Khattak, F. A., Proudler, I., & Weiss, S. (in press). Scalable extraction of analytic eigenvalues from a parahermitian matrix. 1-5.
Paper presented at 32nd European Signal Processing Conference, Lyon, France.

Scalable Extraction of Analytic Eigenvalues from a Parahermitian Matrix

Faizan A. Khattak, 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 order to determine the analytic eigenvalues of a parahermitian matrix, the state-of-the-art algorithm offers proven convergence but its complexity grows factorially with the matrix dimension. Operating in the discrete Fourier transform (DFT) domain, its computational bottleneck is a maximum likelihood (ML) sequence estimation, that investigates a set of paths of likely associations across DFT bins. Therefore, this paper investigates an algorithm that remains covered by its predecessor's proof of convergence but offers a significant reduction in complexity by trading the number of retained paths versus the DFT length. We motivates this, and also introduce an enhanced initialisation point for the ML sequence estimation. The benefits of this proposed scalable analytic extraction algorithm are illustrated in simulations.**

I. Introduction

Polynomial matrix algebra has found diverse application in broadband sensor array applications. Therefore, in last decade we have seen significant development in field of polynomial matrix algebra especially the polynomial eigenvalue decomposition (PEVD), polynomial singular value decomposition (PSVD) and the polynomial QR decomposition (PQRD). With recent proof of analytic EVD existence, an analytic eigenvalues and eigenvectors extraction algorithms [1], [2] were proposed. Subsequently, the analytic SVD existence proof is formally established where the singular value are conventionally restricted to be real-valued on the unit circle similar to the standard SVD.

The SVD of a polynomial matrix can be computed either through the time-domain approximate algorithms or via the application of two analytic EVDs. Both these methods produce singular values that are complex valued on the unit circle which is opposed to the convention of singular values being real-valued on the unit circle.

Therefore, in this document we propose an algorithm that extracts analytic singular values of a polynomial which are real-valued on the unit circle. In addition, the algorithm does not limits the singular values to be positive on the unit circle.

II. Analytic Eigenvalue Decomposition

A. Existence and Uniqueness

An analytic parahermitian matrix $\boldsymbol{R}(z) : \mathbb{C} \to \mathbb{C}^{M\times M}$, that is not connected to multiplexing operations and hence by paraunitary operations cannot be reduced to pseudo-circulant form [3], admits an analytic EVD [4]–[6] as

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

In (1), $\boldsymbol{Q}(z)$ is a paraunitary matrix containing analytic eigenvectors, and $\boldsymbol{\Lambda}(z)$ is a parahermitian diagonal matrix of analytic eigenvalues. The analytic eigenvectors are ambiguous upto an arbitrary allpass functions i.e. if $\boldsymbol{\Phi}(z)$ is a diagonal matrix of allpass factors, then $\boldsymbol{Q}(z)\boldsymbol{\Phi}(z)$ is also a valid eigenvector matrix. In contrast, the analytic eigenvalues in $\boldsymbol{\Lambda}(z)$ are unique up to a permutation.

III. Analytic Eigenvalue Extraction

This section provides the overview of the original algorithm proposed in [1], [7].

A. EVD in Sample Points of $\boldsymbol{R}(z)$ on the Unit Circle

DFT methods such as [1], [2], [7], [8] compute standard EVDs in the sample points $\Omega_k = 2\pi k/K$ of a K-point DFT of $\mathbf{R}[\tau] \circ\!\!-\!\!\bullet\, \boldsymbol{R}(z)$. This yields [9]

$$\boldsymbol{R}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{R}_k = \mathbf{Q}_k\boldsymbol{\Lambda}_k\mathbf{Q}_k^{\mathrm{H}} \text{ for } k = 0,\ldots,(K-1)\ , \tag{2}$$

where $\boldsymbol{\Lambda}_k = \mathrm{diag}\{d_{1,k}^{(K)},\ldots,d_{M,k}^{(K)}\},\ , m = 1,\ldots,(M-1)$ is a real diagonal matrix with bin-wise eigenvalues, and $\mathbf{Q}_k$ represents the matrix holding in its columns their corresponding eigenvectors in the kth sample point. The eigenvalues in $\boldsymbol{\Lambda}_k$ are unique up to a permutation, which can be fixed by ordering in a majorised fashion such that $d_{m,k}^{(K)} \geq d_{m+1,k}^{(K)}$, $\forall\ m = 1,\ldots,(M=1)$ and $k = 0,\ldots,(K-1)$.

The analytic eigenvalues in (1) evaluated at $z = \mathrm{e}^{\mathrm{j}\Omega_k}$ can be related to the bin-wise eigenvalues $\boldsymbol{\Lambda}_k$ through a permutation matrix

$$\boldsymbol{\Lambda}(\mathrm{e}^{\mathrm{j}\Omega_k}) = \mathbf{P}_k\boldsymbol{\Lambda}_k\mathbf{P}_k^{\mathrm{H}}\ . \tag{3}$$

In [8], this permutation of the eigenvalues is pursued by tracking the orthogonality of the eigenvectors. However, challenges arise at algebraic multiplicities, i.e. in bins where multiple eigenvalues are identical [1], [10], [11]. Also, any perturbation of $\boldsymbol{R}(z)$ will typically affect the eigenvectors more strongly than the eigenvalues [1], [12]. Hence we here focus on the eigenvalues alone.

Example 1: To illustration the bin-dependent permutation matrices $\mathbf{P}_k$, $k = 0,\ldots,(K-1)$, in (3), consider the analytic eigenvalues shown in Fig. 1(a). Their $K = 8$ uniformly spaced sample points on the unit circle are shown in Fig. 1(b). The

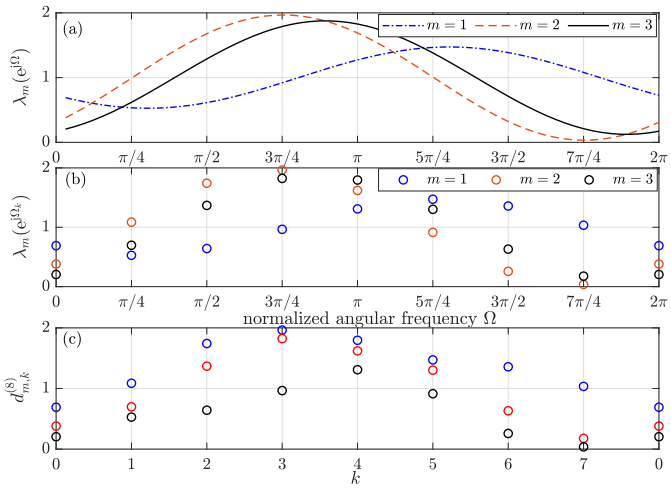

Fig. 1. (a) Analytic eigenvalues evaluated on the unit circle, (b) $K = 8$ sample points of analytic eigenvalues on the unit circle, and (c) spectrally majorised eigenvalues obtained from the binwise EVD of sample points of $\boldsymbol{R}(z)$ on the unit circle. $(d^{(8)}_{m,k}, k = 0, 1, \ldots, 7)$

Fig. 1(c) shows the bin-wise eigenvalues which are ordered in each DFT bin in decreasing order, i.e. they are spectrally majorised [13]. The objective is to permute the bin-wise eigenvalues in each bin via determining $\mathbf{P}_k$ such that the permuted bin-wise eigenvalues associate smoothly across DFT bins as the samples of analytic eigenvalues in Fig. 1(b).

B. Assessing Associations Across Bins

For a spatial dimension M, there are $M!$ possible permutations of arranging the eigenvalues in a single bin. Assuming that the eigenvalues are fixed in the $k = 0$th bin, across all K bins, this leaves $M!^{K-1}$ possible associations to be checked. A suitable metric for smoothness of an association has demonstrated to be the power in the derivatives of a Dirichlet interpolation across the sample point [14], [15]. Minimising this power can be proven to yield the analytic eigenvalues, but unfortunately is unfeasible even for modest values of M and K. Therefore, the proposition in [1] has been to replace an exhaustive search with an iterative ML sequence detection. By starting from the first bin, we

and the goal is to identify the one that results in maximum smoothness of the Dirichlet interpolation performed over the ordered bin-wise eigenvalues [1]. The algorithm works on bin-by-bin basis by appending the Jth bin to the association constructed by permuting previous $(J-1)$ bins denoted as $d^{\{0\to J-1\}}_{m,K}, m = 1, \ldots, M$. Therefore, to identify a possible correct permutation for the Jth bin, it computes the smoothness of an interpolation upto first $J-$bins while assuming the remaining $(K-J)$ samples such that the interpolation across all K samples is maximally smooth [**?**], [**?**], [15]. This accomplished by maximizing

$$\chi^P = \sum_{m=1}^{M} ||\mathbf{R}_2 d^{\{0\to J-1\}}_{m,K}||^2_{\mathrm{F}} \tag{4}$$

where $\mathbf{R}_2 \in \mathbb{J} \times \mathbb{J}$ is the lower-right partition of the upper-right triangular matrix $\mathbf{R}$ obtained through the QR decomposition of $\underline{\mathbf{D}}^{\frac{1}{2}}\mathbf{W}^{\mathrm{H}}$ with

$$\underline{\mathbf{D}} = \begin{cases} \sum_{p=0}^{P} \mathrm{diag}\{0^{2p}, 1^{2p}, ..., L_K^{2p}, L_K^{2p}, ..., 1\}, K \text{ odd} \\ \sum_{p=0}^{P} \mathrm{diag}\{0, 1, ..., (L_K-1), (0.5)^{\frac{1}{2p}} L_K, \\ \qquad (L_K-1), (L_K-2), ..., 1\}^{2p}, K \text{ even}, \end{cases} \tag{5}$$

with P as the maximum derivative order chosen to measure the smoothness of interpolation, and $\mathbf{W}_K$ denotes the $K-$point unitary DFT matrix. The factor $L_K = K/2$ and $(K-1)/2$ for even and odd K, respectively [1].

C. Permutations and Paths

The Jth bin can be appended to the previous $(J-1)$ bins in $M!$ ways, resulting in an equal number of possibilities for associations or paths. For example, considering MK sample points with the first bin set as a reference, there can be $(M!)^{K-1}$ potential associations or paths, with only one being the correct association. However, managing such a vast number of paths becomes computationally impractical, even for smaller values of M. In the remainder of the paper, we refer to incomplete associations as paths. The original algorithm [1] adopts an approach that restricts the displacement of an index at a particular position in permutation generation to a specified maximum limit. For example, if the index displacement distance is restricted to 1 unit, the total number of permutations for $M = 3$ would be $[1, 2, 3]$, $[1, 3, 2]$, $[2, 1, 3]$, where the first position cannot have 3 and the third position cannot expect 1. This approach is applicable due to the bin-wise eigenvalues initially ordered in decreasing order in each bin. The experiment in [1] uses a displacement limit of 3 units, resulting in the total number of permutations per bin varying as $2, 6, 24, 78, 230, 675, 2069, ...$ with $M = 2, 3, 4, 5, 6, 7, 8, ...$ Similarly, the number of paths N_{p} is kept the same as the number of reduced permutations in each bin unless the smoothness metric of any of the paths exceeds the smoothness measure of the majorized path.

D. Iterative Procedure

The algorithm requires two associations at DFT sizes $K/2$ and $K/2$ to determine whether the identified association is correct. It then calculates

$$\xi_1 = \sum_{m=1}^{M} \sum_{k=0}^{K/2-1} |d^{\{2k\}}_{m,K} - d^{\{k\}}_{m,K/2}|^2, \tag{6}$$

and if it is smaller than a certain threshold ϵ (approx. zero), it creates a possibility that the association is correct. However, if the association is correct, the difference between the interpolation through K and $K/2$ sample points should converge to zero for $K \to \infty$ [1] i.e.

$$\Theta_m(e^{j\Omega_)} = d_{(m,K)}(e^{j\Omega_)} - d_{(m,K/2)}(e^{j\Omega_)}, \tag{7}$$

where $d_{m,K}(\mathrm{e}^{\mathrm{j}\Omega})$ is a Dirichlet interpolation over K eigenvalues sample in association $d_{m,K}^{\{0\to K-1\}}$. Hence, the power in its pth derivative is used as a second metric

$$\xi_2^P = \frac{1}{2\pi}\sum_{m=1}^{M}\int_{-\pi}^{\pi}\left|\frac{\mathrm{d}^P}{\mathrm{d}\Omega^P}\Theta_m(\mathrm{e}^{\mathrm{j}\Omega})\right|^2\mathrm{d}\Omega, \tag{8}$$

which, if $\xi_2^P < \epsilon'$, indicates that the association is correct. The time domain sequence $\hat{\lambda}_m[\tau]$ is then obtained via a $K/2-$point IFFT of $d_{m,K/2}^{(0\to K/2-1)}, m = 1,\ldots,M$. On the contrary, if $\xi_1 > \epsilon$, K is doubled, and the process is repeated until the correct association is found, i.e., $(\xi_2^P < \epsilon') \wedge (\xi_1 < \epsilon)$.

E. Scalability with Spatial and Temporal Dimension

This algorithm scale very poorly with spatial and not so well with the temporal dimension $\boldsymbol{R}(z)$. Although, permutations per bin (P_{pb}) are reduced through the maximum displacement approach but for higher spatial dimension, these reduced permutations for large value of M are still considerably high. Additionally, a large number of paths further exponentiate the computational cost, contributing to the complexity growth with respect to the temporal dimension. For instance, if N_{p} paths are kept in the extraction process with the DFT size doubled in each iteration, the total number of permutations in I iterations becomes $2N_{\mathrm{p}}P_{\mathrm{pb}}K\big(1-(\frac{1}{2})^I\big)$ where K is the final DFT size. Since K is dependent upon the temporal dimension of $\mathbf{R}[\tau]$ and the ground-truth eigenvalues, this also prohibits its scalability with temporal dimension.

IV. Modification and Proposed Algorithm

In this section, we propose two important modifications to the DFT-domain approach in [1], which help to extract the analytic eigenvalues from a parahermitian matrix at a significantly reduced cost.

A. Number of Paths versus DFT Size

The ML sequence estimation embedded in the state-of-the-art algorithm in Sec. III uses a number of paths N_{p} to find the correct association across bins, and grows the DFT size K until the association can be established confidently. The role of the DFT size is to allow a sufficiently fine frequency resolution such that the Dirichlet interpolation is close enough to the exact and unique continuous-frequency eigenvalues.

Large numbers of paths often find the correct association within a smaller DFT size. However, if the number of paths is fewer, a larger DFT size may be necessary to identify the correct association. Therefore, a sufficiently large DFT order can render the correct association the smoothest choice among the total $(K-1)M!$ possibilities, even from the starting bins with as few as a single path. This creates the possibility that there may exist a maximum DFT length K_{SP} such that, for $K \geq K_{\mathrm{SP}}$, a single path can easily extract the correct association. It is important to note that K_{SP} will be relative to the temporal and spatial dimensions of the input $\boldsymbol{R}(z)$. Therefore, increasing the DFT length can support fewer paths to extract the correct association. However, indiscriminately increasing the DFT size is not a prudent approach. There are

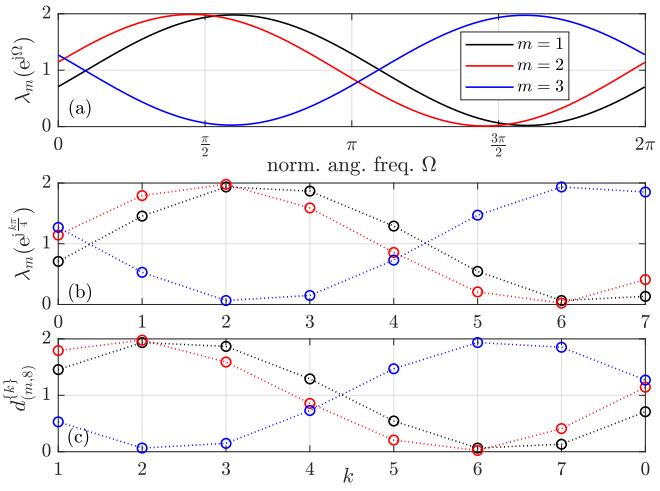

Fig. 2. (a) Analytic eigenvalue example with multiplicity at beginning (b) analytic eigenvalues evaluated at $K = 8$ DFT bins, (c) extracted correct association using single path with $k_{MSEB} = 1$.

cases where a very large DFT length may be required. For example, if there is an algebraic multiplicity near $\Omega = 0$ or in the initial DFT bins, a single path may require a very large DFT order. The same effect is possible for multiple paths with multiple crossings in the starting bins. This effect is illustrated through an example.

Example 2: Lets consider an instance of $\boldsymbol{R}(z) \in \mathbb{C}^{3\times 3}$ with ground truth eigenvalues

$$\lambda_1(z) = \frac{(1.45+\mathrm{j}4.67)z}{10} + 1 + \frac{(1.45-\mathrm{j}4.67)z^{-1}}{10} \tag{9a}$$

$$\lambda_2(z) = \frac{(7.17-\mathrm{j}48.91)z}{100} + 1 + \frac{(7.17-\mathrm{j}48.91)z^{-1}}{100} \tag{9b}$$

$$\lambda_3(z) = \frac{(1.35-\mathrm{j}4.67)z}{10} + 1 + \frac{(1.35+\mathrm{j}4.67)z^{-1}}{10} \tag{9c}$$

illustrated in Fig. 2(a) with $K = 8$ point DFT shown in Fig. 2(b). To extract the correct association with a single path, a DFT size of at least $K = 128$ is required, provided that the DFT size is doubled in each iteration. However, if there are $N_{\mathrm{p}} = 7$ paths, the correct association can be extracted at a DFT size of $K = 8$. This finding suggests that a larger number of paths necessitates a smaller DFT size to achieve the correct association.

B. Maximally Separated Eigenvalues Bin

The Example 2 requires $K = 128$ DFT bins to find the correct association with a single path. The necessity for such a large DFT length is due to the presence of non-trivial algebraic multiplicity in the early bins. If this issue is addressed, it is possible that a smaller number of paths could be used to extract the correct association with a smaller DFT size. For this, we determine the bin where the eigenvalues are maximally separated and assume it as the starting bin for every path. We can do so due to eigenvalues being $2\pi-$periodic, and so in sample points with K as its

period. The sequence in which bins will be visited becomes $[k_{\text{MSEB}}, k_{\text{MSEB}}+1, ..., K-1, 0, 1, ..., k_{\text{MSEB}}-1]$ where k_{MSEB} denotes maximally separated bin. The pseudo-code for finding k_{MSEB} is outlined in Algorithm 1 for total of $K-$bins. This search operation can be performed either K bins or its fraction **round**$(\eta K), 0 < \eta \leq 1$. The search is repeated every time K is varied. By designating k_{MSEB} as the first bin, it facilitates reduced number of paths to capture the appropriate associations in the initial bins, which in turn leads to a comparatively straightforward addition of the subsequent bins.

By employing the maximally separated bin technique, it is possible to extract the correct association with the use of a single path in 8 DFT bins, as demonstrated in Fig. 2 (c), where $k_{\text{MSEB}} = 1$. This finding indicates that, in conjunction with a sufficiently large K, the maximally separated bin approach enables the method of [1] to decompose a higher-dimensional para-Hermitian matrix using reduced number of paths.

Algorithm 1: : Search for k_{MSEB} for Approach-I

```
for k = 0 : K − 1 do
    for m = 1 : M − 1 do
        Δd_m = |d^(K)_{m,k} − d^(K)_{m+1,k}|;
    end
    d_k = min{Δd_1, Δd_2, ..., Δd_{M−1}};
end
k_MSEB = argmax_k {d_k}
```

C. Stopping Criterion

The original method requires computation of two smoothest associations with $K/2$ and $K-$point DFT to calculate ξ_n, where $n = 1, 2$. These associations must be stored separately as $d_{m,K}^{\{0\to K-1\}}$ and $d_{m,2K}^{\{0\to 2K-1\}}$. By examining ξ_1, it is determined whether these two associations are the same. If these are same, ξ_2^P assesses the time domain aliasing to verify the sufficiency of the DFT length for the extraction of analytic eigenvalues. Alternatively, according to Theorem 4 in [1], the convergence of ξ_2^P confirms the correctness of the extracted association.

Given the reliance on ξ_2^p for convergence as per Section IV-C in [1], it's apparent that ξ_1 becomes unnecessary. The condition $\xi_2^P \to 0$ suffices as both a necessary and sufficient criterion for confirming association correctness, serving as a suitable stopping point for the algorithm. Discarding ξ_1 eliminates the need for keeping the previous association constructed at $K/2$ DFT size. Since ξ_2^P requires two association, the association at K DFT size is obtained from the current iteration while for $K/2$ DFT size can be derived from it using even-indexed samples. This enables the computation of $\Theta_m(\mathrm{e}^{\mathrm{j}\Omega})$ for evaluating ξ_2^P via (8) with (7).

D. Convergence

The proposed method maintains reliance on the convergence criterion of ξ_2^p, thereby preserving the validity of the original convergence proof. However, the reduction in the number of paths might impact the algorithm's convergence towards the correct association. To address this, the iterative increment of K contributes to smoothing the eigenvalues, thereby supporting convergence with a reduced number of paths towards the desired solution. Additionally, the concept of maximally separated bins aids in capturing the association's nature in the initial bins, facilitating the addition of later bins with minimal difficulty.

V. Simulation and Results

A. Ensemble Test for Various Paths

An ensemble of 100 instantiations of $\boldsymbol{R}(z) \in \mathbb{C}^{4\times 4}$ is created and analytic eigenvalues are extracted using various number of paths $N_{\text{p}} = \{2, 4, 6, 10\}$. The execution time averaged over the whole ensemble which is given in Table I. It can be seen that for the constructed ensemble, reduced number of paths perform computationally better compared to large number of paths. However, it must be reminded that an optimal number of paths are difficult to determine.

TABLE I
Execution Time Using Different Number Paths

$N_{\text{p}} = 2$	$N_{\text{p}} = 4$	$N_{\text{p}} = 6$	$N_{\text{p}} = 10$
$0.10 \pm 0.07\ s$	$0.12 \pm 0.05\ s$	$0.14 \pm 0.07\ s$	$0.17 \pm 0.11\ s$

B. Comparison With Original Method

The proposed modified approach is compared against the original method in [**?**] through an ensemble comprising 500 $\boldsymbol{R}(z)$ instantiations of spatial dimension $M = 4$ for every order of $\text{Ord}\{\lambda_m(z)\} = \{10, 12, \ldots, 30\}$. These instantiations are constructed through the source model in [**?**], [**?**]. Within the maximum DFT limit of 2^{10}, the proposed modified method couldn't extract eigenvalues through a single path for all instantiations within the constructed ensemble. However, with two paths $N_{\text{p}} = 2$, all instances were easily decomposed resulting in $\xi_2^{p=3} < 10^{-18}$. The extracted eigenvalues matched exactly with the sample points of ground truth eigenvalues and the resulting time domain support was equal to that of the ground-truth as already proved in Fig. 9 of [1] and so we don't re-illustrate it here.

In Figure 3, the execution time for $M = 4$ is plotted against $\text{Ord}\{\lambda_m(z)\}$. Notably, the modified approach exhibits significantly lower execution times for both $N_{\text{p}} = 2$ and $N_{\text{p}} = 4$ compared to prior research [1]. To further illustrate this improvement, we conducted ensemble experiments for higher spatial dimensions, $M \in 5, 6, 7, 8$, and the results are presented in Figure 4. Remarkably, the reduced paths method successfully extracted analytic eigenvalues using only 2 paths within the maximum available DFT limit. Specifically, the proposed modified method extracts eigenvalues of $\mathbb{C}^{6\times 6}$ faster than the original method in [1] extracts that of the $\mathbb{C}^{4\times 4}$ when temporal dimensions are kept same. This suggests the potential to process data from two additional sensors within the same time frame in a broadband array. It must reminded that the

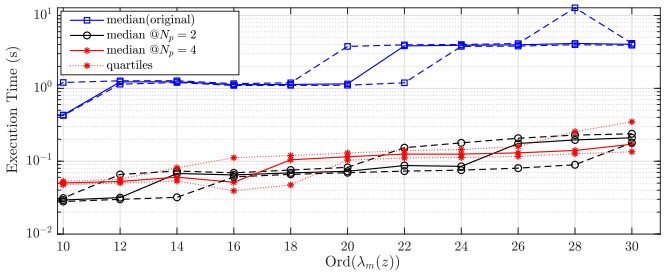

Fig. 3. Execution time of the proposed parallel approach compared to [**?**] over an ensemble containing $\boldsymbol{R}(z) \in \mathbb{C}^{4\text{x}4}$ instantiations.

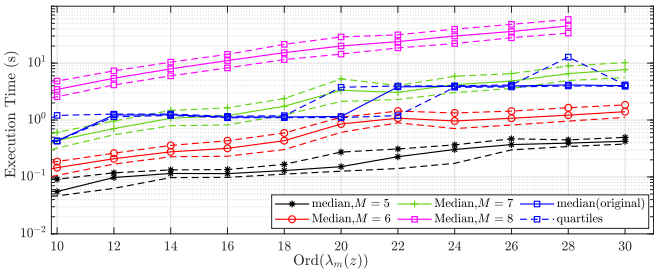

Fig. 4. Execution time of the proposed parallel approach with $N_\text{p} = 2$ over an ensemble containing $\boldsymbol{R}(z) \in \mathbb{C}^{M\text{x}M}, M \in \{5, 6, 7, 8\}$ instantiations.

original method is not simulated for spatial dimensions of $M = \{5, 6, 7, 8\}$ due to very high execution time.

VI. Conclusion

This paper has put forward suggestions for altering the analytic eigenvalue extraction algorithm detailed in [1] in order to decrease its computational expense for diagonalizing higher spatial dimension para-Hermitian matrices. Specifically, the revised method relies on larger DFT sizes to reduce the number of paths required for the extraction of eigenvalues. Additionally, the extraction process is initiated from a bin where the eigenvalues are maximally separated in order to further support the reduced number of paths in succeeding in the extraction. The effectiveness of the proposed modified method is demonstrated by its significant outperformance compared to the original method in terms of computational cost.

References

[1] 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, 2021.

[2] S. Weiss, I. K. Proudler, G. Barbarino, J. Pestana, and J. G. McWhirter, "On properties and structure of the analytic singular value decomposition," *IEEE Transactions on Signal Processing*, 2023, submitted.

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

[4] 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.

[5] 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.

[6] 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.

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

[8] 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. [Online]. Available: http://dx.doi.org/10.1186/1687-6180-2013-93

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

[10] F. Coutts, K. Thompson, S. Weiss, and I. Proudler, "A comparison of iterative and DFT-based polynomial matrix eigenvalue decompositions," in *IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, Curacao, December 2017.

[11] F. K. Coutts, K. Thompson, I. K. Proudler, and S. Weiss, "An iterative dft-based approach to the polynomial matrix eigenvalue decomposition," in *52nd Asilomar Conference on Signals, Systems, and Computers*, Oct 2018, pp. 1011–1015.

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

[13] P. Vaidyanathan, "Theory of optimal orthonormal subband coders," *IEEE Transactions on Signal Processing*, vol. 46, no. 6, pp. 1528–1543, Jun. 1998.

[14] S. Weiss and M. D. Macleod, "Maximally smooth dirichlet interpolation from complete and incomplete sample points on the unit circle," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Brighton, UK, May 2019.

[15] S. Weiss, J. Selva, and M. D. Macleod, "Measuring smoothness of trigonometric interpolation through incomplete sample points," in *28th European Signal Processing Conference*, Amsterdam, Netherlands, Jan. 2021, pp. 2319–2323.