Electronic Transactions on Numerical Analysis.
Volume 7, 1998, pp. 163-181.

ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

RESTARTING TECHNIQUES FOR THE (JACOBI-)DAVIDSON SYMMETRIC EIGENVALUE METHODS *

ANDREAS STATHOPOULOS† AND YOUSEF SAAD‡

Abstract. The (Jacobi-)Davidson method, which is a popular preconditioned extension to the Arnoldi method for solving large eigenvalue problems, is often used with restarting. This has significant performance shortcomings, since important components of the invariant subspace may be discarded. One way of saving more information at restart is through "thick" restarting, a technique that involves keeping more Ritz vectors than needed. This technique and especially its dynamic version have proved very efficient for symmetric cases. A different restarting strategy for the Davidson method has been proposed in [14], motivated by the similarity between the spaces built by the Davidson and Conjugate Gradient methods. For the latter method, a three term recurrence implicitly maintains all required information.

In this paper, we consider the effects of preconditioning on the dynamic thick restarting strategy, and we analyze both theoretically and experimentally the strategy based on Conjugate Gradient. Our analysis shows that, in some sense, the two schemes are complementary, and that their combination provides an even more powerful technique. We also describe a way to implement this scheme without additional orthogonalizations or matrix multiplications.

Key words. Davidson, Jacobi-Davidson, Lanczos, Conjugate Gradient methods, eigenvalue, implicit restarting, deflation, preconditioning.

AMS subject classifications. 65F15.

1. Introduction. Many scientific and engineering applications require the solution of large, sparse, symmetric eigenvalue problems, $Au = \lambda u$, for a few of the lowest or highest (extreme) eigenvalues and eigenvectors (eigenpairs). The Lanczos method and its equivalent in the non-symmetric case, the Arnoldi method, have been used traditionally to solve these problems [18]. However, as the matrix size increases, clustering of the eigenvalues deteriorates the performance of these methods, and because the inverse of A cannot be computed directly, preconditioning becomes necessary to compensate for the loss of efficiency and robustness of iterative methods. The Davidson method and its generalization, the Jacobi-Davidson method [5, 13, 3, 20], are popular extensions to the Arnoldi method. Instead of extracting the eigenvectors from a generated Krylov space, these methods gradually build a different space by incorporating into the existing basis the approximate solution of a correction equation. Procedurally, the two methods are similar to the FGMRES method [19], and in this sense, we refer to the approximate solution of the correction equation as preconditioning.

Despite the benefits of preconditioning, for many hard problems the (Jacobi-)Davidson method may still require a large number of steps. Because the vector iterates must be saved for extracting the eigenvectors, the storage requirements are overwhelming. The problem is actually aggravated in the symmetric case, where the better theoretical framework and software has led researchers to consider matrices of huge size that allow only a few vectors to be stored. Even in the non-preconditioned Lanczos method where a three-term recurrence is known, orthogonality problems and spurious solutions prevent the application of the method for a large number of steps. For these reasons, many restarting variants of the Lanczos and (Jacobi-)Davidson methods are used in practice [4, 17, 21, 1, 7].

For the Lanczos method, the requirement of maintaining the tridiagonal matrix necessitates restarting the iteration with only one vector that is chosen as a linear combination of

* Received February 2, 1998. Accepted for publication August 3, 1998. Recommended by R. Lehoucq. Work supported by NSF grants DMR–9217287 and ASC 95-04038, and the Minnesota Supercomputer Institute

† Department of Computer Science, College of William and Mary, Box 8795, Williamsburg, Virginia 23187-8795, (andreas@cs.wm.edu).

‡ Department of Computer Science, University of Minnesota, 4-192 EE/CSci Bldg., Minneapolis, Minnesota, 55455-0154 (saad@cs.umn.edu.)

ETNA
Kent State University
etna@mcs.kent.edu

the most recent Ritz vectors. This has significant performance shortcomings, since important components of the invariant subspace may be discarded, and the Rayleigh-Ritz procedure cannot minimize over the whole Krylov subspace that had been generated before restarting. An efficient solution to this problem has been given by the Implicitly Restarted Lanczos method (IRL) [21, 1, 9, 10]. IRL provides an implicit way of applying a polynomial filter during restarting and thus removing unwanted spectral information. Implicit restarting provides also an elegant formulation of most previously proposed restarting schemes. For the (Jacobi-)Davidson method, restarting does not exhibit similar difficulties, because all the wanted Ritz vectors can be retained explicitly, and even additional information can be incorporated in the basis at restart. This latter characteristic proves especially useful to the method presented in this paper. Recently, it was shown that without preconditioning, and if the Ritz values are used as shifts in the polynomial filter, the restarted Davidson method builds the same search space as IRL [24, 12].

To reduce the adverse effects of restarting to the above methods, we can save more information at every restart. Recently, we investigated the idea of "thick restarting" [24], which implements this principle by keeping more Ritz vectors than needed. The question to be addressed is which, and how many Ritz vectors to retain at restart. For symmetric, non-preconditioned cases, a dynamic thick restarting scheme that keeps Ritz vectors on both sides of the spectrum has proved extremely efficient. With preconditioning, although it is still efficient, a less expensive scheme might provide similar benefits. For the restarted Davidson method, Murray et al. [14], and Van Lenthe et al. [25] proposed to keep, for each sought eigenvector, the two corresponding Ritz vectors from two successive iterations, the current and the previous one. The motivation stems from the observation that the (Jacobi-)Davidson and the Conjugate Gradient (CG) methods build the same space when solving for the eigenvector of a converged eigenvalue. Although this scheme has certain disadvantages, it has the impressive ability to store most of the required information compactly in one vector.

In this paper we present a theoretical justification for this restarting scheme, we identify its weaknesses which turn out to be complementary to the thick restarting ones, and we obtain a new powerful method by combining the two schemes. Moreover, we present an efficient implementation that does not require any additional orthogonalizations or matrix vector multiplications with the basis vectors. Finally, this is the only restarting scheme known to the authors that cannot be implemented with implicit restarting.

After briefly presenting the Davidson method and the dynamic thick restarting scheme in section 2, we examine the behavior of this scheme with preconditioning. In section 3, we present Murray's CG-based restarting scheme and show its potential benefits. Some theoretical results are given in section 4 to justify the characteristics of the CG-based scheme. In section 5, we confirm the weaknesses hinted by the theory for this scheme through selected numerical results and we propose a combination with thick restarting. Several numerical experiments demonstrate the effectiveness of the new approach, which is further enhanced by an efficient implementation outlined at the end of the section. The paper concludes with some suggestions for incorporating these ideas in eigenvalue software.

2. The Davidson method and thick restarting.

2.1. The Davidson method. Throughout this paper, and unless otherwise stated, we assume that the matrix A is symmetric of order N, with eigenpairs (λ_i, u_i) of which the l lowest (or highest) are sought. The Davidson method first appeared as a diagonally preconditioned version of the Lanczos method for the symmetric eigenproblem. Extensions, to both general preconditioners and to the nonsymmetric case have been given since [12, 3, 7]. The following describes the algorithm for the symmetric case, where the maximum basis size is $m > l$, and

ETNA
Kent State University
etna@mcs.kent.edu

at every restart $p \geq l$ Ritz vectors are retained. $Ortho$ denotes any stable orthogonalization procedure.

ALGORITHM 2.1. **Davidson**

0. Choose initial unit vectors $V_p = \{v_1, \ldots, v_p\}$
1. For $s = 0, 1, \ldots$
2. $w_i = Av_i,\ i = 1, \ldots, p \Leftrightarrow 1$
3. $T_{p-1} = (AV_{p-1}, V_{p-1})$
4. For $j = p, \ldots, m$
5. $w_j = Av_j$
6. $t_{i,j} = (w_j, v_i),\ i = 1, \ldots, j$, *the last column of* T_j
7. Compute some wanted eigenpair, say (μ, c) *of* T_j
8. $x^{(j)} = V_j c$ *and* $r = \mu x^{(j)} \Leftrightarrow Ax^{(j)}$,
the Ritz vector and its residual at the j-th iteration
9. Test $\|r\|$ *for convergence.*
If satisfied target a new vector and return to 7
10. Solve $M_{(s,j)}\delta = r$, *for* δ
11. $v_{j+1} = Ortho(V_j, \delta)$
12. Endfor
13. $V_p = \{x_i,$ *the p lowest Ritz vectors* $\},\ l \leq p < m$. *Restart*
14. Endfor

The preconditioning is performed by solving the equation at step 10, with $M_{(s,j)}$ approximating $(A \Leftrightarrow \mu I)$ in some sense. Originally, Morgan and Scott [13] proposed to solve approximately with some preconditioner the Generalized Davidson correction equation:

$$(A \Leftrightarrow \sigma I)\ \delta = r, \tag{2.1}$$

where σ is an approximation to the sought eigenvalue. In [20], Sleijpen et al. show that for stability, robustness, as well as efficiency, the operator $M_{(s,j)}$ should have a range orthogonal to x (the superscript $^{(j)}$ is implied). The method, called Jacobi-Davidson (JD), solves approximately the projected correction equation:

$$(I \Leftrightarrow xx^T)(A \Leftrightarrow \sigma I)(I \Leftrightarrow xx^T)\ \delta = (I \Leftrightarrow xx^T)(\mu I \Leftrightarrow A)x. \tag{2.2}$$

The projections can be easily applied if an iterative linear solver is used. For preconditioners that approximate A directly, such as incomplete factorizations and approximate inverses, the above orthogonality condition is enforced through an equivalent formulation known as Olsen method. In this paper we use the name "Davidson method" to refer to any of the above different ways (equation (2.1) or (2.2)) to perform the preconditioning step 10.

2.2. Thick restarting. A Davidson step is more expensive than a Lanczos one, requiring a solution of a Rayleigh-Ritz problem, computation of a residual, explicit orthogonalization against all the basis vectors, and additional storage of the vectors $w_i = Av_i$ for all basis vectors. In return, besides the preconditioning benefits, the algorithm is more flexible and robust than Lanczos, it can start with any number of initial vectors, and include in the basis any extra information that is available during the execution. Especially at restart, it can retain any number of approximate eigenvectors without the need of the implicit restarting framework. In [24] we called "thick restarting" a technique in iterative methods that restarts with more Ritz vectors than needed. We showed that in the absence of preconditioning and if exact shifts are used in the IRL, the IRL and the restarted Davidson methods build the same search space.

Reducing the amount of discarded information at every restart improves the convergence of iterative methods. On the other hand, when keeping too many vectors, the Davidson pro-

ETNA
Kent State University
etna@mcs.kent.edu

TABLE 2.1
Comparison of thick (TR(L)) and dynamic thick restarting (Dyn) with original Davidson (TR(5)) on some symmetric Harwell-Boeing matrices, with diagonal and approximate inverse preconditioners. The number of matrix vector multiplications is reported, with a maximum of 5000. Five smallest eigenvalues are sought. Residual norm threshold is 10^{-12} times the Frobenius norm of the matrix. Davidson uses a basis size of 20.

	No preconditioning			Diagonal preconditioning			Approximate Inverse		
Matrix	TR(5)	TR(11)	Dyn	TR(5)	TR(10)	Dyn	TR(5)	TR(12)	Dyn
LUND B	-	-	1347	774	396	349	909	381	338
BCSSTK05	1174	975	612	1322	465	409	358	247	251
BCSSTK16	3962	1333	676	2410	905	663	752	331	317
NOS3	2179	620	458	2096	878	664	524	253	258
NOS5	-	2016	921	2659	1401	819	837	387	354

cess can not effectively build additional basis vectors, and the orthogonalization process becomes a bottleneck. In [24] it was shown that if some Ritz vectors are retained, the Lanczos (non-preconditioned Davidson) process can be approximated gradually by another Lanczos process on a matrix from which the eigenvectors corresponding to these Ritz vectors have been deflated. This provided a criterion for choosing which Ritz vectors to keep at restart; the ones with Ritz values closest to the required eigenvalue, if deflated, increase the gap and thus the convergence rate to the wanted one. When looking for the smallest l eigenpairs, the "thick restarting" scheme retains $p > l$ Ritz pairs and we denote it by TR(p). For symmetric cases however, convergence depends on the gap ratio of the eigenvalues and therefore the other end of the spectrum is also of importance [15]. A more general form of thick restarting keeps the L lowest and R highest Ritz vectors. To choose these two numbers, we proposed a two sided, dynamic thick restarting technique, that captures the trade-off between better error reduction through more, non-restarted Lanczos steps, and larger gap ratios from a thicker restart. Our heuristic minimized the approximate error bound on Ritz values of the Lanczos process, which is described by a Chebyshev polynomial:

$$\frac{1}{C^2_{m-L-R}(1+2\gamma_i)} \approx 2e^{-2(m-L-R)\sqrt{\gamma_i}},$$

where $m \Leftrightarrow L \Leftrightarrow R$ Lanczos steps can be taken before a new restart is necessary, and $\gamma_i = \frac{\mu_i - \mu_{L+1}}{\mu_{L+1} - \mu_{m-R}}$ is the current approximation to the gap ratio of the $i \Leftrightarrow$th eigenvalue.

A few representative results using this strategy on matrices from the Harwell-Boeing collection [6] appear in Table 2.1. For all matrices, we solve for five lowest eigenpairs using a maximum basis size of 20, and we use the Davidson code that was described in [22] together with the extensions proposed in [23]. The number of matrix-vector multiplications decreases significantly when thick restarting of double the number of required eigenpairs is used, and it improves even further with the dynamic thick restarting scheme.

2.3. Dynamic thick restarting and preconditioning. The heuristic for dynamic thick restarting relies on the gap ratios, which govern the convergence of the non-preconditioned Lanczos process. Nevertheless, the dynamic scheme performs equally well when used with preconditioning. In Table 2.1, diagonal and approximate inverse [2] preconditioning have been used to demonstrate the effectiveness of the method.

In Table 2.1, we observe that although dynamic is better than one-sided thick restarting, the improvements diminish with better preconditioners (or easier problems). A closer look at the vectors retained by the dynamic scheme at every step reveals that a significant percentage of the Ritz vectors is retained, specifically the ones closest to the required eigenpairs. In addition, a small (1-3), varying number of highest Ritz pairs is often retained. Thus, it is

ETNA
Kent State University
etna@mcs.kent.edu

reasonable that dynamic outperforms one-sided thick restarting, since it retains at least as much information. On the other hand, by continuously keeping a large number of vectors in the basis, the dynamic scheme is more expensive on the average because of orthogonalization and computation of residuals. These operations can become a bottleneck, if the matrix vector multiplication is cheap. A theoretical explanation of the diminishing difference between thick and dynamic thick restarting in the presence of preconditioning, can be given by interpreting a theorem from Meerbergen ([11], Thm 3.2.2) in the following qualitative way.

THEOREM 2.1. *Let the preconditioned Davidson operator at some iteration be denoted by* $M_C = (M \Leftrightarrow \sigma I)^{-1}(A \Leftrightarrow \mu I)$ *for some shift* σ*. Then, for any eigenpair* (λ_i, u_i) *of* A*, there is an eigenpair* (η_i, w_i) *of* M_C*, and a constant* c*, which depends on* u_i, M_C *and the isolation of* η_i *from other eigenvalues of* M_C*, so that the eigenvector error is bounded by:*

$$\|u_i \Leftrightarrow w_i\|_1 \leq c\frac{|\lambda_i \Leftrightarrow \mu|}{|\lambda_i \Leftrightarrow \sigma|}.$$

According to the theorem, and provided that c and $1/|\lambda_i \Leftrightarrow \sigma|$ stay bounded, if μ is close to some eigenvalue of A, then the corresponding eigenvector of A is approximated by some eigenvector of M_C. Even though the w_i is not computed explicitly, the M_C operator is used for computing the basis vectors in Davidson. As a result, relatively good approximations are obtained for the eigenpairs close to the required one (close to value μ), while eigenvectors on the other side of the spectrum are not approximated well by the basis built with this preconditioner. Therefore, the attempts to deflate the high-end of the spectrum by keeping the highest Ritz vectors are not effective, and this explains the diminishing difference between thick and dynamic thick restarting.

3. Restarting using the CG recurrence. As mentioned above, dynamic thick restarting outperforms one-sided TR because it keeps more information at restart, but it increases the computational expense of every step. The latter can be seen by a simple operation count of the Davidson algorithm. According to notation of the algorithm, the number of floating point operations between two successive restarts is:

$$(m \Leftrightarrow p + 1)\ \mathcal{O}(\text{Matvec+Prec}) + \sum_{j=p}^{m} \mathcal{O}(j^3 + 10jN).$$

Keeping p close to m, as dynamic thick restarting does, increases the average operations per matrix-vector multiplication. For a typical case of $m = 20$, the dynamic scheme is about twice as expensive as the $p = 10$ case. The question is whether the same information can be saved more compactly, and thus reduce these expenses.

For the symmetric case, the three term recurrence of the Lanczos method is an indication that there might be a way to avoid orthogonalization of large bases. But even though orthogonalization is avoided in the Lanczos algorithm, the vectors still need to be stored for the computation of the Ritz vector at the end of the iteration. The Conjugate Gradient (CG) method is equivalent to the Lanczos method for solving linear systems of equations. The CG three term recurrence also achieves an implicit orthogonalization of the basis vectors, but it yields the solution by storing only three vectors and without repeating any of the information seen. The reason for storing all vectors only in the Lanczos method is the non-linearity of the eigenvalue problem, i.e., both the eigenvalue and eigenvector are unknown. However, if the exact eigenvalue is known, the Ritz vectors can be obtained by the CG process [25, 8].

In [25], Van Lenthe et al. also proposed to use the CG method for finding an eigenvector, even when the eigenvalue is not known. This is often justified in computational quantum

ETNA
Kent State University
etna@mcs.kent.edu

chemistry problems, because of the availability of good eigenvalue estimates and the relatively fast convergence of the methods.

Murray et al. [14], suggested a more interesting variant of this idea. They juxtaposed two methods; the Preconditioned CG (PCG) method using a preconditioner M to solve equation (2.1), with $r = (A \Leftrightarrow \mu I)x^{(0)}$ from the zero iteration of Davidson; and the Davidson method starting from $x^{(0)}$ and at every iteration, adding to the basis the preconditioned by M residual of the same targeted vector (e.g., always the lowest eigenpair). They observed that after j iterations, the PCG minimizes the $A\Leftrightarrow$norm of the error on a three-vector space, that is close to the space spanned by $\{x^{(j-1)}, x^{(j)}, r\}$, where $x^{(j-1)}, x^{(j)}$ are successive Ritz vectors from Davidson iterations $j \Leftrightarrow 1$ and j respectively, and r is the residual of $x^{(j)}$. They argued that, were the Davidson method to be restarted at the j-th iteration, it would be beneficial to keep the Ritz vector from the previous iteration ($x^{(j-1)}$) along with the current one. In fact, if the above three-vector spaces from CG and Davidson were identical, there would be no information loss by this restarted Davidson variant —which we call "CG-based"—, because for symmetric matrices the Rayleigh Ritz and Galerkin processes are equivalent. In general, the two spaces are not the same but they are close if the Ritz value does not vary significantly between steps. In the following section we present a theoretical justification of these arguments. Algorithmically, only two Davidson steps need to be modified to implement the CG-based restarting.

ALGORITHM 3.1. **Davidson changes for CG-based restarting**
8.1 $x_{prev} = x^{(j-1)}$*, save the Ritz vector from previous iteration*
8.2 $x^{(j)} = V_j c$ *and* $r = \mu x^{(j)} \Leftrightarrow Ax^{(j)}$
13.1 $V_{p-1} = \{x_i$*, the* $p \Leftrightarrow 1$ *lowest Ritz vectors* $\}$, $l \leq p \Leftrightarrow 1 < m \Leftrightarrow 1$,
13.2 $v_p = Ortho(V_{p-1}, x_{prev})$*, orthogonalize and add the previous Ritz vector.*

In Figure 3.1, we show a typical convergence history of the residual norm for three different restarting strategies. We are interested in the lowest eigenpair of the matrix BCSSTK05 from the Harwell-Boeing collection, and we report the number of matrix vector multiplications for both the diagonally preconditioned and non-preconditioned case. The original Davidson method restarts every 19 steps by keeping only the wanted Ritz pair, while the CG-based scheme restarts with two successive Ritz vectors as mentioned above. The dynamic strategy is applied as described in the previous section.

The convergence behavior is typical in a qualitative sense, since in most of the cases we have tested, dynamic thick restarting is slightly better than the CG-based scheme without preconditioning. Still, the results are impressive. This simple scheme works as well as dynamic thick restarting in the non-preconditioned case, and it is far better than the original Davidson method when preconditioning is used. Moreover, because the matrix size is small and it would be misleading to report times, the figure does not reflect that GD-based restarting is computationally much cheaper per step than dynamic thick restarting. There are also some disadvantages with this scheme that will become clear after the following section, and for which we show that thick restarting is a natural, complementary solution.

4. Theoretical justification. In this section we examine the reasons why the CG-based restarting works and we try to quantify how well a Ritz pair can be approximated through a three term recurrence.

To facilitate presentation clarity, we drop the superscript notation, and we use only a single subscript that denotes the iteration number for any variable, e.g., x_i is the Ritz vector at the i-th iteration. We also let the Davidson method target only the same eigenpair (not necessarily an extreme one) until convergence. We assume that the symmetric matrix A has no multiple, degenerate eigenvalues. We state the following lemma for any matrix, although

ETNA
Kent State University
etna@mcs.kent.edu

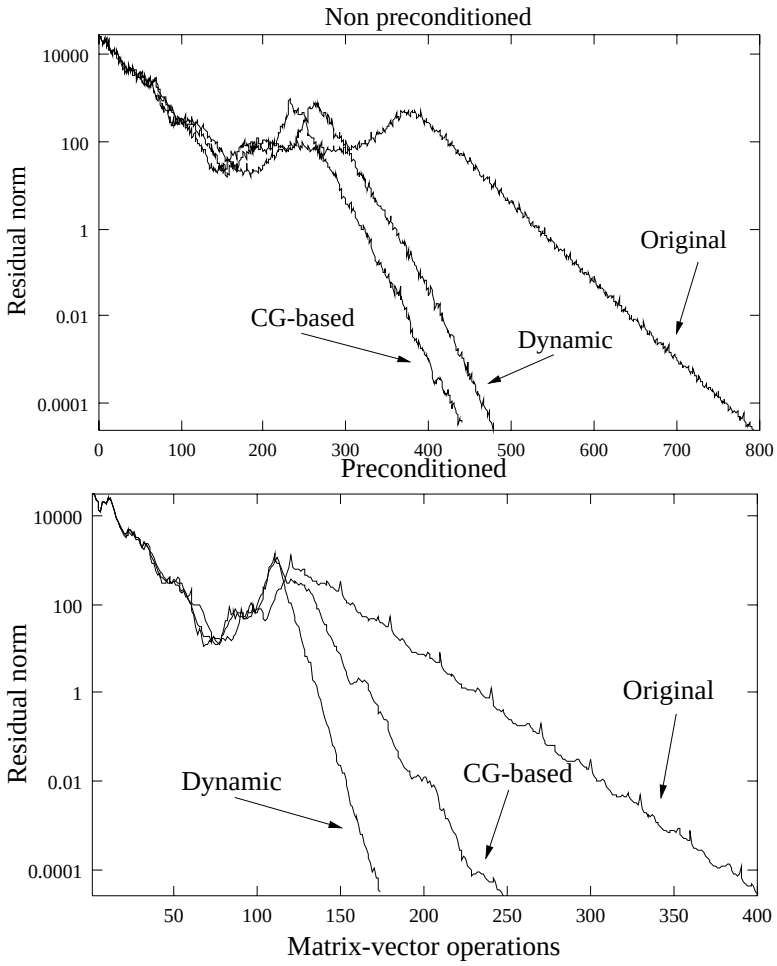

FIG. 3.1. *Residual convergence history for one eigenpair of* BCSSTK05*, for the three different restarting schemes.*

we subsequently use it only for symmetric matrices.

LEMMA 4.1. *Let x_0 a unit-norm vector of $\Re^N$, and $\mu_0 = x_0^T A x_0$. Let $\pi_0 = (I \Leftrightarrow x_0 x_0^T)$ the projector against x_0, and let also the residual of x_0: $r_0 = (A \Leftrightarrow \mu_0 I)x_0 = \pi_0 r_0$. Then for every $k > 1$:*

$$span(\{x_0, Ax_0, \ldots, A^k x_0\}) = span(\{x_0, r_0, (\pi_0 A \pi_0) r_0, \ldots, (\pi_0 A \pi_0)^{k-1} r_0\}).$$

Proof. Let $\mathcal{K}_k = span(\{x_0, Ax_0, \ldots, A^k x_0\})$ denote the Krylov space starting from x_0, and $\mathcal{L}_k = span(\{x_0, r_0, (\pi_0 A \pi_0) r_0, \ldots, (\pi_0 A \pi_0)^{k-1} r_0\})$. Obviously $\mathcal{K}_2 = \mathcal{L}_2$. We assume that $\mathcal{K}_i = \mathcal{L}_i$, for all $i < k$.

Let $x \in \mathcal{L}_k$. There is $\omega_1 \in \mathcal{L}_{k-1} = \mathcal{K}_{k-1}$ and coefficient d_1 such that:

$$x = \omega_1 + d_1 (\pi_0 A \pi_0)^{k-1} r_0 = \omega_1 + \pi_0 A \pi_0 \, \omega_2, \tag{4.1}$$

where $\omega_2 = d_1 (\pi_0 A \pi_0)^{k-2} r_0 \in \mathcal{L}_{k-1} = \mathcal{K}_{k-1}$. Since $\pi_0 = I \Leftrightarrow x_0 x_0^T$ and $A\omega_2 \in \mathcal{K}_k$, we

ETNA
Kent State University
etna@mcs.kent.edu

have:

$$x = \omega_1 + (I \Leftrightarrow x_0 x_0^T) A(\omega_2 \Leftrightarrow (x_0^T \omega_2) x_0) \tag{4.2}$$

$$= \omega_1 + A\omega_2 \Leftrightarrow (x_0^T \omega_2) A x_0 \Leftrightarrow (x_0^T A \omega_2) x_0 + (x_0^T A x_0)(x_0^T \omega_2) x_0 \tag{4.3}$$

$$\in \mathcal{K}_k. \tag{4.4}$$

Thus, $\mathcal{L}_k \subseteq \mathcal{K}_k$. If $\mathcal{L}_k$ is of full dimension, its dimension is $k+1$, the same as $\mathcal{K}_k$, and thus the two spaces must be equal. If $\mathcal{L}_k$ is not of full dimension, then it forms a smaller invariant subspace of dimension $i < k+1$, which is also included in $\mathcal{K}_k$. There are also $\mathcal{L}_{i-1} = \mathcal{L}_k$ and $\mathcal{K}_{i-1} = \mathcal{K}_k$, and from the inductive hypothesis $\mathcal{L}_k = \mathcal{K}_k$, which concludes the proof. □

The above lemma seems to imply that the Generalized Davidson and the Jacobi-Davidson methods (equations (2.1) and (2.2)) are equivalent. However, for both methods the space built by solving the respective correction equation does not include x_0, which is a crucial assumption in the previous lemma. In fact, the following theorem holds only for the Jacobi-Davidson correction equation. Hereafter, we consider A to be symmetric and we use the terms Arnoldi and Lanczos interchangeably.

THEOREM 4.2. *Let x_0 a unit-norm vector of $\Re^N$, $\sigma \in \Re$, and $\mu_0 = x_0^T A x_0$.*
Let (μ_k, x_k) be a Ritz pair obtained after k steps of the un-preconditioned Davidson (Arnoldi) method, with x_0 as a starting vector.
Let $y_k = x_0 + \delta$ be the approximate eigenvector, where δ is the correction obtained by applying k steps of the Conjugate Gradient method to the correction equation (2.2). Then:

$$y_k = x_k \Leftrightarrow \sigma = \mu_k.$$

Proof. Because there is no preconditioning, the space built by Davidson is the Krylov space: $span(\{x_0, Ax_0, \ldots, A^k x_0\})$. The CG method on equation (2.2), starting with zero initial guess builds the space: $span(\{\pi_0 r_0, (\pi_0 A \pi_0 r_0), \ldots, (\pi_0 A \pi_0)^{k-1} r_0\})$. Let V be an orthonormal basis of this CG space. From Lemma 4.1, $\{x_0, V\}$ is also an orthonormal basis of the Davidson space.

If we consider the above basis, and normalize the eigenvector so that the coefficient of x_0 is one, the Rayleigh-Ritz procedure applied at the k-th step of the Davidson method solves the following projected problem:

$$\begin{bmatrix} x_0^T A x_0 & x_0^T A V \\ V^T A x_0 & V^T A V \end{bmatrix} \begin{pmatrix} 1 \\ c_k \end{pmatrix} = \mu_k \begin{pmatrix} 1 \\ c_k \end{pmatrix}, \tag{4.5}$$

or equivalently, the following system, which has k solutions (μ_k, c_k), $c_k \in \Re^k$:

$$\mu_0 + x_0^T A V c_k = \mu_k \tag{4.6}$$

$$V^T A x_0 + V^T A V c_k = \mu_k c_k. \tag{4.7}$$

Since $x_0^T V = 0$, we can choose V as the Lanczos orthogonal vectors generated by x_0, and thus the projected matrix (4.5) becomes tridiagonal. For such matrices it is well known that μ_k cannot be also an eigenvalue of $V^T A V$ (unless $V^T A x_0 = 0$), and therefore we can solve equation (4.7) for c_k:

$$c_k = \Leftrightarrow (V^T A V \Leftrightarrow \mu_k I)^{-1} V^T A x_0. \tag{4.8}$$

Considering the basis V for the Galerkin condition of the CG method, and since $\pi_0 V = V$, and $V^T \pi_0 = V^T$, the projected problem solved is:

$$V^T (A \Leftrightarrow \sigma I) V c_k' = V^T (\mu_0 I \Leftrightarrow A) x_0 \quad \Leftrightarrow$$

$$(V^T A V \Leftrightarrow \sigma I) c_k' = \Leftrightarrow V^T A x_0. \tag{4.9}$$

ETNA
Kent State University
etna@mcs.kent.edu

From equations (4.8) and (4.9), if $\sigma = \mu_k$ the two methods produce the same coefficient vectors and thus same approximate vectors. The converse is also an immediate consequence of equations (4.7) and (4.9). □

In the above theorem, $(A \Leftrightarrow \sigma I)$ does not have to be positive definite. The only requirement is that the Krylov space can be built (no CG breakdown) so that the Galerkin projection can be applied. If in addition, $(A \Leftrightarrow \sigma I)$ is positive definite (or semidefinite) then the energy norm of the error is minimized over V. In general, any Ritz vector can also be obtained from a CG iteration, or more specifically:

COROLLARY 4.3. *If the Ritz value at the k-th step of the Lanczos method is known a priori, there is a three-term recurrence that yields the corresponding Ritz vector. Only the last of the vectors involved in the three-term recurrence is a Ritz vector.*

A similar observation was made by Saad in his thesis at Grenoble [16]. Therein, he identifies a three term recurrence leading to a specific Ritz vector, if the corresponding Ritz value is known. Following our previous notation, the Lanczos process starts from vector x_0 and produces a sequence of tridiagonal matrices $T_i = \text{tridiag}(\beta_i, \alpha_i, \beta_{i+1})$, the Ritz vectors x_i, and their corresponding Ritz values μ_i, $i = 0, \ldots, k$. Initially, we have $T_0 = (\alpha_0)$, $\alpha_0 = \mu_0 = x_0^T A x_0$, and $\beta_1 = \|r_0\|$. Let also $p_i(x)$ denote the characteristic polynomials of the matrices T_i, and define the polynomials:

$$q_i(x) = \frac{p_i(x) \Leftrightarrow p_i(\mu_k)}{(x \Leftrightarrow \mu_k)}, \qquad i = 0, \ldots, k. \tag{4.10}$$

If we set $z_i = q_i(A)x_0$, with the convention $z_{-1} = 0$, then the following holds:

PROPOSITION 4.4.

1. *z_k is a constant multiple of x_k.*
2. *The vectors z_i, $i = 1, \ldots k$, satisfy the following three-term recurrence:*

$$z_i = V_i + (\mu_k \Leftrightarrow \alpha_i)\, z_{i-1} \Leftrightarrow \beta_i^2\, z_{i-2},$$

where V_i, α_i, β_i are the basis vector and matrix coefficients produced by the Lanczos process.

Proof. For the characteristic polynomial of T_k, we have $p_k(\mu_k) = 0$, and (4.10) becomes $q_k(x) = p_k(x)/(x \Leftrightarrow \mu_k)$. It is well known that $q_k(A)x_0 = z_k$ is proportional to x_k, which proves the first part.

For the second part, one can use the determinant definition of the characteristic polynomials of the tridiagonal matrices to verify that the polynomials satisfy the following recurrence:

$$p_i(x) = (x \Leftrightarrow \alpha_i)p_{i-1}(x) \Leftrightarrow \beta_i^2 p_{i-2}(x), \quad i = 2, \ldots, k.$$

Using (4.10) and the above recurrence, we obtain the following after some algebraic manipulations:

$$q_i(A)x_0 = p_{i-1}(A)x_0 + (\mu_k \Leftrightarrow \alpha_i)q_{i-1}(A)x_0 \Leftrightarrow \beta_i^2 q_{i-2}(A)x_0.$$

Since p_{i-1} is the residual polynomial, $p_{i-1}(A)x_0 = V_i$, and the proof follows. □

The following table depicts the effect of the three methods, Lanczos (un-preconditioned Davidson), CG, and the three-term recurrence of proposition 4.4, when started with the same vector x_0, ending in the same Ritz vector x_k at the $k\Leftrightarrow$th step. Intermediate vectors differ in general. The CG eigenvector approximations y_i, and the recurrence vectors z_i are defined in theorem 4.2 and proposition 4.4 respectively.

ETNA
Kent State University
etna@mcs.kent.edu

$$\begin{array}{ccccccccccc} & & x_1 & x_2 & x_3 & \dots & x_{k-2} & x_{k-1} & & & \\ & \nearrow & & & & & & & \searrow & & \\ x_0 & \rightarrow & y_1 & y_2 & y_3 & \dots & y_{k-2} & y_{k-1} & \rightarrow & x_k. \\ & \searrow & & & & & & & \nearrow & & \\ & & z_1 & z_2 & z_3 & \dots & z_{k-2} & z_{k-1} & & & \end{array}$$

From the above we see that when the Ritz value μ_k is known in advance, the negative effects of restarting the Davidson or the Lanczos methods can be canceled by a three term recurrence. However, μ_k is usually computed during the Lanczos (Davidson) procedure. An interesting question is whether the three-term CG recurrence still produced an accurate approximation to the Ritz vector, if an inexact eigenvalue shift is used in equation (2.2). The following lemma attempts to quantify the distance of the vectors x_k and y_k when $\sigma \neq \mu_k$.

LEMMA 4.5. *Let the assumptions and notations of Theorem 4.2 hold, and denote by s_i the eigenvalues of the small matrix $S = V^T AV$, where V is the basis of CG. Then*

$$\frac{\|y_k \Leftrightarrow x_k\|}{\|y_k\|} \leq \frac{|\mu_k \Leftrightarrow \sigma|}{\min_i |s_i \Leftrightarrow \mu_k|}.$$

Proof. From equations (4.7) and (4.9): $V^T(A \Leftrightarrow \sigma I)Vc'_k = \Leftrightarrow V^T r_0 = V^T(A \Leftrightarrow \mu_k I)Vc_k$, and since μ_k is not an eigenvalue of S, as mentioned in the theorem,

$$c_k = (S \Leftrightarrow \mu_k I)^{-1}(S \Leftrightarrow \sigma I)c'_k.$$

From the definitions of x_k, y_k and from the above equation we have:

$$\begin{aligned} \|y_k \Leftrightarrow x_k\| &= \|V(c'_k \Leftrightarrow c_k)\| = \|c'_k \Leftrightarrow c_k\| \\ &\leq \|I \Leftrightarrow (S \Leftrightarrow \mu_k I)^{-1}(S \Leftrightarrow \sigma I)\| \|c'_k\| \leq \frac{\|c'_k\|}{\min_i |s_i \Leftrightarrow \mu_k|} |\mu_k \Leftrightarrow \sigma|. \end{aligned}$$

Finally, notice that $x_0^T V = 0$, and thus $\|y_k\|^2 = \|x_0 + Vc'_k\|^2 = \|x_0\|^2 + \|Vc'_k\|^2 \geq \|Vc'_k\|^2 = \|c'_k\|^2$. A simple substitution completes the proof. □

The lemma implies that when the eigenvalue is almost constant, which usually is the case near convergence or when convergence is slow, the CG and the Davidson/Lanczos method compute almost the same Ritz vectors. The problem with the above lemma is the fact that the denominator $\min_i |s_i \Leftrightarrow \mu_k|$ could be close to zero or even zero for interior eigenvalues. However, for the extreme eigenvalues we can guarantee that the denominator is bounded away from zero, when the Rayleigh quotient of the starting vector x_0 is sufficiently converged. We state first a lemma that relates the eigenvalue convergence of the lowest eigenvalue with the convergence of the corresponding eigenvector. Here (λ_i, u_i) denotes the i_{th} eigenpair of A.

LEMMA 4.6. *Let x_0 a unit-norm vector with Rayleigh quotient μ_0. Let also the gap between the first and the second eigenvalue of A be $\gamma = \lambda_2 \Leftrightarrow \lambda_1$. Then*

$$\sin^2 \theta(x_0, u_1) < \frac{|\mu_0 \Leftrightarrow \lambda_1|}{\gamma}.$$

Proof. See [15] Lemma (11-9-2). □

LEMMA 4.7. *Let the assumptions and notations of Lemma 4.5 hold, and let $\gamma = \lambda_2 \Leftrightarrow \lambda_1$. If the Rayleigh quotient of the initial vector x_0 is close enough to λ_1, i.e., $|\mu_0 \Leftrightarrow \lambda_1| < \epsilon$ for $\epsilon < \gamma/2$, then the denominator in Lemma 4.5 is bounded from below so that:*

$$\frac{\|y_k \Leftrightarrow x_k\|}{\|y_k\|} \leq \frac{|\mu_k \Leftrightarrow \sigma|}{\gamma \Leftrightarrow 2\epsilon}.$$

ETNA
Kent State University
etna@mcs.kent.edu

Proof. From the Cauchy interlace theorem we have that $\lambda_1 \leq \mu_k \leq s_1$, where s_1 is the lowest eigenvalue of S. Therefore, the denominator of Lemma 4.5 becomes $\min_i |s_i \Leftrightarrow \mu_k| = s_1 \Leftrightarrow \mu_k$. Since $\pi_0 V = V$, the eigenvalues of $S = V^T A V$ are the same as those of $V^T \pi_0 A \pi_0 V$. Again from the Cauchy interlace theorem, the smallest non zero eigenvalue $\lambda_1(\pi_0 A \pi_0) \leq \lambda_1(V^T \pi_0 A \pi_0 V) = s_1$, and thus

$$\min_i |s_i \Leftrightarrow \mu_k| \geq \lambda_1(\pi_0 A \pi_0) \Leftrightarrow \mu_k. \tag{4.11}$$

Let x be a vector of unit norm that satisfies $x^T \pi_0 A \pi_0 x = \lambda_1(\pi_0 A \pi_0)$. We can consider x so that $x^T x_0 = 0$, and thus $x^T \pi_0 A \pi_0 x = x^T A x$. Let $x = cu_1 + dy$ be the spectral decomposition of x onto the first eigenvector and the rest of the eigenspace, with $d^2 = 1 \Leftrightarrow c^2$. Let $x_0 = au_1 + bu$ be the spectral decomposition of x_0, with $a^2 = 1 \Leftrightarrow b^2$. From Lemma 4.6 we have $b^2 < \epsilon/\gamma$.

From $x_0^T x = 0$ it follows that $0 = ac + bd(y^T u)$, and thus $a^2 c^2 = b^2 d^2 (y^T u)^2 \leq b^2 d^2$. Using $a^2 = 1 \Leftrightarrow b^2$ and $d^2 = 1 \Leftrightarrow c^2$, we can obtain $c^2 \leq b^2$, and the following inequality:

$$1 \Leftrightarrow c^2 > 1 \Leftrightarrow b^2 > 1 \Leftrightarrow \epsilon/\gamma. \tag{4.12}$$

Representing the Rayleigh quotient of x from its spectral decomposition, we get:

$$x^T A x = c^2 \lambda_1 + (1 \Leftrightarrow c^2) y^T A y = \lambda_1 + (1 \Leftrightarrow c^2)(y^T A y \Leftrightarrow \lambda_1). \tag{4.13}$$

Since $\|y\| = 1$ and $y^T u_1 = 0$, $y^T A y \geq \lambda_2$. From this and inequality (4.12) we have:

$$\lambda_1(\pi_0 A \pi_0) = x^T A x > \lambda_1 + \gamma \Leftrightarrow \epsilon. \tag{4.14}$$

From (4.11) and (4.14), and since μ_k is a decreasing sequence bounded below by λ_1, it follows that:

$$\min_i |s_i \Leftrightarrow \mu_k| > \lambda_1 \Leftrightarrow \mu_k + \gamma \Leftrightarrow \epsilon > \gamma \Leftrightarrow 2\epsilon. \tag{4.15}$$

Note that if $\epsilon < \gamma/2$ the above is bounded away from zero, and for small ϵ the denominator is arbitrary close to γ. □

The lemma can also be extended to interior eigenvalues, provided that all smaller eigenvalues have converged, and that the current basis vectors are orthogonalized against the converged eigenvectors. In this case, we can safely say that the CG on equation (2.2) and the Davidson/Lanczos method compute almost the same Ritz vectors. In a sense, restarting Davidson with the Ritz vector from the previous step attempts to imitate the CG three-term recurrence.

For example, assume that we need to compute the (μ_{i+1}, x_{i+1}) Ritz pair, and that the un-preconditioned Davidson method is restarted in the traditional way after $i \Leftrightarrow 1$ steps, retaining only the current Ritz pair (μ_{i-1}, x_{i-1}). One iteration after restarting, the Davidson method generates (μ_i, x_i) as the next Ritz pair, but after a second iteration the new Ritz pair differs from (μ_{i+1}, x_{i+1}). At this point, the effects of restarting start to show. Consider a hypothetical CG recurrence that uses the unknown μ_{i+1} to produce the wanted Ritz pair in $i+1$ steps. If we apply lemma 4.7 to the vectors y_i of this CG process, but consider instead x_{i-1} and x_i as the endpoints, we get two inequalities:

$$\|y_{i-1} \Leftrightarrow x_{i-1}\|/\|y_{i-1}\| \leq \frac{|\mu_{i+1} \Leftrightarrow \mu_{i-1}|}{\gamma \Leftrightarrow 2\epsilon},$$

$$\|y_i \Leftrightarrow x_i\|/\|y_i\| \leq \frac{|\mu_{i+1} \Leftrightarrow \mu_i|}{\gamma \Leftrightarrow 2\epsilon}.$$

ETNA
Kent State University
etna@mcs.kent.edu

When the Ritz value is almost constant between steps, the Ritz vectors x_{i-1} and x_i approximate the CG iterates for the still uncomputed $i+1$ step. Because x_{i+1} is a linear combination of y_i, y_{i-1} (which cannot be computed), a good restarting basis for the Davidson is one consisting of both Ritz vectors $\{x_{i-1}, x_i\}$. In the case where the matrix is positive definite or semidefinite, this scheme can be thought of as an attempt to maintain the same space where the energy norm of the eigenvector error is minimized. Because of lemma 4.1 and theorem 4.2, the proposed restarting scheme works similarly on the Jacobi-Davidson method (i.e., the Davidson method with the CG method applied to solve equation (2.2)).

Although we have not given approximation bounds for the vectors z_i in proposition 4.4, one could conceive a restarting scheme where in addition to the Davidson basis vectors, the z_i are also generated to be used in restarting instead of the previous Ritz vector. The problem with this approach is that the z_i do not come for free, as the Ritz vector from the previous Davidson iteration does. In addition, the z_i are not normalized and they depend on a three-term recurrence, which makes them prone to instabilities. In our preliminary experiments, the generation of the z_i vectors did not help the convergence of the restarted Davidson method.

There are two problems with extending the above theory to the case of preconditioners that approximate A directly, such as incomplete factorization or approximate inverse. First, even when the preconditioner stays constant between iterations, the space built is not a Krylov subspace because the Ritz value in the residual $(A \Leftrightarrow \mu I)x$ of the correction equation changes. Therefore, this space cannot be replicated by an appropriate CG or PCG method. Second, the Jacobi-Davidson equation involves projections in both sides of the matrix, and it is not obvious how to apply the preconditioned CG method [7].

If we assume that the Ritz value in the residual is kept constant, equal to σ, and that the preconditioner is also constant $(M \Leftrightarrow \sigma I)^{-1}$, at every Davidson iteration the operator is $M_C = (M \Leftrightarrow \sigma I)^{-1}(A \Leftrightarrow \sigma I)$, and a similar argument as in theorem 4.2 can be made (consider Arnoldi on M_C, and PCG on $(A \Leftrightarrow \sigma I)$ with $(M \Leftrightarrow \sigma I)$ as preconditioner). The algorithm described by these assumptions is the restarted preconditioned Arnoldi method by Meerbergen [11], which uses the Arnoldi method to compute a Krylov space of M_C. Then, it uses this space to perform a Rayleigh-Ritz projection with A, and obtain a new σ. In the Davidson method, the latest Ritz estimates are computed at every step and thus it is impractical not to update the shift. However, in view of lemma 4.5, when the eigenvalue is almost constant between steps, the above assumptions describe exactly the Davidson procedure. In that case, the operator can be considered constant and the previous theory applies.

5. Overcoming weaknesses. The results of the previous section provide an explanation for the experiments in section 3 and the original observations in [14], i.e., CG-based restarting is justified when looking for an extreme eigenpair. However, the error bound in lemma 4.5 may be large for interior eigenvalues, even when the distance of the shift from the Ritz value is small. Murray et al. propose to use the Ritz vectors from the previous iteration for all required eigenvectors. Although in practice this improves convergence, it is not theoretically justified except for easy problems, and in our experience we never found it to be competitive with either thick or dynamic thick restarting. To guarantee the benefits of this restarting strategy for a larger number of eigenpairs, we must target one eigenpair at a time, deflating (orthogonalizing against) each one that converges. The question that arises is whether the benefits for the extreme eigenpairs carry over to the interior ones.

Figures 5.1 and 5.2 show the convergence history for each of the four lowest eigenpairs of the BCSSTK05 matrix, without and with diagonal preconditioning, respectively. In Figure 5.1 we see that dynamic thick restarting improves the convergence of all four eigenpairs uniformly. For CG-based restarting, the fast convergence to the first eigenpair postpones convergence to the rest, even though the same restarting scheme is used for all eigenpairs after

ETNA
Kent State University
etna@mcs.kent.edu

lower ones have converged. The fourth eigenpair is found in as many steps as with the original Davidson method, and this trend continues as we move within the spectrum. It is worth mentioning that the convergence rate (slope of the curve) is similar for all eigenpairs and it is close to the rate from dynamic thick restarting. The problem is that these rates are not assumed until an eigenvalue is targeted.

The reasons for this become clear when Figure 5.2 is also taken into account. The fact that the convergence behavior of un-preconditioned Davidson method using CG-based restarting resembles the behavior of the Davidson method with some preconditioner is no surprise. The motivation for this restarting scheme stems from the closeness of the iterates of the Davidson method with the ones from solving the Jacobi-Davidson correction equation (2.2) with CG. As with preconditioning, the convergence is focused on the targeted eigenvector by retaining the information needed to produce the following Ritz vector. At the same time, little or no useful information is kept towards other eigenpairs, causing significant delays to their convergence and loss of superlinear convergence to the targeted one. When preconditioning is used for the dynamic thick restarting, the convergence behavior is similar to the CG-based one. As theorem 2.1 points out, the lower part of the spectrum is better approximated than the higher end, and dynamic thick restarting performs well mainly because it retains a larger number of vectors.

5.1. Combining with thick restarting. Evidently the CG-based scheme is powerful for one extreme eigenpair, but it needs improvement for interior ones. Since its main "drawback" is that it behaves like preconditioning, we can use it together with one-sided thick restarting. While the CG-based restarting would provide fast convergence towards the targeted vector, thick restarting would maintain and build good estimates for the close eigenvectors (see also the discussion in section 2.3). Thus, we propose to restart the Davidson algorithm with the following basis:

$$\{x_1^{(j)}, \ldots, x_p^{(j)}, x^{(j-1)}\},$$

where $x_i^{(j)}$ is the $i\Leftrightarrow$th extreme Ritz vector at step j, and $x^{(j-1)}$ is the Ritz vector from the previous step for the targeted eigenvector. Typically, the thickness of restarting is set to about half the maximum number of basis vectors. It is also possible to monitor eigenvalue convergence, and when interior eigenvalues are near convergence, to retain also their corresponding Ritz vectors from the previous step. In general, if the k lowest eigenvalues have approached convergence, the Davidson method could be restarted with:

$$\{x_1^{(j)}, \ldots, x_p^{(j)}, x_1^{(j-1)}, \ldots, x_k^{(j-1)}\}.$$

We should mention that the above restarting scheme cannot be implemented with implicit restarting because there is no single filtering polynomial that can produce both $x^{(j-1)}$ and $x^{(j)}$.

Figures 5.3 and 5.4 demonstrate how the new combination scheme overcomes the aforementioned problems, and converges faster than dynamic thick restarting especially with preconditioning. For the un-preconditioned case, the lowest two eigenpairs are found faster with the combination scheme than with the dynamic one, but the interior eigenvectors in the dynamic case benefit more from the uniform convergence obtained by the Lanczos gradient (residual of Ritz vectors).

We have performed a series of experiments on matrices from the Harwell-Boeing collection, of which a summary appears in Table 5.1. For all matrices we look for the lowest five eigenpairs, we use diagonal and approximate inverse preconditioning, and we allow a maximum of 5000 matrix-vector multiplications. For all cases our new scheme is consistently

ETNA
Kent State University
etna@mcs.kent.edu

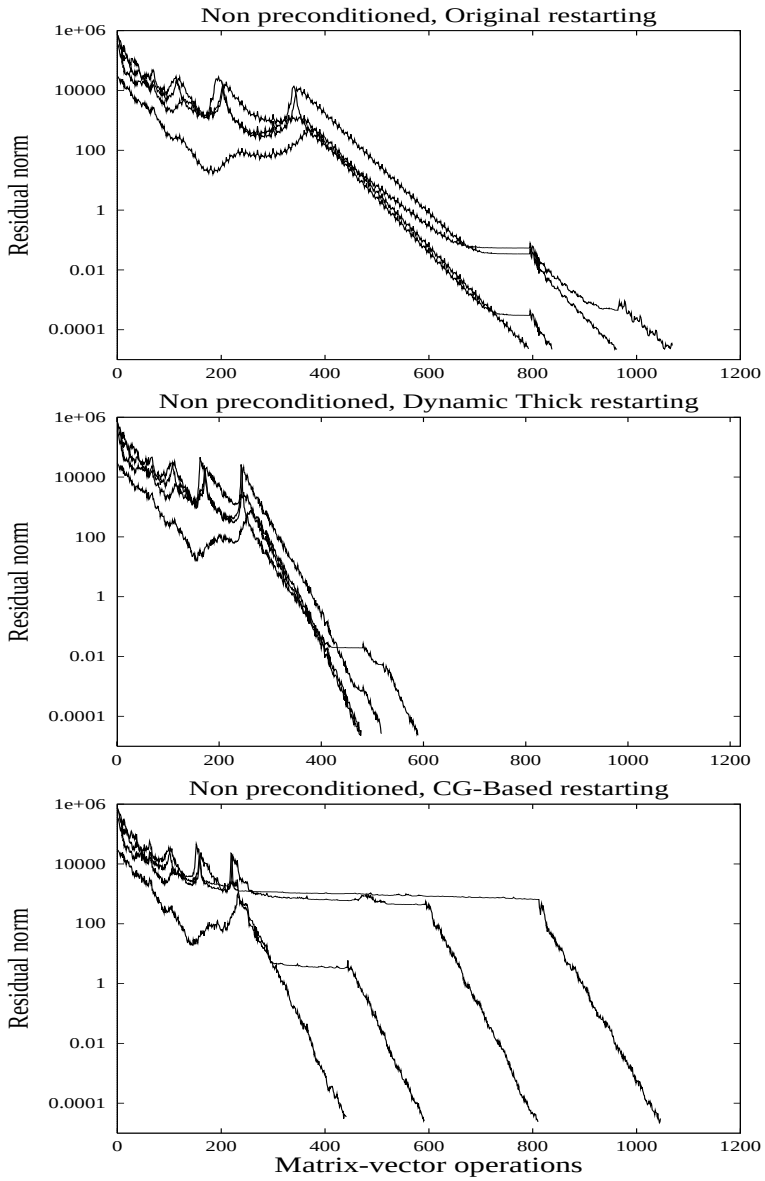

FIG. 5.1. *Residual convergence curves for four eigenpairs of* BCSSTK05*, for the three different restarting schemes without preconditioning.*

better than dynamic thick restarting. Again, since many matrices are of small size we do not report timings. However, timings would be in strong favor of the new scheme, since on average it requires fewer vectors present in the basis than dynamic thick restarting.

5.2. An efficient implementation. The addition of the previous Ritz vector $x^{(j-1)}$ in the basis is not without expense, because it is not orthogonal to the rest of the Ritz vectors $x_i^{(j)}$. In fact $x_1^{(j-1)}$ (assuming we target the lowest eigenpair) is very close to $x_1^{(j)}$ and orthogonalization must be performed in a stable way. Moreover, the scheme seems to require additional storage for keeping the previous Ritz vectors. The Davidson algorithm also needs

ETNA
Kent State University
etna@mcs.kent.edu

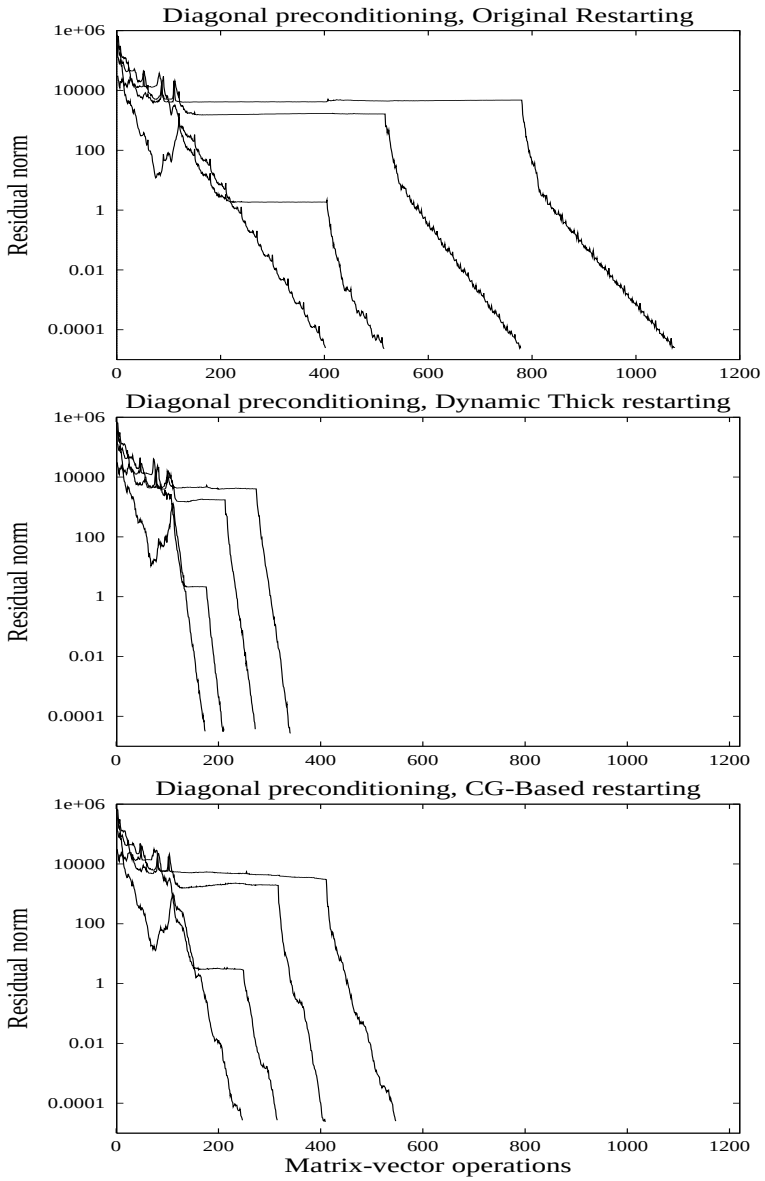

FIG. 5.2. *Residual convergence curves for four eigenpairs of* BCSSTK05*, for the three different restarting schemes with diagonal preconditioning.*

the vectors $w_j = Av_j$, and thus storage for $w^{(j-1)} = Ax_1^{(j-1)}$ should be accommodated for. It is possible that better convergence could be attained if these vectors were used to build a larger basis. Alternatively, we could perform an additional matrix vector multiplication to recompute $w^{(j-1)}$ after orthogonalization, but this could prove expensive.

There is a simple, elegant, and inexpensive way to face the above problems, that is not mentioned in [14]. Instead of keeping large dimension vectors and orthogonalizing these vectors, we could perform all operations and maintain the extra storage in the coefficient space. In the notation of the Davidson algorithm in section 2, the Ritz vector $x^{(j)}$ can be obtained as a linear combination of V with coefficient vector c, $x = Vc$. Thus, instead

ETNA
Kent State University
etna@mcs.kent.edu

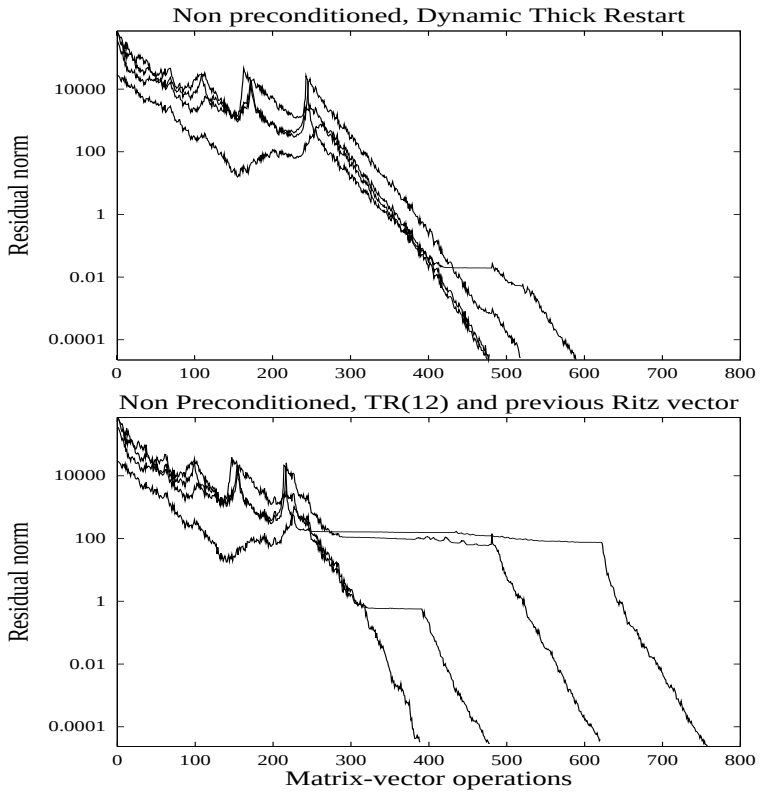

FIG. 5.3. *Residual convergence curves for four eigenpairs of* BCSSTK05*, for the dynamic and the combination restarting schemes without preconditioning.*

TABLE 5.1

Comparison of dynamic thick restarting (Dyn) with the combination of thick and CG-based restarting (TR(p)+1) on symmetric Harwell-Boeing matrices, with diagonal and approximate inverse preconditioners. One Ritz vector from the previous iteration is retained. The number of matrix vector multiplications is reported, with a maximum of 5000. Five smallest eigenvalues are sought. Davidson uses a basis size of 20.

	Diagonal		Approximate Inverse	
	Dyn	TR(10)+1	Dyn	TR(12)+1
BCSSTK01	124	117	108	91
BCSSTK02	190	172	92	87
NOS4	244	184	91	86
BCSSTK22	721	676	300	277
LUND B	349	298	338	327
BCSSTK05	409	370	251	221
NOS1	-	3071	2923	2162
NOS5	819	614	354	336
494 BUS	3456	2098	411	384
662 BUS	902	680	291	269
685 BUS	763	639	267	245
NOS3	664	519	258	242
BCSSTK10	2808	2245	1076	1056
1138 BUS	-	4030	735	655
BCSSTM27	1768	1710	509	502
BCSSTM13	269	263	177	179
BCSSTK21	1141	1062	601	612
BCSSTK16	663	587	317	302

ETNA
Kent State University
etna@mcs.kent.edu

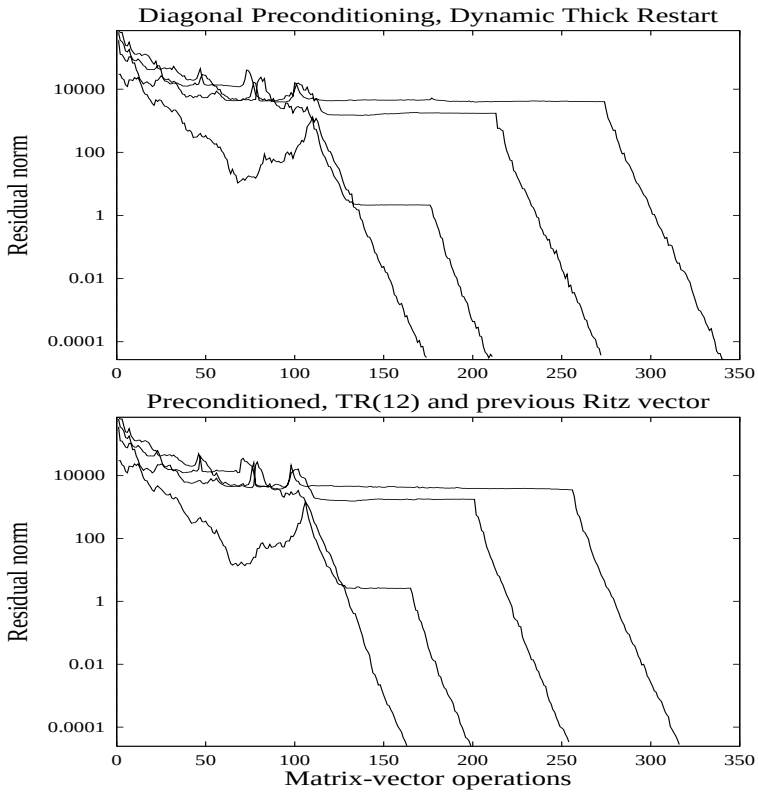

FIG. 5.4. *Residual convergence curves for four eigenpairs of* BCSSTK05*, for the dynamic and the combination restarting schemes with diagonal preconditioning.*

of storing the full-dimension $x^{(j-1)}$ vector, we can store only its coefficients from the last step $c^{(j-1)}$. Since in the following step the basis has one additional vector, we set $c' = [c^{(j-1)T}, 0]^T$, and $x^{(j-1)}$ is computed as:

$$x^{(j-1)} = Vc' = V \begin{bmatrix} c_1^{(j-1)} \\ \vdots \\ c_{m-1}^{(j-1)} \\ 0 \end{bmatrix}.$$

To orthogonalize $x^{(j-1)}$ with some Ritz vector x_i, or $(I \Leftrightarrow x_i x_i^T)x^{(j-1)}$, we notice that $x_i = Vc_i$, $x^{(j-1)} = Vc'$ and that $(I \Leftrightarrow x_i x_i^T)x^{(j-1)} = V(I \Leftrightarrow c_i c_i^T)c'$. Therefore, all orthogonalization can also be performed in the coefficient space, just before we compute the new restarting basis. Similarly, storing $w^{(s-1)}$ vectors is unnecessary, since $w^{(s-1)} = Ax^{(j-1)} = AVc' = Wc'$, and both W and c' are already stored. Finally, it is easy to show that the new projection matrix $V^T AV$ after restarting has the following convenient format:

$$H = \begin{bmatrix} \lambda_1 & & & \\ & \ddots & & \\ & & \lambda_L & \\ & & & c''^T W c'' \end{bmatrix},$$

ETNA
Kent State University
etna@mcs.kent.edu

where c'' is obtained by c' orthogonality against all coefficient vectors of maintained Ritz vectors. Therefore, applying the combination of thick and CG-based restarting schemes can be performed efficiently and without any additional storage. The only expense incurred is similar to thick restarting with a couple of additional Ritz vectors.

6. Conclusions and suggestions. Restarting eigenvalue iterative methods is often a cause of significant convergence deterioration. In this paper, we consider the effects of preconditioning on the dynamic thick restarting strategy, which we proposed previously, and we analyze both theoretically and experimentally an earlier idea based on the Conjugate Gradient method for restarting with two vectors per required eigenvector.

As a conclusion, dynamic thick restarting is both efficient and robust, but its effectiveness is reduced with accurate preconditioning. The CG-based scheme works well and it is also theoretically justified for extreme eigenpairs. For interior ones, as well as for improving convergence to the extreme ones, it should be combined with one-sided thick restarting. An efficient implementation is also possible that makes the combination scheme computationally attractive.

When implementing eigenvalue iterative methods, the kind of problem solved should dictate the choice of restarting strategy. From our experience, dynamic thick restarting performs best, if there is no preconditioning in the Davidson method, or if many eigenpairs are needed. This technique demonstrates similar behavior when interfaced to the ARPACK software [10, 24]. For only one eigenpair or if preconditioning is used, the combination of thick restarting and the CG-based approach is the method of choice. Because of the requirement of explicit restarting, this technique can only be implemented with (Jacobi-)Davidson codes. Finally, the latter scheme is even more attractive when the matrix vector multiplication is inexpensive.

Acknowledgments. The authors would like to thank the referees whose comments helped to improve significantly the clarity and presentation of the paper.

REFERENCES

[1] D. CALVETTI, L. REICHEL, AND D. SORENSEN, *An implicitly restarted Lanczos method for large symmetric eigenvalue problems*, Electronic Trans. Numer. Anal., 2 (1994), pp. 1–21.
[2] E. CHOW AND Y. SAAD, *Approximate inverse preconditioners via sparse-sparse iteration*, SIAM J. Sci. Comput., to appear.
[3] M. CROUZEIX, B. PHILIPPE, AND M. SADKANE, *The Davidson method*, SIAM J. Sci. Comput., 15 (1994), pp. 62–76.
[4] J. CULLUM AND R. A. WILLOUGHBY, *Lanczos algorithms for large symmetric eigenvalue computations*, vol. 2: Programs of Progress in Scientific Computing; v. 4, Birkhauser, Boston, 1985.
[5] 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), pp. 87–94.
[6] I. S. DUFF, R. G. GRIMES, AND J. G. LEWIS, *Sparse matrix test problems*, ACM Trans. Math. Soft., (1989), pp. 1–14.
[7] D. R. FOKKEMA, G. L. G. SLEIJPEN, AND H. A. VAN DER VORST, *Jacobi-Davidson style QR and QZ algorithms for the partial reduction of matrix pencils*, Tech. Rep. 941, Department of Mathematics, University of Utrecht, 1996. To appear in SIAM J. Sci. Comput.
[8] A. V. KNYAZEV, *A preconditioned Conjugate Gradient method for eigenvalue problems and its implementation in a subspace*, International Series of Numerical Mathematics, 96 (1991), pp. 143–154.
[9] R. B. LEHOUCQ, *Analysis and implementation of an implicitly restarted Arnoldi iteration*, PhD thesis, Department of Computational and Applied Mathematics, Rice University, 1995. TR95-13.
[10] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, SIAM, Phildelphia, PA, 1998.
[11] K. MEERBERGEN, *Robust methods for the calculation of rightmost eigenvalues of nonsymmetric problems*, PhD thesis, Department of Computer Science, K. U. Leuven, Heverlee, Belgium, 1996.
[12] R. B. MORGAN, *On restarting the Arnoldi method for large nonsymmetric eigenvalue problems*, Math. Comput., 65 (1996), pp. 1213–1230.

ETNA
Kent State University
etna@mcs.kent.edu

[13] R. B. MORGAN AND D. S. SCOTT, *Generalizations of Davidson's method for computing eigenvalues of sparse symmetric matrices*, SIAM J. Sci. Comput., 7 (1986), pp. 817–825.
[14] C. W. MURRAY, S. C. RACINE, AND E. R. DAVIDSON, *Improved algorithms for the lowest eigenvalues and associated eigenvectors of large matrices*, J. Comput. Phys., 103 (1992), pp. 382–389.
[15] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, SIAM, Philadelphia, PA, 1998.
[16] Y. SAAD, *Computation of eigenvalues of large Hermitian matrices by partitioning techniques*, tech. rep., INPG- University of Grenoble, Grenoble, France, 1974. Dissertation (French).
[17] Y. SAAD, *Chebyshev acceleration techniques for solving nonsymmetric eigenvalue problems*, Math. Comp., 42 (1984), pp. 567–588.
[18] ———, *Numerical methods for large eigenvalue problems*, Manchester University Press, 1992.
[19] ———, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
[20] G. L. G. SLEIJPEN AND H. A. V. DER VORST, *A Jacobi-Davidson iteration method for linear eigenvalue problems*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 401–425.
[21] D. C. SORENSEN, *Implicit application of polynomial filters in a K-step Arnoldi method*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 357–385.
[22] A. STATHOPOULOS AND C. FISCHER, *A Davidson program for finding a few selected extreme eigenpairs of a large, sparse, real, symmetric matrix*, Comput. Phys. Commun., 79 (1994), pp. 268–290.
[23] A. STATHOPOULOS, Y. SAAD, AND C. FISCHER, *Robust preconditioning of large, sparse, symmetric eigenvalue problems*, J. Comput. Appl. Math., 64 (1995), pp. 197–215.
[24] A. STATHOPOULOS, K. WU, AND Y. SAAD, *Dynamic thick restarting of the Davidson, and the implicitly restarted Arnoldi methods*, SIAM J. Sci. Comput., 19 (1998), pp. 227–245.
[25] J. VAN LENTHE AND P. PULAY, *A space-saving modification of Davidson's eigenvector algorithm*, J. Comput. Chem., 11 (1990), pp. 1164–1168.