

A Comparison of Iterative and DFT-Based Polynomial Matrix Eigenvalue Decompositions

Fraser K. Coutts*, Keith Thompson*, Ian K. Proudler*,[†], Stephan Weiss*

* Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow, Scotland

[†] School of Electrical, Electronics & Systems Engineering, Loughborough Univ., Loughborough, UK

{fraser.coutts,keith.thompson,ian.proudler,stephan.weiss}@strath.ac.uk

Abstract—A variety of algorithms have been developed to compute an approximate polynomial matrix eigenvalue decomposition (PEVD). As an extension of the ordinary EVD to polynomial matrices, the PEVD will generate paraunitary matrices that diagonalise a parahermitian matrix. This paper compares the decomposition accuracies of two fundamentally different methods capable of computing an approximate PEVD. The first of these — sequential matrix diagonalisation (SMD) — iteratively decomposes a parahermitian matrix, while the second DFT-based algorithm computes a pointwise in frequency decomposition. We demonstrate through the use of examples that both algorithms can achieve varying levels of decomposition accuracy, and provide results that indicate the type of broadband multichannel problems that are better suited to each algorithm. It is shown that iterative methods, which generate paraunitary eigenvectors, are suited for general applications with a low number of sensors, while a DFT-based approach is useful for fixed, finite order decompositions with a small number of lags.

I. INTRODUCTION

Broadband multichannel problems can be expressed using polynomial matrix representations [1]. Such formulations can be used in a number of areas, including broadband angle of arrival estimation [2], [3], polyphase analysis and synthesis matrices for filter banks [4], and broadband beamforming [5], [6]. These problems typically involve parahermitian polynomial matrices, which are identical to their parahermitian conjugate, i.e., $\mathbf{R}(z) = \tilde{\mathbf{R}}(z) = \mathbf{R}^H(1/z^*)$ [4]. Such a matrix $\mathbf{R}(z)$ can arise as the z -transform of a space-time covariance matrix $\mathbf{R}[\tau]$.

A polynomial matrix eigenvalue decomposition (PEVD), which is an extension of the eigenvalue decomposition to parahermitian matrices, has been defined in [7]. The PEVD uses finite impulse response (FIR) paraunitary matrices [8] to approximately diagonalise a space-time covariance matrix. Given an input parahermitian matrix $\mathbf{R}(z) \in \mathbb{C}^{N \times N}$, and its associated coefficient matrix $\mathbf{R}[\tau]$, PEVD algorithms generate an output diagonal matrix $\mathbf{D}(z)$ containing eigenvalues, and a paraunitary matrix $\mathbf{F}(z)$ containing eigenvectors, such that

$$\mathbf{D}(z) \approx \mathbf{F}(z)\mathbf{R}(z)\tilde{\mathbf{F}}(z). \quad (1)$$

Equation (1) has only approximate equality, as the PEVD of a finite order polynomial matrix is generally not of finite order.

Existing PEVD algorithms include sequential matrix diagonalisation (SMD) [10], second-order sequential best rotation (SBR2) [7], and various evolutions of the algorithm families [11]–[13]. Each of these algorithms uses an iterative approach to approximately diagonalise a parahermitian matrix, and in many cases they encourage spectral majorisation [9], such that the resulting eigenvalues are ordered, as shown for the example of a 2×2 matrix in Fig. 1.

A DFT-based PEVD formulation, which transforms the problem into a pointwise in frequency standard matrix decom-

position, is provided in [14]. The method can either return a spectrally majorised decomposition, akin to the example in Fig. 1(b), or attempt to compute smooth, ideally analytic, eigenvalues as shown in the example of Fig. 1(a). The inherent drawback of a lack of phase-coherence between independent frequency bins [15] is solved via a quadratic nonlinear minimisation problem, which encourages phase alignment between adjacent bins.

As both iterative and DFT-based PEVD algorithms reach a solution to (1) through different methodologies, a comparison of their decomposition performance is of interest. Here, we compare the decompositions of SMD — as a representative of iterative PEVD algorithms — and the DFT-based approach in terms of algorithmic complexity, reconstruction error, eigenvector paraunitarity and order, parahermitian matrix diagonalisation, and spectral majorisation. Through this comparison, we obtain results that indicate the type of broadband multichannel problems that are better suited to each algorithm.

Below, Sec. II and Sec. III will provide a brief overview of the SMD and DFT-based algorithms from [10] and [14]. A comparison of the decomposition accuracies is presented in Sec. IV, with conclusions drawn in Sec. V.

II. SEQUENTIAL MATRIX DIAGONALISATION

A. Overview

The SMD algorithm approximates the PEVD using a series of elementary paraunitary operations to iteratively diagonalise a parahermitian matrix $\mathbf{R}(z) \in \mathbb{C}^{N \times N}$.

Upon initialisation, the algorithm diagonalises the lag-zero coefficient matrix $\mathbf{R}[0]$ by means of its modal matrix $\mathbf{Q}^{(0)}$; i.e., $\mathbf{S}^{(0)}(z) = \mathbf{Q}^{(0)}\mathbf{R}(z)\mathbf{Q}^{(0)H}$. The unitary $\mathbf{Q}^{(0)}$ — obtained from the EVD of the lag-zero slice $\mathbf{R}[0]$ — is applied to all coefficient matrices $\mathbf{R}[\tau] \forall \tau$, and initialises $\mathbf{H}^{(0)}(z) = \mathbf{Q}^{(0)}$.

In the i th step, $i = 1, 2, \dots, I$, the SMD algorithm computes

$$\begin{aligned} \mathbf{S}^{(i)}(z) &= \mathbf{U}^{(i)}(z)\mathbf{S}^{(i-1)}(z)\tilde{\mathbf{U}}^{(i)}(z) \\ \mathbf{H}^{(i)}(z) &= \mathbf{U}^{(i)}(z)\mathbf{H}^{(i-1)}(z) \quad , \end{aligned} \quad (2)$$

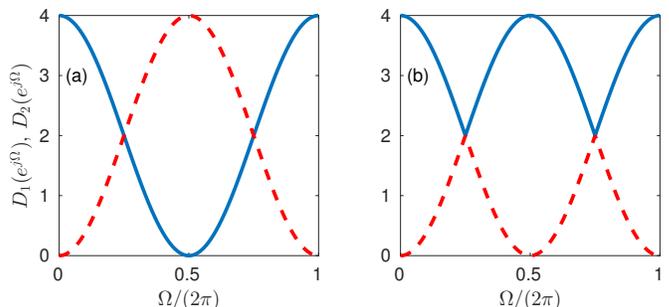


Fig. 1. Example for $\mathbf{D}(z) \in \mathbb{C}^{2 \times 2}$ with (a) analytic and (b) spectrally majorised eigenvalues.

in which

$$\mathbf{U}^{(i)}(z) = \mathbf{Q}^{(i)} \mathbf{\Lambda}^{(i)}(z). \quad (3)$$

The product in (3) consists of a paraunitary delay matrix

$$\mathbf{\Lambda}^{(i)}(z) = \text{diag}\{\underbrace{1 \dots 1}_{k^{(i)}-1} z^{-\tau^{(i)}} \underbrace{1 \dots 1}_{N-k^{(i)}}\}, \quad (4)$$

and a unitary matrix $\mathbf{Q}^{(i)}$, with the result that $\mathbf{U}^{(i)}(z)$ in (3) is paraunitary. For subsequent discussion, it is convenient to define intermediate variables $\mathbf{S}^{(i)'}(z)$ and $\mathbf{H}^{(i)'}(z)$ where

$$\begin{aligned} \mathbf{S}^{(i)'}(z) &= \mathbf{\Lambda}^{(i)}(z) \mathbf{S}^{(i-1)}(z) \tilde{\mathbf{\Lambda}}^{(i)}(z) \\ \mathbf{H}^{(i)'}(z) &= \mathbf{\Lambda}^{(i)}(z) \mathbf{H}^{(i-1)}(z), \end{aligned} \quad (5)$$

and

$$\mathbf{S}^{(i)}(z) = \mathbf{Q}^{(i)} \mathbf{S}^{(i)'}(z) \mathbf{Q}^{(i)H}. \quad (6)$$

Matrices $\mathbf{\Lambda}^{(i)}(z)$ and $\mathbf{Q}^{(i)}$ are selected based on the position of the dominant off-diagonal column in $\mathbf{S}^{(i-1)}(z) \bullet \circ \mathbf{S}^{(i-1)}[\tau]$, as identified by the parameter set

$$\{k^{(i)}, \tau^{(i)}\} = \arg \max_{k, \tau} \|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_2, \quad (7)$$

where

$$\|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_2 = \sqrt{\sum_{m=1, m \neq k}^N |\mathbf{s}_{m,k}^{(i-1)}[\tau]|^2} \quad (8)$$

and $\mathbf{s}_{m,k}^{(i-1)}[\tau]$ represents the element in the m th row and k th column of the coefficient matrix at lag τ , $\mathbf{S}^{(i-1)}[\tau]$.

The shifting process in (5) moves the dominant off-diagonal row and column into the zero lag coefficient matrix $\mathbf{S}^{(i)'}[0]$. The off-diagonal energy in the shifted row and column is then transferred onto the diagonal by the unitary matrix $\mathbf{Q}^{(i)}$ in (6), which diagonalises $\mathbf{S}^{(i)'}[0]$ by means of an ordered EVD.

Iterations continue for I steps until $\mathbf{S}^{(I)}(z)$ is sufficiently diagonalised with dominant off-diagonal column norm

$$\max_{k, \tau} \|\hat{\mathbf{s}}_k^{(I)}[\tau]\|_2 \leq \epsilon, \quad (9)$$

where the value of ϵ is chosen to be arbitrarily small. On completion, SMD generates an approximate PEVD given by

$$\mathbf{D}(z) = \mathbf{S}^{(I)}(z) = \mathbf{F}(z) \mathbf{R}(z) \tilde{\mathbf{F}}(z), \quad (10)$$

where $\mathbf{F}(z)$ is a concatenation of the paraunitary matrices:

$$\mathbf{F}(z) = \mathbf{H}^{(I)}(z) = \mathbf{U}^{(I)}(z) \dots \mathbf{U}^{(0)}(z) = \prod_{i=0}^I \mathbf{U}^{(I-i)}(z).$$

B. Algorithm Complexity

At the i th iteration of SMD, every matrix-valued coefficient in $\mathbf{S}^{(i)'}(z)$ must be left- and right-multiplied with unitary matrix $\mathbf{Q}^{(i)}$; similarly $\mathbf{H}^{(i)'}(z)$ is left-multiplied with $\mathbf{Q}^{(i)}$. A total of $2L^{(i)}$ and $L_H^{(i)}$ matrix multiplications are therefore required to update $\mathbf{S}^{(i)'}(z)$ and $\mathbf{H}^{(i)'}(z)$, which have lengths $L^{(i)}$ and $L_H^{(i)}$. Assuming that the complexity of the multiplication of two $N \times N$ matrices is of order $\mathcal{O}(N^3)$; the complexity of one SMD iteration is of order $\mathcal{O}(N^3(2L^{(i)} + L_H^{(i)})) \approx \mathcal{O}(N^3L)$, if it is assumed that $L^{(i)}$ and $L_H^{(i)}$ are proportional to the length, L , of the input parahermitian matrix.

The update step dominates the complexity of SMD [10]; thus, the algorithm complexity is of order $\mathcal{O}(N^3L)$.

III. DFT-BASED PEVD

A. Overview

The approach in [14] uses a decomposition of the form

$$\mathbf{R}[k] = \mathbf{F}^H[k] \mathbf{D}[k] \mathbf{F}[k], \quad k = 0, 1, \dots, K-1, \quad (11)$$

where $\mathbf{F}[k]$ contains eigenvectors, $\mathbf{D}[k]$ contains eigenvalues, and $\mathbf{R}[k]$ is obtained from the K -point DFT of $\mathbf{R}[\tau]$,

$\mathbf{R}[k] = \mathbf{R}(z)|_{z=w_K^k} = \sum_{\tau_{\min}}^{\tau_{\max}} \mathbf{R}[\tau] w_K^{k\tau}$, $k = 0, 1, \dots, K-1$, where $w_K = e^{-j2\pi/K}$. An approximate PEVD is therefore obtained via K EVDs that are pointwise in frequency.

The PEVD in (1) corresponds to linear convolution in the coefficients domain; however, the decomposition obtained in the DFT domain corresponds to the circular convolution

$$\mathbf{R}[(\tau)_K] = \mathbf{F}^H[(\tau)_K] \otimes \mathbf{D}[(\tau)_K] \otimes \mathbf{F}[(\tau)_K], \quad (12)$$

where \otimes is the circular convolution operator, and $(\tau)_K$ denotes τ modulo K . For (12) to be equivalent to (1), the number of frequency bins must satisfy

$$K \geq (2M + L - 2), \quad (13)$$

where $L = (\tau_{\max} - \tau_{\min} + 1)$ is the length of input parahermitian matrix $\mathbf{R}(z)$, and M is the assumed length of the paraunitary matrix that is non-zero. That is, $\mathbf{F}[\tau] = \mathbf{0}$ for $\tau \geq M$ and $\tau < 0$. Typically, choosing $K = (2M + L - 2)$ is valid, as decomposition accuracy does not increase significantly for larger K [14], but algorithmic computational complexity does.

At each frequency bin, eigenvalues are typically arranged in descending order; this results in approximate spectral majorisation of the polynomial eigenvalues. Sec. III-B discusses the rearrangement of eigenvalues to form a smooth decomposition.

Each eigenvector in a conventional EVD may be influenced by an arbitrary scalar phase angle and still be valid. This ambiguity in phase of each eigenvector can lead to discontinuities in phase between adjacent frequency bins. For a short paraunitary matrix $\mathbf{F}(z)$, these discontinuities must be smoothed. This is achieved through the use of a phase alignment function, described in Sec. III-C, which uses the Dogleg algorithm [16] to solve an unconstrained optimisation problem.

Following phase alignment, $\mathbf{F}[\tau]$ is computed as

$$\mathbf{F}[\tau] = \sum_{k=0}^{K-1} \mathbf{F}[k] w_K^{-k\tau}, \quad \tau = 0, 1, \dots, M-1, \quad (14)$$

and $\mathbf{D}(z)$ is the diagonal elements of $\mathbf{F}(z) \mathbf{R}(z) \tilde{\mathbf{F}}(z)$. Any energy in lags $\tau = M \dots K-1$ of $\mathbf{F}[\tau]$ is ignored.

B. Smooth Decomposition

If strong decorrelation is required for an application, but spectral majorisation is not, then a smooth decomposition may be preferable. In such a decomposition, the eigenvalues — and their eigenvectors — are arranged such that discontinuities between adjacent frequency bins are minimised. Discontinuities occur when the eigenvalues intersect at some frequencies.

For a smooth decomposition, the eigenvectors in adjacent frequency bins are rearranged using the inner product

$$c_{ij}[k] = \mathbf{f}_i[k-1] \mathbf{f}_j^H[k], \quad (15)$$

where, $\mathbf{f}_i[k]$ is the i th row of $\mathbf{F}[k]$. For each eigenvector $\mathbf{f}_i[k-1]$, $i = 1 \dots N$, a subsequent eigenvector $\mathbf{f}_{i'}[k]$ is chosen from an initial set $S = \{1 \dots N\}$ of the rows of $\mathbf{F}[k]$ such that

$$i' = \arg \max_{j \in S} \{|c_{ij}[k]|\}, \quad (16)$$

Once i' is identified, it is removed from the set: $S = S - \{i'\}$, and the next eigenvector is chosen.

The selected eigenvectors are combined in a rearranged matrix $\mathbf{F}'[k] = [\mathbf{f}_1^T[k] \dots \mathbf{f}_N^T[k]]^T$, and $\mathbf{F}[k]$ is set equal to $\mathbf{F}'[k]$. Thus, the eigenvector discontinuity between $\mathbf{F}[k-1]$ and $\mathbf{F}[k]$ has been reduced. This process is completed for $k = 1 \dots K-1$.

C. Phase Alignment

Phase alignment of eigenvectors in adjacent frequency bins is vital for a compact order decomposition. Thus, if a compact order is sought (e.g., here it is desired that only lags $\tau = 0 \dots M-1$ are non-zero), then phase alignment can be achieved by finding the phase changes required for each eigenvector $\mathbf{f}_i[k] \forall i, k$ to enforce this low order.

The phase of the i th eigenvector at frequency bin k can be adjusted by an angle $\theta_i[k]$ according to $\hat{\mathbf{f}}_i[k] = e^{j\theta_i[k]} \mathbf{f}_i[k]$. For the i th polynomial eigenvector $\mathbf{f}_i[\tau]$ to be compact, it is required to find angles $\vec{\theta}_i = [\theta_i[1] \dots \theta_i[K-1]]^T$, that satisfy

$$\mathbf{f}_i[\tau] = \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{f}_i[k] e^{j\theta_i[k]} w_K^{-k\tau} = \mathbf{0}, \quad (17)$$

for $\tau = M, \dots, K-1$. Without loss of generality, let $\theta_i[0] = 0$. These $(K-M)$ -folded equations can be expressed as

$$\mathcal{F}_M(\mathbf{f}_i) \mathbf{x}(\vec{\theta}_i) + f_M(\mathbf{f}_i) = \mathbf{0}, \quad (18)$$

where $\mathbf{x}(\vec{\theta}_i) = [e^{j\theta_i[1]}, e^{j\theta_i[2]}, \dots, e^{j\theta_i[K-1]}]^T$, $f_M(\mathbf{f}_i) = [\mathbf{f}_i[0], \mathbf{f}_i[0], \dots, \mathbf{f}_i[0]]^T$ is a $N(K-M) \times 1$ vector, and

$$\mathcal{F}_M(\mathbf{f}_i) = \begin{bmatrix} \mathbf{f}_i^T[1] w_K^{-M} & \mathbf{f}_i^T[2] w_K^{-2M} & \dots & \mathbf{f}_i^T[K-1] w_K^{-(K-1)M} \\ \mathbf{f}_i^T[1] w_K^{-(M+1)} & \mathbf{f}_i^T[2] w_K^{-2(M+1)} & \dots & \mathbf{f}_i^T[K-1] w_K^{-(K-1)(M+1)} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{f}_i^T[1] w_K^{-(K-1)} & \mathbf{f}_i^T[2] w_K^{-2(K-1)} & \dots & \mathbf{f}_i^T[K-1] w_K^{-(K-1)^2} \end{bmatrix}$$

is a $N(K-M) \times (K-1)$ matrix.

In general, there may exist no phase vector $\vec{\theta}_i$ which satisfies (18). However, by minimising the energy in the coefficients for $\tau = M, \dots, K-1$, some $\vec{\theta}_i$ can be obtained. The energy in these coefficients is therefore used as the objective of the unconstrained minimisation problem

$$J(\vec{\theta}_i) = \|\mathcal{F}_M(\mathbf{f}_i) \mathbf{x}(\vec{\theta}_i) + f_M(\mathbf{f}_i)\|^2, \quad i = 1 \dots N. \quad (19)$$

Thus, $\vec{\theta}_i$ is obtained by solving $\vec{\theta}_i = \arg \min J(\vec{\theta}_i)$. In [14], it was found that the Dogleg method [16] was able to satisfactorily minimise $J(\vec{\theta}_i)$ for $i = 1 \dots N$.

D. Algorithm Complexity

The complexity of each of the N instantiations of the Dogleg method is $\mathcal{O}(K^3)$ due to matrix inversion [14]; thus, the total complexity of the phase alignment step is of order $\mathcal{O}(NK^3)$. Given that K is bounded from below by L in (13) for constant M , $\mathcal{O}(NK^3)$ can be expressed as $\mathcal{O}(NL^3)$. The computation of the frequency domain representation of $\mathbf{R}[\tau] \forall \tau$, the execution of K EVDs, and the smoothing of eigenvalues are of lower complexity than this step; thus, the total complexity of the algorithm is approximately $\mathcal{O}(NL^3)$.

IV. ALGORITHM COMPARISON

A. Metrics

Denote the mean-squared reconstruction error for an approximate PEVD as

$$\text{MSE} = \frac{1}{N^2 L'} \sum_{\tau} \|\mathbf{E}_R[\tau]\|_{\mathbb{F}}^2, \quad (20)$$

where $\mathbf{E}_R[\tau] = \hat{\mathbf{R}}[\tau] - \mathbf{R}[\tau] \forall \tau$, $\hat{\mathbf{R}}(z) = \tilde{\mathbf{F}}(z) \mathbf{D}(z) \mathbf{F}(z)$, L' is the length of $\mathbf{E}_R(z)$, and $\|\cdot\|_{\mathbb{F}}$ denotes the Frobenius norm. Furthermore, define the paraunitarity (PU) error as

$$\eta = \frac{1}{N} \sum_{\tau} \|\mathbf{E}_F[\tau] - \mathbf{I}_N\|_{\mathbb{F}}^2, \quad (21)$$

where $\mathbf{E}_F(z) = \mathbf{F}(z) \tilde{\mathbf{F}}(z)$, and \mathbf{I}_N is an $N \times N$ identity matrix. Finally, define diagonalisation as the ratio of off-diagonal energy in $\mathbf{D}(z)$ to the total energy:

$$E_{\text{diag}} = \frac{\sum_{\tau} \|\bar{\mathbf{D}}[\tau]\|_{\mathbb{F}}^2}{\sum_{\tau} \|\mathbf{D}[\tau]\|_{\mathbb{F}}^2}, \quad (22)$$

where $\bar{\mathbf{D}}[\tau]$ is equal to $\mathbf{D}[\tau]$ but with its diagonal elements set to zero. The output paraunitary matrix $\mathbf{F}(z)$ can be used in signal processing applications. A useful metric for gauging the implementation cost of this matrix is its length, L_F .

B. Approximation of Eigenvalues

The SMD algorithm iteratively diagonalises a parahermitian matrix $\mathbf{R}(z)$; thus, the approximation of the polynomial eigenvalues becomes better with each algorithm iteration. Almost exact diagonalisation of $\mathbf{R}(z)$ typically requires a large number of iterations; this can be problematic, as the parahermitian matrix grows in order at each iteration of SMD [10]. Truncation of the outer lags of the matrix containing lower energy can help mitigate this growth. The polynomial eigenvalues produced by SMD are approximately spectrally majorised [10], and cannot be reordered.

Matrix $\mathbf{D}(z)$ is set to be exactly diagonal in the final step of the DFT-based approach, therefore $E_{\text{diag}} = 0$ for all instances of the algorithm. However, directly setting off-diagonal elements equal to zero in this way negatively impacts the decomposition MSE. The eigenvalue approximation ultimately depends upon the accuracy of the eigenvectors, which increases for increasing M . Note, M is not known a priori.

The DFT-based algorithm naturally produces an approximately spectrally majorised decomposition, but as described in Sec. III-B, the eigenvalues can be reordered to achieve a smooth decomposition. The latter avoids discontinuities at the intersection of eigenvalues in the frequency domain, and typically leads to a more compact (lower order) decomposition.

C. Paraunitarity of Polynomial Eigenvectors

The eigenvectors $\mathbf{F}(z)$ generated by SMD are strictly paraunitary, as they are created as the product of a series of elementary paraunitary matrices. While this is advantageous for some applications, some loss of paraunitarity may be acceptable if other performance gains are made. For example, truncation of the paraunitary matrices within the SMD update step introduces a trade-off between η and E_{diag} for a given paraunitary filter length; i.e., a larger truncation value, μ , sacrifices paraunitarity to reduce the paraunitary filter order required to achieve a certain diagonalisation. Truncation of the paraunitary matrices in SMD is based on a threshold μ , whereby the maximum and minimum lags of $\mathbf{H}^{(i)}(z)$ are reduced from τ_{max} and τ_{min} to $\tilde{\tau}_{\text{max}}$ and $\tilde{\tau}_{\text{min}}$ such that

$$\sum_{\tau=\tilde{\tau}_{\text{max}}+1}^{\tau_{\text{max}}} \|\mathbf{H}^{(i)}[\tau]\|_{\mathbb{F}}^2 < \frac{\mu \sum_{\tau} \|\mathbf{H}^{(i)}[\tau]\|_{\mathbb{F}}^2}{2} > \sum_{\tau=\tau_{\text{min}}}^{\tilde{\tau}_{\text{min}}-1} \|\mathbf{H}^{(i)}[\tau]\|_{\mathbb{F}}^2.$$

The eigenvectors generated by the DFT-based PEVD are only approximately paraunitary [14]. For increasing M , the approximation improves; thus, to achieve a desired level of paraunitarity in an application, an adequate value of M must be determined through experimentation. The required value of M is likely to be lower if a smooth decomposition is used, as discontinuities at eigenvalue intersections are avoided. To represent such discontinuities requires infinitely long polynomials, which do not fit well into the fixed order model of the DFT-based algorithm, as energy from ignored high order

TABLE I. MSE, PU ERROR, DIAGONALISATION, AND FILTER LENGTH COMPARISON FOR FINITE ORDER EXAMPLE.

Method	MSE	η	E_{diag}	L_F
DFT smooth	7.1×10^{-9}	4.9×10^{-9}	0	3
DFT maj.	8.8×10^{-7}	2.4×10^{-3}	0	165
SMD, μ_1	2.6×10^{-25}	1.2×10^{-16}	1×10^{-6}	689
SMD, μ_2	1.7×10^{-10}	4.8×10^{-8}	1×10^{-6}	165

polynomial coefficients may corrupt the extracted coefficients from lags $0 \dots M - 1$.

D. Finite Order Example

Consider the parahermitian matrix

$$\mathbf{R}(z) = \begin{bmatrix} .5z^2 + 3 + .5z^{-2} & -.5z^2 + .5z^{-2} \\ .5z^2 - .5z^{-2} & -.5z^2 + 1 - .5z^{-2} \end{bmatrix}, \quad (23)$$

which has an exact finite order smooth decomposition with

$$\mathbf{F}(z) = \frac{1}{2} \begin{bmatrix} z+1 & -z+1 \\ -z+1 & z+1 \end{bmatrix} \quad \mathbf{D}(z) = \begin{bmatrix} z+2+z^{-1} & 0 \\ 0 & -z+2-z^{-1} \end{bmatrix}$$

where $\mathbf{F}(z)$ contains eigenvectors on its rows, and $\mathbf{D}(z)$ contains analytic eigenvalues on its diagonal, which on the unit circle match the power spectral densities given in Fig. 1(a).

The metrics of Sec. IV-A were calculated for the decomposition of (23) by the DFT-based and SMD algorithms, and can be seen in Tab. I. For the former, both majorised and smooth decompositions were generated to approximate the solutions given in Fig. 1(b) and (a), respectively. SMD truncation parameters of $\mu_1 = 10^{-16}$ and $\mu_2 = 10^{-8}$ were used to demonstrate the trade-off between paraunitarity and diagonalisation for the algorithm. Both instances of SMD were run until $E_{\text{diag}} = 10^{-6}$. For the majorised DFT-based decomposition, M was set equal to 165 ($K = 333$) to allow comparison with SMD when utilising μ_2 .

Given its ability to produce a smooth decomposition, the DFT-based approach is able to almost perfectly approximate the finite order $\mathbf{F}(z)$ and $\mathbf{D}(z)$ for $M = 3$, $K = 9$. In contrast, the SMD algorithm is able to produce a spectrally majorised approximately diagonal $\mathbf{D}(z)$, but the eigenvector matrices are of significantly higher order for both μ_1 and μ_2 . The spectrally majorised DFT-based decomposition has significantly higher MSE and η for the same L_F as SMD with μ_2 . By utilising a higher truncation within SMD, it can be seen that MSE and η have increased, but L_F has decreased.

E. Non-Finite Order Example

As a second example, consider the parahermitian matrix

$$\mathbf{R}(z) = \begin{bmatrix} 2 & z^{-1} + 1 \\ z + 1 & 2 \end{bmatrix}. \quad (24)$$

The eigenvectors

$$\tilde{\mathbf{F}}(z) = \frac{1}{\sqrt{2}} \begin{bmatrix} (z^{-1} + 1)(z^{-1} + 2 + z)^{-1/2} & z^{-1} \\ 1 & -(z^{-1} + 1)(z^{-1} + 2 + z)^{-1/2} \end{bmatrix}$$

and eigenvalues

$$\mathbf{D}(z) = \begin{bmatrix} 2 + (z^{-1} + 2 + z)^{1/2} & 0 \\ 0 & 2 - (z^{-1} + 2 + z)^{1/2} \end{bmatrix}$$

of $\mathbf{R}(z)$ are neither of finite order nor rational; thus, to decompose $\mathbf{R}(z)$ via a PEVD would require polynomial matrices of infinite length for both smooth and majorised decompositions.

The metrics of Sec. IV-A were calculated for the decomposition of (24) by the DFT-based and SMD algorithms, and can be seen in Tab. II. Again, SMD truncation parameters of

TABLE II. MSE, PU ERROR, DIAGONALISATION, AND FILTER LENGTH COMPARISON FOR NON-FINITE ORDER EXAMPLE.

Method	MSE	η	E_{diag}	L_F
DFT smooth	2.1×10^{-5}	2.3×10^{-3}	0	83
DFT maj.	1.4×10^{-6}	2.2×10^{-3}	0	83
SMD, μ_1	4.4×10^{-25}	2.5×10^{-16}	1×10^{-6}	345
SMD, μ_2	2.9×10^{-10}	9.5×10^{-8}	1×10^{-6}	83

$\mu_1 = 10^{-16}$ and $\mu_2 = 10^{-8}$ were used, and both instances of SMD were run until $E_{\text{diag}} = 10^{-6}$. For both DFT-based decompositions, M was set equal to 83 ($K = 167$) to allow comparison with SMD when utilising μ_2 .

The values of MSE and η for both DFT-based PEVDs are significantly higher for this example, while the eigenvectors generated by SMD are of lower order than in Tab. I. This indicates that the DFT-based approach may suffer for similarly complex problems, while SMD is relatively unaffected. For this example, there is actually a slight disadvantage to using a smooth decomposition. Using a higher truncation within SMD has again increased MSE and η , but L_F has decreased.

V. CONCLUSION

In this paper, we have compared the decomposition accuracies of two recent PEVD algorithms. The iterative SMD algorithm has been shown to exhibit significantly lower mean-squared reconstruction error and paraunitarity error than a DFT-based approach; however, SMD does not achieve exact diagonalisation, and its enforcement of spectral majorisation can lead to high polynomial eigenvector orders unless truncation is employed. The ability of the DFT-based method to produce a smooth decomposition can produce extremely compact eigenvectors, but the algorithm's reliance on a fixed eigenvector order can introduce significant paraunitarity error for decompositions where the ground truth is of infinite order.

From an analysis of both algorithms' complexities, it can be determined that SMD becomes significantly more complex for increasing spatial dimension N , while the DFT-based approach becomes significantly more complex for increasing parahermitian matrix length L . Typically, $L > N$ for a parahermitian matrix input to a PEVD algorithm; thus, SMD is likely to offer a lower complexity solution.

When designing PEVD implementations for real applications, both of the algorithms described in this paper could be extremely useful. As a relatively stable algorithm, with typically low reconstruction error, paraunitary eigenvectors, and customisable diagonalisation and eigenvector length, SMD can be deployed in any scenario with reasonably low N . For problems of fixed, finite order, or situations in which a smooth decomposition is preferable or paraunitarity is not required, the DFT-based approach can be used to great effect, provided that L is not too large. However, M not being known a priori is a disadvantage of this method for applications purposes.

ACKNOWLEDGEMENT

Fraser Coutts is the recipient of a Caledonian Scholarship; we would like to thank the Carnegie Trust for their support.

This work was supported in parts by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/K014307/1 and the MOD University Defence Research Collaboration in Signal Processing.

REFERENCES

- [1] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, New York, 1982.
- [2] M. Alrmah, S. Weiss, and S. Lambbotharan. An extension of the music algorithm to broadband scenarios using polynomial eigenvalue decomposition. In *19th European Signal Processing Conference*, pp. 629–633, Barcelona, Spain, Aug. 2011.
- [3] S. Weiss, M. Alrmah, S. Lambbotharan, J. McWhirter, and M. Kaveh. Broadband angle of arrival estimation methods in a polynomial matrix decomposition framework. In *IEEE 5th Int. Workshop Comp. Advances in Multi-Sensor Adaptive Proc.*, St. Martin, pp. 109–112, Dec. 2013.
- [4] P. P. Vaidyanathan. *Multirate Systems and Filter Banks*. Prentice Hall, Englewood Cliffs, 1993.
- [5] S. Weiss, S. Bendoukha, A. Alzin, F. Coutts, I. Proudler, and J. Chambers. MVDR broadband beamforming using polynomial matrix techniques. In *EUSIPCO*, pp. 839–843, Nice, France, Sep. 2015.
- [6] A. Alzin, F. Coutts, J. Corr, S. Weiss, I. K. Proudler, and J. A. Chambers. Adaptive broadband beamforming with arbitrary array geometry. In *IET/EURASIP ISP*, London, UK, Dec. 2015.
- [7] J. G. McWhirter, P. D. Baxter, T. Cooper, S. Redif, and J. Foster. An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE TSP*, 55(5):2158–2169, May 2007.
- [8] S. Icart, P. Comon. Some properties of Laurent polynomial matrices. In *IMA Int. Conf. Math. Signal Proc.*, Birmingham, UK, Dec. 2012.
- [9] P. Vaidyanathan. Theory of optimal orthonormal subband coders. *IEEE TSP*, 46(6):1528–1543, June 1998.
- [10] S. Redif, S. Weiss, and J. McWhirter. Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices. *IEEE TSP*, 63(1):81–89, Jan. 2015.
- [11] 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*, pp. 312–315, Gold Coast, Australia, June 2014.
- [12] Z. Wang, J. G. McWhirter, J. Corr, and S. Weiss. Multiple shift second order sequential best rotation algorithm for polynomial matrix EVD. In *EUSIPCO*, pp. 844–848, Nice, France, Sep. 2015.
- [13] J. Corr, K. Thompson, S. Weiss, J. G. McWhirter, and I. K. Proudler. Causality-Constrained multiple shift sequential matrix diagonalisation for parahermitian matrices. In *EUSIPCO*, pp. 1277–1281, Lisbon, Portugal, Sep. 2014.
- [14] M. Tohidian, H. Amindavar, and A. M. Reza. A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices. *EURASIP J. Adv. Signal Process.*, 2013:93, December 2013.
- [15] R. Klemm. *Space-time Adaptive Processing Principles and Applications*. IEE Radar, Sonar, Navigation Avionics, London, UK, 1998.
- [16] M. J. D. Powell. A new algorithm for unconstrained optimization. *Nonlinear programming*, pp. 31–65, 1970.