A JACOBI-DAVIDSON TYPE SVD METHOD

MICHIEL E. HOCHSTENBACH[1]

Abstract. We discuss a new method for the iterative computation of a portion of the singular values and vectors of a large sparse matrix. Similar to the Jacobi-Davidson method for the eigenvalue problem, we compute in each step a correction by (approximately) solving a correction equation. We give a few variants of this JDSVD method with their theoretical properties. It is shown that JDSVD can be seen as an accelerated (inexact) Newton scheme. We experimentally compare the method with some other iterative SVD methods.

Key words. Jacobi-Davidson, Singular Value Decomposition (SVD), singular values, singular vectors, norm, augmented matrix, correction equation, (inexact) accelerated Newton, improving singular values.

AMS subject classifications. 65F15 (65F35).

1. Introduction. Suppose that we want to compute one or more singular values, and the corresponding singular vectors, of the real $m \times n$ matrix A. This subject has already been studied from a number of different viewpoints [5, 6, 1, 19, 20, 13], for example to determine a few of the largest or smallest singular triples. This partial singular value decomposition can be computed in two different ways using equivalent eigenvalue decompositions.

The first is to compute some eigenvalues and eigenvectors of the $n \times n$ matrix A^TA or the $m \times m$ matrix AA^T. For large (sparse) matrices, direct methods like the QR-method are unattractive, but there exist several iterative methods. In [13], for example, (block) Lanczos [10] and Davidson [2] are applied to A^TA. Another candidate is Jacobi-Davidson [15]. Note that it is in general not advisable (nor necessary) to explicitly form the product A^TA. The nonzero eigenvalues of A^TA and AA^T are the squares of the nonzero singular values of A. This works positively for the separation of large singular values, but it forces a clustering of small ones. Moreover, it can be hard to find very small singular values (relative to the largest singular value) accurately. Apart from this, the approaches via A^TA or AA^T are asymmetric: in the process we approximate only one of the two singular vectors. The second vector can be obtained from the first by a multiplication by A or A^T, but this may introduce extra loss of accuracy. Besides, when we have approximations to both the left and right singular vector, we can use only one of them as starting vector for an iterative method.

A second approach is to compute some eigenvalues and eigenvectors of the *augmented matrix*

$$\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}. \tag{1.1}$$

This approach has its own advantages and disadvantages. The eigenvalues of the augmented matrix are plus and minus the singular values of A, and we can extract the left and right singular vectors from the eigenvectors by just taking the first and second part (see §2). This makes an extra multiplication by A or A^T unnecessary. We do not have the drawback of squaring small singular values. On the negative side, the augmented matrix is larger in size, and the smallest singular values are in the interior of the spectrum.

The Lanczos method for the augmented matrix has been studied by a number of authors [5, 6, 1]. The Lanczos process does not exploit the special (block or 'two-cyclic') structure of the matrix, unless the starting vector is of the form $(u, 0)$ or $(0, v)$. This is essentially Lanczos bidiagonalization of A, see [7, p. 495].

We can also consider the Jacobi-Davidson method [15] for the augmented matrix. This is an efficient method for the computation of a few eigenpairs, and it is of a different nature in comparison to Lanczos. The essence of Jacobi-Davidson is its correction equation, where the shifted operator

[1]Mathematical Institute, Utrecht University, P.O. Box 80.010, NL-3508 TA Utrecht, The Netherlands. `www.math.uu.nl/people/hochstenbach` Email: `hochstenbach@math.uu.nl`
January 2000. Revised December 2000.

is restricted to the subspace orthogonal to the current approximation to an eigenvector. When we solve this equation exactly, then we can show that the updated vector is the same as the one we would get by one step of Rayleigh Quotient Iteration (RQI). But in practice one solves the Jacobi-Davidson correction equation only approximately, and one accelerates the convergence by projecting the matrix onto the subspace spanned by all iterates. Therefore, Jacobi-Davidson can also be viewed as an inexact accelerated RQI.

'Standard' Jacobi-Davidson does not make use of the structure of the augmented matrix. In this paper we propose a Jacobi-Davidson variant that *does* take advantage of the special structure of the matrix. Instead of searching the eigenvector in one subspace, we search the left and right singular vectors in separate subspaces. We still solve a correction equation for the augmented matrix, but we use different projections, and we split the approximate solution of this equation for the expansion of the two search spaces. More similarities and differences are discussed in §7.3.

After some preparations in §2, we introduce the new approach, which we call JDSVD, in §3. In §4, a few variants of the algorithm with their properties are presented. In §5, we show that the JDSVD process can be viewed as an (inexact) accelerated Newton scheme and in §6 we focus on convergence. Various aspects of the method are discussed in §7, and after numerical examples in §8, we finish with conclusions in §9.

2. Preparations. Let A be a real $m \times n$ matrix with singular value decomposition $A = U_* \Sigma V_*^T$ and singular values

$$0 \leq \sigma_{\min} = \sigma_p \leq \sigma_{p-1} \leq \cdots \leq \sigma_2 \leq \sigma_1 = \sigma_{\max},$$

where $p := \min\{m, n\}$. Denote the corresponding left and right singular vectors by $u_{*,j}$ $(1 \leq j \leq m)$, and $v_{*,j}$ $(1 \leq j \leq n)$.

Throughout the paper, $\|\cdot\|$ stands for $\|\cdot\|_2$, we write $\sigma_j(B)$ for the j-th largest singular value of a real matrix B, and simply σ_j for the j-th largest singular value of A. Furthermore, $\Lambda(B)$ is the spectrum of B, and $\Sigma(B)$ is the set of singular values of B. If B is a real symmetric matrix, then $\lambda_j(B)$ denotes the j-th largest eigenvalue of B.

If $a \in \mathbb{R}^m$ and $b \in \mathbb{R}^n$, then, for convenience, we write $\begin{pmatrix} a \\ b \end{pmatrix} \in \mathbb{R}^{m+n}$ also as (a, b). If X is a matrix, then we denote the subspace spanned by the columns of X by $\mathcal{X}$. We use the notation $\mathcal{K}_l(B, x)$ for the Krylov subspace of dimension l generated by B and x.

Definition 2.1. *Let $u \in \mathbb{R}^m$, $v \in \mathbb{R}^n$, $\mathcal{X} \subset \mathbb{R}^m$, and $\mathcal{Y} \subset \mathbb{R}^n$. We say that $\begin{pmatrix} u \\ v \end{pmatrix} \in \mathbb{R}^{m+n}$ is double-orthogonal to the pair of subspaces $(\mathcal{X}, \mathcal{Y})$ if both $u \perp \mathcal{X}$ and $v \perp \mathcal{Y}$, which is denoted by $\begin{pmatrix} u \\ v \end{pmatrix} \perp\perp \begin{pmatrix} \mathcal{X} \\ \mathcal{Y} \end{pmatrix}$. By $(u, v)^{\perp\perp}$ we denote the subspace $\{(a, b) \in \mathbb{R}^m \times \mathbb{R}^n : u^T a = v^T b = 0\}$.*

The following lemma gives a relation between the singular triples of A and the eigenpairs of the augmented matrix.

Lemma 2.2. **(Jordan-Wielandt)** *[18, Th. I.4.2]*
The augmented matrix (1.1) *has eigenvalues*

$$-\sigma_1, \ldots, -\sigma_p, \underbrace{0, \ldots, 0}_{|m-n|}, \sigma_p, \ldots, \sigma_1$$

and eigenvectors

$$\begin{pmatrix} u_{*,j} \\ \pm v_{*,j} \end{pmatrix} \quad (1 \leq j \leq p)$$

corresponding to the $\pm\sigma_j$ and, if $m \neq n$, additionally,

$$\textit{either} \begin{pmatrix} u_{*,j} \\ 0 \end{pmatrix} \ (n+1 \leq j \leq m) \quad \textit{or} \quad \begin{pmatrix} 0 \\ v_{*,j} \end{pmatrix} \ (m+1 \leq j \leq n),$$

depending on whether $m > n$ or $n > m$.

The next definition is the natural analog of a simple eigenvalue (see e.g. [18, p. 15]).

DEFINITION 2.3. *We call σ_i a simple singular value of A if $\sigma_i \neq \sigma_j$ for all $j \neq i$.*

The following lemma gives a link between a simple singular value of A and a simple eigenvalue of A^TA and AA^T.

LEMMA 2.4. *Let $\sigma > 0$. Then σ is a simple singular value of A if and only if σ^2 is a simple eigenvalue of A^TA and AA^T.*

Proof. The nonzero eigenvalues of A^TA and AA^T are just the squares of the nonzero singular values of A (see for example [18, p. 31]). $\square$

Note that the condition $\sigma > 0$ in the previous lemma is necessary. For example, 0 is a simple singular value of the 1×2 matrix $A = (0\ 0)$ but it is not a simple eigenvalue of A^TA.
For future use, we mention the following well-known results.

LEMMA 2.5. **(Weyl)** *[21, pp. 101-102], [18, Cor. IV.4.9], [12, Th. 10.3.1]*
Let B and E be real symmetric $n \times n$ matrices. Then for all $1 \leq j \leq n$

$$\lambda_j(B) + \lambda_n(E) \leq \lambda_j(B+E) \leq \lambda_j(B) + \lambda_1(E).$$

LEMMA 2.6. *[9, (3.3.17)] If B and E are $m \times n$ matrices, then for $1 \leq i, j \leq p$ and $i+j \leq p+1$,*

$$\sigma_{i+j-1}(B+E) \leq \sigma_i(B) + \sigma_j(E).$$

In particular, for $j = 1$ this yields $\sigma_i(B+E) \leq \sigma_i(B) + \sigma_1(E)$, for $i = 1, \ldots, p$.

LEMMA 2.7. *[8, (7.3.8)] Let B and E be real $m \times n$ matrices. Then*

$$\sum_{j=1}^{p} (\sigma_j(B+E) - \sigma_j(B))^2 \leq \|E\|_F^2.$$

LEMMA 2.8. *If U and V are orthogonal $m \times m$ and $n \times n$-matrices, respectively, then for all $1 \leq j \leq p$ we have $\sigma_j(U^TAV) = \sigma_j(A)$. In particular $\|U^TAV\| = \|A\|$.*

Proof. The SVD of U^TAV is just $(U^TU_*)\Sigma(V^TV_*)^T$. The final statement follows from the characterization of the matrix two-norm as the largest singular value. $\square$

LEMMA 2.9. *[9, (3.1.3)] Let B be an $m \times n$ matrix, and let B_l denote a submatrix of B obtained by deleting a total of l rows and/or columns from B. Then*

$$\sigma_j(B) \geq \sigma_j(B_l) \geq \sigma_{j+l}(B)$$

for $1 \leq j \leq p$, where for a $q \times r$ matrix X we set $\sigma_j(X) = 0$ if $j > \min\{q, r\}$.

3. The JDSVD correction equation. Suppose that we have k-dimensional *search spaces* $\mathcal{U} \subset \mathbb{R}^m$ and $\mathcal{V} \subset \mathbb{R}^n$, and *test spaces* $\mathcal{X} \subset \mathbb{R}^m$ and $\mathcal{Y} \subset \mathbb{R}^n$. To determine approximations θ, η to a singular value, and $u \in \mathcal{U}, v \in \mathcal{V}$ (of unit norm) to the corresponding left and right singular vectors, we impose the *double Galerkin condition* with respect to $\mathcal{X}$ and $\mathcal{Y}$ on the *residual* r:

$$r = r(\theta, \eta) := \begin{pmatrix} Av - \theta u \\ A^Tu - \eta v \end{pmatrix} \perp\perp \begin{pmatrix} \mathcal{X} \\ \mathcal{Y} \end{pmatrix}. \tag{3.1}$$

Because $u \in \mathcal{U}$ and $v \in \mathcal{V}$, we can write $u = Uc$ and $v = Vd$, where the columns of the $m \times k$ matrix U and the columns of the $n \times k$ matrix V form bases for $\mathcal{U}$ and $\mathcal{V}$, and $c, d \in \mathbb{R}^k$. Then we want to find θ, η, c, and d that are solutions of

$$\left\{ \begin{array}{rcl} X^T A V d &=& \theta X^T U c \\ Y^T A^T U c &=& \eta Y^T V d, \end{array} \right. \tag{3.2}$$

where X and Y are matrices with columns that form bases for $\mathcal{X}$ and $\mathcal{Y}$. For *test vectors* $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, we have in particular that $r \perp\!\!\perp (x, y)$, so, if $x^T u \neq 0$ and $y^T v \neq 0$,

$$\theta = \frac{x^T A v}{x^T u}, \qquad \eta = \frac{y^T A^T u}{y^T v}. \tag{3.3}$$

This shows that the approximations θ and η may differ. We discuss possible choices for $\mathcal{X}$ and $\mathcal{Y}$, and the resulting relations for u and v in the following section. For now, suppose that we have approximations (u, v, θ, η). We would like to have a double-orthogonal correction $(s, t) \perp\!\!\perp (u, v)$ to (u, v) such that

$$\left\{ \begin{array}{rcl} A(v + t) &=& \sigma(u + s) \\ A^T(u + s) &=& \tau(v + t), \end{array} \right. \tag{3.4}$$

where $\sigma > 0$ and $\tau > 0$ need not to be equal, because the vectors are not normalized. However, since $A^T A(v + t) = \sigma\tau(v + t)$, we have $\sigma\tau = \sigma_i^2$ for some $1 \leq i \leq p$. Equations (3.4) can be rearranged to obtain

$$\begin{pmatrix} -\sigma I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -\begin{pmatrix} Av - \theta u \\ A^T u - \eta v \end{pmatrix} + \begin{pmatrix} (\sigma - \theta)u \\ (\tau - \eta)v \end{pmatrix} = -r + \begin{pmatrix} (\sigma - \theta)u \\ (\tau - \eta)v \end{pmatrix}.$$

Because σ and τ are unknown, we do not know the differences $(\sigma - \theta)u$ and $(\tau - \eta)v$ either. Therefore we can consider the projection of the last equation onto $(x, y)^{\perp\perp}$ along (u, v). This projection is given by

$$\begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix},$$

and it fixes r. Projecting the previous equation we get

$$\begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix} \begin{pmatrix} -\sigma I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r \qquad \text{where } (s, t) \perp\!\!\perp (u, v). \tag{3.5}$$

Since σ and τ are unknown, an obvious choice is to replace them by θ and η. This can be considered as 'throwing away second order terms' ($\sigma - \theta$, $\tau - \eta$, s and t will all be asymptotically small) and suggests that JDSVD is in fact a Newton method, which is true indeed (see §5). Furthermore, since for every $\widetilde{x} \in \mathbb{R}^m$ and $\widetilde{y} \in \mathbb{R}^n$ such that $u^T \widetilde{x} \neq 0$ and $v^T \widetilde{y} \neq 0$

$$\begin{pmatrix} I_m - \frac{\widetilde{x} u^T}{u^T \widetilde{x}} & 0 \\ 0 & I_n - \frac{\widetilde{y} v^T}{v^T \widetilde{y}} \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = \begin{pmatrix} s \\ t \end{pmatrix},$$

(3.5) leads to the following *JDSVD correction equation*:

$$\begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix} \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} I_m - \frac{\widetilde{x} u^T}{u^T \widetilde{x}} & 0 \\ 0 & I_n - \frac{\widetilde{y} v^T}{v^T \widetilde{y}} \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r \tag{3.6}$$

where $(s, t) \perp\!\!\perp (u, v)$. We see that the operator in (3.6) is symmetric if and only if $\widetilde{x}$ and $\widetilde{y}$ are a nonzero multiple of x and y. It maps $(u, v)^{\perp\perp}$ to $(x, y)^{\perp\perp}$. In §5 and §6 we explain why this process may lead to fast convergence, and we meet a generalized version of the correction equation. In the next section we examine several choices for the Galerkin conditions (3.1).

4. Choices for the Galerkin conditions. Consider the eigenvalue problem for a symmetric matrix B, where we have one subspace $\mathcal{W}$ that is used both as search space and test space. If the columns of W form an orthonormal basis for $\mathcal{W}$, then the projected matrix W^TBW has some nice properties, see [12, §11.4]. We will see that searching in two spaces, as in JDSVD, spreads those properties over a few Galerkin choices. In this section we examine some obvious choices.

4.1. The standard choice $\mathcal{X} = \mathcal{U}$ and $\mathcal{Y} = \mathcal{V}$. Consider the situation where the search spaces $\mathcal{U}$ and $\mathcal{V}$ are of equal dimension k. Let us first take the test spaces $\mathcal{X}$ and $\mathcal{Y}$ equal to the search spaces $\mathcal{U}$ and $\mathcal{V}$.

If the columns of U and V form orthonormal bases for $\mathcal{U}$ and $\mathcal{V}$, then with the notation $H := U^TAV$, (3.2) reduces to

$$Hd = \theta c \quad \text{and} \quad H^Tc = \eta d. \tag{4.1}$$

This gives approximations $u = Uc$ and $v = Vd$, where c and d are, respectively, left and right singular vectors of H. With the requirement $\|c\| = \|d\| = 1$ and test vectors $x = u$ and $y = v$, we get

$$\theta = \eta = u^TAv. \tag{4.2}$$

For reasons of symmetry, we choose $\widetilde{x} = x\ (= u)$ and $\widetilde{y} = y\ (= v)$. The resulting algorithm for the computation of $\sigma_{\max}$ is given in Algorithm 4.1.

Input: a device to compute Av and A^Tu for arbitrary u and v, starting vectors u_1 and v_1, and a tolerance ε.
Output: the approximate singular triple (θ, u, v) for the largest singular value $\sigma_{\max}$ and its corresponding singular vectors satisfying $\left\| \begin{pmatrix} Av - \theta u \\ A^Tu - \theta v \end{pmatrix} \right\| \le \varepsilon$.

1. $s = u_1,\ t = v_1,\ U_0 = [\ ],\ V_0 = [\ ]$
   for $k = 1, \ldots$
2. $U_k = \mathrm{RMGS}(U_{k-1}, s)$
   $V_k = \mathrm{RMGS}(V_{k-1}, t)$
3. Compute k-th column of $W_k = AV_k$
   Compute k-th row and column of $H_k = U_k^TAV_k = U_k^TW_k$
4. Compute largest singular triple (θ, c, d) of H_k, $(\|c\| = \|d\| = 1)$
   $u = U_kc,\ v = V_kd$
5. $r = \begin{pmatrix} Av - \theta u \\ A^Tu - \theta v \end{pmatrix} = \begin{pmatrix} W_kd - \theta u \\ A^Tu - \theta v \end{pmatrix}$
6. Stop if $\|r\| \le \varepsilon$
7. Solve (approximately) an $(s, t) \perp\perp (u, v)$ from
   $\begin{pmatrix} I_m - uu^T & 0 \\ 0 & I_n - vv^T \end{pmatrix} \begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix} \begin{pmatrix} I_m - uu^T & 0 \\ 0 & I_n - vv^T \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r$

Alg. 4.1: The JDSVD algorithm for the computation of $\sigma_{\max}(A)$, with Galerkin conditions $\mathcal{X} = \mathcal{U},\ \mathcal{Y} = \mathcal{V}$. The approximations (θ, η, u, v) satisfy $\theta = \eta = u^TAv$.

In step 2 of the algorithm, RMGS stands for repeated modified Gram-Schmidt (see for example [7, pp. 231-232]), used to make s and t orthogonal to U_{k-1} and V_{k-1}, and to expand the search spaces with the normalized vectors. Furthermore, $[\ ]$ stands for the empty matrix, and we omit the index k of all variables that are overwritten in every step. If we are interested in another singular value, for example the smallest, or the one closest to a specific target, we should adjust our choice in step 4 of the algorithm accordingly. The variant of Algorithm 4.1 is the only variant of JDSVD for which the operator in (3.6) is symmetric, and maps $(u, v)^{\perp\perp}$ in itself. Other choices imply that the operator is not symmetric, or maps $(u, v)^{\perp\perp}$ to a different space. See also §7.2.

4.2. Optimality of this choice. The following two results indicate that the method resulting from this standard Galerkin choice is optimal in some sense. Suppose we have an $m \times k$ matrix U and an $n \times k$ matrix V. Then for any $k \times k$ matrices K and L there are associated an $m \times k$ residual matrix $R(K)$, and an $n \times k$ residual matrix $\widetilde{R}(L)$:

$$R(K) := AV - UK \quad \text{and} \quad \widetilde{R}(L) := A^T U - VL.$$

If there exist K and L such that these residual matrices are zero, then we have found left and right singular subspaces, i.e. invariant subspaces of $A^T A$ and AA^T. The following theorem states that if both U and V have orthonormal columns, then $H := U^T AV$ and $H^T = V^T A^T U$ minimize the norm of these residual matrices, a desirable property. It is a generalization of a result in the theory for eigenproblems (see [12, Th. 11.4.2] and [18, Th. IV.1.15]), which deals with residuals of the form $AV - VK$.

THEOREM 4.1. *(Cf. [12, Th. 11.4.2])*
For given $m \times k$ matrix U and $n \times k$ matrix V, let $H = U^T AV$.
(a) If the columns of U are orthonormal, then for all $k \times k$ matrices K we have $\|R(H)\| \le \|R(K)\|$. Moreover, H is unique with respect to the Frobenius norm: $\|R(H)\|_F \le \|R(K)\|_F$ with equality only when $K = H$.
(b) If the columns of V are orthonormal, then $H^T = V^T A^T U$ minimizes the norm of $\widetilde{R}(L)$, and H^T is unique with respect to the Frobenius norm.

Proof. Suppose that the columns of U are orthonormal, then $U^T U = I$, so

$$\begin{aligned} R(K)^T R(K) &= V^T A^T AV + K^T K - K^T H - H^T K \\ &= V^T A^T AV - H^T H + (K-H)^T (K-H) \\ &= R(H)^T R(H) + (K-H)^T (K-H). \end{aligned}$$

Since $(K-H)^T(K-H)$ is positive semi-definite, it follows that

$$\|R(K)\|^2 = \lambda_{\max}(R(K)^T R(K)) \ge \lambda_{\max}(R(H)^T R(H)) = \|R(H)\|^2,$$

where we used Lemma 2.5 in the inequality. For uniqueness, we realize that $\|B\|_F^2 = \mathrm{Trace}(B^T B)$ for every real matrix B. Part (b) can be proved using the same methods. □

PROPOSITION 4.2. *Let u and v be approximations of unit norm. Then*

$$(\theta, \eta) = (u^T Av, u^T Av) \ \textit{minimizes}\ \|r(\theta, \eta)\|.$$

Proof. This can be shown by differentiating $\|r(\theta,\eta)\|^2$ with respect to θ and η. □

Because of these two results, it is a natural idea to take the k singular values $\theta_j^{(k)}$ of H_k as approximations to the singular values of A. When U_k and V_k have orthonormal columns, we see by Lemma 2.8 that these approximations converge in a finite number of steps to the singular values of A. In the following theorem we show that the approximations converge monotonically increasing.

THEOREM 4.3. *Let $\theta_k^{(k)} \le \cdots \le \theta_1^{(k)}$ be the singular values of $H_k := U_k^T AV_k$, where U_k and V_k have orthonormal columns. Then all $\theta_j^{(k)}$ converge monotonically increasing to the σ_j.*

Proof. H_k is a submatrix of H_{k+1}, so according to Lemma 2.9 $\theta_j^{(k+1)} \ge \theta_j^{(k)}$ for $1 \le j \le k$. Because of the orthogonality of U_k and V_k, the $\theta_j^{(k)}$ converge to the σ_j. □

REMARK 4.4. *In practice, one often observes that the $\theta_j^{(k)}$ converge strictly monotonically to the σ_j. With the aid of [21, pp. 94-98], conditions could be formulated under which the convergence is strict.*

Note that the theorem does *not* say that the smallest approximations $\theta_k^{(k)}$ converge monotonically (decreasing) to σ_p, because Lemma 2.9 only gives us $\theta_{k+1}^{(k+1)} \leq \theta_{k-1}^{(k)}$. For example, if $u_k \approx u_{*,p}$ and $v_k \approx v_{*,p-1}$, then $\theta_k^{(k)} \approx 0$, so we see that the smallest approximation can in fact be (much) smaller than σ_p. Experiments show that the convergence of the $\theta_k^{(k)}$ can be irregular and slow (see §8). This is a serious difficulty of working with the augmented matrix, because the smallest singular values are in the interior of its spectrum. We discuss this matter further in §4.3, §4.4, and §7.5. The following theorem gives some relations between the singular values of H_k and those of A. For clarity, we leave out the index k as much as possible.

THEOREM 4.5. *(Cf. [12, Ths. 11.5.1 and 11.5.2], [18, Cor. IV.4.15])*
For $j = 1, \ldots, k$, there exist singular values $\sigma_{j'}$ of A which can be put in one-one correspondence with the singular values θ_j of H in such a way that

$$|\sigma_{j'} - \theta_j| \leq \max\,\{\|R(H)\|, \|\widetilde{R}(H^T)\|\} \qquad (1 \leq j \leq k).$$

Moreover,

$$\sum_{j=1}^{k} (\sigma_{j'} - \theta_j)^2 \leq \|R(H)\|_F^2 + \|\widetilde{R}(H^T)\|_F^2.$$

Proof. Let the columns of $\widetilde{U}$ and $\widetilde{V}$ be orthonormal bases for the orthogonal complements of U and V, respectively. Then both $[U\ \widetilde{U}]$ and $[V\ \widetilde{V}]$ are orthogonal and

$$[U\ \widetilde{U}]^T A [V\ \widetilde{V}] = \begin{pmatrix} H & 0 \\ 0 & \widetilde{U}^T A \widetilde{V} \end{pmatrix} + \begin{pmatrix} 0 & U^T A \widetilde{V} \\ \widetilde{U}^T A V & 0 \end{pmatrix}. \tag{4.3}$$

Using Lemma 2.8 and Lemma 2.6, respectively, we obtain for $1 \leq j \leq p = \min\{m, n\}$

$$\sigma_j(A) = \sigma_j\left([U\ \widetilde{U}]^T A [V\ \widetilde{V}]\right) \leq \sigma_j \begin{pmatrix} H & 0 \\ 0 & \widetilde{U}^T A \widetilde{V} \end{pmatrix} + \sigma_{\max} \begin{pmatrix} 0 & U^T A \widetilde{V} \\ \widetilde{U}^T A V & 0 \end{pmatrix}.$$

Now

$$[U\ \widetilde{U}]^T R(H) = \begin{pmatrix} 0 \\ \widetilde{U}^T A V \end{pmatrix} \qquad \text{and} \qquad [V\ \widetilde{V}]^T \widetilde{R}(H^T) = \begin{pmatrix} 0 \\ \widetilde{V}^T A^T U \end{pmatrix},$$

so, because of the orthogonal invariance of the norm (see Lemma 2.8), $\|R(H)\| = \|\widetilde{U}^T A V\|$ and $\|\widetilde{R}(H^T)\| = \|\widetilde{V}^T A^T U\| = \|U^T A \widetilde{V}\|$. Because

$$\Sigma \begin{pmatrix} H & 0 \\ 0 & \widetilde{U}^T A \widetilde{V} \end{pmatrix} = \Sigma(H)\ \cup\ \Sigma(\widetilde{U}^T A \widetilde{V}),$$

there exist indices j' such that

$$\sigma_{j'} \begin{pmatrix} H & 0 \\ 0 & \widetilde{U}^T A \widetilde{V} \end{pmatrix} = \theta_j.$$

So the theorem's first inequality is obtained by

$$\sigma_{\max} \begin{pmatrix} 0 & U^T A \widetilde{V} \\ \widetilde{U}^T A V & 0 \end{pmatrix} = \max\,\{\|\widetilde{U}^T A V\|, \|U^T A \widetilde{V}\|\} = \max\,\{\|R(H)\|, \|\widetilde{R}(H^T)\|\}.$$

For the second inequality, apply Lemma 2.7 to the splitting of (4.3). □

For the following proposition, we need the minimax theorem for singular values [9, Th. 3.1.2]

$$\sigma_j = \max_{\mathcal{X}^j \subset \mathbb{R}^n} \min_{0 \neq x \in \mathcal{X}} \frac{\|Ax\|}{\|x\|}, \tag{4.4}$$

where $\mathcal{X}^j$ ranges over all subspaces of $\mathbb{R}^n$ of dimension j.

This proposition states that the singular values of $U_k^T A V_k$ are also *not* optimal in another sense:

PROPOSITION 4.6. *Let U_k and V_k have orthonormal columns. For $1 \le j \le k$,*

$$\sigma_j(U_k^T A V_k) \le \sigma_j(A V_k) \qquad \textit{and} \qquad \sigma_j(U_k^T A V_k) \le \sigma_j(A^T U_k).$$

Proof. This follows from (4.4) and the inequalities $\|U_k^T A V_k y\| \le \|A V_k y\|$ and $\|V_k^T A^T U_k x\| \le \|A^T U_k x\|$. □

We have seen that the $\sigma_j(U_k^T A V_k)$ increase monotonically, and that they are bounded above by both $\sigma_j(AV_k) = \lambda_j^{1/2}(V_k^T A^T A V_k)$ and $\sigma_j(A^T U_k) = \lambda_j^{1/2}(U_k^T A A^T U_k)$. This forms one motivation to study other Galerkin choices. A second is the possibly irregular convergence of the smallest singular value of $U_k^T A V_k$.

4.3. Other choices. Suppose that the columns of V form an orthonormal basis for $\mathcal{V}$. By the Galerkin choice $\mathcal{X} = A\mathcal{V}$, $\mathcal{Y} = \mathcal{V}$, with test vectors $x = Av$, $y = v$, and $u = Uc$, $v = Vd$, and $\|v\| = 1$, (3.2) reduces to

$$\left\{ \begin{array}{lcl} V^T A^T A V d & = & \theta\, V^T A^T U c \\ V^T A^T U c & = & \eta\, d \end{array} \right. \tag{4.5}$$

One can check that to satisfy the Galerkin conditions, $(\theta\eta, d)$ should be an eigenpair of $V^T A^T A V$. Now first suppose that $V^T A^T U$ is nonsingular. Note that in this case $\eta \neq 0$, otherwise $V^T A^T U$ would be singular. It follows that $c = \eta\,(V^T A^T U)^{-1} d$, $\eta = v^T A^T u$, and $\theta = v^T A^T A v / v^T A u$. When $V^T A^T U$ is singular, then this construction is impossible, but in this case we can simply restart the process or add extra vectors to the search spaces (see §7.6).

With this Galerkin choice, θ and η do not converge monotonically in general, but we can apply well-known results from eigenvalue theory to assure that their product does converge monotonically to the squares of the singular values, also to the smallest. In §7.2 we discuss the resulting correction equation.

Likewise, if the columns of U form an orthonormal basis for $\mathcal{U}$, the Galerkin choice $\mathcal{X} = \mathcal{U}$, $\mathcal{Y} = A^T\mathcal{U}$ leads to the determination of $(\theta\eta, c)$, an eigenpair of $U^T A A^T U$. These two approaches are natural with respect to minimax considerations, as we will see now.

LEMMA 4.7. *Let $\xi \in [0,1]$. Then we have the following minimax property for singular values:*

$$\sigma_j = \max_{\substack{\mathcal{S}^j \subset \mathbb{R}^m \\ \mathcal{T}^j \subset \mathbb{R}^n}} \min_{\substack{0 \neq s \in \mathcal{S}^j \\ 0 \neq t \in \mathcal{T}^j}} \xi \frac{\|At\|}{\|t\|} + (1-\xi) \frac{\|A^T s\|}{\|s\|} \qquad (1 \le j \le p). \tag{4.6}$$

Proof. This follows from (4.4) and the observation that A and A^T have the same singular values. □

When we have search spaces $\mathcal{U}$ and $\mathcal{V}$, it is a natural idea to substitute $\mathcal{U}$ for $\mathbb{R}^m$ and $\mathcal{V}$ for $\mathbb{R}^n$ in (4.6), as a generalization of a similar idea in the theory of eigenproblems, see [12, p. 236]. This gives the following approximations for the singular values:

$$\tau_j = \max_{\substack{\mathcal{S}^j \subset \mathcal{U} \\ \mathcal{T}^j \subset \mathcal{V}}} \min_{\substack{0 \neq s \in \mathcal{S}^j \\ 0 \neq t \in \mathcal{T}^j}} \xi \frac{\|At\|}{\|t\|} + (1-\xi) \frac{\|A^T s\|}{\|s\|}. \tag{4.7}$$

The following theorem relates these approximations to the Ritz values of A^TA and AA^T:

THEOREM 4.8. $\tau_j = \xi\,(\lambda_j^{1/2}(V^TA^TAV)) + (1-\xi)\,(\lambda_j^{1/2}(U^TAA^TU))$.

Proof. We have that $\mathcal{T}^j \subset \mathcal{V}$ if and only if $\mathcal{T}^j = V\widetilde{\mathcal{T}}^j := \{Vt : t \in \widetilde{\mathcal{T}}^j\}$ and $\widetilde{\mathcal{T}}^j \subset \mathbb{R}^k$. So for the first term of the expression for the τ_j we have that

$$\max_{\mathcal{T}^j\subset\mathcal{V}}\ \min_{0\neq t\in\mathcal{T}^j} \frac{\|At\|^2}{\|t\|^2} = \max_{\widetilde{\mathcal{T}}^j\subset\mathbb{R}^k}\ \min_{0\neq t\in\widetilde{\mathcal{T}}^j} \frac{t^TV^TA^TAVt}{\|t\|^2} = \lambda_j(V^TA^TAV).$$

For the second term we have a similar expression. □

When we take $\xi = 0$ and $\xi = 1$ in Theorem 4.8, we recognize the Galerkin approaches described in (4.5) and the discussion after that. They can essentially be viewed as a two-sided approach to A^TA or AA^T, in the sense that we have approximations to both the left and the right singular vector during the process.

As a generalization, we can consider the test spaces $\mathcal{X}$ and $\mathcal{Y}$ with bases $\alpha u_i + \beta Av_i$ and $\gamma v_i + \delta A^Tu_i$ respectively, where $\alpha^2+\beta^2 = \gamma^2+\delta^2 = 1$. Every choice other than $\alpha = \gamma = 1$ (the standard Galerkin choice discussed in §4.1) involves the computation of additional projected matrices, and more work per iteration.

Another possibility is to take search spaces of unequal dimension, that is, $\mathcal{U}_k$ and $\mathcal{V}_l$ where $k \neq l$. However, in view of the symmetric role of $\mathcal{S}_j$ and $\mathcal{T}_j$ in (4.6), this is probably not very useful.

4.4. Harmonic singular triples. As observed in §4.2, the standard Galerkin choice leads to monotone convergence for the largest singular value, but it can imply irregular behavior for the smallest singular value. A related problem is the selection of the best approximate vectors. Suppose that $u = \sum_{j=1}^m \gamma_j u_{*,j}$ and $v = \sum_{j=1}^n \delta_j v_{*,j}$ are approximate vectors, then $\theta = u^TAv = \sum_{j=1}^p \gamma_j\delta_j\sigma_j$ (we may assume θ non-negative, otherwise take $-u$ instead of u). Now suppose that $\theta \approx \sigma_1$, in the sense that $\sigma_2 < \theta < \sigma_1$ and that $\sigma_1 - \theta$ is (much) smaller than $\theta - \sigma_2$. Then we conclude that $\gamma_1 \approx 1$ and $\delta_1 \approx 1$, so u and v are good approximations to $u_{*,1}$ and $v_{*,1}$. But when $\theta \approx \sigma_p$, then u and v are not necessarily good approximations to $u_{*,p}$ and $v_{*,p}$. For example, u could have a large component of $u_{*,p-1}$ and a small component of $u_{*,1}$, and v a large component of $v_{*,p-2}$ and a small component of $v_{*,1}$. In conclusion, when we search for the largest singular value, it is asymptotically safe to select the largest singular triple of H, but for the smallest singular value the situation is more subtle.

Suppose for the moment that A is square and invertible. If the minimal singular value is the one of interest, the above suggests to study the singular values of A^{-1}. Based on the SVD of A^{-1}:

$$A^{-1} = V_*\Sigma^{-1}U_*^T,$$

we try to find the largest singular values of A^{-1} with respect to certain search spaces $\widehat{\mathcal{U}}, \widehat{\mathcal{V}}$ and test spaces $\widehat{\mathcal{X}}, \widehat{\mathcal{Y}}$. The new Galerkin conditions become (cf. (3.1)):

$$\begin{pmatrix} A^{-1}\widehat{u} - \widehat{\theta}\widehat{v} \\ A^{-T}\widehat{v} - \widehat{\eta}\widehat{u} \end{pmatrix} \perp\!\perp \begin{pmatrix} \widehat{\mathcal{Y}} \\ \widehat{\mathcal{X}} \end{pmatrix}.$$

where $\widehat{u} \in \widehat{\mathcal{U}}$ and $\widehat{v} \in \widehat{\mathcal{V}}$, say $\widehat{u} = \widehat{U}\widehat{c}$ and $\widehat{v} = \widehat{V}\widehat{d}$. To avoid having to work with A^{-1}, we take for the search spaces $\widehat{\mathcal{U}} = A\mathcal{V}$ and $\widehat{\mathcal{V}} = A^T\mathcal{U}$ (cf. [15]). This gives the system

$$\left\{ \begin{array}{rcl} \widehat{Y}^TV\,\widehat{c} &=& \widehat{\theta}\,\widehat{Y}^TA^TU\widehat{d} \\ \widehat{X}^TU\,\widehat{d} &=& \widehat{\eta}\,\widehat{X}^TAV\widehat{c} \end{array} \right. .$$

Taking $\widehat{\mathcal{X}} = \mathcal{U}$ and $\widehat{\mathcal{Y}} = \mathcal{V}$ results in equivalent conditions as in the standard choice (4.1), only now $(\widehat{\eta}, \widehat{\theta})$ and $(\widehat{c}, \widehat{d})$ play the role of (θ^{-1}, η^{-1}) and (d, c). The choices $(\widehat{\mathcal{X}}, \widehat{\mathcal{Y}}) = (A\mathcal{V}, \mathcal{V})$, $(\mathcal{U}, A^T\mathcal{U})$, and $(A\mathcal{V}, A^T\mathcal{U})$ lead to different approximations: to (4.5) and other systems described in §4.3, only the roles of the variables differ. We can call these approximations *harmonic singular triples*, in analogy to the harmonic Ritz pairs in the eigenproblem [11]. So these harmonic approximations have two advantages: the monotone behavior of the approximations to the smallest singular value, and the the selection of a good approximate 'smallest' vector.

The conclusion is that the non-default Galerkin choices, as presented in §4.3, can also be seen as a 'harmonic' approach to the problem. Finally, when A is singular or even nonsquare, we can still consider A^+ with respect to the test and search spaces $A\mathcal{V}$ and $A^T\mathcal{U}$, exploiting the fact that $AA^+A = A$.

5. JDSVD as an (inexact) accelerated Newton scheme. In [16], it is shown that the Jacobi-Davidson method can be interpreted as an inexact accelerated Newton scheme [4] for the eigenvalue problem. Here we show that the same is true for JDSVD applied to the singular value problem. Define $F : \mathbb{R}^m \times \mathbb{R}^n \to \mathbb{R}^m \times \mathbb{R}^n$ as

$$F(u, v) := \begin{pmatrix} Av - \theta u \\ A^T u - \eta v \end{pmatrix},$$

where $\theta = \theta(u, v)$ and $\eta = \eta(u, v)$ are as in (3.3). Thus the function F is nonlinear. Consider the singular value problem where we require the singular vectors u_*, v_* to be scaled such that $u_*^T a = 1$ and $v_*^T b = 1$, for certain vectors $a \in \mathbb{R}^m$ and $b \in \mathbb{R}^n$. So we look for solutions u_*, v_* of the equation $F(u, v) = 0$ in the 'hyperplane'

$$\left\{(u, v) \in \mathbb{R}^m \times \mathbb{R}^n : u^T a = 1,\ v^T b = 1\right\}.$$

We introduce these a and b to derive a more general form of the correction equation (3.6). If (u_k, v_k) are approximations to the singular vectors, then the next Newton approximations (u_{k+1}, v_{k+1}) are given by $(u_{k+1}, v_{k+1}) = (u_k, v_k) + (s_k, t_k)$, where $(s_k, t_k) \perp\perp (a, b)$ satisfies

$$DF(u_k, v_k)(s_k, t_k) = -F(u_k, v_k) = -r_k.$$

Omitting the index k, one may check (remembering that $\theta = \theta(u, v)$ and $\eta = \eta(u, v)$ are as in (3.3)) that the Jacobian $DF(u, v)$ of F is given by

$$DF(u, v) = \begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix} \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix}.$$

Hence, the correction equation of the Newton step is given by

$$\begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix} \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r \qquad \text{where } (s, t) \perp\perp (a, b).$$

For every $\widetilde{x}, \widetilde{y}$ so that $a^T \widetilde{x} \neq 0$ and $b^T \widetilde{y} \neq 0$, this is equivalent to the slightly more general form of the JDSVD correction equation (in comparison with (3.6)):

$$\begin{pmatrix} I_m - \frac{ux^T}{x^T u} & 0 \\ 0 & I_n - \frac{vy^T}{y^T v} \end{pmatrix} \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} I_m - \frac{\widetilde{x}a^T}{a^T \widetilde{x}} & 0 \\ 0 & I_n - \frac{\widetilde{y}b^T}{b^T \widetilde{y}} \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r, \tag{5.1}$$

where $(s, t) \perp\perp (a, b)$. Note that the substitution $a = u$ and $b = v$ gives (3.6).

If we keep $a, b, x, \widetilde{x}, y$, and $\widetilde{y}$ fixed during the process, and if $x^T u_*, y^T v_*, a^T \widetilde{x}$, and $b^T \widetilde{y}$ are nonzero, then Newton iteration produces a series (u_k, v_k) that converges asymptotically quadratically towards (u_*, v_*) if the starting vector (u_1, v_1) is sufficiently close to (u_*, v_*).

But if we take $a, b, x, \widetilde{x}, y$, and $\widetilde{y}$ variable, but converging to certain vectors, such that the denominators in (5.1) do not vanish, we get asymptotically quadratic convergence as well. The choice $a = x = \widetilde{x} = u_k$ and $b = y = \widetilde{y} = v_k$ leads to Algorithm 4.1. With other Galerkin choices described in §4, the test vectors (x, y) are in general not equal to the approximations (u, v), and in this situation the vectors a and b can be useful, see §6 and §7.2.

We see that JDSVD is a Newton scheme, accelerated by the usage of all previous iterates and the projection of A on the subspace that they span. This *subspace acceleration* accelerates the 'pre-quadratic' phase of the method and ensures that we find a singular triple in a finite number of steps. It may be expensive to solve the correction equation exactly. Instead we may solve (5.1) approximately (see §7.1), the resulting method is an inexact accelerated Newton scheme.

In [14], it is proved that if the correction equation is solved exactly, then Jacobi-Davidson applied to a symmetric matrix has asymptotically cubic convergence. Because the augmented matrix (1.1) is symmetric, we expect that JDSVD can also reach cubic convergence. The next section shows that this expectation is correct indeed.

6. Convergence. In the previous section we have already seen that the correction equation represents a Jacobian system in a Newton step. Now we focus on the convergence (see [14] for similar observations for Jacobi-Davidson applied to the eigenvalue problem).

In a mathematical sense, it is not meaningful to speak of asymptotic convergence, because we know that the process ends in a finite number of steps. However, in many practical situations a singular triple will be approximated well, long before the dimension of the search spaces reaches p. At that stage, these approximations display a behavior that looks like a converging infinite sequence close to its limit. When speaking of asymptotic convergence, we think of this situation. In other words, by the word 'asymptotically' we mean the situation where we have a (very) good approximation to the singular triple, rather than the situation where $k \to \infty$.

In the correction equation (5.1), u and v are the current approximations, x and y are test vectors, but we have not said much about choosing $\widetilde{x}$, $\widetilde{y}$, a, and b. They can vary per step. The next lemma and proposition show that exact JDSVD (that is, JDSVD where we solve the correction equation exactly) has asymptotically cubic convergence, for specific choices of the test vectors x and y, and the vectors a and b. To be precise, with ε small enough, if

$$|\angle(u_k, u_*)| = \mathcal{O}(\varepsilon) \text{ and } |\angle(v_k, v_*)| = \mathcal{O}(\varepsilon) \tag{6.1}$$

and if

$$a = x \text{ and } b = y \qquad \text{and} \qquad |\angle(x, u_*)| = \mathcal{O}(\varepsilon) \text{ and } |\angle(y, v_*)| = \mathcal{O}(\varepsilon) \tag{6.2}$$

then $|\angle(u_{k+1}, u_*)| = \mathcal{O}(\varepsilon^3)$ and $|\angle(v_{k+1}, v_*)| = \mathcal{O}(\varepsilon^3)$. Then the approximate singular values (see (3.3)) converge cubically as well.

LEMMA 6.1. *(Cf. [14, Lemma 3.1]) Assume that $Av_* = \sigma u_*$ and $A^T u_* = \tau v_*$, where $\sigma, \tau > 0$, and that $\sqrt{\sigma\tau}$ is a simple singular value of A. Let a, b, x, $\widetilde{x}$, y, and $\widetilde{y}$ be such that $x^T u_*$, $y^T v_*$, $a^T \widetilde{x}$, $b^T \widetilde{y}$, $a^T u_*$, and $b^T v_*$ are all nonzero. Then the map*

$$G := \begin{pmatrix} I_m - \frac{u_* x^T}{x^T u_*} & 0 \\ 0 & I_n - \frac{v_* y^T}{y^T v_*} \end{pmatrix} \begin{pmatrix} -\sigma I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} I_m - \frac{\widetilde{x} a^T}{a^T \widetilde{x}} & 0 \\ 0 & I_n - \frac{\widetilde{y} b^T}{b^T \widetilde{y}} \end{pmatrix}$$

is a bijection from $(a, b)^{\perp\perp}$ onto $(x, y)^{\perp\perp}$.

Proof. Suppose $(z_1, z_2) \perp\perp (a, b)$ and $G(z_1, z_2) = 0$. We show that $z_1 = z_2 = 0$. We have

$$\begin{pmatrix} -\sigma I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \begin{pmatrix} \mu u_* \\ \nu v_* \end{pmatrix},$$

for certain μ, ν. Then

$$\left\{ \begin{array}{lcl} A z_2 & = & \sigma z_1 + \mu u_* \\ A^T z_1 & = & \tau z_2 + \nu v_*. \end{array} \right.$$

Multiplying the first equation by A^T and the second by A we find

$$\left\{ \begin{array}{lcl} (A^T A - \sigma\tau) z_2 & = & (\sigma\nu + \tau\mu) v_* \\ (A A^T - \sigma\tau) z_1 & = & (\sigma\nu + \tau\mu) u_*. \end{array} \right.$$

So both z_1 and u_* belong to the kernel of $(AA^T - \sigma\tau)^2$, and both z_2 and v_* belong to the kernel of $(A^TA - \sigma\tau)^2$. From the simplicity of $\sigma\tau$ using Lemma 2.4, we have that z_1 and z_2 are multiples of u_* and v_* respectively. Because $z_1 \perp a$, $z_2 \perp b$, and $a^T u_* \neq 0$, $b^T v_* \neq 0$ we conclude $z_1 = z_2 = 0$. The bijectivity follows from comparing dimensions. $\square$

PROPOSITION 6.2. *(Cf. [14, Theorem 3.2]) With the assumptions of Lemma* 6.1*, if the initial vectors are close enough to the singular vectors corresponding to a simple nonzero singular value (i.e. if* (6.1) *holds), and if the correction equation is solved exactly, then for fixed vectors x, y, a and b, the JDSVD process has quadratic convergence. Moreover, if* (6.2) *holds then JDSVD has even cubic convergence.*

Proof. For convenience write

$$P = \begin{pmatrix} I_m - \frac{ux^T}{x^Tu} & 0 \\ 0 & I_n - \frac{vy^T}{y^Tv} \end{pmatrix}, \quad B = \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix}, \quad Q = \begin{pmatrix} I_m - \frac{\widetilde{x}a^T}{a^T\widetilde{x}} & 0 \\ 0 & I_n - \frac{\widetilde{y}b^T}{b^T\widetilde{y}} \end{pmatrix}.$$

Then the correction equation (5.1) reads, for $(s,t) \perp\!\!\perp (a,b)$:

$$PBQ(s,t) = PB(s,t) = -r = -B(u,v).$$

Suppose that $\widetilde{u}$ and $\widetilde{v}$ are scalar multiples of the singular vectors u_* and v_* and that $(\widetilde{u},\widetilde{v}) = (u,v) + (e,f)$, where $(e,f) \perp\!\!\perp (a,b)$, and $\|e\| = \mathcal{O}(\varepsilon)$, $\|f\| = \mathcal{O}(\varepsilon)$. Our first goal is to show that $\|(e-s, f-t)\| = \mathcal{O}(\varepsilon^2)$. We know that there are $\sigma, \tau > 0$ such that

$$0 = \begin{pmatrix} -\sigma I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} \widetilde{u} \\ \widetilde{v} \end{pmatrix} = \begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} \widetilde{u} \\ \widetilde{v} \end{pmatrix} - \begin{pmatrix} (\sigma-\theta)\widetilde{u} \\ (\tau-\eta)\widetilde{v} \end{pmatrix}.$$

Therefore, we have

$$\begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} e \\ f \end{pmatrix} = -\begin{pmatrix} -\theta I_m & A \\ A^T & -\eta I_n \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} + \begin{pmatrix} (\sigma-\theta)\widetilde{u} \\ (\tau-\eta)\widetilde{v} \end{pmatrix}. \tag{6.3}$$

We multiply this on the left side by P and use the fact that $PB(u,v) = B(u,v)$:

$$PB(e,f) = -B(u,v) + P((\sigma-\theta)\widetilde{u}, (\tau-\eta)\widetilde{v}). \tag{6.4}$$

Subtracting $PB(s,t) = -B(u,v)$ from (6.4), and noting that $P(u,v) = 0$ we get:

$$PB(e-s, f-t) = P((\sigma-\theta)e, (\tau-\eta)f). \tag{6.5}$$

Multiplying (6.3) on the left by $\begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}^T$ leads to

$$\begin{pmatrix} \sigma-\theta \\ \tau-\eta \end{pmatrix} = \begin{pmatrix} (x^T\widetilde{u})^{-1} & 0 \\ 0 & (y^T\widetilde{v})^{-1} \end{pmatrix} \begin{pmatrix} -\theta x^T & x^TA \\ y^TA^T & -\eta y^T \end{pmatrix} \begin{pmatrix} e \\ f \end{pmatrix}. \tag{6.6}$$

So for fixed x, y, a, and b we have $\|PB(e-s, f-t)\| = \mathcal{O}(\varepsilon^2)$. Using Lemma 6.1 and the assumption that the initial vectors are close enough to the singular vectors, we see that PB in (6.5) is invertible, so $\|(e-s, f-t)\| = \mathcal{O}(\varepsilon^2)$, which implies quadratic convergence. But if additionally (6.2) holds, then

$$\left\| \begin{pmatrix} -\theta x^T & x^TA \\ y^TA^T & -\eta y^T \end{pmatrix} \begin{pmatrix} e \\ f \end{pmatrix} \right\| = \left\| \begin{pmatrix} a^TAf \\ b^TA^Te \end{pmatrix} \right\| = \sigma \left\| \begin{pmatrix} b^Tf \\ a^Te \end{pmatrix} \right\| + \mathcal{O}(\varepsilon^2) = \mathcal{O}(\varepsilon^2),$$

so from (6.6) we see that $\|(\sigma-\theta, \tau-\eta)\| = \mathcal{O}(\varepsilon^2)$. We conclude that in this case the convergence is even cubic. $\square$

On may check that the hypotheses on x and y in the theorem are true if we choose $x_k = u_k$ or $x_k = Av_k$, and $y_k = v_k$ or $y_k = A^Tu_k$ in the process. The cubic convergence can be observed in practice, see §8.

7. Various aspects of the method.

7.1. Solving the correction equation. We now translate a number of observations for Jacobi-Davidson in [15, 14] to the JDSVD context. Consider the situation after k steps of the JDSVD algorithm. For easy reading, we again leave out the index k. In this section we take for simplicity the Galerkin spaces used in §4.1, but most arguments carry over to other choices. First we rewrite the correction equation. Because of $(s,t) \perp\perp (u,v)$, we can eliminate the projections and write (3.6) as

$$\begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r + \begin{pmatrix} \alpha u \\ \beta v \end{pmatrix}, \tag{7.1}$$

where α and β are determined by the requirement that $(s,t) \perp\perp (u,v)$. If we have a nonsingular preconditioner $M \approx \begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix}$ then we can take an approximation

$$(\widetilde{s},\widetilde{t}) = -M^{-1}r + M^{-1}(\alpha u, \beta v). \tag{7.2}$$

1. (cf. [15, p. 406, point 1]) If we approximate (s,t) simply by $\pm r$ (by taking $M = \mp I$ and $\alpha = \beta = 0$), then, because of the orthogonalization at step 2 of Algorithm 4.1, this is equivalent to taking $(\widetilde{s},\widetilde{t}) = (Av, A^T u)$. By induction one can prove that for the special case where we take this simple approximation in every step, we have
$$\mathcal{U}_{2k} = \mathcal{K}_k(AA^T, u_1) \oplus \mathcal{K}_k(AA^T, Av_1), \qquad \mathcal{V}_{2k} = \mathcal{K}_k(A^TA, v_1) \oplus \mathcal{K}_k(A^TA, A^Tu_1),$$
as long as the Krylov subspaces have a trivial intersection. Compare this with bidiagonalization, where
$$\mathcal{U}_k = \mathcal{K}_k(AA^T, Av_1), \qquad \mathcal{V}_k = \mathcal{K}_k(A^TA, v_1).$$
2. (cf. [15, p. 408, point 3]) If θ is not equal to a singular value, then $M = \begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix}$ is nonsingular and $M^{-1}r = (u,v)$. So for the updated vectors $\widetilde{u}, \widetilde{v}$ we have
$$\begin{pmatrix} \widetilde{u} \\ \widetilde{v} \end{pmatrix} = \begin{pmatrix} u+s \\ v+t \end{pmatrix} = \begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix}^{-1} \begin{pmatrix} \alpha u \\ \beta v \end{pmatrix}. \tag{7.3}$$
We conclude that exact JDSVD can be seen as an accelerated scaled RQI.
3. (cf. [15, p. 409, point 4]) If we take $M \neq \begin{pmatrix} -\theta I_m & A \\ A^T & -\theta I_n \end{pmatrix}$, M nonsingular, then with $(\widetilde{s},\widetilde{t}) = M^{-1}(\alpha u, \beta v)$ we obtain an inexact shift and invert method. This may be an attractive alternative if (7.3) is expensive.
4. When we are interested in a singular value close to a specific *target* τ, we can replace this in the left hand side of the correction equation (3.6):
$$\begin{pmatrix} I_m - uu^T & 0 \\ 0 & I_n - vv^T \end{pmatrix} \begin{pmatrix} -\tau I_m & A \\ A^T & -\tau I_n \end{pmatrix} \begin{pmatrix} I_m - uu^T & 0 \\ 0 & I_n - vv^T \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = -r.$$
The advantage of this approach is that we avoid misconvergence to some unwanted singular value 'on the way'. For example, if we want to compute the largest singular value, we can use a known approximation of $\sigma_{\max}$ as target. In practice, $\tau \approx \|A\|_\infty$ may be a good guess (see §8). For the minimal singular value, we can take $\tau = 0$, or a small positive number as target. As soon as we notice that the process starts to converge, we may replace the target in the correction equation by the current approximation to the singular value again.
5. In practice we often solve (5.1) approximately by an iterative method, for example a few steps of GMRES, or MINRES if the operator is symmetric (in case of the standard Galerkin choice). We may use a (projected) preconditioner, see §7.8.

7.2. The correction equation with non-standard Galerkin choices. In the case of non-standard Galerkin choices (see §4.3), we may have the situation that $(x, y) \neq (u, v)$. Now we exploit the flexibility of (a, b) in (5.1): by the choice

$$(a, b) = (x, y) \text{ and } (\widetilde{x}, \widetilde{y}) = (u, v), \tag{7.4}$$

we assure that the operator in (5.1) maps $(x, y)^{\perp\perp}$ onto itself, and that the asymptotic convergence is cubic according to Theorem 6.2 (if the correction equation is solved exactly). Another option is

$$(a, b) = (u, v) \text{ and } (\widetilde{x}, \widetilde{y}) = (x, y) \tag{7.5}$$

to make the operator in (5.1) symmetric. In this case the operator maps $(u, v)^{\perp\perp}$ to $(x, y)^{\perp\perp}$. Therefore, we should use a left 'preconditioner' that maps the image space $(x, y)^{\perp\perp}$ bijectively onto the domain space $(u, v)^{\perp\perp}$ (see also §8 and [14, 17]).

7.3. Comparison with Jacobi-Davidson on the augmented matrix. It is interesting to compare JDSVD with Jacobi-Davidson (JD) on the augmented matrix, starting with the 'same' starting vector $w_1 = (u_1, v_1)/\sqrt{2}$.

There are some analogies between JD and JDSVD. When their correction equations are solved exactly, both converge asymptotically cubically to a simple eigenvalue of the augmented matrix. Moreover, the costs per iteration are almost the same, the only difference is that in each step JDSVD needs a small SVD, while JD needs a small eigenvalue decomposition. The storage requirements are also comparable.

The main difference is the fact that JDSVD, by construction, searches in two (smaller) subspaces, while JD has one search space. If JD solves its correction equation exactly, then in fact it solves (7.3) with $\alpha = \beta$ [15]. This suggests that JDSVD can cope better with 'unbalanced' vectors, that is, vectors (u, v), where $\|u\| \neq \|v\|$. An extreme example of this can be seen by taking a starting vector of the form $(u_*, \delta v_*)$, for a $0 < \delta < 1$. In contrast to JD, JDSVD terminates after computing a zero residual.

Another (mostly theoretical) difference is the fact that JDSVD terminates for every starting vector after at most $\max\{m, n\}$ iterations, and JD on the augmented matrix after at most $m + n$ iterations. In §8, we compare the methods experimentally.

7.4. Refinement procedure. Suppose that we have found an approximate minimal right singular vector $v = (1 - \varepsilon^2)^{1/2} v_{\min} + \varepsilon v_{\max}$ by an iterative method applied to $A^T A$, so that $\sin \angle(v, v_{\min}) = \varepsilon$. Then, in the absence of other information, $u = Av = (1 - \varepsilon^2)^{1/2} \sigma_{\min} u_{\min} + \varepsilon \sigma_{\max} u_{\max}$ is the best approximation to the left singular vector we have to our disposal. But $\tan \angle(u, u_{\min}) \approx \varepsilon \frac{\sigma_{\max}}{\sigma_{\min}} = \kappa(A)\varepsilon$, and this can be large. Moreover, $\|u\|^2 = (1 - \varepsilon^2)\sigma_{\min}^2 + \varepsilon^2 \sigma_{\max}^2$ can be an inaccurate approximation to $\sigma_{\min}^2$, and so may $\|A^T u\|^2 / \|u\|^2$ be.

Hence, the approximations to small singular values, resulting from working with $A^T A$, may be inaccurate. In this situation, we may try to improve the approximate singular triple by a two-sided approach like JDSVD. The following lemma gives a link with [3], where a system with a matrix of the form

$$\begin{pmatrix} -\theta I_m & A & -u & 0 \\ A^T & -\theta I_n & 0 & -v \\ 2u^T & 0 & 0 & 0 \\ 0 & 2v^T & 0 & 0 \end{pmatrix} \tag{7.6}$$

is used for improving an approximate singular triple.

LEMMA 7.1. *(Cf. [14, Theorem 3.5]) The JDSVD correction equation* (5.1) *is equivalent to*

$$\begin{pmatrix} -\theta I_m & A & -u & 0 \\ A^T & -\eta I_n & 0 & -v \\ a^T & 0 & 0 & 0 \\ 0 & b^T & 0 & 0 \end{pmatrix} \begin{pmatrix} s \\ t \\ \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \theta u - Av \\ \eta v - A^T u \\ 0 \\ 0 \end{pmatrix}, \tag{7.7}$$

that is, if (s,t,α,β) *is a solution of* (7.7)*, then* (s,t) *is a solution of the correction equation* (5.1)*, and if* (s,t) *is a solution of* (5.1)*, then there exist unique* α,β *such that* (s,t,α,β) *is a solution of* (7.7).

Proof. We use the same notation as in the proof of Theorem 6.2. The system (7.7) is equivalent to

$$B(s,t) - (\alpha u, \beta v) = -r \qquad \text{and} \qquad (s,t) \perp\perp (a,b).$$

By splitting the first equation in $(x,y)^{\perp\perp}$ and its complement, we obtain

$$\begin{cases} PB(s,t) & = & -r \\ \begin{pmatrix} \alpha \\ \beta \end{pmatrix} & = & \begin{pmatrix} (x^Tu)^{-1} & 0 \\ 0 & (y^Tv)^{-1} \end{pmatrix} \begin{pmatrix} x^T & 0 \\ 0 & y^T \end{pmatrix} B \begin{pmatrix} s \\ t \end{pmatrix} \\ (s,t) & \perp\perp & (a,b) \end{cases}$$

Note that we have used $Pr = r$, $P(\alpha u, \beta v) = 0$, and $r \perp\perp (x,y)$. The first and third equation together are equivalent to the correction equation (5.1), and the second equation determines α,β uniquely. □

REMARK 7.2. *Of course, this equivalence is valid only when both equations* (7.7) *and* (5.1) *are solved exactly, not when we solve them approximately.*

In particular, when we substitute $\eta = \theta$ and $(a,b) = 2(u,v)$, the matrix in (7.7) becomes (7.6).

7.5. Smallest singular value. As mentioned in §4.1, the standard variant of JDSVD may have difficulties with finding the smallest singular value of a matrix. This is not surprising, because the small singular values of A correspond to the interior eigenvalues of the augmented matrix. But in many applications, e.g. the computation of pseudospectra, the smallest singular value is just what we want to compute.

We can use JDSVD with the non-standard Galerkin (harmonic) variants, mentioned in §4.3 and §4.4, starting with zero, or a small positive number as a target, and solve the correction equation rather accurately, possibly with the aid of a preconditioner, see §8. In this way the method is close to a shift and invert iteration, but less expensive. Of course it is hereby advantageous to have a good initial triple (e.g. coming from an iterative method on A^TA), JDSVD (with non-standard Galerkin) can then be used as refinement procedure.

7.6. Restart. A nice property of Jacobi-Davidson is its flexibility in restarting. JDSVD, too, has this advantage: we can restart at every moment in the process with any number of vectors, only keeping those parts of the search spaces that look promising, or possibly adding some extra vectors. All we have to do is compute the new resulting $H = U^TAV$ and continue. This is practical when the search spaces become large, or to avoid a breakdown in case of the non-standard Galerkin choices. Of course, JDSVD can also be *started* with search spaces of dimension larger than one.

7.7. Deflation. We can compute multiple singular triples of A by using a deflation technique. If we have found a singular triple of A, and we want to find another, we can deflate the augmented matrix to avoid finding the same triple again. For JDSVD, this can be done as follows. Suppose that U_f and V_f contain the already found singular vectors. Then it can be checked that, if we found the exact vectors,

$$\begin{pmatrix} I_m - U_fU_f^T & 0 \\ 0 & I_n - V_fV_f^T \end{pmatrix} \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} \begin{pmatrix} I_m - U_fU_f^T & 0 \\ 0 & I_n - V_fV_f^T \end{pmatrix}$$

has the same eigenvalues as the original augmented matrix, except that the found eigenvalues are transformed to zeros. The method can then be restarted with another approximate triple.

7.8. Preconditioning the correction equation. The correction equation of JDSVD can be preconditioned in a manner similar to Jacobi-Davidson (see for example [17]). We use the same notation as in the proof of Theorem 6.2, for the important case $Q = P$. Suppose that we have a preconditioner M for B. For left preconditioning we are given $(s,t) \perp\perp (x,y)$, and we have to solve for $(z_1, z_2) \perp\perp (x,y)$ from

$$PMP(z_1, z_2) = PBP(s,t).$$

Note that we project the preconditioner as well. Hence, for some α, β

$$(z_1, z_2) = M^{-1}B(s,t) - M^{-1}(\alpha u, \beta v),$$

and by using the test vectors we obtain

$$\begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}^T M^{-1} \begin{pmatrix} u & 0 \\ 0 & v \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} x & 0 \\ 0 & y \end{pmatrix}^T M^{-1}B \begin{pmatrix} s \\ t \end{pmatrix}.$$

A recipe for computing (z_1, z_2) is given by the following four steps:

(1) Compute $(\widetilde{u}_1, \widetilde{u}_2) = M^{-1}(u, 0)$ and $(\widetilde{v}_1, \widetilde{v}_2) = M^{-1}(0, v)$.
(2) Compute $(\widetilde{s}, \widetilde{t}) = M^{-1}B(s,t)$.
(3) Compute (α, β) from $\begin{pmatrix} x^T\widetilde{u}_1 & x^T\widetilde{v}_1 \\ y^T\widetilde{u}_2 & y^T\widetilde{v}_2 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} x^T\widetilde{s} \\ y^T\widetilde{t} \end{pmatrix}$.
(4) Compute $(z_1, z_2) = (\widetilde{s}, \widetilde{t}) - \alpha(\widetilde{u}_1, \widetilde{u}_2) - \beta(\widetilde{v}_1, \widetilde{v}_2)$.

An important observation is that step (1), and the computation of the 2×2 matrix in step (3) have to be performed only once at the start of the iterative solution process of the correction equation. The right hand side of the correction equation, minus the residual, is handled similarly.

8. Numerical experiments. Our experiments are coded in MATLAB and executed on a SUN workstation. The following lemma implies that up to rounding errors, it is not a loss of generality to consider (rectangular) diagonal matrices A.

LEMMA 8.1. *If there are no rounding errors, and JDSVD's correction equation* (5.1) *in step k is solved by l_k steps of GMRES, then JDSVD applied to*
(a) $A = U_*\Sigma V_*^T$, with starting vectors u_1 and v_1
(b) Σ, with starting vectors $\widetilde{u}_1 := U_*^T u_1$ and $\widetilde{v}_1 := V_*^T v_1$
gives 'the same' results, that is,

$$\widetilde{\theta}_k = \theta_k \qquad \textit{and} \qquad \|\widetilde{r}_k\| = \|r_k\|.$$

Proof. Define

$$Q = \begin{pmatrix} U_*^T & 0 \\ 0 & V_*^T \end{pmatrix},$$

then Q is orthogonal and one may verify that $(\widetilde{u}_1, \widetilde{v}_1) = Q(u_1, v_1)$, $\widetilde{\theta}_1 := \widetilde{u}_1^T\Sigma\widetilde{v}_1 = u_1^T A v_1 =: \theta_1$, and $\widetilde{r}_1 = Qr_1$. A well-known property of Krylov subspaces assures that (see [12, p. 264]):

$$Q^T\mathcal{K}_l\left(\begin{pmatrix} 0 & \Sigma \\ \Sigma^T & 0 \end{pmatrix}, \widetilde{r}\right) = \mathcal{K}_l\left(Q^T\begin{pmatrix} 0 & \Sigma \\ \Sigma^T & 0 \end{pmatrix}Q, Q^T\widetilde{r}\right) = \mathcal{K}_l\left(\begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}, r\right).$$

With little extra work one can check that the same relation holds for the shifted and projected matrices that are present in the correction equation (5.1), where one should bear in mind that all other vectors involved in the projectors (a, b, x, y, $\widetilde{x}$, and $\widetilde{y}$) must also be altered for the Σ-system in the obvious way. So the approximate solutions from the correction equations satisfy $(\widetilde{s}_1, \widetilde{t}_1) = Q(s_1, t_1)$. By induction we can prove that $\widetilde{U}_k = U_*^T U_k$ and $\widetilde{V}_k = V_*^T V_k$, so the projected matrices are the same in both cases: $\widetilde{H}_k := \widetilde{U}_k^T \Sigma \widetilde{V}_k = U_k^T A V_k = H_k$. In particular, the approximations to the singular values are the same, and the approximations (u_k, v_k) and $(\widetilde{u}_k, \widetilde{v}_k)$ are orthogonal transformations of each other: $(\widetilde{u}_k, \widetilde{v}_k) = Q(u_k, v_k)$ and $\widetilde{r}_k = Q r_k$, so $\|\widetilde{r}_k\| = \|r_k\|$. $\square$

For this reason, we first study some phenomena on $A = \mathrm{diag}([1:100])$ and $A = \mathrm{diag}([1:1000])$. In Figure 8.1(a), the solid line is the convergence history of (the standard variant of Algorithm 4.1 of) JDSVD, for the computation of the largest singular triple of $A = \mathrm{diag}([1:100])$. The starting vectors are the normalized $v_1 = v_{\max} + 0.1r$, where r is a vector with random entries, chosen from a uniform distribution on the unit interval, and $u_1 = Av_1/\|Av_1\|$. The dots represent the error in the approximation $\sigma_{\max} - \theta_k^{(k)}$. In all figures, a horizontal dotted line indicates the stopping tolerance. We solve the correction equation by 200 steps of GMRES. Because the (augmented) matrices in the correction equation (step 7 of Algorithm 4.1) are of size 200×200, this means (theoretically) exactly, so according to Theorem 6.2 we expect cubic convergence. In Figure 8.1(a) we see, for instance, that the error in the approximation in iteration number 5 decreases from $\approx 10^{-2}$ to $\approx 10^{-7}$.

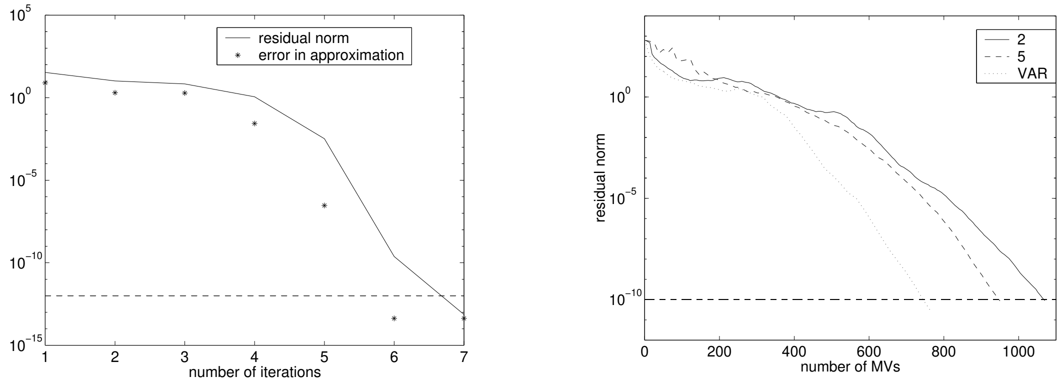

(a) (b)

FIG. 8.1: (a) The convergence history of the exact JDSVD algorithm for diag(1..100) as in Algorithm 4.1: residual norm (solid line) and error in the approximations to $\sigma_{\max}$ (dots). The horizontal dotted line indicates the stopping tolerance. (b) Convergence for diag(1..1000) using respectively 5, 2, and a variable number of GMRES steps to solve the correction equation.

In Figure 8.1(b), we take $A = \mathrm{diag}([1:1000])$, and u_1 and v_1 random vectors (as described above) with unit norm. We experiment with the number of GMRES steps. For the solid line, we solve the correction equation approximately by 5 steps of GMRES, which we denote by GMRES$_5$, for the dashed line by GMRES$_2$, and for the dotted line by a variable number equal to $\max\{2 \cdot (\lceil -\log\|r\|\rceil + 1), 0\}$. Measured in terms of matrix-vector products, the variable choice is best, followed by GMRES$_5$. An explanation of this is that when the initial approximations are not good (as in this case), it is of no use to try hard to solve the correction equation in the beginning. When we are almost converging, it may make sense to solve it more accurately to get fast convergence. See also [17].

In Figure 8.2(a) we compare, for $A = \mathrm{diag}([1:1000])$, standard JDSVD for the three largest singular triples (solid), with Jacobi-Davidson (JD) on the augmented matrix for the computation of the three largest eigenpairs (dashed), each with GMRES$_5$. For JDSVD, we take v_1 a random vector, and $u_1 = Av_1/\|Av_1\|$. For JD we take the 'same' starting vector $(u_1, v_1)/\sqrt{2}$. We see that JDSVD is faster for the first triple, for the second and third we restart with a good approximation, and then the histories are similar.

In Figure 8.2(b), we do the same, but now using GMRES$_2$. For the first two triples, JDSVD is somewhat faster than JD, for the third JDSVD in first instance (mis)converges to the fourth largest singular value 997. Other experiments also suggest that JDSVD is generally (somewhat) faster than JD on the augmented matrix.

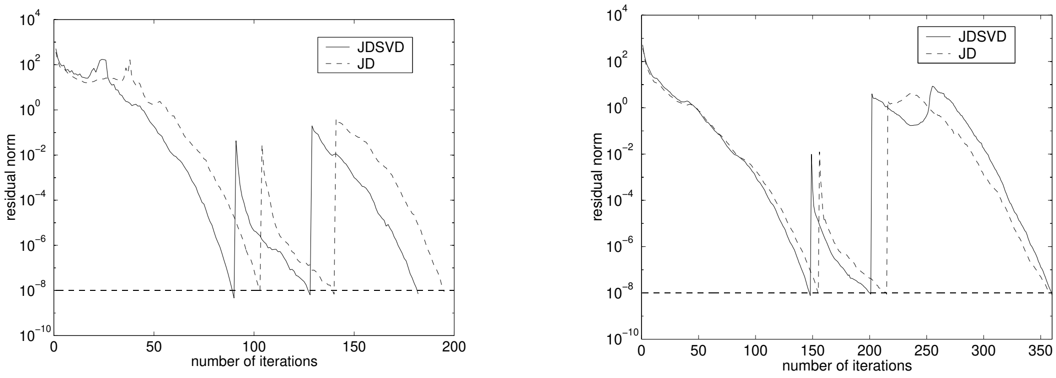

(a) (b)

FIG. 8.2: (a) JDSVD (solid) and JD (dashed) for the three largest σ's of diag(1..1000). (b) The same as Figure 8.2(a), only with GMRES_2 to solve the correction equation.

Next, we take some examples from the Matrix Market (these matrices can be downloaded from `http://math.nist.gov/MatrixMarket`). For Figure 8.3(a), we apply different JDSVD variants to find the smallest singular triple of PDE225, using two random starting vectors and GMRES_{10} (no preconditioning). In all variants, we take initially target 0, but when $\|r\| < 10^{-3}$, we replace the target by the best approximations again (see §7.1 point 4). The solid line is the standard choice, we see an irregular convergence history, as could be expected (see §4). The dashed line represents the Galerkin choice (4.5), where in the correction equation (5.1) we substitute (7.4). Finally, the dash-dotted line is (4.5) with (7.5) substituted in (5.1). In the last case, as seen in §7.2, the operator in (5.1) maps $(u, v)^{\perp\perp}$ to $(x, y)^{\perp\perp}$. Since in this case $v = y$ but $u \neq x$, we use a left 'preconditioner' to handle the correction equation correctly. The preconditioned identity

$$\begin{pmatrix} I_m - \frac{xu^T}{u^Tx} & 0 \\ 0 & I_n \end{pmatrix} I_{m+n} \begin{pmatrix} I_m - \frac{ux^T}{x^Tu} & 0 \\ 0 & I_n \end{pmatrix}$$

maps $(x, y)^{\perp\perp}$ back to $(u, v)^{\perp\perp}$.

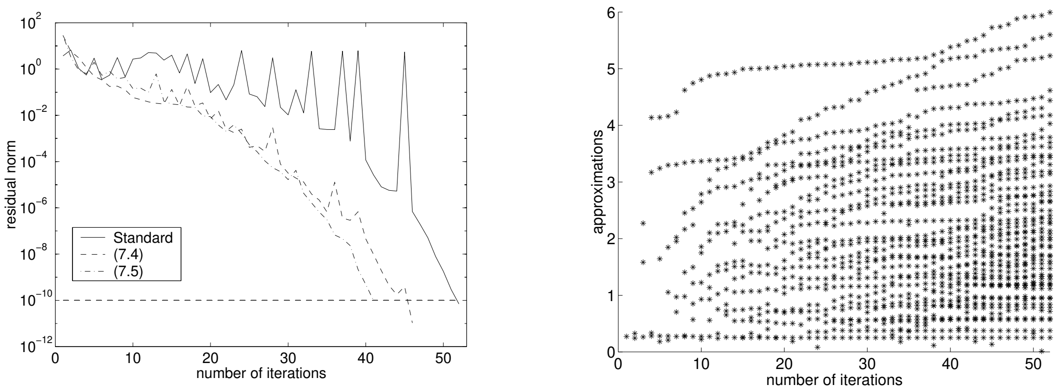

(a) (b)

FIG. 8.3: (a) Three different JDSVD variants for the computation of $\sigma_{\min}$ of PDE225: standard, (4.5)+(5.1)+(7.4), and (4.5)+(5.1)+(7.5) (b) (Non-harmonic) approximations to the singular values by the standard variant.

In Figure 8.3(b), standard JDSVD's approximations to the singular values during this process are plotted. These are 'regular', non-harmonic estimates. Note the monotone convergence of the approximations to the largest singular values, but the irregular behavior of the approximations to the smallest singular value.

Next, we compare JDSVD with Lanczos applied to A^TA for the computation of $\sigma_{\max}$. These methods are of a different nature. The Lanczos method can be viewed as an accelerated power method, while JDSVD can be seen as an an accelerated inexact RQI. An advantage of JDSVD is that we may use preconditioning for the correction equation. Therefore we expect that if we have a reasonable preconditioner, and if preconditioning is relatively cheap in comparison to a multiplication by A or A^T, then JDSVD can be cheaper than Lanczos. On the other hand, if $m \gg n$, or if there is no good or cheap preconditioner available, then we expect that Lanczos will be better. Table 8.1 shows some test results. For JDSVD, we take a target $\tau \approx \|A\|_\infty$, in the hope that $\tau \approx \sigma_{\max}$. We make an incomplete LU-decomposition (using a drop tolerance displayed in the table) of the augmented matrix (1.1) minus τ times the identity, and use $M = LU$ as preconditioner. The starting vector v_1 is the vector with all ones, and then normalized, and u_1 is a random vector. We solve the correction equation by only preconditioning the residual ('0 steps of GMRES'). The process is continued until $\|r\| < 10^{-8}$. Lanczos' method uses v_1 as starting vector, and continues until $\|(A^TA - \theta^2)v\| < 10^{-8}$. The matrix A_1 stands for $\mathrm{diag}(1:100) + 0.1 \cdot \mathrm{rand}(n,n)$, where $\mathrm{rand}(n,n)$ denotes an $n \times n$-matrix with random entries, chosen from a uniform distribution on the unit interval. See [13] for more information on the origin and singular values of the other matrices. For JDSVD, a pair is given, consisting of the number of matrix-vector products (MV's) and the number of systems with L or U. For Lanczos we show the number of MV's.

TABLE 8.1: Some experiments with JDSVD to compute $\sigma_{\max}$, using incomplete LU-factorizations of the shifted augmented matrix. The number of matrix-vector products, and the number of systems with L or U is displayed in the 5th column. The shift (or target) τ (6th column) for the preconditioning is roughly taken to be $\|A\|_\infty$. The last three columns give information on the incomplete LU-factorization: the drop tolerance of ILU, and the resulting number of nonzeros of L and U. We compare JDSVD's results with the MV's of Lanczos applied to A^TA (4th column).

Matrix	size	nnz(A)	Lan(A^TA)	JDSVD	τ	tol	nnz(L)	nnz(U)
HOR131	434 $\times$ 434	4182	30	(28, 65)	0.90	10^{-2}	1792	1792
PORES3	532 $\times$ 532	3474	–	(72, 175)	$2 \cdot 10^5$	10^{-1}	1301	1300
SHERMAN1	1000 $\times$ 1000	3750	74	(24, 66)	5	10^{-2}	4805	4803
A_1	100 $\times$ 100	10000	102	(38, 108)	106	10^{-2}	299	299

For HOR131, the target τ is relatively far from $\sigma_{\max} \approx 0.66$. We see that although JDSVD uses less MV's than Lanczos, Lanczos is cheaper when we take the preconditioning into account. Although for PORES3 ($\sigma_{\max} \approx 1.5 \cdot 10^5$) Lanczos finds a good approximate vector, its residual does not reach the required 10^{-8} due to the high condition number of the matrix. JDSVD does converge, so this is an example of a situation where JDSVD could be used as refinement. For SHERMAN1, the target is a reasonable approximation to $\sigma_{\max} \approx 5.05$. When we take the preconditioning into account, Lanczos is cheaper than JDSVD. The last row of the table is an example where preconditioning is relatively cheap. The reason for this is that we now take the diagonal of A, instead A itself, to form an augmented matrix of the form (1.1) and to make an ILU- decomposition. Using far more MV's, Lanczos is (also counting the preconditioning) more expensive.

Finally, in Table 8.2, we compare JDSVD for the computation of $\sigma_{\min}$ with Lanczos applied to $\left(A^TA\right)^{-1}$. We use the Galerkin choice (4.5) for JDSVD. Note that the comparison with Lanczos is mainly meant to get an idea of how well JDSVD performs. In practice, for large (sparse) A, it is too expensive to work with A^{-1} and A^{-T} or $\left(A^TA\right)^{-1}$. For JDSVD, we take a small target $\tau = 10^{-5}$, drop tolerance 10^{-3}, and again we make an incomplete LU-decomposition based on this target. The starting vectors are the same as for Table 8.1. We solve the correction equation by only preconditioning the residual ('0 steps of GMRES'). Both processes are continued until $\|r\| < 10^{-7}$.

We see that although JDSVD may in general use more MV's, it is much cheaper than Lanczos, due to the sparsity of A, L, and U. For HOR131, Lanczos does not converge to the required 10^{-7}. Again A_1 serves as an example for the situation where preconditioning is relatively cheap, which makes JDSVD attractive. We also tried Lanczos applied to A^TA for the computation of $\sigma_{\min}$, but

TABLE 8.2: Some experiments with JDSVD to compute $\sigma_{\min}$. The numbers of matrix-vector products and systems with L or U (3rd column), and the number of nonzeros of L and U are displayed. We compare JDSVD's results with the number of MV's of Lanczos applied to $(A^TA)^{-1}$.

Matrix	Lan$(A^TA)^{-1}$	JDSVD	nnz(L)	nnz(U)
HOR131	–	(26, 72)	20593	21167
PORES3	12	(36, 108)	3683	5491
SHERMAN1	20	(20, 54)	11575	11738
A_1	14	(28, 78)	200	200

the results were bad (262 MV's for A_1, and more than 500 MV's for the other matrices).

9. Conclusions. We have discussed an alternative approach for the computation of a few singular values and vectors of a matrix. The JDSVD method searches in two separate subspaces, and it can be interpreted as an inexact Newton method for the singular value problem. JDSVD can also be seen as an inexact accelerated scaled RQI method. Therefore the best results may be expected when we have a good initial starting triple (refinement), but we can start with any approximations. While the asymptotic convergence is cubic if the correction equation is solved exactly, in practice we solve it approximately, and then the convergence typically looks (super)linear. Although we discussed mainly the application of JDSVD for the largest and smallest singular value, the method is in principle suitable for all singular values. We may use preconditioning for the solution of the correction equation. This can be a decisive factor for fast convergence. Experiments indicate that JDSVD is a good competitor to other iterative SVD methods, in particular when A is (almost) square and we have a reasonable, relatively cheap preconditioner for the correction equation.

Acknowledgments. The author thanks Henk van der Vorst and Gerard Sleijpen for interesting discussions, and the referees for many helpful comments.

REFERENCES

[1] J. CULLUM, R. A. WILLOUGHBY, AND M. LAKE, *A Lanczos algorithm for computing singular values and vectors of large matrices*, SIAM J. Sci. Stat. Comput., 4 (1983), pp. 197–215.

[2] 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.

[3] J. J. DONGARRA, *Improving the accuracy of computed singular values*, SIAM J. Sci. Stat. Comput., 4 (1983), pp. 712–719.

[4] D. R. FOKKEMA, G. L. G. SLEIJPEN, AND H. A. VAN DER VORST, *Accelerated inexact Newton schemes for large systems of nonlinear equations*, SIAM J. Sci. Comput., 19 (1998), pp. 657–674 (electronic).

[5] G. GOLUB AND W. KAHAN, *Calculating the singular values and pseudo-inverse of a matrix*, J. Soc. Indust. Appl. Math. Ser. B Numer. Anal., 2 (1965), pp. 205–224.

[6] G. H. GOLUB, F. T. LUK, AND M. L. OVERTON, *A block Lanczos method for computing the singular values and corresponding singular vectors of a matrix*, ACM Trans. Math. Software, 7 (1981), pp. 149–169.

[7] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The John Hopkins University Press, Baltimore, London, 3rd ed., 1996.

[8] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1985.

[9] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, 1991.

[10] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, J. Res. Nat. Bur. Stand., 45 (1950), pp. 255–282.

[11] C. C. PAIGE, B. N. PARLETT, AND H. A. VAN DER VORST, *Approximate solutions and eigenvalue bounds from Krylov subspaces*, Num. Lin. Alg. Appl., 2 (1995), pp. 115–133.

[12] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998. Corrected reprint of the 1980 original.

[13] B. PHILIPPE AND M. SADKANE, *Computation of the fundamental singular subspace of a large matrix*, Linear Alg. Appl., 257 (1997), pp. 77–104.

[14] G. L. G. SLEIJPEN, A. G. L. BOOTEN, D. R. FOKKEMA, AND H. A. VAN DER VORST, *Jacobi-Davidson type methods for generalized eigenproblems and polynomial eigenproblems*, BIT, 36 (1996), pp. 595–633. International Linear Algebra Year (Toulouse, 1995).

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

[16] G. L. G. Sleijpen and H. A. van der Vorst, *The Jacobi-Davidson method for eigenvalue problems and its relation with accelerated inexact Newton schemes*, in Iterative Methods in Linear Algebra, II., S. D. Margenov and P. S. Vassilevski, eds., vol. 3 of IMACS Series in Computational and Applied Mathematics, New Brunswick, NJ, U.S.A., 1996, IMACS, pp. 377–389. Proceedings of the Second IMACS International Symposium on Iterative Methods in Linear Algebra, June 17-20, 1995, Blagoevgrad.

[17] G. L. G. Sleijpen, H. A. van der Vorst, and E. Meijerink, *Efficient expansion of subspaces in the Jacobi-Davidson method for standard and generalized eigenproblems*, Electron. Trans. Numer. Anal., 7 (1998), pp. 75–89 (electronic). Large scale eigenvalue problems (Argonne, IL, 1997).

[18] G. W. Stewart and J.-g Sun, *Matrix Perturbation Theory*, Academic Press, San Diego, California, 1990.

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

[20] S. Varadhan, M. W. Berry, and G. H. Golub, *Approximating dominant singular triplets of large sparse matrices via modified moments*, Numer. Algorithms, 13 (1996), pp. 123–152.

[21] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, 1965.