

Performance Trade-Offs in Sequential Matrix Diagonalisation Search Strategies

Jamie Corr*, Keith Thompson*, Stephan Weiss*, John G. McWhirter†, 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—Recently a selection of sequential matrix diagonalisation (SMD) algorithms have been introduced which approximate polynomial eigenvalue decomposition of parahermitian matrices. These variants differ only in the search methods that are used to bring energy onto the zero-lag. Here we analyse the search methods in terms of their computational complexities for different sizes of parahermitian matrices which are verified through simulated execution times. Another important factor for these search methods is their ability to transfer energy. Simulations show that the more computationally complex search methods transfer a greater proportion of the off-diagonal energy onto the zero-lag over a selected range of parahermitian matrix sizes. Despite their higher cost per iteration experiments indicate that the more complex search algorithms still converge faster in real time.

I. INTRODUCTION

Sequential matrix diagonalisation encompasses a family of iterative algorithms that can factorise a parahermitian matrix into an approximate polynomial matrix eigenvalue decomposition (PEVD). The PEVD extends the wide-ranging utility of the EVD from narrowband to broadband problems, and iterative PEVD algorithms have in the past found use in optimal subband [1] and multichannel coding [2]; channel coding [3], transmit and receive beamforming across broadband MIMO channel [4], [5], angle of arrival estimation [6]. It can also provide a preprocessing stage for beamforming by applying denoising [7], decorrelation [8] and optimum subband decompositions [9], or enable novel MVDR beamforming approaches [10].

Parahermitian matrices arise e.g. by including an explicit lag τ to create the space-time covariance $\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^H[n-\tau]\}$. The matrix elements are auto- and cross-correlation sequences that create a symmetry, $\mathbf{R}[\tau] = \mathbf{R}^H[-\tau]$, i.e. a parahermitian matrix is equal to its complex conjugate, time reversed version. Taking the z-transform of the space-time covariance matrix yields the cross spectral density matrix $\mathbf{R}(z) \bullet\!\!\!\circ \mathbf{R}[\tau]$. The parahermitian property is expressed as $\mathbf{R}(z) = \tilde{\mathbf{R}}(z)$, where the parahermitian operator $\{\tilde{\cdot}\}$ implies complex conjugation and time reversal. A polynomial EVD [11]–[13] of such a parahermitian matrix,

$$\mathbf{\Lambda}(z) \approx \mathbf{Q}(z)\mathbf{R}(z)\tilde{\mathbf{Q}}(z) \quad , \quad (1)$$

is claimed to exist in close approximation for FIR paraunitary matrices $\mathbf{Q}(z)$ of sufficiently high order [14].

A number of iterative PEVD algorithms have been introduced, including second order sequential best rotation (SBR2) methods [1], [13], which eliminate the maximum off-diagonal element at every iteration. An approximate PEVD [15] operates on a fixed order paraunitary matrix, but unlike SBR2 has not been proven to converge. More recently, a sequential matrix diagonalisation (SMD) algorithm has been introduced [2], which is also proven to converge but seems capable of attaining better diagonalisation of the space-time covariance matrix $\mathbf{R}(z)$ than SBR2 algorithms [1], [13]. The SMD family has been extended by a multiple-shift version, which has been found to transfer even more energy per iteration [16], with additional searches the only cost increase over SMD. Therefore, the aim of this paper is to explore some of the performance trade-off details w.r.t. computational cost and diagonalisation, between the different SMD algorithms.

The paper is organised as follows. Iterative PEVD algorithms based on the idea of sequential matrix diagonalisation are introduced in Sec. II. The difference in the SMD search steps and their associated cost is investigated in Sec. III. Finally, performance metrics and simulation results are presented in Sec. IV followed by conclusions in Sec. V.

II. SEQUENTIAL MATRIX DIAGONALISATION

The sequential matrix diagonalisation algorithm (SMD) [2] and a number of derivative versions [16]–[18] iteratively diagonalise a parahermitian matrix to approximate its PEVD. The initialisation step of any SMD algorithm fully diagonalises the zero-lag of the parahermitian matrix, $\mathbf{R}[0]$, achieved via the modal matrix $\mathbf{Q}^{(0)}$ of the ordered EVD of $\mathbf{R}[0]$, which is applied to all lags of the parahermitian matrix,

$$\mathbf{S}^{(0)}(z) = \mathbf{Q}^{(0)}\mathbf{R}(z)\mathbf{Q}^{(0)H} \quad . \quad (2)$$

Each iteration of the SMD algorithm includes a shift operation which brings off-diagonal energy onto the zero-lag,

$$\mathbf{S}^{(i)'}(z) = \mathbf{\Lambda}^{(i)}(z)\mathbf{S}^{(i-1)}(z)\tilde{\mathbf{\Lambda}}^{(i)}(z) \quad , \quad i = 1 \dots I \quad . \quad (3)$$

The shift matrix, $\mathbf{\Lambda}^{(i)}(z)$, is determined by the search strategy which varies between SMD versions and will be discussed further in Sec III.

An SMD iteration is then completed by transferring energy from the zero-lag on to the diagonal. Like the initialisation

step, this consists of applying the EVD modal matrix to the entire parahermitian matrix,

$$\mathbf{S}^{(i)}(z) = \mathbf{Q}^{(i)} \mathbf{S}^{(i)'}(z) \mathbf{Q}^{(i)H} \quad . \quad (4)$$

This diagonalises [2], [16]–[18] the zero-lag matrix $\mathbf{S}^{(i)}[0]$. In both EVD steps the application of the modal matrices, $\mathbf{Q}^{(0)}$ & $\mathbf{Q}^{(i)}$, to all lags of the parahermitian matrix represents a large proportion of the cost of the whole SMD algorithm.

The SMD algorithms repeat steps (3) & (4) for either a set number of iterations or until some threshold based on the parahermitian matrix is reached (e.g. magnitude if the max off-diagonal element). Thus after a number of I iterations, the paraunitary matrix which approximately diagonalises $\mathbf{R}(z)$ is obtained by the product of the matrices produced by the initialisation and steps (3) & (4) from each iteration i.e.

$$\hat{\mathbf{Q}}(z) = \mathbf{G}^{(I)}(z) \dots \mathbf{G}^{(1)}(z) \mathbf{G}^{(0)}(z) \quad , \quad (5)$$

where each $\mathbf{G}^{(i)}(z)$ is constructed from the delay and energy transfer matrices from the i th step i.e.

$$\mathbf{G}^{(i)}(z) = \mathbf{Q}^{(i)} \mathbf{\Lambda}^{(i)}(z) \quad . \quad (6)$$

This approximates a PEVD with

$$\mathbf{S}^{(I)}(z) = \hat{\mathbf{Q}}(z) \mathbf{R}(z) \tilde{\hat{\mathbf{Q}}}(z) \quad . \quad (7)$$

Convergence proofs for the various SMD algorithms show that for a sufficiently high I , the off-diagonal energy in $\mathbf{S}^{(I)}(z)$ can be reduced to an arbitrarily low bound.

III. SEARCH METHODS

This section discusses how SMD algorithms identify the elements to be transferred in the first SMD step outlined in Sec. II. We mainly consider the order $\mathcal{O}(\cdot)$ of the computational complexity in terms of multiply-accumulate operations, which depends on the matrix dimension M and the lag dimension L , whereby it must be noted that L grows with every iteration [2], [13] and the extent of the growth varies with both algorithm and input parahermitian matrix $\mathbf{R}(z)$ [19].

A. Column Norm / SMD-Algorithm

The original SMD algorithm [2] in its i th iteration inspects the vectors $\hat{\mathbf{s}}_k^{(i-1)}[\tau]$, which are the columns of $\mathbf{S}^{(i-1)}[\tau]$ but modified by removing its on-diagonal elements. The pair

$$\{k^{(i)}, \tau^{(i)}\} = \arg \max_{k, \tau} \|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_2, \quad i = 1 \dots I, \quad (8)$$

points to the vector $\hat{\mathbf{s}}_k^{(i-1)}[\tau]$ with maximum norm, which is transferred onto the zero-lag matrix and subsequently eliminated. A total of ML column norms of the parahermitian matrix have to be calculated. Each norm requires a squaring of elements, but the square root operation can be omitted as only a comparison of norms but no explicit values are required. Thus, with each column vector having length M , the norm computation is $\mathcal{O}(M^2L)$ followed by a search over $\mathcal{O}(ML)$ elements.

TABLE I
COST COMPARISON OF SMD SEARCH METHODS.

method	norm calc.	comparisons	total
SMD	$\mathcal{O}(M^2L)$	$\mathcal{O}(ML)$	$\mathcal{O}(M^2L)$
ME-SMD	$\mathcal{O}(0)$	$\mathcal{O}(M^2L)$	$\mathcal{O}(M^2L)$
MSME-SMD	$\mathcal{O}(0)$	$\mathcal{O}(M^3L)$	$\mathcal{O}(M^3L)$

B. Maximum Element / ME-SMD Algorithm

Introduced as a simplification to the SMD search in Sec. III-A, the maximum-element SMD (ME-SMD) algorithm [2] replaces the l_2 norm in (8) by the l_∞ norm. Thus in each iteration, the maximum element can be identified without any explicit norm calculation but requires a search over an enlarged set of $\mathcal{O}(M^2L)$ elements. The energy that is transferred in a single step by ME-SMD is always smaller than or equal to that eliminated by the original SMD version, but the algorithm was designed with the expectation of a lower computational complexity.

C. Multiple Shift Maximum Element / MSME-SMD Algorithm

The multiple shift maximum element (MSME) search method, used in the MSME-SMD algorithm [16], initiates every iteration by scanning the entire parahermitian matrix for its maximum off-diagonal element similar to ME-SMD, employing the l_∞ instead of the l_2 norm in (8). However, MSME-SMD does not transfer just one-column into the zero-lag, but will perform a total of $(M-1)$ column shifts to increase the energy transfer in the second step of each iteration. This is achieved by $(M-1)$ searches over increasingly limited search spaces such that previously identified and shifted maxima are not undone by later shifts [16].

This approach requires no norm evaluations but the complexity of the search is $\mathcal{O}(M^3L)$ because each iteration involves searching M^2L elements a total of $M-1$ times (where for asymptotic analysis $M-1$ is simplified to M). An overall cost comparison of the three search methods is provided in Tab. I, with a total search cost order provided on the basis that one comparison for the maximum search is about as expensive as one multiply-accumulate operation. Note the value of L varies with algorithm and grows at each iteration.

IV. RESULTS

A. Simulation Set-Up and Performance Metrics

To assess the proposed search algorithms, we consider an ensemble of 10^3 random $M \times M$ parahermitian matrices $\mathbf{R}(z)$ of order $2L-1$, for $M = 2, 4, \dots, 20$, & $L = 50, 100, \dots, 500$. Each instance of $\mathbf{R}(z)$ is generated as $\mathbf{R}(z) = \mathbf{A}(z) \tilde{\mathbf{A}}(z)$, where $\mathbf{A}(z) \in \mathbb{C}^{M \times M}$ is a random polynomial matrix of order L with independent and identically distributed zero mean and unit variance complex Gaussian entries.

Due to the variation in L , execution time is used to measure the computational complexity of the search methods in Matlab 2014a with the following system specification: Ubuntu 14.04 on a Dell Precision T3610 with Intel® Xeon® E5-1607V2 3.00 GHz x 4 cores and 8 GB RAM.

In addition to execution time we have also investigated the proportion of the off-diagonal energy that is brought onto the zero-lag by each of these search algorithms for the same ensemble but M is restricted to 4, 10 & 20. The proportion of shifted energy, $E_{\text{shift}}^{(m,l)}$, is averaged over the ensemble and calculated as

$$E_{\text{shift}}^{(m,l)} = \frac{\sum_{k=1}^M \|\hat{\mathbf{s}}_k^{(m,l)'}[0]\|_2^2}{\sum_{\tau} \sum_{k=1}^M \|\hat{\mathbf{s}}_k^{(m,l)}[\tau]\|_2^2}, \quad (9)$$

where $\hat{\mathbf{s}}_k^{(m,l)}[\tau]$ is the modified column vector from (8). The numerator in (9) is the off-diagonal energy brought onto the zero-lag, and the denominator is the off-diagonal energy in the entire parahermitian matrix. The algorithm that shifts most energy onto the zero-lag consequently produces the highest $E_{\text{shift}}^{(m,l)}$.

The final test measures diagonalisation, the remaining normalised off-diagonal energy after i iterations,

$$E_{\text{norm}}^{(i)} = \frac{\sum_{\tau} \sum_{k=1}^M \|\hat{\mathbf{s}}_k^{(i)}[\tau]\|_2^2}{\sum_{\tau} \|\mathbf{R}[\tau]\|_F^2}, \quad (10)$$

where $\mathbf{R}[\tau]$ is the initial parahermitian matrix and $\|\cdot\|_F$ the Frobenius norm. Unlike $E_{\text{shift}}^{(m,l)}$, the value for $E_{\text{norm}}^{(i)}$ should ideally be minimised. The matrix dimension M is restricted as above, the initial L is 6 and the ensemble is reduced to 10^2 .

B. Real Time Complexity

The $\mathcal{O}(\cdot)$ notation essentially only shows the shape of the time complexity of these search methods. It is also good to show the real time complexity as there could be hidden constants that dramatically affect real time performance [20], and the length of the parahermitian matrix, L , varies between algorithms and will increase at each iteration.

The real time complexity for the column norm search is given in Fig. 1; this agrees with the complexity analysis of $\mathcal{O}(M^2L)$ with the linear increase with L and shallow but polynomial increase with M . Although thought of as low cost, the real time performance of the maximum element search in Fig. 2 is very similar to that of the column norm approach both with the trends shown with matrix dimensions and the real time performance. Fig. 3 shows the the real time performance of the MSME search. The linear increase with number of lags and quite a steep polynomial increase with the matrix dimension agrees with the complexity analysis of $\mathcal{O}(M^3L)$. Comparing Figs. 1, 2 & 3 it is clear to see that the MSME search is significantly slower for larger matrix sizes than the other two.

C. Energy Transfer

Another important metric for these search algorithms is the amount of energy they bring onto the zero-lag at each iteration. This section will investigate how this varies with matrix dimensions for the various search methods.

The energy transfer for the column norm, maximum element and MSME methods are shown in Fig. 4. The most striking difference is that for the MSME approach the energy transfer




Fig. 1. Column norm search time for varying matrix size.




Fig. 2. Maximum element search time for varying matrix size.




Fig. 3. MSME search time for varying matrix size.

does not degrade as dramatically with the matrix dimension M , this is because as M increases the number of elements brought onto the zero-lag also increases. Comparing only the column norm and maximum element searches we see that the overall trends are very similar and that they degrade at a similar rate. Crucially, overall the column norm approach does indeed tend to transfer more energy than the maximum element method.

D. Real Time Convergence

While the previous two sections have focussed only on the search step of the SMD algorithms, here we show how the algorithms converge in real time over $I = 100$ iterations of each algorithm. Fig. 5 shows a real time convergence example for when $M = 4, 10$ & 20 , with the average final values for L given in Tab. II. Despite its higher computational cost and growth in L the fastest converging algorithm is the MSME-based implementation and the column norm approach converges faster than the algorithm using the maximum element search. When the search algorithms are included in the




Fig. 4. Average energy transferred for varying matrix size and search algorithm. Note the plot MSME, $M = 20$ is behind the Max element, $M = 4$.




Fig. 5. Reduction in off-diagonal energy vs. mean execution time over 100 algorithm iterations.

respective SMD algorithms a large proportion of the cost is dictated by the application of the EVD step in (3). As part of the SMD algorithms the benefit of the simpler search strategies is lost and it becomes more effective to employ the complex approach which transfers more energy.

V. CONCLUSION

The complexities and energy transfer associated with three of the search algorithms used in the SMD family of PEVD algorithms have been investigated in detail. The complexities range from the simplistic maximum element search right up to the more complex multiple shift maximum element search. The MSME search tends to obtain the greatest amount of energy at any iteration; however the multiple shifts cause the complexity to rise significantly with matrix dimensions. The maximum element and column norm searches have a similarly low complexity however this comes at the cost of lower energy transfer. From the results presented the maximum element version does not appear to have any significant benefit over the column norm based search however the column norm search will generally bring more energy onto the zero-lag and hence converge faster. Despite its significantly higher cost, the

TABLE II
AVERAGE FINAL LENGTH, L , FOR PARAHHERMITIAN MATRICES.

method	$M = 4$	$M = 10$	$M = 20$
SMD	74	76	67
ME-SMD	90	181	135
MSME-SMD	98	253	553

MSME search approach has been shown to converge faster in real time than the others for the experiments shown.

REFERENCES

- [1] S. Redif, J. McWhirter, and S. Weiss. Design of FIR paraunitary filter banks for subband coding using a polynomial eigenvalue decomposition. *IEEE Trans. on Sig. Proc.*, 59(11):5253–5264, Nov. 2011.
- [2] S. Redif, S. Weiss, and J. McWhirter. Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices. *IEEE Trans. on Sig. Proc.*, 63(1):81–89, Jan. 2015.
- [3] S. Weiss, C. H. Ta, and C. Liu. A Wiener filter approach to the design of filter bank based single-carrier precoding and equalisation. In *IEEE Int. Symp. on Power Line Comms. and Its Apps.*, pages 493–498, Pisa, Italy, Mar. 26–28 2007.
- [4] M. Davies, S. Lambbotharan, and J. McWhirter. Broadband MIMO beamforming using spatial-temporal filters and polynomial matrix decomposition. In *15th Int. Conf. on DSP*, pages 579–582, Cardiff, UK, Jul. 2007.
- [5] C. H. Ta and S. Weiss. A Design of Precoding and Equalisation for Broadband MIMO Systems. In *15th Int. Conf. on DSP*, pages 571–574, Cardiff, UK, Jul. 2007.
- [6] M. Almah, S. Weiss, and S. Lambbotharan. An extension of the MUSIC algorithm to broadband scenarios using polynomial eigenvalue decomposition. In *19th Euro. Sig. Proc. Conf.*, pages 629–633, Barcelona, Spain, Aug. 2011.
- [7] S. Redif, J. McWhirter, P. Baxter, and T. Cooper. Robust broadband adaptive beamforming via polynomial eigenvalues. In *OCEANS*, pages 1–6, Boston, MA, Sept. 2006.
- [8] C. L. Koh, S. Redif, and S. Weiss. Broadband GSC beamformer with spatial and temporal decorrelation. In *17th Euro. Sig. Proc. Conf.*, pages 889–893, Glasgow, Scotland, Aug. 2009.
- [9] P. Vouras and T. Tran. Robust transmit nulling in wideband arrays. *IEEE Trans. on Sig. Proc.*, 62(14):3706–3719, Jul. 2014.
- [10] S. Weiss, S. Bendoukha, A. Alzin, F. Coumts, I. Proudler, and J. Chambers. MVDR broadband beamforming using polynomial matrix techniques. In *Euro. Sig. Proc. Conf.*, Nice, France, Sept. 2015.
- [11] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, New York, 1982.
- [12] P. P. Vaidyanathan. *Multirate Systems and Filter Banks*. Prentice Hall, Englewood Cliffs, 1993.
- [13] J. G. McWhirter, P. D. Baxter, T. Cooper, S. Redif, and J. Foster. An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Sig. Proc.*, 55(5):2158–2169, May 2007.
- [14] S. Icart and P. Comon. Some properties of Laurent polynomial matrices. In *9th IMA Conf. on Maths. in Sig. Proc.*, Birmingham, UK, Dec. 2012.
- [15] A. Tkachenko. Approximate eigenvalue decomposition of para-hermitian systems through successive FIR paraunitary transformations. In *IEEE Int. Conf. on Acoustics Speech and Sig. Proc.*, pages 4074–4077, Dallas, TX, Mar. 2010.
- [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 Stat. Sig. Proc.*, pages 312–315, Gold Coast, Australia, June 2014.
- [17] J. Corr, K. Thompson, S. Weiss, J. McWhirter, and I. Proudler. Maximum energy sequential matrix diagonalisation for parahermitian matrices. In *48th Asilomar Conf. on Sig., Sys. and Comp.*, pages 470–474, Pacific Grove, CA, USA, Nov. 2014.
- [18] J. Corr, K. Thompson, S. Weiss, J. G. McWhirter, and I. K. Proudler. Causality-Constrained multiple shift sequential matrix diagonalisation for parahermitian matrices. In *22nd Euro. Sig. Proc. Conf.*, pages 1277–1281, Lisbon, Portugal, Sept. 2014.
- [19] J. Corr, K. Thompson, S. Weiss, I. K. Proudler, and J. G. McWhirter. Impact of source model matrix conditioning on PEVD algorithms. In *2nd IET Conf. on Intelligent Sig. Proc.*, London, England, Dec. 2015.
- [20] A. Aho, and J. Ullman. *Foundations of computer science*. W.H. Freeman, New York, 1994.