i$.

A main conclusion from Table 7.1 is that finding the rightmost eigenvalue—although it is located at the exterior of the spectrum—is not an easy task. Standard JD has some limited success for the first two matrices but almost completely fails for the third. The (1,1)-harmonic and (2,2)-harmonic show roughly similar behavior for the first two matrices; in particular, the (2,2)-harmonic approaches show no advantage over the cheaper two competitors.

A big difference is the third test matrix. Standard JD always converges to $\lambda = -1$ while the (1,1)-harmonic converges to the *leftmost* eigenvalue $\lambda = -398$ in all cases. This is caused by the fact that with the choices of the target τ in the table, the $r(\lambda) = p(\lambda)/q(\lambda)$-value of $\lambda = -398$ is very close to 1, which is also the $r(\lambda)$-value of $\lambda = \pm 52i$, while the other eigenvalues have larger $r(\lambda)$-values. Hence, although the (1,1)-harmonic technique finds the *leftmost* instead of the rightmost eigenvalue, it does what it is designed for: finding the eigenvalue with (almost) the smallest $r(\lambda)$-value. The (2,2)-harmonic approach (nearly) always converge to the correct $\lambda = \pm 52i$.

Although of secondary interest to us, we note that the average number of outer iterations for the new harmonic approaches was typically roughly the same ($\pm 10\%$) in comparison to the standard Rayleigh–Ritz extraction, where one takes the rightmost eigenvalue in every step.

Some words about the target placement. From the results in Table 7.1 we see that although the location of the target(s) has some influence, the results are not overly sensitive on the choice of the target. In particular, it seems not essential that the zeros of p and q are close to the spectrum in a relative or absolute sense. On the contrary, it may be advantageous to choose them somewhat remote from the spectrum; see also the contour plot Figure 4.1(a). This is practical in the situation where we do not have a precise idea about the location of the spectrum.

We remark that sometimes the Jacobi–Davidson method is started with a Krylov space of a certain dimension to get a decent initial space. We did several experiments with this technique, and the results were better in some cases, but worse in others (not reported here). We tested the (2,1)-harmonic technique with $p(z) = (z - \tau_1)(z - \tau_2)$ and $q(z) = (z - \tau_3)$ as well, but the results were not competitive to the other approaches, which is not surprising in view of the level curves of Figure 4.2(a).

Our conclusion is that the new approaches may be at least competitive approaches in finding the rightmost eigenvalue, although success is not always guaranteed and keeping in mind that the (2,2)-harmonic approach requires one extra MV per outer step.

EXPERIMENT 7.2. Next, we investigate numerically whether the new approaches can be useful for structured matrices, where we are interested in two targets. We take a random diagonal Hamiltonian matrix of size 1000×1000, of which the spectrum is depicted in Figure 4.1(b). We build up a ten-dimensional search space $\mathcal{U}$ to find the eigenpairs corresponding to the eigenvalues $\lambda_\pm \approx \pm 0.4159 + 0.3966i$ closest to the targets $\tau = \pm 0.4 + 0.4i$. The first two basis vectors of $\mathcal{U}$ are taken to be $u_\pm = x_\pm + \varepsilon w_\pm$, where $x_\pm$ are the eigenvectors corresponding to $\lambda_\pm$ and $w_\pm$ are random vectors of unit length. We complement $\mathcal{U}$ by eight random vectors. For ε we take the values 1, 0.5, and 0.1, resulting in search spaces ranging from worse to better quality.

As stated before, the extraction is especially of importance at the situation of a restart. Therefore, we simulate the situation of a restart by reducing the 10-dimensional search spaces to a 4-dimensional search space by each of the extraction methods. The results are contained in Table 7.2.

TABLE 7.2
Extraction results of the different approaches a restart of a 1000×1000 Hamiltonian matrix from a 10-dimensional to a 4-dimensional subspaces. The 10-dimensional subspaces contains the true eigenvectors perturbed by respectively ε = 1, 0.5, and 0.1.

Method	$\angle(\mathcal{U}, x_1)$	$\angle(\mathcal{U}, x_2)$	$\angle(\mathcal{U}, x_1)$	$\angle(\mathcal{U}, x_2)$	$\angle(\mathcal{U}, x_1)$	$\angle(\mathcal{U}, x_2)$
Before restart	$5.0e-1$	$4.9e-1$	$2.6e-1$	$2.6e-1$	$5.4e-2$	$5.2e-2$
Harmonic	$1.0e+0$	$1.5e+0$	$1.2e+0$	$2.7e-1$	$5.4e-2$	$5.3e-2$
(2,0)-harmonic	$5.5e-1$	$1.5e+0$	$2.7e-1$	$1.6e+0$	$5.4e-2$	$1.6e+0$
(2,1)-harmonic	$1.0e+0$	$5.3e-1$	$2.7e-1$	$2.6e-1$	$5.4e-2$	$5.3e-2$
(2,2)-harmonic	$5.4e-1$	$5.4e-1$	$2.7e-1$	$2.6e-1$	$5.4e-2$	$5.3e-2$
Refined	$5.0e-1$	$1.6e+0$	$2.6e-1$	$1.6e+0$	$5.4e-2$	$1.6e+0$
Double refined	$5.7e-1$	$5.5e-1$	$4.8e-1$	$4.7e-1$	$5.6e-1$	$5.5e-1$

The second row displays the angles of the search space with the eigenvectors before the restart, the other rows the angles after the restart. The (standard) harmonic and refined approach use only $\tau = \tau_+ = 0.4 + 0.4i$. Since an eigenpair is also a harmonic Ritz pair, it is natural that if the search space is good enough, the standard harmonic approach will approximate both eigenvectors well, as can be seen in the case $\varepsilon = 0.1$. However, if the search spaces are less accurate ($\varepsilon = 1$ and 0.5), as will often be the case during the pre-convergence stage in an iterative method, the harmonic approach fails to approximate both eigenvectors well. The (2,2)-harmonic method in particular, and to a lesser extent the (2,1)-harmonic approach and the double refined approach, approximate the eigenvectors from a less accurate search space better. The (standard) refined approach with one target $\tau = \tau_+$ does what it is supposed to do: it approximates one of the two eigenvectors quite well, while the other eigenvector is completely neglected. We plan to incorporate the novel extraction methods in structure preserving iterative methods in future work.

8. Conclusions. We have examined various generalizations of harmonic and refined Rayleigh–Ritz, which can be useful in at least two cases:

- when our eigenvalue problem has structure, and we would like to have a better extraction process than standard Rayleigh–Ritz without sacrificing (implications of) the structure of the problem;
- when we are interested in eigenvalues close to a subset of the complex plane that consists of more than just one point; for instance, the rightmost eigenvalues close to the line $\mathrm{Re}(z) = 0$.

The results of the new methods look promising, both from theoretical and numerical perspective. The costs are relatively modest. When we use two linear polynomials p and q, there are no extra matrix-vector products necessary; if p and/or q has degree two, we need one extra MV per outer iteration step. In view of the number of MVs necessary for the inner iteration, this may be relatively little. The other computations, such as a QR decomposition or a projected eigenvalue problem, are low-dimensional operations.

If the extra computational cost is still considered much, we may execute the new extraction processes at restarts only, since a restart is the moment where a good vector extraction is crucial. This reduces part of the overhead, although we still need an extra MV per step in case $\deg(p) = 2$ or $\deg(q) = 2$.

Also of interest is the stability. For (very) ill-conditioned A, working with polynomials of degree more than one (or two) does not seem attractive because of the high condition of $p(A)$. On the other hand, because of the computation by a QR decomposition, we do not have to fear $[\kappa(p(A))]^2$ in the computations. The pole and zero placement in the harmonic approaches, although not of large influence, is an interesting future question; see [14] for some results for the standard harmonic Ritz method.

This paper focussed on how to extract sensible approximate eigenpairs having a (reasonable) search space. How to generate such a search space is very important but different question. If we are interested in the rightmost eigenvalue, we may determine the best approximate eigenpair by one of the methods in this paper and use Jacobi–Davidson [15] or Arnoldi (see, e.g, [18, Ch. VII]) for the subspace expansion. When we are interested in the eigenvalues of a structured eigenvalue close to one of two targets, we may want to expand the search space in a way that does justice to both targets. We leave this for future work.

Acknowledgments The author would like to thank Bor Plestenjak, Gerard Sleijpen, and Alastair Spence for inspirational discussions and the referees for very constructive and helpful comments.

REFERENCES

[1] C. Beattie, *Harmonic Ritz and Lehmann bounds*, Electron. Trans. Numer. Anal., 7 (1998), pp. 18–39, http://etna.math.kent.edu/vol.7.1998/pp18-39.dir/pp18-39.html.
[2] G. Chen and Z. Jia, *An analogue of the results of Saad and Stewart for harmonic Ritz vectors*, J. Comput. Appl. Math., 167 (2004), pp. 493–498.
[3] M. E. Hochstenbach and G. L. G. Sleijpen, *Harmonic and refined extraction methods for the polynomial eigenvalue problem*, Department of Mathematics, Case Western Reserve University, December 2005. Revised version. Submitted.
[4] Z. Jia, *Refined iterative algorithms based on Arnoldi's process for large unsymmetric eigenproblems*, Linear Algebra Appl., 259 (1997), pp. 1–23.
[5] ———, *The convergence of harmonic Ritz values, harmonic Ritz vectors, and refined harmonic Ritz vectors*, Math. Comp., 74 (2005), pp. 1441–1456.
[6] A. V. Knyazev, V. I. Lebedev, and A. L. Skorokhodov, *The Temple-Lehmann-type methods in iterative algorithms*, Soviet J. Numer. Anal. Math. Modelling, 5 (1990), pp. 265–273.
[7] K. Meerbergen and D. Roose, *Matrix transformations for computing rightmost eigenvalues of large sparse non-symmetric eigenvalue problems*, IMA J. Numer. Anal., 16 (1996), pp. 297–346.
[8] ———, *The restarted Arnoldi method applied to iterative linear system solvers for the computation of rightmost eigenvalues*, SIAM J. Matrix Anal. Appl., 18 (1997), pp. 1–20.
[9] K. Meerbergen and M. Sadkane, *Using Krylov approximations to the matrix exponential operator in Davidson's method*, Appl. Numer. Math., 31 (1999), pp. 331–351.
[10] K. Meerbergen, A. Spence, and D. Roose, *Shift-invert and Cayley transforms for detection of rightmost eigenvalues of nonsymmetric matrices*, BIT, 34 (1994), pp. 409–423.
[11] V. Mehrmann and D. Watkins, *Structure-preserving methods for computing eigenpairs of large sparse skew-Hamiltonian/Hamiltonian pencils*, SIAM J. Sci. Comput., 22 (2001), pp. 1905–1925.
[12] C. C. Paige, B. N. Parlett, and H. A. van der Vorst, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Num. Lin. Alg. Appl., 2 (1995), pp. 115–133.
[13] Y. Saad, *Numerical Methods for Large Eigenvalue Problems*, Manchester University Press, Manchester, UK, 1992.
[14] G. L. G. Sleijpen and J. van den Eshof, *On the use of harmonic Ritz pairs in approximating internal eigenpairs*, Linear Algebra Appl., 358 (2003), pp. 115–137.
[15] G. L. G. Sleijpen and H. A. van der Vorst, *A Jacobi–Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 401–425.
[16] G. L. G. Sleijpen, H. A. van der Vorst, and E. Meijerink, *Efficient expansion of subspaces in the Jacobi–Davidson method for standard and generalized eigenproblems*, Electron. Trans. Numer. Anal., 7 (1998), pp. 75–89, http://etna.math.kent.edu/vol.7.1998/pp75-89.dir/pp75-89.html.
[17] G. W. Stewart, *Matrix algorithms. Vol. II*, SIAM, Philadelphia, PA, 2001.
[18] H. A. van der Vorst, *Computational methods for large eigenvalue problems*, in Handbook of numerical analysis, Vol. VIII, North-Holland, Amsterdam, 2002, pp. 3–179.