



# A Kogbetliantz-type algorithm for the hyperbolic SVD

Vedran Novaković<sup>1</sup> · Sanja Singer<sup>2</sup>

Received: 5 December 2020 / Accepted: 9 September 2021 / Published online: 15 October 2021  
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

## Abstract

In this paper, a two-sided, parallel Kogbetliantz-type algorithm for the hyperbolic singular value decomposition (HSVD) of real and complex square matrices is developed, with a single assumption that the input matrix, of order  $n$ , admits such a decomposition into the product of a unitary, a non-negative diagonal, and a  $J$ -unitary matrix, where  $J$  is a given diagonal matrix of positive and negative signs. When  $J = \pm I$ , the proposed algorithm computes the ordinary SVD. The paper's most important contribution—a derivation of formulas for the HSVD of  $2 \times 2$  matrices—is presented first, followed by the details of their implementation in floating-point arithmetic. Next, the effects of the hyperbolic transformations on the columns of the iteration matrix are discussed. These effects then guide a redesign of the dynamic pivot ordering, being already a well-established pivot strategy for the ordinary Kogbetliantz algorithm, for the general,  $n \times n$  HSVD. A heuristic but sound convergence criterion is then proposed, which contributes to high accuracy demonstrated in the numerical testing results. Such a  $J$ -Kogbetliantz algorithm as presented here is intrinsically slow, but is nevertheless usable for matrices of small orders.

**Keywords** Hyperbolic singular value decomposition · Kogbetliantz algorithm · Hermitian eigenproblem · OpenMP multicore parallelization

**Mathematics Subject Classification (2010)** 65F15 · 65Y05 · 15A18

---

This work is dedicated to the memory of Sanja Singer.

✉ Vedran Novaković  
venovako@venovako.eu

<sup>1</sup> MFBDA project, 10000 Zagreb, Croatia

<sup>2</sup> Faculty of Mechanical Engineering and Naval Architecture, University of Zagreb, Ivana Lučića 5, 10000 Zagreb, Croatia

## 1 Introduction

The Kogbetliantz algorithm [14] is the oldest effective method discovered that computes the singular value decomposition (SVD) of a square matrix  $G$  as  $G = U \Sigma V^*$ , where  $U$  and  $V$  are unitary matrices, while  $\Sigma$  is a diagonal matrix with the non-negative diagonal elements, called the singular values, that are usually ordered non-increasingly, i.e.,  $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$  and  $\sigma_k \geq \sigma_\ell \geq 0$ , for  $1 \leq k < \ell \leq n$ , with  $n$  being the matrix order of  $G$ ,  $U$ ,  $\Sigma$ , and  $V$ .

In this paper, a Kogbetliantz-type algorithm for the hyperbolic singular value decomposition [26] (the HSVD in short) is developed and called the  $J$ -Kogbetliantz algorithm in the following. Since the Kogbetliantz-type algorithms operate only on square matrices, Definition 1.1 of the HSVD is restricted to such a case.

**Definition 1.1** Let  $J$  and  $G$  be two square matrices of order  $n$ , such that  $J$  is a diagonal matrix of signs,  $J = \text{diag}(\pm 1)$ , and  $G$  is a matrix over  $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$  with  $\text{rank}(G J G^*) = \text{rank}(G)$ . A decomposition of  $G$ ,

$$G = U \Sigma V^{-1},$$

such that  $U$  is a unitary matrix over  $\mathbb{F}$ ,  $V$  is a  $J$ -unitary matrix over  $\mathbb{F}$  with respect to  $J$  (i.e.,  $V^* J V = J$ , so  $V$  is hypernormal in the terminology of [4]), and  $\Sigma$  is a real diagonal matrix with a non-negative diagonal, is called the *hyperbolic singular value decomposition* of  $G$ . The diagonal elements of  $\Sigma$  are called the *hyperbolic singular values* of  $G$ , which are assumed to be ordered non-increasingly (non-decreasingly) in any, not necessarily contiguous, range of diagonal elements of  $\Sigma$  for which the corresponding range of diagonal elements of  $J$  contains only positive (negative) signs.

When  $J = \pm I_n$  the HSVD becomes the ordinary SVD, with  $V$  unitary, and the  $J$ -Kogbetliantz algorithm reduces to the ordinary Kogbetliantz algorithm for the SVD.

In Definition 1.1 the assumption that  $\text{rank}(G J G^*) = \text{rank}(G)$  ensures [36] that the HSVD of  $G$  exists, with a diagonal  $\Sigma$ . If the assumption does not hold, or if  $G$  is rectangular, for the  $J$ -Kogbetliantz algorithm  $G$  should be preprocessed by, e.g., the  $J$ -URV factorization [28, 29] to a square matrix  $G_0$  of order  $n_0 \leq n$ , such that  $\text{rank}(G_0 J_0 G_0^*) = \text{rank}(G_0)$ , i.e.,  $G = U_0^* \tilde{G}_0 V_0$ , where  $U_0$  is unitary,  $V_0$  is  $J$ -unitary,  $J_0 = \text{diag}(\pm 1)$  is of order  $n_0$ , and  $G_0 = \tilde{G}_0(1 : n_0, 1 : n_0)$ . Otherwise, let  $n_0 = n$ ,  $G_0 = G$ ,  $V_0 = U_0 = I_n$ , and  $J_0 = J$ .

Besides an obvious application as the main part of a method for solving a Hermitian indefinite eigenproblem [30, 31], the HSVD has also been used in various signal processing applications [26], and recently in a modified approach to the Kalman filtering, especially in the ill-conditioned case [15–17]. The HSVD can be efficiently and accurately computed by the *one-sided* blocked Jacobi-type algorithms [8, 9], even on the massively parallel architectures [21]. A *two-sided*, Kogbetliantz-type method proposed here is not intended to outperform the existing one-sided HSVD algorithms. Instead, it is meant to showcase the tools (e.g., a careful dynamic parallel ordering and a sound convergence criterion) designed to cope with the perils of non-unitary transformations, that empirically happen to be a greater risk for numerical

stability of the two-sided methods—here, for the hyperbolic, but possibly also applicable to the more general, orthosymmetric [18] SVD—than for that of the one-sided algorithms.

The ordinary Kogbetliantz algorithm usually transforms an upper triangular matrix  $R$  (see, e.g., [10]), resulting from the preprocessing of  $G$  by the QR factorization, in order to simplify the transformations and lower their number. If a particular cyclic pivot ordering is applied, at the end of the first cycle a lower triangular matrix is obtained, while after the subsequent cycle the iteration matrix becomes upper triangular again. Unfortunately, a simple generalization of the Kogbetliantz algorithm to a hyperbolic one (by introducing the hyperbolic transformations but leaving the pivot strategy intact) can fail, since even a single hyperbolic transformation from the right, with a sufficiently high condition number, can cause an excessive growth of the off-diagonal elements of the iteration matrix and ruin the convergence of the algorithm. A different, dynamic (data-dependent) pivot strategy is therefore needed to keep the sometimes unavoidable growth of the off-diagonal norm in check. As a consequence, the iteration matrix cannot be assumed to have or retain any particular structure.

The  $J$ -Kogbetliantz algorithm computes the HSVD of  $G_0$  in a sequence of transformations, infinite in general but cut off to a finite leading part when computing in machine precision and a convergence criterion is met, as

$$U^*G_0V \approx \Sigma, \quad U^* = U_N^*U_{N-1}^* \cdots U_1^*, \quad V = V_1V_2 \cdots V_N,$$

where, for each  $k$ ,  $1 \leq k \leq N$ ,  $U_k^*$  and  $V_k$  are the embeddings of the  $2 \times 2$  transformations  $\widehat{U}_k^*$  and  $\widehat{V}_k$ , respectively, into  $I_{n_0}$ , that are applied to the iteration matrix  $G_{k-1}$ , forming a sequence of the iteration matrices

$$G_1, G_2, \dots, G_N \approx \Sigma,$$

where  $G_k = U_k^*G_{k-1}V_k$ . If  $G_0$  is already in the form required of  $\Sigma$ , then  $V = U = I_{n_0}$ ,  $\Sigma = G_0$ , and no transformations take place. Else,  $U_k^*$  is unitary (orthogonal in the real case), while  $V_k$  is  $J_0$ -unitary ( $J_0$ -orthogonal in the real case). Each  $\widehat{U}_k^*$  is unitary (orthogonal), and  $\widehat{V}_k$  is  $\widehat{J}_k$ -unitary ( $\widehat{J}_k$ -orthogonal), where  $\widehat{J}_k = \text{diag}(J_0(p_k, p_k), J_0(q_k, q_k))$  for the  $k$ th pivot index pair  $(p_k, q_k)$ ,  $1 \leq p_k < q_k \leq n_0$ . The decomposition is finalized by forming  $U = (U^*)^* = U_1U_2 \cdots U_N$ , while by multiplying  $V^*J_0V = J_0$  from the left by  $J_0$ , and noting that  $J_0^2 = I_{n_0}$ , it follows that  $V^{-1} = J_0V^*J_0$ .

In the  $k$ th step the elements  $(1, 1)$ ,  $(2, 1)$ ,  $(1, 2)$ , and  $(2, 2)$  of  $\widehat{U}_k^*$  (and  $\widehat{V}_k$ ) are embedded into  $U_k^*$  (and  $V_k$ ) at the pivot positions  $(p_k, p_k)$ ,  $(q_k, p_k)$ ,  $(p_k, q_k)$ , and  $(q_k, q_k)$ , respectively, where the pivot indices are chosen such that for the pivot matrix  $\widehat{G}_{k-1}$ ,

$$\widehat{G}_{k-1} = \begin{bmatrix} G_{k-1}(p_k, p_k) & G_{k-1}(p_k, q_k) \\ G_{k-1}(q_k, p_k) & G_{k-1}(q_k, q_k) \end{bmatrix}, \tag{1.1}$$

holds at least one of

$$\widehat{G}_{k-1}(2, 1) \neq 0, \quad \widehat{G}_{k-1}(1, 2) \neq 0, \quad \widehat{G}_{k-1}(1, 1) \notin \mathbb{R}_0^+, \quad \widehat{G}_{k-1}(2, 2) \notin \mathbb{R}_0^+.$$

When  $\widehat{J}_k(1, 1) = \widehat{J}_k(2, 2)$  (such a case is called *trigonometric* in the following, while the other case is called *hyperbolic*),  $\widehat{G}_{k-1}$  is also a transformation candidate

if it is diagonal, with the non-negative diagonal elements such that  $\widehat{G}_{k-1}(1, 1) < \widehat{G}_{k-1}(2, 2)$  if  $\widehat{J}_k = I_2$ , or  $\widehat{G}_{k-1}(1, 1) > \widehat{G}_{k-1}(2, 2)$  if  $\widehat{J}_k = -I_2$ . This is summarized in Definition 1.2.

**Definition 1.2** In a step  $k \geq 1$ , let  $p_k$  and  $q_k$ ,  $1 \leq p_k < q_k \leq n_0$ , be such indices for which  $\widehat{G}_{k-1}$  from (1.1) does *not* have a form required of a matrix of hyperbolic singular values, as specified in Definition 1.1. Then,  $\widehat{G}_{k-1}$  is called a transformation candidate or a pivot (sub)matrix. If such indices do not exist, there are no transformation candidates in the  $k$ th step. If more than one index pair satisfies the defining condition, one pivot index pair is selected according to the chosen pivot strategy.

If the step index is omitted, then for the pivot index pair  $(p, q)$  the transformation matrices ( $\widehat{U}^*$  and  $\widehat{V}$ ), the pivot matrix ( $\widehat{G}$ ), and the corresponding matrices of signs ( $\widehat{J}$ ) and hyperbolic singular values ( $\widehat{\Sigma}$ ) can be written as

$$\widehat{U}^* = \begin{bmatrix} \widehat{u}_{pp} & \widehat{u}_{qp} \\ \widehat{u}_{pq} & \widehat{u}_{qq} \end{bmatrix}, \quad \widehat{V} = \begin{bmatrix} v_{pp} & v_{pq} \\ v_{qp} & v_{qq} \end{bmatrix}, \quad \widehat{G} = \begin{bmatrix} g_{pp} & g_{pq} \\ g_{qp} & g_{qq} \end{bmatrix}, \quad \widehat{J} = \begin{bmatrix} j_{pp} & 0 \\ 0 & j_{qq} \end{bmatrix}; \quad \widehat{\Sigma} = \begin{bmatrix} \sigma_{pp} & 0 \\ 0 & \sigma_{qq} \end{bmatrix}.$$

With  $\widehat{G}$  given,  $\widehat{U}^*$  and  $\widehat{V}$  are sought for, such that  $\widehat{U}^* \widehat{G} \widehat{V} = \widehat{\Sigma}$ , with  $\sigma_{pp}$  and  $\sigma_{qq}$  being the hyperbolic singular values of  $\widehat{G}$ , and  $\widehat{V}^* \widehat{J} \widehat{V} = \widehat{J}$ .

If there are no transformation candidates in the  $k$ th step, or if the convergence criterion is satisfied, the algorithm stops successfully, with (an approximation of) the HSVD of  $G$  found. Otherwise,  $\widehat{U}_k^*$  and  $\widehat{V}_k$  are computed for one, suitably chosen transformation candidate of  $G_{k-1}$ , and applied to the pivot rows  $p_k$  and  $q_k$  from the left ( $\widehat{U}_k^*$ ) and the pivot columns  $p_k$  and  $q_k$  from the right ( $\widehat{V}_k$ ), in that order for determinacy, to form  $G_k$ . The process is then repeated in the step  $k + 1$ .

Note that several, but at most  $\lfloor n_0/2 \rfloor$  successive steps can be grouped to form a multi-step  $\mathbf{k}$ , where  $0 \leq \|\mathbf{k}\| \leq \lfloor n_0/2 \rfloor$  is the chosen length of  $\mathbf{k}$ , if and only if  $\{p_k, q_k\} \cap \{p_l, q_l\} = \emptyset$  for all  $k \neq l$  such that  $\{k, l\} \subseteq \mathbf{k}$ , i.e., all basic steps within a multi-step can be performed in *parallel* (with  $t \geq 1$  parallel tasks at hand). The number of basic steps may vary from one multi-step to another if, e.g., not enough transformation candidates are available, and may reach zero when no further transformations are possible. With  $n_0 = 6$ , e.g., the multi-step  $\mathbf{1}$  might be  $\{1, 2, 3\}$ ,  $\mathbf{2}$  might be  $\{4, 5\}$ ,  $\mathbf{3}$  might be  $\{6, 7, 8\}$ , continuing until some  $\mathbf{N}$  at which the algorithm halts.

The parallel application of the  $2 \times 2$  transformations has to take into account that all transformations from one side (e.g., the row transformations from the left) have to precede any transformation from the other side (e.g., the column transformations from the right). However, all transformations from the same side can proceed concurrently, and the choice of the first side to be transformed is arbitrary.

To fully describe the  $J$ -Kogbetliantz algorithm, it therefore suffices to specify:

1. a method for computing the HSVD of the pivot matrices of order two,
2. the details of performing the row and the column transformations,
3. a pivot strategy that selects the transformation candidate(s) in a (multi-)step, and
4. a convergence criterion that halts the execution.

The above list guides the organization of this paper as follows. The first item is covered in Section 2, the second one in Section 3, the third one in Section 4, and the

last one in Section 5, containing an overview of the algorithm. The numerical testing results are summarized in Section 6, and the paper is concluded with some remarks on the future work in Section 7. The Appendix contains several technical proofs.

## 2 Computing the HSVD of matrices of order two

In the ordinary Kogbetliantz algorithm, it is preferred that the pivot matrices throughout the process remain (upper or lower) triangular [5, 10] under a cyclic pivot strategy: a serial (e.g., the row or the column cyclic, with  $G_0$  triangular) or a parallel one (e.g., the modulus strategy, with  $G_0$  preprocessed into the butterfly form [11])—a fact relied upon for a simple and accurate computation of the transformation parameters [7, 19, 20]. However, the  $J$ -Kogbetliantz algorithm employs a strategy that has no periodic (cyclic) pattern, so no particular form of the pivot matrices can be guaranteed and no special form of  $G_0$  is presumed. Depending on  $\hat{J}$  and its own structure, a transformation candidate  $\hat{G}$  might have to be preprocessed into a real, triangular form with non-negative elements that is suitable for a novel, numerically stable computation of the transformation matrices  $\hat{U}^*$  and  $\hat{V}$  and the hyperbolic singular values in  $\hat{\Sigma}$ .

### 2.1 The $\hat{J}$ -UTV factorization of $\hat{G}$

The preprocessing of  $\hat{G}$  starts with checking if  $\hat{G}$  is diagonal, because this information will affect the convergence criterion (see Section 5.1) and can simplify an optimized implementation of the  $2 \times 2$  HSVD. Then,  $\hat{G}$  is transformed into a real, triangular matrix  $\hat{T}$  with non-negative elements by the  $\hat{J}$ -UTV factorization that preserves the hyperbolic singular values. The  $\hat{J}$ -UTV factorization is similar to the URV [34] factorization, with the form

$$\check{U}^* \hat{G} \check{V} = \hat{T},$$

where  $\check{U}$  is unitary,  $\check{V}$  is  $\hat{J}$ -unitary, and  $\hat{T}$  is either upper or lower triangular with real, non-negative elements. Moreover,  $\hat{T}$  is upper triangular in the trigonometric case, and  $\hat{t}_{11}^2 \geq \hat{t}_{12}^2 + \hat{t}_{22}^2$ . In the hyperbolic case  $\hat{T}$  is either upper triangular, as described, or lower triangular, with  $\hat{t}_{22}^2 \geq \hat{t}_{21}^2 + \hat{t}_{11}^2$ , depending on whether the first or the second column of  $\hat{G}$  has a greater Frobenius norm, respectively. The  $2 \times 2$   $\hat{J}$ -UTV factorization always exists and is uniquely determined, as demonstrated by Algorithm 1, which, given  $\hat{G}$ , operates on a scaled matrix  $\hat{G}_0 = 2^s \hat{G}$  to avoid the possibility of floating-point overflows. Determination of  $s$  is described in Section 2.1.1, with the only assumption that all (components of) the elements of  $\hat{G}$  are finite. Assume  $\arg(0) = 0$  for simplicity, and let  $P_2$  be the  $2 \times 2$  permutation matrix  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ . Let Boolean flags—those being only true ( $\top$ ) or false ( $\perp$ )—be denoted by small caps (e.g.,  $D$ ,  $H$ , and  $C$ , indicating whether  $\hat{G}_0$  is diagonal, whether the hyperbolic or the trigonometric case is considered, and whether the columns of  $\hat{G}_0$  have been swapped, respectively).

**Algorithm 1** The  $\widehat{J}$ -UTV factorization of a complex  $2 \times 2$  matrix  $\widehat{G}_0$ .

```

if  $\widehat{G}_0$  is diagonal then  $D = \top$  else  $D = \perp$ ; // specifically,  $D = \top$  if  $\widehat{G}_0 = 0$ 
if  $\widehat{J} = I_2$  or  $\widehat{J} = -I_2$  then  $H = \perp$  else  $H = \top$ ; //  $\perp$ : trigonometric,  $\top$ : hyperbolic
case
// STAGE 1: column pivoting of  $\widehat{G}_0$ 
if  $\left\| \begin{bmatrix} (\widehat{G}_0)_{11} \\ (\widehat{G}_0)_{21} \end{bmatrix} \right\|_F < \left\| \begin{bmatrix} (\widehat{G}_0)_{12} \\ (\widehat{G}_0)_{22} \end{bmatrix} \right\|_F$  then  $C = \top$  and  $\check{V}_1 = P_2$  else  $C = \perp$  and  $\check{V}_1 = I_2$ ;
// if  $H \wedge C$ :  $\check{V}_1$  is not  $\widehat{J}$ -unitary but it is canceled by  $\check{V}_5$  in
stage 5L
 $\widehat{G}_1 = \widehat{G}_0 \check{V}_1$ ,  $\check{V} := \check{V}_1$ ; // start accumulating  $\check{V}$ 
// STAGE 2: make the first column of  $\widehat{G}_1$  real and non-negative
 $\check{U}_2^* = \begin{bmatrix} e^{-\arg(\widehat{G}_1)_{11}} & 0 \\ 0 & e^{-\arg(\widehat{G}_1)_{21}} \end{bmatrix}$ ; // applied to get  $(\widehat{G}_2)_{11} = |(\widehat{G}_1)_{11}|$  &
 $(\widehat{G}_2)_{21} = |(\widehat{G}_1)_{21}|$ 
 $\widehat{G}_2 = \check{U}_2^* \widehat{G}_1$ ,  $\check{U}^* := \check{U}_2^*$ ; // start accumulating  $\check{U}^*$ 
// STAGE 3: row pivoting of  $\widehat{G}_2$ 
if  $(\widehat{G}_2)_{11} < (\widehat{G}_2)_{21}$  then  $R = \top$  and  $\check{U}_3^* = P_2$  else  $R = \perp$  and  $\check{U}_3^* = I_2$ ;
 $\widehat{G}_3 = \check{U}_3^* \widehat{G}_2$ ,  $\check{U}^* := \check{U}_3^* \check{U}^*$ ; //  $(\widehat{G}_3)_{11} \geq (\widehat{G}_3)_{21} \geq 0$ 
// STAGE 4: the QR factorization of  $\widehat{G}_3$ 
if  $\neg D$  then // not D implies  $(\widehat{G}_3)_{11} > 0$ 
|  $\tan \phi = (\widehat{G}_3)_{21} / (\widehat{G}_3)_{11}$ ,  $\cos \phi = 1 / \sqrt{1 + \tan^2 \phi}$ ; //  $0 \leq \phi \leq \pi/4$ 
|  $\check{U}_4^* = \cos \phi \begin{bmatrix} 1 & \tan \phi \\ -\tan \phi & 1 \end{bmatrix}$ ; // the Givens rotation  $\check{U}_4^*$  is  $Q^*$  in the QR
else // the QR factorization of a diagonal matrix is trivial
|  $\check{U}_4^* = I_2$ ; //  $\phi = 0$ 
end if
 $\widehat{G}_4 = \check{U}_4^* \widehat{G}_3$ ,  $\check{U}^* := \check{U}_4^*$ ; //  $\widehat{G}_4$  is upper triangular
if  $H \wedge C$  then // this pipeline ends with  $\widehat{T}$  lower triangular
// STAGE 5L: swap the columns of  $\widehat{G}_4$  to cancel  $\check{V}_1$ 
 $\check{V}_3 = P_2$ ,  $\widehat{G}_5 = \widehat{G}_4 \check{V}_3$ ,  $\check{V} := I_2 (= \check{V} \check{V}_3)$ ; // since  $P_2 = P_2^*$ 
// STAGE 6L: swap the rows of  $\widehat{G}_5$ 
 $\check{U}_6^* = P_2$ ,  $\widehat{G}_6 = \check{U}_6^* \widehat{G}_5$ ,  $\check{U}^* := \check{U}_6^* \check{U}^*$ ; //  $\widehat{G}_6$  is lower triangular
// STAGE 7L: make  $(\widehat{G}_6)_{21}$  real and non-negative
 $\check{V}_7 = \begin{bmatrix} e^{-\arg(\widehat{G}_6)_{21}} & 0 \\ 0 & 1 \end{bmatrix}$ ; // applied to get  $(\widehat{G}_7)_{21} = |(\widehat{G}_6)_{21}|$ 
 $\widehat{G}_7 = \widehat{G}_6 \check{V}_7$ ,  $\check{V} := \check{V}_7$ ; //  $\check{V}_7$  is both unitary and  $\widehat{J}$ -unitary
// STAGE 8L: make  $(\widehat{G}_7)_{11}$  real and non-negative
 $\check{U}_8^* = \begin{bmatrix} e^{-\arg(\widehat{G}_7)_{11}} & 0 \\ 0 & 1 \end{bmatrix}$ ; // applied to get  $(\widehat{G}_8)_{11} = |(\widehat{G}_7)_{11}|$ 
 $\widehat{G}_8 = \check{U}_8^* \widehat{G}_7$ ,  $\check{U}^* := \check{U}_8^* \check{U}^*$ ; //  $\widehat{T} = \widehat{G}_8$ 
else // this pipeline ends with  $\widehat{T}$  upper triangular
// STAGE 5U: make  $(\widehat{G}_4)_{12}$  real and non-negative
 $\check{V}_5 = \begin{bmatrix} 1 & 0 \\ 0 & e^{-\arg(\widehat{G}_4)_{12}} \end{bmatrix}$ ; // applied to get  $(\widehat{G}_5)_{12} = |(\widehat{G}_4)_{12}|$ 
 $\widehat{G}_5 = \widehat{G}_4 \check{V}_5$ ,  $\check{V} := \check{V} \check{V}_5$ ; //  $\check{V}_5$  is both unitary and  $\widehat{J}$ -unitary
// STAGE 6U: make  $(\widehat{G}_5)_{22}$  real and non-negative
 $\check{U}_6^* = \begin{bmatrix} 1 & 0 \\ 0 & e^{-\arg(\widehat{G}_5)_{22}} \end{bmatrix}$ ; // applied to get  $(\widehat{G}_6)_{22} = |(\widehat{G}_5)_{22}|$ 
 $\widehat{G}_6 = \check{U}_6^* \widehat{G}_5$ ,  $\check{U}^* := \check{U}_6^* \check{U}^*$ ; //  $\widehat{T} = \widehat{G}_6$ 
end if

```

**Note:** All transitions  $\widehat{G}_l \rightarrow \widehat{G}_{l+1}$  actually happen in-place by transforming the working matrix.

When  $\widehat{T}$  is upper triangular, Algorithm 1 is a refinement of the URV factorization, as described in detail in [22] for the trigonometric case, where  $R = \widehat{T}$  is real and non-negative, with a prescribed relation among the magnitudes of its elements. The  $2 \times 2$   $\widehat{J}$ -UTV factorization guarantees that  $\check{V}$  is both unitary and  $\widehat{J}$ -unitary in all cases. When  $\widehat{G}$  is real, a multiplication of an element  $\hat{g}$  of the working matrix by  $e^{-\arg \hat{g}}$  amounts to a multiplication by  $\text{sign } \hat{g}$ , without any rounding error.

In brief, Algorithm 1 consists of two possible pipelines, having in common the first four stages. The longer pipeline has eight, and the shorter one six stages in total.

In a stage  $s \geq 1$ , the working matrix and either  $\check{U}^*$  or  $\check{V}$  are transformed in place, either by a left multiplication by  $\widehat{U}_s^*$ , or by a right multiplication by  $\widehat{V}_s$ , respectively. The working matrix thus progresses from the state  $\widehat{G}_{s-1}$  to  $\widehat{G}_s$ , until it is eventually reduced to  $\widehat{T}$ , at which point  $\check{U}^*$  and  $\check{V}$  are considered final, i.e., the  $\widehat{J}$ -UTV factorization is computed. Initially,  $\check{U}^*$  and  $\check{V}$  are set to identities, and D and H are determined.

In stage 1, if the column pivoting of  $\widehat{G}_0$  is needed (i.e., its first column is of lower Frobenius norm than its second), C is set to  $\top$  and  $\check{V}_1$  to  $P_2$ , else they become  $\perp$  and  $I_2$ , respectively. Note that  $\check{V}_1$  is not  $\widehat{J}$ -unitary if H and C are  $\top$ , but in that case  $\check{V}_1$  will be canceled by  $P_2 = P_2^*$ , with no intervening right transformations, in stage 5L.

In stage 2, the first column of  $\widehat{G}_1$  is made real and non-negative by  $\check{U}_2^*$ , effectively setting the elements of the first column of  $\widehat{G}_2$  to the magnitudes of those from  $\widehat{G}_1$ . Then, it is possible to compare them in stage 3. If the first is less than the second, the rows of  $\widehat{G}_2$  are swapped by  $\check{U}_3^*$ . In stage 4 (the final one shared among the two pipelines) the QR factorization of  $\widehat{G}_3$  is performed. If  $\widehat{G}_0$  was diagonal,  $\widehat{G}_3$  will come out as such from the previous stage, and the factorization is trivial. Else, the real Givens rotation  $\check{U}_4^*$  (i.e.,  $Q^*$ ) is computed from the elements of the first column of  $\widehat{G}_3$ . Due to the row pivoting, the rotation’s angle  $\phi$  lies between 0 and  $\pi/4$ . Now,  $R$  from the factorization, i.e.,  $\widehat{G}_4$ , is upper triangular, with its second column possibly complex, while  $(\widehat{G}_4)_{11}$  is by construction real and of the largest magnitude in  $\widehat{G}_4$ .

The remaining stages of the longer pipeline, i.e., the one in which H and C are  $\top$ , are described first. In stage 5L  $\check{V}_1 = P_2$  is canceled by  $\check{V}_5 = P_2$ , i.e., the columns of  $\widehat{G}_4$  are swapped to obtain  $\widehat{G}_5$ , which in stage 6L gets its rows swapped by  $\check{U}_6^*$  to become the lower triangular  $\widehat{G}_6$ . Stage 5L is necessary to make the current  $\check{V}$   $\widehat{J}$ -unitary (identity) again. Stage 6L brings the working matrix into a more “familiar” (i.e., lower triangular) form. The working matrix in this pipeline remains lower triangular until the end. To emphasize that, “L” is appended to the stage identifiers. To make the working matrix real,  $\check{V}_7$  (obviously unitary but also  $\widehat{J}$ -unitary, since  $\check{V}_7^* \widehat{J} \check{V}_7 = \widehat{J}$ , where  $\widehat{J}$  has two opposite signs on its diagonal) is applied to it in stage 7L, to turn  $(\widehat{G}_6)_{21}$  into its magnitude. Finally, in stage 8L,  $(\widehat{G}_7)_{11}$  is turned into its magnitude by  $\check{U}_8^*$ .

The shorter pipeline, followed when H or C are  $\perp$ , keeps the working matrix upper triangular, and thus, the stage identifiers have “U” appended to them. Stage 5U turns  $(\widehat{G}_4)_{12}$  into its magnitude by (similarly as above, both unitary and  $\widehat{J}$ -unitary)  $\check{V}_5$ , while stage 6U completes the pipeline by turning  $(\widehat{G}_5)_{22}$  into its magnitude by  $\check{U}_6^*$ .

Figure 1 illustrates the possible progressions of a non-diagonal complex matrix  $\widehat{G}_0$  through the stages of Algorithm 1 in both its pipelines (but only one pipeline is taken with any given input), and as if all optional transformations take place (the column and/or the row pivoting may not actually happen). A diagonal matrix admits a simplified procedure, evident enough for its details to be omitted here for brevity. The state of the working matrix before the operations of each stage commence is shown, while the transformation matrices are implied by the effects of a particular stage.

The numerical issues of computing in the machine precision (in Fortran) are dealt with in Section 2.1.1 for the column norms, in Section 2.1.2 for the polar form of a complex number, and in Section 2.1.3 for the rotation-like transformations. It is assumed that the floating-point arithmetic is not trapping on an exception, that the rounding to nearest (with tie to even) is employed, that the gradual underflow and the fused multiply-add (fma) operation are available, that  $\text{HYPOT}(a, b) = \sqrt{a^2 + b^2}$  intrinsic causes no undue overflow, and that  $\text{HYPOT}(a, b) = 0 \iff a = b = 0$ .


Algorithm 1 is conditionally reproducible, i.e., its results on any given input are bitwise identical in every floating-point environment behaving as described above and employing the same datatype, provided that the HYPOT intrinsic is reproducible.

### 2.1.1 Determination of the scaling parameter s

For a complex number  $z$ ,  $|z|$  can overflow in floating-point arithmetic. The magnitude of  $z$  can also pose a problem, e.g., in computing of the norms of the columns of  $\widehat{G}$  in stage 1 of Algorithm 1 for the column pivoting in the QR factorization. There,

$$\| [a \ b]^T \|_F = \sqrt{|a|^2 + |b|^2} = \text{HYPOT}(\text{HYPOT}(\text{Re}a, \text{Im}a), \text{HYPOT}(\text{Re}b, \text{Im}b)) = c,$$

so it has to be ensured, by an appropriate prescaling of  $\widehat{G}$ , that  $|a|$ ,  $|b|$ , and  $c$  do not overflow. A similar concern is addressed in the xLARTG LAPACK [1] routines, but for efficiency it is advisable to avoid the overhead of function calls in the innermost



**Fig. 1** A schematic representation of the  $\widehat{J}$ -UTV factorization pipelines of a complex  $2 \times 2$  matrix  $\widehat{G}_0$ . The upper pipeline is followed in the trigonometric case, or in the hyperbolic case when no column pivoting was performed in stage 1, resulting in an upper triangular  $\widehat{T}$  after six transformations. Otherwise, the lower pipeline is followed from stage 5 onwards, resulting in a lower triangular  $\widehat{T}$  after eight transformations of the working matrix. The symbols with a black background stand for the complex elements, while those with a white background denote the real non-negative ones. The circled numbers represent the transient values, while the framed operators are reserved for the final ones, such that  $\square^2 \geq \odot^2 + \ast^2$



computational parts. At the same time, to avoid computing with the subnormal (components of) numbers,  $\widehat{G}$  could be upscaled when no overflow can occur. Therefore, to efficiently mitigate as many overflow and underflow issues as practicable, either the following prescaling of  $\widehat{G}$ , and the corresponding postscaling of the resulting  $\widehat{\Sigma}$ , should be employed, or, without any scaling needed, the entire computation of the  $2 \times 2$  HSVD should be performed in a floating-point datatype at least as precise as the type of (the components of) the elements of  $\widehat{G}$ , but with a wider exponent range (e.g., in the Intel’s 80-bit, hardware supported extended precision datatype).

Down- and up-scalings of  $\widehat{G}$  can be made (almost) exact, without any loss of precision in the vast majority of cases, by decreasing, or respectively increasing, the exponents of (the components of) the elements only, i.e., by multiplying the values by  $2^s$  with  $s \in \mathbb{Z}$ , leaving the significands intact (except for a subnormal value and  $s < 0$ , when the  $|s|$  lowest bits of the significand—some of which may not be zero—perish), as explained in [22, subsection 2.1] and summarized here in Fortran parlance.

Let  $\eta(a) = \text{EXPONENT}(a)$  for a floating-point value  $a$ , and let  $\nu = \eta(\Omega)$ , where  $\Omega = \text{HUGE}(0D0)$  is the largest finite DOUBLE PRECISION value. For  $a$  finite, let

$$\chi(a) = \nu - \eta(\max\{|a|, \omega\}) - 2$$

be a measure of how much, in terms of its exponent,  $a$  has to be downscaled, or how much it can be safely upscaled. Here and in the following,  $\omega$  is the smallest non-zero positive subnormal floating-point number,<sup>1</sup> while  $\widehat{\omega} = \text{TINY}(0D0)$  is the smallest positive normal DOUBLE PRECISION value. Two sets of such exponent differences,

$$\begin{aligned} E_{\Re} &= \{\chi(\text{Re}\widehat{g}_{11}), \chi(\text{Re}\widehat{g}_{21}), \chi(\text{Re}\widehat{g}_{12}), \chi(\text{Re}\widehat{g}_{22})\}, \\ E_{\Im} &= \{\chi(\text{Im}\widehat{g}_{11}), \chi(\text{Im}\widehat{g}_{21}), \chi(\text{Im}\widehat{g}_{12}), \chi(\text{Im}\widehat{g}_{22})\}, \end{aligned}$$

are then formed, and the scaling parameter  $s$  is taken to be the minimum of  $E_{\Re} \cup E_{\Im}$ ,

$$s = \min\{\min E_{\Re}, \min E_{\Im}\}. \tag{2.1}$$

Each component  $c$  of the elements of  $\widehat{G}$  is then scaled as  $c_0 = \text{SCALE}(c, s)$  to get  $\widehat{G}_0$ .

Backscaling of the computed hyperbolic singular values by  $2^{-s}$  can cause overflows or underflows (to the subnormal range, or even zero), but they are unavoidable if the scaling should be transparent to the users of the  $2 \times 2$  HSVD. A remedy, discussed in [22], would be to keep the scales in addition to the values to which they apply, in a separate array. Alternatively, computing in a wider datatype avoids both scaling and such issues, at a cost of a slower and not easily vectorizable arithmetic.

In brief, the scaling parameter  $s$  from (2.1) is chosen in such a way that:

1. no column norm of a scaled  $2 \times 2$  real or complex matrix can overflow, and
2. no ordinary singular value of a scaled matrix can overflow [22, Theorem 1], and
3. as many subnormal (components of the) matrix elements as possible are brought up to the normal range, for a faster and maybe more accurate computation.

---

<sup>1</sup>For  $a = 0$ ,  $\eta(a) = 0$  instead of a huge negative integer, so taking  $\max\{|a|, \omega\}$  filters out such  $a$ .

The following Example 2.1 illustrates several reasons for and issues with the scaling.

*Example 2.1* Assume  $\widehat{J}^{[i]} = I_2$ , and consider the following real  $\widehat{G}^{[i]}$  for  $1 \leq i \leq 3$ ,

$$\widehat{G}^{[1]} = \begin{bmatrix} \Omega & \Omega/4 \\ \Omega/8 & \Omega/2 \end{bmatrix}, \quad \widehat{G}^{[2]} = \begin{bmatrix} \Omega/16 & \widehat{\omega}/4 \\ \widehat{\omega}/2 & \widehat{\omega}/8 \end{bmatrix}, \quad \widehat{G}^{[3]} = \begin{bmatrix} \omega & 0 \\ 0 & \Omega \end{bmatrix},$$

alongside their scaled versions  $\widehat{G}_0^{[i]} = 2^{s^{[i]}} \widehat{G}^{[i]}$ ,

$$\widehat{G}_0^{[1]} = \begin{bmatrix} \Omega/4 & \Omega/16 \\ \Omega/32 & \Omega/8 \end{bmatrix}, \quad \widehat{G}_0^{[2]} = \begin{bmatrix} \Omega/4 & \widehat{\omega} \\ 2\widehat{\omega} & \widehat{\omega}/2 \end{bmatrix}, \quad \widehat{G}_0^{[3]} = \begin{bmatrix} \omega/4 = 0 & 0 \\ 0 & \Omega/4 \end{bmatrix},$$

with the scaling parameters, computed from (2.1), being  $s^{[1]} = s^{[3]} = -2$  and  $s^{[2]} = 2$ .

The singular values of  $\widehat{G}^{[1]}$ , computed symbolically, are  $\Omega(\sqrt{145} \pm 5)/16$ , and obviously the larger one would overflow without the scaling. In  $\widehat{G}^{[2]}$  there are three subnormal values. Two of them (the off-diagonal ones) get raised into the normal range, while the diagonal one stays subnormal, after the scaling. Finally,  $\widehat{G}^{[3]}$  is a pathological case of a non-singular matrix that is as ill-conditioned as possible, which becomes singular after the scaling, since  $\omega/4 = 0$  in the rounding-to-nearest mode.

## 2.1.2 Accurate floating-point computation of the polar form of a complex number

Expressing  $z$  as  $|z|e^{i \arg z}$  requires a reliable way of computing  $e^{i \arg z}$ , where

$$e^{i \arg z} = \cos(\arg z) + i \cdot \sin(\arg z).$$

One such method, presented in [22, eq. (1)], is summarized as follows. It is assumed that  $z$  has already been scaled as described in Section 2.1.1, so  $|z|$  cannot overflow.

Let `MIN` be a floating-point minimum function,<sup>2</sup> with a property that if its first argument is a NaN and the second is not, the result is the second argument. For the `IEEE_MIN_NUM` Fortran 2018 [13] intrinsic such a behavior is guaranteed, but is also present with the `MIN` intrinsic of the recent Intel Fortran compilers, and several others. The same property will be required for a maximum function `MAX` later on. Define

$$\cos(\arg z) = \text{MIN} \left( \frac{|\text{Re}z|}{|z|}, 1 \right) \cdot \text{sign}(\text{Re}z), \quad \sin(\arg z) = \frac{\text{Im}z}{\max\{|z|, \omega\}}.$$

When  $|z| = 0$ , `MIN` ensures the correct value of  $\cos(\arg z)$ , because  $\text{MIN}(0/0, 1) = 1$ , while in that case taking the maximum in the denominator of  $\sin(\arg z)$  makes the whole quotient well defined, since  $0/\omega = 0$ . When  $|z| > 0$ , it holds  $\max\{|z|, \omega\} = |z|$ .

## 2.1.3 Performing the rotation-like transformations in floating-point

Some real-valued quantities throughout the paper are obtained by the expressions of the form  $a \cdot b + c$ , and should be computed by a single `fma`, i.e., with a single rounding. One such example is  $\cos \phi = 1/\sqrt{\tan \phi \cdot \tan \phi + 1}$  in stage 4 of Algorithm 1.

<sup>2</sup>For example, the `minimumNumber` operation of the IEEE 754-2019 standard [12, section 9.6].

Regrettably, there is no widespread hardware support for the correctly rounded reciprocal square root (i.e.,  $1/\sqrt{x}$ ) operation. Therefore, the computation of a cosine involves at least three roundings: from the fused multiply-add operation to obtain the argument of the square root, from the square root itself, and from taking the reciprocal value. Instead of the cosines, the respective secants can be computed with two roundings, without taking the reciprocals. Then, in all affected formulas the multiplications by the cosines can be replaced by the respective divisions by the secants. Such an approach is slower than the usual one, but more accurate in the worst case.

There is no standardized way to accurately compute  $d = a \cdot b + c$  with some or all values being complex. However, the real fused multiply-add operation can be employed as in the CUDA [24, `cuComplex.h` header] implementation of the complex arithmetic for an efficient, accurate, and reproducible computation of  $d$  as

$$\begin{aligned} \operatorname{Re}(d) &= \operatorname{fma}(\operatorname{Re}(a), \operatorname{Re}(b), \operatorname{fma}(-\operatorname{Im}(a), \operatorname{Im}(b), \operatorname{Re}(c))), \\ \operatorname{Im}(d) &= \operatorname{fma}(\operatorname{Re}(a), \operatorname{Im}(b), \operatorname{fma}(\operatorname{Im}(a), \operatorname{Re}(b), \operatorname{Im}(c))), \end{aligned}$$

when  $a, b,$  and  $c$  are complex. Otherwise, when  $a$  is real,

$$\operatorname{Re}(d) = \operatorname{fma}(a, \operatorname{Re}(b), \operatorname{Re}(c)), \quad \operatorname{Im}(d) = \operatorname{fma}(a, \operatorname{Im}(b), \operatorname{Im}(c)),$$

and when  $a$  is non-zero and purely imaginary,

$$\operatorname{Re}(d) = \operatorname{fma}(-\operatorname{Im}(a), \operatorname{Im}(b), \operatorname{Re}(c)), \quad \operatorname{Im}(d) = \operatorname{fma}(\operatorname{Im}(a), \operatorname{Re}(b), \operatorname{Im}(c)).$$

Such operations can be implemented explicitly by the `IEEE_FMA` Fortran 2018 [13] intrinsic, or implicitly, by relying on the compiler to emit the appropriate `fma` instructions. By an abuse of notation,  $\operatorname{fma}(a, b, c)$  in the following stand for both the real-valued and the above complex-valued operations, depending on the context.

With  $a$  and  $b$  complex, a multiplication  $a \cdot b$  can be expressed as  $\operatorname{fma}(a, b, c)$  with  $c = 0$  and implemented as such, in a reproducible way, but simplified by converting all real `fma` operations involving a component of  $c$  to real multiplications.

A plane rotation, if it is to be applied from the left, can be written as (see, e.g., [6])

$$\begin{bmatrix} \cos \phi & \pm \sin \phi \\ \mp \sin \phi & \cos \phi \end{bmatrix} = \cos \phi \cdot \begin{bmatrix} 1 & \pm \tan \phi \\ \mp \tan \phi & 1 \end{bmatrix} = C \cdot T,$$

and similarly, if a plane (i.e., trigonometric) rotation is to be applied from the right,

$$T \cdot C = \begin{bmatrix} 1 & \pm \tan \phi \\ \mp \tan \phi & 1 \end{bmatrix} \cdot \cos \phi = \begin{bmatrix} \cos \phi & \pm \sin \phi \\ \mp \sin \phi & \cos \phi \end{bmatrix}.$$

A multiplication by  $T$  can be realized by a single `fma` per an element of the result, while the subsequent scaling by  $C$  can be converted to divisions by the corresponding secants. A similar factorization holds for a hyperbolic rotation, e.g., from the right,

$$T \cdot C = \begin{bmatrix} 1 & \tanh \phi \\ \tanh \phi & 1 \end{bmatrix} \cdot \cosh \phi = \begin{bmatrix} \cosh \phi & \sinh \phi \\ \sinh \phi & \cosh \phi \end{bmatrix}.$$

## 2.2 The HSVD of $\widehat{T}$

Given  $\widehat{T}$  from Algorithm 1, the HSVD of  $\widehat{T}$ ,  $\widetilde{U}^* \widehat{T} \widetilde{V} = \widetilde{\Sigma}$ , now remains to be found.

First,  $\widetilde{U}'^* \widehat{T} \widetilde{V}' = \widetilde{\Sigma}'$  is computed, where (despite the notation) all matrices are real,  $\widetilde{\Sigma}'$  is diagonal with non-negative elements,  $\widetilde{U}'^*$  is orthogonal, and is sought in the form of a plane rotation, while  $\widetilde{V}'$  is  $\widehat{J}$ -orthogonal, and is sought in the form of a plane rotation when  $\widehat{J} = I_2$  or  $\widehat{J} = -I_2$ , and in the form of a hyperbolic rotation otherwise.

If  $\widehat{J} \neq -I_2$ ,  $\widetilde{\Sigma}'$  is equal to the matrix  $\widetilde{\Sigma}$  of the hyperbolic singular values of  $\widehat{T}$ . Else,  $\widetilde{\Sigma}' = P_2 \widetilde{\Sigma} P_2^*$ , i.e., the diagonal elements of  $\widetilde{\Sigma}'$  are in the order opposite to the one prescribed by Definition 1.1, and thus have to be swapped to get  $\widetilde{\Sigma} = P_2^* \widetilde{\Sigma}' P_2$ . In the latter case, when  $\widehat{J} = -I_2$ , let  $\widetilde{U}^*$  and  $\widetilde{V}$  be  $P_2^* \widetilde{U}'^*$  and  $\widetilde{V}' P_2$ , respectively; else, let  $\widetilde{U}^* = \widetilde{U}'^*$  and  $\widetilde{V} = \widetilde{V}'$ . Note that  $\widetilde{\Sigma}$  holds the scaled hyperbolic singular values of  $\widehat{G}$ .

There are four non-disjoint cases for  $\widehat{T}$  and  $\widehat{J}$ , to be considered in the following order, stopping at, and proceeding as in, the first case for which its condition is satisfied:

1. if  $\widehat{T}$  is already diagonal then let  $\widetilde{U}'^* = \widetilde{V}' = I_2$  and  $\widetilde{\Sigma}' = \widehat{T}$ , else
2. if  $\widehat{J} = I_2$  or  $\widehat{J} = -I_2$ ,  $\widehat{T}$  is upper triangular and proceed as in Section 2.2.1, else
3. if  $\widehat{T}$  is upper triangular, then proceed as in Section 2.2.2, else
4.  $\widehat{T}$  is lower triangular, and proceed as in Section 2.2.3.

The first case above corresponds to  $D = \top$ , the second one to  $H = \perp$ , the third one to  $H \wedge \neg C = \top$ , and the last one to  $H \wedge C = \top$ , in terms of the states of Algorithm 1.

In fact, instead of  $\widetilde{\Sigma}'$ , the backscaled  $\widehat{\Sigma}' = 2^{-s} \widetilde{\Sigma}'$  is directly computed for stability, i.e., the hyperbolic singular values of  $\widehat{G}$ , instead of  $\widehat{T}$ , are finally obtained by the backscaling procedure described in a separate paragraph and Algorithm 2 below.

**Backscaling.** A backscaling routine from Algorithm 2 is used in Algorithms 3, 4, and 5. Algorithm 2 distributes the scale  $2^{-s}$ , with  $s$  from (2.1), among one of the values to which it should be applied and the given factor  $f$ ,  $1 \leq f \leq \Omega/4$ . The values  $d_1$  and  $d_2$  correspond to the diagonal elements of  $\widehat{T}$ , in some order. From (2.8), (2.14), and (2.20), it follows that the scaled hyperbolic singular values are of the form  $d_1 f$  and  $d_2/f$ , for a certain  $f$ . The idea behind Algorithm 2 is to upscale  $f$  (since  $s \geq -2$ , it cannot be by more than fourfold) or downscale it as much as possible, while keeping it above the underflow threshold  $\widehat{\omega}$ , to  $f'$ . Any remaining backscaling factor is applied to  $d_1$  before multiplying it by  $f'$  to get  $d'_1$ , while  $d_2$  is divided by  $f$  and the result is fully backscaled by  $2^{-s}$  to obtain  $d'_2$ . With such a distribution of  $2^{-s}$  among  $d_1$  and  $f$ , neither of them should lose accuracy in an intermediate computation leading to  $d'_1$  by being needlessly pushed down to the subnormal range, what could otherwise happen by multiplying any of them by the full backscaling factor.

**Algorithm 2** A safe backscaling of the hyperbolic singular values.

```

subroutine backscale (s, d1, d2, f, d'1, d'2);
Input: s from (2.1), d1 ≥ 0, d2 ≥ 0, 1 ≤ f ≤ Ω/4
Output: d'1 ≥ 0, d'2 ≥ 0
Data: μ = η(ω̂)
Δ = η(f) − μ; // exponents' distance between f and the
underflow threshold
if s > Δ then // partial backscaling of f
| f' = 2−Δ f, ξ = s − Δ;
else // full backscaling of f
| f' = 2−s f, ξ = 0;
end if
d'1 = (2−ξ d1) f', d'2 = 2−s (d2/f);
    
```

For example, let  $f = 1'$ , the immediate floating-point successor of unity, and  $2^{-s} = \widehat{\omega}/2$ . Then,  $2^{-s} f$  would cause the least significant bit of  $f$  to perish, with the result being subnormal. It might also happen that  $2^{-s} d_1$  becomes subnormal. Instead, with Algorithm 2,  $f'$  would be the immediate successor of  $\widehat{\omega}$ , so no accuracy would be lost and normality of the result would be preserved, and  $\xi = 1$ , i.e.,  $2^{-\xi} d_1 = d_1/2$ .

Let in the following  $\alpha$  be  $\text{SQRT}(\Omega) \approx 1.34078079299425956 \cdot 10^{154}$  (when computing in DOUBLE PRECISION), corresponding to  $\sqrt{\text{DBL\_MAX}}$  parameter from [22].

**2.2.1 Trigonometric case with  $\widehat{T}$  upper triangular**

This special case of the ordinary SVD and has been thoroughly covered in [22] (with the trigonometric angles of the opposite signs), but is summarized here for completeness. Many techniques to be introduced here are also used in the hyperbolic case.

The diagonalization requirement  $\widetilde{U}'^* \widehat{T} \widetilde{V}' = \widetilde{\Sigma}'$  can be expressed as

$$\begin{bmatrix} \cos \varphi & \sin \varphi \\ -\sin \varphi & \cos \varphi \end{bmatrix} \begin{bmatrix} \hat{t}_{11} & \hat{t}_{12} \\ 0 & \hat{t}_{22} \end{bmatrix} \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} = \begin{bmatrix} \tilde{\sigma}'_{11} & 0 \\ 0 & \tilde{\sigma}'_{22} \end{bmatrix},$$

where the order of the two matrix multiplications is arbitrary. Let such order be fixed to  $(\widetilde{U}'^* \widehat{T}) \widetilde{V}'$ . By observing that  $\hat{t}_{11} \geq \max\{\hat{t}_{12}, \hat{t}_{22}\} \geq 0$  and  $\hat{t}_{11} > 0$ , and by restricting the ranges for  $\varphi$  and  $\psi$  such that  $\cos \varphi > 0$  and  $\cos \psi > 0$ , it follows that both sides can be scaled by  $1/(\hat{t}_{11} \cos \varphi \cos \psi)$  to obtain a system of equations in  $\tan \varphi$  and  $\tan \psi$  as

$$\begin{bmatrix} 1 & \tan \varphi \\ -\tan \varphi & 1 \end{bmatrix} \begin{bmatrix} 1 & x \\ 0 & y \end{bmatrix} \begin{bmatrix} 1 & -\tan \psi \\ \tan \psi & 1 \end{bmatrix} = \begin{bmatrix} \tilde{\sigma}''_{11} & 0 \\ 0 & \tilde{\sigma}''_{22} \end{bmatrix}, \tag{2.2}$$

where  $0 \leq \min\{x, y\} \leq \max\{x, y\} \leq 1$ . This *crucial* conversion of one arbitrary value to a constant ( $\hat{t}_{11} \rightarrow 1$ ), enabled by the special form of  $\widehat{T}$ , is the key difference from the standard derivation of the Kogbetliantz formulas used in, e.g., xLASV2

LAPACK routines, that makes the ones to be derived here simpler but still highly accurate [22].

Equating the off-diagonal elements of the left- and the right-hand side of (2.2) gives the annihilation conditions

$$x + y \tan \psi - \tan \psi = 0 = \tan \psi (y - x \tan \varphi) - \tan \varphi, \quad (2.3)$$

from which  $\tan \psi$  can be expressed as

$$x + y \tan \varphi = \tan \psi = \frac{\tan \varphi}{y - x \tan \varphi}, \quad (2.4)$$

where the right-hand side is not defined if and only if  $y - x \tan \varphi = 0$ , what would imply, from (2.3), that  $\tan \varphi = 0$  and thus  $y = 0$ , so  $\tan \psi = x$  and the decomposition of  $\widehat{T}$  is complete. In all other cases, (2.4) leads to a quadratic equation in  $\tan \varphi$ ,

$$xy(1 - \tan^2 \varphi) = (x^2 - y^2 + 1) \tan \varphi,$$

or, with the terms rearranged,

$$0 \leq \frac{xy}{x^2 - y^2 + 1} = \frac{\tan \varphi}{1 - \tan^2 \varphi} = \frac{1}{2} \tan(2\varphi),$$

where  $\tan^2 \varphi = 1$  would imply  $x^2 - y^2 + 1 = 0$ , what cannot happen since  $0 \leq y \leq 1$  and  $x > 0$  (due to the assumption that  $\widehat{T}$  is not diagonal). Thus,  $0 \leq \tan(2\varphi) < \infty$ ,

$$\tan(2\varphi) = \frac{2xy}{(x - y)(x + y) + 1}, \quad (2.5)$$

and  $0 \leq \tan \varphi < 1$ . Note that  $y = 0$  implies  $\tan(2\varphi) = 0$ , so  $\tan \varphi = 0$  and no special handling, as above, of this case is actually required. Also, all squares of the elements of the scaled  $\widehat{T}$  have vanished from the denominator in (2.5). Then,

$$\tan \varphi = \frac{\tan(2\varphi)}{1 + \sqrt{1 + \tan^2(2\varphi)}}, \quad \sec \varphi = \sqrt{1 + \tan^2 \varphi}, \quad \cos \varphi = 1/\sec \varphi, \quad (2.6)$$

while, from (2.4),

$$\tan \psi = x + y \tan \varphi, \quad \sec \psi = \sqrt{1 + \tan^2 \psi}, \quad \cos \psi = 1/\sec \psi. \quad (2.7)$$

The following result has been proven as a part of [22, Theorem 1].

**Theorem 2.1** *It holds  $\sqrt{2} \geq \tan \psi \geq \tan \varphi \geq 0$  and  $\tilde{\sigma}'_{11} \geq \tilde{\sigma}'_{22} \geq 0$ , where*

$$\tilde{\sigma}'_{11} = \frac{\sec \psi}{\sec \varphi} \hat{t}_{11} = \frac{\cos \varphi}{\cos \psi} \hat{t}_{11}, \quad \tilde{\sigma}'_{22} = \frac{\sec \varphi}{\sec \psi} \hat{t}_{22} = \frac{\cos \psi}{\cos \varphi} \hat{t}_{22}. \quad (2.8)$$

From Theorem 2.1 it follows that the computed singular values of  $\widehat{T}$  are ordered descendingly. When  $\widehat{J} = -I_2$ , they have to be swapped, as described in Section 2.2.

Algorithm 3 stably computes the derived quantities in floating-point and works also for a diagonal  $\widehat{T}$ . Bounding  $\tan(2\varphi)$  from above by  $\alpha$  is necessary to avoid a possible overflow of the argument of the square root in computing  $\tan \varphi$ , without resorting to HYPOT or losing accuracy of the result (see [22, subsection 2.3]). The expression  $2xy$  is computed as  $(2a)b$ , where  $a = \min\{x, y\}$  and  $b = \max\{x, y\}$ , to prevent the avoidable underflows of  $xy$  (e.g., when  $x$  is just above the underflow threshold and  $y \approx 1/2$ ). No division  $a/b$  should be replaced by  $a(1/b)$ , nor square root inaccurately approximated. Algorithm 3 is reproducible if a method of computing  $1/\sqrt{x}$  is fixed.

---

**Algorithm 3** The SVD of an upper triangular  $\widehat{T}$  in DOUBLE PRECISION.

---

$x = \text{MAX}(\widehat{t}_{12}/\widehat{t}_{11}, 0), \quad y = \text{MAX}(\widehat{t}_{22}/\widehat{t}_{11}, 0); \quad // \quad (2.2) \quad (\text{MAX handles } \widehat{T} = \mathbf{0} \text{ case})$

$\tan(2\varphi) = \min \left\{ \text{MAX} \left( \frac{(2 \min\{x, y\}) \max\{x, y\}}{\text{fma}(x - y, x + y, 1)}, 0 \right), \alpha \right\}; \quad // \quad (2.5)$

$\tan \varphi = \frac{\tan(2\varphi)}{1 + \sqrt{\text{fma}(\tan(2\varphi), \tan(2\varphi), 1)}}, \quad \tan \psi = \text{fma}(y, \tan \varphi, x);$

// (2.6) -- (2.7)

**if** an accurate enough (e.g., correctly rounded)  $\text{rsqrt}(x) = 1/\sqrt{x}$  function is available

**then**

$\cos \varphi = \text{rsqrt}(\text{fma}(\tan \varphi, \tan \varphi, 1)); \quad // \quad (2.6)$

$\cos \psi = \text{rsqrt}(\text{fma}(\tan \psi, \tan \psi, 1)); \quad // \quad (2.7)$

$\text{backscale}(s, \widehat{t}_{11}, \widehat{t}_{22}, \cos \varphi / \cos \psi, \widehat{\sigma}'_{11}, \widehat{\sigma}'_{22}); \quad // \quad (2.8) \text{ \& Algorithm 2}$

**else** // usually,  $\text{rsqrt}(x)$  is not available

$\sec \varphi = \sqrt{\text{fma}(\tan \varphi, \tan \varphi, 1)}, \quad \cos \varphi = 1 / \sec \varphi; \quad // \quad (2.6)$

$\sec \psi = \sqrt{\text{fma}(\tan \psi, \tan \psi, 1)}, \quad \cos \psi = 1 / \sec \psi; \quad // \quad (2.7)$

$\text{backscale}(s, \widehat{t}_{11}, \widehat{t}_{22}, \sec \psi / \sec \varphi, \widehat{\sigma}'_{11}, \widehat{\sigma}'_{22}); \quad // \quad (2.8) \text{ \& Algorithm 2}$

**end if**

---

### 2.2.2 Hyperbolic case with $\widehat{T}$ upper triangular

The diagonalization requirement  $\widetilde{U}'^* \widehat{T} \widetilde{V}' = \widetilde{\Sigma}'$ , scaled by  $1/(\widehat{t}_{11} \cos \varphi \cosh \psi) > 0$  in the hyperbolic case for  $\widehat{T}$  upper triangular, is

$$\begin{bmatrix} 1 & \tan \varphi \\ -\tan \varphi & 1 \end{bmatrix} \begin{bmatrix} 1 & x \\ 0 & y \end{bmatrix} \begin{bmatrix} 1 & \tanh \psi \\ \tanh \psi & 1 \end{bmatrix} = \begin{bmatrix} \widetilde{\sigma}''_{11} & 0 \\ 0 & \widetilde{\sigma}''_{22} \end{bmatrix}, \quad (2.9)$$

from which, similarly to the trigonometric case, the annihilation conditions follow as

$$x + y \tan \varphi + \tanh \psi = 0 = \tanh \psi (y - x \tan \varphi) - \tan \varphi. \quad (2.10)$$

If  $y = 0$  then  $\tan \varphi = 0$  and  $\tanh \psi = -x$  satisfy (2.10), so the decomposition of  $\widehat{T}$  is complete, unless  $x = 1$  and therefore  $\tanh \psi = -1$ , what is impossible. Thus, the HSVD is *not defined* when the columns of  $\widehat{T}$  (equivalently, of  $\widehat{G}$ ) are identical.

In the remaining cases, with  $y > 0$ ,  $\tanh \psi$  can be expressed as

$$x + y \tan \varphi = -\tanh \psi = \frac{\tan \varphi}{x \tan \varphi - y}, \quad (2.11)$$

what gives, in a manner and with caveats similar to the trigonometric case,

$$-\infty < \tan(2\varphi) = \frac{-2xy}{(y-x)(y+x)+1} \leq 0, \quad (2.12)$$

and the remaining functions of  $\varphi$  are computed as in (2.6). From (2.11) it follows

$$\tanh \psi = -(x + y \tan \varphi), \quad \operatorname{sech} \psi = \sqrt{1 - \tanh^2 \psi}, \quad \cosh \psi = 1/\operatorname{sech} \psi. \quad (2.13)$$

Theorem 2.2 is an analogon of Theorem 2.1 in this hyperbolic case.

**Theorem 2.2** *If, in (2.9),  $x \neq 1$ , then  $|\tanh \psi| < 1$  and*

$$\tilde{\sigma}'_{11} = \frac{\operatorname{sech} \psi}{\sec \varphi} \hat{t}_{11} = \frac{\cos \varphi}{\cosh \psi} \hat{t}_{11}, \quad \tilde{\sigma}'_{22} = \frac{\sec \varphi}{\operatorname{sech} \psi} \hat{t}_{22} = \frac{\cosh \psi}{\cos \varphi} \hat{t}_{22}. \quad (2.14)$$

*Proof* From (2.12),  $-1 < \tan \varphi \leq 0$ , and from (2.13),  $|\tanh \psi| = |x + y \tan \varphi|$ . With  $x$  fixed,  $|\tanh \psi|$  attains the maximum value for  $y \tan \varphi$  being either the smallest, i.e., when approaching  $-1$  from the right, or the largest possible, i.e., zero. In the first case,  $|\tanh \psi| < |x - 1| \leq 1$ . In the second,  $|\tanh \psi| = |x| = x$ , so if  $x \neq 1$  then  $|\tanh \psi| < 1$ .

Now (2.14) is proven. From (2.9) and (2.13) it follows

$$\begin{aligned} \tilde{\sigma}''_{11} &= 1 + (x + y \tan \varphi) \tanh \psi = 1 - \tanh^2 \psi = \operatorname{sech}^2 \psi, \\ \tilde{\sigma}''_{22} &= y - x \tan \varphi - \tan \varphi \tanh \psi = y - \tan \varphi (x + \tanh \psi) \\ &= y + y \tan^2 \varphi = y \sec^2 \varphi, \end{aligned}$$

what, after multiplying both equations by  $\hat{t}_{11} \cos \varphi \cosh \psi$ , gives (2.14).  $\square$

Algorithm 4, similarly to Algorithm 3, computes reproducibly and as stably as practicable the HSVD of an upper triangular (including diagonal)  $\widehat{T}$  when  $\widehat{J} \neq \pm I_2$ .



**Algorithm 4** The HSVD of an upper triangular  $\widehat{T}$  in DOUBLE PRECISION.

```

x = MAX( $\hat{t}_{12}/\hat{t}_{11}$ , 0), y = MAX( $\hat{t}_{22}/\hat{t}_{11}$ , 0); // (2.9) (MAX handles
 $\widehat{T} = \mathbf{0}$  case)
tan(2 $\varphi$ ) = -min { MAX (  $\frac{(2 \min\{x, y\}) \max\{x, y\}}{\text{fma}(y - x, y + x, 1)}$ , 0 ),  $\alpha$  }; // (2.12)
tan  $\varphi$  =  $\frac{\tan(2\varphi)}{1 + \sqrt{\text{fma}(\tan(2\varphi), \tan(2\varphi), 1)}}$ , tanh  $\psi$  = -fma(y, tan  $\varphi$ , x);
// (2.6) & (2.13)
if |tanh  $\psi$ |  $\geq \nu$  then return  $\perp$ ; // |tanh  $\psi$ | is unsafe if it is too
close to unity
if an accurate enough (e.g., correctly rounded)  $\text{rsqrt}(x) = 1/\sqrt{x}$  function is
available
then
| cos  $\varphi$  =  $\text{rsqrt}(\text{fma}(\tan \varphi, \tan \varphi, 1))$ ; // (2.6)
| cosh  $\psi$  =  $\text{rsqrt}(\text{fma}(-\tanh \psi, \tanh \psi, 1))$ ; // (2.13)
| backscale (s,  $\hat{t}_{22}$ ,  $\hat{t}_{11}$ , cosh  $\psi$  / cos  $\varphi$ ,  $\hat{\sigma}'_{22}$ ,  $\hat{\sigma}'_{11}$ ); // (2.14) &
| Algorithm 2
else // usually,  $\text{rsqrt}(x)$  is not available
| sec  $\varphi$  =  $\sqrt{\text{fma}(\tan \varphi, \tan \varphi, 1)}$ , cos  $\varphi$  = 1 / sec  $\varphi$ ; // (2.6)
| sech  $\psi$  =  $\sqrt{\text{fma}(-\tanh \psi, \tanh \psi, 1)}$ , cosh  $\psi$  = 1 / sech  $\psi$ ; // (2.13)
| backscale (s,  $\hat{t}_{22}$ ,  $\hat{t}_{11}$ , sec  $\varphi$  / sech  $\psi$ ,  $\hat{\sigma}'_{22}$ ,  $\hat{\sigma}'_{11}$ ); // (2.14) &
| Algorithm 2
end if
return  $\top$ ; // the hyperbolic transformation is defined
and safe

```

Here and in Algorithm 5 from Section 2.2.3 a safety parameter  $\nu \lesssim 1$  is introduced that will be fully explained in Section 2.3. For now, assume  $\nu = 1$ .

### 2.2.3 Hyperbolic case with $\widehat{T}$ lower triangular

The diagonalization requirement  $\widetilde{U}'^* \widehat{T} \widetilde{V}' = \widetilde{\Sigma}'$ , scaled by  $1/(\hat{t}_{22} \cos \varphi \cosh \psi) > 0$  in the hyperbolic case for  $\widehat{T}$  lower triangular, is

$$\begin{bmatrix} 1 & \tan \varphi \\ -\tan \varphi & 1 \end{bmatrix} \begin{bmatrix} y & 0 \\ x & 1 \end{bmatrix} \begin{bmatrix} 1 & \tanh \psi \\ \tanh \psi & 1 \end{bmatrix} = \begin{bmatrix} \tilde{\sigma}''_{11} & 0 \\ 0 & \tilde{\sigma}''_{22} \end{bmatrix}, \tag{2.15}$$

from which the annihilation conditions follow as

$$\tanh \psi (y + x \tan \varphi) + \tan \varphi = 0 = x - y \tan \varphi + \tanh \psi. \tag{2.16}$$

If  $y = 0$  then  $\tan \varphi = 0$  and  $\tanh \psi = -x$  satisfy (2.16), so the decomposition of  $\widehat{T}$  is complete, unless  $x = 1$  and thus  $\tanh \psi = -1$ , what is impossible. The HSVD of  $\widehat{T}$

is not defined, same as in Section 2.2.2, when the columns of  $\widehat{T}$  (i.e.,  $\widehat{G}$ ) are identical.

In the remaining cases, with  $y > 0$ ,  $\tanh \psi$  can be expressed as

$$x - y \tan \varphi = -\tanh \psi = \frac{\tan \varphi}{y + x \tan \varphi}, \tag{2.17}$$

what gives, similarly to the other hyperbolic case,

$$0 \leq \tan(2\varphi) = \frac{2xy}{(y-x)(y+x)+1} < \infty, \tag{2.18}$$

and the remaining functions of  $\varphi$  are computed as in (2.6). From (2.17) it follows

$$\tanh \psi = y \tan \varphi - x, \quad \operatorname{sech} \psi = \sqrt{1 - \tanh^2 \psi}, \quad \cosh \psi = 1/\operatorname{sech} \psi. \tag{2.19}$$

Theorem 2.3 is an analogon of Theorem 2.2 in this hyperbolic case.

**Theorem 2.3** *If, in (2.15),  $x \neq 1$ , then  $|\tanh \psi| < 1$  and*

$$\tilde{\sigma}'_{11} = \frac{\sec \varphi}{\operatorname{sech} \psi} \hat{t}_{11} = \frac{\cosh \psi}{\cos \varphi} \hat{t}_{11}, \quad \tilde{\sigma}'_{22} = \frac{\operatorname{sech} \psi}{\sec \varphi} \hat{t}_{22} = \frac{\cos \varphi}{\cosh \psi} \hat{t}_{22}. \tag{2.20}$$

*Proof* From (2.18)–(2.19),  $0 \leq \tan \varphi < 1$  and  $|\tanh \psi| = |y \tan \varphi - x|$ . With  $x$  fixed,  $|\tanh \psi|$  attains the maximum value for  $y \tan \varphi$  being either the largest, i.e., when approaching unity from the left, or the smallest possible, i.e., zero. In the first case,  $|\tanh \psi| < |1 - x| \leq 1$ . In the second,  $|\tanh \psi| = |-x| = x$ , so  $x \neq 1 \implies |\tanh \psi| < 1$ .

Now (2.20) is proven. From (2.15) and (2.19) it follows

$$\begin{aligned} \tilde{\sigma}''_{11} &= y + x \tan \varphi + \tan \varphi \tanh \psi = y + \tan \varphi(x + \tanh \psi) \\ &= y + y \tan^2 \varphi = y \sec^2 \varphi, \\ \tilde{\sigma}''_{22} &= 1 + (x - y \tan \varphi) \tanh \psi = 1 - \tanh^2 \psi = \operatorname{sech}^2 \psi, \end{aligned}$$

what, after multiplying both equations by  $\hat{t}_{22} \cos \varphi \cosh \psi$ , gives (2.20). □

Algorithm 5, similarly to Algorithm 4, computes reproducibly and as stably as practicable the HSVD of a lower triangular (including diagonal)  $\widehat{T}$  when  $\widehat{J} \neq \pm I_2$ .

Let  $\varepsilon$ , unless noted otherwise, be half the unit in the last place (ulp) of 1, or equivalently,  $\varepsilon = 1 - '1$ , where  $'1$  is the immediate floating-point predecessor of 1 in the chosen datatype; e.g.,  $\varepsilon = 2^{-24}$  and  $\varepsilon = 2^{-53}$  in single and double precision, respectively. In Fortran, the latter is equal to the result of `EPSILON(OD0)/2`.

From (2.13) and (2.19) it follows  $\sqrt{\varepsilon} \leq \operatorname{sech} \psi \leq 1$ , since the argument of the secant's defining square root cannot be less than  $\varepsilon$  whenever  $|\tanh \psi| < 1$  (what is implied by  $\tanh^2 \psi \leq |\tanh \psi| \leq '1$ ). Also, from (2.5), (2.12), and (2.18) it can be concluded that  $|\tan \varphi| < 1$ , so  $1 \leq \sec \varphi \leq \sqrt{2}$  and  $1 \leq \sec \varphi / \operatorname{sech} \psi \leq \sqrt{2/\varepsilon}$ . In the trigonometric case,  $1 \leq \sec \psi \leq \sqrt{3}$  and  $\sec \varphi \leq \sec \psi$ , due to (2.7) and Theorem 2.1, so  $1 \leq \sec \psi / \sec \varphi \leq \sqrt{3} < \sqrt{2/\varepsilon}$ . Hence,  $f \leq \sqrt{2/\varepsilon}$  in the uses of Algorithm 2.

**Algorithm 5** The HSVD of a lower triangular  $\widehat{T}$  in DOUBLE PRECISION.

```

x = MAX( $\hat{t}_{21}/\hat{t}_{22}$ , 0), y = MAX( $\hat{t}_{11}/\hat{t}_{22}$ , 0); // (2.15) (MAX handles
 $\widehat{T} = \mathbf{0}$  case)
tan(2 $\varphi$ ) = min { MAX (  $\frac{(2 \min\{x, y\} \max\{x, y\})}{\text{fma}(y - x, y + x, 1)}$ , 0 ),  $\alpha$  }; // (2.18)
tan  $\varphi$  =  $\frac{\tan(2\varphi)}{1 + \sqrt{\text{fma}(\tan(2\varphi), \tan(2\varphi), 1)}}$ , tanh  $\psi$  = fma(y, tan  $\varphi$ , -x);
// (2.6) & (2.19)
if |tanh  $\psi$ |  $\geq \nu$  then return  $\perp$ ; // |tanh  $\psi$ | is unsafe if it is too
close to unity
if an accurate enough (e.g., correctly rounded) rsqrt(x) = 1/ $\sqrt{x}$  function is
available
then
| cos  $\varphi$  = rsqrt(fma(tan  $\varphi$ , tan  $\varphi$ , 1)); // (2.6)
| cosh  $\psi$  = rsqrt(fma(-tanh  $\psi$ , tanh  $\psi$ , 1)); // (2.19)
| backscale (s,  $\hat{t}_{11}$ ,  $\hat{t}_{22}$ , cosh  $\psi$  / cos  $\varphi$ ,  $\hat{\sigma}'_{11}$ ,  $\hat{\sigma}'_{22}$ ); // (2.20) &
Algorithm 2
else // usually, rsqrt(x) is not available
| sec  $\varphi$  =  $\sqrt{\text{fma}(\tan \varphi, \tan \varphi, 1)}$ , cos  $\varphi$  = 1/sec  $\varphi$ ; // (2.6)
| sech  $\psi$  =  $\sqrt{\text{fma}(-\tanh \psi, \tanh \psi, 1)}$ , cosh  $\psi$  = 1/sech  $\psi$ ; // (2.19)
| backscale (s,  $\hat{t}_{11}$ ,  $\hat{t}_{22}$ , sec  $\varphi$ /sech  $\psi$ ,  $\hat{\sigma}'_{11}$ ,  $\hat{\sigma}'_{22}$ ); // (2.20) &
Algorithm 2
end if
return  $\top$ ; // the hyperbolic transformation is defined
and safe

```

**2.3 The HSVD of  $\widehat{G}$**

The HSVD of  $\widehat{G}$  as  $\widehat{U}^* \widehat{G} \widehat{V} = \widehat{\Sigma}$  is obtained in four phases:

1. computing the scaling parameter  $s$  and  $\widehat{G}_0 = 2^s \widehat{G}$  as in Section 2.1.1,
2. factoring  $\widehat{G}_0$  as  $\widehat{U}^* \widehat{G}_0 \widehat{V} = \widehat{T}$  by Algorithm 1,
3. computing the HSVD of  $\widehat{T}$ ,  $\widetilde{U}^* \widehat{T} \widetilde{V} = \widetilde{\Sigma}$  and  $\widehat{\Sigma} = 2^{-s} \widetilde{\Sigma}$ , as in Section 2.2, and
4. assembling  $\widehat{U}^* = \widetilde{U}^* \widehat{U}^*$  and  $\widehat{V} = \widetilde{V} \widehat{V}$  (and, if required,  $\widehat{V}^{-1} = \widetilde{V} \widehat{V}^{-1}$ ).

The third phase can fail if the HSVD is not defined, or is unsafe, as indicated by the return values of Algorithm 4 and Algorithm 5. Safety (or a lack thereof), parametrized by  $\nu$  in those algorithms, is a user’s notion of how close  $|\tanh \psi|$  can get to unity from below to still define a hyperbolic rotation that is well-enough conditioned to be applied to the pivot columns of the iteration matrix. For the  $2 \times 2$  HSVD in isolation,  $\nu = 1$ , but for the  $n_0 \times n_0$  HSVD it might sometimes be necessary to set it a bit lower if the HSVD process otherwise fails to converge, e.g., to  $\nu = 0.8$ , as in [35], or even lower. There is no prescription for choosing an adequate  $\nu$  in advance; if an execution

takes far more (multi-)steps than expected (see Section 6 for some estimates in terms of cycles), it should be aborted and restarted with a lower  $\nu$ .

If  $\nu < 1$ ,  $|\tanh \psi| < 1$ , and the return value from either algorithm is  $\perp$ , there are two possibilities. First,  $\tanh \psi$  can be set to  $\text{sign}(\tanh \psi)\nu$ , and the hyperbolic transformation can be computed accordingly [32]. However, the off-diagonal elements of the transformed pivot matrix cannot be considered zeros anymore, and they have to be formed in the iteration matrix and included in the weight computations (see Section 3.2). An inferior but simpler solution declares that the  $2 \times 2$  HSVD is not defined, as if  $|\tanh \psi| \geq 1$ , and the pivot pair in question will no longer be a transformation candidate in the current multi-step, in the hope that it will again become one, with a better conditioned hyperbolic rotation, after its pivot row(s) and column(s) have been sufficiently transformed. In the implementation the latter option has been chosen, but the former might lead (not tested) to a faster convergence when it has to be employed.

### 3 Row and column transformations

If the HSVD of  $\widehat{G}_{k-1}$  is not defined, the algorithm stops. Else, having computed  $\widehat{U}_k^*$ ,  $\widehat{\Sigma}_k$ , and  $\widehat{V}_k$  for a transformation candidate with the pivot indices  $p_k$  and  $q_k$ , the  $p_k$ th and the  $q_k$ th row of  $G_{k-1}$  are transformed by multiplying them from the left by  $\widehat{U}_k^*$ ,

$$\begin{bmatrix} G'_{k-1}(p_k, :) \\ G'_{k-1}(q_k, :) \end{bmatrix} = \widehat{U}_k^* \begin{bmatrix} G_{k-1}(p_k, :) \\ G_{k-1}(q_k, :) \end{bmatrix}.$$

Then,  $G_k$  is obtained from  $G'_{k-1}$  after transforming the  $p_k$ th and the  $q_k$ th columns of  $G'_{k-1}$  by multiplying them from the right by  $V_k$ ,

$$\begin{bmatrix} G_k(:, p_k) & G_k(:, q_k) \end{bmatrix} = \begin{bmatrix} G'_{k-1}(:, p_k) & G'_{k-1}(:, q_k) \end{bmatrix} V_k,$$

and setting  $G_k(p_k, p_k)$  to the first diagonal element of  $\widehat{\Sigma}_k$ ,  $G_k(q_k, q_k)$  to the second one (in both cases reusing the possibly more accurate hyperbolic singular values from the HSVD of  $\widehat{G}_{k-1}$  then those computed by the row and the column transformations of  $G_k$ ), while explicitly zeroing out  $G_k(p_k, q_k)$  and  $G_k(q_k, p_k)$ .

If the left and the right (hyperbolic) singular vectors are desired, in a similar way as above the current approximations of  $U^*$  and  $V$  are updated by  $\widehat{U}_k^*$  and  $\widehat{V}_k$ , respectively.

#### 3.1 Effects of a hyperbolic transformation

If  $\widehat{V}_k$  is unitary, the square of the off-diagonal Frobenius norm of the transformed  $G_k$  is reduced by  $|G_{k-1}(q_k, p_k)|^2 + |G_{k-1}(p_k, q_k)|^2 \geq 0$ . Else, if  $\widehat{V}_k$  is  $\widehat{J}_k$ -unitary, the following Lemma 3.1 sets the bounds to the relative change of the square of the Frobenius norm of the  $p_k$ th and the  $q_k$ th columns multiplied from the right by a hyperbolic rotation. Note that a complex  $\widehat{V}_k$  is of the form

$$\widehat{V}_k = \begin{bmatrix} 1 & 0 \\ 0 & e^{-i\beta} \end{bmatrix} \begin{bmatrix} \cosh \psi & \sinh \psi \\ \sinh \psi & \cosh \psi \end{bmatrix} = \begin{bmatrix} \cosh \psi & \sinh \psi \\ e^{-i\beta} \sinh \psi & e^{-i\beta} \cosh \psi \end{bmatrix},$$

which can alternatively be expressed as

$$\begin{bmatrix} \cosh \psi & \sinh \psi \\ e^{-i\beta} \sinh \psi & e^{-i\beta} \cosh \psi \end{bmatrix} = \begin{bmatrix} \cosh \psi & e^{i\beta} \sinh \psi \\ e^{-i\beta} \sinh \psi & \cosh \psi \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & e^{-i\beta} \end{bmatrix},$$

where the rightmost matrix (call it  $B$ ) is unitary and does not change the Frobenius norm of any matrix that multiplies it from the left. Lemma 3.1 is thus stated in full generality, for  $\widehat{V}_k B^*$ , since  $\|X \widehat{V}_k B^*\|_F = \|X \widehat{V}_k B^* B\|_F = \|X \widehat{V}_k\|_F$  for all conformant  $X$ .

**Lemma 3.1** *If  $\mathbf{x}$  and  $\mathbf{y}$  are complex vectors of length  $n$ , such that  $\|[\mathbf{x} \ \mathbf{y}]\|_F > 0$ , and*

$$[\mathbf{x}' \ \mathbf{y}'] = [\mathbf{x} \ \mathbf{y}] \begin{bmatrix} \cosh \psi & e^{i\beta} \sinh \psi \\ e^{-i\beta} \sinh \psi & \cosh \psi \end{bmatrix},$$

then

$$\cosh(2\psi) - |\sinh(2\psi)| \leq \frac{\|[\mathbf{x}' \ \mathbf{y}']\|_F^2}{\|[\mathbf{x} \ \mathbf{y}]\|_F^2} \leq \cosh(2\psi) + |\sinh(2\psi)|.$$

The following Lemma 3.2 refines the bounds stated in Lemma 3.1. Both Lemmas, proved in the [Appendix](#), are used in the proof of Theorem 3.1 in Section 3.2.

**Lemma 3.2** *In the lower or the upper bound established in Lemma 3.1 the equality is attainable if and only if  $\mathbf{y} = \pm e^{i\beta} \mathbf{x}$  or  $\psi = 0$ . The lower bound is always positive but at most unity, and the upper bound is at least unity.*

Another observation is that the norm of the transformed columns depends both on the norm of the original columns, as well as on the hyperbolic transformation applied. Therefore,  $\psi$  of a relatively large magnitude does not by itself pose a problem if the original columns have a modest norm. And contrary, even  $\psi$  of a relatively small magnitude can—and in practice, will—cause the columns’ elements of a huge magnitude, should such exist, to overflow in the finite machine arithmetic.

### 3.2 Weight of a transformation candidate

Let  $\text{off}_F^2(A)$  be the square of the off-diagonal Frobenius norm of  $A \in \mathbb{F}^{n \times n}$ , i.e.,

$$\text{off}_F^2(A) = \sum_{j=1}^n \sum_{\substack{i=1 \\ i \neq j}}^n |a_{ij}|^2 = \|A - \text{diag}(a_{11}, \dots, a_{nn})\|_F^2.$$

The following Theorem 3.1 deals with the amount of change  $\text{off}_F^2(G_{k-1}) - \text{off}_F^2(G_k)$ .

**Theorem 3.1** *For  $k$  such that  $1 \leq k \leq N$  and  $w_k = \text{off}_F^2(G_{k-1}) - \text{off}_F^2(G_k)$  it holds*

$$w_k = |G_{k-1}(q_k, p_k)|^2 + |G_{k-1}(p_k, q_k)|^2 + h_k, \tag{3.1}$$

where  $h_k = 0$  and  $w_k$  is non-negative if  $V_k$  is unitary. Otherwise, for  $h_k$  holds

$$h_k = \sum_{\substack{i=1 \\ i \notin \{p_k, q_k\}}}^n \left( |G_{k-1}(i, p_k)|^2 - |G_k(i, p_k)|^2 + |G_{k-1}(i, q_k)|^2 - |G_k(i, q_k)|^2 \right), \tag{3.2}$$

and  $w_k$  can be negative, positive, or zero.

*Proof* When  $V_k$  is unitary, the statement of the Theorem 3.1 is a well-known property of the Kogbetliantz algorithm, and a consequence of  $U_k^*$  also being unitary, as well as the Frobenius norm being unitary invariant.

Else, if  $V_k$  is not unitary, then observe that the elements of  $G_k$  at the pivot positions (i.e., having their indices taken from the set  $\{p_k, q_k\}$ ) do not contribute to  $\text{off}_F^2(G_k)$ , since the off-diagonal elements at those positions are zero. The change from  $\text{off}_F^2(G_{k-1})$  to  $\text{off}_F^2(G_k)$  is therefore the sum of squares of the magnitudes of those elements, plus any change ( $h_k$ ) happening outside the pivot positions, as in (3.1).

The left transformation is unitary, and therefore the only two rows affected,  $p_k$ th and  $q_k$ th, shortened to have the pivot elements removed, keep their joint Frobenius norm unchanged from  $G_{k-1}$  to  $G'_{k-1}$ . Since the right transformation affects only the  $p_k$ th and the  $q_k$ th columns, either of which intersect the  $p_k$ th and the  $q_k$ th row in the pivot positions only, there is no further change from  $\text{off}_F^2(G'_{k-1})$  to  $\text{off}_F^2(G_k)$  when the off-diagonal norm is restricted to the shortened and transformed rows, so a contribution to  $h_k$  from the left transformation is zero.

Therefore, only the right transformation is responsible for the value of  $h_k$ , which can be bounded by Lemma 3.1, applied to the computed hyperbolic transformation  $\widehat{V}_k$  and the  $p_k$ th and the  $q_k$ th columns with the pivot elements removed from them. The square of the joint Frobenius norm of the shortened columns might either fall or rise after the transformation, due to Lemma 3.2. If it falls,  $h_k$  and thus  $w_k$  is positive. If it rises, depending on the hyperbolic angle  $\psi$  and on the off-diagonal elements,  $h_k$  can become negative and so large in magnitude to push  $w_k$  down to zero or below.

For example, let  $0 < \varepsilon \ll 1$  and observe that in the real case

$$\widehat{G}_{k-1} = \begin{bmatrix} 1 & \varepsilon - 1 \\ 0 & 0 \end{bmatrix} \implies \widehat{V}_k = \cosh \psi \begin{bmatrix} 1 & \varepsilon - 1 \\ 1 - \varepsilon & -1 \end{bmatrix}, \quad \cosh \psi = \frac{1}{\sqrt{\varepsilon(2 - \varepsilon)}},$$

(see Section 2.2.2). Let  $a$  and  $b$  lie in  $p_k$ th and the  $q_k$ th columns, respectively, in a row  $\ell \notin \{p_k, q_k\}$ , and let the other non-pivot elements be zero. Then  $a$  and  $b$  are transformed to  $\tilde{a}$  and  $\tilde{b}$ , respectively, where

$$\begin{bmatrix} \tilde{a} & \tilde{b} \end{bmatrix} = \begin{bmatrix} a & b \end{bmatrix} \widehat{V}_k = \begin{bmatrix} a + (1 - \varepsilon)b & a(\varepsilon - 1) - b \end{bmatrix} / \sqrt{\varepsilon(2 - \varepsilon)}.$$

If  $a = b = 1$ , then  $\tilde{a} = -\tilde{b} = \sqrt{2 - \varepsilon} / \sqrt{\varepsilon}$ , so

$$a^2 - \tilde{a}^2 + b^2 - \tilde{b}^2 = 4(\varepsilon - 1) / \varepsilon \approx -4 / \varepsilon,$$

what is by magnitude far greater than  $0^2 + |\varepsilon - 1|^2 \approx 1$ , thus  $w_k < 0$ . Oppositely,

$$\widehat{G}_{k-1} = \begin{bmatrix} 1 & 1 - \varepsilon \\ 0 & 0 \end{bmatrix} \implies \widehat{V}_k = \cosh \psi \begin{bmatrix} 1 & \varepsilon - 1 \\ \varepsilon - 1 & 1 \end{bmatrix}, \quad \cosh \psi = \frac{1}{\sqrt{\varepsilon(2 - \varepsilon)}},$$

while  $\tilde{a}$  and  $\tilde{b}$  are

$$\tilde{a} = (a + (\varepsilon - 1)b)/\sqrt{\varepsilon(2 - \varepsilon)}, \quad \tilde{b} = (a(\varepsilon - 1) + b)/\sqrt{\varepsilon(2 - \varepsilon)}.$$

If  $a = b = 1$ , then  $\tilde{a} = \tilde{b} = \sqrt{\varepsilon}/\sqrt{2 - \varepsilon}$ , so

$$a^2 - \tilde{a}^2 + b^2 - \tilde{b}^2 = 4(1 - \varepsilon)/(2 - \varepsilon) \approx 2,$$

thus  $w_k > 0$ . Finally, let  $a > 0, b = 0$ , and assume  $\check{V}_k = I_2$  for simplicity. Then

$$\begin{bmatrix} \tilde{a} & \tilde{b} \end{bmatrix} = \begin{bmatrix} a & 0 \end{bmatrix} \begin{bmatrix} \cosh \psi & \sinh \psi \\ \sinh \psi & \cosh \psi \end{bmatrix} = \begin{bmatrix} a \cosh \psi & a \sinh \psi \end{bmatrix},$$

what, together with  $1 = \cosh^2 \psi - \sinh^2 \psi$ , gives

$$a^2 - \tilde{a}^2 + b^2 - \tilde{b}^2 = a^2(1 - \cosh^2 \psi - \sinh^2 \psi) = -2a^2 \sinh^2 \psi,$$

what is equal to  $-2a^2$  for  $|\sinh \psi| = 1$ . Setting

$$\widehat{G}_{k-1} = \begin{bmatrix} 2a & a\sqrt{2} \\ 0 & 0 \end{bmatrix} \implies \widehat{V}_k = \begin{bmatrix} \sqrt{2} & -1 \\ -1 & \sqrt{2} \end{bmatrix}, \quad \sinh \psi = -1.$$

The sum of squares of the off-diagonal elements of  $\widehat{G}_{k-1}$  is  $2a^2$  and thus  $w_k = 0$ .  $\square$

A sequence of matrices  $(G_k)_{k \geq 0}$  converges to a diagonal form if and only if  $(\text{off}_F^2(G_k))_{k \geq 0}$  tends to zero. The sequence  $(\text{off}_F^2(G_k))_{k \geq 0}$  does not have to be monotonically decreasing, i.e.,  $w_k$  can be negative for some  $k$  in the HSVD case, unlike in the ordinary Kogbetliantz algorithm. That significantly complicates any reasoning about convergence in theory, and attainment of (a satisfactory rate of) convergence in practice. To aid the latter, the pivot weights  $w_k$  from (3.1) should be kept as high as possible, by a careful choice of the pivot submatrix among all admissible transformation candidates in each step. In the SVD computation, choosing a candidate with the maximal weight guarantees convergence [25], and is trivially accomplished since the weights are directly computable, due to  $h_k = 0$ . In the hyperbolic case,  $h_k$  from (3.2) cannot be known in advance, without performing the right (column) transformation, and it cannot be estimated (roughly, due to Lemma 3.2) by means of Lemma 3.1 without computing the associated  $\widehat{V}_k$  and occasionally recomputing the column norms, what is of the same linear complexity as transforming those columns.

The computation of  $h_k$  is therefore preferable to an estimation. In each step, it has to be performed for *all* admissible transformation candidates, and non-trivially for all indices  $p$  and  $q$  such that  $p < q, j_{pp} \neq j_{qq}$ , and the associated  $\widehat{V}_k \neq I_2$  (if  $\widehat{V}_k$  is not defined, let  $h_k = -\infty$  instead of halting), i.e., at most  $n_0^2/4$  times, if  $J_0$  has the same number of positive and negative signs. The left transformation by  $\widehat{U}_k^*$  is not needed here, so only the right one by  $\widehat{V}_k$  has to have a virtual variant, that does not change the elements of  $G_{k-1}$ , but for each row  $i \notin \{p, q\}$  computes what would  $G_k(i, p)$  and

$G_k(i, q)$  be from  $G_{k-1}(i, p)$ ,  $G_{k-1}(i, q)$ , and  $\widehat{V}_k$ , and updates  $h_k$  using (3.2). That can be done accurately by applying twice for each  $i \notin \{p, q\}$  the accumulation rule

$$\rho := \rho + (a^2 - \tilde{a}^2) = \text{fma}(a - \tilde{a}, a + \tilde{a}, \rho),$$

once for  $a = |G_{k-1}(i, p)|$  and  $\tilde{a} = |G_k(i, p)|$ , and again for  $a = |G_{k-1}(i, q)|$  and  $\tilde{a} = |G_k(i, q)|$ . If  $\rho$  is initialized to the sum of squares of the magnitudes of the off-diagonal elements of  $\widehat{G}_{k-1}$  instead of to zero, then the final  $\rho$  is  $w_k$  instead of  $h_k$ .

All virtual transformations in a step are mutually independent, so they should be performed in parallel. With enough memory available ( $O(n_0^3)$  in parallel), the computed elements of  $\widehat{U}_k^*$ ,  $\widehat{V}_k$ , and  $G_k$  could be stored, separately for each virtual transformation, and reused if the corresponding candidate has been selected as a pivot. Sequentially, only the data ( $O(n_0)$  values) of a candidate with the maximal weight should be stored. In the tested prototype of the algorithm the virtual transformations are performed in parallel, but all data generated, save the weights, are discarded.

It now emerges that each step requires at least  $O(n_0^2)$ , and at most  $O(n_0^3)$  operations just for computing all  $w_k$ . The complexity of the whole HSVD algorithm is therefore quartic in  $n_0$  in the general case, far worse than the usual cubic algorithms for the ordinary SVD. It is legitimate to ask what impedes development of a Kogbetliantz-type HSVD algorithm with a cubic complexity. For that, the pivot strategy should either ignore the weights and select the pivots in a prescribed order (as, e.g., the row-cyclic or the column-cyclic serial pivot strategies do), or compute all weights in a step with no more than a quadratic number of operations (e.g., by ignoring  $h_k$  in (3.1) and calculating the rest of  $w_k$ ). Either approach works well sometimes, but the former failed in the numerical experiments when  $n_0$  went up to 2000, and the latter even with  $n_0$  around 100, both with a catastrophic increase (overflow) of the off-diagonal norm. It remains an open question whether employing in both cases a much wider floating-point datatype (precision-wise as well as exponent-wise) could save the computation and eventually lead to convergence, but that is of more interest to theory than practice. The pivot strategy motivated here and described in detail in Section 4, however slow, is designed to keep the growth of the off-diagonal norm minimal (at least with a single pivot) when it is unavoidable, and can be used for the matrix orders of up to a few thousands with enough parallelism at hand. A faster but at least equally safe pivot strategy would make the  $J$ -Kogbetliantz algorithm competitive in performance with the pointwise one-sided hyperbolic Jacobi method.

It is remarkable that the extensive numerical tests conducted with many variants of the parallel blocked one-sided hyperbolic Jacobi method, all of them with some prescribed cyclic pivot strategy, either on CPU [30, 31] or on GPU [21, 23], have never shown an indication of a dangerous off-diagonal norm growth, even though the hyperbolic angle of a  $2 \times 2$  transformation might be of a large magnitude there as well, at least in theory. It may be interesting to look further into why that did not happen (and probably is hard to make happen) in practice, unlike with the Kogbetliantz-type algorithm, which requires such a complex pivot strategy to reduce the norm growth.



### 3.3 Floating-point considerations

The sum of squares from (3.1) can overflow, as well as each of the squares, leading to  $w_k = \infty$ . A dynamic rescaling of the whole  $G_{k-1}$  by an appropriate power of two could mitigate that issue, but with a risk that the smallest values by magnitude become subnormal and lose precision. However, if the weights are computed in a wide enough floating-point datatype (e.g., using the Intel's 80-bit extended), no overflow can occur. Similarly, one or both squares can underflow (a fact to be relied upon in Section 5) to a point of becoming zero(s). Then, if  $h_k = 0$ , the only way of avoiding  $w_k = 0$ , without computing in a wider datatype, is scaling  $G_{k-1}$  upwards, thus risking overflow of the largest elements by magnitude. As a partial remedy, the augmented weights from Section 4.1 are always distinct and well ordered, so a deterministic pivot selection is possible even with some (or all) weights being  $\infty$  or underflowing.

There is no rule of thumb how to properly prescale  $G_0$ , so that such issues, as well as the potential overflows due to the hyperbolic transformations, do not needlessly occur. Monitoring the computed weights can indicate should the latter problems be immediately avoided by downscaling  $G_{k-1}$ . In the tested prototype of the algorithm the dynamic scaling of the whole matrix, unlike the scaling from Section 2.1.1, has not been implemented, but should otherwise be if robustness is paramount.

## 4 Dynamic pivot selection based on weights

A dynamic pivot strategy (DPS in short) based on block weights was introduced in [3] for the two-sided block-Jacobi SVD algorithm (as the block-Kogbetliantz algorithm is also called), while the global and the asymptotic quadratic convergence of such a coupling was proven in [25] for the serial (a single block pair per step) and the parallel (multiple block pairs per step) annihilation. As the pointwise Jacobi algorithms are but a special case of the block ones, when the blocks (matrices) contain only one, scalar element, all properties of the dynamic pivoting hold in that context as well.

However, a DPS used in the pointwise Kogbetliantz-type HSVD algorithm differs in several aspects from the one for the SVD. A weight, i.e., the amount of the off-diagonal norm reduction, in the latter is finite (up to a possible floating-point overflow) and non-negative, while in the former it can be of arbitrary sign and infinite. Yet, the goal in both cases is the same: to reduce the off-diagonal norm in each step as much as possible. In the latter the off-diagonal norm growth is impossible, while in the former it is sometimes necessary, but is still kept as low as practicable.

Another important difference is in handling a situation when some or all weights are the same. In the former, a concept of augmented weight is introduced, as follows.

### 4.1 DPS in the sequential case

**Definition 4.1** Let the weight of a  $2 \times 2$  submatrix of  $G_{k-1}$  at the intersection of the  $p$ th and the  $q$ th row with the  $p$ th and the  $q$ th columns be computed according to (3.1) if that submatrix is a transformation candidate. Else, if the submatrix does not need

to be transformed, or cannot be transformed due to at least one its elements being non-finite, define its weight as a quiet NaN. Let a triple  $w_{pq}^{[k]} = (w, p, q)$  be called an *augmented weight*, where  $w$  is the weight of the submatrix induced by  $(p, q)$ . Also, let  $\mathbf{w}_k = \{w_{pq}^{[k]} \mid 1 \leq p < q \leq n_0\}$  be the set of all augmented weights in the  $k$ th step.

For any given  $k$ , a total order  $\leq$  can be defined on the augmented weights that makes all of them distinct, even though the weights themselves may be equal.

**Definition 4.2** Let, for some  $k$ ,  $a$  and  $b$  be two augmented weights in  $\mathbf{w}_k$ , and let them be considered equal, denoted as  $a = b$ , if and only if their corresponding components are equal, i.e.,  $a.w = b.w$ ,  $a.p = b.p$ , and  $a.q = b.q$ . Contrary to the usual definition of NaN, in this context let  $\text{NaN} = \text{NaN}$  and  $\text{NaN} < c$  for any other  $c$ . Let  $\leq$  be the union of the relations  $<$  and  $=$ , where  $a < b$  if and only if

1.  $a.w > b.w$ , or
2.  $a.w = b.w$  and  $a.q - a.p > b.q - b.p$ , or
3.  $a.w = b.w$ ,  $a.q - a.p = b.q - b.p$ , and  $a.q > b.q$ .

**Proposition 4.1** *The relation  $\leq$  from Definition 4.2 makes  $\mathbf{w}_k$  well ordered; specifically, every non-empty subset of  $\mathbf{w}_k$ , including  $\mathbf{w}_k$ , has a unique  $\leq$ -smallest element.*

*Proof* It is easy to verify that  $\leq$  is a total order on  $\mathbf{w}_k$ . Since  $\mathbf{w}_k$  is finite, it is well ordered by  $\leq$ . If all weights in  $S$ ,  $\emptyset \neq S \subseteq \mathbf{w}_k$ , are different, the smallest element is the one with the largest weight (due to condition 1 from Definition 4.2). The quantities  $a.q - a.p$  and  $b.q - b.p$  indicate a band, i.e., a sub/super-diagonal of  $G_{k-1}$  at which  $(q, p)$  and  $(p, q)$  lie, respectively, with the main diagonal being band 0. If several elements of  $S$  have the same maximal weight, the smallest element is the one among them in the farthest band (condition 2). If more than one such element exists, the smallest is the one lying lowest, i.e., with the largest column index (condition 3).  $\square$

The following Corollary 4.1 is a direct consequence of Proposition 4.1 and defines the DPS in the sequential case, i.e., when only one pivot is transformed in each step.

**Corollary 4.1** *Let  $\tilde{\mathbf{w}}_k = \mathbf{w}_k \setminus \{a \mid a.w = \text{NaN}\}$  be a set of augmented weights such that the weights themselves are not NaN. If  $\tilde{\mathbf{w}}_k = \emptyset$ , no transformations are possible and the algorithm stops with  $N = k$ . Else, let  $\hat{a}$  be the  $\leq$ -smallest element of  $\tilde{\mathbf{w}}_k$ . If  $\hat{a}.w = -\infty$ , no transformation is valid and the algorithm halts with an error. Else,  $(\hat{a}.p, \hat{a}.q)$  are the indices of a single pivot to be chosen in the  $k$ th step.*

Finding the smallest element of  $\tilde{\mathbf{w}}_k$  is linear in  $c = |\tilde{\mathbf{w}}_k|$ , i.e., at most quadratic in  $n_0$ , if a naïve method is used. However, any  $t$  disjoint subsets of  $\tilde{\mathbf{w}}_k$ , each of them of size at most  $\lceil c/t \rceil$  and at least one less that, can be linearly searched for their smallest elements, all of them in parallel. The smallest elements thus found can in turn be  $\leq$ -reduced in parallel with  $\lceil \log_2 t \rceil$  complexity to get the smallest element overall.

Furthermore, observe that only the weights in the pivot rows and columns change after a step. Then, in the next step, the weights in the changed positions have to be recomputed and compared with the unchanged weights in the remaining part of the matrix, for which the  $\preceq$ -smallest element can already be found in the previous step. Therefore, in each step two elements of  $\tilde{\mathbf{w}}_k$  have to be found: the  $\preceq$ -smallest one  $a$ , and its closest  $\preceq$ -successor  $b$  such that  $\{a.p, a.q\} \cap \{b.p, b.q\} = \emptyset$ . Finding such  $a$  and  $b$  would be easiest if  $\tilde{\mathbf{w}}_k$  would have already been sorted  $\preceq$ -ascendingly. But if such  $a$  and  $b$  are found, they define two pivots that can both be transformed in parallel, i.e., in a multi-step of length two. Repeating the observation of this paragraph, both a sketch of a method and an argument for the parallel DPS emerges, where a sequence of pivots, all with their indices disjoint, is incrementally built to be transformed in a multi-step. The case of a single pivot per step is here abandoned in favor of the parallel, multi-step case, albeit it can be noticed that some pivots in such a multi-step can lead to the off-diagonal norm growth when the  $\preceq$ -smallest one does not. The sequential case is thus locally (i.e., in each step, but not necessarily globally) optimal with respect to the change of the off-diagonal norm, but the parallel one may not be.

### 4.2 DPS in the multi-step case

For a multi-step  $\mathbf{k}$ , let the augmented weights  $w_{pq}^{[\mathbf{k}]}$  and the set  $\mathbf{w}_{\mathbf{k}}$  of them be defined as in Definition 4.1, with the smallest  $k \in \mathbf{k}$ . Definition 4.2, Proposition 4.1, and Corollary 4.1 are then modified accordingly. Also, let  $\widehat{\mathbf{w}}_{\mathbf{k}}$  be a  $\preceq$ -ascendingly sorted array of the elements of  $\tilde{\mathbf{w}}_{\mathbf{k}} \setminus \{a \mid a.w = -\infty\}$ . An option to get  $\widehat{\mathbf{w}}_{\mathbf{k}}$  from  $\tilde{\mathbf{w}}_{\mathbf{k}}$  is the parallel merge sort. In the prototype implementation, the Baudet–Stevenson odd-even sort with merge-splitting of the subarrays [2] is used, since it is simple and keeps  $\lceil t/2 \rceil$  tasks active at any given time, even though its worst-case complexity is quadratic. Both choices require a work array of  $c$  augmented weights, but that scratch space can be reused elsewhere. For  $t = 1$ , a sequential merge (or quick) sort is applicable.

**Definition 4.3** Let  $\mathbf{S}_{\mathbf{k}}$  be the set of all  $\preceq$ -ascending sequences of length at most  $|\mathbf{k}|$  of the augmented weights from  $\widehat{\mathbf{w}}_{\mathbf{k}}$  with non-intersecting indices, i.e., of all (not necessarily contiguous) subarrays of  $\widehat{\mathbf{w}}_{\mathbf{k}}$  of length at most  $|\mathbf{k}|$ , such that for any two elements  $a$  and  $b$  from a subarray holds  $\{a.p, a.q\} \cap \{b.p, b.q\} = \emptyset$ . Let  $S_{\mathbf{k}} \in \mathbf{S}_{\mathbf{k}}$  be arbitrary,  $m \geq 1$  be the length of  $S_{\mathbf{k}}$ , and define the following functions of  $S_{\mathbf{k}}$ ,

$$w(S_{\mathbf{k}}) = \sum_{\ell=1}^m S_{\mathbf{k}}(\ell).w, \quad \mathbf{o}(S_{\mathbf{k}}) = (l_{\ell} \mid S_{\mathbf{k}}(\ell) = \widehat{\mathbf{w}}_{\mathbf{k}}(l_{\ell}))_{\ell=1}^m,$$

as its weight and as a sequence of indices that its elements have in  $\widehat{\mathbf{w}}_{\mathbf{k}}$ , respectively.

It suffices to restrict Definition 4.3 to the sequences of length  $m > 0$  only, since  $|\widehat{\mathbf{w}}_{\mathbf{k}}| = 0$  implies that no valid transformations are possible, and the execution halts.

**Definition 4.4** Let  $|\mathbf{k}|$ ,  $\widehat{\mathbf{w}}_{\mathbf{k}}$ , and  $\ell$  such that  $1 \leq \ell \leq |\widehat{\mathbf{w}}_{\mathbf{k}}|$  be given. Then, a *parallel ordering*  $O_{\mathbf{k}}^{\ell} \in \mathbf{S}_{\mathbf{k}}$  is the sequence of maximal length, but not longer than  $|\mathbf{k}|$ , such

that  $O_{\mathbf{k}}^\ell(1) = \widehat{\mathbf{w}}_{\mathbf{k}}(\tau_1)$ , with  $\tau_1 = \ell$ , and  $O_{\mathbf{k}}^\ell(l) = \widehat{\mathbf{w}}_{\mathbf{k}}(\tau_l)$  for  $l > 1$ , where  $\tau_l$  is the smallest index of an element of  $\widehat{\mathbf{w}}_{\mathbf{k}}$  such that

$$\{\widehat{\mathbf{w}}_{\mathbf{k}}(\tau_l).p, \widehat{\mathbf{w}}_{\mathbf{k}}(\tau_l).q\} \cap \{\widehat{\mathbf{w}}_{\mathbf{k}}(\tau_i).p, \widehat{\mathbf{w}}_{\mathbf{k}}(\tau_i).q\} = \emptyset, \tag{4.1}$$

for all  $\tau_i$  such that  $1 \leq i < l$ . If  $O_{\mathbf{k}}^\ell$  is of length  $|\mathbf{k}|$ , it is denoted by  $PO_{\mathbf{k}}^\ell$ . A parallel DPS is a pivot strategy that for each  $\mathbf{k}$  finds  $O_{\mathbf{k}}^\ell$ , given an admissible  $\ell$ .

Given an admissible  $\ell$ ,  $O_{\mathbf{k}}^\ell$  from Definition 4.4 exists and is unique. Let its length be  $m_\ell$ . Then,  $\mathbf{o}(O_{\mathbf{k}}^\ell) = (\tau_l)_{l=1}^{m_\ell}$ . Taking the maximal  $m_\ell$  possible reflects an important choice of having most possible pivots per each multi-step transformed in parallel, even though it might imply that  $w(O_{\mathbf{k}}^\ell)$  is smaller than it would have been if only the first  $m'_\ell < m_\ell$  augmented weights were left in  $O_{\mathbf{k}}^\ell$ . Also,  $|\mathbf{k}|$  should be considered to stand for the *desired* number of steps in  $\mathbf{k}$ , until a parallel ordering has been found as described below and the actual, maybe lower, number of steps has been determined.

Algorithm 6 sequentially constructs a parallel ordering from Definition 4.4 for a given index  $\ell$  of  $\widehat{\mathbf{w}}_{\mathbf{k}}$ . It can also be constructed in parallel with Algorithm 7. Definition 4.4 indicates validity of Algorithms 6 and 7. All parallel constructs from here on in the paper are the OpenMP [27] ones, acting on the shared memory.

---

**Algorithm 6** Computing a parallel ordering sequentially.

---

```

Input:  $\widehat{\mathbf{w}}_{\mathbf{k}}, |\mathbf{k}|, n_0$ , and  $\ell, 1 \leq \ell \leq |\widehat{\mathbf{w}}_{\mathbf{k}}|$ .
Output:  $O_{\mathbf{k}}^\ell$  of length at most  $|\mathbf{k}|, 1 \leq |\mathbf{k}| \leq \lfloor n_0/2 \rfloor$ .
Data: a Boolean array R allocated on the thread's stack (faster), with  $n_0$ 
        elements initialized to  $\perp$ .
j :=  $\ell$ ; // start the search within  $\widehat{\mathbf{w}}_{\mathbf{k}}$  from the index  $\ell$ 
for i := 1 to  $|\mathbf{k}|$  do // each iteration appends an augmented
weight to  $O_{\mathbf{k}}^\ell$ 
     $O_{\mathbf{k}}^\ell(i) = \widehat{\mathbf{w}}_{\mathbf{k}}(j), j := j + 1;$  // append the jth augmented
weight from  $\widehat{\mathbf{w}}_{\mathbf{k}}$ 
     $R(O_{\mathbf{k}}^\ell(i).p) := \top, R(O_{\mathbf{k}}^\ell(i).q) := \top;$  // record  $O_{\mathbf{k}}^\ell(i).p$  and
 $O_{\mathbf{k}}^\ell(i).q$  as selected
    while  $j \leq |\widehat{\mathbf{w}}_{\mathbf{k}}|$  do // search the remaining part of  $\widehat{\mathbf{w}}_{\mathbf{k}}$ 
for the next  $\widehat{\mathbf{w}}_{\mathbf{k}}(j)$ 
         $F := R(\widehat{\mathbf{w}}_{\mathbf{k}}(j).p) \vee R(\widehat{\mathbf{w}}_{\mathbf{k}}(j).q);$  // check if  $\widehat{\mathbf{w}}_{\mathbf{k}}(j).p$  or
 $\widehat{\mathbf{w}}_{\mathbf{k}}(j).q$  are selected
        if  $\neg F$  then break; // if not, success; exit the loop
         $j := j + 1;$  // else, try with the next j
    end while
    if  $j > |\widehat{\mathbf{w}}_{\mathbf{k}}|$  then break; // exit if the next  $\widehat{\mathbf{w}}_{\mathbf{k}}(j)$  cannot be
found ( $|\mathbf{k}| := i$ )
end for

```

---

Algorithm 7 is the one chosen for the prototype implementation when  $t > 1$  (as was the case in the tests), with a fallback to Algorithm 6 when  $t = 1$ .

*Example 4.1* Let  $A$  be a matrix representation of the  $\widehat{\mathbf{w}}_k$  computed for some  $G_{k-1}$ ,

$$A = \begin{bmatrix} *_1 & 15 & 10 & 6 & 3 & \boxed{1} \\ 15 & *_2 & 14 & 9 & \boxed{5} & 2 \\ 10 & 14 & *_3 & \boxed{13} & 8 & 4 \\ 6 & 9 & \boxed{13} & *_4 & 12 & 7 \\ 3 & \boxed{5} & 8 & 12 & *_5 & 11 \\ \boxed{1} & 2 & 4 & 7 & 11 & *_6 \end{bmatrix}, \quad \begin{aligned} a_{pq} = l &\iff \widehat{\mathbf{w}}_k(l) = (w_{pq}, p, q); \\ PO_k^1 &= ((w_{16}, 1, 6), (w_{25}, 2, 5), (w_{34}, 3, 4)), \\ w(PO_k^1) &= w_{16} + w_{25} + w_{34}, \\ \mathbf{o}(PO_k^1) &= (1, 5, 13). \end{aligned}$$

---

**Algorithm 7** Computing a parallel ordering with  $t > 1$  tasks.

---

**Use:**  $\text{qNaN}(p)$  returns a quiet NaN with its payload set to  $p$ ;

$\text{isNotNaN}(x) = \top$  iff  $x$  is not a NaN.

**Input:**  $\widehat{\mathbf{w}}_k, |\mathbf{k}|, n_0, t$ , and  $\ell, 1 \leq \ell \leq |\widehat{\mathbf{w}}_k|$ .

**Output:**  $O_k^\ell$  of length at most  $|\mathbf{k}|, 1 \leq |\mathbf{k}| \leq \lfloor n_0/2 \rfloor$ .

```

i :=  $\ell, l := 0; \quad // i$  is the current index into  $\widehat{\mathbf{w}}_k$ 
while  $l < |\mathbf{k}|$  do  $// l$  is the current length of  $O_k^\ell$ 
     $l := l + 1, O_k^\ell(l) = \widehat{\mathbf{w}}_k(i); //$  append the current augmented
    weight to  $O_k^\ell$ 
    if  $l \geq |\mathbf{k}|$  then break;  $//$  return  $O_k^\ell$  if it is complete
     $a_p := O_k^\ell(l).p, a_q := O_k^\ell(l).q; \quad //$  shorthands for the
    indices
     $k := |\widehat{\mathbf{w}}_k| + 1; \quad // k > i$  will be  $\min j$  s.t.  $\widehat{\mathbf{w}}_k(j)$  does not
    collide with  $O_k^\ell(l)$ 
    for  $j = i + 1$  to  $|\widehat{\mathbf{w}}_k|$  do in parallel with threads( $t$ ), shared( $\widehat{\mathbf{w}}_k$ ),
    reduction( $\min : k$ )
        if  $\text{isNotNaN}(\widehat{\mathbf{w}}_k(j).w)$  then  $// \widehat{\mathbf{w}}_k(j)$  is alive if its
        weight is not a NaN
             $b_p := \widehat{\mathbf{w}}_k(j).p, b_q := \widehat{\mathbf{w}}_k(j).q; \quad //$  check if...
             $C := (a_p = b_p) \vee (a_p = b_q) \vee (a_q = b_p) \vee (a_q = b_q); // \widehat{\mathbf{w}}_k(j)$ 
            &  $O_k^\ell(l)$  collide
            if  $C$  then  $\widehat{\mathbf{w}}_k(j).w := \text{qNaN}(j)$  else  $k := \min\{k, j\}; \quad //$  kill
             $\widehat{\mathbf{w}}_k(j)$  if  $C$ 
        end if
    end parallel for
    if  $k > |\widehat{\mathbf{w}}_k|$  then break else  $i := k; \quad //$  take the lowest
    surviving index or stop
end while

```

---

Here,  $a_{pq} = a_{qp}$  for  $1 \leq p < q \leq n_0 = 6$  is the index of  $w_{pq}^{[k]} = (w_{pq}, p, q)$  in  $\widehat{\mathbf{w}}_k$ . Then,  $PO_k^1$  is obtained by either Algorithm 6 or 7, with the pivot index pairs denoted in  $A$  as well with the boxes of diminishing thickness, corresponding to the decreasing weights. Moreover, if all weights are the same (zeros, e.g.), then, by Definition 4.2,  $A$  always has the same form as above, enlarged or shrunk according to  $n_0$ , regardless of the actual elements in  $G_{k-1}$ . Near the end of the Kogbetliantz process,

where most weights are the same, the pivot strategy predictably selects many of the pivots in a pattern resembling the antidiagonal, and expanding towards the SE and the NW corners.

The initial ordering of the elements of  $\mathbf{w}_k$  is irrelevant mathematically. However, to simplify parallelization of the iteration over a triangular index space, i.e., over the indices of the strictly upper triangle of  $G_{k-1}$ , a one-dimensional array  $\mathcal{I}$  of  $\mathbf{n} = n_0(n_0 - 1)/2$  index pairs is pregenerated as follows. Let  $i_+ := 1$  and  $i_- := \mathbf{n}$ . For each index pair  $(p, q)$  in the column-cyclic order (column-by-column, and top-to-diagonal within each column), if  $J_0(p, p) = J_0(q, q)$  let  $\mathcal{I}(i_+) := (p, q)$  and  $i_+ := i_+ + 1$ ; else, let  $\mathcal{I}(i_-) := (p, q)$  and  $i_- := i_- - 1$ . This way  $\mathcal{I}$  is partitioned into two contiguous subarrays,  $\mathcal{I}_+$  and  $\mathcal{I}_-$  (which might be empty), of the index pairs leading to the trigonometric or to the hyperbolic transformations, respectively. One parallel loop over  $\mathcal{I}_+$  then computes all the weights in the trigonometric case, and another parallel loop over  $\mathcal{I}_-$  does the same in the hyperbolic case. Within each of the loops all iterations are of equal complexity, what would not have been the case if the two loops were fused into one, since computing the weights is essentially faster in the trigonometric ( $O(1)$  per weight) than in the hyperbolic ( $O(n_0)$  per weight) case.

## 5 Overview of the algorithm

In this section the convergence criterion, a vital part of the  $J$ -Kogbetliantz algorithm, is discussed in Section 5.1, and the algorithm is summarized in Section 5.2.

### 5.1 Convergence criterion

Traditionally, a convergence criterion for the Jacobi-like (including the Kogbetliantz-like) processes is simple and decoupled from the choice of a pivot strategy. However, in the hyperbolic case, where even a small off-diagonal norm can oscillate from one step to another, a “global” stopping criterion, based on the fall of the off-diagonal norm, relative to the initial one, below a certain threshold, or on stabilization of the diagonal values, e.g., does not suffice alone for stopping the process unattended.

The former approach may stop the process when the off-diagonal norm has relatively diminished, but when there may still be some valid transformations left that can both change the approximate singular values and raise the off-diagonal norm.

On the other hand, if the threshold has been set too low, the process may never (literally or practically) stop and may keep accumulating the superfluous transformations computed almost exclusively from the leftover rounding errors off the diagonal.

If a stopping criterion is solely based on observing that the diagonal elements, i.e., the approximate singular values, have not changed at all in a sequence of successive steps of a certain, predefined length, a few conducted tests indicate that the computed singular values are accurate in a sense of (6.2) and the diagonal has converged to its final value, but the singular vectors are not, with the relative error (6.1) of order  $\sqrt{\varepsilon}$ ,

where  $\varepsilon$  is the machine precision, i.e., the transformations left unperformed would have contributed to the singular vectors significantly, but not to the singular values.

A “local” convergence criterion is thus needed, based on the  $2 \times 2$  transformations in a multi-step belonging to a narrow class of matrices, as shown in Algorithm 8.

---

**Algorithm 8** The convergence criterion, evaluated at each multi-step  $\mathbf{k}$ .

---

```

b := 0; // a counter of the big steps
foreach  $k \in \mathbf{k}$  do // assume the  $k$ th transformations have
  been computed
  if  $\widehat{G}_{k-1}$  was diagonal then  $D = T$  else  $D = \perp$ ;
  if  $\widehat{U}_k^*$  is identity then  $U = T$  else  $U = \perp$ ;
  if  $\widehat{V}_k$  is identity then  $V = T$  else  $V = \perp$ ;
   $S = D \vee (U \wedge V)$ ; //  $S$  is  $T$  if the  $k$ th step is small
  if  $\neg S$  then  $b := b + 1$ ; // else, the  $k$ th step is big
end foreach
if  $b = 0$  then stop; // halt if no big steps in  $\mathbf{k}$ 
    
```

---

The steps of each multi-step  $\mathbf{k}$  are categorized as either *big* or *small*. A step is *big* if its  $2 \times 2$  pivot submatrix is not diagonal, and either the left or the right transformation is not identity; else, it is *small*. A non-trivial *small* step is just a scaling by the factors of unit modulus and/or a swap of the diagonal elements, so it is a heuristic but reasonable expectation that an absence of *big* steps is an indication of convergence.

### 5.2 The $J$ -Kogbetliantz algorithm

The  $J$ -Kogbetliantz algorithm is summarized in Algorithm 9. Note that accumulating the left and the right singular vectors is optional, and that  $\Sigma$  is the diagonal of  $G_N$ .

## 6 Numerical testing

Testing was performed on the Intel Xeon Phi 7210 CPUs, running at 1.3 GHz with TurboBoost turned off in Quadrant cluster mode, with 96 GiB of RAM and 16 GiB of flat-mode MCDRAM (which was not used), under 64-bit CentOS Linux 7.9.2009 with the Intel compilers (Fortran, C), version 19.1.3.304, and the GNU compilers (Fortran, C), version 9.3.1, for the error checking. No BLAS/LAPACK routines from Intel Math Kernel Library were used in the final prototype implementation.

The prototype code has been written in Fortran for the `DOUBLE PRECISION` and `DOUBLE COMPLEX` datatypes, with some auxiliary routines written in C. The real and the complex  $J$ -Kogbetliantz algorithms are implemented as two programs. There are also two error checkers in quadruple precision (Fortran’s `KIND=REAL128`), one which finds the absolute and then the relative normwise error of the obtained HSVD as

$$\|G_0 - U \Sigma V^{-1}\|_F / \|G_0\|_F, \tag{6.1}$$

while the other compares  $\Sigma J_0 \Sigma^T$  with the eigenvalues  $\Lambda$  of  $H = G_0 J_0 G_0^*$ , i.e.,

$$\max_{1 \leq i \leq n_0} |(\lambda'_{ii} - \sigma_{ii}^2 J_{ii}) / \lambda'_{ii}|, \quad \lambda'_{11} \geq \lambda'_{22} \geq \dots \geq \lambda'_{n_0 n_0}, \tag{6.2}$$

where  $\Lambda' = P_\Lambda \Lambda P_\Lambda^T$  ( $P_\Lambda$  being a permutation) has the eigenvalues on the diagonal sorted descendingly to match the ordering of  $\Sigma J_0 \Sigma^T$ . All eigenvalues are non-zero.

---

**Algorithm 9** Overview of the  $J$ -Kogbetliantz algorithm.

---

**Input:**  $G_0$  and  $J_0$ , preprocessed from  $G$  and  $J$ , if necessary.  
**Output:**  $N$ ,  $U$ ,  $\Sigma$ , and  $V^{-1}$ , as described in Section 1.  
 $N := 0;$  // a counter of steps performed in the loop below  
**repeat** // the loop body is a multi-step  $\mathbf{k}$   
    compute the augmented weights  $\mathbf{w}_k$  as described in Section 3.2 and Section 4;  
    filter and sort  $\mathbf{w}_k$  to obtain  $\widehat{\mathbf{w}}_k$  as described in Section 4.2;  
    **if**  $|\widehat{\mathbf{w}}_k| = 0$  **then break;** // terminate early, if possible  
    generate the parallel ordering  $O_k^1$  by Algorithm 7 (or by Algorithm 6 if  $t = 1$ );  
    **foreach**  $k \in \mathbf{k}$  **do** // in parallel with  $t$  threads  
        compute (or reuse) the  $2 \times 2$  HSVD of  $\widehat{G}_{k-1}$ , i.e.,  $\widehat{U}_k^*$ ,  $\widehat{\Sigma}_k$ , and  $\widehat{V}_k$ , as in Section 2;  
        apply  $\widehat{U}_k^*$  from the left to  $U_{k-1}^*$  and  $G_{k-1}$  to obtain  $U_k^*$  and  $G'_{k-1}$ , resp., as in Section 3;  
    **end foreach**  
    **foreach**  $k \in \mathbf{k}$  **do** // in parallel with  $t$  threads  
        apply  $\widehat{V}_k$  from the right to  $G'_{k-1}$  and  $V_{k-1}$  to obtain  $G_k$  and  $V_k$ , resp., as in Section 3;  
    **end foreach**  
     $N := N + |\mathbf{k}|;$  // end of the multi-step  $\mathbf{k}$   
**until** convergence detected by Algorithm 8;  
compute  $U = (U_N^*)^*$  and  $V^{-1} = J_0 V_N^* J_0;$  // optionally

---


The close-to-exact eigenvalues are known since each  $H$  has been generated by taking its double precision eigenvalues  $\Lambda$  pseudorandomly from one of the ranges:

1.  $\lambda \in \langle \varepsilon, 1 \rangle$ , drawn uniformly from  $\langle 0, 1 \rangle$ ,
2.  $|\lambda| \in \langle \varepsilon, 1 \rangle$ , drawn uniformly from  $[-1, 1]$ ,
3.  $|\lambda| \in \langle \varepsilon, 1 \rangle$ , drawn from the normal variable  $\mathcal{N}(\mu = 0, \sigma = 1)$ ,

with a given  $\varepsilon \in \{\varepsilon_1 = 10^{-13}, \varepsilon_2 = 10^{-15}\}$ . Then,  $H = U \Lambda U^*$  (or  $U \Lambda U^T$ ) is formed by applying  $n_0 - 1$  pseudorandom Householder reflectors to  $\Lambda$  in extended precision.

The Hermitian/symmetric indefinite factorization with complete pivoting [29, 33] of  $H$  gives  $J_0$  and  $G'_0$ , which is rounded to a double (complex/real) precision input







**Fig. 2** The relative errors (6.1), in  $\log_{10}$ -scale, in the HSVD computed in `DOUBLE PRECISION` (left) and `DOUBLE COMPLEX` (right) datatypes. The matrix orders on  $x$ -axis are in  $\log_2$ -scale

$G_0$ . For each  $n_0$  twelve pairs  $(G_0, J_0)$  have been generated, six each for the real and the complex case. In each case two pairs come with  $J_0 = I_{n_0}$ , corresponding to the first range above. For a given  $n_0, \varepsilon$ , and a range, the eigenvalues of the real  $H$  are the same as those of the complex  $H$ , due to a fixed pseudo-RNG seed selected for that  $\varepsilon$ .

For each field  $T \in \{\mathbb{R}, \mathbb{C}\}$ , range  $L \in \{1, 2, 3\}$  of the eigenvalues of  $H$ , and  $\varepsilon$  as above, a sequence of test matrices was generated, with their orders ranging from  $n_0 = 4$  to  $n_0 = 2048$  with a variable step: four up to  $n_0 = 128$ , eight up to  $n_0 = 256$ , 16 up to  $n_0 = 512$ , 32 up to  $n_0 = 1024$ , and 256 onwards. For each  $n_0$ , the number of tasks for a run of the  $J$ -Kogbetliantz algorithm was  $t = \min\{64, n_0(n_0 - 1)/2\}$ , since the CPU has 64 cores, and to each core at most one task (i.e., an OpenMP thread) was assigned by setting `OMP_PLACES=CORES` and `OMP_PROC_BIND=SPREAD` environment variables.



**Fig. 3** The maximal relative errors (6.2), in  $\log_{10}$ -scale, in the eigenvalues of  $G_0 J_0 G_0^*$ , with the  $J_0$ -HSVD of  $G_0$  computed in `DOUBLE PRECISION` (left) and `DOUBLE COMPLEX` (right) datatypes. The matrix orders on  $x$ -axis are in  $\log_2$ -scale



**Fig. 4** The number of cycles (6.3) until convergence, when computing in DOUBLE PRECISION (left) and DOUBLE COMPLEX (right) datatypes. The matrix orders on  $x$ -axis are in  $\log_2$ -scale

Let  $\mathbf{N}$ ,  $0 \leq \mathbf{N} \leq N$ , be the number of multi-steps performed until convergence. Then, define  $\mathbf{C}$ , the number of “virtual” sweeps (also called cycles) performed, as

$$\mathbf{C} = \mathbf{N}/(n_0 - 1). \quad (6.3)$$

A “virtual” sweep has at most the same number of steps as would a “real” sweep by a cyclic pivot strategy have, i.e.,  $n_0(n_0 - 1)/2$ , but in it any transformation candidate can be transformed up to  $\lfloor n_0/2 \rfloor$  times. Note that  $\mathbf{C}$  does not have to be an integer.

In each subfigure of Figs. 2, 3, and 4 there are three data series, one for each  $L$ . A data point in a series is the maximum of a value from one run of the  $J$ -Kogbetliantz algorithm on a matrix generated with  $\varepsilon = \varepsilon_1$ , and a value from another run on a matrix generated with  $\varepsilon = \varepsilon_2$ , with all other parameters (i.e.,  $T$ ,  $L$ , and  $n_0$ ) being the same.

A comparison of the relative errors in the decomposition, shown in Fig. 2, leads to a similar conclusion that can be reached by comparing the maximal relative errors in the eigenvalues of  $H$ , shown in Fig. 3. In the real as well as in the complex case a satisfactory accuracy, in a sense of both (6.1) and (6.2), was reached in a reasonably small number of cycles, as shown in Fig. 4.

A sequential (with  $t = 1$ ) variant of the algorithm, performing one step at a time, was compared performance-wise against the parallel multi-step one, with  $\nu = 0.75$  and  $n_0$  going up to 128 and 256 for  $T = \mathbb{C}$  and  $T = \mathbb{R}$ , respectively. The sequential variant was drastically slowed down (up to more than three orders of magnitude) compared to the multi-step one, especially for the “true” HSVD (less so for the “ordinary” SVD), to a point of being totally impractical. The  $J$ -Kogbetliantz algorithm is therefore best run in the multi-step regime, with as much parallelism as possible.

## 7 Conclusions and future work

In this paper an accurate method for computing the  $2 \times 2$  HSVD of real and complex matrices is demonstrated and employed as one of the three major building blocks

of a  $J$ -Kogbetliantz algorithm for general square matrices. The other two important contributions are a heuristic but efficient convergence criterion for all pointwise Kogbetliantz-type processes and a modification of the well-established dynamic pivot strategy that can cope with the pivot weights of arbitrary magnitudes and signs.

To keep the exposition concise, a forward rounding error analysis of the floating-point computation of the  $2 \times 2$  HSVD is left for future work. Furthermore, performing such analysis is slightly impeded by, e.g., a lack of standardized, tight error bounds for the absolute value of a complex number, or equivalently, of the `HYPOT` intrinsic.

The  $J$ -Kogbetliantz algorithm, as presented, is not highly performant even in its parallel form. It is worth exploring if (and what kind of) blocking of the algorithm would be beneficial in terms of performance, without negatively affecting accuracy. A straightforward generalization of the dynamic pivot strategy to block (instead of  $2 \times 2$ ) pivots seems too inefficient, as it would assume at each block-multi-step the full diagonalization of all possible block pivots that require hyperbolic transformations only to compute their weights. Apart from—and complementary to—blocking, there are other options for improving performance, like storing and reusing the  $2 \times 2$  HSVDs computed while forming a multi-step, as explained in Section 3, and implementing a vectorized sorting routine for the suitably represented augmented weights.

The batches of  $2 \times 2$  transformations in each multi-step could be processed in a vectorized way, as in [22], should a highly optimized implementation be required. Such a version of the algorithm would have to resort to a specific vectorized routine in each of the three cases of  $2 \times 2$  transformations (one trigonometric and two hyperbolic, of lower and upper triangular matrices). Consequently, (up to) three disjoint batches would have to be processed separately, with a non-trivial repacking of input data for each of them. A serious practical use-case is required to justify such effort.

Finally, an interesting observation is offered without a proof. For complex matrices, the  $J$ -Kogbetliantz algorithm seems to converge (in limit) not only in the sense of  $\text{off}_F^2(G_k) \rightarrow 0$  and  $\text{diag}(G_k) \rightarrow \Sigma$ , but also with  $\max_{i \neq j} |\arg(G_k)_{ij}| \rightarrow 0$ , for  $k \rightarrow \infty$ .

### Appendix. Proofs of Lemmas 3.1 and 3.2

*Proof (Lemma 3.1)* Let, for  $1 \leq \ell \leq n$ ,  $\gamma_\ell = \arg(x_\ell)$  and  $\delta_\ell = \arg(y_\ell)$ . Then,

$$\begin{bmatrix} x_\ell & y_\ell \end{bmatrix} = \begin{bmatrix} e^{i\gamma_\ell} |x_\ell| & e^{i\delta_\ell} |y_\ell| \end{bmatrix} = e^{i\gamma_\ell} \begin{bmatrix} |x_\ell| & e^{i(\delta_\ell - \gamma_\ell)} |y_\ell| \end{bmatrix} = e^{i\delta_\ell} \begin{bmatrix} e^{i(\gamma_\ell - \delta_\ell)} |x_\ell| & |y_\ell| \end{bmatrix}.$$

Using the second equality, from the matrix multiplication it follows

$$x'_\ell = e^{i\gamma_\ell} (|x_\ell| \cosh \psi + e^{i(\delta_\ell - \gamma_\ell - \beta)} |y_\ell| \sinh \psi),$$

and using the third equality, from the matrix multiplication it follows

$$y'_\ell = e^{i\delta_\ell} (e^{i(\gamma_\ell - \delta_\ell + \beta)} |x_\ell| \sinh \psi + |y_\ell| \cosh \psi).$$

Since  $|e^{-i\gamma_\ell} x'_\ell| = |x'_\ell|$  and  $|e^{-i\delta_\ell} y'_\ell| = |y'_\ell|$  and  $\cos(-\phi) = \cos \phi$ , it holds

$$\begin{aligned} |x'_\ell|^2 &= (|x_\ell| \cosh \psi + \cos(\delta_\ell - \gamma_\ell - \beta) |y_\ell| \sinh \psi)^2 + (\sin(\delta_\ell - \gamma_\ell - \beta) |y_\ell| \sinh \psi)^2 \\ &= |x_\ell|^2 \cosh^2 \psi + |y_\ell|^2 \sinh^2 \psi + \cos(\delta_\ell - \gamma_\ell - \beta) |x_\ell| |y_\ell| 2 \cosh \psi \sinh \psi, \\ |y'_\ell|^2 &= (\cos(\gamma_\ell - \delta_\ell + \beta) |x_\ell| \sinh \psi + |y_\ell| \cosh \psi)^2 + (\sin(\gamma_\ell - \delta_\ell + \beta) |x_\ell| \sinh \psi)^2 \\ &= |x_\ell|^2 \sinh^2 \psi + |y_\ell|^2 \cosh^2 \psi + \cos(\delta_\ell - \gamma_\ell - \beta) |x_\ell| |y_\ell| 2 \cosh \psi \sinh \psi. \end{aligned}$$

After grouping the terms, the square of the Frobenius norm of the new  $\ell$ th row is

$$\begin{aligned} \|[x'_\ell \ y'_\ell]\|_F^2 &= |x'_\ell|^2 + |y'_\ell|^2 = (\cosh^2 \psi + \sinh^2 \psi)(|x_\ell|^2 + |y_\ell|^2) \\ &\quad + 2 \cos(\delta_\ell - \gamma_\ell - \beta) |x_\ell| |y_\ell| 2 \cosh \psi \sinh \psi. \end{aligned} \tag{A.1}$$

Summing the left side of (A.1) over all  $\ell$  one obtains

$$\|[\mathbf{x}' \ \mathbf{y}']\|_F^2 = \sum_{\ell=1}^n (|x'_\ell|^2 + |y'_\ell|^2),$$

what is equal to the right side of (A.1), summed over all  $\ell$ ,

$$\sum_{\ell=1}^n \left( (\cosh^2 \psi + \sinh^2 \psi)(|x_\ell|^2 + |y_\ell|^2) + 2\zeta_\ell |x_\ell| |y_\ell| 2 \cosh \psi \sinh \psi \right),$$

where  $-1 \leq \zeta_\ell = \cos(\delta_\ell - \gamma_\ell - \beta) \leq 1$ , so  $|\zeta_\ell| \leq 1$ . The last sum can be split into a non-negative part and the remaining part of an arbitrary sign,

$$(\cosh^2 \psi + \sinh^2 \psi) \sum_{\ell=1}^n (|x_\ell|^2 + |y_\ell|^2) + 2 \cosh \psi \sinh \psi \sum_{\ell=1}^n 2\zeta_\ell |x_\ell| |y_\ell|.$$

Using the triangle inequality, and observing that  $\sum_{\ell=1}^n (|x_\ell|^2 + |y_\ell|^2) = \|[\mathbf{x} \ \mathbf{y}]\|_F^2$ , this value can be bounded above by

$$(\cosh^2 \psi + \sinh^2 \psi) \|[\mathbf{x} \ \mathbf{y}]\|_F^2 + 2 \cosh \psi |\sinh \psi| \sum_{\ell=1}^n 2|x_\ell| |y_\ell|,$$

and below by

$$(\cosh^2 \psi + \sinh^2 \psi) \|[\mathbf{x} \ \mathbf{y}]\|_F^2 - 2 \cosh \psi |\sinh \psi| \sum_{\ell=1}^n 2|x_\ell| |y_\ell|,$$

where both bounds can be simplified by the identities

$$\cosh^2 \psi + \sinh^2 \psi = \cosh(2\psi), \quad 2 \cosh \psi |\sinh \psi| = |\sinh(2\psi)|.$$

By the inequality of arithmetic and geometric means it holds  $2|x_\ell| |y_\ell| \leq (|x_\ell|^2 + |y_\ell|^2)$ , so a further upper bound is reached as

$$\cosh(2\psi) \|[\mathbf{x} \ \mathbf{y}]\|_F^2 + |\sinh(2\psi)| \|[\mathbf{x} \ \mathbf{y}]\|_F^2,$$

and a further lower bound as

$$\cosh(2\psi) \|[\mathbf{x} \ \mathbf{y}]\|_F^2 - |\sinh(2\psi)| \|[\mathbf{x} \ \mathbf{y}]\|_F^2,$$

what, after grouping the terms and dividing by  $\|[\mathbf{x} \ \mathbf{y}]\|_F^2$ , concludes the proof.  $\square$

*Proof (Lemma 3.2)* Note that  $\cosh(2\psi) + |\sinh(2\psi)| \geq \cosh(2\psi) \geq 1$ , and

$$\begin{aligned} 1 &= \cosh^2(2\psi) - \sinh^2(2\psi) \\ &= (\cosh(2\psi) - |\sinh(2\psi)|) \cdot (\cosh(2\psi) + |\sinh(2\psi)|) \\ &\geq \cosh(2\psi) - |\sinh(2\psi)| > 0. \end{aligned}$$

If  $\psi = 0$ , the equalities in the bounds established in Lemma 3.1 hold trivially. Also, if both equalities hold simultaneously,  $\psi = 0$ .

The inequality of arithmetic and geometric means in the proof of Lemma 3.1 turns into equality if and only if  $|x_\ell| = |y_\ell|$  for all  $\ell$ . When  $|x_\ell||y_\ell| \neq 0$ , it has to hold  $\zeta_\ell = \zeta$ , where  $\zeta = \pm \text{sign}(\sinh \psi)$ , to reach the upper or the lower bound, respectively. From  $\zeta = \pm 1$  it follows  $\delta_\ell = \gamma_\ell + \beta + l\pi$  for a fixed  $l \in \mathbb{Z}$ , i.e.,  $x_\ell = e^{i\gamma_\ell}|x_\ell|$  and  $y_\ell = \pm e^{i\beta} e^{i\gamma_\ell}|x_\ell|$ , so  $y_\ell = \pm e^{i\beta} x_\ell$  for all  $\ell$ . Conversely,  $\mathbf{y} = \pm e^{i\beta} \mathbf{x}$  implies, for all  $\ell$ , that  $|x_\ell| = |y_\ell|$  and  $\zeta_\ell$  is a constant  $\zeta = \pm 1$ , so one of the two bounds is reached.  $\square$

**Acknowledgements** We are much indebted to Saša Singer<sup>†</sup> for his suggestions on the paper’s subject, and to the anonymous referee for significantly improving the presentation of the paper.

**Author contribution** The second author formulated the research topic, reviewed the literature, plotted the figures and proofread the manuscript. The first author performed the rest of the research tasks.

**Funding** This work has been supported in part by Croatian Science Foundation under the project IP–2014–09–3670 “Matrix Factorizations and Block Diagonalization Algorithms” (MFBDA).

**Data availability** Several animations and the full testing dataset with the matrix outputs are available at <http://euridika.math.hr:1846/Jacobi/JKogb/> or by request.

**Code availability** The source code is available in <https://github.com/venovako/JKogb> repository.

**Declarations**

**Conflict of interest** The authors declare no competing interests.

**References**

1. Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D. LAPACK Users’ Guide. Software, Environments and Tools, 3rd edn. Society for Industrial and Applied Mathematics, Philadelphia (1999). <https://doi.org/10.1137/1.9780898719604>
2. Baudet, G., Stevenson, D.: Optimal sorting algorithms for parallel computers. IEEE Trans. Comput. **C-27**(1), 84–87 (1978). <https://doi.org/10.1109/TC.1978.1674957>
3. Bečka, M., Okša, G., Vajteršić, M.: Dynamic ordering for a parallel block-Jacobi SVD algorithm. Parallel Comp. **28**(2), 243–262 (2002). [https://doi.org/10.1016/S0167-8191\(01\)00138-7](https://doi.org/10.1016/S0167-8191(01)00138-7)
4. Bojanczyk, A.W., Onn, R., Steinhardt, A.O.: Existence of the hyperbolic singular value decomposition. Linear Algebra Appl. **185**(C), 21–30 (1993). [https://doi.org/10.1016/0024-3795\(93\)90202-Y](https://doi.org/10.1016/0024-3795(93)90202-Y)

5. Charlier, J.P., Vanbegin, M., Van Dooren, P.: On efficient implementations of Kogbetliantz's algorithm for computing the singular value decomposition. *Numer. Math.* **52**(3), 279–300 (1987). <https://doi.org/10.1007/BF01398880>
6. Drmač, Z.: Implementation of Jacobi rotations for accurate singular value computation in floating point arithmetic. *SIAM J. Sci. Comput.* **18**(4), 1200–1222 (1997). <https://doi.org/10.1137/S1064827594265095>
7. Hari, V., Matejaš, J.: Accuracy of two SVD algorithms for  $2 \times 2$  triangular matrices. *Appl. Math. Comput.* **210**(1), 232–257 (2009). <https://doi.org/10.1016/j.amc.2008.12.086>
8. Hari, V., Singer, S., Singer, S.: Block-oriented  $J$ -Jacobi methods for Hermitian matrices. *Linear Algebra Appl.* **433**(8–10), 1491–1512 (2010). <https://doi.org/10.1016/j.laa.2010.06.032>
9. Hari, V., Singer, S., Singer, S.: Full block  $J$ -Jacobi method for Hermitian matrices. *Linear Algebra Appl.* **444**, 1–27 (2014). <https://doi.org/10.1016/j.laa.2013.11.028>
10. Hari, V., Veselić, K.: On Jacobi methods for singular value decompositions. *SIAM J. Sci. Stat.* **8**(5), 741–754 (1987). <https://doi.org/10.1137/0908064>
11. Hari, V., Zadelj-Martić, V.: Parallelizing the Kogbetliantz method: A first attempt. *J. Numer. Anal. Ind. Appl. Math.* **2**(1–2), 49–66 (2007)
12. IEEE Computer Society: 754-2019 - IEEE Standard for Floating-Point Arithmetic. IEEE, New York. <https://doi.org/10.1109/IEEESTD.2019.8766229> (2019)
13. ISO/IEC JTC1/SC22/WG5 ISO/IEC 1539-1:2018(en) Information Technology — Programming languages — Fortran — Part 1: Base language, 4th edn. ISO, Geneva (2018)
14. Kogbetliantz, E.G.: Solution of linear equations by diagonalization of coefficients matrix. *Quart. Appl. Math.* **13**(2), 123–132 (1955). <https://doi.org/10.1090/qam/88795>
15. Kulikov, G.Yu., Kulikova, M.V.: Hyperbolic-singular-value-decomposition-based square-root accurate continuous-discrete extended-unscented Kalman filters for estimating continuous-time stochastic models with discrete measurements. *Int. J. Robust Nonlinear Control* **30**, 2033–2058 (2020). <https://doi.org/10.1002/rnc.4862>
16. Kulikova, M.V.: Hyperbolic SVD-based Kalman filtering for Chandrasekhar recursion. *IET Control Theory A.* **13**(10), 1525–1531 (2019). <https://doi.org/10.1049/iet-cta.2018.5864>
17. Kulikova, M.V.: Square-root approach for Chandrasekhar-based maximum correntropy Kalman filtering. *IEEE Signal Process Lett.* **26**(12), 1803–1807 (2019). <https://doi.org/10.1109/LSP.2019.2948257>
18. Mackey, D.S., Mackey, N., Tisseur, F.: Structured factorizations in scalar product spaces. *SIAM J. Matrix Anal. and Appl.* **27**(3), 821–850 (2005). <https://doi.org/10.1137/040619363>
19. Matejaš, J., Hari, V.: Accuracy of the Kogbetliantz method for scaled diagonally dominant triangular matrices. *Appl. Math. Comput.* **217**(8), 3726–3746 (2010). <https://doi.org/10.1016/j.amc.2010.09.020>
20. Matejaš, J., Hari, V.: On high relative accuracy of the Kogbetliantz method. *Linear Algebra Appl.* **464**, 100–129 (2015). <https://doi.org/10.1016/j.laa.2014.02.024>
21. Novaković, V.: A hierarchically blocked Jacobi SVD algorithm for single and multiple graphics processing units. *SIAM J. Sci. Comput.* **37**(1), C1–C30 (2015). <https://doi.org/10.1137/140952429>
22. Novaković, V.: Batched computation of the singular value decompositions of order two by the AVX-512 vectorization. *Parallel Process. Lett.* **30**(4), 2050015 (2020). <https://doi.org/10.1142/S0129626420500152>
23. Novaković, V., Singer, S.: A GPU-based hyperbolic SVD algorithm. *BIT* **51**(4), 1009–1030 (2011). <https://doi.org/10.1007/s10543-011-0333-5>
24. NVIDIA Corp.: CUDA C++ Programming Guide v10.2.89. <https://docs.nvidia.com/cuda/cuda-c-programming-guide/> (2019)
25. Okša, G., Yamamoto, Y., Bečka, M., Vajteršič, M.: Asymptotic quadratic convergence of the two-sided serial and parallel block-Jacobi SVD algorithm. *SIAM J. Matrix Anal. and Appl.* **40**(2), 639–671 (2019). <https://doi.org/10.1137/18M1222727>
26. Onn, R., Steinhardt, A.O., Bojanczyk, A.W.: The hyperbolic singular value decomposition and applications. *IEEE Trans. Signal Process.* **39**(7), 1575–1588 (1991). <https://doi.org/10.1109/78.134396>
27. OpenMP ARB: OpenMP Application Programming Interface Version 5.0. <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5.0.pdf> (2018)
28. Singer, S.: Indefinite QR factorization. *BIT* **46**(1), 141–161 (2006). <https://doi.org/10.1007/s10543-006-0044-5>
29. Singer, S., Di Napoli, E., Novaković, V., Čaklović, G.: The LAPW method with eigendecomposition based on the Hari–Zimmermann generalized hyperbolic SVD. *SIAM J. Sci. Comput.* **42**(5), C265–C293 (2020). <https://doi.org/10.1137/19M1277813>

30. Singer, S., Singer, S., Novaković, V., Davidović, D., Bokulić, K., Ušćumlić, A.: Three-level parallel  $J$ -Jacobi algorithms for Hermitian matrices. *Appl. Math. Comput.* **218**(9), 5704–5725 (2012). <https://doi.org/10.1016/j.amc.2011.11.067>
31. Singer, S., Singer, S., Novaković, V., Ušćumlić, A., Dunjko, V.: Novel modifications of parallel Jacobi algorithms. *Numer. Algorithms* **59**(1), 1–27 (2012). <https://doi.org/10.1007/s11075-011-9473-6>
32. Slapničar, I.: Accurate symmetric eigenreduction by a Jacobi method. Ph.D. thesis, FernUniversität-Gesamthochschule, Hagen (1992)
33. Slapničar, I.: Componentwise analysis of direct factorization of real symmetric and Hermitian matrices. *Linear Algebra Appl.* **272**, 227–275 (1998). [https://doi.org/10.1016/S0024-3795\(97\)00334-0](https://doi.org/10.1016/S0024-3795(97)00334-0)
34. Stewart, G.W.: An updating algorithm for subspace tracking. *IEEE Trans. Signal Process.* **40**(6), 1535–1541 (1992). <https://doi.org/10.1109/78.139256>
35. Veselić, K.: A Jacobi eigenreduction algorithm for definite matrix pairs. *Numer. Math.* **64**(1), 241–269 (1993). <https://doi.org/10.1007/BF01388689>
36. Zha, H.: A note on the existence of the hyperbolic singular value decomposition. *Linear Algebra Appl.* **240**, 199–205 (1996). [https://doi.org/10.1016/0024-3795\(94\)00197-9](https://doi.org/10.1016/0024-3795(94)00197-9)

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.