

Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 19/42

Ian N. Zwaan

**Cross product-free matrix pencils for computing
generalized singular values**

December 18, 2019

<http://www.math.uni-wuppertal.de>

Cross product–free matrix pencils for computing generalized singular values*

Ian N. Zwaan[†]

Abstract. It is well known that the generalized (or quotient) singular values of a matrix pair (A, C) can be obtained from the generalized eigenvalues of a matrix pencil consisting of two augmented matrices. The downside of this reformulation is that one of the augmented matrices requires a cross products of the form C^*C , which may affect the accuracy of the computed quotient singular values if C has a large condition number. A similar statement holds for the restricted singular values of a matrix triplet (A, B, C) and the additional cross product BB^* . This article shows that we can reformulate the quotient and restricted singular value problems as generalized eigenvalue problems without having to use any cross product or any other matrix-matrix product. Numerical experiments show that there indeed exist situations in which the new reformulation leads to more accurate results than the well-known reformulation.

Key words. generalized eigenvalue problem, augmented matrix, generalized singular value decomposition, GSVD, quotient singular value decomposition, QSVD, restricted singular value decomposition, RSVD.

AMS subject classification. 65F15; 15A18

1 Introduction. Suppose that $A \in \mathbb{C}^{p \times q}$, then it is well known that the (ordinary) singular values of A can be obtained from the eigenvalues of either of the products

$$(1) \quad \mathcal{A} = A^*A \quad \text{or} \quad \mathcal{A} = AA^*,$$

or from the augmented matrix

$$(2) \quad \mathcal{A} = \begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix}.$$

Likewise, if $C \in \mathbb{C}^{n \times q}$ is a second matrix, then the generalized singular values (precise definitions of the various singular value concepts follow later), also called the quotient singular values [4], of the matrix pair (A, C) can be obtained either from the generalized eigenvalues of the pencil

$$(3) \quad \mathcal{A} - \lambda\mathcal{B} = A^*A - \lambda C^*C$$

or from the augmented pencil

$$(4) \quad \mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix} - \lambda \begin{bmatrix} I & 0 \\ 0 & C^*C \end{bmatrix}.$$

Furthermore, if $B \in \mathbb{C}^{p \times m}$ is a third matrix, then the restricted singular values of the triplet (A, B, C) , can be obtained from the pencil

$$(5) \quad \mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix} - \lambda \begin{bmatrix} BB^* & 0 \\ 0 & C^*C \end{bmatrix}.$$

*This work was supported in part by the Deutsche Forschungsgemeinschaft through the collaborative research centre SFB-TRR55. Version: December 18, 2019.

[†]Faculty of Mathematics and Natural Sciences, Bergische Universität Wuppertal, ianzwaan.com.

Except for (2), all the above (generalized) eigenvalue problems require cross products of the form A^*A , BB^* , or C^*C . Textbooks by, for example, Stewart [20, Sec. 3.3.2], Higham [9, Sec. 20.4], and Golub and Van Loan [8, Sec. 8.6.3] dictate that these types of products are undesirable for poorly conditioned matrices, because condition numbers are squared and accuracy may be lost. This loss of accuracy has also been investigated, for example, by Jia [14] for the SVD, and by and Huang and Jia [13] for the GSVD.

Numerical methods exist, for large and sparse matrices in particular, purposefully designed to avoid explicit use of the unwanted cross products. Examples for the singular value problem include Golub–Kahan–Lanczos bidiagonalization; see, e.g., Demmel [6, Sec. 6.3.3]; and JDSVD by Hochstenbach [10]. Examples for the quotient singular value problem include a bidiagonalization method by Zha [23], JDGSVD by Hochstenbach [11], generalized Krylov methods by Hochstenbach, Reichel, and Yu [12] and Reichel and Yu [18, 19], and a Generalized–Davidson based projection method [25].

The purpose of this article is to show that we can reformulate the ordinary, quotient, and restricted singular value problems as a matrix pencil that consists of two augmented matrices, neither of which require any cross product, or any other matrix-matrix product.

To see precisely how the eigenvalues and eigenvectors of the matrix pencils relate to the singular values and vectors, we can use the Kronecker Canonical Form (KCF). The KCF, detailed in Section 2, is the generalization of the Jordan form of a matrix to matrix pencils, and fully characterises the generalized eigenstructure of pencils. The next step is to reformulate the ordinary singular value decomposition (OSVD) in Section 3, and to analyze the corresponding KCF. This particular reformulation is purely for exposition, since the augmented matrix (2) is already free of cross products. That is, the new reformulation of the OSVD is the simplest case we can consider and the easiest to verify by hand, but already uses the same general approach we use for the other two singular value problems. In the next two sections, Section 4 and 5, we discuss the reformulation of the quotient singular value decomposition (QSVD) and the restricted singular value decomposition (RSVD). The former is better known and more widely used in practice, while the latter is more general. The generality of the RSVD makes it tedious to describe in full detail, and also tedious to get the KCF of its corresponding cross product-free pencil. Still, its treatment is essentially identical to the simpler cases of the OSVD and QSVD. The numerical experiments that follow in Section 7 show that, for some matrix pairs and triplets, we can compute the singular values more accurately from the new cross product-free pencils than from the typical augmented pencils. The new pencils are not without their own downsides, such as an increased problem size and the presence of Jordan blocks, which we further discuss in the conclusion in Section 8.

Throughout this text, I denotes an identity matrix of appropriate size; M^T and M^* denote the transpose and Hermitian transpose, respectively, of a matrix M ; and \otimes denotes the Kronecker product. Furthermore, for some permutation π of length k , the corresponding permutation matrix is given by

$$\Pi = \begin{bmatrix} e_{\pi(1)} & e_{\pi(2)} & \dots & e_{\pi(k)} \end{bmatrix},$$

where the e_j are the j th canonical basis vectors of length k . The notation of permutations and permutation matrices is extended to block matrices, and permute entire blocks of rows or columns at once.

2 The Kronecker canonical form. As mentioned before, the KCF is a generalization of the Jordan form to matrix pencils, and its importance is that it fully describes the generalized eigenvalues and generalized eigenspaces of a matrix pencil. This also means that two matrix pencils are equivalent if-and-only-if they have the same KCF [7, Thm. 5], where the meaning of equivalence is as in the following definition.

Definition 1 (Gantmacher [7, Def. 1]). Two pencils of rectangular matrices $\mathcal{A} - \lambda\mathcal{B}$ and $\mathcal{A}_1 - \lambda\mathcal{B}_1$ of the same dimensions $k \times \ell$ are called strictly equivalent if there exists nonsingular \mathcal{X} and \mathcal{Y} , independent of λ , such that $\mathcal{Y}^*(\mathcal{A} - \lambda\mathcal{B})\mathcal{X} = \mathcal{A}_1 - \lambda\mathcal{B}_1$.

Now we are ready for the definition of the KCF, which is given by the following theorem.

Theorem 2 (Adapted from Kågström [15, Sec. 8.7.2] and Gantmacher [7, Ch. 2]). Let $\mathcal{A}, \mathcal{B} \in \mathbb{C}^{k \times \ell}$, then there exists nonsingular matrices $\mathcal{X} \in \mathbb{C}^{\ell \times \ell}$ and $\mathcal{Y} \in \mathbb{C}^{k \times k}$ such that $\mathcal{Y}^*(\mathcal{A} - \lambda\mathcal{B})\mathcal{X}$ equals

$$\text{diag}(0_{\beta_0 \times \alpha_0}, L_{\alpha_1}, \dots, L_{\alpha_m}, L_{\beta_1}^T, \dots, L_{\beta_n}^T, N_{\gamma_1}, \dots, N_{\gamma_p}, J_{\delta_1}(\zeta_1), \dots, J_{\delta_q}(\zeta_q));$$

where the

$$L_{\alpha_j} = \begin{bmatrix} 0 & 1 & & \\ & \ddots & \ddots & \\ & & 0 & 1 \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & & 1 & 0 \end{bmatrix}$$

are $\alpha_j \times (\alpha_j + 1)$ singular blocks of right (or column) minimal index α_j , the

$$L_{\beta_j}^T = \begin{bmatrix} 0 & & & \\ 1 & \ddots & & \\ & \ddots & 0 & \\ & & & 1 \end{bmatrix} - \lambda \begin{bmatrix} 1 & & & \\ 0 & \ddots & & \\ & \ddots & & 1 \\ & & & & 0 \end{bmatrix}$$

are $(\beta_j + 1) \times \beta_j$ singular blocks of left (or row) minimal index β_j , the

$$N_{\gamma_j} = \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & \ddots & 0 \\ & & & 1 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & & 0 \end{bmatrix}$$

are $\gamma_j \times \gamma_j$ Jordan blocks corresponding to an infinite eigenvalues γ_j , and the

$$J_{\delta_j}(\zeta_j) = \begin{bmatrix} \zeta & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \zeta \end{bmatrix} - \lambda \begin{bmatrix} 1 & 0 & & \\ & \ddots & \ddots & \\ & & \ddots & 0 \\ & & & & 1 \end{bmatrix}$$

are $\delta_j \times \delta_j$ Jordan blocks corresponding to finite eigenvalues $\zeta_j \in \mathbb{C}$.

The meaning of minimal indices is not important for this paper, and an interested reader may refer to the references for more information. Furthermore, the zero block $0_{\beta_0 \times \alpha_0}$ corresponds to α_0 blocks L_0 and β_0 blocks L_0^T by convention. We include it here explicitly for clarity, and because in this paper we have no KCF with any other L_α and L_β^T blocks. Moreover, in this paper all Jordan blocks corresponding to finite nonzero eigenvalues ζ_j will be such that $J_{\delta_j}(\zeta_j) = \zeta_j - \lambda$; that is, $\delta_j = 1$ whenever $\zeta_j \neq 0$.

3 The ordinary singular value decomposition. The augmented matrix (2) does not contain any cross products. Still, we can take the cross product–free matrix pencil for the QSVD and consider the OSVD as a special case. The main reason for us to do so, is to show how we can determine the corresponding KCF for the simplest case that we can consider. That is, this section sets the stage and the reformulations for the QSVD and RSVD and their proofs in the sections to come, follow from the same general idea.

Let us start by recalling the definition of the ordinary singular value decomposition with the following theorem (see, e.g., Golub and Van Loan [8]).

Theorem 3 (Ordinary singular value decomposition (OSVD)). *Let $A \in \mathbb{C}^{p \times q}$; then there exist unitary matrices $U \in \mathbb{C}^{p \times p}$ and $V \in \mathbb{C}^{q \times q}$ such that*

$$\Sigma = U^*AV = \begin{matrix} & q_1 & q_2 \\ p_1 & \begin{bmatrix} D_\sigma & 0 \\ 0 & 0 \end{bmatrix} \\ p_2 & & \end{matrix},$$

where $p = p_1 + p_2$, $q = q_1 + q_2$, and $p_1 = q_1$. Furthermore, $D_\sigma = \text{diag}(\sigma_1, \dots, \sigma_{p_1})$ with $\sigma_j > 0$ for all $j = 1, 2, \dots, p_1$.

Next, let us consider the eigenvalue decomposition of a 4×4 pencil that we can consider as the cross product–free pencil for the 1×1 matrix $\sigma \geq 0$.

Lemma 4. *Suppose σ is a positive real number, and consider the pencil*

$$(6) \quad \mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & \sigma & 0 & 0 \\ \sigma & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

Then the unitary matrices

$$\mathcal{X} = \frac{1}{2} \begin{bmatrix} 1 & -1 & -i & i \\ 1 & -1 & i & -i \\ 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \end{bmatrix} \quad \text{and} \quad \mathcal{Y} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -i & i \\ 1 & -1 & i & -i \end{bmatrix}$$

are such that

$$\mathcal{Y}^*(\mathcal{A} - \lambda\mathcal{B})\mathcal{X} = \text{diag}(\sqrt{\sigma}, -\sqrt{\sigma}, i\sqrt{\sigma}, -i\sqrt{\sigma}) - \lambda I.$$

Proof. The proof is by direct verification. □

With the above definition and lemma, we can state and prove the following theorem and corollary.

Theorem 5. *Let A and its corresponding SVD be as in Theorem 3. Then the KCF of the pencil*

$$(7) \quad \mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & A & 0 & 0 \\ A^* & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \\ I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \end{bmatrix}$$

consists of the following blocks.

1. A series of $p_2 + q_2$ blocks $J_2(0)$.
2. The blocks $J_1(\sqrt{\sigma_1}), \dots, J_1(\sqrt{\sigma_{p_1}}), J_1(-\sqrt{\sigma_1}), \dots, J_1(-\sqrt{\sigma_{p_1}}), J_1(i\sqrt{\sigma_1}), \dots, J_1(i\sqrt{\sigma_{p_1}}), J_1(-i\sqrt{\sigma_1}), \dots, J_1(-i\sqrt{\sigma_{p_1}})$, where $\sigma_1, \dots, \sigma_{p_1}$ are the nonzero singular values of A .

Proof. Let $\mathcal{A}_0 - \lambda\mathcal{B}_0 = \mathcal{A} - \lambda\mathcal{B}$, and define the transformations $\mathcal{X}_0 = \mathcal{Y}_0 = \text{diag}(U, V, U, V)$. Then the pencil $\mathcal{A}_1 - \lambda\mathcal{B}_1 = \mathcal{Y}_0^*(\mathcal{A}_0 - \lambda\mathcal{B}_0)\mathcal{X}_0$ is a square 8×8 block matrix of dimension

$$\underbrace{p_1 + p_2}_{p} + \underbrace{q_1 + q_2}_{q} + \underbrace{p_1 + p_2}_{p} + \underbrace{q_1 + q_2}_{q}.$$

Now let \mathcal{X}_1 and \mathcal{Y}_1 be the permutation matrices corresponding to the permutations

$$\pi_{\mathcal{X}} = (2, 6, 4, 8, \quad 1, 3, 5, 7) \quad \text{and} \quad \pi_{\mathcal{Y}} = (6, 2, 8, 4, \quad 1, 3, 5, 7),$$

respectively. Then the pencil $\mathcal{A}_2 - \lambda\mathcal{B}_2 = \mathcal{Y}_1^*(\mathcal{A}_1 - \lambda\mathcal{B}_1)\mathcal{X}_1$ is block diagonal, and has the following blocks along its diagonal.

1. The $(p_2 + p_2 + q_2 + q_2) \times (p_2 + p_2 + q_2 + q_2)$ block

$$\begin{bmatrix} 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} = \begin{bmatrix} J_2(0) \otimes I & & & \\ & & & \\ & & & J_2(0) \otimes I \\ & & & \end{bmatrix},$$

which yields $p_2 + q_2$ blocks $J_2(0)$ in the KCF of (7) after suitable permutations.

2. The $(p_1 + q_1 + p_1 + q_1) \times (p_1 + q_1 + p_1 + q_1)$ block

$$\begin{bmatrix} 0 & D_\sigma & 0 & 0 \\ D_\sigma & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \\ I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \end{bmatrix},$$

which reduces to a diagonal matrix with the Jordan “blocks” $J_1(\pm\sqrt{\pm\sigma_j})$ after suitable permutations and applying Lemma 4. □

Corollary 6. *The (right) eigenvectors belonging to the eigenvalues $\sqrt{\sigma_j}$, $-\sqrt{\sigma_j}$, $i\sqrt{\sigma_j}$, and $-i\sqrt{\sigma_j}$ are*

$$\begin{bmatrix} u_j \\ v_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -u_j \\ -v_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -iu_j \\ iv_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} iu_j \\ -iv_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix},$$

respectively. Here $u_j = Ue_j$ and $v_j = Ve_j$.

As we will see in the next two sections, we can follow the same general approach with the QSVD and the RSVD. That is, we start with the “obvious” transformation of the pencils, permute the resulting block matrices, and invoke an analogue of Lemma 4.

4 The quotient singular value decomposition. The QSVD can be used to solve, for example, generalized eigenvalue problems of the form (3), generalized total least squares problems, general-form Tikhonov regularization, least squares problems with equality constraints, etc.; see, e.g., Van Loan [21] and Bai [2] for more information. The QSVD is defined the theorem below and comes from Paige and Saunders [17], but has the blocks of the partitioned matrices permuted to be closer to a special case of the RSVD from Section 5.

Theorem 7 (Quotient singular value decomposition). *Let $A \in \mathbb{C}^{p \times q}$ and $C \in \mathbb{C}^{n \times q}$; then there exist a nonsingular matrix $Y \in \mathbb{C}^{q \times q}$ and unitary matrices $U \in \mathbb{C}^{p \times p}$ and $V \in \mathbb{C}^{n \times n}$ such that*

$$U^*AY = \begin{matrix} & q_1 & q_2 & q_3 & q_4 \\ \begin{matrix} p_1 \\ p_2 \\ p_3 \end{matrix} & \begin{bmatrix} 0 & 0 & D_\alpha & 0 \\ 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix} \quad \text{and} \quad V^*CY = \begin{matrix} & q_1 & q_2 & q_3 & q_4 \\ \begin{matrix} n_1 \\ n_2 \\ n_3 \end{matrix} & \begin{bmatrix} 0 & I & 0 & 0 \\ 0 & 0 & D_\gamma & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix},$$

where $n_2 = p_1 = q_3$, $n_1 = q_2$, and $p_2 = q_4$. Furthermore, $D_\alpha = \text{diag}(\alpha_1, \dots, \alpha_{p_1})$ and $D_\gamma = \text{diag}(\gamma_1, \dots, \gamma_{p_1})$ are such that $\alpha_j, \gamma_j > 0$ and $\alpha_j^2 + \gamma_j^2 = 1$ for all $j = 1, 2, \dots, p_1$. These $p_1 = q_3$ pairs (α_j, γ_j) together with q_2 pairs $(0, 1)$ and q_4 pairs $(1, 0)$ are called the nontrivial pairs. The remaining q_1 pairs $(0, 0)$ are called the trivial pairs. Each nontrivial pair (α, γ) corresponds to a quotient singular value $\sigma = \alpha/\gamma$, where the result is ∞ by convention if $\gamma = 0$.

As before, we first consider the one dimensional case and generalize Lemma 4 to the QSVD.

Lemma 8. *Suppose α and γ are positive real numbers and consider the pencil*

$$\mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & \alpha & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & \gamma \\ 1 & 0 & 0 & 0 \\ 0 & \gamma & 0 & 0 \end{bmatrix}.$$

Then the nonsingular matrices

$$\mathcal{X} = \mathcal{Y} = \sigma^{-1/4} \text{diag}(1, \gamma^{-1}, \sqrt{\sigma}, \sqrt{\sigma}),$$

where $\sigma = \alpha/\gamma$, are such that $\mathcal{Y}^*(\mathcal{A} - \lambda\mathcal{B})\mathcal{X}$ is a pencil of the form (6).

Proof. The proof is by direct verification. □

Using the above definition of the QSVD in Theorem 7, and the reduction in Lemma 8, we can state and prove the following theorem and corollary.

Theorem 9. *Let A, C , and their corresponding QSVD be as in Theorem 7. Then the KCF of the pencil*

$$(8) \quad \mathcal{A} - \mathcal{B} = \begin{bmatrix} 0 & A & 0 & 0 \\ A^* & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & I & 0 \\ 0 & 0 & 0 & C^* \\ I & 0 & 0 & 0 \\ 0 & C & 0 & 0 \end{bmatrix}$$

consists of the following blocks.

1. A $q_1 \times q_1$ zero block, which corresponds to quotient singular pairs of the form $(0, 0)$.
2. A series of n_3 blocks N_1 , which correspond to $(1, 0)$ pairs.

3. A series of p_2 blocks N_3 , which correspond to $(1, 0)$ pairs.
4. A series of $p_3 + q_2$ blocks $J_2(0)$, which correspond to $(0, 1)$ pairs.
5. The blocks $J_1(\sqrt{\sigma_1}), \dots, J_1(\sqrt{\sigma_{p_1}}), J_1(-\sqrt{\sigma_1}), \dots, J_1(-\sqrt{\sigma_{p_1}}), J_1(i\sqrt{\sigma_1}), \dots, J_1(i\sqrt{\sigma_{p_1}}), J_1(-i\sqrt{\sigma_1}), \dots, J_1(-i\sqrt{\sigma_{p_1}})$, where $\sigma_1, \dots, \sigma_{p_1}$ are the finite and nonzero quotient singular values of the matrix pair (A, C) .

Proof. Let $\mathcal{A}_0 - \lambda\mathcal{B}_0 = \mathcal{A} - \lambda\mathcal{B}$, and define the transformations $\mathcal{X}_0 = \mathcal{Y}_0 = \text{diag}(U, Y, U, V)$. Then the pencil $\mathcal{A}_1 - \lambda\mathcal{B}_1 = \mathcal{Y}_0^*(\mathcal{A}_0 - \lambda\mathcal{B}_0)\mathcal{X}_0$ is a square 13×13 block matrix of dimension

$$\underbrace{p_1 + p_2 + p_3}_p + \underbrace{q_1 + q_2 + q_3 + q_4}_q + \underbrace{p_1 + p_2 + p_3}_p + \underbrace{n_1 + n_2 + n_3}_n.$$

Now let \mathcal{X}_1 and \mathcal{Y}_1 be the permutation matrices corresponding to the permutations

$$\begin{aligned} \pi_{\mathcal{X}} &= (4, 13, 7, 9, 2, 3, 10, 5, 11, 1, 6, 8, 12), \\ \pi_{\mathcal{Y}} &= (4, 13, 2, 9, 7, 10, 3, 11, 5, 1, 6, 8, 12), \end{aligned}$$

respectively. Then the pencil $\mathcal{A}_2 - \lambda\mathcal{B}_2 = \mathcal{Y}_1^*(\mathcal{A}_1 - \lambda\mathcal{B}_1)\mathcal{X}_1$ is block diagonal, and has the following blocks along its diagonal.

1. A $q_1 \times q_1$ block of zeros.
2. The $n_3 \times n_3$ block $I - \lambda \cdot 0 = N_1 \otimes I$.
3. The $(p_2 + p_2 + q_4) \times (q_4 + p_2 + p_2)$ block

$$\begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & I & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{bmatrix} = N_3 \otimes I,$$

where $p_2 = q_4$.

4. The $(p_3 + p_3 + q_2 + n_1) \times (p_3 + p_3 + n_1 + q_2)$ block

$$\begin{bmatrix} 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} = \begin{bmatrix} J_2(0) \otimes I & \\ & J_2(0) \otimes I \end{bmatrix},$$

where $n_1 = q_2$, which yields $p_3 + q_2$ blocks $J_2(0)$ in the KCF of (8) after suitable permutations.

5. The $(p_1 + q_3 + p_1 + n_2) \times (p_1 + q_3 + p_1 + n_2)$ block

$$\begin{bmatrix} 0 & D_\alpha & 0 & 0 \\ D_\alpha & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & I & 0 \\ 0 & 0 & 0 & D_\gamma \\ I & 0 & 0 & 0 \\ 0 & D_\gamma & 0 & 0 \end{bmatrix},$$

where $n_2 = p_1 = q_3$, which reduces to a diagonal matrix with the Jordan “blocks” $J_1(\pm\sqrt{\pm\sigma_j})$ after suitable permutations and applying Lemma 8.

□

Corollary 10. *The eigenvectors belonging to the nonzero finite eigenvalues $\sqrt{\sigma_j}$, $-\sqrt{\sigma_j}$, $i\sqrt{\sigma_j}$, and $-i\sqrt{\sigma_j}$ are*

$$\begin{bmatrix} u_j \\ \gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -u_j \\ -\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -iu_j \\ i\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} iu_j \\ -i\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix},$$

respectively. Here $u_j = Ue_j$, $v_j = Ve_j$, and $y_j = Ye_{q_1+j}$.

Compared to the OSVD, the KCF for (8) has three extra blocks, nameley the first three blocks in the list of Theorem 9. The first of these blocks is associated with the singularity of the pencil, while the second and third blocks are associated with infinite eigenvalues/singular values.

5 The restricted singular value decomposition. The RSVD is useful for, for example, analyzing structured rank perturbations, computing low-rank approximations of partitioned matrices, minimization or maximization of the bilinear form x^*Ay under the constraints $\|B^*x\|, \|Cy\| \neq 0$, solving matrix equations of the form $BXC = A$, constrained total least squares with exact rows and columns, and generalized Gauss–Markov models with constraints. See, e.g., Zha [22] and De Moor and Golub [5] for more information. The RSVD is defined by the theorem below and comes from [24], but was adapted from the preceding two references. The full definition of the RSVD is quite tedious, but we can follow the same general approach as in the previous two sections.

Theorem 11 (The restricted singular value decomposition). *Let $A \in \mathbb{C}^{p \times q}$, $B \in \mathbb{C}^{p \times m}$, and $C \in \mathbb{C}^{n \times q}$, and define $r_A = \text{rank } A$, $r_B = \text{rank } B$, $r_C = \text{rank } C$, $r_{AB} = \text{rank} \begin{bmatrix} A & B \end{bmatrix}$, $r_{AC} = \text{rank} \begin{bmatrix} A \\ C \end{bmatrix}$, and $r_{ABC} = \text{rank} \begin{bmatrix} A & B \\ C & 0 \end{bmatrix}$. Then the triplet of matrices (A, B, C) can be factorized as $A = X^{-*}\Sigma_\alpha Y^{-1}$, $B = X^{-*}\Sigma_\beta U^*$, and $C = V\Sigma_\gamma Y^{-1}$, where $X \in \mathbb{C}^{p \times p}$ and $Y \in \mathbb{C}^{q \times q}$ are nonsingular, and $U \in \mathbb{C}^{m \times m}$ and $V \in \mathbb{C}^{n \times n}$ are orthonormal. Furthermore, Σ_α , Σ_β , and Σ_γ have nonnegative entries and are such that $\begin{bmatrix} \Sigma_\alpha & \Sigma_\beta \\ \Sigma_\gamma & 0 \end{bmatrix}$ can be written as*

$$(9) \quad \begin{array}{c} \begin{array}{cccccc|cccc} q_1 & q_2 & q_3 & q_4 & q_5 & q_6 & m_1 & m_2 & m_3 & m_4 \\ \hline p_1 & \begin{bmatrix} 0 & 0 & D_\alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & 0 & I \end{bmatrix} & \begin{bmatrix} D_\beta & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ \hline p_5 & \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ \hline n_1 & \begin{bmatrix} 0 & I & 0 & 0 & 0 & 0 \\ 0 & 0 & D_\gamma & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ \hline n_4 & \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} & \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \end{array} \end{array} \quad \begin{array}{l} p_1 = q_3 = r_{ABC} + r_A - r_{AB} - r_{AC} \\ p_2 = q_4 = r_{AC} + r_B - r_{ABC} \\ p_3 = q_5 = r_{AB} + r_C - r_{ABC} \\ p_4 = q_6 = r_{ABC} - r_B - r_C \\ p_5 = r_{AB} - r_A, q_2 = r_{AC} - r_A \\ p_6 = p - r_{AB}, q_1 = q - r_{AC} \\ n_1 = q_2, m_4 = p_5 \\ n_2 = m_1 = p_1 = q_3 \\ n_3 = p_3 = q_5, m_2 = p_2 = q_4 \\ n_4 = n - r_C, m_3 = m - r_B \end{array}$$

where $D_\alpha = \text{diag}(\alpha_1, \dots, \alpha_{p_1})$, $D_\beta = \text{diag}(\beta_1, \dots, \beta_{p_1})$, and $D_\gamma = \text{diag}(\gamma_1, \dots, \gamma_{p_1})$. Moreover, α_j , β_j , and γ_j are scaled such that $\alpha_j^2 + \beta_j^2 + \gamma_j^2 = 1$ for $j = 1, \dots, p_1$. Besides the p_1 triplets $(\alpha_j, \beta_j, \gamma_j)$, there are p_2 triplets $(1, 1, 0)$, p_3 triplets $(1, 0, 1)$, p_4 triplets $(1, 0, 0)$, and $\min\{p_5, q_2\}$ triplets $(0, 1, 1)$.

This leads to a total of $p_1 + p_2 + p_3 + p_4 + \min\{p_5, q_2\} = r_A + \min\{p_5, q_2\} = \min\{r_{AB}, r_{AC}\}$ regular triplets of the form (α, β, γ) with $\alpha^2 + \beta^2\gamma^2 = 1$. Each of these triplets corresponds to a restricted singular value $\sigma = \alpha/(\beta\gamma)$, where the result is ∞ by convention if $\alpha \neq 0$ and $\beta\gamma = 0$. Finally, the triplet has a right (or column) trivial block of dimension $q_1 = \dim(\mathcal{N}(A) \cap \mathcal{N}(C))$, and a left (or row) trivial block of dimension $p_6 = \dim(\mathcal{N}(A^*) \cap \mathcal{N}(B^*))$.

As before, we first consider the one dimensional case and generalize Lemmas 4 and 8 to the RSVD.

Lemma 12. Suppose α, β , and γ are positive real numbers and consider the pencil

$$\mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & \alpha & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & \beta & 0 \\ 0 & 0 & 0 & \gamma \\ \beta & 0 & 0 & 0 \\ 0 & \gamma & 0 & 0 \end{bmatrix}.$$

Then the nonsingular matrices

$$\mathcal{X} = \mathcal{Y} = \sigma^{-1/4} \text{diag}(\beta^{-1}, \gamma^{-1}, \sqrt{\sigma}, \sqrt{\sigma}),$$

where $\sigma = \alpha/(\beta\gamma)$, are such that $\mathcal{Y}^*(\mathcal{A} - \lambda\mathcal{B})\mathcal{X}$ is a pencil of the form (6).

Proof. The proof is by direct verification. □

With the definition of the RSVD in Theorem 11, and the reduction in Lemma 12, we can state and prove the following theorem and corollary.

Theorem 13. Let A, B, C , and their corresponding RSVD be as in Theorem 11. Then the KCF of the pencil

$$(10) \quad \mathcal{A} - \lambda\mathcal{B} = \begin{bmatrix} 0 & A & 0 & 0 \\ A^* & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & B & 0 \\ 0 & 0 & 0 & C^* \\ B^* & 0 & 0 & 0 \\ 0 & C & 0 & 0 \end{bmatrix}$$

consists of the following blocks.

1. A $(p_6 + q_1) \times (p_6 + q_1)$ zero block, which correspond to restricted singular triplets of the form $(0, 0, 0)$.
2. A series of $p_4 + q_6 + m_3 + n_4$ blocks N_1 , which correspond to $(1, 0, 0)$ triplets.
3. A series of p_2 blocks N_3 , which correspond to $(1, 1, 0)$ triplets.
4. A series of p_3 blocks N_3 , which correspond to $(1, 0, 1)$ triplets.
5. A series of $p_5 + q_2 = m_4 + n_1$ blocks $J_2(0)$, which correspond to $(0, 1, 1)$ triplets.
6. The blocks $J_1(\sqrt{\sigma_1}), \dots, J_1(\sqrt{\sigma_{p_1}}), J_1(-\sqrt{\sigma_1}), \dots, J_1(-\sqrt{\sigma_{p_1}}), J_1(i\sqrt{\sigma_1}), \dots, J_1(i\sqrt{\sigma_{p_1}}), J_1(-i\sqrt{\sigma_1}), \dots, J_1(-i\sqrt{\sigma_{p_1}})$, where $\sigma_1, \dots, \sigma_{p_1}$ are the finite and nonzero restricted singular values of the matrix triplet (A, B, C) .

Proof. Let $\mathcal{A}_0 - \lambda\mathcal{B}_0 = \mathcal{A} - \lambda\mathcal{B}$, and define the transformations $\mathcal{X}_0 = \mathcal{Y}_0 = \text{diag}(X, Y, U, V)$. Then the pencil $\mathcal{A}_1 - \lambda\mathcal{B}_1 = \mathcal{Y}_0^*(\mathcal{A}_0 - \lambda\mathcal{B}_0)\mathcal{X}_0$ is a square 20×20 block matrix of dimension

$$\underbrace{p_1 + \cdots + p_6}_p + \underbrace{q_1 + \cdots + q_6}_q + \underbrace{m_1 + \cdots + m_4}_m + \underbrace{n_1 + \cdots + n_4}_n.$$

Now let \mathcal{X}_1 and \mathcal{Y}_1 be permutation matrices corresponding to the permutations

$$\begin{aligned}\pi_{\mathcal{X}} &= (6, 7, \quad 12, 4, 15, 20, \quad 10, 14, 2, \quad 3, 19, 11, \quad 5, 16, 8, 17, \quad 1, 9, 13, 18), \\ \pi_{\mathcal{Y}} &= (6, 7, \quad 4, 12, 15, 20, \quad 2, 14, 10, \quad 11, 19, 3, \quad 16, 5, 17, 8, \quad 1, 9, 13, 18),\end{aligned}$$

respectively. Then the pencil $\mathcal{A}_2 - \lambda\mathcal{B}_2 = \mathcal{Y}_1^*(\mathcal{A}_1 - \lambda\mathcal{B}_1)\mathcal{X}_1$ is block diagonal, and has the following blocks along its diagonal.

1. A $(p_6 + q_1) \times (p_6 + q_1)$ block of zeros.
2. The $(p_4 + q_6 + m_3 + n_4) \times (q_6 + p_4 + m_3 + n_4)$ block

$$\begin{bmatrix} I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = N_1 \otimes I,$$

where $p_4 = q_6$.

3. The $(p_2 + m_2 + q_4) \times (q_4 + m_2 + p_2)$ block

$$\begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & I & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{bmatrix} = N_3 \otimes I,$$

where $m_2 = p_2 = q_4$.

4. The $(q_5 + n_3 + p_3) \times (p_3 + n_3 + q_5)$ block

$$\begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & I & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{bmatrix} = N_3 \otimes I,$$

where $n_3 = p_3 = q_5$.

5. The $(m_4 + p_5 + n_1 + q_2) \times (p_5 + m_4 + q_2 + n_1)$ block

$$\begin{bmatrix} 0 & I & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & I \\ 0 & 0 & 0 & 0 \end{bmatrix} - \lambda \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} = \begin{bmatrix} J_2(0) \otimes I & & & \\ & & & \\ & & & \\ & & & J_2(0) \otimes I \end{bmatrix}$$

where $m_4 = p_5$ and $n_1 = q_2$, which yields $p_5 + q_2$ blocks $J_2(0)$ in the KCF of (10) after suitable permutations.

6. The $(p_1 + q_3 + m_1 + n_2) \times (p_1 + q_3 + m_1 + n_2)$ block

$$\begin{bmatrix} 0 & D_\alpha & 0 & 0 \\ D_\alpha & 0 & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & D_\beta & 0 \\ 0 & 0 & 0 & D_\gamma \\ D_\beta & 0 & 0 & 0 \\ 0 & D_\gamma & 0 & 0 \end{bmatrix},$$

where $m_1 = n_2 = p_1 = q_3$, which reduces to a diagonal matrix with the Jordan “blocks” $J_1(\pm\sqrt{\pm\sigma_j})$ after suitable permutations and applying Lemma 12. □

Corollary 14. *The eigenvectors belonging to the nonzero finite eigenvalues $\sqrt{\sigma_j}$, $-\sqrt{\sigma_j}$, $i\sqrt{\sigma_j}$, and $-i\sqrt{\sigma_j}$ are*

$$\begin{bmatrix} \beta_j^{-1}x_j \\ \gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -\beta_j^{-1}x_j \\ -\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ \sqrt{\sigma_j}v_j \end{bmatrix}, \quad \begin{bmatrix} -i\beta_j^{-1}x_j \\ i\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} i\beta_j^{-1}x_j \\ -i\gamma_j^{-1}y_j \\ \sqrt{\sigma_j}u_j \\ -\sqrt{\sigma_j}v_j \end{bmatrix},$$

respectively. Here $u_j = Ue_j$, $v_j = Ve_{n_1+j}$, $x_j = Xe_j$, and $y_j = Ye_{q_1+q_2+j}$.

A comparison with the KCF of the typical augmented pencil (5) might be insightful. Using the transform $\text{diag}(U, Y)$ and noting that we get a 12×12 block matrix with sizes

$$p + q = p_1 + \cdots + p_6 + q_1 + \cdots + q_6,$$

we can use the permutations

$$\begin{aligned} \pi_x &= (6, 7, \quad 12, 4, \quad 10, 2, \quad 3, 11, \quad 5, 8, \quad 1, 9), \\ \pi_y &= (6, 7, \quad 4, 12, \quad 2, 10, \quad 11, 3, \quad 5, 8, \quad 1, 9). \end{aligned}$$

to see that we get the following blocks for the KCF of (5).

1. A $(p_6 + q_1) \times (p_6 + q_1)$ zero block.
2. A $(p_4 + q_6) \times (q_6 + p_4)$ block $N_1 \otimes I$.
3. A $(p_2 + q_4) \times (q_4 + p_2)$ block $N_2 \otimes I$.
4. A $(p_3 + q_5) \times (q_5 + p_3)$ block $N_2 \otimes I$.
5. A $(p_5 + q_2) \times (p_5 + q_2)$ block $J_1(0) \otimes I$.
6. A $(p_1 + q_3) \times (p_1 + q_3)$ block

$$\begin{bmatrix} 0 & D_\alpha \\ D_\alpha & 0 \end{bmatrix} - \lambda \begin{bmatrix} D_\beta^2 & 0 \\ 0 & D_\gamma^2 \end{bmatrix}.$$

Hence, we see that the KCF of (5) and (10) have comparable blocks. The primary qualitative difference is that the zero singular values do not result in Jordan blocks of size larger than 1.

6 A tree of generalized singular value decompositions. For the OSVD we started with the matrix pencil (7), where \mathcal{A} and \mathcal{B} together have a total of eight nonzero blocks. Two of these blocks are the considered matrix A and its Hermitian transpose A^* , the remaining six blocks are identity matrices. Then, for the QSVD we replaced two of these identity matrices with the considered matrix C and its Hermitian transpose C^* . Next, for the RSVD we replaced two more identity matrices with the considered B and its Hermitian transpose B^* . We now have two identity matrices left, and a natural question to ask is: if we replace these two identities with two other matrices, can we meaningfully interpret the resulting pencil as a generalized SVD? The answer to this question is yes if we take

$$(11) \quad \mathcal{A} - \lambda \mathcal{B} = \begin{bmatrix} 0 & A & 0 & 0 \\ A^* & 0 & 0 & 0 \\ 0 & 0 & D^*D & 0 \\ 0 & 0 & 0 & EE^* \end{bmatrix} - \lambda \begin{bmatrix} 0 & 0 & B & 0 \\ 0 & 0 & 0 & C^* \\ B^* & 0 & 0 & 0 \\ 0 & C & 0 & 0 \end{bmatrix},$$

where $D \in \mathbb{C}^{k \times m}$ and $E \in \mathbb{C}^{n \times l}$ for some $k, l \geq 1$. To see this, suppose for simplicity that D and E are both invertible, then the pencil above corresponds to the RSVD of the triplet $(A, BD^{-1}, E^{-1}C)$. In turn, the restricted singular values of the latter triplet coincide with the ordinary singular values of the product $DB^{-1}AC^{-1}E$ if B and C are also invertible. The usefulness of this QQQQ-SVD is unclear, and the pencil (11), which is no longer cross product-free, is primarily of theoretical interest. For the full definition of this decomposition, see De Moor and Zha [4] for more details.

7 Numerical experiments. The general idea for the experiments in this section is straightforward. Simply generate matrices with sufficiently large condition numbers, form the pencils (4), (8), (5), and (10), and check the accuracy of the computed generalized eigenvalues. The following table summarizes the relation between the various singular value problems, their reformulations as matrix pencils, and the corresponding eigenvalues.

λ	OSVD	QSVD	RSVD
σ^2	(1)	(3)	—
$\pm\sigma$	(2)	(4)	(5)
$\pm\sqrt{\pm\sigma}$	(7)	(8)	(10)

We can generate the matrices in different ways, but the straightforward approaches described next suffice to prove the usefulness of the new pencils. We proceed as follows for the QSVD; given the dimension n and sufficiently large condition numbers κ_Y and κ_Σ :

1. Generate random orthonormal matrices U_Y, V_Y, U and V [16].
2. Compute $\Sigma_Y = \text{diag}(\eta_1, \dots, \eta_n)$ and $Y = U_Y \Sigma_Y V_Y^T$, where $\eta_j = \kappa_Y^{1/2 - (j-1)/(n-1)}$.
3. Let $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$, $\Sigma_\alpha = \text{diag}(\alpha_1, \dots, \alpha_n)$, and $\Sigma_\gamma = \text{diag}(\gamma_1, \dots, \gamma_n)$, where $\sigma_j = \kappa_\Sigma^{1/2 - (j-1)/(n-1)}$, $\alpha_j = \sigma_j(1 + \sigma_j^2)^{-1/2}$, and $\gamma_j = (1 + \sigma_j^2)^{-1/2}$.
4. Compute $A = U \Sigma_\alpha Y^{-1}$ and $C = V \Sigma_\gamma Y^{-1}$.

With the steps above we generate matrices satisfying $\kappa(Y) = \kappa_Y$, $\Sigma = \Sigma_\alpha \Sigma_Y^{-1}$, $\Sigma_\alpha^2 + \Sigma_Y^2 = I$, $\kappa(\Sigma) = \kappa_\Sigma$, and $\kappa(\Sigma_\alpha) = \kappa(\Sigma_\gamma) = \kappa_\Sigma^{1/2}$. To ensure we do not lose precision prematurely, we must do all the computations in higher-precision arithmetic¹ and convert A and C to IEEE 754 double precision at

¹The required random numbers are generated in double precision for simplicity.

the end. Given the dimension n and the condition numbers κ_X , κ_Y , and κ_Σ , we can take a similar approach for the RSVD.

1. Generate U , V , Y , Σ , Σ_α , and Σ_γ as before, and let $\Sigma_\beta = I$.
2. Generate X in the same way as Y , but with the condition number κ_X instead of κ_Y .
3. Let $A = X^{-T}\Sigma_\alpha Y^{-1}$, $B = X^{-T}\Sigma_\beta U^T$, and $C = V\Sigma_\gamma Y^{-1}$.

Again, we have to do all the computations in higher-precision arithmetic and convert A , B , and C to double precision at the end.

Since the generated B and C have full row and column ranks, respectively, the pencils (4) and (5) are Hermitian positive definite generalized eigenvalue problems. This means that we can use specialized solvers to compute the generalized eigenvalues and vectors. The same is not true for (8) and (10), where both A and B are indefinite. Thus, we have to use generic eigensolvers that do not take the special structure of the pencils into account. And as a result, the computed eigenvalues may not be purely real or purely imaginary. Furthermore, suppose that

$$\tilde{\lambda}_1 \approx \sqrt{\sigma}, \quad \tilde{\lambda}_2 \approx i\sqrt{\sigma}, \quad \tilde{\lambda}_3 \approx -\sqrt{\sigma}, \quad \text{and} \quad \tilde{\lambda}_4 \approx -i\sqrt{\sigma},$$

then we have no guarantee that the magnitudes $|\lambda_j|$ match to high accuracy. The latter problem leads to the question: which $|\lambda_j|$ should we pick? Anecdotal observations suggest that the squared (absolute) geometric mean

$$(12) \quad (|\lambda_1| |\lambda_2| |\lambda_3| |\lambda_4|)^{-1/2} \approx \sigma$$

is a reasonable choice.

We choose one small example to illustrate the accuracy lost from working with cross products, and also the approximation picking problem. Suppose that $n = 4$, $\kappa_Y = 10^7$, and $\kappa_\Sigma = 10$; then we get the following results for a randomly generated matrix pair. First the exact quotient singular values rounded to 12 digits after the point:

$$3.162277660168 \quad 1.467799267622 \quad 0.681292069058 \quad 0.316227766017.$$

Taking the square roots of the generalized eigenvalues of the pencil (3) gives:

$$3.\underline{186032382196} \quad 1.\underline{467953130046} \quad 0.\underline{681383277148} \quad 0.\underline{316227883689}.$$

The magnitudes of the generalized eigenvalues of the augmented pencil (4) are:

$$\begin{array}{cccc} 3.\underline{186055633628} & 1.\underline{467953295365} & 0.\underline{681384866293} & 0.\underline{316227814411} \\ 3.\underline{186055633628} & 1.\underline{467953295365} & 0.\underline{681384866293} & 0.\underline{316227814411}. \end{array}$$

The cross product-free pencil gives the squared magnitudes:

$$\begin{array}{cccc} 3.\underline{162277324135} & 1.\underline{467799263849} & 0.\underline{681292066129} & 0.\underline{316227766010} \\ 3.\underline{162277661974} & 1.\underline{467799267618} & 0.\underline{681292069112} & 0.\underline{316227766020} \\ 3.\underline{162277661974} & 1.\underline{467799267618} & 0.\underline{681292069112} & 0.\underline{316227766020} \\ 3.\underline{162278000456} & 1.\underline{467799271389} & 0.\underline{681292072105} & 0.\underline{316227766030}. \end{array}$$

Taking the squared absolute geometric means (12) yields:

$$3.\underline{162277662135} \quad 1.\underline{467799267619} \quad 0.\underline{681292069115} \quad 0.\underline{316227766020}.$$

For comparison, the generalized singular values computed with LAPACK, which uses a Jacobi-type iteration (see, e.g., Bai and Demmel [3] and the routine xTGSJA) are:

$$3.\underline{162277659936} \quad 1.\underline{467799267555} \quad 0.\underline{681292069045} \quad 0.\underline{316227766024}.$$

The underlined digits of the approximations serve as a visual indicator of their accuracy; if d decimals places are underlined, then d is the largest integer for which the absolute error is less than $5 \cdot 10^{-(d+1)}$. The loss of accuracy caused by the cross products in the pencils (3) and (4) is obvious in the results. We can also see that the digits of pairs of eigenvalues of the typical augmented pencil (4) match, while the final two to four digits of the eigenvalues from the new cross product-free pencil (8) differ. Still, the results of the latter are about as accurate as the singular values computed by LAPACK [1]. Hence, users should prefer LAPACK for computing the QSVD of dense matrix pairs, but may consider using (8) in combination with existing solvers for sparse matrix pairs.




Figure 1: The median of the maximum errors in the computed quotient singular values. On the left $\kappa_\Sigma = \kappa_X = 10$ and κ_Y varies, and on the right $\kappa_X = \kappa_Y = 10$ and κ_Σ varies. The dash-dotted line corresponds to the quadratic pencil (3), the solid line to the typical augmented pencil (4), the dashed line to the cross product-free pencil (8), and the dotted line to the LAPACK results.

For a more quantitative analysis, we can generate a large number of matrix pairs and triplets for each combination of n , κ_X , κ_Y , and κ_Σ . For each sample we compute the singular values using the different methods we have available, and use the (squared) absolute geometric mean, when necessary, to average the approximations of each singular value. Then we compute the approximation errors and pick the maximum error for each matrix pair or triplet. And finally, we take the median of all the maximum errors. As a measure for the error we can use the chordal metric

$$\chi(\sigma, \tilde{\sigma}) = \frac{|\sigma - \tilde{\sigma}|}{\sqrt{1 + \sigma^2} \sqrt{1 + \tilde{\sigma}^2}} = \frac{|\sigma^{-1} - \tilde{\sigma}^{-1}|}{\sqrt{1 + \sigma^{-2}} \sqrt{1 + \tilde{\sigma}^{-2}}}.$$

Here σ is an exact (computed with high-precision arithmetic) singular value, and $\tilde{\sigma}$ is a computed approximation.

Figure 1 shows the results for the QSVD with 10000 samples for each combination of parameters. The figure shows that we lose about twice as much accuracy with (3) and (4) than with the new cross product-free pencil. Furthermore, we see that the quotient singular values computed from the new pencil are almost as accurate as the ones computed with LAPACK. When we increase κ_Σ while keeping κ_Y modest, we see that the squared and augmented pencils lost accuracy as κ_Σ increases, while the results from the new pencil and LAPACK remain small.

Likewise, Figure 2 shows the results for the RSVD. Again we see that we lose accuracy about twice as fast when using cross products if κ_Y increases and κ_X and κ_Σ remain modest. Furthermore, we also again see that we lose accuracy with cross products when κ_X and κ_Y remain modest and κ_Σ increases, while the cross product-free approaches keep producing accurate results.




Figure 2: The median of the maximum errors in the computed restricted singular values. On the left $\kappa_X = \kappa_\Sigma = 10$ and κ_γ varies, on the right $\kappa_\Sigma = 10$ and $\kappa_X = \kappa_\gamma$ varies; $n = 10$ in both cases. The solid line corresponds to the typical augmented pencil (5), the dashed line to the cross product-free pencil (10), and the dotted line (on top of the dashed line) to the results from the RSVD algorithm described in [24].

8 Conclusion. We have seen how we can reformulate the quotient and restricted singular value problems as generalized eigenvalue problems, and without using cross products like A^*A , BB^* , or C^*C . Moreover, the numerical examples show the benefits of working with the cross product-free pencils instead of the typical augmented pencils. That is, singular values computed from the former may be more accurate than singular values computed from the latter when B or C are ill-conditioned.

Still, when we use the cross product-free pencils we have to contend with downsides that we have ignored so far. For example, for the QSVD and square A and C , the dimension of the cross product-free pencil (8) is twice as large as the dimension of the typical augmented pencil (4), and four times as large as the dimension of the squared pencil (3). This makes a large computational difference for dense solvers with cubic complexity. Another problem is that we can no longer use specialized solvers for Hermitian-positive-definite problems, even when B is of full row rank and C of full column rank. We also have to contend with extra Jordan blocks, for example for $\sigma = 0$, and the difficulties that numerical software has when dealing with Jordan blocks.

But the above downsides of the cross product-free pencils also provide opportunities for future research. For example, can we reduce the performance overhead by developing solvers (e.g., for large-and-sparse matrices) that exploit the block structure of the pencils? Can we incorporate our knowledge of the structure of the spectrum to improve the computed results?

Acknowledgements. I wish to thank Andreas Frommer for his comments and suggestions.

References.

- [1] E. ANDERSON, Z. BAI, C. BISCHOF, L. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users' Guide*, Society for Industrial and Applied Mathematics, 3 ed., Jan. 1999, <https://doi.org/10.1137/1.9780898719604>.
- [2] Z. BAI, *The CSD, GSVD, their applications and computations*, techreport 958, University of Minnesota. IMA Preprint Series.
- [3] Z. BAI AND J. W. DEMMEL, *Computing the generalized singular value decomposition*, *SIAM J. Sci. Comput.*, 14 (1993), pp. 1464–1486, <https://doi.org/10.1137/0914085>.
- [4] B. DE MOOR AND H. ZHA, *A tree of generalizations of the ordinary singular value decomposition*, *Linear Algebra Appl.*, 147 (1991), pp. 469–500, [https://doi.org/10.1016/0024-3795\(91\)90243-p](https://doi.org/10.1016/0024-3795(91)90243-p).

- Preprint – Preprint – Preprint – Preprint – Preprint
- [5] B. L. DE MOOR AND G. H. GOLUB, *The restricted singular value decomposition: Properties and applications*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 401–425, <https://doi.org/10.1137/0612029>.
 - [6] J. DEMMEL, *Singular value decomposition: Iterative algorithms*, in *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, eds., Society for Industrial and Applied Mathematics. Section 6.3.
 - [7] F. GANTMACHER, *The theory of matrices*, vol. 2, Chelsea Publishing Company.
 - [8] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, 4 ed.
 - [9] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, Society for Industrial and Applied Mathematics, 2 ed., Jan. 2002, <https://doi.org/10.1137/1.9780898718027>.
 - [10] M. E. HOCHSTENBACH, *A Jacobi–Davidson type SVD method*, SIAM J. Sci. Comput., 23 (2001), pp. 606–628, <https://doi.org/10.1137/s1064827500372973>.
 - [11] M. E. HOCHSTENBACH, *A Jacobi–Davidson type method for the generalized singular value problem*, Linear Algebra Appl., 431 (2009), pp. 471–487, <https://doi.org/10.1016/j.laa.2009.03.003>.
 - [12] M. E. HOCHSTENBACH, L. REICHEL, AND X. YU, *A golub–kahan-type reduction method for matrix pairs*, J Sci Comput, 65 (2015), pp. 767–789, <https://doi.org/10.1007/s10915-015-9990-x>.
 - [13] J. HUANG AND Z. JIA, *On choices of formulations of computing the generalized singular value decomposition of a matrix pair*, <https://arxiv.org/abs/1907.10392>.
 - [14] Z. JIA, *Using cross-product matrices to compute the SVD*, Numer Algor, 42 (2006), pp. 31–61, <https://doi.org/10.1007/s11075-006-9022-x>.
 - [15] B. KÅGSTRÖM, *Generalized non-hermitian eigenvalue problems: Singular matrix pencils*, in *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, eds., Society for Industrial and Applied Mathematics. Section 8.7.
 - [16] F. MEZZADRI, *How to generate random matrices from the classical compact groups*, Notices Amer. Math. Soc., 54, pp. 592–604.
 - [17] C. PAIGE AND M. SAUNDERS, *Towards a generalized singular value decomposition*, SIAM J. Numer. Anal., 18 (1981), pp. 398–405, <https://doi.org/10.1137/0718026>.
 - [18] L. REICHEL AND X. YU, *Matrix decompositions for tikhonov regularization*, Electron. Trans. Numer. Anal., 43, pp. 223–243.
 - [19] L. REICHEL AND X. YU, *Tikhonov regularization via flexible arnoldi reduction*, Bit Numer Math, 55 (2015), pp. 1145–1168, <https://doi.org/10.1007/s10543-014-0542-9>.
 - [20] G. STEWART, *Matrix Algorithms*, vol. 2, Society for Industrial and Applied Mathematics, 1 ed., Jan. 2001, <https://doi.org/10.1137/1.9780898718058>.
 - [21] C. F. VAN LOAN, *Generalizing the singular value decomposition*, SIAM J. Numer. Anal., 13 (1976), pp. 76–83, <https://doi.org/10.1137/0713009>.

- [22] H. ZHA, *The restricted singular value decomposition of matrix triplets*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 172–194, <https://doi.org/10.1137/0612014>.
- [23] H. ZHA, *Computing the generalized singular values/vectors of large sparse or structured matrix pairs*, Numerische Mathematik, 72 (1996), pp. 391–417, <https://doi.org/10.1007/s002110050175>.
- [24] I. N. ZWAAN, *Towards a more robust algorithm for computing the restricted singular value decomposition*. In preparation.
- [25] I. N. ZWAAN AND M. E. HOCHSTENBACH, *Generalized Davidson and multidirectional-type methods for the generalized singular value decomposition*, <https://arxiv.org/abs/1705.06120>.