



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Computer Physics Communications 167 (2005) 90–102

Computer Physics
Communications

www.elsevier.com/locate/cpc

Computing eigenvalue bounds for iterative subspace matrix methods

Yunkai Zhou^{a,b}, Ron Shepard^{a,*}, Michael Minkoff^b

^a *Theoretical Chemistry Group, Chemistry Division, Argonne National Laboratory, Argonne, IL 60439, USA*

^b *Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA*

Received 24 November 2004; received in revised form 2 December 2004; accepted 23 December 2004

Abstract

A procedure is presented for the computation of bounds to eigenvalues of the generalized hermitian eigenvalue problem and to the standard hermitian eigenvalue problem. This procedure is applicable to iterative subspace eigenvalue methods and to both outer and inner eigenvalues. The Ritz values and their corresponding residual norms, all of which are computable quantities, are needed by the procedure. Knowledge of the exact eigenvalues is not needed by the procedure, but it must be known that the computed Ritz values are isolated from exact eigenvalues outside of the Ritz spectrum and that there are no skipped eigenvalues within the Ritz spectrum range. A multipass refinement procedure is described to compute the bounds for each Ritz value. This procedure requires $O(m)$ effort where m is the subspace dimension for each pass.

Published by Elsevier B.V.

PACS: 02.10; 02.60; 02.70; 89.80

Keywords: Bounds; Eigenvalue; Subspace; Ritz; Hermitian; Generalized; Gap; Spread

1. Background

The *generalized hermitian eigenvalue problem*

$$(\mathbf{H} - \lambda_k \mathbf{M})\mathbf{z}_k = 0 \quad (1.1)$$

* Corresponding author. Tel.: +630-252-3584, fax: 630-252-4470.

E-mail addresses: zhou@mcs.anl.gov (Y. Zhou), shepard@tcg.anl.gov (R. Shepard), minkoff@mcs.anl.gov (M. Minkoff).

occurs in many applications. In this equation the matrices are of dimension N , $\mathbf{H} = \mathbf{H}^\dagger$, $\mathbf{M} = \mathbf{M}^\dagger$, and \mathbf{M} is positive definite. An important special case is the *standard hermitian eigenvalue problem* for which $\mathbf{M} = \mathbf{1}$. All of the results in this manuscript apply to both the generalized and the standard problems. One approach to the numerical solution of Eq. (1.1) is to expand the approximate eigenvectors in a linearly independent basis $\{\mathbf{x}_j; j = 1 \dots m\}$. These vectors may be collected into the matrix

$$\mathbf{X} = [\mathbf{x}_1 \mid \mathbf{x}_2 \mid \dots \mid \mathbf{x}_m]. \quad (1.2)$$

The projected subspace representation of the matrices \mathbf{H} and \mathbf{M} are denoted $\langle \mathbf{H} \rangle = \mathbf{X}^\dagger \mathbf{H} \mathbf{X}$ and $\langle \mathbf{M} \rangle = \mathbf{X}^\dagger \mathbf{M} \mathbf{X}$. In many cases, the basis vectors in \mathbf{X} will be chosen such that $\langle \mathbf{M} \rangle = \mathbf{1}$, but we consider the general case hereafter. An approximation \mathbf{y} to an eigenvector will be written as a linear combination of the basis vectors, $\mathbf{y} = \mathbf{X} \mathbf{c}$ (for simplicity, the vector norm relation $\mathbf{y}^\dagger \mathbf{M} \mathbf{y} = 1$ is assumed hereafter). A measure of the error of the vector \mathbf{y} is the residual vector

$$\mathbf{r} = (\mathbf{H} - \rho \mathbf{M}) \mathbf{y}. \quad (1.3)$$

$\mathbf{r} = 0$ if and only if \mathbf{y} is an exact eigenvector and ρ is the corresponding exact eigenvalue. Otherwise, as discussed in more detail in Appendix A, the quantity $\sqrt{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}}$, hereafter called the *residual norm*, is a measure of the error in both the eigenvector and the eigenvalue. The expansion coefficients \mathbf{c} are usually determined from the subspace eigenvalue equation

$$(\langle \mathbf{H} \rangle - \rho_j \langle \mathbf{M} \rangle) \mathbf{c}_j = 0. \quad (1.4)$$

It is convenient to assume that the subspace eigenvalues $\{\rho_j; j = 1 \dots m\}$, called the *Ritz values*, and the exact eigenvalues $\{\lambda_k; k = 1 \dots N\}$ are indexed in increasing order. This is not a requirement in order to apply the results of this manuscript, but this assumption simplifies the subsequent notation and discussions. It is convenient to use the Parlett index convention [1] for which negative integers are used to index eigenvalues ordered from highest to lowest (i.e. $\lambda_{-1} \equiv \lambda_N, \lambda_{-2} \equiv \lambda_{N-1}, \dots, \lambda_{-(N-1)} \equiv \lambda_2, \lambda_{-N} \equiv \lambda_1$). We refer to a set of the lowest eigenvalues $\{\lambda_k; k = 1 \dots m\}$, or to a set of the highest eigenvalues $\{\lambda_{-k}; k = 1 \dots m\}$ as *outer* eigenvalues, we refer to either or both of the highest and lowest eigenvalues $\{\lambda_1, \lambda_N\}$ as *extremal* eigenvalues, and we refer to a sequence of eigenvalues that are in the interior of the spectrum $\{\lambda_k; k = m \dots n; 1 < m \leq n < N\}$ as *inner* eigenvalues.

The above choice for the expansion coefficients \mathbf{c} and approximate eigenvalue ρ is optimal in several respects [1,2]. In particular, this choice minimizes the residual norm with respect to variations in ρ and with respect to the vector coefficients. Iterative subspace methods solve Eq. (1.4) several times as the subspace dimension m is increased or decreased and as the basis vectors are expanded, contracted, or otherwise changed during the iterative procedure. The various subspace methods differ in how the individual basis vectors \mathbf{x}_j are generated, in how the basis vectors are

contracted in order to satisfy various resource limitation constraints, and in how preconditioners are used in order to accelerate the convergence of the iterative procedures for the particular eigenpairs of interest. The bounds relations examined in this manuscript apply to all of these various hermitian subspace methods (including the Lanczos [3], Davidson [4], SPAM [5], Generalized Davidson Inverse Iteration [6], Jacobi–Davidson [7], and Generalized Jacobi–Davidson [8] methods).

The present work focuses on assessing the accuracy of the computed eigenvalues. Lower and upper bounds, b_j^- and b_j^+ , respectively, are desired that satisfy

$$b_j^- \leq \lambda_k \leq b_j^+. \quad (1.5)$$

As discussed in more detail below, the mapping of the Ritz value index j and the eigenvalue index k depends on whether the eigenvalues are inner or outer, and on whether the highest or lowest eigenvalues are computed. The further goal is that these bounds may be computed during the iterative procedure, so that they may be used not only to assess the accuracy of the final results, but also to allow the iterative procedure to be terminated when a predetermined accuracy has been achieved in order to avoid unnecessary computational effort.

Several standard inequalities are used to this end. The first is the Ritz Bound from the Ritz variational principle [1,9] (see Eqs. (A.3)–(A.16)).

$$\lambda_j \leq \rho_j, \quad \rho_{-j} \leq \lambda_{-j} \quad \text{for } j = 1 \dots m. \quad (1.6)$$

This bound relation places strict upper bounds on each of the lowest m exact outer eigenvalues, and it places strict lower bounds on the highest m exact outer eigenvalues. No additional information beyond the Ritz values themselves is required, so these bounds are computable.

The second inequality used is the Residual Norm Bound [1] (see Eqs. (A.17)–(A.28)), which requires the Ritz values along with the corresponding residual norms.

$$\sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j} \geq |\lambda_k - \rho_j|. \quad (1.7)$$

These bounds are computable, but they are often very conservative.

The next inequality used in this procedure is the Gap Theorem Bound [1,10] (see Eqs. (A.29)–(A.46)).

$$\frac{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}{\gamma_j} \geq |\lambda_k - \rho_j|. \tag{1.8}$$

The Gap Bound is tighter than the Residual Norm Bound when $\sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j} < \gamma_j$, but it requires knowledge of the exact eigenvalues, and therefore it cannot be computed during the iterative process. However, if a Ritz value of interest ρ_j is separated in the following sense

$$\begin{aligned} \delta_j^- &< (\rho_j - \sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}), \\ (\rho_j + \sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}) &< \delta_j^+ \end{aligned} \tag{1.9}$$

with

$$\delta_j^- = \text{Max}_{k < j} \{b_k^+\}, \quad \delta_j^+ = \text{Min}_{k > j} \{b_k^-\} \tag{1.10}$$

and if it is further assumed that there are no skipped eigenvalues (i.e. eigenvalues that are poorly represented or not represented at all) within the computed spectrum range, then

$$\gamma_j^- \equiv \text{Min} \{ \rho_j - \delta_j^-, \delta_j^+ - \rho_j \} \leq \gamma_j \tag{1.11}$$

is a lower bound to the exact gap γ_j . This separation is shown in Fig. 1. This separation is equivalent to the condition that there exists no expansion vector, or set of expansion vectors, that would move the Ritz value of interest outside of the range $[\rho_j - \sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}, \rho_j + \sqrt{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}]$, or that would reveal some other eigenvalue inside the range $[\delta_j^-, \delta_j^+]$. This computable value

γ_j^- may be combined with Eq. (1.8),

$$\frac{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}{\gamma_j^-} \geq \frac{\mathbf{r}_j^\dagger \mathbf{M}^{-1} \mathbf{r}_j}{\gamma_j} \geq |\lambda_j - \rho_j|. \tag{1.12}$$

When it is necessary to distinguish which bound is being discussed, Eq. (1.8), with the exact eigenvalue gap, will be called the *Exact Gap Bound*, whereas Eq. (1.12), with γ_j^- , will be called the *Computed Gap Bound*. With the above separation conditions, the Computed Gap Bound is tighter than the Residual Norm Bound, although it is not as tight as the Exact Gap Bound.

A fourth bound that will be used is the Spread Bound [1]. This applies only to the extremal eigenvalues (see Eqs. (A.49)–(A.53)).

$$\frac{\mathbf{r}_k^\dagger \mathbf{M}^{-1} \mathbf{r}_k}{\sigma} \leq |\lambda_k - \rho_k| \quad \text{for } k = \pm 1. \tag{1.13}$$

The Spread Theorem requires knowledge of the exact matrix spread, $\sigma = (\lambda_N - \lambda_1)$, which, although computable in some situations, is usually not available. However, if an upper bound σ^+ is available such that $\sigma \leq \sigma^+$, then useful inner bounds can be computed for the two extremal eigenvalues. For example, when computing the largest eigenvalue of a positive definite matrix, the upper bound to the eigenvalue is itself an upper bound to the matrix spread. Combining this upper bound with Eq. (1.13) results in the computable bound

$$\frac{\mathbf{r}_k^\dagger \mathbf{M}^{-1} \mathbf{r}_k}{\sigma^+} \leq \frac{\mathbf{r}_k^\dagger \mathbf{M}^{-1} \mathbf{r}_k}{\sigma} \leq |\lambda_k - \rho_k| \quad \text{for } k = \pm 1. \tag{1.14}$$

Subspace methods are often used in situations in which the matrices \mathbf{H} and \mathbf{M} are not explicitly com-




Fig. 1. The inner Ritz value ρ_j is isolated from the other Ritz values within the subspace by the Residual Norm bound. The exact gap is γ , and the computed lower bound to that gap is γ^- . When the Ritz value is isolated in this manner, the upper and lower bounds to the exact eigenvalue λ_k computed by the Gap Theorem are tighter than those from the Residual Norm Bounds.

puted and stored; instead, the products of these matrices with the basis vectors \mathbf{X} are computed in operator form. Bounds that require information about individual matrix elements could not be used in these important applications because these individual matrix elements are either not available or they would be prohibitively expensive to extract from the matrix–vector product operation. The goal of the present work is to find the best computable eigenvalue bounds given a set of Ritz values, their residual norms, some knowledge of the eigenvalue spacings, and, optionally, an estimate of the matrix spread. This information is available for subspace methods even when the individual matrix elements are not explicitly available for examination. There are several issues that must be addressed toward this goal.

- (1) Under what conditions are the isolated Ritz value and the nonskipped eigenvalue conditions satisfied?
- (2) The error bounds for the different eigenvalues are coupled in the sense that b_j depends on the other b_k bounds for $k \neq j$. How can all bounds be converged self-consistently without violating the bound conditions at intermediate steps during the optimization process or at the final values for these bounds?
- (3) For inner eigenvalues, the upper and lower error bounds on the eigenvalues are symmetric about the computed Ritz value ρ_j . For outer eigenvalues, the upper and lower error bounds are not symmetric. How can this be addressed during the bounds calculations?

These issues are addressed in order. In general, given only the computed Ritz values and their corresponding residual norms, it cannot be guaranteed that there is no skipped eigenvalue. The skipped eigenvalue depends on information outside of the current subspace \mathbf{X} . However, in many applications, the general qualitative structure of the eigenvalues of interest is known beforehand. For example, the eigenpairs may be known at a lower approximation level, on a coarser grid, with a smaller underlying model basis set, at nearby values of some set of external parameters, or from other physical insights known to the user. In addition, the eigenvalue spacings may be known more accurately than the actual eigenvalues

themselves. The matrix may be known to be, for example, positive semidefinite, or positive definite, or negative definite, or other upper or lower bounds may be known based on physical insights of the model upon which the numerical eigenproblem is derived. Furthermore, the rank of any degeneracies may be known for the eigenvalue spectrum based on, for example, physical or mathematical invariance properties, group theory, or other physical or mathematical insights. In these situations, the goal of the numerical calculation is to compute quantitatively the eigenpairs, the qualitative values being already known to some extent. This additional knowledge of the eigenvalue spectrum may be used to allow application of the Gap and the Spread Bounds even though the exact eigenvalues are not known.

When computing the initial bounds from the residual norms, two situations occur for each of the Ritz values: either a Ritz value is separated from the other values, or it is not separated. If it is separated, and if there are no skipped eigenvalues, then one and only one exact eigenvalue will occur in the range $[\rho_k - \sqrt{\mathbf{r}_k^+ \mathbf{M}^{-1} \mathbf{r}_k}, \rho_k + \sqrt{\mathbf{r}_k^+ \mathbf{M}^{-1} \mathbf{r}_k}]$. If two (or more) Ritz values are not separated, then in the general case it is known only that a single eigenvalue exists in the combined residual norm value range. That is, the approximate eigenvectors associated with the two Ritz values may either both be approximating the same exact eigenpair, or they may be approximating two distinct eigenpairs. Consider, for example, the standard eigenvalue problem with

$$\mathbf{H} = \begin{pmatrix} 0 & \varepsilon & 0 \\ \varepsilon & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}. \quad (1.15)$$

For this problem $\lambda = \{-\sqrt{1 + \varepsilon^2}, 0, \sqrt{1 + \varepsilon^2}\}$, $\rho = \{-\varepsilon, \varepsilon\}$, and $|\mathbf{r}_1| = |\mathbf{r}_2| = 1/\sqrt{2}$. Consider first the situation with $\varepsilon \approx 0.1$. Both Ritz values are near the single $\lambda_2 = 0$ exact eigenvalue, and it is only for this single exact eigenvalue that the Residual Norm Bound, for either Ritz value, is seen to apply. The two Ritz values are not separated by the residual norm bound, thus the eigenvalue bounds cannot be refined with the Gap Bound. Furthermore, $\lambda_1 < (\rho_1 - |\mathbf{r}_1|)$ and $\lambda_{-1} > (\rho_{-1} + |\mathbf{r}_{-1}|)$, so it is clear that this subspace representation does not satisfy the nonskipped eigenvalue requirement for either of the nonzero eigenvalues. On the other hand, consider this same problem

with $\varepsilon \approx 10$. In this case the Ritz values are separated by the Residual Norm Bound, and the two Ritz values are good approximations to the two nonzero extremal exact eigenvalues. However, there is a skipped eigenvalue in this case, and the Computed Gap Bounds cannot be applied because lower bounds to the exact gaps cannot be determined from the subspace information. As noted above, it cannot be known simply by examining the residual norms or the subspace eigenvalues that an eigenvalue has been skipped, this information must be determined separately.

The coupling of the bounds for one eigenvalue to the computed bounds for other, usually adjacent, eigenvalues results in the necessity to compute a sequence of bound values. At present, a closed-form solution to this problem is not known to the authors. Consequently, we adopt a straightforward bootstrap approach to this problem. We begin the procedure by assigning Residual Norm Bounds to each of the eigenvalues (along with the Ritz and Spread bounds if appropriate). Then, given the current bounds, all of the bounds for each isolated eigenvalue are refined if possible, and any improvements to the bounds are retained for subsequent passes. Depending on the eigenvalue spacings and the corresponding bounds, the final bounds may be computed in a single refinement pass or they may require multiple passes. The bounds that are computed in this process are valid at any intermediate step, and this allows the process to be truncated while still returning rigorous bounds. For example, once the bounds are below a given threshold, indicating convergence to some level of accuracy of the eigenvalue, the process could be terminated if desired. The application of Eqs. (1.9)–(1.12) appears to require a scan of m elements to compute the upper and lower gap bounds for each Ritz value; this in turn implies $O(m^2)$ effort to refine the bounds for all m Ritz values each pass. We avoid this effort by precomputing the δ^+ array before the process begins, and then updating these elements, along with the corresponding element of δ^- , in decreasing order as each eigenvalue bound is refined. This results overall in only $O(m)$ effort for each pass.

Because the outer eigenvalue bounds are not symmetric about the computed Ritz values, this introduces an asymmetry into the calculation procedure. Consider, for example, the computation of bounds for the lowest few m eigenvalues. First note that the gap as-

sociated with ρ_m cannot be estimated. This is because there is no ρ_{m+1} estimate available that could be used to approximate this gap. The refinement process therefore begins with ρ_{m-1} and proceeds down, in some order, eventually to the bounds for ρ_1 . Secondly, note that the computed gap for ρ_1 depends only on the bound δ_1^+ whereas the computed gaps for the remaining eigenvalues must be estimated by examining both the lower bounds to the higher eigenvalues and the upper bounds for the lower eigenvalues. Finally the upper bounds are given by the Ritz variational principle for all of the eigenvalues except for λ_1 , and for this eigenvalue the computed Spread Bound may be applied optionally (provided that an upper bound to the matrix spread is available), resulting in the situation in which the Ritz value lies outside of the computed eigenvalue range $b_1^- \leq \lambda_1 \leq b_1^+ < \rho_1$. The same considerations apply also to the computation of the highest m eigenvalues.

As a practical matter, the computation of the residual norm $\sqrt{\mathbf{r}_j^T \mathbf{M}^{-1} \mathbf{r}_j}$ for the above bounds calculations requires either the explicit inversion of the metric matrix \mathbf{M} or the solution of the linear equation $\mathbf{M}\mathbf{s} = \mathbf{r}_j$. The inverse may be computed with relatively little effort when the metric \mathbf{M} is a unit matrix (i.e. the standard eigenvalue problem), a diagonal matrix, or a block-diagonal matrix with small diagonal blocks. The linear equation solution is relatively simple in some other applications, including, for example, the tridiagonal or narrow-banded forms that occur in finite difference methods for which sparse LU or sparse Cholesky methods may be applied. In other applications for which the matrix \mathbf{M} has no exploitable structure, the linear equation solution requires a more general iterative approach, the effort for which might range from minimal (e.g., using a preconditioned conjugate gradient approach with a very efficient preconditioner) to something comparable to that required for the eigenproblem itself. It is assumed herein that some suitable method for computing the above residual norm is available for the target problem such that the bounds refinement process described above is practical.

2. Examples

In this section a few examples are given to demonstrate the use of the bounds computations described in

Section 1. The first example is a model with $m = 5$ and an unspecified N . The model Ritz values correspond to the lowest five eigenvalues. The model residual norms are sufficiently small for the Ritz values to be isolated from each other, and it is assumed that these Ritz values are also isolated from any other exact eigenvalues. The specific numerical values are chosen to allow easy computation by hand in order to understand the refinement process in detail. Table 1 summarizes the upper and lower bounds during the refinement process. During Pass 0, the lower bounds are initialized with the Residual Norm Bounds. The upper bound of the lowest eigenvalue is set with the Spread Bound with $\sigma^+ = 10.0$. The remaining upper bounds are initialized with the Ritz bound. During Pass 1, the lower bounds for the lowest four eigenvalues are refined using the Gap Bound expression. Only the bounds that are changed each pass are displayed in

the table. All of the bounds are converged after Pass 1; in the subsequent pass all of the bounds are checked for possible improvement and the refinement process is terminated. It is clear that the final Gap and Spread Bounds are a significant improvement over the original Residual Bounds.

Table 2 shows the refinement process for the same model problem except the Ritz values are assumed to correspond to inner exact eigenvalues. There are assumed to be an undetermined number of eigenvalues below the first Ritz value and above the fifth Ritz value. In Pass 0 both the upper and lower bound corresponding to each Ritz value are initialized with the Residual Norm Bound. As discussed in Section 1, the Gap Bounds cannot be applied to the highest and lowest Ritz values because there is no information available about eigenvalues outside of the Ritz range, so lower bounds to the gaps for these Ritz values can-

Table 1
Bound refinement for the lowest five Ritz values

j	1	2	3	4	5
ρ_j	1.000000	2.000000	3.000000	4.000000	5.000000
$ \mathbf{r}_j $	0.010000	0.010000	0.010000	0.010000	0.010000
Pass 0 b_j^-	0.990000 ⁽¹⁾	1.990000 ⁽¹⁾	2.990000 ⁽¹⁾	3.990000 ⁽¹⁾	4.990000 ⁽¹⁾
b_j^+	0.999990 ⁽³⁾	2.000000 ⁽⁰⁾	3.000000 ⁽⁰⁾	4.000000 ⁽⁰⁾	5.000000 ⁽⁰⁾
Pass 1 b_j^-	0.999900 ⁽²⁾	1.999900 ⁽²⁾	2.999900 ⁽²⁾	3.999899 ⁽²⁾	–
b_j^+	–	–	–	–	–

⁽⁰⁾ Ritz bound.
⁽¹⁾ Residual bound.
⁽²⁾ Gap bound.
⁽³⁾ Spread bound with $\sigma^+ = 10.0$.

Table 2
Bound refinement for five inner Ritz values

j	1	2	3	4	5
ρ_j	1.000000	2.000000	3.000000	4.000000	5.000000
$ \mathbf{r}_j $	0.010000	0.010000	0.010000	0.010000	0.010000
Pass 0 b_j^-	0.990000 ⁽¹⁾	1.990000 ⁽¹⁾	2.990000 ⁽¹⁾	3.990000 ⁽¹⁾	4.990000 ⁽¹⁾
b_j^+	1.010000 ⁽¹⁾	2.010000 ⁽¹⁾	3.010000 ⁽¹⁾	4.010000 ⁽¹⁾	5.010000 ⁽¹⁾
Pass 1 b_j^-	–	1.999899 ⁽²⁾	2.999899 ⁽²⁾	3.999899 ⁽²⁾	–
b_j^+	–	2.000101 ⁽²⁾	3.000101 ⁽²⁾	4.000101 ⁽²⁾	–
Pass 2 b_j^-	–	–	2.999900 ⁽²⁾	–	–
b_j^+	–	–	3.000100 ⁽²⁾	–	–

⁽¹⁾ Residual bound.
⁽²⁾ Gap bound.

not be computed. Consequently, during Pass 1 only the eigenvalue bounds corresponding to the three middle Ritz values are refined. In contrast to the first example in Table 1, another pass is required to refine the middle eigenvalue bound. All of the bounds are converged after Pass 2; in the subsequent pass all of the bounds are checked for possible improvement and the refinement process is terminated. It is clear that the final Gap Bound is a significant improvement over the original Residual Norm Bound for the three inner Ritz values for which they can be applied.

Fig. 2 and Table 3 show a practical application of the bounds calculations to each iteration of a subspace iterative procedure for a standard eigenvalue problem. The matrix has dimension $N = 197,655,128$. This is the real symmetric representation of the quantum mechanical electronic Hamiltonian for the ethylene molecule, C_2H_4 , using a cc-pVTZ orbital basis set computed with the COLUMBUS Program System [11]. This is an example of an application in which individual matrix elements of \mathbf{H} are not available, the required matrix–vector products are computed in operator form. The expansion basis consists of all single and double replacements from an MCSCF reference expansion with 12 active valence orbitals and two frozen-core orbitals (the $1s$ core orbitals of the two carbon atoms). A preliminary approximate B_k calculation [12] was performed on the lowest two eigenpairs in order to generate two qualitatively accurate starting

vectors for the Ritz procedure using the exact matrix. With these two starting vectors, the lowest Ritz value is isolated as shown in Table 3. The expansion vectors used during the iterative procedure are calculated us-

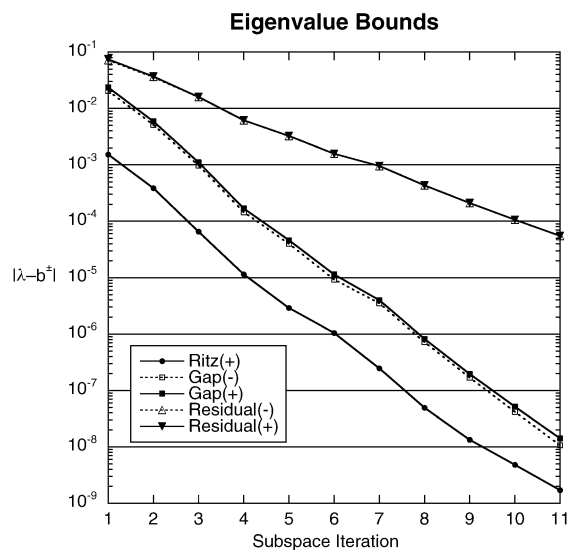


Fig. 2. This graph displays the improvement of the various bounds as a function of the subspace iteration for the lowest eigenvalue of the C_2H_4 hamiltonian matrix. The plotted values are computed as $|\lambda_1 - b_1^\pm|$, where λ_1 is taken as the converged eigenvalue and b_1 is either an upper b_1^+ or lower b_1^- eigenvalue bound for that subspace iteration from the Ritz Variational Theorem, the Residual Norm Theorem, and the Gap Theorem as discussed in the text.

Table 3

Upper and lower bounds for the lowest eigenvalue¹ for each subspace iteration

Sub. iter.	ρ_1	$ \mathbf{r}_1 $	ρ_2	$ \mathbf{r}_2 $	$b_1^+ - b_1^-$	$b_1^+ - \lambda_1$	$\lambda_1 - b_1^-$
1	-78.4232628319	7.307E-02	-78.0877800384	9.372E-02	2.209E-02	1.528E-03	2.056E-02
2	-78.4244055909	3.657E-02	-78.0877800858	9.372E-02	5.505E-03	3.857E-04	5.119E-03
3	-78.4247258142	1.594E-02	-78.0877801056	9.372E-02	1.045E-03	6.548E-05	9.797E-04
4	-78.4247798274	6.168E-03	-78.0877801305	9.372E-02	1.564E-04	1.146E-05	1.449E-04
5	-78.4247883942	3.243E-03	-78.0877801323	9.372E-02	4.323E-05	2.896E-06	4.033E-05
6	-78.4247902447	1.583E-03	-78.0877801336	9.372E-02	1.029E-05	1.046E-06	9.249E-06
7	-78.4247910433	9.545E-04	-78.0877801388	9.372E-02	3.745E-06	2.470E-07	3.498E-06
8	-78.4247912409	4.337E-04	-78.0877801426	9.372E-02	7.730E-07	4.940E-08	7.236E-07
9	-78.4247912769	2.107E-04	-78.0877801451	9.372E-02	1.824E-07	1.340E-08	1.690E-07
10	-78.4247912855	1.066E-04	-78.0877801454	9.372E-02	4.666E-08	4.800E-09	4.186E-08
11	-78.4247912886	5.502E-05	-78.0877801465	9.372E-02	1.244E-08	1.700E-09	1.074E-08

¹ The Ritz values correspond to the lowest two subspace eigenvalues for a matrix of dimension 197,655,128. b_1^+ is computed from the Ritz bound (i.e. $b_1^+ = \rho_1$), and b_1^- is computed from the Gap Bound (i.e. using $\gamma_1^- = \rho_2 - |\mathbf{r}_2| - \rho_1$). The converged Ritz value $\rho_1 = -78.4247912903$ is taken as the exact eigenvalue λ_1 , so the last two columns of this table are calculated only after the iterative procedure has converged.

ing the usual Davidson preconditioned residual [4,12] approach. As can be seen, these expansion vectors improve the lowest Ritz value selectively, and the second Ritz value and its corresponding residual norm, are modified very little during the iterative procedure. The actual subspace dimension is two for the first iteration, increases to six on the fifth iteration, then is contracted to three on the sixth iteration, then increases up to six on the ninth iteration, and is contracted to three on the tenth iteration. These contractions of the expansion subspace (sometimes called *restarts*) are performed in order to reduce the resource requirements [12] (e.g., memory, internode communications, and external disk space), but as seen in Fig. 2 and Table 3, they do not affect significantly the convergence rate for this calculation. Only the lowest two Ritz values and corresponding residual norms are used in the bounds refinement process. This calculation clearly demonstrates how the Computed Gap Bound can be applied in a useful way even when poorly converged Ritz values are used to compute the lower bounds to the gaps.

Fig. 2 shows how the various upper and lower bounds approach the exact eigenvalue as a function of the subspace iteration procedure. It is clear from Fig. 2 that the best upper bound each subspace iteration is given by the Ritz bound and that the best lower bound is given by the Gap Bound; these values are shown in more detail in Table 3. b_1^+ in Table 3 is computed from the Ritz bound and b_1^- in Table 3 is computed from the Gap Bound. As seen in the last two columns, the expected relation $b_1^- \leq \lambda_1 \leq b_1^+$ is satisfied throughout the procedure. The converged Ritz value is taken as the exact eigenvalue λ_1 , so the last two columns of this table are calculated only after the iterative procedure has converged. For this particular calculation, the upper bound for a particular subspace iteration is always closer to the exact eigenvalue than the lower bound for that iteration. It is also interesting to compare the Gap Bounds to the Residual Norm Bounds each iteration. In every iteration, the Gap Bounds are much tighter than the Residual Norm Bounds, the ratio of the two radii approaching 10^4 on the last iteration.

If convergence of the eigenpair to an absolute accuracy of 10^{-4} in the Ritz value were required (equivalent to a relative accuracy of about one part in 10^6 for the lowest eigenvalue), then, with perfect-hindsight knowledge of the final eigenvalue, the iterative process could have been terminated on the fifth iteration. The

$(b_1^+ - b_1^-)$ column of Table 3 is a combination of the Ritz upper bound with the Gap lower bound; these bounds do not require knowledge of the exact eigenvalue and therefore can be computed during the subspace iterative procedure. As seen in Table 3, these bounds would also allow termination on the fifth iteration for this absolute accuracy requirement. In contrast, the Residual Norm Bound would require 11 iterations in order to guarantee convergence to 10^{-4} in the Ritz value. If convergence to an absolute accuracy of 10^{-6} in the Ritz value were required, then the Ritz and Gap Bounds would allow the process to terminate on the eighth iteration, which is again consistent with perfect hindsight knowledge of the eigenvalue; in contrast, the Residual Norm Bound would require about 19 iterations to guarantee convergence of the Ritz value to 10^{-6} . For these types of accurate, large-scale, molecular electronic structure calculations, typical convergence requirements in the Ritz values range from 10^{-4} to 10^{-8} in absolute accuracy.

3. Conclusions

A procedure is presented for the computation of bounds to eigenvalues of the generalized hermitian eigenvalue problem and to the standard hermitian eigenvalue problem. This procedure is applicable to iterative subspace eigenvalue methods and to both outer and inner eigenvalues. The Ritz values and their corresponding residual norms, all of which are computable quantities, are needed by the procedure. Knowledge of the exact eigenvalues is not needed by the procedure, but it must be known that the computed Ritz values are isolated from exact eigenvalues outside of the Ritz spectrum and that there are no skipped eigenvalues within the Ritz spectrum range. A multipass refinement procedure is described to compute the bounds for each Ritz value. This procedure requires $O(m)$ effort where m is the subspace dimension for each pass.

Application of this bounds computation procedure to model problems and to actual production problems demonstrates the usefulness of the procedure. This procedure can be applied during the subspace iterative procedure in order to truncate the iterative process and to avoid unnecessary effort when converging specific eigenvalues to a required target accuracy.

Acknowledgements

This work was supported by the U.S. Department of Energy by the Office of Basic Energy Sciences, Division of Chemical Sciences, and by the Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Science Division, under contract W-31-109-ENG-38.

Appendix A. Bounds

In this appendix, several of the eigenvalue bounds that are used in the algorithm of Section 1 are discussed. These bounds include the Ritz Variational Bound, the Residual Norm bound, the Gap Theorem Bound, and the Spread Theorem bound. These bounds are well known for the standard eigenvalue problem [1,2], and the derivations that are given in this appendix are included primarily for completeness and also to extend the bounds to the generalized eigenvalue problem. To this end, we use the generalized hermitian eigenvalue equation

$$(\mathbf{H} - \lambda_k \mathbf{M})\mathbf{z}_k = 0 \quad (\text{A.1})$$

or in the equivalent matrix form

$$\mathbf{H}\mathbf{Z} = \mathbf{M}\mathbf{Z}\lambda \quad (\text{A.2})$$

with an explicit metric matrix \mathbf{M} , but the proofs for the standard hermitian eigenvalue problem with $\mathbf{M} = \mathbf{1}$ follow directly.

Ritz Variational Bounds. Consider first the situation in which a subspace of dimension m is available. This subspace results in the projected equation

$$\langle \mathbf{H} \rangle^{[m]} \mathbf{C}^{[m]} = \langle \mathbf{M} \rangle^{[m]} \mathbf{C}^{[m]} \rho^{[m]}. \quad (\text{A.3})$$

The superscript denotes the subspace dimension. When a new basis vector is added, the corresponding projected equation is

$$\langle \mathbf{H} \rangle^{[m+1]} \mathbf{C}^{[m+1]} = \langle \mathbf{M} \rangle^{[m+1]} \mathbf{C}^{[m+1]} \rho^{[m+1]}. \quad (\text{A.4})$$

We define the transformation matrix \mathbf{T} in partitioned form as

$$\mathbf{T} = \begin{pmatrix} \mathbf{C}^{[m]} & \boldsymbol{\beta} \\ \mathbf{0} & \alpha \end{pmatrix} \quad (\text{A.5})$$

with

$$\begin{aligned} \langle \mathbf{M} \rangle^{[m+1]} &= \begin{pmatrix} \langle \mathbf{M} \rangle^{[m]} & \mathbf{w} \\ \mathbf{w}^\dagger & x \end{pmatrix}, \\ \mathbf{q} &= (\langle \mathbf{M} \rangle^{[m]})^{-1} \mathbf{w}, \\ \alpha &= (x - \mathbf{w}^\dagger \mathbf{q})^{-1/2}, \\ \boldsymbol{\beta} &= -\alpha \mathbf{q}. \end{aligned} \quad (\text{A.6})$$

This results in

$$\mathbf{T}^{-1} = \begin{pmatrix} (\mathbf{C}^{[m]})^{-1} & (\mathbf{C}^{[m]})^{-1} \mathbf{q} \\ \mathbf{0} & 1/\alpha \end{pmatrix} \quad (\text{A.7})$$

Eq. (A.4) may be transformed as

$$\begin{aligned} (\mathbf{T}^\dagger \langle \mathbf{H} \rangle^{[m+1]} \mathbf{T}) (\mathbf{T}^{-1} \mathbf{C}^{[m+1]}) \\ = (\mathbf{T}^\dagger \langle \mathbf{M} \rangle^{[m+1]} \mathbf{T}) (\mathbf{T}^{-1} \mathbf{C}^{[m+1]}) \rho^{[m+1]} \end{aligned} \quad (\text{A.8})$$

and written in partitioned form as

$$\begin{pmatrix} \rho^{[m]} & \mathbf{h} \\ \mathbf{h}^\dagger & g \end{pmatrix} \begin{pmatrix} \mathbf{U} \\ \mathbf{v} \end{pmatrix} = \begin{pmatrix} \mathbf{U} \\ \mathbf{v} \end{pmatrix} \rho^{[m+1]}. \quad (\text{A.9})$$

Eq. (A.9) has the same eigenvalues as the original equation (A.4). The matrix \mathbf{T} has transformed the generalized eigenvalue equation into a standard eigenvalue equation. There are, of course, an infinite number of such transformations that would achieve that goal. The particular \mathbf{T} chosen above also transforms the first m -dimensional leading sub-block to diagonal form. An arbitrary eigenpair from Eq. (A.9) may then be written

$$\begin{pmatrix} \rho^{[m]} & \mathbf{h} \\ \mathbf{h}^\dagger & g \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ v \end{pmatrix} = \begin{pmatrix} \mathbf{u} \\ v \end{pmatrix} \rho^{[m+1]}. \quad (\text{A.10})$$

The first row can be solved for the vector \mathbf{u} in terms of the scalar v . When substituted into the second row to eliminate \mathbf{u} , the result is

$$g - \mathbf{h}^\dagger (\rho^{[m]} - \rho \mathbf{1})^{-1} \mathbf{h} = \rho. \quad (\text{A.11})$$

The left-hand side of Eq. (A.11) may be written

$$L(\rho) = g - \sum_{j=1}^m \frac{h_j^2}{(\rho_j^{[m]} - \rho)}. \quad (\text{A.12})$$

It is easily verified that $L(\rho)$ has horizontal asymptotes

$$\lim_{\rho \rightarrow -\infty} L(\rho) = \lim_{\rho \rightarrow +\infty} L(\rho) = g. \quad (\text{A.13})$$




Fig. 3. This is an example of the Cauchy Interlace Relation for a matrix of dimension four and for a submatrix of dimension three. The vertical asymptotes (dashed lines) are located at the eigenvalues of the submatrix. The eigenvalues of the full matrix of dimension four, indicated by the open circles, are the intersections of the decreasing branches of the function $L(\rho)$ and of the increasing function $R(\rho) = \rho$. The eigenvalues of the full matrix interlace those of the submatrix.

Furthermore

$$\frac{dL(\rho)}{d\rho} = -\sum_{j=1}^m \frac{h_j^2}{(\rho_j^{[m]} - \rho)^2} \leq 0$$

for $-\infty < \rho < \infty$ and $\rho \neq \rho_j^{[m]}$. (A.14)

Thus $L(\rho)$ is a decreasing function everywhere that it is analytic (differentiable). The vertical asymptotes are the result of the simple pole structure of $L(\rho)$. The analytic segment of $L(\rho)$ that lies between two of the poles is a branch of the function $L(\rho)$, and there are $m + 1$ such branches. An example of this function is plotted for $m = 3$ in Fig. 3. According to Eq. (A.11), the eigenvalues $\rho_k^{[m+1]}$ for $k = 1 \dots m + 1$ are those special values of this function for which $L(\rho) = R(\rho)$ where $R(\rho) = \rho$ is a simple linear function with positive unit slope. The intersections of these two functions are shown in Fig. 3. Because the branches of $L(\rho)$ are decreasing functions, and $R(\rho)$ is an increas-

ing function, there is one and only one intersection for each of the branches. These intersections must therefore satisfy

$$\begin{aligned} \rho_1^{[m+1]} &\leq \rho_1^{[m]}, \\ \rho_j^{[m]} &\leq \rho_{j+1}^{[m+1]} \leq \rho_{j+1}^{[m]} \quad \text{for } j = 1 \dots (m - 1), \\ \rho_m^{[m]} &\leq \rho_{m+1}^{[m+1]}. \end{aligned} \quad (\text{A.15})$$

This is the Cauchy interlace relation [1]. Consider this inequality applied to a particular eigenvalue index j (or $-j$, using the Parlett index convention) as the subspace dimension m approaches the matrix dimension N .

$$\begin{aligned} \lambda_j &= \rho_j^{[N]} \leq \rho_j^{[N-1]} \leq \dots \\ &\leq \rho_j^{[m+2]} \leq \rho_j^{[m+1]} \leq \rho_j^{[m]}, \\ \rho_{-j}^{[m]} &\leq \rho_{-j}^{[m+1]} \leq \rho_{-j}^{[m+2]} \leq \dots \\ &\leq \rho_{-j}^{[N-1]} \leq \rho_{-j}^{[N]} = \lambda_{-j}. \end{aligned} \quad (\text{A.16})$$

This is the Ritz variational principle. In words, the m ordered Ritz values from a subspace of dimension m are each upper bounds to the corresponding lowest m exact eigenvalues of the full matrix, and those Ritz values are simultaneously lower bounds to the highest m exact eigenvalues of the full matrix.

For brevity, we do not consider here the special cases individually that require detailed attention. These include the situations in which $v = 0$ for a particular eigenvector and the situations in which either or both $\rho^{[m]}$ and $\rho^{[m+1]}$ contain degeneracies. We note only that these situations are straightforward to analyze and that the general Ritz variational principle holds in these special situations just as it does in the typical case discussed above.

Residual Norm Theorem. Given an hermitian matrix \mathbf{H} , a positive-definite metric hermitian matrix \mathbf{M} , and vector \mathbf{y} , normalized as $\mathbf{y}^\dagger \mathbf{M} \mathbf{y} = 1$, the Rayleigh quotient is

$$\rho(\mathbf{y}, \mathbf{H}, \mathbf{M}) \equiv \frac{\mathbf{y}^\dagger \mathbf{H} \mathbf{y}}{\mathbf{y}^\dagger \mathbf{M} \mathbf{y}} = \mathbf{y}^\dagger \mathbf{H} \mathbf{y}. \quad (\text{A.17})$$

For brevity in this appendix, we use hereafter the short notation $\rho \equiv \rho(\mathbf{y}, \mathbf{H}, \mathbf{M})$. The following proofs hold for arbitrary Rayleigh quotients; the Ritz values used in Section 1 are Rayleigh quotients that satisfy also

additional conditions. The residual vector $\mathbf{r} \equiv \mathbf{r}(\mathbf{y})$ is

$$\mathbf{r} = (\mathbf{H} - \rho\mathbf{M})\mathbf{y}. \quad (\text{A.18})$$

Note that $\mathbf{y}^\dagger \mathbf{r} = 0$. The vector \mathbf{y} may be expanded in the basis of the exact eigenvectors $\{\mathbf{z}_j; j = 1 \dots N\}$, which may be chosen to be orthonormal, $\mathbf{z}_j^\dagger \mathbf{M} \mathbf{z}_k = \delta_{jk}$.

$$\mathbf{y} = \sum_j \mathbf{z}_j \alpha_j \quad (\text{A.19})$$

with

$$1 = \mathbf{y}^\dagger \mathbf{M} \mathbf{y} = \sum_{j,k} \mathbf{z}_j^\dagger \mathbf{M} \mathbf{z}_k \alpha_j \alpha_k = \sum_j \alpha_j^2. \quad (\text{A.20})$$

This allows the residual to be written in the eigenvector basis in terms of the exact eigenvalues λ_j .

$$\mathbf{r} = \sum_j (\mathbf{H} - \rho\mathbf{M})\mathbf{z}_j \alpha_j = \sum_j (\lambda_j - \rho)\alpha_j \mathbf{M} \mathbf{z}_j. \quad (\text{A.21})$$

Let λ_k be the exact eigenvalue closest to ρ

$$|\lambda_k - \rho| = \text{Min}\{|\lambda_j - \rho|; j = 1 \dots N\}. \quad (\text{A.22})$$

The quantity $(\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r})$ may be written

$$\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r} = \sum_j (\lambda_j - \rho)^2 \alpha_j^2. \quad (\text{A.23})$$

Each term in this summation is nonnegative. The inequality

$$\begin{aligned} \mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r} &\geq \sum_{j=1}^N (\lambda_k - \rho)^2 \alpha_j^2 = (\lambda_k - \rho)^2 \sum_{j=1}^N \alpha_j^2 \\ &= (\lambda_k - \rho)^2 \end{aligned} \quad (\text{A.24})$$

follows by replacing each nonnegative term in the summation in Eq. (A.23) by a nonnegative value of lesser or equal magnitude. This results in the bound

$$\sqrt{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}} \geq |\lambda_k - \rho| \quad (\text{A.25})$$

which may be written in the form

$$\rho - \sqrt{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}} \leq \lambda_k \leq \rho + \sqrt{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}}. \quad (\text{A.26})$$

In the special case of the standard eigenvalue equation, $\mathbf{M} = \mathbf{1}$, these inequalities result in the familiar Residual Norm Bound

$$|\mathbf{r}| \geq |\lambda_k - \rho|, \quad (\text{A.27})$$

$$\rho - |\mathbf{r}| \leq \lambda_k \leq \rho + |\mathbf{r}|. \quad (\text{A.28})$$

There are several other interesting proofs of the Residual Norm Bound (see, for example, Chapter 4 of Parlett [1]). The above proof was chosen for this appendix because it is in the same spirit as the following proof of the Gap Theorem Bound.

It should be emphasized that although the above Residual Norm Bound is rigorous, this does not imply immediately that $\rho_k - \sqrt{\mathbf{r}_k^\dagger \mathbf{M}^{-1} \mathbf{r}_k} \leq \lambda_k \leq \rho_k + \sqrt{\mathbf{r}_k^\dagger \mathbf{M}^{-1} \mathbf{r}_k}$ for a given Ritz and eigenpair index k . The bound only states that, given a Rayleigh quotient ρ , there is at least one eigenvalue index k that satisfies the bound. Additional knowledge that there are no skipped eigenvalues, or knowledge of exactly which eigenvalues are skipped, is required in order to match up the index of a Ritz value ρ_j with an exact eigenvalue λ_k .

The Gap Theorem. The Gap Bound is slightly more involved. We present here the proof given by Parlett [1], but extended in a straightforward manner to the generalized eigenvalue problem. For the Rayleigh quotient $\rho(\mathbf{y}, \mathbf{H}, \mathbf{M})$, the gap $\gamma \equiv \gamma(\rho)$ is defined as

$$\gamma = \text{Min}\{|\lambda_j - \rho|; j = 1 \dots N, j \neq k\}. \quad (\text{A.29})$$

In words, the gap γ is the distance from the Rayleigh quotient ρ to the nearest exact eigenvalue other than λ_k , the eigenvalue that is being approximated by ρ . It is assumed hereafter that λ_k is a nondegenerate eigenvalue and that the gap is strictly positive. The vector \mathbf{y} can be decomposed in the $\mathbf{y}-\mathbf{z}_k$ plane as

$$\mathbf{y} = \mathbf{z}_k \text{Cos } \phi + \mathbf{w} \text{Sin } \phi \quad (\text{A.30})$$

with $\mathbf{w}^\dagger \mathbf{M} \mathbf{w} = 1$ and $\mathbf{w}^\dagger \mathbf{M} \mathbf{z}_k = 0$. We note that the phases of the vectors \mathbf{z}_k and \mathbf{w} are not specified in any way, and the quadrant of the plane rotation angle ϕ is arbitrary to within these phases, but this is unimportant to the proof at hand. This plane decomposition allows the residual to be written

$$\mathbf{r} = (\lambda_k - \rho)\mathbf{M} \mathbf{z}_k \text{Cos } \phi + (\mathbf{H} - \rho\mathbf{M})\mathbf{w} \text{Sin } \phi. \quad (\text{A.31})$$

This results in the relation

$$\begin{aligned} 0 &= \mathbf{y}^\dagger \mathbf{r} \\ &= (\lambda_k - \rho) \text{Cos}^2 \phi + \mathbf{w}^\dagger (\mathbf{H} - \rho\mathbf{M})\mathbf{w} \text{Sin}^2 \phi. \end{aligned} \quad (\text{A.32})$$

It is clear from this expression that the two factors $(\lambda_k - \rho)$ and $(\mathbf{w}^\dagger (\mathbf{H} - \rho\mathbf{M})\mathbf{w})$ have opposite signs.

Eq. (A.32) can be solved for $\text{Cos}^2 \phi$ to give

$$\text{Cos}^2 \phi = -\frac{\mathbf{w}^\dagger(\mathbf{H} - \rho\mathbf{M})\mathbf{w}}{(\lambda_k - \rho)} \text{Sin}^2 \phi. \quad (\text{A.33})$$

Substitution of the identity $\text{Cos}^2 \phi + \text{Sin}^2 \phi = 1$ into Eqs. (A.32) or (A.33) results in

$$\begin{aligned} \text{Sin}^2 \phi &= \frac{-(\lambda_k - \rho)}{\mathbf{w}^\dagger(\mathbf{H} - \rho\mathbf{M})\mathbf{w} - (\lambda_k - \rho)} \\ &= \frac{-(\lambda_k - \rho)}{\mathbf{w}^\dagger(\mathbf{H} - \rho\mathbf{M} - \lambda_k\mathbf{M} + \rho\mathbf{M})\mathbf{w}} \\ &= \frac{-(\lambda_k - \rho)}{\mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M})\mathbf{w}} \\ &= \frac{|\lambda_k - \rho|}{|\mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M})\mathbf{w}|}. \end{aligned} \quad (\text{A.34})$$

The quantity $(\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r})$ may be written

$$\begin{aligned} \mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} &= (\lambda_k - \rho)^2 \text{Cos}^2 \phi + \mathbf{w}^\dagger(\mathbf{H} - \rho\mathbf{M}) \\ &\quad \times \mathbf{M}^{-1}(\mathbf{H} - \rho\mathbf{M})\mathbf{w} \text{Sin}^2 \phi. \end{aligned} \quad (\text{A.35})$$

Substitution of Eq. (A.33) gives

$$\begin{aligned} \mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} &= -\mathbf{w}^\dagger(\mathbf{H} - \rho\mathbf{M})\mathbf{w}(\lambda_k - \rho) \text{Sin}^2 \phi \\ &\quad + \mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M} + \lambda_k\mathbf{M} - \rho\mathbf{M}) \\ &\quad \times \mathbf{M}^{-1}(\mathbf{H} - \rho\mathbf{M})\mathbf{w} \text{Sin}^2 \phi \\ &= \mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M})\mathbf{M}^{-1}(\mathbf{H} - \rho\mathbf{M})\mathbf{w} \text{Sin}^2 \phi. \end{aligned} \quad (\text{A.36})$$

Expanding the vector \mathbf{w} in the eigenvector basis with coefficients χ_j gives

$$\begin{aligned} \mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} &= \text{Sin}^2 \phi \sum_{j(\neq k)} \mathbf{z}_j^\dagger(\mathbf{H} - \lambda_k\mathbf{M}) \\ &\quad \times \mathbf{M}^{-1}(\mathbf{H} - \rho\mathbf{M})\mathbf{z}_j \chi_j^2 \end{aligned} \quad (\text{A.37})$$

$$= \text{Sin}^2 \phi \sum_{j(\neq k)} (\lambda_j - \lambda_k)(\lambda_j - \rho) \chi_j^2 \quad (\text{A.38})$$

$$= \text{Sin}^2 \phi \sum_{j(\neq k)} |\lambda_j - \lambda_k| \cdot |\lambda_j - \rho| \chi_j^2. \quad (\text{A.39})$$

Eq. (A.39) holds because the two factors $(\lambda_j - \lambda_k)$ and $(\lambda_j - \rho)$ are always the same sign for each term in the summation. Replacing each of the positive terms $|\lambda_j - \rho|$ by the smaller value γ results in the inequality

$$\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} \geq \gamma \text{Sin}^2 \phi \sum_{j(\neq k)} |\lambda_j - \lambda_k| \chi_j^2. \quad (\text{A.40})$$

The inequality

$$\begin{aligned} \sum_{j(\neq k)} |\lambda_j - \lambda_k| \chi_j^2 &\geq \left| \sum_{j(\neq k)} (\lambda_j - \lambda_k) \chi_j^2 \right| \\ &= |\mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M})\mathbf{w}| \end{aligned} \quad (\text{A.41})$$

results in the relation

$$\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} \geq \gamma |\mathbf{w}^\dagger(\mathbf{H} - \lambda_k\mathbf{M})\mathbf{w}| \text{Sin}^2 \phi. \quad (\text{A.42})$$

Substitution using Eq. (A.34) gives the inequalities

$$\frac{\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r}}{\gamma} \geq |\lambda_k - \rho|, \quad (\text{A.43})$$

$$\rho - \frac{\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r}}{\gamma} \leq \lambda_k \leq \rho + \frac{\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r}}{\gamma}. \quad (\text{A.44})$$

In the special case of the standard eigenvalue equation, $\mathbf{M} = \mathbf{1}$, these inequalities result in the familiar Gap Theorem Bounds

$$\frac{|\mathbf{r}|^2}{\gamma} \geq |\lambda_k - \rho|, \quad (\text{A.45})$$

$$\rho - \frac{|\mathbf{r}|^2}{\gamma} \leq \lambda_k \leq \rho + \frac{|\mathbf{r}|^2}{\gamma}. \quad (\text{A.46})$$

As for the previous Residual Norm Bound, this relation only states that, given a Rayleigh quotient ρ , there is at least one eigenvalue index k that satisfies the bound. Additional knowledge that there are no skipped eigenvalues, or knowledge of exactly which eigenvalues are skipped, is required in order to match up a Ritz value ρ_j with an exact eigenvalue λ_k . Compared to the Residual Norm Bound in Eq. (A.25), it may be verified that the Gap Bound is tighter than the Residual Norm Bound when $\sqrt{\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r}} < \gamma$. The Gap Bound may be computed even when this relation is not satisfied, but such a bound would be looser than the Residual Norm Bound in this case.

The Spread Theorem. The Spread Theorem applies only to the special cases $k = 1$ or $k = N$ (or equivalently to $k = -1$ or $k = -N$ using the Parlett index convention), that is, to the two extremal eigenvalues only. Replace each of the nonnegative terms $|\lambda_j - \rho|$ in Eq. (A.39) by the larger quantity $\sigma = (\lambda_N - \lambda_1)$. This quantity σ is called the *Spread* of the pencil (\mathbf{H}, \mathbf{M}) . This results in the inequality

$$\mathbf{r}^\dagger\mathbf{M}^{-1}\mathbf{r} \leq \sigma \cdot \text{Sin}^2 \phi \sum_{j(\neq k)} |\lambda_j - \lambda_k| \chi_j^2. \quad (\text{A.47})$$

The equality (compare Eq. (A.41))

$$\begin{aligned} \sum_{j(\neq k)} |\lambda_j - \lambda_k| \chi_j^2 &= \left| \sum_{j(\neq k)} (\lambda_j - \lambda_k) \chi_j^2 \right| \\ &= |\mathbf{w}^\dagger (\mathbf{H} - \lambda_k \mathbf{M}) \mathbf{w}| \end{aligned} \quad (\text{A.48})$$

follows from the fact that each of the terms $(\lambda_j - \lambda_k)$ has the same sign for an extremal eigenvalue. Combining Eqs. (A.34), (A.47), and (A.48) results in the inequality

$$\frac{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}}{\sigma} \leq |\lambda_k - \rho|. \quad (\text{A.49})$$

For $k = 1$ we have $|\lambda_k - \rho| = \rho - \lambda_1$ which results in an upper bound to the lowest eigenvalue

$$\lambda_1 \leq \rho - \frac{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}}{\sigma}. \quad (\text{A.50})$$

For $k = N$ we have $|\lambda_k - \rho| = \lambda_N - \rho$ which results in a lower bound to the highest eigenvalue

$$\rho + \frac{\mathbf{r}^\dagger \mathbf{M}^{-1} \mathbf{r}}{\sigma} \leq \lambda_N. \quad (\text{A.51})$$

The Spread Theorem therefore provides *inner* bounds to the two extremal eigenvalues. In the special case of the standard eigenvalue equation, $\mathbf{M} = \mathbf{1}$, these inequalities result in the inner bounds

$$\lambda_1 \leq \rho - \frac{|\mathbf{r}|^2}{\sigma} \quad (\text{A.52})$$

and

$$\rho + \frac{|\mathbf{r}|^2}{\sigma} \leq \lambda_N. \quad (\text{A.53})$$

From these expressions it is clear than the Spread Bound is always superior to the Ritz bound for any finite value of the matrix spread σ . If no information about the matrix spread is available, then in the limit $\sigma \rightarrow \infty$, the spread bound approaches the simple Ritz variational bound. In this sense, the Spread

Theorem may be thought of as an improvement to the Ritz Bound that utilizes limited additional information about the matrix spectrum and the residual norm.

References

- [1] B.N. Parlett, *The Symmetric Eigenvalue Problem*, Soc. for Industr. and Appl. Math./Prentice-Hall, Englewood Cliffs, NJ, 1998.
- [2] G.W. Stewart, *Eigensystems, Matrix Algorithms*, vol. II, Soc. for Industr. and Appl. Math., Philadelphia, 2001.
- [3] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *J. Res. Nat. Bur. Standards, Sec. B* 45 (1950) 225–280.
- [4] E.R. Davidson, The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices, *J. Comput. Phys.* 17 (1975) 87.
- [5] R. Shepard, A.F. Wagner, J.L. Tilson, M. Minkoff, The Subspace Projected Approximate Matrix (SPAM) Modification of the Davidson method, *J. Comput. Phys.* 172 (2001) 472–514.
- [6] J. Olsen, P. Jørgensen, J. Simons, Passing the one-billion limit in Full Configuration-Interaction (FCI) calculations, *Chem. Phys. Lett.* 169 (1990) 463.
- [7] G.L.G. Sleijpen, H.A. Van der Vorst, A Jacobi–Davidson iteration method for linear eigenvalue problems, *SIAM J. Matrix Anal. Appl.* 17 (1996) 401.
- [8] G.L.G. Sleijpen, A.G.L. Booten, D.R. Fokkema, H.A. Van der Vorst, Jacobi–Davidson type methods for generalized eigenproblems and polynomial eigenproblems, *BIT* 36 (1996) 595.
- [9] W. Ritz, Über eine neue Method zur Lösung Gewisser Variationsprobleme der Mathematischen Physik, *J. Reine Angew. Math.* 135 (1909) 1–61.
- [10] C. Davis, W. Kahan, The rotation of eigenvectors by a perturbation. III, *SIAM J. Numer. Anal.* 7 (1970) 1.
- [11] H. Lischka, R. Shepard, R.M. Pitzer, I. Shavitt, M. Dallos, T. Müller, P.G. Szalay, M. Seth, G.S. Kedziora, S. Yabushita, Z. Zhang, High-level multireference methods in the quantum-chemistry program system COLUMBUS: Analytic MR-CISD and MR-AQCC gradients and MR-AQCC-LRT for excited states, GUGA spin-orbit CI and parallel CI density, *Phys. Chem. Chem. Phys.* 3 (2001) 664–673.
- [12] R. Shepard, I. Shavitt, H. Lischka, Reducing I/O Costs for the eigenvalue procedure in large-scale CI calculations, *J. Comput. Chem.* 23 (2002) 1121–1125.