Linear System Solution by Null-Space Approximation and Projection (SNAP)

M. Ilić [*]   I. W. Turner [†]   Y. Saad[‡]

September 3, 2006

Abstract

Solutions of large sparse linear systems of equations are usually obtained iteratively by constructing a smaller dimensional subspace such as a Krylov subspace. The convergence of these methods is sometimes hampered by the presence of small eigenvalues, in which case some form of deflation can help improve convergence. The method presented in this paper enables the solution to be approximated by focusing the attention directly on the 'small' eigenspace ('singular vector' space). It is based on embedding the solution of the linear system within the eigenvalue problem (singular value problem) in order to facilitate the direct use of methods such as implicitly restarted Arnoldi, or Jacobi-Davidson for the linear system solve. The proposed method, called "Solution by Null-space Approximation and Projection" (SNAP), differs from other similar approaches in that it converts the nonhomogeneous system into a homogeneous one by constructing an annihilator of the right hand side. The solution then lies in the null space of the resulting matrix. We examine the construction of a sequence of approximate null spaces using a Jacobi-Davidson style singular value decomposition method, called restarted SNAP-JD, from which an approximate solution can be obtained. Relevant theory is discussed and the method is illustrated by numerical examples where SNAP is compared with both GMRES and GMRES-IR.

[*]School of Mathematical Sciences, Queensland University of Technology, AUSTRALIA `mailto:m.ilic@qut.edu.au`

[†]School of Mathematical Sciences, Queensland University of Technology, AUSTRALIA `mailto:i.turner@qut.edu.au`

[‡]Department of Computer Science and Engineering, University of Minnesota, USA `mailto:saad@cs.umn.edu`. Work supported by NSF under grant ACI-0305120 and by the Minnesota Supercomputer Institute.

1 Introduction

The most popular general-purpose iterative solution techniques for solving a system of linear equations $A\,x = b$, when the coefficient matrix $A \in \mathbb{R}^{n \times n}$ is large and sparse, is to utilise a combination of preconditioning with Krylov subspace methods, see, e.g., [19]. Krylov methods such as $GMRES$ [20, 19], construct a subspace $\mathcal{K}_m\,(A, b) = span\{b, Ab, \ldots, A^{m-1}b\}$ from which an approximate solution is extracted. If A is non-singular and m is the algebraic grade of A, $\mathcal{K}_m\,(A, b)$ is an invariant subspace on which the system has a unique solution. However, in most cases m is less than the algebraic grade, $\mathcal{K}_m\,(A, b)$ is not invariant, or even an approximately invariant subspace [13] and one requires more complicated procedures such as restarting, deflation or preconditioning, to obtain a good approximate solution.

In general if A has a number of small in magnitude eigenvalues and b contains components of the corresponding eigenvectors, then the Krylov methods may stagnate [14]. Other factors such as negative or complex eigenvalues can also cause problems with solving some particular linear systems and even if the eigenvalue distribution appears reasonable, convergence may still not always occur [11]. It is well known that if a Krylov subspace grows large enough some automatic deflation of small eigenvalues occurs naturally, leading to a superlinear convergence of Krylov methods, see, e.g., [27, 25]. However restarting GMRES prevents this deflation from taking place, resulting in slow convergence and in some instances stagnation.

Although the Krylov subspace $\mathcal{K}_\ell\,(A, b)$, $\ell < m$ is not an approximately invariant subspace [14], it may contain subspaces that are, for example approximate eigenspaces, and some information from these subspaces, particularly the set of eigenvectors associated with the smallest eigenvalues of A, can be retained at the point of restart and used during the next cycle to overcome the problems of slow convergence [6]. GMRES-E [17] and GMRES-IR [18] are two such methods that do exactly that by augmenting a selection of harmonic Ritz vectors associated with the smallest harmonic Ritz values either to the back, or front, of the Krylov subspace respectively. It is noted by Morgan that augmenting approximate eigenvectors to the subspace deflates the corresponding eigenvalues from the spectrum of A. These eigenvector approximations do not need to be fully converged in order to offer some benefits for the deflation process [18].

Another way to overcome the problems of slow convergence, or stagnation, is the construction of a preconditioner that effectively shifts the troublesome small eigenvalues to approximately one, thus enabling the standard method to work [10]. In [1] this idea is improved upon by using implicit restarting to refine these eigenvectors before they are used to construct the preconditioner. Other methods, see, e.g., [9] employ singular vectors instead of eigenvectors.

Since the eigenvalue problem for small eigenpairs has to be solved to improve the convergence rate of the Krylov subspace method for solving the linear system $A\,x = b$, this

strongly suggests embedding the solution of the linear system into an eigenvalue problem, or to be more accurate, a null space vector problem. The linear system $A\,x = b$ can be trivially transformed into such a problem, since the solution x is a null space vector of the matrix $\bar{A} = (I - \frac{bb^T}{\|b\|^2})A$. This new way of thinking enables one to consider powerful strategies utilized in solving eigenvalue problems. For example methods like implicitly restarted Arnoldi [16] or Jacobi-Davidson (JD) [12, 21] become attractive.

This gives the theoretical motivation of our alternative approach for overcoming poor convergence, which we have called *Solution by Null-space Approximation and Projection (SNAP)*. SNAP transforms the problem of solving $A\,x = b$ into a problem of finding the null space of $\bar{A}$, $\mathcal{N}(\bar{A})$. For small size matrices there is a number of methods for computing the null space (see for example [4] and the references within). However, for large sparse matrices, two strategies present themselves. Either we think of zero as an eigenvalue and use methods developed for finding the eigenspace corresponding to zero and small eigenvalues [1, 6, 10, 17, 8]; think of zero as a singular value and use the methods developed for extracting singular vectors [7] associated with 'small singular values' [26]. This latter approach will be the method that we harness when constructing our SNAP-JD algorithms.

This paper is organised as follows. In § 2 we present SNAP within the GMRES theoretical context by constructing a null space vector as the difference between two inhomogeneous solutions, one of which is known at the outset. There is similarity with the discussion for inconsistent systems given in [5]. In § 2.1 we construct the annihilator E as an orthogonal, or oblique, projector and derive the bound on the residual error $\|\,Ax - b\,\|$ in terms of the error $\|\,E\,A\,x\,\| = \sigma$. In § 2.2 we discuss how the initial approximate null space vector can be improved by an iterative procedure, which is implemented in the SNAP-JD algorithms presented in § 4. In § 3 we discuss some relevant theoretical results on approximate null spaces, with Proposition 3 being a key result that motivates the practical ideas. In § 4 the JD analogous subspace expansion method is presented and § 5 sees the theory illustrated by some specific numerical case studies where we compare and comment on the performance of SNAP against variants of GMRES algorithms. In § 6 we summarise the advantages and disadvantages of the SNAP method and point to the future research projects.

To avoid the discussion of inconsistency and deflation procedures when the null space is multidimensional, we will assume that the matrix A is nonsingular and that the right-hand side vector $b \neq 0$. Unless otherwise stated, the norm $\|\,.\,\|$ is the Euclidean norm $\|\,.\,\|_2$. Coordinate vectors of $\mathbb{R}^n$ are denoted by $e_i = (0, \ldots, 1, \ldots, 0)^T$, where 1 is in the ith position.

2   SNAP

Given a linear system $A\,x = b$, the method proposed in this paper starts by building a projector E which we refer to as an *annihilator* of b, such that $E\,b = 0$. Thus, E transforms the original nonhomogeneous system into the homogeneous system $E\,A\,x = 0$. § 2.1 shows three ways to obtain the projector E. We denote by $\bar{A}$ the matrix $\bar{A} = EA$.

The core idea of SNAP lies in its second step where a sequence of what we call approximate null subspaces (ANS) of dimension k and order ε is generated for $\bar{A}$,

$$N(\bar{A}, k, \varepsilon) = \left\{ span\{w_1, w_2, \ldots, w_k\} \mid \|\bar{A}W\| < \varepsilon, \ W = [w_1, w_2, \ldots, w_k] \right\}$$

where the w_i's are orthonormal and ε is decreasing. Note that this space need not be a Krylov subspace and existing methods developed specifically for the eigenvalue problem, see, e.g., [23] or [3] and the references therein, can be employed to obtain the w_i's.

Once this approximate null space is obtained, the third step of the method extracts the approximate solution, which is a scalar multiple of the approximate null space vector. Details of these steps are given in the next sections.

2.1   Formulation of an Annihilator and Error Bound

The annihilator for a given b is not unique. For illustrative purposes we give two main examples that we call an orthogonal projector $E_\perp$ and an oblique projector $E_/$.

Proposition 1 *Define the orthogonal projector $E_\perp = I - \frac{b\,b^T}{\|b\|^2}$ and let a unit vector w be given which satisfies $\|E_\perp\,A\,w\| = \sigma$. Assume that $b^T Aw \neq 0$ and let $\beta = \|b\|^2/(b^T Aw)$ and $x = \beta\,w$. Then the linear system $A\,x = b$ has approximate solution x with residual $r = b - A\,x$ satisfying $\|\,r\,\| = |\beta|\,\sigma$.*

Proof. Clearly:

$$E_\perp\,A\,w = \left(I - \frac{bb^T}{\|b\|^2}\right) A\,w = A\,w - \frac{1}{\beta}\,b = \frac{1}{\beta}\,(A\,x - b)\,.$$

The assumptions now show that $\|r\| = \sigma|\beta|$.   ∎

If $b^T Aw = 0$, then no information on x can be extracted. However, this problem is not too likely to occur in practice. Instead it may happen that w corresponds to a small singular right vector of A in which case $\|Aw\|$ is small, giving rise to a large β.

Another way to interpret the result in Proposition 1 is via the following relative residual:

$$\frac{\|r\|}{\|b\|} = \frac{\|E_\perp\,Aw\|}{\|Aw\|}\frac{1}{|\cos\ \theta|}$$

where θ is the acute angle between b and Aw. This shows that provided $\theta \neq \frac{\pi}{2}$, it is desirable that $E_\perp$ reduces Aw in relative terms.

In the discussion above the relative residual was measured in terms of the 2-norm. If it is more convenient to use the ∞-norm, the oblique projector $E_/ = I - \frac{1}{b_j} b\, e_j^T$, where $|b_j| = \|b\|_\infty$ may be beneficial. In this case, the equivalent of Proposition 1 is

$$\frac{\|r\|}{\|b\|_\infty} = \frac{\|E_/ \, Aw\|}{\|Aw\|} \frac{1}{|\cos\ \theta|}$$

where this time θ is the angle between e_j and Aw.

Yet another example of an oblique projector would be $E = I - \frac{1}{\|b\|_1} b\, e^T$, where e is the vector with components $e_j = \begin{cases} 1 & \text{if } b_j \geq 0 \\ -1 & \text{if } b_j < 0 \end{cases}$ giving in this case

$$\frac{\|r\|}{\|b\|_1} = \frac{\|EAw\|}{\|Aw\|\|e\|} \frac{1}{|\cos\ \theta|}, \quad \theta \text{ being the angle between } e \text{ and } Aw.$$

2.2   Construction of an Approximate Null Space

To obtain an approximate null space vector of $\bar{A}\, x = 0$, where $\bar{A} = E_\perp A$ or $\bar{A} = E_/ A$, we can exploit the GMRES algorithm. Although $\bar{A}$ is singular, the system to be solved is consistent and as stated in [24], GMRES will not break down unless the solution is found.

If v_0 is an arbitrary vector, we can approximately solve the homogenous linear system $\bar{A}x = 0$ by GMRES, starting with the initial guess v_0. This is equivalent to approximately solving the linear system $\bar{A}v = f$, where $f = \bar{A}v_0$ by GMRES starting with a zero initial guess, and then taking $v_0 - v$ as the approximate null space vector. A more detailed description follows.

(i)   Select a unit vector v_0 and set $v_1 = \bar{A}\, v_0$.

(ii)   Process Arnoldi up to some specified point ℓ, resulting in the relation $\bar{A}\, V_l = V_{l+1}\, \bar{H}_l$, $V_l = [v_1, \ldots, v_l]$.

(iii)   Compute the minimizer y_l of $\min_{y \in \mathbb{R}^l} \| \bar{A}\, V_l y - v_1 \|$ as $y_l = \| v_1 \| \bar{H}_l^\dagger e_1$, where $\bar{H}_l^\dagger$ is the pseudoinverse of $\bar{H}_l$.

(iv)   The desired approximate null space vector is given by $w_1 = v_0 - V_l\, y_l$

One can show that $\bar{A}\, w_1 = (I - P)\, v_1$, where $P = V_{l+1}\, \bar{H}_l\, \bar{H}_l^\dagger\, V_{l+1}^T$ is a projector onto the Krylov subspace $\mathcal{K}_l(v_1, \bar{A})$.

In step (i), if initially GMRES was used for solving $A\,x = b$ and stagnation occurred, then the resulting approximate solution obtained at that point would be suitable for use as v_0. Otherwise one can just take a random v_0.

In step (ii), ℓ should not exceed the *analytic grade* defined in [14]. The analytic grade of order t of the matrix $\bar{A}$ with respect to v_1 is defined as the lowest integer ℓ for which $\frac{\|u_\ell - \mathcal{P}_\ell u_\ell\|}{\|u_\ell\|} < 10^{-t}$ where $\mathcal{P}_\ell$ is the orthogonal projector onto the l-th Krylov subspace $\mathcal{K}_l$ and $u_l = \bar{A}^\ell v_1$. The grade can be computed from the Arnoldi algorithm using the matrices $\bar{H}_1, \bar{H}_2, \cdots, \bar{H}_l$ generated during the process. If h_1 is the 1st column of $\bar{H}_1$, and $h_i = \bar{H}_i h_{i-1}$, for $i = 1, \ldots, \ell$ then $\frac{\|u_\ell - \mathcal{P}_\ell u_\ell\|}{\|u_\ell\|} = \frac{|e_{l+1}^T h_l|}{\|h_l\|}$.

At this stage we examine whether $\left\| \bar{A}\,w_1 \right\|$ is small and if that is the case, then we check that $\| r \| = |\beta| \left\| \bar{A}\,w_1 \right\| < \varepsilon$ where β is given in Proposition 1. If the residual is sufficiently small we have an approximate solution of the linear system as $x = \beta w_1$, however this will be rarely true because as was discussed in the introduction, the small singular values of $\bar{A}$ effectively hinder the convergence of GMRES. Nevertheless, it is true that the resulting vector w_1 should be rich in null space vector components and furthermore the fact that $\left\| \bar{A}\,w_1 \right\| \le \| v_1 \| = \left\| \bar{A}\,v_0 \right\|$ makes w_1 a suitable candidate for use in the subspace expansion Jacobi-Davidson like method discussed in §4.

For an initial approximation v_0, the procedure outlined so far in this section produces an approximate null space vector w_1 such that $\left\| \bar{A}\,w_1 \right\| < \varepsilon$, which may not be sufficiently accurate and an iterative procedure has to be established to further refine w_1 to the desired null space vector. Two methods suggest themselves as possible ways for this refinement.

Method 1: Generate a sequence $\{w_m\}$ such that $\left\| \bar{A}\,w_m \right\| \to 0$, then $\{w_m\} \to z \in \mathcal{N}(\bar{A})$. Our basic algorithm SNAP-JD in § 4 (with $k = k_{max}$ fixed) constructs a sequence such that $\|\bar{A}w_{m+1}\| \le \|\bar{A}w_m\|$, to which we refer to as a *monotone decreasing sequence*. The advantage of this method is that an SVD of an upper triangular matrix of dimension no more than $k_{max} \times k_{max}$ is required and the test for terminating the sequence is given by Proposition 1 on the last sequence vector. The disadvantage for choosing k_{max} small is that convergence may be slow.

Method 2: Generate a sequence of approximate null subspaces (ANS), for $\bar{A}$,

$$\left\{ span\{w_1, w_2, \ldots, w_k\} \mid \|\bar{A}W\| < \varepsilon,\ W = [w_1, w_2, \ldots, w_k] \right\}$$

where the w_i's are orthonormal and ε is decreasing. Our algorithm restarted SNAP-JD in § 4 constructs a sequence of subspaces such that $\|\bar{A}W^{(m+1)}\| \le \|\bar{A}W^{(m)}\|$, to which we refer to as a *monotone decreasing sequence of subspaces*. If $\varepsilon < \bar{\sigma}_{k+1}$, then W contains the null space vector (see Proposition 3, next section). However, since the $\bar{\sigma}$'s are unknown, the linear dependence test can be used. If $\{\bar{A}\,w_1, \ldots, \bar{A}\,w_k\}$ are linearly dependent i.e. $\sum_{i=1}^{k} c_i \bar{A}\,w_i = 0$, then $z = \sum_{i=1}^{k} c_i w_i \in \mathcal{N}(\bar{A})$. A test for linear dependence amounts to seeing whether the smallest singular value of $\bar{A}W$ is zero. If w is the vector corresponding to the smallest singular value τ of $\bar{A}W$ i.e.

$\|\bar{A}w\| \leq \tau$, and τ is less than the requirement of Proposition 1, we can terminate the procedure.

3   Approximate Null Spaces

Although the main aim of this paper is to introduce the method of annihilator and its implementation using the SNAP-JD algorithm, this section introduces some relevant results on the theoretical background of what we coin approximate null spaces. Questions such as the following motivate our thinking:

1. If we have a set of vectors satisfying $\|\bar{A}w_i\| < \varepsilon$ as discussed in the previous section do they form a vector subspace? The negative answer to this question forces us to introduce the concept of an approximate null cone (ANC).

2. When will the $span\{w_1, w_2, \ldots, w_k\}$ discussed in the previous section include the null space vector? The answer to this question requires the consideration of $\|\bar{A}W\| < \varepsilon$, which leads to our definition of approximate null subspace (ANS).

3. How do we construct the monotone decreasing sequences mentioned in the previous section?

4. How are the singular values of A, $\bar{A}$ and $\bar{A}W$ related?

Let $B \in \mathbb{R}^{n \times n}$ be a singular matrix with $\mathcal{N}(B) = \{x \in \mathbb{R}^n | Bx = 0\}$ as its null space, which is a vector subspace of $\mathbb{R}^n$.

Definition 1 *An approximate null cone (ANC), $\tilde{N}(B, \varepsilon)$ or just $\tilde{N}(B)$, of B of order ε is defined by*

$$\tilde{N}(B, \varepsilon) = \left\{ x \in \mathbb{R}^n \mid \tfrac{\|Bx\|}{\|x\|} < \varepsilon \right\}.$$

Definition 2 *Let $\{w_1, \ldots, w_k\}$ be k orthonormal (ON) vectors in $\tilde{N}(B, \varepsilon)$ and form the matrix $W = [w_1 \ \cdots \ w_k]$. An approximate null subspace (ANS), $N(B, k, \varepsilon)$ or just $N(B)$ of B of dimension k and order ε is defined by*

$$N(B, k, \varepsilon) = \{span\{w_1, w_2, \ldots, w_k\} \mid \|BW\| < \varepsilon\}.$$

3.1   Properties of ANC and ANS

1. $\tilde{N}(B, \varepsilon)$ is not a vector subspace. To see why, let $x_1 = n + \tfrac{\varepsilon}{2} v_n$, $x_2 = n - \tfrac{\varepsilon}{2} v_n$ where $n \in \mathcal{N}(B)$ and (σ_n, v_n, u_n) is the largest singular triplet of B and $\|n\|^2 + \tfrac{\varepsilon^2}{4} = 1$. Clearly, $x_1 - x_2 \notin \tilde{N}(B, \varepsilon)$, yet $x_1, x_2 \in \tilde{N}(B, \varepsilon)$.

2. $\tilde{N}(B, \varepsilon)$ is a double cone.

3. $N(B, k, \varepsilon)$ is a vector subspace and is a subset of $\tilde{N}(B, \varepsilon)$ i.e. $N(B, k, \varepsilon) \subset \tilde{N}(B, \varepsilon)$.

4. $\tilde{N}(B, \varepsilon)$ contains the null space $\mathcal{N}(B)$ but $N(B, k, \varepsilon)$ may not.

5. $\lim_{\varepsilon \to 0} \tilde{N}(B, \varepsilon) = \mathcal{N}(B)$.

6. $N(B, k, \varepsilon)$ contains $\mathcal{N}(B)$ iff $\inf_{\mathbf{y} \in \mathbb{R}^k} \|BWy\| = 0$.

7. Given k ON vectors $\{w_1, \ldots, w_k\}$ such that $\|\bar{A}w_i\| < \frac{\varepsilon}{\sqrt{k}}, \ i = 1, \ldots, k$, then
   $\text{span}\{w_1, w_2, \ldots, w_k\} \subset N(B, k, \varepsilon) \subset \tilde{N}(B, \varepsilon)$.

8. Let $B = U\Sigma V^T$ be the SVD of B. Partition $V = [V_1 \mid V_2]$ where V_1 corresponds to the smallest k singular values $\sigma_0 = 0, \sigma_1, \ldots, \sigma_{k-1}$, all less than ε and V_2 corresponds to $\sigma_k, \ldots, \sigma_n$ with $\sigma_i \geq \varepsilon, \ i = k, \ldots, n$. Then $\tilde{N}(B, \varepsilon)$ contains all subspaces generated by the singular vectors in V_1.

9. If $x \in \tilde{N}(B, \varepsilon)$, then $\frac{x}{\|x\|}$ is a linear combination $\frac{x}{\|x\|} = n + \sum_i \xi_i v_i$, where $n \in \mathcal{N}(B)$ and v_i is the right singular vector corresponding to σ_i that satisfies $\sigma_i |\xi_i| < \varepsilon$. To see this: $\frac{\|Bx\|}{\|x\|} = \{\sum_i \xi_i^2 \sigma_i^2\}^{\frac{1}{2}} < \varepsilon$. This last statement says that $\tilde{N}(B, \varepsilon)$ contains $\mathcal{N}(B)$, small singular vectors, and can contain small components of large singular vectors. To extract $\mathcal{N}(B)$ from $\tilde{N}(B, \varepsilon)$ one needs to devise strategies to eliminate the small singular vectors and minimise the contribution of small components of large singular vectors.

To monitor how many small singular values of B one could expect to encounter during the construction of the ANS the following proposition is useful. Furthermore, this knowledge can be used to make an informed judgement on whether to restart the iterative process presented in §4.

Proposition 2 *Let $\{w_1, w_2, \ldots, w_k\}$ be an orthonormal basis for $N(B, k, \varepsilon)$ and form the matrix $W = [w_1, w_2, \ldots, w_k] \in \mathbb{R}^{n \times k}$. Assume without loss of generality that the singular values of B are arranged as $\sigma_n \geq \sigma_{n-1} \geq \ldots \geq \sigma_1$. Then the smallest k singular values of B, viz, $\{\sigma_k, \ldots, \sigma_1\}$ are less than ε.*

Proof. Let , $\bar{\sigma}_k = \|BW\|$ so that $\bar{\sigma}_k \leq \varepsilon$. Define the projectors $P = WW^T$ and $Y = I - P$. Clearly, $\text{rank}(P) = k$ and $\text{rank}(Y) = n - k$. Also, $w = Wu \in \mathcal{N}(Y)$. Now,

$$\bar{\sigma}_k = \sup_{\substack{u \in \mathbb{R}^k \\ \|u\|=1}} \|BWu\| = \sup_{\substack{w \in \mathcal{N}(Y) \\ \|w\|=1}} \|Bw\| \geq \sigma_k$$

using Weyl's Theorem (refer [22, P. 67]). $\blacksquare$

Even more revealing is the following proposition which gives some justification for using Definition 2 to define ANS for our theoretical discussion.

Proposition 3 *Let $\sigma_n \geq \sigma_{n-1} \geq \ldots \geq \sigma_1 = 0$ be the singular values of B. Suppose $\|BW\| < \varepsilon$, where W is defined in Definition 2 and $\sigma_k \leq \varepsilon < \sigma_{k+1}$. Then $\mathcal{N}(B) \subset N(B, k, \varepsilon)$.*

Proof. Let P and Y be the projectors as defined in the proof of Proposition 2 so that $P + Y = I$ and $BP + BY = B$. Thus

$$\|B - BY\| = \|BP\| \leq \|BW\| < \varepsilon < \sigma_{k+1}.$$

Hence $rank(BY)$ is at least $n - k$. Otherwise if $rank(BY) = r < n - k$, then by Weyl's Theorem (Note: the statement of the theorem in the given reference uses the descending arrangement of σ_i whereas we use the ascending)

$$\|B - BY\| \geq \sigma_{n-r} \geq \sigma_{k+1} > \varepsilon,$$

since $r < n - k$ implies $r \leq n - k - 1$ or $k + 1 \leq n - r$.
Since $rank(Y) = n - k$ then $rank(BY)$ is at most $n - k$. Therefore $rank(BY) = n - k$. Now $rank(BY) = rank(Y) - \dim \mathcal{N}(B) \cap \mathcal{R}(Y)$. So that $\dim \mathcal{N}(B) \cap \mathcal{R}(Y) = 0$ or $\mathcal{N}(B) \cap N(B, k, \varepsilon)^\perp = \emptyset$. Hence the result. ∎

Another useful result that also follows from Weyl's Theorem relates the singular values of $\bar{A} = EA$ to A and is summarised by the following proposition.

Proposition 4 *Let E be either the orthogonal $E_\perp$ or oblique $E_/$ projector discussed in §2.1 and order the singular values of A as $\sigma_n \geq \sigma_{n-1} \geq \ldots \geq \sigma_1$ and of $\bar{A}$ as $\bar{\sigma}_n \geq \bar{\sigma}_{n-1} \geq \ldots \geq \bar{\sigma}_2 > \bar{\sigma}_1 = 0$. Then, $\sigma_1(A) \leq \bar{\sigma}_2(\bar{A}) \leq \sigma_3(A)$ and in general, $\sigma_i(A) \leq \bar{\sigma}_{i+1}(\bar{A}) \leq \sigma_{i+2}(A)$, $i = 1, 2, \ldots, n - 2$.*

3.2 Construction of Monotone Decreasing Sequence of ANS

The discussion of SNAP in §2 requires the construction of a monotone decreasing sequence of ANS. The iterative procedures outlined in §4 construct such sequences by generating an expanding subspace $\mathcal{X}_k$ of dimension k, over which an SVD of the projected matrix $\bar{A}$ is computed. Key results are summarised by the following proposition and its corollaries.

Proposition 5 *Let $\mathcal{X}_k$ have orthonormal basis $\{x_1, x_2, \ldots, x_k\}$ and define the matrix $X_k \in \mathbb{R}^{n \times k}$ as $X_k = [x_1, \ldots, x_k]$. Consider the subspace $\mathcal{X}_{k+1} = \mathcal{X}_k \oplus \text{span}\{x_{k+1}\}$ with $x_{k+1} \perp \mathcal{X}_k$, $\|x_{k+1}\| = 1$ and define the matrix $X_{k+1} = [x_1, \ldots, x_k, x_{k+1}]$. Then the singular values of $\bar{A}X_{k+1}$ interlace those of $\bar{A}X_k$ i.e.*

$$\sigma_1(\bar{A}X_{k+1}) \leq \sigma_1(\bar{A}X_k) \leq \sigma_2(\bar{A}X_{k+1}) \leq \sigma_2(\bar{A}X_k) \leq \cdots \leq \sigma_k(\bar{A}X_k) \leq \sigma_{k+1}(\bar{A}X_{k+1}).$$
$$(3.1)$$

9

Proof. We use the fact that σ_i is a singular value of $\bar{A}X_k$ iff σ_i^2 is an eigenvalue of the projected matrix $X_k^T B X_k$, $\quad B = \bar{A}^T \bar{A}$. Now

$$X_{k+1}^T B X_{k+1} = \begin{pmatrix} X_k^T B X_k & X_k^T B x_{k+1} \\ x_{k+1}^T B X_k & x_{k+1}^T B x_{k+1} \end{pmatrix}.$$

Cauchy's interlacing theorem (see [22], p72) gives that the eigenvalues of $X_{k+1}^T B X_{k+1}$ and $X_k^T B X_k$ interlace. Hence the result. ∎

If $\{\sigma_i^{(0)}\}_{i=1}^k$ and $\{\sigma_i^{(1)}\}_{i=1}^{k+1}$ denote the singular values of $\bar{A}X_k$ and $\bar{A}X_{k+1}$ respectively, then Proposition 5 gives

$$\sigma_i^{(1)} \le \sigma_i^{(0)}, \quad i = 1, \ldots, k. \tag{3.2}$$

Let $\bar{A}X_k = U_k^{(0)}\Sigma_k^{(0)}V_k^{(0)T}$ and $\bar{A}X_{k+1} = U_{k+1}^{(1)}\Sigma_{k+1}^{(1)}V_{k+1}^{(1)T}$ be the SVD of $\bar{A}X_k$ and $\bar{A}X_{k+1}$ respectively, with singular values ordered increasingly from top left to bottom right. Set

$$w_k = X_k v_1^{(0)}, \quad w_{k+1} = X_{k+1}v_1^{(1)}, \quad \text{and} \quad X_k^{(1)} = X_{k+1}V_k^{(1)}.$$

Corollary 1 *If w_k, w_{k+1}, X_k and $X_k^{(1)}$ are as stated above then*

(i) $\|\bar{A}w_{k+1}\| \le \|\bar{A}w_k\|$;

(ii) $\|\bar{A}X_k^{(1)}\| \le \|\bar{A}X_k\|$.

Proof. These are just the first and the last inequalities in Equation 3.2. ∎

Next let $X_m = [x_1, \ldots, x_k, x_{k+1}, \ldots, x_m]$ be constructed by further expansions and suppose that

$$\bar{A}X_m = U_m \Sigma_m V_m^T$$

is the SVD of $\bar{A}X_m$. Set $w_m = X_m v_1$ and $W_m = X_m V_k$, where v_1 is the first column of V_m and V_k contains the first k columns of V_m. Using the previous result and an induction argument, we have:

Corollary 2 *If w_m, W_m are as stated above then*

(i) *$\{\|\bar{A}w_m\|\}$ is a monotone decreasing sequence as m increases;*

(ii) *$\{\|\bar{A}W_m\|\}$ is a monotone decreasing sequence as m increases for $m > k$ (fixed).*

It is worthwhile to point out that in the same way that the Ritz values are approximations for the eigenvalues of $\bar{A}$, the $\{\sigma_i^{(m)}\}$'s approximate the singular values of $\bar{A}$. If we set $Z_m = X_m V_m$, $(Z_m^T Z_m = I)$, then

$$\bar{A}Z_m = U_m \Sigma_m$$

$$\text{and} \quad (X_m X_m^T)\bar{A}^T U_m = Z_m \Sigma_m^T,$$

which becomes the SVD for $\bar{A}$ when $m = n$, i.e., when $X_m X_m^T = I$.

4 A Jacobi-Davidson Analogue for Linear Systems

The development of the iterative procedures based on Method 1 and Method 2 discussed in §2 requires the construction of vectors w_i. For Method 1, given w_i, one needs a procedure to construct w_{i+1} such that $\|\bar{A}w_{i+1}\| \leq \|\bar{A}w_i\|$. Our basic Algorithm 1 below, takes w_i, generates a subspace W, which usually would not satisfy $\|\bar{A}W\| \leq \varepsilon$, however the vector w corresponding to the smallest singular value of $\bar{A}W$ does. Although one can use Algorithm 1 as described for Method 1 to generate the vectors $\{w_1, \ldots, w_k\}$, a preferred implementation of Method 2 is Algorithm 2. Here, starting with w_1, one generates subspace $W_k^{(1)}$ using Algorithm 1. We then select ℓ smallest singular values of $\bar{A}W_k^{(1)}$ and their corresponding vectors $S_1 = \{w_1^{(1)}, \ldots, w_\ell^{(1)}\}$ to form the ANS $X_\ell^{(1)} = Span(S_1)$, which in turn is expanded in a restarting procedure to $W_k^{(2)}$, from which we select ℓ smallest singular values and their corresponding vectors $S_2 = \{w_1^{(2)}, \ldots, w_\ell^{(2)}\}$ to form the ANS $X_\ell^{(2)} = Span(S_2)$ and so on. At each stage the smallest singular value is used to test for termination.

4.1 Basic SNAP-JD

Corollary 2 of Proposition 5 indicates that merely expanding a subspace by an orthonormal vector will result in an improved convergence. To accelerate the convergence, we describe in this section an algorithm based on the Jacobi-Davidson (JD) idea for choosing the expansion vector. Our basic algorithm (Algorithm 1) starts with an initial vector w_1, constructs a sequence of expanding subspaces, obtains the corresponding SVD, and tests whether the smallest singular value is sufficiently small. The expansion is continued until the test is satisfied or up to some prescribed maximum dimension k_{max}. In the latter case the vector corresponding to the smallest singular value is used to restart the process. This procedure can be repeated, see §4.3 for further details.

Suppose $\mathcal{X}_k$ is the kth constructed subspace with an orthonormal basis $\{x_1, x_2, \ldots, x_k\}$ and define the matrix $X_k \in \mathbb{R}^{n \times k}$ as $X_k = [x_1, \ldots, x_k]$. Let $\bar{A}X_k = Q_k R_k$ be the QR-factorisation of $\bar{A}X_k$, and let $R_k V_k = U_k \Sigma_k$ be the SVD of R_k, with the singular values ordered increasingly from top left to bottom right. Then the smallest singular triplet of the projection of $\bar{A}$ onto $\mathcal{X}_k$ is given by $\{\sigma_1, X_k v_1, Q_k u_1\}$. Set $w_k = X_k v_1$ so that

$$\left\| \bar{A}\,w_k \right\|^2 = \inf_{y \neq 0} \frac{\left\| \bar{A}\,X_k\,y \right\|^2}{\left\| y \right\|^2} = \left\| R_k\,v_1 \right\|^2 = \left\| \sigma_1 u_1 \right\|^2 = \sigma_1^2. \tag{4.1}$$

Suppose now that subspace $\mathcal{X}_k$ is expanded by one orthonormal vector x_{k+1}. Set $X_{k+1} =$

$[x_1, \ldots, x_k, x_{k+1}]$ where $x_{k+1}^T X_k = 0^T$ and write the QR-factorisation of $\bar{A} X_{k+1}$ as

$$\bar{A} X_{k+1} = [Q_k, q_{k+1}] \begin{pmatrix} R_k & Q_k^T \bar{A} x_{k+1} \\ 0^T & \|(I - Q_k Q_k^T) \bar{A} x_{k+1}\| \end{pmatrix}$$

$$= Q_{k+1} R_{k+1}.$$

Let $w_{k+1} = X_{k+1} z$, where $z = \begin{pmatrix} y \\ y_{k+1} \end{pmatrix}$ solves the problem

$$\| \bar{A} w_{k+1} \|^2 = \inf_{z \neq 0} \frac{(y^T, y_{k+1}) \begin{pmatrix} R_k^T R_k & R_k^T r \\ r^T R_k & \alpha \end{pmatrix} \begin{pmatrix} y \\ y_{k+1} \end{pmatrix}}{\| y \|^2 + y_{k+1}^2}$$

$$= \inf_{z \neq 0} \frac{y^T R_k^T R_k y + 2 (r^T R_k y) y_{k+1} + \alpha y_{k+1}^2}{\| y \|^2 + y_{k+1}^2} \qquad (4.2)$$

where $r = Q_k^T \bar{A} x_{k+1}$ and $\alpha = \| \bar{A} x_{k+1} \|^2$. Let y in (4.2) be the vector which produced w_k and choose y_{k+1} to decrease the numerator in (4.2) i.e., minimise

$$f(t) = \| y \|^2 \sigma_1^2 + 2 (r^T R_k y) t + \alpha t^2.$$

This occurs when $t = -\frac{r^T R_k y}{\alpha}$ to give $f(-\frac{r^T R_k y}{\alpha}) = \| y \|^2 \sigma_1^2 - \frac{(r^T R_k y)^2}{\alpha}$, which implies

$$f = \| y \|^2 \left(\sigma_1^2 - \frac{(r^T R_k y)^2}{\alpha \| y \|^2} \right).$$

Now

$$\frac{r^T R_k y}{\| y \|} = \sigma_1 (\bar{A} x_{k+1})^T Q_k u_1 = (\bar{A} x_{k+1})^T \bar{A} w_k.$$

Since $\| \bar{A} w_k \| = \sigma_1$ and $\| \bar{A} x_{k+1} \|^2 = \alpha$ we have

$$f = \| y \|^2 \sigma_1^2 (1 - \cos^2 \theta),$$

where θ is the angle between $\bar{A} w_k$ and $\bar{A} x_{k+1}$. Thus if $\bar{A} w_k$ and $\bar{A} x_{k+1}$ are aligned, i.e. when $\theta = 0$ or π, then $f = 0$. In this case $\| \bar{A} w_{k+1} \| = 0$ and we have the null space vector.

Hence the expansion vector x_{k+1} should be chosen as the (approximate) solution of the correction equation $\bar{A} x_{k+1} = \gamma \varepsilon_k$ and such that $x_{k+1} \perp \mathcal{X}_k$, where γ is some appropriately chosen constant and the null space vector residual is $\varepsilon_k = \bar{A} w_k$.

It can be concluded from the above statement that the expansion vector x_{k+1} could be computed by solving $\bar{A} t = -\varepsilon_k$ directly. However, as noted for the JD method [21], it is preferable to compute the orthogonal correction t for w_k so that we restrict ourselves to the subspace $w_k^\perp$ when determining t. Therefore, to obtain the correction vector the following strategy is recommended:

1. Solve approximately the correction equation

$$(I - w_k\, w_k^T)\, \bar{A}\, (I - w_k\, w_k^T)\, t = -(I - w_k\, w_k^T)\, \varepsilon_k. \tag{4.3}$$

for t.

2. Form

$$x_{k+1} = \frac{t - X_k\, X_k^T\, t}{\| t - X_k\, X_k^T\, t \|}. \tag{4.4}$$

We make two important remarks. First, from Proposition 5 Cor 1, a reasonable expansion vector should suffice. There is no need to solve the correction equation too accurately. We noted above that one needs to monitor $\cos\theta$ and if we approximate the solution of $\bar{A}\, x = u$ by GMRES then $\bar{A}\, V_m = V_{m+1}\, \bar{H}_m$, and if $\bar{H}_m = Q_m\, R_m$ and $\bar{H}_m^\dagger = R_m^T\, (R_m\, R_m^T)^{-1}\, Q_m^T$, $x = \| u \|\, V_m\, \bar{H}_m^\dagger\, V_{m+1}^T\, v_1$, then it can be easily shown that $u^T \bar{A}\, x = \| u \|^2\, \| Q_m^T\, e_1 \|^2$. Therefore,

$$\cos\theta = \frac{\| u \|^2\, \| Q_m^T\, e_1 \|^2}{\| \bar{A}\, x \|} = \| u \|\, \| Q_m^T\, e_1 \|. \tag{4.5}$$

For small θ, $\cos\theta$ varies slowly. As a result, in practice, if θ is small then stop after 2 or 3 additional Arnoldi operations.

A second remark is that the most expensive part of SNAP-JD is in the construction of the expansion vector x_{k+1}. Methods other than Krylov subspace techniques may be suitable for this task.

4.2   Preconditioning

To accelerate the convergence of SNAP-JD it may prove beneficial to use some form of preconditioning on the original nonsingular matrix A. Since by Proposition 4, the smallest nonzero singular values of EA are not less than the smallest singular values of A, any preconditioner that works favourably for A will also work favourably for EA. We advocate the use of a right preconditioning matrix M^{-1} as follows:

$$EAM^{-1}y = 0.$$

The solution of the original system is then, $x = \beta w$, where $w = M^{-1}y$.

A single algorithm encapsulating the SNAP-JD idea using preconiditioning and GMRES to find an approximate expansion vector is given in Algorithm 1. Here m refers to the number of correction equation solves to perform for each expansion vector.

Algorithm 1 : SNAP-JD(m)

Input: $A, M, b, \varepsilon_{null}, k_{max}, m$

Output: $k, \ \|\varepsilon_k\|, \|b\|, x, w_k$

Initialisation:
    Construct annihilator E of b as $E_\perp$ or $E_/$ and set $\bar{A} = EAM^{-1}$

Generate initial null space vector approximation:
    Select random vector v_0, with $\|v_0\| = 1$, and set $v_1 = -\bar{A}v_0$. Approximately solve $\bar{A}t = v_1$ using GMRES(p), where p is the Analytic grade. Set $w_1 = \frac{t+v_0}{\|t+v_0\|}$ as the initial ANS vector
    Set null space residual vector $\varepsilon_1 = \bar{A}w_1$
    Set $x_1 = w_1$ and $X_1 = [x_1]$ then $\mathcal{X}_1 = \text{span}\{x_1\}$; and $k = 1$

Loop for $k = 1, \ldots,$ until $\|\varepsilon_k\| \le \varepsilon_{null}$ or $k = k_{max}$:

    Solve approximately for $t \perp w_k$ from $(I - w_k w_k^T)\bar{A}(I - w_k w_k^T)t = -(I - w_k w_k^T)\varepsilon_k$

    using GMRES(m) (where m equals, e.g., 5 or 10) or by monitoring $\cos\theta$ in (4.5)

    Compute $x_{k+1} = \frac{t - X_k X_k^T t}{\|t - X_k X_k^T t\|}$

    Expand subspace $\mathcal{X}_{k+1} = \mathcal{X}_k \oplus \text{span}\{x_{k+1}\}$ and set $X_{k+1} = [X_k \mid x_{k+1}]$

    Extend QR-decomposition of $\bar{A}X_{k+1} = Q_{k+1}R_{k+1}$

    Obtain the SVD factorization: $R_{k+1}V_{k+1} = U_{k+1}\Sigma_{k+1}$ with smallest singular triplet (σ_0, v_1, u_1)

    Set $w_{k+1} = X_{k+1}v_1$ and $\varepsilon_{k+1} = \sigma_0 Q_{k+1}u_1$, with $\|\varepsilon_{k+1}\| = \sigma_0$

End while Loop

Compute Final Approximate Solution x of linear system:
    $x = \beta w$, $w = M^{-1}w_k$, $\beta = \frac{\|b\|^2}{b^T Aw}$ with residual $\|r\| = |\beta|\|\varepsilon_k\|$

4.3 Restarted SNAP-JD

If the dimension of the subspace $\mathcal{X}_{k+1}$ becomes large and the approximation to the null space vector w is still inaccurate, then SNAP-JD can be restarted. One can perform a *simple restart procedure* as mentioned at the beginning of §4.1 and use a single vector to restart, however a better choice is to use *thick restart*, which is in the spirit of thick restart Krylov methods [6, 17].

This thick restart process uses a number of right singular vectors corresponding to the smallest ℓ singular values. According to Proposition 5, Corollary 2 (with $k = k_{max}$) let $[x_1, \ldots, x_\ell] = X_k^{(1)}V_\ell^{(1)} = W_k^{(1)}$ and $W_k^{(2)} = X_k^{(2)}V_\ell^{(2)}$, where $X_k^{(2)} = [W_k^{(1)}, x_{\ell+1}, \ldots, x_k]$ then $\sigma_\ell^{(2)} = \|\bar{A}W_k^{(2)}\| \le \|\bar{A}W_k^{(1)}\| = \sigma_\ell^{(1)}$. This result ensures that our restarted algorithm proposed below generates a monotone decreasing sequence of ANS.

This procedure is outlined in Algorithm 2 for SNAP-JD(m, k_{max}, ℓ), where m refers to the number of correction equation solves using GMRES, k_{max} is the maximum number

of subspace expansions before restarting and ℓ is the number of right singular vectors retained at the point of restart.

Algorithm 2 : Restarted SNAP-JD(m,k_{max},ℓ)

Input: $A, M, b, \varepsilon_{null}, k_{max}, max\ restarts$

Output: $k,\ \|\varepsilon_k\|, \|b\|, x, w_k$

Initialisation:
  Construct annihilator E of b as $E_\perp$ or $E_/$ and set $\bar{A} = EAM^{-1}$

Generate initial null space vector approximation:
  Select random vector v_0, with $\|v_0\| = 1$, and set $v_1 = -\bar{A}v_0$. Approximately solve $\bar{A}t = v_1$ using GMRES(p), where p is the analytic grade. Set $w_1 = \frac{t+v_0}{\|t+v_0\|}$ as the initial ANS vector
  Set null space residual vector $\varepsilon_1 = \bar{A}\,w_1$
  Set $x_1 = w_1$ then $\mathcal{X}_1 = \mathrm{span}\{x_1\}$; and $k = 1,\ restarts = 1$

Loop while $\|\varepsilon_k\| > \varepsilon_{null}$ and restarts $\leq$ max restarts

  Perform main loop of SNAP-JD(m) from k up to at most k_{max}

  Commence restart procedure if not converged
   Select ℓ smallest singular triplets of R_{k+1} and set $X_\ell = X_{k+1}V_\ell$ and note that $\bar{A}X_\ell = Q_\ell R_\ell$ with $Q_\ell = Q_{k+1}U_\ell,\ R_\ell = \Sigma_\ell$
   Set $w_{\ell+1} = X_{k+1}v_1$ and $\varepsilon_{\ell+1} = \sigma_0 Q_{k+1}u_1$, with $\|\varepsilon_{\ell+1}\| = \sigma_0$
  restarts $=$ restarts $+ 1,\ k = \ell + 1$

End restart while loop

Compute Final Solution
  Approximate solution to linear system is $x = \beta w,\ w = M^{-1}w_k,\ \beta = \frac{\|b\|^2}{b^T A w}$ with residual $\|r\| = |\beta|\,\|\varepsilon_k\|$

Note: The main computational expense for SNAP-JD(m) in Algorithm 1 manifests in performing the SVD of the upper triangular matrix R_{k+1}, which grows in size with the increasing dimension of the subspace $\mathcal{X}_{k+1}$. This expense is reduced by restarting the process in Algorithm 2 for SNAP-JD(m, k_{max}, ℓ). In this manner the SVD of an upper triangular matrix whose maximum dimension attains $k_{max} \times k_{max}$ is computed. The disadvantage of this strategy is that convergence is slowed over the Full SNAP-JD(m) algorithm.

4.4 What if A is singular?

The situation when A is singular and one wishes to obtain an approximate least-squares solution to the linear system $Ax = b$ by GMRES has been carefully studied in [2]. In particular, it is no longer true that break-down of GMRES implies that the underlying approximate solution has converged and is a least-squares solution, in contrast to when A is nonsingular. This is a situation when it is best to work with singular vectors rather than with eigenvectors.

When A is singular then $\mathcal{N}(A) \subset \mathcal{N}(\bar{A})$ and one of the advantages of using the SNAP method is that null vectors of A can be filtered from the ANS to provide the complete solution for $A\,x = b$.

Suppose we have generated a k-dimensional ANS as $N(\bar{A})$ with orthonormal basis given by $\{w_1, w_2, \ldots, w_k\}$ and set $W_k = [w_1, w_2, \ldots, w_k]$. Let the QR-factorisation of $AW_k = Y_k S_k$ and perform SVD of S_k to obtain $S_k V_k = U_k \Sigma_k$ with the singular values on the diagonal of Σ_k arranged in increasing order from top left to bottom right. The solution of $A\,x = b$ can be found using $x = W_k S_k^{\dagger} Y_k^T b$. Furthermore, if the matrices V_k, U_k and Σ_k are partitioned so that

$$S_k\,[V_1, V_2] = [U_1, U_2] \begin{bmatrix} 0 & 0 \\ 0 & \overline{\Sigma} \end{bmatrix}$$

then the null space vectors of A correspond to $W_k V_1$.

The Algorithms described above for SNAP-JD(m) and SNAP-JD(m, k_{max}, ℓ) can also be used to find more than one null space vector by using deflation techniques. This will be explored further in future research work. Finally we note that for singular A the preconditioning strategy given in §4.1 would not be possible and some careful reflection is necessary to uncover an effective preconditioner in this case. We leave this also for a future research topic.

5 Numerical Case Studies

In this section the theory proposed throughout §2-§4 is applied to study the solution of a cross-section of linear systems. To elucidate the issues discussed in the introduction, we give four specific examples where the standard Krylov subspace method (as formulated in the GMRES algorithm) is either not useful because its dimension m has to be very large to capture the solution, or the convergence of restarted GMRES stagnates.

5.1 Example 1

This example is constructed similarly to that proposed in [15] where $A = \text{diag}(0.01, J(1))$, $J(1)$ is an $n-1 \times n-1$ Jordan block and $b = (0, 0, \ldots, 0, 1)^T$. The Arnoldi process requires that $m = n - 1$ for convergence. Clearly this is unsatisfactory and some other solution method must be found. One observes (refer Figure 1) when comparing full GMRES and restarted GMRES(25) for $n = 300$ with initial guess $x_0 = 0$, that restarted GMRES offers very slow convergence towards the solution whilst full GMRES requires a subspace of dimension $m = 299$ to capture the solution.

Figure 1 exhibits the results for SNAP-JD for this example. In this case, due to the nature of the right-hand side vector $b = e_n$, $E_{\perp} = E_/$ and consequently, only the case
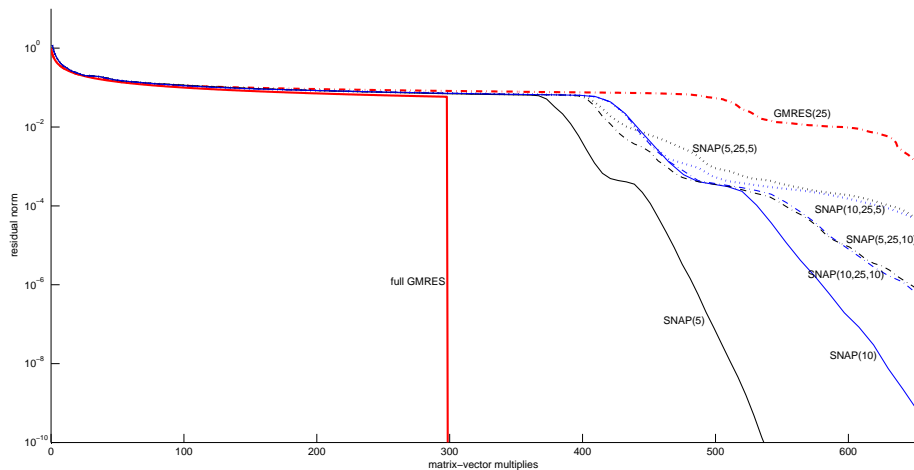
Figure 1: Comparison of full GMRES, restarted GMRES(25), SNAP-JD(5), SNAP-JD(5,25,5), SNAP-JD(5,25,10), SNAP-JD(10), SNAP-JD(10,25,5), SNAP-JD(10,25,10) for Example 1 with $E_\perp$

Table 1: Summary of statistics for the various SNAP-JD methods for Example 1

Method	Subspace Size (k)	Matrix multiplies	σ_0	$\|r\|$
SNAP-JD(5)	88	547	9.584e-013	1.657e-011
SNAP-JD(10)	61	685	3.033e-013	5.244e-012
SNAP-JD(5,25,5)	25	889	4.930e-009	8.525e-008
SNAP-JD(5,25,10)	25	709	8.180e-010	1.415e-008
SNAP-JD(10,25,5)	25	1081	9.639e-013	1.667e-011
SNAP-JD(10,25,10)	25	850	8.482e-013	1.467e-011

$E_\perp$ is presented. The figure shows a comparison of SNAP-JD(m) using $m = 5$ or 10, SNAP-JD(m,25,ℓ) with the number of right singular vectors retained at the point of restart taken as $\ell = 5$ or 10, full GMRES, and restarted GMRES(25). The complete statistics of the SNAP-JD runs are given in Table 1. One notes from the figure and the supporting table that SNAP-JD(5) is the best performer overall in terms of matrix multipies, followed by SNAP-JD(10). The convergence of the restarted methods is slower, but steady, and it appears that retaining ten of the smallest right singular vectors in the subspace performs better than retaining only five. Table 1 highlights that all SNAP-JD variants have converged, however restarting can be costly in terms of the total number of matrix-vector multiplies. The fact that SNAP-JD(5) converges in fewer matrix vector multiplies than SNAP-JD(10) suggests that oversolving the correction equation (4.3) may not offer any significant advantage. Interestingly, although SNAP-JD(5) uses fewer matrix vector multiplications in comparison with SNAP-JD(10), the latter results in a smaller

17

dimensional approximate null subspace on which the solution is found. Finally, it is worth pointing out that with restarting, the maximum dimension of the upper triangular matrix to which the SVD is applied is 25. This is undoubtedly less costly in terms of raw computational cost than working with the matrices of size 61 and 88 for the Full SNAP-JD methods.

5.2   Example 2

The linear system in this example is taken from [18] and concerns the matrix *sherman4* ($A \in \mathbb{R}^{1104 \times 1104}$) from the Harwell-Boeing Sparse Matrix Collection with right-hand side vector the one provided in the collection[1]. For this matrix all eigenvalues are positive and $\sigma(A) \subseteq (0.031, 66.5)$. Figure 3 highlights that standard GMRES performs well and restarted GMRES(25) offers a slow but steady convergence, whilst implicitly restarted GMRES-IR(25,6) and GMRES-IR(25,10) both converge considerably faster than GMRES(25), with a convergence rate similar to that of full GMRES.

The results for Example 2 are shown for SNAP-JD using $E_\perp$ in Figure 2-(a) and $E_/$ in Figure 2-(b), with the full run statistics summarised in Table 2. In this case all SNAP-JD variants perform similarly, converging in less than 200 matrix-vector multiplies. Furthermore, SNAP-JD offered convergence rates consistent with those exhibited by GMRES-IR in Figure 3. There doesn't appear to be any particular advantage of using the orthogonal annihilator $E_\perp$ over the oblique $E_/$.

5.3   Example 3

The linear system in this test concerns the matrix *bfw782a* ($A \in \mathbb{R}^{782 \times 782}$) also from the Matrix Market[1], with right-hand side vector $b = A \times ones(782, 1)$. For this matrix the eigenvalues satisfy $0.007 \le |\lambda_i| \le 11.021$, with 17 eigenvalues less than 0.1 in magnitude. The results for this system given in Figure 5 show that restarted GMRES(25) performs poorly, whilst implicitly restarted GMRES-IR(25,6) and GMRES-IR(25,10) both converge well.

The results for SNAP-JD are exhibited for Example 3 using $E_\perp$ in Figure 4-(a) and $E_/$ in Figure 4-(b). Again, SNAP-JD(m) with $m = 5$ or 10 and SNAP-JD(m,25,ℓ) with $\ell = 5$ or 10, are exhibited along with full GMRES and restarted GMRES(25). The complete statistics of the SNAP-JD runs are given in Table 3. Clearly all SNAP-JD variants are very competitive in comparison to Full GMRES and consistent with GMRES-IR (see Figure 5). The figures show that the orthogonal annihilator slightly out performs the oblique annihilator. Furthermore, retaining a set of the smallest right singular vectors at the point of restart and refining these as discussed for Method 2 in §2.2, offers convergence

---

[1]See the Matrix Market, `http://math.nist.gov/MatrixMarket`.
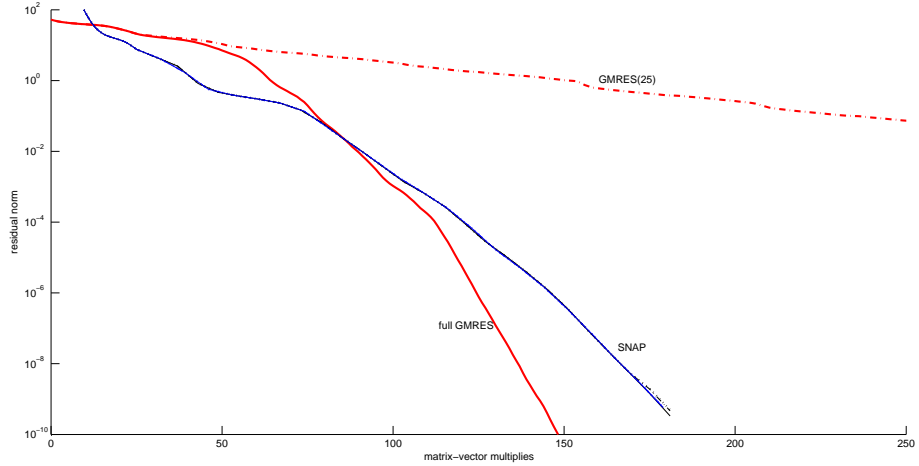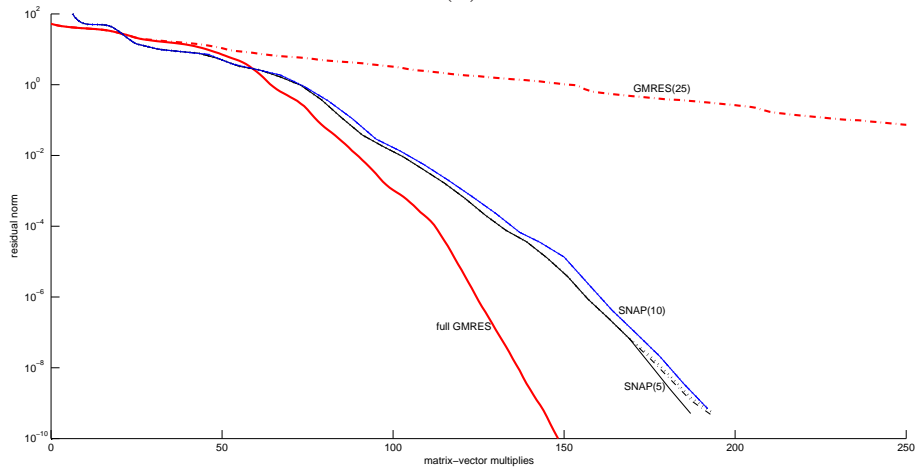
Figure 2: Comparison of full GMRES, restarted GMRES(25), SNAP-JD(5), SNAP-JD(5,25,5), SNAP-JD(5,25,10), SNAP-JD(10), SNAP-JD(10,25,5), SNAP-JD(10,25,10) for *sherman4* in Example 2 for (a) $E_\perp$ and (b) $E_/$
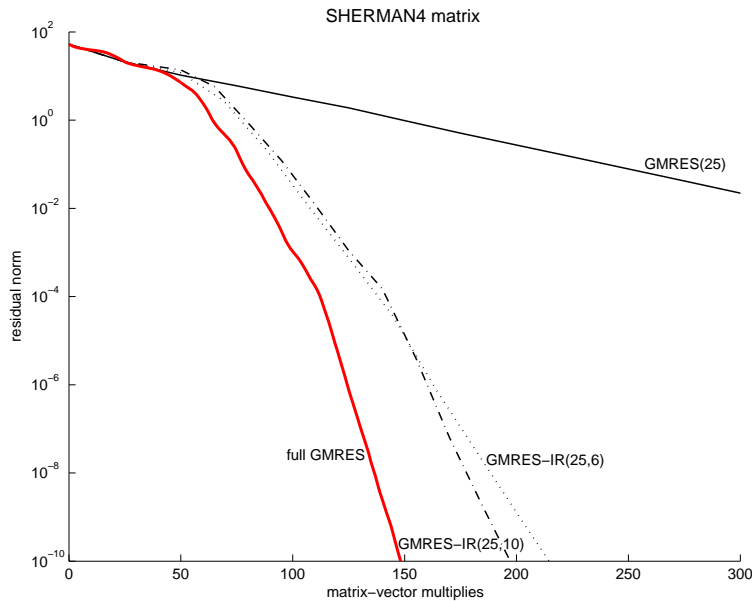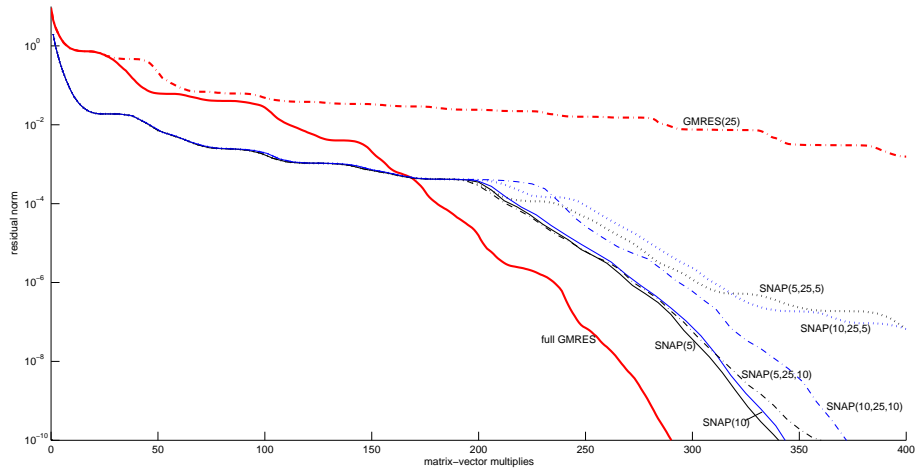
Figure 3: Comparison of full GMRES, restarted GMRES(25), GMRES-IR(25,6) and GMRES-IR(25,10) for the matrix *sherman4* in Example 2.

rates almost the same as full SNAP-JD.

5.4 Example 4

This example concerns the matrix *bcsstk13.mtx* ($A \in \mathbb{R}^{2003 \times 2003}$) taken from set *BC-SSTRUC1* in the Harwell-Boeing collection[1]. The matrix originates from a fluid flow generalised eigenvalue problem and has a condition number estimate of $4.6e10$ and $\|A\|_2 = 3.1e12$. A right-hand side vector b was generated from $b = A \times ones(2003, 1)$. For this matrix (refer Figure 7), ILU preconditioned (with threshold parameter 0.01) GMRES(30) and GMRES(40) could not overcome stagnation, whilst preconditioned Full GMRES required a Krylov subspace of dimension 200 to provide a good approximation to the solution. The results for different variants of SNAP-JD applied to Example 4 are depicted in Figures 6 (a-b) for $E_\perp$ and $E_/$ respectively. For this linear system only preconditioning (ILU with threshold parameter 0.01) is able to make SNAP-JD an effective iterative solution strategy. The figures highlight that SNAP-JD clearly overcomes the stagnation observed for GMRES(30) and GMRES(40). The restarted versions of SNAP-JD, although converging slower than SNAP(5), find the solution in a much smaller dimensional subspace. The pronounced humps in the residual curves evident in both figures can be attributed to the occurrence of a small right singular vector of A within the *ANS*, which causes $\|Aw\|$ to be quite small, hence resulting in an abrupt increase in the residual (refer §2.1

(a)
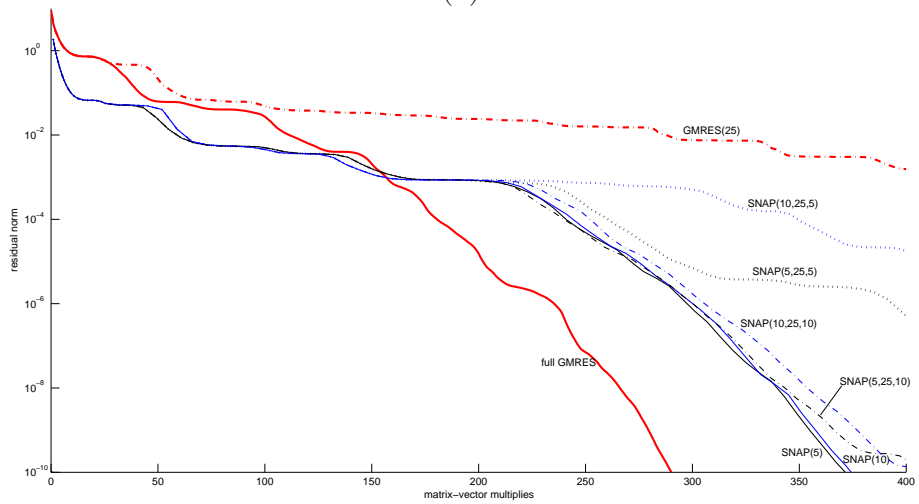
(b)

Figure 4: Comparison of full GMRES, restarted GMRES(25), SNAP-JD(5), SNAP-JD(5,25,5), SNAP-JD(5,25,10), SNAP-JD(10), SNAP-JD(10,25,5), SNAP-JD(10,25,10) for *bfw782a* in Example 3 for (a) $E_\perp$ and (b) $E_/$
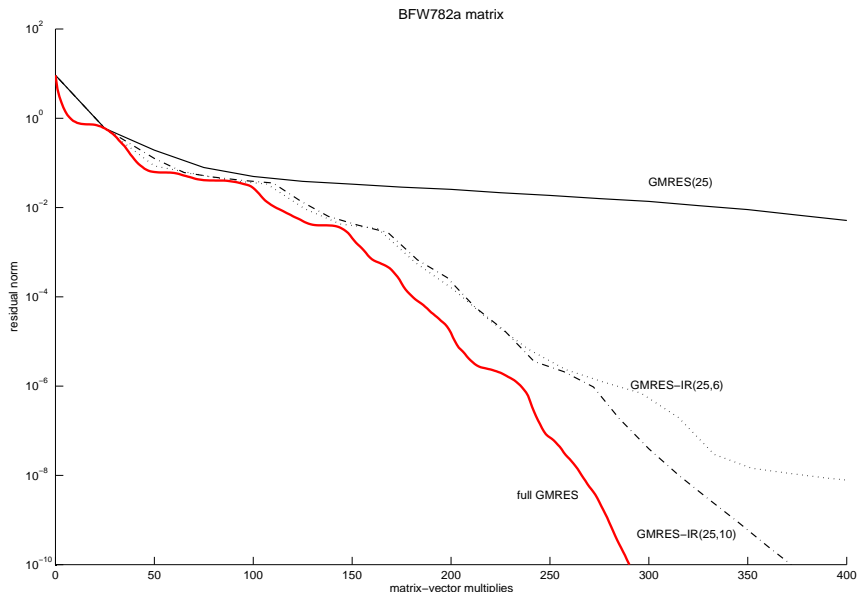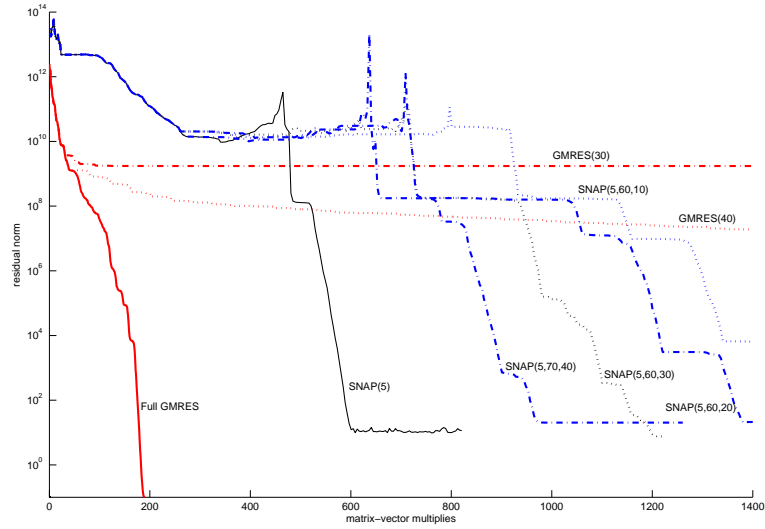
Figure 5: Comparison of full GMRES, restarted GMRES(25), GMRES-IR(25,6) and GMRES-IR(25,10) for the matrix *bfw782a* in Example 3.

for further details). Nevertheless SNAP-JD overcomes this problem, where one notes a steady convergence thereafter. Retaining more singular vectors in the subspace at the point of restart is definitely beneficial. It also appears for this example that the oblique annihilator is performing better than the orthogonal one.

In Figure 7 the performance of implicitly restarted GMRES for Example 4 using the same ILU preconditioning strategy is exhibited. Although GMRES-IR(30,10) offers some improvements over restarted GMRES(30) and GMRES(40) respectively, there is still some stagnation evident. GMRES-IR(40,20) and GMRES-IR(60,30) on the other hand both overcome stagnation, offering a converged solution in around 1100 and 400 matrix-vector multiplies respectively, which again shows a consistent trend with the SNAP-JD variants.

6   Conclusions

A new iterative method has been proposed as an alternative to implicitly restarted GMRES for overcoming the slow convergence, or stagnation, of restarted GMRES when small eigenvalues are present. By transforming the original non-homogeneous system into a homogeneous problem, SNAP-JD focusses the attention on the construction of an ap-

(a)
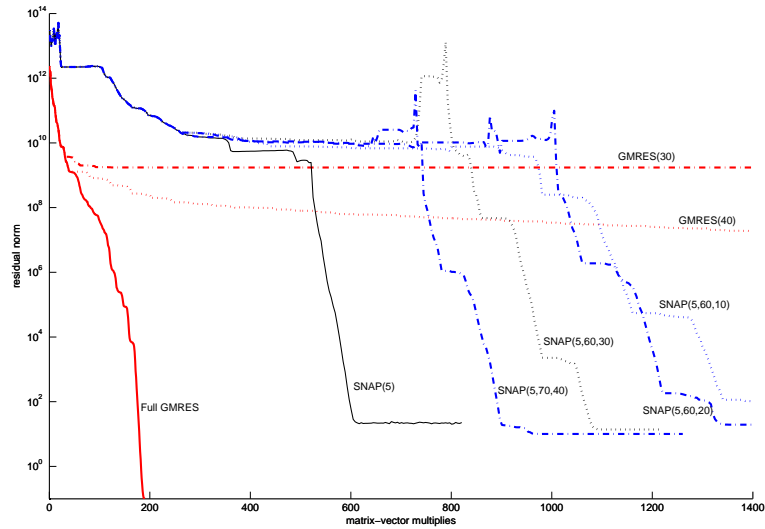
(b)

Figure 6: Comparison of full GMRES, restarted GMRES(30) and GMRES(40), SNAP-JD(5), SNAP-JD(5,60,10), SNAP-JD(5,60,20), SNAP-JD(5,60,30) and SNAP-JD(5,70,40) for *bcsstk13* in Example 4 for (a) $E_{\perp}$ and (b) $E_{/}$
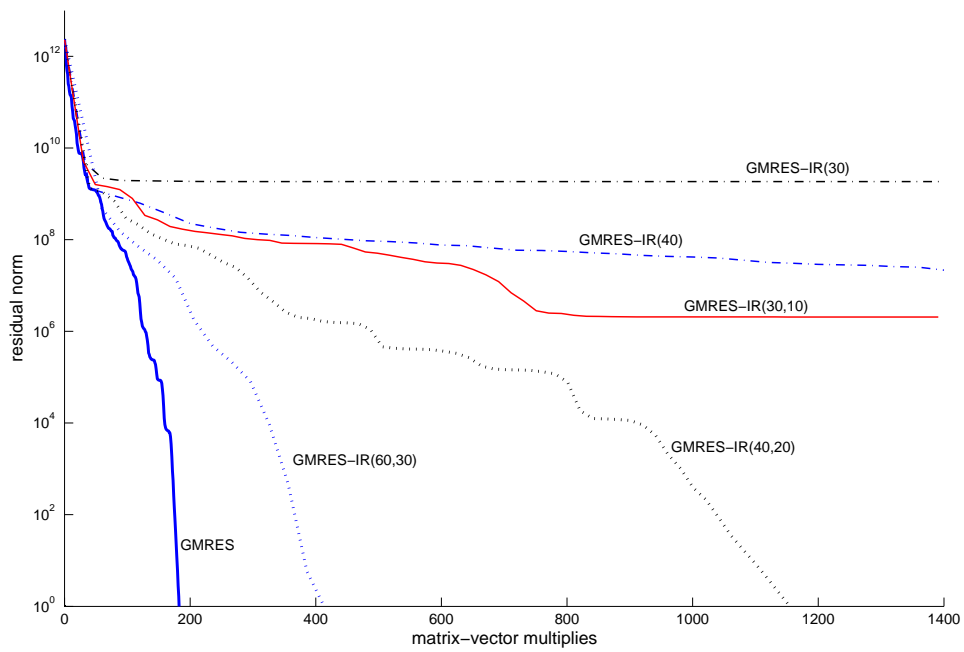
Figure 7: Comparison of full GMRES, restarted GMRES(30) and GMRES(40), GMRES-IR(30,10), GMRES-IR(40,20) and GMRES-IR(60,30) for *bcsstk13* in Example 4

Table 2: Summary of statistics for the various SNAP-JD methods for Example 2

Annihilator	Method	Subspace Size (k)	Matrix multiplies	σ_0	$\|r\|$
$E_\perp$	SNAP-JD(5)	27	181	2.872e-013	3.324e-010
	SNAP-JD(10)	23	179	4.975e-013	5.761e-010
	SNAP-JD(5,25,5)	25	181	4.228e-013	4.894e-010
	SNAP-JD(5,25,10)	25	181	3.994e-013	4.625e-010
	SNAP-JD(10,25,5)	25	179	4.975e-013	5.761e-010
	SNAP-JD(10,25,10)	25	179	4.975e-013	5.761e-010
$E_/$	SNAP-JD(5)	28	187	4.405e-013	5.096e-010
	SNAP-JD(10)	25	192	6.021e-013	6.927e-010
	SNAP-JD(5,25,5)	25	193	4.946e-013	5.454e-010
	SNAP-JD(5,25,10)	25	193	4.023e-013	3.484e-010
	SNAP-JD(10,25,5)	25	192	6.021e-013	6.927e-010
	SNAP-JD(10,25,10)	25	192	6.021e-013	6.927e-010

proximate null space (ANS) from which the solution can be filtered. A JD-analogous method called SNAP-JD is used to expand the ANS in such a way that a better approximation to the desired null space vector results at each expansion. From the results reported here, SNAP-JD appears to be a viable candidate for solving linear systems which cause difficulties to standard iterative solvers. Certainly from the graphical results exhibited in §5 it appears that SNAP-JD performs comparably with implicitly restarted GMRES.

The potential of the new method is that it can be applied to find the solution of rectangular systems and can provide the complete solution for singular systems, where the null space vectors of the coefficient matrix can be determined naturally as part of the overall solution strategy. Both of these topics will be the subject of future research work.

References

[1] J. Baglama, D. Calvetti, G.H.Golub and L. Reichel, (1998), Adaptively Preconditioned GMRES Algorithms, *Siam J. Sci. Comput.*, vol 20, 1, pp 243-269.

[2] P. N. BROWN AND H. F. WALKER, (1997), *Gmres on (nearly) singular systems*, SIAM Journal on Matrix Analysis and Applications, 18 , pp. 37–51.

Table 3: Summary of statistics for the various SNAP-JD methods for Example 3

Annihilator	Method	Subspace Size (k)	Matrix multiplies	σ_0	$\|r\|$
$E_\perp$	SNAP-JD(5)	56	350	7.655e-013	2.141e-011
	SNAP-JD(10)	51	353	6.345e-013	1.775e-011
	SNAP-JD(5,25,5)	25	644	7.625e-013	2.132e-011
	SNAP-JD(5,25,10)	25	392	5.007e-013	1.400e-011
	SNAP-JD(10,25,5)	25	609	9.861e-013	2.757e-011
	SNAP-JD(10,25,10)	25	390	7.518e-013	2.102e-011
$E_/$	SNAP-JD(5)	61	385	4.288e-013	1.194e-011
	SNAP-JD(10)	54	386	5.693e-013	1.593e-011
	SNAP-JD(5,25,5)	25	733	7.600e-013	2.124e-011
	SNAP-JD(5,25,10)	25	421	5.095e-013	1.421e-011
	SNAP-JD(10,25,5)	25	696	6.233e-013	1.736e-011
	SNAP-JD(10,25,10)	25	429	3.717e-013	1.039e-011

[3] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe and H. van der Vorst, editors, (2000) *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide.* SIAM, Philadelphia.

[4] M. Benzi and C. Meyer (1995), A Direct Projection Method for Sparse Linear Systems, SIAM J. Sci. Comput., 16, pp. 1159-1176.

[5] D. Calvetti, B. Lewis and L. Reichel (2000), GMRES-type methods for inconsistent systems,Linear Algebra Appl., 316, pp. 157-169.

[6] A. Chapman and Y.Saad, (1997), Deflated and Augmented Krylov Subspace Techniques, *Numer. Linear Algebra Appl.*, 4, pp. 43-66.

[7] G. H. GOLUB AND C. F. V. LOAN, (1996) *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 3rd ed.

[8] S. A. KHARCHENKO AND A. Y. YEREMIN, (1995), *Eigenvalue translation based preconditioners for the GMRES(k) method*, Numerical Linear Algebra with Applications, 2, pp. 51–70.

[9] E. de Sturler, (1999), Truncation Strategies for optimal Krylov subspace methods, *SIAM J. Matrix Anal. Appl.*, 36, pp.864-889.

[10] J. Erhel, K.Burrage and B.Pohl (1996), Restarted GMRES Preconditioned by Deflation, *J. Comput. Appl. Math*, 69, pp. 303-318.

[11] A. Greenbaum, V. Ptak and Z.Strakos, (1996), Any nonincreasing convergence curve is possible for GMRES. SIAM J.Matrix Anal. Appl., 17:465-469.

[12] M.E. Hochstenbach, (2001), A Jacobi-Davidson type SVD method, SIAM J.Sci. Comput. 23 (2), 606-628.

[13] M.Ilić and I.W.Turner, (2003), Approximately Invariant Subspaces, Proceedings of the Tenth Biennial Computation Techniques and Applications Conference, ANZIAM J., 44 (E), pp C378-C399.

[14] M.Ilić and I.W.Turner, (2005), Krylov Subspaces and the Analytic Grade, Numer. Linear Algebra Appl., 12, 55-76.

[15] I. C. F. Ipsen and C. D. Meyer, (1998) The Idea Behind Krylov Methods, American Mathematical Monthly, 105, 10, pp. 889-899.

[16] R. Lehoucq and D. Sorensen, (2000), Implicitly Restarted Lanczos Method (Section 4.5). In Z. Bai, J. Demmel, J. DOngarra, A. Ruhe and H. van der Vorst, editors, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide.* SIAM, Philadelphia.

[17] R. B. Morgan, (1995), A Restarted GMRES method Augmented with Eigenvectors, *Siam J. Matrix Anal. Appl.*, 16, pp. 1154-1171.

[18] R. B. Morgan, (2000), Implicitly Restarted GMRES and Arnoldi Methods for nonsymmetric systems of equations. SIAM J. Matrix Anal. Appl., 21:1112-1135.

[19] Y. Saad (2003), Iterative Methods for Sparse Linear Systems, second edition, SIAM, Philadelphia.

[20] Y. Saad and M. H. Schultz (1986), GMRES: A Generalized Minimum Residual Algorithm for solving Nonsymmetric Linear Systems, *SIAM J Scientific and Statistical Computing*, 7, pp. 856-869.

[21] G. L. G. Sleijpen and H.A. van der Vorst (1996), A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17, pp 401-425.

[22] G. W. Stewart, (1998), *Matrix Algorithms Volume I:Basic Decompositions*, SIAM, Philadelphia.

[23] G. W. Stewart, (2001), *Matrix Algorithms Volume II:Eigensystems*, SIAM, Philadelphia.

[24] H. Walker and P.N.Brown, (1997), GMRES on (nearly) singular systems, SIAM J. Matrix Anal. Appl., 18, pp. 37-51.

[25] A. van der Sluis and H. A. van der Vorst, (1986), *The rate of convergence of conjugate gradients*, 48, pp. 543–560.

[26] S. Van Huffel, (1991), Iterative algorithms for computing the singular subspace of a matrix associated with its smallest singular values, Linear Algebra Appl., 154/156, pp. 675-709.

[27] H. A. VAN DER VORST AND C. VUIK, (1993), *The superlinear convergence behaviour of GMRES*, Journal of Computational and Applied Mathematics, 48, pp. 327–341.