Compact Order Polynomial Singular Value Decomposition of a Matrix of Analytic Functions

Mohammed A. Bakhit[1], Faizan A. Khattak[1], Ian K. Proudler[1], Stephan Weiss[1], and Garrey W. Rice[2]
[1] Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, Scotland
[2] MathWorks, 180 West George Street Glasgow, G2 2NR, Scotland
{mohammed.bakhit,faizan.khattak,ian.proudler,stephan.weiss}@strath.ac.uk

***Abstract*—This paper presents a novel method for calculating a compact order singular value decomposition (SVD) of polynomial matrices, building upon the recently proven existence of an analytic SVD for analytic, non-multiplexed polynomial matrices. The proposed method calculates a conventional SVD in sample points on the unit circle, and then applies phase smoothing Algorithms to establish phase-coherence between adjacent frequency bins. This results in the extraction of compact order singular values and their corresponding singular vectors. The method is evaluated through experiments conducted on an ensemble of randomised polynomial matrices, demonstrating its superior performance in terms of higher decomposition accuracy and lower polynomial order compared to state-of-the-art techniques.**

I. Introduction

The singular value decomposition (SVD) is a widely used technique in linear algebra that allows the diagonalisation of any matrix $\mathbf{A} \in \mathbb{C}^{M\times N}$ into $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathrm{H}}$, where $\mathbf{U}$ and $\mathbf{V}$ are complex unitary matrices, and $\mathbf{\Sigma}$ is diagonal and real-valued [1, 2]. This decomposition has found extensive utilisation in signal processing generally [3, 4]. The SVD has been particularly attractive for the design of multiple-input multiple-output (MIMO) narrowband communications systems, where $\mathbf{A}$ is the channel matrix of complex gain factors [5–14] . Here, the decomposition afforded by the SVD has been shown to enable solutions for e.g. precoding and equalisation that are optimal in a variety of criteria, such as in the least squares and channel capacity sense.

In broadband systems, the focus shifts from complex-valued gain factors to impulse responses between sources and sensors. Consequently, a MIMO communications system with N transmitters and M receivers can be represented as a matrix $\mathbf{A}[n]$, where every matrix entry $a_{i,j}[n]$, $i = 1,\dots,M$, $j = 1,\dots,N$, is an impulse response in the discrete time index $n \in \mathbb{Z}$. Therefore, its z-transform $\boldsymbol{A}(z) = \sum_n \mathbf{A}[n]z^{-n}$ is now a matrix of transfer functions, or a polynomial matrix. This means that the matrix decomposition achieved by the conventional SVD can only diagonalize the matrix for a particular lag or frequency if operating in the frequency domain. One possible approach is to decompose the broadband problem into discrete Fourier transform (DFT) bins and perform independent processing on each bin. However, this neglects the spectral coherence and often leads to suboptimal solutions [15–17]. To address the above limitations, the polynomial SVD (PSVD) has gained recent attention for general polynomial matrices $\boldsymbol{A}(z) : \mathbb{C} \to \mathbb{C}^{M\times N}$.

The PSVD was often computed via two polynomial eigenvalue decompositions (PEVDs) which have been designed for para-Hermitian matrices. These are supported by theory [18–20] and two families of well-known PEVD algorithms: the second order sequential best rotation (SBR2) [21] [22–24] and sequential matrix diagonalisation (SMD) [25, 26] algorithms. This method of PSVD computation is costly and only provides approximate diagonalisation with low accuracy [27, 28]. For a more accurate PSVD, a polynomial QR decomposition (PQRD) [28] has been utilised but requires repeated application of the PQRD algorithm which again is costly.

Algorithms specifically designed to directly address the PSVD, rather than relying on two PEVDs and/or multiple PQRDs, are the generalised SBR2 (GSBR) [29] and generalised SMD (GSMD) methods. These methods expand upon SBR2 and SMD, extending their applicability from para-Hermitian matrices to general polynomial matrices. While the GSBR2 and GSMD methods are faster compared to previous algorithms, they still suffer from the limitation of providing approximate diagonalisation and likely yield high-order polynomial factors than strictly necessary [30].

In this document, we present a DFT-based PSVD approach which offers accurate diagonalisation. In order to re-establish spectral coherence across DFT bins, the proposed method exploits the existence of an analytic SVD with infinitely differentiable and therefore smooth factors, which exists for almost all analytic matrices $\boldsymbol{A}(z)$ [30]. For this, we first independently calculate conventional SVDs in each DFT bin. To restore the phase-coherence between bin-wise SVD results, phase smoothing is performed in order to obtain compact time-domain support for both paraunitary matrices $\boldsymbol{U}(z)$ and $\boldsymbol{V}(z)$.

This paper is organised as follows: Sec. II explores the existence of the analytic SVD. Sec. III provides a brief overview of existing polynomial SVD algorithms to approximate the analytic SVD. In Sec. IV, the proposed method is presented in detail. A comparison of the proposed method utilising some designed metrics to perform a valid comparison with state-of-the-art methods, is conducted in Sec. V. Finally, the paper concludes with a summary and discussion in Sec. VI.

The work of M. Bakhit is funded by Mathworks and the University of Strathclyde. F. Khattak is the recipient of a scholarship of the Commonwealth Scholarship Commission. The work was also in parts supported by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/S000631/1 and the MOD University Defence Research Collaboration in Signal Processing.

This is a peer-reviewed, accepted author manuscript of the following conference paper: Bakhit, M. A., Khattak, F. A., Proudler, I. K., Weiss, S., & Rice, G. W. Compact order polynomial singular value decomposition of a matrix of analytic functions. 1-5. Paper presented at 9th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Los Suenos, Costa Rica.

II. Analytic Singular Value Decomposition

For a matrix $\boldsymbol{A}(z) : \mathbb{C} \to \mathbb{C}^{M\times N}$ that is analytic in $z \in \mathbb{C}$ on some annulus including at least the unit circle, e.g. a matrix containing transfer functions of stable, causal systems, there exists an analytic SVD

$$\boldsymbol{A}(z) = \boldsymbol{U}(z)\boldsymbol{\Sigma}(z)\boldsymbol{V}^{\mathrm{P}}(z)\,, \tag{1}$$

where $\boldsymbol{\Sigma}(z)$ contains the singular values $\sigma_i(z)$, $i = 1,\ldots,M$ on its diagonal [30]. On the unit circle, we expect $\boldsymbol{\Sigma}(\mathrm{e}^{\mathrm{j}\Omega}) \in \mathbb{R}$, but for analytic singular values these must not be constrained to be non-negative [30–32]. The two paraunitary matrices $\boldsymbol{U}(z) \in \mathbb{C}^{M\times M}$ and $\boldsymbol{V}(z) \in \mathbb{C}^{N\times N}$ i.e. $\boldsymbol{U}(z)\boldsymbol{U}(z)^{\mathrm{P}} = \mathbf{I}$, $\boldsymbol{V}(z)\boldsymbol{V}^{\mathrm{P}}(z) = \mathbf{I}$ carry the left- and right-singular vectors, respectively, in their columns, where the parahermitian operator $\{\cdot\}^{\mathrm{P}}$ implies a Hermitian transposition and time reversal such that $\boldsymbol{U}^{\mathrm{P}}(z) = \{\boldsymbol{U}^{\mathrm{H}}(1/z^*)\}^{\mathrm{H}}$. Regarding the dimensions of $\boldsymbol{A}(z)$, we assume $N \geq M$; otherwise we can consider the analytic SVD of $\boldsymbol{A}^{\mathrm{P}}(z)$.

While the singular values in $\boldsymbol{\Sigma}(z)$ are unique up to a permutation, the singular vectors are ambiguous, such that if $\boldsymbol{u}_i(z)$ is a valid ith singular vector, then so is $\boldsymbol{u}_i'(z) = \boldsymbol{u}_i(z)\phi_i(z)$, where $\phi_i, i = 1,\ldots,M$ is an arbitrary allpass function [30]. Note that for a finite order $\boldsymbol{u}_i(z)$, $\boldsymbol{u}_i'(z)$ will have infinite order unless $\phi_i(z)$ is a simple delay; therefore, the selection of $\phi_i(z)$ has a direct impact of the compactness of the factors. The allpass function $\phi_i(z)$ is common to each pair $\boldsymbol{u}_i(z), \boldsymbol{v}_i(z)$, $i = 1,\ldots,M$; if this coupling is violated, then the corresponding $\sigma_i(z)$ would likely no longer be real on the unit circle [30].

The analyticity of the SVD factors guarantees that while they may have infinite order, an arbitrarily good approximation can be achieved using finite length polynomials, such that with $\boldsymbol{U}(z)$, $\boldsymbol{\Sigma}(z)$, and $\boldsymbol{V}(z)$ now of finite order, the equal sign in (1) is replaced with an approximation sign. This approximation by finite order terms is referred to as a polynomial SVD (PSVD).

If $\boldsymbol{A}(z)$ is estimated, e.g. via system identification from finite data [33], then analogous to the case of the analytic EVD, the singular values are likely to be spectrally majorised [34], i.e. we have that

$$\sigma_i(\mathrm{e}^{\mathrm{j}\Omega}) > \sigma_{i+1}(\mathrm{e}^{\mathrm{j}\Omega}), i = 1,\ldots,(M-1),\ \forall\ \Omega\,. \tag{2}$$

Therefore, in this paper, we assume that $\boldsymbol{A}(z)$ is positive semi-definite and its singular values satisfy (2).

III. Iterative PSVD Algorithms

The Kogbetliantz based PSVD method [29] is an extension of the SBR2 [21] algorithm from para-Hermitian to general polynomial matrices, and therefore here referred to as the GSBR2 algorithm. This extension applies Kogbetliantz transformations [35], which is the extension of the Jacobi transformation to non-symmetric matrices, to eliminate the maximum off-diagonal element at every iteration. It works in two steps: in the first step, the maximum off-diagonal component is shifted to the time-zero component using delay operations. In the second step, a Givens rotation transfers the energy of this element onto the diagonal. These two steps are repeated until the maximum off-diagonal element satisfies a predefined threshold. This method has been proven to converge [35], and demonstrates better diagonalisation than using two PEVDs as suggested in [21].

In contrast, the GSMD algorithm transfers more energy to the diagonal in each iteration [27]. It accomplishes this by selecting the column with the maximum ℓ_2 or ℓ_∞ norm and shifting it to the zero-lag component, followed by performing a complete SVD of this time-zero matrix. As a result, GSMD can achieve faster convergence and yields lower-order paraunitary matrices compared to GSBR2 [27]. Nevertheless, it is important to note that the obtained order tends to remain relatively high when compared to any known ground-truth order of the factors in (1).

Both GSBR2 and GSMD represent time domain approaches that exhibit a propensity to increase the order of the polynomial SVD factors with every iteration step. This tendency becomes more pronounced when striving for complete diagonalisation. For this reason, in the next section we adopt a fundamentally different DFT-domain approach.

IV. Proposed Method

In this section we adopt a DFT-based method that operates similarly to recent analytic eigenvalue decomposition algorithms [36–39]. We first characterise the general approach before in turn extracting the singular values and thereafter the left- and right-singular vectors.

A. Overview of Approach

Because of its analyticity, it suffices to investigate $\mathbf{A}(z)$ on the unit circle for $z = \mathrm{e}^{\mathrm{j}\Omega}$. There, applying a K-point DFT to $\mathbf{A}[n]$, we evaluate an SVD in every frequency bin $\Omega_k = 2\pi k/K$, with $k = 0,\ldots,(K-1)$, such that

$$\boldsymbol{A}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{U}_k\boldsymbol{\Sigma}_k\mathbf{V}_k^{\mathrm{H}} = \mathbf{U}_k\boldsymbol{\Psi}_k\boldsymbol{\Sigma}_k\boldsymbol{\Psi}_k^{\mathrm{H}}\mathbf{V}_k^{\mathrm{H}}\,, \tag{3}$$

where the diagonal unitary matrix $\boldsymbol{\Psi}_k = \mathrm{diag}\{\mathrm{e}^{\mathrm{j}\psi_{1,k}},\ldots,\mathrm{e}^{\mathrm{j}\psi_{M,k}}\}$ can commute with $\boldsymbol{\Sigma}_k$ and therefore expresses the phase ambiguity of the singular vectors [2]. With the restriction of $\mathbf{A}(z)$ to possess non-negative and strictly majorised singular values, motivated by [34], the sample points of the analytic SVD in (1) can be related to (3) as [39]

$$\boldsymbol{\Sigma}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \boldsymbol{\Sigma}_k \tag{4}$$

$$\boldsymbol{U}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{U}_k\boldsymbol{\Psi}_k \tag{5}$$

$$\boldsymbol{V}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{V}_k\boldsymbol{\Psi}_k\,. \tag{6}$$

We now need to exploit the right hand sides of (4)-(6) to determine the sample points of the analytic SVD.

The extraction of the singular values from (4) is straightforward, as long as a sufficient DFT size K is known. Determining K, and therefore extracting sufficiently accurate singular values, will be addressed in Sec. IV-B. We thereafter target the singular vectors, and phase smoothing via the term $\boldsymbol{\Psi}_k$ in (5) and (6), in Sec. IV-C.

B. Singular Value Extraction

With (4), we only need to determine a sufficient DFT size that results in low enough time-domain aliasing when obtaining $\hat{\mathbf{\Sigma}}[n]$ via an inverse DFT from $\mathbf{\Sigma}_k$. This time domain aliasing can be measured as [38]

$$\zeta_{\sigma_m} = \sum_{\tau} \frac{|\hat{\sigma}_m^{(K)}[\tau] - \hat{\sigma}_m^{(K/2)}[\tau]|^2}{|\hat{\sigma}_m^{(K)}[\tau]|^2}, \text{ for } m = 1, \ldots, M \ , \quad (7)$$

where $\hat{\sigma}_m^{(K)}[\tau]$ represents the time-domain equivalent of $[\sigma_{m,1}, \ldots, \sigma_{m,K}]$ obtained via the $K-$point IDFT. The DFT size is doubled until ζ_{σ_m} in (7) falls below some threshold ϵ.

C. Singular Vector Extraction

In order to extract analytic and hence infinitely differentiable singular vectors, we need to determine the phase values $\psi_{m,k}$, $k = 0, \ldots, (K-1)$, $m = 1, \ldots, M$, such that we achieve phase coherence across frequency bins. Note that the phase ambiguity of left- and right-singular values is coupled; we therefore only focus on $\mathbf{U}_k$.

In order to perform phase smoothing over the ith component of the mth singular vector across all K bins, we define

$$\mathbf{r}_{i,m} = [u_{i,m,1}, \ldots, u_{i,m,K}]^{\mathrm{T}} \ . \quad (8)$$

This can be related to phase adjustments as

$$\mathbf{s}_{i,m} = \mathrm{diag}\{\mathbf{a}_m\}\, \mathbf{r}_{i,m} \ , \quad (9)$$

where $\mathbf{s}_{i,m} = [s_{i,m,1}, \ldots, s_{i,m,K}]^{\mathrm{T}}$ with $s_{i,m,k} = \mathrm{e}^{\mathrm{j}\psi_{m,k}} u_{i,m,k}$ and $\mathbf{a}_m = [\mathrm{e}^{\mathrm{j}\psi_{m,1}}, \ldots, \mathrm{e}^{\mathrm{j}\psi_{m,K}}]^{\mathrm{T}}$.

1) Maximally Smooth Interpolation: The primary objective is to achieve a smooth interpolations across the sample points in $\mathbf{s}_{i,m}$ by optimising $\mathbf{a}_m$. To ensure maximum smoothness in terms of the shortest possible support in the time domain, the interpolations, denoted as $\hat{U}_{i,m}^{(K)}(\mathrm{e}^{\mathrm{j}\Omega})$ with K the number of sample points, can be obtained through Dirichlet interpolation [39]:

$$\begin{aligned}\hat{U}_{i,m}^{(K)}(\mathrm{e}^{\mathrm{j}\Omega}) &= \frac{1}{K} \sum_{k=0}^{k-1} s_{i,m,k} \sum_{n=0}^{k-1} \mathrm{e}^{-\mathrm{j}(\Omega - 2\pi k/K)n} \\ &= \frac{1}{\sqrt{K}} \mathbf{e}_K^{\mathrm{H}}(\mathrm{e}^{\mathrm{j}\Omega}) \cdot \mathbf{W}_K^H \cdot \mathbf{s}_{i,m} \ ,\end{aligned}$$

where $\mathbf{e}_K^{\mathrm{H}}(\mathrm{e}^{\mathrm{j}\Omega}) = [1, \mathrm{e}^{-\mathrm{j}\Omega}, ..., \mathrm{e}^{-\mathrm{j}\Omega(K-1)}]^{\mathrm{T}}$, and $\mathbf{W}_K$ is a K-point unitary DFT matrix.

2) Measuring Smoothness of the Interpolation: To measure smoothness of the interpolation $\hat{U}_{i,m}^{(K)}(\mathrm{e}^{\mathrm{j}\Omega})$, we can use the power of its pth derivative, denoted as ξ_p^m, as a suitable metric. This metric therefore quantifies the smoothness of any extracted singular vector. By minimizing ξ_p^m, we thus optimize the phase values in $\mathbf{a}_m$ in (9) for all m. For ξ_p^m, we have

$$\xi_p^m = \sum_{i=1}^{M} \frac{1}{2\pi} \int_0^{2\pi} \left| \frac{\partial^p \hat{U}_{i,m}^{(K)}(\mathrm{e}^{\mathrm{j}\Omega})}{\partial \Omega_p} \right|^2 \mathrm{d}\Omega \ , \quad (10)$$

where the derivative is given as

$$\frac{\partial^p \hat{U}_{i,m}^{(K)}(\mathrm{e}^{\mathrm{j}\Omega})}{\partial \Omega_p} = \frac{1}{\sqrt{K}} \mathbf{e}_K^{\mathrm{H}}(\mathrm{e}^{\mathrm{j}\Omega}) \cdot \mathbf{D}_k^p \cdot \mathbf{W}_K^{\mathrm{H}} . \mathbf{u}'_{i,m}$$

with $\mathbf{D}_k^p = \mathrm{diag}\{0, -\mathrm{j}, -2\mathrm{j}, \ldots, -(K-1)\mathrm{j}\}^p$. Using Parseval's theorem [40], for any vector $\mathbf{x} \in \mathbb{C}^K$, we have

$$\frac{1}{2\pi} \int_0^{2\pi} \left| \mathbf{e}_K^{\mathrm{H}}(\mathrm{e}^{\mathrm{j}\Omega}) \mathbf{x} \right|^2 d\Omega = \mathbf{x}^{\mathrm{H}} \mathbf{x} \ .$$

Therefore, ξ_p^m can be written as

$$\xi_p^m = \frac{1}{K} \sum_{i=1}^{M} \| \mathbf{D}_k^p \cdot \mathbf{W}_K^{\mathrm{H}} \cdot \mathbf{s}_{i,m} \|_2^2 \ , \quad (11)$$

which can be further modified as

$$\xi_p^m = \mathbf{a}_m^{\mathrm{H}} \mathbf{C}_{m,K,p} \mathbf{a}_m \ , \quad (12)$$

where $\mathbf{C}_{m,K,p}$ is a matrix defined as

$$\mathbf{C}_{m,K,p} = \sum_{i=1}^{M} \mathrm{diag}\{\mathbf{r}_{i,m}^{\mathrm{H}}\} \cdot \mathbf{W}_K \cdot \mathbf{D}_k^{2p} \cdot \mathbf{W}_K^{\mathrm{H}} \cdot \mathrm{diag}\{\mathbf{r}_{i,m}\} \ .$$

3) Optimization of the Objective Function: The objective function can be minimized through the Newton Raphson technique. The iterative update for the phase vector $\boldsymbol{\psi}_m = [\psi_{m,1}, \ldots, \psi_{m,K}]^{\mathrm{T}}$ at the ℓth iteration is

$$\boldsymbol{\psi}_m[\ell + 1] = \boldsymbol{\psi}_m[\ell] - \rho \mathbf{H}^{-1} \frac{\partial \xi_p^m(\boldsymbol{\psi}_m[\ell])}{\partial \boldsymbol{\psi}_m} \ . \quad (13)$$

where

$$\frac{\partial \xi_p^m}{\partial \boldsymbol{\psi}_m} = 2\mathrm{Im}\{\mathrm{diag}\{\mathbf{a}_m^*\}\, \mathbf{C}_{m,K,p} \mathbf{a}_m\} \ .$$

The Hessian matrix $\mathbf{H}$, after dropping subscripts m for ease, is

$$2\mathrm{Re}\{\mathrm{diag}\{\mathbf{a}^*\}\, \mathbf{C}_{K,p} \mathrm{diag}\{\mathbf{a}\} - \mathrm{diag}\{\mathrm{diag}\{\mathbf{a}^*\}\, \mathbf{C}_{K,p} \mathbf{a}\}\}$$

and can be made positive definite by dropping the second, generally negligible component [41], thus guaranteeing convergence of (14).

4) Stopping Criterion: At the end of the ℓth iteration, we have $\boldsymbol{\psi}_m[\ell + 1]$, which in turn gives us values for $\mathbf{a}_m$ in (9) and for $\mathbf{\Psi}_k$ in (5), thus obtaining a phase-adjusted left-singular matrics $\hat{\mathbf{U}}_k$, with $\hat{\mathbf{U}}^{(K)}[n]$ its K-point IDFT. With this matrix, we define a combined error in terms of time domain aliasing and lack of orthonormality as

$$\zeta_u = \sum_{\tau \in \mathbb{Z}} \| \hat{\mathbf{U}}^{(K)\mathrm{H}}[-\tau] * \hat{\mathbf{U}}^{(K)}[\tau] - \delta[\tau] \mathbf{I} \|_{\mathrm{F}}^2 \ , \quad (14)$$

where $\| \cdot \|_{\mathrm{F}}$ is the Frobenius norm and $\{*\}$ denotes the convolution operator. If the above metric is below a low threshold ε, the optimization process can be terminated. Then $\hat{\boldsymbol{U}}(z)$ is close to paraunitary, and satisfies this property for all frequencies on the unit circle, and no longer just at the K sample points.

5) Right-Singular Vectors and Reduced SVD: Once the phase is adjusted for the left-singular vectors, for a square matrix $\boldsymbol{A}(z)$ the phase ambiguity and phase smoothing are also settled for its right-singular vectors via (6). For a rectangular matrix $\boldsymbol{A}(z)$ with $N \geq M$, the right-singular vectors $\boldsymbol{v}_i(z),\ i = M+1, \ldots, N$ can form an arbitrary orthonormal basis in the nullspace of $\boldsymbol{A}(z)$. For these columns therefore, across bins we have to first create smooth 1-d subspaces for each of the singular vectors before performing phase smoothing. We are not aware of any algorithm capable of performing this yet. The proposed method therefore only computes a reduced PSVD, where instead of $\hat{\boldsymbol{V}}(z) : \mathbb{C} \to \mathbb{C}^{N\times N}$, we return $\hat{\boldsymbol{V}}'(z) : \mathbb{C} \to \mathbb{C}^{N\times M}$, with $\hat{\boldsymbol{V}}'^{\mathrm{P}}(z)\hat{\boldsymbol{V}}'(z) = \mathbf{I}$.

V. Simulations and Results

A. Ensemble Test Settings

In this section, the proposed compact order PSVD method is compared to the state-of-the-art algorithms GSBR2 [29] and GSMD [27] through an ensemble of randomised, analytic and non-multiplexed polynomial matrices. The ensemble is constructed of 100 instantiations of $\boldsymbol{A}(z) = \boldsymbol{U}(z)\boldsymbol{\Sigma}(z)\boldsymbol{V}^{\mathrm{P}}(z) : \mathbb{C} \to \mathbb{C}^{3\times 5}$ and the test is repeated at various orders of ground-truth singular values and singular vectors, i.e. $\mathrm{Ord}\{\boldsymbol{U}(z)\} = \mathrm{Ord}\{\boldsymbol{V}(z)\} \in \{2, 4, \ldots, 14\}$ with $\mathrm{Ord}\{\boldsymbol{\Sigma}(z)\} \in \{4, 8, \ldots, 28\}$ that is spectrally majorised on the unit circle. The proposed method is executed with $\epsilon = 10^{-8}$. The phase smoothing procedure will continue until either $\varepsilon = 10^{-5}$ is satisfied or maximum of 200 iterations are reached . For the benchmark algorithms GSRB2 and GSMD a maximum of 500 iterations is permitted, unless the magnitude of the maximum off-diagonal element falls below 10^{-6}. Moreover, during the execution, the intermediate paraunitary and diagonalised matrices are order-limited through a truncation threshold of 10^{-5} [42–45].

B. Performance Metrics

For just comparison, apart form considering execution time, the resulting orders of the paraunitary $\hat{\boldsymbol{U}}(z)$ and $\hat{\boldsymbol{V}}(z)$ is also defined as a performance metric. Additionally, for gauging the accuracy of every decomposition method, we define a reconstruction error metric as

$$\xi_R = \frac{\sum_n ||\mathbf{A}[n] - \hat{\mathbf{A}}[n]||_{\mathrm{F}}^2}{\sum_n ||\mathbf{A}[n]||_{\mathrm{F}}^2}\ , \tag{15}$$

where $\hat{\mathbf{A}}[n] = \hat{\mathbf{U}}[n] * \bar{\hat{\mathbf{\Sigma}}}[n] * \hat{\mathbf{V}}^{\mathrm{H}}[-n]$. The matrix of singular values, $\hat{\mathbf{\Sigma}}[n]$ is identical to $\hat{\mathbf{\Sigma}}[n]$ with its off-diagonal entries set to zero in case of the GSBR2 and GSMD algorithms; in case of the proposed method, $\hat{\mathbf{\Sigma}}[n]$ is diagonal by design.

C. Ensemble Results

The ensemble results for all performance metrics are illustrated in Fig. 1. Especially the order of the resulting paraunitary matrix $\hat{\boldsymbol{U}}(z)$ shown in Fig. 1(a), when estimated by the proposed method, is orders of magnitude lower than the state-of-the-art algorithms. Regarding the computational expense, the proposed method outperforms GSBR2, but its performance is nearly comparable or slightly worse than GSMD, as demonstrated in Fig. 1(b); the cost for the proposed approach in this figure increases in steps — this occurs where ever the DFT length needs to be increased by another factor of two. Finally, Fig.1(c) demonstrates that the proposed method achieves the least reconstruction error compared to GSBR2 and GSMD which implies that the proposed method attains a greater degree of diagonalisation with a more compact order for $\hat{\boldsymbol{U}}(z)$.

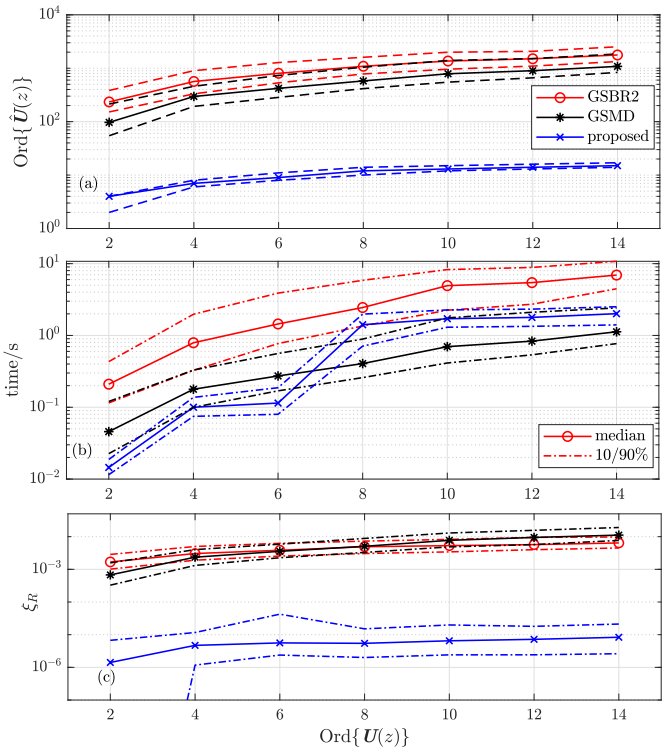

Fig. 1. Ensemble results showing (a) order of the resulting $\hat{\boldsymbol{U}}(z)$, (b) execution time and (c) reconstruction error.

VI. Conclusions

Motivated by the analytic SVD, we have proposed an algorithms that can extract analytic singular values and singular vectors for the case of a general, but spectrally majorised polynomial matrix. This algorithm operates in the DFT domain. Under the restriction of spectral majorisation, the extraction of singular vectors with compact support is straightforward. Determining the singular vectors involved a phase smoothing operation in order to find, within their range of ambiguity, a compact solution. For non-square polynomial matrices, the approach has been restricted to returning a reduced SVD result.

The proposed method has demonstrated substantial advantages over the state-of-the-art GSBR2 and GSMD methods, and in simulations has outperformed the benchmark methods in terms of compact order, improved reconstruction accuracy, and execution time, making it a compelling choice for various practical applications.

REFERENCES

[1] G. Strang, *Linear Algebra and Its Applications*, 2nd ed. New York: Academic Press, 1980.

[2] G.H. Golub and C.F. van Loan, *Matrix Computations*, 3rd ed. Baltimore, Maryland: John Hopkins University Press, 1996.

[3] R.J. Vaccaro, "Svd and signal processing, ii. algorithms, analysis and applications," *Amsterdam: Elsevier*, 1991.

[4] E.F. Deprettere, *SVD and signal processing: algorithms, applications and architectures*. North-Holland Publishing Co., 1989.

[5] C.H. Ta and S. Weiss, "A design of precoding and equalisation for broadband MIMO systems," in *41st Asilomar Conf. on Signals, Systems and Computers*, Pacific Grove, CA, USA, pp. 1616–1620, Nov. 2007.

[6] C.H. Ta and S. Weiss, "A jointly optimal precoder and block decision feedback equaliser design with low redundancy," in *15th European Signal Processing Conference*, pp. 489–492, Poznan, Poland, 2007.

[7] W. Al-Hanafy, A.P. Millar, C.H. Ta, and S. Weiss, "Broadband SVD and non-linear precoding applied to broadband mimo channels," in *42nd Asilomar Conf. Signals, Systems and Computers*, pp. 2053–2057, 2008.

[8] N. Moret, A. Tonello, and S. Weiss, "MIMO precoding for filter bank modulation systems based on PSVD," in *IEEE 73rd Vehicular Technology Conf. — Spring*, pp. 1–5, Budapest, Hungary, May 2011.

[9] S. Weiss, N. Moret, A.P. Millar, A. Tonello, and R.W. Stewart, "Initial results on an MMSE precoding and equalisation approach to MIMO PLC channels," in *2011 IEEE Int. Symposium on Power Line Communications and Its Applications*, pp. 146–152, 2011.

[10] A. Ahrens, A. Sandmann, S. Lochmann, and Z. Wang, "Decomposition of optical MIMO systems using polynomial matrix factorization," in *2nd IET Int. Conference on Intelligent Signal Processing*, pp. 1–6, 2015.

[11] D. Hassan, S. Redif, J.G. McWhirter, and S. Lambotharan, "Polynomial GSVD beamforming for two-user frequency-selective mimo channels," *IEEE Trans. Signal Processing*, vol. 69, pp. 948–959, 2021.

[12] X. Mestre and D. Gregoratti, "A parallel processing approach to filter-bank multicarrier mimo transmission under strong frequency selectivity," in *IEEE Int. Conference on Acoustics, Speech and Signal Processing*, pp. 8078–8082, 2014.

[13] A.I. Perez-Neira, M. Caus, R. Zakaria, D. Le Ruyet, E. Kofidis, M. Haardt, X. Mestre, and Y. Cheng, "Mimo signal processing in offset-QAM based filter bank multicarrier systems," *IEEE Trans. Signal Processing*, vol. 64, no. 21, pp. 5733–5762, 2016.

[14] H. Zamiri-Jafarian and M. Rajabzadeh, "A polynomial matrix SVD approach for time domain broadband beamforming in MIMO-OFDM systems," in *IEEE Vehicular Technology Conference — Spring*,pp. 802–806, 2008.

[15] S. Weiss and I.K. Proudler, "Comparing efficient broadband beamforming architectures and their performance trade-offs," in *14th Int. Conf. on Digital Signal Processing Proceedings*, vol. 1, pp. 417–423, 2002.

[16] W. Kellermann and H. Buchner, "Wideband algorithms versus narrowband algorithms for adaptive filtering in the DFT domain," in *37th Asilomar Conf. Signals, Systems & Computers*, vol. 2 pp. 1278–1282, 2003.

[17] V.W. Neo, S. Redif, J.G. McWhirter, J. Pestana, I.K. Proudler, S. Weiss, and P.A. Naylor, "Polynomial eigenvalue decomposition for multichannel broadband signal processing," *IEEE Signal Processing Mag.*, to appear.

[18] S. Weiss, J. Pestana, and I.K. Proudler, "On the existence and uniqueness of the eigenvalue decomposition of a parahermitian matrix," *IEEE Trans. Signal Processing*, vol. 66, no. 10, pp. 2659–2672, May 2018.

[19] S. Weiss, J. Pestana, I.K. Proudler, and F.K. Coutts, "Corrections to 'On the existence and uniqueness of the eigenvalue decomposition of a parahermitian matrix'," *IEEE Trans. Signal Processing*, vol. 66, no. 23, pp. 6325–6327, Dec. 2018.

[20] G. Barbarino and V. Noferini, "On the Rellich eigendecomposition of para-Hermitian matrices and the sign characteristics of $*$-palindromic matrix polynomials," *Linear Algebra and its Applications*, vol. 672, pp. 1–27, Sep. 2023.

[21] J.G. McWhirter, P.D. Baxter, T. Cooper, S. Redif, and J. Foster, "An EVD algorithm for para-Hermitian polynomial matrices," *IEEE Trans. Signal Processing*, vol. 55, no. 5, pp. 2158–2169, 2007.

[22] S. Redif, J.G. McWhirter, and S. Weiss, "Design of FIR paraunitary filter banks for subband coding using a polynomial eigenvalue decomposition," *IEEE Trans. Signal Processing*, vol. 59, no. 11, pp. 5253–5264, 2011.

[23] S. Redif, S. Weiss, and J.G. McWhirter, "An approximate polynomial matrix eigenvalue decomposition algorithm for para-hermitian matrices," in *11th IEEE Int. Symposium on Signal Processing and Information Technology*, Bilbao, Spain, pp. 421–425, Dec. 2011.

[24] Z. Wang, J.G. McWhirter, J. Corr, and S. Weiss, "Multiple shift second order sequential best rotation algorithm for polynomial matrix EVD," in *23rd European Signal Processing Conference*, Nice, France, pp. 844–848, Sep. 2015.

[25] S. Redif, S. Weiss, and J.G. McWhirter, "Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices," *IEEE Trans. Signal Processing*, vol. 63, no. 1, pp. 81–89, 2015.

[26] J. Corr, K. Thompson, S. Weiss, J. McWhirter, S. Redif, and I. Proudler, "Multiple shift maximum element sequential matrix diagonalisation for parahermitian matrices," in *IEEE Workshop on Statistical Signal Processing*, Gold Coast, Australia, pp. 312–315, June 2014.

[27] F.A. Khattak, I.K. Proudler, and S. Weiss, "Generalised sequential matrix diagonalisation for the SVD of polynomial matrices," in *Int. Conference on Sensor Signal Processing for Defence*, Edinburgh, Scotland, Sep. 2023.

[28] J.A. Foster, J.G. McWhirter, M.R. Davies, and J.A. Chambers, "An algorithm for calculating the QR and singular value decompositions of polynomial matrices," *IEEE Trans. Signal Processing*, vol. 58, no. 3, pp. 1263–1274, 2009.

[29] J.G. McWhirter, "An algorithm for polynomial matrix SVD based on generalised Kogbetliantz transformations," in *18th European Signal Processing Conference*, Aalborg, Denmark, pp. 457–461, Sep. 2010

[30] S. Weiss, I.K. Proudler, G. Barbarino, J. Pestana, and J.G. McWhirter, "On properties and structure of the analytic singular value decomposition," *IEEE Trans. Signal Processing*, submitted.

[31] B. De Moor and S. Boyd, "Analytic properties of singular values and vectors," KU Leuven, Tech. Rep., 1989.

[32] A. Bunse-Gerstner, R. Byers, V. Mehrmann, and N.K. Nicols, "Numerical computation of an analytic singular value decomposition of a matrix valued function," *Numer. Math*, vol. 60, pp. 1–40, 1991.

[33] F.A. Khattak, I.K. Proudler, and S. Weiss, "Enhanced space-time covariance estimation based on a system identification approach," in *Int. Conference on Sensor Signal Processing for Defence*, London, UK, Sep. 2022.

[34] F.A. Khattak, S. Weiss, I.K. Proudler, and J.G. McWhirter, "Space-time covariance matrix estimation: Loss of algebraic multiplicities of eigenvalues," in *56th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 2022.

[35] G.E. Forsythe and P. Henrici, "The cyclic Jacobi method for computing the principal values of a complex matrix," *Transactions of the American Mathematical Society*, vol. 94, no. 1, pp. 1–23, 1960.

[36] S. Weiss, I.K. Proudler, F.K. Coutts, and J. Pestana, "Iterative approximation of analytic eigenvalues of a parahermitian matrix EVD," in *IEEE Int. Conference on Acoustics, Speech and Signal Processing*, Brighton, UK, May 2019.

[37] S. Weiss, I. Proudler, F. Coutts, and J. Deeks, "Extraction of analytic eigenvectors from a parahermitian matrix," in *International Conference on Sensor Signal Processing or Defence*, Edinburgh, UK, 2020.

[38] S. Weiss, I.K. Proudler, and F.K. Coutts, "Eigenvalue decomposition of a parahermitian matrix: extraction of analytic eigenvalues," *IEEE Trans. Signal Processing*, vol. 69, pp. 722–737, 2021.

[39] S. Weiss, I.K. Proudler, F.K. Coutts, and F.A. Khattak, "Eigenvalue decomposition of a parahermitian matrix: Extraction of analytic eigenvectors," *IEEE Trans. Signal Processing*, vol. 71, pp. 1642–1656, 2023.

[40] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*, 2nd ed. Prentice-Hall Englewood Cliffs, 1999.

[41] F.K. Coutts, K. Thompson, J. Pestana, I.K. Proudler, and S. Weiss, "Enforcing eigenvector smoothness for a compact DFT-based polynomial eigenvalue decomposition," in *10th IEEE Workshop on Sensor Array and Multichannel Signal Processing*, pp. 1–5, July 2018.

[42] J. Foster, J.G. McWhirter, and J. Chambers, "Limiting the order of polynomial matrices within the SBR2 algorithm," in *IMA International Conference on Mathematics in Signal Processing*, Cirencester, UK, Dec/ 2006.

[43] C.H. Ta and S. Weiss, "Shortening the order of paraunitary matrices in SBR2 algorithm," in *6th Int. Conference on Information, Communications & Signal Processing*, Singapore, pp. 1–5, Dec. 2007.

[44] J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter, "Row-shift corrected truncation of paraunitary matrices for PEVD algorithms," in *23rd European Signal Processing Conference*, Nice, France, pp. 849–853, Aug. 2015,

[45] ——, "Shortening of paraunitary matrices obtained by polynomial eigenvalue decomposition algorithms," in *Int. Conference on Sensor Signal Processing for Defence*, Edinburgh, Scotland, Sep. 2015.