

Impact of Fast-Converging PEVD Algorithms on Broadband AoA Estimation

Fraser K. Coutts*, Keith Thompson*, Stephan Weiss*, Ian K. Proudler*,†

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

† School of Electrical, Electronics & Systems Engineering, Loughborough Univ., Loughborough, UK

{fraser.coutts,keith.thompson,stephan.weiss,ian.proudler}@strath.ac.uk

Abstract—Polynomial matrix eigenvalue decomposition (PEVD) algorithms have been shown to enable a solution to the broadband angle of arrival (AoA) estimation problem. A parahermitian cross-spectral density (CSD) matrix can be generated from samples gathered by multiple array elements. The application of the PEVD to this CSD matrix leads to a paraunitary matrix which can be used within the spatio-spectral polynomial multiple signal classification (SSP-MUSIC) AoA estimation algorithm. Here, we demonstrate that the recent low-complexity divide-and-conquer sequential matrix diagonalisation (DC-SMD) algorithm, when paired with SSP-MUSIC, is able to provide superior AoA estimation versus traditional PEVD methods for the same algorithm execution time. We also provide results that quantify the performance trade-offs that DC-SMD offers for various algorithm parameters, and show that algorithm convergence speed can be increased at the expense of increased decomposition error and poorer AoA estimation performance.

I. INTRODUCTION

For broadband angle of arrival (AoA) estimation, powerful narrowband methods such as the multiple signal classification (MUSIC) algorithm [1] are not directly applicable. In [2], the polynomial matrix eigenvalue decomposition (PEVD) [3] is used to generalise MUSIC to the case of polynomial space-time covariance matrices, resulting in the development of the spatio-spectral polynomial MUSIC (SSP-MUSIC) algorithm. A comparison in [4] of SSP-MUSIC with an auto-focussing coherent signal subspace AoA estimation approach [5] has found that SSP-MUSIC provides lower AoA estimation performance, but has the advantage of not relying on a-priori spectral information of the sources. Further work in [6] has shown that the accuracy of SSP-MUSIC depends strongly on the efficacy of the PEVD algorithm used.

Existing PEVD algorithms include sequential matrix diagonalisation (SMD) [7], second-order sequential best rotation (SBR2) [3], and various evolutions of the algorithm families [8]–[10]. Work in [11], [12] describes a divide-and-conquer approach for the PEVD. This algorithm — titled divide-and-conquer sequential matrix diagonalisation (DC-SMD) — can be utilised to reduce algorithm complexity and has a framework based on the SMD algorithm.

Here, we compare the AoA estimation performance that DC-SMD offers when paired with SSP-MUSIC with the performance that SBR2 and SMD provide in a broadband AoA estimation scenario. For the same example, we investigate the efficacy of DC-SMD for various algorithm parameters to determine the algorithm’s performance trade-offs. Performance of DC-SMD is measured in terms of AoA estimation accuracy and with several metrics, including algorithm execution time, decomposition mean-squared error, filter paraunitarity, and paraunitary filter length.

A broadband array data model and the SSP-MUSIC broadband AoA algorithm of [2] are outlined in Sec. II. Subse-

quently, Sec. III will provide a brief overview of the DC-SMD algorithm. Simulation results showing the broadband AoA performance of DC-SMD for various algorithm parameters are presented in Sec. IV, with conclusions drawn in Sec. V.

II. BROADBAND SUBSPACE DECOMPOSITION

Based on the signal model for a broadband array described in Sec. II-A, Sec. II-B defines the spatio-spectral polynomial MUSIC algorithm.

A. Broadband Array Data Model

An M -element array of omnidirectional sensors located at positions \mathbf{r}_m , $m = 1 \dots M$ collects broadband data in a vector $\mathbf{x}[n] \in \mathbb{C}^M$. For the l th far-field source $s_l[n]$, the array experiences a planar wavefront with normal \mathbf{k}_l . We are only interested in the relative delay between signals at the sensors, such that contributions to $\mathbf{x}[n]$ in the absence of attenuation due to propagation are

$$\mathbf{x}[n] = \sum_l \sum_{\nu=0}^{\infty} \mathbf{a}_l[\nu] s_l[n - \nu] + \mathbf{v}[n], \quad (1)$$

with the broadband steering vector

$$\mathbf{a}_l[n] = \frac{1}{\sqrt{M}} \begin{bmatrix} f[n - \tau_{l,0}] \\ \vdots \\ f[n - \tau_{l,M-1}] \end{bmatrix} \quad (2)$$

containing fractional delay filters $f[n - \tau_{l,m}]$ with normalised delays $\tau_{l,m} = \frac{\mathbf{k}_l^H \mathbf{r}_m}{cT_s}$, whereby T_s is the sampling period and c the propagation speed in the medium, such that \mathbf{k}_l/c is the l th source’s slowness vector. The vector $\mathbf{v}[n]$ adds spatially and temporally uncorrelated noise with covariance $\mathcal{E}\{\mathbf{v}[n]\mathbf{v}^H[n]\} = \sigma_v^2 \mathbf{I}_M$ to the model in (1), with $\mathcal{E}\{\cdot\}$ the expectation operator and \mathbf{I}_M an $M \times M$ identity matrix. Below, $\mathbf{a}_{\vartheta,\varphi}[n]$ refers to a broadband steering vector determined by \mathbf{k} with azimuth φ and elevation ϑ .

Collected by an M -element broadband array, the data vector $\mathbf{x}[n] \in \mathbb{C}^M$ has a space-time covariance matrix given by

$$\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^H[n - \tau]\}, \quad (3)$$

which forms a transform pair with the cross power spectral density (CSD) matrix $\mathbf{R}(z) \bullet \text{---} \circ \mathbf{R}[\tau]$,

$$\mathbf{R}(z) = \sum_{\tau} \mathbf{R}[\tau] z^{-\tau}. \quad (4)$$

The CSD matrix is parahermitian, i.e., $\mathbf{R}(z) = \tilde{\mathbf{R}}(z) = \mathbf{R}^H(1/z^*)$. A polynomial EVD [3] decouples the parahermitian $\mathbf{R}(z)$ by means of a paraunitary $\mathbf{Q}(z)$,

$$\mathbf{\Lambda}(z) \approx \tilde{\mathbf{Q}}(z) \mathbf{R}(z) \mathbf{Q}(z), \quad (5)$$

such that $\mathbf{\Lambda}(z) = \text{diag}\{\mathbf{\Lambda}_1(z) \mathbf{\Lambda}_2(z) \dots \mathbf{\Lambda}_M(z)\}$ is diagonalised and spectrally majorised with PSDs $\mathbf{\Lambda}_i(e^{j\Omega}) \geq \mathbf{\Lambda}_{i+1}(e^{j\Omega}) \forall \Omega, i = 1 \dots (M - 1)$, with $\mathbf{\Lambda}_i(e^{j\Omega}) = \mathbf{\Lambda}_i(z)|_{z=e^{j\Omega}}$. Equation (5) has only approximate equality, as the PEVD of a finite order polynomial matrix is generally not of finite order.

Thresholding the polynomial eigenvalues $\Lambda_i(z)$ — e.g., extracting eigenvalues with energy $\sum_{\tau} \|\Lambda_i[\tau]\|_{\text{F}}^2 > \lambda$ for some λ , where $\|\cdot\|_{\text{F}}$ is the Frobenius norm — reveals the number of independent broadband sources R contributing to $\mathbf{R}(z)$, and permits a distinction between signal-plus-noise and noise only subspaces $\mathbf{Q}_s(z) \in \mathbb{C}^{M \times R}$ and $\mathbf{Q}_n(z) \in \mathbb{C}^{M \times (M-R)}$,

$$\mathbf{R}(z) = \begin{bmatrix} \mathbf{Q}_s(z) & \mathbf{Q}_n(z) \end{bmatrix} \begin{bmatrix} \Lambda_s(z) & \mathbf{0} \\ \mathbf{0} & \Lambda_n(z) \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{Q}}_s(z) \\ \tilde{\mathbf{Q}}_n(z) \end{bmatrix},$$

where $R < M$, $\Lambda_s(z) \in \mathbb{C}^{R \times R}$ and $\Lambda_n(z) \in \mathbb{C}^{(M-R) \times (M-R)}$.

B. Spatio-Spectral Polynomial MUSIC

The spatio-spectral polynomial MUSIC (SSP-MUSIC) algorithm [2] is an extension of narrowband MUSIC [1] to the broadband case. The idea of the SSP-MUSIC algorithm is to scan the noise-only subspace $\mathbf{Q}_n(z)$,

$$\mathbf{Q}_n(z) = [\mathbf{q}_{R+1}(z) \dots \mathbf{q}_M(z)], \quad (6)$$

which is spanned by eigenvectors corresponding to eigenvalues close to the noise floor, $\Lambda_n(z) \approx \sigma_v^2 \mathbf{I}_{M-R}$. The steering vectors of sources that contribute to $\mathbf{R}(z)$ will define the signal-plus-noise subspace $\mathbf{Q}_s(z)$ and therefore lie in the nullspace of its complement $\mathbf{Q}_n(z)$. As a result, the vector $\tilde{\mathbf{Q}}_n(e^{j\Omega}) \mathbf{a}_{\vartheta, \varphi}(e^{j\Omega})$ has to be close to the origin for $\mathbf{a}_{\vartheta, \varphi}(e^{j\Omega})$ to be a steering vector of a contributing source at frequency Ω , where $\tilde{\mathbf{Q}}_n(e^{j\Omega}) = \tilde{\mathbf{Q}}_n(z)|_{z=e^{j\Omega}}$ and $\mathbf{a}_{\vartheta, \varphi}(e^{j\Omega}) = \mathbf{a}_{\vartheta, \varphi}(z)|_{z=e^{j\Omega}}$. Thus, the SSP-MUSIC algorithm evaluates the reciprocal of the norm of this vector,

$$P_{\text{SSP}}(\vartheta, \varphi, e^{j\Omega}) = \frac{1}{\tilde{\mathbf{a}}_{\vartheta, \varphi}(z) \mathbf{Q}_n(z) \tilde{\mathbf{Q}}_n(z) \mathbf{a}_{\vartheta, \varphi}(z)} \Big|_{z=e^{j\Omega}}, \quad (7)$$

which is large when $\mathbf{a}_{\vartheta, \varphi}(e^{j\Omega})$ is a steering vector of a contributing source. In addition to spatial location of sources in terms of ϑ and φ , $P_{\text{SSP}}(\vartheta, \varphi, e^{j\Omega})$ can determine over which frequency range sources in the direction defined by the steering vector $\mathbf{a}_{\vartheta, \varphi}(z)$ are active.

III. DIVIDE-AND-CONQUER SEQUENTIAL MATRIX DIAGONALISATION

This section outlines the components of the divide-and-conquer sequential matrix diagonalisation (DC-SMD) PEVD algorithm [11]. Following an overview of DC-SMD in Sec. III-A, Sec. III-B and Sec. III-C explain the key stages of this algorithm by detailing the divide and conquer steps, respectively.

A. Divide-and-Conquer Sequential Matrix Diagonalisation

The DC-SMD algorithm approximates the PEVD using a series of elementary paraunitary operations to iteratively diagonalise a parahermitian matrix $\mathbf{R}(z) \in \mathbb{C}^{M \times M}$ and its associated coefficient matrix, $\mathbf{R}[\tau]$. Similarly to other PEVD algorithms, DC-SMD generates an output diagonal matrix $\Lambda(z)$ containing eigenvalues, and a paraunitary matrix $\mathbf{Q}(z)$ containing eigenvectors, which satisfy (5).

While traditional PEVD algorithms — such as SMD [7] — attempt to diagonalise an entire $M \times M$ parahermitian matrix at once, the DC-SMD algorithm first divides the matrix into a number of smaller, independent parahermitian matrices, before diagonalising — or conquering — each matrix separately. For example, a matrix $\mathbf{R}(z) \in \mathbb{C}^{20 \times 20}$ might be brought into block-diagonal form comprising of four 5×5 parahermitian




Fig. 1: (a) Original matrix $\mathbf{R}[\tau] \in \mathbb{C}^{20 \times 20}$, (b) segmented result $\mathbf{R}'[\tau]$, and (c) diagonalised output $\Lambda[\tau]$. N_R , $N_{R'}$, and N_Λ are the maximum lags for matrices $\mathbf{R}[\tau]$, $\mathbf{R}'[\tau]$, and $\Lambda[\tau]$, respectively.

matrices, each of which can be diagonalised independently. Fig. 1 shows the state of the parahermitian matrix at each stage of the process for this example.

If matrix $\mathbf{R}(z)$ is of spatial dimension greater than $\hat{M} \times \hat{M}$ — where \hat{M} is an arbitrary user-defined value — an algorithm named sequential matrix segmentation (SMS) [11] is used to recursively divide the matrix into multiple independent parahermitian matrices. Each parahermitian matrix is then diagonalised in sequence through the use of the SMD algorithm. If $M \leq \hat{M}$, the divide step is skipped, and the input matrix is processed via SMD. A row-shift truncation scheme from [12], [13] is used to shorten the paraunitary matrix output by DC-SMD using a truncation parameter μ_{RST} .

The individual steps of DC-SMD are summarised in more detail in [11].

B. Recursive Polynomial Matrix Segmentation

If $\mathbf{R}(z)$ has spatial dimension $M > \hat{M}$, the divide stage of DC-SMD comes into effect. This stage applies sequential matrix segmentation (SMS) [11] to recursively divide $\mathbf{R}(z)$ into multiple independent parahermitian matrices. SMS is a novel variant of SMD designed to segment an input matrix $\hat{\mathbf{R}}(z) \in \mathbb{C}^{M' \times M'}$ into two independent parahermitian matrices $\hat{\mathbf{R}}_{11}(z) \in \mathbb{C}^{(M'-P) \times (M'-P)}$ and $\hat{\mathbf{R}}_{22}(z) \in \mathbb{C}^{P \times P}$, and two matrices $\hat{\mathbf{R}}_{12}(z) \in \mathbb{C}^{(M'-P) \times P}$ and $\hat{\mathbf{R}}_{21}(z) \in \mathbb{C}^{P \times (M'-P)}$, where $\hat{\mathbf{R}}_{12}(z) = \hat{\mathbf{R}}_{21}(z)$ are approximately zero.

The divide step of DC-SMD operates recursively. In the first recursion, the matrix $\hat{\mathbf{R}}(z)$ input to SMS is equal to $\mathbf{R}(z)$ and $M' = M$. The output matrix $\hat{\mathbf{R}}_{22}(z)$ is stored and subsequently diagonalised during the conquer step. If the second output matrix $\hat{\mathbf{R}}_{11}(z)$ is of spatial dimension greater than $\hat{M} \times \hat{M}$, the second recursion of the divide step uses $\hat{\mathbf{R}}_{11}(z)$ as the input to SMS, and M' is set equal to $M - P$. Recursions continue in this fashion until $(M' - P) \leq \hat{M}$. The dimensions of the smaller matrix produced during division, P , is forced to satisfy $P \leq \hat{M}$.

SMS iteratively minimises the energy in select regions of a parahermitian matrix in an attempt to segment the matrix. Fig. 2 illustrates the segmentation process for $M' = 5$ and $P = 2$.

The SMS algorithm continues operating until I_D iterations have been executed, or when the energy in the targeted regions, $E(\hat{\mathbf{R}}_{12}(z)) + E(\hat{\mathbf{R}}_{21}(z))$, falls below a threshold $2\delta E(\hat{\mathbf{R}}(z))$. Here, δ is some arbitrary value, and $E(\cdot)$ computes the energy in a polynomial matrix.

A parameter μ is used to truncate the parahermitian and paraunitary matrices generated in SMS. More detail on the implementation of this truncation can be found in [14].




Fig. 2: (a) Original matrix $\hat{\mathbf{R}}[\tau] \in \mathbb{C}^{5 \times 5}$, (b) regions (red) to be iteratively driven to zero in SMS for $P = 2$, and (c) segmented result. $N_{\hat{R}}$ and $N_{\hat{R}'}$ are the maximum lags for the original and segmented matrices, respectively.

C. Independent Conquering of Divided Polynomial Matrices

At this stage of DC-SMD, $\mathbf{R}(z)$ has been segmented into multiple independent parahermitian matrices. Each matrix can now be diagonalised individually through the use of the SMD algorithm [7]. Each instance of SMD is provided with a parameter I_C (which defines the maximum possible number of algorithm iterations), a truncation parameter μ , and a convergence parameter ϵ .

IV. RESULTS

This section first defines the performance metrics for evaluating the use of PEVD algorithms within SSP-MUSIC in Sec. IV-A, before setting out two simulation scenarios in Sec. IV-B, over which these metrics will be measured. The first scenario is designed to compare the use of DC-SMD versus SBR2 [3] and SMD [7] within SSP-MUSIC, and the second is formulated to analyse the performance trade-offs of the DC-SMD algorithm in this application. These scenarios are investigated in Sec. IV-C and Sec. IV-D, respectively.

A. Performance Metrics

Since the PEVD algorithms discussed in this paper all iteratively minimise off-diagonal energy, a suitable metric E_{norm} , defined in [7], is used; this metric divides the off-diagonal energy at each iteration of each algorithm by the total energy. During computation of E_{norm} , squared covariance terms are used; therefore a logarithmic notation of $5 \log_{10} E_{\text{norm}}$ is employed.

When truncation is employed, the eigenvectors and eigenvalues output from PEVD algorithms are only able to approximately reconstruct the input matrix. DC-SMD also introduces a segmentation error in its divide step, due to imperfect segmentation in SMS, which is higher for a larger threshold δ . A metric capable of measuring the difference between the original and reconstructed matrices is the mean squared error

$$\text{MSE} = \frac{1}{M^2 L'} \sum_{\tau} \|\mathbf{E}_R[\tau]\|_F^2, \quad (8)$$

where $\mathbf{E}_R[\tau] = \bar{\mathbf{R}}[\tau] - \mathbf{R}[\tau] \forall \tau$, $\bar{\mathbf{R}}(z) = \mathbf{Q}(z)\mathbf{\Lambda}(z)\tilde{\mathbf{Q}}(z)$, and L' is the length of $\mathbf{E}_R(z)$. Note that the length of a polynomial matrix $\mathbf{R}(z)$ is given by the number of lags in its corresponding space-time covariance matrix $\mathbf{R}[\tau]$. Define the paraunitarity (PU) error as

$$\eta = \frac{1}{M} \sum_{\tau} \|\mathbf{E}_Q[\tau] - \mathbf{I}_M\|_F^2, \quad (9)$$

where $\mathbf{E}_Q(z) = \mathbf{Q}(z)\tilde{\mathbf{Q}}(z)$, and \mathbf{I}_M is an $M \times M$ identity matrix. The paraunitary matrix $\mathbf{Q}(z)$ output by a PEVD algorithm is used in SSP-MUSIC; a useful metric for gauging the implementation cost of $\mathbf{Q}(z)$ is its length, L_Q .

B. Simulation Scenarios

In the simulations below, an $M = 12$ element array is illuminated by six broadband sources active over a frequency range $\Omega_i \in [0.1\pi, 0.9\pi]$ with $\varphi_i = 0$, $i = 1 \dots 6$ and different angles of arrival $\vartheta_i \in \{\pm 22.5^\circ; \pm 45^\circ; \pm 63^\circ\}$. The array signals are corrupted by uncorrelated independent and identically distributed complex Gaussian noise at 20 dB SNR. To exclude error sources other than inaccuracies in the subspace identification, the data is modelled as a sum of closely spaced sinusoids, with randomised phases, of length 64000 samples, for whom highly accurate narrowband steering vectors can be used to simulate the data. Space-time covariance matrix $\mathbf{R}[\tau]$ is estimated for $|\tau| \leq 20$. Broadband steering vectors $\mathbf{a}_{\vartheta, \varphi}(z)$ used to scan the noise-only subspace $\mathbf{Q}_n(z)$ are based on fractional delay filters constructed from truncated sinc functions, which can be substantially improved by applying a tapered window [15].

At every iteration step of each PEVD algorithm in the simulation scenarios below, the diagonalisation metric defined in Sec. IV-A is recorded together with the elapsed execution time. The MSE and paraunitarity metrics defined in (8) and (9), and the length of the paraunitary matrix $\mathbf{Q}(z)$ are recorded upon algorithm completion.

Simulations are performed within Matlab R2014a under Ubuntu 16.04 on an MSI GE60-2OE with Intel[®] Core[™] i7-4700MQ 2.40 GHz \times 8 cores and 8 GB RAM.

1) *Scenario 1* : Results in [6] demonstrate that the diagonalisation achieved by the PEVD algorithm paired with SSP-MUSIC is critically important to the quality of AoA estimation; i.e., a $\mathbf{\Lambda}(z)$ that is closer to being diagonal results in a better estimation. In [11], [12], it is shown that DC-SMD is capable of outperforming SMD [7] — particularly for parahermitian matrices of large spatial dimension M — in terms of execution time and complexity. This scenario compares the performance of SSP-MUSIC when paired with the SBR2 [3], SMD, and DC-SMD algorithms.

During iterations of each PEVD algorithm, a convergence parameter of $\epsilon = 5 \times 10^{-3}$ and truncation parameters of $\mu = 10^{-6}$ and $\mu_{\text{RST}} = 10^{-3}$ are used. Information regarding the specific implementation of these parameters within SBR2 and SMD can be found in their respective papers. DC-SMD is executed with input parameters $I_D = 250$, $I_C = 100$, $\delta = 10^{-4}$, $P = 4$, and $\hat{M} = 4$. SMD and SBR2 are run until their execution times match the time taken for the completion of DC-SMD.

2) *Scenario 2* : The DC-SMD algorithm has a number of input parameters that impact convergence. This scenario examines the effect that changing these parameters has on AoA estimation accuracy, algorithm speed, MSE, η , and L_Q .

DC-SMD is executed with input parameters $I_D = 250$, $I_C = 100$, $\epsilon = 5 \times 10^{-3}$, $\mu = 10^{-6}$, and $\mu_{\text{RST}} = 10^{-3}$. The DC-SMD threshold parameter is varied over the range $\delta \in [10^{-5}, 10^{-3}]$ for $P = 4$ and $\hat{M} = 4$, and the division parameter is varied over $P = 2, 3, 4, 6$ for $\hat{M} = P$ and $\delta = 5 \times 10^{-4}$. Execution of DC-SMD is halted when the diagonalisation level reaches $5 \log_{10} E_{\text{norm}} = -15$ dB, to allow direct comparison of results.

C. Comparison with Existing Methods

Obtained for Scenario 1, the plot of Fig. 3 shows the SSP-MUSIC performance at a frequency of $\Omega = 2\pi/3$; from this, it




Fig. 3: Performance of SSP-MUSIC at $\Omega = 2\pi/3$ based on DC-SMD, SBR2, and SMD for a scenario with six independent broadband sources in simulation scenario 1.

TABLE I: Diagonalisation, L_Q , MSE, and η for DC-SMD, SBR2, and SMD in simulation scenario 1.

Algorithm	$5 \log_{10} E_{\text{norm}}$	L_Q	MSE	η
DC-SMD	-17.17	297	1.27×10^{-5}	1.83×10^{-3}
SBR2	-8.035	221	7.19×10^{-9}	4.91×10^{-5}
SMD	-12.87	170	1.15×10^{-7}	1.94×10^{-4}

can be observed that DC-SMD offers superior localisation of the sources. In addition, the results of Fig. 4(a-c) demonstrate that the DC-SMD algorithm is capable of outperforming SBR2 and SMD at most frequencies when each is paired with SSP-MUSIC. This is due to the higher diagonalisation that DC-SMD is able to achieve during the time-limited simulation scenario. Note that each PEVD algorithm was executed for approximately 0.75 seconds.

Tab. I compares the metrics attributed to each decomposition for this simulation scenario. It can be observed that while DC-SMD obtains higher diagonalisation, it also results in higher paraunitary filter length, MSE, and paraunitarity error when compared with SBR2 and SMD.

D. Performance Trade-Offs of DC-SMD

For Scenario 2, Fig. 5 demonstrates the impact of varying the threshold δ for constant P and \hat{M} . For $\delta \leq 0.25 \times 10^{-3}$, all metrics remain reasonably constant. For $\delta > 0.25 \times 10^{-3}$, MSE increases and η decreases with increasing δ , and execution time decays exponentially with increasing δ . No significant trend can be seen between δ and L_Q , though higher δ appears to result in lower L_Q . The plots of Fig. 4(d-f) demonstrate that increasing δ results in poorer AoA estimation; this is likely tied to the increase in MSE observed for large δ . Through an informed choice of δ , a trade-off between MSE, execution time, η , and AoA estimation performance can therefore be reached.

The results of Fig. 4(g-i) show how AoA estimation performance is impacted by varying division parameter P for $\hat{M} = P$ and constant δ . Clearly, better performance is observed for larger P . In Fig. 6, it can be seen that MSE decreases with increasing P , while execution time, η , and L_Q increase. The lower MSE associated with large P again corresponds to increased AoA estimation performance. When choosing a value of P for implementation of DC-SMD, these results indicate that a trade-off between MSE, execution time, η , L_Q , and AoA estimation performance exists. Note that the selection $\hat{M} = P$ is guaranteed to fulfil the condition of $P \leq \hat{M}$ established in Sec. III-B.

V. CONCLUSION

In this paper, we have analysed the performance of a recently developed PEVD algorithm, DC-SMD, in a broadband

AoA estimation scenario. Simulation results have demonstrated that when paired with the SSP-MUSIC broadband AoA estimation algorithm, DC-SMD offers significant performance gains over traditional PEVD algorithms at the expense of increased paraunitary filter length and decomposition error. In addition, an investigation into the performance trade-offs of DC-SMD has shown that through careful choice of algorithm input parameters δ , P , and \hat{M} , a balance can be obtained between decomposition MSE, algorithm execution time, filter paraunitarity, paraunitary filter length, and AoA estimation performance. The presence of these trade-offs is important for the implementation of DC-SMD in other broadband multichannel applications.

ACKNOWLEDGEMENT

Fraser Coutts is the recipient of a Caledonian Scholarship; we would like to thank the Carnegie Trust for their support.

This work was supported in parts 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] R. O. Schmidt. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Trans. Ant. Prop.*, 34(3):276–280, Mar. 1986.
- [2] M. Alrmah, S. Weiss, and S. Lambbotharan. An extension of the music algorithm to broadband scenarios using polynomial eigenvalue decomposition. In *19th European Signal Processing Conference*, pp. 629–633, Barcelona, Spain, Aug. 2011.
- [3] J. G. McWhirter, P. D. Baxter, T. Cooper, S. Redif, and J. Foster. An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE TSP*, 55(5):2158–2169, May 2007.
- [4] S. Weiss, M. Alrmah, S. Lambbotharan, J. McWhirter, and M. Kaveh. Broadband angle of arrival estimation methods in a polynomial matrix decomposition framework. In *IEEE CAMSAP*, St. Martin, pp. 109–112, Dec. 2013.
- [5] P. Pal and P. P. Vaidyanathan. A novel autofocusing approach for estimating directions-of-arrival of wideband signals. In *Asilomar SSC*, Pacific Grove, CA, pp. 1663–1667, Nov. 2009.
- [6] M. Alrmah, J. Corr, A. Alzin, K. Thompson, and S. Weiss. Polynomial subspace decomposition for broadband angle of arrival estimation. In *SSPD*, Edinburgh, UK, pp. 1–5, Sept. 2014.
- [7] S. Redif, S. Weiss, and J. McWhirter. Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices. *IEEE TSP*, 63(1):81–89, Jan. 2015.
- [8] 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*, pp. 312–315, Gold Coast, Australia, June 2014.
- [9] Z. Wang, J. G. McWhirter, J. Corr, and S. Weiss. Multiple shift second order sequential best rotation algorithm for polynomial matrix EVD. In *EUSIPCO*, pp. 844–848, Nice, France, Sep. 2015.
- [10] J. Corr, K. Thompson, S. Weiss, J. G. McWhirter, and I. K. Proudler. Causality-Constrained multiple shift sequential matrix diagonalisation for parahermitian matrices. In *EUSIPCO*, pp. 1277–1281, Lisbon, Portugal, Sep. 2014.
- [11] F. Coutts, J. Corr, K. Thompson, and I. Proudler. Divide-and-Conquer Sequential Matrix Diagonalisation for Parahermitian Matrices. Submitted to *SSPD*, London, UK, Dec. 2017.
- [12] F. Coutts, K. Thompson, S. Weiss, and I. Proudler. Analysing the Performance of Divide-and-Conquer Sequential Matrix Diagonalisation for Large Broadband Sensor Arrays. Submitted to *SiPS*, Lorient, France, Oct. 2017.
- [13] J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter. Row-shift corrected truncation of paraunitary matrices for PEVD algorithms. In *EUSIPCO*, pp. 849–853, Nice, France, Sep. 2015.
- [14] J. Foster, J. G. McWhirter, and J. Chambers. Limiting the order of polynomial matrices within the SBR2 algorithm. In *IMA Int. Conf. Math. Signal Proc.*, Cirencester, UK, Dec. 2006.
- [15] J. Selva. An efficient structure for the design of variable fractional delay filters based on the windowing method. *IEEE Trans. S. P*, 56(8):3770–3775, Aug. 2008.

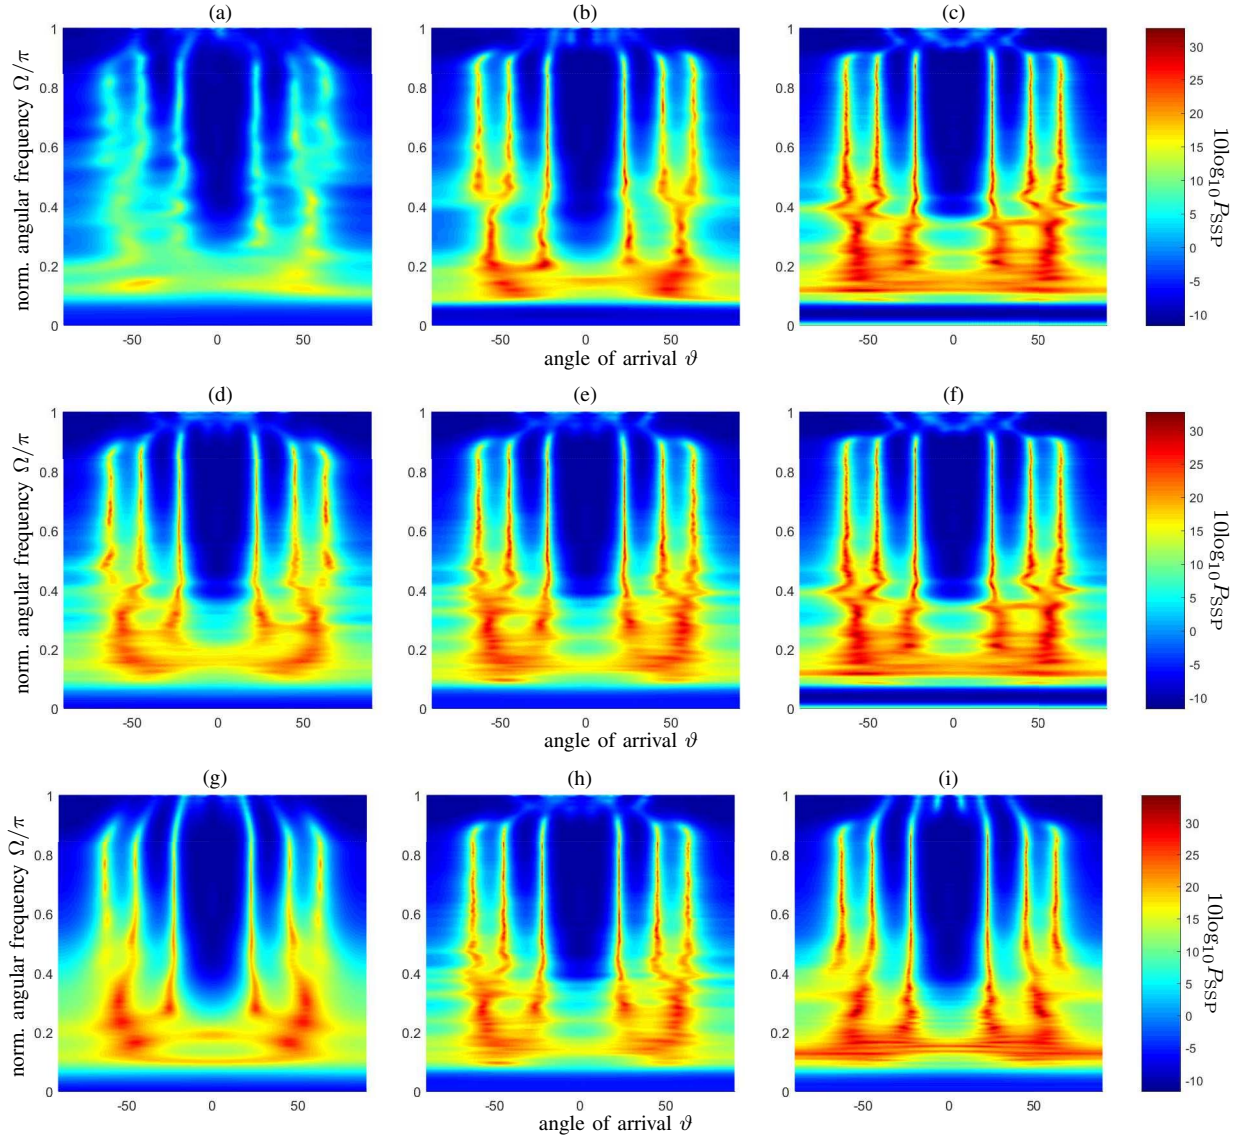


Fig. 4: Performance of SSP-MUSIC based on (a) SBR2, (b) SMD, and (c) DC-SMD for a scenario with six independent broadband sources in simulation scenario 1. Performance of SSP-MUSIC based on DC-SMD with (d) $\delta = 10^{-3}$, (e) $\delta = 5 \times 10^{-4}$, (f) $\delta = 10^{-4}$, (g) $P = 2$ and $\hat{M} = 2$, (h) $P = 4$ and $\hat{M} = 4$, and (i) $P = 6$ and $\hat{M} = 6$ for simulation scenario 2.

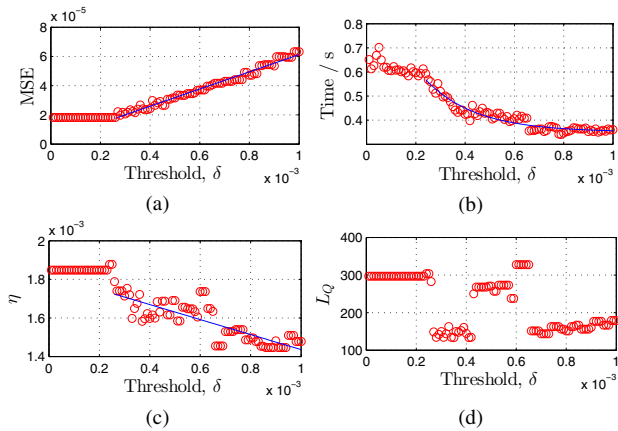


Fig. 5: (a) MSE, (b) execution time, (c) η , and (d) L_Q versus δ for DC-SMD.

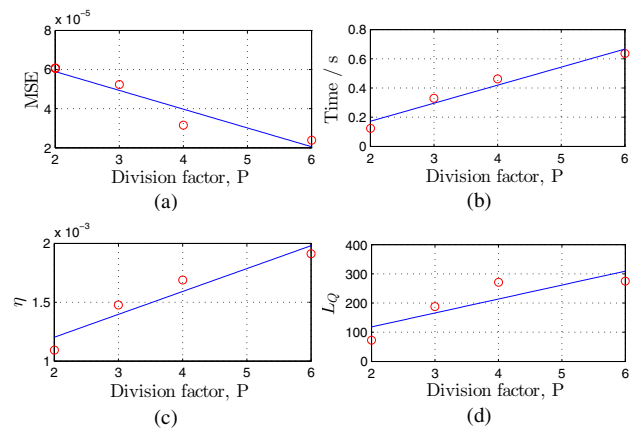


Fig. 6: (a) MSE, (b) time, (c) η , and (d) L_Q versus P for DC-SMD.