

School of Electrical Engineering, Computing and
Mathematical Sciences

Calibration Systematics in the Search for Cosmic
Reionisation

Ronniy Christoffer Joseph

0000-0003-3457-4670

This thesis is presented for the Degree of
Doctor of Philosophy
of
Curtin University

April 2020

To the best of my knowledge and belief this thesis contains no material previously published by any other person except where due acknowledgement has been made. This thesis contains no material which has been accepted for the award of any other degree or diploma in any university.

Ronniy Christoffer Joseph

17 April 2020

“You can use logic to justify almost anything. That’s its power and its flaw.”

— Captain Kathryn Janeway

Acknowledgements

I begin by acknowledging the Traditional Owners of the land on which I have made my home for the past 3.5 years and performed the majority of this work; the Whadjuk people of the Nyoongar Nation, and pay my respects to Elders past and present.

This PhD was an amazing journey and impossible without the many people I met along the way. First of all, thank you Cathryn Trott for taking me on a journey to what used to be the other side of the world. It was inspiring to work alongside you and see your scientific curiosity, drive, and kindness in action. You allowed me to find my own path, and your never-ending support enabled me to shape that path. Thanks to your generosity I have attended many exciting conferences where I met incredible scientists from all over the world. I can not really convey my gratitude for everything you offered me throughout my PhD in this simple paragraph. For my co-supervisor Randall Wayth, your pragmatism and optimism has kept me grounded and my work realistic. You always ensured that I knew where I was going. My meetings with both of you have always been the highlight of my working week, filled with both science and many good laughs.

I am grateful for being part of such a great research group. Locally, I have always been able to ask many questions of Chris Jordan, Steven Murray, Christene Lynch, Jack Line, and Ben McKinley. Beyond our local group I have had many useful questions from Bart Pindor, Nichole Barry, and Aman Choksi. In-

ternationally, I had many thought-provoking discussions with Ruby Byrne, Mike Wilensky, Miguel Morales, Wenyang Li and Jonnie Pober. You have contributed to this work in one meeting or another, but in particular thank you for making me feel welcome to the group, and of course the karaoke sessions. Beyond our EoR group I owe much gratitude to Adrian Sutinjo, Marcin Sokolowski, Daniel Ung, Andrew Williams, Jonathan Sievers, and Hsin Cynthia Chiang.

Special thanks go to ICRAR's amazing outreach team, Greg Rowbotham, Cassandra Rowles, Kirsten Gottschalk and Pete Wheeler for involving me in outreach, to Rich Plotkin and John Morgan for getting me in front of the classroom, and Gemma Anderson for taking me onto DevCom. All of this non-research relief from time to time has kept me sane.

Support comes in various ways and you can find plenty of it in the student cohorts at both nodes of ICRAR. I have had many adventures with my fellow students, especially with Vlad Tudor, Ryan Urquhart, Robin Cook, Ahmed Elagali, Lucie Bakels, Rhys Poulton, Elliot Charlton, Ben Quici, and Kat Ross. We have had some great times together. In particular, I want to thank Pikky Atri for our many (conference) journeys together and Bella Nasirudin for our shared journeys from real to power spectrum space. Your friendships have been invaluable.

I would not be here without my mother's support throughout my adventures. You have done an amazing job at raising both me and my brother, and enabled us to follow our interests. The same goes for my little brother who has been there at every step along the way despite being 14,183 kilometres away, and my friends in various bits of the world. Even though we are continuously changing (continents), our friendship never fades.

Finally, to Robin. Thank you for joining me on this adventure in Australia, the many we had on our way here, and for enduring my ramblings about both actual space and Fourier space. I look forward to our next adventures as we make our way around the world.

Abstract

Detecting a redshifted neutral hydrogen signal from the Epoch of Reionisation (EoR) is one of the most promising probes into the formation history of structure in the Universe. The advent of a new generation of low-frequency radio interferometers has opened this direct window into the EoR by enabling access to the 21-cm signal. However, it is well established that calibration errors have the potential to inhibit an EoR detection by introducing additional spectral features that mimic the structure of EoR signals. In this thesis I investigate the limitations of sky and redundancy based calibration and quantify their impact on the estimation of the 21-cm power spectrum (PS).

Through simulations of the relative calibration step in redundancy based methods, I study the influence of the flux distribution of the radio sky and the impact of antenna position offsets on the complex calibration solutions. I find that the position offsets introduce a bias into the phase component of the calibration solutions. This phase bias increases with the distance between bright radio sources and the pointing centre, and with the flux density of these sources. Our results indicate that redundant calibration outperforms sky-based calibration for target fields that lack bright sources, in part due to the high positional precision of the antenna tiles in the Murchison Widefield Array (MWA).

In addition to these simulations, I present an analytic model to estimate the impact of a wide variety of modelling residuals in EoR PS estimation. I apply

this framework to the specific case of broken dipoles in the MWA to understand its effect and estimate its impact on PS estimation. I find that incorrect beam modelling introduces bias that is three orders of magnitude lower than current lowest limits on the PS. However, I expect that this bias is two orders of magnitude higher than the expected EoR signal. Determining the accuracy of both current beam models and direction dependent calibration pipelines is therefore crucial in our search for an EoR signal.

I conclude this thesis with a comparison of the performance of redundant and sky-based calibration under current sky modelling limits and array deployment tolerances for the MWA. The results for the MWA indicate that in wide-field arrays sky based calibration is marginally more robust against calibration errors. The results for the Hydrogen Epoch of Reionization Array (HERA) indicate that our model underestimates the expected calibration error and hence improvements are necessary to accurately estimate the expected contamination. However, the comparable errors in the MWA indicate that hybrid methods employing both sky based and redundancy based information are the way forward.

Preface

The contents of Chapter 5 are my own work, except for the input and supervision from Dr. Cathryn Trott and Dr. Randall Wayth during the development of the redundant calibration simulations and theoretical analysis of sky based calibration errors.

The contents of Chapter 6 are also my work, except for the input and close supervision from Dr. Cathryn Trott and Dr. Randall Wayth during the development of the analytic description of beam modelling errors in the Murchison Widefield Array. Ainulnabilah Nasirudin provided suggestions that aided in the rapid development of visualisation tools for the results and their interpretation.

Contents

Acknowledgements	vii
Abstract	ix
Preface	xi
List of Figures	xix
1 Introduction	1
1.1 A Brief History of the Universe	1
1.1.1 The Big Bang	2
1.1.2 The Dark Ages	3
1.1.3 A Cosmic Dawn	4
1.2 Rationale	4
1.3 Thesis Outline	5
2 The Epoch of Reionisation	7
2.1 Current constraints	7
2.1.1 Cosmic Microwave Background	8
2.1.2 Quasar Spectra	8

2.1.3	High-redshift galaxies	9
2.2	21-cm Physics	11
2.2.1	The Hyperfine Transition	11
2.2.2	The Spin Temperature	12
2.3	The First Luminous Sources	15
2.3.1	Ultraviolet Sources	15
2.3.2	X-ray Sources	16
2.3.3	The Fiducial Model	16
2.3.4	The Differential Brightness Temperature	17
2.4	21-cm Cosmology	20
2.4.1	The Evolution of the 21-cm signal	20
2.4.2	The 21-cm Power Spectrum	22
2.4.3	Alternative Techniques	23
2.5	Challenges of the EoR Experiment	23
3	Radio Interferometry	25
3.1	The Epoch of Radio Astronomy	25
3.2	Third Generation Radio Telescopes	26
3.3	Interferometry Fundamentals	30
3.3.1	The Measurement Equation	32
3.3.2	<i>uv</i> -coverage	34
3.3.3	Earth Rotation Synthesis	35
3.4	The Primary Beam	36
3.4.1	Sidelobe Response	37
3.4.2	Phased Arrays	38
3.5	Calibration	39

3.5.1	Antenna Gain	40
3.5.2	Bandpass	40
3.5.3	Ionospheric Effects	41
3.6	Sky-Based Calibration	41
3.6.1	Recent Developments	42
3.7	Redundant Calibration	44
3.7.1	Relative Calibration	44
3.7.2	Absolute Calibration	48
3.7.3	Recent Developments	48
3.8	Comparing Calibration Strategies	49
4	Statistical Cosmology	51
4.1	The 21-cm Power Spectrum	51
4.1.1	Mathematical Formalism	52
4.1.2	Estimation Theory	53
4.2	Power Spectrum Estimators	57
4.3	Relating observing and cosmological units	58
4.4	The Window and the Wedge	59
4.5	Power Spectrum Contamination	61
4.5.1	Gridding	61
4.5.2	RFI	62
4.5.3	Foregrounds	62
4.5.4	Calibration	63
4.6	Power Spectrum Limits	64
5	The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions	67

5.1	Introduction	68
5.2	Sky Model Calibration	71
5.2.1	Model Incompleteness Uncertainty	72
5.2.2	Model Incompleteness Bias	74
5.3	Redundant Baseline Calibration	77
5.3.1	Simulating the Bias and Uncertainty	77
5.3.2	The Sky Dependent Uncertainty	78
5.3.3	The Position Offset Bias	82
5.4	Comparing Sky and Redundancy Based Calibration	84
5.5	Discussion	88
5.6	Conclusion and Future Work	95
6	Calibration and 21-cm Power Spectrum Estimation in the Presence of Antenna Beam Variations	97
6.1	Introduction	98
6.2	EoR Signal Estimation with Gain Calibration Errors	100
6.3	The Residual Covariance Matrix	103
6.3.1	Sky Covariance Matrix	104
6.3.2	Beam Covariance Matrix	105
6.4	The Phased Array Beam Model	108
6.4.1	A Broken MWA Dipole	109
6.5	The Gain Error Covariance Matrix	112
6.6	The Gain Corrected Residual Covariance Matrix	115
6.7	Results	117
6.7.1	Sky Model Errors	118
6.7.2	Comparing Sources of Error	119

6.8	Discussion	120
6.9	Conclusion	123
7	Comparing Calibration Systematics in Redundant Arrays Searching for Cosmic Reionisation	125
7.1	Introduction	125
7.2	Generalised Calibration Analysis	128
7.3	Sky Based Calibration Errors	128
7.4	Redundant Calibration Errors	129
7.5	CRLB on Gain Estimates	133
7.5.1	Sky Based Calibration Errors	134
7.5.2	Redundant Calibration Errors	135
7.5.3	Comparison of lower bounds	136
7.6	Spectral Structure of Gain Errors	137
7.6.1	Sky Based Calibration	140
7.6.2	Redundant Calibration	141
7.6.3	Comparing Power Spectrum Contamination	144
7.7	Discussion	147
7.8	Conclusions	150
8	Summary and Conclusions	151
8.1	Position Errors in Redundant Calibration	152
8.2	Beam Variation in Sky Based Calibration	152
8.3	Comparing Redundant and Sky Based Calibration	153
8.4	Conclusion and Outlook	155
8.5	Closing Remarks	156

Appendices	157
A Circumventing Phase Wrapping in Redundant Calibration	159
A.1 The Multi-Frequency Implementation	159
B The Covariance Matrix of Beam Modelling Errors	163
B.1 Propagating Covariance Matrices to Power Spectrum space	163
B.2 The General Beam Covariance Matrix	164
B.3 Beam Variations due to Missing Dipoles	166
B.4 Linearised Gain Error Covariance Matrix	169
C The Covariance Matrix of Non-Redundant Errors	171
C.1 The Covariance Matrix of Non-Redundant Errors	171
D Agreement of Co-authors	177
D.1 Statement of Originality	177
D.2 Co-author Responses	178
E Copyright Information	179
F Chapter Heading	181
Bibliography	185

List of Figures

1.1	The timeline of the Universe	2
2.1	Lyman- α Forest and Gunn-Peterson trough in high- z quasars . .	10
2.2	Neutral hydrogen spin flip	12
2.3	The energy levels involved in the Wouthuysen-Field effect.	14
2.4	The zoo of global Epoch of Reionisation signal models	21
3.1	Schematic overview of the geometry in radio astronomy	30
3.2	Instantaneous uv -coverage for the MWA Phase II	34
3.3	The MWA primary beam: its main lobe and sidelobes	36
3.4	A minimum redundant array and the redundant baselines in this array.	44
4.1	Theoretical 21-cm power spectra for a faint galaxy driven reioni- sation model	54
4.2	A schematic overview of the 2D power spectrum	60
4.3	The latest upper limits on the 21-cm power spectrum	66
5.1	A phasor perspective on errors and models	76
5.2	Calibration solutions from relative calibration as a function of source position	79
5.3	The distribution of phase wraps in relative calibration	81

5.4	Calibration solutions in relative calibration with position errors	83
5.5	The magnitude of the offset residual $\hat{\mathbf{b}}$ for an offset in a given antenna	84
5.6	MWA tile positions and deviations from redundancy	86
5.7	The bias and uncertainty of relative calibration as a function of source flux and positioning error	89
5.8	The bias and uncertainty of relative calibration as a function of source position and positioning error	90
5.9	1-dimensional view on the bias and uncertainty of redundancy and sky based calibration as a function of source flux	91
5.10	1-dimensional view on the bias and uncertainty of redundancy and sky based calibration as a function of source location	92
5.11	The Haslam map at Megahertz frequencies	95
6.1	Broken tiles in the Murchison Widefield Array	110
6.2	Power spectrum of residuals before sky based calibration	111
6.3	Statistical visibility approximation of a stochastic point source sky	114
6.4	Histogram of MWA baselines and power spectrum of the gain errors	116
6.5	Fiducial Epoch of Reionisation signal	118
6.6	Residuals before and after calibration with an incomplete sky model	119
6.7	Residuals after calibration with sky and beam modelling errors	120
6.8	Expected residuals with sky and beam modelling errors for the MWA	121
7.1	CRLB on the gain parameters	138
7.2	Power spectrum of residuals before redundant calibration	141
7.3	Power spectrum of residuals after relative calibration	142
7.4	Power spectrum of residuals after absolute calibration	144
7.5	Power spectrum of residuals after and sky and redundancy based calibration with the MWA	145

7.6	Power spectrum of residuals after sky and redundancy based calibration with HERA	146
A.1	Calibration solutions after multi-frequency relative calibration . .	161

Chapter 1

Introduction

Searching for a needle in a haystack is probably the best way to describe the hunt for signals from the very first stars and galaxies in the Universe. The radio signals from hydrogen gas surrounding those first objects are extremely weak compared to signals from galaxies nearby and far away, our own Milky Way, and artificial radio signals. Nevertheless, driven by our curiosity to understand the history of events that led to our existence, astronomers around the world are undertaking extremely sensitive experiments to detect signals from a poorly understood time in the Universe. This thesis aims to add to the body of knowledge concerning those experiments.

1.1 A Brief History of the Universe

The origin of the Universe is described in terms of the Big Bang. Using the theory of general relativity (Einstein, 1915) and the cosmological principle, i.e. the Universe is homogeneous and isotropic, one can derive the Friedmann–Lemaître–Robertson–Walker metric (Friedmann, 1922; Lemaître, 1927; Robertson, 1935; Walker, 1937). This solution to the field equations of general relativity, and the observations of receding galaxies (Hubble, 1926) led to the description of a homogeneous, isotropic, and expanding Universe. Lemaître realised that reversing

time on this solution implied there once was a moment in time when the entire Universe was compressed into a singularity. A moment in time we now refer to as the Big Bang. In this section we will give a short summary of the history of the Universe.




Figure 1.1: An illustrative timeline of the Universe. The timeline starts on the left with the Big Bang followed by rapid expansion. The next major event is the emission of the Cosmic Microwave background that leads into the Cosmic Dark Ages. The Universe then heats up due to the emergence of the first luminous sources that form ionisation bubbles around them, marking the Epoch of Reionisation. The reionisation process removes neutral hydrogen and transitions into the a structured Universe filled with a diverse zoo of galaxies as observed in the Hubble Ultra Deep Fields. CMB image credit: NASA / WMAP Science Team. HDF image credit: NASA, ESA, HUDF09 Team

1.1.1 The Big Bang

From this singularity the Universe emerges in a hot dense state, and it is not long before inflation causes the Universe to expand quickly. The expansion cools the Universe adiabatically, which starts the formation of matter and anti-matter that continuously annihilate each other. However, due to a small asymmetry be-

tween the two, only matter survives the annihilation process. From the remaining protons, neutrons, electrons and neutrinos, the main constituents of the present day Universe form; hydrogen, helium, and lithium, a process called Big Bang Nucleosynthesis (Gamow, 1946). This all takes a mere 3 minutes after the Big Bang. Although the continuing expansion of the Universe causes it to cool, it takes another 380,000 years before the time of recombination. During this time the primordial gas cools enough, allowing electrons to combine with atomic nuclei to form neutral atoms. As the recombination process decreases the number of free electrons, the efficiency of scattering decreases, and so the Universe becomes transparent to radiation. This transparency allows the photons that were continuously scattered to propagate freely through the Universe, hence we call this moment the time of last scattering. These free photons form the afterglow of the Big Bang we observe today as the Cosmic Microwave Background (CMB).

1.1.2 The Dark Ages

Observations of the CMB revealed that its temperature is extremely uniform across the sky – 2.7 K with spatial fluctuations on the order of $\sim 20 \mu\text{K}$ (Penzias & Wilson, 1965; Mather et al., 1994; Bennett et al., 2013; Planck Collaboration et al., 2016). These small variations in temperature are caused by the density fluctuations in the energy-matter field, the velocity of matter, gravitational redshift as photons climb out of high density regions at the time of last scattering, and by interactions between CMB photons and matter on their way to an observer. The dense regions in the matter field in the early Universe grew in mass under the influence of gravity. Higher density regions attracted more matter than lower density regions. In our current model of the Universe, the ΛCDM model, dark matter collapsed first. Baryonic matter on the other hand cannot collapse because it is halted by the gas pressure of the baryons. This pressure continues to decrease because the gas cools adiabatically while the Universe expands.

1.1.3 A Cosmic Dawn

Once baryonic matter cools sufficiently due to a combination of adiabatic cooling, atomic line cooling, and molecular cooling, it starts to fall into the gravitational wells of the dark matter halos. It is in these halos where the density of matter starts to increase despite the expansion of the Universe. This is where the first seeds of the structures we see around us today are formed; stars, black holes, galaxies and galaxy clusters. The objects formed at this time end the Dark Ages of the Universe as their photons light up the Universe, and their subsequent evolution changes the face of the Universe forever.

1.2 Rationale

It still unclear to us what the exact conditions were of the early Universe and what the nature was of the very first luminous objects. We can form a list with the most likely suspects; black hole seeds that led to the most massive black holes we see today in the centres of galaxies, stars that are more massive than those found in the local Universe, or the remnants of those stars (Barkana & Loeb, 2001). However, many questions surrounding these first sources remain. Whatever these first sources were, their existence changed the appearance of our Universe in a profound way. They ended the dark ages, and their photons also ionized their immediate environments, starting the Epoch of Reionisation (EoR). As more sources are created and as time progresses this ionization progresses until it results in the ionisation of the entire Universe, marking the final and largest phase change of hydrogen gas in the Universe. Studying the EoR allows us to determine what the nature was of the very first luminous objects and how they changed the nearly-uniform, early Universe into the structured one we observe today.

This experiment is however extremely challenging, e.g. due to the large dynamic range between the foreground and the cosmological hydrogen signal, the

ionosphere, the complex signal chain of the latest generation of radio interferometers, and the many analysis steps required to process large amounts of observational data. This makes the calibration and analysis of low frequency radio interferometric data difficult, and in this thesis I will study and discuss some of the challenges suffered by current calibration algorithms.

1.3 Thesis Outline

In this thesis I explore the impact of systematic errors on the calibration of radio telescopes in EoR experiments that aim to measure the 21-cm power spectrum (PS). Chapter 2 reviews current knowledge of this poorly-studied epoch. Chapter 3 discusses the role of, and challenges to experiments trying to detect radio signals from this time. Chapter 4 discusses the 21-cm PS technique, its limitations and reviews current results. Chapter 5 explores the limitations of redundancy based methods as a means to calibrate radio interferometers. Chapter 6 studies the impact of variations of antenna responses in sky-based calibration on the 21-cm PS. Chapter 7 compares systematic errors between redundant and sky based calibration, and I conclude with a summary and outlook in Chapter 8.

The Epoch of Reionisation

The emergence of the very first astrophysical sources of light dramatically changes the face of the Early Universe. Their existence ends the Dark Ages, and causes the last and largest phase change in the Universe. These sources of light emit high energy photons that heat and ionise their surroundings, creating ionised bubbles of gas in the intergalactic galactic medium (IGM). As more sources are created these bubbles grow in number, and as time progresses these bubbles grow in size. Ultimately, they start to overlap with each other, resulting in the ionisation of the entire Universe. The reionisation history of the Universe gives us insight into the nature of these very first sources, shedding light on a poorly understood chapter in the history of the Universe. This chapter discusses the observables of the Epoch of Reionisation (EoR), and how they are shaped by the astrophysics driving it.

2.1 Current constraints

We currently do not know exactly what sources drive the reionisation process. However, we have constraints on when the epoch occurred through substantial effort put into indirect measurements of the EoR. Combining constraints from Cosmic Microwave Background (CMB) measurements, high redshift quasar spec-

tra, and high redshift galaxies the EoR is constrained to start around $z \sim 11$ (Planck Collaboration et al., 2016), and to finish around $z \sim 6$ (Fan et al., 2006; Wang et al., 2019). I now briefly discuss each probe individually and how they contributed to these constraints.

2.1.1 Cosmic Microwave Background

The CMB is a versatile tool to probe the Universe. In the previous chapter I discussed the events that led to the emission of the CMB, and how it carries the imprint of the density field of the very early Universe. However, it does not only contain information about the Universe at the time radiation decoupled from matter. It also contains information about the state of matter throughout cosmic time, because CMB photons inevitably interact with matter as they pass through the Universe. The CMB interacts primarily with free electrons along the line of sight through Thomson scattering (Sunyaev & Zeldovich, 1970). A secondary effect is the Sunyaev-Z'eldovic (SZ) effect that changes the energy of CMB photons. This effect appears most prominently in galaxy clusters, where electrons in the hot intra-cluster gas at $T \sim 10^7$ K upscatter CMB photons through the thermal SZ (tSZ) effect. The effect that enables us to constrain EoR parameters is the kinematic SZ (kSZ) effect that leads to either up or down scattering depending on the direction of the gas velocity with respect to the CMB (Sunyaev & Zeldovich, 1972). The scattering of CMB photons by electrons in the IGM changes the optical depth and the spatial structure of the relic photon field. The most recent measurements of the optical depth, from the Planck mission, indicate that reionisation was nearly complete at $z = 8.8_{-1.4}^{+1.7}$ (Planck Collaboration et al., 2016).

2.1.2 Quasar Spectra

Other probes to study the state of the IGM are high-redshift quasar spectra and the Lyman- α forest. Figure 2.1.2 shows several spectra of high redshift quasars

from Fan et al. (2006). A common feature in these spectra is the sudden decrease in flux at wavelengths shorter than the rest wavelength of the Lyman- α transition from the $n = 1$ to $n = 2$ orbital in neutral hydrogen. Hence, the wavelength at which this drop occurs increases with quasar redshift. This feature, the Gunn-Peterson trough (GP; Gunn & Peterson, 1965), is caused by neutral hydrogen along the line of sight. As photons from the quasar move through the expanding Universe they are continuously redshifted, and when they are redshifted such that the photon energy matches the energy of the Lyman- α line, they can be absorbed. This causes an absorption feature in quasar spectra that is extremely sensitive to trace amounts of neutral hydrogen. The optical depth of the GP quickly saturates when the neutral fraction of hydrogen exceeds $x_{\text{HI}} > 10^{-4}$ (Gunn & Peterson, 1965), and hence the GP is only sensitive to the end of reionisation.

The near flattening of the spectra at redshifts $z > 6$ indicates there is enough neutral hydrogen present to completely absorb the redshifted photons from the quasar along that line of sight. Looking at quasars at lower redshifts, we still see “transmission spikes” bluewards (left) of the Lyman- α line, indicating that the Universe is becoming transparent to quasar spectra due to the reionisation process. The spectra indicate that nearly all hydrogen was ionised around redshift $z = 6$. However, because the reionisation process is patchy, more lines of sight are required for a representative sample of the global properties of the Universe.

2.1.3 High-redshift galaxies

The previous methods probed the state of the IGM, and therefore the effect of reionisation. It is, however, also possible to search for the sources of reionisation directly. One possible source of ionising photons are star forming galaxies, and I will discuss this later in this chapter in more detail. Searches for these high-redshift galaxies, e.g. by Bouwens et al. (2015); Oesch et al. (2015) and Livermore et al. (2017), have revealed galaxies out to redshift $z \sim 11$ (Oesch et al., 2016). The presence of these galaxies indicates the presence of ionising




Figure 2.1: The Lyman- α forest and the Gunn-Peterson trough (GP) in the 19 high-redshift quasar spectra from Fan et al. (2006). The quasars are ordered by redshift showing the GP shifting to shorter wavelengths as a function of redshift. Quasars towards lower redshift show more structure in the GP, than those at higher redshifts. Indicating a lack of neutral hydrogen at lower redshift.

sources, and provide constraints on the UV Luminosity Function (UV LF) at these high redshifts. The UV LF counts the number of galaxies per unit volume as a function of luminosity in the rest frame UV-regime. Similarly, one can measure the luminosity function of Lyman- α emitting galaxies. Both studies find that the luminosity function drops for $z > 6$ indicating the presence of neutral hydrogen that is absorbing photons in this regime. However, it is still unclear whether the change in the luminosity functions is caused by reionisation or due to the evolution of galaxies (Bouwens, 2016; Dijkstra, 2016).

2.2 21-cm Physics

The CMB, Quasars and high z -galaxies provide indirect constraints on the state of the IGM throughout cosmic time. However, the 21-cm line from neutral hydrogen allows us to study that state directly. In this section I discuss the physics that describes this hydrogen emission line and how it relates to the state of the IGM.

2.2.1 The Hyperfine Transition

The emission of a photon at a wavelength of $\lambda_{21} = 21.1061$ cm or a frequency of $\nu_{21} = 1420.4057$ MHz by atomic hydrogen was predicted towards the end of World War II by van de Hulst (1945). It was theorised that a transition in the spin of an electron orbiting a hydrogen nucleus, a proton, would release a photon with a wavelength of roughly 21 cm. Due to the quantum nature of spin, the total spin of an electron and a proton can only have two states: parallel and anti-parallel. However, the proton and electron have opposite charges. This causes the magnetic dipole moment to be anti-aligned, hence having a higher energy, when the spins are parallel. This is not too dissimilar from holding two fridge magnets with anti-aligned magnetic fields close together, this clearly will cost you some effort and energy. When the spin of the electron flips its orientation from parallel to anti-parallel, the difference in energy is carried away by a 21-cm photon (see Figure 2.2).

The probability of the spin flip occurring is staggering low. It has a transition rate (or Einstein A coefficient) of $2.9 \times 10^{-15} \text{ s}^{-1}$. However, despite this low transitional rate, we can still observe the emission of this line due the large number of hydrogen atoms in the Universe. The 21-cm line was first observed by Ewen & Purcell (1951) while measuring the galactic radio spectrum¹.

¹Their description of the detection of galactic hydrogen as one of “the most challenging problem in radio astronomy” should give us all hope for the detection of a signal from reionisation, because they have obviously succeeded.

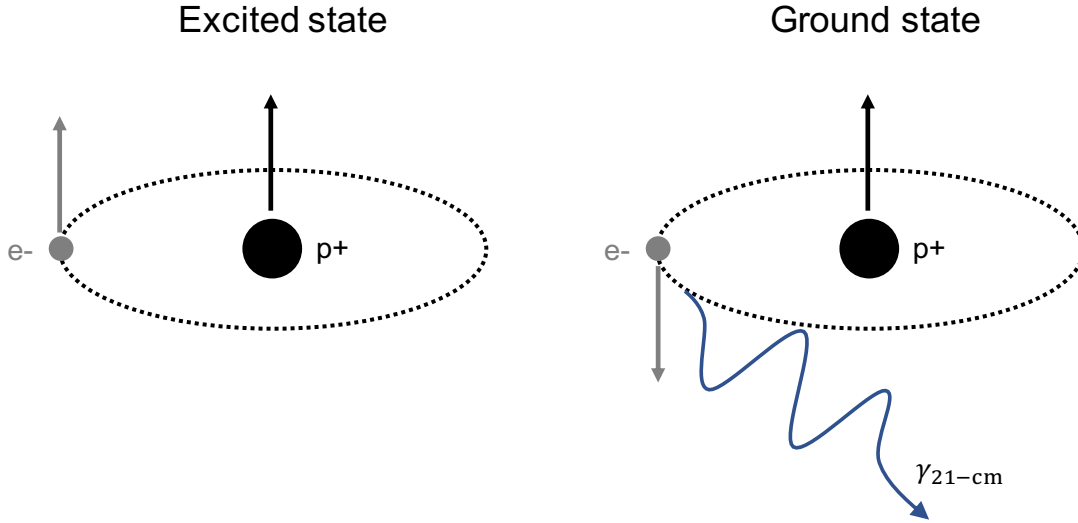


Figure 2.2: Illustration of the spin flip in neutral hydrogen. *left*: The excited state of neutral hydrogen when the spins of the electron and the proton are aligned (and therefore the magnetic moments are anti-aligned). *right*: The ground state of neutral hydrogen when the spins are anti-aligned (and therefore the magnetic moments are aligned). The transition from the excited to the ground state is accompanied by the emission of a photon $\gamma_{21\text{-cm}}$ with a wavelength of 21 cm.

2.2.2 The Spin Temperature

The strength of the 21-cm signal depends on the relative population of atomic hydrogen in the parallel spin state. We can describe the population of hydrogen atoms in the ground state n_0 and the excited state n_1 using the Boltzmann Law

$$\left(\frac{n_1}{n_0}\right) = \left(\frac{g_1}{g_0}\right) \exp\left(-\frac{T_{21}}{T_S}\right), \quad (2.1)$$

where the ratio of statistical weights $g_1/g_0 = 3$, $T_{21} = 68\text{ mK}$ is the temperature corresponding to the energy of the 21-cm transition, and T_S is the spin temperature. The spin temperature does not always equal the kinetic temperature of the hydrogen gas; instead it is a quantity describing the relative atomic population in the ground and excited state, and therefore defines the emission or absorption we observe.

The population of atoms in the parallel state, and therefore the spin tempera-

ture, is controlled by five parameters; the CMB brightness temperature T_γ that describes the intensity of a black body radiation field at that temperature, x_c and x_α the collisional, and the UV-scattering coupling coefficients, T_K the kinetic gas temperature and T_C is the colour temperature of the UV-radiation field. This is summarized in Equation (2.2) (Field, 1958)

$$T_S^{-1} = \frac{T_\gamma^{-1} + x_c T_K^{-1} + x_\alpha T_C^{-1}}{1 + x_c + x_\alpha}. \quad (2.2)$$

The 21-cm line can thus be excited by interactions with CMB photons, collisions between hydrogen and other particles, and finally the scattering of UV photons. The effectiveness of each process is determined by the coupling coefficients.

The collisional coefficient x_c depends on the density of the IGM and particles involved. Initially, this is dominated by hydrogen, but as soon as the reionisation process starts electrons also start playing a role. Below $z < 70$ the density of the IGM is too low for collisions and therefore we would naively expect the CMB temperature T_γ to dominate the spin temperature. Instead, it turns out that the scattering of UV photons from the first luminous sources provides a mechanism to break the coupling to the CMB field.

The scattering of UV-photons couples the spin temperature to the colour temperature of the UV-photon field T_C . Wouthuysen (1952) suggested how the scattering of UV-photons populates the excited spin states of neutral hydrogen indirectly. This mechanism, now called the Wouthuysen-Field effect, plays an important role in the breaking of the coupling between the CMB and the spin temperature. The mechanism relies on the spin excitation of hydrogen using the $n = 2$ level as an intermediate state (see Figure 2.3). Lyman- α photons excite hydrogen from the ground state into that energy level while changing the spin orientation of the electron, because the electric dipole selection rules allow for that. We now end up with a mix of spins in the upper state, which then can decay back into the lowest energy state while preserving that spin state. I will not go into exact details of the quantum mechanics of this spin mixing process,




Figure 2.3: Illustration of the energy levels involved in the Wouthuysen-Field effect. It shows how Lyman- α excitation to the $n = 2$ level acts as an intermediate before the electron cascades into the excited hyperfine spin state. The solid lines indicate transitions relevant to the Wouthuysen-Field effects, the dashed lines indicate other allowed transitions that do not play a role. Reproduced from Pritchard & Furlanetto (2006)

but for the interested reader they can be found in Field (1958). The colour temperature of the UV-photon field is formally defined as

$$\frac{P_{01}}{P_{10}} \equiv \frac{g_1}{g_0} e^{-T_{21}/T_C} \sim 3 \left(1 - \frac{T_{21}}{T_c} \right), \quad (2.3)$$

where P_{01} and P_{10} are the UV-scattering excitation and de-excitation rates, respectively. Hydrogen is optically thick to Lyman- α emission, i.e. a single photon undergoes a large number of scattering events before escaping. This and the transfer of energy from the UV-photon field to the gas through atomic recoils causes the kinetic temperature and the colour temperature to equalise $T_C \rightarrow T_K$.

2.3 The First Luminous Sources

I have discussed what influences the spin temperature, e.g. the kinetic gas temperature T_K , the CMB temperature T_γ , and the UV colour temperature T_C . But these temperatures are simply proxies for the astrophysics we are trying to uncover. In this section I will discuss the impact of different astrophysical sources of radiation.

2.3.1 Ultraviolet Sources

Earlier I have discussed that a Lyman- α field couples the 21-cm spin temperature T_S to the kinetic gas temperature T_K . Star forming galaxies are believed to be sources of Lyman- α radiation (Milosavljević & Safranek-Shrader, 2016). However, it is not just the local galaxies at a certain redshift that determine the local Lyman- α field. Photons from local sources quickly redshift out of the absorption resonance. Hence, UV-sources in the background also matter. Photons from these background sources redshift down to the local Lyman- α resonance and as such contribute to the UV-background at lower redshifts (Furlanetto, 2019). Earlier, I discussed the discovery of galaxies out to redshift $z \sim 11$. However, the *UV LF* at these high redshifts indicates that bright galaxies are most likely not the

dominant cause for reionisation. Their low numbers, either indicate that the bulk of reionisation is driven by other sources of UV-photons, or come from galaxies below our current detection limits (Bouwens, 2016).

2.3.2 X-ray Sources

The kinetic gas temperature is set by ultraviolet photons, shock heating due to structure formation, and most notably due to X-rays (Ciardi & Madau, 2003; Pritchard & Furlanetto, 2007). X-rays are very efficient in penetrating the IGM, and cause multiple ionisations along their path. The liberated electrons interact with the gas through collisions, transferring the energy from the initial X-ray photon to the gas. The total number of X-ray photons and the spectra of the X-rays depend strongly on the astrophysical source.

The lack of metal enriched gas at early times suggests that the first stars had to be significantly more massive than stars in our local Universe to overcome the gas pressure (Milosavljević & Safranek-Shrader, 2016). Their high mass M and therefore short life times (that scale with mass as $\propto M^{-3}$) would rapidly lead to the formation of black holes. Particularly, the formation of black hole binary systems in which a black hole accretes from an orbiting star, would lead to an X-ray Binary (XRB). These are efficient sources of X-ray emission (Mirabel et al., 2011). Isolated stellar mass black holes do not accrete efficiently from a diffuse ISM and hence do not contribute much X-ray flux. Other sources could be supernova remnants that shock heat their surroundings (Oh, 2001), or the first active galactic nuclei (AGN). However, it is hard to form supermassive black holes early in the Universe, unless they were formed through primordial black hole seeds (Madau et al., 2004).

2.3.3 The Fiducial Model

Reionisation is most likely driven by a mix of sources. The current hierarchical model assumes reionisation was initially driven by UV-sources and at later times

reionisation is driven by X-ray sources (Mesinger, 2016). However, there are still many things unclear. The efficiency of each ionising source is still under debate, and the literature is continuously growing with more state of the art simulations, and observational constraints on the astrophysics that drive the EoR. The ionising efficiency of stars at early times depends on the star formation efficiency in the early Universe, the escape fraction of UV-photons from galaxies, and the number of ionising photons produced by the first stars. For XRBs it is still unclear what the formation rate of black holes was, and how many of them would occur in binaries. For primordial black holes it is unclear whether they existed at all. Combined with uncertainties pertaining to the properties of the IGM, e.g. clumping factors, turbulence etc., these parameters impact on the morphology of reionisation (Dayal & Ferrara, 2018), and determine whether reionisation follows an inside-out scenario (high-density regions ionise first), or an outside-in scenario (where low-density regions ionise first) (Lidz, 2016). Therefore, the spatial structure of the reionisation process provides strong constraints on these parameters.

2.3.4 The Differential Brightness Temperature

Now that we understand what drives the spin temperature, we can discuss the actual emission our telescopes receive. To do so I start with the radiative transfer equation of an emission line at frequency ν

$$\frac{dI_\nu}{ds} = \frac{\phi(\nu)h\nu}{4\pi} [n_1 A_{10} - (n_0 B_{01} - n_1 B_{10}) I_\nu], \quad (2.4)$$

where I_ν is the specific intensity at frequency ν , ds is the path length along a ray, $\phi(\nu)$ is the line profile function whose integral is normalised to 1, n_i is the number density of atoms mentioned earlier, A_{ij} and B_{ij} are the Einstein coefficients for the transition from initial state i and final state j . The first term that includes A_{10} describes how much radiation is added to a radiation field

I_ν through spontaneous emission as it passes through a medium. The second term including B_{01} and B_{10} that describe absorption and stimulated emission, respectively, describe how much radiation is absorbed. The coefficients in the first term are often grouped to be the source function S_ν and the second term, and the coefficients in the second term are written as the absorption coefficients κ_ν .

$$\begin{aligned} S_\nu &\equiv \frac{\phi(\nu)c}{4\pi} \frac{n_1 A_{10}}{n_0 B_{01} - n_1 B_{10}} \\ \kappa_\nu &\equiv \frac{h\nu}{c} (n_0 B_{01} - n_1 B_{10}) \end{aligned} \quad (2.5)$$

Replacing the coefficients in Equation (2.4) with these definitions and then solving the differential equation yields

$$I_\nu = S_\nu (1 - e^{-\tau(\nu)}) + I_0(\nu) e^{-\tau(\nu)}. \quad (2.6)$$

Here, I conveniently defined the optical depth $\tau(\nu) = \int \kappa_\nu ds$ as the integral of the absorption coefficient over the line of sight, i.e a measure for the transparency of a medium.

Observing in the radio regime implies we can use the Rayleigh-Jeans Limit that provides a linear relation between the brightness temperature of a black body and its intensity $T_b \sim I_\nu c^2 / 2k_B \nu^2$. This enables us to write the radiative transfer equation in terms of temperatures of the radiation fields involved, i.e. the background CMB field $I_0(\nu)$ defined by T_γ and the 21-cm emission S_ν defined by T_S .

$$T'_b(\nu) = T_S(1 - e^{-\tau\nu}) + T_\gamma(\nu)e^{-\tau\nu} \quad (2.7)$$

The optical depth for the 21-cm line in an expanding Universe filled with hydrogen, and a gas velocity gradient along the line of sight $dv_{\parallel}/dr_{\parallel}$ can be

expressed as follows (Field, 1959)

$$\tau_{21} \approx 0.0092(1 + \delta)(1 + z)^{3/2} \frac{x_{\text{HI}}}{T_S} \left[\frac{H(z)/(1 + z)}{dv_{\parallel}/dr_{\parallel}} \right], \quad (2.8)$$

where δ is the matter overdensity, z is the redshift we are observing, x_{HI} is the fraction of neutral hydrogen and $H(z)$ is the Hubble constant at that redshift. Noting that the observed brightness temperature of the CMB lowers due the expansion of the Universe, and the optical depth at the relevant redshift is typically small, $\tau_{\nu} \ll 1$, I rewrite Equation (2.7) into

$$T_b(\nu) = \frac{T_S - T_{\gamma}(z)}{1 + z} \tau_{\nu}. \quad (2.9)$$

For line emission in radio astronomy, we measure the differential brightness $\delta T_b \equiv T_b - T_{\gamma}$. In other words, we only see 21-cm features when they differ from the background. Usually the CMB acts as a radio background, however, the results from the Experiment to Detect the Global EoR Step (EDGES; Bowman et al., 2008) could indicate the existence of an additional source of diffuse background radiation² (Bowman et al., 2018). Using this and the expression for the optical depth we get the expression for the measured differential brightness temperature

$$\delta T_b = 28 \text{ mK} (1 + \delta) x_{\text{HI}} \left(1 - \frac{T_{\text{CMB}}}{T_{\text{spin}}} \right) \times \left(\frac{\Omega_b h^2}{0.0223} \right) \sqrt{\left(\frac{1 + z}{10} \right) \left(\frac{0.24}{\Omega_m} \right)} \left[\frac{H(z)/(1 + z)}{dv_{\parallel}/dr_{\parallel}} \right]. \quad (2.10)$$

Here, I introduced the cosmic baryon density Ω_b , and h the Hubble parameter. The differential brightness temperature depends on the overdensity δ , the local neutral fraction of hydrogen x_{HI} and the spin temperature. We can infer from the ratio between the CMB temperature and the spin temperature that there is

²At the time of writing this result is still highly debated, and can also possibly be attributed exotic dark matter physics, or to the frequency response of the EDGES instrument.

a saturation limit to observe the 21-cm line in emission, but there is no limit in absorption. This will have an implication for the observable global signal, as discussed in the next section.

2.4 21-cm Cosmology

As discussed earlier, the 21-cm line from neutral hydrogen allows us to probe the astrophysics of the early Universe and constrain the spin temperature, but we can also study more global cosmological parameters (see Equation 2.10). There are three distinct methods that use the 21-cm line as a probe for the early Universe. In this section I discuss these different methods and what information we can extract from them.

2.4.1 The Evolution of the 21-cm signal

Experiments trying to measure the global signal, e.g. EDGES (Bowman et al., 2008), the Shaped Antenna measurement of the background Radio Spectrum (SARAS; Patra et al., 2013), the Large Aperture Experiment to Detect the Dark Age (LEDA; Greenhill & Bernardi, 2012), and the Dark Ages Radio Explorer (DARE; Burns et al., 2012), focus on the sky averaged 21-cm signal as a function of redshift or cosmic time. The advantage of such a global detection is the higher signal to noise ratio, since we are averaging the signal over the entire sky. But a global signal yields only global observables of the EoR. Nevertheless, it allows tighter constraints on the start and duration of the EoR, and the global properties of the first luminous sources.

Figure 2.4 shows a suite of global signal models. Despite the large variety, there is a general trend. Initially, we see the spin temperature slowly moving towards T_γ , or $T_b = 0$. The continuing expansion of the Universe then decreases the matter density rendering the collisional excitation of hydrogen inefficient, and so the spin temperature T_S decouples from the gas temperature T_K . Once the




Figure 2.4: A suite of global signal models for the Epoch of Reionisation. overall they follow a similar trend; the signal goes into absorption, then into emission before it decays to zero due to reionisation. The exact amplitude of the absorption and emission, and the shape of the global signal as a function of frequency depends heavily on the astrophysics driving the heating and reionisation process. Reproduced from (Cohen et al., 2017)

UV sources turn on around $z \sim 25$, the resulting UV-field couples the spin temperature back to gas temperature. The overall gas temperature of the Universe keeps decreasing as the gas cools adiabatically due to expansion. This continues until the first X-ray sources switch on, from this moment forward the gas starts to heat. The heating continues, and we start seeing 21-cm signals in emission until it saturates the differential brightness temperature δT_B . Finally the signal decreases, due to the reionisation that stops hydrogen from emitting a 21-cm signal. A first detection of this global signal has been claimed by EDGES (Bowman et al., 2018), where they have detected the signal in absorption during Cosmic Dawn ($z \sim 15 - 30$). However, the amplitude and shape of the absorption trough have sparked discussion whether this is due to new exotic dark matter physics (Barkana, 2018), an additional source of diffuse background radiation, or due to systematics in the calibration process.

2.4.2 The 21-cm Power Spectrum

CMB experiments demonstrated that power spectra (PS) are useful tools in astronomy. Instead of looking at the global amplitude of the signal, the spatial power spectrum measures fluctuations in the signal amplitude as a function of spatial scale at different redshifts. We know that the signal amplitude fluctuates across the sky, because the heating and ionisation is a patchy process and morphology of reionisation is heavily dependent on astrophysics driving the EoR.

Unlike the global signal, the 21-cm PS gives us insight into the spatial structure of the reionisation process while retaining some of the sensitivity that global signal experiments possess (Furlanetto et al., 2006). I will discuss the details of power spectrum estimation in more detail in Chapter 4.

There are several instruments in the world currently trying to detect the power spectrum from reionisation, and in this thesis I will focus on the efforts using the Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018). We discuss the MWA and other instruments in the next chapter.

2.4.3 Alternative Techniques

We can extend the power spectrum framework to higher orders. The bispectrum (Bharadwaj & Pandey, 2005), i.e. Fourier Transform of the 3-point correlation function, and the trispectrum (Cooray et al., 2008), i.e. the Fourier transform of the 4-point correlation function. Whereas the PS is only sensitive to Gaussianity in the signal as it measures the variance, the bispectrum and trispectrum are sensitive to non-gaussianities that enable us to probe the morphology of reionisation process.

The global signal is an example of a one-point statistic, i.e. the mean. We can also study the global variance, skewness, and kurtosis (Patil et al., 2014; Kittiwisit et al., 2017) as a means to distinguish between different reionisation histories.

The ultimate goal is direct 3-dimensional imaging of the IGM, i.e. tomography (Madau et al., 1997). Observing the differential brightness temperature as a function of position in the Universe, would allow us to directly observe the ionisation bubbles, their structure and clustering as a function of cosmic time. We currently lack telescopes that are sensitive enough to directly image neutral hydrogen in the early Universe. However, the advent of the Square Kilometre Array promises to change that (Koopmans et al., 2015).

2.5 Challenges of the EoR Experiment

The main challenging aspect of the signal from neutral hydrogen is the brightness of the signal. Equation (2.10) is written in terms of millikelvin. Radio foregrounds, however, dominate our measurements because their mean signal is 4-5 orders of magnitude stronger than the EoR. In the variance that reduces to 3-4 orders of magnitudes (Jelic et al., 2008). The foregrounds are a mix of different signals; synchrotron emission from our Milky Way makes up 70% of the radio sky, the sum of all extragalactic emission adds up to 27%, and finally thermal

emission from our galaxy contributes to about 1%. The remainder comes from the CMB (Shaver et al., 1999). The exact mix of foregrounds signals will depend on which parts of the sky are targeted for observations.

The EoR is observed in the MHz regime of the radio spectrum. At these low frequencies turbulence in the ionosphere lowers the quality of our observations. Similar to the effect of the atmosphere on the twinkling of stars in the optical regime, the ionosphere refracts radio waves. This refraction becomes stronger at lower frequencies and affects current low frequency radio observatories, because they have resolutions that are sensitive to these ionospheric distortions. The behaviour of the ionosphere and its impact on low frequency observations is currently being studied (Mevius et al., 2016; Jordan et al., 2017; Albert et al., 2020a) and there is a large investment in techniques to compensate for distortions due to the ionosphere (Mitchell et al., 2008; Yatawatta et al., 2009; Kazemi et al., 2011; Intema et al., 2009; Sullivan et al., 2012; van Weeren et al., 2016; Albert et al., 2020b).

A final major challenge are the instruments themselves. We need many hours of observing time to adequately average out the large radiometric noise, and we need to understand the complex signal chain in these observatories to mitigate any unwanted imprints each instrument leaves on the signal. This thesis aims to bring a deeper understanding to the latter problem. However, I first need to discuss the techniques used by radio telescopes to observe the 21-cm PS.

Radio Interferometry

Radio telescopes reveal a Universe far beyond the optical regime probed by human eyes. Similar to telescopes at other wavelengths, the aperture diameter D of a single radio telescope is limited by the mechanical structure we can build to support it. This puts a limit on the resolution of a single element $\sim \lambda/D$, worsened by the long wavelengths λ of radio signals. The need for higher resolution led to the development of radio interferometry; combining the signals from multiple antennas spaced over larger distances. This allows us to synthesise a telescope that is larger than we can mechanically construct. In this chapter I will give a brief overview of the latest generation of radio telescopes, review the fundamentals of radio interferometry, and conclude with a discussion on existing calibration frameworks upon which this thesis aims to expand.

3.1 The Epoch of Radio Astronomy

The radio view of the Universe started with the accidental discovery of a radio source towards the direction of the Sun by Jansky (1932). This was later refined to be in the direction of the centre of the Milky Way (Jansky, 1935). Nowadays, we know this object as Sagittarius A*; the supermassive black hole hosted by our Milky Way. Although radio astronomy started at what we currently call low

frequencies¹, the discovery of the hydrogen 21-cm line pushed the field towards higher gigahertz frequencies. This also enabled higher resolving power for single dishes, because of the smaller wavelengths involved, and evades refraction by the turbulent ionosphere that increases with wavelength as $\sim \lambda^2$. However, the prospect of detecting the redshifted 21-cm signals from the Epoch of Reionisation (EoR) reinvigorated our interests in megahertz frequencies, and in part motivated the development of low frequency interferometers around the world that unveiled unprecedented views of the low frequency Universe.

3.2 Third Generation Radio Telescopes

Advances in signal processing and computing enabled the construction of the low frequency radio instruments of the 21st century. These advances enabled us to form antenna stations by combining the signals from multiple antennas (see section 3.4.2) creating apertures with larger collecting areas A_{eff} , without resorting to the construction of large mechanical structures. The development of supercomputing facilities enabled the ability to process signals from a much larger number of stations N_{ant} , resulting in a larger total collecting areas A_{tot} , and therefore arrays with higher resolution and sensitivity (Baars et al., 2009). Setting course for the Square Kilometre Array (SKA) many collaborations around the world have developed pathfinder or precursor telescopes to demonstrate the potential success of the SKA. Here, we give a short overview of different low frequency telescopes performing EoR experiments, with a technical summary of each telescope in Table 3.1.

The **Long Wavelength Array**² (LWA; Ellingson et al., 2009) has 3 stations across the U.S.A. Each station consists of 256 dipole antennas randomly distributed within a radius of ~ 100 metres and an outrigger antenna at 300-500

¹Jansky (1932) actually refers to this frequency range as high frequency while discussing the origin of the "cosmic hiss".

²<http://lwa.unm.edu>

metres from the core. Although an interferometer, the array in the Owens Valley (OVRO-LWA) hosts a global signal experiment: the Large Aperture Experiment to Detect the Dark Ages (LEDA; Price et al., 2018), and uses the interferometric mode for calibration. The OVRO-LWA was also used to create an all-sky map of the Northern sky at frequencies between 36-73 MHz using *m*-mode analysis (Eastwood et al., 2018). We briefly discuss *m*-mode analysis in Chapter 4. This map is particularly useful for experiments targeting the Cosmic Dawn of the Universe.

The former Donald C. Backer **Precision Array for Probing the Epoch of Reionization**³ (PAPER; Parsons et al., 2010) was a testing platform for EoR power spectrum (PS) experiments. It was a reconfigurable array with antennas made out of a single dipole on an elevated ground screen. Through its lifetime it has seen various lay-outs ranging from an East-West array to a circular lay-out. PAPER operated as a drift-scan telescope whose pointing of the antennas is fixed to zenith while the sky drifts through the field of view (FoV) as the Earth rotates. In its final design PAPER consisted of 64 elements on a rectangular grid. Initially deployed in Green Bank, U.S.A, it was later moved to the South African SKA site in the Karoo Desert.

The **Giant Metrewave Radio Telescope**⁴ (GMRT; Swarup et al., 1991) follows a more traditional telescope design of 30 dishes with a 45 metre diameter. The dishes are spread out in a Y-shape stretching out to 25 km near Pune, India. Although one of its primary goals is to detect a hydrogen signal at high redshifts, its coverage of a wide range of frequencies from 50 MHz - 1420 MHz services a wide variety of science. This ranges from studies of neutral hydrogen throughout cosmic time to the detection of fast transients.

The **LOW Frequency ARray**⁵ (LOFAR; van Haarlem et al., 2013) is an aperture array system (see section 3.4.2) with two types of antennas. LOFAR has High Band Antennas (HBA), relevant for EoR science, that operate between

³<http://eor.berkeley.edu/>

⁴<http://www.gmrt.ncra.tifr.res.in/>

⁵<http://www.lofar.org/>

120-190 MHz and Low Band Antennas (LBA) that operate between 30-90 MHz. The antennas are organised in so-called stations, where HBA stations are built out of enclosed “tiles” each containing 4×4 dipoles. The sizes of the HBA stations varies throughout the array. The bulk of LOFAR stations are located in The Netherlands. However, LOFAR also has stations across Europe that form the international baselines with lengths up to a few thousands of kilometres. For EoR science only the Dutch stations are used.

The **Murchison Widefield Array**⁶ (MWA; Tingay et al., 2013) is also an aperture array system with 256 tiles located in the Western Australian desert at the Murchison Radio Observatory (MRO), the soon to be home of the low frequency part of the SKA. Similar to LOFAR, each tile is made out of 16 dipoles on a 4×4 grid, but in the MWA they are placed on an open metal mesh and operate as a single unit. In its current deployment, Phase II, the MWA operates in two distinct modes that each use 128 elements (Wayth et al., 2018). In the compact mode, 72 tiles are organised in two highly redundant hexagonal configurations, the “hexes”, and the remaining 52 are distributed pseudo-randomly within a radius of 3 kilometres. In the extended mode, the tiles in the hexes are replaced with 72 antennas that spread out to 6 km providing enhanced imaging capabilities. For the EoR experiment, we use the latter mode to create high fidelity sky models, whereas the compact mode is used to actually detect the EoR (Beardsley et al., 2019).

The **Hydrogen Epoch of Reionization Array**⁷ (HERA; DeBoer et al., 2017) is currently under construction in the Karoo Desert in South Africa. Inspired by the final design of PAPER, it is also a redundant telescope primarily dedicated to perform EoR science. Unlike PAPER, it will be built of zenith-pointed parabolic dishes with feeds suspended above the dish mesh. Similar, to PAPER it will also operate as a drift-scan instrument. In its final form it will have 350 dishes in a closely packed hexagonal formation.

⁶<http://mwatelescope.org>

⁷<https://reionization.org/>

Telescope	N_{ant}	$A_{\text{eff}} [\text{m}^2]$	$\Delta x_{\text{min}} [\text{m}]$	$\Delta x_{\text{max}} [\text{m}]$	$A_{\text{tot}} [\text{m}^2]$
LOFAR	52	$\sim (5.1 - 20) \times 10^2$	68	1.5×10^6	5.5×10^4
LWA	257	7.0	5	1.5×10^3	1.8×10^3
GMRT	30	10^3	100	3.0×10^4	3.0×10^4
MWA	128	16	8	6.0×10^3	3.6×10^3
PAPER	64	18	4	2.1×10^2	1.2×10^3
HERA	350	93	14.6	8.8×10^2	5.4×10^4
SKA-LOW1	512	$\sim 10^3$	35	6.5×10^4	5.0×10^5

Table 3.1: An overview of the technical specifications of radio interferometers performing EoR Science. N_{ant} describes the number of receiver units whose signals are correlated, this is not equal to the number of individual antennas as some telescopes group multiple antennas into a single unit. A_{eff} is the effective area of an element at 150 MHz, Δx_{min} is the shortest baseline length, Δx_{max} is the maximum baseline length, and A_{tot} gives the total collecting area, i.e $N_{\text{ant}} \times A_{\text{eff}}$. This table demonstrates the diversity in design of the various experiments, each with its strengths and weaknesses. Note: the effective area A_{eff} is given at 150 MHz, except for the LWA where we quote the value at 73 MHz.

The **Square Kilometre Array**⁸ (Dewdney et al., 2015) will have its low frequency component SKA-LOW in the Murchison Desert in Western Australia near the MWA site. One of the key science projects of the SKA is to observe the EoR, through both the 21-cm PS and 21-cm tomography (Koopmans et al., 2015). However, it also aims to service a wide range of science goals over a wide range of frequencies through a single receiver element⁹ covering 50-300 MHz. Similar to the MWA and LOFAR, it is an aperture array telescope with of 512 stations in its first deployment, SKA-LOW1. However, unlike LOFAR and the MWA each station will have 256 uniquely random distributed antennas in an attempt to suppress a strong response of the array away from the target pointing (Mort et al., 2017).

⁸<https://www.skatelescope.org/>

⁹At the time of writing the exact design of the dipole is contentious and hence we omit any detail.

3.3 Interferometry Fundamentals

Regardless of their design, location, and primary science goals, the underlying principle is the same for all these arrays. In a radio interferometer with N_{ant} antennas we compute the cross-correlation between the signals measured by a pair of antennas, i.e. we have $N_{\text{ant}}(N_{\text{ant}} - 1)/2$ unique measurements. Both antennas in a pair, or baseline, measure an impeding electric field $\mathbf{E}_\nu(\mathbf{r})$ caused by our cosmic radio sources of interest. Following Clark (1999), the time averaged cross-correlation V_ν between the two electric fields, or visibility, at frequency ν measured by antenna p and antenna q at a location \mathbf{r}_p and \mathbf{r}_q is given by

$$V_\nu(\mathbf{r}_p, \mathbf{r}_q) = \langle \mathbf{E}_\nu(\mathbf{r}_q) \mathbf{E}_\nu^*(\mathbf{r}_p) \rangle. \quad (3.1)$$

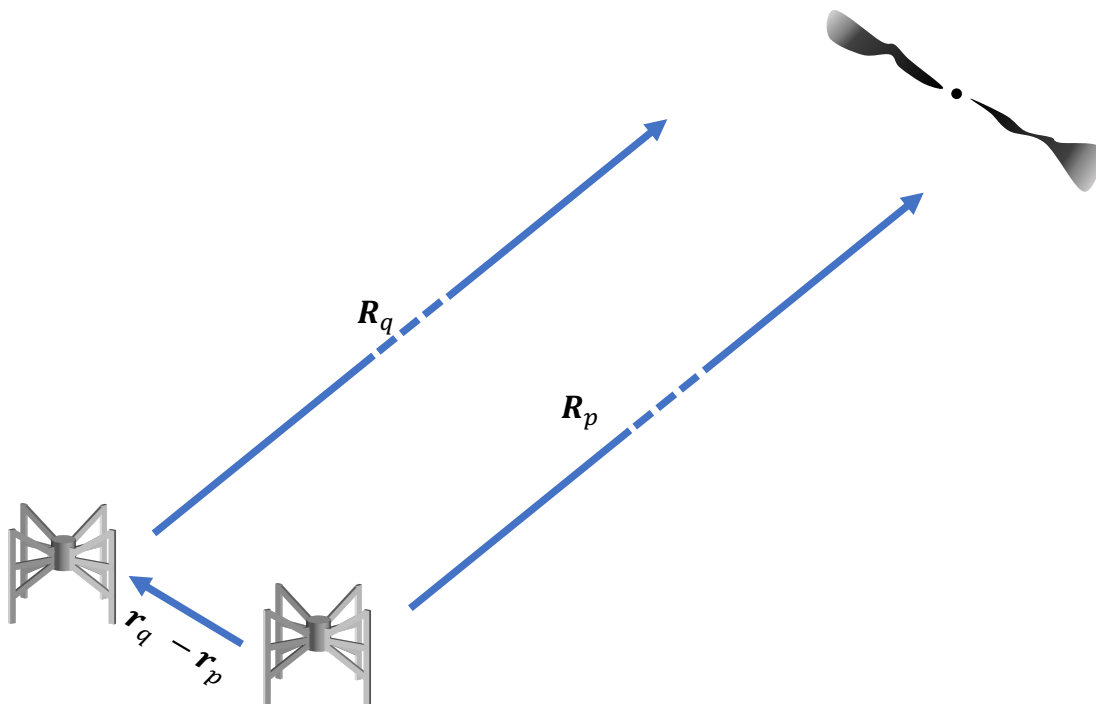


Figure 3.1: Schematic overview of the geometry involved in radio astronomy. Note: vectors representing r_p are omitted, because we assume a sensible choice will be made for the origin of this coordinate system, i.e. near the array.

The asterisk $*$ indicates complex conjugation, and the angle brackets $\langle \rangle$ indicate averaging over time. Astronomers are more interested in intensities I_ν

rather than electric field strengths. We derive the relation between the measured visibilities $V_\nu(\mathbf{r}_p, \mathbf{r}_q)$ as a function of frequency ν and the actual intensity of the sky by explicitly writing the measured electric fields using phasor notation

$$V_\nu(\mathbf{r}_p, \mathbf{r}_q) = \left\langle \int_S \int_S \mathcal{E}_\nu(\mathbf{R}_p) \mathcal{E}_\nu^*(\mathbf{R}_q) \frac{e^{2\pi i \nu |\mathbf{R}_p - \mathbf{r}_p|/c}}{|\mathbf{R}_p - \mathbf{r}_p|} \frac{e^{-2\pi i \nu |\mathbf{R}_q - \mathbf{r}_q|/c}}{|\mathbf{R}_q - \mathbf{r}_q|} dS_p dS_q \right\rangle. \quad (3.2)$$

Here, we defined the phasor amplitude $\mathcal{E}(\mathbf{R})$, with \mathbf{R}_p the separation vector between antenna p and the source of interest, and dS_p is an surface infinitesimal element on the celestial sphere with radius \mathbf{R}_p . Assuming that radio signals in the far-field are not spatially coherent, i.e. the time average of our electric fields is only non-zero when the signals originate from the same point $\mathbf{R} = \mathbf{R}_p = \mathbf{R}_q$, we can replace the double surface integral with a single integral over the full celestial sphere. We then bring the expectation into the integral

$$V_\nu(\mathbf{r}_p, \mathbf{r}_q) = \int_S \langle |\mathcal{E}_\nu(\mathbf{R})|^2 \rangle \frac{e^{2\pi i \nu |\mathbf{R} - \mathbf{r}_p|/c}}{|\mathbf{R} - \mathbf{r}_p|} \frac{e^{-2\pi i \nu |\mathbf{R} - \mathbf{r}_q|/c}}{|\mathbf{R} - \mathbf{r}_q|} dS, \quad (3.3)$$

noting that the squared amplitude of the electric field is equal to the intensity $|\mathcal{E}_\nu(\mathbf{R})|^2 = I_\nu(\mathbf{R})$. If we assume the radio source lies on a celestial sphere with radius R , we can define the components of \mathbf{R} : $R_x = R \cos(\theta_x)$, $R_y = R \cos(\theta_y)$, $R_z = R \cos(\theta_z)$. We then define the ‘‘direction cosines’’

$$\begin{aligned} l &= \cos(\theta_x), \\ m &= \cos(\theta_y), \\ n &= \cos(\theta_z) = \sqrt{1 - l^2 - m^2}, \end{aligned} \quad (3.4)$$

where the angles are measured between the plane of the antennas and \mathbf{R} that points in the direction on which our observation is centred. We refer to this as the pointing centre or *phase centre*¹⁰. Combining this with the assumption that

¹⁰These need not be the same. We can point the antennas in one direction, while arbitrarily adding a phase delay during correlation such that we are phase tracking a different point on the sky.

the distance to our sources is much larger than the separation between antennas $|\mathbf{r}_p/\mathbf{R}| \ll 1$ enables us to rewrite $|\mathbf{R} - \mathbf{r}|$

$$\begin{aligned}
|\mathbf{R} - \mathbf{r}_p| &= \sqrt{(Rl - x_p)^2 + (Rm - y_p)^2 + (Rn - z_p)^2} \\
&= R\sqrt{(l - x_p/R)^2 + (m - y_p/R)^2 + (n - z_p/R)^2} \\
&= R\sqrt{l^2 + m^2 + n^2 - 2/R(lx_p + my_p + nz_p)}
\end{aligned} \tag{3.5}$$

Note that we ignored higher order terms of $|\mathbf{r}_p/\mathbf{R}|$. The direction cosines are elements of the sky direction unit vector $\mathbf{s} = (l, m, n)$, and because the separation between antennas is much smaller than the distance to the source we can use a Taylor expansion to get

$$|\mathbf{R} - \mathbf{r}_p| \sim R - (lx_p + my_p + nz_p). \tag{3.6}$$

We further simplify this expression because astrophysical sources of radio signals are far away. Hence, we rewrite a surface element $dS = R^2 d\Omega$. Combining this with Equation (3.6) and our definition of the intensity we arrive at

$$V_\nu(\mathbf{r}_p, \mathbf{r}_q) = \int I_\nu(\mathbf{s}) e^{-2\pi i \nu \mathbf{s} \cdot (\mathbf{r}_p - \mathbf{r}_q)/c} d\Omega. \tag{3.7}$$

Where $I_\nu(\mathbf{s})$ is the brightness of a patch of sky in the direction of unit vector \mathbf{s} with angular size $d\Omega$. Equation (3.7) relates the visibility V_ν as a function of antenna separation $(\mathbf{r}_p - \mathbf{r}_q)$, and this special form of the *spatial coherence function* is also known as the van Cittert-Zernike Theorem (van Cittert, 1934; Zernike, 1938). Having derived this relation we can discuss how we recover the astronomical quantity of interest; the intensity of the sky I_ν .

3.3.1 The Measurement Equation

Equation (3.7) looks remarkably similar to a Fourier transform, and because these transforms are invertible we can use an inverse Fourier transform to recover the brightness of the sky. We can approximate this by a 2D Fourier transform

in two special cases. Before we discuss those cases we first define the baseline vector $\mathbf{b} = (u, v, w)$ as $\mathbf{b} = (\mathbf{r}_p - \mathbf{r}_q)/\lambda$ and explicitly write the sky vector $\mathbf{s} = \mathbf{s}_0 + \mathbf{s}'$. Here, \mathbf{s}_0 is the unit vector that points towards the phase centre, and $\mathbf{s}' = (l, m, n = \sqrt{1 - l^2 - m^2})$ in Equation (3.7). As mentioned earlier the angles in the direction cosines are measured with respect to the phase center, implying we need to multiply Equation (3.7) with an additional phase factor, $e^{2\pi i \mathbf{s}_0 \cdot \mathbf{b}}$. More importantly, this direction vector defines the orientation of the Cartesian coordinate system that relates sky coordinates (l, m, n) to their Fourier duals (u, v, w) . The directions spanned by $\{u, v\}$ and $\{l, m\}$ are perpendicular to \mathbf{s}_0 . Hence, $\mathbf{s}_0 = (0, 0, 1)^T$, and it only picks up the projected w -component towards phase centre (Thompson et al., 2017).

$$V_\nu(u, v, w) = \int I_\nu(l, m) e^{-2\pi i [ul + vm + w(\sqrt{1 - l^2 - m^2} - 1)]} \frac{dldm}{\sqrt{1 - l^2 - m^2}}. \quad (3.8)$$

It is the w -term in this equation that prevents us from using a 2D Fourier Transform to relate the sky intensity I_ν to the measured visibilities V_ν . However, if we assume that all our array elements lie in a plane, i.e. we assume co-planarity, w becomes zero and Equation (3.8) becomes

$$V_\nu(u, v) = \int \frac{I_\nu(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-2\pi i (ul + vm)} dldm. \quad (3.9)$$

Note: we have ignored the phase tracking term. Alternatively, we can assume we have a small FoV, and restrict ourselves to small values of l and m , i.e. $\sqrt{1 - l^2 - m^2} \approx 1$ resulting in

$$V_\nu(u, v) = \int I_\nu(l, m) e^{-2\pi i (ul + vm)} dldm. \quad (3.10)$$

In both limits we recover the sky intensity by performing a 2D-Fourier transform on the set of visibilities as measured by our interferometer. In general, for modern low frequency radio interferometers these cases have limited applicability.

3.3.2 uv -coverage

Our recovery of the true sky intensity is limited by our ability to sample the uv -plane. A radio interferometer is made out of a finite number of antennas. It has a minimum baseline separation set by the physical size of the antennas, and a maximum extent set by largest achievable antenna separation. Finding a design that provides optimal uv -coverage is a complicated trade-off between costs, that scale with the required computational power for the number of baselines and their maximum length, FoV and sensitivity. Adding to that, different science cases will have different definitions of an optimal uv -coverage. In general, we therefore never fully sample the Fourier transform of the sky intensity (see Figure 3.2) instantaneously. We discuss later how we alleviate this problem.

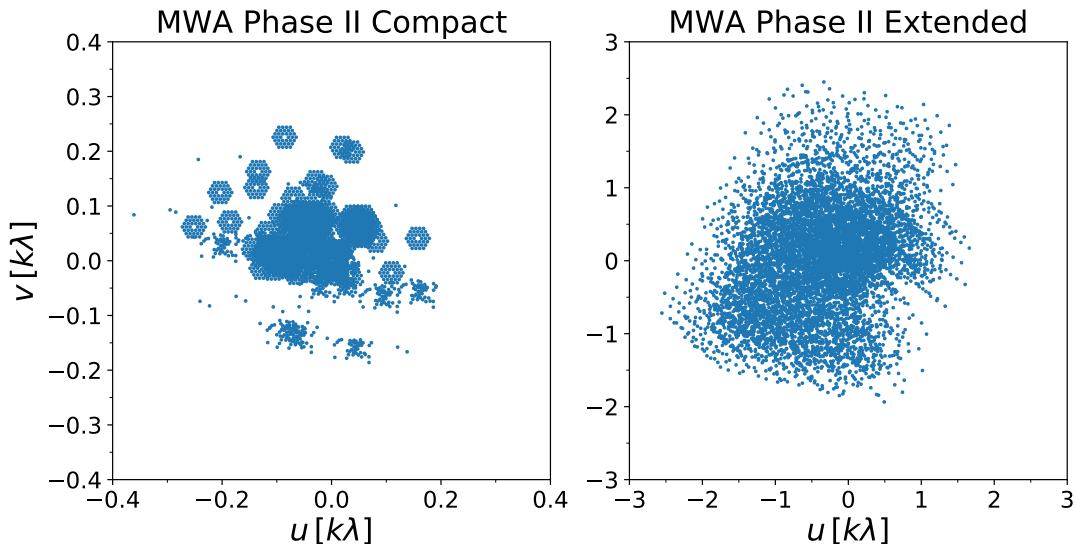


Figure 3.2: The instantaneous uv -coverage for the Murchison Widefield Array Phase II in the compact configuration (*left*), and the extended configuration (*right*). The extended configuration has baselines that are more than 10 times longer than the compact configuration and lends itself for imaging. The compact configuration on the other hand oversamples Fourier modes relevant to the EoR experiment.

The incomplete sampling of the uv -plane can be described by a sampling function $S(\mathbf{b})$, that comes into play when we Fourier transform our measurements to recover the sky intensity:

$$I_{\text{dirty}}(\mathbf{s}, \nu) = \mathcal{F}\{S(\mathbf{b})V_\nu(\mathbf{b})\} \quad (3.11)$$

Here, we introduced the Fourier operator \mathcal{F} , I_{dirty} , and the “dirty” image of the sky. Using the convolution theorem $\mathcal{F}\{X \cdot Y\} = \mathcal{F}\{X\} \otimes \mathcal{F}\{Y\}$ (Bracewell, 2000), where \otimes indicates convolution, we extract the effect of the sampling function more clearly.

$$I_{\text{dirty}}(\mathbf{s}, \nu) = I_{\text{true}}(\mathbf{s}, \nu) \otimes \tilde{S}(\mathbf{s}), \quad (3.12)$$

where $\tilde{S}(\mathbf{s})$ is the Fourier transform of the sampling function. Equation (3.12) shows that by directly Fourier transforming our measurements in the uv -plane into the image plane we do not get our true sky I_{true} . Instead, we get one that is convolved with $\tilde{S}(\mathbf{s})$, hence called the dirty image. This $\tilde{S}(\mathbf{s})$ is also referred to as the synthesised beam, the point spread function (PSF) of the interferometer. The PSF may not be compact, and subsequently imprint structure across the FoV. Recovering the true sky requires deconvolution to mitigate artefacts introduced by the limited sampling of the uv -plane. However, even for non-imaging experiments, e.g. detection of the EoR PS, optimal sampling of the uv -plane remains important because it enables a better detection of an EoR PS signal. This is discussed in Chapter 4.

3.3.3 Earth Rotation Synthesis

We can estimate the 21-cm PS without imaging. However, similarly to imaging our estimate of the 21-cm PS improves if we increase our sampling of the uv -plane. In both radio imaging and 21-cm PS experiments we use rotation synthesis to improve our sampling of the uv -plane. As the Earth rotates, the orientation of

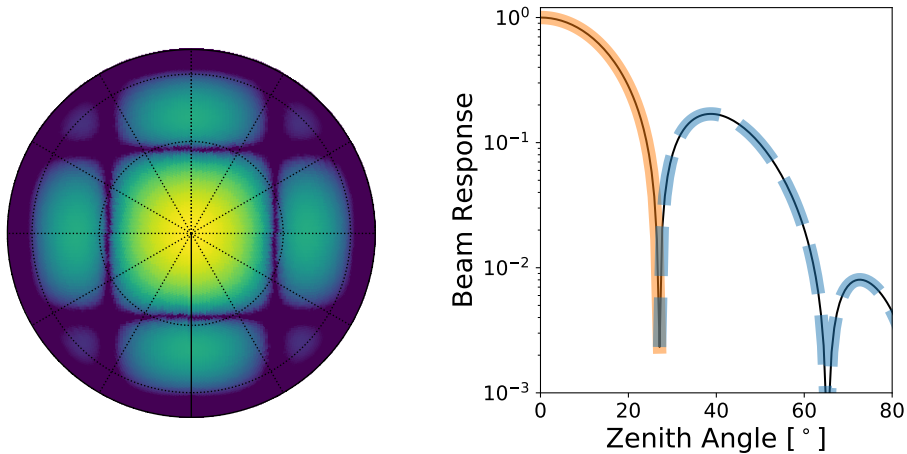


Figure 3.3: An analytic model of the primary beam response for the Murchison Widefield Array at 150 MHz. *left*: A full-sky view on the MWA primary beam. *right*: A slice through the phase center of the MWA responses clearly showing prominent sidelobes.

the sky changes with respect to the array, this also changes the projected baseline separation \mathbf{b} creating elliptical baseline tracks in the uv -plane (Thompson et al., 2017). This allows us to sample more Fourier modes, filling up our uv -plane at the cost of a larger range of w -values.

3.4 The Primary Beam

While revisiting the derivation of the measurement equation we did not consider the response pattern of the antennas. Antennas in general do not have a uniform response to all directions of the sky. Figure 3.3 shows the antenna response of an MWA tile. Including the voltage beam $a(\mathbf{s}, \nu)$ in Equation (3.10) for each antenna yields:

$$V_\nu(\mathbf{b}_{pq}) = \int a_p(l, m, \nu) a_q^*(l, m, \nu) I_\nu(l, m) e^{-2\pi i(u_{pq}l + v_{pq}m)} dl dm. \quad (3.13)$$

It is a common assumption in radio astronomy that all antennas are identical $a_p(l, m, \nu)a_q^*(l, m, \nu) = |a(l, m, \nu)|^2 = A(l, m, \nu)$ (Thompson et al., 2017). The power pattern $A(l, m, \nu)$ that is projected onto the sky is referred to as the primary beam (as opposed to the synthesised beam mentioned earlier). In reality, however, the antenna responses are not the same. This can be due to manufacturing imperfections, hardware failure, or by design. Variations in the tile responses across the MWA were measured by Neben et al. (2015) and Line et al. (2018), and we discuss the implications of these variations in Chapter 6.

There are two related parameters that describe the antenna response; the effective area A_{eff} , and the beam solid angle Ω_A (Napier, 1999). These two parameters are related through (Kraus, 1950)

$$A_{\text{eff}}\Omega_A = \lambda^2. \quad (3.14)$$

If the noise is not sky dominated, the effective area together with the noise of the system define the sensitivity. We measure this as the system equivalent flux density (SEFD). The beam solid angle is a measure of the FoV, which determines the survey capabilities of an array. For an EoR PS measurement it sets a limit on the fluctuation scales we can probe. This adds another trade-off to the design of radio telescopes between instantaneous sensitivity and FoV.

3.4.1 Sidelobe Response

The primary beam can be divided into a main lobe and sidelobes. Ideally, antennas have low sensitivity away from the pointing center. In that case the main lobe describes the bulk of the response of the antenna towards the pointing center. In radio telescopes like the MWA that is not case, and the primary beam also contains relatively strong sidelobes (see Figure 3.3 and the next section). At higher frequencies grating lobes appear where the response is equal to the main lobe. The main lobe is often well characterised, whereas the sidelobe response is not. Although the sidelobes are an order of magnitude weaker than the main lobe

(Line et al., 2018), they cover a large fraction of the sky and hence their effect is non-negligible. Especially, if those parts of the sky contain bright sources, e.g. the Milky Way and Fornax A.

3.4.2 Phased Arrays

We can build a large aperture with phased arrays. The advantage of phased arrays is that it allows us to make more sensitive telescope elements without the use of large mechanical constructions. Also, they are electronically steered rather than mechanically, enabling rapid re-pointing if a transient event occurs. Additionally, if the signals of each element are digitised we can tailor the FoV to various science cases, or even form multiple beams on the sky for rapid scanning.

In a phased array the signals from a cluster of antennas are added coherently by summing the response of different elements multiplied with a delay. These delays point the antenna responses in the desired direction and are described by the array factor, AF ,

$$AF = \sum_{n=1}^{N_{dipoles}} w_n \exp [i(k_x x_n + k_y y_n + k_z z_n)], \quad (3.15)$$

where x_n , y_n , and z_n are the x , y , z coordinates, respectively, of antenna n with respect to the centre of the phased array. w_n is a complex weight with amplitude ranging from 0 to 1. We also define direction cosines for the array factor in spherical coordinates

$$\begin{aligned} k_x &= \frac{2\pi}{\lambda} \sin \theta \sin \phi \\ k_y &= \frac{2\pi}{\lambda} \sin \theta \cos \phi \\ k_z &= \frac{2\pi}{\lambda} \cos \theta, \end{aligned} \quad (3.16)$$

where θ is the zenith angle and ϕ is the azimuth angle (Balanis, 2016). We compute the response for a full MWA tile by multiplying the response of a single dipole on a ground screen with the array factor

$$b_{\text{tile}} = AF \times b_{\text{dipole}} \times b_{\text{ground}}. \quad (3.17)$$

Note: this model does not describe mutual coupling that can significantly impact the realistic response of a phased array. We revisit this in Chapter 6.

Although, the benefits of phased arrays now seem obvious, there are also a few downsides to phased arrays. Particularly, in those with a regular lay-out, i.e. the MWA and to some extent LOFAR. The behaviour of the primary beam response in phased arrays is highly chromatic, as the array factor has a frequency dependence. When the spacing between individual dipoles is larger than $\lambda/2$ the array becomes sparse and its effective area scales as $N \times \lambda^2$. When the spacing is smaller than $\lambda/2$ the array becomes dense and its effective area is capped at $N \times \lambda_{\text{dense}}^2/2$. This sparse regime may seem more beneficial, however, it does suffer from grating lobes that have response away from pointing centre comparable to the main lobe. A final drawback, is that the effective area depends on zenith angle, as the projected surfaces changes (Braun & van Cappellen, 2006).

3.5 Calibration

The ionosphere and the instrumental signal chain leave an imprint on the amplitude and phase of the radio signals we aim to detect. Calibration is the process that edits the data to undo these instrumental or ionospheric propagation effects (see the next sections for examples). In the most general form we describe these effects by including a complex gain factor $g_p(l, m, \nu)$ for every antenna p , to Equation (3.13)

$$V(\mathbf{b}_{pq}, \nu) = \int g_p g_q^* a_p a_q^* I_\nu e^{-2\pi i(u_{pq}l + v_{pq}m)} dl dm. \quad (3.18)$$

Here, all quantities are implicitly dependent on sky coordinates (l, m) and frequency for the sake of space. The complex nature of the gains enables us to describe any changes in amplitude and phase of the incoming radio wave. Cali-

bration solves and corrects for changes to the incoming signals. In essence, it is a parameter estimation problem in which we estimate the gains g_p that impact our true visibilities V_{pq}^{true} based on the measured data V_{pq}^{data} . The most commonly used method to estimate the gains is so called sky-based calibration. However, more recently redundant calibration has undergone a renaissance. We discuss these calibration techniques in more detail in the next section, but first we discuss contributions to the complex gains.

3.5.1 Antenna Gain

One instrumental effect is the overall response of the signal chain from receiver to correlator. This direction independent effect (DIE) can change both the amplitude and phase of the measured signal. This is most often referred to when discussing the gain. However, we can also consider variations in the antenna response that deviate from the modelled primary beam to be antenna gains. This is an example of a direction dependent effect (DDE).

3.5.2 Bandpass

Of particular importance to the EoR experiment is the overall frequency response of the signal chain; the bandpass. Ideally, the relative change in the gains as a function of frequency is small over the bandwidth of an EoR observation. Any small scale variation in the gain can lead to leakage of foreground signals (Barry et al., 2016; Patil et al., 2016; Trott & Wayth, 2016). This has motivated a strong push from various EoR collaborations for spectral smoothness in their instrument design, and developments of calibration algorithms that generate spectrally smooth calibration solutions (see 3.6.1 and Chapter 6).

3.5.3 Ionospheric Effects

In our previous discussion the gain and the bandpass describe the instrumental imprint on the measured visibilities. However, before the signals reach the instrument they pass through the ionosphere. The ionosphere is filled with charged particles due to the interaction between UV and X-ray radiation from the Sun and atoms in the upper layers of the atmosphere. These charged particles refract radio signals of interest and this effect is greater at low frequencies because the refraction angle grows as $\sim \lambda^2$. However, it also varies strongly with the state of the ionosphere (Mevius et al., 2016; Jordan et al., 2017). In contrast with the signal chain of the instrument, the ionosphere primarily adds a phase delay that manifests itself as a refractive shift. For the ionosphere above the MWA a median source offset of $\sim 0.2'$ was found (Jordan et al., 2017).

3.6 Sky-Based Calibration

The most commonly used form of radio telescope calibration uses a sky model based approach. First developed by Cornwell & Wilkinson (1981), sky based calibration starts out with an a priori model for a known source or a set of multiple sources, and then aims to minimise the difference between the measured data V_{pq}^{data} and the model V_{pq}^{model} by varying the antenna gains g_p .

$$\chi^2 = \sum_{pq} w_{pq} |V_{pq}^{\text{data}} - g_p g_q^* V_{pq}^{\text{model}}|^2 \quad (3.19)$$

Traditionally, and still today at higher frequencies, these known sources can be different to the target sources of interest. They merely have to be sufficiently nearby that the calibration solutions obtained can be applied to the actual target. Low-frequency telescopes, however, often have large enough FoVs that encompass both bright calibrator sources and science targets. A major advantage of sky-based calibration techniques is the ability to solve for DDEs. With enough bright sources in different directions, we can calibrate explicitly in the direction of each

source. This enables sky-based calibration to capture deviations from the beam model and variations in the ionosphere.

A major disadvantage, however, is its dependence on an a priori model. Sources present in the data but not included in the model bias the calibration solutions (Grobler et al., 2014). Applying these modified calibration solutions to the data and then attempting to image the data leads to ghost sources, spurious emission, suppression of real emission or the deformation of extended sources Grobler et al. (2014); Wijnholds et al. (2016). Closer to EoR science, calibration solutions that are obtained using incomplete sky models increase the contamination of the 21-cm PS by foreground sources as their spectral structure is imprinted on the gain solutions (Barry et al., 2016; Ewall-Wice et al., 2017; Patil et al., 2016). We discuss this in full detail in Chapters 6 and 7.

3.6.1 Recent Developments

Sky based calibration was extremely successful in the past decades. However, the push in sensitivity by the latest generation of low frequency radio telescopes puts stringent requirements on the level of accuracy of the calibration solutions. This resulted in a major effort over the past decade into calibration schemes that compensate DDEs (Mitchell et al., 2008; Yatawatta et al., 2009; Kazemi et al., 2011; Sullivan et al., 2012; van Weeren et al., 2016). To enable these direction dependent calibration schemes we had to overcome a number of obstacles:

- the computational ability to calibrate on time scales of ionospheric variations
- the computational ability to calibrate in a sufficient number of directions
- taking the curvature of the sky into account (w -term in Equation 3.8) due the large field of view
- modelling the diffuse emission detected by the shortest baselines

These challenges were addressed by different collaborations and resulted in a suite of implementations. Most notable are the efforts of the different LOFAR teams necessitated by LOFAR’s sensitivity to ionospheric distortions. Due to LOFAR’s large FoV its international baselines (> 100 km) are prone to DDEs when those baselines probe scales smaller than the length scale of ionospheric disturbances. Here, we give a brief overview of the different calibration pipelines.

SAGEcal, developed by Yatawatta et al. (2009) and Kazemi et al. (2011) from the LOFAR EoR group solves for the full Jones matrix in ~ 100 different directions. In their implementation they resolve the computational load by applying the SAGE algorithm (Fessler & Hero, 1994) instead of brute force non-linear least squares (Levenberg, 1944; Marquardt, 1963). To constrain the degrees of freedom **SAGEcal** enforces the solutions be spectrally smooth through regularisation, i.e. invoking a penalty for non-smoothness. This relies on the assumption that the frequency dependent response of the instrument is smooth on some scale.

Factor, developed by van Weeren et al. (2016) from the LOFAR Calibration & Imaging Tiger Team, reduces the computational load of calibration in multiple directions by breaking up the sky in facets (Perley, 1989) around bright radio sources, so-called facet calibration. It then proceeds to determine calibration solutions for a facet by calibrating on its central bright source. This relies on the assumption that ionospheric variations occur on large scales (Cohen & Röttgering, 2009), and therefore vary smoothly across the sky. Unlike **SAGEcal** it does not solve for all directions simultaneously, instead it solves for one facet at a time while correcting and imaging that facet.

RTS (Mitchell et al., 2008) relies on the excellent instantaneous uv -coverage of the MWA. During a snapshot the baselines in the uv -plane are nearly co-planar. This enables the use of Equation (3.9) despite the large FoV of the MWA tiles. This alleviates the curved sky problem and enables fast imaging to correct for variations in the ionosphere. Similar to **SAGEcal** it also determines full Jones matrix solutions towards ~ 100 bright radio sources.

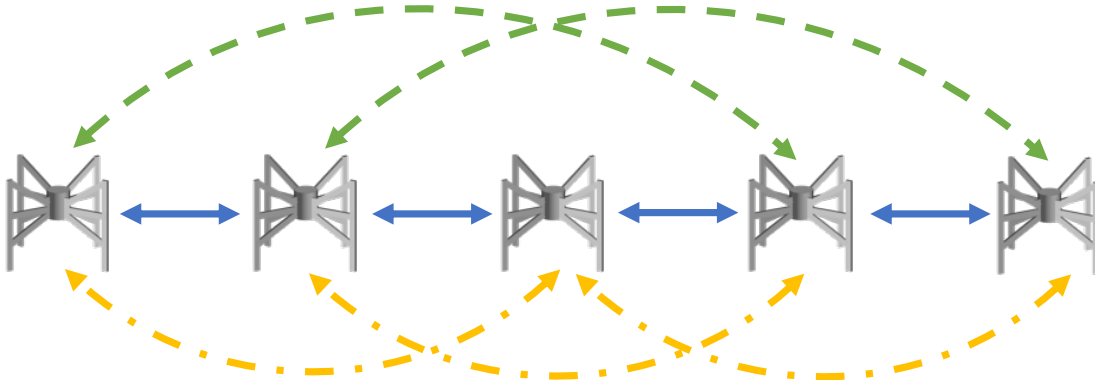


Figure 3.4: A minimum redundant array and the redundant baselines in this array.

3.7 Redundant Calibration

The underpinning principle of redundant calibration is that baselines with the same orientation and length, i.e. redundant baselines, measure exactly the same visibilities. If a radio interferometer measures a sufficiently large enough number of redundant visibilities, we can solve for both the gains and redundant visibilities. This is fundamentally different from sky based calibration, because it does not require prior knowledge of the sky to which our data is matched. We achieve redundancy by creating a highly regular lay-out similar to the MWA hexes, HERA or PAPER. Redundant calibration consists of two steps; *relative calibration* and *absolute calibration*.

3.7.1 Relative Calibration

Because redundant arrays have more measurements than unique baselines we create a closed system of equations for the true visibilities V_{pq}^{true} and the gain solutions g_p if the array is redundant enough (Wieringa, 1992). The simplest redundant array for which such a system is solvable are 5 equidistant antennas on a line (see Figure 3.4).

This array has 10 baselines organised in 3 unique redundant baseline groups: 4 redundant baselines at spacing Δx , 3 redundant baselines at spacing $2\Delta x$, 2 redundant baselines at spacing $3\Delta x$. With 5 unknown antenna gains and 3 unique visibilities there are 8 unknowns for this array, and 10 measurements, i.e. the system is overdetermined. There are two algorithms that solve this system using the assumption of redundancy; `logcal` (Wieringa, 1992) and `lincal` (Liu et al., 2010). We will discuss both algorithms in the context of the 5 element interferometer using only the first two redundant groups.

Logcal

A problem with Equation (??) is that in its current form the gains and visibilities are non-linearly coupled to the data. `logcal` linearises the equation by taking the natural logarithm, as the name suggests. Noting that all complex quantities, can be written as $z = |z|e^{i\phi z}$, we can separate this “logarised” equation into an equation for the amplitudes and one for the phases

$$\begin{aligned}\ln|V_{pq}^{\text{data}}| &= \ln|g_p| + \ln|g_q| + \ln|V_{pq}^{\text{True}}|, \\ \arg|V_{pq}^{\text{data}}| &= \phi_q - \phi_p + \arg|V_{pq}^{\text{True}}|\end{aligned}\tag{3.20}$$

Here, $|g_p|$ is the gain amplitude and ϕ_p is the gain phase of antenna p , $\ln|V_{pq}|$ is the true visibility amplitude and $\arg|V_{pq}|$ is the true visibility phase measured by a baseline pair p and q . Because the amplitude and phase decouple, we can rewrite this into two different matrix equations that can be solved independently.

$$\mathbf{y}_\alpha = \mathbf{A}_\alpha \mathbf{x}_\alpha,\tag{3.21}$$

where the index $\alpha = \{\eta, \phi\}$, η for the amplitude equations and ϕ for the phase equations, \mathbf{y} are the measured data amplitude and phase vectors, \mathbf{x}_η contains the gain amplitude $\ln|g_p|$ and visibility amplitude $\ln|V_{pq}|$, \mathbf{x}_ϕ contains the gain phase ϕ_q and visibility phase $\arg|V_{pq}|$. \mathbf{A} is the matrix that maps the gains and unique

visibilities to the measured data. Equation (3.22) shows this explicitly for the phase

$$\begin{pmatrix} \arg|c_{12}| \\ \arg|c_{23}| \\ \arg|c_{34}| \\ \arg|c_{45}| \\ \arg|c_{13}| \\ \arg|c_{24}| \\ \arg|c_{35}| \end{pmatrix} = \begin{pmatrix} -1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & -1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 & 1 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 & 1 & 0 & 1 \end{pmatrix} \begin{pmatrix} \phi_1 \\ \phi_2 \\ \phi_3 \\ \phi_4 \\ \phi_5 \\ \arg|v_1| \\ \arg|v_2| \end{pmatrix}, \quad (3.22)$$

where $\arg|v_1|$ is the phase of the visibility measured by the Δx spacings and similarly $\arg|v_2|$ for the $2\Delta x$ spacings. This system is, however, not well-behaved because there are 4 intrinsic degeneracies that need to be constrained to yield a unique solution. These four intrinsic degeneracies are the absolute amplitude, the absolute phase and the phase gradient in the x - and y -direction.

- The absolute amplitude: we can arbitrarily increase all gains with a factor A if we simultaneously divide all visibility amplitudes by A^2 . This causes the sky to become brighter or fainter. The simplest constraint we can add is a reference antenna of which we define the gain amplitude $|g_0| = 1$ or $\ln|g_0| = 0$
- The absolute phase: we can add an arbitrary phase term Φ to the gain solutions, that will leave the visibilities unchanged due to the product $g_p^* g_q = |g_p||g_q|e^{i(\phi_p + \Phi - \phi_q - \Phi)}$. We resolve this in a similar fashion by setting the phase of a reference antenna to zero, $\phi_0 = 0$.
- The phase gradient: we can add a phase slope in the x - and y -direction to the array, while simultaneously taking that phase out of our visibilities. This corresponds to tilting the array in either the x - or y -direction and rotating the sky in the opposite direction. We constrain this tip/tilt degeneracy by

setting $\sum_i x_p \phi_p = 0$ and $\sum_p y_p \phi_p = 0$.

Adding these constraints to the system in Equation (3.22) removes the degeneracy. However, these degenerate parameters still need to be fit to create usable solutions at the final stage of calibration. This final step is often referred to as absolute calibration and we will discuss this at the end of this section.

When I took the logarithm of Equation (??), I ignored noise. In reality, Gaussian distributed receiver noise affects our measurements. However, this noise is only Gaussian in real and imaginary space. In amplitude and phase space it is no longer Gaussian and does not have a zero mean. In other words the solutions will be biased. This is a major limitation of `logcal` and this flaw has spawned the creation of `lincal` (Liu et al., 2010).

Lincal

`lincal` was specifically designed as an implementation of redundant calibration that does not suffer from noise bias. In this algorithm the visibility-gain equation is linearised by taking a Taylor expansion of Equation (3.7) with respect to the true solutions g_i around some fiducial guess of the gain g_p^0 and visibility V_{pq}^0 (Liu et al., 2010).

$$V_{pq}^{\text{data}} = g_p^0 g_q^{0*} V_{pq}^0 + g_q^{0*} V_{pq}^0 \Delta g_p + g_p^0 V_{pq}^0 \Delta g_q + g_p^0 g_q^{0*} \Delta V_{pq} \quad (3.23)$$

In this system of equations we want to solve for the corrections Δg_p and ΔV_{pq} that direct us to the true solutions. Because `lincal` is an iterative algorithm it takes a number of iterations before convergence is achieved and this depends highly on the initial guesses g_p^0 and V_p^0 . In practice `logcal` and `lincal` are used in tandem (Zheng et al., 2014). `logcal` produces the initial guess, which is then fed to `lincal` for further refinement. Despite the fact `lincal` works with the full complex gains and visibilities, it is imperative to note that the earlier mentioned degeneracies are still present.

3.7.2 Absolute Calibration

To fix these degeneracy parameters a final calibration step that uses external information is required. This step therefore relies on sky model information and hence is subject to similar problems as normal sky model calibration (Byrne et al., 2019). There are various ways to perform this final step. The total gain g_p , the relative gain h and the absolute gain parameters are related through

$$g_p = Ae^{i[\Delta + \Delta_x x_p + \Delta_y y_p]} \times h_p. \quad (3.24)$$

We can find the absolute gain A by averaging over the sky based calibration solutions and fitting the phase gradient terms over the sky based solutions (Li et al., 2018; Byrne et al., 2019), or by using the relative calibration solutions as a starting point for sky based calibration or vice versa (Li et al., 2019). Alternatively, one can fit for the degeneracy parameters directly by minimising the difference between model visibilities and the data (Kern et al., 2019).

3.7.3 Recent Developments

Efforts into developing publicly available redundant calibration pipelines have been less prominent. Particularly, because redundant arrays are primarily used in EoR science. Both `logcal` and `lincal` are implemented in the publicly available `OMNICAL` (Zheng et al., 2014). This has been used in both in the analysis of PAPER data (Ali et al., 2015) and the MWA Phase II compact data (Li et al., 2018; Li et al., 2019). The HERA team is currently developing their implementation of redundant calibration¹¹.

¹¹https://github.com/HERA-Team/hera_cal

3.8 Comparing Calibration Strategies

Redundant and sky-based calibration are fundamentally different. Each calibration method has its own advantages and drawbacks. Sky-based calibration is generically applicable to any radio interferometric array. A major drawback is that its accuracy relies on the completeness of a sky model. On the other hand, redundant calibration is agnostic of the sky and does not require pre-requisite knowledge. This is advantageous when trying to find an unknown signal in the Universe, and avoiding bias is crucial to success. At first glance, redundant calibration seems perfectly suited for EoR experiments. However, its effectiveness depends on a sufficiently large signal-to-noise ratio across all redundant baselines (see Chapter 5). This means that redundant calibration depends on the structure of the radio sky, even though it may be independent of a sky model. Additionally, it is practically impossible to build a perfectly redundant telescope, and hence its accuracy is determined by the precision at which the array is deployed. In Chapters 5, 6, and 7 I investigate the accuracy and precision of sky based and redundancy based calibration methods, quantify their impact on the 21-cm PS, and motivate the need for further developments of hybrid approaches that combine the best of both techniques.

Chapter 4

Statistical Cosmology

In the absence of sensitive telescopes that directly image the Epoch of Reionisation (EoR), our best hope to understand this early time of the Universe is through a statistical detection of the 21-cm signal. The power spectrum (PS), used for Cosmic Microwave Background (CMB) experiments that study the afterglow of the Big Bang, is a powerful tool that enables us to study the spatial structure of the neutral hydrogen field during reionisation. This chapter discusses the methodology behind the 21-cm power PS, connects it to the physics of the EoR, the fundamentals of radio interferometry, and error propagation tools that form the basis for the research presented in this thesis.

4.1 The 21-cm Power Spectrum

In the early 2000s it became clear that direct tomographic measurements of the EoR, despite the brightness of foregrounds, would be possible due to the smooth frequency structure of both galactic synchrotron emission (Shaver et al., 1999) and extragalactic emission (Di Matteo et al., 2002). It was suggested that multi-frequency angular power spectra could exploit this difference in spectral behaviour (Zaldarriaga et al., 2004), while retaining the ability to discriminate between different reionisation histories (Furlanetto et al., 2004). The angular power spectrum

measures fluctuations in signals that lie on spherical surfaces and uses a spherical harmonic basis set to represent such fluctuations. A multi-frequency adaption of techniques that were originally developed for CMB experiments demonstrated that the angular power spectrum enabled both the separation of foregrounds and EoR signals, and also allowed us to study the evolution of the EoR as a function of cosmic time (Zaldarriaga et al., 2004). Around the same time it was proposed to use the 3-dimensional power spectrum (3D-PS) to probe the 3-dimensional nature of the EoR using Fourier analysis (Bharadwaj & Sethi, 2001; Zaldarriaga et al., 2004; Morales & Hewitt, 2004). Measuring fluctuations in the signal over a volume, rather than a surface, accesses more information and is more sensitive to an EoR signal.

4.1.1 Mathematical Formalism

Formally, the PS is defined as the Fourier transform of the auto-correlation function of the differential brightness temperature field δT . It is a measure of the spatial variance of the 21-cm signal and is given by

$$P(\mathbf{k}) = \int_V \langle \delta T(\mathbf{r}') \delta T(\mathbf{r} + \mathbf{r}') \rangle e^{i\mathbf{k}\cdot\mathbf{r}'} \frac{d^3\mathbf{r}'}{(2\pi)^3}. \quad (4.1)$$

Here, $\delta T(\mathbf{r})$ is the differential brightness temperature at position $\mathbf{r} = (x, y, z)$, \mathbf{k} is the Fourier dual of \mathbf{r} , i.e. a spatial frequency scale, and $\langle \rangle$ denotes averaging over different realisations of the brightness temperature field. If our observational volume is large enough we can assume ergodicity, and approximate an ensemble average with a spatial average¹

$$P(\mathbf{k}) = \frac{1}{V} \int_{V'} \int_V \delta T(\mathbf{r}') \delta T(\mathbf{r} + \mathbf{r}') e^{-2\pi i\mathbf{k}\cdot\mathbf{r}'} d^3\mathbf{r} \frac{d^3\mathbf{r}'}{(2\pi)^3}. \quad (4.2)$$

Analogous to the convolution theorem, we can write the Fourier transform of a cross-correlation as $\mathcal{F}\{f \star g\} = \mathcal{F}\{f\} \cdot \mathcal{F}\{g\}^*$ using the Wiener-Kinchin Theorem,

¹Rather than relying on finding multiple realisations of the Universe.

where the operators \mathcal{F} , \star and $*$ denote a Fourier transform, cross-correlation, and complex conjugation, respectively. Using the theorem we write Equation (4.2) as

$$P(\mathbf{k}) = \frac{1}{(2\pi)^3 V} \tilde{\delta T}(\mathbf{k}) \tilde{\delta T}^*(\mathbf{k}). \quad (4.3)$$

This is also commonly written as

$$\langle \tilde{\delta T}(\mathbf{k}) \tilde{\delta T}^*(\mathbf{k}') \rangle = (2\pi)^3 \delta(\mathbf{k} - \mathbf{k}') P(\mathbf{k}), \quad (4.4)$$

where $\delta(\mathbf{k} - \mathbf{k}')$ is the Dirac delta function.

Under the assumption that the EoR signal is homogeneous and isotropic, a fundamental assumption in cosmology, we can compute the spherically averaged 1-dimensional power spectrum (1D-PS) $P(k)$ from the 3D-PS $P(\mathbf{k})$. Collapsing this 3D-space into 1 dimension enables us to average over all modes \mathbf{k} that fall within a shell with width Δk around $k = |\mathbf{k}|$. This increases the detectability of the signal, because we average over a large set of measurements. It is common to present the ‘dimensionless’ power spectrum, defined as

$$\Delta^2(k) = \frac{k^3}{2\pi^2} P(k). \quad (4.5)$$

The dimensionless power spectrum normalises the power with respect to the number of modes that are measured in the volume². Figure 4.1, shows a set of dimensionless power spectra at different redshifts for a faint galaxy driven reionisation model (Mesinger, 2016).

4.1.2 Estimation Theory

Estimating power spectra has a long history in CMB experiments and galaxy surveys, and those methods have been modified for 21-cm purposes (Liu & Tegmark, 2011; Dillon et al., 2013; Liu et al., 2014; Dillon et al., 2015; Trott et al., 2016).

²Historically, this originates from studies of the matter density power spectrum, where the dimensionless power spectrum is truly dimensionless.




Figure 4.1: Theoretical 21-cm power spectra for a faint galaxy driven reionisation model (Mesinger, 2016). The power spectra for different redshifts show significant change in the shape and amplitude.

The problem we face is that we aim to estimate $P(k)$, a continuous function, from a discretised data vector \mathbf{x} . Each entry in this data vector is a sample of the probability distribution that describes the Fourier transform of the differential brightness temperature field $\delta\tilde{T}$. A common assumption we make is that our visibility data are described by a multivariate Gaussian with zero mean, and covariance described by the matrix \mathbf{C} . The likelihood \mathcal{L} , i.e. the probability that our measurements match our model, for data from such a distribution is given by

$$\mathcal{L}(\mathbf{x}; \mathbf{C}) \propto \frac{1}{\det(\mathbf{C})} e^{-\mathbf{x}^\dagger \mathbf{C}^{-1} \mathbf{x}}. \quad (4.6)$$

Here, \mathbf{x} is a complex data vector containing interferometric measurements, \dagger denotes the Hermitian transpose, and \mathbf{C} is the data covariance matrix. The diagonals of this matrix describe the variance (or power at a particular Fourier mode) and the off-diagonals describe how this power is correlated between different \mathbf{k} -modes sampled by our data vector. The data covariance matrix and the data vector are related through

$$\mathbf{C} \equiv \langle \mathbf{x} \mathbf{x}^\dagger \rangle - \langle \mathbf{x} \rangle \langle \mathbf{x}^\dagger \rangle. \quad (4.7)$$

In a Fourier basis our aim is to estimate the diagonals of the covariance matrix describing the power of 21-cm signal \mathbf{C}_{21} ; they encode the power of the fluctuations at a particular scale size. Although our data are not truly Gaussian, both the EoR signal (Bharadwaj & Pandey, 2005) and foregrounds are non-Gaussian (Ollier et al., 2017), the power spectrum can still prove to be a useful metric, as demonstrated by Figure 4.1. Realistically, our data \mathbf{x} contain more than the 21-cm signal of interest; they also contain foreground sources below the detection threshold, and thermal noise. These different contributions to our data can be described by their own probability distribution, and because they are independent we can write the total covariance of our data as

$$\mathbf{C} = \mathbf{C}_{\text{FG}} + \mathbf{C}_{\text{N}} + \mathbf{C}_{21}, \quad (4.8)$$

where \mathbf{C}_{FG} , \mathbf{C}_{N} , and \mathbf{C}_{21} describe the covariances of the foregrounds, thermal noise, and 21-cm signal, respectively. We will derive the form of similar covariance matrices in Chapter 6. It is important to stress that the full form of \mathbf{C} is not known. Hence, the descriptive power of covariance matrices is limited by our knowledge, and methods employing these matrices are subject to that limitation.

This statistical perspective on our data provides us with powerful tools to study various aspects of EoR power spectrum estimation. A full treatment of finding the optimal power spectrum estimator is beyond the scope of thesis. I will only discuss the general steps taken to derive an estimate for the diagonal entries of \mathbf{C}_{21} . One particular procedure involves finding an estimator p_α of the power spectrum at a particular scale k_α through the principle of maximum likelihood. To do so we take the logarithm of the likelihood, compute its derivative, and solve for its roots. This path often leads to an estimator that takes a quadratic form (Liu & Tegmark, 2011; Dillon et al., 2013; Liu et al., 2014; Dillon et al., 2015; Trott et al., 2016)

$$p_\alpha = (\mathbf{x} - \mathbf{m})^\dagger \mathbf{C}^\alpha (\mathbf{x} - \mathbf{m}) - b_\alpha. \quad (4.9)$$

Here, \mathbf{m} is a model of our data, \mathbf{C}^α is the derivative of the matrix that describes covariances in our data, b_α is a constant for a particular k -bin α . The exact form of \mathbf{C}^α depends on how we treat our data and transform it to \mathbf{k} -space. This matrix notation enables us to describe a wide range of ‘propagation’ effects. Particularly relevant to this thesis is, for instance, the transformation from the measurement domain $\{u, v, f\}$ to the power spectrum domain $\{u, v, \eta\}$. Covariance matrices enable us to do that fairly straightforwardly, because the discrete Fourier transform can be described by a matrix \mathcal{F}

$$\mathbf{C}(u, v, \eta) = \mathcal{F}^\dagger \mathbf{C}(u, v, f) \mathcal{F}. \quad (4.10)$$

4.2 Power Spectrum Estimators

The Universe does not provide us with perfect differential brightness temperature boxes. Instead, we need to extract the information from incomplete radio interferometric data. There are various ways to estimate the 21-cm PS from the data (Liu & Shaw, 2019).

Image Based estimators Fourier transform interferometric visibilities into the image domain to integrate the many hours required to detect the EoR (Jacobs et al., 2016; Patil et al., 2016; Barry et al., 2019a). After sufficiently deep integrations, they can apply Equation (4.3) by computing the 3D Fourier Transform, square that and normalise it by the volume V . In practice, however, the integrated images are Fourier transformed back into the $\{u, v, f\}$ domain before techniques that estimate the Fourier transform over frequency are applied (Patil et al., 2016; Barry et al., 2019a).

Visibility Based estimators (Choudhuri et al., 2014; Trott et al., 2016) remain in the Fourier domain and therefore avoid problems associated with the transformations from the Fourier domain to the image domain, e.g. non-Gaussian noise that arises from the incomplete sampling of the uv -plane I discussed in Chapter 3. If the FoV is small enough, the visibilities can simply be gridded and interpolated on a regular grid in uv -plane before the Fourier transform over the frequency axis is performed. The drawback is that beam corrections are harder, and the w -term I discussed in Chapter 3 is non-negligible. This means that these curvature effects need to be handled explicitly as part of the gridding process (Trott et al., 2016).

Delay Based estimators (Parsons et al., 2012b) Fourier transform the measured visibility over the frequency axis of a particular baseline, rather than gridding baselines onto a discretised uv -plane. Not only does this per-baseline ap-

proach avoid the computational costs of gridding, it also separates foregrounds from EoR signals in so-called delay space and avoids a problem called “mode-mixing”, which we discuss in the next section. The delay transform, however, only approximates the power spectrum for the shortest baselines.

Harmonic Based estimators (Eastwood et al., 2019) use m -mode analysis on the measured visibilities to decompose the sky intensity into spherical harmonic functions (Shaw et al., 2014, 2015). These can be used to compute the angular power spectrum. The relationship between spherical harmonics and plane waves, i.e. the basis functions of the Fourier transform, is then used to convert an angular power spectrum into a spatial power spectrum (Eastwood et al., 2019).

4.3 Relating observing and cosmological units

Throughout this thesis I take the perspective of a visibility based estimator. I discussed in Chapter 3 that radio interferometers naturally measure the Fourier transform of the specific intensity $I(\nu)$. In Chapter 2 I discussed the linear relationship between specific intensity and brightness temperature. We therefore only require to Fourier transform these measured visibilities over frequency for our power spectrum measurements

$$V(u, v, \eta) = \int \Gamma(\nu) V(u, v, \nu) e^{-2\pi i \nu \eta} d\nu. \quad (4.11)$$

The function $\Gamma(\nu)$ is a tapering function applied to our data before the Fourier transform over frequency. The taper suppresses leakage of power due to the finite bandwidth of the data. Throughout this thesis I use a Blackman-Harris window as a taper. It suppresses leakage by 74 dB at the first sidelobe, however, we pay the price for this by a reduction in the effective bandwidth by a factor of 2 (Harris, 1978).

Following Morales & Hewitt (2004) we can relate the observational Fourier modes $\{u, v, \eta\}$ to cosmological Fourier modes $\mathbf{k} = (k_x, k_y, k_{\parallel})$ through

$$k_{\perp} = \frac{2\pi|\mathbf{u}|}{D_M(z)} [\text{Mpc}^{-1}] \quad (4.12)$$

and

$$k_{\parallel} = \frac{2\pi H_0 f_{21} E(z)}{c(1+z)^2} [\text{Mpc}^{-1}]. \quad (4.13)$$

Here, $k_{\perp} = \sqrt{k_x^2 + k_y^2}$ is the modulus of cosmological modes perpendicular to the line of sight, and k_{\parallel} are cosmological modes parallel to the line of sight, $D_M(z)$, H_0 , f_{21} , z are the transverse comoving distance, the Hubble constant, the rest frequency of the 21-cm line, and the observational redshift, respectively. $E(z)$ is given by

$$E(z) \equiv \sqrt{\Omega_M (1+z)^2 + \Omega_k (1+z)^2 + \Omega_{\Lambda}}, \quad (4.14)$$

where I introduced the cosmological density parameters for matter Ω_M , spatial curvature Ω_k , and dark energy Ω_{Λ} (Hogg, 2000).

4.4 The Window and the Wedge

In early work, it was realised that if one includes the primary beam response in the power spectrum methodology, it acts as a window function that introduces correlations between neighbouring cells in $\{u, v, \eta\}$ -space (Morales & Hewitt, 2004)

$$\begin{aligned} V(u, v, \eta) &= \int A(\mathbf{s}, f) I(\mathbf{s}) e^{-2\pi i(\mathbf{s} \cdot \mathbf{u} + \eta f)} d\mathbf{s} df \\ &= \tilde{A}(u, v, \eta) \otimes V_{\text{true}}(u, v, \eta). \end{aligned} \quad (4.15)$$

The obvious effect of this convolution is that it destroys the independence between neighbouring measurements in Fourier space and hence we lose information. Early work that studied the accuracy of calibration algorithms and foregrounds subtraction (Bowman et al., 2009; Datta et al., 2009; Liu et al., 2009) revealed that incorrectly subtracted foregrounds lead to the rise of a wedge-like feature in

the 2D PS (see Figure 4.2).




Figure 4.2: A schematic overview of regions in the 2D-power spectrum. The spectral smoothness of the foregrounds causes them to naturally dominate low k_{\parallel} -modes. However, “mode-mixing” leaks foreground power to higher modes leading to the wedge. The remaining window should be relatively clean from foregrounds contamination. However, this depends highly on the frequency response of the signal chain, and calibration of the data.

Further analysis into the origin of “the Wedge” (Vedantham et al., 2012; Morales et al., 2012; Trott et al., 2012), revealed that contamination of foregrounds arises due to “mode-mixing”. Effectively, mode-mixing describes the impact of the point spread function (PSF) of the instrument, a combination of both the primary beam response and the uv -plane footprint in Fourier space. In Chapter 3, I described the uv -coverage of a radio interferometer with a sampling function $S(u, v, f)$. In that discussion, I did not explicitly state the frequency de-

pendence. However, because baseline length is a function of frequency, the sampling function changes as a function of frequency. In Equation (4.15) I discussed the impact of the primary beam, similarly we can consider the sampling function as another window function applied in the measurement domain. Fourier transforming this over frequency yields another convolution that takes foregrounds that naturally occupy the lowest k_{\parallel} and smear their power out into the wedge.

The wedge naturally limits the number of usable modes in the PS and forces us to work in the EoR “window” (Vedantham et al., 2012). Mode-mixing can be suppressed by appropriately gridding the data in the uv -plane (smoothing the instrument synthesized beam), and by applying a frequency window function or taper before Fourier transforming the data over frequency. The frequency taper smooths the frequency response of the instrument, reducing sidelobes in Fourier space. Murray & Trott (2018) investigate possible array lay-outs that prevent a wedge. They find that small antenna position errors render the wedge in practice unavoidable. However, radially symmetric arrays with a regular array lay-out can suppress the wedge by 3 orders of magnitude.

4.5 Power Spectrum Contamination

Although the wedge arises naturally in a radio interferometric instrument, there are a number of effects that can exacerbate its imprint.

4.5.1 Gridding

Before the measured data are Fourier transformed from the uv -domain to the image plane for coherent addition, or frequency to cosmological Fourier space they need to be placed on a discretised regular grid in $\{u, v, f\}$ -space. Offringa et al. (2019b) studies the impact of gridding on the EoR PS and found that if the gridding kernel is oversampled by a factor of 4000 the noise introduced by the process becomes negligible compared to an EoR signal.

4.5.2 RFI

Of more concern is contamination due to Radio Frequency Interference (RFI). These artificial signals are several orders of magnitude brighter than emission from astrophysical sources. Fortunately, we can excise channels that are occupied by such interference (Offringa et al., 2012). However, more recently have we become aware of low-level RFI that persists in EoR data after excision with traditional routines. Promising steps have been made to isolate and remove this interference, and work in this field is ongoing (Wilensky et al., 2019). Interestingly, the excision of RFI infested frequency channels creates gaps in the data along the frequency direction. Spectral gaps in the data cause contamination in the power spectrum if not accounted for when estimating the Fourier transform along the frequency-direction (Offringa et al., 2019a).

4.5.3 Foregrounds

We have already discussed how, in particular, unsubtracted point source foregrounds and mode-mixing lead to the rise of the wedge in power spectrum space. Although dominated by sources in the main lobe of the primary beam, sources in the side lobes cause additional contamination. Additionally, if foreground sources have intrinsic spectral and spatial structure their contamination can extend well beyond the wedge depending on the particular spatial structure of the foreground emission. Galactic emission and radio galaxies with large angular structure are a source of contamination that extend the point source-driven wedge to larger spatial scales (smaller k_{\perp}) comparable to the EoR (Procopio et al., 2017).

There are various strategies to mitigate foreground contamination that can broadly be categorised as; avoidance, subtraction, and suppression. The delay-based technique is an example of avoidance as it naturally avoids the wedge itself, and enables the choice of particular delay modes that are devoid of foreground contamination. Similarly choosing k -modes beyond the foreground wedge can be classified as such avoidance strategy. Subtraction involves explicit removal

of foreground signals from the data through modeled visibility contributions of foreground sources or non-parametrised techniques that subtract out foreground contamination (Chapman et al., 2012, 2013), or more recently Gaussian Process Regression (Mertens et al., 2018). Often a combination of subtraction and avoidance is used (Kerrigan et al., 2018; Jacobs et al., 2016; Barry et al., 2019a). Finally, suppression techniques are naturally born in studies of optimal power spectrum estimation, and rely on a known covariance matrix that describes the contamination in different \mathbf{k} -modes, i.e. the foreground covariance matrix \mathbf{C}_{FG} in Equation (4.8). This covariance matrix can be used to down weight contaminated k -modes (Liu et al., 2014; Dillon et al., 2015; Trott et al., 2016) in the averaging over shells of $|k|$. However, down weighting measurements based on either empirically estimated or analytically modelled covariance is prone to signal loss, due to incorrect suppression of EoR modes (Cheng et al., 2018). Small errors in the modelled foreground covariance matrix can have large impacts in the estimated EoR-signal due to the large dynamic range involved, and can lead to biased estimates of the signal as demonstrated by the retracted limits from the PAPER-64 experiment (Ali et al., 2018).

4.5.4 Calibration

Finally, one of the primary concerns in EoR experiments, and this thesis, is the impact of calibration on the power spectrum. It was already mused by Morales et al. (2012) that calibration solutions that contain spectral structure have the potential to add additional contamination into the EoR window. This effect was clearly demonstrated by Barry et al. (2016); Ewall-Wice et al. (2017) and Patil et al. (2016) in the context of sky modelling errors. Incomplete source models, and beam models leave residuals on the longest baselines that through calibration contaminate the shorter baselines (horizontal contamination). I revisit this issue in Chapters 6 and 7.

Furthermore, ignoring the shorter baselines during calibration because they

lack sufficiently accurate models that capture the diffuse sky, increases the noise on these baselines when calibration solutions are applied (Patil et al., 2016). Hence, a trade off must be made between the long and the short baselines that are optimal for accurate calibration. Barry et al. (2016) highlighted the importance of spectrally smooth solutions and as such the `FHD/epsilon` (Sullivan et al., 2012; Barry et al., 2019a) pipeline uses low-order polynomial smoothing after per channel calibration. `SAGEcal` on the other hand uses a regulariser to ensure spectral smoothness while calibrating all frequency channels simultaneously, i.e. this regularisation penalises solutions that are non-smooth (Yatawatta, 2015). The impact of calibration errors in redundant calibration are less well characterised. However, we suspect contamination to arise due to inherent non-redundancies in radio interferometers. As discussed in Chapter 3, redundant calibration consists of two steps; relative and absolute calibration. The relative step relies on redundancy and deviations from redundancy could leave an imprint on calibration solutions during the relative calibration step.

In Chapter 5 I study the impact of position errors on relative calibration. Orosz et al. (2019) studies the impact of non-redundancies during relative calibration on the power spectrum, and find indeed that similar to sky-based calibration the deviations from redundancy on the longer baselines introduces contamination on the shorter baselines that are sensitive to an EoR signal. Moreover, sky-based calibration errors creep into the absolute calibration step, which imparts additional spectral structure on the absolute calibration parameters (Byrne et al., 2019). We revisit these issues in Chapter 7, where I discuss the development an analytic model to understand the propagation of these errors into the 21-cm PS.

4.6 Power Spectrum Limits

The complications we discussed prevent us from actually integrating the estimated ~ 1000 hours of data required for a detection of the 21-cm EoR PS. Integrating such a large data set is only useful if the limiting factor is thermal

noise, because that averages down with $\propto 1/t$ in power. However, if we are dominated systematically by the discussed contaminants, integration of a large data set will not yield improved results.

The earliest limits from the GMRT (Paciga et al., 2013), PAPER (Ali et al., 2018), MWA (Dillon et al., 2015; Beardsley et al., 2016) and LOFAR (Patil et al., 2017) only used small amount of data ~ 10 hours resulting in limits on the dimensionless power spectrum between $\Delta_{21}^2 < 10^3 - 10^4 \text{ mK}^2$. A significant effort has been put in analysis, foreground modelling, and RFI mitigation. This resulted into a new suite of upper limits from PAPER (Kolopanis et al., 2019), the MWA (Barry et al., 2019b; Li et al., 2019; Trott et al., 2020), and LOFAR (Mertens et al., 2020). Promisingly, most experiments are now almost reaching $\Delta_{21}^2 < 10^3 \text{ mK}^2$ (see Figure 4.3). However, particularly interesting is the fact that work containing only a few tens of hours (Barry et al., 2019b; Li et al., 2019) are equally competitive with results that have just over a hundred hours (Trott et al., 2020; Mertens et al., 2020). This indicates that there is still quite some improvement to gain from understanding systematic contamination that inhibits the experiment and necessitates the work on calibration systematics presented in the next chapters.

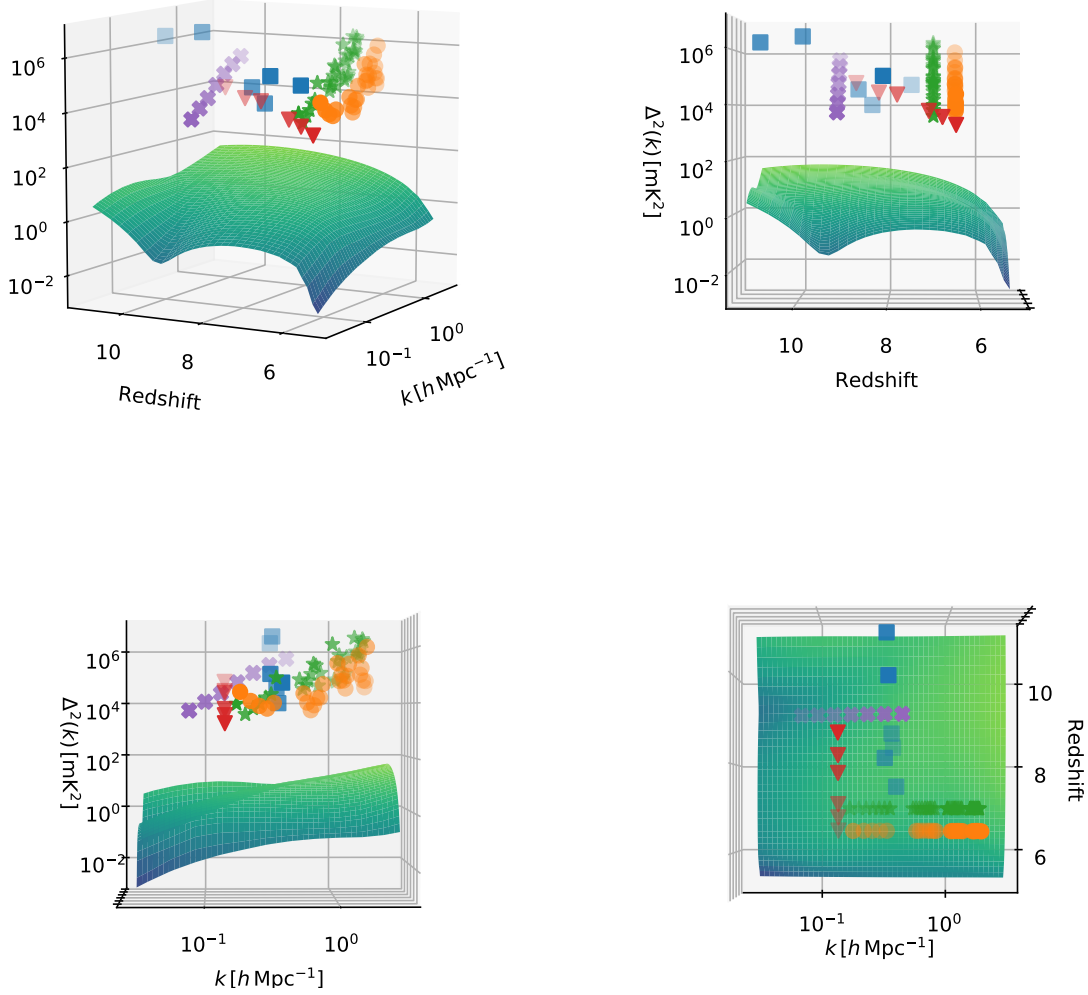


Figure 4.3: The three dimensional space probed by interferometric reionisation experiments. Top left: the dimensionless power spectrum versus redshift, top right: the dimensionless power spectrum versus Fourier scales k , bottom left: a top view of redshift versus fourier modes. The surface is interpolation of a theoretical model of the EoR PS for a faint galaxy driven reionisation process (Mesinger et al., 2016). The blue square are the latest PAPER-64 results (Kolopanis et al., 2019), the orange circles are results from MWA Phase II-compact (Li et al., 2019), the green stars are re-analysed data from MWA Phase I (Barry et al., 2019b), the red triangles are an alternative analysis of MWA data (Trott et al., 2020), and the purple crosses are the latest LOFAR results (Mertens et al., 2020).

The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions

This chapter is a reproduction of Joseph, R. C., Trott, C. M. and Wayth, R.B., “The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions”, 2018, *The Astronomical Journal*, 156, 285, <https://doi.org/10.3847/1538-3881/aaec0b>. It differs from the original in the omission of introductory material that was discussed in Chapter 2, and minor alterations, e.g. referencing to earlier chapters to ensure consistency within this thesis.

Abstract

The advent of a new generation of low-frequency interferometers has opened a direct window into the Epoch of Reionization (EoR). However, key to a detection of the faint 21 cm signal, and reaching the sensitivity limits of these arrays, is a detailed understanding of the instruments and their calibration. In this chapter, we use simulations to investigate the bias and uncertainty of redundancy-based

calibration. Specifically, we study the influence of the flux distribution of the radio sky and the impact of antenna position offsets on the complex calibration solutions. We find that the position offsets introduce a bias into phase component of the calibration solutions. This phase bias increases with the distance between bright radio sources and the pointing center, and with the flux density of these sources. This is potentially problematic for redundant calibration on Murchison Widefield Array (MWA) observations of EoR fields 1 and 2. EoR field 0, however, lacks such sources. We also compared the simulations with theoretical estimates for the bias and uncertainty in sky-model-based calibration on incomplete sky models for the redundant antenna tiles in the MWA. Our results indicate that redundant calibration outperforms sky-based calibration due to the high positional precision of the MWA antenna tiles.

5.1 Introduction

Over the past few years the latest generation of low-frequency interferometers has pushed down the upper limits of the 21-cm power spectrum of the Epoch of Reionisation (EoR) (Beardsley et al., 2016; Patil et al., 2017). However, none of the current instruments, e.g. the Murchison Widefield Array (MWA) (Tingay et al., 2013), the LOw Frequency ARray (LOFAR) (van Haarlem et al., 2013), and the Precision Array for Probing the Epoch of Reionization (PAPER) (Parsons et al., 2010), have detected a signal thus far. The signal, emitted by neutral hydrogen during the EoR, is a direct probe into the state of the Intergalactic Medium (IGM) (Furlanetto et al., 2006) and allows us to directly study the conditions under which the first luminous objects were formed. For more in depth reviews see Morales & Wyithe (2010); Pritchard & Loeb (2012); McQuinn (2016); Furlanetto (2016).

Foreground sources and instrumental effects pose large challenges to the detection of this faint signal. The low frequency foregrounds, e.g. the Milky Way and extragalactic sources, are expected to be 4-5 orders of magnitude stronger

than the neutral hydrogen signal (Furlanetto et al., 2006; Bowman et al., 2006; Morales et al., 2006; Pritchard & Loeb, 2008; Jelic et al., 2008). The removal of these foregrounds requires a detailed understanding of the instrument and its calibration, because the subtraction of these foregrounds, in particular bright compact sources, is sensitive to calibration errors (Datta et al., 2009). Failing to remove bright sources accurately leads to contamination of EoR data, causing the so-called “wedge” feature in the 2D-Power Spectrum. (Datta et al., 2010; Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012) This leakage of bright source residuals into the power spectrum, makes certain scales of the EoR signal inaccessible if not dealt with correctly.

Adequate removal of these foregrounds and extraction of the faint signal from the data puts stringent requirements on our calibration accuracy and precision. Standard calibration schemes correct the sky signal cross-correlations (or ‘visibilities’) measured by radio interferometers using sky models, hereby solving for the gain factors that cause the discrepancy between the modelled visibilities and the measured visibilities (see Rau et al., 2009, for a review). Sky model based calibration has undergone tremendous progress in the past years in order to overcome direction dependent calibration effects, e.g. varying antenna primary beam shapes, and ionospheric distortions, that limit this new generation of instruments. This progress resulted in a large variety of improved sky based calibration implementations, e.g. **RTS** (Mitchell et al., 2008), **SAGEcal** (Yatawatta et al., 2009; Kazemi et al., 2011), **SPAM** (Intema et al., 2009), **FHD** (Sullivan et al., 2012), and facet calibration (van Weeren et al., 2016) to name a few. Nevertheless, at the operating frequency of these new low-frequency interferometers (80–200 MHz), our limited understanding of the sky leads to incomplete models. Model-based calibration with incomplete sky models causes calibration errors that lead to image artifacts, which in turn limit the dynamic range of observations (Grobler et al., 2014; Wijnholds et al., 2016) and, more relevant to EoR science, contaminate the power spectrum (Barry et al., 2016; Ewall-Wice et al., 2017; Trott &

Wayth, 2017). Redundant calibration, however, allows us to escape our ignorance of the low frequency sky because it does not require modelling (Wieringa, 1992). Because of this reason redundant calibration is undergoing a renaissance, resulting in further studies by Noorishad et al. (2012); Liu et al. (2010); Ali et al. (2015); Dillon & Parsons (2016), showing the applicability and some limitations of redundant calibration in low frequency radio telescopes. More recently, redundant calibration was compared to sky model based calibration by Li et al. (2018). Despite the inability of redundant calibration to solve for direction dependent effects, it still remains an interesting alternative to calibrate a radio telescope to first order, where sky based calibration can resolve higher order effects.

In this chapter we study the theoretical performance of redundant calibration. We specifically look at how redundant calibration depends on the flux distribution of the sky and positional errors of the antennas. This allows us to determine which regions of the sky should be calibrated with sky-based calibration or redundant calibration, to yield the most accurate and precise result for a given antenna position precision of the array. We do this by running simulations of redundant calibration in which we calibrate a redundant array with ideal antenna responses on a realistic multi-source sky, while changing the flux and position of a single calibrator source. We compare the distribution of solutions we obtain from these simulations with a theoretical estimate of the sky model calibration bias and an uncertainty due to an incomplete calibration model.

The structure of the chapter is as follows: Section 5.2 discusses sky model calibration and our analytic description of the impact of an incomplete sky model on the bias and uncertainty of the calibration solutions. Section 5.3 describes the set up of the redundant calibration simulations. We discuss the influence of the sky flux distribution on redundant calibration solutions and the impact of position offsets, using a simple 5-element interferometer to demonstrate the fundamental issues of redundant calibration. We conclude our results with a comparison between the bias and uncertainty of redundant calibration, and sky

model based calibration for the redundant MWA tiles in Section 5.4, and we discuss the implications for the MWA in Section 5.5.

5.2 Sky Model Calibration

In this section we describe and derive the impact of an incomplete sky model on the calibration solutions in a sky model based approach. Earlier works studied the effect of calibration on incomplete sky models via analysis and simulations. Salvini & Wijnholds (2014) discuss the statistical performance, Barry et al. (2016) study the impact on EoR power spectrum estimation, and Grobler et al. (2014); Wijnholds et al. (2016) study its impact on the deconvolution of a 2-point source sky. In this work we compare redundant calibration with theoretical estimates for the bias and uncertainty introduced by calibration on an incomplete sky model.

We can write the measured correlation V_{pq}^{data} for a pair of antennas p and q in the absence of noise as a product of the antenna gain factors g_p^* and g_q , and the true visibility V_{pq}^{True}

$$V_{pq}^{\text{data}} = g_p g_q^* V_{pq}^{\text{True}}, \quad (5.1)$$

the superscript ‘*’ indicates complex conjugation. In sky model based calibration we solve for the gains g_i by minimizing the difference between our modelled visibilities m_{ij} and the measured correlations V_{ij} .

$$\min_{\mathbf{g}} \sum_{ij} \|V_{ij} - g_i g_j^* m_{ij}\| \quad (5.2)$$

Here, we write the minimization in the most general form, without explicitly choosing a matrix or vector notation for discussion purposes, as we will switch between those later on. The caveat of this calibration approach is that the signals from unmodelled sources are absorbed into the calibration solutions. To understand how this impacts the solutions, we first derive the uncertainty of sky based calibration solutions due to a stochastic sky of point sources and thermal noise.

We then use this result to derive the bias due to model incompleteness.

5.2.1 Model Incompleteness Uncertainty

To derive the minimum uncertainty on the estimated complex gain solutions g we use the Cramér-Rao Lower Bound (CRLB) on the estimated gain parameters. Throughout this derivation we assume the model used for calibration is a single point source with flux density $S(\mathbf{s}, \nu)$ located at some location \mathbf{s} in the sky. The model visibility for a given baseline \mathbf{b} at frequency ν is then given by:

$$M(\mathbf{b}, \nu) = S(\mathbf{s}, \nu)A(\mathbf{s}) \exp[-2\pi i \mathbf{b} \cdot \mathbf{s}], \quad (5.3)$$

$A(\mathbf{s})$ is the antenna beam response, which we choose to be a Gaussian. We choose an unmodelled source background described by a broken power-law source count distribution dN/dS (Gervasi et al., 2008; Intema et al., 2011; Franzen et al., 2016; Williams et al., 2016)

$$\frac{dN}{dS} = \begin{cases} k_1 S^{-\gamma_1} & \text{if } S_{\text{low}} \leq S < S_{\text{mid}} \\ k_2 S^{-\gamma_2} & \text{if } S_{\text{mid}} \leq S < S_{\text{high}} \end{cases}, \quad (5.4)$$

where dN/dS gives the number of sources per Jy per steradian, and S is the source flux in Jy. Throughout this thesis we use $k_1 = k_2 = 4100$, $\gamma_1 = 1.59$, $\gamma_2 = 2.5$, $S_{\text{low}} = 400$ mJy, $S_{\text{mid}} = 1$ Jy, and $S_{\text{high}} = 5$ Jy. To derive the CRLB on the estimated gain parameters we first compute the Fisher Information Matrix (FIM) \mathcal{I} (Kay, 1993). This takes the following form for a complex multivariate normal distribution with mean \mathbf{m} and gain independent data covariance \mathbf{C}_{data} ;

$$\mathcal{I}_{i,j} = 2\text{Re} \left(\frac{\partial \mathbf{m}^\dagger}{\partial g_i} \mathbf{C}_{\text{data}}^{-1} \frac{\partial \mathbf{m}}{\partial g_j} \right), \quad (5.5)$$

where the superscript ‘ \dagger ’ denotes the Hermitian transpose, the superscript ‘ $^{-1}$ ’ denotes the matrix inverse, \mathbf{m} is a vector where each entry is the mean measured

visibility of a baseline pair pq for a single frequency channel, and \mathbf{C}_{data} is the data covariance matrix. We assume the mean visibility \mathbf{m} is given by the model visibility. The covariance of the data is the sum of thermal noise covariance $\mathbf{C}_{\text{thermal}}$ and the covariance of our stochastic background sky \mathbf{C}_{sky} . For the thermal noise we assume its baseline independent and for the sky covariance we ignore the compact Fourier beam kernel that creates correlations between closely-spaced baselines. We describe the variance of the thermal noise as

$$\Sigma_{\text{thermal}} = \frac{SEFD^2}{Bt}, \quad (5.6)$$

where B is the bandwidth of a single frequency channel, and t is the integration time of the observation. Throughout this thesis we adopt MWA EoR parameters unless stated otherwise, i.e. $SEFD = 10^4$ Jy, $B = 40$ kHz, $t = 120$ s. For these parameters the thermal noise is ~ 9 Jy. We take the expression for the visibility variance for a baseline in a single frequency channel due to a stochastic sky Σ_{sky} , from Trott et al. (2016); Murray et al. (2017)

$$\Sigma_{\text{sky}} = 2\pi \frac{\sigma^2}{2} \left(\frac{k_1}{3 - \gamma_1} [S_{\text{mid}}^{3-\gamma_1} - S_{\text{low}}^{3-\gamma_1}] + \frac{k_2}{3 - \gamma_2} [S_{\text{high}}^{3-\gamma_2} - S_{\text{mid}}^{3-\gamma_2}] \right), \quad (5.7)$$

wherein we assume a flat spectral index of our sources within a single frequency channel, and σ is the frequency dependent beam width. Throughout this thesis we assume a beam width of $\sigma = 30^\circ$, similar to the MWA beam at 150 MHz, resulting in a sky variance of $\Sigma_{\text{sky}} \sim 2.5 \cdot 10^3$ Jy². Because the noise variance Σ_{thermal} and sky variance Σ_{sky} are baseline independent, the total data covariance matrix Σ_{data} is a diagonal matrix. We can therefore rewrite the FIM elements as;

$$\mathcal{I}_{i,j} = 2\text{Re} \left(\sum_n \frac{1}{\Sigma_{\text{data}}} \frac{\partial \mathbf{m}_n^\dagger}{\partial g_i} \frac{\partial \mathbf{m}_n}{\partial g_j} \right), \quad (5.8)$$

where we sum over the data index n . For the CRLB we are only interested in the variance on a gain parameter Σ_g , i.e. we only compute $\mathcal{I}_{m,n}^{-1}$ for $m = n$, which

reduces to

$$\Sigma_g = \frac{\Sigma_{thermal} + \Sigma_{sky}}{2[S(\mathbf{s}, \nu)A(\mathbf{s})]^2(N-1)}, \quad (5.9)$$

where $N - 1$ is the number of baselines formed by an antenna in the array. We note that the variance scales inversely with the number of antennas in the array, and beam-weighted apparent flux density of the modelled source squared. We will use this expression to compare the uncertainty of redundant calibration with sky model based calibration.

5.2.2 Model Incompleteness Bias

To derive an expression for the bias, i.e. the mean deviation from the true solutions introduced by the model incompleteness, we follow Wijnholds et al. (2016) and reformulate Equation (5.2) explicitly in terms of visibility matrices and gain vectors;

$$\min_{\mathbf{g}} \|\mathbf{V} - \mathbf{G}\mathbf{M}\mathbf{G}^\dagger\|. \quad (5.10)$$

\mathbf{V} and \mathbf{M} are matrices containing the measured and modelled visibilities, e.g. \mathbf{V}_{pq} is the measured visibility between antenna p and q , and the matrix $\mathbf{G} = \text{diag}(\mathbf{g})$ contains the complex antenna gains. We ignore the auto-correlations, therefore, the diagonals of \mathbf{M} and \mathbf{V} are zero, and if we ignore the noise we can write the measurements \mathbf{V} in terms of the modelled \mathbf{M} and unmodelled \mathbf{U} sky visibilities. We can also write the gain vector \mathbf{g} as a sum of the true gains \mathbf{g}_t and a deviation introduced by the calibration process $\Delta\mathbf{g}$.

$$\begin{aligned} \mathbf{V} &= \mathbf{G}_t(\mathbf{M} + \mathbf{U})\mathbf{G}_t^\dagger \\ \mathbf{g} &= \mathbf{g}_t + \Delta\mathbf{g} \end{aligned} \quad (5.11)$$

Furthermore we can use the *Hadamard product* \odot , i.e. the element-wise product, to rewrite Equation (5.10) into

$$\min_{\Delta\mathbf{g}} \|\mathbf{g}_t\mathbf{g}_t^\dagger \odot \mathbf{U} - (\mathbf{g}_t\Delta\mathbf{g}^\dagger + \Delta\mathbf{g}\mathbf{g}_t^\dagger) \odot \mathbf{M}\| \quad (5.12)$$

where we have dropped all higher order terms of $\Delta\mathbf{g}$. Wijnholds et al. (2016) derive an approximate closed form solution for $\Delta\mathbf{g}$ by rewriting Equation (5.12) into a least squares form. We will take the solution as the conclusion of this short review, and point the interested reader to their work for the detailed derivation. The closed form solution takes the following form

$$\begin{bmatrix} \Delta\mathbf{g} \\ \Delta\mathbf{g}^* \end{bmatrix} \approx \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{E}\mathbf{g}_t \\ \mathbf{F}\mathbf{g}_t \end{bmatrix}. \quad (5.13)$$

The block matrices are given by

$$\begin{aligned} \mathbf{A} &= \mathbf{M}^*\mathbf{G}_t\mathbf{G}_t^\dagger\mathbf{M}^* \odot \mathbf{I} & \mathbf{B} &= \mathbf{M}^*\mathbf{G}_t \odot \mathbf{G}_t\mathbf{M} \\ \mathbf{C} &= \mathbf{B}^* & \mathbf{D} &= \mathbf{A}^* \\ \mathbf{E} &= \mathbf{M}^*\mathbf{G}_t\mathbf{G}_t^\dagger\mathbf{U}^* \odot \mathbf{I} & \mathbf{F} &= \mathbf{G}_t^\dagger\mathbf{U}^* \odot \mathbf{M}\mathbf{G}_t^\dagger \end{aligned}, \quad (5.14)$$

where $\mathbf{G}_t = \text{diag}(\mathbf{g}_t)$, and \mathbf{I} is the identity matrix. Here, \mathbf{A} and \mathbf{B} encode the total modelled power summed over baselines, and the power in an individual baseline, respectively, whereas \mathbf{E} and \mathbf{F} are the equivalent expressions for the unmodelled power. Intuitively, these matrices describe the additional bias in the solutions from correlations between the model and the residual signal, and the overall power ratio of model to unmodelled sky. Minimising both of these bias terms is desirable for good sky-based calibration. We can use Equation (5.13) to derive the mean gain offset $\langle\Delta\mathbf{g}\rangle$ in the case that our sky model consists of a single point source in the presence of a more complicated sky.

$$\left\langle \begin{bmatrix} \Delta\mathbf{g} \\ \Delta\mathbf{g}^* \end{bmatrix} \right\rangle \approx \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} \begin{bmatrix} \langle\mathbf{E}\rangle\mathbf{g}_t \\ \langle\mathbf{F}\rangle\mathbf{g}_t \end{bmatrix}, \quad (5.15)$$

with

$$\begin{aligned} \langle\mathbf{E}\rangle &= \mathbf{M}^*\mathbf{G}_t\mathbf{G}_t^\dagger\langle\mathbf{U}\rangle^* \odot \mathbf{I} \\ \langle\mathbf{F}\rangle &= \mathbf{G}_t^\dagger\langle\mathbf{U}\rangle^* \odot \mathbf{M}\mathbf{G}_t^\dagger \end{aligned}. \quad (5.16)$$

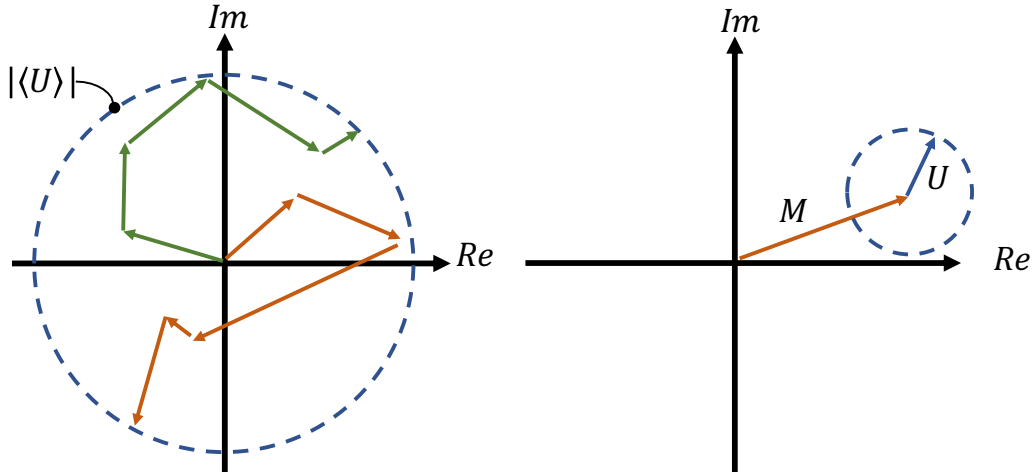


Figure 5.1: *On the left:* an illustration of the random walk through the complex plane due to the addition of different source visibilities of a two stochastic sky realisations. *On the right:* The orientation between the modelled visibility \mathbf{M} and the unmodelled visibility contribution \mathbf{U} .

We can parametrize the mean unmodelled visibility contribution of our stochastic sky $\langle \mathbf{V}_u \rangle$ using the sky visibility variance Σ_{sky} . If we consider the Fourier transform of each point source as a phasor in the complex plane (Re, Im), we can consider a stochastic sky of point sources as a random walk through this plane (see Figure 5.1). Each point source contributes a new complex phasor to our total unmodelled sky phasor. The path length of this random walk, i.e. the total amplitude of our unmodelled visibility, is on average given by the root mean square of distribution from which the phasors are drawn. In our analysis we assume this to be a Gaussian distribution, therefore, the unmodelled visibility amplitude equates to the variance.

Now, we have yet to explore the orientation of the unmodelled visibility component. Every net orientation has equal probability. For our calculations we choose the unmodelled visibility to have an angle of 45° , with respect to the model visibility as a measure for some average offset introduced into our visibility amplitude and phase angle. Even though the true phase angle of the unmodelled visibility is uniformly distributed, we find that this approximation yields comparable results to Monte Carlo simulations with a distribution of phases. We will

use these results in Section 5.4 where we compare the results from our redundant calibration simulations with the theoretical performance of sky model based calibration.

5.3 Redundant Baseline Calibration

Chapter 3 provides a review of Wieringa (1992) and Liu et al. (2010) to highlight the key features of redundant calibration, in particular the difference between `logcal` and `lincal`. Current implementations of redundant calibration use `logcal` to find an initial estimate and further refine the solutions with `lincal` (Zheng et al., 2014). In this work we will do the same.

5.3.1 Simulating the Bias and Uncertainty

To estimate the bias and uncertainty of redundant calibration we simulate the calibration of the antennas in a (nearly) redundant array. In our simulations we define a group of redundant baselines when they lie within $1/6\lambda$ of each other in the uv -plane. This is well within the linear regime of a sinusoid centred at $\exp[2\pi ul]$. However, we will show that deviations from non-redundancy within this threshold impact the calibration accuracy and precision. We also assume a Gaussian beam, similar to our sky model derivation, and assume the beams are identical for each antenna. This is not strictly true for phased arrays (Wijnholds et al., 2010), however, it suffices as a first-order approximation.

We generate a background of radio sources with a flux distribution according to Equation 5.4, and uniformly distribute them over the sky. Finally, we add a source with arbitrary flux and location, similar to the calibrator source in Section 5.2. These sources are gridded onto an (l, m) -grid, and Fourier transformed to generate visibilities using `powerbox` (Murray & Trott, 2018), a tool written to simulate EoR datasets and forward-model them to power spectra. We interpolate the visibilities to produce the measurements for each baseline. Finally, we assume

Gaussian-distributed noise in the real and imaginary components with a variance according to Equation (5.6). These visibilities are then passed to our redundant calibration module, which is a direct implementation of the algorithm described in Section 5.3. The code is publicly available (Joseph et al., 2018).

5.3.2 The Sky Dependent Uncertainty

We first study the influence of the sky, and show how it affects the uncertainty of the estimated gain solutions. We start out with a simple sky model of a statistical background sky while moving a high flux density source with respect to phase centre and trying to calibrate on each realization of the sky. This allows us to study the performance of redundant calibration in drift-scan mode, and simultaneously study the performance of redundant calibration in the MWA EoR fields, which depending on the field have strong in-beam sources. Figure 5.2 shows the dependence of the calibration solutions for an ideal interferometer, i.e. perfect gains $g = 1$ and perfect redundancy, as a function of source position in terms of the direction cosine l , the native interferometry sky coordinate. We show the results for a pure `logcal`-calibration, similar to Wieringa (1992), and for a `logcal + lincal`-calibration.

The results show that overall the solution variance for both the gain amplitude solutions behave better when the strong source is near the centre of the beam, because the signal to noise ratio (S/N) is higher at the pointing centre. We do note that our implementation of the `lincal` algorithm seems slightly biased in the presence of noise, the mean of the solution is 1% below the true value. We filtered out $< 1\%$ of the solution realisations due to bad convergence, i.e. solutions with unrealistically high gain amplitudes.

The gain phase solutions show a similar dependence with some additional structure in the variance due to a problem which is inherent to `logcal`: phase wrapping. The `logcal` implementation can only determine phases between $-\pi \leq \phi \leq \pi$, in which the arctan is defined. When a certain redundant set of baselines




Figure 5.2: Calibration solution from relative calibration as a function of source position for a perfectly redundant array. *top left*: the `logcal` amplitude gain. *top right*: the `logcal + lincal` amplitude gain. *bottom left*: the `logcal` phase gain solutions. *bottom right*: `logcal + lincal` phase gain solutions for a single antenna in a nearly redundant 5 element interferometer as a function of strong source position l . The dark line represents the mean of the solutions, the shaded area indicates the 1-sigma solutions variance. The amplitude solution variance inversely follows the shape of the beam, i.e. the variance increases when the beam response decreases. The mean of the phase solutions fluctuates around $\phi = 0$ and jumps along with the variance at so-called phase wrapping points, which are further explained in the text.

measures a visibility phase of $|\phi_v| = \pi$, due to the location of the dominant source on the sky, the solutions become very sensitive to noise. The visibility phase starts to “jump” between $-\pi$ and π causing large variances in the phase calibration solutions. These phase wrapping points can be determined by solving $2\pi ul = \pi n$ for odd numbers of n , i.e. solving for the source coordinate l when a given baseline with length u measures a phase of π .

We can understand the effect of phase wrapping by adding a noise vector \mathbf{n}_α to the measurement equation, see equation (5.17).

$$\mathbf{c}_\alpha = \mathbf{A}_\alpha \mathbf{x}_\alpha + \mathbf{n}_\alpha \quad (5.17)$$

The phase noise $\mathbf{n}_\phi \propto N/S$ (Liu et al., 2010), however, when the measured visibility phase approaches π , this noise vector diverges $|\mathbf{n}_\phi| \rightarrow 2\pi$. Because the noise of a single baseline is mixed into all solutions when estimating $\hat{\mathbf{x}}$ we get large offsets in the calibration solutions:

$$\hat{\mathbf{x}} = \mathbf{x}_\alpha + [\mathbf{A}^T \mathbf{A}]^{-1} \mathbf{A}^T \mathbf{n}_\alpha. \quad (5.18)$$

Looking at the `logcal + lincal` solutions in Figure 5.2 we see similar behaviour in the phase solutions and amplitude solutions. We also note that our implementation of the algorithm seems to be very sensitive to noise. This results in large variations in the amplitude solutions when the dominant source moves away from phase center. Figure 5.2 shows the mean and the variance of the solutions, even though the solutions do not strictly follow a Gaussian distribution. Figure 5.3 shows the distribution of `logcal` solutions for antenna 2 at the first phase wrapping point. The distribution has 7 distinct peaks, each peak represents a combination of phase wrapping baselines. The first phase wrap occurs in the set of three long redundant baselines, and therefore there are at maximum $\sum_{k=1}^n n!/(k!(n-k)!)$ different combinations and solutions peaks possible. The actual spectrum depends on the array geometry, which is encapsulated in the

matrix **A**. These solutions are the starting point for the `lincal` algorithm, and our implementation of it is not able to recover the true solutions when given a bad starting point.




Figure 5.3: The distribution of gain phase solutions for antenna 2 at the first phase wrapping point $l \sim 0.1$. The discrete set of solutions peaks is caused by phase wrapping in the set of three long baselines. Each peak represents a combination of phase wrapping baselines as marked above the peak

We attempt to circumvent the phase wrapping of a specific baseline in a single channel by extending the `logcal` algorithm to incorporate the measurements of neighbouring frequency channels while assuming the gain solutions remain the same. However, due to the same mixing that takes place in the single frequency channel implementation, phase wrapping will still remain a problem unless a clever selection of frequencies is used to circumvent phase wrapping. We discuss this in Appendix A. We also note that current implementations do apply a pre-calibration step to unwrap the visibility phases by averaging over baselines within a redundant group (Zheng et al., 2014), or by using the products of visibilities to

construct a system equations to solve for the phases (Li et al., 2018).

5.3.3 The Position Offset Bias

In the previous section we described the results for an ideal radio interferometer. However, in reality all antennas will have slight position offsets from their perfectly redundant positions. To understand the impact of positions offset we simulate redundant calibration under the same conditions as before, but now we offsetting one antenna in the x-direction by $\delta x = 20$ cm. The results are shown in Figure 5.4.

We can clearly see that both the amplitude and phase solutions are affected by the position offset. We can understand the oscillatory behaviour of the amplitude solutions by returning to the complex plane. Imagine the complex visibility of the main calibrator as measured by a baseline as a phasor in this plane, we can think of the total sum of background sources as a similar phasor. Each redundant baseline should measure the same amplitude of the sum of these phasors. However, due to the non-redundancy introduced by position offsets the non-redundant pairs measure a different amplitude, this difference propagates through to the solutions. As the primary source moves across the sky its phasor will rotate in the complex plane, constructively and destructively interfering with the background visibility creating this oscillatory behaviour. These oscillations are dampened as the primary source becomes attenuated as it moves outside of the primary beam. The behaviour of the mean phase solutions can be explained in a similar fashion, using the phase of the phasors rather than the amplitude.

For low-N arrays similar to this 5 element toy model the position error propagates to all antennas solutions, due to the coupling of all gain solutions to the visibilities. However, when increasing the number of antennas in the array the coupling becomes weaker as the number of measurements increases.

$$\mathbf{c}_\alpha = \mathbf{A}_\alpha \mathbf{x}_\alpha + \mathbf{n}_\alpha + \mathbf{b}_\alpha \quad (5.19)$$




Figure 5.4: Calibration solution from relative calibration as a function of source position for a perturbed redundant array. *top left*: the `logcal` amplitude gain. *top right*: the `logcal +lincal` amplitude gain. *bottom left*: the `logcal` phase gain solutions. *bottom right*: `logcal +lincal` phase gain solutions for a single antenna in an nearly redundant 5 element interferometer as a function of strong source position l . The dark line represents the mean of the solutions, the shaded area indicates the 1-sigma solutions variance. The amplitude solution variance inversely follows the shape of the beam, i.e. the variance increases when the beam response decreases. The mean of the phase solutions generally fluctuates around $\phi = 0$ and deviates along with the variance at so-called phase wrapping points, which are further explained in the text. The additional position offset, which causes a phase offset from the ideal redundant phase, is absorbed into the solutions causing a slope.




Figure 5.5: The magnitude of the offset residual $\hat{\mathbf{b}}$ for an offset in a given antenna. The first column shows the residual when offsetting the first antenna, i.e. the reference antenna, the second column shows the result for the second antenna etc. We see that offsetting the 3rd antenna, e.g. the middle antenna, impacts only its calibration solution. Its solutions have the strongest constraints due to its baseline participation. It participated in all baselines groups and has a high participation number in each group.

Inverting this equation using the standard least square solution, and not taking into account this extra term, leaves us with an additional residual. We can calculate these residuals for different antenna offsets. Figure 5.5 shows the magnitude of the offset residuals $\hat{\mathbf{b}}$ in the phases of the estimated gain and visibilities when offsetting different antennas by the same amount. These results show that offsetting the antenna with the highest baseline participation does not propagate to all antenna solutions and leaves the visibilities unaffected. Offsetting the reference antenna, in this case antenna 1, has the strongest impact on the solutions of all other antennas. This implies that the choice of reference antenna is not as arbitrary as one might think.

5.4 Comparing Sky and Redundancy Based Calibration

Having varied several parameters within our redundant calibration simulations we can now move forward and apply this formalism to the MWA ‘hexes’. The

hexes contain 72 antenna tiles arranged in 2 hexagons, see Figure 5.6. The shortest baseline, defining the hexagonal lattice, has a length of 14 metres. The hexagons are also placed to be redundant with each other, i.e. they have the same orientation. This provides extra sensitivity on scales relevant for the EoR experiment, and adds redundancy for calibration purposes. Due to a lack of redundant baselines connecting one hex to the reference antenna in the other hex, we either need to invoke another degeneracy parameter to encapsulate a phase offset between the two or calibrate them separately. Li et al. (2018) calibrate the hexes simultaneously, however, for simplicity and speed we calibrate a single hex in our simulations. A single hex forms 630 baselines, of which 601 are redundant, organized in 71 redundant groups¹. Figure 5.6 also shows the non-redundancy of each antenna within the hexes, the antennas are placed with an accuracy on the order of centimetres. This is an order of magnitude below the redundant calibration threshold of $1/6\lambda$ at 150 MHz, i.e. ~ 30 cm.

In our final set of redundant calibration simulations we will offset each tile in a single hex according to Gaussian distributions with mean $\mu_x = 0$, and $\sigma_x = 10^{-4} - 10$ cm. Now we calibrate while moving a 100 Jy source across the sky, and run a separate set of simulations where we fix the location of the source at 3° off-zenith while changing its flux density instead. Figures 5.7 and 5.8 shows the bias and uncertainty we derive from these simulations. To compute these from the distribution of solutions we obtain we calculate the median offset from the true gain solutions for the bias, i.e. $g = 1$ and we take the standard deviation for the uncertainty. All results are averaged over all antennas.

To compare with traditional sky-based calibration we use the bias and uncertainty derived in Section 5.2. The contour lines in Figures 5.7 and 5.8 are the ratios between either the bias or uncertainty of redundant calibration and sky model calibration. To make the comparison slightly easier we take two cuts

¹Theoretically we can also include non-redundant antennas in the calibration, as long as the number of unknowns is lower than the number of measurements. For each redundant hex in the MWA we can add 6 non-redundant tiles before the system becomes unsolvable.

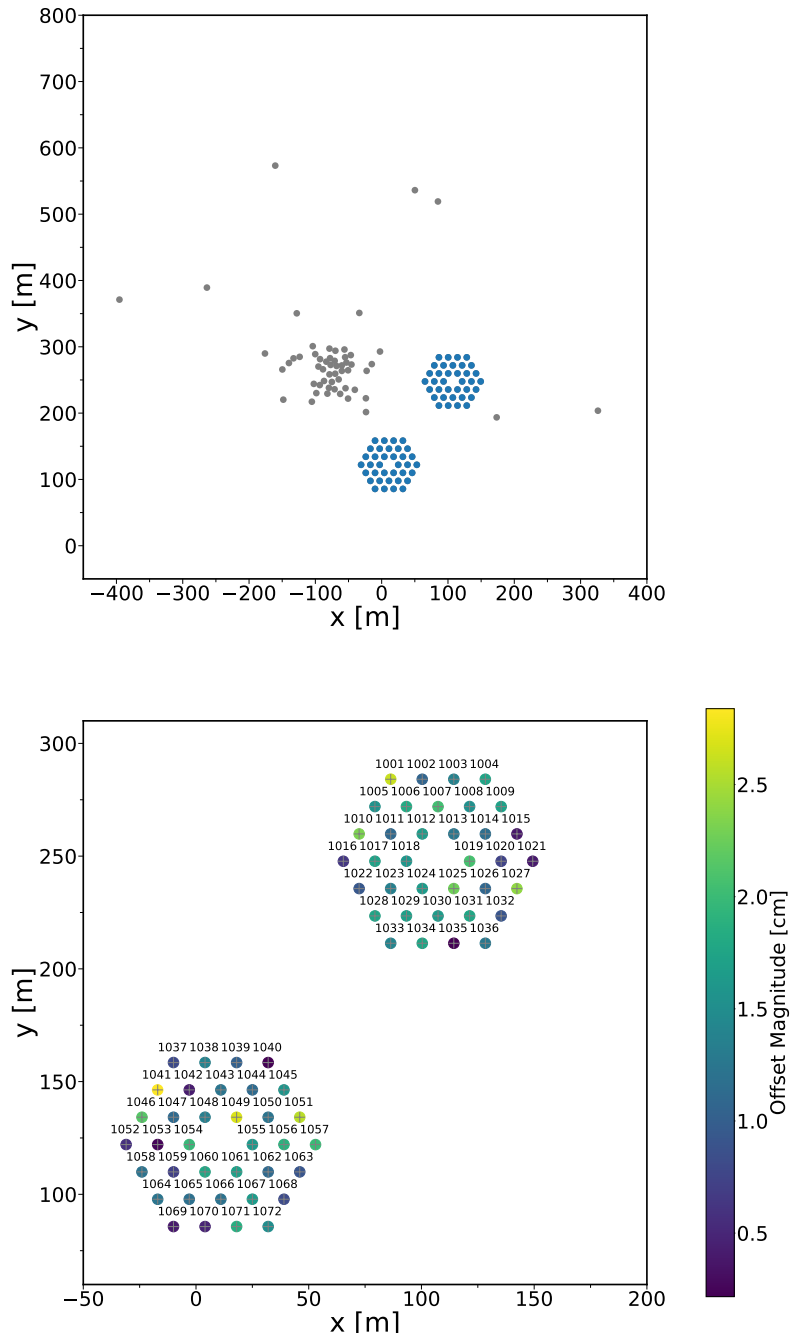


Figure 5.6: *top*: The MWA antenna lay-out, with MWA hex antennas in blue and the random antennas in black. *bottom*: a zoom-in plot on the MWA hex tiles. The colour of each tile indicates its offset from its redundant position. We see that the position deviations from redundancy do not exceed 3 cm.

through the plots in Figures 5.7 and 5.8 at an antenna position precision of $\sigma_x = 0.02$ m and $\sigma_x = 0.10$ m, these cuts are shown in Figures 5.9 and 5.10. From Figures 5.7, 5.8, 5.9 and 5.10 we can conclude the following statements. The amplitude bias depends strongly on the flux density of the primary source and its location on the sky. Redundancy-based calibration has a lower amplitude bias when the sky is dominated by a single point source, and it quickly reaches the accuracy of our implementation of the algorithm as the source moves out of the field of view.

The amplitude uncertainties of redundancy-based and sky-based calibration are comparable. For both redundant and sky model based calibration, they decrease comparably as a function of primary source flux density. However, as a function of source location redundant calibration quickly reaches the noise floor when the source is beyond the FWHM of the primary beam.

Interestingly, the phase bias increases with primary source flux density and distance of the bright source to phase centre. The bias reaches a maximum when the source is at the FWHM of the primary beam, if the source moves beyond this the bias decreases again. When the sky is dominated by a single source that is off-centre, the bias of redundant calibration becomes comparable or larger to that of sky-based calibration. The flux at which the two become comparable is dependent on the magnitude of the position offsets in the array.

The phase uncertainty depends strongly on the source flux density and its location on the sky. When the primary source brightness becomes comparable to the background sky the uncertainties of redundant become larger than that of sky model based calibration (see Figure 5.7). Also note that the behaviour of the uncertainty as a function of source elevation changes for different positional precisions. When the primary source is between phase centre and the FWHM of the primary the uncertainties of redundant calibration are larger for arrays with large positional offsets. Particularly when the brightest source is at the FWHM of the primary, the uncertainties become larger than that of sky model

based calibration (see Figure 5.8). This increase in uncertainty can be explained by phase wrapping as discussed earlier. As the sources moves the measured visibilities phases start wrapping around 2π creating a spectrum of solutions that widens the distribution.

5.5 Discussion

The most notable results from our simulations are the results for the phase bias. Redundant calibration was proposed as an alternative to a sky model based approach because it is agnostic of the sky and therefore it does not suffer from the systematics introduced by an incomplete sky model. However, this work shows that systematics arise in a different way, because we impose the condition that our telescope is perfectly redundant. This manifests itself in systematic phase offsets in our calibration solutions because redundant calibration absorbs antenna position offsets into the calibration solutions. These phase offsets become more prominent when there is a high flux density source away from the pointing center.

Barry et al. (2016); Ewall-Wice et al. (2017); Trott & Wayth (2017) show that calibration on incomplete sky models causes contamination in the EoR power spectrum. Similarly redundant calibration can introduce contamination. The relative position offset changes as a function of wavelength. Therefore, the measured phase offset will therefore also vary as a function of frequency that can introduce a contamination to the EoR power spectrum.

We demonstrated the influence of the sky flux distribution on the performance of redundant calibration. Figure 5.11 shows a map of the radio sky at 408 MHz (Haslam et al., 1982) with the MWA EoR target fields. We can clearly see that these fields are not devoid of high flux density sources. EoR field 1 contains Fornax A and Pictor A, and EoR field 2 contains Hydra A amongst others. However, the results of the redundant calibration simulations for a single MWA hex show that position offsets at the position precision levels of the MWA are not a large source of bias and uncertainty for the phase solutions. Redundant

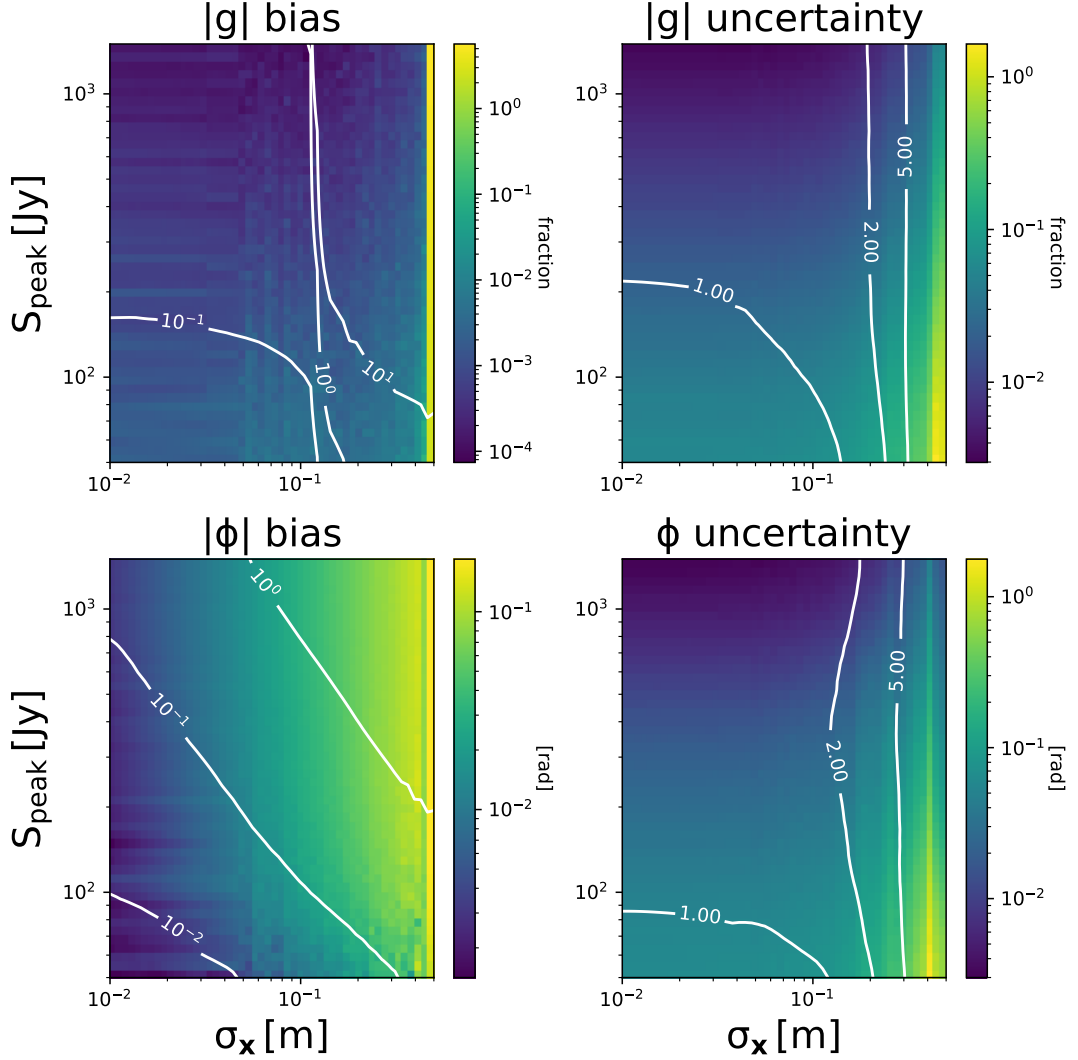


Figure 5.7: The bias (*top left*) and uncertainty (*top right*) of the amplitude solutions $|g|$, and the bias (*bottom left*) and uncertainty (*bottom right*) of the phase solutions $|\phi|$ of redundant calibration on an MWA hex-like array. We vary the position deviations σ_x and source peak flux S_{peak} of a source located 87° above the horizon. All results are averaged over all antennas in the hex. The contour lines are the ratios between the bias and uncertainty of redundancy based and sky model based calibration, e.g. an uncertainty contour line of 2.0 indicates the uncertainty is twice as large for redundant calibration as compared to sky model based calibration.

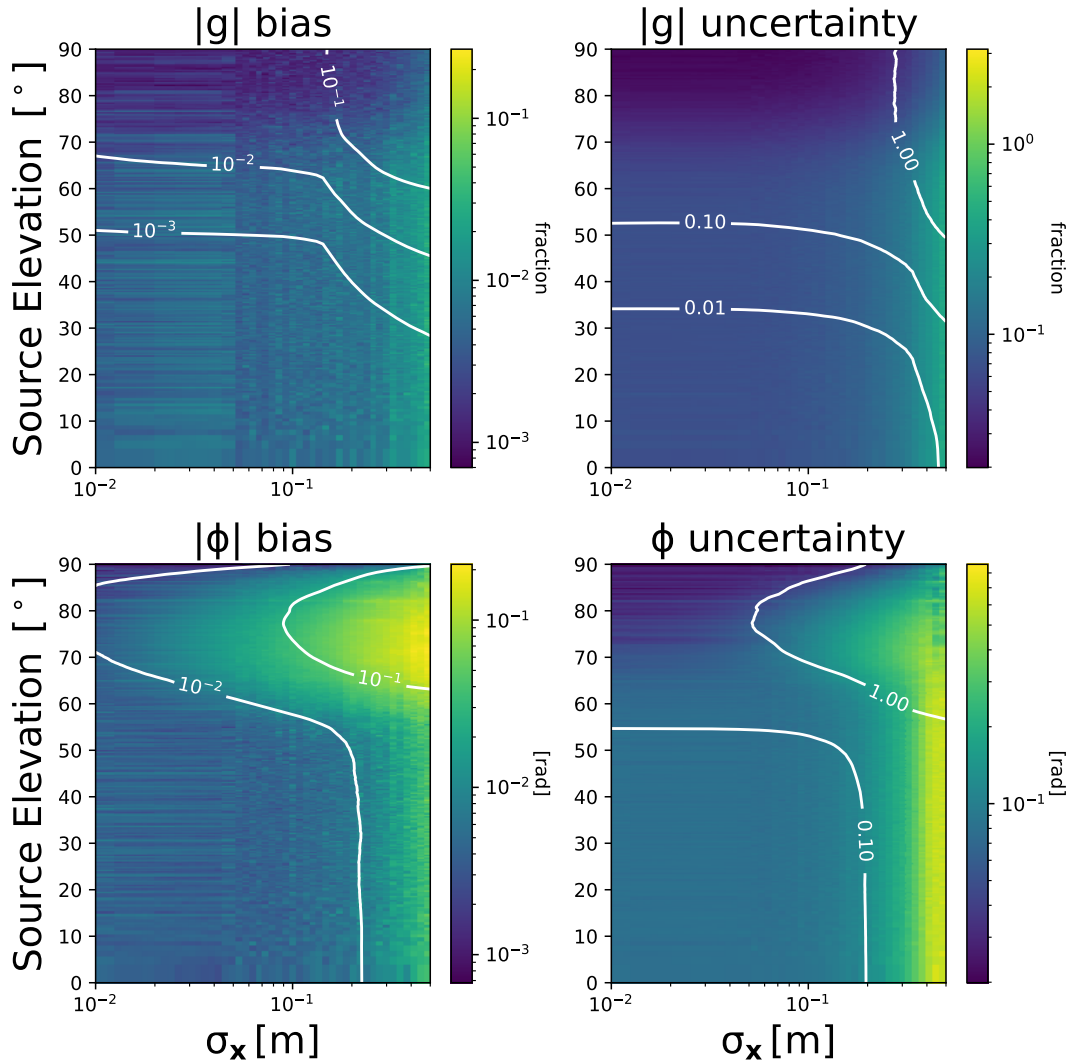


Figure 5.8: The bias (*top left*) and uncertainty (*top right*) of the amplitude solutions $|g|$, and the bias (*bottom left*) and uncertainty (*bottom right*) of the phase solutions $|\phi|$ of redundant calibration on an MWA hex-like array. We vary the position deviations σ_x and source elevation of a 100 Jy source. All results are averaged over all antennas in the hex. The contour lines are the ratios between the bias and uncertainty of redundancy based and sky model based calibration, e.g. an uncertainty contour line of 2.0 indicates the uncertainty is twice as large for redundant calibration as compared to sky model based calibration.




Figure 5.9: A 1-dimensional view on the bias (*top left*) and uncertainty (*top right*) of the amplitude solutions $|g|$, and the bias (*bottom left*) and uncertainty (*bottom right*) of the phase solutions $|\phi|$ of redundant calibration on an imperfectly redundant hex array. We vary the position deviations σ_x and source peak flux S_{peak} of a source located 87° above the horizon. All results are averaged over all antennas in the hex. The blue lines are the bias and uncertainties for an array with positional offsets $\sigma_x = 0.02$ m, in orange the bias and uncertainties for an array with positional offsets $\sigma_x = 0.1$ m, and in green the theoretical estimates for the bias and uncertainty of sky based calibration.

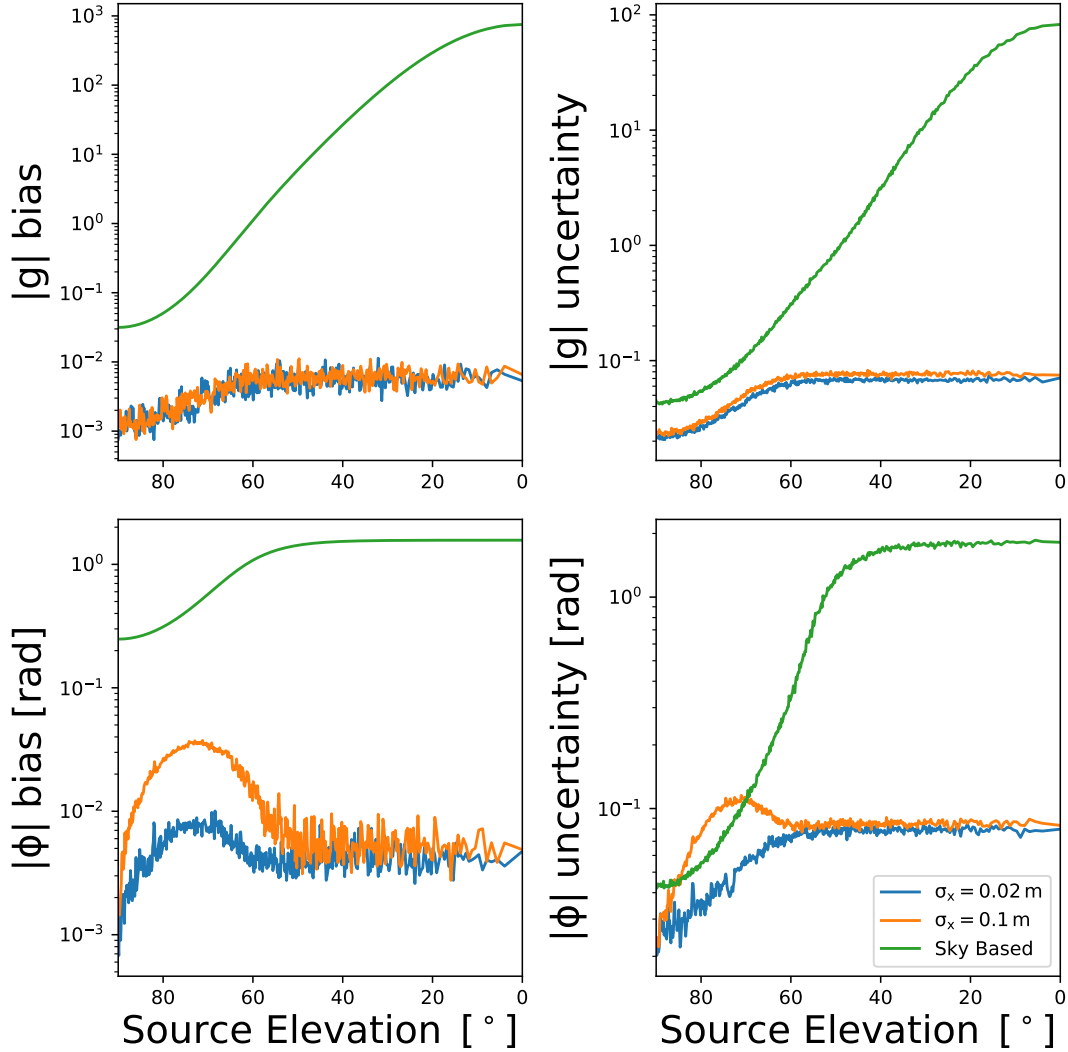


Figure 5.10: A 1-dimensional view on the bias (*top left*) and uncertainty (*top right*) of the amplitude solutions $|g|$, and the bias (*bottom left*) and uncertainty (*bottom right*) of the phase solutions $|\phi|$ of redundant calibration on an imperfectly redundant hex array. We vary the position deviations σ_x and source elevation of a 100 Jy source. All results are averaged over all antennas in the hex. The blue lines are the bias and uncertainties for an array with positional offsets $\sigma_x = 0.02$ m, in orange the bias and uncertainties for an array with positional offsets $\sigma_x = 0.1$ m, and in green the theoretical estimates for the bias and uncertainty of sky based calibration.

calibration even outperforms sky based calibration on a single source. However, if a redundant array, such as HERA, has positional offsets in the order of 10 cm, careful consideration has to be made on when to do redundant calibration. As demonstrated the phase bias can go up to an order of magnitude higher than that of the MWA-like array under these conditions. Fortunately for HERA the primary beam is narrower than that of the MWA, the latter suffers from significant side lobes, this and its large number of redundant baselines makes HERA somewhat robust against positional offsets (Liu et al., 2010). However, the exact trade off is still unclear.

We do note we have simplified the sky model based approach for analytic tractability. In reality a sky model will contain more than 1 calibration source, therefore the bias and uncertainty for a sky based approach will certainly be lower than presented here. But as a general lesson we can conclude that for redundant calibration it is preferable to have strong sources like those present in the EoR1 and EoR2 fields either at the pointing center or at the edge of the beam. EoR field 0 would therefore be an excellent field for redundant calibration.

In this work we have not considered the differences in the antenna response of different antenna tiles. Work by Line et al. (2018)) shows that the tile beam differences are on the order of 10%. This poses most likely the largest hurdle for redundant calibration. Studying the effect of these beam differences and how it impacts the redundancy of the MWA hexes and other radio telescopes is therefore crucial to understand the limitations of redundant calibration in realistic telescopes (see Chapters 6 and 7).

This is also where the true strength of sky based calibration methods comes into play. Because redundant calibration relies on the assumption that each antenna observes the same radio sky, it also is unable to solve for direction dependent effects introduced by different antenna responses and ionospheric distortions. The field of direction dependent calibration faced quite a number of challenges, e.g. diffuse emission detected by shorter baselines, solving for enough different

directions to capture variations in the ionosphere or the primary beam responses, optimizing the calibration time scale to reduce computational load, and the observed curvature of the sky (w-correction) due to the wide FoVs of these new arrays. A large effort has gone into solving these issues, e.g. **SAGEcal** resolved the computational load of solving for a large number of directions by using the SAGE algorithm rather than traditional least squares optimization (Yatawatta et al., 2009), facet calibration divides the sky in facets to reduce the number of parameters to solve for simultaneously (van Weeren et al., 2016), **RTS** employs the MWA’s *uv*-coverage to perform snapshot imaging tackling the *w*-term problem (Mitchell et al., 2008), with the diffuse emission of the Milky Way remaining as a major challenge. We have only mentioned a few implementations available as each science case has its own goal accompanied with its own implementation of sky based calibration. But the result of this large effort are impressive high-fidelity images, required to either study foreground sources or to subtract them. The latter being the goal for the EoR experiment. Solving for these higher-order calibration features is, however, out of reach for standard redundant calibration.

Furthermore, redundant calibration does not truly escape the need for a sky model, because the degenerate parameters discussed in Chapter 3 need to be constrained by external information, i.e sky-based calibration. Li et al. (2018) directly compare redundant calibration using **OMNICAL** and sky model calibration using **FHD** and find that they perform similarly on data from real MWA EoR observations, that include the position offsets and tile beam differences. Li et al. (2018) also investigate the complimentary nature of the two different calibration techniques and find that combining the two methods improves the sensitivity to the EoR power spectrum, demonstrating that “hybrid approaches” are the best way forward. However, this final step can also introduce errors with spectral structure due to an incomplete sky model. Redundant calibration is, at best, a way to add another constraint for a first order calibration step. Higher order effects require a pure sky based calibration that include direction dependent effects.

Nevertheless, redundant calibration can still add useful information if carefully applied. Sievers (2017) propose a calibration algorithm that sits in the middle ground between agnostic redundant calibration and pure sky model based calibration maximally using the information in the data of a generic radio telescopes. This methodology seems to be a promising avenue for the MWA and the future SKA.




Figure 5.11: A view of the radio sky at 408 MHz (Haslam et al., 1982), the plate carrée projection was created by *Skyview*, and the location of the three MWA EoR target fields: EoR0 centered at R.A. 0h and dec. -27° , EoR1 centred at R.A. 4h and dec. -30° and EoR2 centred R.A. 10h and dec. -10° . The circular areas represent the FWHM of the MWA beam at 150 MHz. EoR1 and EoR2 clearly have some strong radio sources away from the point center, e.g. Fornax A and Pictor A in EoR1 and Hydra A in EoR2.

5.6 Conclusion and Future Work

In this work we use a rudimentary implementation of the `logcal` and `lincal` algorithm to understand the fundamental limitations of redundant calibration on nearly redundant telescopes. We simulate redundant calibration under different radio sky conditions and find that the phase solutions are systematically impacted by position offsets in a redundant telescope. Based on our simulations we conclude

the following key statement: redundant calibration performs best when strong radio sources are either at field center or at the edge of the primary beam. We also compare redundant calibration to sky model based calibration and find that for the MWA redundant calibration of the redundant hexes performs better than a sky based approach. However, we require further work to understand the impact of non-redundancies introduced by differences in tile beam responses that may be of larger concern to the MWA EoR experiment. Moreover we also assumed that the sky model consists of only a single source. More work would be required to understand the completeness threshold above which a sky based approach truly outperforms redundant calibration on a nearly redundant array. Finally, to optimally calibrate our radio telescopes we require a hybrid approach that bridges the gap between redundancy based and sky based calibration, and we see this as the way forward in the calibration of the current and next generation of radio telescopes.

Acknowledgments

RCJ thanks Ruby Byrne for useful discussions, and Wenyang Li, Bart Pindor and Christopher Jordan for constructive feedback on this manuscript. This work was supported by the Centre for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), an Australian Research Council Centre of Excellence, funded by grant CE170100013. This research has made use of NASA's Astrophysics Data System. We acknowledge the International Centre for Radio Astronomy Research (ICRAR), a Joint Venture of Curtin University and The University of Western Australia, funded by the Western Australian State government. This scientific work makes use of the Murchison Radio-astronomy Observatory, operated by CSIRO. We acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited.

Chapter 6

Calibration and 21-cm Power Spectrum Estimation in the Presence of Antenna Beam Variations

This chapter is a reproduction of Ronniy C Joseph, C M Trott, R B Wayth, A Nasirudin, Calibration and 21-cm power spectrum estimation in the presence of antenna beam variations, *Monthly Notices of the Royal Astronomical Society*, Volume 492, Issue 2, February 2020, Pages 2017-2028, <https://doi.org/10.1093/mnras/stz3375>. It differs from the origin form due to minor alternations, such a referencing to earlier chapters to enhance the flow of this thesis.

Abstract

Detecting a signal from the Epoch of Reionisation (EoR) requires an exquisite understanding of galactic and extra-galactic foregrounds, low frequency radio instruments, instrumental calibration, and data analysis pipelines. In this work we build upon existing work that aims to understand the impact of calibration errors on 21-cm power spectrum (PS) measurements. It is well established that calibration errors have the potential to inhibit EoR detections by introducing

additional spectral features that mimic the structure of EoR signals. We present a straightforward way to estimate the impact of a wide variety of modelling residuals in EoR PS estimation. We apply this framework to the specific case of broken dipoles in Murchison Widefield Array (MWA) to understand its effect and estimate its impact on PS estimation. Combining an estimate of the percentage of MWA tiles that have at least one broken dipole (15%-40%) with an analytic description of beam errors induced by such dipoles, we compute the residuals of the foregrounds after calibration and source subtraction. We find that that incorrect beam modelling introduces bias in the 2D-PS on the order of $\sim 10^3 \text{ mK}^2 h^{-3} \text{ Mpc}^3$. Although this is three orders of magnitude lower than current lowest limits, it is two orders of magnitude higher than the expected signal. Determining the accuracy of both current beam models and direction dependent calibration pipelines is therefore crucial in our search for an EoR signal.

6.1 Introduction

Detecting a redshifted neutral hydrogen signal from the Epoch of Reionisation (EoR) is one of the most promising probes into formation history of structure in the Universe. The signal enables us to directly observe the state of the intergalactic medium (IGM) over a wide range of cosmic time and indirectly study the sources that impact it. As the very first luminous sources light up the Universe, they heat up the IGM and subsequently reionise it. The redshifted 21-cm line gives us direct insight into the evolution of the IGM temperature and the morphology of the ionisation structures carved out by the first sources of light (Morales & Wyithe, 2010; Pritchard & Loeb, 2012; McQuinn, 2016; Furlanetto, 2016).

However, despite a large international effort by various telescope collaborations around the world, e.g the Murchison Widefield Array (MWA; Tingay et al., 2013; Wayth et al., 2018), the LOw Frequency ARray (LOFAR; van Haarlem et al., 2013) and the Donald C. Backer Precision Array for Probing the Epoch of

Reionization (PAPER; Parsons et al., 2010), the signal has continued to elude a 21-cm power spectrum (PS) detection. The challenges faced by this experiment are large; foregrounds are 4-5 orders of magnitude brighter (Jelic et al., 2008) and the instruments have a complex signal chain. Understanding the behaviour of the latest instruments is ongoing work, and continues to provide crucial input to our calibration strategies. There already exists a large body of work on the residuals after direct subtraction of galactic and extra-galactic foregrounds, and their impact on the 21-cm PS. Most of this work treats the residuals of subtracted foregrounds as a source of Gaussian noise and studies how they affect the 21-cm PS assuming calibration leaves them unchanged (Liu & Tegmark, 2011; Trott et al., 2012; Dillon et al., 2013; Dillon et al., 2015; Trott et al., 2016; Murray et al., 2017). However, it has been well studied that unmodelled foreground noise is non-Gaussian and that outliers in the tail-end of the noise distribution impact calibration on a non-negligible levels (Kazemi & Yatawatta, 2013; Ollier et al., 2017, 2018). In particular, Barry et al. (2016) and Patil et al. (2016) show that sky based calibration in the presence of unmodelled foregrounds imparts additional spectral structure onto data, further inhibiting the detection of an EoR signal. Ewall-Wice et al. (2017) study this effect rigorously with a Gaussian approximation, and derive the imparted spectral structure on residual foregrounds due to modelling errors and found similar results. Similarly for redundant calibration, we know that non-redundancies impart bias onto the calibration solutions. Chapter 5 studies this for position errors on calibration solutions, and Orosz et al. (2019) study the impact of positions errors and beam errors on the 21-cm PS and find that indeed non-redundancies also impart spectral structure that contaminates EoR detections. On top of that, redundant calibration needs external information set by sky based calibration to determine overall calibration parameters, hence the limitations of sky based calibration set a fundamental limit on the calibration accuracy of redundant calibration. This yet again introduces additional spectral structure (Byrne et al., 2019).

In this chapter we build on existing work to further study the impact of modelling errors on sky model calibration, and on 21-cm PS estimation. We derive a framework that enables rather simple propagation of errors into calibration solutions and the EoR power spectrum. We focus our attention on what is undoubtedly the next challenge in the EoR experiment: the performance of individual elements. We describe the errors introduced by broken dipoles in the MWA and make an informed estimate on the contamination we expect. In section 6.2 we discuss calibration, source subtraction and signal estimation; in sections 6.3, 6.5 and 6.6 we derive the covariance matrices that describe our errors, propagate those forward to our gain solutions, and combine the two to derive the frequency structure of the calibrated residuals, respectively. In section 6.7 we present results from the derived framework, we compare sky and beam modelling errors, compare the results to a fiducial EoR signal, and we estimate the impact of broken dipoles in the MWA EoR experiment. We discuss the applicability of this framework to other sources of error, and the implications of these results in an EoR context in section 6.8.

Throughout this chapter, our notation is as follows: lowercase bold letters \mathbf{v} describe vectors, uppercase bold letters \mathbf{C} describe matrices, † denotes the Hermitian transpose, * denotes complex conjugation, and i is the imaginary unit.

6.2 EoR Signal Estimation with Gain Calibration Errors

The aim of calibration is to mitigate all effects that inhibit us from estimating the true sky intensities $I(\mathbf{l}, \nu)$, where \mathbf{l} is the sky coordinate vector and ν is the observing frequency. In the flat sky approximation we can relate the sky intensity $I(\mathbf{l}, \nu)$ to the complex visibilities $V(\mathbf{u}, \nu)$ measured by a pair of antennas in an

interferometer with baseline separation \mathbf{u} through

$$V(\mathbf{u}, \nu) = \int g_p g_q^* b_p(\mathbf{l}, \nu) b_q^*(\mathbf{l}, \nu) I(\mathbf{l}, \nu) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}, \quad (6.1)$$

where g_p is the complex-valued gain of antenna p , and $b_p(\mathbf{l}, \nu)$ is the corresponding beam voltage response. There are currently two popular methods of calibration; sky based calibration and redundant calibration.

Sky based calibration uses a model of the beam and the sky intensity to predict the visibilities V_{pq}^{data} measured by a pair of antennas p and q , and uses these model visibilities V_{pq}^{model} to solve for the unknown antenna gains g_p by minimising the squared differences (L^2 -norm) between the model and the data

$$\chi^2 = \sum_{pq} |V_{pq}^{\text{data}} - g_p g_q^* V_{pq}^{\text{model}}|^2. \quad (6.2)$$

Redundant calibration relies on having multiple identical baselines in arrays with a regular lay-out. These copies measure the same visibility, and therefore minimising the difference between the visibilities by varying the antenna gains in such groups allows us to find both the unknown antenna gains g_p and unknown redundant visibilities V_{α}^{true} without the need for a sky model.

$$\chi^2 = \sum_{\alpha} \sum_{pq, \alpha} |V_{pq}^{\text{data}} - g_p g_q^* V_{\alpha}^{\text{true}}|^2. \quad (6.3)$$

In general, these gains, or antenna responses, are direction dependent, i.e. they can capture deviations from the beam model or distortions by the ionosphere. A direction dependent calibration approach is, however, limited by the number of directions it can solve for. This limit is set by the number of bright sources available for calibration and computational costs. The number of directions ranges from 5-100 for current calibration pipelines. Throughout this work we focus on the direction independent gains as a simplification of the problem, noting that we can describe uncorrected directions by the same perturbative approach we are taking. The direction independent gains describe the global response of

the antenna and signal chain. Redundant calibration is inherently unable to solve for direction dependent effects.

In general, sky models are incomplete, redundant arrays have position errors, and there are variations in the antenna response. We can write our measured complex visibilities V , in the most general way, as a sum of a model m , residuals r that encompass errors on our model due to unmodelled sources, beam response variations, or even low level RFI (Wilensky et al., 2019), EoR signal s , and thermal noise n ;

$$V_{pq} = g_p g_q^* (m_{pq} + r_{pq} + s_{pq}) + n_{pq}. \quad (6.4)$$

When we have data of the form in Equation (6.4) and we calibrate using an incomplete model m , we inherently get incorrect gain estimates $\hat{g}_p = g_p + \delta g_p$ due to the presence of the residuals. When we apply these gain estimates to the data, see Equation (6.5), we get corrected visibilities \hat{V} that contain corruptions that inhibit us from detecting the EoR signal (Barry et al., 2016; Ewall-Wice et al., 2017; Patil et al., 2016). In general, we can ignore the EoR signal at the calibration step, because it is several orders of magnitude weaker than the noise and the foreground residuals.

$$\hat{V}_{pq} = \frac{g_p g_q^* (m_{pq} + r_{pq} + s_{pq}) + n_{pq}}{\hat{g}_p \hat{g}_q^*}. \quad (6.5)$$

After correcting the data, we subtract the sky model. In practice, it is subtracted as part of an iterative calibration process, i.e. “peeling” (Noordam, 2004). This leaves us with estimated data residuals \hat{r} that contain, amongst others, the EoR signal

$$\hat{r}_{pq} = \frac{g_p g_q^* (m_{pq} + r_{pq} + s_{pq}) + n_{pq}}{\hat{g}_p \hat{g}_q^*} - m_{pq}. \quad (6.6)$$

From these residuals we estimate the 21-cm PS. However, Equation (6.6) contains more than the cosmological signal of interest and we will derive the

covariance matrix of the additional residuals to understand their impact on our estimate of the 21-cm PS. We will study these errors from a power spectrum perspective and our approach is as follows:

1. We first compute the data residual covariance matrix $\mathbf{C}_r(u, \nu, \nu')$ within a power spectrum bin u over different frequencies. Specifically, we derive the contribution of beam errors due to broken dipoles (see section 6.3).
2. We then use this to compute an approximation of the gain error covariance matrix \mathbf{C}_g . Instead of computing the gain error per antenna, we compute averaged gain error covariance matrix for each power spectrum bin (see section 6.5).
3. Finally, we combine the two results to derive the covariance matrix of the gain-calibrated and source-subtracted residuals \mathbf{C}_f (see section 6.6).

To estimate how each error contributes to a bias in the EoR power spectrum we propagate these covariances from frequency-space forward to PS space. A Fourier transform over frequency is a linear operation that can be described by a complex matrix \mathbf{F} applied to our complex data vector containing frequency data. Hence, the covariance of the Fourier transformed data can be computed using standard linear error propagation: $\mathbf{F}^\dagger \mathbf{C} \mathbf{F}$. The variance of this propagated covariance matrix describes the added power due to these errors. The off-diagonals describe how this power correlates between different Fourier modes.

6.3 The Residual Covariance Matrix

In this section we derive the different contributions to the residual covariance matrix \mathbf{C}_r . To derive these matrices we start with the general covariance of visibilities,

$$\mathbf{C}_r = \text{Cov}[V(\mathbf{u}, \nu), V(\mathbf{u}', \nu')]. \quad (6.7)$$

To derive the residual covariance \mathbf{C}_r we assume we can separate this into a covariance between different baselines, and a covariance within a given baseline between different frequencies (Liu et al., 2014). In general, baselines with a separation $|\mathbf{u}_1 - \mathbf{u}_2|$ larger than the size of the Fourier transform of the primary beam decorrelate, and it suffices to compute the frequency covariance alone per u -bin.

We consider three contributions to the residual covariance: the covariance due to the sky \mathbf{C}_{sky} that describes the error due to unmodelled sources, the noise covariance \mathbf{C}_n that describes the error due to thermal noise, and the beam covariance \mathbf{C}_{beam} that describes the error due to deviations from the ideal beam model

$$\mathbf{C}_r = \mathbf{C}_{\text{sky}} + \mathbf{C}_{\text{beam}} + \mathbf{C}_n. \quad (6.8)$$

The noise covariance is independent from all other terms, and its structure is well known. We will not discuss it further in this chapter.

6.3.1 Sky Covariance Matrix

The sky covariance matrix \mathbf{C}_{sky} for a baseline at different frequencies has been well studied (Liu & Tegmark, 2011; Dillon et al., 2013; Trott et al., 2016; Murray et al., 2017). It describes the noise due to unmodelled sources and how this noise correlates between different frequency channels. We assume that an infinitesimal patch of sky $d^2\mathbf{l}$ contains a number of sources drawn from a Poisson distribution $\tilde{N} \sim \text{Poisson}\left(\frac{dN}{dS}dSd^2\mathbf{l}\right)$. The intensity of this patch is given by the first moment μ_1 of the source count distribution

$$\mu_n = \int_0^{S_{\text{max}}} S^n \frac{d\tilde{N}}{dS} dS. \quad (6.9)$$

We model the differential source counts $d\tilde{N}/dS$ with a broken power law model (Di Matteo et al., 2002; Ewall-Wice et al., 2017; Murray et al., 2017) to match observations in different flux regimes at low frequencies (Gervasi et al.,

2008; Intema et al., 2011; Franzen et al., 2016; Williams et al., 2016),

$$\frac{dN}{dS} = \begin{cases} k_1 S^{-\beta_1} & \text{if } S_{\text{low}} \leq S < S_{\text{mid}} \\ k_2 S^{-\beta_2} & \text{if } S_{\text{mid}} \leq S < S_{\text{high}} \end{cases}. \quad (6.10)$$

With this broken power law model we want to capture the difference between the distribution of modelled and unmodelled sources. We use $k_1 = k_2 = 4100$, $\beta_1 = 1.59$, $\beta_2 = 2.5$, $S_{\text{low}} = 100 \text{ mJy}$, $S_{\text{mid}} = 1 \text{ Jy}$, and $S_{\text{high}} = 10 \text{ Jy}$. For this model, we assume all sources above 1 Jy are used in calibration and all sources below that threshold are unmodelled. This results in the following expression for the sky covariance matrix,

$$\mathbf{C}_{\text{sky}} = 2\pi(f_0 f'_0)^{-\gamma} \mu_2 \Sigma^2 \exp(-2\pi^2 u^2 \Delta f^2 \Sigma^2). \quad (6.11)$$

Here, μ_2 is the second moment of the source count distribution, γ is the power law index that models the spectral energy distribution of each source, $f_0 = \nu/\nu_0$ where ν_0 is some reference frequency, e.g. the lowest frequency in the bandwidth, $\Delta f = f_0 - f'_0$. Following the notation of Murray et al. (2017), Σ contains products of the beam widths σ at different frequencies

$$\Sigma^2 = \frac{\sigma^2 \sigma'^2}{\sigma^2 + \sigma'^2}. \quad (6.12)$$

6.3.2 Beam Covariance Matrix

Here, we derive the beam perturbation covariance matrix. Similarly to the derivation of the sky covariance, we start out by taking Equation (6.1) under ideal gains $g = 1$, and assume we have a modelled sky intensity I and an unmodelled component δI . We extend this by adding a perturbation δb to the response of one antenna b_p , i.e. $b_p = b + \delta b_p$. We leave the other antenna responses as perfect,

$$V(\mathbf{u}, \nu) = \int b(b + \delta b)^*(I + \delta I) \times e^{-2\pi i \mathbf{u} \cdot \mathbf{1}} d^2 \mathbf{1}. \quad (6.13)$$

We have implicitly written the sky, the beam and their perturbations as functions of sky coordinate \mathbf{l} and frequency ν for brevity. The extra source of noise $\delta V(\mathbf{u}, \nu)$ is the sum of unmodelled components in the visibility

$$\delta V(\mathbf{u}, \nu) = \int \left(b\delta b^* I + bb^* \delta I + b\delta b^* \delta I \right) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}. \quad (6.14)$$

In this derivation, we assume the beam perturbation δb , the modelled sky intensity I , and unmodelled sky δI are random variables. We rewrite equation (6.7) by dividing the sky in voxels and write equation (6.14) as

$$\delta V(\mathbf{u}, \nu) = \sum_p \left(b_p \delta b_p^* I_p + b_p b_p^* \delta I_p + b_p \delta b_p^* \delta I_p \right) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_p} d^2 \mathbf{l}_p. \quad (6.15)$$

Combining equations (6.7) and (6.15), we rewrite the covariance as the sum of the covariances between the Fourier transforms of different sky voxels p and q

$$\mathbf{C}_r = \sum_p \sum_q \mathbf{C}_{pq}. \quad (6.16)$$

where the covariance between different voxels is given by

$$\mathbf{C}_{ij} = \text{Cov} \left[\left(b_p \delta b_p^* I_p + b_p b_p^* \delta I_p + b_p \delta b_p^* \delta I_p \right) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_p}, \right. \\ \left. \left(b'_q \delta b'_q{}^* I'_q + b'_q b'_q{}^* \delta I'_q + b'_q \delta b'_q{}^* \delta I'_q \right) e^{-2\pi i \mathbf{u}' \cdot \mathbf{l}_q} \right]. \quad (6.17)$$

If we extract the constant terms from the covariance, we can focus on the stochastic terms, i.e. the modelled and unmodelled fluxes, and the beam perturbations.

$$\mathbf{C}_{ij} = b_p b'_q{}^* e^{-2\pi i (\mathbf{u} \cdot \mathbf{l}_p - \mathbf{u}' \cdot \mathbf{l}_q)} \times \text{Cov} \left[\delta b_p^* I_p + b_p^* \delta I_p + \delta b_p^* \delta I_p, \right. \\ \left. \delta b'_q{}^* I'_q + b'_q{}^* \delta I'_q + \delta b'_q{}^* \delta I'_q \right]. \quad (6.18)$$

Using the formal definition of the covariance, $\text{Cov}[X, Y] = \langle XY \rangle - \langle X \rangle \langle Y \rangle$, and assuming δb , I , and δI are independent we can expand this further and simplify (see appendix B.2)

$$\begin{aligned}
\mathbf{C}_{ij} &= b_p b_q^* e^{-2\pi i(\mathbf{u}\cdot\mathbf{l}_p - \mathbf{u}'\cdot\mathbf{l}_q)} \times \left(\langle \delta b_p^* \delta b_q' \rangle \text{Cov}[I_p, I_q'] \right. \\
&\quad + \left(b_p^* b_q' + b_p^* \langle \delta b_q'^* \rangle + \langle \delta b_p^* \rangle b_q'^* + \langle \delta b_p^* \delta b_q' \rangle \right) \text{Cov}[\delta I_p, \delta I_q'] \\
&\quad \left. + \left(\langle I_p \rangle \langle I_q \rangle + \langle I_p \rangle \langle \delta I_q' \rangle + \langle \delta I_p \rangle \langle I_q' \rangle + \langle \delta I_p \rangle \langle \delta I_q \rangle \right) \text{Cov}[\delta b_p^*, \delta b_q'^*] \right)
\end{aligned} \tag{6.19}$$

Using the source count distribution that describes the number of sources in a flux bin we can write $\text{Cov}[I_p, I_q]$ as

$$\text{Cov}[I_p, I_q] = (f_0 f_0')^{-\gamma} \int S_p S_q \text{Cov}[N_p, N_q] dS. \tag{6.20}$$

However, because different parts of the sky are independent realisations of a Poisson distribution, $\text{Cov}[N_p, N_q]$ reduces to the mean for the same sky voxel ($\delta_{ij} dN/dS$). The 3rd term in Equation (6.19) then results in the unmodelled sky covariance \mathbf{C}_{sky} , see Equation (6.11). The remaining terms will be grouped in the beam covariance matrix \mathbf{C}_{beam} .

Combining Equations (6.16) and (6.15), and integrating this over the sky and all fluxes S we get

$$\mathbf{C}_r = \mathbf{C}_{\text{sky}} + \mathbf{C}_{\text{beam}} \tag{6.21}$$

where the full form of the beam covariance matrix is given by

$$\begin{aligned}
\mathbf{C}_{\text{beam}} &= (f_0 f_0')^{-\gamma} \mu_{2,m} \int \langle \delta b^*(\mathbf{l}, \nu) \delta b(\mathbf{l}', \nu') \rangle b(\mathbf{l}, \nu) b^*(\mathbf{l}', \nu') e^{-2\pi i(\mathbf{u}-\mathbf{u}')\cdot\mathbf{l}} d^2\mathbf{l} \\
&\quad + (f_0 f_0')^{-\gamma} \mu_{2,r} \int \left(b^* \langle \delta b'^* \rangle + \langle \delta b^* \rangle b'^* + \langle \delta b^* \delta b' \rangle \right) \\
&\quad \quad \quad \times b(\mathbf{l}, \nu) b^*(\mathbf{l}', \nu') e^{-2\pi i(\mathbf{u}-\mathbf{u}')\cdot\mathbf{l}} d^2\mathbf{l} \\
&\quad + (f_0 f_0')^{-\gamma} (\mu_{1,m} + \mu_{1,r})^2 \iint \text{Cov}[\delta b^*(\mathbf{l}, \nu), \delta b^*(\mathbf{l}', \nu')] \\
&\quad \quad \quad \times b(\mathbf{l}, \nu) b^*(\mathbf{l}', \nu') e^{-2\pi i(\mathbf{u}\cdot\mathbf{l} - \mathbf{u}'\cdot\mathbf{l}')} d^2\mathbf{l} d^2\mathbf{l}'.
\end{aligned} \tag{6.22}$$

To summarise, the beam covariance has three contributions:

1. The first contribution comes from the modelled sources and effectively describes the residuals due to subtraction of those sources with an incorrect beam model.
2. The second term describes how the noise from the unmodelled sources \mathbf{C}_{sky} is modified by the beam perturbations.
3. The last term describes the added covariance due to correlations between different parts of the beam.

The flux density of the sky is uncorrelated between different locations on the sky. However, because changes in the beam are in general correlated up to some correlation length, this introduces additional noise set by the mean flux of the sky. This last term is only important on $|u|$ -scales on the order of this correlation length, i.e the diameter of the antenna. In general, \mathbf{C}_{beam} is zero if the antenna response is ideal or more generally when the modelled response is equal to the actual response. The solution to the integrals in Equation (6.22) depends strongly on the form of beam perturbation.

6.4 The Phased Array Beam Model

We now derive the beam covariance matrix \mathbf{C}_{beam} for the “missing” dipole case in an MWA tile. The MWA consists of 128 tiles, and each tile is a 4×4 array of dipoles on a ground screen. The detailed steps of the derivation can be found in Appendix B.3. Here, we will only discuss the important steps for brevity. To derive the covariance of visibilities due to beam perturbations, we start out with the formal description of the beam response of a phased array b_{tile} consisting of N dipoles on a ground screen (see Chapter 3),

$$b_{\text{tile}} = b_{\text{dipole}} \times \sum_{n=0}^N w_n \exp \left[-2\pi i \mathbf{x}_n \cdot \mathbf{l} / \lambda \right]. \quad (6.23)$$

In this description we assume that the individual element electric field responses b_{dipole} are identical, and we can then multiply this single element beam with the array factor to create a compound beam. In this array factor w_n is the element weight. For the zenith pointings considered in this work, these weights range from 0 to 1, however, in general these weights are complex. \mathbf{x}_n is the location of the n th element with respect to the center of the phased array. For simplicity, we approximate the full tile beam as a frequency dependent Gaussian, following

$$b_{\text{tile}} = \exp[-|\mathbf{l}|^2/2\sigma^2(\nu)]. \quad (6.24)$$

We define the width of the voltage beam as $\sigma = \sqrt{2}\epsilon c/D\nu$ with $\epsilon = 0.42$. We re-scale an Airy disk to a Gaussian width using ϵ and assume an MWA tile diameter of $D = 4$ m. The factor $\sqrt{2}$ ensures that the square of the voltage beam is consistent with descriptions of the beam used in the literature.

6.4.1 A Broken MWA Dipole

Although, there are many ways in which we can perturb the beam response, a common and relatively straightforward perturbation is the broken dipole case. Figure 6.1 shows the number of tiles with either one broken dipole in the X- or Y-polarisations, or two broken dipoles, one in each polarisation, for EoR observations over the past years. MWA tiles that have more than one broken dipole in the same polarization will be flagged and their data are not used. At most 50 out of 128 tiles have been marked as having 1 broken dipole. Throughout this chapter we choose a lower limit in this, adopting 25 broken dipoles which corresponds to $\sim 30\%$ of the visibility data.

We describe the broken dipole perturbation to the beam δb by subtracting the contribution of a missing dipole¹

¹Note: we explicitly consider phase offsets due to missing dipoles.

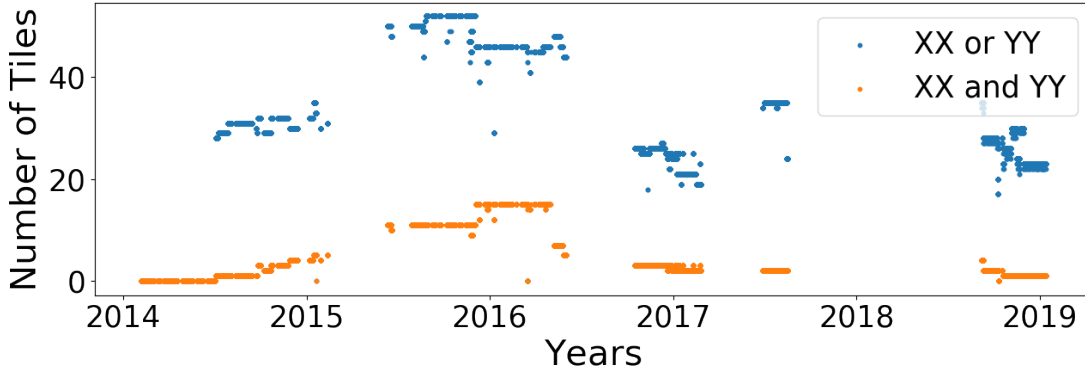


Figure 6.1: The number of MWA tiles with 1 broken dipole in a polarization and the number of tiles with 2 broken tiles in two different polarizations over several years of MWA EoR data.

$$\delta b = -b_{\text{dipole}} \times w_{\text{broken}} \exp \left[-2\pi i \mathbf{x}_{\text{broken}} \cdot \mathbf{l} / \lambda \right]. \quad (6.25)$$

We approximate the response of the missing dipole towards the sky by a Gaussian, with a diameter D that is 1/4 of the full tile, and a weight $w_{\text{broken}} = 1$. This is equivalent to completely removing the dipole. To appropriately account for the contribution of a single dipole relative to the response of an N -element phased array, we normalise the dipole response by N :

$$\delta b = \frac{1}{N} e^{-2\pi i \mathbf{x}_n \cdot \mathbf{l} / \lambda} b_{\text{dipole}}. \quad (6.26)$$

We now have all the tools to derive the specific structure of \mathbf{C}_{beam} for the missing dipole case. We refer interested readers to Appendix B.3 for details of this derivation, here we discuss results and implications for the 2D-PS.

Figure 6.2 shows the 2D-PS for a 30 MHz bandwidth centred at 150 MHz with 251 frequency channels of the unmodelled sky variance, the beam variance and the total variance in cosmological units (see Chapter 2 for the conversion between frequency covariance matrices and the PS). To show all effects due to beam modelling errors, we have extended the range of k_{\perp} beyond the conventional EoR range. The overall structure of the foreground wedge has not changed




Figure 6.2: The unmodelled sky covariance (left), the change in the covariance due to addition of beam perturbations to the covariance (middle), and the resulting total covariance (right). We show the 2D-PS for a 30 MHz bandwidth centred at 150 MHz with 251 frequency channels. Removing dipoles primarily takes away power from the wedge, because the array is less sensitive. However, this leaves us with modelled foreground residual power as those sources are incorrectly subtracted. This effect manifests itself primarily at the edge of the primary beam (or in the sidelobes of realistic phased arrays). This indicates contamination comes primarily from broken dipoles at the edge of a tile. The changing shape of the beam due to broken dipoles is a fairly large scale effect, both spatially and in frequency. Hence, it introduces power only at the the smallest k_{\perp} and k_{\parallel} .

drastically, but if we look at the beam covariance component alone, we see that including beam modelling errors changes the noise in three regions. At small k_{\perp} , corresponding to baseline lengths shorter than the physical dimensions of an MWA tile, we see the contribution due to the correlations from different parts in the beam. We also see that there is overall less power in the wedge, because the missing dipoles decrease the sensitivity of the array as a whole, and we therefore see less of the foregrounds. Finally, we see the residuals of the modelled foregrounds on the edge of the foreground wedge. The largest relevant change due to missing dipoles is in this region, and is dominated by sources in the sidelobes. Note: beam errors do not intrinsically contaminate the EoR window, any visible excess is leakage due to the first sidelobe of the Fourier transform of the Blackman-Harris window.

6.5 The Gain Error Covariance Matrix

Now that we have expressions for the data residuals, we explore how they propagate to the gain solutions during calibration. We derive the covariance matrix of the averaged gain error δg within a PS bin u . Sky model based calibration aims to solve

$$\hat{g}_p \hat{g}_q^* m_{pq} = g_p g_q^* (m_{pq} + r_{pq}). \quad (6.27)$$

We can rearrange this to solve² for the ratios of true and estimated gain solutions $\hat{g}_p = g_p + \delta g_p$ on the left hand side, and total signal over model visibilities on the right hand side

$$\frac{\hat{g}_p \hat{g}_q^*}{g_p g_q^*} = \frac{(m_{pq} + r_{pq})}{m_{pq}}. \quad (6.28)$$

We assume ideal gains $g_p = 1$, and that the gain error is small $\delta g_p \ll g_p = 1$ for all p . Most of the signal is contained in the sky model, and the residuals are much smaller, and we can therefore expect small gain errors:

$$\delta g_p + \delta g_q^* = \frac{r_{pq}}{m_{pq}}. \quad (6.29)$$

Applying this to all baselines and corresponding antenna gain errors yields a system of equations that can be rewritten in matrix form $\mathbf{Ax} = \mathbf{y}$. The vector \mathbf{y} contains the residual-to-model ratios r_{pq}/m_{pq} , and the vector \mathbf{x} contains the gain errors δg_p . The array matrix \mathbf{A} relates a baseline to the antennas that it is made up from. Its (pseudo-)inverse \mathbf{A}^{-1} tells us how much each ratio r_{pq}/m_{pq} in a baseline contributes to an error in a gain solution δg_p . It also implies that the gain error δg_p is a weighted sum of the residual-to-signal ratio in each baseline,

²Note: to solve this system we need to split each entry into its imaginary and real component, and set a reference antenna. We compute all errors relative to that reference antenna, implying that the reference antenna should be error free or close to.

$$\delta g_p = \sum_n w_{pn} \frac{r_n}{m_n}. \quad (6.30)$$

Here, the weights w_{pn} are the entries of the inverse of the array matrix \mathbf{A}^{-1} , where the index n runs over each visibility measurement (instead of ij). To compute the gain error covariance matrix \mathbf{C}_g , we need to compute $\text{Cov}[\delta g, \delta g']$. Noting that Equation (6.30) is a sum over different baselines, that in general fall in different u -bins, and assuming the covariance between different baselines is zero leaves us with

$$\mathbf{C}_g = \sum_n w_n^2 \text{Cov} \left[\frac{r_n}{m_n}, \frac{r'_n}{m'_n} \right]. \quad (6.31)$$

Instead of computing the covariance of a ratio of random variables, we approximate Equation (6.31) by replacing the model signal m with the r.m.s. of the modelled sky $f_0^{-\gamma} \sqrt{\mu_{2,m}}$. This simplifies the expression for the gain error covariance matrix \mathbf{C}_g . Figure 6.3 shows 10 000 visibility amplitude realisations of a stochastic sky, the amplitude of the mean visibility, the mean of the realised visibility amplitudes, and the sky r.m.s. From this we conclude that the sky r.m.s. provides a reasonable approximation to the modelled visibility m_n . We replace the modelled visibility m_n with $f_0^{-\gamma} \sqrt{\mu_{2,m}}$, yielding

$$\mathbf{C}_g = \frac{(f_0 f'_0)^\gamma}{\mu_{2,m}} \sum_n w_n^2 \text{Cov} [r_n, r'_n]. \quad (6.32)$$

We study how data in one bin u_1 contributes to the error in a calibrated bin u_2 ; knowing that each baseline contributes to some gain solutions, each gain solution is applied to the data, and finally these data are binned into $|u|$ -bins. Instead of deriving what the gain error is on a single baseline, we derive how all uncalibrated PS bins change the structure of a calibrated bin through the “averaged gain covariance”.

In our linearised approximation, each antenna gain error is a sum of $(N - 1)$ baseline errors, i.e. all baselines in which that antenna participates. Each




Figure 6.3: 10 000 realisations of a stochastic sky and their amplitude across different u -scales. The average amplitude of all realisation (blue), the amplitude of the average visibilities (orange), the amplitude of each individual realisation (black), and the expected sky r.m.s. for a power law distributed poisson sky (dashed green). On average all baseline do measure some signal, which is reasonably approximated the expected sky r.m.s.

baseline therefore contributes $\sim 1/(N - 1)$ to the gain solutions of the antennas it participates in. We smear this contribution out over all baselines, assuming each baseline contributes $\sim 1/N_b(N - 1)$ to all gain solutions. Two of these antenna errors then propagate to a calibrated baseline. We compute the number of baselines in an uncalibrated bin that see baselines in a calibrated bin and relate that to the error weights $w(u)$, following

$$w(u) = \frac{1}{N - 1} \frac{N_b(u)}{N_b(\text{total})}. \quad (6.33)$$

For an unrealistic array with perfect uniform uv -coverage, the number of baselines drops out, and the weights are $w = 1/N_{\text{bins}}(N - 1)$. For the MWA, we compute the actual distribution of baselines binned in the same way we bin all k_{\perp} .

In Figure 6.4, we show the binned MWA baselines, the variance of the Fourier transformed gain covariance assuming a uniform baseline distribution, and the Fourier transformed gain covariance for MWA Phase II Compact. For a uniform baseline distribution, the variance is the same for each k_{\perp} -bin; however, for the MWA the structure of the variance is different depending on how much it is coupled to other scales. In the next section, we demonstrate that this variance of the gain error effectively becomes a convolution window that smears out power from the foreground wedge into the EoR window.

6.6 The Gain Corrected Residual Covariance Matrix

After obtaining our gain estimates \hat{g} , we apply them to the data and subtract our model visibilities m . This leaves us with our residual estimates \hat{r} from which we aim to detect the EoR signal, see Equation (6.6). We assume the errors on our gains solutions are small $\delta g_p/g_p \ll 1$, enabling us to Taylor-expand ratios between our true gain solutions and the estimates g_p/\hat{g}_p . Grouping terms in products with either the model m or the residuals r results in




Figure 6.4: *left*: A histogram of MWA Phase II compact baselines binned into power spectrum bins. The spikes arise from the redundant hexagons that produce many multiples of the same baselines. *middle*: The Fourier transformed gain variance for an array with uniform uv -coverage with sky-only, and sky and beam errors. *right*: The Fourier transformed gain variance for an MWA Phase II compact uv -coverage for sky model errors only.

$$\hat{r}_{pq} = -(\delta g_p + \delta g_q^*)m_{pq} + (1 - \delta g_p - \delta g_q^*)r_{pq}. \quad (6.34)$$

When we estimate the EoR signal, we grid and average these residual visibilities onto a uv -grid before Fourier transforming along the frequency-direction. We now want to compute the covariance of the averaged gridded residuals to understand the full impact of sky and beam modelling errors on the 21-cm PS

$$\begin{aligned} \mathbf{C}_{\hat{\mathbf{r}}}(u, \nu) &= \text{Cov}[\hat{\mathbf{r}}, \hat{\mathbf{r}}'] \\ &= \langle \hat{\mathbf{r}} \hat{\mathbf{r}}'^{\dagger} \rangle - \langle \hat{\mathbf{r}} \rangle \langle \hat{\mathbf{r}}' \rangle^{\dagger}. \end{aligned} \quad (6.35)$$

Equation (6.34) keeps track off individual antenna p and q and how they impact a baseline. However, we are interested in the covariance of calibrated and model-subtracted data binned at scale u . The details of these final steps can be found in Appendix B.4. Here, we only describe the general assumptions we made. To derive the covariance structure, we see the gains g_p and g_q as two realisations of a random variable within a $|u|$ -bin, i.e. the two gain errors are independent. We also assume the gain error is independent from the model and data residuals of a certain u -bin. We justify the latter assumption because the gain error on a baseline or a u -bin is a linear combination of the residuals of different baselines.

We already assumed that residuals in different u -bins are uncorrelated. This implies that a sum of residuals over different u -bins decorrelates with the residuals of a single u -bin (if we include sufficiently independent residuals). We therefore assume that the gain error is independent from the residuals of the u -bin in which we are computing the covariance of the calibrated residuals $\mathbf{C}_{\hat{\mathbf{r}}}$. When we combine Equation (6.34) and (6.35), we can safely drop the expectation values of the estimated residuals $\langle \hat{\mathbf{r}} \rangle$, because they integrate to zero for baselines longer than the tile diameter. Under these assumptions, the covariance of calibrated data residuals reduces neatly to

$$\mathbf{C}_{\hat{\mathbf{r}}} = 2\mathbf{C}_{\mathbf{g}} \odot \mathbf{C}_{\mathbf{m}} + (1 + 2\mathbf{C}_{\mathbf{g}}) \odot \mathbf{C}_{\mathbf{r}}. \quad (6.36)$$

This final equation compactly describes how calibration propagates errors on our data model into the data product from which we estimate an EoR signal. A product of covariance matrices is a fairly unusual expression. In this specific case, $\mathbf{C}_{\mathbf{g}}$ describes how much of the original covariances – the model or residual covariance – remain and how the correlation between different frequencies is changed. It is a product between the gain error and either the model covariance $\mathbf{C}_{\mathbf{m}}$ or residual covariance $\mathbf{C}_{\mathbf{r}}$ in Fourier space $\{u, v, \nu\}$, and therefore becomes a convolution in PS-space $\{u, v, \eta\}$ that smears out power from the foreground wedge throughout the PS. How much power is actually smeared out depends strongly on the errors that cause this correlation in the gain estimates.

6.7 Results

Now that we have a relatively simple expression that describes the covariance of calibrated data residuals $\mathbf{C}_{\hat{\mathbf{r}}}$, we first apply this to an array with identical tile beams, i.e we consider unmodelled sky noise only, to compare with earlier studies on this problem. We then include beam errors to study how this extra source of modelling errors introduces contamination into the EoR window. We




Figure 6.5: Fiducial EoR PS at redshift 8. We create this 2D PS by deprojecting a 1D EoR PS where the reionisation process is driven by faint galaxies (Mesinger et al., 2016).

then compare these results to a fiducial EoR signal. We have taken a 1D-PS for a faint galaxy driven model at redshift $z \sim 8$ from Mesinger et al. (2016), and deprojected it into a 2D-PS assuming spherical symmetry, see Figure (6.5). The code that generated the results presented here is publicly available (see Joseph, 2019).

6.7.1 Sky Model Errors

Figure 6.6 shows the data residuals after calibration and model subtraction on the left, the difference between uncalibrated and calibrated residuals in the middle, and the ratio between that difference and our fiducial EoR signal on the right. The




Figure 6.6: Comparing calibrated and uncalibrated residuals. Left: Calibrated and subtracted residuals with unmodelled sky noise only. Middle: Difference between calibrated and uncalibrated residuals. Right: Ratio between the difference and a fiducial EoR PS.

difference plot clearly shows the contamination into the EoR window introduced by the structure of the calibration errors. The additional bias we reproduce is similar to the results by Barry et al. (2016) and Ewall-Wice et al. (2017), demonstrating that our formalism neatly reproduces earlier results. Here, we have only reproduced the results that describe the impact of unmodelled sources on a per frequency channel calibration strategy. We have not incorporated the mitigation strategies proposed to suppress excess noise due to these unmodelled sources. However, we note that the expected signal is strongest at the lowest k -modes hence mitigating leakage due to calibration is particularly important at these scales.

6.7.2 Comparing Sources of Error

We now apply this to data residuals including the beam covariance matrix \mathbf{C}_{beam} . Figure 6.7 shows the variance of the sum of calibrated residuals, the difference with calibrated sky model only errors (the left most PS in Figure 6.6), and the ratio of this difference with a fiducial EoR signal. We see that the introduction of beam modelling errors adds power to the edge of the wedge as noted earlier in Figure 6.2. Now that we also include calibration, it smears out power further into the EoR window, causing a relative drop in power in the EoR window. This is not




Figure 6.7: Comparing calibrated residuals with sky only and both sky and beam errors. Left: Calibrated Sky and Beam error residuals. Middle: Difference with Calibrated sky only residuals. Right: Ratio between difference and a fiducial EoR PS. Calibration in the presence of unmodelled sources adds significant contamination into the EoR window.

too surprising because Figure 6.4 shows that beam errors changes the structure of the gain errors, such that they extend to larger k_{\parallel} .

Ultimately, we want to estimate the expected contamination in the EoR window due to these errors under realistic circumstances. Figure 6.1 shows the number of broken MWA tiles during EoR observations, and from this we estimate $\sim 30\%$ of our visibility measurements to be contaminated. We appropriately down weight the beam covariance term C_{beam} with a factor of 0.3^2 ; one factor for both frequencies. We also use the baseline distribution in Figure 6.4 to properly weight the covariance of each u -scale when we compute the gain error covariance matrix. Figure 6.8, shows the expected results for an MWA-like data set. Down weighting the beam covariance changes the structure of the gain covariance in a way that the leakage does not extend that far into the EoR window. Nevertheless, the additional contamination is still on the order of the expected EoR signal.

6.8 Discussion

The structure of the beam errors presented here is dominated by incorrectly subtracted sources towards the horizon. This hints that frequency structure mit-




Figure 6.8: Expected MWA data contamination: Comparing calibrated residuals with sky only and both sky and beam errors. Left: Calibrated Sky and Beam error residuals. Middle: Difference with Calibrated sky only residuals. Right: Ratio between difference and a fiducial EoR PS.

igating techniques should rid us of most of the additional power due to beam modelling errors. Barry et al. (2016) discussed these techniques after finding that incomplete sky models cause foreground power to be convolved with erroneous calibration solutions. They compute an expected level of contamination due to sky modelling error consistent with their simulation on the order of $\sim 10^7 \text{mK}^2 \text{h}^{-3} \text{Mpc}^{-3}$. We estimate contamination due to beam modelling errors to be 3 orders of magnitude lower than this. Patil et al. (2016) suggest multi-frequency calibration as a way around this. Multi-frequency calibration can enforce spectral smoothness on the solutions through a smoothness regulariser (Yatawatta, 2015), and subsequently decrease the variance reducing contamination of the EoR window. However, using all frequency information is computationally challenging and requires appropriate software architecture to overcome limited compute power (Yatawatta et al., 2017).

Ewall-Wice et al. (2017) derived the structure of the gain corrected residuals by keeping track of the baseline ordering. This leads them to directly relate the contamination from the longest baselines into the shorter baselines. When we compute the averaged gain covariance (see Equation 6.30), we take the average of the residuals covariances at different k_{\perp} bins to describe the same effect. They suggest down weighting the longer baselines during calibration, as these baselines

are the source of spectral contamination. On the other hand Patil et al. (2016) suggest excluding the shortest baselines for which we currently lack accurate models of the diffuse foregrounds. However, they also demonstrate that excluded baselines suffer from enhanced noise after calibration. Creating diffuse sky maps similar to Eastwood et al. (2018) for the southern sky between $\nu \sim 35 - 70$ MHz is therefore crucial for accurate calibration. However, in future work we should also consider contamination from short baselines on which diffuse emission from the galaxy dominates due to the lack of such models.

This framework we derived can also be used to study the impact of non-redundancies on redundant calibration, i.e. the antenna position errors and beam variations. However, redundant calibration ultimately needs some external information to set overall gain parameters, e.g. absolute amplitudes and phase gradients see Wieringa (1992); Liu et al. (2010) for more details. Byrne et al. (2019) show that incomplete sky models fundamentally limit the accuracy of redundant calibration solutions due to limitations in sky model based calibration. We expect this to be exacerbated by beam modelling errors that push power further into the window. Orosz et al. (2019) simulate redundant calibration for HERA including non-redundancies. They find that beam variations severely contaminate the EoR window. In future work, we will study what actually poses the largest hurdle to redundant arrays: inherent non-redundancies or incomplete sky models.

Li et al. (2018) study for the first time how redundant calibration and sky model calibration can be used optimally in MWA Phase II compact. They demonstrate that adding redundant calibration improves their results. In an effort to bridge the gap between redundant and sky based calibration, Sievers (2017) developed the “correlation calibration” framework that incorporates uncertainties on calibration models, e.g. sky model incompleteness, position offsets, and beam variations. There are currently tentative results that this is a very promising path forward; however, more work is required to properly compare this to current calibration techniques.

Similar to earlier theoretical work we have not considered the non-Gaussian nature of the noise discussed by Kazemi & Yatawatta (2013); Ollier et al. (2017) and Ollier et al. (2018). In this work we study the impact of the variance as the PS only sensitive to that. It is therefore not unreasonable to expect that the results presented here underestimate the level of contamination. We have also not considered mitigation strategies for EoR window contamination. In our estimation of added contamination due to beam modelling errors, we overestimate the expected errors for the MWA EoR experiment. Both MWA EoR pipelines RTS/CHIPS (Mitchell et al., 2008; Trott et al., 2016) and FHD/ ϵ ppsilon (Sullivan et al., 2012; Jacobs et al., 2016; Barry et al., 2019a) incorporate direction dependent calibration. RTS/CHIPS uses information about broken dipoles explicitly to better model individual MWA tile beams. Similarly, many other calibration pipelines perform direction dependent calibration. However, we need to further quantify how much the responses vary from across an array similar to Line et al. (2018), and how well direction dependent calibration captures variations in the beam to better estimate the expected contamination due beam modelling errors. Given the order of magnitude in which beam modelling errors manifests themselves, it seems very plausible that these errors are potentially our next systematic.

6.9 Conclusion

Inspired by earlier theoretical work in this field, we have derived a relatively intuitive framework that neatly describes contamination in the EoR window due to calibration. We have reproduced earlier results and computed expected errors introduced by beam modelling errors. In this work, we have specifically focused on broken dipoles in the MWA as a perturbation to the model beam because this a relatively straightforward example. However, our results are applicable to a wide range of modelling errors, e.g. more complex beam variations and signal path variations, if analytic descriptions exists for these. We estimate that $\sim 15\% - 40\%$ of the MWA tiles have at least one broken dipole. We have made a rough

estimate on the order of magnitude $\sim 10^3 \text{ mK}^2 h^{-3} \text{ Mpc}^{-3}$ in which contamination by beam modelling errors manifests itself. These numbers are only indicative and we need to further quantify beam variations in-situ and determine how well current calibration techniques are able to account for this. However, we expect that these beam errors could potentially be our next limiting factor in the EoR experiment.

Comparing Calibration Systematics in Redundant Arrays Searching for Cosmic Reionisation

In the previous chapters, I presented an analysis of the impact of errors through a lower bound on the minimum variance of calibration solutions, and the impact of calibration errors on the power spectrum. In this chapter, I will tie these results together to compare the performance of redundant and sky-based calibration under current sky modelling limits and array deployment tolerances for the Murchison Widefield Array (MWA). I use the Cramér-Rao Lower Bound (CRLB) to understand the limitations of sky and redundancy based calibration in the limit of extremely large arrays. I also compare the spectral structure of errors arising in redundant and sky-based calibration and their impact on the 21-cm power spectrum (PS).

7.1 Introduction

Redundancy in radio interferometric arrays has the potential to improve the sensitivity to the 21-cm PS by an order of magnitude (Parsons et al., 2012a). The

regular lay-out of redundant arrays causes oversampling of certain parts of the uv -plane. Because these redundant measurements can be added coherently they provide enhanced sensitivity towards those regions as compared to irregular arrays for the same observing time. We can tune the lay-out of these arrays to improve the sensitivity towards Fourier modes occupied by the EoR signal, and simultaneously avoid modes dominated by Galactic and extragalactic foregrounds. This led to deployment of several EoR experiments with redundancy incorporated into their design, e.g. The Donald C. Packer Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al., 2012b), The Hydrogen Epoch of Reionization Array (HERA; DeBoer et al., 2017), The Murchison Widefield Array Phase II (MWA; Wayth et al., 2018).

Apart from the improvement in sensitivity, we can also use redundancy for calibration (Wieringa, 1992; Liu et al., 2010). As discussed in earlier chapters, redundant calibration does not rely on prior knowledge of the sky and at first glance provides an “unbiased” method of calibration. Philosophically, this aligns with our quest towards a signal that has not been observed before in an experiment that is extremely sensitive to biases. However, on closer inspection, deviations from redundancy do introduce non-negligible biases into the calibration solutions (see Chapter 5), and subsequently impact estimation of the 21-cm PS (Orosz et al., 2019). Liu et al. (2010) provide an augmentation to classic redundant calibration that also solves for non-redundancies to first order at the cost of a higher number of degrees of freedom, however, current implementations lack the proposed feature.

In addition to the bias introduced by non-redundancies, redundant calibration is also not completely free of sky based information. The inherent degeneracies discussed in Chapter 3 need to be resolved through “external calibration”. There are a variety of ways to break the degeneracy using sky based information. We can apply both redundant and sky based calibration and project the solutions of one to the other (Li et al., 2018). Alternatively, we can use the redundant cali-

bration solutions as a starting point for a sky based calibration routine (Li et al., 2019) or vice versa (Kern et al., 2019). We can also forego sky based calibration completely, and simply use visibilities generated from a sky model to directly fix the degeneracy parameters (Kern et al., 2019). However, all methods suffers from the incompleteness of a sky model, and subsequently introduce additional bias into the 21-cm PS (Byrne et al., 2019).

Although redundant arrays naturally lend themselves to redundant calibration, they can also be calibrated using traditional sky based calibration approaches (Li et al., 2018; Kern et al., 2019). As per usual this calibration suffers from the incompleteness of the sky model, limited by the depth of previous surveys of the radio sky. Additionally, redundant arrays that are tuned for EoR experiments are not well suited for imaging due to a lack of uv -coverage. This hinders the implementation of self-calibration to adequately resolve sources on the sky. If the array suffers from a narrow field of view (FoV) it might prove hard to find a sufficient number of bright calibrator sources (Kern et al., 2019).

A recurrent theme throughout this thesis is that both calibration methods have inherent advantages and disadvantages. This has led to some effort to bridge the gap between the two, i.e. hybrid calibration. These hybrid calibration efforts have found to marginally improve the gain solutions (Li et al., 2018; Kern et al., 2019). However, a deeper understanding of the underlying issues in both methods might prove to be insightful in a way forward. In this chapter I investigate the limitations of the two calibration methods in an extreme limit of large-N arrays by computing a lower bound on the variance of the gain solutions. I also compute the contamination expected from redundant and sky based calibration given current array tolerances. I will conclude with a discussion on the implications of these results and the path forward towards a detection of an EoR signal.

Throughout this chapter I adopt the following convention; boldface letters \mathbf{d} denote vector quantities, capitalised boldface letters \mathbf{C} denote matrices, † denotes the Hermitian transpose, and \odot denotes element-wise matrix multiplication.

7.2 Generalised Calibration Analysis

Regardless of the calibration algorithm, the goal is to solve for the gain g_q for an antenna q per frequency and time interval

$$V_{pq}^{\text{data}} = g_p g_q V_{pq}^{\text{true}}, \quad (7.1)$$

where V_{pq}^{true} is the true visibility, and V_{pq}^{data} is the visibility as measured after the signal propagated through the ionosphere and instrument. Both redundant and sky-based calibration assume a model m_{pq} . In redundancy based calibration the model is implicit, i.e. all baseline of the same orientation and length measure the same visibility. In sky-based calibration the model is explicitly based on current knowledge of the radio sky. However, the models do not perfectly describe the underlying true visibility

$$V_{pq}^{\text{true}} = m_{pq} + r_{pq}, \quad (7.2)$$

and the residuals r_{pq} will propagate into the calibration solutions. In the previous chapter I demonstrated the descriptive power of analytic covariance matrices describing the statistical properties of these residuals. Here, I revisit those results for sky based calibration, derive similar descriptions for redundant calibration, and index them with a superscript r for redundant calibration and s for sky based calibration if appropriate.

7.3 Sky Based Calibration Errors

For sky based calibration I consider errors in our sky model due to unmodelled foreground sources $\mathbf{C}_{\text{sky}}^s$ (see Equation 6.11) and errors in the beam model $\mathbf{C}_{\text{beam}}^s$ (see Equation 6.22), and thermal noise \mathbf{C}_N . The general total residual covariance before sky based calibration is given by

$$\mathbf{C}_{\mathbf{r}}^s = \mathbf{C}_{\text{sky}}^s + \mathbf{C}_{\text{beam}}^s + \mathbf{C}_{\text{N}}. \quad (7.3)$$

The thermal noise covariance is diagonal matrix with a variance given by Equation (5.6). Although its contribution is non-negligible, I will ignore it throughout this chapter because it dwarfs the other components on calibration timescales, and does not contain interesting spectral structure.

7.4 Redundant Calibration Errors

Redundant calibration consists of two steps; *relative calibration* and *absolute calibration* (Zheng et al., 2014). Following the notation of Byrne et al. (2019), I write the gain g_p of an antenna p as a product of the absolute gain G and a relative gain h_p

$$g_p = G \times h_p, \quad (7.4)$$

where the absolute gain is defined as

$$G = A \exp[i(\Delta + \Delta_x x_p + \Delta_y y_p)]. \quad (7.5)$$

Here, A is the degenerate absolute gain amplitude, Δ the degenerate phase offset, Δ_x the phase gradient in the x -direction, Δ_y the phase gradient in the y -direction, and x_p and y_p are the x - and y -position of antenna p , respectively (see Chapter 3 for more details). Relative calibration relies on the redundancy of the array to solve for the relative gains. Absolute calibration relies on external information, i.e. sky-based calibration, to solve for the absolute gain (Li et al., 2018; Byrne et al., 2019; Kern et al., 2019). The total gain g_p is a multiplication of the absolute and the relative gain. To understand how the relative calibration error (see Chapter 5 and Orosz et al., 2019) and the absolute calibration error (Byrne et al., 2019) propagate to the total gain, I assume the true gain solutions are

unity, and perturb the solutions in Equation (7.4).

$$\begin{aligned}
g_p &= Gh_p \\
(1 + \delta g_p) &= (1 + \delta G)(1 + \delta h_p) \\
\delta g_p &= \delta G + \delta h_p.
\end{aligned} \tag{7.6}$$

Here I ignored higher order terms between δh_p and δG , the perturbation to the relative and absolute gain, respectively. I relate the covariance of the total gain g_p to the covariances of the absolute and relative gain through

$$\begin{aligned}
\mathbf{C}_g &= \text{Cov}[\delta G + \delta h_p, \delta G' + \delta h'_p] \\
&= \mathbf{C}_G + \mathbf{C}_h,
\end{aligned} \tag{7.7}$$

where I assumed that the error in absolute calibration is independent from the error arising in relative calibration.

The errors that enter the absolute calibration step are exactly the same as for sky based calibration, e.g. sky modelling errors and beam modelling errors. To understand the impact of non-redundancies on redundant calibration I first need an analytic description. The relative calibration step in redundant arrays relies on the fact that identical baselines measure the same visibility. However, in reality it is impossible to position receiver elements on a perfectly regular grid, to build elements with identical responses, and additionally, the various elements might also suffer from pointing errors. Chapter 5 quantifies the positioning error of tiles in the redundant configuration of the MWA to have a variance of $\sigma_x = 2$ cm and Chapter 6 quantifies the fraction of MWA tiles with one broken dipole to vary between 25-40%.

In this analysis I consider these non-redundancies as perturbations to our true redundant visibility. This results in a measured visibility given by

$$V(\mathbf{u}, \nu) = \int Ib(b + \delta b)^* \times e^{-2\pi i(\mathbf{u} + \delta \mathbf{u}) \cdot \mathbf{l}} d^2 \mathbf{l}, \tag{7.8}$$

where δb and δu are the change in antenna response due to broken dipoles and baseline separation due to positioning errors, respectively. I linearise the exponential around \mathbf{u}

$$V(\mathbf{u}, \nu) = \int I b (b + \delta b)^* (1 - 2\pi i \delta \mathbf{u} \cdot \mathbf{l}) \times e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}, \quad (7.9)$$

and extract the perturbation to the visibility in a redundant group while ignoring higher order terms, i.e. products between the beam and position perturbation. This yields a visibility perturbation

$$\delta V(\mathbf{u}, \nu) = \int I (b \delta b^* - 2\pi i b b^* \delta \mathbf{u} \cdot \mathbf{l}) \times e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}. \quad (7.10)$$

For the analysis on the precision of gain estimates, and the impact of these errors on the power spectrum I consider these perturbations as residuals to our “redundancy model” and compute their covariance matrix $\mathbf{C}_{\mathbf{r}}^r(\mathbf{u}, \mathbf{u}')$. Here, I ambiguously choose \mathbf{u}' to denote different baselines at the same frequency, or the same baseline at different frequencies. The covariance matrix of the residuals is given by

$$\mathbf{C}_{\mathbf{r}}^r(\mathbf{u}, \mathbf{u}') = \text{Cov}[\delta V(\mathbf{u}, \nu), \delta V(\mathbf{u}', \nu')]. \quad (7.11)$$

In Appendix C I discuss the details of the derivation of this covariance matrix, and refer the interested reader for the details. Here, I only concern you with the result and interpretation. In general, I can separate the contribution to the error covariance matrix into three components, the first two are derived in Equation (C.10)

$$\mathbf{C}_{\mathbf{r}}^r = \mathbf{C}_{\text{beam}}^r + \mathbf{C}_{\text{position}}^r + \mathbf{C}_{\text{N}}. \quad (7.12)$$

Here, I define the error covariance matrix due to changes in the beam $\mathbf{C}_{\text{beam}}^r$ and due to position errors $\mathbf{C}_{\text{position}}^r$, and I include the thermal noise covariance for

completeness. The error covariance matrix due to changes in the antenna beam response is given by

$$\begin{aligned} \mathbf{C}_{\text{beam}}^r &= \mu_2 (f_0 f'_0)^{-\gamma} \int b^* b' \langle \delta b^* \delta b' \rangle e^{-2\pi i (\mathbf{u} - \mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\ &+ \mu_1^2 (f_0 f'_0)^{-\gamma} \iint b b^* \text{Cov}[\delta b^*, \delta b'^*] e^{-2\pi i (\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}'. \end{aligned} \quad (7.13)$$

where μ_1 and μ_2 are the mean and variance of all point sources above a detection threshold for the stochastic point source sky model, respectively (see Equation 6.9), $f_0 = \nu/\nu_0$, and γ is the index of the power law used to model the spectral energy distribution of each foreground source. Throughout this chapter I adopt a power law index of $\gamma = 0.8$.

The solutions for these integrals for the particular case of broken dipoles in an MWA tile are far from elegant and given in Equations (B.15-B.19). The beam covariance matrix for redundant calibration differs slightly from that of sky based calibration. For redundant calibration I do not distinguish between a modeled component and an unmodeled component. Hence, the second moment of the source counts considers all sources down to the confusion limit as “the signal”. I also do not have a term that tries to reduce the difference between the modelled and unmodelled component. For sky based calibration I derived the beam covariance matrix relative to an unmodelled sky as observed by a fully functional array. In relative calibration no attempt is made to model the sky and hence such a correction is not necessary. The error covariance matrix due to position offsets is in general given by Equation (C.22), and repeated here for the reader;

$$\mathbf{C}_{\text{position}}^r = 16\pi^3 \mu_2 (f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \Sigma_\nu^2 (1 - 2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|^2) e^{-2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|^2}, \quad (7.14)$$

where Σ_ν describes the FoV of the geometric mean of the two primary beam

responses of the two different measurements. I continue with the approximation of a chromatic Gaussian beam re-scaled to an Airy disk response (see Equation 6.24), adopting the MWA-tile size $D = 4$ m as the diameter of the aperture. This enables the relatively straightforward derivation of results as in previous chapters, but also neglects significant sidelobe response present in current low frequency radio interferometers. I omit the second term in Equation (C.22) because it decays quickly with baseline length. For baselines larger than the size of the tile diameter this term drops to zero fairly quickly and can be ignored.

Now that I have descriptions of the errors in both redundant and sky based calibration in place, I have all the tools to study the impact of these errors on our ability to calibrate our instrument and to estimate the 21-cm PS.

7.5 CRLB on Gain Estimates

To compute the error on the complex gains, I revisit the Fisher Information Matrix (FIM) and the Cramér-Rao Lower Bound (CRLB) to compute errors on the estimated antenna gains. I assume we are calibrating data per frequency channel and compute the gain error on a single frequency channel. Formally, the FIM for a multivariate normal distribution is defined as

$$\mathcal{I}_{ij} = 2\text{Re} \left(\frac{\partial \mathbf{m}^\dagger}{\partial \alpha_i} \mathbf{C}_{\text{data}}^{-1} \frac{\partial \mathbf{m}}{\partial \alpha_j} \right) + \text{tr} \left(\mathbf{C}_{\text{data}}^{-1} \frac{\partial \mathbf{C}_{\text{data}}}{\partial \alpha_i} \mathbf{C}_{\text{data}}^{-1} \frac{\partial \mathbf{C}_{\text{data}}}{\partial \alpha_j} \right). \quad (7.15)$$

Here, \mathbf{C}_{data} describes the noise on our data model \mathbf{m} , and how the noise correlates between different baselines. Note: this formal definition of the FIM differs from the form used in Chapter 5. Here, I include a term that accounts for parameters that depend on the noise, i.e. the visibilities in redundant calibration (see Section 7.5.2). The α_i are the parameters that we are trying to estimate. The estimated parameters differ per calibration technique and I discuss them separately for sky and redundancy based calibration in the next section. Having computed the FIM

I compute the CRLB by simply computing the inverse of the FIM

$$\text{Cov}[\alpha_i, \alpha_j] \geq \mathcal{I}_{ij}^{-1}, \quad (7.16)$$

if such an inverse exists (Stoica & Marzetta, 2001). Here, $\text{Cov}[\alpha_i, \alpha_j]$ is the lower bound on the error covariance of the estimated parameters α_i and α_j . The CRLB gives us the variance of the gain errors for an ideal estimator (which may not exist). I ignore that the amplitude and phase are affected differently, and hence the results are only indicative of the global behaviour of the gain solutions.

For the model \mathbf{m} , I revisit the approximation of Chapter 6 and assume that each baseline measures a visibility with amplitude

$$m = f_0^{-\gamma} \sqrt{\mu_2}, \quad (7.17)$$

This threshold differs for sky and redundancy based calibration. In redundant calibration I consider all sources above the confusion noise, whereas in sky based calibration I only consider sources above the sensitivity limit of current surveys. Choosing a sky model description like this effectively describes the entire sky as a single point source, which has an effective flux density equal to the root-mean-square of the total sky flux density. As discussed and demonstrated in the previous chapter, the theoretical mean visibility for such a sky would yield zero for all baselines larger than the tile size, which is nonphysical.

7.5.1 Sky Based Calibration Errors

For sky based calibration the estimation parameters are N gain parameters, i.e. $\alpha = \{\mathbf{g}^T\}$. Here, \mathbf{g} is the vector containing all the antenna gains. The gain parameters are independent of the noise of the unmodelled sources and independent of the beam modelling errors, so I can safely drop the second term in Equation (7.15). In Chapter 6 I chose a slightly pessimistic detection limit of the foreground of 1 Jy. Here, I choose a more optimistic model where calibration is

performed on catalogued sources down to $S_{\min} = 100 \text{ mJy}$ (Carroll et al., 2016; Hurley-Walker et al., 2017). This corresponds to a model brightness of $m \sim 30 \text{ Jy}$.

For a redundant array that is calibrated through sky based calibration I expect significant correlation between visibilities of different baselines due to unmodelled sources. This means that structurally $\mathbf{C}_{\text{sky}}^s$ contains significant off-diagonals that describe the covariance between redundant baselines. For a confusion limit of $\sim 1 \text{ mJy}$ (Franzen et al., 2016; Offringa et al., 2016) and the adopted detection limit, both the diagonal and off-diagonals are about $\sim 4.5 \text{ Jy}^2$. Beam errors are mildly correlated between baselines that share the same broken tile, however, when those baselines are separated by more than the primary beam size in the uv -plane the correlation becomes negligible. I therefore approximate $\mathbf{C}_{\text{beam}}^s$ as a diagonal matrix, i.e. for a redundant array I ignore the correlation with all other baselines. Beam modelling errors contribute an additional $\sim 0.5 \text{ Jy}^2$ to the data covariance matrix. The total data covariance matrix, however, is still very non-diagonal. This prevents the provision of a simple analytic expression of the CRLB, because it depends highly on array lay-out.

7.5.2 Redundant Calibration Errors

I compute the variance for the absolute gain \mathbf{C}_G in the same way as I did for sky based calibration errors, except now I only solve for 1 gain parameter, i.e. the absolute gain amplitude. Byrne et al. (2019) show that the error budget for absolute calibration errors is dominated by the absolute gain amplitude, so this is a fairly safe approximation. I assume the same errors, i.e. sky and beam model based errors, and model signal.

To compute the covariance of relative calibration errors \mathbf{C}_h I use a similar expression for the signal model, however, I now include sources of all fluxes for the modelled sky r.m.s. In practice, we only consider sources beyond the confusion limit and hence going down to $S_{\min} = 1 \text{ mJy}$ suffices. This translates into practically the same sky brightness of 30 Jy . Unlike sky based calibration re-

redundant calibration solves for N gain parameters + M redundant visibilities. So $\alpha = \{\mathbf{g}^T, \mathbf{V}^T\}$, where \mathbf{V} is the vector of unique visibilities. We solve for these parameters subject to position and beam variations, that contribute $\sim 1.5 \text{ Jy}^2$ and $\sim 0.5 \text{ Jy}^2$ to the residual covariance matrix, respectively. Apart from the different covariance matrices involved in the computation of the FIM, a significant difference is the role played by the second term in Equation (7.15). Redundant calibration also solves for the redundant visibilities themselves that depend on the sky signal; their covariance is captured in \mathbf{C} . Hence, to compute the FIM for redundant calibration I need to take this second term into account, without it the FIM is not invertible. The second term is only computed for the elements in the FIM describing the unique redundant visibilities. The covariance matrix is proportional to the square of our model signal, hence, I approximate all relevant entries of $\partial \mathbf{C}_{\text{data}} / \partial \alpha$ as $2f_0^{-\gamma} \sqrt{\mu_2}$ for $\alpha = \{\mathbf{V}^T\}$.

7.5.3 Comparison of lower bounds

Figure 7.1 shows the CRLB computed for a hexagonal array with an increasingly larger antenna number for a fixed hexagonal lattice defined by a shortest spacing of 14 m. I also compute the CRLB for the MWA, two deployment stages of HERA, and the SKA (Dewdney et al., 2015). However, for simplicity I maintain an MWA-like FoV of 30° , fraction of broken antennas of 15%, position error of 2 cm, and sky modelling depth of 100 mJy for all instruments, i.e. these results are only sensitive to the varying uv -coverage and the total number of antennas for the different telescopes. For completeness I include a computation of the CRLB for a thermal noise-only case for both redundant and sky based calibration for observations with a bandwidth of 40 MHz, integration time of 2 minutes, and $SEFD = 20 \cdot 10^3 \text{ Jy}$. The results for the different sources of error all scale with $\sim 1/\text{sqrt}(N_a)$ and clearly illustrate that indeed these errors are much smaller than thermal noise ($\sim 9 \text{ Jy}$) on a calibration time scale of 2 minutes.

For redundant calibration we see that the error of relative calibration is slightly

larger than that of absolute calibration. The difference can be explained as follows. The calibration signal is practically the same for the two calibration steps. However, even though the sky model error in itself is larger than the total error of relative calibration, absolute calibration only solves for one parameter and the CRLB assumes that these errors average down.

The results show that the MWA and both stages of HERA outperform the hexagonal array. The contribution to the total gain error is almost equal for both relative and absolute calibration. Hence, the difference between our real telescope and our hypothetical one can only be explained by the impact of the array lay-out on absolute calibration. Both HERA and MWA have reduced the redundancy compared to the hexagonal array. HERA is divided into three rhombuses and the MWA divides its redundant component over two hexagonal sub-arrays. This reduces the number of off-diagonal terms in the sky covariance matrix, reducing the error in the absolute calibration step.

For sky based calibration, we see a similar trend. HERA and MWA outperform the hexagonal array reinforcing the idea that reducing the redundancy improves any sky based approach. Interestingly that does not seem to be the case for the SKA. The core of the SKA suffers from a large number of similar baselines. These baselines have highly correlated sky model errors that potentially degrade calibration quality. Comparing the errors between redundant calibration and sky based calibration indicates that the errors for sky based calibration are half an order of magnitude lower than that of redundant calibration for fiducial parameters. For the expected sky modelling limits of upcoming surveys with the MWA ($S_{\min} \sim 10$ mJy), this difference becomes even larger.

7.6 Spectral Structure of Gain Errors

The CRLB only gives us an indication of the precision of redundant calibration and sky based calibration. To truly understand the impact of calibration systematics on 21-cm PS estimation, we need to understand their accuracy, and the

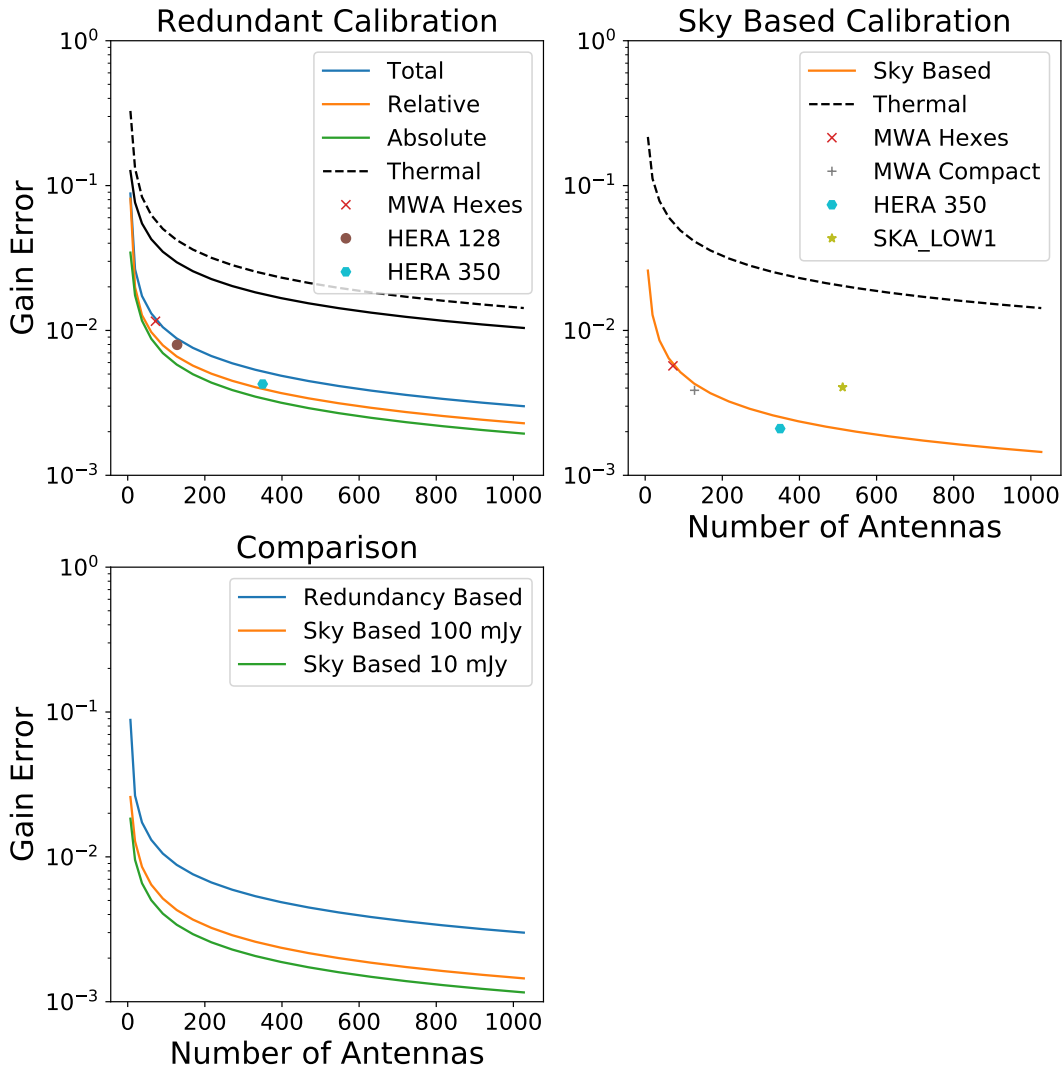


Figure 7.1: *top left*: The standard deviation of the gain errors (square root of the Cramér-Rao Lower Bound) for Redundant Calibration. In green the error from the relative calibration step, in orange the errors from the absolute calibration step, in blue the total error, and the black dashed line indicates the variance due to a thermal noise of 10 Jy alone in the relative calibration step. *top right*: The standard deviation of the gain errors for Sky Based Calibration. In orange the gain error for sky based calibration with a sky model depth of 100 mJy and the black dashed line indicates the expected error for thermal noise alone. *bottom*: A comparison between errors for redundant calibration (blue), sky based calibration with a sky model depth of 100 mJy (orange), and a sky model depth of 10 mJy (green).

resulting spectral structure of gain errors. To do so I revisit the “averaged” gain error concept from chapter 6 as an approximation that enables propagation of error covariance matrices into the PS of residuals after calibration. In the previous chapter I assumed the gain errors are small, and that the error between different baselines is uncorrelated. This formed the basis for the “averaged” gain covariance defined in chapter 6 and repeated here (see Equation 7.18). Instead of computing the gain per antenna, propagating that forward to a per-baseline error, and then into the EoR PS, I compute the gain for a particular u -scale in which several baselines and corresponding antenna gains have been binned. This averaged gain covariance is defined as

$$\mathbf{C}_{\mathbf{g}}(u, \nu, \nu') = \frac{(f_0 f'_0)^\gamma}{\mu_{2,m}} \sum_n w^2(u'_n) \mathbf{C}_{\mathbf{r}}(u'_n, \nu, \nu'). \quad (7.18)$$

The weighting function $w(u_n)$ determines the contribution of uncalibrated measurements at a particular scale u' to the averaged gain for scale u' . After calibration, I apply these gain solutions to our data, and subtract off a sky model. In general, this is the same sky model that is used in sky based calibration. Under the assumption that the errors on our gain solutions are independent of the sky model, and residuals at that particular u -scale¹, I write the covariance of the calibrated residuals $\mathbf{C}_{\hat{\mathbf{r}}}$ in terms of the gain error covariance matrix $\mathbf{C}_{\mathbf{g}}$, sky model covariance matrix $\mathbf{C}_{\mathbf{m}}^s$, and sky residual sky covariance matrix $\mathbf{C}_{\mathbf{r}}^s$ as follows

$$\mathbf{C}_{\hat{\mathbf{r}}} = 2\mathbf{C}_{\mathbf{g}} \odot \mathbf{C}_{\mathbf{m}}^s + (1 + 2\mathbf{C}_{\mathbf{g}}) \odot \mathbf{C}_{\mathbf{r}}^s. \quad (7.19)$$

An important difference between sky based calibration and redundant calibration is the interaction between the various models and residuals. These are discussed in the next sections.

¹Note: although counter-intuitive I made this assumption in Chapter 6 to a) simplify our expression, and b) because the gain error is sum of errors at different u -scales.

7.6.1 Sky Based Calibration

For sky based calibration the unmodeled sky and beam errors that make up the residuals in the averaged gain error, see Equation (7.18), are the same as the residuals that are left in the data after subtraction in Equation (7.19). The 2D-PS structure of these residuals before calibration is shown in Figure 6.2. I then have to assume a weighting function to describe how uncalibrated baselines contaminate their calibrated counterparts and appropriately sum over all u -scales that participate in calibration. I revisit the weighting function used in the previous chapter and detail our reasoning so we can understand how this might differ for redundant calibration in the next section.

In chapter 6 I linearised Equation (7.1), and wrote it in matrix form $\mathbf{y} = \mathbf{A}\mathbf{x}$. Here, \mathbf{y} contains the ratio between residual and model visibilities, \mathbf{x} contains the linearised gain errors, and \mathbf{A} is a matrix that maps each antenna to a baseline. A solution for the gain error can then be obtained using least squares. The matrix \mathbf{A}^{-1} contains the weights describing how the residuals in each baseline contribute to the gain error of each antenna. These weights are on order $\sim 1/(N_a - 1)$, where N_a denotes the number of antennas, N_a . To compute the contribution of an uncalibrated visibility to a calibrated one, I take the baseline weights for each gain error contained in the rows of \mathbf{A}^{-1} and add them in quadrature. The contents of this resulting weight matrix reveal, not unsurprisingly, that each calibrated baseline “sees” the residuals of all baselines formed by the two antennas. I reorder the rows and columns of the resulting weight matrix according to baseline length, and bin them into the same bins used in PS estimation. The results can be reasonably approximated by

$$w(u) = \frac{1}{N_a - 1} \frac{N_b(u)}{N_b(\text{total})}, \quad (7.20)$$

where $N_b(u)$ is number of baselines in bin u .

For pseudo-random arrays we can indeed safely assume that errors between different baselines are uncorrelated. This is not necessarily true for redundant

arrays, as mentioned earlier. Ewall-Wice et al. (2017) investigate the impact of redundancy on the propagation of errors in sky based calibration and find that this correlation between baselines has a small impact on scales relevant to the EoR. Under that assumption I compute the frequency covariance matrix of the averaged gain error through a weighted sum of the residuals of the calibration at each baseline bin, and ignore additional correlation between closely spaced baselines.

7.6.2 Redundant Calibration

I compute the covariance matrix of the relative gain error in the same way as for sky based calibration. The residual covariance matrix \mathbf{C}_r in Equation (7.18) for relative calibration is given by $\mathbf{C}_r = \mathbf{C}_{\text{position}}^r + \mathbf{C}_{\text{beam}}^r$. Figure 7.2 shows the 2D-PS structure of the position error, beam error, and total error covariance matrices for relative calibration. Both contributions to the relative calibration error contain




Figure 7.2: Two-dimensional power spectrum structure of the residuals for relative calibration. *Left*: The 2D-PS of the position error covariance matrix. *Middle*: The 2D-PS of the beam variations error covariance matrix. *Right*: The 2D-PS of the sum of the contributions to the total covariance matrix.

a wedge like feature due to the foregrounds on which these non-redundancies are imprinted. For the position errors the power is concentrated along the edge of the wedge. This hints on the importance of the size of the primary beam and additional contamination due to sidelobes that are not included in this analytic

model. The contribution of beam variations follows a classic wedge like feature, and contains more power than the position errors.

I only need to find the appropriate weighting $w(u)$ to relate uncalibrated baselines to their calibrated counterparts in redundant calibration. For simplicity we choose to adopt the same weighting function as for sky based calibration, the difference being that now we only consider redundant baselines. Figure 7.3 shows the calibrated residuals after relative calibration and sky model subtraction for an array with uniform uv -coverage, and otherwise fiducial MWA parameters.




Figure 7.3: Residuals in the 21-cm power spectrum after relative calibration for an array with uniform uv -coverage, and MWA parameters for FoV, position tolerances and beam variations. *Left*: Residuals after relative calibration with position errors. *Middle*: Residuals after relative calibration with beam variations. *Right*: Residuals after relative calibration with both position errors and beam variations.

The resulting calibration residuals seem very similar to those of sky based calibration presented in the previous chapter. The 2D-PS shows increased power in the foreground wedge, and the window contains residual power that was not there before calibration (see Figure 7.2). A key difference is that the contamination in the window does not extend as far along the k_{\parallel} -axis as compared to our previous sky based results. The introduced power drops off around $k_{\parallel} \sim 0.2 \text{ hMpc}^{-1}$. The amount of power introduced into the window seems to be in line with results from Orosz et al. (2019), however, our analytic model does not capture the same structure of the contamination. Their simulations indicate contamination at both the lowest and highest k_{\perp} -modes hinting at a flaw in our weighting function. I

revisit this issue when I apply our analysis to HERA and the MWA in the next section.

Absolute calibration can be applied in various ways, and one particular strategy involves solving for the absolute parameters using sky based calibration solutions. Byrne et al. (2019) study the approach where the absolute amplitude is found by averaging the gain solutions over the different antennas. Naively, one would expect to compute the covariance matrix of the absolute gain error by simply summing over the averaged gain covariance matrices of sky based calibration in each bin, dividing by the number of bins, and then dividing by the number of antenna gain solutions. The first two operations attempt to capture spectral structure that is expected to differ per antenna, depending heavily on the baselines it forms within the array. Central antennas form many short baselines, whereas remote antennas form many long baselines. The latter division aims to account for the down averaging of the error over antennas. However, if I take this proposed gain covariance matrix to compute the residuals after calibration using Equation (7.19), I find that the residual power in the window is 2 orders of magnitude lower than the result from Byrne et al. (2019). This indicates that the absolute gain errors do not average out, and that the covariance matrix of absolute gain errors is simply a weighted average of the sky based gain errors, without the down averaging over antennas.

Figure 7.4 shows the residual power in the contamination along a 1D-slice at $k_{\perp} = 2 \cdot 10^{-2} h\text{Mpc}^{-1}$ in the 2D-PS due an absolute gain error applied to a baseline of 7 m. The results globally show similar structure to those presented by Byrne et al. (2019). However, the amplitude of the residual power presented here for the error due to sky modelling drops off faster around the $k_{\parallel} = 0.1 h\text{Mpc}^{-1}$ mode and is half an order of magnitude lower. Although I adopt similar sky modelling parameters, this analytic model lacks sidelobe response. Artificially increasing the size of the FoV, or changing the scaling factor ϵ that re-scales a Gaussian beam to the main lobe of an Airy beam increases contaminating power.




Figure 7.4: Residuals in the 21-cm power spectrum at $k_{\perp} = 2 \cdot 10^{-2} \text{ hMpc}^{-1}$ after absolute calibration for an array with uniform uv -coverage, and MWA parameters for the FoV, sky modelling errors, and beam modelling tolerances. *Left*: Residuals after absolute calibration with sky modelling errors. *Middle*: Residuals after absolute calibration with beam modelling errors. *Right*: Residuals after absolute calibration with both sky and beam modelling errors. The black line is a fiducial EoR model from Mesinger (2016).

This indicates that the lack of sidelobes explains the difference in resulting power, and indicates a weakness in this analytic model.

The analytic results presented here broadly align with simulations from (Orosz et al., 2019) and (Byrne et al., 2019). Although there are some difference in the structure or the amplitude of the expected contamination, the analytic model reproduces results from detailed simulations fairly well. These differences are not too surprising as this model does not describe the complexities of sidelobe structure in the antenna response, and the exact mapping of antennas to baselines. With a description for both absolute and relative calibration errors I now have all the ingredients to compare redundant calibration to sky based calibration, and the next section discusses the results when the analytic model is applied to both HERA and the MWA.

7.6.3 Comparing Power Spectrum Contamination

I use the above recipes to compute the expected level of contamination in the 21-cm PS due to calibration errors. The previous results in this chapter assumed an unrealistic uniform uv -coverage. In this section I adopt the design antenna po-

sitions of HERA-128 and deployed antenna positions of MWA Phase II compact, both having 128 antennas, to compute their uv -coverage. For the MWA I adopt the same modelling, and deployment parameters as mentioned before. The same parameters are adopted for HERA, except for the diameter of the aperture. This is changed to match the diameter $D = 14$ m of the HERA dishes. Figures 7.5 and 7.6 show the results for MWA Phase II Compact and HERA-128, respectively. The contours denote ratios between residual power and the fiducial EoR signal discussed in the previous chapter.




Figure 7.5: Expected levels of contamination in the 21-cm EoR PS due to residuals of calibration models for MWA Phase II Compact. *top*: residuals due to sky based calibration. *bottom*: Residuals due to redundant based calibration, including both relative and absolute calibration. Overlaid are ratios between residuals and a fiducial signal, where the residuals have 1, 10, 10^2 times more power than the fiducial signal. For the MWA it seems that the amount of contamination for current modelling and array deployment parameters is comparable.

For the MWA the results are reminiscent of those presented in the previous chapter. Although, for both redundant and sky based calibration the extent of contamination in the window along the k_{\parallel} is significantly reduced as compared to

those earlier results. These changes can be attributed to a change in the sky model depth. Although there is some difference in the structure of the contamination in the window, overall the level and the behaviour of the contamination is very similar between the two calibration methods. Errors in both redundant and sky based calibration reduce the usable region of the power spectrum to lie beyond $k_{\parallel} \sim 2 \times 10^{-1} h\text{Mpc}^{-1}$.




Figure 7.6: Expected levels of contamination in the 21-cm EoR PS due to residuals to calibration models for HERA-128. *top*: residuals due to sky based calibration. *bottom*: Residuals due to redundant based calibration, including both relative and absolute calibration. Overlaid are ratios between residuals and a fiducial signal, where the residuals have 1, 10, 10² times more power than the fiducial signal. I have omitted contour labels because of their strong overlap.

HERA-128, has a comparable number of baselines. However, it differs in uv -coverage and its FoV is 3 times smaller. Overall the wedge for HERA has less power and a lower inclination angle that can both be explained by the small FoV. Contamination in the window is practically absent, apart from contamination due to the first sidelobe of the Blackman-Harris window applied along the frequency direction. This clearly highlights the limitations of our analytic model, in particular our approximation of the beam response with a Gaussian function.

7.7 Discussion

I have presented two analyses of the impact of errors during calibration. The first set of results studies the precision of calibration through the CRLB. The variances presented here may not be representative, because the error contributions, apart from the thermal noise, are not truly Gaussian distributed with a zero mean. The visibility contribution from unmodelled foregrounds is known not to be Gaussian distributed (Ollier et al., 2017). The contribution from beam modelling errors are most likely not Gaussian, and certainly do not have a zero mean. Residuals due to positioning errors are potentially Gaussian distributed. Our approximation in Equation (7.14) shows that actual antenna deployment errors translate linearly into a visibility error. Calibration is therefore subject to errors with distributions that bias the calibration solutions. A more appropriate analysis of the errors in calibration comes through the biased CRLB (Barrett & Myers, 2004). Its drawback is that an expression for the bias needs to be derived in the first place. This is not a trivial task.

The CRLB on the other hand provides a simple and straightforward way to propagate errors under the assumption of a model and provides an indication of the performance of different calibration strategies. The results presented in this chapter indicate that redundant calibration is limited by the relative calibration step. This comes as a mild surprise because the visibility errors involved in relative calibration are smaller in amplitude than for sky model calibration. In ad-

dition the errors in sky based calibration are highly correlated between redundant baselines. This indicates that the smaller number of parameters estimated using sky based calibration offsets the lack of an accurate sky model. However, in my analysis I did not include a description of diffuse foregrounds. These foregrounds contain significant power on baselines relevant for EoR science, and including those would provide a more realistic prediction of calibration precision.

The CRLB indicates that sky based calibration outperforms redundancy based calibration by half an order of magnitude. Sky based calibration has a historical advantage over redundancy based calibration. The fiducial modelling parameters I adopted are the result of more than half a decade of work on the modelling of foregrounds in the Southern sky for the MWA (Carroll et al., 2016; Offringa et al., 2016; Hurley-Walker et al., 2017). The upcoming GaLactic and Extragalactic All-sky MWA Extended Survey (GLEAM-X)² and Long Baseline Epoch of Reionisation Survey (LoBES) that use the enhanced imaging capabilities of the extended configuration of MWA Phase II (Wayth et al., 2018) will continue to provide improvements. Similarly, modelling of the beams of MWA-tiles informed by measurements (Neben et al., 2015; Line et al., 2018) or through advanced electromagnetic simulations (Sokolowski et al., 2017) continue to provide more accurate beam models. These developments signify a difference in the improvements that can be made to sky and redundancy based calibration. Sky and instrument models can be improved through measurements, detailed modelling and arrays expansions. However, there is a limit to the precision at which arrays can be deployed with redundancy. Hence, augmentations to redundant calibration as proposed by Liu et al. (2010) or alternative paths to combine redundant and sky based information proposed by Sievers (2017) are an unavoidable path forward.

Although precision gives an estimate on the performance of different calibration strategies, accuracy is what stands in the way between us and a detec-

²<http://www.mwatelescope.org/gleam-x>

tion of the 21-cm signal. In particular, bias with spectral structure introduces contamination in an otherwise pristine EoR window. To compute the expected contamination for fiducial array deployment, and sky and instrument modelling parameters I made simplifying assumptions resulting in the 2D-PS presented in Figures 7.5 and 7.6. Those results show that the MWA suffers more from calibration errors than HERA. The contour lines indicate ratios between the computed residuals and a fiducial EoR signal. Close to the wedge the contours indicate that the residuals have 100 times more power than the signal and it is only around $k_{\parallel} = 0.2 h\text{Mpc}^{-1}$ that the residuals drop off in power to equal our EoR model signal. This severely reduces the usable region in the window because the contamination in these modes does not average down.

In stark contrast, the results for HERA indicate it is completely robust against calibration errors. Even if I adopt extreme position errors or a higher fraction of broken elements, the results do not change drastically. The planned 128 element deployment of HERA has the same number of antennas as MWA Phase II compact, yet covers a more compact region of the uv -plane than MWA phase II. The latter reduces contamination from long baselines that contain most of the spectral structure (Ewall-Wice et al., 2017). However, a major difference with the MWA it is the narrow FoV of the HERA dishes. The narrow FoV offers some suppression of contamination due to either sky based or redundancy based calibration. However, the lack of any contamination presented here indicate that our results in general are too optimistic.

There are some improvements that can be made to the analytic model. The Gaussian beam approximation strongly suppresses any horizon response. This completely ignores the sidelobe response in real instruments that covers an additional portion of the sky. Modelling the main lobe alone captures only part of the contamination. For the MWA the large main lobe introduces enough spectral structure that it appears in our calibrated residuals. However, the Gaussian approximation seems to break down completely for narrow FoVs for instruments

such as HERA. Orosz et al. (2019) model the HERA beam with an Airy beam which has non-negligible sidelobe response, as they find that a Gaussian unrealistically suppresses the foreground wedge. Including a more realistic yet analytically solvable beam model would significantly improve these results.

Additionally, a more accurate weight function could be devised to approximate the contamination introduced from uncalibrated baselines into their calibrated counterparts. To find an approximation for the weighting function I could consider the `lincal` algorithm from Liu et al. (2010) and its explicit solution in matrix form $\mathbf{y} = \mathbf{A}\mathbf{x}$ from Li et al. (2018). Studying its properties following similar steps as in Chapter 6 could provide the insight that enables a more accurate weight function for redundant calibration. Despite these necessary improvements, the overall framework seems very suitable to study the propagation of errors if analytic expressions for those errors exist. Together with simulations and real data, these analytic models provide us with the deep understanding of the obstacles we need to overcome in our efforts to detect the elusive EoR signal.

7.8 Conclusions

In this chapter I aimed to investigate the limitations of redundant and sky based calibration. I computed both the CRLB and 21-cm PS contamination due to errors in the calibration model. Redundant calibration is primarily limited by relative calibration due to the position precision during array deployment, and variations in the antenna response. Sky based calibration is limited by sky and beam modelling errors. Both the CRLB and the expected PS contamination indicate that in wide-field arrays sky based calibration is marginally more robust against calibration errors. Additionally, sky and beam models can be improved, albeit with significant effort. Non-redundancies on the other hand are inherently built into the the array. Design choices that mitigate non-redundancies are therefore crucial. However, even more important is the development of hybrid calibration methods that employ both redundancy and sky based information.

Summary and Conclusions

The Epoch of Reionisation (EoR) holds the key to understanding the transition from a structureless Early Universe into the diverse and structured Universe we live in today. The redshifted 21-cm line from neutral hydrogen provides us with a direct probe into the state of the intergalactic medium (IGM) during the EoR, and enables us to study both astrophysics and cosmology at the same time. In the absence of extremely large telescopes that provide us with the ability to directly image the 3-dimensional structure of the IGM through cosmic time, we rely on the 21-cm power spectrum (PS). This enables us to further constrain astrophysical and cosmological parameters, and shed light on a poorly studied time in the history of the Universe. The 21-cm PS is an extremely powerful tool for a statistical detection of the redshifted hydrogen signal. It provides increased sensitivity compared to direct imaging, while retaining information about the spatial structure of the signal. However, the extreme dynamic range of the EoR experiment, with foregrounds being 5-6 orders of magnitude brighter than the EoR signal, puts stringent requirements on the amount of control we need to exert on the systematics that plague the experiment. In this thesis I investigated the limits of current calibration techniques subject to realistic systematics based on current experience with the latest generation of low frequency radio interferometers.

8.1 Position Errors in Redundant Calibration

The most common form of antenna gain calibration is the sky based calibration method. However, we now understand that performing calibration in the absence of a complete sky model imprints spectral structure onto the antenna gain solutions that leads to contamination in the 21-cm PS. Redundant calibration aims to provide an alternative to sky based calibration in the absence of a complete model for the foreground sky through so-called relative calibration. By sacrificing instantaneous uv -coverage, we can oversample certain Fourier modes through a regular array lay-out. This oversampling enables us to calibrate a radio interferometer without the need for a sky model. However, building a perfectly redundant telescope is impossible. In Chapter 5 I investigated the impact of position errors on relative calibration algorithms and their ability to reliably constrain the antenna gains. I discussed the results of simulations of relative calibration of the hexagonal substations of the Murchison Widefield Array (MWA) where the tiles are displaced with positions errors similar to the actual displacements in the deployment of MWA phase II. The simulations showed that bright radio sources exacerbate the errors introduced by position errors, and that relative calibration performs optimally when bright sources are located at the pointing centre or in regions on the sky that are suppressed by the primary beam response. I also found that calibrating on the same bright source alone with a sky based approach results in less precise and accurate gain solutions. The results highlighted the importance to understand the sky modelling depths, and array deployment tolerances at which redundancy based calibration outperforms sky based calibration and vice versa.

8.2 Beam Variation in Sky Based Calibration

In Chapter 6 I investigated the impact of beam variations in the MWA on sky based calibration methods. The MWA tiles are built out of 16 dipoles arranged

on a 4×4 grid, and these dipoles can fail over time. I collated data on the number of broken dipoles in the MWA over the course of its lifetime since 2013, and found that the number of tiles with at least one broken dipole varies between 15% to 40%. Based on these results I developed an analytic framework that models the impact of the introduced error when the change in response is not properly taken into account, and translates that into contamination in the 21-cm PS after calibration. After confirming that this framework reproduced work from both earlier simulations and theory, I estimated the expected additional contamination due to beam modelling errors in the MWA. A median value for the fraction of broken tiles translated into 30% of the data having non-ideal beams. The results showed that under these circumstances $\sim 10^3 \text{ mK}^2 h^{-3} \text{ Mpc}^{-3}$ of additional power is added into the EoR window. This is 3 orders of magnitude lower than the current lowest limits, however, it still poses a challenge to the success of the EoR experiment as it limits the amount of usable information in the PS. This contamination is most likely our next limiting factor in the EoR experiment, and motivates further quantification of beam variation through in-situ measurements.

8.3 Comparing Redundant and Sky Based Calibration

In Chapter 7 I tied the results from the two earlier chapters together by providing a comparison between redundant and sky based calibration. I revisited the Cramér-Rao Lower Bound (CRLB) and the analytic framework to estimate power spectrum contamination. This chapter presented a derivation of the covariance matrix describing the residuals due to position errors, and discussed the beam variations covariance matrix in context of redundant calibration. The covariance matrices feed into the CRLB for relative calibration, absolute calibration, and sky based calibration. In the limit of extremely large regular arrays, the CRLB showed that sky based calibration outperforms redundant calibration

for fiducial sky modelling limits, and instrument deployment tolerances of the MWA. However, at best the CRLB provides an estimate of the precision of the different calibration strategies. For the EoR experiment accuracy, in particular spectral accuracy is crucial to its success. It is well known that spectral structure introduces contamination that biases a detection of the 21-cm PS signal. I repurposed the analytic framework from Chapter 6, and employed the covariance matrices that describe sky modelling errors, beam modelling errors, and antenna position errors to compute the spectral structure of the residuals after redundant and sky based calibration. The results showed that for the MWA errors in redundant calibration marginally contaminate more modes than sky based calibration. However, when the model was applied to the Hydrogen Epoch of Reionisation Array (HERA) it indicated that contamination is completely absent. This highlighted a flaw in the Gaussian beam model chosen in Chapters 6 and 7. Although a Gaussian beam describes the main lobe of the beam response fairly well, it lacks sidelobe response. This implies that the framework presented in this thesis underestimates contamination due to calibration. This became evident when I attempted to model the narrow Field of View (FoV) of HERA. The issue can be fairly easily resolved using an Airy function that contains moderate sidelobes, and was used successfully in earlier studies. The results do reiterate a few key points.

- For current sky modelling limits, antenna position precision, and antenna failure rates, sky based and redundancy based calibration seem to introduce contamination on similar levels.
- For redundancy based calibration, beam variations provide most of the contamination.
- The results reinforce the idea that narrow FoV arrays are most robust against calibration errors.
- Finally, the comparable errors in the MWA motivate the development of

calibration avenues that employ both redundancy and sky based information.

8.4 Conclusion and Outlook

Redundancy based and sky based methods provide alternative yet complimentary avenues to calibrate radio interferometric data. At current modelling and deployment parameters for MWA Phase II we expect the calibration error and subsequent power spectrum contamination to be on comparable levels for sky based and redundant calibration. However, this is subject to a few conditions. Redundant calibration performs best in the absence of bright radio sources or when those source are located at phase center. This implies that for the MWA redundant calibration is particularly suited for the EoR0 field, and less so for EoR1 containing Fornax A, and EoR2 that neighbours the Galactic Plane. Sky based calibration outperforms redundancy based calibration in the presence of bright sources, subject to accurate modelling of those sources.

The analytic framework presented in this thesis can be improved with some simple adjustments that could dramatically improve its performance. Inclusion of an Airy beam that describes non-negligible sidelobe response enables the propagation of additional and realistic spectral structure. Improving the weighting function to capture the contribution from various baselines to calibration errors more accurately enables a more reliable prediction of the contamination of relevant power spectrum modes. Finally, a full baseline-frequency covariance treatment and propagation into PS-space enables a derivation of a simple scaling relationship that directly translates modelling and deployment errors into PS contamination.

Combining analytic frameworks, with detailed simulations and real telescope data provides the deepest insights into the limitations that prevents us from detecting a signal from the reionisation epoch. Real data provides us with the problems, simulations and analytic analysis provide models to enable us to reliably

predict contamination and interpret the observed contamination. The results in this thesis further motivate the development of diffuse models for Galactic emission and large extended sources, and work to understand the behaviour of the antennas after deployment. These steps are crucial in an effort to improve our ability to use sky models for calibration. It is intrinsically hard to improve the redundancy of the array. One can reasonably expect errors in the placements of antennas at 1% of the element size, and work in this thesis demonstrates that centimetre precision has the potential to ruin a signal detection. Similarly, enforcing uniformity in the antenna responses of the array is incredibly difficult in the harsh remote locations that host our instruments. Accurate flagging of elements at a daily cadence is required to be able to isolate data that are corrupted, for both sky and redundancy based calibration. However, the most promising avenue forward is an approach that incorporates both sky and redundancy based information. Approaches that account for non-redundancies have been put forward and the next steps are to implement, test and quantify their performance.

8.5 Closing Remarks

The EoR experiment is extremely challenging. We now have over half a decade of experience with the latest generation of radio interferometers, and staggering amounts of data. Unlike the mathematics and analytic calculations presented in this thesis the real world and its data are truly messy. The mess provides challenges we can tackle from many avenues, and with the construction date of the Square Kilometre Array coming closer than ever, we can expect many more petabytes of exciting mess ahead. I hope that this thesis and the work presented here will prove to be useful while we try and make sense of the mess and the beautiful Universe we live in.

Appendices

Circumventing Phase Wrapping in Redundant Calibration

This Appendix is a reproduction of the appendix in Joseph, R. C., Trott, C. M. and Wayth, R.B., ‘The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions’, 2018, *The Astronomical Journal*, 156, 285, <https://doi.org/10.3847/1538-3881/aaec0b>. It differs from the original due to minor alterations, e.g. referencing to earlier chapters to ensure consistency within this thesis.

A.1 The Multi-Frequency Implementation

Redundant calibration is typically presented on a channel-by-channel basis in contrast with standard model based calibration schemes that operate over a range of frequencies. The multi frequency approach uses all of the information available in other frequency channels. In this section we will discuss the multi-frequency implementation of `logcal`, allowing it to benefit from the multi-frequency information available in radio interferometry data. Another motivation is that a multi-frequency implementation has the prospect of resolving the phase wrapping problem. Earlier, we described that phase wrapping occurs when a specific

baseline u observes a source on a specific coordinate $l = 1/2u$, the phase becomes ill-defined at this point. That same baseline should measure a defined phase when observing that same source at a different frequency. This property has motivated us to extend the classical redundant calibration framework by solving for the gain and visibilities simultaneously at different frequency channels. We will assume that the gain is the same at those frequencies, as a first order approximation, but the visibilities are different.

We extend Equation 3.21, by stacking the measurement vectors at different frequencies and adding the visibilities from the additional frequency channels. We then construct the matrix accordingly by realizing that \mathbf{A} can be split into two components:

$$\mathbf{A}_g = \begin{pmatrix} -0 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ -1 & 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 0 \end{pmatrix} \mathbf{A}_v = \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{pmatrix}, \quad (\text{A.1})$$

where \mathbf{A}_g maps the gains onto the measurements and \mathbf{A}_v maps the visibilities onto the measurements. We can construct a multi frequency matrix combining \mathbf{A}_g and \mathbf{A}_v , e.g. for a two-channel solutions estimation.

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_g & \mathbf{A}_v & \mathbf{0} \\ \mathbf{A}_g & \mathbf{0} & \mathbf{A}_v \end{pmatrix} \quad (\text{A.2})$$

Using this extended version of redundant calibration we return to the 5-element interferometer, while varying the number of channels involved to calibrate our antennas. Figure A.1 shows the results when we attempt to calibrate 5 antennas using 2 frequency channels spaced around the actual frequency channel we are

interested in. We observe a significant change in the structure of the variance of the phase solutions. The peaks around the phase wrapping points have severely decreased in width. However, it has not resolved the the phase wrapping point, which was the aim of this multi frequency implementation.

To understand why a multi frequency extension of redundant calibration does not solve the problems immediately we have to return to Equation 3.21. The matrix $[\mathbf{A}^T \mathbf{A}]^{-1} \mathbf{A}^T$ mixes the phase wrapping and the non-phase wrapping channels into the calibration solutions. Although adding extra frequency channels does adds another set of constraints to the calibration solutions, the solutions do no escape the impact of the phase wrapping channel. Adding more channels therefore would decrease variance, but not resolve a phase wrapping point. To really resolve phase wrapping in redundant calibration we require the inclusion of knowledge of the sky.




Figure A.1: From left to right: the amplitude gain, the phase gain solutions for a single antenna in an ideal 5 element interferometer, the visibility amplitude and the visibility phase for one of the two redundant groups as a function of strong source position l . The dark blue line represents the mean of the solutions, the shaded blue area indicates the 1-sigma solutions variance. The amplitude solution variance inversely follows the shape of the beam, i.e. the variance increases when the beam response decreases. The mean of the phase solutions generally fluctuates around $\phi = 0$ and deviates along with the variance at so-called phase wrapping points, which are further explained in the text. The additional position offset, which causes a phase offset from the ideal redundant phase, is absorbed into the solutions causing a slope.

The Covariance Matrix of Beam Modelling Errors

This appendix is a reproduction of the appendix in Ronniy C Joseph, C M Trott, R B Wayth, A Nasirudin, Calibration and 21-cm power spectrum estimation in the presence of antenna beam variations, *Monthly Notices of the Royal Astronomical Society*, Volume 492, Issue 2, February 2020, Pages 2017-2028, <https://doi.org/10.1093/mnras/stz3375>. It differs from the origin form due to minor alternations, such a referencing to earlier chapters to enhance the flow of this thesis.

B.1 Propagating Covariance Matrices to Power Spectrum space

In this paper we derive covariance matrices that describe residuals in PS space. However, we are interested in the structures of residuals in PS space (u, η) . Normally, we apply a frequency taper Γ to our data before Fourier Transforming our data. To compute the covariance in PS space we use the following transformation:

$$\tilde{\mathbf{C}} = \mathbf{F}^\dagger \Gamma(\nu) \mathbf{C} \Gamma(\nu') \mathbf{F}. \quad (\text{B.1})$$

In this work we use a Blackmann-Harris function as a taper. Despite performing extremely well at suppressing sidelobes in PS space. It still has a non-negligible sidelobe within the EoR window, see Figure 6.2. This window function is also used in current MWA EoR pipelines.

B.2 The General Beam Covariance Matrix

Starting from equation (6.18) we can write out the cross-terms in the covariance matrix \mathbf{C}_{ij} .

$$\begin{aligned}
\mathbf{C}_{ij} = & b_p b_q^* e^{-2\pi i(\mathbf{u} \cdot \mathbf{l}_p - \mathbf{u}' \cdot \mathbf{l}_q)} \times \left(b_p^* b_q' \text{Cov}[\delta I_p, \delta I_q'] \right. \\
& + b_p^* \text{Cov}[\delta I_p, \delta b_q^* I_q'] + b_p^* \text{Cov}[\delta I_p, \delta b_q^* \delta I_q'] \\
& + b_q' \text{Cov}[\delta b_p^* I_p, \delta I_q'] + \text{Cov}[\delta b_p^* I_p, \delta b_q^* I_q'] \\
& + \text{Cov}[\delta b_p^* I_p, \delta b_q^* \delta I_q'] + b_q' \text{Cov}[\delta b_p^* \delta I_p, \delta I_q'] \\
& \left. + \text{Cov}[\delta b_p^* \delta I_p, \delta b_q^* I_q'] + \text{Cov}[\delta b_p^* \delta I_p, \delta b_q^* \delta I_q'] \right)
\end{aligned} \tag{B.2}$$

Using the formal definition of the covariance and the properties of the modelled component of the stochastic sky I , the unmodelled component of the sky δI , and our beam perturbations δb , i.e. the the sources that contribute to the intensity about the noise level are independent of the unmodelled sources. So we can further write out equation (B.2).

The first term results into the unmodelled sky covariance matrix \mathbf{C}_{sky} in equation (6.11). The second term becomes zero, because we can separate the averages over the modelled sky, the unmodelled sky, and the beam perturbation, that then cancel each other out.

$$\begin{aligned}
b_p^* \text{Cov}[\delta I_p, \delta b_q^* I_q'] &= b_p^* \left(\langle \delta I_p \delta b_q^* I_q' \rangle - \langle \delta I_p \rangle \langle \delta b_q^* I_q' \rangle \right) \\
&= b_p^* \left(\langle \delta I_p \rangle \langle \delta b_q^* \rangle \langle I_q' \rangle - \langle \delta I_p \rangle \langle \delta b_q^* \rangle \langle I_q' \rangle \right) \\
&= 0
\end{aligned} \tag{B.3}$$

The third term is describes a contribution to the modified sky noise, because the shape of the beam has been changed by a perturbation

$$\begin{aligned} b_p^* \text{Cov}[\delta I_p, \delta b_q^* \delta I'_q] &= b_p^* \left(\langle \delta I_p \delta b_q^* \delta I'_q \rangle - \langle \delta I_p \rangle \langle \delta b_q^* \delta I'_q \rangle \right) \\ &= b_p^* \langle \delta b_q^* \rangle \text{Cov}[\delta I_p, \delta I'_q] \end{aligned} \quad (\text{B.4})$$

The fourth term becomes zero similarly to the second term.

$$b'_q \text{Cov}[\delta b_p^* I_p, \delta I'_q] = 0 \quad (\text{B.5})$$

The fifth term describes the covariance between different parts of the beam, i.e. a beam perturbation in generally changes large portions of the beam, this couples different parts of the sky.

$$\text{Cov}[\delta b_p^* I_p, \delta b_q^* I'_q] = \langle \delta b_p^* \delta b_q^* \rangle \text{Cov}[I_p, I'_q] + \langle I_p \rangle \langle I'_q \rangle \text{Cov}[\delta b_p^*, \delta b_q^*] \quad (\text{B.6})$$

The sixth term describes can also be rewritten as the covariance between different parts of the beam

$$\text{Cov}[\delta b_p^* I_p, \delta b_q^* \delta I'_q] = \langle I_p \rangle \langle \delta I'_q \rangle \text{Cov}[\delta b_p^*, \delta b_q^*] \quad (\text{B.7})$$

The seventh term describes the modification to the sky noise due the perturbation of first frequency

$$b'_q \text{Cov}[\delta b_p^* \delta I_p, \delta I'_q] = \langle \delta b_p^* \rangle b_q^* \text{Cov}[\delta I_p, \delta I'_q] \quad (\text{B.8})$$

The eight' term describes the added covariance due the beam perturbations, similarly to the sixth term.

$$\text{Cov}[\delta b_p^* \delta I_p, \delta b_q^* I'_q] = \langle \delta I_p \rangle \langle I'_q \rangle \text{Cov}[\delta b_p^*, \delta b_q^*] \quad (\text{B.9})$$

And finally the last term, similarly to the fift term, describes the added co-

variance due to different parts of the beam and how it couples different parts of the unmodelled sky together.

$$\begin{aligned} \text{Cov}[\delta b_p^* \delta I_p, \delta b_q^* I_q'] &= \langle \delta b_p^* \delta b_q' \rangle \text{Cov}[\delta I_p, \delta I_q'] \\ &+ \langle \delta I_p \rangle \langle \delta I_q' \rangle \text{Cov}[\delta b_p^*, \delta b_q^*] \end{aligned} \quad (\text{B.10})$$

This leaves us with 6 additional terms on top of the well understood unmodelled sky noise.

$$\begin{aligned} \mathbf{C}_{ij} &= b_p b_q^* e^{-2\pi i(\mathbf{u} \cdot \mathbf{l}_p - \mathbf{u}' \cdot \mathbf{l}_q)} \times \left(\langle \delta b_p^* \delta b_q' \rangle \text{Cov}[I_p, I_q'] \right. \\ &+ (b_p^* b_q' + b_p^* \langle \delta b_q^* \rangle + \langle \delta b_p^* \rangle b_q^* + \langle \delta b_p^* \delta b_q' \rangle) \text{Cov}[\delta I_p, \delta I_q'] \\ &\left. + (\langle I_p \rangle \langle I_q' \rangle + \langle I_p \rangle \langle \delta I_q' \rangle + \langle \delta I_p \rangle \langle I_q' \rangle + \langle \delta I_p \rangle \langle \delta I_q' \rangle) \text{Cov}[\delta b_p^*, \delta b_q^*] \right) \end{aligned} \quad (\text{B.11})$$

B.3 Beam Variations due to Missing Dipoles

First we need to compute the averaged beam perturbation $\langle \delta b(\mathbf{l}, \nu) \rangle$ and the averaged product of perturbations $\langle \delta b(\mathbf{l}, \nu) b(\mathbf{l}', \nu') \rangle$. We can calculate the averaged beam perturbation by simply averaging over the 16 different beam perturbations, in equation (6.26). Although it is possible to write a single expression for these averages using the geometry of a square MWA tile, however, when we Fourier transform this we will end up with a sum over the different broken dipole realisations. Hence, we keep this simple formulation that describes a general tile lay-out.

$$\langle \delta b(\mathbf{l}, \nu) \rangle = -\frac{1}{N^2} b_d \sum_{n=1}^N e^{-2\pi i \mathbf{x}_n \cdot \mathbf{l} / \lambda} \quad (\text{B.12})$$

We can calculate the averaged product of beam perturbations at different frequencies in a similar fashion.

$$\begin{aligned}
\langle \delta b(\mathbf{l}, \nu) \delta b^*(\mathbf{l}', \nu') \rangle &= \frac{1}{N^3} b_d(\mathbf{l}, \nu) b_d^*(\mathbf{l}', \nu') \sum_{n=1}^N e^{2\pi i \mathbf{x}_n \cdot \mathbf{l} / \lambda} e^{2\pi i \mathbf{x}_n \cdot \mathbf{l}' / \lambda'} \\
&= \frac{1}{N^3} b_d(\mathbf{l}, \nu) b_d^*(\mathbf{l}', \nu') \sum_{n=1}^N e^{-2\pi i \mathbf{x}_n \cdot (\mathbf{l} / \lambda - \mathbf{l}' / \lambda')}
\end{aligned} \tag{B.13}$$

Combining equations (B.12), (B.13), and (6.22), we get:

$$\mathbf{C}_{\text{beam}} = \mathbf{C}_A - \mathbf{C}_B - \mathbf{C}_C + \mathbf{C}_D - \mathbf{C}_E \tag{B.14}$$

where all individual covariance components are given by

$$\begin{aligned}
\mathbf{C}_A &= (\mu_{2,m} + \mu_{2,r}) (f_0 f'_0)^{-\gamma} \int \langle \delta b^*(\mathbf{l}, \nu) \delta b(\mathbf{l}, \nu') \rangle b(\mathbf{l}, \nu) b^*(\mathbf{l}, \nu') \\
&\quad \times e^{-2\pi i (\mathbf{u} - \mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
\mathbf{C}_B &= \mu_{2,r} (f_0 f'_0)^{-\gamma} \int b^*(\mathbf{l}, \nu) \langle \delta b^*(\mathbf{l}, \nu') \rangle b(\mathbf{l}, \nu) b^*(\mathbf{l}, \nu') \\
&\quad \times e^{-2\pi i (\mathbf{u} - \mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
\mathbf{C}_C &= \mu_{2,r} (f_0 f'_0)^{-\gamma} \int \langle \delta b^*(\mathbf{l}, \nu) \rangle b^*(\mathbf{l}, \nu') b(\mathbf{l}, \nu) b^*(\mathbf{l}, \nu') \\
&\quad \times e^{-2\pi i (\mathbf{u} - \mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
\mathbf{C}_D &= (\mu_{1,m} + \mu_{1,r})^2 (f_0 f'_0)^{-\gamma} \iint \langle \delta b^*(\mathbf{l}, \nu) \delta b(\mathbf{l}, \nu') \rangle \\
&\quad \times b(\mathbf{l}, \nu) b^*(\mathbf{l}', \nu') e^{-2\pi i (\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}' \\
\mathbf{C}_E &= (\mu_{1,m} + \mu_{1,r})^2 (f_0 f'_0)^{-\gamma} \iint \langle \delta b^*(\mathbf{l}, \nu) \rangle \langle \delta b(\mathbf{l}, \nu') \rangle \\
&\quad \times b(\mathbf{l}, \nu) b^*(\mathbf{l}', \nu') e^{-2\pi i (\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}'.
\end{aligned} \tag{B.15}$$

This leaves us with 6 integrals to solve. We combine the covariance terms $A - D$ in equation (B.15) with (B.12) and (B.13), and write the product of beams, into one single Gaussian. This enables us to use the Hankel transform for each broken dipole realisation.

$$\begin{aligned}
\mathbf{C}_A &= \frac{2\pi\Sigma_A^2(\mu_{2,m} + \mu_{2,r})(f_0f'_0)^{-\gamma}}{N^3} \times \\
&\quad \sum_{n=0}^N \exp\left(-2\pi^2\Sigma_A^2|\mathbf{u} - \mathbf{u}' + \mathbf{x}_n(1/\lambda - 1/\lambda')|^2\right) \\
\mathbf{C}_B &= -\frac{2\pi\Sigma_B^2\mu_{2,r}(f_0f'_0)^{-\gamma}}{N^2} \sum_{n=0}^N \exp\left(-2\pi^2\Sigma_B^2|\mathbf{u} - \mathbf{u}' + \mathbf{x}_n/\lambda|^2\right) \\
\mathbf{C}_C &= -\frac{2\pi\Sigma_C^2\mu_{2,r}(f_0f'_0)^{-\gamma}}{N^2} \sum_{n=0}^N \exp\left(-2\pi^2\Sigma_C^2|\mathbf{u} - \mathbf{u}' + \mathbf{x}_n/\lambda|^2\right)
\end{aligned} \tag{B.16}$$

Here we define:

$$\begin{aligned}
\Sigma_A^2 &= \frac{\sigma^2\sigma'^2\sigma_d^2\sigma_d'^2}{\sigma'^2\sigma_d^2\sigma_d'^2 + \sigma^2\sigma_d^2\sigma_d'^2 + \sigma^2\sigma'^2\sigma_d^2 + \sigma^2\sigma'^2\sigma_d'^2} \\
\Sigma_B^2 &= \frac{\sigma^2\sigma'^2\sigma_d'^2}{2\sigma'^2\sigma_d'^2 + \sigma^2\sigma_d'^2 + \sigma^2\sigma'^2} \\
\Sigma_C^2 &= \frac{\sigma^2\sigma'^2\sigma_d^2}{\sigma'^2\sigma_d^2 + 2\sigma^2\sigma_d^2 + \sigma^2\sigma'^2}
\end{aligned} \tag{B.17}$$

The 4th and fifth integral require us keep the primed and unprimed coordinates separate, we then solve the integrals separately using the same procedure. Rewrite all beams that observe the same sky, either l or l' into a single Gaussian and perform two Hankel transforms one over the unprimed and one over the primed coordinates..

$$\begin{aligned}
\mathbf{C}_D &= (\mu_{1,m} + \mu_{1,r})^2 (f_0 f'_0)^{-\gamma} \frac{2\pi \Sigma_D^2 \Sigma_D'^2}{N^3} \\
&\quad \times \sum_{n=0}^N \exp(-2\pi^2 \Sigma_D^2 |\mathbf{u} - \mathbf{x}_n/\lambda|^2) \exp(-2\pi^2 \Sigma_D'^2 |\mathbf{u}' - \mathbf{x}_n/\lambda'|^2) \\
\mathbf{C}_E &= (\mu_{1,m} + \mu_{1,r})^2 (f_0 f'_0)^{-\gamma} \frac{2\pi^2 \Sigma_D \Sigma_D'}{N^4} \\
&\quad \times \left(\sum_{n=0}^N \exp(-2\pi^2 \Sigma_D^2 |\mathbf{u} - \mathbf{x}_n/\lambda|^2) \right) \\
&\quad \times \left(\sum_{n=0}^N \exp(-2\pi^2 \Sigma_D'^2 |\mathbf{u}' - \mathbf{x}_n/\lambda'|^2) \right)
\end{aligned} \tag{B.18}$$

Where

$$\Sigma_D^2 = \frac{\sigma^2 \sigma_d^2}{\sigma^2 + \sigma_d^2} \tag{B.19}$$

These 6 terms together, effectively averaging the error over 16 dipoles form the covariance for a single baseline between different frequencies. It quantifies the error and how this error is correlated.

B.4 Linearised Gain Error Covariance Matrix

Starting from the definition of covariance we compute $\langle \hat{\mathbf{r}} \hat{\mathbf{r}}' \rangle$ and $\langle \hat{\mathbf{r}} \rangle \langle \hat{\mathbf{r}}' \rangle$ to form the residual covariance matrix. Taking Equation (6.34) we can work out the averaged product assuming the gain error δg is independent from both \mathbf{m} and \mathbf{r} within the same u -bin. This is strictly not true, but because the gain error is a linear combination of several independent u -bins this is a reasonable approximation.

$$\begin{aligned}
\langle \hat{\mathbf{r}} \hat{\mathbf{r}}' \rangle &= \langle (\delta g + \delta g^*) (\delta g' + \delta g^{*'}) \rangle \langle m m^{*'} \rangle \\
&\quad - \langle (\delta g + \delta g^*) (1 - \delta g' - \delta g^{*'}) \rangle \langle m \rangle \langle r^{*'} \rangle \\
&\quad - \langle (1 - \delta g - \delta g^*) (\delta g' + \delta g^{*'}) \rangle \langle m^{*'} \rangle \langle m \rangle \\
&\quad + \langle (1 - \delta g - \delta g^*) (1 - \delta g' - \delta g^{*'}) \rangle \langle r r^{*'} \rangle
\end{aligned} \tag{B.20}$$

The product of averages is fairly straight forward to compute, when we combine both to get the covariance we notice that many terms will drop out leaving us with

$$\begin{aligned}
\mathbf{C}_{\hat{r}} &= \langle (\delta g + \delta g^*)(\delta g' + \delta g'^*) \rangle \langle m m'^* \rangle \\
&\quad - \langle (\delta g + \delta g^*) \rangle \langle (\delta g' + \delta g'^*) \rangle \langle m \rangle \langle m'^* \rangle \\
&\quad + \langle (1 - \delta g - \delta g^*)(1 - \delta g' - \delta g'^*) \rangle \langle r r'^* \rangle \\
&\quad - \langle (1 - \delta g - \delta g^*) \rangle \langle (1 - \delta g' - \delta g'^*) \rangle \langle r \rangle \langle r'^* \rangle
\end{aligned} \tag{B.21}$$

This expression looks fairly similar to the product of gain covariance \mathbf{C}_g with either the model covariance \mathbf{C}_m or the residual covariance \mathbf{C}_r . Because the mean model visibilities $\langle m \rangle$ integrate to zero at u scales beyond the diameter of a tile, and similarly for the residuals $\langle r \rangle$ we can all terms that contain these means. Following a similar argument we can say that $\mathbf{C}_g \sim \langle \delta g \delta g' \rangle$, $\mathbf{C}_m \sim \langle \delta m \delta m' \rangle$, and $\mathbf{C}_r \sim \langle \delta r \delta r' \rangle$. Assuming the gain errors from two antennas are independent we can write $2\mathbf{C}_g = \langle (\delta g + \delta g')(\delta g + \delta g')^* \rangle$

Appendix C

The Covariance Matrix of Non-Redundant Errors

In this appendix we discuss the details over the derivation of the covariance matrices that describe the errors during redundant calibration.

C.1 The Covariance Matrix of Non-Redundant Errors

To derive the covariance matrix of non-redundancies, e.g. position errors and beam perturbations, we start by perturbing the general visibility equation with a beam perturbation δb and a position perturbation $\delta \mathbf{u}$.

$$V(\mathbf{u}, \nu) = \int I b(b + \delta b)^* \times e^{-2\pi i(\mathbf{u} + \delta \mathbf{u}) \cdot \mathbf{l}} d^2 \mathbf{l}. \quad (\text{C.1})$$

We then linearise the exponential around \mathbf{u}

$$V(\mathbf{u}, \nu) = \int I b(b + \delta b)^* (1 - 2\pi i \delta \mathbf{u} \cdot \mathbf{l}) \times e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}. \quad (\text{C.2})$$

If we expand this and ignore higher order terms that contain both the position and beam perturbation, we can extract the perturbation to the redundant

visibility

$$\delta V(\mathbf{u}, \nu) = \int I(b\delta b^* - 2\pi i b b^* \delta \mathbf{u} \cdot \mathbf{l}) \times e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} d^2 \mathbf{l}. \quad (\text{C.3})$$

We want to estimate the covariance \mathbf{C}_r of these "residuals" to our redundant calibration model

$$\begin{aligned} \mathbf{C}_{r,r} &= \sum_{pq} b_p b_q^{*'} e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_p - \mathbf{u}' \cdot \mathbf{l}_q} \Delta \mathbf{l}_p \Delta \mathbf{l}_q \\ &\times \text{Cov}[I_p(\delta b_p^* - 2\pi i b_p^* \delta \mathbf{u} \cdot \mathbf{l}_p), I'_q(\delta b_q^{*'} - 2\pi i b_q^{*'} \delta \mathbf{u}' \cdot \mathbf{l}_q)]. \end{aligned} \quad (\text{C.4})$$

Here, we discretised the sky in pixels with size $\Delta \mathbf{l}$, and taken all non-random quantities out of the covariance. Writing out the cross-terms in the covariance term yields 4 terms, see below

$$\begin{aligned} \mathbf{C}_{r,r} &= \sum_{pq} b_p b_q^{*'} e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_p - \mathbf{u}' \cdot \mathbf{l}_q} \Delta \mathbf{l}_p \Delta \mathbf{l}_q \\ &\times \left(\text{Cov}[I_p \delta b_p^*, I'_q \delta b_q^{*'}] - \text{Cov}[2\pi i b_p^* \delta \mathbf{u} \cdot \mathbf{l}_p, I'_q \delta b_q^{*'}] \right. \\ &\quad \left. - \text{Cov}[I_p \delta b_p^*, 2\pi i b_q^{*'} \delta \mathbf{u}' \cdot \mathbf{l}_q] + \text{Cov}[2\pi i b_p^* \delta \mathbf{u} \cdot \mathbf{l}_p, 2\pi i b_q^{*'} \delta \mathbf{u}' \cdot \mathbf{l}_q] \right). \end{aligned} \quad (\text{C.5})$$

We now discuss each of the four covariance terms in Equation (C.5). We derived the first term in Chapter 6 and Appendix B.2.

$$\text{Cov}[I_p \delta b_p^*, I'_q \delta b_q^{*'}] = \langle \delta b_p \delta b_q' \rangle \text{Cov}[I_p, I'_q] + \langle I_p \rangle \langle I_q' \rangle \text{Cov}[\delta b_p^*, \delta b_q^{*'}]. \quad (\text{C.6})$$

For the second the term we write out the covariance explicitly, $\text{Cov}[X, Y] = \langle XY \rangle - \langle X \rangle \langle Y \rangle$, and separate averages over independent terms and rewrite:

$$\text{Cov}[2\pi i b_p^* \delta \mathbf{u} \cdot \mathbf{l}_p, I'_q \delta b_q^{*'}] = 2\pi b_p \langle \delta \mathbf{u} \cdot \mathbf{l}_p \rangle \langle \delta b_q' \rangle \text{Cov}[I_p, I'_q] \quad (\text{C.7})$$

We assume the baselines offset distribution is Gaussian with a mean of zero and standard deviation $\delta\mathbf{u}$. Hence this second term becomes zero, and similarly for the 3rd term. The final term is slightly more elaborate. We again write out the covariance explicitly take out the constant terms

$$\begin{aligned} \text{Cov}[2\pi i b_p^* \delta\mathbf{u} \cdot \mathbf{l}_p, 2\pi i b_q^* \delta\mathbf{u} \cdot \mathbf{l}_q] &= 4\pi^2 b_p b_q' \text{Cov}[I_p \delta\mathbf{u} \cdot \mathbf{l}_p, I_q' \delta\mathbf{u} \cdot \mathbf{l}_q] \\ &= 4\pi^2 b_p b_q' \left(\langle I_p I_q' \rangle \langle (\delta\mathbf{u} \cdot \mathbf{l}_p) (\delta\mathbf{u}' \cdot \mathbf{l}_q) \rangle \right. \\ &\quad \left. - \langle I_p \rangle \langle I_q' \rangle \langle \delta\mathbf{u} \cdot \mathbf{l}_p \rangle \langle \delta\mathbf{u} \cdot \mathbf{l}_q \rangle \right) \end{aligned} \quad (\text{C.8})$$

Noting again that the second term in the parentheses contains the mean $\langle \delta\mathbf{u} \cdot \mathbf{l}_p \rangle$ which we assume to be zero, we arrive at

$$= 4\pi^2 b_p b_q' \langle (\delta\mathbf{u} \cdot \mathbf{l}_p) (\delta\mathbf{u}' \cdot \mathbf{l}_q) \rangle \times \left(\text{Cov}[I_p, I_q'] + \langle I_p \rangle \langle I_q' \rangle \right). \quad (\text{C.9})$$

Note: in general Equation (C.9) reduces to zero for completely independent baselines. However, when one antenna is shared between two different baselines there also exist some covariance due to the shared position offset of that particular antenna. Our notation that describes the covariance between different $\{u, v, f\}$ -cells does not allow for this. However we will revisit this issue after having simplified these equations. Combining Equations (C.5), (C.6), (C.7) and (C.9) we can start integrating this over the entire sky.

$$\mathbf{C}_r^r = \mathbf{C}_{\text{beam}}^r + \mathbf{C}_{\text{position}}^r \quad (\text{C.10})$$

Where we have defined the first terms from Equation (C.6) as the $\mathbf{C}_{\text{beam},r}$ and the terms from Equation (C.9) as $\mathbf{C}_{\text{position},r}$. Writing out the integrals we get

$$\begin{aligned}
\mathbf{C}_{\text{beam}} &= \mu_2(f_0 f'_0)^{-\gamma} \int b^* b' \langle \delta b^* \delta b' \rangle e^{-2\pi i(\mathbf{u}-\mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
&+ \mu_1^2(f_0 f'_0)^{-\gamma} \iint b b^* \text{Cov}[\delta b^*, \delta b^*'] \times e^{-2\pi i(\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}'
\end{aligned} \tag{C.11}$$

and

$$\begin{aligned}
\mathbf{C}_{\text{position}} &= 4\pi^2 \mu_2(f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \int B B' |\mathbf{l}|^2 e^{-2\pi i(\mathbf{u}-\mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
&+ 4\pi^2 \mu_1^2(f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \iint B B' \mathbf{l} \cdot \mathbf{l}' e^{-2\pi i(\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}'
\end{aligned} \tag{C.12}$$

Here, we have used the fact that for a stochastic sky of point sources, the integral over $\text{Cov}[I_p, I_q]$ is only non-zero when we consider the same part of the sky, i.e. $p = q$. In that case we only consider the variance $\text{Var}[I_p] = \mu_2(f_0 f'_0)^{-\gamma}$, where μ_2 is the second moment of the source count distribution. In Equations C5 and C7 in J19 we solved Equation (7.13) for the 'broken' dipole case in the MWA.

In the appendix we only focus on the position offset covariance matrix. The first integral in Equation (C.12) can be written as a defined Hankel-transform by first integrating over a circle around $2\pi|\mathbf{u} - \mathbf{u}'| \cos \theta$. If we assume a Gaussian beam, the circular symmetry simply adds another factor of 2π

$$\begin{aligned}
&= 4\pi^2 \mu_2(f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \int B B' |\mathbf{l}|^2 e^{-2\pi i(\mathbf{u}-\mathbf{u}') \cdot \mathbf{l}} d^2 \mathbf{l} \\
&= (2\pi)^3 \mu_2(f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \int B B' l^2 e^{-2\pi i|\mathbf{u}-\mathbf{u}'|l} l dl
\end{aligned} \tag{C.13}$$

We can then use the following Hankel transform relation

$$\int_0^\infty r^2 f(r) J_0(kr) r dr = -\frac{d^2 F_0}{dk^2} - \frac{1}{k} \frac{dF_0}{dk} \tag{C.14}$$

where F_0 is the zeroth order Hankel transform of $f(r)$. We first rewrite the product of the two Gaussian primary beam responses as one Gaussian $f(r) = \exp(-l^2/2\Sigma^2)$, where

$$\Sigma^2 = \frac{\sigma^2 \sigma'^2}{\sigma^2 + \sigma'^2}. \quad (\text{C.15})$$

For such a Gaussian the zeroth order Hankel transform becomes

$$F_0 = 2\Sigma_\nu^2 \exp(-2\pi^2 |\mathbf{u} - \mathbf{u}'|^2 \Sigma_\nu^2). \quad (\text{C.16})$$

With these relations we find that the first integral in Equation (C.12) becomes

$$\int BB'l^2 e^{-2\pi i |\mathbf{u} - \mathbf{u}'| l} l dl = 2\Sigma_\nu^2 e^{-2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|^2} (1 - 2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|) \quad (\text{C.17})$$

For the second integral in Equation (C.12) we first solve the Fourier transform over \mathbf{l} and then the integral over \mathbf{l}' . To do so we use the relation $\mathcal{F}\{rf(r)\} = i/(2\pi) d\tilde{B}/dk$ twice, one time per coordinate in $\mathbf{l} = (l, m)$. This results into the following

$$\iint BB' \mathbf{l} \cdot \mathbf{l}' e^{-2\pi i (\mathbf{u} \cdot \mathbf{l} - \mathbf{u}' \cdot \mathbf{l}')} d^2 \mathbf{l} d^2 \mathbf{l}' = -\frac{1}{(2\pi)^2} (\nabla \tilde{B} \cdot \nabla \tilde{B}'^*) \quad (\text{C.18})$$

Here, \tilde{B} denotes the Fourier transform of the antenna response. Conveniently, for a Gaussian beam

$$B(\mathbf{l}) = e^{-|\mathbf{l}|^2 / 2\sigma^2} \quad (\text{C.19})$$

its Fourier transform is rather trivial

$$\tilde{B}(\mathbf{u}) = 2\pi\sigma^2 e^{-2\pi^2\sigma^2 |\mathbf{u}|^2}. \quad (\text{C.20})$$

Taking the gradient of this yields

$$\nabla \tilde{B}(\mathbf{u}) = -(2\pi)^3 \sigma^4 \mathbf{u} e^{-2\pi^2\sigma^2 |\mathbf{u}|^2}. \quad (\text{C.21})$$

The same holds for its complex conjugate except without the negative signs in

the exponent. Combining Equations (C.12), (C.17), (C.18), and (C.21) we get the following expression for the position covariance

$$\begin{aligned} \mathbf{C}_{\text{position}} &= 16\pi^3 \mu_2 (f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle \Sigma_\nu^2 e^{-2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|} (1 - 2\pi^2 \Sigma_\nu^2 |\mathbf{u} - \mathbf{u}'|) \\ &+ (2\pi)^8 \mu_1^2 (f_0 f'_0)^{-\gamma} \langle \delta u \delta u' \rangle (\sigma \sigma')^4 (\mathbf{u} \cdot \mathbf{u}') e^{-2\pi^2 (\sigma |\mathbf{u}|^2 + \sigma' |\mathbf{u}'|^2)} \end{aligned} \quad (\text{C.22})$$

The last term in Equation (C.22) decays to zero for baselines larger than the size of an MWA tile. We are now only left with addressing one issue, the correlation between different baselines. When we consider the same baseline and abuse notation to index antennas with p and q , and write out the baseline perturbation in terms of antenna position perturbations $\delta u_{pq} = (\delta x_p + \delta x_q)/\lambda$. For the same baseline $\langle \delta u_{pq} \delta u'_{pq} \rangle = 2\langle \delta x \rangle^2 / (\lambda \lambda')$, i.e. the variance of the position offset distribution. For baselines that share one antenna this becomes $\langle \delta u_{pq} \delta u'_{ps} \rangle = \langle \delta x \rangle^2 / (\lambda \lambda')$, because only the variance of the shared antenna contributes here. Naturally, for completely independent baselines the term becomes zero.

$$\begin{aligned} \langle \delta u_{pq} \delta u'_{kl} \rangle &= (\delta_{pk} + \delta_{ql} + \delta_{pl} + \delta_{qk}) \langle \delta x \rangle^2 / (\lambda \lambda') \\ &= (\delta_{pk} + \delta_{ql} + \delta_{pl} + \delta_{qk}) \langle \delta x \rangle^2 \frac{\sigma_x^2 \nu_0^2}{c^2} (f_0 f'_0) \end{aligned} \quad (\text{C.23})$$

Agreement of Co-authors

This section is to comply with the requirement that all co-authors of work included in this thesis outline their contributions. All co-authors have read and agreed to the following statement, and their responses can be found below.

D.1 Statement of Originality

Chapter 5 is a reproduction of “The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions”. Its contents are my own work, except for the input and supervision from Dr. Cathryn Trott and Dr. Randall Wayth during the development of the redundant calibration simulations and theoretical analysis of sky based calibration errors. I wrote the the draft of the paper and distributed copies to all co-authors for proofreading. All co-authors provided input during several iterations of feedback until the manuscript was finalised.

Chapter 6 is a reproduction of “Calibration and 21-cm power spectrum estimation in the presence of antenna beam variations”. Its contents are also my work, except for the input and close supervision from Dr. Cathryn Trott and Dr. Randall Wayth during the development of the analytic description of beam modelling errors in the Murchison Widefield Array. Ainulnabilah Nasirudin provided suggestions that aided in the rapid development of visualisation tools for the re-

sults and their interpretation. I wrote the the draft of the paper and distributed copies to all co-authors for proofreading. All co-authors provided input during several iterations of feedback until the manuscript was finalised.

D.2 Co-author Responses

Cathryn Trott <cathryn.trott@curtin.edu.au>

Hi Ronniy,

I agree that the below statements are an accurate representation of my input into the papers.

Cheers,

Cath

Randall Wayth <r.wayth@curtin.edu.au>

Agreed thanks.

Ainulnabilah Nasirudin <a.nasirudin@postgrad.curtin.edu.au>

Sounds good!

I will also send you this sometime in the future.

Cheers,

Bella

Appendix **E**

Copyright Information

AMERICAN ASTRONOMICAL SOCIETY

This agreement must be electronically signed before the American Astronomical Society (AAS) can publish your paper. In the event the article is not judged acceptable for publication in the journal you will be notified in writing and the copyright and all rights conferred by this agreement shall revert to you.

PUBLICATION AND TRANSFER OF COPYRIGHT AGREEMENT Manuscript number: AAS13096

Article title: The Bias and Uncertainty of Redundant and Sky-Based Calibration under Realistic Sky and Telescope Conditions

Names of authors: Ronniy Joseph Cathryn Trott Randall Wayth

Author Rights: AAS grants to the author(s) (or their below-named employers, in the case of works made for hire) the following rights. All copies of the Article made under any of these rights shall include notice of the AAS copyright.

- (1) All proprietary and statutory rights other than copyright, such as patent rights.
- (2) The right after publication by the AAS to grant or refuse permission to third parties to republish all or part of the Article or a translation thereof. In the case of whole articles only, third parties must first obtain permission from the AAS before any right of further publication is granted. The AAS may choose to publish an abstract or portions of the Article before the AAS publishes it in a journal.
- (3) The right to use all or part of the Article in future works and derivative works of their own of any type, and to make copies of all or part of the Article for the authors' use for educational or research purposes.
- (4) In the case of a work made for hire, the right of the employer to make copies of the Article for the employer's internal use, but not for resale.
- (5) In the case of a Gold Open Access work, each of the named authors and, where relevant, the employer consents to the publication of the work under the Creative Commons Attribution 3.0 license (<https://creativecommons.org/licenses/by/3.0/>) or any successor to that license.

Copyright Assignment: Copyright in the Article is hereby transferred to the AAS for the full term of copyright throughout the world, effective as of date of acceptance of the Article for publication in a journal of the AAS. The copyright consists of all rights protected by copyright laws of the United States and of all foreign countries, in all languages and forms of communication, and includes all material to be published as part of the Article in any format or medium. The AAS shall have the right to register copyright to the Article in its name as claimant, whether separately or as part of the journal issue or other medium in which the Article is included and the right to sue, counterclaim, and recover for past, present and future infringement of the rights assigned under this agreement.

This Agreement shall be controlled, construed and enforced in accordance with the laws of the District of Columbia without reference to the conflict of laws provisions thereof. The Parties consent to jurisdiction of the state and federal courts located in the District of Columbia in connection with any proceeding related to this Agreement or its enforcement.

Authorized Signature: Ronniy Joseph Date: 28-08-2018

Certification of Government Employment: An article prepared by a government officer or employee as part of his or her official duties may not be eligible for copyright, if the authors are all employed by one of the governments of Australia, Canada, New Zealand, the UK, or the US. If all the authors of the article are such government employees, one of the authors should sign here. If any of the authors is not such a government employee, do not sign in this box.

Authorized Signature:

Date:

local_p_id: 680745

time: 1535417429

ip address: 134.7.50.214

Journal: Monthly Notices of the Royal Astronomical Society

DOI: 10.1093/mnras/stz3375

Title: Calibration and 21 cm Power Spectrum Estimation in the Presence of Antenna Beam Variations

Standard Licence

You hereby grant to Oxford University Press an exclusive licence for the full period of copyright throughout the world:

- to publish the final version of the Article in the above Journal, and to distribute it and/or to communicate it to the public, either within the Journal, on its own, or with other related material throughout the world, in printed, electronic or any other format or medium whether now known or hereafter devised;
- to make translations and abstracts of the Article and to distribute them to the public;
- to authorize or grant licences to third parties to do any of the above;
- to deposit copies of the Article in online archives maintained by OUP or by third parties authorized by OUP.

You authorize us to act on your behalf to defend the copyright in the Article if anyone should infringe it and to register the copyright of the Article in the US and other countries, if necessary.

In the case of a multi authored article, you confirm that you are authorized by your co-authors to enter the licence on their behalf.

You confirm to OUP that the Article

- is your original work;
- has not previously been published (in print or electronic format), is not currently under consideration by another journal, or if it has already been submitted to other journal, it will be immediately withdrawn;
- will not be submitted for publication to any other journal following acceptance in the above Journal; and
- OUP will be the first publisher of the Article.

You warrant to OUP that

- no part of the Article is copied from any other work,
- you have obtained ALL the permissions required (for print and electronic use) for any material you have used from other copyrighted publications in the Article; and
- you have exercised reasonable care to ensure that the Article is accurate and does not contain anything which is libellous, or obscene, or infringes on anyone's copyright, right of privacy, or other rights.

Further Information

(Full details of OUP's publication rights policies, including author rights can be found at https://academic.oup.com/journals/pages/access_purchase/rights_and_permissions/publication_rights)

Author Self-Archiving Policy

On publication of your Article in the Journal you are not required to remove any previously posted ORIGINAL VERSIONS from your own personal website or that of your employer or free public servers of articles in your subject area, provided (1) you include a link (url) to the VERSION OF RECORD on the Journal's website; AND (2) the Journal is attributed as the original place of publication with the correct citation details given.

You may post the ACCEPTED MANUSCRIPT or the VERSION OF RECORD onto your own website, your institution's website and in institutional or subject-based repositories upon publication, provided that: (1) you include a link (url) to the VERSION OF RECORD on the Journal's website; (2) the Journal is attributed as the original place of publication with the correct citation details given.

Free Link to Published Article

On publication of your article, you will receive a URL, giving you access to the published article on the Journal website, and information on use of this link.

Educational Use

You may use the Article within your employer's institution or company for educational or research purposes only, including use in course-packs, as long as: (1) you do not use it for commercial purposes or re-distribution outside of the institution/company; (2) you

acknowledge the Journal as the original place of publication with the correct citation details given.

ronniy.joseph@icrar.org signed this licence on 2019-12-05 20:30:41 GMT.

Bibliography

- J. G. Albert, et al. (2020a). ‘A probabilistic approach to direction-dependent ionospheric calibration’. *Astronomy and Astrophysics* **633**:A77.
- J. G. Albert, et al. (2020b). ‘Probabilistic direction-dependent ionospheric calibration for LOFAR-HBA’. *Astronomy and Astrophysics* **635**:A147.
- Z. S. Ali, et al. (2018). ‘Erratum: “PAPER-64 Constraints on Reionization: The 21 cm Power Spectrum at $z = 8.4$ ” (2015, ApJ, 809, 61)’. *The Astrophysical Journal* **863**(2):201.
- Z. S. Ali, et al. (2015). ‘PAPER-64 CONSTRAINTS ON REIONIZATION: THE 21 cm POWER SPECTRUM AT $z = 8.4$ ’. *The Astrophysical Journal* **809**(1):61.
- J. W. M. Baars, et al. (2009). ‘Radio Astronomy in the Early Twenty-First Century’. *Proceedings of the IEEE* **97**(8):1377–1381.
- C. Balanis (2016). *Antenna Theory: Analysis and Design*. Wiley.
- R. Barkana (2018). ‘Possible interaction between baryons and dark-matter particles revealed by the first stars’. *Nature* **555**:71–74.
- R. Barkana & A. Loeb (2001). ‘In the beginning: The First sources of light and the reionization of the Universe’. *Phys.Rept.* **349**:125–238.

- H. Barrett & K. Myers (2004). *Foundations of image science*. Wiley series in pure and applied optics. Wiley-Interscience.
- N. Barry, et al. (2019a). ‘The FHD/ ϵ ppsiion Epoch of Reionisation power spectrum pipeline’. *Publications of the Astronomical Society of Australia* **36**:e026.
- N. Barry, et al. (2016). ‘Calibration requirements for detecting the 21 cm epoch of reionization power spectrum and implications for the SKA’. *Monthly Notices of the Royal Astronomical Society* **461**:3135–3144.
- N. Barry, et al. (2019b). ‘Improving the Epoch of Reionization Power Spectrum Results from Murchison Widefield Array Season 1 Observations’. *The Astrophysical Journal* **884**(1):1.
- A. Beardsley, et al. (2019). ‘Science with the Murchison Widefield Array: Phase I results and Phase II opportunities’. *Publications of the Astronomical Society of Australia* **36**.
- A. P. Beardsley, et al. (2016). ‘First Season MWA EoR Power spectrum Results at Redshift 7’. *The Astrophysical Journal* **833**(1):102.
- C. L. Bennett, et al. (2013). ‘Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results’. *The Astrophysical Journal Supplemental* **208**:20.
- S. Bharadwaj & S. K. Pandey (2005). ‘Probing non-Gaussian features in the H i distribution at the epoch of re-ionization’. *Monthly Notices of the Royal Astronomical Society* **358**(3):968–976.
- S. Bharadwaj & S. K. Sethi (2001). ‘HI fluctuations at large redshifts: I-visibility correlation’. *Journal of Astrophysics and Astronomy* **22**(4):293–307.
- R. Bouwens (2016). *High-Redshift Galaxy Surveys and the Reionization of the Universe*, vol. 423 of *Astrophysics and Space Science Library*, p. 111. Springer International Publishing.

- R. J. Bouwens, et al. (2015). ‘UV Luminosity Functions at Redshifts $z \sim 4$ to $z \sim 10$: 10,000 Galaxies from HST Legacy Fields’. *The Astrophysical Journal* **803**:34.
- J. D. Bowman, et al. (2006). ‘The Sensitivity of first generation epoch of reionization observatories and their potential for differentiating theoretical power spectra’. *ApJ* **638**:20–26.
- J. D. Bowman, et al. (2009). ‘Foreground contamination in interferometric measurements of the redshifted 21 cm power spectrum’. *The Astrophysical Journal* **695**(1):183.
- J. D. Bowman, et al. (2008). ‘Toward Empirical Constraints on the Global Redshifted 21 cm Brightness Temperature During the Epoch of Reionization’. *The Astrophysical Journal* **676**:1–9.
- J. D. Bowman, et al. (2018). ‘An absorption profile centred at 78 megahertz in the sky-averaged spectrum’. *Nature* **555**:67–70.
- R. Bracewell (2000). *The Fourier Transform and Its Applications*. Electrical engineering series. McGraw Hill.
- R. Braun & W. van Cappellen (2006). ‘Aperture Arrays for the SKA: Dense or Sparse?’. *arXiv e-prints* pp. astro-ph/0611160.
- J. O. Burns, et al. (2012). ‘Probing the first stars and black holes in the early Universe with the Dark Ages Radio Explorer (DARE)’. *Advances in Space Research* **49**(3):433 – 450.
- R. Byrne, et al. (2019). ‘Fundamental Limitations on the Calibration of Redundant 21 cm Cosmology Instruments and Implications for HERA and the SKA’. *The Astrophysical Journal* **875**:70.

- P. A. Carroll, et al. (2016). ‘A high reliability survey of discrete Epoch of Reionization foreground sources in the MWA EoR0 field’. *Monthly Notices of the Royal Astronomical Society* **461**(4):4151–4175.
- E. Chapman, et al. (2013). ‘The scale of the problem: recovering images of reionization with Generalized Morphological Component Analysis’. *Monthly Notices of the Royal Astronomical Society* **429**(1):165–176.
- E. Chapman, et al. (2012). ‘Foreground removal using FASTICA: a showcase of LOFAR-EoR’. *Monthly Notices of the Royal Astronomical Society* **423**(3):2518–2532.
- C. Cheng, et al. (2018). ‘Characterizing Signal Loss in the 21 cm Reionization Power Spectrum: A Revised Study of PAPER-64’. *The Astrophysical Journal* **868**(1):26.
- S. Choudhuri, et al. (2014). ‘Visibility-based angular power spectrum estimation in low-frequency radio interferometric observations’. *Monthly Notices of the Royal Astronomical Society* **445**:4351–4365.
- B. Ciardi & P. Madau (2003). ‘Probing beyond the Epoch of Hydrogen Reionization with 21 Centimeter Radiation’. *The Astrophysical Journal* **596**:1–8.
- B. G. Clark (1999). ‘Coherence in Radio Astronomy’. In *Synthesis Imaging in Radio Astronomy II*, vol. 180 of *Astronomical Society of the Pacific Conference Series*, p. 1. Astronomical Society of the Pacific.
- A. Cohen, et al. (2017). ‘Charting the parameter space of the global 21-cm signal’. *Monthly Notices of the Royal Astronomical Society* **472**:1915–1931.
- A. S. Cohen & H. J. A. Röttgering (2009). ‘Probing Fine-Scale Ionospheric Structure with the Very Large Array Radio Telescope’. *The Astronomical Journal* **138**:439–447.

- A. Cooray, et al. (2008). ‘Trispectrum of 21-cm background anisotropies as a probe of primordial non-Gaussianity’. *Phys. Rev. D* **77**:103506.
- T. J. Cornwell & P. N. Wilkinson (1981). ‘A new method for making maps with unstable radio interferometers’. *Monthly Notices of the Royal Astronomical Society* **196**:1067–1086.
- A. Datta, et al. (2009). ‘Detection of Signals from Cosmic Reionization Using Radio Interferometric Signal Processing’. *The Astrophysical Journal* **703**:1851–1862.
- A. Datta, et al. (2010). ‘Bright Source Subtraction Requirements for Redshifted 21 cm Measurements’. *The Astrophysical Journal* **724**:526–538.
- P. Dayal & A. Ferrara (2018). ‘Early galaxy formation and its large-scale effects’. *Physics Reports* **780-782**:1 – 64. Early galaxy formation and its large-scale effects.
- D. R. DeBoer, et al. (2017). ‘Hydrogen Epoch of Reionization Array (HERA)’. *Publications of the Astronomical Society of the Pacific* **129**(974):045001.
- P. Dewdney, et al. (2015). ‘SKA1 System BaselineV2 Description’. Tech. rep., SKA Organisation.
- T. Di Matteo, et al. (2002). ‘Radio Foregrounds for the 21 Centimeter Tomography of the Neutral Intergalactic Medium at High Redshifts’. *The Astrophysical Journal* **564**:576–580.
- M. Dijkstra (2016). *Constraining Reionization with Ly α Emitting Galaxies*, vol. 423 of *Astrophysics and Space Science Library*, p. 145. Springer International Publishing.
- J. S. Dillon, et al. (2013). ‘A fast method for power spectrum and foreground analysis for 21 cm cosmology’. *Physical Review D* **87**(4):043005.

- J. S. Dillon, et al. (2015). ‘Empirical covariance modeling for 21 cm power spectrum estimation: A method demonstration and new limits from early Murchison Widefield Array 128-tile data’. *Physical Review D* **91**(12):123011.
- J. S. Dillon & A. R. Parsons (2016). ‘Redundant Array Configurations for 21 cm Cosmology’. *The Astrophysical Journal* **826**:181.
- M. W. Eastwood, et al. (2018). ‘The Radio Sky at Meter Wavelengths: m-mode Analysis Imaging with the OVRO-LWA’. *The Astronomical Journal* **156**:32.
- M. W. Eastwood, et al. (2019). ‘The 21 cm Power Spectrum from the Cosmic Dawn: First Results from the OVRO-LWA’. *The Astronomical Journal* **158**(2):84.
- A. Einstein (1915). ‘Die Feldgleichungen der Gravitation’. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)* **1915**:844–847.
- S. W. Ellingson, et al. (2009). ‘The Long Wavelength Array’. *IEEE Proceedings* **97**:1421–1430.
- A. Ewall-Wice, et al. (2017). ‘The impact of modelling errors on interferometer calibration for 21 cm power spectra’. *Monthly Notices of the Royal Astronomical Society* **470**(2):1849–1870.
- H. I. Ewen & E. M. Purcell (1951). ‘Observation of a Line in the Galactic Radio Spectrum: Radiation from Galactic Hydrogen at 1,420 Mc./sec.’. *Nature* **168**:356.
- X. Fan, et al. (2006). ‘Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with $z \sim 6$ Quasars. II. A Sample of 19 Quasars’. *The Astronomical Journal* **132**:117–136.
- J. A. Fessler & A. O. Hero (1994). ‘Space-alternating generalized expectation-maximization algorithm’. *IEEE Transactions on Signal Processing* **42**(10):2664–2677.

- G. B. Field (1958). ‘Excitation of the Hydrogen 21-CM Line’. *Proceedings of the IRE* **46**(1):240–250.
- G. B. Field (1959). ‘An Attempt to Observe Neutral Hydrogen Between the Galaxies.’. *The Astrophysical Journal* **129**:525.
- T. M. O. Franzen, et al. (2016). ‘The 154 MHz radio sky observed by the Murchison Widefield Array: noise, confusion, and first source count analyses’. *Monthly Notices of the Royal Astronomical Society* **459**:3314–3325.
- A. Friedmann (1922). ‘Über die Krümmung des Raumes’. *Zeitschrift für Physik* **10**:377–386.
- S. Furlanetto, et al. (2006). ‘Cosmology at Low Frequencies: The 21 cm Transition and the High-Redshift Universe’. *Phys.Rept.* **433**:181–301.
- S. R. Furlanetto (2016). ‘The 21-cm Line as a Probe of Reionization’. In A. Mesinger (ed.), *Understanding the Epoch of Cosmic Reionization: Challenges and Progress*, vol. 423 of *Astrophysics and Space Science Library*, p. 247.
- S. R. Furlanetto (2019). ‘Theoretical Framework: The Fundamentals of the 21 cm Line’. In *The Cosmic 21-cm Revolution*, 2514-3433, pp. 1–1 to 1–17. IOP Publishing.
- S. R. Furlanetto, et al. (2004). ‘Statistical probes of reionization with 21 centimeter tomography’. *The Astrophysical Journal* **613**(1):16.
- G. Gamow (1946). ‘Expanding Universe and the Origin of Elements’. *Physical Review* **70**:572–573.
- M. Gervasi, et al. (2008). ‘The Contribution of the Unresolved Extragalactic Radio Sources to the Brightness Temperature of the Sky’. *The Astrophysical Journal* **682**:223–230.

- L. J. Greenhill & G. Bernardi (2012). ‘HI Epoch of Reionization Arrays’ .
- T. L. Grobler, et al. (2014). ‘Calibration artefacts in radio interferometry - I. Ghost sources in Westerbork Synthesis Radio Telescope data’. *Monthly Notices of the Royal Astronomical Society* **439**:4030–4047.
- J. E. Gunn & B. A. Peterson (1965). ‘On the Density of Neutral Hydrogen in Intergalactic Space.’. *The Astrophysical Journal* **142**:1633–1641.
- F. J. Harris (1978). ‘On the use of windows for harmonic analysis with the discrete Fourier transform’. *Proceedings of the IEEE* **66**(1):51–83.
- C. G. T. Haslam, et al. (1982). ‘A 408 MHz all-sky continuum survey. II. The atlas of contour maps.’. *Astronomy and Astrophysics Supplement Series* **47**:1–143.
- D. W. Hogg (2000). ‘Distance measures in cosmology’ .
- E. P. Hubble (1926). ‘Extragalactic nebulae.’. *The Astrophysical Journal* **64**.
- N. Hurley-Walker, et al. (2017). ‘GaLactic and Extragalactic All-sky Murchison Widefield Array (GLEAM) survey - I. A low-frequency extragalactic catalogue’. *Monthly Notices of the Royal Astronomical Society* **464**(1):1146–1167.
- H. T. Intema, et al. (2009). ‘Ionospheric calibration of low frequency radio interferometric observations using the peeling scheme. I. Method description and first results’. *Astronomy and Astrophysics* **501**:1185–1205.
- H. T. Intema, et al. (2011). ‘Deep low-frequency radio observations of the NOAO Boötes field. I. Data reduction and catalog construction’. *Astronomy and Astrophysics* **535**.
- D. C. Jacobs, et al. (2016). ‘The Murchison Widefield Array 21 cm Power Spectrum Analysis Methodology’. *The Astrophysical Journal* **825**:114.
- K. G. Jansky (1932). ‘Directional Studies of Atmospherics at High Frequencies’. *Proceedings of the Institute of Radio Engineers* **20**(12):1920–1932.

- K. G. Jansky (1935). ‘A Note on the Source of Interstellar Interference’. *Proceedings of the Institute of Radio Engineers* **23**(10):1158–1163.
- V. Jelic, et al. (2008). ‘Foreground simulations for the LOFAR - Epoch of Reionization Experiment’. *Mon.Not.Roy.Astron.Soc.* **389**:1319–1335.
- C. H. Jordan, et al. (2017). ‘Characterization of the ionosphere above the Murchison Radio Observatory using the Murchison Widefield Array’. *Monthly Notices of the Royal Astronomical Society* **471**:3974–3987.
- R. Joseph (2019). ‘ronnijoseph/Beam-Perturbations: First Release’.
- R. C. Joseph, et al. (2018). ‘The Bias and Uncertainty of Redundant and Sky-based Calibration Under Realistic Sky and Telescope Conditions’. *The Astronomical Journal* **156**:285.
- S. M. Kay (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- S. Kazemi & S. Yatawatta (2013). ‘Robust radio interferometric calibration using the t-distribution’. *Monthly Notices of the Royal Astronomical Society* **435**(1):597–605.
- S. Kazemi, et al. (2011). ‘Radio interferometric calibration using the SAGE algorithm’. *Monthly Notices of the Royal Astronomical Society* **414**:1656–1666.
- N. S. Kern, et al. (2019). ‘Absolute Calibration Strategies for the Hydrogen Epoch of Reionization Array and Their Impact on the 21 cm Power Spectrum’.
- J. R. Kerrigan, et al. (2018). ‘Improved 21 cm Epoch of Reionization Power Spectrum Measurements with a Hybrid Foreground Subtraction and Avoidance Technique’. *The Astrophysical Journal* **864**(2):131.
- P. Kittiwisit, et al. (2017). ‘Sensitivity of the Hydrogen Epoch of Reionization Array and its build-out stages to one-point statistics from redshifted 21cm

- observations’. *Monthly Notices of the Royal Astronomical Society* **474**(4):4487–4499.
- M. Kolopanis, et al. (2019). ‘A Simplified, Lossless Reanalysis of PAPER-64’. *The Astrophysical Journal* **883**(2):133.
- L. Koopmans, et al. (2015). ‘The Cosmic Dawn and Epoch of Reionisation with SKA’. In *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, p. 1.
- J. Kraus (1950). *Antennas*. Electrical Engineering Series. McGraw-Hill.
- G. Lemaître (1927). ‘Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques’. *Annales de la Société Scientifique de Bruxelles* **47**:49–59.
- K. Levenberg (1944). ‘A method for the solution of certain non-linear problems in least squares’. *Quarterly of applied mathematics* **2**(2):164–168.
- W. Li, et al. (2019). ‘First Season MWA Phase II Epoch of Reionization Power Spectrum Results at Redshift 7’. *The Astrophysical Journal* **887**(2):141.
- W. Li, et al. (2018). ‘Comparing Redundant and Sky-model-based Interferometric Calibration: A First Look with Phase II of the MWA’. *The Astrophysical Journal* **863**(2):170.
- A. Lidz (2016). *Physics of the Intergalactic Medium During the Epoch of Reionization*, pp. 23–63. Springer International Publishing, Cham.
- J. L. B. Line, et al. (2018). ‘In situ measurement of MWA primary beam variation using ORBCOMM’. *Publications of the Astronomical Society of Australia* **35**.
- A. Liu, et al. (2014). ‘Epoch of reionization window. I. Mathematical formalism’. *Phys. Rev. D* **90**(2).
- A. Liu & J. R. Shaw (2019). ‘Data Analysis for Precision 21 cm Cosmology’ .

- A. Liu & M. Tegmark (2011). ‘A method for 21 cm power spectrum estimation in the presence of foregrounds’. *Physical Review D* **83**(10):103006.
- A. Liu, et al. (2010). ‘Precision calibration of radio interferometers using redundant baselines’. *Monthly Notices of the Royal Astronomical Society* **408**(2):10291050.
- A. Liu, et al. (2009). ‘Will point sources spoil 21-cm tomography?’. *Monthly Notices of the Royal Astronomical Society* **394**(3):1575–1587.
- R. C. Livermore, et al. (2017). ‘Directly Observing the Galaxies Likely Responsible for Reionization’. *The Astrophysical Journal* **835**:113.
- P. Madau, et al. (1997). ‘21 Centimeter Tomography of the Intergalactic Medium at High Redshift’. *The Astrophysical Journal* **475**:429–444.
- P. Madau, et al. (2004). ‘Early Reionization by Mini-quasars’. *The Astrophysical Journal* **604**:484–494.
- D. W. Marquardt (1963). ‘An algorithm for least-squares estimation of nonlinear parameters’. *Journal of the society for Industrial and Applied Mathematics* **11**(2):431–441.
- J. C. Mather, et al. (1994). ‘Measurement of the Cosmic Microwave Background Spectrum by the COBE FIRAS Instrument’. *The Astrophysical Journal* **420**:439.
- M. McQuinn (2016). ‘The Evolution of the Intergalactic Medium’. *Annual Review of Astronomy and Astrophysics* **54**:313–362.
- F. Mertens, et al. (2020). ‘Improved upper limits on the 21-cm signal power spectrum of neutral hydrogen at z 9.1 from LOFAR’. *Monthly Notices of the Royal Astronomical Society* .

- F. G. Mertens, et al. (2018). ‘Statistical 21-cm signal separation via Gaussian Process Regression analysis’. *Monthly Notices of the Royal Astronomical Society* **478**(3):3640–3652.
- A. Mesinger (2016). *Understanding the Epoch of Cosmic Reionization*, vol. 423. Springer International Publishing.
- A. Mesinger, et al. (2016). ‘The Evolution Of 21 cm Structure (EOS): public, large-scale simulations of Cosmic Dawn and reionization’. *Monthly Notices of the Royal Astronomical Society* **459**:2342–2353.
- M. Mevius, et al. (2016). ‘Probing ionospheric structures using the LOFAR radio telescope’. *Radio Science* **51**(7):927–941.
- M. Milosavljević & C. Safranek-Shrader (2016). *Star Formation for Predictive Primordial Galaxy Formation*, pp. 65–109. Springer International Publishing, Cham.
- I. F. Mirabel, et al. (2011). ‘Stellar black holes at the dawn of the universe’. *Astronomy and Astrophysics* **528**:A149.
- D. A. Mitchell, et al. (2008). ‘Real-Time Calibration of the Murchison Widefield Array’. *IEEE Journal of Selected Topics in Signal Processing* **2**:707–717.
- M. F. Morales, et al. (2006). ‘Improving Foreground Subtraction in Statistical Observations of 21 cm Emission from the Epoch of Reionization’. *The Astrophysical Journal* **648**:767–773.
- M. F. Morales, et al. (2012). ‘Four Fundamental Foreground Power Spectrum Shapes for 21 cm Cosmology Observations’. *The Astrophysical Journal* **752**:137.
- M. F. Morales & J. Hewitt (2004). ‘Toward Epoch of Reionization Measurements with Wide-Field Radio Observations’. *The Astrophysical Journal* **615**:7–18.

- M. F. Morales & J. S. B. Wyithe (2010). ‘Reionization and Cosmology with 21 cm Fluctuations’. *Ann.Rev.Astron.Astrophys.* **48**:127–171.
- B. Mort, et al. (2017). ‘Analysing the impact of far-out sidelobes on the imaging performance of the SKA-LOW telescope’. *Monthly Notices of the Royal Astronomical Society* **465**:3680–3692.
- S. G. Murray & C. M. Trott (2018). ‘The Effect of Baseline Layouts on the Epoch of Reionization Foreground Wedge: A Semianalytical Approach’. *The Astrophysical Journal* **869**:25.
- S. G. Murray, et al. (2017). ‘An Improved Statistical Point-source Foreground Model for the Epoch of Reionization’. *The Astrophysical Journal* **845**(1):7.
- P. J. Napier (1999). ‘The Primary Antenna Elements’. In *Synthesis Imaging in Radio Astronomy II*, vol. 180 of *Astronomical Society of the Pacific Conference Series*, p. 37.
- A. R. Neben, et al. (2015). ‘Measuring phased-array antenna beampatterns with high dynamic range for the Murchison Widefield Array using 137MHz ORB-COMM satellites’. *Radio Science* **50**(7):614–629.
- J. E. Noordam (2004). ‘LOFAR calibration challenges’. In J. M. Oschmann, Jr. (ed.), *Ground-based Telescopes*, vol. 5489 of *Proceedings of the SPIE*, pp. 817–825.
- P. Noorishad, et al. (2012). ‘Redundancy calibration of phased-array stations’. *Astronomy & Astrophysics* **545**:A108.
- P. Oesch, et al. (2016). ‘A remarkably luminous galaxy at $z= 11.1$ measured with Hubble Space Telescope grism spectroscopy’. *The Astrophysical Journal* **819**(2):129.

- P. A. Oesch, et al. (2015). ‘A Spectroscopic Redshift Measurement for a Luminous Lyman Break Galaxy at $z = 7.730$ Using Keck/MOSFIRE’. *The Astrophysical Journal Letters* **804**:L30.
- A. Offringa, et al. (2019a). ‘The impact of interference excision on 21-cm epoch of reionization power spectrum analyses’. *Monthly Notices of the Royal Astronomical Society* **484**(2):2866–2875.
- A. Offringa, et al. (2019b). ‘Precision requirements for interferometric gridding in the analysis of a 21 cm power spectrum’. *Astronomy & Astrophysics* **631**:A12.
- A. Offringa, et al. (2012). ‘A morphological algorithm for improving radio-frequency interference detection’. *Astronomy & astrophysics* **539**:A95.
- A. R. Offringa, et al. (2016). ‘Parametrizing Epoch of Reionization foregrounds: a deep survey of low-frequency point-source spectra with the Murchison Wide-field Array’. *Monthly Notices of the Royal Astronomical Society* **458**:1057–1070.
- S. P. Oh (2001). ‘Reionization by hard photons. I. X-rays from the first star clusters’. *The Astrophysical Journal* **553**(2):499.
- V. Ollier, et al. (2017). ‘Robust calibration of radio interferometers in non-Gaussian environment’. *IEEE Transactions on Signal Processing* **65**(21):5649–5660.
- V. Ollier, et al. (2018). ‘Robust distributed calibration of radio interferometers with direction dependent distortions’. *Signal Processing* **153**:348–354.
- N. Orosz, et al. (2019). ‘Mitigating the effects of antenna-to-antenna variation on redundant-baseline calibration for 21 cm cosmology’. *Monthly Notices of the Royal Astronomical Society* **487**:537–549.
- G. Paciga, et al. (2013). ‘A refined foreground-corrected limit on the HI power spectrum at $z= 8.6$ from the GMRT Epoch of Reionization Experiment’. *Mon. Not. Roy. Astron. Soc.* **433**(arXiv: 1301.5906):639.

- A. Parsons, et al. (2012a). ‘A SENSITIVITY AND ARRAY-CONFIGURATION STUDY FOR MEASURING THE POWER SPECTRUM OF 21 cm EMISSION FROM REIONIZATION’. *The Astrophysical Journal* **753**(1):81.
- A. R. Parsons, et al. (2010). ‘The Precision Array for Probing the Epoch of Reionization: Eight Station Results’. *The Astronomical Journal* **139**:1468–1480.
- A. R. Parsons, et al. (2012b). ‘A per-baseline, delay-spectrum technique for accessing the 21 cm cosmic reionization signature’. *The Astrophysical Journal* **756**(2):165.
- A. H. Patil, et al. (2017). ‘Upper Limits on the 21 cm Epoch of Reionization Power Spectrum from One Night with LOFAR’. *The Astrophysical Journal* **838**:65.
- A. H. Patil, et al. (2016). ‘Systematic biases in low-frequency radio interferometric data due to calibration: the LOFAR-EoR case’. *Monthly Notices of the Royal Astronomical Society* **463**:4317–4330.
- A. H. Patil, et al. (2014). ‘Constraining the epoch of reionization with the variance statistic: simulations of the LOFAR case’. *Monthly Notices of the Royal Astronomical Society* **443**:1113–1124.
- N. Patra, et al. (2013). ‘SARAS: a precision system for measurement of the cosmic radio background and signatures from the epoch of reionization’. *Experimental Astronomy* **36**:319–370.
- A. A. Penzias & R. W. Wilson (1965). ‘A Measurement of Excess Antenna Temperature at 4080 Mc/s.’. *The Astrophysical Journal* **142**:419–421.
- R. A. Perley (1989). ‘High Dynamic Range Imaging’. In *Synthesis Imaging in Radio Astronomy*, vol. 6 of *Astronomical Society of the Pacific Conference Series*, p. 287.

- Planck Collaboration, et al. (2016). ‘Planck 2015 results. XIII. Cosmological parameters’. *Astronomy and Astrophysics* **594**:A13.
- D. C. Price, et al. (2018). ‘Design and characterization of the Large-aperture Experiment to Detect the Dark Age (LEDA) radiometer systems’. *Monthly Notices of the Royal Astronomical Society* **478**(3):4193–4213.
- J. Pritchard & A. Loeb (2008). ‘Evolution of the 21cm signal throughout cosmic history’. *Physical review D* **78**(10):103511.
- J. R. Pritchard & S. R. Furlanetto (2006). ‘Descending from on high: Lyman-series cascades and spin-kinetic temperature coupling in the 21-cm line’. *Monthly Notices of the Royal Astronomical Society* **367**(3):1057–1066.
- J. R. Pritchard & S. R. Furlanetto (2007). ‘21-cm fluctuations from inhomogeneous X-ray heating before reionization’. *Monthly Notices of the Royal Astronomical Society* **376**:1680–1694.
- J. R. Pritchard & A. Loeb (2012). ‘21 cm cosmology in the 21st century’. *Reports on Progress in Physics* **75**(8):086901.
- P. Procopio, et al. (2017). ‘A High-Resolution Foreground Model for the MWA EoR1 Field: Model and Implications for EoR Power Spectrum Analysis’. *Publications of the Astronomical Society of Australia* **34**:e033.
- U. Rau, et al. (2009). ‘Advances in Calibration and Imaging Techniques in Radio Interferometry’. *Proceedings of the IEEE* **97**(8):14721481.
- H. P. Robertson (1935). ‘Kinematics and World-Structure’. *The Astrophysical Journal* **82**:284.
- S. Salvini & S. J. Wijnholds (2014). ‘Fast gain calibration in radio astronomy using alternating direction implicit methods: Analysis and applications’. *Astronomy and Astrophysics* **571**:A97.

- P. A. Shaver, et al. (1999). ‘Can the reionization epoch be detected as a global signature in the cosmic background?’. *Astronomy and Astrophysics* **345**:380–390.
- J. R. Shaw, et al. (2014). ‘ALL-SKY INTERFEROMETRY WITH SPHERICAL HARMONIC TRANSIT TELESCOPES’. *The Astrophysical Journal* **781**(2):57.
- J. R. Shaw, et al. (2015). ‘Coaxing cosmic 21 cm fluctuations from the polarized sky using m-mode analysis’. *Physical Review D* **91**(8):083514.
- J. L. Sievers (2017). ‘Calibration of Quasi-Redundant Interferometers’ .
- M. Sokolowski, et al. (2017). ‘Calibration and Stokes Imaging with Full Embedded Element Primary Beam Model for the Murchison Widefield Array’. *Publications of the Astronomical Society of Australia* **34**:e062.
- P. Stoica & T. L. Marzetta (2001). ‘Parameter estimation problems with singular information matrices’. *IEEE Transactions on Signal Processing* **49**(1):87–90.
- I. S. Sullivan, et al. (2012). ‘Fast Holographic Deconvolution: A New Technique for Precision Radio Interferometry’. *The Astrophysical Journal* **759**:17.
- R. A. Sunyaev & Y. B. Zeldovich (1970). ‘Small-Scale Fluctuations of Relic Radiation’. *Astrophysics and Space Science* **7**(1):3–19.
- R. A. Sunyaev & Y. B. Zeldovich (1972). ‘The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies’. *Comments on Astrophysics and Space Physics* **4**:173.
- G. Swarup, et al. (1991). ‘The Giant Metre-Wave Radio Telescope’. *Current Science* **60**:95.
- A. R. Thompson, et al. (2017). *Introductory Theory of Interferometry and Synthesis Imaging*, pp. 59–88. Springer International Publishing, Cham.

- S. J. Tingay, et al. (2013). ‘The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies’. *Publications of the Astronomical Society of Australia* **30**.
- C. M. Trott, et al. (2020). ‘Deep multi-redshift limits on Epoch of Reionisation 21 cm Power Spectra from Four Seasons of Murchison Widefield Array Observations’. *Monthly Notices of the Royal Astronomical Society* .
- C. M. Trott, et al. (2016). ‘CHIPS: The Cosmological H i Power Spectrum Estimator’. *The Astrophysical Journal* **818**:139.
- C. M. Trott & R. B. Wayth (2016). ‘Spectral Calibration Requirements of Radio Interferometers for Epoch of Reionisation Science with the SKA’. *Publications of the Astronomical Society of Australia* **33**:e019.
- C. M. Trott & R. B. Wayth (2017). ‘Building Models for Extended Radio Sources: Implications for Epoch of Reionisation Science’. *Publications of the Astronomical Society of Australia* **34**:e061.
- C. M. Trott, et al. (2012). ‘THE IMPACT OF POINT-SOURCE SUBTRACTION RESIDUALS ON 21 cm EPOCH OF REIONIZATION ESTIMATION’. *The Astrophysical Journal* **757**(1):101.
- P. H. van Cittert (1934). ‘Die wahrscheinliche Schwingungsverteilung in einer von einer Lichtquelle direkt oder mittels einer Linse beleuchteten Ebene’. *Physica* **1**(1-6):201–210.
- H. van de Hulst (1945). ‘Radiogolven uit het wereldruim’. *Nederlands Tijdschrift voor Natuurkunde* **11**:210–221.
- M. P. van Haarlem, et al. (2013). ‘LOFAR: The LOW-Frequency ARray’. *Astronomy and Astrophysics* **556**:A2.
- R. J. van Weeren, et al. (2016). ‘LOFAR Facet Calibration’. *The Astrophysical Journal Supplemental* **223**:2.

- H. Vedantham, et al. (2012). ‘Imaging the Epoch of Reionization: Limitations from Foreground Confusion and Imaging Algorithms’. *The Astrophysical Journal* **745**:176.
- A. G. Walker (1937). ‘On Milne’s Theory of World-Structure*’. *Proceedings of the London Mathematical Society* **s2-42**(1):90–127.
- F. Wang, et al. (2019). ‘Exploring Reionization-era Quasars. III. Discovery of 16 Quasars at $6.4 < z < 6.9$ with DESI Legacy Imaging Surveys and the UKIRT Hemisphere Survey and Quasar Luminosity Function at $z \sim 6.7$ ’. *The Astrophysical Journal* **884**(1):30.
- R. B. Wayth, et al. (2018). ‘The Phase II Murchison Widefield Array: Design overview’. *Publications of the Astronomical Society of Australia* **35**.
- M. H. Wieringa (1992). ‘An investigation of the telescope based calibration methods?redundancy? and?self-cal?’. *Experimental Astronomy* **2**(4):203225.
- S. Wijnholds, et al. (2010). ‘Calibration challenges for future radio telescopes’. *IEEE Signal Processing Magazine* **27**:30–42.
- S. J. Wijnholds, et al. (2016). ‘Calibration artefacts in radio interferometry - II. Ghost patterns for irregular arrays’. *Monthly Notices of the Royal Astronomical Society* **457**:2331–2354.
- M. J. Wilensky, et al. (2019). ‘Absolving the SSINS of Precision Interferometric Radio Data: A New Technique for Mitigating Faint Radio Frequency Interference’. *Publications of the Astronomical Society of the Pacific* **131**(1005):114507.
- W. L. Williams, et al. (2016). ‘LOFAR 150-MHz observations of the Boötes field: catalogue and source counts’. *Monthly Notices of the Royal Astronomical Society* **460**:2385–2412.

- S. A. Wouthuysen (1952). ‘On the excitation mechanism of the 21-cm (radio-frequency) interstellar hydrogen emission line.’. *The Astronomical Journal* **57**:31–32.
- S. Yatawatta (2015). ‘Distributed radio interferometric calibration’. *Monthly Notices of the Royal Astronomical Society* **449**:4506–4514.
- S. Yatawatta, et al. (2017). ‘Data multiplexing in radio interferometric calibration’. *Monthly Notices of the Royal Astronomical Society* **475**(1):708–715.
- S. Yatawatta, et al. (2009). ‘Radio Interferometric Calibration Using The SAGE Algorithm’. In *Proc. IEEE 13th Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop*, pp. 150–155.
- M. Zaldarriaga, et al. (2004). ‘21 Centimeter fluctuations from cosmic gas at high redshifts’. *The Astrophysical Journal* **608**(2):622.
- F. Zernike (1938). ‘The concept of partial coherence in optics’. *Physica* **5**:785–795.
- H. Zheng, et al. (2014). ‘MITEoR: a scalable interferometer for precision 21 cm cosmology’. *Monthly Notices of the Royal Astronomical Society* **445**(2):10841103.

Every reasonable effort has been made to acknowledge the owners of copyright material. I would be pleased to hear from any copyright owner who has been omitted or incorrectly acknowledged.