

Fast Givens Rotation Approach to Second Order Sequential Best Rotation Algorithms

Faizan Khattak, Stephan Weiss, Ian K. Proudler

Department of Electronic & Electrical Engineering, University of Strathclyde, Glasgow G1 1XW, Scotland

{faizan.khattak,stephan.weiss,ian.proudler}@strath.ac.uk

Abstract—The second order sequential best rotation (SBR2) algorithm is a popular algorithm to decompose a parahermitian matrix into approximated polynomial eigenvalues and eigenvectors. The work horse behind SBR2 is a Givens rotation interspersed by delay operations. In this paper, we investigate and analyse the application of a fast Givens rotation in order to reduce the computation complexity of SBR2. The proposed algorithm inherits the SBR2’s proven convergence to a diagonalised and spectrally majorised solution for the polynomial eigenvalues. We provide some analysis and examples for the execution speed of this fast Givens-based SBR2 compared to a standard SBR2 implementation.

I. INTRODUCTION

For broadband signals $\mathbf{x}[n] \in \mathbb{C}^M$ acquired by an M -element sensor arrays in discrete time $n \in \mathbb{Z}$, the space-time covariance matrix $\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^H[n-\tau]\}$ captures the complete second order statistics of the data. The explicit lag parameter $\tau \in \mathbb{Z}$ is capable of characterising time delays between the M different sensor elements and hence bears information on aspects such the angle of arrival of a particular source signal. This is different from the narrowband case, where e.g. AoA is resolved simply by phase shifts, and it suffices to consider the instantaneous covariance $\mathbf{R}[0]$. To generalise narrowband optimal solutions, which often are based on the eigenvalue decomposition of $\mathbf{R}[0]$, to the broadband case, techniques have been developed to diagonalise $\mathbf{R}[\tau]$ for every value of τ . Because its z -transform $\mathbf{R}(z) = \sum_{\tau} \mathbf{R}[\tau]z^{-\tau}$ is a polynomial matrix, such techniques are referred to as polynomial matrix EVDs [1], [2].

A polynomial matrix EVD exists in the case of an analytic $\mathbf{R}(z)$ that emerges from unmultiplexed data [2], [3], and two main families of algorithms with proven convergence have emerged over the last decade — sequential matrix diagonalisation (SMD, [4], [5]) and the second order sequential best rotation algorithm (SBR2, [1], [6]). Since applications such as for subband coding [7], beamforming [8], source separation [9] or speech enhancement [10] depend on low computation cost, various efforts have been directed at numerical [11]–[15] and implementational enhancements [15], [16].

Amongst numerical enhancements of polynomial matrix EVD algorithms, particularly the Givens rotation has attracted

attention. The cyclic-by-row approach in [11] limited the number of rotation steps to only implement an approximate EVD within the SMD family of algorithms. Divide-and-conquer schemes [14], [15] have been aiming at reducing the spatial dimension by breaking $\mathbf{R}[\tau]$ into sub-blocks, and hence reducing the cost of matrix-multiplications.

In this paper, we employ the idea of fast Givens rotations. Typically a similarity transform by a Givens rotation requires the modification of each two rows and two columns of the matrix to be transformer. It is well known that square root and division free approaches [17]–[19] lead to simplifications and effectively split the unitary Givens rotation into a simple diagonal matrix and a simplified matrix that will only require have the multiplications and additions of a Givens operation. In [20], this simplification is extended of successive real-valued Givens rotations and is beneficially applied to EVD and QR calculations. Here, we employ a complex-valued extension of this approach in the context of the SBR2 algorithm, so reduce the SBR2’s complexity.

Therefore, in the following, Sec. II will first briefly outline the SBR2 algorithm, followed by the introduction of fast Givens rotations — both as single operations and in sequence — in Sec. III. This approach then drives a fast SBRs version introduced in Sec. IV, which is evaluated in a numerical example and simulations in Sec. V. Finally, conclusions are drawn in Sec. VI.

II. POLYNOMIAL EVD ALGORITHMS

A. Polynomial EVD

The CSD matrix $\mathbf{R}(z)$ admits a parahermitian matrix EVD (PhEVD) $\mathbf{R}(z) = \mathbf{Q}(z)\mathbf{\Lambda}(z)\mathbf{Q}^P(z)$, if $\mathbf{R}(z)$ is analytic and the measurement vector $\mathbf{x}[n]$ does emerge from multiplexed data [2], [3]. In this case, the factor $\mathbf{U}(z)$ is paraunitary, i.e. $\mathbf{Q}(z)\mathbf{Q}^P(z) = \mathbf{I}$, whereby $\mathbf{U}^P(z) = \mathbf{U}^H(1/z^*)$, and contains analytic eigenvectors in its columns. The corresponding analytic eigenvalues form the diagonal parahermitian matrix $\mathbf{\Lambda}(z)$.

A polynomial eigenvalue decomposition [1] is modified version of the PhEVD,

$$\mathbf{R}(z) \approx \mathbf{U}(z)\mathbf{\Gamma}(z)\mathbf{U}^P(z), \quad (1)$$

where the potentially transcendental factors $\mathbf{Q}(z)$ and $\mathbf{\Lambda}(z)$ of the PhEVD are replaced by a polynomial paraunitary $\mathbf{U}(z)$ and a Laurent polynomial, diagonal, and parahermitian $\mathbf{\Gamma}(z)$. Specifically, the analytic $\mathbf{\Lambda}(z)$ is replaced by a spectrally

We gratefully acknowledge support through Commonwealth Scholarship PKCS-2020-518. This work was also supported in parts by the Engineering and Physical Sciences Research Council (EPSRC) Grant number EP/S000631/1 and the MOD University Defence Research Collaboration in Signal Processing.

majorised $\mathbf{\Gamma}(z) = \text{diag}\{\gamma_1(z), \dots, \gamma_M(z)\}$, such that on the unit circle

$$\gamma_{m+1}(e^{j\Omega}) \geq \gamma_m(e^{j\Omega}) \quad \forall \Omega \in \mathbb{R}, \quad m = 1, \dots, (M-1). \quad (2)$$

This spectral majorisation yields eigenvalues that are different from the PhEVD if the analytic eigenvalues intersect. In this case, the polynomial EVD in (1) will require a high polynomial order for a satisfactory approximation, but is important in applications such as subband coding, where the majorisation property in (2) guarantees the maximisation of the coding gain [6], [21].

The PEVD is targeted by a number of algorithms with guaranteed convergence. A family of second order sequential best rotation (SBR2) algorithms [1], [6], [22] aims to iteratively eliminate the largest off-diagonal component at every iteration step, whereby spectral majorisation has been shown to be enforced [23]. More recently, sequential matrix diagonalisation (SMD) algorithms [4], [5] iteratively diagonalises a time-shifted version of the space-time covariance, which converges faster per iteration but at an overall higher computational cost than SBR2.

B. Second Order Sequential Best Rotation Algorithm

SBR2 is a generalisation of the Jacobi method, whereby successive Givens rotations are interspersed by delay operations in order to move maximum off-diagonal elements to the lag zero component, where they then are eliminated.

Starting from $\mathbf{S}_0(z) = \mathbf{R}(z)$, at the i th iteration SBR2 finds the maximum off-diagonal element

$$\{\tau_i, m_i, n_i\} = \arg \max_{\substack{\tau, m, n \\ m \neq n}} |s_{i-1, m, n}[\tau]|, \quad (3)$$

where $s_{i-1, m, n}[\tau]$ is the element in the m th row and n th column of $\mathbf{S}_{i-1}[\tau]$, whereby $\mathbf{S}_{i-1}[\tau] \circ \bullet \mathbf{S}_{i-1}(z)$. A delay (or advance, for $\tau_i < 0$) operation

$$\mathbf{\Delta}_i(z) = \text{diag}\left\{\underbrace{1, \dots, 1}_{n_i-1}, z^{-\tau_i}, \underbrace{1, \dots, 1}_{M-n_i}\right\} \quad (4)$$

then ensures that $s_{i-1, m_i, n_i}[\tau_i]$ as well as $s_{i-1, n_i, m_i}[-\tau_i] = s_{i-1, m_i, n_i}^*[\tau_i]$ are moved to lag zero by the similarity transform $\mathbf{S}_{i-\frac{1}{2}}(z) = \mathbf{\Delta}_i(z)\mathbf{S}_{i-1}(z)\mathbf{\Delta}_i^H(z)$.

The energy of the $s_{i-1, m_i, n_i}[\tau_i]$ and $s_{i-1, n_i, m_i}[-\tau_i]$ is then transferred onto the diagonal via a Givens rotation \mathbf{V}_i ,

$$\mathbf{V}_i = \begin{bmatrix} \mathbf{I}_1 & & & & & & & & \\ & c_i & & & e^{j\gamma_i s_i} & & & & \\ & & & \mathbf{I}_2 & & & & & \\ & -e^{-j\gamma_i s_i} & & & c_i & & & & \\ & & & & & & & & \mathbf{I}_3 \end{bmatrix}, \quad (5)$$

with elements $c_i = \cos(\theta_i)$ and $s_i = \sin(\theta_i)$ only at the intersection of the m_i th and n_i th rows and columns, and \mathbf{I}_i , $i = 1, 2, 3$, identity matrices of appropriate dimensions. The values of θ_i and γ_i depend on $\mathbf{S}_{i-\frac{1}{2}}[0]$, with $\mathbf{S}_{i-\frac{1}{2}}[\tau] \circ \bullet \mathbf{S}_{i-\frac{1}{2}}(z)$ [1], [4]. When completing the i th SBR2 iteration step, therefore

$$\mathbf{S}_i(z) = \mathbf{V}_i \mathbf{S}_{i-\frac{1}{2}}(z) \mathbf{V}_i^H. \quad (6)$$

If after I iterations, the off-diagonal components are suppressed below some given threshold or a maximum number of iterations has been reached, w.r.t. (1) we obtain $\mathbf{\Gamma}(z) = \mathbf{S}_I(z)$ and

$$\mathbf{U}(z) = \mathbf{V}_I \mathbf{\Delta}_I(z) \cdots \mathbf{V}_2 \mathbf{\Delta}_2(z) \cdot \mathbf{V}_1 \mathbf{\Delta}_1(z). \quad (7)$$

Note that in every iteration step, while energy at lag zero is transferred into the diagonal, Givens rotations are also applied at lags $\tau \neq 0$, where some energy may leak back into the off-diagonal components. Nonetheless, SRB2 has been proven to converge [1], [4], [23]. Since the Givens operation is the central algebraic operation in SBR2, in the following we will investigate a fast version derived for the real-valued case in [20].

III. FAST GIVENS ROTATION

A. Givens Rotation

Within SBR2, the Givens rotation is applied in (6), which alternatively in the time domain requires

$$\mathbf{S}_i[\tau] = \mathbf{V}_i^H \mathbf{S}_{i-\frac{1}{2}}[\tau] \mathbf{V}_i. \quad (8)$$

If the support of $\mathbf{S}_{i-\frac{1}{2}}[\tau]$ is restricted to $|\tau_i| \leq T_i$, such that $\mathbf{S}_{i-\frac{1}{2}}[\tau] = \mathbf{0} \quad \forall |\tau| > T_i$, then we require to perform $2T_i + 1$ such similarity transforms. Due to the parahermitian nature of $\mathbf{S}_{i-\frac{1}{2}}(z)$, T of these transforms will be redundant. With the exception of the lag zero matrix $\mathbf{S}_{i-\frac{1}{2}}[0]$, the remaining coefficients $\mathbf{S}_{i-\frac{1}{2}}[\tau]$ will generally not be Hermitian matrices for $\tau \neq 0$. However, it is only the Hermitian zero-lag matrix $\mathbf{S}_{i-\frac{1}{2}}[0]$ where a pair of off-diagonal components must be eliminated.

If initially we focus on a single operation out of the above $T + 1$ similarity transforms for (8) and drop all subscripts, we face typically the problem $\mathbf{S}' = \mathbf{V}^H \mathbf{S} \mathbf{V}$. The similarity transform based on a Givens rotation will modify the m th and the n th rows and columns of \mathbf{S} to obtain \mathbf{S}' , and therefore typically require $4M$ complex-valued multiply accumulate (MAC) operations, which typically translates to $16M$ real-valued MACs.

B. Fast Approach

For a fast implementation of the Givens rotation, termed fast Givens rotation (FGR), [20] contains two interesting ideas. The first is a reduction of computations for a single Givens rotation; a second step, which will be elaborated in Sec. III-C, exploits subsequent savings if several Givens rotations are iterated. The basic idea is to factorise the orthogonal Givens operation into two non-orthogonal ones, where as many elements as possible are set to unity, hence not requiring explicit multiplications.

We first consider a Givens rotation that is applied to a square matrix \mathbf{B}_0 such that $\mathbf{B}_1 = \mathbf{V}_1^H \mathbf{B}_0 \mathbf{V}_1$, with \mathbf{V}_1 a Givens rotation matrix as in (5). Then, according to [20] it is possible to find diagonal matrices \mathbf{D}_0 and \mathbf{D}_1 , such that $\mathbf{B}_0 = \mathbf{D}_0 \mathbf{A}_0 \mathbf{D}_0$, and $\mathbf{B}_1 = \mathbf{D}_1 \mathbf{F}_1 \mathbf{A}_0 \mathbf{F}_1^H \mathbf{D}_1$. One particular and numerically stable choice according to [20] is such that

\mathbf{D}_1 and \mathbf{D}_0 only differ in the m th and n th elements, leading to

$$\mathbf{F}_1 = \begin{bmatrix} \mathbf{I}_1 & & & & & \\ & 1 & & f_{1,1} & & \\ & & \mathbf{I}_2 & & & \\ & f_{1,2} & & 1 & & \\ & & & & \mathbf{I}_3 & \\ & & & & & \end{bmatrix}. \quad (9)$$

The angles ϑ and γ would be calculated from $\mathbf{A}_0 = \mathbf{D}_0^{-1}\mathbf{B}_0\mathbf{D}_0^{-1}$ in the standard Givens fashion, whereby [20] advocates the selection

$$p = \frac{|d_m^2 a_{mm}^2 - d_n^2 a_{nn}^2|}{\sqrt{(d_m^2 a_{mm}^2 - d_n^2 a_{nn}^2)^2 - 4d_m^2 d_n^2 |a_{mn}|^2}}, \quad (10)$$

with $0 \leq p \leq 1$ and a_{mn} the element in the m th row and n th column of \mathbf{A}_0 , such that e.g.

$$\cos \vartheta = \sqrt{\frac{1+p}{2}}. \quad (11)$$

Overall, we may select $\mathbf{D}_0 = \mathbf{I}$. Then for \mathbf{F}_1 to take the form in (9), we require that \mathbf{D}_1 matches \mathbf{D}_0 , but that the m th and n th elements are modified to $d_m \cos \vartheta$ and $d_n \cos \vartheta$, respectively. For the complex-valued case, it can be shown that $\gamma = \angle a_{mn}$ [1]. Further, we have

$$f_{1,1} = e^{j\gamma} \frac{d_n \sin \vartheta}{d_m \cos \vartheta}, \quad f_{1,2} = -e^{-j\gamma} \frac{d_m \sin \vartheta}{d_n \cos \vartheta} \quad (12)$$

for the simplified matrix \mathbf{F}_1 .

The application of \mathbf{F}_1 requires only half the number of MACs compared to \mathbf{V}_1 , but the above approach involves some overheads since $\mathbf{V}_1 = \mathbf{D}_1 \mathbf{F}_1 \mathbf{D}_0^{-1}$, even if $\mathbf{D}_0 = \mathbf{I}$. Nonetheless, computational saving arise [20], and the technique becomes even more powerful in case several Givens rotations need to be executed successively, such as part of an EVD or QR decomposition [24].

C. Successive Application of Fast Givens Rotations

If a number of Givens rotations \mathbf{V}_k , $k = 1, \dots, K$, need to be executed successively, then note from above that [20]

$$\mathbf{V}_K \dots \mathbf{V}_2 \mathbf{V}_1 = \mathbf{D}_K \mathbf{F}_K \dots \mathbf{F}_2 \mathbf{F}_1 \mathbf{D}_0^{-1}. \quad (13)$$

Particularly with $\mathbf{D}_0 = \mathbf{I}$, the evaluation is simple, and the only overhead is to track the modifications of \mathbf{D}_k , $k = 1, \dots, K$ which requires a multiplication of the m th and n th element, as indicated above.

IV. FAST GIVENS ROTATION-BASED SBR2

A. Modified Fast SBR2 Algorithm

Even though the SBR2 algorithm does not comprise of a simple consecutive application of Givens rotations, we can apply the idea of (13). This is due to the fact that the interspersed delay operations $\mathbf{D}(z)$ in (7) are diagonal matrices, and hence substituting $\mathbf{V}_\ell = \mathbf{D}_\ell \mathbf{F}_\ell \mathbf{D}_{\ell-1}^{-1}$, $\ell = 1, \dots, I$, into (7) leads to

$$\mathbf{U}(z) = \mathbf{D}_I \mathbf{F}_I \mathbf{D}_{I-1}^{-1} \Delta_I(z) \dots \mathbf{D}_1 \mathbf{F}_1 \mathbf{D}_0^{-1} \Delta_1(z) \quad (14)$$

$$= \mathbf{D}_I \mathbf{F}_I \Delta_I(z) \dots \mathbf{F}_2 \Delta_2(z) \mathbf{F}_1 \Delta_1(z) \mathbf{D}_0^{-1}. \quad (15)$$

Algorithm 1: Fast Givens Rotation-Based SBR2

-
- 1: inputs: $\mathbf{R}(z)$, δ_{\max} , I_{\max}
 - 2: initialise: $\mathbf{S}'_0(z) = \mathbf{R}(z)$; $\mathbf{D}_0 = \mathbf{I}$; $\mathbf{U}'_0 = \mathbf{I}$; $i = 0$;
 - 3: **repeat**
 - 4: $i \leftarrow i + 1$;
 - 5: find maximum off-diagonal element δ of $\mathbf{D}_{i-1} \mathbf{S}'_{i-1}(z) \mathbf{D}_{i-1}^H$ via (3);
 - 6: determine $\Delta_i(z)$;
 - 7: calculate $\mathbf{S}'_{i-\frac{1}{2}}(z) = \Delta_i(z) \mathbf{S}'_{i-1}(z) \Delta_i^H(z)$;
 - 8: determine $f_{i,1}$ and $f_{i,2}$ based on $\mathbf{S}'_{i-\frac{1}{2}}[0]$;
 - 9: calculate $\mathbf{S}'_i(z) = \mathbf{F}_i \mathbf{S}'_{i-\frac{1}{2}}(z) \mathbf{F}_i^H$;
 - 10: calculate $\mathbf{U}'_i(z) = \mathbf{F}_i \Delta_i(z) \mathbf{U}'_{i-1}(z)$;
 - 11: update \mathbf{D}_i based on \mathbf{D}_{i-1} and $\mathbf{S}'_{i-\frac{1}{2}}[0]$;
 - 12: **until** $(|\delta| < \delta_{\max}) \vee (i \geq I_{\max})$.
 - 13: outputs: $\mathbf{U}(z) = \mathbf{D}_i \mathbf{U}'_i(z)$ and $\Gamma(z) = \mathbf{D}_i \mathbf{S}'_i(z) \mathbf{D}_i^H$
-

Therefore, the overall operation now consists of a sequence of delay operations $\Delta_\ell(z)$ and simplified matrix operations \mathbf{F}_ℓ , whereby a multiplication with \mathbf{F}_ℓ only requires approximately half the MACs as a matrix multiplication with \mathbf{V}_ℓ . Some book keeping is required to update \mathbf{D}_0 , \mathbf{D}_1 all the way up to \mathbf{D}_I , but these changes only ever affect two components of these diagonal quantities at a time.

This results in the modification of the SBR2 algorithm to its fast Givens rotation-based version, with its algorithmic steps outlined in Algorithm 1. The modifications are reflected in the variables $\mathbf{S}'_i(z)$ and $\mathbf{U}'_i(z)$, which are different from $\mathbf{S}_i(z)$ and $\mathbf{U}_i(z)$ in the standard SBR2 algorithm. The steps 8–11 and 13, to determine, apply and propagate the diagonal correction matrix, apply the reduced-cost matrix \mathbf{F}_ℓ , and to finally correct the output also differ from the standard SBR2 method.

B. Convergence

For the convergence of Algorithm 1, the essential detail is that the maximum search in step 5 within $\mathbf{D}_{i-1} \mathbf{S}'_{i-1}(z) \mathbf{D}_{i-1}^H = \mathbf{S}_{i-1}(z)$ is performed over the same quantity as in the SBR2 algorithm in (3). Therefore the convergence proof of the SBR2 algorithm in [1], [6] holds equally for the fast Givens rotation-based SBR2 version. This also implies that Algorithm 1 is guaranteed to converge to a spectrally majorised solution as defined in (2) [23].

Note that it is not necessary to multiply out $\mathbf{D}_{i-1} \mathbf{S}'_{i-1}(z) \mathbf{D}_{i-1}^H$ in Step 5 of Algorithm 1 explicitly; a maximum modulus search can first be performed over the temporal dimension of $\mathbf{S}'_{i-1}[\tau]$, and the matrix-valued result $\mathbf{A}_{\max, \tau}$ can be weighted by $\mathbf{D}_{i-1} \mathbf{A}_{\max, \tau} \mathbf{D}_{i-1}^H$, which due to the Hermitian nature of $\mathbf{A}_{\max, \tau}$ and the exclusion of diagonal terms only takes $\frac{1}{2}M(M-1)$ MACs. The remainder of the maximum search can then be performed over either the lower left or upper right triangular part of $\mathbf{D}_{i-1} \mathbf{A}_{\max, \tau} \mathbf{D}_{i-1}^H$.




Fig. 1. Plots showing $\mathbf{R}[\tau] \in \mathbb{R}^{3 \times 3}$ used as a numerical example; each subplot represents one polynomial entry of $\mathbf{R}[\tau]$.

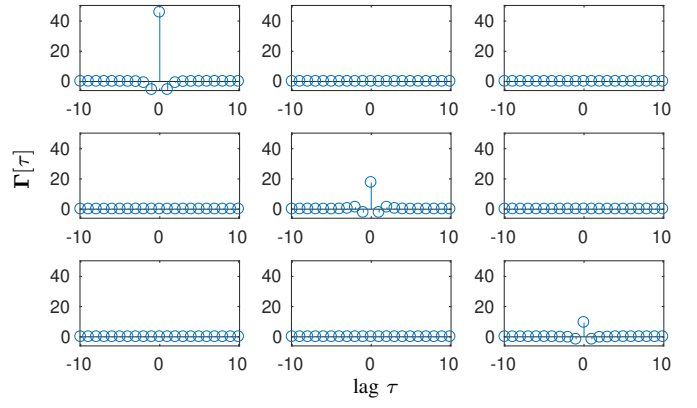


Fig. 2. Approximately diagonalised matrix $\mathbf{\Gamma}[\tau]$, obtained from $\mathbf{R}[\tau]$ in Fig. 1 by the fast Givens rotation-based SBR2 algorithm.

V. SIMULATION AND RESULTS

A. Numerical Example and Convergence

As a numerical example, we utilise the matrix $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{3 \times 3}$ of order 2, such that $\mathbf{R}(z) = \mathbf{R}[-1]z + \mathbf{R}[-0] + \mathbf{R}[1]z^{-1}$. For the matrix-valued coefficients, we have

$$\mathbf{R}[0] = \begin{bmatrix} 14.7 & 0.3 & 2.2 \\ 0.3 & 19.1 & -8.0 \\ 2.2 & -8.0 & 39.0 \end{bmatrix} \quad (16)$$

$$\mathbf{R}[1] = \begin{bmatrix} 0.3 & 4.6 & -7.3 \\ -0.4 & -6.0 & -1.1 \\ 2.5 & 5.9 & -3.7 \end{bmatrix}. \quad (17)$$

For the third coefficient, note that $\mathbf{R}[-1] = \mathbf{R}^H[1]$. This matrix is also characterised in Fig. 1.

Operating on the above $\mathbf{R}(z)$, the fast Givens rotation-based SBR2 algorithm converges in $I = 238$ iterations, and for a maximum off-diagonal component threshold $\delta_{\max} = 1e - 5$ yields the eigenvalues shown in Fig. 2. The extracted eigenvalues $\mathbf{\Gamma}[\tau] \circ \bullet \mathbf{\Gamma}(z)$ are Laurent polynomial decay in both positive and negative lag directions, and Fig. 2 only provides values for the central lags $|\tau| \leq 10$. The spectral majorisation of the eigenvalues as defined in (2), and to which according to [23] and Sec. IV-B the fast Givens rotation-based SBR2 algorithm is guaranteed to converge, is demonstrated in Fig. 3. These power spectral density terms are non-negative real and ordered in power, satisfying (2).

In comparison, for the above $\mathbf{R}(z)$, the standard SBR2 algorithm [1], [6] converges in $I = 235$ iterations. The obtained eigenvalues $\mathbf{\Gamma}_{\text{SBR2}}(z)$ are near-identical to those of the proposed algorithm in Figs. 2 and 3, with an error $\sum_{\tau} \|\mathbf{\Gamma}_{\text{SBR2}}[\tau] - \mathbf{\Gamma}[\tau]\|_{\text{F}}^2$ of -96.9dB , whereby $\|\cdot\|_{\text{F}}$ is the Frobenius norm. The small difference in the number of iterations and in the obtained eigenvalues—which according to [2], [3] are unique—is likely due to the numerical differences between the two algorithm versions, even though the fast Givens approach is claimed to be robust to over- and underflow [20]. The extracted eigenvalues differ by a somewhat larger amount due to their ambiguity w.r.t. multiplications




Fig. 3. Eigenvalues $\mathbf{\Gamma}(z)$ as obtained by the fast Givens rotation-based SBR2 algorithm evaluated on the unit circle, $z = e^{j\Omega}$.

by arbitrary allpass functions [2], [12], and are therefore not shown here.

B. Computational Complexity

In a standard Matlab implementation, where matrix-valued operations are favoured and often there is no relation between the sparseness of a matrix and the execution time for its multiplication, the run time difference (averaged over 500 runs) for the example in Sec. V-A is only 98.1 ms for the standard SBR2 and 92.6 ms for the proposed fast Givens rotation-based SBR2 version. To investigate the potential of the fast Givens approach for speeding up an implementation in a non-Matlab environment, the fast and standard Givens rotation operations have been explicitly implemented in C and operated from within Matlab through pre-compiled MEX files.

Using Matlab's profiler, the MEX-routines are specifically called to perform $I = 150$ iterations on a matrix $\mathbf{R}(z)$ calculated from a ground truth with a diagonal $\mathbf{\Gamma}(z)$ of order 100 and an arbitrary paraunitary matrix $\mathbf{U}(z)$ determined from randomised elementary paraunitary operations [25] of order 50. The execution time averaged over 5000 runs for different spatial dimensions M is summarised in Tab. I. The execution includes memory allocation, maximum searches and various

TABLE I

EXECUTION TIME COMPARISON BETWEEN STANDARD AND FAST GIVENS ROTATION-BASED SBR2 IMPLEMENTATIONS.

method	computation time / [ms]			
	$M = 3$	$M = 4$	$M = 5$	$M = 6$
standard Givens	1.6	2.2	2.9	4.0
fast Givens	1.5	2.0	2.5	3.5

book keeping, and therefore is not as dramatic a reduction as direct comparison in MAC operations might suggest. Nevertheless, as M increases, a substantial gap between the run times of the standard and proposed SBR2 implementations emerges, with a reduction by 1/8 for the largest matrix size of $M = 6$ that is employed here.

VI. CONCLUSION

This paper has exploited a fast Givens rotation trick to reduce the number of multiply-accumulate operations, particularly when operating on a sequence of Givens rotations. This trick has been adopted for a polynomial matrix eigenvalue decomposition technique known as the second order sequential best rotation algorithm, where the interspersing by delay elements can be absorbed. The modified algorithm will minimise then same cost function as the standard SBR2 algorithm, which in every iteration step will eliminate the maximum off-diagonal component. We have shown that due to the maintenance of the cost function and maximum search, the proposed algorithm inherits the convergence proof and properties of the standard SBR2 algorithm to a diagonalised and spectrally majorised solution for the polynomial eigenvalues. Particularly as the spatial dimension—i.e. the number of sensors recording the data—increases, the computational savings can become significant.

REFERENCES

- [1] 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.
- [2] 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.
- [3] S. Weiss, J. Pestana, I. Proudler, and F. 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.
- [4] 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.
- [5] J. Corr, K. Thompson, S. Weiss, J. McWhirter, S. Redif, and I. Proudler, "Multiple shift maximum element sequential matrix diagonalisation for parahermitian matrices," in *IEEE Workshop on Statistical Signal Processing*, Gold Coast, Australia, pp. 312–315, June 2014.
- [6] 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.
- [7] S. Weiss, S. Redif, T. Cooper, C. Liu, P. Baxter, and J. McWhirter, "Paraunitary oversampled filter bank design for channel coding," *EURASIP Journal on Advances in Signal Processing*, vol. 2006, pp. 1–10, 2006.
- [8] 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, pp. 839–843, Sep. 2015.
- [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. W. Neo, C. Evers, and P. A. Naylor, "Speech enhancement using polynomial eigenvalue decomposition," in *2019 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, New Paltz, NY, pp. 125–129, Oct. 2019.
- [11] J. Corr, K. Thompson, S. Weiss, J. McWhirter, and I. Proudler, "Cyclic-by-row approximation of iterative polynomial EVD algorithms," in *Sensor Signal Processing for Defence*, Edinburgh, Scotland, pp. 1–5, Sep. 2014.
- [12] J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter, "Row-shift corrected truncation of paraunitary matrices for PEVD algorithms," in *23rd European Signal Processing Conference*, Nice, France, pp. 849–853, Aug./Sep. 2015.
- [13] —, "Reduced search space multiple shift maximum element sequential matrix diagonalisation algorithm," in *IET/EURASIP Intelligent Signal Processing*, London, UK, Dec. 2015.
- [14] F. Coutts, J. Corr, K. Thompson, I. Proudler, and S. Weiss, "Divide-and-conquer sequential matrix diagonalisation for parahermitian matrices," in *Sensor Signal Processing for Defence Conference*, London, UK, pp. 1–5, Dec. 2017.
- [15] 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.
- [16] S. Kasap and S. Redif, "Novel field-programmable gate array architecture for computing the eigenvalue decomposition of para-hermitian polynomial matrices," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 3, pp. 522–536, Mar. 2014.
- [17] R. W. Stewart, R. Chapman, and T. S. Durrani, "The Square Root In Signal Processing," in *Real-Time Signal Processing XII*, J. P. Letellier, Ed., vol. 1154, International Society for Optics and Photonics. SPIE, pp. 89 – 101, 1989.
- [18] J. Götze and U. Schwiegelshohn, "A square root and division free Givens rotation for solving least squares problems on systolic arrays," *SIAM Journal on Scientific and Statistical Computing*, vol. 12, no. 4, pp. 800–807, 1991.
- [19] M. Moonen and I. Proudler, "Generating 'fast qr' algorithms using signal flow graph techniques," in *Conference Record of The Thirtieth Asilomar Conference on Signals, Systems and Computers*, vol. 1, pp. 410–414, Nov. 1996.
- [20] W. Rath, "Fast Givens rotations for orthogonal similarity transformations," *Numerische Mathematik*, vol. 40, no. 1, pp. 46–56, 1982.
- [21] P. Vaidyanathan, "Theory of optimal orthonormal subband coders," *IEEE Transactions on Signal Processing*, vol. 46, no. 6, pp. 1528–1543, Jun. 1998.
- [22] 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*, , pp. 844–848, Nice, France, Sep. 2015.
- [23] 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.
- [24] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, Maryland: John Hopkins University Press, 1996.
- [25] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs: Prentice Hall, 1993.