



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/244658/>

Version: Accepted Version

Proceedings Paper:

Khattak, F.A., Bakhit, M., Produler, I.K. et al. (Accepted: 2026) Fast Analytic Eigenvector Extraction and Partial Analytic Eigenvalue Decomposition. In: 2026 IEEE High Performance Extreme Computing Conference (HPEC). 30th Annual IEEE High Performance Extreme Computing Virtual Conference, 14-18 Sep 2026, Virtual. IEEE. (In Press)

This is an author produced version of a conference paper accepted for publication in 2026 IEEE High Performance Extreme Computing Conference (HPEC), made available via the University of Leeds Research Outputs Policy under the terms of the Creative Commons Attribution License (CC-BY), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Fast Analytic Eigenvector Extraction and Partial Analytic Eigenvalue Decomposition

Faizan A. Khattak¹, Mohammed Bakhit², Ian K. Proudler², and Stephan Weiss²

¹School of Computer Science, University of Leeds, Leeds LS2 9JT, UK

²Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, Scotland
f.a.khattak@leeds.ac.uk, {mohammed.bakhit,ian.proudler,stephan.weiss}@strath.ac.uk

Abstract—The analytic eigenvalue decomposition (EVD) is a useful tool in many broadband sensor array applications where a parahermitian matrix, such as a space-time covariance matrix, needs to be diagonalised. The eigenvalues are power spectral density terms; if they are spectrally majorised, i.e. strictly ordered at every frequency, existing time domain algorithms can approximate the analytic EVD, but are known to be of limited precision. Recent DFT domain approaches are accurate but can be computationally costly and require the parahermitian matrix to have full spatial rank. In contrast, this paper suggests a partial analytic EVD that extracts accurate analytic eigenpairs and is applicable to low-rank parahermitian matrices. Our approach is based on a phase normalisation of eigenvectors in DFT bins in order to establish phase-coherence across the spectrum, which is simple and is shown to yield an analytic solution. The proposed approach is underpinned by examples and demonstrated in simulations.

I. INTRODUCTION

Broadband multichannel problems are often formulated using a space-time covariance matrix, which contains auto- and cross-correlation sequences as entries. To find solutions to such problems then involves a diagonalisation that holds for all time lags of these sequences, or, in the z -domain, for all values of z [1]. Applications for this include, for example, signal compaction [2], speech enhancement [3], and direction of arrival estimation and source localisation [4], [5]. A key feature of these application is that a limited number of sources result in a low number of dominant analytic eigenvalues, which then motivates subspace decomposition approaches, or low-rank approximations and denoising to suppress unstructured noise.

The diagonalisation of a parahermitian matrix is accomplished by the analytic EVD [6]–[8]. In the case of spectrally majorised eigenvalues, it is approximated by time-domain algorithms such as the second order sequential best rotation (SBR2) [9]–[12] and the sequential matrix diagonalisation (SMD) algorithms [2], [13]. These algorithms iteratively eliminate off-diagonal energy by a sequence of elementary paraunitary operations, but often result in insufficient diagonalisation and high orders of the computed factors of the decomposition. Recent discrete Fourier transform (DFT)-based methods [14]–[17] can accurately retrieve the ground truth analytic EVD factors, but can only do so at high computational cost and for parahermitian matrices that have full spatial rank. The computational cost is governed by the NP-hard problem of phase-smoothing eigenvectors, since DFT-domain approaches

need to re-establish spectral coherence across frequency bins. The limitation to full-rank matrices can be problematic for the above application examples.

In this paper, we address the problem of finding accurate analytic EVD factors for parahermitian matrices of potentially low spatial rank. Firstly we bypass costly phase-smoothing to find a minimum support for eigenvectors by a bin-wise rotation that forces one of the eigenvector components to a target phase across all bins. We show that this has an analytic solution, whereby the target phase in some cases can be zero, i.e. the particular component becomes real-valued, and in other cases may require a phase shift equivalent to a half sample delay. This solution may not have a support that is as compact as when using optimal phase smoothing [16], but the computation is simple and can be readily applied to reduced-rank parahermitian matrices, as long as their eigenvalues are spectrally majorised.

In the following, we review the analytic EVD and its challenges in Sec. II. Based on showing in Sec. III that an analytic solution exists for a phase normalisation of eigenvectors, a simple practical method is outlined in Sec. IV. The extraction of a limited number of eigenpairs — a partial analytic EVD — is discussed in Sec. V, followed by simulations and conclusions in Secs. VI and VII. Files to reproduce all figures and simulations can be found on GitHub¹.

II. ANALYTIC EVD AND COMPUTATIONAL CHALLENGE

A. Analytic Eigenvalue Decomposition

Let us assume that a broadband array signal $\mathbf{y}[n] = [y_1[n], \dots, y_M[n]]^T$ is acquired by M sensors, where $\{\cdot\}^T$ is the transpose operation. Then the space-time covariance matrix of $\mathbf{y}[n]$ can be expressed as [18]

$$\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{y}[n]\mathbf{y}^H[n - \tau]\} , \quad (1)$$

where $\mathcal{E}\{\cdot\}$ denotes expectation, $\{\cdot\}^H$ represents Hermitian transposition, and $\tau \in \mathbb{Z}$ is a lag parameter. The z -transform of $\mathbf{R}(z) = \sum_{\tau} \mathbf{R}[\tau]z^{-\tau}$, or for short $\mathbf{R}(z) \bullet \mathbf{R}[\tau]$, is the cross-spectral density matrix of $\mathbf{y}[n]$. This matrix is a Laurent polynomial matrix [18], satisfying the parahermitian property such that for the parahermitian transposition $\mathbf{R}^P(z) = \{\mathbf{R}(1/z^*)\}^H = \mathbf{R}(z)$ holds.

¹<https://github.com/faiziktk/FastPartialPEVD>

If the data $\mathbf{y}[n]$ is not connected to a multiplexing operation, an analytic EVD

$$\mathbf{R}(z) = \mathbf{Q}(z)\mathbf{\Lambda}(z)\mathbf{Q}^P(z) \quad (2)$$

exist [6]–[8]. In (2), the diagonal parahermitian matrix $\mathbf{\Lambda}(z) = \text{diag}\{\lambda_1(z), \dots, \lambda_M\}$ contains eigenvalues, which, due to $\lambda_m(z) = \lambda_m^P(z)$, satisfy $\lambda(z) \in \mathbb{R} \forall |z| = 1$. The corresponding eigenvectors $\mathbf{q}_m(z)$ can be selected to be orthogonal, and form the columns of $\mathbf{Q}(z)$, which is paraunitary such that $\mathbf{Q}(z)\mathbf{Q}^P(z) = \mathbf{Q}^P(z)\mathbf{Q}(z) = \mathbf{I}$. All factors are analytic, but can be infinite Laurent series.

In some cases, particularly if $\mathbf{R}(z)$ is randomly perturbed [19], the eigenvalues are spectrally majorised [20] s.t. $\lambda_m(e^{j\Omega}) \geq \lambda_{m+1}(e^{j\Omega})$, $m = 1, \dots, (M-1)$, and $\mathbf{\Lambda}(z)$ is unique. As long as the eigenvalues are distinct, the eigenvectors are unique up to an arbitrary allpass $\psi_m(z)$: if $\mathbf{q}_m(z)$ is a valid m -th eigenvector, then so is $\psi_m(z)\mathbf{q}_m(z)$.

Early algorithms to approximate (2) for spectrally majorised eigenvalues include the second order sequential best rotation (SBR2) [10], [11] and sequential matrix diagonalisation (SMD) [2], [12], [13] algorithms, which have found applications in a number of fields, with particular success in the area of speech enhancement [1], [3], [4]. However, SBR2 and SMD algorithms possess only limited accuracy. Recent discrete Fourier transform (DFT)-based methods [14]–[17] are more accurate, but can also incur a significant computational cost. We will explore next why this is the case.

B. Bin-Wise EVD Evaluation

DFT-based algorithms evaluate (2) on a bin-wise basis. If a K -point DFT is applied to $\mathbf{R}[\tau]$, with normalised angular bin frequencies $\Omega_k = \frac{2\pi k}{K}$, $k = 0, \dots, (K-1)$, we obtain Hermitian matrices $\mathbf{R}_k = \mathbf{R}(e^{j\Omega_k})$. These each admit an EVD

$$\mathbf{R}_k = \mathbf{Q}_k \mathbf{\Lambda}_k \mathbf{Q}_k^H. \quad (3)$$

The difficult task is now to associate the factors $\mathbf{Q}_k$ and $\mathbf{\Lambda}_k$ across bins.

Assuming that the eigenvalues in $\mathbf{\Lambda}_k$ are distinct, the factors in (3) relate to sample points of the analytic matrices $\mathbf{Q}(e^{j\Omega})$ and $\mathbf{\Lambda}(e^{j\Omega})$ on the r.h.s. of (2) via [16], [17]

$$\mathbf{\Lambda}_k = \mathbf{P}_k \mathbf{\Lambda}(e^{j\Omega_k}) \mathbf{P}_k^T, \quad (4)$$

$$\mathbf{Q}_k = \mathbf{Q}(e^{j\Omega_k}) \mathbf{\Psi}_k^H \mathbf{P}_k^T. \quad (5)$$

In (4) and (5), $\mathbf{P}_k$ is a permutation matrix responsible for ordering the eigenvalues, and $\mathbf{\Psi}_k$ is a diagonal matrix containing phase shifts that describe the phase ambiguity of the eigenvectors.

Assuming that the eigenvalues are spectrally majorised simplifies the association task since if, in (4), the eigenvalues are also sorted in descending order, then $\mathbf{P}_k = \mathbf{I}$. The remaining task of finding $\mathbf{\Psi}_k$ however requires smoothing of the phases of the eigenvectors across frequency bins. This in principle is an NP-hard quadratically constrained quadratic programming problem, which represents the computational bottleneck in current DFT-based algorithms [14], [16], [17].




Fig. 1. Frequency responses of the elements of analytic eigenvectors obtained by Dirichlet interpolation through $K = 16$ sample points (a) without phase smoothing and (b) with phase smoothing according to [16]. Lines indicate elements $\mu = 1, 2, 3$ in blue, red, and green, respectively, the markers show the values of at the sample points. $\triangle$

Example 1: For an eigenvector $\mathbf{q}(z) : \mathbb{C} \rightarrow \mathbb{C}^3$ of order 4, Fig. 1(a) shows a Dirichlet interpolation through $K = 16$ sample points obtained from a bin-wise EVD. Because the phases of the eigenvectors are not aligned, the recovered frequency responses are not as smooth as the ones in Fig. 1(b), which have been phase-smoothed across bins using the method in [16]. $\triangle$

III. PHASE NORMALISATION

A. Approach

In the remainder of this paper, we target the problem of identifying phase shifts in $\mathbf{\Psi}_k$ in (5) in order to create phase coherence of individual eigenvectors in a computationally less demanding fashion than the phase smoothing in [16]. We representatively address one eigenvector, $\mathbf{q}(z)$. We first show in this section that $\mathbf{q}(z)$ remains analytic if one of its element $q_\mu(z)$, $\mu = 1, \dots, M$, is ‘phase normalised’ by setting its phase to zero or — if necessary — to a linear function corresponding to a half sample delay, with the phases of the remaining elements of $\mathbf{q}(z)$ being altered by the same amount. Sec. IV will then discuss how this phase-normalised, analytic solution can be found from the DFT sample points $\mathbf{q}_k$.

Let ν be the index of the component $q_\nu(z)$ whose phase we want to normalise, and let $\varphi'(z)$ be its phase function. This phase function may not necessarily be analytic in $z \in \mathbb{C}$, indicated by the primed notation, but we will show that an analytic phase function $\varphi(z) = \varphi'(z)z^{-\frac{1}{2}\kappa}$ with either $\kappa = 0$ or $\kappa = 1$ can always be found, i.e. analyticity can be guaranteed by potentially including a half sample delay. Thus, we then create a new eigenvector

$$\bar{\mathbf{q}}(z) = \varphi^P(z)\mathbf{q}(z), \quad (6)$$

in the expectation that the phase normalisation may provide a sufficient level of smoothness once we consider a discretised version across DFT bins. Note that because $\bar{\mathbf{q}}(z)$ is analytic, $\bar{\mathbf{q}}[n] \circ \bar{\mathbf{q}}(z)$ decays at least exponentially [21], and can be arbitrarily closely approximated through delay and truncation by a polynomial [15].

For $\kappa = 0$, the element $\bar{q}_\nu(z)$ is normalised such that it becomes real-valued on the unit circle. We explore this effect in the following example:




Fig. 2. Frequency responses of the elements of the eigenvector $\mathbf{q}(z)$, phase-normalised to its (a) first and (b) second element.

Example 2: For the eigenvector $\mathbf{q}(z)$ in Example 1, we perform phase normalisation w.r.t. its elements $\nu = 1$ and $\nu = 2$ in Fig. 2(a) and (b) respectively; note the ν -th component of the vectors $\bar{\mathbf{q}}(e^{j\Omega})$ are real-valued. While the frequency responses are not as smooth as those in Fig. 1(b), they are analytic and considerably better behaved than those in the intervention-free case of Fig. 1(a). $\triangle$

B. Normalisation

To find the phase function of the ν th element $q_\nu(z)$, which then determines the normalisation, we would like to solve the factorisation problem

$$q_\mu(z) = \bar{q}'_\mu(z)\varphi'(z) \quad (7)$$

$$\text{s.t. } |\varphi'(e^{j\Omega})| = 1, \quad \bar{q}'_\mu(e^{j\Omega}) \in \mathbb{R} \quad \forall \Omega \in [0, 2\pi). \quad (8)$$

This would ensure that $\bar{q}'_\mu(z)$ is real-valued on the unit circle. If $q_\mu(z)$ is interpreted as a 1×1 polynomial matrix, then (7) together with its constraints in (8) represents the analytic singular value decomposition $q_\mu(z) = u'(z)\sigma'(z)v'^P(z)$ [8], [22] with $\bar{q}'_\mu(z) = \sigma'(z)$ and $\varphi'(z) = u'(z)v'^P(z)$. Finally, note that on the unit circle, $\varphi'(z)$ can also be determined via

$$\varphi'(e^{j\Omega}) = e^{j\angle\{q_\mu(e^{j\Omega})\}} \cdot \text{sgn}\{q_\nu(e^{j\Omega})\}, \quad (9)$$

where $\text{sgn}\{x\} = 1$ for $x \geq 0$ and $\text{sgn}\{x\} = 0$ otherwise.

Two cases, indexed by $\kappa \in \{0, 1\}$ [22], need to be considered w.r.t. (7). Case $\kappa = 0$: the singular value $\sigma'(z) = \bar{q}'_\mu(z)$ possesses an even number of zeros on the unit circle. Here the analytic SVD exists with analytic factors. Case $\kappa = 1$: for an odd number of zero crossings of $\bar{q}'_\mu(z)$, the analytic SVD does not exist with factors that are analytic in $z \in \mathbb{C}$ (even though they are analytic in $\Omega \in \mathbb{R}$), indicated by the primed variables in (7). This can be remedied by exchanging a half sample delay between $\bar{q}'_\mu(z)$ and $\varphi'(z)$ [23]. Thus we can obtain an analytic allpass

$$\varphi(z) = z^{-\frac{1}{2}\kappa}\varphi'(z) \quad (10)$$

in all cases, and define as analytic phase normalised eigenvector

$$\bar{\mathbf{q}}(z) = \varphi^P(z)\mathbf{q}(z). \quad (11)$$

Note however that while for $\kappa = 0$, the phase normalisation leads to zero phase for the μ -th component of $\bar{\mathbf{q}}(e^{j\Omega})$, for

$\kappa = 1$, the phase response of $\bar{\mathbf{q}}(e^{j\Omega})$ is non-zero and resembles $e^{-j\Omega/2}$, i.e. a half-sample delay.

Example 3: For a demonstration of phase normalisation, consider $\mathbf{q}(z) = [1+z^{-2}; 1+z^{-1}]^T$, which for simplicity does not possess unit norm and hence is not an eigenvector. We want to phase-normalise w.r.t. the first element $q_1(z) = 1+z^{-2}$. For its analytic SVD, we have

$$q_1(z) = \underbrace{1}_{u(z)} \cdot \underbrace{(z+z^{-1})}_{\sigma(z)=\bar{q}_1(z)} \cdot \underbrace{z^{-1}}_{v^P(z)=\varphi(z)}. \quad (12)$$

Note that $\sigma(z)$ has two zero crossings on the unit circle for $z = e^{j\frac{\pi}{2}}$ and $z = e^{j\frac{3\pi}{2}}$, and hence the analytic SVD exists with analytic factors. With $\varphi(z) = z^{-1}$, we have $\bar{\mathbf{q}}(z) = \mathbf{q}(z) = [z+z^{-1}; z+1]^T$ is the phase-normalised result, with the first element $\bar{q}_1(e^{j\Omega}) = \cos \Omega \in \mathbb{R} \quad \forall \Omega$ as desired. $\triangle$

Example 4: We revisit $\mathbf{q}(z)$ from Example 3, but now we phase normalise w.r.t. $q_2(z) = 1+z^{-1}$. For its analytic SVD, we have

$$q_2(z) = \underbrace{1}_{u(z)} \cdot \underbrace{(z^{\frac{1}{2}} + z^{-\frac{1}{2}})}_{\sigma'(z)=\bar{q}'_1(z)} \cdot \underbrace{z^{-\frac{1}{2}}}_{v^P(z)=\varphi'(z)}, \quad (13)$$

where $\sigma'(z)$ has a single zero crossing at $z = e^{j\pi}$. Therefore, we have $\varphi(z) = z^{-\frac{1}{2}}u'(z)v'^P(z) = z^{-1}$, and $\bar{\mathbf{q}}(z) = \varphi^P(z)\mathbf{q}(z) = \mathbf{q}(z)$. Note that the phase of the first element is $e^{-j\Omega/2}$, i.e. differs from zero phase in order to allow $\varphi(z)$ to be analytic. $\triangle$

IV. FAST ANALYTIC EIGENVECTOR EXTRACTION

A. Bin-Wise Phase Normalisation

To extract a phase-normalised analytic eigenvector $\bar{\mathbf{q}}(z)$ from a set of K sample points $\{\mathbf{q}_k\}$, we need to normalise its phase according to Sec. III. W.l.o.g. assume we want to phase-normalise the first element, $q_{1,k}$, of $\mathbf{q}_k$ such that it can be partitioned as $\mathbf{q}_k = [q_{1,k}, \mathbf{q}_r^T]^T$. We can relate $\mathbf{q}_k$ to the k th sample point of the analytic eigenvector (see (5)) via

$$\mathbf{q}_k = e^{j\psi_k}\mathbf{q}(e^{j\Omega_k}) = e^{j\psi_k} \begin{bmatrix} \varphi'(e^{j\Omega_k})d_k|q_1(e^{j\Omega_k})| \\ \mathbf{q}_r(e^{j\Omega_k}) \end{bmatrix}, \quad (14)$$

and so $e^{j\angle\{q_{1,k}\}} = e^{j\psi_k}\varphi'(e^{j\Omega_k})d_k$, with $d_k = \text{sgn}\{q_1(e^{j\Omega_k})\}$. In order to unambiguously define the phase $\angle\{q_{1,k}\}$, assume that $q_{1,k} \neq 0$ is satisfied². Further, in (14) we have defined an equivalent partitioning of $\mathbf{q}(z) = [q_1(z), \mathbf{q}_r^T(z)]^T$, and exploited $|q_{1,k}| = |q_1(e^{j\Omega_k})|$. With the latter, $d_k|q_{1,k}|$ is a sample point of $\bar{q}'_1(e^{j\Omega})$ in (7).

Example 5: To illustrate the different factors, consider $q_1(z) = 1+z^{-3} = \varphi'(z)\bar{q}_1(z)$. We have $\varphi'(z) = z^{-\frac{3}{2}}$ and $\bar{q}_1(z) = (z^{\frac{3}{2}} + z^{-\frac{3}{2}})$, as indicated in Fig. 3 on the unit circle. Also shown are the values of $\angle\{q_{1,k}\}$ and d_k in (14). $\triangle$

A first step is to phase-normalise $\mathbf{q}_k$ in (14) w.r.t. to the phase $\angle\{q_{1,k}\}$ of its first element,

$$\bar{\mathbf{q}}_k'' = \begin{bmatrix} |q_{1,k}| \\ \varphi'^H(e^{j\Omega_k})d_k\mathbf{q}_r(e^{j\Omega_k}) \end{bmatrix}. \quad (15)$$

²If this assumption is violated, it can be re-established either by shifting [17], [24] or by omitting [25] any frequency bins with singularities.




Fig. 3. Various components in the phase normalisation process of the function $q_1(z)$ of Example 5.

Note that the phase ambiguity by $e^{j\psi_k}$ is removed. However, the set of sample points $\bar{\mathbf{q}}'_k$ do not yet belong to an analytic function, since at any zero crossings of $q_1(z)$ on the unit circle, (i) $|q_{1,k}|$ experiences a phase jump by π , and (ii) in case of an odd number of zero crossings, an additional phase correction according to (10) must be considered. If, by convention, we set $\bar{q}'_1(e^{j\Omega})|_{\Omega=0} > 0$ then we can tie case (i) and (ii) together, as in (10), by recognising that an even or odd number of zero crossings will be determined by d_K , once it is known, being positive or negative. Hence $\kappa = \frac{1-d_K}{2}$, and we now only need to find d_k – the sign or zero crossings of $\bar{q}_1(e^{j\Omega})$.

B. Detecting Zero Crossings on the Unit Circle

Using (15), consider the inner product between consecutive vectors $\bar{\mathbf{q}}''_k$,

$$\bar{\mathbf{q}}_k^{\prime\prime H} \bar{\mathbf{q}}_{k-1}^{\prime\prime} = |q_{1,k}| |q_{1,k-1}| + \varphi'(e^{j\Omega_k}) \varphi'^H(e^{j\Omega_{k-1}}) d_k d_{k-1} \mathbf{q}_r^H(e^{j\Omega_k}) \mathbf{q}_r(e^{j\Omega_{k-1}}). \quad (16)$$

Since $\varphi'(e^{j\Omega})$ is analytic at least in $\Omega \in \mathbb{R}$ and $\mathbf{q}_r(z)$ is analytic in $z \in \mathbb{C}$, for a sufficiently large DFT size K , and we have³

$$\lim_{K \rightarrow \infty} \varphi'(e^{j\Omega_k}) \varphi'^H(e^{j\Omega_{k-1}}) = 1 \quad (17)$$

$$\lim_{K \rightarrow \infty} \mathbf{q}_r^H(e^{j\Omega_k}) \mathbf{q}_r(e^{j\Omega_{k-1}}) = A, \quad (18)$$

with $A \in \mathbb{R}$ and $0 \leq A \leq 1$. If frequency bins k and $k-1$ do not enclose a zero crossing of $q_{1,k}$, then $d_k = d_{k-1}$, and the r.h.s. of (16) equals $\|\mathbf{q}(e^{j\Omega_k})\|_2^2 = 1$. If the bins do enclose a zero crossing, then $|q_{1,k}| \rightarrow 0$ and hence $A \rightarrow 1$, and $d_k = -d_{k-1}$. Therefore, for a sufficiently large DFT size, for (16) we have

$$\bar{\mathbf{q}}_k^{\prime\prime H} \bar{\mathbf{q}}_{k-1}^{\prime\prime} \approx d_k d_{k-1}. \quad (19)$$

Hence starting from $d_0 = 1$, we can recursively evaluate d_k as

$$d_k = d_{k-1} \cdot \text{sgn}\{\Re\{\bar{\mathbf{q}}_k^{\prime\prime H} \bar{\mathbf{q}}_{k-1}^{\prime\prime}\}\}, \quad (20)$$

where $\Re\{\cdot\}$ takes the real part of its argument. With d_k and thus also $\kappa = (1 - d_K)/2$ determined,

$$\bar{\mathbf{q}}_k = d_k e^{-\kappa \frac{\pi k}{K}} \cdot \mathbf{q}_k'' \quad (21)$$

are the sample points of the analytic $\bar{\mathbf{q}}(z)$ on the unit circle.

³Analyticity guarantees smoothness.

C. Iterative Adjustment of the DFT Size

Because of the analyticity of $\bar{\mathbf{q}}(z)$, it is only a matter of the DFT size K being sufficiently large for the sample points $\bar{\mathbf{q}}_k$ in (20) to allow a sufficiently accurate reconstruction. Let $\bar{\mathbf{q}}^{(K)}[n]$ be the inverse DFT applied to the K sample points $\bar{\mathbf{q}}_k$. While the unit norm property of an eigenvector is satisfied in the sample points, the result may deviate at other frequencies. This can be measured as time-domain aliasing via

$$\xi_K = \sum_n |\bar{\mathbf{q}}^{(K)H}[-n] * \bar{\mathbf{q}}^{(K)}[n] - \delta[n]|^2, \quad (22)$$

where $*$ is the convolution operator and $\delta[n]$ the Kronecker function. Note with some small threshold ϵ_{PU} , $\xi_K < \epsilon_{\text{DFT}}$ is a necessary criterion for convergence, but it is not guaranteed to be sufficient.

V. PARTIAL ANALYTIC EIGENVALUE DECOMPOSITION

We now return to the problem of calculating the M analytic eigenvalues and eigenvectors, $\lambda_m(z)$ and $\mathbf{q}_m(z)$, $m = 1, \dots, M$ of $\mathbf{R}(z)$. The eigenvector $\mathbf{q}_m(z)$ can be found as in Sec. IV. To extract the corresponding eigenvalue $\lambda_m(z)$, recall that we have assumed eigenvalues to be spectrally majorised. In (4) we have $\mathbf{P}_k = \mathbf{I}$, and given a sufficiently large DFT size, $\lambda_m[n] \circ \lambda_m(z)$ can be extracted by an inverse DFT from the sample points. Alternatively we can determine

$$\lambda_m(z) = \bar{\mathbf{q}}_m^P(z) \mathbf{R}(z) \bar{\mathbf{q}}_m(z) = \mathbf{q}_m^P(z) \mathbf{R}(z) \mathbf{q}_m(z), \quad (23)$$

since $\bar{\mathbf{q}}_m(z) = \varphi_m^P(z) \mathbf{q}_m(z)$. Because analytic eigenvectors are unique, they are thus independent of the fact that the support of $\bar{\mathbf{q}}_m(z)$ may not be as compact as that of $\mathbf{q}_m(z)$. Thus, an eigenpair $\{\lambda_m(z), \mathbf{q}_m(z)\}$ can be extracted.

If $\mathbf{R}(z)$ has reduced rank or we want a reduced rank approximation, we can extract the first $L < M$ eigenpairs until we find that the power in the $(L+1)$ th eigenvalue falls below a preset threshold. The power in the $(L+1)$ th eigenvalue is given by $\lambda_{L+1}[0] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \lambda_{L+1}(e^{j\Omega}) d\Omega$. Only calculating $L < M$ dominant eigenpairs can be computationally advantageous. Also, current DFT domain methods for the analytic EVD [14]–[17] require $\mathbf{R}(z)$ to have full spatial rank and therefore cannot be employed if $\mathbf{R}(z)$ is low rank.

VI. SIMULATION AND RESULTS

We test the proposed fast partial analytic EVD on an ensemble of random parahermitian matrices $\mathbf{R}(z)$ of size $M = 30$ and of variable order $\text{ord}\{\mathbf{R}(z)\} \in \{8, 16, 24, 32, 40\}$, generated via $\mathbf{R}(z) = \mathbf{A}(z) \mathbf{A}^P(z)$, where $\mathbf{A}(z) : \mathbb{C} \rightarrow \mathbb{C}^{M \times L}$ with $L = 4$ defining the spatial rank, and therefore the number of non-zero analytic eigenvalues, of $\mathbf{R}(z)$. The matrices $\mathbf{A}(z)$ are generated from zero mean unit variance complex Gaussian entries. We compare the proposed method to both SBR2 [9] and the SMD algorithms [2] with a maximum number of 500 and 300 iterations, respectively. To assess the performance of the algorithms, we measure the execution time under Matlab 2025b on an Intel Core i7 (2.3GHz) processor, and the normalised reconstruction error $\zeta = (\sum_{\tau} \|\mathbf{R}[\tau] - \hat{\mathbf{R}}[\tau]\|_F^2) / (\sum_{\tau} \|\mathbf{R}[\tau]\|_F^2)$, between $\mathbf{R}(z)$




Fig. 4. Ensemble result for algorithm comparison: (a) execution time and (b) reconstruction error vs the polynomial order of $R(z)$.

and $\hat{R}(z) = \sum_{m=1}^L \lambda_m(z) \bar{q}_m(z) \bar{q}^P(z)$ based on the extracted eigenpairs $\{\lambda_m(z), \bar{q}_m(z)\}$.

Fig. 4(a) shows the median and 25th and 75th percentiles of the ensemble results in terms execution time. The proposed algorithm is substantially faster than the benchmark methods, which operate on the entire matrix. The ensemble results for the construction error in Fig. 4(b) demonstrates the higher accuracy of proposed algorithm over the benchmark methods. The larger spread of the distribution here is due to the fact that the proposed method was stopped at a maximum DFT length of 2^5 times $\text{ord}\{R(z)\}$ rather than reaching the threshold $\epsilon_{PU} = 10^{-5}$.

VII. CONCLUSION

We have proposed an analytic eigenvector extraction algorithm for a polynomial parahermitian matrix that possesses spectrally majorised eigenvalues and can be rank deficient. Operating in the DFT domain, coherence across bins has to be re-established and we have replaced an expensive phase-smoothing approach by a bin-wise normalisation that provides sample points of an analytic solution. For low-rank parahermitian matrices, it is possible to operate a partial analytic EVD that only extracts a limited number of dominant eigenpairs. We have outlined an algorithm to extract this solution. Compared to benchmark methods that are capable of operating on rank-deficient parahermitian matrices, it can operate with significantly shorter execution time and higher accuracy.

REFERENCES

- [1] V. Neo, C. Evers, S. Weiss, and P.A. Naylor, "Signal compaction using polynomial EVD for spherical array processing with applications," *IEEE/ACM Trans. ASLP*, vol. 31, pp. 3537–3549, 2023.
- [2] S. Redif, S. Weiss, and J.G. McWhirter, "Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices," *IEEE Trans. Signal Process.*, vol. 63, no. 1, pp. 81–89, January 2015.
- [3] V.W. Neo, C. Evers, and P.A. Naylor, "Enhancement of noisy reverberant speech using polynomial matrix eigenvalue decomposition," *IEEE/ACM Trans. ASLP*, vol. 29, pp. 3255–3266, 2021.
- [4] A. Hogg, V. Neo, S. Weiss, C. Evers, and P. Naylor, "A polynomial eigenvalue decomposition MUSIC approach for broadband sound source localization," in *IEEE WASPAA*, New Paltz, NY, Oct. 2021.

- [5] 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 Magazine*, vol. 40, no. 7, pp. 18–37, Nov. 2023.
- [6] S. Weiss, J. Pestana, and I.K. Proudler, "On the existence and uniqueness of the eigenvalue decomposition of a parahermitian matrix," *IEEE Trans. Signal Process.*, vol. 66, no. 10, pp. 2659–2672, May 2018.
- [7] 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 Process.*, vol. 66, no. 23, pp. 6325–6327, Dec. 2018.
- [8] G. Barbarino and V. Noferini, "On the Rellich eigendecomposition of para-Hermitian matrices and the sign characteristics of $\ast$ -palindromic matrix polynomials," *Linear Algebra Appl.*, vol. 672, pp. 1–27, Sep. 2023.
- [9] J.G. McWhirter, P.D. Baxter, T. Cooper, S. Redif, and J. Foster, "An EVD algorithm for para-Hermitian polynomial matrices," *IEEE Trans. Signal Process.*, vol. 55, no. 5, pp. 2158–2169, May 2007.
- [10] 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 Process.*, vol. 59, no. 11, pp. 5253–5264, Nov. 2011.
- [11] 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 Conf.*, Nice, France, Sep. 2015, pp. 844–848.
- [12] V.W. Neo and P.A. Naylor, "Second order sequential best rotation algorithm with householder reduction for polynomial matrix eigenvalue decomposition," in *IEEE ICASSP*, Brighton, UK, May 2019, pp. 8043–8047.
- [13] J. Corr, K. Thompson, S. Weiss, J.G. McWhirter, S. Redif, and I.K. Proudler, "Multiple shift maximum element sequential matrix diagonalisation for parahermitian matrices," in *IEEE Workshop on Statistical Signal Processing*, Gold Coast, Australia, June 2014, pp. 312–315.
- [14] M. Tohidian, H. Amindavar, and A. M. Reza, "A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices," *EURASIP J. Advances in Signal Process.*, vol. 2013, no. 1, pp. 1–16, 2013.
- [15] S. Weiss, I.K. Proudler, and F.K. Coutts, "Eigenvalue decomposition of a parahermitian matrix: extraction of analytic eigenvalues," *IEEE Trans. Signal Process.*, vol. 69, pp. 722–737, 2021.
- [16] 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 Process.*, vol. 71, pp. 1642–1656, Apr. 2023.
- [17] F.A. Khattak, I.K. Proudler, and S. Weiss, "Scalable analytic eigenvalue extraction algorithm," *IEEE Access*, vol. 12, pp. 166 652–166 659, Dec. 2024.
- [18] P.P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs: Prentice Hall, 1993.
- [19] 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 Conf. Signals, Systems, and Computers*, Pacific Grove, CA, Oct. 2022.
- [20] P. Vaidyanathan, "Theory of optimal orthonormal subband coders," *IEEE Trans. Signal Process.*, vol. 46, no. 6, pp. 1528–1543, Jun. 1998.
- [21] L. V. Ahlfors, *Complex analysis: An introduction to the theory of analytic functions of one complex variable*. New York: McGraw-Hill, 1953.
- [22] 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 Process.*, vol. 72, pp. 2260–2275, 2024.
- [23] S. Weiss and G. Barbarino, "Relations between analytic spectral and singular value decompositions," in *33rd European Signal Processing Conf.*, Palermo, Italy, Sep. 2025, pp. 2272–2276.
- [24] F.A. Khattak, I.K. Proudler, S. Weiss, and S. J. Schlecht, "Analytic Gram-Schmidt orthogonalisation and QR decomposition for polynomial matrices," *IEEE Trans. Signal Process.*, 2026, submitted.
- [25] S. Weiss, S.J. Schlecht, and M. Moonen, "Best least squares paraunitary approximation: Analytic procrustes problem," *IEEE Trans. Signal Process.*, to appear Feb. 2026.