Scalable Approach for Analytic Polynomial Subspace Projection Matrices for a Space-Time Covariance Matrix

Faizan A. Khattak, Mohammed Bakhit, Ian K. Proudler, and Stephan Weiss
Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, Scotland
{faizan.khattak,ian.proudler,stephan.weiss}@strath.ac.uk

***Abstract*—In sensor array applications, it can be advantageous to project data onto a given signal subspace, for example, to improve the SNR or as part of direction finding algorithms. In the broadband case, a projection operator can be derived via polynomial matrices and, more specifically, from a space-time covariance matrix. Traditional methods perform a complete polynomial eigenvalue decomposition (PEVD) to achieve this, which can be computationally intensive. We propose a novel method to compute these subspace matrices directly, without the need for a full PEVD. Our approach is evaluated against existing methods using an ensemble of randomized para-Hermitian matrices, demonstrating significant improvements in both accuracy and computation time.**

I. Introduction

Over the past two decades, the extension of ordinary matrix algebra to polynomial matrices has found successful applications in broadband beamforming [1, 2], speech enhancement [3], broadband MIMO design [4–6], subband coding [7], broadband angle of arrival estimation [8], and broadband blind source separation [9, 10], voice activity detection [11], and transient signal detection [12, 13] to name but a few. The majority of these applications require the estimation of a space-time covariance matrix to capture the second-order statistics between sensor signals. To decorrelate these signals or extract dominant signal components, a polynomial eigenvalue decomposition (PEVD) is typically computed. The estimated eigenvectors are then used to compute subspace projection matrices.

Generally speaking, the 'diagonalization' PEVD algorithms, like sequential second order best rotation (SBR2) [7, 14], sequential matrix diagonalization (SMD) [15], and their variants [16, 17] are often employed due to their comparatively lower computational cost compared to the analytic, or compact polynomial order, PEVD algorithms [18–20]. However, these iterative approaches only provide approximate diagonalization while also spectrally majorizing the eigenvalues [14], leading to leakage between subspaces [11]. Although this leakage can be mitigated by using analytic EVD algorithms, such algorithms are currently very computationally expensive even when resorting to optimised implementations [21–23].

In this document, we highlight the need to compute analytic subspace projection matrices without explicitly performing the analytic EVD. We propose a Fourier domain method that computes the ordinary EVD of each of the discrete Fourier transform (DFT) samples of the space-time covariance matrix. The eigenvalues and eigenvectors in each DFT bin are then permuted to create a smooth association across the bins. These eigenvectors are used to estimate the DFT samples of the analytic subspace projection matrices. The DFT size is iteratively increased until the time-domain aliasing incurred in reverting these subspace projection matrices to the time domain becomes negligible.

Mohammed Bakhit is partially funded by MathWorks.

II. Existence of Analytic EVD & Subspace Projection Matrices

A. Existence

A cross-spectral density matrix $\boldsymbol{R}(z)$ is the $z-$transform of a space-time covariance matrix $\mathbf{R}[\tau]$, $\tau \in \mathbb{Z}$, such that $\boldsymbol{R}(z) = \sum_\tau \mathbf{R}[\tau] z^{-\tau}$. It is para-Hermitian symmetric i.e. it is equal to its para-Hermitian conjugate $\boldsymbol{R}^{\mathrm{P}}(z) = \boldsymbol{R}^{\mathrm{H}}(1/z^*)$ i.e. Hermitian conjugation plus time reversal [24]. Any such para-Hermitian matrix $\boldsymbol{R}(z) : \mathbb{C} \to \mathbb{C}^{M\times M}$, which is analytic in nature and does not come from a system with multiplexing, admits an analytic EVD [25, 26] as

$$\boldsymbol{R}(z) = \boldsymbol{Q}(z)\boldsymbol{\Lambda}(z)\boldsymbol{Q}^{\mathrm{P}}(z)\,. \tag{1}$$

The analytic eigenvectors are contained within the columns of $\boldsymbol{Q}(z) = [\boldsymbol{q}_1(z), \ldots, \boldsymbol{q}_M(z)]$, which is a paraunitary matrix, i.e., $\boldsymbol{Q}(z)\boldsymbol{Q}^{\mathrm{P}}(z) = \mathbf{I}$, and $\boldsymbol{\Lambda}(z) = \mathrm{diag}\{\lambda_1(z), \ldots, \lambda_M(z)\}$ is a para-Hermitian diagonal matrix of analytic eigenvalues $\lambda_m(z)$, $m = 1, \ldots, M$. We here assume that the analytic eigenvalues are distinct, such that on the unit circle $\sigma_m(\mathrm{e}^{\mathrm{j}\Omega}) = \sigma_\mu(\mathrm{e}^{\mathrm{j}\Omega})$, $\forall\Omega$, is only possible in the case $m = \mu$.

The analytic eigenvectors $\boldsymbol{q}_m(z)$, $m = 1, \ldots, M$, are ambiguous up to arbitrary all-pass functions, i.e., if $\phi_m(z)$ is an all-pass function, then $\boldsymbol{q}_m(z)\phi_m(z)$ is also a valid mth eigenvector. The polynomial orders of $\boldsymbol{q}_m(z)$ and $\boldsymbol{q}_m(z)\phi_m(z)$ can be radically different. This phenomenon complicates algorithms that seek analytic eigenvectors if they are required to be low-order polynomials as the all-pass function $\phi_m(z)$ that minimises the order of the eigenvectors must be found [20]. In earlier PEVD algorithms [14, 15], this order issue has been coarsely addressed by identifying simple allpass functions in the form of delays [27], and by order truncation [28]; this however affects the precision of the PEVD factorisation.

B. Subspace Projection

The analytic subspace projection matrix corresponding to the mth eigenvector $\boldsymbol{q}_m(z)$ is given by

$$\boldsymbol{\mathcal{P}}_m(z) = \boldsymbol{q}_m(z)\boldsymbol{q}_m^{\mathrm{P}}(z) \; . \tag{2}$$

This Laurent polynomial subspace projection matrix $\boldsymbol{\mathcal{P}}_m(z)$ is para-Hermitian symmetric. As such, the all-pass ambiguity of the eigenvectors cancels, i.e.,

$$\boldsymbol{q}_m(z)\phi_m(z)\left(\boldsymbol{q}_m(z)\phi_m(z)\right)^{\mathrm{P}} = \boldsymbol{q}_m(z)\boldsymbol{q}_m^{\mathrm{P}}(z) \; . \tag{3}$$

The calculation of a suitable allpass function $\phi_m(z)$ for each analytic eigenvector in [20] is generally an NP-hard problem. Hence if it was possibe to calculate $\boldsymbol{\mathcal{P}}_m(z)$ directly, bypassing the explicit computation of a $\phi_m(z)$ that minimises the support of the eigenvectors $\boldsymbol{q}_m(z)$, then the polynomial subspace projection matrices could be calculated at a much lower cost.

III. Proposed Approach

A. Bin-Wise EVD $\mathbf{R}[\tau]$

The proposed approach first computes the DFT sample point $\mathbf{R}_k = \boldsymbol{R}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \boldsymbol{R}(z)\big|_{z=\mathrm{e}^{\mathrm{j}\Omega_k}}$ where $\Omega_k = 2\pi k/K$ for $k = 0, \ldots, (K-1)$, i.e. the result of a K-point DFT applied to $\mathbf{R}[\tau]$. Note that $\mathbf{R}_k$ is Hermitian symmetric and so admits an ordinary EVD producing

$$\boldsymbol{R}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{R}_k = \mathbf{Q}_k\boldsymbol{\Lambda}_k\mathbf{Q}_k^{\mathrm{H}} \quad \text{for} \quad k = 0, \ldots, (K-1) \; , \tag{4}$$

where $\boldsymbol{\Lambda}_k = \mathrm{diag}\{d_{1,k}, \ldots, d_{M,k}\}$, $\quad d_{m,k} \geq d_{m+1,k}, \quad m = 1, \ldots, (M-1)$ is a real diagonal matrix of eigenvalues, and $\mathbf{Q}_k$ represents the eigenvector matrix at the kth sample point. As (4) is computed independently in different DFT bins, we will refer to it as the bin-wise EVD throughout this document.

With the bin-wise EVD computed, the EVD factors in different bins can be related to the respective analytic functions in (1) sampled at $z = \mathrm{e}^{\mathrm{j}\Omega_k}$ akin to [19, 20] as

$$\boldsymbol{\Lambda}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{P}_k\boldsymbol{\Lambda}_k\mathbf{P}_k^{\mathrm{H}} \tag{5}$$

$$\boldsymbol{Q}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{Q}_k\mathbf{A}_k\mathbf{P}_k\boldsymbol{\Phi}_k \tag{6}$$

where $\mathbf{P}_k$ is a permutation matrix and $\boldsymbol{\Phi}_k$ is a diagonal matrix containing phases for establishing phase coherence between adjacent DFT bins. If the k-th bin has non-trivial algebraic multiplicities, the unitary matrix $\mathbf{A}_k$ aligns the 1-d subspaces in which the eigenvectors reside. If there are no non-trivial algebraic multiplicities, $\mathbf{A}_k$ is an identity matrix.

B. Determining Permutation Matrices

To avoid having to calculate $\mathbf{A}_k$ when there are non-trivial algebraic multiplicities, we adopt the approach from [23]. Assume that $\boldsymbol{R}(\mathrm{e}^{\mathrm{j}\Omega_k})$ where $\Omega_k = 2\pi k/K$ has some non-trivial algebraic multiplicities. We replace this sample point with $\Omega_k = 2\pi(k+q_k)/K$ where $0 < q_k \ll 1$ is such that $\boldsymbol{R}(\mathrm{e}^{\mathrm{j}\Omega_k})$ now has no non-trivial algebraic multiplicities. This is possible due the analytic nature of the eigenvalues [23]. This results in a set of bin-wise EVDs that relate to the analytic EVD functions as in (5) and (6) but with $\mathbf{A}_k = I \; \forall k$. The determination of the permutation matrices $\mathbf{P}_k$ are then determined by assessing the orthogonality of the bin-wise eigenvectors between adjacent DFT bins [23]. The penalty for this simplification is that our sample grid is non-uniform which slightly complicates performing an IFFT [29, 30] – see Sec. III-D.

The procedure for determining the permutation $\mathbf{P}_k$ is as follows: with $\mathbf{P}_0 = \mathbf{I}$ and $\bar{\mathbf{Q}}_0 = \mathbf{Q}_0\mathbf{P}_k = [\bar{\mathbf{q}}_{1,0}, \ldots \bar{\mathbf{q}}_{M,0}]$, for $m = 1, \ldots, M$, we determine

$$n_{m,k} = \underset{n}{\mathrm{argmax}} \; |\bar{\mathbf{q}}_{m,k-1}^{\mathrm{H}}\mathbf{q}_{n,k}|, n = 1, \ldots, M \; , \tag{7}$$

starting with $k = 1$. Once for a certain $0 < k \leq (K-1)$, we have determined the quantity $n_{m,k}, m = 1, \ldots, M$, the corresponding permutation matrix in the kth bin is equal to

$$\mathbf{P}_k = [\mathbf{i}_{n_{1,k}}, \ldots, \mathbf{i}_{n_{M,k}}] \; , \tag{8}$$

where $\mathbf{i}_m$ is the mth column of an $M \times M$ identity matrix. Before proceeding to the next bin, we update $\bar{\mathbf{Q}}_k = \mathbf{Q}_k\mathbf{P}_k = [\bar{\mathbf{q}}_{1,k}, \ldots \bar{\mathbf{q}}_{M,k}]$, where $\bar{\mathbf{q}}_{m,k}$ is now the permutation-corrected mth eigenvaue on the kth bin. This process is repeated for subsequent bins $k = 1, \ldots, (K-1)$.

C. Bin-Wise Subspace Projection Matrices

With $\mathbf{P}_k, k = 0, \ldots, (K-1)$ determined and $\mathbf{A}_k = \mathbf{I} \; \forall k$ due to frequency shifts, (6) can be reformulated as

$$\boldsymbol{Q}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \bar{\mathbf{Q}}_k\boldsymbol{\Phi}_k \; . \tag{9}$$

The sample points of the analytic subspace projection matrix in (2) can be related to the factors in (9) as

$$\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_k}) = \boldsymbol{q}_m(\mathrm{e}^{\mathrm{j}\Omega_k})\boldsymbol{q}_m^{\mathrm{H}}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \bar{\mathbf{q}}_{m,k}\bar{\mathbf{q}}_{m,k}^{\mathrm{H}} \; , \tag{10}$$

where the phase ambiguity, $\boldsymbol{\Phi}_k$, cancels out similar to the allpass ambiguity cancellation in (3). Thus samples of the analytic subspace projection matrix can be obtained directly from the bin-wise permuted eigenvector matrix without the need to address the phase ambiguities.

D. Sufficient DFT Size Determination

In general, $\boldsymbol{Q}(z)$ is a matrix valued transcendental function [25] and, therefore, has an infinite Laurent series representation. The same is true for the corresponding subspace projection matrix constructed with $\boldsymbol{Q}(z)$. However, due to analyticity, even the infinite Laurent series must converge absolutely. As a result, these analytic subspace projection matrices can be approximated in the least square sense by para-Hermitian polynomial matrices of finite order through truncation and masking, similar to the Laurent polynomial approximation of analytic eigenvectors [20].

To achieve an acceptable finite order approximation for the analytic subspace projection matrices in the least squares sense, we rely on the time-domain aliasing incurred in the inverse DFT (IDFT) of the sample points $\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_k})$. To capture the time-domain aliasing, we compare the time-domain sequences obtained from two different DFT sizes, K and $K/2$,

i.e. $\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_k})|_{k=0,\ldots,(K-1)}$ and $\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_{2k}})|_{k=0,\ldots,K/2-1}$. The metric used is

$$\gamma_m = \frac{\sum_\tau \|\hat{\mathcal{P}}_m^{(K)}[\tau] - \hat{\mathcal{P}}_m^{(K/2)}[\tau]\|_{\mathrm{F}}^2}{\sum_\tau \|\hat{\mathcal{P}}_m^{(K)}[\tau]\|_{\mathrm{F}}^2}\,, \tag{11}$$

where $\hat{\mathcal{P}}_m^{(K)}[\tau]$ is the time-domain sequence obtained from K-point IDFT, and $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm. The DFT size is iteratively increased until γ_m falls below a sufficiently small non-zero preset threshold ϵ.

Note however that if it was necessary to adjust the positions of the frequency domain sample points to avoid non-trivial algebraic multiplicities as in section III-B the frequency domain samples of the analytic subspace projection matrix $\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_k})|_{k=0,\ldots,(K-1)}$ are on a non-uniform grid. Hence the inverse DFT could, in principal, require inverting a $K \times K$ matrix or simply apply a non-uniform inverse fast Fourier transform in [29, 30]. However, there is a more economical approach for computing the inverse of the non-uniform DFT for the case where only some $k' \ll K$ DFT bins are shifted. Such a low-cost approach can be applied similarly to [23].

E. Overall Procedure

The procedure begins by setting the initial DFT size equal to the next power of 2 higher than the polynomial order of $\boldsymbol{R}(z)$, denoted as $\mathcal{O}\{\boldsymbol{R}(z)\}$. Then EVDs are computed within the DFT bins, with shifts applied to those DFT bins where the minimum eigenvalue distance $|d_{m,k} - d_{m+1,k}|$ in any bin is less than ε. Instead of setting ε to 0, we set it to a non-zero small value of the order 10^{-6} to avoid tightly clustered eigenvalues, which may also lead to discontinuities in the alignment of 1-d subspaces [25, 31]. The permutation matrices $\mathbf{P}_k$ are determined through (7) and (8) bin-by-bin. In each iteration, we updates $\bar{\mathbf{Q}}_k = \bar{\mathbf{Q}}\mathbf{P}_k$, since $\bar{\mathbf{Q}}_{k-1}$ of the previously sorted bin is required to determine the permutation for the new bin. Subsequently, the samples of the analytic subspace projection matrix are computed through (10). Finally, the algorithm computes γ via (11), and if γ is smaller than ϵ, the algorithm terminates; otherwise, the DFT size is doubled. The latter allows to rely on efficient FFT calculations, and to reuse the EVDs in half of those bins where they have been already evaluated previously. This entire procedure is outlined in Algorithm 1, yielding an approximation $\hat{\boldsymbol{\mathcal{P}}}_m(z)$ of the projection matrix in (2).

Algorithm 1: Subspace Projection Matrix Algorithm

Input: $\boldsymbol{R}(z), \epsilon, \varepsilon_\lambda$;
Output: $\hat{\boldsymbol{\mathcal{P}}}_m(z)$;
initialise $K = 2^{\mathrm{ceil}[\log_2(\mathcal{O}\{\boldsymbol{R}(z)\})]},\ \gamma = 1+\epsilon,\ \mathbf{P_0} = \mathbf{I}$;
while $\gamma > \epsilon$ **do**
  determine EVDs in K DFT bins with shifting bin frequency if $\min |d_{m,k} - d_{m+1,k}| < \varepsilon_\lambda$;
  $\bar{\mathbf{Q}}_0 = \mathbf{Q}_0\mathbf{P}_0$;
  for $k = 1 : (K-1)$ **do**
    determine $n_{m,k}$ via (7) for $m = 1, \ldots, M$;
    determine $\mathbf{P}_k$ via $n_{m,k}$ from (8);
    $\bar{\mathbf{Q}}_k = \mathbf{Q}_k\mathbf{P}_k$;
  end
  $\boldsymbol{\mathcal{P}}_m(\mathrm{e}^{\mathrm{j}\Omega_k}) = \bar{\mathbf{q}}_{m,k}\bar{\mathbf{q}}_{m,k}^{\mathrm{H}}$ for $k = 0, \ldots, (K-1)$;
  determine γ_m via (11);
  $K \leftarrow 2K$;
  What about the IDFT?
end
$\hat{\boldsymbol{\mathcal{P}}}_m(z) = \sum_\tau \hat{\mathcal{P}}_m^{(K/2)}[\tau] z^{-\tau}$

IV. Simulations and Results

A. Performance Metrics

The proposed approach is compared against the SBR2 [14] and SMD [15] methods via following performance metrics:

- polynomial order $\mathcal{O}\{\hat{\boldsymbol{\mathcal{P}}}_m(z)\}$ of $\hat{\boldsymbol{\mathcal{P}}}_m(z)$;
- the accuracy of the estimated analytic subspace projection matrix is which measured w.r.t. its difference with the ground truth projection matrices as

$$\xi = \sum_{m=1}^{M} \frac{\sum_\tau \|\mathcal{P}_m[\tau] - \hat{\mathcal{P}}_m[\tau]\|_{\mathrm{F}}^2}{\sum_\tau \|\mathcal{P}_m[\tau]\|_{\mathrm{F}}^2}\,; \tag{12}$$

- algorithm execution time.

B. Ensemble Tests

The proposed approach and benchmark algorithms are compared over an ensemble of both spectrally majorised and unmajorised para-Hermitian matrices constructed through the source model in [7, 15, 19]. This source model is based on generating $\boldsymbol{R}(z)$ from a sequence of elementary paraunitary matrices [24] and innovation filters [32] that shape $\lambda_m(z)$, such that the ground truth factors of (1) are known. We compare the statistics for the above performance metrics over an ensemble of 200 instantiations of $\boldsymbol{R}(z)$ with a spatial dimension of $M = 3$ for each value of $\boldsymbol{\Lambda}(z) = \{4, 8, 12, \ldots, 40\}$ and $\mathcal{O}\{\boldsymbol{Q}(z)\} = \{2, 4, \ldots, 20\}$. Since the performance comparison can vary significantly depending on the nature of the power spectral densities of the broadband sources, we provide results for both overlapping and non-overlapping cases of the sources' power spectral densities.

The SMD and SBR2 algorithms are executed for a maximum of 300 iterations unless the off-diagonal energy falls below 10^{-5}. Intermediate paraunitary and partially diagonalized matrices are truncated via a threshold of 10^{-5} [28]. For lower-order paraunitary matrices, the row-shift corrected truncation approach is used from [27] The proposed approach is executed with $\epsilon = \varepsilon_\lambda = 10^{-6}$. The following simulation results were obtained on an 11th Gen Intel(R) Core(TM) i7-11800H processor running at 2.30GHz with 16 CPUs.

The ensemble test results for spectrally majorised instantiations are illustrated in Fig. 1, where the execution time of the proposed approach in Fig. 1(a) is considerably less than that of the benchmark algorithms. Notably, the proposed approach's execution time for an ensemble consisting of instantiations with a spatial dimension of $M = 9$ is also less than the benchmark algorithms' time for $M = 3$. Similarly, the accuracy of the subspace projection matrices produced by the proposed approach is orders of magnitude higher than the

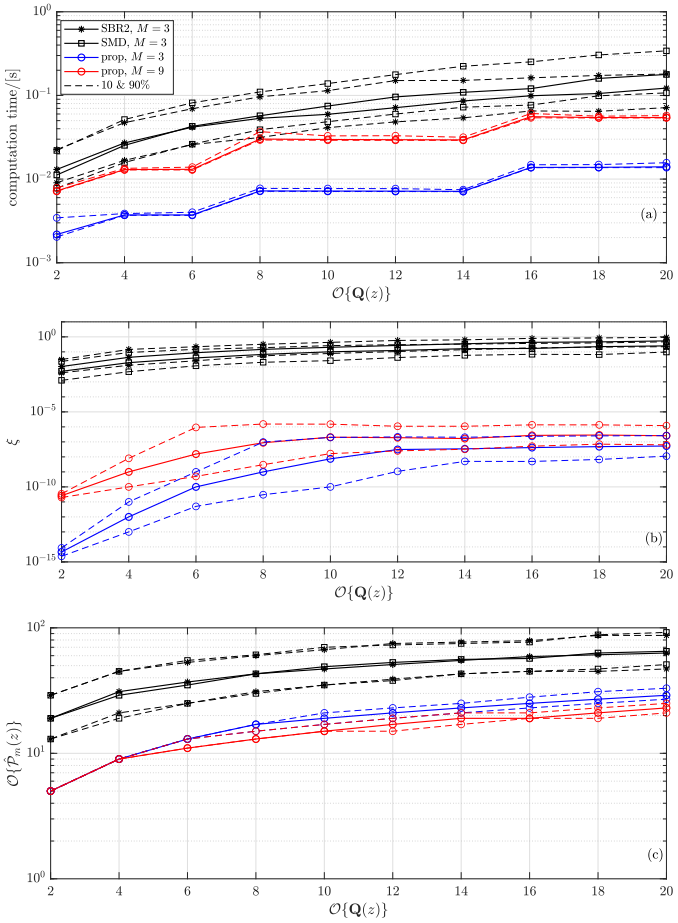

Fig. 1. Ensemble test performed over 200 instantiations of majorised $\boldsymbol{R}(z)$ with (a) showing the execution time, (b) normalized error, and (c) polynomial order of the produced subspace projection matrices.

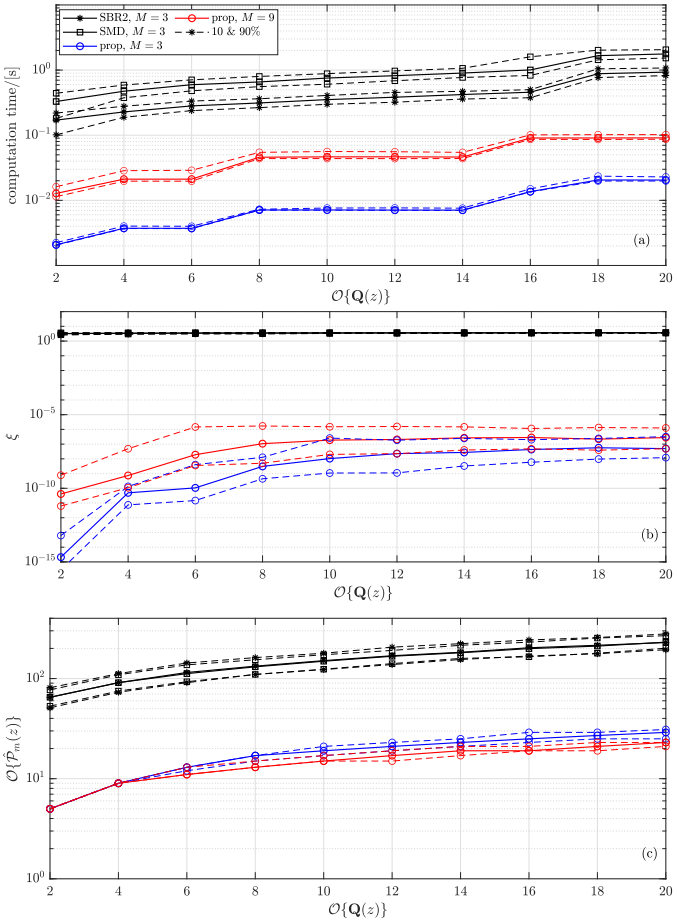

Fig. 2. Ensemble test performed over 200 instantiations of unmajorised $\boldsymbol{R}(z)$ with (a) showing the execution time, (b) normalized error, and (c) polynomial order of the produced subspace projection matrices.

results produced by SBR2 and SMD, as shown in Fig. 1(b). This is due to the fact that both SMD and SBR2 are only able to approximate (1) [25]. Lastly, the polynomial order of the subspace matrices is lower for all values $\mathcal{O}\{\boldsymbol{Q}\}(z)$ for the proposed method, as shown in Fig. 1(c). This reduction in polynomial order results in lower-cost implementation, minimizing resource usage when these filter banks are deployed on DSPs or FPGAs.

For unmajorised instantiations, the benchmark algorithms show increased execution time, error, and polynomial order due to their tendency to converge to a spectrally majorised solution (see Fig. 2). Specifically, the ξ metric is significantly higher for SBR2/SMD, exceeding 1, reflecting the unsuitability of these methods for spectrally unmajorised para-Hermitian matrices. Since SBR2 and SMD encourage or can even be shown to converge towards [33] a spectrally majorised solution, their approximation of (1) becomes even coarser [25]. In contrast, the proposed approach demonstrates orders of magnitude better performance across all three metrics. The scalability of the proposed approach is shown through simulations with a spatial dimension of $M = 9$, where it still considerably outperforms the benchmark algorithms' results for $M = 3$.

V. Conclusion

A novel method for estimating analytic subspace projection matrices has been introduced, which bypasses the need for computing a full analytic eigenvalue decomposition (EVD). This approach involves performing an ordinary EVD within the DFT bins of a para-Hermitian polynomial matrix. It then ensures smooth association across the DFT bins by evaluating the orthogonality of eigenvectors in adjacent bins. Compared to existing algorithms, the proposed method demonstrates superior performance in execution time, accuracy and polynomial order of the subspace projection matrix.

References

[1] S. Weiss, S. Bendoukha, A. Alzin, F. Coutts, I. Proudler, and J. Chambers, "MVDR broadband beamforming using polynomial matrix techniques," in *23rd European Signal Processing Conference*, Nice, France, Sep. 2015, pp. 839–843.

[2] V. Neo, S. Redif, J. McWhirter, J. Pestana, I. Proudler, S. Weiss, and P. Naylor, "Polynomial eigenvalue decomposition for multichannel broadband signal processing: A mathematical technique offering new insights and solutions," *IEEE Signal Processing Magazine*, vol. 40, no. 7, pp. 18–37, Nov. 2023.

[3] V. W. Neo, C. Evers, and P. A. Naylor, "Enhancement of noisy reverberant speech using polynomial matrix eigenvalue decomposition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3255–3266, 2021.

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

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

[6] J. Foster, J. McWhirter, S. Lambotharan, I. Proudler, M. Davies, and J. Chambers, "Polynomial matrix QR decomposition for the decoding of frequency selective multiple-input multiple-output communication channels," *IET Signal Processing*, vol. 6, no. 7, pp. 704–712, Sep. 2012.

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

[8] M. Alrmah, S. Weiss, and S. Lambotharan, "An extension of the MUSIC algorithm to broadband scenarios using polynomial eigenvalue decomposition," in *19th European Signal Processing Conference*, Barcelona, Spain, August 2011, pp. 629–633.

[9] S. Redif, S. Weiss, and J. McWhirter, "Relevance of polynomial matrix decompositions to broadband blind signal separation," *Signal Processing*, vol. 134, pp. 76–86, May 2017.

[10] V. Neo, C. Evers, S. Weiss, and P. Naylor, "Signal compaction using polynomial EVD for spherical array processing with applications," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 3537–3549, 2023.

[11] V. W. Neo, S. Weiss, and P. A. Naylor, "A polynomial subspace projection approach for the detection of weak voice activity," in *Sensor Signal Processing for Defence Conference*, London, UK, Sep. 022, pp. 1–5.

[12] S. Weiss, C. Delaosa, J. Matthews, I. Proudler, and B. Jackson, "Detection of weak transient signals using a broadband subspace approach," in *International Conference on Sensor Signal Processing for Defence*, Edinburgh, Scotland, Sep. 2021, pp. 65–69.

[13] C. A. D. Pahalson, L. H. Crockett, and S. Weiss, "Detection of weak transient broadband signals using a polynomial subspace and likelihood ratio test approach," in *32nd European Signal Processing Conference*, Lyon, France, Aug. 2024.

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

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

[16] 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, June 2014, pp. 312–315.

[17] 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, Sep. 2015, pp. 844–848.

[18] M. Tohidian, H. Amindavar, and A. M. Reza, "A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices," *EURASIP Journal on Advances in Signal Processing*, vol. 2013, no. 1, pp. 1–16, 2013.

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

[20] S. Weiss, I. Proudler, F. Coutts, and F. Khattak, "Eigenvalue decomposition of a parahermitian matrix: extraction of analytic eigenvectors," *IEEE Transactions on Signal Processing*, vol. 71, pp. 1642–1656, Apr. 2023.

[21] F. K. Coutts, I. K. Proudler, and S. Weiss, "Efficient implementation of iterative polynomial matrix evd algorithms exploiting structural redundancy and parallelisation," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 66, no. 12, pp. 4753–4766, Dec. 2019.

[22] F. A. Khattak, S. Weiss, and I. K. Proudler, "Fast givens rotation approach to second order sequential best rotation algorithms," in *International Conference in Sensor Signal Processing for Defence*, Edinburgh, Scotland, Sep. 2021, pp. 40–44.

[23] F. A. Khattak, I. K. Proudler, and S. Weiss, "Scalable analytic eigenvalue extraction algorithm," *IEEE Access*, 2024, submitted.

[24] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs: Prentice Hall, 1993.

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

[26] 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 Transactions on Signal Processing*, vol. 66, no. 23, pp. 6325–6327, Dec. 2018.

[27] J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter, "Shortening of paraunitary matrices obtained by polynomial eigenvalue decomposition algorithms," in *Sensor Signal Processing for Defence*, Edinburgh, Scotland, Sep. 2015.

[28] 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.

[29] D. Potts, G. Steidl, and M. Tasche, *Fast Fourier Transforms for Nonequispaced Data: A Tutorial*. Boston, MA: Birkhäuser Boston, 2001, pp. 247–270.

[30] G. Plonka, D. Potts, G. Steidl, and M. Tasche, *Numerical Fourier Analysis*. Birkhäuser, 2019.

[31] T. Kato, *Perturbation Theory for Linear Operators*. Springer, 1980.

[32] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. New York: McGraw-Hill, 1991.

[33] J. G. McWhirter and Z. Wang, "A novel insight to the SBR2 algorithm for diagonalising para-hermitian matrices," in *11th IMA Conference on Mathematics in Signal Processing*, Birmingham, UK, Dec. 2016.