

Covariance matching techniques for radio astronomy calibration and imaging

Mouri Sardarabadi, Ahmad

DOI

[10.4233/uuid:3eb46b5b-126c-4d95-8821-fb71ce404cac](https://doi.org/10.4233/uuid:3eb46b5b-126c-4d95-8821-fb71ce404cac)

Publication date

2016

Document Version

Final published version

Citation (APA)

Mouri Sardarabadi, A. (2016). *Covariance matching techniques for radio astronomy calibration and imaging*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:3eb46b5b-126c-4d95-8821-fb71ce404cac>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

COVARIANCE MATCHING TECHNIQUES FOR RADIO ASTRONOMY CALIBRATION AND IMAGING

COVARIANCE MATCHING TECHNIQUES FOR RADIO ASTRONOMY CALIBRATION AND IMAGING

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K. C. A. M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 6 juni 2016 om 12:30 uur

door

Ahmad MOURI SARDARABADI

elektrotechnisch ingenieur, Technische Universiteit Delft
geboren te Ahvaz, Iran.

Dit proefschrift is goedgekeurd door de promotor:

Prof. dr. ir. A-J. van der Veen

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. A-J. van der Veen,	Technische Universiteit Delft, promotor

Onafhankelijke commissieleden:

Prof. dr. ir. G.J.T. Leus	Technische Universiteit Delft
Prof. dr. ir. M.H.G. Verhaegen	Technische Universiteit Delft
Prof. dr. G. de Bruyn	Rijksuniversiteit Groningen
Prof. dr. Y. Wiaux	Heriot-Watt University Edinburgh
Prof. dr. A. Leshem	Bar-Ilan University Israël
Dr. ir. S.J. Wijnholds	ASTRON

Copyright © 2016 by A. Mouri Sardarabadi

The research presented in this work has been funded by NWO-TOP GO project "The synthesis of signal processing and radio astronomical calibration and imaging techniques" under number 614.001.005.

ISBN 978-94-6186-670-7

An electronic version of this dissertation is available at

<http://repository.tudelft.nl/>.

CONTENTS

1	Introduction	1
1.1	Motivation	1
1.1.1	Signal Processing Approach To Radio Astronomy	2
1.1.2	RFI Mitigation	3
1.1.3	Gain Calibration	4
1.1.4	Image Reconstruction	4
1.2	Problem Definition	5
1.3	Approach	6
1.3.1	Subspace estimation	6
1.3.2	Kronecker structures and Krylov solvers	7
1.3.3	Constrained Least Squares	7
1.4	Outline	7
1.5	Contributions to Literature	8
2	Preliminaries	11
2.1	Notations	12
2.2	Properties of the Kronecker Product	13
2.3	Covariance Matching Techniques	13
2.3.1	Statistical Sufficiency	14
2.3.2	Maximum Likelihood	14
2.3.3	Method of Scoring	15
2.3.4	Weighted Least Squares	16
2.3.5	Reduced-Gradient method	17
2.4	Krylov Subspace Based Solvers	18
2.4.1	Minimum Polynomial and Krylov subspaces	18
2.4.2	Lanczos algorithm and LSQR	19
2.4.3	Tridiagonalization and Krylov subspaces	21
3	Data Model	23
3.1	Array Response Vector	23
3.1.1	Narrow-Band Assumption	23
3.1.2	Geometrical Delays and the Earth's Rotation	25
3.2	Covariance Model	26
3.3	Gain Model	28
3.4	Non-Parametric Subspace Model	29

4	Complex Factor Analysis	31
4.1	Introduction	31
4.2	Classical Factor Analysis Model	32
4.2.1	Data Model	32
4.2.2	Factor Rotation	34
4.2.3	Scale Invariance	35
4.3	Extensions of the Classical Model	35
4.3.1	Joint Factor Analysis Model	35
4.3.2	Extended and Joint Extended FA model	35
4.4	Cramér-Rao Bound	36
4.4.1	CRB for FA and EFA	36
4.4.2	CRB for JEFA	39
4.5	Estimation Algorithms for Classical FA	40
4.5.1	Ad-hoc Method	41
4.5.2	Maximum Likelihood Estimator	41
4.5.3	Minimum Trace FA	43
4.6	Estimation Algorithms for JEFA	44
4.6.1	Non-linear Weighted Least Squares	45
4.6.2	Krylov-Based Methods	46
4.6.3	Maximum Likelihood for JEFA	48
4.7	Goodness of fit and Detection	49
4.8	Simulations	49
4.8.1	Convergence Speed	51
4.8.2	Subspace Estimation Performance	52
4.8.3	DOA Estimation using JEFA	53
4.8.4	Cramér-Rao Bound simulation	55
4.9	Conclusion	56
4.10	Appendix: Identifiability	56
5	RFI Mitigation Using A Reference Antenna Array	61
5.1	Introduction	61
5.2	Problem statement	64
5.2.1	Data model	64
5.2.2	Covariance model	65
5.3	Existing Spatial Filtering Algorithms	66
5.3.1	Spatial Filtering Using Projections Without Reference Antennas	66
5.3.2	Spatial Filtering Using Reference Antennas	68
5.4	Improved Filtering	69
5.4.1	Factor Analysis	69
5.4.2	Direct ML Estimation Using Extended Factor Analysis	70
5.5	Performance analysis	71
5.5.1	Projection Based Spatial Filtering	71
5.5.2	The CRB of Ψ_{00} and asymptotic statistics of (J)EFA	71
5.5.3	Performance for Long Term Integration (Imaging)	72

5.6	Simulations	73
5.7	Examples on Experimental Data	77
5.7.1	Experiment I	77
5.7.2	Experiment II	79
5.8	Conclusions.	80
6	Direction Independent Gain Calibration	83
6.1	Introduction	83
6.2	Data Model and Problem Definition	84
6.3	Algorithm	84
6.3.1	Covariance Matching	84
6.3.2	Krylov Subspace Based Methods	86
6.3.3	Improving the DOA Estimates	87
6.4	Experimental Data	89
6.5	Conclusion	90
7	Image Reconstruction	93
7.1	Introduction	93
7.2	The Imaging Problem	94
7.2.1	Gridded Imaging Model	95
7.2.2	Unconstrained Least Squares Image	96
7.2.3	Preconditioned Weighted Least Squares Image	97
7.3	Bounds on the Image	98
7.3.1	Matched Filter as Upper Bound	98
7.3.2	Tightest Upper Bound and ASSC	99
7.3.3	Multi-Snapshot MVDR	101
7.3.4	Estimation of the Upper Bound from Noisy Data.	102
7.3.5	Constrained Least Squares Imaging	103
7.4	Constrained optimization.	103
7.4.1	Characterization of the Optimum	104
7.4.2	Gradients	105
7.4.3	Active Set Methods.	106
7.4.4	Strong Off-Grid Sources	108
7.4.5	Boxed Imaging	109
7.5	Implementation using Krylov Subspace based methods	110
7.5.1	Implementation	110
7.6	Simulations	112
7.6.1	Extended Sources	112
7.6.2	Full Sky with 3C Sources	115
7.7	Conclusions.	116
7.8	Appendix: Relation between WLS, Natural and Robust Weighting.	119
7.9	Appendix: Variance of the Dirty Image	120

8	Conclusions and suggestions for further research	123
8.1	Summary of Main Results	123
8.2	Discussion	125
8.2.1	Subspace Estimation	125
8.2.2	Spatial Filtering Using a Reference Array	125
8.2.3	Gain Calibration	125
8.2.4	Imaging	126
8.3	Future research	126
8.3.1	Factor Analysis	126
8.3.2	Calibration	127
8.3.3	Imaging	127
	Bibliography	129
	References	129
	Summary	139
	Samenvatting	141
	Acknowledgment	145
	Biography	147

1

INTRODUCTION

1.1. MOTIVATION

After the first discovery of radio emissions from our galaxy by Karl Jansky in 1933, the field of radio astronomy has provided scientists and the general public with discoveries that have shaped our understanding of the universe. The discovery of quasars and supermassive Black Holes at the center of the galaxies, Cosmic Microwave Background radiation which provided support for the Big Bang model and the Pulsating Radio Sources (Pulsars) are a few contributions from radio astronomy in the last 80 years.

Studying signals from the early universe such as the Epoch of Reionization (EoR) (a period a few hundred million years after the Big Bang when the gasses in the universe started the transition from being completely neutral to completely ionized) demands radio-telescopes with very high sensitivity and resolution. The sensitivity of a radio-telescope increases with its collecting area and its resolution is proportional to its geometrical dimensions and inversely proportional to the wavelength at which the observation are taking place. For studying weak signals at low frequencies (~ 100 MHz) with high resolution (\sim arcsec), the dimension of the telescope will grow beyond 600 km. Fig. 1.1a shows the Arecibo telescope with a diameter of 305m, it is then clear that creating a single dish telescope with dimensions of hundreds of kilometers is practically impossible.

The idea to combine the output of different radio receivers to emulate a larger radio-telescope is mostly contributed to Martin Ryle, who received a Nobel Prize together with Tony Hewish for his work on one the first radio-interferometers called the One-Mile telescope which was constructed at Cambridge, UK and consisted of two receivers. Soon after, in 1970 the 3km Westerbork Synthesis Radio Telescope (WSRT) was constructed in the north of the Netherlands with 14 antennas (Fig. 1.1b) followed by the 36 km Very Large Array (VLA) in USA (1980, Fig. 1.1c) and the 25 km Giant Meter-wave Radio Telescope (GMRT) in India (1988, Fig. 1.1d).

Modern radio-telescopes such as the LOw Frequency ARray (LOFAR) [1] consist of many small receiving elements (~ 3500 in low-band and ~ 28000 in high-band modes



(a) Arecibo



(b) WSRT



(c) VLA



(d) GMRT

Figure 1.1: Various radio-telescopes with historical importance

[2, pp. 8]) with baselines of over 1000 km. Future telescopes such as Square Kilometer Array (SKA) will have order(s) of magnitude more elements and collecting area. However, increasing sensitivity and resolution by itself will not lead to the detection of the desired weak signals. There exists a huge number of foreground radiators including, not only astronomical objects but also strong transmissions from man-made radio systems which must be removed in order to make it possible for the weaker sources to be detected. This is clearly a detection and estimation problem, and such problems are studied extensively in signal processing applications.

1.1.1. SIGNAL PROCESSING APPROACH TO RADIO ASTRONOMY

Considering the the rapid increase in the number of receivers, the dimensions and the field of view (FoV) of the modern telescopes, it is clear that the current algorithms used by astronomers need to be improved. With LOFAR fully operational and the construction of SKA gaining momentum, the demand for automated and scalable algorithms is increasing. At the same time it is important for these algorithms to be based on a sound theoretical foundation in order to make the analysis of the corresponding (large) prob-




Figure 1.2: Schematic representation of main components of imaging pipeline.

lems possible in a systematic manner.

Reformulation of the interferometric measurement equation (e.g. the model presented in [3]) using signal and array processing models, will allow access to the rich set of signal processing tools developed for these fields. A given problem in signal processing is approached by finding the appropriate data models, theoretical bounds needed for performance analysis and algorithms. Signal processing also provides theoretical tools to perform each of these steps in a systematic way, e.g. algorithms can be designed based on proven generic principles such as maximum likelihood (ML) and/or least squares (LS).

The benefit of using the signal processing methods can be seen in several works that have introduced this approach to the radio astronomy such as [4–8]. However, many of the suggested signal processing based approaches to radio astronomy are either unscalable or their performance is not yet tested on large data sets. Another problem is the fact that the results from some of the proposed methods, even though theoretically better motivated, do not provide the desired improvements with respect to classical methods used by astronomers. An example is the least square imaging [9] compared to the classical approaches such as CLEAN (discussed in more detail in Sec. 1.1.4 and Chapter 7). In this thesis we will expand on these models and study how signal processing and linear algebraic approaches can be used to improve the existing algorithms and/or provide new tools capable of addressing some of the issues that new generation of radio telescopes will face.

As illustrated in Fig. 1.2 we can divide the data processing pipeline, from the receivers to the final image, into three closely related main components: radio frequency interference (RFI) mitigation, calibration and imaging. We will discuss some of the problems related to these components in short in the following sections.

1.1.2. RFI MITIGATION

The increasing number of wireless services and applications are rapidly reducing the number of RFI free bands for radio astronomy. Many of these services are moving towards wideband digital systems such as Digital Audio Broadcasting (DAB) currently pop-

ulating the band between 174-240 MHz which also is of interest for radio telescopes such as LOFAR.

The current radio frequency interference (RFI) mitigation techniques which are commonly used by the astronomical community (an overview can be found in [10–12]) use the time-frequency plane to flag and remove the portion of the data that is suspected to be contaminated. Another approach is the use of specialized imaging techniques which are to some degree robust against RFI presence [5]. In situations where there exists a continuously present wide-band RFI, the use of array processing techniques such as spatial filtering could provide access to frequency bands otherwise avoided by astronomers.

An example of such a spatial filtering technique is proposed by [13] and an improved version which uses an auxiliary set of antennas as a reference array to increase the reliability and numerical stability of the algorithm is proposed by [14]. These spatial filtering techniques require the spatial signature of the RFI, which can be found using subspace estimation techniques, in order to suppress them. The eigenvalue decomposition (EVD) is a popular subspace estimation method in signal processing, however it requires the noise on the data to be white or known beforehand (e.g. from calibration) such that the data can be preprocessed by a whitening step. This requirement could be very limiting in practice, especially because the presence of RFI reduces the calibration accuracy. It is then preferred to be able to estimate the required subspace using more generic approaches which will enable us to combine RFI filtering with noise power calibration.

1.1.3. GAIN CALIBRATION

Gain calibration is another important step towards producing high quality images from the measured data [4, 6, 7]. During this step the currently available model for the sky is used to estimate and correct for instrumental effects which will otherwise complicate the imaging algorithms. The data model for gain calibration is non-linear and the number of unknowns grows with both the number of elements in the array and calibration sources. It is then clear that calibration for modern radio-telescopes which consist of a large number of receivers and are sensitive to a much larger number of sources due to their large collecting area and resolution, will grow very rapidly in complexity.

To cope with this increased complexity, current approaches split the unknowns into groups for which a closed-form or simplified solution can be found [4, 15]. Then by applying alternating optimization algorithms such as weighted alternating least squares (WALS), a solution is found. The penalty for approaching this problem by alternating methods is the inherent monotonic convergence of these techniques which could be too slow in some applications. Another way to reduce the growing complexity of the calibration problem is by exploiting structure of the matrices used during the estimation process. Using a signal processing model it can be shown that these matrices possess a strong structure which, combined with the appropriate solvers, can be used to reduce the computational cost and achieve accurate results with faster convergence rate.

1.1.4. IMAGE RECONSTRUCTION

Both classical Fourier based imaging [3] and the parametric model based least square imaging [9] can be divided into two steps. The first step is the construction of the so called “dirty image” followed by a reconstruction step where an estimate of the true

image called the “clean image” is found. For a small FoV and uniform spatial sampling of measurement data, the relation between the dirty image and the true image is a 2D convolution and the recovery of the clean image is a deconvolution process. Even though this exact convolution relation fails to exist for an image with a large FoV and nonuniform spatial samples, the term deconvolution is still used to denote the recovery step. It is well known that the deconvolution step becomes an ill-conditioned problem, if the number of image pixels is large [9]. This has led to development of several deconvolution techniques which use different assumptions on the structure of the image to cope with this problem.

One of the well known and widely used deconvolution techniques is the CLEAN method proposed by Högbom [16], which was subsequently refined and extended in several ways, leading to approaches described in [17–19]. CLEAN and its derivatives can be categorized under sequential source removing techniques, where at each iteration one or a few (point) sources are detected and removed from the dirty image. This process is repeated until the residual (dirty) image can be considered as noise.

Other techniques based on non-negativity of the image [20], sparsity or ℓ_1 constraints on the image [8, 21] and/or a combination of wavelets and sparsity [22, 23] have been proposed to improve the quality of the reconstructed image. Several parameter estimation methods based on array processing techniques have also been proposed, such as the Least Squares (LS) and Least Squares Minimum Variance Imaging (LS-MVI) [9, 24] and maximum likelihood based techniques [25].

The computational complexity of these methods is much higher than classical approaches such as CLEAN and this has formed a barrier for their adoption by astronomers. On the other hand the lack of a comprehensive theoretical motivation for classical approaches makes their analysis difficult. Such analysis is needed and could provide insights which can be used to fill the gap between the parametric imaging techniques and classical ones in order to benefit from both approaches.

1.2. PROBLEM DEFINITION

Consolidating the above discussion, in this thesis the following problem and its related subproblems are addressed: “*How can a signal processing formalism (data model and algorithms) be used to (re)formulate the radio astronomical image formation problems, from the data acquisition at individual receivers to the final image? Which SP tools are appropriate and/or needed?*”. We approach this problem in a hierarchical manner. The possibility of producing an accurate sky-map/model or *imaging* can be studied under the assumption that no RFI is present and that the receivers are ideal and identical elements. Correcting for the deviation between the actual receivers and this ideal array, which is called *calibration*, can be addressed separately by again assuming no RFI presence and some prior knowledge about the sky-model. The next problem that needs to be addressed is the detection and mitigation of RFI.

In the majority of radio astronomical observations, the frequency bands which are contaminated by RFI are thrown away to avoid possible artifacts in the final image. This loss of data results in loss of sensitivity and should be avoided if possible. We will investigate how subspace-based techniques needed in spatial filtering algorithms can be improved to produce RFI free covariance matrices from contaminated measurements

and what tools are needed. The following subproblems form the basis for this study:

- How can we estimate the subspace needed for spatial filtering techniques (and other applications in general) when the noise affecting the array is not white and is unknown (i.e. array is not calibrated)?
- What are the theoretical limitations of these spatial filtering techniques?
- Can the use of a reference array improve the performance of spatial filtering techniques? If so, what are the needed algorithmic modifications?

The work presented in this thesis on calibration focuses on (direction independent) gain calibration and with more emphasis on the computational complexity. For such calibration we usually assume to know the direction of a set of sources based on previous observations or catalogs. The subproblem that is addressed in this case is:

- How can we exploit the inherent structure in the signal processing model for gain calibration to devise a computationally efficient estimation algorithm?

Even though we are assuming an array of ideal and identical elements during the imaging process, the output of the array is affected by the number of elements in the array, the frequency at which the measurements are performed and the array's topology. Also for the majority of the observations the duration is long enough for the measurements to be affected by the Earth's rotation. Taking these effects into consideration the following subproblems/tasks need to be addressed when the imaging problem is studied:

- Starting from the measurement equation, formulate a solid signal processing approach to image formation which is scalable and sufficiently constrained for an underdetermined problem.
- How is this method related to classical CLEAN-based approaches?
- Provide a method to implement the new algorithm on a system with limited resources.

1.3. APPROACH

This section provides a brief overview of the approaches used throughout this thesis to tackle the problems discussed in the previous section.

1.3.1. SUBSPACE ESTIMATION

In order to address the limitation of EVD for subspace estimation for uncalibrated arrays we propose the use of a technique called factor analysis (FA).

FA is a multivariate technique which decomposes a covariance matrix into a low-rank positive semidefinite and a diagonal matrix. The low-rank part contains the signal subspace and the diagonal part is the noise contribution. In many practical scenarios, the diagonal matrix which corresponds to the noise power of each receiver does not change as rapidly between several measurements as the spatial signature of the RFI (the

low-rank part of the decomposition). In these cases a more accurate estimate can be found by jointly decomposing several covariance measurements which share the same diagonal structure. In this thesis possible extensions of FA which allow for estimation of several subspaces jointly are also proposed.

1.3.2. KRONECKER STRUCTURES AND KRYLOV SOLVERS

As will be shown in Chapter 3, the signal processing models used in this thesis are covariance models and hence the problems we encounter can be categorized as covariance matching problems. In many covariance matching problems we need to find the Jacobian of the covariance matrix with respect to a set of unknown parameters which involves reshaping the covariance matrices into vectors. This leads to strong Kronecker or Khatri–Rao structures in the Jacobian matrices (see Chapter 2 for more details).

Finding the unknown parameters then involves finding the (pseudo) inverse of these Jacobian matrices which if performed directly without taking the Kronecker structure into account is prohibitively expensive for large problems (the dimension of the Jacobian increases with both the number of receivers and the number of unknown parameters). However, using the Kronecker and Khatri–Rao structure we can show that performing a matrix vector multiplication involving these Jacobians can be achieved efficiently in both storage and complexity.

The class of Krylov subspace solvers is capable of solving a set of linear equations iteratively by performing a single matrix vector operation (two if the matrix is not symmetric) per iteration. Combining the Kronecker structure to achieve these multiplications efficiently with Krylov subspace methods improves the scalability of various estimation techniques considerably.

In this thesis this approach is worked out for all three major problems we discuss: subspace estimation (Sec. 4.6.2), gain calibration (Sec. 6.3.2) and imaging (Sec. 7.5.1).

1.3.3. CONSTRAINED LEAST SQUARES

Starting from a parametric array processing model it can be shown that in addition to non-negativity as a lower bound on the image we can use beamforming techniques to find an upper bound for the intensity of each pixel in an image. The LS or weighted LS problems with this set of inequality constraints can be solved using a class of optimization techniques called the active set method. In this thesis we demonstrate that this approach leads to a sequential source removing technique similar to CLEAN. Additionally, a recent analysis of LS problems with non-negativity constraints relates this class of problems to sparse reconstruction techniques based on the ℓ_1 constraint [26, 27]. Hence the approach presented in this work links several deconvolution techniques together and provides a theoretical foundation for further analysis.

1.4. OUTLINE

In outline, the topics discussed in each chapter are as follows.

Chapter 2: A brief overview of the known results for Kronecker products, covariance matching techniques and Krylov–subspace based solvers is provided. These techniques and results are used extensively throughout this thesis.

Chapter 3: The covariance model for subspace estimation, calibration and imaging is provided in this chapter. The results are used for the development of the corresponding algorithms in the chapters that follow.

Chapter 4: The relation between subspace estimation and Factor Analysis (FA) is established in this chapter. Also the extension of some known results from real-valued FA to complex-valued data is discussed. New data models denoted as Extended Factor Analysis (EFA) and Joint EFA (JEFA) are introduced and corresponding estimation algorithms are presented. The Cramér–Rao bound for FA, EFA and JEFA is derived and the performance of the proposed algorithms is evaluated against it using simulations. Finally some new results considering the identifiability of the model are presented.

Chapter 5: An overview of spatial filtering techniques for RFI mitigation based on subspace projection is given. The use of a reference antenna to overcome some of the issues with existing spatial filtering technique is also discussed. Using FA for subspace estimation and (J)EFA as alternative to projection methods is proposed. Using simulations it is shown that the newly proposed methods can produce RFI-free covariance matrices from an uncalibrated array. The performance of FA combined with projection based filtering and EFA is demonstrated using data from Westerbork and LOFAR radio-telescopes.

Chapter 6: A new Krylov-subspace based algorithm is developed for gain calibration. The algorithm is tested using data from the LOFAR telescope. Extension of the algorithm to include direction of arrival corrections is also discussed.

Chapter 7: In this chapter it is shown that in addition to non-negativity, the magnitude of each pixel in an image is also bounded from above. Indeed, the classical “dirty image” is an upper bound, but a much tighter upper bound can be formed from the data using array processing techniques. This formulates image formation as a least squares optimization problem with inequality constraints. Solving this constrained least squares problem using active set techniques is proposed, and the steps needed to implement it are described. It is shown that the least squares part of the problem can be efficiently implemented with Krylov-subspace-based techniques. A method for correcting for the possible mismatch between source positions and the pixel grid is also presented. This correction improves both the detection of sources and their estimated intensities. The performance of these algorithms is evaluated using simulations.

Chapter 8: This section provides a summary of the main results and provides a roadmap for future works.

1.5. CONTRIBUTIONS TO LITERATURE

Some of the results presented in this thesis have been accepted for publication in a journal and/or presented to the community during conferences and workshops, as follows.

JOURNAL PAPERS

- A. Mouri Sardarabadi, A.-J. van der Veen, and A.-J. Boonstra, “Spatial Filtering of RF Interference in Radio Astronomy Using a Reference Antenna Array,” *Signal Processing, IEEE Transactions on*, vol. 64, pp. 432–447, Jan 2016
- A. Mouri Sardarabadi, A. Leshem, and A.-J. van der Veen, “Radio Astronomical Image Formation using Constrained Least Squares and Krylov Subspaces,” *Astron-*

omy and Astrophysics, 2015, accepted on 27-10-2015

1

CONFERENCES AND WORKSHOPS

- A. Mouri Sardarabadi and A.-J. van der Veen, “Complex Factor Analysis,” in *Calibration and Imaging (CALIM) Workshop*, (Manchester, UK), July 2011
- A. Mouri Sardarabadi and A.-J. van der Veen, “Subspace estimation using factor analysis,” in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th*, pp. 477–480, June 2012
- A. Mouri Sardarabadi and A.-J. van der Veen, “Constrained imaging for radio astronomy,” in *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013 IEEE 5th International Workshop on*, pp. 344–347, Dec 2013
- A. Mouri Sardarabadi and A.-J. van der Veen, “Application of Extended Factor Analysis for RFI mitigation,” in *International Biomedical and Astronomical Signal Processing (BASP) Frontiers*, (Villars-sur-Ollon, Switzerland), Jan. 2013
- A. Mouri Sardarabadi and A.-J. van der Veen, “Application of Krylov based methods in calibration for radio astronomy,” in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2014 IEEE 8th*, pp. 153–156, June 2014
- A. Mouri Sardarabadi and A.-J. van der Veen, “Joint Extended Factor Analysis,” in *WSA 2015; 19th International ITG Workshop on Smart Antennas; Proceedings of*, pp. 1–8, March 2015
- A. Mouri Sardarabadi, A. Leshem, and A.-J. van der Veen, “Computationally efficient radio astronomical image formation using constrained least squares and the MVDR beamformer,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 5664–5668, April 2015

2

PRELIMINARIES

In this chapter we will review some of the well known and widely used estimation and linear algebraic methods that are extensively used in the following chapters. We will start in Sec. 2.1 by introducing the notation used in this thesis followed by the definition of the Kronecker product and some of its properties in Sec. 2.2. We then discuss maximum-likelihood (ML) and various covariance matching techniques in Sec. 2.3, where we will follow the material covered in [37] closely and then proceed with a short overview of Krylov subspace based techniques for solving linear systems in Sec. 2.4.

2.1. NOTATIONS

x, X	italic lower or upper case characters represent scalars
\mathbf{x}	bold lower case characters represent column vectors
\mathbf{X}	bold upper case characters represent matrices
$x_{ij} = [\mathbf{X}]_{ij}$	corresponds to the entry of \mathbf{X} at the i th row and j th column
$\mathbf{x}_i = [\mathbf{X}]_i$	is the i th column of \mathbf{X}
a_i	is the i th element of the vector \mathbf{a}
\mathbf{I}_P	identity matrix of size $P \times P$ (\mathbf{I} is an identity matrix of appropriate size)
\mathbf{e}_i	a unit vector that has entry i equal to 1
$(\cdot)^T$	transpose operator
$(\cdot)^H$	Hermitian transpose
$(\cdot)^\dagger$	pseudo-inverse
$(\cdot)^*$	complex conjugate
$\text{vect}(\cdot)$	reshapes the argument matrix to a vector by stacking its columns
$\text{unvect}(\cdot)$	inverse of $\text{vect}(\cdot)$ for a matrix with known dimensions
$\text{vectdiag}(\cdot)$	a vector obtained by stacking the diagonal elements of the argument matrix
$\text{diag}(\mathbf{x})$	a diagonal matrix with elements of \mathbf{x} on the main diagonal
$\text{diag}(\mathbf{X})$	is the same as $\text{diag}(\text{vectdiag}(\mathbf{X}))$
$\text{bdiag}(\mathbf{X}_i)$	a block-diagonal matrix from the set of argument matrices
\otimes	Kronecker product
\odot	Hadamard product
\circ	Khatri–Rao product
$\text{tr}(\cdot)$	trace of the argument matrix
\mathbf{K}^{PQ}	a permutation matrix such that $\mathbf{K}^{PQ} \text{vect}(\mathbf{X}) = \text{vect}(\mathbf{X}^T)$ where \mathbf{X} is a $P \times Q$ matrix
$\ \cdot\ _F$	Frobenius norm of a matrix
$\ \cdot\ $ or $\ \cdot\ _2$	Two norm of a vector
$ x $	absolute value of x
$ \mathbf{X} $	determinant of \mathbf{X}
$\mathcal{E}\{\cdot\}$	expectation operator

Additionally, $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ represents the multivariate complex normal distribution with expected value $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. A calligraphic capital letter such as \mathcal{X} represents a set of indices, and $\mathbf{a}_{\mathcal{X}}$ is a column vector constructed by stacking the elements of \mathbf{a} that belong to \mathcal{X} . The corresponding indices are stored with the vector as well (similar to the storage of matlab “sparse” vectors).

2.2. PROPERTIES OF THE KRONECKER PRODUCT

Let \otimes denote the Kronecker product, i.e.,

$$\mathbf{A} \otimes \mathbf{B} := \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \cdots \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

Further, \circ denotes the Khatri-Rao product (column-wise Kronecker product), i.e.,

$$\mathbf{A} \circ \mathbf{B} := [\mathbf{a}_1 \otimes \mathbf{b}_1, \mathbf{a}_2 \otimes \mathbf{b}_2, \dots]$$

and \odot denotes the Schur-Hadamard (element-wise) product.

The following properties hold for matrices and vectors with compatible dimensions:

$$(\mathbf{B}^T \otimes \mathbf{A}) \text{vect}(\mathbf{X}) = \text{vect}(\mathbf{A}\mathbf{X}\mathbf{B}) \quad (2.1)$$

$$(\mathbf{B} \otimes \mathbf{A})^H = (\mathbf{B}^H \otimes \mathbf{A}^H) \quad (2.2)$$

$$(\mathbf{B} \otimes \mathbf{A})^{-1} = (\mathbf{B}^{-1} \otimes \mathbf{A}^{-1}) \quad (2.3)$$

$$(\mathbf{B}^T \circ \mathbf{A})\mathbf{x} = \text{vect}(\mathbf{A} \text{diag}(\mathbf{x})\mathbf{B}) \quad (2.4)$$

$$(\mathbf{B}\mathbf{C} \otimes \mathbf{A}\mathbf{D}) = (\mathbf{B} \otimes \mathbf{A})(\mathbf{C} \otimes \mathbf{D}) \quad (2.5)$$

$$(\mathbf{B}\mathbf{C} \circ \mathbf{A}\mathbf{D}) = (\mathbf{B} \otimes \mathbf{A})(\mathbf{C} \odot \mathbf{D}) \quad (2.6)$$

$$(\mathbf{B}^H \mathbf{C} \odot \mathbf{A}^H \mathbf{D}) = (\mathbf{B} \circ \mathbf{A})^H (\mathbf{C} \odot \mathbf{D}) \quad (2.7)$$

$$\text{vectdiag}(\mathbf{A}^H \mathbf{X} \mathbf{A}) = (\mathbf{A}^* \circ \mathbf{A})^H \text{vect}(\mathbf{X}) \quad (2.8)$$

Additionally for any $P \times Q$ matrix \mathbf{A} there exists a $PQ \times PQ$ permutation matrix \mathbf{K}^{PQ} such that

$$\text{vect}(\mathbf{A}^T) = \mathbf{K}^{PQ} \text{vect}(\mathbf{A}). \quad (2.9)$$

We also have $\mathbf{K}^T \mathbf{K} = \mathbf{I}$, $\mathbf{K}^{P,Q} = (\mathbf{K}^{Q,P})^T$ and hence $\mathbf{K}^{P,Q} \mathbf{K}^{Q,P} = \mathbf{I}$. Using these relations, for any $P \times Q$ matrix \mathbf{A} and $M \times N$ matrix \mathbf{B} we have

$$(\mathbf{A} \otimes \mathbf{B}) \mathbf{K}^{Q,N} = \mathbf{K}^{P,M} (\mathbf{B} \otimes \mathbf{A}) \quad (2.10)$$

$$(\mathbf{A} \circ \mathbf{B}) = \mathbf{K}^{P,M} (\mathbf{B} \circ \mathbf{A}), \quad (2.11)$$

where $Q = N$ for (2.11).

With appropriate use of the identity matrix \mathbf{I} we can use these properties even if we are working with a single matrix or a product of two matrices. For example for a diagonal matrix $\mathbf{D} = \text{diag}(\mathbf{d})$, we have $\text{vect}(\mathbf{D}) = \text{vect}(\mathbf{I} \text{diag}(\mathbf{d}) \mathbf{I}) = (\mathbf{I} \circ \mathbf{I})\mathbf{d}$.

2.3. COVARIANCE MATCHING TECHNIQUES

Estimating physical parameters from noisy data is one of the main activities in signal processing. In this section we will establish the relation between parameter estimation and covariance matching techniques and discuss how various optimization techniques could be used when this problem becomes nonlinear.

2.3.1. STATISTICAL SUFFICIENCY

Let us define \mathbf{y} as a $P \times 1$ vector that represents the output of P receivers and \mathbf{Y} as a $P \times N$ set of independent measurements (or a dataset) from \mathbf{y} where N is the number of measurements. Let $\boldsymbol{\theta}$ be a $n \times 1$ vector of unknown parameters that influence \mathbf{y} . We are interested in finding and estimate for $\boldsymbol{\theta}$ based on the measurements in \mathbf{Y} . Of course we can try to find an estimate for $\boldsymbol{\theta}$ by analyzing \mathbf{Y} directly, however it is also possible to estimate the unknown parameters from a function of the measured data, $t(\mathbf{Y})$, if $t(\cdot)$ is a sufficient statistic [38, pp. 102-105]. Saying $t(\cdot)$ is a sufficient statistic means that we can deduce the same information about $\boldsymbol{\theta}$ from $t(\mathbf{Y})$ with the same accuracy as using the entire dataset. One of the main reasons to use such a statistic is to reduce the size of the measurements without losing any information which will be relevant when estimating the unknown parameters. In this case $t(\cdot)$ can be seen as a loss-less compression.

From this point forward we assume that the samples in \mathbf{Y} have a complex Gaussian distribution. It can be shown that for a zero-mean Gaussian distributed dataset, the sample covariance matrix

$$\hat{\mathbf{R}} = \frac{1}{N} \mathbf{Y} \mathbf{Y}^H,$$

is a sufficient statistic [38, pp. 102-105]. Using this statistics we can reduce the storage from $2PN$ real parameters to P^2 real parameters needed to store the covariance matrix (of course this is a reduction only when $N > P$).

We still need to define a relation between $\boldsymbol{\theta}$ and $\hat{\mathbf{R}}$. We do this by putting a model on the "true" covariance matrix, which is defined as

$$\mathbf{R}(\boldsymbol{\theta}) = \mathcal{E}\{\hat{\mathbf{R}}\} = \mathcal{E}\{\mathbf{y} \mathbf{y}^H\}$$

The relation between \mathbf{R} and $\boldsymbol{\theta}$ is assumed to be known based on the underlying model for \mathbf{y} . For example if we assume that all the elements in \mathbf{y} are independent, \mathbf{R} will become a diagonal matrix and we can model it as $\mathbf{R} = \text{diag}(\boldsymbol{\theta})$ where $\boldsymbol{\theta}$ is a $P \times 1$ vector representing the variances for each element in \mathbf{y} . In this section we will not assume any specific model for $\mathbf{R}(\boldsymbol{\theta})$, however we will assume that such a model exists. In the Chapter 3 we will discuss the data model in more details.

Now that we have established the feasibility of $\hat{\mathbf{R}}$ as a replacement for the original set of measurements \mathbf{Y} , we will discuss various estimation techniques and algorithms to find an estimate for the unknown parameters $\boldsymbol{\theta}$ from $\hat{\mathbf{R}}$.

2.3.2. MAXIMUM LIKELIHOOD

Considering that we have assumed Gaussianity of the data and thus know the likelihood function, we can try to find the ML estimate by means of solving an optimization problem. The aim is to find $\boldsymbol{\theta}$ that maximizes the complex log-likelihood function

$$l(\mathbf{x}; \boldsymbol{\theta}) = N \left[-\log(\pi^P) + \log |\mathbf{R}(\boldsymbol{\theta})|^{-1} - \text{tr}(\mathbf{R}(\boldsymbol{\theta})^{-1} \hat{\mathbf{R}}) \right]. \quad (2.12)$$

One way to achieve this by finding the Fisher score (i.e. the gradient of the log-likelihood function) and setting it equal to zero. The Fisher score for a proper Gaussian distributed

signal is given by [39, p.165]

$$\begin{aligned}
 g(\theta_j) &= \frac{\partial}{\partial \theta_j^*} l(\mathbf{x}; \boldsymbol{\theta}) \\
 &= -N \text{tr} \left[\mathbf{R}^{-1} \left(\frac{\partial \mathbf{R}}{\partial \theta_j} \right)^H \right] \\
 &\quad + N \text{tr} \left[\mathbf{R}^{-1} \left(\frac{\partial \mathbf{R}}{\partial \theta_j} \right)^H \mathbf{R}^{-1} \hat{\mathbf{R}} \right],
 \end{aligned} \tag{2.13}$$

where the partial derivatives are Wirtinger derivatives (i.e. the variable θ and its conjugate are treated as independent variables when derivatives are derived).

By rearranging the results in a vector form we find the following relation for the Fisher score:

$$\mathbf{g}(\boldsymbol{\theta}) = \frac{\partial l(\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \tag{2.14}$$

$$= N \mathbf{J}(\boldsymbol{\theta})^H (\mathbf{R}(\boldsymbol{\theta})^{-T} \otimes \mathbf{R}(\boldsymbol{\theta})^{-1}) \text{vect}(\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})) \tag{2.15}$$

where

$$\mathbf{J}(\boldsymbol{\theta}) = \frac{\partial \text{vect}(\mathbf{R})}{\partial \boldsymbol{\theta}^T} \tag{2.16}$$

To illustrate we revisit our example in the previous section where $\mathbf{R} = \text{diag}(\boldsymbol{\theta})$. In this case $\mathbf{J} = \mathbf{I}_P \circ \mathbf{I}_P$ and we have

$$\begin{aligned}
 \mathbf{0} &= (\mathbf{I}_P \circ \mathbf{I}_P)^T (\text{diag}(\boldsymbol{\theta})^{-1} \otimes \text{diag}(\boldsymbol{\theta})^{-1}) [\text{vect}(\hat{\mathbf{R}}) - (\mathbf{I}_P \circ \mathbf{I}_P) \boldsymbol{\theta}] \\
 &= \text{diag}(\boldsymbol{\theta})^{-2} \text{vect} \text{diag}(\hat{\mathbf{R}}) - \text{diag}(\boldsymbol{\theta})^{-2} \boldsymbol{\theta}
 \end{aligned}$$

and the ML estimate becomes $\hat{\boldsymbol{\theta}} = \text{vect} \text{diag}(\hat{\mathbf{R}})$.

Except for a few simple parameterizations, setting (2.14) equal to zero will lead to a nonlinear equation which has no closed-form solution. In order to find the ML solution in these cases we need to utilize iterative optimization techniques. There are many approaches to ML estimation. However in this section we would like to keep a level of generality and hence limit the discussion of ML estimation to the method of scoring.

2.3.3. METHOD OF SCORING

The method of scoring is a special case of the Newton algorithm where the Hessian of the log-likelihood is replaced by its expected value, assuming the current solution for $\boldsymbol{\theta}$ is the optimal solution. This is exactly how the Fisher Information Matrix (FIM) is defined and hence this procedure is equivalent to replace the Hessian by the FIM [37]. For a complex Gaussian distributed random variable the FIM is given by [40]

$$\mathbf{F}(\boldsymbol{\theta}) = \mathbf{J}(\boldsymbol{\theta})^H (\mathbf{R}(\boldsymbol{\theta})^{-T} \otimes \mathbf{R}(\boldsymbol{\theta})^{-1}) \mathbf{J}(\boldsymbol{\theta}). \tag{2.17}$$

Starting with an initial guess $\hat{\boldsymbol{\theta}}_{(0)}$ for $\boldsymbol{\theta}$, the iterations for scoring method are

$$\hat{\boldsymbol{\theta}}_{(k+1)} = \hat{\boldsymbol{\theta}}_{(k)} + \mu_{(k)} \boldsymbol{\Delta}_{(k)}$$

where $\mu_{(k)}$ is the step-size and Δ is the direction of descent and the solution to the the following system of linear equations

$$\mathbf{F}(\boldsymbol{\theta})\Delta = \mathbf{g}(\boldsymbol{\theta}), \quad (2.18)$$

where $\mathbf{g}(\boldsymbol{\theta})$ is the Fisher score given by (2.14). The iterations are terminated when $\|\mathbf{g}(\hat{\boldsymbol{\theta}})\| \leq \epsilon$ where $\epsilon > 0$ is a small number based on the desired accuracy or if the maximum number of iterations is reached.

Replacing the Hessian with its expected value provides some stability with respect to the Newton method, however the convergence is still sensitive to the initial guess. If the problem is well posed and there are no identification problems then $\Delta = \mathbf{F}(\boldsymbol{\theta})^{-1}\mathbf{g}(\boldsymbol{\theta})$ and \mathbf{F}^{-1} is the Cramér-Rao bound on the variance of the estimates and also the asymptotic covariance for the solution $\hat{\boldsymbol{\theta}}$. In this case the algorithm provides us with the performance statistics for "free" as a by-product. However, inverting both $\mathbf{R}(\boldsymbol{\theta})$ and the FIM at each iteration could be expensive depending on the size of the problem. If N is large enough we could approximate \mathbf{R}^{-1} with $\hat{\mathbf{R}}^{-1}$ in the above iterations. In the next section we will show that by doing so we are solving a special case of the Weighted Least Squares (WLS) problem by a Gauss-Newton iterative procedure.

2.3.4. WEIGHTED LEAST SQUARES

The WLS procedure can be regarded as a true covariance matching technique, because it does not need an underlying statistic like the Gaussianity of the signals. The idea is to match a model covariance matrix to the given sample covariance matrix by minimizing the distance between them with respect to some weighted inner product norm. In general, when we use WLS, we want to minimize a cost function of the form

$$f(\boldsymbol{\theta}) = \mathbf{r}'(\boldsymbol{\theta})^H \mathbf{W} \mathbf{r}'(\boldsymbol{\theta}) \quad (2.19)$$

where $\mathbf{r}'(\boldsymbol{\theta}) = \text{vect}(\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta}))$ is the mismatch between the model and the data, and \mathbf{W} is a positive weighting matrix.

Taking the derivatives with respect to the unknown parameters we see that solving this problem is the same as solving

$$\mathbf{J}^H \mathbf{W} \text{vect}(\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})) = \mathbf{0}. \quad (2.20)$$

where \mathbf{J} is the Jacobian given in (2.16). When the problem is linear then \mathbf{J} is not a function of $\boldsymbol{\theta}$, we can write $\text{vect}(\mathbf{R}) = \mathbf{J}\boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}}$ can be found directly by solving

$$\mathbf{J}^H \mathbf{W} \mathbf{J} \hat{\boldsymbol{\theta}} = \mathbf{J}^H \mathbf{W} \text{vect}(\hat{\mathbf{R}}).$$

If $\mathbf{J}^H \mathbf{W} \mathbf{J}$ is non-singular there exists a unique $\hat{\boldsymbol{\theta}} = (\mathbf{J}^H \mathbf{W} \mathbf{J})^{-1} \mathbf{J}^H \mathbf{W} \text{vect}(\hat{\mathbf{R}})$. However if the problem is non-linear we need an iterative approach to find the solution. In this section we will discuss the Gauss-Newton algorithm [41, pp.134-136].

The iterations for the WLS using Gauss-Newton are

$$\hat{\boldsymbol{\theta}}_{(k+1)} = \hat{\boldsymbol{\theta}}_{(k)} + \mu_{(k)} \Delta_{(k)}$$

where Δ is the solution to the the following system of linear equations

$$\mathbf{J}(\boldsymbol{\theta})^H \mathbf{W} \mathbf{J}(\boldsymbol{\theta}) \Delta = \mathbf{J}(\boldsymbol{\theta})^H \mathbf{W} \text{vect}(\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})), \quad (2.21)$$

and it is called the direction of descent.

We observe that (2.20) and (2.14) are very similar despite having different cost functions. Setting $\mathbf{W} = (\hat{\mathbf{R}}^{-T} \otimes \hat{\mathbf{R}}^{-1})$ as the weight for solving the WLS problem is asymptotically optimal (i.e. converges to ML solution) [37] and is computationally more attractive. Another interesting special case occurs when $\mathbf{W} = \mathbf{I}$ which reduces the problem to non-linear LS.

2.3.5. REDUCED-GRADIENT METHOD

Up to this point we have tried to solve the WLS and ML problems without paying any attention to the identifiability of the problems. Identifiability is strongly model dependent and hence will be addressed when a specific problem is studied. However there are some general remarks and algorithmic procedures that can be discussed without losing generality. For example a necessary condition for identifiability is having more known parameters (e.g. the number of entries in the sample covariance matrix $\hat{\mathbf{R}}$ which is equal to P^2) than unknown parameters which is equal to number of elements in $\boldsymbol{\theta}$, n . We define the degrees of freedom for covariance matching techniques as

$$s = P^2 - n$$

and our first identifiability criteria is that $s \geq 0$. If this is not the case the problem could be remedied by putting L equality constraints on the parameters $\boldsymbol{\theta}$. Let each constraint be formulated as $h_i(\boldsymbol{\theta}) = 0$ for $i = 1, \dots, L$, or in the vector form

$$\mathbf{h}(\boldsymbol{\theta}) = \mathbf{0}.$$

In this case we can redefine the number of degrees of freedom as

$$s = P^2 - n + L,$$

and L should be large enough to guarantee $s \geq 0$. Note that this is a necessary and not a sufficient condition. Having a positive number of degrees of freedom does not guarantee identifiability. However in the problems where we define $\mathbf{h}(\boldsymbol{\theta})$, we assume that this set of constraints can guarantee identification.

If constraints are linear then

$$\mathbf{h}(\boldsymbol{\theta}) = \mathbf{H}\boldsymbol{\theta} = \mathbf{0}. \quad (2.22)$$

It is then clear that any $\boldsymbol{\theta}$ that satisfies these constraints is in the null space of \mathbf{H} . This means that starting from a feasible initial point $\boldsymbol{\theta}_{(0)}$ that satisfies these constants, the search direction toward the solution should be orthogonal to \mathbf{H} . Let \mathbf{U} be a $n \times (n - L)$ matrix forming an orthogonal basis for the null space of \mathbf{H} such that $\mathbf{H}\mathbf{U} = \mathbf{0}$, then in general we can write for the WLS and the scoring updates

$$\Delta = \mathbf{U}(\mathbf{U}^H \mathbf{J}^H \mathbf{W} \mathbf{J} \mathbf{U})^{-1} \mathbf{U}^H \mathbf{J}^H \mathbf{W} \text{vect}(\hat{\mathbf{R}} - \mathbf{R}_{(k)}), \quad (2.23)$$

where $\mathbf{W} = \mathbf{R}(\boldsymbol{\theta})^{-T} \otimes \mathbf{R}(\boldsymbol{\theta})^{-1}$ for scoring method.

However, in general \mathbf{h} is not linear. In these cases we can use a Taylor-series expansion of the constraint function at each iteration such that

$$\mathbf{h}(\boldsymbol{\theta}_{(k)} + \Delta) \approx \mathbf{h}(\boldsymbol{\theta}_{(k)}) + \mathbf{H}(\boldsymbol{\theta}_{(k)})\Delta \quad (2.24)$$

with this approximation we again see that if $\boldsymbol{\theta}_{(k)}$ satisfies the constraints, we can keep the solution feasible by choosing Δ to be orthogonal to \mathbf{H} . If we update \mathbf{U} in (2.23) at each iteration to be an orthogonal basis for the null space of $\mathbf{H}(\boldsymbol{\theta}_{(k)})$, we can use (2.23) again to find the direction of descent as long as the constraint function does not become too non-linear. This approximation is known as the Reduced-Gradient method [41, pp.220-223]. If the non-linearity of the constraint function is too large this algorithm fails to converge. The convergence of this method as with all Newton based methods, depends strongly on the initial guess.

One of the common factors in all the iterative approaches we have discussed so far is the need to solve a linear system of equations. In the context of radio-astronomy the size of the problem we need to solve can become very large and both storage and computational complexity play an important role in choosing the right method for solving these large linear systems. Another factor which will become more apparent when we discuss specific problems such as subspace estimation, RFI migration, calibration and imaging is the inherent Kronecker or Khatri-Rao structure of the Jacobians which is a consequence of the `vect(.)` operator. These properties require us to choose a method which is efficient in storage and can take advantage of the structure in the Jacobians. Krylov subspace based solvers such as Conjugate Gradient, LSQR and MinresQLP form a family of iterative algorithms which are suited to this end and we will discuss them in more detail in the following section.

2.4. KRYLOV SUBSPACE BASED SOLVERS

In this section we will focus on solving a (large) system of linear equations of the form $\mathbf{M}\mathbf{x} = \mathbf{b}$ using Krylov subspace based methods. We will show that we can find a solution by repeated application of matrix vector products of the form $\mathbf{M}\mathbf{v}$ and/or $\mathbf{M}^H\mathbf{u}$ and a few scalar updates. We will start with a general motivation for the method and then focus on the base idea behind the LSQR algorithm in Sec. 2.4.2.

2.4.1. MINIMUM POLYNOMIAL AND KRYLOV SUBSPACES

To illustrate the idea behind Krylov subspace based methods we will assume that \mathbf{M} is a square and non-singular matrix. In this case there exists a unique solution for \mathbf{x} which is given by $\mathbf{x} = \mathbf{M}^{-1}\mathbf{b}$. Using the minimum polynomial of a matrix we can write

$$\mathbf{M}^{-1} = \frac{1}{\gamma_0} \sum_{j=0}^{m-1} \gamma_{j+1} \mathbf{M}^j,$$

where for a diagonalizable matrix \mathbf{M} , m is the number of distinct eigenvalues [42]. Using this polynomial expansion we have for our solution

$$\begin{aligned}\mathbf{x} &= \frac{1}{\gamma_0} \sum_{j=0}^{m-1} \gamma_{j+1} \mathbf{M}^j \mathbf{b} \\ &= [\mathbf{b} \quad \mathbf{M}\mathbf{b} \quad \dots \quad \mathbf{M}^{m-1}\mathbf{b}] \boldsymbol{\gamma}\end{aligned}$$

where

$$\boldsymbol{\gamma} = \frac{1}{\gamma_0} [\gamma_1, \dots, \gamma_m]^T,$$

and $\mathcal{K}_m(\mathbf{M}, \mathbf{b}) = [\mathbf{b}, \mathbf{M}\mathbf{b}, \dots, \mathbf{M}^{m-1}\mathbf{b}]$ is called the Krylov subspace of \mathbf{M} and \mathbf{b} . Krylov subspace based methods compute $\mathcal{K}_n(\mathbf{M}, \mathbf{b})$ iteratively, for $n = 1, 2, \dots$ and find an approximate for \mathbf{x} by means of a projection on this subspace. Updating the subspace involves only a matrix-vector multiplication of the form $\mathbf{M}\mathbf{v}$.

In cases where \mathbf{M} is singular or if it is not a square matrix, another class of Krylov based algorithms such as LSQR can be used which is related to bidiagonalization of the matrix \mathbf{M} . The next section describes the idea behind LSQR and how this helps towards a more scalable implementation of a linear solver.

2.4.2. LANCZOS ALGORITHM AND LSQR

If we are solving a problem of the form $\|\mathbf{b} - \mathbf{M}\mathbf{x}\|_2^2$, we can find a solution by first computing the singular value decomposition (SVD) of \mathbf{M} as

$$\mathbf{M} = \mathbf{U}\mathbf{S}\mathbf{V}^H, \quad (2.25)$$

where \mathbf{U} and \mathbf{V} are unitary matrices and \mathbf{S} is a diagonal matrix with positive singular values. Then the solution \mathbf{x} to $\min \|\mathbf{b} - \mathbf{M}\mathbf{x}\|^2$ is found by solving for \mathbf{y} in

$$\mathbf{S}\mathbf{y} = \mathbf{U}^H \mathbf{b} \quad (2.26)$$

followed by setting

$$\mathbf{x} = \mathbf{V}\mathbf{y}. \quad (2.27)$$

Solving the LS problem with this method is expensive in both number of operations and memory usage, especially if the matrices \mathbf{U} and \mathbf{V} are not needed after finding the solution. As we will see shortly, looking at another matrix decomposition helps us to reduce these costs. For the rest of this section we use the notation given by [43].

The first step in this approach for solving the LS problem is to reduce \mathbf{M} to a lower bidiagonal form as

$$\mathbf{M} = \mathbf{U}\mathbf{B}\mathbf{V}^H, \quad (2.28)$$

where \mathbf{B} is a bidiagonal matrix of the form

$$\mathbf{B} = \left[\begin{array}{cccc|c} \alpha_1 & & & & \\ \beta_2 & \alpha_2 & & & \\ & \ddots & \ddots & & \\ & & \beta_r & \alpha_r & \\ \hline & & & & \mathbf{0} \end{array} \right], \quad (2.29)$$

where $r = \text{rank}(\mathbf{M}) = \text{rank}(\mathbf{B})$ and \mathbf{U}, \mathbf{V} are unitary matrices (different than in (2.25)). This representation is not unique and without loss of generality we could choose \mathbf{U} to satisfy

$$\mathbf{U}^H \mathbf{b} = \beta_1 \mathbf{e}_1 \quad (2.30)$$

where $\beta_1 = \|\mathbf{b}\|_2$ and \mathbf{e}_1 is a unit norm vector with its first element equal to one.

Using \mathbf{B} , forward substitution gives the LS solution efficiently by solving for \mathbf{y} in

$$\mathbf{B}\mathbf{y} = \mathbf{U}^H \mathbf{b} = \beta_1 \mathbf{e}_1 \quad (2.31)$$

followed by

$$\mathbf{x} = \mathbf{V}\mathbf{y}.$$

Using forward substitution we have

$$y_1 = \frac{\beta_1}{\alpha_1} \quad (2.32)$$

$$\mathbf{x}_1 = \mathbf{v}_1 y_1, \quad (2.33)$$

followed by the recursion,

$$y_{n+1} = -\frac{\beta_{n+1}}{\alpha_{n+1}} y_n \quad (2.34)$$

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \mathbf{v}_{n+1} y_{n+1} \quad (2.35)$$

for $n = 1, \dots, M$ where $M < r$ is the iteration at which $\|\mathbf{M}^H(\mathbf{M}\mathbf{x}_n - \mathbf{b})\|^2$ vanishes within the desired precision. We can combine the bidiagonalization and solving for \mathbf{x} and avoid extra storage needed for saving \mathbf{B} , \mathbf{U} and \mathbf{V} . One such algorithm is based on a Krylov subspace method called the Lanczos algorithm [44]. We first initialize with

$$\beta_1 = \|\mathbf{b}\|_2 \quad (2.36)$$

$$\mathbf{u}_1 = \frac{\mathbf{b}}{\beta_1} \quad (2.37)$$

$$\alpha_1 = \|\mathbf{M}^H \mathbf{u}_1\|_2 \quad (2.38)$$

$$\mathbf{v}_1 = \frac{\mathbf{M}^H \mathbf{u}_1}{\alpha_1}. \quad (2.39)$$

The iterations are then given by

$$\begin{aligned} \beta_{n+1} &= \|\mathbf{M}\mathbf{v}_n - \alpha_n \mathbf{u}_n\|_2 \\ \mathbf{u}_{n+1} &= \frac{1}{\beta_{n+1}} (\mathbf{M}\mathbf{v}_n - \alpha_n \mathbf{u}_n) \\ \alpha_{n+1} &= \|\mathbf{M}^H \mathbf{u}_{n+1} - \beta_{n+1} \mathbf{v}_n\|_2 \\ \mathbf{v}_{n+1} &= \frac{1}{\alpha_{n+1}} (\mathbf{M}^H \mathbf{u}_{n+1} - \beta_{n+1} \mathbf{v}_n) \end{aligned} \quad (2.40)$$

for $n = 1, 2, \dots, M$, where $\mathbf{u}_n^H \mathbf{u}_n = \mathbf{v}_n^H \mathbf{v}_n = 1$. This provides us with all the parameters needed to solve the problem.

However, because of finite precision errors, the columns of \mathbf{U} and \mathbf{V} found in this way lose their orthogonality as we proceed. In order to prevent this error propagation into the final solution \mathbf{x} , different algorithms like Conjugate Gradient (CG), MINRES, LSQR, etc. have been proposed. The exact updates for \mathbf{x}_n and stopping criteria to find M depend on the choice of algorithm and are therefore not included in the above iterations.

We will use LSQR mainly in Chapter 7 to solve large image reconstruction problems. In other parts of this thesis we will mainly use MinresQLP which is related to tridiagonalization of the matrix \mathbf{M} and is discussed next.

2.4.3. TRIDIAGONALIZATION AND KRYLOV SUBSPACES

If \mathbf{M} is symmetric we can reduce the number of matrix vector multiplications from two to one by reducing the matrix into tridiagonal form instead of bidiagonal form which was discussed in the previous section. We will follow the discussion provided in [45] closely in this section.

For simplified illustration we assume that \mathbf{M}^{-1} exists. Now let \mathbf{K} be a matrix form of $\mathcal{K}_m(\mathbf{M}, \mathbf{b})$ then we will have

$$\mathbf{MK} = \mathbf{KH}$$

where \mathbf{H} is an upper Hessenberg matrix of the form

$$\mathbf{H} = \left[\begin{array}{ccc|c} 0 & & & \\ 1 & 0 & & \\ & 1 & \ddots & \\ & & \ddots & 0 \\ & & & 1 \end{array} \right] \mathbf{h},$$

and $\mathbf{h} = \mathbf{K}^{-1}\mathbf{M}^m\mathbf{b}$. Now let the QR decomposition of \mathbf{K} be $\mathbf{K} = \mathbf{UZ}$ where \mathbf{U} is a unitary matrix and \mathbf{Z} is an upper triangular matrix. Multiplying from the left with \mathbf{U}^H and from the right with \mathbf{Z}^{-1} we have

$$\mathbf{U}^H\mathbf{MU} = \mathbf{ZHZ}^{-1}.$$

The left-hand-side of the equation is a symmetric matrix. Given the fact that multiplying an upper Hessenberg matrix with an upper triangular matrix must be upper Hessenberg the right-hand-side is an upper Hessenberg matrix and symmetric and hence $\mathbf{T} = \mathbf{ZHZ}^{-1}$ is a tridiagonal matrix. If the system of equations is consistent the tridiagonal matrix \mathbf{T} of \mathbf{M} with $\text{rank}(\mathbf{M}) = r$ has the form

$$\mathbf{T} = \left[\begin{array}{ccc|c} \alpha_1 & \beta_2 & & \\ \beta_2 & \alpha_2 & \beta_3 & \\ & \ddots & \ddots & \beta_r \\ & & \beta_r & \alpha_r \\ \hline & & & \mathbf{0} \end{array} \right]. \quad (2.41)$$

We can modify the Lanczos algorithm from the previous section to find \mathbf{U} and \mathbf{T} in an iterative way. Again without loss of generality we can assume $\mathbf{U}^H\mathbf{b} = \beta_1\mathbf{e}_1$ where $\beta_1 = \|\mathbf{b}\|_2$. Using this relation we can initialize the algorithm with

$$\mathbf{u}_0 = \mathbf{0} \quad (2.42)$$

$$\beta_1 = \|\mathbf{b}\|_2 \quad (2.43)$$

$$\mathbf{u}_1 = \frac{\mathbf{b}}{\beta_1}. \quad (2.44)$$

The iterations for symmetric Lanczos become

$$\begin{aligned} \mathbf{z} &= \mathbf{M}\mathbf{u}_n \\ \alpha_n &= \mathbf{u}_n^H \mathbf{z} \\ \beta_{n+1} &= \|\mathbf{z} - \alpha_n \mathbf{u}_n - \beta_n \mathbf{u}_{n-1}\|_2 \\ \mathbf{u}_{n+1} &= \frac{1}{\beta_{n+1}} (\mathbf{z} - \alpha_n \mathbf{u}_n - \beta_n \mathbf{u}_{n-1}) \end{aligned} \quad (2.45)$$

for $n = 1, 2, \dots, M$, where $M \leq r$.

The solution can then be updated with a single matrix vector product of the form $\mathbf{M}\mathbf{u}$. This procedure forms the basis for the symmetric class of Krylov subspace methods such as MinresQLP.

An overview of Krylov subspace based methods is given by [46, pp.91]. This study shows that LSQR is a good candidate to solve LS problems when we are dealing with an ill-conditioned and non-square matrix and MinresQLP is capable of solving systems that involve symmetric and singular matrices.

When we are solving the direction of descent for a non-linear optimization problem, the matrix \mathbf{M} is a function of the unknown parameters and the matrix vector product $\mathbf{M}\mathbf{u}$ can be performed using the current estimates of these parameters. This makes constructing and storing \mathbf{M} unnecessary and reduces the storage requirements for solving these types problems enormously. In Sec. 4.6.2, 6.3.2 and 7.5.1 we will discuss constructing such a procedure in more detail and we will also show how to take advantage of the Kronecker structure of the Jacobians.

This concludes our discussion on covariance matching and Krylov subspace methods. Both topics are well established and have a rich literature which is beyond the scope of this thesis. The references given in this chapter, especially [37, 45] and [46], provide excellent material for an interested reader.

3

DATA MODEL

In Sec. 2.3 we have seen that for a zero-mean Gaussian distributed data, the sample covariance matrix is a sufficient statistics and that a covariance data model is essential for its analysis. The measured output of the receivers in radio astronomy can also be modeled as zero-mean Gaussian signals and hence the same principles apply. The data model, or measurement equation, defines which (physical) quantities are of interest and describes the relation between them and the collected data. There are various approaches to describing the data model for a set of radio-astronomical measurement. One way is a detailed description of the measurements based on the physical properties of the electromagnetic (EM) fields similar to models presented in [47] and [2, pp. 48-53]. The resulting data model, while general and applicable to any measurement set and instrument, is too complicated from a practical perspective. In this chapter we will motivate the use of signal and array processing models which will enable us to apply the rich set of tools developed for these fields to radio-astronomy.

The method used in this chapter follows the matrix formulation of the data model done in [48] and [2, pp. 57-58]. The resulting data model forms the basis for the algorithms and methods which are discussed in the rest of this thesis.

3.1. ARRAY RESPONSE VECTOR

We consider an instrument where P receivers (stations or antennas) are observing the sky. Fig. 3.1 shows a simple 1-D version of such an array of receivers which we will use to illustrate and motivate some of the assumptions needed for developing our data model. For simplicity we will start with isotropic and identical antennas as receivers.

3.1.1. NARROW-BAND ASSUMPTION

In this section we will discuss the narrow-band assumption with a simplified derivation, using the delay-phase relation of the Fourier transform. A more rigorous derivation can be found in [47].

Considering that EM signals satisfy the wave equation we can assume that the re-




Figure 3.1: Simple 1-D array of receivers observing a source $s(t)$ at an angle ϕ .

ceived signals at different receivers are a delayed version of each other. Let the delay for the p th receiver be defined as τ_p . Let us assume that $\tau_1 = 0$ and that the receivers are isotropic antennas. In this case $y_p(t) = s(t - \tau_p)$ where $s(t)$ is the source signal at the first receiver. We can write this in a simpler way by using the Fourier transform of $y_p(t)$ which we will denote by $Y_p(\omega)$ where ω is the angular frequency. Using this transform we have $Y_p(\omega) = e^{-j\omega\tau_p} S(\omega)$ where $S(\omega)$ is the Fourier transform of $s(t)$ and $j = \sqrt{-1}$. Now we will investigate the effect of replacing $e^{-j\omega\tau_p}$ with a constant phase $e^{-j\omega_0\tau_p}$. We would like the relative error introduced by using this substitution to be very small, i.e.

$$\begin{aligned} \epsilon(\omega)^2 &= \frac{|e^{-j\omega\tau_p} S(\omega) - e^{-j\omega_0\tau_p} S(\omega)|^2}{|e^{-j\omega\tau_p} S(\omega)|^2} \\ &= |e^{-j\omega\tau_p} - e^{-j\omega_0\tau_p}|^2 \\ &= 2 - 2\cos[(\omega - \omega_0)\tau_p] \ll 1. \end{aligned}$$

Using $\cos(-x) = \cos(x)$, a Taylor expansion of $\cos(x) \approx 1 - \frac{x^2}{2}$ and $\omega = 2\pi f$ where f is the frequency in Hz, we have $|(\omega - \omega_0)\tau_p| \ll 1$. In general we can write

$$2\pi\Delta f\tau \ll 1$$

where $\tau = \max_p |\tau_p|$ and Δf is the equivalent base-band bandwidth of the signal. The above relation will hold if $S(\omega - \omega_0)$ is negligible for $|\omega - \omega_0| > 2\pi\Delta f$. From this result we can conclude that if the array's bandwidth-delay product is small enough, we can approximate a delay with a phase shift. By choosing τ equal to the maximum delay on the array we can define the bandwidth for which the narrow-band assumption holds. If

the signal is not narrow-band we will divide it into narrow sub-bands by using an FFT or filter-banks.

Now let us stack the output of the antennas from a single sub-band at ω_0 into a vector \mathbf{y} . Using the phase-delay relation, we obtain

$$\mathbf{y}[n] = \begin{bmatrix} 1 \\ e^{-j\omega_0\tau_2} \\ \vdots \\ e^{-j\omega_0\tau_P} \end{bmatrix} s[n], \quad (3.1)$$

where $\mathbf{y}[n] = \mathbf{y}(T_s n)$ and $s[n]$ are the samples from $\mathbf{y}(t)$ and $s(t)$ with a sampling interval of T_s .

In the next section we will establish the relation between τ and the direction-of-arrival (DOA) of a source and extend the model to multiple sources.

3.1.2. GEOMETRICAL DELAYS AND THE EARTH'S ROTATION

Following the discussion in the previous section, we assume that the received signals at the antennas are sampled and subsequently split into narrow sub-bands. For simplicity, we will consider only a single sub-band in the rest of this section.

We revisit (3.1) and investigate the relation between the delays τ_i , the direction-of-arrival, and the array's geometry. As illustrated in Fig. 3.1 the geometrical delay between the i th and j th element is defined by the distance between them, b_{ij} also known as a baseline, and the cosine of the angle of arrival ϕ or $l = \cos(\phi)$. Because we have chosen $\tau_1 = 0$ as the reference and the array is placed along the x -axis the delay on the p th element is $\tau_p = -x_p \cos(\phi)/c =$ where c is the speed of light. Using $c = \lambda_0 f_0$ where λ is the wavelength and $\omega_0 = 2\pi f_0$ we have $y_p[n] = e^{j2\pi/\lambda_0 x_p l} s[n]$. Let the vector $\boldsymbol{\xi}_p = [x_p, y_p, z_p]^T$ denote the position of the p th receiving element in Earth-bound coordinates. Extending the result from 1-D to 3-D we have

$$\mathbf{y}[n] = \begin{bmatrix} e^{j\frac{2\pi}{\lambda_0} \boldsymbol{\xi}_1^T \mathbf{k}} \\ \vdots \\ e^{j\frac{2\pi}{\lambda_0} \boldsymbol{\xi}_P^T \mathbf{k}} \end{bmatrix} s[n]. \quad (3.2)$$

where

$$\mathbf{k} = \begin{bmatrix} l \\ m \\ n \end{bmatrix} = \begin{bmatrix} \cos(\phi) \cos(\theta) \\ \cos(\phi) \sin(\theta) \\ \sin(\phi) \end{bmatrix}$$

is the (apparent) direction vector with coordinates (l, m, n) . Because \mathbf{k} has unit norm we only need two coordinates (l, m) , while the third coordinate can be calculated using $n = \sqrt{1 - l^2 - m^2}$. (The notation in (3.2) is ambiguous as we use the symbol n both for the sample index and for the 3rd spatial coordinate. However, in this thesis we will not use the 3rd spatial coordinate anymore. The letter n , unless state otherwise, is the sample index.)

Although the sources are considered stationary, because of the Earth's rotation the apparent position of the celestial sources will change with time. For this reason the data

is split into short blocks or “snapshots” of N samples, where the exact value of N depends on the resolution of the instrument.

We stack the output of the P antennas at a single sub-band into a vector $\mathbf{y}_k[n]$, where $n = 1, \dots, N$ denotes the sample index, and $k = 1, \dots, K$ denotes the snapshot index.

The vector in (3.2) which relates $s[n]$ to $\mathbf{y}[n]$ is called the array response vector. To describe this vector while taking the Earth’s rotation into account, we first need to define a fixed coordinate system with respect to celestial sphere. Assume that this fixed coordinate system is based on the right ascension (α) and declination (δ) of a source, and defined as

$$\boldsymbol{\beta} = \begin{bmatrix} \cos(\delta) \cos(\alpha) \\ \cos(\delta) \sin(\alpha) \\ \sin(\delta) \end{bmatrix}.$$

The related Earth-bound direction vector \mathbf{k} with coordinates (l, m) (taking Earth rotation into account) is given by

$$\mathbf{k} = \mathbf{Q}_k(L, B)\boldsymbol{\beta},$$

where $\mathbf{Q}_k(L, B)$ is a 3×3 rotation matrix that accounts for the Earth rotation and depends on the time k and the observer’s longitude L and latitude B .

From this point forward, we will drop the subscript $[\cdot]_0$ for frequency related quantities such as λ_0 . These quantities are assumed to be calculated at the center of corresponding narrow sub-bands. For the q th source with coordinates (l_q, m_q) at the k th snapshot, the direction vector is \mathbf{k}_q . Stacking the coefficients for $p = 1, \dots, P$ into a vector $\mathbf{a}_{k,q} = \mathbf{a}_k(\mathbf{k}_q)$, we obtain the array response vector, which thus has the model

$$\mathbf{a}_{k,q} = \mathbf{a}_k(\mathbf{k}_q) = \frac{1}{\sqrt{P}} e^{\frac{j2\pi}{\lambda} \boldsymbol{\Xi}^T \mathbf{k}_q} = \frac{1}{\sqrt{P}} e^{\frac{j2\pi}{\lambda} \boldsymbol{\Xi}^T \mathbf{Q}_k(L, B) \boldsymbol{\beta}_q} \quad (3.3)$$

where $\boldsymbol{\Xi}$ is a $3 \times P$ matrix containing the positions $\boldsymbol{\xi}_p$ of the P receiving elements. We introduced a scaling by $1/\sqrt{P}$ as a normalization constant such that $\|\mathbf{a}_k(\mathbf{k}_q)\| = 1$. The entries of the array response vector are connected to the Fourier Transform coefficients familiar in radio astronomy models.

3.2. COVARIANCE MODEL

Assuming an array that is otherwise calibrated, the received antenna signals (generalizing (3.1)) $\mathbf{y}_k[n]$ can be modeled as

$$\mathbf{y}_k[n] = \mathbf{A}_k \mathbf{s}[n] + \mathbf{n}_k[n], \quad n = 1, \dots, N; \quad k = 1, \dots, K \quad (3.4)$$

where \mathbf{A}_k is a $P \times Q$ matrix whose columns are the array response vectors $[\mathbf{A}_k]_q = \mathbf{a}_{k,q}$, $\mathbf{s}[n]$ is a $Q \times 1$ vector representing the signals from the sky, and $\mathbf{n}_k[n]$ is a $P \times 1$ vector modeling the noise.

From the data, the system estimates covariance matrices of the input vector at each snapshot $k = 1, \dots, K$, as

$$\hat{\mathbf{R}}_k = \frac{1}{N} \sum_{n=1}^N \mathbf{y}_k[n] \mathbf{y}_k[n]^H, \quad k = 1, \dots, K. \quad (3.5)$$

Since the received signals and noise are Gaussian, these covariance matrix estimates form sufficient statistics for the problems we will address in this thesis [25]. The covariance matrices are given by

$$\mathbf{R}_k = \mathcal{E}\{\mathbf{y}_k \mathbf{y}_k^H\} \quad (3.6)$$

for which the model is

$$\mathbf{R}_k = \mathbf{A}_k \mathbf{\Sigma} \mathbf{A}_k^H + \mathbf{R}_{n,k}, \quad (3.7)$$

where $\mathbf{\Sigma} = \mathcal{E}\{\mathbf{s} \mathbf{s}^H\}$ and $\mathbf{R}_{n,k} = \mathcal{E}\{\mathbf{n}_k \mathbf{n}_k^H\}$ are the source and noise covariance matrices, respectively. We have assumed that sky sources are stationary, and if we also assume that they are independent, we can model $\mathbf{\Sigma} = \text{diag}(\boldsymbol{\sigma})$ where

$$\boldsymbol{\sigma} = [\sigma_1 \quad \dots \quad \sigma_Q]^T \quad (3.8)$$

represents the intensity of the sources. To connect the covariance data model (3.7) to language more familiar to radio astronomers, let us take a closer look at the elements of the matrix \mathbf{R}_k . Temporarily ignoring the noise covariance matrix $\mathbf{R}_{n,k}$ we note that

$$\begin{aligned} [\mathbf{R}_k]_{ij} &= \frac{1}{P} \sum_{q=1}^Q \sigma_q a_{kqi} a_{kqj}^* \\ &= \frac{1}{P} \sum_{q=1}^Q \sigma_q e^{j \frac{2\pi}{\lambda} (\boldsymbol{\xi}_i - \boldsymbol{\xi}_j)^T \mathbf{k}_q} \\ &= \frac{1}{P} \sum_{q=1}^Q \sigma_q e^{j \frac{2\pi}{\lambda} [(x_i - x_j)l_q + (y_i - y_j)m_q + (z_i - z_j)\sqrt{1 - l_q^2 - m_q^2}]} \end{aligned} \quad (3.9)$$

If we denote $\frac{1}{\lambda} [x_i - x_j, y_i - y_j, z_i - z_j]^T \mathbf{Q}_k(L, B) = [u_{ij}, v_{ij}, w_{ij}]^T$, then we can write $[\mathbf{R}_k]_{ij} \equiv V(u_{ij}, v_{ij}, w_{ij})$, where $V(u, v, w)$ is the visibility function, and (u, v, w) are the spatial frequencies [10]. In other words, the entries of the covariance matrix \mathbf{R}_k are samples of the visibility function at a given frequency and time arranged in a matrix, and (3.7) represents the measurement equation in matrix form.

We can write this equation in several other ways. By vectorizing both sides of (3.7) and using the properties of Kronecker products (2.4), we obtain

$$\mathbf{r}_k = (\mathbf{A}_k^* \circ \mathbf{A}_k) \boldsymbol{\sigma} + \mathbf{r}_{n,k} \quad (3.10)$$

where $\mathbf{r}_k = \text{vect}(\mathbf{R}_k)$ and $\mathbf{r}_{n,k} = \text{vect}(\mathbf{R}_{n,k})$. After stacking the vectorized covariances for all of the snapshots we obtain

$$\mathbf{r} = \boldsymbol{\Psi} \boldsymbol{\sigma} + \mathbf{r}_n \quad (3.11)$$

where

$$\mathbf{r} = \begin{bmatrix} \mathbf{r}_1 \\ \vdots \\ \mathbf{r}_K \end{bmatrix}, \quad \boldsymbol{\Psi} = \begin{bmatrix} \mathbf{A}_1^* \circ \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_K^* \circ \mathbf{A}_K \end{bmatrix}, \quad \mathbf{r}_n = \begin{bmatrix} \mathbf{r}_{n,1} \\ \vdots \\ \mathbf{r}_{n,K} \end{bmatrix}. \quad (3.12)$$

Similarly we vectorize and stack the sample covariance matrices as

$$\hat{\mathbf{r}}_k = \text{vect}(\hat{\mathbf{R}}_k), \quad \hat{\mathbf{r}} = \begin{bmatrix} \hat{\mathbf{r}}_1 \\ \vdots \\ \hat{\mathbf{r}}_K \end{bmatrix}. \quad (3.13)$$

This collects all the available covariance data into a single vector. This is the model which we will use in Chapter 7 to estimate the intensity distribution of the sources from the measured covariance matrices.

3

3.3. GAIN MODEL

So far we have assumed an array of identical and isotropic antennas. However, the array which is used during the actual measurements will deviate from this assumption in various ways. The process of compensating for these deviations is called calibration. Depending on which effects we can compensate for, we might need to extend the data model to incorporate new correction (or calibration) parameters. It is important to realize that even if the antennas are of high quality material and design, the wiring and other environmental and instrumental effects make it necessary to perform a calibration step. Using a narrow-band assumption we can model the instrumental effects for each receiver as a complex factor g_p for $p = 1, \dots, P$ which we will call the antenna gain. If the gain is direction dependent, i.e. $g_p(\mathbf{k})$, we have a gain factor per antenna per direction. In this case we need to change the model for array response matrix \mathbf{A} to $\mathbf{G}_k \odot \mathbf{A}$ where

$$[\mathbf{G}_k]_{pq} = g_p(\mathbf{k}_q).$$

The covariance model (for a single snapshot) that includes this change is

$$\mathbf{R} = (\mathbf{G}_k \odot \mathbf{A}) \boldsymbol{\Sigma} (\mathbf{A} \odot \mathbf{G}_k)^H + \mathbf{R}_n.$$

Using the identity $(\mathbf{D}\mathbf{A}) \odot \mathbf{B} = \mathbf{D}(\mathbf{A} \odot \mathbf{B})$ where \mathbf{D} is a diagonal matrix, we can decompose the gain matrix as $\mathbf{G}_k = \mathbf{G}\mathbf{G}_0$ where \mathbf{G} is the diagonal direction independent gain of the antennas and \mathbf{G}_0 is the direction dependent part. This leads to the following model

$$\mathbf{R} = \mathbf{G}(\mathbf{G}_0 \odot \mathbf{A}) \boldsymbol{\Sigma} (\mathbf{A} \odot \mathbf{G}_0)^H \mathbf{G}^H + \mathbf{R}_n.$$

This model is used by Van der Tol and Van der Veen [7] for direction dependent calibration and its derivation (based on physical models for EM fields) can also be found in [2, pp. 52]. Taking all the elements in \mathbf{G}_0 as unknown parameters causes an unidentifiability problem. The number of unknowns in \mathbf{G}_0 can be reduced by using spatiotemporal continuity of its elements as suggested by [7]. In addition to instrumental causes, ionospheric effects also contribute to the direction dependent gains [49]. Treatment of direction dependent gains is an active field of research and goes beyond the scope this work.

For the rest of this discussion, we will assume that except for some direction independent effects, the antennas have the same directional behavior. This will allow us to model \mathbf{G}_0 also as a diagonal matrix (i.e. a shared gain per direction for all antennas) which will reduce the data model to

$$\mathbf{R} = \mathbf{G}\mathbf{A}\mathbf{G}_0\boldsymbol{\Sigma}\mathbf{G}_0^H\mathbf{A}^H\mathbf{G}^H + \mathbf{R}_n.$$

Any phase change caused by \mathbf{G}_0 will be compensated by \mathbf{G}_0^H due to the diagonal structure of both \mathbf{G}_0 and $\mathbf{\Sigma}$, and we can absorb $|\mathbf{G}_0|^2$ into the matrix $\mathbf{\Sigma}$. This will give us the simplified and commonly used data model for the array

$$\mathbf{R} = \mathbf{G}\mathbf{A}\mathbf{\Sigma}\mathbf{A}^H\mathbf{G}^H + \mathbf{R}_n$$

which we will use in Chapter 6 for calibration. After the calibration step the model defined in the previous section can be used for imaging purposes.

3.4. NON-PARAMETRIC SUBSPACE MODEL

In addition to the parametric modeling of the covariance matrices presented in the previous sections, in some applications it is beneficial to use a general linear algebraic approach to analyze the data. Matrix decomposition techniques such as QR, Singular Value Decomposition (SVD) and Eigenvalue Decomposition (EVD) form a powerful set of tools that can be used for data analysis. In this section we show how information can be obtained from a covariance data by using the EVD. The subspace based model presented here plays an important role in many detection and estimation methods used in signal processing [50–52]. We extend this model in Chapter 4 and use it in Chapter 5 to develop spatial filtering techniques.

Given the fact that the power of even the strongest celestial sources is ~ 15 dB below the noise level, for short integration intervals we can assume that the covariance matrix is dominated by the noise and man-made radio-frequency (RF) signals (if present). Even for RF free frequency bands, the data is dominated by a few strong sources and the contribution of the remaining sources is negligible. In this regime we can assume that the number of sources Q is smaller than the number of receivers (i.e. $Q < P$) and that the array response matrix \mathbf{A} is low-rank (and tall).

Let us assume that $\mathbf{R}_n = \sigma_n^2 \mathbf{I}_P$ where σ_n^2 is the noise power. In this case we can write

$$\mathbf{R} = \mathbf{A}\mathbf{\Sigma}\mathbf{A}^H + \sigma_n^2 \mathbf{I}_P.$$

Considering that $\mathbf{R}_0 = \mathbf{A}\mathbf{\Sigma}\mathbf{A}^H$ is a Hermitian matrix, we can diagonalize it using a unitary matrix \mathbf{U} and find its eigenvalue decomposition

$$\mathbf{R}_0 = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^H = [\mathbf{U}_0 \quad \mathbf{U}_1] \begin{bmatrix} \mathbf{\Lambda}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{U}_0^H \\ \mathbf{U}_1^H \end{bmatrix}$$

where $\mathbf{\Lambda}_0$ is a $Q \times Q$ diagonal matrix of positive (and non-zero) eigenvalues of \mathbf{R}_0 , \mathbf{U}_0 is a $P \times Q$ matrix with the eigenvectors corresponding to $\mathbf{\Lambda}_0$ and \mathbf{U}_1 forms a unitary basis for the null space of \mathbf{A} . By using $\mathbf{U}\mathbf{U}^H = \mathbf{I}_P$ we have $\sigma_n^2 \mathbf{I}_P = \sigma_n^2 \mathbf{U}\mathbf{U}^H$ and

$$\mathbf{R} = \mathbf{A}\mathbf{\Sigma}\mathbf{A}^H + \sigma_n^2 \mathbf{U}\mathbf{U}^H = [\mathbf{U}_0 \quad \mathbf{U}_1] \begin{bmatrix} \mathbf{\Lambda}_0 + \sigma_n^2 \mathbf{I}_Q & \mathbf{0} \\ \mathbf{0} & \sigma_n^2 \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{U}_0^H \\ \mathbf{U}_1^H \end{bmatrix}$$

is the eigenvalue decomposition of \mathbf{R} . This relation enables us to find two subspaces spanned by \mathbf{U}_0 and \mathbf{U}_1 from eigenvalue decomposition of \mathbf{R} . Because the contribution of the subspace spanned by \mathbf{U}_1 to the covariance matrix is entirely defined by the noise

power, this subspace is denoted as the noise subspace. In contrast the subspace defined by \mathbf{U}_0 is called the signal subspace.

The assumption that the noise covariance matrix is a scaled identity matrix is very limiting for practical scenarios. In Chapter 4 we will start by modeling the noise covariance matrix as an unknown diagonal matrix which will lead to the technique known as Factor Analysis and then extend it to more general models.

4

COMPLEX FACTOR ANALYSIS

4.1. INTRODUCTION

In this chapter we focus on a generic covariance decomposition technique called factor analysis (FA). This technique will be used as a common tool to address several radio astronomical problems in subsequent chapters. FA is a multivariate technique with the assumption that a covariance matrix \mathbf{R} can be modeled as

$$\mathbf{R} = \mathbf{A}\mathbf{A}^H + \mathbf{D}$$

where \mathbf{A} is a low ranked matrix and \mathbf{D} is a positive diagonal matrix. Real-valued FA was first introduced by Spearman [53] in 1904 to find a quantitative measure for intelligence, given a series of test results. However, it was the work done by Lawley, Anderson, Jöreskog and others between 1940 and 1970 [54–56] which developed FA as an established multivariate technique. Currently, FA is an important and popular part of the latent variable analysis with many applications in various fields of science [57].

Despite the popularity of FA in many fields, its application within the signal processing community has been limited. This is mainly due to the popularity of the Eigenvalue Decomposition (EVD) which is at the heart of many subspace based signal processing techniques. However, the application of EVD is limited to systems perturbed by white noise. In other words we must be able to model the noise covariance matrix as $\sigma^2 \mathbf{I}$ where σ^2 is the variance of the noise. If the noise covariance matrix is known, calibration and whitening techniques can be used as pre-processing procedures to make EVD applicable for systems where this assumption does not hold. However a more preferable approach is to develop techniques that can replace EVD for more practical and generic data models.

Using FA (or classical FA) for array processing has been suggested by [58] to address the case where the noise is unknown, independent and different for each element in the array i.e. the noise covariance matrix is a diagonal matrix with unknown elements. For cases where the noise covariance matrix is no longer diagonal but has a known structure, in this thesis we propose Extended FA (EFA).

The mentioned techniques work on a single covariance matrix. However in many applications the desired subspace changes rapidly which means that a series of short-term covariance matrices are available requiring an extension towards “joint FA”. In this chapter an overview of the classical FA is given and the changes needed when extending this model to complex numbers are discussed. Further it is shown that applying subspace estimation techniques for each short-term covariance leads to sub-optimal estimates, since the stationarity of the diagonal (or extended structure) is not used. To address this issue an estimation technique is developed that estimates the desired subspaces jointly and is flexible enough to include generic noise models.

Estimating the unknown parameters leads to a non-linear optimization problem. Various optimization techniques for FA, EFA and Joint EFA (JEFA) will be addressed here. In particular a Gauss-Newton-Krylov based technique that solves the non-linear optimization efficiently in both memory usage and complexity will be presented. The application of EFA and JEFA for RFI mitigation and array calibration are discussed in the chapters that follow.

The setup of this chapter is as follows: In Sec. 4.2.1 we discuss the data and covariance models for FA and its extensions, in Sec. 4.4 we derive the Cramér-Rao bound for the estimated parameters, Sec. 4.5 gives a short overview of the algorithm used for classical FA, Sec. 4.6 describes the algorithm we use to estimate the parameters for JEFA, various detection methods are discussed in Sec. 4.7 and in Sec. 4.8 we use simulations to evaluate the performance of the proposed methods.

4.2. CLASSICAL FACTOR ANALYSIS MODEL

4.2.1. DATA MODEL

Here we consider a system which has P receiving elements that are exposed to Q sources with complex Gaussian distribution. Each element could have different gain and noise level. We assume that the noise is a proper complex Gaussian process and uncorrelated between different receiving elements. Stacking the received signals from each antenna, we could model the system as

$$\mathbf{y}(t) = \mathbf{A}_0 \mathbf{x}(t) + \mathbf{n}(t), \quad (4.1)$$

where \mathbf{y} is a $P \times 1$ vector of received signals, \mathbf{A}_0 is a $P \times Q$ array response containing all the gains and spatial signatures of the sources, \mathbf{x} is a $Q \times 1$ vector representing source signals and \mathbf{n} is a $P \times 1$ vector modeling noise contributions in the system.

With the assumption that the sources and noise contributions are uncorrelated, we can write the covariance matrix for \mathbf{y} as

$$\begin{aligned} \mathbf{R} &= \mathcal{E}\{\mathbf{y}\mathbf{y}^H\} \\ &= \mathbf{A}_0 \mathbf{R}_x \mathbf{A}_0^H + \mathbf{R}_n \\ &= \mathbf{R}_0 + \mathbf{R}_n. \end{aligned} \quad (4.2)$$

The noise-free covariance matrix, \mathbf{R}_0 , is a Hermitian positive semi-definite matrix of rank Q and as such it could be written in Gramian form as $\mathbf{R}_0 = \mathbf{A}\mathbf{A}^H$ where \mathbf{A} is a $P \times Q$ matrix and $\text{rank}(\mathbf{A}) = Q$. The noise covariance matrix $\mathbf{R}_n = \mathcal{E}\{\mathbf{n}\mathbf{n}^H\}$ and we assume that \mathbf{R}_n^{-1} exists. In the case where $\mathbf{R}_n = \sigma^2 \mathbf{I}$, the column span of \mathbf{A} can be estimated using

an eigenvalue decomposition of \mathbf{R} . In other words let the (economical) singular value decomposition (SVD) of \mathbf{A} be

$$\mathbf{A} = \mathbf{U}_0 \mathbf{\Sigma} \mathbf{V}^H$$

where \mathbf{U}_0 is a semi-unitary matrix of size $P \times Q$ forming an orthonormal basis for the column space of \mathbf{A} , \mathbf{V} is a $Q \times Q$ unitary matrix forming an orthonormal basis for the row space of \mathbf{A} and $\mathbf{\Sigma}$ is a $Q \times Q$ diagonal matrix containing the singular values of \mathbf{A} in descending order. Similarly let the eigenvalue decomposition of \mathbf{R} be

$$\mathbf{R} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^H, \quad (4.3)$$

where $\lambda_1 \geq \dots \geq \lambda_P$. Then for the case where $\mathbf{R}_n = \sigma^2 \mathbf{I}_P = \sigma^2 \mathbf{U} \mathbf{U}^H$ we have

$$\mathbf{U} = [\mathbf{U}_0 \quad \mathbf{U}_1].$$

This makes it possible to find the column span of \mathbf{A} from \mathbf{R} , using the eigenvalue decomposition of \mathbf{R} . A property which has been used in many subspace based array processing techniques such as ESPRIT [59] and MUSIC [60]. However this technique fails when the noise is not white and \mathbf{R}_n takes another model.

For classical FA we assume that the noise between the elements is independent, but not necessarily identical and as such $\mathbf{R}_n = \mathbf{D}$, is a diagonal matrix. Rewriting the model as

$$\mathbf{R} = \mathbf{A} \mathbf{A}^H + \mathbf{D}, \quad (4.4)$$

we arrive at the classical FA model [56, 61] with the exception that (4.4) contains complex numbers.

As discussed in Sec. 2.3 we rarely have direct access to the covariance matrix \mathbf{R} and we need to construct an estimate of this matrix from measured samples. Here we assume to have access to N samples from \mathbf{y} and construct the sample covariance matrix as

$$\hat{\mathbf{R}} = \frac{1}{N} \sum_{n=0}^{N-1} \mathbf{y}[n] \mathbf{y}[n]^H.$$

Given $\hat{\mathbf{R}}$ we are interested in finding estimates $\hat{\mathbf{A}}$ and $\hat{\mathbf{D}}$ for the FA model.

Various problems arise when we try to find the model parameters \mathbf{A} and \mathbf{D} . For example consider the following scenario with

$$\mathbf{A}_1 = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ \vdots & \vdots \\ 1 & 1 \end{bmatrix}$$

then for any $\mathbf{R} = \mathbf{A}_1 \mathbf{A}_1^H + \mathbf{D}_1$ we also have $\mathbf{R} = \mathbf{A}_2 \mathbf{A}_2^H + \mathbf{D}_2$ where

$$\mathbf{A}_2 = \sqrt{2} \begin{bmatrix} 1/2 \\ 1 \\ \vdots \\ 1 \end{bmatrix}, \quad \mathbf{D}_2 = \mathbf{D}_1 + \frac{1}{2} \mathbf{e}_1 \mathbf{e}_1^T$$

and \mathbf{e}_i is the i th column of the identity matrix. This example illustrates that there is an identifiability problem that needs to be addressed. FA provides a complete set of tools and techniques that address these problems. In the rest of this chapter we will discuss if and how these problems change when we extend the model first to complex numbers and then to Extended FA. In the following sections, with the exception of Sec. 4.10 where we discuss the identification problem in more detail, we assume that the matrix \mathbf{D} can be identified uniquely.

4.2.2. FACTOR ROTATION

Given a FA model $\mathbf{R} = \mathbf{A}_1 \mathbf{A}_1^H + \mathbf{D}$ there is also another equivalent model with $\mathbf{R} = \mathbf{A}_2 \mathbf{A}_2^H + \mathbf{D}$ where $\mathbf{A}_2 = \mathbf{A}_1 \mathbf{Q}$ and \mathbf{Q} is a unitary matrix. It is then clear that the problem of finding \mathbf{A} is not unique and requires proper identifiability conditions (or constraints). Note that this type of identifiability problem has no effect on the matrix \mathbf{D} .

One way to produce unique model solutions is by setting certain elements in \mathbf{A} equal to zero or forcing them to be purely imaginary or purely real. We will refer to this type of constraints as "structural constraints". An example of such a constraint is forcing \mathbf{A} to be lower triangular with the elements along the diagonal to be purely real and positive numbers. From linear algebra we know this as LQ factorization of $\mathbf{A} = \mathbf{LQ}$ (or QR factorization of \mathbf{A}^H). Indeed the lower triangular matrix \mathbf{L} for a nonsingular matrix is unique and the LQ factorization of any \mathbf{A} that satisfies the FA model differs only in the rotation matrix \mathbf{Q} .

The second type of constraint does not involve the structure of \mathbf{A} but forces the columns of this matrix to be orthogonal with respect to a certain weighting matrix \mathbf{W} . This can be written as

$$\mathbf{A}^H \mathbf{W} \mathbf{A} = \mathbf{\Gamma} \quad (4.5)$$

where $\mathbf{\Gamma}$ is a real diagonal matrix with $\gamma_1 > \dots > \gamma_Q$. Note that $\mathbf{\Gamma}$ is unknown and the constraint follows from its diagonal structure or in terms of the columns of \mathbf{A} , $\mathbf{a}_i^H \mathbf{W} \mathbf{a}_j = 0$ for $i \neq j$ and $\mathbf{a}_i^H \mathbf{W} \mathbf{a}_i$ is real. For real FA this is a sufficient constraint, however for complex FA additional phase constraints are required to find a unique \mathbf{A} . To illustrate let us consider $\mathbf{A}' = \mathbf{A} \mathbf{\Phi}$ where $\mathbf{\Phi}$ is a complex diagonal matrix which has elements of the form $e^{j\phi}$, or

$$\mathbf{\Phi} = \begin{bmatrix} e^{j\phi_1} & & \\ & \ddots & \\ & & e^{j\phi_Q} \end{bmatrix}$$

then both \mathbf{A} and \mathbf{A}' satisfy the same constraints (4.5). Hence, when we are dealing with complex numbers there are Q additional phases that need to be chosen. This is also true for EVD and SVD when applied to complex numbers. This phase ambiguity can be addressed by forcing the diagonal part of $\mathbf{A}^T \mathbf{A}$ to be real. This is obviously the case when real FA model is considered.

Two popular choices for \mathbf{W} are $\mathbf{W} = \mathbf{I}$ or $\mathbf{W} = \mathbf{D}^{-1}$ [55]. If we choose $\mathbf{W} = \mathbf{I}$ then $\mathbf{U}_0 = \mathbf{A} \mathbf{\Gamma}^{-1/2}$ has orthogonal columns and we have

$$\mathbf{R}_0 = \mathbf{U}_0 \mathbf{\Gamma} \mathbf{U}_0^H. \quad (4.6)$$

In other words, by solving \mathbf{A} with this constraint we find the EVD for \mathbf{R}_0 with $\mathbf{\Gamma}$ as eigenvalues.

The choice of $\mathbf{W} = \mathbf{D}^{-1}$ is advantageous when we need to compute \mathbf{R}^{-1} . Using Woodbury's identity we have

$$\mathbf{R}^{-1} = \mathbf{D}^{-1} - \mathbf{D}^{-1} \mathbf{A} (\mathbf{I}_Q + \mathbf{A}^H \mathbf{D}^{-1} \mathbf{A})^{-1} \mathbf{A}^H \mathbf{D}^{-1}. \quad (4.7)$$

Using an eigenvalue decomposition of $(\mathbf{A}^H \mathbf{D}^{-1} \mathbf{A} + \mathbf{I}_Q) = \mathbf{Q}(\mathbf{\Gamma} + \mathbf{I}_Q)\mathbf{Q}^H$ we can find a \mathbf{Q} such that $\mathbf{A}' = \mathbf{A}\mathbf{Q}$ satisfies the constraints, and reduce the computational complexity of full inversion from $O(P^3)$ to a diagonal inversion and $O(Q^3)$ for the eigenvalue decomposition. In other words if we need to compute \mathbf{R}^{-1} we can get \mathbf{Q} required for this type of constraints for free.

4.2.3. SCALE INVARIANCE

If we assume that $\mathbf{R}_1 = \mathbf{A}_1 \mathbf{A}_1^H + \mathbf{D}_1$ and hence satisfies the FA model, then any $\mathbf{R}_2 = \mathbf{G} \mathbf{R}_1 \mathbf{G}^H$ where \mathbf{G} is a diagonal (complex) matrix also satisfies FA model with $\mathbf{A}_2 = \mathbf{G} \mathbf{A}_1$ and real diagonal matrix $\mathbf{D}_2 = \mathbf{G} \mathbf{D}_1 \mathbf{G}^H$. This means that unknown antenna gains do not affect the model and that FA can be used on an uncalibrated array.

It is worth noting that $\mathbf{A}_1^H \mathbf{D}_1^{-1} \mathbf{A}_1 = \mathbf{A}_2^H \mathbf{D}_2^{-1} \mathbf{A}_2 = \mathbf{\Gamma}$ is scale invariant and $\text{tr}(\mathbf{\Gamma})$ could be defined as a metric for average SNR.

4.3. EXTENSIONS OF THE CLASSICAL MODEL

4.3.1. JOINT FACTOR ANALYSIS MODEL

Now we extend the model by assuming access to a series of sample covariance matrices $\hat{\mathbf{R}}_m$ where $m = 1, \dots, M$, from M independent "snapshots", each containing N samples. We assume the following model for $\mathcal{E}\{\hat{\mathbf{R}}_m\} = \mathbf{R}_m$:

$$\mathbf{R}_m = \mathbf{A}_m \mathbf{A}_m^H + \mathbf{\Psi}, \quad m = 1, \dots, M \quad (4.8)$$

where \mathbf{A}_m is a low rank matrix of size $P \times Q_m$ with $Q_m < P$ for all $m = 1, \dots, M$ and $\mathbf{\Psi}$ is a positive-definite matrix common among all the M models. Depending on the application we are interested in \mathbf{A}_m and $\mathbf{\Psi}$ or only one of them. In many applications we are just interested in the column span of \mathbf{A}_m .

In common applications, $\mathbf{\Psi}$ represents the noise covariance matrix of the system, \mathbf{R}_n . For a system which is calibrated or has identical components, it is common to assume that $\mathbf{\Psi} = \sigma^2 \mathbf{I}_P$. However, other models where $\mathbf{\Psi}$ is not the noise covariance matrix will also be discussed in the following sections.

In the following sections we will allow more general models for $\mathbf{\Psi}$ and discuss techniques for finding \mathbf{A}_m and $\mathbf{\Psi}$ from noisy estimates $\hat{\mathbf{R}}_m$. We will also estimate \mathbf{A}_m and $\mathbf{\Psi}$ jointly. When $\mathbf{\Psi} = \mathbf{D}$ (i.e. a diagonal matrix) we call this model Joint Factor Analysis (JFA).

4.3.2. EXTENDED AND JOINT EXTENDED FA MODEL

Factor Analysis can be extended to a more general model where a certain structure is assumed to be known for $\mathbf{\Psi}$. Here we consider $\mathbf{\Psi}$ of the form

$$\mathbf{\Psi} = \mathbf{M} \odot \mathbf{\Psi}$$

where \mathbf{M} is a symmetric matrix containing only ones and zeros. We call \mathbf{M} a mask matrix. We can model various types of covariance matrices using this approach (for example: block-diagonal matrices, band matrices, sparse matrices, etc.). We assume \mathbf{M} to be known based on the application. The Extended FA model then becomes

$$\mathbf{R} = \mathbf{A}\mathbf{A}^H + \mathbf{M} \odot \mathbf{\Psi}. \quad (4.9)$$

Similar to JFA we propose Joint Extended FA (JEFA), where we have

$$\mathbf{R}_m = \mathbf{A}_m \mathbf{A}_m^H + \mathbf{M} \odot \mathbf{\Psi}. \quad (4.10)$$

and we are interested in estimating $\hat{\mathbf{\Psi}}$ and $\hat{\mathbf{A}}_m$ jointly.

Note that for $\mathbf{M} = \mathbf{I}$ where \mathbf{I} is an identity matrix of appropriate size, this model reduces to JFA (and for $M = 1$ to classical FA).

4

4.4. CRAMÉR-RAO BOUND

Before considering algorithms for estimating the model parameters \mathbf{A}_m and $\mathbf{\Psi}$, we first derive the theoretical performance bound on the variance of their estimators which will allow us to evaluate the performance of the algorithms presented in the following sections.

Suppose that $\hat{\boldsymbol{\theta}}$ is the maximum likelihood estimator of $\boldsymbol{\theta}_0$, where $\boldsymbol{\theta}_0$ is the true parameter vector. Then the asymptotic distribution of $(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$ is $\mathcal{N}(\mathbf{0}, \mathbf{C})$ where \mathbf{C} is the inverse of the Fisher information matrix \mathbf{F} . \mathbf{C} is the Cramér-Rao lower bound (CRB) for an unbiased estimator [62]. For normally distributed data with covariance matrix \mathbf{R} , the Fisher information matrix is [40]

$$\mathbf{F} = N \mathbf{J}^H (\mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \mathbf{J}$$

where $\mathbf{J} = \partial \text{vect}(\mathbf{R}) / \partial \boldsymbol{\theta}^T$ the Jacobian matrix.

If the model parameters are not unique, \mathbf{F} is singular, then for identifiability we need to pose additional constraints on $\boldsymbol{\theta}$, say $\mathbf{h}(\boldsymbol{\theta}_0) = \mathbf{0}$, where $\mathbf{h}(\boldsymbol{\theta})$ is a vector of functions. Let the Jacobian of $\mathbf{h}(\boldsymbol{\theta})$ be given by

$$\mathbf{H}(\boldsymbol{\theta}) := \frac{\partial \mathbf{h}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T}.$$

The constrained CRB, \mathbf{C} , is then given by [63]

$$\mathbf{C} = \mathbf{U}(\mathbf{U}^H \mathbf{F}(\boldsymbol{\theta}_0) \mathbf{U})^{-1} \mathbf{U}^H \quad (4.11)$$

where \mathbf{U} is a semi-unitary matrix, and the columns of \mathbf{U} form an orthonormal basis for the null-space of \mathbf{H} such that $\mathbf{H}(\boldsymbol{\theta}_0) \mathbf{U}(\boldsymbol{\theta}_0) = \mathbf{0}$. The constraints should be chosen such that $\mathbf{U}^H \mathbf{F}(\boldsymbol{\theta}_0) \mathbf{U}$ is invertible.

4.4.1. CRB FOR FA AND EFA

We will now apply these results to our situation. First we consider only a single snapshot \mathbf{R}_m with model $\mathbf{R}_m = \mathbf{A}_m \mathbf{A}_m^H + \mathbf{\Psi}$, as given by (4.8). Using only single snapshot is equivalent to finding the CRB for classical FA and EFA. The structural constraint $\mathbf{\Psi} = \mathbf{M} \odot \mathbf{\Psi}$ in (4.9) is satisfied if we will parametrize $\text{vect}(\mathbf{\Psi})$ as

$$\text{vect}(\mathbf{\Psi}) = \mathbf{S}_U \boldsymbol{\psi} + \mathbf{S}_L \boldsymbol{\psi}^* + (\mathbf{I}_P \odot \mathbf{I}_P) \mathbf{d},$$

where \mathbf{S}_L and \mathbf{S}_U are suitable selection matrices based on the structure of \mathbf{M} , the entries of the strictly upper-triangular part of Ψ are stacked into the vector ψ , its diagonal entries $\mathbf{d} = \text{vectdiag}(\Psi)$. The unknown complex parameters are stacked into a vector θ_m ,

$$\theta_m = \begin{bmatrix} \text{vect}(\mathbf{A}_m) \\ \text{vect}(\mathbf{A}_m^*) \\ \psi \\ \psi^* \\ \mathbf{d} \end{bmatrix} \quad (4.12)$$

which has number of entries

$$n = 2PQ_m + \text{tr}(\mathbf{M}^2), \quad (4.13)$$

where Q_m is the number of sources at m th snapshot.

To derive the Fisher information matrix, we partition the corresponding Jacobians $\mathbf{J}_m = \partial \text{vect}(\mathbf{R}) / \partial \theta_m^T$, to conform with the partitioning of θ_m such that

$$\mathbf{J}_m = [\mathbf{J}_{\mathbf{A}_m}, \mathbf{J}_{\mathbf{A}_m^*}, \mathbf{J}_\psi, \mathbf{J}_{\psi^*}, \mathbf{J}_\mathbf{d}].$$

Using Wirtinger derivatives we find

$$\begin{aligned} \mathbf{J}_{\mathbf{A}_m} &= (\mathbf{A}_m^* \otimes \mathbf{I}_P) \\ \mathbf{J}_{\mathbf{A}_m^*} &= (\mathbf{I}_P \otimes \mathbf{A}_m) \mathbf{K}^{PQ_m} \\ \mathbf{J}_\psi &= \mathbf{S}_U \\ \mathbf{J}_{\psi^*} &= \mathbf{S}_L \\ \mathbf{J}_\mathbf{d} &= (\mathbf{I}_P \circ \mathbf{I}_P) \end{aligned} \quad (4.14)$$

where \mathbf{K}^{PQ} is a permutation matrix such that $\text{vect}(\mathbf{X}^T) = \mathbf{K}^{PQ} \text{vect}(\mathbf{X})$ for \mathbf{X} a $P \times Q$ matrix. We also used the relation

$$\begin{aligned} \text{vect}(\mathbf{R}_m) &= (\mathbf{A}_m^* \otimes \mathbf{I}_P) \text{vect}(\mathbf{A}_m) + \text{vect}(\Psi) \\ &= (\mathbf{I}_P \otimes \mathbf{A}_m) \mathbf{K}^{PQ_m} \text{vect}(\mathbf{A}_m^*) + \text{vect}(\Psi). \end{aligned}$$

The unconstrained Fisher information \mathbf{F}_m is singular, because as mentioned in Sec. 4.2.2 the FA model is invariant with respect to a multiplication of the matrix \mathbf{A}_m with a unitary matrix \mathbf{Q} at the right, which has Q_m^2 unknown elements (there are $2Q_m^2$ real unknowns, however $\mathbf{Q}^H \mathbf{Q} = \mathbf{I}_Q$ puts Q^2 constraints on the unknowns reducing the total to Q_m^2). As a result for identifiability, we need to pose Q_m^2 constraints on the matrix \mathbf{A}_m . Without loss of generality we choose $\mathbf{A}_m^H \mathbf{A}_m$ to be diagonal (which poses $Q_m(Q_m - 1)$ real-valued constraints), and $\text{diag}(\mathbf{A}_m^T \mathbf{A}_m)$ to be real (which poses another Q_m constraints). A complete discussion of constraints on \mathbf{A}_m is given in Sec. 4.2.2.

To write this as a function $\mathbf{h}(\theta_m) = \mathbf{0}$, let $\mathbf{E}_1 = (\mathbf{I}_{Q_m} \circ \mathbf{I}_{Q_m})^T$ and let \mathbf{E}_2 be a complementary set of rows such that $[\mathbf{E}_1^T, \mathbf{E}_2^T]^T$ is a permutation matrix. Then $\mathbf{E}_1 \text{vect}(\mathbf{A}_m^T \mathbf{A}_m)$ selects the diagonal elements of $\mathbf{A}_m^T \mathbf{A}_m$, and $\mathbf{E}_2 \text{vect}(\mathbf{A}_m^H \mathbf{A}_m)$ selects the off-diagonal entries of $\mathbf{A}_m^H \mathbf{A}_m$. The constraint function $\mathbf{h} = [\mathbf{h}_1^T, \mathbf{h}_2^T]^T$ is then

$$\mathbf{h}(\theta_m) = \begin{bmatrix} \mathbf{h}_1(\theta_m) \\ \mathbf{h}_2(\theta_m) \end{bmatrix} = \begin{bmatrix} \mathbf{E}_1 \text{vect}(\mathbf{A}_m^T \mathbf{A}_m - \mathbf{A}_m^H \mathbf{A}_m^*) \\ \mathbf{E}_2 \text{vect}(\mathbf{A}_m^H \mathbf{A}_m) \end{bmatrix} = \mathbf{0}$$

and its Jacobian $\mathbf{H}(\boldsymbol{\theta}_m)$ is

$$\mathbf{H}(\boldsymbol{\theta}_m) = [\mathbf{H}_{\mathbf{A}_m} \quad \mathbf{0} \quad \mathbf{0} \quad \mathbf{0}]$$

where

$$\mathbf{H}_{\mathbf{A}_m} = \begin{bmatrix} \frac{\partial \mathbf{h}_1}{\partial \text{vect}^T(\mathbf{A}_m)} & \frac{\partial \mathbf{h}_1}{\partial \text{vect}^T(\mathbf{A}_m^*)} \\ \frac{\partial \mathbf{h}_2}{\partial \text{vect}^T(\mathbf{A}_m)} & \frac{\partial \mathbf{h}_2}{\partial \text{vect}^T(\mathbf{A}_m^*)} \end{bmatrix}, \quad (4.15)$$

and the trailing zeros correspond to derivatives of $\mathbf{h}(\boldsymbol{\theta}_m)$ with respect to $\boldsymbol{\psi}$, $\boldsymbol{\psi}^*$ and \mathbf{d} . The needed derivatives are given by

$$\begin{aligned} \frac{\partial \mathbf{h}_1}{\partial \text{vect}^T(\mathbf{A}_m)} &= \mathbf{E}_1 [(\mathbf{I}_Q \otimes \mathbf{A}_m^T) + (\mathbf{A}_m^T \otimes \mathbf{I}_Q) \mathbf{K}^{P,Q}] \\ \frac{\partial \mathbf{h}_1}{\partial \text{vect}^T(\mathbf{A}_m^*)} &= -\mathbf{E}_1 [(\mathbf{I}_Q \otimes \mathbf{A}_m^H) + (\mathbf{A}_m^H \otimes \mathbf{I}_Q) \mathbf{K}^{P,Q}] \\ \frac{\partial \mathbf{h}_2}{\partial \text{vect}^T(\mathbf{A}_m)} &= \mathbf{E}_2 (\mathbf{I}_Q \otimes \mathbf{A}_m^H) \\ \frac{\partial \mathbf{h}_2}{\partial \text{vect}^T(\mathbf{A}_m^*)} &= \mathbf{E}_2 (\mathbf{A}_m^T \otimes \mathbf{I}_P) \mathbf{K}^{P,Q} \end{aligned}$$

Using QR or SVD on $\mathbf{H}(\boldsymbol{\theta}_m)$, we can find a basis \mathbf{U}_m for the null-space of $\mathbf{H}(\boldsymbol{\theta}_m)$, and calculate the Constrained CRB \mathbf{C}_m for a single measurement.

As shown above, the constraint function $\mathbf{h}(\boldsymbol{\theta}_m)$ depends only on \mathbf{A}_m , and its derivatives with respect to $\boldsymbol{\theta}_\Psi$ are zero. This allows us to partition \mathbf{U}_m as

$$\mathbf{U}_m = \begin{bmatrix} \mathbf{U}_{\mathbf{A}_m} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}.$$

Using this partitioning and (4.11) we have

$$\begin{aligned} \mathbf{C}_m &= \begin{bmatrix} \mathbf{C}_{\mathbf{A}_m \mathbf{A}_m} & \mathbf{C}_{\mathbf{A}_m \Psi} \\ \mathbf{C}_{\mathbf{A}_m \Psi}^H & \mathbf{C}_{m, \Psi \Psi} \end{bmatrix} \\ &= \frac{1}{M} \mathbf{U}_m \begin{bmatrix} \mathbf{U}_{\mathbf{A}_m}^H \mathbf{F}_{\mathbf{A}_m \mathbf{A}_m} \mathbf{U}_{\mathbf{A}_m} & \mathbf{U}_{\mathbf{A}_m}^H \mathbf{F}_{\mathbf{A}_m \Psi} \\ \mathbf{F}_{\mathbf{A}_m \Psi}^H \mathbf{U}_{\mathbf{A}_m} & \mathbf{F}_{m, \Psi \Psi} \end{bmatrix}^{-1} \mathbf{U}_m^H \end{aligned} \quad (4.16)$$

and hence

$$\mathbf{C}_{m, \Psi \Psi} = \frac{1}{N} \left[\mathbf{F}_{m, \Psi \Psi} - \mathbf{F}_{\mathbf{A}_m \Psi}^H \mathbf{U}_{\mathbf{A}_m} (\mathbf{U}_{\mathbf{A}_m}^H \mathbf{F}_{\mathbf{A}_m \mathbf{A}_m} \mathbf{U}_{\mathbf{A}_m})^{-1} \mathbf{U}_{\mathbf{A}_m}^H \mathbf{F}_{\mathbf{A}_m \Psi} \right]^{-1}. \quad (4.17)$$

Suppose we solve a joint (E)FA problem in a non-joint fashion. The performance of applying (E)FA on each snapshot separately to find $\hat{\boldsymbol{\Psi}}$ (or $\hat{\mathbf{D}}$) and averaging these results, follows by assuming that the estimates $\hat{\boldsymbol{\theta}}$ are independent, then for $\hat{\boldsymbol{\Psi}} = 1/M \sum_m \hat{\boldsymbol{\Psi}}_m$ the performance bound becomes

$$\mathbf{C}_\Psi = 1/M^2 \sum_m \mathbf{C}_{m, \Psi \Psi}. \quad (4.18)$$

This bound is higher than the bound we will derive for the entire dataset shortly.

4.4.2. CRB FOR JEFA

Using the results of the previous section, we have the Fisher Information for each snapshot and we can use them to find the CRB of the entire dataset. Because the time samples are independent, the joint loglikelihood of the entire dataset becomes

$$l(\boldsymbol{\theta}) = \sum_m l_m(\boldsymbol{\theta}_m). \quad (4.19)$$

For a simpler representation we define

$$\boldsymbol{\theta}_{\mathbf{A}_m} = \begin{bmatrix} \text{vect}(\mathbf{A}_m) \\ \text{vect}(\mathbf{A}_m^*) \end{bmatrix}. \quad (4.20)$$

and

$$\boldsymbol{\theta}_{\Psi} = \begin{bmatrix} \boldsymbol{\psi} \\ \boldsymbol{\psi}^* \\ \mathbf{d} \end{bmatrix}. \quad (4.21)$$

The total number of elements for $\boldsymbol{\theta}$ becomes

$$n = 2P \sum_{m=1}^M Q_m + \text{tr}(\mathbf{M}^2). \quad (4.22)$$

Now we need to find the gradients of the new loglikelihood. First we take derivatives with respect to $\boldsymbol{\theta}_{\mathbf{A}_m}$. If we assume that all \mathbf{A}_m are independent i.e. $\partial l_m(\boldsymbol{\theta}_{\mathbf{A}_m}) / \partial \boldsymbol{\theta}_{\mathbf{A}_j} = \mathbf{0}$ for $m \neq j$ and we find

$$\frac{\partial l(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_{\mathbf{A}_m}} = \frac{\partial l_m(\boldsymbol{\theta}_m)}{\partial \boldsymbol{\theta}_{\mathbf{A}_m}} \quad (4.23)$$

which is the same as for estimating the ML separately. However, for Ψ we have

$$\frac{\partial l(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_{\Psi}} = \sum_m \frac{\partial l_m(\boldsymbol{\theta}_m)}{\partial \boldsymbol{\theta}_{\Psi}}. \quad (4.24)$$

The Fisher information can also be written as

$$\mathbf{F}_m = \mathcal{E} \left\{ \frac{\partial}{\partial \boldsymbol{\theta}_m^T} \left(\frac{\partial l_m(\boldsymbol{\theta}_m)}{\partial \boldsymbol{\theta}_m} \right) \right\},$$

and

$$\mathbf{F}_{total} = \mathcal{E} \left\{ \frac{\partial}{\partial \boldsymbol{\theta}^T} \left(\frac{\partial l(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right) \right\}.$$

Combining this relation with (4.23) and (4.24), we can write the Fisher information for the entire dataset as

$$\mathbf{F}_{total, \mathbf{A}_m \mathbf{A}_m} = \mathbf{F}_{m, \mathbf{A}_m \mathbf{A}_m} \quad (4.25)$$

$$\mathbf{F}_{total, \mathbf{A}_m \mathbf{A}_j} = \mathbf{0} \quad m \neq j \quad (4.26)$$

$$\mathbf{F}_{total, \mathbf{A}_m \Psi} = \mathbf{F}_{m, \mathbf{A}_m \Psi} \quad (4.27)$$

$$\mathbf{F}_{total, \Psi \Psi} = \sum_m \mathbf{F}_{m, \Psi \Psi} \quad (4.28)$$

or in matrix form

$$\mathbf{F}_{total} = N \left[\begin{array}{ccc|c} \mathbf{F}_{1,A_1A_1} & \mathbf{0} & \dots & \mathbf{F}_{1,A_1\Psi} \\ \mathbf{0} & \mathbf{F}_{2,A_2A_2} & \dots & \mathbf{F}_{2,A_2\Psi} \\ \vdots & \ddots & \ddots & \vdots \\ \hline \mathbf{F}_{1,A_1\Psi}^H & \mathbf{F}_{2,A_2\Psi}^H & \dots & \sum_m^M \mathbf{F}_{m,\Psi\Psi} \end{array} \right]. \quad (4.29)$$

where the \mathbf{F}_m can be calculated using (4.14) as

$$\mathbf{F}_{m,A_mA_m} = \begin{bmatrix} \mathbf{J}_{A_m^m}^H \\ \mathbf{J}_{A_m^*}^H \end{bmatrix} (\mathbf{R}_m^{-T} \otimes \mathbf{R}_m^{-1}) [\mathbf{J}_{A_m} \quad \mathbf{J}_{A_m^*}] \quad (4.30)$$

$$\mathbf{F}_{m,A_m\Psi} = \begin{bmatrix} \mathbf{J}_{A_m^m}^H \\ \mathbf{J}_{A_m^*}^H \end{bmatrix} (\mathbf{R}_m^{-T} \otimes \mathbf{R}_m^{-1}) \mathbf{J}_\Psi \quad (4.31)$$

$$\mathbf{F}_{total,\Psi\Psi} = \mathbf{J}_\Psi^H \left[\sum_m^M (\mathbf{R}_m^{-T} \otimes \mathbf{R}_m^{-1}) \right] \mathbf{J}_\Psi \quad (4.32)$$

where

$$\mathbf{J}_\Psi = [\mathbf{S}_U \quad \mathbf{S}_L \quad (\mathbf{I}_P \circ \mathbf{I}_P)].$$

Because the constraint matrix is only a function of $\boldsymbol{\theta}_{A_m}$, the constraint matrix for the entire dataset becomes

$$\mathbf{H}_{total} = \begin{bmatrix} \mathbf{H}_{A_1} & \mathbf{0} & \dots & \mathbf{0} & & & \\ \mathbf{0} & \mathbf{H}_{A_2} & \dots & \vdots & & & \\ \vdots & \ddots & \ddots & \mathbf{0} & & & \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{H}_{A_m} & & & \\ & & & & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

where \mathbf{H}_{A_m} is given by (4.15). Because \mathbf{H}_{total} is very sparse and the block diagonal structure decouples the constraints, we can use efficient QR decomposition algorithms to find a unitary basis, \mathbf{U}_{total} for its null space efficiently.

The Constrained CRB can now be found, similar as before, using \mathbf{F}_{total} and \mathbf{U}_{total} . Using the matrix inversion lemma on the final result we obtain the expression for $\mathbf{C}_{\Psi\Psi}$ as

$$\mathbf{C}_{\Psi\Psi} = \frac{1}{N} \left[\sum_m \mathbf{F}_{m,\Psi\Psi} - \sum_m \mathbf{F}_{A_m\Psi}^H \mathbf{U}_{A_m} (\mathbf{U}_{A_m}^H \mathbf{F}_{A_mA_m} \mathbf{U}_{A_m})^{-1} \mathbf{U}_{A_m}^H \mathbf{F}_{A_m\Psi} \right]^{-1}. \quad (4.33)$$

This expression replaces the bound (5.16) and incorporates the benefits of jointly processing M FA models with a shared Ψ .

4.5. ESTIMATION ALGORITHMS FOR CLASSICAL FA

The classical FA problem was introduced in 1904 [53] and several algorithms were proposed [64–66], all for real data matrices (although readily extended to the complex case). In this section we will review some of these approaches.

4.5.1. AD-HOC METHOD

The estimation problem could be approached as a two stage minimization problem [67]. In this approach we try to minimize the LS cost function defined as

$$\min_{\mathbf{A}, \mathbf{D}} \|\hat{\mathbf{R}} - \mathbf{A}\mathbf{A}^H - \mathbf{D}\|_F^2 \quad (4.34)$$

where $\|\cdot\|_F$ is the Frobenius norm, in two stages, i.e. the Alternating LS approach. First for a given \mathbf{A} , (4.34) is minimized with respect to \mathbf{D} and in the next stage, \mathbf{D} is held constant and a new \mathbf{A} is found.

The iteration steps are

$$\hat{\mathbf{D}}_{(k+1)} = \text{diag}(\hat{\mathbf{R}} - \hat{\mathbf{A}}_{(k)}\hat{\mathbf{A}}_{(k)}^H) \quad (4.35)$$

$$\hat{\mathbf{A}}_{(k+1)} = \hat{\mathbf{D}}_{(k+1)}^{\frac{1}{2}} \mathbf{U}_{(k)} \mathbf{L}_{(k)}^{\frac{1}{2}} \quad (4.36)$$

where $\mathbf{L}_{(k)}$ and $\mathbf{U}_{(k)}$ are the Q largest eigenvalues and the corresponding eigenvectors of the matrix $\hat{\mathbf{D}}_{(k+1)}^{-1/2} \hat{\mathbf{R}} \hat{\mathbf{D}}_{(k+1)}^{-1/2} - \mathbf{I}_P$.

This method will monotonically approach a local minimum. Its rate of convergence is slow but it is robust against initialization and because of this, various hybrid approaches have been suggested in the literature. E.g., [68] suggests that at each iteration only the eigenvalues (and the corresponding eigenvectors) which are bigger than 1 be used to estimate \mathbf{A} . In this way instead of choosing the largest Q eigenvalues and eigenvectors, a set with $Q' \leq Q$ is chosen. This approach should accelerate the convergence [68].

4.5.2. MAXIMUM LIKELIHOOD ESTIMATOR

Using an ML approach, the aim is to find \mathbf{A} and \mathbf{D} that maximize the complex log-likelihood function

$$l(\hat{\mathbf{R}}; \mathbf{A}, \mathbf{D}) = N \left[-\log(\pi^P) + \log|\mathbf{R}^{-1}| - \text{tr}(\mathbf{R}^{-1}\hat{\mathbf{R}}) \right], \quad (4.37)$$

where $\mathbf{R} = \mathbf{A}\mathbf{A}^H + \mathbf{D}$ and $\hat{\mathbf{R}}$ is sample covariance matrix formed from N independent samples. To achieve this we find the gradient of this likelihood function (also known as the Fisher score) and set it equal to zero. The Fisher score for a proper Gaussian distributed signal is given by [39, p.165]

Using the results from Sec. 2.3.2 and 4.4, the vector form of the Fisher score can be written as

$$\mathbf{g}(\boldsymbol{\theta}) = N\mathbf{J}^H (\mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \quad (4.38)$$

where

$$\begin{aligned} \mathbf{J} &= \left[\frac{\partial \text{vect}(\mathbf{R})}{\partial \text{vect}^T(\mathbf{A})}, \frac{\partial \text{vect}(\mathbf{R})}{\partial \text{vect}^T(\mathbf{A}^*)}, \frac{\partial \text{vect}(\mathbf{R})}{\partial \text{vect}^T(\mathbf{D})} \right] \\ &= [\mathbf{J}_A, \mathbf{J}_{A^*}, \mathbf{J}_D]. \end{aligned} \quad (4.39)$$

Using these results and (4.38), the Fisher score in matrix form could be found to be

$$\begin{aligned} \text{unvect}(\mathbf{g}_A) &= N \text{unvect} \left[(\mathbf{A}^* \otimes \mathbf{I}_P)^H (\mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \right] \\ &= N \text{unvect} \left[(\mathbf{A}^T \mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \right] \\ &= N \mathbf{R}^{-1} (\hat{\mathbf{R}} - \mathbf{R}) \mathbf{R}^{-1} \mathbf{A}, \end{aligned} \quad (4.40)$$

and in the same way

$$\text{unvect}(\mathbf{g}_D) = N \text{diag}(\mathbf{R}^{-1}(\hat{\mathbf{R}} - \mathbf{R})\mathbf{R}^{-1}). \quad (4.41)$$

The ML technique requires us to set (4.40) and (4.41) equal to zero, unfortunately this will not produce a closed-form solution. As a result different iterative techniques such as the scoring method and EM based approaches have been suggested in the literature that we will discuss here.

THE SCORING METHOD

The scoring algorithm is a variant of the Newton–Raphson algorithm where the gradient and the Hessian are replaced by the Fisher score and Fisher information matrix respectively. The resulting iterative method is called the *method of scoring* or the *scoring method* [38, p.180]. The scoring iterations are:

$$\boldsymbol{\theta}_{(k+1)} = \boldsymbol{\theta}_{(k)} + \mu_{(k)} \Delta, \quad (4.42)$$

where

$$\Delta = \mathbf{F}_{(k)}^{\dagger} \mathbf{g}_{(k)}$$

is the direction of descent and

$$\mathbf{g} = \begin{bmatrix} \mathbf{g}_A \\ \mathbf{g}_{A^*} \\ \mathbf{g}_D \end{bmatrix}.$$

is the Fisher score.

By close inspection of the ML method described in [65], we observe that the method described there is an approximation of the scoring method where $\mathbf{F}_{(k)}^{\dagger}$ is approximated by $\text{diag}(\mathbf{F}_{(k)})^{-1}$. Here we suggest using a block-diagonal approximation of the Fisher information instead of a diagonal one which will allow us to solve the direction of descent for \mathbf{A} and \mathbf{D} separately. Using \mathbf{F}_{AA} and $\mathbf{F}_{\Psi\Psi}$ from the results in Sec. 4.4 we have

$$\begin{aligned} \Delta_A &= \mathbf{F}_{AA}^{-1}(\mathbf{A}^T \otimes \mathbf{I})(\mathbf{R}^{-T} \otimes \mathbf{R}^{-1})\text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \\ &= [(\mathbf{A}^T \otimes \mathbf{I})(\mathbf{R}^{-T} \otimes \mathbf{R}^{-1})(\mathbf{A}^* \otimes \mathbf{I})]^{-1}(\mathbf{A}^T \otimes \mathbf{I})(\mathbf{R}^{-T} \otimes \mathbf{R}^{-1})\text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \\ &= [(\mathbf{A}^T \mathbf{R}^{-T} \mathbf{A}^*) \otimes \mathbf{R}^{-1}]^{-1}(\mathbf{A}^T \otimes \mathbf{I})(\mathbf{R}^{-T} \otimes \mathbf{R}^{-1})\text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \\ &= [(\mathbf{A}^T \mathbf{R}^{-T} \mathbf{A}^*)^{-1} \otimes \mathbf{R}](\mathbf{A}^T \mathbf{R}^{-T} \otimes \mathbf{R}^{-1})\text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \\ &= [(\mathbf{A}^T \mathbf{R}^{-T} \mathbf{A}^*)^{-1} \mathbf{A}^T \mathbf{R}^{-T} \otimes \mathbf{I}]\text{vect}(\hat{\mathbf{R}} - \mathbf{R}) \\ &= \text{vect}[(\hat{\mathbf{R}} - \mathbf{R})\mathbf{R}^{-1}\mathbf{A}(\mathbf{A}^H \mathbf{R}^{-1}\mathbf{A})^{-1}] \end{aligned} \quad (4.43)$$

and thus

$$\mathbf{A}_{(k+1)} = \mathbf{A}_{(k)} + \mu_{(k)}(\hat{\mathbf{R}} - \mathbf{R}_{(k)})\mathbf{R}_{(k)}^{-1}\mathbf{A}_{(k)}(\mathbf{A}_{(k)}^H \mathbf{R}_{(k)}^{-1}\mathbf{A}_{(k)})^{-1}.$$

When we use the constraint that $\mathbf{A}^H \mathbf{D}^{-1} \mathbf{A} = \mathbf{\Gamma}$ we also have $\mathbf{A}^H \mathbf{R}^{-1} \mathbf{A} = \mathbf{\Gamma}(\mathbf{I}_Q + \mathbf{\Gamma})^{-1}$ is diagonal (this follows from (4.7)) and $(\mathbf{A}^H \mathbf{R}^{-1} \mathbf{A})^{-1} = \mathbf{\Gamma}^{-1}(\mathbf{I}_Q + \mathbf{\Gamma})$ and hence we can reduce the computations for the direction of descent for \mathbf{A} by using the following relation

$$\begin{aligned} \mathbf{R}^{-1}\mathbf{A}(\mathbf{A}^H \mathbf{R}^{-1}\mathbf{A})^{-1} &= \mathbf{D}^{-1}\mathbf{A}(\mathbf{I}_Q - (\mathbf{I}_Q + \mathbf{\Gamma})^{-1}\mathbf{\Gamma})\mathbf{\Gamma}^{-1}(\mathbf{I}_Q + \mathbf{\Gamma}) \\ &= \mathbf{D}^{-1}\mathbf{A}[\mathbf{\Gamma}^{-1}(\mathbf{I}_Q + \mathbf{\Gamma}) - \mathbf{I}_Q] \\ &= \mathbf{D}^{-1}\mathbf{A}\mathbf{\Gamma}^{-1}. \end{aligned} \quad (4.44)$$

Following the same procedure for \mathbf{D} we find

$$\mathbf{d}_{(k+1)} = \mathbf{d}_{(k)} + \mu_{(k)} (\mathbf{R}_{(k)}^{-T} \odot \mathbf{R}_{(k)}^{-1})^{-1} \text{vectdiag}[\mathbf{R}_{(k)}^{-1} (\hat{\mathbf{R}} - \mathbf{R}_{(k)}) \mathbf{R}_{(k)}^{-1}]$$

where we have used $\mathbf{F}_{\Psi\Psi} = \mathbf{F}_{\mathbf{D}\mathbf{D}} = \mathbf{R}^{-T} \odot \mathbf{R}$.

As the original approximation by [65] is equivalent to the technique better known as the Jacobi iterative approach to a linear system, the new proposed approach is a Block-Jacobi method. Further, if instead of $\mathbf{R}_{(k)}$ in the iteration for \mathbf{d} we use $\mathbf{A}_{(k+1)} \mathbf{A}_{(k+1)}^H + \mathbf{D}_{(k)}$ this method will become equivalent to a Block-Gauss-Seidel iteration [45].

As with any Newton based method, this approach requires a good initial point in order to converge to a desirable solution. The Ad-hoc method (see 4.5.1) can be used to provide an approximate solution as the initial guess for this algorithm.

EM BASED ALGORITHM

The expectation maximization technique is a general tool for finding an extremum of the likelihood function [38]. Its application to FA was first proposed by [69]. An overview of the original method and several of its derivatives can be found in [68]. Because of the simplicity of the final iterations it is straightforward to extend these algorithms to the complex domain. We will limit ourselves to a somewhat more recent algorithm which uses the Kullback–Leibler divergence (KLD) in order to find the ML solution. In the case of Gaussian distributed families the KLD will reduce to a likelihood ratio function which has been used for ML estimation by [56].

For a given set of samples maximizing the likelihood function is the same as minimizing the likelihood ratio function

$$F(\mathbf{A}, \mathbf{D}) = \text{tr}(\mathbf{R}^{-1} \hat{\mathbf{R}}) - \log |\mathbf{R}^{-1} \hat{\mathbf{R}}| - P \quad (4.45)$$

which is the same as the KLD when the divergence is calculated between a Gaussian distribution described by $\hat{\mathbf{R}}$ and another Gaussian distribution parametrized by \mathbf{A} and \mathbf{D} . Even though both functions are the same, an information geometric perspective on the KLD provides some easy to implement algorithms which avoids taking derivatives and converges to a local minimum [70]. The following iterations summarize the KLD FA [70]: Let

$$\Phi(\mathbf{A}, \mathbf{D}) = \mathbf{I}_Q - \mathbf{A}^H \mathbf{R}^{-1} \mathbf{A} + \mathbf{A}^H \mathbf{R}^{-1} \hat{\mathbf{R}} \mathbf{R}^{-1} \mathbf{A},$$

then

$$\hat{\mathbf{A}}_{(k+1)} = \hat{\mathbf{R}} \mathbf{R}_{(k)}^{-1} \hat{\mathbf{A}}_{(k)} \Phi(\hat{\mathbf{A}}_{(k)}, \hat{\mathbf{D}}_{(k)})^{-1} \quad (4.46)$$

$$\hat{\mathbf{D}}_{(k+1)} = \text{diag}(\hat{\mathbf{R}} - \hat{\mathbf{A}}_{(k+1)} \hat{\mathbf{A}}_{(k+1)}^H \mathbf{R}_{(k)}^{-1} \hat{\mathbf{R}}). \quad (4.47)$$

We will evaluate the performance of this technique using simulations (see Sec. 4.8).

4.5.3. MINIMUM TRACE FA

Minimum Trace FA tries to find the unknown parameters by formulating the problem as a minimization of a trace function. The matrix for which the trace is minimized differentiates different approaches. We will distinguish between two classes of minimum trace FA (MTFA), one that is equivalent to Weighted Least Squares (WLS) and one related to

minimum rank FA (MRFA). The WLS for classical FA is a special case of the non-linear WLS algorithm that will be discussed in Sec. 4.6.1 and is not addressed here.

Revisiting the FA problem, ideally we would like to find $\mathbf{R}_0 = \mathbf{A}\mathbf{A}^H$ and \mathbf{D} directly by solving

$$\begin{aligned} \min_{\mathbf{D}} \quad & \text{rank}(\mathbf{R} - \mathbf{D}) \\ \text{s.t.} \quad & \mathbf{R} - \mathbf{D} \geq \mathbf{0} \\ & \mathbf{D} > \mathbf{0}. \end{aligned}$$

However there are two major issues to this direct approach. The first issue follows from the fact that minimizing the rank is an NP-hard problem [71] and second problem follows from the fact that we only have access to $\hat{\mathbf{R}}$ which puts a relatively high lower bound on $\text{rank}(\hat{\mathbf{R}} - \mathbf{D})$ [64] (i.e. the reduction in the rank of $\hat{\mathbf{R}}$ is very limited when we just modify its diagonal elements). MRFA has been suggested very early in the development of factor analysis by Ledermann [64]. Ledermann suggested using $\text{tr}(\mathbf{R} - \mathbf{D})$ as the objective function to avoid the NP-hardness of the rank function. This problem has been later revisited by Shapiro [72, 73] and most recently by [74] where the following optimization problem is proposed:

$$\begin{aligned} \min_{\mathbf{R}_0} \quad & \text{tr}(\mathbf{R}_0) \\ \text{s.t.} \quad & \mathbf{R}_0 \geq \mathbf{0} \\ & \mathbf{R}_0 + \mathbf{D} = \mathbf{R} \\ & \mathbf{D} > \mathbf{0}. \end{aligned}$$

Using the convexity of the trace function it can be shown that the solution to the problem as stated above provides the correct FA model with high probability if $Q < P/2$ and if there is access to \mathbf{R} [74]. No results are known when $\hat{\mathbf{R}}$ is used, however this approach seems to be applicable in practice when we replace the equality constraint $\mathbf{R}_0 + \mathbf{D} = \mathbf{R}$ with $\|\hat{\mathbf{R}} - \mathbf{R}_0 - \mathbf{D}\|_F \leq \epsilon$. The recovered rank for \mathbf{R}_0 depends heavily on the choice of ϵ . The rank recovered by MTFA can be used as an initial guess for the sequential detection described in Sec. 4.7.

From empirical research done on applying each of these methods extensively, use of the ad-hoc or KLD method to form an initial guess followed by the scoring method appears to be the most robust way to find the model parameters. For more results see Sec. 4.8.

4.6. ESTIMATION ALGORITHMS FOR JEFA

So far we have only discussed existing estimation techniques for the classical FA and minor changes that are needed when we extend the model to complex-valued data. In this section we will focus on developing estimation techniques for JEFA.

Estimating the parameters for JEFA leads to a non-linear optimization problem that we will solve using a variation of a Jacobian-free Newton-Krylov (JFNK) technique [75] and a matrix-free Gauss-Newton-Krylov (MFGNK) [76]. The main idea behind the Newton-Krylov technique is to solve the linear system needed to find the direction of descent using Krylov subspace based solvers. As discussed in Sec. 2.4, Krylov-subspace based algorithms find the solution to a linear system such as $\mathbf{B}\mathbf{x} = \mathbf{y}$ by repeated calculation of the matrix vector products of the form $\mathbf{B}\mathbf{v}$. In many applications, and as we will demonstrate in our case, \mathbf{B} is related to the Jacobians and the multiplications can be performed using

these Jacobians. The JFNK and MFGNK techniques avoid storing the Jacobian by using a Taylor expansion to approximate the needed matrix vector products [75, 76]. The Kronecker structure of the Jacobians derived in the previous sections allows us to develop a method that also avoids storing the Jacobians but does exact computation of the matrix vector product.

We will discuss Non-linear Weighted Least Squares (NLWLS) and the Maximum Likelihood (ML) for finding $\hat{\Psi}$ and $\hat{\mathbf{A}}_m$.

4.6.1. NON-LINEAR WEIGHTED LEAST SQUARES

We start by vectoring and stacking all the covariance matrices to form a measurement vector

$$\hat{\mathbf{r}} = [\text{vect}^T(\hat{\mathbf{R}}_1) \quad \dots \quad \text{vect}^T(\hat{\mathbf{R}}_m)]^T, \quad (4.48)$$

and similarly

$$\mathbf{r}(\boldsymbol{\theta}) = [\text{vect}^T(\mathbf{R}_1(\boldsymbol{\theta})) \quad \dots \quad \text{vect}^T(\mathbf{R}_m(\boldsymbol{\theta}))]^T, \quad (4.49)$$

where

$$\boldsymbol{\theta} = [\boldsymbol{\theta}_{\mathbf{A}_1}^T \quad \dots \quad \boldsymbol{\theta}_{\mathbf{A}_m}^T \quad \boldsymbol{\theta}_{\Psi}^T]^T. \quad (4.50)$$

We can estimate the unknown parameters in $\boldsymbol{\theta}$ using NLWLS defined as

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \|\mathbf{W}^{1/2} [\hat{\mathbf{r}} - \mathbf{r}(\boldsymbol{\theta})]\|_2^2 \quad (4.51)$$

where \mathbf{W} is a weighting matrix. The optimum weighting matrix is the covariance matrix of the entire vectorized dataset, $\mathbf{W} = \text{bdiag}(\mathbf{R}_m^{-T} \otimes \mathbf{R}_m^{-1})$, however because we have only access to $\hat{\mathbf{R}}_m$ and we use

$$\mathbf{W} = \begin{bmatrix} \hat{\mathbf{R}}_1^{-T} \otimes \hat{\mathbf{R}}_1^{-1} & & \\ & \ddots & \\ & & \hat{\mathbf{R}}_M^{-T} \otimes \hat{\mathbf{R}}_M^{-1} \end{bmatrix} \quad (4.52)$$

which will give an asymptotically optimal solution for a Gaussian distributed data matrix [37]. We would like to note here that the WLS cost function (4.51) with weight (4.52) and $M = 1$ can be formulated as a minimum trace problem

$$\begin{aligned} \|\mathbf{W}^{1/2} [\hat{\mathbf{r}} - \mathbf{r}(\boldsymbol{\theta})]\|_2^2 &= \|\mathbf{R}^{1/2} (\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})) \mathbf{R}^{1/2}\|_F^2 \\ &= \text{tr}\{[\hat{\mathbf{R}}^{-1} (\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta}))]^2\} \\ &= \text{tr}[(\mathbf{I} - \hat{\mathbf{R}}^{-1} \mathbf{R}(\boldsymbol{\theta}))^2] \end{aligned}$$

and the real version of this cost function for classical FA has been proposed by [66].

A very common iterative technique for solving nonlinear optimization problems is the Gauss-Newton algorithm, where the Hessian is replaced by the Gramian of the Jacobians [41]. The updates are similar to Newton updates and are given by

$$\hat{\boldsymbol{\theta}}_{(k+1)} = \hat{\boldsymbol{\theta}}_{(k)} + \mu_{(k)} \boldsymbol{\Delta} \quad (4.53)$$

where Δ is the direction of descent. For $M = 1$ and $\Psi = \mathbf{D}$, $\hat{\theta}$ is the same as $\hat{\theta}$ used for classical FA. To find Δ we need to solve

$$\mathbf{B}(\theta)\Delta = \mathbf{g}(\theta) \quad (4.54)$$

where

$$\mathbf{B}(\theta) = \mathbf{J}^H(\theta)\mathbf{W}\mathbf{J}(\theta) \quad (4.55)$$

is the Gramian and $\mathbf{g}(\theta)$ is the gradient of the NLWLS given by

$$\mathbf{g}(\theta) = \mathbf{J}^H(\theta)\mathbf{W}[\hat{\mathbf{r}} - \mathbf{r}(\theta)]. \quad (4.56)$$

We will continue the iterations given by (4.53) until $\|\mathbf{g}(\hat{\theta}_{(k)})\|_2 < \epsilon$ where $\epsilon > 0$ depends on the desired accuracy. This concludes the Gauss-Newton algorithm. The key step is solving the linear system in (4.54). We will discuss a Krylov based method for solving this system to find Δ .

4.6.2. KRYLOV-BASED METHODS

There are many approaches to solving a large system of linear equations. The choice of the solver depends heavily on the structure the matrix involved. For example if the matrix has a Toeplitz or circulant structure, fast methods based on FFT can be applied. In our case the Gramian matrix $\mathbf{B}(\theta)$ has and $\mathbf{J}(\theta)$ poses Kronecker structure. For these matrices we propose using a Krylov subspace based solver and demonstrate how it can reduce storage and/or complexity.

There are various Krylov subspace based solvers, an overview of these solvers can be found for example in [77]. We know from our study of the Fisher information that the solution to the problem is not unique, this means that the Jacobians and hence \mathbf{B} are singular. One possible Krylov solver that is capable of finding a solution for singular matrices is the MinresQLP algorithm [78] and for this reason we have chosen this solver for our iterative approach.

MinresQLP is a standard iterative solver that we do not discuss in detail here. The solver requires the availability of a subroutine that can perform a matrix vector multiplication of the form $\mathbf{u} = \mathbf{B}\mathbf{v}$. Other operations in MinresQLP have negligible complexity compared to this matrix–vector multiplication. We will show how we can perform this multiplication without needing to store the Jacobians using the Kronecker structure.

We will drop the dependency on θ from the notation and write only \mathbf{J} and \mathbf{r} because θ does not change while we are solving for Δ . In order to calculate $\mathbf{u} = \mathbf{B}\mathbf{v}$ for \mathbf{B} given in (4.55), we define an intermediate result $\mathbf{z} = \mathbf{J}\mathbf{v}$. Given the fact that \mathbf{v} , \mathbf{u} have the same dimensions as θ and \mathbf{z} has the same dimensions as \mathbf{r} we are going to partition them in the same manner using (4.50), (4.21) and (4.20) as

$$\mathbf{v} = \begin{bmatrix} \text{vect}(\mathbf{V}_{\mathbf{A}_1}) \\ \text{vect}(\mathbf{V}_{\mathbf{A}_1^*}) \\ \vdots \\ \mathbf{S}_U^T \text{vect}(\mathbf{V}_{\Psi}) \\ \mathbf{S}_L^T \text{vect}(\mathbf{V}_{\Psi}) \\ (\mathbf{I}_P \circ \mathbf{I}_P)^T \text{vect}(\mathbf{V}_{\Psi}) \end{bmatrix}, \quad \mathbf{u} = \begin{bmatrix} \text{vect}(\mathbf{U}_{\mathbf{A}_1}) \\ \text{vect}(\mathbf{U}_{\mathbf{A}_1^*}) \\ \vdots \\ \mathbf{S}_U^T \text{vect}(\mathbf{U}_{\Psi}) \\ \mathbf{S}_L^T \text{vect}(\mathbf{U}_{\Psi}) \\ (\mathbf{I}_P \circ \mathbf{I}_P)^T \text{vect}(\mathbf{U}_{\Psi}) \end{bmatrix} \quad (4.57)$$

and

$$\mathbf{z} = \begin{bmatrix} \text{vect}(\mathbf{Z}_1) \\ \vdots \\ \text{vect}(\mathbf{Z}_m) \end{bmatrix} \quad (4.58)$$

where \mathbf{S}_L and \mathbf{S}_U are selection matrices based on the mask matrix \mathbf{M} . To find \mathbf{u} we will compute $\mathbf{U}_{\mathbf{A}_m}$, $\mathbf{U}_{\mathbf{A}_m}^*$ and \mathbf{U}_{Ψ} . We note that if $\mathbf{V}_{\mathbf{A}_m}^* = \mathbf{V}_{\mathbf{A}_m}^*$ then $\mathbf{U}_{\mathbf{A}_m}^* = \mathbf{U}_{\mathbf{A}_m}^*$ which means that only $\mathbf{U}_{\mathbf{A}_m}$ needs to be calculated.

The Jacobian for the entire dataset is given by

$$\mathbf{J} = \frac{\partial \text{vect}(\mathbf{R})}{\partial \boldsymbol{\theta}^T} = \begin{bmatrix} \mathbf{J}_{\mathbf{A}_1} & \mathbf{J}_{\mathbf{A}_1}^* & \dots & \mathbf{0} & \mathbf{J}_{\Psi} \\ \mathbf{0} & & \dots & \mathbf{0} & \mathbf{J}_{\Psi} \\ \mathbf{0} & \ddots & \ddots & \mathbf{0} & \mathbf{J}_{\Psi} \\ \mathbf{0} & \dots & \mathbf{J}_{\mathbf{A}_m} & \mathbf{J}_{\mathbf{A}_m}^* & \mathbf{J}_{\Psi} \end{bmatrix} \quad (4.59)$$

and hence using $\mathbf{z} = \mathbf{J}\mathbf{v}$, (4.58) and (4.14):

$$\begin{aligned} \text{vect}(\mathbf{Z}_m) &= (\mathbf{A}_m^* \otimes \mathbf{I}_P) \text{vect}(\mathbf{V}_{\mathbf{A}_m}) + (\mathbf{I}_P \otimes \mathbf{A}_m) \mathbf{K}^{PQ} \text{vect}(\mathbf{V}_{\mathbf{A}_m}^*) + \text{vect}(\mathbf{V}_{\Psi}) \\ &= \text{vect}(\mathbf{V}_{\mathbf{A}_m} \mathbf{A}_m^H + \mathbf{A}_m \mathbf{V}_{\mathbf{A}_m}^H + \mathbf{M} \odot \mathbf{V}_{\Psi}) \end{aligned}$$

where we have used $\mathbf{V}_{\mathbf{A}_m}^* = \mathbf{V}_{\mathbf{A}_m}^*$. It follows directly from unvectorizing both sides that

$$\mathbf{Z}_m = \mathbf{V}_{\mathbf{A}_m} \mathbf{A}_m^H + \mathbf{A}_m \mathbf{V}_{\mathbf{A}_m}^H + \mathbf{M} \odot \mathbf{V}_{\Psi}. \quad (4.60)$$

This means that we can calculate $\mathbf{J}\mathbf{v}$ by only reshaping the vector \mathbf{v} to appropriate matrices $\mathbf{V}_{\mathbf{A}_m}$ and \mathbf{V}_{Ψ} and applying (4.60). The variables \mathbf{A}_m are the current estimates of unknown parameters and hence require no additional storage.

The next matrix vector multiplications that we need is $\mathbf{z}_W = \mathbf{W}\mathbf{z}$. Using the properties of Kronecker products it is straightforward to show that \mathbf{z}_W can be calculated using

$$\mathbf{Z}_{W_m} = \hat{\mathbf{R}}_m^{-1} \mathbf{Z}_m \hat{\mathbf{R}}_m^{-1} \quad (4.61)$$

and

$$\mathbf{z}_W = \begin{bmatrix} \text{vect}(\mathbf{Z}_{W_1}) \\ \vdots \\ \text{vect}(\mathbf{Z}_{W_m}) \end{bmatrix}. \quad (4.62)$$

The final product we need to calculate is $\mathbf{u} = \mathbf{J}^H \mathbf{z}_W$. From the structure of (4.59), we see that

$$\begin{aligned} \text{vect}(\mathbf{U}_{\mathbf{A}_m}) &= \mathbf{J}_{\mathbf{A}_m}^H \text{vect}(\mathbf{Z}_{W_m}) \\ \text{vect}(\mathbf{U}_{\mathbf{A}_m}^*) &= \mathbf{J}_{\mathbf{A}_m}^H \text{vect}(\mathbf{Z}_{W_m}). \end{aligned}$$

Unvectorizing both sides and applying (4.14) we find

$$\begin{aligned} \mathbf{U}_{\mathbf{A}_m} &= \mathbf{Z}_{W_m} \mathbf{A}_m \\ \mathbf{U}_{\mathbf{A}_m}^* &= \mathbf{Z}_{W_m}^T \mathbf{A}_m^*. \end{aligned} \quad (4.63)$$

The remaining term \mathbf{U}_Ψ is given by:

$$\mathbf{U}_\Psi = \sum_{m=1}^M \mathbf{M} \odot \mathbf{Z}_{\mathbf{W}_m}. \quad (4.64)$$

To summarize, in order to calculate $\mathbf{B}\mathbf{v}$ we reshape \mathbf{v} into $\mathbf{V}_{\mathbf{A}_m}$ and \mathbf{V}_Ψ and use (4.60), (4.61), (4.63) and (4.64) to find the result. The gradient \mathbf{g} can be calculated in a similar manner by using $\mathbf{Z}_m = \hat{\mathbf{R}}_m - \mathbf{R}_m$. The procedure that performs these steps is provided to MinresQLP which then solves for Δ .

By assuming $\mathbf{V}_{\mathbf{A}_m}^* = \mathbf{V}_{\mathbf{A}_m}^*$ we showed in (4.60) that \mathbf{Z}_m is Hermitian, and because $\hat{\mathbf{R}}_m^{-1}$ is Hermitian so is $\mathbf{Z}_{\mathbf{W}_m}$. From the properties of Hermitian matrices we have $\mathbf{Z}_{\mathbf{W}_m}^T = \mathbf{Z}_{\mathbf{W}_m}^*$ and thus $\mathbf{U}_{\mathbf{A}_m}^* = \mathbf{U}_{\mathbf{A}_m}^*$. We still need to show that the assumption about $\mathbf{V}_{\mathbf{A}_m}^*$ is valid. It can be shown that MinresQLP provides \mathbf{v} that has the property $\mathbf{V}_{\mathbf{A}_m}^* = \mathbf{V}_{\mathbf{A}_m}^*$ when solving $\mathbf{B}\Delta = \mathbf{g}$ if \mathbf{g} has this property. Calculating \mathbf{g} is achieved by setting $\mathbf{Z}_m = \hat{\mathbf{R}}_m - \mathbf{R}_m$ and following the procedure above. Because both $\hat{\mathbf{R}}_m$ and \mathbf{R}_m are Hermitian, it follows that the needed property holds for \mathbf{g} and hence for \mathbf{v} .

It is also worth noting that when $\mathbf{W} = \mathbf{I}$ we have $\mathbf{Z}_{\mathbf{W}_m}\mathbf{A}_m = \mathbf{Z}_m\mathbf{A}_m$. This allows us to compute $\mathbf{U}_{\mathbf{A}_m}$ using $\mathbf{A}_m(\mathbf{V}_{\mathbf{A}_m}^H\mathbf{A}_m)$ and $\mathbf{V}_{\mathbf{A}_m}(\mathbf{A}_m^H\mathbf{A}_m)$ which has a complexity of $O(PQ^2)$ instead of the original $O(P^2Q)$. Also the term $\mathbf{A}_m^H\mathbf{A}_m$ can be calculated outside the MinresQLP iterations, saving more computations.

4.6.3. MAXIMUM LIKELIHOOD FOR JEFA

An alternative to the NLWLS solved in Sec. 4.6.1 is to maximize the likelihood function (4.19). A Hessian-free (similar to Jacobian and matrix free) variation in combination with Krylov solvers has been suggested in [79]. However by using the results of the previous section we will develop a method based on the scoring method, where the Hessian is replaced by the Fisher information matrix [38]. We already have derived the Fisher information for the entire dataset. We will denote this method Scoring-Krylov (SK).

The scoring method as presented by [38] can be summarized as

$$\hat{\boldsymbol{\theta}}_{(k+1)} = \hat{\boldsymbol{\theta}}_{(k)} + \mu_{(k)}\Delta$$

where Δ is the solution to

$$\mathbf{F}_{total}(\boldsymbol{\theta})\Delta = \mathbf{g}(\boldsymbol{\theta})$$

where \mathbf{g} is the Fisher score. Using the same technique as we have done to find the derivatives for NLWLS we find the derivative of the likelihood as

$$\mathbf{g}(\boldsymbol{\theta}) = \mathbf{J}^H(\boldsymbol{\theta})\mathbf{W}(\boldsymbol{\theta})(\hat{\mathbf{r}} - \mathbf{r}(\boldsymbol{\theta}))$$

where $\mathbf{W}(\boldsymbol{\theta}) = \text{bdiag}(\mathbf{R}_m^{-T}(\boldsymbol{\theta}) \otimes \mathbf{R}_m^{-1}(\boldsymbol{\theta}))$.

Given the definition of the gradient and the Fisher information, we observe that the same techniques used for NLWLS can be applied to ML. We only need to replace the weighting matrix to be the current best estimate of $\mathbf{R}(\boldsymbol{\theta})$. However, this approach requires the inversion of the covariance matrices at each iteration. If based on the structure of the masking matrix \mathbf{M} , inversion of Ψ is computationally more accommodating, then the Woodbury matrix identity can be used to find the inverse with less complexity.

When using MinresQLP for the SK technique the only step that needs modification is (4.61), where we replace $\hat{\mathbf{R}}_m$ with $\mathbf{R}_m(\hat{\mathbf{A}}_{m,(k)}, \hat{\mathbf{\Psi}}_{(k)})$ (note that k remains the same during the MinresQLP iterations).

This will allow us to find the ML solution using the advantages of the Krylov based solvers without the need to store the Fisher information matrix or the Jacobians.

4.7. GOODNESS OF FIT AND DETECTION

One of the parameters that needs to be found for the FA model is the number of underlying common factors, Q (i.e. $\text{rank}(\mathbf{A})$). In the concept of array processing it is the same as detecting the number of sources that the array is exposed to.

Finding the true value of Q is part of multiple hypothesis testing and model order selection problem and there exists an extensive literature on this subject (an overview can be found in [80, pp. 222-223][81]). In this section we will limit the discussion to a general likelihood ratio test (GLRT) which is used to decide whether the FA model fits a given sample covariance matrix. If the model fits, it can be seen as detecting Q sources.

Two hypotheses are tested against each other. H_0 assumes that there is an FA model underlying the data, while H_1 assumes no structure. Consider the following test

$$\zeta = \frac{L_1^*}{L_0^*} \geq \gamma$$

where L_1^* is the maximum value of the likelihood when H_1 is true, and L_0^* is the maximum value of the likelihood for a FA model. Taking the natural logarithm from both sides we see that it reduces to (4.45) and we can write

$$\begin{aligned} \lambda &= 2 \log(\zeta) \\ &= 2N [\text{tr}(\mathbf{R}^{-1} \hat{\mathbf{R}}) - \log |\mathbf{R}^{-1} \hat{\mathbf{R}}| - P]. \end{aligned} \quad (4.65)$$

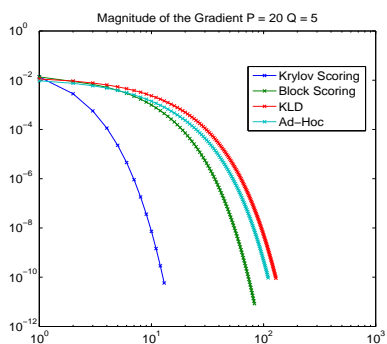
From [67, p.267][61, p.281] we know that this statistic has an asymptotic χ_s^2 distribution under H_0 where s is the degree of freedom defined by (4.66) later in Sec. 4.10 (appendix to this chapter). We can use this statistic to find a false alarm ratio detector. In the special case where $Q = 0$ this test indicates whether there are any sources active during the measurement.

If the GLRT passes for a given estimate \hat{Q}_0 it will also pass for any $\hat{Q} > \hat{Q}_0$ and if it fails, it also fails for any $\hat{Q} < \hat{Q}_0$. Therefore, instead of a linear search for \hat{Q} we propose to use a binary search. In this case the number of needed FA estimates is on average $\log_2(Q_{\max}) + 1$ where Q_{\max} is the maximum number of possible sources for FA given by (4.67) and follows the relation $Q_{\max} < P - \sqrt{P}$.


4.8. SIMULATIONS

We will evaluate the performance of the proposed models and algorithms using a series of simulations. We will start by evaluating the convergence speed of different algorithms in Sec. 4.8.1, we will then evaluate the quality of the estimated subspace using classical and Joint FA in Sec. 4.8.2, in Sec. 4.8.3 we use simulation to show how JEFA can be used to improve DOA estimation and finally in Sec. 4.8.4 we show that the proposed algorithm for JEFA converges to the CRB as number of samples becomes larger.


4




(a) Convergence speed no finite sample noise for $Q = 5$




(b) Convergence speed no finite sample noise for $Q = 12$



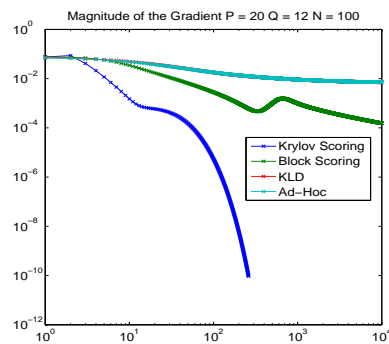
(c) Convergence speed $N = 1000, Q = 5$



(d) Convergence speed $N = 1000, Q = 12$



(e) Convergence speed $N = 100, Q = 5$



(f) Convergence speed $N = 100, Q = 12$

Figure 4.1: Convergence for $P = 20$. Vertical axis is the magnitude of the gradient and the horizontal axis is the number of iterations.

Table 4.1: Complexity of each iteration for different algorithms

Approach	flops (order)
Ad-Hoc	$P^3 + P^2Q$
KLD	$P^2Q + PQ^2 + Q^3$
Krylov NLLS	$P^2Q + I_K(PQ^2)$
Krylov NLWLS	$P^2Q + I_K(P^2Q + PQ^2)$
Joint Krylov NLWLS	$P^2 \sum_m^M Q_m + I_K(P^2 \sum_m^M Q_m + P \sum_m^M Q_m^2)$
Krylov Scoring	$P^2 \sum_m^M Q_m + \sum_m^M Q_m^3 + I_K(P^2 \sum_m^M Q_m + P \sum_m^M Q_m^2)$

4.8.1. CONVERGENCE SPEED


In this section we will evaluate the speed of convergence for different algorithms using simulations. Table 4.1 shows the complexity for a single iteration of the algorithms listed. I_K is the number of iterations needed for the Krylov solver to converge. This number can be chosen to be very small depending on how much improvement is desired with respect to the descend direction provided by the gradient. In the simulations presented here we allow the solver to fully converge based on the default convergence criteria of MinresQLP.

Using different simulations we illustrate that not only the complexity but also the convergence speed of the different algorithms depends heavily on Q and the level of finite sample noise. An array with $P = 20$ elements is simulated. The matrix \mathbf{A} is chosen randomly with a standard complex Gaussian distribution (i.e. each element is distributed as $\mathcal{N}(0, 1)$) and \mathbf{D} is chosen randomly with a uniform distribution between 1 and 5. The same initial point is chosen for all the algorithms.


Fig. 4.1 shows the convergence rate of different ML algorithms based on the magnitude of the gradient. For the simulated array the maximum number of sources is given by $Q_{max} = 15$. Hence the simulation results illustrated in the first column of Fig. 4.1 with $Q = 5$ are representative for low rank cases and the second column with $Q = 12$ for high rank cases of the classical FA model.

Remarks:

- The Krylov Scoring outperforms all the other presented algorithms based on the number of iterations needed for convergence.
- For high Q a low number of samples (Fig. 4.1e) can be catastrophic for most of the algorithms.
- The Block Scoring algorithm also performs well with respect to the ad-hoc and KLD, however the number of iterations is still very large for large Q and low number of samples.
- Scoring, Block Scoring and NLWLS based algorithms are sensitive to the initial guess, however in practice the result from a few steps of the KLD algorithm (e.g. 10) has been observed to be a reliable initial guess.



(a) Attenuation as the function of SNR for different techniques.



(b) Angle difference between the estimated subspace and true subspace.

Figure 4.2: Subspace estimation performance of EVD, FA and JFA for various SNR and white noise.

4.8.2. SUBSPACE ESTIMATION PERFORMANCE

In this section we study the performance of FA and JFA where we take $\Psi = \sigma^2 \mathbf{I}_P$, so we can compare the performance to EVD.

For this simulations we have chosen $Q_m = 2$, $P = 5$, $M = 5$ and $\sigma = 1$ is the noise power. We study the subspace estimation performance for various signal to noise ratios (SNR) ranging from -5 dB to 20 dB. Each sample covariance matrix is generated using $N = 100$ samples and \mathbf{A}_m is generated as a random complex matrix.

Two metrics are used to measure performance of the estimated subspace. We use the estimated subspaces to find a projection matrix into the null-space of $\hat{\mathbf{A}}_m$ which we will denote by $\hat{\mathbf{P}}_m$ and we measure

$$\text{Subspace error} = \frac{\|\hat{\mathbf{P}}_m \mathbf{A} \mathbf{A}^H \hat{\mathbf{P}}_m\|_F}{\|\mathbf{A} \mathbf{A}^H\|_F}.$$

In Fig. 4.2a the result of this simulations is presented. As FA and JFA have to estimate more parameters, we expect a drop in performance compared to EVD. This simulation shows that this occurs for FA at low SNR. JFA exploits the stationarity of the noise component and hence has a quite small performance penalty with respect to EVD.

The other metric we use is the angle between two subspaces calculated using MATLAB command *subspace*. This result is shown in Fig. 4.2b. The subspace angle difference between the true and estimated subspaces decreases as SNR increases. JFA follows the performance of EVD with a very small gap.

Because JFA has a more general model, it is applicable in many practical situations and we have shown that applying this technique in classical scenarios where $\Psi = \sigma^2 \mathbf{I}_P$ does not result in a significant performance loss. I.e. the accuracy of the subspace estimated using JFA is very close to that of EVD even though for a model with $\Psi = \sigma^2 \mathbf{I}_P$ the number of unknowns in JFA model is larger.




Figure 4.3: Result of matched filter beamforming as a function of time index k with and without interference.

4

4.8.3. DOA ESTIMATION USING JEFA

In this scenario we use the estimated subspace from EVD, EJFA and EFA as the input to a DOA estimator based on ESPRIT [59]. We simulate $M = 10$ with $Q_m = 2$ targets moving along the tracks T_1 and T_2 as illustrated in Fig. 4.4 between the snapshots. We have a uniform linear array with $P = 7$ receivers that observe the targets, however $P_0 = 5$ of these receivers are contaminated with unknown interfering signals. We will model the interfering signal as a stationary unknown colored noise which leads to a mask matrix defined by

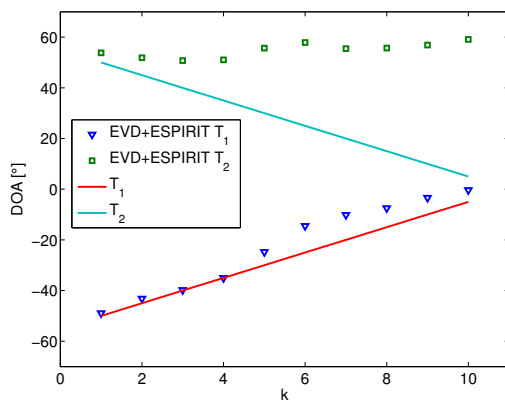
$$\mathbf{M} = \begin{bmatrix} \mathbf{1}_{P_0} \mathbf{1}_{P_0}^T & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_2 \end{bmatrix}$$

where $\mathbf{1}_{P_0}$ is a $P_0 \times 1$ vector with all entries equal to unity. The sample covariance matrix for each snapshot is obtained using $N = 100$ samples.


Fig. 4.3a shows the result of matched filter beamforming on the simulated data when there is no interferer present and Fig. 4.3b shows the effect of the interfering signals. Because of the limited resolution of the device, the beamformer cannot differentiate the two signals in the last snapshot.

We present the Monte-Carlo (MC) results of ESPRIT for each snapshot based on the subspace estimated by various algorithms in Fig. 4.4.


- Because $\Psi \neq \sigma^2 \mathbf{I}$, EVD is not able to recover the correct subspace and hence as illustrated in Fig. 4.4a the estimated subspace is biased, (note that the bias is different between each snapshot and is a function of both \mathbf{A}_m and Ψ and the wild behavior shown in this figure does not disappear by increasing the number of MC runs).
- Fig. 4.4b shows the result obtained by applying the EFA separately on each snapshot followed by ESPRIT. Because the resolution decreases for higher angles (as can be seen in Fig. 4.3a) and because not the entire dataset is used the variance of the DOA estimates is higher for the first few snapshots, also as the targets get closer it is more difficult to differentiate their subspace. Both problems affect the performance of EFA.



(a) DOA results for EVD + ESPRIT



(b) DOA results for EFA + ESPRIT



(c) DOA results for JEFA + ESPRIT

Figure 4.4: Comparison of subspace based DOA estimation for EVD, EFA and JEFA.




Figure 4.5: Performance of the diagonal estimates compared to the CRB

- The performance of JEFA is illustrated in Fig. 4.4c. Because both the correct data model has been used and estimation is done over the entire dataset, JEFA is able to recover the subspaces and hence the DOA estimates more accurately.

4.8.4. CRAMÉR-RAO BOUND SIMULATION

In this part we investigate the performance of the proposed algorithm using the Cramér-Rao bound. We use a setup with $P = 5$, $Q_m = 2$, $\Psi = \mathbf{D}$ with diagonal elements ranging from 0.5 to 1.5. Two different approaches are compared. The first approach is to apply FA separately and then use $\hat{\mathbf{D}} = 1/K \sum_m \hat{\mathbf{D}}_m$. The other approach is to estimate using JFA.

We use

$$\begin{aligned} \mathcal{E}\{\|\hat{\mathbf{D}} - \mathbf{D}\|_F^2\} &= \mathcal{E}\{\text{vect}(\hat{\mathbf{D}} - \mathbf{D})^H \text{vect}(\hat{\mathbf{D}} - \mathbf{D})\} \\ &= \text{tr}[\mathcal{E}\{\text{vect}(\hat{\mathbf{D}} - \mathbf{D}) \text{vect}(\hat{\mathbf{D}} - \mathbf{D})^H\}] \geq \text{tr}(\mathbf{C}_\Psi) \end{aligned}$$

where \mathbf{C}_Ψ is the sub-matrix of CRB corresponding to Ψ , to measure performance. We estimate $\mathcal{E}\{\|\hat{\mathbf{D}} - \mathbf{D}\|_F^2\}$ using Monte Carlo simulations. Fig. 4.5 shows the result of this simulation. This figure clearly illustrates that the proposed joint estimation reaches the CRB asymptotically and that applying the estimation separately followed by an averaging results in a sub-optimal estimation with higher variance.

4.9. CONCLUSION

We have provided a method for jointly estimating the non-stationary low-rank and stationary structured part of a series of covariance matrices by developing efficient algorithms based on Newton-Krylov optimization techniques. An algorithm to find the ML estimates has also been presented.

The Cramér-Rao bound for the entire dataset has been provided and the performance of the algorithm have been illustrated using simulations.

The general structure of the data model should make application of this technique possible in a wide range of signal processing applications. In Chapter 5 we will use the results in this section extensively to spatial filtering of RFI and in particular, spatial filtering with a reference antenna array. We will also use the Cramér-Rao bound derived in this chapter to find the performance bound for these techniques.

4

4.10. APPENDIX: IDENTIFIABILITY

One of the challenges with the FA models is the problem of identifiability. There are two identification problems that need to be addressed. As in [82] we call two solutions, θ_1 and θ_2 , observationally equivalent if for a set of observations with probability density $p(\mathbf{x}; \theta)$, we would have $p(\mathbf{x}; \theta_1) = p(\mathbf{x}; \theta_2)$. The problem is called (globally) identifiable if for a solution θ , there are no observationally equivalent solutions on the entire solution space Θ . The identification problem for the FA model can now be formulated as

- Given a Hermitian matrix \mathbf{R} with the decomposition $\mathbf{R} = \mathbf{R}_0 + \mathbf{D}$ and Q , are \mathbf{R}_0 and \mathbf{D} , as defined in previous sections, identifiable?

The early results on this identification problem are published in [55]. Later work on this subject has been summarized by [83], and [84] gives a recent overview of important theorems on this subject. However only for cases where Q is one or two (very small ranks) do these theorems provide both sufficient and necessary conditions of identifiability. Here we use the results provided by [82] to formulate a necessary and sufficient condition to answer this identification question.

Of course a necessary condition for identifiability is that the number of knowns exceeds the number of unknowns. This puts limits on the rank of \mathbf{A} . To find this limit we will study the degree of freedom we have for the estimation parameter based on a given sample covariance matrix.

From the sample covariance matrix we have $P(P-1)/2$ complex and P real known parameters which are in total P^2 real knowns and from the FA model and the constraints we have PQ complex parameters in \mathbf{A} , P real parameters in \mathbf{D} and Q^2 constraints. As such the total degree of freedom becomes

$$\begin{aligned} s &= P^2 - 2PQ - P + Q^2 \\ &= (P - Q)^2 - P. \end{aligned} \quad (4.66)$$

In order for the FA model to be identifiable, $s > 0$ is a necessary condition. Now solving for Q we find that the maximum number of sources that could theoretically be detected by FA is

$$Q < P - \sqrt{P} \quad (4.67)$$

And following the same procedure for EFA we find

$$Q < P - \sqrt{\text{tr}(\mathbf{M}^2)}. \quad (4.68)$$

In order to find such a bound on Q for JEFA we need to assume that Q_m is constant. In this case we find the following bound

$$Q < P - \sqrt{\frac{\text{tr}(\mathbf{M}^2)}{M}}. \quad (4.69)$$

To be able to use known literature on identifiability it is useful to choose a parametrization of the unknowns in term of real values. let $\boldsymbol{\theta}_R$ denote such a parametrization. One way to define $\boldsymbol{\theta}_R$ for classical FA is

$$\boldsymbol{\theta} = \mathbf{T}\boldsymbol{\theta}_R, \quad (4.70)$$

where

$$\mathbf{T} = \frac{1}{\sqrt{2}} \begin{bmatrix} \mathbf{I}_{PQ} & j\mathbf{I}_{PQ} & \mathbf{0} \\ \mathbf{I}_{PQ} & -j\mathbf{I}_{PQ} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \sqrt{2}\mathbf{I}_P \end{bmatrix}. \quad (4.71)$$

It is straightforward to show that \mathbf{T} is a unitary transformation and hence does not change the number of real unknowns. We will also assume the use of structural constraints on \mathbf{A} which leads to linear constraints (e.g. LQ constraint). Let the minimum constraints on $\boldsymbol{\theta}_R$ be written in the form of a function $\mathbf{h}(\boldsymbol{\theta}_R) = \mathbf{0}$ and solution space be $\boldsymbol{\Theta}_h$. Let the Jacobian of this function be

$$\mathbf{H}_R = \frac{\partial \mathbf{h}}{\partial \boldsymbol{\theta}_R^T} \quad (4.72)$$

and let $\mathbf{F}_R = \mathbf{T}^H \mathbf{F} \mathbf{T}$ be the (real) Fisher Information matrix. Now we reformulate Theorem 2 of [82] as :

Suppose $\boldsymbol{\theta}_0 \in \boldsymbol{\Theta}_h$ is a regular point of $\mathbf{H}_R(\boldsymbol{\theta}_R)$, $\mathbf{V}(\boldsymbol{\theta}_R) = \begin{pmatrix} \mathbf{F}_R \\ \mathbf{H}_R \end{pmatrix}$ and $\text{rank}(\mathbf{A}) = Q$. Then $\boldsymbol{\theta}_0$ is locally identifiable if and only if $\text{rank}[\mathbf{V}(\boldsymbol{\theta}_0)] = 2PQ + P$.

If $\text{rank}(\mathbf{V}) < 2PQ + P$ it means that there is another parametrization $\mathbf{R} = \mathbf{R}_0 + \mathbf{D} = \mathbf{R}_1 + \mathbf{D}_1$ such that $\text{rank}(\mathbf{R}_1) \leq \text{rank}(\mathbf{R}_0)$ and $\mathbf{D} \neq \mathbf{D}_1$. This means that the matrix \mathbf{D} can not be uniquely estimated. If the application is rank reduction for encoding or storage this is not catastrophic. However if the model parameters are used for explaining physical phenomena this problem should be complemented with constraints on the matrix \mathbf{D} itself. For example in array processing if the array signature combined with the noise covariance matrix are unidentifiable then \mathbf{D} will also contain part of the signal power and one of the signal subspaces will be lost.

In this chapter, we assume that the signal and noise have a proper complex Gaussian distribution. This can be used to simplify the identification criteria. Using Bang's formula we can write the FIM as

$$\mathbf{F}_R = \mathbf{J}_R^H (\mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \mathbf{J}_R, \quad (4.73)$$

where $\mathbf{J}_R = \mathbf{J}\mathbf{T}$. Considering that $(\mathbf{R}^{-T} \otimes \mathbf{R}^{-1})$ is a positive definite matrix, $\mathbf{H} = \mathbf{H}_R \mathbf{T}^H$ and $\mathbf{H}^H \mathbf{H}$ has the same row space as \mathbf{H} we have

$$\begin{aligned} \text{rank}(\mathbf{V}) &= \text{rank} \left[\begin{pmatrix} \mathbf{F} \\ \mathbf{H}^H \mathbf{H} \end{pmatrix} \right] \\ &= \text{rank} \left[\begin{pmatrix} \mathbf{J}^H & \mathbf{H}^H \end{pmatrix} \begin{pmatrix} \mathbf{R}^{-T} \otimes \mathbf{R}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{m^2} \end{pmatrix} \begin{pmatrix} \mathbf{J} \\ \mathbf{H} \end{pmatrix} \right] \\ &= \text{rank} \left[\begin{pmatrix} \mathbf{J} \\ \mathbf{H} \end{pmatrix} \right]. \end{aligned} \quad (4.74)$$

This means that by studying the rank of the Jacobian we can establish the identifiability of the problem. However we are interested in finding identifiability conditions that do not depend on \mathbf{R} and \mathbf{D} or the power of the sources. This means that we would like to have conditions that solely depend on the subspace of \mathbf{A} which is \mathbf{U}_0 . To this end we will take a close look at the Jacobians.

Our first observation is that the Jacobian depends only on the structure of \mathbf{D} being diagonal and does not depend on the actual value of its elements. This already shows that studying the Jacobian instead of the total Fisher information matrix is a better way to establish identifiability. For the next part we use the fact that for any matrix $\text{rank}(\mathbf{J}) = \text{rank}(\mathbf{J}^H)$ and we use the (economical) singular value decomposition of $\mathbf{A} = \mathbf{U}_0 \mathbf{\Gamma}^{1/2} \mathbf{Q}^H$ to obtain

$$\begin{aligned} \mathbf{J}^H &= \begin{bmatrix} \mathbf{A}^T \otimes \mathbf{I}_P \\ (\mathbf{A}^H \otimes \mathbf{I}_P) \mathbf{K}^{P,P} \\ (\mathbf{I}_P \circ \mathbf{I}_P)^H \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{Q}^* \mathbf{\Gamma}^{1/2} \mathbf{U}_0^T \otimes \mathbf{I}_P \\ (\mathbf{Q} \mathbf{\Gamma}^{1/2} \mathbf{U}_0^H \otimes \mathbf{I}_P) \mathbf{K}^{P,P} \\ (\mathbf{I}_P \circ \mathbf{I}_P)^H \end{bmatrix} \\ &= \mathbf{\Phi}^H \begin{bmatrix} (\mathbf{\Gamma}^{1/2} \otimes \mathbf{I}_P) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & (\mathbf{\Gamma}^{1/2} \otimes \mathbf{I}_P) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{P.} \end{bmatrix} \tilde{\mathbf{U}} \end{aligned} \quad (4.75)$$

where

$$\tilde{\mathbf{U}} = \begin{bmatrix} \mathbf{U}_0^T \otimes \mathbf{I}_P \\ (\mathbf{U}_0^H \otimes \mathbf{I}_P) \mathbf{K}^{P,P} \\ \mathbf{I}_P \circ \mathbf{I}_P \end{bmatrix} \quad (4.76)$$

It is straightforward to show that the first matrix is a unitary full rank matrix, the second matrix containing $\mathbf{\Gamma}$ is also a diagonal matrix and is full rank which means that rank of the Jacobians depends solely on \mathbf{U}_0 and diagonal structure of \mathbf{D} which is what we needed to show. So the identifiability can be established by examining $\tilde{\mathbf{U}}$. The constraints and hence the matrix \mathbf{H} fix the Q^2 degrees of freedom we have for the matrix \mathbf{Q} which means that in order for the problem to be (locally) identifiable we need

$$\text{rank}(\tilde{\mathbf{U}}) = 2PQ + P - Q^2. \quad (4.77)$$

This result is easily extended to EFA by replacing $(\mathbf{I}_P \circ \mathbf{I}_P)^H$ in (4.76) by \mathbf{J}_{Ψ}^H . The identifiability criterium for EFA becomes

$$\text{rank}(\tilde{\mathbf{U}}) = 2PQ + \text{tr}(\mathbf{M}^2) - Q^2. \quad (4.78)$$

To conclude, we have used the identifiability problem to find the maximum number of sources that can be modeled using (E)FA. We have also shown that the local identifiability of the FA is completely defined by the signal subspace and the structure of the Jacobians with respect to the noise covariance matrix. This structure is completely defined by the masking matrix \mathbf{M} in (4.9).

5

RFI MITIGATION USING A REFERENCE ANTENNA ARRAY

5.1. INTRODUCTION

Radio astronomical observations are increasingly contaminated by man-made RF interference. In bands below 2 GHz, we find TV and radio signals, mobile communication (GSM), radar, satellite communication (Iridium) and localization beacons (GPS, Glonass), etc. Although some bands are specifically reserved for astronomy, the stop-band filters of some communication systems are not always adequate. Moreover, scientifically relevant observations are not limited to these bands. Hence, there is a growing need for interference cancellation techniques.

Current systems for interference cancellation mostly operate at the post-correlation level, by rejecting suspect correlation products in the time-frequency plane, or by specialized imaging algorithms. Spatial filtering at shorter time scales (pre-correlation) is not commonly applied, but would offer interesting possibilities in the first-stage suppression of continuously present wide-band interference in bands that are currently avoided by astronomers. An example is the band between 174-240 MHz, which is currently being populated by Digital Audio Broadcast (DAB) transmissions but is also of interest for the LOFAR radio telescope.

Depending on the interference and the type of instrument, several kinds of RFI mitigation techniques are applicable. Overviews can be found in [10–12, 85, 86]. E.g., intermittent interference such as radar pulses can be detected using short-term Fourier transforms and the contaminated time-frequency cells omitted during long-term integration to order 10 s [10]. Similarly, during postprocessing we can suppress intermittent signals using time-frequency blanking, where detection can be based on anomalous power or higher order spectral kurtosis [87–89]. However, many communication signals

1. Results presented in this chapter are partially published in [28].

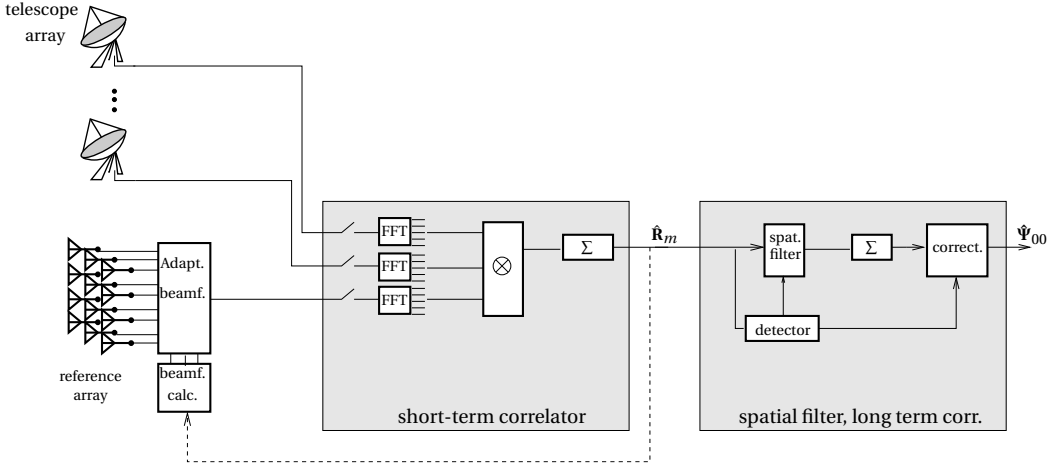


Figure 5.1: Telescope array augmented with a reference phased array. A “telescope” could also be a beam-formed “station” output, where a station consists of an array of antenna elements.

are continuous in time. For a single-dish single-feed telescope, there are not many other options than to consider an extension by a reference antenna which picks up only the interference. In this case LMS-type adaptive cancellation techniques have already been proposed by [92–94].

With an array of P telescope dishes (an interferometer), spatial filtering techniques are applicable as well. The desired instrument outputs in this case are $P \times P$ correlation matrices, integrated to order 10 s (more generally: the time over which astronomical array signals can be considered stationary, also taking the rotation of the Earth into account). Based on short-term correlation matrices (integration to e.g., 10 ms) and narrow subband processing, the array signature vector of an interferer can be estimated and subsequently projected out. The resulting long-term averages of these matrices are mostly interference-free, but they are biased because of the missing dimensions. Such a projection operation also affects the sensitivity and beam-shape of the array [95]. If the projection vector was sufficiently varying, the bias can be corrected for [13, 96]. For stationary interferers (e.g., TV stations or geostationary satellites), this might not work very well, and the correction has to be done during image formation [5]. A special case of a “stationary” interferer is interference entering on only a single telescope dish. The projections will simply remove that channel, and the information can never be recovered. A third limitation is that for relatively weak interference the estimate of the signature vector will not be very accurate so that it will not be perfectly canceled.

To improve on these aspects, we consider in this chapter to extend the telescope array with one or more reference antennas. These might be simple omnidirectional an-

Exceptions are techniques based on higher-order statistics [90] or estimation of outliers in variance [91].

tennas, located close to an interfering source (e.g., a poorly shielded computer in the observatory), or a satellite dish pointing into the direction of a geostationary satellite. Most flexibility is obtained by using a phased array which can adaptively be pointed towards the strongest interferers. In the experiment in Sec. 5.7, we have used a focal plane array that was mounted on one of the telescope dishes, pointing to zenith.

In the context of phased array telescopes consisting of stations that each form beams on the sky, such as LOFAR or SKA, the equivalent of a “telescope” is a station. The reference array may then be a separate array, or remaining degrees of freedom of the stations (e.g., independent beams). It may also consist of the individual station antenna outputs, if we have access to them, or a subset of these.

The generic set-up considered here is shown in Fig. 5.1. The telescope signals (or station beams) are split into narrow sub-bands and correlated to each other over short time intervals. The reference signals are correlated along with the telescope signals as if they were additional telescopes, and spatial filtering algorithms that project out contaminated dimensions can be applied to the resulting short-term integrated covariance matrices. These matrices could also be used to adaptively beam-steer the reference array towards an interferer. The output of the spatial filter is long-term integrated, and formally we have to apply a correction matrix to correct for the projected dimensions.

In the literature, several papers have appeared which propose to apply some form of spatial filtering on extended arrays. Briggs et al. [97] consider a single dual-polarized telescope (two channels), augmented with two reference antennas. With their technique a single interferer can be cancelled; it is not immediately obvious how it can be extended to more general cases (more antennas, more interferers). Kocz et al. [98] propose a projection based spatial filter specifically for multibeam receivers and show its application for detecting pulsars. Jeffs et al. [99, 100] propose spatial filtering algorithms along the lines of [13, 96]; we will summarize their approach in Sec. 5.3.1 and subsequently make extensions which may improve the performance. Hellbourn et al. [101] use the cyclostationarity of RFI signals to improve the projection estimation while using the same projection correction technique used in [13, 96]. The improvements discussed in this Chapter are hence equally applicable to their technique.

The work presented here is the continuation of the previous work done by Albert-Jan Boonstra and Alle-Jan van der Veen [14]. For sake of completeness we repeat their approach in Sec. 5.3.1 and 5.3.2 and then extend it with the use of factor analysis (FA), which has been introduced in Chapter 4. For the complete version of the combined work we refer the reader to [28].

The structure of this chapter is as follows. In Sec. 5.2, we define the data model and state the problem. In Sec. 5.3.1, we present an existing spatial filtering algorithms that does not make use of reference antennas. In Sec. 5.3.2, we extend on this approach, and also present a more generic (Maximum Likelihood) approach based on FA and EFA. Sec. 5.5 discusses the theoretical performance of these algorithms. Sec. 5.6 shows simulation results, and Sec. 5.7 shows results on experimental data.

5.2. PROBLEM STATEMENT

5.2.1. DATA MODEL

Assume we have a telescope array (primary array) with P_0 elements, and a reference array with P_1 elements. The total number of elements is $P = P_0 + P_1$.

We consider complex signals $y_p(t)$ received at the antennas $p = 1, \dots, P$ in a sufficiently narrow subband. For the interference free case the primary array output vector $\mathbf{y}_0(t)$ is modeled in complex baseband form as

$$\mathbf{y}_0(t) = \mathbf{v}_0(t) + \mathbf{n}_0(t)$$

where $\mathbf{y}_0(t) = [y_1(t), \dots, y_{P_0}(t)]^T$ is the $P_0 \times 1$ vector of telescope signals at time t , $\mathbf{v}_0(t)$ is the received sky signal, assumed on the time scale of 10 s to be a zero-mean stationary complex Gaussian vector with covariance matrix $\mathbf{R}_{v,0}$ (the astronomical ‘visibilities’), and $\mathbf{n}_0(t)$ is the $P_0 \times 1$ noise vector. In the general case for an uncalibrated antenna array where the noise on each element is independent and Gaussian, the noise covariance matrix is a diagonal matrix $\mathbf{\Sigma}_0 = \text{diag}(\boldsymbol{\sigma}_0)$, where $\boldsymbol{\sigma}_0$ is a $P_0 \times 1$ vector of noise powers on each element, and if the array has identical elements or is whitened after calibration this simplifies to $\sigma_0^2 \mathbf{I}$. The astronomer is interested in $\mathbf{R}_{v,0}$.

If an interferer is present the primary array output vector is modeled as

$$\mathbf{y}_0(t) = \mathbf{v}_0(t) + \mathbf{a}_0(t)s(t) + \mathbf{n}_0(t)$$

where $s(t)$ is the interferer signal with spatial signature vector $\mathbf{a}_0(t)$ which is assumed stationary only over short time intervals. Without loss of generality, we can absorb the unknown amplitude of $s(t)$ into $\mathbf{a}_0(t)$ and thus set the power of $s(t)$ to 1.

Consider now that we also have a reference antenna array. The outputs of the p_1 reference antennas are stacked into a vector $\mathbf{y}_1(t)$, modeled as

$$\mathbf{y}_1(t) = \mathbf{a}_1(t)s(t) + \mathbf{n}_1(t).$$

It is assumed here that the contribution of the astronomical sources to the reference signals is negligible. The noise on the reference antenna array is assumed to be independent and Gaussian with a diagonal covariance matrix. For an uncalibrated array $\mathbf{\Sigma}_1 = \text{diag}(\boldsymbol{\sigma}_1)$, and for a calibrated array with identical elements it is $\mathbf{\Sigma}_1 = \sigma_1^2 \mathbf{I}$.

Stacking all antenna signals in a single vector

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}_0 \\ \mathbf{y}_1 \end{bmatrix}$$

and similarly for

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}_0 \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{a} = \begin{bmatrix} \mathbf{a}_0 \\ \mathbf{a}_1 \end{bmatrix}, \quad \mathbf{n} = \begin{bmatrix} \mathbf{n}_0 \\ \mathbf{n}_1 \end{bmatrix},$$

we obtain

$$\mathbf{y}(t) = \mathbf{v}(t) + \mathbf{a}(t)s(t) + \mathbf{n}(t). \quad (5.1)$$

We make the following additional assumptions on this model:

In subsequent notation, the subscript ‘0’ will generally refer to the primary array and ‘1’ to the reference array.

- (A1) The noise covariance matrices are unknown diagonal matrices.
- (A2) $\mathbf{R}_{v,0} \ll \mathbf{\Sigma}_0$. This is reasonable as even the strongest sky sources are about 15 dB under the noise floor most phased array systems. (This assumption can be violated by a few strong sources which are observed by a dish).
- (A3) The processing bandwidth is sufficiently narrow, so that possible multipath propagation of the interferer will add up to a single signature vector $\mathbf{a}(t)$ and the interferer is seen as a single source. For this it is at least required that the maximal propagation delay along the telescope array is small compared to the inverse bandwidth.
- (A4) The interferer signature $\mathbf{a}(t)$ is stationary over short processing times (say less than 10 ms). It may or may not vary over longer periods. Note that even interferers fixed on Earth will appear to move as the Earth rotates and the telescopes track a source in different direction. This effect depends on the look direction and the maximal baseline length of the telescopes. (The earlier mentioned window of order 10 s over which $\mathbf{R}_{v,0}$ is stationary is derived from this as well.) The rotation of the telescopes and the associated delay compensation ('fringe stopping') introduced to keep the astronomical signals coherent, give rise to phase changes of the entries of $\mathbf{a}_0(t)$ which for long baselines are significant already over short time intervals. The amplitudes will change because the interferer is usually received via the side lobes of the telescope antenna response, which are highly non-constant and cause temporal variations for tracking dishes or beamformed stations.

The model (5.1) with a calibrated array was considered in [13]. The model is easily extended to multiple interfering sources, in which case we obtain

$$\begin{aligned} \mathbf{y}_0(t) &= \mathbf{v}_0(t) + \mathbf{A}_0(t)\mathbf{s}(t) + \mathbf{n}_0(t) \\ \mathbf{y}_1(t) &= \mathbf{A}_1(t)\mathbf{s}(t) + \mathbf{n}_1(t) \end{aligned}$$

or equivalently

$$\mathbf{y}(t) = \mathbf{v}(t) + \mathbf{A}(t)\mathbf{s}(t) + \mathbf{n}(t)$$

where $\mathbf{A}: P \times Q$ has Q columns corresponding to Q interferers, and $\mathbf{s}(t)$ is a vector with Q entries.

5.2.2. COVARIANCE MODEL

Let be given observations $\mathbf{y}[n] := \mathbf{y}(nT_s)$, where T_s is the sampling period. We assume that $\mathbf{A}(t)$ is stationary at least over intervals of NT_s , and construct short-term covariance estimates $\hat{\mathbf{R}}_m$,

$$\hat{\mathbf{R}}_m = \frac{1}{N} \sum_{n=(m-1)N}^{mN-1} \mathbf{y}[n]\mathbf{y}[n]^H, \quad m = 1, \dots, M$$

where N is the number of samples per short-term average and M is the total number of "snapshots". All interference filtering algorithms in this paper are based on applying operations to each $\hat{\mathbf{R}}_m$ to remove the interference, followed by further averaging over the M resulting matrices to obtain a long-term average.

Considering the $\mathbf{A}_m := \mathbf{A}(mNT_s)$ as deterministic, the expected value of each $\hat{\mathbf{R}}_m$ is denoted by \mathbf{R}_m , which can be written in block-partitioned form as

$$\mathbf{R}_m = \begin{bmatrix} \mathbf{R}_{00,m} & \mathbf{R}_{01,m} \\ \mathbf{R}_{10,m} & \mathbf{R}_{11,m} \end{bmatrix}$$

According to the assumptions, \mathbf{R}_m has model

$$\begin{aligned} \mathbf{R}_m &= \mathbf{A}_m \mathbf{A}_m^H + \mathbf{\Psi} \\ &= \left[\begin{array}{c|c} \mathbf{R}_{v,0} + \mathbf{A}_{0,m} \mathbf{A}_{0,m}^H + \mathbf{\Sigma}_0 & \mathbf{A}_{0,m} \mathbf{A}_{1,m}^H \\ \hline \mathbf{A}_{1,m} \mathbf{A}_{0,m}^H & \mathbf{A}_{1,m} \mathbf{A}_{1,m}^H + \mathbf{\Sigma}_1 \end{array} \right] \end{aligned} \quad (5.2)$$

where $\mathbf{\Psi} := \mathbf{R}_v + \mathbf{\Sigma}$ is the interference-free covariance matrix, $\mathbf{R}_v := \text{bdiag}[\mathbf{R}_{v,0}, \mathbf{0}]$ contains the astronomical visibilities, and $\mathbf{\Sigma} := \text{bdiag}[\mathbf{\Sigma}_0, \mathbf{\Sigma}_1]$ is the diagonal noise covariance matrix. The objective is to estimate the interference-free covariance submatrix (long-term estimate) $\mathbf{\Psi}_{00} := \mathbf{R}_{v,0} + \mathbf{\Sigma}_0$.

5

5.3. EXISTING SPATIAL FILTERING ALGORITHMS

5.3.1. SPATIAL FILTERING USING PROJECTIONS WITHOUT REFERENCE ANTENNAS

In [13], a spatial filtering algorithm based on projections was introduced, and subsequently analyzed in [96]. Although that algorithm did not assume the presence of reference antennas, it can also be used in our current situation. We will first discuss the case where the spatial signature of the interferers are deterministic or known, then we will generalize it to the case where it is estimated from the data.

DETERMINISTIC OR KNOWN SPATIAL SIGNATURE

Suppose that an orthogonal basis of the subspace spanned by interferer spatial signatures $\text{span}(\mathbf{A}_m)$ is known. Let the basis vectors be the columns of a matrix \mathbf{U}_m . We can then form a spatial projection matrix \mathbf{P}_m ,

$$\mathbf{P}_m := \mathbf{I} - \mathbf{U}_m \mathbf{U}_m^H \quad (5.3)$$

which is such that $\mathbf{P}_m \mathbf{A}_m = \mathbf{0}$. When this spatial filter is applied to the data covariance matrix,

$$\hat{\mathbf{Q}}_m := \mathbf{P}_m \hat{\mathbf{R}}_m \mathbf{P}_m$$

then all the energy due to the interferers will be nulled:

$$\mathcal{E}\{\hat{\mathbf{Q}}_m\} = \mathbf{P}_m \mathbf{\Psi} \mathbf{P}_m.$$

If we subsequently average the modified covariance matrices $\hat{\mathbf{Q}}_m$, we obtain a long-term estimate

$$\hat{\mathbf{Q}} := \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{Q}}_m = \frac{1}{M} \sum_{m=1}^M \mathbf{P}_m \hat{\mathbf{R}}_m \mathbf{P}_m. \quad (5.4)$$

$\hat{\mathbf{Q}}$ is an estimate of $\mathbf{\Psi}$, but it is biased due to the projection. To correct for this we first write the two-sided multiplication as a single-sided multiplication employing the matrix identity $\text{vect}(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A})\text{vect}(\mathbf{B})$. This gives

$$\text{vect}(\hat{\mathbf{Q}}) = \frac{1}{M} \sum_{m=1}^M \mathbf{C}_m \text{vect}(\hat{\mathbf{R}}_m) \quad (5.5)$$

where

$$\mathbf{C}_m := \mathbf{P}_m^T \otimes \mathbf{P}_m.$$

If the interference was completely removed then

$$\mathcal{E}\{\text{vect}(\hat{\mathbf{Q}})\} = \frac{1}{M} \sum_{m=1}^M \mathbf{C}_m \text{vect}(\mathbf{\Psi}) = \mathbf{C} \text{vect}(\mathbf{\Psi}) \quad (5.6)$$

where

$$\mathbf{C} := \frac{1}{M} \sum_{m=1}^M \mathbf{C}_m.$$

In view of this, we can apply a correction \mathbf{C}^{-1} to $\hat{\mathbf{Q}}$ to obtain the corrected estimate $\hat{\mathbf{\Psi}}$

$$\hat{\mathbf{\Psi}} := \text{unvect}(\mathbf{C}^{-1} \text{vect}(\hat{\mathbf{Q}})). \quad (5.7)$$

If the interference was completely projected out then $\hat{\mathbf{\Psi}}$ is an unbiased estimate of the covariance matrix without interference. This algorithm was introduced in [13] and its performance was discussed in [96].

The main computational complexity is in constructing \mathbf{C} and inverting it, as this is generally a very large matrix ($P^2 \times P^2$). Inversion of this would require $O(P^6)$ operations, but because the inverse is applied to only a single vector this can be reduced to $O(P^4)$ using numerical techniques. Note in this respect that, for large M and sufficiently varying \mathbf{A}_m , \mathbf{C} is usually quite close to an identity matrix, and the hope would be that the correction can be omitted or highly simplified under such conditions.

The reconstructed covariance matrix $\hat{\mathbf{\Psi}}$ is size $P \times P$. In the present case, we are only interested in the submatrix corresponding to the primary antennas. Hence, the estimate produced by the algorithm is the $P_0 \times P_0$ submatrix in the top-left corner, $\hat{\mathbf{\Psi}}_{00}$.

UNKNOWN SPATIAL SIGNATURE, KNOWN NOISE COVARIANCE

The spatial signatures of the interferers are generally unknown, but if the noise covariance $\mathbf{\Sigma}$ is known the interfering subspace can be estimated from an eigen-analysis of the sample covariance matrices $\hat{\mathbf{R}}_m$. If $\mathbf{\Sigma}$ is not a multiple of \mathbf{I} , then we first have to whiten the noise to make the noise powers on all antennas the same. This is done by working with $\mathbf{\Sigma}^{-1/2} \hat{\mathbf{R}}_m \mathbf{\Sigma}^{-1/2}$. Without interference and assuming \mathbf{R}_v is negligible compared to $\mathbf{\Sigma}$, all eigenvalues of this matrix are expected to be close to 1. With Q interferers, Q eigenvalues become larger, and the eigenvectors corresponding to these eigenvalues are an estimate of $\text{span}(\mathbf{\Sigma}^{-1/2} \mathbf{A}_m)$.

Remarks:

In comparison, image formation techniques work with correlation matrices of size $P \times P$ and have a complexity of $O(P^2)$.

1. The algorithm relies on the invertibility of \mathbf{C} , which is constructed from projection matrices. Each projection matrix is rank deficient. Hence, \mathbf{C} is invertible only if the spatial signature vectors which are projected out are sufficiently varying. In [96] it was noted that for $Q = 1$, usually already 3 different projections are sufficient to guarantee that \mathbf{C} is full rank.
2. The algorithm is inefficient in the sense that it first reconstructs the complete covariance matrix, then selects the submatrix of interest. Since more parameters (the complete covariance) are estimated, the performance (estimation accuracy) is reduced.
3. If the noise covariance $\mathbf{\Sigma}$ is not known, then the eigenvalue decomposition can be replaced by a more general Factor Analysis decomposition, see Sec. 5.4.1.
4. Regarding the subspace estimation, the maximum number of interferers is constrained by $Q < P$. Each \mathbf{C}_m has size $P^2 \times P^2$ and rank $(P - Q)^2$; invertibility requires at least $M(P - Q)^2 \geq P^2$ (in case the projected subspaces are completely arbitrary).

In summary, this spatial filtering algorithm does not really take advantage of the reference antennas. In the processing, it treats them like ordinary antennas. The only benefit obtained from them is that, with an improved INR, the estimate of the interference subspace will be better, so that the interference can be filtered out better. The performance is then limited by the conditioning of \mathbf{C} (thus the variability of the spatial signature vectors).

5.3.2. SPATIAL FILTERING USING REFERENCE ANTENNAS

IMPROVED SPATIAL FILTER WITH PROJECTIONS

Taking the above remarks into account, we derive an improved algorithm. Compute the projections and long-term average of the projected estimates $\hat{\mathbf{Q}}$ as before in (5.4). Then (5.6) applies:

$$\mathcal{E}\{\text{vect}(\hat{\mathbf{Q}})\} = \mathbf{C}\text{vect}(\mathbf{\Psi}).$$

Based on this, we previously set $\text{vect}(\hat{\mathbf{\Psi}}) = \mathbf{C}^{-1}\text{vect}(\hat{\mathbf{Q}})$, which is the solution in LS sense of the covariance model error minimization problem, $\|\text{vect}(\hat{\mathbf{Q}}) - \mathbf{C}\text{vect}(\hat{\mathbf{\Psi}})\|^2$. Now, instead of this, partition $\mathbf{\Psi}$ as in (5.2) into 4 submatrices. Since we are only interested in recovering $\mathbf{\Psi}_{00}$, the other submatrices in $\hat{\mathbf{\Psi}}$ are replaced by their expected values, respectively $\mathbf{\Psi}_{01} = \mathbf{0}$, $\mathbf{\Psi}_{10} = \mathbf{0}$, $\mathbf{\Psi}_{11} = \mathbf{\Sigma}_1$. This corresponds to solving the reduced-size covariance model error minimization problem,

$$\hat{\mathbf{\Psi}}_{00} = \arg \min_{\mathbf{\Psi}_{00}} \left\| \text{vect}(\hat{\mathbf{Q}}) - \mathbf{C}\text{vect} \left(\left[\begin{array}{c|c} \mathbf{\Psi}_{00} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{\Sigma}_1 \end{array} \right] \right) \right\|^2.$$

The solution of this problem reduces to a standard LS problem after separating the knowns from the unknowns. Thus, rearrange the entries of $\text{vect}(\mathbf{\Psi})$ into

$$\begin{bmatrix} \text{vect}(\mathbf{\Psi}_{00}) \\ \sigma_1 \\ \mathbf{0} \end{bmatrix}$$

where $\sigma_1 = \text{vectdiag}(\Sigma_1)$, and repartition \mathbf{C} accordingly, to obtain the equivalent problem

$$\begin{aligned}\hat{\psi}_{00} &= \arg \min_{\psi_{00}} \left\| \text{vect}(\hat{\mathbf{Q}}) - [\mathbf{C}_1 \ \mathbf{C}_2 \ \mathbf{C}_3] \begin{bmatrix} \psi_{00} \\ \sigma_1 \\ \mathbf{0} \end{bmatrix} \right\|^2 \\ &= \arg \min_{\psi_{00}} \|(\text{vect}(\hat{\mathbf{Q}}) - \mathbf{C}_2 \sigma_1) - \mathbf{C}_1 \psi_{00}\|^2 \\ &= \mathbf{C}_1^\dagger (\text{vect}(\hat{\mathbf{Q}}) - \mathbf{C}_2 \sigma_1),\end{aligned}\tag{5.8}$$

where $\hat{\psi}_{00} = \text{vect}(\hat{\Psi}_{00})$ and $\psi_{00} = \text{vect}(\Psi_{00})$. The advantage compared to the preceding algorithm is that \mathbf{C}_1 is a tall matrix, and better conditioned than \mathbf{C} . This improves the performance of the algorithm in cases where \mathbf{C} is ill-conditioned.

Remarks:

1. The subspace estimation has not changed, and the maximum number of interferers is still constrained by $Q < P$. Now \mathbf{C}_1 has size $P^2 \times P_0^2$ and invertibility requires at least $M(P - Q)^2 \geq P_0^2$, in case the projected subspaces are completely arbitrary.
2. Even if the interferers are located stationary (\mathbf{A}_m constant), \mathbf{C}_1 is expected to have full column rank and hence the improved algorithm can estimate the astronomical covariance (provided $Q \leq P_1$).
3. The same advantage holds in case an interferer only contaminates one of the primary antennas ($\mathbf{a}_{i,0}$ has only one nonzero entry). Without reference antenna, the projection is always the same and cannot be corrected: the correlations corresponding to that antenna are lost. With a reference antenna, they can be recovered.
4. If the array is calibrated then we can assume Σ_1 to be known from calibration. For an uncalibrated array we suggest using FA to estimate Σ_1 , as described next in Sec. 5.4.1.

5.4. IMPROVED FILTERING

5.4.1. FACTOR ANALYSIS

If we ignore the astronomical correlations \mathbf{R}_v , then the short-term covariance matrices are modeled as in (5.2),

$$\mathbf{R}_m = \mathbf{A}_m \mathbf{A}_m^H + \Sigma, \tag{5.9}$$

where $\Sigma = \text{bdiag}[\Sigma_0, \Sigma_1]$ is a diagonal matrix. This is clearly the same model as the classical FA in Chapter 4. Using FA we can find $\mathbf{P}_m = \mathbf{I} - \mathbf{A}_m (\mathbf{A}_m^H \mathbf{A}_m)^{-1} \mathbf{A}_m^H$ and $\sigma_1 = \text{diag}(\Sigma_1)$ in (5.8), which are needed components for the algorithms described in the previous sections.

Several estimation algorithms and their complexities are listed in table 4.1. In general, the complexity is of $O(P^2 Q)$ per short-term correlation matrix. Detection of the number of interferers Q was discussed in Sec 4.7.

5.4.2. DIRECT ML ESTIMATION USING EXTENDED FACTOR ANALYSIS

As we will show here, the data model (5.2) satisfies the (J)EFA model which we introduced in Sec. 4.3.2. This will allow us to directly find a Maximum Likelihood (ML) estimate for Ψ_{00} .

The covariance model (5.2) is

$$\mathbf{R}_m = \mathbf{A}_m \mathbf{A}_m^H + \Psi = \mathbf{A}_m \mathbf{A}_m^H + \left[\begin{array}{c|c} \Psi_{00} & \mathbf{0} \\ \hline \mathbf{0} & \Sigma_1 \end{array} \right]. \quad (5.10)$$

where we are interested in estimating the unknown square matrix Ψ_{00} and, for an uncalibrated array, Σ_1 is unknown. Thus, the appropriate masking matrix \mathbf{M} such that $\Psi = \mathbf{M} \odot \Psi$ is

$$\mathbf{M} = \left[\begin{array}{c|c} \mathbf{1}\mathbf{1}^T & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{I} \end{array} \right].$$

If we replace Ψ by Ψ_m (i.e., for each snapshot \mathbf{R}_m we estimate an independent Ψ_m), then we can apply EFA on each snapshot. Of course, as discussed in the previous chapter, we can also estimate \mathbf{A}_m and a stationary Ψ jointly using JEFA. We will first discuss the application of EFA on each snapshot.

Each $\hat{\mathbf{R}}_m$ will give us an estimate $\hat{\Psi}_m$, and $\hat{\Psi}_{00,m}$ is simply the upper left sub-block of this matrix. The long-term estimate is given by

$$\hat{\Psi}_{00} = \frac{1}{M} \sum_{m=1}^M \hat{\Psi}_{00,m}. \quad (5.11)$$

A necessary condition for identification is that the degrees of freedom $s > 0$. We use (4.68) and $\text{tr}(\mathbf{M}^2) = P_0^2 + P_1$ and we find

$$s = P^2 + Q^2 - 2PQ - (P_0^2 + P_1).$$

Solving for the number of reference antennas P_1 we find

$$P_1 > Q - P_0 + \frac{1}{2} + \sqrt{P_0^2 + Q - P_0 + \frac{1}{4}}. \quad (5.12)$$

Thus, if P_0 is small, we need $P_1 > Q + \sqrt{Q}$, and if P_0 is large, this reduces to $P_1 > Q$.

Similarly for JEFA we have

$$P_1 > Q - P_0 + \frac{1}{2M} + \sqrt{P_0^2 + Q - P_0 + \frac{1}{4M^2}}. \quad (5.13)$$

This shows that the number of snapshots has little effect on the number of reference antenna needed.

Remarks:

1. For each snapshot $2PQ + P_0^2 + P_1$ unknowns are estimated. The computational complexity of the Krylov Scoring is $O(P^2Q + Q^3 + I_K(P^2Q + PQ^2))$ where I_K is the number of Krylov iterations. Note that no expensive post processing as in Eq. (5.7) or (5.8) is needed.

2. Although each snapshot is processed with a Maximum Likelihood estimator, the overall algorithm is not maximum likelihood as Ψ_{00} is estimated using an average of the $\hat{\Psi}_{00,m}$ in (5.11). This is improved by using JEFA.
3. If not all primary antennas receive the RFI, then P_0 should be replaced by \tilde{P}_0 in (5.12), where \tilde{P}_0 is the number of primary antennas that receive RFI.
4. This approach assumes that the number of interferers is known. If this is not the case we can follow a similar approach as was suggested in Sec. 5.4.1.

5.5. PERFORMANCE ANALYSIS

5.5.1. PROJECTION BASED SPATIAL FILTERING

The performance of the projection based spatial filtering as described in Sec. 5.3.1 is discussed in [96]. We only restate the final result without derivation here. The performance of the estimated RFI covariance matrix $\hat{\Psi}$ is given by

$$\text{Cov}\{\hat{\Psi}\} = \mathbf{C}^{-1} \text{Cov}\{\hat{\mathbf{Q}}\} \mathbf{C}^{-1} \approx \frac{\sigma^4}{MN} \mathbf{C}^{-1}, \quad (5.14)$$

where we have assumed $\Sigma = \sigma^2 \mathbf{I}$ and $\text{Cov}\{\mathbf{X}\}$ is the covariance matrix of $\text{vect}(\mathbf{X})$. The final estimate $\hat{\Psi}_{00}$ is a submatrix of $\hat{\Psi}$. Its performance is a submatrix of $\text{Cov}\{\hat{\Psi}\}$. Compared to the RFI case (where $\text{Cov}\{\hat{\Psi}\} \approx \frac{\sigma^4}{MN} \mathbf{I}$), this indicates that \mathbf{C}^{-1} determines the relative performance of the spatial filtering algorithm of Sec. 5.3.1. The conditioning of \mathbf{C}^{-1} depends on the variability of \mathbf{A}_m , the spatial signatures of the interferer. For large M and sufficiently varying \mathbf{A}_m , $\mathbf{C} \rightarrow \mathbf{I}$ and the performance is expected to be similar to the interference-free case.

Based on the derivation presented in [96], the performance of the extended algorithm with reference antenna can also be found to be [28]

$$\text{Cov}\{\hat{\Psi}_{00}\} = \mathbf{C}_1^\dagger \text{Cov}\{\hat{\mathbf{Q}}\} \mathbf{C}_1^{\dagger H} \approx \frac{\sigma^4}{MN} \mathbf{C}_1^\dagger \mathbf{C} \mathbf{C}_1^{\dagger H} \quad (5.15)$$

It is known that for any tall matrix \mathbf{J} for which \mathbf{CJ} is full column rank

$$(\mathbf{CJ})^\dagger \mathbf{C}(\mathbf{CJ})^{\dagger H} \leq \mathbf{J}^H \mathbf{C}^{-1} \mathbf{J}$$

([102, lemma 3.1]). Choosing \mathbf{J} a selection matrix such that $\mathbf{C}_1 = \mathbf{CJ}$, it can be deduced that the algorithm of Sec. 5.3.2 is always more efficient than the algorithm of Sec. 5.3.1.

5.5.2. THE CRB OF Ψ_{00} AND ASYMPTOTIC STATISTICS OF (J)EFA

As a maximum likelihood technique, the EFA algorithm applied to a single short-term estimate is unbiased and will asymptotically reach the CRB, hence its asymptotic performance is given by the appropriate submatrix of (4.17). In order to avoid confusion between the CRB and the matrix \mathbf{C} in previous sections we will denote the CRB derived in Sec. 4.4 by $\mathbf{\Gamma}$ instead of \mathbf{C} .

The long-term estimate $\hat{\Psi}_{00}$ in EFA is obtained by simply averaging the short-term estimates (assuming the estimates are independent), so that its performance is given by the covariance matrix (using the results from Sec. 4.4)

$$\Gamma_{00,efa} = 1/M^2 \sum_m \Gamma_{00,m}, \quad (5.16)$$

similar to (5.16). Using simulations we will show that this asymptotic performance of EFA is achieved for a moderate number of samples and/or INR.

In reality, the short-term estimates are not independent as the data model shows that they have Ψ in common, and we should use the JEFA algorithm to estimate the RFI free covariance matrix over the entire data set of M covariance matrices. This approach improves the performance considerably. The CRB on the long-term estimate for Ψ is also derived in Sec. 4.4 and the corresponding bound for $\hat{\Psi}_{00}$ is given by the submatrix of the result given by (4.33) and is denoted by Γ_{00} .

5.5.3. PERFORMANCE FOR LONG TERM INTEGRATION (IMAGING)

So far we have discussed the statistics for a single long-term estimate $\hat{\Psi}_{00}$. In many astronomical applications we need to combine a large number of these estimates in order to boost the desired signals to detectable levels. One example of such application is producing a two-dimensional image of the sky.

Fourier based imaging (called the ‘dirty’ image in astronomy, i.e., prior to further deconvolution) can be viewed as computing a weighted average of the entries of the long-term covariance estimates [5]. If K estimates are averaged, then the variance of the estimates is scaled by $1/K$. Without RFI, the individual estimates have a variance given by $\text{Cov}\{\hat{\mathbf{R}}_m\} = (\mathbf{R}_m^T \otimes \hat{\mathbf{R}}_m)/N$. With RFI removal using projections, the estimates have a slightly higher variance given by (5.14) or (5.15). The performance penalty corresponds to the missing data in the projected dimensions, which is natural and acceptable.

The main worry for astronomers would come from any bias that is present in the long-term covariance estimates. The following remarks can be made.

- As a maximum likelihood technique, the EFA is not biased. However, a bias can be present in case a weak interferer is present but not detected (i.e., model mismatch in the EFA). This is a natural limitation in any interference removal technique. The residual interference must be detected and removed after further averaging.
- The projection techniques have a bias, but this bias is present on the diagonal entries of the long-term covariance estimates [96]. Many imaging techniques routinely omit these diagonal entries (the auto-correlations of the antennas) because they are dominated by the system noise. Alternatively, it is possible to correct for the bias to a certain extent [96]. A second source of bias corresponds to RFI that is present but not detected (which as mentioned above is a common problem for RFI mitigation techniques).

In any case, the averaging inherent in the imaging process has a tendency to wash out any residual interference.

In the next section we use the CRB to show that the fundamental bound on the total variance of the estimated covariance is very close to the RFI free case, and we also

show that the proposed algorithms are close to this bound provided that the RFI is strong enough to be detected.

5.6. SIMULATIONS

We first test the performance of the algorithms in a simulation set-up. We use $P = 7$ antennas, with $P_0 = 5$ primary antennas (telescopes) and $P_1 = 2$ reference antennas. For simplicity, the array is a uniform linear array with half-wavelength spacing and the same noise power on all antennas.

The astronomical source is simulated by a source with a constant direction-of-arrival of 10° with respect to array broadside. The source has $\text{SNR}_0 = -20$ dB with respect to each primary array element, and $\text{SNR}_1 = -40$ dB for the reference antenna.

The interferer is simulated by a source with a randomly generated and varying complex \mathbf{a}_m , and varying INRs. This corresponds to a Rayleigh fading interferer. A GLRT is performed with a false alarm probability of 0.1 to detect the interfering signal.

The following algorithms are compared:

- the traditional subtraction method [28], is denoted by ‘Trad Filter’,
- the spatial filtering algorithm using projections and eigenvalue computations, Sec. 5.3.1, denoted ‘eig-ref’,
- the improved spatial filtering algorithm with reduced-size covariance reconstruction, Sec. 5.3.2, denoted ‘eig-ref-red’ and for Factor Analysis version ‘fa-ref’,
- the version that uses Extended Factor Analysis is denoted as ‘EFA’, Sec. 5.4.2, and the joint version is denoted as ‘JEFA’,
- for comparison, the spatial filtering technique without reference antenna, denoted ‘eig-no-ref’, the covariance estimate without RFI (‘RFI free’), and the estimate obtained without any filtering (‘no-filter’).

Fig. 5.2(a) shows the relative mean-squared-error (MSE) of the primary filtered covariance estimate compared to the theoretical value $\mathbf{R}_{v,0} + \sigma_0^2 \mathbf{I}$, for varying interferer powers INR_0 on the primary array. Here, we took $N = 5000$ short-term samples and $M = 2$ long-term averages, which is unrealistically small but serves to illustrate the effect of limited variability of \mathbf{a}_m (only two different vectors). The interferer array gain was $\text{INR}_1 - \text{INR}_0$ of 5 dB. Similarly, Fig. 5.2(b) shows the MSE for varying INR difference and an INR_0 of 10 dB. The amplitudes and phases are varying after each short-term averaging period (i.e., the interferer is Rayleigh fading).

In Fig. 5.3, we consider a case where the RFI enters the primary array on only a single element. In Fig. 5.3(a), the INR_0 is varied, while the INR difference is 0 dB and we consider a shorter short-term integration time $N = 200$ and a longer long-term averaging time $M = 5$ than before. Similarly, Fig. 5.3(b) shows the MSE for varying short-term integration samples N , for an INR_0 of 10 dB.

As reference line, we show the CRB Γ_{00} derived in (4.33). Because this is a matrix, we

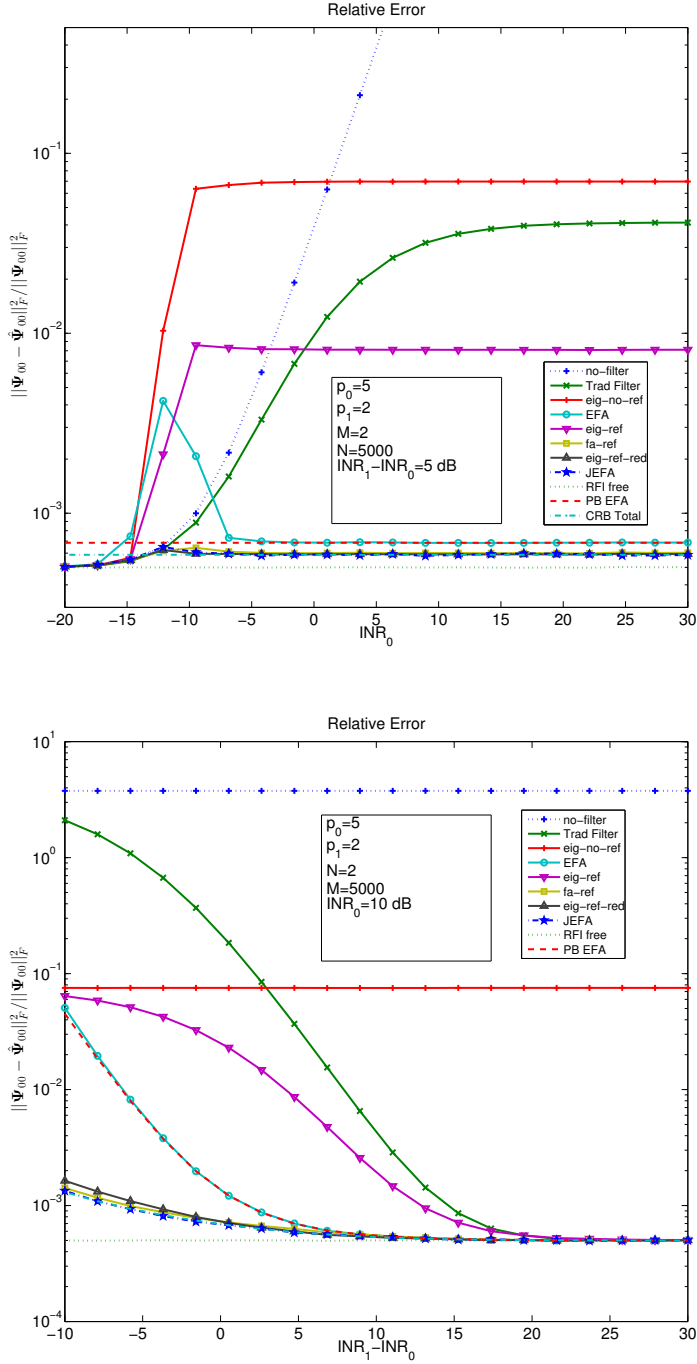


Figure 5.2: Simulation with limited variability of the interferer array response vector ($N = 2$). Relative MSE (a) as function of interferer power (in dB) at the primary array elements, (b) as function of the interferer power difference (in dB) between the reference elements and the primary elements.

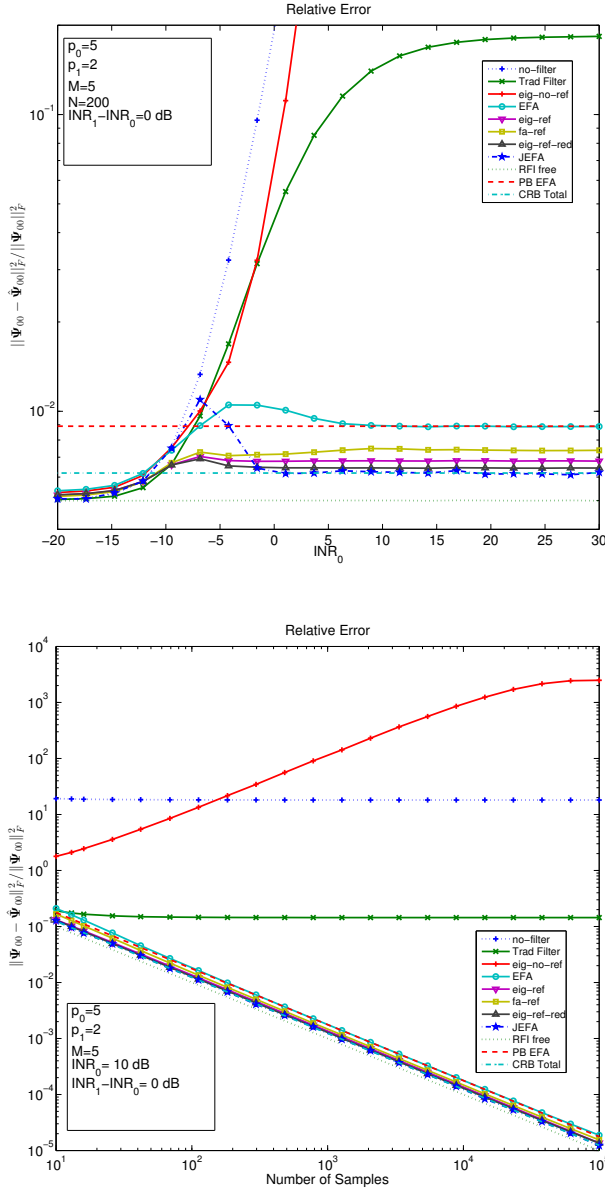


Figure 5.3: (a) Simulation with the interferer entering on only a single primary antenna. Relative MSE (a) as function of interferer power (in dB) at the primary elements, (b) as a function of short-term integration samples N .

use the following relation between the MSE and the trace of the CRB:

$$\begin{aligned}\mathcal{E} \|\mathbf{\Psi}'_{00}\|_F^2 &= \mathcal{E} \text{vect}(\mathbf{\Psi}'_{00})^H \text{vect}(\mathbf{\Psi}'_{00}) \\ &= \mathcal{E} \text{tr}(\text{vect}(\mathbf{\Psi}'_{00}) \text{vect}(\mathbf{\Psi}'_{00})^H) \\ &\geq \text{tr}(\mathbf{\Gamma}_{00})\end{aligned}\tag{5.17}$$

hence $\text{tr}(\mathbf{\Gamma}_{00})$ is a bound on the MSE performance of the proposed algorithms. The MSE is estimated using Monte-Carlo runs.

Observations are:

- The new algorithms that use projections with a reference antenna array (eig-ref-red and fa-ref) operate close to the CRB and have a great advantage over the spatial filtering algorithm without reference antenna (eig-no-ref) in case the \mathbf{a}_m -vector is not sufficiently varying [see Fig. 5.2(a)]. The MSE performance is flat for varying INR and INR difference, which is very desirable. Moreover, it is very close to the RFI-free case. Using FA to find the projections does not noticeably degrade the performance of the filter even though more parameters are estimated.
- The CRB is generally close to the RFI-free case. For low INR, the performance can be better than the CRB because the RFI is not detected and the CRB does not take this model mismatch into account.
- The EFA method also performs well for reasonable INR difference. It operates close to its theoretical performance bound unless the RFI is not detected. However, this bound is seen to be appreciably higher than the CRB in some simulations. This is because the EFA estimates the parameters of each short-term covariance matrix separately, whereas they have parameters in common (i.e., $\mathbf{\Psi}$). As suggested using JEFA improves the performance significantly.
- The new algorithms are often better than the subtraction technique (Trad Filter). The subtraction is only accurate if the INR difference is large compared to the INR at the primary array. If the INR difference is small, or if the INR at the primary array is relatively large, then the subtraction technique fails. This is probably caused by the bias in the inverted term (power of the interferer, with added noise power). It makes the algorithm not reliable to use. As is illustrated in Fig. 5.3(b) the traditional subtraction does not improve with a higher number of short-term samples which indicates that this is not an efficient estimator.
- If the interferer enters only on one telescope and on the reference antenna, as in Fig. 5.3, then the algorithm without a reference antenna is performing poorly: it cannot reconstruct the contaminated dimension. The algorithms with reference antennas perform fine.

In summary, based on these simulations, we recommend to consider JEFA and ‘fa-ref’. In these set of simulation we have used a noise realization which is white and hence the performance of ‘eig-ref-red’ (Sec. 5.3.2) is similar with ‘fa-ref’. We recommend this approach for systems with a known or white noise.




Figure 5.4: Reference focal-plane array mounted on a dish.

5.7. EXAMPLES ON EXPERIMENTAL DATA

5.7.1. EXPERIMENT I

To test the algorithm on actual data, we have made a short observation of the strong astronomical source 3C48 contaminated by Afristar satellite signals. The set-up follows Fig. 5.1. The primary array consists of $P_0 = 3$ of the 14 telescope dishes of the Westerbork Synthesis Radio Telescope (WSRT), located in The Netherlands. As reference signals we use $P_1 = 27$ of 52 elements of a focal-plane array that is mounted on another dish of the WSRT which is set off-target (see Fig. 5.4) such that it has no dish gain towards the astronomical source nor to the interferer.

We recorded 13.4 seconds of data with 80 MS/s, and processed these offline. Using short-term windowed Fourier transforms, the data was first split into 8192 frequency bins (from which we used 1537), and subsequently correlated and averaged over $N = 4048$ samples to obtain $M = 64$ short-term covariance matrices.

Fig. 5.5(a) shows a few autocorrelations and crosscorrelations on the primary antennas and Fig. 5.5(b) shows the autocorrelation of 6 reference antennas. The interference is clearly seen in the spectrum. The interference consists of a lower and higher frequency part. The low frequency part is stronger on the reference antenna and the higher part stronger on the primary antenna. However, because of a relatively high number of reference antennas the total INR, as we will see, is high enough for the algorithms to be effective.

Because no calibration step has been performed we use a generalized likelihood ratio test (GLRT) [103] to detect if each frequency bin is contaminated with RFI and then we use FA to estimate the noise powers and the signal spatial signature. The result of whitening the spectrum with the estimated result of FA is shown in Fig. 5.6(a).

The resulting auto- and crosscorrelation spectra after filtering are shown in Fig. 5.7. The autocorrelation spectra are almost flat, and close to 1 (the whitened noise power). The cross-correlation spectra show that the spatial filtering with reference antenna has removed the RFI within the sensitivity of the telescope. Also it shows the power of using FA and EFA at this stage in the processing chain, as they do not require the array to be calibrated.




Figure 5.5: Observed spectrum from (a) the primary telescopes, (b) 6 of the reference antennas

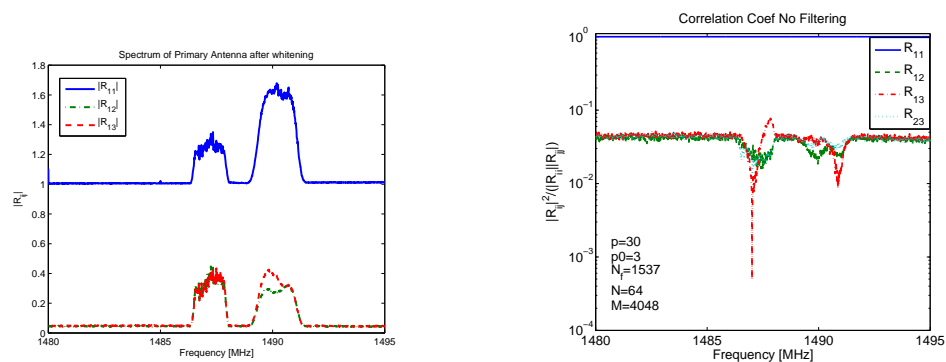


Figure 5.6: (a) Spectra of primary antenna after whitening, (b) Average normalized correlation coefficients without filtering

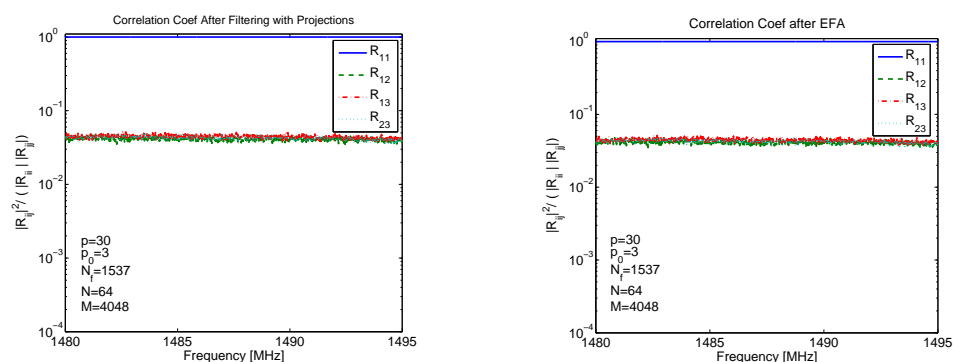


Figure 5.7: Averaged Normalized correlation coefficients (a) after filtering using method in section 5.3.2, (b) after using EFA.




Figure 5.8: Spectrum received at a LOFAR HBA station

5.7.2. EXPERIMENT II

In a second experiment, we use raw data from the LOFAR station RS409 in HBA mode 5 (100-200 MHz), acquired via the transient buffer board. Data from the 46 (out of 48) x-polarization receiving elements are sampled with a frequency of 200 MHz and correlated. Samples are then divided into 1024 subbands with the help of tapering and an FFT. From these samples we form $M = 4$ covariance matrices with an integration time of 19 ms ($N = 1862$) for each subband. No calibration was done on the resulting covariance matrices.

The LOFAR HBA has a hierarchy of antennas, where a single receiving element output is the result of analog beamforming on 16 antennas (4×4) in a tile. During the measurements the analog beamformers were tracking the strong astronomical source Cyg A.

The received spectrum is shown in Fig. 5.8. Above 174 MHz, the spectrum is heavily contaminated by wideband DAB transmissions.

We have used 6 of the 46 receiving elements as reference array for our filtering techniques and the rest as primary array. Because we do not have dedicated reference antennas and that the data is already beamformed the assumption that the source is too weak at each short integration time (19 ms) is not completely valid. Also the assumption that the sky sources are much weaker on the reference antennas is not valid in this case because the reference array elements are also following Cyg A. Finally, we have the same exposure to the RFI on the secondary array as we have on the primary so there is no additional RFI gain for the secondary array.

This experiment was conducted on March 12th 2013

To illustrate the performance of the filtering technique we produce snapshot images of the sky (i.e., images based on a single covariance matrix). For an uncontaminated image, we have chosen subband 250 at 175.59 MHz, see Fig. 5.9(a), while for RFI-contaminated data we take subband 247 at 175.88 MHz, see Fig. 5.9(b). These two subbands have been chosen because they are close to each other (in frequency) and we expect that the astronomical images for these bands would be similar. Subband 247 is heavily contaminated and has a 10 dB flux increase on the auto-correlations and a 20 dB increase on the cross-correlations.

The repeated source visible in Fig. 5.9(a) is Cyg A; the repetition is due to the spatial aliasing which occurs at these frequencies (the tiles are separated by more than half a wavelength). The contaminated image in Fig. 5.9(b) shows no trace of Cyg A; note the different amplitude scale which has been increased by a factor 100.

Fig. 5.10 (a) shows the image after filtering the RFI using the algorithm with FA and projections ('fa-ref') as presented in Sec. 5.3.2, and Fig. 5.10 (b) shows the image after using EFA (Sec. 5.4.2). Both images are nearly identical, and very similar to the clean image in Fig. 5.9(a).

Remarks:

- This data shows an example where the contaminated portion of the spectrum is broad, limiting the applicability of time-frequency blanking (post-processing).
- Both filtering techniques appear robust against the modeling errors implicit in this experiment setup.
- The resulting covariance estimates produces snapshot images comparable to an RFI-free channel.

Unfortunately the available data collection system at LOFAR did not allow us to create images with longer integration times.

5.8. CONCLUSIONS

Spatial filtering algorithms for removing RFI on covariance matrix estimates using reference antennas have been proposed, applicable to both calibrated and uncalibrated arrays. For the uncalibrated case, Factor Analysis is used to estimate the interference subspace. An algorithm to estimate RFI-free covariance matrices directly using Extended Factor Analysis (EFA) has also been presented. The statistical performance of the proposed algorithms has been evaluated and the CRB for the entire dataset is presented.

These algorithms generalize previously proposed spatial filtering algorithms that did not use a reference array. Simulations show that using a reference array is beneficial even if the reference antennas receive less interference power than the primary antennas. Another advantage of a reference array is that the algorithms are applicable even if the interference enters on only a single primary antenna, which was not the case for the previously proposed projection algorithm. The algorithms for uncalibrated arrays based on FA and EFA have also been tested on experimental data from astronomical instruments to illustrate their applicability and performance in real-world scenarios; the results are very encouraging.

A disadvantage of the projection techniques is that they require a computationally




Figure 5.9: (a) Clean subband 250, (b) Contaminated subband 247

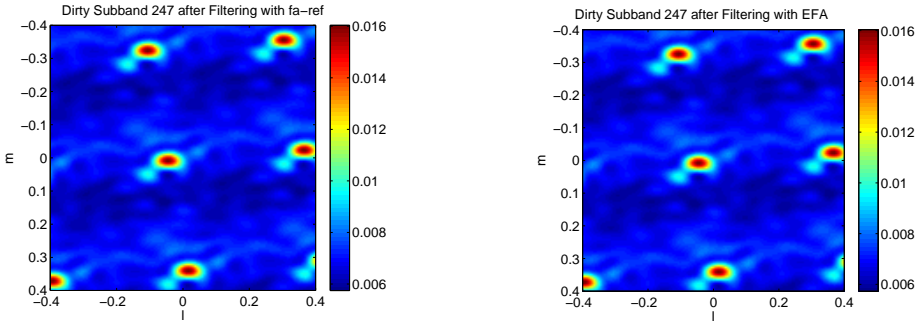


Figure 5.10: (a) Result of filtering using 'fa-ref', (b) Result of filtering using EFA

unattractive matrix inversion which is needed to correct the covariance estimates for the missing (projected) dimensions. The EFA technique is a direct technique which averages maximum likelihood estimates of cleaned short-term covariance matrices; it is computationally more attractive. Unfortunately, the simulations indicate that the EFA method has a lower performance for low INR and/or low number of samples. The key of the problem is that EFA processes short-term covariance matrices independently and does not exploit that they have a common term Ψ . The solution is in joint processing of these matrices using JEFA, as has been demonstrated in the simulations.

6

DIRECTION INDEPENDENT GAIN CALIBRATION

6.1. INTRODUCTION

As explained in Sec. 3.3, to model the signals coming from different directions using phase–delay relations we need to make sure that any other phase change that is not related to source’s position is calibrated for. This is very important for retrieving the correct shape of an object when producing accurate skymaps. The model, as will be presented in the next section, is non–linear and the number of unknowns grows with the number of elements in the array and calibration sources. This means that for very large arrays we are dealing with a complex optimization problem. The authors in [15] and [4] overcome this problem by splitting the unknowns into groups for which a closed formed solution could be found. Then by applying alternating optimization algorithms such as alternating least squares (ALS) or alternating weighted least squares (WALS), a monotonic convergence to the solution could be achieved. Also in a more general case where the polarization is also taken into account, like the case studied by [104], in order to reduce the computation costs the unknowns are split into smaller sub–sets and each set is updated in an alternating fashion.

In this chapter we will show that the matrices used during the estimation process posses a strong Khatri–Rao structure which, combined with Krylov subspace based methods, like minresQLP [78], can be exploited to reduce the computational costs and achieve accurate results with low complexity and fast convergence rate, without using an alternating approach. Another advantage of the proposed method is a tremendous reduction in the memory usage which could be desired in some applications.

This chapter is organized in the following sections. In Sec. 6.2 we will introduce the data model then in Sec. 6.3 we will present the algorithm and show how the structure

1. Results presented in this chapter are partially published in [34].

of the Jacobian can be used to achieve the final result with low complexity, in Sec 6.4 we show some numerical results obtained from experimental data produced by the LOFAR radio telescope.

6.2. DATA MODEL AND PROBLEM DEFINITION

An array of P elements with known locations, $\boldsymbol{\xi}_p = [x_p, y_p, z_p]^T$, $p = 1, \dots, P$, is exposed to Q calibration sources with known spatial coordinates \mathbf{k}_q , $q = 1, \dots, Q$. The elements of the steering matrix, \mathbf{A} consist of the geometric phase delays,

$$a_{p,q} = \frac{1}{\sqrt{P}} e^{j \frac{2\pi}{\lambda} \boldsymbol{\xi}_p^T \mathbf{k}_q}, \quad (6.1)$$

where λ is the wavelength and we have assumed that the narrowband assumption holds such that the delays translate into phase changes. Now we will stack the received signal from each antenna into a vector called $\mathbf{y}(t)$ and obtain the following data model

$$\mathbf{y}(t) = \mathbf{G}\mathbf{A}\mathbf{s}(t) + \mathbf{n}(t), \quad (6.2)$$

where $\mathbf{G} = \text{diag}(\mathbf{g})$ is a diagonal matrix modeling the gain of each element, \mathbf{s} is a $Q \times 1$ vector representing the signal from calibration sources and \mathbf{n} is a $P \times 1$ vector representing the noise in the system. This is a commonly used model for array calibration [4].

We assume that the noise and the sources are independent Gaussian processes which allows us to write the model for the covariance matrix of the array vector as

$$\mathbf{R} = \mathcal{E}\{\mathbf{y}\mathbf{y}^H\} = \mathbf{G}\mathbf{A}\mathbf{R}_s\mathbf{A}^H\mathbf{G}^H + \mathbf{R}_n, \quad (6.3)$$

where $\mathbf{R}_s = \text{diag}(\boldsymbol{\sigma})$ is the covariance matrix of the sources, $\boldsymbol{\sigma}$ is a $Q \times 1$ vector which represent the power of each calibration source and \mathbf{R}_n is the covariance matrix of the system and the sky noise.

Using the above model, we will now define the calibration problem. We want to estimate \mathbf{g} , $\boldsymbol{\sigma}$ and the unknown parameters in \mathbf{R}_n , when we have available to us a sample covariance matrix measured from N samples that is defined as

$$\hat{\mathbf{R}} = \frac{1}{N} \sum_n \mathbf{y}[n]\mathbf{y}[n]^H, \quad (6.4)$$

where $\mathbf{y}[n] := \mathbf{y}(nT_s)$ and T_s is the sampling period. This formulation of the problem leads to covariance matching techniques like the one described in [37] and discussed in Sec 2.3.

6.3. ALGORITHM

6.3.1. COVARIANCE MATCHING

It is desirable to find statistically efficient estimates for the calibration problem. The Maximum Likelihood methods (ML) are in this case very popular. For these methods the cost function is the log-likelihood. For N samples this function is given by

$$l(\boldsymbol{\theta}) = N \left[-\log |\pi^P| + \log |\mathbf{R}^{-1}(\boldsymbol{\theta})| - \text{tr}(\mathbf{R}^{-1}(\boldsymbol{\theta})\hat{\mathbf{R}}) \right]. \quad (6.5)$$

where $\boldsymbol{\theta}$ is a vector containing all the unknowns which we will define shortly.

However when a large number of samples is available and a suitable weighting is applied optimal results can also be found using WLS [37]. In the case of WLS we have

$$f_{WLS}(\boldsymbol{\theta}) = \|\mathbf{W}^{1/2}(\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta}))\mathbf{W}^{1/2}\|_F^2. \quad (6.6)$$

The WLS reduces to LS if $\mathbf{W} = \mathbf{I}_P$. To have a solution that approaches ML asymptotically we choose, $\mathbf{W} = \hat{\mathbf{R}}^{-1}$ [37]. One method for finding a $\boldsymbol{\theta}$ that minimizes/maximizes these functions is the descend algorithm where the solution is updated as

$$\boldsymbol{\theta}_{(k+1)} = \boldsymbol{\theta}_{(k)} + \mu_{(k)} \boldsymbol{\Delta}_{(k)}, \quad (6.7)$$

where $\boldsymbol{\Delta}$ is the direction of descent and μ is the step size. At each iteration we need to find the direction of descent using the Jacobian of our cost function. As was discussed in Sec. 2.3.4 finding $\boldsymbol{\Delta}$ can be achieved by solving the following system at each iteration

$$\mathbf{J}^H (\mathbf{W}^* \otimes \mathbf{W}) \mathbf{J} \boldsymbol{\Delta} = \mathbf{J}^H (\mathbf{W}^* \otimes \mathbf{W}) \text{vect} [\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})], \quad (6.8)$$

where $\mathbf{J} = \frac{\partial \text{vect}(\mathbf{R})}{\partial \boldsymbol{\theta}^T}$ is the Jacobian. For WLS this approach is equivalent to Gauss–Newton and in the case of ML where the optimal weighting matrix is $\mathbf{W}_{(k)} = \mathbf{R}^{-1}(\boldsymbol{\theta}_{(k)})$, this approach is equivalent to the scoring method where

$$\mathbf{F} = \mathbf{J}^H (\mathbf{R}^{-T} \otimes \mathbf{R}^{-1}) \mathbf{J}, \quad (6.9)$$

is the Fisher information matrix [105].

In our application we can define $\boldsymbol{\theta}$ as

$$\boldsymbol{\theta} = [\mathbf{g}^T \quad \mathbf{g}^H \quad \boldsymbol{\sigma}^T \quad \boldsymbol{\sigma}_n^T]^T, \quad (6.10)$$

where $\boldsymbol{\sigma}_n = \mathbf{S}^H \text{vect}(\mathbf{R}_n)$ and \mathbf{S} is a selection matrix. If \mathbf{R}_n is assumed to be a diagonal matrix, then $\mathbf{S} = (\mathbf{I}_P \circ \mathbf{I}_P)$ and $\boldsymbol{\sigma}_n = \text{vectdiag}(\mathbf{R}_n)$.

If we partition $\boldsymbol{\theta}$ in this way we can also partition the Jacobian as

$$\mathbf{J} = [\mathbf{J}_g, \mathbf{J}_{g^*}, \mathbf{J}_\sigma, \mathbf{J}_{\sigma_n}], \quad (6.11)$$

where

$$\begin{aligned} \mathbf{J}_g &= \frac{\partial \text{vect}(\mathbf{R})}{\partial \mathbf{g}^T} = (\mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \mathbf{A}^T \otimes \mathbf{I}_P) (\mathbf{I}_P \circ \mathbf{I}_P) \\ &= \mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \mathbf{A}^T \circ \mathbf{I}_P, \end{aligned} \quad (6.12)$$

$$\begin{aligned} \mathbf{J}_{g^*} &= \frac{\partial \text{vect}(\mathbf{R})}{\partial \mathbf{g}^H} = (\mathbf{I}_P \otimes \mathbf{G} \mathbf{A} \mathbf{R}_s \mathbf{A}^H) (\mathbf{I}_P \circ \mathbf{I}_P) \\ &= \mathbf{I}_P \circ \mathbf{G} \mathbf{A} \mathbf{R}_s \mathbf{A}^H, \end{aligned} \quad (6.13)$$

$$\begin{aligned} \mathbf{J}_\sigma &= \frac{\partial \text{vect}(\mathbf{R})}{\partial \boldsymbol{\sigma}^T} = (\mathbf{G}^* \mathbf{A}^* \otimes \mathbf{G} \mathbf{A}) (\mathbf{I}_P \circ \mathbf{I}_P) \\ &= \mathbf{G}^* \mathbf{A}^* \otimes \mathbf{G} \mathbf{A}, \end{aligned} \quad (6.14)$$

$$\mathbf{J}_{\sigma_n} = \frac{\partial \text{vect}(\mathbf{R})}{\partial \boldsymbol{\sigma}_n^T} = \mathbf{S}. \quad (6.15)$$

In order to solve (6.8) we define the matrix

$$\mathbf{B} = \mathbf{J}^H (\mathbf{W}^* \otimes \mathbf{W}) \mathbf{J}, \quad (6.16)$$

and the gradient vector

$$\mathbf{b} = \mathbf{J}^H (\mathbf{W}^* \otimes \mathbf{W}) \text{vect} [\hat{\mathbf{R}} - \mathbf{R}(\boldsymbol{\theta})], \quad (6.17)$$

such that at each iteration we need to solve $\mathbf{B}\boldsymbol{\Delta} = \mathbf{b}$. In contrast with the rest of this thesis we have used letter \mathbf{b} for the gradient to avoid confusion with the gain vector \mathbf{g} . The dimensions of \mathbf{B} depend on the number of unknowns

$$n = 2P + Q + \|\text{vect}(\mathbf{S})\|_1 \quad (6.18)$$

and for a large array it could become a problem to store it in memory. Also because we cannot assume any sparse, circular or Toeplitz structure in \mathbf{B} , except that it is Hermitian, solving this problem with a direct method has a cubic complexity which must be repeated at each iteration. However if we use a solver based on the Krylov subspace method we can use the Khatri–Rao structure of the Jacobian matrices. This will reduce the complexity and memory usage which is important for very large arrays like SKA.

6

6.3.2. KRYLOV SUBSPACE BASED METHODS

As discussed in Sec. 2.4 the Krylov subspace based methods solve $\mathbf{B}\boldsymbol{\Delta} = \mathbf{b}$ by using matrix–vector products of the form $\mathbf{B}\mathbf{v}$ repeatedly. If this operation can be performed in an efficient way, then application of these methods are preferred to other methods. Especially because we can define a procedure that performs the matrix–vector product, the matrix \mathbf{B} does not need to be stored in memory. This makes Krylov subspace based methods very suitable for situations where \mathbf{B} is very large. We will now show how this matrix vector product can be performed in an efficient way. We have chosen minresQLP because it is capable of handling singular matrices which adds robustness during the iterations [78].

We will split the operation of $\mathbf{B}\mathbf{v}$ into three steps. First we will calculate two intermediate results

$$\mathbf{c} = \mathbf{J}\mathbf{v}, \quad (6.19)$$

and $\mathbf{c}_W = (\mathbf{W}^* \otimes \mathbf{W})\mathbf{c}$. Using these intermediate results we then calculate the final result

$$\mathbf{B}\mathbf{v} = \mathbf{J}^H \mathbf{c}_W. \quad (6.20)$$

Now we will show how each of these steps use the Khatri–Rao structure and can be done efficiently. In order to calculate \mathbf{c} we partition \mathbf{v} , in the same way we have partitioned $\boldsymbol{\theta}$, as

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}_g^T & \mathbf{v}_{g^*}^T & \mathbf{v}_\sigma^T & \mathbf{v}_{\sigma_n}^T \end{bmatrix}^T, \quad (6.21)$$

then we have

$$\mathbf{c} = \mathbf{J}\mathbf{v} = \mathbf{J}_g \mathbf{v}_g + \mathbf{J}_{g^*} \mathbf{v}_{g^*} + \mathbf{J}_\sigma \mathbf{v}_\sigma + \mathbf{J}_{\sigma_n} \mathbf{v}_{\sigma_n}. \quad (6.22)$$

If we unvectorize both sides and substitute the definition of each Jacobian we find

$$\begin{aligned}
\mathbf{C} &= \text{unvect}(\mathbf{c}) \\
&= \text{diag}(\mathbf{v}_g) \mathbf{A} \mathbf{R}_s \mathbf{A}^H \mathbf{G}^H + \mathbf{G} \mathbf{A} \mathbf{R}_s \mathbf{A}^H \text{diag}(\mathbf{v}_{g^*}) \\
&\quad + \mathbf{G} \mathbf{A} \mathbf{V}_{\sigma} \mathbf{A}^H \mathbf{G}^H + \mathbf{V}_{\sigma_n} \\
&= (\mathbf{v}_g \mathbf{g}^H + \mathbf{g} \mathbf{v}_{g^*}^T) \odot \mathbf{R}_0 + \mathbf{G} \mathbf{A} \mathbf{V}_{\sigma} \mathbf{A}^H \mathbf{G}^H + \mathbf{V}_{\sigma_n},
\end{aligned} \tag{6.23}$$

where \odot is the Hadamard or element-wise multiplication, $\mathbf{R}_0 = \mathbf{A} \mathbf{R}_s \mathbf{A}^H$, $\mathbf{V}_{\sigma} = \text{diag}(\mathbf{v}_{\sigma})$ and $\mathbf{V}_{\sigma_n} = \text{unvect}(\mathbf{S} \mathbf{v}_{\sigma_n})$. If \mathbf{R}_n is a diagonal matrix then $\mathbf{V}_{\sigma_n} = \text{diag}(\mathbf{v}_{\sigma_n})$.

Because \mathbf{G} and all of the \mathbf{V} matrices are diagonal, the computational complexity for calculating \mathbf{C} is very low. Also because \mathbf{C} has the same dimensions as a covariance matrix calculating $\mathbf{C}_W = \text{unvect}(\mathbf{c}_W)$ becomes simply

$$\mathbf{C}_W = \mathbf{W} \mathbf{C} \mathbf{W}^H. \tag{6.24}$$

Finally we need to calculate

$$\mathbf{B} \mathbf{v} = \mathbf{J}^H \mathbf{c}_W = \begin{bmatrix} [(\mathbf{R}_0^T \mathbf{G}) \odot \mathbf{C}_W] \mathbf{1} \\ [\mathbf{C}_W^T \odot (\mathbf{R}_0 \mathbf{G}^H)] \mathbf{1} \\ [(\mathbf{C}_W \mathbf{G} \mathbf{A})^T \odot (\mathbf{A}^H \mathbf{G}^H)] \mathbf{1} \\ \mathbf{S}^H \mathbf{u}_W \end{bmatrix}. \tag{6.25}$$

The first three operations consist of an element-wise multiplication and summation of the columns of each row, which are computationally cheap operations. Only \mathbf{R}_0 and $\mathbf{G} \mathbf{A}$ need to be calculated and saved. The third term can also be calculated efficiently if it is viewed as a beamforming operation done on \mathbf{C}_W with columns of $\mathbf{G} \mathbf{A}$ as the beamformer vectors. The last operation is just a selection operation and if \mathbf{R}_n is diagonal it is equal to $\text{vectdiag}(\mathbf{C}_W)$. If we replace \mathbf{C} by $\hat{\mathbf{R}} - \mathbf{R}$ in (6.24), the same procedure can be used to calculate \mathbf{b} .

In conclusion, to calculate $\mathbf{B} \mathbf{v}$ we perform (6.23), (6.24) and (6.25). The procedure that does these operations is given to the minresQLP along with the gradient \mathbf{b} to produce Δ . Computationally this means that per iteration we have a $O(P^2 Q)$ complexity to update all the unknowns and because of the fast convergence of the Newton based method a few number of iterations are needed to converge to an accurate solution. This comparable to other alternating methods that have a $O(P^2)$ complexity such as [106] and [104], however the convergence rate of Newton based methods is in general faster than alternating algorithms. For ML, \mathbf{R} needs to be inverted at each iteration which increases the complexity, however if \mathbf{R}_n is diagonal the inversion can be reduced from $O(P^3)$ to $O(Q^3)$ using the Woodbury matrix identity [107].

6.3.3. IMPROVING THE DOA ESTIMATES

In many cases the direction of arrival (DOA) of the sources is only known from previous (usually lower resolution) catalogs or a (residual) dirty image. We can attempt to refine

the estimates and recover the DOA by adding the spatial coordinate of the sources to the vector of unknowns. A popular spatial coordinate system in radio–astronomy is

$$\mathbf{k} = \begin{bmatrix} l \\ m \\ n \end{bmatrix} = \begin{bmatrix} \cos(\phi) \cos(\theta) \\ \cos(\phi) \sin(\theta) \\ \sin(\phi) \end{bmatrix}. \quad (6.26)$$

We can either solve directly for ϕ and θ or solve for l, m subject to constraint that

$$l^2 + m^2 + n^2 = 1.$$

We will discuss the latter choice in more detail.

JACOBIANS FOR IMAGE COORDINATES

In this section we will derive the Jacobians needed for the covariance matching techniques presented in the previous sections. Let the l coordinate for each of the Q sources be stacked in a $Q \times 1$ vector denoted by \mathbf{l} and let \mathbf{m} be defined similarly for their m coordinates. First we will derive the Jacobian for \mathbf{l} and show that it also has a Khatri–Rao structure. The Jacobian for \mathbf{m} is derived in similar way.

$$\begin{aligned} \mathbf{J}_l &= \frac{\partial \text{vect} \mathbf{R}}{\partial \mathbf{l}^T} \\ &= \frac{\partial \text{vect} \mathbf{R}}{\partial \text{vect}^T(\mathbf{A})} \frac{\partial \text{vect} \mathbf{A}}{\partial \mathbf{l}^T} + \frac{\partial \text{vect} \mathbf{R}}{\partial \text{vect}^T(\mathbf{A}^*)} \frac{\partial \text{vect}(\mathbf{A}^*)}{\partial \mathbf{l}^T} \end{aligned} \quad (6.27)$$

Next, we need to find the derivatives of the steering vectors with respect to the spatial coordinates. It follows that

$$\frac{\partial \text{vect} \mathbf{A}}{\partial \mathbf{l}^T} = j \frac{2\pi}{\lambda} (\mathbf{I}_Q \circ \mathbf{X} \mathbf{A} - \mathbf{I}_Q \circ \mathbf{Z} \mathbf{A} \mathbf{L}) \quad (6.28)$$

and

$$\frac{\partial \text{vect} \mathbf{A}^*}{\partial \mathbf{l}^T} = -j \frac{2\pi}{\lambda} (\mathbf{I}_Q \circ \mathbf{X} \mathbf{A}^* - \mathbf{I}_Q \circ \mathbf{Z} \mathbf{A}^* \mathbf{L}) \quad (6.29)$$

where $\mathbf{X} = \text{diag}([x_1, x_2, \dots, x_P]^T)$ is a diagonal matrix with its entries the x –coordinates of the receiving elements, $\mathbf{Z} = \text{diag}([z_1, z_2, \dots, z_P]^T)$ is defined similarly for the z –coordinates of the receivers and

$$\mathbf{L} = \begin{bmatrix} \frac{l_1}{\sqrt{1-l_1^2-m_1^2}} & & \\ & \ddots & \\ & & \frac{l_Q}{\sqrt{1-l_Q^2-m_Q^2}} \end{bmatrix}.$$

We also have

$$\frac{\partial \text{vect} \mathbf{R}}{\partial \text{vect}^T(\mathbf{A})} = \mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \otimes \mathbf{G} \quad (6.30)$$

$$\frac{\partial \text{vect} \mathbf{R}}{\partial \text{vect}^T(\mathbf{A}^*)} = (\mathbf{G}^* \otimes \mathbf{G} \mathbf{A} \mathbf{R}_s) \mathbf{K}^{PQ}. \quad (6.31)$$

Combining these results we find

$$\begin{aligned}
 \mathbf{J}_1 &= -j \frac{2\pi}{\lambda} [-(\mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \otimes \mathbf{G})(\mathbf{I}_Q \circ \mathbf{X}\mathbf{A} - \mathbf{I}_Q \circ \mathbf{Z}\mathbf{A}\mathbf{L}) + (\mathbf{G}^* \otimes \mathbf{G}\mathbf{A}\mathbf{R}_s) \mathbf{K}_{p,q}(\mathbf{I}_Q \circ \mathbf{X}\mathbf{A}^* - \mathbf{I}_Q \circ \mathbf{Z}\mathbf{A}^* \mathbf{L})] \\
 &= -j \frac{2\pi}{\lambda} [(\mathbf{G}^* \mathbf{X}\mathbf{A}^* \circ \mathbf{G}\mathbf{A}\mathbf{R}_s) - (\mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \circ \mathbf{G}\mathbf{X}\mathbf{A}) - (\mathbf{G}^* \mathbf{Z}\mathbf{A}^* \mathbf{L} \circ \mathbf{G}\mathbf{A}\mathbf{R}_s) + (\mathbf{G}^* \mathbf{A}^* \mathbf{R}_s \circ \mathbf{G}\mathbf{Z}\mathbf{A}\mathbf{L})] \\
 &= -j \frac{2\pi}{\lambda} [(\mathbf{G}^* \mathbf{X}\mathbf{A}^* \circ \mathbf{G}\mathbf{A}) - (\mathbf{G}^* \mathbf{A}^* \circ \mathbf{G}\mathbf{X}\mathbf{A}) - (\mathbf{G}^* \mathbf{Z}\mathbf{A}^* \mathbf{L} \circ \mathbf{G}\mathbf{A}) + (\mathbf{G}^* \mathbf{A}^* \circ \mathbf{G}\mathbf{Z}\mathbf{A}\mathbf{L})] \mathbf{R}_s \\
 &= -j \frac{2\pi}{\lambda} [\mathbf{G}^* (\mathbf{X}\mathbf{A}^* - \mathbf{Z}\mathbf{A}^* \mathbf{L}) \circ \mathbf{G}\mathbf{A} - \mathbf{G}^* \mathbf{A}^* \circ \mathbf{G}(\mathbf{X}\mathbf{A} - \mathbf{Z}\mathbf{A}\mathbf{L})] \mathbf{R}_s
 \end{aligned} \tag{6.32}$$

where we have used the identity $\mathbf{K}^{P,Q}(\mathbf{I}_Q \circ \mathbf{B}) = (\mathbf{B} \circ \mathbf{I}_Q)$. In the same way we can find

$$\mathbf{J}_m = -j \frac{2\pi}{\lambda} [\mathbf{G}^* (\mathbf{Y}\mathbf{A}^* - \mathbf{Z}\mathbf{A}^* \mathbf{M}) \circ \mathbf{G}\mathbf{A} - \mathbf{G}^* \mathbf{A}^* \circ \mathbf{G}(\mathbf{Y}\mathbf{A} - \mathbf{Z}\mathbf{A}\mathbf{M})] \mathbf{R}_s \tag{6.33}$$

where $\mathbf{Y} = \text{diag}([y_1, y_2, \dots, y_P]^T)$ contains the y -coordinates of the array elements and

$$\mathbf{M} = \begin{bmatrix} \frac{m_1}{\sqrt{1-l_1^2-m_1^2}} & & \\ & \ddots & \\ & & \frac{m_q}{\sqrt{1-l_Q^2-m_Q^2}} \end{bmatrix}$$

If an accurate initial guess for l_Q and m_Q is known, the estimation algorithm in the previous section can be adapted by adding the terms $\mathbf{J}_1 \mathbf{v}_1$ and $\mathbf{J}_m \mathbf{v}_m$ to \mathbf{c} as defined in (6.22) where we can use the properties of the Khatri–Rao product in the same way we have done to calculate (6.23) and (6.25). This DOA correction improves the accuracy of the calibration techniques by providing a more accurate array response matrix \mathbf{A} .

6.4. EXPERIMENTAL DATA

We have used a measurement set from the LOFAR radio telescope to test our method. The sample covariance matrix of $P = 273$ dipoles (from 6 LOFAR inner core stations) is available to us from a single channel with a central frequency of 58.98 MHz and a bandwidth of 195 kHz which is sampled at the Nyquist rate. The integration time for this covariance matrix is 1 second.

The proposed method is used to calibrate this array for three different cases. In each case we assume to know the position of one, two and three sources such that $Q = 1, 2$ and 3. For each case we look at the norm for the residual \mathbf{E} defined as

$$\|\mathbf{E}\|_F = \|\hat{\mathbf{R}} - \mathbf{R}(\hat{\boldsymbol{\theta}})\|_F,$$

and the norm of the gradient defined by (6.17). These results are illustrated in Figure 6.1a and 6.1b. From the norm of the gradient we know that in all three cases the algorithm

This data is the courtesy of ASTRON (<http://www.astron.nl>) and is provided to the authors as part of a collaboration within the NWO TOP project. We also would like to acknowledge the help of Stefan Wijnholds and Peeyush Prasad for obtaining this data.

converges after a few iterations. However, for $Q = 1$ the final residual is much higher than the other two and it takes the algorithm much longer to converge. In order to explain this we use the calibration results to make a full sky image (Figure 6.1c). In this image we observe that two strong point sources, Cygnus A and Cassiopeia A, are visible to the array and also there is strong background radiation from the Milky-Way (going from top to the bottom of the image). In order to obtain better calibration results we use the fact that the extended sources like the Milky-Way mostly affect the shorter baselines [2]. Based on this knowledge we have chosen the selection matrix \mathbf{S} to include the baselines smaller than 25 times the wavelength into \mathbf{R}_n as part of the noise covariance structure. In this case the total number of unknowns, n , is 23393. The proposed method has a complexity of $O(P^2Q)$ while a naive and direct approach to the Gauss–Newton algorithm would have had a complexity of $O(n^3)$ which would be extremely expensive.

By adding Cassiopeia A and the Milky-Way to the model we have achieved a much better result for $Q = 2$. In order to verify this we also made Figure 6.1d by imaging the residual, \mathbf{E} , which shows how weaker sources (three orders of magnitude lower than the strong sources) can now be detected from this residual.



Finally, for the case $Q = 3$ we have added the next brightest source to the model, in this case we also converge and have a smaller residual. However because the sources are much weaker the difference is harder to visualize.

6


6.5. CONCLUSION

We have shown that the covariance data model for calibrating a radio telescope has a strong Khatri–Rao structure. We have shown that this structure could be used to perform fast matrix–vector computations which is the building block of the Krylov subspace based methods. Finally we have used the proposed method to calibrate real measurement set from the LOFAR radio telescope with low computational complexity.

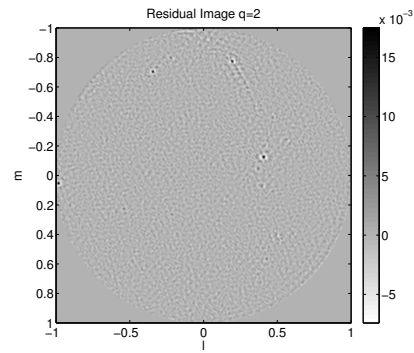
After we estimate the gains we can correct for them by multiplying both side of $\hat{\mathbf{R}}$ by \mathbf{G}^{-1} . The result is then the calibrated version of $\hat{\mathbf{R}}$ for direction independent gains. In Chapter 7 we assume that this step has already been performed and that we have access to calibrated data.

(a) Residual norm for $Q = 1, 2$ and 3 


(b) Convergence of the gradient to zero




(c) Full sky image after gain correction



(d) Full sky image after removal of Cas A and Cyg A



(e) Cyg A without calibration



(f) Cyg A with calibration

Figure 6.1: Results from calibrating LOFAR data

7

IMAGE RECONSTRUCTION

7.1. INTRODUCTION

Image formation for radio astronomy can be defined as estimating the spatial intensity distribution of celestial sources over the sky. The measurement equation (“data model”) is linear in the source intensities, and the resulting least squares problem has classically been implemented in two steps: formation of a “dirty image”, followed by a deconvolution step. In this process, an implicit model assumption is made that the number of sources is discrete, and subsequently the number of sources has been replaced by the number of image pixels (assuming each pixel may contain a source).

The deconvolution step becomes ill-conditioned if the number of pixels is large [9]. Alternatively, the directions of sources may be estimated along with their intensities, but this is a complex non-linear problem. Classically, this has been implemented as an iterative subtraction technique, wherein source directions are estimated from the dirty image, and their contribution is subtracted from the data. This mixed approach is the essence of the CLEAN method proposed by Högbom [16], which was subsequently refined and extended in several ways, leading to the widely used approaches described in [17–19].

The conditioning of the image deconvolution step can be improved by incorporating side information such as non-negativity of the image [20], source model structure beyond simple point sources (e.g., shapelets and wavelets [108]), sparsity or ℓ_1 constraints on the image [8, 21] or a combination of both wavelets and sparsity [22, 23]. Beyond these, some fundamental approaches based on parameter estimation techniques have been proposed, such as the Least Squares Minimum Variance Imaging (LS-MVI) [24], maximum likelihood based techniques [25] and Bayesian based techniques [109, 110]. Computational complexity is a concern and this has not been addressed in these approaches.

1. Results presented in this chapter are published in [29].

New radio telescopes such as the Low Frequency Array (LOFAR), the Allen Telescope Array (ATA), Murchison Widefield Array (MWA) and the Long Wavelength Array (LWA) are composed of many stations (each station made up of multiple antennas that are combined using adaptive beamforming), and the increase in number of antennas and stations continues in the design of the square kilometer array (SKA). These instruments have or will have a significantly increased sensitivity and a larger field of view compared to traditional telescopes, leading to many more sources that need to be taken into account. They also need to process larger bandwidths to reach this sensitivity. Besides the increased requirements on the performance of imaging, the improved spatial resolution leads to an increasing number of pixels in the image, and the development of computationally efficient techniques is critical.

To benefit from the vast literature related to solving least square problems, but also to gain from the non-linear processing offered by standard deconvolution techniques, we propose to reformulate the imaging problem as a parameter estimation problem described by a weighted least squares optimization problem with several constraints. The first is a non-negativity constraint, which would lead to the non-negative least squares algorithm (NNLS) proposed in [20]. But we show that the pixel values are also bounded from above. A coarse upper bound is provided by the classical dirty image, and a much tighter bound is the “minimum variance distortionless response” (MVDR) dirty image that was proposed in the context of radio astronomy in [25].

We propose to solve the resulting constrained least squares problems using an active set approach. This results in a computationally efficient imaging algorithm that is closely related to existing non-linear sequential source estimation techniques such as CLEAN with the benefit of accelerated convergence due to tighter upper bounds on the intensity over the complete image. Because the constraints are enforced over the entire image, this eliminates the inclusion of negative flux sources and other anomalies that appear in some existing sequential techniques.

To further reduce the computational complexity we show that the data model has a Khatri-Rao structure. This can be exploited to significantly improve the data management and parallelism compared to general implementations of least squares algorithms.

Using the model introduced in Sec. 3.1 we describe the image formation problem in Sec. 7.2. A constrained least squares problem is formulated, using various intensity constraints that take the form of dirty images. The solution of this problem using active set techniques in Sec. 7.4 generalizes the classical CLEAN algorithm. In Sec. 7.5 we discuss the efficient implementation of a key step in the active set solution using Krylov subspaces. We end up with some simulated experiments demonstrating the advantages of the proposed technique and conclusions regarding future implementation.

7.2. THE IMAGING PROBLEM

Using the results from Sec. 3.1 and in particular (3.11), we have the following model for K snapshots

$$\mathbf{r} = \Psi\boldsymbol{\sigma} + \mathbf{r}_n \quad (7.1)$$

where

$$\mathbf{r} = \begin{bmatrix} \text{vect}(\mathbf{R}_1) \\ \vdots \\ \text{vect}(\mathbf{R}_K) \end{bmatrix}, \quad \Psi = \begin{bmatrix} \mathbf{A}_1^* \circ \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_K^* \circ \mathbf{A}_K \end{bmatrix}, \quad \mathbf{r}_n = \begin{bmatrix} \mathbf{r}_{n,1} \\ \vdots \\ \mathbf{r}_{n,K} \end{bmatrix}. \quad (7.2)$$

The imaging problem is to find the intensity, σ , of the sources, along with their directions represented by the matrices \mathbf{A}_k , from given sample covariance matrices $\hat{\mathbf{R}}_k$, $k = 1, \dots, K$. As the source locations are generally unknown, this is a complicated (non-linear) direction-of-arrival estimation problem.

The usual approach in radio astronomy is to define a grid for the image, and to assume that each pixel (grid location) contains a source. In this case the source locations are known, and estimating the source intensities is a linear problem, but for high-resolution images the number of sources may be very large. The resulting linear estimation problem is often ill-conditioned unless additional constraints are posed.

7.2.1. GRIDDED IMAGING MODEL

After defining a grid for the image and assuming that a source exists for each pixel location, let I (rather than Q) denote the total number of sources (pixels), σ an $I \times 1$ vector containing the source intensities, and \mathbf{A}_k ($k = 1, \dots, K$) the $P \times I$ array response matrices for these sources. Note that the \mathbf{A}_k are known, and that σ can be interpreted as a vectorized version of the image to be computed. (To distinguish the gridded source locations and source powers from the “true” sources, we will later denote parameters and variables that depend on the Q true sources by a tilde.)

We can also use the independence between the time samples to write the aggregate data model from Sec. 3.2 as

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{R}_K \end{bmatrix} = \sum_{q=1}^Q \sigma_q (\mathbf{I}_K \circ \mathbf{A}^q) (\mathbf{I}_K \circ \mathbf{A}^q)^H + \mathbf{R}_n, \quad (7.3)$$

where

$$\mathbf{R}_n = \begin{bmatrix} \mathbf{R}_{n,1} & \dots & \mathbf{0} \\ \vdots & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{R}_{n,K} \end{bmatrix}, \quad (7.4)$$

$$\mathbf{A}^q = [\mathbf{a}_{1,q} \quad \dots \quad \mathbf{a}_{K,q}], \quad q = 1, \dots, Q. \quad (7.5)$$

For a given observation $\hat{\mathbf{r}}$, image formation amounts to the estimation of σ . For a sufficiently fine grid, σ approximates the solution of the discrete source model. However, as we will discuss later, working in the image domain leads to a gridding related noise floor. This is solved by fine adaptation of the locations of the sources and estimating the true locations in the visibility domain.

A consequence of using a discrete source model in combination with sequential source removing techniques such as CLEAN is the modeling of extended structures in the image by many point sources. As we will discuss in Sec. 7.6, this also holds for the algorithms proposed in this chapter.

7.2.2. UNCONSTRAINED LEAST SQUARES IMAGE

If we ignore the term \mathbf{r}_n , then (3.11) directly leads to Least Squares (LS) and Weighted Least Squares (WLS) estimates of $\boldsymbol{\sigma}$ [9]. In particular, solving the imaging problem with LS leads to the minimization problem

$$\min_{\boldsymbol{\sigma}} \frac{1}{2K} \|\hat{\mathbf{r}} - \boldsymbol{\Psi}\boldsymbol{\sigma}\|^2, \quad (7.6)$$

where the normalization factor $2K$ is introduced to simplify the expression for the gradient and does not affect the solution. It is straightforward to show that the solution to this problem is given by any $\boldsymbol{\sigma}$ that satisfies

$$\mathbf{H}_{\text{LS}}\boldsymbol{\sigma} = \hat{\boldsymbol{\sigma}}_{\text{MF}} \quad (7.7)$$

where we define the “matched filter” (MF, also known as the classical “direct Fourier transform dirty image”) as

$$\hat{\boldsymbol{\sigma}}_{\text{MF}} = \frac{1}{K} \boldsymbol{\Psi}^H \hat{\mathbf{r}} = \frac{1}{K} \sum_k \text{vectdiag}(\mathbf{A}_k^H \hat{\mathbf{R}}_k \mathbf{A}_k), \quad (7.8)$$

and the deconvolution matrix \mathbf{H}_{LS} as

$$\mathbf{H}_{\text{LS}} = \frac{1}{K} \boldsymbol{\Psi}^H \boldsymbol{\Psi} = \frac{1}{K} \sum_k (\mathbf{A}_k^T \mathbf{A}_k^*) \odot (\mathbf{A}_k^H \mathbf{A}_k), \quad (7.9)$$

where we have used the definition of $\boldsymbol{\Psi}$ from (3.12) (with tilde removed) and properties of the Kronecker and Khatri-Rao products. Similarly we can define the WLS minimization as

$$\min_{\boldsymbol{\sigma}} \frac{1}{2K} \|\mathbf{W}^{1/2}(\hat{\mathbf{r}} - \boldsymbol{\Psi}\boldsymbol{\sigma})\|^2, \quad (7.10)$$

where

$$\mathbf{W} = \text{bdiag}(\mathbf{R}_k^{-T} \otimes \mathbf{R}_k^{-1}) = \begin{bmatrix} \mathbf{R}_1^{-T} \otimes \mathbf{R}_1^{-1} & & \\ & \ddots & \\ & & \mathbf{R}_K^{-T} \otimes \mathbf{R}_K^{-1} \end{bmatrix}$$

and the weighting assumes Gaussian distributed observations. The weighting improves the statistical properties of the estimates, and $\hat{\mathbf{R}}$ is used instead of \mathbf{R} because it is available and gives asymptotically the same optimal results, i.e., convergence to maximum likelihood estimates [37]. The solution to this optimization is similar to the solution to the LS problem and is given by any $\boldsymbol{\sigma}$ that satisfies

$$\mathbf{H}_{\text{WLS}}\boldsymbol{\sigma} = \hat{\boldsymbol{\sigma}}_{\text{WLS}}, \quad (7.11)$$

where

$$\hat{\boldsymbol{\sigma}}_{\text{WLS}} = \frac{1}{K} \boldsymbol{\Psi}^H \mathbf{W} \hat{\mathbf{r}} \quad (7.12)$$

is the “WLS dirty image” and

$$\mathbf{H}_{\text{WLS}} = \frac{1}{K} \boldsymbol{\Psi}^H \mathbf{W} \boldsymbol{\Psi} \quad (7.13)$$

is the associated deconvolution operator.

A connection to beamforming is obtained as follows. The i th pixel of the “Matched Filter” dirty image in equation (7.8) can be written as

$$\hat{\sigma}_{\text{MF},i} = \frac{1}{K} \sum_k \mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k \mathbf{a}_{k,i}$$

and if we replace $\mathbf{a}_{k,i}/\sqrt{K}$ by a more general “beamformer” $\mathbf{w}_{k,i}$, this can be generalized to a more general dirty image

$$\sigma_{\mathbf{w},i} = \sum_k \mathbf{w}_{k,i}^H \hat{\mathbf{R}}_k \mathbf{w}_{k,i} \quad (7.14)$$

Here, $\mathbf{w}_{k,i}$ is called a beamformer because we can consider that it acts on the antenna vectors $\mathbf{y}_k[n]$ as $z_{k,i}[n] = \mathbf{w}_{k,i}^H \mathbf{y}_k[n]$, where $z_{k,i}[n]$ is the output of the beamformer, and $\sigma_{\mathbf{w},i} = \sum_k \mathcal{E}\{|z_{k,i}|^2\}$ is interpreted as the total output power of the beamformer, summed over all snapshots. We will encounter several such beamformers in the rest of this chapter. Most of the beamformers discussed here include the weighted visibility vector $\mathbf{W}\mathbf{r}$. The relation between this weighting and more traditional weighting techniques, such as Natural and Robust weighting, is discussed in Sec. 7.8.

7.2.3. PRECONDITIONED WEIGHTED LEAST SQUARES IMAGE

If Ψ has full column rank then \mathbf{H}_{LS} and \mathbf{H}_{WLS} are non-singular and there exists a unique solution to LS and WLS. For example the solution to (7.7) becomes

$$\boldsymbol{\sigma} = \mathbf{H}_{\text{LS}}^{-1} \hat{\boldsymbol{\sigma}}_{\text{MF}}. \quad (7.15)$$

Unfortunately, if the number of pixels is large then \mathbf{H}_{LS} and \mathbf{H}_{WLS} become ill-conditioned or even singular, so that (7.7) and (7.11) have an infinite number of solutions [9]. Generally, we need to improve the conditioning of the deconvolution matrices and to find appropriate regularizations.

One way to improve the conditioning of a matrix is by applying a preconditioner. The most widely used and simplest preconditioner is the Jacobi preconditioner [111] which, for any matrix \mathbf{M} , is given by $[\text{diag}(\mathbf{M})]^{-1}$. Let $\mathbf{D}_{\text{WLS}} = \text{diag}(\mathbf{H}_{\text{WLS}})$, then by applying this preconditioner to \mathbf{H}_{WLS} we obtain

$$[\mathbf{D}_{\text{WLS}}^{-1} \mathbf{H}_{\text{WLS}}] \boldsymbol{\sigma} = \mathbf{D}_{\text{WLS}}^{-1} \hat{\boldsymbol{\sigma}}_{\text{WLS}}. \quad (7.16)$$

We take a closer look at $\mathbf{D}_{\text{WLS}}^{-1} \hat{\boldsymbol{\sigma}}_{\text{WLS}}$ for the case where $K = 1$. In this case

$$\begin{aligned} \mathbf{H}_{\text{WLS}} &= (\mathbf{A}_1^* \circ \mathbf{A}_1)^H (\hat{\mathbf{R}}_1^{-T} \otimes \hat{\mathbf{R}}_1^{-1}) (\mathbf{A}_1^* \circ \mathbf{A}_1) \\ &= (\mathbf{A}^T \hat{\mathbf{R}}_1^{-T} \mathbf{A}_1^*) \odot (\mathbf{A}_1^H \hat{\mathbf{R}}_1^{-1} \mathbf{A}_1) \end{aligned}$$

and

$$\mathbf{D}_{\text{WLS}}^{-1} = \begin{bmatrix} \frac{1}{(\mathbf{a}_{1,1}^H \hat{\mathbf{R}}_1^{-1} \mathbf{a}_{1,1})^2} & & \\ & \ddots & \\ & & \frac{1}{(\mathbf{a}_{1,I}^H \hat{\mathbf{R}}_1^{-1} \mathbf{a}_{1,I})^2} \end{bmatrix}.$$

This means that

$$\begin{aligned}\mathbf{D}_{\text{WLS}}^{-1} \hat{\boldsymbol{\sigma}}_{\text{WLS}} &= \mathbf{D}_{\text{WLS}}^{-1} (\hat{\mathbf{R}}_1^{-T} \otimes \hat{\mathbf{R}}_1^{-1}) (\mathbf{A}_1^* \circ \mathbf{A}_1)^H \hat{\mathbf{r}}_1 \\ &= (\hat{\mathbf{R}}_1^{-T} \mathbf{A}_1^* \mathbf{D}_{\text{WLS}}^{-1/2} \circ \hat{\mathbf{R}}_1^{-1} \mathbf{A}_1 \mathbf{D}_{\text{WLS}}^{-1/2})^H \hat{\mathbf{r}}_1\end{aligned}$$

which is equivalent to a dirty image that is obtained by applying a beamformer of the form

$$\mathbf{w}_i = \frac{1}{\mathbf{a}_{1,i}^H \hat{\mathbf{R}}_1^{-1} \mathbf{a}_{1,i}} \hat{\mathbf{R}}_1^{-1} \mathbf{a}_{1,i} \quad (7.17)$$

to both sides of $\hat{\mathbf{R}}_1$ and stacking the results, $\hat{\boldsymbol{\sigma}}_i = \mathbf{w}_i^H \hat{\mathbf{R}}_1 \mathbf{w}_i$, of each pixel into a vector. This beamformer is known in array processing as the Minimum Variance Distortionless Response (MVDR) beamformer [112], and the corresponding dirty image is called the MVDR dirty image and was introduced in the radio astronomy context in [25]. This shows that preconditioned WLS image (motivated from its connection to the maximum likelihood) is expected to exhibit the features of high-resolution beamforming associated with the MVDR. Examples of such images are shown in Sec. 7.6.

7.3. BOUNDS ON THE IMAGE

Another approach to improve the conditioning of a problem is to introduce appropriate constraints on the solution. Typically, image formation algorithms exploit external information regarding the image in order to regularize the ill-posed problem. For example maximum entropy techniques [113, 114] impose a smoothness condition on the image while the CLEAN algorithm [16] exploits a point source model wherein most of the image is empty, and this has recently been connected to sparse optimization techniques [8].

A lower bound on the image is almost trivial: each pixel in the image represents the intensity at a certain direction, hence is non-negative. This leads to a lower bound $\boldsymbol{\sigma} \geq \mathbf{0}$. Such a non-negativity constraint has been studied for example in [20], resulting in a non-negative LS (NNLS) problem

$$\begin{aligned}\min_{\boldsymbol{\sigma}} \quad & \frac{1}{2K} \|\hat{\mathbf{r}} - \boldsymbol{\Psi} \boldsymbol{\sigma}\|^2 \\ \text{subject to} \quad & \mathbf{0} \leq \boldsymbol{\sigma}\end{aligned} \quad (7.18)$$

A second constraint follows if we also know an upper bound $\boldsymbol{\gamma}$ such that $\boldsymbol{\sigma} \leq \boldsymbol{\gamma}$, which will bound the pixel intensities from above. We will propose several choices for $\boldsymbol{\gamma}$.

7.3.1. MATCHED FILTER AS UPPER BOUND

By closer inspection of the i th pixel of the MF dirty image $\hat{\boldsymbol{\sigma}}_{\text{MF}}$, we note that its expected value is given by

$$\sigma_{\text{MF},i} = \frac{1}{K} \sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k \mathbf{a}_{k,i}.$$

Using

$$\mathbf{a}_i = \text{vect}(\mathbf{A}^i) = \begin{bmatrix} \mathbf{a}_{1,i}^T & \dots & \mathbf{a}_{I,K}^T \end{bmatrix}^T, \quad (7.19)$$

and the normalization $\mathbf{a}_{k,i}^H \mathbf{a}_{k,i} = 1$, we obtain

$$\sigma_{\text{MF},i} = \frac{1}{K} \mathbf{a}_i^H \mathbf{R} \mathbf{a}_i = \sigma_i + \frac{1}{K} \mathbf{a}_i^H \mathbf{R}_r \mathbf{a}_i, \quad (7.20)$$

where

$$\mathbf{R}_r = \sum_{j \neq i} \sigma_j (\mathbf{I}_K \circ \mathbf{A}^j) (\mathbf{I}_K \circ \mathbf{A}^j)^H + \mathbf{R}_n \quad (7.21)$$

is the contribution of all other sources and the noise. Note that \mathbf{R}_r is positive-(semi)definite. Thus, (7.20) implies $\sigma_{\text{MF},i} \geq \sigma_i$ which means that the expected value of the MF dirty image forms an upper bound for the desired image, or

$$\sigma \leq \sigma_{\text{MF}}. \quad (7.22)$$

As indicated in Sec. 7.2.2, we can extend this concept to a more general beamformer \mathbf{w}_i . The output power of this beamformer, in the direction of the i th pixel, becomes

$$\sigma_{\mathbf{w},i} = \mathbf{w}_i^H \mathbf{R} \mathbf{w}_i = \sigma_i \mathbf{w}_i^H (\mathbf{I}_K \circ \mathbf{A}^i) (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{w}_i + \mathbf{w}_i^H \mathbf{R}_r \mathbf{w}_i. \quad (7.23)$$

If we require that

$$\mathbf{w}_i^H (\mathbf{I}_K \circ \mathbf{A}^i) (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{w}_i = 1 \quad (7.24)$$

we have

$$\sigma_{\mathbf{w},i} = \sigma_i + \mathbf{w}_i^H \mathbf{R}_r \mathbf{w}_i. \quad (7.25)$$

As before, the fact that \mathbf{R}_r is positive definite implies that

$$\sigma_i \leq \sigma_{\mathbf{w},i}. \quad (7.26)$$

We can easily verify that $\mathbf{w}_{\text{MF},i}$ satisfies (7.24) and hence $\sigma_{\text{MF},i}$ is a specific upper bound.

7.3.2. TIGHTEST UPPER BOUND AND ASSC

Using the relation between the MF dirty image and beamformers as discussed in Sec. 7.2.2 we will answer the following question: What is the tightest upper bound for σ_i that we can construct using linear beamforming?

We can translate the problem of finding the tightest upper bound to the following optimization question:

$$\begin{aligned} \sigma_{\text{opt},i} &= \min_{\mathbf{w}_i} \mathbf{w}_i^H \mathbf{R} \mathbf{w}_i \\ \text{s.t. } &\mathbf{w}_i^H (\mathbf{I}_K \circ \mathbf{A}^i) (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{w}_i = 1 \end{aligned} \quad (7.27)$$

where $\sigma_{\text{opt},i}$ would be this tightest upper bound.

To solve this optimization problem we follow standard optimization techniques and define the Lagrangian and take derivatives with respect to \mathbf{w} and the Lagrange multiplier μ . This leads to the following system

$$\mathbf{w} = \mu \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{w} \quad (7.28)$$

$$1 = \mathbf{w}^H (\mathbf{I}_K \circ \mathbf{A}_i) (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{w} \quad (7.29)$$

Because \mathbf{R} is full-rank and (7.29) we can model \mathbf{w} as

$$\mathbf{w} = \mu \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x}. \quad (7.30)$$

Filling back into (7.28) we have

$$\begin{aligned} & \mu \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x} \\ &= \mu^2 \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x} \end{aligned} \quad (7.31)$$

and

$$\begin{aligned} & (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x} \\ &= \mu (\mathbf{I}_K \circ \mathbf{A}_i) (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x} \end{aligned} \quad (7.32)$$

multiplying both sides by $(\mathbf{I}_K \circ \mathbf{A}_i)^H$ we get

$$\mathbf{x} = \mu (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x}. \quad (7.33)$$

Doing the same for (7.29) we have

$$\begin{aligned} & \mu^2 \mathbf{x}^H (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \mathbf{x} \\ &= 1. \end{aligned} \quad (7.34)$$

Now we use (7.33) and we find

$$\mathbf{x}^H \mathbf{x} = 1 \quad (7.35)$$

which makes finding \mathbf{x} an eigenvalue problem. By taking a closer look at the matrix $(\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i)$ we find that this matrix is diagonal

$$\begin{aligned} & (\mathbf{I}_K \circ \mathbf{A}_i)^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}_i) \\ &= \begin{bmatrix} \mathbf{a}_{1,i}^H \mathbf{R}_1^{-1} \mathbf{a}_{1,i} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{a}_{2,i}^H \mathbf{R}_2^{-1} \mathbf{a}_{2,i} & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \mathbf{a}_{K,i}^H \mathbf{R}_K^{-1} \mathbf{a}_{K,i} \end{bmatrix} \end{aligned} \quad (7.36)$$

and hence $\mathbf{x} = \mathbf{e}_m$ is an elementary vector with all entries equal to zero except for m th entry which equals unity. m is the index corresponding to largest eigenvalue, λ_{\max} , and from (7.33) we have $\mu = 1/\lambda_{\max}$. Filling back for \mathbf{w} we find

$$\mathbf{w}_{i,\text{opt}} = \frac{1}{\mathbf{a}_{m,i}^H \mathbf{R}_m^{-1} \mathbf{a}_{m,i}} \mathbf{R}^{-1} (\mathbf{e}_m \otimes \mathbf{a}_{m,i}) \quad (7.37)$$

and the output of the beamformer

$$\begin{aligned} \sigma_{\text{opt}} &= \mathbf{w}_{i,\text{opt}}^H \mathbf{R} \mathbf{w}_{i,\text{opt}} \\ &= \frac{\mathbf{a}_{m,i}^H \mathbf{R}_m^{-1} \mathbf{a}_{m,i}}{(\mathbf{a}_{m,i}^H \mathbf{R}_m^{-1} \mathbf{a}_{m,i})^2} \\ &= \frac{1}{\mathbf{a}_{i,m}^H \mathbf{R}_m^{-1} \mathbf{a}_{i,m}} \\ &= \min_k \left(\frac{1}{\mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i}} \right) \end{aligned} \quad (7.38)$$

Here $\sigma_{\text{opt},i}$ is the tightest upper bound and the beamformer that achieves this bound is called the adaptive selective sidelobe canceller (ASSC) [115].

One problem with using this result in practice is that $\sigma_{\text{opt},i}$ depends on a single snapshot. Actual dirty images are based on the sample covariance matrix $\hat{\mathbf{R}}$ and hence they are random variables. If we use a sample covariance matrix $\hat{\mathbf{R}}$ instead of the true covariance matrix \mathbf{R} in (7.38), the variance of the result can be unacceptably large. An analysis of this problem and various solutions for it are discussed in [115].

7.3.3. MULTI-SNAPSHOT MVDR

We would like to find a beamformer that exhibits the same averaging behavior as MF beamformer while being as tight as possible. To this end we will tolerate an increase of the bound with respect to the tightest (resulting from ASSC beamformer). We suggest to find a beamformer that instead of (7.24) satisfies the slightly different normalization constraint

$$\mathbf{w}_i^H \mathbf{a}_i = \sqrt{K}. \quad (7.39)$$

We will show that the expected value of the resulting dirty image constitutes a larger upper bound than the ASSC (7.38), but because the output power of this beamformer depends on more than one snapshot it will have a lower variance than ASSC, so that it is more robust in practice.

With this constraint, the beamforming problem is

$$\begin{aligned} \mathbf{w}_i &= \arg \min_{\mathbf{w}_i} \mathbf{w}_i^H \mathbf{R} \mathbf{w}_i \\ \text{s.t. } \mathbf{w}_i^H \mathbf{a}_i &= \sqrt{K} \end{aligned} \quad (7.40)$$

which is recognized as the classical MVDR beamforming problem [112]. Thus, the solution is given in closed form as

$$\mathbf{w}_{\text{MVDR},i} = \frac{\sqrt{K}}{\mathbf{a}_i^H \mathbf{R}^{-1} \mathbf{a}_i} \mathbf{R}^{-1} \mathbf{a}_i \quad (7.41)$$

To demonstrate that this image is still an upper bound we show that

$$\alpha := \mathbf{w}_i^H (\mathbf{I}_K \circ \mathbf{A}^i) (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{w}_i \geq 1. \quad (7.42)$$

Indeed, inserting (7.41) into this inequality gives

$$\begin{aligned} & K \frac{\mathbf{a}_i^H \mathbf{R}^{-1} (\mathbf{I}_K \circ \mathbf{A}^i) (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{R}^{-1} \mathbf{a}_i}{(\mathbf{a}_i^H \mathbf{R}^{-1} \mathbf{a}_i)^2} \\ &= K \frac{\sum_k (\mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i})^2}{\left(\sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i} \right)^2} \\ &= K \frac{\mathbf{h}^T \mathbf{h}}{\mathbf{h}^T \mathbf{1}_K \mathbf{1}_K^T \mathbf{h}} \geq K \frac{1}{\lambda_{\max}(\mathbf{1}_K \mathbf{1}_K^T)} = 1, \end{aligned} \quad (7.43)$$

where $\mathbf{h} = (\mathbf{I}_K \circ \mathbf{A}^i)^H \mathbf{R}^{-1} \mathbf{a}_i$ is a $K \times 1$ vector with entries $h_k = \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i}$ and $\lambda_{\max}(\cdot)$ is the largest eigenvalue of the argument matrix. Hence, a similar reasoning as in (7.23) gives

$$\sigma_{\text{MVDR},i} = \alpha \sigma_i + \mathbf{w}_{\text{MVDR},i}^H \mathbf{R}_r \mathbf{w}_{\text{MVDR},i} \geq \sigma_i.$$

Note that $\mathbf{w}_{\text{MF},i}$ also satisfies the constraint in (7.40), i.e. $\mathbf{w}_{\text{MF},i}^H \mathbf{a}_i = \sqrt{K}$, but does not necessarily minimize the output power $\mathbf{w}_i^H \mathbf{R} \mathbf{w}_i$, therefore the MVDR dirty image is smaller than the MF dirty image: $\sigma_{\text{MVDR}} \leq \sigma_{\text{MF}}$. Thus it is a tighter upper bound. This relation also holds if \mathbf{R} is replaced by the sample covariance $\hat{\mathbf{R}}$.

$$\sigma_{\text{MVDR},i} = \frac{1}{\frac{1}{K} \sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i}}, \quad (7.44)$$

satisfies $\sigma_i \leq \sigma_{\text{MVDR},i} \leq \sigma_{\text{MF},i}$ and produces a very tight bound. This leads to the following constraint

$$\boldsymbol{\sigma} \leq \boldsymbol{\sigma}_{\text{MVDR}}. \quad (7.45)$$

Interestingly, for $K = 1$ the MVDR dirty image is the same image as we obtained earlier by applying a Jacobi preconditioner to the WLS problem. For this case it is also the tightest upper bound because it is identical to the results obtain from applying the ASSC beamformer.

7.3.4. ESTIMATION OF THE UPPER BOUND FROM NOISY DATA

The upper bounds (7.22) and (7.45) assume that we know the true covariance matrix \mathbf{R} . However in practice we only measure $\hat{\mathbf{R}}$ which is subject to statistical fluctuations. Choosing a confidence level of 6 times the standard deviation of the dirty images ensures that the upper bound will hold with probability 99.9%. This leads to an increase of the upper bound by a factor $1 + \alpha$ where $\alpha > 0$ is chosen such that

$$\boldsymbol{\sigma} \leq (1 + \alpha) \hat{\boldsymbol{\sigma}}_{\text{MF}}. \quad (7.46)$$

Similarly, for the MVDR dirty image the constraint based on $\hat{\mathbf{R}}$ is

$$\boldsymbol{\sigma} \leq (1 + \alpha) \hat{\boldsymbol{\sigma}}_{\text{MVDR}} \quad (7.47)$$

where

$$\hat{\sigma}_{\text{MVDR},i} = \frac{C}{\frac{1}{K} \sum_k \mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k^{-1} \mathbf{a}_{k,i}} \quad (7.48)$$

is an unbiased estimate of the MVDR dirty image, and

$$C = \frac{N}{N - p} \quad (7.49)$$

is a bias correction constant. With some algebra the unbiased estimate can be written in vector form as

$$\hat{\boldsymbol{\sigma}}_{\text{MVDR}} = \mathbf{D}^{-1} \boldsymbol{\Psi}^H \mathbf{W} \hat{\mathbf{f}}, \quad (7.50)$$

where

$$\mathbf{D} = \frac{1}{KC} \text{diag}^2(\mathbf{A}^H \hat{\mathbf{R}}^{-1} \mathbf{A}), \quad (7.51)$$

and

$$\begin{aligned} \mathbf{A} &= [\mathbf{A}_1^T \quad \dots \quad \mathbf{A}_K^T]^T \\ &= [\mathbf{a}_1 \quad \dots \quad \mathbf{a}_I]. \end{aligned} \quad (7.52)$$

The exact choice of α and C are discussed in Sec. 7.9.

7.3.5. CONSTRAINED LEAST SQUARES IMAGING

Now that we have lower and upper bounds on the image, we can use these as constraints in the LS imaging problem to provide a regularization. The resulting constrained LS (CLS) imaging problem is

$$\begin{aligned} \min_{\boldsymbol{\sigma}} \quad & \frac{1}{2K} \|\hat{\mathbf{r}} - \boldsymbol{\Psi}\boldsymbol{\sigma}\|^2 \\ \text{s.t.} \quad & \mathbf{0} \leq \boldsymbol{\sigma} \leq \boldsymbol{\gamma} \end{aligned} \quad (7.53)$$

where $\boldsymbol{\gamma}$ can be chosen either as $\boldsymbol{\gamma} = \boldsymbol{\sigma}_{\text{MF}}$ for the MF dirty image or $\boldsymbol{\gamma} = \boldsymbol{\sigma}_{\text{MVDR}}$ for the MVDR dirty image (or their sample covariance based estimates given by (7.46) and (7.47)).

The improvements to the unconstrained LS problem that were discussed in Sec. 7.2.2 are still applicable. The extension to WLS leads to the cost function

$$f_{\text{WLS}}(\boldsymbol{\sigma}) = \frac{1}{2} \|\mathbf{W}^{1/2} (\hat{\mathbf{r}} - \boldsymbol{\Psi}\boldsymbol{\sigma})\|^2. \quad (7.54)$$

The constrained WLS problem is then given by

$$\begin{aligned} \min_{\boldsymbol{\sigma}} \quad & f_{\text{WLS}}(\boldsymbol{\sigma}) \\ \text{s.t.} \quad & \mathbf{0} \leq \boldsymbol{\sigma} \leq \boldsymbol{\gamma}. \end{aligned} \quad (7.55)$$

We also recommend to include a preconditioner which, as was shown in Sec. 7.2.3, relates the WLS to the MVDR dirty image. However, because of the inequality constraints, (7.55) does not have a closed form solution and it is solved by an iterative algorithm. In order to have the relation between WLS and MVDR dirty image during the iterations we introduce a change of variable of the form $\check{\boldsymbol{\sigma}} = \mathbf{D}\boldsymbol{\sigma}$, where $\check{\boldsymbol{\sigma}}$ is the new variable for the preconditioned problem and the diagonal matrix \mathbf{D} is given in (7.51). The resulting constrained preconditioned WLS (CPWLS) optimization problem is

$$\begin{aligned} \check{\boldsymbol{\sigma}} = \quad & \arg \min_{\check{\boldsymbol{\sigma}}} \frac{1}{2} \|\mathbf{W}^{1/2} (\hat{\mathbf{r}} - \boldsymbol{\Psi}\mathbf{D}^{-1}\check{\boldsymbol{\sigma}})\|^2 \\ \text{s.t.} \quad & \mathbf{0} \leq \check{\boldsymbol{\sigma}} \leq \mathbf{D}\boldsymbol{\gamma} \end{aligned} \quad (7.56)$$

and the final image is found by setting $\boldsymbol{\sigma} = \mathbf{D}^{-1}\check{\boldsymbol{\sigma}}$. (Here we used that \mathbf{D} is a positive diagonal matrix so that the transformation to an upper bound for $\check{\boldsymbol{\sigma}}$ is correct.) Interestingly, the dirty image that follows from the (unconstrained) Weighted Least Squares part of the problem is given by the MVDR image $\hat{\boldsymbol{\sigma}}_{\text{MVDR}}$ in (7.50).

7.4. CONSTRAINED OPTIMIZATION

The constrained imaging formulated in the previous section requires the numerical solution of the optimization problems (7.53) or (7.56). The problem is classified as a positive definite quadratic program with simple bounds, this is a special case of a convex optimization problem with linear inequality constraints, and we can follow standard approaches to find a solution [41, 116].

For an unconstrained optimization problem, the gradient of the cost function calculated at the solution must vanish. If in an iterative process we are not yet at the optimum, the gradient is used to update the current solution. For constrained optimization,

the constraints are usually added to the cost function using (unknown) Lagrange multipliers that need to be estimated along with the solution. At the solution, part of the gradient of the cost function is not zero but related to the nonzero Lagrange multipliers. For inequality constraints, the sign of the Lagrange multipliers plays an important role.

As we will show, these characteristics of the solution (based on the gradient and the Lagrange multipliers) can be used to develop an algorithm called the active set method, which is closely related to the sequential source removing techniques such as CLEAN.

In this Section, we use the active set method to solve the constrained optimization problem.

7.4.1. CHARACTERIZATION OF THE OPTIMUM

Let $\bar{\sigma}$ be the solution to the optimization problem (7.53) or (7.56). An image is called feasible if it satisfies the bounds $\sigma \geq \mathbf{0}$ and $-\sigma \geq -\gamma$. At the optimum, some pixels may satisfy a bound with equality, and these are called the “active” pixels.

We will use the following notation. For any feasible image σ , let

$$\mathcal{L}(\sigma) = \{i \mid \sigma_i = 0\} \quad (7.57)$$

$$\mathcal{U}(\sigma) = \{i \mid \sigma_i = \gamma_i\} \quad (7.58)$$

$$\mathcal{A}(\sigma) = \mathcal{L}(\sigma) \cup \mathcal{U}(\sigma) \quad (7.59)$$

$$\mathcal{F}(\sigma) = \mathcal{I} \setminus \mathcal{A}(\sigma). \quad (7.60)$$

$\mathcal{I} = \{1, \dots, I\}$ is the set of all pixel indices, $\mathcal{L}(\sigma)$ is the set where the lower bound is active, i.e., the pixel value is 0. $\mathcal{U}(\sigma)$ is the set of pixels which attain the upper bound. $\mathcal{A}(\sigma)$ is the set of all pixels where one of the constraints is active, these are the active pixels. Finally, the free set $\mathcal{F}(\sigma)$ is the set of pixels i which have values strictly between 0 and γ_i . Further, for any vector $\mathbf{v} = [v_i]$, let $\mathbf{v}_{\mathcal{F}}$ correspond to the subvector with indices $i \in \mathcal{F}$, and similarly define $\mathbf{v}_{\mathcal{L}}$ and $\mathbf{v}_{\mathcal{U}}$. We will write $\mathbf{v} = \mathbf{v}_{\mathcal{F}} \oplus \mathbf{v}_{\mathcal{L}} \oplus \mathbf{v}_{\mathcal{U}}$.

Let $\bar{\sigma}$ be the optimum, and let $\bar{\mathbf{g}} = \mathbf{g}(\bar{\sigma})$ be the gradient of the cost function at this point. Define the free sets and active sets $\mathcal{F}, \mathcal{L}, \mathcal{U}$ at $\bar{\sigma}$. We can write $\bar{\mathbf{g}} = \bar{\mathbf{g}}_{\mathcal{F}} \oplus \bar{\mathbf{g}}_{\mathcal{L}} \oplus \bar{\mathbf{g}}_{\mathcal{U}}$. Associated with the active pixels of $\bar{\sigma}$ is a vector $\bar{\boldsymbol{\lambda}} = \bar{\boldsymbol{\lambda}}_{\mathcal{L}} \oplus \bar{\boldsymbol{\lambda}}_{\mathcal{U}}$ of Lagrange multipliers. Optimization theory [41] tells us that the optimum $\bar{\sigma}$ is characterized by the following conditions:

$$\mathbf{g}_{\mathcal{F}}(\bar{\sigma}) = \mathbf{0} \quad (7.61)$$

$$\bar{\boldsymbol{\lambda}}_{\mathcal{L}} = \bar{\mathbf{g}}_{\mathcal{L}} \geq \mathbf{0} \quad (7.62)$$

$$\bar{\boldsymbol{\lambda}}_{\mathcal{U}} = -\bar{\mathbf{g}}_{\mathcal{U}} \geq \mathbf{0}. \quad (7.63)$$

Thus, the part of the gradient corresponding to the free set is zero, but the part of the gradient corresponding to the active pixels is not necessarily zero. Since we have simple bounds, this part becomes equal to the Lagrange multipliers $\bar{\boldsymbol{\lambda}}_{\mathcal{L}}$ and $-\bar{\boldsymbol{\lambda}}_{\mathcal{U}}$ (the negative sign is caused by the condition $-\sigma_{\mathcal{U}} \geq -\gamma_{\mathcal{U}}$). The condition $\bar{\boldsymbol{\lambda}} \geq \mathbf{0}$ is crucial: a negative Lagrange multiplier would indicate that there exists a feasible direction of descent \mathbf{p} for which a small step into that direction, $\bar{\sigma} + \mu\mathbf{p}$, has a lower cost and still satisfies the constraints, thus contradicting optimality of $\bar{\sigma}$ [41].

“Active set” algorithms consider that if the true active set at the solution would be known, the optimization problem with inequality constraints reduces to an optimization with equality constraints,

$$\mathbf{z} = \underset{\boldsymbol{\sigma}}{\operatorname{argmin}} f(\boldsymbol{\sigma}) \quad (7.64)$$

$$\text{s.t. } \boldsymbol{\sigma}_{\mathcal{L}} = \mathbf{0}, \boldsymbol{\sigma}_{\mathcal{U}} = \boldsymbol{\gamma}_{\mathcal{U}}.$$

Since we can substitute the values of the active pixels into $\boldsymbol{\sigma}$, the problem becomes a standard unconstrained LS problem with a reduced dimension: only $\tilde{\boldsymbol{\sigma}}_{\mathcal{F}}$ needs to be estimated. Specifically, for CLS the unconstrained subproblem is formulated as

$$f(\boldsymbol{\sigma}) = \frac{1}{2K} \|\mathbf{b}_{\text{LS}} - \boldsymbol{\Psi}_{\mathcal{F}} \boldsymbol{\sigma}_{\mathcal{F}}\|^2 \quad (7.65)$$

where

$$\mathbf{b}_{\text{LS}} = \hat{\mathbf{r}} - \boldsymbol{\Psi}_{\mathcal{U}} \boldsymbol{\sigma}_{\mathcal{U}}. \quad (7.66)$$

Similarly for CPWLS we have

$$f(\check{\boldsymbol{\sigma}}) = \frac{1}{2} \|\mathbf{b}_{\text{PWLS}} - \mathbf{W}^{1/2} (\boldsymbol{\Psi} \mathbf{D}^{-1})_{\mathcal{F}} \check{\boldsymbol{\sigma}}_{\mathcal{F}}\|^2 \quad (7.67)$$

where

$$\mathbf{b}_{\text{PWLS}} = \mathbf{W}^{1/2} (\hat{\mathbf{r}} - (\boldsymbol{\Psi} \mathbf{D}^{-1})_{\mathcal{U}} \check{\boldsymbol{\sigma}}_{\mathcal{U}}) \quad (7.68)$$

In both cases, closed form solutions can be found, and we will discuss a suitable Krylov-based algorithm for this in Sec. 7.5.

Hence the essence of the constrained optimization problem is to find \mathcal{L} , \mathcal{U} and \mathcal{F} . In the literature algorithms for this are called *active set methods*, and we propose a suitable algorithm in Sec. 7.4.3.

7.4.2. GRADIENTS

We first derive expressions for the gradients required for each of the unconstrained subproblems (7.65) and (7.67). Generically, a WLS cost function (as function of a real-valued parameter vector $\boldsymbol{\theta}$) has the form

$$f(\boldsymbol{\theta})_{\text{WLS}} = \beta \|\mathbf{G}^{1/2} \mathbf{c}(\boldsymbol{\theta})\|^2 = \beta \mathbf{c}(\boldsymbol{\theta})^H \mathbf{G} \mathbf{c}(\boldsymbol{\theta}) \quad (7.69)$$

where \mathbf{G} is a Hermitian weighting matrix and β is a scalar. The gradient of this function is

$$\mathbf{g}(\boldsymbol{\theta}) = 2\beta \left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\theta}^T} \right)^H \mathbf{G} \mathbf{c}. \quad (7.70)$$

For LS we have $\boldsymbol{\theta} = \boldsymbol{\sigma}$, $\mathbf{c} = \hat{\mathbf{r}} - \boldsymbol{\Psi} \boldsymbol{\sigma}$, $\beta = \frac{1}{2K}$ and $\mathbf{G} = \mathbf{I}$. This leads to

$$\begin{aligned} \mathbf{g}_{\text{LS}}(\boldsymbol{\sigma}) &= -\frac{1}{K} \boldsymbol{\Psi}^H (\hat{\mathbf{r}} - \boldsymbol{\Psi} \boldsymbol{\sigma}) \\ &= \mathbf{H}_{\text{LS}} \boldsymbol{\sigma} - \hat{\boldsymbol{\sigma}}_{\text{MF}}. \end{aligned} \quad (7.71)$$

For PWLS, $\theta = \check{\sigma}$, $\mathbf{c} = \hat{\mathbf{r}} - \Psi \mathbf{D}^{-1} \check{\sigma}$, $\beta = \frac{1}{2}$ and $\mathbf{G} = \mathbf{W}$. Substituting into (7.70) we obtain

$$\begin{aligned} \mathbf{g}_{\text{PWLS}}(\check{\sigma}) &= -\mathbf{D}^{-1} \Psi^H \mathbf{W} (\hat{\mathbf{r}} - \Psi \mathbf{D}^{-1} \check{\sigma}) \\ &= \mathbf{H}_{\text{PWLS}} \check{\sigma} - \hat{\sigma}_{\text{MVDR}} \end{aligned} \quad (7.72)$$

where

$$\mathbf{H}_{\text{PWLS}} = \mathbf{D}^{-1} \Psi^H \mathbf{W} \Psi \mathbf{D}^{-1}, \quad (7.73)$$

and we used (7.50).

An interesting observation is that the gradients can be interpreted as residual images obtained by subtracting the dirty image from a convolved model image. This will at a later point allow us to relate the active set method to sequential source removing techniques.

7.4.3. ACTIVE SET METHODS

In this section, we describe the steps needed to find the sets \mathcal{L} , \mathcal{U} and \mathcal{F} , and the solution. We follow the template algorithm proposed in [41]. The algorithm is an iterative technique where we gradually improve on an image. Let the image at iteration j be denoted by $\sigma^{(j)}$ where $j = 1, 2, \dots$, and we always ensure this is a feasible solution (satisfies $\mathbf{0} \leq \sigma^{(j)} \leq \gamma$). The corresponding gradient is the vector $\mathbf{g} = \mathbf{g}(\sigma^{(j)})$, and the current estimate of the Lagrange multipliers λ is obtained from \mathbf{g} using (7.62), (7.63). The sets \mathcal{L} , \mathcal{U} and \mathcal{F} are current estimates that are not yet necessarily equal to the true sets.

If this image is not yet the true solution, it means that one of the conditions in (7.61)–(7.63) is violated. If the gradient corresponding to the free set is not yet zero ($\mathbf{g}_{\mathcal{F}} \neq \mathbf{0}$), then this is remedied by recomputing the image from the essentially unconstrained subproblem (7.64). It may also happen that some entries of λ are negative. This implies that we do not yet have the correct sets \mathcal{L} , \mathcal{U} and \mathcal{F} . Suppose $\lambda_i < 0$. The connection of λ_i to the gradient indicates that the cost function can be reduced in that dimension without violating any constraints [41], at the same time making that pixel not active anymore. Thus we remove the i th pixel from the active set, add it to the free set, and recompute the image with the new equality constraints using (7.64). As discussed later, a threshold ϵ is needed in the test for negativity of λ_i and therefore this step is called the “detection problem”.

Table 7.1 summarizes the resulting active set algorithm and describes how the solution \mathbf{z} to the subproblem is used at each iteration. Some efficiency is obtained by not computing the complete gradient \mathbf{g} at every iteration, but only the parts corresponding to \mathcal{L} , \mathcal{U} , when they are needed. For the part corresponding \mathcal{F} , we use a flag that indicates whether $\mathbf{g}_{\mathcal{F}}$ is zero or not.

In line 1, the iterative process is initialized. This can be done in many ways. As long as the initial image lies within the feasible region ($\mathbf{0} \leq \sigma^{(0)} \leq \gamma$), the algorithm will converge to a constrained solution. We can simply initialize by $\sigma^{(0)} = \mathbf{0}$.

Line 3 is a test for convergence, corresponding to the conditions (7.61)–(7.63). The loop is followed while a condition is violated.

If $\mathbf{g}_{\mathcal{F}}$ is not zero, then in line 5 the unconstrained subproblem (7.64) is solved. If this solution \mathbf{z} satisfies the feasibility constraints, then it is kept, the image is updated

Table 7.1: Constrained LS Imaging Using Active Sets

```

1: Initialize: set the initial image  $\sigma^{(0)} = \mathbf{0}$ ,  $j = 0$ , set the free set  $\mathcal{F} = \emptyset$ , and  $\mathcal{L}, \mathcal{U}$  ac-
   accordingly
2: Set the flag Freegradient-isnotzero := True
3: while Freegradient-isnotzero or  $\lambda_{\min} < 0$  do
4:   if Freegradient-isnotzero then
5:     Let  $\mathbf{z}$  be the solution of the unconstrained subproblem (7.64)
6:     if  $\mathbf{z}$  is feasible then
7:       Update the image:  $\sigma_{\mathcal{F}}^{(j+1)} = \mathbf{z}$ 
8:       Set Freegradient-isnotzero := False
9:       Compute the “active” part of the gradient and estimate the Lagrange multi-
       pliers
10:      Let  $\lambda_{\min}$  be the smallest Lagrange multiplier and  $i_{\min}$  the corresponding pixel
       index
11:    else
12:      Compute the direction of descent  $\mathbf{p} = \mathbf{z} - \sigma_{\mathcal{F}}^{(j)}$ 
13:      Compute the maximum feasible nonnegative step-size  $\mu_{\max}$  and let  $i$  be the
       corresponding pixel index that will attain a bound
14:      Update the image:  $\sigma_{\mathcal{F}}^{(j+1)} = \sigma_{\mathcal{F}}^{(j)} + \mu_{\max} \mathbf{p}$ 
15:      Add a constraint: move  $i$  from the free set  $\mathcal{F}$  to  $\mathcal{L}$  or  $\mathcal{U}$ 
16:      Set Freegradient-isnotzero := True
17:    end if
18:    Increase the image index:  $j := j + 1$ 
19:  else
20:    Delete a constraint: move  $i_{\min}$  from  $\mathcal{L}$  or  $\mathcal{U}$  to the free set  $\mathcal{F}$ 
21:    Set Freegradient-isnotzero := True
22:  end if
23: end while

```


accordingly, and the gradient is estimated at the new solution (only $\lambda_{\min} = \min(\boldsymbol{\lambda})$ is needed, along with the corresponding pixel index).

If \mathbf{z} is not feasible, then in line 12-16 we try to move into the direction of \mathbf{z} as far as possible. The direction of descent is $\mathbf{p} = \mathbf{z} - \boldsymbol{\sigma}_{\mathcal{F}}^{(j)}$, and the update will be $\boldsymbol{\sigma}_{\mathcal{F}}^{(j+1)} = \boldsymbol{\sigma}_{\mathcal{F}}^{(j)} + \mu \mathbf{p}$, where μ is a non-negative step size. The i th pixel will hit a bound if either $\sigma_i^{(j)} + \mu p_i = 0$ or $\sigma_i^{(j)} + \mu p_i = \gamma_i$, i.e., if

$$\mu_i = \max \left(-\frac{\sigma_i^{(j)}}{p_i}, \frac{\gamma_i - \sigma_i^{(j)}}{p_i} \right) \quad (7.74)$$

(note that μ_i is non-negative). Then the maximal feasible step size towards a constraint is given by $\mu_{\max} = \min(\mu_i)$, for $i \in \mathcal{F}$. The corresponding pixel index is removed from \mathcal{F} and added to \mathcal{L} or \mathcal{U} .

If in line 3 the gradient satisfied $\mathbf{g}_{\mathcal{F}} = \mathbf{0}$ but a Lagrange multiplier was negative, we delete the corresponding constraint and add this pixel index to the free set (line 20). After this, the loop is entered again with the new constraint sets.

Suppose we initialize the algorithm with $\boldsymbol{\sigma}^{(0)} = \mathbf{0}$, then all pixel indices will be in the set \mathcal{L} , and the free set is empty. During the first iteration $\boldsymbol{\sigma}_{\mathcal{F}}$ remains empty but the gradient is computed (line 9). Equations (7.71) and (7.72) show that it will be equal to the negated dirty image. Thus the minimum of the Lagrange multipliers λ_{\min} will be the current strongest source in the dirty image and it will be added to the free set when the loop is entered again. This shows that the method as described above will lead to a sequential source removal technique similar to CLEAN. In particular, the PWLS cost function (7.72) relates to LS-MVI [24], which applies CLEAN-like steps to the MVDR dirty image.

In line 3, we try to detect if a pixel should be added to the free set ($\lambda_{\min} < 0$). Note that $\boldsymbol{\lambda}$ follows from the gradient, (7.71) or (7.72), which is a random variable. We should avoid the occurrence of a “false alarm”, because it will lead to overfitting the noise. Therefore, the test should be replaced by $\lambda_{\min} < -\epsilon$, where $\epsilon > 0$ is a suitable detection threshold. Because the gradients are estimated using dirty images, they share the same statistics (the variance of the other component in (7.71) and (7.72) is much smaller). To reach a desired false alarm rate, we propose to choose ϵ proportional to the standard deviation of the i th pixel on the corresponding dirty image for the given cost function. (How to estimate the standard deviation of the dirty images and the threshold is discussed in Sec. 7.9). Choosing ϵ to be 6 times the standard deviation ensures a false alarm of $< 0.1\%$ over the complete image.

The use of this statistic improves the detection and hence the estimates greatly, however the correct detection also depends on the quality of the estimates in the previous iterations. If a strong source is off-grid, the source is usually underestimated, this leads to a biased estimation of the gradient and the Lagrange multipliers, which in turn leads to inclusion of pixels that are not real sources. In the next section we describe one possible solution for this case.

7.4.4. STRONG OFF-GRID SOURCES

In this section, we use a tilde to indicate “true” source parameters (as distinguished from the gridded source model). E.g., $\tilde{\boldsymbol{\sigma}}$ indicates the vector with the true source intensities

and $\tilde{\Sigma}$ the corresponding diagonal matrix, $\Psi_{k,q}$ indicates their array response vectors and $\tilde{\Psi}_k$ the corresponding matrix. The versions without tilde will refer to the I gridded sources.

The mismatch between Ψ and the unknown $\tilde{\Psi}$ results in an underestimation of source intensities, which means that the remaining contribution of that source produces bias and possible artifacts in the image. In order to achieve high dynamic ranges we suggest finding a grid correction for the pixels in the free set \mathcal{F} .

Let $\mathbf{a}_{k,i}$ have the same model as $\Psi_{k,q}$ with β_i pointing towards the center of the i th pixel. When a source is within a pixel but not exactly in the center we can model this mismatch as

$$\begin{aligned}\Psi_{k,q} &= \frac{1}{\sqrt{P}} e^{j\frac{2\pi}{\lambda} \Xi^T \mathbf{Q}_k (\beta_i + \delta_i)} \\ &= \mathbf{a}_{k,i} \odot e^{j\frac{2\pi}{\lambda} \Xi^T \mathbf{Q}_k \delta_i}\end{aligned}$$

where $\delta_i = \tilde{\beta}_q - \beta_i$ and $i \in \mathcal{F}$. Because both β_i and $\tilde{\beta}_q$ are 3×1 unit vectors, each has only two degrees of freedom. This means that we can parameterize the unknowns for the grid correcting problem using coefficients $\delta_{1,i}$ and $\delta_{i,2}$. We will assume that when a source is added to the free set, its actual position is very close to the center of the pixel on which it was detected. This means that $\delta_{1,i}$ and $\delta_{i,2}$ are within the pixel's width, denoted by W , and height, denoted by H . In this case we can replace (7.64) by a non-linear constrained optimization,

$$\begin{aligned}\min_{\delta, \sigma} & \frac{1}{2} \|\mathbf{b} - \Psi(\delta)_{\mathcal{F}} \sigma_{\mathcal{F}}\|_2^2 \\ \text{s.t.} & -W/2 < \delta_{1,i} < W/2 \\ & -H/2 < \delta_{i,2} < H/2\end{aligned}\tag{7.75}$$

where $\Psi(\delta)_{\mathcal{F}}$ contains only the columns corresponding to the set \mathcal{F} , δ_j is a vector obtained by stacking $\delta_{i,j}$ for $j = 1, 2$ and

$$\mathbf{b} = \hat{\mathbf{r}} - \Psi_{\mathcal{U}} \sigma_{\mathcal{U}}.\tag{7.76}$$

This problem can also be seen as a direction of arrival (DOA) estimation which is an active research area and out of the scope of this work. A good review of DOA mismatch correction for MVDR beamformers can be found in [117] and [118] proposed a correction method which is specifically applicable to the radio astronomical context.

Besides solving (7.75) instead of (7.64) in line 5 of the active set method we will also need to update the upper bounds and the standard deviations of the dirty images at the new pixel positions that are used in the other steps (e.g., line 3, 6 and 13), the rest of the steps remain the same. Because we have a good initial guess to where each source in the free set is, we propose a Newton based algorithm to do the correction.

7.4.5. BOXED IMAGING

A common practice in image deconvolution techniques like CLEAN is to use a-priori knowledge and narrow the searching area for the sources to a certain region of the image,

called CLEAN-boxes. Because the contribution of the sources (if any) outside these boxes is assumed to be known, we can subtract them from the data such that we can assume that the intensity outside the boxes is zero.

In order to include these boxes in the optimization process of the active set algorithm it is sufficient to make sure that the value of the pixels not belonging to these boxes do not change and remain zero. This is equivalent to replacing Ψ with $\Psi_{\mathcal{B}}$, where \mathcal{B} is the set of indices belonging to the boxes, before we start the optimization process. However as we will explain in the next section, we will avoid storing the matrix Ψ in memory by exploiting its Khatri-Rao structure. We address this implementation issue by replacing (7.60) with

$$\mathcal{F}(\sigma) = (\mathcal{I} \setminus \mathcal{A}(\sigma)) \cap \mathcal{B} \quad (7.77)$$

which makes sure that the values of the elements outside of the boxes do not change. This has the same effect as removing the columns not belonging to \mathcal{B} from Ψ . Of course we have to make sure that these values are initialized to zero. By choosing $\sigma^{(0)} = \mathbf{0}$ this is automatically the case. The only problem with this approach is that the values outside the box remain in the set \mathcal{L} which is used for estimating the Lagrange variables, resulting in expensive calculation that are not needed. This problem is easily solved by calculating the gradient only for the pixels belonging to \mathcal{B} . The a-priori non-zero values of the pixels (that were not in the boxes and were removed from the data) are added to the solution when the optimization process is finished.

7.5. IMPLEMENTATION USING KRYLOV SUBSPACE BASED METHODS

7

From the active set methods described in the previous section, we know that we need to solve (7.65) or (7.67) at each iteration. In this section we describe how to achieve this efficiently and without the need of storing the whole convolution matrix in memory.

During the active set updates we need to solve linear equations of the form $\mathbf{M}\mathbf{x} = \mathbf{b}$. However there are cases where we do not have direct access to the elements of the matrix \mathbf{M} . This can happen for example when \mathbf{M} is too large to fit in memory. There are also cases where \mathbf{M} (or \mathbf{M}^H) are implemented as subroutines that produce the result of the matrix vector multiplication $\mathbf{M}\mathbf{v}$ for some input vector \mathbf{v} . For example for $\mathbf{M} = \Psi$ the operation $\Psi^H \mathbf{v}$ generates a dirty image. An equivalent (and maybe optimized) implementation of such imaging subroutine might be already available to the user. In these scenarios it is necessary or beneficial to be able to solve the linear systems, using only the available matrix vector multiplication or the equivalent operator. A class of iterative solvers that can solve a linear system by only having access to the result of the multiplications with the matrix \mathbf{M} are the Krylov subspace based Methods.

7.5.1. IMPLEMENTATION

During the active set iteration we need to solve (7.65) and (7.67) where the matrix \mathbf{M} in LSQR is replaced by $\Psi_{\mathcal{F}}$ and $bW^{1/2}(\Psi\mathbf{D}^{-1})_{\mathcal{F}}$ respectively. Because Ψ has a Khatri-Rao structure and selecting and scaling a subset of columns does not change this, $\Psi_{\mathcal{F}}$ and $(\Psi\mathbf{D}^{-1})_{\mathcal{F}}$ also have a Khatri-Rao structure. Here we will show how to use this structure to implement (2.40) in parallel and with less memory usage.

Note that the only time the matrix \mathbf{M} enters the algorithm is via the matrix-vector multiplications $\mathbf{M}\mathbf{v}_n$ and $\mathbf{M}^H\mathbf{u}_{n+1}$. As an example we will use $\mathbf{M} = \Psi_{\mathcal{F}}$ for solving (7.65). Let $\mathbf{k}_n = \Psi_{\mathcal{F}}\mathbf{v}_n$. We partition \mathbf{k}_n as Ψ into

$$\mathbf{k}_n = [\mathbf{k}_{1,n}^T \quad \dots \quad \mathbf{k}_{K,n}^T]^T. \quad (7.78)$$

Using the definition of Ψ in (3.12), the operation $\mathbf{k}_n = \Psi_{\mathcal{F}}\mathbf{v}_n$ could also be performed using

$$\mathbf{K}_{k,n} = \sum_{i \in \mathcal{F}} v_{i,n} \mathbf{a}_{k,i} \mathbf{a}_{k,i}^H. \quad (7.79)$$

and subsequently setting

$$\mathbf{k}_{k,n} = \text{vect}(\mathbf{K}_{k,n}). \quad (7.80)$$

This process can be highly parallelized because of the independence between the correlation matrices of each time snapshot. The matrix $\mathbf{K}_{k,n}$ can then be used to find the updates in (2.40).

The operation $\mathbf{M}^H\mathbf{u}$ in (2.40), is implemented in a similar way. Using the beamforming approach (similar to Sec. 7.3), this operation can also be done in parallel for each pixel and each snapshot.

In both cases the calculations can be formulated as correlations and beamforming of parallel data paths which means that efficient hardware implementations are feasible. Also we can consider traditional LS or WLS solutions as a special case when all the pixels belong to the free set which means that those algorithms can also be implemented efficiently in hardware in the same way. During the calculations we work with a single beamformer at the time, and the matrix Ψ need not to be pre-calculated and stored in memory. This makes it possible to apply image formation algorithms for large images when there is a memory shortage.

The computational complexity of the algorithm is dominated by the transformation between the visibility domain and image domain (correlation and beamforming). The dirty image formation and correlation have a complexity of $O(KP^2I)$. This means that the worst case complexity of the active set algorithm is $O(TMKP^2I)$ where T is the number of active set iterations and M is the maximum number of Krylov iterations. A direct implementation of CLEAN for solving the imaging problem presented in Sec. 7.2 in a similar way would have a complexity of $O(TKP^2I)$. Hence the proposed algorithm is order M times more complex, essentially because it recalculates the flux for all the pixels in the free-set while CLEAN only estimates the flux of a newly added pixel. Considering that (for a well posed problem) solving $\mathbf{M}\mathbf{x} = \mathbf{b}$ using LSQR is algebraically equivalent to solving $\mathbf{M}^H\mathbf{M}\mathbf{x} = \mathbf{M}^H\mathbf{b}$ using CG [119], we can use the convergence properties of CG [45] to obtain an indication of the required number of Krylov iterations M . It is found that M is of the order $O(\sqrt{\text{card}(\mathcal{F})})$ where $\text{card}(\mathcal{F})$ is the cardinality of the free set which is equal to the number of pixels in the free set.

In practice, many implementations of CLEAN use the FFT instead of a DFT (matched filter) for calculating the dirty image. Extending the proposed method to use similar techniques is possible and will be presented in future works.




Figure 7.1: Contoured true source in dB scale

7

7.6. SIMULATIONS

The performance of the proposed methods are evaluated using simulations. Because the active set algorithm adds a single pixel to the free set at each step, it is important to investigate the effect of this procedure on extended sources and noise. For this purpose, in our first simulation set-up we will use a high dynamic range simulated image with a strong point source and two weaker extended sources in the first part of the simulations. In a second set-up we will make a full sky image using sources from the 3C catalog.

Following the discussion in Sec. 7.4.2 we define the residual image for CLS and CLEAN as

$$\sigma_{res} = \Psi^H (\hat{\mathbf{r}} - \Psi \sigma - \mathbf{r}_n),$$

and for CPWLS we use

$$\sigma_{res} = \mathbf{D}^{-1} \Psi^H \mathbf{W} (\hat{\mathbf{r}} - \Psi \mathbf{D}^{-1} \check{\sigma} - \mathbf{r}_n)$$


where we assume to know the noise covariance matrix \mathbf{R}_n .

7.6.1. EXTENDED SOURCES


An array of 100 dipoles ($P = 100$) with random distribution is used with the frequency range of 58-90 MHz from which we will simulate three equally spaced channels. Each channel has a bandwidth of 195 kHz and is sampled at Nyquist-rate. These specifications are consistent with the LOFAR telescope in LBA mode [1]. LOFAR uses 1 second

Figure 7.2: Contoured dirty images in dB scale

(a) MF dirty image



(b) MVDR dirty image



snapshots and we will simulate using only two snapshots, i.e., $K = 2$. We use spectrally white sources for the simulated frequency channels which allows us to extend the data model to one containing all frequency data by simply stacking the individual $\hat{\mathbf{r}}$ for each frequency into a single vector. Likewise we stack the individual Ψ into a single matrix. Since the source intensity vector σ is common for all frequencies, the augmented data model has the same structure as before.

The simulated source is a combination of a strong point source with intensity 40 dB and two extended structures with intensities of 0 dB. The extended structures are composed from seven nearby Gaussian shaped sources, one in the middle and 6 on a hexagon around it. This configuration is selected to generate an easily reproducible example. Figure 7.1 shows the simulated image in dB scale. The background noise level that is added is at -10 dB which is also 10 dB below the the extended sources. This is equivalent to a dynamic range of 50 dB and a minimum SNR of 10dB.

Figures 7.2a and 7.2b show the matched filter and MVDR dirty images respectively. The first column of Figure 7.3 shows the final result of the CLEAN, CLS with the MF dirty image as upper bound, CLS with the MVDR dirty image as upper bound, and CPWLS with the MVDR dirty image as upper bound without the residual images. For each image, the extracted point sources have been convolved with a Gaussian beam to smoothen the image. We used a Gaussian beam that has the same main beamwidth as the MF dirty image. The second column of Figure 7.3 shows the corresponding residual images as defined before, and the last column shows a cross section parallel to the β_2 axis going through the sources at the center of the image.

Remarks are:

- As expected the MVDR dirty image has a much better dynamic range (≈ 50 dB) and lower side-lobes compared to the MF dirty image (≈ 15 dB dynamic range).
- Due to a better initial dirty image and upper bound, the CPWLS deconvolution gives a better reconstruction of the image.




Figure 7.3: Extended Source Simulations. Units for first and third column are in dB. Linear scale is used for residual images (second column).

- The cross sections show the accuracy of the estimated intensities. This shows that not only the shape but also the magnitude of the sources are better estimated using CPWLS.
- Using the MVDR upper bound for CLS improves the estimate, illustrating the positive effect of using a proper upper bound.
- The residual image for CPWLS is almost two orders of magnitude lower than the residual images for CLEAN and CLS.
- While the residual image of the CLS algorithm appears very similar to the CLEAN reconstruction, CLS can guarantee that these values are inside the chosen confidence interval of 6 standard deviations of each pixel, while CLEAN does not provide such a guarantee.

7.6.2. FULL SKY WITH 3C SOURCES

In a second simulation set-up, we construct an all-sky image with sources from the 3C catalog. The array configuration is the same as before with the same number of channels and snapshots. A background noise level of 0 dB (with respect to 1 Jansky) is added to the sky.

We first checked which sources from the 3C catalog are visible at the simulated date and time. From these we have chosen 20 sources that represent the magnitude distribution on the sky and produce the highest dynamic range available in this catalog. Table 7.2 shows the simulated sources with corresponding parameters. The coordinates are the (l, m) coordinates at the first snapshot. Because the sources are not necessarily on the grid points, we have combined the active set deconvolution with the grid corrections on the free set as described in Sec. 7.4.4.

Figure 7.4a shows the true and estimated positions for the detected sources. Because the detection mechanism was able to detect the correct number of sources, we have included the estimated fluxes also in Table 7.2 for easier comparison. Figure 7.4b shows the full sky MF dirty image. Figure 7.5a shows the final reconstructed image with the residual added to it (with grid corrections applied), and Figure 7.5c shows the same result for CLEAN.

Remarks:

- The active set algorithm with grid corrections automatically stops after adding the correct number of sources based on the detection mechanism we have incorporated in the active set method;
- Because of the grid correction no additional sources are added to compensate for incorrect intensity estimates on the grids;
- All 20 sources are visible in the final reconstructed image and no visible artifacts are added to the image.
- CLEAN also produces a reasonable image with all the sources visible. However, a few hundred point sources have been detected during the CLEAN iteration, most of which are the result of the strong sources that are not on the grid. Some clear

Table 7.2: Simulated Sources from 3C Catalog


Names	l	m	Flux (Jy)	Est. Flux (Jy)
3C 461	-0.30485	0.19131	11000	10997.61
3C 134	0.59704	-0.02604	66	65.92
3C 219	0.63907	0.6598	44	44.07
3C 83.1	0.28778	-0.13305	28	27.97
3C 75	0.30267	-0.684	23	23.02
3C 47	-0.042882	-0.51909	20	19.97
3C 399.2	-0.97535	0.20927	19	18.97
3C 6.1	-0.070388	0.47098	16	15.99
3C 105	0.57458	-0.60492	15	15.10
3C 158	0.9017	-0.12339	14	14.01
3C 231	0.28956	0.72005	13	13.02
3C 303	-0.1511	0.95402	12.5	12.51
3C 277.1	0.12621	0.93253	12	12.03
3C 320	-0.3597	0.93295	11.5	11.62
3C 280.1	0.15171	0.98709	11	10.95
3C 454.2	-0.29281	0.31322	10.5	10.48
3C 458	-0.61955	-0.56001	10	10.01
3C 223.1	0.67364	0.68376	9.5	9.63
3C 19	-0.23832	-0.30028	9	8.87
3C 437.1	-0.83232	-0.24924	5	4.99

artifacts are introduced (as seen in the residual image) which are also the result of the incorrect subtraction of off-grid sources.


- Fig. 7.5b shows that the residual image using active set and grid corrections contains a “halo” around the position of the strong source—the residual image is not flat. In fact, the detection mechanism in the active set algorithm (with a threshold of 6 times the standard deviation) has correctly not considered this halo as a source. The halo is a statistical artifact due to finite samples and will be reduced in magnitude by longer observations, with a rate proportional to $1/\sqrt{NK^2}$.
- The CLEAN algorithm requires more than 100 sources to model the image. This is mainly because of the the strong off-grid source (Cassiopeia A). This illustrates that while CLEAN is less complex than the proposed method when the number of detected sources are equal, in practice CLEAN might need many more sources to model the same image.

7.7. CONCLUSIONS

Based on a parametric model and constraints on the intensities, we have formulated image deconvolution as a weighted least squares optimization problem with inequality constraints. We showed that the classical (matched filter) dirty image is an upper bound,




(a) Position Estimates

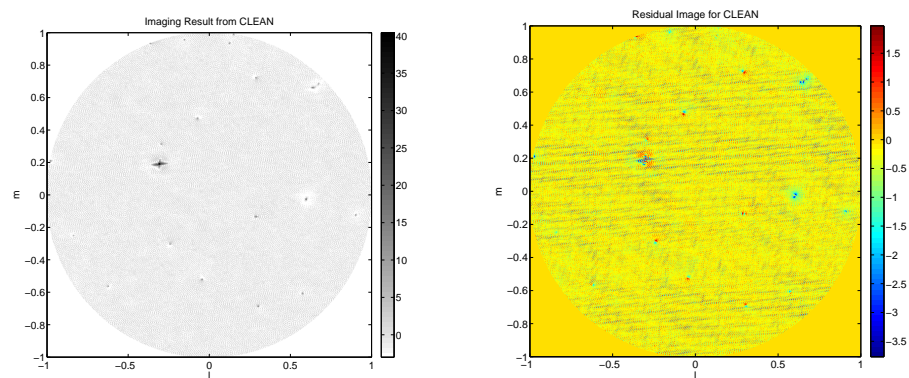


(b) Full Sky MF Dirty Image in dB (with respect to 1 Jy)

Figure 7.4: Point source simulations



(a) Reconstructed image with grid correction plus (b) Residual image using active set and grid correction residual image



(c) Reconstructed image with CLEAN plus residual im- (d) Residual image using CLEAN
age

Figure 7.5: Reconstructed images in dB (with respect to 1 Jy) scale and residual images in linear scale

but a much tighter upper bound is provided by the “MVDR dirty image”. The conditioning of the problem can be improved by a preconditioning step, which is also related to the MVDR dirty image.

Secondly, the constrained least squares problem is solved using an active set based method. The relation between the resulting method and sequential source removing techniques such as CLEAN is explained. The theoretical background of the active set methods can be used to gain better insight into how the sequential techniques work. In particular, the active set algorithm uses a detection threshold with a known false alarm, that can be set such that no false sources appear in the image, and we have shown that by introducing a grid correcting step into the active set iterations we can improve both the detection of the sources and the estimation of their intensities.

Thirdly, the Khatri-Rao structure of the data model is used in combination with Krylov based techniques to solve the linear systems involved in the deconvolution process with less storage and complexity. The complexity of the algorithm is higher than that of classical sequential source removing techniques (by a factor proportional to the square root of the detected number of sources), because the detected source intensities are re-estimated by the Krylov subspace technique after each step of the active set iteration. However the proposed algorithm has a better detection mechanism compared to classical CLEAN which leads to a lower number of sources to model the image. As a result the overall complexity is expected to be comparable. We also expect that the performance of the algorithm can be readily improved because the updates by the active set iterations are one-dimensional (one source is added or removed), and this can be exploited to update the Krylov subspaces accordingly, rather than computing them each time from scratch. This is left as future work.

The simulations show that the proposed CPWLS algorithm provides improved spatial structure and improved intensity estimates compared to CLEAN based deconvolution of the classical dirty image. A particularly attractive aspect is the demonstrated capability of the algorithm to perform automated source detection, which will be of interest for upcoming large surveys.

7

7.8. APPENDIX: RELATION BETWEEN WLS, NATURAL AND ROBUST WEIGHTING

Natural weighting is a technique to improve the detection of weak sources by promoting the visibility values that have a better signal-to-noise-ratio [20]. This is done by dividing each visibility sample by the variance of noise on that sample (while assuming that the noise on each sample is independent). Considering that the visibility samples are the elements of the covariance matrix $\hat{\mathbf{R}}_k$ we can model the sample visibilities as

$$\hat{\mathbf{r}}_k = \mathbf{r}_k + \boldsymbol{\epsilon} \quad (7.81)$$

where $\boldsymbol{\epsilon}$ is the complex noise on the samples. As discussed in Sec. 7.9, $\hat{\mathbf{R}}$ has a Wishart distribution and for a large number of samples N we have $\hat{\mathbf{r}}_k \sim \mathcal{N}(\mathbf{r}_k, (\mathbf{R}_k^T \otimes \mathbf{R}_k)/N)$. This means that $\boldsymbol{\epsilon} = \hat{\mathbf{r}}_k - \mathbf{r}_k$ has a complex Gaussian distribution $\mathcal{N}(\mathbf{0}, (\mathbf{R}_k^T \otimes \mathbf{R}_k)/N)$. Because astronomical sources are usually much weaker than the system noise, it is common to use the approximation $\mathbf{R}_k \approx \mathbf{R}_{n,k}$. With this approximation and using the independence

of system noise on each receiving element (antenna or station), we can assume that $\mathbf{R}_{n,k}$ is diagonal and that $(\mathbf{R}_k^T \otimes \mathbf{R}_k)/N \approx (\mathbf{R}_{n,k}^T \otimes \mathbf{R}_{n,k})/N$ is also a diagonal approximation of the noise covariance matrix on the visibility samples. With this framework we can write the natural weighting as

$$\hat{\mathbf{r}}_{\text{natural}} = N(\mathbf{R}_{n,k}^{-T} \otimes \mathbf{R}_{n,k}^{-1})\hat{\mathbf{r}}_k \quad (7.82)$$

This shows that natural weighting is a very reasonable approximation of the weighting used when solving (7.11) for WLS (except for a factor N that drops out from both sides).

Next, we relate WLS to Robust Weighting [20] by assuming slightly different simplifications. Let us assume that $\mathbf{R}_{n,k} = \sigma_n^2 \mathbf{I}$ and let us consider a single source with intensity σ then we have for

$$\begin{aligned} \mathbf{R}_k^{-1} &= (\mathbf{R}_{n,k} + \sigma \mathbf{\Psi}_k \mathbf{\Psi}_k^H)^{-1} \\ &= \mathbf{R}_{n,k}^{-1} - \frac{\sigma \mathbf{R}_{n,k}^{-1} \mathbf{\Psi}_k \mathbf{\Psi}_k^H \mathbf{R}_{n,k}^{-1}}{1 + \sigma \mathbf{\Psi}_k^H \mathbf{R}_{n,k}^{-1} \mathbf{\Psi}_k} \\ &= \frac{1}{\sigma_n^2} \left(\mathbf{I} - \frac{\mathbf{\Psi}_k \mathbf{\Psi}_k^H}{1 + \frac{\sigma_n^2}{\sigma}} \right). \end{aligned} \quad (7.83)$$

Compared to natural weighting, now not only the noise power but also the available signal power is taken into account for the weighting. The term $1/(1 + \sigma_n^2/\sigma)$ is the same as the parametric Wiener filter in the Fourier domain as given by [20] which relates Robust Weighting to standard signal processing concepts. However Robust Weighting also takes the visibility sampling of gridded uv-plane into account when calculating the weights, which is not explained in the derivation above. Hence the exact relation between Robust Weighting and WLS is still missing. This relation is interesting and should be addressed in future works.

7.9. APPENDIX: VARIANCE OF THE DIRTY IMAGE

To find the confidence intervals for the dirty images we need to find estimates for the variance of both matched filter and MVDR dirty images. In our problem the sample covariance matrix is obtained by squaring samples from a Gaussian process. This means that $N\hat{\mathbf{R}} \sim \mathcal{W}_p(\mathbf{R}, N)$ where $\mathcal{W}_p(\mathbf{R}, N)$ is the Wishart distribution function of order p with expected value equal to \mathbf{R} and N degrees of freedom. For any deterministic vector $\boldsymbol{\zeta}$,

$$N\boldsymbol{\zeta}^H \hat{\mathbf{R}} \boldsymbol{\zeta} \sim \boldsymbol{\zeta}^H \mathbf{R} \boldsymbol{\zeta} \chi^2(N). \quad (7.84)$$

where $\chi^2(N)$ is the standard χ^2 distribution with N degrees of freedom. In radio astronomical applications N is usually very large and we can approximate this χ^2 distribution with a Gaussian such that $\boldsymbol{\zeta}^H \hat{\mathbf{R}} \boldsymbol{\zeta} \sim \mathcal{N}(\boldsymbol{\zeta}^H \mathbf{R} \boldsymbol{\zeta}, (\boldsymbol{\zeta}^H \mathbf{R} \boldsymbol{\zeta})^2/N)$. The variance of the matched filter dirty image is given by

$$\text{Var}(\sigma_{\text{MF},i}) = \frac{1}{NK^2} \sum_k (\mathbf{a}_{k,i}^H \mathbf{R} \mathbf{a}_{k,i})^2$$

Using this result we can find the $x\%$ confidence interval which results in an increase of the upper bound such that

$$\sigma \leq \hat{\sigma}_{\text{MF}} + \alpha \sqrt{\text{Var}(\hat{\sigma}_{\text{MF}})} \quad (7.85)$$

where α is chosen depending on x . Requiring at most a single false detection on the entire image translate into $\alpha \approx 6$.

When we estimate the MVDR dirty image from sample covariance matrices we need to be more careful, mainly because the result is biased and we need to correct for that bias. For each pixel of the MVDR dirty image obtained from sample covariance matrices we have

$$\hat{\sigma}_{\text{MVDR},i} = K g(Z) = \frac{K}{\sum_k \mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k^{-1} \mathbf{a}_{k,i}} \quad (7.86)$$

where $g(Z) = 1/Z$ and $Z = \sum_k \mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k^{-1} \mathbf{a}_{k,i}$. Using a perturbation model $Z = Z_0 + \Delta Z$ and a Taylor approximation we find

$$\begin{aligned} g(Z) &\approx \frac{1}{Z_0} - \frac{1}{Z_0^2} \Delta Z \\ &\approx \frac{1}{Z_0^2} (Z_0 - \Delta Z). \end{aligned} \quad (7.87)$$

Let $Z_0 = \mathcal{E}\{Z\}$ then $\mathcal{E}\{\Delta Z\} = 0$ and $\mathcal{E}\{g(Z)\} \approx 1/Z_0$. We would like this estimate to be unbiased which means that we want

$$\mathcal{E}\{g(Z)\} \approx \frac{1}{\sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i}} \quad (7.88)$$

however we have,

$$\begin{aligned} Z_0 &= \sum_k \mathbf{a}_{k,i}^H \mathcal{E}\{\hat{\mathbf{R}}_k^{-1}\} \mathbf{a}_{k,i} \\ &= \sum_k \mathbf{a}_{k,i}^H \frac{N \mathbf{R}_k^{-1}}{N-p} \mathbf{a}_{k,i} \\ &= \frac{N}{N-p} \sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i} \end{aligned} \quad (7.89)$$

where we have used $\mathcal{E}\{\hat{\mathbf{R}}^{-1}\} = \frac{N}{N-p} \mathbf{R}^{-1}$ [120]. So in order to remove this bias we need to scale it by a correction factor

$$C = \frac{N}{N-p} \quad (7.90)$$

and

$$\hat{\sigma}_{\text{MVDR},i} = CK g(Z). \quad (7.91)$$

Now we need to find an estimate for the variance of the MVDR dirty image. Using (7.87) we see that the first order approximation of $\text{Var}(g(Z)) \approx \text{Var}(Z)/Z_0^4$. We find $\text{Var}(Z)$ using the independence of each snapshot so we can write

$$\text{Var}(Z) = \sum_k \text{Var}(\mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k^{-1} \mathbf{a}_{k,i}). \quad (7.92)$$

In order to find $\text{Var}(\mathbf{a}_{k,i}^H \hat{\mathbf{R}}_k^{-1} \mathbf{a}_{k,i})$ we need to use some properties of the complex inverse Wishart distribution. A matrix has complex inverse Wishart distribution if its inverse has a complex Wishart distribution [120]. Let us define an invertible matrix \mathbf{B} as

$$\mathbf{B} = [\mathbf{a}_{k,i} \quad \mathbf{B}_1] \quad (7.93)$$

then $\mathbf{X} = (\mathbf{B} \hat{\mathbf{R}}^{-1} \mathbf{B}^H)/N$ has an inverse Wishart distribution because $\mathbf{X}^{-1} = N(\mathbf{B}^{-H} \hat{\mathbf{R}} \mathbf{B}^{-1})$ has a Wishart distribution. In this case $\mathbf{X}_{11} = (\mathbf{a}_{k,i}^H \hat{\mathbf{R}}^{-1} \mathbf{a}_{k,i})/N$ also has an inverse Wishart distribution with less degrees of freedom. The covariance of an inverse Wishart matrix is derived in [120], however because we are dealing only with one element, this results simplifies to

$$\text{Var}(N\mathbf{X}_{11}) = \frac{N^2}{(N-p)^2(N-p-1)} (\mathbf{a}_{k,i}^H \mathbf{R}^{-1} \mathbf{a}_{k,i})^2. \quad (7.94)$$

The variance of the unbiased MVDR dirty image is thus given by

$$\begin{aligned} \text{Var}(\hat{\sigma}_{\text{MVDR},i}) &= \text{Var}(CKg(Z)) \\ &\approx \frac{K^2}{(N-p-1)} \frac{\sum_k (\mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i})^2}{(\sum_k \mathbf{a}_{k,i}^H \mathbf{R}_k^{-1} \mathbf{a}_{k,i})^4}. \end{aligned}$$

Now that we have the variance we can use the same method that we used for MF dirty image to find α and

$$\sigma \leq \hat{\sigma}_{\text{MVDR}} + \alpha \sqrt{\text{Var}(\hat{\sigma}_{\text{MVDR}})} \quad (7.95)$$

8

CONCLUSIONS AND SUGGESTIONS FOR FURTHER RESEARCH

8.1. SUMMARY OF MAIN RESULTS

In this thesis we set out to address the applicability of the signal processing formalism to a set of radio astronomical problems. We considered three related main subproblems: RFI mitigation, calibration and imaging. For each of these problems the generic signal processing approach, i.e. modeling, performance analysis and algorithm design, has been demonstrated to be a viable methodology.

Using array processing concepts we have been able to reformulate the radio astronomical measurement equation as a covariance data model. This has enabled us to use the familiar signal processing tools such as subspace estimation and covariance matching techniques as a basis for both analysis and algorithm design. We have also been confronted with the necessity of developing new tools, or tailoring the existing tools to take full advantage of the astronomical data model. Examples of new approaches used in this work are factor analysis (FA) which is used as a basis for developing new subspace estimation techniques and the combination of Kronecker structures and Krylov subspace based methods to reduce the storage and complexity of non-linear optimization problems. These new methods can be added to the the set of tools used in signal processing. We will now give a more detailed summary of the these approaches.

In Chapter 4 we discussed the classical FA and introduced a new set of tools such as Extended FA (EFA) and Joint Extended FA (JEFA). A new estimation algorithm based on Krylov subspace solvers has been proposed for finding the unknown parameters for these new models. Considering that classical FA and EFA can be formulated as an especial cases of JEFA, the newly proposed algorithm is also capable of estimating FA model parameters in an efficient way.

In Chapter 5 we used data from WSRT and LOFAR to demonstrate the potential of spatial filtering and Extended FA to remove RFI and restore contaminated frequency bands which are otherwise unusable for observations. By finding the Cramér–Rao Bound

(CRB) for JEFA, it has become possible to study the statistics for these spatial filtering techniques and to demonstrate that the theoretical bound for the variance of a filtered band is very close to a RFI free band. Using simulation and real data the performance of the new algorithm has been demonstrated which achieves the theoretical limit based on the CRB. It has also been shown using simulations that JEFA can be applied in cases where the Eigen–Value Decomposition (EVD) is also applicable without any significant loss in quality of the recovered subspace, which makes this technique a suitable replacement of EVD for a majority of practical scenarios. Based on the results presented in this thesis, it has been shown that the (local) identifiability of the FA model is completely defined by the signal subspace and structure of non–zero elements of the noise covariance matrix and does not depend on the actual source and noise power.

In Chapter 6 it has been shown that the Khatri–Rao structure of the Jacobians resulting from the vectorization of covariance matching problems can be efficiently combined with Krylov subspace based solvers to solve the direction independent gain calibration problem. Using the same approach it has been shown that the Jacobian with respect to direction of arrival (DOA) also possesses a Khatri–Rao structure, allowing for DOA corrections with the same strategy. In comparison to other techniques such as alternating least squares, using the proposed technique can reduce the complexity of the algorithm per iteration from $O(P^3)$ to $O(P^2)$ (where P is the number of receivers), while also reducing the number of iterations.

In Chapter 7 the relation between the optimally weighted least squares (WLS) problems and the MVDR beamformer was established through a Jacobi preconditioner which demonstrates that data dependent high resolution imaging can be achieved while at the same time improving the statistical performance of the reconstructed image. Using the positive (semi) definiteness of the covariance matrices it was shown that an upper bound on the image values can be found using beamforming techniques. Based on the multi-snapshot data model for the measured covariance data a closed-form solution for the tightest upper bound for an image was derived and its relation with the ASSC beamformer was established. Estimating a high resolution upper bound from noisy data has been shown to be possible by using a multi-snapshot MVDR beamformer.

Revisiting the LS and WLS with the additional non–negativity and upper bound constraints resulted in drastic algorithmic changes to the way these problems are solved. Using these conceptually simple constraints has led to the application of an active set optimization technique which is much closer to sequential source removing methods like CLEAN from an algorithmic point of view than to the traditional (W)LS problems.

The use of factor analysis (FA), application of Krylov solvers in covariance matching problems and constrained LS problems form the core of the approaches in presented work. Starting from these general concepts an attempt has been made to strike a balance between the development of generic signal processing tools and detailed worked out algorithms tailored to radio astronomical problems. Hence in both aspects the results presented in this work could be refined further. The limitations and needed refinements of these methods are discussed next.

8.2. DISCUSSION

8.2.1. SUBSPACE ESTIMATION

In contrast to EVD which is a linear algebraic decomposition, FA is a multivariate model matching technique. Even though we have demonstrated that in models for radio astronomy (and array processing in general) the expected value of the covariance matrices have such a model, when the SNR and number of samples are low the high variance of the sample covariance matrices causes a large deviation from the model which in turn causes difficulties for estimating the FA parameters, while the estimation process for EVD does not change. This does not mean that the subspace estimated by EVD in these situations is accurate, but from an algorithmic point of view this could be seen as a reliability concern when we replace EVD with FA.

Considering that this is only a problem when the subspace of a (weak) signal is changing too rapidly such that long term integration is not possible, we can still use JEFA as an effective way of estimating the noise covariance matrix by buffering several short snapshots and use the results for whitening subsequent snapshots. EVD can then be used to find the subspace for whitened snapshots. In this scenario JEFA can be seen as a noise calibration step instead of a subspace estimation one. However if the noise covariance matrix changes too rapidly as well then other methods than (JE)FA should be considered.

Another complication in the application of (JE)FA is the detection problem which requires several estimation and hypothesis testing steps. The statistics of the sequential hypothesis testing mechanics that has been proposed in Sec. 4.7 is currently unknown (and in general difficult to find). As a result it is not possible to guarantee a certain detection performance (such as a false alarm probability).

8.2.2. SPATIAL FILTERING USING A REFERENCE ARRAY

The performance of spatial filtering with a reference array and its robustness against model mismatches makes it a promising approach for RFI mitigation. As such use of this technique is highly recommended for the new generation of radio telescopes. Considering that a relatively high time resolution and access to station level correlations are needed, using this technique requires new hardware and data-path modifications for radio telescopes which are already in operation. On the other hand the fact that these modifications can be done at station level means that such modifications are only required at those stations which are highly contaminated with RFI.

It is also important to note that the results presented in this work not only illustrate the performance of the particular technique used here, but also highlights the fact that signal processing tools can be used to remove RFI to acceptable levels for radio astronomical applications. This should encourage the adaptation of RFI mitigation techniques to replace/complement traditional approaches based on flagging.

8.2.3. GAIN CALIBRATION

Using covariance matching techniques combined with Krylov subspace solvers, we have demonstrated how a scalable gain calibration algorithm can be constructed. This technique in its current formulation is suitable for station level calibration and can be considered as an alternative to currently used algorithms.

Direction dependent and polarimetric calibration are essential for full instrument calibration. The method presented here needs further refinements and extensions in order to be suitable for this purpose.

8.2.4. IMAGING

As we have seen, confining the pixel intensity within a region in the positive orthant has profound consequences for the way in which we solve LS image reconstruction problems. Similarly choosing the initial point for starting the algorithm and the way the free set is constructed have been shown to have significant meaning by relating this algorithm to sequential source removing techniques. However, now that such a relation is established and the mathematical background for some of the choices in algorithms such as CLEAN has been found, we could again add conceptually small modifications which will have a larger impact on the way a solution can be found. For example, we can replace the initial image from the lower bound (an empty image) by any other feasible starting point (e.g. the dirty image). Even though the algorithm is unchanged, the interpretation in this case is different than sequential source removing techniques and in particular the way these algorithms handle the extended structures in the image. As a consequence, the active set algorithm should not simply be categorized as a sequential source removing technique. A more comprehensive (physical) analysis is needed in order to fine-tune the generic setup of this algorithm to some of the more specific imaging applications in particular cases with extended sources.

8.3. FUTURE RESEARCH

8.3.1. FACTOR ANALYSIS

In this thesis we have seen that FA and its derivatives form a set of powerful tools and, even though FA is popular in many fields, its application in (array) signal processing has been limited. We have also seen that these techniques can be used in practical signal processing scenarios with reliable results. However as was illustrated, the performance of the algorithms relies heavily on the number of samples and/or SNR available. In low SNR and low number of sample scenarios, finding an accurate initial guess might require a large number of iterations and the probability of getting stuck on a local minimum becomes larger. Other considerations like ensuring that the noise covariance matrix is positive definite also become more dominant in this regime making the estimation process more involved and complex. In this regard some problems remain open:

- If the global minimizer for WLS and ML can be found in low SNR/sample cases, is it close enough to the true solution to be of any practical value?

This problem is partially addressed by studying the CRB bound in Sec. 5.6 specifically figures 5.2 and 5.3 where it is clear that under a certain signal power the CRB is higher than the variance of the unprocessed signal. In these scenarios even if we could find a solution, the variance of the estimates would be of very poor quality. However there is a region where according to the CRB using spatial filtering would be beneficial but the (E)FA algorithm cannot converge to the right solution. This is the region of interest for future analysis.

- Can the result of the local identifiability for the signal subspace and structure of the noise covariance matrix be extended to a global condition? Can an example be generated where the problem is locally identifiable but not unique up to a rotation? Can such a counterexample be found for any P and Q (except for real FA when $P \leq 5$ and $Q \leq 2$ for which it is established by Anderson [55] that the local identifiability leads to global identifiability) and what is the probability of encountering such a case?
- Can a general result be stated about the (local) identifiability of subspaces spanned by the columns of the array response matrix?

8.3.2. CALIBRATION

The work presented in this thesis on calibration is limited to improvements regarding the computational complexity of direction-independent gain calibration. The extension of this strategy to polarimetric and direction-dependent gain calibration is an important step needed towards next generations of radio telescopes such as SKA.

The model structures that enable efficient implementation of the Gauss-Newton-Krylov technique are based on point source models. It is important to extend these models in such a way that a computationally efficient implementation can be found for extended structures. This is a crucial step as the number of sources that can be categorized as extended will increase rapidly with the higher resolutions of new telescopes.

8.3.3. IMAGING

EXTENDED STRUCTURES IN IMAGES

The data model in Chapter 7 assumes that each pixel in the image is a point source. This data model in combination with CLS and the active set algorithm has led to a sequential source removing technique similar to CLEAN. This similarity is also true for the way that these algorithms handle extended structures in the sky which is done by modeling them as many point sources. This results in producing a less accurate sky model and increasing the overall computational complexity. To counter this issue we need to update the data model to include these extended structures which leads to the following questions:

- Is it possible to extend the CLS imaging to include more general sky models? If so can we improve the lower and upper bounds based on these new models?
- How does the accuracy of such an approach compare to existing algorithms such as sparse [8] or Bayesian [109] techniques?
- How does the complexity change with this approach? Does the gain in the accuracy justify such a modification?

UPDATED KRYLOV SUBSPACES

One of the most important features of Krylov subspace based methods is the storage reduction needed to find a solution to a linear system of equations. For example for a $M \times N$ image, storing the deconvolution matrix in memory (in double precision) alone requires $4 \times MN(MN - 1)$ bytes which for a relatively small image of size 500×500 is over

200 GB of storage. Using the Krylov based method enables us to perform the desired deconvolution in a system with less than 4 GB of memory. However during the active set iteration, we add and remove a single row and column to the deconvolution matrix and update the estimates. Currently there is no computationally efficient way to combine these small perturbations with Krylov subspace methods. Other solvers such as the ones based on QR decomposition are capable of performing this update but require similar storage requirements as the original deconvolution matrix. Finding methods that can benefit from both the storage capacity of the Krylov subspace methods and small computational cost for row/column up- and downdates is important for deconvolving very high resolution images expected for next generation radio-telescopes.

PRECONDITIONING

In this thesis the relation between the preconditioning of the system of linear equations that results from solving weighted least squares problems and high resolution beam-forming techniques such as MVDR has been demonstrated. However there are many preconditioning techniques that remain unexplored. Knowing that this relation exists, new approaches to creating both dirty and clean images based on various preconditioners should be investigated.

BIBLIOGRAPHY

REFERENCES

- [1] M. P. van Haarlem, M. W. Wise, A. W. Gunst, *et al.*, “LOFAR: The LOw-Frequency ARray,” *A&A*, vol. 556, p. A2, 2013.
- [2] S. J. Wijnholds, *Fish-Eye Observing with Phased Array Radio Telescope*. PhD thesis, Delft University of Technology, 2010.
- [3] A. Thompson, J. Moran, and G. Swenson, eds., *Interferometry and Synthesis in Radio astronomy*. John Wiley and Sons, 1986.
- [4] S. Wijnholds and A.-J. van der Veen, “Multisource self-calibration for sensor arrays,” *Signal Processing, IEEE Transactions on*, vol. 57, pp. 3512–3522, Sept 2009.
- [5] A. Leshem and A.-J. van der Veen, “Radio-astronomical imaging in the presence of strong radio- interference,” *IEEE Trans. Informat. Theory*, pp. 1730–1747, August 2000.
- [6] A.-J. Boonstra and A.-J. van der Veen, “Gain calibration methods for radio telescope arrays,” *IEEE Tr. Signal Processing*, vol. 51, pp. 25–38, 2003.
- [7] S. van der Tol, B. Jeffs, and A.-J. van der Veen, “Self-calibration for the lofar radio astronomical array,” *Signal Processing, IEEE Transactions on*, vol. 55, pp. 4497–4510, Sept 2007.
- [8] Y. Wiaux, L. Jacques, G. Puy, A. Scaife, and P. Vandergheynst, “Compressed sensing imaging techniques for radio interferometry,” *Monthly Notices of the Royal Astronomical Society*, vol. 395, no. 3, pp. 1733–1742, 2009.
- [9] S. Wijnholds and A.-J. van der Veen, “Fundamental imaging limits of radio telescope arrays,” *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, no. 5, pp. 613–623, 2008.
- [10] A. Leshem, A.-J. van der Veen, and A. J. Boonstra, “Multichannel interference mitigation techniques in radio-astronomy,” *The Astrophysical Journal Supplements*, pp. 355–373, November 2000.
- [11] P. Fridman and W. Baan, “RFI mitigation methods in radio astronomy,” *Astronomy and Astrophysics*, vol. 378, pp. 327–344, 2001.
- [12] A.-J. Boonstra, *Radio frequency interference mitigation in radio asyronomy*. PhD thesis, Delft University of Technology, The netherlands, 2005.

- [13] J. Raza, A.-J. Boonstra, and A.-J. van der Veen, "Spatial filtering of RF interference in radio astronomy," *IEEE Signal Processing Letters*, vol. 9, Mar. 2002.
- [14] A.-J. Boonstra, A.-J. van der Veen, and J. Raza, "Spatial filtering of continuous interference in radio astronomy," in *Proc. IEEE ICASSP*, vol. 3, (Orlando (FL)), pp. 2933–2936, IEEE, May 2002.
- [15] S. Wijnholds, S. van der Tol, R. Nijboer, and A.-J. van der Veen, "Calibration challenges for future radio telescopes," *Signal Processing Magazine, IEEE*, vol. 27, pp. 30–42, Jan 2010.
- [16] J. A. Högbom, "Aperture synthesis with nonregular distribution of interferometer baselines," *Astron. Astrophys. Suppl.*, vol. 15, pp. 417–426, 1974.
- [17] T. Cornwell, "Multiscale CLEAN deconvolution of radio synthesis images," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, pp. 793–801, Oct. 2008.
- [18] U. Rau, S. Bhatnagar, M. Voronkov, and T. Cornwell, "Advances in calibration and imaging techniques in radio interferometry," *Proceedings of the IEEE*, vol. 97, pp. 1472–1481, Aug 2009.
- [19] S. Bhatnager and T. Cornwell, "Adaptive scale sensitive deconvolution of interferometric images I. Adaptive Scale Pixel (ASP) decomposition," *Astronomy and Astrophysics*, vol. 426, pp. 747–754, 2004.
- [20] D. S. Briggs, *High fidelity deconvolution of moderately resolved sources*. PhD thesis, The New Mexico Institute of Mining and Technology, Socorro, New Mexico, 1995.
- [21] R. Levanda and A. Leshem, "Radio astronomical image formation using sparse reconstruction techniques," *Electrical and Electronics Engineers in Israel, 2008. IEEEI 2008. IEEE 25th Convention of*, pp. 716–720, Dec. 2008.
- [22] R. Carrillo, J. McEwen, and Y. Wiaux, "Sparsity averaging reweighted analysis (sara): a novel algorithm for radio-interferometric imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 426, no. 2, pp. 1223–1234, 2012.
- [23] R. E. Carrillo, J. D. McEwen, and Y. Wiaux, "Purify: a new approach to radio-interferometric imaging," *Monthly Notices of the Royal Astronomical Society*, vol. 439, no. 4, pp. 3591–3604, 2014.
- [24] C. Ben-David and A. Leshem, "Parametric high resolution techniques for radio astronomical imaging," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, pp. 670–684, Oct. 2008.
- [25] A. Leshem and A.-J. van der Veen, "Radio-astronomical imaging in the presence of strong radio interference," *IEEE Trans. on Information Theory, Special issue on information theoretic imaging*, pp. 1730–1747, August 2000.
- [26] A. Bruckstein, M. Elad, and M. Zibulevsky, "On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations," *Information Theory, IEEE Transactions on*, vol. 54, no. 11, pp. 4813–4820, 2008.

- [27] N. Meinshausen, "Sign-constrained least squares estimation for high-dimensional regression," *Electron. J. Statist.*, vol. 7, pp. 1607–1631, 2013.
- [28] A. Mouri Sardarabadi, A.-J. van der Veen, and A.-J. Boonstra, "Spatial Filtering of RF Interference in Radio Astronomy Using a Reference Antenna Array," *Signal Processing, IEEE Transactions on*, vol. 64, pp. 432–447, Jan 2016.
- [29] A. Mouri Sardarabadi, A. Leshem, and A.-J. van der Veen, "Radio Astronomical Image Formation using Constrained Least Squares and Krylov Subspaces," *Astronomy and Astrophysics*, 2015.
- [30] A. Mouri Sardarabadi and A.-J. van der Veen, "Complex Factor Analysis," in *Calibration and Imaging (CALIM) Workshop*, (Manchester, UK), July 2011.
- [31] A. Mouri Sardarabadi and A.-J. van der Veen, "Subspace estimation using factor analysis," in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th*, pp. 477–480, June 2012.
- [32] A. Mouri Sardarabadi and A.-J. van der Veen, "Constrained imaging for radio astronomy," in *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2013 IEEE 5th International Workshop on*, pp. 344–347, Dec 2013.
- [33] A. Mouri Sardarabadi and A.-J. van der Veen, "Application of Extended Factor Analysis for RFI mitigation," in *International Biomedical and Astronomical Signal Processing (BASP) Frontiers*, (Villars-sur-Ollon, Switzerland), Jan. 2013.
- [34] A. Mouri Sardarabadi and A.-J. van der Veen, "Application of Krylov based methods in calibration for radio astronomy," in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2014 IEEE 8th*, pp. 153–156, June 2014.
- [35] A. Mouri Sardarabadi and A.-J. van der Veen, "Joint Extended Factor Analysis," in *WSA 2015; 19th International ITG Workshop on Smart Antennas; Proceedings of*, pp. 1–8, March 2015.
- [36] A. Mouri Sardarabadi, A. Leshem, and A.-J. van der Veen, "Computationally efficient radio astronomical image formation using constrained least squares and the MVDR beamformer," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 5664–5668, April 2015.
- [37] B. Ottersten, P. Stoica, and R. Roy, "Covariance matching estimation techniques for array signal processing applications," *Digital Signal Processing*, vol. 8, no. 3, pp. 185–210, 1998.
- [38] S. M. Kay, *Fundamentals of Statistical Signal Processing, Estimation theory*, vol. Volume I. Prentice Hall, 1993.
- [39] P. J. Schreier, *Statistical Signal Processing of Complex-Valued Data*. Cambridge University Press, 2010.

- [40] W. Bangs, *Array processing with generalized beamformers*. PhD thesis, Yale University, 1971.
- [41] P. E. Gill, W. Murray, and M. H. Wright, *Practical optimization*. London: Academic Press Inc. [Harcourt Brace Jovanovich Publishers], 1981.
- [42] I. C. F. Ipsen and C. D. Meyer, "The idea behind krylov methods," *The American Mathematical Monthly*, vol. 105, no. 10, pp. 889–899, 1998.
- [43] C. C. Paige and M. A. Saunders, "LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares," *ACM Trans. Math. Softw.*, vol. 8, pp. 43–71, Mar. 1982.
- [44] G. Golub and W. Kahan, "Calculating the Singular Values and Pseudo-Inverse of a Matrix," *Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis*, vol. 2, no. 2, pp. 205–224, 1965.
- [45] J. W. Demmel, *Applied Numerical Linear Algebra*. SIAM, 1997.
- [46] S.-C. T. Choi, *Iterative methods for singular linear equations and least-squares problems*. PhD thesis, Stanford University, 2006.
- [47] A. Thompson, "The response of radio-astronomy synthesis array to interfering signals," *IEEE Trans. on Antennas and Propagation*, vol. 30, pp. 450–456, May 1982.
- [48] A.-J. van der Veen, A. Leshem, and A.-J. Boonstra, "Array signal processing for radio astronomy," *Experimental Astronomy (EXPA)*, vol. 17, no. 1-3, pp. 231–249, 2004. ISSN 0922-6435.
- [49] S. van der Tol, *Bayesian Estimation for Ionospheric Calibration in Radio Astronomy*. PhD thesis, Delft University of Technology, 2009.
- [50] M. Viberg, B. Ottersten, and T. Kailath, "Detection and estimation in sensor arrays using weighted subspace fitting," *Signal Processing, IEEE Transactions on*, vol. 39, pp. 2436–2449, Nov 1991.
- [51] A.-J. van der Veen, S. Talwar, and A. Paulraj, "A subspace approach to blind space-time signal processing for wireless communication systems," *Signal Processing, IEEE Transactions on*, vol. 45, pp. 173–190, Jan 1997.
- [52] A.-J. van der Veen, M. Vanderveen, and A. Paulraj, "Joint angle and delay estimation using shift-invariance techniques," *Signal Processing, IEEE Transactions on*, vol. 46, pp. 405–418, Feb 1998.
- [53] C. Spearman, "The proof and measurement of association between two things," *The American Journal of Psychology*, vol. 15, pp. 72–101, Jan 1904.
- [54] D. N. Lawley, "Vi.—the estimation of factor loadings by the method of maximum likelihood," *Proceedings of the Royal Society of Edinburgh*, vol. 60, no. 01, pp. 64–82, 1940.

- [55] T. W. Anderson and H. Rubin, "Statistical inference in factor analysis," *In Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, vol. 5, pp. 111 – 150, 1956.
- [56] K. G. Jöreskog, "A general approach to confirmatory maximum likelihood factor analysis," *PSYCHOMETRIKA*, vol. 34, no. 2, 1969.
- [57] D. J. Bartholomew, M. Knott, and I. Moustaki, *Latent Variable Models and Factor Analysis: A Unified Approach*. John Wiley and Sons, 2011.
- [58] A.-J. van der Veen, A. Leshem, and A.-J. Boonstra, "Signal processing for radio astronomical arrays," in *Sensor Array and Multichannel Signal Processing Workshop Proceedings, 2004*, pp. 1–10, July 2004.
- [59] R. Roy and T. Kailath, "ESPRIT-Estimation of signals parameters via rotational invariance techniques," *IEEE Trans. Acoust. Speech Signal Processing*, vol. 38, pp. 984–995, July 1989.
- [60] R. Schmidt, *A Signal Subspace Approach to Multiple Emitter Location and Spectral Estimation*. PhD thesis, Stanford university, 1981.
- [61] T. Anderson, *An Introduction to Multivariate Statistical Analysis*. Wiley, third edition ed., 2003.
- [62] S. Kay, *Fundamentals of statistical signal processing: Estimation theory*. PTR, Prentice Hall, 1993.
- [63] A. K. Jagannatham and B. D. Rao, "Cramer-rao lower bound for constrained complex parameters," *IEEE SIGNAL PROCESSING LETTERS*, vol. 11, NOVEMBER 2004.
- [64] W. Ledermann, "On a problem concerning matrices with variable diagonal elements," *Proceedings of the Royal Society of Edinburgh*, vol. 60, pp. 1–17, 1 1940.
- [65] K. G. Jöreskog, "A general approach to confirmatory maximum likelihood factor analysis," *ETS Research Bulletin Series*, vol. 1967, no. 2, pp. 183–202, 1967.
- [66] S. Y. Lee, "The gauss-newton algorithm for the weighted least squares factor analysis," *Journal of the Royal Statistical Society*, vol. 27, June 1978.
- [67] K. Mardia, J. Kent, and J. Bibby, *Multivariate Analysis*. ACADEMIC PRESS, 1997.
- [68] J.-H. Zhao, P. Yu, and Q. Jiang, "ML estimation for factor analysis: EM or non-EM?," *Statistics and Computing*, vol. 18, pp. 109–123, 2008. 10.1007/s11222-007-9042-y.
- [69] D. Rubin and D. Thayer, "EM algorithms for ML factor analysis," *Psychometrika*, vol. 47, pp. 69–76, 1982. 10.1007/BF02293851.
- [70] A.-K. Seghouane, "An iterative projections algorithm for ml factor analysis," in *IEEE Workshop on Machine Learning for Signal Processing*, pp. 333 –338, oct. 2008.

- [71] J. F. Buss, G. S. Frandsen, and J. O. Shallit, "The computational complexity of some problems of linear algebra," *Journal of Computer and System Sciences*, vol. 58, no. 3, pp. 572 – 596, 1999.
- [72] A. Shapiro, "Weighted minimum trace factor analysis," *Psychometrika*, vol. 47, no. 3, pp. 243–264, 1982.
- [73] A. Shapiro, "Rank-reducibility of a symmetric matrix and sampling theory of minimum trace factor analysis," *Psychometrika*, vol. 47, no. 2, pp. 187–199, 1982.
- [74] J. Saunderson, V. Chandrasekaran, P. A. Parrilo, and A. S. Willsky, "Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting," *SIAM Journal on Matrix Analysis and Applications*, vol. 33, no. 4, pp. 1395–1416, 2012.
- [75] D. Knoll and D. Keyes, "Jacobian-free Newton–Krylov methods: a survey of approaches and applications," *Journal of Computational Physics*, vol. 193, no. 2, pp. 357 – 397, 2004.
- [76] V. Akcelik, G. Biros, and O. Ghattas, "Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation," in *Supercomputing, ACM/IEEE 2002 Conference*, pp. 41–41, Nov 2002.
- [77] S.-C. T. Choi, *Iterative methods for singular linear equations and least-squares problems*. PhD thesis, Stanford University, 2006.
- [78] S. Choi, C. Paige, and M. Saunders, "MINRES-QLP: A Krylov subspace method for indefinite or singular symmetric systems," *SIAM Journal on Scientific Computing*, vol. 33, no. 4, pp. 1810–1836, 2011.
- [79] O. Vinyals and D. Povey, "Krylov subspace descent for deep learning," *arXiv preprint arXiv:1111.4259*, 2011.
- [80] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*. Prentice Hall PTR, Jan. 1998.
- [81] P. Stoica and Y. Selen, "Model-order selection: a review of information criterion rules," *Signal Processing Magazine, IEEE*, vol. 21, pp. 36 – 47, July 2004.
- [82] T. J. Rothenberg, "Identification in parametric models," *Econometrica*, vol. 39, no. 3, pp. 577–591, 1971.
- [83] A. Shapiro, "Identifiability of factor analysis: some results and open problems," *Linear Algebra and its Applications*, vol. 70, no. 0, pp. 1 – 7, 1985.
- [84] K. Hayashi and G. A. Marcoulides, "TEACHER'S CORNER: Examining identification issues in factor analysis," *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 13, no. 4, pp. 631–645, 2006.

- [85] A.-J. van der Veen, A. Leshem, and A.-J. Boonstra, "Array signal processing for radio astronomy," in *The Square Kilometre Array: An Engineering Perspective* (P. Hall, ed.), pp. 231–249, Dordrecht: Springer, 2005. ISBN 1-4020-3797-x. Reprinted from *Experimental Astronomy*, 17(1-3), 2004.
- [86] A.-J. Boonstra and S. van der Tol, "Spatial filtering of interfering signals at the initial low frequency array (lofar) phased array test station," *Radio Science*, vol. 40, no. RS5S09, p. (16 pages), 2005. doi:10.1029/2004RS003135.
- [87] B. Guner, J. Johnson, and N. Niamsuwan, "Time and frequency blanking for radio-frequency interference mitigation in microwave radiometry," *IEEE Tr. Geoscience and Remote Sensing*, vol. 45, no. 11, pp. 3672–3679, 2007.
- [88] A. Offringa, A. de Bruyn, M. Biehl, S. Zaroubi, G. Bernardi, and V. Pandey, "Post-correlation radio frequency interference classification methods," *Monthly Notices of the Royal Astronomical Society*, vol. 405, no. 1, pp. 155–167, 2010.
- [89] G. Nita, D. Gary, Z. Liu, G. Hurford, and S. White, "Radio frequency interference excision using spectral-domain statistics," *Publications of the Astronomical Society of the Pacific*, vol. 119, no. 857, pp. 805–827, 2007.
- [90] P. Fridman, "RFI excision using a higher order statistics analysis of the power spectrum," *Astronomy and Astrophysics*, vol. 368, pp. 369–376, 2001.
- [91] P. Fridman, "Statistically stable estimates of variance in radio-astronomy observations as tools for radio-frequency interference mitigation," *The Astronomical Journal*, vol. 135, no. 5, p. 1810, 2008.
- [92] C. Barnbaum and R. Bradley, "A new approach to interference excision in radio astronomy: Real time adaptive cancellation," *The Astronomical Journal*, vol. 115, pp. 2598–2614, 1998.
- [93] S. Ellingson, J. Bunton, and J. Bell, "Cancellation of GLONASS signals from radio astronomy data," in *Proc. SPIE, Vol. 4015, Radio Telescopes* (H. R. Butcher, ed.), vol. 4015, pp. 400–407, 2000.
- [94] L. Li and B. Jeffs, "Analysis of adaptive array algorithm performance for satellite interference cancellation in radio astronomy," *URSI General Assembly*, Aug. 2002.
- [95] C. Hansen, K. Warnick, B. Jeffs, J. Fisher, and R. Bradley, "Interference mitigation using a focal plane array," *Radio science*, vol. 40, no. 5, 2005.
- [96] S. van der Tol and A.-J. van der Veen, "Performance analysis of spatial filtering of RF interference in radio astronomy," *Signal Processing, IEEE Transactions on*, vol. 53, pp. 896–910, March 2005.
- [97] F. Briggs, J. Bell, and M. Kesteven, "Removing radio interference from contaminated astronomical spectra using an independent reference signal and closure relations," *The Astronomical Journal*, vol. 120, pp. 3351–3361, 2000.

- [98] J. Kocz, F. Briggs, and J. Reynolds, "Radio frequency interference removal through the application of spatial filtering techniques on the Parkes multibeam receiver," *The Astronomical Journal*, vol. 140, no. 6, p. 2086, 2010.
- [99] B. Jeffs, K. Warnick, and L. Li, "Improved interference cancellation in synthesis array radio astronomy using auxiliary antennas," in *IEEE ICASSP*, (Hong Kong), Apr. 2003.
- [100] B. Jeffs, L. Li, and K. Warnick, "Auxiliary antenna-assisted interference mitigation for radio astronomy arrays," *IEEE Tr. Signal Processing*, vol. 53, pp. 439–451, Feb. 2005.
- [101] G. Hellbourg, R. Weber, C. Capdessus, and A.-J. Boonstra, "Cyclostationary approaches for spatial RFI mitigation in radio astronomy," *Comptes Rendus Physique*, vol. 13, no. 1, pp. 71 – 79, 2012. The next generation radiotelescopes / Les radiotélescopes du futur.
- [102] B. Porat, *Digital processing of random signals—Theory and methods*. Englewood Cliffs (NJ): Prentice Hall, 1994.
- [103] A. Leshem and A.-J. van der Veen, "Multichannel detection of Gaussian signals with uncalibrated receivers," *IEEE Signal Processing Letters*, vol. 8, no. 4, pp. 120–122, 2001.
- [104] S. Kazemi, S. Yatawatta, and S. Zaroubi, "Radio interferometric calibration via ordered-subsets algorithms: OS-LS and OS-SAGE calibrations," *Monthly Notices of the Royal Astronomical Society*, vol. 434, no. 4, pp. 3130–3141, 2013.
- [105] M. R. Osborne, "Fisher's method of scoring," *International Statistical Review / Revue Internationale de Statistique*, vol. 60, no. 1, pp. pp. 99–117, 1992.
- [106] Salvini, Stefano and Wijnholds, Stefan J., "Fast gain calibration in radio astronomy using alternating direction implicit methods: Analysis and applications," *A&A*, vol. 571, p. A97, 2014.
- [107] M. A. Woodbury, *Inverting Modified Matrices*. No. 42 in Statistical Research Group Memorandum Reports, Princeton, NJ: Princeton University, 1950.
- [108] R. Reid, "Smear fitting: a new image-deconvolution method for interferometric data," *Monthly Notices of the Royal Astronomical Society*, vol. 367, no. 4, pp. 1766–1780, 2006.
- [109] H. Junklewitz, M.R. Bell, and T. Enßlin, "A new approach to multifrequency synthesis in radio interferometry," *A&A*, 2015.
- [110] M. Lochner, I. Natarajan, J. T. Zwart, O. Smirnov, B. A. Bassett, N. Oozeer, and M. Kunz, "Bayesian inference for radio observations," *Monthly Notices of the Royal Astronomical Society*, vol. 450, no. 2, pp. 1308–1319, 2015.

- [111] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. V. der Vorst, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, 2nd Edition*. Philadelphia, PA: SIAM, 1994.
- [112] J. Capon, "High resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE*, pp. 1408–1418, 1969.
- [113] B. Frieden, "Restoring with maximum likelihood and maximum entropy," *Journal of the Optical Society of America*, vol. 62, pp. 511–518, 1972.
- [114] S. Gull and G. Daniell, "Image reconstruction from incomplete and noisy data," *Nature*, vol. 272, pp. 686–690, 1978.
- [115] R. Levanda and A. Leshem, "Adaptive selective sidelobe canceller beamformer with applications to interference mitigation in radio astronomy," *Signal Processing, IEEE Transactions on*, vol. 61, pp. 5063–5074, Oct 2013.
- [116] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [117] C.-Y. Chen and P. Vaidyanathan, "Quadratically constrained beamforming robust against direction-of-arrival mismatch," *Signal Processing, IEEE Transactions on*, vol. 55, pp. 4139–4150, Aug 2007.
- [118] Y. Gu and A. Leshem, "Robust adaptive beamforming based on interference covariance matrix reconstruction and steering vector estimation," *Signal Processing, IEEE Transactions on*, vol. 60, pp. 3881–3885, July 2012.
- [119] D. C.-I. Fong, *Minimum-Residual Methods for Sparse Least-Squares Using Golub-Kahan Bidiagonalization*. PhD thesis, Stanford University, 2011.
- [120] P. Shaman, "The inverted complex wishart distribution and its application to spectral estimation," *Journal of Multivariate Analysis*, vol. 10, no. 1, pp. 51 – 59, 1980.

SUMMARY

The search for the answer to one of the most fundamental scientific questions, “How was the universe formed?”, requires us to study very weak radio signals from the early universe. In the last eighty years, radio astronomers have been able to use radio frequency observations for significant discoveries such as quasars, supermassive Black Holes and the Cosmic Microwave Background radiation. Radio astronomers use a radio telescope to study the cosmos. A radio telescope usually consists of an *array* of radio receivers (antennas) and supporting hardware/software to produce synthesized images of the sky. While the earlier generation of the radio telescopes such as the Westerbork Synthesis Radio Telescope (WSRT), the Very Large Array (VLA) and the Giant Meter-wave Radio Telescope (GMRT) consisted of 14-45 receivers separated a few kilometers (3-25 km baselines), the next generation of radio telescopes such as LOFAR and SKA have thousands of receivers which cover distances of over 1000 km. This massive increase in the number of receivers and the geometric dimensions is a consequence of the required (high) resolution and sensitivity for modern scientific studies and while it is necessary, it does not guarantee the desired results without the appropriate data and signal processing.

The main challenges in radio astronomy can be divided in three closely related problems: mitigation of man-made radio frequency interference, calibration and image formation. The main goal of this thesis is to investigate how the signal processing formalism can be used to systematically model and analyze these three problems and what signal processing tools are needed for addressing them.

The number of RFI free bands is diminishing rapidly as a consequence of the increased number of wireless services and applications. The shift towards wideband digital systems has created new problems which are not sufficiently addressed by currently implemented RFI detection and mitigation systems. For this class of continuously present wide-band RFI, the use of array processing techniques such as spatial filtering could provide access to frequency bands otherwise unusable by astronomers. Such a spatial filtering can be achieved by estimating and removing the subspace that the interfering signal is occupying. Many signal processing algorithms use the eigenvalue decomposition (EVD) for estimating the signal subspace. However the use of EVD is limited to systems where the noise is white or known from calibration. This requirement is a limiting factor for applying these techniques to uncalibrated arrays with unknown noise models. In these situations a more generic approaches which allows for combined RFI filtering and noise power calibration is preferred. In this thesis factor analysis (FA) is proposed as suitable substitution for EVD.

FA is a technique that allows for the decomposition of the signal into a low-rank part corresponding to the signal and a diagonal part which represents the covariance of the noise on the receivers. Because the diagonal elements can be different this technique can be used when the noise is not white and forms a generalization of the EVD. In RFI mitigation applications the signal part of the data is dominated by RFI and changes more

rapidly than the noise. Estimating the noise covariance which is shared by several measurements jointly allows for a more accurate estimation. As a result extensions to the classical FA are proposed to improve the estimates for the diagonal part of the decomposition in a joint fashion. Even a diagonal noise structure can be limiting in some applications. For example the contribution of the Milky Way affects the short baselines which can be modeled by using a non-diagonal covariance matrix. The FA model can be extended for this type of signals. An extension to FA called Extended FA (EFA) is used to allow for capturing such structures into the model. Similar to JFA we can also estimate the parameters in EFA jointly, and the resulting method is denoted by Joint EFA (JEFA). Using nonlinear optimization techniques combined with Krylov subspace based solvers an scalable algorithm is developed. The statistical efficiency of this algorithm is shown by comparing its results to the Cramér–Rao bound and its application in RFI mitigation has been demonstrated on measurements from the WSRT and LOFAR.

Antenna gain calibration is an essential step in producing accurate images. Using common array processing data models, gain calibration is formulated as a nonlinear covariance matching problem. In this thesis we show that the matrices involved in this estimation problem are highly structured and that the system of equations involving these matrices can be efficiently solved using Krylov subspace based solvers (similar to JEFA). The resulting calibration algorithm is scalable and requires a low number of iterations in order to converge which makes it an attractive alternative to currently available techniques.

Both classical and parametric based image formations consist of two steps. First a “dirty” image is constructed from the measurements and then an improved estimate is found by performing a deconvolution step. When the number of pixels on the image becomes large, the deconvolution step becomes an ill-posed problem. In this thesis we show that image values are bounded from below by a nonnegativity constraint and above by the dirty image. Using beamforming techniques, we show that tighter upper bounds can be constructed using the MVDR beamformer. These bounds allow us to regularize the deconvolution problem by a set of inequality constraints. Following a signal processing model, the image formation is then formulated as a parameter estimation problem with inequality constraints. This optimization problem can be solved using an active set algorithm. We show that, with the right initialization, the active set steps are very similar to sequential source removing techniques such as CLEAN. This connection between classical approaches and parametric imaging techniques provides the necessary theoretical basis for further analysis and allows for improving both methods.

Based on the results presented in this thesis we can conclude that signal processing methodologies can provide new solutions to the radio astronomical problems and also shed light on the inner working of the classical techniques. Hence, a signal processing approach is extremely beneficial in tackling the problems that the next generation of radio telescopes will face.

SAMENVATTING

De zoektocht naar het antwoord op een van de meest fundamentele wetenschappelijke vragen, “Hoe is het heelal ontstaan?”, vraagt ons om zeer zwakke radiosignalen van het vroege heelal te bestuderen. In de afgelopen tachtig jaar, zijn de radio-astronomen in staat geweest om radiofrequentie waarnemingen te gebruiken voor het doen van belangrijke ontdekkingen zoals quasars, superzwarte zwart gaten en de kosmische achtergrondstraling. Radio astronomen maken gebruik van een radiotelescoop om de kosmos te bestuderen. Een radiotelescoop bestaat meestal uit een reeks van radio-ontvangers (antennes) en ondersteunende hardware / software om gesynthetiseerde beelden van de hemel te produceren. Terwijl de eerdere generatie van de radiotelescopieën zoals de Westerbork Synthesis Radio Telescoop (WSRT), de Very Large Array (VLA) en de Giant Meter-wave Radio Telescoop (GMRT) uit 14-45 ontvangers met een onderlinge afstand van een paar kilometer (3-25 km basislijnen) bestonden, de volgende generatie radiotelescopieën zoals LOFAR en SKA hebben duizenden ontvangers die verspreid zijn over afstanden tot meer dan 1000 km. Deze sterke toename van dimensie en het aantal ontvangers is het gevolg van de vereiste (hoge) resolutie en gevoeligheid die voor moderne wetenschappelijke studies noodzakelijk zijn. Echter kan deze vergroting van de basislijnen en het aantal ontvangers de gewenste resultaten zonder de juiste data en signaalverwerking niet garanderen.

De belangrijkste uitdagingen in radioastronomie kunnen worden onderverdeeld in drie nauw verwante problemen: wegfilteren van de door mens gemaakte radiofrequentieinterferenties (RFI), kalibratie en beeldvorming. Het belangrijkste doel van dit proefschrift is om te onderzoeken hoe de signaalverwerking formalisme kan worden gebruikt om op een systematische wijze deze drie problemen aan te pakken en welke gereedschappen uit signaalverwerking nodig zijn voor het modelleren en analyseren van deze problemen.

Het aantal frequentiekanalen die geen last hebben van RFI neemt als gevolg van de toename van het aantal draadloze diensten en toepassingen snel af. De migratie naar breedbandige digitale systemen heeft nieuwe problemen gecreëerd die niet voldoende door de huidige implementaties van RFI detectie- en mitigatie systemen kunnen worden opgelost. Voor de klasse van continu aanwezig breedbandige RFI kan het gebruiken van matrix verwerkingstechnieken zoals spatiele filtering toegang geven tot frequentiebanden die anderszins voor astronomen onbruikbaar zijn. Een dergelijke spatiele filtering kan worden bereikt door het schatten en verwijderen van de deelruimte waar het interfererende signaal zich in bevindt. Veel signaalverwerkingsalgoritmen maken gebruik van de eigenwaardenontbinding (EWO) voor het schatten van de signaal-deelruimte. Echter het gebruik van EWO is beperkt tot systemen waarop de ruis wit of bekend is uit de kalibratie. Deze eis is een beperkende factor voor de toepassing van deze technieken op niet gekalibreerde systemen die een onbekende ruismodel hebben. In deze situaties is het gebruiken van methodes die een combinatie van RFI filtering en ruiskalibratie moge-

lijk maken meer gewenst. In dit proefschrift wordt Factor Analyse (FA) als een geschikte vervanging voor de EWO voorgesteld.

FA is een techniek waarmee een covariantiematrix wordt opgesplitst in een lage-rang deel en een diagonaal deel. Het deel met de lage rang kan de deelruimte van het gewenste signaal modelleren en het diagonale deel kan worden gebruikt voor het modelleren van ruis. Omdat de diagonale elementen verschillend kunnen zijn, kan deze techniek worden gebruikt in situaties waar ruis niet wit is en dus FA kan als een veralgemenisering van EWO gezien worden. Bij RFI gerelateerde toepassingen wordt de signaal-deelruimte gedomineerd door RFI die in meeste gevallen veel sneller verandert dan de ruis. Hierdoor wordt de ruis-covariantiematrix gedeeld tussen meerdere metingen. Het gezamenlijk schatten van deze matrix zorgt voor een nauwkeurigere schatting. In dit proefschrift is een uitbreiding van de klassieke FA voorgesteld die zo'n gezamenlijke schatting mogelijk maakt. Ook het modelleren van de ruis-covariantiematrix met een diagonale matrix kan beperkend zijn voor sommige toepassingen. Bijvoorbeeld de bijdrage van de Melkweg beïnvloedt de korte basislijnen en kan door een diagonale covariantiematrix niet worden gemodelleerd. Het FA model kan worden uitgebreid voor dit soort signalen. Wij stellen voor om een uitbreiding van FA, genoemd Extended FA (EFA), te gebruiken voor het modelleren van dergelijke structuren. Net als bij JFA kunnen we ook de parameters in EFA op een gezamenlijke wijze schatten. De methode die hiervoor ontwikkeld is noemen wij Joint EFA (JEFA). Met behulp van niet-lineaire optimalisatietechnieken gecombineerd met op Krylov deelruimte gebaseerde algoritmes hebben wij een schaalbare techniek ontwikkeld voor het oplossen van de onbekenden in het (JE)FA model. De statistische efficiëntie van dit algoritme is aangetoond door het vergelijken van de resultaten met de Cramér-Rao grens en de toepassing ervan is gedemonstreerd door middel van het wegfilteren van RFI uit metingen die door de WSRT en LOFAR gemaakt zijn.

Antenneversterking kalibratie is een essentiële stap in het maken van nauwkeurige beelden. Door het gebruiken van standaard signaalbewerking modellen, hebben wij de antenneversterking kalibratie geformuleerd als een niet-lineaire covariantie matching probleem. In dit proefschrift laten we zien dat de matrices die bij dit schattingsprobleem voorkomen zeer gestructureerd zijn en dat het stelsel van lineaire vergelijkingen van deze matrices efficiënt kan worden oplost door het gebruik maken van op Krylov deelruimte gebaseerde methoden (vergelijkbaar met JEFA). Het resulterende kalibratie-algoritme is schaalbaar en vereist een klein aantal iteraties om te convergeren en daardoor vormt dit een aantrekkelijk alternatief voor huidige technieken.

Zowel klassieke technieken als op parametrische model-schatting gebaseerde technieken voor beeldvorming bestaan uit twee stappen. Eerst wordt een "dirty image" opgebouwd uit de metingen en vervolgens wordt een verbeterde schatting gevonden door het uitvoeren van een deconvolutie stap. Wanneer het aantal pixels in het beeld groot is, wordt de deconvolutie stap een slecht-gesteld probleem. In dit proefschrift laten we zien dat de intensiteit van pixels begrensd is tussen nul en de waarde in de dirty image. Door het gebruik maken van bundelvormingstechnieken tonen we aan dat een zeer strakke bovengrens kan worden geconstrueerd met de MVDR bundelvormer. Met behulp van deze grenzen kunnen we het deconvolutieprobleem regulariseren door een set van ongelijkheden. Door middel van een signaalverwerkingsmodel wordt de beeldvorming geformuleerd als een optimalisatieprobleem met ongelijkheden. Dit optimalisatiepro-

bleem kan worden opgelost met behulp van een active-set algoritme. We tonen aan dat met de juiste initialisatie, de active-set stappen vergelijkbaar zijn met een sequentiële bronverwijderingstechniek zoals CLEAN. Deze relatie tussen klassieke benaderingen en op parametrische schatting gebaseerde beeldvormingstechnieken geeft de benodigde theoretische basis voor verdere analyse en verbetering van beide methoden.

Op basis van de resultaten in dit proefschrift kunnen we concluderen dat signaalverwerkingsmethodologie nieuwe oplossingen kan bieden aan radio-astronomische problemen en ook licht kan werpen op de interne werking van de klassieke technieken. Daarom is een signaalverwerkingsaanpak zeer gunstig voor het oplossen van de problemen waar de volgende generatie van radiotelescopen mee zal worden geconfronteerd.

ACKNOWLEDGMENT

I would like to use this opportunity to express my gratitude to the people without whom this work would not have been possible. Special thanks go to my supervisor Professor Alle-Jan van der Veen for his great wisdom, endless patience, support, and understanding throughout my journey in putting together this thesis. I would also like to thank Dr. Stefan Wijnholds, Prof. Amir Leshem, Dr. Rob Remis and Prof. Geert Leus for their invaluable ideas, comments, and feedback.

I have been lucky and privileged to collaborate daily with the brilliant members of the CAS group. It has been a deep pleasure to work with Seyran Khademi, Hadi Jamali Rad, Shahrzad Naghibzade, Raj Rajan, Hamid Ramezani, Sundeep Chepuri, Venkat Roy, Jorge Castaneda, Andrea Simonetto and all other colleagues at TU Delft.

I owe my existence and hence this work to my loving, dearest, and selfless parents Golbahar Tahmasebi Alijanvand and Bahman Mouri Sardarabadi for whom I will forever be grateful. I extend this gratitude to my beloved sisters Leila and Mehrnoosh for their unconditional love and support. Lastly, I would like to thank my lovely partner Shima whose encouraging words and constant support gave me the strength I needed.

BIOGRAPHY



Ahmad Mouri Sardarabadi was born in Ahvaz, Iran in 1985. He received the B.Sc. degree (cum laude) in 2009 from TU Delft, The Netherlands. In 2011, he received the M.Sc. degree (cum laude) from TU Delft for a study on the applications of Factor Analysis in array and signal processing. Since 2011 he has been pursuing his Ph.D. degree at TU Delft, working on signal processing for radio astronomy. His research interests are signal processing and multivariate analysis, with applications to radio astronomy.