

Cyclic-by-Row Approximation of Iterative Polynomial EVD Algorithms

Jamie Corr*, Keith Thompson*, Stephan Weiss*, John G. McWhirter†, and Ian K. Proudler‡

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

†School of Engineering, Cardiff University, Cardiff, Wales, UK

‡School of Electrical, Electronics and Systems Engineering, Loughborough University, Loughborough, UK

{jamie.corr,keith.thompson,stephan.weiss}@strath.ac.uk

Abstract—A recent class of sequential matrix diagonalisation (SMD) algorithms have been demonstrated to provide a fast converging solution to iteratively approximating the polynomial eigenvalue decomposition of a parahermitian matrix. However, the calculation of an EVD, and the application of a full unitary matrix to every time lag of the parahermitian matrix in the SMD algorithm results in a high numerical cost. In this paper, we replace the EVD with a limited number of Givens rotations forming a cyclic-by-row Jacobi sweep. Simulations indicate that a considerable reduction in computational complexity compared to SMD can be achieved with a negligible sacrifice in diagonalisation performance, such that the benefits in applying the SMD are maintained.

I. INTRODUCTION

For many signal processing and in particular array problems, the eigenvalue decomposition (EVD) of the data's covariance matrix is an essential tool. The EVD permits e.g. subspace decompositions or the construction of optimum linear processing for MIMO systems. For the broadband case with a multichannel data vector $\mathbf{x}[n]$, the use of narrowband phase shifts are not sufficient, and the time delays have to be explicitly addressed. The space-time covariance matrix $\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^H[n-\tau]\}$, where $\mathcal{E}\{\cdot\}$ is the expectation operator and $\{\cdot\}^H$ denotes the Hermitian transpose, thus arises. Its z -transform is the cross-spectral density (CSD) matrix $\mathbf{R}(z) = \sum_{\tau} \mathbf{R}[\tau]z^{-\tau}$, which is a polynomial matrix that cannot be decomposed by the narrowband EVD.

A polynomial EVD (PEVD) was defined in the context of subband coding [2], which can be applied to a parahermitian matrix $\mathbf{R}(z) = \tilde{\mathbf{R}}(z)$ such that

$$\mathbf{R}(z) \approx \mathbf{H}(z)\mathbf{D}(z)\tilde{\mathbf{H}}(z) \quad , \quad (1)$$

with the parahermitian operator $\tilde{\mathbf{R}}(z) = \mathbf{R}^H(z^{-1})$. The factors on the r.h.s. of (1) include a paraunitary matrix $\mathbf{H}(z)$ such that $\mathbf{H}(z)\tilde{\mathbf{H}}(z) = \tilde{\mathbf{H}}(z)\mathbf{H}(z) = \mathbf{I}$, and a parahermitian $\mathbf{D}(z)$, which is diagonal,

$$\mathbf{D}(z) = \text{diag}\{D_0(z) D_1(z) \dots D_{M-1}(z)\} \quad . \quad (2)$$

The matrix $\mathbf{D}(z)$ is spectrally majorised, such that the power spectral densities $D_m(e^{j\Omega}) = D_m(z)|_{z=e^{j\Omega}}$ satisfy $D_{m+1}(e^{j\Omega}) \geq D_m(e^{j\Omega}) \forall \Omega, m = 0 \dots (M-1)$. In [2], (1) is stated with equality for the optimum subband coder, even though the existence of an exact PEVD with FIR paraunitary

matrices is not guaranteed [3]. However, the close approximation in (1) is supported in [4] for $\mathbf{H}(z)$ of sufficiently high order.

To calculate the PEVD, a number of iterative algorithms have been developed. The second order sequential best rotation (SBR2) algorithm [3] is a generalisation of the Jacobi method whereby in every step the off-diagonal element of maximum absolute value — referred to below as the dominant element — of $\mathbf{R}(z)$ is eliminated by an elementary paraunitary operation. In this operation, a time shift matrix brings this dominant off-diagonal element into the zero lag matrix, where its energy is then transferred onto the diagonal by a Givens rotation. An alternative approach in [5] uses a paraunitary filter bank of fixed order with a lower computational complexity than SBR2, but unlike SBR2 is not guaranteed to converge. A modified SBR2 version, optimised for the coding gain problem, was proposed in [6], which identifies and eliminates the maximum normalised off-diagonal element at every iteration step.

A new family of algorithms, termed sequential matrix diagonalisation (SMD), have been introduced in [7]–[9]. Different from SBR2, these algorithms do not only transfer the dominant element, but instead eliminate all off-diagonal elements in the zero lag plane by means of a (scalar) EVD. SMD algorithms have been demonstrated to diagonalise a parahermitian matrix in fewer iterations, with higher precision, and with a lower order paraunitary matrix $\mathbf{H}(z)$ than SBR2, which makes SMD very attractive for applications that e.g. require accurate broadband subspace decompositions [10]. A number of SMD algorithms will be reviewed later in this paper; while their application is attractive, the computational complexity of performing the actual decomposition is significantly higher than SBR2.

This paper proposes a new low cost approximation for the SMD family of algorithms which significantly reduces the computational complexity but maintains algorithm performance. Based on a review of iterative PEVD algorithms in Sec. II, Sec. III outlines the proposed approach, which replaces the EVD at every iteration by a single Jacobi sweep implemented by a fixed number of Givens rotations in a cyclic-by-row arrangement [11], [12]. Simulation results are presented in Sec. IV to highlight the performance of cyclic-by-row approximations of the SMD benchmarked against state-of-the-art algorithms. Finally, conclusions are drawn in Sec. V.

because more energy is transferred from off-diagonal to on-diagonal elements. However, an EVD has to be calculated at every iteration, and the modal matrix $\mathbf{Q}^{(i)}$ has no longer the sparse structure of the Givens rotation in (10), requiring a full matrix multiplication at every lag τ .

While the EVD and full matrix multiplications makes the SMD algorithm much more costly to calculate compared to SBR2, two benefits can be noted:

- SMD can reach levels of diagonalisation that SBR2 cannot achieve;
- SMD generally leads to lower-order paraunitary matrices $\mathbf{H}(z)$, which have a lower cost compared to those obtainable by SBR2.

Therefore, for applications such as those requiring broadband signal subspace decompositions, SMD permits better performance with lower order paraunitary filter banks compared to SBR2.

D. Multiple Shift Maximum Element SMD Algorithm

To simplify the search phase of the SMD algorithm and avoid the calculation of modified column norms at every iteration, a maximum element SMD (ME-SMD) algorithm has been proposed in [9]. Replacing the l_2 norm in (13) by the l_∞ norm, the parameter search and delay operation within a single iteration is identical to SBR2. However, the zero-lag slice $\mathbf{S}^{(i)}[0]$ is still diagonalised at every step, such that from iteration to iteration step, ME-SMD differs from SBR2. The ME-SMD algorithm has a very similar behaviour to the standard SMD algorithm, in terms of convergence, but also in terms of complexity, which is dominated by the EVD calculation and modal matrix application, rather than the parameter search.

Realising that the main advantage of SMD over SBR2 comes from the enhanced reduction in off-diagonal energy at every iteration, a multiple shift ME-SMD algorithm has been proposed [7], which shifts not one but several columns with elements onto the lag zero slice. As a column operation will partially undo previous row operations and vice versa, the sequence of these shifts is important. This sequence is guided by the definition of search spaces $\mathcal{S}^{(i,j)}$ to which the parameter search is limited during the j th step of the i th iteration [7], [8], with $\mathcal{S}^{(i,j)}$ being a function of both $\mathcal{S}^{(i,j-1)}$ and $k^{(i,j-1)}$. Based on $\mathcal{S}^{(i,j)}$, the identified parameter set will be

$$\{k^{(i,j)}, \tau^{(i,j)}\} = \arg \max_{k \in \mathcal{S}^{(i,j)}, \tau} \|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_\infty, \quad (14)$$

Different search spaces have been defined to ensure that the maximal number $J = M - 1$ columns can be brought onto the zero lag matrix [7], [8]. All of the above algorithms have in common a considerably enhanced diagonalisation over the basic SMD algorithm as discussed in Sec. II-C.

III. CYCLIC-BY-ROW ITERATIVE PEVD

Sec. II-C identified the computational cost to compute an approximate PEVD via SMD algorithms as a potential obstacle. In experiments using Matlab's profiler, the calculation

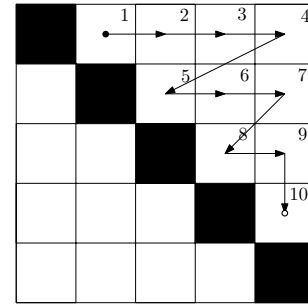


Fig. 1. Cyclic-by-row execution of Givens rotations implementing one Jacobi sweep, exemplified for a 5×5 matrix with start \bullet and end point \circ .

and application of an EVD per iteration step was singled out as the major contributor to this high cost. Therefore, this section proposes an inexpensive numerical approximation of the EVD by a Jacobi sweep consisting of a limited number of Givens rotations in a cyclic-by-row approach [12]. Below, Sec. III-A motivates the approach and Sec. III-B outlines the general procedure, which is then applied to a number of SMD algorithms in Sec. III-C.

A. EVD Approximation

A number of variations exist to implement the EVD of a Hermitian matrix [11], [12]. An iterative approach to an approximate EVD is the classical Jacobi algorithm [12], which consists of a sequence of Givens rotations targeting the maximum off-diagonal element at each iteration step — note that the SBR2 algorithm is a generalisation of this technique to the parahermitian case [3]. A simpler alternative which does not require a maximum search at every step is the cyclic-by-row algorithm [11], [12], which is outlined below

The cyclic-by-row approach uses a sequence of so-called Jacobi sweeps until off-diagonal elements are suppressed below a given threshold or a predefined number of sweeps have been executed. A Jacobi sweep for an $M \times M$ matrix consists of a fixed number of $(M^2 - M)/2$ Givens rotations, which are applied in a cyclic row approach as highlighted in Fig. 1. Each Givens rotation as defined in (10) will transfer the energy of an off-diagonal element onto the diagonal while undoing some of the work of previous Givens rotations. However, over the course of one Jacobi sweep, the off-diagonal energy is reduced. By performing more Jacobi sweeps, the accuracy of the EVD approximation is improved.

B. Cyclic-by-Row Algorithm

Iterative approximate PEVD algorithms such as SBR2 and SMD minimise off-diagonal energy until a predefined threshold ρ is reached, as described by (11). Therefore, within one iteration step of SMD, a full EVD with a suppression of off-diagonal energy to the numerical equivalent of zero appears to be an overkill, and a lower precision with a limited number of Jacobi sweeps will very likely suffice to achieve the task of reducing off-diagonal energy below the value ρ .

Experimentation has shown that for the approximation of the SMD algorithms detailed below, a single cyclic-by-row

Jacobi sweep proved sufficient and provided the best cost-performance trade-off, as we will detail in Sec. IV. With this approach, the unitary Jacobi rotation matrix $\mathbf{Q}^{(i)}$ from SBR2 in (10) now becomes the product of $N = (M^2 - M)/2$ Givens rotations,

$$\mathbf{Q}^{(i)} = \prod_{n=1}^N \mathbf{Q}^{(i,n)}, \quad (15)$$

where $\mathbf{Q}^{(i,n)}$ is the n th Givens rotation used in the i th iteration of an iterative PEVD algorithm using the single sweep cyclic-by-row approach.

C. Cyclic-by-Row SMD Algorithms

The cyclic-by-row single Jacobi sweep approximation of the EVD can be embedded in all algorithms of the SMD family. It may be argued that the term sequential matrix *diagonalisation* is no longer appropriate, as the approximate EVD also results in only an approximate diagonalisation, and algorithms will therefore share some properties of SBR2, where only part of the off-diagonal energy of the lag zero matrix is transferred onto the diagonal. However, we assume that the approximation is within the bound ρ for off-diagonal energy, and that therefore the term *diagonalisation* is justified within the SMD family's limited, pre-defined accuracy of decomposition.

All SMD algorithms perform an initial diagonalisation by an EVD according to (12), which in the cyclic-by-row version is approximated by a single Jacobi sweep. The EVD in subsequent iterations is also replaced by a single Jacobi sweep, and the unitary matrix (15) as applied in (6) can be implemented as a sequence of Givens rotations rather than a full matrix multiplication. The specific SMD family versions therefore consistently apply the single Jacobi sweep approach, and only differ in the way columns and rows are identified for transfer to the lag zero matrix using (5) at the i th iteration:

- SMD [9]: in its original form, the sequential matrix diagonalisation algorithm transfers the column with the largest off-diagonal column norm onto the lag zero matrix;
- ME-SMD [9]: with a simplified search compared to SMD, the column containing the maximum off-diagonal element is transferred to the lag zero matrix;
- MSME-SMD [7]: this multiple-shift version transfers $(M - 1)$ columns identified by their maximum elements;
- C-MSME-SMD [8]: a causally-constrained multiple-shift version, with a more restricted search space for maxima as compared to MSME-SMD.

The cyclic-by-row approximations of these algorithms will be compared to their standard versions as well as to SBR2 in the next section.

IV. RESULTS

A. Simulation Set-Up and Performance Metrics

To assess the proposed iterative PEVD algorithms, we consider the diagonalisation performance over an ensemble of 10^2 random 5×5 parahermitian matrices $\mathbf{R}(z)$ of order 11. Each instance of $\mathbf{R}(z)$ is generated as $\mathbf{R}(z) = \mathbf{A}(z)\tilde{\mathbf{A}}(z)$,

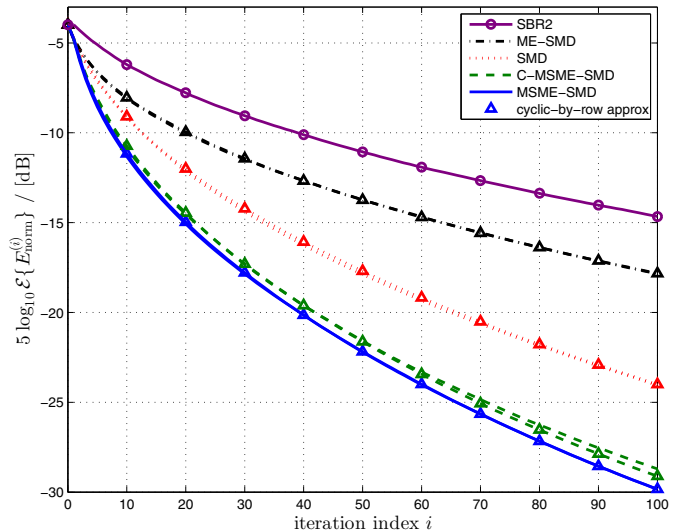


Fig. 2. Normalised remaining off-diagonal energy $E_{\text{norm}}^{(i)}$ according to (16) for various iterative PEVD algorithms versus iterations.

where $\mathbf{A}(z) \in \mathbb{C}^{5 \times 5}$ is a random polynomial matrix of order 6 with independent and identically distributed zero mean unit variance complex Gaussian entries. - To measure diagonalisation, the remaining normalised off-diagonal energy after i iterations,

$$E_{\text{norm}}^{(i)} = \frac{\sum_{\tau} \sum_{k=1}^M \|\hat{\mathbf{s}}_k^{(i)}[\tau]\|_2^2}{\sum_{\tau} \|\mathbf{R}[\tau]\|_F^2}, \quad (16)$$

is considered, where $\hat{\mathbf{s}}_k^{(i)}[\tau]$ is the modified column vector of (9) and $\|\cdot\|_F$ the Frobenius norm. The computational complexity of various algorithms is measured by the execution time in Matlab R2013a under Ubuntu 12.04 on a Dell Precision T3610 with Intel® Xeon® E5-1607V2 3.00 GHz x 4 cores and 8 GB RAM.

B. Convergence Speed

The diagonalisation performance versus iterations is shown in Fig. 2 for the various algorithms. The SMD family converges generally significantly faster than SBR2, as also highlighted in [7]–[9], with the multiple-shift versions performing best. Interestingly, the cyclic-by row single Jacobi sweep approximations of the EVD lead to no noticeable performance degradation for the SMD family of algorithms, thereby confirming the single sweep selection in Sec. III-B.

C. Calculation Cost

The computation time required for 100 iterations of the different iterative PVD algorithms is plotted in Fig. 3. Compared to SBR2, the SMD algorithms require much more processing time. However, a significant reduction in cost can be noticed for the cyclic-by-row approximations of SMD algorithms. These are still more costly than the SBR2 algorithm, which only requires a single Givens rotation per iteration and therefore is guaranteed to have a lower complexity.

More interesting than the cost per iteration is the required execution time to reach a specific level of diagonalisation.

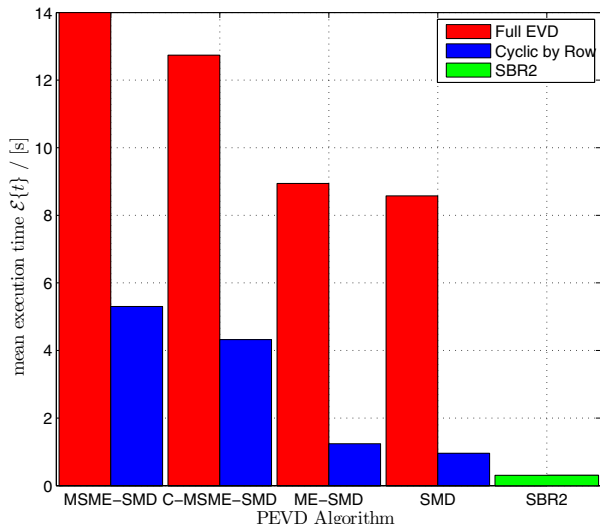


Fig. 3. Mean execution time for 100 iterations of the various PEVD algorithms.

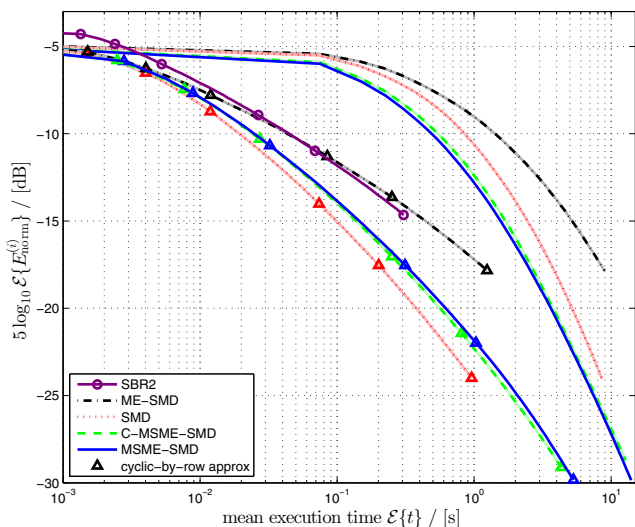


Fig. 4. Normalised remaining off-diagonal energy versus mean execution time for various iterative PEVD algorithms.

Fig. 4 shows the normalised remaining off-diagonal energy as a function of the time taken to calculate this specific decomposition. This graph is obtained by merging the information of Fig. 2 and 3. Using a full EVD, the SMD family of algorithms are inferior to SBR2 when a low level of diagonalisation suffices. For high values of diagonalisation, SBR2 cannot provide the required diagonalisation, and the computationally expensive SMD family is on the only option, with particular benefit for multiple-shift versions as established in [7], [8].

Using the cyclic-by-row approximation, the actual cost to reach a specific level of diagonalisation is reduced below even what is required for SBR2. Therefore, the cyclic-by-row approximation of SMD algorithms does not only yield lower application cost — as established in [7], [9], SMD decompositions yield paraunitary $H(z)$ of lower order — but in terms of total cost per decomposition is also less expensive to calculate than SBR2.

V. CONCLUSION

This paper has presented a cyclic-by-row approximation of sequential matrix diagonalisation algorithms. SMD algorithms have previously been shown to be attractive for application because of their excellent diagonalisation performance and lower order paraunitary filter banks compared to sequential best rotation algorithms; however the SMD family has been previously been shown to be computationally costly to calculate due an EVD required at every iteration. This costly step has been replaced by a single Jacobi sweep using a fixed number of Givens rotations in a cyclic-by-row approach.

Simulation results indicate that the algorithm performance is not compromised by this approximation, such that SMD algorithms retain their superior features for application. However, the computational cost is reduced such that a significantly lower complexity to calculate a decomposition compared to standard SMD implementation is achieved. This holds also for lower levels of diagonalisation, where SBR2 previously retained an advantage in terms of calculation cost, such that the proposed cyclic-by-row approximations of SMD algorithms now globally offer the best cost-performance trade-off both in terms of calculating and applying an iterative PEVD.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] P. Vaidyanathan. Theory of optimal orthonormal subband coders. *IEEE Trans SP*, 46(6):1528–1543, June 1998.
- [2] P. Vaidyanathan. *Multirate Systems and Filter Banks*. Prentice Hall, New Jersey, 1993.
- [3] J.G. McWhirter, P.D. Baxter, T. Cooper, S. Redif, and J. Foster. An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans SP*, 55(5):2158–2169, May 2007.
- [4] S. Icart and P. Comon. Some properties of Laurent polynomial matrices. In *9th IMA Conf. Maths in Sig. Proc.*, Birmingham, UK, Dec. 2012.
- [5] A. Tkachenko. Approximate eigenvalue decomposition of para-hermitian systems through successive fir paraunitary transformations. In *IEEE Int. Conf. Acoustics Speech Signal Proc.*, pp. 4074–4077, Mar. 2010.
- [6] S. Redif, J.G. McWhirter, and S. Weiss. Design of FIR paraunitary filter banks for subband coding using a polynomial eigenvalue decomposition. *IEEE Trans SP*, 59(11):5253–5264, Nov. 2011.
- [7] J. Corr, K. Thompson, S. Weiss, J.G. McWhirter, S. Redif, I.K. Proudler Multiple Shift Maximum Element Sequential Matrix Diagonalisation for Parahermitian Matrices. *IEEE Workshop on Statistical Signal Proc.*, Gold Coast, Australia, June. 2014.
- [8] J. Corr, K. Thompson, S. Weiss, J. G. McWhirter, I. K. Proudler Causal Multiple Shift Maximum Element Sequential Matrix Diagonalisation for Parahermitian Matrices. *22nd European Signal Proc. Conf.*, Lisbon, Portugal, Sept. 2014.
- [9] S. Redif, S. Weiss, and J.G. McWhirter. Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans SP*, submitted, 2014.
- [10] M. Alrmah, S. Weiss, and S. Lambotharan. An extension of the MUSIC algorithm to broadband scenarios using polynomial eigenvalue decomposition. In *19th European Signal Proc. Conf.*, pp. 629–633, Barcelona, Spain, Aug. 2011.
- [11] J. Götzte Monitoring the stage of diagonalization in Jacobi-type methods *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Adelaide, Australia, vol. III, pp. 441–444, Apr. 2014.
- [12] G.H. Golub and C.F. Van Loan. *Matrix Computations*. John Hopkins University Press, 3rd ed., 1996.