

School of Electrical Engineering, Computing and Mathematical Sciences

Advanced statistical methods for detecting the Epoch of Reionisation

Jaiden Hue Cook

0000-0001-8001-0791

This thesis is presented for the Degree of

Doctor of Philosophy (Physics)

of

Curtin University

October 2023

DECLARATION

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.

Jaiden Hue Cook

13 October 2023



“Strength does not come from winning. Your struggles develop your strengths. When you go through hardships and decide not to surrender, that is strength.”

— Arnold Schwarzenegger

We acknowledge that Curtin University works across hundreds of traditional lands and custodial groups in Australia, and with First Nations people around the globe. We wish to pay our deepest respects to their ancestors and members of their communities, past, present, and to their emerging leaders. Our passion and commitment to work with all Australians and peoples from across the world, including our First Nations peoples are at the core of the work we do, reflective of our institutions' values and commitment to our role as leaders in the Reconciliation space in Australia.

ACKNOWLEDGEMENTS

There are a lot of people I would like to thank, and acknowledge throughout the course of my PhD. Firstly I would like to thank my supervisors Cath and Jack. I could not have asked for a better combination of supervisors, and I am grateful for their combined patience and support throughout my journey. I want to thank Cath in particular for choosing me as a PhD student, I am honoured to have worked with one the greats in this field, and look forward to future collaborations. Furthermore, Cath's empathy and support has meant a lot to me, and I appreciate her immense level of care for all the students at CIRA, we are lucky to have her. I want to thank Jack, to whom I say "know your body". Jack is one of the most hard working individuals I have ever met, and is always willing to help others when they are need, often to his detriment. I will miss him when he leaves.

I would like to thank the whole EoR group at CIRA for being a welcoming and fun group of people, without whom academia and science would be a less welcoming and interesting place! To mention them by name: Aishwarya, Anshu, Ben, Chris, Christene, Dev, Himanshu, Kariuki, Jishnu, Nichole, Randall, Ravi, and Ridhima.

There are a few special mentions. I am contractually obliged to mention Mike, to whom I owe my first born child, and possibly my soul. My only regret in our transaction is that I did not agree to more stroopwaffles. I would also like to mention my fellow PhD student Sandro, whose cheeky grin, and proclivity towards mischief

never failed to both inconvenience me, and lighten my mood. I will always reflect fondly upon memories of our shared summer studentships, and the late nights working on our astrostats assignments. I would like to make a special mention to Susmita, who provided emotional support throughout the last part of our shared PhD journey. Thanks for being a good friend, and being someone I could talk to, the short stack would not be complete without you. I would also like to thank Callan for being a good deskmate, our time together was too short, however we will always have Stefanie's. I won't however miss the casual consumption of tomatoes as if they are apples (madness!). I want to thank Mawson, for accompanying me on numerous climbing expeditions, and hearing me complain about my various troubles.

I would also like to mention a few past and present CIRA students, who I am thankful to have shared my journey with: Angie, Ben Q, Danica, Elliott, Ferry, Kat, Keegan, Nick, Pikky, Ronniy, Silvia, Stefan, Steve, Torrance, and Tyrone.

I would also like to thank my parents, Cinnamon and Don, for both supporting and encouraging me through my early undergraduate career. I also want to honour my younger siblings Jessica and Jake, for whom I sacrificed time with to complete this work. To those friends and family not mentioned by name, I thank you as well for being a part of my journey.

Last but definitely not least, I would like to thank my loving partner and best friend Chantelle, who has supported me throughout my entire PhD journey through all my ups and downs. I appreciate her unconditional support and love, especially the past few months, where she has helped support me through her acts of kindness to support me through my final PhD push. I look forward to our continued journey together.

The research conducted in this thesis was made possible by the support of the Australian Government Research Training (RTP) Scholarship.

ABSTRACT

The Epoch of Reionisation (EoR) is the period in cosmic history when the first luminous objects reionised the predominantly neutral hydrogen intergalactic medium. This period in cosmic time contains information about these sources, and the evolution of cosmic structures in the Universe. Direct observations during the EoR are limited, and thus the Universe during and before the EoR is poorly understood. The spin flip 21cm transition line of neutral hydrogen offers a potential probe of the Universe during these periods. However, it is too faint to be directly observed with the current generation of radio telescopes. Most scientific effort has therefore been focused on measuring the statistics of the signal, through either the global signal (eg. EDGES, SARAS) or the two-point spatial statistics (eg. MWA, LOFAR, HERA). The spatial statistic experiments measure the power spectrum (Fourier transform of the two-point correlation function), which probes the spatial variance (Gaussianity) of the 21cm brightness temperature. The measured 21cm signal however is contaminated by Galactic and extra-galactic foregrounds which are up to five orders of magnitude brighter than the expected signal. To mitigate these effects the Murchison Widefield Array (MWA) observes fields with relatively few foregrounds that can be subtracted from the data. However, widefield sources in the highly attenuated MWA primary beam sidelobes such as Centaurus A, and Galactic plane supernova remnants are bright enough to still be significant. These sources are additionally in highly chromatic parts of the

primary beam, which contaminate the line-of-sight power spectrum modes. The cosmological 21cm signal is additionally expected to be highly non-Gaussian prior to and during reionisation. Some of the non-Gaussianities manifest from heating due to X-ray sources, and from the topology of ionisation. Higher order statistics therefore encode astrophysical information about the first luminous sources not contained in the power spectrum. Much effort in the literature has been to characterise the third order statistical moment by calculating the bispectrum (Fourier transform of the three-point correlation function) of simulated data. However, the bispectrum is computationally expensive to calculate and typically has a lower signal to noise than the power spectrum.

This thesis investigates the widefield foreground contamination, and the non-Gaussianities of both the expected 21cm signal and foreground sources. To facilitate this I developed multi-component 2D Gaussian sky-models for the extended widefield sources, and a visibility simulation and power spectrum estimation pipeline. Using this pipeline, I determined that widefield sources will need to be removed from some MWA observations to detect the expected cosmological 21cm signal. I then characterised the non-Gaussianities of the 21cm signal, by calculating the skew spectrum. The skew spectrum is the cross spectrum of the squared temperature field with the temperature field. We generated a set of seven 21cm simulations, which varied the halo mass cut off for ionising photons, the X-ray luminosity per unit star formation rate, and the X-ray energy threshold. To isolate the non-Gaussianities present in the skew spectrum, I normalise it by the power spectrum, effectively removing the Gaussian amplitude contribution. We find unique features in the normalised skew spectrum that coincide with the characteristic bubble size during reionisation, and the first luminous sources during the epoch of heating. Furthermore, I develop an analytical framework for estimating the instrumental MWA skew spectrum. We apply this method to a fiducial simulation of the 21cm signal, and to realistic simulations of foreground MWA EoR fields. We compare both instrumental skew spectra to estimates derived from the foreground lightcone, and the fiducial 21cm lightcone respectively. We find we can retrieve the non-Gaussianity for both the foregrounds and the 21cm signal, however the instrument impart chromatic

structure, similar to the leakage found in the power spectrum.

CONTENTS

1	Introduction	1
1.1	Cosmology	2
1.2	Epoch of Reionisation	3
1.3	Constraints on the EoR	4
1.3.1	Cosmic Microwave Background	4
1.3.2	Quasars	6
1.3.3	Observations of High Redshift Galaxies	6
1.4	Sources of Ionisation	7
1.5	Thesis Motivation & Outline	9
2	21cm Line Physics	11
2.1	Radiative Transfer	12
2.2	21cm Differential Temperature Brightness	15
2.3	Spin Temperature	18
2.4	Heating of the Intergalactic Medium	21
2.5	Summary	23
3	Radio Interferometry & Simulation	25
3.1	Spatial Coherence Function	26

3.1.1	Synthesis Imaging	30
3.1.2	Sampling	31
3.1.3	Primary Beam Response	32
3.1.4	Polarimetry and the Measurement Equation	33
3.2	Murchison Widefield Array	37
3.2.1	Fully Element Embedded (FEE) Primary Beam	37
3.3	Simulating Visibilities	39
3.3.1	Array Geometry and Baseline Distribution	40
3.3.2	WODEN	42
3.4	Summary	47
4	Stochastic Processes	49
4.1	Ensemble Average	50
4.2	Correlation Functions	51
4.3	Stationary and Ergodic Processes	52
4.3.1	Ergodic Processes	53
4.3.2	Power Spectral Density	53
4.3.3	Two Point Fourier Correlation	55
4.3.4	Higher Order Poly Spectra	56
4.4	Summary	57
5	21cm Statistics	59
5.1	Cosmological Fields	60
5.2	21cm Power Spectrum	61
5.2.1	Astrophysics of the 21cm Power Spectrum	62
5.2.2	21cm Power Spectrum During EoR	64
5.3	21cm Bispectrum	66
5.3.1	Understanding the Bispectrum	66
5.3.2	21cm Bispectrum During EoR	68
5.3.3	Skew Spectrum	69

5.4	Interferometry and Power Spectrum	71
5.4.1	Power Spectrum	72
5.4.2	Foreground Contamination and Wedge	72
5.5	21cm Power Spectrum Experiments	77
5.5.1	MWA	77
5.5.2	LOFAR	77
5.5.3	HERA	78
5.6	Summary	78
6	SNRs, Centaurus A and the 21cm EoR Signal	79
6.1	Introduction	81
6.2	Methodology	84
6.2.1	Fourier Sky Cube	85
6.2.2	Simulating Visibilities	86
6.2.3	Gridding	88
6.2.4	Calculating the 1D and 2D Power Spectra	90
6.2.5	The Fiducial 21cm Signal	91
6.3	Data & Morphological Models	93
6.3.1	Centaurus A	94
6.3.2	Supernova Remnants	98
6.3.3	Constructing Sky-Models	103
6.4	Results	106
6.4.1	Sidelobe and Null Test Observations	106
6.4.2	Partial Sky Models	110
6.5	Discussion	114
6.5.1	Future Work	116
6.6	Conclusions	118
6.A	Cosmological Conversion	120
6.A.1	Conversion Factor	121
6.B	Thermal Noise	121

6.C	2D Gaussian Projection Approximation	122
7	21cm Skew Spectrum	125
7.1	Introduction	127
7.2	Power Spectrum and Skew Spectrum	131
7.2.1	Power Spectrum	131
7.2.2	Bispectrum and Skew Spectrum	132
7.3	MERAXES	133
7.3.1	A Realistic Galaxy Population	133
7.3.2	IGM Ionisation State	135
7.3.3	21cm Signal	136
7.3.4	Simulations	138
7.4	Results	140
7.4.1	21cm Lightcones	141
7.4.2	Neutral Fraction and Ionisation	143
7.4.3	21cm Statistics	144
7.5	Power Spectrum and Skew Spectrum During Reionisation	154
7.5.1	Normalised Skew Spectrum	154
7.5.2	Power Spectrum, Skew Spectrum, and the Normalised Skew Spectrum	155
7.6	Detectability of the Normalised Skew Spectrum	159
7.6.1	Cosmic Variance of the Power Spectrum and the Skew Spectrum	160
7.6.2	Uncertainty in the Normalised Skew Spectrum	162
7.7	Discussion and Conclusion	166
8	Instrumental Foreground & 21cm Skew Spectrum	171
8.1	Introduction	172
8.2	Foreground Sky-Model	173
8.2.1	GLEAM Year 1	175
8.2.2	LoBES	176

8.2.3	PUMA	177
8.3	Fiducial 21cm Model	178
8.3.1	WODEN Simulation	179
8.4	Instrumental Skew Spectrum	180
8.4.1	Apparent Quadratic Sky Brightness	181
8.4.2	Skew Visibilities	182
8.4.3	Estimating the Skew Spectrum	185
8.5	Foreground Spectra	187
8.5.1	Foreground Power Spectrum	187
8.5.2	Foreground Skew Spectrum	190
8.5.3	Foreground Normalised Skew Spectrum	192
8.6	Fiducial 21cm Spectra	193
8.6.1	Power Spectrum	194
8.6.2	Skew Spectrum	196
8.6.3	Normalised Skew Spectrum	199
8.7	Foregrounds and 21cm Poly Spectra	201
8.7.1	Window Power Spectra	201
8.7.2	Window Skew Spectra	201
8.7.3	Window Normalised Skew Spectra	203
8.8	Discussion and Conclusion	206
8.A	Analytic Toy Model	209
8.A.1	Analytic Power Spectrum	209
8.A.2	Analytic Skew Spectrum	211
8.A.3	Analytic Normalised Skew Spectrum	213
8.A.4	Toy Model Test	215
8.B	2D Radial Blackman-Harris Window	216
8.B.1	Top Hat Radial Blackman-Harris Window	217
8.C	Skew Spectrum Conversion Calculation	217
9	Summary and Conclusion	219

LIST OF FIGURES

1.1	The spectra of the 19 quasars observed by Fan et al. (2006).	5
2.1	Diagram of the hyperfine transition of an electron in a spin parallel (excited state) to a spin anti-parallel (ground state) for the hydrogen atom.	12
2.2	Redshift evolution of the spin temperature (solid red line), the kinetic temperature (dotted blue line), and the background radiation temperature (dashed green line).	20
2.3	Shows how the Wouthuysen-Field effect mixes the spin states through the hyperfine splitting of the $1S$ and $2P$ states.	21
3.1	Coordinate system for the electric field measured at $\mathbf{E}(\mathbf{r})$ as emitted from some source at \mathbf{R} from the origin.	27
3.2	The coordinate system for a single baseline $\mathbf{r}_1 - \mathbf{r}_2$	28
3.3	The coordinate system for a single baseline.	30
3.4	Example MWA primary beam response projected onto the (l, m) orthographic image plane.	38
3.5	Right-handed Cartesian coordinate system defined relative Hour angle and declination.	41
3.6	The relationship between the (u, v, w) coordinate system, the $(X_\lambda, Y_\lambda, Z_\lambda)$ coordinate system, and the celestial sphere (H, δ)	42

3.7	MWA Phase I, zenith phased (u, v) baseline distribution for a single time step and frequency (~ 180 MHz) channel.	43
5.1	Example 1D spherically averaged dimensionless power spectrum of a simulated 21cm signal taken from Balu et al. (2023).	62
5.2	Example of the global signal temperature $\delta\bar{T}_b$ and the power at two fixed spatial scales $k = 0$	63
5.3	An example of the bispectrum for the equilateral triangle configuration, taken from Lewis (2011).	67
5.4	Example of the 2D power spectrum, showing the different regions affected by instrument chromaticity.	73
5.5	Upper limits measured from ongoing and past radio interferometric EoR experiments (references therein).	76
6.1	Simulation pipeline flow chart.	86
6.2	Comparison of the sampled fiducial signal (solid black line), and the reconstructed EoR window fiducial signal from a spherically averaged 1D power spectrum (dashed red line).	92
6.3	Model SIN projected image of Centaurus A fit to the deep multi-scale image from (McKinley et al., 2021).	95
6.4	GLEAM cutout images of Puppis A (Subfigure (a)) and Vela (Subfigure (b)) at 200 MHz.	100
6.5	Example apparent sky-model images for sidelobe sky-model (Subfigure (a)) and the null sky-model (Subfigure (b)).	105
6.6	2D power spectra for the sidelobe, null simulation sky-models, and the fiducial 21 cm 2D power spectrum.	108
6.7	2D Power spectrum of the sidelobe simulation without CenA in the sky-model.	109
6.8	1D power spectra for a series of partial input sky-models, showing linearly spaced k bin widths.	112

6.9	(l,m) plane of the visible celestial sphere.	123
6.10	Ellipse in the non-offset rotated frame.	123
7.1	21cm differential brightness temperature lightcone slices for each simulation volume as a function of redshift from $z = 20 - 5$	141
7.2	The ionisation history of the average neutral hydrogen (neutral fraction \bar{x}_{HI}) IGM calculated for each simulation (labelled).	142
7.3	Mean brightness temperature for all simulations, calculated from the coeval boxes as a function of redshift.	144
7.4	The PS as a function of redshift for all simulations at a fixed spatial scale.	145
7.5	The SS as a function of redshift for all simulations at a fixed spatial scale.	147
7.6	The dimensionless SS plotted against the dimensionless PS as a function of redshift for each simulation at a fixed spatial scale $k \sim 0$	150
7.7	PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), mid- M_h (double dotted dash line), and the high- M_h (dash dotted line).	153
7.8	PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), low- L_X (double dotted dash line), and the high- L_X (dash dotted line).	155
7.9	PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), low- E_0 (double dotted dash line), and the high- E_0 (dash dotted line).	156
7.10	The numerically estimated statistical uncertainties on PS (a) and the SS (b) (solid line), compared to the expected Gaussian uncertainties (dashed line), and the estimated non-Gaussian component (dotted line).	161
7.11	Correlation coefficient between the Fiducial PS and SS for the $x_{\text{HI}} \sim 0$.	163
7.12	The first order propagated uncertainties on the normalised SS (solid black line), compared to the Gaussian model uncertainties (dashed line), and the estimated non-Gaussian component (dotted line), for the Fiducial simulation at $\bar{x}_{\text{HI}} = 0$	164

7.13	Signal to noise ratio of $\gamma(k)$ for the Fiducial simulation scaled by the expected SKA_LOW observing comoving volume.	166
8.1	EoR0 (panel (a)) and EoR1 (panel (b)) WODEN simulated foreground 2D PS.	188
8.2	The 1D PS for the EoR0 (dashed dot line) and the EoR1 (dashed line) fields.	189
8.3	EoR0 (a) and EoR1 (b) 2D foreground SS with the primary beam attenuation.	190
8.4	The 1D SS for the EoR0 (dashed dot line) and the EoR1 (dashed line) fields.	191
8.5	The normalised SS for the EoR0 (dash dotted line) and the EoR1 (dashed line) fields.	192
8.6	Spherically averaged 1D PS of the Fiducial 21cm signal.	193
8.7	The 2D PS for the Fiducial signal (panel (a)), the instrumental WODEN simulated 2D PS (panel (b)), the colourbar scale is the same for panels (a) and (b).	195
8.8	Mean 2D PS ratio for the fiducial 21cm signal, and the instrumental simulated 21cm signal, this is shown with the solid black line.	196
8.9	Spherically averaged 1D SS of the Fiducial 21cm signal.	197
8.10	The 2D SS for the Fiducial signal (panel (a)), the instrumental simulation 2D SS (panel (b)), the colourbar scale is the same for panels (a) and (b).	198
8.11	Mean 2D SS ratio for the fiducial 21cm signal, and the instrumental simulated 21cm signal, this is shown as the solid black line.	199
8.12	Normalised SS of the fiducial 21cm signal.	200
8.13	WODEN simulated EoR0, EoR1, and 21cm 2D PS with wedge cuts applied at $k_{\perp} \leq 0$	202
8.14	Compares the EoR0 (dash dotted line), EoR1 (dashed line) and 21cm (solid black line) window power.	203
8.15	Shows the SS for the EoR0 (dash dotted line), the EoR1 (dashed line), and the simulated 21cm signal (solid line).	204

8.16 Shows the normalised SS for the EoR0 (dash dotted line), the EoR1 (dashed line), and the simulated 21cm signal (solid line). 205

8.17 Expected point source normalised SS for random 50 point source model (solid black), compared to the toy model (cyan). 216

LIST OF TABLES

7.1	Astrophysical parameter summary for the seven MERAXES simulations. . . .	138
-----	--	-----

CHAPTER 1

INTRODUCTION

By happenstance, or virtue of the anthropic principle, we are privileged to live in a Universe where we can peer back in time to the earliest epochs of cosmic history. The earliest observable epoch shows the Universe is dominated by a hot dense almost uniform plasma, suggesting the Universe underwent a phase of rapid exponential expansion known as cosmic inflation. Approximately 300,000 years after cosmic inflation, the hot dense plasma condensed to form neutral hydrogen, in a process known as recombination. The Universe following recombination is opaque to some forms of electromagnetic radiation due to the abundance of neutral hydrogen, thus beginning the *cosmic dark ages*. From the dark ages emerge the first stars and galaxies formed through gravitational collapse of primordial hydrogen, signalling the *cosmic dawn*. These sources grow in number, illuminating and eventually reionising the entire Universe. The rapid transition of the Universe from a neutral to ionised state is called the Epoch of Reionisation (EoR). Observations of neutral hydrogen during the EoR promise to elucidate our understanding of the first stars, galaxies, and the cosmos. However, detecting the signals from early hydrogen is a monumental challenge, one that requires high precision experiments in the presence of significant systematics. Numerous methods, primarily statistical ones, along with many current and future instruments have

taken up the mantle for detecting the signal.

This thesis aims to understand the systematics, by modelling astrophysical sources of contamination, and through developing new statistics, that we can test with simulations of the neutral hydrogen signal. We apply these methods in the context of ongoing instruments, to better understand what astrophysical information we are sensitive to during the EoR.

1.1 Cosmology

Our current understanding of the Universe and its evolution through cosmic time is built upon the framework of Einstein's Theory of General Relativity (GR) (Einstein, 1915). GR describes how matter and energy affect the curvature of spacetime, which in turn dictates how matter moves through spacetime. Friedmann (1922); Lemaître (1927); Robertson (1935); Walker (1937) all independently solved the equations of GR for the entire Universe under the assumptions of homogeneity and isotropy. The resulting solution is called the FLRW-metric, and describes how the Universe evolves for a given matter/energy distribution, and geometry.

During the early period of the field of modern cosmology, the Universe was thought to be in a steady state, which famously resulted in Einstein introducing the cosmological constant Λ to 'correct' his equations of GR. However, observations of distant galaxies by Hubble (1929) revealed they were receding at a velocity proportional to their distance (known as Hubble's law), thus implying that the Universe was expanding. Furthermore, the discovery of the cosmic microwave background (CMB) by Penzias and Wilson (1965) found the sky to have an incredibly uniform temperature of 2.7 K that was close to a perfect blackbody. The uniformity of the CMB suggested a rapid expansion phase at early times in the Universe. Starobinsky (1980) and Guth (1981) developed the inflationary model, by linking the expansion of the Universe to the collapse of an initial quantum state.

The CMB and inflation however only capture part of the cosmological framework. Observations of the rotation curves of galaxies by Rubin and Ford (1970); Rubin et al.

(1978), revealed the existence of *dark matter*. Dark matter does not appear to interact with ordinary matter except through gravitation, and overwhelmingly makes up the bulk of matter in the Universe. Furthermore, independent observations of distant type Ia supernovae by Riess et al. (1998) and Perlmutter et al. (1999) revealed the accelerating expansion of the Universe, and thus the existence of *dark energy*, named as such due to the unknown mechanism driving the expansion.

The standard Λ CDM (CDM stands for cold dark matter) model of cosmology describes the Universe through the FLRW-metric, in the context of the energy densities of dark matter, dark energy, and baryonic matter. These different contributions parameterise the model, and can be inferred through observations of the CMB, and Type Ia supernovae. Observations of the anisotropies of the CMB by the Planck Collaboration et al. (2020) indicate that dark energy accounts for $\sim 69\%$, dark matter for $\sim 26\%$, and baryonic matter $\sim 5\%$ of the energy density respectively. We assume this standard cosmological model throughout this work.

1.2 Epoch of Reionisation

The EoR is the period in cosmic history when the first stars, galaxies, and compact objects ionised the predominantly neutral hydrogen inter galactic medium (IGM). This represents the last significant state change of the Universe after recombination, to the current era, and occurred from the redshift of $z \sim 10$ to $z \sim 5$ (~ 0.7 Gyr), approximately one billion years after cosmic inflation (Peacock, 1998). During this period, ionising radiation, predominantly UV photons around the first stars, carved out large ionised bubbles in the surrounding neutral hydrogen. These bubbles with sizes of $\sim 1 - 10$ Mpc grew until they eventually overlapped, fully ionising the remaining neutral hydrogen (Furlanetto et al., 2004b). The history, morphology and timing of reionisation encodes a plethora of valuable astrophysical and cosmological information (Zaroubi, 2013; Furlanetto, 2019; Wise, 2019). As such, understanding the sources responsible for reionisation, and the structure formation of the Universe that led to the formation of these sources, is currently a paramount scientific pursuit. In this section

we discuss the constraints, probes, and sources of ionisation during the EoR.

1.3 Constraints on the EoR

1.3.1 Cosmic Microwave Background

CMB photons interact with free electrons through Thomson scattering as they propagate through the Universe over cosmological distances (Zeldovich and Sunyaev, 1969). In hot cluster environments, these CMB photons can be inverse Compton scattered by hot inter cluster electrons, through a process known as the thermal Sunyaev-Zel'dovich (SZ) effect (Sunyaev and Zeldovich, 1972). This occurs because these electrons are at extremely high temperatures ($\sim 10^7$ K); this is known as the thermal SZ effect. The kinetic SZ effect occurs when the velocity streams of ionised bubbles Doppler shifts the CMB photons. This in turn embeds information regarding reionisation, and can be used to constrain the duration of reionisation as well as the evolution of the ionisation fraction (Mesinger et al., 2012). The primary action of this process is that the CMB photons are affected by the total column density of free electrons along the line of sight. This effect is parameterised by the Thomson scattering optical depth τ . Measurements of the reionisation optical depth from CMB observation by Planck Collaboration et al. (2021) imply a midpoint of reionisation at $z = 7.68 \pm 0.79$ for a 68% confidence interval, with an optical depth of $\tau = 0.062 \pm 0.007$ (as updated from Heinrich and Hu (2021)). However, there is disagreement in the optical depth and timing of reionisation, with earlier measurements from WMAP¹ nine year data indicating $\tau = 0.089 \pm 0.014$, and $z = 10.36$ (Hinshaw et al., 2013). Work by Weiland et al. (2018) and Millea, Marius and Bouchet, François (2018) show that systematics, and assumptions in the priors of the cosmological parameter fits, can bias the determination of the optical depth.




Figure 1.1: The spectra of the 19 quasars observed by Fan et al. (2006). Sources are ordered from highest to lowest redshift (top to bottom). The Gunn-Peterson trough is visible for each source. This Figure was retrieved from Fan et al. (2006)

1.3.2 Quasars

Quasars can be used to constrain the end of reionisation through the Lyman- α (from hereon $\text{Ly}\alpha$) forest. $\text{Ly}\alpha$ emission is the transition from the $n = 2$ to the $n = 1$ energy level in neutral hydrogen. Quasars (or quasi stellar objects) are incredibly luminous active galactic nuclei, powered by the accretion of the central super massive black hole. These luminous sources can be seen out to high redshift, some during the EoR. As ultra-violet (UV) photons emitted from these objects propagate throughout the Universe, they undergo cosmological redshift through the $\text{Ly}\alpha$ line. Ground state hydrogen atoms in clouds along the line of sight absorb these photons. Absorption can happen for any neutral hydrogen cloud along the line of sight. Since these clouds are at different distances, the resulting absorption lines appear at different wavelengths; this is referred to as the $\text{Ly}\alpha$ forest. The Gunn-Peterson trough occurs because the optical depth of $\text{Ly}\alpha$ photons is relatively high for neutral hydrogen, and is quickly suppressed when the neutral fraction of hydrogen is $x_{\text{HI}} \geq 10^{-4}$ (Gunn and Peterson, 1965). $\text{Ly}\alpha$ saturates easily when there is only a small amount of neutral hydrogen, resulting in a trough at the redshift of the emitted light from that particular quasar. As such $\text{Ly}\alpha$ is a good probe of the end of reionisation.

Evidence of quasars with this trough at $z > 6$, and the absence of this trough for most quasars below $z < 6$, is evidence of the Universe undergoing a state transition from one that is predominantly neutral, to an ionised one. In Figure 1.1 taken from Fan et al. (2006), we see examples of quasars during the EoR, with the characteristic Gunn-Peterson trough. Fan et al. (2006) used 19 quasars to constrain the end of reionisation to $z \sim 6$. Bosman et al. (2022) updated this with a total of 67 quasars, suggesting a later reionisation of $z = 5.3$.

1.3.3 Observations of High Redshift Galaxies

The first stars and galaxies are considered to be the predominant contributor of ionising radiation, responsible for the EoR (Loeb and Barkana, 2001). Observational constraints

¹Wilkinson Microwave Anisotropy Probe

from the CMB and Gunn-Peterson trough constrain the timeline of reionisation, but not necessarily the source of ionising photons. Therefore direct observations of star forming high redshift galaxies are key to understanding the UV emissivity during the EoR (for example Oesch et al., 2015). Understanding the UV luminosity function of high redshift ($z > 6$) star forming galaxies provides one such probe. From the luminosity function, the star formation rate and UV emissivity can be inferred, allowing for estimates of the contribution of ionisation from early galaxies (Bouwens et al., 2015). Additionally, observations of the high redshift luminosity function show drops in amplitude compared to the luminosity function at lower redshifts ($z < 6$). This indicates the presence of neutral hydrogen at higher redshifts.

With the advent of the James Webb Space Telescope (JWST), the high redshift star forming galaxies, and therefore the UV luminosity function, is more accessible than ever before at higher redshifts ($z = 8 - 15$) (Donnan et al., 2022). Recent spectroscopic observations with the JWST NIRSpec instrument, have confirmed the existence of galaxies with $z > 10$ (for example; Bunker, Andrew J. et al., 2023; Wang et al., 2023). In light of these discoveries, high redshift galaxies have been used to determine the $z \geq 8$ UV luminosity function (for example; Donnan et al., 2022; Bouwens et al., 2023). These early results indicate some redshift evolution in the UV luminosity function (Donnan et al., 2022), indicating a higher than expected luminosity density. However, these results are still uncertain, and require additional high redshift galaxies to constrain the star formation rates and efficiencies at early times. Parameter inference from ongoing and future radio interferometric observations, hope to offer a complimentary constraint on the star formation efficiency.

1.4 Sources of Ionisation

Ionisation of neutral hydrogen originates mainly from two sources: UV photons which are primarily produced from high mass star formation; X-rays primarily produced from the remnants of massive star formation. UV ionisation is intrinsically linked to the star formation history, the mass of the stars, and the UV luminosity function (Mirocha,

2019). Through this relationship, the emission is parameterised through the number of ionising photons per stellar baryon N_{ion} . Barkana and Loeb (2005a) find a fiducial value of $N_{\text{ion}} = 4000$, which is expected to increase if the progenitors are population III stars (Tumlinson and Shull, 1999; Bromm et al., 2001; Schaerer, D., 2002). Population III stars are the first stars, and form in the pristine (metal free) environments of the early Universe from molecular hydrogen (H_2). As a result of these conditions, the first stars are expected to be massive with masses in the range of $30 - 300 M_{\odot}$ (Mirocha, 2019). Only a small fraction of the ionising photons will escape the galaxy and interact with the IGM; this is parameterised by the quantity known as the escape fraction f_{esc} . This is typically expected to have values of $f_{\text{esc}} \sim 10 - 20\%$ based on current constraints (Robertson et al., 2015).

X-ray emission, like UV emission, is correlated with star formation. However, X-ray emission primarily comes from the remnants of high mass star formation, particularly from stellar mass black holes. Furthermore, X-ray emission from these remnants comes from accretion which occurs most strongly in binary systems, where the compact remnant accretes material from a donor star. These sources are known as X-ray binaries and there are two classes depending on the donor star responsible for the accretion. Low mass X-ray binaries (LMXBs) accrete matter from a low mass donor star through Roche-Lobe overflow, whereas high mass X-ray binaries (HMXBs) accrete material from the wind/outflows from high mass donor stars (for a review on these source see Remillard and McClintock, 2006). Observations from nearby galaxies inform our understanding of the X-ray luminosity function, which is overwhelmingly dominated by HMXBs (Gilfanov et al., 2004; Fabbiano, 2006; Mineo et al., 2012). Additional sources of X-ray emission come from shocks as a result of supernova explosions from high mass stars. Supernovae eject vast amounts of material and energy into the inter stellar medium. The resulting relativistic electrons produced in supernovae cool via inverse Compton scattering and bremsstrahlung resulting in X-ray emission (Oh, 2001). X-ray emission has been determined to only be at most responsible for $\sim 10\%$ of ionisation (Mesinger et al., 2013). X-ray emission however plays an important role

in heating the IGM, which we will discuss in more detail in the following Chapters 2 and 7.

1.5 Thesis Motivation & Outline

The EoR encodes a wealth of astrophysical information about the first astrophysical sources, and early cosmological structures. Current efforts are focused on detecting the statistical moments of the neutral hydrogen 21cm hyperfine spin flip transition line during the EoR. In particular, measurements of the spatial statistical moments of the 21cm signal encode information about the topology of ionisation, and the distribution of neutral matter. Most focus has been on detecting the second order spatial statistical moment through the power spectrum. The power spectrum assumes that the underlying statistics of the signal are Gaussian in nature. The signal however is expected to be highly non-Gaussian due to non-linear structure formation, reionisation topology, and heating from the first astrophysical sources. Higher order statistics such as the bispectrum have been proposed to measure the non-Gaussianity of the 21cm signal. However, the bispectrum is even more challenging to estimate than the power spectrum.

Detecting the EoR with the current and future generations of radio interferometers is complicated by the presence of systematic effects, primarily from astrophysical foregrounds. Astrophysical foregrounds can be mitigated to an extent through avoidance techniques. However, the chromatic nature of radio interferometers, and the wide fields of view of low frequency interferometers means that these systematics will continue to complicate efforts to detect the signal. In particular, widefield extended sources present in the sidelobes of the primary beams of radio interferometers can contaminate statistical measurements, impeding the detection of the signal.

To understand these effects, In Chapter 2 I explore the fundamental physics of 21cm emission from neutral hydrogen, and how it propagates throughout the Universe from the earliest epochs. In Chapter 3, I outline the instrumentation and the underlying theory of radio interferometry, and introduce the Murchison Widefield Array (MWA). In Chapter 4 I develop the statistical framework for interpreting the data from radio

interferometers, and how it is applied to stochastic processes such as those that generate the 21cm signal. In Chapter 5, I then discuss the statistics from Chapter 4 for the 21cm signal in the context of astrophysics discussed in Chapter 2 and the instrumentation of Chapter 3.

This thesis builds on the existing efforts to mitigate foreground effects, and develops statistical techniques for detecting the non-Gaussianity of the neutral hydrogen during the EoR. We investigate the contamination of widefield extended source models by fitting morphological models to the brightness distribution of Galactic Plane supernova remnants (SNRs) and the bright galaxy Centaurus A. In Chapter 6 we investigate the contamination of these models to the 21cm power spectrum. We also investigate the skew spectrum of the 21cm signal, which is the integral of the bispectrum. We perform a set of seven cosmological 21cm simulations, varying the different astrophysical model parameters responsible for the ionisation topology, and heating; we discuss this project in Chapter 7. Building on these two projects, we develop a method for measuring the skew spectrum from radio interferometer observations. In Chapter 8 we use the foreground models developed in this work along with existing all-sky foreground models, to predict the instrumental skew spectrum for EoR observations with the MWA. We then perform this same method for an all-sky fiducial 21cm signal, and compare it to the expected skew spectrum. In Chapter 9 I conclude with a summary of results across the three projects, and the outlook for this work in the future.

CHAPTER 2

FUNDAMENTAL PHYSICS OF THE 21CM LINE

With the exception of observations of high redshift galaxies during the EoR, the aforementioned constraints only probe the duration, the mid point, and the neutral fraction of the EoR. Perhaps the most promising probe of the EoR is the hyperfine spin flip transition of neutral hydrogen. This transition occurs because of the interaction between the intrinsic magnetic dipole moment of the proton in the hydrogen atom and the intrinsic spin state of the electron. This causes the ground state energy level to split into two via the Zeeman effect. The excited state occurs when the electron spin is aligned with the proton spin. The excited state will eventually decay to the ground state, and in the process emit a single photon with a wavelength of 21 cm. A diagram of the process is shown in Figure 2.1.

After recombination the baryonic matter in the Universe is predominantly neutral hydrogen, which will emit 21cm radiation. As the Universe expands and the radiation propagates, it is cosmologically redshifted to longer wavelengths (or lower frequencies). The cosmological 21cm signal therefore tracks the evolution of the Universe in cosmic time, and by measuring the redshifted signal it is possible to construct a tomographic

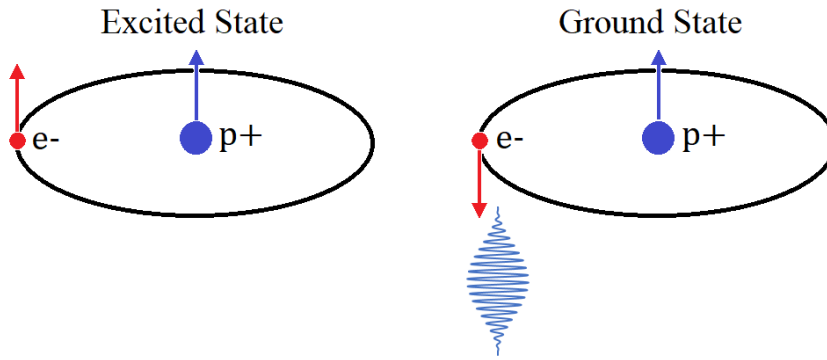


Figure 2.1: Diagram of the hyperfine transition of an electron in a spin parallel (excited state) to a spin anti-parallel (ground state) for the hydrogen atom. The transition results in the emission of a single photon.

map of the signal's evolution. Furthermore, neutral hydrogen and its 21cm emission is a sensitive calorimeter of the astrophysical processes of the first stars, galaxies, and compact objects. In this section we will outline the physics of the 21cm line, and how it couples a wealth of astrophysical and cosmological information about the first electromagnetic sources in the Universe.

2.1 Radiative Transfer

Radiative transfer describes how energy is transported in the form of electromagnetic radiation over some distance. Radiative transfer fundamentally assumes that radiation can be described as a collection of propagating rays, ignoring the wave like nature of photons. Fundamentally, radiation that propagates from some origin through some medium will interact with that medium through absorption and stimulation. The fundamental quantity of radiative transfer is the radiation intensity, which is the amount of energy received per unit area, per unit solid angle, per unit time, in the frequency range $\nu + d\nu$:

$$dE = I_\nu dA d\Omega dt d\nu. \quad (2.1)$$

I_ν incorporates the intensity from all processes, emission and absorption, and is the quantity that some observer (or instrument) would measure here on Earth. From I_ν we can define important quantities such as the emission coefficient j_ν . This is the rate of

energy emitted by some medium per unit volume, per unit time, per unit solid angle, and between the frequencies $\nu + d\nu$:

$$dI_\nu = j_\nu ds. \quad (2.2)$$

Similarly we can also describe the absorption of incident radiation by some medium with the coefficient α_ν . This differs from the definition of the emission coefficient however, where the absorption coefficient is described by the fraction of the total intensity dI_ν/I_ν which decreases per unit distance ds and has units of m^{-1} :

$$dI_\nu = -\alpha_\nu I_\nu ds. \quad (2.3)$$

From Equations 2.2 and 2.3 we can write a total expression for the equation of radiative transfer, accounting for both the emission and absorption mechanisms along the line of sight:

$$\frac{dI_\nu}{ds} = j_\nu - \alpha_\nu I_\nu. \quad (2.4)$$

From the absorption coefficient we can define the dimensionless optical depth $\tau_\nu = \int ds \alpha_\nu$, where $d\tau_\nu = \alpha_\nu ds$. Dividing through by α_ν therefore yields:

$$\frac{dI_\nu}{d\tau_\nu} = S_\nu - I_\nu, \quad (2.5)$$

where S_ν is known as the source function and is described by $S_\nu = j_\nu/\alpha_\nu$. The general solution to the radiative transfer equation is then:

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

We now consider radiative transfer for the neutral hydrogen spin flip transition line. Some proportion of hydrogen atoms will be in the spin parallel state (we denote 1) and some will be in the anti-parallel spin state (or the ground state, which we denote 0). Figure 2.1 demonstrates the situation, where a hydrogen atom in state 1 transitions to

state 0 by spontaneously emitting a 21cm photon. This emission is characterised by the Einstein coefficient $A_{10} = 2.85 \times 10^{-15} \text{ s}^{-1}$, which describes the expected lifetime of state 1. The expected decay time is $\sim 10^7$ yrs, which is short relative to the time period between recombination and the appearance of the first astrophysical sources. Therefore we can treat neutral hydrogen in the IGM as being in local thermal equilibrium. This means that the number of hydrogen atoms in some arbitrary state i , is given by the Boltzmann distribution:

$$\frac{n_i}{n} = g_i \frac{e^{-E_i/k_b T}}{Z}, \quad (2.7)$$

where E_i is the energy of the i th state, g_i is the degeneracy of states, k_b is the Boltzmann constant, T is the temperature of the ensemble, Z is the partition function, n_i is the number density of the i th state, and n is the number density of all particles in the ensemble. The proportion of neutral hydrogen atoms in state 1, versus state 0 is given by the ratio of their Boltzmann distributions:

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} \exp \left\{ -\frac{\Delta E_{21}}{k_b T_S} \right\}. \quad (2.8)$$

$\Delta E_{21} = h\nu_{21}$ is the energy difference between the spin parallel and anti-parallel state, where h is Planck's constant, and T_S is the spin temperature. We can also express this energy difference in terms of a thermodynamic temperature T_\star , which has the value 68 mK; in almost all circumstances $T_\star \ll T_S$ ¹. The ratio of the degeneracies is $g_1/g_0 = 3$ which results from the number of states for the spin triplet compared to the ground state.

In addition to the spontaneous emission of photons, hydrogen atoms can absorb radiation from some background radiation field. Background radiation can also stimulate emission if the radiation is near the resonant frequency of the atom². This results in the emission of a photon in the same direction and with the same polarisation as the inci-

¹This occurs due to the coupling of the spin temperature with various physical phenomena (CMB background, and kinetic gas temperature), which set T_S to an amplitude of $1 - 10^3$ K (see Zaroubi, 2013, for a review).

²This was playfully described to me as an incident photon tickling an atom near its resonance frequency.

dent photon that stimulated the emission. In these cases the absorption and stimulated emission are described by the Einstein coefficients $B_{01}\bar{J}$ and $B_{10}\bar{J}$ respectively, where $\bar{J} = \int_0^\infty d\nu J_\nu \phi(\nu)$ is the mean intensity of the incident background radiation field, and $\phi(\nu)$ is the line profile for the 21cm line.

For a system that is in local thermodynamic equilibrium, the number of states that change from a lower to higher state and vice versa must be equal; this is known as the principle of detailed balance:

$$n_0 B_{01} \bar{J} = n_1 A_{10} + n_1 B_{10} \bar{J}. \quad (2.9)$$

By assuming a Planck spectrum ($\bar{J} = B_\nu(T)$), we can describe the Einstein coefficients in terms of the macroscopic properties of the medium:

$$g_0 B_{01} = g_1 B_{10} \quad (2.10)$$

$$A_{10} = \frac{2h\nu^3}{c^2} B_{10}. \quad (2.11)$$

From these definitions we can recast the emission and the absorption coefficients in Equation 2.4 in terms of the Einstein coefficients, where $j_\nu = (h\nu/(4\pi))n_1 A_{10}\phi(\nu)$, and similarly for the absorption coefficient $\alpha_\nu = (h\nu/(4\pi))(n_0 B_{01} - n_1 B_{10})\phi(\nu)$:

$$\frac{dI_\nu}{ds} = \frac{\phi(\nu)h\nu}{4\pi} (n_1 A_{10} - (n_0 B_{01} - n_1 B_{10})I_\nu). \quad (2.12)$$

2.2 21cm Differential Temperature Brightness

We now turn our attention back to Equation 2.6, which for a system in thermodynamic equilibrium we can express in terms of the temperature brightness, where the Rayleigh-Jeans limit is given by $T'_b(\nu) \approx I_\nu c^2 / 2k_b \nu^2$:

$$T'_b(\nu) = T_S(1 - e^{-\tau_\nu}) + T'_R(\nu)e^{-\tau_\nu}. \quad (2.13)$$

$T'_R(\nu)$ is the brightness temperature of the background radiation field which is incident on the cloud along the path of the ray. Equation 2.13 describes the radiative transfer along a line of sight through a cloud of uniform excitation temperature T_S . The optical depth τ_ν for the 21cm line is expected to be small such that $e^{-\tau_\nu} \sim 1 - \tau_\nu$:

$$T'_b(\nu) = (T_S - T'_R(\nu)) \tau_\nu + T'_R(\nu). \quad (2.14)$$

What we actually measure is the differential temperature brightness, therefore we subtract $T'_R(\nu)$ from both sides of Equation 2.14. The differential temperature brightness is therefore $\delta T'_b(\nu) = T'_b(\nu) - T'_R(\nu)$. As a result of the cosmological distances these rays travel, the apparent differential temperature brightness observed here on Earth would be $\delta T_b(\nu) = \delta T'_b(\nu_0)/(1+z)$, where the observed frequency is $\nu = \nu_0/(1+z)$. Finally, we note that the background radiation field is in almost all cases the CMB, denoted by T_{CMB} . Since the CMB is almost a perfect black body, the temperature is described by $T_{\text{CMB}} = 2.73(1+z)$ K. Substituting these changes to Equation 2.14 yields:

$$\delta T_b(z) = \frac{T_S - T_{\text{CMB}}(z)}{1+z} \tau_\nu. \quad (2.15)$$

The optical depth encodes the physics of the 21cm temperature brightness, and is of particular interest. Here we derive the 21cm optical depth in terms of the absorption coefficient:

$$\begin{aligned} \tau_\nu &= \int ds \alpha_\nu \\ &= \int ds \frac{h\nu}{4\pi} (n_0 B_{01} - n_1 B_{10}) \phi(\nu). \end{aligned} \quad (2.16)$$

Substituting for the Einstein coefficient B_{10} :

$$\tau_\nu = \int ds \frac{h\nu}{4\pi} n_0 B_{01} \frac{g_0}{g_1} \left(\frac{g_1}{g_0} - \frac{n_1}{n_0} \right) \phi(\nu). \quad (2.17)$$

Next we substitute in Equation 2.8 and A_{10} :

$$\tau_\nu = \int ds \frac{3h c^2 n_0 A_{10}}{8\pi \nu k_b T_S} \phi(\nu). \quad (2.18)$$

We assume that most of the terms in Equation 2.18 are constants with respect to the proper length element ds . We then perform the integral with respect to n_0 , where $\int ds n_0 = N_{\text{HI}}/4$, the column density of HI. The 1/4 term represents the fraction of HI atoms that are in the hyperfine singlet state.

The line profile $\phi(\nu)$ includes effects from many different sources, such as thermal effects and bulk motion. The velocity broadening of the line profile over a line of sight element s is given by $\Delta v \sim s dv_{\parallel}/dr_{\parallel}$, where the term $dv_{\parallel}/dr_{\parallel}$ is the total velocity gradient profile along the line of sight. For the line of sight velocity gradient the peculiar velocity and the Hubble expansion are the dominant terms. The line profile is then given by $\phi(\nu) \sim c/(\Delta v \nu)$. Furthermore, we can express the column density in terms of the same line of sight segment s and the neutral fraction of hydrogen atoms x_{HI} . Here $N_{\text{HI}} = x_{\text{HI}} n_H(z) s$, where $n_H(z)$ is the redshift dependent number density of hydrogen atoms, which can be expressed as $n_H(z) = n_H/(1+z)^3$. Performing these substitutions we get the final expression for the optical depth:

$$\tau_\nu = \frac{3}{32\pi} \frac{hc^3}{\nu_0^2} \frac{A_{10}}{k_b T_S} \frac{x_{\text{HI}} n_H}{(1+z) dv_{\parallel}/dr_{\parallel}}. \quad (2.19)$$

We also substitute for the emitted frequency ν_0 , which removes two factors of $(1+z)$ from the denominator. This derivation follows that of Furlanetto (2016), which was first shown by Gunn and Peterson (1965) for Ly α , and made possible by Field (1958).

The optical depth is often recast in terms of the Hubble constant $H(z)$ at z , and the fractional overdensity $(1+\delta) = \rho/\bar{\rho}$, where ρ is the matter density, and $\bar{\rho}$ is the mean density (for example Furlanetto, 2016):

$$\tau_\nu \approx 9.2 \times 10^{-3} (1+z)^{3/2} (1+\delta) \frac{x_{\text{HI}}}{T_S} \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right]. \quad (2.20)$$

We can substitute Equation 2.20 into the differential temperature brightness in Equation 2.15:

$$\delta T_b(z) \approx 9x_{\text{HI}}(1 + \delta)(1 + z)^{1/2} \left[1 - \frac{T_{\text{CMB}(z)}}{T_S} \right] \left[\frac{H(z)/(1 + z)}{dv_{\parallel}/dr_{\parallel}} \right] \text{ mK.} \quad (2.21)$$

Equation 2.21 shows the final expression for the differential temperature brightness (also known as the temperature contrast). Immediately we see two important regimes. The first occurs when the spin temperature is saturated $T_S \gg T_{\text{CMB}}$; in this case $\delta T_b(z)$ is effectively independent of the spin temperature. In this scenario, the density contrast and the neutral fraction become the primary driver for δT_b signal. In the other regime, $T_{\text{CMB}} \gg T_S$ which indicates that the signal is in absorption relative to the background radiation. It is clear from Equation 2.21 that T_S is key in determining the observability of the 21cm signal, whether in absorption or emission. In the following section we discuss the different mechanisms that set the spin temperature.

2.3 Spin Temperature

The spin temperature is set by three different processes; particle collisions, interactions with the background CMB photons, and the scattering of UV photons, primarily $\text{Ly}\alpha$ photons through the Wouthuysen-Field effect (Wouthuysen, 1952; Field, 1958). Each process establishes thermodynamic equilibrium with the neutral hydrogen, and is therefore described by their own temperatures through the principle of detailed balance:

$$n_1(c_{10} + P_{10} + A_{10} + B_{10}I_{\text{CMB}}) = n_0(c_{01} + P_{01} + A_{01} + B_{01}I_{\text{CMB}}), \quad (2.22)$$

P_{10} and C_{10} are the Einstein coefficients for the $\text{Ly}\alpha$ and collisional processes. Each of the different processes can be described in terms of their temperature:

$$\frac{P_{01}}{P_{10}} = 3 \left(1 - \frac{T_{\star}}{T_{\alpha}} \right) \quad (2.23)$$

$$\frac{C_{01}}{C_{10}} = 3 \left(1 - \frac{T_{\star}}{T_K} \right). \quad (2.24)$$

Here T_α is the colour temperature for the Ly α scattering, and T_K is the kinetic gas temperature for neutral hydrogen. The ratios of the Einstein coefficients are determined by taking the ratio of the Boltzmann distributions for each process. We can describe T_S as a weighted average of all of these processes (Field, 1958, 1959a):

$$T_S^{-1} = \frac{T_{\text{CMB}}^{-1} + x_\alpha T_\alpha^{-1} + x_c T_K^{-1}}{1 + x_\alpha + x_c}, \quad (2.25)$$

where x_α and x_c are the coupling weights for these two processes.

As the Universe expands, both the hydrogen gas and the CMB radiation temperatures reduce. The neutral hydrogen however cools faster due to the adiabatic expansion of the Universe. As a result, when the background CMB radiation density is high enough, absorption and stimulated emission from the CMB radiation couples to the spin temperature via scattering through the hyperfine levels. This drives the spin temperature towards the CMB temperature ($T_S \rightarrow T_{\text{CMB}}$) (Furlanetto, 2019). We can see these effects in Figure 2.2 which shows the redshift evolution of the spin temperature (T_S), the CMB temperature (given as T_γ in this example), and the kinetic gas temperature (T_K).

The second process driving T_S is collisions with particles within the IGM; these interactions include electron-hydrogen, proton-hydrogen, and hydrogen-hydrogen interactions (Zygelman, 2005; Furlanetto and Furlanetto, 2006). These collisions break the coupling of T_S to the CMB and drive the spin temperature back towards the kinetic gas temperature T_K . Calculations of the collision excitation and de-excitation rates are determined by the quantum mechanical cross sections for the different particle interactions. Zygelman (2005) and Furlanetto and Furlanetto (2006) found that although the atomic cross sections for hydrogen-hydrogen interactions are small (relative to the other interactions), these collisions dominate in a relatively unperturbed IGM when the neutral fraction is high. Free electrons become important when the gas is partially ionised, which we discuss in the following section. Collisional coupling in this regime is important when the IGM density is high enough, which occurs at $z < 70$ (see Figure 2.2). As the IGM expands, and the density drops, $T_S \rightarrow T_{\text{CMB}}$ by $z \sim 7$, rendering




Figure 2.2: Redshift evolution of the spin temperature (solid red line), the kinetic temperature (dotted blue line), and the background radiation temperature (dashed green line). In this Figure, the green line represents the CMB temperature. Figure retrieved from Mesinger et al. (2011).

the differential 21cm signal effectively invisible (Furlanetto, 2019).

After collisional coupling becomes insignificant, Ly α coupling through the WF-effect breaks the spin/CMB coupling through the mixing of the hyperfine ground states (Wouthuysen, 1952; Field, 1958). Figure 2.3 shows a diagram of mixing mechanism, where due to the dipole selection rule only transitions that conserve total angular momentum are allowed (Pritchard and Furlanetto, 2006). As a result electrons can be promoted to either of the $n = 2$ states, which can result in the change of the spin of the promoted electron. These electrons can then decay back to the $n = 1$ ground state with effectively randomised (mixed) spin states. In essence, Ly α absorption acts as an intermediary, by which the hyperfine spin states of neutral hydrogen can be mixed, breaking the CMB coupling in the process.

The significance of the WF-effect depends on the rate at which Ly α photons scatter per atom through a gas of neutral hydrogen. This scattering rate has been shown to be $\sim 10^5$ before the Ly α photon redshifts out of resonance (Field, 1959b). Other Lyman




Figure 2.3: Shows how the Wouthuysen-Field effect mixes the spin states through the hyperfine splitting of the $1S$ and $2P$ states. The solid lines indicate transitions that participate in mixing the hyperfine ground state levels. The dashed lines indicate allowed transitions; these do not contribute to the ground state mixing. This figure was retrieved from Pritchard and Furlanetto (2006).

series photons can also participate in this mode mixing, however, the $\text{Ly}\alpha$ scattering has been shown to be the dominant mechanism due to the large number of scatterings, suppressing the other transitions (Pritchard and Furlanetto, 2006). Notably the colour temperature for $\text{Ly}\alpha$ scattering can be shown to be $T_\alpha \approx T_K$. This is a direct result of the large number of scatterings which force the $\text{Ly}\alpha$ profile to become a blackbody with temperature T_K near the line centre. This profile does however redden as the $\text{Ly}\alpha$ photons lose energy through recoil via their large number of scatterings (Field, 1958).

2.4 Heating of the Intergalactic Medium

Both collisional coupling and the WF-effect drive the spin temperature to the kinetic temperature. Therefore, processes that heat the IGM are key to understanding the 21cm brightness temperature. There are three primary heating mechanisms that set the kinetic temperature of the hydrogen gas: $\text{Ly}\alpha$ background emission; CMB radiation background; the X-ray background.

The $\text{Ly}\alpha$ background contributes to the heating of the neutral hydrogen through recoil from scattering. However Chen and Miralda-Escudé (2004) and Furlanetto and

Pritchard (2006) show that this mechanism in isolation is unimportant, impacting the kinetic temperature on the scale of $\sim 1\%$. This occurs because the energy gained by recoil from Ly α photons is balanced by the scattering diffusivity in equilibrium. However, the heating effect of Ly α can be enhanced when scattering from CMB photons is included. Venumadhav et al. (2018) showed that when the two effects are considered in tandem, CMB scattering amplifies the Ly α heating, resulting in a modification of the kinetic temperature by 9 – 15%. CMB photons scatter through the hyperfine levels, which results in a net heating effect above the expected adiabatic temperature. This extra energy is transferred to the Ly α photons as they mix the hyperfine levels through the WF-effect. This then results in heat being deposited through recoil, resulting in a net heating effect due to the large number of scatterings Ly α undergo before redshifting out of resonance.

The most important heating mechanism comes from X-ray emissions from both galaxies and quasars (Madau et al., 1997). The resulting X-ray photons have mean free paths that are directly proportional to their energy with the highest energy X-rays having mean free paths larger than the Hubble length (Oh, 2001; McQuinn, 2012). Some fraction of these high energy X-rays will deposit their energy uniformly as heat in the IGM. Lower energy X-rays on the other hand will deposit their energy in more localised regions in the IGM, converting much of this to heat.

X-rays deposit heat into the IGM by photoionising helium or hydrogen atoms. The resulting free electrons are called primaries. The remaining photon energy not required to liberate these electrons is converted into kinetic energy. These primary electrons will interact with the IGM through three main processes. The first process is collisional ionisations which generate secondary electrons that scatter throughout the IGM redistributing the kinetic energy. The second process is collisional excitations with either HI or HeI, where HI will emit Ly α photons contributing to the Ly α background; HeI on the other hand will emit ionising photons that will go on to produce more secondary electrons. The final interaction is collision with other free electrons in the IGM through Coulomb interactions, distributing the kinetic energy. The photoelectrons eventually

cool below the Ly α threshold, and join the other electrons in thermal equilibrium. These electrons then tend to thermal equilibrium with the rest of the neutral matter through elastic collisions on times scales less than the Hubble time (Furlanetto, 2019).

2.5 Summary

As we have seen in the previous sections the neutral hydrogen 21cm encodes a lot of information about the different astrophysical processes, at different cosmological periods in the Universe. It is a potential sensitive calorimeter of the Universe during the cosmic dark ages, and at the first light from the first stars, galaxies, and compact objects. In Chapter 5 we revisit the 21cm line, how we can measure it with radio interferometers, and what astrophysical information we are sensitive to.

CHAPTER 3

RADIO INTERFEROMETRY & SIMULATION

Radio interferometry is a key pillar of modern astronomy, with applications varying from investigating the neutral hydrogen in the early cosmos, to the imaging of the event horizons of supermassive black holes. The first interferometers were optical instruments. Michelson (1920); Michelson and Pease (1921) used one of the first optical interferometers to measure the angular sizes of nearby bright stars. This simple set up was constructed from two optical telescopes separated by some distance. Incident radiation from a source in the far-field was offset of by some angle θ . The two waves were then combined, and the amplitudes of the fringes measured to determine the *visibility*. The fringes were separated by $\Delta\theta \approx \lambda/D$, where λ is the observing wavelength, and D is the distance between the two optical detector. If the angular size of the source was smaller than the offset angle θ , then ratio of the fringe amplitudes was one. For sources with angular sizes larger than this separation the visibility was less than one. Early radio interferometers built on these principles taking inspiration from optical interferometers.

Ryle and Vonberg (1946) constructed one of the first radio interferometers to investigate radio emission from the sun. This early radio interferometer was a dipole antenna array that was East-West facing oriented, and pointed at the meridian. This array scanned the sky in right ascension as the Earth rotated. McCready et al. (1947) was one of the first to realise that the interferometer was measuring a Fourier component of the brightness distribution. McCready et al. (1947) provided one of the first descriptions of Fourier synthesis and how an image could be produced from visibility measurements from multiple baselines. A big development came in 1952 with Ryle (1952), who introduced the phase switching into interferometers, which removed unwanted components in the signal combination. This led to the development of the first correlator, establishing the modern conventions for radio interferometry.

In this chapter we will establish a simple derivation of the Fourier relationship between the brightness intensity, and the fundamental measurement of radio interferometers: the visibility. The visibility effectively measures the correlation of the electric field as measured by two elements separated by some distance. From this measurement under certain conditions we can reconstruct the brightness distribution of the sky, as well as probe the statistics of the 21cm signal, as we will see in Chapter 5.

3.1 Spatial Coherence Function

In this section we follow the derivation of the spatial coherence function from Clark (1999). Fundamentally, radio interferometers relate the electric field $\mathbf{E}_v(\mathbf{R})$ emitted at some point \mathbf{R} to the electric field $\mathbf{E}_v(\mathbf{r})$ measured or induced at some detector (antenna) located at \mathbf{r} . The electric field $\mathbf{E}_v(\mathbf{r})$ is the integral over all electric fields with some separation \mathbf{R} :

$$\mathbf{E}_v(\mathbf{r}) = \int_{\mathbb{R}^3} P_v(\mathbf{R}, \mathbf{r}) \mathbf{E}_v(\mathbf{R}) d^3\mathbf{x}. \quad (3.1)$$

Performed over all space, $P_v(\mathbf{R})$ is the propagator, and describes how the field emitted at \mathbf{R} propagates to position \mathbf{r} . Electric fields are vector fields, and can be decomposed into linearly independent components. For now, we assume that the field measured at




Figure 3.1: Coordinate system for the electric field measured at $\mathbf{E}(\mathbf{r})$ as emitted from some source at \mathbf{R} from the origin. The source is indicated by the blue ellipse, and is emitting incoherent and isotropic radiation in the far-field relative to the point \mathbf{r} .

\mathbf{r} can be treated as a scalar, which simplifies Equation 3.1 to a more manageable form. We come back to polarimetry in Section 3.1.4 of this chapter. For now we build our understanding from simple scalar fields.

The integral in Equation 3.1 integrates over all space. If $|\mathbf{R}| \gg |\mathbf{r}|$, we can assume that all radiation is emitted on a spherical surface at distance \mathbf{R} from the origin; no radiation is emitted within the sphere. This is the far-field approximation, which allows for the assumption that radiation emitted at some point \mathbf{R} is effectively a plane wave when received at position \mathbf{r} , with the plane wave described by $e^{2\pi i \nu |\mathbf{R} - \mathbf{r}| / c}$. Figure 3.3 demonstrates the setup in this diagram. The distribution of the scalar electric field emitted at the spherical surface is then expressed as $\mathcal{E}(\mathbf{R})$. Taking these assumptions into account, $E_\nu(\mathbf{r})$ becomes:

$$E_\nu(\mathbf{r}) = \int_{\mathbb{R}^2} \mathcal{E}_\nu(\mathbf{R}) \frac{e^{2\pi i \nu |\mathbf{R} - \mathbf{r}| / c}}{|\mathbf{R} - \mathbf{r}|} dS. \quad (3.2)$$

Since emission is confined to the surface of a sphere, the integral over all space




Figure 3.2: The coordinate system for a single baseline $\mathbf{r}_1 - \mathbf{r}_2$. The blue ellipse is some source emitting isotropic radiation in the far-field.

reduces to one over the surface of the sphere, with the surface area element dS . In the literature this equation is referred to as the general form of the quasi-monochromatic component of the electric field at the frequency (ν) from all cosmic electromagnetic radiators.

Up to this point we have described the electric field that can be measured at some point \mathbf{r} , and the contributions from all celestial sources. Fundamentally radio interferometers measure a spatial correlation of cosmic signals. What we are interested in is the correlation between the electric fields measured at points \mathbf{r}_1 and \mathbf{r}_2 :

$$\mathcal{V}_\nu(\mathbf{r}_1, \mathbf{r}_2) \equiv \langle \mathbf{E}_\nu(\mathbf{r}_1) \mathbf{E}_\nu^*(\mathbf{r}_2) \rangle, \quad (3.3)$$

The superscript $*$ indicates the complex conjugate, since the electric field of the plane wave is a fundamentally complex valued function. Substituting Equation 3.2 into

Equation 3.3 yields:

$$\mathcal{V}_v(\mathbf{r}_1, \mathbf{r}_2) = \left\langle \int_{\mathbb{R}^2} \int_{\mathbb{R}^2} \mathcal{E}_v(\mathbf{R}_1) \mathcal{E}_v^*(\mathbf{R}_2) \frac{e^{2\pi i v |\mathbf{R}_1 - \mathbf{r}_1|/c}}{|\mathbf{R}_1 - \mathbf{r}_1|} \frac{e^{2\pi i v |\mathbf{R}_2 - \mathbf{r}_2|/c}}{|\mathbf{R}_2 - \mathbf{r}_2|} dS_1 dS_2 \right\rangle. \quad (3.4)$$

In this regime both $\mathcal{E}(\mathbf{R}_1)$ and $\mathcal{E}(\mathbf{R}_2)$ are confined to the celestial sphere, thus as defined from the origin $|\mathbf{R}_1| = |\mathbf{R}_2|$. Notably, in the far-field approximation the exponential terms in Equation 3.5 are effectively constants since $\mathbf{R} \gg \mathbf{r}$, therefore $|\mathbf{R} - \mathbf{r}|$ is relatively unchanged around the celestial sphere. Notably, we can exchange the expectation and integral operators in Equation 3.5, therefore, the expectation occurs only over the product $\langle \mathcal{E}_v(\mathbf{R}_1) \mathcal{E}_v^*(\mathbf{R}_2) \rangle$. Again, since $|\mathbf{R}|$ is sufficiently large, we can assume that points across the celestial sphere are not correlated: that is, they are not spatially coherent. From a causality perspective, since sources are separated by a distance $\delta\theta|\mathbf{R}|$, the travel time of light to causally link events is large enough that we can consider them incoherent/independent. The exception perhaps being in very local regions about a source/event¹. Therefore when $\mathbf{R}_1 \neq \mathbf{R}_2$ we have $\langle \mathcal{E}_v(\mathbf{R}_1) \mathcal{E}_v^*(\mathbf{R}_2) \rangle = 0$. This is effectively a delta function. Integrating over this yields

$$\mathcal{V}_v(\mathbf{r}_1, \mathbf{r}_2) = \int_{\mathbb{R}^2} \langle \mathcal{E}_v(\mathbf{R}) \mathcal{E}_v^*(\mathbf{R}) \rangle |\mathbf{R}|^2 \frac{e^{2\pi i v |\mathbf{R} - \mathbf{r}_1|/c}}{|\mathbf{R} - \mathbf{r}_1|} \frac{e^{2\pi i v |\mathbf{R} - \mathbf{r}_2|/c}}{|\mathbf{R} - \mathbf{r}_2|} dS. \quad (3.5)$$

We can further simplify $\mathcal{V}_v(\mathbf{r}_1, \mathbf{r}_2)$, by substituting $dS = |\mathbf{R}|^2 d\Omega$, and $I(\mathbf{s}) = \langle \mathcal{E}_v(\mathbf{R}) \mathcal{E}_v^*(\mathbf{R}) \rangle |\mathbf{R}|^2$, which is the brightness intensity distribution. Here $\mathbf{s} = \mathbf{R}/|\mathbf{R}|$ is the unit vector to the radiating point on the celestial sphere. Returning to the far-field approximation, we can write $|\mathbf{R} - \mathbf{r}| \approx (\mathbf{R} - \mathbf{r}) \cdot \mathbf{s}$ (the dot product divided by the magnitude of the radius)². Thus $|\mathbf{R} - \mathbf{r}_1| - |\mathbf{R} - \mathbf{r}_2| \approx -\mathbf{s} \cdot (\mathbf{r}_1 - \mathbf{r}_2)$, resulting in:

$$\mathcal{V}_v(\mathbf{r}_1, \mathbf{r}_2) \approx \int_{\mathbb{R}^2} I_v(\mathbf{s}) e^{-2\pi i v \mathbf{s} \cdot (\mathbf{r}_1 - \mathbf{r}_2)/c} d\Omega. \quad (3.6)$$

Equation 3.6 is the spatial coherence function of $E_v(\mathbf{r})$, and only depends on the

¹From an instrument perspective the scale at which a source or event is coherent is likely smaller than the instrument resolution.

²We can make this assumption because the angle between $(\mathbf{R} - \mathbf{r})$ and \mathbf{s} in the dot product is effectively infinitesimal, thus $\lim_{\delta\theta \rightarrow 0} \cos \delta\theta = 1$




Figure 3.3: The coordinate system for a single baseline. Taken from Thompson et al. (2017)

separation of points measured by the vector $(\mathbf{r}_1 - \mathbf{r}_2)$. This is the visibility and is the fundamental quantity measured by interferometers, where we see that the correlation of the induced electric fields at points \mathbf{r}_1 and \mathbf{r}_2 is related to the brightness intensity on the celestial sphere. This visibility has units of $\text{W m}^{-2} \text{Hz}^{-1}$, which is often in radio astronomy, and in this work, converted into a new unit called the *Jansky* $\text{Jy} = 10^{-26} \text{W m}^{-2} \text{Hz}^{-1}$.

3.1.1 Synthesis Imaging

Under certain conditions/constraints we can treat Equation 3.6 as the Fourier transform of the brightness distribution $I_\nu(\mathbf{s})$. In this section we look at an example, called planar analysis, where the brightness distribution is projected onto a plane from the celestial sphere. This configuration and its relation to the interferometer baseline can be seen in Figure 3.3. Next we transform to the planar coordinate system, where the unit vector

\mathbf{s} can be decomposed into the component direction cosines (l, m, n) . Additionally, we transform the baseline vector $\lambda(\mathbf{r}_1 - \mathbf{r}_2) = (u, v, w)$ where $\lambda = \nu/c$ is the wavelength, and the baseline vector is in units of wavelengths. Expressing Equation 3.6 in terms of the coordinate transformation yields

$$\mathcal{V}_\nu(u, v, w) = \iint \frac{I_\nu(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-2\pi i(ul + vm + w(n-1))} dl dm, \quad (3.7)$$

where the Jacobian that transforms to the coordinate system from angular coordinates to the direction cosines (l, m) is $d\Omega = dl dm / \sqrt{1 - l^2 - m^2}$. If the interferometer elements are co-planar, that is $w = 0$, Equation 3.7 becomes:

$$\mathcal{V}_\nu(u, v) \approx \iint \frac{I_\nu(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-2\pi i(ul + vm)} dl dm. \quad (3.8)$$

This approximation works well for instruments that have small fields of view, where $(l^2 + m^2 \ll 1)$, thus:

$$\mathcal{V}_\nu(u, v) \approx \iint I_\nu(l, m) e^{-2\pi i(ul + vm)} dl dm. \quad (3.9)$$

If these conditions are satisfied then the visibility function is related to the sky intensity distribution by a 2D Fourier transform with respect to (l, m) . In this work we deal with an aperture array, which has both a large field of view, and is not co-planar. We come back to this case later when we discuss how to simulate the visibilities measured by an interferometer.

3.1.2 Sampling

In practice we can only measure $\mathcal{V}_\nu(u, v)$ at discrete points of (u, v) for each baseline. We can describe this process through the sampling function $S(u, v)$, which is zero when there is no measurement. Performing the Fourier inversion with the sampling function

included we get the following expression:

$$I_v^D(l, m) = \iint \mathcal{V}_v(u, v) S(u, v) e^{2\pi i(ul+vm)} du dv. \quad (3.10)$$

$I_v^D(l, m)$ is often referred to as the dirty image, because it is the real sky intensity distribution convolved with the point spread function (PSF) $I_{\text{psf}}(l, m)$:

$$I_v^D(l, m) \equiv I_v(l, m) * I_{\text{psf}}(l, m), \quad (3.11)$$

The sign $*$ indicates convolution. The dirty image demonstrates our ignorance of the full sky brightness distribution. The PSF is often removed through an iterative process called CLEANing when constructing science quality images from radio interferometry data (Högbom, 1974). This iterative process reconstructs the "true sky" image by assuming a restoring kernel, that is typically a Gaussian PSF. There are many different types of deconvolution algorithms available for synthesis imaging; CLEANing and its many derivatives are the most common. In this work we do not invert the visibilities to get the image and as such we limit our discussion of the different deconvolution methods. We instead refer the interested reader to the literature (see Offringa et al., 2014; Offringa and Smirnov, 2017, for examples.).

3.1.3 Primary Beam Response

In addition to sampling, radio interferometers are constructed from receiving elements which are not uniformly sensitive to all emission in the sky. A given antenna element in a baseline will have a response pattern that is projected onto the sky, which is called the primary beam pattern, and can be projected onto the same plane as the sky brightness distribution. We denote the primary beam function as $\mathcal{B}(l, m)$. What the instrument actually measures is the apparent sky brightness $I_{A_v}(l, m) = \mathcal{B}(l, m)I_v(l, m)$. Inserting this definition into Equation 3.7:

$$\mathcal{V}_v(u, v, w) = \iint \frac{\mathcal{B}(l, m)I_v(l, m)}{\sqrt{1-l^2-m^2}} e^{-2\pi i(ul+vm+w(n-1))} dl dm. \quad (3.12)$$

In this case we can subsume the $1/\sqrt{1-l^2-m^2}$ into the primary beam. The primary beam function typically drops off rapidly as a function of distance from the phase centre (this is where the instrument is pointing), heading towards zero, and so $\mathcal{B}(l, m)/\sqrt{1-l^2-m^2} \approx \mathcal{B}(l, m)$ is a fine assumption. For most arrays, it is the primary beam that allows the planar assumption, because the instrument is only sensitive to emission from a small region in the sky. This is not the case with the instrumental simulations we perform in this work. We discuss this more in the following sections.

Typically in radio interferometric imaging the primary beam is assumed to be the same for each element³, and is usually known analytically to within a few percent. In these cases the primary beam can be divided out after the interferometric image is created through Fourier inversion. This does however impact the noise further from the centre of the primary beam through the division of small number, thus there is usually a field of view (FoV), which is defined by the primary beam.

3.1.4 Polarimetry and the Measurement Equation

At the start of this section we considered only a scalar monochromatic electric field measured by a two element interferometer. In reality, the electric field of a plane wave is a vector quantity, which is describe by the *Jones* vector (Jones, 1941):

$$\mathbf{E}_v = \begin{pmatrix} E_X \\ E_Y \end{pmatrix}. \quad (3.13)$$

Plane waves oscillate transverse to their direction of motion; for the electric field component of a plane wave $\mathbf{E}_v(t)$ this can be described by two linearly independent components E_X and E_Y . The geometric orientation of these components, and how they change with time, is called the polarisation. Each individual plane wave has its own polarisation, however, radiation emitted from an incoherent source emits radiation with random polarisations. Some astrophysical processes do emit polarised radiation,

³With the exception of very long baseline interferometry (VLBI), which constructs an array out of different radio telescopes across the world.

one example being synchrotron radiation, emitted by relativistic electrons gyrating in a magnetic field. The resulting emission is partially polarised due the orientation of the electron relative to the magnetic field of the source. Other sources of polarised emission are pulsars and masers. Understanding the polarisation of the emission provides insight into the in situ magnetic field of these sources.

The polarisation state of a plane wave is typically described by the *Stokes* parameters, I , Q , U , and V (McMaster, 1954). I is the total intensity of the electromagnetic radiation, Q and U are the linear polarisation components, and V is the circular polarisation. The different parameters are related to the amplitudes of the linear plane wave components through the following relations:

$$I = |E_X|^2 + |E_Y|^2 \quad (3.14)$$

$$Q = |E_X|^2 - |E_Y|^2 \quad (3.15)$$

$$U = 2\Re\{E_X E_Y^\dagger\} \quad (3.16)$$

$$V = -2\Im\{E_X E_Y^\dagger\}. \quad (3.17)$$

The \dagger symbol indicates the complex conjugate transpose, and the \Re and \Im are the real and imaginary components of some complex quantity. The Stokes parameters in Equation 3.14 can equally be defined in a circular polarisation basis with independent right-hand (E_R) and left-hand (E_L) circular polarisation components. We choose the linear polarisation basis here for consistency reasons that we explore in more detail in Section 3.3.

3.1.4.1 Measurement Equation

The process outlined in Section 3.1 describes the measurement for the amplitude of one of the linear components. Now we describe the full measurement of the individual polarisation visibilities, which the literature calls the radio interferometer measurement equation (RIME) (Thompson et al., 2017; Taylor et al., 1999). The concepts of polarisation and visibilities have been around since the early parts of radio

interferometry, however the RIME was not formulated in this description until a series of papers by (Hamaker et al., 1996; Sault et al., 1996; Hamaker and Bregman, 1996; Hamaker, 2000). These works describe the visibility measurements as a matrix or vector using Jones and Mueller matrix representations. Smirnov (2011) revised the formulation of the RIME, and it is this work that we use to derive the RIME in this section. These works additionally focus on the calibration of radio interferometers, and some other practical considerations associated with radio interferometric data. In this work we neglect these discussions, since we work purely with simulated visibilities, often constructed to provide the ideal case. These choices are discussed further in Section 3.3 where we lay out the methods we used to simulate visibilities.

Now consider an antenna which is designed to measure the North-South (X) and East-West (Y) linear polarisation components of incoming radiation⁴. The electric field \mathbf{E}'_v received by the antenna is related to the electric field emitted by the source on the celestial sphere by:

$$\mathbf{E}'_v = \mathbf{J}\mathbf{E}_v. \quad (3.18)$$

\mathbf{J} is the Jones matrix (Jones, 1941), which is a 2×2 complex matrix that describes the propagation effects of the plane wave from emission to reception; this is analogous to the propagator function in Equation 3.1. From transmission to reception, there is any number of propagation effects that can transform the signal, from the Faraday rotation of the ionosphere, to reception by the instrument. Each of these effects can be described by a Jones matrix \mathbf{J}_i , and can be combined as the product into a single total Jones matrix $\mathbf{J} = \mathbf{J}_n \mathbf{J}_{n-1} \cdots \mathbf{J}_1$. The measured polarisation components E'_X and E'_Y of the instrument do not necessary match the polarisation of incident radiation. From these components the Stokes parameters can still be derived. We can describe the voltage induced at the antenna from the emitted radiation as follows:

$$\mathbf{v} = \begin{pmatrix} v_X \\ v_Y \end{pmatrix} = \mathbf{J}\mathbf{E}_v. \quad (3.19)$$

⁴The choice of direction for X and Y is arbitrary, the MWA defines X as the East-West direction, however IAU conventions define it to be North-South. This results in a sign change in U Stokes parameter.

Now consider a two element interferometer containing antennas 1 and 2⁵. The voltages induced in antennas 1 and 2 are given by \mathbf{v}_1 and \mathbf{v}_2 . Like in Equation 3.3 we correlate the two signals together:

$$\mathcal{V}_{12} = 2\langle \mathbf{v}_1 \mathbf{v}_2^\dagger \rangle. \quad (3.20)$$

The dagger corresponds to the conjugate transpose⁶. Equation 3.22 describes the voltage correlations for each linear combination of instrument polarisations $\langle \mathbf{v}_{1X} \mathbf{v}_{2X}^\dagger \rangle$, $\langle \mathbf{v}_{1X} \mathbf{v}_{2Y}^\dagger \rangle$, $\langle \mathbf{v}_{1Y} \mathbf{v}_{2X}^\dagger \rangle$, and $\langle \mathbf{v}_{1Y} \mathbf{v}_{2Y}^\dagger \rangle$. We can relate the correlation functions to the Stokes parameters via

$$2\langle \mathbf{E}_1 \mathbf{E}_2^\dagger \rangle = \mathbf{B} = \begin{pmatrix} I + Q & U + iV \\ U - iV & I - Q \end{pmatrix}. \quad (3.21)$$

\mathbf{B} is the brightness matrix, and has the same units as intensity ($\text{W m}^{-2} \text{Hz}^{-1} \text{sr}^{-1}$). Substituting this into Equation 3.20:

$$\mathcal{V}_{12} = \mathbf{J}_1 \mathbf{B} \mathbf{J}_2^\dagger. \quad (3.22)$$

Equation 3.22 describes the correlations of the voltages from an incident plane wave received from some arbitrary position on the celestial sphere. Integrating this expression as a function of solid angle across the whole sky gives the full visibility description of the RIME:

$$\mathcal{V}_{12} = \int_{4\pi} \mathbf{J}_1(\mathbf{s}) \mathbf{B}(\mathbf{s}) \mathbf{J}_2^\dagger(\mathbf{s}) d\Omega. \quad (3.23)$$

In Equation 3.12 \mathbf{s} is the unit vector pointing to some source location, and the subscript 4π indicates the integral of the solid angle across the whole sky. The Jones matrix in Equation 3.23 contains a phase term (Smirnov, 2011), which is equivalent to the exponential from Equation 3.12. The integral is performed for each independent correlation, as a function of solid angle across the sky, and is the same operation performed in Equation 3.12 for the scalar field example.

⁵All $N > 2$ element radio interferometers are an ensemble of two element interferometers.

⁶The factor of two is a convention chosen by (Smirnov, 2011), it is valid to redefine the system without this.

3.2 Murchison Widefield Array

In this chapter so far we have provided a general description of how a basic two element radio interferometer works. In this section we discuss the Murchison Widefield Array (MWA), which is the primary science instrument used in this work. One of the main science goals for the MWA is to detect the 21cm signal during the EoR (Bowman et al., 2013; Wayth et al., 2018). The MWA is a low frequency aperture array, which is located in the Murchison Radio astronomy observatory (MRO; traditional name Inyarrimanha Ilgari Bundara), which is a special radio quiet zone (Bowman et al., 2013; Ord et al., 2010). The Phase I configuration of the MWA array contained 128 tile stations with a maximum baseline length of approximately 3 km (required for the arcminute level resolution). Following the Phase II upgrade, the array comprises a total of 256 tile stations, which can be configured into two sub array configurations – the extended and the compact – both containing 128 tiles. Each tile in the array is comprised of 16 X (North-South) and Y (East-West) linear polarisation dipole antennas in a 4×4 configuration (Wayth et al., 2018; Bowman et al., 2013). The MWA operates in the frequency range from 70 – 320 MHz with an instantaneous bandwidth of 30.72 MHz which can be divided into at most 10 kHz channels. In this work we only simulate the Phase I MWA configuration (Tingay et al., 2013). Most of the EoR observations are taken using this configuration (Bowman et al., 2013).

3.2.1 Fully Element Embedded (FEE) Primary Beam

As mentioned in Section 3.1.3 the instrument applies a frequency dependent spatial taper to the intensity distribution called the primary beam response $B(l, m, \nu)$. The MWA primary beam pattern is a complex structure which is a super position of all the individual dipole beams for a given tile. Each dipole element and the total beam can be described by individual polarisation components. For the MWA these are the linear X (North-South) and Y (East-West) components. Although the MWA dipoles are fixed to the ground, it is possible to electronically point the beam by introducing a delay in the dipole signal chain in the beam forming step. This differs from radio




Figure 3.4: Example MWA primary beam response projected onto the (l, m) orthographic image plane. This beam model was created at 180 MHz using the FEE beam model with the delays set to a zenith pointing. The colourbar shows the beam response in a log scale to highlight the distribution of the MWA primary beam grating sidelobes.

dish telescopes which mechanically slew to their targets. The benefit of being able to physically point at the target is the beam shape is consistent. The electronic delays introduced to observe different parts of the sky have different beam projections, which makes beam modelling and calibrating MWA observations more challenging than that for dish interferometers.

Models of the MWA primary beam are necessary to calibrate the MWA visibility data, and for forward modelling expectations of EoR observations. Of particular importance is understanding the chromatic behaviour of the beam. In particular, how the primary beam sidelobes interact with foreground sources, and by consequence affect the 21cm power spectrum; we revisit this in Chapter 6. Sutinjo et al. (2015) motivated the need for a physical MWA beam model that could explain a large fraction of the instrumental Stokes Q leakage seen in observations. This leakage can reach the 20% of the Stokes I level. Sutinjo et al. (2015) suggested the fully element embedded (FEE) model, which simulates each dipole element pattern independently and then combines them. Using the reciprocity theorem, which defines the response of a dipole as equal to its emitted radiation, the response for each dipole can be generated. For the FEE

beam model, each dipole radiation pattern is simulated separately, taking into account how this emitted radiation stimulates the neighbouring dipoles. By construction this includes the mutual coupling of the individual dipoles; these patterns are then combined together. Sokolowski et al. (2017) describes this process through a spherical wave expansion model. The downside of the FEE model is it requires significantly more computation than the simpler less physical Hertzian dipole model of Ord et al. (2010). As a result of the computational cost, these models were only generated for each coarse channel (1.28 MHz resolution).

Throughout this work we use the FEE beam model to simulate the MWA tile beam response. Figure 3.4 shows an example of a zenith-pointed MWA tile beam at a frequency of $\nu = 150$ MHz. To get estimates for each fine channel (80 kHz channel resolution) we interpolate the beam as a function of frequency, typically interpolating over a larger frequency range than the bandwidth. Notably, we assume that MWA tile beam is the same for each tile. In reality this is not the case. Often individual dipoles within a tile can be inactive⁷, and additionally small variations in position, gain and inclination for each dipole within a tile affect the tile beams as well.

These differences in the FEE tile beam model to real MWA tiles were measured by Line et al. (2018) and Chokshi et al. (2021) using ORBCOMM satellites. They found good comparison with the FEE beam model, however, higher order perturbations and asymmetries particularly in the sidelobes were present. These differences were attributed to environmental effects, particularly from partial concealment of the ground plane from soil, rocks, and foliage (Chokshi et al., 2021). Addressing these concerns are outside the scope of this work. We return to the MWA primary beam in the following section.

3.3 Simulating Visibilities

Throughout this work we use simulated data, primarily simulated visibilities from observations conducted with the MWA. In principle visibility simulation is the estimate

⁷Kangaroos have been known to damage dipoles, but typically LNAs just fail.

of Equation 3.12 using a known model of the sky intensity distribution $I_v(l, m)$, the instrument primary beam pattern $\mathcal{B}(l, m)$, and the array baseline positions (u, v, w) .

In this work we use two different methods to estimate Equation 3.12. The first method Fourier transforms the sky brightness distribution in the image plane (l, m) , and samples the resulting (u, v) plane at the baseline (u, v) coordinates. The curvature and projection term related to the w coordinate are incorporated through a method called w -projection (Cornwell et al., 2008). We discuss this method in detail in Chapter 6, where it is a large component of the method in that project. The other method takes advantage of the linear nature of the Fourier transform, and directly calculates the analytic Equation 3.12 for each source model. This method describes the individual source intensity distributions with analytic models in the image plane, which have corresponding analytic Fourier transforms. Due to the linear nature of the Fourier transform, this can be computed for each source separately and in parallel, at the cost of great computational expense. We discuss this method in detail in Section 3.3.2.

In both cases the sky can be modelled as a linear combination of individual source intensity distributions:

$$I_{\text{sky}}(l, m) = \sum_{i=1}^{N_{\text{source}}} I_i(l, m), \quad (3.24)$$

where $I_i(l, m)$ is the individual source intensity distribution function. We discuss the types of source distribution models used by WODEN ⁸ in Section 3.3.2.

3.3.1 Array Geometry and Baseline Distribution

To simulate radio interferometric visibilities, we need to know the (u, v, w) coordinates for each baseline for a given observation. These coordinates depend on the position of the individual array elements relative to each other, their position on the Earth's surface, and the phase reference centre of the array. The (u, v, w) coordinates are defined relative to a Cartesian coordinate system, and we use the Cartesian coordinates as defined in Thompson et al. (2017). In this system, the X and Y axes are defined in

⁸Documentation and installation instructions for WODEN can be found <https://woden.readthedocs.io/en/latest/installation/installation.html>




Figure 3.5: Right-handed Cartesian coordinate system defined relative Hour angle and declination. The axis alignment with the Hour angle and declination values is shown. Image taken from Thompson et al. (2017).

a plane parallel to the Earth's equator, where the Z axis points towards the North pole, Y points East, and X is parallel to the meridian plane. Figure 3.5 shows the geometry of the coordinate system and how the different axes relate to the Hour angle (H) and declination (δ) coordinate system.

The $(X_\lambda, Y_\lambda, Z_\lambda)$ coordinates for a given baseline can be determined by the baseline separation D_λ , Azimuth (\mathcal{A}), and the elevation (\mathcal{E}) of the baseline. Additionally, the latitude (\mathcal{L}) of the array also needs to be known:

$$\begin{pmatrix} X_\lambda \\ Y_\lambda \\ Z_\lambda \end{pmatrix} = D_\lambda \begin{pmatrix} \cos \mathcal{L} \sin \mathcal{E} - \sin \mathcal{L} \cos \mathcal{E} \cos \mathcal{A} \\ \cos \mathcal{E} \sin \mathcal{A} \\ \sin \mathcal{L} \sin \mathcal{E} + \cos \mathcal{L} \cos \mathcal{E} \cos \mathcal{A} \end{pmatrix} \quad (3.25)$$

The positions of the baselines are dependent on where the array phase reference centre is located. This determines the baseline projection, which can then be decoupled into the individual components (u, v, w) . Figure 3.6 demonstrates the geometric relationship between the (u, v, w) coordinate system and the $(X_\lambda, Y_\lambda, Z_\lambda)$ coordinate system. The relationship is determined by a rotation in the spherical coordinate system:




Figure 3.6: The relationship between the (u, v, w) coordinate system, the $(X_\lambda, Y_\lambda, Z_\lambda)$ coordinate system, and the celestial sphere (H, δ) . In this case w points towards the source S . Image taken from Thompson et al. (2017).

$$\begin{pmatrix} u \\ v \\ w \end{pmatrix} = \begin{pmatrix} \sin H & \cos H & 0 \\ -\sin \delta \cos H & \sin \delta \cos H & \cos \delta \\ \cos \delta \cos H & -\cos \delta \sin H & \sin \delta \end{pmatrix} \begin{pmatrix} X_\lambda \\ Y_\lambda \\ Z_\lambda \end{pmatrix} \quad (3.26)$$

The MWA Phase I has excellent snapshot (u, v) coverage due to the large number of baselines, where the number of baselines is $N(N - 1)$, where N is the number of tiles. Figure 3.7 shows the baseline distribution for a zenith pointed phase reference centre. Notably, $\mathcal{V}(u, v, w)$ is Hermitian which means $\mathcal{V}(-u, -v, -w) = \mathcal{V}^*(u, v, w)$, which results in a symmetric baseline distribution about $u = v = 0$.

3.3.2 WODEN

WODEN is a radio interferometric simulation tool, designed to directly compute Equation 3.12, from analytic expressions for source models (Line, 2022). WODEN simulates the full RIME as a function of position (l, m) as described in Section 3.1.4 for the




Figure 3.7: MWA Phase I, zenith phased (u, v) baseline distribution for a single time step and frequency (~ 180 MHz) channel. The Hermitian baseline values are also plotted.

MWA. The instrumental Jones matrix for the RIME is defined below:

$$\mathbf{J}(l, m) = \begin{pmatrix} g_X(l, m) & D_X(l, m) \\ D_Y(l, m) & g_Y(l, m) \end{pmatrix}, \quad (3.27)$$

g_X is the gain term for the linear X polarisation, and g_Y is the gain term for the linear Y polarisation, these have the same definition as Section 3.1.4. The off-diagonal terms $D_X(l, m)$ and $D_Y(l, m)$ are the instrumental polarisation leakage terms for their respective polarisations.

In Section 3.1.4 we derived the relationship between the instrumental polarisations and the Stokes parameters using a Jones matrix formalism. However, an equally valid representation can be made using Mueller matrices (Sault et al., 1996). For completeness, Smirnov (2011) includes this derivation. This is the formalism used by WODEN. In the Mueller matrix formalism the measurement equation for the visibilities

measured by antennas 1 and 2 are defined below:

$$\begin{pmatrix} \mathcal{V}_{12,XX}(l, m) \\ \mathcal{V}_{12,XY}(l, m) \\ \mathcal{V}_{12,YX}(l, m) \\ \mathcal{V}_{12,YY}(l, m) \end{pmatrix} = \mathbf{J}_1(l, m) \otimes \mathbf{J}_2(l, m) \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & i \\ 0 & 0 & 1 & -i \\ 1 & -1 & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathcal{V}_{12,I}(l, m) \\ \mathcal{V}_{12,Q}(l, m) \\ \mathcal{V}_{12,U}(l, m) \\ \mathcal{V}_{12,V}(l, m) \end{pmatrix} \quad (3.28)$$

The Jones matrix $\mathbf{J}(l, m)$ describes the primary beam pattern of an antenna, which is assumed to be the same for all MWA tiles. The matrix on the right-hand side of Equation 3.28 relates the Stokes parameters to the instrumental polarisations, this is the vector representation of the brightness matrix in Equation 3.21. Equation 3.28 is calculated for all directions l_j, m_j in the sky-model, for each baseline, frequency, and time step. This is then summed over all j , producing the linear Stokes polarisation visibilities, per baseline, frequency, and time.

For a given frequency and time snapshot, the MWA has 8128 baselines (16256 including the Hermitian pairs), each with four linear polarisation XX , XY , YX , and YY . This is calculated for N_{sources} ; the computation scales quickly with the number of sources. WODEN takes advantage of the parallel computing power of Graphical Processing Units (GPUs), which are well suited to these tasks. A typical 2 minute snapshot observation, containing 10^5 sources on a Pawsey supercomputing cluster takes ~ 10 minutes to complete.

3.3.2.1 WODEN Source Models

For a given polarisation and for analytically determined source models Equation 3.12 can be written as the sum of all sources j :

$$\mathcal{V}_S(u_i, v_i, w_i) = \sum_j \xi_j(u_i, v_i) S(l_j, m_j) \exp[-2\pi i(u_i l_j + v_i m_j + w_i(n_j - 1))], \quad (3.29)$$

j indicates the source number, and $\xi_j(u_i, v_i)$ is the source envelope function, which is the amplitude of the Fourier transform of the source distribution function. This is

computed at the (u_i, v_i) for each source j .

The source envelope function is a general expression that describes the analytic image space Fourier transform for a source intensity distribution. WODEN uses multiple source envelope functions: point sources or Dirac delta functions $\delta(l - l_j, m - m_j)$; elliptical Gaussian functions; shapelets (Refregier, 2003). The amplitude of the Fourier transform of a Dirac delta function is a constant. If the Dirac delta is offset from the coordinate origin there is an exponential term $e^{-2\pi i(u_i l_j + v_i m_j)}$ included. Point sources are the simplest models to simulate. To model the more complex morphology of extended sources, a collection of point sources can be used, or they can be modelled with other basis functions such as elliptical Gaussians, and two dimensional shapelets. A more complex source can be constructed from an ensemble of elliptical Gaussians, where an individual elliptical Gaussian function is defined by:

$$G(l, m) = \exp \left(-4 \ln(2) \left[\frac{(l - l_0)^2}{\theta_{\text{maj}}^2} + \frac{(m - m_0)^2}{\theta_{\text{min}}^2} \right] \right), \quad (3.30)$$

l_0, m_0 is the central location of the Gaussian in the image plane, θ_{min} and θ_{maj} are the minor and major elliptical axes of the Gaussian, which are determined as the Full Width Half Maximum (FWHM) of the 1D Gaussian along the minor and major axes. The Fourier transform of this in the (u, v) plane is also a 2D elliptical Gaussian:

$$\xi_j = \exp \left(-\frac{\pi^2}{4 \ln(2)} \left(k_x^2 \theta_{\text{maj}}^2 + k_y^2 \theta_{\text{min}}^2 \right) \right) \quad (3.31)$$

$$k_x = \cos(\phi_{\text{PA}})v_i + \sin(\phi_{\text{PA}})u_i \quad (3.32)$$

$$k_y = -\sin(\phi_{\text{PA}})v_i + \cos(\phi_{\text{PA}})u_i \quad (3.33)$$

Another useful set of basis functions are shapelets, which are the combination of a Hermite polynomial with a Gaussian function (Refregier, 2003). Riding et al. (2017) introduced the use of shapelets for modelling extended radio sources for MWA sources, which was adapted and improved upon for MWA EoR observations by Line et al. (2020). Notably, Line et al. (2020) demonstrated the improvement of shapelet models over existing multi-component Gaussian models for the bright extended source Fornax

A. Shapelets like the Gaussian basis also have analytic Fourier transforms. The one dimensional shapelet function defined for l is given by:

$$B_p(l; \beta) \equiv \beta^{-1/2} [2^p \pi^2 p!]^{-1/2} H_p(\beta^{-1} l) \exp\left(-\frac{l^2}{2} \beta^2\right). \quad (3.34)$$

β is a scaling factor and $H_p(x)$ is a Hermite polynomial of order p where p is a positive integer. The two dimensional image plane shapelet function is described by combining two orthogonal basis shapelet functions:

$$B_{p_1, p_2}(l, m; \beta_1, \beta_2) = B_{p_1}(l; \beta_1) B_{p_2}(m; \beta_2). \quad (3.35)$$

The analytical Fourier transform of the shapelet is given by:

$$\tilde{B}_{p_1, p_2}(k_x, k_y) = i^{p_1 + p_2} B_{p_1, p_2}(k_x, k_m; \beta_1^{-1}, \beta_2^{-1}). \quad (3.36)$$

A shapelet model for a given source can be constructed from a number of components, each being independent from the other. This is the same method used for constructing multi-component Gaussian models. The coordinate system in Fourier space and therefore the shapelet envelope function is defined as:

$$\xi_j = \sum_{k, l}^{p_k + p_l \leq p_{\max}} C_{p_k, p_l} \tilde{B}_{p_k, p_l}(k_x, k_y), \quad (3.37)$$

$$k_x = \frac{\pi}{\sqrt{2 \ln(2)}} [\cos(\phi_{\text{PA}}) v_i + \sin(\phi_{\text{PA}}) u_i], \quad (3.38)$$

$$k_y = \frac{\pi}{\sqrt{2 \ln(2)}} [-\sin(\phi_{\text{PA}}) v_i + \cos(\phi_{\text{PA}}) u_i]. \quad (3.39)$$

Equations 3.31 and 3.37 describe the analytic Fourier transforms of the source morphologies. However, the flux density for each source, and for each independent component S_{ji} varies as a function of frequency. The spectral energy distribution (SED) of a radio source is typically described by a power law, where the flux density is proportional to ν^α . The spectral index is given by α which typically has a value of $\alpha \sim -0.7$ for synchrotron emission (Conway et al., 1963; Kardashev, 1962). The full

power law description is given below:

$$S_i = S_0 \left(\frac{\nu_i}{\nu_0} \right)^\alpha, \quad (3.40)$$

ν_0 is the reference frequency, and S_0 is the flux density at the reference frequency. However, some sources demonstrate more complex frequency dependent behaviour, most notably turning points in their SED. For synchrotron emission, these turning points are caused by synchrotron self absorption (van der Laan, 1966). Turning points are typically modelled using a curved power law, which is simply a parabola in log space (Blundell et al., 1999):

$$S_i = S_0 \left(\frac{\nu_i}{\nu_0} \right)^\alpha e^{q(\ln \frac{\nu_i}{\nu_0})^2} \quad (3.41)$$

α is still the spectral index at the reference frequency, and q is the curvature term defined at the reference frequency.

3.4 Summary

In this chapter we have developed the theoretical framework by which radio interferometers measure the brightness temperature of the celestial sphere. Furthermore, we have discussed two methods for how to simulate radio interferometer visibilities, with particular focus on the MWA. This work is fundamental to the projects presented in Chapter 6 and Chapter 8, where we utilise simulated visibilities to estimate the foreground power spectrum and skew spectrum respectively.

CHAPTER 4

STOCHASTIC PROCESSES

As we will see in Chapter 5 practical measurements with existing instruments are focused on various statistics of the cosmological 21cm signal. To understand, and correctly interpret the astrophysical and cosmological information encoded in the 21cm statistical signal, we must first understand the nature of stochastic (or random) processes. This chapter develops that mathematical framework and concepts following from the derivations of Chapter 2 from Mandel and Wolf (1995) and Chapter 9 from Papoulis (1991).

A stochastic process can be thought of as some function $f(x)$ that does not depend in a deterministic manner on the variable x . In this case we consider x to be a spatial variable (units of length), however it is standard practice to use time or t as the dependent variable (as is the case in Mandel and Wolf (1995) and Papoulis (1991)). The choice is however arbitrary, and the concepts are generalised to any space. In Chapter 5 we will see that x can also be a vector associated with an arbitrary number of spatial coordinates. This is of particular importance when dealing with random fields that have two or three spatial dimensions. In this case for simplicity and without the loss of generality we deal with only a single spatial dimension.

4.1 Ensemble Average

An important aspect of random or stochastic processes, is that they can only be described by their statistical properties. For the process $f(x)$ this is a direct consequence of the character of variations being independent of the position x . In this section we describe the idea of the ensemble average, which will be critical to understanding stochastic processes for cosmological 21cm statistics. Later in this chapter, we will see how under the right conditions we can accurately estimate the ensemble average. This will be important in Chapter 5 when we discuss the statistics of the 21cm signal.

Consider the random process $f(x)$. This function is effectively a random variate, that is drawn from some domain. This domain is described by a probability density function $p(f(x))$, or probability distribution function $p(f(x))$, which satisfies the condition $\int p(f(x)) df(x) = 1$. We define the expectation value of $f(x)$ at the position x :

$$\langle f(x) \rangle = \int f(x) p(f(x)) df(x). \quad (4.1)$$

Another definition is to consider the set of all realisations of $f(x)$ for a fixed x , where the set is defined as $\{(^{(1)}f(x), ^{(2)}f(x), \dots, ^{(n)}f(x))\}$. The superscript preceding $f(x)$ here assigns an index for the i th realisation in the set. This set is the ensemble and as $n \rightarrow \infty$ it encapsulates the totality of the random process for the fixed x . We can then define the ensemble average in the limit $N \rightarrow \infty$:

$$\langle f(x) \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N ^{(i)}f(x). \quad (4.2)$$

Both definitions outlined in Equation 4.1 and 4.2 describe the ensemble average or the expectation of $f(x)$ because we sum over all possible realisations defined in the space $p(f(x))$ ¹. From this one operation, we can build all the important statistical tools and relationships contained in $f(x)$.

¹In this chapter we use the term expectation value, and ensemble average interchangeably.

4.2 Correlation Functions

In the previous section we discussed the expectation value of $f(x)$ at a fixed arbitrary point x . In totality, $p(f(x))$ describes an infinite family of probability densities as a function of x , since the expectation of $f(x_1)$ is not necessarily equal to that of $f(x_2)$. Although $p(f(x))$ contains information about the expectation value at x , it does not however elucidate any potential correlations that might exist between $f(x_1)$ and $f(x_2)$. This information is instead contained in the joint probability distribution function $p_2(f(x_1), f(x_2))$. Additionally, the joint probability distribution contains all the information present in $p_1(f(x))$. This means that it satisfies the condition that if we integrate the joint distribution over $f(x_2)$, we should retrieve $p_1(f(x_1))$:

$$\int p_2(f(x_1), f(x_2)) df(x_2) = p_1(f(x_1)), \quad (4.3)$$

The subscript denotes the order of the probability distribution. From the joint probability distribution we can calculate the two point spatial correlation function by calculating the ensemble average:

$$\xi_2(x_1, x_2) \equiv \langle f(x_1)f(x_2) \rangle = \int f(x_1)f(x_2)p_2(f(x_1), f(x_2)) df(x_1)df(x_2). \quad (4.4)$$

The two-point correlation function contains the next highest order information after the mean $\langle f(x) \rangle$, and is the quantity that describes correlations and how they extend over distance (in the one dimensional case). For random processes in higher dimensional spaces, this contains information about correlations between position vectors. The two-point correlation function and the joint probability distribution still do not fully describe the random process. We can similarly describe a three-point correlation function $\xi_3(x_1, x_2, x_3) = \langle f(x_1)f(x_2)f(x_3) \rangle$ (which will be important later), that requires a three-point joint probability distribution $p_3(f(x_1), f(x_2), f(x_3))$. In fact, there is an infinite number of higher order correlation functions ξ_n (the subscript indicates the correlation order) and their associated joint distributions $p_n(f(x_n), f(x_{n-1}), \dots, f(x_1))$,

each containing more information about the random process than the last. That is, they satisfy the same condition presented in Equation 4.3.

4.3 Stationary and Ergodic Processes

Now we consider a special class of random processes that have statistical properties that are invariant as a function of x . In this case the statistics of $f(x)$ are identical under translation $f(x+r)$, and stochastic processes that satisfy this condition are said to be stationary. A stochastic process is said to have n -order stationarity if its statistics satisfy the following relation:

$$p_n(f(x_n), f(x_{n-1}), \dots, f(x_1)) = p_n(f(x_n+r), f(x_{n-1}+r), \dots, f(x_1+r)). \quad (4.5)$$

Wide sense stationary processes only satisfy the relation in Equation 4.5 up to $n = 2$. For this work, we assume strict stationarity (up to arbitrary order). For a more in depth discussion on the topic we suggest chapter 9 from (Papoulis, 1991).

From Equation 4.5 it is evident that the ensemble average $\langle f(x) \rangle = \text{const} \forall x$. Similarly for the two-point correlation function, we see that we can rewrite Equation 4.4 in terms of the difference between x_2 and x_1 :

$$\xi_2(x_1, x_2) = \int f(x_1)f(x_2)p_2(f(x_1), f(x_2)) df(x_1)df(x_2) \quad (4.6)$$

$$= \int f(x_1)f(x_2)p_2(f(x+x_2-x_1), f(x)) df(x_1)df(x_2) \quad (4.7)$$

$$= \langle f(x)f(x+x_2-x_1) \rangle, \quad (4.8)$$

Therefore $\xi_2(x_1, x_2) = \xi_2(x_2 - x_1)$ and only depends on the difference between x_1 and x_2 , which we now represent with the arbitrary variable r . It should also be noted that $\xi_2(r) = \xi_2(-r)$ is symmetric, or Hermitian for complex processes. For higher order stationarity this can be generalised, where for the three-point correlation function $\xi_3(x_1, x_2, x_3)$ only depends on the difference between $r = x_2 - x_1$, $s = x_3 - x_1$, and therefore $s - r$.

4.3.1 Ergodic Processes

The ensemble average or expectation of a stochastic process cannot always be practically measured. In real cases we do not always know $p(f(x))$, or the more general joint probability distributions. Additionally we do not have the infinite ensemble with which to measure the statistics. Ergodic stochastic processes are those where the statistics of $f(x)$ can be determined from a sufficiently large observation. In this case we can specify the mean of some sample:

$$\bar{f}(x) = \lim_{L \rightarrow \infty} \frac{1}{L} \int_{-L/2}^{L/2} f(x) dx, \quad (4.9)$$

where in the limit of $L \rightarrow \infty$, $\bar{f}(x) = \langle f(x) \rangle$. In reality $\bar{f}(x)$ is itself a random variable and not a constant, where the variance in $\bar{f}(x)$ depends on the size of L , and should tend to zero as L tends to infinity. Likewise ergodicity can be defined for the other moments, and correlation functions as well. The mean of ergodic random processes is analogous to the sample mean, and how the sample mean for a large set of number tends to population mean (law of large numbers). Stationarity and ergodicity are the two fundamental concepts required for the determination of the cosmological 21cm signal. We revisit these ideas, and how they manifest on cosmological scales in Chapter 5.

4.3.2 Power Spectral Density

One of the most important attributes of an ergodic stationary random process $f(x)$ is its power spectral density (or the *power spectrum*) $P(k)$, which is defined as $\langle |\hat{f}(k)|^2 \rangle$. If the Fourier transform $\hat{f}(k)$ of $f(x)$ exists, then the power spectral density is the normalised Fourier transform of the two-point correlation function, as defined by the Wiener-Khinchine theorem (Wiener, 1930; Khintchine, 1934):

$$P(k) = \frac{1}{L} \int_{\mathbb{R}} dr e^{-ikr} \xi_2(r). \quad (4.10)$$

In this case we define the forward and the backward Fourier transforms in a similar

manner to Peacock (1998) and Obreschkow et al. (2012):

$$\hat{f}(k) = \frac{1}{L} \int_{\mathbb{R}} dx e^{-ikx} f(x) \quad (4.11)$$

$$f(x) = \frac{L}{(2\pi)} \int_{\mathbb{R}} dk e^{ikx} \hat{f}(k). \quad (4.12)$$

The power spectrum, effectively measures the amplitude of fluctuations (variance) for a particular spatial scale k , which in this case has units of inverse length.

We can show the relationship between $P(k)$ and $\xi(r)$ by starting with $\langle f(x)f(x') \rangle$ and substituting for $f(x)$ from Equation 4.11:

$$\xi_2(x, x') = \frac{L^2}{(2\pi)^2} \left\langle \iint_{\mathbb{R}^2} dk dk' e^{ikx} e^{ik'x'} \hat{f}(k) \hat{f}(k') \right\rangle. \quad (4.13)$$

As defined in Section 4.1 for an ergodic process the ensemble average is the integral over a large enough length (area, or volume for higher dimensions) $\langle \rangle = 1/L \int dx$ as defined in Equation 4.9. Additionally, for an ergodic stationary process we can assume that $r = x - x'$, thus Equation 4.13 becomes:

$$\xi_2(r) = \frac{L}{(2\pi)^2} \iiint_{\mathbb{R}^3} dr dk dk' e^{-ik'r} e^{ix(k'+k)} \hat{f}(k) \hat{f}(k'). \quad (4.14)$$

Substituting for the Dirac delta function which is defined as $\delta(x-r) = 1/(2\pi) \int_{\mathbb{R}} e^{ip(x-r)} dp$, and integrating over k' , yields:

$$\xi_2(r) = \frac{L}{(2\pi)} \int_{\mathbb{R}} dk e^{irk} \hat{f}(k) \hat{f}^*(k), \quad (4.15)$$

where $\hat{f}^*(k) = \hat{f}(-k)$ is the conjugate of $\hat{f}(k)$, which is Hermitian. Therefore, $\hat{f}(k) \hat{f}^*(k) = |\hat{f}(k)|^2$, and thus:

$$\xi_2(r) = \frac{L}{(2\pi)} \int_{\mathbb{R}} dk e^{irk} P(k). \quad (4.16)$$

Through Fourier inversion we retrieve the expression in the Wiener-Khinchine theorem in Equation 4.10.

There is an important caveat to the derivations in Equations 4.13 and 4.16; we assume that the continuous random process $f(x)$ is both integrable and square integrable. However, random processes do not tend to zero as $x \rightarrow \pm\infty$; this results in divergent integrals. Wiener (1930) and Khintchine (1934) show that this is not an impediment so long as the two-point correlation function exists, is monotonic and finite. If so there exists a Fourier relationship where power spectral density exists.

4.3.3 Two Point Fourier Correlation

For the stochastic process $f(x)$ consider some realisation $^{(j)}f(x)$, which has the Fourier pair $^{(j)}\hat{f}(k)$. This naturally implies that the function $\hat{f}(k)$ is itself a stochastic process. We can equivalently define a two-point correlation function for $\hat{f}(k)$, where the general form of the two point correlation function for complex valued stochastic process is the ensemble average of the product $\hat{f}(k)\hat{f}^*(k')$:

$$\langle \hat{f}(k)\hat{f}^*(k') \rangle = \left\langle \frac{1}{L^2} \iint_{\mathbb{R}^2} dx dx' f(x)f(x')e^{ikx}e^{-ik'x'} \right\rangle, \quad (4.17)$$

where the ensemble average and the integrals are interchangeable:

$$\langle \hat{f}(k)\hat{f}^*(k') \rangle = \frac{1}{L^2} \iint_{\mathbb{R}^2} dx dx' \langle f(x)f(x') \rangle e^{ikx}e^{-ik'x'}. \quad (4.18)$$

For a stationary and ergodic process we have the definition of the two-point correlation function in the space x as $\langle f(x)f(x') \rangle = \xi_2(r)$. In this case without loss of generality we define $r = x' - x$, therefore $dr = dx'$:

$$\langle \hat{f}(k)\hat{f}^*(k') \rangle = \frac{1}{L^2} \iint_{\mathbb{R}^2} dx dr \xi_2(r)e^{ix(k-k')}e^{-ikr}. \quad (4.19)$$

Substituting for the Dirac delta function, and Fourier transforming $\xi_2(r)$:

$$\langle \hat{f}(k)\hat{f}^*(k') \rangle = \frac{(2\pi)}{L} \delta_D(k - k')\hat{\xi}_2(k), \quad (4.20)$$

where from Equation 4.10 $\hat{\xi}_2(k) = P(k)$, therefore:

$$\langle \hat{f}(k) \hat{f}^*(k') \rangle = \frac{(2\pi)}{L} \delta_D(k - k') P(k). \quad (4.21)$$

Equation 4.21 demonstrates that the Fourier components of an ergodic stationary process are uncorrelated. We will see the definition of Equation 4.21 as it is often used as another way of defining the power spectrum.

4.3.4 Higher Order Poly Spectra

To study the 21cm signal in reality, we have to work with three dimensional data, so it is important to extend the discussion here to multiple dimensions. The expressions derived in the previous sections have equivalent counterparts for higher order spatial dimensions. From hereon we express these terms for three-dimensional quantities where $\mathbf{x} = [x, y, z]$ and $\mathbf{k} = [k_x, k_y, k_z]$.

Stochastic signals which are fully described by their mean and covariance (two-point correlation) are said to be Gaussian in nature. That is, they can be fully described by a random Gaussian function $\mathcal{N}(\mu, \mathbf{C})$ with mean μ and covariance matrix \mathbf{C} . Processes that contain non-Gaussianity require higher order statistics to characterise their stochastic nature. Equations 4.10 and 4.21 can be generalised for higher order correlation functions, such as the three and four-point correlation functions. For arbitrary order correlation function there is an equivalent arbitrary order poly spectrum (Peebles, 1980; Obreschkow et al., 2012). Of particular importance is the bispectrum, which is the Fourier transform of the three-point correlation function:

$$B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv \frac{1}{V^2} \iint_{\mathbb{R}^6} d^3\mathbf{r} d^3\mathbf{s} e^{-i\mathbf{k}_1 \cdot \mathbf{r}} e^{-i\mathbf{k}_2 \cdot \mathbf{s}} \xi_3(\mathbf{r}, \mathbf{s}, \mathbf{r} - \mathbf{s}) \quad (4.22)$$

$$\langle \hat{f}(\mathbf{k}_1) \hat{f}(\mathbf{k}_2) \hat{f}^*(\mathbf{k}_3) \rangle = \frac{(2\pi)^3}{V} \delta_D(\mathbf{k}_1 + \mathbf{k}_2 - \mathbf{k}_3) B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3). \quad (4.23)$$

In Equations 4.22 we define the bispectrum ($B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3)$) from the three-dimensional three-point correlation function ($\xi_3(\mathbf{r}, \mathbf{s}, \mathbf{r} - \mathbf{s})$), where the Fourier transform is per-

formed over the \mathbf{x} vector space. The Dirac delta in this case specifies that all Fourier modes which do not form a closed triangle in k -space are uncorrelated. Obreschkow et al. (2012) demonstrates the derivation of the arbitrary order poly spectrum for an D -dimensional scalar random field.

4.4 Summary

This chapter develops the fundamental statistical processes, and assumptions that underpin all the 21cm statistical analysis in the following chapters. In the next chapter we apply these assumptions to the cosmological 21cm differential temperature brightness, and we use them to establish the 21cm power spectrum and bispectrum.

CHAPTER 5

21CM STATISTICS

The current generation of radio instruments is not sensitive enough to directly image the cosmological 21cm signal during reionisation and earlier epochs. However, Zaldarriaga et al. (2004) showed that the statistics of the 21cm fluctuations should be accessible. Importantly, the statistical impact of the topology of ionisation of hydrogen on the 21cm signal, should be distinguishable from that of the matter distribution (Zaldarriaga et al., 2004). Furthermore, Furlanetto et al. (2004b) showed that the power spectrum of the 21cm fluctuations could be used to differentiate different reionisation histories. Morales and Hewitt (2004) discuss the practical measurement of the 21cm signal with low frequency radio interferometers such as the MWA, utilising the full three dimensional Fourier space probed by interferometers to measure the 21cm signal. This method naturally separates the continuous and smooth foreground contamination from the 21cm signal, which is expected to be highly varying with frequency.

However, the power spectrum only contains information about the spatially dependent variance of the signal, effectively assuming the underlying distribution is Gaussian. Morales and Hewitt (2004) showed that the 21cm signal is expected to be highly non-Gaussian prior to and during reionisation. Prior to reionisation, during the epoch of heating (EoH), non-Gaussianity originates from the inhomogeneous heating of the

IGM, primarily from X-ray emission from the first luminous sources (Furlanetto, 2006). As X-ray heating progresses, it eventually saturates the IGM, and the non-Gaussianities are driven toward the matter density non-Gaussianities (Watkinson et al., 2018). During reionisation the IGM is highly heated above the CMB temperature, and becomes saturated ($T_S \gg T_{\text{CMB}}$) (Ciardi and Madau, 2003; Chen and Miralda-Escudé, 2004; Furlanetto et al., 2006a). During this period, the non-Gaussianities become dominated by the ionisation topology (Hutter et al., 2019), and the matter density (Lidz et al., 2007). To characterise the non-Gaussianities of the cosmological 21cm signal, higher order spectra such as the bispectrum and the trispectrum are required.

In this chapter we summarise the astrophysical and cosmological information contained in the statistics of the 21cm signal. We then show their connection to radio interferometric measurements, and outline some of the challenges associated with these measurements. We then discuss the ongoing and future radio interferometric experiments.

5.1 Cosmological Fields

In general the temperature brightness of some distribution of baryonic matter can be described as a scalar Eulerian field $T(\mathbf{x})$ which is a function of position \mathbf{x} . In this case, T is a general description for any arbitrary baryon distribution. Equally as valid is the matter distribution (baryon or dark matter), often described by the density field $\rho(\mathbf{x})$. Of particular interest is the temperature fluctuation (perturbation) field, which is defined as (Peacock, 1998):

$$\delta_T(\mathbf{x}) \equiv \frac{T(\mathbf{x}) - \langle T \rangle}{\langle T \rangle}, \quad (5.1)$$

where the angular brackets indicate the ensemble average. Naturally, $\langle \delta_T(\mathbf{x}) \rangle = 0$.

The field $\delta_T(\mathbf{x})$ can essentially be considered a random field, that on the cosmological scales is statistically homogeneous and isotropic (Peacock, 1998). These assumptions are critical because they relate to the concepts of stationary and ergodic

processes discussed in Chapter 4. Homogeneity and isotropy ensure that the statistics of the $\delta_T(\mathbf{x})$ as a function of translation $\mathbf{x} + \mathbf{r}$ and rotation $R\mathbf{x}$ (R represents some rotation matrix) are invariant. Importantly we assume the process is ergodic, an important consideration for we do not have access to an infinite ensemble of Universes; a reasonable assumption. Assuming ergodicity, we can define the ensemble average $\langle T(\mathbf{x}) \rangle$ as the volume average \bar{T} :

$$\langle T(\mathbf{x}) \rangle = \bar{T} = \frac{1}{V} \int_V T(\mathbf{x}) dV. \quad (5.2)$$

We can equally describe the 21cm temperature brightness difference fluctuations:

$$\delta_{21}(\mathbf{x}) \equiv \frac{\delta T_b(\mathbf{x}) - \delta \bar{T}_b}{\delta \bar{T}_b}. \quad (5.3)$$

From Equation 5.3 we can derive all the fundamental statistical quantities such as the power spectrum and bispectrum, from the Fourier transform of δ_T as demonstrated in Chapter 4. In the next sections we will see this description in practice, and what information it contains about astrophysics and cosmology.

5.2 21cm Power Spectrum

As mentioned in Section 4.3.2 the power spectrum is the Fourier transform of the two point correlation function, and encodes structural information about the 21cm signal by measuring the strength of fluctuations as a function of spatial scale.

Radio interferometers measure the Fourier component of the temperature field, and thus are sensitive to the spatial fluctuations of the 21cm signal. The power spectrum is therefore straightforward to estimate from radio interferometers. We discuss this more in Section 5.4.1 of this chapter, and we go into specific details about how we estimate the power spectrum from visibilities in Chapter 6. From Equation 4.21 in Chapter 4 we have the expression for the power spectrum, computed as the angle-averaged Fourier




Figure 5.1: Example 1D spherically averaged dimensionless power spectrum of a simulated 21cm signal taken from Balu et al. (2023). This figure shows the evolution of the power spectrum as a function of redshift and ionisation state.

transform of $\delta_{21}(\mathbf{x})$:

$$\langle \hat{\delta}_{21}(\mathbf{k}) \hat{\delta}_{21}^*(\mathbf{k}') \rangle = \delta \bar{T}_b^2 (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') P_{21}(k) [\text{mK}^2 \text{Mpc}^3]. \quad (5.4)$$

For an isotropic universe the temperature field can not have a preferred direction, thus $\langle P(\mathbf{k}) \rangle = P(k)$, where $P(k)$ is the isotropic power spectrum. $P(k)$ typically has units of $\text{mK}^2 \text{Mpc}^3$. It is therefore standard practice to normalise out the spatial scales, and thus we define the dimensionless power spectrum as

$$\Delta_{21}^2(k) \equiv \frac{k^3}{2\pi^2} P_{21}(k) [\text{mK}^2]. \quad (5.5)$$

In Figure 5.1 we show an example of the dimensionless power spectrum, for a Fiducial 21cm signal at different redshifts and ionisation states during reionisation.

5.2.1 Astrophysics of the 21cm Power Spectrum

The power spectrum contains a wealth of astrophysical information regarding the 21cm signal. This is apparent from panel (b) in Figure 5.2, which shows the power spectrum




Figure 5.2: Example of the global signal temperature $\delta\bar{T}_b$ and the power at two fixed spatial scales $k = 0.5 \text{ Mpc}^{-1}$ (dotted) and $k = 0.1 \text{ Mpc}^{-1}$ (solid line). Taken from Mesinger et al. (2016).

for a simulated 21cm signal taken from Mesinger et al. (2016) as a function of redshift at two spatial scales, $k = 0.1 \text{ Mpc}^{-1}$ (solid line) and $k = 0.5 \text{ Mpc}^{-1}$ (dotted line). In both panels (a) and (b) of Figure 5.2 we see numerous peaks and turning points as a function of redshift (frequency). These features are coupled to the evolution of the spin temperature, and reveal different pieces of information regarding the cosmological evolution of the 21cm signal, and thus the Universe as a function of redshift. We list the different features in detail:

- i Dark Ages: at these higher redshifts the neutral IGM is dense enough that the spin temperature is coupled to the kinetic gas temperature through collisions T_K . As the IGM adiabatically expands it cools $T_K \propto (1+z)^2$, faster than the CMB temperature as $T_{\text{CMB}} \propto (1+z)$ (Furlanetto, 2019). This drives the differential temperature contrast into absorption, resulting in a peak in the power spectrum, as the fluctuations are driven by the gas density distribution. By $z \sim 30$ the IGM density drops, collisions become less important, and the spin temperature couples again with the CMB (Furlanetto, 2019).
- ii Ly α Coupling: As the first stars form, they couple the spin temperature to the kinetic temperature through Ly α pumping via the WF-effect (Wouthuysen, 1952; Field, 1958). This leads to a deep absorption trough of the differential temperature brightness relative to the CMB. This drives the amplitude fluctuations, and thus results in a peak in the power spectrum.

- iii X-ray Heating: In this period X-rays from the first astrophysical sources, in particular high mass X-ray binaries, heat the IGM through secondary collisions from ionisation. This drives the IGM into emission relative to the CMB. This results in another peak in the power spectrum around $z \sim 15$. Preionisation of the IGM occurs in this period.
- iv Reionisation: In this period the IGM begins to undergo reionisation predominantly from UV photons. The temperature reaches its maximum during the midpoint of reionisation, and then drops to zero as reionisation consumes the remaining neutral hydrogen gas. This results in a peak in the power spectrum during the mid point of reionisation.

It is clear from Figure 5.2 that the amplitude and frequency (redshift) of the mean 21cm brightness and power spectrum depend highly on underlying astrophysical processes. Each of the different turning points is coupled to the dominant astrophysical processes during that period. During the EoR as the spin temperature saturates $T_S \gg T_{\text{CMB}}$, the differential temperature brightness reduces to an approximate form that is proportional to the ionisation fraction. Therefore, this can be used to estimate the timing and span of reionisation, as well as the number density of ionising photons.

5.2.2 21cm Power Spectrum During EoR

During the EoR, the spin temperature can become saturated $T_S \gg T_{\text{CMB}}$ if heated to well above the background CMB radiation field (Ciardi and Madau, 2003; Chen and Miralda-Escudé, 2004; Furlanetto et al., 2006a,b). In this case, the 21cm temperature brightness is independent of the spin temperature in Equation 2.21, and thus reduces to (Barkana and Loeb, 2005b; Furlanetto et al., 2006a; Lidz et al., 2007):

$$\delta T_{21}(\mathbf{x}) = T_0 \bar{x}_{\text{HI}} (1 + \delta_x(\mathbf{x})) (1 + \delta_\rho(\mathbf{x})), \quad (5.6)$$

where δ_x is the fractional fluctuation in the hydrogen neutral fraction $\delta_x = (x_{\text{HI}} - \bar{x}_{\text{HI}})/\bar{x}_{\text{HI}}$, δ_ρ is the gas density, and T_0 is a normalisation factor calculated from Equa-

tion 5.6 (Lidz et al., 2007). Expanding Equation 5.6 and calculating the two-point correlation function yields:

$$\begin{aligned}
\langle \delta_{21}(\mathbf{x}_1)\delta_{21}(\mathbf{x}_2) \rangle &= T_0^2 \langle x_{\text{HI}}(\mathbf{x}) \rangle^2 \left[\langle \delta_x(\mathbf{x}_1)\delta_x(\mathbf{x}_2) \rangle \right. \\
&\quad + 2\langle \delta_x(\mathbf{x}_1)\delta_\rho(\mathbf{x}_2) \rangle + \langle \delta_\rho(\mathbf{x}_1)\delta_\rho(\mathbf{x}_2) \rangle \\
&\quad + 2\langle \delta_x(\mathbf{x}_1)\delta_\rho(\mathbf{x}_1)\delta_x(\mathbf{x}_2) \rangle \\
&\quad + 2\langle \delta_x(\mathbf{x}_1)\delta_\rho(\mathbf{x}_1)\delta_\rho(\mathbf{x}_2) \rangle \\
&\quad \left. + \langle \delta_x(\mathbf{x}_1)\delta_\rho(\mathbf{x}_1)\delta_x(\mathbf{x}_2)\delta_\rho(\mathbf{x}_2) \rangle \right] \\
&\quad + \text{constant.}
\end{aligned} \tag{5.7}$$

Equation 5.7 demonstrates the separable nature of the 21cm signal in the saturated spin temperature limit. It also demonstrates the presence of higher order statistical moments from the cross-correlations of the density and neutral fraction fields (as also seen by Cooray (2005)). The power spectrum can then be defined by taking the Fourier transform of Equation 5.7:

$$\begin{aligned}
P_{21}(k) &= T_0^2 \langle x_{\text{HI}}(\mathbf{x}) \rangle^2 \left[P_{xx}(k) + 2P_{x\rho}(k) + P_{\rho\rho}(k) \right. \\
&\quad \left. + 2P_{x\rho,x}(k) + 2P_{x\rho,\rho}(k) + P_{x\rho,x\rho}(k) \right]
\end{aligned} \tag{5.8}$$

Since the Fourier transform is a linear operation we see that in the saturated spin temperature limit, that the power spectrum is the sum of the gas density, neutral fraction and cross terms spectra. There exists a higher order moment $P_{x\rho,x\rho}$ which probes the non-Gaussian nature of the 21cm signal. We discuss this in more detail in Chapter 7. The separable components of the power spectrum in this limit have been studied extensively (see Cooray, 2005; Furlanetto et al., 2006a,b; Lidz et al., 2007, for examples).

5.3 21cm Bispectrum

Non-Gaussianity, and by extension higher order statistics, have been shown to be important for constraining the astrophysics of the 21cm signal, as well as for confirming its detection (Shimabukuro et al., 2017; Tiwari et al., 2022). The lowest order statistic sensitive to non-Gaussianity is the three-point correlation function, which we described in Chapter 4:

$$\xi_3(\mathbf{r}, \mathbf{s}, \mathbf{r} - \mathbf{s}) \equiv \langle \delta(\mathbf{x} + \mathbf{r})\delta(\mathbf{x} + \mathbf{s})\delta(\mathbf{x}) \rangle. \quad (5.9)$$

As in the case of the auto-correlation function, since the process is stationary and ergodic, the three-point correlation function only depends on the differences \mathbf{r} , \mathbf{s} , and $\mathbf{r} - \mathbf{s}$. Likewise, the Fourier transform of the three-point correlation function with respect to \mathbf{r} and \mathbf{s} yields the bispectrum, which is often defined as:

$$\langle \tilde{\delta}_T(\mathbf{k}_1)\tilde{\delta}_T(\mathbf{k}_2)\tilde{\delta}_T^*(\mathbf{k}_3) \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3)B(k_1, k_2, k_3), \quad (5.10)$$

where the delta function requires that the wave-vectors $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$ form a closed triangle. The bispectrum probes the degree of coherence for a random field for those three wave-vectors.

5.3.1 Understanding the Bispectrum

The bispectrum probes a large parameter space of potential triangle configurations, encoding different types of structural correlations. The simplest triangle configuration is the equilateral configuration where $|\mathbf{k}_1| = |\mathbf{k}_2| = |\mathbf{k}_3|$, which describes the bispectrum with a single wave-vector amplitude k , much like the power spectrum. In Figure 5.3 we see an example of the equilateral triangle configuration taken from Lewis (2011).

The wave-vectors combine in real space to probe circularly symmetric filamentary structures in the random field. Notably, the equilateral configuration typically has the highest amplitude, and as such has been extensively studied in the 21cm literature (Bharadwaj and Saiyad Ali, 2005; Yoshiura et al., 2015; Shimabukuro et al., 2016,

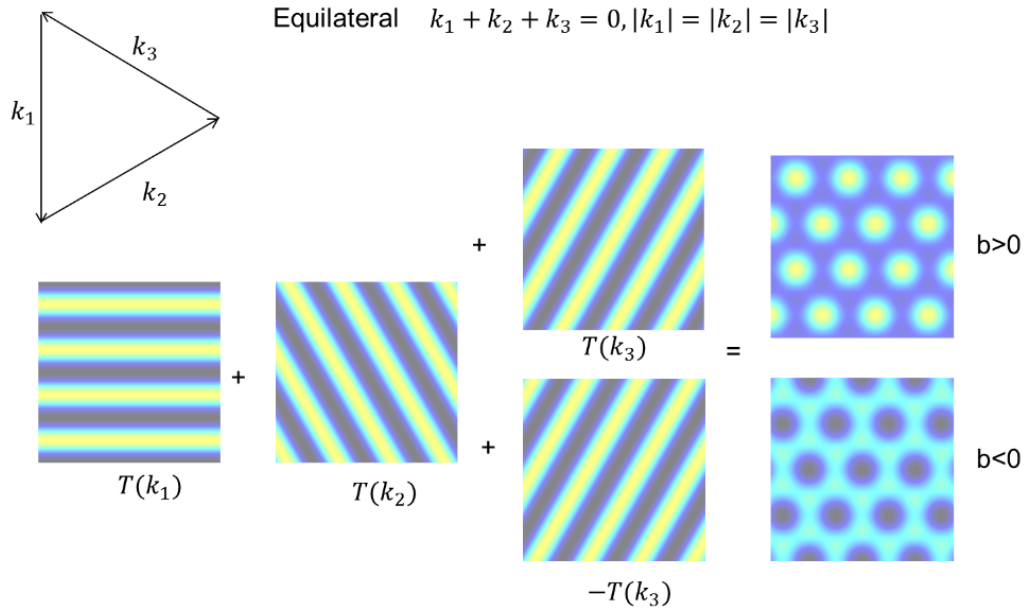


Figure 5.3: An example of the bispectrum for the equilateral triangle configuration, taken from Lewis (2011). The fields $T(k_1)$, $T(k_2)$ and $T(k_3)$ represent the different 2D wavemodes projected onto a 2D image. These add together to in two configurations $b > 0$ for positive amplitude, and for $b < 0$ for negative amplitude. Values of the bispectrum greater amplitude ($b > 0$) represent circular cross-sections with above average amplitude (circularly symmetric clumping). Negative values ($b < 0$) represent circularly symmetric voids.

2017; Watkinson et al., 2017, 2018; Majumdar et al., 2018; Hutter et al., 2019; Trott et al., 2019). The equilateral configuration effectively probes the spherical symmetry of the field, where positive values ($b > 0$ in Figure 5.3) indicate an above average number of spherical regions, and negative values indicate an above average number of spherically symmetric void structures (HII bubbles during reionisation for example). As such, the bispectrum, unlike the power spectrum, can be negative. The sign of the bispectrum has been shown to be an important potential probe of reionisation topology (Majumdar et al., 2018; Hutter et al., 2019).

Other triangle configurations are also of interest, in particular squeezed isosceles triangles $k_1 \sim k_2 \gg k_3$ (Lewis, 2011). In the squeezed configuration, two wave modes are much larger than the remaining mode. This is essentially equivalent to the small scale power being modulated by larger scale structures. In this case a positive squeezed bispectrum indicates that large scale over densities will likely have more small scale structure, and vice versa for the negative case.

Other studied configurations include the folded triangle $k_1 = 2k_2 = 2k_3$, and the elongated triangle $k_1 = k_2 + k_3$, which are more sensitive to planar or flattened structures (Lewis, 2011). These modes become significant when the topology of the field deviates from spherical symmetry. This deviation results from either heating from X-ray emission (Watkinson et al., 2018), or from reionisation bubbles (Majumdar et al., 2018; Hutter et al., 2019).

5.3.2 21cm Bispectrum During EoR

Similar to Section 5.2.2, the bispectrum can be described in the saturation limit during the mid to late points of reionisation. In Shimabukuro et al. (2016) and Majumdar et al. (2018) they perform a similar decomposition of the bispectrum as seen in Equation 5.8 for the power spectrum:

$$\begin{aligned}
 B_{21}(k_1, k_2, k_3) = T_0^3 \langle x_{\text{HI}}(\mathbf{x}) \rangle^3 & \left[B_x(k_1, k_2, k_3) + B_\rho(k_1, k_2, k_3) \right. \\
 & + \text{cross - correlation terms} \\
 & \left. + \text{higher order terms} \right].
 \end{aligned}
 \tag{5.11}$$

This relationship is similar to the power spectrum expansion, and shows the potential of the bispectrum for characterising the ionisation field, and the matter distribution. Due to the large number of cross terms and higher order terms present in Equation 5.11, Shimabukuro et al. (2016) only consider the auto bispectrum terms. Shimabukuro et al. (2016) shows that during reionisation, the matter and neutral fraction auto bispectrum predominantly contribute to the total bispectrum. The neutral fraction during reionisation is dominant compared to the matter density on all scales. Majumdar et al. (2018) performs a similar analysis to Shimabukuro et al. (2016), and finds that the neutral fraction and the matter density dominate the total bispectrum at different stages of reionisation. Moreover, Majumdar et al. (2018) improves upon Shimabukuro et al. (2016) who only considers the absolute value of the bispectrum. Majumdar et al. (2018) show that the sign of the bispectrum could be an important indicator of reionisation, where a negative amplitude would be a direct result of ionised regions.

5.3.3 Skew Spectrum

There are several unique challenges associated with measuring the bispectrum in comparison to the power spectrum. Most notably, the bispectrum has lower signal to noise than the power spectrum, since the bispectrum sensitivity scales with the number of available triangle modes. Coupled with the challenge of determining which modes to study, and their resulting interpretation, the bispectrum has received less attention as a probe of the 21cm signal than the power spectrum. These challenges motivated the CMB cosmology community to develop the skew spectrum (Szapudi and Szalay, 1997; Munshi et al., 1998; Cooray, 2001), which integrates the bispectrum over all triangle modes for a fixed size, effectively collapsing the bispectrum into a single wave mode (Regan, 2017):

$$S_\gamma(\mathbf{k}) = \frac{\delta\bar{T}_b^3}{(2\pi)^3} \int_{\mathbb{R}} d^3\mathbf{q} B(\mathbf{k}, \mathbf{q}, \mathbf{k} + \mathbf{q}) [\text{mK}^3 \text{Mpc}^3]. \quad (5.12)$$

It can be shown that Equation 5.12 is equivalent to the cross spectrum of the quadratic temperature fluctuation field ($\delta_T^2(\mathbf{x})$), with the temperature fluctuation field. Starting from Equation 5.12 and expressing the bispectrum as the product of the Fourier transform of the temperature fluctuation field we can prove this relationship:

$$S_\gamma(\mathbf{k}) = \frac{\delta\bar{T}_b^3}{(2\pi)^3} \hat{\delta}_T(\mathbf{k}) \int_{\mathbb{R}^3} d^3\mathbf{q} \hat{\delta}_T(\mathbf{q}) \hat{\delta}_T(-\mathbf{k} - \mathbf{q}). \quad (5.13)$$

We define the Fourier transform as $\hat{\delta}_T(\mathbf{k}) = 1/V \int_{\mathbb{R}^3} d^3\mathbf{x} e^{-i\mathbf{k}\cdot\mathbf{x}} \delta_T(\mathbf{x})$, which we substitute into Equation 5.13:

$$S_\gamma(\mathbf{k}) = \frac{\delta\bar{T}_b^3}{(2\pi)^3} \hat{\delta}_T(\mathbf{k}) \int_{\mathbb{R}^3} d^3\mathbf{q} \frac{1}{V^2} \iint_{\mathbb{R}^6} d^3\mathbf{x} d^3\mathbf{x}' e^{-i\mathbf{q}\cdot\mathbf{x}} e^{i(\mathbf{k}+\mathbf{q})\cdot\mathbf{x}'} \delta_T(\mathbf{x}) \delta_T(\mathbf{x}'). \quad (5.14)$$

Substituting in the Dirac delta function:

$$S_\gamma(\mathbf{k}) = \delta\bar{T}_b^3 \hat{\delta}_T(\mathbf{k}) \frac{1}{V^2} \iint_{\mathbb{R}^6} d^3\mathbf{x} d^3\mathbf{x}' e^{i\mathbf{k}\cdot\mathbf{x}'} \delta_D(\mathbf{x}' - \mathbf{x}) \delta_T(\mathbf{x}) \delta_T(\mathbf{x}'). \quad (5.15)$$

Thus:

$$S_\gamma(\mathbf{k}) = \delta\bar{T}_b^3 \hat{\delta}_T(\mathbf{k}) \frac{1}{V^2} \int_{\mathbb{R}^3} d^3\mathbf{x} e^{i\mathbf{k}\cdot\mathbf{x}} \delta_T^2(\mathbf{x}). \quad (5.16)$$

Here we define $\hat{\delta}_{T^2}(\mathbf{k}) = \mathcal{F}\{\delta^2(\mathbf{x})\}$, where $\hat{\delta}_{T^2}^*(\mathbf{k}) = \mathcal{F}\{\delta_T^2(\mathbf{x})\}^*$, and again the super script * indicates the conjugate. Substituting this into Equation 5.17:

$$S_\gamma(\mathbf{k}) = \delta\bar{T}_b^3 \frac{\hat{\delta}_T(\mathbf{k})\hat{\delta}_{T^2}^*(\mathbf{k})}{V} [\text{mK}^3 \text{Mpc}^3]. \quad (5.17)$$

We equally define the isotropic skew spectrum $S_\gamma(k) = \langle S_\gamma(\mathbf{k}) \rangle_{k=|\mathbf{k}|}$. In this definition the skew spectrum units differ from the power spectrum by only a factor of mK. Similarly to Equation 5.4, we can also define the Skew spectrum as:

$$\langle \hat{\delta}_{T^2}(\mathbf{k})\hat{\delta}_T(\mathbf{k}') \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') S_\gamma(k) [\text{mK}^3 \text{Mpc}^3]. \quad (5.18)$$

The skew spectrum, like the power spectrum, also has a dimensionless definition:

$$\Delta_{T^2,T}^2(k) \equiv \frac{k^3}{2\pi^2} S_\gamma(k) [\text{mK}^3]. \quad (5.19)$$

The skew spectrum collapses all information about the non-Gaussianities present in the bispectrum as a function of a single wave-mode k . From this we can make direct comparisons with the power spectrum. However, in the process of performing this operation we lose some of the information contained in the three-point correlation function, in particular the bispectrum phase information. The bispectrum phase information has shown to be sensitive to cosmological structure (for example see Obreschkow et al., 2012). In Chapter 7 we investigate the skew spectrum of the 21cm signal, and its sensitivity to the ionisation, and X-ray heating effects during the EoH and the EoR. In Chapter 8 we develop a method to determine the instrumental skew spectrum. We use this to assess the foreground skew spectrum, and a fiducial 21cm skew spectrum for simulated MWA observations.

5.4 Interferometry and Power Spectrum

The connection between the three dimensional Fourier transform of the 21cm temperature brightness and visibilities is well known (Zaldarriaga et al., 2004; Morales and Hewitt, 2004; Bharadwaj and Saiyad Ali, 2005). In this section we outline the fundamental relationship between visibilities and the power spectrum. In Chapter 6 we go into more detail about the power spectrum estimation method used in this work.

The visibilities $\mathcal{V}(u, v, \nu)$ as measured by an interferometer vary as a function of frequency (ν). To relate visibility measurements to the Fourier transform of the temperature fluctuation field we perform a Fourier transform along the frequency domain:

$$\tilde{\mathcal{V}}(u, v, \eta) = \int_{\mathbb{R}} d\nu e^{-2\pi i u \eta} W(\nu) \mathcal{V}(u, v, \nu), \quad (5.20)$$

where η is the Fourier pair of ν and has units of seconds. In practice the Fourier operation is performed on discrete data through an FFT. It is standard practice to taper the discrete data as a function of frequency with a spectral window function $W(\nu)$. In this work we use a Blackman-Harris window (Harris, 1978). Spectral tapering results in a loss of sensitivity due to a reduced effective bandwidth. However, the upside is a suppression in leakage and aliasing as a function of η , which is important when considering the relative faintness of the 21cm signal to contaminants. We will discuss this further later in this section.

The Fourier modes (u, v, η) can be directly related to the cosmological units (k_x, k_y, k_z) through the following unit conversions (Morales and Hewitt, 2004):

$$k_x = \frac{2\pi u}{D_M(z)} h \text{ Mpc}^{-1}; \quad (5.21)$$

$$k_y = \frac{2\pi v}{D_M(z)} h \text{ Mpc}^{-1}; \quad (5.22)$$

$$k_{\parallel} = \frac{2\pi H_0 f_{21} E(z) \eta}{c(1+z)^2} h \text{ Mpc}^{-1}. \quad (5.23)$$

H_0 is the Hubble constant, f_{21} is the 21 cm frequency, z is the redshift, and $E(z)$ is the cosmological function given by $E(z) = \sqrt{\Omega_M(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda}$. $D_M(z)$

is the co-moving transverse distance, which is given by Hogg (1999):

$$D_M(z) = D_H \int_0^z \frac{dz'}{E(z')}, \quad (5.24)$$

and has units of h^{-1} Mpc. This transforms our signal into cosmological units.

5.4.1 Power Spectrum

From Equation 5.20 we see that $\tilde{\mathcal{V}}(\mathbf{k}) \propto \hat{\delta}_T(\mathbf{k})$. Thus, the isotropic power spectrum can then be estimated through a spherical average in k -space:

$$P(k) = C \langle \tilde{\mathcal{V}}(\mathbf{k}) \tilde{\mathcal{V}}^*(\mathbf{k}) \rangle_{k \in |\mathbf{k}|}. \quad (5.25)$$

Likewise the 2D cylindrically averaged power spectrum can be determined:

$$P(k_\perp, k_\parallel) = C \langle \tilde{\mathcal{V}}(k_\perp, k_\parallel) \tilde{\mathcal{V}}^*(k_\perp, k_\parallel) \rangle, \quad (5.26)$$

Where $k_\perp = \sqrt{k_x^2 + k_y^2}$, and $k_z = k_\parallel$ is the line of sight mode. In both Equations 5.25 and 5.26, C is the cosmological conversion factor which converts from units of $\text{Jy}^2 \text{Hz}^2$ to units of $\text{mK}^2 \text{Mpc}^3$. We go into more detail regarding this calculation in the context of radio interferometric measurements in Chapter 6.

The cylindrical PS is useful as a diagnostic tool, where the spectrally smooth foreground sources should be isolated in the lower line of sight k_\parallel modes, relative to the 21cm signal. This should create a relatively foreground free window from which to estimate the 21cm power spectrum.

5.4.2 Foreground Contamination and Wedge

One of the biggest challenges in measuring the expected cosmological 21cm signal during reionisation is intervening foreground sources. These sources are primarily from diffuse Galactic synchrotron emission, and extra-galactic radio sources, typically from AGN. Foreground emission from these sources is between four to five orders of




Figure 5.4: Example of the 2D power spectrum, showing the different regions affected by instrument chromaticity. Figure taken from Barry et al. (2016). As defined in the text, k_{\parallel} is derived from the Fourier transform along frequency, and k_{\perp} is the average of k -modes from Fourier transforms over (u, v) . This manifests in spectrally smooth sources towards the bottom, spectrally structured at the top, large spatial scales at left, and small spatial scales at the right.

magnitude brighter than the 21cm signal (see references in Furlanetto, 2016). The intervening astrophysical foregrounds however are expected to be spectrally smooth in contrast to the 21cm signal (Shaver et al., 1999). This allows for the separation in Fourier space between the signal and foregrounds.

Figure 5.4 shows a cartoon example of the 2D power spectrum for a typical MWA EoR observation (taken from Barry et al. (2016)). In the absence of instrumental effects, due to the smooth spectral characteristics of astrophysical foregrounds, we would expect the foreground signal to be isolated to the lowest line of sight modes (k_{\parallel}). However, two effects couple together to form a wedge like structure and result in foreground spectral leakage. The primary effect is the frequency dependent baseline distribution. From Equation 3.23 we see that the (u, v) values for each visibility depend on frequency, with longer baselines being more frequency dependent. This results in a radial migration of visibilities as a function of frequency and effectively imprints the instruments PSF onto the power spectrum (Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012). Additionally, the primary beam through convolution couples the visibility to a local region in the uv -plane. Both of these effects in tandem with each other result in frequency dependent mode-mixing, which leak foreground power throughout the 2D power spectrum. This renders some modes inoperable, and leaves only a small EoR window from which to estimate the 1D power spectrum.

5.4.2.1 Astrophysical Foregrounds

Sources of foreground contamination come from two primary sources, extra-galactic radio sources, and Galactic radio sources. Galactic foreground contamination can be further split into a diffuse Galactic continuum synchrotron emission component, and more compact extended sources such as supernova remnants and HII regions. Extra-galactic foregrounds with the exception of some large and bright extended sources (Fornax A and Centaurus A for example) are typically point sources as seen by radio interferometers. The Galactic foreground contamination on the other hand is far more extended and affects the shorter baselines more than the extended baselines. This

asymmetry in baseline sensitivity complicates calibration (Byrne et al., 2021). As a result, MWA EoR observing strategies seek to mitigate the effects of the Galactic foregrounds by observing away from the Galactic plane, in relatively quiescent fields dominated by extra-galactic point sources (Trott et al., 2020).

The MWA however has a large field of view (~ 20 deg) (Tingay et al., 2013), which results in numerous sources, of both extra-galactic and Galactic origin being present in the primary beam sidelobes. Pober et al. (2016a) demonstrate that including point sources in the sidelobes when subtracting a foreground model reduces power in the foreground wedge, and result in less spectral leakage into the EoR window. However, numerous extended foreground sources present in the sidelobes of some MWA observations are ignored. These include supernova remnants and Centaurus A. We discuss the impact of these widefield extended sources on the MWA power spectrum in Chapter 6.

5.4.2.2 Foreground Mitigation

Numerous methods have been developed that are designed to mitigate foreground contamination, either through avoidance, subtraction, or suppression (for a review see Chapman, 2017). Both Galactic and extra-galactic foregrounds have well known smoothly varying spectra (e.g Jelić et al., 2008). This contrasts with the 21cm signal which is expected to vary rapidly with frequency. This difference allows for the separation of the foreground signal, and the 21cm signal, in both Fourier space, and in frequency space. Methods to remove the spectrally smooth foregrounds, involve creating parametric models (e.g Morales et al., 2006; Jelić et al., 2008; Liu et al., 2009; Bowman et al., 2009; Datta et al., 2010), and non-parametric models (e.g Harker et al., 2009; Chapman et al., 2012a,b). However, systematic errors in the foreground models, either from poorly understood instrumental effects, or from modelling techniques, introduce bias across all spatial scales (Chapman, 2017). Avoidance techniques seek to measure the 21cm power spectrum from the relatively clean modes in the EoR window (Datta et al., 2010; Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012).




Figure 5.5: Upper limits measured from ongoing and past radio interferometric EoR experiments (references therein). These are compared to an expected fiducial 21cm model (solid black), as well as to three other models faint galaxies (dashed line), bright galaxies (dash dot line) and noheating. This plot was taken from Shimabukuro et al. (2022).

However, instrumental leakage as a function of frequency and instrumental polarisation leakage (both from assuming an idealised beam model), couple to the foreground power, contaminating the EoR window. Moreover, missing frequency channels result in the foreground wedge aliasing through higher k_{\parallel} modes, further contaminating the EoR window. Spectral tapering can reduce the impact of spectral leakage, but this alone is not enough to retrieve the signal. Suppression seeks to down weight the Fourier modes which are most affected by foreground power (e.g Trott et al., 2016). In reality some combination of all of these methods is likely required to detect the signal.

Other techniques try to take advantage of the statistical independence of the 21cm signal relative to the instrument noise, calibration errors, and foregrounds, using Gaussian process regression (Mertens et al., 2018). These techniques are imperfect, and thus some foreground residuals will remain in the data. These residuals couple with the instrumental chromaticity and produce a wedge like structure in the 2D power spectrum (Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012).

5.5 21cm Power Spectrum Experiments

21cm power spectrum experiments are performed using low frequency radio interferometers such as the MWA. Broadly there are two approaches; the image based approach, and the delay based approach. In both cases measurements of the Fourier components of the sky brightness distribution are made at different frequencies. The distribution of these samples as measured by the instrument is frequency dependent. The image based approach grids the measured data onto a rectilinear grid, as a function of frequency, and then applies a Fourier transform as a function of frequency to derive the full Fourier representation of the sky/field (see Chapter 6 for more detail). In the delay approach the frequency dependence of the Fourier component sampling is ignored. This approach is valid for small baselines $\lesssim 200$ m (for example Thyagarajan et al., 2015). A Fourier transform is then performed as a function of frequency for each individual baseline to get the delay space representation of the Fourier sky cube.

In this section we briefly outline the ongoing power spectrum experiments and their current results for three of the principal instruments in the field. Figure 5.5 we show the current upper limits of the ongoing and past experiments, as compared to a fiducial 21cm model.

5.5.1 MWA

The MWA was discussed in Section 3.2; here we discuss the results of the MWA EoR experiments. From 110 hours of observations Trott et al. (2020) find an upper limit of $\Delta_{21}^2 \leq (43)^2 \text{ mK}^2$ at $k = 0.14 h \text{ Mpc}^{-1}$, at redshift $z = 6.5$. Barry et al. (2019) additionally finds an upper limit of $\Delta_{21}^2 \leq (62.4)^2 \text{ mK}^2$ at $k = 0.2 h \text{ Mpc}^{-1}$, at redshift $z = 7$ from 21 hours of data.

5.5.2 LOFAR

The LOw Frequency ARay (LOFAR; van Haarlem, M. P. et al., 2013) is an aperture array with stations located primary in the Netherlands, and across large separations

in Europe. There are two types of antennas in LOFAR; high band antennas with a frequency range of 120 – 190 MHz, and low band antennas with a frequency range of 30 – 90 MHz. The high band station tiles are used for EoR science. From 141 hours of LOFAR data Mertens et al. (2020) finds an upper limit of $\Delta_{21}^2 \leq (73)^2 \text{ mK}^2$ at $k = 0.075 \text{ h Mpc}^{-1}$, at redshift $z = 9.1$.

5.5.3 HERA

The Hydrogen Epoch of Reionisation Array (HERA; DeBoer et al., 2017) is a low frequency aperture array, located in the Karoo Desert in South Africa. HERA is configured in a redundant array layout similar to the compact core of the MWA phase II array (Wayth et al., 2018). Constructed from 14 meter zenith pointed mesh dishes, HERA observes in a drift scan mode. The full HERA instrument will contain 350 elements, with an observing frequency range from 50 – 250 MHz, designed with the goal of probing the 21cm emission in the redshift range $z = 6 - 12$. From ~ 36 hours of observations, Abdurashidova et al. (2022) report upper limits of $\Delta_{21}^2 \leq (30.76)^2 \text{ mK}^2$ at $k = 0.192 \text{ h Mpc}^{-1}$, at redshift $z = 7.9$, and $\Delta_{21}^2 \leq (95.74)^2 \text{ mK}^2$ at $k = 0.256 \text{ h Mpc}^{-1}$, at redshift $z = 10.4$. These upper limits were performed with HERA Phase I configuration.

5.6 Summary

In this chapter we have developed the power spectrum, and the skew spectrum for a cosmological field like the 21cm signal. We discussed what astrophysical information the power spectrum and the bispectrum of the 21cm signal is sensitive to. We then derived the skew spectrum in context of the bispectrum. Finally, we conclude with the current results of 21cm power spectrum experiments with low frequency interferometers. This chapter establishes the foundational context for the projects we present in Chapters 6, 7 and 8. In particular the skew spectrum and the non-Gaussianity it is sensitive to is the fundamental question for the works in Chapters 7 and 8.

CHAPTER 6

INVESTIGATING THE CONTRIBUTION OF EXTENDED RADIO SOURCES TO THE EPOCH OF REIONISATION POWER SPECTRUM

Adapted from:

Jaiden H. Cook, Cathryn M. Trott, and Jack L. B. Line (2022), "Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum", *Monthly Notices of the Royal Astronomical Society*, Volume 514, Issue 1, pp. 790-805, DOI: 10.1093/mnras/stac1330

This chapter is my own work, I created the method including the OSIRIS simulation pipeline, and the Gaussian fitting code. Furthermore, I performed all the simulations and data analysis in this work, with the input of my supervisors C. M. Trott and J. L. B. Line. The idea for this project came from C. M. Trott, following discussions with myself regarding extended sources in the MWA primary beam. Additionally, the experience and support of C. M. Trott, was fundamental in creating the gridding and power spectrum components of the OSIRIS pipeline. The visibility and gridding simulation component of OSIRIS, was greatly influenced by previous work performed by J. L. B. Line. The support and guidance of J. L. B. Line was instrumental in the development of OSIRIS. I was the primary author on the paper, I drafted the manuscript, which was distributed to C. M. Trott and J. L. B. Line, who provided proof reading, and feedback until the final version was produced.

Abstract

We investigate the contribution of extended radio sources such as Centaurus A, and Galactic supernova remnants (SNRs) to our ability to detect the statistical 21 cm signal from the Epoch of Reionisation (EoR) with the Murchison Widefield Array (MWA). These sources are typically ignored because they are in highly attenuated parts of the MWA primary beam, however in aggregate these sources have apparent flux densities of 10 Jy on angular scales we expect to detect the 21 cm signal. We create bespoke multi-component 2D Gaussian models for Galactic SNRs and for Centaurus A, and simulate the visibilities for two MWA snapshot observations. We grid those visibilities and then Fourier transform them with respect to frequency, averaging them both spherically and cylindrically to produce the 1D and 2D power spectra. We compare the simulated 1D power spectra to the expected 21 cm power spectrum. We find that although these extended sources are in highly attenuated parts of the MWA primary beam pattern, collectively they have enough power ($\sim 10^4 - 10^5 \text{ mK}^2 h^{-3} \text{ Mpc}^3$) on EoR

significant modes ($|\mathbf{k}| \lesssim 0.1 h \text{ Mpc}^{-1}$) to prohibit detection of the 21 cm signal ($\sim 10^4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$). We find that 50 – 90% of sources must be removed in order to reduce leakage to a level of $\sim 10 - 20\%$ of the 21 cm power spectrum on EoR significant modes. The effects of widefield extended sources will have implications on the detectability of the 21 cm signal for the MWA and with the future Square Kilometre Array (SKA).

6.1 Introduction

Radio observations of the redshifted 21 cm neutral hydrogen emission line have the capability to reveal underlying astrophysical formation mechanisms during the cosmic dawn, and the Epoch of Reionisation (EoR) (Furlanetto et al., 2006b). The EoR is the period of cosmic time where the predominantly neutral hydrogen inter-galactic medium (IGM), transitioned to a fully ionised state after the formation of the first stars, galaxies, and black holes. Observations of quasars (Fan et al., 2006) and the anisotropies in the Cosmic Microwave Background through the Sunyaev–Zel’dovich effect (Mesinger et al., 2012), have constrained the EoR to a redshift range of $5.4 \lesssim z \lesssim 10$. The cosmological nature of the 21cm emission line allows for the direct observation of the full reionisation history. The future Square Kilometre Array (SKA) promises to directly image the redshifted 21 cm signal during the EoR (Koopmans et al., 2015).

The current generation of low frequency radio instruments lack the sensitivity to directly image the 21 cm signal, and are thus focused on estimating the 21 cm statistics as a function of spatial scale by calculating the 21 cm power spectrum. The 21 cm statistics have the potential to differentiate between different reionisation scenarios, and therefore provide an insight into the underlying astrophysical reionisation mechanisms (see Furlanetto et al., 2006b; Morales and Wyithe, 2010; Pritchard and Loeb, 2012; Furlanetto, 2016, for comprehensive reviews). The current generation of radio instruments includes the Murchison Widefield Array (MWA, Bowman et al., 2013; Tingay et al., 2013; Wayth et al., 2018); Low-Frequency Array (LOFAR, van Haarlem, M. P. et al., 2013); the Precision Array for Probing the Epoch of Reionization (PA-

PER, Parsons et al., 2010); Hydrogen Epoch of Reionization Array (HERA, DeBoer et al., 2017); The Amsterdam–ASTRON Radio Transients Facility and Analysis Center (AARTFAARC, Prasad et al., 2016); the New extension in Nancay upgrading LOFAR (NenuFAR, Zarka et al., 2012). The MWA is a 256 element interferometer, with 128 operational at any one time in a compact or extended configuration (Wayth et al., 2018). Measuring the statistical 21 cm signal from the EoR is one of the main science goals of the MWA (Bowman et al., 2013).

Foreground Galactic and extragalactic radio sources at redshifted 21 cm frequencies pose a fundamental problem for detecting the 21 cm signal during the EoR. These foreground sources can be $10^4 - 10^5$ times brighter than the underlying cosmological 21 cm signal (Furlanetto et al., 2006b). The frequency structure of the 21 cm signal varies rapidly with frequency when compared to foreground emission (Shaver et al., 1999). Foreground emission is proportional to a power law distribution, and varies relatively smoothly over frequency. Therefore foreground power is expected to be primarily isolated to low line of sight k Fourier modes compared to the 21 cm EoR signal (Morales and Hewitt, 2004; Bowman et al., 2009). However instrumental chromaticity imparts highly varying spectral structure which leaks power into prospective EoR modes through a process known as mode mixing (Bowman et al., 2009; Datta et al., 2010). One way to avoid some of these effects is through the 2D power spectrum, which separates the power spectrum modes into line of sight modes k_{\parallel} and perpendicular angular modes k_{\perp} in units of Mpc^{-1} (Morales et al., 2006; Datta et al., 2010). Radio interferometers sparsely sample the uv plane (which is proportionate to k_{\perp}), however baseline length is wavelength dependent and so introduces frequency structure into the foreground emission. As a result of this structure, foreground emission leaks into higher k_{\parallel} modes as a function of k_{\perp} (Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012), resulting in a wedge-shaped foreground-dominated area.

Most of the foreground power is contained in the wedge, leaving a relatively clean ‘EoR window’ (Vedantham et al., 2012). However, calibration errors and primary beam chromaticity can cause leakage from the foreground wedge into the EoR window

(Morales et al., 2012; Trott et al., 2012; Barry et al., 2016). This problem is compounded for sources further from the centre of the field, as the primary beam changes more with frequency the further away from the point of maximum sensitivity. Pober et al. (2016a) analysed the effects of including source subtraction from the sidelobes of the MWA primary beam when calculating the 2D power spectrum. They found that sources further from the centre of the field leaked more power from the foreground wedge into the window. The MWA primary beam spectral structure for different EoR fields is shown in Figures 27 and 28 in Trott et al. (2020). At the edges of the sidelobes, and away from the main lobe, the MWA primary beam spectral index is steep, introducing rapidly changing spectral structure to sources in these locations. Furthermore, Pober et al. (2016a) found that including these extragalactic sources located in the beam sidelobes during foreground removal reduced the power in the EoR window by a few percent.

Poher et al. (2016a) was only concerned with point sources in the sidelobes, however, in the EoR 2¹ field there are several exceptionally bright extended sources, which due to their low apparent surface brightness are generally not included in MWA EoR processing pipelines. Primarily this field contains Centaurus A (CenA), which is the brightest radio galaxy in the southern hemisphere spanning 4×8 deg with a brightness of ~ 4000 Jy at 183 MHz (Alvarez et al., 2000; McKinley et al., 2013). CenA is often present or at the edge of one of the MWA primary beam sidelobes for EoR 2 field pointings. As a result CenA is highly attenuated, but has an apparent brightness on the order of 10 Jy. Additionally, the complex spectral structure of the MWA primary beam at the sidelobes imprints frequency structure that can lead to leakage in the EoR window. Leakage at this apparent brightness can still be orders of magnitude brighter than the expected 21 cm signal.

In addition to CenA the Galactic plane also appears in one or several of the MWA primary beam sidelobes. The Galactic plane is populated by a large number of bright supernova remnants (SNRs) as well as large scale diffuse radio emission. SNRs

¹EoR 2 field coordinates: RA=10.3 h, Dec=-10°

themselves have flux densities that range from 1 – 1000 Jy, and have angular extents that are similar in scale to the expected 21 cm reionisation bubbles (Wyithe and Loeb, 2004; Furlanetto and Oh, 2005). Likewise, these sources are in complex parts of the MWA primary beam, which can cause leakage from the foreground wedge into the EoR window. Further complications occur as these extended sources rotate through the MWA primary beam, imparting varying spectral structure in the process. Their extended nature also means the spectral structure imparted by the beam changes across the source, and can vary significantly depending on the location of the source within the primary beam.

The effect of these attenuated but complex sources at the field edge has not been established for 21 cm EoR science. To investigate the amount of leakage caused by these sources in the EoR window, in this work we create a sky-model which contains morphological models of CenA and Galactic plane SNRs. The modelling of the morphological models for Galactic SNRs and CenA is described in Section 6.3. We then run various sky-models through a simulation pipeline (described in Section 6.2) which calculates the 1D and 2D power spectrum with a fiducial 21 cm signal (via 21CMFAST Mesinger et al., 2011). We then look at how much of the sky-model needs to be subtracted to retrieve the 21 cm signal (Section 6.4). In this work we perform all cosmological calculations with the Planck Collaboration et al. (2020) cosmology, where $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

6.2 Methodology

To test the leakage of Galactic Plane SNRs and CenA into the EoR window, we developed a method which simulates the contribution of extended radio sources to the visibilities measured by the MWA. Briefly, we describe the steps of the method here, going into more detail in the subsequent subsections. The first step generates a sky-model image cube $I(\mathbf{l}, \nu)$ as a function of frequency. These sky-model cubes are constructed from multi-component 2D Gaussian models of CenA and Galactic plane SNRs; for details on how the sky-model cubes and the 2D Gaussian model components were cre-

ated, refer to Section 6.3. The sky-model cube is Fast Fourier Transformed (FFT) into the Fourier sky-cube $\tilde{I}(\mathbf{u}, \nu)$. The visibilities $\mathcal{V}(\mathbf{u}, \nu)$ are simulated by sampling the Fourier sky-cube using the MWA (u, v) distribution. The sampling process incorporates the FFT of the MWA primary beam, effectively simulating MWA measurements. The sampled visibilities are then gridded onto the uv -plane reconstructing the Fourier sky-cube which is denoted by $\tilde{\tilde{I}}(\mathbf{u}, \nu)$. An FFT is then performed with respect to the frequency axis to retrieve the reconstructed Fourier sky-cube $\tilde{\tilde{I}}(\mathbf{u}, \eta)$ as a function of the line of sight mode η . $\tilde{\tilde{I}}(\mathbf{u}, \eta)$ is then averaged both spherically and cylindrically to calculate the 1D and 2D power spectra respectively.

For comparison a fiducial simulated 21 cm signal is added to a noise-only reconstructed Fourier sky-cube $\tilde{\tilde{I}}_{\mathcal{N}}(\mathbf{u}, \eta)$. This is then spherically and cylindrically averaged to calculate the 1D and 2D noise plus 21 cm signal 1D and 2D power spectra. We then compare the 21 cm signal power spectra to the simulated widefield extended power spectra to determine the significance of leakage at EoR k -modes of interest. The fiducial 21 cm signal was generated using 21CMFAST power spectrum simulations, and is taken from Mesinger et al. (2011).

To simulate MWA observations we created a simulation pipeline called Observational Supernova-remnant Instrumental Reionisation Investigative Simulator (OSIRIS)². The core interferometric simulation functions are based on the MAJICK software package (Line, 2017). The general process of the OSIRIS pipeline is described by the flow chart in Figure 6.1.

6.2.1 Fourier Sky Cube

Radio interferometers measure a complex coherence term known as the visibility $\mathcal{V}(\mathbf{u})$, which is the cross-correlation between two antenna elements. The visibility is described by the measurement equation, which relates the sky-brightness distribution $I(\mathbf{l})$ to the visibility as a function of \mathbf{u} (Taylor et al., 1999):

²<https://github.com/JaidenCook/OSIRIS>

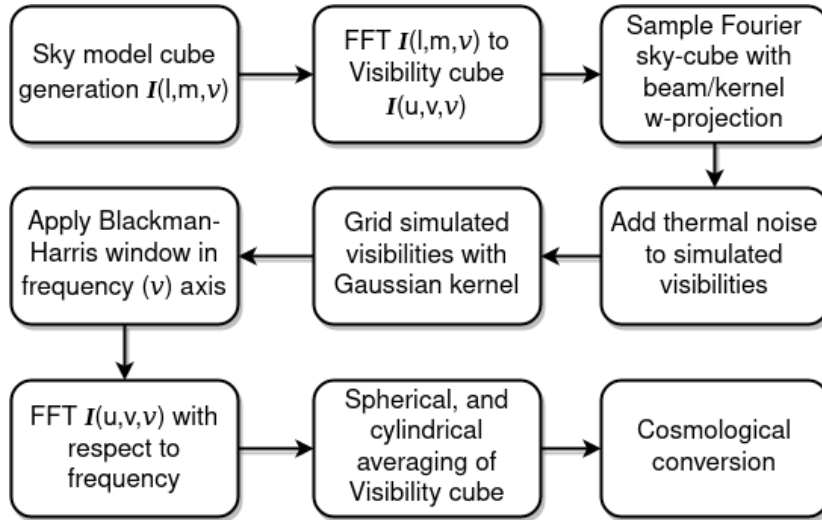


Figure 6.1: Simulation pipeline flow chart. Shows all the different steps from the sky-model generation to the power spectrum calculation and cosmological conversion.

$$\mathcal{V}(\mathbf{u}, w) = \int_{-\infty}^{\infty} \frac{B(\mathbf{l})I(\mathbf{l})}{n} e^{-2\pi i(w(n-1))} e^{-2\pi i(\mathbf{u}\cdot\mathbf{l})} d\mathbf{l} \quad (6.1)$$

$B(\mathbf{l})$ is the primary beam as projected onto the celestial sphere, and n is the direction cosine along the phase centre, which is defined by $n = \sqrt{1 - |\mathbf{l}|^2}$. The vector \mathbf{u} represents the physical displacement of the tiles on a plane in units of wavelengths and is represented by the coordinates (u, v) ; the vector \mathbf{l} contains the direction cosines (l, m) which are defined on the image plane. The w -terms encapsulate the curvature of the sky, and are significant because the MWA has a large field of view (> 10 deg). OSIRIS takes an input sky-model cube $I(\mathbf{l}, \nu)$, which is then Fourier transformed with respect to \mathbf{l} via an FFT. The resulting FFT produces the Fourier sky cube $\tilde{I}(\mathbf{u}, \nu)$, mapped to a regular (u, v) grid, defined by the extent of the input image (l, m) grid.

6.2.2 Simulating Visibilities

Simulating the instrumentally-measured visibilities is performed by sampling the uv -plane with a kernel that incorporates the MWA primary beam response $B(\mathbf{l}, \nu)$ and the curvature of the sky through the w -kernel³. This process samples the Fourier space for each baseline as a function of frequency. The baseline coordinates (u, v, w) for

³The w -kernel is defined as $e^{-2\pi i(w(n-1))}$.

each frequency slice are determined using the MWA Phase I array layout, and a set of MAJICK functions. These functions use the array (east, north, height) and pointing centre to determine the baseline coordinates in meters (u, v, w) , which are converted into units of wavelength for each frequency channel. In this work we use a zenith pointed array, since we consider extended radio models of the entire sky. The sampling kernel $\tilde{K}(\mathbf{u}, w_i, \nu)$ for a given baseline at a particular frequency, is the convolution of the FFT of the MWA primary beam, and the FFT of the w -kernel:

$$\tilde{K}(\mathbf{u}, w, \nu) = \int_{-\infty}^{\infty} B(\mathbf{l}, \nu) e^{-2\pi i(w(n-1))} e^{-2\pi i(\mathbf{u}\cdot\mathbf{l})} d\mathbf{l}. \quad (6.2)$$

The curvature term w is lost in the 2D FFT from image space (l, m) to Fourier space (u, v) . The w -kernel reincorporates the curvature of the sky through a process called w -projection (see Cornwell et al. (2008) for further details). Each baseline has a different w -term, and as such there is a unique sampling kernel for every baseline.

The MWA primary beam $B(\mathbf{l}, \nu)$ is generated using the Fully Element Embedded (FEE) model described by Sokolowski et al. (2017). The FEE beam model only has a frequency resolution of 1.28 MHz, however the channel resolution of the Fourier sky-cube is $\Delta\nu_f = 80$ kHz, thus the FEE beam model requires interpolation as a function of frequency. Without interpolation, the coarse band structure of the beam will be present in the EoR window when we Fourier transform with respect to frequency. Before the OSIRIS pipeline performs the interpolation, the FEE beam model is generated for 36 coarse channels (bandwidth 1.28 MHz) spanning the frequency range 147.2 – 193.3 MHz. The resulting beam cube is interpolated as a function of frequency using cubic splines. The observations simulated in this work have a bandwidth of $\Delta\nu = 15.36$ MHz, therefore each simulated observation has 192 channels. A primary beam model is generated for each channel using the interpolated FEE beam model.

Using the frequency interpolated FEE beam model, and the w_i term for the i th baseline, the OSIRIS pipeline generates a unique sampling kernel for that baseline. The simulated MWA visibility for the i th baseline \mathbf{u}_i , is determined by taking the sampling kernel weighted average of the $I(\mathbf{u}_j, \nu)$ grid pixels (the subscript j denotes the pixel

index) centred at \mathbf{u}'_i :

$$\mathcal{V}(\mathbf{u}_i, w_i, \nu) = \frac{\sum_{j=0}^N \tilde{K}(\mathbf{u}_j - \mathbf{u}'_i, w_i, \nu) \tilde{I}(\mathbf{u}_j, \nu)}{\sum_{j=0}^N \tilde{K}(\mathbf{u}_j - \mathbf{u}'_i, w_i, \nu)}. \quad (6.3)$$

The sampling kernel $\tilde{K}(\mathbf{u}_j - \mathbf{u}'_i, w_i)$ determines the weight for the \mathbf{u}_j j th grid point. For each frequency channel there are 8128 baselines. Each baseline has a complex conjugate pair where $\mathcal{V}(\mathbf{u}) = \mathcal{V}^\dagger(-\mathbf{u})$, for a total of 16256 simulated visibilities for each frequency. To minimise computation we use a coarse kernel size of (91×91) pixels, where each pixel has size 0.5λ . The number of operations per baseline is proportional to N^2 , however the accuracy of the sampling kernel is asymptotic. This is a reasonable trade-off in accuracy for computational efficiency.

Once the visibilities have been sampled, Gaussian thermal noise is added using the radiometer equation for a single baseline (see the appendix section 6.B). The noise level for a given baseline is determined by the observing frequency, the channel width ($\Delta\nu_f$) and the observation time length Δt . In this work Δt was used to control the noise level; we set $\Delta t = 10^4$ hours to ensure that the 21 cm signal could be detected in a single snapshot observation. This allows for a quantitative analysis of our ability to detect the 21 cm EoR signal in the presence of the intervening extended foreground objects. A more realistic approach would be to simulate the full 10^4 hours of observations incorporating rotation synthesis. This is however computationally expensive, and this level of complexity is not required to answer the underlying question in this paper. We will further discuss rotation synthesis in Sections 6.4.1 and 6.5.1.3.

6.2.3 Gridding

Gridding is the process by which the Fourier sky-cube is reconstructed from the visibilities; this is the first step in calculating the power spectrum. Gridding reconstructs the Fourier sky-cube as a function of frequency, by distributing the frequency dependent measured visibilities onto the (u, v) plane via a smooth gridding kernel. This is im-

portant because the contributions to a single visibility come from a region of the (u, v) space. Each grid point \mathbf{u}_j is the weighted average of all the baselines $\mathcal{V}(\mathbf{u}_i)$ multiplied by some weight $W(\mathbf{u}_j - \mathbf{u}_i)$ determined at the j th grid point via

$$\tilde{\mathcal{I}}(\mathbf{u}_j, \nu) = \frac{\sum_{i=0}^{N_{bl}} W(\mathbf{u}_j - \mathbf{u}_i) \mathcal{V}(\mathbf{u}_i, \nu)}{\sum_{i=0}^{N_{bl}} W(\mathbf{u}_j - \mathbf{u}_i)}. \quad (6.4)$$

The weights are determined by a smooth tapered gridding kernel function. In this work we use a Gaussian kernel defined as

$$W(\mathbf{u}_j - \mathbf{u}_i) = \frac{1}{2\pi\sigma^2} \exp\left\{-\frac{|\mathbf{u}_j - \mathbf{u}_i|^2}{2\sigma^2}\right\}. \quad (6.5)$$

The Gaussian kernel has a width of $\sigma = 2\lambda$, and a kernel window pixel size of (91×91) , where each pixel has size 0.5λ . Smooth tapered gridding kernels help to reduce leakage into the Fourier k -modes ($|k| > 0.1 h \text{ Mpc}^{-1}$) of interest for detecting the 21 cm EoR signal. Once the Fourier sky-cube has been reconstructed via the gridding process, we perform an FFT with respect to frequency to produce the reconstructed Fourier sky-cube as a function η :

$$\tilde{\mathcal{I}}(\mathbf{u}, \eta) = \int_{-\infty}^{\infty} \tilde{\mathcal{I}}(\mathbf{u}, \nu) e^{-2\pi i(\nu \cdot \eta)} d\nu \text{ Jy Hz}. \quad (6.6)$$

Prior to the FFT we spectrally taper the reconstructed Fourier sky-cube with a Blackman-Harris window. This tapering reduces spectral leakage introduced by aliasing from the bandwidth limited FFT in the frequency axis. Aliasing introduces a sinc function which spreads power from foreground wedge modes into higher $k_{||}$ parallel modes in the EoR window.

6.2.4 Calculating the 1D and 2D Power Spectra

The power spectrum provides information on how Gaussian the perturbations in the 21 cm brightness temperature are as a function of the spatial k -modes, which have units of ($h \text{ Mpc}^{-1}$) (Morales and Hewitt, 2004; Furlanetto et al., 2006b), and is the main output product of MWA EoR science (Bowman et al., 2013). The k -modes can be converted from the Fourier modes (u, v, η) into the components (k_x, k_y, k_{\parallel}). These conversions are outlined in Morales and Hewitt (2004), and are performed using Equations 6.20 outlined in the appendix. The power spectrum as a function of the k -modes is determined by averaging the product of $\tilde{\mathcal{I}}(\mathbf{k})$ and its conjugate $\tilde{\mathcal{I}}^\dagger(\mathbf{k})$ (denoted by the \dagger):

$$P(\mathbf{k}) = \delta_D(\mathbf{k} - \mathbf{k}') \frac{1}{\Omega_V} \langle \tilde{\mathcal{I}}^\dagger(\mathbf{k}) \tilde{\mathcal{I}}(\mathbf{k}) \rangle, \quad (6.7)$$

where Ω_V is the solid angle of the field of view; the Dirac delta (δ_D) and angular brackets represent the ensemble average over the field. Equation 6.7 is equivalent to the three dimensional Fourier transform of the two point correlation function. Due to the effective isotropy of the 21 cm signal (Furlanetto et al., 2006b), the power spectrum represents the variance of a random Gaussian field as a function of k -mode. For the 1D spherically averaged power spectrum we average spherical shells:

$$P(\mathbf{k}) = \frac{\sum_{i \in |\mathbf{k}|} \tilde{\mathcal{I}}^\dagger(\mathbf{k}_i) \tilde{\mathcal{I}}(\mathbf{k}_i) \tilde{W}(\mathbf{k}_i)}{\sum_{i \in |\mathbf{k}|} \tilde{W}(\mathbf{k}_i)} \text{Jy}^2 \text{Hz}^2 \quad (6.8)$$

where $|\mathbf{k}| = \sqrt{k_x^2 + k_y^2 + k_{\parallel}^2}$. The 2D cylindrically averaged power spectrum instead averages rings of $k_{\perp} = \sqrt{k_x^2 + k_y^2}$ as a function of k_{\parallel} :

$$P(k_{\perp}, k_{\parallel}) = \frac{\sum_{i \in k_{\perp}} \tilde{\mathcal{I}}^\dagger(\mathbf{k}_i) \tilde{\mathcal{I}}(\mathbf{k}_i) \tilde{W}(\mathbf{k}_i)}{\sum_{i \in k_{\perp}} \tilde{W}(\mathbf{k}_i)} \text{Jy}^2 \text{Hz}^2. \quad (6.9)$$

Throughout the gridding process, the accumulated Gaussian weights for each \mathbf{u}_j

grid point were stored in a weights array $W(\mathbf{u})$. The new Fourier weights $\tilde{W}(\mathbf{k})$ are the frequency average of the accumulated Gaussian weights $W(\mathbf{u})$.

6.2.5 The Fiducial 21cm Signal

For comparison with the SNR and CenA sky-model power spectra, we create a noise only reconstructed Fourier sky-cube $\tilde{\mathcal{I}}_{\mathcal{N}}(\mathbf{k})$ with an added fiducial simulated 21 cm signal. Using the radiometer equation (Equation 6.29 in the appendix), we generate random noise for the real and imaginary components for each visibility as a function of frequency. These visibilities are then gridded and Fourier transformed to create the noise-only reconstructed Fourier sky-cube. Since the power spectrum is a measure of the variance of the underlying visibility distributions at different k -modes (Section 6.2.4), we use simulated models of the 21 cm power spectrum to generate random Gaussian fields as a function of $|\mathbf{k}|$. These random Gaussian fields can then be added to $\tilde{\mathcal{I}}_{\mathcal{N}}(\mathbf{k})$, approximating a full 21 cm simulation without foregrounds. However, to properly simulate the signal we might detect with the MWA, a more accurate method would be to use a simulated 21 cm image cube as input into the pipeline. This would capture any potential signal loss as a result of the pipeline.

In this work we use a fiducial 21 cm power spectrum model created by Mesinger et al. (2011) using the software simulation package 21CMFAST. 21CMFAST is a semi-numerical modelling package which uses astrophysical approximations to efficiently simulate the cosmological 21 cm signal. The generated 21 cm power spectrum from 21CMFAST has been shown to be accurate to within $\sim 10\%$ of more complex hydrodynamical numerical simulations (Trac et al., 2008) on spatial scales of ≥ 1 Mpc.

The fiducial 21 cm 1D power spectrum we use in this work is calculated at a redshift of $z = 7.171$ which is approximately the redshift at the centre of the simulation observing band for the EoR 2 field ($\nu = 183$ MHz). The fiducial 21 cm power spectrum is then interpolated as a function of $|\mathbf{k}|$. The interpolated power spectrum is then converted from units of mK^2 to units of $\text{Jy}^2 \text{Hz}^2$:




Figure 6.2: Comparison of the sampled fiducial signal (solid black line), and the reconstructed EoR window fiducial signal from a spherically averaged 1D power spectrum (dashed red line). Both lines are in good agreement except at low k modes. There are less samples in these modes.

$$\sigma^2(\mathbf{k}) = \frac{2\pi^2}{k^3} \frac{\Delta^2(\mathbf{k})}{C} \text{Jy}^2 \text{Hz}^2 \quad (6.10)$$

$\Delta(\mathbf{k})$ is the power spectrum which has not been volume normalised. C is a cosmological unit conversion factor which converts the power spectrum from cosmological units to $\text{Jy}^2 \text{Hz}^2$ (given by Equation 6.28 in the appendix). Using Equation 6.10 and the interpolated 21 cm power spectrum we calculate a $\sigma(\mathbf{k})$ cube for each k -mode, using the k -mode grid corresponding to the simulated visibilities. These sigma values are then used to sample a random normal distribution for both the real and imaginary components of the complex visibility. The resulting random Gaussian complex cube is our 21 cm Fourier sky-cube as a function of k -modes which can be added to $\tilde{I}_{\mathcal{N}}(\mathbf{k})$.

To test whether the noise plus random Gaussian 21 cm Fourier sky-cube with the gridded Gaussian weights generates the expected power spectrum, we calculate the spherically averaged 1D power spectrum. Figure 6.2 shows the fiducial 1D power spectrum signal in black, and the expected 21 cm signal in the dashed red line. Only at the lowest k -modes do we not fully retrieve the expected signal, due to the relatively poor sampling at the shortest (< 100 baselines below $k \sim 0.01 h \text{Mpc}^{-1}$) baselines.

6.3 Data & Morphological Models

Extended radio sources such as CenA typically have angular sizes on the order of $\sim 1^\circ$ or larger. Most extended radio source modelling tools such as PyBDSF (Mohan and Rafferty, 2015), primarily use generalised 2D Gaussian functions to fit source flux density at different angular scales. 2D Gaussian functions have great utility because they have analytical Fourier transforms, and require less components than Dirac delta models, which essentially model each pixel as an independent radio source. In this work we similarly use generalised 2D Gaussians defined below:

$$G(x, y) = G_0 e^{-\left(a(x-x_0)^2 + 2b(x-x_0)(y-y_0) + c(y-y_0)^2\right)} \quad (6.11)$$

where a , b , and c are parameters that simplify the expression:

$$a = \frac{\cos^2 \theta_p}{2\sigma_x^2} + \frac{\sin^2 \theta_p}{2\sigma_y^2} \quad (6.12)$$

$$b = -\frac{\sin 2\theta_p}{4\sigma_x^2} + \frac{\sin 2\theta_p}{4\sigma_y^2} \quad (6.13)$$

$$c = \frac{\sin^2 \theta_p}{2\sigma_x^2} + \frac{\cos^2 \theta_p}{2\sigma_y^2} \quad (6.14)$$

x_0 and y_0 are the x -axis and y -axis positions of the centre of the Gaussian, θ_p is the position angle the major axis of the Gaussian makes relative to y -axis. σ_x is the Gaussian width in the x -axis, and σ_y is the Gaussian width in the y -axis.

To correctly model the different angular scales of morphological features, we can construct a function which is a summation of Gaussians of varying sizes for the different angular scales:

$$I_{\text{Source}}(x, y; \hat{\theta}) = \sum_{i=0}^{N_{\text{gauss}}} G(x, y; \hat{\theta}_i) \quad (6.15)$$

In this instance $\hat{\theta}_i = (x_0, y_0, \sigma_x, \sigma_y, \theta_p, G_0)_i$ is the vector of parameters for the i th component Gaussian, and $\hat{\theta}$ represents the matrix of vectors with $(N_{\text{gauss}} \times 6)$

elements. To fit the multi-component Gaussian model we minimise the square residuals $(I_{\text{Source}}(x, y; \hat{\theta}) - I_{\text{data}})^2$, with the Python package `scipy` (Virtanen et al., 2020). This method performs well if the boundary conditions for the parameter space and the initial conditions are chosen well. Peak detection methods (discussed further in Sections 6.3.1.1 and 6.3.2.1), instrumental resolution, and known source sizes help to restrict the total number of components, as well as provide good initial guesses on the fit parameters.

6.3.1 Centaurus A

In this work we utilise the best available MWA image of Cen A (McKinley et al., 2021), taken at 185 MHz with an observing bandwidth of 30.72 MHz. McKinley et al. (2021) observed CenA using Phase I MWA data and Phase II extended MWA baseline configuration data. The final image has an rms background noise level of approximately 4 mJy/beam with a peak brightness of 202 Jy/beam in the inner lobes, giving the image a dynamic range of approximately 50000. This image is free of significant artefacts, and provides the most accurate detailed representation of CenA at these radio frequencies to date (McKinley et al., 2021).

6.3.1.1 Centaurus A Morphological Model

The image in Figure 1 from McKinley et al. (2021) was used to create a bespoke morphological model of CenA, by fitting 2D Gaussians to prominent peaks. Since this image is large (1258×2452 pixels), and has four orders of magnitude in dynamic range, it was split into different regions which could be individually modelled. In particular the compact smaller scale structures of CenA such as the inner lobes, the Northern Middle Lobe (NML), and the background galaxy MRC1318-434B were separated into different images.

The bespoke fitting process for these three images was the same; we used the Python package `skimage` to perform local peak detection with the function `peak_local_max` (van der Walt et al., 2014). The peak detection parameters were manually adjusted to




Figure 6.3: Model SIN projected image of Centaurus A fit to the deep multi-scale image from (McKinley et al., 2021). The leftmost panel shows the full extent of Cen A, with the different morphological regions labelled. The inner lobes and the NML in particular are shown in the solid coloured boxes. The top rightmost panel is a close up image of the inner lobes, where the Northern inner lobe (NIL) and the Southern inner lobe (SIL) are labelled separately. The bottom rightmost panel is a close up image of the NML model. The background galaxy MRC1318-434B is shown in the SOL.

choose an appropriate number of peaks for each image. An estimate of the appropriate number of peaks was determined by looking at the images with overlaid contours. We then used the flood fill algorithm from `skimage` to create a cutout island. Islands are subsets of the data on which 2D Gaussian fitting is performed. The flood fill parameters were manually fine tuned until the diffuse emission of each feature was almost entirely encapsulated. Peaks that lay outside of the island were removed. For the inner lobes image we identified 17 peaks, 15 for the NMLs image, and 12 for the background galaxy image.

Before each image was fitted, the background flux density was estimated by calculating the median pixel value of all the pixels outside of the island mask. The background was assumed to be constant throughout each image. This median background flux was then subtracted from the island removing the flux density offset introduced by the outer lobes of CenA. Using the island mask and peak locations, we then fitted the N number of 2D Gaussians to each image using the `scipy.optimize` function `curve_fit()` (Virtanen et al., 2020). We restricted the minimum Gaussian size to have the same parameters as the Gaussian restoring beam for the image (McKinley et al., 2021). The resulting fit for the inner lobes can be seen on the top right hand panel of Figure 6.3, and the resulting fit to the NML can be seen on the bottom right hand panel of Figure 6.3.

Once successful fits to the image were obtained, the models were subtracted from the main CenA image. The source finding algorithm `Aegean` (see Hancock et al. (2012) and Hancock et al. (2018) for details) was then applied to the residual image to identify point sources that might be present in the outer lobes and the periphery. 1034 points sources were found and subtracted from the residual CenA image; most of these sources fell outside of the outer lobes due to the lower background flux density. With the new residual image we used the `astropy` function `block_reduce` to down sample the image by a scale factor of 19. The reduction of the residual image scale reduces overall computational load. The new image had angular pixel sizes of ~ 5 arcmins. The function `block_reduce` can conserve the summation of the flux density in the

down sampling process, which we use here. The Northern Outer Lobe (NOL) and the Southern Outer Lobe (SOL) were then separated into two cropped images, and the same source finding and fitting process applied to the inner lobes and NML was applied to the reduced outer lobe images. In total 9 peaks were found for the NOL, approximately half of which corresponded to the large scale diffuse emission from the NML. A total of 8 peaks were identified for the SOL. The Gaussian fits to these peaks were not restricted to a minimum size, since the pixel size is larger than the PSF in the down sampled images. A total of 61 Gaussians (including the 12 from the background galaxy) were fitted to the CenA image, ranging in size from the the Gaussian restoring beam PSF to $\sim 2^\circ$.

The total CenA model image can be seen in Figure 6.3, which is separated into three panels. The large left hand side panel illustrates the entire 61 component CenA model, with the main features such as the inner lobes and the outer lobes labelled. The smaller right hand side panels illustrate the compact models of the inner lobes⁴ and the NML respectively. The main morphological features are labelled in black text.

6.3.1.2 Centaurus A Spectral Model

In addition to the morphology of CenA we require a spectral model at low radio frequencies to capture the spectral structure of CenA in the power spectrum. For this purpose we use the spectral index map shown in Figure 4 of McKinley et al. (2018) as a guide. The spectral index distribution of CenA has been thoroughly investigated in the literature Alvarez et al. (2000); McKinley et al. (2013, 2018). At low radio frequencies the spectral index distribution of CenA is relatively uniform with a spectral index range of -0.5 to -0.8 , and an average spectral index of ~ -0.7 across the entire source. There is small scale regional variation, particularly at the edge of the outer lobes and in the inner lobes (McKinley et al., 2018). For this work following the suggestions from McKinley (private communication) we assign a flatter spectral index of $\alpha = -0.5$ to the inner lobes, and we assign the rest of CenA an approximate median spectral index

⁴The inner lobes are often separated into the Northern Inner Lobes (NIL) and the Southern Inner Lobes (SIL). This convention is retained in the top rightmost pannel.

of $\alpha = -0.7$. For the purposes of this work, a relatively simple spectral behaviour is adequate.

Using the spectral index and the derived flux density for each component of the CenA model, we compare the total integrated flux density from our CenA model⁵ to the measured total integrated flux density from the literature (Alvarez et al., 2000; McKinley et al., 2013). We rescaled literature flux densities by a spectral index of $\alpha = -0.7$ to a frequency of $\nu = 184.95$ MHz. The total integrated model CenA flux density is 4096 ± 274 Jy compared to 5538.8 ± 817.8 Jy for Alvarez et al. (2000), and 4832 ± 1066 Jy for McKinley et al. (2013). The model recovers most of the flux density, with some flux density missing on intermediate and small scales in the outer lobes. The difference of $\Delta S_{\text{tot}\pm} \sim 15\%$ with our model compared to McKinley et al. (2013) does not affect our ability to answer the question as to whether or not CenA causes leakage into the EoR window for EoR 2 observations. Additionally the flux scale uncertainty for the total CenA flux density calculated by McKinley et al. (2013) were $\sim 20\%$, so for all applied purposes in this paper the model CenA flux scale is adequate.

6.3.2 Supernova Remnants

The cataclysmic end to a massive star's life ejects material at high speeds into the surrounding inter stellar medium. Relativistic electrons accelerated at the shock boundaries of SNRs produce synchrotron radiation as they interact with the local magnetic field (Berezhko and Völk, 2004). This emission is dominant at radio wavelengths particularly around 1 GHz (Stafford et al., 2019). Known Galactic SNRs in the low frequency radio regime have been extensively studied (see Dubner and Giacani, 2015, for a review), and have been collated into a comprehensive catalogue (Green, 2019). This catalogue provides information about the position in RA and DEC, as well the major and minor elliptical sizes of each SNR. Additionally the catalogue provides the expected 1 GHz flux density and spectral index derived from the literature where possible (see Green, 2019, for references).

⁵This is including the background galaxy as the comparison is made to measurements made at low resolution which confuse the background galaxy with the diffuse emission of the outer lobes.

Green (2019) SNR catalogue contains 294 Galactic SNRs, 269 of which have 1 GHz flux density values. 25 SNRs either had no 1 GHz flux density estimates, or only had upper limits, and were removed from the catalogue. Of the remaining 269 sources only 218 had spectral index measurements, some of which are dubious (Green, 2019). For the 51 SNRs that did not have spectral index values they were assigned the population median spectral index value of $\alpha \sim -0.5$ as a placeholder. The SNR flux densities were then scaled from 1 GHz flux to a frequency of 183 MHz, which corresponds to the frequency at the centre of the simulated EoR 2 field observations.

Further subsetting of the SNR catalogue is performed using major axis size of the remaining SNRs. A cutoff size of ≥ 23 arcminutes is applied since this is twice the size of the $> 300\lambda$ (~ 11.5 arcminutes) uv -cutoff. This cutoff is applied in uv -space to the visibilities because the 21 cm signal power is expected to be the greatest at larger spatial scales (Furlanetto et al., 2006b). After applying the major-axis size condition, the SNR catalogue only has 101 remaining SNRs. Additional subsetting is performed for SNRs below a declination of $+30$ deg, of which there are 73. Sources above this cutoff are not contained in The GaLactic and Extra-galactic All-sky Murchison Widefield Array (GLEAM) survey. GLEAM was an all sky survey that observed the southern sky below declinations of $+30$ deg using the MWA (Wayth et al., 2015), images from GLEAM are publicly available through the GLEAM VO server (Hurley-Walker et al., 2017)⁶. For each of these sources we download 200 MHz cutout images from the GLEAM VO server. The 200 MHz wideband GLEAM image is the most sensitive with an angular resolution of ~ 2 arcminutes (Hurley-Walker et al., 2017). The 2D Gaussian component fitting to these images is described in the following section.

6.3.2.1 SNR Morphological Models

For some SNRs which have relatively low surface brightness, island fitting methods such as Aegean and PyBDSF (Mohan and Rafferty, 2015) have a tendency to over-fit the wide-band 200 MHz GLEAM cutout image backgrounds. Due to the relatively

⁶http://gleam-vo.icrar.org/gleam_postage/q/form

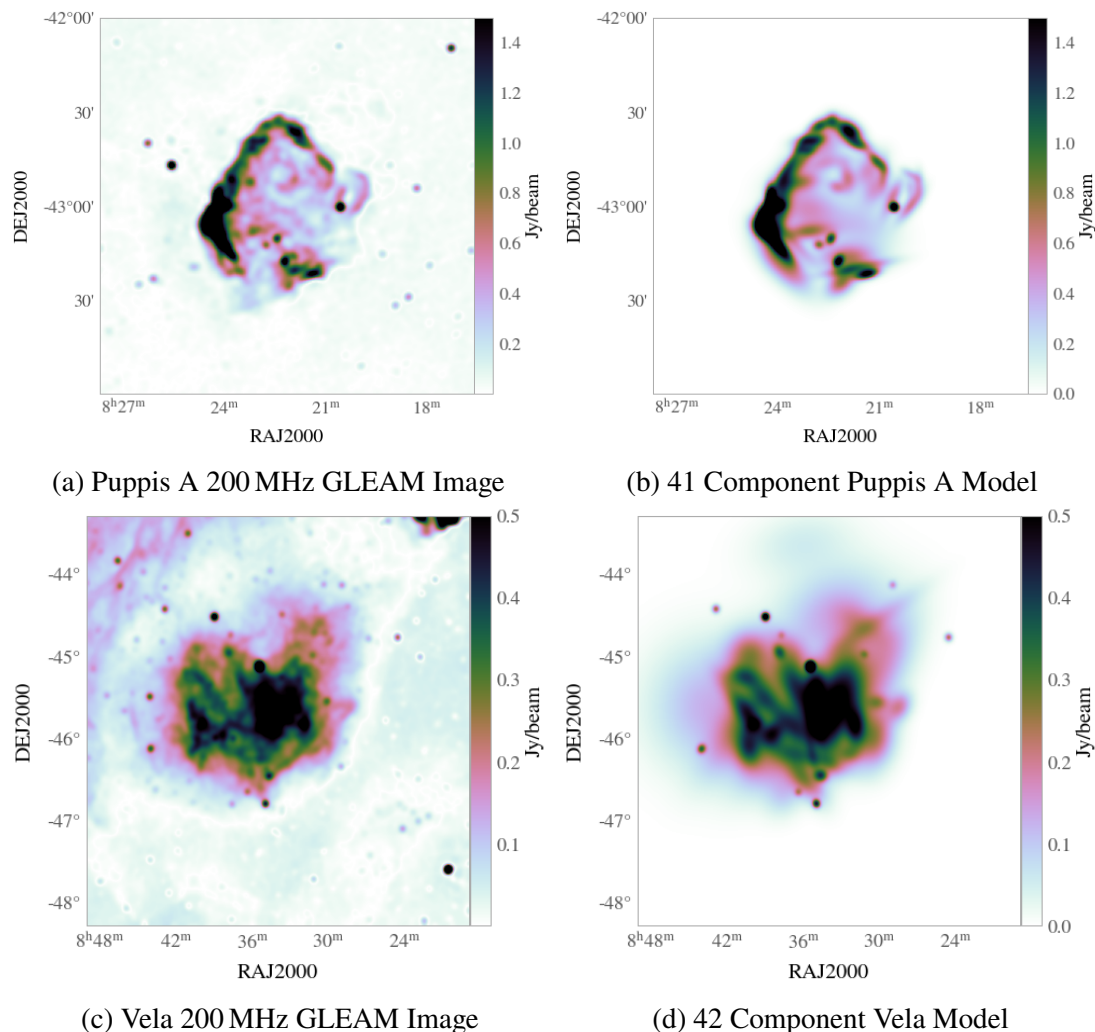


Figure 6.4: GLEAM cutout images of Puppis A (Subfigure (a)) and Vela (Subfigure (b)) at 200 MHz. The Puppis A image has a peak of 4.50 [Jy/beam], and is convolved with a Gaussian restoring beam with a major and minor size of $a = 2.23$ [arcmin], $b = 2.14$ [arcmin], and a position angle of $\sim 315^\circ$ relative to North. The Vela image has a peak flux density of 1.62 [Jy/beam], and is convolved with a Gaussian restoring beam with a major and minor size of $a = 2.23$ [arcmin], $b = 2.14$ [arcmin], and a position angle of $\sim 351^\circ$ relative to North. Due to the size (5×5 deg) of the Vela image, it is further convolved with a Gaussian of size $a_{\text{maj}} \sim 5.41$ [arcmin]. The lower resolution allowed for the fit of fewer components to the Vela image. This does not affect the sky-models in this work, since the sky-model image cube resolution is ~ 11 arcmin. Subfigure (b) is the 41 component Gaussian model for Puppis A, fit only to an 1° circular cutout. The Puppis A model image has a peak flux density of 4.48 [Jy/beam]. Subfigure (c) is the 42 component model of Vela, which was fit to a $\sim 4^\circ$ circular cutout of Vela, and has a peak flux density of 1.60 [Jy/beam].

large number of GLEAM cutout images ($N = 73$), we instead opted to develop an automated fitting method which utilises prior information about the size, and location of each SNR. The prior information is taken from the SNR catalogue, where the major axis and the centroid RA and DEC position for the SNR is used to create an island mask.

The fitting method employed to fit each SNR was similar to the bespoke method developed for CenA, with some key differences. In particular we took a more accurate approach in calculating the image background. This is particularly important for SNRs that have a low surface brightness compared to the image background. The GLEAM SNR cutout images do not have the large dynamic range of the CenA image from (McKinley et al., 2021). For the SNRs the background emission was determined through an iterative approach, where the pixels outside the island were averaged. The fitting algorithm then calculates the root mean squared (rms) of the masked image (island pixels set to NaN). We use a default rms threshold of 2.5σ above the median background to mask potential point sources. The median background and rms are then recalculated and further thresholding performed. Convergence to a single background noise level for each cutout image was quick, typically taking a max number of five iterations, this was set as the default.

Once the background has been calculated it is subtracted from the island image. We then perform peak detection using the `skimage` function `blob_dog()`. This method blurs the image with increasing standard deviations (in terms of pixel coordinates), and calculates the difference between successive images which are then stacked into a difference image cube. Blobs or peaks are identified as local maximums in the data cube. This allows for the detection of different scales of peaks (van der Walt et al., 2014).

After peak detection, we then fit 2D Gaussians using the `scipy.optimize` function `curve_fit()`, as we did when fitting CenA. The fitting parameter space is restricted by constraining the maximum Gaussian fit size to a fraction of the SNR major axis

(default fraction is 1/8)⁷. The fitting space is also restricted to be within the island, minimising spurious fits outside the island. Additionally, the minimum 2D Gaussian size is restricted to match the image restoring beam.

To test the validity of the multi-component fit model, we also fit a single 2D Gaussian to each SNR image. For some filled type SNRs this model might be a more accurate representation of the morphology, additionally allowing for an automated comparison which can distinguish between potentially real and spurious fits. However, many fits still had to be assessed by eye to ensure the multi-component models were not fitting noise, or image artefacts. The single 2D Gaussian fit only has two free parameters, the Gaussian amplitude and the position angle. The major and minor axis sizes are fixed from the information from SNR catalogue. To compare the multi-component fits to the single Gaussian fit we utilise the Bayesian Information Criterion (BIC) (Schwarz, 1978):

$$\text{BIC} = \chi^2 + k \log n, \quad (6.16)$$

where χ^2 is the sum of the squared residuals which have been normalised by the squared image rms, k is the number of model fit parameters, and n is the number of data points. The model with the lower BIC is the preferred fit (Schwarz, 1978), which for most SNRs is typically a multi-component model. Some sources were too faint to be present in the GLEAM 200 MHz images, and peaks were fit to sidelobe confusion noise, or to artefacts. In these cases we replaced these fits with the single Gaussian fit. In total out of the 73 fit candidates 24 had a preferential single Gaussian fit.

To determine the accuracy of the SNR models to the expected flux density, the total integrated model flux density for each SNR was compared to the expected flux density provided by Green (2019). The median ratio for all SNRs was $\sim 1.1 \pm 0.4$, with one outlier the Vela SNR model having a ratio of 17.9. The expected flux density for Vela as quoted in Green (2019) was determined from single dish Parkes observations made

⁷The 1/8 size constraint was found to be reasonable, since most observed SNR morphologies are generally dominated by smaller scale filament like structures (Dubner and Giacani, 2015)

by Milne (1968). The GLEAM images are missing baselines below 60 m and thus large scale flux density from Vela.

Figure 6.4 shows example fit models of Puppis A, and Vela compared to their corresponding GLEAM images. The left hand panels are the original GLEAM images, with Puppis A on the top row and Vela on the bottom row. The model images are on the right hand side with Puppis A on top row and Vela on the bottom row.

6.3.3 Constructing Sky-Models

The model fit parameters for CenA and the Galactic Plane SNRs were collated into a FITS table which contains the RA, DEC, spectral index, the total model integrated 200 MHz flux density, the major and minor axes, as well as the position angle for each component. Using this table, models of the entire sky in image space can be generated. For a single frequency slice the sky-model image array can be described as the aggregate of all of the model sources:

$$I_{\text{sky}}(\mathbf{l}) = \sum_{i=1}^{N_{\text{source}}} I_{\text{source},i}(\mathbf{l}) \quad (6.17)$$

This aggregate modelling approach is useful, because it allows for the creation of partial sky-models, effectively simulating source subtraction. This can be used to determine how much of the Galactic Plane SNRs and CenA need to be removed in order to retrieve the 21 cm signal in the power spectrum. For a given observation time, we calculate the Azimuth and Altitude for each source and its model components using *astropy* (Astropy Collaboration et al., 2013, 2018). Sources which are below the observation horizon ($\theta_{\text{alt}} < 0$) are ignored. Substituting Equation 6.15 into Equation 6.17 generalises the description of the total sky model to the aggregate of all the model 2D Gaussian components:

$$I_{\text{sky}}(\mathbf{l}) = \sum_{i=1}^{N_{\text{source}}} \sum_{j=1}^{N_{i,\text{gauss}}} G_j(\mathbf{l}; \hat{\theta}_j), \quad (6.18)$$

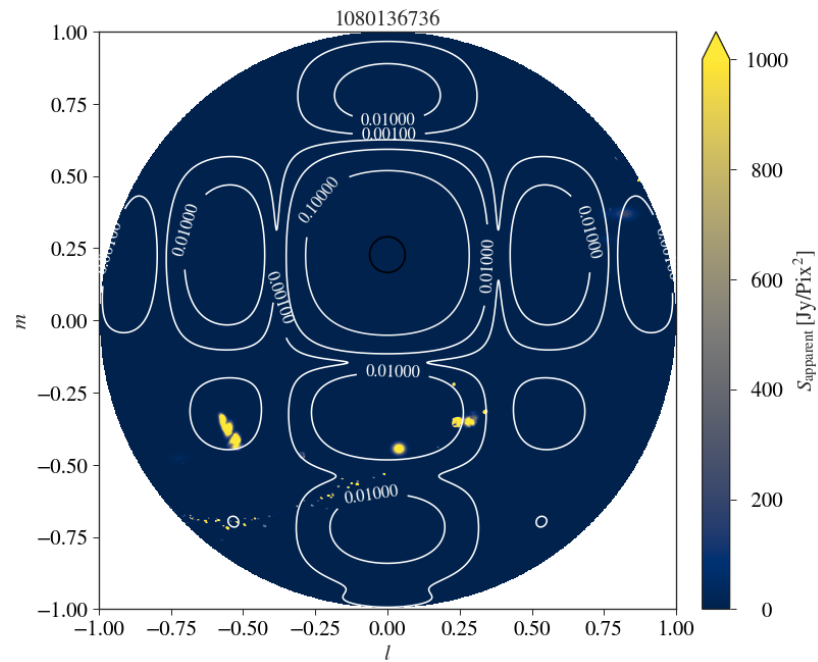
where the j th source has $N_{j,\text{gauss}}$ Gaussian components, with each component having

$\hat{\theta}_j$ model parameters. For a zenith phase centre, the (l, m) plane is an orthographic projection. Due to the small angle approximation, the Gaussian models were defined in a 2D plane, however when placing them in the (l, m) frame they will need to be correctly projected. The major and minor axes for all Gaussians are recalculated as a function of their Altitude angle. This conserves the total flux density of the source. The projection effect is continuous, however to simplify calculations we use an approximation. For more details on how the projection is calculated refer to the Appendix Section 6.C.

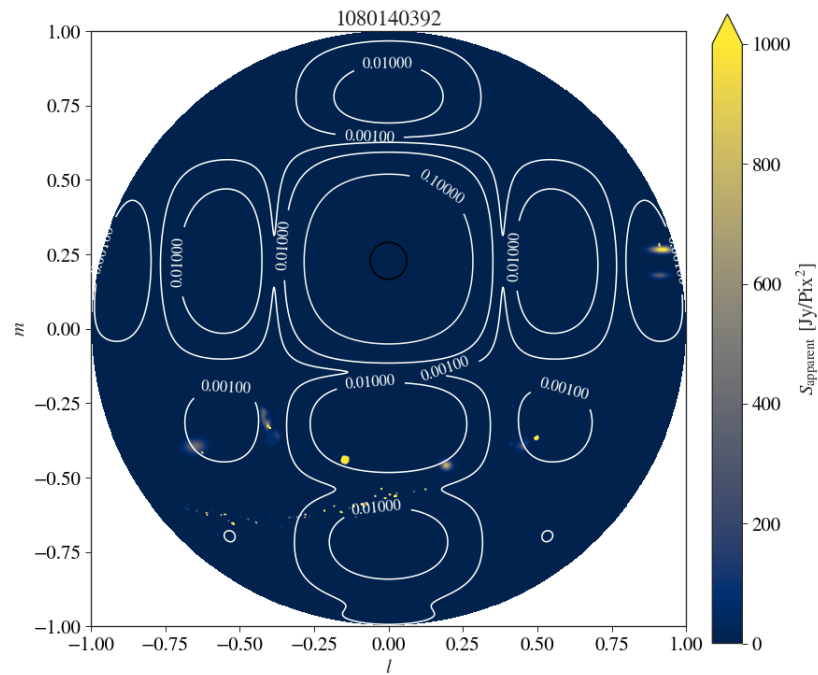
The OSIRIS pipeline accepts a sky-model cube $I(\mathbf{l}, \nu)$ which varies as a function of frequency. In this work we assume that the source morphology does not evolve with frequency across the simulated observation bandwidth (15.36 MHz). This is a reasonable assumption since we fit wideband images (≥ 30.72 MHz) of SNRs and CenA. We also assume that the spectral behaviour of the source components can be modelled with a power law $I \propto \nu^\alpha$, where α is the spectral index. This simplifies the calculation of the sky-model cube, since the OSIRIS pipeline only calculates a template Gaussian which can be scaled as a function of frequency. The iterative sum for each Gaussian model component j for the i th source is described below:

$$I_{\text{sky}}(\mathbf{l}, \nu) = \sum_{i=1}^{N_{\text{source}}} \left(\frac{\nu}{\nu_0} \right)^{\alpha_i} \sum_{j=1}^{N_{i,\text{gauss}}} G_j(\mathbf{l}; \hat{\theta}_j). \quad (6.19)$$

Some Gaussians have $\sigma \ll \Delta l$ (pixel size), and therefore are not properly sampled by the coarse pixel grid. One solution is to increase the grid size to effectively sample the smallest Gaussian model, however this drastically increases the required computational resources. Furthermore, we are not interested angular scales less than ~ 10 arcmin. Instead we set the minimum angular major and minor axis size to be equal to the pixel size (which is ~ 8.4 arcmin), which conserves flux density and effectively sets these smaller components as point sources.



(a) Sidelobe Sky-model



(b) Null Sky-model

Figure 6.5: Example apparent sky-model images for sidelobe sky-model (Subfigure (a)) and the null sky-model (Subfigure (b)). The solid white contours show the MWA primary beam with lines at levels $[10^{-3}, 10^{-2}, 10^{-1}, 0.9]$. In Subfigure (a) Centaurus A can be seen in a sidelobe at $l \sim -0.5$ and $m \sim -0.35$. The Galactic Plane SNR sources can be seen in an arc intersecting several sidelobes, Vela and Puppis A are both visible at $l \sim 0.25$, and $m \sim -0.35$. In Subfigure (b) Centaurus A and Vela have migrated out of their respective sidelobes and into primary beam nulls. Puppis A in Subfigure (b) has migrated into another sidelobe.

6.4 Results

EoR 2 field MWA observations have CenA positioned in one of the MWA primary beam sidelobes, which is a concern for EoR science. The contribution of CenA to the power spectrum is expected to be greater than the 21 cm signal on degree size scales that are important for EoR science. Earth rotation synthesis will mitigate some of the power of CenA as it rotates from the sidelobe into a primary beam null. However a full simulation of hundreds of hours of MWA observations for the EoR 2 field is not necessary to determine whether CenA and Galactic Plane SNRs introduce leakage into the EoR window. Therefore the `OSIRIS` pipeline only simulates a single time step, and thus does not incorporate Earth rotation synthesis. In conjunction with CenA a procession of Galactic Plane SNRs rotates through one of the MWA primary beam sidelobes for the EoR 2 field. The aggregate power of the Galactic Plane SNRs will not be as strongly affected by Earth rotation synthesis, but will however vary as different sources become more prominent. Equation 6.17 allows for the construction of partial sky-models which simulate the subtraction of CenA and Galactic Plane SNRs. In this section we investigate the 2D and 1D power spectrum of several input sky-models of the EoR 2 field. In particular we look at two distinct observations to analyse the different spectral characteristics, and how the resulting leakage affects the detectability of the 21 cm signal.

6.4.1 Sidelobe and Null Test Observations

To characterise the effects of rotation synthesis we simulate two sky-models of the Galactic Plane and CenA separated by one hour in time. The first observation has CenA situated in a sidelobe of the MWA primary beam (herein referred to as the sidelobe observation), and the second observation has CenA situated in a null of the MWA primary beam (herein referred to as the null observation). Figures 6.5a and 6.5b show the average apparent sky-models across the entire observing bandwidth, where the sky-model cube was attenuated by the FEE MWA primary beam model, and averaged as a function of frequency. The average MWA primary beam pattern

across the bandwidth is shown with the solid white contours. Subfigure 6.5a shows the sidelobe sky-model with CenA clearly visible in the sidelobe. Subfigure 6.5b shows the null sky-model with CenA rotated into the primary beam null.

In addition to the sidelobe and null observation simulations, we perform a third simulation of the sidelobe observation without CenA where the model just contains the Galactic Plane SNRs. By comparing the relative difference in the magnitude of the resulting 2D power spectrum we can determine what effect rotation synthesis may have on these observations for different k -modes. We can also compare this to the expected 21 cm power expected on these modes. Figure 6.6 shows the resulting 2D power spectrum for the sidelobe observation (6.6a), the null observation (6.6b), the fiducial 21 cm 2D power spectrum (6.6c), and the ratio of the sidelobe and null 2D power spectrum (6.6d).

The solid and dashed black lines in Figure 6.6 show the expected horizon for the entire sky, and the edge of the field of view (Morales et al., 2012; Trott et al., 2012). The horizon line demarcates the bright foreground wedge from the relatively clean EoR window. To assess the level of leakage we compare the average power in a small window defined by $k_{\perp} \in [0.01, 0.03]$ and $k_{\parallel} \in [0.1, 0.3]$ for the sidelobe, null, and 21 cm 2D power spectra. The average window power in the 21 cm 2D power spectrum is $1.8 \times 10^4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$ compared to $3.44 \times 10^5 \text{ mK}^2 h^{-3} \text{ Mpc}^3$ for the sidelobe 2D power spectrum, and $3.5 \times 10^4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$ for the null 2D power spectrum. The sidelobe observation is ~ 20 times greater than the expected 21 cm signal in the window, compared to a factor of ~ 2 greater for the null observation. For comparison the average window power for a sidelobe observation which contains only CenA is $3.36 \times 10^5 \text{ mK}^2 h^{-3} \text{ Mpc}^3$, clearly showing that CenA is the dominant source of leakage for the sidelobe observation. Subfigure 6.6d shows the excess power of the sidelobe observation compared to the null observation. The largest ratio values (of order 10^3) are mostly confined to the foreground wedge and at higher k_{\perp} , which corresponds to smaller spatial scales. The median ratio in the EoR window is 8.2, which is indicative of the order of magnitude difference in leakage through the EoR window.

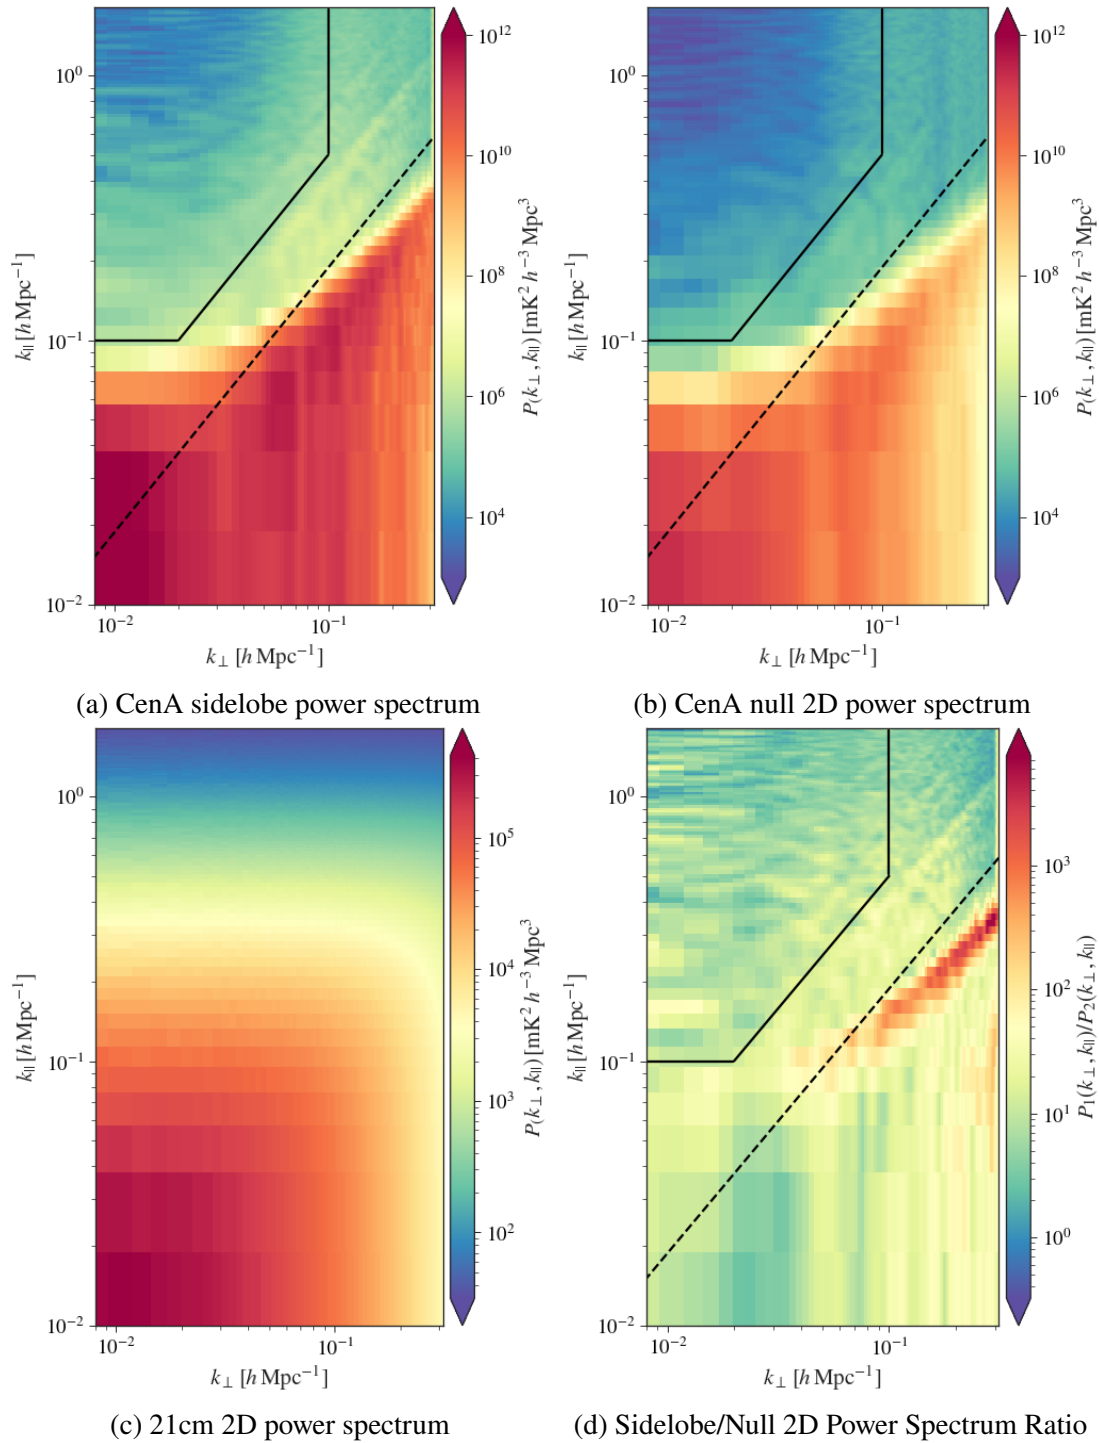


Figure 6.6: 2D power spectra for the sidelobe, null simulation sky-models, and the fiducial 21 cm 2D power spectrum. Subfigure (a) is the 2D power spectrum for the sidelobe case, the solid black line indicates the wedge cut used to calculate the 1D power spectrum in Figure 6.8, the gradient of the solid black line indicates the horizon. The dashed black line indicates the gridding kernels field of view. Subfigure (b) is the 2D power spectrum for the null simulation. Subfigure (c) is the 2D power spectrum of the fiducial 21 cm signal. Subfigure (d) is the ratio of the sidelobe 2D power spectrum to the null 2D power spectrum simulation. Subfigure (a) and (b) have the same colourbar scale.




Figure 6.7: 2D Power spectrum of the sidelobe simulation without CenA in the sky-model. The colourbar scale is the same as the 2D power spectrum in Subfigure 6.6a. There is a clear difference between this 2D power spectrum and that shown in Subfigure 6.6a, with this 2D power spectrum resembling the null 2D power spectrum.

We perform a similar assessment of leakage for a single zenith flat spectrum point source, with an apparent flux density of 10.2 Jy (this is the same as CenA for the sidelobe observation). In this case we perform a noiseless simulation and remove the primary beam, only keeping the spectral tapering. The spectral tapering with the Blackman-Harris window will have sidelobes that will contribute leakage into the window. Performing the same window calculation as per the CenA simulation, we find the median power in the window for the flat spectrum source is $22.4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$, this is ~ 3 orders of magnitude less than the expected 21 cm signal. Therefore we conclude that the Blackman-Harris sidelobes are not the primary contributor to the leakage seen in the EoR window.

We also investigated the 2D power spectrum of the sidelobe simulation without CenA in the sky-model, serving as a useful comparison to the null observation. The resulting 2D power spectrum is shown in Figure 6.7; the colourbar is the same scale as those in Figure 6.6. The average power in the window for the sidelobe minus CenA 2D power spectrum is $2.7 \times 10^4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$. This is a similar level of power compared to the null observation, however the only contribution to leakage in the window is from Galactic SNRs in this case. The similarity between the null simulation and the sidelobe minus CenA simulation may indicate a potential mitigation strategy for reducing the contribution from CenA in EoR 2 observations. However the leakage from Galactic Plane SNRs is still significant, and the change in the spectral properties and intensities of SNRs as the Galactic Plane rotates through the primary beam could be significant.

6.4.2 Partial Sky Models


Figure 6.7 demonstrates that even without CenA in the input sky-model, the leakage of power into the EoR window from Galactic plane SNRs is on the order of the expected fiducial 21 cm signal power. In this section we assess how much of the SNRs need to be subtracted from the sidelobe and null sky-models in order to significantly recover the 21 cm signal. To test this we generated a series of partial sky-model simulations for both the sidelobe and null sky-models without the 21 cm signal. The sky-model catalogue was ordered by the apparent flux density from the faintest to the brightest source; the fractional total apparent flux density for each source was then calculated. We then generated three sky-models for each observation with upper limits of 10%, 50% and 90% of the total apparent sky-model flux density. We shall refer to these as the deep, the medium, and the shallow partial apparent sky-models respectively. The partial sky-model method assumes an ideal case where we can subtract 100% of a sources total flux density. However, in reality this is not possible; simulating source subtraction errors (position or amplitude errors specifically) will not affect the main question of this paper. The partial sky-models along with the total SNR sky-model, and the CenA only sky-model for both observations were run through the OSIRIS pipeline. The 1D

power spectrum was then calculated from window modes defined by $k_{\parallel} > 0.1 h \text{ Mpc}^{-1}$, $k_{\perp} < 0.1 h \text{ Mpc}^{-1}$, and $(k_{\perp}, k_{\parallel})$ modes above the horizon⁸. We also calculated the 1D power spectrum for the fiducial 21 cm signal plus the simulation noise (\mathcal{N}). The resulting 1D power spectrum for both observations and the respective partial and total sky-models can be seen in Figure 6.8.

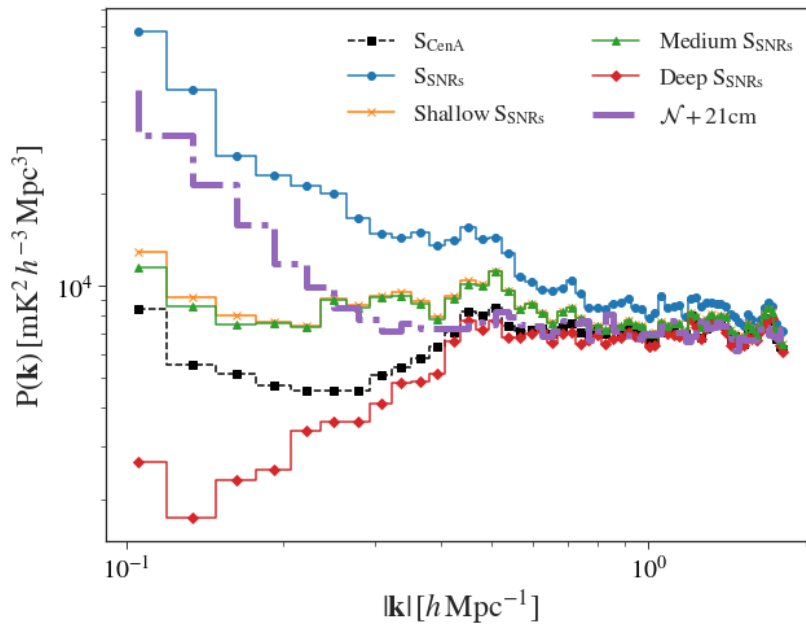
The orange crosses, solid green triangles, and the solid red diamonds show the deep (90%), the medium (50%), and the shallow (10%) upper limit partial sky-model power spectrum for both the sidelobe and null observations in Figure 6.8. Since the partial sky-models are discretised by source and ordered from faintest to brightest, the relative percentages for the deep, medium and shallow partial sky-models are different for the sidelobe and the null observations. For the sidelobe observations the relative percentages approximately are 10%, 36%, and 76% for the deep, medium and shallow partial sky-models. For the null observation the relative percentages are approximately 10%, 50%, and 74% respectively. For reference, the total SNR sky-model power spectrum and the CenA only sky-model power spectrum are shown with the solid blue circles and the solid black squares respectively. The dash dot purple line is the fiducial 21 cm signal with a 1000 hr noise level.

The sidelobe and null observations have a similar total apparent brightness ($\sim 8 \text{ Jy}$ for both), however in Figure 6.8 there is significant difference in the total 1D power spectrum. The null and sidelobe observations are separated by one hour in time and therefore most of the the SNRs in the model are the same, but in different parts of the MWA primary beam. For small and faint SNRs this has little impact on the power spectrum, as can be seen from the similarities in structure and power for the deep and medium upper limit partial sky-models for the sidelobe and null observations. However, this matters for the brightest most prominent sources which affect the shallow partial sky-model and the total SNR sky model. The difference between the medium, the shallow and the total 1D power spectra for both the null and the sidelobe observations

⁸The horizon k -mode cut is defined by the relationship: $k_{\parallel} > \frac{\pi}{2} \frac{D_M E(z)}{D_H (1+z)} k_{\perp}$ (Morales et al., 2012), where D_M is the co-moving distance, D_H is the Hubble distance, $\pi/2$ is the radius of the sky in radians, and the function $E(z)$ is defined by Hogg (1999).



(a) Sidelobe Partial 1D Power Spectra



(b) Null Partial 1D Power Spectra

Figure 6.8: 1D power spectra for a series of partial input sky-models, showing linearly spaced k bin widths. The black square markers with the dashed black line represent the CenA only sky-model, the solid blue circles represent the total SNR sky-model. The orange crosses, the green triangles and the red diamonds are the 1D power spectra are for partial sky-models with upper limit 10% (deep), 50% (medium) and 90% (shallow) total model apparent flux density. The actual percentages for the deep, medium and shallow partial sky-models are $\sim 10\%$, $\sim 36\%$ and $\sim 76\%$ for the sidelobe observation, and $\sim 10\%$, $\sim 50\%$ and $\sim 75\%$ for the null observation. The dash dot purple line with no markers is the fiducial 21 cm signal. Subfigure (a) shows the partial sky-models for the sidelobe observation. The medium sidelobe partial sky-model is on the order of the fiducial 21 cm power spectrum, the deep partial sky-model is below the fiducial 21 cm power spectrum. Subfigure (b) shows the partial sky-models for the null observation.

are typically one or two bright extended sources; their morphology and the primary beam spectral structure imparted upon them, has the biggest impact on leakage in the 1D power spectrum.

For the sidelobe observation the total sky-model and the shallow partial sky-model are the same order of magnitude as the fiducial 21 cm signal, indicating significant contamination of the signal. In contrast the null observation shallow partial sky-model is significantly below the expected 21 cm signal on modes $|\mathbf{k}| < 0.3 h \text{ Mpc}^{-1}$. Additionally, there is little difference between the null shallow and medium partial sky-model 1D power spectra. The difference between the shallow and medium sky-models is two exceptionally large ($\sim 3 \text{ deg}$) SNRs G205.5+00.5, and G330.0+15.0⁹. Together their apparent brightness is $\sim 2 \text{ Jy}$. Due to their large degree-scale sizes these sources did not have the surface brightness to be fitted by the GLEAM cutout image method in Section 6.3.2.1. Therefore, these sources are modelled by single component Gaussians and are missing the smaller scale structures present in their morphology. The large single Gaussian components act as a spatial filter in the k_{\perp} axis, modulating and restricting all leakage to $k_{\perp} < 0.02 h \text{ Mpc}^{-1}$ modes. When calculating the 1D power spectrum, the relatively few number of modes $k_{\perp} < 0.02 h \text{ Mpc}^{-1}$ with significant leakage are averaged over, reducing their contribution to the window. This demonstrates the sensitivity of this type of analysis to the details of extended source morphology, and why accurate SNR subtraction is crucial. Future work will better model large single component SNRs to more accurately investigate their leakage.

For both the sidelobe and null observations the 21 cm signal has a power ratio of $\sim 5 - 10$ at $|\mathbf{k}| \in [0.1, 0.3] h \text{ Mpc}^{-1}$ for the deep partial sky-model. For the medium partial sky-model ($\sim 50\%$) the signal to model power ratio is approximately $\sim 2 - 3$, requiring at least 90% subtraction of the SNRs from the sidelobe and null sky-models in this simulation to retrieve a significant detection of the 21 cm signal.

⁹The SNRs G205.5+00.5, and G330.0+15.0 are also known as the Monoceros Nebula, and the Lupus Loop.

6.5 Discussion

Using an input sky-model of Galactic Plane SNRs and CenA processed through an MWA simulation and power spectrum pipeline, we demonstrate that extended radio sources in the sidelobes of EoR 2 observations introduce leakage up to an order of magnitude greater than the 21 cm signal into the EoR 2D power spectrum window. This work shows that almost all of these widefield extended sources must be removed from the visibilities, in order to reduce contamination on EoR significant k -modes ($|\mathbf{k}| \lesssim 0.1 h \text{ Mpc}^{-1}$), down to $\sim 10 - 20\%$ of the 21 cm signal power. Additionally, the position of sources in the MWA primary beam matters for the overall level of leakage expected in the EoR window, as the spectral behaviour of the primary beam varies dramatically across the sky. This effect can be seen in Trott et al. (2020), which demonstrated the chromatic effects of the MWA primary beam as a function of angular position by calculating the beam spectral index across a 30.72 MHz observing band. Figures 27 to 29 from Trott et al. (2020) demonstrate the steep changes at the edges of sidelobes which have spectral indices that range from -30 to 30 . These Figures only capture the first order changes in the beam as a function of frequency. From Cook et al. (2021) Figure 6 we see that for a fixed angular position the MWA primary beam can have complex polynomial like structure, not easily described by a simple power law. This spectral structure far from the main lobe of the primary beam is imparted onto radio sources, varying their spectra more rapidly with frequency. This changing structure of the MWA primary beam with position and frequency is primarily responsible for the leakage seen in the EoR window in this work.

We can assess the level of spectral leakage from CenA into the EoR window for the sidelobe observation, by comparing the expected DC power level of CenA to the power level measured in the EoR window. The apparent brightness of CenA for the sidelobe observations is $\sim 10.2 \text{ Jy}$, which leads to an expected DC mode power of $2.54 \times 10^{13} \text{ mK}^2 h^{-3} \text{ Mpc}^3$, after applying the appropriate conversions. The power at $k_{\perp} = 0.01 h \text{ Mpc}^{-1}$, $k_{\parallel} = 0.1 h \text{ Mpc}^{-1}$ is $6.93 \times 10^5 \text{ mK}^2 h^{-3} \text{ Mpc}^3$, which is a level of leakage on the order of 0.01% . The apparent flux density of the SNR only sky-

models for the sidelobe and null observations is comparable to the CenA apparent flux density. However, there is an order of magnitude less leakage. Performing the same calculation for the sidelobe observation with only SNRs we find a power level at $k_{\perp} = 0.01 h \text{ Mpc}^{-1}$, $k_{\parallel} = 0.1 h \text{ Mpc}^{-1}$ of $7.2 \times 10^4 \text{ mK}^2 h^{-3} \text{ Mpc}^3$, for approximately 0.005% leakage. Modelling and removing these sources will yield improvements by reducing leakage. This has implications for MWA EoR observations at certain pointings (not just the EoR 2 field). In particular the EoR 1 highband field observation from Trott et al. (2020) in Figure 14, clearly has sidelobes intersecting the Galactic plane. However, this part of the Galactic Plane is not as dominated by SNRs as the part visible in the EoR2 field observations.

One important consideration is determining what the expected leakage might be for SKA-LOW observations. The individual SKA-LOW stations will have pseudo random distributed antennas to reduce the average sidelobe gain for all the station tiles (Dewdney et al., 2013). However, as a result of the pseudo random antenna distribution, the station primary beam has two distinct regions outside the main lobe. One region with regular sidelobes close to the main lobe called the coherent region, and another region $\sim 0.3\sqrt{N}$ (N is the number of antennas per station) sidelobes away from the main lobe with randomly distributed sidelobes, this is called the incoherent region Mort et al. (2016). Assuming we have a similar observation of the EoR 2 field with the future SKA-LOW array, due to the smaller field of view, CenA and the Galactic SNRs find themselves in the incoherent part of the SKA-LOW primary beam (> 30 deg from the main lobe). The incoherent part of the SKA-LOW primary beam has an expected power proportional to $\sim 1/N = 0.004$. This is confirmed for the average SKA-LOW station beam through OSKAR (Mort et al., 2010) simulations of the SKA-LOW primary beam at 180 MHz (assuming an analytic log-dipole antenna model with no mutual coupling). The expected beam power in the incoherent region of the OSKAR simulated average primary beam was found to be 0.003. This is coincidentally approximately the same beam power as the MWA sidelobe CenA occupies in the sidelobe observation. If we assume similar beam spectral behaviour, we would find a similar level of leakage in

the EoR window for future SKA-LOW EoR 2 field observations. Analysing how the SKA-LOW station beam changes with frequency is outside the scope of this work, however the chromatic nature of the station tiles, and the bright extended nature of radio sources in the incoherent region, will require consideration in future SKA-LOW EoR observations.

6.5.1 Future Work

In the process of investigating and fitting SNRs using the GLEAM cutout images, we noticed there are numerous HII regions which are bright at MWA radio frequencies. These regions also have similar sizes and scales to SNRs, and thus to the 21 cm ionisation bubbles. Similarly to Green (2019) there is a comprehensive Galactic HII catalogue containing 1442 HII regions (Paladini et al., 2003). This catalogue provides diameters, and flux densities at 2.7 GHz. HII regions are relatively opaque at the lower frequencies which the MWA observes for the EoR fields. However, there are still HII regions which are bright enough to be detected at MWA frequencies and have been observed by GLEAM (Wayth et al., 2015). A similar method can be applied to model the HII regions using the catalogue information as a prior.

The 1D power spectrum of the CenA only null observation in Figure 6.8b, demonstrates a potential observation strategy for the EoR 2 field, where CenA is strategically placed in a null. Morgan et al. (2019) developed a method for determining the best MWA primary beam projection to place the sun in a null for a particular pointing. This could be a useful observing strategy for the EoR 2 field going forward. This however will not be effective for Galactic Plane SNRs, since the Galactic Plane SNRs span the entire breadth of the sky.

6.5.1.1 Morphological Models

The morphological models presented in this work are a good first step to removing their contribution from the visibilities of EoR observations, particularly for the EoR 2 field. The CenA and Galactic SNR models have a ~ 1.5 arcmin angular resolution, which

corresponds to $k_{\perp} = 2.4 h \text{ Mpc}^{-1}$. We perform a 300λ cutoff effectively smoothing over angular scales smaller than ~ 11.5 arcmin. However, accurate models of these smaller scale components are still important. Errors on the order of a few percent for smaller scale components will be averaged over larger angular scales, introducing leakage into k -modes less than $0.3 h \text{ Mpc}^{-1}$.

Improvements to the morphological model fitting on all relevant angular scales can be made, especially for the largest and brightest sources. Other basis functions for fitting the morphological structure besides Gaussians exist, such as shapelets (Refregier, 2003) which are an orthonormal set of functions based on Hermite polynomials. Line et al. (2020) compared morphological Gaussian component models and shapelet models of the extended complex radio galaxy Fornax A. Shapelets performed better at modelling the complex smaller scale angular structure ($\theta < 11.5$ arcmin) of Fornax A, and could prove useful in modelling the complex structure of SNRs, as well as the intermediate scales of CenA.

6.5.1.2 Centaurus A Model

There are some important caveats regarding the morphological model of CenA, in particular the larger scale components of the outer lobes. Referring to Figure 1 of McKinley et al. (2021), the outer lobes of CenA contain complex structure from arcminute to degree size scales. Due to the large extent of the image, the larger scales were down sampled by a factor of 19, conserving the flux summation. This effectively removed angular structures on scales of less than 5 – 10 arcminutes. This reduces the complexity of the model at the cost of accuracy. As a result our model of CenA under predicts the flux density of CenA. For the model of CenA presented in this work to be useful for further EoR science the intermediate angular scales will need to be modelled appropriately.

6.5.1.3 The OSIRIS Pipeline

The OSIRIS pipeline developed for this work is self consistent, and compares well to a similar pipeline MAJICK (Line, 2017). However there are several areas in which the OSIRIS pipeline can be improved. Currently OSIRIS accepts a sky-model cube, which is then Fourier transformed via an FFT to derive the Fourier sky-cube. Since Gaussians have analytic Fourier transforms it is possible to generate a Fourier sky-cube without performing an FFT. Analytic Fourier transforms of Gaussian component image cubes, would allow for a nominal speed boost, and would reduce FFT related errors (Lanman et al., 2022). However, the benefit of using an FFT is any sky-model can be input into OSIRIS. This could be incorporated as a future feature to OSIRIS, where a user can choose to perform an FFT or analytically determine the Fourier sky-cube.

The OSIRIS pipeline could also incorporate rotation synthesis. This would allow for more accurate simulations of snapshot observations; with better (u, v) plane coverage. Additionally, several processes of the OSIRIS pipeline can be made parallel to increase simulation speed, which would be necessary if we were to upgrade OSIRIS to incorporate rotation synthesis. These upgrades may be unnecessary with MWA simulation packages such as (WODEN; Line, 2022). In future work we plan to incorporate WODEN simulations when generating observation model visibilities.

6.6 Conclusions

In this work we simulate all-sky images containing only extended radio sources such as CenA and Galactic SNRs. We use these models to determine the level of leakage in the EoR window for the MWA EoR2 field. We find that up to $\sim 50 - 90\%$ of the complex extended sources need be subtracted from the visibilities in order to reduce leakage to a level of $\sim 10 - 20\%$ of the expected 21 cm signal; this is in addition to the compact point sources which are already subtracted. The leakage from these extended sources is primarily caused by widefield chromatic effects of the MWA primary beam far from the main lobe. Additionally, we find that although the future SKA-LOW primary beam is

an improvement compared to the MWA, chromatic effects and leakage from widefield sources will still affect extended widefield sources. Extended widefield sources will likely need to be subtracted in order to perform EoR science with the SKA.

Acknowledgements

This research was partly supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project number CE170100013. JHC is supported by a Research Training Program scholarship. CMT is supported by an ARC Future Fellowship under grant FT180100321. The International Centre for Radio Astronomy Research (ICRAR) is a Joint Venture of Curtin University and The University of Western Australia, funded by the Western Australian State government. The MWA Phase II upgrade project was supported by Australian Research Council LIEF grant LE160100031 and the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto. 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. We acknowledge the Pawsey Supercomputing Centre which is supported by the Western Australian and Australian Governments.

Data Availability

The MWA Centaurus A radio image taken from McKinley et al. 2021 is available through the Strasbourg Astronomical Data Center (CDS) via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5) or via <https://cdsarc.unistra.fr/viz-bin/cat/J/other/NatAs>. The cutout GLEAM images used to model supernova remnants are publicly available through the GLEAM VO server http://gleam-vo.icrar.org/gleam_postage/q/form Hurley-Walker et al. 2017. The sky-models, the re-

sulting visibilities and their power spectra were simulated via the pipeline Observational Supernova-remnant Instrumental Reionisation Investigative Simulator (OSIRIS), which is publicly available at <https://github.com/JaidenCook/OSIRIS>. Examples of how to replicate the sky-model and power spectrum output data arrays used in this work are available in the Github documentation. These simulations model Murchison Widefield Array Phase I data (MWA; Tingay et al. 2013), available at <https://asvo.mwatelescope.org/>. The OSIRIS pipeline uses MWA observation metafits files to generate the primary beam for simulations, these can be downloaded at <https://asvo.mwatelescope.org/>.

6.A Cosmological Conversion

To meaningfully understand the cosmological significance of the EoR signal we convert the (u, v, η) coordinates and the power to be in terms of cosmological coordinates. This cosmological conversion is described by Morales and Hewitt (2004):

$$k_x = \frac{2\pi u}{D_M(z)} h \text{ Mpc}^{-1} \quad (6.20)$$

$$k_y = \frac{2\pi v}{D_M(z)} h \text{ Mpc}^{-1} \quad (6.21)$$

$$k_{||} = \frac{2\pi H_0 f_{21} E(z) \eta}{c(1+z)^2} h \text{ Mpc}^{-1} \quad (6.22)$$

H_0 is the Hubble constant, f_{21} is the 21 cm frequency, z is the redshift, and $E(z)$ is the cosmological function given by $E(z) = \sqrt{\Omega_M(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda}$. $D_M(z)$ is the co-moving transverse distance, which is given by Hogg (1999):

$$D_M(z) = D_H \int_0^z \frac{dz'}{E(z')} \quad (6.23)$$

This is the co-moving distance and has units of h^{-1} Mpc. This transforms our signal into cosmological units.

6.A.1 Conversion Factor

We can describe S_η in terms of the temperature brightness using Rayleigh-Jeans law:

$$S_\eta = \Omega \Delta \nu_f \frac{2k_b}{\lambda_o^2} T_b \text{ Jy Hz}, \quad (6.24)$$

T_b is the temperature brightness, $\Delta \nu_f$ is the channel width in Hz, Ω is the field of in steradians. We square Equation 6.24, and then normalise by the volume $\Omega \Delta \nu$, where $\Delta \nu$ is the observation bandwidth. We can relate $\Omega \Delta \nu = \theta_x \theta_y \Delta \nu$, where θ_x and θ_y are both defined in Morales and Hewitt (2004). Morales and Hewitt (2004) provides a conversion for θ_x and θ_y in terms of cosmological parameters:

$$\Omega \Delta \nu = \frac{r_x r_y \Delta r_z}{D_M^2(z) D_H} \frac{\nu_{21} E(z)}{(1+z)^2} \text{ sr Hz}. \quad (6.25)$$

Note that $r_x r_y \Delta r_z = \Delta V_C$ our co-moving volume element. It can then be shown that:

$$\frac{\lambda_o^4}{4k_b^2} \frac{S_\eta^2}{\Omega \Delta \nu} = \frac{\Delta \nu_f^2}{\Delta \nu^2} \frac{\Delta V_C}{D_M^2(z) D_H} \frac{\nu_{21} E(z)}{(1+z)^2} T_b^2 \text{ K}^2 \text{ sr Hz}, \quad (6.26)$$

Rearranging we obtain our final expression:

$$N_c^2 (1+z)^2 \frac{D_M^2(z) D_H}{\nu_{21} E(z)} \frac{\lambda_o^4}{4k_b^2} \frac{S_\eta^2}{\Omega \Delta \nu} = \Delta V_C T_b^2 \text{ K}^2 \text{ Mpc}^3. \quad (6.27)$$

From Equation 6.27 we can define the cosmological unit conversion factor from $\text{Jy}^2 \text{ Hz}^2$ to $\text{Mpc}^3 \text{ mK}^2 \text{ Jy}^{-2} \text{ Hz}^{-2}$:

$$C = (1+z)^2 \frac{D_M^2(z) D_H}{\nu_{21} E(z)} \frac{\lambda_o^4}{4k_b^2} \frac{N_c^2}{\Omega \Delta \nu} \times 10^6 \text{ Mpc}^3 \text{ mK}^2 \text{ Jy}^{-2} \text{ Hz}^{-2} \quad (6.28)$$

6.B Thermal Noise

The radiometer equation for a single baselines is given by Taylor et al. (1999):

$$\sigma = 2 \frac{k_b T_{\text{sys}}(\nu)}{A_{\text{eff}} \sqrt{\Delta\nu\Delta t}}, \quad (6.29)$$

$k_b = 1380.648 \text{ Jy K}^{-1} \text{ m}^2$ is Boltzmann's constant, $A_{\text{eff}} = 21.5 \text{ m}^2$ is the effective area of the MWA tile, $T_{\text{sys}}(\nu)$ is the system temperature:

$$T_{\text{sys}}(\nu) = 50 + 228 \left(\nu / 150 \text{ MHz} \right)^{-2.53} \text{ K} \quad (6.30)$$

6.C 2D Gaussian Projection Approximation

For orthographic projections of the celestial sphere circular Gaussians will be compressed as a function of their Altitude/Zenith angle. This can be generalised in the case of an elliptical Gaussian where we have an exaggerated representation of the problem in Figure 6.9. In the case of Figure 6.9 the coordinate system is the (l, m) plane. The red ellipse will have some semi-major and semi-minor axis sizes (a, b) , a centre positioned at (l_0, m_0) , an azimuth angle ϕ_0 relative to the m -axis, and a position angle θ_{pa} relative to the non-rotated reference frame of the ellipse.

Compression of the ellipse happens only along the radial direction, for convenience we work in the rotated reference frame which aligns with the radial direction (l', m') , which is rotated with respect to the azimuth angle θ_0 . In this case our ellipse is rotated with respect to the m' axis by the position angle θ_{pa} . The non-rotated reference frame of the ellipse is denoted by (l'', m'') . An example of this can be seen in Figure 6.10.

Compression of the Gaussian is a fundamentally continuous process that occurs as a function of $\cos \theta$, where θ is the zenith angle. Since most Gaussians in astronomy are small in angular scale we can approximate the compression, by compressing the entire m' axis by the value $\cos \theta_0$. We can then use Pythagoras theorem to determine an approximation of what the new semi-major and minor axes size will be:




Figure 6.9: (l, m) plane of the visible celestial sphere. An ellipse in red offset from the centre is located at an azimuth angle of ϕ_0 .




Figure 6.10: Ellipse in the non-offset rotated frame. Here the ellipse is rotated by the intrinsic position angle θ_{pa}

$$a' = \sqrt{\delta l_a^2 + (\delta m_a \cos \theta_0)^2} \quad (6.31)$$

$$a' = a \sqrt{\sin^2 \theta_{pa} + \cos^2 \theta_{pa} \cos^2 \theta_0} \quad (6.32)$$

$$b' = \sqrt{\delta l_b^2 + (\delta m_b \cos \theta_0)^2} \quad (6.33)$$

$$b' = b \sqrt{\cos^2 \theta_{pa} + \sin^2 \theta_{pa} \cos^2 \theta_0} \quad (6.34)$$

Where $\delta l_a = a \sin \theta_{pa}$, $\delta m_a = a \cos \theta_{pa}$. and Where $\delta l_b = b \cos \theta_{pa}$, and $\delta m_b = b \sin \theta_{pa}$. These components are described by the uncompressed components which are derived in an uncompressed flat plane.

CHAPTER 7

THE IMPACT OF IONISATION MORPHOLOGY AND X-RAY HEATING ON THE COSMOLOGICAL 21CM SKEW SPECTRUM

Adapted from:

Jaiden H. Cook, Sreedhar Balu, Bradley Greig, Cathryn M. Trott, and Jack L. B.

Line, Yuxiang Qin, J. Stuart B. Wyithe, (2023), "The Impact of ionization
Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum",
submitted to Monthly Notices of the Royal Astronomical Society

This chapter is primarily my own work except for the simulations which were contributed by S. Balu. From the simulated data I calculated the power spectrum and the skew spectrum. Furthermore, I developed the normalised skew spectrum, and performed calculations regarding the cosmic variance for the power spectrum, skew spectrum and normalised skew spectrum. The idea for this chapter came from discussions with B. Greig, C. M. Trott, J. L. B. Line, and myself, regarding non-Gaussianity and the 21cm signal, following prior work performed by B. Greig. Importantly, discussion and feedback contributed by B. Greig were instrumental in analysing and interpreting the results. Further discussion and insight from S. Balu, Y. Qin, and J. Stuart B. Wyithe contributed important input into this work. Additionally, important discussions regarding the interpretation of the results were contributed by my supervisors C. M. Trott and J. L. B. Line.

I was the primary author on the paper, I drafted the majority of the manuscript, with the exception of Section 7.3 which was drafted and edited by S. Balu. The manuscript was distributed to co-authors, who provided feedback, and proof reading, until the final version was produced.

Abstract

The cosmological 21cm signal offers a potential probe of the early Universe and the first luminous and ionising sources. Current experiments probe the spatially-dependent variance (Gaussianity) of the signal through the power spectrum (PS). The signal however is expected to be highly non-Gaussian due to the complex topology of reionisation and X-ray heating. We investigate the non-Gaussianities of X-ray heating and reionisation, by calculating the skew spectrum (SS) of the 21cm signal using `MERAXES`, which couples a semi-analytic galaxy population with semi-numerical reionisation simulations. The SS is the cross-spectrum of the quadratic temperature brightness field with itself. We generate a set of seven simulations from $z = 30$ to $z = 5$, varying the halo mass threshold for hosting star formation, the X-ray luminosity

per star formation rate, and the minimum X-ray energy escaping host galaxies. We calculate the SS for each simulation, finding the SS is predominantly negative as a function of redshift, and transitions to positive towards the end of heating and early stage of reionisation, remaining positive throughout reionisation. This culminates in a peak at the mid-point of reionisation in both the PS and the SS. We then normalise the SS by the PS during the early, mid and late periods of reionisation to isolate the non-Gaussianities. The normalised SS has a trough ($k \sim 0.1 \text{ Mpc}^{-1}$) and peak ($k \sim 0.4 - 1 \text{ Mpc}^{-1}$) structure during the mid and late points of reionisation, which correlate to the ionisation topology, and the neutral islands in the IGM. We then calculate the cosmic variance of the SS and normalised SS, and find there is high enough signal to noise to detect the trough/peak features by future SKA_LOW observations in the absence of thermal noise and foregrounds.

7.1 Introduction

The cosmological 21cm neutral hydrogen line promises to be an insightful probe of the first luminous sources and the structure of the Universe during early cosmic time. The first luminous sources (stars, galaxies, compact objects) heat and ionise the surrounding intergalactic medium (IGM), through the cumulative emission of ultraviolet (UV) and X-ray photons (see the following review papers: Barkana and Loeb, 2001; Morales and Wyithe, 2010; Pritchard and Loeb, 2012; Furlanetto, 2016).¹ These ionised bubbles grow and eventually overlap, culminating in the end of reionisation by redshift ~ 5.3 (Bosman et al., 2022). These bubbles encode information about these sources onto the cosmological 21cm temperature brightness signal (Furlanetto and Oh, 2005). These luminous sources also heat the neutral hydrogen medium through X-ray emission, which encodes additional information about these sources (Pritchard and Furlanetto, 2007; Furlanetto, 2016). The cosmological 21cm signal is measured relative to the cosmic microwave background (CMB), and can be either in relative emission or absorption. By

¹Reionisation predominately occurs due to UV photons, with some contribution from X-ray emission (up to ~ 10 percent Mesinger et al. (2013)).

measuring the 21cm signal we can construct the spatial, and line of sight distributions of neutral hydrogen. This will allow for the properties of the first luminous sources to be inferred through their influence on the cosmological 21cm signal.

Most of the focus in the 21cm cosmological community has been on measuring either the one or two-point statistics of the signal. The one point statistic experiments determine the sky averaged quantities (for example the global mean temperature). For example: The Shaped Antenna measurement of the background Radio Spectrum 3 telescope (SARAS3, Nambissan et al., 2021); the Experiment to Detect the Global EoR Signature (EDGES, Bowman et al., 2018). The two point statistic experiments are primarily measured by radio interferometers. The current generation of radio interferometers includes the Murchison Widefield Array (MWA, Tingay et al., 2013; Wayth et al., 2018); Low-Frequency Array (LOFAR, van Haarlem, M. P. et al., 2013); Hydrogen Epoch of Reionisation Array (HERA, DeBoer et al., 2017); the New extension in Nancay upgrading LOFAR (NenuFAR, Zarka et al., 2012).

The two point statistic experiments calculate the PS of the 21cm signal which is the Fourier transform of the two point correlation function. This measures the Gaussianity or the variance of the signal as a function of comoving spatial scale. If the signal is entirely Gaussian this would capture all the information about the 21cm signal within cosmic variance². The signal is however expected to be highly non-Gaussian as it evolves during the Epoch of Heating (EoH) and the Epoch of reionisation (EoR) (Wyithe and Morales, 2007; Lidz et al., 2007). In the former the non-Gaussianities are driven by the appearance of the first luminous sources which heat the neutral IGM primarily through X-ray emission (Furlanetto, 2006). During this period, the strong emission from the first luminous sources drive above average temperature contrasts relative to the IGM. Eventually as X-ray heating progresses the medium saturates driving the non-Gaussianities to the matter density (Watkinson et al., 2018). During the latter stages of reionisation the spin temperature of the neutral hydrogen is expected to be saturated

²Since we cannot truly measure the ensemble average power spectrum, we can only estimate it over some volume. Each independent realisation therefore is a random sample of the true PS with some *cosmic* variance.

($T_S \gg T_{\text{CMB}}$), and therefore the non-Gaussianities are largely driven by the ionisation topology (Hutter et al., 2019). Analytical estimates of the characteristic size of the ionisation topology (bubbles) around individual luminous sources are ~ 10 cMpc during the late time period of reionisation (Wyithe and Loeb, 2004; Furlanetto and Oh, 2005; Zahn et al., 2007). However, Lin et al. (2016) showed that the characteristic size is underestimated and is closer to $\sim 20 - 100$ cMpc. Lin et al. (2016) and Giri et al. (2017) demonstrate the difficulty of determining the characteristic size from the complex 3D ionisation topology during reionisation.

Non-Gaussianity has been shown to be important in constraining the 21cm signal, in particular during reionisation, and could be important for confirming a detection of the 21cm signal (Shimabukuro et al., 2017). Non-Gaussianity has primarily been investigated by calculating the expected 21cm bispectrum. The bispectrum is the Fourier transform of the three point correlation function (Peebles, 1980), like the PS it probes the central third order moment as a function of spatial scales. The bispectrum offers a complementary picture of the 21cm signal especially during the EoR (Bharadwaj and Pandey, 2005; Majumdar et al., 2018). Watkinson et al. (2018) investigated the non-Gaussianity due to X-ray heating from stellar sources and high mass X-ray binaries. Hutter et al. (2019) investigated the ionisation morphology and the effects on the non-Gaussianity of the 21cm signal during reionisation. Numerous studies have been conducted on the bispectrum and its sensitivity during reionisation (Bharadwaj and Pandey, 2005; Yoshiura et al., 2015; Shimabukuro et al., 2016, 2017; Watkinson et al., 2017; Majumdar et al., 2018; Mondal et al., 2021; Majumdar et al., 2020). Trott et al. (2019) measure the bispectrum of MWA data, looking at a gridded and non-gridded estimator, establishing upper limits. Watkinson et al. (2020) looked at the expected foreground bispectrum, commenting on the detectability of the 21cm bispectrum in the presence of foreground systematics. More recently Tiwari et al. (2022) showed that the bispectrum can help constrain reionisation parameters.

The bispectrum has low signal to noise relative to the PS, and is computationally intensive to measure, even with the fast Fourier transform method of Watkinson et al.

(2017). To mitigate the difficulties related to computation and sensitivity, much focus has been on calculating the equilateral bispectrum (for example: Bharadwaj and Pandey, 2005; Yoshiura et al., 2015; Watkinson et al., 2018), as well as the squeezed bispectrum, which compresses one of the triangle mode sides (Chiang et al., 2014; Mondal et al., 2021). Other work has investigated the higher order one point statistics of the simulated 21cm signal, due to X-ray heating and Ly- α coupling (for example: Watkinson and Pritchard, 2015; Ross et al., 2019). The CMB cosmology community has investigated alternatives that probe the non-Gaussianity through the cross spectrum of quadratic temperature fields with the temperature field (Cooray, 2001). This is called the skew spectrum (SS), and is a collapsed form of the bispectrum, compressing the information into a pseudo PS as a function of one wavenumber (Fourier modes (k)) (Regan, 2017). Generalised in Szapudi and Szalay (1997); Munshi et al. (1998) and first used by Cooray (2001), it is now gaining interest in the 21cm community with the release of Ma and Peng (2023) at the time of writing this paper. Again drawing on the CMB cosmology community for inspiration, Dai et al. (2020) investigated what information can be gained by combining the PS and the SS. They found that the SS in conjunction with the PS offered increased constraints on cosmological parameters. The SS promises to have better signal to noise than the bispectrum because it integrates over all bispectrum triangle configurations for a given Fourier mode k . Additionally the SS can be directly compared to the PS because they can be measured at the same Fourier modes. The quadratic field cross correlation approach also makes it easy to measure the SS from simulations without having to first calculate the bispectrum.

In this work we use the updated version of `MERAXES`, which couples a semi-analytic galaxy formation model with a semi-numerical reionisation simulation to provide a realistic population of galaxies which can interact with the IGM through a variety of feedback effects. These feedback effects include supernovae, AGN, and photoheating, along with the infall/accretion of gas (Mutch et al., 2016; Qin et al., 2017; Qiu et al., 2019). Balu et al. (2023) updated `MERAXES` to include X-ray heating and spin temperature evolution for the semi-analytic galaxy formation model. Additionally

in Balu et al. (2023), the halo merger trees for the $210 h^{-1}$ Mpc simulations were augmented to include all atomically cooled galaxies out to $z = 20$ ($\sim 2 \times 10^7 h^{-1} M_{\odot}$). We build on the earlier work of Balu et al. (2023) by performing additional simulations varying the ionisation morphology through changing the minimum mass threshold for galaxies hosting star-formation, the X-ray luminosity and minimum energy threshold for X-rays escaping their host galaxies. Combined, these simulations enable the exploration of the cosmic evolution of the 21cm signal using a realistic population of galaxies from the cosmic dawn down to the completion of the EoR. These are ideal for studying the non-Gaussianity of the EoH and the EoR using the SS.

The paper is outlined as follows; in section 7.2 we define the PS and SS. In section 7.3 we briefly describe the simulations performed in this work. Section 7.4 presents the thermal and ionisation history of each simulation as well as the statistics as a function of redshift. Section 7.5 presents the PS, the SS and the normalised SS during the EoR for each simulation. Section 7.6 discusses the detectability of the normalised SS for the future SKA_LOW radio interferometer. We discuss and draw conclusions from the results in Section 7.7. The cosmology used in this work is defined by Planck Collaboration et al. (2021): $h = 0.68$, $\Omega_m = 0.31$, $\Omega_b = 0.048$, $\Omega_{\Lambda} = 0.69$, $\sigma_8 = 0.81$, and $n_s = 0.96$. All cosmological scales are in comoving units.

7.2 Power Spectrum and Skew Spectrum

In this section we review the PS and the SS, and how they are calculated from simulation volumes.

7.2.1 Power Spectrum

The PS is the Fourier transform of the two point correlation function, and probes the Gaussianity of a random field as a function of spatial scale k (Peebles, 1980):

$$\langle \tilde{\delta}_T(\mathbf{k}) \tilde{\delta}_T(\mathbf{k}') \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') P(k). \quad (7.1)$$

In Equation 7.1 the angular brackets $\langle \rangle$ denote the ensemble average over different realisations of the Universe. The Dirac delta $\delta_D(\mathbf{k} + \mathbf{k}')$ restricts the average to uncorrelated modes, and $P(k)$ is the spherically averaged power spectrum. $\delta_T(\mathbf{x})$ is the brightness temperature (or density) contrast $\delta_T(\mathbf{x}) = (T(\mathbf{x}) - \bar{T}(\mathbf{x}))/\bar{T}(\mathbf{x})$, and $\bar{T}(\mathbf{x})$ is the mean temperature. $\tilde{\delta}_T(\mathbf{k})$ is the three dimensional Fourier transform of $\delta_T(\mathbf{x})$, defined by:

$$\tilde{\delta}_T(\mathbf{k}) = \frac{V}{N_{\text{pix}}} \sum \delta_T(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}}. \quad (7.2)$$

We define the dimensionless PS:

$$\Delta_T^2(k, z) \equiv \frac{k^3}{(2\pi^2)V} \overline{\delta T_b^2} \langle \tilde{\delta}_T(\mathbf{k}, z) \tilde{\delta}_T^*(\mathbf{k}, z) \rangle \text{mK}^2, \quad (7.3)$$

where the angular brackets now denote the incoherent average in a spherical shell of width $\Delta \log k$, and $\tilde{\delta}_T^*(\mathbf{k}, z)$ is the conjugate transpose of $\tilde{\delta}_T(\mathbf{k}, z)$.

7.2.2 Bispectrum and Skew Spectrum

The bispectrum of the 21cm brightness temperature fluctuations is defined as:

$$\langle \tilde{\delta}_T(\mathbf{k}_1) \tilde{\delta}_T(\mathbf{k}_2) \tilde{\delta}_T(\mathbf{k}_3) \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B(k_1, k_2, k_3), \quad (7.4)$$

again the angular brackets and the Dirac delta function denote the ensemble average. For the bispectrum the ensemble average is over all triplet values that satisfy the closed triangle condition $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$. Without loss of generality we let $\mathbf{k}_1 = \mathbf{k}$, $\mathbf{k}_2 = \mathbf{q}$, and $\mathbf{k}_3 = -(\mathbf{k} + \mathbf{q})$. The SS is the integral of the bispectrum $B(k, q, |\mathbf{k} + \mathbf{q}|)$ over all possible triangle configurations, for a fixed triangle side k :

$$S_\gamma(k) = \frac{1}{(2\pi)^3} \int_{\mathbb{R}} d^3q B(k, q, |\mathbf{k} + \mathbf{q}|). \quad (7.5)$$

It can be shown that Equation 7.5 is equivalent to the cross spectrum of the mean subtracted squared temperature field, to the temperature field:

$$\langle \tilde{\delta}_{T^2}(\mathbf{k}) \tilde{\delta}_T(\mathbf{k}') \rangle \equiv (2\pi)^3 \delta_D(\mathbf{k} + \mathbf{k}') S_\gamma(k), \quad (7.6)$$

Similarly to the PS definition in Equation 7.1 the delta function and the angular brackets denote the ensemble average. We define the Fourier transform of the squared temperature field $\tilde{\delta}_{T^2}(\mathbf{k})$ below:

$$\tilde{\delta}_{T^2}(\mathbf{k}) = \frac{V}{N_{\text{pix}}} \sum (\delta_T(\mathbf{x}))^2 e^{-i\mathbf{k}\cdot\mathbf{x}}. \quad (7.7)$$

Similar to Equation 7.3 we can define the dimensionless SS:

$$\Delta_{T^2,T}^2(k, z) \equiv \frac{k^3}{(2\pi^2)V} \overline{\delta T_b^3} \langle \tilde{\delta}_{T^2}(\mathbf{k}, z) \tilde{\delta}_T^*(\mathbf{k}, z) \rangle. \quad (7.8)$$

We calculate the SS by performing a three dimensional Fourier transform of the mean subtracted and squared temperature field. We then take the product of this with the conjugate of the Fourier transform of the temperature field, and then average in spherical shells of width $\Delta \log k$. For consistency and for comparison we use the same bins for calculating the SS and the PS throughout this work.

7.3 MERAXES

In order to simulate the cosmic 21cm signal, we use the MERAXES (Mutch et al., 2016) semi-analytical galaxy formation and evolution model. In this section, we give a brief summary of MERAXES and refer the reader to other relevant works for further details.

7.3.1 A Realistic Galaxy Population

We make use of the L210_N4320 dark matter-only simulation of the GENESIS suite of N -body simulations (Power et al. in prep). L210_N4320 has 4320^3 dark matter particles in a cubical volume of side length $L = 210h^{-1}$ Mpc achieving a halo mass resolution of $\sim 5 \times 10^8 M_\odot$.

The halo merger trees from L210_N4320 were further ‘augmented’ to a halo mass

resolution of $\sim 3 \times 10^7 M_{\odot}$, the atomic cooling limit at $z = 20$, using the Monte-Carlo algorithm code `DARKFOREST` (Qiu et al., 2020). This is achieved by sampling low-mass haloes from a conditional halo mass function that is based on the extended Press-Schechter theory (Bond et al., 1991; Bower, 1991; Lacey and Cole, 1993), after modifications to match the N -body simulations’ halo mass functions (HMFs). These haloes are then ‘grafted’ onto the L210_N4320 merger tree in a manner such that the final augmented HMFs agree with those from high resolution N -body simulations (see Figure 2 of Balu et al., 2023). `DARKFOREST` also assigns and evolves the positions of the newly added haloes using the local halo density field and the linear continuity equation (see Qiu et al., 2020; Balu et al., 2023, for further details). We therefore effectively create an N -body simulation that has a statistically complete galaxy population down to redshift $z = 20$; we deploy `MERAXES` on this augmented simulation.

The goal of `MERAXES` is to simulate the growth and evolution of galaxies during the EoR in a self-consistent manner. This is achieved through detailed and physically motivated prescriptions for varied astrophysical phenomena such as radiative cooling of gas, star formation, supernovae (SNe) and active galactic nuclei feedback, and mergers (Mutch et al., 2016; Qin et al., 2017; Qin et al., 2017; Qiu et al., 2019).

For each simulation snapshot, a dark matter halo increases its baryonic mass in proportion to the universal cosmic baryonic fraction. This mass is added to a ‘hot gas’ reservoir from where it can cool down to form a ‘cold gas disk’. Following a star formation prescription based on the Kennicutt-Schmidt law (Kennicutt, 1998), stars are created out of this cold disk when a cold mass threshold is reached. The cadence of our simulation is constructed so that the longest time-step is ~ 16 Myr. Hence newly formed massive stars can go SNe in the same time-step and less massive stars can survive for a few snapshots. `MERAXES` therefore, has implementations for both instantaneous and delayed SNe feedback. The primary impact of SNe is to heat up the cold gas in a galaxy. SNe therefore move a portion of the cold gas to the hot halo and in very extreme energetic cases can even remove the gas from the galaxy altogether.

The amount of stellar mass in a galaxy fixes the amount of ionising UV and X-ray

photons that it produces. Once the local environment of a galaxy is ionised, the cooling properties of the IGM are affected. MERAXES couples reionisation feedback and galaxy growth by self-consistently modifying the amount of gas that is accreted onto a galaxy depending on the local UV background and the local IGM ionisation state. The UV escape fraction $f_{\text{esc}} (\leq 1)$ of the galaxies is a power-law in redshift z (also see Section 7.3.4.1):

$$f_{\text{esc}} = f_{\text{esc},0} \left(\frac{1+z}{6} \right)^{\alpha_{\text{esc}}}, \quad (7.9)$$

where $f_{\text{esc},0}$ is the escape fraction normalisation and $\alpha_{\text{esc}} = 0.20$ is the power-law index.

In this work, we adopt the same fiducial simulation as Balu et al. (2023), L210_AUG (hereon labelled Fiducial). This simulation has been calibrated with respect to the UV luminosity functions and the colour-magnitude relation within a rigorous Bayesian framework (Qiu et al., 2019), as well as the stellar mass functions (Balu et al., 2023), at $z \sim 4 - 10$. Reionisation parameters were tuned such that the reionisation history is consistent with existing measurements of the IGM neutral fraction and the CMB optical depth (see Balu et al., 2023, in particular Figure 3 and Table 2).

7.3.2 IGM Ionisation State

MERAXES computes the IGM ionisation, following the semi-numerical code 21CMFAST (Mesinger et al., 2011), via an excursion-set formalism (Furlanetto et al., 2004a). First, we grid the simulation volume and assign galaxies to the voxels based on their positions. We subdivide our simulation volume into 1024^3 cells, corresponding to a cell size of $\sim 0.2 h^{-1}$ Mpc. In spheres of decreasing radii, we compare the number of ionising photons from the stellar baryons and the total baryons in the IGM. After accounting for recombinations that can happen in the densest parts of the IGM, we flag a cell as ionised when the number of ionising photons is higher than that of the neutral baryons.

$$N_{\text{b}*}(\mathbf{x}, z|R) N_{\gamma} f_{\text{esc}} \geq N_{\text{atom}}(\mathbf{x}, z|R) (1 + \bar{n}_{\text{rec}}) (1 - \bar{x}_e), \quad (7.10)$$

where $N_{b*}(\mathbf{x}, z|R)$ is the number of stellar baryons in a sphere of radius R centred at \mathbf{x} and redshift z , $N_\gamma = 4000$ is the number of UV ionising photons per baryon (Barkana and Loeb, 2007), N_{atom} is the number of neutral HI in the same volume, \bar{n}_{rec} is the average number of recombinations in the IGM (Sobacchi and Mesinger, 2014), and \bar{x}_e is the mean electron fraction accounting for the secondary ionisations caused by X-ray photons. Motivated by the mean-free path of a typical UV photon in the IGM (Songaila and Cowie, 2010), we set the maximum of $R = 50$ Mpc and decrease it successively down to the size of a cell.

7.3.3 21cm Signal

The differential brightness temperature of the 21cm emission from a cloud of HI gas illuminated by CMB radiation of temperature T_{CMB} is given by:

$$\begin{aligned} \delta T_b(\nu) &= \frac{T_S - T_{\text{CMB}}}{1+z} (1 - e^{-\tau_{\nu_0}}) \\ &\approx 27 x_{\text{HI}} (1 + \delta_{\text{nl}}) \left(\frac{H}{dv_r/dr + H} \right) \left(1 - \frac{T_{\text{CMB}}}{T_S} \right) \\ &\quad \times \left(\frac{1+z}{10} \frac{0.15}{\Omega_M h^2} \right)^{1/2} \left(\frac{\Omega_b h^2}{0.023} \right) \text{ mK}, \end{aligned} \quad (7.11)$$

where T_S is the IGM spin temperature which determines the energy level populations of the HI hyperfine states, τ_{ν_0} is the optical depth, $\delta_{\text{nl}} \equiv \rho/\bar{\rho} - 1$ is defined as the evolved Eulerian density contrast (ρ is the density), $H(z)$ is the Hubble parameter, dv_r/dr is the line-of-sight co-moving velocity gradient, and x_{HI} is the neutral fraction. MERAXES sources the density and the velocity fields from the N-body simulations and creates self-consistent T_S and x_{HI} fields.

7.3.3.1 Spin Temperature

As can be seen from the Equation (7.11), the spin temperature T_S of the IGM plays a major role in the 21cm signal. The level populations of the HI hyperfine states depend on a number of physical processes in the IGM, including the amount and the energy of

the UV and X-ray photons. T_S is related to the UV and X-ray emission via:

$$T_S^{-1} = \frac{T_{\text{CMB}}^{-1} + x_\alpha T_\alpha^{-1} + x_c T_K^{-1}}{1 + x_\alpha + x_c}, \quad (7.12)$$

where x_α and x_c are the Wouthuysen-Field coupling (Wouthuysen, 1952; Field, 1958) and the collisional coupling coefficients respectively. x_c is computed by taking into account the collisions of HI atoms amongst themselves as well as with free electrons and protons in the IGM (Zygelman, 2005; Furlanetto and Furlanetto, 2006). x_α depends on the local Ly α background flux and closely follows the implementation in Mesinger et al. (2011). T_α is the ‘colour’ temperature, T_K is the kinetic temperature of the IGM, and we assume $T_\alpha = T_K$ (Field, 1959a).

The spin temperature field is therefore very sensitive to the T_K , which is impacted by X-ray heating. The evolution of the T_K depends on the angle-averaged X-ray intensity $J(\mathbf{x}, E, z)$ which is computed as a function of the position \mathbf{x} , X-ray photon energy E , and redshift z :

$$J(\mathbf{x}, E, z) = \frac{(1+z)^3}{4\pi} \int_z^\infty dz' \frac{cdt}{dz'} \epsilon_X e^{-\tau}, \quad (7.13)$$

where we have integrated the comoving X-ray emissivity ϵ_X back along the lightcone, and $e^{-\tau}$ accounts for the probability that an X-ray photon emitted at redshift z' survives till z . We compute ϵ_X as a function of the position \mathbf{x} , X-ray photon energy $E_e = E(1+z')/(1+z)$ at the emitted redshift z' :

$$\epsilon_X(\mathbf{x}, E_e, z') = \frac{L_X}{\text{SFR}} \times \text{SFRD}(\mathbf{x}, E_e, z'), \quad (7.14)$$

where L_X/SFR is the galaxies’ specific X-ray luminosity per unit star formation rate (SFR), and SFRD is the star formation rate density. We assume a power-law in X-ray photon energy E , $L_X/\text{SFR} \propto E^{-\alpha_X}$, with $\alpha_X = 1$ which is consistent with observations of high mass X-ray binaries in the local Universe (Mineo et al., 2012; Fragos et al., 2013; Pacucci et al., 2014), and is normalised:

$$L_{X < 2 \text{ keV}}/\text{SFR} = \int_{E_0}^{2 \text{ keV}} dE_e L_X/\text{SFR}, \quad (7.15)$$

Table 7.1: Astrophysical parameter summary for the seven MERAXES simulations. See text for details.

Name	Minimum Halo Mass (M_{thresh}) [M_{\odot}]	$f_{\text{esc},0}$	X-ray Luminosity ($L_{X<2\text{keV}}/\text{SFR}$) [$\text{erg s}^{-1} M_{\odot}^{-1} \text{yr}$]	X-ray threshold (E_0) [keV]	Comments
Fiducial	3×10^7	0.14	3.16×10^{40}	0.5	Fiducial simulation
mid_ M_h	10^9	0.25	3.16×10^{40}	0.5	Intermediate halo mass threshold simulation
high_ M_h	10^{10}	0.45	3.16×10^{40}	0.5	High halo mass threshold simulation
low_ L_X	10^8	0.14	3.16×10^{38}	0.5	Low X-ray luminosity simulation
high_ L_X	10^8	0.14	3.16×10^{42}	0.5	High X-ray luminosity simulation
low_ E_0	10^8	0.14	3.16×10^{40}	0.2	Low X-ray energy threshold simulation
high_ E_0	10^8	0.14	3.16×10^{40}	1	High X-ray energy threshold simulation

where $L_{X<2 \text{ keV}}/\text{SFR}$ is the soft-band X-ray luminosity per SFR in units of ($\text{erg s}^{-1} M_{\odot}^{-1} \text{yr}$), and E_0 fixes the minimum energy for an X-ray photon so that it is not absorbed within the galaxy.

7.3.4 Simulations

To aid our physical interpretation of the features present in the SS we run a further set of six simulations with MERAXES, in addition to our fiducial simulation.

In particular, we are interested in the physical processes which impact the morphology of the 21cm signal. To explore the impact of the ionisation morphology we vary the minimum mass for halos hosting star formation. Setting the UV escape fraction to zero in galaxies below a given mass threshold alters the size and distribution of the ionised regions (i.e. produces larger, more isolated bubbles for an increasing mass threshold). With regard to the heating morphology, we vary the X-ray luminosity and the minimum energy threshold for X-rays escaping their host environment. Increasing the X-ray energy threshold decreases the prevalence of bubbles of heated IGM gas transitioning toward an effective uniform background of IGM heating. Table 7.1 summarises the simulations used in this work along with the values of the parameters that are varied.

7.3.4.1 Halo Mass Threshold

To explore the impact of the minimum halo mass on the EoR morphology, we modify the f_{esc} (also see Equation 7.9 and Section 7.3.2) prescription as follows:

$$f_{\text{esc}} = \begin{cases} f_{\text{esc},0} \left(\frac{1+z}{6} \right)^{\alpha_{\text{esc}}} & , \quad M_{\text{halo}} \geq M_{\text{thresh}} \\ 0, & M_{\text{halo}} < M_{\text{thresh}}. \end{cases} \quad (7.16)$$

We run simulations with mass thresholds $M_{\text{thresh}} = 10^9$ and $10^{10} M_{\odot}$, and label them *mid_* M_h and *high_* M_h respectively (we point out that our Fiducial simulation contains all haloes down to $3 \times 10^7 M_{\odot}$). By considering an increasing halo mass threshold, we effectively decrease the total number of galaxies capable of contributing to reionisation. To compensate for the loss of ionising sources we increase the UV escape fraction of those remaining star-forming galaxies to ensure a reionisation history consistent with our observational constraints. We therefore, increase the UV escape fraction normalisation $f_{\text{esc},0}$ (see Equation 7.9) to $[0.25, 0.45]$ respectively (see Fig 7.2). These simulations can thus probe the impact of ionisation morphology in the PS and SS. An increasing mass threshold should result in larger ionised regions (changing the physical location of features in the PS/SS). We emphasise that we still populate and evolve the galaxies in the haloes below the mass threshold, and these galaxies can start contributing to the UV ionisation budget when their host halo mass passes the threshold. We point out that we do not suppress emission of X-ray photons by these galaxies. In this manner, we fix the X-ray background across these simulations (Fiducial, *mid_* M_h and *high_* M_h) to be the same. Keeping the X-ray background constant for each simulation was a deliberate design choice. This isolates the changes in ionisation morphology to the halo mass threshold (larger halos resulting in larger features for example). With a consistent ionisation history, and X-ray heating background, the only change between simulations, and therefore on the 21cm statistics comes from the differences in the ionisation morphology.

7.3.4.2 X-Ray Luminosity

We also consider two simulations with a lower and higher X-ray luminosity $L_{X < 2 \text{ keV}} / \text{SFR} = [3.16 \times 10^{38}, 3.16 \times 10^{42}] \text{ ergs s}^{-1} M_{\odot}^{-1} \text{ yr}$ as compared to the fiducial value of $3.16 \times$

10^{40} ergs $s^{-1} M_{\odot}^{-1}$ yr. These two simulations, labelled low_L_X and $high_L_X$ correspond to the L210_AUG_LOWX and L210_AUG_HIGHX simulations in Balu et al. (2023). These simulations cover a range of X-ray luminosities per SFR, one order of magnitude broader than what is observed in the local soft band X-ray luminosity (Mineo et al., 2012; Fialkov et al., 2016) based on the range adopted in Greig and Mesinger (2017).

For low_L_X , X-ray heating is inefficient and the IGM ionises before it is heated (21cm signal always remains in absorption). This produces large temperature contrasts between the ionised and neutral regions resulting in much higher amplitudes for the 21cm statistics.

7.3.4.3 X-ray Energy Threshold

We explore the impact of the X-ray photon energy threshold E_0 by producing two simulations with $E_0 = 0.2$ keV and $E_0 = 1$ keV compared to $E_0 = 0.5$ keV for the Fiducial simulation. Decreasing the energy threshold, coupled with our power-law X-ray spectral energy distribution, results in a higher fraction of softer X-ray photons. As softer photons have shorter mean free paths, more heat energy is deposited closer to the host galaxies resulting in more prevalent bubbles of heating around the first galaxies. Increasing the energy threshold removes this heating morphology as the X-ray photons now penetrate much deeper into the IGM before depositing their heat energy resulting in an effective uniform background of heating. In effect, varying this energy threshold will alter the amplitude of the 21cm statistics during the heating epoch (see e.g. Pacucci et al., 2014; Greig and Mesinger, 2017).

7.4 Results

In this section, we analyse the thermal cosmological history of the 21cm brightness temperature signal for the simulation sets. We plot 2D slices of the lightcone boxes as a function of redshift for each simulation. We also calculate the neutral fraction (\bar{x}_{HI}) for each simulation coeval box as a function of redshift, and compare the results of




Figure 7.1: 21cm differential brightness temperature lightcone slices for each simulation volume as a function of redshift from $z = 20 - 5$. The top slice is the Fiducial simulation (labelled), and the bottom slice is the high_ E_0 simulation. The colour bar is a symmetric log-scale, where blue indicates absorption and red is emission relative to the CMB. Each lightcone slice is fixed to the same temperature scale.

each simulation. We then calculate the PS and SS for each simulation as a function of redshift for a set of three different spatial scales from large to small scales, comparing the results.

7.4.1 21cm Lightcones

Fig 7.1 shows a lightcone slice of each simulation as a function redshift. In descending order the panels show the Fiducial, mid_ M_h , high_ M_h , low_ L_X , high_ L_X , low_ E_0 , and high_ E_0 simulations. The colour bar is a log symmetric colour map, where blue indicates absorption, and red indicates emission relative to the CMB. Orange indicates zero temperature difference. During reionisation $\delta T_b = 0$ is typically associated with




Figure 7.2: The ionisation history of the average neutral hydrogen (neutral fraction \bar{x}_{HI}) IGM calculated for each simulation (labelled).

regions that are ionised.

The same N -body dark matter particle `GENESIS` simulations are used to generate each of the different `MERAXES` simulations. Therefore each simulation has the same dark matter halo distribution. In Fig 7.1 this is evident at high redshifts ($z \lesssim 20$) and during reionisation ($z \lesssim 8$) in the location and approximate size of the first ionisation regions. There are some obvious differences in the temperature contrast due to the different X-ray heating parameters. Of note, we see that the `low_Lx` is always in absorption, even during reionisation, and the `high_Lx` simulation is heavily preheated at high redshift. The `low_E0` simulation has small regions of localised heated gas that appear in emission at $z > 15$. In contrast, the `high_E0` simulation results in a more uniform heating of the IGM and thus the brightness temperature is relatively featureless. There is also a clear difference in the size of ionisation regions between the Fiducial, `mid_Mh`, and `high_Mh` simulations, with the size increasing from Fiducial to `high_Mh` at fixed redshift. We discuss these features in the context of the statistics in the following subsections.

7.4.2 Neutral Fraction and Ionisation

In Fig 7.2 we show the average neutral fraction (\bar{x}_{HI}) calculated for each simulation coeval box as a function of redshift. The Fiducial model was calibrated to match existing observational constraints (see Balu et al. (2023) for details). The halo mass threshold simulations have increased UV escape fractions as a function of halo mass to ensure similar reionisation histories for easier comparison of the ionisation morphology. The solid black line is the average neutral fraction for the Fiducial simulation. By $z = 10$ the Fiducial simulation is already partially ionised at the ~ 5 percent level. We note that the ionisation history for the low_ E_0 (purple dashed line), the high_ E_0 (light purple dash dotted line), and the low_ L_X (light green dash dotted line) simulations are effectively identical to the Fiducial case. This is expected, since ionisation is predominantly driven by UV photons not X-ray emission, additionally these simulations also have the same halo mass thresholds and escape fraction ($f_{\text{esc},0}$) as the Fiducial simulation (see Table 7.1). In contrast the high_ L_X (dark green dashed line) simulation undergoes reionisation early relative to the Fiducial, due to the increase in the number of ionisations following secondary collisions of the X-ray photons. The high_ L_X simulation has a ionisation fraction of ~ 10 percent reionisation by $z = 10$. This is not unexpected, as shown by Mesinger et al. (2013) X-ray emission can result in a maximum of up to 10 percent ionisation before the EoR. For the high_ L_X simulation, which has the largest X-ray luminosity emission, we would expect to see a high pre-ionisation of the IGM, prior to the EoR.

The mid_ M_h (dark blue dashed line) and the high_ M_h (light blue dash dotted line) simulations begin reionisation later than the Fiducial simulation, as it takes longer for haloes to gravitationally grow in excess of their respective mass thresholds to emit ionising UV photons. Nevertheless, the mid_ M_h and the high_ M_h simulations neutral fraction profiles result in a similar ionisation history to the Fiducial. This is due to the effective parameterisation of the escape fraction relative to the halo mass threshold. To compensate for the loss of ionising sources as a function of increasing the halo mass threshold, the UV escape fraction was increased proportional to the halo mass




Figure 7.3: Mean brightness temperature for all simulations, calculated from the coeval boxes as a function of redshift.

threshold. The $\text{mid-}M_h$ and $\text{high-}M_h$ simulations have UV $f_{\text{esc},0}$ values of 0.25, and 0.45, compared to the Fiducial with 0.14. Therefore, more ionising UV photons escape per unit mass from the same higher mass halos in $\text{high-}M_h$ and $\text{mid-}M_h$, for the same amount of star formation, compared to the Fiducial simulation. As reionisation progresses, star formation increases, this results in a more rapid (sharper) reionisation relative to the Fiducial simulation.

7.4.3 21cm Statistics

7.4.3.1 Mean Brightness Temperature

This section follows a similar approach to Balu et al. (2023), where we calculate the mean temperature for each coeval box as a function of redshift. Fig 7.3 shows the mean temperature for each simulation. As previously mentioned, the Fiducial model, $\text{low-}L_X$ and $\text{high-}L_X$ simulations are taken from Balu et al. (2023). We see that the mean temperature for these three simulations agree with those shown in Figure 8 of Balu et al. (2023), for more details on the $\text{low-}L_X$ and $\text{high-}L_X$ simulations we refer the reader to this work. For all simulations we see a characteristic absorption feature which occurs when Ly- α emission couples the spin temperature to the gas temperature.




Figure 7.4: The PS as a function of redshift for all simulations at a fixed spatial scale. Panel (a) shows the redshift evolution at $k \sim 0.05 \text{ Mpc}^{-1}$, panel (b) shows the redshift evolution at $k \sim 0.1 \text{ Mpc}^{-1}$, and panel (c) shows the redshift evolution at $k \sim 1 \text{ Mpc}^{-1}$.

As the gas expands adiabatically it cools relative to the CMB, increasing the relative absorption. For most simulations this absorption trough occurs at approximately $z \sim 15$ (with the exception of the `high_LX` simulation). The timing of the trough depends on X-ray heating which eventually drives the IGM into emission (with the exception of the `low_LX` simulation), occurring during reionisation at $z \gtrsim 10$. This culminates in a peak roughly at the midpoint of reionisation in the redshift range of $z \sim 6 - 8$.

As mentioned in the previous section the halo mass simulations both start reionisation later, but conclude earlier than the Fiducial. This delay results in higher heating and reionisation rates for the halo mass simulations. The higher ionisation rate and heating being directly associated with the larger escape fraction. This directly effects the mean temperature and the amplitude of the PS (as seen in Fig 7.4).

The X-ray energy threshold simulations follow a similar evolution to the Fiducial. The biggest difference is between the amplitude and onset of X-ray heating between $10 < z < 15$. Here the `low_E0` (purple dashed line) undergoes heating earlier than the Fiducial and the `high_E0` (light purple dash dotted line) simulations. This should be expected since there are relatively more lower energy soft X-rays available to heat the IGM. Additionally the timing of heating should be earlier for the `low_E0` simulation because the mean free path of X-ray photons is proportional to their energy, meaning softer X-rays deposit their energy into the IGM before harder X-rays.

7.4.3.2 21cm Power Spectrum

Fig 7.4 shows the PS calculated at spatial scales $k \sim 0.05$ (panel (a)), 0.1 (panel (b)), and 1.0 Mpc^{-1} (panel (c)), as a function of redshift for each coeval box, and each simulation. The lines in Fig 7.4 correspond to the same simulations in Fig 7.3. The features discussed in Fig 7.3 are broadly mirrored in Fig 7.4 (a) for most simulations. Most simulations have the characteristic three peak structure due to Ly- α pumping (Wouthuysen, 1952; Field, 1958), the EoH and the EoR (Furlanetto, 2006; Mesinger et al., 2013).

We find the Fiducial simulation in panel (a) of Fig 7.4 at $k \sim 0.05 \text{ Mpc}^{-1}$ has a peak at $z \sim 17$ during Ly- α pumping, a peak at $z \sim 12$ during the EoH, and the third and final peak at $z \sim 6.5$ during the mid point of reionisation. In panel (a) most simulations have the characteristic three peak structure with different timings $z \pm 1$, with the exception of the low_ L_X and high_ L_X simulations. In panel (b) at $k \sim 0.1 \text{ Mpc}^{-1}$ the peak during the EoH has mostly merged with the EoR but is still weakly present for some simulations.

We find the same features for the low_ L_X and high_ L_X simulations for panel (b) and (c) of Fig 7.4, as seen by Balu et al. (2023) for the L210_AUG_lowX and L210_AUG_highX simulations. The high amplitude of the low_ L_X simulation is due to the higher temperature contrast that results from the colder IGM (lower X-ray luminosity). The lower amplitude of the high_ L_X simulation is due to the higher X-ray luminosity heating, this reduces the temperature fluctuations on all scales. The inefficient heating due to lower X-ray luminosity in the low_ L_X simulation results in the merging of the EoH and EoR peaks. In the high_ L_X simulation we have four peaks, an early peak at $z \sim 22$ coincident with the relatively weak absorption trough (Ly- α pumping), a peak at $z \sim 18$ which corresponds to the EoH heating, with the IGM being in emission at this stage. There is a peak during the midpoint of reionisation at $z \sim 7$. In addition to the expected three peaks there is an additional peak at $z \sim 10$. This is due to the first ionisation sources, with the high_ L_X simulation having ~ 10 percent ionisation by redshift 10. The four peak structure of the high_ L_X simulation is weakly




Figure 7.5: The SS as a function of redshift for all simulations at a fixed spatial scale. Panel (a) shows the redshift evolution at $k \sim 0.05 \text{ Mpc}^{-1}$, panel (b) shows the redshift evolution at $k \sim 0.1 \text{ Mpc}^{-1}$, and panel (c) shows the redshift evolution at $k \sim 1 \text{ Mpc}^{-1}$.

maintained in panel (b) of Fig 7.4 at $k \sim 0.1 \text{ Mpc}^{-1}$.

For the high_ E_0 case at $k \sim 0.05 \text{ Mpc}^{-1}$ the first peak occurs at $z \sim 15$, which is roughly coincident with the absorption trough which reaches a minimum at $z \lesssim 15$. The second less prominent peak occurs during the EoH ($z \sim 8.5$). The lower amplitude and later occurrence of the peak during the EoH are a direct result of the higher X-ray energy threshold, which produce less structure in the IGM due to the more uniform heating of the longer mean-free path X-ray photons. The first patches of emission that are correlated on the largest scales appear later relative to the other simulations, we see similar behaviour in Figure 1 of Greig and Mesinger (2017) (bottom row of Figure 1). This delayed and more uniform heating means the Ly- α pumping dominates the amplitude of the PS for a longer period. The final peak occurs during reionisation at $z \sim 6$. At $k \sim 0.1 \text{ Mpc}^{-1}$ and $k \sim 1 \text{ Mpc}^{-1}$ we do not see the peak during the EoH, at these scales the peak has merged with the EoR due to the delayed start to heating and less efficient heating at smaller scales. Additionally, we have a lower amplitude during the EoR due to the more uniform IGM temperature contrast.

For the low_ E_0 simulation at a $k \sim 0.05 \text{ Mpc}^{-1}$ and $k \sim 0.1 \text{ Mpc}^{-1}$ we see the three peaked structure. The peak due to Ly- α occurs early at $z \sim 17$, with the peak during the EoH occurring at $z \sim 12$ which coincides with the pockets of emission seen in Fig 7.1. The amplitude of the EoH is greater for the low_ E_0 simulation due to the lower X-ray energy threshold. These lower energy X-rays efficiently heat the local medium around the first luminous sources, producing higher temperature contrasts in

the IGM due to the more inhomogeneous heating. This produce more inhomogeneous structures, increasing the overall power. These correlate on the largest scales resulting in a strong peak. The high temperature contrast regions are the first to ionise, thus during the EoR, the relative amplitude of the EoR peak returns to that of the Fiducial model.

The mid_ M_h and the high_ M_h simulations in panel (a) have broadly the same PS, differing at most in amplitude during the EoH and the EoR, with the high_ M_h simulation peaking first. The mid_ M_h and high_ M_h simulations compared to the Fiducial have reionisation topologies driven by larger (and more biased) sources. This results in an increase in the amplitude of the 21cm PS. The PS of both simulations are broadly the same as the Fiducial simulation, only significantly deviating at $z \sim 10$ when reionisation begins to become more significant for the Fiducial simulation. For the mid_ M_h and the high_ M_h simulation the power stays relatively flat around $z \sim 10$, this is reflected by the delayed start to reionisation in Fig 7.2 for these two simulations.

7.4.3.3 21cm Skew Spectrum

Fig 7.5 shows the SS calculated at $k \sim 0.05$ (panel (a)), 0.1 (panel (b)), and 1.0 Mpc^{-1} (panel (c)), as a function of redshift for each coeval box, and each simulation. The lines correspond to the same simulations in Fig 7.3 and 7.4. There are some interesting features in panel (a) at $k \sim 0.05 \text{ Mpc}^{-1}$ that appear during the EoH for the Fiducial, low_ E_0 , mid_ M_h and high_ M_h simulations. This peak occurs at $z \sim 12.5$ for the Fiducial, mid_ M_h and high_ M_h simulations, and at $z \sim 14$ for the low_ E_0 simulation, and each peak has a width of $\Delta z \approx 1.5^3$. We also see a similar peak at $z \sim 8$ for the low_ L_X simulation during the EoH/EoR which overlap due to the low X-ray heating efficiency. Notably, this peak has a shorter width of $\Delta z \lesssim 1$ but a relatively higher amplitude, likely due to the larger contrast in temperature since the low_ L_X simulation is in absorption.

The peaks are correlated with the appearance of localised regions of above average

³We note that we do not expect to detect this peaked feature at $k \sim 0.05 \text{ Mpc}^{-1}$ with the current and future radio interferometers without performing wedge removal.

temperature contrast around the first luminous sources. This is most apparent in the low_ E_0 simulation, where in Fig 7.1 at $z \sim 14$ there are small patches (on the scale of a few pixels) close to the zero temperature, which is large compared to the mean IGM temperature at this redshift (which is in absorption). Since the first luminous sources are the most massive they are distributed on the largest scales, and thus this feature is only seen at $k \sim 0.05 \text{ Mpc}^{-1}$. The timing of the peak depends on the X-ray energy threshold E_0 . Since softer X-rays deposit energy closer to the host galaxy, the peak appears earlier for lower E_0 values. The slight difference in timing between the Fiducial and the halo mass simulations can not be related to X-ray emission. For the halo mass simulations, galaxies in halos below the threshold still emit X-ray emission. Therefore, the X-ray background is essentially constant for the Fiducial, and halo mass simulations. The observed difference is the result of early star formation, and extra heating from UV emission around the lower mass haloes in the Fiducial simulation. For the low_ L_X simulation, the peak does not occur until significantly later, as the contrast must occur between the ionised sources and the cold IGM. We do not see this peak in panels (b) or (c). A similar peak during the EoH is seen in Figure 5 of Watkinson et al. (2018), where they investigated the normalised equilateral bispectrum of high mass X-ray binaries and stellar sources. They find a peak at $z \sim 20$ and $z \sim 16$ in the PS, and a peak at $z \sim 18$ in the normalised equilateral bispectrum, at the scale $k \sim 0.05 \text{ Mpc}^{-1}$.

In Ma and Peng (2023) they investigate the SS during the EoH and the EoR with 21CMFAST. In Figure 3 and 4 of Ma and Peng (2023), they display the SS as a function of redshift at different fixed scales. They find two peaks in the SS, one associated with X-ray heating which couples the spin temperature to the matter density, and a final peak during reionisation from the remaining neutral islands that are in emission. They find a negative dip between both peaks which they associate with the appearance of the first bubbles, which correlates with the ionisation fraction (neutral fraction).

Our simulations show different features. We find three peaks for the Fiducial, low_ E_0 , mid_ M_h and high_ M_h simulations at $k \sim 0.05 \text{ Mpc}^{-1}$ ⁴ in panel (a) of Fig

⁴We note that we do not expect to detect this peaked feature at $k \sim 0.05 \text{ Mpc}^{-1}$ with the current and future radio interferometers without performing wedge removal.




Figure 7.6: The dimensionless SS plotted against the dimensionless PS as a function of redshift for each simulation at a fixed spatial scale $k \sim 0.1 \text{ Mpc}^{-1}$. Each curve indicates the co-evolution of the PS and SS amplitude as a function of redshift, with the right pointed triangle, indicating the start point at $z = 20$, and the left pointing triangle indicating the endpoint at $z = 5$. The grey line indicates the expected evolution if the Fiducial simulation $S_\gamma(k) = (-P^{3/2}(k)k^{3/2}, P^{3/2}(k)k^{3/2})$.

7.4. Importantly we do not find these features on the smaller scales in panels (b) and (c), instead we see a single peak and trough for the Fiducial, mid_ M_h , high_ M_h and high_ E_0 simulations. There is a weaker trough and peak structure for the low_ E_0 , low_ L_X and high_ L_X simulations, which correspond to the other peaks/troughs in panel (a). We do not find a negative SS at the start of reionisation as suggested by Majumdar et al. (2018); Ma and Peng (2023). Majumdar et al. (2018) investigated the ionisation bubble bispectrum, which showed a negative amplitude for their toy model. Instead of a negative dip, we do however find a dip at $z \sim 7$ for the Fiducial simulation which is approximately 30 percent ionised at this point. We likewise see similar dips for the other simulations at similar redshifts during the EoR.

The key distinction between this work and Ma and Peng (2023), is that `MERAXES` incorporates a semi-analytic galaxy formation model with reionisation. In this prescription galaxies have a finite lifetime, and can cease star formation due to a variety of feedback mechanisms. In `21CMFAST` however, once a cell contains sufficient mass to contribute to star formation, it remains on indefinitely. This results in a slower build up of heating and ionisation in `MERAXES` compared to `21CMFAST`. Thus the EoH occurs more slowly and later than `21CMFAST`, effectively decreasing the time between the EoH and the EoR. As a result of the slower rate of heating, the mean temperature in Ma and Peng (2023) is larger in amplitude between absorption and emission compared to our

mean temperature profile. The slower rate of change in the mean temperature (slower heating) in our simulations, leads to the blending or overlapping as seen in other works (for example Greig and Mesinger, 2017).

Another interesting feature in Fig 7.5 is the rapid transition to a positive SS, occurring within one coeval redshift box (a cosmic blink). This happens in all simulations except low_ L_X which is always in absorption. Interestingly this transition for almost all simulations does not coincide with the mean temperature transitioning to positive values in Fig 7.3. Looking at Fig 7.1 there is a clear explanation. During the transition period between the EoH and the EoR ($z \sim 10$), the star formation rate increases. The X-ray luminosity is proportional to the star formation rate. This results in the appearance of heated islands that are in emission relative to the rest of the IGM which is undergoing more uniform and less efficient heating. These isolated heated islands drive the SS into positive territory, eventually tapering as the IGM saturates and the heated regions overlap. This results in the relative flattening of the SS as the temperature distribution on large scales saturates.

During reionisation, the ionisation morphology appears to dominate the signal in the SS, with a peak associated with the midpoint of reionisation, with the exception of the low_ L_X simulation which has a trough (effectively mirrored about the x-axis). As the neutral fraction drops below 50 percent most of the medium is ionised, and this tapers the SS, much like the PS, driving the signal to zero as reionisation progresses. Like the PS, the SS amplitude is largest for the high_ M_h simulation at the largest scale due to the fact that reionisation is driven by larger, more biased galaxies. However, the high_ L_X simulation has the highest overall amplitude in panel (b) and (c) Fig 7.5. The efficient X-ray heating in the high_ L_X simulation drives the non-Gaussianities, and results in an earlier reionisation, and therefore peak at ($z \sim 7$) compared to the other simulations. The earlier conclusion of reionisation results in a noisy SS at the lowest redshifts. The noise increase is due to the relatively small number of neutral voxels, and is dominated by small number statistics.

In general (particularly the small scales) we find that the SS follows $\overline{\delta T_b^3}$ as a function

of redshift. This indicates that the non-Gaussianities at the smaller scales are driven largely by the matter distribution. Whereas the deviations at large scales from $\overline{\delta T_b^3}$ are primarily due to non-Gaussianities introduced from the X-ray heating around the first luminous sources. This is most obvious in the timing of the peak in panel (a) between the low_ E_0 and the Fiducial simulation, where the only difference in the simulation is the X-ray energy threshold.

7.4.3.4 Skew Spectrum and Power Spectrum Co-evolution

Fig 7.6 shows the dimensionless SS and PS of each simulation as a function of redshift. Each simulation set is separated into individual panels, with the mid_ M_h , and high_ M_h in panel (a), the low_ L_X and high_ L_X in panel (b), and the low_ E_0 , and high_ E_0 simulations in panel (c). Each curve is taken at $k \sim 0.1 \text{ Mpc}^{-1}$, the colour coded triangles show the direction the curve travels as a function of redshift. The start point is indicated by the right facing triangle (relative to the peak), and the endpoint being the left facing triangle.

The curves in Fig 7.6 demonstrate how the SS and PS co-evolve as a function of redshift. For the Fiducial simulation we see the PS increases as the SS becomes more negative, this occurs during absorption at early redshifts. Then as the IGM heats due to X-ray emission, the PS amplitude decreases, and the SS amplitude increases towards zero, until eventually becoming positive at a fixed PS amplitude. The PS and SS then both increase as reionisation and heating occur in tandem, eventually culminating in a downward trend towards zero after the midpoint of reionisation. This evolution is broadly mirrored by the other simulations with the exception of the low_ L_X and high_ L_X simulations. For the low_ L_X simulation the SS is always negative, there is also a fairly flat SS from $z = 10$ to $z = 6$, with a small peak. These corresponds to the turning points in the curve, similar to the Fiducial simulation. The high_ L_X simulation demonstrates deviation from the Fiducial evolution, with the turning point from negative to positive SS amplitude happening at low PS amplitude. We see more turning points with increasing PS amplitude for the high_ L_X simulation which are correlated with the




Figure 7.7: PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), mid_ M_h (double dotted dash line), and the high_ M_h (dash dotted line). Each figure from left to right is the $\bar{x}_{\text{HI}} \sim 0.75$, 0.5 , and 0.25 for each simulation.

two peaks seen between $10 < z < 20$ in Fig 7.4 panel (b). These finally culminate in a turning point at high SS and PS amplitude during the peak of reionisation, transitioning to the zero amplitude for both.

For the low_ E_0 and high_ E_0 simulations, the evolution is notably different. For the low_ E_0 simulation, the structure mirrors the Fiducial simulation, with the sharp transition in amplitude for the SS happening earlier and at higher PS amplitudes due to the increased heating morphology. The more uniform heating of the IGM in the case of the high_ E_0 simulation leads to less variation, and thus to less rapid growth of the SS and PS amplitude with respect to the other simulations. Notably, the transition from negative to positive SS amplitude occurs over a longer period for high_ E_0 . The PS amplitude in this case varies as the SS amplitude transitions from negative to positive; this is a key difference between the high_ E_0 on other simulations.

For the halo mass simulations, the co-evolution of the SS and PS as a function of redshift is practically identical compared to the Fiducial simulation. The deviations occur due to the differences between the PS and SS amplitudes. During the EoH

since all simulations have the same X-ray background, the main difference is due to heating from UV emission. For the Fiducial simulation this can occur at lower halo masses. The heating leads to a reduction in the PS amplitude for the Fiducial simulation when the IGM is in absorption; this effect is however minimal compared to X-ray heating (Furlanetto et al., 2006b). Once reionisation starts, the $\text{mid-}M_h$ and $\text{high-}M_h$ simulations can emit more UV photons per unit mass per unit star formation, this results in relatively more UV heating. When the IGM is in emission, this leads to higher SS and PS amplitudes for the $\text{mid-}M_h$ and $\text{high-}M_h$ simulations compared to the Fiducial.

7.5 Power Spectrum and Skew Spectrum During Reionisation

In this section we calculate the spherically averaged PS and the spherically averaged SS for each simulation during the epoch of reionisation. We investigate what information we are sensitive to during reionisation, and what measuring the PS and the SS together can reveal about the physics of reionisation. We investigate this by measuring the normalised SS, discussed in the following section [7.5.1](#).

7.5.1 Normalised Skew Spectrum

The SS like the bispectrum is a measure of the non-Gaussianity through the central third order moment statistics of the temperature field. Watkinson et al. (2018) found significant fluctuations in the expected bispectrum around the zero point. These large fluctuations are due to the PS amplitude present in the statistic. Inspired by Eggemeier and Smith (2016), Watkinson et al. (2018) normalises the bispectrum to a unitless ‘normalised’ bispectrum, which is normalised by the PS and the k -modes. This normalisation is akin to measuring skewness, which is the central third order moment normalised by the cube of the standard deviation (the variance to the power of $3/2$). In this work we perform a similar normalisation to remove the Gaussian component




Figure 7.8: PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), low_ L_X (double dotted dash line), and the high_ L_X (dash dotted line). Each figure from left to right is the $\bar{x}_{\text{HI}} \sim 0.75$, 0.5 , and 0.25 for each simulation.

of the amplitude in the SS. We normalise the SS by the PS taken to the power of $3/2$, providing a unitless quantity $\gamma(k)$, referred to as the normalised SS:

$$\gamma(k) = \frac{\Delta_{T^2T}^2}{(\Delta_T^2)^{3/2}}. \quad (7.17)$$

In Equation 7.17 Δ_T^2 and $\Delta_{T^2T}^2$ are the dimensionless spherically averaged PS and SS, which have units of mK^2 and mK^3 respectively. The subscripts T and T^2T indicate the dimensionless PS and SS respectively. Deviations from a flat distribution as a function of spatial scale will be indicative of when the PS of the temperature fluctuations is more or less significant relative to the SS.

7.5.2 Power Spectrum, Skew Spectrum, and the Normalised Skew Spectrum

In this section we focus our investigation on what the PS and SS look like for each simulation at different stages of reionisation, specifically at neutral fractions of $x_{\text{HI}} \sim$




Figure 7.9: PS (row one), SS (row two) and normalised SS (row three), of the Fiducial (solid black line), low_ E_0 (double dotted dash line), and the high_ E_0 (dash dotted line). Each figure from left to right is the $\bar{x}_{\text{HI}} \sim 0.75, 0.5, \text{ and } 0.25$ for each simulation.

0.75, 0.5, and 0.25. This allows for a direct comparison of the state of the IGM for each simulation set as compared to the Fiducial simulation. In addition to the PS and SS at these states of the IGM, we also calculate the normalised SS ($\gamma(k)$) from Equation 7.17 for each simulation. $\gamma(k)$ allows for the isolation of the non-Gaussianity of the signal by normalising out the PS amplitude. In Figs 7.7, 7.8 and 7.9 we show the spherically averaged PS (first row), the spherically averaged SS (second row), and $\gamma(k)$ (third row) as a function of spatial scale, with the different neutral fractions of $\sim 0.75, \sim 0.5$ and ~ 0.25 representing the first, second and third columns of each figure.

In Fig 7.7 we look at the PS, SS and the normalised SS of the halo mass simulation set. In Fig 7.7 the Fiducial simulation is the solid black line, the mid_ M_h the double dot dashed line, and the high_ M_h is the dashed dot line. For both the PS and the SS we see a flattening of the spectra ($k < 1 \text{ Mpc}^{-1}$) as reionisation progresses. This is typically related to the ionisation morphology (Zaldarriaga et al., 2004). Notably, the PS and SS are very similar as a function of k , this is due to the Gaussian component present in the SS amplitude.

In the normalised SS we see an interesting feature develop in all three simulations, there is a local minima and local maxima in $\gamma(k)$. The local minima is caused by the flattening of the PS due to the ionisation morphology. The flattening is more prevalent in the PS than the SS, and therefore a minima appears in $\gamma(k)$ at the characteristic scales of ionising regions. The minima only becomes prevalent during the mid to late stages of reionisation once the ionised regions have percolated. The local maxima on the other hand is due to small scale structures. At $\bar{x}_{\text{HI}} \sim 0.75$ the amplitude is tied to the non-linear clustering of the ionised sources and their individual ionised bubbles. This is evident due to the increased amplitude for the Fiducial simulation relative to the mid- M_h and high- M_h simulations. The Fiducial simulation contains many more smaller mass galaxies increasing the non-linear amplitude on these scales. As reionisation proceeds, the maxima grows in amplitude. At $\bar{x}_{\text{HI}} \sim 0.25$ the maxima is significantly larger for the Fiducial simulation relative to the mid- M_h and high- M_h simulations. The peak is now driven by the prevalence of neutral islands in the IGM (see e.g. Fig 7.1). Due to the smaller escape fraction, the Fiducial simulation contains the largest number of neutral islands. On the other hand, the high- M_h with its much larger escape fraction ionises a larger volume per ionising source, preventing the appearance of neutral islands. Thus the high- M_h simulation does not exhibit a strong local maxima.

The location of the local minimum for the Fiducial simulation changes with decreasing neutral fraction from $k \sim 0.2 \text{ Mpc}^{-1}$ to $\sim 0.08 \text{ Mpc}^{-1}$. The characteristic scales of ionising regions is expected to increase with redshift and decreasing (increasing) neutral (ionisation) fraction (Giri et al., 2017). The local minima at $\bar{x}_{\text{HI}} \sim 0.5$ appears at $k \sim 0.1 \text{ Mpc}^{-1}$ for all simulations. The expected size of the ionising regions ranges from 20 – 100 Mpc (Wyithe and Loeb, 2004; Zaldarriaga et al., 2004; Lin et al., 2016), with the minima corresponding to size scales from $\sim 31 - 79 \text{ Mpc}$. The local maxima on the other hand appears at $k \sim 1 \text{ Mpc}^{-1}$ for the Fiducial, $k \sim 0.6 \text{ Mpc}^{-1}$ for the mid- M_h and $k \sim 0.4 \text{ Mpc}^{-1}$ for the high- M_h simulation. These correspond to sizes 6.3 Mpc, 3.8 Mpc, and 1.9 Mpc for the Fiducial, mid- M_h and the high- M_h simulations

respectively.

Fig 7.8 has the same layout as Fig 7.7, here low_L_X is the double dot dashed line, and the dash dot line is the $high_L_X$ simulation. Here the low_L_X simulation power spectra differs from the Fiducial and $high_L_X$ at all three ionisation states. This difference is driven by the larger absolute temperature difference of the low_L_X simulation compared to the other two simulations. The flattening of the low_L_X PS at $k < 1 \text{ Mpc}^{-1}$ is still prevalent, since reionisation still proceeds the same for the low_L_X , $high_L_X$ and Fiducial simulation. The differences at high k are due to the signal temperature on small scales, in the Fiducial and the $high_L_X$ simulations the signal is in emission, with the $high_L_X$ simulation having a larger emission temperature which causes larger temperature offsets with the ionised regions and thus produces a higher-amplitude 21cm PS. This is not the case for the low_L_X simulation which remains in absorption as reionisation progresses.⁵

When we look at the SS we see a similar trend, with the exception that low_L_X is negative compared to the Fiducial and $high_L_X$ scenarios. The shape of low_L_X relative to the other simulations is broadly mirrored about the k -axis, and shows a similar morphology to the PS, as is similar with the Fiducial and the $high_L_X$ simulations. The normalised SS shows similar features to the Fiducial for both the low_L_X (albeit negative) and for the $high_L_X$. There is a peak/trough at $k \sim 0.1 \text{ Mpc}^{-1}$, and a subsequent peak/trough at $k \sim 1 \text{ Mpc}^{-1}$. However, this is less pronounced in the low_L_X case, and this feature has more of a flattening from $k = 1 - 10 \text{ Mpc}^{-1}$, which is also seen in the PS. Due to the similarity in the ionisation morphology between the Fiducial, low_L_X and the $high_L_X$ simulations especially at $x_{\text{HI}} \sim 0.5$, and 0.25 , the amplitudes of the higher modes ($k > 1 \text{ Mpc}^{-1}$) appear to be affected more by the X-ray heating, than the ionisation morphology itself.

Fig 7.9 shows the X-ray threshold simulation PS, SS and normalised SS. Here the double dotted dashed line is the low_E_0 simulation, and $high_E_0$ is the dashed dotted line. In all cases the SS and the PS of the Fiducial and the low_E_0 simulations are

⁵It should be noted that at smaller scales $k > 1 \text{ Mpc}^{-1}$ approximations in the construction of the simulation yield numerical artefacts that could potentially impact the results.

almost identical during reionisation, differing at most at $k > 1 \text{ Mpc}^{-1}$. Again, the ionisation morphology for all three simulations is practically identical because they contain the same halo mass distribution and threshold.

In the case of the high_ E_0 simulation, increasing E_0 removes the lowest energy photons, this moves the peak of the X-ray distribution to higher energies. This results in a higher relative fraction of harder X-ray photons. Since harder X-rays heat on larger scales (more uniform heating), this appears to reduce the amplitude on all scales for both the PS and SS respectively. However the reduction seems to impact the SS more than the PS. Fluctuations and non-Gaussianities during reionisation are largely driven by ionisation morphology at large scales ($k < 1 \text{ Mpc}^{-1}$) and by the gas density on small scales ($k \geq 1 \text{ Mpc}^{-1}$) (Watkinson et al., 2018; Majumdar et al., 2018; Ma and Peng, 2023).

7.6 Detectability of the Normalised Skew Spectrum

There is a clear potential advantage to the normalised SS ($\gamma(k)$) compared to the SS. The trough at $k \sim 0.1 \text{ Mpc}^{-1}$ and the peak at $k \sim 1 \text{ Mpc}^{-1}$ demonstrate the sensitivity of the normalised SS to non-Gaussianity in the 21cm signal. Additionally, calculating the SS essentially requires the calculation of the PS, it is therefore straightforward to construct the normalised SS by dividing out the Gaussian amplitude. Naturally the detectability of the features present in $\gamma(k)$, particularly the trough, is important to estimate if there is any potential for its use as a probe of reionisation. In this section we investigate the statistical uncertainty of the PS and the SS, where we use these errors to propagate the expected statistical uncertainty on the normalised SS. We save the discussion of system noise from interferometers and foregrounds for future work; these effects deserve independent consideration.

First, we consider the statistical uncertainty of the 21cm PS, also known as the cosmic variance. Therefore of a random Gaussian field, the PS and by extension the variance, describe all the information contained in the field. In this case the uncertainty on the PS in the absence of thermal or instrumental noise is determined by the Poisson

sampling:

$$\sigma(\Delta_T^2)(k) = \Delta_T^2(k) \sqrt{\frac{(2\pi)^2}{Vk^2\Delta k}}. \quad (7.18)$$

The Poisson sampling error is proportional to one over the square root of the number of Fourier modes for a given spherical shell with width Δk ⁶. The number of modes is proportional to the sampled co-moving volume of space V used to calculate the spectrum. Equation 7.18 assumes Gaussianity, however the signal becomes non-Gaussian as reionisation progress (Cooray, 2005; Furlanetto, 2006; Wyithe and Morales, 2007). These non-Gaussianities correlate the signal on different Fourier modes (Mondal et al., 2015b), which in turn introduces a noise floor to the expected PS cosmic variance (Mondal et al., 2015a, 2016):

$$\sigma(\Delta_T^2)(k) = \Delta_T^2(k) \sqrt{\frac{(2\pi)^2}{Vk^2\Delta k}} + \sqrt{\frac{T(k, k)}{V}} \quad (7.19)$$

The non-Gaussianities in the PS are proportional to the Trispectrum $T(k, k)$ and $V^{-1/2}$. The effect of these non-Gaussianities is to flatten the signal to noise providing a fundamental limit to the signal detectability (Mondal et al., 2015b). A calculation of the analytic cosmic variance of the SS is outside of the scope of this work. In this section we assume that the relative uncertainties in the skew-spectrum follow a similar relationship to the PS. Since the SS contains a Gaussian amplitude component, this seems reasonable, and we will demonstrate this in the following sections.

7.6.1 Cosmic Variance of the Power Spectrum and the Skew Spectrum

To investigate the cosmic variance during the EoR, we follow the same method outlined in Balu et al. (2023). We split each of the simulation volumes in Figs 7.7, 7.8 and 7.9 into 27 sub volumes each with side length $70 h^{-1}$ Mpc. We then calculate the PS and SS for each sub volume. The cosmic variance for the PS and the SS is numerically

⁶ Δk is not a constant since logarithmic bins are typically used to calculate the PS.



(a) $\Delta_T^2(k)$ Uncertainty(b) $\Delta_{T^2T}^2(k)$ Uncertainty

Figure 7.10: The numerically estimated statistical uncertainties on PS (a) and the SS (b) (solid line), compared to the expected Gaussian uncertainties (dashed line), and the estimated non-Gaussian component (dotted line).

estimated by calculating the variance with respect to the mean power and SS for each k -mode. Subfigures 7.10a and 7.10b show the sample error (solid black line) for the PS and the SS, compared to the expected Gaussian uncertainty (dashed line) for the Fiducial model at a neutral fraction of $x_{\text{HI}} \sim 0.5$. We also estimate the Non-Gaussian component (dotted line) by subtracting the expected Gaussian errors from the sampled errors, these results are similar to those shown in Figure 3 of Greig et al. (2022a). We find qualitative agreement with Mondal et al. (2015a, 2016); Greig et al. (2022a), where we find an earlier transition to non-Gaussianities in Subfigure 7.10a in agreement with that seen by Balu et al. (2023).

The uncertainties at large scales are dominated by the Gaussian component at ($k \lesssim 0.3\text{Mpc}^{-1}$) for the PS and ($k \lesssim 0.2\text{Mpc}^{-1}$) for the SS. This transition is seen at ($k \lesssim 0.5\text{Mpc}^{-1}$) in Mondal et al. (2015b) and Greig et al. (2022a). Balu et al. (2023) attributes the earlier transition to the spin temperature evolution and detailed physical prescriptions for the higher non-Gaussianity in the L210 box. Overall, we find the assumption that the uncertainties in the SS have a similar form as the PS to be a good one.

7.6.2 Uncertainty in the Normalised Skew Spectrum

To estimate the uncertainty in the normalised SS we calculate the linear first order error propagation of equation 7.17:

$$R_y^2(k) = \frac{9}{4}R_T^2(k) + R_{T^2T}(k) - 3R_T(k)R_{T^2T}(k)\rho(k), \quad (7.20)$$

($R(k) = \sigma_x(k)/X(k)$) is the relative error for either the PS or the SS, labelled with the subscripts T and T^2T respectively. $\rho(k)$ is the Pearson correlation coefficient of the dimensionless PS and SS as a function of spatial scale. We find significant correlation between the PS and the SS as calculated from the sub volumes in the previous section.

Fig 7.11 shows the correlation for the $\bar{x}_{\text{HI}} \sim 0.75, 0.5$, and the 0.25 Fiducial simulation coeval boxes as a function of k . The average correlation is 0.82, 0.88 and 0.92 respectively for each of the coeval boxes. Considering the morphological




Figure 7.11: Correlation coefficient between the Fiducial PS and SS for the $x_{\text{HI}} \sim 0.75$ (solid), the 0.5 (dash dotted line), and the 0.25 (double dot dashed line) neutral fractions.

similarity between the PS and the SS during reionisation this result is not surprising. This demonstrates that the SS during reionisation has significant contribution from the PS amplitude. Of note is the lower correlation at large scales for the $\bar{x}_{\text{HI}} 0.75$ simulation subvolumes. There are two effects which likely result in the lower correlation at large scales. Firstly, the variance on the PS and SS is higher at larger spatial scales for the simulation subvolumes due to fewer k -mode bins at these scales. Secondly, ionisation bubbles at $\bar{x}_{\text{HI}} 0.75$ are more likely to be isolated (~ 10 Mpc), than to be merged (as they are at smaller neutral fractions) with other bubbles to form larger structures (> 10 Mpc). This results in some of the subvolume PS and SS at larger scales being more or less dominated by the ionisation morphology or the matter density. Whereas at small scales the matter density dominates in all ionisation scenarios. This is also true for the ionisation morphology at large scales when $\bar{x}_{\text{HI}} \leq 0.5$. The correlation of the matter density with the ionisation morphology at the largest scales results in a lower correlation.

We can derive an expression for the Gaussian component errors in the normalised skew spectrum if we assume that Trispectrum component of R_T and R_{T^2T} is zero ($T(k, k) = 0$). In this case the relative errors for the PS and SS are equal ($R_T = R_{T^2T}$)

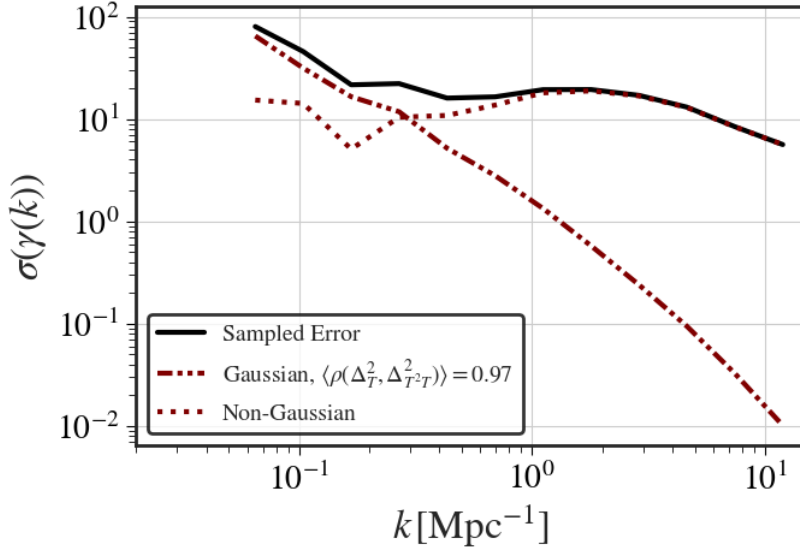


Figure 7.12: The first order propagated uncertainties on the normalised SS (solid black line), compared to the Gaussian model uncertainties (dashed line), and the estimated non-Gaussian component (dotted line), for the Fiducial simulation at $\bar{x}_{\text{HI}} = 0.5$.

and only depend on the volume V and the shell volume $k^2 \Delta k$. Equation 7.20 therefore simplifies to:

$$R_\gamma(k) = \frac{R_T(k)}{2} \sqrt{13 - 12\rho(k)}. \quad (7.21)$$

We use Equation 7.21 as a model for the Gaussian component of the uncertainties in $\gamma(k)$. Fig 7.12 shows the sample estimated errors (solid black line) calculated from Equation 7.20. The red double dot dashed line shows the estimated Gaussian uncertainties estimated from Equation 7.21. Finally the non-Gaussian component (dotted line) was likewise estimated by subtracting the Gaussian uncertainty model at $\bar{x}_{\text{HI}} = 0.5$ from the full uncertainty estimation. The transition from Gaussian dominated to non-Gaussian dominated uncertainties occurs at $k \lesssim 0.3 \text{Mpc}^{-1}$. Interestingly, we find that the Non-Gaussian uncertainties are roughly flat as a function of k .

7.6.2.1 Detection Predictions

In this section we perform a rudimentary signal to noise (S/N) estimate for the normalised SS, for future SKA_LOW observations. For the estimate we assume the Fiducial simulation as the 21cm signal, and we consider the neutral fractions $\bar{x}_{\text{HI}} = 0.25$,

$\bar{x}_{\text{HI}} = 0.5$, and $\bar{x}_{\text{HI}} = 0.75$, and their respective redshifts 7.4, 6.6 and 6.1.

The signal to noise ratio can be defined as the inverse of the relative uncertainty ($S/N(k) = 1/R_\gamma(k)$). In this case we assume the full relative uncertainty for $\gamma(k)$, which includes the non-Gaussianities in Equation 7.20. Notably, $R_\gamma \propto V^{-1/2}$ ($S/N \propto V^{1/2}$) therefore, to estimate the S/N for an SKA_LOW observation, we can replace the simulation comoving volume ($V_{\text{sub}}^{1/2}$) by the SKA_LOW comoving volume ($V_{\text{SKA}}^{1/2}$). This is done by first calculating the S/N for the Fiducial simulation, then normalising out $V_{\text{sub}}^{1/2}$, and finally scaling by $V_{\text{SKA}}^{1/2}$.

To determine the SKA_LOW comoving volume for an observation, we first need to know the field of view Ω_f and the observing bandwidth $\Delta\nu$. From these values we can determine the comoving volume for each redshift (neutral fraction) (Hogg, 1999). The angular width of the main lobe of the primary beam for an interferometer is approximately given by $\theta \sim \lambda/D$, where λ is the observing wavelength, and D is the station diameter which we assume is 35 m (Turner, 2015). Using the observing wavelength for each neutral fraction we calculate the field of view to be 2.87^2 deg^2 , 2.6^2 deg^2 , and 2.44^2 deg^2 , for each \bar{x}_{HI} respectively. Finally assuming an observing bandwidth of $\sim 30 \text{ MHz}$ for each neutral fraction, we then determine the comoving volume using Equation 28 from Hogg (1999). We then calculate the S/N for each neutral fraction for the Fiducial simulation by taking the ratio of $\gamma(k)/\sigma(\gamma(k))$. The S/N is then scaled by $\sqrt{V_{\text{SKA}}/V_{\text{sub}}}$.

Fig 7.13 shows the resulting estimated S/N for $\gamma(k)$ during reionisation for $\bar{x}_{\text{HI}} = 0.25$ (dashed dot line), $\bar{x}_{\text{HI}} = 0.5$ (dash double dot line), and $\bar{x}_{\text{HI}} = 0.75$ (solid line). The tapering of the S/N as a function of k is characteristic of the non-Gaussian component in the cosmic variance and is also seen in Figure 4 from Mondal et al. (2015b). We see that for all neutral fractions for all spatial scales the $S/N > 10$, with a max signal to noise of ~ 300 for $\bar{x}_{\text{HI}} = 0.75$. Thus, in the absence of thermal noise, foregrounds and systematics, the features present in $\gamma(k)$ should be detectable. Ma and Peng (2023) perform their own S/N analysis for the SKA_LOW with the addition of 1000h of thermal noise. They find a S/N of ~ 20 for their SS estimates. In the ideal case




Figure 7.13: Signal to noise ratio of $\gamma(k)$ for the Fiducial simulation scaled by the expected SKA_LOW observing comoving volume. The solid curve corresponds to the Fiducial neutral fractions $x_{\text{HI}} \sim 0.75$, 0.5 for the dash dotted line, and 0.25 for the double dot dashed line.

assuming foregrounds and systematic noise can be removed, we should expect to be sensitive to the normalised SS trough at $k = 0.1 \text{ Mpc}^{-1}$ and peak at $k \leq 1 \text{ Mpc}^{-1}$.

7.7 Discussion and Conclusion

This work investigates the non-Gaussianities resulting from the ionisation morphology and from X-ray heating in the neutral hydrogen signal as a function of redshift. We calculate the spherically averaged PS and SS of the expected 21cm signal. These simulations were generated for a set of seven simulations in the redshift range of 30 – 5, produced with the semi-analytic simulation package MERAXES. We vary the ionisation morphology through the halo mass threshold, the X-ray luminosity as a function of star formation, and the X-ray energy threshold. The ionisation morphology introduces intrinsic non-Gaussianities by effectively eliminating emission around ionisation sources. The size and distribution of ionisation topologies is therefore important for understanding non-Gaussianities during the EoR. X-ray heating, and in particular the energy distribution of X-rays emitted by luminous sources also affect the non-Gaussianities of the 21cm signal. Softer X-rays heat the IGM on more local scales, by varying this cutoff we effectively control how patchy the heating is; these non-Gaussianities

are most significant during the EoH. We also vary the X-ray luminosity as a function of star formation rate. This affects the amplitude of the signal and the sign of the non-Gaussianities.

We find some interesting features in the SS during the EoH for most of the simulations. These features signal the emergence of the first localised regions of above average temperature contrast around the large luminous sources. These features are only observed on large scales $k \sim 0.05 \text{ Mpc}^{-1}$, and drive the 21cm SS into the positive regime. We do not expect to be sensitive to these features with instruments such as the SKA_LOW due to foregrounds. We find that during the reionisation the SS signal is dominated by the ionisation morphology, where the X-ray heating primarily modulates the amplitude and smaller spatial scales ($k \gtrsim 0.5 \text{ Mpc}^{-1}$) of the SS and PS.

We further investigate the ionisation state and statistics of the IGM during the EoR. We look at the spherically averaged PS, SS and the normalised SS for each simulation set compared to the Fiducial at $\bar{x}_{\text{HI}} \sim 0.75, 0.5$ and 0.25 . In all simulations we find a local minimum at $k \sim 0.1 \text{ Mpc}^{-1}$ during the midpoint of reionisation. This minimum corresponds to the characteristic ionisation topology (bubble) size during the EoR (Wyithe and Loeb, 2004; Furlanetto and Oh, 2005; Lin et al., 2016), and we see the evolution of the minima to larger scales with decreasing neutral fraction (increasing ionisation fraction). We expect that this feature should be detectable at $k = 0.1 \text{ Mpc}^{-1}$ in the absence of instrumental, thermal noise and foreground contamination for the SKA_LOW and by extension current interferometric experiments. We also see evidence of a local maxima in the halo mass simulations, that shifts to larger scales as a function of the halo mass threshold. This corresponds to small ionised or hot regions around these halos. Our study highlights the importance of higher order statistics, and what can be gained from calculating both the SS and the PS.

The halo mass threshold simulations display the importance of ionisation topology as demonstrated in Fig 7.7. This clearly has the biggest impact on the structure present in the normalised SS during reionisation. There are however some important limitations and caveats related to the halo mass simulations. These simulations varied the ionisation

morphology by restricting the halo mass threshold above which galaxies could emit ionising UV photons, and scaling the amount of UV emission as a function of the halo mass threshold. Galaxies below the threshold however still produced X-ray emission, and thus contribute to heating the IGM. Although this model is nonphysical, it allowed for the halo mass threshold simulations to have a comparable X-ray background to the Fiducial. This effectively isolated the impact of the ionisation topology on the 21cm statistics independent of X-ray heating. Changing the X-ray emission in the same manner would delay heating, and have an undesirable impact on the 21cm statistics.

Future work should consider how the ionisation topology is quantitatively linked to the trough feature seen in the halo mass simulations. This could be investigated by measuring the characteristic scales of ionisation for the coeval boxes in this work using the methods outlined in Lin et al. (2016), in particular the mean free path method. Understanding the characteristic ionisation scale and variance, and comparing them to the trough position and width are important for quantitatively understanding how the ionisation topology affects the higher order statistics. This could additionally be performed on a simple toy model similar to the one used in Majumdar et al. (2018), where the characteristic scales can be varied to understand their impact on the observed features in the normalised SS, independent of other affects such as X-ray heating.

Following from the bispectrum work in Shimabukuro et al. (2017) and Tiwari et al. (2022), the SS as a method for further constraining the EoR model parameters should be investigated. Shimabukuro et al. (2017) perform a Fisher analysis on the PS and bispectrum, and shows that the bispectrum is more sensitive to the EoR model parameters in comparison to the PS. Tiwari et al. (2022) show that using the bispectrum in conjunction with the PS improves the constraints on parameters by up a factor of 4 than by the PS alone. One caveat however, is the SS is the integrated bispectrum. Through the process of integration we lose some of the information that is contained in the bispectrum. Thus, we might expect the SS to perform poorly compared to the bispectrum, especially when considering the high degree of correlation between the SS and the PS during reionisation. The normalised SS which should be more sensitive to

the non-Gaussianity of the 21cm signal might offer a better constraint than what can be gained from the SS alone.

Future work will consider the practicality of calculating the SS with the current and future generations of radio interferometric instruments. The squared sky temperature can not be measured directly, so we must estimate it from interferometric visibilities derived from observations. This ultimately involves a convolution of the signal in Fourier space, or a multiplication in image space and subsequent inversion back to Fourier space. Each additional step in the process propagates systematic and instrumental effects, and spreads them across different Fourier modes. These effects already impact the PS for numerous experiments which are systematics limited rather than thermal noise limited. Further investigation is needed to understand how these effects propagate through the SS, and whether they render a realistic measurement impractical.

CHAPTER 8

INSTRUMENTAL FOREGROUND & 21CM SKEW SPECTRUM

This chapter is primarily my own work, with the exception of the 21cm WODEN simulations, which were performed by J. L. B. Line. The idea for this chapter came from C. M. Trott. I performed the WODEN simulations for the EoR0 and EoR1 fields. I constructed the all-sky model, from which I calculated the power spectrum, the skew spectrum and the normalised skew spectrum. I developed the mathematical framework for the skew visibilities with input from C. M. Trott. I developed the toy-model, which I use in the appendix to analyse the normalised skew spectrum. Discussions with C. M. Trott and J. L. B. Line. contributed to the analysis in this work. I wrote the chapter, with feedback and proof reading contributed by my supervisors C. M. Trott and J. L. B. Line.

In the previous chapter, we found that the normalised skew spectrum (SS) contained astrophysical information about the ionisation morphology during the EoR. Practical measurements of the normalised SS require estimation from radio interferometric visibilities. This poses some unique challenges due to the chromatic nature of radio interferometric instruments, and has not been attempted before. In addition, astrophys-

ical foregrounds will be an issue for any instrumentally estimated SS and therefore require investigation to characterise their effects. In this chapter we investigate the foreground SS, and the SS for a fiducial 21cm signal simulated to match the MWA's widefield of view. We develop an analytical framework for estimating the instrumental SS from the radio interferometric visibilities. We perform `WODEN` simulations for both the foreground and 21cm models. From the `WODEN` simulated visibilities we estimate the instrumental foreground and 21cm SS. We find that the amplitude of the normalised SS for the fiducial 21cm signal in the window is one to two orders of magnitude greater than the foregrounds. We find good agreement between the instrumental and expected 21cm spectra, however, chromatic effects such as mode mixing shift signal to smaller spatial scales in the instrumental estimation.

8.1 Introduction

Astrophysical foregrounds are up to five orders of magnitude brighter than the expected cosmological 21cm signal (Furlanetto, 2016, references therein). Outright removal of all foreground signal is practically impossible, therefore most 21cm experiments employ a mix of avoidance and mitigation. Radio interferometers however, are chromatic instruments. This chromaticity leads to foreground contamination across the line of sight and angular Fourier modes (known as mode mixing), forming a wedge-shaped feature in the 2D PS (Morales et al., 2012; Trott et al., 2012; Vedantham et al., 2012). Foreground contamination in the 21cm PS from astrophysical sources (point sources, extended sources, and diffuse sources) is a thoroughly-studied topic (e.g. Bowman et al., 2009; Datta et al., 2010; Procopio et al., 2017; Pober et al., 2016a; Line et al., 2020; Byrne et al., 2021; Cook et al., 2022). However, how foreground contamination affects higher order statistics is less well understood. Watkinson et al. (2020) investigate the foreground contamination for the MWA bispectrum to see if there is an equivalent wedge structure. They find foregrounds in the bispectrum to be a challenging obstacle for the measurement of the 21cm bispectrum.

The SS as estimated from a radio interferometer will suffer from similar chromatic

foreground contamination like the PS. Therefore, this chapter is focused on investigating these chromatic effects on both foregrounds, and a fiducial 21cm signal. To test these effects we develop an analytical method for estimating the instrumental SS from radio interferometric visibilities. We apply this method `WODEN` simulated visibilities for two EoR fields EoR0¹ and EoR1², and a fiducial 21cm signal. We additionally develop an analytical point source toy model that describes the effects of spectral tapering, and spatial tapering on the expected normalised SS; this is discussed in the Appendix Section 8.A.

The `WODEN` simulated foreground visibility models are generated for each 80 kHz fine channel in the 30.72 MHz bandwidth centred at 180 MHz. We also use model visibilities of a fiducial 21cm signal (Greig et al., 2022b) during the mid to late period of reionisation to model the expected 21cm SS. We then compare the foreground PS, SS, and normalised SS to the expected 21cm signal.

This chapter is structured as follows: In section 8.2 we discuss the foreground catalogues used to construct the foreground model; in Section 8.3 we discuss the fiducial 21cm model and the `WODEN` simulated visibilities; in Section 8.4 we derive the SS from MWA visibilities; in Section 8.5 we show the resulting foreground PS, SS and normalised SS; Section 8.6 we demonstrate the fiducial 21cm PS, SS, normalised SS, and compare it to the `WODEN` simulated PS, SS and normalised SS; In Section 8.7 we compare the foreground and 21cm window PS, SS and normalised SS; we summarise and discuss the results in Section 8.8.

8.2 Foreground Sky-Model

In order to understand foreground contamination, and to ultimately remove it from EoR data, sky-models of the foregrounds are required. Typically, these are constructed from radio interferometric surveys (catalogues) and images. In this work we categorise foreground emission into compact, extended, or diffuse. Compact objects are point

¹EoR0 field coordinates: RA=0 h, Dec=-27°

²EoR1 field coordinates: RA=4 h, Dec=-27°

source-like objects that have angular extents close to or below the instrument resolution (PSF); these account for the majority of astrophysical sources. Extended objects have larger angular extents up to several degrees. These can be of Galactic or extragalactic origin; examples include Fornax A, Centaurus A, Galactic plane SNRs and HII regions. Diffuse emission is predominantly composed of the background diffuse Galactic synchrotron emission, and covers a large portion of the sky.

Diffuse emission is difficult to measure and to model. Recent work by Byrne et al. (2021) and Kriele et al. (2022) to create diffuse sky maps for MWA EoR science reveal numerous challenges, most of them involve unique systematic effects which are still to be determined. Byrne et al. (2021) imaged the Southern hemisphere with the MWA, discovering large polarisation leakage across the EoR0 field; it is unclear if this emission is real or a systematic. Efforts to incorporate this map into the MWA EoR calibration and data reduction pipeline have yielded little improvement in MWA EoR foreground mitigation (MWA EoR group private communications). Kriele et al. (2022) on the other hand uses the m -mode spherical harmonic decomposition developed by Shaw et al. (2014) to measure the flux density of the Southern hemisphere. However, although promising this work only considers a single frequency channel, and does not remove all of the compact sources. In comparison with other maps such as those in (Haslam et al., 1982) and Eastwood et al. (2018), Kriele et al. (2022) finds several systematics present in their map. These are low level, but large enough to impact EoR science. In addition to the systematics present in both maps, modelling the images requires either modelling all pixels as point sources, or decomposing the maps into a set of Fourier analytic basis functions (e.g. Gaussians or shapelets). The latter method is complicated by projection effects, and the former requires excessive computation on the order of tens to millions of point source calculations per visibility per time and frequency bin³. In light of these difficulties, and considering that the principal EoR fields are pointed away from the brighter parts of the diffuse emission, we consider the inclusion of the diffuse emission outside the scope of this work.

³For a WODEN simulation this takes several hours to complete a snapshot simulation.

Point sources are the simplest model, and at a single frequency can be described by three parameters; their brightness, and their position co-ordinates in RA and DEC. Extended sources in contrast require complex multi-component models that are fitted with kernels to model their brightness morphology. These kernels can either be point sources, Gaussians or some other basis function (see Chapter 6 and 3 for a review).

The brightness of foreground sources additionally changes as a function of frequency. Frequency-dependent emission is typically modelled using a power law where $S \propto \nu^{-\alpha}$ (van der Laan, 1966; Conway et al., 1963). The spectral index α is determined by the underlying astrophysical processes, where synchrotron emission is the dominant emission mechanism. The spectral characteristics of these sources are often more complicated than a simple power law, with many sources having turnovers in their spectra primarily due to synchrotron self absorption (Blundell et al., 1999).

To accurately model the foreground contamination the sky-model therefore needs to contain morphological and spectral information about each source. The foreground sky-model is stitched together from several large sky-surveys with different sensitivities, angular resolutions and completeness. The MWA EoR sky-model is constructed primarily from a base of MWA survey catalogues, cross-matched with higher frequency (> 231 MHz) surveys; we describe these in Section 8.2.3. Additional work has focused on more accurately modelling sources in the principal EoR fields such as Procopio et al. (2017) and Lynch et al. (2021) (discussed further in section 8.2.2); these catalogues are contained in the EoR sky-model. In this section we briefly describe the different MWA surveys that form the basis of the sky-model, and how they are stitched together to create a multi-frequency sky-catalogue and model.

8.2.1 GLEAM Year 1

We described the GLEAM (Wayth et al., 2015) survey in Section 6.3.2; here we briefly describe the extra galactic data release from the first year of data collection (Hurley-Walker et al., 2017, from hereon GLEAMYR1). This data release covers $24,831 \text{ deg}^2$ of the Southern sky (below $+30 \text{ deg}$), for a total of 20 subbands each of width 7.68 MHz

covering the frequency range 72 – 231 MHz, taken with the Phase I MWA configuration (Tingay et al., 2013). In total the first data release catalogue contains 307, 455 detected radio sources, with angular resolution of ~ 2 arcmin. GLEAMYR1 is missing some key regions, in particular a 10 deg region around the Galactic plane due to calibration and source finding difficulties from diffuse synchrotron emission (Hurley-Walker et al., 2017). The Galactic plane data was later released by Hurley-Walker et al. (2019). Additionally, the Magellanic clouds are missing, for similar reasons to the Galactic plane, and were later released by For et al. (2018). The three final missing regions correspond 9 arcmin around Centaurus A; an observation where Centaurus A was located in a primary beam sidelobe, and a wedge region that could not be calibrated as a result of poor ionospheric conditions (Hurley-Walker et al., 2017).

8.2.2 LoBES

The MWA Long Baseline Epoch of Reionisation Survey (Lynch et al., 2021, LoBES) was designed to improve upon the resolution and sensitivity of GLEAMYR1, for the purpose of constructing high quality foreground models for the MWA EoR0 and EoR1 fields. LoBES was conducted using the extended configuration of the MWA Phase II array (Wayth et al., 2018), which has approximately double the resolution of the Phase I configuration (Tingay et al., 2013). The higher resolution⁴ allows for better sources models which reduces confusion noise. Additionally the extended Phase II configuration has more uniform uv -plane coverage (Wayth et al., 2018), which improves sidelobe confusion noise, through an improved synthesised beam (PSF). In addition to the EoR0 and EoR1 fields, the survey covers the surrounding eight fields to capture the sources present in the MWA primary beam sidelobes. Pober et al. (2016b) showed that point sources in the first MWA primary beam sidelobes contribute significantly to the foreground contamination to the PS. Therefore it is important that these are included in any foreground models.

⁴Less than an arcminute at 200 MHz, a factor of 2 better than GLEAMYR1.

Lynch et al. (2021) present a catalogue that covers 3069 deg^2 , with an average rms of $2.1 \text{ mJy beam}^{-1}$, containing 80,824 sources. There are 16 spectral measurements between 100 and 230 MHz, and 78% of sources have spectral modelling (as described in (Lynch et al., 2021)). Over the survey region the catalogue is 90% complete at 32 mJy, and 70% complete at 10.5 mJy.

8.2.3 PUMA

The Positional Update and Matching Algorithm (Line et al., 2017, PUMA) was created to combine catalogues with different sensitivities and resolutions. PUMA uses a positional Bayesian probabilistic approach along with spectral matching criteria to cross-match sources. PUMA first cross-matches sources based on their Bayesian probability, utilising a base catalogue (usually the one with the lowest angular resolution). Source matches are grouped into three categories: `isolated`, `dominant`, or `multiple`, depending on the probabilities. In the case of `isolated` sources there is only one match. High positional probability matches are accepted at face value; low probability matches must also satisfy a spectral criteria. `dominant` sources are the matches that best fit a spectral model compared to all the other possible combinations of sources. The final category `multiple` occurs if no dominant source is found; in this case the source flux densities are combined, and a power law is fitted to determine the spectral properties. If the fit is good, the match is accepted. If the source fails this step, it is flagged for visual inspection.

Using Line et al. (2017), the base of the EoR sky model was created by cross-matching `GLEAMR1` with the following surveys: the 74 MHz Very Large Array Low Frequency Sky Survey redux (VLSSr; Lane et al., 2014), TGSS (Intema et al., 2017), the 843 MHz Sydney University Molonglo Sky Survey (SUMSS; Bock et al., 1999), NVSS (Condon et al., 1998). The resulting sky-model contains 308,584 sources and covers frequencies from 72 MHz to 1.4 GHz, including (where possible) the full GLEAM bands from 72 MHz to 231 MHz. Additionally, the MWA EoR sky-model incorporates the LoBES catalogue, and the EoR1 sky-model from Procopio et al. (2017). We

additionally incorporate the SNR models and the Centaurus A model from Chapter 6 into the total sky-model.

Using the EoR sky-model, we perform two WODEN simulations for the EoR0 and EoR1 fields. Both simulations are zenith phased, 2 minute snapshot observations, centred at 180 MHz, with fine channel size of 80 kHz.

8.3 Fiducial 21cm Model

In this section we describe a fiducial 21cm signal generated by Greig et al. (2022b), that we simulate through WODEN. Greig et al. (2022b) constructed the Gpc scale 21cm lightcone used in this work from the 21cm simulation software 21CMFAST (Mesinger et al., 2011; Murray et al., 2020). The simulation lightcone covers the redshift range $z = 6.2 - 7.5$ (~ 180 MHz central frequency), with a transverse comoving size of 7.5 Gpc, and a line of sight comoving size of 703.1 Mpc. The lightcone is constructed from a total of $6400 \times 6400 \times 600$ voxels with a resolution Δx of ~ 1.17 Mpc, with each voxel having some temperature contrast with units mK. The lightcone was stitched together from 2D slices take from the total simulation volume, that were interpolated linearly with respect to cosmic time.

For a more detailed review on the simulation software and the excursion set formalism (Furlanetto et al., 2004a) used by 21CMFAST, we refer the reader to Chapter 7. Here we briefly describe the simulation parameters used to generate the 21cm signal. The simulation was performed assuming the spin temperature was saturated $T_S \gg T_{\text{CMB}}$, which greatly reduces the computational complexity, because the temperature brightness in this regime is effectively independent of the spin temperature. The 21cm signal simulation is parameterised using power law functions that vary as a function of the dark matter halo mass. The parameter f_* represents the fraction of gas contained in stars for a given galaxy, which is a power law with index α_* . The amount of ionising UV radiation produced by a galaxy that can escape is described by the escape fraction f_{esc} , this is also assumed to be a power law with index α_{esc} . Both power laws are normalised by a halo mass of $10^{10} M_{\odot}$. Feedback effects and inefficient cooling pro-

cesses restrict smaller mass haloes from containing active star forming galaxies. This is parameterised by an exponential function which describes the fraction of dark matter halos that cannot host galaxies. The suppression scale of the exponential is given by the parameter M_{turn} . Greig et al. (2022b) performed the simulation with the parameters $f_{*,10} = 0.05$, $\alpha_* = 0.5$, $f_{\text{esc},10} = 0.08$, $\alpha_{\text{esc}} = -0.5$, $t_* = 0.5$ and $M_{\text{turn}} = 10^{8.7}$, which are consistent with recent observational and reionisation constraints (Park et al., 2019). From hereon we refer to the simulation lightcone as the Fiducial lightcone. This is not to be confused with the Fiducial signal considered in Chapter 7, which was constructed with a different formalism and parameter set.

8.3.1 WODEN Simulation

In this work we use WODEN to simulate the Fiducial 21cm signal. The construction and simulation of the Fiducial 21cm visibilities is described in (Line et al. in prep); here we give an overview of the process used to simulate the signal into WODEN. As discussed in Chapter 3 WODEN simulates visibilities by directly calculating the measurement equation from a series of analytic functions that represent the source intensity distribution on the sky. The 21cm signal is effectively a random field, and changes drastically with frequency. To effectively simulate the signal each pixel has to be modelled as a point source in WODEN, each with a frequency sample at every observed frequency.

Modelling each pixel as a point source however is not without challenge. The Fiducial lightcone box is a collection of 2D Cartesian projections of the 21cm temperature brightness. Each slice corresponds to a given redshift (frequency), with the spacing roughly corresponding to $\Delta\nu \sim 80$ kHz (Greig et al., 2022b); this spacing however is not uniform in frequency. Furthermore the transverse comoving size of each redshift slice is the same. To simulate the point sources each pixel needs to be converted to an RA and DEC position, and a flux density in units of Jy, therefore the angular size of each pixel needs to be calculated. The pixel angular size however is redshift (frequency) dependent, and is given by $\Delta\theta = \Delta x / D_M(z)$, where $D_M(z)$ is the line of sight

comoving distance. The non-uniformity of the frequency spacing, and the frequency dependent angular size of the pixels require the the Fiducial lightcone be interpolated spatially and as a function of frequency to model the point sources in WODEN .

Line et al. (in prep) chose to interpolate all the slices to the finest angular resolution. This corresponds to the redshift slice $z = 7.57$, with a pixel angular resolution of ~ 27 arcsec. These were interpolated onto a TAN FITS projection (Calabretta and Greisen, 2002), centred at $RA = 0^h$ and $DEC = -30$ deg. Line et al. (in prep) experimented with different cubic and bi-linear interpolation methods, however, these all resulted in signal loss at small and large scales in the 21cm PS. Line et al. (in prep) found that the nearest neighbour interpolation was the only interpolation method that did not result in significant signal loss. It was also found that nearest neighbour interpolation was required for the frequency slices as well.

The resulting angular and frequency interpolated lightcone was simulated through WODEN for a zenith phased observation for a two minute snapshot at 8 s time resolution and 80 kHz frequency resolution. In this work we use this Λ CDM $h = 0.68$ cosmological model, with the following parameters, $\Omega_M = 0.31$, $\Omega_\Lambda = 0.69$, and $\Omega_b = 0.048$ from Planck Collaboration et al. (2021).

8.4 Instrumental Skew Spectrum

In this section we describe a method for calculating the SS from radio interferometric visibilities. There are two related approaches that can be employed to perform this calculation. The first method is the most straightforward, and involves imaging the visibilities for each frequency channel. These would then be squared, and Fourier transformed as a function of frequency to get the resulting $\tilde{I}(\mathbf{u}, \eta)$ cube. This would still require accurate foreground subtraction. Additionally, images would likely need to go through a deconvolution process such as CLEAN to identify and remove any remaining sources. However, limitations would result from remaining low signal to noise sources, these will contribute sidelobe confusion noise, which correlates the image noise statistics over large angular separations. Furthermore, imaging and CLEANing

all fine channels would be computationally expensive. This might be mitigated by performing a multi-frequency CLEAN, which would also increase the signal to noise for faint sources, at the cost of spectral structure for these sources. The second method involves a convolution of the visibility terms with each other and some kernel function. We detail the latter method and discuss the former at the end of this chapter.

8.4.1 Apparent Quadratic Sky Brightness

To estimate the SS, one of the fundamental quantities we need to determine is the Fourier transform of the quadratic sky intensity distribution $\tilde{I}^2(\mathbf{l})$. However, radio interferometers have a finite sensitivity response to the sky, and actually measure the apparent sky brightness intensity $I_A(\mathbf{l}) = \mathcal{B}(\mathbf{l})I(\mathbf{l})$, where $\mathcal{B}(\mathbf{l})$ is the instrument's primary beam response, and $I(\mathbf{l})$ is the true sky brightness intensity for a frequency between ν and $\nu + d\nu$. The process outlined in the rest of this section is independent of frequency, therefore we ignore the frequency dependence of the aforementioned terms. Due to the primary beam response we are fundamentally restricted to measuring $I_A^2(\mathbf{l}) = \mathcal{B}^2(\mathbf{l})I^2(\mathbf{l})$; taking the Fourier transform:

$$\tilde{I}_A^2(\mathbf{u}) \equiv \int_{\mathbb{R}^2} d^2\mathbf{l} \mathcal{B}^2(\mathbf{l}) I^2(\mathbf{l}) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}}. \quad (8.1)$$

This quantity cannot be measured directly by an interferometer, which only measures the Fourier transform of the apparent sky brightness (recall from Chapter 3):

$$\tilde{I}_A(\mathbf{u}) \equiv \int_{\mathbb{R}^2} d^2\mathbf{l} \mathcal{B}(\mathbf{l}) I(\mathbf{l}) e^{-2\pi i \mathbf{u} \cdot \mathbf{l}}. \quad (8.2)$$

We can however relate Equation 8.2 and Equation 8.1, through the convolution theorem:

$$\tilde{I}_A^2(\mathbf{u}) = \tilde{\mathcal{B}}(\mathbf{u}) * \tilde{I}(\mathbf{u}), \quad (8.3)$$

Likewise we can apply the convolution theorem to Equation 8.1:

$$\tilde{I}_A^2(\mathbf{u}) = (\tilde{\mathcal{B}} * \tilde{\mathcal{B}} * \tilde{I} * \tilde{I})(\mathbf{u}). \quad (8.4)$$

Convolution is a commutative and associative operation, meaning we can rearrange Equation 8.4 in the following manner:

$$\tilde{I}_A^2(\mathbf{u}) = (\tilde{\mathcal{B}} * \tilde{I})(\mathbf{u}) * (\tilde{\mathcal{B}} * \tilde{I})(\mathbf{u}), \quad (8.5)$$

therefore:

$$\tilde{I}_A^2(\mathbf{u}) = \tilde{I}_A(\mathbf{u}) * \tilde{I}_A(\mathbf{u}). \quad (8.6)$$

We can therefore describe the Fourier transform of the quadratic apparent sky brightness as the convolution of the Fourier transform of the apparent sky brightness with itself.

8.4.2 Skew Visibilities

Real instruments do not measure $\tilde{I}_A(\mathbf{u})$ in its entirety. Instead radio interferometers measure visibilities that are discrete samples of $\tilde{I}_A(\mathbf{u})$ in the uv -plane:

$$\mathcal{V}(\mathbf{u}) = S(\mathbf{u}) (\tilde{\mathcal{B}}(\mathbf{u}) * \tilde{I}(\mathbf{u})), \quad (8.7)$$

where $S(\mathbf{u})$ is the sampling function:

$$S(\mathbf{u}) = \sum_i^{N_{\text{vis}}} \delta(\mathbf{u} - \mathbf{u}_i), \quad (8.8)$$

\mathbf{u}_i indicates the i th baseline vector, and δ is the Dirac delta function. The measured visibilities can therefore be described as:

$$\mathcal{V}(\mathbf{u}) = \sum_i^{N_{\text{vis}}} \mathcal{V}_i \delta(\mathbf{u} - \mathbf{u}_i), \quad (8.9)$$

where $\mathcal{V}_i = |\tilde{I}_A(\mathbf{u})|e^{i\phi}$ is the complex visibility measured by the instrument. In imaging we reconstruct the Fourier transform of the apparent sky brightness through the gridding process. The gridding process approximates the Fourier transform by convolving the visibility with some gridding kernel, which is fit to the Full Width Half Maximum (FWHM) of the primary beam response:

$$\tilde{I}_A(\mathbf{u}) \approx (\tilde{G} * \mathcal{V})(\mathbf{u}). \quad (8.10)$$

In Equation 8.10 $\tilde{G}(\mathbf{u})$ represents the gridding kernel, which for simplicity we assume is a Gaussian $\tilde{G}(\mathbf{u}; \sigma_u) = \exp\left(-\frac{1}{2}(\mathbf{u})^2/\sigma_u^2\right)$. In this case the width of the Gaussian σ_u is inversely proportional to the FWHM of the main lobe of the primary beam response. The gridding procedure is necessary, because the measured visibility is effectively a weighted average of the uv -plane with the Fourier transform of the primary beam. Gridding mimics this instrumental effect by spreading the measured visibility with some kernel that has parameters proportional to the primary beam. We can substitute Equation 8.10 into Equation 8.6 to get an approximation of the Fourier transform of the quadratic apparent sky brightness:

$$\tilde{I}_A^2(\mathbf{u}) \approx (\tilde{G} * \mathcal{V})(\mathbf{u}) * (\tilde{G} * \mathcal{V})(\mathbf{u}). \quad (8.11)$$

Again we take advantage of the commutative properties of convolution and we rewrite Equation 8.11:

$$\tilde{I}_A^2(\mathbf{u}) = (\tilde{G} * \tilde{G})(\mathbf{u}) * (\mathcal{V} * \mathcal{V})(\mathbf{u}), \quad (8.12)$$

where the convolution of two identical Gaussians is a Gaussian $\tilde{G}(\mathbf{u}; \sigma_u) * \tilde{G}(\mathbf{u}; \sigma_u) = \tilde{G}(\mathbf{u}; \sqrt{2}\sigma_u)$, where the width has increased by a factor of $\sqrt{2}$. Next we perform the convolution of $(\mathcal{V} * \mathcal{V})(\mathbf{u})$:

$$(\mathcal{V} * \mathcal{V})(\mathbf{u}) = \sum_i^{N_{\text{vis}}} \sum_j^{N_{\text{vis}}} \mathcal{V}_i \mathcal{V}_j \int_{\mathbb{R}^2} d^2\mathbf{u}' \delta(\mathbf{u}' - \mathbf{u}_i) \delta(\mathbf{u} - \mathbf{u}' - \mathbf{u}_j). \quad (8.13)$$

Performing the integral yields:

$$(\mathcal{V} * \mathcal{V})(\mathbf{u}) = \sum_i^{N_{\text{vis}}} \sum_j^{N_{\text{vis}}} \mathcal{V}_i \mathcal{V}_j \delta(\mathbf{u} - \mathbf{u}_i - \mathbf{u}_j). \quad (8.14)$$

We now perform the convolution of Equation 8.14 with $\tilde{G}(\mathbf{u}; \sqrt{2}\sigma_u)$:

$$\tilde{I}_A^2(\mathbf{u}) = \sum_i^{N_{\text{vis}}} \sum_j^{N_{\text{vis}}} \mathcal{V}_i \mathcal{V}_j \int_{\mathbb{R}^2} d^2\mathbf{u}' \delta(\mathbf{u}' - \mathbf{u}_i - \mathbf{u}_j) \tilde{G}(\mathbf{u} - \mathbf{u}'; \sqrt{2}\sigma_u), \quad (8.15)$$

Integrating over the delta function:

$$\tilde{I}_A^2(\mathbf{u}) = \pi\sigma_u^2 \sum_i^{N_{\text{vis}}} \sum_j^{N_{\text{vis}}} \mathcal{V}_i \mathcal{V}_j \exp\left(-\frac{(\mathbf{u} - \mathbf{u}_i - \mathbf{u}_j)^2}{4\sigma_u^2}\right). \quad (8.16)$$

$\tilde{I}_A^2(\mathbf{u})$ is a continuous function on the uv -plane, however, visibilities are only measured at a single point. As such, we only calculate Equation 8.16 at locations where baselines are measured. Therefore we define what we imaginatively call the *skew visibility*, which is $\mathcal{V}_\gamma(\mathbf{u}) \equiv \tilde{I}_A^2(\mathbf{u})$. Where $\mathcal{V}_\gamma(\mathbf{u})$ is Equation 8.16 evaluated for a specific baseline location \mathbf{u} . From hereon all references to Equation 8.16 are in the context of the skew visibility $\mathcal{V}_\gamma(\mathbf{u})$.

In Equation 8.16 the dual sum requires N_{vis}^2 evaluations for each baseline, which is computationally expensive. However, the expression can be further simplified thanks to the Gaussian kernel. The product of $\mathcal{V}_i \mathcal{V}_j$ is highly attenuated when $|\mathbf{u} - \mathbf{u}_i - \mathbf{u}_j| \geq 5\sqrt{2}\sigma_u$. This is equivalent to a five sigma cut which attenuates the product by approximately six orders of magnitude. From the geometry of the 2D convolution, we can break this condition into two regions, one centred around the uv -plane origin, and another centred around some arbitrary baseline with position vector \mathbf{u}_k . In this case only visibilities with position vectors $|\mathbf{u} - \mathbf{u}_k| \leq 5\sqrt{2}\sigma_u$ or $|\mathbf{u}| \leq 5\sqrt{2}\sigma_u$ will have any significant contribution to the convolution at \mathbf{u}_k . This greatly reduces the number of required computations for a given baseline, and renders the computation tractable.

Another interesting feature from Equation 8.16 is the constant factor $\pi\sigma_u^2$. Since $\sigma_u \propto \sigma_l^{-1}$, where σ_l is directly proportional to the angular size in radians of the primary

beam, this implies that $\pi\sigma_u^2 \propto \Omega^{-1}$, where Ω is the field of view. This is what we would expect, since the units of $\tilde{I}_A^2(\mathbf{u})$ should be $\text{Jy}^2 \text{Sr}^{-1}$.

There is an important caveat to the method outlined in this section; the skew visibilities cannot be accurately estimated without the zero-spacing mode (DC mode) of the array. This is measured by the auto-correlations of the individual MWA tiles. The auto-correlations are the self correlation of an MWA tiles time varying complex voltage with itself. The auto-correlation is real valued, and represents the primary beam weighted average sky brightness as seen by each tile. Since the DC mode is real, and typically larger in value (with the exception of zero mean fields) than other modes, it has the largest contribution to the sum in $\mathcal{V}_\gamma(\mathbf{u})$. WODEN can simulate the auto-correlations, and the MWA measures the auto-correlations, therefore we include them in our simulations. There are however some important considerations to this requirement, which we discuss at the end of this Chapter.

8.4.3 Estimating the Skew Spectrum

Once the skew visibilities are determined, we can estimate the SS. We employ effectively the same estimation method as outlined for the PS calculation described in the OSIRIS pipeline in Chapter 6. Here we outline the same process but updated for the SS.

The first step requires gridding the skew visibilities onto a regular uv -grid. We use a grid with $\Delta v = \Delta u = 0.5 \lambda$, where the gridded value is determined as the weighted average of the visibilities with some kernel function:

$$\tilde{I}^2(\mathbf{u}_j, \nu) = \frac{\sum_{i=0}^{N_{\text{vis}}} W_{\text{BH}}(\mathbf{u}_j - \mathbf{u}_i; \sqrt{2}\sigma_u) \mathcal{V}_\gamma(\mathbf{u}_i, \nu)}{\sum_{i=0}^{N_{\text{vis}}} W_{\text{BH}}(\mathbf{u}_j - \mathbf{u}_i; \sqrt{2}\sigma_u)}. \quad (8.17)$$

The gridding kernel or the weights W_{BH} is a 2D Blackman-Harris gridding kernel, that has a width inversely proportional relative to the FWHM of a Gaussian fit to the MWA primary beam. For the skew visibilities the width is $\sqrt{2}\sigma_u$, with the extra factor of $\sqrt{2}$ because the squared primary beam is reduced in angular size by the same

factor. The 2D Blackman-Harris gridding kernel is calculated as the outer product of two orthogonal 1D Blackman-Harris functions. This has more of a square shape than the radial Blackman-Harris function, and is closer to the square main lobe of the MWA primary beam. The square Blackman-Harris gridding kernel improves upon the Gaussian kernel used in Chapter 6, and is similar to the gridding approach employed by the MWA EoR PS calculation pipeline CHIPS (Trott et al., 2016).

Following the gridding process we perform a spectral taper with a 1D Blackman-Harris that is centred at the central frequency:

$$\tilde{I}^2(\mathbf{u}, \eta) = \int_{\mathbb{R}} W_{\text{BH}}(\nu) \tilde{I}^2(\mathbf{u}, \nu) e^{-2\pi i(\nu \cdot \eta)} d\nu. \quad (8.18)$$

After Fourier transforming with respect to frequency, the gridded weights $W_{\gamma G}(\mathbf{u}, \nu)$ (which were collected during the gridding process) are averaged as a function of frequency:

$$\bar{W}_{\gamma G}(\mathbf{k}_{\perp}) = \frac{1}{\Delta\nu} \int d\nu W_{\gamma G}(\mathbf{k}_{\perp}, \nu). \quad (8.19)$$

Here we have implicitly performed the coordinate transform from \mathbf{u} to \mathbf{k}_{\perp} . The weights for each k_{\parallel} channel are defined as $\bar{W}_{\gamma G}(\mathbf{k}_{\perp}, k_{\parallel}) = \bar{W}_{\gamma G}(\mathbf{k}_{\perp}) \forall k_{\parallel}$. Gridding is also performed for the visibilities, and the frequency-averaged gridded weights are defined as \bar{W}_G .

The 1D spherically averaged SS is then calculated as the weighted average, considering the gridded weights from both the visibilities and the skew visibilities:

$$S_{\gamma}(k) = \frac{\sum_{i \in |\mathbf{k}|} \tilde{I}^2(\mathbf{k}_i) \tilde{I}^*(\mathbf{k}_i) \sqrt{\bar{W}_{\gamma G}(\mathbf{k}_i) \bar{W}_G(\mathbf{k}_i)}}{\sum_{i \in |\mathbf{k}|} \sqrt{\bar{W}_{\gamma G}(\mathbf{k}_i) \bar{W}_G(\mathbf{k}_i)}}. \quad (8.20)$$

The 2D cylindrically-averaged SS is likewise defined:

$$S_\gamma(k_\perp, k_\parallel) = \frac{\sum_{i \in k_\perp} \tilde{I}^2(\mathbf{k}_i) \tilde{I}^*(\mathbf{k}_i) \sqrt{\bar{W}_{\gamma G}(\mathbf{k}_i) \bar{W}_G(\mathbf{k}_i)}}{\sum_{i \in |\mathbf{k}|} \sqrt{\bar{W}_{\gamma G}(\mathbf{k}_i) \bar{W}_G(\mathbf{k}_i)}}. \quad (8.21)$$

In Equations 8.20 and 8.21 the factor $\sqrt{\bar{W}_{\gamma G}(\mathbf{k}_i) \bar{W}_G(\mathbf{k}_i)}$ is the geometric mean of the frequency averaged gridded weights. The geometric mean arises in the unit conversion performed in the appendix Section 8.C, where we show the unit conversion calculation for the SS from $\text{Jy}^3 \text{ Hz}^2 \text{ Sr}^{-1}$ to $\text{mK}^3 h^{-3} \text{ Mpc}^3$. In the conversion the field of view probed by the SS is the geometric mean of the skew visibility field of view Ω_γ and the visibility field of view Ω (refer to the derivation for further details). Furthermore, we found that the geometric mean of the weights yielded the best agreement with the expected SS, compared to the product of the weights. This is due to the weights being the convolution of the gridding kernel with the baseline distribution; the gridding kernel area is inversely proportional to the field of view. The SS calculation described in Equations 8.20 and 8.21 has been implemented into the OSIRIS pipeline.

8.5 Foreground Spectra

Using WODEN and the sky-model described in Section 8.2, we simulate the visibilities for the EoR0 and EoR1 fields. With the simulated visibilities, and the method described in Section 8.4.3, we calculate the spherically, and cylindrically averaged PS, SS and normalised SS for both fields, using OSIRIS .

8.5.1 Foreground Power Spectrum

Figure 8.1 shows the 2D PS estimated from the WODEN simulated visibilities for the EoR0 and EoR1 fields (Subfigures 8.1a and 8.1b, respectively). We see the characteristic foreground wedge structure, where the main component of the wedge feature is due to point sources in the main lobe of the MWA primary beam (Morales et al., 2012;




Figure 8.1: EoR0 (panel (a)) and EoR1 (panel (b)) WODEN simulated foreground 2D PS. Characteristic foreground wedge structure is present, along with some wedge artefacts that propagate through the 2D PS line of sight modes at higher k_{\perp} . Leakage can also be seen in the EoR window.




Figure 8.2: The 1D PS for the EoR0 (dashed dot line) and the EoR1 (dashed line) fields.

Vedantham et al., 2012; Trott et al., 2012). Leakage can be seen at higher k_{\perp} modes across all the line of sight modes. This leakage results from the patchy nature of the MWA uv -coverage. Furthermore, leakage into the EoR window can be seen from the foreground wedge, which is due to the chromaticity of the primary beam. Leakage is more substantial at modes greater than $k_{\perp} \gtrsim 0.06 \text{ Mpc}^{-1}$, where the lower density of visibilities leads to gaps in the uv -plane.

Figure 8.2 shows the 1D spherically averaged PS for the WODEN simulated EoR0 and EoR1 fields. The EoR1 field has a higher amplitude due to the presence of Fornax A at the edge of the main lobe of the primary beam. This is most notable at larger more extended scales $k < 0.1 \text{ Mpc}^{-1}$. In contrast the EoR0 field does not contain any bright extended sources, therefore the structure of the EoR0 PS is similar to the expected PS for a point source. The sharp drop at $k > 0.3, \text{ Mpc}^{-1}$ corresponds to the 300λ cutoff in uv -space. Beyond this point only $k_{\parallel} > 0$ contribute to the PS, resulting in a drop in power.




Figure 8.3: EoR0 (a) and EoR1 (b) 2D foreground SS with the primary beam attenuation. Grey pixels have negative values. The characteristic foreground wedge is apparent in both cases, with additional wedge fringes visible. These wedge fringes are sidelobes generated by the spectral blackman-harris taper.

8.5.2 Foreground Skew Spectrum

In this section we calculate the SS from the simulated EoR0 and EoR1 visibilities, and the skew visibilities calculated using the method described in Section 8.4.2.

Figure 8.3 shows the 2D SS for the WODEN simulated EoR0 and EoR1 fields (Sub-figures 8.3a and 8.3b, respectively). In both cases we see a foreground wedge, similar to the 2D WODEN PS case in Figure 8.1. Notably, we see less spectral leakage than the 2D PS, in part due to the quadratic primary beam attenuation of the squared temperature field. Sources are significantly more attenuated, resulting in less spectral leakage across the line of sight modes.

However, there is a clear difference, the SS outside of the wedge appears to oscillate between negative and positive values. Particularly modes with $k_{\perp} \gtrsim 0.1 \text{ Mpc}^{-1}$ appear to be dominated by sidelobe features. The attenuation of these sidelobes relative to




Figure 8.4: The 1D SS for the EoR0 (dashed dot line) and the EoR1 (dashed line) fields.

the foreground wedge is too low to be from the blackman-harris spectral tapering⁵. Notably changing the gridding kernel to a Gaussian function changes the structure of these sidelobes. Furthermore, the parallel nature of these sidelobes with the foreground wedge, indicates this is likely a result of some gridding kernel artefact. Since this effect is not seen with the PS at this level, it could be related to the different gridding kernel sizes required to calculate the SS. It is possible that when multiplying the Fourier grids of the quadratic temperature field and the temperature field, that a discontinuity manifests as these sidelobes for small relative values of the SS. Importantly, this effect is small, and does not significantly impact the results of this work. Understanding the origin of these sidelobes is therefore the focus of future work.

Figure 8.4 shows the 1D spherically averaged SS for the EoR0 and EoR1 fields. Again we see the EoR1 field has a higher amplitude than the EoR0 field. Interestingly, the sidelobes observed in Figure 8.3 do not appear to impact the 1D SS, and the morphology of the spectra seem to closely resemble that of the PS. This is unsurprising, since we saw in Chapter 7, that the SS amplitude is modulated by the PS.

⁵Blackman-harris first sidelobe attenuation is -90 dB or 9 orders of magnitude from peak amplitude.




Figure 8.5: The normalised SS for the EoR0 (dash dotted line) and the EoR1 (dashed line) fields.

8.5.3 Foreground Normalised Skew Spectrum

In this section we describe the normalised SS denoted by $\gamma(k)$ for the EoR0 and EoR1 foregrounds models. We recall from Chapter 7 that the normalised SS is defined as the ratio of the dimensionless SS and PS:

$$\gamma(k) = \frac{\Delta_{T^2T}^2(k)}{(\Delta_T^2(k))^{3/2}}, \quad (8.22)$$

where Δ^2 denotes the dimensionless spectrum:

$$\Delta_x^2(k) = \frac{k^3}{2\pi^2} P_x(k). \quad (8.23)$$

The subscripts T and T^2T denote the PS and the SS respectively, each with units of mK^2 and mK^3 .

Figure 8.5 shows the normalised SS for the EoR0 and EoR1 fields. Interestingly, the EoR0 field has a higher normalised SS than the EoR1 field. This contrasts with the PS and SS where the trend is reversed. This results from the asymmetry in the




Figure 8.6: Spherically averaged 1D PS of the Fiducial 21cm signal. The solid black line is the lightcone estimated PS, and the dashed dotted line is the `wODEN` instrumentally simulated PS.

difference in amplitude between the SS and PS for both fields, where the SS amplitude difference is smaller. Since $\gamma(k)$ is the SS normalised by the PS, this results in a larger rescaling for the EoR1 field. The morphology of both spectra are notably similar, with some differences in the location of the knee like structure at $k \sim 0.6 \text{ Mpc}^{-1}$.

8.6 Fiducial 21cm Spectra

In this section we describe the calculation and estimation of the fiducial 21cm signal described in Section 8.3. For comparison with the instrumentally simulated spectra we present in this section, we also determine the PS, SS and therefore the normalised SS, directly from the simulation lightcone. We directly compare the instrumental simulated spectra to the lightcone estimated spectra. Using the auto-correlations and the `wODEN` simulated visibilities, we estimate the skew visibilities and therefore the SS of the signal.


8.6.1 Power Spectrum

Figure 8.6 shows the spherically averaged PS for the fiducial lightcone signal (solid black), and the instrumentally simulated signal (dashed dotted line). The visibility simulation was performed by WODEN which we described in Section 8.3.1. From the visibilities we calculate the instrumental PS using the PS component of the OSIRIS pipeline described in Chapter 6.

In general there is good agreement between the PS in Figure 8.6, however signal loss is present on most scales, primarily on the largest scales ($k < 0.1 h \text{Mpc}^{-1}$). To investigate this signal loss we also calculate the cylindrically averaged 2D PS for both the lightcone and the visibility simulation.


Figure 8.7 shows the 2D PS for the fiducial lightcone (Subfigure 8.7a), the visibility simulation (Subfigure 8.7b), and the ratio between the two (Subfigure 8.7c). In Subfigure 8.7b mode mixing is apparent, with there being relatively more power at higher k_{\perp} and k_{\parallel} values than in Subfigure 8.7a. The Ratio subplot in Subfigure 8.7c, shows that there is more power in the instrumental simulation in the smaller angular modes ($0.06 h \text{Mpc}^{-1} < k_{\perp} < 3 h \text{Mpc}^{-1}$), and less power at the larger scale angular modes ($k_{\perp} < 0.06 h \text{Mpc}^{-1}$). In both cases the power difference is fairly constant with k_{\parallel} , except for the highest k_{\perp} modes which have ratio values that decrease when $k_{\parallel} > 0.1 h \text{Mpc}^{-1}$. The difference in the instrumental 21cm signal and the fiducial signal is primarily dependent on the k_{\perp} .

Figure 8.8 shows the mean ratio as a function of k_{\perp} (black), compared to unity (dashed grey line), with the mean ratio displayed with the solid blue line. We see in Figure 8.8 that the best agreement occurs between $k_{\perp} = 0.05 - 0.2$. The ratio at large scales $k_{\perp} \leq 0.05 h \text{Mpc}^{-1}$ coincides with the increase in baseline density on the shortest baselines. The gridding kernels on the smallest baselines significantly overlap, which results in decorrelation. This leads to an underestimate in signal and thus to an increased ratio. The decorrelation effects that result from the baseline distribution are known (Morales et al., 2012), and the signal loss is corrected for by multiplying by a decorrelation factor which is ~ 2 . To test whether the signal loss is specific to



(a) 21cm Fiducial 2D PS

(b) 21cm Instrumental 2D PS



(c) 21cm 2D PS Ratio

Figure 8.7: The 2D PS for the Fiducial signal (panel (a)), the instrumental WODEN simulated 2D PS (panel (b)), the colourbar scale is the same for panels (a) and (b). Panel (c) shows the ratio of the fiducial to instrumental minus one. The colourbar scale is log symmetric, with zero representing good agreement between both PS. Red indicates more signal in the Fiducial PS, and blue indicates more signal in the instrumental PS.




Figure 8.8: Mean 2D PS ratio for the fiducial 21cm signal, and the instrumental simulated 21cm signal, this is shown with the solid black line. The dashed grey horizontal line at 1 marks a 2D ratio of unity. The solid blue line marks the average of the ratio.

OSIRIS Line et al. (in prep) computed the PS from the WODEN visibilities using the CHIPS PS estimation pipeline. The resulting PS compared to the lightcone estimated PS did not show signal loss on the large scales like the OSIRIS estimated PS. These results are in preparation, and are not included here, but highlight the importance of comparison with other pipelines in determining systematic errors. We will return to these effects in the discussion.

8.6.2 Skew Spectrum

Figure 8.9 shows the spherically averaged SS for the Fiducial lightcone signal (solid black), and the instrumentally simulated signal (dashed dotted line). We see remarkably good agreement between the instrumental SS and the fiducial SS at scales $k > 0.1 \text{ Mpc}^{-1}$. The SS however is over estimated at larger scales, showing the opposite systematic compared to the instrumental PS.

Figure 8.10 shows the 2D SS for the Fiducial lightcone (Subfigure 8.10a), the WODEN simulation (Subfigure 8.10b), and the ratio minus one between the two (Subfig-




Figure 8.9: Spherically averaged 1D SS of the Fiducial 21cm signal. The solid black line is the lightcone estimate SS, and the dashed dotted line is the instrumentally simulated SS.

ure 8.10c). There is a concentration of signal at the low k_{\perp} and k_{\parallel} for the instrumental SS. Furthermore, there appears to be an overestimate of signal at $k_{\parallel} \sim 0.3 h \text{ Mpc}^{-1}$ for small k_{\perp} . Additionally we see evidence of mode mixing at higher k_{\perp} and k_{\parallel} for the instrumental SS, similar to the PS. The trends in the SS ratio are effectively the reverse of those seen in Subfigure 8.7c. The SS Ratio likewise is fairly constant with k_{\parallel} , and changes mostly as a function of k_{\perp} . Notably, there is a region at $k_{\parallel} \sim 0.3 h \text{ Mpc}^{-1}$ which appears to have relatively flat ratio as a function of k_{\perp} , this appears as a minimum at the $k_{\perp} > 0.1 h \text{ Mpc}^{-1}$ modes. A similar structure does appear to be present in Subfigure 8.7c for the 2D PS ratio, however it is less prominent than in the SS case.

Figure 8.11 shows the mean 2D SS ratio as a function of k_{\perp} averaged across the line of sight modes. For the instrumental SS we assume a similar decorrelation factor, however this assumption is likely inaccurate. Since CHIPS does not currently have the capability to calculate the SS, no comparison yet exists. However, systematics that can be identified and fixed in the PS calculation through a CHIPS comparison, will likely positively impact the SS estimation. Further work is required to determine the decorrelation effects for the SS compared to the PS to understand the source of the





Figure 8.10: The 2D SS for the Fiducial signal (panel (a)), the instrumental simulation 2D SS (panel (b)), the colourbar scale is the same for panels (a) and (b). Panel (c) shows the Fiducial to instrumental ratio minus one. The colourbar scale is log symmetric, with zero representing good agreement between both SS. Red indicates more signal in the Fiducial SS, and blue indicates more signal in the instrumental SS.




Figure 8.11: Mean 2D SS ratio for the fiducial 21cm signal, and the instrumental simulated 21cm signal, this is shown as the solid black line. The grey dashed horizontal line marks unity, and the solid blue line marks the mean ratio.

systematics. We discuss this further at the end of this chapter.

8.6.3 Normalised Skew Spectrum

Figure 8.12 shows the normalised SS calculated from the fiducial lightcone (solid black), and from the instrumental SS and PS (dashed dot line). For the fiducial lightcone we see a peak and trough structure, with a minimum occurring around $k \sim 0.2 \text{ Mpc}^{-1}$, and a maxima occurring between $1 \text{ Mpc}^{-1} < k < 2 \text{ Mpc}^{-1}$. This structure is similar to the one observed in Chapter 7, however we note that the types of simulations performed in that work are different from the one used in this chapter. Notably, the fiducial 21cm simulation assumes that $T_S \gg T_{\text{CMB}}$, which simplifies the calculation of δT_b (Greig et al., 2022b). Furthermore, the normalised SS considered in Chapter 7 was calculated for a coeval box at a fixed redshift. The simulations in this work use lightcones, which would have signal evolution as a function of frequency (redshift). This will affect the observed features. Analysing the lightcone effects on the normalised SS is outside the scope of this work.




Figure 8.12: Normalised SS of the fiducial 21cm signal. The solid black line is the lightcone estimate of the normalised SS, and the dashed dotted line is the WODEN simulated instrumental normalised SS.

The discrepancies between the instrumental and fiducial lightcone normalised SS arise from the deviations between the PS and SS estimates. The largest difference occurs at large spatial scales or low k , where there is significant signal loss in the PS. This signal loss is magnified when calculating the normalised SS, because the PS is raised to the power of $3/2$. To demonstrate this impact, let us consider some mode which has signal loss in the PS of 50%. Raising this mode to the power of $3/2$ this becomes ~ 0.35 . This implies, that signal loss of 50% in the PS for a given mode, leads to an increase of ~ 2.83 in the normalised SS for the same mode. Since signal loss occurs where we expect to see the minima in the normalised SS, this results in a shift of the location of the minima in the instrumental case ($\sim 0.3 \text{ Mpc}^{-1}$). Similarly, for the instrumental SS which underestimates the fiducial lightcone SS at small spatial scales, we see the peak position migrate inward to larger scales. The systematic effects that impact both the SS, and the PS will both propagate through the normalised SS. Thus, the normalised SS is inherently more sensitive to instrumental effects than its constituent spectra. We leave this analysis to future work.

8.7 Foregrounds and 21cm Poly Spectra

In the previous sections we calculated the PS, SS and normalised SS for the foreground and 21cm WODEN simulated data. In this section we compare the foreground spectra and 21cm spectra. In particular, we employ the foreground avoidance method where we estimate the expected PS, SS and norm SS from the EoR window. In this chapter we do not simulate subtraction of the sky-model from the visibilities, we leave this to future work, we simply compare the amplitudes. The EoR window in this work is defined by all k modes with $k_{\perp} \leq 0.06 \text{ Mpc}^{-1}$ and $k_{\parallel} \geq 0.1 \text{ Mpc}^{-1}$. Additionally we employ a horizon cut defined by Morales et al. (2012). Figure 8.13 shows an example of the wedge cut 2D PS for the EoR0, EoR1 and 21cm simulated visibilities.

8.7.1 Window Power Spectra

Figure 8.14 shows the window power calculated from the wedge cuts shown in Figure 8.13 for the EoR0 (dash dotted line), the EoR1 (dashed line) and the 21cm simulation (solid black line). As expected the foreground leakage is significant without the subtraction of foreground sources, resulting in a power difference with the 21cm signal on the order of $10^3 - 10^4$ for all modes. We note that the EoR1 simulation has more power than the EoR0 simulation.

8.7.2 Window Skew Spectra

Figure 8.15 shows the window SS calculated for the EoR0, EoR1 and 21cm simulations, with the same lines corresponding to Figure 8.14. Subfigure 8.15a shows the SS in an asymmetric log scale with the linear region of the axis ranging from -10^8 to 10^8 . As a result, since the amplitude of the 21cm SS is significantly less than the foreground SS it appears as a flat straight line. Therefore, to better compare the differences in amplitude in Subfigure 8.15b we show the absolute SS for all three simulations. From Subfigure 8.15b we see that the amplitude of the 21cm SS is at most six orders of magnitude less than the foreground SS amplitude for all modes.

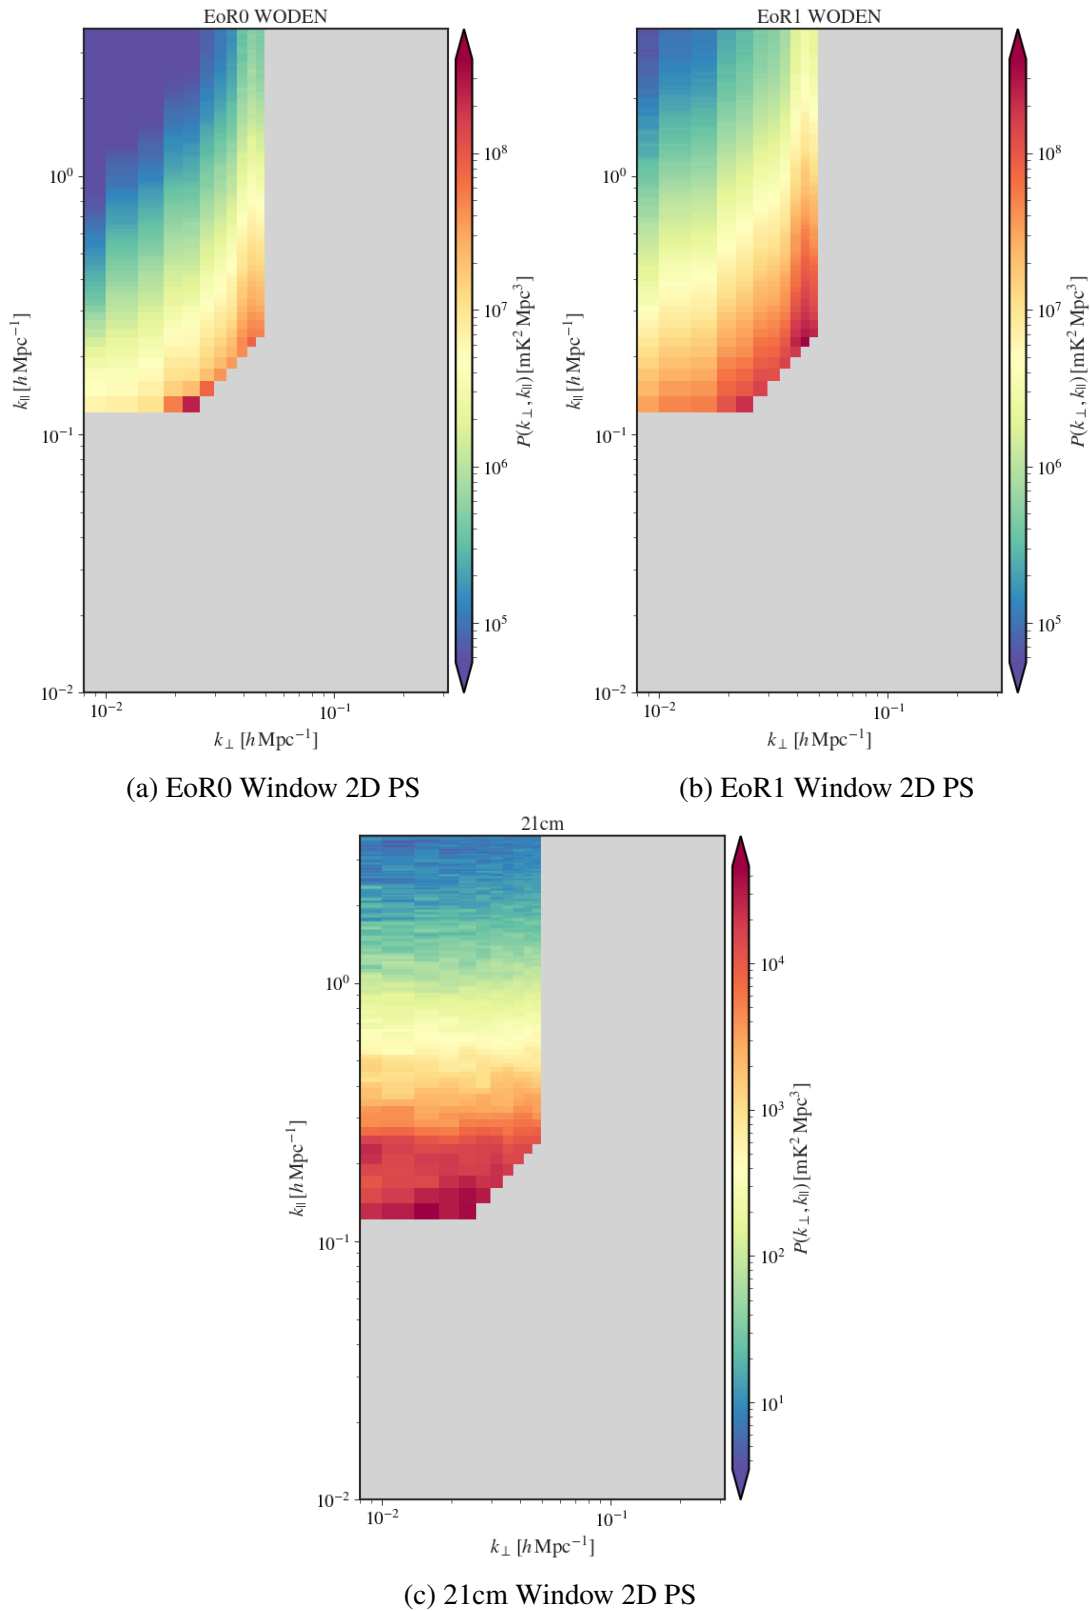


Figure 8.13: wODEN simulated EoR0, EoR1, and 21cm 2D PS with wedge cuts applied at $k_{\perp} \leq 0.06 \text{ Mpc}^{-1}$ and $k_{\parallel} \geq 0.1 \text{ Mpc}^{-1}$. The colour bar scale for Subfigure 8.13a and 8.13b is the same. The grey pixels indicate the data that has been set to NaN in the wedge cut.





Figure 8.14: Compares the EoR0 (dash dotted line), EoR1 (dashed line) and 21cm (solid black line) window power.


Notably, the foreground SS in Figure 8.15 appear to contain sidelobes as a function of k . It is not clear if these sidelobes are the result of the blackman-harris spectral taper, or are also a result of the gridding artefacts discussed in Section 8.5.2. In contrast to the sidelobes in Figure 8.3, these sidelobes are > 9 orders of magnitude (or -90 dB) less than the foreground wedge. This is the expected level of attenuation for the blackman-harris spectral taper. However, we do not see these sidelobes in the 1D window PS. A possible explanation, is that the skewness of the foregrounds is close to zero. The SS is a measure of the central third order moment as a function of spatial scale. We can therefore describe the central third order moment as the skewness modulated by the cubed standard deviation (or the variance to the power of $3/2$). Thus, large values can obscure a small central third order moment. We will see this normalisation in practice in the following section.

8.7.3 Window Normalised Skew Spectra

In this section we calculate the normalised window SS for the EoR0, EoR1 and 21cm simulated visibilities, using the same method outlined in Section 8.5.3. Figure 8.16




(a) Wedge cut SS




(b) Absolute wedge cut SS

Figure 8.15: Shows the SS for the EoR0 (dash dotted line), the EoR1 (dashed line), and the simulated 21cm signal (solid line). Subfigure 8.15a is in the asymmetrical log scale, where the linear part of the axis ranges from -10^8 to 10^8 . Subfigure 8.15b shows the absolute SS for all three simulations.



(a) EoR0 2D SS WODEN



(b) EoR1 2D SS WODEN

Figure 8.16: Shows the normalised SS for the EoR0 (dash dotted line), the EoR1 (dashed line), and the simulated 21cm signal (solid line). Subfigure 8.16a is in the asymmetrical log scale, where the linear part of the axis ranges from -1 to 1 . Subfigure 8.15b shows the absolute normalised SS for all three simulations.

shows the normalised SS in Subfigure 8.16a and the absolute normalised SS in Subfigure 8.16b. In Subfigure 8.16a the linear region of the asymmetric log-axis ranges from -1 to 1 . We see a similar relationship as the SS, however, the amplitude of the foreground models is less than the 21cm signal for almost all spatial modes except the largest ones. Additionally, the amplitude of the absolute normalised SS for the foreground spectra rapidly approaches zero, changing by 4 – 5 orders of magnitude. In contrast the 21cm normalised SS is relatively constant across all modes. Furthermore, we do not see the sidelobe structures in the 21cm signal. Notably, the 21cm signal has the characteristic trough and peak structure seen in the normalised SS throughout this chapter and the previous chapter.

Figure 8.16 naively suggests that the 21cm normalised SS is detectable relative to the foreground normalised SS. It is true that the skewness in the 21cm temperature field is greater than the foreground fields for these modes. However, real data contains both foregrounds, noise, and 21cm signal, all of which couple together. Furthermore, the relationship between the different components of the normalised SS is not linear, since it requires normalisation by the PS. Future work must consider the effects of noise, and residual foregrounds to accurately determine the viability of detecting the normalised 21cm SS.

8.8 Discussion and Conclusion

In this chapter, we simulated the visibilities for a set of two foreground observations, and a fiducial 21cm simulation using `WODEN`. We then develop a mathematical framework for estimating the instrumental SS, by calculating an intermediate product we call the skew visibilities. From the skew visibilities, and the visibilities we estimate the PS, SS, and normalised SS for the foreground and 21cm simulations. We then compare the PS, SS and normalised SS for the window modes for all simulations. Leakage from the foreground wedge into the window modes is significant, and is three orders of magnitude greater than the 21cm PS. Likewise, the foreground SS is dominated by the same leakage and is approximately six orders of magnitude greater than the 21cm

SS. We find that the amplitude of the normalised SS for the fiducial 21cm signal in the window is one to two orders of magnitude greater than the foregrounds. Additionally, we compared the PS, SS and normalised SS calculated directly from the 21cm simulated lightcone image, with the WODEN visibility simulated output. We found good agreement between the two, however, systematic and instrumental effects resulted in signal loss for the WODEN simulated 21cm SS. Furthermore, systematic signal loss was present on most modes for the instrumental PS, signal loss not seen with the CHIPS PS estimation.

There are several important caveats to the method in this work. In particular, the skew visibility calculation method is limited by the requirement of the $(\mathbf{u}) = (0, 0)$ mode (DC mode) to accurately estimate $\tilde{I}_A^2(\mathbf{u})$. This mode corresponds to the auto-correlations measured by the individual MWA tiles, which in principle should be the same. Without the auto-correlations simulated by WODEN, we found the SS for the foreground and the 21cm models oscillated from negative to positive across all modes, not just the low signal window modes. This occurs, because the DC mode in the convolution, is multiplied by the visibility for that baseline, and is therefore the dominant term in the convolution. The requirement to use the auto-correlation limits this method because of two main effects. First, each tile will have different receiver and antenna noise components which are additive to the auto-correlation. These terms disappear in the cross-correlation required to calculate the visibility. Averaging the auto-correlations of all tiles together will help minimise these effects, but not remove them. If the auto-correlations are to be used with real data, a robust method for estimating the noise power will be required. Secondly, the auto-correlations are effectively the beam-weighted average (or sum depending on the normalisation) of the sky brightness, and therefore see all the foregrounds (including the difficult-to-model diffuse emission). Accurate modelling and subtraction of the foregrounds is therefore paramount for determining the 21cm SS directly from visibilities. A potential solution could be to use a pseudo auto-correlation. This could be constructed by using a large enough constant in lieu of the real instrumental auto-correlations. This could be determined in an iterative process, or approximated using knowledge of the instrument,

and the sky-model. More work is needed to develop such a method, but a pseudo-SS could be a viable alternative if usage of the real auto-correlations proves untenable.

In this work we do not consider thermal noise and how this impacts the expected instrumental foreground and 21cm SS. For a single baseline with Gaussian thermal noise, we would expect the corresponding skew visibility noise to have a chi squared distributed component. This however, could potentially be approximated with a Gaussian distribution. As the number of visibility product terms increases in the skew visibility convolution, the resulting chi-squared distribution becomes more Gaussian. The situation is perhaps not so straightforward, and requires careful consideration, and is therefore the focus of planned future work.

A realistic test of the skew visibility calculation method would be to observe a bright calibrator source (Pictor A, 3C444, for example) with the MWA. Notably, as a result of the GLEAM survey there is a large volume of archival Phase I MWA data of calibrator sources (Wayth et al., 2015). A comparison of a real observation with a WODEN simulation for the SS, with realistic noise would be a good test for the method outlined in this work. For a bright source, this test could be performed on a relatively small observation.

The normalised SS is more affected by systematic effects than the SS or the PS as evidenced in Figure 8.12. In particular the largest and smallest scales are most impacted by systematic signal loss or over estimation for both the PS or SS (Figures 8.6 and 8.9). However, we have accurate estimation in the spatial modes $0.1 h \text{ Mpc}^{-1} < k < 1 h \text{ Mpc}^{-1}$. This is the regime of interest for measuring the 21cm signal, since large-scale modes are dominated by foregrounds, which are typically avoided and removed when calculating the 1D PS. The biggest improvement can be found by understanding the systematic signal loss in the PS; this has the biggest overall effect on the 21cm normalised SS. As mentioned, this signal loss at large scales is not seen in CHIPS (Line et al. in prep). Due to the simplicity of the SS method, and the considerable overlap with the PS calculation, future should consider expanding CHIPS to incorporate the SS calculation. This will provide an independent comparison that will be useful for

probing the systematics seen in this work.

8.A Analytic Toy Model

From Figure 8.5 it is clear that tapering performed as a function of (\mathbf{u}, ν) affects the structure of the normalised SS. This couples with the intrinsic structure from both point sources and extended sources. To understand these effects, and how they impact the normalised foreground SS, we develop a point source model that is used as a comparative tool to analyse the results in Figure 8.5. To build this model, we start by defining the point source sky brightness distribution $I(\mathbf{l}, \nu)$:

$$I(\mathbf{l}, \nu) = \sum_{i=1}^{N_s} I_i \delta(\mathbf{l} - \mathbf{l}_i), \quad (8.24)$$

where I_i is the i th source intensity, N_s is the number of sources, \mathbf{l}_i is the vector of the source direction cosines, ν is the frequency, and δ is the Dirac delta function. For simplicity we assume $I_i(\nu) = I_i \forall \nu$ (constant). We also do not consider any primary beam attenuation. A natural extension would be to consider a spectrally flat primary beam function, which would effectively be a different constant factor multiplied by each source.

8.A.1 Analytic Power Spectrum

Using the point source model described in Equation 8.24, we can derive an analytical estimate of the visibilities by substituting this expression into the measurement equation (3.9):

$$\mathcal{V}(\mathbf{u}) = \sum_{i=1}^{N_s} I_i \int_{\mathbb{R}^2} d^2\mathbf{l} e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} \delta(\mathbf{l} - \mathbf{l}_i). \quad (8.25)$$

In Equation 8.25 we have exchanged the order of the integral and the summation; this is possible because the integral is a linear operation. Performing the integration:

$$\mathcal{V}(\mathbf{u}) = \sum_{i=1}^{N_s} S_i e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_i}. \quad (8.26)$$

Here we have replaced I_i by S_i , where I_i has units of Jy/Sr and S_i has units of Jy. This conversion happens because the Fourier transform of the sky-brightness temperature integrates over all solid angles, dropping the per steradian dependence.

In reality with discrete data, we perform Fourier transforms over limited ranges, assuming periodic boundary conditions. It is standard practice to taper data with a window function; in general this window function can be a function of (\mathbf{u}, ν) . Applying the window function $W(\mathbf{u}, \nu)$ we write the frequency-dependent visibility function:

$$\mathcal{V}(\mathbf{u}, \nu) = W(\mathbf{u}, \nu) \sum_{i=1}^{N_s} S_i e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_i}. \quad (8.27)$$

In Equation 8.27 we are assuming that we sample the whole uv -plane, however, real radio interferometers have frequency dependent uv -samples. We discuss this at the end of the chapter. Next we Fourier transform $\mathcal{V}(\mathbf{u}, \nu)$ with respect to frequency:

$$\tilde{\mathcal{V}}(\mathbf{u}, \eta) = \Delta \nu \tilde{W}(\mathbf{u}, \eta) \sum_{i=1}^{N_s} S_i e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_i}, \quad (8.28)$$

where $\tilde{W}(\mathbf{u}, \eta)$ is the Fourier transform of the window function, and $\Delta \nu$ is the frequency bandwidth of the observation. In the absence of tapering, $\tilde{W}(\mathbf{u}, \eta) \propto \text{sinc}(\eta)$ since the window is a boxcar function. We then define the power function as $P(\mathbf{u}, \eta) = \tilde{\mathcal{V}}(\mathbf{u}, \eta) \tilde{\mathcal{V}}^*(\mathbf{u}, \eta)$, where the superscript $*$ denotes the complex conjugate:

$$P(\mathbf{k}) = C_P |\tilde{W}(\mathbf{k})|^2 \sum_{i=1}^{N_s} \sum_{j=1}^{N_s} S_i S_j e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta \mathbf{l}_{ij}}. \quad (8.29)$$

Here C_P is the conversion factor from $\text{Jy}^2 \text{Hz}^2$ to $\text{mK}^2 \text{Mpc}^3 \text{h}^{-3}$, and C_k converts from \mathbf{u} to \mathbf{k}_\perp . In Equation 8.29 we switch to vector notation where $\mathbf{k} = [k_x, k_y, k_z]$, $\mathbf{k}_\perp = [k_x, k_y]$, and $\Delta \mathbf{l}_{ij} = [l_i - l_j, m_i - m_j]$. $P(\mathbf{k})$ is a real quantity since $\Delta \mathbf{l}_{ji} = -\Delta \mathbf{l}_{ij}$, and $S_j S_i = S_i S_j$, however, we will keep the exponential notation for now, for consistency

with the following section. We can express Equation 8.29 in an expanded form:

$$P(\mathbf{k}) = C_P |\tilde{W}(\mathbf{k})|^2 \left[\sum_{n=1}^{N_s} S_n^2 + \sum_{i=1}^{N_s} \sum_{j=1, j \neq i}^{N_s-1} S_i S_j e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta \mathbf{l}_{ij}} \right], \quad (8.30)$$

where we separate the terms where $i = j$ thus $\Delta \mathbf{l}_{ii} = 0$. We define the spherical average of $P(\mathbf{k})$ to be $P(k) = \langle P(\mathbf{k}) \rangle_{|\mathbf{k}| \in k}$. Performing the spherical average in Equation 8.30:

$$P(k) = C_P \langle |\tilde{W}(\mathbf{k})|^2 \rangle \sum_{n=1}^{N_s} S_n^2 + C_P \left\langle |\tilde{W}(\mathbf{k})|^2 \sum_{i=1}^{N_s} \sum_{j=1, j \neq i}^{N_s-1} S_i S_j e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta \mathbf{l}_{ij}} \right\rangle. \quad (8.31)$$

We have two terms, one containing all the cross terms, and one containing all the independent quadratic terms. To simplify the model in the following sections we assume the cross terms are much smaller than the squared terms. This effectively assumes that the point sources are spatially uncorrelated, and therefore independent.

8.A.2 Analytic Skew Spectrum

To derive the point source SS, we employ a similar method as the previous section. We define the *skew visibilities* as the Fourier transform of the squared sky brightness distribution:

$$\mathcal{V}_\gamma(\mathbf{u}) = \int_{\mathbb{R}^2} d^2 \mathbf{l} e^{-2\pi i \mathbf{u} \cdot \mathbf{l}} \left(\sum_{i=1}^{N_s} I_i \delta(\mathbf{l} - \mathbf{l}_i) \right)^2. \quad (8.32)$$

Equivalently we can describe \mathcal{V}_γ through the convolution of the visibilities:

$$\mathcal{V}_\gamma(\mathbf{u}) \equiv \Delta \Omega (\mathcal{V}(\mathbf{u}) * \mathcal{V}(\mathbf{u})), \quad (8.33)$$

where $\Delta \Omega$ is the solid angle of the resolution element, ensuring that the units are $\text{Jy}^2 \text{Sr}^{-1}$ as expected from Equation 8.32. The squared term in Equation 8.32 can be

expanded using the identity $\left(\sum_{n=1}^N a_n\right)^2 = \sum_{n=1}^N a_n^2 + 2 \sum_{j=1}^N \sum_{i=1}^{j-1} a_i a_j$:

$$\begin{aligned} \left(\sum_{i=1}^{N_s} I_i \delta(\mathbf{1} - \mathbf{l}_i)\right)^2 &= \sum_{n=1}^{N_s} I_n^2 \delta(\mathbf{1} - \mathbf{l}_n) \\ &+ 2 \sum_{j=1}^{N_s} \sum_{i=1}^{j-1} I_i I_j \delta(\mathbf{1} - \mathbf{l}_i) \delta(\mathbf{1} - \mathbf{l}_j), \end{aligned} \quad (8.34)$$

where $\delta^2(\mathbf{1} - \mathbf{l}_i) = \delta(\mathbf{1} - \mathbf{l}_i)$. We also note that $\delta(x - a)\delta(x - b) = 0$ when $a \neq b$, therefore:

$$\left(\sum_{i=1}^{N_s} I_i \delta(\mathbf{1} - \mathbf{l}_i)\right)^2 = \sum_{i=1}^{N_s} I_i^2 \delta(\mathbf{1} - \mathbf{l}_i). \quad (8.35)$$

Substituting Equation 8.35 into 8.32 and performing the integral:

$$\mathcal{V}_\gamma(\mathbf{u}, \nu) = \frac{1}{\Delta\Omega} \sum_{i=1}^{N_s} S_i^2 e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_i}, \quad (8.36)$$

where we noted that the Fourier transform of I^2 has units of Jy^2/Sr , and the flux density S has units of Jy , so we multiply by a factor of $1/\Delta\Omega$ to ensure that the units are correct for the discrete data. This allows us to express Equation 8.36 in terms of the integrated flux density S_i .

In the same process as the previous section $\mathcal{V}_\gamma(\mathbf{u}, \nu) = W(\mathbf{u}, \nu) \mathcal{V}_\gamma(\mathbf{u})$. $\mathcal{V}_\gamma(\mathbf{u}, \nu)$ is Fourier transformed as a function of frequency to obtain:

$$\tilde{\mathcal{V}}_\gamma(\mathbf{u}, \eta) = \frac{\Delta\nu}{\Delta\Omega} \tilde{W}(\mathbf{u}, \eta) \sum_{i=1}^{N_s} S_i^2 e^{-2\pi i \mathbf{u} \cdot \mathbf{l}_i}. \quad (8.37)$$

$\tilde{\mathcal{V}}_\gamma(\mathbf{u}, \eta)$ has the same form as equation 8.28. The SS is defined as $S_\gamma(\mathbf{u}, \eta) \equiv \tilde{\mathcal{V}}_\gamma(\mathbf{u}, \eta) \tilde{\mathcal{V}}_\gamma^*(\mathbf{u}, \eta)$:

$$S_\gamma(\mathbf{k}) = C_\gamma |\tilde{W}(\mathbf{k})|^2 \sum_{i=1}^{N_s} \sum_{j=1}^{N_s} S_i^2 S_j^2 e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta \mathbf{l}_{ij}}. \quad (8.38)$$

Here, C_γ is the conversion from $\text{Jy}^3 \text{Hz}^2 \text{Sr}^{-1}$ to $\text{mK}^3 \text{Mpc}^3 \text{h}^{-3}$, and C_k converts \mathbf{u} to

\mathbf{k}_\perp . We expand Equation 8.38 for all the $\Delta\mathbf{l}_{ij} = 0$:

$$S_\gamma(\mathbf{k}) = C_\gamma |\tilde{W}(\mathbf{k})|^2 \left[\sum_{n=1}^{N_s} S_n^3 + \sum_{i=1}^{N_s} \sum_{j=1, j \neq i}^{N_s-1} S_i^2 S_j e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta\mathbf{l}_{ij}} \right]. \quad (8.39)$$

Next we take the spherical average of $S_\gamma(\mathbf{k})$:

$$S_\gamma(k) = C_\gamma \langle |\tilde{W}(\mathbf{k})|^2 \rangle \sum_{n=1}^{N_s} S_n^3 + C_\gamma \left\langle |\tilde{W}(\mathbf{k})|^2 \sum_{i=1}^{N_s} \sum_{j=1, j \neq i}^{N_s-1} S_i^2 S_j e^{-2\pi i C_k \mathbf{k}_\perp \cdot \Delta\mathbf{l}_{ij}} \right\rangle. \quad (8.40)$$

Equation 8.40 defines the point source SS. We note that the second term in the right hand side of Equation 8.40 is real, since the spherical average sums over all the conjugate terms in the radius k and $k + \Delta k$. Additionally, we expect this term to be negligible due to the point sources being spatially uncorrelated; we explore this in the following section.

8.A.3 Analytic Normalised Skew Spectrum

We now have expressions for the point source PS and SS. To calculate the normalised SS, we substitute Equation 8.23 into 8.22:

$$\gamma(k) = (2\pi^2)^{1/2} k^{-3/2} \frac{S_\gamma(k)}{P^{3/2}(k)}. \quad (8.41)$$

Before substituting in the SS and PS, we further simplify Equations 8.31 and 8.40 by assuming the cross terms are negligible:

$$P(k) \approx C_P \langle |\tilde{W}(\mathbf{k})|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2} \sum_{n=1}^{N_s} S_n^2 \quad (8.42)$$

$$S_\gamma(k) \approx C_\gamma \langle |\tilde{W}(\mathbf{k})|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2} \sum_{n=1}^{N_s} S_n^3. \quad (8.43)$$

We make this assumption because point sources are independent, and we expect the spherical average of the sum of the cross terms to be much less than the $i = j$ terms.

Substituting in the simplified expressions for the point source power and SS yields:

$$\gamma(k) = (2\pi^2)^{1/2} C_\gamma C_P^{-3/2} \frac{\sum_{n=1}^{N_s} S_n^3}{\left(\sum_{n=1}^{N_s} S_n^2\right)^{3/2}} k^{-3/2} \langle |\tilde{W}(\mathbf{k})|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2}. \quad (8.44)$$

In Equation 8.44 we see that the structure of $\gamma(k)$ for the toy model only depends on the $k^{-3/2}$ term, and the square root of the spherically averaged window power function $\langle |\tilde{W}(\mathbf{k})|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2}$ (for lack of a better name). The remaining terms are constants, and can be collapsed into a more compact general form. We evaluate the normalised SS at k_0 :

$$\gamma(k_0) = (2\pi^2)^{1/2} k_0^{-3/2} C_\gamma C_P^{-3/2} \frac{\sum_{n=1}^{N_s} S_n^3}{\left(\sum_{n=1}^{N_s} S_n^2\right)^{3/2}}, \quad (8.45)$$

where $\langle |\tilde{W}(k_0)|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2} = 1$ (assuming the filter is normalised). Multiplying through by $k_0^{3/2}$ we are left with a constant term:

$$\gamma_0 k_0^{3/2} = (2\pi^2)^{1/2} C_\gamma C_P^{-3/2} \frac{\sum_{n=1}^{N_s} S_n^3}{\left(\sum_{n=1}^{N_s} S_n^2\right)^{3/2}}. \quad (8.46)$$

Performing the final substitution for $\gamma_0 k_0^{3/2}$:

$$\gamma(k) = \gamma_0 \left(\frac{k}{k_0}\right)^{-3/2} \langle |\tilde{W}(\mathbf{k})|^2 \rangle_{|\mathbf{k}| \in k}^{-1/2}. \quad (8.47)$$

Equation 8.47 is the point source toy model for the normalised SS. The term γ_0 encapsulates the value of the normalised SS $\gamma(k)$ assuming a point source only model. The model is a product of a power law with a slope of $-3/2$ and the square root of the spherically averaged three dimensional windowing function. The generalised expression in Equation 8.47 is useful, because the constant $\gamma_0 k_0^{3/2}$ can be fit to any normalised SS. Any deviation from the measured normalised SS to the toy model would be an indication of non-Gaussianities from extended sources, or from other effects.

There is one important caveat for this toy model; it does not include a description for

the chromatic primary beam. A non-chromatic model would be trivial to include, and would be subsumed into the constant factors in the toy model. However, a chromatic beam with spectral behaviour would affect the frequency dependence of both the PS and the SS. However, we expect this effect to be smaller than the windowing on the structure of the normalised SS, and we discuss it more in the following section.

8.A.4 Toy Model Test

In this section we compare the toy model of Equation 8.47 to the expected point source model of Equation 8.41. We create a random sample of 50 point sources with uniformly distributed positions $l \in \mathcal{U}(-0.5, 0.5)$ and $m \in \mathcal{U}(-0.5, 0.5)$. This ensures that all point sources lie above the horizon. Each point source has a uniformly generated flux density $S_i \in \mathcal{U}(0, 1) \text{ Jy}$ ⁶. We calculate the expected power cube from Equation 8.30, and skew power cube from Equation 8.39. The sum is performed for each (k_x, k_y) pixel in the $k_{\parallel} = 0$ mode, and is saved for a cube of size $1291 \times 1291 \times 192$. In Equations 8.30 and 8.39 the sum only needs to be performed for every \mathbf{k}_{\perp} once, with the window function being a multiplicative factor; this greatly simplifies the computation. Due to computational constraints of calculating the sum, we consider a small sky-model containing only 50 sources for this example. For consistency, we perform the same tapering as a function of (\mathbf{u}, ν) as described in Section 8.5.1. This tapering is applied to a cube of ones with pixel dimensions of $1291 \times 1291 \times 384$, defining the window function $W(\mathbf{k}_{\perp}, \nu)$, which we then Fourier transform as a function of frequency to derive $\tilde{W}(\mathbf{k})$. This is then multiplied by the power or skew cube. The spherically averaged power and SS is then calculated using the `OSIRIS AVGSPHERICAL` function. The normalised SS is then calculated using Equation 8.41.

For the toy model we calculate the spherically averaged window function, and then take the square root. The constants γ_0 and k_0 are then determined from the first value of the previously calculated expected normalised SS, where k_0 is the first k mode in the average, and γ_0 is the corresponding value.

⁶Although nonphysical, this can be generalised for negative point sources.




Figure 8.17: Expected point source normalised SS for random 50 point source model (solid black), compared to the toy model (cyan).

Figure 8.17 shows the expected point source normalised SS calculated from Equation 8.41, compared to the toy model calculated from Equation 8.47. The bottom panel shows the residuals, which are on the order of a percent or less, which demonstrates good agreement. This result clearly demonstrates the impact the windowing function has on the data.

8.B 2D Radial Blackman-Harris Window

The radial two Blackman-Harris window is equivalent to the 1D definition, where the position is now the radial distance:

$$\begin{aligned}
 W_{BH}(r(\mathbf{u}, \mathbf{u}'); D) = & a_0 - a_1 \cos\left(\frac{2\pi(r(\mathbf{u}, \mathbf{u}') + \frac{D}{2})}{D}\right) + a_2 \cos\left(\frac{4\pi(r(\mathbf{u}, \mathbf{u}') + \frac{D}{2})}{D}\right) \\
 & - a_3 \cos\left(\frac{6\pi(r(\mathbf{u}, \mathbf{u}') + \frac{D}{2})}{D}\right),
 \end{aligned}
 \tag{8.48}$$

where D is the size of the window, and the radius is determined relative to some position vector \mathbf{u}' :

$$r(\mathbf{u}, \mathbf{u}') = |\mathbf{u} - \mathbf{u}'|. \quad (8.49)$$

8.B.1 Top Hat Radial Blackman-Harris Window

Using Equation 8.48 we can define a radial top hat Blackman-Harris window:

$$W_T(r(\mathbf{u}, \mathbf{u}')) = \begin{cases} 1, & |\mathbf{u}| \leq R \lambda \\ W_{BH}(r(\mathbf{u}, \mathbf{u}'); D), & \text{elsewhere,} \end{cases} \quad (8.50)$$

R is the threshold radius, which is typically set at 300λ

8.C Skew Spectrum Conversion Calculation

The cosmological unit conversion calculation for the SS, follows the same process as the PS calculation discussed in the appendix of Chapter 6. There are however some subtle differences that are important. In this case we have two equations, one defined by the Fourier transform of the quadratic temperature brightness, and another defined by the Fourier transform of the temperature field:

$$\Omega_\gamma \tilde{I}_b^2 = \frac{\lambda_o^4}{4k_b^2} \tilde{I}_\gamma^2 \quad [\text{K}^2 \text{ Sr}] \quad (8.51)$$

$$\Omega \tilde{T}_b = \frac{\lambda_o^2}{2k_b} \tilde{T} \quad [\text{K Sr}]. \quad (8.52)$$

The Fourier transform for both fields acquires a solid angle term, thus $\mathcal{F}[T_b] \propto \text{K Sr}$ and $\mathcal{F}[T_b^2] \propto \text{K}^2 \text{ Sr}$. We note that the tilde implies a Fourier transform, where $\tilde{I}_\gamma^2 \propto \text{Jy}^2 \text{ Sr}^{-1}$ and $\tilde{T}_\nu \propto \text{Jy}$. We can combined the two expressions above:

$$\tilde{T}_b^2 \tilde{T}_b \Omega \Omega_\gamma = \frac{\lambda_o^6}{8k_b^3} \langle \tilde{I} \tilde{I}_\gamma^2 \rangle [\text{K}^3 \text{Sr}^2], \quad (8.53)$$

Taking the expectation of both sides:

$$\langle \tilde{T}_b^2 \tilde{T}_b \rangle \bar{\Omega}^2 = \frac{\lambda_o^6}{8k_b^3} \langle \tilde{I} \tilde{I}_\gamma^2 \rangle [\text{K}^3 \text{Sr}^2], \quad (8.54)$$

Where $\bar{\Omega} = \sqrt{\Omega \Omega_\gamma}$ is the geometric mean field of view, we have $\Omega_\gamma = \Omega/2$, thus $\bar{\Omega} = \Omega/\sqrt{2}$. From Hogg (1999) we have an expression for comoving volume, which we rearrange for the solid angle $\bar{\Omega}$:

$$\bar{\Omega} = \frac{1}{\Delta\nu} \frac{\Delta V_C}{D_M^2(z) D_H} \frac{\nu_{21} E(z)}{(1+z)^2} [\text{Sr}], \quad (8.55)$$

Therefore:

$$10^9 \times \Delta\nu (1+z)^2 \frac{D_M^2(z) D_H}{\nu_{21} E(z)} \frac{\lambda_o^6}{8k_b^3} \frac{\langle \tilde{I} \tilde{I}_\gamma^2 \rangle}{\bar{\Omega}} = \Delta V_C \langle \tilde{T}_b^2 \tilde{T}_b \rangle [\text{mK}^3 \text{Mpc}^3] \quad (8.56)$$

Finally, the conversion factor can be written as:

$$C_{\text{skew}} = (1+z)^2 \frac{D_M^2(z) D_H}{\nu_{21} E(z)} \frac{\lambda_o^6}{8k_b^3} \frac{\Delta\nu}{\bar{\Omega}} \times 10^9 [\text{mK}^3 \text{Jy}^{-3} \text{Mpc}^3] \quad (8.57)$$

CHAPTER 9

SUMMARY AND CONCLUSION

The Epoch of Reionisation (EoR) is a key period in cosmological history, and marks the last great phase transition of the Universe. During this period, the Universe transitioned from a relatively uniform state, to the highly structured formation of galaxies and clusters we see today. The 21cm neutral hydrogen signal from the cosmological intergalactic medium encodes information about the sources of reionisation, as well as the structure and distribution of matter in the early Universe. Probing the statistics of the cosmological 21cm signal is the focus of current and future EoR experiments. Measurements of the 21cm power spectrum will probe the spatially variant amplitude of the signal, promising to reveal a trove of astrophysical and cosmological information. The 21cm signal is expected to be highly non-Gaussian during the EoR; information not accessible by the power spectrum. Further measurements of higher order spectra such as the bispectrum promise to unlock the non-Gaussian nature of the 21cm signal, providing a complimentary picture of Universe during the EoR.

Measurements of the 21cm power spectrum require high levels of precision, and as such face significant systematic challenges. As a result, to date only upper limits on the 21cm signal have been established, and efforts are ongoing to improve the systematics. The primary challenge comes from astrophysical foregrounds, which are up to 5 orders

of magnitude brighter than the expected 21cm signal during reionisation. Subtraction and avoidance techniques mitigate spectral leakage from the spectrally smooth foregrounds. However, the chromaticity of the instrument couples to unsubtracted sources, resulting in spectral leakage across the Fourier modes, contaminating the power spectrum. Bright extended sources such as Galactic Plane supernova remnants, typically located in the primary beam sidelobes, are particularly sensitive to chromatic effects. Due to the difficulty of modelling these sources, and the high attenuation of the primary beam sidelobes, they have been largely ignored. Higher order spectra, in particular the bispectrum, will likewise be foreground-limited. Additionally, higher order statistics have lower relative signal to noise compared to the power spectrum. This further compounds the difficulty of measuring higher spectra, making measurements of the 21cm bispectrum inherently more challenging than the power spectrum. Nevertheless, despite the challenges faced by both power spectrum and higher order spectra measurements, the detection of the 21cm signal more than warrants the required scientific effort.

This thesis aims to address some of these challenges by improving our understanding of foreground contamination, and exploring other statistics where the systematics may be improved. It understands the impact of widefield extended sources to the expected 21cm power spectrum. By modelling the Galactic Plane supernova remnants and Centaurus A, we found that these sources contributed significant contamination to the power spectrum. Additionally, this thesis simplifies the estimation of spatially variant non-Gaussian 21cm signal, by integrating the bispectrum into a pseudo-power spectrum, known as the skew spectrum. Through cosmological simulations, it analysed the sensitivity of the 21cm skew spectrum to changes in the ionisation, and X-ray heating topologies during the Epoch of Heating (EoH), and the EoR. We find that the skew spectrum is sensitive to these changes, but is contaminated with the Gaussian (power spectrum) statistical component. To isolate the non-Gaussianities in the 21cm signal, we normalise the skew spectrum by the power spectrum, and find the normalised skew spectrum is sensitive to the ionisation topology during the EoR, with a trough correlated

to the characteristic ionisation scale. Building on these efforts, we develop an analytical framework for estimating the radio interferometric skew spectrum, and applied this to simulated Phase I MWA data. In particular, we analyse the non-Gaussianity for a fiducial 21cm signal, and for realistic foreground models which include the widefield extended sources. We successfully recreate the expected foreground and 21cm skew spectrum using the analytical framework, finding the skew spectrum, like the power spectrum, is susceptible to the same chromatic instrumental effects such as mode mixing.

This thesis has established the skew spectrum as an impactful probe of the non-Gaussian nature of the cosmological 21cm signal. Furthermore, it has demonstrated the importance of widefield extended sources in the MWA primary beam sidelobes, as well as characterised the non-Gaussianity of relevant foreground fields. Importantly, it demonstrates that radio interferometers, in principle, can estimate the skew spectrum. However, additional work is required before the instrumental skew spectrum can be applied to real data. The biggest challenges are understanding the noise behaviour of the instrumental skew spectrum, and removing the noise power and foregrounds from the auto-correlations (which are required to calculate the skew spectrum). Regardless of these challenges, the scientific value from the potential applications of an interferometric skew spectrum estimator are apparent. Furthermore, the analytical framework laid out in this thesis, can be extended to other higher order spectra such as the trispectrum, therefore multiplying the scientific benefits. This thesis therefore lays the foundation for future high order pseudo-power spectra work with radio interferometers.

COPYRIGHT INFORMATION

We include the copyright permission from the Oxford journals below. Of particular note is the section titled '**After Publication**', which states that published journal articles can be included in a thesis for reuse without having to obtain formal permission. A formal request for permission is currently pending.

Author reuse and self-archiving

[Scope of the policy](#)

[Rights granted](#)

[Uncovered uses](#)

[Sharing](#)

[How to request permissions through PLS Clear](#)

Oxford University Press (OUP) makes a substantial investment in the editorial review, copyediting, typesetting, design, printing, coding for electronic publication, marketing, and distribution of the works we publish. Reuse permissions enable that investment and protect against piracy and plagiarism.

As an author of a title published by OUP Academic, Trade, Reference, Science and/or Medical books groups, you have certain rights and obligations with regard to reuse and self-archiving. This policy sets out the way in which you may reuse pre- and post-publication versions of your work for your own teaching, sharing, and self-archiving purposes, without first obtaining written permission from OUP.

Scope of the policy

The specific policy parameters set out here relate *ONLY* to titles published by the following Academic book publishing groups in the UK, US, India, and Canada:

- Academic
- Trade
- Science
- Medical
- Reference

These policies *do not* cover publications from the following publishing groups, which are handled on a case-by-case basis by each book's editor:

- Higher Education
- Law
- Education - Schools/K-12 (Canada)
- ELT (Canada)

Authors wishing to reuse work published with [English Language Teaching \(ELT\)](#) or [Education](#) divisions must contact the respective rights teams directly to submit permissions requests.

OUP's journals have unique author reuse guidelines which may be found on the homepage of each journal.

If you are unsure about which group is responsible for your publication, please [ask your editorial contact](#).

Rights Granted

Prior to publication

Prior to publication, you retain the right to make one chapter (or article, as appropriate) of the original pre-copyedited version of your work available on any or all of the following:

- your personal website
- your employer's website
- on pre-print servers (i.e., free public servers of original version articles or other content in your subject area)

These uses are permitted if you add an acknowledgement that the content has been accepted for publication, such as:

This is a draft of a chapter/article that has been accepted for publication by Oxford University Press in the forthcoming book [title] by/edited by (Author/editor) due for publication in [year].

After your work publishes, please make sure to update this record with accurate details and add links to the OUP catalogue and webpage.

After publication

After publication, you may reuse portions of your content in *only* the following ways without obtaining formal permission:

- one chapter, up to 10% of the total content, and/or three figures/illustrations /tables from a single book with one or more authors
- a maximum of one chapter/article from your contribution to an edited volume or collection (e.g., Oxford Handbooks)
- a maximum of one chapter/article of your contribution to an online-only or digital-original publication

Additionally, OUP is pleased to grant permission for the following uses:

- posting on your personal website or in an institutional or subject-based repository after a
 - 12-month period for *Science and Medical* titles
 - 24-month period for *Academic, Trade, and Reference* titles
- inclusion in scholarly, not-for-profit derivative reuses (e.g., the extension of your contribution to a book-length work, or inclusion in an edited collection of your own work, or any work of which you are an author or editor)
- reproduction within course-packs or e-course-packs for your own teaching purposes, but only if the cost to students does not exceed the cost of reproduction
- inclusion within your thesis or dissertation
- Permission for these reuses is granted on the following conditions:
 - the material you wish to reuse is your own work and has already been published by OUP
 - the intended reuse is for scholarly (non-commercial) purposes, for publication by a not-for-profit publisher
 - full acknowledgement is made of the original publication stating the specific material reused [pages, figure numbers, etc.], [Title] by/edited by [Author/editor], [year of publication], reproduced by permission of Oxford University Press [link to OUP catalogue if available, or OUP website] request formal permission when you are reusing a passage from your previously OUP-published work via plsclear.com (see '[How to request permission through PLS Clear](#)'). Permission is typically granted without charge. You will need to do this even if you retain the copyright in your contract, as your contract grants OUP the exclusive rights to publish your work include the following credit when you are use less than 10% of your previous OUP-published work in a new publication: "This material was originally published in [Title] by / edited by [Author / Editor] and has been reproduced by permission of Oxford University Press [link to book within an OUP online product and/or <http://global.oup.com/academic>]. For permission to reuse this material, please visit <http://global.oup.com/academic/rights>"
- let your OUP editor know if you plan to use a portion of your OUP publication in your thesis or dissertation
- you have obtained permission from your co-authors for reuse/republication of your joint-authored work
- any reuse on personal websites and institutional or subject based repositories includes a link to the published work online (e.g. on Oxford Academic) and/or to the OUP online catalogue entry
- the material is not distributed under any kind of Open Access license (e.g., Creative Commons) that would counter the terms of your license with OUP

Uncovered Uses

CO-AUTHOR PERMISSIONS

Here we include the co-author responses to the originality statements in Chapters 6, 7 and 8. Please refer to these chapters for the associated originality statements. Below are the signed responses of the co-authors.

To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.



(Name of co-author)



(Signature of co-author)

To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.




(Name of co-author)



(Signature of co-author)

To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.




(Name of co-author)



(Signature of co-author)

To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.



(Name of co-author)





To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.



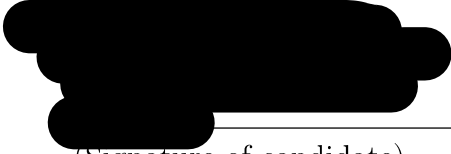
(Name of co-author)



(Signature of co-author)

To whom it may concern,

I, Jaiden Cook, have outline my contributions and the contributions of relevant co-authors to the chapters in this thesis that have been adapted from published or submitted papers Chapter 6: “Investigating the Contribution of Extended Radio Sources to the Epoch of Reionisation Power Spectrum”, and Chapter 7: “The Impact of ionisation Morphology and X-ray Heating on the Cosmological 21cm Skew Spectrum”.



(Signature of candidate)

I, as a co-author, endorse that the level of contribution indicated by the candidate as in the “Statement of contribution by Co-authors” is appropriate.



(Name of co-author)



(Signature of co-author)

BIBLIOGRAPHY

Abdurashidova, Z., Aguirre, J. E., Alexander, P., Ali, Z. S., Balfour, Y., Beardsley, A. P., and Bernardi, G. (2022). First results from hera phase i: Upper limits on the epoch of reionization 21 cm power spectrum. *The Astrophysical Journal*, 925(2):221.

Alvarez, H., Aparici, J., May, J., and Reich, P. (2000). The radio continuum spectrum of Centaurus A's large-scale components. *A* A, 355:863–872.

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., Greenfield, P., Droettboom, M., Bray, E., and Aldcroft, T. a. (2013). Astropy: A community Python package for astronomy. *A* A, 558:A33.

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., Lim, P. L., Crawford, S. M., and Conseil, S. a. (2018). The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. *AJ*, 156(3):123.

Balu, S., Greig, B., Qiu, Y., Power, C., Qin, Y., Mutch, S., and Wyithe, J. S. B. (2023). Thermal and reionization history within a large-volume semi-analytic galaxy formation simulation. *MNRAS*, 520(3):3368–3382.

- Barkana, R. and Loeb, A. (2001). In the beginning: the first sources of light and the reionization of the universe. *Physics Reports*, 349(2):125–238.
- Barkana, R. and Loeb, A. (2005a). Detecting the earliest galaxies through two new sources of 21 centimeter fluctuations. *The Astrophysical Journal*, 626(1):1.
- Barkana, R. and Loeb, A. (2005b). A method for separating the physics from the astrophysics of high-redshift 21 centimeter fluctuations. *The Astrophysical Journal*, 624(2):L65.
- Barkana, R. and Loeb, A. (2007). The physics and early history of the intergalactic medium. *Reports on Progress in Physics*, 70(4):627.
- Barry, N., Hazelton, B., Sullivan, I., Morales, M. F., and Pober, J. C. (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(3):3135–3144.
- Barry, N., Wilensky, M., Trott, C. M., Pindor, B., Beardsley, A. P., Hazelton, B. J., and Sullivan, I. S. (2019). Improving the epoch of reionization power spectrum results from Murchison widefield array season 1 observations. *The Astrophysical Journal*, 884(1):1.
- Berezhko, E. G. and Völk, H. J. (2004). The theory of synchrotron emission from supernova remnants. *A*, 427:525–536.
- Bharadwaj, S. and Saiyad Ali, S. (2005). On using visibility correlations to probe the HI distribution from the dark ages to the present epoch – I. Formalism and the expected signal. *MNRAS*, 356(4):1519–1528.
- Bharadwaj, S. and Pandey, S. K. (2005). Probing non-Gaussian features in the HI distribution at the epoch of re-ionization. *MNRAS*, 358(3):968–976.

- Blundell, K. M., Rawlings, S., and Willott, C. J. (1999). The Nature and Evolution of Classical Double Radio Sources from Complete Samples. *AJ*, 117(2):677–706.
- Bock, D. C.-J., Large, M. I., and Sadler, E. M. (1999). SUMSS: A Wide-Field Radio Imaging Survey of the Southern Sky. I. Science Goals, Survey Design, and Instrumentation. *AJ*, 117:1578–1593.
- Bond, J. R., Cole, S., Efstathiou, G., and Kaiser, N. (1991). Excursion Set Mass Functions for Hierarchical Gaussian Fluctuations. *ApJ*, 379:440.
- Bosman, S. E. I., Davies, F. B., Becker, G. D., Keating, L. C., Davies, R. L., Zhu, Y., and Eilers, A.-C. (2022). Hydrogen reionization ends by $z = 5.3$: Lyman- optical depth measured by the XQR-30 sample. *Monthly Notices of the Royal Astronomical Society*, 514(1):55–76.
- Bouwens, R., Illingworth, G., Oesch, P., Stefanon, M., Naidu, R., van Leeuwen, I., and Magee, D. (2023). UV luminosity density results at $z > 8$ from the first JWST/NIRCam fields: limitations of early data sets and the need for spectroscopy. *Monthly Notices of the Royal Astronomical Society*, 523(1):1009–1035.
- Bouwens, R. J., Illingworth, G. D., Oesch, P. A., Caruana, J., Holwerda, B., Smit, R., and Wilkins, S. (2015). REIONIZATION AFTER PLANCK: THE DERIVED GROWTH OF THE COSMIC IONIZING EMISSIVITY NOW MATCHES THE GROWTH OF THE GALAXY UV LUMINOSITY DENSITY. *The Astrophysical Journal*, 811(2):140.
- Bower, R. G. (1991). The evolution of groups of galaxies in the Press–Schechter formalism. *Monthly Notices of the Royal Astronomical Society*, 248(2):332–352.
- Bowman, J. D., Morales, M. F., and Hewitt, J. N. (2009). Foreground Contamination in Interferometric Measurements of the Redshifted 21 cm Power Spectrum. *ApJ*, 695(1):183–199.

- Bowman, J. D., Cairns, I., Kaplan, D. L., Murphy, T., Oberoi, D., Staveley-Smith, L., Arcus, W., et al. (2013). Science with the Murchison Widefield Array. *PASA*, 30:e031.
- Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., and Mahesh, N. (2018). An absorption profile centred at 78 megahertz in the sky-averaged spectrum. *Nature*, 555(7694):67–70.
- Bromm, V., Kudritzki, R. P., and Loeb, A. (2001). Generic spectrum and ionization efficiency of a heavy initial mass function for the first stars. *The Astrophysical Journal*, 552(2):464.
- Bunker, Andrew J., Saxena, Aayush, Cameron, Alex J., Willott, Chris J., Curtis-Lake, Emma, Jakobsen, Peter, and Carniani, Stefano (2023). Jades nirspec spectroscopy of gn-z11: Lyman- emission and possible enhanced nitrogen abundance in a $z = 10.60$ luminous galaxy. *AA*, 677:A88.
- Byrne, R., Morales, M. F., Hazelton, B., Sullivan, I., Barry, N., Lynch, C., Line, J. L. B., and Jacobs, D. C. (2021). A map of diffuse radio emission at 182MHz to enhance epoch of reionization observations in the Southern hemisphere. *Monthly Notices of the Royal Astronomical Society*, 510(2):2011–2024.
- Calabretta, M. R. and Greisen, E. W. (2002). Representations of celestial coordinates in FITS. *A*, 395:1077–1122.
- Chapman, E., Abdalla, F. B., Harker, G., Jelić, V., Labropoulos, P., Zaroubi, S., Brentjens, M. A., de Bruyn, A. G., and Koopmans, L. V. E. (2012a). Foreground removal using fastica : a showcase of LOFAR-EoR . *Monthly Notices of the Royal Astronomical Society*, 423(3):2518–2532.
- Chapman, E., Abdalla, F. B., Bobin, J., Starck, J.-L., Harker, G., Jelić, V., and Labropoulos, P. (2012b). 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.
- Chapman, E. (2017). Foreground mitigation in the epoch of reionization. *Proceedings of the International Astronomical Union*, 12(S333):261–268.
- Chen, X. and Miralda-Escudé, J. (2004). The Spin-Kinetic Temperature Coupling and the Heating Rate due to Ly α Scattering before Reionization: Predictions for 21 Centimeter Emission and Absorption. *ApJ*, 602(1):1–11.
- Chiang, C.-T., Wagner, C., Schmidt, F., and Komatsu, E. (2014). Position-dependent power spectrum of the large-scale structure: a novel method to measure the squeezed-limit bispectrum. *Journal of Cosmology and Astroparticle Physics*, 2014(05):048.
- Chokshi, A., Line, J. L. B., Barry, N., Ung, D., Kenney, D., McPhail, A., Williams, A., and Webster, R. L. (2021). Dual polarization measurements of MWA beam patterns at 137 MHz. *MNRAS*, 502(2):1990–2004.
- Ciardi, B. and Madau, P. (2003). Probing beyond the epoch of hydrogen reionization with 21 centimeter radiation. *The Astrophysical Journal*, 596(1):1.
- Clark, B. G. (1999). Coherence in Radio Astronomy. In Taylor, G. B., Carilli, C. L., and Perley, R. A., editors, *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, page 1.
- Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., and Broderick, J. J. (1998). The NRAO VLA Sky Survey. *ApJ*, 115:1693–1716.
- Conway, R. G., Kellermann, K. I., and Long, R. J. (1963). The radio frequency spectra of discrete radio sources. *MNRAS*, 125:261.
- Cook, J. H., Seymour, N., and Sokolowski, M. (2021). A calibration and imaging strategy at 300 mhz with the murchison widefield array (mwa). *Publications of the Astronomical Society of Australia*, 38:e063.

- Cook, J. H., Trott, C. M., and Line, J. L. B. (2022). Investigating the contribution of extended radio sources to the Epoch of Reionization power spectrum. *Monthly Notices of the Royal Astronomical Society*, 514(1):790–805.
- Cooray, A. (2001). Squared temperature-temperature power spectrum as a probe of the cmb bispectrum. *Phys. Rev. D*, 64:043516.
- Cooray, A. (2005). Large-scale non-Gaussianities in the 21-cm background anisotropies from the era of reionization. *MNRAS*, 363(3):1049–1056.
- Cornwell, T. J., Golap, K., and Bhatnagar, S. (2008). The noncoplanar baselines effect in radio interferometry: The w-projection algorithm. *IEEE Journal of Selected Topics in Signal Processing*, 2(5):647–657.
- Dai, J.-P., Verde, L., and Xia, J.-Q. (2020). What can we learn by combining the skew spectrum and the power spectrum? *Journal of Cosmology and Astroparticle Physics*, 2020(08):007.
- Datta, A., Bowman, J. D., and Carilli, C. L. (2010). Bright Source Subtraction Requirements for Redshifted 21 cm Measurements. *ApJ*, 724(1):526–538.
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., Alexander, P., Ali, Z. S., Beardsley, A. P., Bernardi, G., and Bowman, J. D. (2017). Hydrogen epoch of reionization array (HERA). *Publications of the Astronomical Society of the Pacific*, 129(974):045001.
- Dewdney, P. E., Turner, W., Millenaar, R., McCool, R., Lazio, J., and Cornwell, T. J. (2013). Ska technical document.
- Donnan, C. T., McLeod, D. J., Dunlop, J. S., McLure, R. J., Carnall, A. C., Begley, R., and Cullen (2022). The evolution of the galaxy UV luminosity function at redshifts $z = 8 - 15$ from deep JWST and ground-based near-infrared imaging. *Monthly Notices of the Royal Astronomical Society*, 518(4):6011–6040.
- Dubner, G. and Giacani, E. (2015). Radio emission from supernova remnants. , 23:3.

- Eastwood, M. W., Anderson, M. M., Monroe, R. M., Hallinan, G., Barsdell, B. R., Bourke, S. A., and Clark, M. A. (2018). The radio sky at meter wavelengths: m-mode analysis imaging with the ovro-lwa. *The Astronomical Journal*, 156(1):32.
- Eggemeier, A. and Smith, R. E. (2016). Cosmology with phase statistics: parameter forecasts and detectability of BAO. *MNRAS*, 466(2):2496–2516.
- Einstein, A. (1915). Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pages 844–847.
- Fabbiano, G. (2006). Populations of x-ray sources in galaxies. *Annual Review of Astronomy and Astrophysics*, 44(1):323–366.
- Fan, X., Strauss, M. A., Becker, R. H., White, R. L., Gunn, J. E., Knapp, G. R., and Richards, G. T. (2006). Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with $z \sim 6$ Quasars. II. A Sample of 19 Quasars. *AJ*, 132(1):117–136.
- Fialkov, A., Cohen, A., Barkana, R., and Silk, J. (2016). Constraining the redshifted 21-cm signal with the unresolved soft X-ray background. *Monthly Notices of the Royal Astronomical Society*, 464(3):3498–3508.
- Field, G. B. (1958). Excitation of the Hydrogen 21-CM Line. *Proceedings of the IRE*, 46:240–250.
- Field, G. B. (1959a). The Spin Temperature of Intergalactic Neutral Hydrogen. *ApJ*, 129:536.
- Field, G. B. (1959b). The Time Relaxation of a Resonance-Line Profile. *ApJ*, 129:551.
- For, B. Q., Staveley-Smith, L., Hurley-Walker, N., Franzen, T., Kapińska, A. D., Filipović, M. D., and Collier, J. D. (2018). A multifrequency radio continuum study of the Magellanic Clouds - I. Overall structure and star formation rates. *MNRAS*, 480(2):2743–2756.

- Fragos, T., Lehmer, B., Tremmel, M., Tzanavaris, P., Basu-Zych, A., Belczynski, K., Hornschemeier, A., Jenkins, L., Kalogera, V., Ptak, A., and Zezas, A. (2013). X-ray binary evolution across cosmic time. *The Astrophysical Journal*, 764(1):41.
- Friedmann, A. (1922). Über die Krümmung des Raumes. *Zeitschrift für Physik*, 10:377–386.
- Furlanetto, S. R., Zaldarriaga, M., and Hernquist, L. (2004a). The growth of H II regions during reionization. *The Astrophysical Journal*, 613(1):1.
- Furlanetto, S. R., Zaldarriaga, M., and Hernquist, L. (2004b). Statistical probes of reionization with 21 centimeter tomography. *The Astrophysical Journal*, 613(1):16.
- Furlanetto, S. R. and Oh, S. P. (2005). Taxing the rich: recombinations and bubble growth during reionization. *MNRAS*, 363(3):1031–1048.
- Furlanetto, S. R., McQuinn, M., and Hernquist, L. (2006a). Characteristic scales during reionization. *MNRAS*, 365(1):115–126.
- Furlanetto, S. R., Peng Oh, S., and Briggs, F. H. (2006b). Cosmology at low frequencies: The 21cm transition and the high-redshift universe. *Physics Reports*, 433(4):181–301.
- Furlanetto, S. R. and Furlanetto, M. R. (2006). Spin-exchange rates in electron–hydrogen collisions. *Monthly Notices of the Royal Astronomical Society*, 374(2):547–555.
- Furlanetto, S. R. (2006). The global 21-centimeter background from high redshifts. *MNRAS*, 371(2):867–878.
- Furlanetto, S. R. and Pritchard, J. R. (2006). The scattering of Lyman-series photons in the intergalactic medium. *Monthly Notices of the Royal Astronomical Society*, 372(3):1093–1103.
- Furlanetto, S. R. (2016). *The 21-cm Line as a Probe of Reionization*, pages 247–280. Springer International Publishing, Cham.

- Furlanetto, S. R. (2019). Theoretical framework: The fundamentals of the 21 cm line. In *The Cosmic 21-cm Revolution*, 2514-3433, pages 1–1 to 1–17. IOP Publishing.
- Gilfanov, M., Grimm, H.-J., and Sunyaev, R. (2004). Lx-SFR relation in star-forming galaxies. *Monthly Notices of the Royal Astronomical Society*, 347(3):L57–L60.
- Giri, S. K., Mellema, G., Dixon, K. L., and Iliev, I. T. (2017). Bubble size statistics during reionization from 21-cm tomography. *Monthly Notices of the Royal Astronomical Society*, 473(3):2949–2964.
- Green, D. A. (2019). A revised catalogue of 294 Galactic supernova remnants. *Journal of Astrophysics and Astronomy*, 40(4):36.
- Greig, B. and Mesinger, A. (2017). Simultaneously constraining the astrophysics of reionization and the epoch of heating with 21CMMC. *Monthly Notices of the Royal Astronomical Society*, 472(3):2651–2669.
- Greig, B., Wyithe, J. S. B., Murray, S. G., Mutch, S. J., and Trott, C. M. (2022a). Generating extremely large-volume reionization simulations. *Monthly Notices of the Royal Astronomical Society*, 516(4):5588–5600.
- Greig, B., Wyithe, J. S. B., Murray, S. G., Mutch, S. J., and Trott, C. M. (2022b). Generating extremely large-volume reionization simulations. *Monthly Notices of the Royal Astronomical Society*, 516(4):5588–5600.
- Gunn, J. E. and Peterson, B. A. (1965). On the Density of Neutral Hydrogen in Intergalactic Space. *ApJ*, 142:1633–1636.
- Guth, A. H. (1981). Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356.
- Hamaker, J. P., Bregman, J. D., and Sault, R. J. (1996). Understanding radio polarimetry. I. Mathematical foundations. *A&A*, 117:137–147.
- Hamaker, J. P. and Bregman, J. D. (1996). Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters. *A&A*, 117:161–165.

- Hamaker, J. P. (2000). Understanding radio polarimetry. IV. The full-coherency analogue of scalar self-calibration: Self-alignment, dynamic range and polarimetric fidelity. *A&A*, 143:515–534.
- Hancock, P. J., Murphy, T., Gaensler, B. M., Hopkins, A., and Curran, J. R. (2012). Compact continuum source finding for next generation radio surveys. *MNRAS*, 422:1812–1824.
- Hancock, P. J., Trott, C. M., and Hurley-Walker, N. (2018). Source finding in the era of the ska (precursors): Aegean 2.0. *PASA*, 35:e011.
- Harker, G., Zaroubi, S., Bernardi, G., Brentjens, M. A., De Bruyn, A. G., Ciardi, B., Jelić, V., and Koopmans, L. V. E. (2009). Non-parametric foreground subtraction for 21-cm epoch of reionization experiments. *Monthly Notices of the Royal Astronomical Society*, 397(2):1138–1152.
- Harris, F. (1978). On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83.
- Haslam, C. G. T., Salter, C. J., Stoffel, H., and Wilson, W. E. (1982). A 408 MHz all-sky continuum survey. II. The atlas of contour maps. , 47:1–143.
- Heinrich, C. and Hu, W. (2021). Reionization effective likelihood from planck 2018 data. *Phys. Rev. D*, 104:063505.
- Hinshaw, G., Larson, D., Komatsu, E., Spergel, D. N., Bennett, C. L., Dunkley, J., Nolte, M. R., Halpern, M., Hill, R. S., Odegard, N., Page, L., Smith, K. M., Weiland, J. L., Gold, B., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., Wollack, E., and Wright, E. L. (2013). Nine-year wilkinson microwave anisotropy probe (wmap) observations: Cosmological parameter results. *The Astrophysical Journal Supplement Series*, 208(2):19.
- Högbom, J. A. (1974). Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. , 15:417.

- Hogg, D. W. (1999). Distance measures in cosmology. *arXiv e-prints*, pages astro-ph/9905116.
- Hubble, E. (1929). A Relation between Distance and Radial Velocity among Extragalactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168–173.
- Hurley-Walker, N., Callingham, J. R., Hancock, P. J., Franzen, T. M. O., Hindson, L., Kapińska, A. D., Morgan, et al. (2017). GaLactic and Extragalactic All-sky Murchison Widefield Array (GLEAM) survey - I. A low-frequency extragalactic catalogue. *MNRAS*, 464:1146–1167.
- Hurley-Walker, N., Hancock, P. J., Franzen, T. M. O., Callingham, J. R., Offringa, A. R., Hindson, L., and Wu, C. (2019). GaLactic and Extragalactic All-sky Murchison Widefield Array (GLEAM) survey II: Galactic plane $345^\circ \leq l \leq 67^\circ$, $180^\circ \leq b \leq 240^\circ$. *PASA*, 36:e047.
- Hutter, A., Watkinson, C. A., Seiler, J., Dayal, P., Sinha, M., and Croton, D. J. (2019). The 21cm bispectrum during reionization: a tracer of the ionization topology. *MNRAS*, 492(1):653–667.
- Intema, H. T., Jagannathan, P., Mooley, K. P., and Frail, D. A. (2017). The gmrt 150 mhz all-sky radio survey - first alternative data release tgss adr1. *A&A*, 598:A78.
- Jelić, V., Zaroubi, S., Labropoulos, P., Thomas, R. M., Bernardi, G., Brentjens, M. A., De Bruyn, A. G., Ciardi, B., Harker, G., Koopmans, L. V. E., Pandey, V. N., Schaye, J., and Yatawatta, S. (2008). Foreground simulations for the LOFAR–epoch of reionization experiment. *Monthly Notices of the Royal Astronomical Society*, 389(3):1319–1335.
- Jones, R. C. (1941). A new calculus for the treatment of optical systems. description and discussion of the calculus. *J. Opt. Soc. Am.*, 31(7):488–493.
- Kardashev, N. S. (1962). Nonstationarity of Spectra of Young Sources of Nonthermal Radio Emission. *Sov. Astron.*, 6:317.

- Kennicutt, R. C. (1998). The global schmidt law in star-forming galaxies. *The Astrophysical Journal*, 498(2):541.
- Khintchine, A. (1934). Korrelationstheorie der stationären stochastischen prozesse. *Mathematische Annalen*, 109:604–615.
- Koopmans, L., Pritchard, J., Mellema, G., Aguirre, J., Ahn, K., Barkana, R., van Bemmell, I., and Bernardi, G. a. (2015). The Cosmic Dawn and Epoch of Reionisation with SKA. *PoS, AASKA14:001*.
- Kriele, M. A., Wayth, R. B., Bentum, M. J., Juswardy, B., and Trott, C. M. (2022). Imaging the southern sky at 159 mhz using spherical harmonics with the engineering development array 2. *Publications of the Astronomical Society of Australia*, 39:e017.
- Lacey, C. and Cole, S. (1993). Merger rates in hierarchical models of galaxy formation. *Monthly Notices of the Royal Astronomical Society*, 262(3):627–649.
- Lane, W. M., Cotton, W. D., van Velzen, S., Clarke, T. E., Kassim, N. E., Helmboldt, J. F., Lazio, T. J. W., and Cohen, A. S. (2014). The Very Large Array Low-frequency Sky Survey Redux (VLSSr). *MNRAS*, 440:327–338.
- Lanman, A. E., Murray, S. G., and Jacobs, D. C. (2022). Validation Solutions to the Full-sky Radio Interferometry Measurement Equation for Diffuse Emission. *ApJS*, 259(1):22.
- Lemaître, G. (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.
- Lewis, A. (2011). The real shape of non-gaussianities. *Journal of Cosmology and Astroparticle Physics*, 2011(10):026.
- Lidz, A., Zahn, O., McQuinn, M., Zaldarriaga, M., Dutta, S., and Hernquist, L. (2007). Higher order contributions to the 21 cm power spectrum. *The Astrophysical Journal*, 659(2):865.

- Lin, Y., Oh, S. P., Furlanetto, S. R., and Sutter, P. M. (2016). The distribution of bubble sizes during reionization. *Monthly Notices of the Royal Astronomical Society*, 461(3):3361–3374.
- Line, J. (2017). *PUMA and MAJICK: cross-matching and imaging techniques for a detection of the epoch of reionisation*. PhD thesis, The school of Physics, The University of Melbourne.
- Line, J. L. B., Webster, R. L., Pindor, B., Mitchell, D. A., and Trott, C. M. (2017). PUMA: The Positional Update and Matching Algorithm. *PASA*, 34:e003.
- Line, J. L. B., McKinley, B., Rasti, J., Bhardwaj, M., Wayth, R. B., Webster, R. L., Ung, D., Emrich, D., Horsley, L., Beardsley, A., and et al. (2018). In situ measurement of mwa primary beam variation using orbcomm. *Publications of the Astronomical Society of Australia*, 35:e045.
- Line, J. L. B., Mitchell, D. A., Pindor, B., Riding, J. L., McKinley, B., Webster, R. L., Trott, C. M., Hurley-Walker, N., and Offringa, A. R. (2020). Modelling and peeling extended sources with shapelets: A fornax a case study. *Publications of the Astronomical Society of Australia*, 37:e027.
- Line, J. L. B. (2022). ‘woden’: A cuda-enabled package to simulate low-frequency radio interferometric data. *Journal of Open Source Software*, 7(69):3676.
- Liu, A., Tegmark, M., and Zaldarriaga, M. (2009). Will point sources spoil 21-cm tomography? *Monthly Notices of the Royal Astronomical Society*, 394(3):1575–1587.
- Loeb, A. and Barkana, R. (2001). The Reionization of the Universe by the First Stars and Quasars. *ARA*
A, 39:19–66.
- Lynch, C. R., Galvin, T. J., Line, J. L. B., Jordan, C. H., Trott, C. M., Chege, J. K., McKinley, B., et al. (2021). The mwa long baseline epoch of reionisation survey—i.

- improved source catalogue for the eor 0 field. *Publications of the Astronomical Society of Australia*, 38:e057.
- Ma, Q.-B. and Peng, L. (2023). Skew spectrum and smoothed skewness of 21-cm signals from epoch of reionization. *MNRAS*, 523(1):640–645.
- Madau, P., Meiksin, A., and Rees, M. J. (1997). 21 centimeter tomography of the intergalactic medium at high redshift. *The Astrophysical Journal*, 475(2):429.
- Majumdar, S., Pritchard, J. R., Mondal, R., Watkinson, C. A., Bharadwaj, S., and Mellema, G. (2018). Quantifying the non-Gaussianity in the EoR 21-cm signal through bispectrum. *MNRAS*, 476(3):4007–4024.
- Majumdar, S., Kamran, M., Pritchard, J. R., Mondal, R., Mazumdar, A., Bharadwaj, S., and Mellema, G. (2020). Redshifted 21-cm bispectrum – I. Impact of the redshift space distortions on the signal from the Epoch of Reionization. *MNRAS*, 499(4):5090–5106.
- Mandel, L. and Wolf, E. (1995). *Random (or stochastic) processes*, page 41–91. Cambridge University Press.
- McCready, L. L., Pawsey, J. L., and Payne-Scott, R. (1947). Solar Radiation at Radio Frequencies and Its Relation to Sunspots. *Proceedings of the Royal Society of London Series A*, 190(1022):357–375.
- McKinley, B., Briggs, F., Gaensler, B. M., Feain, I. J., Bernardi, G., Wayth, R. B., and Johnston-Hollitt, M. a. (2013). The giant lobes of Centaurus A observed at 118 MHz with the Murchison Widefield Array. *MNRAS*, 436(2):1286–1301.
- McKinley, B., Tingay, S. J., Carretti, E., Ellis, S., Bland-Hawthorn, J., Morganti, R., and Line, J. a. (2018). The jet/wind outflow in Centaurus A: a local laboratory for AGN feedback. *MNRAS*, 474(3):4056–4072.
- McKinley, B., Tingay, S. J., Gaspari, M., Kraft, R. P., Matherne, C., Offringa, A. R., McDonald, M., Calzadilla, M. S., Veilleux, S., Shabala, S. S., Gwyn, S. D. J., Bland-

- Hawthorn, J., Crnojevic, D., Gaensler, B. M., and Johnston-Hollitt, M. (2021). Multi-scale feedback and feeding in the closest radio galaxy centaurus a.
- McMaster, W. H. (1954). Polarization and the Stokes Parameters. *American Journal of Physics*, 22(6):351–362.
- McQuinn, M. (2012). Constraints on X-ray emissions from the reionization era. *Monthly Notices of the Royal Astronomical Society*, 426(2):1349–1360.
- Mertens, F. G., Ghosh, A., and Koopmans, L. V. E. (2018). Statistical 21-cm signal separation via Gaussian Process Regression analysis. *Monthly Notices of the Royal Astronomical Society*, 478(3):3640–3652.
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., Offringa, A. R., Mellema, G., Zaroubi, S., and Brentjens, M. A. (2020). Improved upper limits on the 21cm signal power spectrum of neutral hydrogen at $z \approx 9.1$ from LOFAR. *MNRAS*, 493(2):1662–1685.
- Mesinger, A., Furlanetto, S., and Cen, R. (2011). 21cmfast: a fast, seminumerical simulation of the high-redshift 21-cm signal. *Monthly Notices of the Royal Astronomical Society*, 411(2):955–972.
- Mesinger, A., McQuinn, M., and Spergel, D. N. (2012). The kinetic Sunyaev–Zel’dovich signal from inhomogeneous reionization: a parameter space study. *Monthly Notices of the Royal Astronomical Society*, 422(2):1403–1417.
- Mesinger, A., Ferrara, A., and Spiegel, D. S. (2013). Signatures of X-rays in the early Universe. *MNRAS*, 431(1):621–637.
- Mesinger, A., Greig, B., and Sobacchi, E. (2016). The Evolution Of 21cm Structure (EOS): public, large-scale simulations of Cosmic Dawn and reionization. *MNRAS*, 459(3):2342–2353.
- Michelson, A. A. (1920). On the Application of Interference Methods to Astronomical Measurements. *Proceedings of the National Academy of Science*, 6(8):474–475.

- Michelson, A. A. and Pease, F. G. (1921). Measurement of the Diameter of α Orionis with the Interferometer. *ApJ*, 53:249–259.
- Millea, Marius and Bouchet, François (2018). Cosmic microwave background constraints in light of priors over reionization histories. *AA*, 617:A96.
- Milne, D. K. (1968). Radio emission from the supernova remnant Vela-X. *Australian Journal of Physics*, 21:201.
- Mineo, S., Gilfanov, M., and Sunyaev, R. (2012). X-ray emission from star-forming galaxies - I. High-mass X-ray binaries. *MNRAS*, 419(3):2095–2115.
- Mirocha, J. (2019). Astrophysics from the 21 cm background. In *The Cosmic 21-cm Revolution*, 2514-3433, pages 2–1 to 2–39. IOP Publishing.
- Mohan, N. and Rafferty, D. (2015). PyBDSF: Python Blob Detection and Source Finder.
- Mondal, R., Bharadwaj, S., and Majumdar, S. (2015a). Statistics of the epoch of reionization 21-cm signal – I. Power spectrum error-covariance. *Monthly Notices of the Royal Astronomical Society*, 456(2):1936–1947.
- Mondal, R., Bharadwaj, S., Majumdar, S., Bera, A., and Acharyya, A. (2015b). The effect of non-Gaussianity on error predictions for the Epoch of Reionization (EoR) 21-cm power spectrum. *Monthly Notices of the Royal Astronomical Society: Letters*, 449(1):L41–L45.
- Mondal, R., Bharadwaj, S., and Majumdar, S. (2016). Statistics of the epoch of reionization (EoR) 21-cm signal – II. The evolution of the power-spectrum error-covariance. *Monthly Notices of the Royal Astronomical Society*, 464(3):2992–3004.
- Mondal, R., Mellema, G., Shaw, A. K., Kamran, M., and Majumdar, S. (2021). The Epoch of Reionization 21-cm bispectrum: the impact of light-cone effects and detectability. *MNRAS*, 508(3):3848–3859.

- Morales, M. F. and Hewitt, J. (2004). Toward epoch of reionization measurements with wide-field radio observations. *The Astrophysical Journal*, 615(1):7.
- Morales, M. F., Bowman, J. D., and Hewitt, J. N. (2006). Improving foreground subtraction in statistical observations of 21 cm emission from the epoch of reionization. *The Astrophysical Journal*, 648(2):767–773.
- Morales, M. F. and Wyithe, J. S. B. (2010). Reionization and Cosmology with 21-cm Fluctuations. *ARA* A, 48:127–171.
- Morales, M. F. and Wyithe, J. S. B. (2010). Reionization and cosmology with 21-cm fluctuations. *Annual Review of Astronomy and Astrophysics*, 48(1):127–171.
- Morales, M. F., Hazelton, B., Sullivan, I., and Beardsley, A. (2012). FOUR FUNDAMENTAL FOREGROUND POWER SPECTRUM SHAPES FOR 21 cm COSMOLOGY OBSERVATIONS. *The Astrophysical Journal*, 752(2):137.
- Morgan, J. S., Macquart, J.-P., Chhetri, R., Ekers, R. D., Tingay, S. J., and Sadler, E. M. (2019). Interplanetary scintillation with the murchison widefield array v: An all-sky survey of compact sources using a modern low-frequency radio telescope. *Publications of the Astronomical Society of Australia*, 36:e002.
- Mort, B., Dulwich, F., Razavi-Ghods, N., de Lera Acedo, E., and Grainge, K. (2016). Analysing the impact of far-out sidelobes on the imaging performance of the SKA-LOW telescope. *Monthly Notices of the Royal Astronomical Society*, 465(3):3680–3692.
- Mort, B. J., Dulwich, F., Salvini, S., Adami, K. Z., and Jones, M. E. (2010). Oskar: Simulating digital beamforming for the ska aperture array. In *2010 IEEE International Symposium on Phased Array Systems and Technology*, pages 690–694.
- Munshi, D., Melott, A. L., and Coles, P. (1998). Generalised cumulant correlators and hierarchical clustering.

- Murray, S. G., Greig, B., Mesinger, A., Muñoz, J. B., Qin, Y., Park, J., and Watkinson, C. A. (2020). 21cmfast v3: A python-integrated c code for generating 3d realizations of the cosmic 21cm signal. *Journal of Open Source Software*, 5(54):2582.
- Mutch, S. J., Geil, P. M., Poole, G. B., Angel, P. W., Duffy, A. R., Mesinger, A., and Wyithe, J. S. B. (2016). Dark-ages reionization and galaxy formation simulation – III. Modelling galaxy formation and the epoch of reionization. *Monthly Notices of the Royal Astronomical Society*, 462(1):250–276.
- Nambissan, J. T., Subrahmanyam, R., Somashekar, R., Udaya Shankar, N., Singh, S., Raghunathan, A., Girish, B. S., Srivani, K. S., and Sathyanarayana Rao, M. (2021). Saras 3 cd/eor radiometer: design and performance of the receiver. *Experimental Astronomy*, 51.
- Obreschkow, D., Power, C., Bruderer, M., and Bonvin, C. (2012). A robust measure of cosmic structure beyond the power spectrum: Cosmic filaments and the temperature of dark matter. *The Astrophysical Journal*, 762(2):115.
- Oesch, P. A., van Dokkum, P. G., Illingworth, G. D., Bouwens, R. J., Momcheva, I., Holden, B., Roberts-Borsani, G. W., Smit, R., Franx, M., Labbé, I., González, V., and Magee, D. (2015). A spectroscopic redshift measurement for a luminous lyman break galaxy at $z = 7.730$ using keck/mosfire. *The Astrophysical Journal Letters*, 804(2):L30.
- Offringa, A. R., McKinley, B., Hurley-Walker, N., Briggs, F. H., Wayth, R. B., Kaplan, D. L., Bell, M. E., et al. (2014). wsclean: an implementation of a fast, generic wide-field imager for radio astronomy. *MNRAS*, 444(1):606–619.
- Offringa, A. R. and Smirnov, O. (2017). An optimized algorithm for multiscale wide-band deconvolution of radio astronomical images. *MNRAS*, 471(1):301–316.
- Oh, S. P. (2001). Reionization by hard photons. i. x-rays from the first star clusters. *The Astrophysical Journal*, 553(2):499.

- Ord, S. M., Mitchell, D. A., Wayth, R. B., Greenhill, L. J., Bernardi, G., Gleadow, S., and Edgar, R. G. (2010). Interferometric imaging with the 32 element Murchison wide-field array. *Publications of the Astronomical Society of the Pacific*, 122(897):1353.
- Pacucci, F., Mesinger, A., Mineo, S., and Ferrara, A. (2014). The X-ray spectra of the first galaxies: 21cm signatures. *Monthly Notices of the Royal Astronomical Society*, 443(1):678–686.
- Paladini, R., Burigana, C., Davies, R. D., Maino, D., Bersanelli, M., Cappellini, B., Platania, P., and Smoot, G. (2003). A radio catalog of Galactic HII regions for applications from decimeter to millimeter wavelengths. *A*, 397:213–226.
- Papoulis, A. (1991). *Probability, Random Variables, and Stochastic Processes*. Electrical Engineering Series. McGraw-Hill.
- Park, J., Mesinger, A., Greig, B., and Gillet, N. (2019). Inferring the astrophysics of reionization and cosmic dawn from galaxy luminosity functions and the 21-cm signal. *Monthly Notices of the Royal Astronomical Society*, 484(1):933–949.
- Parsons, A. R., Backer, D. C., Foster, G. S., Wright, M. C. H., Bradley, R. F., Gugliucci, N. E., Parashare, C. R., Benoit, E. E., Aguirre, J. E., Jacobs, D. C., Carilli, C. L., Herne, D., Lynch, M. J., Manley, J. R., and Werthimer, D. J. (2010). THE PRECISION ARRAY FOR PROBING THE EPOCH OF RE-IONIZATION: EIGHT STATION RESULTS. *The Astronomical Journal*, 139(4):1468–1480.
- Peacock, J. A. (1998). *Cosmological density fields*, page 495–552. Cambridge University Press.
- Peebles, P. J. E. (1980). *The large-scale structure of the universe*.
- Penzias, A. A. and Wilson, R. W. (1965). A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421.

- Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, and Project, T. S. C. (1999). Measurements of Ω and Λ from 42 High-Redshift Supernovae. *ApJ*, 517(2):565–586.
- Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., and Banday, A. J. (2020). Planck 2018 results - vi. cosmological parameters. *A&A*, 641:A6.
- Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A. J., Barreiro, R. B., Bartolo, N., Basak, S., Battye, R., Benabed, K., Bernard, J. P., and Bersanelli (2021). Planck 2018 results. VI. Cosmological parameters (Corrigendum). *A A*, 652:C4.
- Pober, J. C., Hazelton, B. J., Beardsley, A. P., Barry, N. A., Martinot, Z. E., Sullivan, I. S., Morales, M. F., and Bell, M. E. (2016a). THE IMPORTANCE OF WIDE-FIELD FOREGROUND REMOVAL FOR 21 cm COSMOLOGY: A DEMONSTRATION WITH EARLY MWA EPOCH OF REIONIZATION OBSERVATIONS. *The Astrophysical Journal*, 819(1):8.
- Pober, J. C., Hazelton, B. J., Beardsley, A. P., Barry, N. A., Martinot, Z. E., Sullivan, I. S., Morales, M. F., et al. (2016b). THE IMPORTANCE OF WIDE-FIELD FOREGROUND REMOVAL FOR 21 cm COSMOLOGY: A DEMONSTRATION WITH EARLY MWA EPOCH OF REIONIZATION OBSERVATIONS. *The Astrophysical Journal*, 819(1):8.
- Prasad, P., Huizinga, F., Kooistra, E., van der Schuur, D., Gunst, A., Romein, J., Kuiack, M., Molenaar, G., Rowlinson, A., Swinbank, J. D., and Wijers, R. A. M. J. (2016). The aartfaac all-sky monitor: System design and implementation. *Journal of Astronomical Instrumentation*, 05(04):1641008.
- Pritchard, J. R. and Furlanetto, S. R. (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.
- Pritchard, J. R. and Furlanetto, S. R. (2007). 21-cm fluctuations from inhomogeneous X-ray heating before reionization. *Monthly Notices of the Royal Astronomical Society*, 376(4):1680–1694.
- Pritchard, J. R. and Loeb, A. (2012). 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901.
- Procopio, P., Wayth, R. B., Line, J., Trott, C. M., Intema, H. T., Mitchell, D. A., Pindor, B., Riding, J., Tingay, S. J., Bell, M. E., and 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.
- Qin, Y., Mutch, S. J., Poole, G. B., Liu, C., Angel, P. W., Duffy, A. R., Geil, P. M., Mesinger, A., and Wyithe, J. S. B. (2017). Dark-ages reionization and galaxy formation simulation - X. The small contribution of quasars to reionization. *MNRAS*, 472(2):2009–2027.
- Qin, Y., Mutch, S. J., Poole, G. B., Liu, C., Angel, P. W., Duffy, A. R., Geil, P. M., Mesinger, A., and Wyithe, J. S. B. (2017). Dark-ages reionization and galaxy formation simulation – X. The small contribution of quasars to reionization. *Monthly Notices of the Royal Astronomical Society*, 472(2):2009–2027.
- Qiu, Y., Mutch, S. J., da Cunha, E., Poole, G. B., and Wyithe, J. S. B. (2019). Dark-age reionization and galaxy formation simulation - XIX. Predictions of infrared excess and cosmic star formation rate density from UV observations. *MNRAS*, 489(1):1357–1372.
- Qiu, Y., Mutch, S. J., Elahi, P. J., Poulton, R. J. J., Power, C., and Wyithe, J. S. B. (2020). An efficient hybrid method to produce high-resolution large-volume dark matter simulations for semi-analytic models of reionization. *Monthly Notices of the Royal Astronomical Society*, 500(1):493–505.

- Refregier, A. (2003). Shapelets — I. A method for image analysis. *MNRAS*, 338(1):35–47.
- Regan, D. (2017). An inventory of bispectrum estimators for redshift space distortions. *Journal of Cosmology and Astroparticle Physics*, 2017(12):020–020.
- Remillard, R. A. and McClintock, J. E. (2006). X-ray properties of black-hole binaries. *Annual Review of Astronomy and Astrophysics*, 44(1):49–92.
- Riding, J. L., Mitchell, D. A., and Webster, R. L. (2017). Shapelets for the MWA. In Lorente, N. P. F., Shorridge, K., and Wayth, R., editors, *Astronomical Data Analysis Software and Systems XXV*, volume 512 of *Astronomical Society of the Pacific Conference Series*, page 257.
- Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., and Gilliland (1998). Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *AJ*, 116(3):1009–1038.
- Robertson, B. E., Ellis, R. S., Furlanetto, S. R., and Dunlop, J. S. (2015). Cosmic reionization and early star-forming galaxies: A joint analysis of new constraints from planck and the hubble space telescope. *The Astrophysical Journal Letters*, 802(2):L19.
- Robertson, H. P. (1935). Kinematics and World-Structure. *ApJ*, 82:284.
- Ross, H. E., Dixon, K. L., Ghara, R., Iliev, I. T., and Mellema, G. (2019). Evaluating the QSO contribution to the 21-cm signal from the Cosmic Dawn. *Monthly Notices of the Royal Astronomical Society*, 487(1):1101–1119.
- Rubin, V. C. and Ford, W. Kent, J. (1970). Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *ApJ*, 159:379.
- Rubin, V. C., Ford, W. K., J., and Thonnard, N. (1978). Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa -> Sc. *ApJ*, 225:L107–L111.

- Ryle, M. and Vonberg, D. D. (1946). Solar Radiation on 175 Mc./s. *Nature*, 158(4010):339–340.
- Ryle, M. (1952). A New Radio Interferometer and Its Application to the Observation of Weak Radio Stars. *Proceedings of the Royal Society of London Series A*, 211(1106):351–375.
- Sault, R. J., Hamaker, J. P., and Bregman, J. D. (1996). Understanding radio polarimetry. II. Instrumental calibration of an interferometer array. *A&A*, 117:149–159.
- Schaerer, D. (2002). On the properties of massive population iii stars and metal-free stellar populations. *A&A*, 382(1):28–42.
- Schwarz, G. (1978). Estimating the dimension of a model. *Ann. Stat.*, 6(2):461–464.
- Shaver, P. A., Windhorst, R. A., Madau, P., and de Bruyn, A. G. (1999). Can the reionization epoch be detected as a global signature in the cosmic background? *A A*, 345:380–390.
- Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., and Sitwell, M. (2014). All-sky interferometry with spherical harmonic transit telescopes. *The Astrophysical Journal*, 781(2):57.
- Shimabukuro, H., Yoshiura, S., Takahashi, K., Yokoyama, S., and Ichiki, K. (2016). 21cm line bispectrum as a method to probe cosmic dawn and epoch of reionization. *MNRAS*, 458(3):3003–3011.
- Shimabukuro, H., Yoshiura, S., Takahashi, K., Yokoyama, S., and Ichiki, K. (2017). Constraining the epoch-of-reionization model parameters with the 21-cm bispectrum. *MNRAS*, 468(2):1542–1550.
- Shimabukuro, H., Hasegawa, K., Kuchinomachi, A., Yajima, H., and Yoshiura, S. (2022). Exploring the cosmic dawn and epoch of reionization with the 21cm line. *Publications of the Astronomical Society of Japan*, 75(Supplement₁) : S1 – –S32.

- Smirnov, O. M. (2011). Revisiting the radio interferometer measurement equation. I. A full-sky Jones formalism. *A&A*, 527:A106.
- Sobacchi, E. and Mesinger, A. (2014). Inhomogeneous recombinations during cosmic reionization. *Monthly Notices of the Royal Astronomical Society*, 440(2):1662–1673.
- Sokolowski, M., Colegate, T., Sutinjo, A. T., Ung, D., Wayth, R., Hurley-Walker, N., Lenc, E., and Pindor, B. (2017). Calibration and Stokes Imaging with Full Embedded Element Primary Beam Model for the Murchison Widefield Array. *PASA*, 34:e062.
- Songaila, A. and Cowie, L. L. (2010). The evolution of Lyman limit absorption systems to redshift six*. *The Astrophysical Journal*, 721(2):1448.
- Stafford, J. N., Lopez, L. A., Auchettl, K., and Holland-Ashford, T. (2019). The age evolution of the radio morphology of supernova remnants. *The Astrophysical Journal*, 884(2):113.
- Starobinsky, A. (1980). A new type of isotropic cosmological models without singularity. *Physics Letters B*, 91(1):99–102.
- Sunyaev, R. A. and Zeldovich, Y. B. (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.
- Sutinjo, A., O’Sullivan, J., Lenc, E., Wayth, R. B., Padhi, S., Hall, P., and Tingay, S. J. (2015). Understanding instrumental Stokes leakage in Murchison Widefield Array polarimetry. *Radio Science*, 50:52–65.
- Szapudi, I. and Szalay, A. S. (1997). A new class of estimators for the n-point correlations.
- Taylor, G., Carilli, C., Perley, R., and (U.S.), N. R. A. O. (1999). *Synthesis Imaging in Radio Astronomy II: A Collection of Lectures from the Sixth NRAO/NMIMT Synthesis Imaging Summer School Held at Socorro, New Mexico, USA, 17-23 June, 1998*. Astronomical Society of the Pacific conference series. Astronomical Society of the Pacific.

- Thompson, A. R., Moran, J. M., and Swenson, George W., J. (2017). *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*.
- Thyagarajan, N., Jacobs, D., Bowman, J., Barry, N., Beardsley, A., Bernardi, G., and Briggs, F. (2015). Confirmation of wide-field signatures in redshifted 21 cm power spectra. *The Astrophysical Journal*, 807.
- Tingay, S. J., Goeke, R., Bowman, J. D., Emrich, D., Ord, S. M., Mitchell, D. A., Morales, M. F., et al. (2013). The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *PASA*, 30:e007.
- Tiwari, H., Shaw, A. K., Majumdar, S., Kamran, M., and Choudhury, M. (2022). Improving constraints on the reionization parameters using 21-cm bispectrum. *Journal of Cosmology and Astroparticle Physics*, 2022(04):045.
- Trac, H., Cen, R., and Loeb, A. (2008). Imprint of Inhomogeneous Hydrogen Reionization on the Temperature Distribution of the Intergalactic Medium. *ApJ*, 689(2):L81.
- Trott, C. M., Wayth, R. B., and Tingay, S. J. (2012). THE IMPACT OF POINT-SOURCE SUBTRACTION RESIDUALS ON 21 cm EPOCH OF REIONIZATION ESTIMATION. *The Astrophysical Journal*, 757(1):101.
- Trott, C. M., Pindor, B., Procopio, P., Wayth, R. B., Mitchell, D. A., McKinley, B., Tingay, S. J., and Barry, N. (2016). CHIPS: THE COSMOLOGICAL h i POWER SPECTRUM ESTIMATOR. *The Astrophysical Journal*, 818(2):139.
- Trott, C. M., Watkinson, C. A., Jordan, C. H., Yoshiura, S., Majumdar, S., Barry, N., Byrne, R., Hazelton, B. J., Hasegawa, K., Joseph, R., and et al. (2019). Gridded and direct epoch of reionisation bispectrum estimates using the murchison widefield array. *Publications of the Astronomical Society of Australia*, 36:e023.
- Trott, C. M., Jordan, C. H., Midgley, S., Barry, N., Greig, B., Pindor, B., and Cook, J. H. (2020). Deep multiredshift limits on Epoch of Reionization 21 cm power spectra from four seasons of Murchison Widefield Array observations. *MNRAS*, 493(4):4711–4727.

- Tumlinson, J. and Shull, J. M. (1999). Zero-metallicity stars and the effects of the first stars on reionization. *The Astrophysical Journal*, 528(2):L65.
- Turner, W. (2015). *SKA-TEL-SKO-0000008, SKA Phase 1 System Level 1 Requirements*, (SKA-TEL-SKO-0000008):40.
- van der Laan, H. (1966). A Model for Variable Extragalactic Radio Sources. *Nature*, 211:1131–1133.
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., and Yu, T. a. (2014). scikit-image: image processing in python. *PeerJ*, 2:e453.
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., Heald, G., McKean, J. P., Hessels, J. W. T., de Bruyn, A. G., and Nijboer, R. (2013). Lofar: The low-frequency array. *A&A*, 556:A2.
- Vedantham, H., Shankar, N. U., and Subrahmanyam, R. (2012). IMAGING THE EPOCH OF REIONIZATION: LIMITATIONS FROM FOREGROUND CONFUSION AND IMAGING ALGORITHMS. *The Astrophysical Journal*, 745(2):176.
- Venumadhav, T., Dai, L., Kaurov, A., and Zaldarriaga, M. (2018). Heating of the intergalactic medium by the cosmic microwave background during cosmic dawn. , 98(10):103513.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., and Peterson, P. a. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272.
- Walker, A. G. (1937). On Milne’s Theory of World-Structure. *Proceedings of the London Mathematical Society*, 42:90–127.
- Wang, B., Fujimoto, S., Labbé, I., Furtak, L. J., Miller, T. B., Setton, D. J., and Zitrin, A. (2023). Uncover: Illuminating the early universe—jwst/nirspec confirmation of z gt; 12 galaxies. *The Astrophysical Journal Letters*, 957(2):L34.

- Watkinson, C. A. and Pritchard, J. R. (2015). The impact of spin-temperature fluctuations on the 21-cm moments. *Monthly Notices of the Royal Astronomical Society*, 454(2):1416–1431.
- Watkinson, C. A., Majumdar, S., Pritchard, J. R., and Mondal, R. (2017). A fast estimator for the bispectrum and beyond – a practical method for measuring non-Gaussianity in 21-cm maps. *MNRAS*, 472(2):2436–2446.
- Watkinson, C. A., Giri, S. K., Ross, H. E., Dixon, K. L., Iliev, I. T., Mellema, G., and Pritchard, J. R. (2018). The 21-cm bispectrum as a probe of non-Gaussianities due to X-ray heating. *MNRAS*, 482(2):2653–2669.
- Watkinson, C. A., Trott, C. M., and Hothi, I. (2020). The bispectrum and 21-cm foregrounds during the Epoch of Reionization. *MNRAS*, 501(1):367–382.
- Wayth, R. B., Lenc, E., Bell, M. E., Callingham, J. R., Dwarakanath, K. S., Franzen, T. M. O., For, et al. (2015). GLEAM: The GaLactic and Extragalactic All-Sky MWA Survey. *PASA*, 32:e025.
- Wayth, R. B., Tingay, S. J., Trott, C. M., Emrich, D., Johnston-Hollitt, M., McKinley, B., Gaensler, B. M., et al. (2018). The Phase II Murchison Widefield Array: Design overview. *PASA*, 35:e033.
- Weiland, J. L., Osumi, K., Addison, G. E., Bennett, C. L., Watts, D. J., Halpern, M., and Hinshaw, G. (2018). Effect of template uncertainties on the wmap and planck measures of the optical depth due to reionization. *The Astrophysical Journal*, 863(2):161.
- Wiener, N. (1930). Generalized harmonic analysis. *Acta Mathematica*, 55(none):117 – 258.
- Wise, J. H. (2019). Cosmic reionisation. *Contemporary Physics*, 60(2):145–163.
- Wouthuysen, S. A. (1952). On the excitation mechanism of the 21-cm (radio-frequency) interstellar hydrogen emission line. *AJ*, 57:31–32.

- Wyithe, J. S. B. and Loeb, A. (2004). A characteristic size of ~ 10 Mpc for the ionized bubbles at the end of cosmic reionization. *Nature*, 432(7014):194–196.
- Wyithe, J. S. B. and Loeb, A. (2004). A characteristic size of 10Mpc for the ionized bubbles at the end of cosmic reionization. *Nature*, 432:194–196.
- Wyithe, J. S. B. and Morales, M. F. (2007). Biased reionization and non-Gaussianity in redshifted 21-cm intensity maps of the reionization epoch. *MNRAS*, 379(4):1647–1657.
- Yoshiura, S., Shimabukuro, H., Takahashi, K., Momose, R., Nakanishi, H., and Imai, H. (2015). Sensitivity for 21cm bispectrum from Epoch of Reionization. *MNRAS*, 451(1):266–274.
- Zahn, O., Lidz, A., McQuinn, M., Dutta, S., Hernquist, L., Zaldarriaga, M., and Furlanetto, S. R. (2007). Simulations and analytic calculations of bubble growth during hydrogen reionization. *The Astrophysical Journal*, 654(1):12.
- Zaldarriaga, M., Furlanetto, S. R., and Hernquist, L. (2004). 21 centimeter fluctuations from cosmic gas at high redshifts. *The Astrophysical Journal*, 608(2):622.
- Zarka, P., Girard, J. N., Tagger, M., and Denis, L. (2012). LSS/NenuFAR: The LOFAR Super Station project in Nançay. In Boissier, S., de Laverny, P., Nardetto, N., Samadi, R., Valls-Gabaud, D., and Wozniak, H., editors, *SF2A-2012: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics*.
- Zaroubi, S. (2013). *The Epoch of Reionization*, pages 45–101. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Zeldovich, Y. B. and Sunyaev, R. A. (1969). The Interaction of Matter and Radiation in a Hot-Model Universe. *ApSS*, 4(3):301–316.
- Zygelman, B. (2005). Hyperfine Level-changing Collisions of Hydrogen Atoms and Tomography of the Dark Age Universe. *ApJ*, 622(2):1356–1362.

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.