Scalable Library for Eigenvalue Problem Computations SLEPc

SLEPc Technical Report **STR-10**
Available at http://slepc.upv.es

Davidson Type Subspace Expansions for the Linear Eigenvalue Problem

M. E. Hochstenbach
E. Romero
J. E. Roman

Last update: November, 2012 (SLEPc 3.3)
Previous updates: –

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

Contents

1 Introduction

We consider the generalized eigenvalue problem (GEP)

$$A\mathbf{x} = \lambda B\mathbf{x}, \tag{1}$$

where A and B are $n \times n$ matrices, and are interested in interior eigenvalues close to a given target $\tau \in \mathbb{C}$. Efficient computation of these eigenvalues is a hard task that generally requires both a suitable subspace extraction process (often harmonic Rayleigh–Ritz is used, see, e.g., [Stewart, 2001]) and a quality subspace expansion method. This subspace expansion in turn generally requires a (good) preconditioner and/or many steps of an iterative linear solver, depending on the complexity of the problem at hand. In this paper we assume that we have a preconditioner M, which, for instance, may be an inexact LU-decomposition of $A - \tau B$.

Iterative methods based on Krylov subspaces (for instance, Lanczos for Hermitian problems, and Arnoldi and Krylov-Schur for non-Hermitian problems) are widely used to compute the eigenvalues in the extremes of the spectrum of standard eigenvalue problems. However, Davidson methods may present better performance computing interior eigenvalues and/or in generalized eigenproblems when exact solves with $A - \tau B$ are unaffordable [Davidson, 1975; van Lenthe and Pulay, 1990; Crouzeix *et al.*, 1994; Heuveline *et al.*, 1997; Arbenz *et al.*, 2006; Genseberger, 2010].

Starting with the introduction of the Davidson method [Davidson, 1975], there have been a wide variety of developments in the subspace expansion. Generalized Davidson (GD) [Morgan and Scott, 1986] introduces the first expansion that uses an arbitrary preconditioner, which is applied to the residual to try to enrich the approximation in the direction of the desired eigenvector. However, it may very well occur that the resulting vector is almost collinear to the approximated eigenvector, leading to the stagnation of the method. The Olsen variant [Olsen

et al., 1990] attempts to avoid this situation by working on the orthogonal complement subspace to the approximated eigenvector.

Jacobi–Davidson (JD) [Sleijpen and van der Vorst, 1996, 2000] also seeks to avoid the stagnation, but differs from the previous methods in the fact that the convergence of JD may depend less on the quality of the preconditioner. The JD expansion results from the approximate solution of a linear system called the correction equation; the quality or the efficiency of the computation may be enhanced by a preconditioner.

GD and Olsen are attractive because of their straightforward implementation and good performance for easier problems. A major challenge in JD is the adaptive determination of parameters such as the number of inner steps; see [Stathopoulos, 2007; Hochstenbach and Notay, 2009] for recent progress in this direction. A rule-of-thumb is that JD may be necessary for harder problems.

We now first review some existing methods to expand the search space before we propose a new method in Section 2. For interior eigenvalues subspace expansion methods include (inexact) Rayleigh quotient iteration (RQI), (inexact) inverse iteration, or Jacobi–Davidson [Sleijpen *et al.*, 1996]. Let $(\theta, \mathbf{u}) \approx (\lambda, \mathbf{x})$ be an approximate eigenpair, where $\mathbf{u}$ is in the search space $\mathcal{U}$. In the Jacobi–Davidson method, a possible correction equation is

$$\left(I - \frac{B\mathbf{u}\mathbf{u}^*}{\mathbf{u}^*B\mathbf{u}}\right)(A - \theta B)\,\mathbf{t} = -(A\mathbf{u} - \theta B\mathbf{u}), \qquad \mathbf{t} \perp \mathbf{u}, \tag{2}$$

where this $\mathbf{t}$ is used to expand the search space. With the (standard) projected preconditioning we solve $\mathbf{t} \perp B\mathbf{u}$ from

$$\left(I - \frac{M^{-1}B\mathbf{u}\mathbf{u}^*}{u^*M^{-1}B\mathbf{u}}\right)M^{-1}(A - \theta B)\mathbf{t} = -\left(I - \frac{M^{-1}B\mathbf{u}\mathbf{u}^*}{\mathbf{u}^*M^{-1}B\mathbf{u}}\right)M^{-1}(A\mathbf{u} - \theta B\mathbf{u}).$$

First we note that we may approximate the solution $\mathbf{t}$ by just taking the right-hand side

$$\mathbf{t}_{\text{olsen}} = -\left(I - \frac{M^{-1}B\mathbf{u}\mathbf{u}^*}{\mathbf{u}^*M^{-1}B\mathbf{u}}\right)M^{-1}(A\mathbf{u} - \theta B\mathbf{u}).$$

This is a linear combination of $M^{-1}(A\mathbf{u} - \theta B\mathbf{u})$ and $M^{-1}B\mathbf{u}$, orthogonal to $\mathbf{u}$. The subscript reflects that fact that it is a generalization for the GEP of the approach advocated by Olsen *et al.* [1990], who proposed the expansion

$$-\left(I - \frac{M^{-1}\mathbf{u}\mathbf{u}^*}{\mathbf{u}^*M^{-1}\mathbf{u}}\right)M^{-1}(A\mathbf{u} - \theta\mathbf{u})$$

for the standard eigenvalue problem $A\mathbf{x} = \lambda\mathbf{x}$.

We can also precondition (2) by a regular (unprojected) preconditioner (cf. [Hochstenbach and Notay, 2009])

$$M^{-1}\left(I - \frac{B\mathbf{u}\mathbf{u}^*}{\mathbf{u}^*B\mathbf{u}}\right)(A - \theta B)\,\mathbf{t} = M^{-1}(A\mathbf{u} - \theta B\mathbf{u}), \qquad \mathbf{t} \perp B\mathbf{u}.$$

Again, we may approximate the solution $\mathbf{t}$ of (2) by just taking the right-hand side

$$\mathbf{t}_{\mathrm{GD}} = M^{-1}(A\mathbf{u} - \theta B\mathbf{u});$$

this approach is called generalized Davidson (GD) or also preconditioned inverse iteration (PINVIT) in the literature (see, e.g., [Neymeyr, 2001]).

We now make some comparisons between $\mathbf{t}_{\mathrm{olsen}}$ and $\mathbf{t}_{\mathrm{GD}}$. Suppose that M is a preconditioner of good quality, for the moment we will assume that $M^{-1} = (A - \tau B)^{-1}$ is an exact inverse. Then, for this special case,

$$\mathbf{t}_{\mathrm{olsen}} = -\mathbf{u} + (\mathbf{u}^*(A - \tau B)^{-1}B\mathbf{u})^{-1}\,(A - \tau B)^{-1}B\mathbf{u}$$

and

$$\mathbf{t}_{\mathrm{GD}} = \mathbf{u} + (\tau - \theta)(A - \tau B)^{-1}B\mathbf{u}.$$

If θ is very close to τ, which for instance may be possible if the target is quite accurate, it is clear that $\mathbf{t}_{\mathrm{GD}}$ may degenerate. Indeed, the case that $\theta = \tau$ suffers from the well-known "Davidson paradox": the perfect preconditioner gives no subspace expansion. The Olsen approach does not share this disadvantage; note that

$$\|(\mathbf{u}^*(A - \tau B)^{-1}B\mathbf{u})^{-1}\,(A - \tau B)^{-1}B\mathbf{u}\| \geq 1.$$

However, a clear disadvantage of $\mathbf{t}_{\mathrm{olsen}}$ is that, unless M^{-1} is an exact inverse of $A-\tau B$ as above, this approach spends two actions of the preconditioner ($M^{-1}A\mathbf{u}$ and $M^{-1}B\mathbf{u}$, or $M^{-1}B\mathbf{u}$ and $M^{-1}(A\mathbf{u} - \theta B\mathbf{u})$); while $\mathbf{t}_{\mathrm{GD}}$ spends only one. In the next section we propose a new approach that attempt to turn this fact into a strength.

2 A double subspace expansion approach

Based on the observations of the previous section, we conclude that $\mathbf{t}_{\mathrm{GD}}$ is a comparatively cheap approach and may be sensible in particular if the preconditioner is of good quality and $M^{-1}(A-\theta B)$ is not close to the identity. On the other hand, the expansion $\mathbf{t}_{\mathrm{olsen}}$ may be more robust in general, but is twice as expensive in terms of actions with the preconditioner.

We now propose a new subspace expansion that may combine the strengths of the two approaches: we will expand the search space by *both* $M^{-1}A\mathbf{u}$ and $M^{-1}B\mathbf{u}$. While these vectors will asymptotically be collinear, generally this will not be the case until very late in the process.

We note that this subspace expansion process has a number of potential favorable properties. First, the expansion includes both $\mathbf{t}_{\mathrm{olsen}}$ and $\mathbf{t}_{\mathrm{GD}}$. Second, by expanding the space by more than one vector per outer iteration, effectively one (harmonic) Rayleigh–Ritz extraction process is avoided, saving computational costs. Third, the new expansion relies on the robustness of the extraction process to select the best combination of $M^{-1}A\mathbf{u}$ and $M^{-1}B\mathbf{u}$.

3 Comparative analysis

We introduce a general scheme that considers both the original GD expansion and the double expansion, but without taking into account the subspace acceleration, as shown in Algorithm 1.

Algorithm 1 (Simplified Davidson for finding the eigenvalue closest to τ)
Input: initial eigenvector approximation $\mathbf{x}^{(0)}$.
Output: $\theta^{(k)}$ and $\mathbf{x}^{(k)}$ from the last iteration k.
Compute $\theta^{(0)} = \rho(\mathbf{x}^{(0)})$
For $i = 0, 1, 2, \ldots$

1. Compute the residual $\mathbf{r}^{(i)} = (A - \theta^{(i)} B)\,\mathbf{x}^{(i)}$.
2. Test for convergence.
3. Compute $\mathbf{t}^{(i)} \perp B^* B\mathbf{x}^{(i)}$ such that $\|(A - \tau B)\mathbf{t}^{(i)} - \mathbf{r}^{(i)}\| \leq \xi_1^{(i)} \|\mathbf{r}^{(i)}\|$.
4. Set $\mathbf{x}^{(i+1)} = \|\mathbf{x}^{(i)} - \mathbf{t}^{(i)}\|^{-1}(\mathbf{x}^{(i)} - \mathbf{t}^{(i)})$ and $\theta^{(i+1)} = \rho(\mathbf{x}^{(i+1)})$.

Without loss of generality, for the following discussion we compute the approximate eigenvalue associated to an approximate eigenvector as the related generalized Rayleigh quotient

$$\rho(\mathbf{x}) = \frac{\mathbf{x}^* B^* A \mathbf{x}}{\mathbf{x}^* B^* B \mathbf{x}}, \tag{3}$$

and we consider that the approximate eigenvectors $\mathbf{x}$ are normalized so that $\|B\mathbf{x}\| = 1$.

Assume we have the approximate eigenvector $\mathbf{x}^{(i)}$ and an expansion

$$M^{-1}\tilde{\mathbf{t}}^{(i)} = M^{-1}(A + \beta^{(i)} B)\mathbf{x}^{(i)},$$

computed by GD (with $\beta^{(i)} = -\theta^{(i)}$) or GD2. Consider $\mathbf{t}^{(i)}$ in Algorithm 1 as a projection of the expansion $M^{-1}\tilde{\mathbf{t}}^{(i)}$ where the $B^* B\mathbf{x}^{(i)}$ direction has been removed. We characterize the quality of the expansion by $\mathbf{d}^{(i)}$,

$$(A - \tau B)\mathbf{t}^{(i)} = \mathbf{r}^{(i)} + \mathbf{d}^{(i)}, \text{ for } \mathbf{t}^{(i)} \perp B^* B\mathbf{x}^{(i)}. \tag{4}$$

We can write Eq. (4) as

$$(A - \tau B)M^{-1}\tilde{\mathbf{t}}^{(i)} = \mathbf{r}^{(i)} + \alpha B\mathbf{x}^{(i)} + \mathbf{d}^{(i)},$$

where

$$\alpha = \frac{\mathbf{x}^{(i)^*} B^* B \left(M^{-1}\tilde{\mathbf{t}}^{(i)} - (A - \tau B)^{-1}(\mathbf{r}^{(i)} + \mathbf{d}^{(i)})\right)}{\mathbf{x}^{(i)^*} B^* B (A - \tau B)^{-1} B\mathbf{x}^{(i)}}.$$

This results in

$$Q(A - \tau B)M^{-1}\tilde{\mathbf{t}}^{(i)} = (A - \tau B)\tilde{Q}M^{-1}\tilde{\mathbf{t}}^{(i)} = Q\mathbf{r}^{(i)} + Q\mathbf{d}^{(i)}, \tag{5}$$

where

$$Q = I - \frac{B\mathbf{x}^{(i)}\mathbf{x}^{(i)^*}B^*B(A-\tau B)^{-1}}{\mathbf{x}^{(i)^*}B^*B(A-\tau B)^{-1}B\mathbf{x}^{(i)}} \quad \text{and} \quad \tilde{Q} = I - \frac{(A-\tau B)^{-1}B\mathbf{x}^{(i)}\mathbf{x}^{(i)^*}B^*B}{\mathbf{x}^{(i)^*}B^*B(A-\tau B)^{-1}B\mathbf{x}^{(i)}}.$$

Hence we take $\mathbf{t}^{(i)} := \tilde{Q}M^{-1}\tilde{\mathbf{t}}^{(i)}$, satisfying $\tilde{Q}M^{-1}\tilde{\mathbf{t}}^{(i)} \perp B^*B\mathbf{x}^{(i)}$.

Then we approximate the distance to the best expansion by $\delta^{(i)} := \|\mathbf{d}^{(i)}\|$ taking $\mathbf{t}^{(i)} = \tilde{Q}M^{-1}\tilde{\mathbf{t}}^{(i)}$. For GD, that is for $\tilde{\mathbf{t}}^{(i)} = \mathbf{r}^{(i)}$, this distance is

$$\delta_{GD}^{(i)} := \|\mathbf{d}_{GD}\| = \left\| \left(Q(A-\tau B)M^{-1} - I\right)\mathbf{r}^{(i)} \right\|.$$

Note that $\delta_{GD}^{(i)} = 0$ if $M = A - \tau B$. In the case of GD2, $\beta^{(i)}$ is a free parameter, and it is possible to find the optimal value $\beta_{opt}^{(i)}$ that minimizes $\delta^{(i)}$. Furthermore, we will assume that GD2 selects $\beta_{opt}^{(i)}$ when computing the correction $\mathbf{t}^{(i)}$.

Proposition 1 *If GD2 selects $\beta^{(i)} = \beta_{opt}^{(i)}$, then in general $\delta_{GD}^{(i)} \geq \delta_{GD2}^{(i)}$ and the equality holds at every iteration if $M = A - \tau B$.*

Proof If we rewrite Eq. (5) as

$$\begin{aligned}\mathbf{d}^{(i)} &= (Q(A-\tau B)M^{-1} - I)\mathbf{r}^{(i)} + Q(A-\tau B)M^{-1}B\mathbf{x}^{(i)}(\beta^{(i)} - \theta^{(i)}) \\ &= \mathbf{d}_{GD}^{(i)} + Q(A-\tau B)M^{-1}B\mathbf{x}^{(i)}(\beta^{(i)} - \theta^{(i)}),\end{aligned} \tag{6}$$

then the application of the least squares leads to bound $\delta_{GD2}^{(i)}$ as

$$\delta_{GD2}^{(i)} = \min_{\beta} \delta^{(i)} = \|(Q(A-\tau B)M^{-1}B\mathbf{x}^{(i)})^{\perp}\mathbf{d}_{GD}^{(i)}\| \leq \|\mathbf{d}_{GD}^{(i)}\| := \delta_{GD}^{(i)},$$

where $\mathbf{v}^{\perp} = I - \mathbf{v}(\mathbf{v}^*\mathbf{v})^{-1}\mathbf{v}^*$. Of course, when $M = A - \tau B$, $\mathbf{d}_{GD}^{(i)} = \mathbf{0}$, and hence $\mathbf{d}_{GD2}^{(i)} = \mathbf{0}$.

From Corollary 1 it is clear that the possible utility of the double expansion is in problems with preconditioners M^{-1} far from $(A - \tau B)^{-1}$.

3.1 Convergence analysis

We approach the analysis of the convergence of GD by relating the Davidson iterations with inexact inverse iterations.

Algorithm 2 (Inexact Inverse Iteration)
Input: initial approximate eigenvector $\mathbf{x}^{(0)}$.
Output: $\theta^{(k)}$ and $\mathbf{x}^{(k)}$ from the last iteration k.
For $i = 1, 2, \ldots$

1. Choose the shift $\sigma^{(i)}$ and the tolerance $\xi^{(i)}$.

2. Find $\mathbf{y}^{(i)}$ such that

$$\|(A-\sigma^{(i)}B)\,\mathbf{y}^{(i)}-B\mathbf{x}^{(i-1)}\|\le\xi^{(i)}\|B\mathbf{x}^{(i-1)}\|. \tag{7}$$

3. Set $\mathbf{x}^{(i)}=\mathbf{y}^{(i)}\|B\mathbf{y}^{(i)}\|^{-1}$ and $\theta^{(i)}=\rho(\mathbf{x}^{(i)})$.

4. Test for convergence.

A quite general convergence theory of Algorithm 2 is presented in [Freitag and Spence, 2007], for the computation of a finite eigenvalue θ_1 and the corresponding right eigenvector $\mathbf{x}_1$ of a generalized nonsymmetric eigenvalue problem. Consider the block factorization of $A-\theta B$ as

$$U^{-1}(A-\theta B)X=\begin{bmatrix}t_{11}&\mathbf{0}^*\\ \mathbf{0}&T_{22}\end{bmatrix}-\theta\begin{bmatrix}s_{11}&\mathbf{0}^*\\ \mathbf{0}&S_{22}\end{bmatrix}, \tag{8}$$

with nonsingular square matrices U and X of size n. In [Freitag and Spence, 2007], $\mathbf{x}^{(i)}$ is decomposed as

$$\mathbf{x}^{(i)}=\alpha^{(i)}(\mathbf{x}_1\zeta^{(i)}+X_2\mathbf{p}^{(i)}),$$

for some $\zeta^{(i)}\in\mathbb{C}$ and $\mathbf{p}^{(i)}\in\mathbb{C}^{n-1}$, where $X=[\mathbf{x}_1\ X_2]$ and $\alpha^{(i)}$ is chosen so that $\|B\mathbf{x}^{(i)}\|=1$. Then the quotient

$$\pi^{(i)}=\frac{\|S_{22}\mathbf{p}^{(i)}\|}{|s_{11}\zeta^{(i)}|}$$

is introduced as a measure for convergence, since it can be interpreted as a generalized tangent of the angle between $\mathbf{x}^{(i)}$ and the eigenvector $\mathbf{x}_1$. The following theorem shows the conditions guaranteeing that $\pi^{(i)}$ decreases linearly.

Theorem 1 *Let $(\theta_1,\mathbf{x}_1)$ be an algebraically simple eigenpair of (1) and let the decomposition (8) be induced by $\mathbf{x}_1$, with $\theta_1=t_{11}/s_{11}$. Assume that the initial guess $\mathbf{x}^{(0)}$ satisfies $0<\|S_{22}\mathbf{p}^{(0)}\|<1$ and $\sigma^{(i)}\notin\lambda(T_{22},S_{22})$. If $\sigma^{(i)}$ and $\tau_{(i)}$ are chosen in Algorithm 2 so that*

$$|\theta_1-\sigma^{(i)}|<\frac{\|(T_{22}-\theta_1S_{22})^{-1}\|^{-1}}{2\,\|S_{22}\|}\|S_{22}\mathbf{p}^{(i)}\|,$$

and

$$\xi^{(i)}<\frac{\alpha^{(i)}}{\|B\mathbf{x}^{(i)}\|\,\|\mathbf{u}_1\|}\,\beta\,|s_{11}\zeta^{(i)}|,$$

with $0\le2\beta<1-\pi^{(0)}$, then Algorithm 2 converges linearly.

Proof See Theorem 3.4 in [Freitag and Spence, 2007].

Using the convergence theory of Inexact Inverse Iteration, we shall prove the linear convergence of GD2. For that we rewrite the step of the computed correction $\mathbf{t}^{(i)}$ (step 3 in Algorithm 1) represented by Eq. (4) in the form of Eq. (7),

$$(A-\tau B)(\mathbf{t}^{(i)}-\mathbf{x}^{(i)})(\tau-\theta^{(i)})^{-1}-B\mathbf{x}^{(i)}=\mathbf{d}^{(i)}_{GD2}(\tau-\theta^{(i)})^{-1}.$$

Hence if the computed correction obtained in Algorithm 1, $\mathbf{t}^{(i)}$, is used in Algorithm 2 as

$$\mathbf{y}^{(i)} = (\mathbf{t}^{(i)} - \mathbf{x}^{(i)})(\tau - \theta^{(i)})^{-1},$$

we can bound $\delta_{III}^{(i)}$ assuming that the step 3 constraint in Algorithm 1 holds every iteration, as

$$\delta_{III}^{(i)} = \|(A - \sigma^{(i)}B)\,\mathbf{y}^{(i)} - B\mathbf{x}^{(i-1)}\| = \|\mathbf{d}_{GD2}^{(i)}\||\tau - \theta^{(i)}|^{-1} \leq \xi_1\|\mathbf{r}^{(i)}\||\tau - \theta^{(i)}|^{-1}. \tag{9}$$

Corollary 1 *Let $(\lambda_1, \mathbf{x}_1)$ be an algebraically simple eigenpair of (1). Assume that the initial vector $\mathbf{x}^{(0)}$, and the values of $\sigma^{(i)} = \tau$ and $\xi^{(i)} := \xi_{III}^{(i)}$, where*

$$\xi_{III}^{(i)} = \xi_1\|\mathbf{r}^{(i)}\||\tau - \theta^{(i)}|^{-1}, \tag{10}$$

satisfy the conditions of Theorem 1. Then Algorithm 1 converges linearly to $(\lambda_1, \mathbf{x}_1)$.

Proof Note that $\|B\mathbf{x}^{(i)}\| = 1$ and Eq. (9) yield

$$\delta_{III}^{(i)} \leq \xi_1\|\mathbf{r}^{(i)}\||\tau - \theta^{(i)}|^{-1} \leq \xi_{III}^{(i)}\|B\mathbf{x}^{(i)}\|.$$

Note that the suggested value of $\xi_{III}^{(i)}$ at Eq. (10) is similar to the proposed one at Remark 3.5 in [Freitag and Spence, 2007].

4 Algorithmic details and numerical experiments

4.1 Block versions

As a consequence of expanding the search subspace with two vectors, the matrix-vector product, the preconditioner application and the orthogonalization of the search subspace double their cost, at least, per iteration with respect to the single vector expansions. Rearranging the operations in blocks may improve the data locality and reduce the time spent per vector.

Some high-performance libraries provide multivector sparse matrix-dense vector product and are available even for new computer architectures such as multicore processors and GPUs [Williams *et al.*, 2009]. Unfortunately, the availability of multivector preconditioners is quite reduced. Still, in very large problems, consecutively applying the preconditioner to the vectors may imply a time reduction since the second application can reuse the preconditioner data already available in faster memory level. An example of this can be found in [Romero and Roman, 2011].

The SVQB [Stathopoulos and Wu, 2002] method computes an orthogonal basis of a set of vectors, but instead of orthogonalizing vector by vector, it can work with several vectors employing matrix-matrix operations almost exclusively. An interesting candidate for updating the search subspace V with new vectors W is the GS-SVQB that combines a Gram–Schmidt procedure to orthogonalize the new vectors against V, $W \leftarrow (I - VV^*)W$, and SVQB for the inner orthogonalization of the new vectors. Like (classical or modified) Gram–Schmidt, an iterative version of GS-SVQB that achieves good orthogonality levels is available.

Table 1: Iterations spent by the solvers versus the preconditioner quality

α	$10^{\frac{0}{3}}$	$10^{\frac{1}{3}}$	$10^{\frac{2}{3}}$	$10^{\frac{3}{3}}$	$10^{\frac{4}{3}}$	$10^{\frac{5}{3}}$	$10^{\frac{6}{3}}$
$\gamma(M_\alpha)$	1.2	1.1	2.5	9.6	31.3	251.6	20.5
GD	13	16	22	34	65	174	–
GD2	18	22	28	34	50	78	126

Remark 1 *As mentioned before,* $A\mathbf{u}$ *and* $B\mathbf{u}$ *will asymptotically become linear dependent. Therefore, it may be sensible to first determining an orthonormal basis for* $[A\mathbf{u}\ B\mathbf{u}]$ *before applying the preconditioner* M*. However, we note that in experiments we did not encounter an example in which the orthogonalization of* $[A\mathbf{u}\ B\mathbf{u}]$ *decreased the number of iterations significantly.*

4.2 Convergence study

First we illustrate the potential of the new expansion by computing interior eigenvalues of a diagonal generalized eigenvalue problem of order $n = 200$, formed by diagonal matrices A and B, with $a_{i,i} = i$ and $b_{i,i} = n - i + 1$. We study the impact of the quality of the preconditioner in the convergence. For that, we use a preconditioner with a configurable quality,

$$M_\alpha^{-1} = (A - \tau B + \alpha E)^{-1}, \tag{11}$$

where E is a diagonal matrix whose diagonal elements are random numbers uniformly distributed in the interval $[-1, 1]$. The quality of a preconditioner M for $A - \tau B$ may be estimated by its difference relative to M:

$$\gamma(M) = \|M^{-1}(M - (A - \tau B))\| = \|I - M^{-1}(A - \tau B)\|.$$

For the above eigenproblem, we obtained larger values of $\gamma(M_\alpha)$ for larger values of α, when α has a relatively low value (see Table 1). We present results up to this tendency is dramatically inverted, because for large values of α, $\gamma(M_\alpha)$ lowers to 1:

$$\lim_{\alpha\to\infty} \|I - M_\alpha^{-1}(A - \tau B)\| = \lim_{\alpha\to\infty} \|(A - \tau B + \alpha E)^{-1}\alpha E\| = 1.$$

We look for the eigenvalues closest to τ, which is set to the arithmetic mean of the eigenvalues of the pair (A, B). The tolerance on the residual norm for the convergence of the eigenpair is 10^{-10}. The search subspace is bounded to 50 vectors and the solvers restart with 25 vectors. The harmonic Rayleigh-Ritz procedure extracts the approximate eigenpairs from the search subspace. In Figure 1 we have plotted the residual norm against the number of applications of the preconditioner (which is the same as the application of the matrix A, and also for matrix B), until one eigenpair converges. We can infer from the left plot that, in the case of using the high quality preconditioner M_0, the extra vector of the new expansion does not accelerate the convergence, compared with GD. However, using the low quality preconditioner M_α with

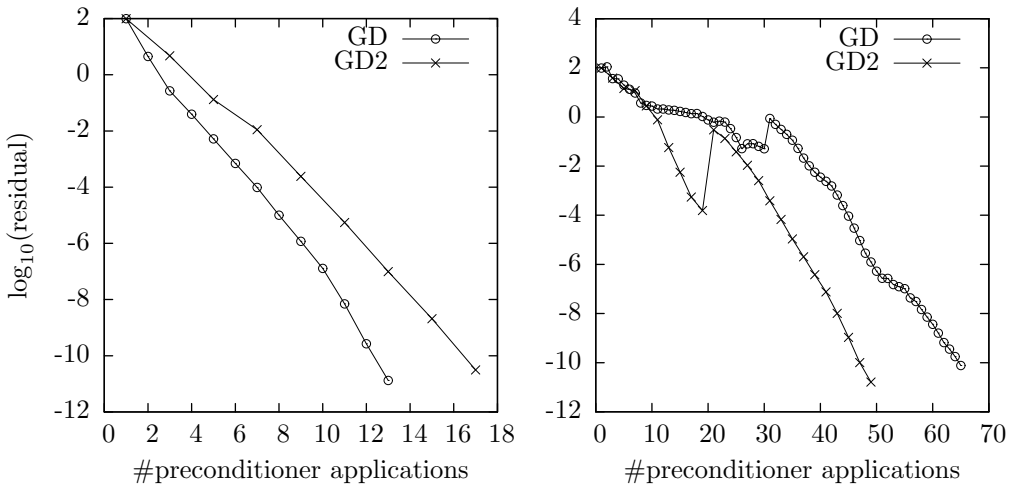

Figure 1: Residual norm against preconditioner applications spent by single expansion (GD) and the new expansion (GD2) using the high quality preconditioner M_0 (left plot) and the low quality preconditioner M_α with $\alpha = 10^{4/3}$ (right plot).

$\alpha = 10^{\frac{4}{3}}$ the acceleration of the new expansion is evident. Table 1 shows the progressive effect of the preconditioner quality on the total number of preconditioner applications required by the solvers. One sees that the new expansion is less sensitive to the lack of preconditioner quality and its performance is dramatically better in (very) low quality cases (in which the GD required more than 1000 iterations to converge).

We tested the approach also on standard eigenvalue problems, in particular with a diagonal matrix A with entries $a_{i,i} = i/(n-i+1)$, that has the same solutions as the generalized example above. In this case, we found that the convergence history is very similar to the one shown in Fig. 1.

We have checked the conclusion inferred in the previous simple case by testing the new expansion in a collection of 262 problems, both standard and generalized. The problem matrices were taken from the real and complex matrices available in the University of Florida Sparse Matrix Collection[1], disregarding their original application. The targets have been set to the arithmetic mean of the eigenvalues of the problem, which guarantees that the obtained approximate eigenvalues are interior. The tolerance on the residual norm for the converged eigenpair was $10^{-7}\|A\|_2$ and $10^{-7}(\|A\|_2 + |\theta|\|B\|_2)$, respectively, for standard and generalized eigenproblems. These stopping criteria come from the backward stability theory applied to eigenvalue problems and is useful to automatically set sensible tolerances considering the conditioning of

[1] http://www.cise.ufl.edu/research/sparse/matrices/

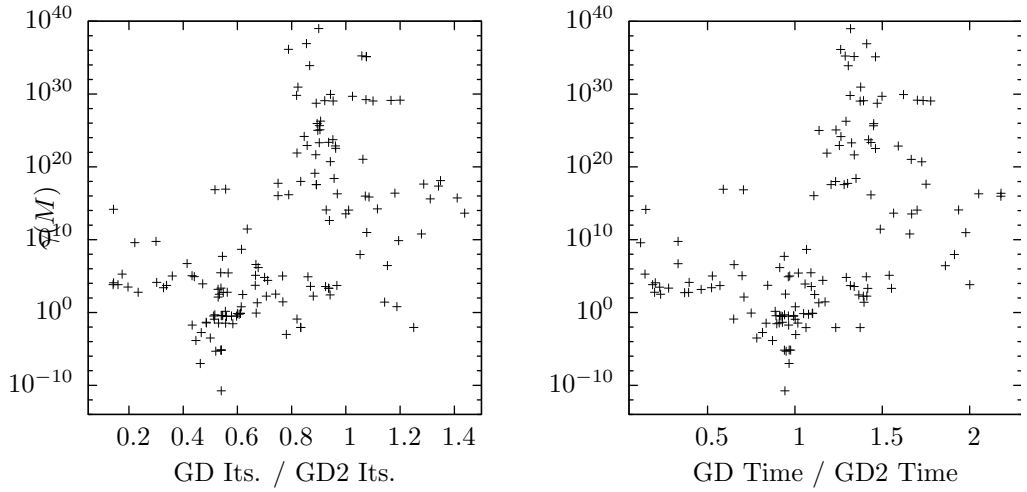

Figure 2: Gain of the new expansion (GD2) over the single expansion (GD) in number of preconditioner applications (left plot) and time (right plot), versus the quality of the ILU(0) decomposition as preconditioner.

the problem (see [Higham and Higham, 1998, Thm 2.1]). The preconditioner used was the standard incomplete LU factorization without fill-in, ILU(0), provided by the Matlab function `ilu`. The search subspace is bounded to 100 vectors and the solvers were configured to restart with 50 vectors. All the problems are solved five times with different random initial vectors (but the same vectors for both methods).

We present the time and the iterations spent by an experimental code in Matlab with harmonic Rayleigh-Ritz procedure to extract the eigenvalues close to the target, thick restart and the possibility of selecting between the single expansion (GD) or the double expansion (GD2) approach. The executions were carry out on a machine consisting of 256 JS20 blade computing nodes, each of them with two 64-bit PowerPC 970+ processors running at 2.2 GHz. The interpreter of the Matlab code was Octave 3.3.54.

Considering an iteration as adding one vector to the search subspace, the left plot in Figure 2 represents the gain in iterations of the new expansion over the single expansion against the quality of the ILU(0) preconditioner. It is possible to observe two clusters of points roughly separated by the gain 0.7 and the quality 10: in general the preconditioner quality of the points with a gain lower than 0.7 is less than 10, and the preconditioner quality of the points with a gain higher than 0.7 is more than 10. This tendency strengthens our hypothesis that new expansion is more suitable for low quality preconditioners.

Finally, we present some timing results. Our code uses the Matlab multivector matrix-vector product and the multivector preconditioner application. We note that instead of a block orthogonalization procedure (discussed in the previous section), the simpler repeated classical Gram–Schmidt procedure is used. The right plot in Figure 2 represents the gain in time of the new expansion over the single expansion. We can observe that the quality-gain pattern is similar to the pattern shown by the left plot: the single expansion is faster in 115 out of 169 (68%) problems with preconditioner quality less than 10, while the double expansion is faster in 85 out of 93 (91%) with preconditioner quality greater than 10.

5 The SLEPc Implementation

The GD2 method described in section 2 has been added to SLEPc in version 3.3, in particular as a variant of the generalized Davidson `EPSGD` solver.

5.1 User Options

In order to use the GD2 variant when running the `EPSGD` solver, a flag must be activated with

```
EPSGDSetDoubleExpansion(EPS eps,PetscBool use_gd2)
```

or alternatively via the command-line the option `-eps_gd_double_expansion`.

When the double expansion flag has been activated, the GD solver expands the search subspace with the two vectors $KA\mathbf{u}$ and $KB\mathbf{u}$ (instead of the single vector $K\mathbf{r}$), where K is the preconditioner, $\mathbf{u}$ is the selected approximate eigenvector and $\mathbf{r}$ its associated residual vector.

References

Arbenz, P., M. Becka, R. Geus, U. Hetmaniuk, and T. Mengotti (2006). On a parallel multilevel preconditioned Maxwell eigensolver. *Parallel Comput.*, 32(2):157–165.

Crouzeix, M., B. Philippe, and M. Sadkane (1994). The Davidson Method. *SIAM J. Sci. Comput.*, 15(1):62–76.

Davidson, E. R. (1975). The Iterative Calculation of a Few of the Lowest Eigenvalues and Corresponding Eigenvectors of Large Real-Symmetric Matrices. *J. Comput. Phys.*, 17(1):87–94.

Freitag, M. A. and A. Spence (2007). Convergence Theory for Inexact Inverse Iteration Applied to the Generalised Nonsymmetric Eigenproblem. *Electron. Trans. Numer. Anal.*, 28:40–64.

Genseberger, M. (2010). Improving the Parallel Performance of a Domain Decomposition Preconditioning Technique in the Jacobi-Davidson Method for Large Scale Eigenvalue Problems. *App. Numer. Math.*, 60(11):1083–1099.

Heuveline, V., B. Philippe, and M. Sadkane (1997). Parallel Computation of Spectral Portrait of Large Matrices by Davidson Type Methods. *Numer. Algorithms*, 16(1):55–75.

Higham, D. J. and N. J. Higham (1998). Structured Backward Error and Condition of Generalized Eigenvalue Problems. *SIAM J. Matrix Anal. Appl.*, 20(2):493–512.

Hochstenbach, M. E. and Y. Notay (2009). Controlling Inner Iterations in the Jacobi–Davidson Method. *SIAM J. Matrix Anal. Appl.*, 31(2):460–477.

Morgan, R. B. and D. S. Scott (1986). Generalizations of Davidson's method for computing eigenvalues of sparse symmetric matrices. *SIAM J. Sci. Statist. Comput.*, 7(3):817–825.

Neymeyr, K. (2001). A geometric theory for preconditioned inverse iteration I: Extrema of the Rayleigh quotient. *Linear Algebra Appl.*, 322(1-3):61–85.

Olsen, J., P. Jørgensen, and J. Simons (1990). Passing the One-Billion Limit in Full Configuration-Interaction (FCI) Calculations. *Chem. Phys. Lett.*, 169(6):463–472.

Romero, E. and J. E. Roman (2011). Computing subdominant unstable modes of turbulent plasma with a parallel Jacobi–Davidson eigensolver. *Concur. Comput.: Pract. Exp.*, 23(17):2179–2191.

Sleijpen, G. L. G., A. G. L. Booten, D. R. Fokkema, and H. A. van der Vorst (1996). Jacobi-Davidson Type Methods for Generalized Eigenproblems and Polynomial Eigenproblems. *BIT*, 36(3):595–633.

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

Sleijpen, G. L. G. and H. A. van der Vorst (2000). A Jacobi–Davidson Iteration Method for Linear Eigenvalue Problems. *SIAM Rev.*, 42(2):267–293.

Stathopoulos, A. (2007). Nearly Optimal Preconditioned Methods for Hermitian Eigenproblems under Limited Memory. Part I: Seeking One Eigenvalue. *SIAM J. Sci. Comput.*, 29(2):481–514.

Stathopoulos, A. and K. Wu (2002). A Block Orthogonalization Procedure With Constant Synchronization Requirements. *SIAM J. Sci. Comput.*, 23(6):2165–2182.

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

van Lenthe, J. H. and P. Pulay (1990). A Space-Saving Modification of Davidson's Eigenvector Algorithm. *J. Comput. Chem.*, 11(10):1164–1168.

Williams, S., L. Oliker, R. Vuduc, J. Shalf, K. Yelick, and J. Demmel (2009). Optimization of sparse matrix-vector multiplication on emerging multicore platforms. *Parallel Comput.*, 35(3):178–194.