

Algorithmic Enhancements to Polynomial Matrix Factorisations

Fraser Kenneth Coutts, Stephan Weiss

Material adapted from PhD Thesis of the same name.

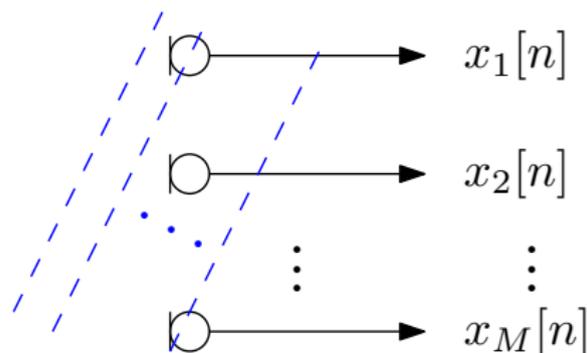
This thesis available online via EURASIP's library: <https://tinyurl.com/couttsphd>.

22nd June 2023

Presentation Overview

1. Polynomial Matrix Background & Motivation
2. Polynomial Matrix EVD (PEVD)
3. Existing Methods to Compute the PEVD
4. Algorithmic Improvements to Existing Methods
5. Novel Algorithms and Approaches
6. Conclusion

Summary of Background



- ▶ Cross-spectral density

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

is a polynomial matrix.

- ▶ Parahermitian:

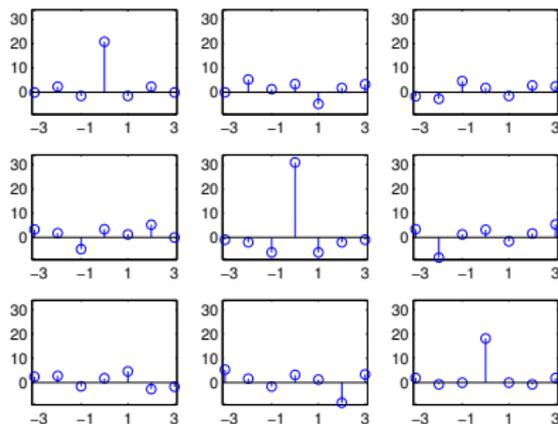
$$\mathbf{R}^P(z) = \mathbf{R}^H(1/z^*) = \mathbf{R}(z)$$

- ▶ Space-time covariance matrix:

$$\mathbf{R}[\tau] = \mathcal{E}\{\mathbf{x}[n]\mathbf{x}^H[n-\tau]\}$$

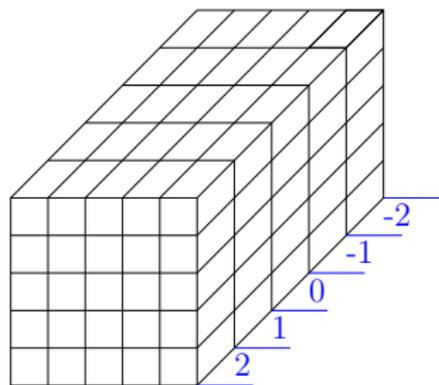
- ▶ Matrix of auto- & cross-correlation sequences

- ▶ Symmetry $\mathbf{R}[\tau] = \mathbf{R}^H[-\tau]$



Summary of Background

- ▶ $\mathbf{R}(z)$ is a matrix with (Laurent) polynomial entries or alternatively a polynomial with matrix-valued coefficients.
- ▶ Can be interpreted as a three-dimensional matrix.



$$\mathbf{R}(z) = \begin{bmatrix} R_{11}(z) & R_{12}(z) & \cdots & R_{1M}(z) \\ R_{21}(z) & R_{22}(z) & & \vdots \\ \vdots & & \ddots & \vdots \\ R_{M1}(z) & \cdots & \cdots & R_{MM}(z) \end{bmatrix}$$

$$= \sum_{\tau} \mathbf{R}[\tau] z^{-\tau}$$

Summary of Background

- ▶ Eigenvalue decomposition (EVD) of a Hermitian covariance matrix \mathbf{R} offers optimality for many narrowband problems:

$$\mathbf{R} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^H$$

$\mathbf{\Lambda}$ is diagonal

\mathbf{Q} is unitary, i.e. $\mathbf{Q}\mathbf{Q}^H = \mathbf{I}$

- ▶ The polynomial matrix EVD (PEVD) is an extension to parahermitian matrices $\mathbf{R}(z)$:

$$\mathbf{R}(z) \approx \mathbf{Q}(z)\mathbf{\Lambda}(z)\mathbf{Q}^P(z) \quad \mathbf{\Lambda}(z) \text{ is diagonal}$$

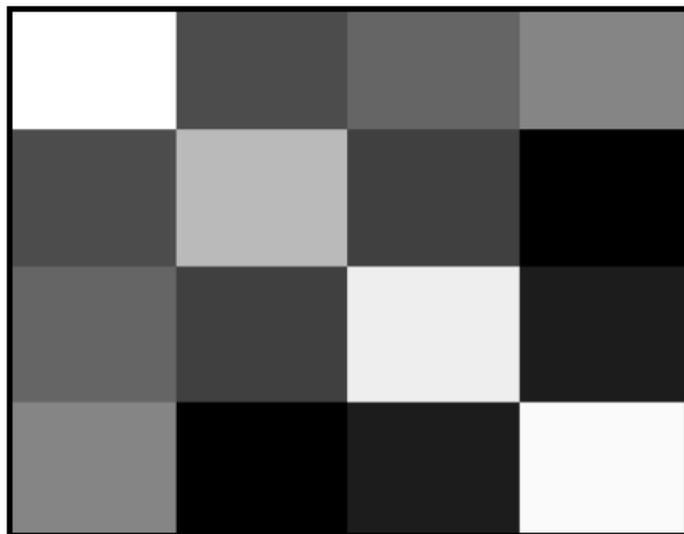
$\mathbf{Q}(z)$ is paraunitary, $\mathbf{Q}(z)\mathbf{Q}^P(z) = \mathbf{I}$

- ▶ Diagonalisation of $\mathbf{\Lambda}(z)$ is important for, e.g., decoupling of broadband MIMO systems.
- ▶ Polynomial subspace decomposition (i.e. $\mathbf{Q}(z)$) used in, e.g., broadband AoA estimation and beamforming.

How to Factorise a Polynomial Matrix?

- Iteratively minimise off-diagonal energy in $\mathbf{R}(z)$ and collect 'diagonalising' operations in $\mathbf{Q}(z)$.

$$\mathbf{R}(z) \rightarrow \mathbf{\Lambda}(z)$$

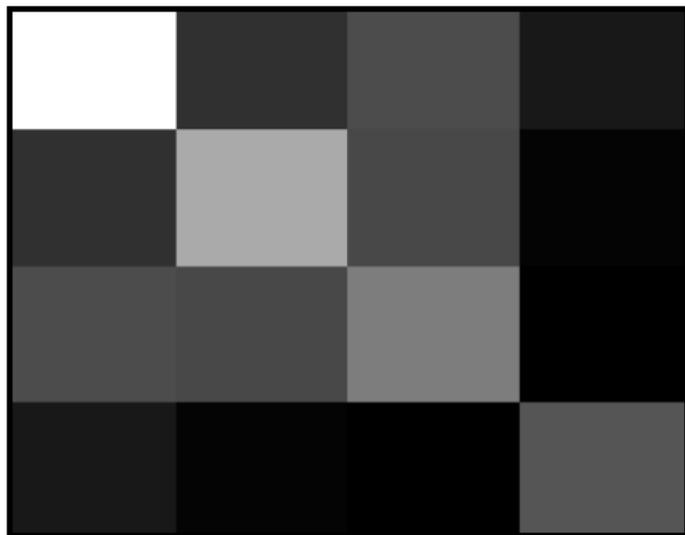


Energy in 'flattened' parahermitian matrix. White \Rightarrow high energy.

How to Factorise a Polynomial Matrix?

- ▶ Iteratively minimise off-diagonal energy in $\mathbf{R}(z)$ and collect 'diagonalising' operations in $\mathbf{Q}(z)$.

$$\mathbf{R}(z) \rightarrow \mathbf{\Lambda}(z)$$

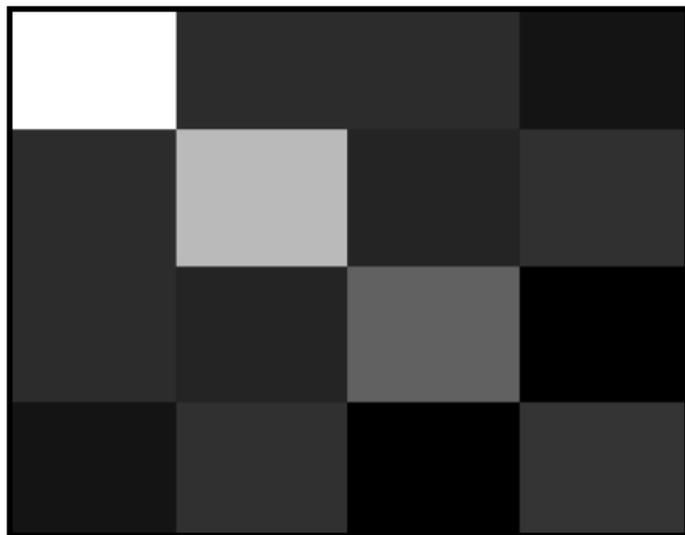


Energy in 'flattened' parahermitian matrix. White \Rightarrow high energy.

How to Factorise a Polynomial Matrix?

- ▶ Iteratively minimise off-diagonal energy in $\mathbf{R}(z)$ and collect 'diagonalising' operations in $\mathbf{Q}(z)$.

$$\mathbf{R}(z) \rightarrow \mathbf{\Lambda}(z)$$

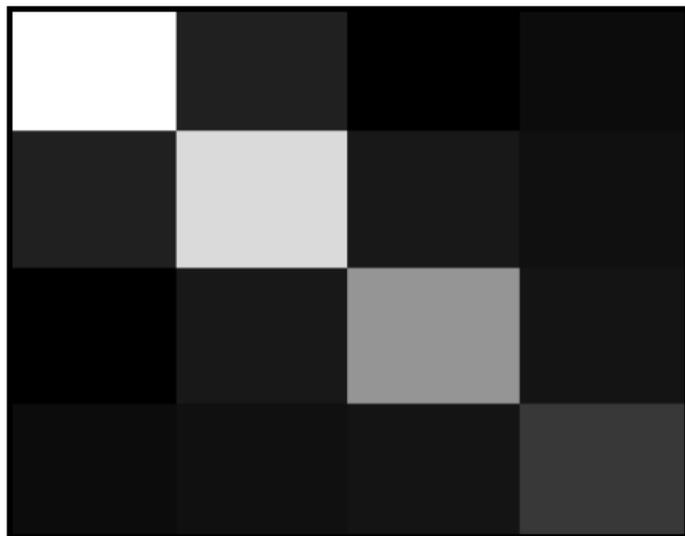


Energy in 'flattened' parahermitian matrix. White \Rightarrow high energy.

How to Factorise a Polynomial Matrix?

- ▶ Iteratively minimise off-diagonal energy in $\mathbf{R}(z)$ and collect 'diagonalising' operations in $\mathbf{Q}(z)$.

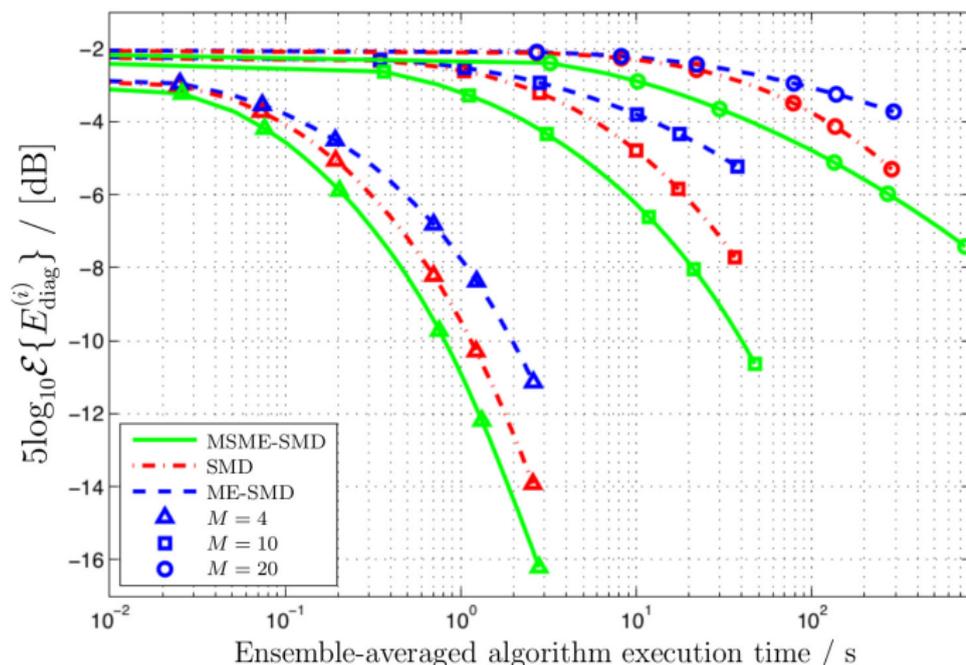
$$\mathbf{R}(z) \rightarrow \mathbf{\Lambda}(z)$$



Energy in 'flattened' parahermitian matrix. White \Rightarrow high energy.

How to Factorise a Polynomial Matrix?

- ▶ Iteratively minimise off-diagonal energy in $\mathbf{R}(z)$ and collect 'diagonalising' operations in $\mathbf{Q}(z)$.
- ▶ Key performance metric: off-diagonal energy vs execution time.



How to Factorise a Polynomial Matrix?

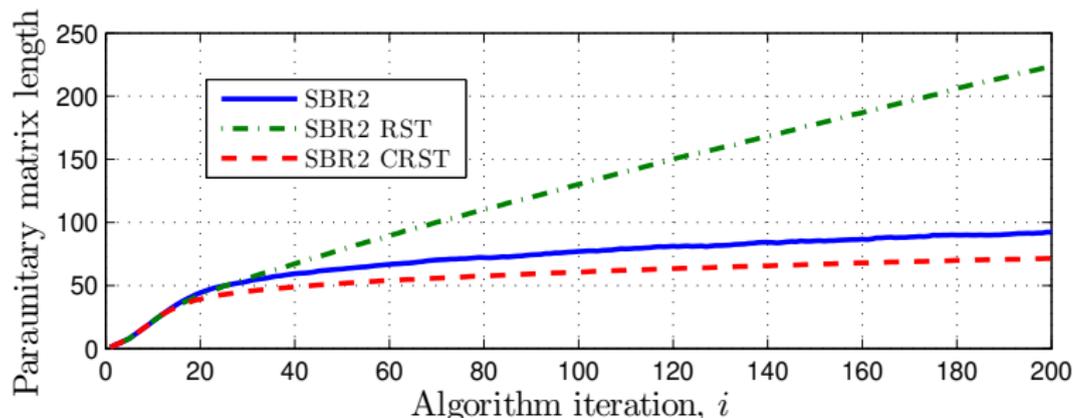
- ▶ Several iterative PEVD algorithms successfully minimise off-diagonal energy.
- ▶ Most important:
 - ▶ Second order sequential best rotation (SBR2) [1];
 - ▶ Sequential matrix diagonalisation (SMD) [2].
- ▶ SBR2 cheap to implement but can require many iterations.
- ▶ SMD more expensive but can converge in fewer iterations.

[1] J. G. McWhirter et al., An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Signal Process.*, 55(5):2158–2169, May 2007.

[2] S. Redif et al., Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans. on Signal Process.*, 63(1):81–89, Jan. 2015.

Requirement for Truncation in Iterative PEVD Algorithms

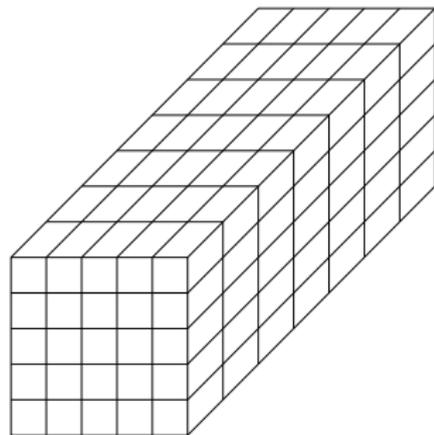
- ▶ The 'shift' operations performed in SBR2 and SMD lead to an increase in the order of the paraunitary matrix $\mathbf{Q}(z)$ at **each iteration**.
- ▶ If the order of $\mathbf{Q}(z)$ is not restricted in some way, memory usage and computational complexity will scale linearly with the order.



State-of-the-Art Truncation

[Ta *et al.*, ICIS&SP TSP 2007]

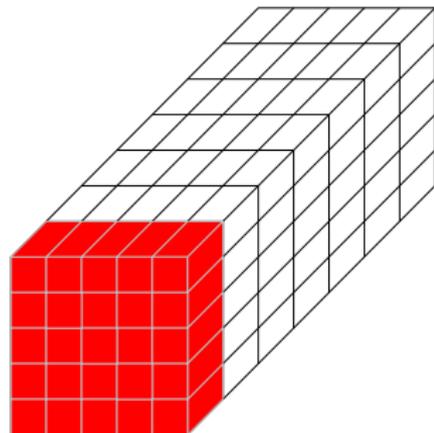
- ▶ Used to reduce the polynomial order of the paraunitary $Q(z)$.



State-of-the-Art Truncation

[Ta *et al.*, ICIS&SP TSP 2007]

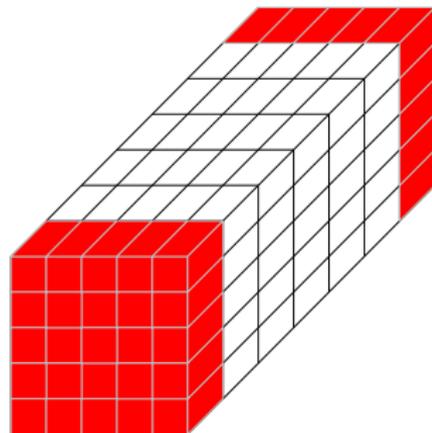
- ▶ Used to reduce the polynomial order of the paraunitary $Q(z)$.
- ▶ Successively removes the outermost lag of $Q(z)$ with the lowest energy.



State-of-the-Art Truncation

[Ta *et al.*, ICIS&SP TSP 2007]

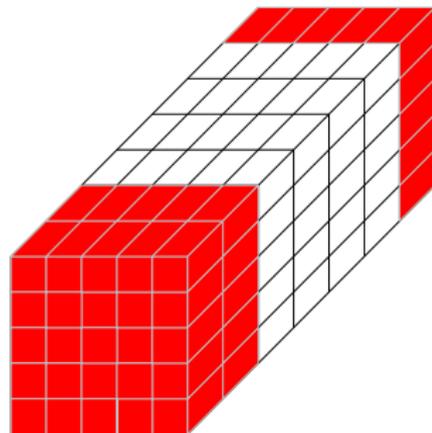
- ▶ Used to reduce the polynomial order of the paraunitary $Q(z)$.
- ▶ Successively removes the outermost lag of $Q(z)$ with the lowest energy.
- ▶ Lags can be removed from either side of $Q(z)$.



State-of-the-Art Truncation

[Ta *et al.*, ICIS&SP TSP 2007]

- ▶ Used to reduce the polynomial order of the paraunitary $Q(z)$.
- ▶ Successively removes the outermost lag of $Q(z)$ with the lowest energy.
- ▶ Lags can be removed from either side of $Q(z)$.



PEVD Ambiguity

- ▶ The paraunitary matrix in the PEVD is not unique.
- ▶ The same diagonalised parahermitian matrix, $\mathbf{\Lambda}(z)$, can be obtained using different paraunitary $\mathbf{F}(z) = \mathbf{Q}^P(z)$,

$$\mathbf{R}(z) \approx \mathbf{F}^P(z)\mathbf{\Lambda}(z)\mathbf{F}(z) = \hat{\mathbf{F}}^P(z)\mathbf{\Lambda}(z)\hat{\mathbf{F}}(z) \quad .$$

- ▶ Note that polynomial eigenvectors are in the **rows** of $\mathbf{F}(z)$.
- ▶ Using a modifying matrix, $\mathbf{\Gamma}(z)$, we can go from $\mathbf{F}(z)$ to $\hat{\mathbf{F}}(z)$,

$$\hat{\mathbf{F}}(z) = \mathbf{\Gamma}(z)\mathbf{F}(z) \quad .$$

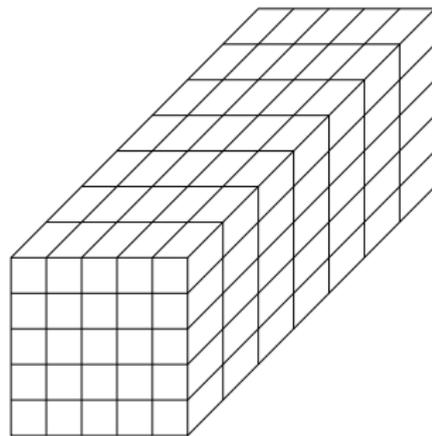
- ▶ For the row-shift truncation we simply use,

$$\mathbf{\Gamma}(z) = \text{diag}\{z^{-\tau_1} \ z^{-\tau_2} \ \dots \ z^{-\tau_M}\} \quad ,$$

which individually delays or advances each of the M rows of $\mathbf{F}(z)$.

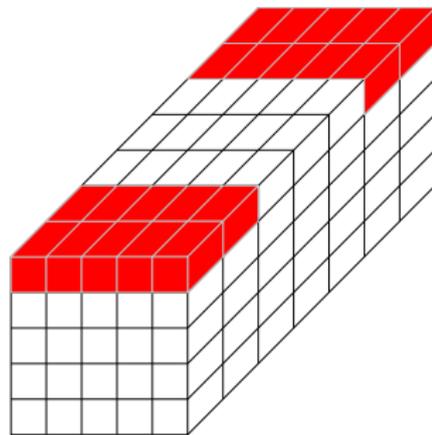
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.



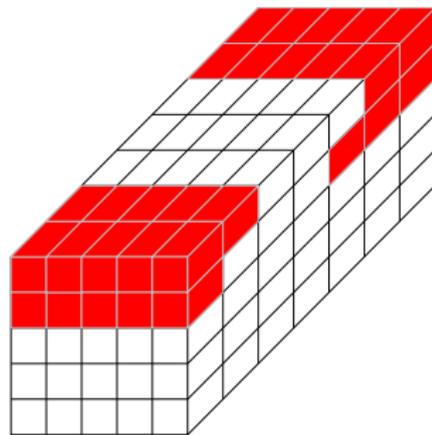
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.



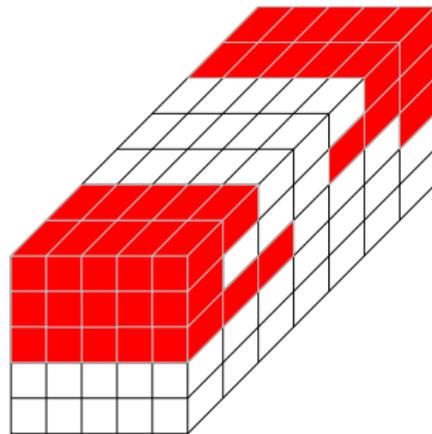
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.
- ▶ Each row can be truncated by a different amount.



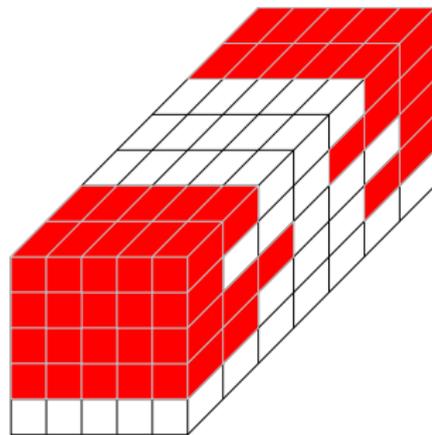
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.
- ▶ Each row can be truncated by a different amount.



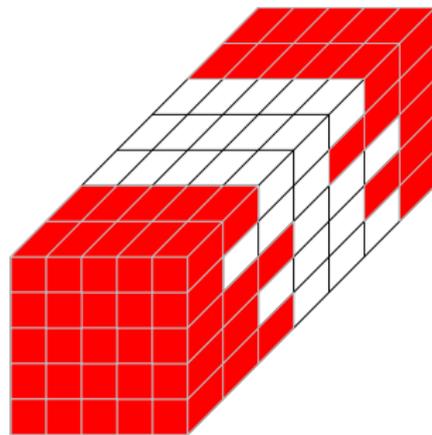
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.
- ▶ Each row can be truncated by a different amount.



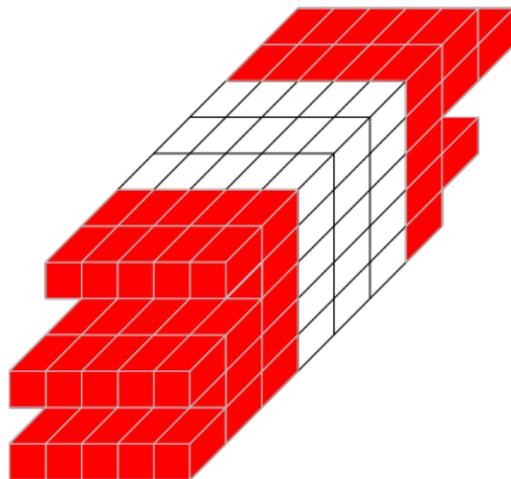
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.
- ▶ Each row can be truncated by a different amount.



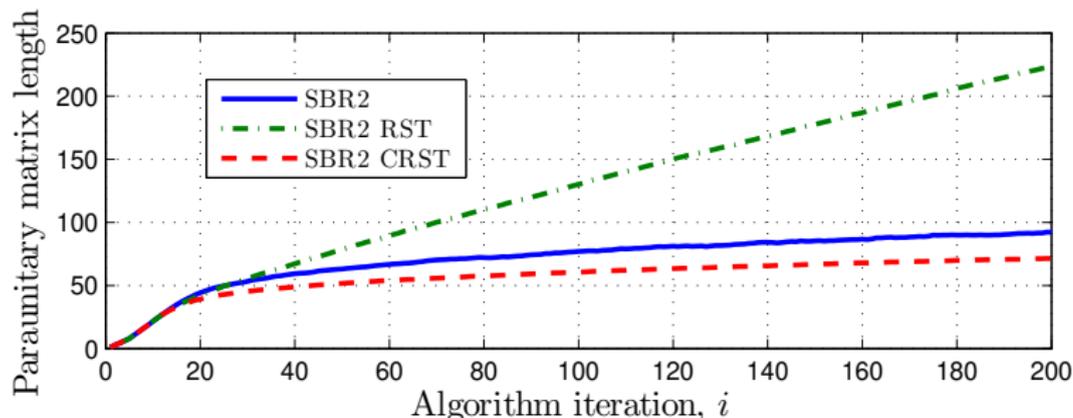
Row-Shift Truncation

- ▶ Used to reduce the polynomial order of the paraunitary $\mathbf{F}(z)$.
- ▶ Works the same as the state-of-the-art but applied to each row individually.
- ▶ Each row can be truncated by a different amount.
- ▶ Final step aligns rows using $\mathbf{\Gamma}(z)$.



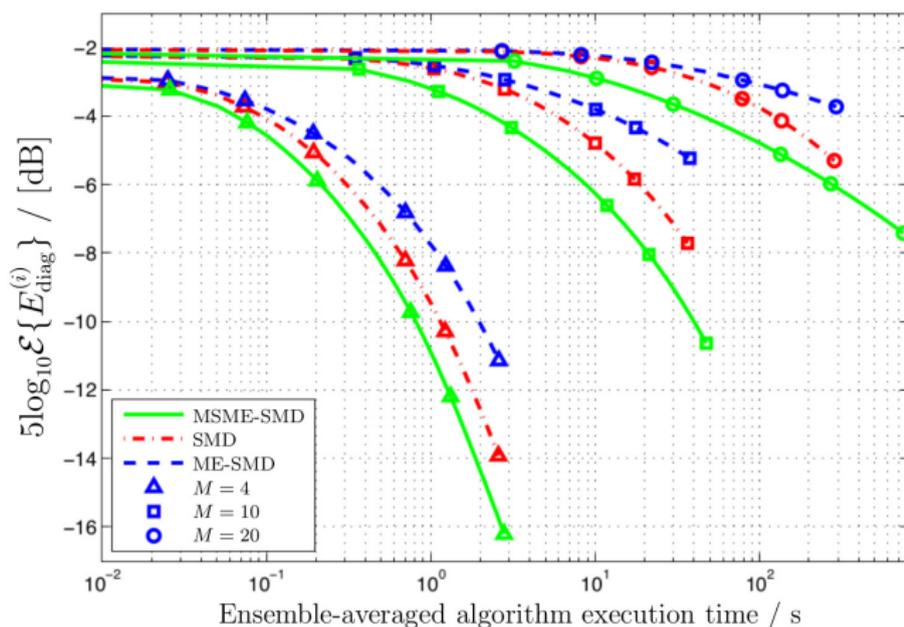
Requirement for Truncation in Iterative PEVD Algorithms

- ▶ The 'shift' operations performed in SBR2 and SMD lead to an increase in the order of the paraunitary matrix $Q(z)$ at **each iteration**.
- ▶ If the order of $Q(z)$ is not restricted in some way, memory usage and computational complexity will scale linearly with the order.



Motivation for my Research

- ▶ Existing iterative algorithms are slow to converge and not feasible to implement — worse for large spatial dimension M .
- ▶ Truncation reduces complexity but introduces error if excessive.



Motivation for my Research

- ▶ Existing iterative algorithms are slow to converge and not feasible to implement — worse for large spatial dimension M .
- ▶ Truncation reduces complexity but introduces error if excessive.
- ▶ Aims:
 - ▶ Develop techniques to lower complexity (and memory requirements) of existing methods.
 - ▶ Design novel, fast algorithms with improved scalability.
 - ▶ Investigate frequency-based methods to compute PEVD.

Algorithmic Improvements

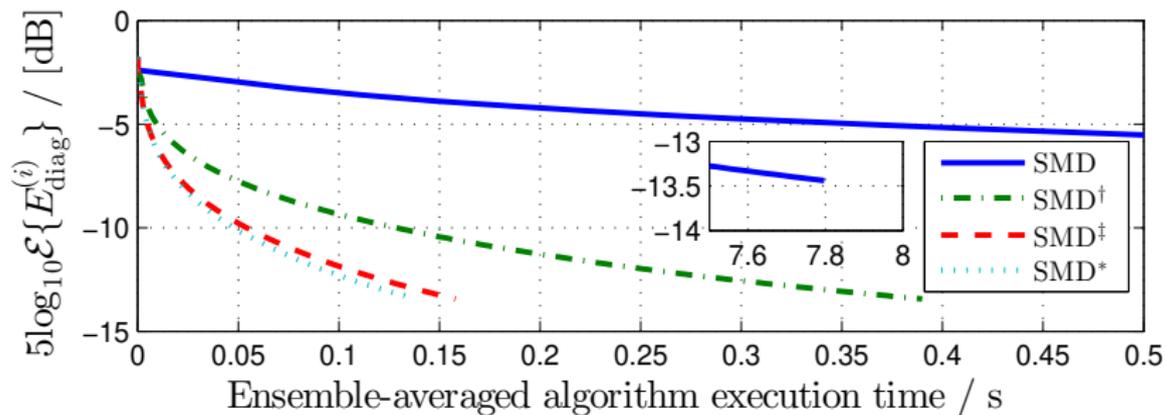
- ▶ First step: code profiling and optimisation.

Before optimisation

Line Number	Code	Calls	Total Time	% Time	Time Plot
137	[R,H] = DiagZL_slow(R,H);	1000	43.432 s	97.9%	
99	[R,H] = DiagZL_slow(R,H);	10	0.449 s	1.0%	
121	[R,H,stopcrit] = SMDZL_compMea...	1000	0.208 s	0.5%	

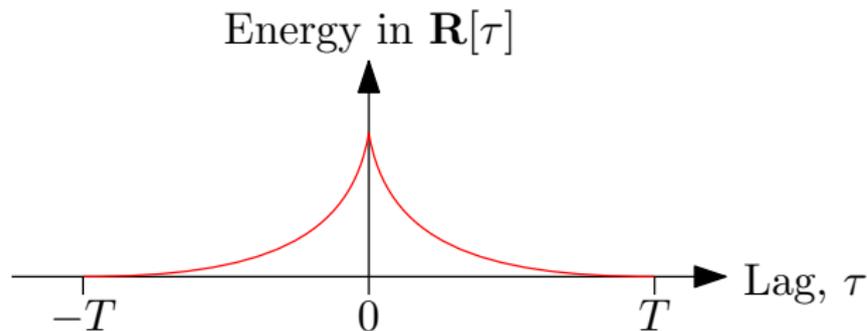
After optimisation

Line Number	Code	Calls	Total Time	% Time	Time Plot
121	[R,H,stopcrit] = SMDZL_compMea...	1000	0.159 s	32.6%	
137	[R,H] = DiagZL(R,H);	1000	0.077 s	15.7%	
142	[H,PUref] = TruncatePU_2(H,Mu,...	1000	0.066 s	13.5%	



Exploiting Parahermitian Matrix Structure

- ▶ Due to symmetry, the energy in a parahermitian matrix $\mathbf{R}[\tau]$ is generally concentrated at its centre.



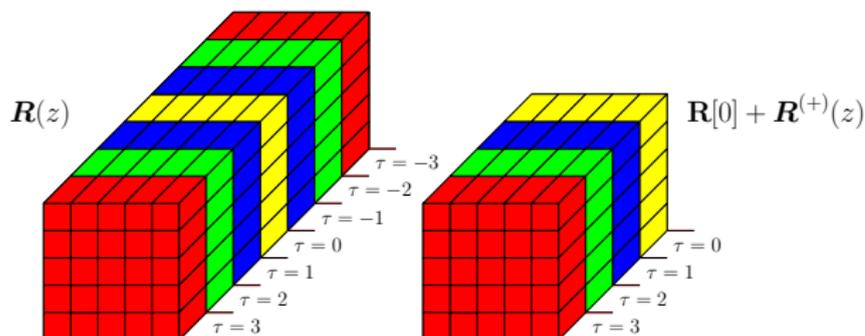
- ▶ With this knowledge, we investigated:
 - ▶ Methods that exploit matrix symmetry to approximately halve computational complexity and memory requirements.
 - ▶ Ways to focus our efforts on the most important central lags of the parahermitian matrix.

Reduced Parahermitian Matrix Representation

- By segmenting a parahermitian matrix $\mathbf{R}(z)$, we can write

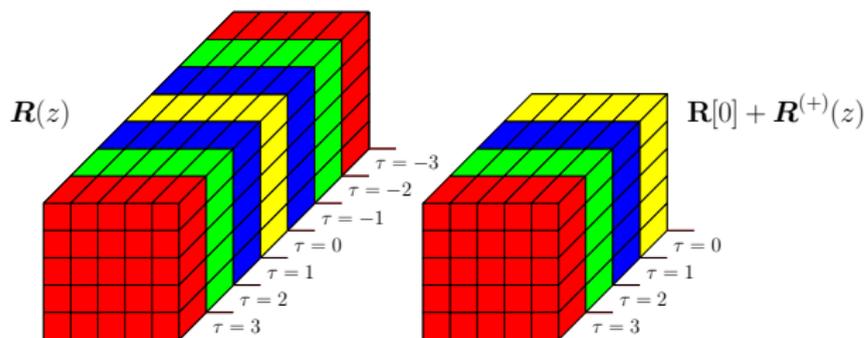
$$\mathbf{R}(z) = \mathbf{R}^{(-)}(z) + \mathbf{R}[0] + \mathbf{R}^{(+)}(z).$$

- $\mathbf{R}[0]$ is the zero lag matrix, $\mathbf{R}^{(+)}(z)$ contains terms for positive lag elements only, and $\mathbf{R}^{(-)}(z) = \mathbf{R}^{(+),P}(z)$.
- It is therefore sufficient to record a 'half-matrix' version of $\mathbf{R}(z)$.



Reduced Parahermitian Matrix Representation

- ▶ Columns beyond lag zero ($\tau = 0$) have been discarded.
- ▶ Modifications have to be made to the search and shift stages.
- ▶ Both columns **and** rows in the reduced matrix are searched.
- ▶ A 'cyclic shift' approach is employed.

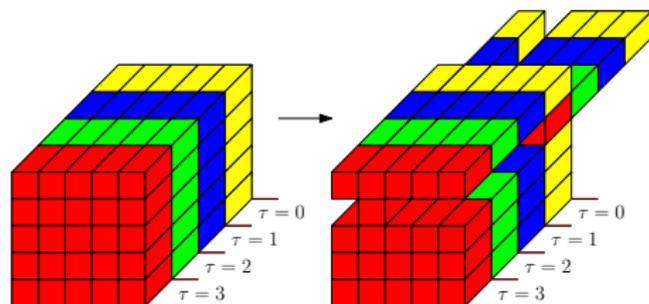


Reduced Parahermitian Matrix Representation

- ▶ An example of the shift operation is depicted for the case of $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ with parameters $k^{(i)} = 2$ and $\tau^{(i)} = -3$.

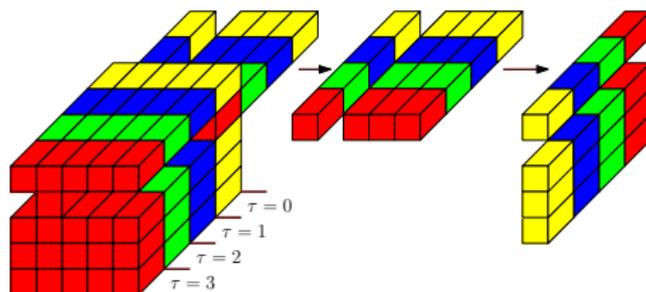
Reduced Parahermitian Matrix Representation

- ▶ An example of the shift operation is depicted for the case of $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ with parameters $k^{(i)} = 2$ and $\tau^{(i)} = -3$.
- ▶ Shifting procedure:
 1. The row is shifted.



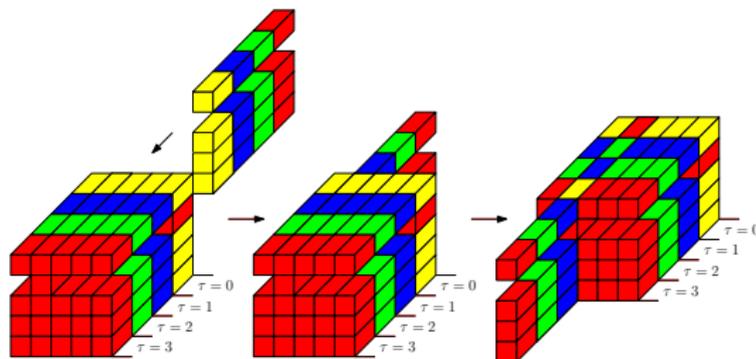
Reduced Parahermitian Matrix Representation

- ▶ An example of the shift operation is depicted for the case of $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ with parameters $k^{(i)} = 2$ and $\tau^{(i)} = -3$.
- ▶ Shifting procedure:
 1. The row is shifted.
 2. Non-diagonal elements in the $k^{(i)}$ th row past lag zero are extracted and parahermitian transposed.



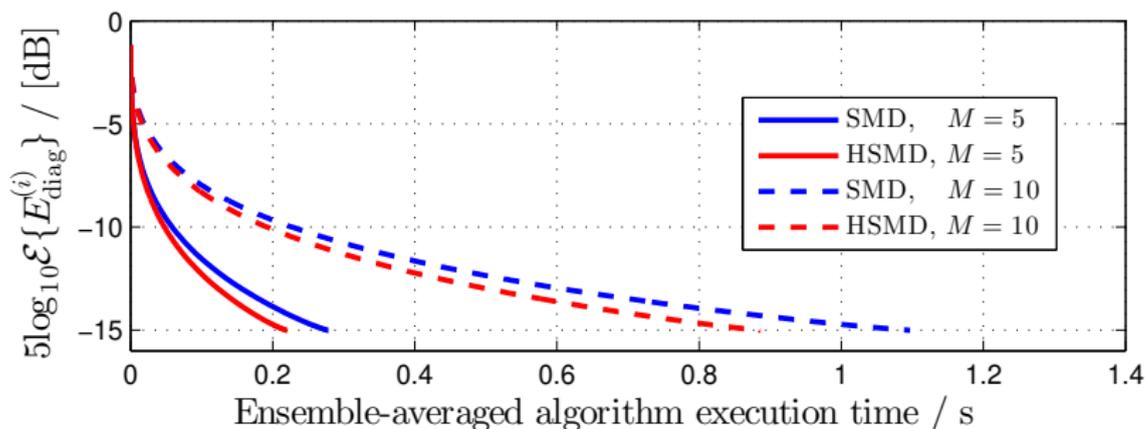
Reduced Parahermitian Matrix Representation

- ▶ An example of the shift operation is depicted for the case of $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ with parameters $k^{(i)} = 2$ and $\tau^{(i)} = -3$.
- ▶ Shifting procedure:
 1. The row is shifted.
 2. Non-diagonal elements in the $k^{(i)}$ th row past lag zero are extracted and parahermitian transposed.
 3. These elements are appended to the $k^{(i)}$ th column at lag zero and this column is shifted in the opposite direction.



Reduced Parahermitian Matrix Representation

- ▶ Off-diagonal energy versus algorithm execution time for standard SMD algorithm and 'half-matrix' SMD (HSMD) implementation.
- ▶ M is spatial dimension of parahermitian matrix, e.g., number of array elements.



Reduced Parahermitian Matrix Representation

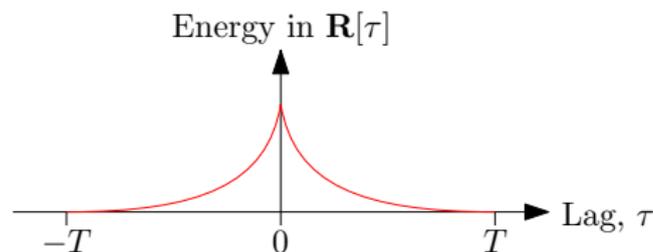
- ▶ Approximate resource requirements of standard (SMD) and proposed (HSMD) representations of an $M \times M \times (2N^{(i)} + 1)$ parahermitian matrix at the i th iteration.

Method	Complexity	Storage	Memory Moves
SMD	$4N^{(i)} M^3$	$2N^{(i)} M^2$	$4N^{(i-1)} M$
HSMD	$2N^{(i)} M^3$	$N^{(i)} M^2$	$2N^{(i-1)} M$

- ▶ All resource requirements are approximately halved using the proposed approach.

Restricted Update SMD (RU-SMD)

- ▶ Restricting the search space of iterative PEVD algorithms to a subset of lags around lag zero of a parahermitian matrix can bring performance gains. [Corr *et al.*, ISP'15 & Coutts *et al.*, Asilomar SSC'16]
- ▶ The update step of SMD is its most computationally costly operation. [Redif *et al.*, IEEE TSP 2015]
- ▶ The restricted update SMD (RU-SMD) algorithm therefore restricts the search and update spaces of the SMD algorithm.
- ▶ Aim of RU-SMD is to be **less expensive** and **faster** than SMD.

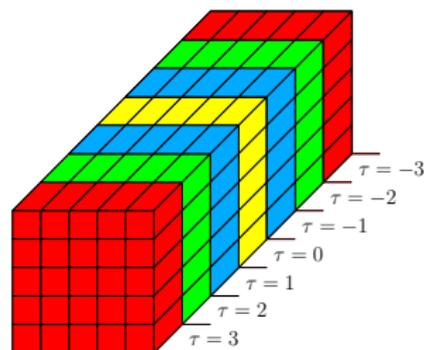


RU-SMD Overview

- ▶ The RU-SMD algorithm computes the PEVD of a parahermitian matrix $\mathbf{R}(z)$ over $i = 0 \dots I$ iterations.
- ▶ RU-SMD has two main steps:
 1. Restricted update: iteratively diagonalise parahermitian matrix while monotonically contracting the search and update space.
 2. Matrix regeneration: regenerate parahermitian matrix when search and update space has maximally contracted.
- ▶ Steps 1 and 2 are repeated for index $\alpha = 0 \dots \beta$.
- ▶ The space restriction in 1 limits the number of search operations, and reduces the computations required to update the parahermitian matrix.

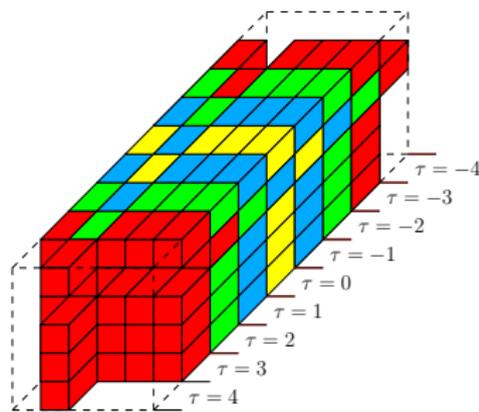
Restricted Update Approach

- ▶ Matrix $\mathbf{S}^{(i-1)}(z) = \mathbf{R}_{(\alpha)}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ with maximum lag $\tau_{\max}^{(i)} = 3$ input to restricted update step. Note: $\mathbf{R}_{(0)}(z) = \mathbf{R}(z)$.
- ▶ Treat lags $|\tau| > \tau_{\max}^{(i)}$ as invalid.



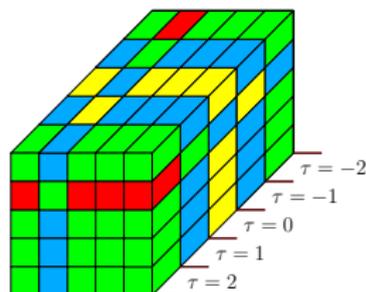
Restricted Update Approach

- ▶ Shifting of $k^{(i)}$ th row and column energy to lag zero from lags $\pm\tau^{(i)}$ ($k^{(i)} = 2$, $\tau^{(i)} = -1$).
- ▶ Invalid values from lags $|\tau| > \tau_{\max}^{(i)}$ are shifted towards lag zero.



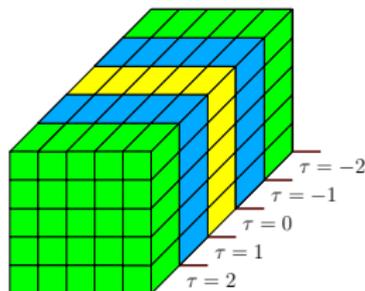
Restricted Update Approach

- ▶ Valid central matrix with maximum lag $(\tau_{\max}^{(i)} - |\tau^{(i)}|) = 2$, $\mathbf{S}^{(i)''}(z)$, is extracted.
- ▶ Lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}|)$ are invalid.



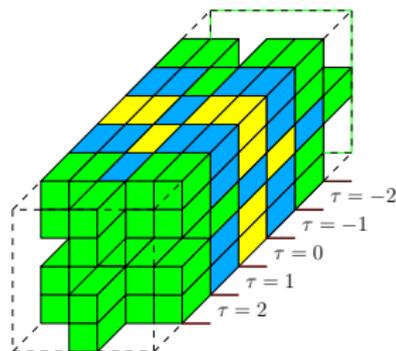
Restricted Update Approach

- ▶ Update step: $\mathcal{S}^{(i)}(z) = \mathbf{Q}^{(i)} \mathcal{S}^{(i)''}(z) \mathbf{Q}^{(i),H}$. Matrix $\mathbf{Q}^{(i)}$ is obtained from EVD of lag zero.
- ▶ Lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}|)$ are invalid.



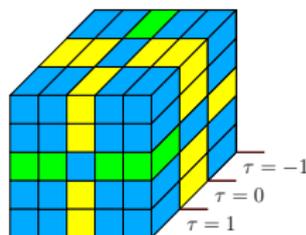
Restricted Update Approach

- ▶ Shifting of $k^{(i+1)}$ th row and column energy to lag zero from lags $\pm\tau^{(i+1)}$ ($k^{(i+1)} = 3$, $\tau^{(i+1)} = -1$).
- ▶ Invalid values from lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}|)$ are shifted towards lag zero.



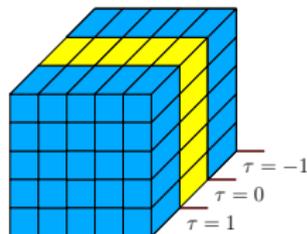
Restricted Update Approach

- ▶ Valid central matrix with maximum lag $(\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}|) = 1$, $\mathbf{S}^{(i+1)''}(z)$, is extracted.
- ▶ Lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}|)$ are invalid.



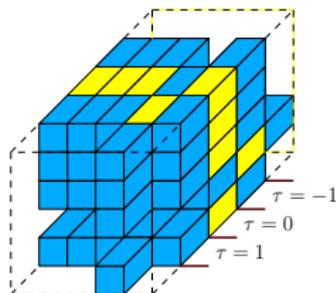
Restricted Update Approach

- ▶ Update step: $\mathbf{S}^{(i+1)}(z) = \mathbf{Q}^{(i+1)} \mathbf{S}^{(i+1)H}(z) \mathbf{Q}^{(i+1),H}$. Matrix $\mathbf{Q}^{(i+1)}$ is obtained from EVD of lag zero.
- ▶ Lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}|)$ are invalid.



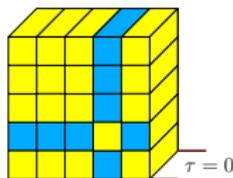
Restricted Update Approach

- ▶ Shifting of $k^{(i+2)}$ th row and column energy to lag zero from lags $\pm\tau^{(i+2)}$ ($k^{(i+2)} = 4$, $\tau^{(i+2)} = -1$).
- ▶ Invalid values from lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}|)$ are shifted towards lag zero.



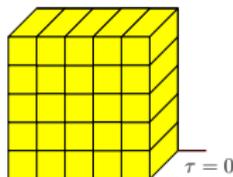
Restricted Update Approach

- ▶ Valid central matrix with maximum lag $(\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}| - |\tau^{(i+2)}|) = 0$, $\mathbf{S}^{(i+2)''}(z)$, is extracted.
- ▶ Lags $|\tau| > (\tau_{\max}^{(i)} - |\tau^{(i)}| - |\tau^{(i+1)}| - |\tau^{(i+2)}|)$ are invalid.



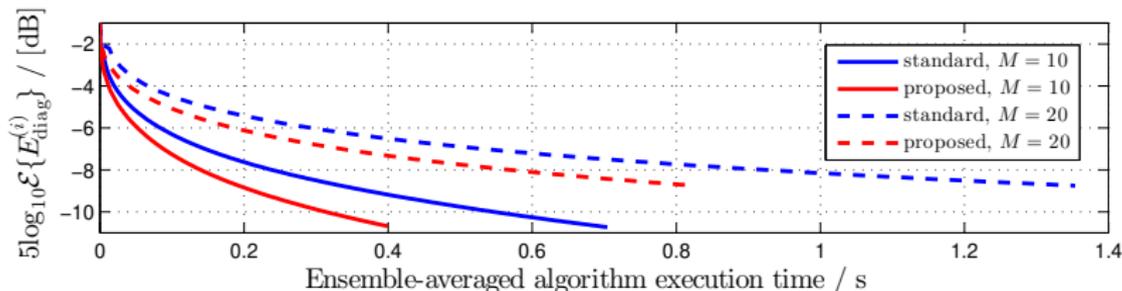
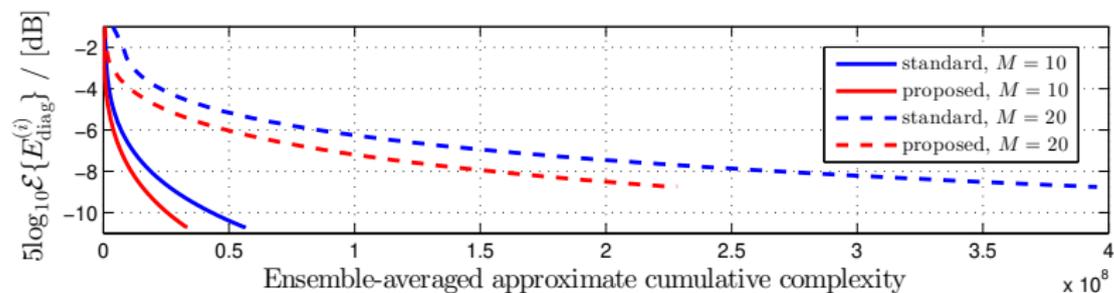
Restricted Update Approach

- ▶ Update step: $\mathbf{S}^{(i+2)}(z) = \mathbf{Q}^{(i+2)} \mathbf{S}^{(i+2)H}(z) \mathbf{Q}^{(i+2),H}$. Matrix $\mathbf{Q}^{(i+2)}$ is obtained from EVD of lag zero.
- ▶ Only zero lag remains: matrix must now be regenerated.



RU-SMD Performance

- ▶ The restricted update process reduces the complexity required at each iteration, but requires a matrix regeneration step.
- ▶ Fortunately, overall algorithm execution time and complexity requirements are reduced.

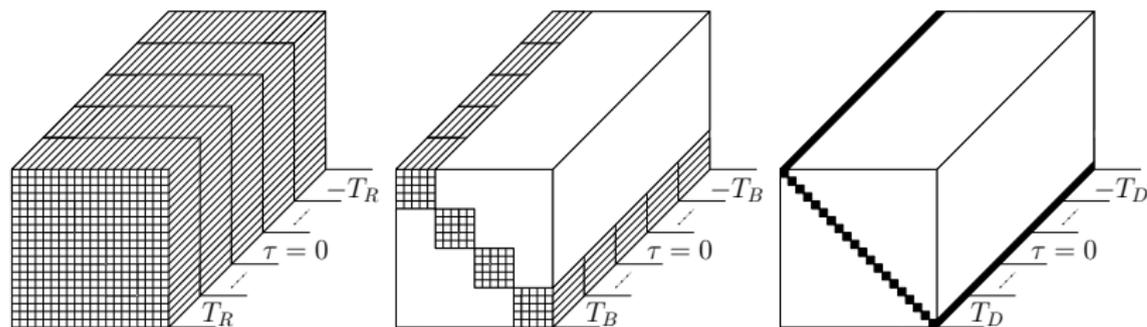


Summary of Algorithmic Improvements

- ▶ Improvements to algorithm implementation efficiency.
- ▶ Half-matrix form for parahermitian matrix.
- ▶ Restricted update approach.
- ▶ For the interested reader, see my thesis!
- ▶ **Problem:** none of these address increased algorithmic complexity (proportional to M^3) as spatial dimension M increases.

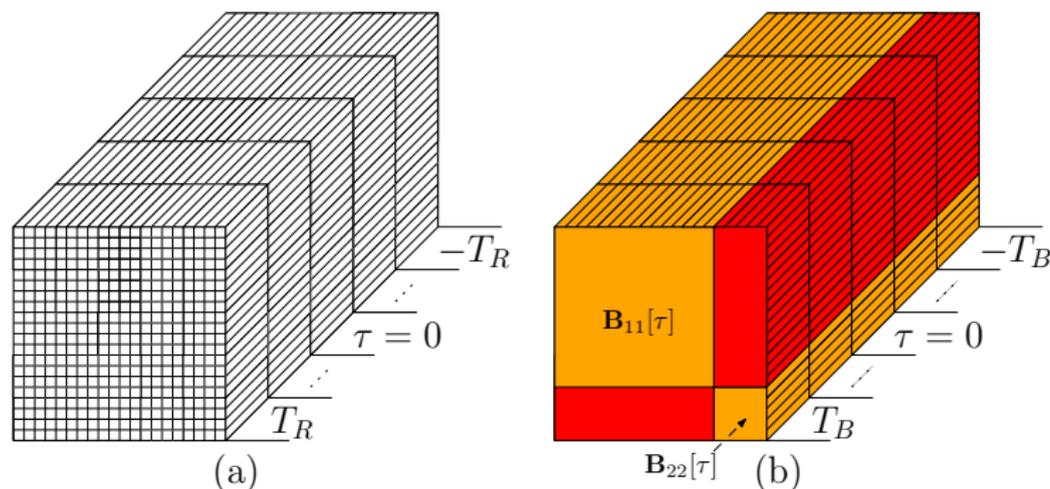
'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Traditional PEVD algorithms are tasked with diagonalising an entire $M \times M$ parahermitian matrix via sequential operations.
- ▶ DaC method first 'divides' the matrix into a number of smaller, independent parahermitian matrices, before diagonalising — or 'conquering' — each matrix separately.
- ▶ DaC scheme can substantially reduce PEVD complexity, which is typically proportional to M^3 .



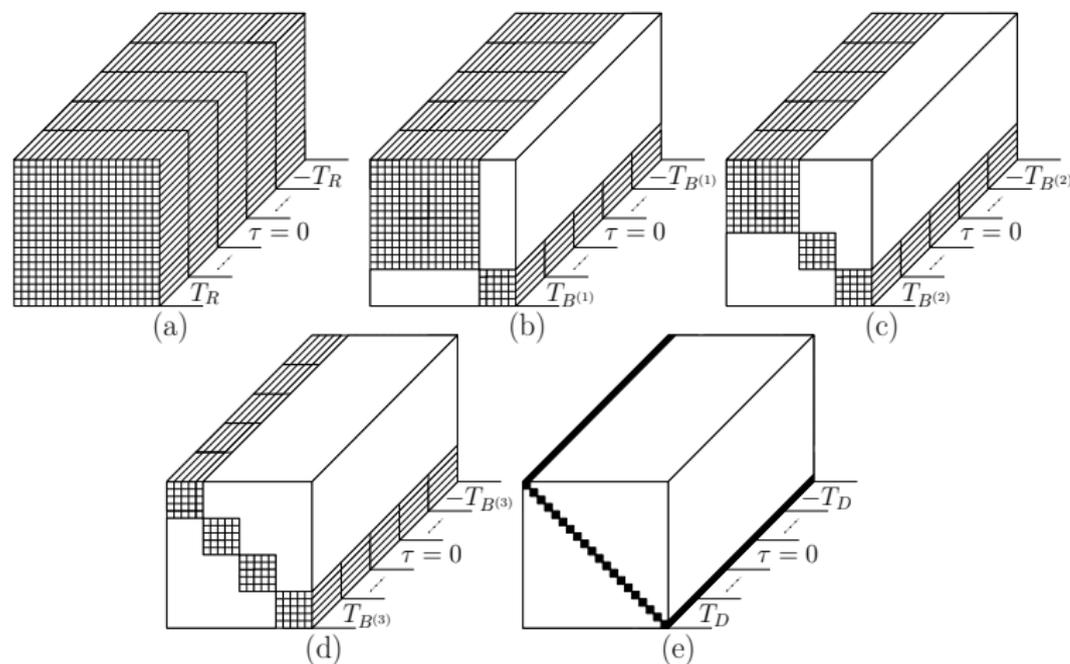
'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Iteratively minimising energy in red regions yields a block diagonal parahermitian matrix.
- ▶ Remaining $\mathbf{B}_{11}(z)$ and $\mathbf{B}_{22}(z)$ are independent parahermitian matrices and can be diagonalised separately.



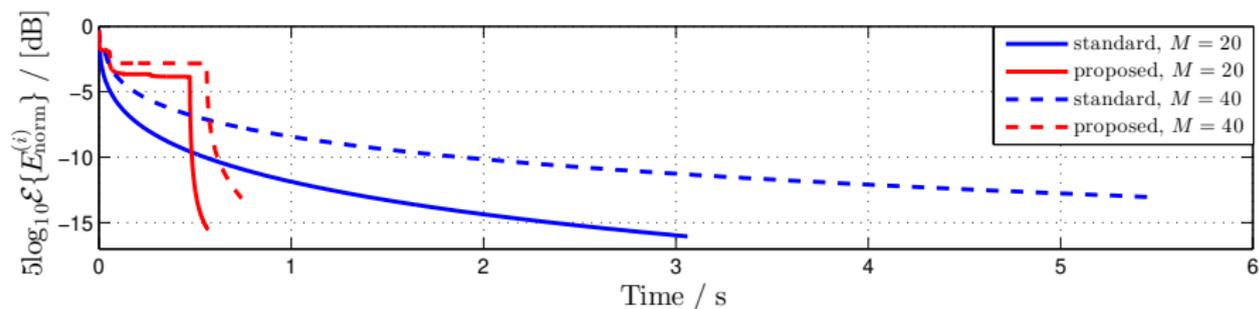
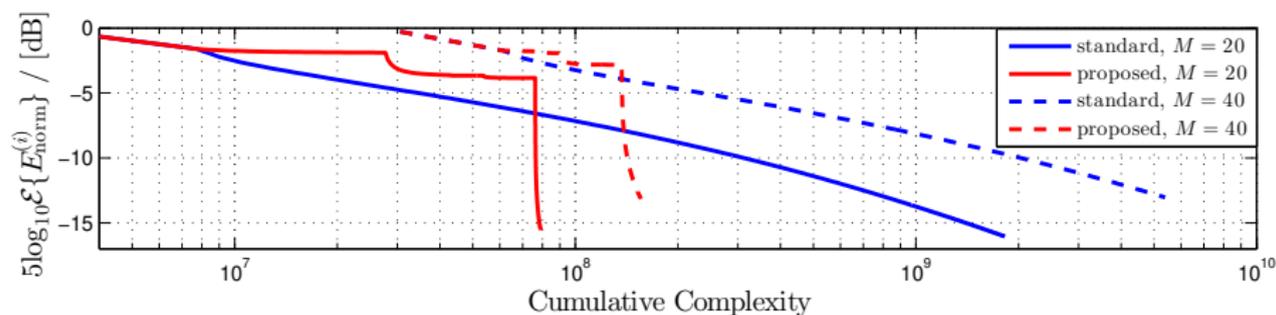
'Divide-and-Conquer' (DaC) Approach to the PEVD

- Several block diagonalisation steps yield a block diagonal matrix.



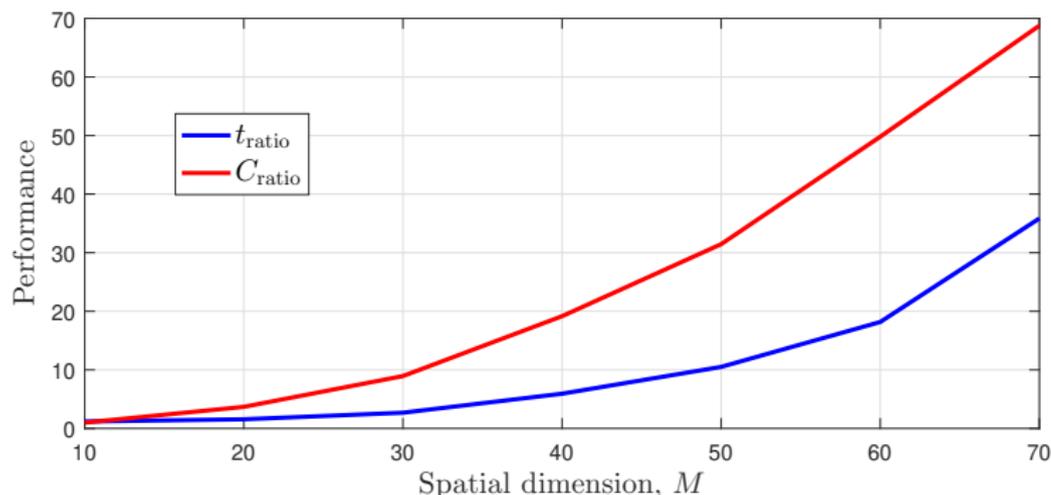
'Divide-and-Conquer' (DaC) Approach to the PEVD

- SMD (standard) versus DaC SMD (DC-SMD, proposed).



'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Divide-and-conquer strategy becomes increasingly useful as spatial dimension M increases.

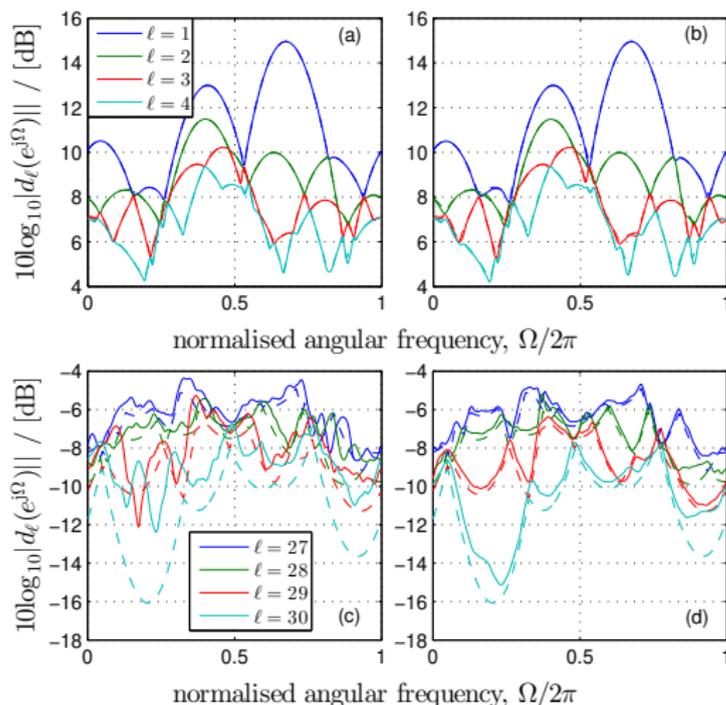


$$t_{\text{ratio}} = \frac{\mathcal{E}\{\text{Execution time}_{\text{SMD}, -10 \text{ dB}, M}\}}{\mathcal{E}\{\text{Execution time}_{\text{DC-SMD}, -10 \text{ dB}, M}\}}$$

$$C_{\text{ratio}} = \frac{\mathcal{E}\{\text{Complexity}_{\text{SMD}, -10 \text{ dB}, M}\}}{\mathcal{E}\{\text{Complexity}_{\text{DC-SMD}, -10 \text{ dB}, M}\}}$$

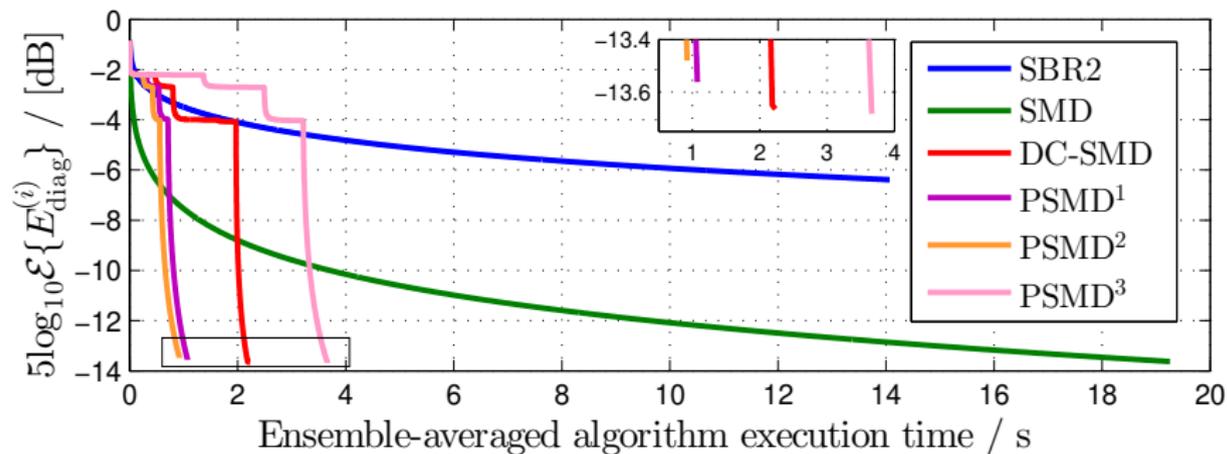
'Divide-and-Conquer' (DaC) Approach to the PEVD

- Power spectral densities of the (a,b) first and (c,d) last four eigenvalues obtained from (a,c) SMD and (b,d) DC-SMD.



'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Independent parahermitian matrices \Rightarrow parallel processing.
- ▶ Task: combine parallelised DaC strategies for the PEVD with developed complexity and memory reduction techniques.
- ▶ **Parallel-Sequential Matrix Diagonalisation (PSMD).**



'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Independent parahermitian matrices \Rightarrow parallel processing.
- ▶ Task: combine parallelised DaC strategies for the PEVD with developed complexity and memory reduction techniques.
- ▶ **Parallel-Sequential Matrix Diagonalisation (PSMD).**

Method	E.val. Res.	MSE	Paraun. Err.	L_Q
SBR2 [1]	1.1305	1.293×10^{-6}	2.448×10^{-8}	133.8
SMD [2]	0.0773	3.514×10^{-6}	6.579×10^{-8}	165.5
DC-SMD	0.0644	6.785×10^{-6}	1.226×10^{-14}	360.4
PSMD ¹	0.0658	6.918×10^{-6}	4.401×10^{-15}	279.3
PSMD ²	0.0661	8.346×10^{-6}	1.303×10^{-8}	156.0
PSMD ³	0.0245	7.618×10^{-7}	1.307×10^{-14}	307.6

-
- [1] J. G. McWhirter et al., An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Signal Process.*, 55(5):2158–2169, May 2007.
- [2] S. Redif et al., Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans. on Signal Process.*, 63(1):81–89, Jan. 2015.
- [3] F. K. Coutts et al., Efficient Implementation of Iterative Polynomial Matrix EVD Algorithms Exploiting Structural Redundancy and Parallelisation. *IEEE Trans. on Circuits and Systems*, 66(12):4753–4766, Dec. 2019.

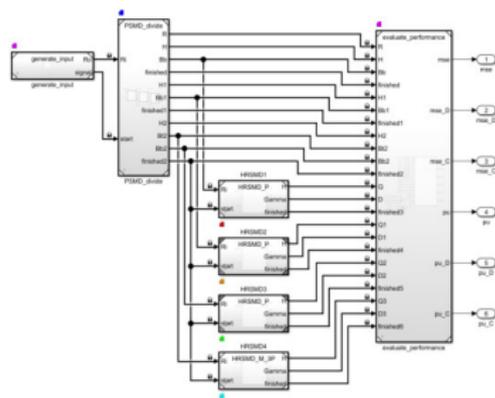
'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Independent parahermitian matrices \Rightarrow parallel processing.
- ▶ Task: combine parallelised DaC strategies for the PEVD with developed complexity and memory reduction techniques.
- ▶ **Parallel-Sequential Matrix Diagonalisation (PSMD).**
- ▶ MATLAB's Simulink Embedded Coder software used to convert PSMD algorithm to parallelised C-based implementation.
- ▶ Efficient implementations on quad-core Intel i7-4700MQ CPU and quad-core Raspberry Pi 3 Model B+.



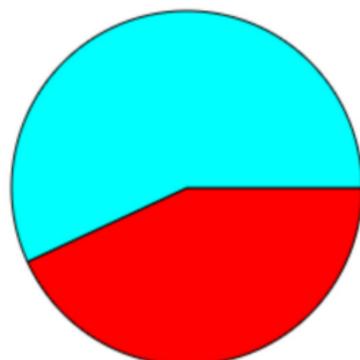
'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ Independent parahermitian matrices \Rightarrow parallel processing.
- ▶ Task: combine parallelised DaC strategies for the PEVD with developed complexity and memory reduction techniques.
- ▶ **Parallel-Sequential Matrix Diagonalisation (PSMD)**.
- ▶ MATLAB's Simulink Embedded Coder software used to convert PSMD algorithm to parallelised C-based implementation.
- ▶ Efficient implementations on quad-core Intel i7-4700MQ CPU and quad-core Raspberry Pi 3 Model B+.

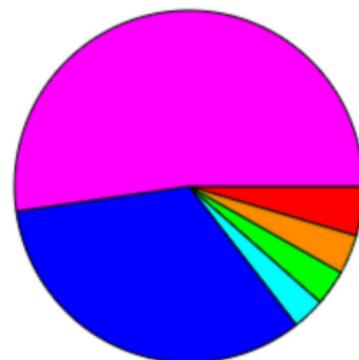


'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ MATLAB's Simulink Embedded Coder software used to convert PSMD algorithm to parallelised C-based implementation.
- ▶ Quad-core Intel i7-4700MQ CPU: parallel implementation 16.4% faster than serial implementation.



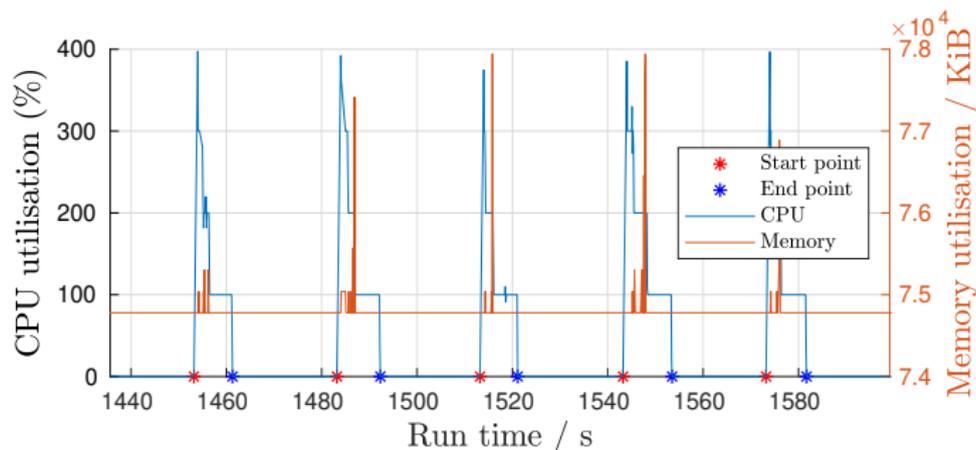
PSMD (771733.077)
System (1017050.410)



HRSMD1 (79740.723) HRSMD4 (52535.875)
HRSMD2 (63042.116) PSMD_divide (582895.842)
HRSMD3 (57544.391) System (915395.153)

'Divide-and-Conquer' (DaC) Approach to the PEVD

- ▶ MATLAB's Simulink Embedded Coder software used to convert PSMD algorithm to parallelised C-based implementation.
- ▶ Quad-core Intel i7-4700MQ CPU: parallel implementation 16.4% faster than serial implementation.
- ▶ Quad-core Raspberry Pi 3 Model B+: parallel implementation 19.6% faster than serial implementation.



DFT-Based PEVD Algorithms

- ▶ Demonstrated one of the two main categories of PEVD algorithm:

1. Iterative time (lag) based methods.

- ▶ Directly manipulate polynomial coefficients and seek to iteratively diagonalise parahermitian matrix $\mathbf{R}(z)$.
- ▶ Encourage spectral majorisation of eigenvalues.
- ▶ Established algorithms: SBR2 [1] and SMD [2].
- ▶ Recent low-complexity, parallelisable, divide-and-conquer algorithm in [3] for large arrays.

[1] J. G. McWhirter et al., An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Signal Process.*, 55(5):2158–2169, May 2007.

[2] S. Redif et al., Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans. on Signal Process.*, 63(1):81–89, Jan. 2015.

[3] F. K. Coutts et al., Efficient Implementation of Iterative Polynomial Matrix EVD Algorithms Exploiting Structural Redundancy and Parallelisation. *IEEE Trans. on Circuits and Systems*, 66(12):4753–4766, Dec. 2019.

DFT-Based PEVD Algorithms

- ▶ However, there is an alternative type of PEVD algorithm:

2. Fixed order frequency based methods.

- ▶ Transform the problem into a pointwise-in-frequency standard matrix decomposition.
- ▶ Able to provide a spectrally majorised decomposition, or attempt to approximate maximally smooth, analytic eigenvalues.
- ▶ Formulation in [4] performs well for finite order problems [5], but requires an a priori guess of the polynomial order of $Q(z)$.

[4] M. Tohidian et al., A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices. *EURASIP J. Adv. Signal Process.*, 2013:93, December 2013.

[5] F. K. Coutts et al., A Comparison of Iterative and DFT-Based Polynomial Matrix Eigenvalue Decompositions. *In Proc. IEEE 7th Int. Workshop Comp. Advances in Multi-Sensor Adaptive Process.*, Dec. 2017.

Fixed Order Frequency Based Methods

- ▶ An existing approach obtains an approximate PEVD via K independent EVDs in the discrete frequency domain:

$$\mathbf{R}[k] = \mathbf{Q}[k]\mathbf{\Lambda}[k]\mathbf{Q}^H[k], \quad k = 0 \dots K - 1,$$

$$\mathbf{R}[k] = \mathbf{R}(z)|_{z=e^{j\Omega_k}} = \sum_{\tau} \mathbf{R}[\tau]e^{-j\Omega_k\tau}, \quad \Omega_k = 2\pi k/K$$

- ▶ Can rearrange the eigenvalues and -vectors at each frequency bin.
- ▶ Ambiguity in eigenvector phase leads to discontinuities in phase between eigenvectors in adjacent frequency bins.
- ▶ Phase alignment step alters phases to minimise discontinuities.
- ▶ Following the permutation (if desired) and phase alignment of $\mathbf{Q}[k]$, $\mathbf{Q}[\tau]$ and $\mathbf{\Lambda}[\tau]$ are computed via the inverse DFT.

Fixed Order Frequency Based Methods

- ▶ While analytic polynomial eigenvalues and eigenvectors have been shown to exist as absolutely convergent Laurent series in [6] there is currently no way of knowing the length of the series a priori.
- ▶ When converting $\mathbf{\Lambda}[k]$ and $\mathbf{Q}[k]$ to the lag domain, the order of the IDFT restricts the series' length to K .
- ▶ For unsufficiently large K , this can result in energy from ignored high order polynomial coefficients corrupting the fixed set of K coefficients (i.e., time domain aliasing).

[6] S. Weiss et al., On the Existence and Uniqueness of the Eigenvalue Decomposition of a Parahermitian Matrix. *IEEE Trans. on Signal Process.*, 66(10):2659–2672, May 2018.

Proposed Approach Overview

- ▶ We propose an iterative frequency-based scheme:

$$\mathbf{R}[k] = \mathbf{Q}[k]\mathbf{\Lambda}[k]\mathbf{Q}^H[k], \quad k = 0 \dots K_i - 1.$$

- ▶ K_i is increased within the set $K_i \in \{2^{l+i} \mid l, i \in \mathbb{Z}^+, 2^l > L\}$.
- ▶ L is the polynomial order of $\mathbf{R}(z)$.
- ▶ $i = 0, 1, \dots, I - 1$ records current iteration.
- ▶ Use method for reordering the eigenvalues and -vectors from [4].
- ▶ Use a phase alignment function from [7], which uses Powell's 'dogleg' algorithm [8] to maximise eigenvector smoothness.
- ▶ Iterations continue while decomposition MSE above threshold ϵ .

[4] M. Tohidian et al., A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices. *EURASIP J. Adv. Signal Process.*, 2013:93, December 2013.

[7] F. K. Coutts et al., Enforcing Eigenvector Smoothness for a Compact DFT-based Polynomial Eigenvalue Decomposition. *In Proc. IEEE Workshop on Sensor Array and Multichannel Sig. Process.*, July. 2018.

[8] M. J. D. Powell, A new algorithm for unconstrained optimization. *Nonlinear programming*, 31–65, 1970.

Proposed Approach Overview

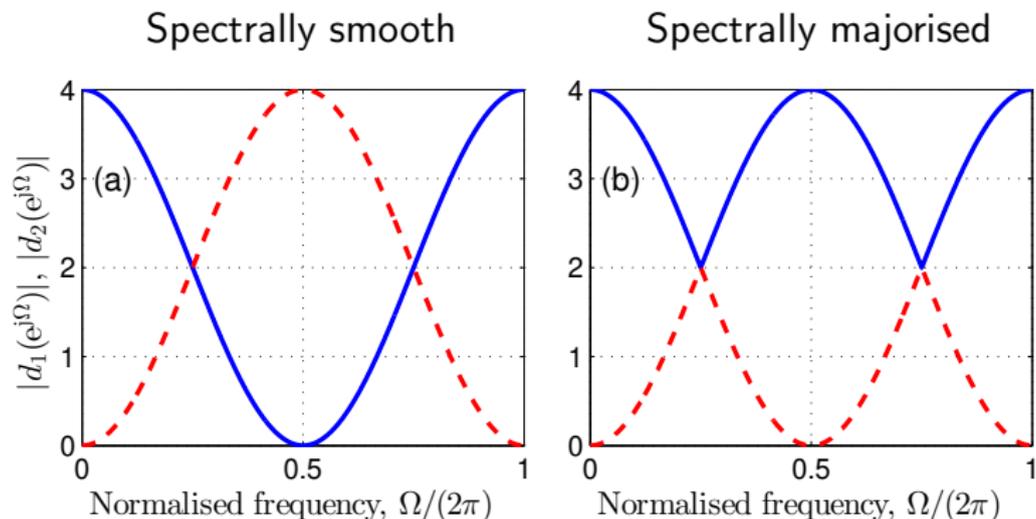
- ▶ We propose an iterative frequency-based scheme:

$$\mathbf{R}[k] = \mathbf{Q}[k]\mathbf{\Lambda}[k]\mathbf{Q}^H[k], \quad k = 0 \dots K_i - 1.$$

- ▶ K_i is increased within the set $K_i \in \{2^{l+i} \mid l, i \in \mathbb{Z}^+, 2^l > L\}$.
- ▶ In iteration i , $\mathbf{Q}[k]$ and $\mathbf{\Lambda}[k]$ for $k = 0, 1, \dots, K_i - 1$ are identical for $k = 0, 2, 4, \dots, K_{i+1} - 2$ in the $(i + 1)$ th iteration.
- ▶ Phase alignment step exploits this to aid optimisation.
- ▶ $\mathbf{Q}[\tau]$ and $\mathbf{\Lambda}[\tau]$ are computed via the inverse DFT following the permutation (if desired) and phase alignment of $\mathbf{Q}[k]$.
- ▶ MSE computed between $\hat{\mathbf{R}}(z) = \mathbf{Q}(z)\mathbf{\Lambda}(z)\mathbf{Q}^P(z)$ and $\mathbf{R}(z)$.

Smooth Decomposition

- ▶ In a smooth (analytic) decomposition, the eigenvalues and -vectors are arranged such that discontinuities between adjacent frequency bins are minimised.



Smooth Decomposition

- ▶ In a smooth (analytic) decomposition, the eigenvalues and -vectors are arranged such that discontinuities between adjacent frequency bins are minimised.
- ▶ For a smooth decomposition, the eigenvectors in adjacent frequency bins are rearranged using the inner product

$$c_{mn}[k] = \mathbf{q}_m^H[k-1]\mathbf{q}_n[k],$$

where, $\mathbf{q}_m[k]$ is the m th column of $\mathbf{Q}[k]$.

- ▶ $c_{mn}[k] \approx 1$ if $\mathbf{q}_m[k-1]$ & $\mathbf{q}_n[k]$ aligned; $c_{mn}[k] \approx 0$ otherwise.
- ▶ Goal: reorder columns of $\mathbf{Q}[k]$ such that $c_{mm}[k] \approx 1 \forall m$.
- ▶ $\mathbf{\Lambda}[k]$ rearranged according to the reordering of the eigenvectors.

Phase Alignment

- ▶ Phase alignment of eigenvectors in adjacent frequency bins is vital for a compact-order decomposition.
- ▶ A matrix $\mathbf{C}^{(P)}$ [7] can be used to calculate the total power in the derivatives of a function up to and including the P th derivative.
 - ▶ Function ‘smoothness’ is characterised by a low total power.
- ▶ Phase of the m th eigenvector at frequency bin k can be adjusted by an angle θ_k according to $\mathbf{q}_m[k] \leftarrow e^{j\theta_k} \mathbf{q}_m[k]$.
- ▶ We compute a vector of phases $\boldsymbol{\theta} = [\theta_0, \dots, \theta_{K_i-1}]^T$ s.t. the m th eigenvector $\mathbf{q}_m[k] \forall k$ is maximally smooth.

[7] F. K. Coutts et al., Enforcing Eigenvector Smoothness for a Compact DFT-based Polynomial Eigenvalue Decomposition. *In Proc. IEEE Workshop on Sensor Array and Multichannel Sig. Process.*, July. 2018.

Phase Alignment

- ▶ An objective function has been derived in [7] that measures the smoothness of all elements of $\mathbf{q}_m[k]$ and takes the form

$$f(\boldsymbol{\theta}) = \mathbb{R}\{\mathbf{u}^H \boldsymbol{\Gamma} \mathbf{u}\} .$$

- ▶ $\mathbf{u}^H = [e^{j\theta_0}, \dots, e^{j\theta_{K_i-1}}]$, $\boldsymbol{\Gamma} = \sum_{n=0}^{M-1} \text{diag}\{\mathbf{v}_n\} \mathbf{C}_{(P)} \text{diag}\{\mathbf{v}_n^H\}$, and $\mathbf{v}_n = [\mathbf{q}_{m,n}[0], \dots, \mathbf{q}_{m,n}[K_i - 1]]$.
- ▶ $\mathbf{q}_{m,n}[k]$ denotes the n th element (row) of eigenvector $\mathbf{q}_m[k]$.
- ▶ We employ relatively low cost Powell's iterative 'dogleg' trust region strategy [8] for the unconstrained minimisation of $f(\boldsymbol{\theta})$.
- ▶ For $i > 0$, can use previous $\boldsymbol{\theta}$ to give a more informed initial guess.

[7] F. K. Coutts et al., Enforcing Eigenvector Smoothness for a Compact DFT-based Polynomial Eigenvalue Decomposition. *In Proc. IEEE Workshop on Sensor Array and Multichannel Sig. Process.*, July. 2018.

[8] M. J. D. Powell, A new algorithm for unconstrained optimization. *Nonlinear programming*, 31–65, 1970.

DFT-Based PEVD Performance

- ▶ DFT-based PEVD algorithm capable of outperforming existing methods.
- ▶ Suitable for scenarios with a high number of sensors.
- ▶ Example below has $\mathbf{R}(z) : \mathbb{C} \rightarrow \mathbb{C}^{5 \times 5}$ of order 38 with ground truth polynomial eigenvalues that are **not** spectrally majorised.

Method	MSE	Paraun. Err.	E_{diag}	Time / s	L_Q	Complexity
proposed	5.750×10^{-29}	2.887×10^{-22}	0	0.08854	64	$\mathcal{O}(ML^3)$
SBR2 [1]	1.815×10^{-6}	3.303×10^{-8}	10^{-6}	37.64	600.0	$\mathcal{O}(M^2L)$
SMD [2]	9.321×10^{-7}	3.847×10^{-8}	10^{-6}	11.34	357.9	$\mathcal{O}(M^3L)$

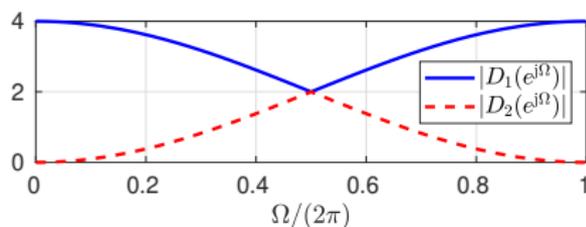
[1] J. G. McWhirter et al., An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Signal Process.*, 55(5):2158–2169, May 2007.

[2] S. Redif et al., Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans. on Signal Process.*, 63(1):81–89, Jan. 2015.

DFT-Based PEVD Performance

- Decompose the theoretical parahermitian matrix

$$\mathbf{R}(z) = \begin{bmatrix} 2 & z^{-1} + 1 \\ z + 1 & 2 \end{bmatrix}.$$



- Eigenvectors & eigenvalues are neither of finite order nor rational.
- To decompose $\mathbf{R}(z)$ via an exact PEVD would require polynomial matrices of infinite length.

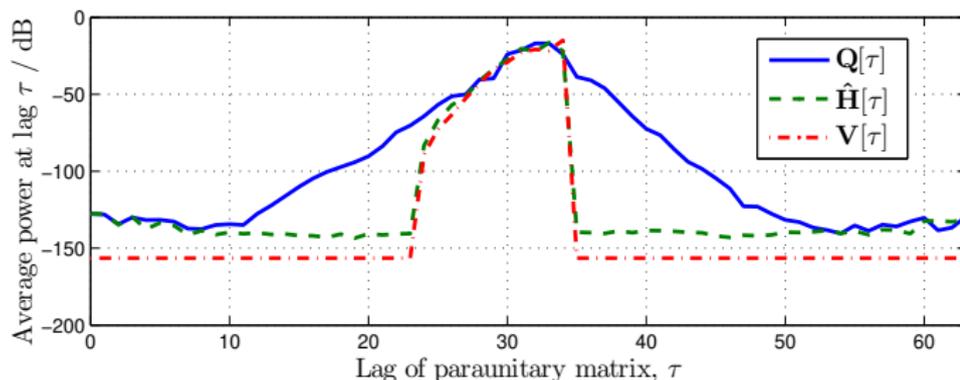
Method	MSE	Paraun. Err.	E_{diag}	Time / s	L_Q
proposed	7.077×10^{-9}	1.381×10^{-4}	0	0.1196	64
SMD, μ_1	4.362×10^{-25}	2.466×10^{-16}	10^{-6}	0.6256	345
SMD, μ_2	2.909×10^{-10}	9.546×10^{-8}	10^{-6}	0.1995	83
SBR2	2.909×10^{-10}	9.546×10^{-8}	10^{-6}	0.1724	83

Approximating a minimum-order solution to the PEVD

$$\mathbf{R}(z) = \mathbf{V}(z)\mathbf{\Lambda}(z)\mathbf{V}^P(z) \approx \mathbf{Q}(z)\mathbf{D}(z)\mathbf{Q}^P(z)$$

$$\mathbf{Q}(z) = \mathbf{H}(z)\mathbf{U}(z)$$

- ▶ Task: find all-pass filter bank $\mathbf{U}(z) = \text{diag}\{u_1(z), \dots, u_M(z)\}$.
- ▶ $u_m(z)$ defined by the greatest common divisor [9] of $\mathbf{q}_m(z)$.



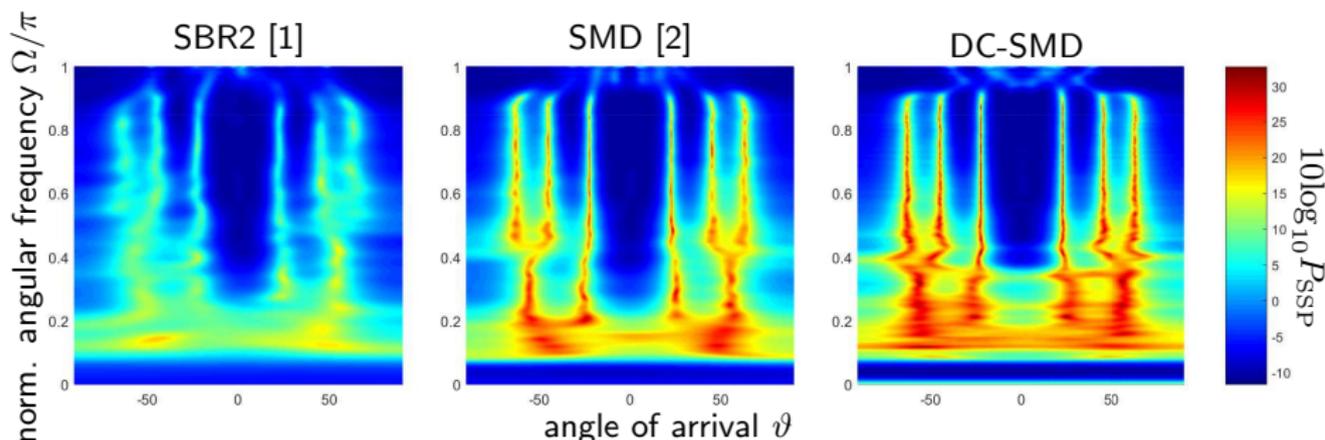
[9] F. C. Chang, Polynomial GCD derived through monic polynomial subtractions. *ISRN Applied Mathematics*, 2011.

Angle of Arrival Estimation (AoA)

- ▶ Polynomial MUSIC algorithm requires an accurate estimate of the noise-only subspace.
 - ▶ Accuracy strongly depends on PEVD quality.
- ▶ Each PEVD algorithm will produce a different level of AoA estimation performance.
 - ▶ More PEVD iterations and a longer algorithm runtime will generally increase performance at the cost of decreased flexibility.
- ▶ Divide-and-conquer strategies offer very fast diagonalisation performance, and are able to resolve weaker polynomial eigenvalues.
 - ▶ Does this translate to better AoA estimation performance if simulation time is fixed?

Angle of Arrival Estimation Results

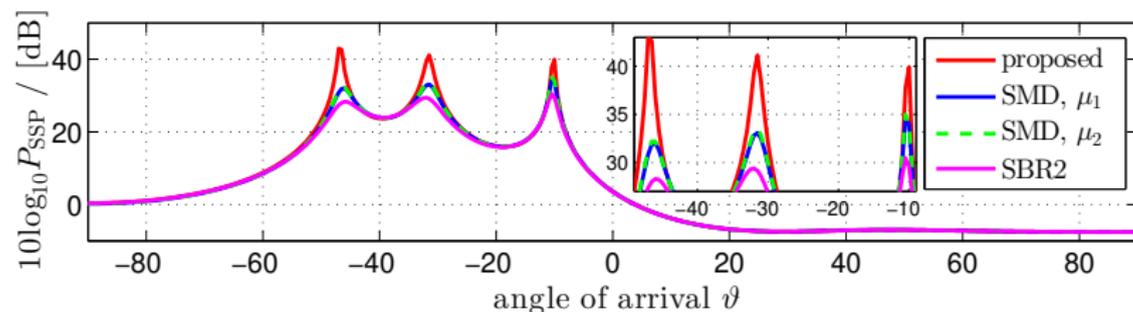
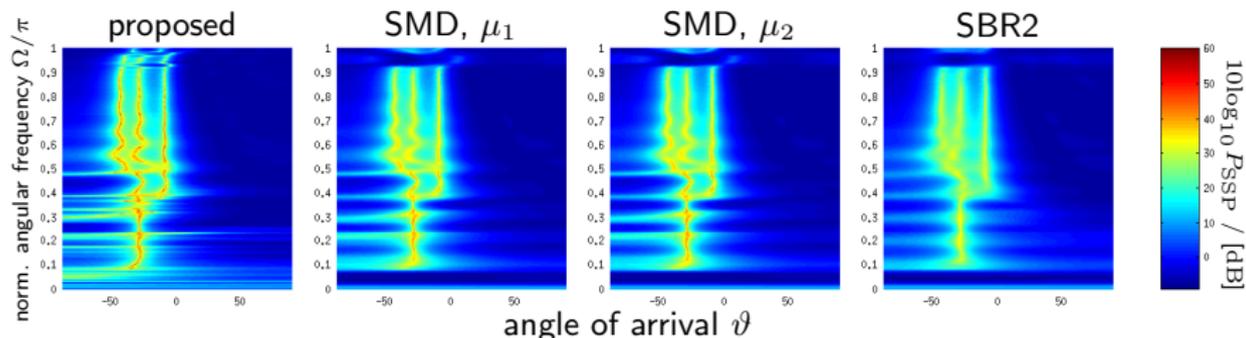
- ▶ ‘Divide-and-conquer’ (DaC) approach to the PEVD:
- ▶ 6 sources sharing frequency range $\Omega \in [0.1\pi, 0.9\pi]$, 20 dB SNR.



-
- [1] J. G. McWhirter et al., An EVD Algorithm for Para-Hermitian Polynomial Matrices. *IEEE Trans. on Signal Process.*, 55(5):2158–2169, May 2007.
 - [2] S. Redif et al., Sequential Matrix Diagonalisation Algorithms for Polynomial EVD of Parahermitian Matrices. *IEEE Trans. on Signal Process.*, 63(1):81–89, Jan. 2015.

Angle of Arrival Estimation Results

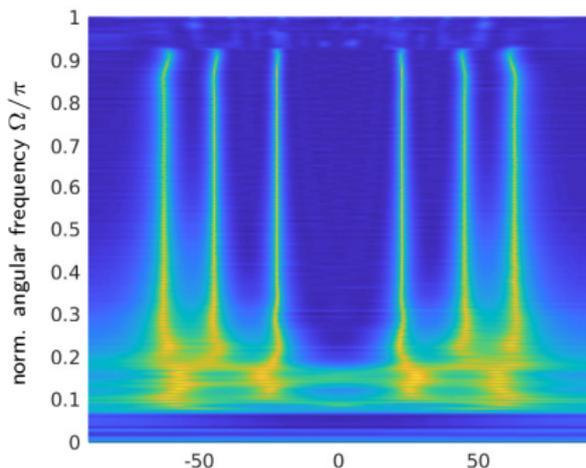
- ▶ Frequency-Based PEVD Algorithms:
- ▶ 3 sources with different frequency ranges, 20 dB SNR.



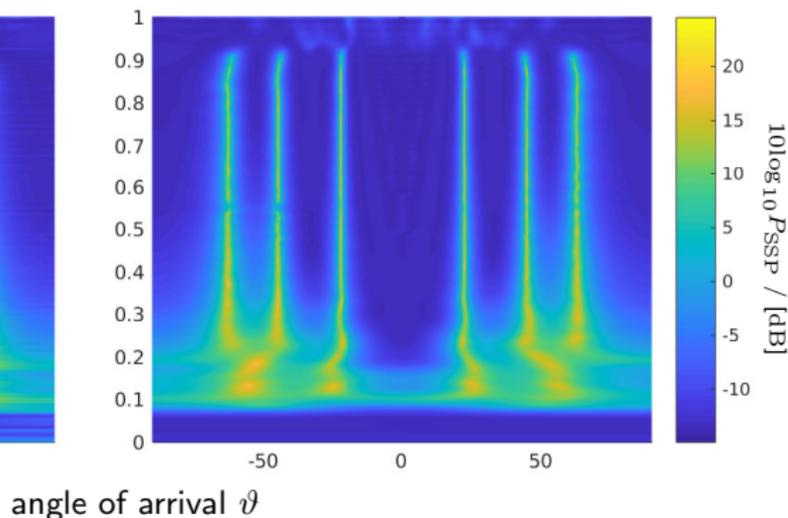
AoA Comparison for DFT-Based and PSMD Algorithms

- ▶ Both algorithms executed for 1.5 seconds
- ▶ $M = 24$
- ▶ 6 sources

DFT-based

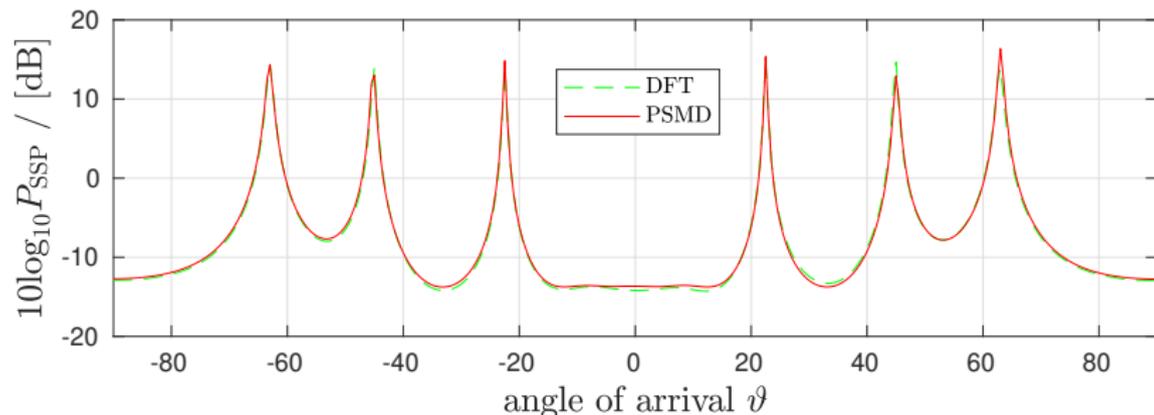


PSMD



AoA Comparison for DFT-Based and PSMD Algorithms

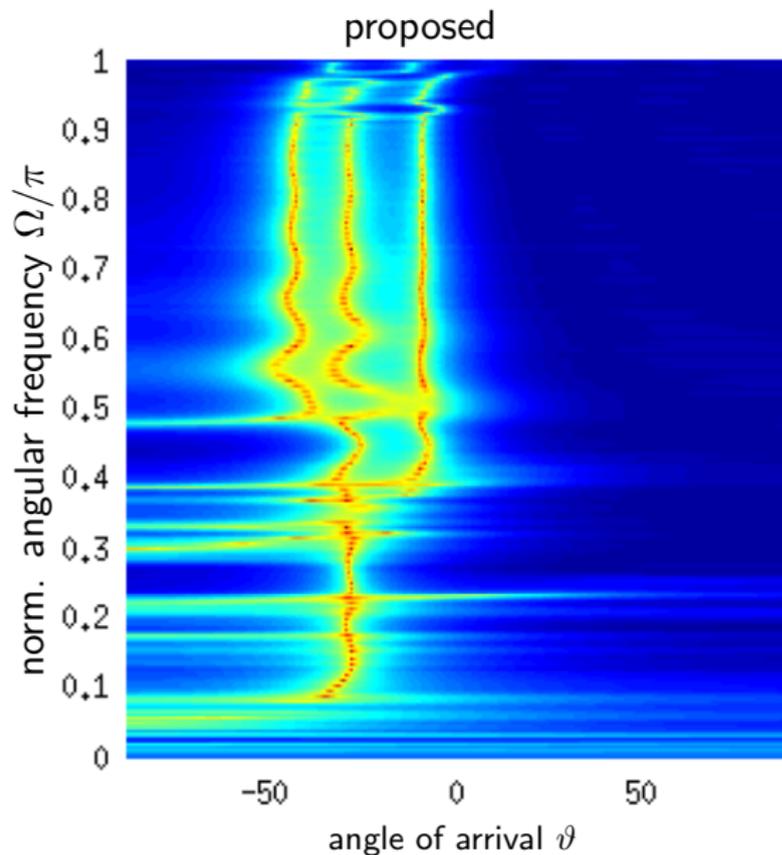
- ▶ Both algorithms executed for 1.5 seconds
- ▶ $M = 24$ sensors
- ▶ 6 sources
- ▶ Evaluated at $\Omega = \pi/2$
- ▶ Similar levels of performance



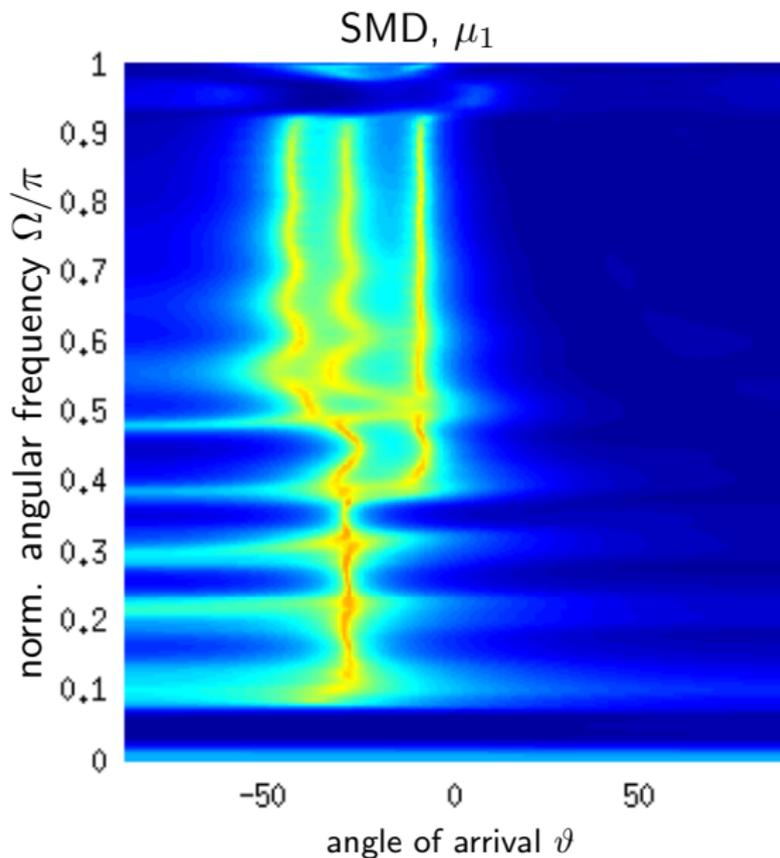
Summary

- ▶ Summarised existing iterative PEVD algorithms.
- ▶ Introduced some methods used to lower the computational cost of these algorithms.
- ▶ Gave overview of parallelised 'divide-and-conquer' approach to the PEVD and highlighted implementation benefits.
- ▶ Explained DFT-based PEVD approaches and their advantages.
- ▶ Demonstrated the effects of PEVD algorithm choice on AoA estimation results.
- ▶ `fraser.coutts@ed.ac.uk` | <https://tinyurl.com/couttsphd>

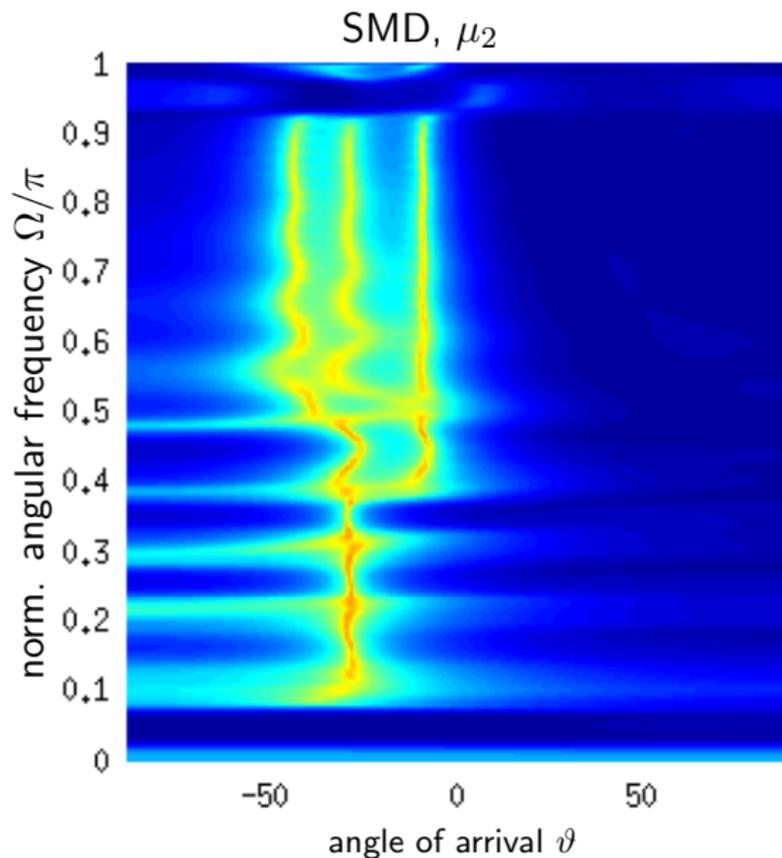
Simulation Scenario 3 Zoomed



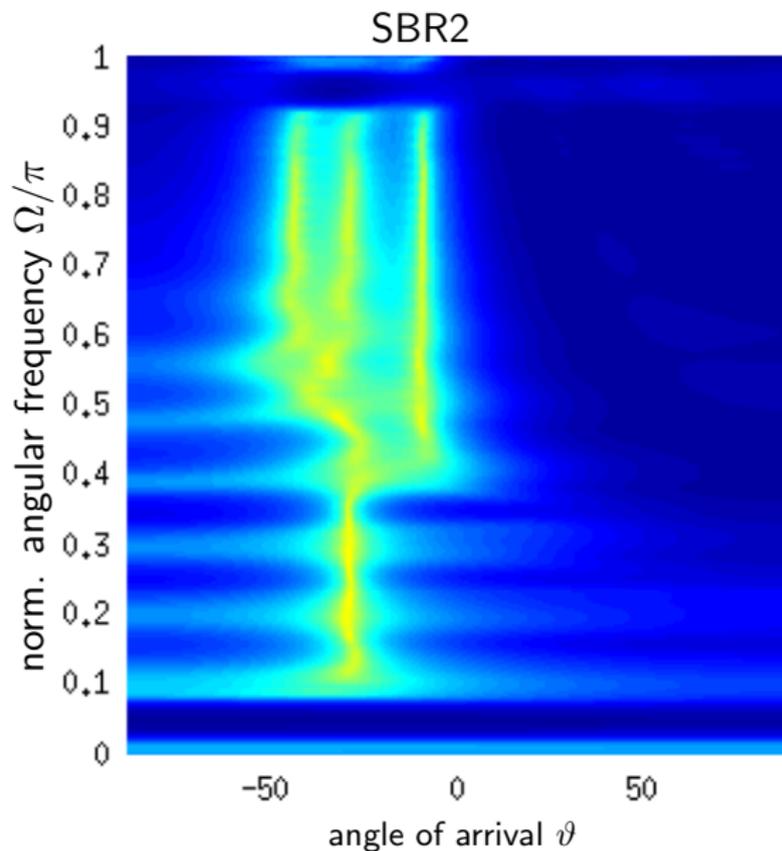
Simulation Scenario 3 Zoomed



Simulation Scenario 3 Zoomed



Simulation Scenario 3 Zoomed

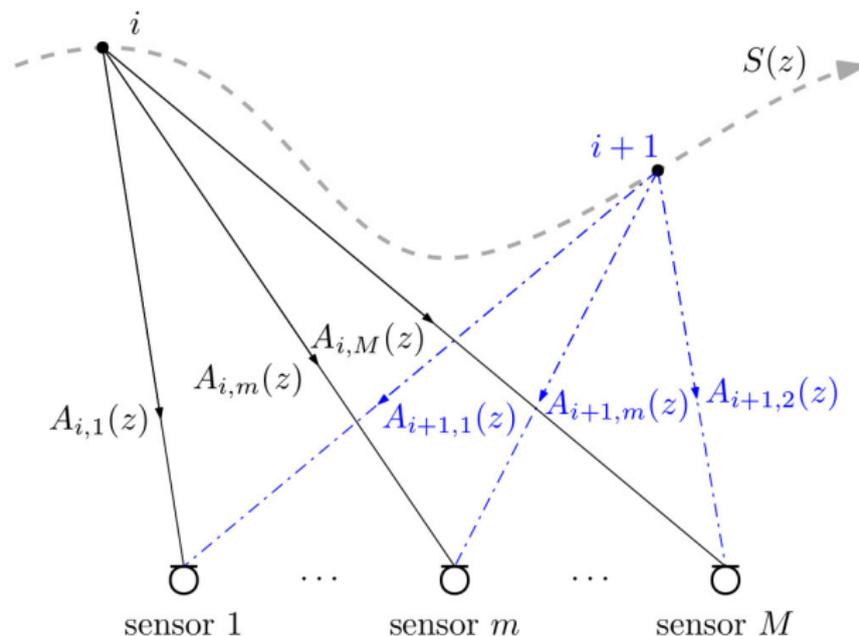


Future Work

- ▶ Impact of divide-and-conquer algorithms on spectral majorisation.
- ▶ Investigate scenarios with ‘naturally’ (up to permutations) block diagonal matrices.
- ▶ Alternative permutation/optimisation strategies for the DFT-based PEVD.
- ▶ Minimum-order solutions and impulse response estimation.
- ▶ Deploying developed algorithms in the real world.

Future Work

- ▶ Minimum-order solutions and impulse response estimation:



-
- [9] F. C. Chang, Polynomial GCD derived through monic polynomial subtractions. *ISRN Applied Mathematics*, 2011.
- [10] S. Weiss et al., Identification of broadband source-array responses from sensor second order statistics. *IEEE Sensor Signal Processing for Defence Conference*, Dec. 2017

Future Work

- ▶ Minimum-order solutions and impulse response estimation:
 - ▶ Cross-spectral density matrix $\mathbf{R}_i(z)$ defined by transfer functions, $\mathbf{a}_i(z)$, and source power spectral density, $\mathbf{S}(z)$:

$$\mathbf{R}_i(z) = \mathbf{a}_i(z)S(z)\mathbf{a}_i^P(z) + \sigma_n^2\mathbf{I}_M \approx \mathbf{q}_i(z)d_i(z)\mathbf{q}_i^P(z) + \sigma_n^2\mathbf{I}_M$$

$$\mathbf{a}_i^P(z)\mathbf{a}_i(z) = A_i(z) = A_i^{(+)}(z)A_i^{(+),P}(z), \quad \mathbf{a}_{i,\text{norm}}(z) = \frac{\mathbf{a}_i(z)}{A_i^{(+)}(z)}$$

$$\mathbf{R}_i(z) \approx \mathbf{a}_{i,\text{norm}}(z)A_i^{(+)}(z)S(z)A_i^{(+),P}(z)\mathbf{a}_{i,\text{norm}}^P(z) + \sigma_n^2\mathbf{I}_M$$

$$\mathbf{q}_i(z) = \frac{\mathbf{a}_i(z)}{A_i^{(+)}}, \quad d_i(z) = A_i^{(+)}(z)S(z)A_i^{(+),P}(z)$$

[9] F. C. Chang, Polynomial GCD derived through monic polynomial subtractions. *ISRN Applied Mathematics*, 2011.

[10] S. Weiss et al., Identification of broadband source-array responses from sensor second order statistics. *IEEE Sensor Signal Processing for Defence Conference*, Dec. 2017

Future Work

- ▶ Minimum-order solutions and impulse response estimation:

$$\mathbf{q}_i(z) = \frac{\mathbf{a}_i(z)}{A_i^{(+)}}, \quad d_i(z) = A_i^{(+)}(z)S(z)A_i^{(+)\text{P}}(z)$$

$$\hat{S}(z) = \text{GCD} \{d_1(z), \dots, d_I(z)\}$$

- ▶ Phase ambiguity in $\mathbf{q}_i(z) = \hat{\mathbf{q}}_i(z)u_i(z)$ can be eliminated through determination of greatest common divisor (GCD) $u_i(z)$ [9].

$$\hat{\mathbf{a}}_i(z) = \hat{A}_i^{(+)}(z)\hat{\mathbf{q}}_i(z)$$

- ▶ Magnitude **and** phase of transfer functions recovered.
- ▶ Next problem: identifying GCD of noisy eigenvalues.

[9] F. C. Chang, Polynomial GCD derived through monic polynomial subtractions. *ISRN Applied Mathematics*, 2011.

[10] S. Weiss et al., Identification of broadband source-array responses from sensor second order statistics. *IEEE Sensor Signal Processing for Defence Conference*, Dec. 2017

References I

- [1] J. G. McWhirter, P. D. Baxter, T. Cooper, S. Redif, and J. Foster.
An EVD Algorithm for Para-Hermitian Polynomial Matrices.
IEEE Transactions on Signal Processing, 55(5):2158–2169, May 2007.
- [2] 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.
- [3] F. K. Coutts, I. K. Proudler, and S. Weiss.
Efficient implementation of iterative polynomial matrix evd algorithms exploiting structural redundancy and parallelisation.
IEEE Transactions on Circuits and Systems I: Regular Papers, 66(12):4753–4766, Dec. 2019.
- [4] M. Tohidian, H. Amindavar, and A. M. Reza.
A DFT-based approximate eigenvalue and singular value decomposition of polynomial matrices.
EURASIP Journal on Advances in Signal Processing, 2013(1):1–16, 2013.
- [5] F. Coutts, K. Thompson, S. Weiss, and I. Proudler.
A comparison of iterative and DFT-based polynomial matrix eigenvalue decompositions.
In *IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, Curacao, December 2017.
- [6] S. Weiss, J. Pestana, and I. K. Proudler.
On the existence and uniqueness of the eigenvalue decomposition of a parahermitian matrix.
IEEE Transactions on Signal Processing, 66(10):2659–2672, May 2018.
- [7] F. K. Coutts, K. Thompson, J. Pestana, I. Proudler, and S. Weiss.
Enforcing eigenvector smoothness for a compact DFT-based polynomial eigenvalue decomposition.
In *10th IEEE Workshop on Sensor Array and Multichannel Signal Processing*, pages 1–5, July 2018.
- [8] M. Powell.
A new algorithm for unconstrained optimization.
In J. Rosen, O. Mangasarian, and K. Ritter, editors, *Nonlinear Programming*, pages 31–65. Academic Press, 1970.

References II

- [9] F. C. Chang.
Polynomial GCD derived through monic polynomial subtractions.
ISRN Applied Mathematics, 2011(714102), 2011.
- [10] S. Weiss, N. J. Goddard, S. Somasundaram, I. K. Proudler, and P. A. Naylor.
Identification of broadband source-array responses from sensor second order statistics.
In Sensor Signal Processing for Defence Conference, pages 1–5, London, UK, December 2017.
- [11] C. H. Ta and S. Weiss.
Shortening the order of paraunitary matrices in SBR2 algorithm.
In 6th International Conference on Information, Communications & Signal Processing, pages 1–5, Singapore, Dec. 2007.
- [12] S. Redif, S. Weiss, and J. G. McWhirter.
An approximate polynomial matrix eigenvalue decomposition algorithm for para-hermitian matrices.
In 11th IEEE International Symposium on Signal Processing and Information Technology, pages 421–425, Bilbao, Spain, December 2011.
- [13] J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter.
Reduced search space multiple shift maximum element sequential matrix diagonalisation algorithm.
In IET/EURASIP Intelligent Signal Processing, London, UK, December 2015.
- [14] F. K. Coutts, J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. G. McWhirter.
Memory and complexity reduction in parahermitian matrix manipulations of PEVD algorithms.
In 24th European Signal Processing Conference, Budapest, Hungary, Aug. 2016.
- [15] S. Weiss, J. Pestana, I. K. Proudler, and F. K. Coutts.
Corrections to “on the existence and uniqueness of the eigenvalue decomposition of a parahermitian matrix”.
IEEE Transactions on Signal Processing, 66(23):6325–6327, Dec 2018.
- [16] S. Weiss, I. K. Proudler, F. K. Coutts, and J. Pestana.
Iterative approximation of analytic eigenvalues of a parahermitian matrix EVD.
In IEEE International Conference on Acoustics, Speech and Signal Processing, Brighton, UK, May 2019.

References III

- [17] S. Weiss, I. Proudler, F. Coutts, and J. Deeks.
Extraction of analytic eigenvectors from a parahermitian matrix.
In International Conference on Sensor Signal Processing or Defence, Edinburgh, UK, 2020.
- [18] S. Weiss, I. K. Proudler, and F. K. Coutts.
Eigenvalue decomposition of a parahermitian matrix: extraction of analytic eigenvalues.
IEEE Transactions on Signal Processing, 69:722–737, 2021.
- [19] S. Weiss, I. K. Proudler, G. Barbarino, J. Pestana, and J. G. McWhirter.
On properties and structure of the analytic singular value decomposition.
IEEE Transactions on Signal Processing, 2023.
to be submitted.
- [20] F. Coutts, J. Corr, K. Thompson, S. Weiss, I. Proudler, and J. McWhirter.
Complexity and search space reduction in cyclic-by-row PEVD algorithms.
In 50th Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, November 2016.
- [21] F. Coutts, K. Thompson, S. Weiss, and I. Proudler.
Impact of fast-converging PEVD algorithms on broadband AoA estimation.
In Sensor Signal Processing for Defence Conference, pages 1–5, London, UK, December 2017.
- [22] F. Coutts, J. Corr, K. Thompson, I. Proudler, and S. Weiss.
Divide-and-conquer sequential matrix diagonalisation for parahermitian matrices.
In Sensor Signal Processing for Defence Conference, pages 1–5, London, UK, December 2017.
- [23] F. Coutts, J. Corr, K. Thompson, S. Weiss, P. I.K., and J. McWhirter.
Multiple shift QR decomposition for polynomial matrices.
In IMA International Conference on Mathematics in Signal Processing, Birmingham, UK, December 2016.
- [24] F. Coutts, K. Thompson, I. Proudler, and S. Weiss.
Restricted update sequential matrix diagonalisation for parahermitian matrices.
In IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Curacao, December 2017.

References IV

- [25] F. K. Coutts, K. Thompson, J. Pestana, I. Proudler, and S. Weiss.
Enforcing eigenvector smoothness for a compact DFT-based polynomial eigenvalue decomposition.
In 10th IEEE Workshop on Sensor Array and Multichannel Signal Processing, pages 1–5, July 2018.
- [26] F. K. Coutts, K. Thompson, I. K. Proudler, and S. Weiss.
An iterative dft-based approach to the polynomial matrix eigenvalue decomposition.
In 52nd Asilomar Conference on Signals, Systems, and Computers, pages 1011–1015, Oct 2018.
- [27] F. K. Coutts, K. Thompson, S. Weiss, and I. K. Proudler.
Analysing the performance of divide-and-conquer sequential matrix diagonalisation for large broadband sensor arrays.
In 2017 IEEE International Workshop on Signal Processing Systems, pages 1–6, Lorient, France, October 2017.
- [28] F. Coutts, J. Corr, K. Thompson, I. Proudler, and S. Weiss.
Divide-and-conquer sequential matrix diagonalisation for parahermitian matrices.
In Sensor Signal Processing for Defence Conference, pages 1–5, London, UK, December 2017.
- [29] F. Coutts, J. Corr, K. Thompson, S. Weiss, P. I.K., and J. McWhirter.
Multiple shift QR decomposition for polynomial matrices.
In IMA International Conference on Mathematics in Signal Processing, Birmingham, UK, December 2016.
- [30] F. Coutts, K. Thompson, I. Proudler, and S. Weiss.
Restricted update sequential matrix diagonalisation for parahermitian matrices.
In IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, Curacao, December 2017.
- [31] F. K. Coutts, K. Thompson, J. Pestana, I. Proudler, and S. Weiss.
Enforcing eigenvector smoothness for a compact DFT-based polynomial eigenvalue decomposition.
In 10th IEEE Workshop on Sensor Array and Multichannel Signal Processing, pages 1–5, July 2018.
- [32] F. K. Coutts, K. Thompson, I. K. Proudler, and S. Weiss.
An iterative dft-based approach to the polynomial matrix eigenvalue decomposition.
In 52nd Asilomar Conference on Signals, Systems, and Computers, pages 1011–1015, Oct 2018.