

Chasing Ancient Demons: Tools for measuring 21 cm Fluctuations before Reionization

by

Aaron Ewall-Wice

B.A., University of Chicago (2011)


Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

© Massachusetts Institute of Technology 2017. All rights reserved.



ARCHIVES

Author . **Signature redacted**
Department of Physics
May 23rd, 2017

Certified by..... **Signature redacted**
Jacqueline Hewitt
Professor of Physics
Thesis Supervisor

Accepted by ... **Signature redacted**
Nergis Mavalvala
Professor of Physics
Associate Department Head for Education



77 Massachusetts Avenue
Cambridge, MA 02139
<http://libraries.mit.edu/ask>

DISCLAIMER NOTICE

Due to the condition of the original material, there are unavoidable flaws in this reproduction. We have made every effort possible to provide you with the best copy available.

Thank you.

***Some pages in the original document contain text that runs off the edge of the page.**

(See pages 203 & 303)

***Best image quality available.**

Chasing Ancient Demons: Tools for measuring 21 cm Fluctuations before Reionization

by

Aaron Ewall-Wice

Submitted to the Department of Physics
on May 23rd, 2017, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis, we take the first steps towards measuring the fluctuations in HI emission before reionization which carry information on the first X-ray emitting compact objects and hot interstellar gas heated by the deaths of the first stars (ancient demons). First, we show that existing and planned interferometers are sensitive enough to place interesting constraints on the astrophysics of X-ray heating. Second, we obtain first upper limits on the pre-reionization fluctuations with the Murchison Widefield Array. We also use these measurements to explore the impact of low-frequency systematics, such as increased foreground brightness and the ionosphere.

We discover that contamination by fine-scale frequency structure introduced by the instrument is the leading obstacle to measuring the 21 cm power spectrum before reionization. This motivates the design of a next-generation experiment, HERA, with acceptable levels of intrinsic spectral structure. We also perform a careful examination of whether traditional calibration strategies are sufficient to suppress instrumental spectral structure. We find that while existing calibration techniques have critical flaws, there exist promising strategies to overcome these deficiencies which we are now pursuing.

Thesis Supervisor: Jacqueline Hewitt

Title: Professor of Physics

I became aware of the old island here that flowered once for Dutch sailors' eyes – a fresh, green breast of the new world. Its vanished trees, the trees that had made way for Gatsby's house, had once pandered in whispers to the last and greatest of all human dreams; for a transitory enchanted moment man must have held his breath in the presence of this continent, compelled into an aesthetic contemplation he neither understood nor desired, face to face for the last time in history with something commensurate to his capacity for wonder.

F. SCOTT FITZGERALD

The Great Gatsby, 1925

Contents

Acknowledgments	49
1 Introduction	51
1.1 The History of the Universe as Revealed through astronomical Observations	51
1.1.1 The Early Universe and the Cosmic Microwave Background	52
1.1.1.1 The Bulk Properties of the Universe	53
1.1.1.2 Grand Unification Epoch and Inflation	54
1.1.1.3 Baryogenesis, Nucleosynthesis, and Recombination	56
1.1.2 The Modern Universe	58
1.1.3 The Dark Ages and the Cosmic Dawn	60
1.1.3.1 The Dark Ages and the First Small Structures	61
1.1.3.2 The First Star, Galaxies, and Supermassive Black Holes	62
1.1.3.3 Reheating and Reionization of the Intergalactic Medium	64
1.1.3.4 A Summary of Some Unanswered Questions.	65
1.1.4 Expected Progress in the Coming Decade	67
1.1.4.1 IR Galaxy Surveys	67
1.1.4.2 The Cosmic Microwave Background	69
1.1.4.3 Optical Surveys	70
1.1.4.4 Gravitational Wave Astronomy	70
1.2 The 21 cm Line as a Cosmological Probe.	71
1.2.1 Astrophysics from the 21 cm Transition	72
1.2.2 Observables	73
1.3 Observing 21 cm from the Cosmic Dawn	77

1.3.1	Current Experiments	77
1.3.2	Foregrounds and Array Design	79
1.3.3	The state of 21 cm power spectrum measurements at the time of this Thesis.	85
1.4	This Thesis	86
1.4.0.1	The Sensitivity of Interferometers to the Power Spec- trum and astrophysics before Reionization.	87
2	Detecting the peaks of the cosmological 21cm signal	91
2.1	Introduction	91
2.2	Cosmological signal	94
2.2.1	Warm Dark Matter Models	98
2.3	Instrument sensitivity	100
2.3.1	Calculation of Thermal Noise	100
2.3.2	Array Models	103
2.4	Results	106
2.4.1	Physical insight into the signal	106
2.4.2	Detectability of the peak power	110
2.4.3	General detectability of the cosmic 21cm signal	113
2.4.4	Reionization or X-ray heating?	115
2.5	Considering more pessimistic foregrounds	117
2.6	Conclusions	118
3	Detecting the 21 cm Forest in the 21 cm Power Spectrum	121
3.1	INTRODUCTION	121
3.2	THEORETICAL BACKGROUND	125
3.2.1	Notation	125
3.2.2	The Forest's Modification of the Brightness Temperature	126
3.2.3	The Morphology of the Forest Power Spectrum	128
3.3	SIMULATIONS	130
3.3.1	Simulations of the Optical Depth of the IGM	130

3.3.2	The Model of the Radio Loud Source Distribution	133
3.3.2.1	Review of Constraints and Predictions of High Redshift Radio Counts	133
3.3.2.2	Our Choice of Population Model	134
3.3.3	Adding Sources to the Simulation	135
3.4	SIMULATION RESULTS	136
3.4.1	Computing Power Spectra	136
3.4.2	Simulation Output and the Location of the Forest in k -space	138
3.4.3	The Morphology of the Simulation results.	140
3.5	PROSPECTS FOR DETECTION WITH AN MWA-LIKE ARRAY	144
3.5.1	Power Spectrum Estimation Methods	144
3.5.2	Noise and Foreground Models	147
3.5.3	Detectability Results	149
3.5.4	Distinguishability Results	152
3.6	The Detectability of the Forest over a Broad Parameter Space	157
3.7	Conclusions and Future Outlook	159
3.A	Appendix: A Derivation of the morphology of P_f	161
3.A.1	Appendix: The Suppression of the Cross terms from LoS Cross Power Spectra	162
3.A.2	Appendix: Supression of the Cross Terms from Summing the Random Source Phases	165
3.B	Appendix: A Comparison between two source models	168
4	Constraining High Redshift X-ray Sources	171
4.1	Introduction	171
4.2	From Measurements to Constraints	174
4.3	Simulations of the 21 cm Signal	176
4.3.1	Semi-Numerical Simulations	177
4.3.2	Astrophysical Parameters and their Fiducial Values	178
4.4	Instrument and Foreground Models	183

4.5	Power Spectrum Derivatives and Their Physical Origin	188
4.5.1	How Reionization Parameters Affect the 21 cm Power Spectrum.	188
4.5.2	How X-ray Spectral Properties Affect the Power Spectrum. . .	191
4.6	Constraints from Heating Epoch Observations	193
4.6.1	Degeneracies Between Parameters	193
4.6.2	How well can Epoch of Reionization Measurements Constrain X-ray Spectral Properties?	195
4.6.3	How well do Epoch of X-ray Heating measurements improve Constraints on Reionization?	198
4.6.4	Overall Parameter Constraints	199
4.7	Conclusions	200
5	First Limits on the 21 cm EoX Power Spectrum	205
5.1	Introduction	206
5.2	Observing and Initial Data Reduction	210
5.2.1	Observations with the MWA	210
5.2.2	RFI flagging and Initial Calibration	211
5.2.3	MFS Imaging and Flux Scale Corrections	216
5.2.4	Data Cubes for Power Spectrum Analysis	218
5.3	Addressing the Challenges of Low Frequency Observing	223
5.3.1	Radio-Frequency Interference	224
5.3.2	Ionospheric Contamination	226
5.3.3	Instrumental Spectral Structure	234
5.3.4	Calibration with Autocorrelations	241
5.3.5	The time dependence of residual structure.	245
5.4	Power Spectrum Results	247
5.4.1	Systematics in the 2d Power Spectrum.	248
5.4.2	Comparing Calibration Techniques	253
5.4.3	Power Spectra Comparison Between Nights of Varying Iono- spheric Activity	257

5.4.4	First Upper Limits on the 21 cm Power Spectrum During the Pre-Reionization Epoch	258
5.4.5	The outlook for EoR Measurements on the MWA.	260
5.5	Conclusions and Future Experimental Considerations	262
5.A	Appendix: The Effect of Cable Reflections on Tile Gains	265
5.B	Appendix: The Power Spectrum of ionospheric phase fluctuations from measurements of Differential Refraction.	267
5.C	Appendix: The Amplitude of Scintillation Noise in MWA Visibilities .	268
6	he HERA Dish II: Electromagnetic Simulations and Science Implications	271
6.1	Introduction	272
6.2	The Impact of Reflections on Delay-Transform Power Spectra	276
6.3	Electromagnetic Simulations of the HERA dish element	280
6.3.1	The Simulations	281
6.3.2	Deconvolving the Response Function	283
6.4	Simulation Results	284
6.4.1	The Time Domain Response of the HERA Dish	285
6.4.2	The Delay Response of Subbands	285
6.4.3	The Origin of the Knee.	287
6.4.4	Tradeoffs in Beam Properties.	290
6.4.5	Verifying Our Framework with S_{11} Measurements and Simulations	293
6.5	The Effect of the Spectral Structure due to Reflections and Resonances in the HERA dish on Foreground Leakage and Sensitivity	296
6.5.1	Extrapolating the Bandpass and Power Kernel	296
6.5.2	The Impact of delay response of the HERA Antenna on Foreground Contamination	298
6.5.3	The Implications of Reflections and Resonances on EoR Science	302
6.6	Conclusions	310
7	Calibration Modeling Errors in 21 cm Power Spectra	313

7.1	Introduction	313
7.2	Formalism	318
7.2.1	The Statistics of Unmodeled Source Visibilities	319
7.2.1.1	Unmodeled Point Sources	319
7.2.1.2	Diffuse Galactic Emission	320
7.2.2	Frequency Domain Calibration Errors	321
7.2.3	The Impact of Gain Errors on the 21 cm Power Spectrum	325
7.2.4	Calibration Bias for a Simplified Model	330
7.3	Modeling Noise in Existing Arrays	333
7.3.1	Instrumental Models	333
7.3.2	Modeling Results	335
7.3.3	The Impact of Primary Beam Modeling Errors	338
7.3.4	The Dependence of Modeling Noise on Array and Catalog Properties	341
7.3.4.1	Catalog Depth	341
7.3.4.2	Time Averaging	342
7.3.4.3	Array Configuration	343
7.4	Eliminating Modeling Noise with Baseline Weighting	345
7.4.1	Gaussian Weighting for Sky-Based Calibration	347
7.4.2	The Impact of Inverse Baseline Weighting on Power Spectrum Sensitivity	349
7.5	Conclusions	355
7.A	Appendix: The Impact of Redundancy	357
7.B	Appendix: The Point Source Approximation For Modeled Foregrounds.	360
7.C	Appendix: Expressions for Second Moments of Delay-Transformed Calibration Errors.	363
7.C.1	Derivation of Equation 7.29	365
7.C.2	Derivation of Equation 7.30	366
7.C.3	Derivation of Equation 7.31	366
7.C.4	Derivation of Equation 7.32	366

7.C.5	Derivation of Equation 7.33	367
7.D	Appendix: Components of Λ and Ψ for Non-Redundant, Uniformly Weighted Calibration Solutions	367
7.D.1	Equation 7.36	367
7.D.2	Equation 7.37	369
7.E	Appendix: A Simplified Expression for Minimal Accessible Line-of- Sight Modes.	370
8	Conclusion	373
	Bibliography	378

List of Figures

- 1-1 The history of our Universe as revealed by cosmological observations. To date, our direct observations are limited to the “modern universe” (roughly 1 billion years after the big bang) and a thin slice of the early Universe (the CMB). The critical chapters in structure formation, in which the first baryonic objects collapsed and the first stars ignited (the “Dark Ages” and the “Cosmic Dawn”) remains mostly beyond the reach of today’s instrumentation. 52
- 1-2 Left: The number density and apparent IR magnitudes at galaxies during and before reionization along with the coverage of planned surveys by JWST and WFIRST from (Mason et al., 2015). Without the assistance of strong gravitational lensing, it will be difficult for IR surveys to look beyond reionization. Right: The UV luminosity function predicted by the DRAGONS simulations at several different redshifts for (Liu et al., 2016). The turnover at $M_{AB} = -12$ corresponds to the luminosity of atomic cooling halos and is significantly below the detection threshold of near-term IR surveys. 69

- 1-3 Top: A two-dimensional slice of the brightness temperature field from the Evolution of 21 cm Structure (EoS) 2016 data release (Mesinger et al., 2016). Middle: the corresponding global signal being targeted by single dipole experiments. Bottom: The power spectrum of brightness temperature fluctuations at the $k = 0.1 h\text{Mpc}^{-1}$ scale. The global signal exhibits a $z \approx 100$ dip which occurs when the rate of collisions between CMB photons and HI atoms drops due to the Universe’s expansion and rises again as collisional coupling to the HI kinetic temperature becomes similarly inefficient. Ly- α photons from the first stars introduce a dramatic falloff at $z \approx 30$, strongly coupling T_s to T_k . $T_s = T_k$ rises dramatically between $10 \lesssim z \lesssim 20$ as X-rays heat the IGM. As this heating proceeds, high-contrast exists between hot and cold IGM patches introduces significant fluctuations and a maximum in the power spectrum at $z \approx 15$. Fluctuations from the ionization field take over at $z \lesssim 10$ before reionization eliminates the neutral hydrogen in the IGM, bringing both the power spectrum and the global signal to ~ 0 75
- 1-4 Several examples of existing and upcoming interferometry experiments pursuing a detection of cosmological 21 cm fluctuations. From left to right and top to bottom: LOFAR, the GMRT, the MWA, HERA, PAPER, the LWA, CHIME, HIRAX, and the SKA-1 LOW. Several of the arrays (“imaging arrays”) have their antennas distributed in a random fashion while others arrange their antenna elements in regular patterns (“redundant arrays”). 80
- 1-5 While the foregrounds in 21 cm experiments are a factor of $\sim 10^4$ times brighter, they are expected to evolve smoothly in frequency and are thus characterizable by a power law. Cosmological 21 cm emission is expected to have complex frequency structure. Thus the foregrounds and signal can be isolated by a Fourier transform in the frequency direction. 81

1-6	A cartoon representation of the “EoR window” from Dillon et al. (2015b). In the absence of instrumental effects, foregrounds would possess smooth frequency evolution and the EoR window would extend to the same small k_{\parallel} for all k_{\perp} . The Fourier exponent in equation 1.8 gives rise to the “wedge”, denoted by the dark red and orange features. Chromaticity in the primary beam, arising from equation 1.9 adds additional contamination which is usually accounted for by a buffer denoted by the yellow region.	83
1-7	Left: Greig and Mesinger (2017)’s estimate for the $1 - \sigma$ confidence region for the evolution of the neutral fraction of the intergalactic medium as a function of redshift. Right: Greig et al. (2016)’s calculation of existing limits on the spin temperature and neutral fraction the intergalactic medium at redshift 8.4 (Greig et al., 2016) including (top-left) recent 21 cm upper limits by (Ali et al., 2015), 21 cm and McGreer et al. (2015) dark-pixel statistics from high redshift quasars (top right), 21 cm and CMB observations (bottom left)	86
1-8	All of the upper limits on the 21 cm power spectrum at the time of this thesis between $z \approx 6$ and $z \approx 12$. (Jacobs, Private Communication). The sensitivity of the next generation HERA instrument is shown as a yellow line.	87
2-1	Poisson (cosmic variance) component of the S/N, i.e. $\sqrt{N_k}$, for our fiducial observational strategy and $k \approx 0.1 \text{ Mpc}^{-1}$	100
2-2	Time spent in each uv cell for a 2000h observation at $z = 10$ for PAPER, MWA128T, LOFAR, HERA, SKA, (<i>left to right</i>).	103

2-3 *Top panels:* Amplitude of the 21cm power at $k = 0.1 \text{ Mpc}^{-1}$ in various models. We also plot the (1σ) sensitivity curves corresponding to a 2000h observation with: MWA128T, LOFAR, SKA on the left, and MWA256T, PAPER, and the proposed HERA instrument on the right. The recent upper limit from Parsons et al. (2014) is shown at $z = 7.7$. *Bottom panels:* The corresponding average 21cm brightness temperature offset from the CMB. 105

2-4 Various quantities evaluated at the redshift where the $k = 0.1 \text{ Mpc}^{-1}$ power is the largest in each model. Shown are power amplitude, redshift, mean neutral fraction, and mean brightness temperature, clockwise from top left. Overlaid are the Thompson scattering optical depth, τ_e , contours corresponding to 1 and 2σ constraints from the 9yr release of WMAP data (WMAP9; Hinshaw et al. 2013). The right-side y-axis shows the corresponding values of the WDM particle mass, m_{wdm} , computed according to the approximation in eq. (2.5). Regions of the parameter space in which the peak power occurs during the reionization or X-ray heating epochs are demarcated with the appropriate labels in some panels. 107

2-5 Cumulative distributions of T_γ/T_S (see eq. 2.1), at $z = 14, 16, 18$ (spanning the X-ray peak) for a fiducial, $T_{\text{vir}} = 10^4 \text{ K}$, $f_X = 1$ model. As seen above, the power peaks when the distribution of T_γ/T_S is the broadest. 110

2-6	S/N of the detection of the $k = 0.1 \text{ Mpc}^{-1}$ 21cm peak power (i.e. computed at the redshift of the maximum signal). All maps show the same range in S/N to highlight differences between instruments, and are computed assuming a total integration time of 2000h. Due to our fiducial observing strategy which minimizes thermal noise at the expense of cosmic variance, some high S/N regions ($S/N \gtrsim 50$) are limited by cosmic (Poisson) variance for the case of LOFAR and SKA, which have smaller beams than MWA and PAPER (see Fig. 2-1). This Poisson variance limit can be avoided with a different observing strategy; hence we caution the reader not to lend weight to the apparent better performance of MWA and PAPER in the upper left region of parameter space.	111
2-7	Maximum S/N possible with current and upcoming interferometers after 2000h (considering all redshifts instead of just the peak signal as in Fig. 2-6).	113
2-8	The S/N vs redshift evolution for the “fiducial” model, with $M_{\min} = 10^8 M_{\odot}$, $f_X = 1$	115
2-9	The 21cm power spectrum amplitude at $k = 0.1 \text{ Mpc}^{-1}$ (<i>left</i>), $\delta\bar{T}_b$ (<i>center</i>), and \bar{x}_{HI} (<i>right</i>), all evaluated at the redshift when S/N is the largest, assuming MWA128 sensitivities.	116
2-10	The maximum S/N for MWA-128T at $k = 0.2\text{Mpc}^{-1}$	118
3-1	The mean thermal evolution of our IGM simulations for our three models. “cool IGM”- solid lines, “fiducial IGM”- dashed-dotted lines, and “hot IGM” - dashed lines. $\langle T_s \rangle$ is plotted in lavender. Varying f_X effectively shifts $\langle T_s \rangle$ in redshift.	132
3-2	The 21 cm forest is dominated by sources in the 1-10 mJy flux range. We plot the sum of fluxes squared, in Equation (3.16), for $S < S_{\nu}$. A detection of P_f would constrain the high redshift source counts at these flux intervals.	134

- 3-3 For every heating scenario we study, there is some redshift and region within the EoR window for which 21 cm forest dominates the power spectrum. Here we show the fractional difference between the power spectrum with, (P'_b), and without (P_b) the forest for the redshifts (top to bottom) 9.2, 11.2, 12.2, 15.4, and 17.5. The diagonal lines denote the location of the “wedge”. By $z \gtrsim 12.2$ there is a substantial region ($k_{\parallel} \gtrsim 0.5 \text{ Mpc}^{-1}$) of the Fourier volume that our simulations cover in which the forest dominates P_b by a factor of a few. 137
- 3-4 The 21 cm forest dominates the spherically averaged power spectrum for $k \gtrsim 0.5 \text{ Mpc}^{-1}$. Plotted is the spherically averaged power spectrum with (dashed lines) and without (solid lines) the presence of the 21 cm forest. In our cool model, the forest causes a significant power increase at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ at redshifts as low as $z = 11.2$. At $z = 15.4$ we see a significant feature in all thermal scenarios. Our cool IGM model experiences a reduction in the power spectrum amplitude at $z \gtrsim 17.5$ as it passes through the X-ray heating peak. 139
- 3-5 We plot the magnitude of the difference between the 21 cm power spectrum with and without the presence of the 21 cm forest including the auto-power and cross power terms of Equation (3.13). At high redshifts and low f_X , there is little k_{\perp} structure in $P'_b - P_b$, indicating that P_f is the significant contributor. At lower redshifts and higher f_X , we see significant k_{\perp} structure, indicating that in a heated IGM, $P'_b - P_b$ is dominated by $P_{f,b}$ which is somewhat spherically symmetric and negative at large k . The trough in the low redshift plots marks the region where $P_f - 2\text{Re}(P_{f,b})$ transitions from negative (for small k) to positive (for large k). 140

3-6 Our semi-analytic prediction agrees well with unclustered simulation results. The semi-analytic prediction of Equation (3.15) is plotted with dashed lines and $\Delta_f^2(k)$ computed directly from our simulation without clustering in solid lines. This demonstrates that for $k \gtrsim 10^{-1} \text{ Mpc}^{-1}$, the cross terms in Equation (3.14) may be ignored and P_f may be well approximated by the LoS power spectrum of τ_{21} multiplied by the summed squared fluxes for sources lying in and behind the data cube. 142

3-7 We see that for a fixed quasar distribution, the magnitude of P_f can be parameterized by $\langle T_s \rangle$ and that the amplitude is consistent with a simple power law. Here, we plot $P_f(k_{\parallel})$ at $k_{\parallel} = 0.5 \text{ Mpc}^{-1}$ vs. $\langle T_s \rangle$ for all considered redshifts and f_X . The black line is the power law $\langle T_s \rangle^{-2}$ as one might expect for an amplitude set by $\langle \tau_{21} \rangle^2$ (Equation (3.16)). Inasmuch of this simple trend, a modest spread in heating models gives us a decent understanding of the behavior of the amplitude for P_f . This relation holds for the quasar population considered here because the integral over the luminosity function does not change significantly over the redshifts we consider. 143

3-8 The cross power spectrum, $\text{Re}(P_{f,b})$'s, sign is determined by the anti-correlation of x_{HI} and T_s during the pre-heating epoch and by x_{HI} after heating has taken place. Here we show the sign of $\text{Re}(P_{f,b})$ for our three different heating models as a function of redshift. At pre-heating redshifts, T_s is small and x_{HI} is relatively uniform so that T_b and T_f primarily depend on T_s and anti-correlate so that $\text{Re}(P_{f,b})$ is negative. At low redshifts, T_b is independent of T_s and fluctuations are primarily sourced by x_{HI} so that T_b and T_f are correlated and $\text{Re}(P_{f,b})$ is positive. Furthermore, heating proceeds in an "inside-out" manner so that the smallest scales become correlated first. 145

3-9 Array layouts that we use to determine the detectability and distinguishability of the 21 cm forest power spectrum signature. We chose to study two moderate extensions of the MWA-128T: MWA-256T and MWA-512T. In addition we study a 4096T array that is representative of a HERA scale instrument with ~ 400 times the collecting area of the MWA. Tile locations are drawn randomly from a distribution that is constant for the inner 50 m and drops as r^{-2} for larger radii. 148

3-10 Detections (black dots) and upper limits (red triangles) of the 21 cm Power Spectrum at $z=11.2$ for all of our arrays and heating models in the presence of 21 cm forest absorption from background RL sources. The grey fill denotes the 2σ region around the measured power spectrum with no RL sources present. To determine whether we can detect the forest imprint we ask, "do the points and their error bars lie outside the gray shaded region?" MWA-256T and MWA-512T would be capable of distinguishing power spectra with or without sources in our cool IGM model, however only 4096T is consistently sensitive to the $k \gtrsim 0.5 \text{ Mpc}^{-1}$ region where the forest dominates. Only for our cool IGM model, MWA-512T would sufficient to detect this upturn as well. Hence a moderate MWA extension would likely be able to constrain some RL populations given a cooler heating scenario while a HERA scale instrument will be able to constrain the W08 RL population using the Forest power spectrum even for more emissive heating scenarios. Note that the upturn in the gray region is not from increased power at high k but larger error bars. 150

3-11 These plots are identical to Figure 3-10 except the array is fixed to be MWA-4096T, representative of a HERA generation instrument, and redshift is varied. A HERA class instrument is able to resolve the upturn at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ that distinguishes the forest, and should be able to detect the 21 cm forest feature considered in this work for a variety of heating scenarios. The thermal noise error bars are too small to resolve by eye in most of these plots. 153

3-12 The significance of distinguishability across all measured k bins (Equation (3.29)) for all arrays, redshifts, and IGM heating models for a 1000 hour observation. An extension of MWA-128T is capable of distinguishing models with and without the 21 cm forest from the W08 RL population in our cool and fiducial heating scenarios. MWA-512T and HERA scale MWA-4096T are capable of distinguishing the forest in the power spectrum in all heating models considered in this work. 155

3-13 The 21 cm Forest moderately enhances the distinguishability between thermal scenarios and MWA scale interferometers can distinguish between the power spectra for reasonable X-ray heating histories. Here we show the cumulative z -score described in Equation (3.29), except now applied to the difference between different IGM heating models, for all arrays and redshifts. At low redshift, the forest decreases the distinguishability of different X-ray heating scenarios by subtracting from the higher amplitude model. When the positive auto-term dominates at high redshift, the forest increases the contrast between given heating models. 156

3-14 The ratio of the 21 cm Forest power spectrum, $P_f(k = 0.5 \text{ Mpc}^{-1})$ to thermal noise for 1000 hours of observation on a HERA scale interferometer, extrapolated over a large range of X-ray efficiencies and flux squared densities. Vertical dashed black and white lines indicate the value of the simulation by (Wilman et al., 2008) while the horizontal black and white lines indicate the f_X efficiencies that we explicitly simulate in this paper. At the highest redshifts, $\langle T_s \rangle$ levels off and the detectability of the signal is independent of redshift. At late preionization redshifts, we see that the 21 cm Forest will only be detectable for heating efficiencies $\lesssim 1$ 158

3-15 The LoS cross power spectra between spatially separated lines of sight are on the order of $\sim 100 - 1000$ times smaller than auto power spectra. In the left figure, we plot the ratio of the real cross power spectra between lines of sight separated by 24 Mpc to auto power spectra, and on the right we show the ratio of the imaginary cross power spectrum to the auto power spectrum. In both cases, for $k_{\parallel} \gtrsim 10^{-1} \text{ Mpc}^{-1}$, the cross power spectra are on the order of 10-1000 times smaller. The real cross power spectrum becomes non negligible on scales comparable to the separation between the two lines of sight. 165

3-16 Here we see that $P[\Sigma_{cross}(u_{\perp}, \Theta, N)]$ is invariant in N , Θ , and u_{\perp} , and N with a mean of approximately zero. The lines which indicate, $P[\Sigma_{cross}(u_{\perp}, \Theta, N)]$, are estimated from 10000 draws. Since $\langle \Sigma_{cross} \rangle \approx 0$ we expect the cross terms to contribute negligibly to P_f in 3D Fourier Space. 167

3-17 Left: The percentage of the integrated luminosity function in Equation (3.16) as a function of the source fluxes at 5GHz for comparison to the catalogue of H04. We see that most contributions to the forest power spectrum come in between $S_{5\text{ GHz}} = 10\ \mu\text{Jy}$ and $S_{5\text{ GHz}} = 10\ \text{mJy}$. Right: The ratio of the number of sources with redshift greater than z between $S_{5\text{ GHz}} = 10\ \mu\text{Jy}$ and $10\ \text{mJy}$ as predicted by the W08 and H04. The W08 simulation over predicts the counts in H04 by a factor ten at $z \gtrsim 12$ and nearly 80 at $z \gtrsim 16$, emphasizing the importance of exploring this widely unconstrained parameter space in future work. 170

4-1 Top: The evolution of the density weighted average of the neutral fraction. Middle: The evolution of the kinetic temperature (T_k), T_s , T_{cmb} . Bottom: the redshift evolution of the mean 21 cm brightness temperature, $\langle\delta T_b\rangle$. All averages are taken over volume. 182

4-2 Δ_{21}^2 as a function of redshift for two different 1d k bins, $k = 0.1h\text{Mpc}^{-1}$ (black lines) and $k = 0.4h\text{Mpc}^{-1}$ (red lines). For each parameter, we show the power spectrum at $\theta = \pm 0.1$ (thick solid and thick dashed lines respectively) along with the difference (thin solid lines). Note that our parameterization defines θ as the fractional difference of each parameter from its fiducial. The first two peaks of the three-peaked structure, discussed in Pritchard and Furlanetto (2007); Santos et al. (2008); Baek et al. (2010) and Mesinger et al. (2013), is clearly visible representing the epochs of reionization and X-ray heating. With the exception of R_{mfp} , parameter changes affect a broad range of redshifts. 184

- 4-3 Top: The power spectrum of 21 cm fluctuations (solid red lines) over numerous redshifts. Filled regions denote the 1σ errors for the instruments considered in this paper with moderate foregrounds. Bottom: The derivatives of the 21 cm power spectrum with respect to the astrophysical parameters considered in this work as a function of k at various redshifts. Derivatives are substantial over all redshifts except for R_{mfp} which only affects the end of reionization. Notably, $\partial_{\zeta}\Delta_{21}^2$ is negative on small scales at high redshift, a signature of the beginnings of inside out reionization while X-ray spectral parameters follow very similar redshift trends, indicative of degeneracy. 189
- 4-4 Left: The logarithm of the probability distribution function (PDF) of pixels at $z = 12.4$ for our fiducial model as a function of T_s vs. x_{HI} . Even before the majority of reionization, early HII bubbles ionize the hottest points in the IGM, leading to a pileup of high T_s pixels at $x_{HI} = 0$ and reducing the contrast in δT_b between hot and cold regions. Right: PDFs of δT_b with and without x_{HI} manually set to unity everywhere. The presence of ionization during X-ray heating leads to a decrease in the large T_s wing, near 30 mK, and a spike at 0 mK, leading to a reduction in the dynamic range of the field and an overall decrease in power. 191
- 4-5 Plotting $w_i(k, z)$ versus redshift for several different co-moving scales gives us a sense of the covariances between various parameters. Here we assume 1080 hours of observation on HERA-331 and the moderate foreground model. Since the thermal noise on Δ^2 increases rapidly with k , w_i is maximized at larger spatial scales. As we might expect from our fiducial model, w_i for reionization parameters is maximized at lower redshifts while w_i for X-ray spectral parameters is significant over the heating epoch. $T_{\text{vir}}^{\text{min}}$ affects both heating and reionization and has a broad redshift distribution. Vertical lines indicate the location of each of the three power spectrum peaks at $k = 0.1h \text{ Mpc}^{-1}$ 194

- 4-6 Plotting $w_i(k, z)$ at a single cosmological Fourier mode for all of our parameters on the same panel facilitates direct comparison. Many of the parameters have similar redshift evolutions that differ by a sign, making their effects on the power spectrum degenerate. 195
- 4-7 The fractional errors on astrophysical parameters as a function of maximal observed redshift for HERA-331 with moderate foreground contamination. From low redshift measurements, ν_{\min} and f_X are constrained to within $\approx 40\%$ though the spectral index α_X remains highly uncertain. Measurements at $z \gtrsim 10$ allow for $\lesssim 10\%$ limits on X-ray spectral parameters, including α_X and a factor of two improvement in constraints on T_{vir}^{\min} and reionization. Inability to observe within the FM radio-band (pink shaded region) raises the errors on heating parameters by a factor of two. The fall of the error bars with redshift bottoms out at high z due to increasing thermal noise. 196
- 4-8 95% confidence regions for the X-ray spectral properties of early galaxies, marginalizing over $(\zeta, R_{\text{mfp}}, T_{\text{vir}}^{\min})$ from measurements on HERA-331 with our moderate foregrounds scenario. At low redshift ($z \leq 10$), hardly any signature of α_X is present, leading to large error bars on the α_X axis. Because ν_{\min} and f_X incur very similar changes on the power spectrum during the beginning of reionization (Fig. 4-5), they are highly degenerate. Observations of the heating peak break these degeneracies. 197

- 4-9 Confidence ellipses (95%) for $T_{\text{vir}}^{\text{min}}$ and reionization parameters. By comparing the ellipses resulting from fixing our heating history and only observing at low redshift and the ellipses resulting from marginalizing over all parameters but including heating epoch measurements, one can see that a significant fraction of the gains in reionization uncertainties at high redshift come from breaking degeneracies with heating parameters rather than the direct signatures of reionization. This also shows that not marginalizing over heating parameters leads to over-optimistic predictions of reionization uncertainties. 199
- 4-10 95% confidence ellipses for our moderate and optimistic foreground models assuming 1080 hours of drift-scan observations on HERA-331. Heating parameters tend to be highly degenerate with each other but independent of reionization parameters. As the majority of the information heating comes from higher redshifts where thermal noise is much higher, their uncertainty regions tend to be several times larger. 204
- 5-1 A radio map at 408 MHz (Haslam et al., 1982) sin-projected over the region of the sky observed in this paper. Cyan through magenta contours indicate the total fraction of observation time weighted by our primary beam gain for our three hours of observation at 83 MHz. Red contours indicate R.A.-decl. lines. Observation tracked the position (R.A.(J2000) = $4^{\text{h}}0^{\text{m}}0^{\text{s}}$, decl.(J2000) = $-30^{\circ}0'0''$) on a region of the sky with relatively little galactic contamination and dominated by the resolved sources Fornax A and Pictor A. The galactic anticentre and bright diffuse sources, such as the Gum Nebula, are below 1% bore-site gain. 212

- 5-2 The autocorrelation spectrum of a single tile, showing the MWA band-pass, is plotted here (solid black line) along with the frequency ranges over which our data was taken (gray striped rectangles). Observations were performed simultaneously on two non-contiguous bands located on either side of the 88-108 MHz FM band (red shaded region) to both assess conditions within the FM band and to preserve some usable bandwidth should it have proven overly contaminated by RFI. 212
- 5-3 A false-color plot of the fractional change in our calibration amplitudes over each pointing (in which beamformer delay settings are fixed). Pointing changes are marked by the solid white horizontal lines. We see that the calibration amplitudes vary within a pointing on the order of several percent with little systematic variation. There are several coarse channel scale jumps on September 5th that correspond to observations in which the number of sources identified within a snapshot image were reduced (see Fig. 5-9). We found that these jumps corresponded to excess flagging from `cotter`, indicating high levels of RFI or other bad data, and dropped them from our analysis. Vertical lines from RFI are visible at 98 MHz and 107 MHz, especially on September 6th. 215
- 5-4 A deep image of the MWA "EoR-01" field centered at (R.A.(J2000) = $4^h0^m0^s$ and decl.(J2000) = $-30^\circ0'0''$), derived by stacking restored multifrequency synthesis Stokes XX and YY images produced by `WSClean` on Band 1. The dominant source in our field is Fornax A (detailed in the inset) whose structure is well recovered in imaging. Pictor A is also present at the center of the image (at $\sim 30\%$ beam) along with the diffuse Puppis and Vela supernova remnants on the left. 218

5-5 Left: The standard deviation over uv cells as a function of frequency for an even/odd time step difference cube after three hours of integration. The blue line is derived from data, while the green line is a model with a system temperature of 470 K at 150 MHz and a spectral index of -2.6 . Spikes in the standard deviation are present at the center of each coarse channel since the center channel has one half of the data due to flagging the center channel which is contaminated by a digital artifact before averaging from 40 to 80 kHz. Right: The ratio of variance taken over frequency in each uv cell and our variance model using the same system temperature as on the left. The ratio between our model and observed variance is close to unity across the uv plane. White cells indicate uv voxels that were flagged at all frequencies due to poor sampling. All data in this figure is from Band 1. 224

5-6 The percentage of visibilities flagged for RFI by `cotter` as a function of time and frequency. White regions indicate missing data including the coarse band edges and blue-dashed vertical lines indicate the edge of the FM band. While Band 1 remains predominantly clear with a few sparse events within the lower end of the FM band, contamination is significantly greater over Band 2. Even in the FM band, RFI events are either isolated in frequency or time, allowing us to flag them. A handful of observations in Band 1 on September 5th are missing entire coarse channels which we also discard. Bar plots on the bottom and right show the averages of the RFI flagging fraction over time and frequency respectively. 225

5-7 Differential refraction derived from source pairs within 30 minute bins on September 5th, 2013 (top row) and September 6th, 2013 (bottom row). Band 1 (black points) scaled by the ratio of the band center frequencies square (solid black line) very nicely predicts the differential refraction in Band 2 (red points), indicating that the refraction we are measuring here is indeed due to ionospheric fluctuations. The magnitude of ionospheric activity differs significantly between September 5th and 6th and peaks over the last observations taken on the 6th. We also show fits to an isotropic power spectrum model of differential refraction at 83 MHz (dashed black line) and print the inferred diffractive scale. 232

5-8 The dimensionless integral, $F_n(x)$ normalized to unity at $F_n(10)$, given in equation 5.20. For small values of x , $F_n(x)$ is well approximated by a power law, but flattens out towards $x = 1$. Hence, the structure function of observed source offsets levels out at the outer energy injection scale of the turbulence. 232

5-9 The number of sources identified in 112 s multifrequency synthesis images of Band 1 as a function of time over both nights of observing. On September 5th, the source counts increase with primary beam gain, until transit (vertical gray line) before decreasing. On September 6th, when more severe ionospheric refraction was observed, the source counts remain significantly lower. Fewer sources were identified in a handful of September 5th snapshots corresponding to observations in which calibration and flagging anomalies were observed (Figs. 5-3 and 5-6). We exclude these snapshots from our analysis. 233

5-10 The regions of the EoR window contaminated by foregrounds due to uncalibrated cable reflections for several different redshifts. Dark gray regions denote contamination from first order cable reflections assuming a wedge out to the first null of the primary beam plus the $0.15 h \text{ Mpc}^{-1}$ at $z = 8.5$ buffer observed in Pober et al. (2013a). Since the buffer is associated with the intrinsic spectral structure of foregrounds, it lives in delay space. Dark gray regions denote foreground contamination within the wedge which exists even without instrumental spectral structure. At $z = 7$, a representative EoR redshift, the contaminated regions remain at relatively high k_{\parallel} and have smaller widths due to the smaller primary beam and the scaling of k_{\perp} and k_{\parallel} with z . While regions exist between the first order reflections that are somewhat wider at lower redshift, second order reflections can still potentially contaminate nearly all of the EoR window in which interferometers are supposed to be sensitive (light grey shaded regions). Second order reflections are below the sensitivity level of this study but will also pose an obstacle to longer integration unless calibrated out. 239

5-11 We show the amplitude of a calibration gain averaged over a fifteen minute pointing (black circles) along with the square root of our autocorrelations which have been scaled by a third order polynomial and a single seven meter reflection to match the calibration solution (red line). After multiplying the autocorrelations by a smooth function, they are brought into good agreement with the calibration gains. . . . 242

5-12 Left: In order to obtain reflection parameters, we divide our scaled autocorrelation (magenta circles) by a smooth function consisting of a third order polynomial and large scale reflections (green line). Right: We fit this ratio (magenta circles) to a reflection function (green line) and are left with $\sim 10\%$ residuals (grey points). 242

- 5-13 Histograms of fitted cable reflection amplitudes for Band 1 (blue) and Band 2 (green) obtained from fits to autocorrelations for three different cable lengths between MWA receivers and tiles. The reflection amplitudes range from 0.2 – 1% making them difficult to fit using the noisy self calibration solutions. Reflection amplitudes in Band 2 are systematically larger than Band 1 for all cable lengths, indicative of non-trivial frequency evolution in the reflection parameters. 244
- 5-14 Dynamic spectra of the square root of a representative tile autocorrelation. Note the different color bars for the two frequency bands since Band 1 evolves more steeply in frequency than Band 2. The autocorrelations exhibit repetitive structure in time from night to night with smooth time variations occurring as the sky rotates overhead and steep transitions occurring every ≈ 30 minutes due to changes in the analogue beamformer settings as the antennas track the sky. Limited RFI is plainly visible within the FM band, especially in Band 2, and the events are consistent with the flagging events identified by `cotter` shown in Fig. 5-6. 245
- 5-15 Left: The residuals to fitting reflection functions in our autocorrelations for all two minute time-steps in our analysis for a representative tile with a 90 meter beamformer to receiver connection (light grey points). While some scatter exists in the residuals due to fitting noise, they average to non-zero values on the order of $\sim 10^{-3}$ (black dots). These residuals are due to mismodeling the reflections and at a lower level potentially arise from digital artifacts. Right: The same as the left but for a 150 meter cable whose reflection coefficient is $\sim \times 2$ as large as the 90 meter cable, leading to larger residuals due to mismodeling. 246

5-16 The absolute value of our cylindrical power spectrum estimate from our three nights of observing on Band 2 (left) and Band 1 (right). We overplot the locations of the primary beam (dash-dotted), horizon (dashed), and horizon plus a $0.1 h \text{ Mpc}^{-1}$ buffer (solid black) wedges. We see that the foregrounds are primarily contained within the wedge and that the EoR window is, for the most part, noise-like. There is some low SNR structure below $k_{\parallel} \approx 0.5 h \text{ Mpc}^{-1}$, corresponding to k_{\parallel} modes contaminated by cable reflections. The amplitude in power rises very quickly due to an increase in thermal noise which rises very quickly at large k_{\parallel} due to a rapid falloff in uv coverage beyond $k_{\perp} \sim 0.2 h \text{ Mpc}^{-1}$. 249

5-17 $P(k)$ over Band 2 (left) and Band 1 (right) with a color scale that highlights cells with positive or negative values. We expect regions that are thermal noise dominated to contain an equal number of positive and negative estimates and regions that are dominated by foreground leakage to be entirely positive. We observe significant foreground contamination outside of the wedge up to $k_{\parallel} \approx 0.5 h \text{ Mpc}^{-1}$ in both bands. 250

5-18 The errors on \hat{p} arising from residual foregrounds and thermal noise are determined by looking at even/odd difference cubes and foreground-subtracted residual cubes using the method of Dillon et al. (2015a). We show the error bars on our cylindrical power spectrum here, seeing that errors arising from foregrounds are contained within the wedge. These foreground errors are maximized at the smallest and largest k_{\perp} arising from large power in diffuse emission and increasing thermal noise from a dropoff in baseline density respectively. 251

5-19 The foreground contamination within the wedge along with residual detections due to miscalibrated fine frequency features in the bandpass are especially clear in plots of the ratio between power and the error bars estimated by the empirical covariance method of D15. We overplot the wedge with a $0.1 h \text{ Mpc}^{-1}$ buffer along with the wedge translated to cable reflection delays of our 90 and 150 m receiver to beamformer cables to highlight the effect of this systematic. 252

5-20 Here we show the ratio between our 2D power spectrum and the error bars estimated by the empirical covariance method of Dillon et al. (2015a). On the top left, we show our data calibrated using our initial calibration (see § 5.2.2) with no attempt made to correct for standing wave structure in the MWA bandpass. Bright, band-like structures are clearly visible at the delays associated with reflections. On the top right, we show a first attempt to correct for cable reflections by fitting a sinusoidal model to rather noisy calibration solutions that had been integrated over a night of observing (1.5 hours each night). While the bands appear weaker, they are still quite visible above the noise. In the bottom right panel, we show the same plot with calibration solutions using scaled autocorrelations described in § 5.3.4. In the lower left panel we show a power spectrum with calibration solutions using autocorrelations for the amplitudes but without any attempt to correct reflections in the phase solutions. Pronounced reflection features are visible in this power spectrum, indicating that any mismodeled reflection structure in the phases will contaminate our measurement. 254

5-21 The level of power at a fixed k_{\parallel} corresponding to the delay of reflections from our 150 m cable (left), and comparing it to a value of k_{\parallel} unaffected by cable reflections (right). The blue line shows the power spectrum level for calibration in which the bandpass is modeled as a polynomial with no attempt to correct fine frequency scale reflections. We see that power is on the order of ~ 50 times the thermal noise level (green-dashed line). Attempting to fit the reflections to calibration solutions integrated over each night gives us an improvement in the power level by roughly an order of magnitude (orange solid line). Using calibration solutions derived from autocorrelations brings down the reflection power by another factor of a few (purple solid line) but is still unable to bring the majority of measurements below the $\sim 1\sigma$ level. While we think that the autocorrelations accurately capture the fine frequency structure of the gains, we are still forced to model this fine frequency structure and predict it in the phases. Residual power is likely due to inaccuracies in this modeling. The right hand panel shows all data below the estimated noise level. This is due to the fact that in (Dillon et al., 2014) it is shown that the method for calculating error bars layed-out in Liu and Tegmark (2011); Dillon et al. (2013) slightly over-estimates the noise. 256

5-22 The Band 1, 1d power spectra from our two nights of observing: September 5th, 2013 (black) and September 6th, 2013 (red). We saw in Fig. 5-7, that the magnitude of refractions on September 6th were on average twice as severe. The two power spectra nearly indistinguishable (within error bars) despite the significant differences in conditions, indicating that ionospheric systematics do not have a significant effect after three hours of integration, even at these low frequencies. 258

5-23 Dimensionless 1d power spectra derived by Integrating spherical shells excluding the foreground contaminated wedge region with a $0.1 h \text{ Mpc}^{-1}$ buffer. Black dots indicate the mean estimated from the weighted average in each bin. Vertical error bars denote the 2σ uncertainties while horizontal error bars indicate the width of window functions. We also shade regions of k-space that we expect to have some level of foreground contamination due to uncalibrated cable reflection structure. Gray shaded regions clearly correspond to regions in which our power spectrum measurements are not consistent with thermal noise. We note that where our upper limits do agree with thermal noise, the power spectrum is on the order of ~ 100 times larger than the upper limits set with the MWA at $\approx 180 \text{ MHz(D15)}$. This factor is reasonable given that the sky noise (noise power spectrum) scales with $\sim f^{-2.6}$ ($f^{-5.2}$) and the primary beam solid angle increases as $\sim f^2$ 260

6-1 The HERA antenna element (bottom) uses a parabolic dish to achieve an order of magnitude increase in collecting area over the PAPER antenna (top). The sleeved dipole in the center of the PAPER backplane is identical to the sleeved dipole being suspended under the cylindrical skirt over the vertex of the HERA dish. The suspended feed arrangement has the potential to introduce intra-antenna reflections which we explore in this work. 275

6-2 A cartoon demonstration of the impact on foregrounds of the frequency dependent beam. Left: The location of three, spectrally flat, sources in delay space assuming a frequency independent beam (no reflections in the antenna element). Right: the presence of chromaticity due to delayed signal within the antenna smears the source in delay with the kernel given by equation 6.7. Since the frequency response of the dish is a function of direction on the sky, the shape of the delay kernel is different for each source line. We see that this smearing can lead to substantial supra-horizon emission. In this paper, we consider a direction independent delay-kernel that is source primarily by reflections while a more general direction-dependent kernel is explored in Thyagarajan et al. (2016) 280

6-3 An illustration of our simulation products and their origin in the HERA antenna geometry. A plane wave is injected from above the feed. The electric field of the plane wave at the feed terminals (red line) along with the voltage output is recorded (black line). The feed in our simulation is situated 5 m above the bottom of the dish, hence there is a ≈ 30 ns delay between when the plane wave passes the terminal for the first time (A) and when it is first absorbed in the dipole (B), leading to the voltage response. Of concern to 21 cm experiments are the subsequent reflections between the feed and the dish (C) which can lead to larger delays that contaminate the EoR window. 282

6-4 The absolute value of the Fourier transform of the voltage output from our dish simulations (green line) and the input wave (blue line), normalized to the amplitude of the input wave at 150 MHz. The ratio between input and output is plotted as a red line. Since our input is limited to frequencies between ≈ 20 and 280 MHz, there are significant numerical artifacts in the ratio that causes divergence towards the plot edges. To eliminate this noise, we multiply by a Blackman-Harris window between 100 and 200 MHz and set our estimate to zero elsewhere (cyan line). 284

6-5 The Fourier transform of the simulated voltage response of the HERA Dish (solid black line) and the PAPER antenna element (dashed black line). Reflections in the HERA dish element lead to significantly enhanced power above ~ 50 ns. Since negative delays should be devoid of signal, they allow us to determine the dynamic range of our simulations which have a numerical noise/sidelobes floor of -60 dB. 286

6-6 The absolute value of the power kernel for the HERA dish (solid black line) and for the PAPER antenna element (dashed black line) calculated using equation 6.7. While an antenna can only physically have a voltage response at positive delays, the delay kernel is formed from the convolution of one antenna with the time reversed conjugate response function of the other. Hence, the power kernel for two identical antennas will have $\tilde{R}(\tau) = \tilde{R}^*(-\tau)$ 286

- 6-7 The power kernel for the three subbands discussed in § 6.4.2 along with the kernel for the full bandwidth response function. While the long term falloff from reflections is prominent between 130 – 160 MHz, it appears at a much lower level in the other two subbands which fall below the central subband by ~ 20 dB at ~ 300 ns. k_{\parallel} values for each delay are computed at 150 MHz. The wider central lobe below 150 ns for the subband gains is due to the lower delay resolution from the smaller bandwidth. We also show the delay kernel for 100 MHz bandwidth (black thick line). 288
- 6-8 Residuals on the absolute value of the gain over several subbands after fitting to a sixth order polynomial. Consistent with our findings in Fig. 6-7, the fine frequency residuals in the 145-155 MHz subband are over an order of magnitude greater than those in the other subbands. 288
- 6-9 A closeup rendering of the HERA feed which is suspended over the reflector, illustrating the cylindrical skirt, the backplane, and the dipole. Long time-scale spectral structure arises from electrical oscillations within the cylindrical cavity. 290
- 6-10 The absolute value of the time-domain voltage response of the cylindrical dipole feed compared to the absolute value of the voltage response of the feed suspended over the dish. As we might expect, the ~ 35 nm lobed structures associated with feed-dish reflections are absent from the simulation of the feed only. However, the knee like feature after ~ 100 ns is. This indicates that the most severe spectral contamination in the current HERA design does not originate in reflections between the feed and the dish but rather within the feed itself. 291

6-11 The time-domain response of the HERA antenna towards a plane wave incident from zenith for a variety of termination impedances. As we vary the termination impedance, the structure which is dominated by feed-dish reflections, below ≈ 150 ns, varies significantly but leaves delays greater than ≈ 200 ns unchanged. Only in the extreme, 500Ω case do the reflections extend to large delays. Since structure above 250 ns is primarily responsible for contaminating the EoR window (6-14), the termination impedance has a relatively small impact on HERA's overall sensitivity. 291

6-12 A comparison between time-domain simulations (black line) and measurements (grey line) of S_{11} for the HERA dish. We also show an S_{11} measurement with the cables leading from the VNA to the feed terminated by an open circuit which allows us to probe the dynamic range of the measurement. We use the standard deviation of the open measurement (grey dashed line) between 200 and 400 ns as our systematic floor (grey shaded region). We find very good agreement between our S_{11} measurement and the simulation, validating the predictions of our simulations. Both the simulations and measurements in this figure were derived from delay transforms over 100 MHz. 294

6-13 The absolute magnitude of a 100-200 MHz delay transformed visibility from a 14-meter baseline (blue line) compared to the same visibility (green line) contaminated by the delay-response observed in our simulations of the HERA dish. We see that the extended delay kernel smooths out structure, originating from foregrounds, within the horizon. For HERA, we expect to use the delay-CLEAN to remove foregrounds. However, the depth of CLEANing is limited by the noise level on a single baseline. We show the foreground residuals from a CLEAN down to the 5σ noise level after 20 minutes of integration. Since CLEANing cannot distinguish between foregrounds and signal, it should only be performed within a narrow region of delay-space, close to the horizon and cannot remove the broad wings leaked by the resonance unless it is accurately modeled. 301

6-14 The delay transform power spectrum for different baselines over several redshifts with (green solid line) and without (blue solid line) the presence of the simulated chromaticity in the HERA dish. Each estimate is computed using the square of a Blackman-Harris windowed delay transform (equation 6.23) over a noise equivalent bandwidth of 10 MHz. We also show the power spectra of foreground residuals after CLEANing the entire 100 MHz band with (dashed cyan line) and without (red dashed line) the simulated spectral structure. For all subbands, except $z = 8.5$, we find that the delay response of the HERA antenna has a negligible effect on the k_{\parallel} where the foregrounds drop below the signal level. This is the same subband where enhanced spectral structure due to resonance in the feed cylinder is observed in our simulations (Figs. 6-7 and 6-8). Outside of the resonance, the chromaticity of the antenna, which is primarily sourced by feed-dish reflections, is practically identical to the achromatic model. It follows that feed-dish reflections have a negligible impact on HERA's 21 cm power spectrum sensitivity. For the achromatic beam, power outside of the horizon is leaked by the window-function of the finite-band Fourier transform. Thus, reducing power inside of the horizon does reduce the level of contamination outside of the window except when the structure is intrinsic to the frequency gain of the antenna as is the case at $z=8.5$ where structure is present at super-horizon delays due to the feed resonance. 303

6-15 The 1σ thermal noise levels achieved by HERA-350 at three different redshifts with (grey line) and without (black line) the presence of beam chromaticity due to the chromaticity observed in this work. We compare these noise levels to the fiducial power spectrum signal (red line). We saw in Fig. 6-14 that with reflections, foregrounds exceed the signal level out to $k = 0.23 h\text{Mpc}^{-1}$ at $z = 8.5$ which we assume are unusable, forcing us to ignore modes out to 350 ns beyond the horizon, leading to the sensitivity projected in the red curve. The absence of these reflections allows us to work within 250 ns of the horizon (green curve), leading to an increase in sensitivity by a factor of ≈ 1.5 307

6-16 95% confidence regions for reionization parameters assuming 1000 hours of observation on HERA-350 between the redshifts between 7.5 and 12. The presence of the reflections leads to an increase in the major axes of these confidence regions by a factor of one to two. 308

6-17 Fractional Errors on reionization and heating parameters as a function of maximal observed redshift out to the low end of HERA-350's initial observing band at $z = 12$. The presence of strong reflections contained within a small subband at $z = 8.5$ has a minimal impact on our overall constraints on reionization parameters. If these reflections are not localized they can worsen our sensitivity to some parameters such as $T_{\text{vir}}^{\text{min}}$ by a factor of two. 308

7-1 Top: The power spectrum residuals computed using equation 7.35 for the MWA, LOFAR, HERA, and the SKA-1 LOW designs with sources modeled down to the array confusion limit. Unmodeled foregrounds are contained within the wedge, which is demarked by dashed black lines at the first primary beam null and solid black lines at the horizon. The calibration errors introduced by these foregrounds bleed out of the wedge into the EoR window. The narrower central lobes (larger stations) employed by LOFAR and the SKA help to significantly reduce the leakage at large k_{\parallel} that exists for the MWA. Contours where the signal, from a 21cmFAST simulation, is equal to unity, five, and ten times the calibration noise are indicated by cyan, orange, and red lines respectively. Bottom: The same as the top for optimistic scenarios. The optimistic scenario for LOFAR and the SKA involves complete modeling of point sources down to 0.1 mJy using additional long baselines. For HERA and the MWA, the optimistic scenario assumes that the 10 mJy source catalog from the TGSS is used for calibration. If long baselines can faithfully model the sources down to 0.1 mJy, modeling noise does not appear to limit LOFAR and the SKA. Sky-based calibration with HERA is improved significantly by using a deep source catalog from a complementary array. The vertical stripe in the LOFAR figure at $k_{\perp} \sim 0.6h\text{Mpc}^{-1}$ arises from the arrangement of the HBA antennas in short spaced pairs so that even the outrigger antennas, which are heavily contaminated, participate in a single short baseline. . . . 337

7-2 Top: Same as Fig. 7-1 except now we consider a perfect calibration catalog with calibration errors arising only from mismodeling the primary beam at the 1% level in the main-lobe and the 10% level in the side-lobes. With the exception of HERA, foreground residuals arising from primary beam modeling errors dominate the signal in the entire EoR window. Bottom: the same as above, but with a uniform modeling accuracy of 1%. Much of the EoR window is still contaminated for LOFAR and the MWA while significant bias exists in much of the EoR window for the SKA. 340

7-3 Using equation 7.51, we show the smallest k_{\parallel}^{\min} that is not dominated by modeling noise errors for arrays with Gaussian primary beams, a random circular Gaussian distribution of N_{ant} antennas with standard deviation σ_{ant} , each with diameter d_{ant} . The area below the white dashed line is where side-lobes render the assumption of Gaussian beams inaccurate. Beige regions on the lower-right hand corner of each plot denote unphysically high packing densities. We see that for all existing instrument designs, calibration noise extends to large k_{\parallel} values that will reduce their sensitivity to the 21 cm signal. HERA will benefit greatly from the fact that it can be calibrated redundantly with minimal reliance on a sky model. 346

7-4 Top: Residual power spectra with each visibility weighted equally in determining the calibration solutions (\mathbf{W} set to the identity matrix). Middle row: the same but now weighting visibilities with a Gaussian function of baseline length (equation 7.52). Much of the EoR window is cleared of contamination from calibration residuals. However pronounced stripes of contamination still exist, especially for LOFAR and the MWA. These stripes arise from short baselines formed from antennas involved in no other short baselines. In order to solve for both antenna gains, they must use information from long baselines, resulting in significant chromaticity on the few short baselines to which the problematic antenna gains are applied. Bottom: flagging visibilities after calibration until all gains participate in $n_{\text{eff}} \geq 2$ baselines, we find the EoR window free of these stripes. To reiterate, solid lines demarcate regions where the fiducial EoR signal is 1, 5, or 10 times the power of the calibration modeling error. The dashed diagonal line indicates the location of the wedge associated with the first null of the primary beam; the solid line indicates the horizon wedge. 350

7-5 Top: Cylindrically binned power spectra of calibration errors due to beam modeling errors at the level of 1% in the main-lobe and 10% in the side-lobes (equation 7.49). Each visibility has been weighted equally in determining the calibration solutions. Bottom row: The level of cylindrically binned power spectrum residuals from the primary-beam modeling errors in the top row but now with calibration solutions derived from visibilities that are weighted with equation 7.52. Short baselines contributing to antennas with $n_{\text{eff}} \leq 2$ have also been flagged from the calibration fit. Weighting with a Gaussian is capable for removing calibration modeling errors due to beam mismodeling at the level that we see in today's experiments (Neben et al., 2015). 351

7-6 The ratio between thermal noise errors on a spherically averaged power spectrum estimate with Gaussian weighting and uniform weighting of visibilities in calibration. We show this ratio for LOFAR where $\sigma_w = 100$ m and the MWA with $\sigma_w = 50$ m. In both cases, the reduction in sensitivity to the power spectrum is by a factor less than two for small k where the interferometers have maximum sensitivity. Under the Gaussian weighting scheme, antennas with fewer short baselines have increased thermal noise in their gains. Increasingly large k -bins include larger numbers of visibilities formed from antennas with fewer short baselines (small n_{eff}) which have large increases in their thermal noise, leading to a trend of increasing sensitivity loss with increasing k . Since the MWA has a narrower weighting function, with $\sigma_w = 50$ m, this increase occurs faster than for LOFAR. 354

7-7 We compare the amplitude of modeling noise on a short baseline when \mathbf{R} is assumed to be diagonal (orange line) and the off-diagonal terms of \mathbf{R} are explicitly included for a 91-element (dashed lines) and 331-element (solid lines) hexagonally packed array of 14 m apertures. We find that even in a maximally redundant array, the effect of redundancy is to change the overall amplitude of the modeling noise by a factor of a few. This only has a small impact ($\lesssim 10\%$) on $k_{\parallel}^{\text{min}}$, the smallest k_{\parallel} where the 21 cm signal (black line) dominates over the modeling noise, as computed from equation 7.51 which ignores the effect of redundancy. 359

7-8 Comparisons between calibration modeling noise with realistic modeled foregrounds (grey lines) and the point source foregrounds used throughout this paper (black lines) with (solid) and without (dashed) Gaussian baseline weighting for two different LSTs and baseline lengths on the MWA from 10 MHz noise equivalent bandwidth centered at 150 MHz. The red line denotes the amplitude of the HI power spectrum generated with 21cmFAST. We find that fully modeled foregrounds change the overall amplitude of the of the calibration noise since the amplitude of a particular modeled visibility does not necessarily equal the amplitudes of every other modeled visibility. Chromatic y_{α} s also introduce some additional spectral structure which results in a larger width of calibration errors in k_{\parallel} . The overall impact on the LoS mode where modeling noise bias falls below the 21 cm signal is only on the order of 10% with Gaussian weighting. 362

List of Tables

3.1	IGM Heating Parameters	132
3.2	Instrumental Parameters	148
4.1	The 1σ error forecasts for reionization and heating parameters on the instruments studied in this paper assuming 1080 hours of drift-scan observations on $\Delta z = 0.5$ co-eval bands over all redshifts between 5 and 25.	203
5.1	Bulk ionospheric and solar weather properties on the two nights of observing presented here. $\langle \text{TEC} \rangle$ indicates the mean total electron content over the entire night. We also show the diffractive scale calculated from all source separations on both nights which differ by a factor of two.	234
5.2	There are N cables of each length (L) between the MWA receivers and beamformers with associated round-trip delay times (τ_r). In the $f_{30\lambda}$ column, we list the percentage of baselines within 30λ at 83 MHz (where the majority of the MWA's power spectrum sensitivity lies) that are formed from at least one tile with the given cable length. We also list the k_{\parallel} of each delay given by equation (5.24) for three different redshifts. Cable reflections that are significantly above the k_{\parallel} values where we expect to obtain maximum sensitivity to the power spectrum at EoR redshifts ($z \approx 7$) move into the maximum sensitivity region at EoX redshifts ($z \approx 16$). Higher order reflections will also contaminate multiples of and differences between the delays and k_{\parallel} values listed in this table (thought at a lower level).	240

6.1	Ellipticity and effective area of the HERA antenna for different feed configurations.	293
7.1	Inspection of the equations in § 7.2 yields a number of analytic and qualitative relationships between the properties of an array and modeling catalog.	342

Acknowledgments

I would like to extend my gratitude to my advisor Jacqueline whose support, insight, healthy skepticism, and skill for asking insightful questions help me grow enormously over my PhD. Over the course of graduate school, I have benefited gratefully from my discussions and collaborations with Joshua Dillon, Adrian Liu, Chris Williams, Lu Feng, Andrei Mesinger, Jonathan Pober, Max Tegmark, Abraham Neben, Danny Jacobs, Rich Bradley, Cathryn Trott, Bryna Hazelton, Nichole Barry, Miguel Morales, Saul Kohn, Aaron Parsons, and James Aguirre.

As an undergraduate, I was fortunate enough to have several fantastic teachers advisors and mentors: Steven Meyer, Clarence Cheng, Lindsey Bleem, Tylor Natoli, Abigail Crites, Kyle Story, Nina Heinrichs, Henry Frisch, and Ronnie Haddani. My success and motivation in college level math and science classes also came in large part through the support and companionship of Sean Weinberg, Louis Smith, Grace Evans, Brian Dressner, Rachel Meng, Tote Hughes, and Guillerme Nesssheheim.

As a member of the astronomical community, I have the extreme privilage of being able to devote my labors towards investigating some of the most fanatstic and beautiful questions facing humankind. I would not have this fortune were it not for the influence and mentorship of several extraoridinary teachers that I was lucky enough to have early in life.

In high school I benefited greatly from the inspiration and encouragement from Carolyn Stewart who provided me with the foundantional tools in computer science and 3d calculus which gave me a huge advantage at the beginning of my college level physics classes and research. A strong interest and conceptual understanding of computers, networking, and programming, which also gave me a head-start in

research was provided by my high-school friends Michael Rauscher, Adam Wilson, and Josh Alberts with whom I would spend long hours building computers, taking apart phones, and installing linux on Xboxes.

My first and greatest teacher is my mother, Nina Ewall, who home schooled me until 7th grade, bringing me up in an environment where the classroom had no compact-support but rather existed in all places at all times. My parent's love and respect of the environment and the natural world coupled with their insistant sharing it with us through countless outdoor hikes that were filled with discussions of geology, botony, and ecology filled me with a desire to examine and understand nature and instilled in me the simple joy of observing. Their trips often brought us to landscapes of stunning vastness, beauty and solitude which awoke in me a passion for observing the Universe, the grandest landscape of all.

While these trips set the kindling, it was my parents' copy of Carl Sagan's *Cosmos, a Personal Journey* which started the fire. I will never forget finding that book in their bookshelf back in 2001 and being moved by Sagan's pros which convinced me that I wanted to devote my life to scientific enquiry.

Chapter 1

Introduction

1.1 The History of the Universe as Revealed through Astronomical Observations

The sky is a time machine. When we point a telescope into the sky, we are dipping a bucket into a torrent of photons emanating from many different streams and tributaries along the river of time. These photons inform a picture of our universe's nature and history. Humanity's unveiling of this picture is a towering achievement but much remains to be explored. This thesis contributes to a new technique for exploring our past, known as 21 cm tomography. The contribution of 21 cm is best appreciated in the context of what we've learned so far. Hence, we shall begin with an overview of our universe and its history as revealed through astronomical observations and what mysteries remain.

Fig. 1-1 is an artistic sketch of the universe's history. One of the great unknown chapters in our timeline is the very beginning (if it is a beginning). Modern physics allows us to extrapolate down to $t \gtrsim 10^{-32}$ seconds and there exist families of plausible models that can explore within $\gtrsim 10^{-43}$ seconds of the greatest mysteries in physics. The other two mostly unobserved chapters are the "Dark Ages", and the "Cosmic Dawn" (lasting from $\sim 400,000$ to ~ 1 billion years after the start of our time line. During these chapters, the first stars, galaxies, and astrophysical black holes formed




Figure 1-1: The history of our Universe as revealed by cosmological observations. To date, our direct observations are limited to the “modern universe” (roughly 1 billion years after the big bang) and a thin slice of the early Universe (the CMB). The critical chapters in structure formation, in which the first baryonic objects collapsed and the first stars ignited (the “Dark Ages” and the “Cosmic Dawn”) remains mostly beyond the reach of today’s instrumentation.

and the Universe underwent a dramatic chemical and structural transformation. We have some ideas of what happened from the laws of physics and requirements of continuity with the other chapters, but the details of the Universe’s early childhood remain shrouded in mystery. What of the other chapters? We discuss first, 400,000 years after the big-bang, which we observe in a thin surface of reflected light known as the “Cosmic Microwave Background” (CMB) (§ 1.1.1) before covering our modern Universe after the Cosmic Dawn (§ 1.1.2).

1.1.1 The Early Universe and the Cosmic Microwave Background

In 1929, Edwin Hubble observed that Galaxies have recessional velocities that increase with distance (Hubble, 1929), establishing that the universe is expanding. However, this expansion alone is consistent with two extremely different cosmic histories. The first is that the Universe expands indefinitely, but the density of matter remains

constant (with new matter being created to accommodate the increasing volume). In such a scenario, the Universe's state, on cosmological scales, is constant in time (Bondi and Gold, 1948). Alternatively, we can imagine that at some point in the distant past the Universe was much denser and much hotter, in a state colloquially known of as the "big bang". This model predicts a CMB and makes specific predictions about the abundances of elements (Alpher et al., 1948) at redshifts before star formation.

The Big Bang triumphed in 1965 when the predicted CMB was discovered by Penzias and Wilson (1965). This photon background, reflected off free electrons and protons just before they combined into neutral hydrogen (HI), has transformed our understanding of the early Universe. Subsequent measurements of CMB anisotropies by experiments including the Cosmic Background Explorer (COBE) (Smoot et al., 1992; Bennett et al., 1996), the Wilkinson Microwave Anisotropy Probe (WMAP) (Hinshaw et al., 2003; Spergel et al., 2003), and Planck (Planck Collaboration, 2013) combined with theoretical modeling of how initial conditions evolve up to the 400,000 year period observed in the CMB (Seljak and Zaldarriaga, 1996) allows us to constrain our Universe's bulk properties and the events that occurred at earlier times. From these CMB observations, along with measurements of primordial chemical abundances (Walker et al., 1991; Fields et al., 2014), galaxy surveys (York et al., 2000; Percival et al., 2001; Peacock et al., 2001; Eisenstein et al., 2011; Anderson et al., 2012), quasar absorption measurements (Croft et al., 2002), gravitational lensing (Wittman et al., 2000; Munshi et al., 2008), and the relationship between the luminosity distances and redshifts of supernovae (Riess et al., 1998; Perlmutter et al., 1999), we have deduced the following about our Universe's bulk properties.

1.1.1.1 The Bulk Properties of the Universe

On scales greater than ~ 100 Mpc, our Universe is well described by a spatially homogenous and isotropic Friedmann-Robertson-Lamaitre-Walker (FRLW) metric whose evolution is governed by the Friedmann equations (Friedmann, 1922). Within experimental error bars, the total energy content of the Universe is equal to the "critical density", which implies that it is either infinite or has a much greater radius of

curvature then our observable volume. What makes up this energy content? In units of the critical density, our Universe has a non-relativistic matter content of $\Omega_m \approx 0.31$ but a baryon content of $\Omega_b \approx 0.05$. The rest of Ω_m is in the form of “dark matter”, $\Omega_c \approx 0.26$, which does not interact with anything electromagnetically. While a multitude of astronomical observations are best explained by non-luminous dark matter, we don’t know what it is. Microlensing measurements (Alcock et al., 2000; Tisserand et al., 2007) rule out significant contributions from baryonic massive compact halo objects (such as black holes and brown dwarfs). Thus, our best candidate for dark matter is a family of one or more weakly interacting massive particles (WIMP). Beyond the gravitational signature of dark matter in the CMB, large scale structure, and galaxy dynamics, the products of WIMP annihilations may have already been detected in positrons (Abdo et al., 2009; Accardo et al., 2014) and X-rays (Bulbul et al., 2014). Our universe contains ≈ 3 relativistic species of neutrinos whose sum of masses is constrained to be $\lesssim 194$ eV. The total energy contribution from these neutrinos plus photons adds up to $\Omega_r \approx 10^{-5}$. While negligible today, the radiation content scales with redshift as $\Omega_r(z) \propto (1+z)^4$ and played an extremely important role during the time before the CMB ($z \gtrsim 3400$). The majority of the Universe’s energy content today is composed of repulsive *dark energy*, $\Omega_\Lambda \approx 0.69$. Since it came to dominate ~ 8 billion years ago, dark energy’s negative pressure and constant density now drives our Universe towards exponential expansion. However, simple models predict its impact was negligible before the CMB.

1.1.1.2 Grand Unification Epoch and Inflation

The unprobed history of the *very* early Universe ($t \lesssim 10^{-32}$) is in many respects more uncertain than the Dark Ages and Cosmic Dawn. Here we overview what is known about these early times. For a more detailed discussion of this time-period along with the epochs covered in § ??, the reader should consult Kolb and Turner (1990); Weinberg (2008). Describing times before $\lesssim 10^{-43}$ seconds requires a theory of quantum gravity which eludes contemporary theory. Physics between 10^{-43} and 10^{-36} seconds may be described by general relativity plus a grand unified theory (GUT) involving

the unification of the electromagnetic and strong/weak nuclear forces. A number of physical theories claim to describe the GUT epoch; however experimental tests now elude us since the energy scales approach $\sim 10^{16}$ GeV, far beyond the capabilities of today's particle accelerators. The GUT epoch ended with the strong nuclear force becoming distinct from the unified electroweak force at $\sim 10^{-33}$ seconds.

The frontier of our experimentally verified understanding of the early Universe is the inflationary epoch at $\sim 10^{-33} - 10^{-32}$ seconds. The inflationary paradigm is the most widely accepted family of theories explaining how our Universe evolved within its first $10^{-36} - 10^{-32}$ sec and elegantly solves three puzzles that exist in cosmology (see Kolb and Turner (1990); Liddle and Lyth (2000); Weinberg (2008)). One: There are coherent regions of the CMB that are so large that light could not have traveled across them by the time of their observation, $\sim 400,000$ years after the big bang. How could these structures come into equilibrium if they were causally disconnected? This is known as the “horizon problem”. Two: Today we observe a universe that is spatially flat (with the energy density equal to the critical density). However, for the universe to appear flat within the limits of our current measurements requires $\Omega_k \lesssim 10^{-60}$ around the time of the big bang: a conceptually unappealing “fine-tuning” of initial conditions. This is known as the “flatness problem”. Three: GUTs predict the production of stable exotic particles, such as magnetic monopoles in the extremely hot and dense conditions of the early universe. No such particles have ever been observed. This is known as the “exotic relics” problem.

Guth (1981) realized that the last problem could be fixed by positing that the very early universe was dominated by a field with a negative equation of state that causes the universe to expand by roughly e^{60} within its first 10^{-32} seconds (inflation), diluting the density of any exotic relics to below one per Hubble volume. Inflation was later explained by a “slow-rolling” scalar field (Linde, 1982; Albrecht and Steinhardt, 1982) that eventually decayed into the standard model particles that we observe today. It was quickly realized that inflation also solves the horizon problem by allowing different parts of the observable Universe to come into equilibrium before blowing them up to beyond the horizon scale, and the flatness problem, by smoothing out any

pre-existing curvature to below e^{-60} . Inflation also gives an elegant explanation for the origin of structure, from the fluctuations in the CMB to the galaxies in Fig. 1-1) (see, for example Liddle (1999)). The structures that we observe today started out as Gaussian fluctuations in the quantum fields that existed before inflation whose statistics are entirely described by their *power spectrum*. The definition for the power spectrum of an isotropic and homogeneous scalar field, $\delta(\mathbf{x})$ is,

$$(2\pi)^3 \delta_D^{(3)}(\mathbf{k} - \mathbf{k}') P(k) = \langle \tilde{\delta}(\mathbf{k}) \tilde{\delta}^*(\mathbf{k}') \rangle, \quad (1.1)$$

where δ_D is the Dirac-delta function and the brackets denote an average over an ensemble of universes with identical statistical properties. When inflation begins, the microscopic quantum fluctuations described by equation 1.1 are blown up to cosmological scales and “frozen in” as they cross the causal horizon, setting the initial conditions for all future evolution.

Observations of the CMB and large scale structure constrain the initial primordial power spectrum of scalar perturbations at $k_* = 0.05 \text{ Mpc}^{-1}$ to be a power-law, $P_s(k) \sim A_s \left(\frac{k}{k_*}\right)^{n_s + \frac{1}{2} \frac{dn_s}{d \ln k} + \dots}$ with $A_s \approx 2.1 \times 10^{-9}$, $n_s \approx 0.97$, and $dn/d \ln k \approx 0$ (Planck Collaboration et al., 2016b). Inflation is also expected to produce tensor perturbations (gravitational waves) (Starobinskii, 1979) that are observable as B-mode polarization in the CMB (Kamionkowski et al., 1997; Zaldarriaga and Seljak, 1998). Upper limits and/or measurements of these B-modes are expected to reveal the energy scale of inflation and information on the GUT physics that ruled the Universe. The Planck collaboration (Planck Collaboration et al., 2016b) has established upper limits on the relative amplitude of the tensor perturbations to scalar perturbations of $r \lesssim 0.1$.

1.1.1.3 Baryogenesis, Nucleosynthesis, and Recombination

The decay of the inflation field leaves us with a Universe filled with the standard model particles that we are familiar with today. Between $10^{-12} - 10^{-6}$ sec, electroweak symmetry breaking occurred, separating the electromagnetic force from the weak nuclear force. At $\sim 10^{-5}$ sec, average particle energies passed into the realm of

the highest energies that can be probed with particle accelerators on Earth ($\sim 10^{13}$ eV). After inflation, quarks, anti-quarks, and gluons dominated the Universe's energy content, existing freely in a plasma, rather than being bound up in the hadrons (e.g. neutrons and protons). As the temperature fell below the binding energy of these systems, hadrons formed and annihilated. A slight, one part in $\sim 10^{10}$ asymmetry in the baryon-anti-baryon content of the Universe results in some hadrons being left over (roughly 10^{-10} for every photon). Some of these will eventually become human beings. At ~ 1 sec, neutrinos decouple, forming a background that may someday be used to probe the time period before the CMB. As the temperature drops below the mass of the electron (≈ 511 keV), e^\pm pairs drop out of equilibrium. Protons and neutrons are kept in equilibrium with neutrinos through various beta/inverse beta decay interactions until the Universe's expansion freezes them out. At this point, neutrons either decay after ~ 10 min or are incorporated into ^2H , ^3He , ^7Li , and mostly ^4He ; the first atomic nuclei besides H .

After ~ 20 min, the Universe is filled with atomic nuclei, but not yet atoms (these require lower-energy electrons and fewer ionizing photons). Roughly 25% of these nuclei (by mass) are ^4He and $\approx 75\%$ are H . Relatively tiny amounts of ^7Li , ^3He , and other light nuclei accompany the ^4He and H , scattering through a sea of free electrons and photons. No heavier nuclei form since not enough time is spent at the required high temperatures and densities which will only exist after the advent of stars.

At $\sim 47,000$ years, the energy density of matter (dark and baryonic) exceeds that of photons and after $\sim 380,000$ years, e-P recombinations outpace ionizations, allowing nearly all of the electrons in the universe to combine with nuclei into atoms (recombination). This has several important consequences. (1) Atoms form. (2) With no free electrons obstructing them, photons are finally able to travel freely over large distances. After scattering off free electrons before they combine into atoms, the photons stream across ~ 13.8 billion light years to be observed as the CMB. (3) Recombination eliminates photon pressure that had prevented the gravitational collapse of baryonic structures. It is shortly after recombination that the first collapsed clusters of baryons could condense into the first progenitors of the first galaxies. (4)

Sound waves traveling in the plasma are frozen into the large scale-structure as baryonic acoustic oscillations (BAO) which manifest themselves as ring-like structures in the large-scale galaxy distribution. The relatively simple physics of the photon-baryon fluid that preceded the CMB allows cosmologists to infer, with high precision, the bulk properties of the Universe just 380,000 years after its birth.

1.1.2 The Modern Universe

The modern Universe in Fig. 1-1 is clearly a very different place from the CMB. Instead of a relatively featureless fluid that is uniform to one part in 10^5 , things are much clumpier; the contrast between overdense regions (like the Earth) and underdense regions (like cosmic voids) is much larger than during the CMB. In addition, space is filled with luminous baryonic structures such as Galaxies that host supermassive black holes in their centers and are filled with stars and planets. The chemical composition has also changed. At the time of the CMB and before, hardly any elements heavier than ${}^4\text{He}$ existed while today many parts of the Universe (such as the Earth) contain elements as heavy as ${}^{238}\text{U}$. The transformation from the CMB to our modern Universe took place during the unobserved chapters in our cosmic timeline.

In the local universe we find a space filled with luminous stars and gas enriched with metals. Most stars exist in gravitationally bound systems of baryons and dark matter, ranging from globular clusters with masses of $\sim 10^2 M_\odot$ to galaxies $\sim 10^{12} M_\odot$ to clusters of galaxies $\sim 10^{14} M_\odot$. The mass fractions of stars within these systems varies significantly from nearly 100% in globular clusters to less than 1% in some dwarf galaxies and clusters. The remaining mass (detectable through dynamics and gravitational lensing) is typically attributed to dark matter. Only $\sim 10\%$ of baryons are observed within galaxies (in gas and stars) (Persic and Salucci, 1992; Bristow and Phillipps, 1994; Fukugita et al., 1998; Shull, 2003; Fukugita and Peebles, 2004). The remaining 90% of baryons, which must exist given primordial abundances and CMB observations, reside in a diffuse filamentary gas between the galaxies, known as the intergalactic medium (IGM) (Cen and Ostriker, 1999; Davé et al., 2001). Much of this gas ($\sim 50\%$) is thought to have been shock heated to $\sim 10^6$ K while collapsing onto

virialized¹ structures forming at $z \lesssim 3$ (Cen and Ostriker, 2006) while the rest exists at $\sim 10^4$ K due to photionization at the end of the cosmic dawn ($z \sim 6$). The complex combination of astrophysical effects impacting the $z \sim 0$ IGM makes it difficult to obtain constraints on any heating in the distant past. While hydrogen and helium after the CMB was primarily neutral, most of the hydrogen and helium in the IGM today is ionized.

Roughly 8 billion years ago ($z \sim 1$), the Universe's energy budget became dominated by dark energy which is causing the expansion of the universe to accelerate (Riess et al., 1998; Perlmutter et al., 1999). This accelerating expansion has the important effect of halting the growth of perturbations and effectively shutting down the formation of any new collapsed structures. We have no idea what this dark energy is. One possibility is the ground state of our Universe's quantum fields (a cosmological constant). However, attempts to compute Λ within this framework vastly over-predicts it (by $\sim 10^{100}$). Other possibilities are that dark energy is a manifestation of a more correct theory of gravity (modified General Relativity) or a new dynamical field with its own structure.

At $z \sim 0$, the observed distribution of virialized objects on scales greater than our Galaxy is consistent with the standard cosmological model dominated by cold dark matter and dark energy (Λ CDM). However, on smaller scales, there are several unresolved discrepancies. First, Λ CDM simulations tend to produce significantly greater numbers of sub-haloes in Milky-way like galaxies than are actually observed (Moore et al., 1999a) (the "missing satellite problem"). Second, the "cuspy" cores of dark matter halos observed in Λ CDM simulations (Navarro et al., 1997) over-predict the concentration of mass in galaxy centers (Flores and Primack, 1994). One possible explanation of these discrepancies is energy injection from vigorous star formation and supernovae explosions. These stellar processes could have driven star forming gas out of low mass halos and/or heated the surrounding environment to prevent them from accreting and forming stars, rendering them dark and unobservable today

¹A structure is said to be virialized when a system becomes stable with its particles obeying the *virial theorem*.

(Bullock et al., 2000; Benson et al., 2002; Somerville, 2002). This same feedback may have dumped energy into the dark matter cores, puffing them up and eliminating “cusps”. In recent years, an additional complication has arisen (“too big to fail”) where the most massive sub-halos of Milky-Way like galaxies, with very deep potential wells that should be able to retain baryons against stellar feedback (Garrison-Kimmel et al., 2013), are overpredicted in Λ CDM simulations (Boylan-Kolchin et al., 2011). This has renewed interest in modifying our cold, collisionless dark matter model through either collisional self-interacting dark matter (Spergel and Steinhardt, 2000; Vogelsberger et al., 2012; Rocha et al., 2013; Peter et al., 2013) and/or relativistic, dark matter that can free-stream and erase substructure (Polisensky and Ricotti, 2011; Lovell et al., 2012; Macciò et al., 2012). A number of purely baryonic factors, such as tidal and ram pressure stripping of dwarf-spheroidals, may also explain too big to fail (e.g. Brooks et al. 2013; Arraki et al. 2014).

1.1.3 The Dark Ages and the Cosmic Dawn

Future observations of the Dark Ages and the Cosmic Dawn, coupled with more complete observations of intermediate redshifts where dark energy came to dominate (Fig. 1-1) not only fills in the narrative of our cosmic history, but will help resolve some of the outstanding problems confronting our observations of the modern and very early universe. For example, the apparent discrepancies between Λ CDM simulations and the distribution of galaxies at $z \sim 0$ might be explained by stellar feedback during these unobserved chapters. The nature of dark energy can be better understood by extending our observations of large-scale structure from $z \lesssim 1$ (covered by optical surveys) to $z \sim 3$. The enormous volume of pristine density fluctuations, uncontaminated by any stellar astrophysics, that existed during the dark ages, contain an enormous amount of cosmological information, dwarfing that of the CMB by many orders of magnitude (Loeb and Zaldarriaga, 2004). One day, measuring these fluctuations could provide us with unparalleled constraints on inflation and fundamental physics (Chen et al., 2016). While direct observations have not yet been realized, this has not stopped us from arriving at a rough picture of what occurred between the

CMB and modern universe. In this section, we overview theoretical progress towards what we think happened over the Dark Ages and Cosmic Dawn.

1.1.3.1 The Dark Ages and the First Small Structures

Were it not for the 10^{-5} level fluctuations in the energy density, possibly blown up by inflation, this Universe would expand indefinitely at all scales with photons and matter forever cooling adiabatically as $(1+z)^{-1}$ and $(1+z)^{-2}$ respectively. Instead of a problem for the distant future, the “cold death” would reign at all times. Fortunately, this is not the case. Instead, primordial fluctuations grew through gravitational attraction. Since these inhomogeneities initially started small ($\delta \sim 10^{-5}$), it is possible to track their growth analytically with linearized versions of the Euler equations in comoving cosmological coordinates (Peebles, 1980). When $\delta \lesssim 1$, one can continue to track the growth of these perturbations using the Zel’dovich approximation (Zel’dovich, 1970).

Eventually, perturbations in the dark matter reached a critical overdensity at which they collapsed into virialized halos (Zel’dovich, 1970) with abundances predicted analytically Press and Schechter (1974); Bond et al. (1991) and through simulations (e.g. Sheth et al. 2001; Jenkins et al. 2001). These non-linear dark matter structures built up hierarchically with less massive virialized halos merging to form larger structures as time progressed. Unlike dark matter, baryons experience collisions amongst themselves and Cosmic Microwave Background photons. Until recombination, baryonic structures were prevented from collapsing by the pressure from frequent collisions with CMB photons. After recombination, the baryons remained thermally coupled to the CMB, cooling as $(1+z)^{-1}$ until $z \sim 150$ after which they cooled as $(1+z)^{-2}$. Since baryons exert pressure through collisions, their non-linear collapse onto dark matter halos requires that gravitational attraction overcomes their pressure support which requires the mass of a collapsed object to exceed a threshold known as the “Jean’s mass” (Jeans, 1902). The Jean’s mass depends on the time-dependent temperature, density, and ionization state of the IGM. The heating and ionization of the Universe, would therefore also affect the future large-scale distribution of baryonic objects. In simulations, the masses of the earliest collapsed dark matter

halos to accrete a significant fraction of baryons is $\sim 10^4 M_\odot$ (Naoz et al., 2011). It is within these relatively tiny dark matter halos, only one hundred millionth the mass of our Milky-Way Galaxy that the first stars formed.

1.1.3.2 The First Star, Galaxies, and Supermassive Black Holes

While baryonic structures exceeding the Jeans mass were able to collapse into virialized dark-matter halos, without some way to further cool, they would be doomed to exist as diffuse clouds of gas at roughly the virial temperatures of their host halos ($\sim 10^3$ K). Today collapsing pre-stellar clouds are cooled by line emission from metals. With only hydrogen, helium, and trace light elements, the primary cooling channel available during the Dark Ages was molecular hydrogen, H_2 , (Haiman et al., 1996a; Abel et al., 1997; Tegmark et al., 1997a) which could form in sufficient quantities to collapse early baryonic halos in the presence of the residual electrons (McDowell, 1961) left over from recombination (Haiman et al., 1996a). Significant progress has been made over the past decade in simulating the formation of the first stars: a task requiring enormous dynamic range to resolve very fine cosmological scales in both space (\sim pc) and time (fractions of a year) (see Barkana and Loeb (2001); Bromm and Yoshida (2011); Barkana (2016) for reviews of our theoretical understanding of the first stars). Initial studies indicated that the first stars might be hundreds of solar masses (Bromm et al., 1999; Bromm, 2000; Abel et al., 2000). More sophisticated modeling has since revealed that the collapsing gas clouds form disks that can fragment into $\sim 10 M_\odot$ stellar masses (Stacy et al., 2010), some of which may merge into more massive stars as time progresses (Greif et al., 2012). While ensemble simulations have yielded primordial mass functions much more top-heavy than today (Hirano et al., 2015), this mass distribution remains wildly uncertain and requires further improvements in modeling precision (Hirano and Bromm, 2016).

The $\sim 10^5 M_\odot$ halos that probably hosted the first stars may not have constituted the “first galaxies” where prolonged stellar populations were maintained (Bromm and Yoshida, 2011). Stellar ignition produced photoionization fronts (Kitayama et al., 2004; Whalen et al., 2004) and luminous supernovae (Bromm et al., 2003; Greif et al.,

2007) which could eject gas from these halos, shutting down any future star formation. The radiative backgrounds from the first stars were also capable of dissociating H_2 . Without H_2 cooling, stars could likely only form in halos with $T_{\text{vir}} \gtrsim 10^4$ K, the threshold for atomic Hydrogen cooling. The same supernovae that sterilized the first stellar birth-places seeded the surrounding IGM with metals which allowed for next-generation Pop-II stars to form. The timing and morphology of IGM enrichment depends on the unknown Pop-III mass distribution. For example, very massive Pop-III stars may have ended their lives in ultra-luminous pair instability supernovae (PISN) (Woosley et al., 2002), flinging their metals far into the intergalactic medium and rapidly altered the IGM's chemical state (Yoshida et al., 2004). Less massive Pop-III stars would have died in hypernovae with substantially lower metal yields (Nomoto et al., 2006) while the very most massive Pop-III stars probably collapsed directly into black holes (Fryer et al., 2001; ?), producing little enrichment (see tables I/II in (Karlsson et al., 2013)). Wise et al. (2012) find that while PISN can quickly raise the IGM metallicity to above 10^{-3} solar, a lower mass Pop-III Initial Mass Function (the mass distribution of new stars) (IMF) can lower this floor by a factor of ~ 10 . The change in the metallicity of the Universe caused its star formation to enter a new regime, triggering later generations of Pop-II stars, which are substantially less massive and longer lived than most Pop-IIIs. Recent simulations indicate that this transition may have happened more slowly than previously expected, with Pop-III stars continuing to form at relatively late times ($z \sim 7$) (Xu et al., 2016). The minimum mass of halos hosting the first long-term stellar populations and contributing significantly to the evolution of the IGM is typically believed to be $\sim 10^8 M_{\odot}$ (with a virial temperature of $\sim 10^4$ K). These halos coalesced from the less massive molecular cooling halos at $z \sim 10$ (Barkana, 2016), had deep enough potential wells to maintain star forming gas against radiative feedback, and were hot enough for atomic hydrogen cooling to serve as a mechanism for stellar collapse.

The formation of the first supermassive black holes also probably occurred during the epoch of the first stars. Luminous quasars arising from accretion onto supermassive black holes with masses $\gtrsim 10^9 M_{\odot}$ are observed within the first billion years

of our Universe’s history (e.g. Mortlock et al. (2011); Wu et al. (2015)) and it is an outstanding puzzle as to how these systems grew over such a short time-scale. There are two popular contenders for the progenitors of such systems during the epoch of the first stars. The first are very massive Pop-III star with $M \gtrsim 100 M_{\odot}$ which would need to accrete close to the maximally allowed “Eddington” luminosity (Madau and Rees, 2001; Volonteri and Rees, 2006). The second is the direct collapse of $M \sim 10^5 M_{\odot}$ clouds into black holes. This special scenario could occur if fragmentation into smaller components is suppressed by the absence of metals and external, H_2 dissociating radiation backgrounds (Loeb and Rasio, 1994; Eisenstein and Loeb, 1995).

1.1.3.3 Reheating and Reionization of the Intergalactic Medium

The optical depth to the cosmic microwave background (Planck Collaboration et al., 2016a), patchy kinetic Sunyaev Zeldovich effect (McQuinn et al., 2005; Aghanim et al., 2008; Reichardt et al., 2012; Mesinger et al., 2012), absorption in the spectra of high-redshift quasars (Fan et al., 2001; White et al., 2003; Mesinger et al., 2004; Carilli et al., 2010; McGreer et al., 2015), gamma-ray bursts (Chornock et al., 2013), and obscuration of Ly- α emitting galaxies (Stark et al., 2010; Pentericci et al., 2011; Caruana et al., 2014), indicate that between the redshifts of ~ 20 and ~ 6 , radiation from the earliest generations of stars, compact objects, and supernovae reionized and heated (to $\gtrsim 10^4$ K) the hydrogen in the IGM (see Furlanetto et al. 2006a; Morales and Wyithe 2010; Pritchard and Loeb 2012; McQuinn 2015; Furlanetto 2016 for reviews along with Greig et al. 2016 for the most up-to-date constraints on reionization). In Fig. 1-7, we show Greig et al. (2016)’s estimates for our existing uncertainty of how the neutral fraction of the IGM evolved over reionization. Two components of the stellar radiation had a crucial role in this transformation: X-rays produced by accreting stellar remnants and/or hot interstellar gas, and ultraviolet (UV) photons. UV photons with energies greater than 13.6 eV ionized the neutral hydrogen that they encountered in the IGM. X-rays were also capable of producing ionization and had much longer mean free paths than UV photons, allowing them

to travel far from their sources. The large energies of X-rays causes them to create a high energy “primary” electron which then distributes the X-rays energy through a cascade of secondary ionizations, HI/HeI excitations, and coulomb collisions with other electrons. The number of UV photons per stellar baryon is expected to be on the order of thousands (Furlanetto and Loeb, 2004) while the number of X-ray photons per stellar baryon is typically $\sim 10^{-1}$ (Mesinger et al., 2013). In addition, the fraction of energy that an X-ray deposits in ionizations drops rapidly beyond $(1 - x_{HI}) \gtrsim 0.1$, with most of it going into heating from coulomb collisions (Shull and van Steenberg, 1985; Furlanetto and Stoever, 2010). Thus UV photons are expected to have been the primary drivers of reionizing the IGM except in extreme models (see Mesinger et al. (2013)). In many scenarios, X-rays, able to deposit large amounts of energy across very large cosmological distances, pre-heated the IGM before the majority of ionization took place (Madau et al., 1997; Chen and Miralda-Escudé, 2004; Furlanetto, 2006b). While stars are thought to have been the dominant sources of UV photons during reionization (Madau et al., 1999; Meiksin, 2005; Faucher-Giguère et al., 2009), some posit that micro-quasars dominated the ionization budget (Madau and Haardt, 2015). The nature of the X-ray sources is also an open question. Canonical models indicate that X-rays were sourced by accreting high-mass X-ray binaries (Baek et al., 2010; Mirabel et al., 2011), hot interstellar medium components ejected by early supernovae (Pacucci et al., 2014), and possibly micro-quasars (Fialkov et al., 2014). How much any one of these X-ray sources contributed to heating is currently unknown. It is also possible that X-rays released by annihilating dark matter (Valdés et al., 2013; Evoli et al., 2014; Lopez-Honorez et al., 2016, 2017) and/or cosmic rays (Sazonov and Sunyaev, 2015; Leite et al., 2017) provided a significant fraction of the pre-heating. In the absence of heating from X-rays, hydrogen atoms would have been heated to 10^4 K by the same UV photons doing the ionization. As mentioned above, $\sim 50\%$ of all baryons in the IGM today underwent additional heating to $\sim 10^6$ K in virial shocks at $z \lesssim 3$ (Cen and Ostriker, 1999), thus the nature of pre-heating (if any) is best probed by direct observations of the cosmic dawn.

1.1.3.4 A Summary of Some Unanswered Questions.

Having reviewed our existing knowledge of the Universe's history and what we think might have happened over its unexplored chapters, we are confronted with two frontiers in physics and astrophysics. Furthering these frontiers will not only further our fundamental picture of reality but also fill in the missing chapters of our cosmic story. Here are some of the unanswered questions from our summary. We begin with the astrophysical questions.

1. What were the properties of the first stars (e.g. their masses) and what environments did they form in?
2. Did subsequent generations of stars form in the same small dark matter halos as the first generation, or were these halos shut down by astrophysical feedback?
3. How long did it take to complete the transition from zero-metallicity Pop-III stars to Pop-II stars?
4. How did supermassive black holes form so quickly, within the first several hundred million years of the Universe?
5. Did stars reionize the Universe or were quasars responsible?
6. Did large or small galaxies contribute to the reionization and heating of the IGM?
7. How did the reionization and heating of the IGM affect the formation of future stars and galaxies?
8. How much was the heating of the IGM contributed to by hot interstellar medium, accreting black holes, annihilating dark matter, and cosmic rays?
9. Can reionization and baryonic physics explain the seeming inconsistencies between simulated and observed halo substructure at the present day or do we need to change the Λ CDM picture?

10. Are there Pop-III stars in the Universe today? What about their black hole remnants?

And here are some of the unknowns in fundamental physics.

1. What is causing the Universe's expansion to accelerate?
2. What precisely is dark matter? What are the masses and interaction cross sections of its constituents?
3. What is the origin of structure in the Universe?
4. Did Inflation happen?
5. If it did happen, what caused inflation?
6. What grand unified theory existed from the Planck time up to the inflation epoch?
7. Are there only three species of relativistic neutrinos in the Universe?
8. Do the properties of the Universe change significantly beyond our observable volume? (see (Tegmark, 2003) for a review of possible scenarios)
9. What is the theory of quantum gravity that can allow us to understand the evolution of the Universe into the Planck Epoch?

A family of cosmological observables should be accessible within the next decade and will help answer these questions. Before focusing on 21 cm experiments, we will review these other techniques in order to understand their strengths, limitations, and where 21 cm can aid them as a powerful complimentary probe.

1.1.4 Expected Progress in the Coming Decade

Besides 21 cm experiments, a number of observational programs are coming online within the decade which promise to revolve many of the above questions. These include

1.1.4.1 Infrared Galaxy Surveys

Since most photons blueward of 912 Å are absorbed by HI, infrared (IR) surveys are an important tool for exploring reionization and beyond. The James-Webb Space Telescope (JWST) is scheduled to launch in 2018 with a detection threshold of $\lesssim 32$ mag between $0.9 - 2.5\mu\text{m}$ (Gardner et al., 2006; Windhorst et al., 2008) while the Wide Field Infrared Survey Telescope (WFIRST) should operate by the mid 2020s (Spergel et al., 2013). These surveys will probe the cosmic dawn through photometric and spectroscopic observations of Galaxies during reionization, performing detailed characterization of the brightest sources in the early Universe. To give a sense of what types of galaxies WFIRST and JWST will investigate, in Fig. 1-2, we show Mason et al. (2015)'s plot of predicted counts for sources with given apparent IR magnitudes along with the sensitivity regions for various surveys. It is clear that both surveys will be capable of studying the details of luminous galaxies within the reionization epoch at $z \lesssim 10$. However, it will be difficult for such IR surveys to probe earlier times. An apparent magnitude of ≈ 30 corresponds to an absolute magnitude of ≈ -20 at $z \approx 10$. It will also be difficult for these IR surveys to probe the low-mass end of the galaxy luminosity function, to determine whether stellar populations were sustained in halos below the atomic cooling threshold. To illustrate this, we show (Liu et al., 2016)'s luminosity functions generated in the Dark-ages Reionization and Galaxy-formation Observables Numerical Simulation (DRAGONS). A turnover in the galaxy counts occurs at $M_{AB} \sim -12$, which corresponds to the atomic cooling threshold which, in this scenario, sets the lower limit on the galaxy masses contributing to reionization. Unfortunately, this faint threshold is ≈ 4 magnitudes below JWST's detection limit. 21 cm promises to supplement IR surveys by probing a signal that is sensitive to the low-mass galaxy cutoff and with upcoming telescopes, potentially observable at redshifts beyond reionization. A second probe of the Cosmic Dawn epoch promised by JWST and WFIRST are constraints on the rates of ultraluminous pair-instability supernovae (Woosley et al., 2002) that may have been generated by the death's of Pop-III Stars, helping to constrain their mass distribution and metal




Figure 1-2: Left: The number density and apparent IR magnitudes at galaxies during and before reionization along with the coverage of planned surveys by JWST and WFIRST from (Mason et al., 2015). Without the assistance of strong gravitational lensing, it will be difficult for IR surveys to look beyond reionization. Right: The UV luminosity function predicted by the DRAGONS simulations at several different redshifts for (Liu et al., 2016). The turnover at $M_{AB} = -12$ corresponds to the luminosity of atomic cooling halos and is significantly below the detection threshold of near-term IR surveys.

enrichment channels (Weinmann and Lilly, 2005; Hummel et al., 2012; de Souza et al., 2013).

1.1.4.2 The Cosmic Microwave Background

Measurements of the anisotropies in the cosmic microwave background have become increasingly more sensitive, detecting the faint B-mode polarization from gravitational lensing (Hanson et al., 2013), while reaching to ever smaller angular scales (higher ℓ) to constrain secondary anisotropies sensitive to the nature of dark energy, dark matter, and reionization. Over the next decade, a series of ground-based “stage IV” (SIV) experiments are expected to come online with $\sim 10^5$ detectors (see Abazajian et al. 2016). By observing or putting upper limits on the B-mode polarization of the CMB, SIV experiments are predicted to constrain the ratio between tensor and scalar perturbations to $r \lesssim 0.001$ (Abazajian et al., 2015a), the sum of neutrino masses to $\lesssim 16$ meV, and the number of relativistic species to $\sigma(N_{eff}) \lesssim 0.020$ (Abazajian et al., 2015b). CMB experiments are thus poised at being able to probe inflationary physics at the grand unification energy scale while providing a wealth of information

on the nature of neutrinos and dark matter.

1.1.4.3 Optical Surveys

Several optical experiments for imaging, including the Dark Energy Survey (DES) (The Dark Energy Survey Collaboration, 2005) (SIII), Euclid (Laureijs et al., 2011), and eventually LSST (LSST Science Collaboration et al., 2009) (SIV), are coming on-line to probe dark energy and the large scale structure in the local universe through the galaxy distribution, weak lensing, and supernovae. The successor to the SIII Baryonic Oscillation Spectroscopic Survey (BOSS) (Eisenstein et al., 2011) is the Dark Energy Spectral Instrument (DESI) (DESI Collaboration et al., 2016). DESI is expected to precisely measure the Universe’s expansion from spectroscopy of galaxies out to $z \lesssim 2$ and beyond. Optical surveys promise to vastly improve our understanding of dark energy through detailed mapping of the “epoch of acceleration” but are susceptible to systematics from astrophysics and the non-linear evolution. Another valuable output from these surveys will be detailed observations of the properties (temperatures, metallicities, surface gravities, ages, etc..) and kinematics of billions of nearby stars within the Milky-Way, useful for near-field cosmology. Detailed measurements of the kinematics of nearby stars will allow us to better map the Galaxy’s gravitational potential, probing the nature of dark matter on small scales. Joint kinematical/age analyses promise to shed light on how galaxies similar to our own assemble. Besides 21 cm, near field studies to uncover the formation histories of nearby dwarf Galaxies are perhaps one of the few direct means that we have of determining whether very low-mass halos participated in reionization and allow us to determine how metals were distributed in the early universe (Karlsson et al., 2013; Frebel and Norris, 2015).

1.1.4.4 Gravitational Wave Astronomy

Massive population III stars will end their fate as black holes, which, in binary systems will eventually inspiral and merge, potentially being detected by already operating gravitational wave observatories such as the advanced Laser Interferometric Gravitational-wave Observatory (aLIGO). Hartwig et al. (2016) note that black hole-

black hole inspirals with $M_{BH} \gtrsim 200M_{\odot}$ must be unambiguously primordial, rather than originating from Pop-I/II stars. These authors find that the rates of such mergers detectable by aLIGO could be $\gtrsim 1$ per-decade albeit with enormous model-dependent uncertainties. Detections or non-detections of such massive gravitational-wave inspirals may therefore help to constrain the primordial IMF.

1.2 The 21 cm Line as a Cosmological Probe.

The advances listed in § 1.1.4 are expected to revolutionize our understanding of the early Universe. However, substantial gaps still exist in both volume and time. As mentioned above, upcoming IR surveys are limited to observing only the bright galaxies during reionization but not much beyond that (which includes the Dark Ages, the epoch of the first stars, and the X-ray heating of the IGM). Even mapping reionization will leave a large fraction of our Hubble volume unexplored, throwing away large amounts of cosmological information.

A promising technique complementing those listed in § 1.1.4 is 21 cm tomography. Most of the baryons in the Universe are neutral and ionized hydrogen. Neutral Hydrogen (HI) possesses a hyperfine splitting of its ground state associated with the spin-spin coupling of the proton and the electron. First predicted by Van de Hulst in 1944 (van Woerden and Strom, 2006) and detected by Edwin Purcell (Ewen and Purcell, 1951), a transition from the triplet to the singlet hyperfine states produces a photon with a wavelength of ≈ 21 cm or a frequency of 1.4 GHz. This line is optically thin. Hence, while it is faint in emission, it does not saturate at small concentrations of HI (unlike the Ly- α line). In this section, we will discuss how the observable brightness temperature of 21 cm emission encodes the physical properties of high-redshift HI gas § 1.2.1, how 21 cm emission is expected to evolve through the Dark-Ages, Cosmic Dawn, and epoch of acceleration (§ 1.2.2) and what we can learn from it. In § ?? we will give a brief overview of existing efforts to detect 21 cm from different epochs our Universe's history.

1.2.1 Astrophysics from the 21 cm Transition

Following standard arguments (e.g. Rybicki and Lightman (1979); Morales and Hewitt (2004); Furlanetto et al. (2006a)), one can show that the optical depth for 21 cm from a patch of hydrogen gas at redshift, z , is given by

$$\tau \approx 9.2 \times 10^{-3} (1 + \delta)(1 + z)^{3/2} \frac{x_{HI}}{T_s} \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right], \quad (1.2)$$

where x_{HI} is the fraction of Hydrogen atoms that are neutral, δ is the baryonic overdensity, $H(z)$ is the Hubble factor at redshift z , and $dv_{\parallel}/dr_{\parallel}$ is the gradient of the radial velocity of the hydrogen gas. The spin temperature, T_s , is defined by the Boltzmann factor associated with the relative number of atoms in the triplet and singlet state (n_1/n_0),

$$\frac{n_1}{n_0} = \frac{g_1}{g_0} \exp\left(-\frac{h\nu_0}{k_B T_s}\right) = 3 \exp\left(-\frac{h_P \nu_0}{k_B T_s}\right), \quad (1.3)$$

where h_P is Planck's constant and K_B is Boltzmann's constant. The physical quantity that is measured in 21 cm observations is the *brightness temperature* in contrast against the Cosmic Microwave Background. This is given by (Furlanetto, 2006b)

$$\begin{aligned} \delta T_b(\nu) &= \frac{T_s - T_{\text{CMB}}}{1+z} (1 - e^{-\tau}) \approx \frac{T_s - T_{\text{CMB}}}{1+z} \tau \\ &\approx 9x_{HI}(1+\delta)(1+z)^{1/2} \left[1 - \frac{T_{\text{CMB}}}{T_s} \right] \left[\frac{H(z)/(1+z)}{dv_{\parallel}/dr_{\parallel}} \right] \text{mK}. \end{aligned} \quad (1.4)$$

While the spin temperature increases, so too does the intensity of emission per unit of optical depth. At the same time, the optical depth decreases as T_s^{-1} , so that the contribution to observed 21 cm brightness from emission is independent of T_s (in part because the relative populations of triplet and singlet states given by equation 1.3 is saturated at typical $T_s \gg h_P \nu_{21}/k_B$. Absorption of the CMB radiation, on the other hand, introduces a factor of $\frac{T_{\text{CMB}}}{T_s}$ which is only significant with $T_s \lesssim T_{\text{CMB}}$. Hence, the spin-temperature only has a significant impact on observed brightness temperature of 21 cm emission when it is relatively cool.

x_{HI} depends on the relative rates of ionizing photons and recombinations. What about T_s ? The spin transition is excited (or de-excited) by three types of interactions and is hence coupled to three different astrophysical temperatures. These are,

1. Stimulated emission and absorption from CMB photons, coupling $T_s \rightarrow T_{\text{CMB}}$.
2. Collisions with hydrogen atoms, helium, deuterium, electrons, and protons, coupling $T_s \rightarrow T_K$, the kinetic temperature of the gas.
3. Interactions with UV photons through the *Wouthuysen Field Effect* (Wouthuysen, 1952; Field, 1959) which couples $T_s \rightarrow T_c$, the Ly- α backgrounds color temperature.

The dependence of T_s on these three temperatures; T_k , T_c , and T_{CMB} , is given by

$$T_s^{-1} = \frac{T_{\text{CMB}}^{-1} + x_c T_k^{-1} + x_\alpha T_c^{-1}}{1 + x_c + x_\alpha} \quad (1.5)$$

where x_c and x_α are coupling constants to collisions and *UV* photons respectively. These are determined by the physical state of the gas and the radiation backgrounds. As mentioned above, the large optical depth of the Lyman- α transition typically sets $T_c \rightarrow T_k$. We have established how the physical state of the hydrogen gas affects the observable brightness temperature. Now, let us explore how this brightness temperature evolves through the Dark Ages and the cosmic dawn and what we can learn from this evolution.

1.2.2 Observables

The ultimate goal of 21 cm tomography is a three-dimensional map of HI emission beginning from the time the spin-temperature decoupled from the cosmic microwave background, all the way to the present epoch. However, due to the faint nature of the signal, we only expect to realize detections of statistical quantities at all but the very lowest redshifts over the next several years. Two quantities are being focused on in particular:

1. The brightness temperature of 21 cm emission averaged over all angular directions, as a function of frequency (the "global" signal), $\langle \delta T_b \rangle$.
2. The power spectrum of spatial brightness temperature fluctuations, $P_{21}(k)$, defined by.

$$\langle \tilde{\delta}_{21}(\mathbf{k}) \tilde{\delta}_{21}^*(\mathbf{k}') \rangle = (2\pi)^3 P_{21}(k) \delta_D^{(3)}(\mathbf{k} - \mathbf{k}') \quad (1.6)$$

where $\delta_{21} = \delta T_b - \langle \delta T_b \rangle$. Note that we leave in the dimensions for temperature in δ_{21} . While the statistical distribution of spatial fluctuations evolves as a function of redshift, it is typical to choose a redshift interval over which the statistics of the signal are approximately constant to estimate the power spectrum (Bowman et al., 2006). Work is underway to produce estimators are able to estimate multiple redshifts simultaneously (Trott, 2016).

Unlike the CMB, the 21 cm signal is not expected to be Gaussian. Thus, while the global signal and power spectrum are the primary targets of experiments within the next several years, future measurements of higher order moments will be necessary to recover all of the information available (Cooray, 2005; Pillepich et al., 2007; Yoshiura et al., 2015). How do astrophysics influence these observables? We will delve into this question in greater detail in chapters 2,3, and 4, but to illustrate the richness of the signal, we show the evolution of both the global signal and the amplitude of fluctuations at the $0.1 h\text{Mpc}^{-1}$ scale, with redshift, in Fig. 1-3. Let's discuss briefly how the astronomical events of the Dark Ages and Cosmic Dawn influence these observables.

Beyond $z \gtrsim 200$, collisions with CMB photons and gas particles are frequent enough to keep $T_k = T_{\text{CMB}} = T_s$. Thus no observable global signal exists below $\sim 7\text{MHz}$. As the universe expands, CMB photon scatterings become increasingly rare, and the gas begins to cool adiabatically with $T_k \propto (1+z)^2$. Collisions bring $T_{\text{CMB}} \gtrsim T_s \gtrsim T_k$ for a time, leading to an absorption feature between $200 \gtrsim z \gtrsim 40$. Regions with higher density have both greater optical depth and higher rates of collisions between atoms, coupling T_s more strongly to T_k and resulting in greater absorption than underdense regions (Zaldarriaga et al., 2004; Pritchard and Furlanetto, 2006).




Figure 1-3: Top: A two-dimensional slice of the brightness temperature field from the Evolution of 21 cm Structure (EoS) 2016 data release (Mesinger et al., 2016). Middle: the corresponding global signal being targeted by single dipole experiments. Bottom: The power spectrum of brightness temperature fluctuations at the $k = 0.1 \text{ hMpc}^{-1}$ scale. The global signal exhibits a $z \approx 100$ dip which occurs when the rate of collisions between CMB photons and HI atoms drops due to the Universe's expansion and rises again as collisional coupling to the HI kinetic temperature becomes similarly inefficient. Ly- α photons from the first stars introduce a dramatic falloff at $z \approx 30$, strongly coupling T_s to T_k . $T_s = T_k$ rises dramatically between $10 \lesssim z \lesssim 20$ as X-rays heat the IGM. As this heating proceeds, high-contrast exists between hot and cold IGM patches introduces significant fluctuations and a maximum in the power spectrum at $z \approx 15$. Fluctuations from the ionization field take over at $z \lesssim 10$ before reionization eliminates the neutral hydrogen in the IGM, bringing both the power spectrum and the global signal to ~ 0 .

The result is a vast volume of fluctuations whose relationship to the primordial density fluctuations set in by inflation is relatively simple to model, unless some exotic energy injection mechanism (such as annihilating dark matter) exerted an influence during these times in which case the Dark Ages could serve as an exquisite probe of new energy sources (Furlanetto et al., 2006b; Valdés et al., 2013; Evoli et al., 2014). Below $z \lesssim 40$, the gas becomes diffuse enough that collisions can no longer keep T_s coupled to T_k and it drifts back towards T_{CMB} , leading the fluctuations and mean amplitude of δT_b to drop back to zero.

The ignition of the first stars at $z \lesssim 40$ floods the universe with Ly- α photons (Pritchard and Furlanetto, 2006). The large optical depth of the Ly- α transition rapidly brings the color temperature of these photons, T_c into equilibrium with T_k while strongly coupling T_s to $T_c = T_k$ (Field, 1958). This results in a rapid drop in $\langle T_s \rangle$ and $\langle \delta T_b \rangle$. Fluctuations in this Ly- α background are also present, boosting the power spectrum into a “first stars” peak at $z \approx 20$. Ly- α quickly saturates, leading to a drop in the fluctuation amplitude. The deaths of the first stars are followed by the first generation of High-Mass X-ray binaries and hot interstellar medium which emit X-rays that heat the interstellar gas, raising T_k and T_s . A dearth of X-ray emitting sources (inefficient star formation) or hard X-ray spectra can delay heating (Mesinger et al., 2013; Fialkov et al., 2014) and lead to an even deeper absorption trough than the one shown in Fig. 1-3. Alternatively, significantly more efficient X-ray emission (as we might expect for a lower metallicity IMF (Mirabel et al., 2011)), or additional heating from shocks (Gnedin and Shaver, 2004) might reduce the amplitude of the absorption feature as can heating from annihilations in certain dark matter models (Evoli et al., 2014). Recent attempts to calibrate heating to the high- z UV-luminosity function observed by the Hubble Space Telescope yield less efficient heating than previously expected (Mirocha et al., 2017). In many models, the appearance of hot emitting regions within the highly absorptive cool IGM, during the rise in the global temperature, leads to very dramatic fluctuations in the brightness temperature (note that the $(1 - T_{\text{CMB}}/T_s)$ term in equation 1.4 has a very large dynamic range between $-\infty$ and 0). Eventually, T_s exceeds T_{CMB} and δT_b passes through zero after which the

HI appears in emission. In standard “inside-out” models of reionization, UV photons eventually carve out bubbles around early ionizing sources and after T_s saturates, dominate the fluctuations in δT_b . The statistics and morphologies of these bubbles inform us on the nature of the sources contributing to ionization (i.e. What were the masses of the galaxies? How much were photons self absorbed? Were the sources mainly Pop-II stars, Pop-III stars, or perhaps even quasars? How far could the ionizing photons travel within an HII bubble before being absorbed by dense self-shielded Lyman- α systems (Sobacchi and Mesinger, 2014)?) During reionization, it is also possible to probe cosmology either directly through redshift-space distortions (Barkana and Loeb, 2005a), or indirectly by better constraining the optical depth to the CMB, τ which breaks degeneracies with fundamental cosmology parameters such as the amplitude of scalar fluctuations (Liu et al., 2016; Liu and Parsons, 2016). The constraints on τ obtainable by an interferometer are model dependent while a model-independent estimate of τ is possible if heating was sufficient to saturate T_s (Fialkov and Loeb, 2016). The ionized bubbles eventually merge and the mean neutral fraction drops to zero, along with δT_b . After reionization, neutral hydrogen primarily resides in galaxies where the recombination rate still exceeds ionizing flux. During the epoch of acceleration, 21 cm emission traces the distribution of dark matter (Pritchard and Loeb, 2008), allowing for precision cosmology from intermediate redshift 21 cm measurements.

1.3 Observing 21 cm from the Cosmic Dawn

1.3.1 Current Experiments

While this thesis focuses on attempts to detection the power spectrum of 21 cm fluctuations, a number of experiments are underway to detect 21 cm emission from reionization, the cosmic dawn, and the epoch of acceleration. These include single-antenna *global signal* experiments which aim to detect the global absorption signal during the cosmic dawn: LEDA (Bernardi et al., 2015), SARAS (Patra et al., 2013), SCIHI

(Voytek et al., 2014), EDGES-LOW (Monsalve et al., 2017), and DARE (Burns et al., 2012). BIGHORNS (Sokolowski et al., 2015) and EDGEs (Bowman and Rogers, 2010; Monsalve et al., 2017) target the rise and fall in the global signal during reionization.

Probing the spatial fluctuations in the 21 cm emission can either be accomplished by steerable dish telescope (such as the Green Bank Telescope (GBT) which is being used to probe low redshift galaxies during the epoch of acceleration (Masui et al., 2013), or with an interferometer. Interferometers consist of arrays of antennas. To form images of the sky, an interferometer cross-correlates the voltages from different antenna pairs to form *visibilities*. The visibility formed from the i^{th} and j^{th} antenna of an interferometer is (ignoring polarization for pedagogical reasons),

$$V_{ij}(\nu) = \langle v_i(\nu)v_j^*(\nu) \rangle_T, \quad (1.7)$$

where $\langle \cdot \rangle$ denotes an average in time. Each measured visibility is related to the brightness distribution on the sky, $I(\nu, \hat{\mathbf{s}})$, through (Thompson et al., 1986),

$$V_{ij}(\nu) = g_i(\nu)g_j^*(\nu) \int d\Omega e^{-2\pi i \mathbf{b}_{ij} \cdot \hat{\mathbf{s}} \nu/c} A_{ij}(\nu, \hat{\mathbf{s}}) I(\nu, \hat{\mathbf{s}}), \quad (1.8)$$

where \mathbf{b}_{ij} is the displacement between the two antennas, $A_{ij}(\hat{\mathbf{s}})$ is the product of the voltage pattern beam of antenna i with the voltage pattern beam of antenna j , and g_i/g_j are the complex gains of each antenna arising from the analog signal path. Assuming that the gains are corrected, V_{ij} is essentially the Fourier transform of the intensity distribution of the sky multiplied by the antenna beams and a geometric factor that can be ignored for small fields of view or absorbed into the beams. Thus, one can obtain maps of the intensity distribution on the sky by measuring a large set of visibilities with different \mathbf{b}_{ij} and inverting the Fourier transform. The output from a global-signal dipole can be determined by taking $\mathbf{b}_{ij} \rightarrow 0$ (i.e. the correlation of the voltages from a single antenna with itself).

Interferometers (shown in Fig. 1-4) targeting the Epoch of Reionization (EoR) include the Giant Metrewave Telescope (GMRT) (Paciga et al., 2013), the Murchison Widefield Array (MWA) (Tingay et al., 2013a), the Low Frequency Array (LOFAR)

(van Haarlem et al., 2013), the Precision Array for Probing the Epoch of Reionization (PAPER) (Parsons et al., 2010), and the Hydrogen Epoch of Reionization Array (HERA) (DeBoer et al., 2016). Both the MWA and HERA are potentially capable of probing higher redshifts, $z \lesssim 20$ and much of this thesis will discuss the efforts to accomplish this. The Long Wavelength Array (LWA) (Ellingson et al., 2009) covers a frequency range from 30 MHz to 82 MHz also aims to probe the cosmic dawn signal from the first stars X-ray heating. A number of interferometric facilities are also coming online to probe HI in galaxies during the epoch of acceleration. These include the Canadian Hydrogen Intensity Mapping Experiment (CHIME) (Bandura et al., 2014) and the Hydrogen Intensity and Realtime Analysis eXperiment (HIRAX) (Newburgh et al., 2014). One can see that the arrangement of antennas varies dramatically from array to array. While PAPER, HERA, HIRAX, and CHIME all have significantly different antenna elements, they are similar in that their antennas are arranged in a regular pattern on the ground while the MWA, LOFAR, and the GMRT arrange their antennas randomly. The differences in these two array designs reflect two different strategies in detecting 21 cm intensity fluctuations in the presence of vastly brighter foregrounds which we now discuss in § 1.3.2.

1.3.2 Foregrounds and Array Design

Today, a number of experiments exist with enough sensitivity to have already detected the 21 cm signal in the absence of foregrounds (Pober et al., 2014). Over the range of frequencies targeted by 21 cm cosmology experiments, from the cosmic dawn to the epoch of acceleration, synchrotron emission from continuum processes in our Galaxy and the multitude of galaxies between us and the cosmic dawn give rise to radio foregrounds that are $\sim 10^4$ times brighter (Wang et al., 2006; Di Matteo et al., 2002). Both global signal and fluctuation experiments aim to take advantage of the fact that the intensity synchrotron foregrounds is known to evolve smoothly, with negligible power at small line of sight scales. For the signal, each observed frequency corresponds to a different patch of Hydrogen at a different distance along the line of sight with a different density, ionized fraction, and spin temperature, thus there




Figure 1-4: Several examples of existing and upcoming interferometry experiments pursuing a detection of cosmological 21 cm fluctuations. From left to right and top to bottom: LOFAR, the GMRT, the MWA, HERA, PAPER, the LWA, CHIME, HIRAX, and the SKA-1 LOW. Several of the arrays (“imaging arrays”) have their antennas distributed in a random fashion while others arrange their antenna elements in regular patterns (“redundant arrays”).




Figure 1-5: While the foregrounds in 21 cm experiments are a factor of $\sim 10^4$ times brighter, they are expected to evolve smoothly in frequency and are thus characterizable by a power law. Cosmological 21 cm emission is expected to have complex frequency structure. Thus the foregrounds and signal can be isolated by a Fourier transform in the frequency direction.

should be significant power from 21 cm on small LoS scales while the foregrounds should have negligible power (Fig. 1-5).

Inspecting equation 1.8, one sees that if we are interested in measuring the power spectrum, an interferometer already performs the angular Fourier transform for us. Thus, for small fields of view where we can ignore the effects of the curved sky, b_{ij}/c is proportional to a cosmological wave-number perpendicular to the line-of-sight, $b_{ij}/c \propto k_{\perp}$ (Morales and Hewitt, 2004; Parsons et al., 2012a). For small frequency intervals, ν is proportional to the comoving radial distance, r . Thus, the Fourier dual to ν , delay (τ) is proportional to the comoving wavenumber parallel to the line of sight. $\tau \propto k_{\parallel}$. If the only frequency structure in our measurement was from $I(\nu, \hat{s})$ in equation 1.8, then filtering foregrounds and measuring the power-spectrum from 21 cm emission would be a simple matter of Fourier transforming in the frequency direction and multiplying by the appropriate constants to convert from native interferometry quantities ($\text{Jy}^2 \text{Hz}^2$) to cosmological power spectrum quantities ($\text{mK}^2 \text{Mpc}^3$) (Parsons

et al., 2012a).

It is clear from equation 1.8 that this is not the case. There are three different ways that an interferometry measurement introduces additional frequency structure into the otherwise smoothly varying foregrounds, all of which are apparent in equation 1.8.

1. The Fourier exponent. The same phase term that preserves spatial information and enables synthesis imaging introduces frequency structure. Holding the frequency structure for all other terms constant, we see in equation 1.8, that each line-of-sight is multiplied by complex sinusoid with a delay of $\tau = \mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c$ (Dillon et al., 2015b). The maximal frequency line of sight contributing to a visibility with fixed \mathbf{b}_{ij} is in the direction $\hat{\mathbf{s}} \parallel \mathbf{b}_{ij}$ and at the horizon, $\tau_{max} = b_{ij}/c$. Thus, foregrounds in a visibility are leaked from $\tau = 0$ to $\tau_{max} = b_{ij}/c$. The longer a baseline, the larger the contaminated delay. Hence, the foregrounds fill in a “wedge” (Fig. 1-6) in $b_{ij}/c - \tau$ or $k_{\perp} - k_{\parallel}$ space (Datta et al., 2010; Vedantham et al., 2012; Parsons et al., 2012b; Morales et al., 2012; Trott et al., 2012; Thyagarajan et al., 2013; Liu et al., 2014a,b). A single antenna measurement is obtained for $b_{ij} \rightarrow 0$. Thus, the wedge does not directly impact global-signal experiments though frequency structure in $A_{ij}(\nu, \hat{\mathbf{s}})$ arises through similar considerations.
2. The Primary beam. The response of an antenna's primary beam, $A(\nu, \hat{\mathbf{s}})$ is related to the antenna aperture $B(\nu, \mathbf{x})$ through (Kraus, 1973)

$$A(\hat{\mathbf{s}}, m) = \left| \frac{\nu^2}{c^2} \int_{\text{aperture}} d^2\mathbf{x} e^{-2\pi i \mathbf{x} \cdot \hat{\mathbf{s}} \nu/c} B(\nu, \mathbf{x}) \right|^2 \quad (1.9)$$

Thus the chromatic complex exponent that sources the wedge in interferometry measurements also sets a maximal chromaticity for an antenna aperture $\tau_{max} = r_{max}/c$. Additional chromaticity in $B(\nu, \mathbf{x})$ can be introduced by reflections within the antenna structure (Ewall-Wice et al., 2016c; Patra et al., 2017). This source of frequency structure plagues global signal measurements just as much as those by an interferometer.




Figure 1-6: A cartoon representation of the “EoR window” from Dillon et al. (2015b). In the absence of instrumental effects, foregrounds would possess smooth frequency evolution and the EoR window would extend to the same small k_{\parallel} for all k_{\perp} . The Fourier exponent in equation 1.8 gives rise to the “wedge”, denoted by the dark red and orange features. Chromaticity in the primary beam, arising from equation 1.9 adds additional contamination which is usually accounted for by a buffer denoted by the yellow region.

3. The analog signal path. Additional frequency structure is introduced through reflections, filters, and digitization. Since these effects appear down-stream of the antenna, they can be described as a frequency dependent multiplicative gain.

Global signal experiments must either ensure that their instrumental response (items 2 and 3) does not introduce frequency structure into the foregrounds that mimics the global signal evolution (in particular dips with frequency scales similar to the absorption trough or reionization drop). Interferometers must do the same with the additional systematic of the “wedge” (item 1). Several strategies are being pursued to mitigate instrumental frequency structure and drives the dichotomy of regular versus random antenna arrangements in Fig. 1-4.

The regular configurations are colloquially known as “redundant” arrays. In the absence of gain and beam-pattern differences, the quantity measured by a baseline only depends on the relative displacement of the two antennas. Thus, regular arrays of identical antennas obtain many redundant measurements of the same intrinsic visibility. Redundancy allows such instruments to employ a form of calibration known as “redundant calibration” which uses the repeated measurements to over-constrain and correct the gains of each antenna (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014). The regular sampling function of such arrays tends to produce grating side-lobes in the image domain, which are degrade the limit to which point-source foregrounds can be modeled and subtracted. Thus, some redundant arrays (namely PAPER and HERA) give up on trying to measure the 21 cm signal within the wedge and focus on maximizing the number of measurements within the “EoR window” by packing their antennas close together (providing lots of short baselines) (DeBoer et al., 2016). Other redundant arrays probing the epoch of acceleration hope to invert the wedge using inverse covariance estimators and high-precision characterization of the antenna beams and gains (Shaw et al., 2014, 2015; Newburgh et al., 2014).

The arrays with random antenna spacings obtain more complete sampling of the Fourier-sky and a superior ability to model and subtract foregrounds. Larger antenna spacings are also employed to increase the resolution of these arrays. This drives down

the confusion limit; the flux limit at which more than one point source falls within the array’s point spread function (Condon, 1974). The preponderance of long baselines in an imaging array represents a gamble on foreground subtraction since it leaves fewer measurements within the EoR window. If the foregrounds cannot be removed within the horizon wedge, one is left with a large number of foreground contaminated modes and significantly reduced sensitivity (Pober et al., 2014). In the absence of redundant measurements, calibration of these “imaging” arrays requires that we provide a detailed model of the sky to bring the number of degrees of freedom (N_{ant} complex gains) below the number of measurements. Examples of such algorithms include *stefcal* (Salvini and Wijnholds, 2014) and peeling (Mitchell et al., 2008; Ord et al., 2010).

1.3.3 The state of 21 cm power spectrum measurements at the time of this Thesis.

At the writing of this thesis, instruments exist (the MWA and LOFAR) with sufficient sensitivity to obtain a low SNR detection of 21 cm fluctuations after ~ 1000 hours of integration but none has demonstrated the suppression of foreground leakage outside the wedge to below thermal sensitivity limits within the EoR window. In Fig. 1-8, we show existing upper limits on the 21 cm power spectrum over reionization to date. Results from the MWA (Dillon et al., 2013, 2015a; Jacobs et al., 2016; Trott et al., 2016; Beardsley et al., 2016), the GMRT (Paciga et al., 2013), and LOFAR (Patil et al., 2017) are limited by systematic detections outside of the wedge, some of which are clearly related to reflections within the analog signal paths of antennas. PAPER, so far, has achieved the most stringent limits on the 21 cm signal, roughly 400 mK^2 at $z = 8.4$, one to two orders of magnitude below the next best limits from competing experiments and it too is limited by marginal detections which are likely caused by foreground leakage, although particularly inefficient heating scenarios are capable of generating large 21 cm fluctuations of this amplitude (see Mesinger et al. 2013). This has allowed PAPER to rule out no pre-heating of the IGM before reionization




Figure 1-7: Left: Greig and Mesinger (2017)’s estimate for the $1 - \sigma$ confidence region for the evolution of the neutral fraction of the intergalactic medium as a function of redshift. Right: Greig et al. (2016)’s calculation of existing limits on the spin temperature and neutral fraction the intergalactic medium at redshift 8.4 (Greig et al., 2016) including (top-left) recent 21 cm upper limits by (Ali et al., 2015), 21 cm and McGreer et al. (2015) dark-pixel statistics from high redshift quasars (top right), 21 cm and CMB observations (bottom left)

(?Greig et al., 2016) (see Fig. 1-7). These results indicate that unacceptably large fine-scale frequency structures are imprinted by today’s instruments and that existing calibration techniques are insufficient in suppressing these structures. Several of the chapters in this thesis take on the problem of instrumental frequency structure. In addition, no experiment has attempted to measure 21 cm fluctuations beyond $z \approx 11$ (Jacobs et al., 2015). Measurements at higher redshifts should prove invaluable since they probe the epoch in which fluctuations were sourced by spin temperature (X-ray heating) fluctuations rather than ionization fluctuations. Such measurements will allows us to 1) learn about high redshift X-ray sources and 2) break degeneracies between these two astrophysical processes.

1.4 This Thesis

The work presented in this thesis contributes to the field of 21 cm cosmology in three major ways.




Figure 1-8: All of the upper limits on the 21 cm power spectrum at the time of this thesis between $z \approx 6$ and $z \approx 12$. (Jacobs, Private Communication). The sensitivity of the next generation HERA instrument is shown as a yellow line.

1.4.0.1 The Sensitivity of Interferometers to the Power Spectrum and astrophysics before Reionization.

First, through several theoretical works, it explores to what extent measurements with an interferometer at redshifts beyond reionization ($10 \lesssim z \lesssim 20$) can help us constrain the properties of X-ray emitting high energy processes and the population of radio-loud quasars at very high redshifts (chapters 2,3, 4).

Chapter 2 is a study of what regions of astrophysical parameter space might be detectable by 21 cm power spectrum measurements during and before reionization. In this work, I collaborated with Andrei Mesinger who ran suites of semi-numerical cosmological simulations to determine the amplitude of the 21 cm over parameter space while I undertook calculations of the sensitivities of existing and upcoming arrays.

Chapter 3 presents a novel technique for constraining the population of radio loud quasars before reionization by measuring the signature of HI absorption of their spectra in 21 cm maps. In this work, I developed a formalism for describing the contributions of the radio-loud quasar population to the power spectra of these maps, com-

bined semi-numerical simulations of the 21 cm signal, provided by Andrei Mesinger, with models of the radio-loud galaxy population to compute the power spectrum contribution from their absorption.

In chapter 4, we use the Fisher-Matrix formalism to determine what constraints on astrophysical parameters can be obtained with next-generation measurement of HI fluctuations before reionization. In this paper, I wrote a program to compute the derivatives of the power spectrum with respect to different astrophysical parameters using ensembles of 21cmFAST(Mesinger et al., 2011) simulations along with a pipeline to combine these derivatives with calculations of interferometer sensitivities to obtain confidence ellipses on the astrophysical parameters. Significant advice on the structure of this work and the models to address was provided by Andrei Mesinger, Joshua Dillon, Jonathan Pober, and Adrian Liu.

While Chapters 2, 3, and 4 build our motivation for extending our measurement of HI fluctuations before reionization with existing and upcoming interferometers, the systematic obstacles to a pre-reionization measurement (which are also faced by reionization measurements) are significantly more challenging. The diffuse plasma that surrounds the Earth, known as the ionosphere, absorbs and refracts electromagnetic waves with a severity that increases towards lower frequencies. The amplitude of foreground emission increases dramatically towards lower frequencies. For fixed antenna geometries, the primary beam increases in size, introducing emission closer to the horizon and the edge of the EoR window. Fixed baseline lengths sample larger angular scales, increasing the impact of the exceedingly bright diffuse contamination from within our own galaxy. At the same time, fluctuations in ionospheric Faraday rotation may lead to enhanced suppression of polarized foregrounds at low frequencies, aiding in a power-spectrum detection. In order to determine the impact of these systematic challenges and to obtain a first upper limit on the 21 cm power spectrum beyond $z \approx 11$, we conduct an observing campaign on the MWA between 75 and 113 MHz. The results of this campaign are reported in Chapter 5. With the three hours of data that we analyze, we find our measurements to be limited by frequency-dependent structure in the instrumental gains and are unable to

remove these structures with existing sky-based calibration procedures. The results of this analysis deliver a stark message: To make further progress, radio interferometers must either be designed so that they are characterized by smooth frequency dependent gains or calibration methods must be developed that can robustly correct antenna gains. In this work, I built the pipeline for calibrating the visibilities taken by the MWA and converted the calibrated products into maps. I conducted the analyses of these maps to determine ionospheric conditions, and produced data-cube inputs for a power-spectrum pipeline managed by Joshua Dillon. I also conducted a detailed analysis of calibration solutions to determine the sources of foreground leakage in our power spectra and designed a calibration method using information from autocorrelations that improved our sensitivity.

The last two chapters of this thesis (6 & 7) establish promising new techniques for overcoming the difficulties encountered in Chapter 5. Chapter 6 describes electromagnetic simulation studies of the next generation Hydrogen Epoch of Reionization Array (HERA) antenna. HERA greatly enhances the collecting area of 21 cm experiments over previous generations by attaching parabolic dishes to the front-end of PAPER. These dishes introduce an additional risk of analog reflections which we study in this chapter. By combining our antenna modeling with realistic simulations of radio foregrounds, we determine that even in the presence of these reflections, HERA should still be able to deliver reionization science. The antenna designs in this work were arrived upon by Rich Bradley and David DeBoer and Rich Bradley constructed the initial electromagnetic models and carried out measurements with a vector network analyzer on a prototype antenna. I modified the electromagnetic models, conducted the analysis of the simulation outputs and measurements and determined the impact that the antenna performance, observed in these simulations, would have on HERA's ability to constrain reionization.

In Chapter 7, we use analytic techniques to understand why sky-based calibration has so far failed in correcting instrumental gains to the precision required by 21 cm science. In particular, we find that necessarily incomplete sky models introduce errors into calibration gain solutions which bleed power out of the wedge. In understand-

ing how this contamination scales with array and foreground model properties, we suggest a promising solution of down-weighting long baselines in calibration; a technique that we are currently pursuing in observations and simulations. In this work, I applied the analytic formalism in (Liu et al., 2010) to the problem of sky-based calibration and extended it further to describe the impact of calibration errors on the power spectrum. After exploring several potential solutions including long-baseline flagging (which often lead to problematic singular matrices), I evaluated and settled on the convenient solution suggested by Joshua Dillon of using weighting (rather than flagging) to eliminate the leakage from long to short baselines.

We re-summarize the findings of this thesis and conclude in Chapter 8

Chapter 2

Reionization and Beyond: detecting the peaks of the cosmological 21cm signal

The content of this chapter was published in *Mesinger, A., Ewall-Wice, A. and Hewitt, J. (2014), 'Reionization and beyond: detecting the peaks of the cosmological 21 cm signal', MNRAS 439, 3262–3274.*

2.1 Introduction

The dawn of the first stars and black holes of our Universe is at the forefront of modern cosmological research. The redshifted 21cm line from neutral hydrogen will arguably provide the largest insights into these epochs. The 21cm signal is sensitive to the ionization and thermal state of the gas, and is therefore a powerful probe of both the intergalactic medium (IGM; where most of the baryons reside), as well as the first galaxies (whose radiation governs the evolution of the IGM). Since it is a line transition, the 21cm signal can tell us about the three dimensional structure of cosmic gas, and its evolution. First generation interferometers, like the Low Frequency Array (LOFAR; van Haarlem et al. 2013)¹, Murchison Wide Field Array (MWA; Tingay

¹<http://www.lofar.org/>

et al. 2013a)², and the Precision Array for Probing the Epoch of Reionization (PA-PER; Parsons et al. 2010)³ are coming on-line, with second-generation instruments such as the Square Kilometre Array (SKA; Mellema et al. 2013)⁴ soon to follow, offering the promise of full tomographical imaging of the early Universe.

Given that initial interferometric measurements will likely be noise-limited, the first-generation instruments are focusing on statistical detections: going after the large-scale, spherically-averaged 21cm power spectrum. Furthermore, efforts have mostly focused on the reionization epoch. However, it is highly likely that the peak in the amplitude of large-scale fluctuations occurred during the preceding epoch when X-rays began heating the cold IGM (e.g. Pritchard and Furlanetto 2007; Mesinger and Furlanetto 2007; Baek et al. 2010; Santos et al. 2010; McQuinn and O’Leary 2012; Mesinger et al. 2013). Sourced by strong absorption of cold gas against the CMB and large temperature fluctuations in the IGM, the 21cm power during X-ray heating is expected to be at least an order of magnitude higher than that during reionization. As we shall see below, in many cases this increase is large enough to compensate for the increase in the thermal noise of the interferometer at the corresponding lower frequencies. Considering sensitivity alone, detecting the heating epoch in 21cm interferometry can therefore be of comparable difficulty to detecting reionization, though the potential challenges of radio frequency interference (RFI; the heating epoch extends through the FM band) and calibrating a larger beam might pose additional challenges.

Heating is expected to be dominated by the X-rays from early astrophysical sources, most likely X-ray binaries (XRB; e.g. Mirabel et al. 2011; Fragos et al. 2013). However, some classes of popular annihilating dark matter (DM) models can also imprint a strong signature in the IGM thermal evolution, which is not degenerate with that of the astrophysical X-rays (e.g. Chuzhoy 2008; Valdés et al. 2013; Evoli et al. 2014)⁵ Therefore the 21cm power spectrum during the heating regime encodes

²<http://www.mwatelescope.org/>

³<http://eor.berkeley.edu>

⁴<http://www.skatelescope.org/>

⁵ Initially, other sources of heating were thought to be important, sourced by the Ly α background (Madau et al., 1997) and structure formation shocks (e.g. Gnedin and Shaver 2004). However,

valuable astrophysical and even cosmological insight.

In this work, we quantify the detectability of *both* X-ray heating *and* reionization with upcoming and future 21cm interferometers. We perform an astrophysical parameter study, exploring different minimum DM halo masses required to host galaxies, M_{\min} , as well as the galactic X-ray emissivity. We also discuss the observability of the signal in terms of popular warm dark matter models, recasting M_{\min} to an analogous warm dark matter particle mass, m_{wdm} . During the completion of this work, a similar study was presented by Christian and Loeb (2013). Our work extends their results, including broader astrophysical parameter space exploration, furthering physical intuition, and incorporating sensitivity models of several upcoming interferometers.

We focus on detecting the power spectrum in a single k -bin, centered around $k = 0.1 \text{ Mpc}^{-1}$, which Pober et al. (2013a) identifies as being relatively clean of foregrounds. Digging out larger scales will likely require extensive foreground cleaning, while smaller scales are quickly drowned by instrument noise (we find a factor of ~ 5 increase in the rms noise going to $k = 0.2 \text{ Mpc}^{-1}$). Hence detections with the first generation instruments might have only a relatively narrow window available in k space.

This paper is organized as follows. In §2.2 we describe our simulations of the cosmological signal, while in §2.3 we discuss the adopted telescope sensitivities. In §2.4, we present our results, including the detectability of the peak signal and peak S/N across our parameter space. In §2.5 we briefly consider the potential impact of foreground contamination. Finally we conclude in §2.6.

Unless stated otherwise, we quote all quantities in comoving units. We adopt the background cosmological parameters: $(\Omega_{\Lambda}, \Omega_{\text{M}}, \Omega_b, n, \sigma_8, H_0) = (0.68, 0.32, 0.049, 0.96, 0.83, 67 \text{ km s}^{-1} \text{ Mpc}^{-1})$, consistent with recent results from the Planck mission

these are now thought to be sub-dominant to X-rays (e.g. Chen and Miralda-Escudé 2004; Rybicki 2006; Furlanetto and Loeb 2004; McQuinn and O’Leary 2012). In particular, McQuinn and O’Leary (2012) perform convergence tests, using both grid and smoothed particle hydrodynamics simulations, quantifying the importance of shock heating in their Appendix A. They find that shock heating only boosts the mean temperature by a few percent at $z \gtrsim 12$, thus having a negligible impact on our conclusions below. Extreme models in which heating and reionization occurs very late ($z < 10$), could have a somewhat larger contribution from shock heating ($\sim 10\%$), and we caution the reader not to over-interpret the precise values of the signal in this, admittedly unlikely, regime.

(Planck Collaboration, 2013).

2.2 Cosmological signal

The 21cm signal is usually represented in terms of the offset of the 21cm brightness temperature from the CMB temperature, T_γ , along a line of sight (LOS) at observed frequency ν (c.f. Furlanetto et al. 2006a):

$$\begin{aligned} \delta T_b(\nu) &= \frac{T_S - T_\gamma}{1+z} (1 - e^{-\tau_{\nu_0}}) \approx \\ & 27 x_{\text{HI}} (1 + \delta_{\text{nl}}) \left(\frac{H}{dv_r/dr + H} \right) \left(1 - \frac{T_\gamma}{T_S} \right) \\ & \times \left(\frac{1+z}{10} \frac{0.15}{\Omega_M h^2} \right)^{1/2} \left(\frac{\Omega_b h^2}{0.023} \right) \text{mK}, \end{aligned} \quad (2.1)$$

where T_S is the gas spin temperature, τ_{ν_0} is the optical depth at the 21-cm frequency ν_0 , $\delta_{\text{nl}}(\mathbf{x}, z) \equiv \rho/\bar{\rho} - 1$ is the evolved (Eulerian) density contrast, $H(z)$ is the Hubble parameter, dv_r/dr is the comoving gradient of the line of sight component of the comoving velocity, and all quantities are evaluated at redshift $z = \nu_0/\nu - 1$.

To simulate the 21cm signal, we use a parallelized version of the publicly available 21CMFAST code⁶. 21CMFAST uses perturbation theory (PT) and excursion-set formalism to generate density, velocity, source, ionization, and spin temperature fields. For further details and tests of the code, interested readers are encouraged to see Mesinger and Furlanetto (2007), Zahn et al. (2011), Mesinger et al. (2011) and Mesinger et al. (2013). Here we outline our simulation set-up and the free parameters in our study.

Our simulation boxes are 600 Mpc on a side, with a resolution of 400^3 . Ionizations by UV photons are computed in an excursion-set fashion (Furlanetto et al., 2004), by comparing the local number of ionizing photons to neutral atoms. The cumulative number of ionizing photons is given by multiplying the fraction of mass collapsed in halos more massive than some threshold mass, $f_{\text{coll}}(> M_{\text{min}})$, by an ionizing efficiency

⁶<http://homepage.sns.it/mesinger/Sim.html>

which can be written as:

$$\zeta_{\text{UV}} = 30 \left(\frac{N_\gamma}{4400} \right) \left(\frac{f_{\text{esc}}}{0.1} \right) \left(\frac{f_\star}{0.1} \right) \left(\frac{1.5}{1 + \bar{n}_{\text{rec}}} \right), \quad (2.2)$$

where f_\star is the fraction of gas converted into stars, N_γ is the number of ionizing photons per stellar baryon, f_{esc} is the fraction of UV ionizing photons that escape into the IGM, and \bar{n}_{rec} is the mean number of recombinations per baryon. Here we fix the ionizing efficiency to $\zeta_{\text{UV}} = 30$, which agrees with the measured electron scattering optical depth in the fiducial model (defined below), and instead vary the X-ray luminosity of galaxies. Although there is uncertainty in the value and evolution of ζ_{UV} , by varying M_{min} we reasonably capture the redshift evolution of the reionization peak (corresponding to $\bar{x}_{\text{HI}} \sim 0.5$), which is the dominant factor in its detectability (along with the offset of the reionization and heating epochs). Hence, even though our main focus here is the 21cm peak power which is very insensitive to changes in ζ_{UV} (e.g. Mesinger et al. 2013; Christian and Loeb 2013), we expect that our range of reionization S/N estimates to also be robust.

The comoving X-ray emissivity in our models can be expressed as

$$\epsilon_{h\nu}(\nu_e, \mathbf{x}, z) = \alpha h \frac{N_X}{\mu m_p} \left(\frac{\nu_e}{\nu_0} \right)^{-\alpha} \left[\rho_{\text{crit},0} \Omega_b f_\star (1 + \bar{\delta}_{\text{nl}}) \frac{df_{\text{coll}}(> M_{\text{min}})}{dt} \right], \quad (2.3)$$

where N_X is the number of X-ray photons per stellar baryon, μm_p is the mean baryon mass, $\rho_{\text{crit},0}$ is the current critical density, f_\star is fraction of baryons converted into stars (we take $f_\star = 0.1$), $\bar{\delta}_{\text{nl}}$ is the mean non-linear overdensity. The quantity in the brackets is the comoving star formation rate density (SFRD). We assume the same M_{min} for both UV and X-ray sources. Following previous works (e.g. Pritchard and Furlanetto 2007; Santos et al. 2008; Baek et al. 2010; Mesinger et al. 2011), we take a spectral (energy) index of $\alpha = 1.5$, and assume photons below $h\nu_0 = 300$ eV are obscured, these having optical depths exceeding unity for $N_{\text{HI}} \gtrsim 10^{21.5} \text{ cm}^{-2}$, consistent with the column densities seen in high-redshift gamma-ray bursts (GRBs; Totani et al.

2006; Greiner et al. 2009). In addition to heating by X-rays, we include Compton heating, adiabatic cooling/heating, and heating through changing ionization species (Mesinger et al., 2011).

We compute the Wouthuysen-Field (WF; Wouthuysen 1952; Field 1958) coupling (i.e. Ly α pumping; when the Ly α background from the first stars couples the spin temperature to the gas temperature) using the Lyman resonance backgrounds from both X-ray excitation of HI, and direct stellar emission. The later is found to dominate by two orders of magnitude in our fiducial models. For the direct stellar emission, we assume standard Population II spectra from (Barkana and Loeb, 2005a), and sum over the Lyman resonance backgrounds (Mesinger et al., 2011). This fiducial spectrum results in $\sim 10^4$ rest-frame photons between Ly α and the Lyman limit.

Our models have two free parameters:

1. f_X – *the X-ray efficiency of galaxies*. Our fiducial choice of $f_X \equiv (N_X/0.25) = 1$ corresponds to $N_X = 0.25$ X-ray photons per stellar baryon. This choice results from a total X-ray luminosity above $h\nu_0 = 0.3$ keV of $L_{X,0.3+\text{keV}} \sim 10^{40}$ erg s $^{-1}$ (M_\odot yr $^{-1}$) $^{-1}$, using our spectral energy index of $\alpha = 1.5$.⁷ This choice is consistent with (a factor of ~ 2 higher than) an extrapolation from the 0.5–8 keV measurement of Mineo et al. (2012a), $L_{X,0.5-8\text{keV}} \approx 3 \times 10^{39}$ erg s $^{-1}$ (M_\odot yr $^{-1}$) $^{-1}$. It is highly uncertain how the X-ray luminosity evolves towards higher redshifts, although several studies argue that the higher binary fraction expected in the first galaxies results in more XRBs (e.g. Mirabel et al. 2011; Fragos et al. 2013). There is also tentative evidence from the Chandra Deep Field-South (CDF-S; e.g. Xue et al. 2011) that the X-ray luminosity to star formation rate is increasing out to $z \sim 4$ (Basu-Zych et al., 2013). Extreme evolution is limited by the 2σ upper limits from the $z \sim 6$ galaxy sample of Cowie et al. (2012),

⁷It is more common in the literature to parameterize X-ray efficiency by the ratio of the X-ray luminosity to star formation rate, generally measured for star burst galaxies. However, this number depends on the choice of bandwidth over which the X-ray luminosity is measured (e.g. Mineo et al. 2012a). We also note that many of the observationally quoted X-ray luminosities are sensitive only to energies high enough to interact little with the IGM (for example, photons with energies $\gtrsim 2$ keV have mean free paths greater than the Hubble length at $z \sim 15$, even through a neutral Universe; Baek et al. 2010; McQuinn 2012). Hence the value of f_X for even low redshift galaxies is very uncertain.

which admittedly still allow $N_X \sim 1000$, e.g. assuming a SFR of $\sim 0.1 M_\odot \text{ yr}^{-1}$, and our fiducial choice of $h\nu_0 = 0.3 \text{ keV}$ and $\alpha = 1.5$. *Below, we explore the reasonable range $10^{-3} \lesssim f_X \lesssim 10^3$.*

2. M_{min} – the minimum mass of DM halos which host star-forming galaxies. M_{min} can be expressed as:

$$\begin{aligned}
 M_{\text{min}} &= 10^8 h^{-1} \left(\frac{\mu}{0.6} \right)^{-3/2} \left(\frac{\Omega_m \Delta_c}{\Omega_m^z 18\pi^2} \right)^{-1/2} \times \\
 &\times \left(\frac{T_{\text{vir}}}{1.98 \times 10^4 \text{ K}} \right)^{3/2} \left(\frac{1+z}{10} \right)^{-3/2} M_\odot \simeq \\
 &\simeq 10^8 \left(\frac{1+z}{10} \right)^{-3/2} M_\odot
 \end{aligned} \tag{2.4}$$

where μ is the mean molecular weight, $\Omega_m^z = \Omega_m (1+z)^3 / [\Omega_m (1+z)^3 + \Omega_\Lambda]$ and $\Delta_c = 18\pi^2 + 82d - 39d^2$ with $d = \Omega_m^z - 1$.⁸ As a fiducial choice, we take $M_{\text{min}} \sim 10^8 M_\odot$, corresponding to the atomic cooling threshold at $z \sim 10$. The first galaxies were likely hosted by less massive, molecularly cooled halos, $M_{\text{min}} \sim 10^{6-7} M_\odot$ (e.g., Haiman et al. 1996a; Abel et al. 2002; Bromm et al. 2002). However, star formation inside such small halos was likely inefficient (with a handful of stars per halo), and was eventually suppressed by the heating from X-rays themselves or other feedback processes (Haiman et al., 2000; Ricotti et al., 2001; Mesinger et al., 2006; Haiman and Bryan, 2006). M_{min}

⁸Another common approach is to argue that efficient cooling of gas at a redshift-independent temperature sets the threshold for hosting star forming galaxies. This motivates using a fixed halo virial temperature, T_{vir} , as a fundamental parameter, effectively introducing a redshift dependence to M_{min} according to eq. 2.4. However, using a fixed M_{min} facilitates a more straightforward mapping to a particle mass in popular warm dark matter cosmologies, as we shall see below. In any case, both a fixed M_{min} or fixed T_{vir} are oversimplifications, since feedback physics, either by SNe (e.g. Springel and Hernquist 2003) or the X-ray and UV backgrounds themselves (e.g. Ricotti and Ostriker 2004; Kuhlen and Madau 2005; Mesinger et al. 2013), will likely govern the redshift evolution of M_{min} . We are mostly interested in the value of M_{min} during the X-ray heating phase. It is unlikely there is dramatic evolution of M_{min} during this relatively rapid epoch (compare for example the fiducial model in Mesinger et al. 2013 which uses a fixed T_{vir} , with the one in the top panels of Fig. 2-3, which uses a fixed M_{min}). We note however that during the very early stages, a fixed M_{min} model shows a more rapid evolution impacting the depth of the mean absorption trough (bottom panel of Fig. 2-3) compared with a fixed T_{vir} model.

could also have been larger than the atomic cooling threshold due to feedback processes (e.g. Springel and Hernquist 2003; Okamoto et al. 2008; Pawlik and Schaye 2009; Sobacchi and Mesinger 2013b,a). It is unlikely that M_{\min} was larger than $\sim 10^{10-11}$ since these values approximately latch onto the steeply-rising faint end of the observed galaxy luminosity functions at $z \sim 6-8$ (e.g., Salvaterra et al. 2011; Finlator et al. 2011). Furthermore, it would be difficult to complete reionization by $z \sim 5-6$ without a contribution from galaxies hosted by smaller halos (e.g. Kuhlen and Faucher-Giguere 2012; Choudhury et al. 2008). *Below, we explore the reasonable range $10^7 \lesssim M_{\min}/M_{\odot} \lesssim 10^{10}$.*

For the purposes of this work, it is useful to keep in mind that increasing f_X has the effect of shifting the X-ray heating epoch (and associated peak in power) towards higher redshifts, while increasing M_{\min} has the effect of shifting *all* astrophysical epochs towards lower redshifts. Furthermore, increasing M_{\min} has the additional impact of speeding up cosmic evolution, as structures form more rapidly on the high-mass end of the mass function.

2.2.1 Warm Dark Matter Models

Our framework also allows us to estimate the 21cm signal in warm dark matter (WDM) cosmologies. WDM models with particle masses of order $m_{\text{wdm}} \sim \text{keV}$ became popular as a cosmological way of alleviating small-scale problems of CDM, such as a dearth of locally observed dwarf galaxies and flattened rotation curves in galaxy centers (e.g. Moore et al. 1999b). Current measurements place limits of $m_{\text{wdm}} \gtrsim 1-3$ keV (Barkana and Loeb, 2001; de Souza et al., 2013; Kang et al., 2013; Pacucci et al., 2013; Viel et al., 2013), with various degrees of astrophysical degeneracy. Due to the hierarchal nature of structure formation, the impact of WDM (or any model with a dearth of small-scale power) is larger at higher redshifts, with the Universe becoming increasingly empty. Therefore, a detection of the pre-reionization 21cm signal could strengthen limits on m_{wdm} (Sitwell et al., 2014).

Structure formation in WDM models is suppressed through (i) particle free-streaming, and (ii) residual velocity dispersion of the particles. Effect (i) can be included by suppressing the standard matter transfer function below the free-streaming scale (e.g. Bode et al. 2001), while effect (ii) acts as an effective pressure, slowing the early growth of perturbations (e.g. Barkana and Loeb 2001). Effect (ii) is generally ignored as it is difficult to include in N -body simulations, since it translates to an *intra*-particle dispersion in the codes. By an analogy to a baryonic Jeans mass, Barkana and Loeb (2001) derived a critical WDM halo mass for thermal relics, below which structure formation is suppressed due to the particle velocity dispersion. de Souza et al. (2013) empirically found (see their Fig. 1) that a step-function suppression of halos smaller than ~ 60 times this critical Jeans mass, results in collapse fractions which are very close to the full random walk procedure of Barkana and Loeb (2001). A step-function suppression allows us to relate the WDM particle mass, m_{wdm} to an “effective” M_{min} :

$$M_{\text{min,eff}} \approx 2 \times 10^{10} M_{\odot} \left(\frac{\Omega_{\text{wdm}} h^2}{0.15} \right)^{1/2} \left(\frac{m_{\text{wdm}}}{1 \text{keV}} \right)^{-4} \left(\frac{1+z_I}{3500} \right)^{3/2}, \quad (2.5)$$

where z_I corresponds to the redshift of matter-radiation equality. With this casting we can present our results in terms of m_{wdm} as well as M_{min} , without running additional dedicated simulations. There is however an important caveat: our simulations use a standard CDM transfer function (Eisenstein and Hut, 1998), without the WDM cut-off. The empirical calibration in eq. 2.5 included the proper WDM transfer function (Bode et al., 2001). However, due to the sharpness of the barrier transition and strength of the effective pressure effect, it is not unreasonable to assume that most of the small-mass suppression is included, even with the CDM transfer function. We check the accuracy of this mapping below (see Fig. 2-3) for fiducial astrophysical parameters. We find that our simple prescription which underestimates the suppression by using the CDM power spectrum shifts the evolution of the 21cm signal to higher redshifts by a modest $\Delta z \approx 1$, with the peak power relatively unaffected (compare the $m_{\text{wdm}} = 2\text{keV}$ magenta curve, with the $M_{\text{min}} = M_{\text{min,eff}} = 10^9 M_{\odot}$ red curve; the




Figure 2-1: Poisson (cosmic variance) component of the S/N, i.e. $\sqrt{N_k}$, for our fiducial observational strategy and $k \approx 0.1 \text{ Mpc}^{-1}$.

former includes the WDM transfer function of Bode et al. 2001). Therefore, we find this simple mapping of eq.(2.5) reasonable, and include the corresponding m_{wdm} values on the right vertical axes of our plots. However, we caution that this conversion should only be treated as approximate.

Furthermore, we stress that it is the *maximum* value of M_{min} , set by *either* cosmology *or* cooling physics, which regulates galaxy formation. Hence if the gas cannot efficiently cool to form stars in halos down to the $M_{\text{min,eff}}$ threshold from eq. (2.5), we cannot use this signal to probe the WDM particle mass.

2.3 Instrument sensitivity

Throughout this work, we use as the observable quantity the spherically averaged power spectrum, $P_{21} \equiv k^3/(2\pi^2V) \delta\bar{T}_b(z)^2 \langle |\delta_{21}(\mathbf{k}, z)|^2 \rangle_k$, where $\delta_{21}(\mathbf{x}, z) \equiv \delta T_b(\mathbf{x}, z)/\delta\bar{T}_b(z) - 1$. Furthermore, we focus on the large-scale signal at $k = 0.1 \text{ Mpc}^{-1}$, lying in the “sweet spot” of 21cm interferometry: large enough for the cosmic signal not to be removed in the foreground cleaning process, yet small enough to have high S/N with upcoming instruments (e.g. Lidz et al. 2008; Dillon et al. 2013; Pober et al. 2013a). Our default power spectrum bin width is $d \ln k = 0.5$.

2.3.1 Calculation of Thermal Noise

To compute the thermal noise variance for each array, we perform rotation synthesis of 6 hours per night about zenith on a uv-plane whose resolution is set by the inverse of the array’s primary beam FWHM. With the flat sky approximation, and assuming

small baselines and bandwidth, the amplitude of the noise power spectrum in each uv-cell is given by (e.g. Morales 2005; Parsons et al. 2012a):

$$P_N \approx D_M(z)^2 Y \frac{k^3}{2\pi^2} \frac{\Omega'}{2t} T_{sys}^2, \quad (2.6)$$

where Ω' is a beam dependent factor described in Parsons et al. (2014), $D_M(z)$ is a constant factor that converts between a transverse angle on the sky and comoving distance units, Y is a factor that converts between frequency and radial comoving distance, and t is the total time spent by all baselines in the uv cell during the aperture synthesis. T_{sys} is the system temperature which is the sum of T_{rec} , the receiver noise temperature and T_{sky} , the sky temperature. For T_{sky} we use the measurement of Rogers and Bowman (2008) of $T_{sky} = 237 \left(\frac{\nu}{150\text{MHz}}\right)^{-2.5} \text{K}^9$. We set the receiver temperature to $T_{rec} = 50\text{K} + 0.1T_{sky}$ (Dewdney et al., 2013)¹⁰.

Chromatic effects due to the dependence of an array's uv coverage with frequency are approximated by averaging the coverage over the data cube and using the T_{sys} at the data cube's center frequency. While mean noise power can be removed from the data by computing power spectra estimates, \hat{P}_k , from data interleaved in time (Dillon et al., 2013), the variance of the thermal noise power spectrum is expected to be the leading contribution to measurement uncertainty within the "EoR window". We assume that a power spectrum estimate is computed by taking an inverse variance weighted average of all u-v cells within a k-bin. In performing inverse variance weighting, the variance of a \hat{P}_k is $\sigma_N^2 = 1/\sum_j \sigma_j^{-2}$.

To account for foregrounds, we exclude all u-v cells lying within the "wedge": the region of k-space contaminated by foregrounds which are thrown to larger k_{\parallel} modes by the chromaticity of the interferometer. The maximum k_{\parallel} contaminated by this

⁹We choose the minimal value measured at the galactic pole, where presumably power spectrum studies would focus. Measurements at hotter galactic latitudes could reduce our S/N estimates by a factor of 2-4.

¹⁰While receiver temperature can vary from instrument to instrument, it is likely dominated by T_{sky} .

mechanism, at a fixed k_{\perp} is given by (Morales et al., 2012; Parsons et al., 2012a):

$$k_{\parallel}^{max} = \sin \Theta_{max} \left(\frac{D_M(z)E(z)}{D_H(1+z)} \right) k_{\perp} + k_{intr} \quad (2.7)$$

where D_H is the Hubble distance, $E(z) = H(z)/H_0$, and Θ_{max} is the maximum angle on the sky from which foregrounds enter the beam. k_{intr} is an offset to account for the intrinsic ‘‘spectral unsmoothness’’ of the foregrounds. We take $k_{intr} = 0.02 \text{ Mpc}^{-1}$ and Θ_{max} equal to one half the FWHM of the primary beam. While this allows for possibly larger S/N at even larger scales ($k \sim 0.04 \text{ Mpc}^{-1}$; e.g. Beardsley et al. 2013), it is very uncertain whether this region will be clean of foregrounds. Hence we choose to compromise, working at what is more likely to be a foreground free scale, $k = 0.1 \text{ Mpc}^{-1}$ (e.g. Pober et al. 2013a). In §2.5 we examine sensitivity at smaller spatial scales with an even larger k_{intr} .

Although 8 MHz bands are a common choice in the literature, this bandwidth is large enough to average over signal evolution at very high redshifts. We therefore consider a band corresponding to $\Delta z = 0.5$ for all redshifts. In fixing the redshift span, we impose a minimal k_{\parallel} resolution sampled by the instrument which becomes dramatically worse at high z . At $z \gtrsim 20$, this minimal k_{\parallel} becomes so large that our instruments do not sample any modes at $k = 0.1 \text{ Mpc}^{-1}$ (the sharp drop in the number of modes at $z > 20$ is evident in Fig. 2-1).¹¹

We also include the Poisson noise of the cosmic signal, $P_{21}/\sqrt{N_k}$, where N_k is the instrument-dependent number of modes in our power spectrum bin. This Poisson term, corresponding to the maximum achievable S/N (i.e. if thermal noise is zero), is plotted in Fig. 2-1. Due to the large values of $N_k \sim 10^3 - 10^4$, the Poisson (cosmic variance) noise only dominates high-sigma detections, generally achievable only with the SKA. We note that our observational strategy is chosen to minimize the thermal noise. A different strategy (sampling more independent fields) would lower

¹¹Although necessary for sensitivity estimates, a bandwidth choice is relatively arbitrary. Hence we do not additionally smooth the cosmic signal, showing its intrinsic value. As already mentioned, our fiducial choice is motivated by the negligible evolution of the signal over the bandwidth, with the X-ray heating power near the peak evolving by only a few percent. Nevertheless, we highlight that a wider bandwidth could extend SKA sensitivities to $z > 20$.




Figure 2-2: Time spent in each uv cell for a 2000h observation at $z = 10$ for PAPER, MWA128T, LOFAR, HERA, SKA, (left to right).

the Poisson noise and increase the thermal noise; such a strategy can be employed to avoid the cosmic variance limited regime of SKA, if extremely high S/N is desired. For simplicity, we do not perform inverse weighing on the Poisson noise, noting that this has a minor impact (on very high S/N values). Therefore, adding the noise terms in quadrature, our total S/N can be expressed as:

$$S/N = \frac{P_{21}}{\sqrt{P_{21}^2/N_k + \sigma_N^2}}. \quad (2.8)$$

2.3.2 Array Models

In this work, we consider four existent and planned arrays: MWA (Tingay et al., 2013c), LOFAR (van Haarlem et al., 2013), PAPER (Parsons et al., 2010), and the SKA (Dewdney et al., 2013). Here we briefly describe the model for each instrument.

MWA128T/256T: Our model of the MWA contains 128 tiles whose locations are given in Beardsley et al. (2012). Each antenna element is modeled by a 4x4 grid of short dipoles 0.3 m above an infinite conducting plane and spaced 1.1 m apart. The primary beam FWHM is computed from the radiation pattern of this arrangement. The bandpass of the MWA cuts off at ≈ 75 MHz so we set $P_N \rightarrow \infty$ below this frequency. A frequency resolution of 40 kHz is used. We also consider a possible extension to 256 tiles (MWA256T), for which the infrastructure is already in place. The additional tiles are placed randomly (with a 5 meter minimum separation), drawing from a uniform distribution within 50

meters of the center and then $1/r^2$ for $r > 50$ meters.

PAPER: We use the maximally redundant configuration of PAPER 128 with a band pass ranging from 100 to 200 MHz and a frequency resolution of 48 kHz. The primary beam is fixed to be 0.72 sr over all redshifts along with $\Omega' = 1.69sr$ (Parsons et al., 2014). PAPER is a drift scan instrument, hence it observes a larger number of fields with a shorter integration time per field. We assume that 2000 hours of observation (an optimistically high choice corresponding to $\sim 2-3$ observing seasons) are spread over observations of 3 different subfields with 2 hours of integration per field per night, increasing the Poisson counts significantly (Fig. 2-1).

LOFAR: LOFAR is comprised of two different sub-arrays. We use the high and low band antenna locations in the 40 core stations along with the beams and effective areas described in van Haarlem et al. (2013). For the low band antennas we assume the “inner configuration”, interpolating values for the beam FWHM and effective areas between those given in Table B.1 in van Haarlem et al. (2013). We treat the region between the high and low bandpasses (80-110 MHz) as unobservable. A frequency resolution of 10 kHz is used.

SKA: Antenna locations for our SKA model are based on the SKA Low Phase 1 design described in Dewdney et al. (2013). 866 station locations are drawn from a Gaussian distribution with 75 % falling within 1000m of the center. Each station is modeled as a 17x17 array of log-periodic dipoles whose effective areas and beams are given in Table 3 and Appendix A of Dewdney et al. (2013). The frequency resolution is 1 kHz.

In addition to the above fiducial instruments, we briefly present preliminary noise estimates from the proposed, second-generation Hydrogen Epoch of Reionization Arrays (HERA; <http://reionization.org/>). HERA is a proposed instrument comprised of 547 antennas with a hexagonal packing configuration (Pober et al., in prep). Each of the static dish antennae is modeled as a 14 meter filled aperture. Because HERA




Figure 2-3: *Top panels:* Amplitude of the 21cm power at $k = 0.1 \text{ Mpc}^{-1}$ in various models. We also plot the (1σ) sensitivity curves corresponding to a 2000h observation with: MWA128T, LOFAR, SKA on the left, and MWA256T, PAPER, and the proposed HERA instrument on the right. The recent upper limit from Parsons et al. (2014) is shown at $z = 7.7$. *Bottom panels:* The corresponding average 21cm brightness temperature offset from the CMB.

would be a drift scan instrument with a narrow field of view, the observation strategy we adopt is to observe, on each night, 9 different ~ 10 degree fields for 45 minutes each for a total of 6 hours per night. We adopt a frequency resolution of 98 kHz. The configuration is optimized for statistical detections.

It is illustrative to compare the uv coverage of each array, especially within the compact core sourcing the sensitivity for the cosmological signal. We show in Figure 2-2, the time spent in each uv cell at $z = 10$ for PAPER, MWA128T, LOFAR, HERA, SKA, (*left to right*), over the course of 2000 hours of observing. Similar plots for each individual array can be found in Beardsley et al. (2013) (MWA), van Haarlem et al. (2013) (LOFAR) and Dewdney et al. (2013) (SKA). PAPER and HERA, which are optimized for EoR measurements show compact uv distributions, while multi-purpose instruments like MWA, LOFAR, and the SKA have broad uv coverage. The MWA has a very wide field of view, resulting in fine pixelization. LOFAR, HERA, and SKA have relatively narrow primary beams, translating into a more coarse resolution in the uv plane.

2.4 Results

2.4.1 Physical insight into the signal

In the top panels of Fig. 2-3, we plot the evolution of the $k = 0.1 \text{ Mpc}^{-1}$ 21cm power in various models. The black solid curves corresponds to a “fiducial model”, with $f_X = 1$ and $M_{\text{min}} = 10^8 M_{\odot}$ (approximately the atomic cooling threshold at these redshifts).

Note that most models exhibit the familiar three peak structure (e.g. Pritchard and Furlanetto 2007; Mesinger et al. 2011; Baek et al. 2010). In order of decreasing redshift, these peaks correspond to the three astrophysical (radiation-driven) epochs of the 21cm signal: (i) WF coupling; (ii) X-ray heating; and (iii) reionization.

Extreme models can avoid having a three-peaked structure by merging the reionization and X-ray heating peaks. In these cases, the X-ray background is faint enough ($f_X \lesssim 10^{-2}$; McQuinn and O’Leary 2012; Christian and Loeb 2013), that it is unable to heat the IGM prior to the completion of reionization. The resulting contrast between cold neutral and ionized regions can drive up the reionization peak considerably. We show one such model with the blue curves in Fig. 2-3.

We also overlay the sensitivity curves of MWA128T¹², LOFAR and SKA (*left panel*), as well as MWA256T, PAPER and the proposed HERA instrument (*right panel*). The first generation instruments will have difficulty detecting the X-ray heating peak in the fiducial model. As the X-ray efficiency is decreased, or galaxies are hosted by more massive later-appearing halos, the heating peak moves to lower redshifts, making it more easily observable with the MWA (LOFAR unfortunately has a band gap in this regime, and the PAPER bandpass cuts-off beyond $z \approx 13$). The SKA, within its bandpass, easily has the required sensitivity to detect all reasonable models, and is in fact limited by cosmic variance for our observational strategy, as we

¹²Our MWA128T noise estimates are consistent with those in (Beardsley et al., 2013) at $z = 8$, when accounting for their different choice of system temperature. However, our noise curves are a factor of ~ 10 higher than the ones in Christian and Loeb (2013). This is due primarily to the fact that they neglect to evolve the system temperature, using the $z = 8$ value at all redshifts. As the system temperature is expected to scale as $\propto (1+z)^{2.5}$, neglecting its evolution results in a dramatic underestimate of the noise at high- z




Figure 2-4: Various quantities evaluated at the redshift where the $k = 0.1 \text{ Mpc}^{-1}$ power is the largest in each model. Shown are power amplitude, redshift, mean neutral fraction, and mean brightness temperature, clockwise from top left. Overlaid are the Thomson scattering optical depth, τ_e , contours corresponding to 1 and 2 σ constraints from the 9yr release of WMAP data (WMAP9; Hinshaw et al. 2013). The right-side y-axis shows the corresponding values of the WDM particle mass, m_{wdm} , computed according to the approximation in eq. (2.5). Regions of the parameter space in which the peak power occurs during the reionization or X-ray heating epochs are demarcated with the appropriate labels in some panels.

shall see below.¹³

In Fig. 2-4, we plot the power amplitude, redshift, mean neutral fraction, and mean brightness temperature (clockwise from top left), all evaluated at the redshift where the $k = 0.1 \text{ Mpc}^{-1}$ power is the largest in each model.

The peak power in our parameter space spans the range $30 \lesssim P_{21}/\text{mK}^2 \sim 5000$. However, for a large-swath of reasonable models ($10^{-2} \lesssim f_X \lesssim 10^2$), the peak power is remarkably constant at a few hundred mK^2 . For most models, the power peaks

¹³We caution that the effective bandpass for first generation instruments is limited by data transfer and processing, and will be narrower than the full range shown in Fig. 2-3. For example, MWA has an effective bandpass of 32MHz, allowing simultaneous observations over, e.g. $z = 8-10$ and $z = 12-17$. Below, for simplicity and in order to be conservative, we only present S/N estimates in a single $\Delta z = 0.5$ frequency bin. With wider frequency coverage, the S/N could be increased by a factor of a few by averaging over the effective bandpass.

when the mean brightness temperature is ~ -100 mK. The fluctuations in T_γ/T_S are maximized at this time, when areas surrounding sources have been heated to above the CMB temperature, and yet the bulk of the IGM is still cold, seen in absorption. This can be seen in Fig. 2-5, where we plot the cumulative distribution functions (CDFs) of T_γ/T_S at three redshifts spanning the X-ray peak of the fiducial model. When the amplitude of the power is largest, only a few percent of the IGM is seen in emission (middle curve in Fig. 2-5). Shortly afterwards, the distribution of T_γ/T_S piles up around zero, and the temperature fluctuations cease being important. Hence, this process is self-similar for many models, especially given that the bias of the halos, which would impact the temperature fluctuations, does not evolve dramatically over our chosen mass range when compared at the same astrophysical epoch¹⁴ (e.g. McQuinn et al. 2007). Therefore, the lack of notable change in the X-ray peak height is understandable. However, we caution that the precise peak height and power spectrum evolution is likely affected by the spectral energy distribution of X-ray sources; we postpone an investigation of this to future work.

Understandably, the strongest trend in peak power is with f_X ; however both extrema are found in models with high values of M_{\min} . This is because the growth of structures is both delayed and more rapid in models with higher M_{\min} ; we elaborate more on these trends below.

The weakest signal is found in models with a high f_X . In these models, X-ray heating starts at a high redshift, closely following the onset of Ly α pumping. As a result, the spin temperature does not have time to couple strongly to the kinetic temperature, before the later is driven up by X-ray heating (see the damping of the global absorption trough of the green curve in the bottom panel of Fig. 2-3). This is clearly evident in lower left panel of Fig. 2-4: models with high X-ray emissivities have much higher (less negative) values of the mean brightness temperature, $\delta\bar{T}_b$, with the peak power in very high $f_X \sim 10^3$ models occurring during the “emission” regime

¹⁴For fixed astrophysical parameters, such as f_X and f_{UV} , a higher value of M_{\min} delays the milestones in the signal. Hence the same astrophysical epoch, such as reionization or X-ray heating, corresponds to a lower redshift when the halo mass function has already evolved. Instead, when comparing models at the same redshift (and analogously the same mass function), the imprint of the different halo bias resulting from different choices of M_{\min} values is more notable.

(when $\bar{\delta T}_b$ is positive). This overlap of WF coupling and X-ray heating increases with increasing M_{\min} (or analogously decreasing m_{wdm}), due to the rapid growth of structures on the high-mass tail of the mass function. Although not the focus of this study, we note that weaker Ly α pumping could also result in an overlap of these two epochs. In particular, if the direct stellar emission was 100 times weaker than assumed, then the 21cm peak power would be decreased by a factor of few for even the fiducial value of $f_X \sim 1$.

The strongest signal on the other hand, is found in models with a low f_X . In these models, heating (or reionization) starts later when the δT_b contrast is larger due to the evolution of the ratio T_γ/T_S (eq. 2.1; note that due to the expansion of the Universe, T_γ/T_S evolves roughly as $\propto (1+z)^{-1}$ prior to heating). Again, the signal is even stronger at high values of M_{\min} (or analogously low values of m_{wdm}), which further shift the evolution to lower redshifts.

The allowed peak amplitude decreases somewhat, when considering only models which are within the 2σ constraints on τ_e from WMAP9 (demarcated with the solid black curves)¹⁵. Limits on τ_e generally rule out late reionization models, which for our choice of ζ_{UV} correspond to high values of $M_{\min} \gtrsim 10^{9.5} M_\odot$. However, at efficiencies $f_X \gtrsim 10$, X-rays begin to contribute to reionization at the $\gtrsim 10\%$ level (c.f. McQuinn and O’Leary 2012; Mesinger et al. 2013), and the WMAP9 τ_e isocontours curve upward, including high M_{\min} models.

We can also see from the lower right panel in Fig. 2-4 that for $f_X \gtrsim 10^{-2}$ the Universe is mostly neutral when the 21cm power peaks. This means that the peak power indeed occurs during the X-ray heating epoch, before reionization. In contrast, the reionization peak in power (c.f. top panel of Fig. 2-3), should occur at $\bar{x}_{\text{HI}} \sim 0.5$

¹⁵We compute τ_e from the evolutions of the average neutral ionized fraction, $\langle x_i \rangle$, and density, $\langle n \rangle$. Strictly speaking, the correlations between these two fields should be taken into account, i.e. $\langle x_i \times n \rangle \neq \langle x_i \rangle \times \langle n \rangle$. Mesinger et al. (2013) note that the fact that reionization is “inside-out” on large scales results in a slightly higher value of τ_e , than estimated ignoring correlations. We account for this bias by multiplying our τ_e estimates by 1.04, the size of the bias in fiducial, UV driven reionization scenarios. We further assume that reionization has completed by our last redshift output at $z = 7$ (i.e. imposing $\bar{x}_{\text{HI}} = 0$ at $z < 7$). This means that τ_e is overestimated somewhat for models with high $M_{\min} \gtrsim 10^9 M_\odot$ which do not complete reionization by this redshift. Hence the upper 2σ contour is conservatively broad.




Figure 2-5: Cumulative distributions of T_γ/T_S (see eq. 2.1), at $z = 14, 16, 18$ (spanning the X-ray peak) for a fiducial, $T_{\text{vir}} = 10^4$ K, $f_X = 1$ model. As seen above, the power peaks when the distribution of T_γ/T_S is the broadest.

(e.g. Lidz et al. 2008; Friedrich et al. 2011; Mesinger et al. 2013). In our models, this occurs for values of $f_X \lesssim 10^{-2}$, consistent with simple analytic estimates (McQuinn and O’Leary, 2012). In these models, the X-ray background is too weak to heat the IGM before reionization. The resulting contrast between the (very) cold neutral and ionized patches drives the reionization power to values of $\gtrsim 10^3$ mK² (c.f. Parsons et al. 2014).

2.4.2 Detectability of the peak power

We now include our 2000h sensitivity estimates from §2.3, in order to predict the detectability of the X-ray heating peak with current and future interferometers. In Fig. 2-6, we plot the S/N according to eq. (2.8), with which the peak power can be detected. Regions in white correspond to areas of parameter space where the signal to noise is less than unity.

The isocontours of S/N generally follow the diagonal trend of the redshift isocontours from the top right panel in Fig. 2-4. This is due to the fact the sensitivities of the interferometers (noise) vary more strongly with redshift than does the amplitude of the peak power (signal). The exception to this trend is the strip at $f_X \lesssim 10^{-2}$, corresponding to reionization in a cold IGM, when the power can jump to $P_{21} \gtrsim 10^3$ mK². This “cold-reionization peak” is strongly detectable by all instruments. Therefore, it should be noted that *there is no such thing as an “X-ray heating peak” for*




Figure 2-6: S/N of the detection of the $k = 0.1 \text{ Mpc}^{-1}$ 21cm peak power (i.e. computed at the redshift of the maximum signal). All maps show the same range in S/N to highlight differences between instruments, and are computed assuming a total integration time of 2000h. Due to our fiducial observing strategy which minimizes thermal noise at the expense of cosmic variance, some high S/N regions ($S/N \gtrsim 50$) are limited by cosmic (Poisson) variance for the case of LOFAR and SKA, which have smaller beams than MWA and PAPER (see Fig. 2-1). This Poisson variance limit can be avoided with a different observing strategy; hence we caution the reader not to lend weight to the apparent better performance of MWA and PAPER in the upper left region of parameter space.

$$f_X \lesssim 10^{-2}.$$

The peak power of a fiducial model ($M_{\min} \approx 10^8 M_{\odot}$, $f_X \approx 1$) lies at the edge of detectability for MWA128T. However, if X-ray heating is delayed, either by lowering f_X or increasing M_{\min} by a factor of ~ 10 , 128T MWA can detect the X-ray peak at the $S/N \sim \text{few}-10$ level.

The bandpass coverage means that the parameter space region in which S/N is greater than unity (at peak signal) does not evolve much when upgrading the MWA to 256 tiles. However, the S/N of a detection increases dramatically for MWA256, with the X-ray peak being detectable at the $S/N \gtrsim 10$ level throughout the bandpass.

In the case of LOFAR, we find that the peak signal is detectable at $S/N > 1$ only with the high-band antennas (we keep the blank LOFAR low panel in Fig. 2-6 for the sake of symmetry). Only late heating models are therefore detectable, corresponding to the high- M_{\min} , low- f_X corner¹⁶. If one discounts the “cold-reionization” $f_X \lesssim 10^{-2}$ strip, as well as the low- τ_e region, it is apparent that LOFAR is unlikely to detect the peak power during X-ray heating. However, even when the actual peak signal lies beyond the band pass, it is possible that the X-ray heating power might still extend into the LOFAR low bandpass, allowing for low S/N detections (see for example the fiducial model in Fig. 2-3 in which the power peaks outside the bandpass at $z \sim 16$, but is still detectable at $z \sim 17$). We will quantify this below when we plot the peak S/N for LOFAR low.

The results for PAPER are quite similar to those of LOFAR. However, the slightly wider PAPER bandpass extending to lower frequencies allows for a slightly wider strip of X-ray heating detections in Fig. 2-6.

Finally, it is evident from the bottom right panel of Fig. 2-6 that the second generation interferometer, SKA, will be a huge improvement over the first generation instruments. SKA should be able to detect all X-ray heating models we consider, with the exception of the lower right corner in which the peak signal extends beyond our fiducial coverage, $z \gtrsim 20$ (we remind the reader that this is not a fundamental

¹⁶The lower S/N of these detections compared with the analogous ones by the MWA results from the Poisson (cosmic variance) limit (see Fig. 2-1). LOFAR and SKA have a narrower beam than the MWA.




Figure 2-7: Maximum S/N possible with current and upcoming interferometers after 2000h (considering all redshifts instead of just the peak signal as in Fig. 2-6).

limitation of SKA, and can be avoided with a wider bandwidth choice). In fact, all of the SKA detections are Poisson noise limited (c.f. Fig. 2-1). Our observational strategy was chosen to minimize thermal noise at the expense of Poisson noise. If however even stronger detections are desired with the SKA, one could observe more fields for a shorter period of time (however the benefits of a $S/N \gtrsim 500$ detection compared with a $S/N \gtrsim 50$ one, are not immediately obvious).

2.4.3 General detectability of the cosmic 21cm signal

Thus far we have focused on the maximum amplitude (over all redshifts) of the 21cm power spectrum at $k = 0.1 \text{ Mpc}^{-1}$. This peak power generally corresponds to the epoch of X-ray heating ($f_X \gtrsim 10^{-2}$), or “cold-reionization” ($f_X \lesssim 10^{-2}$). We now ask instead how detectable is the cosmic 21cm signal, *regardless of the epoch*. Due to the increase in instrument noise towards higher redshifts, in many models the reionization peak is more detectable than the X-ray one, even though the cosmic signal is weaker

(see Fig. 2-3).

In Fig. 2-7, we show S/N plots as in Fig. 2-6, but computed at the maximum S/N (instead of the maximum signal). Understandably, the detectable region of parameter space broadens for all instruments, as models in which the peak signal fell out of the observable bands are now again considered. This broadening of detectable parameter space is most dramatic for LOFAR and PAPER, whose band passes are optimized for reionization.

As mentioned above, we see that LOFAR low could catch the pre-reionization signal (albeit at $S/N \sim$ unity), even though it could not detect the actual peak of the signal. Indeed our “fiducial model” (c.f. the black curve in the left panel of Fig. 2-3) falls into this category, with the *pre*-peak power extending into the LOFAR-low detection region. It is also interesting to note that the upper strip in the LOFAR low panel of Fig. 2-7 corresponds to the earlier Ly α pumping peak, before X-ray heating, which is also detectable at $S/N \sim$ unity. This strip includes models in which the X-ray heating is sufficiently delayed to allow the Ly α pumping epoch to extend into the LOFAR-low detection region (c.f. the blue curve in the left panel of Fig. 2-3).

In Fig. 2-8, we show the evolution of the S/N for the “fiducial” model: $M_{\min} = 10^8 M_{\odot}$, $f_X = 1$. We can see that reionization is detectable at $\approx 4\sigma$ with MWA128T and PAPER, and at $\approx 20\text{--}30\sigma$ with LOFAR and the potential MWA256T. MWA128T and LOFAR might be able to detect X-ray heating in the fiducial model at $\approx 1\text{--}2\sigma$. The SKA is cosmic variance limited throughout.

It should be noted that in this work, we compute the S/N from a single k -bin and frequency (i.e. redshift) bin. In principle, one can boost the S/N of the detection by a weighted sum over both the available k -modes as well the available frequency bins. Depending on the foreground smoothness, the first generation instruments might only have a narrow window in k -space to make the measurement. However, the signal could still have a $S/N \gtrsim 1$ over several frequency (redshift) bins (see Fig. 2-8). Hence, by summing over these bins, the S/N of the detection could be boosted by an additional factor of \sim few, provided there is contiguous frequency coverage surrounding the




Figure 2-8: The S/N vs redshift evolution for the “fiducial” model, with $M_{\min} = 10^8 M_{\odot}$, $f_X = 1$.

relevant epochs.

It is important to also note that for continuous and relatively wide band pass interferometers like the MWA, the change between the S/N computed at max signal vs that computed at max S/N (Fig. 2-6 vs Fig. 2-7) is not very dramatic: mostly low S/N detections extend into the parameter space where the peak X-ray heating signal occurs too early to make MWA’s $z \approx 18$ band cut. We can also see this explicitly for the fiducial model in Fig. 2-8: the S/N of X-ray heating is only a factor of ≈ 2 less than that of reionization for MWA128T. This is suggestive that for these instruments, the detectability of reionization and X-ray heating are roughly comparable from a thermal noise perspective¹⁷. We explore this in the following section.

2.4.4 Reionization or X-ray heating?

We can now ask the question: “does the maximum S/N correspond to the reionization or the X-ray heating epoch?” Aside from the bandpass limitations of LOFAR and PAPER, this distinction is not so clear.

To further quantify this, in Fig. 2-9 we plot the 21cm power spectrum amplitude (*left*), $\delta\bar{T}_b$ (*center*), and the neutral fraction (*right*), all computed at the redshift of peak S/N, assuming MWA128T sensitivities. Although we do not include the plots, we note that the trends for continuous frequency coverage instruments like the SKA are the same, in the regime when the detections are not cosmic variance limited

¹⁷Additional observational challenges associated with the lower frequencies of the X-ray heating epoch such as calibration and RFI are not considered in this work.




Figure 2-9: The 21cm power spectrum amplitude at $k = 0.1 \text{ Mpc}^{-1}$ (left), $\delta\bar{T}_b$ (center), and \bar{x}_{HI} (right), all evaluated at the redshift when S/N is the largest, assuming MWA128 sensitivities.

(i.e. excluding the Poisson contribution to the noise). From Fig. 2-9, we see that there is a clear separation between the regimes where the peak S/N is achieved during reionization vs during X-ray heating.

As already noted, during the reionization epoch (more precisely, when the power spectrum of the ionization field is dominating the total 21cm signal), the large-scale 21cm power peaks at the mid-point of reionization (e.g. Lidz et al. 2008; Friedrich et al. 2011), except for extreme models with $f_X > 10^3$ (Mesinger et al., 2013). The peak of $k = 0.1 \text{ Mpc}^{-1}$ power during reionization is roughly $\sim 10 \text{ mK}^2$, more than an order of magnitude less than the peak during the X-ray heating epoch. From Fig. 2-9, we see that models with $f_X \gtrsim 1$ indeed follow these reionization epoch predictions¹⁸.

Instead for $10^{-1.5} \lesssim f_X \lesssim 10^{-0.5}$, the Universe is mostly neutral at the highest S/N, with peak power amplitudes of $\gtrsim 100 \text{ mK}^2$. This corresponds to the epoch of X-ray heating, signifying that for $10^{-1.5} \lesssim f_X \lesssim 10^{-0.5}$ the X-ray peak becomes more detectable than the reionization peak. This is understandable by considering the positions of the reionization and X-ray heating peaks. We can already see from

¹⁸If X-rays contribute significantly to reionization (e.g. very high values of f_X), then the large-scale power during the advanced stages of reionization can peak at lower neutral fractions (later stages). This is due to the suppression of ionization power from X-rays, whose mean free paths result in a smoother reionization (Mesinger et al., 2013). As reionization progresses however, a smaller fraction of the photon energy goes into ionizations, and the relative contribution of UV photons increases.

Fig. 2-3, that for a fiducial model ($f_X \sim 1$), the increase in the interferometer noise between reionization and X-ray heating epochs is almost (not quite) compensated by the increase in the signal. As the X-ray efficiency is lowered, the heating epoch moves to lower redshifts and becomes more detectable than the reionization epoch.

As the X-ray efficiency is lowered even further, $f_X \lesssim 10^{-2}$, we enter into the regime of “cold reionization”: with X-rays unable to heat the IGM before the completion of reionization. The corresponding contrast between the cold neutral and ionized IGM drives the 21cm power to amplitudes in excess of 10^3 mK². In this regime, there is effectively no “X-ray heating epoch”, and the highest S/N is again achieved during reionization.

2.5 Considering more pessimistic foregrounds

Preliminary observations by Pober et al. (2013a) suggest that our fiducial choice of $k = 0.1$ Mpc⁻¹ is reasonably free of foregrounds, though slightly larger scales lie securely in the “wedge”. Here we briefly consider a more pessimistic scenario, in which the frequency structure in the foregrounds contaminates modes even out to $k = 0.2$ Mpc⁻¹.

Generally speaking, the resulting noise levels at $k = 0.2$ Mpc⁻¹ for our sensitivity curves are approximately five times higher than at $k = 0.1$ Mpc⁻¹ for all instruments considered. Since the shape of the 21cm power spectrum is flat at the peak amplitude ($dP_{21}/d \ln k \approx 0$; Mesinger et al. 2013), we expect the S/N in the thermal noise dominated regime to be a factor of five times less at $k = 0.2$ Mpc⁻¹ than at $k = 0.1$ Mpc⁻¹.

In Figure 2-10 we show the maximum S/N obtainable with MWA-128T at $k = 0.2$ Mpc⁻¹ (to be compared with the top left panel of Fig. 2-7). We see that the detectable parameter space shrinks, and the broad $S/N \gtrsim 10$ region is now only marginally detectable at $S/N \sim$ unity. This confirms that we have a relatively narrow k -space “window” with first generation instruments; their detection of X-ray heating and reionization might depend on our ability to avoid or mitigate foregrounds at $k \approx 0.1$




Figure 2-10: The maximum S/N for MWA-128T at $k = 0.2\text{Mpc}^{-1}$.

Mpc^{-1} .

2.6 Conclusions

Upcoming statistical and (eventual) tomographical studies of the early Universe with 21cm interferometry should dramatically increase our understanding of both astrophysics and cosmology. The potential of this probe has not yet been fully explored.

Here we perform an astrophysical parameter study, exploring different minimum DM halo masses required to host galaxies, M_{\min} , as well as the galactic X-ray emissivity, normalized to present-day values, f_X . We also discuss the signal in terms of popular warm dark matter models, recasting M_{\min} to an analogous warm dark matter particle mass, m_{wdm} . We study the detectability of the 21cm power spectrum with current and upcoming interferometers. We quantify the peak S/N as well as the S/N at the peak signal, generally corresponding to the epoch of X-ray heating.

Our two free parameters, M_{\min} and f_X , are fundamental in controlling the timing and relative offset of reionization and X-ray heating. The resulting plane of models in parameter space spans the maximum variation in the S/N. Hence, although we focus on the under-appreciated X-ray heating epoch, we also expect our predictions for reionization and the overall achievable S/N in Fig. 2-7 to be robust.

For values of $10^{-2} \lesssim f_X \lesssim 10^2$, the peak amplitude of the 21cm power at $k = 0.1 \text{Mpc}^{-1}$ is roughly constant at a few hundred mK^2 . In this regime, the peak power

occurs during X-ray heating, when just a few percent of the IGM is in emission and the fluctuations in T_γ/T_S are maximized. Stronger X-ray backgrounds instead heat the IGM before its spin temperature has efficiently coupled to the gas temperature, resulting in a weak signal, $P_{21} \sim 10 \text{ mK}^2$. On the other hand, weaker X-ray backgrounds are insufficient to heat the IGM before the completion of reionization. Reionization in these scenarios proceeds in a cold IGM, with the resulting contrast driving the 21cm power to values in excess of thousands of mK^2 .

Aside from bandpass limitations, in “reasonable” models (within an order of magnitude of fiducial values) X-ray heating is detectable at roughly comparable signal-to-noise to reionization. The increase in the signal during the X-ray heating epoch can approximately compensate for the increase in thermal noise going to lower frequencies. A stronger detection is achievable if X-ray heating occurred late, driven by either X-ray faint galaxies ($f_X \lesssim 1$) or those hosted by halos more massive than the cooling threshold, $M_{\text{min}} > 10^8 M_\odot$ (or analogously if $m_{\text{wdm}} < 3.6 \text{ keV}$).

For reasonable models, it is unclear if all first generation interferometers will detect reionization with a 2000h observation. For MWA128T and PAPER, the peak S/N should be of order unity. Robust detections are only likely if we can effectively mitigate foregrounds on large ($k < 0.1 \text{ Mpc}^{-1}$) scales, or if contiguous frequency coverage allows us to sum detections over several frequency (i.e. redshift) bins. Stronger detections, with $S/N \gtrsim 10$ are likely with LOFAR. On the other hand, the continuous bandpass extending to $z \approx 18$ allows MWA to detect a broader range of X-ray heating models, compared to the more limited bandpasses of PAPER and LOFAR. Reasonable models of X-ray heating could be detectable at S/N of unity. The prospects for detecting both reionization and X-ray heating are much improved for MWA with an extension to 256 tiles. The SKA will be a huge improvement over first generation instruments.

We thank A. Liu for providing the 128 element PAPER specifications, HERA specifications, as well as for comments on a draft version of this work. We also thank

A. Loeb for constructive comments on the draft, as well as J. Dillon and L. Koopmans for fruitful conversations. We thank F. Pacucci for assistance with matplotlib. A. E.-W. and J. H. thank the Massachusetts Institute of Technology School of Science for support. This material is based upon work partially supported by the National Science Foundation Graduate Research Fellowship (to A. E.-W.) under Grant No. 1122374.

Chapter 3

Detecting the 21 cm Forest in the 21 cm Power Spectrum

The content of this chapter was published in *Ewall-Wice, A., Dillon, J. S., Mesinger, A. and Hewitt, J. (2014), 'Detecting the 21 cm forest in the 21 cm power spectrum', MNRAS 441, 2476–2496.*

3.1 INTRODUCTION

Observations of emission and absorption at 21 cm from the neutral intergalactic medium (IGM) at high redshift will offer an unprecedented glimpse into the cosmic dark-ages up through the epoch of reionization (EoR), constraining both fundamental cosmological parameters and the properties of the first stars and galaxies (Furlanetto et al., 2006a; Morales and Wyithe, 2010; Pritchard and Loeb, 2012, for reviews). Direct mapping of the 21 cm signal during the EoR is likely a decade or more away, requiring projected instruments such as the Square Kilometer Array (SKA). However, a first generation of experiments attempting to detect the power spectrum are already underway. These include the Low Frequency ARray (van Haarlem et al., 2013, LO-FAR), the Murchison Widefield Array (Tingay et al., 2013a, MWA), the Precision Array for Probing the Epoch of Reionization (Parsons et al., 2014, PAPER), and the Giant Metre-wave Telescope (Paciga et al., 2013, GMRT). The MWA, PAPER, and

LOFAR have the potential to achieve statistical detections of brightness temperature fluctuations within the next several years (Bowman et al., 2006; McQuinn et al., 2006; Beardsley et al., 2013; Mesinger et al., 2013)

Most theoretical investigations of observing neutral hydrogen in the EoR have focused on IGM emission and absorption against the Cosmic Microwave Background (CMB). It has also been recognized by Carilli et al. (2002); Furlanetto and Loeb (2002); Xu et al. (2009); Mack and Wyithe (2012); Ciardi et al. (2013) that the 21 cm forest, HI absorption in the spectra of background radio-loud (RL) active galactic nuclei (AGN), can be used to probe the IGM's thermal state.

Studies of the forest have focused on its detection in the frequency spectra of a known RL source to glean information on the thermal properties of the absorbing IGM. The possibility for such a study depends on the existence of high redshift RL sources. As of 2013, the RL source distribution is only well constrained out to $z \sim 4$ (see de Zotti et al. (2010) for review). Theoretical work suggests that at 100 MHz hundreds of $S \sim 1$ mJy sources with redshifts greater than 10 might exist within one of the (30^2 fields of view (FoV) offered by existing and upcoming wide field interferometers (Haiman et al., 2004; Wilman et al., 2008) (hereafter H04 and W08 respectively). However the discovery of a suitable source at high redshift entails an extensive follow up program to measure photometric redshifts of radio selected candidates.

Should sufficiently RL sources exist, a line of sight (LoS) detection of individual absorption features will require large amounts of integration time on a radio telescope with the collecting area comparable to the Square Kilometer Array (SKA). At reionization redshifts, (Mack and Wyithe, 2012) find that a 5σ detection of an individual absorption feature with a $z \approx 9$ Cygnus A type source¹ would require years of integration on an SKA-like instrument. Ciardi et al. (2013) find that after 1000 hours of integration only 0.1% of the LoS in an IGM simulation box contained regions of large enough optical depth to produce absorption features² observable by LOFAR. Hence

¹ flux density at 151 MHz of $S_{151} \approx 20$ mJy and spectral index of $\alpha \approx 1.05$

² against a $S_{129} \approx 50$ mJy source at $z \sim 7$

a detection of the forest with a present day interferometer would require a very rare juxtaposition of an extremely loud RL source with an outlying optical depth feature. Even if this detection were achieved, it is unlikely that significant inferences on the thermal history could be made from only a handful of such observations.

While detecting individual absorption features presents an enormous challenge, statistical methods have been demonstrated to reduce the necessary integration times. One target for a statistical detection is the increased variance in flux, along the LoS. It is shown in (Mack and Wyithe, 2012) that the integration time required for detecting this variance increase for a Cygnus A source, is only a few weeks with an SKA-like telescope, as apposed to the decades needed for detecting a single feature. Ciardi et al. (2013) find that LOFAR could detect the global suppression in the spectrum of a 50 mJy source at $z \sim 12$ with a 1000 hour integration, though they note that a detection by LOFAR is unlikely due to excessive RFI in the FM band ($80 \text{ MHz} \lesssim \nu \lesssim 108 \text{ MHz}$).

The possibility of a statistical detection of the forest using information from the wide FoV available to the current and upcoming generations of experiments has not yet been investigated. Observing the forest signature in the 21 cm power spectrum would integrate the signal from many high redshift sources within a FoV, reducing the sensitivity requirements of the instrument. Also, a power spectrum detection does not require a priori knowledge of high redshift sources. Hence the technique we describe can put constraints on both the properties of the IGM, such as the heating and reionization history, and the population of high redshift RL sources. It is likely that 21cm forest absorption features could be fruitfully explored using high-order statistical measures as well, but we do not consider those in this paper.

In this proof-of-concept, we begin to explore the characteristics and observability of the forest in the 21 cm power spectrum. We derive analytically the features that the global forest should introduce to the power spectrum and confirm their existence by combining semi-numerical simulations of the IGM, computed with 21cmFAST (Mesinger et al., 2011), with the semi-empirical model of the high redshift population of RL sources from W08. We find that in all heating scenarios studied, the contri-

bution to the 21 cm fluctuations by the absorption of our RL sources is comparable to or dominates the contribution from the brightness temperature on small spatial scales ($k \gtrsim 0.50 \text{ Mpc}^{-1}$). To determine the detectability of the forest in the power spectrum, we perform sensitivity calculations for several radio arrays with designs similar to the MWA, including a future array with a collecting area of $\sim 0.1 \text{ km}^2$, similar to the planned Hydrogen Epoch of Reionization Array (HERA). In order to give the reader a sense of how the strength of this signal scales across a large range of radio loud source populations, we extrapolate the expected S/N of the Forest using our analytic expression for the signal strength.

This paper is organized as follows. In Section 3.2 we provide the theoretical background and use a toy model to derive the morphology of the 21 cm forest power spectrum; relating its shape and amplitude to the optical depth power spectrum and the radio luminosity function. In Section 3.3 we describe the semi-numerical simulations of the IGM along with the semi-empirical RL source distribution of W08 and how we combine them to simulate the wide field forest. In Section 3.4 we discuss our results and identify the separate regions of k -space that may be used to independently constrain the thermal history of the IGM and the high redshift RL distribution. In Section 3.5 we explore the prospects for detecting the forest in spherically averaged power spectrum measurements considering the sensitivity of current and future radio arrays. In Section 3.6 we extrapolate our detectability results across a broad range of source populations and X-ray heating scenarios.

Throughout this work we assume a flat universe with the cosmological parameters $h = 0.7$, $\Omega_\Lambda = 0.73$, $\Omega_M = 0.27$, $\Omega_b = 0.082$, $\sigma_8 = 0.82$, and $n = 0.96$ as determined by the WMAP 7-year release (Komatsu et al., 2011). All cosmological distances are in comoving units unless stated otherwise.

3.2 THEORETICAL BACKGROUND

In this section we establish our notation and present a basic mathematical description of how forest absorption modifies the 21 cm brightness temperature signal.

3.2.1 Notation

We adopt the Fourier transform convention

$$\tilde{f}(\mathbf{k}) = \int d^3x e^{-i\mathbf{k}\cdot\mathbf{x}} f(\mathbf{x}). \quad (3.1)$$

In addition, we often refer to cylindrical Fourier coordinates where $k_\perp \equiv \sqrt{k_x^2 + k_y^2}$ and $k_\parallel \equiv |k_z|$. The power spectrum of a field A over a comoving volume V is defined as

$$P_A = \frac{1}{V} \langle |\widetilde{\Delta A}|^2 \rangle \quad (3.2)$$

and the cross power spectrum between fields A and B over V is given by

$$P_{A,B} = \frac{1}{V} \langle \widetilde{\Delta A} \widetilde{\Delta B}^* \rangle \quad (3.3)$$

where

$$\Delta A = A - \langle A \rangle \quad (3.4)$$

and $\langle A \rangle$ is defined as the ensemble average of A though in practice it is computed by averaging over some spatial or Fourier volume. In our discussion, we will also be referring to the one dimensional LoS power spectrum (not to be confused with the 1D spherical power spectrum) of a field A along a LoS column of comoving length L .

$$P_A^{LoS}(k_z) = \frac{1}{L} \int dz dz' \Delta A(z) \Delta A(z') e^{ik_z(z-z')} \quad (3.5)$$

Finally, we use Δ^2 to denote the dimensionless power spectrum

$$\Delta^2(k) \equiv \frac{k^3}{2\pi^2} P(k) \quad (3.6)$$

3.2.2 The Forest's Modification of the Brightness Temperature

The forest absorption traces the optical depth of the IGM and will therefore introduce a signal on similar spatial scales as the 21 cm brightness temperature. We now discuss this signal in detail. The optical depth of a high redshift HI cloud is given by (Furlanetto et al., 2006a)

$$\tau_{21} \approx .0092(1 + \delta)(1 + z)^{3/2} \frac{x_{HI}}{T_s} \left[\frac{H(z)/(1 + z)}{dv_{\parallel}/dr_{\parallel}} \right]. \quad (3.7)$$

δ is the fractional baryonic over-density, $H(z)$ is the Hubble factor, $dv_{\parallel}/dr_{\parallel}$ is the velocity gradient along the LoS (including the Hubble expansion), and x_{HI} is the neutral hydrogen fraction. The numerical factor in front of Equation (3.7) is computed from fundamental constants and is independent of cosmology. The spin temperature, T_s is defined by the relative population densities of the two hyperfine energy levels, n_1 and n_0 (Field, 1958)

$$\frac{n_1}{n_0} = 3 \exp\left(-\frac{h\nu_{21}}{k_B T_s}\right). \quad (3.8)$$

Where, h is Plank's constant, k_B is the Boltzmann constant, and $\nu_{21} = 1420.41\text{MHz}$ is the rest frame frequency of the hyperfine transition radiation.

Prior works on 21 cm tomography assume that the sky temperature at $\nu = \nu_{21}/(1+z)$ in the direction of an HI cloud is given by

$$T_{sky} = \frac{T_s}{(1+z)}(1 - e^{-\tau_{21}}) + \frac{T_{CMB}}{(1+z)}e^{-\tau_{21}} + T_{fg}. \quad (3.9)$$

where T_{CMB} is the comoving temperature of the cosmic microwave background radiation and T_{fg} is the temperature of foreground emission including synchrotron radiation of the Galaxy, resolved point sources, free-free emission, and radio emission from unresolved point sources below the confusion limit (Di Matteo et al., 2002; Jelić et al., 2008; de Oliveira-Costa et al., 2008; Wang et al., 2006).

The first term in Equation (3.9) includes both the 21 cm emission and self absorp-

tion of the HI cloud, hence it is multiplied by a factor of $(1 - e^{-\tau_{21}})$. The second term describes the observed intensity of a background source shining through the cloud so its temperature is attenuated by $e^{-\tau_{21}}$. The third term describes radiation emitted by sources closer than the cloud so its intensity unaffected by τ_{21} .

21 cm experiments seek to measure the difference between the first two terms of Equation (3.9) and T_{CMB} . This difference is often referred to as the ‘‘differential brightness temperature’’ and is given by (Furlanetto et al., 2006a)

$$T_b = \frac{(T_s - T_{CMB})}{(1+z)}(1 - e^{-\tau_{21}}) \approx \frac{T_s - T_{CMB}}{(1+z)}\tau_{21}. \quad (3.10)$$

We depart from previous work by considering the effect of radio loud sources behind the HI cloud whose combined observed³ brightness temperature we denote as T_{RL} . Including these background sources, Equation (3.9) becomes

$$T'_{sky} = \frac{T_s}{(1+z)}(1 - e^{-\tau_{21}}) + \frac{T_{CMB}}{(1+z)}e^{-\tau_{21}} + T_{RL}e^{-\tau_{21}} + T_{fg} \quad (3.11)$$

T_{fg} and T_{RL} are expected to have predominantly smooth spectra which reside within a limited region of Fourier space known as the ‘‘wedge’’ (Datta et al., 2010; Morales et al., 2012; Vedantham et al., 2012). Smooth spectrum components may be removed by filtering (Parsons et al., 2012a) or subtraction (Bowman et al., 2009; Liu et al., 2009; Dillon et al., 2014), both employing the separation of the foregrounds and signal in the Fourier domain.

We will focus on the fluctuating signal, assuming that the smooth spectrum components of the foregrounds and background sources are properly avoided and/or subtracted. The effective differential brightness temperature now includes a contribution from the forest absorption features.

$$T_b \rightarrow T'_b \approx T_b - T_f \quad (3.12)$$

³In accordance with much of the literature, we use the observed temperature for T_{RL} and T_{fg} , rather than the comoving temperature as we have for T_s and T_{CMB} . As a result, there are no factors of $(1+z)$ under T_{RL} or T_{fg} .

where $T_f = T_{RL}\tau_{21}$ is the “forest temperature”. We can see how the power spectrum is transformed by the inclusion of T_f by inserting Equation (3.12) into Equation (3.2)

$$P_b \rightarrow P'_b = P_b + P_f - 2\text{Re}(P_{f,b}) \quad (3.13)$$

Where $P_b \equiv P_{T_b}$, $P'_b \equiv P_{T'_b}$, $P_f \equiv P_{T_f}$ and $P_{f,b} \equiv P_{f,T_b}$. Equation (3.13) sums up how the forest modifies the power spectrum that we expect to observe in upcoming 21 cm observations. Essentially, smooth spectrum power from T_{RL} is leaked from the largest spatial modes to those occupied by T_b via a convolution with the power spectrum of the optical depth field. The magnitude of this leakage will increase with the magnitude of the optical depth.

3.2.3 The Morphology of the Forest Power Spectrum

The first thing one might ask concerning the forest contribution described in Equation (3.13) is how the magnitudes of the two contributions compare to each other and what their qualitative features are. While we will answer these questions with simulations it is useful to gain as much insight as we can through analytic methods. We start with P_f which can be decomposed (see Appendix 3.A for a derivation) into a sum of auto power spectra P_j originating from each individual RL source behind or within an imaged volume of IGM and their cross power spectra, $P_{j,k}$.

$$P_f = \frac{1}{V} \left\langle \left| \widetilde{T_{RL}\tau_{21}} \right|^2 \right\rangle = \sum_j P_j + 2\text{Re} \left(\sum_{j<k} P_{j,k} \right) \quad (3.14)$$

If all of the background sources are unresolved⁴ then each P_j is the absolute magnitude of the Fourier transform of a function that is a delta function in the perpendicular to LoS directions. As a result, each P_j in Equation (3.14) is constant in k_\perp . The cross multiplying $P_{j,k}$ terms are not so simple; however, we show in Appendix 3.A that in the absence of clustering, the cross sum only contributes to P_f at the 10% level for

⁴a fair assumption given the large synthesized beams of interferometers and small angular extent of high redshift sources

$k_{\parallel} \gtrsim 0.1 \text{ Mpc}^{-1}$. At these scales, P_f only has considerable structure along k_{\parallel}

$$P_f(\mathbf{k}) \approx \sum_j P_j(k_{\parallel}) = \frac{D_M^2 \lambda^4}{4k_B^2 \Omega_{cube}} P_{\tau_{21}}^{LoS}(k_{\parallel}) \langle \sum_j s_j^2 \rangle \quad (3.15)$$

where $\lambda = \lambda_{21}(1+z)$ is the observed wavelength of 21 cm light emitted from the center of the imaged volume, D_M is the comoving distance to the data cube, and Ω_{cube} is the solid angle subtended by the cube. In the second step, we have expressed each P_j in terms of the flux of each source, s_j , and the 1D power spectrum along the line-of-sight to that source, $P_{\tau_{21}}^{LoS}$. In addition P_f is positive so that it will always add to the power spectrum amplitude

We can convert the sum in Equation (3.15) to an integral over the radio luminosity function

$$P_f \approx \frac{c D_M^2 \lambda^4}{4k_B^2} P_{\tau_{21}}^{LoS}(k_{\parallel}) \int_z^{\infty} \int_0^{\infty} s'^2 \rho(z, z', s') \frac{D_M^2(z')}{H(z)} dz' ds' \quad (3.16)$$

where $\rho(z, z', s') \equiv \frac{dN}{ds' dV_c}$ is the differential number of radio loud sources per comoving volume at redshift z' per flux bin at observed frequency $\nu_{21}/(1+z)$ and s' is the flux at $\nu = \nu_{21}/(1+z)$.

Equation (3.16) tells us that the amplitude of the forest power spectrum is set by the integral over the high redshift radio luminosity function multiplied by the average optical depth squared⁵ while the shape of the forest power spectrum is set by the 1D LoS power spectrum of optical depth fluctuations.

$P_{f,b}$ does not separate so conveniently but we can gain insight into whether it adds or subtracts to Equation (3.13) by considering the physical phenomena that govern T_f and T_b . Expanding Equation (3.10) and T_f in terms of the IGM properties using Equation (3.7) one can see that $P_{f,b}$ is the cross power spectrum between the two quantities:

$$T_b \approx 9x_{HI}(1+\delta)(1+z)^{1/2} \left[1 - \frac{T_{CMB}}{T_s} \right] \left[\frac{H(z)}{dv_{\parallel}/dr_{\parallel}} \right] \text{ mK} \quad (3.17)$$

⁵By our definition, the power spectrum is the Fourier transform squared of $\Delta\tau_{21}$, not $\delta\tau_{21} = \Delta\tau_{21}/\langle\tau_{21}\rangle$ which is often used in other work. Hence our power spectrum amplitude is set by $\langle\tau_{21}\rangle^2$

and

$$T_f \approx 0.009 x_{HI} (1 + \delta) (1 + z)^{1/2} \frac{T_{RL}}{T_s} \left[\frac{H(z)}{dv_{\parallel}/dr_{\parallel}} \right] \quad (3.18)$$

Before the reionization era, x_{HI} is relatively homogenous so that fluctuations in T_b are governed primarily by those in T_s . Regions of the IGM with larger T_s will have more positive T_b but smaller T_f . Because of this anti-correlation between T_f and T_b , $\text{Re}(P_{f,b})$ is negative during the pre-reionization era and the net effect will be for it to increase the power spectrum amplitude through its negative contribution in Equation (3.13). At lower redshifts, after X-rays have heated the IGM, $T_s \gg T_{CMB}$, and T_b becomes independent of T_s . As a result, T_b is always positively correlated with x_{HI} as is T_f . $\text{Re}(P_{f,b})$ is positive with a net effect of subtracting from the power spectrum amplitude. We are unable to make any more progress analytically, but we will reexamine the cross power term in our simulation results below.

We now move on to describe our simulations. We will return to our discussion of the power spectrum morphology in the context of our simulation results in Section 3.4.

3.3 SIMULATIONS

In this section we describe the semi-numerical simulations that we use to explore a range of IGM thermal histories along with the the semi-empirical RL source model that we employ to add the 21 cm forest signal.

3.3.1 Simulations of the Optical Depth of the IGM

Our IGM simulations are run using a parallelized version of the public, semi-numerical 21cmFAST code⁶ described in Mesinger et al. (2011). Tests of the code can be found in Mesinger and Furlanetto (2007); Zahn et al. (2011); Mesinger et al. (2011). The simulation box is 750 Mpc on a side, with resolution of 500^3 . Different scenarios for τ_{21} can be obtained by exploring histories of the spin temperature, T_s and/or the

⁶<http://homepage.sns.it/mesinger/Sim>

neutral fraction, x_{HI} .

21cmFAST includes sources of both UV ionizing photons and X-rays. The former dominate reionization (i.e. x_{HI}), except for extreme scenarios we do not consider in this work (Furlanetto, 2006b; McQuinn and O’Leary, 2012; Mesinger et al., 2013). Since a full parameter study is beyond the scope of this work, and since the bulk of the relevant signal is likely during the pre-reionization epoch, we fix the ionizing emissivity of galaxies (and hence the reionization history), to agree with the Thompson scattering optical depth from WMAP (Komatsu et al., 2011). Instead we focus on the X-ray emissivity and its impact on T_s .

T_s is affected by a variety of processes. These include Ly- α photons which couple to the hyperfine transition through the Wouthuysen-Field effect (Wouthuysen, 1952; Field, 1958), particle collisions, and emission or absorption of CMB photons. The coupling of T_s to these processes is described by (e.g. Furlanetto et al. 2006a):

$$T_s^{-1} = \frac{T_{CMB}^{-1} + x_c T_k^{-1} + x_\alpha T_c^{-1}}{1 + x_c + x_\alpha} \quad (3.19)$$

where T_k is the kinetic temperature of the HI gas, T_c is the color temperature of Ly- α photons, and x_c and x_α are the collisional and Ly- α coupling constants. Due to the high optical depth of the neutral IGM to Ly- α photons, the color temperature is very closely coupled to the kinetic temperature, $T_c \approx T_k$ (Wouthuysen, 1952; Hirata, 2006)

Although the self-annihilation of some dark matter candidates can contribute significantly (Valdés et al., 2013), in fiducial models T_k is predominantly determined by X-ray heating (e.g. Furlanetto et al. 2006a). Hence, we explore a range of optical depth histories by running simulations for different galactic X-ray emissivities.

We use the fiducial model of X-ray heating described in Mesinger et al. (2013), adopting a spectral energy index of $\alpha = 1.5$ and an obscuration threshold of 300 eV. We parameterize the X-ray luminosity by a dimensionless efficiency parameter, f_X . Our fiducial model, $f_X = 1$ corresponds to 0.2 photons per stellar baryon, or a total X-ray luminosity above $h\nu_0 = 0.3$ keV of $L_{X,0.3+\text{keV}} \approx 10^{40}$ erg s $^{-1}$ (M_\odot yr $^{-1}$) $^{-1}$. This

Table 3.1: IGM Heating Parameters

Number	Name	f_X
	Hot IGM	5.0
	Fiducial IGM	1.0
	Cool IGM	0.2




Figure 3-1: The mean thermal evolution of our IGM simulations for our three models. “cool IGM”- solid lines, “fiducial IGM”- dashed-dotted lines, and “hot IGM” - dashed lines. $\langle T_s \rangle$ is plotted in lavender. Varying f_X effectively shifts $\langle T_s \rangle$ in redshift.

choice is consistent with (a factor of ~ 2 higher than) an extrapolation from the 0.5–8 keV measurement of Mineo et al. (2012a) that $L_{X,0.5-8\text{keV}} \approx 3 \times 10^{39} \text{ erg s}^{-1} (M_\odot \text{ yr}^{-1})^{-1}$.

Summarized in Table 3.1 are our three values of f_X : a “fiducial IGM” model with $f_X = 1$ corresponding to the fiducial value in Mesinger et al. (2013), a “hot IGM” model with $f_X = 5$, and a “cool IGM” model with $f_X = 0.2$. In Figure 3-1 We show the evolution of the mean spin and brightness temperatures from our simulations. Over the range of emissivities considered, the effect of varying f_X is to shift the evolution of $\langle T_s \rangle$ in redshift. Because P_f varies as $\langle \tau_{21} \rangle^2 \sim \langle T_s \rangle^{-2}$ and f_X simply shifts $\langle T_s \rangle$ in redshift, this relatively modest spread in f_X is sufficient to understand a broader range of expected outcomes, as we shall see below.

3.3.2 The Model of the Radio Loud Source Distribution

We now review present constraints on the RL source distribution and describe the semi-empirical radio luminosity function that we use to simulate the global 21 cm forest. To gain perspective of how our choice of population model might compare to other theoretical work we determine which flux ranges are relevant to the sum in Equation (3.15) and compare the counts of sources in W08 to those in H04. We also describe our method for combining the simulated radio sources with our simulations of the IGM.

3.3.2.1 Review of Constraints and Predictions of High Redshift Radio Counts

Constraints on the luminosity function of the most luminous radio loud sources are presently limited to $z \sim 4$ (de Zotti et al., 2010). Confirmed in these works, is that the comoving density of ultra steep spectrum sources peaks at $z \sim 2$ with little evidence for evolution out to $z \gtrsim 4.5$.

To model the abundance of RL quasars with $6 \lesssim z \lesssim 20$ one must rely on theoretical extrapolations. Haiman et al. (2004) give estimates of source counts by assigning black hole masses to a halo mass function using the black hole mass-velocity dispersion relation of Wyithe and Loeb (2003). The RL fraction is derived assuming Eddington accretion, and the RL-i band luminosity correlation observed by Ivezić et al. (2002).

More sophisticated attempts at predicting the bolometric luminosities of high redshift quasars up to $z = 11$ have been undertaken using hydrodynamic simulations with self consistent models for black hole growth and feedback (DeGraf et al., 2012). Even with a more nuanced treatment of the luminosity distribution, the RL fraction at high redshift still remains a wide open question. Indeed, the purpose of this work is to propose a technique for determining this population by showing that an empirically motivated RL population can have significant and observable features in the power spectrum for a range of thermal scenarios.




Figure 3-2: The 21 cm forest is dominated by sources in the 1-10 mJy flux range. We plot the sum of fluxes squared, in Equation (3.16), for $S < S_\nu$. A detection of P_f would constrain the high redshift source counts at these flux intervals.

3.3.2.2 Our Choice of Population Model

We choose to work with the RL AGN population described in W08 in which sources are generated by sampling extrapolated radio luminosity functions biased to structure from a CDM simulation. Specifically, the radio luminosity function used is that “Model C.” from Willott et al. (2001) which describes the high and low luminosity populations of AGN as Schechter functions. The redshift evolution of the low luminosity population is modeled as a power law in redshift while the high luminosity component as a gaussian with a mean of $z \approx 1.9$. Lists of source positions, fluxes, and morphologies from the Wilman simulation are downloadable through a web interface⁷.

Having chosen our population model, we can employ our formalism from Section 3.2 to understand which sub population of the luminosity function contributes most to P_f . In Figure 3-2, we plot the percent contribution of sources below a threshold, S_ν , to P_f from the flux squared sum in Equation (3.15). One can see that roughly 75% of the contribution to P_f comes from sources with fluxes between 1 – 10 mJy at 80 – 115 MHz. At lower redshifts, the integral curves are increasingly dominated by higher fluxes as the sources with the greatest fluxes increase in number. The detection or lack of detection of the features we find using this simulation would either confirm or reject the W08 model for sources with S_ν between 1 and 10 mJy. While this paper

⁷http://s-cubed.physics.ox.ac.uk/s3_sex

is a study of observability for one model, in future work we will determine what range of RL population this technique can constrain.

It is worth getting an order of magnitude idea of how our choice of the W08 semi-empirical model might compare to other theoretical predictions of the radio luminosity function. In Appendix 3.B we compare the source counts in our semi-empirical prediction to the more physically motivated bottom up model in H04. The counts of W08 sources contributing to the bulk of P_f tend to be more numerous than those in H04 by a factor of ≈ 10 at $z \sim 12$ to ≈ 80 for $z \sim 15 - 20$, underscoring the need for a full parameter space study. Even though such a study is beyond the scope of this paper, our extrapolated results in Figure 3-14 show that the range of populations that the power spectrum can constrain depends heavily on the IGM's thermal history.

3.3.3 Adding Sources to the Simulation

We simulate the theoretical power spectra accessible to upcoming observations by drawing 36 random sub-fields from the W08 simulations and combining them with 36 random 8MHz slices from our IGM simulations. The number of subfields is chosen to roughly correspond to the $\sim (30^2)$ FoV of the MWA.

While our analytic approach in Section 3.2 does not account for sources within the imaged volume, we incorporate them into our simulation by determining the location of DM halos down to masses of $5 \times 10^9 M_\odot$ through the excursion-set + perturbation theory approach outlined in Mesinger and Furlanetto (2007). We then populate these dark matter halos with RL sources, monotonically assigning the most luminous sources at 151 MHz⁸ to the most massive halos. Sources falling behind the cubes retain their original positions. All W08 sources are unresolved in our IGM simulation; hence, for each pixel the fluxes for all sources behind that pixel are summed together to give S_{pix} . This flux cube is converted to temperature using the

⁸We order sources by their luminosity at *observed* frequency of 151 MHz at regardless of their redshift which varies very little over the span of an 8 MHz data cube so that we are approximately comparing their rest frame luminosities.

Rayleigh-Jeans equation,

$$T_{pix} = \frac{\lambda^2 S_{pix}}{2k_B \Omega_{pix}}, \quad (3.20)$$

where Ω_{pix} is the solid angle subtended by each simulation pixel⁹. Finally we introduce quasar absorption by multiplying this source cube by our τ_{21} cube $T_f \approx T_{pix} \tau_{21}$.

3.4 SIMULATION RESULTS

In this section we present the results of our combined IGM-RL population model by computing the spherical and cylindrical power spectrum, $P(k)$, averaged over our 36 sub-cubes. We identify the regions of k -space in which the forest is dominant and might be used to constrain the high redshift radio luminosity function and discuss the morphology of the observed power spectra, verifying the essential results of Section 3.2.

3.4.1 Computing Power Spectra

Power spectra are computed using a direct Fast Fourier transform of each data cube multiplied by a kaiser window along the LoS with attenuation parameter $\beta = 3.5$. In averaging over bins of our spherical power spectra, we exclude the “wedge”, the region of k -space heavily contaminated by foregrounds given by (Vedantham et al., 2012; Morales et al., 2012)

$$k_{\parallel} \leq \sin \frac{\Theta}{2} \left(\frac{D_M(z)}{D_H} \frac{E(z)}{(1+z)} \right) k_{\perp} \quad (3.21)$$

where z is the redshift of a data cube’s center frequency, $D_M(z)$ is the comoving distance, $E(z) = H(z)/H_0$, and Θ is the FWHM of the primary beam which we calculate using a short dipole model of the MWA antenna element. Table 3.2 gives the FWHM value of our primary beam model for several different frequencies.

⁹We show in Appendix 3.A, that the choice of pixel solid angle does not effect P_f




Figure 3-3: For every heating scenario we study, there is some redshift and region within the EoR window for which 21 cm forest dominates the power spectrum. Here we show the fractional difference between the power spectrum with, (P'_b), and without (P_b) the forest for the redshifts (top to bottom) 9.2, 11.2, 12.2, 15.4, and 17.5. The diagonal lines denote the location of the “wedge”. By $z \gtrsim 12.2$ there is a substantial region ($k_{\parallel} \gtrsim 0.5 \text{ Mpc}^{-1}$) of the Fourier volume that our simulations cover in which the forest dominates P_b by a factor of a few.

3.4.2 Simulation Output and the Location of the Forest in k -space

We now discuss the power spectra output by our simulations and the significant features produced by the forest.

To isolate the the effect of the forest and to compare its significance to the brightness temperature power spectrum, P_b , we plot the fractional difference between P'_b , the power spectrum with the forest, and P_b in Figure 3-3. We see that the forest introduces a significant feature, especially at the smallest scales. This feature is most prominent at high redshifts and less emissive heating models, when the IGM is cool. For our cool model, the forest feature dominates P_b by over a factor of 100 for a wide range of redshifts. In the fiducial model, the dominant region is primarily at larger values of k_{\parallel} , though dominance by a factor of a few is visible at $z = 12.2$ and $z = 17.5$. In our hot model, a significant feature is visible only for $z \gtrsim 12.2$.

For all heating scenarios, there are redshifts $z \gtrsim 12.2$ in which the same region of Fourier space contains a strong forest signal that dominates P_b by a factor of at least a few. Fortunately for those interested in the brightness temperature signal, the region $k \lesssim 10^{-1} \text{ Mpc}^{-1}$ remains dominated by P_b . Hence at pre-reionization redshifts, $k \lesssim 10^{-1} \text{ Mpc}^{-1}$ can still be used to constrain cosmology and the thermal history of the IGM. With the thermal properties of the IGM determined, one may constrain the high redshift RL population using the forest power spectrum signal at $k \gtrsim 0.5 \text{ Mpc}^{-1}$.

The first generation of interferometers will not be sensitive enough to measure the cylindrical power spectrum with high S/N but will rather measure the spherically averaged power spectrum. We compute spherically averaged power spectra from data cubes with and without the presence of forest absorption and excluding the wedge. We plot these power spectra in Figure 3-4. In all of the heating scenarios considered, the forest introduces significant power at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ for $z \gtrsim 15.4$. Hence, it is in principle possible to constrain the distribution of RL AGN at high redshift for a range of heating scenarios.

We note that the high- k region extends into our simulations' Nyquist frequency of




Figure 3-4: The 21 cm forest dominates the spherically averaged power spectrum for $k \gtrsim 0.5 \text{ Mpc}^{-1}$. Plotted is the spherically averaged power spectrum with (dashed lines) and without (solid lines) the presence of the 21 cm forest. In our cool model, the forest causes a significant power increase at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ at redshifts as low as $z = 11.2$. At $z = 15.4$ we see a significant feature in all thermal scenarios. Our cool IGM model experiences a reduction in the power spectrum amplitude at $z \gtrsim 17.5$ as it passes through the X-ray heating peak.




Figure 3-5: We plot the magnitude of the difference between the 21 cm power spectrum with and without the presence of the 21 cm forest including the auto-power and cross power terms of Equation (3.13). At high redshifts and low f_X , there is little k_\perp structure in $P'_b - P_b$, indicating that P_f is the significant contributor. At lower redshifts and higher f_X , we see significant k_\perp structure, indicating that in a heated IGM, $P'_b - P_b$ is dominated by $P_{f,b}$ which is somewhat spherically symmetric and negative at large k . The trough in the low redshift plots marks the region where $P_f - 2\text{Re}(P_{f,b})$ transitions from negative (for small k) to positive (for large k).

2.1 Mpc^{-1} . We ensure that the forest dominance is not an aliasing effect by running simulations on a 125 Mpc cube with six times higher resolution. The results in the overlapping k -space regions agree well with these larger volume, lower resolution simulations.

3.4.3 The Morphology of the Simulation results.

We now explain the morphology of our simulation results and verify our analytic predictions in Section 3.2.

We noted in Figure 3-3 that the 21 cm forest dominates the power spectrum both

at large k_{\perp} and k_{\parallel} . The former observation is consistent with a forest power spectrum that is uniform in k_{\perp} . In Figure 3-5 we show $|P'_b - P_b|$ and see that at high redshift and cool heating models, the forest power spectrum is mostly uniform in \perp though at lower redshifts and hotter IGM, there is significant k_{\perp} structure. Since in section 3.2 we showed that P_f only varies along k_{\parallel} , this suggests that the cross power spectrum, $P_{f,b}$ is the prime contributor to $P'_b - P_b$ in a hot IGM, while P_f is in a cool one. The trough at lower redshifts, at $k \sim 0.5 \text{ Mpc}^{-1}$ is caused by the fact that $-2P_{f,b}$ is negative as we shall see below.

A potentially interesting consequence of the auto-terms invariance in k_{\perp} is a potential for contaminating the separation of powers analysis advocated in Barkana and Loeb (2005a). We may Taylor expand P_f

$$P_f(k_{\parallel}) = P_f(k\mu) = \sum_{n=1}^{\infty} \frac{1}{n!} \frac{\partial P_f}{\partial (\mu k)}_{\mu k=0} (\mu k)^n, \quad (3.22)$$

so P_f introduces signal over a wide range of powers of μ and has the potential to contaminate the cosmological μ^4 and μ^6 components of the brightness temperature power spectrum. On the other hand, the small k , where the perturbative expansion is most accurate, is dominated by the diffuse brightness temperature emission. In all but the coolest heating models, contamination will likely be small, since we can see in Figure 3-3 that $P_f \lesssim 0.1P_b$ at $k \lesssim 0.1 \text{ Mpc}^{-1}$.

Decomposing the forest signal into powers of μ may be another way of distinguishing it from the brightness temperature. Even within the "IGM dominated" region. Detailed analysis on contamination of the cosmological signal and additional distinguishability offered by the angular dependence is beyond the scope of this paper will be the subject of future work.

To be more quantitative, we turn our attention to right hand side of Equation (3.15) and verify our decomposition of the forest power spectrum into $P_{\tau_{21}}^{LoS}$ and the sum of background source fluxes. To do this, we find the summed squares of the fluxes (at the center frequency of the observation) of all sources falling in or behind our data cubes at several redshifts, multiply by the 1D LoS power spectrum of τ_{21}




Figure 3-6: Our semi-analytic prediction agrees well with unclustered simulation results. The semi-analytic prediction of Equation (3.15) is plotted with dashed lines and $\Delta_f^2(k)$ computed directly from our simulation without clustering in solid lines. This demonstrates that for $k \gtrsim 10^{-1} \text{ Mpc}^{-1}$, the cross terms in Equation (3.14) may be ignored and P_f may be well approximated by the LoS power spectrum of τ_{21} multiplied by the summed squared fluxes for sources lying in and behind the data cube.

and compare with Δ_f^2 computed from our simulation as outlined above. We find that Equation (3.15) consistently underpredicts the simulation amplitude by a factor of 2. However, when we remove the clustering of sources by randomly assigning source positions (rather than using the dark matter biased positions), Equation (3.15) agrees with simulation output within 5 – 20% over the studied redshifts. Hence we rewrite Equation (3.16) as

$$P_f \approx A_{cl} \frac{c D_M^2 \lambda^4}{4 k_B^2} P_{\tau_{21}}^{LoS} \int_z^\infty \int_0^\infty s^2 \rho(z, z', s) \frac{D_M^2(z')}{H(z')} dz' ds \quad (3.23)$$

Where A_{cl} is a constant of order unity that accounts for the boost in power due to clustering. We briefly explain this power boost in Appendix 3.A. In Figure 3 – 6 we show the power spectrum, Δ_f^2 computed from our simulation and the prediction from Equation (3.15) for several redshifts in our fiducial heating model. For $k \gtrsim 10^{-1} \text{ Mpc}^{-1}$, Equation (3.15) agrees with our simulation at the 10% level, indicating that we can ignore the cross terms in Equation (3.14) and consider the forest power spectrum as the simple product of the 1D τ_{21} power spectrum and the integrated radio luminosity function.




Figure 3-7: We see that for a fixed quasar distribution, the magnitude of P_f can be parameterized by $\langle T_s \rangle$ and that the amplitude is consistent with a simple power law. Here, we plot $P_f(k_{\parallel})$ at $k_{\parallel} = 0.5 \text{ Mpc}^{-1}$ vs. $\langle T_s \rangle$ for all considered redshifts and f_X . The black line is the power law $\langle T_s \rangle^{-2}$ as one might expect for an amplitude set by $\langle \tau_{21} \rangle^2$ (Equation (3.16)). Inasmuch of this simple trend, a modest spread in heating models gives us a decent understanding of the behavior of the amplitude for P_f . This relation holds for the quasar population considered here because the integral over the luminosity function does not change significantly over the redshifts we consider.

A striking feature of Figure 3-5 is the apparent similarity of P_f along diagonal sets of different redshifts and models. For example, the “Cool IGM” model at $z = 12.2$ is very similar to the “Fiducial IGM” result at $z = 15.4$ and the “Hot IGM” at $z = 17.5$. It is suggestive that one can obtain the results of one particular thermal model by simply shifting another model in redshift, this translational invariance in redshift demonstrates that we may not need to simulate a broad range of heating models to understand the evolution of the forest power spectrum. Indeed, given our decomposition in Equation (3.15) where the amplitude of P_f is proportional to $\langle \tau_{21} \rangle^2 \propto \langle T_s^{-1} \rangle^2$, we should expect $\langle T_s \rangle$ to be a more generally applicable parameterization than f_X and redshift during the pre-reionization epoch. To show the importance of $\langle T_s \rangle$ as a parameter, we plot, in Figure 3-7, the amplitude of P_f at $k_{\parallel} = 0.5 \text{ Mpc}^{-1}$ as a function of $\langle T_s \rangle$ for our three heating scenarios and redshifts. Across all thermal models and redshifts, the amplitude of P_f is well described by a power law of $\langle T_s \rangle^{-2}$, consistent with the normalization predicted in Equation (3.15).

Verifying our prediction on the sign of $P_{f,b}$ is our next task; we plot this quantity in Figure 3-8 for all models and redshifts. At high redshift, $P_{f,b}$ is entirely negative due to the anti-correlation between T_f and T_b and adds to the total amplitude of P'_b . As heating takes place, T_s drops out of T_b and fluctuations in T_b are sourced predominantly by variations in x_{HI} leading to positive correlation between T_b and T_f for positive $P_{f,b}$. As we see from the figures, this process is “inside-out”, with large scales remaining anti-correlated longer than the small scales. Heating proceeds in an “inside-out” manner, and since there is an overlap between the completion of heating and onset of reionization, temperature fluctuations remain important on large scales (Pritchard and Furlanetto, 2007; Mesinger et al., 2013).

3.5 PROSPECTS FOR DETECTION WITH AN MWA-LIKE ARRAY

We now turn to addressing the detectability of the power spectrum signature of the forest and its distinguishability from the power spectrum, P_b . Our strategy is to combine our simulations with random realizations of instrumental noise and galactic and extragalactic foregrounds. With data cubes containing both our simulated signals and our random contaminants, we can then take advantage of the full quadratic estimator formalism developed by Tegmark (1997a), adapted for 21 cm tomography by Liu and Tegmark (2011); hereafter LT11, and accelerated for large data sets by Dillon et al. (2013); hereafter D13. In this section, we will explain those techniques and show what results when our simulations of the forest are added to realistic foregrounds and instrumental noise.

3.5.1 Power Spectrum Estimation Methods

To estimate the power spectrum of the forest, we apply the quadratic estimator formalism (Tegmark, 1997a). This formalism has the advantage that, in the approximation of foregrounds and noise that are completely described by their covariances,




Figure 3-8: The cross power spectrum, $\text{Re}(P_{f,b})$'s, sign is determined by the anti-correlation of x_{HI} and T_s during the pre-heating epoch and by x_{HI} after heating has taken place. Here we show the sign of $\text{Re}(P_{f,b})$ for our three different heating models as a function of redshift. At pre-heating redshifts, T_s is small and x_{HI} is relatively uniform so that T_b and T_f primarily depend on T_s and anti-correlate so that $\text{Re}(P_{f,b})$ is negative. At low redshifts, T_b is independent of T_s and fluctuations are primarily sourced by x_{HI} so that T_b and T_f are correlated and $\text{Re}(P_{f,b})$ is positive. Furthermore, heating proceeds in an “inside-out” manner so that the smallest scales become correlated first.

all cosmological information is preserved in going from three-dimensional data cubes to power spectra. This formalism was adapted by LT11 for 21 cm power spectrum estimation and further refined and accelerated by D13.¹⁰

In essence, the method relies an optimal and unbiased estimator of band powers in the k_{\perp} - k_{\parallel} plane, $\hat{\mathbf{p}}$, defined as

$$\hat{p}^{\alpha} = \sum_{\beta} M^{\alpha\beta} (\mathbf{x}^{\top} \mathbf{C}^{-1} \mathbf{Q}^{\beta} \mathbf{C}^{-1} \mathbf{x} - b^{\beta}). \quad (3.24)$$

where \mathbf{x} is a vector containing mean-subtracted data, \mathbf{C} is the covariance of \mathbf{x} , including noise and contaminants, \mathbf{Q} is a matrix that encodes the Fourier transforming, squaring, and binning necessary to calculate a band power, and \mathbf{b} is the bias term. The normalization matrix \mathbf{M} is related to the Fisher information matrix \mathbf{F} . Both \mathbf{F} and \mathbf{b} can be calculated via a Monte Carlo using the fact that

$$b^{\beta} = \langle \mathbf{x}^{\top} \mathbf{C}^{-1} \mathbf{Q}^{\beta} \mathbf{C}^{-1} \mathbf{x} \rangle \equiv \langle \hat{q}^{\beta} \rangle \quad (3.25)$$

and that

$$\mathbf{F} = \text{Cov}(\hat{\mathbf{q}}). \quad (3.26)$$

The ensemble average of each band power is related to the true band power \mathbf{p} by a window function matrix, $\mathbf{W} = \mathbf{M}\mathbf{F}$,

$$\langle \hat{\mathbf{p}} \rangle = \mathbf{W}\mathbf{p}. \quad (3.27)$$

The error on true band powers is also related to \mathbf{M} and \mathbf{F} through

$$\text{Cov}(\hat{\mathbf{p}}) = \mathbf{M}\mathbf{F}\mathbf{M}^{\top}. \quad (3.28)$$

Each quadratic estimator can thus be thought of as a weighted average of the true band powers with potentially correlated errors, both of which depend on one's choice

¹⁰For further details on this particular implementation of the quadratic estimator method, the reader is referred to D13.

of \mathbf{M} . Though any choice of \mathbf{M} that makes \mathbf{W} a properly normalized weighted average is reasonable, we adopt a form of \mathbf{M} that makes the errors on $\hat{\mathbf{p}}$ uncorrelated. Dillon et al. (2014), argue that this choice of \mathbf{M} dramatically reduces the contamination of the EoR window by residual foregrounds. It also provides a set of band power estimates which can be considered both mutually exclusive and collectively exhaustive because they cover the whole k_{\perp} - k_{\parallel} plane while not containing any overlapping information.

3.5.2 Noise and Foreground Models

The method outlined above requires model means and covariances of the contaminants that contribute to \mathbf{x} , like noise and foregrounds. Our model of the instrumental noise depends, first and foremost, on the design of the interferometer. In this paper, we consider the MWA with 128 tiles whose locations are detailed in Beardsley et al. (2012) as representative of the current generation of low frequency interferometers. Additionally, we consider possible realizations of double and quadruple sized instruments (MWA-256T and MWA-512T, respectively), as representative of extensions to current generation interferometers or next generation, $A_{eff} \sim 0.1\text{km}^2$, arrays such as the Hydrogen Epoch of Reionization Array (HERA) (Backer et al., 2009). As we will show, we generally do not need a square kilometer scale instrument to see the statistical effects of the forest.

To generate our MWA-256T and MWA-512T designs with maximum sensitivity to 21 cm cosmology, we add antenna tiles to the current MWA-128T design within a dense core 900 m in radius. These are drawn blindly from a probability distribution similar to that in Bowman et al. (2007): uniform for $r < 50$ m and decreasing as r^{-2} above 50 m. The tile locations of the arrays we use are shown in Figure 3-9.

Our model for the noise is adapted from D13¹¹. In it, we incorporate observation times calculated in each uv -cell from 1000 hours of rotation synthesis at the latitude

¹¹The method of D13 is adapted with one correction: the form of the noise power spectrum adapted from (Tegmark and Zaldarriaga, 2009) does not include the assumption that the field and beam sizes are the same.




Figure 3-9: Array layouts that we use to determine the detectibility and distinguishability of the 21 cm forest power spectrum signature. We chose to study two moderate extensions of the MWA-128T: MWA-256T and MWA-512T. In addition we study a 4096T array that is representative of a HERA scale instrument with ~ 400 times the collecting area of the MWA. Tile locations are drawn randomly from a distribution that is constant for the inner 50 m and drops as r^{-2} for larger radii.

Table 3.2: Instrumental Parameters

f (MHz)	FWHM (deg)	A_{eff} (m ²)	T_{sys} (K)
150	23	23	290
120	30	24	490
100	34	24	760
80	39	27	1300

of the MWA. The effective area of each tile is computed using a crossed dipole model while the system temperature is treated as the sum of receiver temperature, given by a power law fitted to two data points appearing in Tingay et al. (2013a), and sky temperature, measured in Rogers and Bowman (2008). In Table 3.2 we give our instrumental parameters at several different frequencies.

Similarly, our model of the foregrounds is the one application of the model developed by LT11 and D13. For the sake of simplicity,¹² we model extragalactic foregrounds as a random field of point sources with fluxes up to 200 Jy. They have an

¹²Breaking extragalactic foregrounds into a bright “resolved” population and a confusion-limited “unresolved” population only improves the error bars (D13), so our efficient choice is also a conservative one.

average spectral index of 0.5 and variance in their spectral indices of 0.5. Their clustering has a correlation length scale of $7'$. Likewise, we model Galactic synchrotron radiation as a random field with an amplitude of 335.4 K at 150 MHz, a coherence length scale of 30° , and a mean spectral index of 0.8 with an uncertainty in that index of 0.1.

As we have previously discussed, we conservatively cut out the region of k_\perp - k_\parallel space that lies below the wedge. Once the wedge has been excised, we optimally bin from 2D to 1D Fourier space with the inverse covariance weighted technique described by D13.

To create simulated observations, we divide our simulated volumes into 36 fields, each 750 Mpc on a side, which roughly fill the primary beam of our antenna tiles. We add random noise and foregrounds to each field independently, taking advantage of the fast technique for foreground and noise simulations developed by D13. Finally, we take the sample variance of the cosmic signal into account by using our power spectrum results from Section 3.4 and by counting the number of independent modes probed by the instrument at each k scale.

3.5.3 Detectability Results

We now present the results of our sensitivity calculation. We demonstrate that, given prior knowledge¹³ of the X-ray heating history, a power spectrum measurement with a modest expansion of an MWA-like instrument is sufficient to distinguish between scenarios with or without the forest in our fiducial and cool heating models. Since the forest signal is detectable with smaller arrays only at smaller k , where P_b dominates, its effect is likely degenerate with diffuse IGM emission. Observing this region for all considered models will require a HERA scale instrument with $A_{eff} \sim 0.1\text{km}^2$.

In order to determine the array size necessary to resolve the forest power spectrum, we first focus on $z = 11.2$, the lowest redshift considered where there is significant signal for one of our thermal models and quasar counts are relatively high. In Figure

¹³Here, "prior knowledge" means that we know what the IGM power spectrum without the 21 cm forest to within the error bars of our thermal sensitivity.




Figure 3-10: Detections (black dots) and upper limits (red triangles) of the 21 cm Power Spectrum at $z=11.2$ for all of our arrays and heating models in the presence of 21 cm forest absorption from background RL sources. The grey fill denotes the 2σ region around the measured power spectrum with no RL sources present. To determine whether we can detect the forest imprint we ask, “do the points and their error bars lie outside the gray shaded region?” MWA-256T and MWA-512T would be capable of distinguishing power spectra with or without sources in our cool IGM model, however only 4096T is consistently sensitive to the $k \gtrsim 0.5 \text{ Mpc}^{-1}$ region where the forest dominates. Only for our cool IGM model, MWA-512T would be sufficient to detect this upturn as well. Hence a moderate MWA extension would likely be able to constrain some RL populations given a cooler heating scenario while a HERA scale instrument will be able to constrain the W08 RL population using the Forest power spectrum even for more emissive heating scenarios. Note that the upturn in the gray region is not from increased power at high k but larger error bars.

3-10 we shade the 2σ region for a detection of $\Delta^2(k)$ with no 21 cm forest absorption present and mark detections of $\Delta^2(k)$ with 21 cm forest absorption with black dots. The 2σ vertical error bars, given by the diagonal elements of Equation (3.28) are marked in red. Also marked in red are the horizontal error bars which are given by the 20th and 80th percentiles of the window functions. To determine whether we can detect the forest imprint, we ask “are the points consistent with the gray shaded region?”

We see that MWA-256T and MWA-512T can distinguish cool models with and without the forest at greater than 2σ . However these detections are not within the region of Fourier space where the forest dominates P'_b . As a result, though MWA expansions can resolve two models with or without the forest, it is unlikely that they will be able to distinguish a model with the forest from one with a slight variation in heating. If an independent measure of the global spin temperature can be obtained, the radio luminosity function might be constrained with a modest MWA extension. We note that MWA-4096T is only able to detect the forest in our cool model at $z = 11.2$ since the optical depth in our more X-ray emissive models is far too small at this time.

To see more broadly what might be achieved by the next generation, we show in Figure 3-11 the error bars and detections with and without the forest across all considered f_X and z for our HERA scale model. We find that $z=15.4$ is our “sweet spot” for the W08 distribution. 4096T is able to resolve the $k \gtrsim 0.5 \text{ Mpc}^{-1}$ forest region for all of the IGM heating models that we investigate. For our cool and fiducial models, 4096T is also able to observe the forest region for a range of redshifts. These results show that a HERA scale array has the potential to constrain the IGM state by measuring Δ^2 for $k \lesssim 10^{-1} \text{ Mpc}^{-1}$, where the brightness temperature dominates, and the RL distribution in observing the region $k \gtrsim 0.5 \text{ Mpc}^{-1}$ where the forest has a significant contribution.

Over the course of the IGM’s evolution, there are times where the 21-cm power spectrum becomes particularly steep; for example, during the era immediately before the X-ray heating peak. As a result, observing excess power at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ for

a single redshift alone will likely not be sufficient to constrain the radio luminosity function. However, discerning the IGM thermal history with measurements of the power spectrum amplitude at $k \sim 0.1 \text{ Mpc}^{-1}$ and observing an absence of flattening at high k , over the range of redshifts after the X-ray heating peak as shown in Figure 3-4 should allow for constraints to be placed on the high-redshift radio luminosity function.

3.5.4 Distinguishability Results

In order to quantify how distinguishable our simulations with the forest are from our simulations without the forest for a given instrument, redshift, and heating model, we calculate the standard score of the χ^2 sum of the power spectrum values across all k -bins,

$$Z \equiv \frac{\chi^2 - N_k}{\sqrt{2N_k}}, \quad (3.29)$$

where N_k is the number of k bins, $\chi^2 \equiv \sum_k \left(\frac{P_b(k) - P_b(k)}{\sigma_k} \right)^2$, and σ_k are the diagonal elements of Equation (3.28) for each model without the 21 cm forest present. Assuming statistical independence between k bins, Z is the number of standard deviations at which we can distinguish a model with the 21 cm forest from a model without it using the χ^2 statistic. Unfortunately, this measure is somewhat naive since it does not account for potential degeneracies in the power spectrum amplitude from different thermal histories. However it enables us to quantitatively compare outlooks across the numerous dimensions of redshift, array, and heating history. We consider a $Z \gtrsim 10$ to indicate significant distinguishability.

In Figure 3-12 we show the value of Equation (3.29) for all models and arrays. Our first observation is that MWA-128T is not capable of distinguishing a model with the forest from a model without the forest for any of the considered f_X . MWA-256T would be capable of distinguishing the forest at all considered $z \gtrsim 9.2$ for our cool X-ray heating model at greater than 5σ and in our fiducial heating model only at the highest considered redshift (which is near the X-ray heating peak). MWA-512T would be capable of resolving the forest at the two highest redshifts for our




Figure 3-11: These plots are identical to Figure 3-10 except the array is fixed to be MWA-4096T, representative of a HERA generation instrument, and redshift is varied. A HERA class instrument is able to resolve the upturn at $k \gtrsim 0.5 \text{ Mpc}^{-1}$ that distinguishes the forest, and should be able to detect the 21 cm forest feature considered in this work for a variety of heating scenarios. The thermal noise error bars are too small to resolve by eye in most of these plots.

fiducial model and at all considered redshifts for our cool model. The hot model remains unobservable for all MWA expansion arrays but is accessible to a HERA scale instrument.

How the distinguishability between different heating models is affected by the presence of the 21 cm forest is explored in Figure 3-13. In our 128T table, we see that a detection of the IGM and constraints on low X-ray emissive histories are possible with the current generation of EoR experiments. There are several caveats worth noting however. First, the high S/N distinctions at $z = 9.2$ are due to a detection of the reionization peak at redshifts in which reionization physics such as the uv-efficiency (which we have assumed fixed) become significant. However, we note that this result contradicts the marginal detectability claimed in Mesinger et al. (2013) primarily due to the fact that we include bins with $k < 0.1 \text{ Mpc}^{-1}$ in our standard score. Though these bins have large S/N they may be contaminated by more pessimistic foreground leakage than we consider here such as what is observed by Pober et al. (2013a). We also note that the increased sensitivity of combining k-bins allow for constraints on the fiducial X-ray model at $z \sim 15$. The peaks in detectability at $z \approx 9$ and $z \approx 17$ arise from the two peaked structure of the power spectrum in redshift with the low redshift peak corresponding the reionization, and the high redshift peak corresponding to x-ray heating (Pritchard and Furlanetto, 2007). We see that the forest introduces a small enhancement to the distinguishability between hot and cool heating models. Since the forest adds positively to the power spectrum of a cool, optically thick IGM, its presence enhances the distinguishability between vigorous and cool heating. We find that a modest extension to the MWA can distinguish between hot and fiducial models over a wider range of redshifts and MWA-4096T is able to distinguish between all models over our entire considered redshift range.




Figure 3-12: The significance of distinguishability across all measured k bins (Equation (3.29)) for all arrays, redshifts, and IGM heating models for a 1000 hour observation. An extension of MWA-128T is capable of distinguishing models with and without the 21 cm forest from the W08 RL population in our cool and fiducial heating scenarios. MWA-512T and HERA scale MWA-4096T are capable of distinguishing the forest in the power spectrum in all heating models considered in this work.




Figure 3-13: The 21 cm Forest moderately enhances the distinguishability between thermal scenarios and MWA scale interferometers can distinguish between the power spectra for reasonable X-ray heating histories. Here we show the cumulative z-score described in Equation (3.29), except now applied to the difference between different IGM heating models, for all arrays and redshifts. At low redshift, the forest decreases the distinguishability of different X-ray heating scenarios by subtracting from the higher amplitude model. When the positive auto-term dominates at high redshift, the forest increases the contrast between given heating models.

3.6 The Detectability of the Forest over a Broad Parameter Space

For the sake of simplicity, we focus on the detectability of the 21 cm Forest power spectrum from the single population model considered in Wilman et al. (2008). In doing this, it is unclear over what range of radio loud populations the signal is observable. Fortunately, thanks to Equation (3.23), we can give order of magnitude estimates of how the detectability of the Forest power spectrum scales with the radio loud source population and the heating history. According to Equation (3.23), the amplitude of the forest power spectrum, at pre-reionization redshifts, scales as

$$P_f \propto \frac{1}{\langle T_s \rangle^2} \frac{\sum_i s_i^2(> z)}{\Omega} \quad (3.30)$$

where $\sum_i s_i^2(> z)\Omega^{-1}$ is the average sum of source fluxes squared per solid angle. We will call this quantity the *flux squared density* of the source population. We take advantage of the simple scaling in Equation (3.30) to extrapolate the amplitude of the Forest signal over a large range of heating models and redshifts. At each redshift, with our fiducial heating model and source population, we obtain a normalization factor for P_f at a single mode, $k = 0.5 \text{ Mpc}^{-1}$. We then compute $\langle T_s \rangle$ for a large number of lower resolution, $(600 \text{ Mpc})^3$ 21cmFAST simulations with 400^3 pixels, varying the f_X parameter by three orders of magnitude from $f_X = 10^{-2} - 10^1$. In Figure 3-14, we show the ratio of P_f to the amplitude of thermal noise as a function of f_X and the flux squared density of sources, marking the predicted flux squared density of Wilman et al. (2008) by a dashed black and white line. We find that the detectability of the forest power spectrum at $z \sim 10$ depends strongly on the thermal state of the IGM, with models significantly fainter than Wilman et al. (2008) undetectable except for cool heating histories with $f_X \lesssim 10^{-1}$. On the other hand, for $z \gtrsim 15$, X-rays in all models have not had sufficient time to heat the IGM above the adiabatic cooling floor and the detectability of P_f becomes significantly less dependent on f_X , allowing for a broader range of populations to be probed at higher f_X .




Figure 3-14: The ratio of the 21 cm Forest power spectrum, $P_f(k = 0.5 \text{ Mpc}^{-1})$ to thermal noise for 1000 hours of observation on a HERA scale interometer, extrapolated over a large range of X-ray efficiencies and flux squared densities. Vertical dashed black and white lines indicate the value of the simulation by (Wilman et al., 2008) while the horizontal black and white lines indicate the f_X efficiencies that we explicitly simulate in this paper. At the highest redshifts, $\langle T_s \rangle$ levels off and the detectability of the signal is independent of redshift. At late preionization redshifts, we see that the 21 cm Forest will only be detectable for heating efficiencies $\lesssim 1$.

3.7 Conclusions and Future Outlook

Using semi-numerical simulations of the thermal history of the IGM, and a semi-empirical RL source distribution, we have shown that the 21 cm forest imprints a distinctive feature in the power spectrum that is, for the most part, invariant in k_{\perp} and, depending on the RL population and thermal history, potentially dominates over the cosmological 21 cm power spectrum at $k_{\parallel} \gtrsim 0.5 \text{ Mpc}^{-1}$. We have also derived a simple semi-analytic equation that directly relates the forest power spectrum of τ_{21} and the radio luminosity function.

Using realistic simulations of power spectrum estimation and including the effects of foregrounds and noise, we have shown that a moderate extension of the MWA-128T instrument has the thermal sensitivity to detect the forest feature in the power spectrum for the W08 RL population with an X-ray efficiency of $f_x \lesssim 1$. For more vigorous heating scenarios, a HERA scale array will have the sensitivity to distinguish this feature. Our simulations also support the results of Christian and Loeb (2013) and Mesinger et al. (2013), that low emissivity heating scenarios can be constrained with existing arrays and an extensive examination of the heating history will be possible in the future with larger instruments.

Signal-to-noise considerations alone do not tell us whether we will be able to distinguish the forest signal from the effects of IGM physics on the power spectrum, especially at small k where a slight change in f_x might shift the power spectrum amplitude up or down, mimicking the shift from the 21 cm forest. Fortunately, the region, $k \gtrsim 0.5 \text{ Mpc}^{-1}$ is dominated by the forest power spectrum, P_f , for a range of redshifts in all of our heating models. Specifically, the 21 cm forest removes the $k \gtrsim 0.5 \text{ Mpc}^{-1}$ flattening that occurs after the X-ray heating peak. Observations of the power spectrum over a range of redshifts, with a sensitivity similar to HERA or the SKA should be able to isolate the thermal history at $k \lesssim 0.1 \text{ Mpc}^{-1}$ and constrain RL populations similar to that of W08 at $k \gtrsim 0.5 \text{ Mpc}^{-1}$.

While this paper is a proof of concept, considering a single fiducial RL source distribution, it is possible that measurements with current generation instruments,

or moderate extensions, can put constraints on more optimistic scenarios. On the other hand, there are many steep decline scenarios whose power spectrum signatures will be inaccessible even to future arrays. In section 3.6 we illustrate the scaling of the detectability of the signal with source flux squared density and X-ray emissivity, finding that populations with order of magnitude smaller flux squared densities than W08 will require a relatively cool preionization IGM to be detectable. In particular, we note that the H04 simulation is one to two orders of magnitude more pessimistic than the predictions of W08 at the highest considered redshifts and would not be detectable in the forest dominated region if $f_X \gtrsim 10^{-1}$. However, higher resolution simulations of the IGM indicate that $\Delta_{\tau_{21}}^2$ continues to climb to $k \sim 10 \text{ Mpc}^{-1}$ while P_b remains flat. Hence the result of a fainter radio luminosity function would be to shift the region of forest dominance to higher k rather than eliminating it, leaving the possibility of detection for a more powerful instrument such as the SKA. There also exists the possibility of separating P_f using its LoS symmetry which might be exploited at $k \sim 0.1 \text{ Mpc}^{-1}$ where EoR interferometers are most sensitive. Finally, we have not considered the absorption of mini halos which Mack and Wyithe (2012) show to substantially increase the variance along the line of sight towards sources (see their Figure 11). Since this variance is an integral of the power spectrum we are being conservative in neglecting them. The sensitivity of future instruments to the forest can be enhanced by increased frequency resolution, allowing them to probe the higher k_{\parallel} modes where the forest is especially strong. The parameter space of radio loud quasars is greatly unconstrained and the disparity between W08 and H04 simply underscores the need for future studies to explore this parameter space. The exploration of a range of RL populations for fixed arrays is left for future work.

In summary, we have shown that the 21 cm power spectrum not only contains information on the IGM in absorption and emission against the CMB but also includes detectible, and in many cases non-negligible signatures of the 21 cm forest. This absorption may be used to constrain the high redshift RL population and IGM thermal history with upcoming interferometers.

%sectionAcknowledgements

3.A Appendix: A Derivation of the morphology of P_f

In Section 3.2 we present a formula, Equation (3.15), for the the 21 cm forest power spectrum that is the sum of the auto power spectra along the line of sight to each background source. This equation is particularly convenient because it can easily be decomposed into an integral of the radio luminosity function and the optical depth power spectrum. In addition, its k -space morphology, which includes no structure in k_{\perp} , is relatively simple. In this appendix we derive Equation (3.15) by applying an analytic toy model to the auto and cross power spectrum contributions to P_f described in Equation (3.13). For the sake of analytic tractability, we invoke a number of approximations. However our results describe P_f very well for $k \gtrsim 10^{-1} \text{ Mpc}^{-1}$. Our assumptions are

1. The sources all have the same flux. The W08 simulation includes sources ranging from 1 nJy to ~ 10 mJy over the redshifts of interest. We see in Figure 3-2 that the integral of the source fluxes squared is dominated (at the 10% level) by sources with S_{ν} between 1 – 10 mJy so modeling our population as having equal flux gives a decent order of magnitude approximation.
2. The sources are spatially uncorrelated. Clustering from the W08 dark matter bias is actually significant and boosts the results of our simulation, relative to Equation (3.15), by a factor of two without changing P_f 's predicted shape. We will thus absorb this clustering boost into a multiplicative factor of order unity.
3. The sources are unresolved. This will almost certainly be true in all interesting cases given the large synthesized beams of radio interferometers and the extreme distances to the sources.
4. Source spectra are flat over the frequency interval of a data cube. This is true on the 10% level over a ~ 8 MHz band for $S \sim \nu^{-.75}$ sources. Because this slow variation gives a very narrowly peaked convolution kernel in k -space, power spectra are not noticeably effected by this assumption.

5. The source positions are completely uncorrelated with the cube optical depth field. In reality, the sources that fall within a data cube should be correlated with τ_{21} . We find that correlating or not correlating in cube sources only changes the simulation output by approximately 10%.

We start by reiterating Equation (3.14) where P_f may be written as

$$\begin{aligned}
P_f &= \frac{1}{V} \left\langle \left| \widetilde{\Delta T_{RL} \tau_{21}} \right|^2 \right\rangle \\
&= \sum_j P_j + 2\text{Re} \left(\sum_{j < k} P_{j,k} \right) \\
&\equiv \Sigma_{auto} + \Sigma_{cross},
\end{aligned} \tag{3.31}$$

where $P_j = \frac{1}{V} \langle |\widetilde{\Delta T_j \tau_{21}}|^2 \rangle$ and $P_{j,k} = \frac{1}{V} \langle \widetilde{\Delta T_j \tau_{21}} \widetilde{\Delta T_k \tau_{21}}^* \rangle$. The first term in Equation (3.31) sums the power spectra of each of the absorbed background sources which is positive and the second term is the sum of their cross power spectra.

We will show that for the range of spatial scales perpendicular to the LoS, accessed by EoR interferometers, the auto power terms in Equation (3.31) dominate the cross power ones at $k_{\parallel} \gtrsim 10^{-1} \text{ Mpc}^{-1}$. We show that the suppression of cross terms is due to two mechanisms: (1) the cross terms are proportional to the cross power spectra between widely separated lines of sight and (2) the cross terms are multiplied by randomly phased sinusoids which cancel out when summed.

3.A.1 Appendix: The Suppression of the Cross terms from LoS Cross Power Spectra

To relate the sum in Equation (3.31) to the spectra and locations of the background sources, we assume that all sources are unresolved so that T_j is a delta-function in the plane perpendicular to the LoS. Here, as in McQuinn et al. (2006), we will adopt observers coordinates (ℓ, m, ν) , rather than comoving coordinates (x, y, z) , to emphasize the fact that the the broad-spectrum source does not physically occupy a range of positions along the LoS. In such coordinates, the temperature field of each

source can be written as $T_j(\ell, m, \nu)$ where ℓ and m are the direction cosines from the north-south and east-west directions, and ν is the difference from the data cube's central frequency. $\tau_{21}T_j(\ell, m, \nu)$ is given by

$$\tau_{21}T_j(\ell, m, \nu) = \Omega_{pix}\delta(\ell - \ell_j)\delta(m - m_j)\tau_{21}(\ell_j, m_j, \nu)T_j \quad (3.32)$$

where Ω_{pix} is the solid angle of a map pixel and $\delta(\dots)$ is the Dirac delta function. For notational simplicity, we will use vector notation to denote direction cosines, $\ell = (\ell, m)$ and their Fourier duals, $\mathbf{u} = (u, v)$. Taking the Fourier transform of $\tau_{21}T_j(\ell, \nu)$ and summing over all sources gets

$$\begin{aligned} \widetilde{T}_f(\mathbf{u}, \eta) &= \sum_j \widetilde{\tau_{21}T_j}(\mathbf{u}, \eta) \\ &= \Omega_{pix} \sum_j T_j e^{2\pi i(\ell_j \cdot \mathbf{u})} \int \tau_{21}(\ell_j, \nu) e^{-2\pi i\eta\nu} d\nu. \end{aligned} \quad (3.33)$$

We take the modulus squared of Equation (3.33) and multiply by the cosmology dependent variables, $D_M^2 Y$ (Parsons et al., 2012a) that relate observers coordinates to the cosmological comoving coordinates that we've used to define our power spectrum in Equation (3.2). We find that the sum of the auto terms in Equation (3.31) is

$$\Sigma_{auto} = \frac{D_M^2 \Omega_{pix}^2}{\Omega_{cube}} P_{\tau_{21}}^{LoS}(k_{\parallel}) \left\langle \sum_j T_j^2 \right\rangle. \quad (3.34)$$

The sum of cross terms is

$$\begin{aligned}
\Sigma_{cross} &= 2 \frac{D_M^2 \Omega_{pix}^2}{\Omega_{cube}} \sum_{j < k} T_j T_k \left[\right. \\
&\quad \text{Re} \left(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}) \right) \langle \cos[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \\
&\quad \left. + \text{Im} \left(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}) \right) \langle \sin[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right] \\
&= 2 P_{\tau_{21}}^{LoS}(k_{\parallel}) \frac{D_M^2 \Omega_{pix}^2}{\Omega_{cube}} \sum_{j < k} T_j T_k \left[\right. \\
&\quad \frac{\text{Re} \left(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}) \right)}{P_{\tau_{21}}^{LoS}(k_{\parallel})} \langle \cos[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle + \\
&\quad \left. \frac{\text{Im} \left(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}) \right)}{P_{\tau_{21}}^{LoS}(k_{\parallel})} \langle \sin[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right], \tag{3.35}
\end{aligned}$$

where $\Delta\ell_{j,k} = \ell_j - \ell_k$. Here, we define the cross power spectrum between two lines of sight to be

$$P_{\tau_{21};j,k}^{LoS}(k_z) = \frac{1}{L} \int dz dz' e^{ik_z(z-z')} \Delta\tau_{21}(\ell_j, z) \Delta\tau_{21}(\ell_k, z'). \tag{3.36}$$

It is clear from Equation (3.35) that each summand in Σ_{cross} is smaller than each term in Σ_{auto} by a factor of the ratio between the LoS cross power spectra of spatially separated lines of sight and the LoS auto power spectrum. If lines of sight to each source are sufficiently separated, this ratio should be very small. In Figure 3-15 we show the ratios of $\text{Re} \left(P_{\tau_{21};j,k}^{LoS} \right) / P_{\tau_{21}}^{LoS}$ and $\text{Im} \left(P_{\tau_{21};j,k}^{LoS} \right) / P_{\tau_{21}}^{LoS}$ from our fiducial model at $z = 12.2$, separated by $L_{\perp} = 24$ Mpc which is the mean distance in our data cube between 1000 background sources. Because two sufficiently separated lines of sight should be statistically independent except on the largest spatial scales, these ratios are on the order of $10^{-2} - 10^{-3}$ for $k_{\parallel} \gtrsim 10^{-1} \text{ Mpc}^{-1}$.




Figure 3-15: The LoS cross power spectra between spatially separated lines of site are on the order of $\sim 100 - 1000$ times smaller than auto power spectra. In the left figure, we plot the ratio of the real cross power spectra between lines of site separated by 24 Mpc to auto power spectra, and on the right we show the ratio of the imaginary cross power spectrum to the auto power spectrum. In both cases, for $k_{\parallel} \gtrsim 10^{-1} \text{ Mpc}^{-1}$, the cross power spectra are on the order of 10-1000 times smaller. The real cross power spectrum becomes non negligible on scales comparable to the separation between the two lines of site.

3.A.2 Appendix: Suppression of the Cross Terms from Summing the Random Source Phases

The factor of 100-1000 introduced by the ratio of the cross spectra to the auto spectra would be enough to suppress the cross terms if the number of sources were reasonably small. However the number of cross terms relative to the number of auto terms in Equation (3.31) goes as $(N - 1)/2$ where N is the number of contributing sources. Thus, even though the cross power spectrum between individual LoS pairs is small, naively summing 100-500 sources could still yield a significant contribution. We now show that summing over many randomly distributed source angles suppresses this.

Since $\text{Im} \left(\frac{P_{\tau_{21};j,k}^{LoS}}{P_{\tau_{21}}^{LoS}} \right)$ is on the same order of, or smaller than $\text{Re} \left(\frac{P_{\tau_{21};j,k}^{LoS}}{P_{\tau_{21}}^{LoS}} \right)$, we will use the real term on both the cosine and sine terms in Equation (3.35) to give an upper bound. Assuming that all sources have the same temperature, $T_j = T_k = T_0$,

we may write

$$\begin{aligned} \Sigma_{cross} \approx 2 \sum_{j < k} T_0^2 P_{\tau_{21};j,k}^{LoS}(k_{\parallel}) \left[\langle \cos[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right. \\ \left. + \langle \sin[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right]. \end{aligned} \quad (3.37)$$

Similarly,

$$\Sigma_{auto} \approx NT_0^2 P_{\tau_{21}}^{LoS} \quad (3.38)$$

Hence the ratio between Σ_{cross} and Σ_{auto} is given by

$$\begin{aligned} \Sigma_{cross}/\Sigma_{auto} \approx 2 \frac{\text{Re}(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}))}{NP_{\tau_{21}}^{LoS}(k_{\parallel})} \sum_{j < k} \left[\langle \cos[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right. \\ \left. + \langle \sin[2\pi(\mathbf{u} \cdot \Delta\ell_{j,k})] \rangle \right] \end{aligned} \quad (3.39)$$

Because of the cylindrical symmetry, we need only concern ourselves with a uv cell at $v = 0$ and simply write

$$\begin{aligned} \Sigma_{cross}/\Sigma_{auto} \approx 2 \frac{\text{Re}(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}))}{NP_{\tau_{21}}^{LoS}(k_{\parallel})} \sum_{j < k} \left[\langle \cos[2\pi u_{\perp} \Delta\ell_{j,k}] \rangle \right. \\ \left. + \langle \sin[2\pi u_{\perp} \Delta\ell_{j,k}] \rangle \right] \\ = \frac{\text{Re}(P_{\tau_{21};j,k}^{LoS}(k_{\parallel}))}{P_{\tau_{21}}^{LoS}(k_{\parallel})} \langle \Sigma_{cos}(u_{\perp}, \Theta, N) \rangle, \end{aligned} \quad (3.40)$$

where

$$\Sigma_{cos}(u_{\perp}, \Theta, N) \equiv \frac{2}{N} \sum_{j < k} \cos[2\pi u_{\perp} \Delta\ell_{j,k}] + \sin[2\pi u_{\perp} \Delta\ell_{j,k}]. \quad (3.41)$$

We can easily compute this ensemble average for any u_{\perp} by drawing N different source positions distributed randomly over the angular span of the field, Θ , and summing over the sines and cosines of pair-wise angle differences. In Figure 3-16 we show $P[\Sigma_{cos}(u_{\perp}, \Theta, N)]$ for randomly distributed $\Delta\ell_{j,k}$ for a variety of N , u_{\perp} , and Θ where the minimal u_{\perp} is set by the maximal scale accessible by an interferometers primary beam, $\sim 1/\Theta$. We calculate these distributions from 10000 random realizations. We




Figure 3-16: Here we see that $P[\Sigma_{cross}(u_{\perp}, \Theta, N)]$ is invariant in N , Θ , and u_{\perp} , and N with a mean of approximately zero. The lines which indicate, $P[\Sigma_{cross}(u_{\perp}, \Theta, N)]$, are estimated from 10000 draws. Since $\langle \Sigma_{cross} \rangle \approx 0$ we expect the cross terms to contribute negligibly to P_f in 3D Fourier Space.

see that the distribution of Σ_{cos} is independent of N , Θ , and u_{\perp} and has a mean of ≈ 0 (which is the quantity that sets the amplitude of Σ_{cross}). As long as sources are randomly distributed, we can expect LoS cross power spectra to suppress the cross terms sum to below the 10% level at $k_{\parallel} \gtrsim 10^{-1} \text{ Mpc}^{-1}$, regardless of the number of terms.

We may finally write.

$$\begin{aligned}
 P_f(\mathbf{k}) &\approx \sum_j P_j(k_{\parallel}) \\
 &= \frac{D_M^2 \Omega_{pix}^2}{\Omega_{cube}} \sum_j T_j^2 P_{\tau_{21}}^{LoS} \\
 &= \frac{D_M^2 \lambda^4}{4k_B^2 \Omega_{cube}} \sum_j s_j^2 P_{\tau_{21}}^{LoS} \tag{3.42}
 \end{aligned}$$

where $\lambda = \lambda_{21}(1+z)$ is the wavelength at the center of the data cube, P_j is the absorption power spectrum for the j^{th} source, s_j and T_j are the flux and temperatures of the j^{th} source, Ω_{cube} is the solid angle subtended by the observed volume, and $P_{\tau_{21}}^{LoS}$ is the 1D LoS power spectrum.

We may therefor consider the absorption power spectrum resulting from the forest as simply the sum of the absorption power spectra of each individual source in the background of the source cube. Since all quantities in this sum are positive, we

see that the amplitude of the power spectrum increases linearly with the number of sources present behind an observed volume. Because the power spectra for unresolved sources are constant in k_{\perp} , P_f will have a structure that is nearly constant in k_{\perp} .

Hence, for $k \gtrsim 10^{-1} \text{ Mpc}^{-1}$, Equation (3.14) simplifies to a sum of the auto power spectra along the LoS to each source. We finish by briefly commenting on the effect of clustering which we have ignored but we find (after comparing Equation (3.15) to our simulations) is still significant. Clustering will cause a disproportionate number of sources to reside in close proximity on the sky. The effect of this is two fold. First, the clustered sources will tend to be behind correlated optical depth columns so that the cross terms between such sources will be better described by auto power spectra. Second, the phases between such sources will be small so that they will not sum to zero. In addition, they will not introduce significant k_{\perp} structure except at the smallest perpendicular scales. Hence the cross terms introduced by clustered sources will closely resemble the k_{\parallel} invariant auto terms and simply increase the overall amplitude of P_f . We treat this increase by introducing a multiplicative constant of order unity, A_{cl} , in Equation (3.23).

3.B Appendix: A Comparison between two source models

In this paper, we choose to work with the semi-empirical source population in the simulation by Wilman et al. (2008). This choice was in part motivated by the lack of constraints at high redshift and the ease which we could use data from the W08 simulation using its online interface. Another prediction in the literature for the high redshift radio luminosity function is made by Haiman et al. (2004). This model, like the one in W08, relies on a number of uncertain assumptions but is a more physically motivated bottom up approach which is derived from the cold dark matter power spectrum and assumptions about the black hole-halo mass relation and radio loud fraction. In this appendix, we attempt to understand how our choice of the Wilman

source population compares to that in H04. To do this, we attempt to compare the source counts from W08 that contribute the most to P_f to those of H04 who provide cumulative flux counts for 1 – 10 GHz as a function of redshift. To compare the W08 sources, we compute the percentage of the radio luminosity function integral in Equation (3.16) as a function of the extrapolated $S_{5 \text{ GHz}}$. On the left, in Figure 3-17, a large fraction of P_f is determined by W08 sources with 5 GHz fluxes between 10 μJy and 10 mJy. We show, in Figure 3-17, the ratio of W08 and H04 source counts with $S_{5 \text{ GHz}}$ between 10 μJy and 10 mJy. The H04 counts fall much faster with redshift than those of W08. At $z \sim 10 - 12$ the number of contributing sources is larger by a factor of ≈ 10 and ≈ 80 by $z \sim 16$.

This comparison is very approximate since different spectral indices are assumed in H04 and W08. However, we emphasize that the observability claims we make in this paper would not apply accurately to the H04 prediction. A more extensive exploration of parameter space will be necessary to determine what range of radio loud source populations may be constrained by the power spectrum technique.

Since P_b is observed to be flat out to $k \approx 10 \text{ Mpc}^{-1}$ while P_f climbs, a more pessimistic source scenario has the effect of pushing the forest dominant region to higher k_{\parallel} which does not preclude detection with a more powerful telescope such as the SKA.




Figure 3-17: Left: The percentage of the integrated luminosity function in Equation (3.16) as a function of the source fluxes at 5GHz for comparison to the catalogue of H04. We see that most contributions to the forest power spectrum come in between $S_{5\text{GHz}} = 10 \mu\text{Jy}$ and $S_{5\text{GHz}} = 10 \text{mJy}$. Right: The ratio of the number of sources with redshift greater than z between $S_{5\text{GHz}} = 10 \mu\text{Jy}$ and 10mJy as predicted by the W08 and H04. The W08 simulation over predicts the counts in H04 by a factor ten at $z \gtrsim 12$ and nearly 80 at $z \gtrsim 16$, emphasizing the importance of exploring this widely unconstrained parameter space in future work.

Chapter 4

Constraining High Redshift X-ray Sources with Next Generation 21 cm Power Spectrum Measurements

The content of this chapter was published in *Ewall-Wice, A., Hewitt, J., Mesinger, A., Dillon, J. S., Liu, A. and Pober, J. (2016), 'Constraining high-redshift X-ray sources with next generation 21-cm power spectrum measurements', MNRAS 458, 2710–2724.*

4.1 Introduction

Observations of 21 cm emission from neutral hydrogen at high redshift will provide us with a unique window into the cosmic dark ages, the birth of the first stars, and the reionization of the IGM (see Furlanetto et al. (2006a); Morales and Wyithe (2010); Pritchard and Loeb (2012) for reviews). Current experiments aiming to detect the cosmological 21 cm line currently follow two tracks. The first involves accessing the sky averaged “global signal” with a single dipole which is being attempted by experiments such as EDGES (Bowman and Rogers, 2010), LEDA (Greenhill and Bernardi, 2012; Bernardi et al., 2015), DARE (Burns et al., 2012), SciHi (Voytek et al., 2014), and BIGHORNS (Sokolowski et al., 2015). Alternatively, one can observe spatial fluctuations in emission with a radio interferometer and new generation of instruments

are already collecting data with the aim of making a first statistical detection. These include the MWA (Tingay et al., 2013a), PAPER (Parsons et al., 2010), LOFAR (van Haarlem et al., 2013), and the GMRT (Paciga et al., 2013). Already, a number of upper limits have been established over the redshifts at which reionization is expected to have occurred (Dillon et al., 2014; Parsons et al., 2014; Jacobs et al., 2015; Ali et al., 2015; Dillon et al., 2015a). This first generation of instruments may possess the sensitivity to make a low signal to noise detection of the power spectrum during the Epoch of Reionization (EoR). However, planned experiments such as the Square Kilometer Array (SKA) and the Hydrogen Epoch of Reionization Array (HERA) will be capable of constraining the parameters in reionization models to several percent precision (Pober et al., 2014; Greig and Mesinger, 2015).

The 21 cm line is an extremely rich observable, extending far beyond the start of reionization into the dark ages. The EoR is only the final chapter in the remarkable story that it encodes. Before reionization, X-rays emitted from the first generations of high mass X-ray binaries (HMXB) (Mirabel et al., 2011) and/or the hot interstellar medium (ISM) (Pacucci et al., 2014) reversed the adiabatic cooling of the IGM, likely bringing it into emission against the cosmic microwave background (CMB). Experiments will observe the brightness temperature of 21 cm emission against the CMB which is given by (Furlanetto et al., 2006a)

$$\delta T_b \approx 9.2 x_{HI} (1 + \delta) \left[1 - \frac{T_{cmb}}{T_s} \right] \left[\frac{H(z)/(1+z)^{1/2}}{dv_{\parallel}/dr} \right] \text{mK}, \quad (4.1)$$

where x_{HI} is the neutral fraction, δ is the fractional baryon overdensity, T_{cmb} is the temperature of the CMB at the redshift of emission (z), T_s is the spin temperature of the HI gas, $H(z)$ is the hubble parameter at redshift z , and dv_{\parallel}/dr is the radial velocity gradient. At the end of adiabatic cooling, the $(1 - T_{cmb}/T_s)$ factor is relatively large and negative (Furlanetto, 2006b). The first sources heating the IGM with X-rays introduce large dynamic range in this factor that is significantly greater in absolute separation than the permitted $[0, 1]$ in x_{HI} . This leads to a power spectrum that is roughly an order of magnitude larger than during reionization except in extreme

heating models (Pritchard and Furlanetto, 2007; Santos et al., 2008; Baek et al., 2010; Mesinger et al., 2013).

Recent work has shown that inefficient heating scenarios might be constrained with lower redshift measurements on current experiments (Christian and Loeb, 2013; Mesinger et al., 2014). Indeed, the latest PAPER upper limits have ruled out an IGM with a spin temperature below ≈ 10 K (Pober et al., 2015; Greig et al., 2016). While X-rays have some effect on the ionization field (Mesinger et al., 2013), their primary impact is encoded in T_s in all but the most extreme cases. Assuming that the linear relation between X-ray luminosity and star formation rate is similar to what is observed locally. T_s is expected to be much greater than T_{cmb} , during reionization, eliminating its impact on δT_b (Furlanetto, 2006b). Hence, to precisely determine the physics of IGM heating, measurements at redshifts higher than reionization may be needed. What limits on heating are possible at low redshift and to what extent measurements of the pre-reionization epochs are necessary to establish precision limits are open questions.

In this paper, we use semi-numerical simulations and the Fisher matrix formalism to determine the limits that next generation experiments will put on the properties of X-ray sources during the cosmic dawn. In particular, we focus on what we might expect to learn from observations of the reionization epoch alone and what measurements at higher redshifts might add to our knowledge. We also explore what additional reionization constraints might exist at pre-reionization redshifts. Earlier studies of constraints on reionization parameters (Pober et al., 2014; Greig and Mesinger, 2015) assume that $T_s \gg T_{\text{cmb}}$ which Mesinger et al. (2013) have shown can significantly under-predict the power spectrum amplitude early on and before reionization. Back-lit by a negative $(1 - T_{\text{cmb}}/T_s)$, there may exist additional detectable reionization signatures at higher redshift.

This paper is organized as follows. In § 4.2 we describe the Fisher matrix formalism which we use to connect uncertainties in power spectrum measurements error bars on IGM heating properties. In § 4.3 we describe our semi-numerical simulations along with our model parameters and their fiducial values. We describe our calculations

of thermal noise in § 4.4 and discuss the arrays examined. In § 4.5 we examine the information encoded in the derivative of the power spectrum with respect to each astrophysical parameter. We discuss the degeneracies between parameters and our predictions for model limits in § 4.6 and conclude in § 4.7.

4.2 From Measurements to Constraints

Today, the quantity that interferometers are attempting to estimate is $\Delta^2(k)$, the power spectrum of 21 cm brightness temperature fluctuations. The power spectrum at comoving wavenumber $k = |\mathbf{k}|$ is defined in terms of the spatial Fourier transform of the brightness temperature field, $\tilde{\delta}_b(\mathbf{k})$.

$$\langle \tilde{\delta}_b(\mathbf{k}) \tilde{\delta}_b(\mathbf{k}') \rangle = \delta_D(\mathbf{k} - \mathbf{k}') \frac{2\pi^2}{k^3} \Delta^2(k), \quad (4.2)$$

where $\langle \rangle$ denotes the ensemble average over all realizations of the brightness field and δ_D is the Dirac delta function.

21 cm observations of Δ^2 can be connected to theory and simulations by performing a maximum likelihood (ML) analysis to obtain the best-fit parameters for an astrophysical model. The uncertainty on the parameters obtained from such an analysis may be approximated using the Fisher matrix formalism which we now describe.

The Fisher matrix describes the amount of information, contained within a data set, on the parameters of a model. It can be defined through the derivatives of the ln-likelihood about the ML parameter values (Fisher, 1935).

$$F_{ij} \equiv - \left\langle \frac{\partial^2 \ln \mathcal{L}}{\partial \theta_i \partial \theta_j} \right\rangle. \quad (4.3)$$

Here, \mathcal{L} is the likelihood of observing the outcome of a measurement given a set of model parameters θ . In our case the measurement is the set of power spectrum values in each Fourier and redshift bin while θ is the set of parameters in our astrophysical model. Intuitively, we see in equation 4.3 that the largest amount of information exists for parameters that cause the greatest change in the likelihood.

If the likelihood function is Gaussian, which is usually a good approximation in the case of small errors about the ML point, then the covariance matrix of these parameters, C_{ij} , is simply the inverse of the Fisher matrix,

$$C_{ij} = (\mathbf{F}^{-1})_{ij}. \quad (4.4)$$

The more strongly \mathcal{L} depends on θ_i and θ_j , the smaller the covariance matrix values and the error bars on our estimates.

Judicious choices in constructing the estimate of the power spectrum from visibility data can ensure that the Fourier modes and redshift bins at which the power spectrum is estimated are uncorrelated (Liu and Tegmark, 2011; Dillon et al., 2013). Assuming that these errors are Gaussian, the ln-likelihood is given by

$$\ln \mathcal{L}(\mathbf{x}, \boldsymbol{\theta}) = - \sum_{\beta} \frac{1}{2\sigma_{\beta}^2} (x_{\beta} - \Delta_{\beta}^2(\boldsymbol{\theta}))^2, \quad (4.5)$$

where β indexes each (k,z) bin and σ_{β} is the standard deviation of the power spectrum in the β^{th} bin. Taking the derivatives with respect to θ_i and θ_j , one obtains

$$F_{ij} = \sum_{\beta} \frac{1}{\sigma_{\beta}^2} \frac{\partial \Delta_{\beta}^2}{\partial \theta_i} \frac{\partial \Delta_{\beta}^2}{\partial \theta_j} + \frac{2}{\sigma_{\beta}^2} \frac{\partial \sigma_{\beta}}{\partial \theta_i} \frac{\partial \sigma_{\beta}}{\partial \theta_j}. \quad (4.6)$$

The first term in this formula is identical to the equation appearing in Pober et al. (2014) and arises from the mean term in the power spectrum estimate. The second term involves the derivative of the standard deviation of each power spectrum bin and accounts for the fact that cosmic variance is proportional to the power spectrum itself.¹

To elucidate the relative contribution from the second term in equation 4.6 to F_{ij} , we may assume that our measurement of Δ_{β}^2 is dominated by sample variance which

¹Formulas similar to equation 4.6 have been derived in the context of CMB analysis (Bunn et al., 1995; Vogeley and Szalay, 1996; Tegmark et al., 1997b). In the CMB case, the observable is usually taken to be the complex $a_{\ell m}$ which are distributed as a zero-mean Gaussian. In this case, the information in the variance of these terms leads to a formula very similar to the first sum in equation 4.6 (e.g. equation 18 in Tegmark et al. (1997b)) while the means contribute nothing.

is a decent approximation for the modes next generations arrays such as HERA-331 and the SKA will be most sensitive to (Mesinger et al., 2014). One can choose bin-sizes small enough such that Δ^2 is effectively constant over each bin, hence $\sigma_\beta = \Delta_\beta^2 / \sqrt{N_\beta}$ where N_β is the number independent measurements in the β^{th} bin. With this approximation, equation 4.6 becomes

$$F_{ij} = \sum_{\beta} \frac{N_{\beta}}{\Delta_{\beta}^4} \left(1 + \frac{2}{N_{\beta}} \right) \frac{\partial \Delta_{\beta}^2}{\partial \theta_i} \frac{\partial \Delta_{\beta}^2}{\partial \theta_j}. \quad (4.7)$$

Hence the contribution to the Fisher information from sample variance is only significant for small N_β . For the arrays studied in this paper, we find that the variance term in equation 4.6 tends to only contribute at 1% the level of the mean term so we ignore it and obtain an identical expression to the one appearing in Pober et al. (2014). Dropping the variance information allows us to write equation 4.6 as the dot product between two vectors used in Pober et al. (2014)

$$F_{ij} \equiv \mathbf{w}_i \cdot \mathbf{w}_j, \quad (4.8)$$

where

$$w_i(k, z) = \frac{1}{\sigma(k, z)} \frac{\partial \Delta^2(k, z)}{\partial \theta_i}. \quad (4.9)$$

Inspecting w_i allows us to determine the contribution of each (k, z) bin to the Fisher information on each parameter.

4.3 Simulations of the 21 cm Signal

The formalism that we set out in § 4.2 requires two ingredients: simulations of the brightness temperature field as a function of redshift and calculations of the error bars on each power spectrum measurement. In this section we address the first ingredient. We first briefly describe the publicly available 21cmFAST² code which we use to simulate the signal (§ 4.3.1). We then describe the specific reionization and heating

²http://homepage.sns.it/mesinger/DexM__21cmFAST.html

parameters that we choose to vary and our choices for their fiducial values (§ 4.3.2).

4.3.1 Semi-Numerical Simulations

We generate realizations of δT_b using the 21cmFAST code for which a detailed description is available in Mesinger and Furlanetto (2007) and Mesinger et al. (2011). The simulation volume is 750 Mpc on a side with 400^3 cells. Here we give a brief overview of its treatment of the heating and reionization physics.

A density field is calculated at each redshift by evolving an initial Gaussian random field via the Zeldovich approximation. An ionization field is then computed from the density fluctuations using the excursion set formalism of Furlanetto and Loeb (2004). Cells of the density field, filtered at a comoving scale R , are determined to be ionized if the fraction of mass that has collapsed into virialized structures, f_{coll} , exceeds the inverse of an ionizing efficiency parameter, defined as ζ . R is varied between the pixel size up to the mean free path (MFP) of ultraviolet (UV) photons in the HII regions which we denote as R_{mfp} .

The HI spin temperature couples to both the kinetic temperature of the gas (through collisions) and the Lyman- α flux from the first generations of stars via the Wouthuysen-Field effect (Field, 1959). In most scenarios, Lyman- α coupling saturates early on while the high optical depth to Lyman- α absorption couples the color temperature of UV photons to the kinetic temperature of the HI gas. Hence, T_s primarily reflects the thermal state of the HI gas and astrophysical processes affecting it. The impact on the kinetic temperature from X-rays is determined by integrating the X-ray specific emissivity along a light cone for each cell. The specific emissivity is assumed to be dominated by HMXBs or hot ISM and hence proportional to the star formation rate (Mineo et al., 2012a,b). The full expression for the emissivity in each simulation cell used in 21cmFAST is (Mesinger et al., 2013)

$$\epsilon_{h_p\nu}(\nu, \mathbf{x}, z') = \alpha_X h_p \frac{f_X}{10} \left(\frac{\nu}{\nu_{\min}} \right)^{-\alpha_X} \rho_{\text{SFR}}(\mathbf{x}, z') \quad (4.10)$$

For $\nu \geq \nu_{\min}$ and 0 otherwise. Here, h_p is the planck constant, ν_{\min} is a low-energy

obscuration threshold, α_X is the spectral index of X-ray emission and f_X is known as the ‘‘X-ray efficiency’’ which serves as an overall normalization parameter with $f_X \approx 1$ corresponding to 0.1 X-ray photon per baryon involved in star formation. Note that the factor of ν_{\min} within the exponent removes its effect on the overall X-ray luminosity. The star formation rate density, $\rho_{SFR}(\mathbf{x}, z')$, is approximated within each voxel with the equation (Mesinger et al., 2013)

$$\rho_{SFR}(\mathbf{x}, z') = \left[\langle \rho_b \rangle f_* (1 + \delta_{nl}(\mathbf{x})) \frac{df_{coll}}{dt}(\mathbf{x}) \right] \quad (4.11)$$

where $\langle \rho_b \rangle$ is the mean baryon density, f_* is the fraction of baryons in stars (assumed to be 0.1), δ_{nl} is the non-linear overdensity averaged over all smoothing scales, and df_{coll}/dt the derivative of the fraction of mass collapsed into virialized objects with virial temperatures greater than a certain threshold (which we explain below), It is computed using the hybrid prescription of Barkana and Loeb (2004).

4.3.2 Astrophysical Parameters and their Fiducial Values

In this work, we explore our ability to constrain a six-parameter model of reionization and heating which accounts for the major astrophysical degrees of freedom, $(\zeta, R_{\text{mfp}}, T_{\text{vir}}^{\min}, f_X, \nu_{\min}, \alpha_X)$ which are defined below. We choose a dimensionless parameterization by letting θ_i be the fractional difference of the parameter from its fiducial values. For example, $\theta_\zeta \equiv (\zeta - \zeta^{\text{fid}})/\zeta^{\text{fid}}$. We label the set of θ s for $(\zeta, R_{\text{mfp}}, T_{\text{vir}}^{\min}, f_X, \nu_{\min}, \alpha_X)$ as $(\theta_\zeta, \theta_R, \theta_T, \theta_f, \theta_\nu, \theta_\alpha)$.

The sensitivity of 21 cm experiments to the first three of these parameters was previously considered in Pober et al. (2014) and Greig and Mesinger (2015). These works only considering reionization redshifts and ignored the contribution to δT_b from spin temperature fluctuations. Mesinger et al. (2013) show that f_X and other heating parameters have an impact during the early stages of reionization, potentially introducing previously ignored degeneracies. 21 cm measurements during the heating epoch may therefor enhance our understanding of reionization. We now describe our choices for the fiducial value at which we compute the derivatives in equation 4.6 for

each parameter.

- ζ : The ionization efficiency describes the number of ultraviolet photons per unit time that enter the IGM from galaxies. It primarily affects the timing of reionization and can be written as a degenerate combination of parameters, $\zeta = f_{esc} f_* N_{\gamma/b} (1 + n_{rec})^{-1}$ (Furlanetto and Loeb, 2004), where f_{esc} is the UV escape fraction, $N_{\gamma/b}$ is the number of ionizing photons per baryon in stars and n_{rec} is the number of times a typical HI atom undergoes recombination. Combinations of the optical depth to the CMB, quasar absorption data, the kinetic Sunyaev-Zeldovich effect allow some constraints on ζ (Mesinger et al., 2012) and is typically thought to lie between 5 and 50. We choose a fiducial value of $\zeta = 20$ which, combined with our other fiducial choices, predicts a mean optical depth to the CMB of $\tau_e \approx 0.08$, and 50% reionization at $z_{re} = 8.5$ which is in line with current constraints from Planck Collaboration et al. (2015).
- R_{mfp} : The MFP of UV photons in HII regions, this parameter determines the maximum HII bubble size and primarily determines the location of the “knee” in the power spectrum. Physically, R_{mfp} is set by the number density and optical depth of Lyman-limit systems. R_{mfp} is highly unconstrained with limits from observations of Lyman- α systems at $z \sim 6$ allowing for values between 3 and 80 Mpc (Songaila and Cowie, 2010). Subgrid modeling of inhomogeneous recombination point towards a smaller R_{mfp} between 5 and 20 Mpc (Sobacchi and Mesinger, 2014) leading us to choose a fiducial value of $R_{mfp} = 15$ Mpc.
- T_{vir}^{min} : The emission lines of H_2 are expected to serve as the predominant cooling pathway for primordial gas, allowing for the collapse and fragmentation necessary for star formation. While radiative cooling is possible in halos with $T_{vir} \gtrsim 100K$ (Haiman et al., 1996b; Tegmark et al., 1997a), photodissociation induced by UV background set up by the first sources suppresses cooling in less massive halos without sufficient self shielding (Haiman et al., 1996c, 1997). Even if sufficient cooling occurs, feedback can also suppress star formation (Springel and Hernquist, 2003; Mesinger and Dijkstra, 2008; Okamoto et al., 2008). Stud-

ies indicate a broad range of plausible minimum virial temperatures between $T_{\text{vir}} = 10^2 - 10^5$ K corresponding to halo masses between $10^6 - 10^8 M_{\odot}$ at $z \approx 10$. We adopt a fiducial $T_{\text{vir}}^{\text{min}}$ of 1.5×10^4 K which corresponds to the threshold for atomic line cooling. $T_{\text{vir}}^{\text{min}}$ directly enters our simulations by determining the minimum mass above which f_{coll} and df_{coll}/dt is calculated. While thermal or mechanical feedback has the potential to change $T_{\text{vir}}^{\text{min}}$ as a function of redshift (Mesinger et al., 2013), we assume that all the halos involved in heating the IGM also take part in reionizing it and hold $T_{\text{vir}}^{\text{min}}$ constant. We will sometimes group $T_{\text{vir}}^{\text{min}}$ with other “reionization” parameter due to its inclusion in the three-parameter model of previous works that only address reionization (Mesinger et al., 2012; Pober et al., 2014; Greig and Mesinger, 2015; Greig et al., 2016). However, we emphasize that it determines the minimal masses of the halos driving X-ray heating as well and is just as much an “X-ray heating parameter” as any of the parameters determining the X-ray SEDs of early galaxies which we now list.

- f_X : Our fiducial value of $f_X = 1$ is chosen to give an integrated 0.5-8 keV luminosity of $\approx 5 \times 10^{39}$ erg s $^{-1}$ M_{\odot} yr $^{-1}$, consistent with the $\approx 3 \times 10^{39}$ erg s $^{-1}$ M_{\odot} yr $^{-1}$ observed at $z=0$ by Mineo et al. (2012a) and corresponds to 0.1 X-ray photons per stellar baryon. While $f_X = 1$ matches local observations, there are reasons to expect different values at high redshift. For example, the decrease in metallicity with redshift might increase the rate of very X-ray luminous black hole X-ray binaries, boosting f_X (Mirabel et al., 2011). Observations out to $z \sim 4$ on *Chandra Deep Field-South* (Xue et al., 2011) have been interpreted with conflicting results. Cowie et al. (2012) do not find evolution in the X-ray luminosity to star formation rate out to redshift ~ 4 while Basu-Zych et al. (2013) claim to observe weak evolution consistent with population synthesis models (Fragos et al., 2013) after adjusting for dust extinction.
- ν_{min} : Absorption by the ISM in early galaxies will cause the emergent spectral energy distribution to differ from the ones intrinsic to the sources. In particular,

large ν_{\min} can lead to a particularly hard X-ray spectrum that delays heating and reduces the contrast between hot and cold patches during heating (Fialkov et al., 2014). The degree to which absorption is expected to be present depends critically on both the column density and composition of the ISM of the host galaxies with metals absorbing X-rays with energies above 0.5 keV and helium primarily absorbing softer X-rays. Our fiducial choice of the X-ray obscuration threshold, $\nu_{\min} = 0.3$ keV is identical to that used in Pacucci et al. (2014) and describes an ISM with a similar column density to the Milky Way but with low metallicity.

- α_X : The spectral index of X-ray emission from early galaxies. X-ray emission from local galaxies is observed to originate from two different sources: X-ray binaries (XRB) and the diffuse hot ISM. XRB spectra generally follow a power law of $\alpha \sim 1.0$ between 0.5 and 2 keV (Mineo et al., 2012a). Emission from the hot ISM originates from metal line cooling and thermal bremsstrahlung in the plasma generated by supernovae and stellar winds. The diffuse emission observed by Mineo et al. (2012b) has been found by Pacucci et al. (2014) to be well approximated by a power law of $\alpha \sim 3.0$ between $\sim 0.5 - 10$ keV. These authors also observed that the maximum amplitude of the power spectrum during X-ray heating is highly sensitive to α_X . A steeper spectrum resulting from ISM dominated emission is abundant in soft X-rays and leads to significantly higher contrast between hot and cold regions due to their short MFP. This leads to a boost in power spectrum amplitude by a factor of ~ 3 . We choose a fiducial α_X of 1.2.

For each astrophysical parameter, we run six simulations varying θ_i by $\pm 1\%$, $\pm 5\%$, and $\pm 10\%$ of its fiducial value and a linear fit of $\Delta^2(k, z)$ is used to compute $\partial\Delta^2/\partial\theta_i$. Inspection shows that the power spectrum is well described by a linear trend over this range of parameter values. In Fig. 4-1, we show the evolution of the density-weighted-average of the neutral fraction as a function of redshift for our fiducial model which predicts 50% reionization at $z_{re} \approx 8.5$. We also display the volume averaged




Figure 4-1: Top: The evolution of the density weighted average of the neutral fraction. Middle: The evolution of the kinetic temperature (T_k), T_s , T_{cmb} . Bottom: the redshift evolution of the mean 21 cm brightness temperature, $\langle \delta T_b \rangle$. All averages are taken over volume.

evolution of the kinetic, spin, and 21 cm brightness temperatures compared to the evolution of the CMB temperature. We see that $\langle T_s \rangle$ exceeds T_{cmb} at $z \approx 12.5$ but is within an order of magnitude of T_{cmb} out to a redshift of $z \approx 9$, hence we can expect some signatures of spin temperature fluctuations to be present in our signal early on in reionization. In Fig. 4-2 we show our power spectra as a function of redshift at two different comoving scales at $\pm 10\%$ of their fiducial parameter values. The evolution of the power spectrum on large scales follows the three peaked structure noted in Pritchard and Furlanetto (2007); Santos et al. (2008); Baek et al. (2010); Mesinger et al. (2013) corresponding, in order of redshift, to reionization, X-ray heating, and Lyman- α coupling. The redshifts of these peaks depends on the comoving scale. At $k = 0.1$, the peaks occur at $z = 8.5, 15,$ and 22.5 .

4.4 Instrument and Foreground Models

The next ingredient for our Fisher matrix analysis is a set of error bars on each of the observed cosmological modes. The error on each spherically averaged power spectrum estimate not only depends on both instrumental parameters such as the collecting area, uv -coverage, and observing time, but also strongly on our ability to mitigate foreground contamination.

We investigate two large experiments: the Hydrogen Epoch of Reionization Array (HERA), which is being commissioned in South Africa, and the Square Kilometer Array (SKA). For each instrument, we derive the power spectrum sensitivity using the public `21cmSense` code³ (Pober et al., 2013b, 2014) assuming 6 hours of observations per night over 180 days and a spectral resolution of 100 kHz. For simplicity, we also assume that all arrays perform drift-scan observations. While the SKA is not a drift-scan instrument, the sensitivity difference between a drift scan and tracking a single eight degree field for six hours per night was only found to be on the order of $\sim 10 - 20\%$ (Pober, 2015) depending on the degree of foreground contamination. We assume that observations are taken simultaneously at all redshifts between 5 and 30

³<https://github.com/jpober/21cmSense>




Figure 4-2: Δ_{21}^2 as a function of redshift for two different 1d k bins, $k = 0.1 h\text{Mpc}^{-1}$ (black lines) and $k = 0.4 h\text{Mpc}^{-1}$ (red lines). For each parameter, we show the power spectrum at $\theta = \pm 0.1$ (thick solid and thick dashed lines respectively) along with the difference (thin solid lines). Note that our parameterization defines θ as the fractional difference of each parameter from its fiducial. The first two peaks of the three-peaked structure, discussed in Pritchard and Furlanetto (2007); Santos et al. (2008); Baek et al. (2010) and Mesinger et al. (2013), is clearly visible representing the epochs of reionization and X-ray heating. With the exception of R_{mfp} , parameter changes affect a broad range of redshifts.

and that power spectra are estimated from sub-bands with a co-evolution bandwidth of $\Delta z = 0.5$, to avoid strong evolution effects (Mao et al., 2008). The redshift interval assumes observing over the frequency band between 57 and 237 MHz and we note that simultaneous observations over such a large band may not be possible with a single feed, raising the observing time by a factor of two. For each array, we consider a discrete uv plane upon which measurements are gridded with a cell size set by the instrument's antennae footprint and using rotation synthesis, we compute the number of seconds of observing performed in each cell $\tau(\mathbf{k})$. We emphasize that $\tau(\mathbf{k})$ is different for each observed frequency due to the fact that each instrumental baseline has a fixed physical length and antenna size while the uv cell a baseline occupies is set by the number of wavelengths between its two antennas. The uv cell size also depends on frequency since the number of wavelengths spanned by the physical antenna changes. The variance of the power spectrum estimate within each uv cell is given by

$$\sigma^2(\mathbf{k}) = \left[X^2 Y \frac{\Omega' T_{sys}^2}{2\tau(\mathbf{k})} + P_{21}(\mathbf{k}) \right]^2, \quad (4.12)$$

where Ω' is the ratio between the solid angle integral of the primary beam squared (Parsons et al., 2014) and the solid angle integral of the beam while T_{sys} is the sum of the sky and receiver temperatures whose values we choose to be $T_{rec} = 100$ K and $T_{sky} = 60\lambda^{2.55}$ K (Fixsen et al., 2011). X is the comoving angular diameter distance and Y is a linear conversion factor between frequency and radial distance given by (Morales and Hewitt, 2004)

$$Y = \frac{c(1+z)^2}{H_0 f_{21} E(z)}, \quad (4.13)$$

where $f_{21} \approx 1420$ MHz is the frequency of 21 cm radiation, c is the speed of light, H_0 is the hubble parameter at $z = 0$, and $E(z) = H(z)/H_0$. At each observed frequency and array, $\tau(\mathbf{k})$ is computed by dividing the uv plane into cells with diameter D/λ and adding up the cumulative time within each cell occupied by each array's antennas after rotation synthesis. Ω' is computed assuming a gaussian beam with $\sigma = 0.45\lambda/D$ where D is the diameter of the antenna element, an approximation that ensures the volume of the central lobe of the airy disk for the aperture matches that of a Gaussian.

The sensitivity within each k -bin is computed by taking the inverse variance weighted average of all uv cells within the bin that are not contaminated by foregrounds.

While foregrounds are expected to dominate the signal by a factor of $\gtrsim 10^5$, they are also spectrally smooth and only occupy a limited region of k -space known as the “wedge” (Datta et al., 2010; Parsons et al., 2012a; Morales et al., 2012; Vedantham et al., 2012; Thyagarajan et al., 2013; Trott et al., 2012; Hazelton et al., 2013; Liu et al., 2014a,b; Thyagarajan et al., 2015a,c). The degree to which we might be able to observe inside (and close to) this foreground contaminated region will depend crucially on our ability to characterize the foregrounds and our instrument. Since the extent of the wedge corresponds to the angular offset of sources from the phase center, our ability to characterize and subtract sources from the primary beam sidelobes will determine what uncontaminated modes will be available for power spectrum estimation as seen in Thyagarajan et al. (2015a,c) and Pober et al. (2016). We consider the three different foreground scenarios from Pober et al. (2014) to describe the efficacy of foreground isolation.

- **Optimistic (Foreground Subtraction):** All modes outside the full width half power of the primary beam are sufficiently decontaminated as to be used in power spectrum estimation. We also include a small buffer of $k_{\parallel} = 0.05 h\text{Mpc}^{-1}$ to account for intrinsic spectral structure in the foregrounds and/or the instrument. This buffer is significantly smaller than the supra horizon emission observed in Pober et al. (2013a) out to $\approx 0.1 h\text{Mpc}^{-1}$. Ionospheric diffraction, whose severity runs inversely with frequency will likely increase the difficulty of foreground subtraction at X-ray heating redshifts.
- **Moderate (Foreground Avoidance):** In this scenario, we assume that all modes within the wedge are unusable and supra-horizon emission extends to $0.1 h\text{Mpc}^{-1}$ beyond the wedge, in line with observations (Pober et al., 2013a; Parsons et al., 2014; Ali et al., 2015; Dillon et al., 2015a).
- **Pessimistic (Instantly Redundant Delay Transform Power Spectrum):** This scenario is almost identical to our moderate foregrounds scenario except

that only baselines that are instantaneously redundant in local sidereal time are added coherently. This is the sensitivity achievable using the current delay power spectrum technique (Parsons et al., 2012a) which thus far has led to the most stringent upper limits in the field (Ali et al., 2015). We note that there is no fundamental reason for the delay transform technique to not coherently add partially redundant visibilities which is an ongoing topic of research.

Both Ω' and $\tau(\mathbf{k})$ in equation 4.12 depend on our instrument. In this work, we consider the following experimental configurations:

- **HERA-127/331:** The Hydrogen Epoch of Reionization Array (HERA) is an experiment undergoing commissioning now, to detect the 21 cm brightness temperature fluctuations during and before the EoR. Focusing on the foreground avoidance approach that has thus far proved successful for PAPER, HERA is designed to maximize collecting area outside of the wedge by filling the uv plane with short baselines. The antenna layout for HERA involves 331 hexagonally-packed, 14 m diameter dishes. A staged buildout is expected to occur with the penultimate and ultimate stages comprising of a 127 and 331 dish core. HERA will also contain an additional 21 outrigger antennas to assist in imaging and foreground characterization. However, we do not include these outriggers in our analysis since they do not contribute significantly to HERA's sensitivity which is derived primarily from its short, core baselines.
- **SKA-1 LOW:** We base our model of the SKA-LOW instrument on the description in Dewdney et al. (2013). We also reduce the antenna count by 50% to reflect the recent rebaselining, making it nearly identical to the proposed design #1 in Greig et al. (2015). The array is comprised of 446, 35 m diameter phased arrays of log-periodic dipole antennas. These stations are distributed in radius as a Gaussian with 75% of antennae falling within a 1 km radius.

4.5 Power Spectrum Derivatives and Their Physical Origin

Having described our simulations of the signal and instrumental noise, we are in a position to discuss the two stages of our results which include the derivatives of the power spectrum with respect to each parameter and the resulting covariances. In this section we provide physical intuition for the outputs of our derivative calculations and the nature of the information on each quantity that is available at different redshifts.

We show our fiducial power spectrum along with the 1σ uncertainty regions for the arrays studied in this paper at the top of Fig. 4-3. As observed in numerous previous works (e.g. Pritchard and Furlanetto (2007); Christian and Loeb (2013); Mesinger et al. (2014)), the 21 cm signal is detectable out to $z \approx 21$ due to the larger contrast available between cold and hot regions of the IGM during heating (Mesinger et al., 2013). We show $\partial\Delta_{21}^2/\partial\theta_i$ for our astrophysical parameters at the bottom of Fig. 4-3 and see that with the exception of R_{mfp} , the derivative of the power spectrum with respect to the ionization parameters is substantial out to redshifts beyond the typical range associated with reionization. While we compute our Fisher matrix from redshift bins of $\Delta z = 0.5$, for legibility, the panels in Fig. 4-3 are shown for intervals of $\Delta z = 3.0$. In addition, In the next two sections, we discuss, in detail, the origins of the trends in the derivatives related to both reionization (§ 4.5.1) and X-ray heating (§ 4.5.2).

4.5.1 How Reionization Parameters Affect the 21 cm Power Spectrum.

We now describe the trends in the derivatives associated with reionization, several of which have already been discussed in the literature (Mesinger et al., 2012; Pober et al., 2014; Sobacchi and Mesinger, 2014) as well as new signatures present at high redshift that are only detectable with the inclusion of the large negative $(1 - T_{\text{cmb}}/T_s)$ term supplied by the spin temperature calculation.




Figure 4-3: Top: The power spectrum of 21 cm fluctuations (solid red lines) over numerous redshifts. Filled regions denote the 1σ errors for the instruments considered in this paper with moderate foregrounds. Bottom: The derivatives of the 21 cm power spectrum with respect to the astrophysical parameters considered in this work as a function of k at various redshifts. Derivatives are substantial over all redshifts except for R_{mfp} which only affects the end of reionization. Notably, $\partial_{\zeta}\Delta_{21}^2$ is negative on small scales at high redshift, a signature of the beginnings of inside out reionization while X-ray spectral parameters follow very similar redshift trends, indicative of degeneracy.

The primary effect of increasing ζ is to accelerate the time of reionization, shifting the peak to higher redshift. Thus we see a negative derivative to the left and the positive derivative to the right of the reionization peak. It is interesting to note that the derivative with respect to ζ remains significant (and primarily negative) far beyond the rise of the reionization peak and well into the heating epoch. This is due to the fact that the small precursor ionization bubbles exist out to high redshift, occupying the same voxels with the largest over-densities and greatest T_s . These HII bubbles set δT_b to 0 at the hottest points in the IGM and reduce the contrast between hot spots and the cold background (Fig. 4-4). Because these bubbles occur on small spatial scales, this leads to a reduction in the power spectrum amplitude with increasing ζ at large k . At large scales, we see a positive derivative at the rise of the heating peak and a negative derivative at its fall as we might expect if the heating rate were increased (Fig. 4-2). An increased ionization fraction decreases the optical depth of X-rays from the photoionization of HI, HeI, and HeII, providing such a rate increase. The falling edge of the Lyman- α peak is similarly affected, perhaps due to an increase in the number of Lyman- α photons arising from X-ray excitations.

Increasing the MFP of ionizing photons is known to shift the “knee” of the power spectrum to larger comoving scales. The diameter of the regions corresponding to our fiducial R_{mfp} of 15 Mpc^{-1} correspond to $k \approx 0.3 \text{ hMpc}^{-1}$, hence the positive derivative below $k \approx 0.3 \text{ hMpc}^{-1}$ and the negative derivative at smaller scales. Since the MFP in HII regions only affects the brightness temperature fluctuations once the HII bubbles themselves have had time to grow out to this scale, we see no significant effect of R_{mfp} on the early reionization power spectrum beyond $z \approx 9$. Mesinger et al. (2012) note that a smaller R_{mfp} has the effect of delaying the end of reionization, explaining the negative derivative across all scales at the lowest redshifts.

$T_{\text{vir}}^{\text{min}}$ affects both reionization and heating, however we will discuss its affect in this section. Since increasing $T_{\text{vir}}^{\text{min}}$ delays heating and reionization, its clearest signature is to shift the peaks towards low redshift, leading to positive differences to the left of each peak and negative differences to the right (Fig. 4-2). Smaller comoving scales ($k \gtrsim 0.4 \text{ hMpc}^{-1}$) transition through the peaks at earlier times than larger scales (see




Figure 4-4: Left: The logarithm of the probability distribution function (PDF) of pixels at $z = 12.4$ for our fiducial model as a function of T_s vs. x_{HI} . Even before the majority of reionization, early HII bubbles ionize the hottest points in the IGM, leading to a pileup of high T_s pixels at $x_{HI} = 0$ and reducing the contrast in δT_b between hot and cold regions. Right: PDFs of δT_b with and without x_{HI} manually set to unity everywhere. The presence of ionization during X-ray heating leads to a decrease in the large T_s wing, near 30 mK, and a spike at 0 mK, leading to a reduction in the dynamic range of the field and an overall decrease in power.

Fig. 4-2). At $z = 9$ and 18, small comoving scales are at the fall of the reionization and heating peaks respectively (hence the positive derivatives) and at $z = 21$ at the rise of the heating peak (rather than the fall of Lyman- α) leading to the negative derivative at large k in Fig. 4-3.

4.5.2 How X-ray Spectral Properties Affect the Power Spectrum.

We now describe and provide physical intuition for the derivatives with respect to the X-ray spectral properties of galaxies before reionization.

As increasing f_X raises the heating rate, the most obvious consequence is to shift the X-ray heating peak to higher redshift (Mesinger et al., 2013; Christian and Loeb, 2013; Mesinger et al., 2014; Pacucci et al., 2014). In Fig. 4-3 this trend is clearly observable in the positive derivative at the rising edge of the heating peak and the negative derivative at the falling edge. There is also a significant positive derivative early in reionization and a slightly negative one at its conclusion. The reionization

peak is impacted by a number of competing effects related to f_X . Since X-rays also generate secondary ionizations and have a longer MFP than UV photons, highly emissive scenarios produce a partially ionized haze that reduces the contrast between ionized and neutral patches and the amplitude of the reionization power spectrum (Mesinger et al., 2013). Secondary ionizations also have the effect of shifting the reionization peak to higher redshift. As the power spectrum maximum occurs $z \approx z_{re}$, this causes an increase in power at the start of reionization and a decrease at the tail end. Finally, increasing f_X raises the spin temperature during reionization. Since we are in the regime where $T_s > T_{\text{cmb}}$ over the reionization peak, increasing f_X leads to an increase in the $(1 - T_{\text{cmb}}/T_s)$ factor in δT_b , leading to an overall increase in the reionization power spectrum before the spin temperature's impact has saturated. In Figs 4-2 and 4-3 we see the difference is positive during the onset of reionization and negative during the fall; indicating that the rise in spin temperature from increased f_X and the shift to higher redshifts dominates the onset of reionization but that the direct spin temperature effects are saturated by the end.

We next examine the derivatives with respect to the obscuration threshold, ν_{min} . Raising ν_{min} hardens the X-ray spectrum of the first galaxies. Since hard X-rays have significantly longer MFPs, this delays their absorption, leading to relatively uniform late heating (Fialkov et al., 2014; Pacucci et al., 2014). This has the effect of shifting the minimum between the reionization and heating peaks when $\langle \delta T_b \rangle \approx 0$ to lower redshift while suppressing the amplitude of heating fluctuations. In addition to the longer MFP, the heating delay is also due to harder X-rays depositing a larger fraction energy into ionizations rather than heating (Furlanetto and Stoever, 2010). The increased ionization energy fraction also leads to a slight reduction in the mean neutral fraction across redshift. One might also expect a slight shift in the reionization peak to high redshift as well and a decrease in amplitude from the reduced spin temperature and lower contrast in the ionization field from increased X-ray ionizations. While a shift and amplitude reduction in the heating peak is clear, we do not see a decrease in the reionization peak. In fact, inspection of Figs 4-2 and 4-3 indicates that the amplitude of the peak actually increases. An explanation for this behavior is that the

ionization and $(1 - T_{\text{cmb}}/T_s)$ fields anti-correlate, leading to a negative contribution to the overall power spectrum amplitude. A larger ν_{min} leads to a decrease in the contrast between hot and cold patches which reduces the overall amplitude of this negative cross correlation and causes an increase in the reionization power spectrum. While the reduction in contrast raises the power spectrum amplitude in the neighborhood of our fiducial model, the trends towards a smaller reionization power spectrum dominate at much larger ν_{min} (Mesinger et al., 2013; Fialkov et al., 2014).

As noted in Pacucci et al. (2014), increasing α_X reduces the mean free path of the X-rays, amplifying the contrast between hot and cool patches and leading to an increase in the heating power spectrum amplitude over the entire peak. The increased spin temperature also drives up the amplitude of the ionization power spectrum and causes a slight shift towards higher redshifts as well. α_X has the weakest signature during the reionization epoch, perhaps in part due to the amplification of the anticorrelation between spin temperature and ionization fraction canceling out the increase in T_s . We note that in models with much lower heating efficiency that an increased α_X leads to a noticeable dip during the rise of the reionization power spectrum (e.g. Fig. 8 in Christian and Loeb (2013) or Fig. 5 in Pacucci et al. (2014)) due to the enhancement of the T_s - x_{HI} anti-correlation.

4.6 Constraints from Heating Epoch Observations

We now turn to the results of our Fisher matrix calculation. We focus our discussion on degeneracies and the dependence of overall constraints on the range of observed redshifts.

4.6.1 Degeneracies Between Parameters

Combining our power spectrum derivatives with our sensitivity estimates, we calculate $w_i(k, z)$ which indicates the relative contribution of each 1d k bin and redshift to the Fisher information on each parameter. In addition, we can easily spot the sources of covariance between the different parameters by looking for similar k and z evolution.




Figure 4-5: Plotting $w_i(k, z)$ versus redshift for several different co-moving scales gives us a sense of the covariances between various parameters. Here we assume 1080 hours of observation on HERA-331 and the moderate foreground model. Since the thermal noise on Δ^2 increases rapidly with k , w_i is maximized at larger spatial scales. As we might expect from our fiducial model, w_i for reionization parameters is maximized at lower redshifts while w_i for X-ray spectral parameters is significant over the heating epoch. $T_{\text{vir}}^{\text{min}}$ affects both heating and reionization and has a broad redshift distribution. Vertical lines indicate the location of each of the three power spectrum peaks at $k = 0.1 h \text{Mpc}^{-1}$

If at the same redshifts and k values, two parameters have opposite (equal) signs in w_i , then a positive change in the first parameter can be compensated for by a positive (negative) change in the other, leading to degeneracy. This degeneracy can be broken if additional redshifts and Fourier modes are added in which the parameters do not have similar evolution. We denote the w_i s for $(\zeta, R_{\text{mfp}}, T_{\text{vir}}^{\text{min}}, f_X, \nu_{\text{min}}, \alpha_X)$ as $(w_\zeta, w_R, w_T, w_f, w_\nu, w_\alpha)$.

To understand specific sources of covariance, we plot w_i vs. z at several different comoving scales (Fig. 4-5) and directly compare the redshift evolution between all w_i s at a single mode in Fig. 4-6. Here we assume the thermal noise from HERA-331 and our moderate foregrounds scenario. Redshifts where $|w_i(k, z)|$ is largest indicate where our measurements will have maximal sensitivity to each θ_i .

As we might expect, the greatest information on ζ is obtained over reionization.




Figure 4-6: Plotting $w_i(k, z)$ at a single cosmological Fourier mode for all of our parameters on the same panel facilitates direct comparison. Many of the parameters have similar redshift evolutions that differ by a sign, making their effects on the power spectrum degenerate.

Because w_ζ follows essentially the opposite trend of w_T , even out to higher redshifts, there is extensive degeneracy between the two parameters. Since w_R is only significant during the end of reionization, R_{mfp} 's degeneracies with ζ and $T_{\text{vir}}^{\text{min}}$ are broken with the inclusion of higher redshift observations. Turning to the X-ray spectral parameters; the evolution of w_f and w_ζ follows very similar trends during the fall of the heating peak and reionization. While w_ν and w_f have unique trends over the entire duration of reionization, their evolution is very similar during the end of the heating peak where the power spectrum is more sensitive to them. Because its impact on reionization is negligible, α_X 's covariance with the reionization parameters will be very small. Over heating, w_α follows a similar and opposite trend to w_ν .

4.6.2 How well can Epoch of Reionization Measurements Constrain X-ray Spectral Properties?

We now determine what constraints on X-ray spectral properties can actually be obtained by measurements of the reionization peak which, in our model, extends to roughly $z \lesssim 10$. In Fig. 4-7 we show the marginalized 1σ error bars as a function of the maximal redshift included in power spectrum observations on HERA-331 with moderate foregrounds. We see that below the onset of the reionization peak at $z \approx 10$, the error bars are at $\approx 40\%$ for f_x and ν_{min} while the error on α_x exceeds 100%. The latter is understandable given that w_α is very small below $z = 10$ relative to the other




Figure 4-7: The fractional errors on astrophysical parameters as a function of maximal observed redshift for HERA-331 with moderate foreground contamination. From low redshift measurements, ν_{\min} and f_X are constrained to within $\approx 40\%$ though the spectral index α_X remains highly uncertain. Measurements at $z \gtrsim 10$ allow for $\lesssim 10\%$ limits on X-ray spectral parameters, including α_X and a factor of two improvement in constraints on T_{vir}^{\min} and reionization. Inability to observe within the FM radio-band (pink shaded region) raises the errors on heating parameters by a factor of two. The fall of the error bars with redshift bottoms out at high z due to increasing thermal noise.




Figure 4-8: 95% confidence regions for the X-ray spectral properties of early galaxies, marginalizing over $(\zeta, R_{\text{mfp}}, T_{\text{vir}}^{\text{min}})$ from measurements on HERA-331 with our moderate foregrounds scenario. At low redshift ($z \leq 10$), hardly any signature of α_X is present, leading to large error bars on the α_X axis. Because ν_{min} and f_X incur very similar changes on the power spectrum during the beginning of reionization (Fig. 4-5), they are highly degenerate. Observations of the heating peak break these degeneracies.

parameters.

We delve into the source of heating uncertainty in ν_{min} and f_X over reionization by plotting the 95% confidence ellipses of our heating parameters, in Fig. 4-8, from observations over reionization ($5 \lesssim z \lesssim 10$), heating ($10 \lesssim z \lesssim 25$), and both ($5 \lesssim z \lesssim 25$). We see that during reionization, there is a large negative correlation between f_X and ν_{min} (w_f in Fig. 4-5 follows very similar trends to w_ν at the end of reionization). When we fix ν_{min} at its fiducial value, we obtain several percent constraints on f_X (instead of 40%) and vice versa. α_X , covaries weakly with the other heating parameters over reionization, but has very large error bars. Eliminating it reduces the errors on ν_{min} and f_X by a factor of two. The inclusion of heating measurements in addition to reionization, removes much of the $f_X - \nu_{\text{min}}$ degeneracy, bringing their fractional errors to within $\approx 6\%$.

While signatures of heating are present in the early stages of reionization, degeneracies between heating parameters prevent precision constraints. All constraints on

α_X come from direct measurements of the heating epoch at $z \gtrsim 15$.

4.6.3 How well do Epoch of X-ray Heating measurements improve Constraints on Reionization?

We now determine what measurements during the Epoch of X-ray heating ($z \gtrsim 10$) can teach us about reionization. In Fig. 4-7 we see that the errors on $T_{\text{vir}}^{\text{min}}$ and ζ are reduced roughly by a factor of two when data from the heating epoch is included. Where do these improvements originate from? We know from our discussion in § 4.5 that information on $T_{\text{vir}}^{\text{min}}$ and ζ extends to higher redshift, providing one possible explanation. On the other hand, inspection of Fig. 4-6 indicates that there are also degeneracies between the reionization and heating parameters during the beginning of reionization and higher redshift measurements add information by breaking these degeneracies. We determine the impact of these two sources of information by comparing the 95% confidence ellipses for reionization parameters derived from reionization observations in which heating parameters have been fixed and the confidence ellipses when all parameters have been marginalized over with heating observations included (Fig. 4-9). While fixing heating parameters leads to a significant drop in the areas of the confidence regions, the inclusion of high redshift measurements provides additional improvements. We conclude that the improvements in reionization constraints arise through both mechanisms; breaking low redshift degeneracies with the heating parameters and obtaining additional information present in early HII bubbles. This also means that not marginalizing over heating parameters (as is done in Pober et al. (2014) and Greig and Mesinger (2015)) leads to overly-optimistic predictions of reionization constraints. We note that the heating-reionization degeneracies arise primarily at the beginning of reionization when $T_s \sim T_{\text{cmb}}$ (Fig. 4-1). In a more efficient heating scenario, we would expect their contribution to be reduced.




Figure 4-9: Confidence ellipses (95%) for $T_{\text{vir}}^{\text{min}}$ and reionization parameters. By comparing the ellipses resulting from fixing our heating history and only observing at low redshift and the ellipses resulting from marginalizing over all parameters but including heating epoch measurements, one can see that a significant fraction of the gains in reionization uncertainties at high redshift come from breaking degeneracies with heating parameters rather than the direct signatures of reionization. This also shows that not marginalizing over heating parameters leads to over-optimistic predictions of reionization uncertainties.

4.6.4 Overall Parameter Constraints

We observe overall degeneracies and error bars in the 95% confidence ellipses derived from the inverse of the Fisher matrix for data between $z = 5$ and $z = 25$ in Fig. 4-10. Here we assume that all frequencies, including those within the FM are accessible to observations. Relatively little degeneracy exists between heating and reionization parameters, though we would not expect this to be the case for a model in which heating were delayed.

In Fig. 4-7 we show the 1σ uncertainty on each astrophysical parameter as a function of the maximal redshift included in the Fisher matrix analysis. While R_{mfp} reaches its minimal error of several percent by $z \approx 9$, the error bar on ζ and $T_{\text{vir}}^{\text{min}}$ can be improved by nearly a factor of two by including power spectrum measurements of heating. Measurements at $z \gtrsim 10$ bring the error bars on all heating parameters below $\approx 6\%$ for HERA-331 while reionization measurements alone yield 40% errors. In Fig. 4-7 we also show the effect of the FM on heating parameter error

bars. Observations within the FM lead to a factor of two improvement in limits on heating, showing that the RFI environment of the observatory, within the FM, will have an important effect on the science that can be performed. It is found in Ewall-Wice et al. (2016b) that after three hours of integration, FM is not a limiting systematic for pre-reionization observations at the Murchison Radio Observatory in Western Australia (where SKA-1 is planned). However, the amount of lower level RFI that might become a limiting obstacle after the hundreds of hours of observation necessary for a detection is unknown. We note that the impact of the FM on our constraints is model dependent and that the heating peak for the scenario considered in this work occurs at $z \approx 15$, right in the middle of the FM (88-108 MHz). Models with a different fiducial f_X would produce a power spectrum peak at higher or lower redshifts, making the scenario that we consider here pessimistic with respect to the FM's impact on science.

We generalize our discussion to additional instrument and foreground scenarios discussed in § 4.4 by displaying forecasts of 1σ fractional parameter uncertainties in Table 4.1. In the pessimistic scenario, the baselines that are not instantaneously redundant are never added coherently, leading to a very significant reduction in the SKA's sensitivity and preventing it from placing any significant limits on X-ray heating. In the foreground avoidance scenario, SKA-1 and HERA-127 both place $\approx 10 - 15\%$ constraints on heating parameters while HERA-331 obtains $\approx 6\%$ error bars. Should we obtain sufficient characterization of foregrounds as to allow us to subtract them and work within the wedge, then several percent to sub-percent constraints are possible with HERA-331 and the SKA which is far beyond the the modeling uncertainties in our semi-numerical framework.

4.7 Conclusions

Measurements of the 21 cm power spectrum during the EoR are poised to put significant limits on the properties of the UV sources that ionized the IGM. At higher redshifts, the power spectrum is heavily influenced by X-rays from accretion onto the

first stellar mass black holes and ISM heated by the first supernovae. While reionization has a number of complementary probes, observations of the 21 cm global signal and power spectrum at these higher redshifts provide us with what is likely the only means of obtaining detailed knowledge on the earliest X-ray sources and their impact on future generations of galaxies.

In this paper we have used the Fisher matrix formalism and semi-numerical simulations to take a first step in quantifying the accuracy with which upcoming experiments will constrain the properties of the first X-ray sources. Our analysis also aims to understand what additional constraints on reionization parameters exist at higher redshift when the spin temperature calculation is included and whether higher redshift observations might break degeneracies between reionization parameters such as the degeneracy between ζ and $T_{\text{vir}}^{\text{min}}$.

We have found that the detectable impact of the ionization efficiency is manifested in the form of early HII holes around IGM hotspots and that the inclusion of the spin temperature calculation and the additional heating parameters increases our uncertainty of reionization parameters through new degeneracies. Observations of the heating epoch reduce the errors on reionization parameters by a factor of two by accessing the information in early HII bubbles and breaking degeneracies between ionization and heating during the beginning of reionization. Since previous works ignored the degeneracies of reionization with heating the predictions in these works are therefore optimistic by a factor of ~ 2 .

Though heating does have an effect on the reionization power spectrum as noted by Mesinger et al. (2013) and Fialkov et al. (2014) and clearly visible in the non-zero derivatives at $z = 9$ in Fig. 4-3 in our paper, the effects of different heating parameters are highly degenerate leading to $\gtrsim 40\%$ fractional error bars unless higher redshift observations are folded in. Information on the detailed spectral properties of the sources, which would enable us to discriminate between hot ISM or HMXB heating as well as precision constraints on other heating parameters will likely require power spectrum measurements at $z \gtrsim 10$. In the model we study here, HERA and SKA-low will be able to place $\approx 6 - 10\%$ constraints on heating parameters even if observations

in the FM band are not possible.

In this analysis, we have chosen to examine a single model in a large allowed parameter space. We do not think our predictions will change in different models by more than an order of magnitude based on trends that are well documented in the literature. It is shown in Pober et al. (2014) that HERA is capable of detecting the reionization peak over several orders of magnitude in $T_{\text{vir}}^{\text{min}}$ and a wide range of ζ and R_{mfp} values. The height of the heating peak is constant through 3-4 orders of magnitude in f_X (Mesinger et al., 2014) while the redshift of the peak remains approximately between 10 and 20 (Pacucci et al., 2014), hence the SNR on heating should not vary by more than an order of magnitude. An α_X that is larger than our fiducial value by a factor of 2 – 3 is shown, in Pacucci et al. (2014), to boost the amplitude of the heating peak by a factor of $\approx 2 - 3$ which we would expect to improve our constraints at a similar level. A much harder X-ray spectrum such as that discussed in Fialkov et al. (2014) leads to a reduction in the power spectrum amplitude by a factor of ≈ 2 but also shift the peak to lower redshifts. We would expect the combination of these effects to give results within an order of magnitude of those presented here.

In adopting the Fisher matrix technique, we have assumed that our likelihood function is Gaussian. Though this is a reasonable approximation about the ML point for small error bars, a more robust approach would be an MCMC calculation such as that presented in Greig and Mesinger (2015). It is for this reason that we do not give projections for the sensitivity of current arrays since projections of large error bars using the Fisher matrix are not self consistent. As of now, calculations of heating are not sufficiently rapid to allow for MCMC sampling of the likelihood function. Speeding up the heating calculation for suitability in MCMC is the subject of ongoing work.

Finally, with the exception of RFI, we have not attempted to directly address the fact that known systematic obstacles to 21 cm observing, such as the ionosphere, foreground brightness, and the increasing extent of the primary beam, become worse at lower frequencies and may pose challenges in addition to the the ones we address:

Parameters	Pessimistic Foregrounds			Moderate Foregrounds			Optimi
	HERA-127	HERA-331	SKA-1	HERA-127	HERA-331	SKA-1	HERA-127
$\Delta\zeta/\zeta^{\text{fid}}$	0.04	0.02	0.16	0.03	0.01	0.03	0.01
$\Delta R_{\text{mfp}}/R_{\text{mfp}}^{\text{fid}}$	0.03	0.02	0.10	0.02	0.02	0.04	0.01
$\Delta T_{\text{vir}}^{\text{min}}/T_{\text{vir}}^{\text{min,fid}}$	0.05	0.03	0.23	0.04	0.02	0.05	0.02
$\Delta f_X/f_X$	0.18	0.07	0.79	0.15	0.06	0.13	0.06
$\Delta\nu_{\text{min}}/\nu_{\text{min}}^{\text{fid}}$	0.19	0.07	0.79	0.15	0.06	0.12	0.07
$\Delta\alpha_X/\alpha_X^{\text{fid}}$	0.15	0.07	0.75	0.13	0.06	0.13	0.06

Table 4.1: The 1σ error forecasts for reionization and heating parameters on the instruments studied in this paper assuming 1080 hours of drift-scan observations on $\Delta z = 0.5$ co-eval bands over all redshifts between 5 and 25.

namely RFI and increased thermal noise. Observations in Ewall-Wice et al. (2016d) find that while ionospheric refraction does not appear to impact the level of foreground leakage beyond the wedge, the brighter foregrounds extended primary beam heighten the severity of any uncalibrated structure in the instrumental bandpass.

Our results indicate that precision measurements of the first high energy galactic processes can be expected from the upcoming generation of power spectrum experiments, provided that they exploit the information in the redshifts typically considered to precede reionization.




Figure 4-10: 95% confidence ellipses for our moderate and optimistic foreground models assuming 1080 hours of drift-scan observations on HERA-331. Heating parameters tend to be highly degenerate with each other but independent of reionization parameters. As the majority of the information heating comes from higher redshifts where thermal noise is much higher, their uncertainty regions tend to be several times larger.

Chapter 5

First Limits on the 21 cm Power Spectrum during the Epoch of X-ray heating.

*The content of this chapter was published in Ewall-Wice, A., Dillon, J. S., Hewitt, J. N., Loeb, A., Mesinger, A., Neben, A. R., Offringa, A. R., Tegmark, M., Barry, N., Beardsley, A. P., Bernardi, G., Bowman, J. D., Briggs, F., Cappallo, R. J., Carroll, P., Corey, B. E., de Oliveira-Costa, A., Emrich, D., Feng, L., Gaensler, B. M., Goeke, R., Greenhill, L. J., Hazelton, B. J., Hurley-Walker, N., Johnston-Hollitt, M., Jacobs, D. C., Kaplan, D. L., Kasper, J. C., Kim, H., Kratzenberg, E., Lenc, E., Line, J., Lonsdale, C. J., Lynch, M. J., McKinley, B., McWhirter, S. R., Mitchell, D. A., Morales, M. F., Morgan, E., Thyagarajan, N., Oberoi, D., Ord, S. M., Paul, S., Pindor, B., Pober, J. C., Prabu, T., Procopio, P., Riding, J., Rogers, A. E. E., Roshi, A., Shankar, N. U., Sethi, S. K., Srivani, K. S., Subrahmanyam, R., Sullivan, I. S., Tingay, S. J., Trott, C. M., Waterson, M., Wayth, R. B., Webster, R. L., Whitney, A. R., Williams, A., Williams, C. L., Wu, C. and Wyithe, J. S. B. (2016), ‘First limits on the 21 cm power spectrum during the Epoch of X-ray heating’, MNRAS **460**, 4320–4347.*

5.1 Introduction

Mapping the 21 cm transition of neutral hydrogen at high-redshift promises to revolutionize our knowledge on the first generations of stars and galaxies and to provide a unique probe of the “Dark Ages” preceding this first generation of luminous objects (see Barkana and Loeb (2001); Furlanetto et al. (2006a); Morales and Wyithe (2010) for reviews). Planned instruments such as the Square Kilometre Array (SKA) (Koopmans et al., 2015) and the Hydrogen Epoch of Reionization Array (HERA) (Pober et al., 2014) are expected to elucidate the formation of the first luminous structures and to place strict constraints on the properties of the sources that reionized the intergalactic medium (Pober et al., 2014; Greig and Mesinger, 2015). A number of experiments including the Giant Metrewave Telescope (GMRT) (Paciga et al., 2013), the Low-Frequency Array (LOFAR) (van Haarlem et al., 2013), the Murchison Wide-field Array (MWA) (Bowman et al., 2013; Tingay et al., 2013a), and the Precision Array for Probing the Epoch of Reionization (PAPER) (Parsons et al., 2010) are already underway to explore the challenges of separating the faint cosmological signal from bright foregrounds and to attempt a first detection of the power spectrum of the cosmological 21 cm emission line.

Thus far, these experiments have targeted redshifts between 6 and 12. During this Epoch of Reionization (EoR), ultraviolet photons from the first generations of luminous sources transformed the intergalactic medium (IGM) from predominantly neutral to ionized. Over the past several years, deep integrations have placed significant upper limits on the power spectrum during reionization (Paciga et al., 2013; Dillon et al., 2014; Parsons et al., 2014; Ali et al., 2015; Dillon et al., 2015a; Trott et al., 2016). The best upper limits of 502 mK^2 at $z \sim 8.4$ (Ali et al., 2015) have begun to rule out scenarios where the neutral IGM experiences little or no heating (Pober et al., 2015). Integrations at comparatively high redshifts have also been carried out. Dillon et al. (2014) put an upper limit on the power spectrum at $z = 11.7$ using the 32-tile MWA pathfinder. A much deeper integration at redshift 10.3 was performed with PAPER’s 32-element configuration (Jacobs et al., 2015), though it was limited

by residual foregrounds at the edge of the instrumental bandpass.

While observations of the power spectrum during the EoR alone will shed light on the sources and astrophysics that drove reionization, it is only the final milestone in the evolution of the neutral IGM. Before reionization, the gas was heated, most likely by the first generations of high mass X-ray binaries (HMXB) (Mirabel et al., 2011) and/or hot interstellar medium (ISM) (Pacucci et al., 2014). Brightness temperature fluctuations from inhomogenous heating at these early times can yield power spectrum amplitudes that are over an order of magnitude larger than those expected during reionization (Pritchard and Furlanetto, 2007; Mesinger et al., 2013). Even before the X-ray heating, fluctuations in the brightness temperature were likely sourced by fluctuations in the Lyman- α flux field from the first stars (Barkana and Loeb, 2005b; Pritchard and Furlanetto, 2006).

The ultimate goal of 21 cm cosmology is a three dimensional map of the entire IGM between $z \approx 200$ and reionization since, at least in principle, the 21 cm line is a cosmological observable accessible all the way back through the dark ages to the decoupling of the spin temperature from the cosmic microwave background (CMB) (Furlanetto et al., 2006a). Even though the ionosphere obscures extraterrestrial radio signals below about 30 MHz (Jester and Falcke, 2009), it is expedient to use ground based experiments to cover as great a redshift span as possible. Because foreground amplitudes and ionospheric effects grow progressively at lower frequency, the most reasonable next step after reionization in our march into the dark ages is the Epoch of X-ray heating (EoX). The exact redshift range for the EoX depends on the astrophysical model (see for example Mesinger et al. (2013, 2014); Pacucci et al. (2014)), but a reasonable range, targeted in this work, is $z=11.6$ (113 MHz) to $z=17.9$ (75 MHz) .

First experimental 21 cm constraints on the thermal history of the IGM come from Pober et al. (2015) who used upper limits on the power spectrum at reionization redshifts to rule out inefficient heating histories. These constraints arise from the fact that the observable brightness temperature difference from the CMB depends on the spin temperature as

$$\Delta T_b \propto \left(1 - \frac{T_\gamma}{T_s}\right), \quad (5.1)$$

where T_γ is the temperature of the CMB and T_s is the spin temperature of the gas which is expected to be closely coupled to the gas kinetic temperature before substantial heating takes place (Furlanetto et al., 2006a). For a cold IGM, $1 - T_\gamma/T_s$ is large and negative leading to large amplitude contrasts between neutral and ionized regions.

However, assuming that the number of X-rays per baryon involved in star formation is the same as what is observed in nearby star forming galaxies (Mineo et al., 2012a), the HI spin temperature is expected to be heated well above the CMB by the time reionization begins, saturating the effect of heating on equation 5.1 (Furlanetto, 2006b). Hence, direct measurements of the 21 cm line during the EoX will be necessary if we want to learn about the detailed properties of the thermal history and the astrophysical phenomena that influence it. Recent work has shown that if X-ray heating proceeds inefficiently, the current generation of interferometers will be sensitive enough to detect the power spectrum sourced by spin temperature fluctuations at $z \approx 12$ (Christian and Loeb, 2013). Next generation of 21 cm observatories will detect the heating power spectrum for a wide range of heating scenarios out to redshifts as high as 20 (Mesinger et al., 2014) and place percent level constraints on the properties of the earliest X-ray sources (Ewall-Wice et al., 2016d). While pre-reionization measurements are expected to shed light on the first stellar mass black holes or the hot ISM, they may also offer us insights into other astrophysical processes. It is possible for dark matter annihilation (Valdés et al., 2013) and the existence of warm dark matter (Sitwell et al., 2014) to create observable impacts on the IGM thermal history. Finally, the IGM is especially cool and optically thick during the beginning of the heating process, making it ideal for 21 cm forest (Furlanetto and Loeb, 2002; Carilli et al., 2002; Furlanetto, 2006a; Mack and Wyithe, 2012; Ciardi et al., 2013) studies should any radio loud sources exist at those redshifts. It is also possible to constrain the source population itself by detecting its signature in the 21 cm power spectrum (Ewall-Wice et al., 2014).

Complementary observations of the sky-averaged (the “global”) 21 cm signal with a single dipole can also explore the reionization and pre-reionization epochs and exper-

iments such as EDGES (Bowman and Rogers, 2010), LEDA (Greenhill and Bernardi, 2012), DARE (Burns et al., 2012), SARAS (Patra et al., 2013), SciHI (Voytek et al., 2014), and BIGHORNS (Sokolowski et al., 2015) are beginning to take data. While demanding much greater sensitivity than global signal experiments, power spectrum measurements with an interferometer probe fine frequency scales while foregrounds occupy a limited region of Fourier-space, known as the “wedge” (Datta et al., 2010; Parsons et al., 2012a; Morales et al., 2012; Vedantham et al., 2012; Trott et al., 2012; Hazelton et al., 2013; Thyagarajan et al., 2013; Liu et al., 2014a,b; Thyagarajan et al., 2015a,c). The region of Fourier space outside of the wedge, in principle free of foregrounds and therefore having greater sensitivity to brightness temperature fluctuations, is known as the “EoR window” (henceforth “window”). foreground modeling and calibration errors that are smooth in frequency should have limited impact within the window.

In this paper we assess the levels of systematic errors that are especially severe at the lower EoX frequencies (relative to those typical of EoR studies) including the ionosphere, radio-frequency interference (RFI), and the enhanced noise and foregrounds from a sky that is both intrinsically brighter at lower frequencies and observed with a larger primary beam. In § 5.2 we describe the MWA, our observations, and data reduction. In § 5.3 we address the systematic errors that are especially challenging below EoR frequencies and our efforts to mitigate them. The limiting systematic error that we encounter is fine frequency structure in the instrumental bandpass due to standing wave reflections in the cables between the MWA’s beamformers and receivers. After making a reasonable assumption about the relationship between our autocorrelations and the gain amplitudes, we achieve notable improvement in calibration but we are still left with significant foreground contamination.

Power spectrum upper limits are derived (§ 5.4) which are broadly consistent with thermal noise except in several regions of Fourier space corresponding to the light travel time delays of the reflections. We expect that refined calibration techniques employing better foreground models (Carroll et al., 2016) and redundant baselines (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014) can improve the removal of this

contamination. In order to avoid signal loss and the introduction of spurious spectral structure, we have been conservative in the number of free parameters allowed in our gain solutions; increasing these may also resolve this problem. Reduced cable lengths expected in upcoming experiments such as HERA and the SKA will ameliorate the problem of reflections.

5.2 Observing and Initial Data Reduction

We begin our discussion with an overview of our observations and our data reduction procedure. Our analysis yields two different image products with 112 s cadence: high resolution continuum images created from bandwidth multifrequency synthesis (MFS) (with $\approx 6'$ resolution) where baselines across all fine frequencies are combined into a single image, and naturally weighted multifrequency data cubes, where each fine frequency channel is imaged separately and integrated over three hours. We use the MFS images to evaluate ionospheric conditions, and we use the multifrequency data cubes in our power spectrum analysis. We note that with 112 s averaging, we are performing significant averaging over fine time-scale ionospheric effects which for the MWA baselines have a typical coherence time of $\approx 10 - 44$ s (Thyagarajan et al., 2015b) (henceforth V15a). After outlining the instrument and our observing strategy (§ 5.2.1), we discuss our initial calibration procedure (§ 5.2.2) and finish with the production MFS images (§ 5.2.3) and data cubes (§ 5.2.4) which serve as the input to our power spectrum pipeline.

5.2.1 Observations with the MWA

The MWA (Lonsdale et al., 2009; Tingay et al., 2013a) is a 128 antenna interferometer located at the Murchison Radio Observatory (MRO) in Western Australia (26.70°S, 116.67° E) with an analog bandpass of 80-300 MHz. Each correlated antenna tile consists of 16 dual polarization dipole elements arranged in a four-by-four grid. The phased output of the dipoles on each tile is summed together in an analog beam-former and delivered to one of sixteen different receiver units where 30.72 MHz of bandwidth

is digitized before correlation in an on-site building. We refer the reader to Prabu et al. (2015) and Ord et al. (2015) for a detailed discussion of the MWA’s receivers and correlator, respectively. The instrument is designed to achieve a diverse set of science goals (Bowman et al., 2013) including a first detection of the 21 cm power spectrum during the EoR, detecting and monitoring transients, pulsars (Tremblay et al., 2015), solar and heliospheric science (Tingay et al., 2013b), and a low-frequency survey of the sky below DEC=+25° (Wayth et al., 2015).

Observations of a field centered at R.A.(J2000) = 4^h0^m0^s and decl.(J2000)=−30°0′0″ were carried out for 4.13 hours over two nights on September 5th and 6th, 2013, with primary beams of the 128 MWA antenna elements (“tiles”) formed at five different altitude/azimuth pointings each night to track the field. After flagging for RFI and anomalous behavior that we will discuss in detail in § 5.3.1, our total observation time for our power spectrum upper limit is 3.08 hours. In Fig. 5-1 we show the relative integration time on the sky weighted by the primary beam over all observations. We observed with 40 kHz spectral resolution simultaneously over two contiguous bands; a 16.64 MHz interval between 75.52 MHz and 90.88 MHz (Band 1) and a 14.08 MHz band between 98.84 MHz and 112.64 MHz (Band 2). Both of these sub-bands overlap with the FM band (88-108 MHz). In Fig. 5-2 we show our observed bands superimposed on the autocorrelation spectrum of a single MWA tile. Observations in Band 1 took place right on the edge of the analogue cutoff of the MWA, making its shape relatively complicated to model in calibration (our calibration is direction independent and described in § 5.2.2 and § 5.3.4). Band 2, while in a flatter region of the bandpass, has a larger overlap with the FM band. Observations were divided into 112 s snapshot with data averaged into 0.5 s integration intervals by the correlator.

5.2.2 RFI flagging and Initial Calibration

The data were first flagged for RFI contamination. An optimized version of `aoflagger` (Offringa et al., 2012), called `cotter` (Offringa et al., 2015), was run on each snapshot with automatic RFI identification performed only on the visibility cross correlations. Additional flags were applied to the center and 40 kHz edges of each 1.28 MHz reso-




Figure 5-1: A radio map at 408 MHz (Haslam et al., 1982) sin-projected over the region of the sky observed in this paper. Cyan through magenta contours indicate the total fraction of observation time weighted by our primary beam gain for our three hours of observation at 83 MHz. Red contours indicate R.A.-decl. lines. Observation tracked the position (R.A.(J2000) = $4^h0^m0^s$, decl.(J2000) = $-30^\circ0'0''$) on a region of the sky with relatively little galactic contamination and dominated by the resolved sources Fornax A and Pictor A. The galactic anticentre and bright diffuse sources, such as the Gum Nebula, are below 1% bore-site gain.




Figure 5-2: The autocorrelation spectrum of a single tile, showing the MWA bandpass, is plotted here (solid black line) along with the frequency ranges over which our data was taken (gray striped rectangles). Observations were performed simultaneously on two non-contiguous bands located on either side of the 88-108 MHz FM band (red shaded region) to both assess conditions within the FM band and to preserve some usable bandwidth should it have proven overly contaminated by RFI.

lution coarse spectral channel to remove digital artifacts arising from the two stage channelization scheme used in the MWA. After flagging, visibilities were averaged in time from 0.5 to 2 s and in frequency from 40 to 80 kHz. The averaged visibilities were then converted to Common Astronomy Software Applications Package (CASA) measurement sets (McMullin et al., 2007) which served as the inputs to all ensuing steps in our pipeline. The percentage of all data flagged by `cotter` was approximately 0.75% for Band 1 and 2% for Band 2. As described in § 5.3.1, we also implement additional (and highly conservative) flagging of observations by inspecting autocorrelations which leads us to discard 25% of our data. Our initial calibration was divided into three steps: A preliminary complex antenna gain solution using an approximate sky model, one iteration of self calibration, and polynomial fitting to reduce noise and limit fine frequency scale systematics that might arise from our incomplete sky model.

For the first step, our sky model combined a list of point sources with images of the two bright resolved sources in our field: Fornax A and Pictor A. For the point source model we included the 200 brightest sources at our frequencies based on extrapolated power law fits to data from the Coolgoora survey (Slee, 1995), the Molongolo Reference Catalogue (Large et al., 1981) and more recent measurements by PAPER at 145 MHz (Jacobs et al., 2011). Fornax A is the brightest extended source in our field and is highly resolved, so we modeled it with a VLA image taken by Fomalont et al. (1989) at 1.4 GHz, scaled to match the flux density and spectral index measured by Bernardi et al. (2013) at 180 MHz, and extrapolated to our band with a spectral index of -0.88 . For Pictor A, we used a VLA image at 333 MHz by Perley et al. (1997) and extrapolated to our band with a spectral index of -0.71 . The model components for our initial calibration extended down to an apparent flux density of ≈ 5 Jy which is comparable to the flux uncertainties in the brightest sources in the initial catalog. Due to the high uncertainty in the sky at our frequencies, this model was updated by a round of self calibration which we describe shortly.

The CASA bandpass function was used to obtain a first set of best fit calibration gains, averaging over 32 fine channels for each solution. Since our starting model was

uncertain at the 10% level, an iteration of self calibration was performed by MFS imaging and deconvolving 10^4 components with WSClean (Offringa et al., 2014). The CLEAN components were used as a model for a second run of `bandpass` where we relaxed channel averaging so that the complex gain for each 80 kHz fine channel was found independently. In Fig. 5-3 we show the fractional change in calibration amplitude for one antenna tile over time intervals in which beamformer settings were held constant (pointings). Over our two nights we find that variations are on the order of a few per cent and dominated by uniform amplitude jumps across the entire frequency range that are strongly correlated between all antennas. Observations of the autocorrelations do not show such behavior within each pointing so these gain jumps must arise from the calibration routine itself. Possible sources of these jumps could be variations in the the overall amplitudes in self-calibration which occurs if the cleaning step does not recover all of the flux density on large scales, unmodeled sources moving through the sidelobes of the primary beam, or the result of a varying sky in the presence of beam modeling errors. Since these features do not introduce fine frequency structure, we do not think they are an impediment to the power spectrum analysis in this work.

We attempted to apply calibration solutions in which each frequency channel was allowed to vary freely but we found significant power was introduced on fine scales, contaminating the EoR window. This could be due to unmodeled sources adding spectral structure from long to short baselines or insufficient SNR on the solutions themselves. There is also possibility for signal loss given the large number of degrees of freedom. The degree to which calibration can remove signal and how unmodeled sources with enhanced spectral structure at long baselines can be mixed into short baselines are open questions that are beyond the scope of this analysis and are being investigated by Barry et al. (2016); Trott and Wayth (2016); Ewall-Wice et al. (2016c). With these concerns in mind, we erred on the side of caution and fit each j^{th} antenna gain with the product of three smooth functions in frequency, f ,

$$g_j(f) = P_j(f)R_{j,7m}(f)B(f), \quad (5.2)$$




Figure 5-3: A false-color plot of the fractional change in our calibration amplitudes over each pointing (in which beamformer delay settings are fixed). Pointing changes are marked by the solid white horizontal lines. We see that the calibration amplitudes vary within a pointing on the order of several percent with little systematic variation. There are several coarse channel scale jumps on September 5th that correspond to observations in which the number of sources identified within a snapshot image were reduced (see Fig. 5-9). We found that these jumps corresponded to excess flagging from `cotter`, indicating high levels of RFI or other bad data, and dropped them from our analysis. Vertical lines from RFI are visible at 98 MHz and 107 MHz, especially on September 6th.

where $P_j(f)$ is a third order polynomial in amplitude and first order polynomial in phase; $B(f)$ is a median bandpass that accounts for the coarse band shape determined by taking the median of calibration amplitudes across all tiles and polarizations; and $R_{j,7m}(f)$ is a reflection function that accounts for standing wave ripples in the bandpass arising from impedance mismatches in the 7-meter low noise amplifier (LNA) to beamformer cable connections. It is straightforward to show (see Appendix 5.A) that a cable with length $L(j)$ introduces a multiplicative reflection term to the overall gain,

$$R_j(f) = \frac{1}{1 - r_j e^{i(2\pi f \tau_{L(j)} + \phi_j)}}, \quad (5.3)$$

where $r_j e^{i\phi_j}$ is a complex coefficient that is a function of the impedance of the cable and its connections at either end and $\tau_{L(j)}$ is the time for a signal to travel from one end of the cable and back. We note that this multiplicative term is derived from the infinite sum of reflections occurring from each n^{th} round trip of the reflected wave and hence accounts for higher order reflections, not just the first round trip contribution.

When we formed power spectra from visibilities calibrated by this initial method, we found that our beamformer to receiver cables introduced spectral structure at the $\lesssim 1\%$ level into our instrumental bandpass which were not removed by this smooth model (we return to this issue in § 5.3.3). The effect of this spectral structure on the MFS images, used to measure ionospheric refraction in § 5.3.2, was negligible and we therefore employed this calibration for their production. A more refined calibration procedure, which we describe in § 5.3.4, was used for the data cubes and our power spectra.

5.2.3 MFS Imaging and Flux Scale Corrections

To form MFS images, we averaged the antenna phases over each night and held them constant for every snapshot calibration solution to average over time-variability introduced by the ionosphere. This was done in order to ensure that short time-scale snapshot-to-snapshot variations were not due to time variations in the calibration solutions caused by the ionosphere. A multifrequency synthesis image was created

from each snapshot, band, and polarization. From each XX and YY polarization snapshot which we call I_{XX} and I_{YY} respectively, we created a Stokes I snapshot corrected for the primary beam using an analytic dipole model of the MWA primary beam (which we denote as B_{XX} and B_{YY}).

$$I_I(\theta, \phi) = \frac{I_{XX}(\theta, \phi)B_{XX}(\theta, \phi) + I_{YY}(\theta, \phi)B_{YY}(\theta, \phi)}{B_{XX}^2(\theta, \phi) + B_{YY}^2(\theta, \phi)}. \quad (5.4)$$

Sources were identified using the *Aegean* source finder (Hancock et al., 2012). An overall flux scale for each stokes I snapshot was set following the technique, described in Jacobs et al. (2013), in which the flux scale for all sources is simultaneously fit to catalog flux densities using a Markov chain Monte Carlo method. We used the ten highest signal-to-noise point sources in each field and catalog flux densities interpolated between 74 MHz measurements from the Very Large Array Low-Frequency Sky Survey Redux (Lane et al., 2014) and 80 and 160 MHz measurements from the Culgoora catalog (Slee, 1995). Since a flux scale error has no frequency dependence and the errors themselves evolve slowly in time near the center of the primary beam, we do not think that such mismodeling will result in frequency dependent errors. The dominant uncertainty in our flux scale is the systematic uncertainty in the model source fluxes themselves which are on the order of $\approx 20\%$ (Jacobs et al., 2013) while uncertainties in the beam model contribute at the several percent level (Neben et al., 2015). On September 6th, we observed systematically smaller source counts than on September 5th (Fig. 5-9). We attribute this difference to greater ionospheric turbulence on September 6th; and discuss this result further in § 5.3.2.

In Fig. 5-4 we show a deep, primary beam corrected, integration of a portion of our field formed from a MFS image of Band 1. Known sources are well reproduced. The diffuse structures of the Vela and Puppis supernova remnants are clearly visible along with the fine scale structure of Fornax A.

Similar to previous upper limits on the power spectrum (Dillon et al., 2014; Parsons et al., 2014; Dillon et al., 2015a; Ali et al., 2015; Trott et al., 2016), our analysis does not consider the cross polarization products from the interferometer, which




Figure 5-4: A deep image of the MWA “EoR-01” field centered at (R.A.(J2000) = $4^h0^m0^s$ and decl.(J2000) = $-30^\circ0'0''$), derived by stacking restored multifrequency synthesis Stokes XX and YY images produced by `WSClean` on Band 1. The dominant source in our field is Fornax A (detailed in the inset) whose structure is well recovered in imaging. Pictor A is also present at the center of the image (at $\sim 30\%$ beam) along with the diffuse Puppis and Vela supernova remnants on the left.

would require an additional calibration step to solve for the arbitrary phase difference between the X and Y polarized arrays (Cotton, 2012; Moore et al., 2017).

5.2.4 Data Cubes for Power Spectrum Analysis

The inputs for our power spectrum analysis are multifrequency data cubes further calibrated with autocorrelation data from each tile. We describe the autocorrelation calibration technique in § 5.3.4 where we address systematic errors revealed by our first look at data cubes and power spectra. In this section we describe our procedure for building the data cubes. While the MWA’s long baselines, extending out to 2864 m, are useful for gauging ionospheric conditions and other systematics, the uv plane is completely filled out to $\lesssim 160$ m after ≈ 3 hours of rotation synthesis. Since forming image power spectra using data with incomplete uv coverage leads to unwanted mode-

mixing and spectral artifacts (Hazelton et al., 2013), we threw away the sparse regions of the uv plane and reprocessed the data for our power spectrum analysis at much lower angular resolution than our MFS images. For each 112s snapshot, we divided the data into even and odd time step visibility sets formed from every other two second integration step, which we later cross multiply to form power spectra without noise bias (Dillon et al., 2014). A naturally weighted, dirty snapshot cube was produced for each set with 80 kHz spectral resolution, 80 pixels on a side and 1.0° (0.75°) resolution for Band 1 (2).

To form a power spectrum we will have to uniformly weight the sum of these cubes, hence the array point spread function (PSF), $\mathbf{s}_{j,even/odd}$, which is the 2d Fourier transform of the number of samples within each uv cell; and the analytic primary beam matrix, \mathbf{B}_j , were also saved for each snapshot, polarization and interleaved time step. The power spectrum formalism of (Dillon et al., 2013) assumes the flat sky approximation and requires a sufficiently small field of view to be computationally tractable. Hence, before creating a uniformly weighted data cube by stacking the naturally weighted images in the uv plane and dividing by their cumulative sampling function, we cropped each snapshot from 80×80 to 24×24 pixels.

We created even and odd, uniformly weighted temperature cubes by summing all naturally weighted snapshots across both the XX and YY polarizations and dividing by the sum of their cumulative sampling function in the uv plane,

$$\hat{\mathbf{x}}_{even/odd} = \frac{\lambda^2}{2k_b \Omega_{pix}} \mathcal{F}_2^{-1} \left[\frac{\sum_j \mathcal{F}_2 \mathbf{B}_j \hat{\mathbf{x}}_{j,even/odd}}{\sum_j \mathcal{F}_2 \mathbf{B}_j^2 \mathbf{s}_j} \right], \quad (5.5)$$

where the division of the two sums is to be understood as element-wise division. \mathcal{F}_2 denotes the two dimensional Fourier transform matrix. Indexing by angle cosines, ℓ, m , their duals, u and v , and frequency, f , we may write \mathcal{F}_2 and its inverse as,

$$[\mathcal{F}_2]_{uvf' \ell mf} = \Omega_{pix} e^{-2\pi i(\ell u + mv)} \delta_{f'f} \quad (5.6)$$

$$[\mathcal{F}_2^{-1}]_{\ell m f uv f'} = \frac{1}{N_\ell N_m} \Omega_{pix} e^{2\pi i(\ell u + mv)} \delta_{f'f}, \quad (5.7)$$

where Ω_{pix} is the solid angle of each pixel. Note that we are dividing by the convolution of the sampling function with the beam squared in the uv plane so we do not have to worry about dividing by the beam nulls in image space. The prefactor at the front of equation 5.5 converts from brightness to temperature units.

By dividing by the point spread function in the uv plane, equation 5.5 is essentially the application of uniform weighting to our data with some additional factors of the primary beam which warrant explanation. An additional factor of the primary beam was included in the sum to upweight regions of the field to which the MWA has the greatest gain, and is equivalent to optimal mapmaking (Tegmark, 1997b; Dillon et al., 2015b). Other 21 cm pipelines, notably Fast Holographic Deconvolution (Sullivan et al., 2012) perform a similar upweighting by gridding visibilities with the primary beam while the fringe rate filtering procedure (Parsons et al., 2016b) weights without gridding or imaging at all.

The impact of multiplying by the two factors of the beam in 5.5 has the effect of convolving the true visibilities the uv space beam convolved with its complex conjugate. Since our uv cells are quite large, this is well approximated by multiplication of the visibilities by the convolution of the uv beam with its complex conjugate. We therefore also include the factor \mathbf{B}^2 in the denominator to correctly normalize the data in the uv plane.

Our $\hat{\mathbf{x}}$ estimate of \mathbf{I} in equation 5.5 is similar to the Stokes-visibility \mathbf{I} approximated in previous power spectrum analyses (Dillon et al., 2014; Parsons et al., 2014; Dillon et al., 2015a; Ali et al., 2015). Visibility stokes \mathbf{I} is susceptible to leakage from visibility stokes $Q \equiv \frac{1}{2}(XX - YY)$ due to beam ellipticity which has the potential to introduce fine frequency structure, caused by Faraday rotation, into the EoR window (Jelić et al., 2010; Geil et al., 2011; Moore et al., 2013; Jelić et al., 2014; Moore et al., 2017; Asad et al., 2015). We can obtain an upper limit on polarization leakage in our power spectrum estimate by considering the best upper limits to date of the polarized Q, U , and V visibility power spectra that might leak into our \mathbf{I} visibilities, measured by Moore et al. (2017) over the large field of view of PAPER. In this analysis the authors place limits of $\approx 5 \times 10^4 \text{ mK}^2$ on Q, U , and V at $\approx 120 \text{ MHz}$. Scaling by

the frequency dependence of the sky, the polarization power spectrum levels at our frequencies should be below $\approx 5 \times 10^4 \text{ mK}^2 (80 \text{ MHz}/120 \text{ MHz})^5 \approx 3.7 \times 10^5 \text{ mK}^2$. The leakage from Q/U to I is given by equations 15 and 16 in Moore et al. (2017), and is equal to the product of the polarized power spectrum and the ratio between the integral of the differences of the X and Y polarized beams squared and the integral of the sum of the polarized beams squared. Using a short dipole model of our primary beam, we find that this ratio for the MWA beam is $\approx 5 \times 10^{-3}$. We therefore estimate an upper limit on the stokes Q,U to I leakage in our power spectra to be $\approx 5 \times 10^{-3} \times 3.7 \times 10^5 \text{ mK}^2 \approx 1.5 \times 10^3 \text{ mK}^2$. This is slightly larger than the EoX power spectrum which is anticipated to be several hundred mK^2 (Pritchard and Furlanetto, 2007; Santos et al., 2008; Baek et al., 2010; Mesinger et al., 2013) so it is still possible that polarized leakage may limit a detection unless direction-dependent polarization corrections are applied. However, this number is an upper limit and the actual leakage is probably lower. As of now, the most sensitive limits on the EoR power spectrum formed from XX and YY visibilities (Ali et al., 2015) limit $(Q,U) \rightarrow I$ leakage from the similarly elliptical PAPER beam to below the level of $\approx 500 \text{ mK}^2$ between $k_{\parallel} \approx 0.2 - 0.5 \text{ hMpc}^{-1}$ at 150 MHz while Asad et al. (2015) predict stokes polarized power spectrum from observations of the 3C196 field at $\approx 142 \text{ MHz}$ to be at the level of only $10^2 - 10^3 \text{ mK}^2$ at $k_{\perp} \lesssim 0.1 \text{ hMpc}^{-1}$ (see their figure 12, panel a). Scaling this up by $(140 \text{ MHz}/80 \text{ MHz})^5 \approx 16$ to account for the increasing sky temperature and applying our ellipticity factor of 5×10^{-3} gives a polarization leakage of $\approx 8 - 80 \text{ mK}^2$ which is still several times smaller than the predicted amplitude of the EoX power spectrum.

In order to reduce artifacts from aliasing at the edges of the coarse channels due to the two-stage polyphase filter bank, we flagged 240 kHz at each side. Finally, uv cells with poor sampling were flagged since sampling and noise in these cells can change rapidly with frequency, leading to fine frequency artifacts (Hazelton et al., 2013).

To check our flux scale, we estimate the level of thermal noise and the system temperature from the difference of our even/odd interleaved cubes, $\hat{\mathbf{x}}_{\text{even}} - \hat{\mathbf{x}}_{\text{odd}}$. Since the PSF is virtually identical between 2 second time steps, each pixel of $\hat{\mathbf{x}}_{\text{even}} - \hat{\mathbf{x}}_{\text{odd}}$

in the (u, v, f) basis has zero mean and a variance of (Thompson et al., 1986)

$$\langle |[\mathcal{F}_2 \widehat{\mathbf{x}}]_{uvf}|^2 \rangle = \frac{\lambda^4 T_{sys}^2(f)}{2A_e^2(f)t(u, v, f)df}, \quad (5.8)$$

where df is the fine channel frequency width, T_{sys} is the system temperature, and $t(u, v, f)$ is the total integration time in the uv cell equal to the sampling function in the uv plane multiplied by the 2 second integration time step dt ,

$$t(u, v, f) = \sum_j [\mathcal{F}_2 \mathbf{B}_j^2 \mathbf{s}_j]_{uvf} dt. \quad (5.9)$$

$A_e(f)$ is the effective area of the MWA tile at frequency f computed from an analytic dipole model. We may determine T_{sys} by taking the ratio of the variance across uv cells at each frequency in $\widehat{\mathbf{x}}_2 - \widehat{\mathbf{x}}_1$ and the average across uv cells of our model variance predicted by equation (5.8) without the T_{sys} factor. We find that $T_{sys}(f) = 2091\text{K}$ (1139 K) at 83 (106) MHz with an error of $\approx 20\%$ which is dominated by the systematic uncertainty in the fluxes of the sources used to set our flux scale. Assuming a spectral index of -2.6 (Rogers and Bowman, 2008; Fixsen et al., 2011), these values imply a system temperature of $T_{sys} \approx 470\text{K}$ at 150 MHz, consistent with what is found at higher frequencies in (Dillon et al., 2015a).

In Fig. 5-5 we show the standard deviation across all uv cells at each frequency in our Band 1 difference cube along with the square root of the mean of our model variances at each frequency assuming $T_{sys} \propto f^{-2.6}$ and normalized to the center channel. We see that they are in good agreement. We also find that the standard deviation across frequency in each uv cell is consistent with the square root of the mean across frequency of our model variances.

An interesting question is whether or not our determination of T_{sys} is contaminated by ionospheric scintillation noise. For baselines within the Fresnel radius, $r_F = \sqrt{\lambda h / (2\pi)}$, of an ionospheric plasma screen of height h (which is the case for the MWA core), V15a determine the coherence time for scintillation noise to be set by the length of time it takes for overhead plasma, traveling at velocity v , to cross the Fresnel radius, $2r_F/v$. For typical plasma velocities of $\approx 100 - 500 \text{ km s}^{-1}$ and

$h \approx 600$ km, consistent with measurements (Loi et al., 2015b), we obtain coherence times between 10 and 44 seconds at 83 MHz, which is significantly greater than the two second interleaving of our data cubes. Hence, ionospheric fluctuations are likely subtracted away in our differencing on two second intervals. For a much lower altitude of 100 km, the correlation time is still ≈ 4 seconds at 83 MHz. Even if there was still some variation between the time slices, we can put an upper limit on the variation relative to thermal noise by comparing the amplitude of scintillation noise we might expect given our primary beam and the parameters of the phase power spectrum of the ionospheric fluctuations that we determine in § 5.3.2. We find that the level of scintillation noise on a single visibility in a two second snapshot (appendix 5.C) is only $\lesssim 2\%$ the thermal noise level. Further suppression of the scintillations comes from the fact that we approximate T_{sys} at each frequency by taking the standard deviation across the uv plane within the Fresnel zone, in which the noise is expected to be coherent and would not contribute to such a standard deviation. The coherence in frequency of the ionospheric fluctuations (V15a,b) would also suppress their contribution to the standard deviation across frequency in each uv cell. For these reasons, we expect scintillation noise to have a very small contribution to our determination of T_{sys} at or below the 1% level.

5.3 Addressing the Challenges of Low Frequency Observing

A number of systematics associated with observing the EoR become dramatically more challenging as one moves to higher redshift. Because the Epoch of X-ray heating (EoX) spans the FM band, we expect enhanced RFI contamination. The ionosphere’s influence on electromagnetic wave propagation increases with wavelength, though its smooth evolution in frequency should cause its impact on source mis-subtraction and calibration to be contained within the wedge (Trott and Tingay, 2015; Vedantham and Koopmans, 2016) (henceforth V15b). Moving down in frequency, the larger




Figure 5-5: Left: The standard deviation over uv cells as a function of frequency for an even/odd time step difference cube after three hours of integration. The blue line is derived from data, while the green line is a model with a system temperature of 470 K at 150 MHz and a spectral index of -2.6 . Spikes in the standard deviation are present at the center of each coarse channel since the center channel has one half of the data due to flagging the center channel which is contaminated by a digital artifact before averaging from 40 to 80 kHz. Right: The ratio of variance taken over frequency in each uv cell and our variance model using the same system temperature as on the left. The ratio between our model and observed variance is close to unity across the uv plane. White cells indicate uv voxels that were flagged at all frequencies due to poor sampling. All data in this figure is from Band 1.

primary beam extends foreground emission to higher delays and hence larger k_{\parallel} while the foregrounds increase rapidly in brightness temperature as $\approx f^{-2.6}$ (Rogers and Bowman, 2008; Fixsen et al., 2011). Finally, spectral structure in our gains at fixed delays move down in k_{\parallel} at higher redshift and, due to the increased primary beam width, occupy a greater extent in k -space as well (Thyagarajan et al., 2015a). In this section we determine the impact (if any) of each these obstacles on our power spectrum analysis and describe our strategies for mitigating them. We deal with RFI, ionosphere, and spectral structure in § 5.3.1, § 5.3.2, and § 5.3.3/5.3.4 respectively.

5.3.1 Radio-Frequency Interference

As explained above, automated RFI detection and flagging was performed using `cotter` on the 0.5 s, 40 kHz resolution cross correlations before time and frequency averaging. To illustrate the time-frequency structure of RFI contamination, we plot the fraction of visibilities flagged at each fine frequency channel and 112 s snapshot




Figure 5-6: The percentage of visibilities flagged for RFI by `cotter` as a function of time and frequency. White regions indicate missing data including the coarse band edges and blue-dashed vertical lines indicate the edge of the FM band. While Band 1 remains predominantly clear with a few sparse events within the lower end of the FM band, contamination is significantly greater over Band 2. Even in the FM band, RFI events are either isolated in frequency or time, allowing us to flag them. A handful of observations in Band 1 on September 5th are missing entire coarse channels which we also discard. Bar plots on the bottom and right show the averages of the RFI flagging fraction over time and frequency respectively.

interval in Fig. 5-6. One can see that the majority of Band 1 is clear of RFI. Even within the region overlapping with the FM, events are sparse in time and frequency. In Band 2 we see significantly greater interference, especially in the two lowest coarse channels. There are clearly a greater number of events contained within the FM band however they only appear intermittently with the exception of a handful of 40 kHz fine-channels. The existence of intermittent FM signals, even in a radio quiet site such as the MRO is possibly due to signals from over the horizon, reflected off of the bottom of the ionosphere

It is possible for interference that is present for extended periods of time but weak enough to remain below the 112s noise floor over which RFI is flagged to contaminate our data. RFI can also make it past flagging through calibration solutions which are

derived from the autocorrelations on which `cotter` does not perform flagging (§5.3.4). When we first created integrated data cubes there were clear signs in Band 2 that some low level RFI contamination remained in the form of ripples in the power spectrum and spikes in the frequency domain of our gridded visibility cubes. We identified and discarded observations that appeared to contain increased flagging for a wide range of frequencies and completely flagged any channels that contained spikes in our final data cubes.

The lower two coarse channels in Band 2 were contaminated for a wide range of times (Fig. 5-6) and were thus excluded from our power spectrum cubes entirely. In addition, we flagged a total of 9, 80 kHz channels (7 % of our data) which appeared to be contaminated by RFI at contiguous intervals over a significant number of observations; 104.08, 104.48, 106.08, 106.24, 106.32, 107.44, 107.60, and 107.84 MHz. The rest of the FM band appears clean after routine `cotter` flagging is applied. After these channels are discarded, we see no evidence that our three hour power spectrum results are limited by RFI (see the end of § 5.4.1 for further discussion).

5.3.2 Ionospheric Contamination

The refraction induced by gradients in the total electron content (TEC) of the ionosphere scales as λ^2 and is therefore expected to be more severe at our frequencies than those associated with the EoR. In this section we quantify the severity of ionospheric conditions over our observations by measuring the differential refraction of point sources relative to known catalog positions. We find that the ionospheric gradients change considerably over the duration of our observations, despite nearly constant and mild solar weather indicators (Table 5.1). However, when we form power spectra from data with nearly a factor of two difference in the observed gradients, we see no effect on the power spectrum within the EoR window (Fig. 5-22). The analysis presented here is meant to assess the level of refraction and, in § 5.4.3, its impact on the power spectrum. Readers who are interested in a more detailed analysis of TEC gradients over the MRO and an interpretation of their physical origin or their impact on time domain astrophysics should consult Loi et al. (2015b,a).

Radiation passing through a plasma of electrons with spatial density $N_e(r)$ acquires a phase given by (Rybicki and Lightman, 1979)

$$\phi \approx \frac{e^2}{cm_e f} \int N_e(r) dr, \quad (5.10)$$

where c is the speed of light, e is the electron charge, and m_e is the electron mass. If we assume the ionosphere is a flat screen of plasma at height h and that the TEC changes linearly over scales comparable to the separation between antenna pairs in an interferometer, $\mathbf{x}_i - \mathbf{x}_j = \mathbf{b}_{ij}$, the visibilities formed by cross multiplying and averaging the electric field measured by two antennas is given by (see V15a for a derivation)

$$\begin{aligned} V_{ij} &= \langle E(\mathbf{x}_i) E^*(\mathbf{x}_j) \rangle \\ &\approx \int d^2 \mathbf{s} I(\mathbf{s}) e^{i \nabla \phi(\mathbf{s}) \cdot \mathbf{b}_{ij}} e^{-2\pi i \mathbf{s} \cdot \mathbf{b}_{ij} f / c}, \end{aligned} \quad (5.11)$$

This is the standard equation for radio interferometric visibilities (Thompson et al., 1986) where the angular positions \mathbf{s} of the sources with intensity $I(\mathbf{s})$ have been modified to be

$$\mathbf{s} \rightarrow \mathbf{s}' = \mathbf{s} + \frac{c}{2\pi f} \nabla \phi(\mathbf{s}). \quad (5.12)$$

where $\nabla \phi(\mathbf{s})$ is the gradient of the phase screen with respect to the E-W and N-S directions. Hence one can measure gradients in the TEC by observing offsets in the positions of point sources. Note that the gradients themselves are a function of position $\nabla \phi(\mathbf{s})$. While we focus on refractive effects as a proxy for ionospheric gradients, we also note that a significant number of the MWA baselines are within the Fresnel radius at these frequencies and experience significant fluctuations in the source amplitudes as well, but we do not attempt a detailed analysis of these amplitude fluctuations in this work.

Cohen and Röttgering (2009) have observed ionospheric TEC gradients with the VLA using the differential refraction statistic and Helmboldt and Intema (2012) have measured the 2d power spectrum of spatial and temporal fluctuations in TEC over the VLA. A similar power spectrum analysis exploiting the MWA's much larger field of

view was recently carried out by Loi et al. (2015b,a). Because the ionosphere is curved and the derivation of equation (5.11) relies on the Fresnel approximation this model is only strictly valid for small fields of view so we only measure source shifts within 15 degrees of the phase center. To obtain a global picture of ionospheric conditions, we turn to the differential refraction statistic employed in Cohen and Röttgering (2009) which we now briefly describe.

For an ensemble of source pairs with an angular separation of θ , the one dimensional differential refraction statistic, $D(\theta)$, is defined as

$$D(\theta) = \langle |\Delta\theta_1 - \Delta\theta_2|^2 \rangle, \quad (5.13)$$

where $\Delta\theta_{1/2}$ is the measured offset of source (1/2) from its known catalog position.

If $\Delta\theta = \Delta\theta_1 - \Delta\theta_2 = (\Delta\alpha, \Delta\delta)$ is a two-dimensional vector with each component distributed with standard deviation σ , then the probability density function of the amplitude square is given by an exponential distribution. We compute an estimate of $D(\theta)$, $\hat{D}(\theta)$, within each angular separation bin by fitting a histogram of the lower 80% of source separations to an exponential distribution in order to eliminate potentially spurious outliers. $D(\theta)$, as will be explained, is directly related to the power spectrum of phase fluctuations whose properties we will determine below.

Each $\Delta\theta$ is the sum of an ionospheric offset and a noise term arising from random errors in determining the position of the source,

$$\Delta\theta_{1/2} = \Delta\mathbf{I}_{1/2} + \Delta\mathbf{N}_{1/2}, \quad (5.14)$$

where $\Delta\mathbf{I}$ is the contribution to position offset due to TEC gradients and $\Delta\mathbf{N}$ is the contribution from position errors. Expanding equation (5.13),

$$\begin{aligned} D(\theta) &= \langle |\Delta\mathbf{I}_1|^2 \rangle + \langle |\Delta\mathbf{I}_2|^2 \rangle - 2 \langle \Delta\mathbf{I}_1 \cdot \Delta\mathbf{I}_2 \rangle \\ &+ \langle |\Delta\mathbf{N}_1|^2 \rangle + \langle |\Delta\mathbf{N}_2|^2 \rangle - 2 \langle \Delta\mathbf{N}_1 \cdot \Delta\mathbf{N}_2 \rangle. \end{aligned} \quad (5.15)$$

For separations greater than several times the width of the synthesized beam, the

background noise is uncorrelated. In our analysis we only consider separations that are greater than $\approx 8.4'$ while our synthesized beam has a diameter of $4.2'$. In this regime, we can ignore the cross term in equation (5.15). Furthermore, it is roughly stationary over the center of the primary beam lobe so that the noise terms add a θ independent offset to equation (5.15). If θ is small enough that both sources fall behind a single isoplanatic ionosphere patch but large enough such that the synthesized beams do not significantly overlap, $\langle |\Delta \mathbf{I}_1|^2 \rangle = \langle |\Delta \mathbf{I}_2|^2 \rangle = \langle \Delta \mathbf{I}_1 \cdot \Delta \mathbf{I}_2 \rangle$ so the ionospheric terms in equation (5.15) cancel out and we are left with only the noise bias terms. We may therefore determine the noise bias from smallest non-zero separation bin and subtract it. Our estimate of the square root of the structure function of the ionospheric fluctuations is

$$\widehat{R}(\theta) = \sqrt{\widehat{D}(\theta) - \widehat{D}(\epsilon)}, \quad (5.16)$$

where ϵ is the median angular separation of our smallest bin which is $30'$.

In Fig. 5-7 we show the differential refraction computed from all differential source pair separations over 30 minutes on September 5th and 6th, 2013 for both of our observing bands. We see that the level of fluctuations recorded in both bands scales as λ^2 indicating that it indeed originates from ionospheric effects. Kassim et al. (1993) use a similar comparison to confirm ionospheric refraction as a source of variation in visibility phases. On September 6th, the levels of refraction are approximately a factor of two greater than those observed on September 5th, peaking at the end of the night.

We can relate our differential refraction measurements to the underlying power spectrum of phase fluctuations. In appendix 5.B), we derive the relationship

$$D(\theta) = \frac{2}{2\pi} \left(\frac{c}{2\pi f} \right)^2 \int dk k^3 (1 - J_0(kh\theta)) P(k) \quad (5.17)$$

where h is the height of the plasma screen, \mathbf{k} is the 2d wave vector, and J_0 is the Bessel function of the first kind. We parameterize $P(k)$ as a generalization of the form given in V15a which describes fluctuations that level out at some outer scale

$$r_0 = 2\pi/k_0.$$

$$P(k) = \phi_0^2 \frac{4\pi(n-1)}{k_0^2} \left[\left(\frac{k}{k_0} \right)^2 + 1 \right]^{-n}. \quad (5.18)$$

Substituting this form of the power spectrum into equation 5.17, we obtain

$$D(\theta) = 4(n-1) \left(\frac{c}{2\pi f} \right)^2 \phi_0^2 k_0^2 F_n(k_0 h \theta), \quad (5.19)$$

where $F_n(x)$ is a dimensionless integral.

$$F_n(x) = \int dq q^3 (1 - J_0(qx)) [q^2 + 1]^{-n}. \quad (5.20)$$

In Fig. 5-8 we show several examples of $F_n(x)$, noting that it exhibits power law behavior for small values of x and levels off towards $x = 1$. Assuming a plasma height of $h = 600$ km and fitting our structure functions to equation 5.19, we obtain values for the power spectrum normalization, ϕ_0 , the spectral index, n , and the outer energy injection scale, r_0 . We show our fits to our 83 MHz band 1 as dashed black lines in Fig. 5-7. Fitted values for r_0 were on the order of several hundred kilometers, n ranged between 2.3 and 2.7, and ϕ_0 between 4 rad on the quietest times and 45 rad during the most severe refraction at the end of Sept 6th. The reported n values are somewhat steeper than $n = 11/6 \approx 1.8$ for Kolmogorov turbulence. However, the spectral index of ionospheric fluctuations has been found to vary significantly (Rufenach, 1972; Cohen and Röttgering, 2009). We check the slopes of $F_n(x)$ for small x for our fitted indices and find that they lie within the slopes of the power law fits derived from various time ranges by Cohen and Röttgering (2009). We note that there are small systematic departures from the smooth behavior described by $F_n(x)$ in most of our intervals. These are possibly due to anisotropies in the phase fluctuation fields and departure from turbulent behavior over the short 30-minute time-scales due to transient phenomena such as traveling ionospheric disturbances.

Given our fitted values, we can also compute the diffractive scale of the ionospheric fluctuations r_{diff} which gives the scale at which the structure function of the phase field, ϕ reaches unity. Phase fluctuations that have large amplitudes and decorrelate

rapidly with separation give a smaller diffractive scale, allowing it to serve as a single number indicator of the severity of the fluctuations. We determine r_{diff} for each 30-minute interval by computing the structure function from each fitted power spectrum and numerically solving for r_{diff} . We obtain error bars by computing r_{diff} for 1000 instances of (r_0, n, ϕ_0) drawn from a multivariate gaussian distribution whose covariance is the estimate of the fitted parameter covariances given by the Levenberg-Marquardt method as implemented in `scipy`¹. We use the 16 and 84% percentile values of the resulting distributions to obtain 1σ upper and lower bounds. On September 5th, the median r_{diff} , at 83 MHz, ranged between $\approx 11 - 13$ km while on September 6th, it ranged between $\approx 4 - 6$ km. We report these values and their errors in Fig. 5-7. We also show the values of r_{diff} derived from all source pairs on each night in table 5.1. The frequency scaling of the diffractive scale depends on the spectral slope of $P(k)$, n . The n -values in our analyses typically fall between 2.3 and 2.7 which yields a $r_{\text{diff}} \propto f$ scaling². Hence, our measurements imply diffractive scales of ≈ 20 km and ≈ 10 km at 150 MHz on September 5th and 6th respectively. These values are within the range of typical scales (5 – 50 km) described in V15a.

Changes in the refractive index over a source on time scales shorter than the snapshot integration can cause source smearing in image space, resulting in a reduction in the observed peak brightness and reducing the number of source detections (Kassim et al., 2007). In Fig. 5-9 we compare the number of sources identified in each snapshot over September 5th and 6th. On September 5th, when the level of refraction is significantly lower, we observe that the number of sources increases as the field approaches transit, corresponding to the pointing at which the beam has maximal gain and then turns over. On September 6th, the night in which significantly greater refraction was observed, the number of sources identified stays relatively constant and significantly lower than any of the source counts observed on September 5th. There are also noticeable drops in the source counts on a handful of observations on September 5th which we found to correspond to flagging events in which an entire coarse band was

¹<https://github.com/scipy/scipy>

² Kolmogorov turbulence, with a spectral index of $n = 11/6$ in the V15a parameterization scales in a very similar way; $r_{\text{diff}} \propto 6/5 \approx 1$




Figure 5-7: Differential refraction derived from source pairs within 30 minute bins on September 5th, 2013 (top row) and September 6th, 2013 (bottom row). Band 1 (black points) scaled by the ratio of the band center frequencies square (solid black line) very nicely predicts the differential refraction in Band 2 (red points), indicating that the refraction we are measuring here is indeed due to ionospheric fluctuations. The magnitude of ionospheric activity differs significantly between September 5th and 6th and peaks over the last observations taken on the 6th. We also show fits to an isotropic power spectrum model of differential refraction at 83 MHz (dashed black line) and print the inferred diffractive scale.




Figure 5-8: The dimensionless integral, $F_n(x)$ normalized to unity at $F_n(10)$, given in equation 5.20. For small values of x , $F_n(x)$ is well approximated by a power law, but flattens out towards $x = 1$. Hence, the structure function of observed source offsets levels out at the outer energy injection scale of the turbulence.




Figure 5-9: The number of sources identified in 112s multifrequency synthesis images of Band 1 as a function of time over both nights of observing. On September 5th, the source counts increase with primary beam gain, until transit (vertical gray line) before decreasing. On September 6th, when more severe ionospheric refraction was observed, the source counts remain significantly lower. Fewer sources were identified in a handful of September 5th snapshots corresponding to observations in which calibration and flagging anomalies were observed (Figs. 5-3 and 5-6). We exclude these snapshots from our analysis.

eliminated with significant RFI detections on the edges (Fig. 5-6). We drop these snapshots from the rest of our analysis.

While refraction varied significantly over both nights and within each night, the bulk solar weather and geomagnetic conditions are nearly identical. In Table 5.1 we also list several bulk statistics such as K_p index, 10.7 cm flux (F10.7)³, and mean TEC at the MRO⁴. The solar flux at 10.7 cm is an often used index of solar activity and is known to correlate with the emission of the UV photons responsible for generating the free electrons that impact ionospheric radio propagation (Yeh and Flaherty, 1966; Titheridge, 1973; da Rosa et al., 1973) with values ranging between 50 and 300 SFU (1 SFU= 10^{-22} W m⁻²). The K_p index (Bartels et al., 1939), quantifies the severity of geomagnetic activity and combines many local measurements of the maximal horizontal displacement of the Earth’s magnetic field. Values for K_p range from 0 to 9 with any values above 5 indicating a geomagnetic storm. We find that the bulk values are very similar between nights, indicating that ionospheric indicators derived from

³ K_p and 10.7 cm flux values were obtained from <http://spaceweather.com/archive.php>.

⁴Values obtained by averaging global TEC maps downloaded from the MIT Haystack Madrigal database (Rideout and Coster, 2006).

Night	K_p	$\langle \text{TEC} \rangle$ (TECU)	F10.7 (sfu)	$\langle r_{\text{diff}} \rangle$ (km)
Sep 5th	2	10.6	10.9	$10.4 \pm_{0.2}^{0.4}$
Sep 6th	2	11.4	11.0	$5.20 \pm_{.05}^{.04}$

Table 5.1: Bulk ionospheric and solar weather properties on the two nights of observing presented here. $\langle \text{TEC} \rangle$ indicates the mean total electron content over the entire night. We also show the diffractive scale calculated from all source separations on both nights which differ by a factor of two.

the observations themselves are a much better way of assessing data quality. Indeed, Helmboldt et al. (2012) also found that the levels of ionospheric turbulence and the incidence of traveling ionospheric disturbances did not appear to correlate strongly with bulk ionosphere and solar weather statistics. A systematic study of ionospheric conditions at the MRO site and correlations with bulk ionosphere statistics at higher frequencies is currently underway (Loi et al., in preparation). Studies incorporating information from ancillary probes such as GPS stations are also being carried out (Arora et al., 2015).

Having established that one of our observing nights had twice the level of refraction than the other, we can compare power spectra from each night to gauge the impact of ionospheric fluctuations on the 21 cm power spectrum. In § 5.4.3, we find that the two power spectra appear indistinguishable.

5.3.3 Instrumental Spectral Structure

At low frequencies, the combination of wider primary beams and intrinsically brighter foregrounds causes leakage into the EoR window to corrupt a wider range of delays which are related to cosmological Fourier modes. In addition, the cosmological modes occupied by a fixed delay depends on redshift, causing features that contaminate low-sensitivity regions of k -space at one redshift to contaminate scientifically important regions at another.

When we form a power spectrum from visibilities calibrated by the techniques described in § 5.2.2, we are immediately confronted with detections of striped artifacts in the EoR window at discrete delays (see the top left panel of Fig. 5-20). These

delays correspond to the round-trip travel times on the lengths of cable that connect the MWA's receivers and beamformers. We can get a rough understanding of how miscalibrated cable reflections affect the power spectrum by considering the effect of reflections with delay, τ_j and complex amplitude \tilde{r}_j on the j^{th} tile. To first order in $|\tilde{r}_{i/j}|$, the effect of uncalibrated cable reflection is to multiply a visibility by a reflection factor.

$$V_{ij} \rightarrow V_{ij} (1 + \tilde{r}_i e^{2\pi i \tau_i f} + \tilde{r}_j^* e^{-2\pi i \tau_j f} + \mathcal{O}(\tilde{r}^2)) \quad (5.21)$$

The power spectrum is formed, roughly, by taking a Fourier Transform of the gridded visibilities in frequency and squaring. The square of the Fourier transform of a visibility with uncorrected reflections becomes,

$$\begin{aligned} |\tilde{V}_{ij}(\tau)|^2 &\rightarrow |\tilde{V}_{ij}(\tau)|^2 + \\ &2\text{Re} \left[\tilde{r}_i \tilde{V}_{ij}(\tau) \tilde{V}_{ij}(\tau - \tau_i) \right] + \\ &2\text{Re} \left[\tilde{r}_j^* \tilde{V}_{ij}(\tau + \tau_j) \tilde{V}_{ij}(\tau) \right] + \\ &2\text{Re} \left[\tilde{r}_i \tilde{r}_j^* \tilde{V}_{ij}(\tau - \tau_i) \tilde{V}_{ij}^*(\tau + \tau_j) \right] + \\ &2\text{Re} \left[\tilde{r}_i \tilde{r}_j^* \tilde{V}_{ij}(\tau - \tau_i + \tau_j) \tilde{V}_{ij}^*(\tau) \right] + \\ &2\text{Re} \left[\tilde{r}_i^2 \tilde{V}_{ij}(\tau - 2\tau_i) \tilde{V}_{ij}^*(\tau) \right] + \\ &2\text{Re} \left[\tilde{r}_j^{*2} \tilde{V}_{ij}(\tau + 2\tau_j) \tilde{V}_{ij}^*(\tau) \right] + \\ &|\tilde{r}_i|^2 |\tilde{V}_{ij}(\tau - \tau_i)|^2 + |\tilde{r}_j|^2 |\tilde{V}_{ij}(\tau + \tau_j)|^2 + \mathcal{O}(\tilde{r}^3). \end{aligned} \quad (5.22)$$

All terms in equation 5.22 involve the cross multiplications of $\tilde{V}_{ij}(\tau + \Delta\tau) \tilde{V}_{ij}^*(\tau + \Delta\tau')$ and a coefficient on the order of \tilde{r}^n . The cable delays on the MWA are all $\gtrsim 90$ m, corresponding to round trip delays significantly beyond the wedge for the short baseline lengths considered in our power spectrum analysis. Hence, if $|\tilde{r}|$ is greater than the ratio between the signal and foreground amplitudes, R_{fg} , terms with $\Delta\tau = \Delta\tau'$ dominate equation 5.22. As a consequence, the first and last lines in equation 5.22 dominate all others. The first ($\mathcal{O}(0)$) is the foregrounds in the absence of reflections and exceeds the signal by a factor of 10^{10} but is also contained within the

wedge. The last line contaminates $\tau = \tau_i, -\tau_j$ at the level of $\tilde{r}^2 \times 10^{10}$ the signal level. All other lines in equation 5.22 exceed the signal level by $\tilde{r}^2 10^5$ but also contaminate a greater range of delays: $\tau = \tau_i, -\tau_j, 2\tau_i, -2\tau_j, \tau_i - \tau_j$.

As we will discuss below, we find that $\tilde{r} \lesssim 10^{-2}$ (Fig. 5-9), hence our analysis is only sensitive to the first and last lines arising from first order reflections. The fourth, fifth and sixth terms in equation 5.22, which couple second order reflections and beats will be above the level of the cosmological signal twice the round trip delay times and their differences. However, if the $\lesssim 1\%$ reflections can be corrected to be below the level of $\lesssim 10^{-3}$, these cross terms will only appear at the 10^{-3} level of the signal and not impede a detection (though they may introduce some bias).

Since the $\mathcal{O}(\tilde{r}^2)$ terms in the last line of equation 5.22 dominate the others outside of the wedge, the lowest order effect of a reflection is to multiply our foregrounds by the reflection coefficient squared and translate them outside of the low-delay region in which they are usually confined (the wedge) to the round-trip delay of our cable which for the MWA is outside of the EoR window. Unless $|\tilde{r}|$ can be brought below the ratio between the foregrounds and the signal itself ($|\tilde{r}_i| \lesssim 10^{-5}$), the modes within a wedge translated to τ_i will remain unusable. Uncorrected reflections at the 10^{-2} level will also contaminate higher order harmonics and the differences between the delays.

For \tilde{r} significantly larger than R_{fg} (as is the case here), delays corresponding to higher order reflections will also be contaminated. These higher order terms are well below the sensitivity of this analysis but they can still potentially pose an obstacle to the detection of the signal. Hence, it is worth commenting on the level that \tilde{r}^3 and \tilde{r}^4 terms contaminate our data. First we address \tilde{r}^3 . Since every contribution of order $\mathcal{O}(\tilde{r}^3)$ in equation 5.22 involves the product of an $\mathcal{O}(\tilde{r}^3)$ coefficient with delay transformed visibilities translated to two different delays ($\Delta\tau \neq \Delta\tau'$), these terms will contribute at the level of $\tilde{r}^3 \times 10^5$ the level of the foregrounds. Uncorrected third order terms with $\tilde{r} \lesssim 10^{-2}$ will contaminate our data at 10^{-1} the level of the signal and we do not consider them a serious issue, especially if the reflections are corrected to the $\approx 10^{-3}$ level.

Writing down all \tilde{r}^4 order terms in equation 5.22 is straightforward but not terribly enlightening. We can obtain the leading contributions to the $\mathcal{O}(\tilde{r}^4)$ terms in equation 5.22 by ignoring the products of delay visibilities translated by different amounts, $\Delta\tau \neq \Delta\tau'$. The $\mathcal{O}(\tilde{r}^4)$ contributions with $\Delta\tau = \Delta\tau'$ are

$$\begin{aligned} & \dots + |\tilde{r}_i|^4 |\tilde{V}_{ij}(\tau - 2\tau_i)|^2 + |\tilde{r}_j|^4 |\tilde{V}_{ij}(\tau + 2\tau_j)|^2 \\ & + |\tilde{r}_i|^2 |\tilde{r}_j|^2 |V_{ij}(\tau - \tau_i + \tau_j)|^2 + \dots \end{aligned} \quad (5.23)$$

We conclude that the sub-percent reflections observed in our data will introduce $\mathcal{O}(\tilde{r}^4)$ terms to the power spectrum at $\tilde{r}^4 \times 10^{10}$ times the amplitude of the signal at twice the fundamental delays and their differences. Second order reflections at the $\lesssim 1\%$ level will therefore dominate the signal by two orders of magnitude, impeding a detection. Fortunately, the \tilde{r}^4 dependence of these second order terms greatly amplifies even modest improvements in correcting the reflections. For example, if the reflections are brought to below the 0.1% level, the $\mathcal{O}(\tilde{r}^4)$ terms will be brought to below 10^{-2} the level of the 21 cm signal. We are able to bring the reflections down to ≈ 0.002 so they are not a problem in our data.

What cosmological wave-vectors in our measurements are contaminated by reflections? In cosmological coordinates, a fixed delay corresponds to a line-of-sight wave-number of k_{\parallel} (in units of $h \text{ Mpc}^{-1}$) of

$$k_{\parallel} \approx \frac{2\pi H_0 f_{21} E(z)}{ch(1+z)^2} \tau, \quad (5.24)$$

(Morales and Hewitt, 2004) where c is the speed of light, $H_0/h = 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, f_{21} is the hyperfine emission frequency, and $E(z) = \sqrt{\Omega_M(1+z)^3 + \Omega_k(1+z)^2 + \Omega_\Lambda}$. For constant τ , the k_{\parallel} center of the translated delay wedge will decrease with increasing redshift.

In Table 5.2 we list the lengths of cable between the MWA receivers and beamformers along with their round-trip delays, and their corresponding k_{\parallel} at the center redshift of our two bands along with a redshift typical of an EoR measurement. Be-

cause $P(k)$ decreases rapidly with increasing k , interferometers are expected to have the highest signal to noise at the smallest delay that is uncontaminated by foregrounds (280 ns, corresponding to $k_{\parallel} \approx 0.1 - 0.2 h \text{ Mpc}^{-1}$ at EoX to EoR frequencies (Pober et al., 2013a)). Assuming that reflections can be corrected to be below the 10^{-3} level so that only the last line in equation 5.22 is above the signal, they should be benign as long as they remain at sufficiently large or small k_{\parallel} that they don't leak into the region of maximal sensitivity. Because standing waves translate the entire wedge up to their delay, reflections located at the edge of the wedge will result in excess supra-horizon emission while cable reflections outside of the wedge will contaminate a finite chunk of k_{\parallel} , not just the delay of the reflection itself.

We see in Table 5.2 that the minimum k_{\parallel} associated with a cable ripple on the MWA at $z = 7$ is $0.42 h \text{ Mpc}^{-1}$ which, even if we allow for a delay width of 280 ns (the approximate width of the wedge with a supra-horizon emission buffer at $k_{\perp} = 0$), is above the region of maximal sensitivity. However, at $z = 16$, the approximate center redshift of our Band 1, this cable ripple lies at $0.27 h \text{ Mpc}^{-1}$, which can leak a significant amount of power into the sensitivity sweet spot due to its finite k_{\parallel} width. This effect is illustrated in Fig. 5-10. At $z = 7$, the smallest reflection delay on the MWA introduces foreground contamination down to $k_{\parallel} \approx 0.3 h \text{ Mpc}^{-1}$, leaving a small foreground-free window in which cosmological measurements may be performed. On the other hand, at $z = 16$, this window becomes smaller with the other reflections filling the EoR window up to $k_{\parallel} \sim 0.8 h \text{ Mpc}^{-1}$. If the corrected reflections are several times larger than 10^{-3} , higher order terms below the 90 m delay may be comparable to the level of the 21cm signal (shown as light grey regions and centered on light-gray dashed lines). However, only the peak of this foreground power will be near the signal level and the broad wings caused by beam chromaticity and will be well below the signal.

This effect is purely geometrical in that while the mapping from instrumental delays to k-space varies, the number of measurements in (u, v, η) cells which are uncontaminated stays constant. However, as we go to higher redshift the increasing width of the primary beam increases foreground power at supra-horizon delays, effectively




Figure 5-10: The regions of the EoR window contaminated by foregrounds due to uncalibrated cable reflections for several different redshifts. Dark gray regions denote contamination from first order cable reflections assuming a wedge out to the first null of the primary beam plus the $0.15 h \text{ Mpc}^{-1}$ at $z = 8.5$ buffer observed in Pober et al. (2013a). Since the buffer is associated with the intrinsic spectral structure of foregrounds, it lives in delay space. Dark gray regions denote foreground contamination within the wedge which exists even without instrumental spectral structure. At $z = 7$, a representative EoR redshift, the contaminated regions remain at relatively high k_{\parallel} and have smaller widths due to the smaller primary beam and the scaling of k_{\perp} and k_{\parallel} with z . While regions exist between the first order reflections that are somewhat wider at lower redshift, second order reflections can still potentially contaminate nearly all of the EoR window in which interferometers are supposed to be sensitive (light grey shaded regions). Second order reflections are below the sensitivity level of this study but will also pose an obstacle to longer integration unless calibrated out.

L (m)	N	τ_r (μ s)	$f_{30\lambda}$ %	$k_{\parallel}(z = 7)$ ($h \text{ Mpc}^{-1}$)	$k_{\parallel}(z = 12)$ ($h \text{ Mpc}^{-1}$)	$k_{\parallel}(z = 16)$ ($h \text{ Mpc}^{-1}$)
90	19	.74	40	.42	.31	.27
150	31	1.2	80	.70	.53	.50
230	23	1.9	29	1.1	.83	.70
320	8	2.6	0.4	1.5	1.2	1.0
400	17	3.3	1.7	1.9	1.4	1.3
524	30	4.3	1.9	2.5	1.9	1.7

Table 5.2: There are N cables of each length (L) between the MWA receivers and beamformers with associated round-trip delay times (τ_r). In the $f_{30\lambda}$ column, we list the percentage of baselines within 30λ at 83 MHz (where the majority of the MWA’s power spectrum sensitivity lies) that are formed from at least one tile with the given cable length. We also list the k_{\parallel} of each delay given by equation (5.24) for three different redshifts. Cable reflections that are significantly above the k_{\parallel} values where we expect to obtain maximum sensitivity to the power spectrum at EoR redshifts ($z \approx 7$) move into the maximum sensitivity region at EoX redshifts ($z \approx 16$). Higher order reflections will also contaminate multiples of and differences between the delays and k_{\parallel} values listed in this table (thought at a lower level).

reducing the number of usable (u, v, η) cells. As mentioned above, the cosmological power spectrum decreases significantly with increasing k so the fact that smaller k are contaminated by cable reflections at higher redshifts hurts our sensitivity disproportionately. We also note that any partial reflections from kinks and bends within the cable itself can lead to contamination of additional delays below the round-trip travel time on the cable.

We make a first attempt to remove these reflections by fitting reflection functions (equation 5.3) to our self calibration solutions divided by the smooth fit in equation 5.2. Since our per snapshot calibration solutions are too noisy to obtain good fits for the beamformer-receiver reflection, we fit these on calibration solutions that are averaged over each night of observing. We find that this method is of limited efficacy in removing the receiver to beamformer ripples (Fig. 5-20, top right).

Since, to avoid cosmological signal removal and spurious frequency structure due to mis-modeled foregrounds, we are attempting to model our bandpass with a small number of parameters we are unable to capture the full spectral structure of the cable reflections. As we see in Fig. 5-13 the reflection parameters are frequency dependent

and we do not have a clear picture of their precise evolution. Calibration exploiting redundant baselines might be able to make headway on the problem since it is not sensitive to unmodeled signal except in for deriving an overall phase and amplitude scaling that averages over all baselines (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014). The MWA’s baselines are designed for very low redundancy, making the technique unusable here. Future upgrades to the MWA are expected to include a significant number of redundant baselines (Tingay, private communication).

5.3.4 Calibration with Autocorrelations

Confronted with the problem of reflections contaminating the power spectrum, we apply an alternative approach that uses tile autocorrelations to obtain calibration amplitude gains with sub-percent level accuracy.

What information is encoded in the autocorrelations? Let $I(\mathbf{s}, f)$ be the brightness distribution on the sky at frequency f and direction \mathbf{s} . Consider the sky signal entering the antenna and traveling through a signal chain in which the m^{th} successive element applies a multiplicative complex gain $g_{j,m}$ and adds a zero mean noise component with variance $N_{j,m}^2$. Correlating the output at the M^{th} gain element with itself to form an autocorrelation yields

$$V_{jj}(f) = |g_j|^2 \left[\sum_{m=0}^M \frac{N_{j,m}^2(f)}{\prod_{n=0}^m |g_{j,n}|^2} + \int d\Omega A_j(\mathbf{s}, f) I(\mathbf{s}, f) \right]. \quad (5.25)$$

Here, $g_j = \prod_{m=0}^M g_{j,m}$ is the net gain and $A_j(\mathbf{s}, f)$ is the antenna beam. Using simulations of diffuse emission and an analytic model of the MWA primary beam, we find that $\int A_j I(\mathbf{s}, f) d\Omega$ is fit at the 10^{-5} - 10^{-6} level by a third order polynomial while $N_{j,m}^2(f)$, which is due to noise in analogue electronics should also vary smoothly with frequency. Hence $\sqrt{V_{jj}}$ is well approximated by the product of $|g_j|$ multiplied by a smooth factor which may be modeled by a low-order polynomial.

To remove this multiplicative factor, we need a model for the ratio of the square root of each autocorrelation to the amplitude of the calibration solution. We use the product of a third order polynomial and a 7 m cable reflection (modeling the LNA-




Figure 5-11: We show the amplitude of a calibration gain averaged over a fifteen minute pointing (black circles) along with the square root of our autocorrelations which have been scaled by a third order polynomial and a single seven meter reflection to match the calibration solution (red line). After multiplying the autocorrelations by a smooth function, they are brought into good agreement with the calibration gains.




Figure 5-12: Left: In order to obtain reflection parameters, we divide our scaled autocorrelation (magenta circles) by a smooth function consisting of a third order polynomial and large scale reflections (green line). Right: We fit this ratio (magenta circles) to a reflection function (green line) and are left with $\sim 10\%$ residuals (grey points).

beamformer cables) and fit it to the ratio, averaged over 112s intervals, using our noisy initial calibration solutions.

We then use the square root of the autocorrelation divided by this polynomial as our calibration amplitude. In Fig. 5-11 we demonstrate the validity of this technique by comparing the smoothly corrected autocorrelations for a single snapshot with a calibration amplitude that has been averaged over a single pointing and see that they are in very good agreement.

While the autocorrelations can be used for our amplitude calibration, there still remains the problem of adding the reflection ripple to phase calibration. We can

use autocorrelations to predict this ripple in the phases. We obtain its parameters following the same fitting procedure described in §5.2.2 except this time we fit the scaled autocorrelations, with a smooth polynomial divided out, to the amplitude of the reflection term in equation 5.3,

$$|R_j(f)| = \frac{1}{\sqrt{1 - 2r_j \cos(2\pi f\tau_{L(j)} + \phi_j) + r_j^2}}. \quad (5.26)$$

In Fig. 5-12 we illustrate the fitting procedure by showing the autocorrelations divided by the smooth fit along with the best fit model reflection. One can see that the residuals in the reflection fit tend to be on the order of 10% hence there is some fine scale structure at the 10^{-3} level that we are still unable to model. Since our model includes the impact of higher order reflections, we think that these residuals arise from unmodeled frequency dependence in the reflection coefficients, sub-reflections in the cables, and digital artifacts present in the autocorrelations. Recalling our discussion in § 5.3.3, 10^{-3} -level residuals will leave contamination in our power spectra at the level of 10^4 times the signal level due the $|\tilde{r}|^2$ terms in the last line of equation 5.22 but suppress all higher order reflections to below the signal level.

From $\tau_{L(j)}$, ϕ_j , and r_j , we add the reflection's additive contribution to the gain phase

$$\begin{aligned} \text{Arg}(g_j) &\rightarrow \text{Arg}(g'_j) \\ &= \text{Arg}(g_j) + \tan^{-1} \left[\frac{-r_j \sin(2\pi f\tau_{L(j)} + \phi_j)}{1 - r_j \cos(2\pi f\tau_{L(j)} + \phi_j)} \right]. \end{aligned} \quad (5.27)$$

Since tiles with 320, 400, and 524 m cables only contribute to $\sim 4\%$ of our sensitive baselines, we discard them entirely. In Fig. 5-13 we show the distributions of the reflection amplitudes fitted from autocorrelations (averaged over the night of September 5th) inferred for our 90, 150, and 230 m cables for both the high and low bands. One can see that the reflection coefficients are on the order of fractions of a percent and vary significantly from cable length to cable length. This is reasonable since the cable impedance, which determines the reflection amplitude, is a function




Figure 5-13: Histograms of fitted cable reflection amplitudes for Band 1 (blue) and Band 2 (green) obtained from fits to autocorrelations for three different cable lengths between MWA receivers and tiles. The reflection amplitudes range from 0.2 – 1% making them difficult to fit using the noisy self calibration solutions. Reflection amplitudes in Band 2 are systematically larger than Band 1 for all cable lengths, indicative of non-trivial frequency evolution in the reflection parameters.

of both its geometry and dielectric properties (with equal length cables likely formed from cable batches of similar dielectric properties). In addition, frequency evolution of the reflection amplitude is apparent by comparing the fit in Bands 1 and 2 implying that a single delay standing wave is not quite the correct model to use in our phases.

Autocorrelations are particularly susceptible to RFI and potential contamination due to cross talk and other artifacts. In Fig. 5-11, we saw that after flagging the channel edges, the spectral structure in the autocorrelations was consistent with our calibration solutions up to a smooth polynomial factor. In Fig. 5-14, we inspect for artifacts and RFI in a typical tile autocorrelation as a function of time. We see




Figure 5-14: Dynamic spectra of the square root of a representative tile autocorrelation. Note the different color bars for the two frequency bands since Band 1 evolves more steeply in frequency than Band 2. The autocorrelations exhibit repetitive structure in time from night to night with smooth time variations occurring as the sky rotates overhead and step transitions occurring every ≈ 30 minutes due to changes in the analogue beamformer settings as the antennas track the sky. Limited RFI is plainly visible within the FM band, especially in Band 2, and the events are consistent with the flagging events identified by `cotter` shown in Fig. 5-6.

that RFI is present at similar times that were flagged in autocal (Fig. 5-6). We also see that the time evolution of each autocorrelation is consistent between the two nights with rapid 10% transitions occurring at ≈ 30 minute intervals when the analogue beamformer settings are changed to track the sky. Ripples in frequency are also visible, corresponding to the structure in the standing wave reflections. It is difficult to pick out small artifacts in this dynamic spectrum view unless more large-scale smooth structure is fitted out, as is done in our calibration procedure. The residuals after this fitting give a better picture of what fine spectral features exist in the dynamic spectra of the autocorrelations which we discuss in the next section.

5.3.5 The time dependence of residual structure.

We noted in § 5.3.4 that our fits to reflections tended to have $\sim 10\%$ residuals. Since we rely on these fits to predict the reflections in our gain phases, we expect residuals




Figure 5-15: Left: The residuals to fitting reflection functions in our autocorrelations for all two minute time-steps in our analysis for a representative tile with a 90 meter beamformer to receiver connection (light grey points). While some scatter exists in the residuals due to fitting noise, they average to non-zero values on the order of $\sim 10^{-3}$ (black dots). These residuals are due to mismodeling the reflections and at a lower level potentially arise from digital artifacts. Right: The same as the left but for a 150 meter cable whose reflection coefficient is $\sim \times 2$ as large as the 90 meter cable, leading to larger residuals due to mismodeling.

of these fits that are also present in the phases to contribute reflection power at a similar level. Our residuals could arise from thermal noise in the autocorrelations and calibration solutions. If this were the case we might expect them to average down with time. On the other hand, these residuals might also arise from mismodeling of the reflections themselves and would not average down with time. The result would be a systematic floor which can only be overcome by finding the correct model of the reflections or removing them from the signal path.

Plotting the fit residuals for two representative 90 m and 150 m tiles over the low band (Fig. 5-15), we find them to be at the $\sim 10^{-3}$ level. While there is some scatter in these residuals due to fitting noise, their frequency dependent shape is relatively constant. As a consequence, the residuals average to a spectrum with frequency structure. These residuals are likely due to mismodeling of the frequency dependent amplitude, phase, and period of the reflections but at a lower level may have some contributions from digital artifacts and cross-talk present in the autocorrelations. Because the component in these residuals that is sourced by reflections is also present in the phases which we are trying to model, there remains an uncorrected component to the gains that we are not calibrating out and does not average down with time.

While calibration with autocorrelations still appears to be limited by fine frequency artifacts arising from reflections, the high SNR of the reflections in the autocorrelations does offer significant improvement over fitting the reflections in the calibration solutions themselves. We provide a more quantitative look at the improvement achieved using calibration with auto-correlations in § 5.4.2.

5.4 Power Spectrum Results

We can now present our power spectrum results and the first upper limits on the Epoch of X-ray heating power spectrum. We form cross power spectra of the even and odd timestep data cubes through the empirical covariance modeling technique developed in (Dillon et al., 2015a) (D15). In this procedure, the foreground residual model used in the inverse-covariance weighted quadratic power spectrum estimates and in the associated error statistics is derived from the data. It assumes that foreground residuals are correlated in frequency but uncorrelated in the uv plane and depend only on frequency and $|\mathbf{u}|$. We refer the reader to D15 and its predecessors (Tegmark, 1997a; Liu and Tegmark, 2011; Dillon et al., 2013, 2014) for a thorough discussion of how this technique works. Along with estimates of the power spectrum amplitude, our pipeline outputs error bars and window functions which describe the mixing of the true power spectrum values into each estimate. We form 1d power spectra by binning our 2d power spectra using the optimal estimator formalism of Dillon et al. (2014) with the weights of all modes lying outside of the EoR window or with k_{\parallel} values showing consistent cable reflection contamination set to zero (D15).

First we will examine our two dimensional power spectra for Bands 1 and 2, derived from ≈ 15 MHz of bandwidth each, and comment on systematics (§5.4.1) and how well our calibration techniques mitigate them (§ 5.4.2). We finish by presenting our spherically binned 1d power spectra, our most sensitive data product. We use our 1d power spectra to compare foreground contamination from ionospheric systematics on both nights (§ 5.4.3) and determine our best upper limits (§ 5.4.4).

5.4.1 Systematics in the 2d Power Spectrum.

The absolute values of our two dimensional power spectrum estimates using data calibrated with auto-correlations are shown in Fig. 5-16 for both bands. The distinctive “wedge” confines the majority of our foreground power with some supra-horizon emission clearly present out to $\sim 0.1 h \text{ Mpc}^{-1}$ as was found in observations of foreground contamination with a similarly large PAPER primary beam (Pofer et al., 2013a). Smooth frequency calibration errors, arising from foreground mismodeling, may also contribute to the supra-horizon emission along with intrinsic chromaticity in the primary beam itself. As we expected, the level of foregrounds and thermal noise is noticeably higher in our measurement of Band 1. We note that at the edge of our k_{\perp} range, there is a significant increase in power which is due to a rapid increase in thermal noise from the drop-off in the uv -coverage of our instrument. Though somewhat hard to see by eye, there are signs of coherent non-noise-like structures in both bands below $k_{\parallel} \approx 0.5 h \text{ Mpc}^{-1}$.

We confirm these faint $k_{\parallel} \lesssim 0.5 h \text{ Mpc}^{-1}$ structures as systematic contamination by inspecting the sign of our power spectrum estimate over the k_{\perp} - k_{\parallel} plane. While the expected value of the even/odd cross power spectrum is always positive, k -bins that are dominated by noise have an equal probability of being positive or negative. Regions in which band powers are predominately positive are detections of foregrounds or systematics. In Fig. 5-17 we show $P(k)$ from data calibrated with auto-correlations for both bands with an inverse hyperbolic sine color scale to highlight regions of k -space that have positive or negative values. It is clear that the region of with $k_{\parallel} \lesssim 0.5 h \text{ Mpc}^{-1}$ is not well described by thermal noise.

Detections of foregrounds and systematics are especially visible in the ratio between the power spectrum and error bars predicted by our empirical covariance method (Fig. 5-18). In Fig. 5-19 we observe excess power at the $\sim 2\sigma$ level. While this is not a significant excess on a per cell basis, we detect this same power at high significance when we average in bins of constant $k \equiv \sqrt{k_{\perp}^2 + k_{\parallel}^2}$.

Because this excess power is present at similar levels over both of our observing




Figure 5-16: The absolute value of our cylindrical power spectrum estimate from our three nights of observing on Band 2 (left) and Band 1 (right). We overplot the locations of the primary beam (dash-dotted), horizon (dashed), and horizon plus a $0.1 h \text{ Mpc}^{-1}$ buffer (solid black) wedges. We see that the foregrounds are primarily contained within the wedge and that the EoR window is, for the most part, noise-like. There is some low SNR structure below $k_{\parallel} \approx 0.5 h \text{ Mpc}^{-1}$, corresponding to k_{\parallel} modes contaminated by cable reflections. The amplitude in power rises very quickly due to an increase in thermal noise which rises very quickly at large k_{\parallel} due to a rapid falloff in uv coverage beyond $k_{\perp} \sim 0.2 h \text{ Mpc}^{-1}$.




Figure 5-17: $P(k)$ over Band 2 (left) and Band 1 (right) with a color scale that highlights cells with positive or negative values. We expect regions that are thermal noise dominated to contain an equal number of positive and negative estimates and regions that are dominated by foreground leakage to be entirely positive. We observe significant foreground contamination outside of the wedge up to $k_{\parallel} \approx 0.5 h \text{ Mpc}^{-1}$ in both bands.




Figure 5-18: The errors on \hat{p} arising from residual foregrounds and thermal noise are determined by looking at even/odd difference cubes and foreground-subtracted residual cubes using the method of Dillon et al. (2015a). We show the error bars on our cylindrical power spectrum here, seeing that errors arising from foregrounds are contained within the wedge. These foreground errors are maximized at the smallest and largest k_{\perp} arising from large power in diffuse emission and increasing thermal noise from a dropoff in baseline density respectively.




Figure 5-19: The foreground contamination within the wedge along with residual detections due to miscalibrated fine frequency features in the bandpass are especially clear in plots of the ratio between power and the error bars estimated by the empirical covariance method of D15. We overplot the wedge with a $0.1 h \text{ Mpc}^{-1}$ buffer along with the wedge translated to cable reflection delays of our 90 and 150 m receiver to beamformer cables to highlight the effect of this systematic.

subbands, one of which has a significantly greater overlap with the FM, we cannot attribute this excess to RFI. In our 1d power spectra we also find that excess power is detected in our highest redshift bin which is outside of the FM entirely (Fig. 5-23, right panel). The best explanation we have for this leakage is the residual structure in the MWA’s bandpass caused by standing wave reflections on the beamformer to receiver cables. To demonstrate the plausibility of this explanation, we overlay the wedge translated to the k_{\parallel} modes corresponding to the delays of our 90 and 150 meter cables. For clarity, we do not show the 230 meter cable reflections in this overlay since their amplitudes and the number of tiles affected is comparatively small. We also observe this reflection in the 1d power spectrum (Fig. 5-23) which has higher signal to noise. We find that the region where one might expect contamination from a cable reflection is in good agreement with the observed excess power.

5.4.2 Comparing Calibration Techniques

Having formed 2d power spectra and estimates of the vertical error bars, we are in a position to assess the performance of our calibration solution in removing systematics. By inspecting the signal to error ratio in the EoR window, we compare our different calibration techniques. In Fig. 5-20 we show the ratio of $P(k)$, binned over annuli, to the error bars in Band 1 for the calibration techniques discussed in this work. For all calibration methods, the majority of foreground detections are contained within the wedge with a $\sim 0.1 h \text{Mpc}^{-1}$ buffer, indicating that all perform at a similar level in removing smooth gain structure within the wedge.

We first inspect a power spectrum derived from data calibrated using the initial method described in § 5.2.2 in which coarse band structure is removed by averaging over tiles, the per tile amplitudes and phases of each antennas are fit to smooth polynomials, and no attempt is made to model the beamformer-receiver reflections (top left corner). Significant foreground power is visible beyond the wedge to $k \sim 0.5 h \text{Mpc}^{-1}$ and is especially bright at the delays corresponding to the k_{\parallel} values of the cable reflections in Table 5.2. The fact that the 150 m delay dominates the others stems from the fact that most of our short baselines are formed from 150 m cables




Figure 5-20: Here we show the ratio between our 2D power spectrum and the error bars estimated by the empirical covariance method of Dillon et al. (2015a). On the top left, we show our data calibrated using our initial calibration (see § 5.2.2) with no attempt made to correct for standing wave structure in the MWA bandpass. Bright, band-like structures are clearly visible at the delays associated with reflections. On the top right, we show a first attempt to correct for cable reflections by fitting a sinusoidal model to rather noisy calibration solutions that had been integrated over a night of observing (1.5 hours each night). While the bands appear weaker, they are still quite visible above the noise. In the bottom right panel, we show the same plot with calibration solutions using scaled autocorrelations described in § 5.3.4. In the lower left panel we show a power spectrum with calibration solutions using autocorrelations for the amplitudes but without any attempt to correct reflections in the phase solutions. Pronounced reflection features are visible in this power spectrum, indicating that any mismodeled reflection structure in the phases will contaminate our measurement.

and that the amplitude of the reflections in the 150 m cables is larger compared to the 90 and 230 m cable lengths (Fig. 5-13). We next show a first attempt to fit out the reflections by averaging all calibration amplitudes in a night, dividing out a polynomial, and fitting equation 5.26. While the power in the bands is reduced significantly, residuals remain at the $2-10\sigma$ level, especially in the reflection bands. Since our initial calibration solutions are so noisy, it makes sense that they are difficult to fit.

We finally inspect results from calibrations derived from the autocorrelations described in §5.3.4 (lower right). While there is significant reduction compared to the amplitude on the top right corner, there still exist residuals outside of the window at the $\sim 1 - 2\sigma$ level. We think that these residuals arise from imperfect modeling of the reflection coefficients in the autocorrelation amplitudes, which will leave some reflection structure in the visibility phases. To demonstrate the impact of unmodeled reflection structure in the phases, we leave the phases of our auto-calibration solutions uncorrected for any fitted reflection coefficients (lower left) and find that significant power is reintroduced into the window.

We can get a more quantitative view of how much autocorrelations can improve calibration by taking a slice through the cylindrical power spectrum at the k_{\parallel} of our 150 m cable reflection (Fig. 5-21) where we see that fitting the calibration solutions was able to remove roughly an order of magnitude of the power in the reflection while AutoCal removes a factor of ~ 20 . Since the power spectrum is proportional to the square of the visibilities which are primarily contaminated by first order reflection contributions, this corresponds to an accuracy of $\approx 20\%$ in removing the reflections in the visibilities and is consistent with the residuals observed in Fig. 5-15. Such inaccuracy likely arises from our inability to model the precise frequency dependence of the reflection parameters in the phases and is on a similar order to the residuals observed in Fig. 5-12. Since the reflections are removed to this accuracy in the visibility, we can briefly comment on how the relative contribution of second order reflections (which are below our noise floor even without any calibration). Since the second order reflections appear in the data at the \tilde{r}^4 level and we have reduce their




Figure 5-21: The level of power at a fixed k_{\parallel} corresponding to the delay of reflections from our 150 m cable (left), and comparing it to a value of k_{\parallel} unaffected by cable reflections (right). The blue line shows the power spectrum level for calibration in which the bandpass is modeled as a polynomial with no attempt to correct fine frequency scale reflections. We see that power is on the order of ~ 50 times the thermal noise level (green-dashed line). Attempting to fit the reflections to calibration solutions integrated over each night gives us an improvement in the power level by roughly an order of magnitude (orange solid line). Using calibration solutions derived from autocorrelations brings down the reflection power by another factor of a few (purple solid line) but is still unable to bring the majority of measurements below the $\sim 1\sigma$ level. While we think that the autocorrelations accurately capture the fine frequency structure of the gains, we are still forced to model this fine frequency structure and predict it in the phases. Residual power is likely due to inaccuracies in this modeling. The right hand panel shows all data below the estimated noise level. This is due to the fact that in (Dillon et al., 2014) it is shown that the method for calculating error bars layed-out in Liu and Tegmark (2011); Dillon et al. (2013) slightly over-estimates the noise.

amplitude in the data from $\lesssim 0.01$ to $\lesssim 0.003$, they will enter the power spectrum at the level of $\lesssim 10 \times 10^{-1}$ the level of the 21 cm signal.

We attempted to better model the reflections by allowing for frequency evolution of the amplitudes but found little improvement in the power spectrum. We also found that we are able to obtain better fits of the autocorrelations by adding additional smooth reflections terms to equation 5.2 which could be important if unmodeled large scale structures bias our fits of small scale ones. However, using more complicated fits of the large scale structure, we did not observe significant improvement in power spectrum contamination. The solutions that we ultimately settled on in this analysis allow for a power law evolution of the reflection amplitude and add an additional

small delay reflection term to equation 5.2 which lies well within the wedge. In most of our autocorrelation fits, residuals remained at the $\sim 10\%$ level, some of which may arise from secondary reflections in bent or kinked cables. While these residuals were clearly present at high SNR in the autocorrelations, we have not found a way to sufficiently model the contribution of these low level structures to our phases.

While using the autocorrelations has allowed us to characterize and subtract the fine spectral structure in the instrument better, it may not be a viable long term solution, even in the regions of the EoR window that currently appear foreground free. RFI contamination and digital artifacts are known to contaminate autocorrelations and likely exist below our current noise level.

5.4.3 Power Spectra Comparison Between Nights of Varying Ionospheric Activity

An open question is whether or not the ionosphere will significantly hamper measurements of the power spectrum. The fact that the severity of ionospheric effects increase with λ^2 makes the question especially pertinent at low frequency. Changes in foreground emission induced by ionospheric effects can enter the power spectrum in two ways: through calibration and through the foreground residuals themselves. We check whether either of these potential error sources have an observable effect on our 1d power spectrum in Fig. 5-22 by comparing power spectra derived from 1.4 hours of Band 1 data on September 5th, over which ionospheric activity was comparatively mild to the same number of hours of Band 1 data on September 6th where differential refraction was approximately twice as severe.

We find that the power spectra, which are estimated from data outside of the wedge, are consistent with each other. This result confirms the intuitive idea that since ionospheric errors in the foreground model are spectrally smooth (evolving as $\sim \lambda^2$), they should be contained within the wedge. We also extended our 1d power spectrum estimation into the wedge to see whether the foreground detections appeared to be significantly different and find that they are not. This suggests that the random errors




Figure 5-22: The Band 1, 1d power spectra from our two nights of observing: September 5th, 2013 (black) and September 6th, 2013 (red). We saw in Fig. 5-7, that the magnitude of refractions on September 6th were on average twice as severe. The two power spectra nearly indistinguishable (within error bars) despite the significant differences in conditions, indicating that ionospheric systematics do not have a significant effect after three hours of integration, even at these low frequencies.

induced by the ionosphere average down with time. It is important to keep in mind that the spatial scales being probed in our analysis are relatively large, on the order of $\gtrsim 2.5^\circ$, while ionospheric refraction at these frequencies effects sub-arcminute scales. Hence the contamination that we might expect from ionospheric refraction should be small. Amplitude scintillation effects are prominent on short baselines (V15a) and likely dominate any contamination, however their spectral coherence still constrains them to be predominantly within the wedge (V15b).

5.4.4 First Upper Limits on the 21 cm Power Spectrum During the Pre-Reionization Epoch

We limit our 1d power spectra to redshift widths of $\Delta z \sim 1.5$ to minimize effects from cosmic evolution. A redshift interval of $\Delta z \sim 0.5$ is the range most cited in the literature over which the statistics of the brightness temperature field are expected to be stationary (Mao et al., 2008). However, at higher redshift, the frequency range corresponding to $\Delta z \approx 0.5$ decreases as $(1+z)^{-2}$ with $\Delta z = 0.5$ corresponding

to a bandwidth of only 2.45 MHz by $z = 16$. Reducing our bandwidth to such a small interval leads to poor k resolution which we prefer to maintain for assessing systematics. Since we are far from a detection, we opt for a larger redshift interval than we would otherwise use if we were actually observing the cosmological signal. In Fig. 5-23 we show 1d power spectra derived from our three hours of observing. Vertical error bars give 2σ uncertainties and the horizontal error bars give the width of our window functions. The amplitudes of our power spectrum values are consistent with thermal noise except for the regions of k-space below $k_{\parallel} \lesssim 0.5 h \text{Mpc}^{-1}$. At $k \approx 1 h \text{Mpc}^{-1}$, where our measurements are well described by thermal noise, our upper limits are on the order of 100 times higher than the results presented in D15 in which a similar three hour upper limit was established at ≈ 180 MHz. This factor of ≈ 100 is consistent with what we expect from equation 5.8. The sky temperature increases with decreasing frequency as $f^{-2.6}$ leading to a factor of ≈ 30 from T_{sys}^2 while λ^4/A_e^2 introduces an additional factor of 4-10.

The detections at small k are many orders of magnitude larger than the expected cosmological signal from a 21cmFAST simulation (Mesinger et al., 2011) (blue solid lines) so they cannot possibly originate from the redshifted HI emission. Instead, these detections are most likely the miscalibrated reflection structure observed in our 2D power spectra. We shade out regions of the k axis in which we expect contamination given the reflections discussed above and find that they correspond to the same modes where detections are observed. These systematic detections occupy the regions of Fourier space where our interferometer has the greatest sensitivity to the cosmic signal. Since we do not expect the systematics dominated regions to integrate down, a detection with the MWA in its current state using the techniques presented in this work would have to take place at $k \gtrsim 0.5 h \text{Mpc}^{-1}$, requiring over 10^5 hours of integration—a rather infeasible time scale. Thus, in order to probe the pre-reionization epoch, improvements in calibration and/or changes in the hardware of the MWA will have to be implemented. We note that at lower redshifts, the primary beam is smaller and the k-modes occupied by reflections are farther away from the sensitivity sweet spot, so it is less likely that this problem will prevent the MWA




Figure 5-23: Dimensionless 1d power spectra derived by Integrating spherical shells excluding the foreground contaminated wedge region with a $0.1 h \text{ Mpc}^{-1}$ buffer. Black dots indicate the mean estimated from the weighted average in each bin. Vertical error bars denote the 2σ uncertainties while horizontal error bars indicate the width of window functions. We also shade regions of k -space that we expect to have some level of foreground contamination due to uncalibrated cable reflection structure. Gray shaded regions clearly correspond to regions in which our power spectrum measurements are not consistent with thermal noise. We note that where our upper limits do agree with thermal noise, the power spectrum is on the order of ~ 100 times larger than the upper limits set with the MWA at $\approx 180 \text{ MHz}$ (D15). This factor is reasonable given that the sky noise (noise power spectrum) scales with $\sim f^{-2.6}$ ($f^{-5.2}$) and the primary beam solid angle increases as $\sim f^2$

from detecting the EoR power spectrum.

Our best upper limits fall within the region of Fourier space with systematic errors and, while we do not expect them to integrate down with more observing time, we can infer that $\Delta^2(k)$ is less than $2.5 \times 10^7 \text{ mK}^2$ at $k = 0.18 h \text{ Mpc}^{-1}$ and $z = 12.2$, $8.3 \times 10^7 \text{ mK}^2$ at $k = 0.21 h \text{ Mpc}^{-1}$ and $z = 15.35$, and $2.7 \times 10^8 \text{ mK}^2$ at $k = 0.22 h \text{ Mpc}^{-1}$ and $z = 17.05$, all at 95% confidence.

5.4.5 The outlook for EoR Measurements on the MWA.

A pertinent question arising from our analysis is how much the observed reflections impact or limit observations with the MWA of the Epoch of Reionization power spec-

trum at higher frequencies. The answer depends significantly on the calibration and reduction approach and as of now, several different efforts using alternative calibration and reduction schemes are being undertaken (Jacobs et al., 2016). The analyses in Dillon et al. (2015a) and Beardsley et al. (2016) are calibrated in a way similar to this work, employing limited calibration parameters to avoid detrimental modeling errors. After an integration time of ≈ 3 hours, Dillon et al. (2015a) also observe the cable reflections above the thermal noise level, however the smallest k mode occupied by the shortest 90 m cable lies at $k \approx 0.4 h\text{Mpc}^{-1}$ while the delay width is narrower due to the smaller primary beam, causing the wedge to occupy fewer k_{\parallel} modes. As a result, there are regions of k -space below the shortest reflection that are still consistent with noise. Such miscalibrated structure should be highly detectable after $\approx 10 - 30$ hours of integration but the results of such an analysis are still forthcoming (Beardsley et al., 2016). If the reflections can be corrected to the $\sim 10^{-3}$ level as was done in this analysis, the region below the first reflection should remain free of contamination from the beam-former to receiver reflections.

Additional calibration pipelines, which include far greater degrees of freedom, such as the Real Time System (RTS) (Mitchell et al., 2008; Ord et al., 2010) and the reduction pipeline discussed in Offringa et al. (2016) include direction dependent calibration, ionospheric phase fitting, and greater frequency resolution, are also being applied to MWA data sets. A recent upper limit at 180 MHz derived from RTS calibrated data and the CHIPS power spectrum estimator did not show evidence of the cable reflections being present (Trott et al., 2016). It is likely that the enhanced degrees of freedom allowed by RTS calibration did a better job at removing the structure from reflections but it is difficult to tell given that the error bars due to thermal noise at the comoving scales relevant to these reflections are an order of magnitude larger than those in Dillon et al. (2015a) because of the shorter integration time. Ultimately, the increase in the number of fitting parameters may enhance the removal of instrumental chromaticity in the EoR window, however simulations by Barry et al. (2016) show that small errors in ones calibration model will introduce power into the window in excess of the 21 cm signal unless the intrinsic bandpass is smooth enough

to be modeled by a small number of parameters or a source model exists with an accuracy significantly beyond what is currently available. Ultimately, the increase in fitting parameters may enhance the removal of instrumental chromaticity in the EoR window. Whether they can be introduced without adding power into the EoR window in excess of the signal, due to small errors in source modeling is still an open question that is currently being investigated.

5.5 Conclusions and Future Experimental Considerations

In this paper, we have presented low frequency radio observations with the MWA at unprecedentedly high redshifts between 11.6 and 17.9. Our goals in conducting these observations were to place upper limits on the 21 cm power spectrum during the Epoch of X-ray heating and to assess the levels of systematics which are expected to be generally worse than at EoR frequencies. These systematics include ionospheric effects, RFI (due to the FM band) and increased thermal noise. We need to control these systematics if we are to learn the detailed properties of the sources that heated the IGM; be they the first generation of stellar mass black holes, the hot interstellar medium left over from the first supernovae explosions in the universe, or dark matter annihilation.

With regards to RFI, we have found after three hours of integration that existing algorithms are sufficient to flag RFI below the FM band. Within the FM band, we have found that only a handful of channels are contaminated continuously and that after discarding them our power spectra do not show any evidence of RFI contamination. This bodes well for future planned 21 cm experiments at the MRO such as the SKA-low which is expected to make high signal to noise detections of the power spectrum (Koopmans et al., 2015). However, we are still many orders of magnitude above the level of a detection and reducing the thermal noise through longer integrations may reveal lower level RFI.

Over two nights of observing, we encounter different ionospheric conditions, observed quantitatively using the differential refraction metric described in Cohen and Röttgering (2009). We establish that ionospheric fluctuations are the source of observed position shifts by comparing the level of refraction in our two observing bands and find that they exhibit the expected λ^2 evolution. Diffractive scales on the second night of ≈ 5 km are a factor of two shorter than the first night, indicative of more severe ionospheric activity. When we compare the 1d power spectra derived from an equal amount of data on each night, we find that they are very similar to each other, lending support to the idea that since ionospheric effects on calibration and foreground residuals are spectrally smooth, they should not contaminate the EoR window.

While the majority of foreground power is contained within the wedge, we find high-significance foreground detections within the EoR window out to a $k_{\parallel} \lesssim 0.5 h \text{ Mpc}^{-1}$. These contaminated regions are consistent with miscalibrated cable reflections. We are able to obtain an order of magnitude improvement on removing the worst of these features using fits to autocorrelations, however they still limit our sensitivity at the 2-5 σ level. In addition, since auto-correlations are generally contaminated by RFI and digital artifacts, it is likely that in reducing the dominant obstacle in our data, we have introduced additional features that are below the noise level of this analysis. Since the reflections occupy the regions of k -space where we would otherwise expect the greatest cosmological sensitivity, our best upper limits are a factor of a few larger than the limits we would obtain if we were thermal noise limited. Cable reflections are especially pernicious at higher redshifts because the increasing primary beam width adds foreground power to delays ever closer to the horizon. While supra horizon emission off of the wedge moves up in k_{\parallel} , the modes occupied by cable reflections move down, increasing in width. The EoR window is crushed between the shortest reflection mode and the top of the wedge.

While our observations on the MWA will not integrate down below $\approx 10^8 \text{ mK}^2$ at $k \lesssim 0.5 h \text{ Mpc}^{-1}$ and is limited by the intrinsic spectral structure of the instrument, the systematics encountered in this analysis do not prevent 21 cm observations at

high redshift in general. A robust source catalog, that includes emission all the way down to the horizon along with precise models of the primary beam will lead to less foreground power bleeding from the edge of the wedge, (Thyagarajan et al., 2015a,c; Pober et al., 2016) and potentially open up a foreground free region under the first cable reflection. Resolving the question of cosmological signal loss and mixing of foreground spectral structure from large to short baselines may enable us to calibrate with more free parameters, better capturing the spectral structure of the bandpass. More robust calibration of these features may also be obtainable with a redundant array (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014). The 128-tile MWA has very little redundancy by design, however an additional 128 tile expansion is expected to introduce two highly redundant, hex-packed, subarrays (Tingay, private communication). The final plan for HERA, which is currently under construction, is dominated by 331 hexagonally packed dishes. Its layout is designed to take advantage of redundant calibration as well (Pober et al., 2014). Finally, calibration using injected signals (Patra et al., 2015) can also be employed to make high precision measurements of the bandpass.

The most sure way of eliminating reflection features is to remove them in hardware either by ensuring better impedance matching on the cable connections, changing the cable lengths to move reflections out of the window, or early digitization. The current HERA design employs cables no longer than 35m in length, translating to $k_{\parallel} = 0.09 h \text{ Mpc}^{-1}$ at $z = 16$ and ensures that reflections within the dish are below an acceptable level (Ewall-Wice et al., 2016a; Patra et al., 2017; Thyagarajan et al., 2016), while the planned MWA phase III upgrade and the SKA are considering digitization at the beamformers (Tingay, private communication), eliminating reflections altogether.

While measurements of the 21 cm line at EoR frequencies can teach us about the nature of UV photon sources and constrain cool thermal histories, a significant number of scenarios predict saturation of heating's contribution to brightness temperature fluctuations during reionization. In order to learn of the detailed properties of the sources that heated the IGM and to exploit the full potential of the 21 cm line as a

cosmological and astrophysical probe, we will invariably want to extend our search to as low a frequency as possible. In this work we have obtained a first look at the systematics facing us in this high redshift realm and have found that most of them are navigable. As of now, our primary limitation lies in the design of our instrument and calibration, both of which can be dramatically improved on relatively short time-scales. Ultimately, we expect a combination of improvements in instrumental design including shorter/no cables to keep reflections inside of the wedge and redundant baseline layouts allowing for more robust calibration to allow for much deeper integrations in the near future.

5.A Appendix: The Effect of Cable Reflections on Tile Gains

Throughout this work, we reference several expressions dealing with standing waves on cables that arise from mismatched impedances at their connections. In this section we derive these expressions for the reader's convenience. Discussions of this problem can be found in most elementary electricity and magnetism texts.

An voltage signal, $A(x, t)$ incident on the end of a transmission line with impedance Z_0 and length L that is terminated by some resistance R_L will be partially reflected $B(x, t)$ and transmitted $C(x, t)$. The amplitudes of the reflected and transmitted components can be found by enforcing continuity in the voltage across the connection and are given by

$$B(L, t) = \frac{Z_0 - R_L}{Z_0 + R_L} A(L, t) \equiv \tilde{R}A(L, t) \quad (5.28)$$

$$C(L, t) = \frac{2Z_0}{Z_0 + R_L} A(L, T) \equiv \tilde{T}A(L, t) \quad (5.29)$$

The impedance of a length L coaxial line is given by

$$Z_0 = R_0 + i \left(2\pi f \ell_0 L - \frac{1}{2\pi f c_0 L} \right) \quad (5.30)$$

where c_0 is the capacitance per unit length and ℓ_0 is the inductance per unit length.

A ubiquitous undergraduate electricity and magnetism exercise involves finding these quantities for a coaxial cable filled with a dielectric of permittivity ϵ and permeability μ (Griffiths, 2013), yielding

$$c_0 = \frac{2\pi\epsilon}{\ln \frac{d_o}{d_i}} \quad (5.31)$$

and

$$\ell_0 = \frac{\mu}{2\pi} \ln \frac{d_o}{d_i}. \quad (5.32)$$

Here d_i is the radius of the inner wire of the coaxial cable and d_o is the radius of the outer shell. It is clear from equation 5.30 that the reflection coefficients are dependent on frequency in a way that is influenced by the cable geometry and dielectric properties.

Now we consider the coaxial cable terminated on both ends with reflection coefficients \tilde{R}_0 and \tilde{R}_L . A monochromatic voltage signal with frequency f entering the cable at $x = 0$ with amplitude $s(f)$ will travel to the end of the cable ($x = L$) where part of it will be transmitted and the other part reflected. The complex amplitude of the transmitted component is $\tilde{T}_L(f)s(f)e^{\pi i \tau f}$ while the reflected component has complex amplitude $\tilde{R}_L s(f)e^{\pi i \tau f}$, where τ is the time it takes for the signal to propagate down the length of the cable and back. The reflected component will travel back down to $x = 0$ and be re-reflected and transmitted with an amplitude of $\tilde{T}_L(f)\tilde{R}_0\tilde{R}_L s(f)e^{3\pi i \tau f}$. We may compute the total output at $x = L$ as a series of transmitted waves where the n^{th} summand has gone through n partial reflections,

$$s_{eff}(f) = \tilde{T}_0 e^{\pi i \tau f} \sum_{n=0}^{\infty} \left(\tilde{R}_0 \tilde{R}_L e^{2\pi i \tau f} \right)^n \quad (5.33)$$

$$= s_{eff}(f) \tilde{T}_0 e^{\pi i \tau f} \frac{1}{1 - \tilde{R}_0 \tilde{R}_L e^{2\pi i \tau f}}. \quad (5.34)$$

The term $\tilde{T}_0 e^{\pi i \tau f}$ has a phase and amplitude that evolves gradually with frequency so we may treat it as part of a smooth complex gain $g(f)$ which will include the contributions from all other steps in the signal path. The gain of the tile in the

presence of reflections becomes

$$g(f) \rightarrow g(f)' = g(f) \frac{1}{1 - r e^{i(2\pi\tau + \phi)}} \quad (5.35)$$

where $r e^{i\phi} = \tilde{R}_0 \tilde{R}_L$, both terms potentially evolving with frequency.

5.B Appendix: The Power Spectrum of ionospheric phase fluctuations from measurements of Differential Refraction.

In this section, we derive the relationship between the structure function of source offsets and the underlying power spectrum of ionospheric phase fluctuations. We will adopt the common assumption that the TEC above the MWA, and hence the phases added to transiting electromagnetic waves are described approximately by a Gaussian random field (Rufenach, 1972; Singleton, 1974) whose power spectrum we denote as $P(\mathbf{k})$. In § 5.3.2, we measure the differential refraction of source positions which we may express in terms of the gradients of the phase screen.

$$\begin{aligned} D(\theta) &= 2 \left(\frac{c}{2\pi f} \right)^2 \langle \nabla\phi(\mathbf{r}_0) \cdot \nabla\phi(\mathbf{r}_0) - \nabla\phi(\mathbf{r}_0) \cdot \nabla\phi^*(\mathbf{r}_0 + \mathbf{r}) \rangle \\ &= 2 \left(\frac{c}{2\pi f} \right)^2 [\rho_{\nabla}(0) - \rho_{\nabla}(\mathbf{r})], \end{aligned} \quad (5.36)$$

where $\rho_{\nabla}(\mathbf{r})$ is the correlation function of the ionospheric gradients. We can write $\rho_{\nabla}(\mathbf{r})$ in terms of the power spectrum by expanding $\nabla\phi(\mathbf{r})$ in terms of its Fourier components

$$\nabla\phi(\mathbf{r}) = \frac{i}{(2\pi)^2} \int d^2\mathbf{k} \tilde{\phi}(\mathbf{k}) \mathbf{k} e^{i\mathbf{k}\cdot\mathbf{r}} \quad (5.37)$$

Hence,

$$\rho_{\nabla}(\mathbf{r}) = \frac{1}{(2\pi)^2} \int d^2\mathbf{k} k^2 e^{-i\mathbf{k}\cdot\mathbf{r}} P(\mathbf{k}), \quad (5.38)$$

where we have used the definition of the power spectrum,

$$\langle \tilde{\phi}(\mathbf{k})\tilde{\phi}(\mathbf{k}') \rangle = (2\pi)^2 P(\mathbf{k})\delta_D^{(2)}(\mathbf{k} - \mathbf{k}'). \quad (5.39)$$

If we assume isotropy of the field, we have

$$\rho_{\nabla}(r) = \frac{1}{2\pi} \int dk k^3 P(k) J_0(kr). \quad (5.40)$$

Thus, by measuring the structure function of source offsets, we effectively measure the power spectrum of the ionospheric fluctuations.

5.C Appendix: The Amplitude of Scintillation Noise in MWA Visibilities

In this section, we estimate the amplitude of scintillation noise present in each of the two second time steps that we interleave to estimate the system temperature. The time between the interleaved steps used to compute our system temperature is smaller than the coherence time given in V15a. However, computing the amplitude of scintillation noise, assuming that it is entirely decorrelated between our two-second time steps allows us to place an upper limit on what systematic bias in T_{sys} that might arise. We estimate the level of scintillation noise a baseline with length \mathbf{b} arising from a source population with a number density per solid angle and intrinsic flux bin given by,

$$\frac{d^2 N(S_t, f)}{dS_t d\Omega} = C S_t^{-\alpha} f^{-\beta}, \quad (5.41)$$

using equation (2.7) in V15b

$$\begin{aligned} \sigma_{\text{scint}}^2[V(\mathbf{b})] &= 4S_{eff}^2 \int d^2\mathbf{k} P(\mathbf{k}) \sin^2(\pi\lambda h\mathbf{k}^2 - \pi\mathbf{b} \cdot \mathbf{q}) \\ S_{eff}^2 &\approx \frac{CB_{eff}(f)f^{-\beta}}{3-\alpha} S_{max}^{3-\alpha}. \end{aligned} \quad (5.42)$$

Here $P(\mathbf{k})$ is the power spectrum of ionospheric phase fluctuations and S_{max} is the maximal apparent source flux for which ionospheric effects have not been calibrated out. B_{eff} is the effective primary beam of the instrument and can be computed from the equation

$$B_{eff}(f) = \int d\Omega B^{\alpha-1}(f, \boldsymbol{\ell}) \quad (5.43)$$

where $B(f, \ell)$ is the antenna primary beam. This equation is derived assuming a small field of view $\lesssim 10^\circ$. However, V15b find that it is accurate to within $\approx 10\%$ for substantially larger fields such as the MWA's.

Since we do not attempt to calibrate out the fluctuations on 2s intervals, this source flux can be obtained by setting the number of sources in the field of view of the instrument with fluxes equal to S_{max} to one (V15a),

$$S_{max}(f) = \left(\frac{\alpha - 1}{C f^{-\beta} B_{eff}(f)} \right)^{1/(1-\alpha)} s \quad (5.44)$$

For our source population, we use fits by Di Matteo et al. (2002) to the source counts observed in the 6C survey (Hales et al., 1988) at 151 MHz,

$$\frac{d^2 N}{dS_t d\Omega}(S_t, f_0 = 150 \text{ MHz}) = k \left(\frac{S_t}{.880 \text{ Jy}} \right)^{-\gamma} \quad (5.45)$$

where $k = 4000 \text{ sr}^{-1} \text{ Jy}^{-1}$ and $\gamma = 2.5$. Assuming that all of the sources have a spectral index close to the observed mean of $\delta = 0.8$, we determine the frequency dependence of the source counts by setting the number of sources with fluxes above flux S_t at $f_0 = 150 \text{ MHz}$ equal to the number of sources at f with fluxes greater than $S'_t = S_t(f/f_0)^{-\delta}$. Doing this, we obtain

$$\frac{d^2 N(S_t, f)}{dS_t d\Omega} = k \left(\frac{f}{f_0} \right)^{\delta(1-\gamma)} S_t^{-\gamma} \quad (5.46)$$

which is similar to the expression in Trott and Tingay (2015) except for an order-unity difference in the frequency power law which was neglected in that work since $f/f_0 \approx 1$ and here, where $f/f_0 \approx 1/2$ accounts for an $\approx 25\%$ enhancement in the

source counts.

We substitute $\beta = \delta(\gamma - 1) = 1.2$, $C = kf_0^{-\delta(1-\gamma)} = 1.6 \times 10^6 \text{Jy}^{-1} \text{sr}^{-1} \text{MHz}^{-2.5}$, $\alpha = \gamma = 2.5$, and an effective beam area of 5.41. Using the short dipole model of the MWA beam, we compute a $B_{eff}(83 \text{ MHz})$ of 0.33 sr. From these numbers, we obtain $S_{eff} = 212.3 \text{ Jy}$.

The final ingredient is $P(\mathbf{k})$ which we compute from our fits of our differential refraction measurements described in appendix 5.B and using the functional form in equation 5.19. Applying equation 5.42, we obtain values for $\sigma_{scint}[V(\mathbf{b})]$ between $\approx 4 - 6 \text{ Jy}$ on the 30 minute intervals on September 5th and $\approx 1 - 2 \text{ Jy}$ on the 30 minute intervals on September 6th at 83 MHz.

We estimate the noise on a single antenna for each two second interleaved time interval is given by (Morales and Hewitt, 2004)

$$\sigma_v = \frac{k_b T_{\text{sys}}}{A_e \sqrt{2df\tau}} \quad (5.47)$$

where df is the channel width, A_e is the effective area of the tile, and $T_{\text{sys}} \approx T_{\text{sky}} = 60(\lambda/\text{meter})^{-2.6} \text{ K}$ (Rogers and Bowman, 2008; Fixsen et al., 2011) which dominates the MWA's system temperature at lower frequency. Using $df = 80 \text{ kHz}$ and $\tau = 2 \text{ s}$, we obtain $\sigma_v \approx 315 \text{ K}$. Hence, on a single two second integration for each of our visibilities, scintillation noise contributes at the level of $\lesssim 2\%$ relative to the system noise during the most severe times and $\lesssim 0.3\%$ during the calmest intervals.

Chapter 6

The Hydrogen Epoch of Reionization Array Dish II: Characterization of Spectral Structure with Electromagnetic Simulations and its science Implications.

The content of this chapter was published in *Ewall-Wice, A., Bradley, R., Deboer, D., Hewitt, J., Parsons, A., Aguirre, J., Ali, Z. S., Bowman, J., Cheng, C., Neben, A. R., Patra, N., Thyagarajan, N., Venter, M., de Lera Acedo, E., Dillon, J. S., Dickenson, R., Doolittle, P., Egan, D., Hedrick, M., Klima, P., Kohn, S., Schaffner, P., Shelton, J., Saliwanchik, B., Taylor, H. A., Taylor, R., Tegmark, M. and Wirt, B. (2016), 'The Hydrogen Epoch of Reionization Array Dish. II. Characterization of Spectral Structure with Electromagnetic Simulations and Its Science Implications.'*, *ApJ* **831**, 196.

6.1 Introduction

Observations of the redshifted 21 cm radiation from neutral hydrogen in the intergalactic medium (IGM) have the potential to illuminate the hitherto unobserved *dark ages* and *cosmic dawn*, revolutionizing our understanding of the first UV and X-ray sources in the universe and how they influenced galactic evolution (see Furlanetto et al. (2006a), Morales and Wyithe (2010), and Pritchard and Loeb (2012) for reviews). Major experimental endeavors are underway to detect the 21 cm signal, with most focusing on the epoch of reionization (EoR) during which UV photons from early galaxies converted the hydrogen in the universe from neutral to ionized. One approach involves measuring the sky-averaged global signal and is being pursued by experiments such as EDGES (Bowman and Rogers, 2010), LEDA (Greenhill and Bernardi, 2012; Bernardi et al., 2015), DARE (Burns et al., 2012), SciHi (Voytek et al., 2014), ZEBRA (Ekers and Subrahmanyan, 2012), SARAS (Patra et al., 2015), and BIGHORNS (Sokolowski et al., 2015) either in their planning stages or already taking data. The global signal is also potentially observable with a zero-spacing interferometer (Presley et al., 2015; Singh et al., 2015; Venumadhav et al., 2015). Another strategy is to observe spatial fluctuations in the 21 cm emission using radio interferometers. A first generation of such experiments are attempting to obtain upper limits or a first statistical detection of the power spectrum of 21 cm brightness temperature fluctuations. These include the Giant Metrewave Telescope (GMRT) (Paciga et al., 2013), the Low Frequency ARray (LOFAR), (van Haarlem et al., 2013), the Murchison Widefield Array (MWA) (Tingay et al., 2013a), the MIT Epoch of Reionization Experiment (MITEoR) (Zheng et al., 2013), and the Precision Array for Probing the Epoch of Reionization (PAPER) (Parsons et al., 2010). Already, many of these experiments are beginning to yield upper limits on the 21 cm signal (Dillon et al., 2013; Parsons et al., 2014; Jacobs et al., 2016; Dillon et al., 2015a; Trott et al., 2016; Ewall-Wice et al., 2016d) and significant scientific results. The most stringent power-spectrum upper limit of $\approx 500 \text{ mK}^2$ by PAPER (Ali et al., 2015) is able to rule out a number of scenarios in which the intergalactic medium received little or no heating

from X-rays (Pober et al., 2015; Greig et al., 2016). As these current observatories reach their sensitivity limits, work is beginning on the next generation of interferometers, which will have the sensitivity required for a robust detection of the 21 cm power spectrum. These include the Square Kilometer Array (SKA-1 Low) (Mellema et al., 2013) and the Hydrogen Epoch of Reionization Array (HERA)(Pober et al., 2014; DeBoer et al., 2016; Dillon and Parsons, 2016) .

The primary challenge to obtaining a high redshift detection of the cosmological signal through both of these methods is the existence of foregrounds that are $\approx 10^5$ times brighter (Bernardi et al., 2009; Pober et al., 2013a; Dillon et al., 2014). Fortunately, it is expected that the foregrounds are spectrally smooth. In power-spectrum experiments, smooth foregrounds are naturally contained to a finite region of Fourier space, corresponding to large line of sight scales, known as the *wedge* (Datta et al., 2010; Vedantham et al., 2012; Parsons et al., 2012b; Thyagarajan et al., 2013; Liu et al., 2014a,b).

Since the location of each foreground on the sky determines its position in the wedge, with sources near the horizon being introduced at line of sight scales closest to the *EoR window*, the angular response of an instrument has significant implications on the amount of side-lobe power that is leaked into the EoR window. In Thyagarajan et al. (2015a,c) and Pober et al. (2016), it is found that the response of the antenna, near the horizon, has a significant impact on foreground contamination and that centralized beams with highly suppressed sidelobes are preferable.

Beyond leakage from wedge sidelobes, any structure in the frequency response of the instrument is imprinted on the foregrounds and has the potential to leak power into the EoR window at small line of sight scales, masking the signal. Indeed, sub-percent spectral features in the analogue and digital signal chains on the initial build-out of the MWA are proving to be a significant calibration challenge (Dillon et al., 2015a; Ewall-Wice et al., 2016d; Beardsley et al., 2016).

In principle, spectral structure in the bandpass of an instrument may be removed in calibration. However, many approaches rely on detailed models of the foregrounds themselves and imperfections in these models may introduce spurious spectral struc-

ture (e.g. Barry et al. (2016); Patil et al. (2016); Ewall-Wice et al. in preparation). Redundant calibration (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014) avoids relying on a detailed sky models but in its current state would be unable to correct direction dependent chromaticity that varies from antenna to antenna. Furthermore, direction dependent calibration requires imaging which is not performed in the delay power spectrum technique (Parsons et al., 2012b) that HERA is designed for. Because of our limited knowledge of low frequency foregrounds and the fidelity of calibration algorithms, it is important to design experiments whose bandpasses are as devoid of spectral structure as possible.

Drawing from the lessons of the PAPER, MWA, and MITEoR experiments, HERA (Pober et al., 2014; DeBoer et al., 2016; Dillon and Parsons, 2016) is a next generation 21 cm experiment designed to achieve a two-orders of magnitude improvement in sensitivity over current efforts, allowing it to make a robust detection of the 21 cm power spectrum during the EoR. Much of this sensitivity increase is enabled by moving the collecting area of the instrument into short baselines, with more modes outside of the wedge, and a switch from PAPER’s skirted dipoles and the MWA’s phased dipole arrays to an antenna element that consists of a feed suspended over a large reflecting parabolic dish. In Fig. 6-1 we show a PAPER antenna and one of the initial 19 HERA dishes currently being deployed in South Africa at the same site. A central requirement for HERA’s dish design is that the antenna have a response that is sufficiently smooth in frequency and a narrow beam, leaving the EoR window free of supra-horizon emission.

This paper and its companions (Neben et al., 2016; Patra et al., 2017; Thyagarajan et al., 2016) describe a multi-pronged campaign to assess the requirements and performance of the HERA dish for isolating foregrounds within the wedge. We accomplish this by establishing antenna specifications with simulations of foregrounds using the Precision Radio Interferometry Simulator (PRISim¹)(Thyagarajan et al., 2016) and verifying that the HERA primary antenna element meets these specifications with reflectometry (Patra et al., 2017) and ORBCOMM beam mapping (Neben

¹<https://github.com/nithyanandan/PRISim>




Figure 6-1: The HERA antenna element (bottom) uses a parabolic dish to achieve an order of magnitude increase in collecting area over the PAPER antenna (top). The sleeved dipole in the center of the PAPER backplane is identical to the sleeved dipole being suspended under the cylindrical skirt over the vertex of the HERA dish. The suspended feed arrangement has the potential to introduce intra-antenna reflections which we explore in this work.

et al., 2016). In this work, we present the results of time-domain electromagnetic simulations that are intended to predict the degree of spectral structure due to reflections in the HERA dish, assess the impact of this spectral structure on the leakage of foregrounds into the EoR window, and to verify consistency with the reflectometry measurements taken in the field.

This paper is organized as follows. In § 6.2 we lay out our analytic framework, describing the impact of signal-path delays (from reflections or otherwise) on foreground leakage in delay-transform power spectra. In § 6.3 we describe our time-domain electromagnetic simulations of the HERA dish element and how we extract the voltage response function. In § 6.4 we describe the results and their origin in the antenna geometry. We also verify our simulation framework by comparing its prediction for the S_{11} reflection coefficient of the HERA dish to direct field measurements described in Patra et al. (2017). In § 6.5 we apply our electromagnetic simulation results to a foreground model to determine the extent that the HERA dish’s chromatic structure compromises the EoR window and the impact on HERA’s ability to constrain reionization parameters. We summarize our conclusions in § 6.6.

6.2 The Impact of Reflections on Delay-Transform Power Spectra

In this section, we show how delayed signals in the analog signal path of an antenna lead to foreground contamination of the EoR window. Intuitively, any reflections in the signal path introduce frequency ripples in the gain of the instrument. Since time delay is the Fourier dual to frequency (Parsons et al., 2012b), reflections with larger delays introduce ripples with shorter periods. Any high-delay frequency structure imprinted on the much brighter foregrounds has the potential to mimic and swamp the signal unless it is brought below the ratio between the foregrounds and the weak cosmological signal. Equations describing the effect of direction independent reflections in an interferometers signal chain downstream of the feed are derived in Ewall-Wice

et al. (2016d). We now extend this analysis by considering the direction dependent reflections that can occur within the antenna element. We start by denoting the electric field of radiation arriving from direction $\hat{\mathbf{s}}$ at the i^{th} antenna element at position \mathbf{x}_i with angular frequency ω as $s(\hat{\mathbf{s}}, \mathbf{x}_i, \omega)$ and the time-domain field as $\tilde{s}(\hat{\mathbf{s}}, \mathbf{x}_i, t)$. Using the fact that $\tilde{s}(\hat{\mathbf{s}}, \mathbf{x}_i, t) = \tilde{s}(\hat{\mathbf{s}}, \mathbf{0}, t - \tau_i) \equiv \tilde{s}(\hat{\mathbf{s}}, t - \tau_i)$, where $\tau_i = \hat{\mathbf{s}} \cdot \mathbf{x}_i/c$, we omit \mathbf{x}_i from our notation, referring to $\tilde{s}(\hat{\mathbf{s}}, \mathbf{x}_i, t)$ as $\tilde{s}(\hat{\mathbf{s}}, t - \tau_i)$.

Reflections within the signal chain of each antenna are generally described by a complex direction-dependent reflection coefficient, $r_i(\hat{\mathbf{s}}, \omega)$ which we refer to as $\tilde{r}_i(\hat{\mathbf{s}}, \tau)$ in the delay domain. The effect of each reflection with delay τ is to add the signal to itself multiplied by \tilde{r}_i and delayed by τ . The voltage signal measured at the i^{th} antenna element, \tilde{v}_i , is the integral over solid angle of the electric fields arriving from all directions. The presence of reflections introduces a convolution of the electric field entering the antenna (delayed by τ_i) with $\tilde{r}_i(\hat{\mathbf{s}}, \tau)$:

$$\tilde{v}_i(t) = \int d\Omega \int d\tau \tilde{r}_i(\hat{\mathbf{s}}, \tau) \tilde{s}(\hat{\mathbf{s}}, t - \tau_i - \tau). \quad (6.1)$$

Applying the Fourier convolution theorem, the Fourier transform of this equation gives $v_i(\omega)$ as the simple angular integral of the product of $s_i(\hat{\mathbf{s}}, \omega)$ and $r_i(\hat{\mathbf{s}}, \omega)$.

$$v_i(\omega) = \int d\Omega r_i(\hat{\mathbf{s}}, \omega) s(\hat{\mathbf{s}}, \omega) e^{-i\omega\tau_i}. \quad (6.2)$$

The correlator of a radio interferometer records the time-averaged product of the Fourier transformed voltage streams of the i^{th} and j^{th} antennas. The time averaged correlation between the two antennas is

$$\begin{aligned} V'_{ij}(\omega) &= \langle v_i(\omega) v_j^*(\omega) \rangle_t \\ &= \int d\Omega r_i(\hat{\mathbf{s}}, \omega) r_j^*(\hat{\mathbf{s}}, \omega) I(\hat{\mathbf{s}}, \omega) e^{-i\omega\Delta\tau_{ij}} \\ &= \int d\Omega r_i(\hat{\mathbf{s}}, \omega) r_j^*(\hat{\mathbf{s}}, \omega) I(\hat{\mathbf{s}}, \omega) e^{-i\omega\mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c}, \end{aligned} \quad (6.3)$$

where $\Delta\tau_{ij} = \tau_i - \tau_j = (\mathbf{x}_i - \mathbf{x}_j) \cdot \hat{\mathbf{s}}/c$, $I(\hat{\mathbf{s}}, \omega) = \langle |s(\hat{\mathbf{s}}, \omega)|^2 \rangle_t$ is the intensity, $\mathbf{b}_{ij} =$

$(\mathbf{x}_i - \mathbf{x}_j)$, and the 'prime' on V_{ij} indicates that this is the measured visibility with antenna gains applied rather than the intrinsic visibility of the sky. Here, we have invoked the fact that electromagnetic waves arriving from different directions are incoherent (Thompson et al., 1986). This equation is the familiar interferometry equation (Thompson et al., 1986), allowing us to see that $r_i(\hat{\mathbf{s}}, \omega)$ is mathematically equivalent to the voltage beam of the antenna. This can be explained more intuitively by the fact that any component of the instrument that multiplies the signal in the frequency domain will convolve the signal in the delay domain. Convolutions of the signal in the time-domain are mathematically identical to the effect of reflections. Setting a specification on reflections is hence equivalent to setting a specification on the spectral smoothness of the antenna's voltage beam.

In order to filter spectrally smooth foregrounds from the signal, many experiments are employing the *delay transform* over frequency, defined as (Parsons et al., 2012b)

$$\tilde{V}_{ij}(\tau) = \frac{1}{2\pi} \int d\omega V_{ij}(\omega) e^{i\omega\tau}. \quad (6.4)$$

Applying this to equation 6.3, we obtain

$$\tilde{V}_{ij}(\tau) = \frac{1}{2\pi} \int d\Omega \int d\omega r_i(\hat{\mathbf{s}}, \omega) r_j^*(\hat{\mathbf{s}}, \omega) I(\hat{\mathbf{s}}, \omega) e^{-i\omega(\mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c - \tau)}. \quad (6.5)$$

Let us examine the quantity within the angular integral. Assuming that the voltage beam is perfectly flat in its spectrum and ignoring the smooth spectral structure in $I(\hat{\mathbf{s}}, \omega)$ we see that for fixed $\hat{\mathbf{s}}$, the ω integral will result in a dirac-delta function in delay, centered at $\tau = \mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c$. hence each source located at $\hat{\mathbf{s}}$ on the sky is mapped to a line $\tau = \mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c$, resulting in the much discussed "wedge" (Datta et al., 2010; Vedantham et al., 2012; Parsons et al., 2012b; Morales et al., 2012; Thyagarajan et al., 2013; Liu et al., 2014a,b).

The presence of a realistic frequency dependent beam causes each source line to be convolved in delay with the direction dependent kernel. To see this, we carry out

the ω integral in equation 6.5 and apply the Fourier Convolution theorem.

$$\tilde{V}_{ij}(\tau) = \int d\Omega \int d\tau' \tilde{I}(\hat{\mathbf{s}}, \mathbf{b}_{ij} \cdot \hat{\mathbf{s}}/c - \tau - \tau') \tilde{R}(\hat{\mathbf{s}}, \tau'), \quad (6.6)$$

where

$$\tilde{R}_{ij}(\hat{\mathbf{s}}, \tau) = \int d\tau' \tilde{r}_i(\hat{\mathbf{s}}, \tau' - \tau) \tilde{r}_j^*(\hat{\mathbf{s}}, \tau'). \quad (6.7)$$

Thus, each source line in the wedge is convolved with an antenna kernel, \tilde{R}_{ij} , which is itself the convolution of the time-reversed delay response of the voltage beam of antenna i with the complex conjugate of voltage beam of antenna j . In the remainder of this paper, we will often refer to \tilde{R}_{ij} as the *power kernel* applied to a visibility. We demonstrate the effect of foreground smearing in Fig. 6-2 for a simple model with only three sources. Without reflections, the sources would form lines intersecting zero in the two dimensional space with the baseline length, b on the x-axis and the delay, τ on the y-axis. With the reflections, the sources are smeared out, leading to supra-horizon contamination. For the sake of simplicity, we now consider the case where the beam can be separated into independent angular and frequency components, $r_i(\hat{\mathbf{s}}, \omega) = g_i(\omega) a_i(\hat{\mathbf{s}})$. For such a case, every line in Fig. 6-2 would be convolved with the same delay dependent shape, normalized to the gain of $a_i(\hat{\mathbf{s}})$. In this situation, we have

$$\tilde{V}_{ij}(\tau) = \int d\tau' \int d\tau'' \tilde{g}_i(\tau' - \tau'') \tilde{g}_j^*(\tau'') \tilde{V}_{ij}^a(\tau - \tau'), \quad (6.8)$$

where V_{ij}^a is the visibility for the achromatic voltage pattern, $a_i(\hat{\mathbf{s}})$. We can gain further insight into the behavior of the delay kernel arising from chromaticity by assuming that $\tilde{g}_i(\tau = 0) \gg \tilde{g}_i(\tau > 0)$, which should be the case at large delays for the smooth bandpasses our instruments are designed to have.

$$\begin{aligned} \tilde{V}_{ij}(\tau) &\approx \tilde{g}_i(0) \int d\tau' \tilde{g}_j^*(\tau') \tilde{V}_{ij}^a(\tau - \tau') \\ &\quad + \tilde{g}_j^*(0) \int d\tau' \tilde{g}_i(-\tau') \tilde{V}_{ij}^a(\tau - \tau'). \end{aligned} \quad (6.9)$$

Hence, to first order, the impact of reflections is to convolve the delay-transformed




Figure 6-2: A cartoon demonstration of the impact on foregrounds of the frequency dependent beam. Left: The location of three, spectrally flat, sources in delay space assuming a frequency independent beam (no reflections in the antenna element). Right: the presence of chromaticity due to delayed signal within the antenna smears the source in delay with the kernel given by equation 6.7. Since the frequency response of the dish is a function of direction on the sky, the shape of the delay kernel is different for each source line. We see that this smearing can lead to substantial supra-horizon emission. In this paper, we consider a direction independent delay-kernel that is source primarily by reflections while a more general direction-dependent kernel is explored in Thyagarajan et al. (2016)

visibility with the voltage beam of the instrument. This may be a somewhat counter-intuitive result since naive dimensional analysis might predict that the power-kernel is proportional to the square of the delay-response. This linear relation requires that the voltage beam fall roughly five orders of magnitude (the same as the dynamic range between foregrounds and signal) in the regions of delay space that we want to measure the signal.

6.3 Electromagnetic Simulations of the HERA dish element

Having formally derived the impact of reflections on foreground visibilities, we are in a position to investigate their existence in the HERA dish. In this paper, we assume a separable beam kernel whose response is given at zenith and whose high-delay components are sourced by intra-dish reflections (equation 6.9). The impact of

the direction dependent kernel which includes side-lobe chromaticity is investigated in (Thyagarajan et al., 2016). In this section, we describe the setup and parameters of our simulations (§ 6.3.1), and how we extract the voltage response function of the dish (§ 6.3.2).

6.3.1 The Simulations

The time-dependent electromagnetic behavior of the antenna was modeled using Microwave Studio, a commercial numerical simulation package produced by Computer Simulation Technology (CST). The model consists of an idealized 14 m diameter paraboloid reflector with a feed structure placed at the focus (4.5 m above the surface). The feed is a dual linear polarized sleeved dipole identical to that used as an element of PAPER but with a cylindrical skirt used in place of the angled planar reflectors. The dipole arms are modeled as copper tubing with aluminum disks and reflector surfaces. For simulation, the entire model was encapsulated in vacuum dielectric with radiative boundaries and discrete ports of $125\ \Omega$ impedance were defined at the terminals of the orthogonal dipoles. The simulation is conducted in a box encapsulating the dish geometry plus three wavelengths (at 100 MHz) on all sides with open boundary conditions.

Microwave Studio divides the model into approximately 94 million discrete hexahedral cells where the field is calculated over time using Finite Integration Technique incorporating the Perfect Boundary Approximation² (Workflow and Solver Overflow Document, CST Microwave Studio, 2015, Chapter 3.) in response to a 100-200 MHz broadband pulse of about 40 ns in duration. The model was excited in two ways: 1) via a E-W polarized plane wave entering the model along the bore sight, and 2) from one of the discrete terminals (whose analysis we focus on in § 6.4.5). The voltage responses at the terminals are monitored for these two cases over a duration of 500 ns as the excitation pulse travels throughout the model.

In Fig. 6-3 we show the geometry of the electromagnetic simulation. A 150 MHz plane wave with a Gaussian envelope is initialized above the dish vertex traveling in

²This proprietary technique is used by CST to reduce the number of simulation cells.




Figure 6-3: An illustration of our simulation products and their origin in the HERA antenna geometry. A plane wave is injected from above the feed. The electric field of the plane wave at the feed terminals (red line) along with the voltage output is recorded (black line). The feed in our simulation is situated 5 m above the bottom of the dish, hence there is a ≈ 30 ns delay between when the plane wave passes the terminal for the first time (A) and when it is first absorbed in the dipole (B), leading to the voltage response. Of concern to 21 cm experiments are the subsequent reflections between the feed and the dish (C) which can lead to larger delays that contaminate the EoR window.

the $-z$ direction. It is reflected by the dish before entering the dipole feed, hidden below the cylindrical skirt in this figure. We record the electric field voltage of the plane wave at the feed output terminals as a function of time, plotted as a red line in Fig. 6-3 along with the output voltage at the terminals for the feed polarized parallel to the plan wave (black line). The delay between the central envelope of the plane wave and the voltage output of the feed is ≈ 30 ns which corresponds to the round trip travel time from the feed to the dish vertex and back. However, while the input plane wave, modulated by a gaussian, falls off rapidly after the first ≈ 20 ns after its peak, we see that the voltage output decays far more slowly due to reflections and resonances within the antenna structure. We are able to get a qualitative feel for the amplitude of the reflections by inspecting the falloff of the time domain voltage response and see that after 60 ns it reaches ≈ -25 dB.

6.3.2 Deconvolving the Response Function

We can do much better than this. From equation 6.1, we know that the voltage output results from the convolution of the plane wave input with the voltage gain of the antenna. Since we know the input wave, a straightforward application of the Fourier convolution theorem allows us to determine the voltage response.

Since our simulation is sampled in finite time steps, we will adopt discretized notation for this section. In particular, our simulation consists of N samples, evenly spaced by $d\tau$ at times $\tau_n = n \times d\tau$. We denote the output voltage at the feed terminals at time τ_n as \tilde{v}_n . Rewriting the convolution in equation 6.1 in discrete notation, we have

$$\tilde{v}_n(\hat{\mathbf{s}}) = \sum_m \tilde{r}_m(\hat{\mathbf{s}}) \tilde{s}_{n-m}(\hat{\mathbf{s}}). \quad (6.10)$$

We may undo this convolution by taking a discrete Fourier transform (DFT) of both $\tilde{\mathbf{v}}$ and $\tilde{\mathbf{s}}$ in time, dividing them in Fourier space, and taking an inverse DFT back. Symbolically,

$$\tilde{r}(\hat{\mathbf{s}}) = \mathcal{F}^{-1} \left[\frac{\mathcal{F}\tilde{\mathbf{v}}(\hat{\mathbf{s}})}{\mathcal{F}\tilde{\mathbf{s}}(\hat{\mathbf{s}})} \right], \quad (6.11)$$

where \mathcal{F} is the Fourier transform matrix for a 1d vector of length N .

$$\mathcal{F}_{mn} = e^{-2\pi i mn/N}. \quad (6.12)$$

In Fig. 6-4 we show the amplitude of the Fourier transform of our Gaussian input, centered at 150 MHz along with the voltage response. Since our input only has support between ≈ 20 and 280 MHz, the direct ratio of our voltage response and input wave is dominated by numerical noise outside of this range. We eliminate these artifacts by multiplying our ratio by a Blackman-Harris window between 100 MHz and 200 MHz and set our estimate to zero elsewhere. From a physical standpoint, this is sensible since 21 cm experiments only observe a limited bandwidth. PAPER's correlator, which will initially serve as the HERA backend samples over a 100 MHz instantaneous frequency interval and analog filtering is applied to prevent aliasing.




Figure 6-4: The absolute value of the Fourier transform of the voltage output from our dish simulations (green line) and the input wave (blue line), normalized to the amplitude of the input wave at 150 MHz. The ratio between input and output is plotted as a red line. Since our input is limited to frequencies between ≈ 20 and 280 MHz, there are significant numerical artifacts in the ratio that causes divergence towards the plot edges. To eliminate this noise, we multiply by a Blackman-Harris window between 100 and 200 MHz and set our estimate to zero elsewhere (cyan line).

6.4 Simulation Results

We now discuss the results of our simulations. We focus on the time-domain die off of the voltage response and the resulting power kernel (§ 6.4.1), comparing it to an identical time-domain simulation of the skirted dipole antenna used by PAPER. In § 6.4.2 we investigate the dependence of the power kernel on frequency to determine whether specific parts of the HERA bandpass are more affected by reflections than others. In § 6.4.3 we determine the origin of excess response at long delays in our simulations. In § 6.4.4, we discuss potential trade-offs in eliminating this structure from the antenna by determining the impact of removing components of the HERA feed that are responsible for the contamination. Finally, in § 6.4.5 we verify our simulation framework by comparing a separate time domain simulation of S_{11} of the HERA dish to a direct field measurement with a vector network analyzer (VNA).

6.4.1 The Time Domain Response of the HERA Dish

Applying equation 6.11 to our simulation, we obtain estimates of the time-domain voltage response of the HERA dish towards zenith which we plot in Fig. 6-5. We also conduct a time-domain simulation of the voltage output in response to an identical input plane wave for the skirted dipole PAPER antenna (pictured above the HERA dish in Fig. 6-1) in order to determine whether the presence of the parabolic dish introduces reflections and spectral structure in excess of previous successful antenna designs.³ We inspect the absolute value of \tilde{v} for both the PAPER and HERA antennas in Fig. 6-5. Since non-zero values of \tilde{v} at negative delays violates causality, we assume such features are sourced by artifacts such as side-lobes of the zero-delay peak and/or numerical precision noise which sets a limit of ≈ -60 dB on the dynamic range of our method. We first note peaks in the HERA curve that are spaced by ≈ 35 ns and are absent in the PAPER simulation. Another significant difference between the two curves is a knee in the HERA gain at ≈ 120 ns that is not present in PAPER, leading to an increase in gain by approximately 20 dB at 200 ns. Because this knee does not exhibit nodes at 35 ns intervals it is probably not sourced by reflections in the antenna-feed geometry but by reflections involving a geometric length that is not resolved by the 100 MHz bandwidth of the simulation.

Our next step is to compute the power kernel, \tilde{R} , given by equation 6.7 by performing a convolution of the time-reversed voltage response with the complex conjugate voltage response. In Fig. 6-6 we show the resulting power kernel for PAPER and HERA. Since both voltage gains drop rapidly with increasing delay, the approximation in equation 6.9 holds quite well and we see that the kernels fall off at a rate similar to the voltage response.

6.4.2 The Delay Response of Subbands

Because the 21 cm brightness temperature fluctuations evolve over redshift intervals of $\Delta z \gtrsim 0.5$ (Zaldarriaga et al., 2004), experiments will sub-divide their bands into

³To date, PAPER has produced the most stringent limits on the 21 cm power spectrum leading us to use its design as a standard.




Figure 6-5: The Fourier transform of the simulated voltage response of the HERA Dish (solid black line) and the PAPER antenna element (dashed black line). Reflections in the HERA dish element lead to significantly enhanced power above ~ 50 ns. Since negative delays should be devoid of signal, they allow us to determine the dynamic range of our simulations which have a numerical noise/sidelobes floor of -60 dB.




Figure 6-6: The absolute value of the power kernel for the HERA dish (solid black line) and for the PAPER antenna element (dashed black line) calculated using equation 6.7. While an antenna can only physically have a voltage response at positive delays, the delay kernel is formed from the convolution of one antenna with the time reversed conjugate response function of the other. Hence, the power kernel for two identical antennas will have $\tilde{R}(\tau) = \tilde{R}^*(-\tau)$.

≈ 10 MHz intervals for power spectrum estimates at multiple redshifts. Thus the localization of reflections within the HERA band will determine whether all or only a subset of redshifts are affected. To determine whether the reflections are localized in frequency, we compute the voltage delay response and power kernel for three different subbands: 100 – 130 MHz, 130 – 160 MHz, and 160 – 190 MHz. In order to maintain decent resolution of the kernel itself, we use frequency ranges that are several times larger than the actual subbands that will be used for EoR power spectrum estimation (≈ 10 MHz). The resulting power kernels are plotted in Fig. 6-7. The central lobe of the subband kernels is significantly wider due to the wider window functions incurred by the reduced bandwidth. We find that the shallow long-term falloff is only visible within the central 130 – 160 MHz band, indicating that long term reflections are isolated around 150 MHz.

To further illustrate the isolation of fine frequency structure in the center of the bandpass and to verify that our observations are not an artifact of our reduction of the simulation outputs, we fit 10 MHz intervals of the absolute value of the simulated gains to a sixth order polynomial and inspect the residuals in Fig. 6-8. We find that the gain residuals on the sixth order fit are an order of magnitude greater on the 145 – 155 MHz subband than any other frequency interval.

6.4.3 The Origin of the Knee.

In Fig. 6-5 we see that the long-term delay response of the HERA antenna differs from the PAPER design in two regards, the existence of node-like structures spaced every ~ 35 ns, corresponding to the round trip delay between the dish and the feed, and a long-term knee that dominates the response function after 100 ns but does not exhibit any periodicity that might be associated with the feed-dish geometry. Here we establish the origins of the long time-scale structure using simulations. Reflectometry measurements conducted on the isolated feed in several different configurations (Patra et al., 2017) independently discover and verify these origins. To determine whether this knee-like structure originates from feed-dish reflections, we perform a simulation of a plane wave incident along the bore-sight the HERA feed without the dish and




Figure 6-7: The power kernel for the three subbands discussed in § 6.4.2 along with the kernel for the full bandwidth response function. While the long term falloff from reflections is prominent between 130 – 160 MHz, it appears at a much lower level in the other two subbands which fall below the central subband by ~ 20 dB at ~ 300 ns. k_{\parallel} values for each delay are computed at 150 MHz. The wider central lobe below 150 ns for the subband gains is due to the lower delay resolution from the smaller bandwidth. We also show the delay kernel for 100 MHz bandwidth (black thick line).




Figure 6-8: Residuals on the absolute value of the gain over several subbands after fitting to a sixth order polynomial. Consistent with our findings in Fig. 6-7, the fine frequency residuals in the 145-155 MHz subband are over an order of magnitude greater than those in the other subbands.

compare it with the time-domain response for the feed-and-dish system in Fig. 6-10. The feed only simulation lacks the reflections associated with the dish and feed configuration (present at small delays) but retains the large-delay knee. Thus, the dominant large time-scale contamination of the HERA band-pass is not intrinsic to the arrangement of a feed suspended over the dish.

We isolate the source of the knee within the feed structure (shown in Fig. 6-9) by running simulations of the dipole feed with only a back-plane along with a sleeved dipole in isolation. We show both of these voltage responses in Fig. 6-10. When the cylindrical skirt is removed from the feed, the knee vanishes. Because of the narrow band nature of the knee (§ 6.4.2), its presence only when the cylinder is attached, and its exponential falloff (characteristic of a damped harmonic oscillator), we conclude that the cylindrical skirt is behaving as a resonant cavity and the resulting stored oscillating fields are responsible for the majority of the fine-scale spectral structure in the HERA antenna's gain. We estimate the quality factor of the resonance to be ≈ 6.5 from the time constant of the exponential falloff which yields a resonance width of ≈ 20 MHz.

Since the termination of the dipole could be significantly different from the termination used in the simulation (100Ω), we also investigate the impact of the termination impedance of the dipole on the delay-response. we vary the termination impedance between 50 to 500Ω for several plane-wave simulations and show the resulting voltage responses in Fig. 6-11. We find that changing the termination impedance has a significant impact on the time-domain response below ≈ 150 ns in the region that is dominated by dish-feed reflections. Since the termination impedance does not affect resonance of the feed cavity, its effect is small beyond ≈ 200 ns. Only in the extreme, 500Ω case do the reflections extend to appreciable delays. Since delays over ≈ 250 ns are responsible for contaminating the EoR window, termination impedance will only impact HERA's sensitivity for an extremely poor match.

Simulations combining the dish with the analog signal-chain show that matching networks can reduce the levels feed-dish reflections below those observed in this paper,




Figure 6-9: A closeup rendering of the HERA feed which is suspended over the reflector, illustrating the cylindrical skirt, the backplane, and the dipole. Long time-scale spectral structure arises from electrical oscillations within the cylindrical cavity.

even under realistic termination conditions (Fagnoni and de Lera Acedo, 2016)⁴.

Since the feed properties dominate the antenna’s long-term response, we investigate the voltage response (Fig. 6-5) and the power kernel (Fig. 6-6) of the HERA antenna without the cylindrical skirt and a bare dipole. Removing the cylinder leads to a ≈ 10 dB improvement in the antenna response after 100 ns and eliminating the backplane as well yields a 20 to 30 dB improvement. Hence, the delay response of the HERA antenna can potentially be improved by modifying the feed. There are several trade-offs in polarization and sensitivity that are made when we remove feed components which we now discuss.

6.4.4 Tradeoffs in Beam Properties.

We determined above that the long-delay *knee* in the gain of the HERA dish is primarily caused by resonance in the cylindrical skirt and backplane that surround the PAPER dipole. These structures are added to enhance the symmetry of its beam, to improve the effective area of the antenna, and reduce cross coupling (DeBoer et al., 2016). It is therefore worth examining the impact that removing these components has on the properties of the HERA beam.

We first examine the trade-off that is made in polarization when we remove elements of the original design. Faraday rotation is capable of imparting fine-scale

⁴See Fig. 4. This paper investigates a different variant of the HERA feed in which the resonance is absent but whose primary beam properties are still being investigated.




Figure 6-10: The absolute value of the time-domain voltage response of the cylindrical dipole feed compared to the absolute value of the voltage response of the feed suspended over the dish. As we might expect, the ~ 35 nm lobed structures associated with feed-dish reflections are absent from the simulation of the feed only. However, the knee like feature after ~ 100 ns is. This indicates that the most severe spectral contamination in the current HERA design does not originate in reflections between the feed and the dish but rather within the feed itself.




Figure 6-11: The time-domain response of the HERA antenna towards a plane wave incident from zenith for a variety of termination impedances. As we vary the termination impedance, the structure which is dominated by feed-dish reflections, below ≈ 150 ns, varies significantly but leaves delays greater than ≈ 200 ns unchanged. Only in the extreme, 500Ω case do the reflections extend to large delays. Since structure above 250 ns is primarily responsible for contaminating the EoR window (6-14), the termination impedance has a relatively small impact on HERA's overall sensitivity.

frequency structure into the Q/U polarized Stokes visibilities. Antennas in which the X and Y polarized beams are not identical can leak stokes Q/U visibilities and the spectral structure they contain into the Stokes-I visibilities from which delay power spectra are formed, potentially masking the 21 cm signal (Jelić et al., 2010; Moore et al., 2013, 2017). It is shown in Moore et al. (2013) and Moore et al. (2017) that this leakage is given approximately by

$$P_{leak} \approx \frac{A_-}{A_+} P_P \equiv \xi P_P, \quad (6.13)$$

where P_P is the linearly polarized power spectrum, and

$$A_{\pm} = \int d\Omega |A_{xx}(\hat{\mathbf{s}}) \pm A_{yy}(\hat{\mathbf{s}})|^2, \quad (6.14)$$

where A_{xx} is the x polarized beam and A_{yy} is the the y polarized beam. In Table 6.1 we give ξ at 150 MHz, computed for the three configurations of the HERA feed discussed in § 6.4.3. We see that the feed with the cylinder has the smallest ellipticity with $\xi \approx 3 \times 10^{-3}$ over much of the band. Polarization measurements from Asad et al. (2015) indicate a polarized power spectra of $P_P \approx 10^3 \text{ mk}^2$ within the EoR window. Thus, an antenna with $\xi \approx 3 \times 10^{-3}$ would produce leakage of several mk^2 which is roughly an order of magnitude below typical reionization peak amplitudes of several tens of mK^2 . However, the feed variant in which the cylinder has been removed has an ellipticity that is roughly $\approx 4\times$ larger and a bare dipole, $\approx 10\times$ larger. These enhanced ellipticities would bring the level of the polarization leakage to or above the level predicted in Asad et al. (2015). On the other hand, it is possible the LST binning and averaging over many nights with different rotation measures can suppress polarization leakage by a factor of $\approx 10^3$ in the power spectrum (Moore et al., 2017).

The second impact that we consider, is the effective area of the antenna, given by (Wilson et al., 2009),

$$A_{\text{eff}} = \frac{\lambda^2}{\int d\Omega A(\hat{\mathbf{s}})}. \quad (6.15)$$

In Table 6.1, we see that the effective area at 150 MHz is actually enhanced by the

Table 6.1: Ellipticity and effective area of the HERA antenna for different feed configurations.

Feed Configuration	Ellipticity, ξ	Effective Area, A_{eff} , (m^2)
Cylindrical and Backplane	3.5×10^{-3}	56.7
No Cylinder	1.5×10^{-2}	97.2
No Cylinder or Backplane	2.8×10^{-2}	41.9

removal of the cylinder but is negatively impacted by the removal of the backplane which introduces a significant back-lobe, reducing the beam’s directivity. As detailed in Neben et al. (2016), this reduction in the effective area of the dish lowers HERA’s sensitivity by a factor of order unity.

The development of the HERA feed is still underway and it is possible that improvements in the feed will allow for the 150 MHz resonance to be eliminated while maintaining the illumination and beam-symmetry of the original feeds.

6.4.5 Verifying Our Framework with S_{11} Measurements and Simulations

We now assess the accuracy of our time-domain simulation framework by comparing simulations to measurements of the S_{11} parameter of the HERA dish. Up until now, our simulations have derived the voltage response of the dish using simulations of an incoming plane wave as is the case for radio signals arriving from objects at cosmological distances. It is possible to probe the gain of the dish using objects in the far field such as known radio sources (Thyagarajan et al., 2011; Pober et al., 2012; Colegate et al., 2015) or constellations of ORBCOMM satellites (Neben et al., 2015, 2016). However natural radio sources are too weak to probe the dish response at the $\lesssim 10^{-4}$ level necessary to verify our simulations and the ORBCOMM technique can only be used to map the gain at the 137 MHz ORBCOMM transmission frequency. Work is currently underway to use broad-band transmitters flown into the far field of the dish by drones (Jacobs et al., 2017) but this system is still under development. Reflectometry of the dish with a VNA is a straightforward alternative used in Patra et al. (2017) to estimate the gains directly. Rather than comparing their estimate of




Figure 6-12: A comparison between time-domain simulations (black line) and measurements (grey line) of S_{11} for the HERA dish. We also show an S_{11} measurement with the cables leading from the VNA to the feed terminated by an open circuit which allows us to probe the dynamic range of the measurement. We use the standard deviation of the open measurement (grey dashed line) between 200 and 400 ns as our systematic floor (grey shaded region). We find very good agreement between our S_{11} measurement and the simulation, validating the predictions of our simulations. Both the simulations and measurements in this figure were derived from delay transforms over 100 MHz.

the gain with our predicted gains (which is done in their paper), we set up a time domain simulation to compute the S_{11} parameter of HERA antenna and compare it to direct S_{11} measurements.

S_{11} refers to the complex ratio between a voltage signal transmitted into the feed terminals and the voltage reflected back as a function of frequency, $S_{11}(f) \equiv v_{\text{trans}}(f)/v_{\text{recv}}(f)$. Measurements, described in further detail in Patra et al. (2017), of S_{11} were taken on a prototype HERA dish at the National Radio Astronomy Observatory Green Bank facility using an Agilent 8753D VNA. The VNA was connected to the antenna’s balanced feed terminal via a 30.5 m length of RG-8X-LL coaxial cable and M/A Com HH-128 180 degree hybrid junction. The VNA was calibrated at the balanced end of the hybrid junction using a set of termination standards having SMA connectors. A small adapter from SMA to the post terminals on the feed was not included in the calibration. The VNA excites the terminals with a band-limited (100-200 MHz) pulse, and the complex reflected voltage at the calibration plane is measured as a function of frequency. The reflection for the case of open terminals was

measured to confirm the dynamic range and resolution of the measurement. While our simulations are terminated with an impedance of $125\ \Omega$, the VNA measurements are terminated at $100\ \Omega$ which only has a small effect on the large-delay response of the antenna (Fig. 6-11).

Because S_{11} is defined as a ratio in the frequency domain in a form identical to our voltage gains, we may run time-domain simulations similar to those described in § 6.3.1 except rather than simulating an incoming plane wave, we simulate the excitation of the feed terminals by a delta-gap impedance port and record output voltage as a function of time. We calculate S_{11} from the simulation in the same manner that we obtained \tilde{r} using equation 6.11 with \tilde{v}_{recv} taking the place of \tilde{v} in the numerator inside the Fourier transform and \tilde{v}_{trans} taking the place of \tilde{s} in the denominator.

In Fig. 6-12 we show the simulated amplitude of S_{11} as a function of frequency for the HERA dish, observing a distinctive two peaked structure before a steep die off in delay that transitions to a shallower falloff at ≈ 150 ns. The first peak is due to the reflection of the input wave off of the back of the feed while the second, roughly 35 ns later arises from the transmitted component of the input wave reflecting off of the dish and arriving back at the feed. The ensuing long term die off arises from reflections within the feed and dish structure and for reasons that will be elaborated on in Patra et al. (2017) corresponds very closely to our simulations of the dish gain itself (compare with Fig. 6-5). We get a sense of the dynamic range of the measurement by unhooking the SMA adapter that attaches the sleeved dipole feed to the cable from our VNA, forming an open circuit that should ideally give a reflection coefficient of ≈ 1 at zero delay with no reflections at any other times. In this measurement, we find noise-like structure at ≈ -50 dB. Below ≈ 500 ns, we found that this structure does not integrate down with time, leading us to conclude that it is caused by systematics, likely uncalibrated low level reflections in the VNA-feed cables. The level of this noise sets the systematic floor in our measurements which we show as a grey shaded region in Fig. 6-12. We find that in the region where the S_{11} measurement is above the systematics floor, there is good agreement with our simulations (within several dB).

6.5 The Effect of the Spectral Structure due to Reflections and Resonances in the HERA dish on Foreground Leakage and Sensitivity

We can now explore the impact of the dish’s performance on the leakage of foregrounds out of the wedge, and into the EoR window. Beyond the delay kernels considered in this paper and Patra et al. (2017), the extent of leakage will depend both on the angular structure of the primary beam, which is established through measurements and simulations in Neben et al. (2016) and the brightness of the foregrounds themselves Thyagarajan et al. (2016). In this section, we investigate the amplitude of foreground leakage as a function of delay given the angular primary beam model and our simulation of the dish’s spectral structure. We start by performing a physically motivated extrapolation of the delay structure observed in our simulations to cover the range of comoving scales relevant to power spectrum measurements (§ 6.5.1). In § 6.5.2, we combine this extrapolated delay response with simulations of foregrounds to determine the overall level of foreground power in HERA’s visibilities. This leakage will cause large-scale line-of-sight Fourier modes to be contaminated by foregrounds and hence inaccessible to the foreground filtering approach. Since the signal-to-noise ratio is maximized at the smallest k values, the loss of these modes will reduce the significance of the power spectrum detection and negatively impact the overall bottom line of the science that HERA can accomplish. We explore the impact of the chromaticity due to reflections on science using the Fisher matrix formalism in § 6.5.3.

6.5.1 Extrapolating the Bandpass and Power Kernel

Our simulations of the dish response only extend to ≈ 400 ns. However, interferometers are expected to be most sensitive to the 21 cm power spectrum at comoving scales between $k \approx 0.1 - 0.5 h \text{ Mpc}^{-1}$, corresponding to the delays between 180 and 900 ns at $z = 8$. To extrapolate out to the higher delays in this range, we assume that the response function is dominated by a sum of field oscillations within the feed’s

cylindrical skirt and reflections between the feed and the dish.

$$\tilde{r} = \tilde{r}_o + \tilde{r}_r \quad (6.16)$$

The long term falloff after the knee, which is dominated by oscillations within the cylinder, appears as a line on a linear-log plot (Fig. 6-5). This is the response we would expect for a damped harmonic oscillator. We thus model the voltage response from the cylindrical skirt as an exponential

$$\tilde{r}_o = A_o X_o^{-\frac{\tau 150 \text{ MHz}}{2Q}} \quad (6.17)$$

where Q is the quality factor of the cavity which we determine to be ≈ 6.5 .

We see in Fig. 6-5 that when the cylinder and backplane are removed, the response of the antenna falls off much faster, in line with the response below ≈ 130 ns. Like the response from the cylinder, this falloff also goes exponentially which we now show is also indicative of reflections.

We let Γ_d represent the reflection coefficient of the dish vertex and Γ_f represent the reflection coefficient of the feed in the presence of each other. An electromagnetic wave incident on the feed, at $t = 0$, is accepted with amplitude $(1 + \Gamma_f)$. The reflected component travels back to the dish and acquires an amplitude of $(\Gamma_f \Gamma_d)$ before returning at time, τ_d later where $(1 + \Gamma_f)$ will be accepted and Γ_f will be reflected back towards the dish. Summing the infinite series of reflections, the time dependent voltage at the feed is

$$\tilde{v}_r(t) = (1 + \Gamma_f) \sum_m (\Gamma_f \Gamma_d)^m \tilde{s}(t - m\tau_d), \quad (6.18)$$

which implies that

$$\tilde{r}_r(\tau) = (1 + \Gamma_f) \sum_m (\Gamma_f \Gamma_d)^m \delta_D(\tau - m\tau_d). \quad (6.19)$$

In practice, each delta-function in equation 6.21 is convolved with the window function

arising from the finite bandwidth of the instrument. For our simulations, which have a bandwidth of 100 MHz, the width of the window function is ≈ 10 ns, partially filling in the gaps between each reflection. Since the number of reflections is $m = t/\tau_d$, than we can write the long-term delay response in discrete form as

$$\tilde{r}(nd\tau) \approx (\Gamma_f \Gamma_d)^{nd\tau/\tau_d} \quad (6.20)$$

which is an exponential in time. We thus model the response due to reflections between the feed and the dish as an exponential

$$\tilde{r}_r = A_r X_r^{(\tau/30\text{ns})}. \quad (6.21)$$

To extrapolate beyond 400 ns, we fit the voltage response to the sum of two exponentials for \tilde{r}_o and \tilde{r}_r and set the feed response beyond the maximum time our simulation to this power law sum. In our pedagogical treatment, we have assumed that the reflection coefficients are frequency independent which is not the case in real life. The impact of frequency dependent reflection coefficients is to replace the summands in equation 6.19 with the Fourier transform of $(\Gamma_f \Gamma_d)^m$ evaluated at $t = (\tau - m\tau_d)$. As long as $(\Gamma_f \Gamma_d)^m$ is compact in delay space, which is the case if the reflection coefficients evolve smoothly with frequency, than our pedagogical approximation still yields the power law in equation 6.21.

6.5.2 The Impact of delay response of the HERA Antenna on Foreground Contamination

Given the chromaticity due to reflections and resonances in the HERA antenna, what Fourier modes will still be accessible with the delay filtering technique? To answer this question, we combine our extrapolated simulations of the HERA dish's spectral structure with simulations of foregrounds.

The foreground model is discussed in detail in Thyagarajan et al. (2015a) but we describe them briefly here for the reader's convenience. It consists of two major com-

ponents: diffuse synchrotron emission from our Galaxy whose structure is described by the Global Sky Model (GSM) of de Oliveira-Costa et al. (2008) and a population of point sources including objects from the NRAO Sky Survey (NVSS) (Condon et al., 1998) at 1.4 GHz and the Sydney University Mologolo Sky Survey (SUMMS) (Bock et al., 1999) at 843 MHz. Their fluxes are extrapolated to the observed 100-200 MHz band using a spectral index of $\langle\alpha\rangle = -0.83$ determined in Mauch et al. (2003). Visibilities are computed from the diffuse and point source models assuming an achromatic angular response for the HERA beam at 137 MHz described in Neben et al. (2016). We compute two sets of visibilities: one in which the spectral structure of the dish is assumed to be completely flat, and another in which the beam is multiplied by the frequency dependent gain at zenith determined by our simulations.

The foreground filtering procedure employed by PAPER and HERA involves delay transforming the visibilities and performing a 1D CLEAN (Parsons and Backer, 2009; Parsons et al., 2012b) which discovers and subtracts foregrounds within the horizon plus a small buffer, allowing for the suppression of foreground side-lobes in delay space. The level of foreground subtraction possible by this procedure is limited by the thermal noise level on the visibility, which in turn depends on the number of time steps and redundant baselines that are averaged before performing the CLEANing step. In this work, we assume that each visibility is CLEANed independently with a twenty minute cadence. Each short baseline is coherent for ≈ 10 s per night (Ali et al., 2015) so this is equivalent to each local sidereal time (LST) being integrated for ≈ 120 nights before CLEANing. If we were to average over redundant baselines this number would drop enormously. The standard deviation on the real and imaginary part of a single delay transformed visibility is (Morales and Hewitt, 2004)

$$\sigma_V = \frac{\sqrt{2B}k_B T_{sys}}{A_e \sqrt{\tau}}, \quad (6.22)$$

where A_e is the effective area of the dish, B is the bandwidth, T_{sys} is the system temperature, τ is the integration time, and k_B is the Boltzmann constant. For T_{sys} we use the equation $T_{sys} = 100\text{K} + T_{sky}$ where 100 K is the temperature of the PAPER

receiver and $T_{sky} = 60(\lambda/1 \text{ meter})^{2.55}$ is the sky temperature (Rogers and Bowman, 2008; Fixsen et al., 2011). For A_e we use the value of 98 m^2 determined in Neben et al. (2016). We CLEAN down to five times the thermal noise level. Subtracting a model of the sky is an alternative to CLEANing that could reach below the noise level on a single visibility but also assumes a priori knowledge of the foregrounds and the instrumental response. In Fig 6-13 we compare the delay transform of visibilities before and after CLEANing at the LST of 4 hours. While CLEANing is able to remove structure within the horizon, it does not reduce any of the power leaked outside of the horizon by the chromaticity of the dish. If we knew the spectral response of the dish perfectly, then we might be able to CLEAN with this kernel and remove the supra-horizon structure though the true direction-dependent nature of the chromaticity would complicate this task.

To form estimates of the 21 cm power spectrum, we split each visibility into Blackman-Harris windowed sub-bands centered at redshift intervals of $\Delta z = 0.5$ and each with a noise equivalent bandwidth of 10 MHz^5 . The flat sky approximation allows us to Fourier transform each interval in frequency, square, and multiply by a set of prefactors to obtain a power spectrum estimate (Parsons et al., 2012a, 2014),

$$\hat{P}(\mathbf{k}) = \left(\frac{2k_B}{\lambda^2} \right)^2 \frac{X^2 Y}{B_{pp} \Omega_{pp}} |\tilde{V}(\mathbf{u})|^2. \quad (6.23)$$

Here, λ is the central wavelength of the observation, k_B is the Boltzmann constant, B_{pp} is the integral of the square of the bandpass, and Ω_{pp} is the integral of the primary beam squared over solid angle. X and Y are linear factors converting between native interferometry and cosmological coordinates, defined through the relation $2\pi \mathbf{u} = 2\pi(u, v, \eta) = (Xk_x, Xk_y, Yk_z)$.

As a drift scan instrument, HERA will observe the sky at many LSTs within the declination stripe that passes through its primary beam, averaging over the power spectrum estimate at each LST. It is well documented that foreground power varies

⁵The end to end width of each Fourier transformed interval is 20 MHz, however the noise equivalent width of this interval (given by the integral of the bandpass squared) is only ≈ 0.5 the full width since the Blackman-Harris window suppresses the edge channels.




Figure 6-13: The absolute magnitude of a 100-200 MHz delay transformed visibility from a 14-meter baseline (blue line) compared to the same visibility (green line) contaminated by the delay-response observed in our simulations of the HERA dish. We see that the extended delay kernel smooths out structure, originating from foregrounds, within the horizon. For HERA, we expect to use the delay-CLEAN to remove foregrounds. However, the depth of CLEANing is limited by the noise level on a single baseline. We show the foreground residuals from a CLEAN down to the 5σ noise level after 20 minutes of integration. Since CLEANing cannot distinguish between foregrounds and signal, it should only be performed within a narrow region of delay-space, close to the horizon and cannot remove the broad wings leaked by the resonance unless it is accurately modeled.

significantly over LST (Thyagarajan et al., 2015a), hence such an estimate will either filter or weight LSTs in a way that minimizes the impact of the most contaminated observations. For our analysis, we focus on a single LST of 4 hours which is a patch of sky with low Galactic contamination. Such a patch is representative of the kind that HERA will focus most of its observing time on.

Computing the power spectra, we inspect the amplitude of foregrounds given the delay-response and angular pattern of the HERA dish for baselines of two different lengths in Fig. 6-14. In both baselines, we find that the residuals after CLEANing tend to be at similar levels except at the subband centered at $z = 8.5$ (150 MHz) where the cylinder resonance is present. In bands outside of the resonance, foreground power outside of the horizon is dominated by side-lobes of the finite window function used

in the Fourier transform. Hence, if we remove power inside of the horizon through cleaning, the level of this leakage is reduced. Power that is not reduced by the 1d clean is caused by intrinsic super-horizon structure in the antenna gain. Over the resonance, foreground residuals remain above the signal level out to $k_{\parallel} = 0.23 h\text{Mpc}^{-1}$. Several other baselines especially those oriented entirely in the E-W direction, which we do not show, contain up to two orders of magnitude greater galactic contamination as is noted in Thyagarajan et al. (2015c). In data analysis, these baselines would be down-weighted or discarded so that they do not bias our final estimate. Outside of the resonance, beam chromaticity is dominated by beam-antenna reflections and since the level of foreground contamination is nearly identical to the achromatic beam, we conclude that feed-dish reflections do not pose a significant limitation HERA’s ability to measure the EoR power spectrum. Dishes are therefore a viable strategy for scaling the collecting area of 21 cm interferometry experiments at low cost.

The level of the foregrounds in Fig. 6-14 is conservative since no attempt has been made to apply inverse covariance weighting techniques (Tegmark, 1997a; Liu and Tegmark, 2011; Dillon et al., 2013; Parsons et al., 2014; Liu et al., 2014a,b; Dillon et al., 2015b,a; Trott et al., 2016) or fringe rate filtering (Parsons et al., 2016a). Applications of inverse covariance filters in recent PAPER observations have yielded reductions in foreground power by greater than an order of magnitude (Ali et al., 2015), hence our simulations show that even with the presence of the enhanced spectral structure from reflections, HERA will be able to isolate foregrounds well below the level of thermal noise in most of the EoR window.

6.5.3 The Implications of Reflections and Resonances on EoR Science

A primary near-term goal of 21 cm EoR observations is to obtain information about the nature of the sources that drove reionization. Since the amplitude of the 21 cm signal is largest at smaller k values, a loss of large scale signal due to foreground contamination eliminates the modes on which HERA would otherwise have the greatest




Figure 6-14: The delay transform power spectrum for different baselines over several redshifts with (green solid line) and without (blue solid line) the presence of the simulated chromaticity in the HERA dish. Each estimate is computed using the square of a Blackman-Harris windowed delay transform (equation 6.23) over a noise equivalent bandwidth of 10 MHz. We also show the power spectra of foreground residuals after CLEANing the entire 100 MHz band with (dashed cyan line) and without (red dashed line) the simulated spectral structure. For all subbands, except $z = 8.5$, we find that the delay response of the HERA antenna has a negligible effect on the k_{\parallel} where the foregrounds drop below the signal level. This is the same subband where enhanced spectral structure due to resonance in the feed cylinder is observed in our simulations (Figs. 6-7 and 6-8). Outside of the resonance, the chromaticity of the antenna, which is primarily sourced by feed-dish reflections, is practically identical to the achromatic model. It follows that feed-dish reflections have a negligible impact on

signal to noise detections, reducing its overall sensitivity. In this section, we use the Fisher Matrix formalism to estimate the impact of the HERA’s intra-dish reflections on its sensitivity and ability to constrain the astrophysics of reionization. The Fisher Matrix allows us to forecast the covariances and errors on reionization parameters given errors on power spectrum observations due to the uncertainties caused by cosmic variance and thermal noise which is in turn determined by the uv coverage and observing time of the interferometer. The covariance between the parameters, θ , of an astrophysical model is given by the inverse of the Fisher matrix, \mathbf{F} which for Gaussian and independently determined power spectrum bins may be written approximately as (Pober et al., 2014),

$$F_{ij} \approx \sum_{k,z} \frac{1}{\sigma^2(k,z)} \frac{\partial \Delta^2(k,z)}{\partial \theta_i} \frac{\partial \Delta^2(k,z)}{\partial \theta_j}, \quad (6.24)$$

where $\Delta^2(k,z)$ is the power spectrum amplitude for some $k-z$ bin and $\sigma^2(k,z)$ is the variance of the power spectrum estimate. We write this equation as approximate since it ignores additive terms arising from the dependence of $\sigma^2(k,z)$ on model parameters due to cosmic variance which only contribute at the $\approx 1\%$ level (Ewall-Wice et al., 2016b).

To simulate Δ^2 , we use the publicly available 21cmFAST⁶ code (Mesinger et al., 2011) which generates realizations of the 21 cm brightness temperature field using the excursion set formalism of Furlanetto and Loeb (2004). We employ a popular three parameter model of reionization (Mesinger et al., 2012) with the following variables

- ζ : The “ionization efficiency” is defined in Furlanetto and Loeb (2004) to be the inverse of the mass collapse fraction necessary to ionize a region and is computed from a number of other physical parameters including the fraction of collapsed baryons that form stars and the UV photon escape fraction. Because ζ acts as an efficiency parameter, its primary effect is to change the timing of reionization. We choose a fiducial value of $\zeta = 20$, though expected values range anywhere between 5 and 50 (Songaila and Cowie, 2010).

⁶http://homepage.sns.it/mesinger/DexM__21cmFAST.html

- R_{mfp} : The presence of Lyman limit systems and other potential absorbers within HII regions causes UV photons to have a finite mean free path denoted by R_{mfp} . In the 21cmFAST framework, HII regions cease to grow after reaching the radius of R_{mfp} , primarily impacting the morphology of the signal. We choose a fiducial value of $R_{\text{mfp}} = 15$ Mpc which is in line with recent simulations accounting for the subgrid physics of absorption (Sobacchi and Mesinger, 2014).
- $T_{\text{vir}}^{\text{min}}$: The minimal mass of dark matter halos that hosted ionizing sources. While in principle, halos with virial temperatures as small as 10^2 K are thought to be able to form stars (Haiman et al., 1996b; Tegmark et al., 1997a), thermal and mechanical feedback have been seen to raise this limit to as high as 10^5 K (Springel and Hernquist, 2003; Mesinger and Dijkstra, 2008; Okamoto et al., 2008). We choose a fiducial value of $T_{\text{vir}}^{\text{min}} = 1.5 \times 10^4$ K which is set by the atomic line cooling threshold.

In order to account for the degeneracies in the power spectrum between heating from X-rays and reionization from UV photons, we also marginalize over three additional parameters that describe the impact of heating from early X-ray sources as explored in (Ewall-Wice et al., 2016b). These are the X-ray heating efficiency, f_X ; the maximal energy of X-ray photons that are self absorbed by the ISM of early galaxies, ν_{min} ; and the spectral slope, α which are taken to have fiducial values of 1, 0.3 keV, and -1.2 respectively. We choose to parameterize our model in terms of the fractional difference of each variable from its fiducial value so that, for example, $\theta_\zeta = (\zeta - \zeta_{\text{fid}})/\zeta_{\text{fid}}$ and compute the derivatives in equation 6.24 by performing a linear fit to realizations of the 21 cm power spectrum calculated by 21cmFAST at $\theta_i = \pm 10^{-2}, \pm 5 \times 10^{-2}, \pm 10^{-1}$, and $\pm 2 \times 10^{-2}$.

For each measurement in the uv plane, the standard deviation of a power spectrum measurement ($\sigma^2(k, z)$) is given by the direct sum of sample variance and thermal noise (McQuinn et al., 2006) which in turn depends on the primary beam of the instrument and the time spent sampling each uv cell. For our analysis, we assume that the uv plane is sampled by circular apertures with effective areas of 98 m^2 and that

$\tau(\mathbf{k})$ is determined by a drift scan in which non-instantaneously redundant baselines are combined within each uv cell. We compute the standard deviation of each (k, z) bin for the proposed 350-element deployment of HERA (HERA-350) using the public `21cmSense`⁷ code (Pober et al., 2013b, 2014).

We have seen in Fig. 6-14 that the simulated chromaticity of the dish leaks foregrounds beyond the horizon to varying degrees depending on the subband with the resonance occurring at $z = 8.5$ (150 MHz). In this subband, foregrounds exceed the signal out to ~ 380 ns which for a 14.6 m baseline is ≈ 330 ns beyond the horizon. At all other redshifts, the leakage due to reflections only extends to ≈ 250 ns beyond the horizon. To determine the impact of the observed reflections on HERA’s ability to constrain the astrophysics of reionization, we vary the cutoff in delay below which cosmological Fourier modes will be unobservable. To determine this cutoff, we consider three different scenarios for beam chromaticity that capture a range of possibilities informed by our simulations.

- **Optimistic: The cylinder resonance is mitigated.** In our most optimistic scenario, we assume that a more advanced feed design is able to mitigate the 150 MHz resonance in the cylinder while preserving polarization symmetry and effective area. In this case the foregrounds pass below the level of the signal at 250 ns beyond the horizon at all redshifts between $z = 7$ and 12, consistent with what is observed in the bands where the reflections are less severe or when they are not present at all. Since the achromatic beam stays above the foregrounds below ≈ 250 ns as well, this number is set by the width of the Blackman-Harris window function and more optimal estimators of the power spectrum may reduce it substantially (see Ali et al. (2015)). Hence, even our optimistic scenario is on the conservative side, making it more pessimistic than previous Fisher analyses (Pober et al., 2014; Liu et al., 2016; Liu and Parsons, 2016; Ewall-Wice et al., 2016b). In this scenario, we assume that only modes within 250 ns of the horizon are unobservable.

⁷<https://github.com/jpober/21cmSense>




Figure 6-15: The 1σ thermal noise levels achieved by HERA-350 at three different redshifts with (grey line) and without (black line) the presence of beam chromaticity due to the chromaticity observed in this work. We compare these noise levels to the fiducial power spectrum signal (red line). We saw in Fig. 6-14 that with reflections, foregrounds exceed the signal level out to $k = 0.23 h\text{Mpc}^{-1}$ at $z = 8.5$ which we assume are unusable, forcing us to ignore modes out to 350 ns beyond the horizon, leading to the sensitivity projected in the red curve. The absence of these reflections allows us to work within 250 ns of the horizon (green curve), leading to an increase in sensitivity by a factor of ≈ 1.5 .




Figure 6-16: 95% confidence regions for reionization parameters assuming 1000 hours of observation on HERA-350 between the redshifts between 7.5 and 12. The presence of the reflections leads to an increase in the major axes of these confidence regions by a factor of one to two.




Figure 6-17: Fractional Errors on reionization and heating parameters as a function of maximal observed redshift out to the low end of HERA-350’s initial observing band at $z = 12$. The presence of strong reflections contained within a small subband at $z = 8.5$ has a minimal impact on our overall constraints on reionization parameters. If these reflections are not localized they can worsen our sensitivity to some parameters such as $T_{\text{vir}}^{\text{min}}$ by a factor of two.

- **Moderate: The Simulations accurately capture the chromaticity of the dish.** In this case, we assume that the impedance match is sufficient for the contamination from reflections to pass below the level of the signal at 250 ns beyond the horizon except at redshift 8.5 where the feed resonance causes foregrounds to extend above the signal to 350 ns beyond the horizon. In this scenario, we assume that modes within 250 ns of the horizon are unobservable except for redshift 8.5 where modes below 350 ns beyond the horizon are unobservable. This is the scenario that would follow directly from our simulation results.
- **Pessimistic: Structure is present in all sub-bands.** In this scenario, we assume that the resonance is still present in the neighborhood of 150 MHz and that a very poor impedance match (similar to the $500\ \Omega$ case in Fig. 6-11) is obtained, allowing for significant contamination from reflections over the entire band, beyond the resonance. For this case, structure would be present out to ≈ 350 ns over the entire 100-200 MHz frequency range covered by HERA. In this scenario, we assume that modes within 350 ns of the horizon are unobservable across all redshift intervals.

In Fig. 6-15, we compare the level of 1σ thermal noise for our optimistic and pessimistic scenarios to the amplitude of the 21 cm signal at several different redshifts. The smallest k modes lost to foreground contamination are the highest signal to noise measurements that HERA is expected to obtain. Their absence impacts sensitivity in two ways. By leading to a reduction in the maximal signal to noise ratio of ≈ 1.5 and a reduction in the total number of modes that the instrument is able to measure.

Folding our calculations of thermal noise and the derivatives of Δ^2 into equation 6.24 and inverting, we obtain the covariance matrix for model parameters. We show the 95% confidence ellipses for the reionization parameters in Fig. 6-16. The presence of the resonance within a limited sub-band about $z = 8.5$ leads to an almost negligible increase in the extent of our confidence intervals while its presence across the entire band causes the lengths and widths of our confidence ellipses to increase by

a factor of ≈ 2 . The diagonal elements of our covariance give us error bars on each parameter which we plot in Fig. 6-17. We see that similar to our confidence regions, the error bars on reionization parameters for our optimistic and moderate scenarios are nearly indistinguishable. In our worst case where the resonance is present across the entire band, we see an increase in our error bars by a factor of ≈ 2 .

The error bars on reionization parameters, even for our most optimistic model, are a factor of a few larger than the “moderate” errors predicted in previous works using likelihood analyses such as Pober et al. (2014), Greig and Mesinger (2015); Greig et al. (2016, 2015) Ewall-Wice et al. (2016b), Liu et al. (2016); Liu and Parsons (2016). There are several reasons for this. Firstly, in our most optimistic scenario, we assumed that foregrounds cause the signal to be inaccessible below 250 ns beyond the edge of the wedge which corresponds to $k_{\min} \approx 0.15 h\text{Mpc}^{-1}$ at $z = 8.5$ while in previous works a minimal comoving $k_{\min} \approx 0.1 h\text{Mpc}^{-1}$ was used. Secondly, previous studies assumed a fully illuminated HERA aperture, which for a 14 m diameter dish predicts an effective area of $\approx 155 \text{ m}^2$. Electromagnetic simulations and ORBCOMM mapping of the angular beam pattern of the HERA dish show that the effective area of the antenna element is actually $\approx 98 \text{ m}^2$ at 137 MHz which leads to an increase in the overall thermal noise levels by a factor of ≈ 1.5 (Neben et al., 2016).

Although conservative, our analysis shows that the level of the HERA antenna’s delay response is not an insurmountable obstacle for the 21 cm power spectrum measurement. Even if the additional systematics introduced by the feed resonance are present over the entire band, HERA will obtain a $\gtrsim 10\sigma$ detection of the power spectrum and be capable of establishing precision constraints on the properties of the sources that reionized the IGM.

6.6 Conclusions

In this paper, we have derived the impact of spectral structure from reflections in the analog signal chain of an interferometer on foreground contamination of the 21 cm signal. We have used simulations of electromagnetic waves incident from zenith on

the primary antenna element on HERA to determine the extent to which spectral structure due to reflections in the dish leak foregrounds into the EoR window. The results of our simulations of the dish’s voltage response are broadly consistent with reflectometry measurements (Patra et al., 2017) and can be summarized in the following points.

- Additional spectral structure, in excess of that observed in the antenna design of HERA’s predecessor, PAPER, is observed in our simulations. Because this structure is present in simulations without the dish and is dramatically reduced when the feed’s cylindrical enclosure is removed, we conclude that this structure arises primarily from a resonance within the feed cylinder. On the other hand, contamination from reflections between the feed and the parabolic dish, which are present outside of the 150 MHz resonance, are relatively minor. Such reflections were perceived as a significant risk to the use of dish antennas in EoR studies and our simulations show that their impact on array sensitivity is small and 21 cm experiments can be economically scaled up through the use of dishes. Further adjustments to the feed design aim to remove the cylindrical resonance while preserving collecting area and polarization match.
- The resonance which is responsible for the majority of contamination has a finite width of ≈ 10 MHz of the HERA band. Because estimates of the power spectrum are obtained from sub-intervals of ≈ 10 MHz, the structure that we have simulated will only impact a single $\Delta z = 0.5$ redshift interval.
- Combining simulations of the HERA antenna’s delay-response with the foreground model of Thyagarajan et al. (2016), we find that a resonance in the feed extends foregrounds above the level of the cosmological signal to ≈ 350 ns beyond the horizon while without the resonance (but with reflections), foregrounds extend above the signal to ≈ 250 ns beyond the horizon (Fig. 6-14). These forecasts are conservative in that we do not attempt to inversely weight the foregrounds by their covariances or apply additional mitigation algorithms such as delay rate filtering.

- If beam chromaticity is contained within a 10 MHz subband around 150 MHz, then the overall constraints that HERA will be able to place on the astrophysics of reionization are minimally impacted. If the resonance were to extend across the entire band, our constraints on reionization parameters will suffer a two-fold increase in uncertainty but still remain on the order of 10%.

Our results are very encouraging since they show that the HERA design is capable of isolating the foregrounds to sufficiently low delays to make a high significance detection of the 21 cm power spectrum during the EoR. Since reflections are not a major source of spectral structure, feeds suspended over dishes are a viable strategy for increasing the collecting area of HI interferometers at a modest cost. Our results also emphasize the importance of mitigating any spectral structure in the analog signal chain. The contamination from the feed resonance masking our signal over some redshifts is mathematically identical to Fourier modes introduced by any instrumental spectral structure. Hence it is important that reflections be stringently suppressed at all steps in the signal chain beyond the dish to ensure that bright foregrounds fall below the EoR signal. A systematic definition of specifications on the spectral structure necessary to observe the EoR power spectrum at different k values is presented in Thyagarajan et al. (2016).

Chapter 7

The Impact of Modeling Errors on Interferometer Calibration for 21 cm Power Spectra

The content of this chapter was submitted the *Monthly Notices of the Royal Astronomical Society* in October 2017. A preprint has been published online as Ewall-Wice, A., Dillon, J. S., Liu, A. and Hewitt, J. (2016), ‘The Impact of Modeling Errors on Interferometer Calibration for 21 cm Power Spectra’, *ArXiv e-prints 1610.02689*.

7.1 Introduction

Observations of redshifted 21 cm emission are poised to unveil the properties of the earliest luminous sources in the universe, their impact on the global state of the intergalactic medium and how they affected the subsequent generations of stars and galaxies (see McQuinn 2015; Furlanetto 2016 for recent reviews).

One approach to detecting the cosmological 21 cm signal is to measure the fluctuations in the brightness temperature which can be mapped tomographically with a radio interferometer. To enhance the significance of a detection, most experiments are attempting to measure the spherically averaged power spectrum of these fluctuations. The mitigation of foregrounds that are four to five orders of magnitude brighter than

the signal itself is a central challenge that 21 cm experiments must overcome but is greatly aided by the spectral smoothness of these foregrounds (Di Matteo et al., 2002; Oh and Mack, 2003; Morales and Hewitt, 2004; Zaldarriaga et al., 2004). An interferometer measures the brightness distribution on the sky by cross correlating the outputs from many pairs of antennas. Flat-spectrum radio waves from a single point source, at a given time of observation, appear at a fixed time delay in the correlation between two antennas. Since the delay between two correlated antenna outputs is the Fourier dual to frequency, each fixed-delay source introduces a sinusoidal ripple as a function of frequency with a period that is inversely proportional to the difference of the arrival times of that source at the two correlated antennas. This sinusoid in frequency will correspond to a single comoving cosmological mode. In the absence of reflections, the maximal delay between signals arriving from a source on the sky (corresponding to the maximal line-of-sight (LoS) cosmological Fourier mode that is contaminated) occurs when the source is located along the separation of the antennas, at the horizon. Hence, as viewed by an interferometer, the spectrally smooth foregrounds are naturally contained within a region of Fourier space known as the *wedge* (Datta et al., 2010; Vedantham et al., 2012; Parsons et al., 2012b; Morales et al., 2012; Thyagarajan et al., 2013; Liu et al., 2014a,b) which is given by the horizon delay for each baseline separation and increases with that separation.

It is also possible for the signal chain of the instrument to imprint spectral structure into the measured visibilities. For example, a reflection within the signal path can delay the correlated signal. Hence, longer delays in the signal path contaminate finer frequency scales and are capable of leaking significant power outside of the wedge (Ewall-Wice et al., 2016d,a; Beardsley et al., 2016). Digital artifacts can also introduce fine spectral features such as those introduced by the polyphase filter bank on the Murchison Widefield Array (MWA) (Offringa et al., 2016). For each delay that is contaminated by structure in the antenna gains, an attenuated copy of the foregrounds, which are $\sim 10^4 - 10^5$ times larger than the signal, is introduced. Using foreground simulations, Thyagarajan et al. (2016) establish that in order to avoid contaminating the comoving LoS scales of several $h^{-1}\text{Mpc}$ or smaller, which are tar-

geted by 21 cm experiments, instrumental chromaticity beyond a 250 ns delay must be suppressed to the ≈ -50 dB level. Thus (a), the smoothness of the instrumental gain must meet this specification, or (b), calibration methods must be capable of suppressing any instrumental spectral structure to be within these limits.

Interferometric experiments have taken several distinct approaches to calibrating out instrumental spectral structure. Experiments focusing on imaging, such as the Murchison Widefield Array (MWA) (Tingay et al., 2013a), the Low Frequency Array (LOFAR) (van Haarlem et al., 2013), early deployments of the Precision Array for Probing the Epoch of Reionization (PAPER) (Jacobs et al., 2011, 2013; Kohn et al., 2016) and the Giant Metrewave telescope (GMRT) (Paciga et al., 2013)¹ calibrate their gains on a model of the sky that is usually iteratively improved with self-calibration (where observed sources are fed into an updated sky-model which is used to obtain more accurate gain solutions). Pipelines such as the MWA’s real time system (RTS), (Mitchell et al., 2008), Fast Holographic Deconvolution (FHD) (Sullivan et al., 2012), and `sageCAL` (Kazemi et al., 2011; Kazemi and Yatawatta, 2013; Kazemi et al., 2013), rely on the modeling approach which we refer to as sky-based calibration. An alternative route is to constrain the instrumental gains using many redundant measurements of the same visibility (Wieringa, 1992; Liu et al., 2010; Zheng et al., 2014). This strategy was implemented by the MIT EoR (MITEoR) array (Zheng et al., 2014, 2017b,a), the latest configurations of PAPER (Parsons et al., 2014; Ali et al., 2015), and the Hydrogen Epoch of Reionization Array (HERA) (DeBoer et al., 2016; Dillon and Parsons, 2016) that is now being commissioned in South Africa². Finally, *in-situ* calibration can be obtained using the injection of known signals (Patra et al., 2015). The Canadian Hydrogen Intensity Mapping Experiment (CHIME) (Newburgh et al., 2014) is employing a combination of redundant calibration, signal injection, and

¹In the GMRT’s case, a single, well known, pulsar is used while the rest of the sky is eliminated by difference time-steps that correspond to the pulsar’s “on” and “off” states.

²HERA is designed to be fully redundantly calibratable but it is useful to assess the performance of sky-based calibration as an alternative with potentially different systematics. Since redundant calibration does not rely as much on a model of the sky, there exists the possibility of this array outperforming any of the predictions in this paper.

pulsar holography to correct for instrumental gains.

Recent analyses of MWA data, using sky-based calibration have been contaminated by intrinsic chromaticity in the signal chain at the $\lesssim -20$ dB level (Dillon et al., 2015a; Ewall-Wice et al., 2016d; Jacobs et al., 2016; Beardsley et al., 2016) out to less than a delay of 2×10^3 ns, arising from a combination of reflections in the beam-former to receiver cables and digital artifacts. Increasing the frequency degrees of freedom within sky-based calibration is a potential solution as the gains are permitted to absorb fine-scale instrumental frequency structure at high delays (Offringa et al., 2016) and improvements in features such as cable reflections were noted in power spectra calibrated with additional parameters (Trott et al., 2016).

While calibration solutions with fine frequency degrees of freedom are able to model the detrimental spectral features in an instrumental bandpass, they are susceptible to absorbing the imperfections in any sky-model used for calibration. Naively, errors in a smooth foreground model should not impart spectrally complex errors into a gain solution. However, because every gain participates in many baselines with varying lengths and (due to the wedge) intrinsic chromaticities, calibration can imprint the frequency-dependent errors of the longest baselines in which an antenna participates into its gain solution. The application of this gain solution on the short baselines that the antenna participates in will mix contamination from long to short baselines, potentially contaminating the EoR window. Recent studies by Barry et al. (2016) (henceforth B16) and Patil et al. (2016) have demonstrated the existence of these errors in simulations of the special cases of the MWA and LOFAR with specific point source realizations. It has not yet been established how these errors scale with the properties of the instrument and the source catalog and whether they will pose a fundamental limitation to upcoming 21 cm experiments that expect to rely on sky-based calibration such as the Square Kilometre Array (SKA). Although B16 proposes a low-order-polynomial-based method to mitigate these effects, it generally relies upon intrinsically spectrally-smooth antenna bandpasses, which may not be the case for many interferometers.

In this paper, we employ linearized approximations of the calibration equations

developed by Wieringa (1992) (W92) and Liu et al. (2010) (L10) to investigate the amplitude of errors arising from incomplete calibration catalogs. Since these faint unmodeled sources can be described statistically (Liu and Tegmark, 2011, 2012; Trott et al., 2012; Dillon et al., 2013, 2014, 2015b; Trott et al., 2016), we will address the ensuing errors as a type of correlated noise which we will hereafter refer to as *modeling noise*. Unlike its thermal counterpart, modeling noise does not integrate down with observing time, biasing any power spectrum estimate. Since interleaved times in this noise are correlated, this bias cannot be eliminated (unlike thermal noise) by the technique of cross-multiplying interleaved time integrations (e.g. Dillon et al. 2014). We will derive equations describing the amplitude of modeling noise and its dependence on the properties of a radio interferometer such as the antenna count, distribution, and element size along with the depth of the calibration catalog. We use these equations to approximate the level of modeling noise in the existing MWA and LOFAR experiments (finding that our analytic results are in broad agreement with the simulation results in B16) along with the expected contamination in the upcoming instruments SKA-1 LOW and HERA. This contamination arises fundamentally from the chromaticity on long baselines, hence it can be eliminated by down-weighting long baselines in calibration, a strategy that we develop and verify in this paper.

We take an analytic approach in order to illuminate the origins of modeling noise in 21 cm power spectrum measurements and guide future array design and calibration strategies. For analytic tractability, we make a number of assumptions, which we attempt to describe clearly in the text, but do not necessarily hold for all observing scenarios. Thus, our quantitative results should be understood as accurate only to within an order of magnitude, illustrating how modeling noise scales with the properties of the sky catalog and instrumental parameters. Relaxing the assumptions in this paper for more accurate predictions is the subject of ongoing simulation work.

This paper is organized as follows. In § 7.2, we introduce our analytic framework, based on W92 and L10, for describing the impact of calibration modeling errors on the 21 cm power spectrum and discuss its dependence on array and catalog properties. In § 7.3 we apply this formalism to predict the amplitude of calibration errors relative to

21 cm fluctuations in current and upcoming experiments given our current knowledge of foregrounds and primary beams. In § 7.4, we explore a strategy for eliminating this noise through inverse baseline-length weighting. We conclude in § 7.5.

7.2 Formalism

In this paper, we model gain errors as a statistical noise arising from the myriad of faint unmodeled sources. Such sources are not precisely modeled in calibration.

Since baselines and visibilities are formed from two antennas, we index them with a greek index and we index antennas with lower-case latin indices. We will also sometimes explicitly write a baseline index as a 2-tuple of antenna indices (e.g. $\alpha = (i, j)$). We describe the residual, $c_\alpha(\nu)$, between the true visibility formed from antennas i and j , $v_\alpha^{\text{true}}(\nu)$, and the model visibility, $y_\alpha(\nu)$ as a random variable with a mean $\langle c_\alpha(\nu) \rangle$ and covariance $\mathbf{C}_{\alpha\beta}(\nu, \nu') \equiv \langle [c_\alpha(\nu) - \langle c_\alpha(\nu) \rangle] [c_\beta(\nu') - \langle c_\beta(\nu') \rangle]^* \rangle$. We assume that c_α is composed of the sum of the 21 cm signal, s_α , unmodeled foregrounds, r_α , and a component arising from thermal noise, n_α whose impact on calibration is explored in Trott and Wayth (2016). The true visibility is the sum between the modeled and unmodeled component.

$$v_\alpha^{\text{true}} = y_\alpha + c_\alpha = y_\alpha + r_\alpha + n_\alpha + s_\alpha \quad (7.1)$$

Each unmodeled component is statistically independent which means that

$$\text{Cov}[\mathbf{c}, \mathbf{c}^\dagger] \equiv \mathbf{C} = \mathbf{R} + \mathbf{N} + \mathbf{S}, \quad (7.2)$$

where \dagger denotes the conjugate transpose of a vector or matrix. In § 7.2.1 we discuss expressions for the amplitude and frequency coherence of r_α in terms of a parameterized point source population and diffuse galactic emission. This “noise” will be imprinted on the calibration solutions in a way that, for sufficiently small errors, is analytically tractable and can be described using the matrix formalism of W92 and L10, which we overview in § 7.2.2. We derive expressions for the impact of these

errors on the 21 cm power spectrum in § 7.2.3 and the degree to which each visibility covariance contributes in 7.2.4. Using the expressions we derive, we discuss the scaling of modeling noise with the properties of the source-catalog and array in § 7.3.4. Since both the 21 cm signal and thermal noise terms are already well considered in the literature (W92, L10, Trott and Wayth 2016), we will focus on the contribution from r_α .

7.2.1 The Statistics of Unmodeled Source Visibilities

Extensive work exists on statistical models of faint point sources in the power spectrum (e.g. Wang et al. 2006; Liu and Tegmark 2011; Trott et al. 2012; Dillon et al. 2013, 2015a; Trott et al. 2016) and we take an approach similar to these papers and assume the sources have uniform spectral structure (described by a single power law) that can be factored out of the visibilities and is far less significant than the frequency dependence introduced by the interferometric point-spread function. We now give an overview of our characterization of the unmodeled point sources along with the diffuse emission from the Galaxy. Since residual Galactic emission is, for the most part, uncorrelated with residual point-source emission, the covariance of unmodeled emission on each baseline is given by the sum of the covariance of each source,

$$\mathbf{R} = \mathbf{R}^P + \mathbf{R}^G, \quad (7.3)$$

where \mathbf{R}^P is the covariance due to unmodeled point sources and \mathbf{R}^G is the covariance of Galactic emission. We now describe our model of the covariances for these two emission sources.

7.2.1.1 Unmodeled Point Sources

With the MWA, point sources are completely sampled down to $S_{\min} \approx 50 - 80$ mJy (Carroll et al., 2016; Hurley-Walker et al., 2017; Line et al., 2016) and on LOFAR down to the $S_{\min} \approx 0.1$ mJy level (Williams et al., 2016) within the primary beam. We represent these sources with an achromatic version of the model from Liu and

Tegmark (2011). At these faint fluxes, the sources are isotropically distributed and the number of sources with fluxes between S and $S + dS$ within an infinitesimal solid angle $d\Omega$ is well described by a random Poisson process with a power law mean (Di Matteo et al., 2002).

$$\frac{d^2 N}{dS d\Omega} = k \begin{cases} \left(\frac{S}{S_*}\right)^{-\gamma_1} & S \leq S_* \\ \left(\frac{S}{S_*}\right)^{-\gamma_2} & S > S_* \end{cases}, \quad (7.4)$$

where $k = 4000 \text{Jy}^{-1} \text{sr}^{-1}$, $\gamma_1 = 1.75$, $\gamma_2 = 2.5$, and $S_* = 0.88 \text{Jy}$.

Consider a visibility, v_α , formed by antennas i and j , that are separated by baseline \mathbf{b}_α . The covariance between two baselines $v_\alpha(\nu)$ and $v_\beta(\nu')$ at two frequencies, ν and ν' , assuming un-clustered and flat-spectrum sources is

$$\begin{aligned} R_{\alpha\beta}^P(\nu, \nu') &= \int_0^{S_{\min}} dS \frac{dN}{dS d\Omega} S^2 \int d\Omega |A(\hat{\mathbf{s}})|^2 e^{-2\pi i \hat{\mathbf{s}} \cdot (\mathbf{b}_\alpha \nu - \mathbf{b}_\beta \nu')/c} \\ &= \sigma_r^2(S_{\min}) \int d\Omega |A(\hat{\mathbf{s}})|^2 e^{-2\pi i \hat{\mathbf{s}} \cdot (\mathbf{b}_\alpha \nu - \mathbf{b}_\beta \nu')/c}, \end{aligned} \quad (7.5)$$

where $A(\hat{\mathbf{s}})$ is the primary beam of each antenna which we assume are identical. The Fourier convolution theorem tells us that the last integral in equation 7.5 is equal to the convolution of the Fourier transform of the beam with itself evaluated at $(\mathbf{b}_\alpha \nu/c - \mathbf{b}_\beta \nu'/c)$ in the uv plane. This quantity falls to zero when $|\mathbf{b}_\alpha \nu/c - \mathbf{b}_\beta \nu'/c|$ is larger than the diameter of the antenna aperture. Thus as long as two baselines are separated by a distance greater than the antenna aperture diameter, $R_{\alpha\neq\beta}^P \approx 0$. We may therefore ignore off diagonal terms in the residual covariance matrix for minimally redundant arrays. It turns out that the diagonal covariance assumption gives similar results, even for maximally redundant arrays (see Appendix 7.A).

7.2.1.2 Diffuse Galactic Emission

Diffuse Galactic emission is correlated on large angular scales. We may construct a simple model of this emission using the same steps we used to obtain equation ?? and assuming that uv power spectrum of the diffuse emission does not evolve significantly over an antenna footprint. Under these assumptions, one can show that the covariance

between two visibilities from diffuse emission is

$$R_{\alpha\beta}^G(\nu, \nu') \approx P_G(\mathbf{b}_\alpha/\lambda_0) \int d\Omega |A(\hat{\mathbf{s}})|^2 e^{-2\pi i \hat{\mathbf{s}} \cdot (\mathbf{b}_\alpha \nu - \mathbf{b}_\beta \nu')/c} \quad (7.6)$$

where $P_G(\mathbf{u})$ is the power spectrum of diffuse galactic emission in the uv plane and λ_0 is the wavelength of the center of the interferometer's band. To model $P_G(\mathbf{u})$ we use an empirical power law fit to the two dimensional power spectrum of a desourced and destriped (Remazeilles et al., 2015) Galactic emission map (Haslam et al., 1982) centered at RA=60°, DEC=-30° and scaled from 408 to 150 MHz using a frequency power law with a spectral index of -0.6 (Rogers and Bowman, 2008; Fixsen et al., 2011)³. We find that the angular power spectrum of galactic emission at 150 MHz is well modeled by a power law in $u = |\mathbf{b}|/\lambda_0$, $P_G(\mathbf{u}) = 6 \times 10^{11} u^{-5.7} \text{Jy}^2 \text{Sr}^{-1}$. Throughout this paper, we will assume that the model used for calibration completely ignores diffuse emission so that all of $R_{\alpha\beta}^G$ is included in the co-variance of residual visibilities.

7.2.2 Frequency Domain Calibration Errors

So far we have a model of the discrepancies between true and modeled interferometer visibilities. Given this model, what are the statistics of the errors in our frequency dependent gain solutions? Our goal in this subsection is to derive the covariances of errors in gain parameters in terms of the covariances of the visibility residuals discussed in § 7.2.1.

We will start by writing down the system of equations that calibration algorithms attempt to solve, and, following W92 and L10, we will reduce this system to a set of linear equations that are valid in the regime of small calibration errors which is the case for errors generated by the faintest sources on the sky. This approximation holds when the gains are nearly correct after large gain variations are removed by a first iteration of calibration using a reasonably accurate calibration catalog. Writing down these systems in matrix form, the covariances of the least-squares solutions for these linear systems are readily obtained in the same manner as W92 and L10.

³This power law is for spectral radiance. For brightness temperature, the spectral index is -2.6.

We start by writing down the equations that calibration must solve. In line with the notation of L10, we parameterize the small gain and phase of the i^{th} antenna after rough calibration as the exponent of a complex number,

$$g_i(\nu) = e^{\eta_i(\nu) + i\phi_i(\nu)} \approx 1 + \eta_i(\nu) + i\phi_i(\nu), \quad (7.7)$$

where η_i is the amplitude of the gain and ϕ_i is the phase. In calibration, one attempts to solve the set of equations

$$g_i(\nu)g_j^*(\nu)v_{ij}^{\text{true}}(\nu) = v_{ij}^{\text{meas}}(\nu) \quad (7.8)$$

where $v_{ij}^{\text{meas}}(\nu)$ is the measured visibility. If we divide by y_{ij} on both sides, we have

$$g_i g_j^* \left(1 + \frac{c_{ij}}{y_{ij}} \right) = \frac{v_{ij}^{\text{meas}}}{y_{ij}}. \quad (7.9)$$

Recall that c_{ij} represents the sum of unmodeled components of a visibility (equation 7.1) while y_{ij} represents the modeled component. For analytic tractability, we will linearize these equations by working to first order in c_{ij}/y_{ij} , η_i , and ϕ_i . With this approximation, equation 7.9 becomes

$$\frac{v_{ij}^{\text{meas}}}{y_{ij}} \approx (1 + \eta_i + i\phi_i)(1 + \eta_j - i\phi_j) \left(1 + \frac{c_{ij}}{y_{ij}} \right) \quad (7.10)$$

$$\approx 1 + \eta_i + \eta_j + i\phi_i - i\phi_j + \frac{c_{ij}}{y_{ij}} \quad (7.11)$$

Separating the real and imaginary parts gives us two systems of linear equations;

$$\eta_i + \eta_j + \text{Re} \left(\frac{c_{ij}}{y_{ij}} \right) \approx \text{Re} \frac{v_{ij}^{\text{meas}}}{y_{ij}} - 1 \equiv M_{ij}^{\eta} \quad (7.12)$$

and

$$\phi_i - \phi_j + \text{Im} \left(\frac{c_{ij}}{y_{ij}} \right) \approx \text{Im} \frac{v_{ij}^{\text{meas}}}{y_{ij}} \equiv M_{ij}^{\phi}. \quad (7.13)$$

Since residual foregrounds may be described statistically, we treat c_{ij} as a noise term in the same way that thermal noise is treated in L10. Unlike thermal noise, which

is typically uncorrelated in frequency and ideally has the same variance across baselines, the correlation properties of modeling noise are those of the unmodeled sources discussed in § 7.2.1. We can write the system of equations given by equation 7.12 in matrix form⁴,

$$\begin{pmatrix} M_{12}^\eta \\ M_{23}^\eta \\ \vdots \\ M_{N-1N}^\eta \\ M_{13}^\eta \\ \vdots \\ M_{N-2N}^\eta \\ \vdots \\ M_{1N}^\eta \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & \dots & 0 \\ 0 & 1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ 1 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & \dots & 1 \end{pmatrix} \begin{pmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \\ \vdots \\ \eta_N \end{pmatrix} + \text{Re} \left(\frac{\mathbf{c}}{\mathbf{y}} \right), \quad (7.14)$$

which which we write more compactly as $\mathbf{M}^\eta = \mathbf{A}\boldsymbol{\eta} + \text{Re}(\mathbf{c}/\mathbf{y})$. The same can be done for equation 7.13.

$$\begin{pmatrix} M_{12}^\phi \\ M_{23}^\phi \\ \vdots \\ M_{N-1N}^\phi \\ M_{13}^\phi \\ \vdots \\ M_{N-2N}^\phi \\ \vdots \\ M_{1N}^\phi \\ 0 \end{pmatrix} = \begin{pmatrix} 1 & -1 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & -1 \\ 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & \dots & -1 \\ 1 & 1 & \dots & 1 \end{pmatrix} \begin{pmatrix} \phi_1 \\ \phi_2 \\ \vdots \\ \phi_N \end{pmatrix} + \text{Im} \left(\frac{\mathbf{c}}{\mathbf{y}} \right), \quad (7.15)$$

⁴The system of equations used here only attempts to solve for the gains. In redundant calibration, the number of unique true visibilities is reduced to a point where one can also solve for them as-well. This leads to different forms for the matrix equations (see W92 and L10 for examples).

where the last row in the matrix arises from imposing the constraint (L10) that $\sum_j \phi_j = 0$ ⁵. We write the imaginary equation as $\mathbf{M}^\phi = \mathbf{B}\phi + \text{Im}(\mathbf{c}/\mathbf{y})$.

Given a model and measurements of \mathbf{M}^ϕ and \mathbf{M}^η , a least squares estimator that applies weights of \mathbf{W} to each measurement will arrive at solutions for $\boldsymbol{\eta}$ and ϕ given by

$$\hat{\boldsymbol{\eta}} = (\mathbf{A}^\top \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{W} \mathbf{M}^\eta \equiv \boldsymbol{\Lambda} \mathbf{M}^\eta \quad (7.16)$$

and

$$\hat{\phi} = (\mathbf{B}^\top \mathbf{W} \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{W} \mathbf{M}^\phi \equiv \boldsymbol{\Psi} \mathbf{M}^\phi. \quad (7.17)$$

We emphasize that $\hat{\boldsymbol{\eta}}$ and $\hat{\phi}$ are estimates of the true values, $\boldsymbol{\eta}$ and ϕ . The covariance of these estimates,

$$\mathbf{C}_\eta(\nu, \nu') = \langle \hat{\boldsymbol{\eta}}(\nu) \hat{\boldsymbol{\eta}}^\top(\nu') \rangle - \langle \hat{\boldsymbol{\eta}}(\nu) \rangle \langle \hat{\boldsymbol{\eta}}^\top(\nu') \rangle \quad (7.18)$$

$$\mathbf{C}_\phi(\nu, \nu') = \langle \hat{\phi}(\nu) \hat{\phi}^\top(\nu') \rangle - \langle \hat{\phi}(\nu) \rangle \langle \hat{\phi}^\top(\nu') \rangle, \quad (7.19)$$

is given by

$$\mathbf{C}_\eta(\nu, \nu') = \boldsymbol{\Lambda} \text{Cov} \left[\text{Re} \left(\frac{\mathbf{c}}{\mathbf{y}} \right), \text{Re} \left(\frac{\mathbf{c}}{\mathbf{y}} \right)^\top \right] \boldsymbol{\Lambda}^\top \quad (7.20)$$

$$\mathbf{C}_\phi(\nu, \nu') = \boldsymbol{\Psi} \text{Cov} \left[\text{Im} \left(\frac{\mathbf{c}}{\mathbf{y}} \right), \text{Im} \left(\frac{\mathbf{c}}{\mathbf{y}} \right)^\top \right] \boldsymbol{\Psi}^\top \quad (7.21)$$

Thus, we have arrived at expressions for the covariances of errors in the gain parameters in terms of the covariances of the real and imaginary components of the unmodeled visibilities. Equations 7.20 and 7.21 show that the covariance of any given gain solution is the linear combination of the covariances of every visibility in the array. Thus, the application of a gain solution (derived from an incomplete sky model) to a short baseline introduces the fine-frequency errors from long baselines.

⁵The arbitrary phase reference is often set in sky-based calibration by defining the phases as the differences between each antenna phase and that of an arbitrarily chosen reference antenna. This constraint can be written as, $\phi_{\text{ref}} = 0$ and would modify the last row of \mathbf{B} to be zero except for the index of the reference antenna (rather than all ones as we have written it). While choosing the reference antenna form of the phase constraint affects the details of some of the expressions in this paper, it results in the same scaling relationships and has a negligible effect on quantitative results.

Our next step is to determine the impact of this leakage on the power spectrum.

7.2.3 The Impact of Gain Errors on the 21 cm Power Spectrum

We now propagate the frequency dependent errors in each gain solution into the delay power spectrum. Calibration gives us an estimate of the gains,

$$\widehat{g}_i = e^{\widehat{\eta}_i + i\widehat{\phi}_i} \approx 1 + \widehat{\eta}_i + \widehat{\phi}_i. \quad (7.22)$$

whose deviations from the true gains (η_i and ϕ_i) have covariances given by equations 7.20 and 7.21. The corrected, model-subtracted visibilities obtained from calibration are given by

$$\begin{aligned} V_{ij} &= \frac{g_i g_j^*}{\widehat{g}_i \widehat{g}_j^*} (y_{ij} + c_{ij}) - y_{ij} \\ &\approx (y_{ij} + c_{ij}) \times \\ &\quad [1 + (\eta_i - \widehat{\eta}_i) + (\eta_j - \widehat{\eta}_j) + i(\phi_i - \widehat{\phi}_i) - i(\phi_j - \widehat{\phi}_j)] - y_{ij} \end{aligned} \quad (7.23)$$

The delay transform (Parsons et al., 2012b) is a popular estimate of the power spectrum in which visibilities are Fourier transformed from frequency into delay. Delay can be mapped approximately to Fourier modes along the LoS while the uv coordinates of the visibility can be mapped to Fourier modes perpendicular to the LoS. The delay-transform is given by,

$$\widetilde{V}_{ij}(\tau) = \int d\nu e^{2\pi i\nu\tau} V_{ij}(\nu), \quad (7.24)$$

which we can apply to the gain-corrected and foreground subtracted visibility in equation 7.23. Taking the delay transform of equation 7.23 and setting $\boldsymbol{\eta}' \equiv \widehat{\boldsymbol{\eta}} - \boldsymbol{\eta}$ and $\boldsymbol{\phi}' \equiv \widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}$ we have,

$$\begin{aligned} \widetilde{V}_{ij}(\tau) &\approx -y_{ij} \star (\widetilde{\eta}'_i + \widetilde{\eta}'_j + i\widetilde{\phi}'_i - i\widetilde{\phi}'_j) \\ &\quad + \widetilde{c}_{ij} - \widetilde{c}_{ij} \star (\widetilde{\eta}'_i + \widetilde{\eta}'_j + i\widetilde{\phi}'_i - \widetilde{\phi}'_j), \end{aligned} \quad (7.25)$$

where \star denotes a convolution in delay space. For the sake of analytic tractability, we will ignore the chromaticity of y_α and set all $y_\alpha = S_0$, essentially assuming that that the modeled visibilities are dominated by a single source near the phase center which exceeds the flux of all other sources by a factor of several. Even with chromatic y_α , per-frequency inverse covariance weighting, which multiplies each α^{th} weight by $|y_\alpha|^2$ (L10) removes some of this structure. In Appendix 7.B we explore the impact of relaxing this assumption and find that our achromatic y_{ij} model predicts the LoS wave numbers at which the modeling noise drops below the 21 cm power spectrum to within $\approx 10\%$ of what we find with chromatic y_{ij} s obtained from a realistic sky model. Still, this dramatic assumption limits the accuracy of our specific quantitative predictions and we are exploring its impact in full calibration simulations.

The cosmological 21 cm power spectrum, $P(\mathbf{k})$, is well approximated by the mean amplitude square of the delay-transformed visibility multiplied by linear factors given in Parsons et al. (2012a)

$$P(\mathbf{k}) \approx \left(\frac{c^2}{2k_B^2 \nu_0^2} \right)^2 \frac{X^2(\nu_0)Y(\nu_0)}{B_{pp}\Omega_{pp}} \langle |\tilde{V}(\mathbf{u}, \eta)|^2 \rangle \quad (7.26)$$

where $\Omega_{pp} = \int d\Omega |A(\hat{\mathbf{s}})|^2$ and $B_{pp} = \int df |B(\nu)|^2$ are respectively the integrals of the squares of the beam and bandpass, ν_0 is the center frequency of the observation, k_B is the Boltzmann constant, and (X, Y) are multiplicative factors converting between native interferometry coordinates and comoving cosmological coordinates, $2\pi(u, v, \eta) = (Xk_x, Xk_y, Yk_z)$. In line with Parsons et al. (2012a), η is used to denote the Fourier dual to frequency at fixed $|\mathbf{u}|$ and τ to denote the frequency Fourier transform of a visibility which integrates over a slanted line in u - ν space. While this slanted integral introduces non-negligible mode-mixing (namely the wedge), it is a decent approximation for the range of η s probed by current and next-generation experiments.

Therefore, we can estimate the power spectrum from calibration-modeling errors by cross-multiplying \tilde{V}_α with its complex conjugate. If we denote the expectation

value

$$\langle \tilde{V}_\alpha(\tau) \tilde{V}_\beta^*(\tau') \rangle \equiv P_{\alpha\beta}(\tau, \tau'), \quad (7.27)$$

the bias from visibility residuals, is equal to $P_{\alpha\alpha}(\tau, \tau)$ multiplied by the constant prefactors in equation 7.26. While we only need $P_{\alpha\alpha}(\tau, \tau)$ for the bias, we will need off diagonal terms to calculate the variances of binned and averaged power spectrum estimates. We first write down $P_{\alpha\beta}(\tau, \tau')$ to second order in \mathbf{c}/\mathbf{y} with baseline α formed from antennas i and j and baseline β formed from antennas m and n ,

$$\begin{aligned} & \langle \tilde{V}_\alpha(\tau) \tilde{V}_\beta^*(\tau') \rangle \equiv P_{\alpha\beta} \\ & \approx S_0^2 \left[\langle \tilde{\eta}'_i \tilde{\eta}'_{m*} \rangle + \langle \tilde{\eta}'_i \tilde{\eta}'_{n*} \rangle + \langle \tilde{\eta}'_j \tilde{\eta}'_{m*} \rangle + \langle \tilde{\eta}'_j \tilde{\eta}'_{n*} \rangle \right. \\ & \quad + \langle \tilde{\phi}'_i \tilde{\phi}'_{m*} \rangle - \langle \tilde{\phi}'_i \tilde{\phi}'_{n*} \rangle - \langle \tilde{\phi}'_j \tilde{\phi}'_{m*} \rangle + \langle \tilde{\phi}'_j \tilde{\phi}'_{n*} \rangle \\ & \quad - i \langle \tilde{\eta}'_i \tilde{\phi}'_{m*} \rangle + i \langle \tilde{\eta}'_i \tilde{\phi}'_{n*} \rangle - i \langle \tilde{\eta}'_j \tilde{\phi}'_{m*} \rangle + i \langle \tilde{\eta}'_j \tilde{\phi}'_{n*} \rangle \\ & \quad \left. - i \langle \tilde{\phi}'_i \tilde{\eta}'_{m*} \rangle + i \langle \tilde{\phi}'_i \tilde{\eta}'_{n*} \rangle - i \langle \tilde{\phi}'_j \tilde{\eta}'_{m*} \rangle - i \langle \tilde{\phi}'_j \tilde{\eta}'_{n*} \rangle \right] \\ & \quad - S_0 \left[\langle \tilde{\eta}'_i \tilde{c}_\beta^* \rangle - \langle \tilde{\eta}'_j \tilde{c}_\beta^* \rangle - i \langle \tilde{\phi}'_i \tilde{c}_\beta^* \rangle + i \langle \tilde{\phi}'_j \tilde{c}_\beta^* \rangle \right. \\ & \quad \left. - \langle \tilde{c}_\alpha \tilde{\eta}'_{m*} \rangle - \langle \tilde{c}_\alpha \tilde{\eta}'_{n*} \rangle + i \langle \tilde{c}_\alpha \tilde{\phi}'_{m*} \rangle - i \langle \tilde{c}_\alpha \tilde{\phi}'_{n*} \rangle \right] \\ & \quad + \langle \tilde{c}_\alpha \tilde{c}_\beta^* \rangle. \end{aligned} \quad (7.28)$$

For the sake of space, we do not explicitly write τ or τ' in every term but understand that every complex conjugated term in each product is a function of τ' and every non-conjugated term is a function of τ .

Equation 7.28 involves six types of terms; those involving cross-multiples of $\tilde{\eta}'$, cross-multiples of $\tilde{\phi}'$, cross-multiples of $\tilde{\eta}'$ and $\tilde{\phi}'$, products between $\tilde{\eta}$ and $\tilde{\phi}$ with \tilde{c} , and finally the covariances of the residuals themselves. In Appendix 7.C we obtain approximate expressions for each of the first five terms when the baseline separation

is longer than the antenna diameter,

$$\begin{aligned}
\langle \tilde{\eta}'_i(\tau) \tilde{\eta}'_{m*}(\tau') \rangle &\approx \int d\nu d\nu' e^{2\pi i(\nu\tau - \nu'\tau')} [\mathbf{C}_\eta(\nu, \nu')]_{im} \\
&\equiv \frac{S_0^{-2}}{2} [\boldsymbol{\Lambda} \tilde{\mathbf{C}}(\tau, \tau') \boldsymbol{\Lambda}^\top]_{im} \\
&= \frac{S_0^{-2}}{2} \Lambda_{i\gamma} \Lambda_{\delta m}^\top \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau')
\end{aligned} \tag{7.29}$$

$$\begin{aligned}
\langle \tilde{\phi}'_i(\tau) \tilde{\phi}'_{m*}(\tau') \rangle &\approx \int d\nu d\nu' e^{2\pi i(\nu\tau - \nu'\tau')} [\mathbf{C}_\phi(\nu, \nu')]_{im} \\
&= \frac{S_0^{-2}}{2} [\boldsymbol{\Psi} \tilde{\mathbf{C}}(\tau, \tau') \boldsymbol{\Psi}^\top]_{im} \\
&= \frac{S_0^{-2}}{2} \Psi_{i\gamma} \Psi_{\delta m}^\top \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau')
\end{aligned} \tag{7.30}$$

$$\begin{aligned}
\langle \tilde{c}_\alpha(\tau) \tilde{\eta}'_{m*}(\tau') \rangle &\approx \frac{1}{2} \Lambda_{m\gamma} \int d\nu d\nu' e^{2\pi i(\nu\tau - \nu'\tau')} [\mathbf{C}(\nu, \nu')]^{\gamma\alpha} \\
&= \frac{1}{2} \Lambda_{i\gamma} \tilde{\mathbf{C}}^{\gamma\alpha}(\tau, \tau')
\end{aligned} \tag{7.31}$$

$$\begin{aligned}
\langle \tilde{c}_\alpha(\tau) \tilde{\phi}'_{m*}(\tau') \rangle &\approx \frac{i}{2} \Psi_{m\gamma} \int d\nu d\nu' e^{2\pi i(\nu\tau - \nu'\tau')} [\mathbf{C}(\nu, \nu')]^{\gamma\alpha} \\
&= \frac{i}{2} \Psi_{m\gamma} \tilde{\mathbf{C}}^{\gamma\alpha}(\tau, \tau')
\end{aligned} \tag{7.32}$$

$$\langle \tilde{\eta}'_i(\tau) \tilde{\phi}'_{m*}(\tau') \rangle \approx 0, \tag{7.33}$$

where we used Einstein-notation with repeated raised and lowered indices to denote summation and have defined $\tilde{\mathbf{C}}(\tau, \tau')$ as the delay-transform of the \mathbf{C} matrix. $\tilde{\mathbf{C}}_{\alpha\beta}(\tau, \tau') \equiv \int d\nu d\nu' e^{-2\pi i(\tau\nu - \tau'\nu')} \mathbf{C}_{\alpha\beta}(\nu, \nu')$. We also denote the delay-transform of the \mathbf{N} , \mathbf{R} , and \mathbf{S} matrices in a similar way as $\tilde{\mathbf{N}}$, $\tilde{\mathbf{R}}$, and $\tilde{\mathbf{S}}$. The final term in equation 7.28 is simply the covariance matrix of the delay-transformed residual visibilities,

$\tilde{\mathbf{C}}_{\alpha\beta}(\tau, \tau')$. Using the above identities, we may write equation 7.28 as

$$\begin{aligned}
\mathbf{P}_{\alpha\beta}(\tau, \tau') &= \frac{1}{2} \left[\Lambda_{i\gamma} \Lambda_{\delta m}^\top + \Lambda_{i\gamma} \Lambda_{\delta n}^\top + \Lambda_{j\gamma} \Lambda_{\delta m}^\top + \Lambda_{j\gamma} \Lambda_{\delta n}^\top + \right. \\
&\quad \left. \Psi_{i\gamma} \Psi_{\delta m}^\top - \Psi_{i\gamma} \Psi_{\delta n}^\top - \Psi_{j\gamma} \Psi_{\delta m}^\top + \Psi_{j\gamma} \Psi_{\delta n}^\top \right] \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau') \\
&\quad - \frac{1}{2} (\Lambda_{i\gamma} + \Lambda_{j\gamma} - \Psi_{i\gamma} + \Psi_{j\gamma}) \tilde{\mathbf{C}}^{\gamma\beta}(\tau, \tau') \\
&\quad - \frac{1}{2} (\Lambda_{m\gamma} + \Lambda_{n\gamma} + \Psi_{m\gamma} - \Psi_{n\gamma}) \tilde{\mathbf{C}}^{\gamma\alpha}(\tau, \tau) \\
&\quad + \tilde{\mathbf{C}}_{\alpha\beta}(\tau, \tau'), \tag{7.34}
\end{aligned}$$

The power-spectrum bias in delay-transform estimates is given by $\mathbf{P}_{\alpha\alpha}(\tau, \tau)$ ($i = m$ and $j = n$),

$$\begin{aligned}
\mathbf{P}_{\alpha\alpha}(\tau, \tau) &= \frac{1}{2} (\Lambda_{i\gamma} \Lambda_{\delta i}^\top + 2\Lambda_{i\gamma} \Lambda_{\delta j}^\top + \Lambda_{j\gamma} \Lambda_{\delta j}^\top) \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau) \\
&\quad + \frac{1}{2} (\Psi_{i\gamma} \Psi_{\delta i}^\top - 2\Psi_{i\gamma} \Psi_{\delta j}^\top + \Psi_{j\gamma} \Psi_{\delta j}^\top) \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau) \\
&\quad - (\Lambda_{i\gamma} + \Lambda_{j\gamma}) \tilde{\mathbf{C}}^{\gamma\alpha}(\tau, \tau) + \tilde{\mathbf{C}}_{\alpha\alpha}(\tau, \tau). \tag{7.35}
\end{aligned}$$

Equations 7.35 and 7.34 show how calibration leaks unmodeled structure in every visibility, including the highly chromatic ones, into the power spectrum of otherwise smooth short baselines. The last term in equation 7.35 is the power spectrum of unmodeled foregrounds, noise, and the signal itself. Recall that since foregrounds are naturally contained within the horizon delay of b_α , it does not contribute power into the EoR window. The sums in the first two lines, on the other hand, mix the chromaticity of foregrounds on all baselines into the delay power spectrum of the α^{th} visibility. Baselines that are longer than b_α contribute emission at delays below their individual horizon-delays which are greater than the horizon delay of b_α , allowing for contamination of the EoR window.

Typically, an estimate of $\mathbf{P}_{\alpha\alpha}$ is obtained by cross-multiplying integration over independent time intervals and since noise within each interval is independent, $\mathbf{N}_{\alpha\beta}(\nu, t; \nu', t') = 0$ (Dillon et al., 2014) and we can ignore the thermal noise contribution to the bias given by equation 7.34. However, a subtlety introduced by calibration errors is that

if calibration solutions for the cross-multiplied visibilities are not derived from complementary time intervals, there will still exist a thermal noise bias arising from all but the last term in 7.34. This is the case in Dillon et al. (2015a), Ewall-Wice et al. (2016d), and Beardsley et al. (2016) where 0.5 s time-steps are used for interleaving visibilities but 112 s non-interleaved time-steps are used for calibration. This bias can also survive cross multiplying different redundant measurements of the same visibility as is done with PAPER (Parsons et al., 2014; Ali et al., 2015).

7.2.4 Calibration Bias for a Simplified Model

How do the covariances between pairs of visibilities contribute to the final power spectrum? We showed in § 7.2.1 that \mathbf{R} , the covariance matrix of un-modeled foreground visibilities, is well approximated as a diagonal in minimally redundant arrays. The same is true for thermal noise which arises from independent fluctuations at each antenna and the 21 cm signal. Thus, the first two lines of equation 7.35 are formed from the weighted sum of the variance of the $N_{\text{ant}}(N_{\text{ant}} - 1)/2$ visibilities (where N_{ant} is the number of antenna elements) with $\gamma = \delta$ where the weight of each variance is given by $\Lambda_{i\gamma}\Lambda_{\gamma j}^T$ and $\Psi_{i\gamma}\Psi_{\gamma j}^T$. These values depend crucially on our choice of visibility weighting, \mathbf{W} , but it is highly instructive to examine the case where \mathbf{W} is equal to the identity. In Appendix 7.D we use matrix algebra to that in the case of \mathbf{W} equal to the identity,

$$\Lambda_{i\gamma} = \begin{cases} \frac{1}{N_{\text{ant}}-1} & i \in \gamma \\ \frac{-1}{2(N_{\text{ant}}-1)(N_{\text{ant}}-2)} & i \notin \gamma \end{cases} \quad (7.36)$$

and

$$\Psi_{i\gamma} = \begin{cases} \frac{1}{N_{\text{ant}}} & \gamma = (i, \cdot) \\ -\frac{1}{N_{\text{ant}}} & \gamma = (\cdot, i) \\ 0 & i \notin \gamma \end{cases} \quad (7.37)$$

where we denote $\gamma = (i, \cdot)$ to denote any visibility with i as the non-conjugated antenna and $\gamma = (\cdot, i)$ to be any baseline with i as the conjugated antenna. Equation 7.36 makes intuitive sense if we recall that $\Lambda_{i\gamma}$ is the linear weight multiplied by

each M_γ^η that is summed to form the i^{th} gain solution. Inspecting equation 7.12, we see that summing all $N_{\text{ant}} - 1$ $\Lambda_{i\gamma} M_\gamma^\eta$, that antenna i participates in gives us

$$\sum_{\gamma \ni i} \Lambda_{i\gamma} M_\gamma^\eta = \eta_i + \frac{1}{N_{\text{ant}} - 1} \sum_{k \neq i} \eta_k \quad (7.38)$$

To remove the extra sum, and isolate η_i , we must subtract the sum all M^η that do not include the i^{th} antenna, divided by $N_{\text{ant}} - 1$. For each $k \neq i$, there are $N_{\text{ant}} - 2$ baselines that involve k but not i , so we must also divide each term by $N_{\text{ant}} - 2$. This gives us the weights for baselines not involving the i^{th} antenna in equation 7.36. We can apply similar logic to equation 7.37 by inspecting equations 7.13 and 7.37.

The weight with which the covariance between each pair of measurements contributes to the total covariance of $\hat{\boldsymbol{\eta}}$ and $\hat{\boldsymbol{\phi}}$ is just the product of the weight with which each measurement is linearly summed.

$$\Lambda_{i\gamma} \Lambda_{\delta j}^\top = \begin{cases} \frac{1}{(N_{\text{ant}} - 1)^2} & i \in \gamma \text{ and } j \in \delta \\ \frac{-1}{(N_{\text{ant}} - 2)(N_{\text{ant}} - 1)^2} & i \in \gamma \text{ and } j \notin \delta \\ \frac{-1}{(N_{\text{ant}} - 2)(N_{\text{ant}} - 1)^2} & i \notin \gamma \text{ and } j \in \delta \\ \frac{1}{(N_{\text{ant}} - 1)^2 (N_{\text{ant}} - 2)^2} & i \notin \gamma \text{ and } j \notin \delta \end{cases} \quad (7.39)$$

and

$$\Psi_{i\gamma} \Psi_{\delta j}^\top = \begin{cases} \frac{1}{N_{\text{ant}}^2} & \gamma = (i, \cdot), \delta = (j, \cdot) \text{ or } \gamma = (\cdot, i), \delta = (\cdot, j) \\ -\frac{1}{N_{\text{ant}}^2} & \gamma = (i, \cdot), \delta = (\cdot, j) \text{ or } \gamma = (\cdot, i), \delta = (j, \cdot) \\ 0 & \text{otherwise.} \end{cases} \quad (7.40)$$

For non-redundant arrays, $\mathbf{C}(\nu, \nu')$ is diagonal and we can focus on $\gamma = \delta$ terms. From equations 7.39 and 7.40 we see that when $i = j$, each visibility variance is weighted by $\sim N_{\text{ant}}^{-2}$ when $i \in \gamma$ and at most by N_{ant}^{-4} when $i \notin \gamma$. Since there are $\sim N_{\text{ant}}$ visibilities with antenna i and $\sim N_{\text{ant}}^2$ visibilities without antenna i , $i = j$ terms contributing to $\mathbf{P}_{\alpha\alpha}$ are given roughly by the average of visibility variances not involving i divided by $\sim N_{\text{ant}}^{-2}$ plus the average of visibility variances involving antenna

i divided by $\sim N_{\text{ant}}$.

$$\begin{aligned} \Lambda_{i\gamma} \Lambda_{\delta i}^T \tilde{\mathbf{C}}^{\gamma\delta}(\tau, \tau) &\approx \frac{1}{N_{\text{ant}}^2} \langle \tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) \rangle_{i \notin \delta} + \frac{1}{N_{\text{ant}}} \langle \tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) \rangle_{i \in \delta} \\ &\approx \frac{1}{N_{\text{ant}}} \langle \tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) \rangle_{i \in \delta} \end{aligned} \quad (7.41)$$

where the $\langle \rangle_{i \in \delta}$ indicate an average over the set of baselines that include antenna i and $\langle \rangle_{i \notin \delta}$ denotes an average over baselines that are not formed using antenna i . The same equation holds for the Ψ sums. Considering how elements of $\mathbf{C}(\nu, \nu')$ scale with N_{ant} , we can see the average of visibilities involving antenna i dominate $\mathbf{P}_{\alpha\alpha}$ by a factor of N_{ant} .

For $i \neq j$, there is exactly one visibility that involves both antennas and will be weighted at most by N_{ant}^{-2} . The $\sim N_{\text{ant}}$ visibilities formed from i XOR j are weighted by N_{ant}^{-3} and the $\sim N_{\text{ant}}^2$ visibilities that involve neither i nor j are weighted by N_{ant}^{-4} . Thus all terms with $i \neq j$ in equation 7.35 give contributions on the order of the average of the visibility variances divided by N_{ant}^2 .

It follows that if $\mathbf{C}(\nu, \nu')$ is diagonal for all ν, ν' and \mathbf{W} is equal to the identity, $i = j$ sums in equation 7.35 dominate by $\sim N_{\text{ant}}$ and are well approximated by the average visibility variance involving antenna i or j divided by N_{ant} . The overall level of foreground contamination from calibration errors therefore goes as N_{ant}^{-1} with the details of the extent in delay contamination depending on the antenna distribution and primary beam. Replacing each $i = j$ sum in equation 7.35 with an average over visibility covariances involving i and j and ignoring $i \neq j$ sums we arrive at an approximate formula that can be readily used to estimate $\mathbf{P}_{\alpha\alpha}$.

$$\begin{aligned} \mathbf{P}_{\alpha\alpha}(\tau, \tau) &\approx \frac{1}{N_{\text{ant}}} \left[\langle \tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) \rangle_{i \in \delta} + \langle \tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) \rangle_{j \in \delta} \right] \\ &\quad + \tilde{\mathbf{C}}_{\alpha\alpha}(\tau, \tau) \end{aligned} \quad (7.42)$$

The two assumptions going into this formula are that for each ν and ν' , $\mathbf{C}(\nu, \nu')$ is diagonal (minimal redundancy) and that \mathbf{W} is set equal to unity. Equation 7.42

illustrates how the bias of a power spectrum estimate depends on both the covariance of the individual baseline (the second term) and the covariances of the baselines that share common antennas. In other words, the measurement of the power spectrum for a given baseline and delay depends on both on the residual foregrounds, noise, and signal at that baseline and delay and, suppressed by a factor of N_{ant} , that of all other baselines at that delay that share an antenna with it.

7.3 Modeling Noise in Existing Arrays

Having developed our formalism in § 7.2, we may obtain order-of-magnitude estimates for the visibility modeling noise using equation 7.34 for four existing or planned arrays; LOFAR, MWA, HERA, and the re-baselined SKA-LOW. We discuss our models of each instrument (§ 7.3.1). We then determine the level of modeling noise in (§ 7.3.2) and the impact of beam modeling errors (§ 7.3.3). Equation 7.42 can be used to provide us with some intuition for how the properties of the noise scales with those of the array and catalog. In § 7.3.4, we discuss these scalings and to what extent they may be used to reduce the amplitude of modeling noise. In each simulation, we assume that the foreground model, used for calibration and subtraction, contains point sources modeled down to some minimal flux level S_{min} and that the true sky contains both the diffuse emission and all point sources.

7.3.1 Instrumental Models

For all arrays, we consider an Airy beam for an aperture with diameter d_{ant} ,

$$A(\hat{\mathbf{s}}) = A(\theta) = \left(2 \frac{J_1(\pi d_{\text{ant}} \cos \theta / \lambda_0)}{\pi d_{\text{ant}} \cos \theta \lambda_0} \right)^2, \quad (7.43)$$

where θ is the arc length from the beam pointing centre. An Airy beam has the virtue of a simple analytic expression that, unlike a Gaussian beam, exhibits realistic side-lobe structure which in turn affects foreground contamination near the edge of the wedge (Thyagarajan et al., 2015a; Pober et al., 2016).

Strictly speaking, the primary beam evolves with frequency; however, we find in numerical calculations that allowing for this variation has a negligible impact on our results. In order to expedite the computation of calibration noise, especially for the arrays with large numbers of antennas, such as HERA and the SKA-1, we also assume that \mathbf{R} is diagonal. This is clearly not the case for the highly redundant HERA layout but we find (Appendix 7.A) that this only impacts the amplitude of the modeling noise by a factor of order unity and has a negligible impact on which modes are contaminated. In all arrays, we assume that every baseline is given equal weighting of unity. Note that for the SKA and LOFAR, we do not explicitly include outrigger antennas in our model of calibration. Our models for each individual instrument are as follows.

- **The Murchison Widefield Array** For the MWA, we use the 128 tile layout described in Beardsley et al. (2012) and Tingay et al. (2013a). Antennas are modeled as 4 m diameter circular apertures. We assume a flux limit of 86 mJy, which is the limit for the array’s naturally weighted point spread function at 150 MHz and similar to limits obtained in Carroll et al. (2016). Other analyses have obtained complete samples down to 35 – 50 mJy (Hurley-Walker et al., 2017) but this order unity change in S_{\min} does not significantly impact the modeling noise level which scales as $\sim S_{\min}^{1.25}$ (see § 7.3.4). The deeper TIFR GMRT Sky Survey (TGSS) covers a significant portion of the MWA’s field of view and is complete down to 10 mJy (Intema et al., 2017). We therefore also consider an optimistic scenario where a deep TGSS catalog is used to calibrate the instrument.
- **The Low Frequency Array** We model LOFAR as the 48-element high band core described in van Haarlem et al. (2013), with 30 m diameter circular stations. The confusion limit for the naturally weighted core is ≈ 35 mJy at 150 MHz. However, the use of LOFAR’s extended baselines measures source catalogs that are complete down to $S_{\min} \approx 0.1$ mJy (Williams et al., 2016). While the (Williams et al., 2016) survey is over a $\approx 4^\circ$ field of view, the catalog

we consider here covers the entire sky. Such a catalog would involve accurately characterizing ~ 27 million sources and may not happen before the SKA but we consider it as a very optimistic bracket on LOFAR’s performance.

- **The Hydrogen Epoch of Reionization Array** For our model of HERA, we use the 331 element hexagonally packed core described in Pober et al. (2014). Each element is modeled as a 14 m diameter circular aperture. HERA is designed to be calibrated redundantly (Dillon and Parsons, 2016); hence, the power spectrum estimates it obtains will not directly be affected by the modeling errors we explore in this paper. We choose to include HERA in order to assess the performance of compact cores in sky-based calibration and to explore sky-based calibration as an alternative to redundant calibration. HERA’s confusion limit is $S_{\min} \approx 11$ Jy. However, the $\text{dec} \approx -30^\circ$ stripe that it will scan is also covered by the TGSS survey which is complete down to ≈ 10 mJy. We therefore also consider an optimistic scenario in which the TGSS catalog is used for calibration.
- **The Square Kilometre Array** We investigate the level of modeling noise in the SKA-1 Low design proposed in Dewdney et al. (2013) but scaled down to half of the described collecting area due to recent rebaselining. The array consists of 497, 30 m stations with a number density distributed as a Gaussian in radius where 75 % of antennas fall within 1 km of the center, corresponding to a standard deviation of $\sigma_{\text{ant}} \sim 600$ m. The confusion limit of the SKA’s core is ≈ 27 mJy, however the inclusion of extended baselines out to ≈ 100 km will bring the confusion limit at 150 MHz to ≈ 0.1 mJy (Prandoni and Seymour, 2015) which we also consider as an optimistic case.

7.3.2 Modeling Results

In Fig. 7-1, we plot cylindrically binned and averaged delay-transform power spectra of residual visibilities from unmodeled foregrounds, calculated using equation 7.35 for the MWA, LOFAR, HERA, and SKA-1. We explore two different S_{\min} values for each

array. As we might expect, the majority of residual power is contained within the wedge, arising from the last term in equation 7.35. This term is the power spectrum of the unmodeled residual sources and would exist in the absence of calibration errors. For the MWA, which has a smaller aperture, hence a wider primary beam, the wedge of unmodeled sources extends to larger k_{\parallel} values. Beyond the wedge extends the power spectrum of calibration errors which exist at the level of $10^6 - 10^8 h^3 \text{Mpc}^{-3} \text{mK}^2$; one to two orders of magnitude greater than the 21 cm signal. For the MWA, the level of contamination inside of the EoR window is within an order of magnitude of the simulated errors encountered in B16 who consider a calibration catalog that is incomplete to a similar depth. It is apparent that the calibration errors experience a sharp cutoff at the k_{\parallel} corresponding to the delay of the edge of the main lobe on the longest baselines of the array. A vertical stripe of additional contamination appears in the LOFAR plot at $k_{\perp} \approx 0.6 h \text{Mpc}^{-1}$ which corresponds to separation scale for the HBA antenna pairs. Since even the longest outriggers participate in a short baseline with this length, more significant contamination is introduced at the corresponding Fourier mode.

We also estimate the region of k -space in which the 21 cm signal will be accessible by computing the ratio between the 2d power spectrum of residual visibilities and a representative signal computed using 21cmFAST⁶ (Mesinger and Furlanetto, 2007; Mesinger et al., 2011). The reionization parameters are set to $T_{\text{vir}}^{\text{min}} = 2 \times 10^4 K$, $\zeta = 20$, and $R_{\text{mfp}} = 15 \text{Mpc}$, yielding a redshift of 50% reionization of ≈ 8.5 . For fiducial catalog limits, we see that the entire EoR window is unusable for LOFAR and the MWA while the SKA is only able to detect the signal at large $k \gtrsim 0.4 h \text{Mpc}^{-1}$. If LOFAR and the SKA use their extended baselines to obtain deep source catalogs and calibrate on these catalogs *with only their core antennas*, they will be able to isolate modeling errors to be contained primarily within the wedge. By calibrating on a deep 10 mJy catalog such as the TGSS, HERA can rely on traditional sky-based calibration as a potential alternative to its primary redundant strategy.

⁶http://homepage.sns.it/mesinger/DexM__21cmFAST.html




Figure 7-1: Top: The power spectrum residuals computed using equation 7.35 for the MWA, LOFAR, HERA, and the SKA-1 LOW designs with sources modeled down to the array confusion limit. Unmodeled foregrounds are contained within the wedge, which is demarked by dashed black lines at the first primary beam null and solid black lines at the horizon. The calibration errors introduced by these foregrounds bleed out of the wedge into the EoR window. The narrower central lobes (larger stations) employed by LOFAR and the SKA help to significantly reduce the leakage at large k_{\parallel} that exists for the MWA. Contours where the signal, from a 21cmFAST simulation, is equal to unity, five, and ten times the calibration noise are indicated by cyan, orange, and red lines respectively. Bottom: The same as the top for optimistic scenarios. The optimistic scenario for LOFAR and the SKA involves complete modeling of point sources down to 0.1 mJy using additional long baselines. For HERA and the MWA, the optimistic scenario assumes that the 10 mJy source catalog from the TGSS is used for calibration. If long baselines can faithfully model the sources down to 0.1 mJy, modeling noise does not appear to limit LOFAR and the SKA. Sky-based calibration with HERA is improved significantly by using a deep source catalog from a complementary array. The vertical stripe in the LOFAR figure at $k_{\perp} \sim 0.6 h\text{Mpc}^{-1}$ arises from the arrangement of the HBA antennas in short spaced pairs so that even the outrigger antennas, which are heavily contaminated, participate in a single short baseline.

7.3.3 The Impact of Primary Beam Modeling Errors

So far, we have assumed that the antenna primary beam is known perfectly. Here we examine the impact of an imperfectly modeled beam on calibration noise. It is worth noting that while we only examine the impact of beam errors on calibration, errors in beam modeling can affect other aspects of the analysis (for example the power spectrum normalization equation in equation 7.26). Recent in-situ measurements with Orbcomm satellites (Neben et al., 2015, 2016) indicate that electromagnetic modeling of instrumental primary beams may only be accurate to the 1% level within the central lobe and only to the 10% level within the side-lobes. Even if a complete model of the sky exists, systematic errors in the apparent flux of these sources will cause calibration errors similar to those encountered in § 7.3.1. We describe beam-modeling errors as an angle dependent function, $D(\hat{\mathbf{s}})$ that is added to the known component of the beam $B(\hat{\mathbf{s}})$.

$$A(\hat{\mathbf{s}}) = B(\hat{\mathbf{s}}) + D(\hat{\mathbf{s}}). \quad (7.44)$$

For the purposes of this section only, we take the optimistic case that we have a perfect external catalog and that all calibration modeling error comes from an incorrect model of the primary beam. A true visibility in the presence of these errors is

$$\begin{aligned} v_\alpha^{\text{true}} &= \int d\Omega [B(\hat{\mathbf{s}}) + D(\hat{\mathbf{s}})] I(\hat{\mathbf{s}}) e^{-2\pi i \mathbf{b}_\alpha \nu / c} \\ &= y_\alpha + \int d\Omega D(\hat{\mathbf{s}}) I(\hat{\mathbf{s}}) e^{-2\pi i \mathbf{b}_\alpha \nu / c}. \end{aligned} \quad (7.45)$$

Our new calibration residual, r_α takes on the form:

$$r_\alpha \rightarrow \int d\Omega D(\hat{\mathbf{s}}) I(\hat{\mathbf{s}}) e^{-2\pi i \mathbf{b}_\alpha \nu / c}. \quad (7.46)$$

This leads to a new form of $R_{\alpha\beta}$ as well:

$$R_{\alpha\beta} \rightarrow \int_0^{S_{\max}} dS \frac{d^2 N}{dS d\Omega} S^2 \int d\Omega |D(\hat{\mathbf{s}})|^2 e^{-2\pi i(b_\alpha \nu/c - b_\beta \nu'/c)}, \quad (7.47)$$

where S_{\max} is the flux of the highest flux source in the field of view which is obtained by setting the number of sources with intrinsic flux greater than S_{\max} equal to unity,

$$S_{\max} = S_* \left[k \int d\Omega A(\hat{\mathbf{s}}) \right]^{1/\gamma}. \quad (7.48)$$

Since the literature typically reports fractional errors in beam-modeling, we describe $D(\hat{\mathbf{s}})$ as the true beam $A(\hat{\mathbf{s}})$ multiplied by a fractional error function $D(\hat{\mathbf{s}}) = f(\hat{\mathbf{s}})A(\hat{\mathbf{s}})$ where we parameterize $f(\hat{\mathbf{s}})$ as the following piecewise function,

$$f(\hat{\mathbf{s}}) = \begin{cases} A [1 - (1 - e_z) \exp(-\cos^2 \theta / 2\sigma_e^2)] & |\cos \theta| < s_1 \\ A [1 - (1 - e_z) \exp(-s_1^2 / 2\sigma_e^2)] & |\cos \theta| \geq s_1, \end{cases} \quad (7.49)$$

where Ae_z is the fractional beam-modeling error at the pointing centre, s_1 is the angular distance of the pointing center to the first side-lobe and A, σ_e may be adjusted to give different fractional modeling errors in the side-lobes. This function allows us to assign an arbitrary modeling uncertainty to the zenith and side-lobes. We compute the level of beam-modeling noise in the 21 cm power spectrum for two different scenarios, one in which the beam is known to 1% at zenith and 10% in the side-lobes, which is consistent with the precision reported in Neben et al. (2015). We also consider a scenario in which an order of magnitude improvement in beam modeling has been achieved and the beam is known to 1% in both the side-lobes and the main lobe which is the target precision for in-development drone experiments (Jacobs et al., 2017). We note that our model describes beam modeling errors that are completely correlated between antennas. It is possible that uncorrelated errors (which we might expect to arise from imperfections in the construction of each station) will integrate down differently from the modeling errors we consider here.

Plotting the resulting residual power spectra from equation 7.34 in Fig. 7-2, we




Figure 7-2: Top: Same as Fig. 7-1 except now we consider a perfect calibration catalog with calibration errors arising only from mismodeling the primary beam at the 1% level in the main-lobe and the 10% level in the side-lobes. With the exception of HERA, foreground residuals arising from primary beam modeling errors dominate the signal in the entire EoR window. Bottom: the same as above, but with a uniform modeling accuracy of 1%. Much of the EoR window is still contaminated for LOFAR and the MWA while significant bias exists in much of the EoR window for the SKA.

see that with the current precision of primary beam models, the calibration noise masks the power spectrum across all of k -space for the MWA, LOFAR, and the SKA. Even with an order of magnitude improvement in our modeling, a significant foreground bias $\gtrsim 20\%$ of the signal amplitude will be present in measurements by LOFAR and SKA-1. Thus, even under the optimistic foreground modeling scenarios considered above, foreground errors will still contaminate the EoR window unless significant improvements in beam modeling are made. HERA's compact layout limits the impact of beam-errors to small delays so that a significant portion of the EoR window is accessible even in the fiducial beam-modeling scenario.

7.3.4 The Dependence of Modeling Noise on Array and Catalog Properties

We can use the equations developed in § 7.2.3 to determine the impact of array configuration and catalog depth on the power spectrum bias $P_{\alpha\alpha}$. Since we are interested in the contribution from modeling errors which, unlike thermal noise, do not average down with integration time, we will let $\mathbf{C} = \mathbf{R}$ in equations 7.28-7.34. From equation 7.42, we list the effects of changing various parameters in the instrument design and source catalog in Table 7.1. There are a number of adjustments in the array layout that can be made to reduce the amplitude of the errors. Several of these adjustments have multiple effects that work against each other.

7.3.4.1 Catalog Depth

Ignoring diffuse emission, the power spectrum of modeling errors in equation 7.42 is proportional to $\sigma_r^2 = \int_0^{S_{\min}} S^2 S^{-\gamma} dS \propto S_{\min}^{3-\gamma}$. With the power law index of 1.75 for faint source populations, the noise level will scale as $S_{\min}^{1.25}$. Hence, clearing a contaminated region requires improvements in catalog depth on the same order of magnitude as the ratio of the bias to the expected signal.

Strategy	Impact on Error $P(\mathbf{k})$
Reduce S_{\min} .	Reduce amplitude as $S_{\min}^{3-\gamma}$.
Reduce the standard deviation of antenna positions, σ_{ant} .	(a) Increase amplitude as $\sigma_a^{2(3-\gamma)/(1-\gamma)}$. (b) Reduce maximum $k_{\parallel}^{\min} \sim \sigma_a$ of errors.
Increase N_{ant} .	(a) Reduce amplitude as N_{ant}^{-1} . (b) Decrease S_{\min} . [†]
Increase aperture diameter, d_{ant} .	(a) Reduces $k_{\parallel}^{\min} \sim d_{\text{ant}}^{-1}$. (b) Requires larger σ_{ant} leading to larger k_{\parallel}^{\min} .

[†] Depends on the distribution of additional antennas.

Table 7.1: Inspection of the equations in § 7.2 yields a number of analytic and qualitative relationships between the properties of an array and modeling catalog.

7.3.4.2 Time Averaging

If the instrumental gains are stable in time, modeling noise can be suppressed by averaging over LSTs. We investigate the level of suppression that is possible for *non fringe-stopped* baselines using multi-field averaging by calculating the temporal coherence of the modeling noise over some time interval, Δt . After a time Δt has passed, the primary beam of the instrument that had a gain of $A(\hat{\mathbf{s}})$ towards the brightness field at $I(\hat{\mathbf{s}})$ at time t will now have a gain of $A(\hat{\mathbf{s}})$ towards $I(\hat{\mathbf{s}} + \Delta\hat{\mathbf{s}})$ at time $t + \Delta t$,

$$\begin{aligned}
R_{\alpha\alpha}(\nu, t; \nu, t + \Delta t) &= \left(\int d\Omega d\Omega' e^{-2\pi i \nu b_{\alpha} \cdot (\hat{\mathbf{s}} - \hat{\mathbf{s}}')/c} \right) \times \\
&\quad \text{Cov}[I_r(\hat{\mathbf{s}}), I_r^*(\hat{\mathbf{s}}' + \Delta\hat{\mathbf{s}})] \\
&= (\sigma_r^2(S_{\min}) e^{-2\pi i \nu b_{\alpha} \cdot \Delta\hat{\mathbf{s}}/c}) \times \\
&\quad \int d\Omega A(\hat{\mathbf{s}}) A^*(\hat{\mathbf{s}} - \Delta\hat{\mathbf{s}}). \tag{7.50}
\end{aligned}$$

When $\Delta\hat{\mathbf{s}}$ is larger than the extent of the beam on the sky, the integral in equation 7.50 is close to zero. Hence a baseline is temporally coherent with itself when $\Delta\hat{\mathbf{s}}$ is small enough that its fields of view at the different times overlap. If the gains are stable over time, one can calibrate on multiple fields and reduce the power spectrum of calibration modeling errors by a factor of N_p , where N_p is the number of non-overlapping pointings. More significant suppression can arise from the oscillating

term in equation 7.50 which arises from our assumption that the sky has moved by $\Delta\hat{\mathbf{s}}$ and would not appear in the covariance between non fringe-stopped baselines. Averaging this oscillatory term over multiple LSTs can potentially lead to a significant reduction in the amplitude of modeling noise and is the subject of future work.

7.3.4.3 Array Configuration

There are three primary ways of changing the array configuration to affect modeling errors.

1. **Antenna Distribution:** Reducing the length of baselines involved in calibration reduces the chromaticity of gain errors and thus the smallest Fourier mode, k_{\parallel}^{\min} , that is not dominated by modeling noise. On the other hand, the array point spread function (PSF), and hence the minimal flux that an array can model for self-calibration is also set by its compactness. If the antennas are distributed as a Gaussian with standard deviation σ_a , then the naturally weighted PSF can be approximated by a Gaussian with standard deviation, $\sigma_p = \lambda_0/(2\pi\sigma_a)$. Condon (1974) determine that the confusion limit of an array, S_{\min} , depends on the PSF as $\sigma_p^{2/(\gamma-1)} \propto \sigma_a^{2/(1-\gamma)}$. Since the amplitude of the calibration noise is proportional to $S_{\min}^{3-\gamma}$, the overall normalization of calibration noise will scale with the standard deviation of the antenna distribution as $\sigma_a^{2(3-\gamma)/(1-\gamma)} \sim \sigma_a^{-3.33}$. At a glance, this is a very steep change in amplitude which might counteract the decrease in chromaticity. However, will find below that the impact of chromaticity is much more important.
2. **Antenna Count:** Increasing the number of antennas will cause the amplitude of the modeling noise power spectrum to reduce as $\sim N_{\text{ant}}$ but larger numbers of antennas will also force the array to be less compact, potentially increasing k_{\parallel}^{\min} while driving down the confusion limit.
3. **Antenna Size:** Increasing the size of each antenna reduces the primary beam width and hence the contamination from foregrounds at delays near the horizon but also drives up the minimal baseline size.

The scaling of the noise with the array characteristics listed above can be illuminated with some further simplifying assumptions. In particular, if all of the stations have Gaussian beams with angular standard deviations of $\sigma_b \approx \epsilon\lambda/d_{\text{ant}}$ where d_{ant} is the antenna diameter and $\epsilon \approx 0.45$ and that the antennas are Gaussian distributed with a standard deviation of σ_{ant} , equation 7.42 allows us to derive a closed-form prediction of the minimal k_{\parallel} in such an interferometer that is not contaminated by foregrounds (Appendix 7.E),

$$\begin{aligned}
k_{\parallel}^{\text{min}} \approx & 1.24 h\text{Mpc}^{-1} \sqrt{\frac{1+z}{10}} \left(\frac{\sigma_{\text{ant}}}{1\text{km}}\right) \left(\frac{d_{\text{ant}}}{10\text{m}}\right)^{-1} \times \\
& \left[1 + 0.35 \log\left(\frac{1+z}{10}\right) - 0.04 \log\left(\frac{\Omega_m}{0.27}\right) \right. \\
& + 0.1 \log\left(\frac{S_{\text{min}}}{10\text{mJy}}\right) - 0.08 \log\left(\frac{P_{21}}{10^4 \text{mK}^2 h^{-3} \text{Mpc}^3}\right) \\
& - 0.08 \log\left(\frac{N_{\text{ant}}}{100}\right) - 0.08 \log\left(\frac{\sigma_{\text{ant}}}{1\text{km}}\right) \\
& \left. + \log\left(\frac{d_{\text{ant}}}{10\text{m}}\right) \right] \tag{7.51}
\end{aligned}$$

where P_{21} is the amplitude of the 21 cm power spectrum. This formula can be used to get a quick order-of-magnitude sense as to whether a mode will be accessible to an instrument however it is very optimistic in that it assumes a Gaussian primary beam. While it also strictly assumes that the antennas are distributed as a Gaussian, we have found that it holds to 10% accuracy for non-Gaussian arrays (such as LOFAR and the MWA) as well.

From equation 7.51, we see that the extent of modeling noise contamination depends primarily on $\sigma_{\text{ant}}/d_{\text{ant}}$ while other quantities, such as N_{ant} and S_{min} , are contained within a logarithm have a much weaker impact on $k_{\parallel}^{\text{min}}$. This proportionality makes sense intuitively since larger apertures have smaller primary beams, suppressing emission at large zenith angles and larger delay. In close packed arrays, the $\sigma_{\text{ant}}/d_{\text{ant}}$ proportionality can be saturated so that the σ_r^2 inside of the logarithm will matter. While this equation ignores the existence of side-lobes, it gives us an order of magnitude estimate of how modeling noise scales with array properties. In Fig. 7-

3, the k_{\parallel}^{\min} values predicted from the naturally-weighted confusion limits of various planned arrays exceeds $k_{\parallel}^{\min} \gtrsim 0.2h\text{Mpc}^{-1}$, including for the SKA-1 core. Since interferometers such as the SKA and LOFAR focus most of their sensitivity at small k_{\parallel} values, their ability to detect the 21 cm signal will be heavily impacted by foreground modeling errors (Pober et al., 2014).

We can get a conservative sense for how side-lobes extend k_{\parallel}^{\min} beyond the values predicted in equation 7.51 by setting the amplitude of the modeling noise at zero-delay, multiplied by square of the side-lobe amplitudes (for an Airy beam, equal to -13dB) equal to the 21 cm signal (see Appendix 7.E). We denote the region of instrumental parameters space that is affected by side-lobes in Fig 7-3 with a grey overlay. Since all planned instruments fall within this region, the k_{\parallel}^{\min} predictions in this figure are actually optimistic. For these arrays, a more detailed calculation of equation 7.34 with realistic side-lobes is necessary. We found in § 7.3.2, with more realistic side-lobes considered, that the k_{\parallel}^{\min} obtained is indeed significantly larger than predicted by equation 7.51.

7.4 Eliminating Modeling Noise with Baseline Weighting

While optimistic scenarios in foreground characterization may be precise enough to suppress calibration modeling noise below the 21 cm signal, elimination of this contamination will also require beam characterization that is beyond the current state of the art. Enabling a power spectrum detection in existing sky-based calibrated experiments calls for an alternate strategy. Redundant calibration is one existing, and so-far successful alternative though it can only be applied to regularly spaced arrays. Though redundant calibration does not rely on a detailed sky model, it is possible that antenna-to-antenna beam variations and position errors can violate the assumption of redundancy and introduce chromatic artifacts that are similar to the ones we have found for sky-based calibration, a potential shortcoming that is being




Figure 7-3: Using equation 7.51, we show the smallest k_{\parallel}^{\min} that is not dominated by modeling noise errors for arrays with Gaussian primary beams, a random circular Gaussian distribution of N_{ant} antennas with standard deviation σ_{ant} , each with diameter d_{ant} . The area below the white dashed line is where side-lobes render the assumption of Gaussian beams inaccurate. Beige regions on the lower-right hand corner of each plot denote unphysically high packing densities. We see that for all existing instrument designs, calibration noise extends to large k_{\parallel} values that will reduce their sensitivity to the 21 cm signal. HERA will benefit greatly from the fact that it can be calibrated redundantly with minimal reliance on a sky model.

investigated. One approach is to ensure that the instrument contains no structure in region of k -space relevant for 21 cm studies, allowing for smooth fits that do not contaminate the EoR window (Barry et al., 2016). This is one of the approaches being adopted by HERA (Neben et al., 2016; Ewall-Wice et al., 2016a; Thyagarajan et al., 2016; Patra et al., 2017) and an upgrade to the MWA. In this section, we explore an alternative strategy that can be used even when the bandpass is not already intrinsically smooth. By exponentially suppressing long baselines, sky-based calibration is able to remove fine-frequency structure while avoiding contamination within the EoR window.

Supra-horizon contamination from calibration noise arises from the inclusion of longer baselines in calibrating gain solutions that are applied to short baselines, leaking power from large too small k_{\perp} . One way of mitigating this source of contamination is to weight the visibilities contributing to each gain solution in a way that dramatically up-weights short baselines over long ones. This can be accomplished by choosing

an appropriate \mathbf{W} matrix in equations 7.16 and 7.17. In § 7.4.1 we explore the efficacy of using a specific form of baseline weighting to eliminate modeling noise. The use of non-unity weights will result in an increase in thermal noise which we discuss in § 7.4.2.

7.4.1 Gaussian Weighting for Sky-Based Calibration

We explore the performance of a \mathbf{W} matrix that downweights long baselines with the functional form

$$W_{\alpha\beta} = \begin{cases} \exp\left(-\frac{b_{\alpha}^2}{2\sigma_w^2}\right) & \alpha = \beta \\ 0 & \alpha \neq \beta. \end{cases} \quad (7.52)$$

This function can result in weights that vary over a range beyond what is allowed for by numerical precision. In order to avoid poorly conditioned matrices, a regularization term is also added equal to the identity multiplied by 10^{-6} , which is large enough to avoid numerical precision errors, but also small enough such that the weights on long baselines are negligible compared to the short ones (and below the dynamic range between foregrounds and signal). With this weighting, core-antennas participating in many short baselines will have their gain solutions dominated by relatively achromatic core visibilities. Meanwhile, outrigger antennas that participate in only long-baselines will derive their solutions from many baselines with similarly small weights. In both cases, a normalization step of $(\mathbf{AWA}^T)^{-1}$ corrects for the fact that these weights do not sum to unity. Thus, long and short baselines are both effectively calibrated using the Gaussian weighting scheme while the leakage of chromatic errors on long baselines into gain solutions being applied to short baselines is stymied.

We calculate $P_{\alpha\alpha}$ given by equation 7.34 for the arrays considered in this paper with different values of σ_w . For LOFAR and the SKA we use $\sigma_w = 100$ m. For the MWA and HERA, whose cores are especially compact and have larger fields of view than LOFAR and the SKA, we apply more aggressive weighting with $\sigma_w = 50$ m. We compare the cylindrically binned and averaged results in the middle row of Fig. 7-4 with cylindrical power spectra with \mathbf{W} equal to the identity (top row) and find that

most of the EoR window is now free of foreground contamination with the power spectrum accessible at $k_{\parallel} \gtrsim 0.1 h\text{Mpc}^{-1}$ for most arrays.

However, stripes of foreground contamination still extend into the EoR window at distinct k_{\perp} values in the MWA, LOFAR, and to a lesser degree for the SKA. Isolating these baselines in the uv plane, we find that this contamination arises from antennas that are associated with less than two baselines that receive significant weighting. Define $n_{\text{eff}}(i)$ for the i^{th} antenna to be equal to the sum of the weights of all visibilities that include this antenna divided by their maximum value. As far as calibration is concerned, $n_{\text{eff}}(i)$ describes the effective number of baselines that an antenna participates in. If the number of effective baselines that are used to derive gain solutions is too small, the system is under-constrained and a degeneracy exists between possible solutions for the antenna gains. Two antennas have two gains to solve for, but only one visibility between them. The estimator is forced to break these degeneracies by up-weighting the contribution from the long baselines. An example where n_{eff} is smaller than two, in our Gaussian weighting scheme, would be for an antenna that is extremely far away from all but one other antenna. Only a single baseline associated with this antenna has significant weight while the rest are downweighted to zero.

We identify these problematic baselines by calculating $n_{\text{eff}}(i)$ for each antenna. We then flag and exclude from the fit the highest weighted visibilities on all antennas with $n_{\text{eff}} \leq 2$ until all n_{eff} are greater than 2. Flagging these visibilities leads to a loss in $\approx 6\%$ of visibilities for LOFAR, 1.3% for the MWA, 0.1% for the SKA, and no visibilities for HERA. The high n_{eff} for HERA antennas is something we would expect given its compact configuration (every antenna has many short baselines associated with it). Similarly, the SKA core we model is a compact Gaussian with few isolated antennas. LOFAR, on the other hand, has antennas that are arranged in pairs that are separated by short distances so that all of the isolated outriggers have a single short baseline associated with them (which results in the vertical stripe at $k_{\perp} \approx 0.4 h\text{Mpc}^{-1}$ in the second row of Fig. 7-4). The MWA lacks these pairs, and as a result has fewer low n_{eff} antennas which tend to lie in the transition between its compact core and

extended outriggers. We show cylindrically binned power spectra formed from the delay transform residuals of unflagged visibilities in the bottom row of Fig. 7-4, finding that upon flagging this small population of problematic baselines, the EoR window is almost entirely clear above $0.1 h\text{Mpc}^{-1}$ for all arrays studied. We also show the delay-transformed power spectrum estimates of visibilities contaminated by primary beam modeling errors of 1% at zenith and 10% in the side-lobes with and without Gaussian visibility weighting applied in the calibration solutions (Fig. 7-5). With Gaussian weighting, we are also able to mitigate contamination with the current level of primary beam modeling errors.

7.4.2 The Impact of Inverse Baseline Weighting on Power Spectrum Sensitivity

For an interferometer with identical antenna elements, the thermal noise level on every baseline is the same and \mathbf{N} is proportional to the identity matrix. For the point source approximation of the modeled foregrounds, the optimal weighting minimizing the errors due to thermal noise in each gain solution is therefore also the identity matrix. Because of its departure from identity weights, the Gaussian weighting that we proposed in the previous section has the effect of increasing thermal noise uncertainties in both the gains and the final power spectrum estimate. In order to see how Gaussian weighting increases the variance due to thermal noise in the gain solutions, one can consider the fact that the variance of the gain solutions goes as N_{ant}^{-1} (equation 7.42). For a particular antenna gain, Gaussian weighting reduces the effective number of visibilities whose noises are averaged over in each gain solution so that the variance of the antenna gain is now $\sim n_{\text{eff}}^{-1}$ rather than N_{ant}^{-1} . In the weighting schemes employed in § 7.4.1, n_{eff} goes down by a factor of order 1 – 10, remaining between 10 – 100 for LOFAR and the MWA.

Assuming Gaussian errors, the covariance between the square of two delay-transformed




Figure 7-4: Top: Residual power spectra with each visibility weighted equally in determining the calibration solutions (\mathbf{W} set to the identity matrix). Middle row: the same but now weighting visibilities with a Gaussian function of baseline length (equation 7.52). Much of the EoR window is cleared of contamination from calibration residuals. However pronounced stripes of contamination still exist, especially for LOFAR and the MWA. These stripes arise from short baselines formed from antennas involved in no other short baselines. In order to solve for both antenna gains, they must use information from long baselines, resulting in significant chromaticity on the few short baselines to which the problematic antenna gains are applied. Bottom: flagging visibilities after calibration until all gains participate in $n_{\text{eff}} \geq 2$ baselines, we find the EoR window free of these stripes. To reiterate, solid lines demarcate regions where the fiducial EoR signal is 1, 5, or 10 times the power of the calibration modeling error. The dashed diagonal line indicates the location of the wedge associated with the first null of the primary beam; the solid line indicates the horizon wedge.




Figure 7-5: Top: Cylindrically binned power spectra of calibration errors due to beam modeling errors at the level of 1% in the main-lobe and 10% in the side-lobes (equation 7.49). Each visibility has been weighted equally in determining the calibration solutions. Bottom row: The level of cylindrically binned power spectrum residuals from the primary-beam modeling errors in the top row but now with calibration solutions derived from visibilities that are weighted with equation 7.52. Short baselines contributing to antennas with $n_{eff} \leq 2$ have also been flagged from the calibration fit. Weighting with a Gaussian is capable for removing calibration modeling errors due to beam mismodeling at the level that we see in today's experiments (Neben et al., 2015).

visibilities is given by

$$\begin{aligned}\sigma_{\alpha\beta}^2 &= \left\langle |\tilde{V}_\alpha(\tau)|^2 |\tilde{V}_\beta(\tau)|^2 \right\rangle - \left\langle |\tilde{V}_\alpha(\tau)|^2 \right\rangle \left\langle |\tilde{V}_\beta(\tau)|^2 \right\rangle \\ &= (\mathbf{P}_{\alpha\beta}^N + \mathbf{P}_{\alpha\beta}^R + \mathbf{P}_{\alpha\beta}^S)^2\end{aligned}\tag{7.53}$$

where $\mathbf{P}_{\alpha\beta}^N \equiv \left\langle \tilde{V}_\alpha^N(\tau) \tilde{V}_\beta^{N*}(\tau) \right\rangle$ is the covariance matrix of the thermal noise component of delay-transformed visibilities and $\mathbf{P}_{\alpha\beta}^{R(S)}$ are the covariances of the delay-transformed residual foreground (signal) visibilities. While the residual foreground component can contribute significantly, it is only of concern in the regions of k -space where the amplitude of the foreground modeling noise is comparable to or greater than the level of the 21 cm signal. Since we are interested in how the thermal noise increases in the region of k -space where we have reduced foreground bias to well below the signal level, we will focus our attention on the thermal noise component and ignore the sample variance from modeling noise and signal for the remainder of this discussion.

We may compute $\mathbf{P}_{\alpha\beta}^N$ using equation 7.34 with $\mathbf{C} \rightarrow \mathbf{N}$. Typically, thermal noise is uncorrelated between baselines so $\tilde{\mathbf{N}}^{\alpha\beta}$ is diagonal. In the absence of calibration errors, the covariance between the squares of different delay-transformed visibility products arising from thermal noise would therefore also be zero. The presence of calibration errors introduces additional components to the thermal noise in all but the last term of equation 7.34 that are correlated from baseline to baseline. For identity weights and a diagonal noise-covariance, the off-diagonal terms in $\mathbf{P}_{\alpha\beta}$ go roughly as $\sim N_{\text{ant}}^{-2} \sim N_{\text{vis}}^{-1}$ compared to the diagonal terms (which have the order unity contribution that does not arise from calibration). Thus, for $\alpha \neq \beta$, $\sigma_{\alpha\beta}^2 \sim N_{\text{vis}}^{-2} \sigma_{\alpha\alpha}^2$ and has, so far, been ignored in other sensitivity calculations

(e.g. McQuinn et al. 2006, Parsons et al. 2012a, Beardsley et al. 2013, and Pober et al. 2014).

In order to obtain enough sensitivity for a detection, interferometry experiments are expected to perform spherical binning and averaging in k -space to obtain power-spectrum estimates, \hat{p}_A whose covariance we denote as Σ_{AB} (denoting band-powers with upper-case latin subscripts). The variance of a binned and averaged power spectrum estimate with identity weights is given by

$$\begin{aligned}\Sigma_{AA} &= N_A^{-2} \left(\sum_{\alpha \in A} \sigma_{\alpha\alpha}^2 + \sum_{\alpha \in A} \left[\sum_{\beta \in A; \alpha \neq \beta} \sigma_{\alpha\beta}^2 \right] \right) \\ &\sim N_A^{-2} \left(\sum_{\alpha \in A} \sigma_{\alpha\alpha}^2 + \frac{N_A}{N_{\text{vis}}^2} \sum_{\alpha \in A} \sigma_{\alpha\alpha}^2 \right).\end{aligned}\quad (7.54)$$

Thus, the contribution to Σ_{AA} from off-diagonal elements of the noise-covariance is sub-dominant to the contribution from diagonal elements as $\sim N_A/N_{\text{vis}}^2$ where N_A is the number of visibilities averaged within the A^{th} bin.

Non-uniform weighting in calibration decreases the effective number of visibilities in calibration, increasing the off-diagonal terms in equation 7.34. This in turn leads to an increase in the overall error bar on each spherically binned and averaged power spectrum estimate. We compute the degree to which Gaussian weighting degrades sensitivity to the spherically binned power spectrum by comparing Σ_{AA} for both uniform and Gaussian weighting within a single LST. While calibration correlates the noise on different squared visibilities in the same power spectrum bin, we can minimize the extra error by inverse-covariance weighting them before averaging.

We perform this averaging and report how the Gaussian down-weighting of long baselines affects the thermal noise on the final power spectrum estimate in Fig. 7-6. Because the covariance matrices for HERA and the SKA are very-large and would require significant computation to invert, we only perform this calculation for LOFAR and the MWA. We also assume that each power spectrum estimate only incorporates visibilities outside of the wedge. The proportion of long baselines which tend to be formed from antennas with smaller n_{eff} increases with each k -bin. Hence, the decrease




Figure 7-6: The ratio between thermal noise errors on a spherically averaged power spectrum estimate with Gaussian weighting and uniform weighting of visibilities in calibration. We show this ratio for LOFAR where $\sigma_w = 100$ m and the MWA with $\sigma_w = 50$ m. In both cases, the reduction in sensitivity to the power spectrum is by a factor less than two for small k where the interferometers have maximum sensitivity. Under the Gaussian weighting scheme, antennas with fewer short baselines have increased thermal noise in their gains. Increasingly large k -bins include larger numbers of visibilities formed from antennas with fewer short baselines (small n_{eff}) which have large increases in their thermal noise, leading to a trend of increasing sensitivity loss with increasing k . Since the MWA has a narrower weighting function, with $\sigma_w = 50$ m, this increase occurs faster than for LOFAR.

in sensitivity increases with k . Since the MWA weighting function is more compact, with $\sigma_w = 50$ m, the increase in the error ratio goes faster than for LOFAR which has a wider weighting function with $\sigma_w = 100$ m. Within the region that instruments are expected to be sensitive to the 21 cm signal, the error bars only go up by less than two. Gaussian weighting increases the thermal noise in the power spectrum measurement, but only by a level similar to intrinsic thermal noise that would be present even if calibration were perfect. Gaussian weighting can therefore allow us to circumvent the problem of foreground modeling noise in calibration while only sacrificing a small amount of sensitivity to the 21 cm power spectrum.

While baseline-dependent weighting is able to clear the EoR window, it does not necessarily allow any instrument to work within the wedge. Figs. 7-8 and 7-2 show that this would still require superb foreground models accurate to the 0.1 mJy level and modeling of the primary beam to the 10^{-3} level in the main lobe and 10^{-2} level in the side-lobes. Until these milestones are achieved, extended arrays will suffer a

disproportionate reduction in delay power-spectrum sensitivity relative to compact arrays like HERA (Pober et al., 2014).

7.5 Conclusions

In this work, we derived expressions for the amplitude of the power spectrum bias arising from the imprint of foreground modeling errors on calibration. These expressions assumed that calibration errors are small enough such that their solutions are obtained through a linear set of equations, which is the case in the final stages of iterative, sky-based calibration schemes when the errors in the foreground model are small. Using these equations we are able to explain the amplitude of the biases that have been simulated for the special cases of the MWA (B16) and LOFAR (Patil et al., 2016) and to predict the amplitude of modeling noise in the power spectrum for the SKA-1 and HERA (which does not actually rely on this approach). We performed this analysis in a variety of foreground and beam modeling scenarios. We also use our formalism to determine the dependence of modeling noise on the parameters of the array and the accuracy of the calibration catalog. These results do not apply to the redundant calibration strategies used by HERA and PAPER, although errors introduced by deviations from redundancy still have the potential to contaminate the window in a similar way. Our analysis also reveals that noise bias exists in current power spectrum estimates where separate calibration solutions are not obtained for interleaved data sets. Whether this bias limits 21 cm experiments requires further analysis but it can easily be avoided by obtaining independent calibration solutions for cross multiplied data.

This paper aimed to illuminate the source of calibration errors within the EoR window. In order to make our analysis analytically tractable, we employed a number of assumptions. These include assuming that the array is minimally redundant so that we can ignore off diagonal elements of the visibility covariance matrix, and that the sources themselves are flat-spectrum. A more significant assumption that will not hold in many observing scenarios is that we ignored the chromaticity of modeled

foregrounds, which holds approximately when the modeled fluxes are dominated by a source at the phase center that exceeds the flux of the next brightest source by a factor of a few. We also assumed that our instruments had Airy-beams, that sources could be characterized down to a fixed flux-level across the entire sky, ignored ionospheric effects and polarization (Sault et al., 1996; Jelić et al., 2010; Moore et al., 2013; Asad et al., 2015; Kohn et al., 2016; Moore et al., 2017) which is especially severe on the large spatial scales (Lenc et al., 2016) that we suggest should be relied upon in calibration strategies. Hence, specific quantitative predictions in this paper should be regarded as accurate to within an order of magnitude and on the optimistic side. In validating the design of a future instruments, full end-to-end simulations should be employed, though this is left to future work.

Our calculations indicate that for current catalog limits presented in Carroll et al. (2016); Hurley-Walker et al. (2017), and Williams et al. (2016) both the MWA and LOFAR will observe an EoR window that is heavily contaminated by chromatic calibration errors due to unmodeled sources. Since the chromaticity of these errors increases with the length of baselines involved in calibration, removing inner baselines from calibration, as is required to avoid signal loss with direction dependent calibration (Patil et al., 2016) will only exacerbate these chromatic errors and is probably the source of the systematics floor observed by LOFAR in Patil et al. (2017) (these authors note that calibration errors as a likely culprit but not that the use of long baselines is exacerbating the problem). Our analytic treatment suggests that instead, sky-based experiments should use their short baselines to calibrate power-spectrum data which may preclude the direction-dependent approach to avoid signal loss and will likely require more accurate models of diffuse emission. LOFAR may also be able reduce the amplitude of calibration errors below the power spectrum, at large spatial scales, by averaging over multiple fields of view (if its gains are temporally stable) and/or by building a source catalog complete down to $\approx 100 \mu\text{Jy}$ across the entire sky. Even if such a catalog is constructed, beam modeling precision will also need to be improved by an order of magnitude over what has been achieved in the literature. The large field of view on the MWA decreases the number of fields that can be aver-

aged over and increases the k_{\parallel} values contaminated by modeling errors, making the path to removing this noise with extant methods considerably more difficult than for LOFAR.

Our analysis motivates a potential solution to the problem of modeling noise in sky-based calibration. Since contamination within the EoR window arises from the coupling of long baseline errors into the calibration solutions on short ones, our proposed strategy is to down-weight the contribution of long baselines to the gain solutions that are applied to short baselines. The linear least-squares estimator formalism employed in this paper provides a natural framework for incorporating such weights. Experimenting with a Gaussian weighting scheme, we find that down-weighting long baselines should allow for both existing and future arrays to correct fine-frequency bandpass structures without introducing chromatic sky-modeling errors. While such weighting will increase the level of thermal noise present in calibration solutions, we find that this noise increase will only result in power spectrum error bars that are $\approx 1 - 1.5$ times larger than the case where all visibilities are weighted identically. This method prevents calibration errors from limiting the foreground avoidance approach, which seeks to detect the 21 cm signal within the EoR window and thus requires the calibrated instrumental response to be spectrally smooth. This method is not sufficient to enable *foregrounds subtraction*; accessing the signal inside of the wedge. Working within the delay-wedge will require significant improvements in foreground and primary beam modeling.

7.A Appendix: The Impact of Redundancy

Throughout this paper, we ignored the impact of redundancy between visibilities, letting $R_{\alpha\beta}$ be diagonal when calculating modeling noise. However, redundancy is significant in highly compact arrays, such as HERA. Here we argue that the impact of redundancy on the modeling noise levels, calculated in this work, is to multiply the overall noise level by a factor of order unity which only has a small effect on the extent of contaminated modes in k -space. We also verify this argument with a

numerical calculation.

When \mathbf{C} is diagonal, the sum in equation 7.35 is only over terms with $\gamma = \delta$. The existence of redundant baselines introduces non-negligible off-diagonal terms in the visibility covariance matrix \mathbf{R} . For each $\gamma = \delta$ term in the non-redundant sum, we can consider the additional summands, with $\gamma \neq \delta$ that are introduced for each ii/jj term and ij/ji term. We start with ii/jj .

For a fixed baseline γ that involves antenna i , there will be at most $\sim N_{\text{ant}}$ additional baselines that are redundant with γ and do not involve gain i . From equations 7.39 and 7.40, the weighting of the covariance between two different baselines in which only one involves antenna i goes as $\sim N_{\text{ant}}^{-3}$. Thus, the presence of redundancy adds no more than $\sim N_{\text{ant}}$ terms involving antenna i but not antenna j and vice versa, for each ii summand in equation 7.35. Multiplying this overall factor of N_{ant}^{-2} by N_{ant} to account for the N_{ant} different ii sums leads to a contribution to the noise amplitude on the order of $\sim N_{\text{ant}}^{-1}$, similar to the level of the noise without redundancy. As a result, redundancy changes the modeling error amplitude by a factor of order unity in the diagonal terms. Next we consider the ij terms in equation 7.35.

For a given γ and $i \neq j$, there will be at most $\sim N_{\text{ant}}$ redundant baselines that do not involve the i^{th} or j^{th} gains, causing the weighting of each unique variance term to go as N_{ant}^{-3} rather than N_{ant}^{-4} in the non-redundant case. Since there are $\sim N_{\text{ant}}^2$ unique baselines that do not involve i or j , the overall sum of these terms goes as N_{ant}^{-1} . As a result, the ij terms in equation 7.35 will have a similar magnitude as the ii/jj terms but the overall impact on the amplitude of the modeling noise described in equation 7.35 still changes the amplitude by a factor of order unity.

We confirm these arguments with a numerical comparison between the amplitude of the modeling noise with and without redundancy taken into account for two redundant arrays of 91 and 331 hexagonally packed 14 m apertures and S_{min} equal to the naturally-weighted confusion limit. We compute the off-diagonal elements of \mathbf{R} by numerically computing the beam integral in equation 7.5 for all $R_{\alpha\beta}$ with Airy beams and perform the full matrix inversions prescribed in equation 7.16 and 7.17. We compare our results to the same calculation where all off-diagonal elements of




Figure 7-7: We compare the amplitude of modeling noise on a short baseline when \mathbf{R} is assumed to be diagonal (orange line) and the off-diagonal terms of \mathbf{R} are explicitly included for a 91-element (dashed lines) and 331-element (solid lines) hexagonally packed array of 14 m apertures. We find that even in a maximally redundant array, the effect of redundancy is to change the overall amplitude of the modeling noise by a factor of a few. This only has a small impact ($\lesssim 10\%$) on k_{\parallel}^{\min} , the smallest k_{\parallel} where the 21 cm signal (black line) dominates over the modeling noise, as computed from equation 7.51 which ignores the effect of redundancy.

\mathbf{R} are set to zero (Fig. 7-7) and find that the difference in amplitude is essentially a factor of order unity, leading to a negligible increase in the effective k_{\parallel}^{\min} . This calculation confirms our argument for HERA-scale arrays.

7.B Appendix: The Point Source Approximation For Modeled Foregrounds.

For analytic tractability, we assumed that the modeled component of foregrounds were well characterized by a flat-spectrum point source at zenith, whose visibilities are achromatic. Throughout the paper, the rest of the unmodeled foregrounds considered in our calculations were not assumed to be a single point source and are characterized by chromatic visibilities (see § 7.2.1.1-7.2.1.2). In this appendix we explore the consequences of relaxing this assumption.

A significant consequence of the foregrounds not being dominated by a single point source at the phase center is that for a fixed frequency, y_α 's amplitude will vary significantly from baseline to baseline, often approaching zero where source fringes destructively interfere. As a result, $\text{Cov} \left[\frac{c_\alpha}{y_\alpha}, \frac{c_\alpha^*}{y_\alpha^*} \right]$ can vary rapidly in frequency where y_α approaches zero. Thus, any weighting scheme that does not take these nulls into account will experience calibration error chromaticity in large excess of what we have found so far.

Instead, it is typical for calibration solutions to be obtained for each frequency through inverse covariance weighting. Since the thermal-noise covariance matrix is usually proportional to the identity, per-frequency inverse covariance weights are proportional to $|y_\alpha|^2$. Under this scheme, we may employ a weights matrix that is frequency dependent.

$$W_{\alpha\alpha} \rightarrow W'_{\alpha\alpha} = W_{\alpha\alpha} |y_\alpha(\nu)|^2 \quad (7.55)$$

which leads Λ and Ψ to be frequency dependent as well and we can no longer separate them from \tilde{C} in the delay-transform. The delay-transform visibility in equation 7.25 becomes

$$\tilde{V}_\alpha \approx \int d\nu e^{2\pi i \nu \tau} \left[y_\alpha \left(\tilde{\eta}_i + \tilde{\eta}_j + i\tilde{\phi}'_i - i\tilde{\phi}'_j \right) + c_\alpha \right] \quad (7.56)$$

where every term, including y_α is a function of frequency. The expectation value for the Delay-transformed product of \tilde{V}_α with its complex conjugate to second order in

\mathbf{c}/\mathbf{y} (equation 7.28) is now,

$$\begin{aligned}
P_{\alpha\beta} = \int d\nu d\nu' e^{2\pi i\tau(\nu-\nu')} & [y_\alpha y_\beta^* \langle \eta'_i \eta_\ell'^* \rangle + y_\alpha y_\beta^* \langle \eta'_j \eta_m'^* \rangle \\
& + y_\alpha y_\beta^* \langle \eta'_j \eta_\ell'^* \rangle + y_\alpha y_\beta^* \langle \eta'_i \eta_m'^* \rangle \\
& + y_\alpha y_\beta^* \langle \phi'_i \phi_\ell'^* \rangle + y_\alpha y_\beta^* \langle \phi'_j \phi_m'^* \rangle \\
& - i y_\alpha y_\beta^* \langle \phi'_i \phi_m'^* \rangle - i y_\alpha y_\beta^* \langle \phi'_j \phi_\ell'^* \rangle \\
& + y_\alpha \langle \eta'_i c_\beta^* \rangle + y_\alpha \langle \eta'_j c_\beta^* \rangle \\
& + i y_\alpha \langle \phi'_i c_\beta^* \rangle - i y_\alpha \langle \phi'_j c_\beta^* \rangle \\
& + y_\beta^* \langle \eta_\ell^* c_\alpha \rangle + y_\beta^* \langle \eta_m^* c_\alpha \rangle \\
& - i y_\beta^* \langle \phi_\ell^* c_\alpha \rangle + i y_\beta^* \langle \phi_m^* c_\alpha \rangle \\
& + \langle c_\alpha c_\beta^* \rangle.] \tag{7.57}
\end{aligned}$$

where every complex conjugated quantity is a function of ν' and every non-conjugated quantity is a function of ν . Since the weight and design matrices are no longer frequency independent, second order moments cannot be separated into frequency dependent and independent components as we did with the point source approximation. In order to compute $P_{\alpha\beta}$, we must calculate all second order moments with a given source model and design matrix and take the Fourier transforms.

For realistic y_α , we use simulations of point source foregrounds obtained from the PRISim software package (Thyagarajan et al., 2015a) for the MWA-128T array layout with antennas modeled as 4 m diameter dishes. For each 100 kHz channel over a 20 MHz band, we use a weights matrix $W'_{\alpha\alpha}(\nu) = W_{\alpha\alpha} |y_\alpha(\nu)|^2$ and compute the two-dimensional Fourier transform in equation 7.57 to obtain $P_{\alpha\alpha}$ for several baselines. We take the MWA to be pointing at a declination equal to its latitude of -26.701° (Tingay et al., 2013a) at LST=0 and 4 hr. We run two different simulations, one in which $W_{\alpha\alpha}$ is set to unity (and the weights matrix $W'_{\alpha\alpha} = W_{\alpha\alpha} |y_\alpha|^2$) and the other where $W_{\alpha\alpha}$ is given by equation 7.52 with $\sigma_w = 50$ m.

In Fig. 7-8, we show the amplitude of calibration modeling noise on a short (8 m) and long (174 m) MWA baseline with and without fully modeled foregrounds for




Figure 7-8: Comparisons between calibration modeling noise with realistic modeled foregrounds (grey lines) and the point source foregrounds used throughout this paper (black lines) with (solid) and without (dashed). Gaussian baseline weighting for two different LSTs and baseline lengths on the MWA from 10 MHz noise equivalent bandwidth centered at 150 MHz. The red line denotes the amplitude of the HI power spectrum generated with 21cmFAST. We find that fully modeled foregrounds change the overall amplitude of the of the calibration noise since the amplitude of a particular modeled visibility does not necessarily equal the amplitudes of every other modeled visibility. Chromatic y_α s also introduce some additional spectral structure which results in a larger width of calibration errors in k_{\parallel} . The overall impact on the LoS mode where modeling noise bias falls below the 21 cm signal is only on the order of 10% with Gaussian weighting.

several different LSTs. In all cases, we see that the fully modeled foregrounds extend the width of the foreground noise to larger k_{\parallel} , something we would expect to occur with the additional spectral structure they introduce. In addition, the amplitude of the modeling noise is modified since the multiplication by the modeled foregrounds on a particular baseline, $y_{\alpha}y_{\beta}^*$ (equation 7.57) does not necessarily cancel out the modeled foregrounds in the numerator of each summed $\langle c_{\gamma}c_{\delta}^* \rangle / (y_{\gamma}y_{\delta}^*)$ as they do when y_{α} is constant. Despite these differences, we find that over the range of LSTs and baselines studied, the overall impact on the minimal LoS wavenumber of modes that can be observed at a particular uv point is only on the order of $\approx 10\%$. Thus, the approximation of the modeled foregrounds as a point-source at the phase center does not have a significant impact on range of modes that are masked by foreground modeling errors.

7.C Appendix: Expressions for Second Moments of Delay-Transformed Calibration Errors.

In this section, we derive the approximate expressions for the second moments that we use to go from equation 7.28 to equation 7.34. To derive equations 7.29 through 7.33, we first note that

$$\begin{aligned} \text{Cov}[\text{Re}(\mathbf{c}), \text{Re}(\mathbf{c})^{\top}] &\approx \frac{1}{2} \mathbf{C}. \\ \text{Cov}[\text{Im}(\mathbf{c}), \text{Im}(\mathbf{c})^{\top}] &\approx \frac{1}{2} \mathbf{C}. \end{aligned} \tag{7.58}$$

This assertion is true for thermal noise, \mathbf{n} , since both the real and imaginary components of the thermal noise are given by identical, zero-mean normal distributions. We need only show that this assertion holds for the unmodeled foregrounds \mathbf{r} . We

start by writing

$$\begin{aligned}
& \text{Cov}[\text{Re}(\mathbf{r}), \text{Re}(\mathbf{r})^\top]_{\alpha\beta} \\
& \propto \int d\Omega \cos\left(\frac{2\pi\nu\mathbf{b}_\alpha \cdot \hat{\mathbf{s}}}{c}\right) \cos\left(\frac{2\pi\nu'\mathbf{b}_\beta \cdot \hat{\mathbf{s}}}{c}\right) |A(\hat{\mathbf{s}})|^2 \\
& = \frac{1}{4} \int d\Omega \left[e^{2\pi i(\nu\mathbf{b}_\alpha + \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} + e^{-2\pi i(\nu\mathbf{b}_\alpha + \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} \right] |A(\hat{\mathbf{s}})|^2 \\
& + \frac{1}{4} \int d\Omega \left[e^{2\pi i(\nu\mathbf{b}_\alpha - \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} + e^{-2\pi i(\nu\mathbf{b}_\alpha - \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} \right] |A(\hat{\mathbf{s}})|^2, \tag{7.59}
\end{aligned}$$

where we dropped the multiplicative instrument-independent terms in the covariance (equations 7.5 & 7.6) in favor of a proportionality sign. All of the terms in equation 7.59 will integrate to zero unless $|\mathbf{b}_\alpha \pm \mathbf{b}_\beta| \lesssim d_{\text{ant}}$ (less than one fringe fits within the primary beam main-lobe) which is only true if $\mathbf{b}_\alpha \approx \pm \mathbf{b}_\beta$ where the negative case causes exponential terms in the first line of equation 7.59 to be non-zero and the positive case causes the second line to be non-zero. We may choose baseline indexing such that we never have $\mathbf{b}_\alpha \approx -\mathbf{b}_\beta$, by having antenna numbers increase with increasing E-W and than N-S position. With this indexing,

$$\begin{aligned}
& \text{Cov}[\text{Re}(r_\alpha), \text{Re}(r_\beta)] \\
& \propto \frac{1}{4} \int d\Omega \left[e^{2\pi i(\nu\mathbf{b}_\alpha - \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} + e^{-2\pi i(\nu\mathbf{b}_\alpha - \nu'\mathbf{b}_\beta) \cdot \hat{\mathbf{s}}/c} \right] |A(\hat{\mathbf{s}})|^2 \\
& = \frac{1}{4} [\text{Cov}(r_\alpha, r_\beta^*) + \text{Cov}(r_\alpha^*, r_\beta)]. \tag{7.60}
\end{aligned}$$

For beams that are symmetric around the phase center, $\text{Cov}[r_\alpha, r_\beta^*]$ is real and $\text{Cov}[r_\alpha, r_\beta^*] = \text{Cov}[r_\alpha^*, r_\beta]$, proving our assertion that

$$\text{Cov}[\text{Re}(r_\alpha), \text{Re}(r_\beta)] = \frac{1}{2} \mathbf{R}_{\alpha\beta}. \tag{7.61}$$

A very similar set of steps with identical assumptions yields

$$\text{Cov}[\text{Im}(r_\alpha), \text{Im}(r_\beta)] = \frac{1}{2} \mathbf{R}_{\alpha\beta}. \tag{7.62}$$

Next, we show that

$$C_{\alpha\beta} \gg \langle c_\alpha \rangle \langle c_\beta^* \rangle. \quad (7.63)$$

This can be seen by writing the product of the averages

$$\begin{aligned} \langle c_\alpha \rangle \langle c_\beta^* \rangle &= \int d\Omega e^{-2\pi i \nu \mathbf{b}_\alpha \cdot \hat{\mathbf{s}}/c} A(\hat{\mathbf{s}}) \langle I(\hat{\mathbf{s}}) \rangle \\ &\times \int d\Omega' e^{2\pi i \nu' \mathbf{b}_\beta \cdot \hat{\mathbf{s}}'/c} A(\hat{\mathbf{s}}') \langle I(\hat{\mathbf{s}}') \rangle \end{aligned} \quad (7.64)$$

both integrate to zero when $b_{\alpha/\beta} \gtrsim d_{\text{ant}}$ and $\langle I(\hat{\mathbf{s}}) \rangle$ is smooth as a function of position (which is typically true of foreground residuals and signal).

We can now show derivations for equations 7.29 through 7.33.

7.C.1 Derivation of Equation 7.29

We start on the left hand side with

$$\begin{aligned} \langle \tilde{\eta}_i \tilde{\eta}_j^* \rangle &= \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} [[\mathbf{C}_\eta]_{ij} + \langle \eta_i \rangle \langle \eta_j^* \rangle] \\ &\approx S_0^{-2} \Lambda_{i\gamma} \Lambda_{\delta j}^\top \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} \text{Cov}[\text{Re}(c^\gamma), \text{Re}(c^\delta)] \\ &= \frac{S_0^{-2}}{2} \Lambda_{i\gamma} \Lambda_{\delta j}^\top \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} C^{\gamma\delta}(\nu, \nu') \\ &= \frac{S_0^{-2}}{2} \Lambda_{i\gamma} \Lambda_{\delta j}^\top \tilde{C}^{\gamma\delta} \end{aligned} \quad (7.65)$$

In going from the first to the second line, we threw away the product of the means (equation 7.63). Going from the second to the third line, we used equation 7.61.

7.C.2 Derivation of Equation 7.30

Following the same procedure for equation 7.29,

$$\begin{aligned}
\langle \phi'_i \tilde{\phi}'_j{}^* \rangle &= \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} [[\mathbf{C}_\phi]_{ij} + \langle \phi_i \rangle \langle \phi_j^* \rangle] \\
&\approx S_0^{-2} \Psi_{i\gamma} \Psi_{\delta j}^\top \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} \text{Cov}[\text{Im}(c^\gamma), \text{Im}(c^\delta)] \\
&= \frac{S_0^{-2}}{2} \Psi_{i\gamma} \Psi_{\delta j}^\top \int d\nu d\nu' e^{-2\pi i \tau(\nu - \nu')} C^{\gamma\delta}(\nu, \nu') \\
&= \frac{S_0^{-2}}{2} \Psi_{i\gamma} \Psi_{\delta j}^\top \tilde{C}^{\gamma\delta}.
\end{aligned} \tag{7.66}$$

7.C.3 Derivation of Equation 7.31

Starting with the left-hand side of equation 7.31,

$$\begin{aligned}
\langle \tilde{c}_\alpha \tilde{\eta}_i^* \rangle &= \frac{S_0^{-1}}{2} \int d\nu d\nu' e^{-2\pi i(\nu - \nu')\tau} \Lambda_{i\gamma} \langle c_\alpha (c^\gamma + c^{\gamma*}) \rangle \\
&\approx \frac{S_0^{-1}}{2} \Lambda_{i\gamma} \int d\nu d\nu' e^{-2\pi i(\nu - \nu')\tau} C_\alpha^\gamma(\nu, \nu') \\
&= \frac{S_0^{-1}}{2} \Lambda_{i\gamma} \tilde{C}_\alpha^\gamma.
\end{aligned} \tag{7.67}$$

To go from the first to second line here, we used the fact that $\langle c_\alpha c^\gamma \rangle \propto \int d\Omega e^{-2\pi \hat{\mathbf{s}} \cdot (\mathbf{b}_\alpha \nu + \mathbf{b}_\beta \nu')/c} |A(\hat{\mathbf{s}})|^2$ which, as discussed above, integrates to zero for $b_{\alpha/\beta} \gtrsim d_{\text{ant}}$.

7.C.4 Derivation of Equation 7.32

Following the same steps used for equation 7.31,

$$\begin{aligned}
\langle \tilde{c}_\alpha \tilde{\phi}_i^* \rangle &= \frac{S_0^{-1}}{2i} \int d\nu d\nu' e^{-2\pi i(\nu - \nu')\tau} \Psi_{i\gamma} \langle c_\alpha (c^\gamma - c^{\gamma*}) \rangle \\
&\approx \frac{i S_0^{-1}}{2} \Psi_{i\gamma} \int d\nu d\nu' e^{-2\pi i(\nu - \nu')\tau} C_\alpha^\gamma(\nu, \nu') \\
&= S_0^{-1} \frac{i}{2} \Psi_{i\gamma} \tilde{C}_\alpha^\gamma.
\end{aligned} \tag{7.68}$$

7.C.5 Derivation of Equation 7.33

We may show this last identity by expanding the real and imaginary components of c .

$$\begin{aligned}
\langle \tilde{\eta}_i \tilde{\phi}_j^* \rangle &= S_0^{-2} \int d\nu d\nu' e^{2\pi i(\nu - \nu')\tau} \Lambda_{i\gamma} \Psi_{\delta\mathbf{a}\mathbf{e}}^\top \langle \text{Re}(c^\gamma) \text{Im}(c^\delta)^* \rangle \\
&= \frac{S_0^{-2}}{4i} \int d\nu d\nu' e^{2\pi i(\nu - \nu')\tau} \Lambda_{i\gamma} \Psi_{\delta\mathbf{a}\mathbf{e}}^\top \langle (c^\gamma + c^{\gamma*})(c^\delta - c^{\delta*}) \rangle \\
&\approx 0
\end{aligned} \tag{7.69}$$

We obtain the last line approximately equal to zero due to the fact that $\langle (c^\gamma + c^{\gamma*})(c^\delta - c^{\delta*}) \rangle = \langle c^\gamma c^\delta \rangle + \langle c^{\gamma*} c^{\delta*} \rangle + \langle c^\gamma c^{\delta*} \rangle - \langle c^{\gamma*} c^\delta \rangle$. The first two terms evaluate to zero since they involve integrals over $e^{\pm 2\pi i \hat{\mathbf{s}} \cdot (\mathbf{b}_\gamma \nu + \mathbf{b}_\delta \nu')}$ and the last two terms are equal to each other so they subtract to give 0.

7.D Appendix: Components of Λ and Ψ for Non-Redundant, Uniformly Weighted Calibration Solutions

In this section, we derive equations 7.36 and 7.37 which are valid when the weights matrix is equal to unity. While of limited applicability, they provide us with insight into the scaling of modeling noise with properties of the source catalog and array and allow us to identify the degree to which any visibility covariance contributes to the covariances of gain solutions.

7.D.1 Equation 7.36

We wish to evaluate

$$\Lambda_{i\gamma} = [(\mathbf{A}\mathbf{A}^\top)^{-1}\mathbf{A}^\top]_{i\gamma}. \tag{7.70}$$

We start with $(\mathbf{AA}^\top)_{ij}$. Evaluating this matrix product for a non-redundant array is straightforward since each element is given by the dot-product of the i^{th} column of \mathbf{A} with the j^{th} column. Since a given column is equal to unity at the indices of visibilities in which that antenna participates in and zero otherwise, the dot product of columns is equal to $N_{\text{ant}} - 1$ if $i = j$ and equal to unity if $i \neq j$.

$$\mathbf{AA}^\top = \begin{pmatrix} N_{\text{ant}} - 1 & 1 & 1 & \dots & 1 \\ 1 & N_{\text{ant}} - 1 & 1 & \dots & 1 \\ 1 & 1 & N_{\text{ant}} - 1 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & N_{\text{ant}} - 1 \end{pmatrix}. \quad (7.71)$$

\mathbf{AA}^\top can be decomposed into the sum of a diagonal matrix and a matrix formed from an outer product,

$$\mathbf{AA}^\top = \begin{pmatrix} N_{\text{ant}} - 2 & 0 & 0 & \dots & 0 \\ 0 & N_{\text{ant}} - 2 & 0 & \dots & 0 \\ 0 & 0 & N_{\text{ant}} - 2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & N_{\text{ant}} - 2 \end{pmatrix} + \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & 1 & 1 & \dots & 1 \\ 1 & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 1 \end{pmatrix}, \quad (7.72)$$

and thus can be inverted using the Sherman-Morrison formula.

$$(\mathbf{AA}^\top)_{ij}^{-1} = \begin{cases} \frac{1}{N_{\text{ant}} - 2} \left[1 - \frac{1}{2(N_{\text{ant}} - 1)} \right] & i = j \\ \frac{-1}{2(N_{\text{ant}} - 1)(N_{\text{ant}} - 2)} & i \neq j. \end{cases} \quad (7.73)$$

We can now evaluate $(\mathbf{A}\mathbf{A}^\top)_{ij}^{-1}(\mathbf{A}^\top)^j_\gamma$ which is the sum of the entries in the i^{th} row of $(\mathbf{A}\mathbf{A}^\top)^{-1}$ that correspond to antennas that participate in the γ^{th} baseline. If $i \in \gamma$, we add an entry in $(\mathbf{A}\mathbf{A}^\top)^{-1}$ where $i = j$ to an entry where $i \neq j$. If $i \notin \gamma$ we add two entries in $(\mathbf{A}\mathbf{A}^\top)^{-1}$ where $i \neq j$. For these two cases we get,

$$\Lambda_{i\gamma} = [(\mathbf{A}\mathbf{A}^\top)^{-1}\mathbf{A}^\top]_{i\gamma} = \begin{cases} \frac{1}{N_{\text{ant}}-1} & i \in \gamma \\ \frac{-1}{(N_{\text{ant}}-1)(N_{\text{ant}}-2)} & i \notin \gamma, \end{cases} \quad (7.74)$$

which completes the proof.

7.D.2 Equation 7.37

We begin evaluating $\Psi_{i\gamma} = (\mathbf{B}\mathbf{B}^\top)^{-1}\mathbf{B}^\top$ with the product $(\mathbf{B}\mathbf{B}^\top)_{ij}$ which is the dot product of the i^{th} column of \mathbf{B} with the j^{th} column. Each i^{th} column contains $N_{\text{ant}} - 1$ non-zero rows that are 1 when the i^{th} antenna is the non-conjugated participant in the baseline and -1 when the antenna is the conjugated participant. The last row of \mathbf{B} is composed entirely of ones. Thus the dot product of any column with itself is N_{ant} and the dot product of a column with any other is equal to zero. Thus,

$$(\mathbf{B}\mathbf{B}^\top)_{ij} = \begin{cases} N_{\text{ant}} & i = j \\ 0 & i \neq j \end{cases}, \quad (7.75)$$

whose inverse is trivial. $(\mathbf{B}\mathbf{B}^\top)^{-1}(\mathbf{B}^\top)^j_\gamma$ is the sum of each element of the i^{th} row of $(\mathbf{B}\mathbf{B}^\top)^{-1}$ that participates the γ^{th} visibility. Since $(\mathbf{B}\mathbf{B}^\top)^{-1}$ is diagonal, this sum is only non-zero when $i = j$. If i is the non-conjugated antenna in the visibility, than \mathbf{B}^\top multiplies by 1 and if i is the conjugated antenna in γ , \mathbf{B}^\top multiplies by -1 . We obtain

$$\Psi_{i\gamma} = \begin{cases} \frac{1}{N_{\text{ant}}} & \gamma = (i, \cdot) \\ -\frac{1}{N_{\text{ant}}} & \gamma = (\cdot, i) \\ 0 & i \notin \gamma \end{cases}, \quad (7.76)$$

completing our proof.

7.E Appendix: A Simplified Expression for Minimal Accessible Line-of-Sight Modes.

In this section, we derive equation 7.51 from equation 7.42 with the additional assumptions that the array has a Gaussian beam with standard deviation $\sigma_b = \epsilon\lambda_0/d_{\text{ant}}$ and that its antennas are arranged in a Gaussian configuration with standard deviation σ_{ant} . We also ignore the contribution from thermal noise, which integrates down with time, and assume that $\mathbf{C} = \mathbf{R}$. We first compute $\tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau)$ for a Gaussian beam

$$\begin{aligned}\tilde{\mathbf{C}}_{\delta\delta}(\tau, \tau) &= \sigma_r^2 \int d\nu d\nu' \int d\Omega e^{2\pi i\tau(\nu-\nu')} e^{-2\pi i\mathbf{b}_\delta \cdot \hat{\mathbf{s}}(\nu-\nu')/c} |A(\hat{\mathbf{s}})|^2 \\ &\approx \sigma_r^2 \int d\nu d\nu' \int d\Omega e^{2\pi i\tau(\nu-\nu')} e^{-2\pi i\mathbf{b}_\delta \cdot \hat{\mathbf{s}}(\nu-\nu')/c} e^{-s^2/2\sigma_b^2} \\ &\approx \sigma_r^2 B \int d\Delta\nu \int d\Omega e^{2\pi i\tau\Delta\nu} e^{-2\pi i\mathbf{b}_\delta \cdot \hat{\mathbf{s}}\Delta\nu/c} e^{-s^2/\sigma_b^2} \\ &\approx \sigma_r^2 B \frac{c}{b_\delta} \sqrt{\pi} \sigma_b \exp\left(-\frac{c^2\tau^2}{b_\delta^2\sigma_b^2}\right).\end{aligned}\tag{7.77}$$

To derive the last line, we used the flat-sky approximation, letting the angular integral run over infinity. We also approximate the bandwidth as infinite. Thus,

$$\langle \tilde{\mathbf{C}}_{\delta\delta} \rangle_{i \in \delta} \approx \sigma_r^2 c B \sqrt{\pi} \sigma_b \left\langle \frac{1}{b_\delta} \exp\left(-\frac{c^2\tau^2}{b_\delta^2\sigma_b^2}\right) \right\rangle_{i \in \delta}.\tag{7.78}$$

Since the chromaticity increases monotonically with increasing baseline length and the antennas with the largest numbers of short baselines are at the center of the array, the minimal $k_{\parallel}^{\text{min}}$ accessible by an interferometer will occur on a short baseline formed from two antennas near the core of the array. With the core antenna positions equal to $r_i \approx 0$ so that $b_\delta = |\mathbf{r}_i - \mathbf{r}_k| \approx r_k$, the average of a function of the length of baselines that a core antenna participates in is equal to the average of that function over antenna positions

$$\langle \tilde{\mathbf{C}}_{\delta\delta} \rangle_{i \in \delta} \approx \sigma_r^2 c B \sqrt{\pi} \sigma_b \left\langle \frac{1}{r_k} \exp\left(-\frac{c^2\tau^2}{r_k^2\sigma_b^2}\right) \right\rangle_{r_k}.\tag{7.79}$$

We can compute this average analytically if the antennas are distributed as a Gaussian

with standard deviation σ_{ant} .

$$\begin{aligned} \left\langle \frac{1}{r_k} \exp\left(-\frac{c^2 \tau^2}{r_k^2 \sigma_b^2}\right) \right\rangle_{r_k} &= \frac{1}{2\pi \sigma_{\text{ant}}^2} \int d^2 r_k r_k^{-1} e^{-\frac{c^2 \tau^2}{r_k^2 \sigma_b^2} - \frac{r_k^2}{2\sigma_{\text{ant}}^2}} \\ &= \sqrt{\frac{\pi}{2}} \frac{1}{\sigma_{\text{ant}}} \exp\left(-\frac{\sqrt{2} c \tau}{\sigma_b \sigma_{\text{ant}}}\right). \end{aligned} \quad (7.80)$$

It follows, that for i and j antennas close to the core, the averages over baselines evaluate to

$$\langle \tilde{C}_{\delta\delta} \rangle_{i \in \delta} \approx \langle \tilde{C}_{\delta\delta} \rangle_{j \in \delta} \approx \sigma_r^2 c B \frac{\pi}{\sqrt{2}} \frac{\sigma_b}{\sigma_{\text{ant}}} \exp\left(-\frac{\sqrt{2} c \tau}{\sigma_b \sigma_{\text{ant}}}\right). \quad (7.81)$$

Thus, for two core antennas with Gaussian beams in a Gaussian antenna distribution, the contamination from calibration errors in equation 7.42 reduces to

$$P_{\alpha\alpha} - \tilde{C}_{\alpha\alpha} \approx 2\sigma_r^2 B \frac{c}{N_{\text{ant}}} \frac{\pi}{\sqrt{2}} \frac{\sigma_b}{\sigma_{\text{ant}}} \exp\left(-\frac{\sqrt{2} c \tau}{\sigma_b \sigma_{\text{ant}}}\right) \quad (7.82)$$

The minimal delay where the signal can be measured, τ_{min} , is set by where the calibration noise passes below the signal. Thus, we obtain τ_{min} by setting equation 7.82, multiplied by the prefactors in equation 7.26 that convert from $\text{Jy}^2 \text{Hz}^2$ to $\text{mK}^2 h^{-3} \text{Mpc}^3$ equal to the 21 cm power spectrum,

$$\begin{aligned} P_{21} &\approx \left(\frac{\lambda_0^2}{2k_B}\right)^2 \frac{X^2 Y}{B_{pp} \Omega_{pp}} \left(P_{\alpha\alpha}(\tau_{\text{min}}, \tau_{\text{min}}) - \tilde{C}_{\alpha\alpha}(\tau_{\text{min}}, \tau_{\text{min}})\right) \\ &\approx \left(\frac{\lambda_0^2}{2k_B}\right)^2 \frac{X^2 Y}{B_{pp} \Omega_{pp}} \sqrt{2\pi} \sigma_r^2 B \frac{c}{N_{\text{ant}}} \frac{\sigma_b}{\sigma_{\text{ant}}} \exp\left(-\frac{\sqrt{2} c \tau_{\text{min}}}{\sigma_b \sigma_{\text{ant}}}\right), \end{aligned} \quad (7.83)$$

and invert it.

$$\begin{aligned} \tau_{\text{min}} &\approx \frac{\sigma_b \sigma_{\text{ant}}}{\sqrt{2} c} \log\left(\frac{\lambda_0^4 X^2 Y \sqrt{2\pi} \sigma_r^2 c \sigma_b}{4k_B^2 P_{21} \Omega_{pp} B_{pp} N_{\text{ant}} \sigma_{\text{ant}}}\right) \\ &= \frac{\epsilon}{\sqrt{2} \nu_0} \frac{\sigma_{\text{ant}}}{d_{\text{ant}}} \log\left(\frac{\lambda_0^4 X^2 Y \sigma_r^2 B \nu_0}{B_{pp} 2\sqrt{2} k_B^2 N_{\text{ant}} \epsilon P_{21} \sigma_{\text{ant}}}\right) \end{aligned} \quad (7.84)$$

Using the fact that $k_{\parallel}^{\min} = 2\pi\tau_{\min}/Y$, we arrive at equation 7.51.

$$k_{\parallel}^{\min} = \frac{\epsilon\sqrt{2\pi}\sigma_{\text{ant}}}{Y\nu_0 d_{\text{ant}}} \log\left(\frac{\lambda_0^4 X^2 Y \sigma_r^2 B \nu_0 d_{\text{ant}}}{B_{pp} 2\sqrt{2} k_B^2 N_{\text{ant}} \epsilon P_{21} \sigma_{\text{ant}}}\right). \quad (7.85)$$

Checking this approximate expression against direct calculation for arrays with Gaussian beams using equation 7.34 yields an accuracy of $\approx 10\%$, even in arrays that are not strictly Gaussian such as the MWA, LOFAR, and HERA where σ_{ant} is the standard deviation of the non-Gaussian antenna distribution.

The primary shortcoming of equation 7.85 is that it assumes a Gaussian primary beam which only accounts for the delay at which the contamination from the main-lobe falls beneath the signal. Since side-lobes can easily enter at the $\gtrsim 5\%$ level, it is possible for them to contaminate the EoR window at much larger k_{\parallel} than the k_{\parallel}^{\min} predicted in equation 7.85. While the contribution of side-lobes for different baselines will fall at different delays and will not add coherently when averaging over the antenna distribution, we can assume that they add directly to obtain a conservative upper bound on when their contribution will affect k_{\parallel}^{\min} . If the side-lobes added directly in the antenna average, than their contribution to the amplitude of foreground-residuals would be on the order of f_{sl}^2 the level of the foreground residuals at zero delay, where f_{sl} is the ratio between the gain of the side-lobe and the gain at bore-sight. A conservative estimate of when side-lobes are at the level of the 21 cm signal can be obtained by setting the right hand side of equation 7.83 at zero-delay multiplied f_{sl}^2 equal to the 21 cm power spectrum.

$$\begin{aligned} P_{21} &\approx f_{sl}^2 \left(\frac{\lambda_0^2}{2k_B}\right)^2 \frac{X^2 Y}{B_{pp} \Omega_{pp}} \sqrt{2\pi} \sigma_r^2 B \frac{c}{N_{\text{ant}}} \frac{\sigma_b}{\sigma_{\text{ant}}} \\ &\approx f_{sl}^2 \left(\frac{\lambda_0^2}{2k_B}\right)^2 \frac{X^2 Y}{B_{pp} \pi} \sqrt{2\pi} \sigma_r^2 B \frac{c}{N_{\text{ant}}} \frac{d_{\text{ant}}}{\epsilon \sigma_{\text{ant}} \lambda_0} \end{aligned} \quad (7.86)$$

We use this condition to denote the white-dashed region of parameter space in Fig. 7-3 where side-lobes may render the predictions of equation 7.51 inaccurate.

Chapter 8

Conclusion

...but that's no matter – tomorrow we will run faster, stretch out our arms farther...

And then one fine morning –

F. SCOTT FITZGERALD

The Great Gatsby, 1925

Our investigation into the properties of errors in sky-based interferometric calibration brings us to the end of this thesis. A detection of 21 cm fluctuations and the ensuing cosmological revolution has not yet happened but significant progress has been made on our journey towards the light of the cosmic dawn. While obstacles have been encountered, none has been without a clear path forward and dramatic improvements in our upper limits with interesting constraints on reionization and reheating are tantalizingly close. Let us summarize what we have learned in this thesis.

In Chapter 2, we used semi-analytic models to determine the detectability of 21 cm fluctuations from before and during reionization over an unprecedented swath of parameter space, varying the efficiency with which baryons in stars are converted into X-ray photons and the minimal mass of dark matter halos contributing to X-ray heating and reionization. We find that existing and upcoming interferometers, such as the MWA and HERA, not only have the capacity to probe the reionization epoch that they were designed to explore, but are also capable of obtaining detections and scientifically interesting limits on the 21 cm fluctuations sourced by X-ray heating that are

most prominent before reionization. In Chapter 4, we determined that HERA, which is now being commissioned, can obtain $\sim 10\%$ constraints on the properties of the first X-ray sources in the Universe. HERA is capable of distinguishing heating scenarios dominated by high-mass X-ray binaries from heating from hot interstellar gas generated by the first supernovae. In addition, a potential inability to work within the highly contaminated FM band does not appear to be an insurmountable obstacle. Finally, we find that pre-reionization measurements will help to reduce uncertainties in ionization astrophysics by breaking degeneracies with heating astrophysics in the power spectrum. Our analysis in Chapter 3 teaches us that the relatively large optical depth of the cool gas before reionization may offer an opportunity for pre-reionization 21 cm power-spectrum measurements to constrain the population of radio-loud quasars as well. These chapters provide a strong argument for pursuing observations at high redshifts ($10 \lesssim z \lesssim 20$) that were previously overlooked for the current generation of experiments and HERA.

A number of systematic challenges in the power-spectrum measurement become worse at higher redshift (lower frequency). These include the fact that foreground intensities nearly double from ≈ 150 to 80 MHz, and the ionosphere's impact on radio propagation becomes more severe. Additional instrumental effects complicate pre-reionization power-spectrum measurements including the fact that the primary beam becomes wider for a fixed antenna geometry, enhancing high-delay horizon contamination and fixed baseline lengths sample larger angular scales, increasing diffuse contamination from our Galaxy. Exploring the impact of these systematics requires both end-to-end simulations and direct observations. We pursue the latter in Chapter 5 of this thesis. Using several hours of observations with the MWA, we find that at the level of the sensitivities reached there does not appear to be a correlation between the severity of ionospheric refraction and foreground contamination. We also discover that FM-band RFI at the Murchison Radio Observatory is sparse in both frequency and time, potentially allowing for a future detection of 21 cm fluctuations within the FM band. The ultimate obstacle that we encounter in our campaign are fine-scale frequency structures arising in the MWA's analogue signal path at the same LoS

scales as the cosmological signal. We determine that these frequency structures arise from reflections in the coaxial cables between the MWA’s antenna beamformers and receivers. These structures might be absorbed by increasing the frequency resolution of our calibration solutions. However when we attempt to increase the calibration degrees of freedom, we find increased systematics at different fine comoving scales. We are left attempting to model cable reflections with a small number of degrees of freedom and make some progress with information from the autocorrelations. Unfortunately, our attempts to model the frequency dependent reflection structures fall short and we are left with measurements limited by systematics that are a factor of a few above the thermal noise floor.

With this result, we are faced with two possible paths forward. 1) Understand why calibration is incapable of removing the fine-frequency structure in the gains. If possible, create a calibration procedure that avoids the existing shortcomings and/or 2) design future experiments so that their analogue signal paths are devoid of such structure. We make progress on path 1) in Chapter 6 by evaluating the design of the next-generation HERA antenna, finding that reflections between the parabolic reflecting dish (necessary to achieve HERA’s sensitivity gains) and the dipole feed do not pose a significant obstacle to its ability to do reionization science. Similar simulation work has been completed in designing a low-frequency HERA feed and a prototype is currently undergoing field testing. HERA will avoid introducing higher delay reflections by using RF over fiber in its signal path which has much smaller reflection coefficients. Thus, the problems that limited the MWA are not likely to hinder HERA.

Is a smooth instrumental bandpass actually necessary? why couldn’t sky-based calibration be employed to remove frequency structure? In Chapter 7, we perform a careful analysis of how errors in an incomplete model of the sky propagate into and interferometer’s gain solutions and into the measurement of the power-spectrum itself. What we find, is that the chromatic errors in the sky-model leak power from long-baselines (where the wedge extends foreground contamination to large delays) into the same large delays on short-baselines, spoiling the EoR window. It is for

this reason that we are not able to fit arbitrarily fine frequency degrees of freedom during calibration. While such fits correct the detrimental small-scale structure in the gains, they introduce new structure arising from the unmodeled components of the sky on long baselines. In comprehending the nature of this obstacle, we have also found a promising solution; to downweight long baselines in calibration itself. While this down-weighting leads to moderate increase in thermal noise, it successfully suppresses fine-frequency calibration errors. Thus, it may still be possible to measure the power spectrum of 21 cm fluctuations with existing experiments known to have non-smooth bandpasses (such as the MWA and LOFAR). Significant progress might also be achieved by averaging gain solutions in time.

We have obtained a clearer understanding of the systematic obstacles to a 21 cm detection and more importantly, that concrete solutions to these obstacles exist. We are implementing these fixes in new arrays such as HERA, upgrades to the MWA, and the SKA all of which have more than enough raw sensitivity to make a detection. The next several years could hardly be any more exciting.

...the universe is grand but life is grander...

LIU CHIXIN
Death's End, 2016

Bibliography

- K. N. Abazajian, K. Arnold, J. Austermann, B. A. Benson, C. Bischoff, J. Bock, J. R. Bond, J. Borrill, I. Buder, D. L. Burke, E. Calabrese, J. E. Carlstrom, C. S. Carvalho, C. L. Chang, H. C. Chiang, S. Church, A. Cooray, T. M. Crawford, B. P. Crill, K. S. Dawson, S. Das, M. J. Devlin, M. Dobbs, S. Dodelson, O. Doré, J. Dunkley, J. L. Feng, A. Fraisse, J. Gallicchio, S. B. Giddings, D. Green, N. W. Halverson, S. Hanany, D. Hanson, S. R. Hildebrandt, A. Hincks, R. Hlozek, G. Holder, W. L. Holzapfel, K. Honscheid, G. Horowitz, W. Hu, J. Hubmayr, K. Irwin, M. Jackson, W. C. Jones, R. Kallosh, M. Kamionkowski, B. Keating, R. Keisler, W. Kinney, L. Knox, E. Komatsu, J. Kovac, C.-L. Kuo, A. Kusaka, C. Lawrence, A. T. Lee, E. Leitch, A. Linde, E. Linder, P. Lubin, J. Maldacena, E. Martinec, J. McMahon, A. Miller, V. Mukhanov, L. Newburgh, M. D. Niemack, H. Nguyen, H. T. Nguyen, L. Page, C. Pryke, C. L. Reichardt, J. E. Ruhl, N. Sehgal, U. Seljak, L. Senatore, J. Sievers, E. Silverstein, A. Slosar, K. M. Smith, D. Spergel, S. T. Staggs, A. Stark, R. Stompor, A. G. Vieregg, G. Wang, S. Watson, E. J. Wollack, W. L. K. Wu, K. W. Yoon, O. Zahn, and M. Zaldarriaga. Inflation physics from the cosmic microwave background and large scale structure. *Astroparticle Physics*, 63:55–65, March 2015a. doi: 10.1016/j.astropartphys.2014.05.013.
- K. N. Abazajian, K. Arnold, J. Austermann, B. A. Benson, C. Bischoff, J. Bock, J. R. Bond, J. Borrill, E. Calabrese, J. E. Carlstrom, C. S. Carvalho, C. L. Chang, H. C. Chiang, S. Church, A. Cooray, T. M. Crawford, K. S. Dawson, S. Das, M. J. Devlin, M. Dobbs, S. Dodelson, O. Doré, J. Dunkley, J. Errard, A. Fraisse, J. Gallicchio, N. W. Halverson, S. Hanany, S. R. Hildebrandt, A. Hincks, R. Hlozek, G. Holder, W. L. Holzapfel, K. Honscheid, W. Hu, J. Hubmayr, K. Irwin, W. C. Jones, M. Kamionkowski, B. Keating, R. Keisler, L. Knox, E. Komatsu, J. Kovac, C.-L. Kuo, C. Lawrence, A. T. Lee, E. Leitch, E. Linder, P. Lubin, J. McMahon, A. Miller, L. Newburgh, M. D. Niemack, H. Nguyen, H. T. Nguyen, L. Page, C. Pryke, C. L. Reichardt, J. E. Ruhl, N. Sehgal, U. Seljak, J. Sievers, E. Silverstein, A. Slosar, K. M. Smith, D. Spergel, S. T. Staggs, A. Stark, R. Stompor, A. G. Vieregg, G. Wang, S. Watson, E. J. Wollack, W. L. K. Wu, K. W. Yoon, and O. Zahn. Neutrino physics from the cosmic microwave background and large scale structure. *Astroparticle Physics*, 63:66–80, March 2015b. doi: 10.1016/j.astropartphys.2014.05.014.
- K. N. Abazajian, P. Adshead, Z. Ahmed, S. W. Allen, D. Alonso, K. S. Arnold, C. Baccigalupi, J. G. Bartlett, N. Battaglia, B. A. Benson, C. A. Bischoff, J. Borrill,

V. Buza, E. Calabrese, R. Caldwell, J. E. Carlstrom, C. L. Chang, T. M. Crawford, F.-Y. Cyr-Racine, F. De Bernardis, T. de Haan, S. di Serego Alighieri, J. Dunkley, C. Dvorkin, J. Errard, G. Fabbian, S. Feeney, S. Ferraro, J. P. Filippini, R. Flauger, G. M. Fuller, V. Gluscevic, D. Green, D. Grin, E. Grohs, J. W. Henning, J. C. Hill, R. Hlozek, G. Holder, W. Holzappel, W. Hu, K. M. Huffenberger, R. Keskitalo, L. Knox, A. Kosowsky, J. Kovac, E. D. Kovetz, C.-L. Kuo, A. Kusaka, M. Le Jeune, A. T. Lee, M. Lilley, M. Loverde, M. S. Madhavacheril, A. Mantz, D. J. E. Marsh, J. McMahon, P. D. Meerburg, J. Meyers, A. D. Miller, J. B. Munoz, H. N. Nguyen, M. D. Niemack, M. Peloso, J. Peloton, L. Pogosian, C. Pryke, M. Raveri, C. L. Reichardt, G. Rocha, A. Rotti, E. Schaan, M. M. Schmittfull, D. Scott, N. Sehgal, S. Shandera, B. D. Sherwin, T. L. Smith, L. Sorbo, G. D. Starkman, K. T. Story, A. van Engelen, J. D. Vieira, S. Watson, N. Whitehorn, and W. L. Kimmy Wu. CMB-S4 Science Book, First Edition. *ArXiv e-prints*, October 2016.

A. A. Abdo, M. Ackermann, M. Ajello, W. B. Atwood, M. Axelsson, L. Baldini, J. Ballet, G. Barbiellini, D. Bastieri, M. Battelino, B. M. Baughman, K. Bechtol, R. Bellazzini, B. Berenji, R. D. Blandford, E. D. Bloom, G. Bogaert, E. Bonamente, A. W. Borgland, J. Bregeon, A. Brez, M. Brigida, P. Bruel, T. H. Burnett, G. A. Caliandro, R. A. Cameron, P. A. Caraveo, P. Carlson, J. M. Casandjian, C. Cecchi, E. Charles, A. Chekhtman, C. C. Cheung, J. Chiang, S. Ciprini, R. Claus, J. Cohen-Tanugi, L. R. Cominsky, J. Conrad, S. Cutini, C. D. Dermer, A. de Angelis, F. de Palma, S. W. Digel, G. di Bernardo, E. Do Couto E Silva, P. S. Drell, R. Dubois, D. Dumora, Y. Edmonds, C. Farnier, C. Favuzzi, W. B. Focke, M. Frailis, Y. Fukazawa, S. Funk, P. Fusco, D. Gaggero, F. Gargano, D. Gasparini, N. Gehrels, S. Germani, B. Giebels, N. Giglietto, F. Giordano, T. Glanzman, G. Godfrey, D. Grasso, I. A. Grenier, M.-H. Grondin, J. E. Grove, L. Guillemot, S. Guiriec, Y. Hanabata, A. K. Harding, R. C. Hartman, M. Hayashida, E. Hays, R. E. Hughes, G. Jóhannesson, A. S. Johnson, R. P. Johnson, W. N. Johnson, T. Kamae, H. Katagiri, J. Kataoka, N. Kawai, M. Kerr, J. Knödlseeder, D. Kocevski, F. Kuehn, M. Kuss, J. Lande, L. Latronico, M. Lemoine-Goumard, F. Longo, F. Loparco, B. Lott, M. N. Lovellette, P. Lubrano, G. M. Madejski, A. Makeev, M. M. Massai, M. N. Mazziotta, W. McConville, J. E. McEnery, C. Meurer, P. F. Michelson, W. Mitthumsiri, T. Mizuno, A. A. Moiseev, C. Monte, M. E. Monzani, E. Moretti, A. Morselli, I. V. Moskalenko, S. Murgia, P. L. Nolan, J. P. Norris, E. Nuss, T. Ohsugi, N. Omodei, E. Orlando, J. F. Ormes, M. Ozaki, D. Paneque, J. H. Panetta, D. Parent, V. Pelassa, M. Pepe, M. Pesce-Rollins, F. Piron, M. Pohl, T. A. Porter, S. Profumo, S. Rainò, R. Rando, M. Razzano, A. Reimer, O. Reimer, T. Reposeur, S. Ritz, L. S. Rochester, A. Y. Rodriguez, R. W. Romani, M. Roth, F. Ryde, H. F.-W. Sadrozinski, D. Sanchez, A. Sander, P. M. Saz Parkinson, J. D. Scargle, T. L. Schalk, A. Sellerholm, C. Sgrò, D. A. Smith, P. D. Smith, G. Spandre, P. Spinelli, J.-L. Starck, T. E. Stephens, M. S. Strickman, A. W. Strong, D. J. Suson, H. Tajima, H. Takahashi, T. Takahashi, T. Tanaka, J. B. Thayer, J. G. Thayer, D. J. Thompson, L. Tibaldo, O. Tibolla, D. F. Torres, G. Tosti, A. Tramacere, Y. Uchiyama, T. L. Usher, A. van Etten, V. Vasileiou, N. Vilchez, V. Vitale, A. P. Waite, E. Wallace, P. Wang, B. L. Winer, K. S. Wood, T. Ylinen,

and M. Ziegler. Measurement of the Cosmic Ray e^+e^- Spectrum from 20GeV to 1TeV with the Fermi Large Area Telescope. *Physical Review Letters*, 102(18): 181101, May 2009. doi: 10.1103/PhysRevLett.102.181101.

T. Abel, P. Anninos, Y. Zhang, and M. L. Norman. Modeling primordial gas in numerical cosmology. *New A*, 2:181–207, August 1997. doi: 10.1016/S1384-1076(97)00010-9.

T. Abel, G. L. Bryan, and M. L. Norman. The Formation and Fragmentation of Primordial Molecular Clouds. *ApJ*, 540:39–44, September 2000. doi: 10.1086/309295.

T. Abel, G. L. Bryan, and M. L. Norman. The Formation of the First Star in the Universe. *Science*, 295:93–98, January 2002.

L. Accardo, M. Aguilar, D. Aisa, B. Alpat, A. Alvino, G. Ambrosi, K. Andeen, L. Arruda, N. Attig, P. Azzarello, A. Bachlechner, F. Barao, A. Barrau, L. Barrin, A. Bartoloni, L. Basara, M. Battarbee, R. Battiston, J. Bazo, U. Becker, M. Behlmann, B. Beischer, J. Berdugo, B. Bertucci, G. Bigongiari, V. Bindi, S. Bizzaglia, M. Bizzarri, G. Boella, W. de Boer, K. Bollweg, V. Bonnivard, B. Borgia, S. Borsini, M. J. Boschini, M. Bourquin, J. Burger, F. Cadoux, X. D. Cai, M. Capell, S. Caroff, G. Carosi, J. Casaus, V. Cascioli, G. Castellini, I. Cernuda, D. Cerreta, F. Cervelli, M. J. Chae, Y. H. Chang, A. I. Chen, H. Chen, G. M. Cheng, H. S. Chen, L. Cheng, A. Chikanian, H. Y. Chou, E. Choumilov, V. Choutko, C. H. Chung, F. Cindolo, C. Clark, R. Clavero, G. Coignet, C. Consolandi, A. Contin, C. Corti, B. Coste, Z. Cui, M. Dai, C. Delgado, S. Della Torre, M. B. Demirköz, L. Derome, S. Di Falco, L. Di Masso, F. Dimiccoli, C. Díaz, P. von Doetinchem, W. J. Du, M. Duranti, D. D’Urso, A. Eline, F. J. Eppling, T. Eronen, Y. Y. Fan, L. Farnesini, J. Feng, E. Fiandrini, A. Fiasson, E. Finch, P. Fisher, Y. Galaktionov, G. Gallucci, B. García, R. García-López, H. Gast, I. Gebauer, M. Gervasi, A. Ghelfi, W. Gillard, F. Giovacchini, P. Goglov, J. Gong, C. Goy, V. Grabski, D. Grandi, M. Graziani, C. Guandalini, I. Guerri, K. H. Guo, D. Haas, M. Habiby, S. Haino, K. C. Han, Z. H. He, M. Heil, R. Henning, J. Hoffman, T. H. Hsieh, Z. C. Huang, C. Huh, M. Incagli, M. Ionica, W. Y. Jang, H. Jinchi, K. Kanishev, G. N. Kim, K. S. Kim, Th. Kirn, R. Kossakowski, O. Kounina, A. Kounine, V. Koutsenko, M. S. Krafczyk, S. Kunz, G. La Vacca, E. Laudi, G. Laurenti, I. Lazzizzera, A. Lebedev, H. T. Lee, S. C. Lee, C. Leluc, G. Levi, H. L. Li, J. Q. Li, Q. Li, Q. Li, T. X. Li, W. Li, Y. Li, Z. H. Li, Z. Y. Li, S. Lim, C. H. Lin, P. Lipari, T. Lippert, D. Liu, H. Liu, M. Lolli, T. Lomtadze, M. J. Lu, Y. S. Lu, K. Luebelmeyer, F. Luo, J. Z. Luo, S. S. Lv, R. Majka, A. Malinin, C. Mañá, J. Marín, T. Martin, G. Martínez, N. Masi, F. Massera, D. Maurin, A. Menchaca-Rocha, Q. Meng, D. C. Mo, B. Monreal, L. Morescalchi, P. Mott, M. Müller, J. Q. Ni, N. Nikonov, F. Nozoli, P. Nunes, A. Obermeier, A. Oliva, M. Orcinha, F. Palmonari, C. Palomares, M. Paniccia, A. Papi, M. Pauluzzi, E. Pedreschi, S. Pensotti, R. Pereira, R. Pilatrini, F. Pilo, A. Piluso, C. Pizzolotto, V. Plyaskin, M. Pohl, V. Poireau, E. Postaci,

- A. Putze, L. Quadrani, X. M. Qi, P. G. Rancoita, D. Rapin, J. S. Ricol, I. Rodríguez, S. Rosier-Lees, L. Rossi, A. Rozhkov, D. Rozza, G. Rybka, R. Sagdeev, J. Sandweiss, P. Saouter, C. Sbarra, S. Schael, S. M. Schmidt, D. Schuckardt, A. Schulz von Dratzig, G. Schwering, G. Scolieri, E. S. Seo, B. S. Shan, Y. H. Shan, J. Y. Shi, X. Y. Shi, Y. M. Shi, T. Siedenburger, D. Son, F. Spada, F. Spinella, W. Sun, W. H. Sun, M. Tacconi, C. P. Tang, X. W. Tang, Z. C. Tang, L. Tao, D. Tescaro, Samuel C. C. Ting, S. M. Ting, N. Tomassetti, J. Torsti, C. Türkoğlu, T. Urban, V. Vagelli, E. Valente, C. Vannini, E. Valtonen, S. Vaurynovich, M. Vecchi, M. Velasco, J. P. Vialle, V. Vitale, G. Volpini, L. Q. Wang, Q. L. Wang, R. S. Wang, X. Wang, Z. X. Wang, Z. L. Weng, K. Whitman, J. Wienkenhöver, H. Wu, K. Y. Wu, X. Xia, M. Xie, S. Xie, R. Q. Xiong, G. M. Xin, N. S. Xu, W. Xu, Q. Yan, J. Yang, M. Yang, Q. H. Ye, H. Yi, Y. J. Yu, Z. Q. Yu, S. Zeissler, J. H. Zhang, M. T. Zhang, X. B. Zhang, Z. Zhang, Z. M. Zheng, F. Zhou, H. L. Zhuang, V. Zhukov, A. Zichichi, N. Zimmermann, P. Zucco, and C. Zurbach. High statistics measurement of the positron fraction in primary cosmic rays of 0.5–500 gev with the alpha magnetic spectrometer on the international space station. *Phys. Rev. Lett.*, 113:121101, Sep 2014. doi: 10.1103/PhysRevLett.113.121101. URL <https://link.aps.org/doi/10.1103/PhysRevLett.113.121101>.
- N. Aghanim, S. Majumdar, and J. Silk. Secondary anisotropies of the CMB. *Reports on Progress in Physics*, 71(6):066902, June 2008. doi: 10.1088/0034-4885/71/6/066902.
- A. Albrecht and P. J. Steinhardt. Cosmology for grand unified theories with radiatively induced symmetry breaking. *Physical Review Letters*, 48:1220–1223, April 1982. doi: 10.1103/PhysRevLett.48.1220.
- C. Alcock, R. A. Allsman, D. R. Alves, T. S. Axelrod, A. C. Becker, D. P. Bennett, K. H. Cook, N. Dalal, A. J. Drake, K. C. Freeman, M. Geha, K. Griest, M. J. Lehner, S. L. Marshall, D. Minniti, C. A. Nelson, B. A. Peterson, P. Popowski, M. R. Pratt, P. J. Quinn, C. W. Stubbs, W. Sutherland, A. B. Tomaney, T. Van dehei, and D. Welch. The MACHO Project: Microlensing Results from 5.7 Years of Large Magellanic Cloud Observations. *ApJ*, 542:281–307, October 2000. doi: 10.1086/309512.
- Z. S. Ali, A. R. Parsons, H. Zheng, J. C. Pober, A. Liu, J. E. Aguirre, R. F. Bradley, G. Bernardi, C. L. Carilli, C. Cheng, D. R. DeBoer, M. R. Dexter, J. Grobbelaar, J. Horrell, D. C. Jacobs, P. Klima, D. H. E. MacMahon, M. Maree, D. F. Moore, N. Razavi, I. I. Stefan, W. P. Walbrugh, and A. Walker. PAPER-64 Constraints on Reionization: The 21 cm Power Spectrum at $z = 8.4$. *ApJ*, 809:61, August 2015. doi: 10.1088/0004-637X/809/1/61.
- R. A. Alpher, H. Bethe, and G. Gamow. The Origin of Chemical Elements. *Physical Review*, 73:803–804, April 1948. doi: 10.1103/PhysRev.73.803.
- L. Anderson, E. Aubourg, S. Bailey, D. Bizyaev, M. Blanton, A. S. Bolton, J. Brinkmann, J. R. Brownstein, A. Burden, A. J. Cuesta, L. A. N. da Costa,

- K. S. Dawson, R. de Putter, D. J. Eisenstein, J. E. Gunn, H. Guo, J.-C. Hamilton, P. Harding, S. Ho, K. Honscheid, E. Kazin, D. Kirkby, J.-P. Kneib, A. Labatie, C. Loomis, R. H. Lupton, E. Malanushenko, V. Malanushenko, R. Mandelbaum, M. Manera, C. Maraston, C. K. McBride, K. T. Mehta, O. Mena, F. Montesano, D. Muna, R. C. Nichol, S. E. Nuza, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, J. Parejko, I. Pâris, W. J. Percival, P. Petitjean, F. Prada, B. Reid, N. A. Roe, A. J. Ross, N. P. Ross, L. Samushia, A. G. Sánchez, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, R. A. Skibba, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, L. Verde, C. Wagner, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, X. Xu, C. Yèche, I. Zehavi, and G.-B. Zhao. The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Release 9 spectroscopic galaxy sample. *MNRAS*, 427:3435–3467, December 2012. doi: 10.1111/j.1365-2966.2012.22066.x.
- B. S. Arora, J. Morgan, S. M. Ord, S. J. Tingay, N. Hurley-Walker, M. Bell, G. Bernardi, N. D. R. Bhat, F. Briggs, J. R. Callingham, A. A. Deshpande, K. S. Dwarakanath, A. Ewall-Wice, L. Feng, B.-Q. For, P. Hancock, B. J. Hazelton, L. Hindson, D. Jacobs, M. Johnston-Hollitt, A. D. Kapińska, N. Kudryavtseva, E. Lenc, B. McKinley, D. Mitchell, D. Oberoi, A. R. Offringa, B. Pindor, P. Procopio, J. Riding, L. Staveley-Smith, R. B. Wayth, C. Wu, Q. Zheng, J. D. Bowman, R. J. Cappallo, B. E. Corey, D. Emrich, R. Goeke, L. J. Greenhill, D. L. Kaplan, J. C. Kasper, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, M. F. Morales, E. Morgan, T. Prabu, A. E. E. Rogers, A. Roshi, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, and C. L. Williams. Ionospheric Modelling using GPS to Calibrate the MWA. I: Comparison of First Order Ionospheric Effects between GPS Models and MWA Observations. *PASA*, 32:e029, August 2015. doi: 10.1017/pasa.2015.29.
- K. S. Arraki, A. Klypin, S. More, and S. Trujillo-Gomez. Effects of baryon removal on the structure of dwarf spheroidal galaxies. *MNRAS*, 438:1466–1482, February 2014. doi: 10.1093/mnras/stt2279.
- K. M. B. Asad, L. V. E. Koopmans, V. Jelić, V. N. Pandey, A. Ghosh, F. B. Abdalla, G. Bernardi, M. A. Brentjens, A. G. de Bruyn, S. Bus, B. Ciardi, E. Chapman, S. Daiboo, E. R. Fernandez, G. Harker, I. T. Iliev, H. Jensen, O. Martinez-Rubi, G. Mellema, M. Mevius, A. R. Offringa, A. H. Patil, J. Schaye, R. M. Thomas, S. van der Tol, H. K. Vedantham, S. Yatawatta, and S. Zaroubi. Polarization leakage in epoch of reionization windows - I. Low Frequency Array observations of the 3C196 field. *MNRAS*, 451:3709–3727, August 2015. doi: 10.1093/mnras/stv1107.
- Donald Backer, James Aguirre, Judd Bowman, Richard Bradley, C Carilli, S Furlanetto, L J Greenhill, Jacqueline N Hewitt, C J Lonsdale, S M Ord, Aaron Parsons, and A R Whitney. HERA – Hydrogen Epoch of Reionization Arrays. pages 1–46, July 2009.

- S. Baek, B. Semelin, P. Di Matteo, Y. Revaz, and F. Combes. Reionization by UV or X-ray sources. *A&A*, 523:A4, November 2010. doi: 10.1051/0004-6361/201014347.
- K. Bandura, G. E. Addison, M. Amiri, J. R. Bond, D. Campbell-Wilson, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, M. Dobbs, M. Fandino, K. Gibbs, A. Gilbert, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Höfer, P. Klages, T. L. Landecker, K. Masui, J. Mena Parra, L. B. Newburgh, U.-l. Pen, J. B. Peterson, A. Recnik, J. R. Shaw, K. Sigurdson, M. Sitwell, G. Smecher, R. Smegal, K. Vanderlinde, and D. Wiebe. Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder. In *Ground-based and Airborne Telescopes V*, volume 9145 of Proc. SPIE, page 914522, July 2014. doi: 10.1117/12.2054950.
- R. Barkana. The rise of the first stars: Supersonic streaming, radiative feedback, and 21-cm cosmology. *Phys. Rep.*, 645:1–59, July 2016. doi: 10.1016/j.physrep.2016.06.006.
- R. Barkana and A. Loeb. In the beginning: the first sources of light and the reionization of the universe. *Phys. Rep.*, 349:125–238, July 2001. doi: 10.1016/S0370-1573(01)00019-9.
- R. Barkana and A. Loeb. Unusually Large Fluctuations in the Statistics of Galaxy Formation at High Redshift. *ApJ*, 609:474–481, July 2004. doi: 10.1086/421079.
- R. Barkana and A. Loeb. A Method for Separating the Physics from the Astrophysics of High-Redshift 21 Centimeter Fluctuations. *ApJ*, 624:L65–L68, May 2005a. doi: 10.1086/430599.
- R. Barkana and A. Loeb. Detecting the Earliest Galaxies through Two New Sources of 21 Centimeter Fluctuations. *ApJ*, 626:1–11, June 2005b. doi: 10.1086/429954.
- N. Barry, B. Hazelton, I. Sullivan, M. F. Morales, and J. C. Pober. Calibration requirements for detecting the 21 cm epoch of reionization power spectrum and implications for the SKA. *MNRAS*, 461:3135–3144, September 2016. doi: 10.1093/mnras/stw1380.
- J. Bartels, N. H. Heck, and H. F. Johnston. The three-hour-range index measuring geomagnetic activity. *Terrestrial Magnetism and Atmospheric Electricity (Journal of Geophysical Research)*, 44:411, 1939. doi: 10.1029/TE044i004p00411.
- A. R. Basu-Zych et al. The X-Ray Star Formation Story as Told by Lyman Break Galaxies in the 4 Ms CDF-S. *ApJ*, 762:45, January 2013. doi: 10.1088/0004-637X/762/1/45.
- A. P. Beardsley, B. J. Hazelton, I. S. Sullivan, P. Carroll, N. Barry, M. Rahimi, B. Pindor, C. M. Trott, J. Line, D. C. Jacobs, M. F. Morales, J. C. Pober, G. Bernardi, J. D. Bowman, M. P. Busch, F. Briggs, R. J. Cappallo, B. E. Corey, A. de Oliveira-Costa, J. S. Dillon, D. Emrich, A. Ewall-Wice, L. Feng, B. M. Gaensler, R. Goeke,

- L. J. Greenhill, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H. S. Kim, E. Kratzenberg, E. Lenc, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, E. Morgan, A. R. Neben, N. Thyagarajan, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Rosh, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, M. Tegmark, S. J. Tingay, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. First Season MWA EoR Power spectrum Results at Redshift 7. *ApJ*, 833:102, December 2016. doi: 10.3847/1538-4357/833/1/102.
- A P Beardsley et al. A new layout optimization technique for interferometric arrays, applied to the Murchison Widefield Array. *Monthly Notices of the Royal Astronomical Society*, 425(3):1781–1788, 2012.
- A. P. Beardsley et al. The EoR sensitivity of the Murchison Widefield Array. *MNRAS*, 429:L5–L9, February 2013. doi: 10.1093/mnras/sls013.
- C. L. Bennett, A. J. Banday, K. M. Gorski, G. Hinshaw, P. Jackson, P. Keegstra, A. Kogut, G. F. Smoot, D. T. Wilkinson, and E. L. Wright. Four-Year COBE DMR Cosmic Microwave Background Observations: Maps and Basic Results. *ApJ*, 464:L1, June 1996. doi: 10.1086/310075.
- A. J. Benson, C. S. Frenk, C. G. Lacey, C. M. Baugh, and S. Cole. The effects of photoionization on galaxy formation - II. Satellite galaxies in the Local Group. *MNRAS*, 333:177–190, June 2002. doi: 10.1046/j.1365-8711.2002.05388.x.
- G. Bernardi, A. G. de Bruyn, M. A. Brentjens, B. Ciardi, G. Harker, V. Jelić, L. V. E. Koopmans, P. Labropoulos, A. Offringa, V. N. Pandey, J. Schaye, R. M. Thomas, S. Yatawatta, and S. Zaroubi. Foregrounds for observations of the cosmological 21 cm line. I. First Westerbork measurements of Galactic emission at 150 MHz in a low latitude field. *A&A*, 500:965–979, June 2009. doi: 10.1051/0004-6361/200911627.
- G. Bernardi, L. J. Greenhill, D. A. Mitchell, S. M. Ord, B. J. Hazelton, B. M. Gaensler, A. de Oliveira-Costa, M. F. Morales, N. Udaya Shankar, R. Subrahmanyam, R. B. Wayth, E. Lenc, C. L. Williams, W. Arcus, B. S. Arora, D. G. Barnes, J. D. Bowman, F. H. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. Deshpande, L. deSouza, D. Emrich, R. Goeke, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, D. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, E. Morgan, D. Oberoi, J. Pathikulangara, T. Prabu, R. A. Remillard, A. E. E. Rogers, A. Rosh, J. E. Salah, R. J. Sault, K. S. Srivani, J. Stevens, S. J. Tingay, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, and J. S. B. Wyithe. A 189 MHz, 2400 deg² Polarization Survey with the Murchison Widefield Array 32-element Prototype. *ApJ*, 771:105, July 2013. doi: 10.1088/0004-637X/771/2/105.
- G. Bernardi, M. McQuinn, and L. J. Greenhill. Foreground Model and Antenna Calibration Errors in the Measurement of the Sky-averaged λ_{21} cm Signal at $z \sim 20$. *ApJ*, 799:90, January 2015. doi: 10.1088/0004-637X/799/1/90.

- D. C.-J. Bock, M. I. Large, and E. M. Sadler. SUMSS: A Wide-Field Radio Imaging Survey of the Southern Sky. I. Science Goals, Survey Design, and Instrumentation. *AJ*, 117:1578–1593, March 1999. doi: 10.1086/300786.
- P. Bode, J. P. Ostriker, and N. Turok. Halo Formation in Warm Dark Matter Models. *ApJ*, 556:93–107, July 2001. doi: 10.1086/321541.
- J. R. Bond, S. Cole, G. Efstathiou, and N. Kaiser. Excursion set mass functions for hierarchical Gaussian fluctuations. *ApJ*, 379:440–460, October 1991. doi: 10.1086/170520.
- H. Bondi and T. Gold. The Steady-State Theory of the Expanding Universe. *MNRAS*, 108:252, 1948. doi: 10.1093/mnras/108.3.252.
- J. D. Bowman and A. E. E. Rogers. A lower limit of $\Delta z > 0.06$ for the duration of the reionization epoch. *Nature*, 468:796–798, December 2010. doi: 10.1038/nature09601.
- J. D. Bowman, M. F. Morales, and J. N. Hewitt. The Sensitivity of First-Generation Epoch of Reionization Observatories and Their Potential for Differentiating Theoretical Power Spectra. *ApJ*, 638:20–26, February 2006. doi: 10.1086/498703.
- J. D. Bowman, M. F. Morales, and J. N. Hewitt. Constraints on Fundamental Cosmological Parameters with Upcoming Redshifted 21 cm Observations. *ApJ*, 661:1–9, May 2007. doi: 10.1086/516560.
- J. D. Bowman, M. F. Morales, and J. N. Hewitt. Foreground Contamination in Interferometric Measurements of the Redshifted 21 cm Power Spectrum. *ApJ*, 695:183–199, April 2009. doi: 10.1088/0004-637X/695/1/183.
- J. D. Bowman, I. Cairns, D. L. Kaplan, T. Murphy, D. Oberoi, L. Staveley-Smith, W. Arcus, D. G. Barnes, G. Bernardi, F. H. Briggs, S. Brown, J. D. Bunton, A. J. Burgasser, R. J. Cappallo, S. Chatterjee, B. E. Corey, A. Coster, A. Deshpande, L. deSouza, D. Emrich, P. Erickson, R. F. Goeke, B. M. Gaensler, L. J. Greenhill, L. Harvey-Smith, B. J. Hazelton, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, L. D. Matthews, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. H. Morgan, S. M. Ord, J. Pathikulangara, T. Prabu, R. A. Remillard, T. Robishaw, A. E. E. Rogers, A. A. Roshi, J. E. Salah, R. J. Sault, N. U. Shankar, K. S. Srivani, J. B. Stevens, R. Subrahmanyan, S. J. Tingay, R. B. Wayth, M. Waterson, R. L. Webster, A. R. Whitney, A. J. Williams, C. L. Williams, and J. S. B. Wyithe. Science with the Murchison Widefield Array. *PASA*, 30:e031, April 2013. doi: 10.1017/pas.2013.009.
- M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat. Too big to fail? The puzzling darkness of massive Milky Way subhaloes. *MNRAS*, 415:L40–L44, July 2011. doi: 10.1111/j.1745-3933.2011.01074.x.

- P. D. Bristow and S. Phillipps. On the Baryon Content of the Universe. *MNRAS*, 267:13, March 1994. doi: 10.1093/mnras/267.1.13.
- V. Bromm. *Star formation in the early universe*. PhD thesis, Yale University, 2000.
- V. Bromm and N. Yoshida. The First Galaxies. *ARA&A*, 49:373–407, September 2011. doi: 10.1146/annurev-astro-081710-102608.
- V. Bromm, P. S. Coppi, and R. B. Larson. Forming the First Stars in the Universe: The Fragmentation of Primordial Gas. *ApJ*, 527:L5–L8, December 1999. doi: 10.1086/312385.
- V. Bromm, P. S. Coppi, and R. B. Larson. The Formation of the First Stars. I. The Primordial Star-forming Cloud. *ApJ*, 564:23–51, January 2002.
- V. Bromm, N. Yoshida, and L. Hernquist. The First Supernova Explosions in the Universe. *ApJ*, 596:L135–L138, October 2003. doi: 10.1086/379359.
- A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper. A Baryonic Solution to the Missing Satellites Problem. *ApJ*, 765:22, March 2013. doi: 10.1088/0004-637X/765/1/22.
- E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall. Detection of an Unidentified Emission Line in the Stacked X-Ray Spectrum of Galaxy Clusters. *ApJ*, 789:13, July 2014. doi: 10.1088/0004-637X/789/1/13.
- J. S. Bullock, A. V. Kravtsov, and D. H. Weinberg. Reionization and the Abundance of Galactic Satellites. *ApJ*, 539:517–521, August 2000. doi: 10.1086/309279.
- E. F. Bunn, D. Scott, and M. White. The COBE normalization for standard cold dark matter. *ApJ*, 441:L9–L12, March 1995. doi: 10.1086/187776.
- J. O. Burns, J. Lazio, S. Bale, J. Bowman, R. Bradley, C. Carilli, S. Furlanetto, G. Harker, A. Loeb, and J. Pritchard. Probing the first stars and black holes in the early Universe with the Dark Ages Radio Explorer (DARE). *Advances in Space Research*, 49:433–450, February 2012. doi: 10.1016/j.asr.2011.10.014.
- C. L. Carilli, N. Y. Gnedin, and F. Owen. H I 21 Centimeter Absorption beyond the Epoch of Reionization. *ApJ*, 577:22–30, September 2002. doi: 10.1086/342179.
- C. L. Carilli, R. Wang, X. Fan, F. Walter, J. Kurk, D. Riechers, J. Wagg, J. Hennawi, L. Jiang, K. M. Menten, F. Bertoldi, M. A. Strauss, and P. Cox. Ionization Near Zones Associated with Quasars at $z \sim 6$. *ApJ*, 714:834–839, May 2010. doi: 10.1088/0004-637X/714/1/834.
- P. A. Carroll, J. Line, M. F. Morales, N. Barry, A. P. Beardsley, B. J. Hazelton, D. C. Jacobs, J. C. Pober, I. S. Sullivan, R. L. Webster, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, B. E. Corey, A. de Oliveira-Costa, J. S. Dillon, D. Emrich, A. Ewall-Wice, L. Feng, B. M. Gaensler, R. Goeke, L. J. Greenhill,

- J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H. Kim, E. Kratzenberg, E. Lenc, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, E. Morgan, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, N. U. Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, M. Tegmark, N. Thyagarajan, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. A high reliability survey of discrete Epoch of Reionization foreground sources in the MWA EoR0 field. *MNRAS*, 461:4151–4175, October 2016. doi: 10.1093/mnras/stw1599.
- J. Caruana, A. J. Bunker, S. M. Wilkins, E. R. Stanway, S. Lorenzoni, M. J. Jarvis, and H. Ebert. Spectroscopy of $z \sim 7$ candidate galaxies: using Lyman alpha to constrain the neutral fraction of hydrogen in the high-redshift universe. *MNRAS*, 443:2831–2842, October 2014. doi: 10.1093/mnras/stu1341.
- R. Cen and J. P. Ostriker. Where Are the Baryons? *ApJ*, 514:1–6, March 1999. doi: 10.1086/306949.
- R. Cen and J. P. Ostriker. Where Are the Baryons? II. Feedback Effects. *ApJ*, 650:560–572, October 2006. doi: 10.1086/506505.
- X. Chen and J. Miralda-Escudé. The Spin-Kinetic Temperature Coupling and the Heating Rate due to $\text{Ly}\alpha$ Scattering before Reionization: Predictions for 21 Centimeter Emission and Absorption. *ApJ*, 602:1–11, February 2004. doi: 10.1086/380829.
- X. Chen, C. Dvorkin, Z. Huang, M. H. Namjoo, and L. Verde. The future of primordial features with large-scale structure surveys. *J. Cosmology Astropart. Phys.*, 11:014, November 2016. doi: 10.1088/1475-7516/2016/11/014.
- R. Chornock, E. Berger, D. B. Fox, R. Lunnan, M. R. Drout, W.-f. Fong, T. Laskar, and K. C. Roth. GRB 130606A as a Probe of the Intergalactic Medium and the Interstellar Medium in a Star-forming Galaxy in the First Gyr after the Big Bang. *ApJ*, 774:26, September 2013. doi: 10.1088/0004-637X/774/1/26.
- T. R. Choudhury, A. Ferrara, and S. Gallerani. On the minimum mass of reionization sources. *MNRAS*, 385:L58–L62, March 2008. doi: 10.1111/j.1745-3933.2008.00433.x.
- P. Christian and A. Loeb. Measuring the X-ray background in the reionization era with first generation 21 cm experiments. *J. Cosmology Astropart. Phys.*, 9:014, September 2013. doi: 10.1088/1475-7516/2013/09/014.
- L. Chuzhoy. Impact of Dark Matter Annihilation on the High-Redshift Intergalactic Medium. *ApJ*, 679:L65–L68, June 2008. doi: 10.1086/589504.
- B. Ciardi, P. Labropoulos, A. Maselli, R. Thomas, S. Zaroubi, L. Graziani, J. S. Bolton, G. Bernardi, M. Brentjens, A. G. de Bruyn, S. Daiboo, G. J. A. Harker,

- V. Jelic, S. Kazemi, L. V. E. Koopmans, O. Martinez, G. Mellema, A. R. Offringa, V. N. Pandey, J. Schaye, V. Veligatla, H. Vedantham, and S. Yatawatta. Prospects for detecting the 21 cm forest from the diffuse intergalactic medium with LOFAR. *MNRAS*, 428:1755–1765, January 2013. doi: 10.1093/mnras/sts156.
- A. S. Cohen and H. J. A. Röttgering. Probing Fine-Scale Ionospheric Structure with the Very Large Array Radio Telescope. *AJ*, 138:439–447, August 2009. doi: 10.1088/0004-6256/138/2/439.
- T. M. Colegate, A. T. Sutinjo, P. J. Hall, S. K. Padhi, R. B. Wayth, J. G. Bij de Vaate, B. Crosse, D. Emrich, A. J. Faulkner, N. Hurley-Walker, E. de Lera Acedo, B. Juswardy, N. Razavi-Ghods, S. J. Tingay, and A. Williams. Antenna array characterization via radio interferometry observation of astronomical sources. *ArXiv e-prints*, January 2015.
- J. J. Condon. Confusion and Flux-Density Error Distributions. *ApJ*, 188:279–286, March 1974. doi: 10.1086/152714.
- J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick. The NRAO VLA Sky Survey. *AJ*, 115:1693–1716, May 1998. doi: 10.1086/300337.
- A. Cooray. Large-scale non-Gaussianities in the 21-cm background anisotropies from the era of reionization. *MNRAS*, 363:1049–1056, November 2005. doi: 10.1111/j.1365-2966.2005.09506.x.
- W. D. Cotton. a New Method for Cross Polarized Delay Calibration of Radio Interferometers. *Journal of Astronomical Instrumentation*, 1:1250001, December 2012. doi: 10.1142/S2251171712500018.
- L. L. Cowie, A. J. Barger, and G. Hasinger. The Faintest X-Ray Sources from $z = 0$ TO 8. *ApJ*, 748:50, March 2012. doi: 10.1088/0004-637X/748/1/50.
- R. A. C. Croft, D. H. Weinberg, M. Bolte, S. Burles, L. Hernquist, N. Katz, D. Kirkman, and D. Tytler. Toward a Precise Measurement of Matter Clustering: Ly α Forest Data at Redshifts 2-4. *ApJ*, 581:20–52, December 2002. doi: 10.1086/344099.
- A. V. da Rosa, H. Waldman, J. Bendito, and O. K. Garriott. Response of the ionospheric electron content to fluctuations in solar activity. *Journal of Atmospheric and Terrestrial Physics*, 35:1429–1442, 1973.
- A. Datta, J. D. Bowman, and C. L. Carilli. Bright Source Subtraction Requirements for Redshifted 21 cm Measurements. *ApJ*, 724:526–538, November 2010. doi: 10.1088/0004-637X/724/1/526.
- R. Davé, R. Cen, J. P. Ostriker, G. L. Bryan, L. Hernquist, N. Katz, D. H. Weinberg, M. L. Norman, and B. O’Shea. Baryons in the Warm-Hot Intergalactic Medium. *ApJ*, 552:473–483, May 2001. doi: 10.1086/320548.

- A. de Oliveira-Costa, M. Tegmark, B. M. Gaensler, J. Jonas, T. L. Landecker, and P. Reich. A model of diffuse Galactic radio emission from 10 MHz to 100 GHz. *MNRAS*, 388:247–260, July 2008. doi: 10.1111/j.1365-2966.2008.13376.x.
- R. S. de Souza, A. Mesinger, A. Ferrara, Z. Haiman, R. Perna, and N. Yoshida. Constraints on warm dark matter models from high-redshift long gamma-ray bursts. *MNRAS*, 432:3218–3227, July 2013. doi: 10.1093/mnras/stt674.
- G. de Zotti, M. Massardi, M. Negrello, and J. Wall. Radio and millimeter continuum surveys and their astrophysical implications. *A&A Rev.*, 18:1–65, February 2010. doi: 10.1007/s00159-009-0026-0.
- D. R. DeBoer, A. R. Parsons, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardley, G. Bernardi, J. D. Bowman, R. F. Bradley, C. L. Carilli, C. Cheng, E. de Lera Acedo, J. S. Dillon, A. Ewall-Wice, G. Fadana, N. Fagnoni, R. Fritz, S. R. Furlanetto, B. Glendenning, B. Greig, J. Grobbelaar, B. J. Hazelton, J. N. Hewitt, J. Hickish, D. C. Jacobs, A. Julius, M. Kariseb, S. A. Kohn, T. Lekalake, A. Liu, A. Loots, D. MacMahon, L. Malan, C. Malgas, M. Maree, N. Mathison, E. Matsetela, A. Mesinger, M. F. Morales, A. R. Neben, N. Patra, S. Pieterse, J. C. Pober, N. Razavi-Ghods, J. Ringuette, J. Robnett, K. Rosie, R. Sell, C. Smith, A. Syce, M. Tegmark, N. Thyagarajan, P. K. G. Williams, and H. Zheng. Hydrogen Epoch of Reionization Array (HERA). *ArXiv e-prints*, June 2016.
- C. DeGraf, T. Di Matteo, N. Khandai, and R. Croft. Growth of Early Supermassive Black Holes and the High-redshift Eddington Ratio Distribution. *ApJ*, 755:L8, August 2012. doi: 10.1088/2041-8205/755/1/L8.
- DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, and et al. The DESI Experiment Part I: Science, Targeting, and Survey Design. *ArXiv e-prints*, October 2016.
- P. Dewdney et al. SKA1 SYSTEM BASELINE DESIGN. *SKA1 SYSTEM BASELINE DESIGN*
http://www.skatelescope.org/wp-content/uploads/2013/03/SKA-TEL-SKO-DD-001-1_BaselineDesign1.pdf, pages 1–98, March 2013.
- T. Di Matteo, R. Perna, T. Abel, and M. J. Rees. Radio Foregrounds for the 21 Centimeter Tomography of the Neutral Intergalactic Medium at High Redshifts. *ApJ*, 564:576–580, January 2002. doi: 10.1086/324293.
- J. S. Dillon and A. R. Parsons. Redundant Array Configurations for 21 cm Cosmology. *ApJ*, 826:181, August 2016. doi: 10.3847/0004-637X/826/2/181.
- J. S. Dillon, A. Liu, and M. Tegmark. A fast method for power spectrum and foreground analysis for 21 cm cosmology. *Phys. Rev. D*, 87(4):043005, February 2013. doi: 10.1103/PhysRevD.87.043005.

- J. S. Dillon, A. Liu, C. L. Williams, J. N. Hewitt, M. Tegmark, E. H. Morgan, A. M. Levine, M. F. Morales, S. J. Tingay, G. Bernardi, J. D. Bowman, F. H. Briggs, R. C. Cappallo, D. Emrich, D. A. Mitchell, D. Oberoi, T. Prabu, R. Wayth, and R. L. Webster. Overcoming real-world obstacles in 21 cm power spectrum estimation: A method demonstration and results from early Murchison Widefield Array data. *Phys. Rev. D*, 89(2):023002, January 2014. doi: 10.1103/PhysRevD.89.023002.
- J. S. Dillon, A. R. Neben, J. N. Hewitt, M. Tegmark, N. Barry, A. P. Beardsley, J. D. Bowman, F. Briggs, P. Carroll, A. de Oliveira-Costa, A. Ewall-Wice, L. Feng, L. J. Greenhill, B. J. Hazelton, L. Hernquist, N. Hurley-Walker, D. C. Jacobs, H. S. Kim, P. Kittiwisit, E. Lenc, J. Line, A. Loeb, B. McKinley, D. A. Mitchell, M. F. Morales, A. R. Offringa, S. Paul, B. Pindor, J. C. Pober, P. Procopio, J. Riding, S. Sethi, N. U. Shankar, R. Subrahmanyan, I. Sullivan, N. Thyagarajan, S. J. Tingay, C. Trott, R. B. Wayth, R. L. Webster, S. Wyithe, G. Bernardi, R. J. Cappallo, A. A. Deshpande, M. Johnston-Hollitt, D. L. Kaplan, C. J. Lonsdale, S. R. McWhirter, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, K. S. Srivani, A. Williams, and C. L. Williams. Empirical covariance modeling for 21 cm power spectrum estimation: A method demonstration and new limits from early Murchison Widefield Array 128-tile data. *Phys. Rev. D*, 91(12):123011, June 2015a. doi: 10.1103/PhysRevD.91.123011.
- J. S. Dillon, M. Tegmark, A. Liu, A. Ewall-Wice, J. N. Hewitt, M. F. Morales, A. R. Neben, A. R. Parsons, and H. Zheng. Mapmaking for precision 21 cm cosmology. *Phys. Rev. D*, 91(2):023002, January 2015b. doi: 10.1103/PhysRevD.91.023002.
- D. J. Eisenstein and P. Hut. HOP: A New Group-Finding Algorithm for N-Body Simulations. *ApJ*, 498:137–+, May 1998. doi: 10.1086/305535.
- D. J. Eisenstein and A. Loeb. Origin of quasar progenitors from the collapse of low-spin cosmological perturbations. *ApJ*, 443:11–17, April 1995. doi: 10.1086/175498.
- D. J. Eisenstein, D. H. Weinberg, E. Agol, H. Aihara, C. Allende Prieto, S. F. Anderson, J. A. Arns, É. Aubourg, S. Bailey, E. Balbinot, and et al. SDSS-III: Massive Spectroscopic Surveys of the Distant Universe, the Milky Way, and Extra-Solar Planetary Systems. *AJ*, 142:72, September 2011. doi: 10.1088/0004-6256/142/3/72.
- R. Ekers and R. Subrahmanyan. Core, zebra, saras; the zero spacing interferometers. http://www.caastro.org/files/30/1496724333/ekers_eor.pdf, 2012.
- S. W. Ellingson, T. E. Clarke, A. Cohen, J. Craig, N. E. Kassim, Y. Pihlstrom, L. J. Rickard, and G. B. Taylor. The Long Wavelength Array. *IEEE Proceedings*, 97: 1421–1430, August 2009. doi: 10.1109/JPROC.2009.2015683.
- C. Evoli, A. Mesinger, and A. Ferrara. Unveiling the nature of dark matter with high redshift 21 cm line experiments. *J. Cosmology Astropart. Phys.*, 11:024, November 2014. doi: 10.1088/1475-7516/2014/11/024.

- A. Ewall-Wice, J. S. Dillon, A. Mesinger, and J. Hewitt. Detecting the 21 cm forest in the 21 cm power spectrum. *MNRAS*, 441:2476–2496, July 2014. doi: 10.1093/mnras/stu666.
- A. Ewall-Wice, R. Bradley, D. Deboer, J. Hewitt, A. Parsons, J. Aguirre, Z. S. Ali, J. Bowman, C. Cheng, A. R. Neben, N. Patra, N. Thyagarajan, M. Venter, E. de Lera Acedo, J. S. Dillon, R. Dickenson, P. Doolittle, D. Egan, M. Hedrick, P. Klima, S. Kohn, P. Schaffner, J. Shelton, B. Saliwanchik, H. A. Taylor, R. Taylor, M. Tegmark, and B. Wirt. The Hydrogen Epoch of Reionization Array Dish. II. Characterization of Spectral Structure with Electromagnetic Simulations and Its Science Implications. *ApJ*, 831:196, November 2016a. doi: 10.3847/0004-637X/831/2/196.
- A. Ewall-Wice, J. S. Dillon, J. N. Hewitt, A. Loeb, A. Mesinger, A. R. Neben, A. R. Offringa, M. Tegmark, N. Barry, A. P. Beardsley, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, D. Emrich, L. Feng, B. M. Gaensler, R. Goeke, L. J. Greenhill, B. J. Hazelton, N. Hurley-Walker, M. Johnston-Hollitt, D. C. Jacobs, D. L. Kaplan, J. C. Kasper, H. Kim, E. Kratzenberg, E. Lenc, J. Line, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, N. Thyagarajan, D. Oberoi, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, N. U. Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. First limits on the 21 cm power spectrum during the Epoch of X-ray heating. *MNRAS*, 460:4320–4347, August 2016b. doi: 10.1093/mnras/stw1022.
- A. Ewall-Wice, J. S. Dillon, A. Liu, and J. Hewitt. The Impact of Modeling Errors on Interferometer Calibration for 21 cm Power Spectra. *ArXiv e-prints*, October 2016c.
- A. Ewall-Wice, J. Hewitt, A. Mesinger, J. S. Dillon, A. Liu, and J. Pober. Constraining high-redshift X-ray sources with next generation 21-cm power spectrum measurements. *MNRAS*, 458:2710–2724, May 2016d. doi: 10.1093/mnras/stw452.
- H. I. Ewen and E. M. Purcell. Observation of a Line in the Galactic Radio Spectrum: Radiation from Galactic Hydrogen at 1,420 Mc./sec. *Nature*, 168:356, September 1951. doi: 10.1038/168356a0.
- N. Fagnoni and E. de Lera Acedo. The Hydrogen Epoch of Reionization Array (HERA). Improvement of the antenna response with a matching network and scientific impacts. *ArXiv e-prints*, June 2016.
- X. Fan, V. K. Narayanan, R. H. Lupton, M. A. Strauss, G. R. Knapp, R. H. Becker, R. L. White, L. Pentericci, S. K. Leggett, Z. Haiman, J. E. Gunn, Z. Ivezić, D. P. Schneider, S. F. Anderson, J. Brinkmann, N. A. Bahcall, A. J. Connolly, I. Csabai, M. Doi, M. Fukugita, T. Geballe, E. K. Grebel, D. Harbeck, G. Hennessy, D. Q.

- Lamb, G. Miknaitis, J. A. Munn, R. Nichol, S. Okamura, J. R. Pier, F. Prada, G. T. Richards, A. Szalay, and D. G. York. A Survey of $z > 5.8$ Quasars in the Sloan Digital Sky Survey. I. Discovery of Three New Quasars and the Spatial Density of Luminous Quasars at $z \sim 6$. *AJ*, 122:2833–2849, December 2001. doi: 10.1086/324111.
- C.-A. Faucher-Giguère, A. Lidz, M. Zaldarriaga, and L. Hernquist. A New Calculation of the Ionizing Background Spectrum and the Effects of He II Reionization. *ApJ*, 703:1416–1443, October 2009. doi: 10.1088/0004-637X/703/2/1416.
- A. Fialkov and A. Loeb. Precise Measurement of the Reionization Optical Depth from the Global 21 cm Signal Accounting for Cosmic Heating. *ApJ*, 821:59, April 2016. doi: 10.3847/0004-637X/821/1/59.
- A. Fialkov, R. Barkana, and E. Visbal. The observable signature of late heating of the Universe during cosmic reionization. *Nature*, 506:197–199, February 2014. doi: 10.1038/nature12999.
- G. B. Field. . *Proceedings of the Institute of Radio Engineers*, 46:240, 1958.
- G. B. Field. The Spin Temperature of Intergalactic Neutral Hydrogen. *ApJ*, 129:536–+, May 1959. doi: 10.1086/146653.
- B. D. Fields, P. Molaro, and S. Sarkar. Big-Bang Nucleosynthesis. *ArXiv e-prints*, December 2014.
- K. Finlator, R. Davé, and F. Özel. Galactic Outflows and Photoionization Heating in the Reionization Epoch. *ApJ*, 743:169, December 2011. doi: 10.1088/0004-637X/743/2/169.
- R. A. Fisher. The logic of inductive inference. *Journal of the Royal Statistical Society*, 98(1):pp. 39–82, 1935. ISSN 09528385. URL <http://www.jstor.org/stable/2342435>.
- D. J. Fixsen, A. Kogut, S. Levin, M. Limon, P. Lubin, P. Mirel, M. Seiffert, J. Singal, E. Wollack, T. Villela, and C. A. Wuensche. ARCADE 2 Measurement of the Absolute Sky Brightness at 3-90 GHz. *ApJ*, 734:5, June 2011. doi: 10.1088/0004-637X/734/1/5.
- R. A. Flores and J. R. Primack. Observational and theoretical constraints on singular dark matter halos. *ApJ*, 427:L1–L4, May 1994. doi: 10.1086/187350.
- E. B. Fomalont, K. A. Ebner, W. J. M. van Breugel, and R. D. Ekers. Depolarization silhouettes and the filamentary structure in the radio source Fornax A. *ApJ*, 346:L17–L20, November 1989. doi: 10.1086/185568.
- T. Fragos et al. X-Ray Binary Evolution Across Cosmic Time. *ApJ*, 764:41, February 2013. doi: 10.1088/0004-637X/764/1/41.

- A. Frebel and J. E. Norris. Near-Field Cosmology with Extremely Metal-Poor Stars. *ARA&A*, 53:631–688, August 2015. doi: 10.1146/annurev-astro-082214-122423.
- A. Friedmann. Über die Krümmung des Raumes. *Zeitschrift für Physik*, 10:377–386, 1922. doi: 10.1007/BF01332580.
- M. M. Friedrich, G. Mellema, M. A. Alvarez, P. R. Shapiro, and I. T. Iliev. Topology and sizes of H II regions during cosmic reionization. *MNRAS*, 413:1353–1372, May 2011. doi: 10.1111/j.1365-2966.2011.18219.x.
- C. L. Fryer, S. E. Woosley, and A. Heger. Pair-Instability Supernovae, Gravity Waves, and Gamma-Ray Transients. *ApJ*, 550:372–382, March 2001. doi: 10.1086/319719.
- M. Fukugita and P. J. E. Peebles. The Cosmic Energy Inventory. *ApJ*, 616:643–668, December 2004. doi: 10.1086/425155.
- M. Fukugita, C. J. Hogan, and P. J. E. Peebles. The Cosmic Baryon Budget. *ApJ*, 503:518–530, August 1998. doi: 10.1086/306025.
- S. R. Furlanetto. The 21-cm forest. *MNRAS*, 370:1867–1875, August 2006a. doi: 10.1111/j.1365-2966.2006.10603.x.
- S. R. Furlanetto. The global 21-centimeter background from high redshifts. *MNRAS*, 371:867–878, September 2006b. doi: 10.1111/j.1365-2966.2006.10725.x.
- S. R. Furlanetto. The 21-cm Line as a Probe of Reionization. In A. Mesinger, editor, *Astrophysics and Space Science Library*, volume 423 of *Astrophysics and Space Science Library*, page 247, 2016. doi: 10.1007/978-3-319-21957-8_9.
- S. R. Furlanetto and A. Loeb. The 21 Centimeter Forest: Radio Absorption Spectra as Probes of Minihalos before Reionization. *ApJ*, 579:1–9, November 2002. doi: 10.1086/342757.
- S. R. Furlanetto and A. Loeb. Large-Scale Structure Shocks at Low and High Redshifts. *ApJ*, 611:642–654, August 2004.
- S. R. Furlanetto and S. J. Stoeber. Secondary ionization and heating by fast electrons. *MNRAS*, 404:1869–1878, June 2010. doi: 10.1111/j.1365-2966.2010.16401.x.
- S. R. Furlanetto, M. Zaldarriaga, and L. Hernquist. The Growth of H II Regions During Reionization. *ApJ*, 613:1–15, September 2004. doi: 10.1086/423025.
- S. R. Furlanetto, S. P. Oh, and F. H. Briggs. Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe. *Phys. Rep.*, 433:181–301, October 2006a. doi: 10.1016/j.physrep.2006.08.002.
- S. R. Furlanetto, S. P. Oh, and E. Pierpaoli. Effects of dark matter decay and annihilation on the high-redshift 21cm background. *Phys. Rev. D*, 74(10):103502, November 2006b. doi: 10.1103/PhysRevD.74.103502.

- J. P. Gardner, J. C. Mather, M. Clampin, R. Doyon, M. A. Greenhouse, H. B. Hammel, J. B. Hutchings, P. Jakobsen, S. J. Lilly, K. S. Long, J. I. Lunine, M. J. McCaughrean, M. Mountain, J. Nella, G. H. Rieke, M. J. Rieke, H.-W. Rix, E. P. Smith, G. Sonneborn, M. Stiavelli, H. S. Stockman, R. A. Windhorst, and G. S. Wright. The James Webb Space Telescope. *Space Sci. Rev.*, 123:485–606, April 2006. doi: 10.1007/s11214-006-8315-7.
- S. Garrison-Kimmel, M. Rocha, M. Boylan-Kolchin, J. S. Bullock, and J. Lally. Can feedback solve the too-big-to-fail problem? *MNRAS*, 433:3539–3546, August 2013. doi: 10.1093/mnras/stt984.
- P. M. Geil, B. M. Gaensler, and J. S. B. Wyithe. Polarized foreground removal at low radio frequencies using rotation measure synthesis: uncovering the signature of hydrogen reionization. *MNRAS*, 418:516–535, November 2011. doi: 10.1111/j.1365-2966.2011.19509.x.
- N. Y. Gnedin and P. A. Shaver. Redshifted 21 Centimeter Emission from the Pre-Reionization Era. I. Mean Signal and Linear Fluctuations. *ApJ*, 608:611–621, June 2004. doi: 10.1086/420735.
- L. J. Greenhill and G. Bernardi. HI Epoch of Reionization Arrays. *ArXiv e-prints*, January 2012.
- T. H. Greif, J. L. Johnson, V. Bromm, and R. S. Klessen. The First Supernova Explosions: Energetics, Feedback, and Chemical Enrichment. *ApJ*, 670:1–14, November 2007. doi: 10.1086/522028.
- T. H. Greif, V. Bromm, P. C. Clark, S. C. O. Glover, R. J. Smith, R. S. Klessen, N. Yoshida, and V. Springel. Formation and evolution of primordial protostellar systems. *MNRAS*, 424:399–415, July 2012. doi: 10.1111/j.1365-2966.2012.21212.x.
- B. Greig and A. Mesinger. 21CMMC: an MCMC analysis tool enabling astrophysical parameter studies of the cosmic 21 cm signal. *MNRAS*, 449:4246–4263, June 2015. doi: 10.1093/mnras/stv571.
- B. Greig and A. Mesinger. The global history of reionization. *MNRAS*, 465:4838–4852, March 2017. doi: 10.1093/mnras/stw3026.
- B. Greig, A. Mesinger, and L. V. E. Koopmans. Optimal core baseline design and observing strategy for probing the astrophysics of reionization with the SKA. *ArXiv e-prints*, September 2015.
- B. Greig, A. Mesinger, and J. C. Pober. Constraints on the temperature of the intergalactic medium at $z = 8.4$ with 21-cm observations. *MNRAS*, 455:4295–4300, February 2016. doi: 10.1093/mnras/stv2618.
- J. Greiner et al. GRB 080913 at Redshift 6.7. *ApJ*, 693:1610–1620, March 2009. doi: 10.1088/0004-637X/693/2/1610.

- David J Griffiths. *Introduction to electrodynamics; 4th ed.* Pearson, Boston, MA, 2013. URL <https://cds.cern.ch/record/1492149>.
- A. H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, January 1981. doi: 10.1103/PhysRevD.23.347.
- Z. Haiman and G. L. Bryan. Was Star Formation Suppressed in High-Redshift Mini-halos? *ApJ*, 650:7–11, October 2006. doi: 10.1086/506580.
- Z. Haiman, A. A. Thoul, and A. Loeb. Cosmological Formation of Low-Mass Objects. *ApJ*, 464:523–+, June 1996a. doi: 10.1086/177343.
- Z. Haiman, A. A. Thoul, and A. Loeb. Cosmological Formation of Low-Mass Objects. *ApJ*, 464:523, June 1996b. doi: 10.1086/177343.
- Z. Haiman, A. A. Thoul, and A. Loeb. Cosmological Formation of Low-Mass Objects. *ApJ*, 464:523, June 1996c. doi: 10.1086/177343.
- Z. Haiman, M. J. Rees, and A. Loeb. Destruction of Molecular Hydrogen during Cosmological Reionization. *ApJ*, 476:458–463, February 1997.
- Z. Haiman, T. Abel, and M. J. Rees. The Radiative Feedback of the First Cosmological Objects. *ApJ*, 534:11–24, May 2000.
- Z. Haiman, E. Quataert, and G. C. Bower. Modeling the Counts of Faint Radio-Loud Quasars: Constraints on the Supermassive Black Hole Population and Predictions for High Redshift. *ApJ*, 612:698–705, September 2004. doi: 10.1086/422834.
- S. E. G. Hales, J. E. Baldwin, and P. J. Warner. The 6C survey of radio sources. II - The zone $\delta = 30\text{--}51$ deg, $\alpha = 08\text{h}30\text{m}\text{--}17\text{h}30\text{m}$. *MNRAS*, 234:919–936, October 1988. doi: 10.1093/mnras/234.4.919.
- P. J. Hancock, T. Murphy, B. M. Gaensler, A. Hopkins, and J. R. Curran. Compact continuum source finding for next generation radio surveys. *MNRAS*, 422:1812–1824, May 2012. doi: 10.1111/j.1365-2966.2012.20768.x.
- D. Hanson, S. Hoover, A. Crites, P. A. R. Ade, K. A. Aird, J. E. Austermann, J. A. Beall, A. N. Bender, B. A. Benson, L. E. Bleem, J. J. Bock, J. E. Carlstrom, C. L. Chang, H. C. Chiang, H.-M. Cho, A. Conley, T. M. Crawford, T. de Haan, M. A. Dobbs, W. Everett, J. Gallicchio, J. Gao, E. M. George, N. W. Halverson, N. Harrington, J. W. Henning, G. C. Hilton, G. P. Holder, W. L. Holzapfel, J. D. Hrubes, N. Huang, J. Hubmayr, K. D. Irwin, R. Keisler, L. Knox, A. T. Lee, E. Leitch, D. Li, C. Liang, D. Luong-Van, G. Marsden, J. J. McMahon, J. Mehl, S. S. Meyer, L. Mocuano, T. E. Montroy, T. Natoli, J. P. Nibarger, V. Novosad, S. Padin, C. Pryke, C. L. Reichardt, J. E. Ruhl, B. R. Saliwanchik, J. T. Sayre, K. K. Schaffer, B. Schulz, G. Smecher, A. A. Stark, K. T. Story, C. Tucker, K. Vanderlinde, J. D. Vieira, M. P. Viero, G. Wang, V. Yefremenko, O. Zahn, and M. Zemcov. Detection of B-Mode Polarization in the Cosmic Microwave Background with

- Data from the South Pole Telescope. *Physical Review Letters*, 111(14):141301, October 2013. doi: 10.1103/PhysRevLett.111.141301.
- T. Hartwig, M. Volonteri, V. Bromm, R. S. Klessen, E. Barausse, M. Magg, and A. Stacy. Gravitational waves from the remnants of the first stars. *MNRAS*, 460:L74–L78, July 2016. doi: 10.1093/mnras/slw074.
- C. G. T. Haslam, C. J. Salter, H. Stoffel, and W. E. Wilson. A 408 MHz all-sky continuum survey. II - The atlas of contour maps. *A&AS*, 47:1, January 1982.
- B. J. Hazelton, M. F. Morales, and I. S. Sullivan. The Fundamental Multi-baseline Mode-mixing Foreground in 21 cm Epoch of Reionization Observations. *ApJ*, 770:156, June 2013. doi: 10.1088/0004-637X/770/2/156.
- J. F. Helmboldt and H. T. Intema. Very large array observations of disturbed ion flow from the plasmasphere to the nighttime ionosphere. *Radio Science*, 47:RS0K03, January 2012. doi: 10.1029/2012RS004979.
- J. F. Helmboldt, W. M. Lane, and W. D. Cotton. Climatology of midlatitude ionospheric disturbances from the Very Large Array Low-frequency Sky Survey. *Radio Science*, 47:RS5008, October 2012. doi: 10.1029/2012RS005025.
- G. Hinshaw, D. N. Spergel, L. Verde, R. S. Hill, S. S. Meyer, C. Barnes, C. L. Bennett, M. Halpern, N. Jarosik, A. Kogut, E. Komatsu, M. Limon, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: The Angular Power Spectrum. *ApJS*, 148:135–159, September 2003. doi: 10.1086/377225.
- G. Hinshaw et al. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *ApJS*, 208:19, October 2013. doi: 10.1088/0067-0049/208/2/19.
- S. Hirano and V. Bromm. Formation and survival of Population III stellar systems. *ArXiv e-prints*, December 2016.
- S. Hirano, T. Hosokawa, N. Yoshida, K. Omukai, and H. W. Yorke. Primordial star formation under the influence of far ultraviolet radiation: 1540 cosmological haloes and the stellar mass distribution. *MNRAS*, 448:568–587, March 2015. doi: 10.1093/mnras/stv044.
- C. M. Hirata. Wouthuysen-Field coupling strength and application to high-redshift 21-cm radiation. *MNRAS*, 367:259–274, March 2006. doi: 10.1111/j.1365-2966.2005.09949.x.
- E. Hubble. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15:168–173, March 1929. doi: 10.1073/pnas.15.3.168.

- J. A. Hummel, A. H. Pawlik, M. Milosavljević, and V. Bromm. The Source Density and Observability of Pair-instability Supernovae from the First Stars. *ApJ*, 755:72, August 2012. doi: 10.1088/0004-637X/755/1/72.
- N. Hurley-Walker, J. R. Callingham, P. J. Hancock, T. M. O. Franzen, L. Hindson, A. D. Kapińska, J. Morgan, A. R. Offringa, R. B. Wayth, C. Wu, Q. Zheng, T. Murphy, M. E. Bell, K. S. Dwarakanath, B. For, B. M. Gaensler, M. Johnston-Hollitt, E. Lenc, P. Procopio, L. Staveley-Smith, R. Ekers, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, L. Greenhill, B. J. Hazelton, D. L. Kaplan, C. J. Lonsdale, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, R. L. Webster, A. Williams, and C. L. Williams. GaLactic and Extragalactic All-sky Murchison Widefield Array (GLEAM) survey - I. A low-frequency extragalactic catalogue. *MNRAS*, 464:1146–1167, January 2017. doi: 10.1093/mnras/stw2337.
- H. T. Intema, P. Jagannathan, K. P. Mooley, and D. A. Frail. The GMRT 150 MHz all-sky radio survey. First alternative data release TGSS ADR1. *A&A*, 598:A78, February 2017. doi: 10.1051/0004-6361/201628536.
- Ž. Ivezić, K. Menou, G. R. Knapp, M. A. Strauss, R. H. Lupton, D. E. Vanden Berk, G. T. Richards, C. Tremonti, M. A. Weinstein, S. Anderson, N. A. Bahcall, R. H. Becker, M. Bernardi, M. Blanton, D. Eisenstein, X. Fan, D. Finkbeiner, K. Finlator, J. Frieman, J. E. Gunn, P. B. Hall, R. S. J. Kim, A. Kinkhabwala, V. K. Narayanan, C. M. Rockosi, D. Schlegel, D. P. Schneider, I. Strateva, M. SubbaRao, A. R. Thakar, W. Voges, R. L. White, B. Yanny, J. Brinkmann, M. Doi, M. Fukugita, G. S. Hennessy, J. A. Munn, R. C. Nichol, and D. G. York. Optical and Radio Properties of Extragalactic Sources Observed by the FIRST Survey and the Sloan Digital Sky Survey. *AJ*, 124:2364–2400, November 2002. doi: 10.1086/344069.
- D. C. Jacobs, J. E. Aguirre, A. R. Parsons, J. C. Pober, R. F. Bradley, C. L. Carilli, N. E. Gugliucci, J. R. Manley, C. van der Merwe, D. F. Moore, and C. R. Parashare. New 145 MHz Source Measurements by PAPER in the Southern Sky. *ApJ*, 734:L34, June 2011. doi: 10.1088/2041-8205/734/2/L34.
- D. C. Jacobs, A. R. Parsons, J. E. Aguirre, Z. Ali, J. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, P. Klima, D. H. E. MacMahon, J. R. Manley, D. F. Moore, J. C. Pober, I. I. Stefan, and W. P. Walbrugh. A Flux Scale for Southern Hemisphere 21 cm Epoch of Reionization Experiments. *ApJ*, 776:108, October 2013. doi: 10.1088/0004-637X/776/2/108.
- D. C. Jacobs, J. C. Pober, A. R. Parsons, J. E. Aguirre, Z. S. Ali, J. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, P. Klima, A. Liu, D. H. E. MacMahon, J. R. Manley, D. F. Moore, I. I. Stefan, and W. P. Walbrugh. Multiredshift Limits on the 21 cm Power Spectrum from PAPER. *ApJ*, 801:51, March 2015. doi: 10.1088/0004-637X/801/1/51.

- D. C. Jacobs, B. J. Hazelton, C. M. Trott, J. S. Dillon, B. Pindor, I. S. Sullivan, J. C. Pober, N. Barry, A. P. Beardsley, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, D. Emrich, A. Ewall-Wice, L. Feng, B. M. Gaensler, R. Goetze, L. J. Greenhill, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H. Kim, E. Kratzenberg, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, A. R. Neben, N. Thyagarajan, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, T. Prabu, P. Procopio, J. Ridding, A. E. E. Rogers, A. Roshi, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, M. Tegmark, S. J. Tingay, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. The Murchison Widefield Array 21 cm Power Spectrum Analysis Methodology. *ApJ*, 825:114, July 2016. doi: 10.3847/0004-637X/825/2/114.
- D. C. Jacobs, J. Burba, J. D. Bowman, A. R. Neben, B. Stinnett, L. Turner, K. Johnson, M. Busch, J. Allison, M. Leatham, V. Serrano Rodriguez, M. Denney, and D. Nelson. First Demonstration of ECHO: an External Calibrator for Hydrogen Observatories. *PASP*, 129(3):035002, March 2017. doi: 10.1088/1538-3873/aa56b9.
- J. H. Jeans. The Stability of a Spherical Nebula. *Philosophical Transactions of the Royal Society of London Series A*, 199:1–53, 1902. doi: 10.1098/rsta.1902.0012.
- V. Jelić, S. Zaroubi, P. Labropoulos, R. M. Thomas, G. Bernardi, M. A. Brentjens, A. G. de Bruyn, B. Ciardi, G. Harker, L. V. E. Koopmans, V. N. Pandey, J. Schaye, and S. Yatawatta. Foreground simulations for the LOFAR-epoch of reionization experiment. *MNRAS*, 389:1319–1335, September 2008. doi: 10.1111/j.1365-2966.2008.13634.x.
- V. Jelić, S. Zaroubi, P. Labropoulos, G. Bernardi, A. G. de Bruyn, and L. V. E. Koopmans. Realistic simulations of the Galactic polarized foreground: consequences for 21-cm reionization detection experiments. *MNRAS*, 409:1647–1659, December 2010. doi: 10.1111/j.1365-2966.2010.17407.x.
- V. Jelić, A. G. de Bruyn, M. Mevius, F. B. Abdalla, K. M. B. Asad, G. Bernardi, M. A. Brentjens, S. Bus, E. Chapman, B. Ciardi, S. Daiboo, E. R. Fernandez, A. Ghosh, G. Harker, H. Jensen, S. Kazemi, L. V. E. Koopmans, P. Labropoulos, O. Martinez-Rubi, G. Mellema, A. R. Offringa, V. N. Pandey, A. H. Patil, R. M. Thomas, H. K. Vedantham, V. Veligatla, S. Yatawatta, S. Zaroubi, A. Alexov, J. Anderson, I. M. Avruch, R. Beck, M. E. Bell, M. J. Bentum, P. Best, A. Bonafede, J. Bregman, F. Breitling, J. Broderick, W. N. Brouw, M. Brüggen, H. R. Butcher, J. E. Conway, F. de Gasperin, E. de Geus, A. Deller, R.-J. Dettmar, S. Duscha, J. Eislöffel, D. Engels, H. Falcke, R. A. Fallows, R. Fender, C. Ferrari, W. Frieswijk, M. A. Garrett, J. Grießmeier, A. W. Gunst, J. P. Hamaker, T. E. Hassall, M. Haverkorn, G. Heald, J. W. T. Hessels, M. Hoeft, J. Hörandel, A. Horneffer, A. van der Horst, M. Iacobelli, E. Juette, A. Karastergiou, V. I. Kondratiev, M. Kramer, M. Kuniyoshi, G. Kuper, J. van Leeuwen, P. Maat, G. Mann, D. McKay-Bukowski,

- J. P. McKean, H. Munk, A. Nelles, M. J. Norden, H. Paas, M. Pandey-Pommier, G. Pietka, R. Pizzo, A. G. Polatidis, W. Reich, H. Röttgering, A. Rowlinson, A. M. M. Scaife, D. Schwarz, M. Serylak, O. Smirnov, M. Steinmetz, A. Stewart, M. Tagger, Y. Tang, C. Tasse, S. ter Veen, S. Thoudam, C. Toribio, R. Vermeulen, C. Vocks, R. J. van Weeren, R. A. M. J. Wijers, S. J. Wijnholds, O. Wucknitz, and P. Zarka. Initial LOFAR observations of epoch of reionization windows. II. Diffuse polarized emission in the ELAIS-N1 field. *A&A*, 568:A101, August 2014. doi: 10.1051/0004-6361/201423998.
- A. Jenkins, C. S. Frenk, S. D. M. White, J. M. Colberg, S. Cole, A. E. Evrard, H. M. P. Couchman, and N. Yoshida. The mass function of dark matter haloes. *MNRAS*, 321:372–384, February 2001. doi: 10.1046/j.1365-8711.2001.04029.x.
- S. Jester and H. Falcke. Science with a lunar low-frequency array: From the dark ages of the Universe to nearby exoplanets. *New A Rev.*, 53:1–26, May 2009. doi: 10.1016/j.newar.2009.02.001.
- M. Kamionkowski, A. Kosowsky, and A. Stebbins. Statistics of cosmic microwave background polarization. *Phys. Rev. D*, 55:7368–7388, June 1997. doi: 10.1103/PhysRevD.55.7368.
- X. Kang, A. V. Macciò, and A. A. Dutton. The Effect of Warm Dark Matter on Galaxy Properties: Constraints from the Stellar Mass Function and the Tully-Fisher Relation. *ApJ*, 767:22, April 2013. doi: 10.1088/0004-637X/767/1/22.
- T. Karlsson, V. Bromm, and J. Bland-Hawthorn. Pregalactic metal enrichment: The chemical signatures of the first stars. *Reviews of Modern Physics*, 85:809–848, April 2013. doi: 10.1103/RevModPhys.85.809.
- N. E. Kassim, R. A. Perley, W. C. Erickson, and K. S. Dwarakanath. Subarcminute resolution imaging of radio sources at 74 MHz with the Very Large Array. *AJ*, 106:2218–2228, December 1993. doi: 10.1086/116795.
- N. E. Kassim, T. J. W. Lazio, W. C. Erickson, R. A. Perley, W. D. Cotton, E. W. Greisen, A. S. Cohen, B. Hicks, H. R. Schmitt, and D. Katz. The 74 MHz System on the Very Large Array. *ApJS*, 172:686–719, October 2007. doi: 10.1086/519022.
- S. Kazemi and S. Yatawatta. Robust radio interferometric calibration using the t -distribution. *MNRAS*, 435:597–605, October 2013. doi: 10.1093/mnras/stt1347.
- S. Kazemi, S. Yatawatta, S. Zaroubi, P. Lampropoulos, A. G. de Bruyn, L. V. E. Koopmans, and J. Noordam. Radio interferometric calibration using the SAGE algorithm. *MNRAS*, 414:1656–1666, June 2011. doi: 10.1111/j.1365-2966.2011.18506.x.
- S. Kazemi, S. Yatawatta, and S. Zaroubi. Clustered calibration: an improvement to radio interferometric direction-dependent self-calibration. *MNRAS*, 430:1457–1472, April 2013. doi: 10.1093/mnras/stt018.

- T. Kitayama, N. Yoshida, H. Susa, and M. Umemura. The Structure and Evolution of Early Cosmological H II Regions. *ApJ*, 613:631–645, October 2004. doi: 10.1086/423313.
- S. A. Kohn, J. Aguirre, C. Nunhokee, G. Bernardi, J. Pober, Z. Ali, R. Bradley, C. Carilli, D. DeBoer, N. Gugliucci, D. Jacobs, P. Klima, D. MacMahon, J. Manley, D. Moore, A. Parsons, I. Stefan, and W. Walbrugh. Constraining Polarized Foregrounds for EoR Experiments I: 2D Power Spectra from the PAPER-32 Imaging Array. *ArXiv e-prints*, February 2016.
- E. W. Kolb and M. S. Turner. *The early universe*. 1990.
- E. Komatsu, K. M. Smith, J. Dunkley, C. L. Bennett, B. Gold, G. Hinshaw, N. Jarosik, D. Larson, M. R. Nolta, L. Page, D. N. Spergel, M. Halpern, R. S. Hill, A. Kogut, M. Limon, S. S. Meyer, N. Odegard, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation. *ApJS*, 192:18, February 2011. doi: 10.1088/0067-0049/192/2/18.
- L. Koopmans, J. Pritchard, G. Mellema, J. Aguirre, K. Ahn, R. Barkana, I. van Bemmel, G. Bernardi, A. Bonaldi, F. Briggs, A. G. de Bruyn, T. C. Chang, E. Chapman, X. Chen, B. Ciardi, P. Dayal, A. Ferrara, A. Fialkov, F. Fiore, K. Ichiki, I. T. Illiev, S. Inoue, V. Jelic, M. Jones, J. Lazio, U. Maio, S. Majumdar, K. J. Mack, A. Mesinger, M. F. Morales, A. Parsons, U. L. Pen, M. Santos, R. Schneider, B. Semelin, R. S. de Souza, R. Subrahmanyan, T. Takeuchi, H. Vedantham, J. Wagg, R. Webster, S. Wyithe, K. K. Datta, and C. Trott. The Cosmic Dawn and Epoch of Reionisation with SKA. *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, art. 1, April 2015.
- J. D. Kraus. *Radio astronomy*. 1973.
- M. Kuhlen and C.-A. Faucher-Giguere. Concordance models of reionization: implications for faint galaxies and escape fraction evolution. *MNRAS*, 423:862–876, June 2012. doi: 10.1111/j.1365-2966.2012.20924.x.
- M. Kuhlen and P. Madau. The first miniquasar. *MNRAS*, 363:1069–1082, November 2005. doi: 10.1111/j.1365-2966.2005.09522.x.
- W. M. Lane, W. D. Cotton, S. van Velzen, T. E. Clarke, N. E. Kassim, J. F. Helmboldt, T. J. W. Lazio, and A. S. Cohen. The Very Large Array Low-frequency Sky Survey Redux (VLSSr). *MNRAS*, 440:327–338, May 2014. doi: 10.1093/mnras/stu256.
- M. I. Large, B. Y. Mills, A. G. Little, D. F. Crawford, and J. M. Sutton. The Molonglo Reference Catalogue of Radio Sources. *MNRAS*, 194:693, February 1981. doi: 10.1093/mnras/194.3.693.

- R. Laureijs, J. Amiaux, S. Arduini, J. . Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, and et al. Euclid Definition Study Report. *ArXiv e-prints*, October 2011.
- N. Leite, C. Evoli, M. D'Angelo, B. Ciardi, G. Sigl, and A. Ferrara. Do Cosmic Rays Heat the Early Intergalactic Medium? *ArXiv e-prints*, March 2017.
- E. Lenc, B. M. Gaensler, X. H. Sun, E. M. Sadler, A. G. Willis, N. Barry, A. P. Beardsley, M. E. Bell, G. Bernardi, J. D. Bowman, F. Briggs, J. R. Callingham, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, A. A. Deshpande, J. S. Dillon, K. S. Dwarkanath, D. Emrich, A. Ewall-Wice, L. Feng, B.-Q. For, R. Goeke, L. J. Greenhill, P. Hancock, B. J. Hazelton, J. N. Hewitt, L. Hindson, N. Hurley-Walker, M. Johnston-Hollitt, D. C. Jacobs, A. D. Kapinska, D. L. Kaplan, J. C. Kasper, H.-S. Kim, E. Kratzenberg, J. Line, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, J. Morgan, T. Murphy, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, L. Staveley-Smith, R. Subrahmanyam, I. S. Sullivan, M. Tegmark, N. Thyagarajan, S. J. Tingay, C. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, J. S. B. Wyithe, and Q. Zheng. Low frequency observations of linearly polarized structures in the interstellar medium near the south Galactic pole. *ArXiv e-prints*, July 2016.
- A. R. Liddle. An Introduction to Cosmological Inflation. In A. Masiero, G. Senjanovic, and A. Smirnov, editors, *High Energy Physics and Cosmology, 1998 Summer School*, page 260, 1999.
- A. R. Liddle and D. H. Lyth. *Cosmological Inflation and Large-Scale Structure*. June 2000.
- A. Lidz, O. Zahn, M. McQuinn, M. Zaldarriaga, and L. Hernquist. Detecting the Rise and Fall of 21 cm Fluctuations with the Murchison Widefield Array. *ApJ*, 680:962–974, June 2008. doi: 10.1086/587618.
- A. D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108:389–393, February 1982. doi: 10.1016/0370-2693(82)91219-9.
- J. Line et al. In Preparation. *In Preparation*, 2016.
- A. Liu and A. R. Parsons. Constraining cosmology and ionization history with combined 21 cm power spectrum and global signal measurements. *MNRAS*, 457:1864–1877, April 2016. doi: 10.1093/mnras/stw071.
- A. Liu and M. Tegmark. A method for 21 cm power spectrum estimation in the presence of foregrounds. *Phys. Rev. D*, 83(10):103006, May 2011. doi: 10.1103/PhysRevD.83.103006.

- A. Liu and M. Tegmark. How well can we measure and understand foregrounds with 21-cm experiments? *MNRAS*, 419:3491–3504, February 2012. doi: 10.1111/j.1365-2966.2011.19989.x.
- A. Liu, M. Tegmark, J. Bowman, J. Hewitt, and M. Zaldarriaga. An improved method for 21-cm foreground removal. *MNRAS*, 398:401–406, September 2009. doi: 10.1111/j.1365-2966.2009.15156.x.
- A. Liu, M. Tegmark, S. Morrison, A. Lutomirski, and M. Zaldarriaga. Precision calibration of radio interferometers using redundant baselines. *MNRAS*, 408:1029–1050, October 2010. doi: 10.1111/j.1365-2966.2010.17174.x.
- A. Liu, A. R. Parsons, and C. M. Trott. Epoch of reionization window. I. Mathematical formalism. *Phys. Rev. D*, 90(2):023018, July 2014a. doi: 10.1103/PhysRevD.90.023018.
- A. Liu, A. R. Parsons, and C. M. Trott. Epoch of reionization window. II. Statistical methods for foreground wedge reduction. *Phys. Rev. D*, 90(2):023019, July 2014b. doi: 10.1103/PhysRevD.90.023019.
- Adrian Liu, Jonathan R. Pritchard, Rupert Allison, Aaron R. Parsons, Uroš Seljak, and Blake D. Sherwin. Eliminating the optical depth nuisance from the cmb with 21 cm cosmology. *Phys. Rev. D*, 93:043013, Feb 2016. doi: 10.1103/PhysRevD.93.043013. URL <http://link.aps.org/doi/10.1103/PhysRevD.93.043013>.
- C. Liu, S. J. Mutch, P. W. Angel, A. R. Duffy, P. M. Geil, G. B. Poole, A. Mesinger, and J. S. B. Wyithe. Dark-ages reionization and galaxy formation simulation - IV. UV luminosity functions of high-redshift galaxies. *MNRAS*, 462:235–249, October 2016. doi: 10.1093/mnras/stw1015.
- A. Loeb and F. A. Rasio. Collapse of primordial gas clouds and the formation of quasar black holes. *ApJ*, 432:52–61, September 1994. doi: 10.1086/174548.
- A. Loeb and M. Zaldarriaga. Measuring the Small-Scale Power Spectrum of Cosmic Density Fluctuations through 21cm Tomography Prior to the Epoch of Structure Formation. *Physical Review Letters*, 92(21):211301, May 2004. doi: 10.1103/PhysRevLett.92.211301.
- S. T. Loi, T. Murphy, M. E. Bell, D. L. Kaplan, E. Lenc, A. R. Offringa, N. Hurley-Walker, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, B. E. Corey, A. A. Deshpande, D. Emrich, B. M. Gaensler, R. Goetze, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, J. C. Kasper, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, A. E. E. Rogers, A. Rosi, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, and C. L. Williams. Quantifying ionospheric effects on time-domain astrophysics with the Murchison Widefield Array. *MNRAS*, 453:2731–2746, November 2015a. doi: 10.1093/mnras/stv1808.

- S. T. Loi, T. Murphy, I. H. Cairns, F. W. Menk, C. L. Waters, P. J. Erickson, C. M. Trott, N. Hurley-Walker, J. Morgan, E. Lenc, A. R. Offringa, M. E. Bell, R. D. Ekers, B. M. Gaensler, C. J. Lonsdale, L. Feng, P. J. Hancock, D. L. Kaplan, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, R. B. Wayth, R. L. Webster, A. Williams, and C. L. Williams. Real-time imaging of density ducts between the plasmasphere and ionosphere. *Geophys. Res. Lett.*, 42:3707–3714, May 2015b. doi: 10.1002/2015GL063699.
- C. J. Lonsdale, R. J. Cappallo, M. F. Morales, F. H. Briggs, L. Benkevitch, J. D. Bowman, J. D. Bunton, S. Burns, B. E. Corey, L. Desouza, S. S. Doeleman, M. Derome, A. Deshpande, M. R. Gopala, L. J. Greenhill, D. E. Herne, J. N. Hewitt, P. A. Kamini, J. C. Kasper, B. B. Kincaid, J. Kocz, E. Kowald, E. Kratzenberg, D. Kumar, M. J. Lynch, S. Madhavi, M. Matejek, D. A. Mitchell, E. Morgan, D. Oberoi, S. Ord, J. Pathikulangara, T. Prabu, A. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. U. Shankar, K. S. Srivani, J. Stevens, S. Tingay, A. Vaccarella, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, and C. Williams. The Murchison Widefield Array: Design Overview. *IEEE Proceedings*, 97:1497–1506, August 2009. doi: 10.1109/JPROC.2009.2017564.
- L. Lopez-Honorez, O. Mena, Á. Moliné, S. Palomares-Ruiz, and A. C. Vincent. The 21 cm signal and the interplay between dark matter annihilations and astrophysical processes. *J. Cosmology Astropart. Phys.*, 8:004, August 2016. doi: 10.1088/1475-7516/2016/08/004.
- L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz, and P. Villanueva Domingo. Constraints on warm dark matter from the ionization history of the Universe. *ArXiv e-prints*, March 2017.
- M. R. Lovell, V. Eke, C. S. Frenk, L. Gao, A. Jenkins, T. Theuns, J. Wang, S. D. M. White, A. Boyarsky, and O. Ruchayskiy. The haloes of bright satellite galaxies in a warm dark matter universe. *MNRAS*, 420:2318–2324, March 2012. doi: 10.1111/j.1365-2966.2011.20200.x.
- LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al. LSST Science Book, Version 2.0. *ArXiv e-prints*, December 2009.
- A. V. Macciò, S. Paduroiu, D. Anderhalden, A. Schneider, and B. Moore. Cores in warm dark matter haloes: a Catch 22 problem. *MNRAS*, 424:1105–1112, August 2012. doi: 10.1111/j.1365-2966.2012.21284.x.
- K. J. Mack and J. S. B. Wyithe. Detecting the redshifted 21 cm forest during reionization. *MNRAS*, 425:2988–3001, October 2012. doi: 10.1111/j.1365-2966.2012.21561.x.

- P. Madau and F. Haardt. Cosmic Reionization after Planck: Could Quasars Do It All? *ApJ*, 813:L8, November 2015. doi: 10.1088/2041-8205/813/1/L8.
- P. Madau and M. J. Rees. Massive Black Holes as Population III Remnants. *ApJ*, 551:L27–L30, April 2001. doi: 10.1086/319848.
- P. Madau, A. Meiksin, and M. J. Rees. 21 Centimeter Tomography of the Intergalactic Medium at High Redshift. *ApJ*, 475:429–+, February 1997. doi: 10.1086/303549.
- P. Madau, F. Haardt, and M. J. Rees. Radiative Transfer in a Clumpy Universe. III. The Nature of Cosmological Ionizing Sources. *ApJ*, 514:648–659, April 1999. doi: 10.1086/306975.
- Y. Mao, M. Tegmark, M. McQuinn, M. Zaldarriaga, and O. Zahn. How accurately can 21cm tomography constrain cosmology? *Phys. Rev. D*, 78(2):023529, July 2008. doi: 10.1103/PhysRevD.78.023529.
- C. A. Mason, M. Trenti, and T. Treu. The Galaxy UV Luminosity Function before the Epoch of Reionization. *ApJ*, 813:21, November 2015. doi: 10.1088/0004-637X/813/1/21.
- K. W. Masui, E. R. Switzer, N. Banavar, K. Bandura, C. Blake, L.-M. Calin, T.-C. Chang, X. Chen, Y.-C. Li, Y.-W. Liao, A. Natarajan, U.-L. Pen, J. B. Peterson, J. R. Shaw, and T. C. Voytek. Measurement of 21 cm Brightness Fluctuations at $z \sim 0.8$ in Cross-correlation. *ApJ*, 763:L20, January 2013. doi: 10.1088/2041-8205/763/1/L20.
- T. Mauch, T. Murphy, H. J. Buttery, J. Curran, R. W. Hunstead, B. Piestrzynski, J. G. Robertson, and E. M. Sadler. SUMSS: a wide-field radio imaging survey of the southern sky - II. The source catalogue. *MNRAS*, 342:1117–1130, July 2003. doi: 10.1046/j.1365-8711.2003.06605.x.
- M. R. C. McDowell. On the formation of H₂ in H I regions. *The Observatory*, 81: 240–243, December 1961.
- I. D. McGreer, A. Mesinger, and V. D’Odorico. Model-independent evidence in favour of an end to reionization by z approximately 6. *MNRAS*, 447:499–505, February 2015. doi: 10.1093/mnras/stu2449.
- J. P. McMullin, B. Waters, D. Schiebel, W. Young, and K. Golap. CASA Architecture and Applications. In R. A. Shaw, F. Hill, and D. J. Bell, editors, *Astronomical Data Analysis Software and Systems XVI*, volume 376 of *Astronomical Society of the Pacific Conference Series*, page 127, October 2007.
- M. McQuinn. Constraints on X-ray emissions from the reionization era. *MNRAS*, 426:1349–1360, October 2012. doi: 10.1111/j.1365-2966.2012.21792.x.
- M. McQuinn. The Evolution of the Intergalactic Medium. *ArXiv e-prints*, November 2015.

- M. McQuinn and R. M. O’Leary. The Impact of the Supersonic Baryon-Dark Matter Velocity Difference on the $z \sim 20$ 21 cm Background. *ApJ*, 760:3, November 2012. doi: 10.1088/0004-637X/760/1/3.
- M. McQuinn, S. R. Furlanetto, L. Hernquist, O. Zahn, and M. Zaldarriaga. The Kinetic Sunyaev-Zel’dovich Effect from Reionization. *ApJ*, 630:643–656, September 2005. doi: 10.1086/432049.
- M. McQuinn, O. Zahn, M. Zaldarriaga, L. Hernquist, and S. R. Furlanetto. Cosmological Parameter Estimation Using 21 cm Radiation from the Epoch of Reionization. *ApJ*, 653:815–834, December 2006. doi: 10.1086/505167.
- M. McQuinn, A. Lidz, O. Zahn, S. Dutta, L. Hernquist, and M. Zaldarriaga. The morphology of HII regions during reionization. *MNRAS*, 377:1043–1063, May 2007. doi: 10.1111/j.1365-2966.2007.11489.x.
- A. Meiksin. Constraints on the ionization sources of the high-redshift intergalactic medium. *MNRAS*, 356:596–606, January 2005. doi: 10.1111/j.1365-2966.2004.08481.x.
- G. Mellema, L. V. E. Koopmans, F. A. Abdalla, G. Bernardi, B. Ciardi, S. Daiboo, A. G. de Bruyn, K. K. Datta, H. Falcke, A. Ferrara, I. T. Iliev, F. Iocco, V. Jelić, H. Jensen, R. Joseph, P. Labropoulos, A. Meiksin, A. Mesinger, A. R. Offringa, V. N. Pandey, J. R. Pritchard, M. G. Santos, D. J. Schwarz, B. Semelin, H. Vedantham, S. Yatawatta, and S. Zaroubi. Reionization and the Cosmic Dawn with the Square Kilometre Array. *Experimental Astronomy*, 36:235–318, August 2013. doi: 10.1007/s10686-013-9334-5.
- A. Mesinger and M. Dijkstra. Ultraviolet radiative feedback during the advanced stages of reionization. *MNRAS*, 390:1071–1080, November 2008. doi: 10.1111/j.1365-2966.2008.13776.x.
- A. Mesinger and S. Furlanetto. Efficient Simulations of Early Structure Formation and Reionization. *ApJ*, 669:663–675, November 2007. doi: 10.1086/521806.
- A. Mesinger, Z. Haiman, and R. Cen. Probing the Reionization History Using the Spectra of High-Redshift Sources. *ApJ*, 613:23–35, September 2004. doi: 10.1086/422898.
- A. Mesinger, G. L. Bryan, and Z. Haiman. Ultraviolet Radiative Feedback on High-Redshift Protogalaxies. *ApJ*, 648:835–851, September 2006. doi: 10.1086/506173.
- A. Mesinger, S. Furlanetto, and R. Cen. 21CMFAST: a fast, seminumerical simulation of the high-redshift 21-cm signal. *MNRAS*, 411:955–972, February 2011. doi: 10.1111/j.1365-2966.2010.17731.x.
- A. Mesinger, M. McQuinn, and D. N. Spergel. The kinetic Sunyaev-Zel’dovich signal from inhomogeneous reionization: a parameter space study. *MNRAS*, 422:1403–1417, May 2012. doi: 10.1111/j.1365-2966.2012.20713.x.

- A. Mesinger, A. Ferrara, and D. S. Spiegel. Signatures of X-rays in the early Universe. *MNRAS*, 431:621–637, May 2013. doi: 10.1093/mnras/stt198.
- A. Mesinger, A. Ewall-Wice, and J. Hewitt. Reionization and beyond: detecting the peaks of the cosmological 21 cm signal. *MNRAS*, 439:3262–3274, April 2014. doi: 10.1093/mnras/stu125.
- A. Mesinger, B. Greig, and E. Sobacchi. The Evolution Of 21 cm Structure (EOS): public, large-scale simulations of Cosmic Dawn and reionization. *MNRAS*, 459: 2342–2353, July 2016. doi: 10.1093/mnras/stw831.
- S. Mineo, M. Gilfanov, and R. Sunyaev. X-ray emission from star-forming galaxies - I. High-mass X-ray binaries. *MNRAS*, 419:2095–2115, January 2012a. doi: 10.1111/j.1365-2966.2011.19862.x.
- S. Mineo, M. Gilfanov, and R. Sunyaev. X-ray emission from star-forming galaxies - II. Hot interstellar medium. *MNRAS*, 426:1870–1883, November 2012b. doi: 10.1111/j.1365-2966.2012.21831.x.
- I. F. Mirabel, M. Dijkstra, P. Laurent, A. Loeb, and J. R. Pritchard. Stellar black holes at the dawn of the universe. *A&A*, 528:A149, April 2011. doi: 10.1051/0004-6361/201016357.
- J. Mirocha, S. R. Furlanetto, and G. Sun. The global 21-cm signal in the context of the high- z galaxy luminosity function. *MNRAS*, 464:1365–1379, January 2017. doi: 10.1093/mnras/stw2412.
- D. A. Mitchell, L. J. Greenhill, R. B. Wayth, R. J. Sault, C. J. Lonsdale, R. J. Cappallo, M. F. Morales, and S. M. Ord. Real-Time Calibration of the Murchison Widefield Array. *IEEE Journal of Selected Topics in Signal Processing*, 2:707–717, November 2008. doi: 10.1109/JSTSP.2008.2005327.
- R. A. Monsalve, A. E. E. Rogers, J. D. Bowman, and T. J. Mozdzen. Calibration of the EDGES High-band Receiver to Observe the Global 21 cm Signature from the Epoch of Reionization. *ApJ*, 835:49, January 2017. doi: 10.3847/1538-4357/835/1/49.
- B. Moore, S. Ghigna, F. Governato, G. Lake, T. Quinn, J. Stadel, and P. Tozzi. Dark Matter Substructure within Galactic Halos. *ApJ*, 524:L19–L22, October 1999a. doi: 10.1086/312287.
- B. Moore, T. Quinn, F. Governato, J. Stadel, and G. Lake. Cold collapse and the core catastrophe. *MNRAS*, 310:1147–1152, December 1999b. doi: 10.1046/j.1365-8711.1999.03039.x.
- D. F. Moore, J. E. Aguirre, A. R. Parsons, D. C. Jacobs, and J. C. Pober. The Effects of Polarized Foregrounds on 21 cm Epoch of Reionization Power Spectrum Measurements. *ApJ*, 769:154, June 2013. doi: 10.1088/0004-637X/769/2/154.

- D. F. Moore, J. E. Aguirre, S. A. Kohn, A. R. Parsons, Z. S. Ali, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, D. C. Jacobs, P. Klima, A. Liu, D. H. E. MacMahon, J. R. Manley, J. C. Pober, I. I. Stefan, and W. P. Walbrugh. Limits on Polarized Leakage for the PAPER Epoch of Reionization Measurements at 126 and 164 MHz. *ApJ*, 836:154, February 2017. doi: 10.3847/1538-4357/836/2/154.
- M. F. Morales. Power Spectrum Sensitivity and the Design of Epoch of Reionization Observatories. *ApJ*, 619:678–683, February 2005. doi: 10.1086/426730.
- M. F. Morales and J. Hewitt. Toward Epoch of Reionization Measurements with Wide-Field Radio Observations. *ApJ*, 615:7–18, November 2004. doi: 10.1086/424437.
- M. F. Morales and J. S. B. Wyithe. Reionization and Cosmology with 21-cm Fluctuations. *ARA&A*, 48:127–171, September 2010. doi: 10.1146/annurev-astro-081309-130936.
- M. F. Morales, B. Hazelton, I. Sullivan, and A. Beardsley. Four Fundamental Foreground Power Spectrum Shapes for 21 cm Cosmology Observations. *ApJ*, 752:137, June 2012. doi: 10.1088/0004-637X/752/2/137.
- D. J. Mortlock, S. J. Warren, B. P. Venemans, M. Patel, P. C. Hewett, R. G. McMahon, C. Simpson, T. Theuns, E. A. González-Solares, A. Adamson, S. Dye, N. C. Hambly, P. Hirst, M. J. Irwin, E. Kuiper, A. Lawrence, and H. J. A. Röttgering. A luminous quasar at a redshift of $z = 7.085$. *Nature*, 474:616–619, June 2011. doi: 10.1038/nature10159.
- D. Munshi, P. Valageas, L. van Waerbeke, and A. Heavens. Cosmology with weak lensing surveys. *Phys. Rep.*, 462:67–121, June 2008. doi: 10.1016/j.physrep.2008.02.003.
- S. Naoz, N. Yoshida, and R. Barkana. The non-linear evolution of baryonic overdensities in the early universe: initial conditions of numerical simulations. *MNRAS*, 416:232–241, September 2011. doi: 10.1111/j.1365-2966.2011.19025.x.
- J. F. Navarro, C. S. Frenk, and S. D. M. White. A Universal Density Profile from Hierarchical Clustering. *ApJ*, 490:493–508, December 1997. doi: 10.1086/304888.
- A. R. Neben, R. F. Bradley, J. N. Hewitt, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, A. A. Deshpande, R. Goeke, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, D. L. Kaplan, C. J. Lonsdale, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, R. B. Wayth, R. L. Webster, A. Williams, and C. L. Williams. Measuring phased-array antenna beam patterns with high dynamic range for the Murchison Widefield Array using 137 MHz ORBCOMM satellites. *Radio Science*, 50:614–629, July 2015. doi: 10.1002/2015RS005678.

- A. R. Neben, R. F. Bradley, J. N. Hewitt, D. R. DeBoer, A. R. Parsons, J. E. Aguirre, Z. S. Ali, C. Cheng, A. Ewall-Wice, N. Patra, N. Thyagarajan, J. Bowman, R. Dickenson, J. S. Dillon, P. Doolittle, D. Egan, M. Hedrick, D. C. Jacobs, S. A. Kohn, P. J. Klima, K. Moodley, B. R. B. Saliwanchik, P. Schaffner, J. Shelton, H. A. Taylor, R. Taylor, M. Tegmark, B. Wirt, and H. Zheng. The Hydrogen Epoch of Reionization Array Dish. I. Beam Pattern Measurements and Science Implications. *ApJ*, 826:199, August 2016. doi: 10.3847/0004-637X/826/2/199.
- L. B. Newburgh, G. E. Addison, M. Amiri, K. Bandura, J. R. Bond, L. Connor, J.-F. Cliche, G. Davis, M. Deng, N. Denman, M. Dobbs, M. Fandino, H. Fong, K. Gibbs, A. Gilbert, E. Griffin, M. Halpern, D. Hanna, A. D. Hincks, G. Hinshaw, C. Höfer, P. Klages, T. Landecker, K. Masui, J. M. Parra, U.-L. Pen, J. Peterson, A. Recnik, J. R. Shaw, K. Sigurdson, M. Sitwell, G. Smecher, R. Smegal, K. Vanderlinde, and D. Wiebe. Calibrating CHIME: a new radio interferometer to probe dark energy. In *Ground-based and Airborne Telescopes V*, volume 9145 of Proc. SPIE, page 91454V, July 2014. doi: 10.1117/12.2056962.
- K. Nomoto, N. Tominaga, H. Umeda, C. Kobayashi, and K. Maeda. Nucleosynthesis yields of core-collapse supernovae and hypernovae, and galactic chemical evolution. *Nuclear Physics A*, 777:424–458, October 2006. doi: 10.1016/j.nuclphysa.2006.05.008.
- A. R. Offringa, J. J. van de Gronde, and J. B. T. M. Roerdink. A morphological algorithm for improving radio-frequency interference detection. *A&A*, 539:A95, March 2012. doi: 10.1051/0004-6361/201118497.
- A. R. Offringa, B. McKinley, N. Hurley-Walker, F. H. Briggs, R. B. Wayth, D. L. Kaplan, M. E. Bell, L. Feng, A. R. Neben, J. D. Hughes, J. Rhee, T. Murphy, N. D. R. Bhat, G. Bernardi, J. D. Bowman, R. J. Cappallo, B. E. Corey, A. A. Deshpande, D. Emrich, A. Ewall-Wice, B. M. Gaensler, R. Goeke, L. J. Greenhill, B. J. Hazelton, L. Hindson, M. Johnston-Hollitt, D. C. Jacobs, J. C. Kasper, E. Kratzenberg, E. Lenc, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, N. Kudryavtseva, D. Oberoi, S. M. Ord, B. Pindor, P. Procopio, T. Prabu, J. Riding, D. A. Roshi, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, and C. L. Williams. WSCLEAN: an implementation of a fast, generic wide-field imager for radio astronomy. *MNRAS*, 444:606–619, October 2014. doi: 10.1093/mnras/stu1368.
- A. R. Offringa, R. B. Wayth, N. Hurley-Walker, D. L. Kaplan, N. Barry, A. P. Beardsley, M. E. Bell, G. Bernardi, J. D. Bowman, F. Briggs, J. R. Callingham, R. J. Cappallo, P. Carroll, A. A. Deshpande, J. S. Dillon, K. S. Dwarakanath, A. Ewall-Wice, L. Feng, B.-Q. For, B. M. Gaensler, L. J. Greenhill, P. Hancock, B. J. Hazelton, J. N. Hewitt, L. Hindson, D. C. Jacobs, M. Johnston-Hollitt, A. D. Kapińska, H.-S. Kim, P. Kittiwisit, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, J. Morgan, A. R. Neben, D. Oberoi, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu,

- P. Procopio, J. Riding, N. Udaya Shankar, S. Sethi, K. S. Srivani, L. Staveley-Smith, R. Subrahmanyan, I. S. Sullivan, M. Tegmark, N. Thyagarajan, S. J. Tingay, C. M. Trott, R. L. Webster, A. Williams, C. L. Williams, C. Wu, J. S. Wyithe, and Q. Zheng. The Low-Frequency Environment of the Murchison Widefield Array: Radio-Frequency Interference Analysis and Mitigation. *PASA*, 32:e008, March 2015. doi: 10.1017/pasa.2015.7.
- A. R. Offringa, C. M. Trott, N. Hurley-Walker, M. Johnston-Hollitt, B. McKinley, N. Barry, A. P. Beardsley, J. D. Bowman, F. Briggs, P. Carroll, J. S. Dillon, A. Ewall-Wice, L. Feng, B. M. Gaensler, L. J. Greenhill, B. J. Hazelton, J. N. Hewitt, D. C. Jacobs, H.-S. Kim, P. Kittiwisit, E. Lenc, J. Line, A. Loeb, D. A. Mitchell, M. F. Morales, A. R. Neben, S. Paul, B. Pindor, J. C. Pober, P. Procopio, J. Riding, S. K. Sethi, N. U. Shankar, R. Subrahmanyan, I. S. Sullivan, M. Tegmark, N. Thyagarajan, S. J. Tingay, R. B. Wayth, R. L. Webster, and J. S. B. Wyithe. Parametrising Epoch of Reionization foregrounds: A deep survey of low-frequency point-source spectra with the MWA. *MNRAS*, February 2016. doi: 10.1093/mnras/stw310.
- S. P. Oh and K. J. Mack. Foregrounds for 21-cm observations of neutral gas at high redshift. *MNRAS*, 346:871–877, December 2003. doi: 10.1111/j.1365-2966.2003.07133.x.
- T. Okamoto, L. Gao, and T. Theuns. Mass loss of galaxies due to an ultraviolet background. *MNRAS*, 390:920–928, November 2008. doi: 10.1111/j.1365-2966.2008.13830.x.
- S. M. Ord, D. A. Mitchell, R. B. Wayth, L. J. Greenhill, G. Bernardi, S. Gleadow, R. G. Edgar, M. A. Clark, G. Allen, W. Arcus, L. Benkevitch, J. D. Bowman, F. H. Briggs, J. D. Bunton, S. Burns, R. J. Cappallo, W. A. Coles, B. E. Corey, L. Desouza, S. S. Doeleman, M. Derome, A. Deshpande, D. Emrich, R. Goeke, M. R. Gopalakrishna, D. Herne, J. N. Hewitt, P. A. Kamini, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, J. Kocz, E. Kowald, E. Kratzenberg, D. Kumar, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, S. Madhavi, M. Matejek, M. F. Morales, E. Morgan, D. Oberoi, J. Pathikulangara, T. Prabu, A. E. E. Rogers, A. Rosh, J. E. Salah, A. Schinkel, N. Udaya Shankar, K. S. Srivani, J. Stevens, S. J. Tingay, A. Vaccarella, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, and C. Williams. Interferometric Imaging with the 32 Element Murchison Wide-Field Array. *PASP*, 122:1353–1366, November 2010. doi: 10.1086/657160.
- S. M. Ord, B. Crosse, D. Emrich, D. Pallot, R. B. Wayth, M. A. Clark, S. E. Tremblay, W. Arcus, D. Barnes, M. Bell, G. Bernardi, N. D. R. Bhat, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. A. Deshpande, L. deSouza, A. Ewell-Wice, L. Feng, R. Goeke, L. J. Greenhill, B. J. Hazelton, D. Herne, J. N. Hewitt, L. Hindson, N. Hurley-Walker, D. Jacobs, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, N. Kudryavtseva, E. Lenc, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A.

- Mitchell, M. F. Morales, E. Morgan, D. Oberoi, A. Offringa, J. Pathikulangara, B. Pindor, T. Prabu, P. Procopio, R. A. Remillard, J. Riding, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya Shankar, K. S. Srivani, J. Stevens, R. Subrahmanyan, S. J. Tingay, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The Murchison Widefield Array Correlator. *PASA*, 32:e006, March 2015. doi: 10.1017/pasa.2015.5.
- G. Paciga, J. G. Albert, K. Bandura, T.-C. Chang, Y. Gupta, C. Hirata, J. Odegova, U.-L. Pen, J. B. Peterson, J. Roy, J. R. Shaw, K. Sigurdson, and T. Voytek. A simulation-calibrated limit on the H I power spectrum from the GMRT Epoch of Reionization experiment. *MNRAS*, 433:639–647, July 2013. doi: 10.1093/mnras/stt753.
- F. Pacucci, A. Mesinger, and Z. Haiman. Focusing on warm dark matter with lensed high-redshift galaxies. *MNRAS*, 435:L53–L57, August 2013. doi: 10.1093/mnras/slt093.
- F. Pacucci, A. Mesinger, S. Mineo, and A. Ferrara. The X-ray spectra of the first galaxies: 21 cm signatures. *MNRAS*, 443:678–686, September 2014. doi: 10.1093/mnras/stu1240.
- A. Parsons, J. Pober, M. McQuinn, D. Jacobs, and J. Aguirre. A Sensitivity and Array-configuration Study for Measuring the Power Spectrum of 21 cm Emission from Reionization. *ApJ*, 753:81, July 2012a. doi: 10.1088/0004-637X/753/1/81.
- A. R. Parsons and D. C. Backer. Calibration of Low-Frequency, Wide-Field Radio Interferometers Using Delay/Delay-Rate Filtering. *AJ*, 138:219–226, July 2009. doi: 10.1088/0004-6256/138/1/219.
- A. R. Parsons, J. C. Pober, J. E. Aguirre, C. L. Carilli, D. C. Jacobs, and D. F. Moore. A Per-baseline, Delay-spectrum Technique for Accessing the 21 cm Cosmic Reionization Signature. *ApJ*, 756:165, September 2012b. doi: 10.1088/0004-637X/756/2/165.
- A. R. Parsons, A. Liu, Z. S. Ali, and C. Cheng. Optimized Beam Sculpting with Generalized Fringe-rate Filters. *ApJ*, 820:51, March 2016a. doi: 10.3847/0004-637X/820/1/51.
- A. R. Parsons, A. Liu, Z. S. Ali, and C. Cheng. Optimized Beam Sculpting with Generalized Fringe-rate Filters. *ApJ*, 820:51, March 2016b. doi: 10.3847/0004-637X/820/1/51.
- A. R. Parsons et al. New Limits on 21 cm Epoch of Reionization from PAPER-32 Consistent with an X-Ray Heated Intergalactic Medium at $z = 7.7$. *ApJ*, 788:106, June 2014. doi: 10.1088/0004-637X/788/2/106.
- Aaron R. Parsons et al. The Precision Array for Probing the Epoch of Re-ionization: Eight Station Results. *AJ*, 139:1468–1480, April 2010. doi: 10.1088/0004-6256/139/4/1468.

- A. H. Patil, S. Yatawatta, S. Zaroubi, L. V. E. Koopmans, A. G. de Bruyn, V. Jelić, B. Ciardi, I. T. Iliev, M. Mevius, V. N. Pandey, and B. K. Gehlot. Systematic biases in low frequency radio interferometric data due to calibration: the LOFAR EoR case. *ArXiv e-prints*, May 2016.
- A. H. Patil, S. Yatawatta, L. V. E. Koopmans, A. G. de Bruyn, M. A. Brentjens, S. Zaroubi, K. M. B. Asad, M. Hatef, V. Jelić, M. Mevius, A. R. Offringa, V. N. Pandey, H. Vedantham, F. B. Abdalla, W. N. Brouw, E. Chapman, B. Ciardi, B. K. Gehlot, A. Ghosh, G. Harker, I. T. Iliev, K. Kakiichi, S. Majumdar, G. Mellema, M. B. Silva, J. Schaye, D. Vrbanec, and S. J. Wijnholds. Upper Limits on the 21 cm Epoch of Reionization Power Spectrum from One Night with LOFAR. *ApJ*, 838:65, March 2017. doi: 10.3847/1538-4357/aa63e7.
- N. Patra, R. Subrahmanyam, A. Raghunathan, and N. Udaya Shankar. SARAS: a precision system for measurement of the cosmic radio background and signatures from the epoch of reionization. *Experimental Astronomy*, 36:319–370, August 2013. doi: 10.1007/s10686-013-9336-3.
- N. Patra, J. D. Bray, P. Roberts, and R. D. Ekers. Bandpass calibration of a wideband spectrometer using coherent pulse injection. *ArXiv e-prints*, February 2015.
- N. Patra, A. R. Parsons, D. R. DeBoer, N. Thyagarajan, A. Ewall-Wice, G. Hsyu, T. Kuk Leung, C. K. Day, J. E. Aguirre, P. Alexander, Z. S. Ali, A. P. Beardsley, J. D. Bowman, R. F. Bradley, C. L. Carilli, C. Cheng, E. de Lera Acedo, J. S. Dillon, G. Fadana, N. Fagnoni, R. Fritz, S. R. Furlanetto, B. Glendenning, B. Greig, J. Grobbelaar, B. J. Hazelton, J. N. Hewitt, D. C. Jacobs, A. Julius, M. Kariseb, S. A. Kohn, A. Lebedeva, T. Lekalake, A. Liu, A. Loots, D. MacMahon, L. Malan, C. Malgas, M. Maree, Z. Martinot, N. Mathison, E. Matsetela, A. Mesinger, M. F. Morales, A. R. Neben, S. Pieterse, J. C. Pober, N. Razavi-Ghods, J. Ringuette, J. Robnett, K. Rosie, R. Sell, C. Smith, A. Syce, M. Tegmark, P. K. G. Williams, and H. Zheng. The Hydrogen Epoch of Reionization Array Dish III: Measuring Chromaticity of Prototype Element with Reflectometry. *ArXiv e-prints*, January 2017.
- A. H. Pawlik and J. Schaye. Photoheating and supernova feedback amplify each other’s effect on the cosmic star formation rate. *MNRAS*, 396:L46–L50, June 2009. doi: 10.1111/j.1745-3933.2009.00659.x.
- J. A. Peacock, S. Cole, P. Norberg, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. D. Cannon, M. Colless, C. Collins, W. Couch, G. Dalton, K. Deeley, R. De Propris, S. P. Driver, G. Efstathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, W. J. Percival, B. A. Peterson, I. Price, W. Sutherland, and K. Taylor. A measurement of the cosmological mass density from clustering in the 2dF Galaxy Redshift Survey. *Nature*, 410:169–173, March 2001.
- P. J. E. Peebles. *The large-scale structure of the universe*. 1980.

- L. Pentericci, A. Fontana, E. Vanzella, M. Castellano, A. Grazian, M. Dijkstra, K. Boutsia, S. Cristiani, M. Dickinson, E. Giallongo, M. Giavalisco, R. Maiolino, A. Moorwood, D. Paris, and P. Santini. Spectroscopic Confirmation of $z \sim 7$ Lyman Break Galaxies: Probing the Earliest Galaxies and the Epoch of Reionization. *ApJ*, 743:132, December 2011. doi: 10.1088/0004-637X/743/2/132.
- A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *ApJ*, 142:419–421, July 1965. doi: 10.1086/148307.
- W. J. Percival, C. M. Baugh, J. Bland-Hawthorn, T. Bridges, R. Cannon, S. Cole, M. Colless, C. Collins, W. Couch, G. Dalton, R. De Propriis, S. P. Driver, G. Efsthathiou, R. S. Ellis, C. S. Frenk, K. Glazebrook, C. Jackson, O. Lahav, I. Lewis, S. Lumsden, S. Maddox, S. Moody, P. Norberg, J. A. Peacock, B. A. Peterson, W. Sutherland, and K. Taylor. The 2dF Galaxy Redshift Survey: the power spectrum and the matter content of the Universe. *MNRAS*, 327:1297–1306, November 2001. doi: 10.1046/j.1365-8711.2001.04827.x.
- R. A. Perley, H.-J. Roser, and K. Meisenheimer. The radio galaxy PictorA – a study with the VLA. *A&A*, 328:12–32, December 1997.
- S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *ApJ*, 517:565–586, June 1999. doi: 10.1086/307221.
- M. Persic and P. Salucci. The baryon content of the universe. *MNRAS*, 258:14P–18P, September 1992. doi: 10.1093/mnras/258.1.14P.
- A. H. G. Peter, M. Rocha, J. S. Bullock, and M. Kaplinghat. Cosmological simulations with self-interacting dark matter - II. Halo shapes versus observations. *MNRAS*, 430:105–120, March 2013. doi: 10.1093/mnras/sts535.
- A. Pillepich, C. Porciani, and S. Matarrese. The Bispectrum of Redshifted 21 Centimeter Fluctuations from the Dark Ages. *ApJ*, 662:1–14, June 2007. doi: 10.1086/517963.
- + Planck Collaboration. Planck 2013 results. XVI. Cosmological parameters. *ArXiv e-prints:1303.5076*, March 2013.
- Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. Planck 2015 results. XIII. Cosmological parameters. *ArXiv e-prints*, February 2015.
- Planck Collaboration, R. Adam, N. Aghanim, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak,

R. Battye, K. Benabed, J.-P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, E. Calabrese, J.-F. Cardoso, J. Carron, H. C. Chiang, L. P. L. Colombo, C. Combet, B. Comis, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, E. Di Valentino, C. Dickinson, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, F. Elsner, T. A. Enßlin, H. K. Eriksen, E. Falgarone, Y. Fantaye, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, A. Gruppuso, J. E. Gudmundsson, F. K. Hansen, G. Helou, S. Henrot-Versillé, D. Herranz, E. Hivon, Z. Huang, S. Ilić, A. H. Jaffe, W. C. Jones, E. Keihänen, R. Keskitalo, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lähteenmäki, J.-M. Lamarre, M. Langer, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. López-Caniego, Y.-Z. Ma, J. F. Macías-Pérez, G. Maggio, A. Mangilli, M. Maris, P. G. Martin, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M.-A. Miville-Deschênes, D. Molinari, A. Moneti, L. Montier, G. Morgante, A. Moss, P. Naselsky, P. Natoli, C. A. Oxborrow, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, L. Patrizii, O. Perdereau, L. Perotto, V. Pettorino, F. Piacentini, S. Plaszczynski, L. Polastri, G. Polenta, J.-L. Puget, J. P. Rachen, B. Racine, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, M. Rossetti, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, G. Sirri, R. Sunyaev, A.-S. Suur-Uski, J. A. Tauber, M. Tenti, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, J. Valiviita, F. Van Tent, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, A. Zacchei, and A. Zonca. Planck intermediate results. XLVII. Planck constraints on reionization history. *A&A*, 596:A108, December 2016a. doi: 10.1051/0004-6361/201628897.

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, and et al. Planck 2015 results. XX. Constraints on inflation. *A&A*, 594:A20, September 2016b. doi: 10.1051/0004-6361/201525898.

J. C. Pober, A. R. Parsons, D. C. Jacobs, J. E. Aguirre, R. F. Bradley, C. L. Carilli, N. E. Gugliucci, D. F. Moore, and C. R. Parashare. A Technique for Primary Beam Calibration of Drift-scanning, Wide-field Antenna Elements. *AJ*, 143:53, February 2012. doi: 10.1088/0004-6256/143/2/53.

J. C. Pober, A. R. Parsons, J. E. Aguirre, Z. Ali, R. F. Bradley, C. L. Carilli, D. DeBoer, M. Dexter, N. E. Gugliucci, D. C. Jacobs, P. J. Klima, D. MacMahon, J. Manley, D. F. Moore, I. I. Stefan, and W. P. Walbrugh. Opening the 21 cm Epoch of Reionization Window: Measurements of Foreground Isolation with PAPER. *ApJ*, 768:L36, May 2013a. doi: 10.1088/2041-8205/768/2/L36.

- J. C. Pober, A. R. Parsons, D. R. DeBoer, P. McDonald, M. McQuinn, J. E. Aguirre, Z. Ali, R. F. Bradley, T.-C. Chang, and M. F. Morales. The Baryon Acoustic Oscillation Broadband and Broad-beam Array: Design Overview and Sensitivity Forecasts. *AJ*, 145:65, March 2013b. doi: 10.1088/0004-6256/145/3/65.
- J. C. Pober, A. Liu, J. S. Dillon, J. E. Aguirre, J. D. Bowman, R. F. Bradley, C. L. Carilli, D. R. DeBoer, J. N. Hewitt, D. C. Jacobs, M. McQuinn, M. F. Morales, A. R. Parsons, M. Tegmark, and D. J. Werthimer. What Next-generation 21 cm Power Spectrum Measurements can Teach us About the Epoch of Reionization. *ApJ*, 782:66, February 2014. doi: 10.1088/0004-637X/782/2/66.
- J. C. Pober, Z. S. Ali, A. R. Parsons, M. McQuinn, J. E. Aguirre, G. Bernardi, R. F. Bradley, C. L. Carilli, C. Cheng, D. R. DeBoer, M. R. Dexter, S. R. Furlanetto, J. Grobbelaar, J. Horrell, D. C. Jacobs, P. J. Klima, S. A. Kohn, A. Liu, D. H. E. MacMahon, M. Maree, A. Mesinger, D. F. Moore, N. Razavi-Ghods, I. I. Stefan, W. P. Walbrugh, A. Walker, and H. Zheng. PAPER-64 Constraints On Reionization. II. The Temperature of the $z = 8.4$ Intergalactic Medium. *ApJ*, 809:62, August 2015. doi: 10.1088/0004-637X/809/1/62.
- J. C. Pober, B. J. Hazelton, A. P. Beardsley, N. A. Barry, Z. E. Martinot, I. S. Sullivan, M. F. Morales, M. E. Bell, G. Bernardi, N. D. R. Bhat, J. D. Bowman, F. Briggs, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, A. A. Deshpande, J. S. Dillon, D. Emrich, A. M. Ewall-Wice, L. Feng, R. Goeke, L. J. Greenhill, J. N. Hewitt, L. Hindson, N. Hurley-Walker, D. C. Jacobs, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H.-S. Kim, P. Kittiwisit, E. Kratzenberg, N. Kudryavtseva, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, E. Morgan, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, S. K. Sethi, N. Udaya Shankar, K. S. Srivani, R. Subrahmanyam, M. Tegmark, N. Thyagarajan, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The Importance of Wide-field Foreground Removal for 21 cm Cosmology: A Demonstration with Early MWA Epoch of Reionization Observations. *ApJ*, 819:8, March 2016. doi: 10.3847/0004-637X/819/1/8.
- Jonathan Pober. Updated 21 cm Experiment Sensitivities. http://reionization.org/wp-content/uploads/2015/05/HERA4_sensecalc.pdf, 2015.
- E. Polisensky and M. Ricotti. Constraints on the dark matter particle mass from the number of Milky Way satellites. *Phys. Rev. D*, 83(4):043506, February 2011. doi: 10.1103/PhysRevD.83.043506.
- T. Prabu, K. S. Srivani, D. A. Roshi, P. A. Kamini, S. Madhavi, D. Emrich, B. Crosse, A. J. Williams, M. Waterson, A. A. Deshpande, N. U. Shankar, R. Subrahmanyam, F. H. Briggs, R. F. Goeke, S. J. Tingay, M. Johnston-Hollitt, G. M. R, E. H. Morgan, J. Pathikulangara, J. D. Bunton, G. Hampson, C. Williams, S. M. Ord, R. B. Wayth, D. Kumar, M. F. Morales, L. deSouza, E. Kratzenberg, D. Pallot,

- R. McWhirter, B. J. Hazelton, W. Arcus, D. G. Barnes, G. Bernardi, T. Boller, J. D. Bowman, R. J. Cappallo, B. E. Corey, L. J. Greenhill, D. Herne, J. N. Hewitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, C. J. Lonsdale, M. J. Lynch, D. A. Mitchell, D. Oberoi, R. A. Remillard, A. E. Rogers, J. E. Salah, R. J. Sault, J. B. Stevens, S. Tremblay, R. L. Webster, A. R. Whitney, and S. B. Wyithe. A digital-receiver for the MurchisonWidefield Array. *Experimental Astronomy*, 39: 73–93, March 2015. doi: 10.1007/s10686-015-9444-3.
- I. Prandoni and N. Seymour. Revealing the Physics and Evolution of Galaxies and Galaxy Clusters with SKA Continuum Surveys. *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, art. 67, April 2015.
- M. E. Presley, A. Liu, and A. R. Parsons. Measuring the Cosmological 21 cm Monopole with an Interferometer. *ApJ*, 809:18, August 2015. doi: 10.1088/0004-637X/809/1/18.
- W. H. Press and P. Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. *ApJ*, 187:425–438, February 1974. doi: 10.1086/152650.
- J. R. Pritchard and S. R. Furlanetto. Descending from on high: Lyman-series cascades and spin-kinetic temperature coupling in the 21-cm line. *MNRAS*, 367:1057–1066, April 2006. doi: 10.1111/j.1365-2966.2006.10028.x.
- J. R. Pritchard and S. R. Furlanetto. 21-cm fluctuations from inhomogeneous X-ray heating before reionization. *MNRAS*, 376:1680–1694, April 2007. doi: 10.1111/j.1365-2966.2007.11519.x.
- J. R. Pritchard and A. Loeb. Evolution of the 21cm signal throughout cosmic history. *Phys. Rev. D*, 78(10):103511, November 2008. doi: 10.1103/PhysRevD.78.103511.
- J. R. Pritchard and A. Loeb. 21 cm cosmology in the 21st century. *Reports on Progress in Physics*, 75(8):086901, August 2012. doi: 10.1088/0034-4885/75/8/086901.
- C. L. Reichardt, L. Shaw, O. Zahn, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H. M. Cho, T. M. Crawford, A. T. Crites, T. de Haan, M. A. Dobbs, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, Z. Hou, J. D. Hrubes, M. Joy, R. Keisler, L. Knox, A. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, T. Natoli, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, K. K. Schaffer, E. Shirokoff, H. G. Spieler, Z. Staniszewski, A. A. Stark, K. Story, A. van Engelen, K. Vanderlinde, J. D. Vieira, and R. Williamson. A Measurement of Secondary Cosmic Microwave Background Anisotropies with Two Years of South Pole Telescope Observations. *ApJ*, 755:70, August 2012. doi: 10.1088/0004-637X/755/1/70.

- M. Remazeilles, C. Dickinson, A. J. Banday, M.-A. Bigot-Sazy, and T. Ghosh. An improved source-subtracted and destriped 408-MHz all-sky map. *MNRAS*, 451: 4311–4327, August 2015. doi: 10.1093/mnras/stv1274.
- M. Ricotti and J. P. Ostriker. X-ray pre-ionization powered by accretion on the first black holes - I. A model for the WMAP polarization measurement. *MNRAS*, 352: 547–562, August 2004. doi: 10.1111/j.1365-2966.2004.07942.x.
- M. Ricotti, N. Y. Gnedin, and J. M. Shull. Feedback from Galaxy Formation: Production and Photodissociation of Primordial H₂. *ApJ*, 560:580–591, October 2001. doi: 10.1086/323051.
- William Rideout and Anthea Coster. Automated gps processing for global total electron content data. *GPS Solutions*, 10(3):219–228, 2006. ISSN 1080-5370. doi: 10.1007/s10291-006-0029-5. URL <http://dx.doi.org/10.1007/s10291-006-0029-5>.
- A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *AJ*, 116:1009–1038, September 1998. doi: 10.1086/300499.
- M. Rocha, A. H. G. Peter, J. S. Bullock, M. Kaplinghat, S. Garrison-Kimmel, J. Oñorbe, and L. A. Moustakas. Cosmological simulations with self-interacting dark matter - I. Constant-density cores and substructure. *MNRAS*, 430:81–104, March 2013. doi: 10.1093/mnras/sts514.
- Alan EE Rogers and Judd D Bowman. Spectral index of the diffuse radio background measured from 100 to 200 MHz. *The Astronomical Journal*, 136(2):641, 2008.
- C. L. Rufenach. Power-law wavenumber spectrum deduced from ionospheric scintillation observations. *J. Geophys. Res.*, 77:4761, 1972. doi: 10.1029/JA077i025p04761.
- G. B. Rybicki. Improved Fokker-Planck Equation for Resonance-Line Scattering. *ApJ*, 647:709–718, August 2006. doi: 10.1086/505327.
- G. B. Rybicki and A. P. Lightman. *Radiative processes in astrophysics*. 1979.
- R. Salvaterra, A. Ferrara, and P. Dayal. Simulating high-redshift galaxies. *MNRAS*, 414:847–859, June 2011. doi: 10.1111/j.1365-2966.2010.18155.x.
- S. Salvini and S. J. Wijnholds. Fast gain calibration in radio astronomy using alternating direction implicit methods: Analysis and applications. *A&A*, 571:A97, November 2014. doi: 10.1051/0004-6361/201424487.

- M. G. Santos, A. Amblard, J. Pritchard, H. Trac, R. Cen, and A. Cooray. Cosmic Reionization and the 21 cm Signal: Comparison between an Analytical Model and a Simulation. *ApJ*, 689:1–16, December 2008. doi: 10.1086/592487.
- M. G. Santos, L. Ferramacho, M. B. Silva, A. Amblard, and A. Cooray. Fast large volume simulations of the 21-cm signal from the reionization and pre-reionization epochs. *MNRAS*, 406:2421–2432, August 2010. doi: 10.1111/j.1365-2966.2010.16898.x.
- R. J. Sault, J. P. Hamaker, and J. D. Bregman. Understanding radio polarimetry. II. Instrumental calibration of an interferometer array. *A&AS*, 117:149–159, May 1996.
- S. Sazonov and R. Sunyaev. Preheating of the Universe by cosmic rays from primordial supernovae at the beginning of cosmic reionization. *MNRAS*, 454:3464–3471, December 2015. doi: 10.1093/mnras/stv2255.
- U. Seljak and M. Zaldarriaga. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *ApJ*, 469:437, October 1996. doi: 10.1086/177793.
- J. R. Shaw, K. Sigurdson, U.-L. Pen, A. Stebbins, and M. Sitwell. All-sky Interferometry with Spherical Harmonic Transit Telescopes. *ApJ*, 781:57, February 2014. doi: 10.1088/0004-637X/781/2/57.
- J. R. Shaw, K. Sigurdson, M. Sitwell, A. Stebbins, and U.-L. Pen. Coaxing cosmic 21 cm fluctuations from the polarized sky using m -mode analysis. *Phys. Rev. D*, 91(8):083514, April 2015. doi: 10.1103/PhysRevD.91.083514.
- R. K. Sheth, H. J. Mo, and G. Tormen. Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes. *MNRAS*, 323:1–12, May 2001. doi: 10.1046/j.1365-8711.2001.04006.x.
- J. M. Shull. The Inventory and Distribution of Baryons at Redshift $z=0$. In J. L. Rosenberg and M. E. Putman, editors, *The IGM/Galaxy Connection. The Distribution of Baryons at $z=0$* , volume 281 of *Astrophysics and Space Science Library*, page 1, April 2003. doi: 10.1007/978-94-010-0115-1_1.
- J. M. Shull and M. E. van Steenberg. X-ray secondary heating and ionization in quasar emission-line clouds. *ApJ*, 298:268–274, November 1985. doi: 10.1086/163605.
- S. Singh, R. Subrahmanyam, N. Udaya Shankar, and A. Raghunathan. On the Detection of Global 21-cm Signal from Reionization Using Interferometers. *ApJ*, 815:88, December 2015. doi: 10.1088/0004-637X/815/2/88.
- D.G. Singleton. Power spectra of ionospheric scintillations. *Journal of Atmospheric and Terrestrial Physics*, 36(1):113 – 133, 1974. ISSN 0021-9169. doi: [http://dx.doi.org/10.1016/0021-9169\(74\)90071-3](http://dx.doi.org/10.1016/0021-9169(74)90071-3). URL <http://www.sciencedirect.com/science/article/pii/0021916974900713>.

- M. Sitwell, A. Mesinger, Y.-Z. Ma, and K. Sigurdson. The imprint of warm dark matter on the cosmological 21-cm signal. *MNRAS*, 438:2664–2671, March 2014. doi: 10.1093/mnras/stt2392.
- O. B. Slee. Radio sources observed with the Culgoora circular array. *Australian Journal of Physics*, 48:143–186, 1995. doi: 10.1071/PH950143.
- G. F. Smoot, C. L. Bennett, A. Kogut, E. L. Wright, J. Aymon, N. W. Boggess, E. S. Cheng, G. de Amici, S. Gulkis, M. G. Hauser, G. Hinshaw, P. D. Jackson, M. Janssen, E. Kaita, T. Kelsall, P. Keegstra, C. Lineweaver, K. Loewenstein, P. Lubin, J. Mather, S. S. Meyer, S. H. Moseley, T. Murdock, L. Rokke, R. F. Silverberg, L. Tenorio, R. Weiss, and D. T. Wilkinson. Structure in the COBE differential microwave radiometer first-year maps. *ApJ*, 396:L1–L5, September 1992. doi: 10.1086/186504.
- E. Sobacchi and A. Mesinger. The depletion of gas in high-redshift dwarf galaxies from an inhomogeneous reionization. *MNRAS*, 432:L51, May 2013a. doi: 10.1093/mnrasl/slt035.
- E. Sobacchi and A. Mesinger. How does radiative feedback from an ultraviolet background impact reionization? *MNRAS*, 432:3340–3348, July 2013b. doi: 10.1093/mnras/stt693.
- E. Sobacchi and A. Mesinger. Inhomogeneous recombinations during cosmic reionization. *MNRAS*, 440:1662–1673, May 2014. doi: 10.1093/mnras/stu377.
- M. Sokolowski, S. E. Tremblay, R. B. Wayth, S. J. Tingay, N. Clarke, P. Roberts, M. Waterson, R. D. Ekers, P. Hall, M. Lewis, M. Mossammaparast, S. Padhi, F. Schlagenhafer, A. Sutinjo, and J. Tickner. BIGHORNS - Broadband Instrument for Global HyDrOgen ReioNisation Signal. *PASA*, 32:e004, February 2015. doi: 10.1017/pasa.2015.3.
- R. S. Somerville. Can Photoionization Squelching Resolve the Substructure Crisis? *ApJ*, 572:L23–L26, June 2002. doi: 10.1086/341444.
- A. Songaila and L. L. Cowie. The Evolution of Lyman Limit Absorption Systems to Redshift Six. *ApJ*, 721:1448–1466, October 2010. doi: 10.1088/0004-637X/721/2/1448.
- D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, J. Centrella, W. Traub, C. Baltay, J. Colbert, D. Bennett, A. Kiessling, B. Macintosh, J. Merten, M. Mortonson, M. Penny, E. Rozo, D. Savransky, K. Stapelfeldt, Y. Zu, C. Baker, E. Cheng, D. Content, J. Dooley, M. Foote, R. Goullioud, K. Grady, C. Jackson, J. Kruk, M. Levine, M. Melton, C. Peddie, J. Ruffa, and S. Shaklan. WFIRST-2.4: What Every Astronomer Should Know. *ArXiv e-prints*, May 2013.

- D. N. Spergel and P. J. Steinhardt. Observational Evidence for Self-Interacting Cold Dark Matter. *Physical Review Letters*, 84:3760–3763, April 2000. doi: 10.1103/PhysRevLett.84.3760.
- D. N. Spergel, L. Verde, H. V. Peiris, E. Komatsu, M. R.olta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, J. L. Weiland, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. *ApJS*, 148:175–194, September 2003. doi: 10.1086/377226.
- V. Springel and L. Hernquist. The history of star formation in a Λ cold dark matter universe. *MNRAS*, 339:312–334, February 2003. doi: 10.1046/j.1365-8711.2003.06207.x.
- A. Stacy, T. H. Greif, and V. Bromm. The first stars: formation of binaries and small multiple systems. *MNRAS*, 403:45–60, March 2010. doi: 10.1111/j.1365-2966.2009.16113.x.
- D. P. Stark, R. S. Ellis, K. Chiu, M. Ouchi, and A. Bunker. Keck spectroscopy of faint $3 < z < 7$ Lyman break galaxies - I. New constraints on cosmic reionization from the luminosity and redshift-dependent fraction of Lyman alpha emission. *MNRAS*, 408:1628–1648, November 2010. doi: 10.1111/j.1365-2966.2010.17227.x.
- A. A. Starobinskiĭ. Spectrum of relict gravitational radiation and the early state of the universe. *Soviet Journal of Experimental and Theoretical Physics Letters*, 30:682, December 1979.
- I. S. Sullivan, M. F. Morales, B. J. Hazelton, W. Arcus, D. Barnes, G. Bernardi, F. H. Briggs, J. D. Bowman, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. Deshpande, L. deSouza, D. Emrich, B. M. Gaensler, R. Goeke, L. J. Greenhill, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, E. Morgan, D. Oberoi, S. M. Ord, J. Pathikulangara, T. Prabu, R. A. Remillard, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya Shankar, K. S. Srivani, J. Stevens, R. Subrahmanyam, S. J. Tingay, R. B. Wayth, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. Fast Holographic Deconvolution: A New Technique for Precision Radio Interferometry. *ApJ*, 759:17, November 2012. doi: 10.1088/0004-637X/759/1/17.
- M. Tegmark. How to measure CMB power spectra without losing information. *Phys. Rev. D*, 55:5895–5907, May 1997a. doi: 10.1103/PhysRevD.55.5895.
- M. Tegmark. How to Make Maps from Cosmic Microwave Background Data without Losing Information. *ApJ*, 480:L87–L90, May 1997b. doi: 10.1086/310631.
- M. Tegmark. Parallel Universes. *Scientific American*, 288(5):40–51, May 2003. doi: 10.1038/scientificamerican0503-40.

- M. Tegmark and M. Zaldarriaga. Fast Fourier transform telescope. *Phys. Rev. D*, 79 (8):083530, April 2009. doi: 10.1103/PhysRevD.79.083530.
- M. Tegmark, J. Silk, M. J. Rees, A. Blanchard, T. Abel, and F. Palla. How Small Were the First Cosmological Objects? *ApJ*, 474:1, January 1997a. doi: 10.1086/303434.
- M. Tegmark, A. N. Taylor, and A. F. Heavens. Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? *ApJ*, 480:22–35, May 1997b.
- The Dark Energy Survey Collaboration. The Dark Energy Survey. *ArXiv Astrophysics e-prints*, October 2005.
- A. R. Thompson, J. M. Moran, and G. W. Swenson. *Interferometry and synthesis in radio astronomy*. 1986.
- N. Thyagarajan, D. J. Helfand, R. L. White, and R. H. Becker. Variable and Transient Radio Sources in the FIRST Survey. *ApJ*, 742:49, November 2011. doi: 10.1088/0004-637X/742/1/49.
- N. Thyagarajan, N. Udaya Shankar, R. Subrahmanyan, W. Arcus, G. Bernardi, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, L. deSouza, D. Emrich, B. M. Gaensler, R. F. Goeke, L. J. Greenhill, B. J. Hazelton, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. H. Morgan, D. Oberoi, S. M. Ord, J. Pathikulangara, R. A. Remillard, A. E. E. Rogers, D. Anish Roshi, J. E. Salah, R. J. Sault, K. S. Srivani, J. B. Stevens, P. Thiagaraj, S. J. Tingay, R. B. Wayth, M. Waterson, R. L. Webster, A. R. Whitney, A. J. Williams, C. L. Williams, and J. S. B. Wyithe. A Study of Fundamental Limitations to Statistical Detection of Redshifted H I from the Epoch of Reionization. *ApJ*, 776:6, October 2013. doi: 10.1088/0004-637X/776/1/6.
- N. Thyagarajan, D. C. Jacobs, J. D. Bowman, N. Barry, A. P. Beardsley, G. Bernardi, F. Briggs, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, J. S. Dillon, D. Emrich, A. Ewall-Wice, L. Feng, R. Goeke, L. J. Greenhill, B. J. Hazelton, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H.-S. Kim, P. Kittiwisit, E. Kratzenberg, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyan, I. S. Sullivan, M. Tegmark, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. Foregrounds in Wide-field Redshifted 21 cm Power Spectra. *ApJ*, 804:14, May 2015a. doi: 10.1088/0004-637X/804/1/14.
- N. Thyagarajan, D. C. Jacobs, J. D. Bowman, N. Barry, A. P. Beardsley, G. Bernardi, F. Briggs, R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, J. S. Dillon,

- D. Emrich, A. Ewall-Wice, L. Feng, R. Goeke, L. J. Greenhill, B. J. Hazelton, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, H.-S. Kim, P. Kittiwisit, E. Kratzenberg, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, A. E. E. Rogers, A. Roshi, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, M. Tegmark, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. Foregrounds in Wide-field Redshifted 21 cm Power Spectra. *ApJ*, 804:14, May 2015b. doi: 10.1088/0004-637X/804/1/14.
- N. Thyagarajan, D. C. Jacobs, J. D. Bowman, N. Barry, A. P. Beardsley, G. Bernardi, F. Briggs, R. J. Cappallo, P. Carroll, A. A. Deshpande, A. de Oliveira-Costa, J. S. Dillon, A. Ewall-Wice, L. Feng, L. J. Greenhill, B. J. Hazelton, L. Hernquist, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. L. Kaplan, H.-S. Kim, P. Kittiwisit, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, A. R. Neben, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, M. Tegmark, S. J. Tingay, C. M. Trott, R. B. Wayth, R. L. Webster, A. Williams, C. L. Williams, and J. S. B. Wyithe. Confirmation of Wide-field Signatures in Redshifted 21 cm Power Spectra. *ApJ*, 807:L28, July 2015c. doi: 10.1088/2041-8205/807/2/L28.
- N. Thyagarajan, A. R. Parsons, D. R. DeBoer, J. D. Bowman, A. M. Ewall-Wice, A. R. Neben, and N. Patra. Effects of Antenna Beam Chromaticity on Redshifted 21 cm Power Spectrum and Implications for Hydrogen Epoch of Reionization Array. *ApJ*, 825:9, July 2016. doi: 10.3847/0004-637X/825/1/9.
- S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Boller, B. Crosse, R. B. Wayth, C. J. Lonsdale, S. Tremblay, D. Palot, T. Colegate, A. Wicenc, N. Kudryavtseva, W. Arcus, D. Barnes, G. Bernardi, F. Briggs, S. Burns, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. Deshpande, L. Desouza, B. M. Gaensler, L. J. Greenhill, P. J. Hall, B. J. Hazelton, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, M. J. Lynch, B. Mckinley, S. R. Mcwhirter, E. Morgan, D. Oberoi, J. Pathikulangara, T. Prabu, R. A. Remillard, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya-Shankar, F. Schlagenhauer, K. S. Srivani, J. Stevens, R. Subrahmanyam, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *PASA*, 30:e007, January 2013a. doi: 10.1017/pasa.2012.007.
- S. J. Tingay, D. Oberoi, I. Cairns, A. Donea, R. Duffin, W. Arcus, G. Bernardi, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. Desh-

- pande, L. deSouza, D. Emrich, B. M. Gaensler, G. R. L. J. Greenhill, B. J. Hazelton, D. Herne, J. N. Hewitt, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, J. A. Kennewell, B. B. Kincaid, R. Koenig, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, S. M. Ord, J. Pathikulangara, T. Prabu, R. A. Remillard, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya-Shankar, K. S. Srivani, J. Stevens, R. Subrahmanyam, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The Murchison Widefield Array: solar science with the low frequency SKA Precursor. In *Journal of Physics Conference Series*, volume 440 of *Journal of Physics Conference Series*, page 012033, June 2013b. doi: 10.1088/1742-6596/440/1/012033.
- S. J. Tingay et al. The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *PASA*, 30:e007, January 2013c. doi: 10.1017/pasa.2012.007.
- P. Tisserand, L. Le Guillou, C. Afonso, J. N. Albert, J. Andersen, R. Ansari, É. Aubourg, P. Bareyre, J. P. Beaulieu, X. Charlot, C. Coutures, R. Ferlet, P. Fouqué, J. F. Glicenstein, B. Goldman, A. Gould, D. Graff, M. Gros, J. Haissinski, C. Hamadache, J. de Kat, T. Lasserre, É. Lesquoy, C. Loup, C. Magneville, J. B. Marquette, É. Maurice, A. Maury, A. Milsztajn, M. Moniez, N. Palanque-Delabrouille, O. Perdureau, Y. R. Rahal, J. Rich, M. Spiro, A. Vidal-Madjar, L. Vigroux, S. Zylberajch, and EROS-2 Collaboration. Limits on the Macho content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *A&A*, 469: 387–404, July 2007. doi: 10.1051/0004-6361:20066017.
- J. E. Titheridge. The electron content of the southern mid-latitude ionosphere, 1965 - 1971. *Journal of Atmospheric and Terrestrial Physics*, 35:981–1001, May 1973.
- T. Totani, N. Kawai, G. Kosugi, K. Aoki, T. Yamada, M. Iye, K. Ohta, and T. Hattori. Implications for Cosmic Reionization from the Optical Afterglow Spectrum of the Gamma-Ray Burst 050904 at $z = 6.3$. *PASJ*, 58:485–498, June 2006.
- S. E. Tremblay, S. M. Ord, N. D. R. Bhat, S. J. Tingay, B. Crosse, D. Pallot, S. I. Oronsaye, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, B. E. Corey, A. A. Deshpande, D. Emrich, R. Goeke, L. J. Greenhill, B. J. Hazelton, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, T. Prabu, A. E. E. Rogers, A. Roshi, N. Udaya Shankar, K. S. Srivani, R. Subrahmanyam, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, and C. L. Williams. The High Time and Frequency Resolution Capabilities of the Murchison Widefield Array. *PASA*, 32:e005, February 2015. doi: 10.1017/pasa.2015.6.
- C. M. Trott. Exploring the evolution of reionization using a wavelet transform and the light cone effect. *MNRAS*, 461:126–135, September 2016. doi: 10.1093/mnras/stw1310.

- C. M. Trott and S. J. Tingay. The Effect of Interplanetary Scintillation on Epoch of Reionization Power Spectra. *ApJ*, 814:27, November 2015. doi: 10.1088/0004-637X/814/1/27.
- C. M. Trott and R. B. Wayth. Spectral Calibration Requirements of Radio Interferometers for Epoch of Reionisation Science with the SKA. *PASA*, 33:e019, May 2016. doi: 10.1017/pasa.2016.18.
- C. M. Trott, R. B. Wayth, and S. J. Tingay. The Impact of Point-source Subtraction Residuals on 21 cm Epoch of Reionization Estimation. *ApJ*, 757:101, September 2012. doi: 10.1088/0004-637X/757/1/101.
- C. M. Trott, B. Pindor, P. Procopio, R. B. Wayth, D. A. Mitchell, B. McKinley, S. J. Tingay, N. Barry, A. P. Beardsley, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, P. Carroll, A. de Oliveira-Costa, J. S. Dillon, A. Ewall-Wice, L. Feng, L. J. Greenhill, B. J. Hazelton, J. N. Hewitt, N. Hurley-Walker, M. Johnston-Hollitt, D. C. Jacobs, D. L. Kaplan, H. S. Kim, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. F. Morales, E. Morgan, A. R. Neben, N. Thyagarajan, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, J. C. Pober, T. Prabu, J. Ridding, N. Udaya Shankar, S. K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, M. Tegmark, R. L. Webster, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. CHIPS: The Cosmological H I Power Spectrum Estimator. *ApJ*, 818:139, February 2016. doi: 10.3847/0004-637X/818/2/139.
- M. Valdés, C. Evoli, A. Mesinger, A. Ferrara, and N. Yoshida. The nature of dark matter from the global high-redshift H I 21 cm signal. *MNRAS*, 429:1705–1716, February 2013. doi: 10.1093/mnras/sts458.
- M. P. van Haarlem et al. LOFAR: The LOw-Frequency ARray. *A&A*, 556:A2, August 2013. doi: 10.1051/0004-6361/201220873.
- H. van Woerden and R. G. Strom. The beginnings of radio astronomy in the Netherlands. *Journal of Astronomical History and Heritage*, 9:3–20, June 2006.
- H. Vedantham, N. Udaya Shankar, and R. Subrahmanyam. Imaging the Epoch of Reionization: Limitations from Foreground Confusion and Imaging Algorithms. *ApJ*, 745:176, February 2012. doi: 10.1088/0004-637X/745/2/176.
- H. K. Vedantham and L. V. E. Koopmans. Scintillation noise power spectrum and its impact on high-redshift 21-cm observations. *MNRAS*, 458:3099–3117, May 2016. doi: 10.1093/mnras/stw443.
- T. Venumadhav, T.-C. Chang, O. Doré, and C. M. Hirata. A practical theorem on using interferometry to measure the global 21-cm signal. *ArXiv e-prints*, December 2015.
- M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt. Warm Dark Matter as a solution to the small scale crisis: new constraints from high redshift Lyman-alpha forest data. *ArXiv e-prints:1306.2314*, June 2013.

- M. S. Vogele and A. S. Szalay. Eigenmode Analysis of Galaxy Redshift Surveys. I. Theory and Methods. *ApJ*, 465:34, July 1996. doi: 10.1086/177399.
- M. Vogelsberger, J. Zavala, and A. Loeb. Subhaloes in self-interacting galactic dark matter haloes. *MNRAS*, 423:3740–3752, July 2012. doi: 10.1111/j.1365-2966.2012.21182.x.
- M. Volonteri and M. J. Rees. Quasars at $z=6$: The Survival of the Fittest. *ApJ*, 650:669–678, October 2006. doi: 10.1086/507444.
- T. C. Voytek, A. Natarajan, J. M. Jáuregui García, J. B. Peterson, and O. López-Cruz. Probing the Dark Ages at $z \sim 20$: The SCI-HI 21 cm All-sky Spectrum Experiment. *ApJ*, 782:L9, February 2014. doi: 10.1088/2041-8205/782/1/L9.
- T. P. Walker, G. Steigman, H.-S. Kang, D. M. Schramm, and K. A. Olive. Primordial nucleosynthesis redux. *ApJ*, 376:51–69, July 1991. doi: 10.1086/170255.
- X. Wang, M. Tegmark, M. G. Santos, and L. Knox. 21 cm Tomography with Foregrounds. *ApJ*, 650:529–537, October 2006. doi: 10.1086/506597.
- R. B. Wayth, E. Lenc, M. E. Bell, J. R. Callingham, K. S. Dwarakanath, T. M. O. Franzen, B.-Q. For, B. Gaensler, P. Hancock, L. Hindson, N. Hurley-Walker, C. A. Jackson, M. Johnston-Hollitt, A. D. Kapińska, B. McKinley, J. Morgan, A. R. Ofringa, P. Procopio, L. Staveley-Smith, C. Wu, Q. Zheng, C. M. Trott, G. Bernardi, J. D. Bowman, F. Briggs, R. J. Cappallo, B. E. Corey, A. A. Deshpande, D. Emrich, R. Goeke, L. J. Greenhill, B. J. Hazelton, D. L. Kaplan, J. C. Kasper, E. Kratzenberg, C. J. Lonsdale, M. J. Lynch, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, A. E. E. Rogers, A. Rosh, N. U. Shankar, K. S. Srivani, R. Subrahmanyam, S. J. Tingay, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, and C. L. Williams. GLEAM: The GaLactic and Extragalactic All-Sky MWA Survey. *PASA*, 32:e025, June 2015. doi: 10.1017/pasa.2015.26.
- S. Weinberg. *Cosmology*. Oxford University Press, 2008.
- S. M. Weinmann and S. J. Lilly. The Number and Observability of Population III Supernovae at High Redshifts. *ApJ*, 624:526–531, May 2005. doi: 10.1086/428106.
- D. Whalen, T. Abel, and M. L. Norman. Radiation Hydrodynamic Evolution of Primordial H II Regions. *ApJ*, 610:14–22, July 2004. doi: 10.1086/421548.
- R. L. White, R. H. Becker, X. Fan, and M. A. Strauss. Probing the Ionization State of the Universe at $z>6$. *AJ*, 126:1–14, July 2003. doi: 10.1086/375547.
- M. H. Wieringa. An investigation of the telescope based calibration methods 'redundancy' and 'self-cal'. *Experimental Astronomy*, 2:203–225, 1992. doi: 10.1007/BF00420576.

- W. L. Williams, R. J. van Weeren, H. J. A. Röttgering, P. Best, T. J. Dijkema, F. de Gasperin, M. J. Hardcastle, G. Heald, I. Prandoni, J. Sabater, T. W. Shimwell, C. Tasse, I. M. van Bemmell, M. Brüggel, G. Brunetti, J. E. Conway, T. Enßlin, D. Engels, H. Falcke, C. Ferrari, M. Haverkorn, N. Jackson, M. J. Jarvis, A. D. Kapinska, E. K. Mahony, G. K. Miley, L. K. Morabito, R. Morganti, E. Orrù, E. Retana-Montenegro, S. S. Sridhar, M. C. Toribio, G. J. White, M. W. Wise, and J. T. L. Zwart. LOFAR 150-MHz observations of the Bo\`otes field: Catalogue and Source Counts. *ArXiv e-prints*, May 2016.
- C. J. Willott, S. Rawlings, K. M. Blundell, M. Lacy, and S. A. Eales. The radio luminosity function from the low-frequency 3CRR, 6CE and 7CRS complete samples. *MNRAS*, 322:536–552, April 2001. doi: 10.1046/j.1365-8711.2001.04101.x.
- R. J. Wilman, L. Miller, M. J. Jarvis, T. Mauch, F. Levrier, F. B. Abdalla, S. Rawlings, H.-R. Klöckner, D. Obreschkow, D. Olteanu, and S. Young. A semi-empirical simulation of the extragalactic radio continuum sky for next generation radio telescopes. *MNRAS*, 388:1335–1348, August 2008. doi: 10.1111/j.1365-2966.2008.13486.x.
- T. L. Wilson, K. Rohlfs, and S. Hüttemeister. *Tools of Radio Astronomy*. Springer-Verlag, 2009. doi: 10.1007/978-3-540-85122-6.
- R. A. Windhorst, N. P. Hathi, S. H. Cohen, R. A. Jansen, D. Kawata, S. P. Driver, and B. Gibson. High resolution science with high redshift galaxies. *Advances in Space Research*, 41:1965–1971, 2008. doi: 10.1016/j.asr.2007.07.005.
- J. H. Wise, M. J. Turk, M. L. Norman, and T. Abel. The Birth of a Galaxy: Primordial Metal Enrichment and Stellar Populations. *ApJ*, 745:50, January 2012. doi: 10.1088/0004-637X/745/1/50.
- D. M. Wittman, J. A. Tyson, D. Kirkman, I. Dell’Antonio, and G. Bernstein. Detection of weak gravitational lensing distortions of distant galaxies by cosmic dark matter at large scales. *Nature*, 405:143–148, May 2000.
- S. E. Woosley, A. Heger, and T. A. Weaver. The evolution and explosion of massive stars. *Reviews of Modern Physics*, 74:1015–1071, November 2002. doi: 10.1103/RevModPhys.74.1015.
- S. A. Wouthuysen. On the excitation mechanism of the 21-cm (radio-frequency) interstellar hydrogen emission line. *AJ*, 57:31–32, 1952. doi: 10.1086/106661.
- X.-B. Wu, F. Wang, X. Fan, W. Yi, W. Zuo, F. Bian, L. Jiang, I. D. McGreer, R. Wang, J. Yang, Q. Yang, D. Thompson, and Y. Beletsky. An ultraluminous quasar with a twelve-billion-solar-mass black hole at redshift 6.30. *Nature*, 518:512–515, February 2015. doi: 10.1038/nature14241.
- J. S. B. Wyithe and A. Loeb. Self-regulated Growth of Supermassive Black Holes in Galaxies as the Origin of the Optical and X-Ray Luminosity Functions of Quasars. *ApJ*, 595:614–623, October 2003. doi: 10.1086/377475.

- H. Xu, M. L. Norman, B. W. O'Shea, and J. H. Wise. Late Pop III Star Formation During the Epoch of Reionization: Results from the Renaissance Simulations. *ApJ*, 823:140, June 2016. doi: 10.3847/0004-637X/823/2/140.
- Y. Xu, X. Chen, Z. Fan, H. Trac, and R. Cen. The 21 cm Forest as a Probe of the Reionization and The Temperature of the Intergalactic Medium. *ApJ*, 704: 1396–1404, October 2009. doi: 10.1088/0004-637X/704/2/1396.
- Y. Q. Xue et al. The Chandra Deep Field-South Survey: 4 Ms Source Catalogs. *ApJS*, 195:10, July 2011. doi: 10.1088/0067-0049/195/1/10.
- K. C. Yeh and B. J. Flaherty. Ionospheric electron content at temperate latitudes during the declining phase of the sunspot cycle. *J. Geophys. Res.*, 71:4557–4570, October 1966. doi: 10.1029/JZ071i019p04557.
- D. G. York, J. Adelman, J. E. Anderson, Jr., S. F. Anderson, J. Annis, N. A. Bahcall, J. A. Bakken, R. Barkhouser, S. Bastian, E. Berman, W. N. Boroski, S. Bracker, C. Briegel, J. W. Briggs, J. Brinkmann, R. Brunner, S. Burles, L. Carey, M. A. Carr, F. J. Castander, B. Chen, P. L. Colestock, A. J. Connolly, J. H. Crocker, I. Csabai, P. C. Czarapata, J. E. Davis, M. Doi, T. Dombeck, D. Eisenstein, N. Ellman, B. R. Elms, M. L. Evans, X. Fan, G. R. Federwitz, L. Fiscelli, S. Friedman, J. A. Frieman, M. Fukugita, B. Gillespie, J. E. Gunn, V. K. Gurbani, E. de Haas, M. Haldeman, F. H. Harris, J. Hayes, T. M. Heckman, G. S. Hennessy, R. B. Hindsley, S. Holm, D. J. Holmgren, C.-h. Huang, C. Hull, D. Husby, S.-I. Ichikawa, T. Ichikawa, Z. Ivezić, S. Kent, R. S. J. Kim, E. Kinney, M. Klaene, A. N. Kleinman, S. Kleinman, G. R. Knapp, J. Korienek, R. G. Kron, P. Z. Kunzst, D. Q. Lamb, B. Lee, R. F. Leger, S. Limmongkol, C. Lindenmeyer, D. C. Long, C. Loomis, J. Loveday, R. Lucinio, R. H. Lupton, B. MacKinnon, E. J. Mannery, P. M. Mantsch, B. Margon, P. McGehee, T. A. McKay, A. Meiksin, A. Merelli, D. G. Monet, J. A. Munn, V. K. Narayanan, T. Nash, E. Neilsen, R. Neswold, H. J. Newberg, R. C. Nichol, T. Nicinski, M. Nonino, N. Okada, S. Okamura, J. P. Ostriker, R. Owen, A. G. Pauls, J. Peoples, R. L. Peterson, D. Petravick, J. R. Pier, A. Pope, R. Pordes, A. Prosapio, R. Rechenmacher, T. R. Quinn, G. T. Richards, M. W. Richmond, C. H. Rivetta, C. M. Rockosi, K. Ruthmansdorfer, D. Sandford, D. J. Schlegel, D. P. Schneider, M. Sekiguchi, G. Sergey, K. Shimasaku, W. A. Siegmund, S. Smee, J. A. Smith, S. Snedden, R. Stone, C. Stoughton, M. A. Strauss, C. Stubbs, M. SubbaRao, A. S. Szalay, I. Szapudi, G. P. Szokoly, A. R. Thakar, C. Tremonti, D. L. Tucker, A. Uomoto, D. Vanden Berk, M. S. Vogeley, P. Waddell, S.-i. Wang, M. Watanabe, D. H. Weinberg, B. Yanny, N. Yasuda, and SDSS Collaboration. The Sloan Digital Sky Survey: Technical Summary. *AJ*, 120:1579–1587, September 2000. doi: 10.1086/301513.
- N. Yoshida, V. Bromm, and L. Hernquist. The Era of Massive Population III Stars: Cosmological Implications and Self-Termination. *ApJ*, 605:579–590, April 2004. doi: 10.1086/382499.

- S. Yoshiura, H. Shimabukuro, K. Takahashi, R. Momose, H. Nakanishi, and H. Imai. Sensitivity for 21 cm bispectrum from Epoch of Reionization. *MNRAS*, 451:266–274, July 2015. doi: 10.1093/mnras/stv855.
- O. Zahn, A. Mesinger, M. McQuinn, H. Trac, R. Cen, and L. E. Hernquist. Comparison of reionization models: radiative transfer simulations and approximate, seminumeric models. *MNRAS*, 414:727–738, June 2011. doi: 10.1111/j.1365-2966.2011.18439.x.
- M. Zaldarriaga and U. Seljak. Gravitational lensing effect on cosmic microwave background polarization. *Phys. Rev. D*, 58(2):023003, July 1998. doi: 10.1103/PhysRevD.58.023003.
- M. Zaldarriaga, S. R. Furlanetto, and L. Hernquist. 21 Centimeter Fluctuations from Cosmic Gas at High Redshifts. *ApJ*, 608:622–635, June 2004. doi: 10.1086/386327.
- Y. B. Zel’dovich. Gravitational instability: An approximate theory for large density perturbations. *A&A*, 5:84–89, March 1970.
- H. Zheng, M. Tegmark, V. Buza, J. S. Dillon, H. Gharibyan, J. Hickish, E. Kunz, A. Liu, J. Losh, A. Lutomirski, S. Morrison, S. Narayanan, A. Perko, D. Rosner, N. Sanchez, K. Schutz, S. M. Tribiano, M. Zaldarriaga, K. Zarb Adami, I. Zelko, K. Zheng, R. Armstrong, R. F. Bradley, M. R. Dexter, A. Ewall-Wice, A. Magro, M. Matejek, E. Morgan, A. R. Neben, Q. Pan, C. M. Peterson, M. Su, J. Villasenor, C. L. Williams, H.-I Yang, and Y. Zhu. Mapping our Universe in 3D with MITEoR. *ArXiv e-prints*, September 2013.
- H. Zheng, M. Tegmark, V. Buza, J. S. Dillon, H. Gharibyan, J. Hickish, E. Kunz, A. Liu, J. Losh, A. Lutomirski, S. Morrison, S. Narayanan, A. Perko, D. Rosner, N. Sanchez, K. Schutz, S. M. Tribiano, M. Valdez, H. Yang, K. Z. Adami, I. Zelko, K. Zheng, R. P. Armstrong, R. F. Bradley, M. R. Dexter, A. Ewall-Wice, A. Magro, M. Matejek, E. Morgan, A. R. Neben, Q. Pan, R. F. Penna, C. M. Peterson, M. Su, J. Villasenor, C. L. Williams, and Y. Zhu. MITEoR: a scalable interferometer for precision 21 cm cosmology. *MNRAS*, 445:1084–1103, December 2014. doi: 10.1093/mnras/stu1773.
- H. Zheng, M. Tegmark, J. S. Dillon, D. A. Kim, A. Liu, A. R. Neben, J. Jonas, P. Reich, and W. Reich. An improved model of diffuse galactic radio emission from 10 MHz to 5 THz. *MNRAS*, 464:3486–3497, January 2017a. doi: 10.1093/mnras/stw2525.
- H. Zheng, M. Tegmark, J. S. Dillon, A. Liu, A. R. Neben, S. M. Tribiano, R. F. Bradley, V. Buza, A. Ewall-Wice, H. Gharibyan, J. Hickish, E. Kunz, J. Losh, A. Lutomirski, E. Morgan, S. Narayanan, A. Perko, D. Rosner, N. Sanchez, K. Schutz, M. Valdez, J. Villasenor, H. Yang, K. Zarb Adami, I. Zelko, and K. Zheng. Brute-force mapmaking with compact interferometers: a MITEoR northern sky map from 128 to 175 MHz. *MNRAS*, 465:2901–2915, March 2017b. doi: 10.1093/mnras/stw2910.