Impact of source model matrix conditioning on PEVD algorithms

J Corr[*], K Thompson[*], S Weiss[*], I K Proudler[†], J G McWhirter[‡]

[*]*Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, Scotland*
[†]*School of Electronic, Electrical and Systems Engineering, Loughborough University, Loughborough, UK*
[‡]*School of Engineering, Cardiff University, Cardiff, Wales, UK*
{jamie.corr,keith.thompson,stephan.weiss}@strath.ac.uk

Keywords: broadband array processing; space-time covariance matrix; polynomial matrix; matrix factorisation; polynomial eigenvalue decomposition.

Abstract

Polynomial parahermitian matrices can accurately and elegantly capture the space-time covariance in broadband array problems. To factorise such matrices, a number of polynomial EVD (PEVD) algorithms have been suggested. At every step, these algorithms move various amounts of off-diagonal energy onto the diagonal, to eventually reach an approximate diagonalisation. In practical experiments, we have found that the relative performance of these algorithms depends quite significantly on the type of parahermitian matrix that is to be factorised. This paper aims to explore this performance space, and to provide some insight into the characteristics of PEVD algorithms.

1. Introduction

Parahermitian polynomial matrices can compactly characterise quantities such as space-time covariance matrices in broadband array problems. Based on a data vector $\mathbf{x}[n] \in \mathbb{C}^M$, the space-time covariance $\boldsymbol{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^{\mathrm{H}}[n]\}$, with $\mathcal{E}\{\cdot\}$ the expectation operator, leads to a polynomial matrix representation for its z-transform, $\boldsymbol{R}(z) = \sum_\tau \boldsymbol{R}[\tau]z^{-\tau}$. This cross-spectral density matrix $\boldsymbol{R}(z)$ is parahermitian, i.e. $\tilde{\boldsymbol{R}}(z) = \boldsymbol{R}^{\mathrm{H}}(z^{-1}) = \boldsymbol{R}(z)$, where the parahermitian operator $\{\tilde{\cdot}\}$ performs a complex conjugate transpose and time reveral of all matrix entries.

To extend the utility of the eigenvalue (EVD) and singular value decompositions (SVD) [1] to general broadband problems, a polynomial EVD (PEVD [2–4]) has been defined. Given a parahermitian $\boldsymbol{R}(z)$, the PEVD

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

results in paraunitary factors $\boldsymbol{Q}(z)$ and a diagonal, spectrally majorised and parahermitian $\boldsymbol{\Lambda}(z)$. The latter contains the polynomial eigenvalues,

$$\boldsymbol{\Lambda}(z) = \mathrm{diag}\{\Lambda_0(z)\ \Lambda_1(z)\ \ldots\ \Lambda_{M-1}(z)\} \quad . \tag{2}$$

with spectral majorisation enforcing an ordering such that

$$\Lambda_{m+1}(e^{j\Omega}) \geq \Lambda_m(e^{j\Omega}), \ \forall\ \Omega,\ m = 0 \ldots M-1 \quad . \tag{3}$$

Paraunitarity of $\boldsymbol{Q}(z)$ implies that $\boldsymbol{Q}(z)\tilde{\boldsymbol{Q}}(z) = \tilde{\boldsymbol{Q}}(z)\boldsymbol{Q}(z) = \boldsymbol{I}$. While equality in (1) is not guaranteed, the approximation has been suggested to hold close for sufficiently high orders of $\boldsymbol{Q}(z)$ [5].

A number of PEVD algorithms have been introduced [4,6–10], and offer various performance characteristics. The algorithms in [4,6,10] have been demonstrated on parahermitian matrices $\boldsymbol{R}(z) \in \mathbb{C}^{M \times M}$ derived from arbitrary $\boldsymbol{A}(z) \in \mathbb{C}^{M \times K}$ as $\boldsymbol{R}(z) = \boldsymbol{A}(z)\tilde{\boldsymbol{A}}(z)$, with $K \geq M$ to enable full rank. In [7,8], subband coding was considered as an application, and the parahermitian matrices that need to be factorised by the algorithms are based on auto-regressive functions generating auto-correlation functions with infinite support but potentially permitting finite order paraunitary factors (for a justification, see the factorisation in Sec. IV.B.3 in [8]). In [9], a source model convolutively mixes spectrally majorised sources by means of an arbitrary paraunitary matrix, such that the ground truth PEVD with finite order factors and equality in (1) is guaranteed. Since these publications use differently conditioned problems, a direct comparison between algorithms proposed in individual papers is not always straightforward.

In this paper, we generalise the source model idea in [9] to carefully control of the conditioning of the parahermitian matrix. This includes a definition of the dynamic range of the underlying source, which can be linked to the condition number or eigenvalue spread of a parahermitian matrix by generalisation from the field of non-polynomial matrices. Besides the dynamic range, we also define different relation between the sources' PSDs. These may be

- not spectrally majorised (i.e. with overlapping PSDs);
- spectrally majorised, with '≥' in (3), or
- strictly spectrally majorised, with '>' in (3).

A ensemble of randomised parahermitian matrices with different dynamic ranges and types of majorisations are factorised by a number of PEVD algorithms belonging to the second order sequential best rotation (SBR2, [4,8] and the sequential matrix diagonalisation (SMD, [9,10]) families.

In the following, Sec. 2 briefly details the PEVD algorithms belonging to the SBR2 and SMD families. Sec. 3 shows the

impact of the type of majorisation on the order to the factors of the ground-truth PEVD, and introduces the condition number as a metric for the dynamic range of a parahermitian matrix. Experimental results for applying the various PEVD algorithms to differently conditioned parahermitian matrices are discussed in Sec. 4, followed by conclusions in Sec. 5.

2. PEVD Algorithms

2.1 General Anatomy

The current most popular PEVD algorithms [4,8–10] have the goal of diagonalising a parahermitian matrix $\boldsymbol{R}(z)$ starting from an initial approximation $\boldsymbol{S}^{(0)}(z)$. The ith iteration of all algorithms consists of three common steps operating on $\boldsymbol{S}^{(i-1)}(z)$, which vary with implementation.

In the first step of the i-th iteration, the remaining off-diagonal elements of the parahermitian matrix $\boldsymbol{S}^{(i-1)}(z)$ are searched. Part of the off-diagonal energy is then transferred onto the zero lag in the second step using a paraunitary shift matrix,

$$\boldsymbol{S}^{(i)\prime}(z) = \tilde{\boldsymbol{\Lambda}}^{(i)}(z)\boldsymbol{S}^{(i-1)}(z)\boldsymbol{\Lambda}^{(i)}(z) , \quad i = 1\ldots I \quad . \quad (4)$$

The search step and therefore the construction of the shift matrix, $\boldsymbol{\Lambda}^{(i)}(z)$, depend on the particular PEVD implementation to be detailed below. The final step in in the ith iteration is to bring the off-diagonal energy, found in step one and shifted in step two, onto the diagonal of the zero lag matrix. This is accomplished by means of a unitary matrix, $\mathbf{Q}^{(i)}$, which is applied to all lags in the parahermitian, $\boldsymbol{S}^{(i)\prime}(z)$, such that

$$\boldsymbol{S}^{(i)}(z) = \mathbf{Q}^{(i)\mathrm{H}}\boldsymbol{S}^{(i)\prime}(z)\mathbf{Q}^{(i)} \quad . \quad (5)$$

Like the shift matrix, $\boldsymbol{\Lambda}^{(i)}(z)$, the construction of the unitary energy transfer matrix, $\mathbf{Q}^{(i)}$, depends on the specific PEVD algorithm.

The PEVD algorithm is complete when either a set number of iterations, I, have been carried out or the search step returns an amount of energy lower than a predefined threshold. Upon completion the PEVD algorithm returns the approximate polynomial eigenvalues in the diagonalised parahermitian $\boldsymbol{S}^{(I)}(z)$ and the approximate polynomial eigenvectors in $\boldsymbol{Q}^{(I)}(z)$. The polynomial eigenvectors are simply the product of the unitary energy transfer matrices, $\mathbf{Q}^{(i)}$, and paraunitary shift matrices, $\boldsymbol{\Lambda}^{(i)}(z)$, from each of the I iterations,

$$\boldsymbol{Q}^{(I)}(z) = \prod_{i=1}^{I} \mathbf{Q}^{(i)}\boldsymbol{\Lambda}^{(i)}(z) \quad . \quad (6)$$

To reduce the computational cost of applying the paraunitary matrix, $\boldsymbol{Q}^{(I)}(z)$, a paraunitary trim function is used to significantly reduce the polynomial order of $\boldsymbol{Q}^{(I)}(z)$. In this paper we use the recently developed row-shift corrected truncation method [11], this approach takes advantage of an ambiguity in the paraunitary matrix to further reduce its polynomial order.

2.2 Second Order Sequential Best Rotation

With the initialisation $\boldsymbol{S}^{(0)}(z) = \boldsymbol{R}(z)$, the first step of the SBR2 algorithm [4] at the ith iteration utilises a search for the off-diagonal element with the largest modulus,

$$\{k^{(i)}, \tau^{(i)}\} = \arg\max_{k,\tau} \|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_\infty , \quad i = 1\ldots I , \quad (7)$$

where the modified column vector, $\hat{\mathbf{s}}_k^{(i-1)}[\tau]$, contains all elements apart from the on-diagonal entries. Based on the column and lag indices, $k^{(i)}$ and $\tau^{(i)}$ respectively, the paraunitary shift matrix, $\boldsymbol{\Lambda}^{(i)}(z)$, is then generated as

$$\boldsymbol{\Lambda}^{(i)}(z) = \mathrm{diag}\{\underbrace{1 \ldots 1}_{k^{(i)}-1} \; z^{-\tau^{(i)}} \; \underbrace{1 \ldots 1}_{M-k^{(i)}}\} \quad . \quad (8)$$

The maximum element found in (7) and shifted onto the zero lag using (8) is transferred onto the diagonal using a Jacobi transformation for $\mathbf{Q}^{(i)}$ in (5). The sparse nature of the Jacobi transformation means that rather applying a full matrix multiplication to each lag in the parahermitian matrix, only two rows and columns of $\boldsymbol{S}^{(i)\prime}(z)$ are affected across all its lags.

2.3 Sequential Matrix Diagonalisation

The SMD algorithm [9] includes an initialisation step which diagonalises the zero lag of the parahermitian matrix,

$$\mathbf{S}^{(0)}[0] = \mathbf{Q}^{(0)\mathrm{H}}\mathbf{R}[0]\mathbf{Q}^{(0)} \quad . \quad (9)$$

In (9) the unitary matrix, $\mathbf{Q}^{(0)}$, is the modal matrix from the EVD of $\mathbf{R}[0]$ which brings all the energy in the zero lag onto the diagonal, zeroing the off-diagonal elements. As with $\mathbf{Q}^{(i)}$ in (5), $\mathbf{Q}^{(0)}$ is applied to all lags of the parahermitian matrix, such that $\boldsymbol{S}^{(0)}(z) = \mathbf{Q}^{(0)\mathrm{H}}\boldsymbol{R}(z)\mathbf{Q}^{(0)}$.

The i-th iteration of the SMD algorithm starts with the search for the maximum column norm,

$$\{k^{(i)}, \tau^{(i)}\} = \arg\max_{k,\tau} \|\hat{\mathbf{s}}_k^{(i-1)}[\tau]\|_2 , \quad i = 1\ldots I . \quad (10)$$

Using the l_2 norm differs from (7), which extracts the maximum element using the l_∞ norm. Like SBR2, the shift step in the SMD approach uses (8) to construct the paraunitary shift matrix $\boldsymbol{\Lambda}^{(i)}(z)$.

Rather than transferring the energy from a single element onto the diagonal like SBR2, the SMD algorithm uses the modal matrix of the EVD of the zero lag, similar to (9), to construct $\mathbf{Q}^{(i)}$ and transfer all the zero lag energy onto the diagonal. Typically the SMD algorithm will transfer a greater amount of energy than SBR2 at each iteration. The main drawback of the SMD algorithm is the cost associated with applying the non-sparse $\mathbf{Q}^{(i)}$ to the entire parahermitian matrix at each iteration.

With the addition of the initialisation step the calculation of the paraunitary matrix, $\boldsymbol{Q}^{(I)}(z)$, in (6) must be modified to

$$\boldsymbol{Q}^{(I)}(z) = \mathbf{Q}^{(0)} \prod_{i=1}^{I} \mathbf{Q}^{(i)}\boldsymbol{\Lambda}^{(i)}(z) \quad , \quad (11)$$

in order to reflect the initial diagonalisation by $\mathbf{Q}^{(0)}$.

2.4 Multiple Shift Maximum Element SMD

The MSME-SMD algorithm [10] employs the same initialisation step as the SMD algorithm above to bring the zero lag energy onto the diagonal. At each iteration, the SMD's l_2 column norm search is replaced by a maximum element search as in (7). Whereas the SMD algorithm immediately diagonalises the energy brought onto the zero lag matrix, the MSME-SMD algorithm uses a set of reduced search spaces, detailed in [10], to bring a total of $M-1$ maximum elements onto the zero lag at each iteration, where M is the spatial dimension of the parahermitian matrix. The reduced search spaces have a dual purpose: firstly they ensure that previous maxima transferred onto the zero lag are not undone by subsequent shifts; secondly they are designed to guarantee a total of $M-1$ elements are brought onto the zero lag at each iteration.

To bring the $M-1$ maximum elements onto the zero lag, the paraunitary delay matrix, $\mathbf{\Lambda}^{(i)}(z)$, must be modified to be

$$\mathbf{\Lambda}^{(i)}(z) = \text{diag}\left\{ z^{-\tau_1^{(i)}} \ z^{-\tau_2^{(i)}} \ \ldots \ z^{-\tau_M^{(i)}} \right\} \quad . \quad (12)$$

The paraunitary delay matrix in (12) allows each of the M rows and columns of the parahermitian matrix to be advanced or delayed.

The MSME-SMD algorithm uses the same technique as SMD to transfer energy from all the elements shifted onto the zero lag onto the diagonal. Using the multiple shifts the MSME-SMD algorithm will bring more energy onto the zero lag at each iteration than the SMD equivalent. The computational cost of one MSME-SMD iteration is slightly higher than SMD but with the high cost of applying the EVD modal matrix, $\mathbf{Q}^{(i)}$, to all parahermitian lags the higher energy transfer of MSME-SMD is more beneficial for real time convergence. A drawback of MSME-SMD compared to both SBR2 and SMD is that the order of the paraunitary and parahermitian matrices associated with the PEVD will grow faster.

3. Source Model Conditioning

For the analysis and simulations in this paper, we assume that the parahermitian matrices have a known ground truth decomposition. This enables us to control the condition of the problem that is addressed by the various PEVD algorithms, and also assess and compare the solution that is reached.

3.1 Source Model

The general model is depicted in Fig. 1. A total of L independent source signals with individual power spectral densities (PSDs) $F_l(z)\tilde{F}_l(z)$, $l = 1 \ldots L$, are generated by exciting innovation filters $F_l(z)$ with unit variance zero-mean uncorrelated complex Gaussian sources $u_l[n]$ [12]. The order of the innovation filters $F_l(z)$ is P, and careful control of the filter gain and the maximum radius of zeros can determine the dynamic range of the source and whether they e.g. are spectrally majorised as in [9]. Convolutive mixing of the source signals is performed by a random paraunitary matrix $\boldsymbol{A}(z) \in$
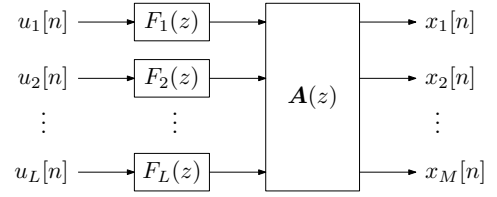
Figure 1. Source model with L unit variance zero mean uncorrelated complex Gaussian excitations $u_l[n]$, innovation filters with transfer functions $F_l(z)$, $l = 1 \ldots L$, followed by a paraunitary convolutive mixing system $\boldsymbol{A}(z)$.

$\mathbb{C}^{M \times L}$ of order K, with $M \geq L$. This matrix is determined by extracting L columns from

$$\boldsymbol{A}'(z) = \prod_{p=1}^{K} (\mathbf{I} - \mathbf{v}_k \mathbf{v}_k^{\text{H}} + \mathbf{v}_k \mathbf{v}_k^{\text{H}} z^{-1}) , \quad (13)$$

which is a product of K elementary paraunitary matrices [3], with $\mathbf{v}_k$, $k = 1 \ldots P$, being random unit norm vectors.

The space-time covariance matrix constructed from the output $\mathbf{x}^{\text{T}}[n] = [x_1[n] \ \ldots x_M[n]]$ is therefore given as

$$\boldsymbol{R}(z) = \sum_{\tau} \mathcal{E}\left\{\mathbf{x}[n]\mathbf{x}^{\text{H}}[n-\tau]\right\} z^{-\tau} \quad (14)$$

$$= \boldsymbol{A}(z)\boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)\tilde{\boldsymbol{A}}(z) . \quad (15)$$

The diagonal matrix $\boldsymbol{F}(z) = \text{diag}\{F_1(z) \ldots F_L(z)\}$ contains the L innovation filters.

3.2 Polynomial Eigenvalue Decomposition

Given that the parahermitian matrix in (15) is factorised into paraunitary and diagonal parahermitian matrices, it bears close relation with the PEVD (1) of $\boldsymbol{R}(z)$. If $\boldsymbol{F}(z)$ is spectrally majorised, then indeed the PEVD $\boldsymbol{R}(z) = \boldsymbol{Q}(z)\mathbf{\Lambda}(z)\tilde{\boldsymbol{Q}}(z)$ exists with equality and is given by $\boldsymbol{Q}(z) = \boldsymbol{A}(z)$ and $\mathbf{\Lambda}(z) = \boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)$.

If $\boldsymbol{F}(z)$ is not spectrally majorised, then a PEVD satisfying both diagonalisation and spectral majorisation could be derived by dividing and re-assigning spectral components of $\boldsymbol{F}(z)$ via a paraunitary matrix $\boldsymbol{U}(z)$ such that $\boldsymbol{U}(z)\boldsymbol{F}(z)$ is spectrally majorised. Then $\mathbf{\Lambda}(z) = \boldsymbol{U}(z)\boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)\tilde{\boldsymbol{U}}(z)$, and $\boldsymbol{U}(z)$ can be absorbed into $\boldsymbol{A}(z)$ to yield the polynomial modal matrix $\boldsymbol{Q}(z) = \boldsymbol{A}(z)\tilde{\boldsymbol{U}}(z)$. Since an ideal $\boldsymbol{U}(z)$ providing a binary spectral mask will require infinite support, the order of the factors $\boldsymbol{Q}(z)$ and $\mathbf{\Lambda}(z)$ is likely to be much higher than in the spectrally majorised case.

Example. Let $L = M = 2$ with a diagonal $\boldsymbol{F}(z) = \text{diag}\{1 + z^{-1}; \ 1 - z^{-1}\}$ generating the unmajorised PSDs in Fig. 2. If $U_{ij}(z)$, $i, j \in \{1, 2\}$, are the elements of a matrix $\boldsymbol{U}(z)$ to enforce spectral majorisation, then $U_{11}(z)$ and $U_{22}(z)$ must be halfband lowpass and $U_{12}(z)$ and $U_{21}(z)$ halfband highpass filters. If selected as quadrature mirror filters with $\boldsymbol{U}_{21}(z) = -\tilde{\boldsymbol{U}}_{12}(z)$ and $\boldsymbol{U}_{22}(z) = \tilde{\boldsymbol{U}}_{11}(z)$, the condition of paraunitarity reduces to demand power complementarity,

$$\boldsymbol{U}_{11}(z)\tilde{\boldsymbol{U}}_{11}(z) + \boldsymbol{U}_{22}(z)\tilde{\boldsymbol{U}}_{22}(z) = 1 . \quad (16)$$

3
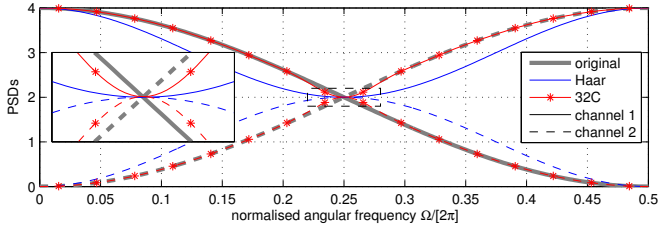
Figure 2. PSDs of unmajorised sources, and after frequency-reassignment using paraunitary matrices based on Haar [13] and 32C filters [14].
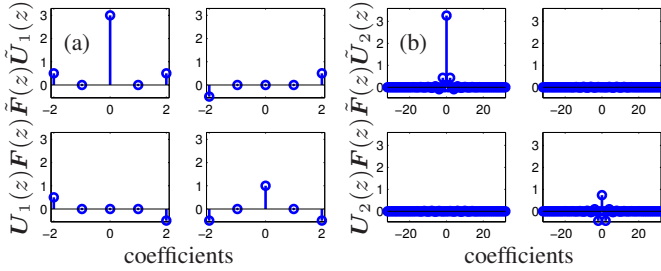
Figure 3. Approximately diagonalised matrices for paraunitary systems based on (a) Haar [13] and (b) 32C filters [14].

For $\boldsymbol{U}(z)\boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)\tilde{\boldsymbol{U}}(z)$ to retain a diagonal structure, it can be shown that

$$U_{11}(z)U_{12}(z)F_1(z)\tilde{F}_1(z) = U_{11}(z)U_{12}(z)F_2(z)\tilde{F}_2(z) \quad (17)$$

is also required. This can be achieved only if $U_{11}(z)U_{12}(z) = 0$, i.e. they are ideal, complementary, infinite length halfband lowpass and highpass filters.

Using a Haar filter [13] of order 1 to construct $\boldsymbol{U}_1(z)$, the PSDs along the diagonal are now spectrally majorised as evident from Fig. 2. However, inspecting $\boldsymbol{U}_1(z)\boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)\tilde{\boldsymbol{U}}_1(z)$ in Fig. 3(a), (17) is violated resulting in off-diagonal terms. Using the filter 32C from [14] to construct an approximately paraunitary $\boldsymbol{U}_2(z)$, the higher order of 31 now results in an approximately diagonalised $\boldsymbol{U}_2(z)\boldsymbol{F}(z)\tilde{\boldsymbol{F}}(z)\tilde{\boldsymbol{U}}_2(z)$ as demonstrated in Fig. 3(b), which is also spectrally majorised according to Fig. 2.

Therefore if sources contributing to $\boldsymbol{R}(z)$ are not spectrally majorised, a PEVD of $\boldsymbol{R}(z)$ in the sense of the definition in (1)–(3) requires higher order polynomial matrix factors than for a case where sources are spectrally majorised.

3.3 Eigenvalue Spread

Since PEVD algorithms have a stopping criterion that is tied to a threshold for off-diagonal values, the resolution of sources depends on the dynamic range of the source signals. Therefore, this dynamic range can be defined as the ratio between the maximum and minimum value across all source PSDs and frequencies,

$$\gamma = \frac{\max_{\Omega,l} |F_l(e^{j\Omega})|^2}{\min_{\Omega,l} |F_l(e^{j\Omega})|^2} . \quad (18)$$

For $M = K$, even in the case where sources are not spectrally majorised, (18) represents a polynomial matrix condition num-

ber,

$$\gamma = \frac{\max_{\Omega} \Lambda_1(e^{j\Omega})}{\min_{\Omega} \Lambda_M(e^{j\Omega})} , \quad (19)$$

as after re-assigning frequency bands between channels, the minimum and maximum PSD values remain unaltered as demonstrated in Sec. 3.2.

4. Results

The following subsections present the details of the simulation scenario followed by the performance metrics used to compare the different source models and PEVD algorithms. The final three subsections present and analyse the results of the simulations.

4.1 Performance Metrics

To assess the impact of source model conditioning on PEVD performance we use the following metrics. First the convergence of the PEVD algorithms is monitored via the normalised off-diagonal energy at the i-th iteration,

$$E_{\mathrm{norm}}^{(i)} = \frac{\sum_\tau \sum_{k=1}^M \|\hat{\mathbf{s}}_k^{(i)}[\tau]\|_2^2}{\sum_\tau \|\mathbf{R}[\tau]\|_F^2} , \quad (20)$$

where $\hat{\mathbf{s}}_k^{(i)}[\tau]$ is the same modified column vector used in (7) and the denominator consists of the sum of Frobenius norms, $\|\cdot\|_F$, of the initial parahermitian matrix, $\mathbf{R}[\tau]$, for each of the τ lags.

As well as noting $E_{\mathrm{norm}}^{(i)}$ for every iteration, the order of the truncated paraunitary matrices is recorded to show how the source model affects the growth of the paraunitary matrix, which directly represents the implementation cost of this lossless filter bank.

To compare the diagonal matrices produced by the PEVD to the ground truth of the source model we use the PSDs. Ideally the PSDs extracted by PEVD algorithms should match those of the source model, bar any frequency-reassignments in the case of spectrally unmajorised sources.

4.2 Simulation Scenarios

The first two sets of simulations present the results from 500 iterations of the PEVD algorithms outlined in Sec. 2 for the spectrally majorised and ummajorised examples over an ensemble of 10^2 random instantiations. With $L = 4$ sources acquired by $M = 4$ sensors, for each instantiation the source model produces a distinct parahermitian $\boldsymbol{R}(z) \in \mathbb{C}^{4 \times 4}$. With $P = K = 30$, the order of $\boldsymbol{R}(z)$ is 120. For each ensemble, restrictions on the radii of zeros in the innovation filters create an average dynamic range of either 10 or 20 dB.

The final set of results demonstrate example PSDs produced after 100 SMD iterations compared to the original spectrally majorised and unmajorised source models. These simulations use a single source model for each combination of majorisation and dynamic range.
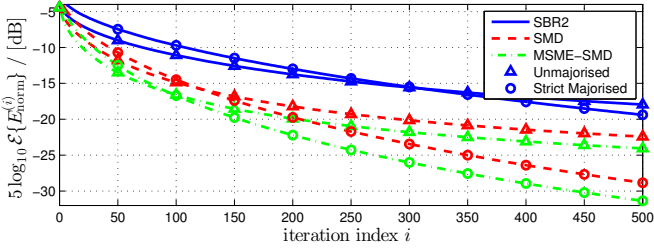
Figure 4. Reduction in off diagonal energy for different source models with a dynamic range of 10 dB for a selection of PEVD algorithms.
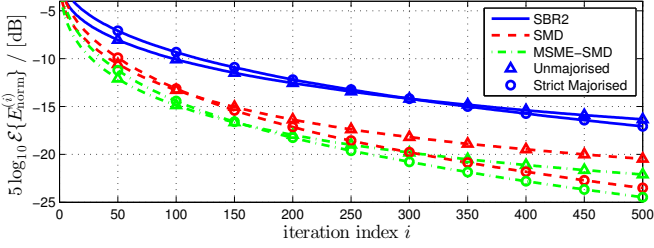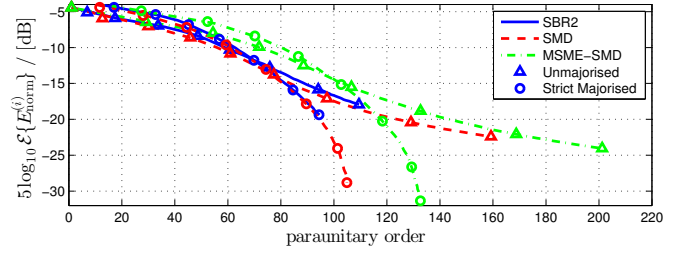

Figure 5. Reduction in off diagonal energy for different source models with a dynamic range of 20 dB for a selection of PEVD algorithms.
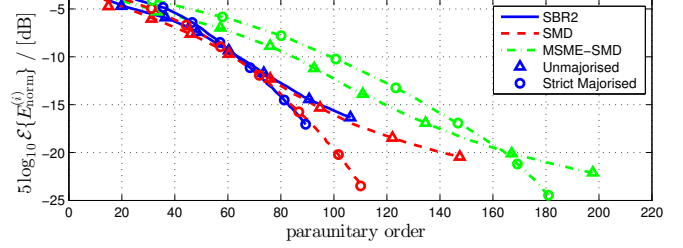

Figure 6. Paraunitary matrix order for different source models with a dynamic range of 10 dB for a selection of PEVD algorithms.

Figure 7. Paraunitary matrix order for different source models with a dynamic range of 20 dB for a selection of PEVD algorithms.

4.3 Algorithm Convergence

Figs. 4 and 5 show how the different algorithms converge for the two source models identified in Sec. 3 for dynamic ranges of 10 dB and 20 dB respectively. Across both dynamic ranges and source models the MSME-SMD algorithm offers the fastest reduction in off-diagonal energy followed by SMD and then the SBR2 algorithm. In both Figs. 4 and 5 all algorithms initially converge faster for the unmajorised source but as the number of iterations increases, these curves slow down and are overtaken by the strictly majorised sources. After 500 iterations there is a noticeable difference between the two source models, with the strictly majorised being better; this is apparent for both dynamic ranges and all three PEVD algorithms. With the higher dynamic range in Fig. 5 we can see that the curves all appear worse than their counterparts in Fig. 4 and end up closer together.

4.4 Paraunitary Order

The growth in paraunitary order for the PEVD methods using the unmajorised and strictly majorised sources at 10 dB is shown in Fig. 6 with the larger dynamic range of 20 dB depicted in Fig. 7. In both Figs. 6 and 7 the SMD and SBR2 algorithms perform similarly but the multiple shifts of the MSME-SMD algorithm cause the paraunitary order to grow faster. The paraunitary order for the MSME-SMD algorithm is also affected more when the dynamic range of the source increases. For all the algorithms over both dynamic ranges we see that the paraunitary orders for the unmajorised sources tends be higher than the strictly majorised source. The main exception to this is the MSME-SMD with the strictly majorised source where it mostly performs worse than its unmajorised equivalent.

4.5 Power Spectral Densities

This section investigates four example source models which have had the SMD algorithm applied for 100 iterations each. PSDs of the source models are shown in Figs. 8,9,10 and 11, first showing a 10 dB dynamic range for the strictly majorised source then the unmajorised equivalent followed by the same sources with a 20 dB dynamic range. Like the simple example in Fig.3 the unmajorised sources in Figs. 9 and 11 are approximately majorised by channel permutations for different frequency bands. Comparing the two types of majorisation we can see that the unmajorised sources appear to be modelled better by the SMD algorithm than the strictly majorised sources. When the dynamic range of the source is increased from 10 dB to 20 dB the SMD algorithm does not achieve the same level of accuracy.

The performance metrics studied in the previous subsections are shown in Tab. 1 for the source decompositions in Figs. 8 – 11. It is interesting to notice that for the majorised source the 20 dB PEVD has a better diagonalisation measure yet the source representation appears worse. The parameters in Tab. 1 fall very near the cross-over points in Figs. 4 – 7 so the fact that for 10 dB the unmajorised case has better diagonalisation and paraunitary order and for 20 dB has worse diagonalisation and paraunitary order is not surprising. Running the simulations over 500 iterations yields the results in brackets in Tab. 1 which match the final trends shown in Figs. 4 – 7.

5. Conclusion

This paper has investigated how the conditioning of the parahermitian matrix can affect the performance of a PEVD algorithm. Using the proposed source model, properties of the parahermitian matrix can be carefully controlled. A number of PEVD algorithms have been compared for different conditions of this source model.
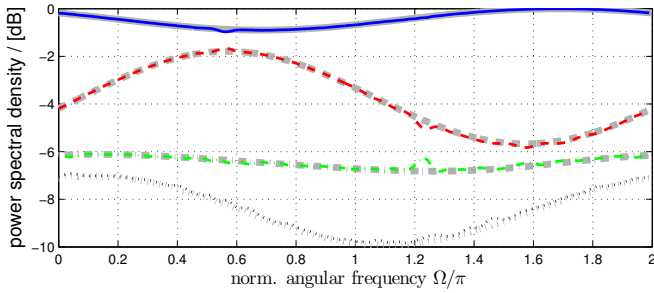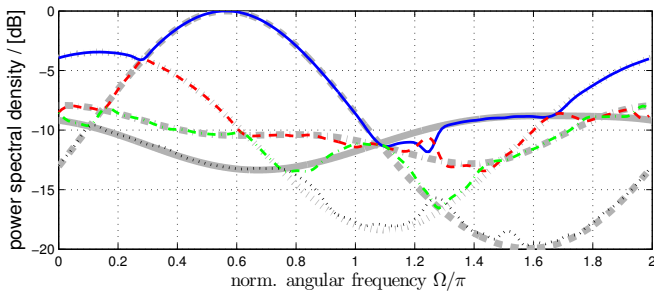
Figure 8. PSD shown for a strictly majorised source model with 10 dB dynamic range overlaid with SMD decomposition after 100 iterations.
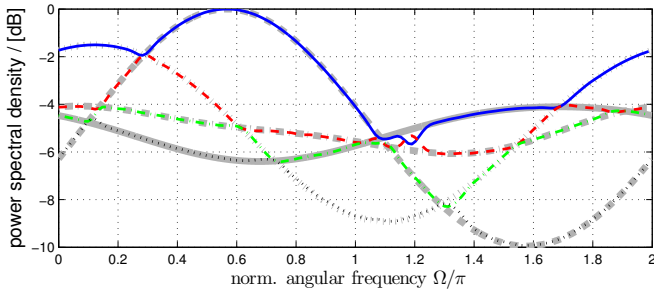
Figure 9. PSD shown for a unmajorised source model with dynamic range of 10 dB overlaid with SMD decomposition after 100 iterations.
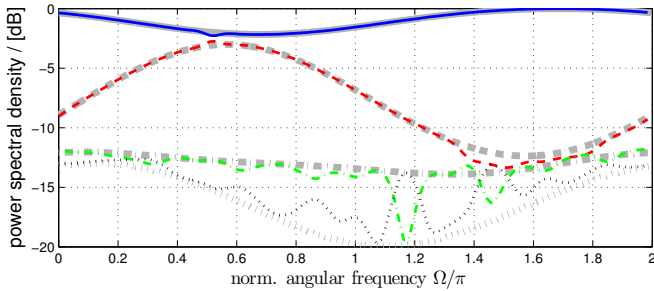
Figure 10. PSD shown for a strictly majorised source model with 20 dB dynamic range overlaid with SMD decomposition after 100 iterations.

Figure 11. PSD shown for a unmajorised source model with 20 dB dynamic range overlaid with SMD decomposition after 100 iterations.

Table 1. Performance metrics for source model PSDs after 100 (and 500) SMD iterations .

Source Model	Diag. Meas. (dB)	PU Order
Strict 10 dB	$-13.11\ (-29.90)$	88 (123)
Unmajorised 10 dB	$-14.69\ (-22.35)$	80 (151)
Strict 20 dB	$-13.31\ (-25.40)$	66 (100)
Unmjorised 20 dB	$-12.81\ (-20.18)$	84 (138)

The results show that the speed of convergence is related to the source model used, in particular the dynamic range and the ordering of the eigenvalues. From the results presented in this paper a higher dynamic range will typically cause the PEVD algorithms to converge more slowly in terms of reducing off-diagonal energy; although it has minimal affect on the paraunitary orders for SBR2 and SMD algorithms the orders in case of MSME-SMD tend to grow faster. When the ordering of the polynomial eigenvalues is changed, i.e. majorised vs. unmajorised, the ordered or majorised version will converge faster and to a better level of diagonalisation with a lower order paraunitary matrix independent of the PEVD algorithm.

Acknowledgement

References

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

[2] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, New York, 1982.

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

[4] 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*, 55(5):2158–2169, May 2007.

[5] S. Icart and P. Comon. Some properties of Laurent polynomial matrices. In *9th IMA Conference on Mathematics in Signal Processing*, Birmingham, UK, December 2012.

[6] J. G. McWhirter and P. D. Baxter. A Novel Technqiue for Broadband SVD. In *12th Annual Workshop on Adaptive Sensor Array Processing*, MIT Lincoln Labs, Cambridge, MA, 2004.

[7] A. Tkacenko and P. Vaidyanathan. Iterative greedy algorithm for solving the fir paraunitary approximation problem. *IEEE Transactions on Signal Processing*, 54(1):146–160, Jan. 2006.

[8] S. Redif, J. McWhirter, and S. Weiss. Design of FIR parauniftary filter banks for subband coding using a polynomial eigenvalue decomposition. *IEEE Transactions on Signal Processing*, 59(11):5253–5264, November 2011.

[9] S. Redif, S. Weiss, and J. McWhirter. Sequential matrix diagonalization algorithms for polynomial EVD of parahermitian matrices. *IEEE Transactions on Signal Processing*, 63(1):81–89, January 2015.

[10] 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*, pages 312–315, Gold Coast, Australia, June 2014.

[11] 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, August/September 2015.

[12] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, 3rd edition, 1991.

[13] N. J. Fliege. *Multirate Digital Signal Processing: Multirate Systems, Filter Banks, Wavelets*. John Wiley & Sons, Chichester, 1994.

[14] R. E. Crochiere and L. R. Rabiner. *Multirate Digital Signal Processing*. Prentice Hall, Englewood Cliffs, NJ, 1983.