The Murchison Widefield Array 21cm Epoch of Reionization Experiment: Design, Construction, and First Season Results

Adam Beardsley

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2015

Reading Committee:

Miguel F. Morales, Chair

Thomas R. Quinn

Matthew J. McQuinn

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

The Murchison Widefield Array 21cm Epoch of Reionization Experiment: Design, Construction, and First Season Results

Adam Beardsley

Chair of the Supervisory Committee:
Professor Miguel F. Morales
Physics

The Cosmic Dark Ages and the Epoch of Reionization (EoR) remain largely unexplored chapters in the history and evolution of the Universe. These periods hold the potential to inform our picture of the cosmos similar to what the Cosmic Microwave Background has done over the past several decades. A promising method to probe the neutral hydrogen gas between early galaxies is known as 21cm tomography, which utilizes the ubiquitous hyper-fine transition of HI to create 3D maps of the intergalactic medium.

The Murchison Widefield Array (MWA) is an instrument built with a primary science driver to detect and characterize the EoR through 21cm tomography. In this thesis we explore the challenges faced by the MWA from the layout of antennas, to a custom analysis pipeline, to bridging the gap with probes at other wavelengths. We discuss many lessons learned in the course of reducing MWA data with an extremely precise measurement in mind, and conclude with the first deep integration from array. We present a 2-σ upper limit on the EoR power spectrum of $\Delta^2(k) < 1.25 \times 10^4$ mK^2 at cosmic scale $k = 0.236$ h Mpc^{-1} and redshift $z = 6.8$. Our result is a marginal improvement over previous MWA results and consistent with the best published limits from other instruments. This result is the deepest imaging power spectrum to date, and is a major step forward for this type of analysis. While our limit is dominated by systematics, we offer strategies for improvement for future analysis.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

BAO: Baryon Acoustic Oscillation

CFT: Continuous Fourier Transform

CHIME: Canadian Hydrogen Intensity Mapping Experiment

CHIPS: Cosmological HI Power Spectrum

CLEAN: A deconvolution algorithm used in radio astronomy. Despite its stylization, it is not an acronym.

CMB: Cosmic Microwave Background

DARE: Dark Ages Radio Explorer

DEC.: Declination

DFT: Discrete Fourier Transform

EDGES: Experiment to Detect the Global EoR Signal

EOR: Epoch of Reionization

EOR0, EOR1, EOR2: Names for three sky fields chosen by the MWA to observe the EoR

εppsilon: Error Propagated Power Spectrum with Interleaved Observed Noise

FFT: Fast Fourier Transform

FHD: Fast Holographic Deconvolution

GMRT: Giant Metrewave Radio Telescope

GPU: Graphic Processing Unit

GSM: Global Sky Model

HEALPIX: Hierarchical Equal Area isoLatitude Pixelization

HERA: Hydrogen Epoch of Reionization Array

IGM: Intergalactic Medium

JWST: James Webb Space Telescope

LOFAR: LOw Frequency ARray

MRC: Molonglo Reference Catalog

MWA: Murchison Widefield Array

MWACS: Murchison Widefield Array Commissioning Survey

PAPER: Precision Array for Probing the Epoch of Reionization

PDF: Probability Distribution Function

PSF: Point Spread Function

RA: Right Ascension

RMS: Root Mean Square

RFI: Radio Frequency Interference

RTS: Real Time System

SKA: Square Kilometre Array

SNR: Signal to Noise Ratio

SUMSS: Sydney University Molonglo Sky Survey

UVFITS: UV Flexible Image Transport System

WMAP: Wilkinson Microwave Anisotropy Probe

ΛCDM: Lambda Cold Dark Matter

ACKNOWLEDGMENTS

This thesis is the culmination of my work over several years, and I could not have done it without the love and support of many people to whom I owe many thanks. First I would like to thank my family, especially my parents, Peter and Crystal, and four brothers, Brian, Kevin, Erik, and Ben, who have supported me while I pursue a career in science.

Next I would like to thank my advisor, Miguel Morales, for teaching me the art of radio interferometry. I can only hope to ever attain an intuition for radio instruments that he seemingly effortlessly possesses. He also holds an amazing aptitude for understanding the struggle of the graduate student and enabled me to produce my best work while not losing my sanity.

Thank you to the postdocs I have worked with – Bryna Hazelton, Ian Sullivan, Jonnie Pober, and Danny Jacobs – who have been infinitely patient and helpful with my endless questions. It has been a pleasure to work with the members of the UW radio cosmology group past and present: Patti Carroll, Christian Boutan, Alex Bastidas Fry, Nichole Barry, Brynn MacCoy, Lisa McBride, Arielle Steger, Jess Murray, Zach Martinot, Zac Banks, and Nancy Sackman. I am also greatly appreciative for the opportunity to work the MWA collaboration as a whole.

I am extremely grateful to have found a very supportive network of friends in Seattle. I have enjoyed many eating adventures, climbing sessions, volleyball games, and grumpy hours with Jennie Chen, Max Hansen, Matt Hoffman, Tom Noel, Brian Burkholder, and Carolyn Auchter. In the depths of the basement I was lucky to find refuge with the Blinov Atomic Coffee Drinkers Guild, which allowed me to keep up spirits. Thank you to Amy Robertson for helping me find a non-academic outlet my first year with volunteer opportunities, and

helping me secure a teaching experience at Seattle Pacific University. A special thank you to fellow Badgers that found their way out to Seattle, and made holidays feel a little more like home: Matt Hartig, Greta Block, Dhan Rao, Elise Gale, and Megan Neitzel.

The MWA collaboration acknowledges the Wajarri Yamatji people as the traditional owners of the observatory site.

DEDICATION

To my wonderful fiancée Jackie Malsch. She has given me unyielding support and understanding as I venture on this journey. I cannot wait to begin the next chapter of our lives together.

Chapter 1

INTRODUCTION

"Begin at the beginning," the King said, very gravely, "and go on till you come to the end: then stop."

– Lewis Carroll, *Alice in Wonderland*

Our story begins at the end of a major phase change of the Universe. Approximately 370,000 years after the big bang, the Universe was hot, dense, and extremely ionized. Photons were coupled to matter making the pre-galactic medium completely opaque. A cosmic instant later the Universe's expansion caused it to cool sufficiently to the point where nearly all protons recombined with electrons forming neutral hydrogen, decoupling the photons which could then propagate freely. Those photons that were released at the so-called surface of last scattering are observed today, more than thirteen billion years later, as the cosmic microwave background (CMB). Studies of the CMB over the last several decades have been extremely productive in refining our model of how the Universe began, and how large scale structure formed from tiny fluctuations imprinted in the early Universe. However, the CMB is the last direct signal that has been detected to date from this period of cosmic history.

The time between the recombination event and the galaxies we see around us today is often categorized into two distinct eras – the Dark Ages and the Epoch of Reionization. Both periods remain largely unexplored, and hold the potential to inform our picture of the cosmos in ways analogous to what the CMB has done over the past several decades.

Immediately following decoupling of the CMB photons, the Universe was transparent and dark – photons were free to propagate, but few sources existed to create more radiation.

The small fluctuations we see in the CMB evolved linearly through gravitation. Due to few active components and relatively simple physics, structure evolution during these so-called Dark Ages can be well modeled [26]. Indeed direct observations of this era would further constrain cosmological models and offer immediate tests of fundamental physics. However, because the physics is so simple, direct detection of this era will be extremely challenging because there were no bright emission sources to detect.

When over densities grew sufficiently large, structure formation transitioned from linear to non-linear dynamics. Stars turned on, galaxies formed, and the radiation emitted from these new sources reionized the hydrogen residing in the intergalactic medium (IGM). This era is referred to as the Epoch of Reionization (EoR), and represents a major gap in our picture of the cosmic history.

1.1 The Epoch of Reionization

While the Dark Ages can be well modeled with standard cosmological parameters, the non-linear nature of the EoR makes it incredibly difficult to model and is yet very much unconstrained. The generally accepted picture of how the EoR unfolded is illustrated in Figure 1.1. When the first galaxies formed, they began emitting ultra-violet (UV) radiation, which ionized the IGM immediately surrounding themselves. As the immediate neighborhood of these over densities became ionized, it also became transparent to the UV photons, allowing them to propagate farther and create expanding “bubbles" centered around the matter over densities where they originated. The bubbles continued to grow until they overlapped with other bubbles, and eventually permeated the entirety of the IGM. What remains today is a fully ionized IGM with only small pockets of neutral gas residing in galaxies themselves.

The details of this picture are largely unconstrained – When did reionization begin? How long did it take to fully ionize the IGM? Did many small bubbles grow at once, or did few large bubble dominate the process? Did ionization affect galaxy formation?

By addressing these observational question, we can unlock the answers to more fundamental astrophysical questions such as: Were the reionizing sources stars, galaxies, quasars,

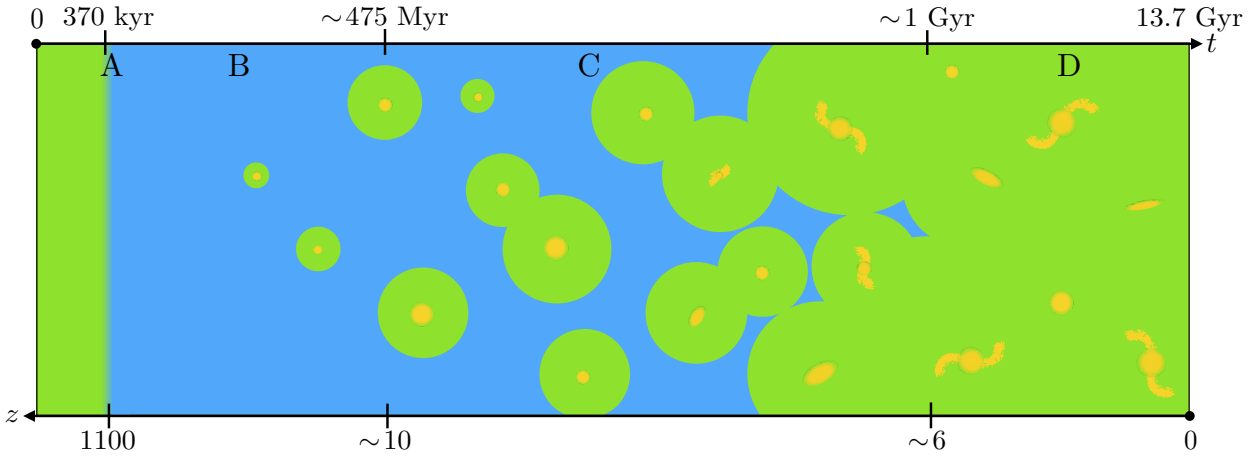

Figure 1.1: A cartoon visualization of major structure formation events in the history of the Universe. The top axis shows time since the Big Bang, increasing to the right. The bottom axis shows redshift, increasing to the left. Axes are highly non-linear, and stretched to highlight the periods of interest. Green represents ionized hydrogen while blue represents neutral. Before ∼370 thousand years, the Universe was entirely ionized, but at event (A) the pre-galactic medium cooled sufficiently to decouple the photons from baryonic matter, releasing what is observed today as the CMB. The time immediately following is referred to as the Dark Ages (B). Around 475 million years, stars and galaxies are ionizing the IGM around them initiating the Epoch of Reionization (C). The ionized bubbles continued to grow and overlap until the IGM became completely ionized around one billion years after the Big Bang. This results in the highly ionized Universe we see today with pockets of neutral gas residing in galaxies (D).

or some combination? What types of stars formed first in the Universe, and are there any fossils remaining today? and What role did feedback into the IGM play in galaxy formation?

Studies of the polarization of the CMB can place integrated constraints on the timing of reionization. Because the CMB is used as a backlight to the EoR, a simple model is used where reionization was instantaneous, and the optical depth can be translated into a reionization redshift. Using this model, the WMAP nine year results placed a constraint on the redshift of reionization at $z_{re} = 10.6 \pm 1.1$ [6]. The Planck satellite, launched in 2009, has been producing improved maps of the CMB, and released a constraint of $z_{re} = 11.1 \pm 1.1$ in 2013 [71]. This result, however, utilized the polarization maps from WMAP as the Planck maps were yet preliminary. The latest release from Planck includes more mature polarization maps, and the result is completely independent of WMAP polarization, with a reionization redshift of $z_{re} = 8.8^{+1.7}_{-1.4}$ [72].

Observations of highly redshifted quasars using the technique described in [32] have placed upper bounds on redshifts by which reionization is complete. Fan et al. [24] showed that reionization must be complete by $z \approx 6$, and Choudhury et al. [17] corroborated this low reionization redshift by observing a steep decline in Lyα emitting galaxies in the redshift range $6 \lesssim z \lesssim 8$. Deep optical and infrared galaxy surveys are beginning to reach the redshifts necessary to further constrain reionization (e.g. [10]), and the launch of the James Webb Space Telescope in 2018 should allow for very sensitive studies to constrain the EoR [27].

1.2 21cm Tomography

A promising direct probe of the EoR and Dark Ages is known as 21cm tomography – the science of exploiting the 21cm hyperfine transition of neutral hydrogen to create three dimensional maps of the universe (see [26] and [57] for reviews of the theory and observational prospects, respectively). There are two major advantages to this method – the atomic transition is very narrow and well understood, so the observed redshift maps directly to the Universal scale factor at emission, and therefore a line of sight distance; and hydrogen is

ubiquitous and dominates the baryonic matter.

The physics behind the 21cm hyperfine line is a subject covered in most undergraduate quantum mechanics classes. Due to the magnetic dipole interaction between the proton and electron in neutral hydrogen, the ground state is split into symmetric and anti-symmetric energy levels. The higher energy symmetric state can spontaneously decay to the lower energy anti-symmetric state, emitting a photon with wavelength 21.106 cm. This transition is extremely rare with a lifetime of $\sim 10^7$ years, but we are saved by the sheer amount of hydrogen that exists essentially everywhere in the Universe.

It has been shown that during the redshifts of interest ($z \sim 6-9$), the spin temperature of hydrogen, $T_{\rm S}$, will largely trace the kinetic temperature of the IGM (e.g. [78]), which in turn is much larger than the mean CMB temperature, $T_{\rm CMB}$. Under this assumption we can express the 21cm brightness temperature contrasted from the CMB at a given location in space $\mathbf{r}$ in terms of the local baryon over density, δ_ρ, the local ionization fraction, x_i, and a velocity distortion-sourced fluctuation, δ_v, which accounts for the peculiar velocity along the line of sight. These parameters combine to give the brightness temperature (see [26] for a complete derivation):

$$\delta T_{21}(\mathbf{r}) \equiv \frac{T_{\rm S} - T_{\rm CMB}(z)}{1+z}\left(1 - e^{-\tau_\nu}\right) \approx T_0(1 - x_i(\mathbf{r}))(1 + \delta_\rho(\mathbf{r}))(1 - \delta_v(\mathbf{r})) \qquad (1.1)$$

where τ_ν is the optical depth at frequency ν, and $T_0 = 28[(1+z)/10]^{1/2}$ mK. The fundamental aim of 21cm tomography is to create a three dimensional map of the IGM brightness temperature, which in turn probes the parameters above.

There are a couple qualitative features of Equation 1.1 that are worth mentioning before moving on. First, the baryon over density term shows a higher fluctuation for higher density (more hydrogen means more radiation). However, the most dense regions are likely to be the first to form galaxies and become ionized, causing the ionization term to vanish. Thus a brightness temperature map, though brighter at higher densities, is truly a map of neutral hydrogen and essentially a “negative" of what might be discovered through galaxy surveys.

Second, the physics involved with the 21cm fluctuations is isotropic, meaning that there

is no preferred direction. This symmetry is nearly broken for two reasons. The velocity distortions introduce a term which is dependent on the line of sight direction. However, we will see in Chapter 4 that we can account for this effect and retain the symmetry. Additionally, the fact that our line of sight probes not only a third spatial dimension, but also the time coordinate (farther away means earlier times), observations must be restricted to a cosmologically homogeneous line of sight distance range. With this restriction we can leverage the isotropic symmetry which will become invaluable in the following chapters.

The expected 21cm brightness temperature is extremely faint (on order tens of mK), especially when compared to the sky thermal noise (on order hundreds of K) [12]. Consequently, the first generation experiments will not have the sensitivity to produce high fidelity temperature maps, but have instead opted for a statistical measurement in the form of a power spectrum. We define the Fourier transform[1] of the spatial brightness temperature as

$$\widetilde{T}_{21}(\mathbf{k}) \equiv \mathcal{FT}\left[\delta T_{21}(\mathbf{r})\right] = \int \mathrm{d}^3\mathbf{r}\, \delta T_{21}(\mathbf{r}) e^{-i\mathbf{k}\cdot\mathbf{r}} \tag{1.2}$$

where the integral is done over all space, and $\mathbf{k}$ is the three dimensional wavenumber Fourier dual to the position vector $\mathbf{r}$. The power spectrum is then defined by

$$\left\langle \widetilde{T}_{21}(\mathbf{k}) \widetilde{T}^*_{21}(\mathbf{k}') \right\rangle \equiv (2\pi)^3 \delta_D(\mathbf{k}-\mathbf{k}') P_{21}(\mathbf{k}), \tag{1.3}$$

where the angle brackets, $\langle ... \rangle$, represent an ensemble average over many realizations of the Universe, and δ_D is the Dirac delta function.

Observationally we cannot take an ensemble average of many Universes, but we can take advantage of the symmetry discussed above. Because the physics is isotropic the power spectrum is dependent on the magnitude k only ($P_{21}(\mathbf{k}) \rightarrow P_{21}(k)$). Therefore, one can average the measured power spectrum in spherical shells of equal magnitude k to reduce sample variance.

---

[1]There are many conventions for the Fourier transform, and indeed several are chosen throughout the literature. I will attempt to remain consistent throughout this work, however there are places where my definition will need to change to match appropriate coordinates. Please see Appendix A for a summary of conventions and how to normalize with discrete transform numerical packages.

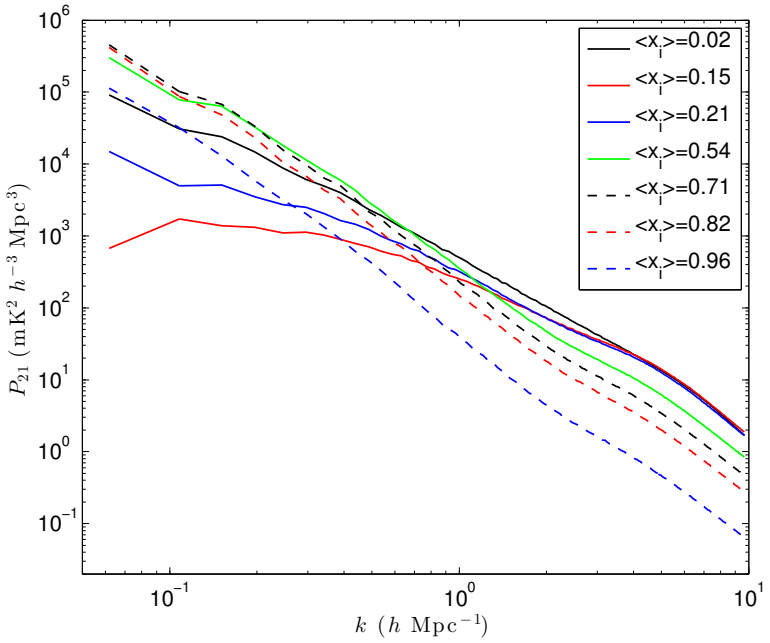

Figure 1.2: The evolution of the 21cm power spectrum as a function of mean ionization fraction, adapted from [46]. The signal is easily strongest at low k, or large spatial scales. When reionization begins the power spectrum takes a quick dip as the coherent large scale modes are wiped out by small bubbles (solid red line). As the bubbles grow, the power spectrum rises again and peaks around $\langle x_i \rangle \sim 0.71$, only to ultimately fall as the amount of hydrogen drops. The dimensionless Hubble parameter, h, in our units is defined such that the Hubble constant is $H_0 = h \times 100$ km s^{-1} Mpc^{-1}. Although the value of h is fairly well known to be about 0.68 [72], it is often left arbitrary in the literature and we do the same here.

Figure 1.2 shows the expected evolution of the 21cm power spectrum, as seen by simulations [46]. We can see that the power is highest at small k, or large spatial scales. The signal peaks in time when the mean ionization fraction is around 0.71. Simulations such as this help guide the experimental strategy to optimize both instrument design and observing choices when hunting for the EoR.

1.3 Challenges

The main challenges facing an EoR 21cm power spectrum detection fall under one theme: *precision*. The signal we seek is minuscule compared to both the thermal noise and the sea of foregrounds that cover the sky. Here I describe a few key challenges and our strategy to overcome them.

As already alluded to, the thermal sky noise at the relevant frequencies ($\sim$150-200 MHz) is several orders of magnitude brighter than the highly redshifted 21cm signal. This can be addressed in essentially three ways: perform a statistical power spectrum measurement, integrate a measurement for a very long time, and repeat the measurement many times. The latter two strategies are the focus of Chapters 3 and 4.

Foregrounds include synchrotron and electron free-free radiation from extragalactic radio galaxies as well as our own Milky Way, and are 4-5 orders of magnitude brighter than the expected signal[13]. Unlike the thermal noise, foregrounds will not integrate down and so a different tact is necessary. There are two ways of overcoming this challenge: foreground avoidance and foreground subtraction. For this work we implement both strategies. The regions of sky in which we will attempt to observe the EoR were chosen as relatively quiet in terms of foregrounds. However, as we will see in the coming chapters, we are still dominated by foregrounds in every single pixel of our maps, in just two minutes of integration. Further foreground avoidance is possible by taking advantage of symmetries and filtering sources out of the data by selecting pixels in Fourier space which are uncontaminated. This will be discussed in detail in Chapter 2. Foreground subtraction involves the high-precision modeling of sources and subtracting them directly from the data. The details of our subtraction is

described in Chapters 6 and 7.

The final challenge I will mention here is in the analysis. Because the measurement is so precise, any errors in the analysis pipeline can be catastrophic for the result. It is therefore incredibly important to use extreme caution along the way and have tests throughout to assure errors are not being injected into the data. It turns out that the power spectrum itself is an extremely sensitive diagnostic to such errors, and so a thorough understanding on how errors manifest there is very valuable in eliminating these systematics. This method of refining our analysis is illustrated in Chapter 7.

1.4 The Murchison Widefield Array

With an understanding of the Epoch of Reionization and a promising technique for detection, we turn to experiments underway. Several radio instruments have recently come online designed to make an EoR detection, including the GMRT (Giant Metrewave Radio Telescope [65]), LOFAR (LOw Frequency ARray [95]), PAPER (Precision Array for Probing the Epoch of Reionization [69]), and the MWA (Murchison Widefield Array [90]). Each of these experiments approaches the problem in a unique way, and all will inform future generations of instruments, such as HERA (Hydrogen Epoch of Reionization Array, DeBoer et al. 2015, in prep). This work focuses on early power spectrum results specifically from the MWA.

The Murchison Widefield Array (MWA) is a radio interferometer built in the Western Australian outback with the primary science goal of detecting the 21cm EoR power spectrum. While primarily an EoR instrument, the MWA also serves as a general observatory for several science programs including galactic and extragalactic surveys, time domain astrophysics, solar monitoring, and ionospheric science. The technical design of the array is detailed in Tingay et al. [90] and the science capabilities are described in Bowman et al. [14].

The remote location of the array was chosen to avoid human-made radio frequency interference (RFI). The FM radio band, as well as several TV station bands lie within the MWA bandwidth, so it was necessary to locate the array away from populated areas. A recent study of the RFI environment in Murchison can be found in Offringa et al. [63].

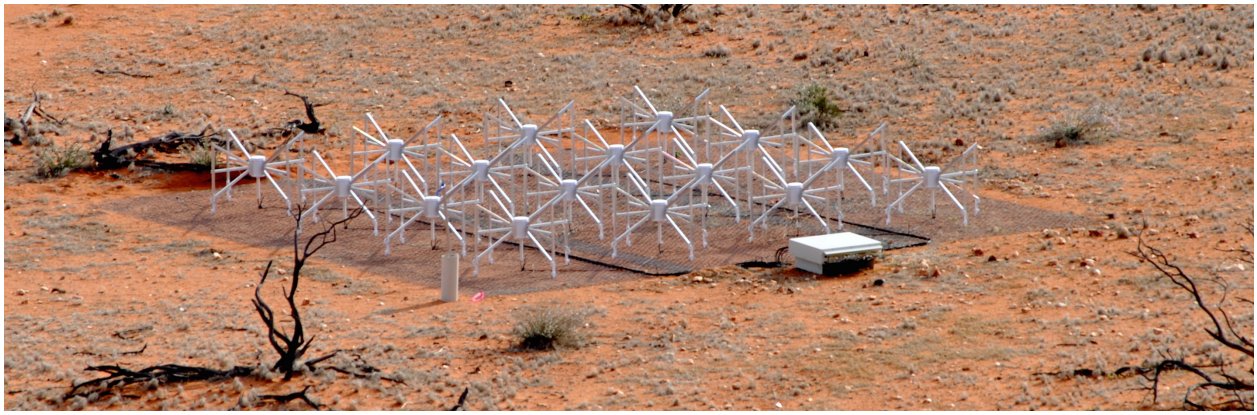

Figure 1.3: One of 128 MWA tiles placed in the Australian outback. The signals from the sixteen dual polarization antennas are combined in the analog beamformer shown on the right edge of the tile. The signal is then passed to the digital receivers for further processing.

The collecting elements of the MWA – referred to as antennas or “tiles" – comprise 16 dual polarization dipoles placed on a regular grid (spaced by 1.1 m), lying on a 5m×5m ground screen. Each dipole is sensitive to the entire sky and is optimized to operate in the frequency range 80-300 MHz. The radio frequency (RF) signal from each dipole for a given polarization is combined in an analog beamformer which uses physical delays to “point" the tile to different regions of the sky. The pointing of the tile both directs the response to a region of interest and reduces the field of view. The response function of the combined signals is referred to as the primary beam. The relatively small aperture to wavelength ratio means that the telescope has a wide field of view – about 25 degrees at 150 MHz.

The MWA contains 128 such tiles, which are placed in a pseudorandom way to optimize for science and obey physical constraints. The algorithm used for tile placement is the topic of Chapter 3. The center of the array contains a tightly packed core about 50m in radius, which is necessary for a sensitive EoR measurement. The remainder of the tiles extend out to radius of 1.5 km to provide higher resolution for calibration and source subtraction, as well as benefits to other science programs.

The RF signals from each tile are transmitted to digital receivers in the field which service eight tiles each (for a total of 16 receivers) [76]. Here the signal is digitized and filtered into 24 $\times$ 1.28 MHz frequency channels. These 24 sub-bands (total of 30.72 MHz) are chosen when the observation is scheduled, and do not need to be contiguous (though usually are). The signals from the 24 sub-bands $\times$ 8 tiles $\times$ 2 polarizations for each receiver are then transmitted to the correlator over fiber optic link.

The correlator performs a secondary frequency channelization into 10 kHz channels, then cross-multiplies the signal between all tiles for each channel. These cross correlations are the *visibilities* which is the fundamental measurement of a radio interferometer. The correlator next performs a frequency and time average before writing the visibilities to disk. The details of this averaging are determined by the observer, but a typical cadence is 2 second integration at 40 kHz frequency resolution.

The data are then written to disk and stored in two mirrored archives – one connected to site with a dedicated 10-Gbps optical fiber link to the Pawsey High Performance Computing Centre for SKA science in Perth, Western Australia, and the other on a dedicated archive and computing cluster at MIT in Cambridge, MA. To date the MWA has collected nearly 1,700 hours of EoR observations, amounting to about three petabytes of raw data.

Armed with a telescope that has been observing for a year and a half, it's time to dig into the details and pull out a limit on the Epoch of Reionization 21cm power spectrum.

Chapter 2

MEASUREMENT THEORY

"When you see someone putting on his Big Boots, you can be pretty sure that an Adventure is going to happen."

– A.A. Milne, *Winnie-the-Pooh*

Before we can proceed on our EoR adventure, we must equip ourselves appropriately. It is beneficial to have a discussion on the fundamentals of radio interferometers, which will serve to both set up a framework for the analysis, and to establish notation.

2.1 Response of Radio Interferometers

As mentioned in the previous chapter, the fundamental measurement of an interferometer is the visibilities – the cross correlation of channelized electric signal from all pairs of antennas.

$$V_{ij}(\nu,\tau) \equiv \left\langle E_i(\nu,t)E_j^*(\nu,t)\right\rangle_t \tag{2.1}$$

Here we use subscripts to denote the antennas, ν is the frequency, and the angle brackets represent a time average. Although we are averaging in time, the visibilities are read out on a finite cadence, and they still have a (coarser) time dependence denoted by τ. The time dependence of visibilities are due to thermal noise, the rotation of the earth, and any instrumental time dependence such as gain changes. For this chapter we will neglect noise and assume that our observations are sufficiently short in time to omit the other factors.

The response of the interferometer to a sky intensity, $I(\boldsymbol{\theta},\nu)$, can be expressed as [87]

$$V_{ij}(\nu) = \int \mathrm{d}^2\boldsymbol{\theta}\, A_{ij}(\boldsymbol{\theta},\nu)I(\boldsymbol{\theta},\nu)e^{-2\pi i\nu(\mathbf{r}_i-\mathbf{r}_j)\cdot\boldsymbol{\theta}/c}, \tag{2.2}$$

where c is the speed of light, $A_{ij}(\boldsymbol{\theta}, \nu)$ is the primary beam (angular and frequency response) associated with the antenna pair, and $\mathbf{r}_i$ and $\mathbf{r}_j$ are the positions of the antennas on the ground. The absolute positions of the collecting elements never appear in the measurement equation, and so it is convenient to define the *baseline* as the relative position between any pair of antennas ($\mathbf{b}_{ij} \equiv \mathbf{r}_i - \mathbf{r}_j$).

As a conceptual exercise, we will consider the case where $A_{ij} = 1$ for all (i, j), $\boldsymbol{\theta}$, and ν. In this case, Equation 2.2 becomes

$$V_{ij}(\nu) \to \int \mathrm{d}^2\boldsymbol{\theta}\, I(\boldsymbol{\theta}, \nu) e^{-2\pi i \nu \mathbf{b}_{ij} \cdot \boldsymbol{\theta}/c}. \tag{2.3}$$

At this stage we should notice that a given visibility is simply a sample of a Fourier transform of the sky intensity, albeit with a different convention from the one defined in Chapter 1. The relationship between the sky and the response of visibilities is sometimes referred to as the "miracle of radio interferometry", and it is the basis of synthesis imaging. If the interferometer can sufficiently sample the Fourier transform of the sky, we can invert the process and form an image.

We are now motivated to define a couple more coordinates and functions. First we define the coordinates that are the Fourier dual to $\boldsymbol{\theta}$. By inspection of Equation 2.3, we can see that the appropriate choice is $\mathbf{u} \equiv \mathbf{b}/\lambda$ (where $\lambda = c/\nu$ is the observed wavelength). Often the individual vector components are written as $\mathbf{u} = (u, v, w)$ or in the case of a planar array $\mathbf{u} \to (u, v)$, or simply uv. We can then define the Fourier transform of the sky intensity, which is referred to as the coherence function.

$$\widetilde{I}(\mathbf{u}, \nu) \equiv \int \mathrm{d}^2\boldsymbol{\theta}\, I(\boldsymbol{\theta}, \nu) e^{-2\pi i \mathbf{u} \cdot \boldsymbol{\theta}} \tag{2.4}$$

We can now think of the visibilities as a sampling function of $\widetilde{I}$, where a pair of antennas sample the position $\mathbf{b}_{ij} = \mathbf{u}\lambda$. We make careful note here that the sampling position depends on the wavelength, or frequency, of the visibility, which we will soon see leads to mode mixing in our analysis.

Next we aim to reintroduce the primary beam. Using the same physics as the classic single and double slit interference experiments, it can be shown that the response of a single

antenna on the sky is the Fourier transform of its response to electric field incident on the ground, which to first order can be thought of as the physical shape of the antenna (though it turns out this is not nearly a good enough approximation for the precision necessary to detect the EoR). It follows that the the pair-wise primary beam, which is the product of the two antennas' sky responses, is the Fourier transform of the convolution of their ground-space responses. We can then express our sampling of the uv plane as a convolution of the pair-wise primary beam with the coherence function.

$$V_{ij}(\nu) = \int \mathrm{d}^2\mathbf{u}\, \widetilde{A}_{ij}(\mathbf{u} - \mathbf{u}_{ij}, \nu)\widetilde{I}(\mathbf{u}, \nu) \tag{2.5}$$

Another way to think about this operation is that the pair of antennas integrates the uv space over which it resides in order to estimate the value at its center, $\mathbf{u}_{ij}$.

With a set of measured visibilities, one can now create a rough image of the sky. Though there are better methods available, as will be shown in Section 6.2, we can make a first approximation to a sky here. We first note that the visibilities are measurements of the coherence function at the location of the baseline.

$$V_{ij}(\nu) \approx \widetilde{I}'(\mathbf{u}_{ij}, \nu) \tag{2.6}$$

I will use primes to denote measured values. We can then invert the Fourier transform in Equation 2.4.

$$I'(\boldsymbol{\theta}, \nu) \equiv \mathcal{F}\mathcal{T}^{-1}\left[\widetilde{I}'(\mathbf{u}_{ij}, \nu)\right] \approx \sum_{ij} V_{ij}(\nu) e^{2\pi i \mathbf{u}_{ij}\cdot\boldsymbol{\theta}} \tag{2.7}$$

While this imaging procedure is sufficient to introduce the concept of synthesis imaging, it is far from the best method. A direct discreet Fourier transform of the visibilities will have a "natural weighting" (heavier weight on uv modes that are more heavily sampled) which is difficult to adjust, it is not properly normalized, and is computationally very expensive. For the analysis pipeline described in Chapter 6, we will use the Fast Holographic Deconvolution [83] package, which includes an imaging strategy based on the optimal map-making techniques described in Morales and Matejek [56] and Bhatnagar, S. et al. [8], which in turn are based on CMB analysis [85, 86].

Our general strategy for forming a power spectrum will be to create the three dimensional map using the "miracle of radio interferometry" for the angular directions, and the frequency-to-redshift relationship for the third dimension. We can then perform a 3D Fourier transform, square, and average in spherical bins to arrive at the power spectrum. What will follow will be a hurricane of transforms and coordinate changes. It is helpful to spend some time describing the various spaces we will encounter and their relationships.

2.2 Spaces

The following discussion will step through several transformations, coordinate changes, and spaces involved in the analysis of an EoR power spectrum measurement. Figure 2.1 is offered as a visual guide through the Fourier labyrinth.

Let us take as our fundamental space a region of interest in the Universe where we hope to measure the intensity of 21cm radiation. We express the brightness temperature within this region as $I(\mathbf{r})$, where $\mathbf{r}$ are the normal cartesian coordinates (x, y, z), and have units of length. We choose to align the coordinate system such that the z axis is along our line of sight to the observation, often called the parallel direction and is denoted by $r_{||}$. Thus x and y are perpendicular to the line of sight, and are often referred to as such, denoted as $\mathbf{r}_{\perp}$. A convenient unit for $\mathbf{r}$ is the mega-parsec, Mpc, and $I(\mathbf{r})$ is typically expressed in mK. These coordinates and units are useful for cosmological parameters, but are not typically used in traditional radio astronomy. It is therefore helpful to translate these coordinates to a space more suited for the instrument when we do our measurement.

We can easily express the $\mathbf{r}$ coordinates in terms of angles on the sky, $\boldsymbol{\theta}$, and frequency of observation, ν [59].

$$x = D_M(z)\theta_x \tag{2.8a}$$

$$y = D_M(z)\theta_y \tag{2.8b}$$

$$\Delta z = \frac{c(1+z)^2}{H_0 \nu_{21} E(z)} \Delta\nu \tag{2.8c}$$

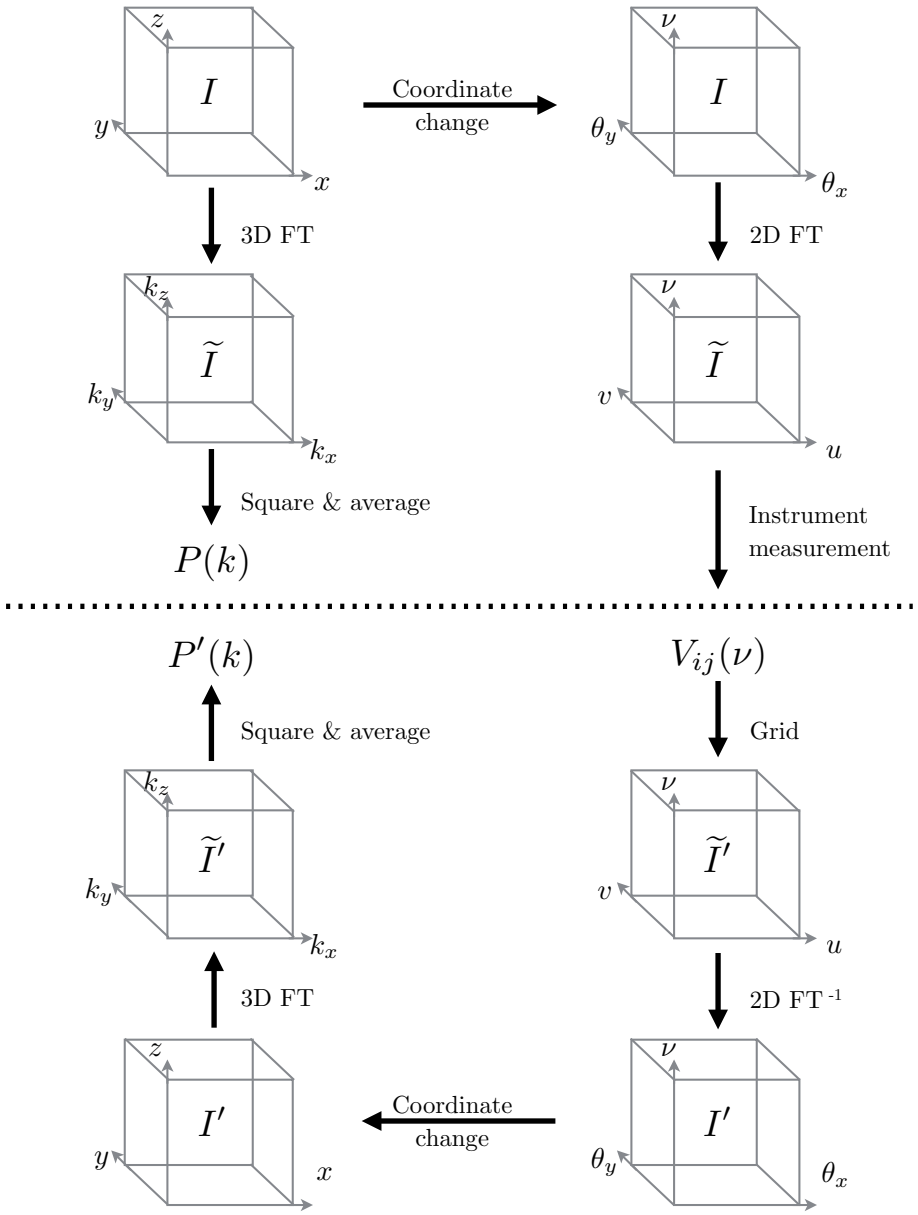

Figure 2.1: The spaces of an EoR power spectrum measurement. We begin with the representation of the true cube of sky in the upper left. We can follow the path straight down to form a power spectrum of the true sky, or we can follow the long path clockwise to send the the sky through the instrument to create a measured power spectrum. Everything above the dotted line is “true", and has not been measured, while everything below the dotted line has been through the instrument, and represents real data products. The left column is in **r** or **k** cosmological coordinates, while the right column is in coordinates that interface well with the instrument.

Here $D_M(z)$ is the comoving distance to the observed redshift, H_0 is the Hubble constant, ν_{21} is the emitted frequency of the 21cm transition, and $E(z) \equiv [\Omega_M(1+z)^3+\Omega_k(1+z)^2+\Omega_\Lambda]^{1/2}$ which is parameterized by the composition of the universe into matter (Ω_M), curvature (Ω_k), and dark energy (Ω_Λ). A pedagogical review of various coordinates for cosmology is available from Hogg [34].

We can also convert the units of temperature to flux density. The unit of choice here is Jy/str, where the conversion is

$$1\,\mathrm{mK} = 10^{-23}\frac{c^2\ \mathrm{str}}{2\nu^2 k_B}\ \mathrm{Jy/str}. \tag{2.9}$$

The result is the exact same brightness function we had before, but simply in different units and in a different coordinate system, $I(\boldsymbol{\theta}, \nu)$. But this version is ready to be absorbed into the measurement by the instrument.

Next we Fourier transform to the coherence function, $\widetilde{I}(\mathbf{u}, \nu)$. Note that even though I use a tilde to denote a Fourier space, this particular space is really a hybrid – we have not touched the line of sight direction, which remains in frequency, or equivalently image, space.[1] The uv coordinates were derived by considering the separation between two antennas, and thus the units can be thought of as a length, but measured in units of wavelengths. Equivalently, one can consider $\mathbf{u}$ as the Fourier dual to $\boldsymbol{\theta}$, which will lead to units of rad^{-1}. These are simply two interpretations of the same quantity and do not require any rescaling to express one or the other. The units of $\widetilde{I}(\mathbf{u}, \nu)$ are Jy.

We have already seen the process by which the baselines integrate a region of uv space to create visibilities, which also have units of Jy. Note that the integral of Equation 2.5 then implies that $\widetilde{A}_{ij}(\mathbf{u}, \nu)$ has units of λ^{-2} or str. Indeed, in our framework we force the normalization such that $\int \mathrm{d}^2\mathbf{u}\, \widetilde{A}_{ij}(\mathbf{u}, \nu) = 1$ for a zenith pointed beam, which enforces the desirable property that a 1 Jy point source at center will produce 1 Jy visibilities for all baselines.

[1] Perhaps it would be more appropriate to use a smaller tilde, $\tilde{I}(\mathbf{u}, \nu)$. But somehow I suspect that notation would be less than useful.

If we were to create an image at this point as in Equation 2.7, we would find that the intensity on the sky follows the response of the beam due to the convolution in Equation 2.5. This is referred to as the “apparent sky frame" as it is the sky that the instrument “sees" in the sense that it reflects how the sky intensity actually coupled into the collecting elements. The true sky frame can then be estimated by simply dividing by the primary beam in image space.

Returning to the visibilities, the next step is to make a better estimate of the coherence function with the measured visibilities. This step is often referred to as “gridding" because we will project the visibilities onto a regular grid so we can use the Fast Fourier Transform to create images, which will save orders of magnitude in computation compared to Equation 2.7. In line with optimal mapmaking, we will use the primary beam itself as the gridding kernel [85, 86, 56, 8]. This has the effect of convolving the data with another factor of the beam and is referred to as the holographic frame.

Regardless of the gridding details chosen, the result is the measured coherence function, $\widetilde{I}'(\mathbf{u}, \nu)$, which retains the units of Jy. We can now inverse Fourier transform to create our measured image $I'(\boldsymbol{\theta}, \nu)$. However, we are certainly not back to where we started. That tiny prime carries a lot of differences between the true sky image and our measured holographic sky image. We have two factors of the primary beam, which we can remove assuming a) we know the primary beam perfectly, and b) all antennas are exactly the same. In addition, when we sampled the uv plane with our visibilities, we did it imperfectly. It would be extremely costly to have a filled uv distribution, and so any realistic instrument will contain gaps and have finite extent. The result of this sampling is a point spread function (PSF), which effectively convolves the sky with the Fourier transform of the sampling function. This instrument-convolved image is called a dirty image.

When we discuss foreground subtraction we will utilize a sky model which will be subtracted from our data. A sky model typically consists of a list of point sources with location and flux, and a continuous map of diffuse emission. We can treat the model sky in the same fashion as the true sky, and simulate the measurement step (Equation 2.5) numerically using

a model of the instrument. One can then subtract the model visibilities from the data visibilities to create residual visibilities. However, each of the steps discussed are linear operations and the subtraction can happen at any point in the analysis (up to the squaring step below). Often it is useful to carry the model through the rest of the analysis to use as a diagnostic for systematic effects, as will be demonstrated in Chapter 7.

Next we can return to the cosmology-friendly coordinates, $I'(\boldsymbol{\theta}, \nu) \rightarrow I'(\mathbf{r})$, by simply inverting the process described at the beginning of this section. If we are satisfied with our measured image, we can next perform a three dimensional Fourier transform to wavenumber, $\mathbf{k}$, space.[2]

$$\widetilde{I}'(\mathbf{k}) \equiv \mathcal{FT}\left[I'(\mathbf{r})\right] = \int \mathrm{d}^3\mathbf{r}\, I'(\mathbf{r}) e^{-i\mathbf{k}\cdot\mathbf{r}} \tag{2.10}$$

The wavenumber $\mathbf{k}$ is the Fourier dual to $\mathbf{r}$ and has units of Mpc^{-1}, while $\widetilde{I}'(\mathbf{k})$ has units of mK Mpc^3.

There is an important point to make here that should not be overlooked. The baseline coordinates, $\mathbf{u}$, were Fourier dual to $\boldsymbol{\theta}$, which is a simple scaling to perpendicular position, $\mathbf{r}_\perp$, which in turn is Fourier dual to the perpendicular wavenumber, $\mathbf{k}_\perp$. All that is to say that the baselines of an interferometer actually probe specific wavenumber modes. In fact we can express a similar relationship as Equation 2.8.

$$\mathbf{k}_\perp = \frac{2\pi}{D_M(z)}\mathbf{u} \tag{2.11}$$

This fact is crucial to 21cm cosmological measurements, and drives not only instrument design (e.g. [3, 12, 51, 68, 75]), but also an entire school of analysis (e.g. [66, 70]).

The final step is to relate $\widetilde{I}'(\mathbf{k})$ to the power spectrum. Following Morales and Wyithe [57],

$$\left\langle \left|\widetilde{I}'(\mathbf{k})\right|^2 \right\rangle = \frac{1}{(2\pi)^3} \int \mathrm{d}^3\mathbf{k}'\, P(\mathbf{k}') \left|\widetilde{A}(\mathbf{k}-\mathbf{k}')\right|^2 . \tag{2.12}$$

[2]"But wait!" screams the attentive reader, "We just did an inverse Fourier transform! Why are we going back to Fourier space again?" Indeed at this stage it appears we are spinning our wheels to show off. But in practice the imaging step is helpful for combining observations from different times. We can account for the array rotating under the motion of the earth by gridding to a new uv plane on a relatively quick cadence (2 minutes in the case of the MWA), and combining the observations in image space.

We assume that the primary beam is compact and sharply peaked in k-space, so we approximate the power spectrum to be constant over the integral (e.g. [12]). We then arrive at our estimate of the power spectrum.

$$P'(\mathbf{k}) = \frac{\left\langle \left| \widetilde{I}'(\mathbf{k}) \right|^2 \right\rangle}{\frac{1}{(2\pi)^3} \int \mathrm{d}^3\mathbf{k}' \left| \widetilde{A}(\mathbf{k}-\mathbf{k}') \right|^2} \tag{2.13}$$

The integral of the primary beam is determined by the bandwidth and field of view of the instrument (See e.g. [58, 59, 12]). We can write down a back-of-the-envelope approximation by using Parseval's theorem.

$$\frac{1}{(2\pi)^3} \int \mathrm{d}^3\mathbf{k}' \left| \widetilde{A}(\mathbf{k}-\mathbf{k}') \right|^2 = \int \mathrm{d}^3\mathbf{r}' \left| A(\mathbf{r}-\mathbf{r}') \right|^2 \approx D_M(z)^2 \Omega \Delta D \tag{2.14}$$

Here Ω is the solid angle of the observation, and ΔD is the extent of the observation in the line-of-sight direction. This integral has units of Mpc^3, thus our units for the power spectrum in Equation 2.13 are mK^2 Mpc^3. For a detailed analytic calculation of this integral for a realistic model of the MWA beam, see Appendix B.

2.3 The EoR Window

Because we are making a measurement with a physical instrument, we are limited to certain regions of our spaces that we can access. Here we walk through several of those limitations, and motivate the concept of a so-called "EoR window" which will emerge as the region of k space where we can perform our measurement.

First, the $k_\perp$ modes which can be accessed are limited by the baselines of the interferometer. As seen in Equation 2.11, the baselines directly probe specific $k_\perp$ modes given by the antenna pair separation. Therefore, the minimum and maximum baselines impose a limitation on the minimum and maximum $k_\perp$ modes that can be accessed – represented by the vertical exclusion regions in Figure 2.2.

Next we turn to the $k_{||}$ dimension. The upper limit here is given by the frequency resolution of the measurement. The frequency resolution gives the smallest scales accessible

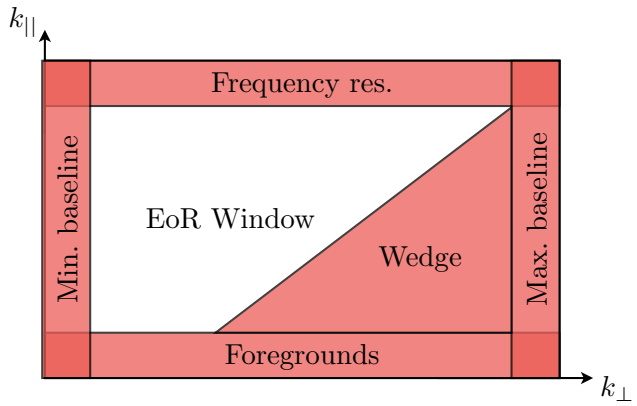

Figure 2.2: A schematic of the EoR window. The minimum and maximum $k_{\perp}$ modes accessible are determined by the baseline distribution of the array. The maximum $k_{||}$ mode is set by the frequency resolution of the observation. Bright spectrally smooth foregrounds contaminate the lowest $k_{||}$, as well as a characteristic “wedge" shape due to the instrumental mode mixing described in Section 2.4. These exclusions leave a region, dubbed the “EoR window" where an EoR measurement may be achievable.

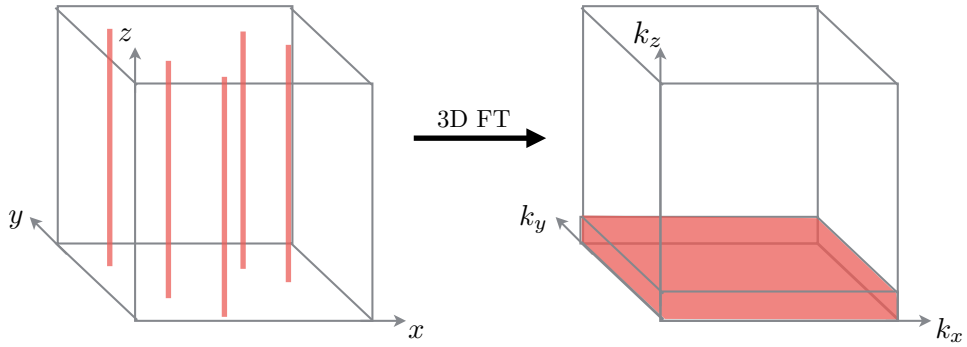

Figure 2.3: Spectrally smooth foregrounds mapped to k-space. The left shows an image space representation of point sources, where the smooth spectral structure is interpreted as long columns in the line of sight direction. When Fourier transformed, the δ functions in the perpendicular dimension become extended in $k_\perp$ and the columns in z are restricted to low $k_{||}$ modes.

in the line of sight, or equivalently limits the highest $k_{||}$ mode, given by

$$k_{||,\max} = \frac{\pi}{\Delta z} \approx \frac{\pi H_0 \nu_{21} E(z)}{c(1+z)^2 \Delta\nu}. \tag{2.15}$$

Likewise, the bandwidth of the observation, B, gives the resolution of $k_{||}$.

$$\Delta k_{||} = \frac{2\pi}{\Delta D} \approx \frac{2\pi H_0 \nu_{21} E(z)}{c(1+z)^2 B} \tag{2.16}$$

However, the lower end of accessible $k_{||}$ modes is limited by foreground contamination. The foregrounds present in a 21cm EoR measurement are extragalactic radio sources and synchrotron emission from our own Milky Way galaxy. These foregrounds are many orders of magnitude brighter than the expected signal, and fill every pixel of a continuum map. However, because these sources are spectrally smooth, they are limited to only a few low $k_{||}$ modes – illustrated in Figure 2.3.

2.4 The Wedge

The final exclusion region of Figure 2.2 is a topic of ongoing investigation and warrants a full discussion. The fact that baselines probe specific $k_\perp$ modes which are proportional to their baseline length measured in wavelengths leads to instrumental mode mixing. This mixing throws the smooth spectrum foreground power into higher $k_{||}$, producing a characteristic "wedge" shape in the 2D $(k_\perp, k_{||})$ space. The effect was seen in early simulations of realistic instruments [21], and has subsequently been a hot topic of discussion in the literature [33, 49, 50, 61, 70, 74, 89, 91, 92].

The wedge can be understood by considering a simple flat spectrum point source in the sky, defined by a delta function intensity.

$$I_{\rm pt}(\boldsymbol{\theta}, \nu) = I_0 \delta^2(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \tag{2.17}$$

The coherence function for such a sky is simply a frequency-independent plane wave, or "fringe", in uv space, where the frequency of the plane wave is proportional to θ_0.

$$\widetilde{I}_{\rm pt}(\mathbf{u}, \nu) = I_0 e^{-2\pi i \mathbf{u}\cdot\boldsymbol{\theta}_0} \tag{2.18}$$

Although this point source has no inherent frequency dependence, the baselines themselves migrate in frequency (see Figure 2.4) introducing frequency structure to the measurement.

$$V_{ij,{\rm pt}}(\nu) = \int {\rm d}^2\mathbf{u} \widetilde{A}_{ij}(\mathbf{u} - \nu \mathbf{b}_{ij}/c, \nu) I_0 e^{-2\pi i \mathbf{u}\cdot\boldsymbol{\theta}_0} \tag{2.19}$$

In the limiting case of a δ-function primary beam, this simplifies to

$$V_{ij,{\rm pt}}(\nu) \to I_0 e^{-2\pi i \nu \mathbf{b}_{ij}\cdot\boldsymbol{\theta}_0/c} \tag{2.20}$$

Using the frequency-space wavelength of this oscillating function, along with the coordinate mapping of Equation 2.8, we can write down the specific $k_{||}$ mode contaminated by this particular point source and measured by this baseline.

$$k_{||,{\rm pt}} \approx \frac{2\pi H_0 \nu_{21} E(z) \mathbf{b}_{ij}\cdot\boldsymbol{\theta}_0}{c^2(1+z)^2} \approx \frac{H_0 E(z) \mathbf{k}_\perp \cdot \boldsymbol{\theta}_0}{c(1+z)} \tag{2.21}$$

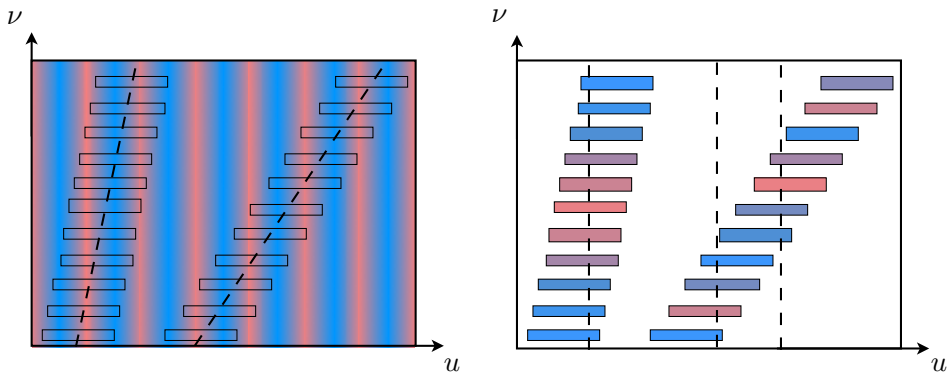

Figure 2.4: A schematic showing the mode mixing that manifests as the “wedge" feature shown in Figure 2.2. On the left the red/blue oscillations represent the frequency-independent fringes from a single flat spectrum point source described by Equation 2.18. The diagonal dashed lines show two examples of the frequency dependence of the u value probed by given baselines. The more separated the antenna pair (larger u), the more quickly the baseline migrates. The boxes show the integration region for an example primary beam. On the right we show the gridding step where the phases of the integrated visibilities are shown by colored boxes. When Fourier transforming to $k_{||}$ space, the transform is taken along the vertical dashed lines. From this picture we can see that the instrument has introduced spectral structure where none existed in the true sky frame. The larger baseline shows faster frequency oscillation due to its quicker migration, and thus mixes the foreground to higher $k_{||}$ modes.

In the second equality above we approximated the baseline as probing a single $\mathbf{k}_\perp$ mode for illustrative purposes.

Under this approximation we can see how the characteristic wedge shape manifests in $(k_\perp, k_{||})$ space. The contaminated mode is proportional to the baseline length, or $k_\perp$ – creating a line with slope given by the cosmological factors, and a maximum θ_0. What the maximum θ_0 exactly is determines the height of the wedge and different choices are used for projected sensitivity calculations. In principle an absolute maximum is set by the horizon ($\theta_{\rm max} = \pi/2$), but a more optimistic perspective may lead one to choose the field of view of the instrument, with the assumption that emission from beyond the first null of the primary beam does not significantly couple into the instrument. Indeed a recent study by Thyagarajan et al. [88] demonstrated that for the case of the MWA, near-horizon emission from our galaxy produces very strong wedge structure, while an instrument like HERA better rejects the horizon emission due to its smaller field of view.

In reality the edge of the wedge is not as sharp of a cut off as has been described so far. Real sources do contain some frequency structure, and the convolution of the primary beam also softens the edge. This is why utilizing the EoR window to employ foreground avoidance, while necessary, is not a sufficient strategy, and foreground subtraction is also needed. By modeling foregrounds with a realistic instrument simulation that incorporates the wedge mode mixing, we can lower the overall foreground power, and hope to achieve an EoR window clean enough for a cosmological measurement.

Chapter 3

A NEW LAYOUT OPTIMIZATION TECHNIQUE FOR INTERFEROMETRIC ARRAYS, APPLIED TO THE MWA

"In your language you have a form of poetry called the sonnet... There are fourteen lines, I believe, all in iambic pentameter. That's a very strict rhythm or meter... And each line has to end with a rigid pattern. And if the poet does not do it exactly this way, it is not a sonnet... But within this strict form the poet has complete freedom to say whatever he wants... You're given the form, but you have to write the sonnet yourself. What you say is completely up to you."

– Madeleine L'Engle, *A Wrinkle in Time*

Faced with the freedom to place MWA antennas anywhere, within certain rules and obligations, we chose a scientific approach to filling in the lines. The following chapter is based on Beardsley et al. [3], with minor changes for coherency within this dissertation and updated information about the MWA.

3.1 Layout Motivation

Antenna placement is a critical design criterion for any interferometric array as it determines the baseline distribution and thus the angular distribution of the point spread function of the radio telescope. Nearly all observatory sites have areas where antennas cannot be placed. Buildings, roads, runways, power and data access, land use and ownership issues, endangered flora and fauna, flood zones, elevation, and ground stability are but a few of the common issues that constrain the placement of antennas. Even in remote desert locations a flat and barren terrain can quickly become dotted with exclusion zones where antennas cannot be placed.

This is of particular concern for 21 cm Cosmology telescopes targeting the Epoch of Reionization (EoR) and Baryon Acoustic Oscillation (BAO) dark energy measurements, as the quality of the monochromatic PSF is directly related to the ability to subtract foreground contamination [92, 60, 21, 48, 7]. Antenna exclusion zones can introduce asymmetries in the baseline distribution which limit the angular dynamic range and thus achievable level of foreground subtraction (see Morales and Wyithe [57] for a recent review of foreground subtraction for 21 cm Cosmology).

There is a long history of array configuration studies, including optimization of arrays with cost constraints [e.g. 18], simulated annealing for small N arrays [19], optimization to reduce the peak sidelobe levels [39, 40], or optimization to match a particular baseline distribution with and without ground constraints [42, 9]. Our particular concern is situations in which some areas cannot be used (exclusion zones), a particular radial baseline distribution must be met, and a very high angular dynamic range must be achieved. While our problem is similar to that of Boone [9], we find that the figure of merit used in that work does not sufficiently capture large scale structure in the baseline distribution. We develop an alternative figure of merit, which naturally leads to a new optimization method.

In Section 3.2 we explore the effect of exclusion zones on the baseline distribution, develop a new spatially sensitive figure of merit, and present our new optimization method. We then apply our method to placing the MWA antennas in Section 3.3, and present a 512 antenna layout for the MWA. Shortly after this work was completed the MWA was rescoped as a 128 antenna instrument. The MWA that was actually built consists of a 128 antenna subset of the layout presented here, where our algorithm was adapted to select a subset of antennas which optimized our figure of merit.

3.2 Array Layout Comparisons and a New Technique

Proposed and future large N radio arrays will face the challenge of placing hundreds to thousands of antennas to optimize scientific goals, while obeying numerous constraints. While most physical constraints exist on the antenna locations (areas of exclusion on the ground),

science capabilities are optimized in the *uv* plane for an interferometric measurement, and hence these arrays should match the ideal *baseline* distribution as closely as possible. This makes the problem very non-linear because any one antenna placement affects $N-1$ baselines, and it is not immediately obvious how a constraint such as an exclusion area will affect the baseline distribution.

In our analysis, we explored three array layout methods. The first method is random and with no exclusion areas ("random unmasked"), in which antennas are placed randomly with a weighted radial distribution. Algorithmically, for each antenna a radius is first drawn from a distribution that matches the desired radial density profile, then azimuthal angles are chosen at random until one is found that does not overlap with previously placed antennas. The second array generation method is also random but incorporates exclusion areas ("random masked"). This method is identical to the random unmasked method, with the addition of avoiding exclusion areas by the use of a mask that is checked in the same step as checking for overlap with previous antennas. The third and final method is the algorithm that we developed ("active method"), that actively minimizes spatial structure in the baseline distribution and is detailed later in this section.

We assume the scientifically desired *uv* or antenna distribution of the array is known. For our examples we use the specifications for the original MWA telescope which has a dense core optimized for EoR measurements, with a smooth radial dependence. More details of the MWA distribution are given in Section 3.3, but for a full motivation, see Bowman et al. [12]. While we use this example for our work, our methods are generalizable to any large N array.

Figure 3.1 shows three examples of baseline distributions generated by the random unmasked, random masked, and active masked methods respectively. The left pane shows the baseline distribution on a logarithmic scale, while the right pane shows the difference from the ideal smooth analytic function to accentuate undesired structures in the *uv* distributions. These three examples are representative of the over four thousand array layouts we have hand graded to arrive at our conclusions.

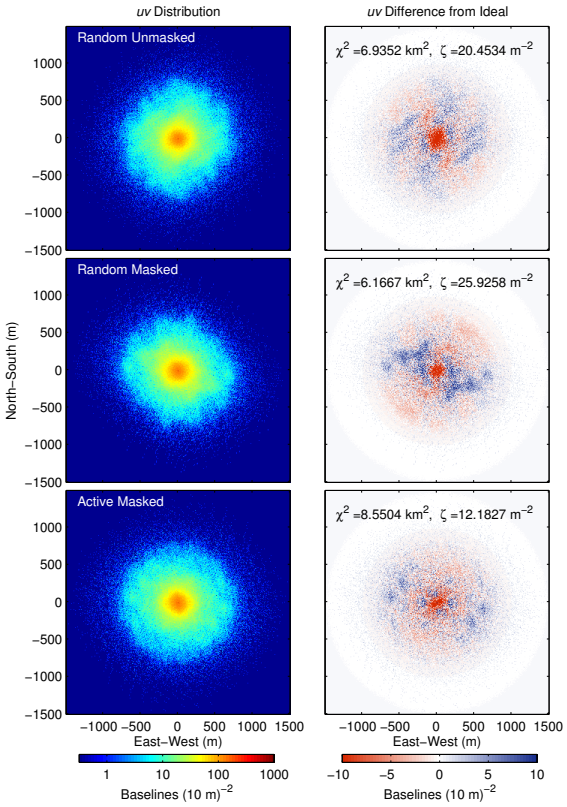

Figure 3.1: Example *uv* distributions for array layouts generated by (top to bottom) the random unmasked, random masked, and active masked methods. The left panes show the snapshot single frequency baseline distributions on a log scale, while the right panes show the difference of this distribution from the smooth analytic ideal. The small scale fuzzy noise is equally present in all array realizations and is due to the finite number of antennas. However, the large scale structure varies greatly from array to array. The exclusion areas have introduced significant asymmetries in the baseline distribution of the random masked method (middle row). The active masked method (bottom row) is able to highly suppress this structure, even beyond the level of the *unconstrained* random method (top row). Furthermore, we see the χ^2 values for these sample arrays show no correlation with the quality of the arrays, however the figure of merit, ζ, accurately reflects the amount of azimuthal structure in the distributions (see Figure 3.2).

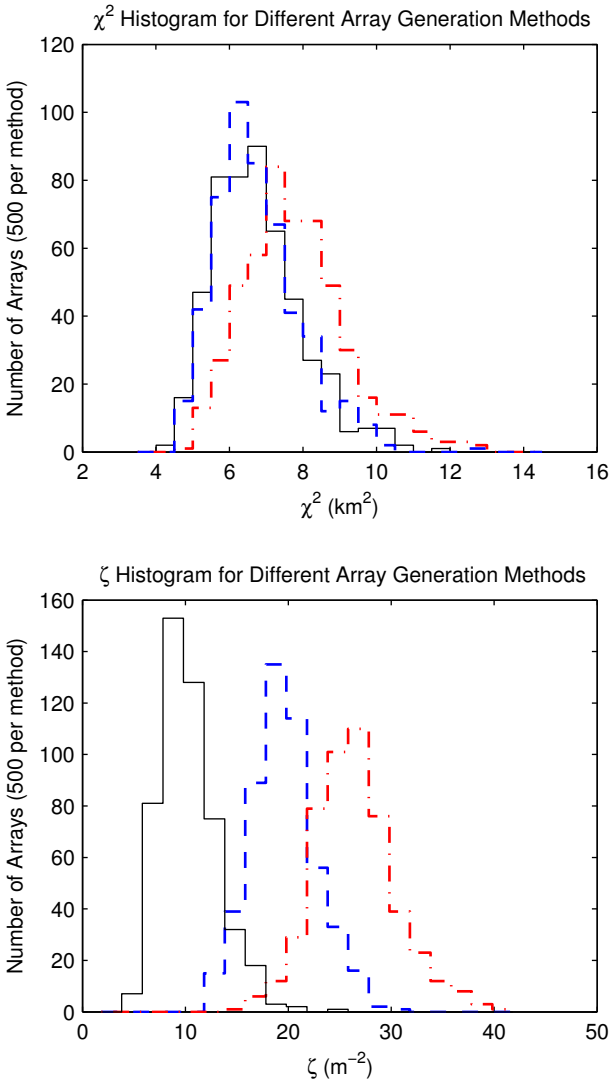

Figure 3.2: χ^2 and ζ histograms for random unmasked, random masked, and active masked methods as denoted by the thick dashed blue, thick dot-dashed red, and thin solid black lines respectively. While the χ^2 values do not distinguish the quality of the different realizations (a conclusion firmly supported by our hand grading), the ζ values strongly separate the realizations based on asymmetry.

All of these images have fuzzy small scale noise due to the finite number of antennas. However, there is another more insidious artifact present in the masked baseline distribution (middle row of Figure 3.1) – large scale structure imprinted by the antenna exclusion areas. In all of the masked random array realizations there are significant regions of over and under densities in the baseline distribution which translate directly into unwanted PSF features.

To understand the effect of baseline over and under densities, consider a nearly perfect *uv* distribution with a small region of excess baselines. This region of *uv* over density can be viewed as a 'wave packet' of baselines at similar spatial frequencies. In the wave packet picture there is a fundamental corrugation in the PSF given by the location of the center of the excess region. However, the nearby modes in the packet beat in and out of phase with the fundamental corrugation. When the wave numbers are all in phase the amplitude is very high—the number of excess baselines in the region—but they quickly dephase only to rephase again some distance further along in the PSF. The undesirable "features" seen in most PSFs are the periodic signature of a wavepacket beating across the PSF. Over (or under) dense regions that cover a large portion of the *uv* plane will quickly damp down (wide bandwidth), though they often have a lot of power due to the large number of baselines involved, and correspond to large close-in sidelobes. Smaller features in the *uv* plane damp more slowly and repeat many times across the PSF leading to the small far sidelobes. A smooth *uv* distribution necessarily leads to a smooth PSF, and the PSF sidelobe structure is dominated not by the unavoidable fuzzy noise but instead the larger regions of over and under density in the *uv* plane.

Our first approach to quantify the deviations from the desired *uv* distributions was to consider χ^2. This was calculated by gridding the *uv* distribution and integrating the square of the difference from the ideal, weighted by the variance in each pixel from 500 random unmasked realizations. However, χ^2 is not a spatially aware function—any deviation from the ideal is weighted the same regardless of where in the *uv* plane the deviation occurs. Because of this lack of spatial information, χ^2 does not capture the large scale structure that is important to choosing an array. The examples in Figure 3.1 vary quite a bit in quality, however, the

associated χ^2 values do not reliably reflect the degree of spatial structure. The insensitivity of χ^2 to array quality is demonstrated again in the histogram in Figure 3.2(a). Despite a clear qualitative difference between the masked and unmasked random configurations (dashed blue and dash-dot red) the distributions of χ^2 are very similar.

With the spatial dependence in mind, our next step was to develop a figure of merit based on a Bessel decomposition. The residual uv distribution ($D(r,\phi)$, difference between actual and desired) can be decomposed into Bessel modes

$$D(r,\phi) = \sum_{n=1}^{\infty}\sum_{m=0}^{\infty} J_m\left(\frac{x_{mn}r}{R}\right)\left(A_{mn}\sin(m\phi) + B_{mn}\cos(m\phi)\right), \tag{3.1}$$

where x_{mn} is the n^{th} zero of the m^{th} Bessel function, and R is the maximum allowed baseline length. Using the orthogonality of the Bessel functions, we can determine A_{mn} and B_{mn} [e.g. 36].

$$A_{mn} = \frac{2}{\pi R^2 J_{m+1}^2(x_{mn})}\int_0^{2\pi} d\phi \int_0^R dr\ rD(r,\phi)J_m(\frac{x_{mn}r}{R})\sin(m\phi) \tag{3.2a}$$

$$B_{mn} = \frac{2}{\pi R^2 J_{m+1}^2(x_{mn})}\int_0^{2\pi} d\phi \int_0^R dr\ rD(r,\phi)J_m(\frac{x_{mn}r}{R})\cos(m\phi) \tag{3.2b}$$

The amplitudes of the asymmetric Bessel coefficients ($A_{mn},\ B_{mn},\ m > 0$) reflect the asymmetric spatial over and under densities in the uv plane. We then define a figure of merit ζ as the sum of these Bessel coefficients

$$\zeta \equiv \sum_{n=1}^{\max(n)}\sum_{m=1}^{\max(m)} \sqrt{A_{mn}^2 + B_{mn}^2}\ . \tag{3.3}$$

The double sum is not infinite for two reasons: computational time, and higher modes describe smaller scale structure. Because we are interested in suppressing large scale structure in the uv, we can truncate the sum. For our work, we found that values of $\max(n) = 10$ and $\max(m) = 20$ were computationally feasible and provided sufficient information to suppress the large scale structure.

Smaller ζ corresponds to less spatial structure, and hence a more desirable layout. The right hand panel of Figure 3.2 shows the ζ distributions for the same arrays, clearly separating

the masked and unmasked random arrays, and in Figure 3.1 the ζ values accurately track the quality of the array realizations.

To minimize our figure of merit, we created an active masked algorithm based on the ζ figure of merit. For computational reasons we first place a subset of the antennas (350 of 496 for our example) using the random masked method. Then for each remaining antenna, we first choose a weighted random radius, r, and many candidate azimuthal locations (angular spacing of 10 m in our example). We then select the location with the smallest ζ value, and repeat until all N antennas are placed.

The result is clear in Figures 3.1 & 3.2. Despite having the additional constraint of the exclusion areas, the active masked method produces more symmetric baseline distributions than either the random masked (expected) or random unmasked methods (unexpected). We can see this qualitatively by comparing the baseline distributions in the three examples. In the thousands of arrays we examined by hand we observed a very strong correlation between small ζ and spatial symmetry. The success of our algorithm is shown statistically by the distribution of ζ values for the three methods in Figure 3.2(b).

It is of interest to note that significant baseline asymmetry arises even in the unmasked random array realizations (Figure 3.1 top row, no exclusion areas). These asymmetries are due to shot noise in the random antenna placement. Conceptually, as the last few antennas are added images of the entire array are added to the uv plane at that distance from the center. For centrally condensed arrays this can produce lumps in the uv plane. Alternatively, one can consider moving a single antenna on the ground which coherently changes $N - 1$ baselines. Small random associations can thus make significant correlations in the baseline distribution.

The new active method based on the Bessel decomposition figure of merit produces arrays which are superior to even an unconstrained random algorithm, even in the face of significant exclusion areas.

3.3 Original MWA Layout

We have used our new algorithm to determine the originally planned 512 antenna layout of the MWA. The MWA currently consists of 128 antennas, with the capacity to easily expand to 256 antennas given modest additional funding [90]. The original concept for the MWA was 512 antennas [51], though expansion to 512 antennas will require significant additional investment in MWA infrastructure. In this work we assess the originally envisaged 512 antenna concept for the MWA, and the current MWA consists of a 128 antenna subset of the array presented here. The currently built MWA is described in Tingay et al. [90], and its cosmological sensitivity is discussed in Chapter 4.

In our work the majority of the 512 antennas (496) will be distributed over a 1.5 km diameter core, with the remaining 16 antennas at a $\sim$3 km diameter to provide higher angular resolution for solar measurements. The 16 "outliers" are placed by hand, while we implement our algorithm for the 496 core antennas. The antenna density distribution will be constant within a central 50m radius, and have a r^{-2} dependence beyond (see Lonsdale et al. [51] for a full description of the original concept for the MWA instrument and Tingay et al. [90] for a description of the built MWA). The smooth ideal uv distribution is an auto-convolution of this antenna density distribution, and is shown in Figure 3.3.

Several parameters of the algorithm were explored to further improve the quality of the arrays being generated. For example, we varied the number of random antennas placed before initializing the active phase of our algorithm. We found that placing 350 random antennas (according to the ideal antenna distribution described above) provided sufficiently unconstrained initial conditions to proceed with the active phase. Running the algorithm in this mode 500 times provided a good sampling of the phase space.

We also investigated several array center locations within a few hundred meters of the nominal array location. Due to the irregular distribution of avoidance areas on the ground, choosing different centers did have an effect on the quality of the best arrays generated by our algorithm. In particular, a center near a high concentration of vegetation or rock outcrops

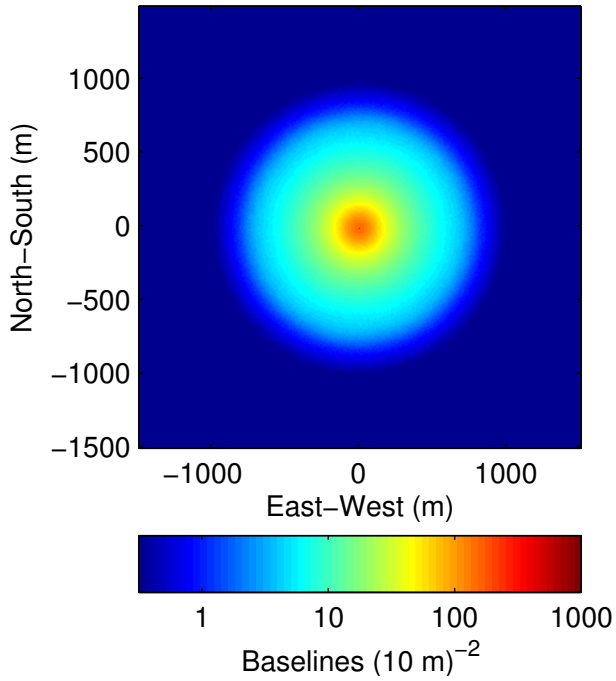

Figure 3.3: Ideal MWA *uv* distribution. This distribution was used in the active masked method for generating the original MWA layout. Compare to top left pane of Figure 3.5

usually results in a deficit of short baselines. We used this information, along with feedback from a ground truth survey in February 2011 to determine our final array center.

After generating 500 candidate arrays for each potential location, we used ζ as a guideline for selecting the highest quality array layouts, backed by hand grading. The result is the final location and layout of the MWA. Figure 3.6 shows an illustration the final array overlaid on an aerial photo of the site. The locations of 496 core antennas, along with the 16 hand placed outliers, are available in the electronic supplement to Beardsley et al. [3]. Figures 3.4 & 3.5 show the corresponding uv distribution and PSF. There is essentially no asymmetry in the final array—all of the large scale structure is greatly suppressed, providing a very smooth uv sampling. The small residual ripples in the PSF are on the order of the background noise we expect due to the finite number of antennas. Following the discussion of Morales [58] and using the parameters from Bowman et al. [12], the thermal noise uncertainty on the EoR power spectrum can be calculated in the bottom panel of Figure 3.4. With the exception of the small deviation at very low cosmological wavenumber, k, the thermal noise for our proposed layout very nearly traces that of the ideal array.

3.4 Layout Optimization Conclusions

While we have been motivated by the need to explore array configurations for the MWA with exquisite smoothness in the PSF despite significant exclusion areas, we hope our method will be useful for determining the antenna layouts of other large N arrays. In particular we have developed a new figure of merit based on Bessel decomposition that is sensitive to large scale over and under densities in the uv plane. We have shown that algorithms based on this figure of merit can achieve extremely smooth baseline distributions while avoiding areas where antennas cannot be placed.

We have used this new algorithm to generate a 512 antenna layout of the MWA. The currently built stage of the MWA consists of a 128 antenna subsample of the array configuration shown in this paper.

Our figure of merit is similar, but different from the traditional approach of minimizing

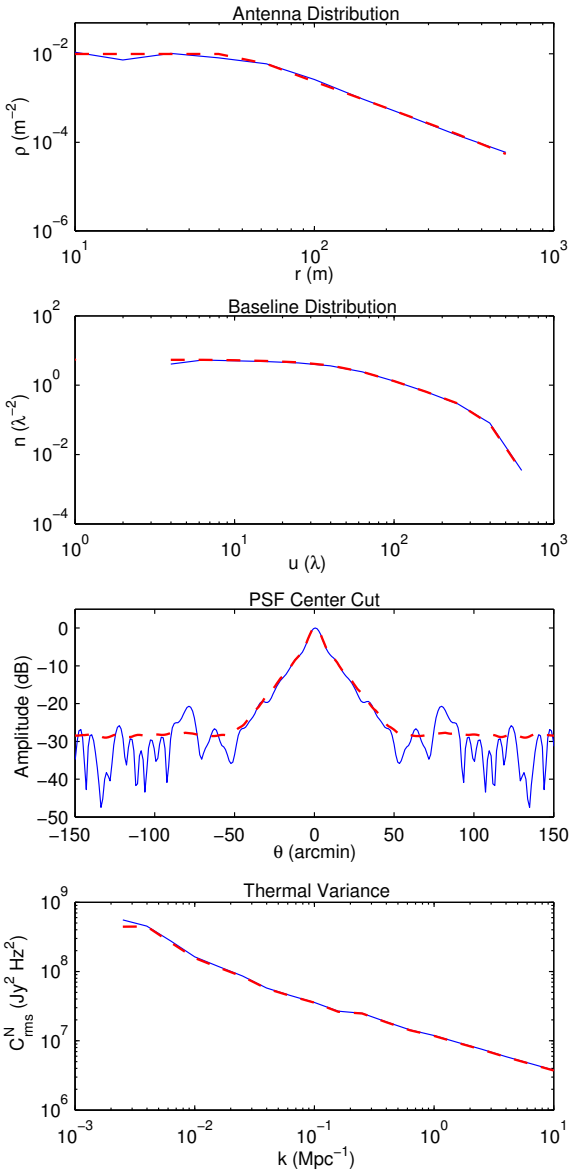

Figure 3.4: In each figure the original MWA array layout (thin solid blue) is compared to the analytic ideal baseline distribution (thick dashed red) at 150 MHz observing frequency. Top to bottom these are the radial antenna distribution, the radial baseline distribution, a cut through the snapshot single frequency PSF, and the thermal noise as a function of cosmological wavenumber k for a fiducial EoR measurement (following Bowman et al. [12]). In all aspects the original array very nearly traces the ideal array.

the maximum sidelobe. Both figures will result in low unwanted structure in the point spread function. While minimizing the maximum sidelobe focuses on a hard constraint of the maximum peak, ζ captures a more holistic picture of the distribution with a softer constraint. Further work would be necessary to more precisely characterize the differences between the two.

Further work could also explore the possibility of implementing genetic algorithms or stimulated annealing with our figure of merit, ζ. For our present work, however, a trial-and-error method sufficed, and avoided questions of convergence while sampling a sufficient area of phase space.

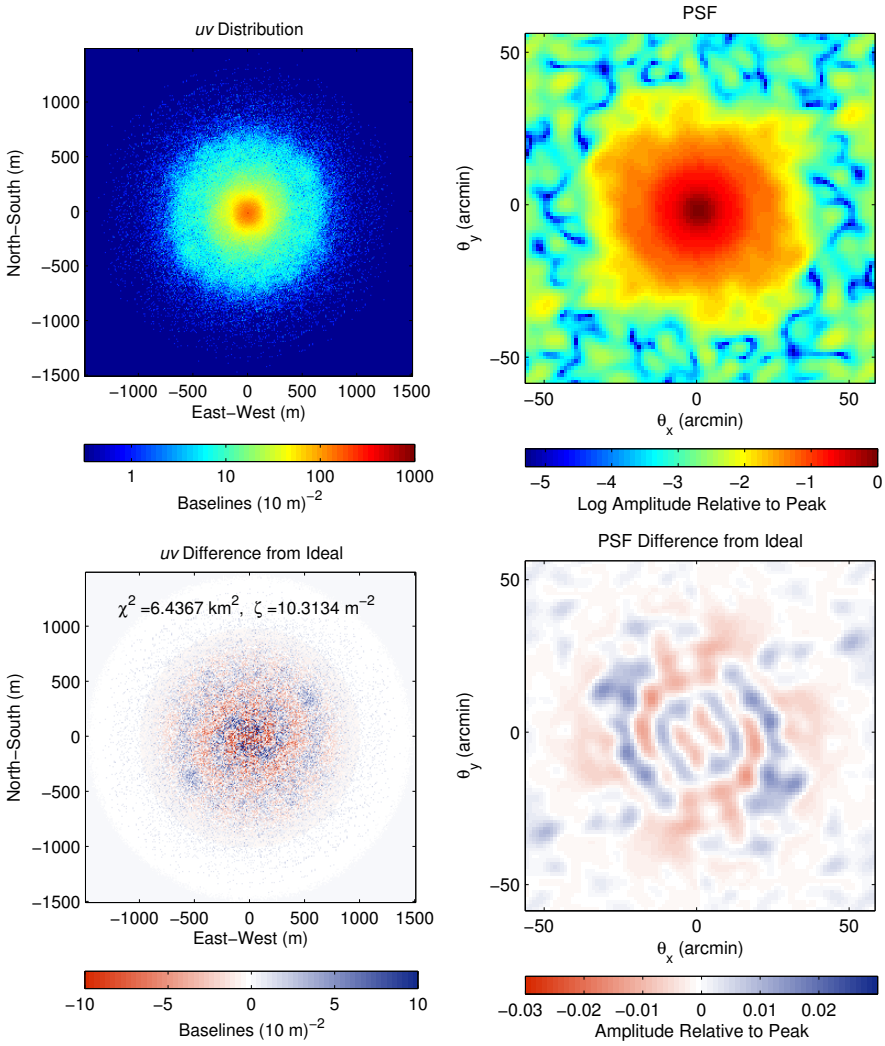

Figure 3.5: Baseline distribution and point spread function of original MWA array layout. The baseline distribution is shown in the left column, with the difference from ideal on the bottom. The azimuthal structure is nearly completely suppressed, and only small scale noise remains. The snapshot PSF for 150 MHz at zenith is shown in the right column. The sensitivity relative to the peak is shown on the top, while the difference from ideal is on the bottom. The residual ripples in the PSF difference are $\approx 1\%$ of the peak, which is on order with the background ripples expected due to our finite number of antennas.

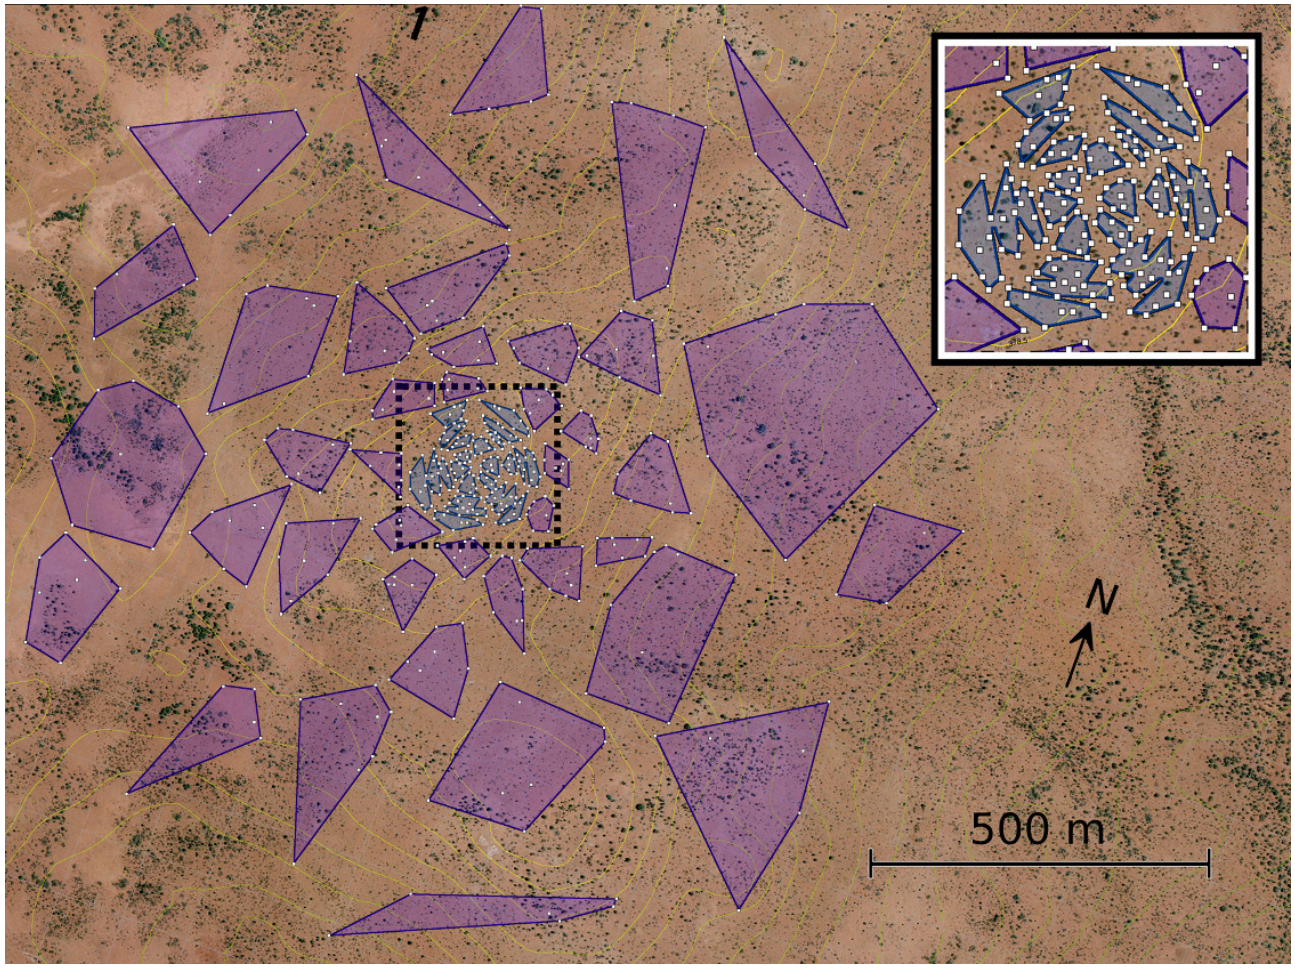

Figure 3.6: An aerial photo of the MWA site with the original 496 antenna core array superimposed. White squares represent antennas (to scale). The colored polygons depict a possible receiver scheme. Each polygon outlines a receiver set's electrical footprint (8 antennas per polygon except for a few outer receiver sets that will service some of the 16 outlier antennas not shown here). Inset: An enlarged view of the center of the array.

Chapter 4

THE EOR SENSITIVITY OF THE MURCHISON WIDEFIELD ARRAY

"Experience is the only thing that brings knowledge, and the longer you are on earth the more experience you are sure to get."

– L. Frank Baum, *The Wonderful Wizard of Oz*

In the case of EoR experiments, time brings sensitivity, and the longer we integrate the more sensitivity we get. It is useful to know how long the MWA will need to collect data in order to have the sensitivity needed to detect the EoR. This chapter is based on Beardsley et al. [4], with minor changes for coherency within this dissertation and updated information about the MWA. Since publication, this calculation has become the reference sensitivity for the MWA collaboration and has been used to inform observing strategies and potential future build-outs of the array.

4.1 Sensitivity Background

The MWA has been built in the radio quiet Murchison Radio Observatory in Western Australia, aiming to measure the EoR power spectrum via the 21 cm signal over a large range of redshifts. The originally planned MWA was to consist of 512 antennas, distributed over a circular region of radius 1.5 km [51]. With available funding the instrument was re-scoped to 128 antennas, but will have similar layout characteristics to the originally planned 512 antenna array. A full description of the 128 antenna instrument is presented in Tingay et al. [90], and a thorough description of the science capabilities will be presented in Bowman et al. [14].

Here we calculate the MWA's expected sensitivity to the EoR signal using the physical

antenna locations. A given baseline (the separation vector between any two antennas) is sensitive to a particular angular Fourier mode on the sky, so the baseline distribution is directly related to the EoR sensitivity of an array [58]. The MWA baseline distribution has a dense core for EoR sensitivity and a smooth extended radial profile for calibration and foreground subtraction purposes [12]. The locations of the 128 antennas for the MWA were optimized using the algorithm presented in Beardsley et al. [3] and shown in Figure 4.1. A table of the locations of all 128 antennas is available with our originally published article [4].

We use a fiducial model to calculate the MWA's sensitivity and provide the information needed to quickly apply any model in tables attached to Beardsley et al. [4]. The EoR observing plan for the MWA is to track fields when they are above 45 degrees elevation and the sun and galactic center are below the horizon. Over an annual cycle, this yields a full observational season of 900 hours integration on a primary field (dubbed "EoR1") and 700 hours on a second field ("EoR0"). For the fiducial model, we find that with a full season of observation the MWA will be capable of a 14σ power spectrum detection, along with constraints on the slope.

Throughout this chapter we use a ΛCDM cosmology with $\Omega_m = 0.73$, $\Omega_\Lambda = 0.27$, and $h = 0.7$, consistent with WMAP seven year results [41]. All distances and wavenumbers are in comoving coordinates.

4.2 EoR Sensitivity

The power spectrum measurement of the sky temperature is done in three dimensions (two angular directions, and the line-of-sight direction achieved through redshift), so we must find the uncertainty in each three dimensional voxel in cosmological wavenumber ($\mathbf{k}$) space, then perform a weighted average in spherical bins to arrive at one-dimensional sensitivity (following Morales [58], McQuinn et al. [53], and Morales and Wyithe [57]).

The fundamental visibility measurement of an interferometer is done in (u, v, ν) space, where u and v are the baseline coordinates (measured in wavelengths), and ν is the frequency

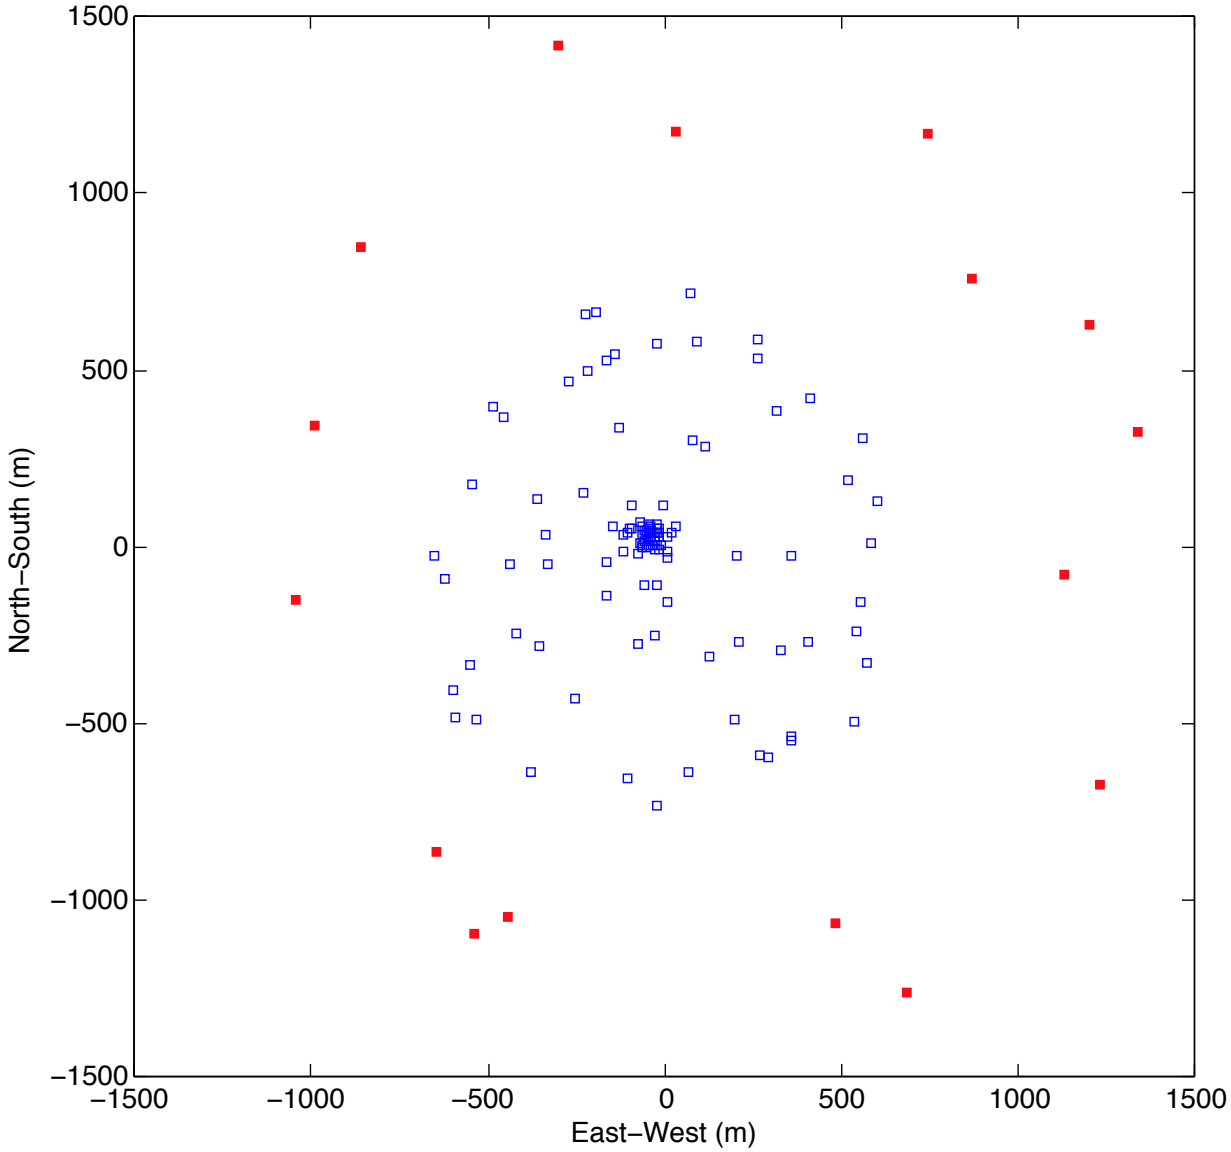

Figure 4.1: The antenna locations for the 128 antenna MWA. Positions are measured relative to -26° 42' 4.396" Latitude, 116° 40' 13.646" Longitude. The blue squares show the core 112 antennas which will be integrated for an EoR measurement. The solid red squares represent the outlier antennas used for solar measurements, but are not used for EoR measurements. While the antennas are indeed square, the squares shown here are not to scale.

Table 4.1: Observational parameters for sensitivity estimation

Parameters	Values
No. of antennas	112*
Central frequency	158 MHz (z $\sim$ 8)
Field of view	31°
Effective area per antenna	14.5 m^2
Total bandwidth	8 MHz
T_{sys}	440 K
Channel width	40 kHz
Latitude	-26.701°
Primary Field RA	6^h
Secondary Field RA	0^h

*Sixteen of the 128 antennas are not integrated for EoR measurements and are not included here.

of the observation. The thermal uncertainty on the visibility measurement is given by

$$V_{\rm rms}(u,v,\nu) = \frac{c^2 T_{\rm sys}}{\nu^2 A_e \sqrt{\Delta\nu\tau}}, \tag{4.1}$$

where $T_{\rm sys}$ is the system temperature, A_e is the effective collecting area per antenna, $\Delta\nu$ is the frequency channel width, and τ is the total integration time for the mode including redundant baselines [57]. Observational parameters for our calculation are shown in Table 4.1. The system temperature is dominated by galactic foreground emission, and redshift dependence is discussed in Bowman et al. [12]. Here we assume a constant system temperature over the observational bandwidth.

To determine τ, the integration time per (u,v,ν) voxel, we use the surveyed antenna locations, and perform an aperture rotation for 3 hours on either side of zenith. We approxi-

mate chromatic effects by calculating the baseline migration along the frequency dimension, then averaging. This includes chromatic effects while avoiding a full covariance calculation (Hazelton, et al. 2015, in preparation).

The sampling matrix for one day of observation on a single EoR field is shown in Figure 4.2. The MWA will have a very dense, highly redundant uv core, with a smooth radial profile extending to 1.5 km. The large number of baselines in the core will beat down the thermal variance for those modes because the effective observing time is the sum of all the baselines observing the mode.

As discussed in Chapter 2, the (u, v) coordinates map directly to the transverse cosmological wavenumber by Equation 2.11. The observing frequency dimension maps to the line-of-sight direction, and must be Fourier transformed to the $k_{||}$ dimension. Once these conversions are done, our data is in three dimensional k-space, and we square to reach the power spectrum. Propagating errors, the thermal uncertainty per k-space bin is given by

$$C^N(\mathbf{k}) = T_{\text{sys}}^2 \left(\frac{D_M^2 \lambda^2}{A_e} \right) \left(\frac{\Delta D}{B} \right) \frac{1}{\tau}. \tag{4.2}$$

The second term can be thought of as converting the uv bin size (A_e) to cosmological wavenumber space and has units of Mpc^2. The third term converts the width of the observation from bandwidth to line of sight spatial extent and has units of Mpc s (for flat space the line-of-sight and transverse distances are equivalent), and τ is the integration time for the k-space bin (in seconds). Inserting the values from Table B.1 for all terms except the integration time gives

$$C^N(\mathbf{k}) = \frac{6.95 \times 10^{13}}{\tau} \text{ mK}^2\text{Mpc}^3. \tag{4.3}$$

There is also a sample variance contribution to the uncertainty. Assuming the distribution is Gaussian, the sample variance per three dimensional voxel is given by the power spectrum itself [53]. Combining the thermal and sample uncertainties gives the total variance per 3D k-space voxel

$$\sigma_P^2(\mathbf{k}) = \left(P_{21}(\mathbf{k}) + C^N(\mathbf{k}) \right)^2. \tag{4.4}$$

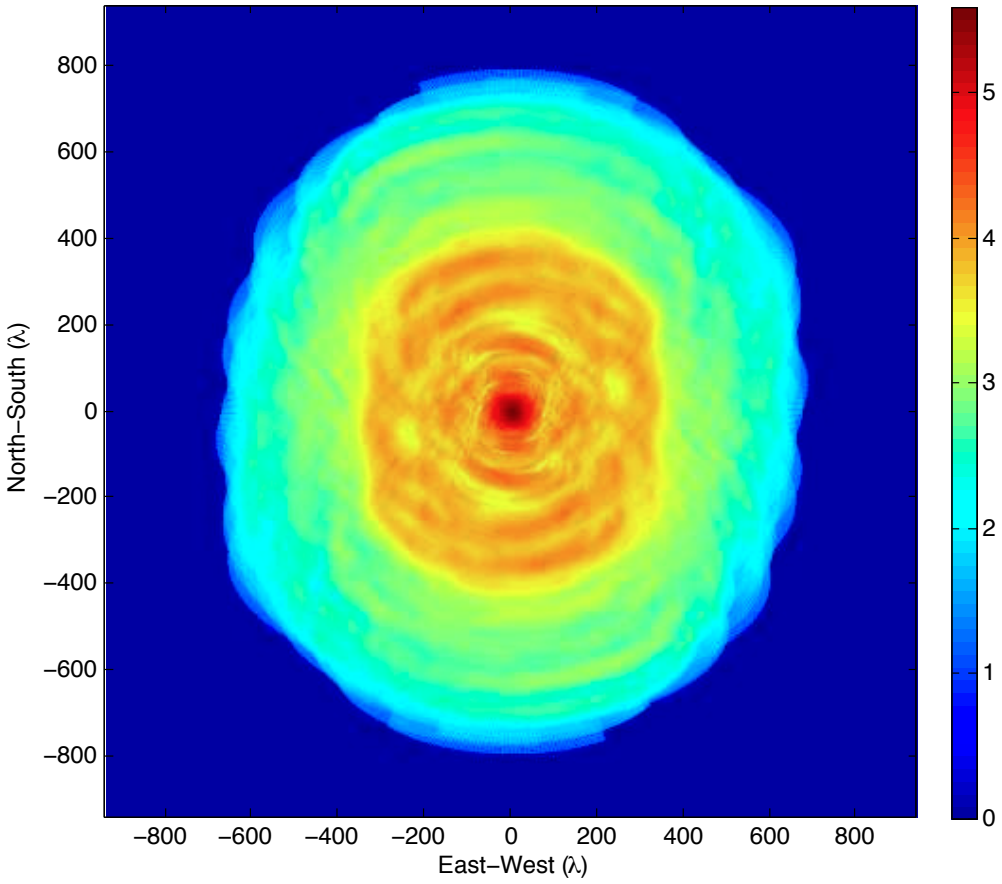

Figure 4.2: Effective integration time per day per (u, v) cell including rotation synthesis for 128 antenna MWA at $z \sim 8$ ($\lambda = 1.89$ m). The color scale units are the logarithm of effective seconds observed per day per angular mode, assuming six hours of integration per day on one EoR field. Note that the total number of seconds observed per day is 21600, but an angular mode can be effectively observed longer due to redundant baselines. The most observed mode in this array is $\sim 4 \times 10^5$ sec/day. The uv cell size is dictated by the size of the instrumental window function. Following Bowman et al. [12] we used a cell size of (8.3 m)^2. Data for this figure is available as an attachment to Beardsley et al. [4] in a machine readable table to easily plug into Equation 4.3.

Because of the sample variance term, the calculated sensitivity of an array depends on one's choice of theoretical EoR model. While surveying the landscape of EoR models is beyond the scope of this paper, we have included a table of the effective seconds observed per day per (u, v) cell in the electronic supplement to Beardsley et al. [4] (data for Figure 4.2). The seconds per day can be combined with the observing strategy to calculate the integration time per cell, τ in Equation 4.3, and combined with the theoretical model in Equation 4.4 to accurately determine the sensitivity of the MWA for any proposed model. The coefficient values in Equation 4.3 and the coordinates of the supplemental table can be scaled to different redshifts with $\sim 5\%$ error on the resulting sensitivity, or the antenna locations and synthesis rotation can be used to recalculate the integration time per bin using the supplemental table as a cross-check. In the remainder of this chapter we use the fiducial model of a fully neutral IGM [26][1] as an example of how to accurately calculate the EoR power spectrum sensitivity.

The underlying EoR fluctuations are assumed to be isotropic, however velocity distortions will amplify the signal in the line-of-sight direction on relevant large scales [2]. For our fiducial model this angular dependence is given by $P_{21}(\mathbf{k}) = (1 + 2\mu^2 + \mu^4)P_{21}(k)$, where $\mu = k_{||}/|\mathbf{k}|$. This effect depends on whether dark matter or ionizing sources are sourcing the fluctuations. Throughout reionization both sources will be relevant and the above expression will depend on the cross-power spectrum between the fluctuations. Because our fiducial model is a fully neutral IGM we can use this simplified relation. In addition the MWA is sensitive to much smaller $k_\perp$ modes compared to $k_{||}$ modes, so this effect is significant for the dark matter sourced fiducial model.

Figure 4.3 shows the signal to noise per voxel in a slice through the 3D k-space. At low k, a large signal and dense baseline distribution result in a signal to noise approaching 1. Moving up in $k_{||}$, the signal diminishes, but the baseline density remains constant, so the sensitivity drops relatively slowly. Moving up in $k_\perp$, however, both the signal and the baseline density drop, resulting in a more drastic drop in sensitivity.

---

[1] Available online at www.astro.ucla.edu/ sfurlane/21cm_pk.htm

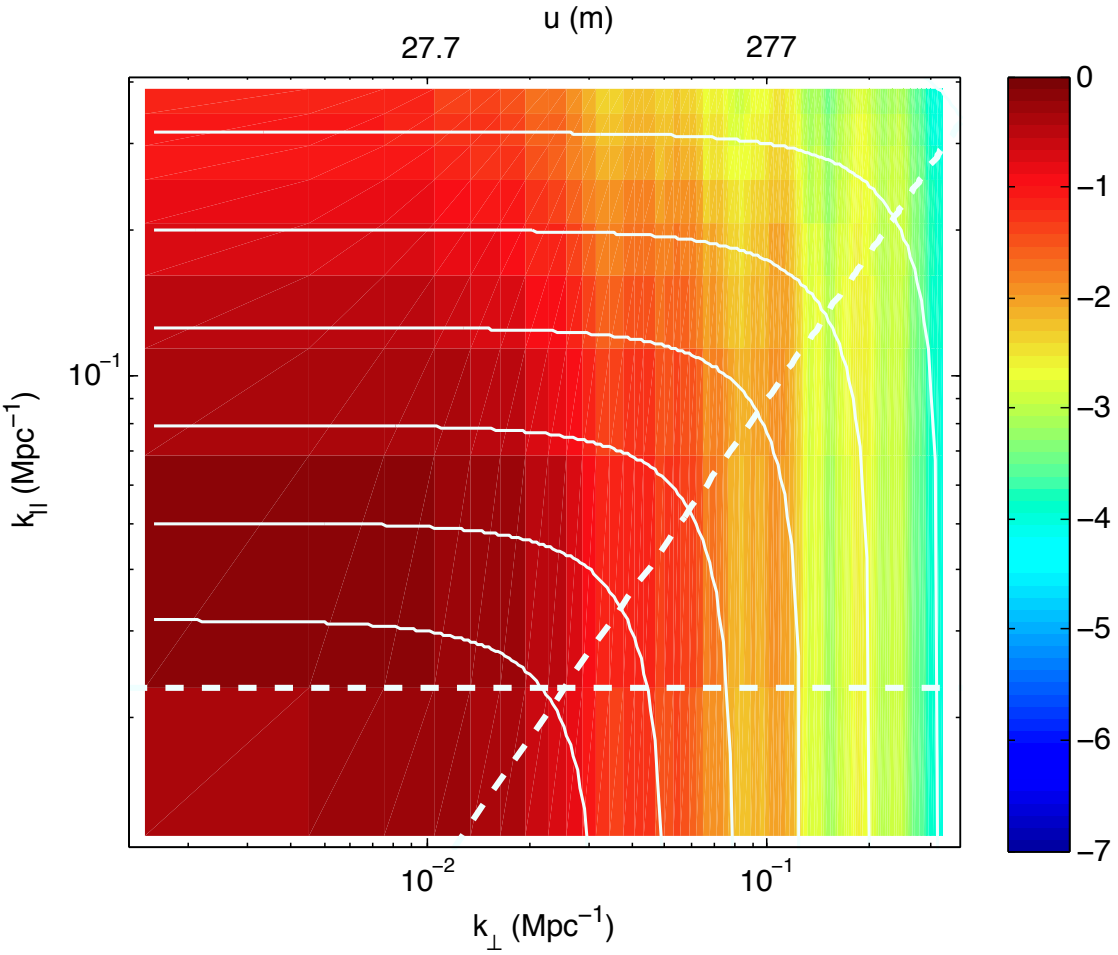

Figure 4.3: Estimated power spectrum sensitivity to EoR signal per voxel for the MWA. The quantity plotted is $\log_{10}(P_{21}(\mathbf{k})/(\sigma_P(\mathbf{k}))$ for a two dimensional slice of the three dimensional data cube with 900 hours of integration. The white curved lines show the bin edges used for the one dimensional plot (Figure 4.4). The data below the horizontal dashed line and to the right of the diagonal dashed line will be contaminated by foregrounds. Only data within the EoR window (the upper left) is used to calculate the sensitivity in Fig. 4.4. For reference, the corresponding baseline lengths are given on the top axis.

As discussed in Sections 2.3 & 2.4, foreground contamination limits the observability of the EoR. Fortunately, the contamination is localized in 3D k-space, leaving a relatively uncontaminated EoR window [92, 61]. The spectrally smooth foregrounds are fit to low order polynomials over the full 30.72 MHz instrument bandwidth, contaminating low line-of-sight wavenumbers [13]. However, an individual observation is limited to ~ 8 MHz due to cosmic evolution, so only our $k_{||} = 0$ bin will be contaminated. This exclusion zone is shown in Figure 4.3 by the region below the horizontal white line. In addition, mode mixing effects will throw power higher in $k_{||}$, creating a wedge shape of contamination [33, 49, 50, 61, 70, 74, 89, 91, 92]. The location of this contamination is indicated in Figure 4.3 by the region below the diagonal line. The 'EoR window' is to the left of the diagonal line and above the horizontal line. In this calculation we only use modes within the EoR window.

The next step is to perform a weighted average to condense the three dimensional data into a one dimensional power spectrum. The underlying power spectrum is expected to be isotropic, so averaging in spherical shells of constant $|\mathbf{k}|$ is appropriate. As discussed earlier, the velocity distortion terms cause the power spectrum to be anisotropic, but can be remedied by dividing the signal and noise by the angular dependence, $(1 + 2\mu^2 + \mu^4)$ in our case. Then voxels within a constant k shell have the same power spectrum signal, and can be averaged weighting by the uncertainty per voxel.

Figure 4.4 shows the sensitivity of the MWA to this EoR power spectrum. The theoretical one dimensional spherically averaged power spectrum (dotted blue line) and the uncertainty per bin (various step lines) are plotted. The uncertainty is plotted as a step function to show the binning used in the spherical average with the edges of the steps corresponding to the white curved lines in Figure 4.3.

The uncertainty is shown for 450 and 900 hours on one field, as well as a two field observation with 900 hours on one field (RA = 6^{h}) and 700 hours on a second (RA = 0^{h}), corresponding to one full season of observation. The lowest k bin approaches the sample variance limit as the signal to noise per voxel reaches ~ 1 and the array begins to image the largest EoR scales. The higher k bins, however, are thermal noise dominated at 900 hours.

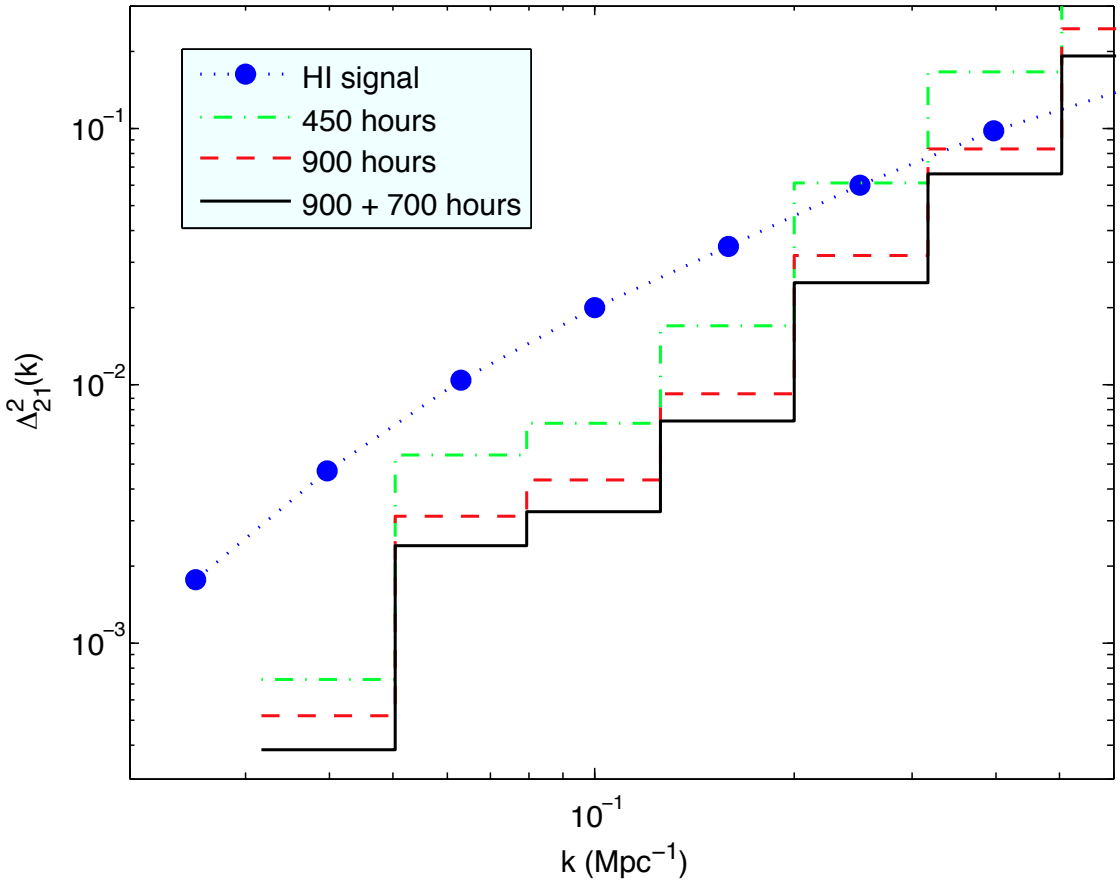

Figure 4.4: Estimated 1D sensitivity for the MWA, for various integration scenarios. The dotted blue line is the theoretical spherically averaged power spectrum [26], where $\Delta^2_{21}(k) = P_{21}(k)k^3/(2\pi^2 T_0^2)$ and $T_0 = 28[(1+z)/10]^{1/2}$ mK ≈ 26.6 mK. The several step functions represent the uncertainty per bin, with the edges of the steps corresponding to the edges of the bins when averaging (white curves in Figure 4.3). Single field observations are shown for 450 hours (dash-dot green) and 900 hours (dashed red) of integration. The solid black line corresponds to 900 hours on a primary EoR field, combined with 700 hours on a secondary field. This averaging excluded any data that would be contaminated by foreground subtraction (below the horizontal line, and to the right of the wedge in Figure 4.3).

We can also follow Lidz et al. [46] and fit an amplitude and slope to $\ln \Delta^2_{21}(k)$ in $\ln(k)$,

$$\ln \Delta^2_{21}(k) = \ln \Delta^2_{21}(k = k_p) + \alpha \ln(k/k_p), \tag{4.5}$$

where k_p is a fixed pivot wavenumber. The uncertainty on the amplitude depends on the pivot wavenumber, and we choose $k_p = 0.06$ Mpc^{-1}. The uncertainly is estimated assuming Gaussian statistics, and we fit directly in the 3D k-space to avoid binning effects and biases. For a full season of observation (900 hours on a primary field, 700 hours on a secondary), we predict a SNR of 14 on the amplitude and 10.9 on the slope (α) for the fiducial model. This does not take into account instrument downtime due to inevitable maintenance, nor loss of data for unforeseen reasons. With a more conservative observation time of 450 hours on a single field, we expect a SNR of 7.1 on the amplitude and 5.0 on the slope. Even with less than half a full observing season, the MWA has the potential for an EoR detection.

This calculation does not account for systematic biases from calibration and foreground subtraction errors. Efforts are underway to understand these affects and to achieve this level of sensitivity [91].

4.3 Sensitivity Conclusions

Using the layout of the 128 antenna MWA, we have estimated the instrument sensitivity to a model EoR power spectrum, taking into account synthesis rotation, chromatic and asymmetrical baseline effects, and excluding modes that are contaminated by foreground subtraction. We provide the tools required to calculate the MWA sensitivity for any model. With an optimistic full season of observation, we would expect to detect the fiducial power spectrum amplitude with SNR ~ 14, and constrain the slope with SNR ~ 10.9.

Chapter 5

ADDING CONTEXT TO JWST SURVEYS WITH CURRENT AND FUTURE 21CM RADIO OBSERVATIONS

"If you shut your eyes and are a lucky one, you may see at times a shapeless pool of lovely pale colours suspended in the darkness; then if you squeeze your eyes tighter, the pool begins to take shape, and the colours become so vivid that with another squeeze they must go on fire."

– J.M. Barrie, *Peter and Wendy*

No matter how hard they squeeze, the first generation of radio EoR experiments will not have the sensitivity to make full images of the highly redshifted IGM. However, they may be able to provide context for other wavelength observations such as the infrared galaxy surveys by the James Webb Space Telescope. This chapter was adapted from Beardsley et al. [5] and includes minor changes for coherency with this dissertation.

5.1 Motivation for Imaging

Due to an extremely faint EoR signal and high thermal noise, the first generation of radio experiments was designed to perform a statistical measurement of the EoR in the form of the cosmological power spectrum. The second generation of such experiments is just on the horizon with the Hydrogen Epoch of Reionization Array (HERA[1]), which is also designed to perform a power spectrum measurement but with much higher sensitivity and design decisions informed by first generation experiments (DeBoer, et al. 2014, in prep). While the first generation of radio arrays will not have sufficient sensitivity to produce a full image of

[1] http://reionization.org

the EoR, recent studies of instrumental effects and sensitivity have motivated the possibility of imaging at very large scales [52, 16].

Meanwhile, the JWST (James Webb Space Telescope[2]) is in development to explore the EoR in the infrared regime with a much higher resolution and smaller field of view compared to the ground based radio projects [27]. This instrument will have sufficient sensitivity and resolution to identify and study individual galaxies during the EoR (e.g. [96]). Due to its small field of view ($\sim$ 2 arcmin), the JWST cannot perform a comprehensive survey of the EoR. Instead it will need to focus on small patches of the sky to collect a representative sample of the Universe's evolution. A 21cm map of the IGM will then allow JWST to correlate the ionization fraction of the gas with properties such as luminosity functions, spectral energy distributions, morphologies, and the emission strengths of lines like Ly-α or Hα. These correlations in turn can answer questions about how the environment of galaxies affect how they form and evolve.

Below we demonstrate the potential to produce such a map using a current radio array (MWA) and a future array (HERA) to provide information about the IGM environment of the galaxies which will be studied by the JWST. We use a realistic HI signal simulation and take into account several dominant instrumental effects when creating our map. There is potential for similar cross-studies with other experiments such as the LSST deep drilling fields [44], a WFIRST Guest Observer program [82], or existing Hubble Space Telescope fields [28]. Lidz et al. [45] explored the merit of a cross power spectrum between 21cm maps and an extended Subaru survey [38], however Subaru has very small overlapping survey area with selected fields for deep integration by the MWA, and no overlap with HERA. We focus on the JWST as one of its core science programs is to study the EoR, and it will produce very deep pointed images at the redshifts of interest.

The remainder of this chapter is organized as follows: in Section 5.2 we describe the simulation used, in Section 5.3 we describe our instrumental model including thermal noise

[2]http://www.jwst.nasa.gov/

and foreground contamination, in Section 5.4 we show the results of our imaging plan, and we conclude and discuss future work in Section 5.5.

Throughout this chapter we use a ΛCDM cosmology with $\Omega_m = 0.73$, $\Omega_\Lambda = 0.27$, and $h = 0.7$, consistent with WMAP seven year results [41]. All distances and wavenumbers are in comoving coordinates.

5.2 The 21cm Signal

To show the capability of the MWA and HERA to image the EoR we use a simulation of the 21 cm signal. We make use of a "semi-numeric" reionization simulation (e.g. [97, 54]) based on excursion set modeling of the EoR [25]. Specifically, we use the models from Malloy and Lidz [52], which track the 21cm brightness temperature field across 512^3 cells in a periodic, 1 co-moving Gpc/h simulation box.

The contrast between 21 cm brightness temperature and the CMB at a given location, $\mathbf{r}$, can be approximately expressed in terms of the fractional baryon over density, δ_ρ, and the local ionization fraction, x_i, as

$$\delta T_b(\mathbf{r}) = T_0(1 - x_i(\mathbf{r}))(1 + \delta_\rho(\mathbf{r})) \tag{5.1}$$

where $T_0 = 28[(1+z)/10]^{1/2}$ mK. We will use a simulation cube of δT_b as the signal to be detected by our instrument, and the corresponding cube of x_i to inform us of the underlying ionization environment. We also use this simulation to make a power spectrum to use in our filter (Equation 5.4). The simulation box is 1 h^{-1}Gpc on a side, while the MWA field of view at $z_{\rm fid} = 6.9$ is much larger. To account for this we tile the simulation cubes. The periodic boundary conditions ensure no artificial discontinuities are introduced.

For this study we use a mean ionization fraction $\bar{x}_i = 0.79$ as a benchmark because the bubbles are sufficiently large to do some imaging, but there is still enough neutral gas around to leave some contrast in the brightness temperature field. A slice of constant $r_{||}$ (line of sight) of the brightness temperature cube is shown in Figure 5.3(a). The characteristic bubble size in this image is significantly smaller than the expected resolution of the MWA

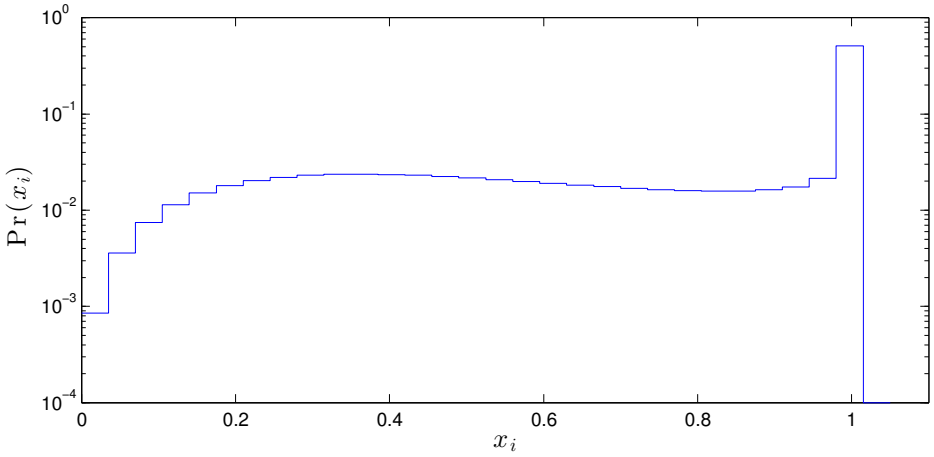

Figure 5.1: Histogram of ionization fraction x_i for the fiducial input simulation. More than half (51%) of the pixels are fully ionized ($x_i > 0.98$), while the remaining pixels have a very broad and flat distribution ranging from fully neutral to nearly ionized.

at this sensitivity. However, we aim to extract meaningful information about ionization environment without explicitly identifying bubbles or characterizing their sizes.

As a reference for upcoming plots, a histogram of the ionization fraction per pixel of the fiducial input simulation is shown in Figure 5.1. By a happy coincidence, at the redshifts studied here, the simulation pixel size is approximately the same size as the field of view of JWST (1.1 arcmin and 2.2 arcmin respectively). The per pixel ionization fraction should be interpreted as a smearing to the scale of the simulation pixel size, or half of a JWST field. With sufficient resolution (scale of individual galaxies), every pixel would either be fully ionized or fully neutral, and the distribution would simply have 79% pixels fully ionized. Because of our finite sized pixels, we instead have two distributions. Approximately half of pixels reside in bubbles larger than the pixel size and are fully ionized ($> 98\%$, last bin), while the majority of the remaining pixels have a very broad distribution spanning fully neutral ($x_i = 0$) to nearly completely ionized ($x_i \lesssim 0.98$).

A blind JWST survey would probe galaxies that reside in IGM with the population

shown in Figure 5.1. While any random pointing without any prior information has about a 51% chance of being completely ionized and 23% chance of being mostly neutral, the galaxy surveys will not be able to differentiate between IGM environments. Our goal is to distinguish the regions of fully ionized gas from those less than half ionized even without fully resolving the reionization bubble distribution, providing two separate populations for the JWST to correlate against.

5.3 Instrument Model

We use a realistic model of the current MWA taking into account several instrumental effects to model the thermal noise and the foreground contamination. We outline the highlights of this model here, and a detailed description can be found in Beardsley et al. [4]. The observing parameters used can be found in Table 5.1.

Each visibility measurement of an interferometer is sensitive to a Fourier mode perpendicular to the line of sight, given by $\mathbf{k}_\perp = 2\pi\mathbf{u}/D_m(z)$ where $\mathbf{u}$ is the baseline vector in units of wavelengths and $D_m(z)$ is the transverse comoving distance to the observation at redshift z [34]. Additionally, the frequency dimension maps to $r_{||}$ through the cosmological redshift due to the expansion of the Universe. The thermal noise on each visibility is

$$V_{\rm rms}(\mathbf{k}_\perp, r_{||}) = \frac{c^2 T_{\rm sys}}{\nu^2 A_e \sqrt{\Delta\nu\tau}} \tag{5.2}$$

where $T_{\rm sys}$ is the system temperature, ν is the frequency of observation, A_e is the effective collecting area per antenna, $\Delta\nu$ is the frequency channel width, and τ is the integration time. This can easily be expanded to a long tracked observation by allowing τ to represent the total time observing a given $(\mathbf{k}_\perp, r_{||})$ bin including redundant baselines and accounting for rotation of the earth.

We can then create a noisy image by adding Gaussian random noise at the level of Equation 5.2 to the Fourier transform of the simulation brightness temperature cube, where τ is independently calculated for each pixel.

$$\delta\tilde{T}_b'(\mathbf{k}) = \delta\tilde{T}_b(\mathbf{k}) + \tilde{n}(\mathbf{k}) \tag{5.3}$$

Table 5.1: Observational parameters for imaging simulation

Parameter	MWA	HERA
No. of antennas	112*	37, 127, 331
Central frequency	180 MHz (z $\sim$ 6.9)	180 MHz
Field of view	27.2°	7.3°
Effective area per antenna	20 m^2	154 m^2
Total bandwidth	6 MHz	6 MHz
T_{sys}	315.5 K	315.5 K
Channel width	80 kHz	96 kHz
Latitude	-26.701°	-30.0°
Integration per day	6 hr	32 min
Total integration	1,000 hr	120 hr

*Sixteen of the 128 MWA antennas lie significantly far from the core to provide high resolution. However, in the very low signal to noise regime these antennas offer very little sensitivity and greatly increase the size of the uv plane, so are omitted from this work.

Here we use a prime to represent a measurement of a signal and a tilde to represent a variable in Fourier space. In the step from Equation 5.2 to Equation 5.3, we implicitly performed a Fourier transform of the noise along the line of sight. This is a straightforward operation with uniform noise in the direction of the transform (see Morales [58] for a full derivation).

Next we apply a filter to the measurement which will do two things: it will mask out modes which are expected to be contaminated by foregrounds, and it will down-weight low signal to noise modes.

Foregrounds are expected to be a major challenge for radio observations of the EoR, however the effects of foreground contamination have been studied in great detail and shown to be restricted to low $k_{||}$ modes and a so-called "wedge" (see, e.g. [49, 50, 74, 33, 89, 91, 61, 92, 21, 13, 60]). How far the wedge bleeds into high $k_{||}$ and whether cosmological information can be extracted is an active topic of investigation. We adopt a fairly optimistic fiducial and use the field of view as the limit of foreground contamination, but explore other choices in the following section. The foreground mask is shown in Figure 5.2(a).

To down weight low signal to noise modes we use the Wiener filter, which is defined by the power spectra of the signal, $P_{21}(\mathbf{k})$, and the noise, $P_N(\mathbf{k})$. The signal power spectrum will be measured from actual observation data, but here is estimated from the simulation cube itself. It is important to note here that the form of our filter, and thus subsequent analysis, is dependent on the HI power spectrum. For our fiducial test, the signal power spectrum is calculated from the simulation image shown in Figure 5.3(a). Figure 1 of Lidz et al. [46] explores the form and evolution of the HI power spectrum. While the methodologies, assumptions, and simulation size used there are different than our current simulation, the power spectra are qualitatively similar and the differences have minimal impact on our results. We explore the dependence of our results on the ionization fraction (which in turn can be used as a proxy for bubble sizes) in Section 5.4. The noise power spectrum comes from

propagating Equation 5.2 to power spectrum space. The resulting filter can be expressed as

$$F(\mathbf{k}) = \begin{cases} 0 & \text{if contaminated} \\ P_{21}(\mathbf{k})/(P_{21}(\mathbf{k}) + P_N(\mathbf{k})) & \text{otherwise.} \end{cases} \tag{5.4}$$

Our estimate of the EoR image can then be written as the inverse Fourier transform of the filter multiplied by our k space measurement.

$$\delta\hat{T}_b(\mathbf{r}) = N \times \mathcal{F}\mathcal{T}^{-1}\left[F(\mathbf{k})\delta\tilde{T}'_b(\mathbf{k})\right] \tag{5.5}$$

We also introduce a normalization factor, N, to set the scale of the resulting image. Because interferometers produce zero-mean images, and much of the power has been removed due to our filter, we normalize each filtered noisy image to have mean of zero and range of one with arbitrary units. This forces values to be between roughly -0.5 and 0.5, and allows us to plot images with very different filters on the same axes (see Section 5.4).

The shape of this filter in k space is shown in Figure 5.2(a). The horizontal dashed white line represents the low $k_{||}$ contaminated by foregrounds. Because the MWA observes with an instantaneous bandwidth of 30.72 MHz but the cosmological measurement is restricted to 6 MHz, the ~ 4 modes of expected contamination are restricted to only the lowest $k_{||}$ bin in our case. The right-most diagonal dashed white line represents the wedge characterized by the field of view of the MWA. All bins below the horizontal and to the right of the diagonal dashed line are masked out, and shown as dark blue.

The majority of the sensitivity lies at low $k_\perp$ and low $k_{||}$. This is due to a combination of lower thermal noise due to a dense array core and the signal power spectrum being high for small k. The sharp drop off of sensitivity at low $k_\perp$ is due to not having antennas placed physically closer than 7.5 m apart. However, there is *some* sensitivity there due to projection from earth rotation.

An image space representation of our filter is shown in Figure 5.2(b). The slice shown is for constant r_x=0 to show the shape in $r_{||}$ and $r_\perp$. In this space the field of view is much larger than the line of sight extent, so the image has been cropped in $r_\perp$ but not $r_{||}$. Because

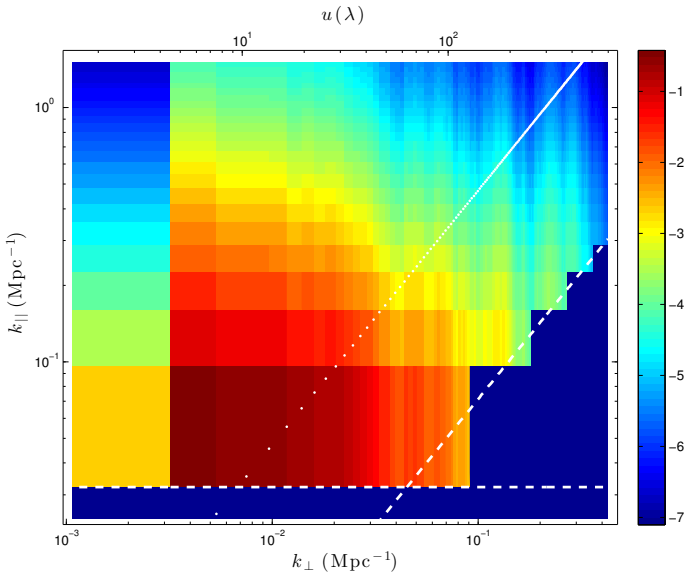

(a)

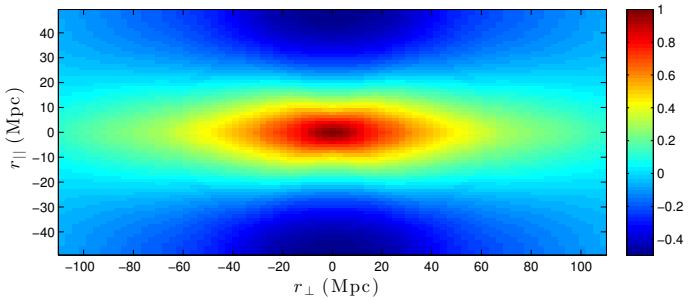

(b)

Figure 5.2: The filter used for our fiducial test. *Top:* A slice of the filter, $\log_{10}(F(k_\perp, k_{||}))$, in wavenumber space. A log color scale was used to emphasize the bins contributing to the filter. The bins below the horizontal white dashed line and to the right of the diagonal white dashed line are masked out due to foreground contamination, and shown in dark blue. Most of the sensitivity of our measurement lies at low $k_\perp$, low $k_{||}$. The sharp drop in sensitivity at the lowest $k_\perp$ bin is due to the minimum spacing of antennas in the MWA. *Bottom*: The image space representation of the filter, $F(r_\perp, r_{||})$, or the inverse Fourier transform of the top image. Shown is a slice of constant $r_x = 0$. Along the line of sight direction the filter ranges from negative response to a positive peak back to negative. The response is elongated in the perpendicular direction due to the smaller $k_\perp$ modes contributing.

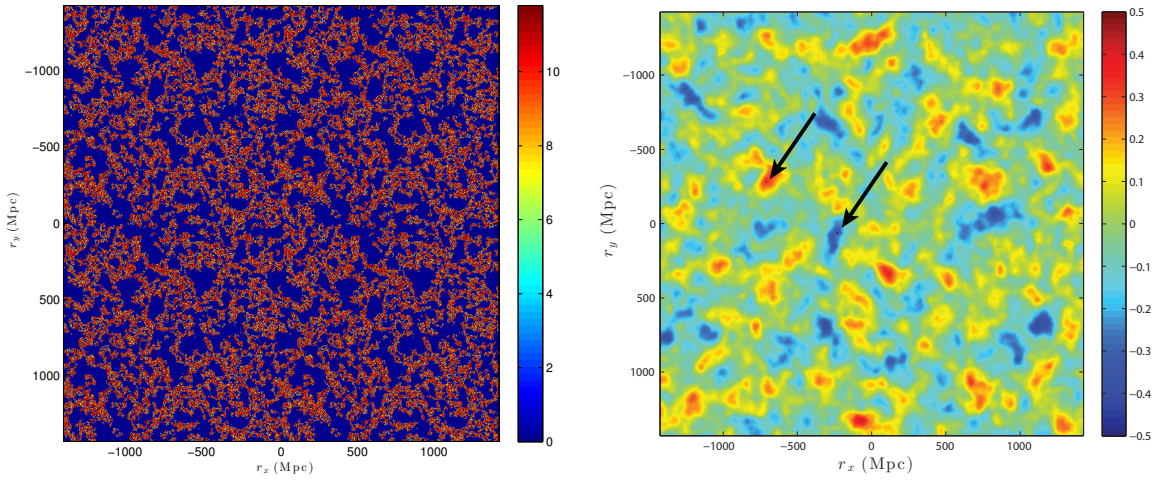

(a) (b)

Figure 5.3: *Left:* The input simulation brightness temperature, $\delta T_b(\mathbf{r})$, in mK. The dark blue regions are zero temperature and represent fully ionized bubbles. *Right:* The filtered noisy image produced by our fiducial MWA instrument model, in arbitrary units. Owing to the peculiar shape of the filter (Figure 5.2), the structure of the filtered image can look very different than the input image. The arrows are pointing to tiny black boxes approximately the size of the JWST field of view. A full infrared survey would probe many such regions of varying IGM ionization fractions, informed by the probability distributions shown in Figure 5.6.

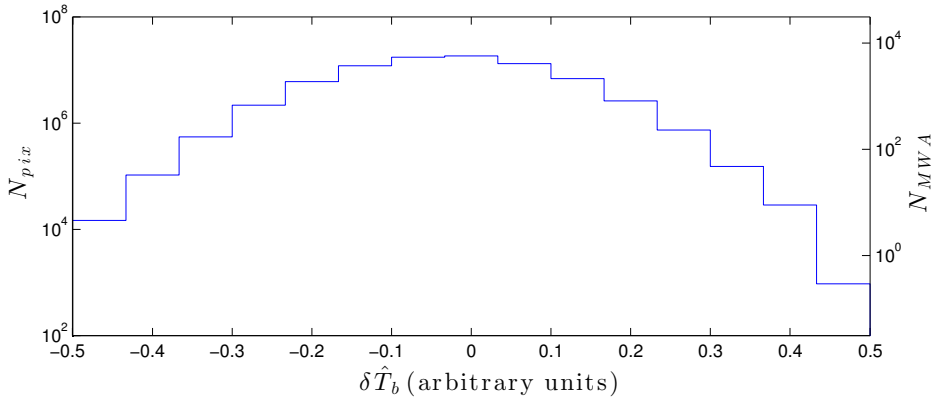

Figure 5.4: Histogram of the values in the filtered noisy image shown in Figure 5.3(b), using our fiducial MWA instrument model. The total number of pixels in the simulation is $\sim 8 \times 10^7$. The second vertical axis provides a reference to the approximate number of independent regions in the map, given the filter size for the MWA.

the majority of line of sight sensitivity comes from the first non-zero $k_{||}$ mode, the line of sight response resembles a single mode curve moving from negative response at the near edge, peaking in the center, and returning to negative at the far edge. In addition the response in the perpendicular direction is much more broad than the line of sight.

An example of a filtered noisy image is shown in Figure 5.3(b). Due to the peculiar filter shape and natural sensitivity of the instrument, the images contain very non-intuitive structures. Indeed it can be extremely difficult to gauge the quality of the images without a rigorous metric. In the following section we show that despite low signal to noise and a highly irregular response, we can retrieve information about the underlying ionization field.

For reference we show a histogram of the measured temperature values in Figure 5.4. For each temperature value, $N_{\rm pix}$ is the number of simulation pixels measured at that value. The pixels in our filtered map are not all independent owing to the relatively large point spread function of our filter, so we provide a second vertical axis ($N_{\rm MWA}$) showing the approximate number of independent regions in the map, determined by the size of the MWA filter.

5.4 Imaging Results

In this section we attempt to answer a simple question: Given a measured value in a noisy filtered 21 cm image, what, if anything, can we say about the ionization fraction of the IGM for the galaxies that JWST would observe at that location?

To answer this question, we first bin the values of the filtered noisy image into fifteen equal sized bins. Then for each binned temperature value, we construct the probability distribution of the underlying ionization fraction by histogramming the x_i values corresponding to the selected image pixels. The result for our fiducial MWA observation are shown in Figure 5.5.

Due to the extremely peaked input distribution (Figure 5.1), the resulting correlations are also heavily skewed toward high x_i. In other words, no matter what value is measured in our image, there is a significant probability it is a fully ionized region simply because most of the image is fully ionized, and our fiducial instrument does not have the resolution to decouple small pockets of ionized gas from neutral clouds.

But not all is lost. Two encouraging features can be seen in Figure 5.5. The first is that the highly ionized bin on the far right drops by about a half order of magnitude (note the log color scale - the corresponding values are 0.8 to 0.2) moving up in the figure. This means that a low observed temperature is much more likely to be fully ionized than a high value. The second feature is the gradual increase in probability of low ionization moving to higher brightness temperatures. However, like the input image, the distribution is quite broad and not peaked in any one place. In an attempt to quantify these features, we further bin the ionization fractions. We wish to ignore the broad, mildly ionized pixels and focus on two distinct sub-populations of pixels: those that are fully ionized ($x_i > 0.98$) and those that are less than half ionized ($x_i < 0.5$). These populations account for 51% and 23% of the input simulation pixels, respectively. We will find that division leads to a clearer distinction in the underlying x_i probability distributions.

In Figures 5.6-5.8 the solid line of a given color is the probability distribution of fully ionized pixels, while the dashed line of the same color corresponds to less than half ionized

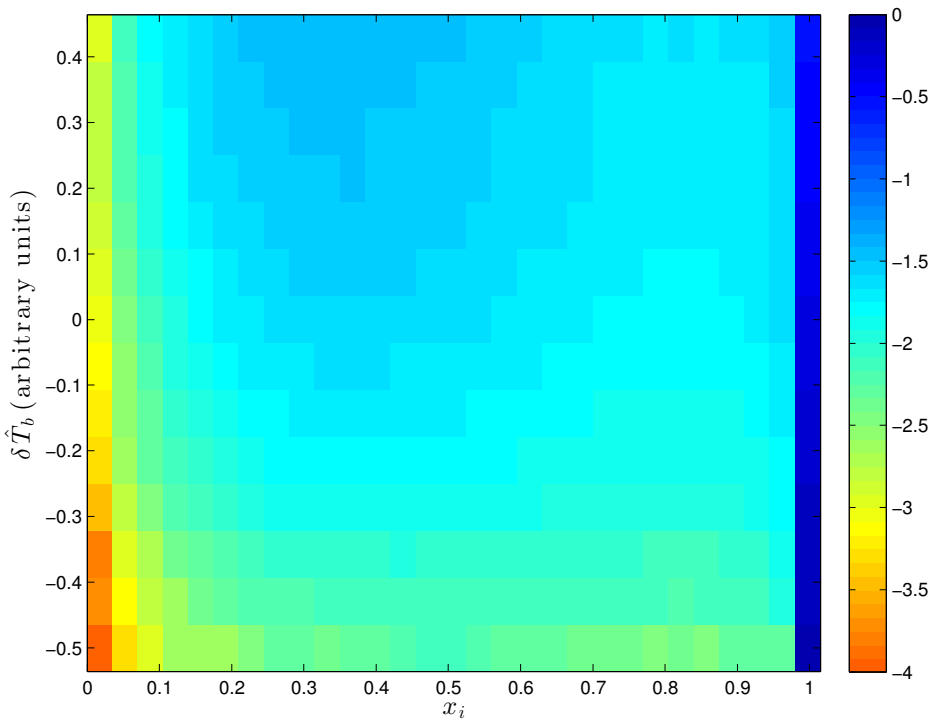

Figure 5.5: Image correlations with x_i (log color scale). For each binned observed temperature, $\delta\hat{T}_b$, we histogram the underlying pixels in the simulation x_i cube and normalize each row independently to a sum of one. Each row of the image can be viewed as the probability distribution of ionization fraction given a observed temperature. Most of the distribution lies at high x_i due to the highly peaked native distribution (Figure 5.1). Nevertheless, a correlation can be seen between measured value and input ionization fraction.

pixels. The colors then represent different trials which we explain in turn.

First let us examine the fiducial observation indicated by the black lines in Figure 5.6 and reproduced in the subsequent figures for reference. This trial was done using the MWA observation parameters listed in Table 5.1, a wedge defined by the field of view, and a mean ionization fraction of 79%. The most striking feature is the large separation in probabilities at low observed temperature, with the probability of fully ionized far exceeding the “blind" guess of 51%. Also encouraging to note is that moving up in filtered temperature, the solid and dashed black lines cross and separate again. This means that we can potentially distinguish two regions of ionization levels with our filtered noisy image.

We will see in the following sub-sections that the success of our filter to distinguish ionization regions is closely tied to the size of the point spread function compared to the size of the bubbles in the simulation. When the filtered instrument has sensitivity on order of the size of ionized bubbles, it tends to have success in identifying ionized regions which have low temperature due to the lack of neutral hydrogen. Similarly, where the instrument can resolve large regions of neutral hydrogen, we have success on the right side of the plot. It is not surprising that we tend to have higher certainty for ionized regions than we do for neutral regions because ionized bubbles tend to saturate with extended regions of purely ionized gas, while neutral clouds contain an admixture of neutral and partially ionized regions that current radio experiments cannot hope to resolve.

While we cannot claim with certainty the ionization fraction of any given pixel, we can produce a partial sky map of ionized and non-ionized regions based on a given probability threshold. With such a map, the JWST can form an observing plan to probe regions of high and low IGM ionization fraction, allowing the study of galactic properties in different environments. In Figure 5.3 we show schematically a couple fields that JWST could observe to probe a highly ionized region and a mostly neutral region. A comprehensive survey would involve many such fields, statistically incorporating our probability distribution functions.

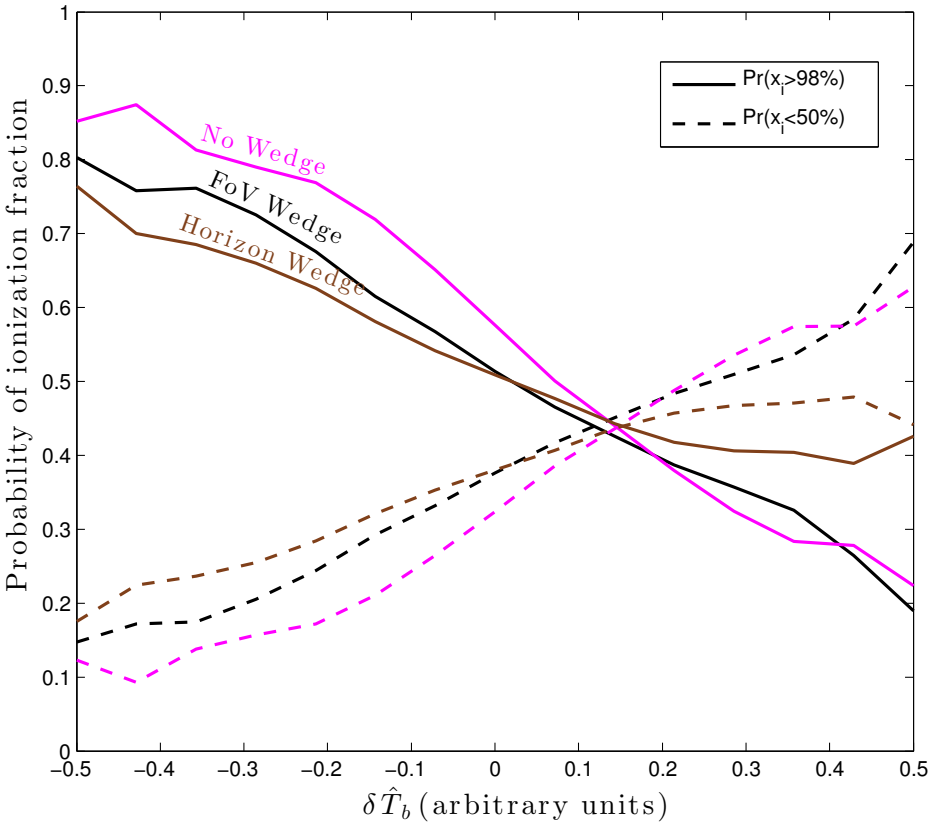

Figure 5.6: The probability of fully ionized pixels (solid lines) and mostly neutral pixels (dashed lines) as a function of measured temperature, $\delta\hat{T}_b$. The separation, cross, and re-separation of the solid and dashed lines show that our method is able to distinguish between the two populations of ionized environments. The black lines represent our fiducial MWA instrument with mean ionization fraction of 0.79 and a foreground wedge defined by the field of view. Also plotted are the very optimistic no wedge (magenta lines) and the pessimistic horizon wedge (brown lines).

5.4.1 The Influence of the Wedge

Many factors can influence the effectiveness of this strategy. Here we explore a few and leave a comprehensive study to future work. The first variable we explore is a choice of the wedge. Our fiducial choice is fairly optimistic and relies on no foreground contamination outside the field of view of the instrument.

Figure 5.6 shows the probability lines for two other choices as well. First is the extremely optimistic choice of there being no foregrounds at all. While there is no realistic expectation that current or future radio observations will be able to completely remove the wedge and recover all cosmological information from those modes, we include this line to show the effect of going from no foregrounds whatsoever to including a more realistic wedge. Indeed the effect is quite minimal. While the no wedge line shows a slightly higher probability of full ionization at low measured values, it has a crossing at roughly the same location as the field of view wedge line, and has very similar values on the right side of the plot.

Although quite unexpected, this surprisingly small effect can be understood by examining the filter shown in Figure 5.2(a). While there are many bins excluded below the white dashed lines, the strong majority of our sensitivity actually lies safely outside the wedge at low $k_\perp$. So while adding information will almost always help, it is easy to see why in this case the difference between a field of view wedge and no wedge at all is marginal.

Next we consider a pessimistic wedge defined by the horizon. The dotted white diagonal line in Figure 5.2(a) represents the cut made in Fourier space for this choice. Anything below that line is assumed to be contaminated by foregrounds and omitted. Here we can see the wedge digs into the $k_\perp$ sensitivity of the instrument, and so we expect a poorer resolution. In Figure 5.6 we see the effect on the correlation with x_i values. The separation at low observed temperatures is marginally worse than the field of view wedge case. However, the separation at high temperature values is essentially nonexistent. The poor resolution due to foreground contamination has washed out any ability of our instrument to distinguish small pockets of neutral gas.

Even with a pessimistic outlook on foreground mitigation, a separation of distributions at low $\delta\hat{T}_b$ offers valuable information for a JWST survey. But to get the most out of these images, it is necessary to continue investigation into foreground subtraction and avoidance to try to recover as much of the wedge as possible.

5.4.2 Mean Ionization Fraction

Next we turn to uncertainty in the ionization history. Although the power spectrum measurement will be able to constrain models of reionization, it is not clear how well the mean ionization fraction $\bar{x}_i$ for a given redshift will be known (see Pober, et al. 2015, in prep., for a recent discussion). We explore this uncertainty by producing images for simulation cubes at mean ionization fractions of $\bar{x}_i = 0.68$ and 0.89 in addition to our fiducial value of 0.79.

The results are shown in Figure 5.7. The effect of a higher than expected ionization fraction is a wider separation at low $\delta\hat{T}$, but no cross over between populations. This can be explained with the fact that the input image is dominated by large ionized bubbles. Our instrument has low resolution, so is better at picking out large bubbles (low observed temperatures), but the remaining neutral clouds are limited to small isolated regions to which our instrument is blind.

On the other hand, a lower mean ionization yields a small separation at low $\delta\hat{T}_b$ and large separation at large values. The explanation for this effect is the same as the latter but run in reverse. At low ionization fraction the bubbles are smaller and our instrument has a harder time detecting them. But the neutral clouds are large and easily distinguished.

From this test we learn the importance of considering the uncertainty of ionization fraction. While any one of these lines can provide valuable IGM information to JWST surveys, a full marginalization over uncertainty should be performed to ensure accurate statistics. The details however, are outside the scope of this paper and ultimately depend on the detected power spectrum and the uncertainty it yields on ionization fraction. To rough approximation, if the ionization fraction can be known to ± 0.1, the red and blue curves can be interpreted as the error bars on the black curve.

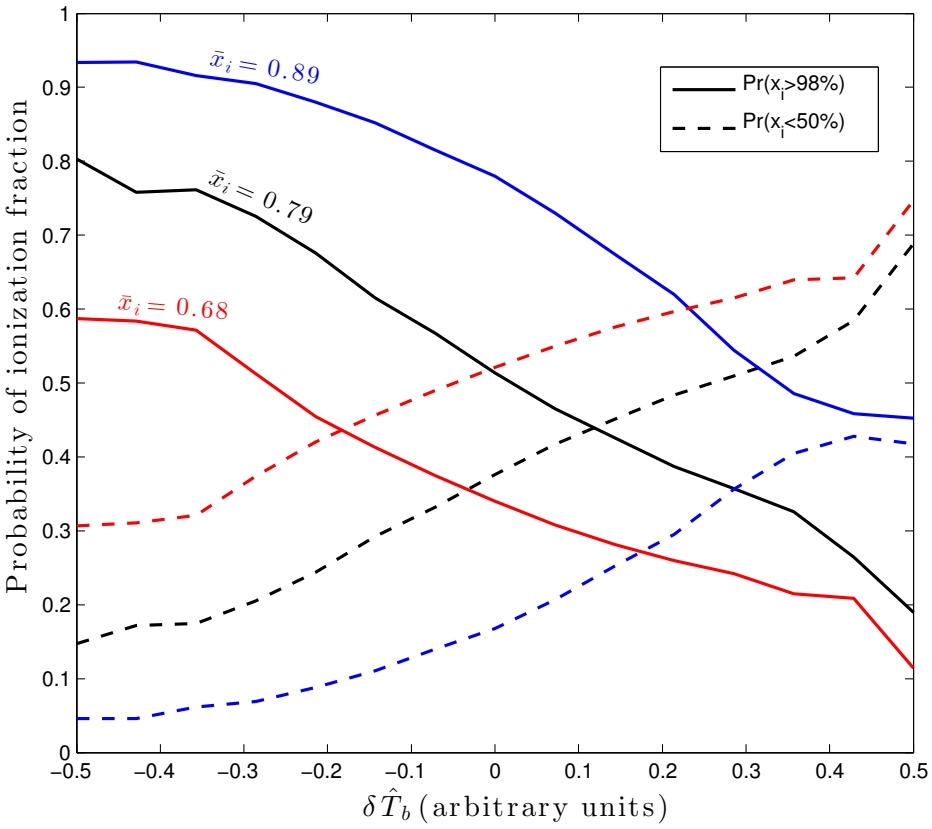

Figure 5.7: Probabilities of ionization populations for varying mean ionization fraction. The fiducial lines are reproduced in black, while the blue lines show a more ionized Universe ($\bar{x}_i = 0.89$) and the red lines represent the less ionized $\bar{x}_i = 0.68$. As in Figure 5.6, the solid lines represent fully ionized regions while the dashed lines represent less than half ionized.

5.4.3 Future Observations with HERA

Finally, we set our sights on the next generation of radio EoR experiments. HERA was designed to have much larger collecting area per element as well as highly redundant baselines for increased sensitivity. The larger area per element results in a more narrow field of view. Additionally, the elements cannot point and thus HERA is naturally a drift instrument resulting in far less integration time per day on a single field. The dishes are arranged in a closely packed hexagonal configuration. Current preliminary funding for HERA is sufficient to build 37 dishes in South Africa, while additional funding could expand to 331 dishes. The observing parameters for our HERA model are shown in the right column of Table 5.1. Note the drastically smaller integration time used for HERA compared to the fiducial MWA.

The results of our HERA simulation are shown in Figure 5.8. With only 120 hours of integration, the currently funded 37 dish instrument performs very comparable to the fiducial MWA with 1,000 hours. Adding more dishes substantially improves the sensitivity. With an intermediate 127 dishes, HERA will be able to identify ionized regions with near certainty, and an improved ability to distinguish neutral regions (high values of $\delta\hat{T}_b$). This may seem counter-intuitive due to the shorter baselines of the tightly-packed HERA compared to the MWA. But due to the large elements and redundant configuration, HERA actually has high sensitivity out to larger baselines than the MWA, resulting in a better resolution after the filter is applied. Therefore it is more capable of distinguishing ionized from neutral regions.

Finally, as expected the fully proposed 331 dish HERA will perform the best of our tests. Not only can it identify fully ionized regions with near 100% certainty like the 127 dish version, but it can also identify mostly neutral regions to $\gtrsim$80% confidence. The increased confidence on either end of the distribution will help constrain the correlations with JWST galaxy surveys.

Our HERA calculations have been performed on a single field of view region of the sky. Because HERA is a drift instrument, it will observe many such fields every day resulting in a strip of sky at this level of sensitivity.

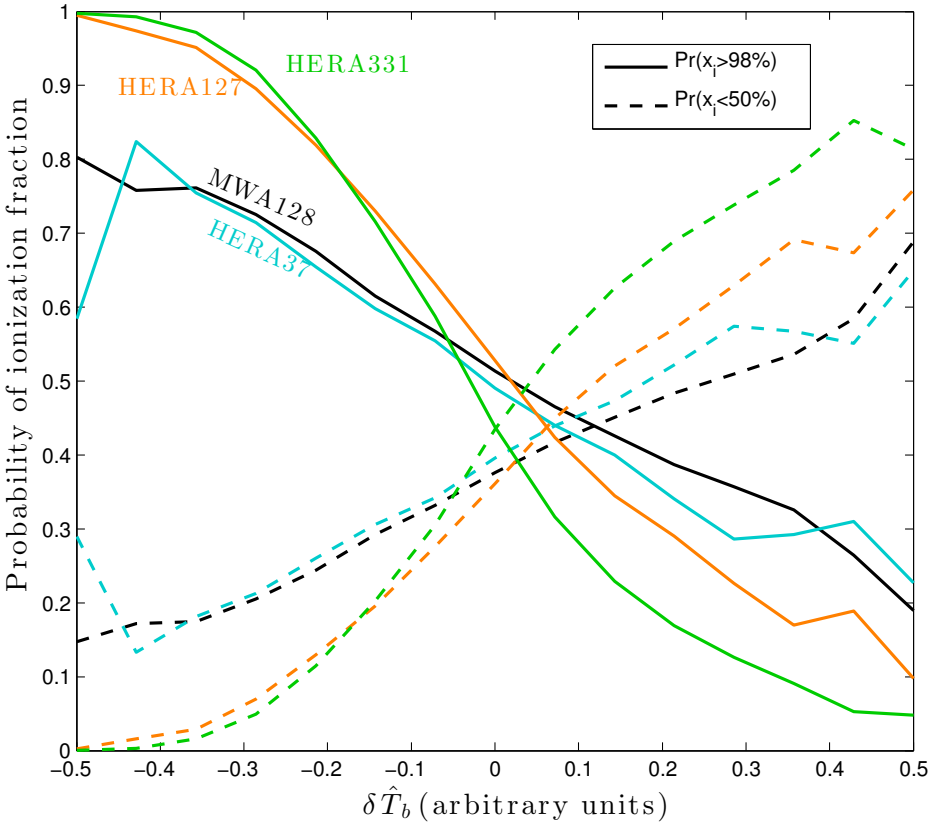

Figure 5.8: Probabilities of ionization populations for the HERA instrument. The currently funded 37 dish HERA (cyan lines) perform comparable to the fiducial MWA, while further buildouts of 127 dishes (orange lines) and 331 dishes (green lines) strongly outperform our fiducial and offer the most information for correlations with JWST.

5.5 Imaging Conclusions and Further Work

We have shown that current radio interferometers, despite being designed for power spectrum measurements, can produce useful image information that can provide context to future galaxy surveys by instruments like the JWST. The probability distribution functions we have shown in Figures 5.6-5.8 will allow for a statistical study of local ionization fraction against galaxy properties such as luminosity functions, spectral energy distributions, and morphologies probed by infrared surveys. Such studies will enable us to answer questions such as:

- How does the apparent age of stars in galaxies depend on ionization environment?
- What is the dust content of galaxies residing in neutral environments versus ionized?
- How did reionization affect the shapes and sizes of galaxies forming from ionized gas?

In our analysis the foreground contaminated "wedge" proved to be a less significant factor than expected. While there exist gains between the horizon wedge and the field of view wedge, the difference between field of view and no wedge at all is only marginal. The highly active investigations of foreground contamination will hopefully allow imaging to push down to the field of view.

Our technique is also robust against uncertainties in mean ionization fraction. A more comprehensive marginalization over $\bar{x}_i$ is needed before correlations with future surveys can be performed. However, the marginalization will depend strongly on the constraints provided by power spectrum measurements and other experiments.

Our method may be improved by considering other filters and imaging strategies. The Wiener filter was used here as a benchmark to yield high signal to noise, however further optimization may be achieved by considering the probability distributions of Figures 5.6-5.8 as a figure of merit. Here we avoided attempts to make claims on the sizes of the ionized

bubbles and instead focused on an instrument-driven filter, but additional work with matched filters akin to [52] will yield additional scientifically interesting information.

An additional followup to this work would be to correlate the filtered noisy images with other galaxy properties in the simulations. For example interesting cosmological and astrophysical information is contained in the actual timing of reionized bubbles (z_{re}) or the number of galaxies with halos over some threshold mass within ionized regions. Measurements of this variety would give more handles on the underlying reionization models.

The MWA has completed its first year of EoR observing, and has recently begun its second year which will nearly triple the total observing time. The work currently being done to eliminate systematics and reduce the data to a power spectrum measurement is essential in the path towards our proposed imaging project. Meanwhile HERA is pushing forward and will begin construction in 2015. This next generation instrument will build on the lessons learned thus far and provide a much more sensitive power spectrum measurement, as well as potential for imaging. Observations from these instruments will allow us to provide IGM context to the deep galaxy surveys of instruments like the JWST which will study the objects reionizing the Universe.

Chapter 6
ANALYSIS PIPELINE

The MWA has been performing science observations for a year and a half, and has accumulated just shy of 1,700 hours on EoR fields, amounting to about three petabytes of raw data. Our next challenge is to deal with this large amount of data and reduce it to a cosmological power spectrum. There are many steps in this process and this chapter is dedicated to outlining our strategy.

At a top level, the pipeline can be described with the following steps.

1. *Preprocessing.* The raw correlator data is reformatted, flagged for RFI contamination, and averaged.

2. *Calibration.* We adjust the gains of the antennas based on a sky model.

3. *Imaging.* Frequency-dependent images of the sky are formed on a snapshot basis. There exist analysis strategies which bypass this step, however it is an essential step in the pipeline for this work.

4. *Integration.* The snapshot images are accumulated. While a simple step, I enumerate it here to mark the point where any data selection must occur.

5. *Forming a power spectrum.* Finally, we Fourier transform the integrated image and form a power spectrum.

The international MWA EoR collaboration has chosen to tackle this challenge by developing two independent pipelines which are able to exchange data products at different points in the pipes for quality assurance and cross checking. The details of this strategy are outlined

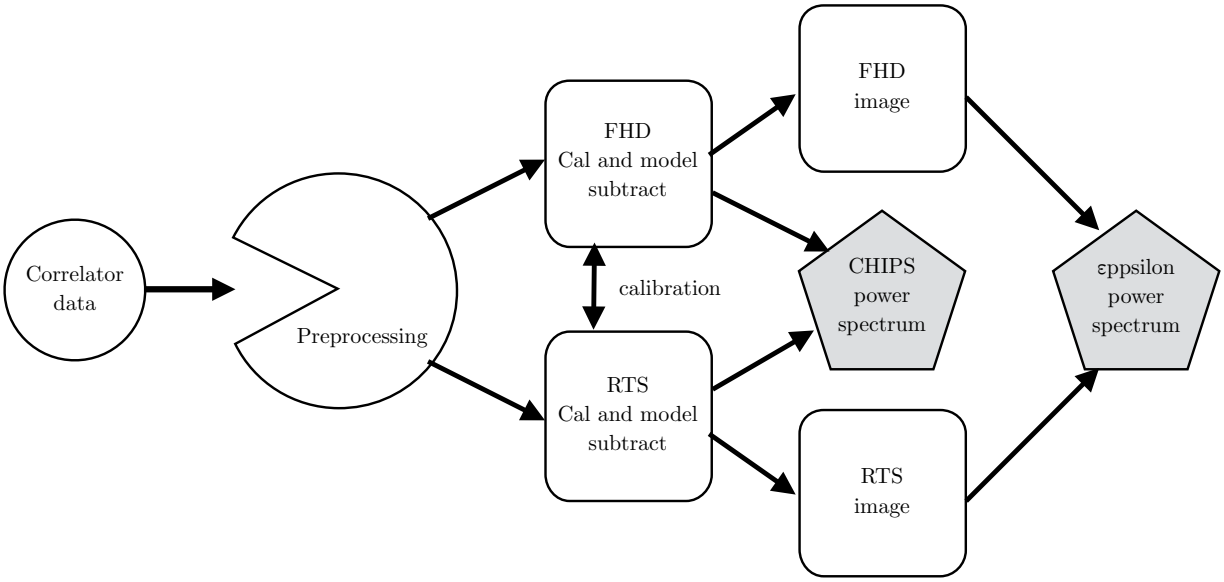

Figure 6.1: The international MWA EoR analysis strategy (adapted from D.C. Jacobs et al. 2015, in preparation). Two reference pipelines have been identified as the main thrust of effort for the collaboration. The U.S. pipe (upper path) involves calibrating and imaging with the FHD package, and forming power spectra from output images using the εppsilon pipe. The Australian pipe (lower path) uses the RTS for calibration, and CHIPS to form power spectra directly from calibrated visibilities. The pipes also have several points to compare data products and exchange outputs.

in D.C. Jacobs et al. 2015 (in preparation). Here I reproduce the key strategy with Figure 6.1. The top path, which follows FHD to εppsilon, is primarily developed by U.S. partners, and is the focus of the pipe described here. The lower path, following RTS to CHIPS, is primarily driven by Australian institutions.

The pipes share a common processing step, which is described in detail in Section 6.1. The pipes then diverge with the U.S. pipe utilizing the FHD framework for calibration, model subtraction, and imaging, and is described in Section 6.2. For these steps the Australian

pipeline uses the Real Time System (RTS) which was originally developed to calibrate and image MWA data, as the name suggests, in real time [55, 64, 93]. With the rescoping of the instrument to 128 antennas, the RTS was repurposed as an offline analysis tool but retains many of the design features originally proposed.

The calibration and imaging pipes then hand the data off to the power spectrum pipe segments. The Cosmological HI Power Spectrum (CHIPS) pipeline is based on an optimal inverse covariance weighting technique, and requires the calibrated visibilities (C. M. Trott et al. 2015, in preparation). The Error Propagated Power Spectra with Interleaved Observed Noise (stylized εppsilon) pipeline uses the frequency dependent images along with several other data products to form power spectra and propagate the error bars directly through the pipe (Hazelton et al. 2015, in preparation). The FHD to εppsilon path is the focus of this chapter.

In order to develop the various pieces of the pipeline, an early set of data was identified to represent a “golden set." The golden set is about three hours of data (94 112 second snapshots) taken on August 23, 2013, and was chosen for a few reasons. This was one of the first days of science operations where no major hardware issues were known, the specific snapshots chosen were pointed at the “EoR0" field, which is relatively devoid of dominant point sources or galactic structure, the field was relatively high elevation, and the bulk of the galactic disk was set below the horizon. Choosing such a data set helped to ensure many of the challenges to come would be minimized while we hone the analysis pipelines. It is also helpful to identify such a standard set so the collaboration is all working on the same data for comparisons and cross-checks. Many of the sample plots and figures to come are taken from the golden set.

The remainder of this chapter is organized as follows. In Section 6.1 I will describe the preprocessing step and present statistics on the RFI flagging done on MWA data. In section 6.2 I will describe the calibration, model subtraction, and imaging done with the FHD framework. Then in Section 6.3 I will describe the εppsilon pipeline which produces cosmological power spectra.

6.1 Preprocessing

The first step in analysis is to convert the raw correlator data to a more useful format, perform RFI flagging, and average to reduce data volume. We perform these steps with the `COTTER`[1] package, which in turn calls the `AOFLAGGER`[2] to do RFI flagging [62].

The base unit of data out of the correlator is a single observation, or a snapshot. In the observing mode for this work, a snapshot is 112 seconds of visibility data at 0.5 second, 40 kHz resolution. The raw data is in a non-standard memory dump format from the graphical processing units (GPUs) used in the correlator to perform the cross multiply step. Each 1.28 MHz frequency sub-band ("coarse channel") is written to a separate data file which contains all the visibilities for that coarse channel for a snapshot.

We run `COTTER` on all data ingested by the MIT archive using a semi-automated daemon (dubbed CHOMPY) which identifies newly arriving data, locates all files, launches `COTTER`, and populates a local metadata and quality control database. This database contains any information about observation settings or any quality metrics we find useful for making data selection cuts later in the pipe.

`COTTER` reads in the GPU files and performs RFI flagging. A full exploration of the RFI environment at the MWA site is presented in Offringa et al. [63]. The amount of flagging done on the data for this work is shown in Figure 6.2. The benefit of building the MWA in a remote location is clear – most snapshots contain $\lesssim 3\%$ RFI contaminated data. The outliers in the RFI distribution are mostly due to other problems with the instrument at certain times. Once our final data selection is made (Section 8.1), the maximum RFI flagging is 2.7%.

In addition we flag two 40 kHz channels at the bottom and top ends (total of four channels) of each coarse channel. This is due to known aliasing from the bandpass filters applied in the receivers (A.M. Levine, internal memo). In principle this effect may be calibrated

[1] `http://mwa-lfd.haystack.mit.edu/twiki/bin/view/Main/Cotter` (requires MWA collaboration login)

[2] `http://aoflagger.sourceforge.net/`

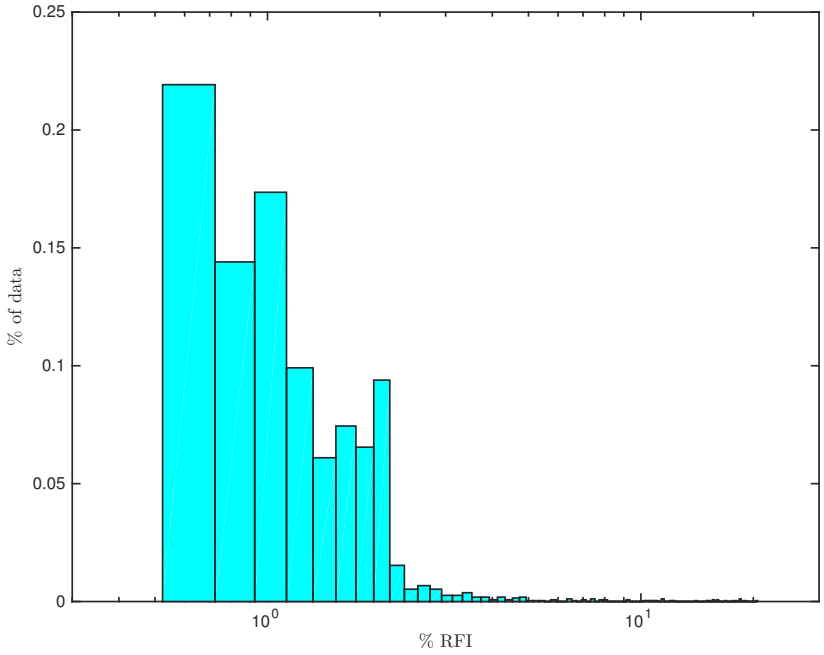

Figure 6.2: This histogram shows the amount of data flagged due to RFI contamination in all the data considered for this work. The fraction of flagging is calculated for each snapshot, then histogrammed to show the distribution. Almost all snapshots have $\lesssim 3\%$ RFI contamination. The outliers extending to tens of percent are later excluded from the analysis in Section 8.1. A full analysis of the RFI environment at the MWA is presented in Offringa et al. [63].

out in the future, but for early analysis it is safer to remove the offending channels. Due to early hardware bugs in the telescope, tiles sometimes were not pointed correctly at the beginning of observations. We thus also flag the initial two seconds of each observation to avoid mis-pointed data. This problem has since been solved, but exists in all data for this work. The central channel of each coarse band is also flagged because it corresponds to the DC mode of that sub-band, which can easily float and does not contain useful information.

Next, the data is averaged to 80 kHz frequency resolution and a 2 second cadence to reduce data volume then written to disk as UVFITS files [31], which is a visibility storage format based on the Flexible Image Transport System (FITS, [94]). Each UVFITS file contains the visibility data for a single snapshot (all frequency channels now), and is about 7.8 GB in size. At this stage the data is ready to be passed onto the calibration and imaging pipeline.

6.2 Calibration and Imaging

The Fast Holographic Deconvolution (FHD[3], [83]) framework was developed as a highly efficient method to identify and model sources in interferometric data. The algorithm is based on the A-projection/software-holography framework and has many features useful for cosmological observations with the MWA including wide-field polarimetry, antenna dependent beam models, and its ability to model realistic instrumental response. The software package has expanded beyond a deconvolution algorithm to include calibration, point source and diffuse modeling, and exporting of data products necessary for power spectrum estimation.

Full deconvolution is a very computationally expensive task, and performing it on every observation would be prohibitive for the large data set needed for an EoR measurement. Instead we have opted to use a subset of data to construct a sky model, which we then treat as an input to a “First Pass" analysis mode (see Figure 6.3). First Pass uses the sky model to create model visibilities, accounting for instrument measurement effects. These model

[3] `https://github.com/miguelfmorales/FHD`

visibilities are used for the calibration step as well as foreground subtraction.

6.2.1 Sky Model

The sky model is both an input and an output of our calibration and imaging pipeline, and thus we take an iterative approach. To provide a starting point, we used the MWA Commissioning Survey (MWACS) catalog which was created using mosaicked images from data while the MWA was being commissioned [35]. This catalog of point and extended sources served as our initial input to the pipeline shown in Figure 6.3.

With an initial sky model, we could create model visibilities. This is done by direct Fourier transforming the point sources of the catalog to a uv grid, and "measuring" the resulting coherence function by convolving with the MWA primary beam at all baseline locations. The MWACS catalog consists of over 14,000 point sources and it is impractical to perform the DFT for the entire catalog – especially when the process must be repeated for every snapshot as the earth rotates and the uv plane changes. We thus select sources within a certain threshold of our primary beam (chosen to be 5% of the maximum) to include in our model, reducing the number of sources to model to about 2,400 depending on the specific snapshot. An example of this selection is shown in Figure 6.4. We will see in Section 7.2 that it is also beneficial to include sources in the first sidelobes of the primary beam. This modeling process is also extended to include diffuse sky emission by over-resolving the sky beyond the resolution of the instrument and treating it as a series of point sources.

Already we see that a model of the primary beam is necessary to model the instrument response. We initially used an analytic beam model based on a short dipole response and the geometry of the MWA tile (see Appendix B). However, it has been shown that this model is insufficient for precision cosmology experiments, and we have since adopted a more robust model which includes simulated dipole response and mutual coupling between dipoles [84]. Further empirical models are under investigation, but have not yet matured enough to be included in our analysis (see A.R. Neben et al. 2015, submitted to Radio Science).

With a set of model visibilities we calibrate the measured visibilities using the algorithm

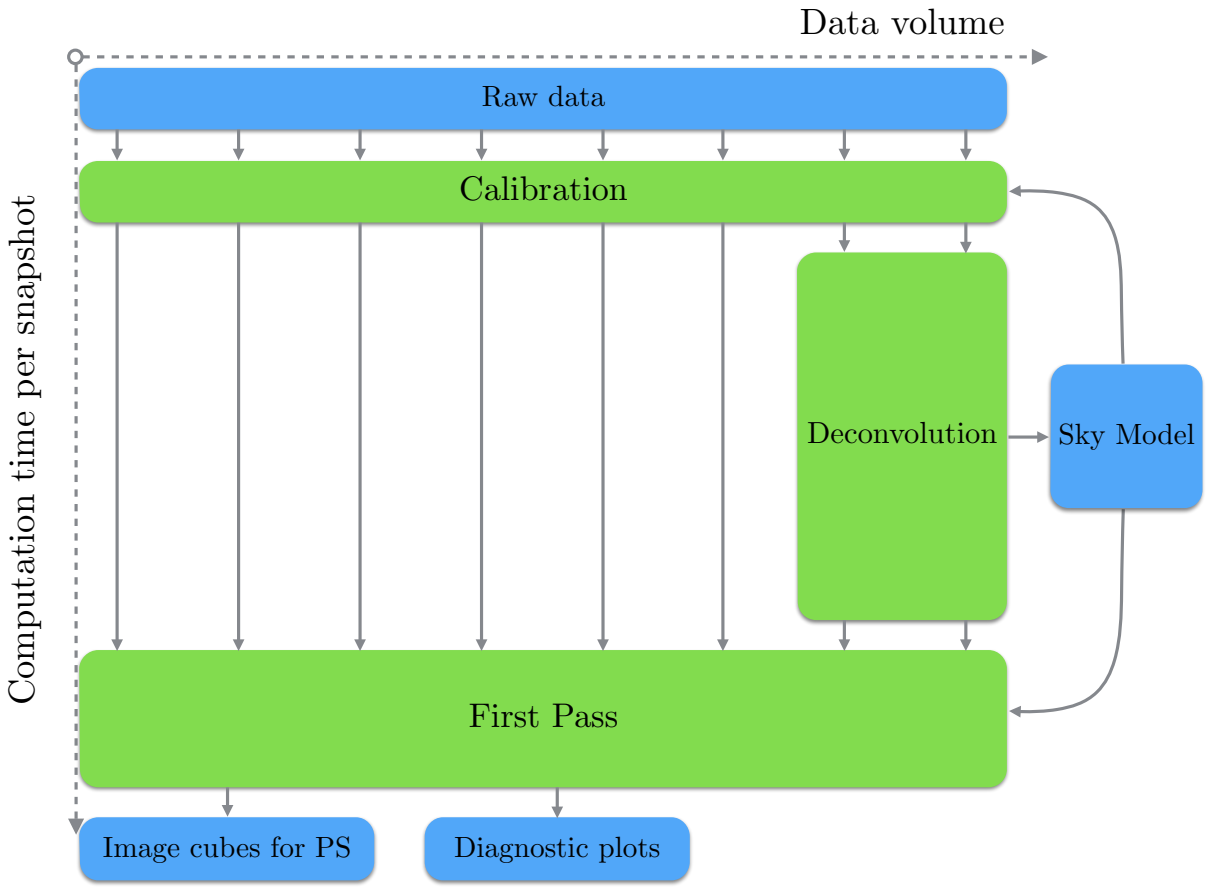

Figure 6.3: Schematic of the calibration and imaging pipeline. Blue boxes show data products, and green show computational steps. The horizontal pseudo axis illustrates the data volume used in the different steps, and the vertical axis shows the computation required per snapshot. All the EoR data is calibrated and put through First Pass. By only using a subset of data to deconvolve and form a sky model, we are able to reduce a huge amount of computational strain on our analysis pipeline. The sky model to be used is a large topic of investigation, as it both serves as an input to calibration and as a foreground model to subtract from the data. The components of our model are discussed in Chapter 7.

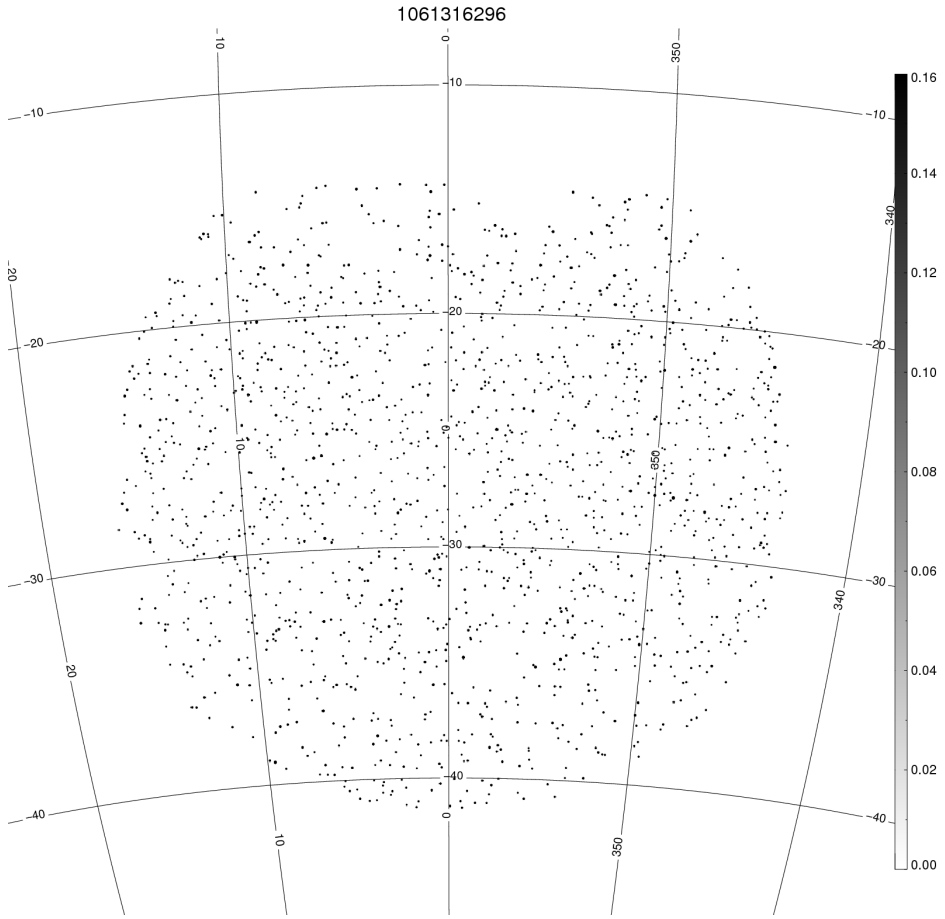

Figure 6.4: An example of sources used as the sky model for a first round of calibration and deconvolution and the EoR0 field. These sources are based on the MWACS catalog, clipped to a primary beam threshold of 5%. While the MWACS catalog covers much of the EoR0 field, the maximum declination cutoff is apparent at the top of the image. While not ideal, this selection of sources was plenty sufficient as an initial model for calibration. This particular model was used for a zenith-pointed snapshot.

described in the next subsection. We then used the calibrated visibilities from the golden data set to run through the Fast Holographic Deconvolution package in its originally intended mode to deconvolve sources and further refine our sky model. The details of this process and the resulting catalog will be presented in a series of papers by P. A. Carroll, 2015 (in preparation). The process FHD uses to create a sky model is similar to the CLEAN algorithm in fitting for sources and subtracting the PSF convolved response, but it has several benefits beyond this more traditional method. FHD is extremely fast because it makes use of the so-called holographic mapping function, which essentially maps every input uv pixel to gridded uv pixel, accounting for the instrument convolution and gridding steps in one large but sparse matrix. It also centroids sources when fitting, and thus allows for floating point source locations which yields angular position accuracy far beyond the nominal resolution of the instrument.

Example images produced by this process are shown in Figures 6.5 & 6.6. The left panel of Figure 6.5 shows a dirty image which is used as input the the deconvolution step. FHD identifies peaks in the image and builds a model of the point sources. The resulting model is shown in the right panel of Figure 6.5, after the instrument convolved sources have been subtracted out and the model has been added back in without the PSF sidelobes. This can be thought of as our best estimate of the true sky, after we have removed instrumental effects. Figure 6.6 then shows FHD operating in firstpass mode, where the model built through deconvolution serves as an input and is directly subtracted from the visibilities without fitting beyond calibration. The residual image is mostly devoid of point sources, with a handful of mis-subtractions. These leftover contaminates will continue to be reduced as we improve the model and the calibration. In just two minutes of integration, the continuum image (averaged across the full 30.72 MWA bandwidth) is signal dominated with a signal to noise ratio (SNR) of about 10 in every pixel. This is why it is necessary to leverage the symmetries involved with foregrounds versus cosmological signal to perform the measurement in the EoR window.

We continue to iterate and improve on the sky model. In Section 7.2 I show the effect of

an improved sky model on the power spectrum measurement.

6.2.2 Calibration

With a sky model in hand we can calibrate the EoR data. The purpose of this step is to account for the instrument's response to the sky intensity. This response includes the actual coupling of the dipoles to the sky as well as any amplifications, attenuations, or delays in the signal chain. We accomplish this within FHD in two steps – first by matching the raw visibilities to the model visibilities with minimal assumptions about the response of the instrument, then by enforcing restrictions on the structure of the gain solutions motivated by physical effects. The second step is necessary to increase our signal to noise on our calibration solutions, and to avoid overfitting and removing cosmological signal.

We will allow our complex gains to account for any direction-independent response on a per antenna, per frequency, per polarization basis. To define our gains in this way, we have made two simplifications. First we push all direction-dependence of the antenna response into the model of the primary beam. FHD is capable of using separate models for each antenna, though this has not yet been implemented for the MWA due to a lack of individual models. The second simplification is to ignore any cross terms between our different axes, for example a mutual coupling between two antennas. In principle these terms would introduce *baseline* dependent gains, rather than antenna dependent. Luckily the MWA has not yet seen evidence of these terms, and so we currently neglect them in our solutions.

Under these assumptions we can model the complex gains as a multiplicative term applied to the electric field signal of each antenna.

$$E_i(\nu, t) \rightarrow g_i(\nu, t) E_i(\nu, t) \tag{6.1}$$

We will treat each instrumental polarization separately, and allow the subscript, i, to run over both X and Y. Because the fundamental measurement of an interferometer is the visibility, we must work with that.

$$V_{ij}(\nu, \tau) \equiv \left\langle E_i(\nu, t) E_j^*(\nu, t) \right\rangle_t \rightarrow g_i(\nu, \tau) g_j^*(\nu, \tau) V_{ij}(\nu, \tau) \tag{6.2}$$

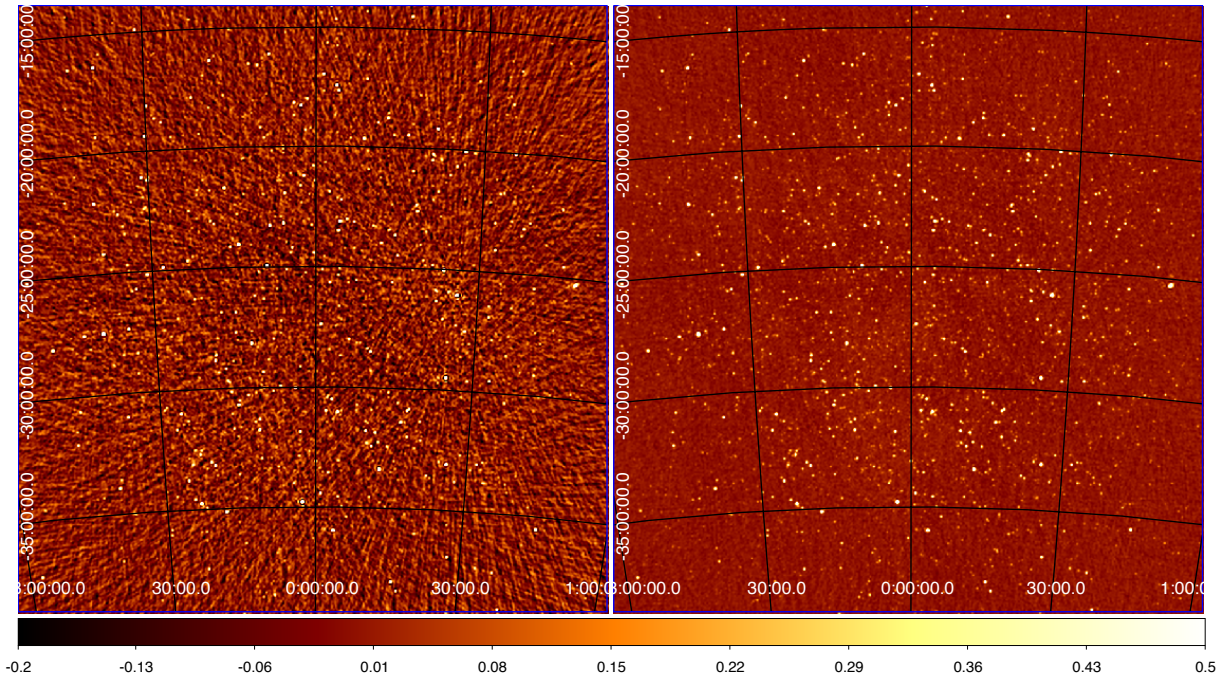

Figure 6.5: *Left:* A dirty image of the EoR0 field made by the MWA. Right ascension (RA) is on the horizontal axis, declination is on the vertical, and the color scale is in Jy. The field is dominated by bright point sources (white dots) and their PSF convolved sidelobes (ripples throughout the image). FHD fits for these sources and subtracts model visibilities from the raw data to create residual visibilities. *Right:* Restored image of the same field as on the left. The restored image is created by first subtracting an instrument convolved model of the sky, then adding the model sources back into the image without the PSF sidelobes. This is our best estimate of what the sky actually looks like, where the points are extragalactic radio sources.

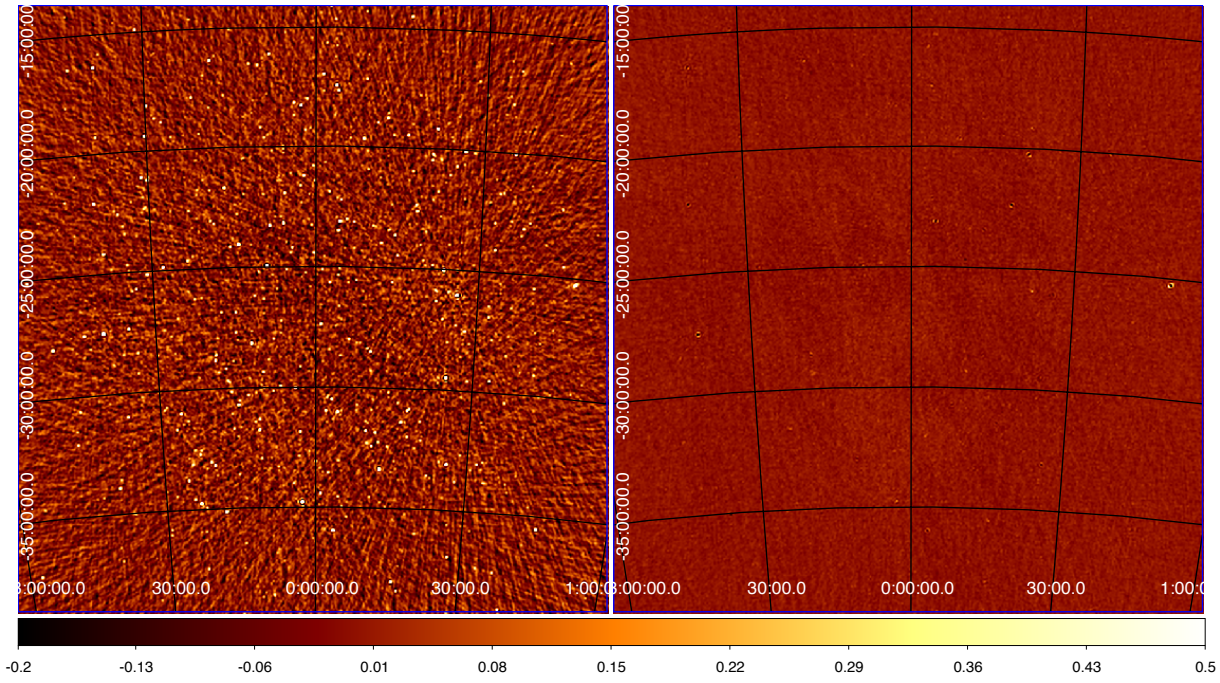

Figure 6.6: *Left:* The same dirty image as shown in Figure 6.5, shown again for reference. *Right:* Residual image created by subtracting instrument convolved model visibilities from the data. This residual image was made in "firstpass" mode, where a model is the input to FHD and is subtracted from the data, with no fitting beyond the calibration step. We see that the vast majority of the foreground power has been removed, though small mis-subtractions still exist. Continuing to improve the model and calibration will reduce these remaining contaminates. These images were made from the full 30.72 MHz MWA bandwidth and are very much signal dominated in every pixel in just two minutes of integration (SNR ~ 10). This is why the EoR window is key to detecting the faint cosmological signal behind these bright foregrounds.

Our aim is to solve for the gains and remove them to recover the calibrated visibilities. Once again we will be working on a snapshot basis, and will ignore any time dependence of the visibilities or the gains. Because we solve for the gains independently for each snapshot, we are automatically allowing for gain variation on timescales greater than 2 minutes.

We proceed to the first step in our calibration process, which is to use model visibilities to solve for the gains, treating all antennas, frequencies, and polarizations independently. Traditionally, radio interferometers use dedicated calibration observations which point at bright known sources to calibrate. This allows the observer to approximate the sky as a single point source, making the model visibilities easy to calculate. However, the MWA has a very large field of view making it difficult to point at single point sources. Instead we realized early on that the high signal to noise in even the relatively quiet EoR fields was sufficient for calibration given a good model of the entire field. So we opt to model the foregrounds on the actual field where we wish to do our measurement. This makes the modeling more difficult but has the benefit of being able to use the same model for foreground subtraction, and allows us to calibrate on every single snapshot rather than repointing the telescope every so often.

We can express the measured, uncalibrated visibilities as

$$V'_{ij}(\nu) \approx g_i(\nu) g_j^*(\nu) V_{ij}(\nu) \tag{6.3}$$

where we return to the prime notation for measured values, and let the unprimed V_{ij} denote the model visibilities (with assumed unity gains). With N_a antennas and N_ν frequency channels, we have $N_\nu \times N_a$ unknowns and $N_\nu \times N_a(N_a - 1)/2$ independent equations, so this in principle should be solvable. However, there is approximately zero chance our model visibilities are perfect so this problem is over constrained. We instead turn to a goodness of fit, χ^2.

$$\chi^2 = \sum_{ij} \left| V'_{ij}(\nu) - g_i(\nu) g_j^*(\nu) V_{ij}(\nu) \right|^2 \tag{6.4}$$

Now our job is to minimize χ^2 with respect to the gains.

FHD utilizes the StEFCal algorithm described in Salvini and Wijnholds [77] to minimize χ^2. This algorithm has been shown to operate with a numerical complexity scaling as $\mathcal{O}(P^2)$ where P is the number of receiver paths ($N_\nu \times N_a$ in our case). This is the optimal scaling assuming use of all available visibility data. In addition the convergence properties of the algorithm have been studied and it has been shown to operate at or near the Cramer-Rao Bound, the theoretical limit to parameter estimation.

The result is estimated gains for every antenna, frequency channel, and polarization. However, with certain known properties of the antenna responses we can reduce the number of free parameters in our solution and avoid fitting out true cosmological signal. This takes us to the second step of calibration.

We next aim to reduce the number of free parameters in the calibration. This will both increase the signal to noise on the solutions, and prevent us from fitting out any cosmological signal as long as we make wise choices about our model. Currently we have solved for $2 \times N_\nu \times N_a \times N_{\text{pol}}$ parameters, where the two is for the real and imaginary parts of the gain. For the MWA this amounts to about 196,608 parameters.

We model our gains as an amplitude bandpass common to all antennas with a given beamformer to receiver cable length, $B_\alpha(\nu)$, and an antenna dependent low order polynomial in frequency. In addition we have found it necessary to solve for a cable reflection term, which will be discussed in detail in Section 7.1.1. We can then express our restricted gain as

$$\hat{g}_i(\nu) = B_\alpha(\nu)(P_i(\nu) + R_i(\nu)), \tag{6.5}$$

where the polynomial and reflection mode can be further decomposed as

$$P_i(\nu) = (A_{i,0} + \nu A_{i,1} + \nu^2 A_{i,2}) e^{i(\phi_{i,0} + \nu\phi_{i,1})} \tag{6.6a}$$

$$R_i(\nu) = R_{i,0} e^{-2\pi i \tau_i \nu}. \tag{6.6b}$$

The coefficients $A_{i,n}$, $\phi_{i,n}$, and τ_i are real, while R_0 is allowed to be complex. The parameter τ_i is the time required for the signal to traverse the cable twice.

We first solve for the bandpass term, $B_\alpha(\nu)$. The motivation for this term is that all the antennas share a broad frequency-dependent response due to the physical dipoles and

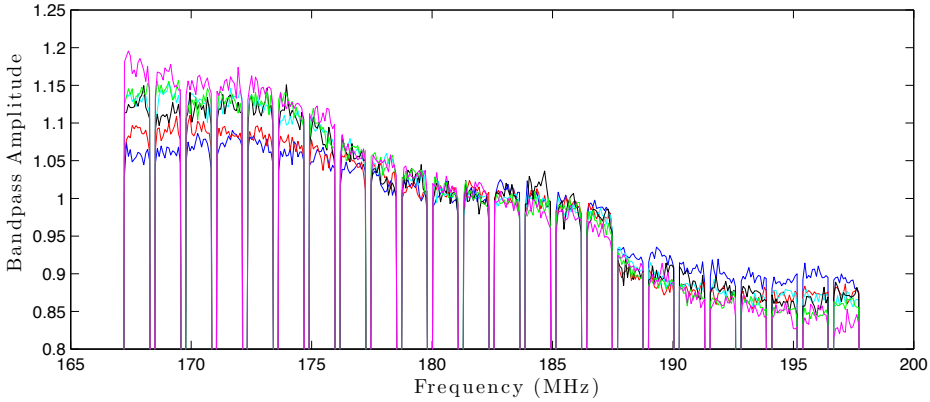

Figure 6.7: The per-cable length bandpasses, $B_\alpha(\nu)$, for a single zenith snapshot on August 23, 2013. The colors correspond to 90 (blue), 150 (red), 230 (cyan), 320 (black), 400 (green), and 524 (magenta) meter cable lengths. Only the X polarization is shown for clarity. The broad slope is due to the response of the dipoles that make up the antenna. The sharp periodic dips are due to the masked channels from known aliasing. The sudden jump around 187 MHz is due to a digital gain discontinuity in the receiver chain. These bandpasses are used to smooth all antenna gains before fitting for a polynomial.

various filters on the signal chain. Included in this is the polyphase filter applied in the receivers when creating the coarse channels. This filter has a very characteristic shape with a sharp drop-off at the edges, which results in a periodic structure in the gains. We group tiles into their cable lengths because the attenuation in the cable, and the different types of cables used for different lengths, can introduce differences in the bandpass. The bandpass is solved for using the amplitudes of all the unrestricted gains in the first step of calibration. For each frequency, polarization, and cable length, we use the median of the amplitudes of these gains as the bandpass value. An example result of this process is shown in Figure 6.7.

Next we solve for the polynomial term of the restricted gains. This is done by first dividing the unrestricted gains by the bandpass, then for each tile fitting a low order polynomial. We have found that a quadratic amplitude and a linear phase fit is sufficient to capture most

response structure in MWA antennas.

Finally, we find the reflection mode. This is done for antennas with suspected reflections in their cables between the beamformers and the receivers. The most offensive reflection was seen to exist in those antennas whose cables are 150 meters[4]. A reflected signal appears in the response of the instrument as a delay, which when Fourier transformed to frequency space is a single oscillating mode. We can solve for the term by simply performing a direct Fourier transform to the mode of interest, however we found that the mode of interest is not exactly known due to uncertainty in cable length (plus or minus several centimeters), and in the velocity of the signal in the cable. As a result, we must also solve for the reflection time, τ_i, itself. This is currently done by Fourier transforming to a highly over resolved frequency grid and selecting the mode where the reflection amplitude is largest.

The results of an example calibration are shown in Figures 6.8 – 6.10. First is the calibration amplitude in Figure 6.8. This is the bandpass term, multiplied by the amplitude component of the polynomial fit, and modulated by the reflection mode where appropriate. We can see the general shape of the bandpass common to all antennas, with minor differences due to the offset, slope, and quadratic terms of the polynomial fit. The effect of the reflection mode is minuscule on these plots.

Figure 6.9 shows the phase of the gains. By definition, these are linear with the exception of the tiny reflection mode. Most of the tiles are very flat in phase, with the exception of tile 138. The calibration algorithm flagged tile 138 due to its large slope in phase, though it's not clear this is actually justified – large slope could be purely due to a larger delay in the signal chain than expected. However, for the time being it is safer to flag suspicious data until we can determine the actual cause.

Figure 6.10 shows the residual gains when subtracting the restricted gains from unrestricted, $g_i(\nu) - \hat{g}_i(\nu)$. We can see most of the tiles have largely noise-like residuals. This

[4]For logistical purposes the beamformer to receiver cables were purchased in six set lengths – 90, 150, 230, 320, 400, and 524 meters. Because of the cable length quantization, reflections appeared in power spectra early in analysis. Perhaps this was actually a feature of the instrument, because it is better to find these systematics early than when we try to claim a detection.

Figure 6.8: Calibration amplitudes for single zenith snapshot from the golden data set. We can see the common bandpass in all the tiles, with minor differences from the polynomial fit. Blue lines correspond to the XX polarization, and red are YY. Tile 105 was flagged in the preprocessing step, and so was never calibrated. Tiles 11, 138, 142, and 151 were flagged due to poor calibration (see Figures 6.9 & 6.10).

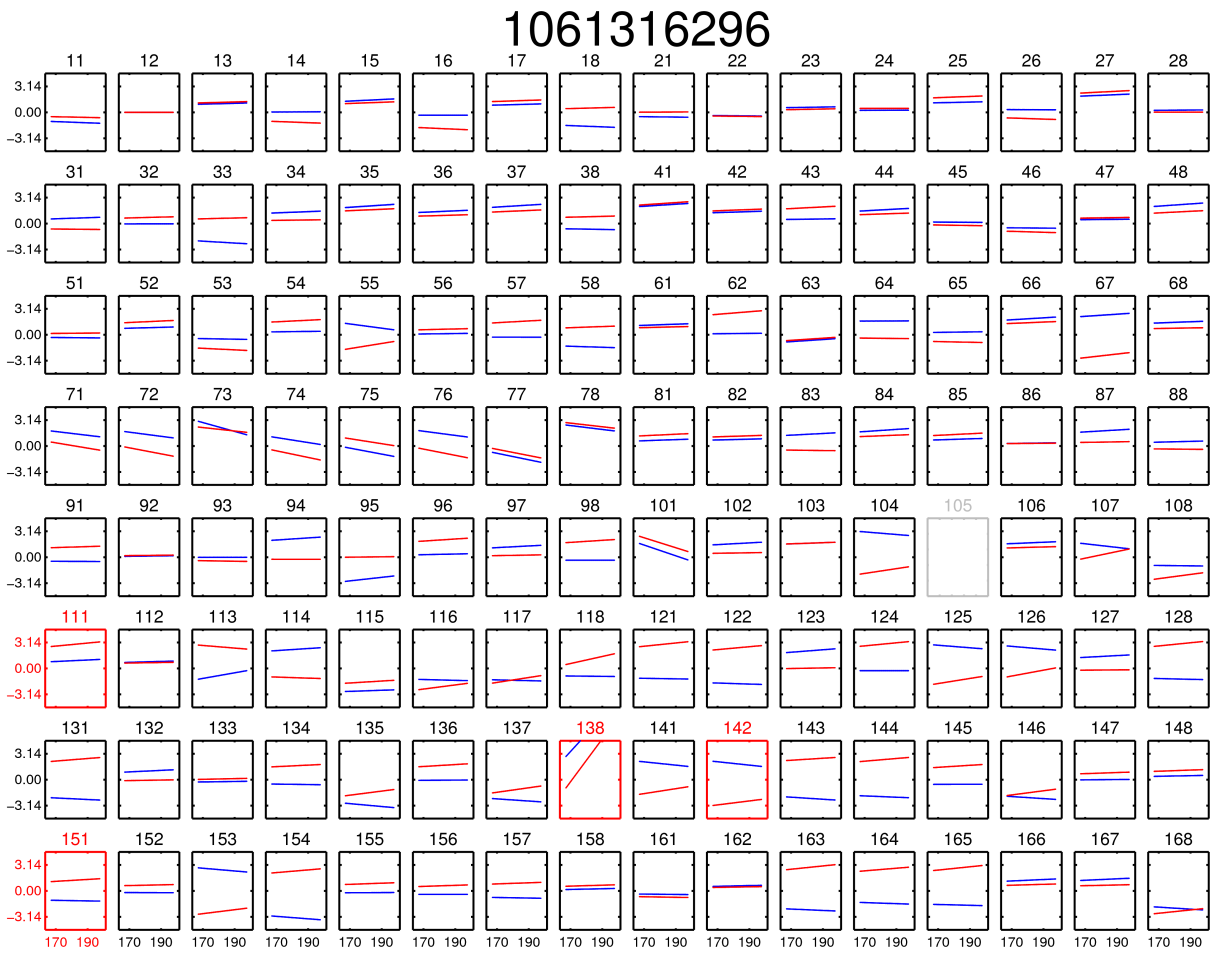

Figure 6.9: Calibration phases for single zenith snapshot from the golden data set. The slopes are generally flat with the notable exception of tile 138, which was flagged for its large phase slope. When solving for calibration gains, the absolute phase is a completely free parameter because it falls out when forming visibilities. We therefore have the freedom to choose a reference antenna and set its phase to zero for all frequencies and polarizations. For historical reasons we use tile 12 for this purpose.

reassures our strategy of fitting for fewer parameters to increase signal to noise. Tiles 111, 142, and 151 were flagged due to large residuals, indicating a potentially poor fit. Some larger structure can be seen, for example the large waves in tile 45. This is likely due to power in the actual visibilities that we did not model, or modeled poorly. Another benefit of the parameter fit is to smooth out this structure that artificially arose in the unrestricted gain step due to an imperfect model.

We can now enumerate the number of parameters solved for with the restricted gain model. For each polarization we have: a bandpass parameter at each frequency for each of six cable lengths, five polynomial coefficients for each antenna, and three reflection parameters for antennas with suspected reflections. The total number of free parameters is then $N_{\rm pol} \times (6N_\nu + 5N_a + 3N_r)$, where N_r is the number of antennas with a cable reflection term. For the MWA this is approximately 6,074 free parameters – a significant reduction from the unrestricted gains.

The final step of calibration is to apply the gain solutions to the visibility data. We simply divide all visibilities by the associated gains.

$$V'_{ij}(\nu) \rightarrow \frac{V'_{ij}(\nu)}{\hat{g}_i(\nu)\hat{g}^*_j(\nu)} \approx V_{ij} \tag{6.7}$$

This entire calibration process is done for every snapshot of EoR data, providing both a two-minute resolution time dependence of calibration solutions, as well as model visibilities which are used for foreground subtraction and diagnostic purposes.

6.2.3 Image Cubes and Other Data Products

The next step in the pipeline is to form images. This is again done on a snapshot cadence. We first grid the visibilities using the primary beam as the gridding kernel.

$$\widetilde{I}'(\mathbf{u},\nu) = \sum_{ij} \widetilde{A}^*(\mathbf{u}-\mathbf{u}_{ij},\nu)V'_{ij} \tag{6.8}$$

At this stage we also average in frequency by a factor of two by gridding pairs of frequency channels to the same uv plane. This results in a frequency resolution of 160 kHz. The choice

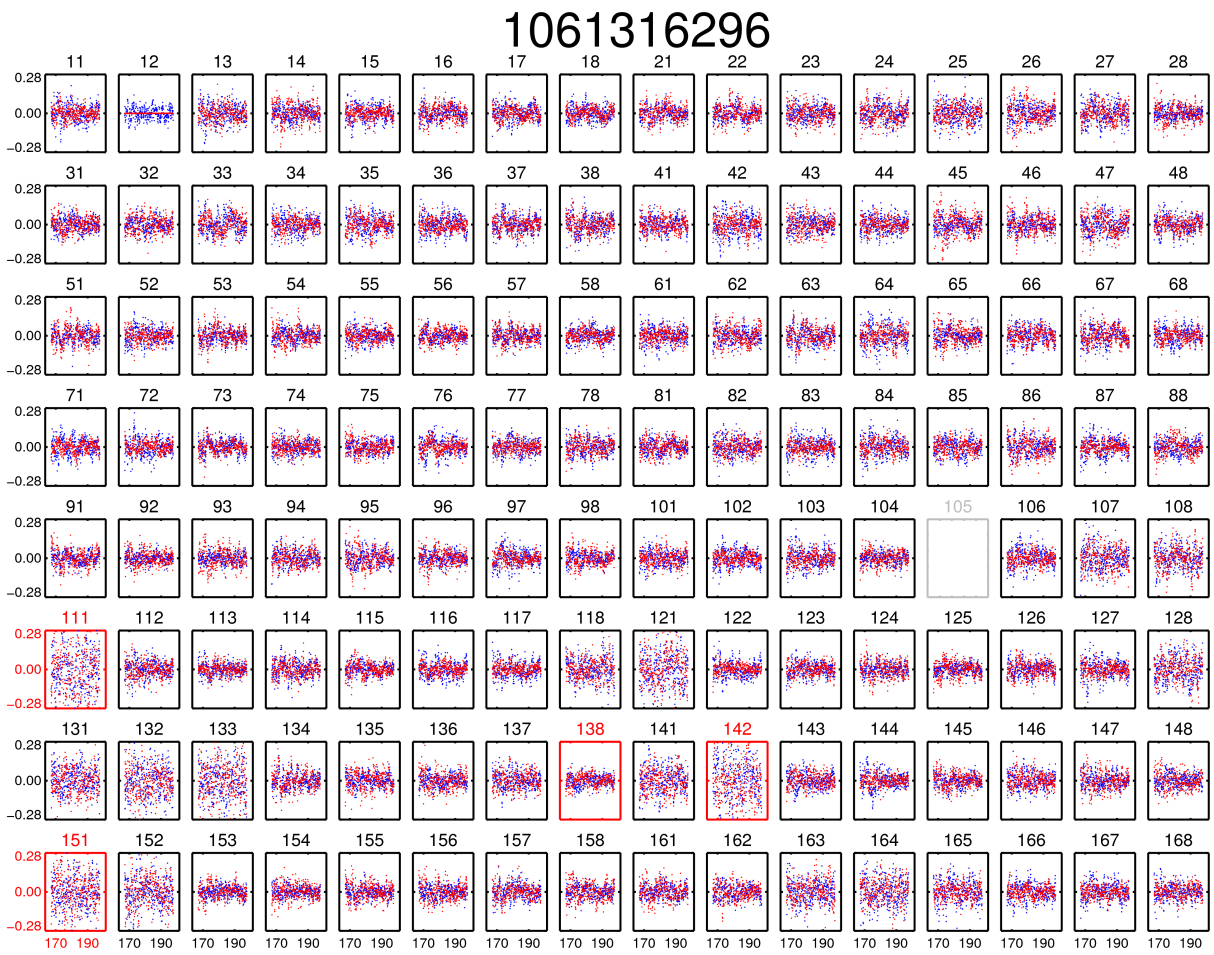

Figure 6.10: Calibration residuals for single zenith snapshot from the golden data set, for the XX polarization. Unlike the previous figures, this entire figure is the XX polarization which was chosen for clarity. The blue lines are the real part of the residual, and the red lines are the imaginary part. Tile 12 has zero imaginary residual because it is chosen to be completely real. Tiles 111, 142, and 151 are flagged due to their high level of residual gain.

of using the beam as the gridding kernel puts us in the holographic frame, which has been shown to be the optimal frame for combining images in the sense that it preserves all information for parameter estimation [56, 8]. This frame is useful not only for gridding visibilities within a snapshot, but is also optimal for combining snapshots for long integrations.

However, as alluded to earlier, the array is constantly rotating and so the uv plane is not stationary. One can account for this in multiple ways. The W-projection algorithm was developed to handle non-coplanar baselines by accounting for fresnel diffraction effects when propagating the electric field waves to a common uv plane [20]. Alternatively we can bypass the unfortunate non-coplanarity of baselines from different snapshots by first Fourier transforming to the sky frame and combining observations there. Each technique has its trade-offs, but as FHD is an imager at heart we opted for the imaging method.

We form images by inverse Fourier transforming our estimate of the coherence function.

$$I'(\boldsymbol{\theta}, \nu) = \mathcal{FT}^{-1}\left[\widetilde{I}'(\mathbf{u}, \nu)\right] \tag{6.9}$$

While the sky frame is appealing, we are not quite in the clear. Because each snapshot's uv plane is different, the corresponding Fourier domains ($\boldsymbol{\theta}$) differ (see Appendix A). We thus need to interpolate to a common set of sky pixels before combining images. We do this using the Hierarchical Equal Area isoLatitude Pixelization (HEALPix) framework [29]. HEALPix provides a set of equal-area pixels on the curved sky which will remain the same for each snapshot. The resolution of the HEALPix coordinates is quantized by the parameter N_{side}, and the total number of pixels on the sky is given by $N_{\text{pix}} = 12 \times N_{\text{side}}^2$. We choose a resolution comparable to the MWA instrument resolution, which yields $N_{\text{side}} = 1024$, and a resolution of about 3.4 arcmin.

Early on we realized that choices of the uv gridding plane affected the quality of the images and thus the power spectra. First, when making images we found that if our gridding resolution was too coarse (imaging small field of view), foregrounds from sidelobes would alias into the primary field of view. We solved this problem by imaging out 90° from phase center (gridding resolution of a half wavelength), and cropping the image. We also found that

if we select the cropping region based on beam value, we get hard edges in the integrated images from different snapshots having beams pointed in slightly different directions. To solve this we select a predetermined set of HEALPix pixels to interpolate to, and use the same set for all snapshots on a given field. The final cropped field of view for this work is a 21° square centered at RA $= 0$ h, dec $= -27^\circ$.

The imaging step is done for each pair of frequency channels in the bandwidth, resulting in a three dimensional cube. A slice of a snapshot cube is shown in Figure 6.11. This image is still in the holographic frame so that it can be added with other snapshots, so it is tapered by two factors of the primary beam. It is also naturally weighted (uv modes weighted by the number of baselines sampling them) to retain the full sensitivity of the instrument. These factors are accounted for in the power spectrum segment of the pipeline (Section 6.3).

While the primary output of our calibration and imaging step is image cubes, there are actually several other data products that are necessary for power spectrum estimation and diagnostic purposes. I will list these products here.

Diagnostic plots. The FHD package outputs many plots that can be used for quality assurance. A few of these have already been shown here (Figures 6.8–6.10). Also available are continuum images in several polarizations, images of the beams, and images of the uv weights.

Calibrated visibilities. As mentioned earlier, the CHIPS pipeline requires calibrated visibilities for power spectrum estimation, and data-crossovers are a piece of the international collaboration strategy. FHD thus exports the calibrated visibilities, as well as model visibilities for this purpose.

Weights cubes. The holographic frame in some sense is a weighted frame where modes with more measurements are weighted higher. When we actually form power spectra, we will need to divide by these weights to return to a sensible mean. Looking back to our gridding equation (Equation 6.8), we can see that the appropriate weights are simply the gridded

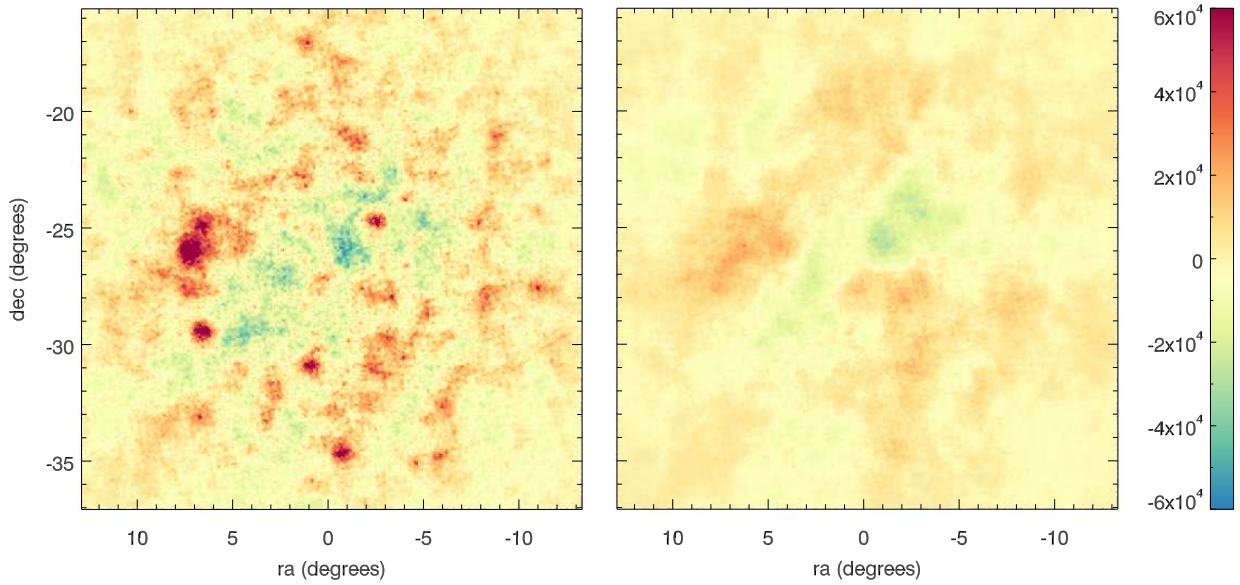

Figure 6.11: Sample dirty and residual image slices taken from a single zenith snapshot from the golden data set. These slices are at 182 MHz and width 160 kHz. On the left is the dirty image contaminated with lots of foregrounds. The right panel shows the residual image after foregrounds have been subtracted. Much of the structure from the left slice is gone on the right, and indication of a good foreground model and good calibration. These images are in the holographic frame, and so are weighted by two factors of the beam. They are also naturally weighted which emphasizes the modes which MWA is most sensitive to (large scales). Thus the units are not quite in sensible temperature units, and should only be viewed for relative scales.

beams.

$$W(\mathbf{u},\nu) = \sum_{ij} \widetilde{A}^*(\mathbf{u} - \mathbf{u}_{ij}) \tag{6.10}$$

These weights should be added across multiple snapshots in the way that the visibilities are added. Thus we perform the same imaging and interpolation to HEALPix steps as we did with the data to integrate the weights. This provides us with weights cubes which we will use within εppsilon.

Variance cubes. In the same way that we want the weights for each snapshot, the variance in each pixel is also necessary to propagate errors through our pipeline. Under the assumption that each visibility within a given frequency channel has the same level of noise, $V_{\rm rms}(\nu)$, we can write down the variance in each $uv\nu$ pixel, again using the gridding equation.

$$\left(\widetilde{I}'_{\rm rms}(\mathbf{u},\nu)\right)^2 = \left(V_{\rm rms}(\nu)\right)^2 \sum_{ij} \left|\widetilde{A}(\mathbf{u} - \mathbf{u}_{ij},\nu)\right|^2 \tag{6.11}$$

We note that although it would be nice if this was simply the square of the weights, because the square is before the sum, we are not so lucky. Instead, we must compute this cube separately and carry it through the analysis.

We measure $V_{\rm rms}(\nu)$ directly from the calibrated visibilities. This is done for each frequency channel by first separating the data into interleaved time stamps (at the two second resolution), labeling the two groups as even and odd time samples. We subtract the even times from the odd to remove any signal in the visibilities, then calculate the noise across all visibilities for a given frequency channel. The result of this process for an example snapshot is shown in Figure 6.12.

Even/odd interleaved cubes. Each of the cubes described so far (dirty, model, weights, variance) are exported as both even and odd time samples. These are useful as a diagnostic tool for checking systematics, but also fundamental to the εppsilon pipeline. Summing even and odd data cubes is equivalent to integrating the full snapshot, while differencing them removes the signal leaving only noise. This difference cube is used to check the error propagation (from the variance cube) throughout the pipeline, and to subtract the noise bias off the power spectrum in εppsilon.

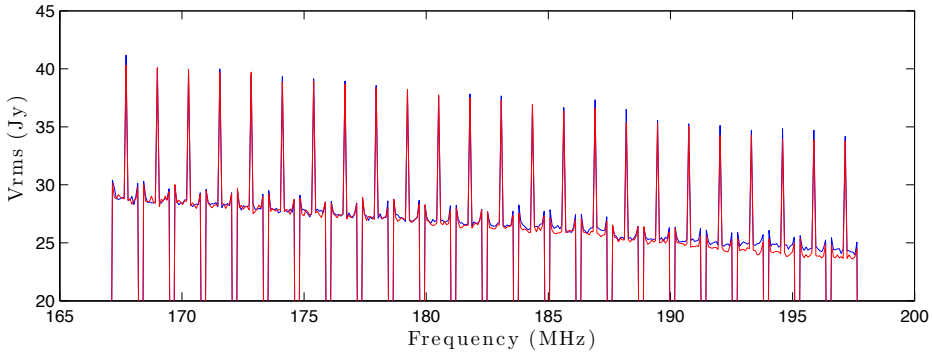

Figure 6.12: The noise spectrum for the MWA for 2 second integrated visibilities. A full two minutes snapshot is used to the calculate the noise level, but the values displayed are for the two second cadence read in from UVFITS files. We can see periodic dips due to the missing data at coarse channel edges. The spikes in the center of each coarse channel are due to the fact that the center 40 kHz channel is flagged, but then averaged with the adjacent channel. This results in half as much data in this channel than the rest, and thus more noise (a factor of $\sqrt{2}$ higher).

Polarizations. All the above cubes are exported in instrumental polarizations XX and YY. In principle FHD is capable of a full polarization output (including XY and YX) which allows for a full Stokes reconstruction, however this is omitted for the present analysis. Instead we only retain the XX and YY which is sufficient to produce Stokes I and Q.

6.3 εppsilon

The final step in our analysis pipeline is to integrate our image cubes and form power spectra using the Error Propagated Power Spectrum with Interleaved Observed Noise (εppsilon) package[5].

6.3.1 Integration

The preceding steps were carefully honed in order to make the integration step as simple as possible. The image cubes are already in the holographic frame with predetermined pixels, and we have propagated weights and variance cubes in such a way that we simply add cubes together and divide by the weights later.

While the integration step is extremely simple by design, it is significant in the pipeline because it marks our last opportunity to exclude data at the snapshot ($\sim$ 2 minute) time scale. This step has been crucial when considering various data cuts and jackknife tests. Computationally, the majority of analysis time is spent calibrating and imaging the data as it is done on a snapshot basis. The integration step also scales with the number of snapshots, but only requires about a minute per snapshot to read in the images (which can then be spread across a compute cluster). The power spectrum step is then completely agnostic to the integration time from a computational point of view, and is relatively fast ($\sim$ 3 hours) compared to the previous steps. Therefore, not only does the integration step mark the last place in the analysis for data cuts, but it is also a convenient marker as a reasonable place to repeat analysis with minor changes – easily allowing for various jackknifes.

[5] `https://github.com/miguelfmorales/PS`

6.3.2 Power Spectrum and Error Propagation

After integration, the final step is to form a power spectrum from the integrated cubes. Because the data are in HEALPix format in the angular dimensions, and not on a regular grid, we cannot use the Fast Fourier Transform, but instead must use the generic Discrete Fourier Transform (DFT). However, we must ensure that the $k_\perp$ space we transform to is a match to the region of sky where we have data. The $k_\perp$ resolution is given by the field of view of our selected sky,

$$\Delta k_\perp = \frac{2\pi}{\Delta D_\perp}, \tag{6.12}$$

where $\Delta D_\perp$ is the perpendicular span of our observation in cosmological distance units (Mpc). Likewise, the $k_\perp$ extent is determined by the perpendicular resolution in image space, though this is less important for our purposes because our sensitivity is much higher at low $k_\perp$. We also need to match the *shape* of the $k_\perp$ space to the image space. The two dimensional Fourier transform only produces an orthogonal space of modes when the shape of the spaces are square[6]. This fact influenced our choice of HEALPix pixels when making our images. The choice was made so that the region is roughly square, has a field of view that produces a $k_\perp$ resolution the size of the coherence of our visibilities, and a resolution approximately equal to that of the instrument.

After the DFT we are in the (k_x, k_y, z) hybrid Fourier space. Because the images are necessarily real, the full $k_\perp$ plane is redundant over a 180° rotation. We therefore cut half of the plane to ensure we are not double counting information.

Before we perform the last Fourier transform along the line of sight, we divide by the weights. This has the effect of removing one factor of the primary beam, and removing the natural weighting which emphasizes shorter baselines. In other words, we are completing the weighted average after first performing the weighted sum. This is the appropriate frame in which to divide by the weights because it is the frame where the weights were first applied

[6]One could use a different shape region of sky, but would then need a corresponding transform. For example a circular region would use a Bessel decomposition as was used in Chapter 3. For simplicity we opted to use the Fourier transform and square image and $k_\perp$ spaces.

(in the gridding step). The units of the data at this stage are mK Mpc^2. We must also divide the variance by the weights squared to propagate the error properly.

In order to achieve a large dynamic range in the line of sight transform, we apply a Blackman-Harris window function along the frequency[7]. This has the effect of reducing sidelobes in the PSF of the transform, and thus helping us separate the foregrounds from the cosmological signal. As a consequence it also reduces our effective bandwidth by a factor of two.

The line of sight transform can be complicated by missing data in the frequency dimension (due to RFI flagging, or masked out channels due to known aliasing). We have adopted the Lomb-Scargle least-squares method of best determining the power in a given frequency mode [79]. This technique accounts for the fact that the sampled space may not be equally sensitive to sine and cosine modes, but instead determines the truly orthogonal eigenfunctions in the space.

The next step in the pipeline is to square and form a power spectrum estimate. Up until this stage we have carried around two interleaved cubes – the even time steps and the odd time steps. Now is the time to examine how to use these cubes to both produce an unbiased power spectrum estimate, and measure the noise in the data. For this section we will be working in purely **k**-space, and so I will omit it from the notation where possible. The result of the Lomb-Scargle process is two modes for each **k** pixel (analogous to sine and cosine). However, we can treat them completely independently in the math, so I will leave them separate until we are ready to average bins together. We start by forming sum and subtraction cubes.

$$\widetilde{I}'_{\rm sum} = \frac{\widetilde{I}'_e/\sigma_e^2 + \widetilde{I}'_o\sigma_o^2}{1/\sigma_e^2 + 1/\sigma_o^2}\ ; \quad \widetilde{I}'_{\rm sub} = \frac{\widetilde{I}'_e/\sigma_e^2 - \widetilde{I}'_o\sigma_o^2}{1/\sigma_e^2 + 1/\sigma_o^2} \tag{6.13}$$

I have included weights (now the inverse variance) when forming the sum and subtraction, though we have found that having unequal weights leaves residual signal in the subtraction data cube. We therefore enforce flagging such that the data in each interleaved cube is

[7]Wikipedia actually has a great discussion on various spectral windows and their purposes in signal processing. `http://en.wikipedia.org/wiki/Window_function`

exactly symmetric and so $\sigma_e^2 \to \sigma_o^2$. We will still reintroduce the these weights when it comes time to average pixels. The variances in the summed and subtracted cubes are

$$\sigma^2_{\text{sum}} = \sigma^2_{\text{sub}} = \frac{\sigma_e^2}{2}. \tag{6.14}$$

Down the line it will be beneficial to have the probability distribution function (PDF) for these cubes to keep track of expected values and variances. We make the assumption that noise is a random realization of a symmetric gaussian distribution. The math gets a little messy, so I will work with an arbitrary complex gaussian number, c, with expected value μ, and variance σ^2. The PDF of c is

$$\Pr(c) = \frac{1}{\pi\sigma^2} e^{-|c-\mu|^2/\sigma^2}. \tag{6.15}$$

Next we find the PDF of $|c|$ which is calculated by integrating a ring of constant magnitude.

$$\Pr(|c|) = \int_0^{2\pi} \mathrm{d}\phi \Pr(c)|c| = \frac{2|c|}{\sigma^2} e^{-(|c|^2+|\mu|^2)/\sigma^2} I_0\left(\frac{2|c||\mu|}{\sigma^2}\right) \tag{6.16}$$

Here I_0 is the 0^{th} modified Bessel function. A change of variables takes us to the PDF for $|c|^2$.

$$\Pr(|c|^2) = \frac{\mathrm{d}|c|}{\mathrm{d}|c|^2}\Pr(|c|) = \frac{1}{\sigma^2} e^{-(|c|^2+|\mu|^2)/\sigma^2} I_0\left(\frac{2|c||\mu|}{\sigma^2}\right) \tag{6.17}$$

The stats we care most about at this stage are the expectation value and variance of $|c|^2$.

$$\left\langle |c|^2 \right\rangle = |\mu|^2 + \sigma^2 \tag{6.18a}$$

$$\mathrm{var}(|c|^2) = \sigma^2(2|\mu|^2 + \sigma^2) \tag{6.18b}$$

We can now return to the sum and difference cubes and determine the expected values and variances. We take S to be the signal in our given $\mathbf{k}$ pixel, but note that this is not yet the estimate of the power – we have yet to divide by a window function as shown in Equation 2.13.

$$\left\langle |\widetilde{I}'_{\text{sum}}|^2 \right\rangle = |S|^2 + \sigma^2_{\text{sum}} \tag{6.19a}$$

$$\left\langle |\widetilde{I}'_{\text{sub}}|^2 \right\rangle = \sigma^2_{\text{sub}} \tag{6.19b}$$

$$\mathrm{var}\left(|\widetilde{I}'_{\mathrm{sum}}|^2\right) = \sigma^2_{\mathrm{sum}}(2|S|^2 + \sigma^2_{\mathrm{sum}}) \tag{6.20a}$$

$$\mathrm{var}\left(|\widetilde{I}'_{\mathrm{sub}}|^2\right) = \sigma^4_{\mathrm{sub}} \tag{6.20b}$$

Because our sum term has a noise bias in its expected value, we must subtract the pure noise term to arrive at an estimate of the coherence function squared.[8]

$$|\widetilde{I}'|^2 = |\widetilde{I}'_{\mathrm{sum}}|^2 - |\widetilde{I}'_{\mathrm{sub}}|^2 \tag{6.21}$$

This gives us a variance of

$$\sigma^2_{\widetilde{I}'} = \mathrm{var}\left(|\widetilde{I}'_{\mathrm{sum}}|^2\right) + \mathrm{var}\left(|\widetilde{I}'_{\mathrm{sub}}|^2\right) = 2\sigma^2_{\mathrm{sum}}(|\widetilde{I}|^2 + \sigma^2_{\mathrm{sum}}). \tag{6.22}$$

The variance of our measurement actually contains the signal itself. However, in our limit of very large noise and small signal in each individual pixel, we can largely ignore that term.

At this stage we can finally form an estimate of our three dimensional power spectrum. We simply divide by the integral of the beam squared, as shown in Equation 2.13. In principle the beam integral can be performed nearly anywhere in the pipeline by using Parseval's theorem. In practice we calculate this integral within the imaging step where the highest resolution and most accurate beam model is available. We perform the integral for each snapshot, then perform a weighted average of all snapshots. The framework of FHD also allows us to adopt an average beam within a snapshot, for example if different antennas have different beams, but this has not been used yet as we have not seen evidence for varying beams.

The result is an estimate of the power spectrum at each $\mathbf{k}$ pixel, $P'(\mathbf{k})$. We also have the variance on each pixel by dividing the variance on the coherence function squared by the

[8]Often the cross power between the even and odd samples is used to remove the noise bias. It can easily be shown that for two complex number, c_1 and c_2, our method of squaring sums and differences then subtracting is related to the cross power by the relationship $|c_1 + c_2|^2 - |c_1 - c_2|^2 = 4\mathrm{Re}(c_1 c_2^*)$.

same window function, but squared.

$$\sigma_P^2(\mathbf{k}) = \frac{\sigma_{\widetilde{I}'}^2(\mathbf{k})}{\left(\frac{1}{(2\pi)^3}\int \mathrm{d}^3\mathbf{k}' \left|\widetilde{A}(\mathbf{k}-\mathbf{k}')\right|^2\right)^2} \tag{6.23}$$

Our next step is to average the measured power spectrum – usually either to a 2D or 1D power spectrum – to reduce statistical noise and sample variance. The details of which pixels to average will be discussed later, but for now we will show the error propagation through the averaging process. We use as our weights the inverse variance on each pixel. We denote a band power (any number of pixels averaged together) with a greek letter, and use latin indices to denote the individual pixels that contribute to the band power.

$$P'_\alpha = \frac{\sum P'(\mathbf{k}_i)/\sigma_P^2(\mathbf{k}_i)}{\sum 1/\sigma_P^2(\mathbf{k}_i)} \tag{6.24}$$

The variance on the averaged band powers is then given by

$$\sigma_\alpha^2 = \frac{1}{\sum 1/\sigma_P^2(\mathbf{k}_i)}. \tag{6.25}$$

One of the key features of εppsilon is that it propagates the noise through the entire pipeline, which can be useful for diagnostics. However, interpreting the differenced cube after squaring is slightly non-trivial and warrants a small discussion. If we return to our PDF of a complex gaussian, this time with zero mean (as should be the case with the difference cube), we observe that the variance of the magnitude square is simply the square of the expected value (Equation 6.18 when $\mu = 0$). However in the averaging step, the variance and expected values diverge. To save on notation, let $\lambda_i \equiv 1/\sigma_i^2$.

$$|c|_\alpha^2 = \frac{\sum \lambda_i^2 |c_i|^2}{\sum \lambda_i^2} \tag{6.26a}$$

$$\left\langle |c|_\alpha^2 \right\rangle = \frac{\sum \lambda_i^2 \left\langle |c_i|^2 \right\rangle}{\sum \lambda_i^2} = \frac{\sum \lambda_i}{\sum \lambda_i^2} \tag{6.26b}$$

$$\mathrm{var}\left(|c|_\alpha^2\right) = \frac{1}{\sum \lambda_i^2} \tag{6.26c}$$

This difference between expected noise value and the variance of the noise leads to a distinction between *measured noise* and the actual error bars on the power spectrum. This will be evident in the plots to come.

My final comment on the noise propagation is to show the actual PDF of the weighted average. The sum of exponentially distributed numbers with equal parameters λ_i yields an Erlang distributed number. The weighted sum above does not exactly yield an Erlang distribution, but if the weights are reasonably uniform (within an order of magnitude), we can well approximate the resulting distribution as Erlang with effective parameters. The PDF of an Erlang is given by

$$\Pr(z) = \frac{z^{n-1}e^{-z/\beta}}{\beta^n \Gamma(n)} \tag{6.27}$$

where n would be the number of pixels that went into the average if all weights were the same, and β would be $1/\lambda$. Because our weights are not all the same some pixels contribute more than others, which leads to an effective number that went into the average less than the true number, and a variance more heavily weighted by the small variance pixels. We can solve for the effective parameters by writing down the mean and variance of this distribution, and setting them equal to those for our weighted average.

$$\langle z \rangle = n_{\text{eff}}\beta_{\text{eff}} = \frac{\sum \lambda_i}{\sum \lambda_i^2} \tag{6.28a}$$

$$\text{var}(z) = n_{\text{eff}}\beta_{\text{eff}}^2 = \frac{1}{\sum \lambda_i^2} \tag{6.28b}$$

Solving these equations yields

$$n_{\text{eff}} = \frac{\left(\sum \lambda_i\right)^2}{\sum \lambda_i^2} \tag{6.29a}$$

$$\beta_{\text{eff}} = \frac{1}{\sum \lambda_i}. \tag{6.29b}$$

6.3.3 Power Spectrum Plots

The primary output of εppsilon is power spectra, and power spectra plots. These take some deciphering, so I will present a standard set of plots made from the golden data set.

The first set of plots is shown in Figure 6.13. This is a six panel plot showing the dirty, model, and residual two dimensional power spectra in polarizations XX and YY. The dirty spectra are formed from the calibrated visibilities, the model from the simulated visibilities

from our foreground model, and the residual is formed from subtracting model visibilities from dirty visibilities. Note that this is *not* the same as subtracting the middle panel from the left panel as the subtraction is done before the square. We arrive at two dimensional power spectra by averaging in $k_\perp$ rings from the full three dimensional space. This is a useful plot because it maintains the symmetries involved in both the signal (spherical) and foreground (cylindrical). We have found this power spectrum space to be extremely useful is diagnosing problems with the data and the analysis itself. The power spectrum, by design, is extremely sensitive, and so it displays errors from any point in the pipeline.

To help orient the user, the axes of these plots are shown in cosmological units of $k_\perp$ (bottom horizontal axis) and $k_{||}$ (left vertical axis) as well as instrumental units of baseline length u (top horizontal axis) and delay mode (right vertical axis). Recall that simple mappings relate these units (Chapter 2).

Several of the features already described are evident in these plots – the low $k_{||}$ (or delay) modes are completely contaminated by foreground power, the wedge is marked by the solid black line, the coarse band harmonics appear at periodic $k_{||}$ stripes, the low uv coverage at higher $k_\perp$ (or u) results in a contaminated region on the right side of the plots, and an EoR window is evident at low $k_\perp$ and high $k_{||}$ where not otherwise contaminated. Power oscillating between positive and negative (dark blue) is a good indication that the power spectrum is noise dominated in those regions. Negative power is slightly counter-intuitive, but nonetheless expected because we subtracted the noise bias. The noise dominated regions show the benefit of the EoR window concept. We saw that in just two minutes of integration we were signal dominated in every single pixel of an image. Here we see after three hours of integration we have regions of k-space which are $\sim$6 orders of magnitude lower than the foregrounds, and noise dominated. As we integrate longer we expect the noise to drop and further systematics to potentially emerge. Because the data set for this work is much less than needed for an EoR detection, a success will be to have regions of power spectrum space dominated by noise.

Figure 6.14 shows the propagated error in the same two dimensional space – both the

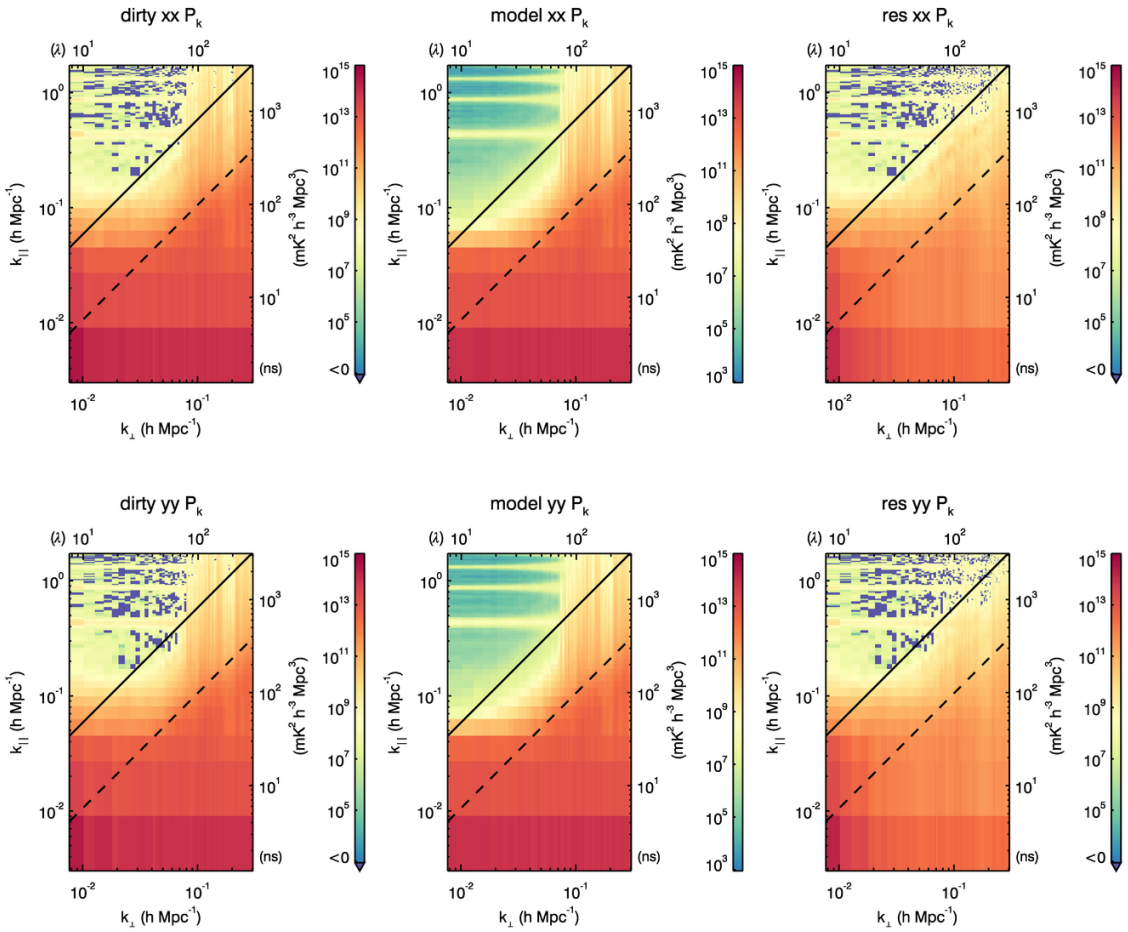

Figure 6.13: Two dimensional power spectra from the three hour integrated golden data set. The columns are the dirty, model, and residual power spectra. Many of the features discussed in the preceding sections and chapters can be seen. The low $k_{||}$ modes are heavily contaminated by foregrounds. Although this contamination drops by about two orders of magnitude in the residual, it is still much higher than the expected signal, and thus we must turn to the EoR window for a detection. We can see the contamination thrown up into the wedge feature, illustrated with the diagonal black solid (horizon) and dashed (field of view) lines. The poor uv coverage at relatively large $k_\perp$ values causes additional power to be thrown higher in $k_{||}$. We can also see the coarse band harmonics at periodic $k_{||}$ modes. The noisy regions in the EoR window is where we hope to place an upper limit on the cosmological power spectrum.

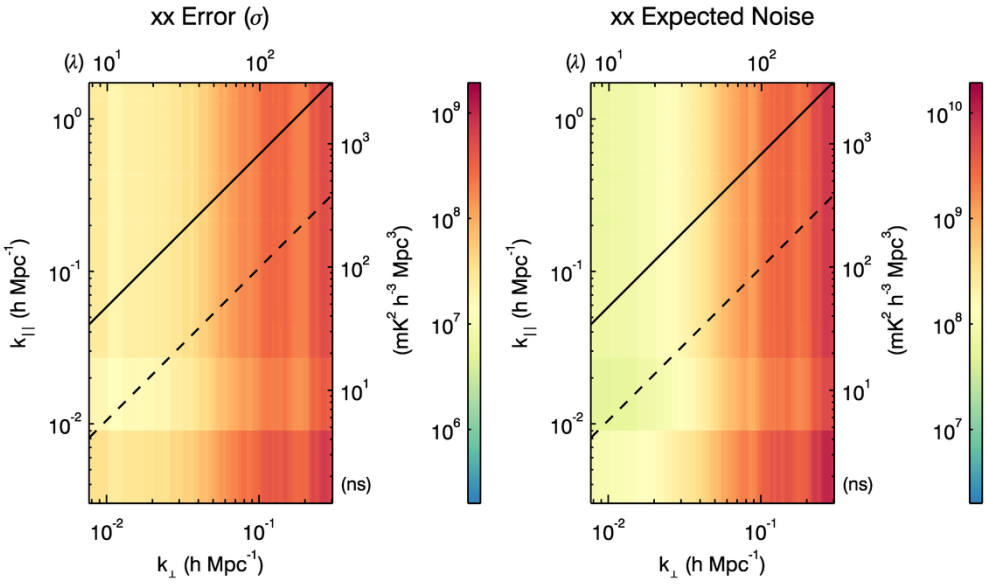

Figure 6.14: *Left:* The propagated error bar for the three hour golden data set integrated power spectra shown in Figure 6.13. The high baseline density at low $k_\perp$ results in lower error bars as expected. *Right:* The expected noise which has been propagated through the entire end-to-end pipeline. This is equivalent to the uncorrected noise bias of the data power spectrum, or the expected power spectrum of the difference cube.

error bar on the power spectrum (σ) and the expected noise values. Though closely related, these plots differ due to the effect discussed in the previous subsection. We see that the error is lowest at low $k_\perp$ as expected due to the large density of short baselines in the MWA. The expected noise is the noise bias that we would expect without accounting for it with the sum and difference cubes. These two plots are propagated through the analysis and never contact the actual data after the initial calculation of noise on the calibrated visibilities.

Figure 6.15 compares the expected noise with what is actually observed by forming a power spectrum from the difference cube. The left panel looks very similar to the right panel of Figure 6.14 except that it is not as smooth because each pixel is a realization of the noise with a standard deviation given by the left panel of Figure 6.14. The right panel of

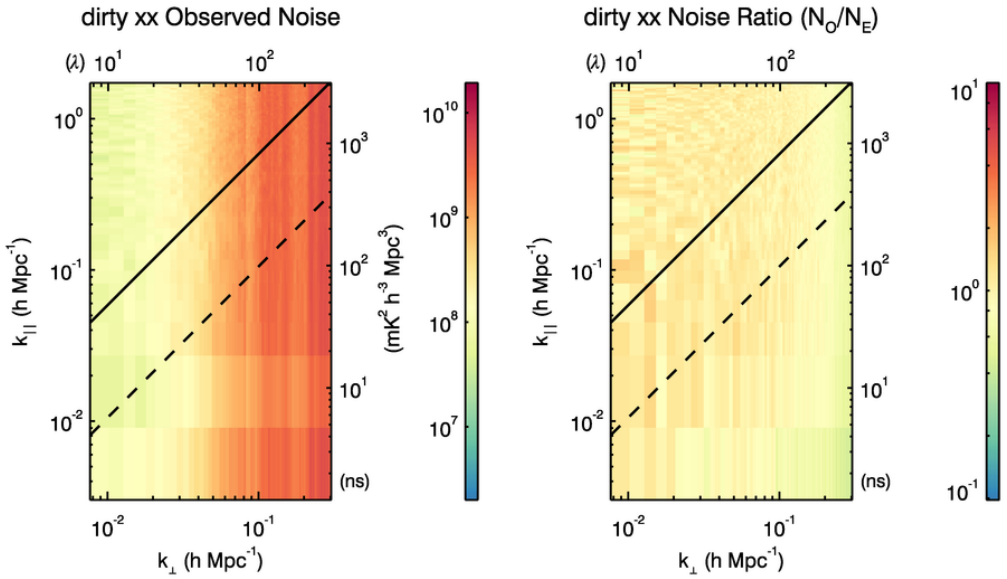

Figure 6.15: *Left:* The observed noise from the three hour integrated golden data set. This plot is the power spectrum of the difference cube, and can be thought of as the observed noise bias which we subtract from the data to recover an unbiased power spectrum estimator. *Right:* The ratio of the observed noise to the expected noise (right panel of Figure 6.14). This diagnostic plot is an end-to-end check on our error propagation as the expected noise does not touch the data after the noise is initially measured from the calibrated visibilities.

Figure 6.15 is the ratio of the observed noise to the expected noise. We see that the ratio is centered at 1 and has small fluctuations due to the randomness of actual noise. This plot has been extremely useful in diagnosing our pipeline, and provides an end-to-end check on our error propagation.

Finally Figure 6.16 shows an example one dimensional power spectrum. Here we have averaged in spherical shells of constant k magnitude, and made no effort to exclude foreground contaminated regions. Even without masking out the wedge or low $k_{||}$, we can see the benefit of foreground subtraction. The residual power spectra are significantly lower than

their corresponding dirty spectra. Some relatively large k modes even approach the noise level. We will improve on this greatly in Chapter 8 when we select regions of k-space to include in our final power spectrum.

The plots we have shown in this section are not the cutting edge best power spectra available. However, they do provide a convenient reference when experimenting with improved analysis methods, as we will explore in the next chapter.

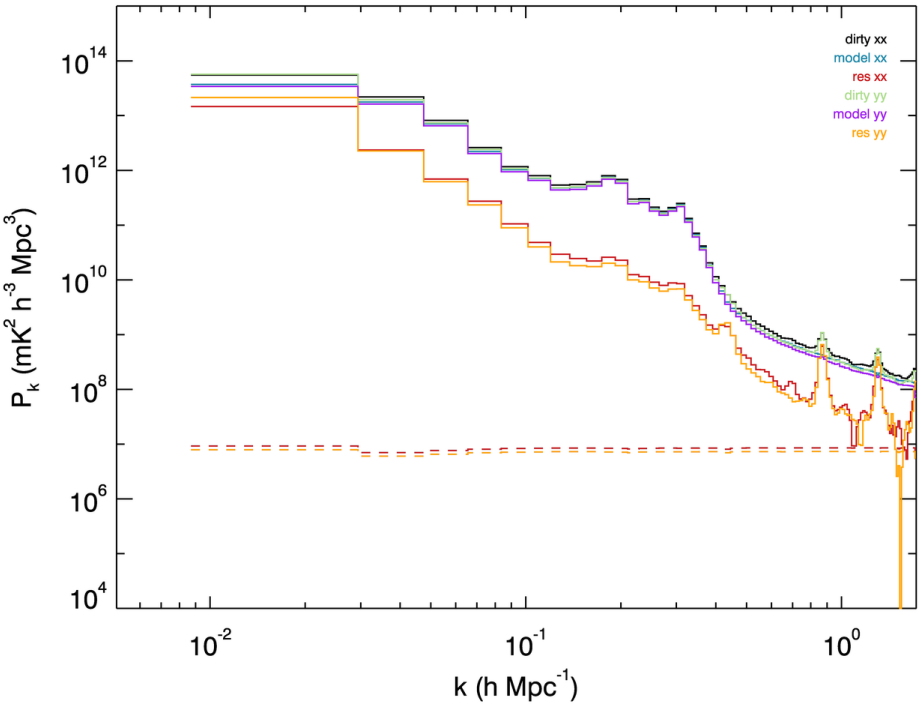

Figure 6.16: The 1D power spectrum from three hours of golden data. We have not made any effort yet to restrict the pixels which contribute to the binning, so this power spectrum includes much foreground contaminated signal. Data selection will be the focus of Section 8.1. Even without a rigorous selection we can see that higher k modes approach the noise level (dashed lines), achieving many orders of magnitude dynamic range when compared to the bright foregrounds at low k. We can also see the coarse bands in this plot at $k \sim 0.9$ and 1.3 h Mpc^{-1}.

Chapter 7

CLEANING THE EOR WINDOW

"There is nothing like looking, if you want to find something. You certainly usually find something, if you look, but it is not always quite the something you were after."

– J.R.R. Tolkein, *The Hobbit*

When looking for the EoR via the 21cm power spectrum, it is very easy to find all sorts of systematics creeping in that were never expected. It is essential to be diligent about maintaining a clean EoR window to have any chance at an eventual detection. In this chapter I highlight a handful of improvements we have made to the analysis to remove systematics and foregrounds. These examples are not a full catalog of effects seen, but provide the reader with a flavor for the problems involved in a cosmological measurement.

The power spectrum itself is an extremely sensitive tool for both making a measurement, but also for diagnosing systematics. Many troublesome bugs and issues cannot be seen in images, but become very apparent when we form power spectra. Because a clean power spectrum is our eventual goal, we have adopted a policy to always inspect the impact on the power spectrum before making changes to the analysis. This has resulted in the need for hundreds of iterations of our pipeline, and because this is simply impractical to do over all the data, we choose to focus on the three hour golden data set until we believe we cannot learn more without going deeper. Therefore, the plots shown in this chapter are all taken from the golden data set.

7.1 The Rogue Gallery

While developing the calibration, imaging, and power spectrum pipelines, many strange and unexpected "features" arose. Many of these effects turned out to be the result of analysis choices made in earlier radio astronomy pipelines that were insufficient for the precision cosmology we are attempting, others are unique properties of the MWA instrument. Here I document three such features and how we fixed them in our analysis.

7.1.1 The "Fourth Line"

When making power spectra in our early analysis we were confident we understood the various structures in and outside the EoR window. We have the low $k_{||}$ foreground contamination, the mode-mixed wedge, poor uv coverage throwing power higher in $k_{||}$ at large $k_{\perp}$, and the harmonics of the coarse band shape. The locations of the coarse band harmonics are easily predictable, given our observing parameters and the coarse band width. With a central frequency at 182 MHz and coarse bands being 1.28 MHz, we expect these bands to show at $k_{||,\mathrm{coarse}} \approx 0.44n$ h Mpc^{-1}, where n is an integer greater than zero.

Within our typical $k_{||}$ range we expect to see three coarse bands, and for some time we always saw three coarse bands in the power spectra. However, after we implemented the second step of our calibration to enforce smooth gains, a mysterious "fourth line" revealed itself (see left panel of Figure 7.1). Because this line was absent when we made no restrictions on the gains, and only showed up when we required the gains be smooth in frequency, we concluded that the line existed in the actual data and we were previously calibrating it out. However, the nature and origin of the line was yet unknown.

Partially influenced by the delay spectrum style of analysis, we began to think about the line not as a cosmological scale in terms of $k_{||}$ value, but as a delay time. Given the location of the line, we can calculate the delay time to be $\tau \approx 1.23$ μs. This corresponds almost exactly to twice the signal travel time through our 150 meter beamformer to receiver cables (with a velocity $0.81c$). We therefore suspected what we were seeing was a signal reflection

in the cables.

If our suspicion was correct, we can easily model the reflection term in the calibration gains. We first write down the signal in the receiver in terms of the incident signal and the reflected signal (ignoring other gain terms for now).

$$E_i'(t) = E_i(t) + R_{i,0}E_i(t-\tau) \tag{7.1}$$

We have lumped all phase and attenuation accumulated through the reflection process into one complex parameter, $R_{i,0}$. Next we consider a single frequency channel, and propagate the reflected signal.

$$E_i'(t,\nu) = E_i(t,\nu) + R_{i,0}E_i(t-\tau,\nu) \tag{7.2a}$$

$$= E_i(t,\nu) + R_{i,0}E_i(t,\nu)e^{-2\pi i\tau\nu} \tag{7.2b}$$

$$= E_i(t,\nu)(1 + R_{i,0}e^{-2\pi i\tau\nu}) \tag{7.2c}$$

We an see that the reflection manifests as a ripple term in the response of the antenna. In other words, the instrument introduces a small ripple into the frequency structure of the natively spectrally smooth foregrounds. The level of this reflection is below the signal to noise necessary to see it in images, but is easily observed in the highly sensitive power spectrum. Before we enforced smooth calibration solutions, we fit each frequency independently and inherently fit this term out. Now that we understand the nature of the line, we can fit for the reflection parameter, $R_{i,0}$, without needing to return to the fully unrestricted gains. After implementing a fitting scheme, we see that the "fourth line" disappears from our three hour power spectrum (right panel of Figure 7.1).

When we inspected the line even closer we found that the reflection time, τ, was not precisely known. This is due to uncertainty in the velocity in the cable and variations in the cable lengths. We therefore also allow for a fit of the reflection time for each antenna individually. This is done by performing a DFT of our unrestricted gains to a finely resolved grid of reflection times after removing the bandpass and polynomial terms. We then find the reflection time which maximizes $|R_{0,i}|$ and select it as the time parameter for that antenna.

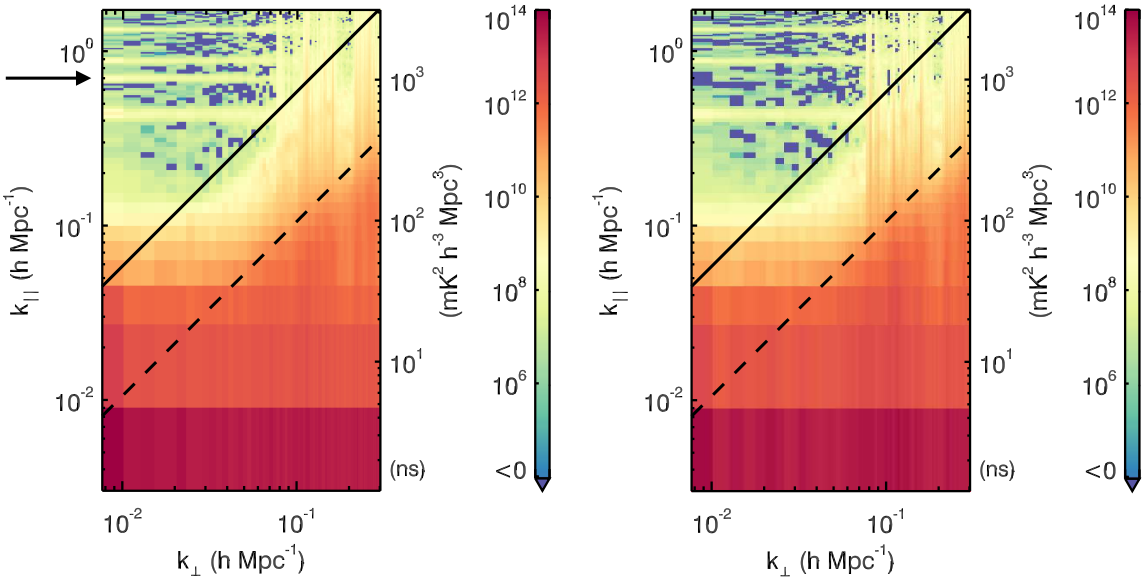

Figure 7.1: Power spectra demonstrating the "fourth line". *Left:* A three hour dirty power spectrum made from the golden data set, before we implemented the cable reflection fitting into the calibration step. The black arrow points to a horizontal band at $k_{||} \approx 0.6$ h Mpc^{-1}, which cannot be accounted for by the coarse bands. This line turned out to correspond nearly exactly to a delay mode with delay equal to twice the signal travel time in a 150 meter cable. *Right:* After we implement the cable reflection fitting in our calibration solutions, we see the fourth line disappears.

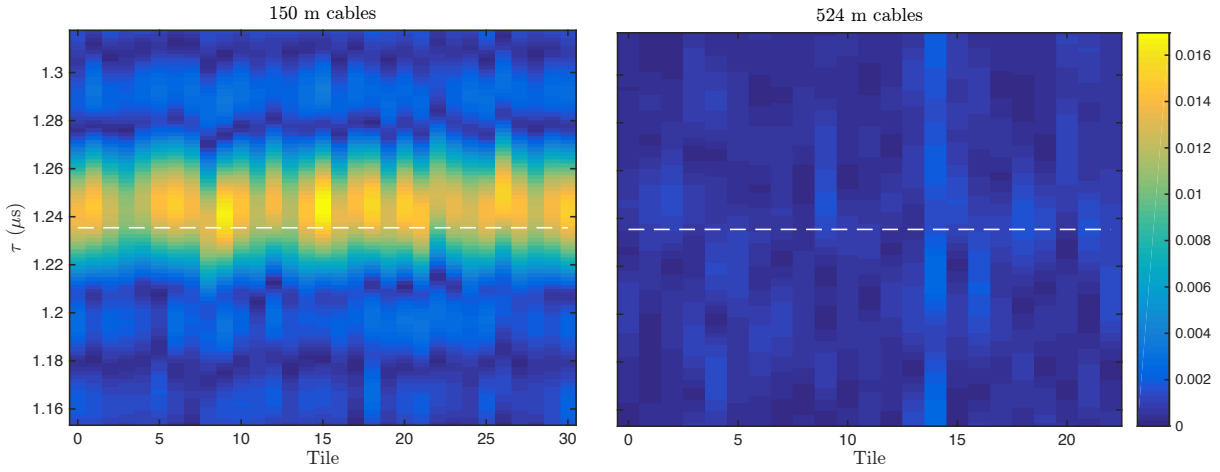

Figure 7.2: Reflection mode fitting. We DFT the unrestricted calibration solutions (after removing the bandpass and polynomial fits) to a highly over-resolved delay grid (vertical axis). We then select the bin which maximizes the reflection coefficient to use in our restricted gain solution. On the left I show the procedure for all tiles with 150 meter cables – those expected to contain a reflection term contributing to the fourth line. The white dashed line corresponds to the nominal delay for exactly 150 meter cable and a velocity factor of 0.81. We see that most tiles have a delay mode slightly greater than the nominal value, and there is variation between the tiles. For reference I show the same calculation done for tiles with 524 meter cables on the right. Here the reflection terms are completely absent.

An example of this process, integrated over three hours for clarity, is shown in Figure 7.2. For comparison, I also show the same modes for tiles with 524 meter cables. It is clear only the 150 meter cables are contaminating this particular region of $k_{||}$ modes.

This reflection line is an excellent example of the utility of the power spectrum as a diagnostic, and the importance of understanding the data through the entire pipeline. Only by understanding a physical problem in the hardware, and being able to propagate its effect through to the power spectrum, were we able to diagnose and resolve this issue.

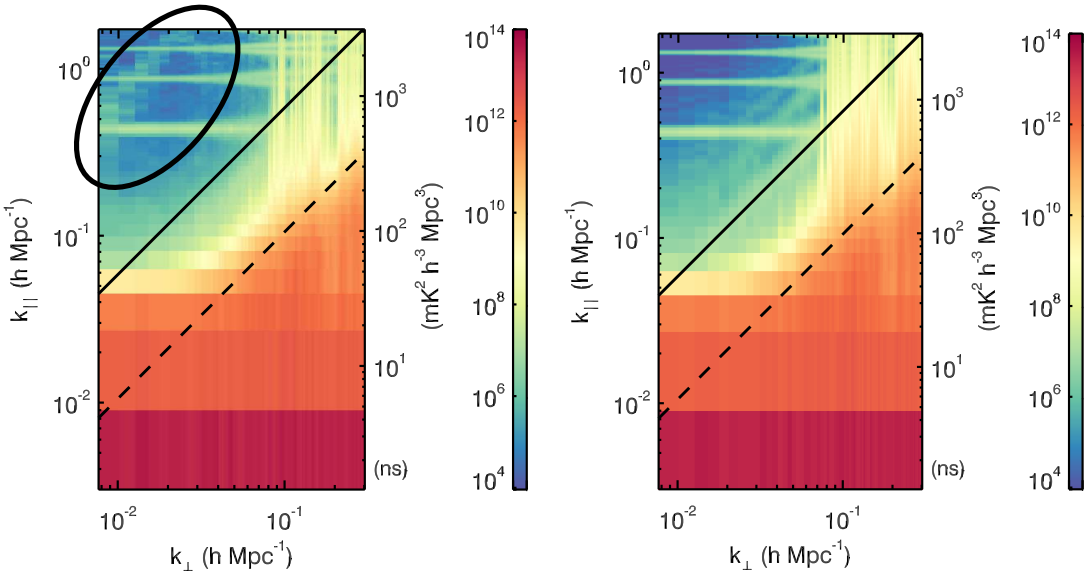

Figure 7.3: *Left*: Model power spectrum demonstrating the appearance of the "ghost line". Circled in black is the faint line with slope the same as the wedge, which then bounces off the top of the plot and is aliased back down. *Right:* Model power spectrum after increasing the beam model resolution to 0.007 wavelengths. The ghost line is effectively gone. In the process of fixing the ghost line, however, we introduced a beam clipping bug which created the periodic lines seen just above the solid black line. This bug has also since been fixed.

7.1.2 The "Ghost Line"

The reader may have noticed that the collaboration often adopts names for various features we see in the power spectrum. The "ghost line" refers to a very faint line we saw in the *model* power spectrum at high $k_{||}$, shown in the left panel of Figure 7.3. This line was particularly perplexing because it was evident in the model, which should only contain power we include. Because our model is spectrally smooth, and we have very good uv coverage at low $k_{\perp}$, it was hard to explain how the line showed up at high $k_{||}$ – especially an arcing line that did not look like any structure we understood before.

We shed some light on the situation by forming a power spectrum at higher frequency resolution. This meant a higher range of $k_{||}$, and we saw that the ghost line did not actually turn over, but instead continues up. It only appears to arc in our normal power spectra due to it being aliased down. A straight line with a slope the same as the wedge is a little more understandable. We simply needed to find a step in our analysis that would create a wedge-like feature, but at a much higher mixing.

The culprit ended up being the resolution at which we model the primary beam in uv (the Fourier transform of the image-space primary beam). In order to save on computation, we create an over-resolved model of the primary beam, then use a nearest-neighbor selection to choose the model grid points which most closely line up with the uv grid point when gridding the visibilities (or model visibilities in this case). Because the baselines migrate in frequency and we are selecting from a predetermined set of model points, we are effectively introducing discrete shifts in the beam. This can also be thought of as an error in baseline location that changes as a function of frequency, throwing power up in $k_{||}$ when we perform the line of sight Fourier transform. The uv grid is at half wavelength resolution (to image horizon to horizon), and we were modeling the beam at a resolution of 0.04 wavelengths, which we believed to be sufficient. However, this undesirable effect in the power spectrum proved that we must over-resolve the primary beam even further. While the most accurate answer is to model at infinite resolution, or to implement interpolation between grid points, it is not feasible to do so computationally. Instead we choose a resolution at which the effect no longer impacts our power spectrum – both because the power is thrown into higher $k_{||}$, and because the amplitude is decreased with better models. With experimentation we found the sweet spot to be a beam model resolution of about 0.007 wavelengths. The result is shown in the right panel of Figure 7.3.

This exercise has shown the value in retaining the model through the entire pipeline. Because the model does not contain any noise, it is helpful to identify problems with our analysis that contaminate the EoR window without having to integrate deeply. We would not have been able to see the ghost line in data until we integrated much deeper, which

would have resulted in a very slow debugging cycle. Processing is limited by the number of snapshots we put through the imaging pipeline, and the number we would need to see the ghost line would take more than a month to process each time we wanted to test a solution. Our aim with the model power spectrum is to maintain a clean and quiet EoR window.

7.1.3 Compression of MWA EoR Data

In order to save valuable storage space on the MWA archives, a compression of all raw GPU files was proposed. The input of the cross-multiply leg of the correlator is integers, which are cast into floating point to leverage the software libraries available for GPUs. The cross multiply maintains integer values, although represented as floats. Then the correlator averages in frequency by a factor of four in the case of EoR data. The compression scheme multiplies by four then recasts to integers which are substantially lower cost to store given the dynamic ranges in our data. Overall the compression saves about a factor of two in storage space – a huge gain if permissible. This procedure should be information lossless, but upon inspection we saw that the actual visibility data values differed by tiny amounts (about a part in 10^7). For more traditional radio astronomy this difference is tolerable, but when attempting to detect a signal many orders of magnitude below the noise, it is important to study the impact before committing to a compression of all data. In particular, if the errors in the visibility data have a systematic bias they can skew our precision measurements, ruining our hope for an accurate power spectrum detection.

Spoiler alert: The following analysis was done under the assumption that the compressed data was less correct than the uncompressed. However, we later discovered that the compression procedure actually *fixes* a small error in the data due to the way the correlator converts from integer to float at the front end. This will be explained in more detail at the end of the section, but should be mentioned here to help the reader remember that the compressed data is actually an improvement.

To study the effect of the data compression, we used the GPU files from a single observation from the golden data set. We used an uncompressed and a compressed copy of

Table 7.1: Table of compression variables

Variable	Description
d_u	Raw visibility data, from uncompressed GPU file
d_c	Raw visibility data, from compressed GPU file
δ	Difference between raw data, $d_u - d_c$
N	Number of data points in a set
$\langle\delta\rangle_N$	Average of δ over N data points
$\sigma_{\delta,N}$	Standard deviation of δ over N data points

each file for our analysis, and were agnostic to real/imaginary components of visibilities, but instead treated all data as simply numbers. Table 7.1 summarizes shorthand variables used for different data products that were read or computed.

First we look at the convergence of the difference between the uncompressed and compressed data. With an unbiased error between the two data sets, we expect the average of the difference, $\langle\delta\rangle_N$, to be on order $\sigma_{\delta,N}/\sqrt{N}$. Figure 7.4(a) shows the running average of the difference over entire GPU files. Each GPU box seems to have a bias, but that bias differs depending on the GPU box. Recall that the different GPU boxes correspond to different coarse channels in our frequency band, so one can easily see how structure in these biases could introduce artificial spectral structure in our data. We chose three GPU boxes that cover the range of the biases that we see to show the convergence in Figure 7.4(a), while Figure 7.4(b) shows this convergent value for each GPU box. Notably, all the biases are positive. There is a clear jump from GPU box 8 to 9, which corresponds to the digital gain jump.

One question that arises from seeing systematic biases like in Figure 7.4 is whether the bias is an additive one (independent of data value) or multiplicative (dependent on data value). To answer this question we histogram the number of data points versus data value,

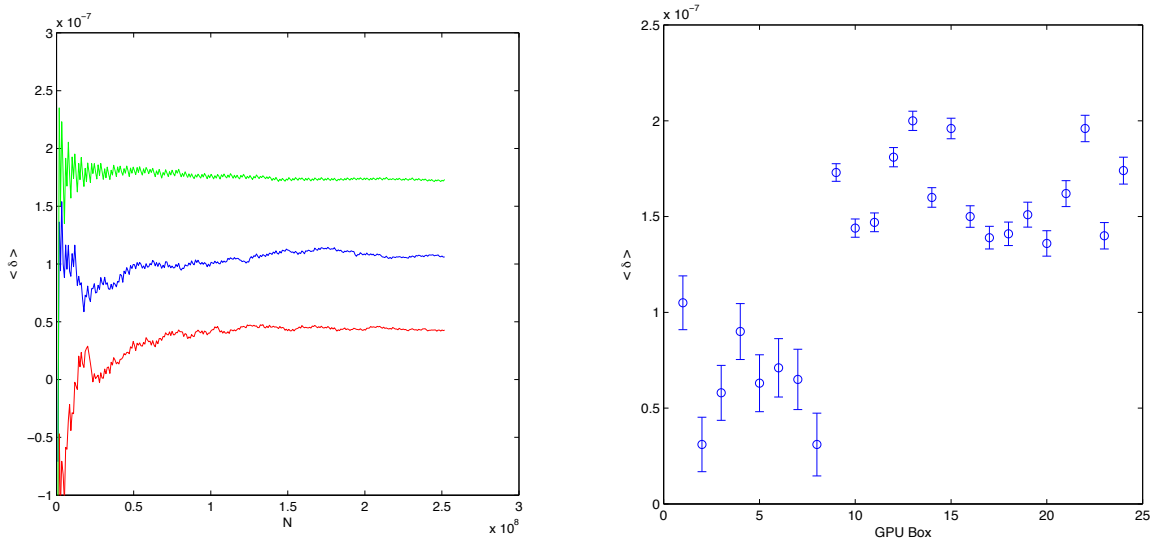

Figure 7.4: Compression bias per GPU box. *Left:* The running average of δ converges on non-zero values. Shown are examples for GPU boxes 1 (blue), 2 (red), and 12 (green). This color scheme is the same for all the remaining plots in this section. *Right:* The compression bias for each GPU box. The error bars are the expected magnitude of convergence, given by $\sigma_\delta/\sqrt{N_{tot}}$. The exceptionally large data values from the autocorrelations were removed from the data set for these plots.

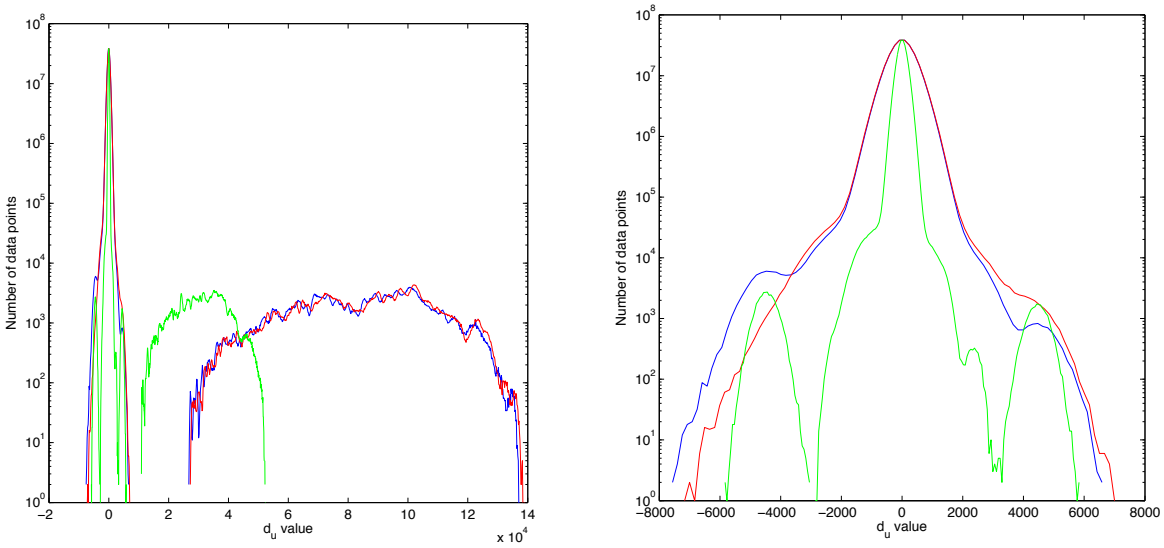

Figure 7.5: Histogram of data values in a few GPU files. Colors correspond to the same GPU boxes as in Figure 7.4. *Left:* The full data range, including large positive values from autocorrelations. *Right:* Autocorrelations trimmed out by restricting values to less than 8000.

and show the sum of the differences in each data value bin. The results are shown in Figure 7.5. Two clear distributions are apparent in the left pane - the data from cross-correlations centered on zeros, and the data from autocorrelations at large positive values. These large, systematically positive values may introduce problems for the compression, and they do not contribute to power spectrum analysis. Therefore we remove them from the data set.

Finally, we look at the the compression bias, $\langle\delta\rangle$ as a function of data value. For a multiplicative bias, we expect a linear relationship between $\langle\delta\rangle$ and d_u. If this were the case, we could imagine removing the bias with a simple calibration as all the data would be affected in the same way. If the data were additive we would expect a flat bias as a function of d_u. Again, this could be easily removed, this time by subtracting the bias out of the data.

However, in Figure 7.6 we see a complicated dependence. For small values of d_u, the bias

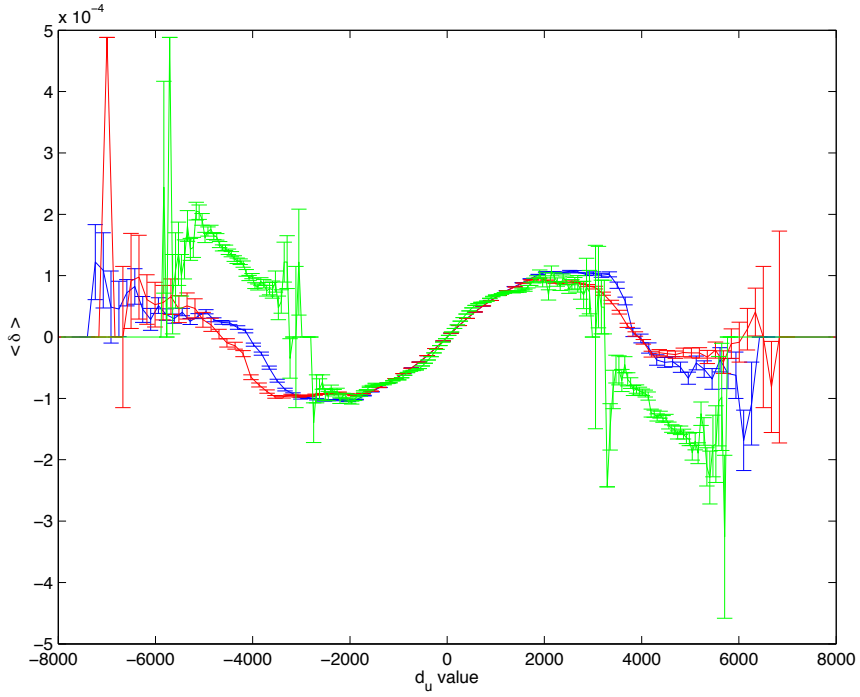

Figure 7.6: The compression bias as a function of data value. A strong positive correlation between bias and data value exists for small values. But at large values the dependence becomes much more complicated.

seems to be multiplicative. But at some point (dependent on GPU box), the bias turns over and becomes negatively correlated with d_u. This complicated structure is harder to model and take out of the data without a better understanding of the origin and cause of the bias.

This bias we see in the difference between uncompressed and compressed data only has real consequence if it has the potential to corrupt the subsequent analysis. We ran the compressed and uncompressed data through our entire analysis pipeline using identical code, identical flags, and identical calibrations. A power spectrum from the uncompressed data is shown in Figure 7.7(a).

While the two dimensional power spectrum is often useful to see structures and other

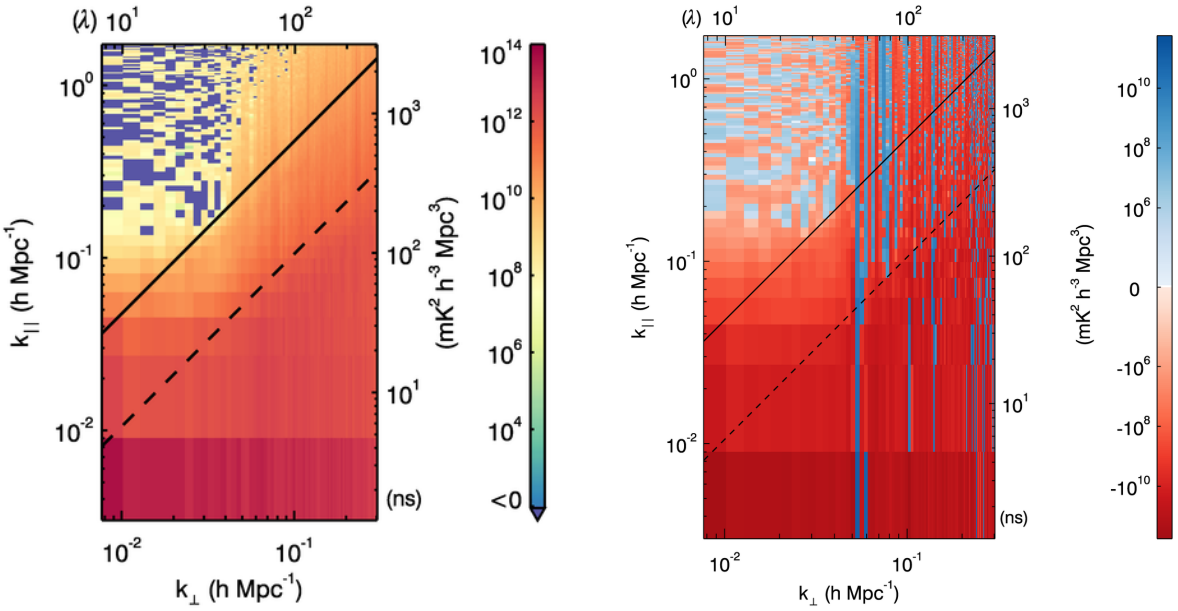

Figure 7.7: *Left:* Snapshot power spectrum created from uncompressed GPU files. *Right:* Difference between power spectrum created from uncompressed files and that created from compressed files. The difference is about a part in 10^4 in the wedge, and about 5 orders of magnitude lower in the EoR window – comparable to the expected cosmological signal.

effects, this particular difference was small enough that it would not be visible by comparing power spectra side by side. We instead introduce another diagnostic plot which is useful to highlight differences between two different analysis runs – the difference power spectrum, shown in Figure 7.7(b). Note that this difference is not created by subtracting images then forming a power spectra in the way that a residual power spectrum is formed. Instead this difference is the difference of the actual power spectra. The color scale therefore is signed, and we have adopted a special logarithmic scale which allows for negative values as well.

While often times insightful, the difference power spectrum can be difficult to interpret without careful consideration of what is "good" or "bad". In the case of compression, we simply wished to quantify the magnitude of the difference. The difference is several orders of

magnitude lower than the power spectrum itself, but not as low as we hoped. The difference in the foreground region (low $k_{||}$) is about a part in 10^4, which is substantially higher than the expected cosmological signal. In the window the differences appear noise-like, but when inspected closer it is clear that the hash follows the same patterns as the raw power spectrum, and therefore is the noise of the data, not the noise of the compression. We concluded that the compression will indeed impact the power spectrum measurement, and we needed to investigate the cause of the differences before we could sign off on a full compression of the data archive.

The origin of the data deviation turned out to be at the front end of the cross-multiply leg of the correlator. The input of the this step is 4 bit integers (ranged -7 to 7), which are scaled and cast to float (ranged -1 to 1). For this promotion, the correlator uses the GPUs' auto promotion in texture memory, which requires a division by 127. This division is not fully representable in binary as 127 is a prime number, so this recast truncates at a low bit level. Thus a small error is introduced.

The bit precision from this step on is deep enough that the small errors introduced cannot obscure the true values one would obtain with integer arithmetic[1]. Therefore, a rounding step (as is done with the compression) actually undoes the small error introduced in the correlator, and the compressed files can be considered to be more accurate than the uncompressed.

This was a rare example of a truly free lunch. We not only saved about a factor of two in archive space, but also identified and removed a potentially harmful problem with our data at the same time. This further demonstrates the importance of understanding the data at all levels of the analysis in order to maintain the pristine data necessary for our precision measurement.

---

[1] As the correlator accumulates in time the potential error grows as square root of the time steps integrated. For the specific bit depths and time accumulation relevant for the MWA there is no danger of this error growing to an integer level. However, other data flow paradigms (such as LST binning) may see this error become problematic if they accumulate significantly more data before correcting.

7.2 Foreground Model Testing

The process of forming power spectra to evaluate our analysis is good for both hunting bugs and systematics, but also for determining how changes in our foreground model affect our final results. Here I will discuss three major improvements we have made to our model beyond the MWACS catalog.

7.2.1 Point Sources

The most obvious place for improvement in the model is the point sources in our main lobe. The commissioning survey was performed with subarrays of 32 antennas and very few long baselines, so we expect we can improve on this with data from the full 128 antenna MWA. Using the strategy outlined in Figure 6.3, we deconvolve point sources with FHD on a subset of data. In fact we use a subset of the golden data set for this purpose.

Currently FHD can deconvolve a single snapshot at a time, resulting in a separate catalog of point sources for each snapshot. We next combine the catalogs by identifying repeated sources and determining reliability based on detection fraction, signal to noise, beam values, and other statistics. We want to ensure that sources we include in our model are true sources because they both influence the calibration, and are subtracted from all our data. P.A. Carroll has developed a machine learning algorithm to identify reliable sources to include in our catalog, and will detail the process in a series of upcoming papers.

The resulting catalog consists of 2,690 point sources in the EoR0 field within the MWA primary beam at 182 MHz. The number of sources is only a mild improvement over the same sky region with the MWACS catalog (about 2,400 sources). However, with the improved array and statistical analysis, we expect the newer catalog to be an improvement. As we have seen repeatedly, the real test is the change in the power spectrum.

Figure 7.8 shows the result of comparing power spectra created with the MWACS catalog and with our improved catalog using FHD on the golden data set. In this difference plot we are showing the power spectrum using MWACS minus the power spectrum with the new

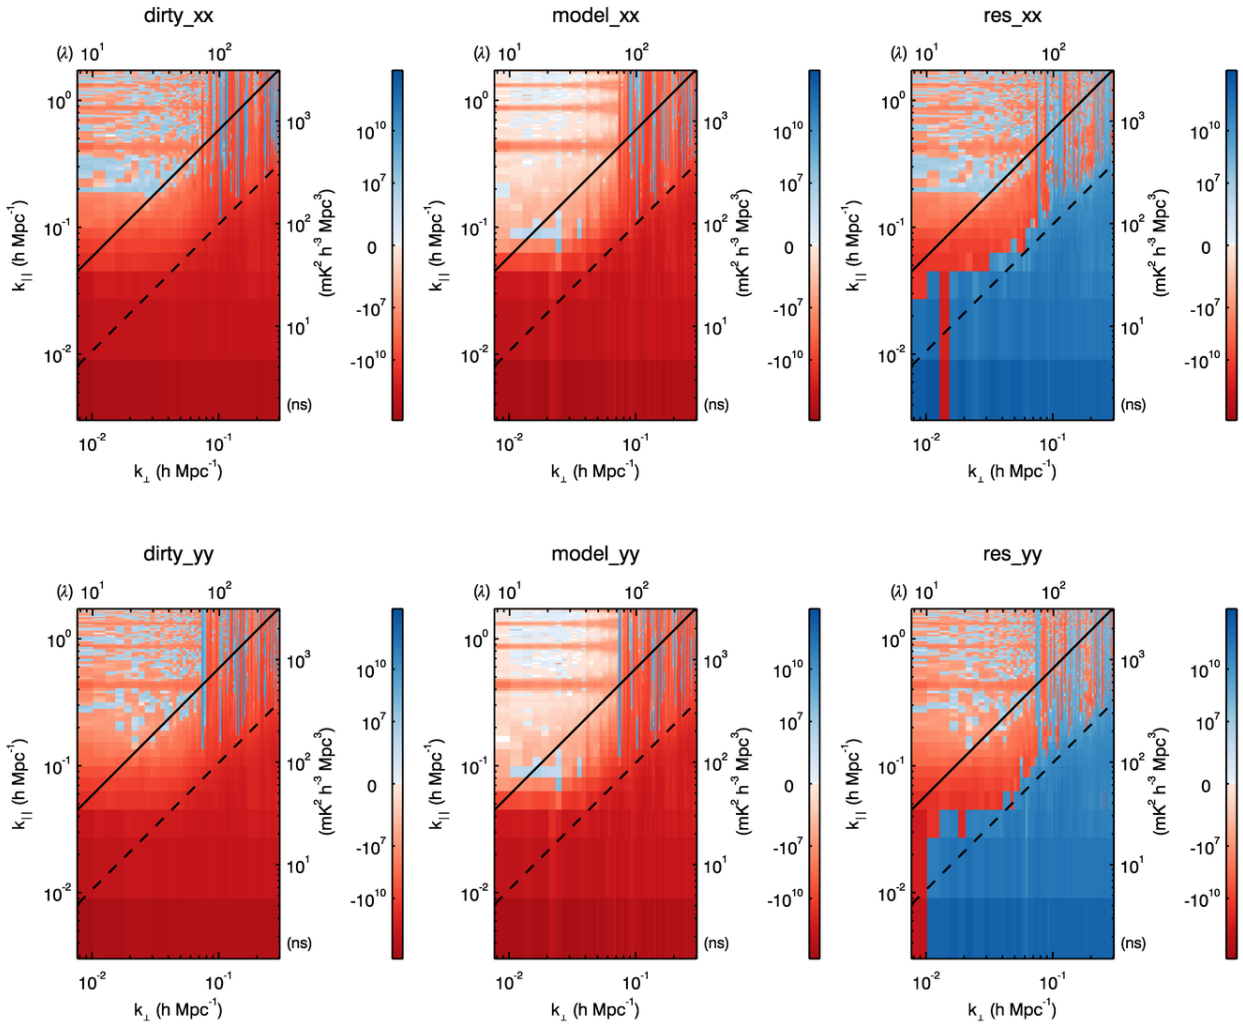

Figure 7.8: Golden data set power spectrum differences showing the effect of a new point source catalog created using FHD. The dirty power spectrum has higher amplitude with the new catalog due to an improved calibration. The model is higher due to more point sources in the catalog. The residual is lower in the wedge indicating better subtraction of foregrounds.

catalog, so red indicates higher power with the new catalog.[2] To understand this figure we walk from left to right. The dirty power difference is heavily red, indicating more power when using the new catalog. This is due to the calibration changing. When a catalog more closely matches the sky, the calibration better locks in bringing to focus the sources that are already in the data, increasing the total power.

The model difference is also very red, which is expected due to more sources in the new catalog. Notice though that the sources are in the primary field of view, thus lie mostly below the dashed black line.

Finally we turn to the residual difference. Here the wedge is very blue – indicating less power in the residual when using the new catalog. Despite the calibration increasing the power overall, the sources better subtract with the new catalog. While very encouraging not all is good news. The large red region above the wedge (within the boundaries $0.05\ h\ \mathrm{Mpc}^{-1} < k_{||} < 0.2\ h\ \mathrm{Mpc}^{-1}$, and $k_{\perp} \lesssim 0.6\ h\ \mathrm{Mpc}^{-1}$) is indication that there is foreground contamination at higher $k_{||}$ – perhaps due to spectral structure of the point sources, or from sources outside the field of view. However, the clues we have seen so far indicate that indeed our calibration has improved and that despite the power going up in this region, it is more correct and we should not try to avoid it. Instead we will need to tackle this problem from another approach.

7.2.2 Sidelobe Sources

Way back in Chapter 2 we saw that the wedge is created by the mode mixing of many foregrounds into higher $k_{||}$ modes than they natively reside. We saw with Equation 2.21 that the extent a single point source mixes up into $k_{||}$ is proportional to the angle on the sky – meaning that sources farther from center mix higher. This means that despite the main field of view being where we hope to make our measurement, perhaps the greatest risk for

[2]We will attempt to adhere to the convention of always using a standard minus a test to keep some consistency in the direction of the difference plots. However, the reader should keep in mind that this does not always imply red is good or bad. One should always consider what the colors mean in the context of the test being performed.

foreground contaminating the EoR window is sources from way outside the field of view. The effect of sources and galactic structure near the horizon on EoR measurements has been shown in Thyagarajan et al. [88]. We therefore explore the possibility of subtracting point sources from the sidelobes of the MWA. This work will be described in J.C. Pober et al. 2015 (in preparation).

Because FHD cannot reliably deconvolve point sources in the sidelobes, we turn to other catalogs for sidelobe sources. In the simple test here we simply use the MWACS sources where available. This results in incomplete coverage of the sidelobes (because MWACS does not cover the full sky), but does provide enough to observe the effect. In order to keep the test as simple as possible, we do not use the sidelobe sources for calibration. This way the only difference is in the model which is subtracted.

The difference power spectra (no sidelobes minus sidelobes) are shown in Figure 7.9. First notice that there is no dirty difference. This is because the calibrations were identical and therefore the dirty power spectra are identical. The strongest difference in the model is seen where we expect sidelobe sources to manifest – the strip above the dashed black line. We have added power in this region, so it is consistently red. This same region is mostly blue in the residual, indicating a successful subtraction. By improving our model in this region we should be able to decrease the contamination power and open the EoR window.

There are currently plans to obtain observations pointed at the sidelobe regions in order to perform deconvolution and the reliability statistics we have used for the primary field. However these regions will contain added complications we have avoided in the EoR0 field which was purposely chosen as devoid of large sources or galactic structure. The sidelobe fields are significantly more complicated, and deconvolving them will be a major feat for a future graduate student.

In the mean time we can construct a better catalog using data available to us. P. A. Carroll has created a hierarchical catalog using source lists from FHD, MWACS, Culgoora [81], and the Molonglo Reference Catalog (MRC, [43]). We prioritize the source lists in this order based on our confidence in their predicted 182 MHz flux. The source lists are first clustered

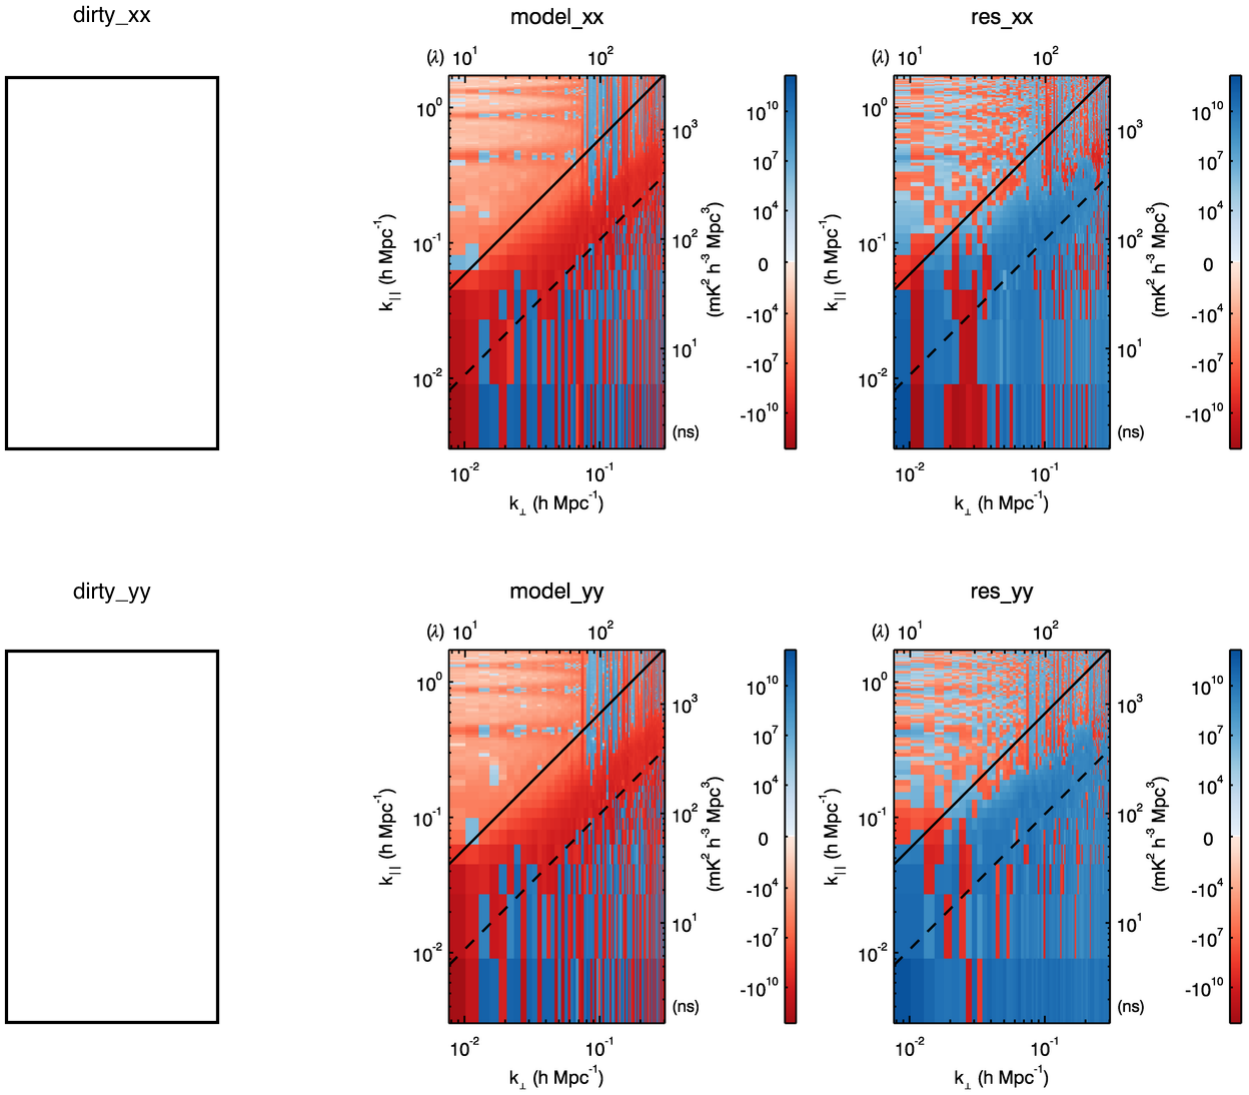

Figure 7.9: Golden data set power spectra differences demonstrating the effect of subtracting sidelobe sources. Because we held the calibration constant in this test the dirty power spectra are identical and the axes simply serve as place holders. The model power spectrum has higher amplitude just above the dashed black line when including sidelobe sources, as expected. The residual power spectrum is lower in the same region due to more power being subtracted. By subtracting sources in the sidelobes, we hope to unlock a larger window.

to avoid repeated sources, using a 3.5 arcmin neighborhood radius. For each cluster the flux from the highest priority catalog was used. Spectral index values were taken from the MWACS and Culgoora catalogs where available. If no spectral index was given, a two-point spectral index was estimated for Culgoora-MRC matches or MRC-SUMSS matches. All other sources were given a spectral index of -0.8. The spectral index was used to extrapolate the flux from the catalog frequency to 182MHz. The resulting master catalog is shown in Figure 7.2.2. The EoR0 field corresponds to the red FHD source patch. In future work we will aim to fill out more regions of the sky with the methods used to create the FHD catalog. In the mean time this aggregate catalog will suffice to fill the sidelobes for foreground subtraction.

7.2.3 Diffuse Emission

Finally we turn to a different type of foreground. The sources we have been subtracting thus far are radio galaxies which can be treated as point sources by the relatively low resolution MWA. There is also significant foreground power in our data due to emission from our own galaxy, which is seen as large diffuse clouds. For computational purposes we have divided the galactic emission into two regimes – the faint clouds in our main lobe, and the bright plane and other structures in the sidelobes. While the latter has been studied by many people and a Global Sky Model (GSM) is readily available [22], the computational obstacle of simulating the instrument response to a full sky diffuse model is yet prohibitive. We instead focus here on diffuse structure within the primary field and leave the full-sky model to future iterations of analysis.

Our method of creating a diffuse model is very simple – we use the MWA to measure it. We simply combine the residual images from the three hour golden data set to create an integrated image devoid of most point sources (and the ones remaining will get subtracted along with the diffuse). There are only a couple minor caveats to this. When normally integrating images we do so in the holographic, naturally weighted frame to minimize signal to noise. But we need this image to be uniformly weighted so it is an accurate representation

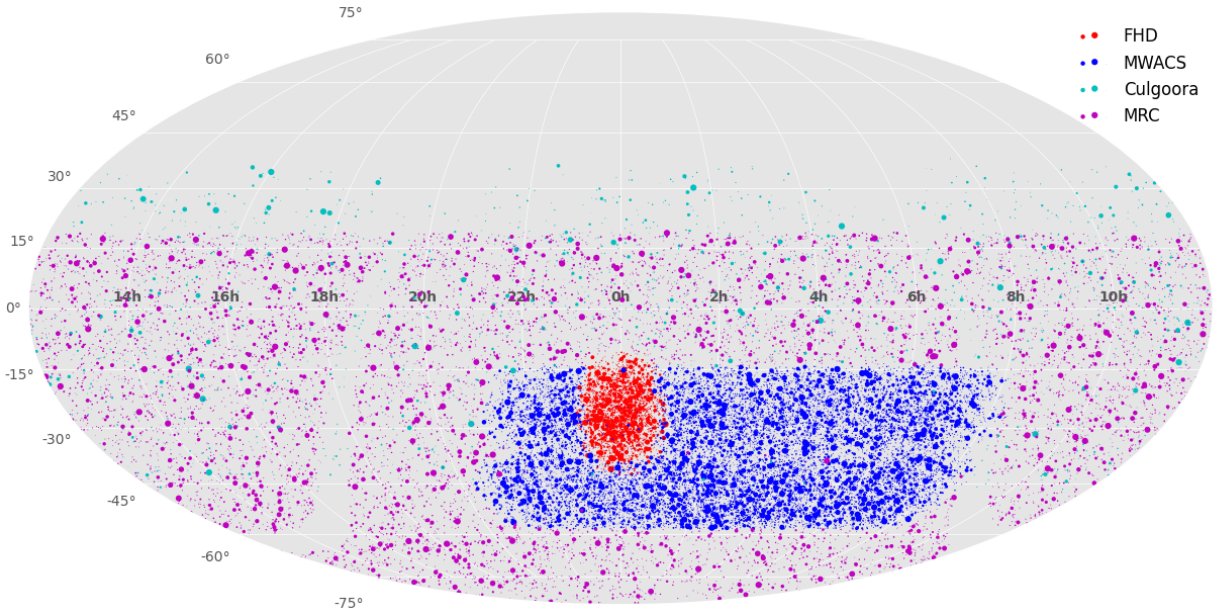

Figure 7.10: Master catalog used for foreground subtraction. This catalog combines the source lists from our own FHD catalog, the MWACS catalog, Culgoora sources, and the MRC, prioritizing in that order. The EoR0 field corresponds to the red FHD patch, while we use the other catalogs to fill in the sidelobes of the MWA. The size of each dot is proportional to the 182 MHz flux of the source, clipped at 20 Jy. Figure courtesy of P.A. Carroll.

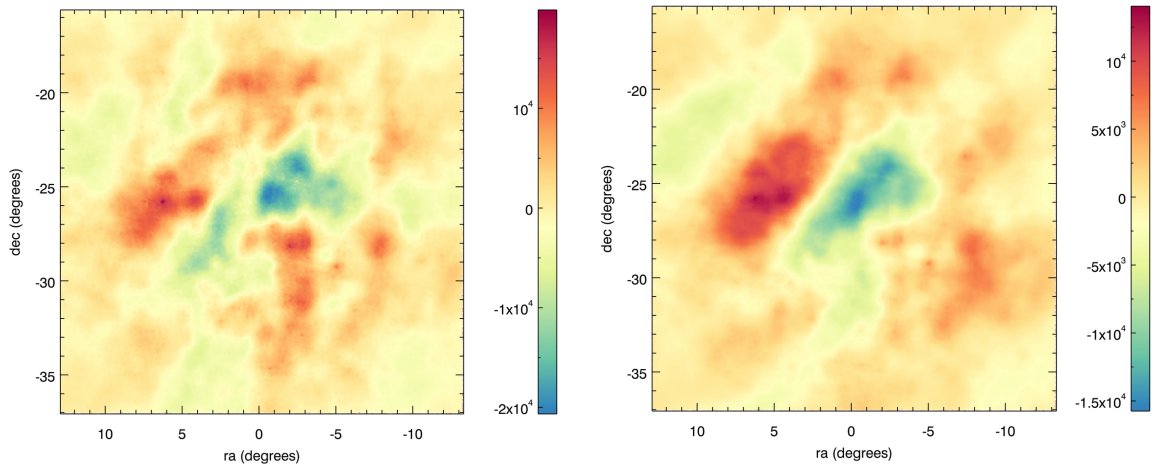

Figure 7.11: *Left:* The diffuse model within the EoR0 field used for foreground subtraction. This model was created using residual images from the golden data set. The image shown is naturally weighted to highlight the large scale structure, but the actual model used for subtraction is uniformly weighted. *Right:* The diffuse model after it has been convolved by a single snapshot point spread function.

of the actual sky (we do not wish to send the image through the instrument twice). After integration we form a pseudo Stokes I image by adding the XX and YY polarizations weighted by the respective beams. The final note is that because we are only using a frequency summed (continuum) image, we are in no danger of subtracting cosmological signal, other than what we have already lost due to foreground contamination.

A naturally weighted version of the diffuse model is shown in Figure 7.11. The model is stored in HEALPix format for easy input to FHD. Although the actual model used is uniformly weighted, I choose to show the naturally weighted version to highlight the galactic cloud structure we are subtracting.

The next step is to import the diffuse model into FHD to form model visibilities. In principle the HEALPix map can be treated as a series of point sources and we could perform

a DFT as we do for the point sources. However, the map has over 400,000 pixels, and would be computationally impossible to DFT for every snapshot. Instead we interpolate the map to a regular sky grid determined in a way to be a Fourier match to the uv grid for each snapshot. After interpolating, we can use an FFT, saving orders of magnitude in computation. Once we have a uv plane for the diffuse model we can integrate to visibilities as usual.

The result of sending the diffuse model through the instrument for a single snapshot is shown in the right panel of Figure 7.11. The point spread function from a single snapshot is significantly more poor than the three hour rotated beam, so the structure has been smoothed out from the left panel.

As has become our custom, we next turn to the power spectrum. The difference between subtracting our diffuse model and not is shown in Figure 7.12. We did not use the diffuse model for calibration, so again the dirty power spectra were identical. We see that our model power spectrum that included the diffuse component contains significantly more power than without, and we see in the residual difference that it subtracted out quite nicely. Note that the color scale is a lot larger than in previous difference plots – this is a major change. Indeed about 90% of the total power in our images has been subtracted when we use the diffuse model. The right panel in Figure 7.12 shows the improvement primarily in the wedge, with seemingly noise dominating the EoR window. Because the diffuse model does have significantly more power in the window (left panel), we expect there to be a positive effect in that region when the noise is integrated down.

We have made many improvements to our foreground subtraction, and I have highlighted a few examples here. Yet there still remains the region at low $k_\perp$ and $k_{||} \approx 0.1$ h Mpc^{-1} where we are dominated by foregrounds and have not been able to subtract. We are still uncertain about exactly what causes this contamination, but the likely suspects are spectral structure in the point sources or the diffuse model, polarization terms in the diffuse, low elevation emission coming from far out sidelobes, inaccuracies in our beam model, or more than likely a combination of all these. Understanding and removing this contamination will be an active topic of research in the near future so that current and next generation EoR

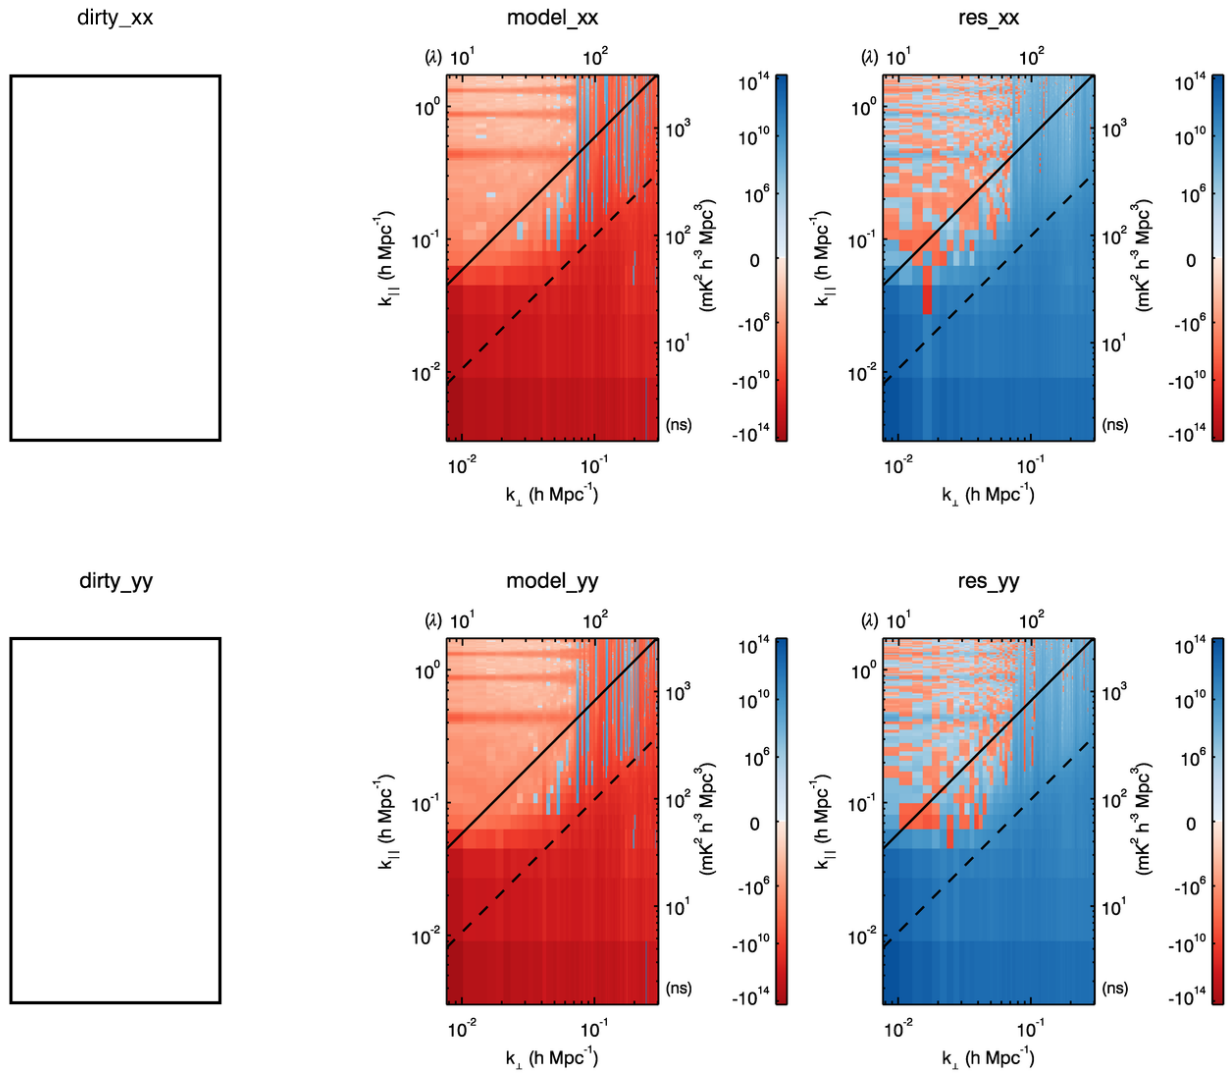

Figure 7.12: Golden data set power spectra differences showing the effect of subtracting a diffuse foreground model. Once again the calibration is the same so the dirty power spectrum differences are zero. The diffuse model adds an extremely large amount of power to the model, and the residual is lower in the wedge indicating that it subtracts successfully.

experiments can realize their full sensitivity.

Chapter 8

EOR POWER SPECTRUM LIMIT

It is now time to combine all that we have learned and apply it to the first semester of EoR data from the MWA. This data set is not nearly large enough to hope to detect the EoR power spectrum, but instead serves as a milestone as the first deep integration with MWA data. Our aim is to rigorously reduce this data such that we maintain a clean EoR window, and potentially identify systematics for future processing. We start by processing the entire data set through the preprocessing, calibration, and imaging steps described in Chapter 6 while incorporating all the improvements described in Chapter 7. The foreground model used for calibration includes points sources both in the main primary beam lobe, but also in the first sidelobes, using the master catalog described in Section 7.2. While not used for calibration, we also use the diffuse model described in the same section for foreground subtraction.

8.1 Data Selection

The data set for this work is all the EoR0 high band snapshots from the first semester of science observations from the MWA. The observing semester starts with the golden data set from August 23, 2013, and extends through to the end of 2013. This consists of 2,780 snapshots, or about 86.5 hours, corresponding to about 125 TB of raw data. However, not nearly all of this data is currently useable for various reasons. We next describe our process of cutting data while retaining high quality data.

8.1.1 Observation Selections

A natural resolution at which to make data cuts is on the two minute snapshot cadence of the telescope. The snapshots are treated independently through the calibration and imaging steps of the pipeline, so we can make cuts at this level without a need to reprocess with FHD (the time dominant component of our analysis). We have employed several methods for selecting data, and I describe them in turn.

First we used a power spectrum jackknife method – by breaking the data into smaller sections we hope to learn patterns about what makes data good or bad. One such powerful slice was to create "per-pointing" power spectra, where we organize the observations into unique day and pointing groups. A pointing is defined as a set of beamformer delays which point the tiles to different regions of the sky. On a given day the MWA tracks the EoR field across several pointings, depicted in Figure 8.1. Our aim with this cut was to determine if certain days or pointings were bad, perhaps due to particularly bad RFI or poor beam models for off-zenith pointings.

While it is impractical to show all the per pointing power spectra we created in this medium, I show a representative example day in Figure 8.2. The most obvious feature shown here is the poor spectra early on in the night. Pointings -5 and -4 are quite terrible and obviously contaminated. The trained eye will notice that pointing -3 is also significantly more contaminated than -2 on. Comparing -3 and +3 we can conclude that this is not simply a beam model issue (the symmetric pointings would be equally troublesome). Instead we have concluded that the contamination we see is due to very strong galactic emission near the horizon early in the night. This is corroborated by the strong line of power seen along the solid black horizon line in pointings -4 and -3. The galactic disk is so terribly bright that even though it is low in our sidelobes, it completely ruins our power spectra.

It is possible that by modeling and subtracting the galaxy near the horizon we may be able to recover this data, but for now it is lost. While it is difficult to see in the two dimensional power spectrum, pointing +4 is also slightly contaminated in the same way,

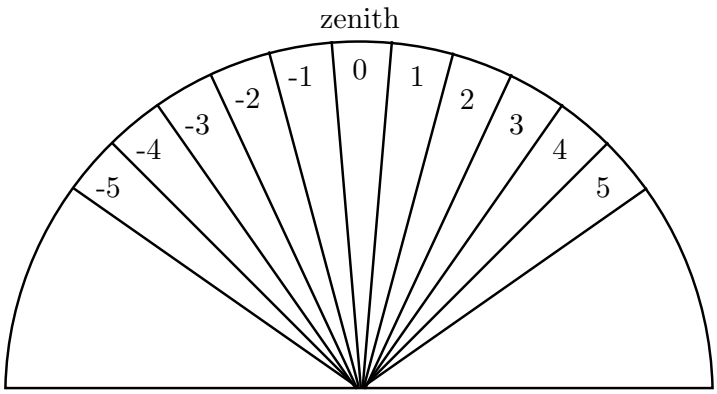

Figure 8.1: A cartoon representation of the MWA pointings in the EoR "drift and shift" observing mode. The MWA will begin tracking the EoR field as early as 5 discrete pointings before zenith, which we call pointing -5. The sky rotates for about thirty minutes before the instrument repoints onto the field. Over the course of a track the instrument will follow the field overhead and down to the large positive pointings.

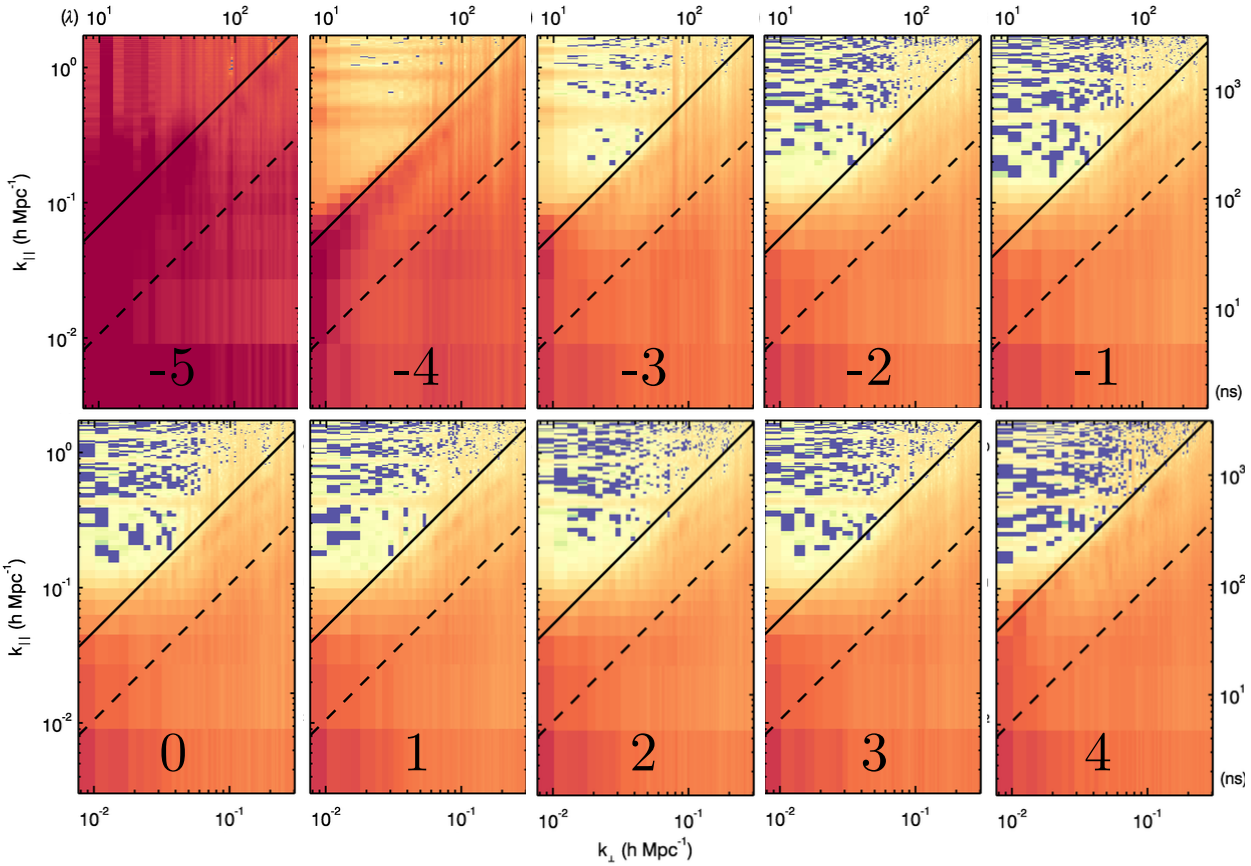

Figure 8.2: An example jackknife test. For this test we divided the data into days and pointings. This is an example array of power spectra (residual XX) for a single day, August 26, 2013. The early pointings are heavily contaminated by the galaxy in the sidelobes, the window becomes more clear near zenith, and we can see trace contamination at the end of the night (pointing +4) when the galaxy has risen again.

due to the galaxy rising at the end of the night. While this contamination is not nearly as offensive as early in the night, there are very few observations in this pointing, and it is better to be on the safe side, so we also cut the +4 pointing.

Already we have seen great power in these jackknives. We have removed data contaminated by bright galactic emission, and what remains is relatively well behaved data. However, we hope to improve on this even further with refined tools.

Motivated by the lessons we learn from jackknife power spectra, we developed a tool to quickly predict the power spectrum quality on an individual snapshot basis. This is done by creating a delay spectrum, which allows us to bypass the computationally expensive gridding step. After forming the delay spectrum, we integrate the total power above the wedge and below the first coarse band harmonic, calling this the total window power. By comparing this with many other statistical metrics we found it is an excellent indicator of the data quality on a snapshot basis. We also found that this metric works whether we calculate it before or after calibration, with a scaling factor relating the two. We show the uncalibrated data here as it will be useful in the future to use as a metric *before* the FHD steps to save computation time. The units of the window power are Jy^2, but because the data is uncalibrated it can instead be thought of as arbitrary units.

The window power for all 2,780 snapshots is shown in the left panel of Figure 8.3. The data is plotted against local sidereal time (LST) in degrees, with the colors representing different observing days, and the vertical dashed lines represent shifts in pointing. We can immediately see the success of this statistic to recreate our choice to remove pointings -5 through -3. We also see a significant number of individual snapshots with excess power that may contribute to contamination. The snapshots with high window power were seen to correlate with poor pointing jackknife power spectra, confirming that this statistic is accurately predicting the actual power spectrum quality. We remove all snapshots with window power greater than 1.2 or less than 0.8. The result is shown in the right panel of Figure 8.3, including the pointing cuts discussed earlier.

Next we consider the window power in different polarizations. In the left panel of Fig-

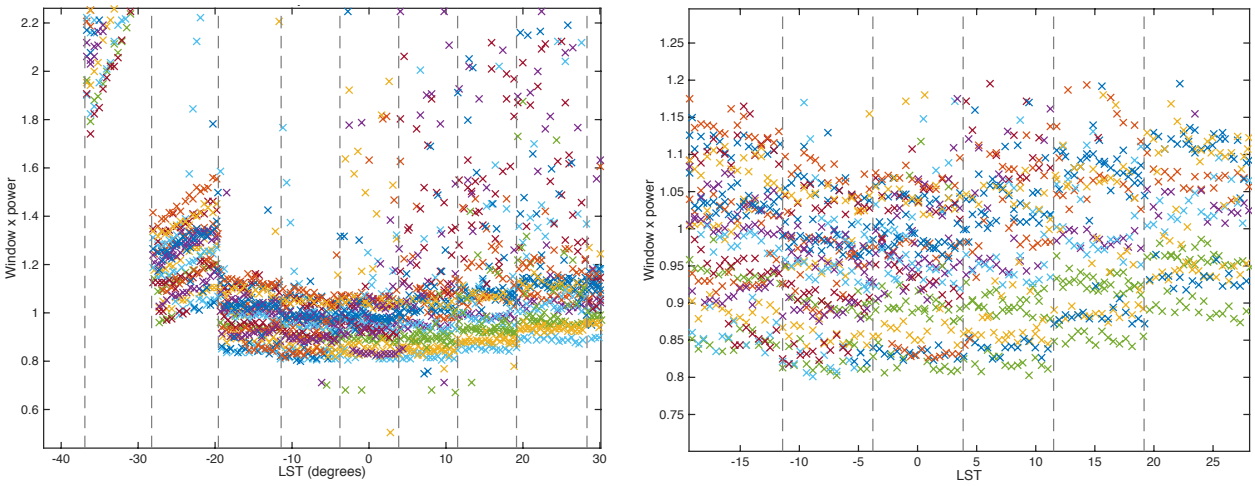

Figure 8.3: Window power for each observation in our data set. The window power is calculated from the delay spectrum of uncalibrated data as a fast quality metric. Each color is a distinct day of observing, and the vertical dashed lines represent pointing shifts. The right panel shows the same data after we have cut snapshots due to high or low window power.

ure 8.4 we plot the ratio of the of window powers in YY and XX polarizations, after the above cuts have been made. We see that pointing +3 has excess power in the Y polarization, potentially due to galaxy emission or RFI in a specific direction and more apparent in Y. Because the different polarizations have different primary beams, especially the sidelobes of off-zenith pointings, it is not surprising that the galaxy may contaminate the sidelobe of one polarization more than another. However, it is likely that if we can see this in one polarization in a two minute snapshot we will likely see it in both polarizations as we integrate down. Again, to be safe, we cut the pointing for both polarizations.

We also trim any snapshots with this ratio greater than 1.13 or less than 1.04. When one polarization has intermittently higher power than the other, it is likely that we are seeing low level RFI that was missed by `COTTER`, or some instrumental issue such as mis-pointings of tiles. We were reassured in this cut by returning to the per-pointing, per-day jackknifes discussed above. By cutting individual observations based on window power ratios, we were able to restore previously bad pointings. This allows us to cut only the contaminated observations without needing to cut a full pointing. The resulting observations are shown in the right panel of Figure 8.4.

We make one final cut from inspecting the residual images output from FHD. We calculate a fractional residual flux RMS in the following way. For each snapshot continuum image, we find the residual flux in the pixels of all sources greater than 0.5 Jy subtracted within half-beam power and scale by the integrated flux of the source. This gives us a list of fractional residual fluxes for the snapshot, which we then use to calculate an RMS. We found that most snapshots had residual flux RMS $< 10\%$ with some outliers, which we will cut from the data. This cut largely overlapped with previous cuts we made, but removed 95 additional observations.

The final snapshot selection includes 1,029 snapshots, or just over 32 hours of data. The integrated 2D power spectrum is shown in Figure 8.5. A few features can be seen in this figure. The first and most worrisome is that the window below the first coarse band harmonic is completely signal dominated. This is likely due to unaccounted for spectral structure of

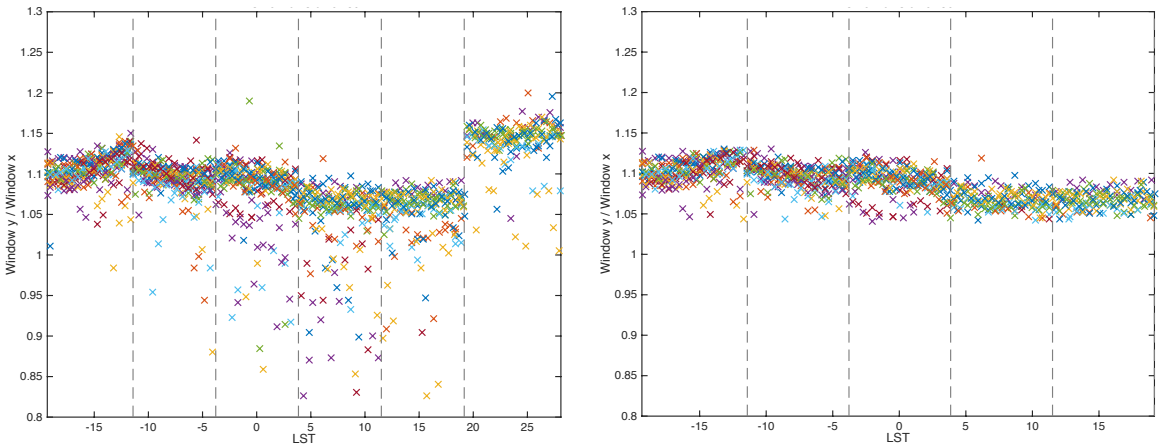

Figure 8.4: Data cut based on window power ratio. For each snapshot that passed the cut shown in Figure 8.3, we plot the ratio of window powers in the YY to XX polarizations. Clear outliers can be seen, including the whole of pointing +3. The right panel shows the remaining observations after this cut.

foregrounds, leaking of galactic emission, or errors in the beam model. Losing these bins to contamination is especially troublesome because the cosmological signal is expected to be largest at low k. Attempting to recover this region of the power spectrum will be a main focus of future analysis.

Another feature we can see in Figure 8.5 is the reemergence of the "fourth line". It is likely that our snapshot based reflection fitting had insufficient signal to noise to fully remove the line. We can study the evolution of the reflection fits for an example tile in Figure 8.6. These plots are drawn from tile 13, but are representative of all tiles with 150 meter cables. The top panel shows the evolution of the reflection amplitude fit over the course of the data set, after all our snapshot based data cuts. We can see that the fits are quite noisy, and decrease slightly over time. A smoothed average may serve to better remove the reflection effect in the future. The bottom panel shows the actual delay mode used in each snapshot. It is not surprising that the noise can cause the selected mode to vary, but the large outliers

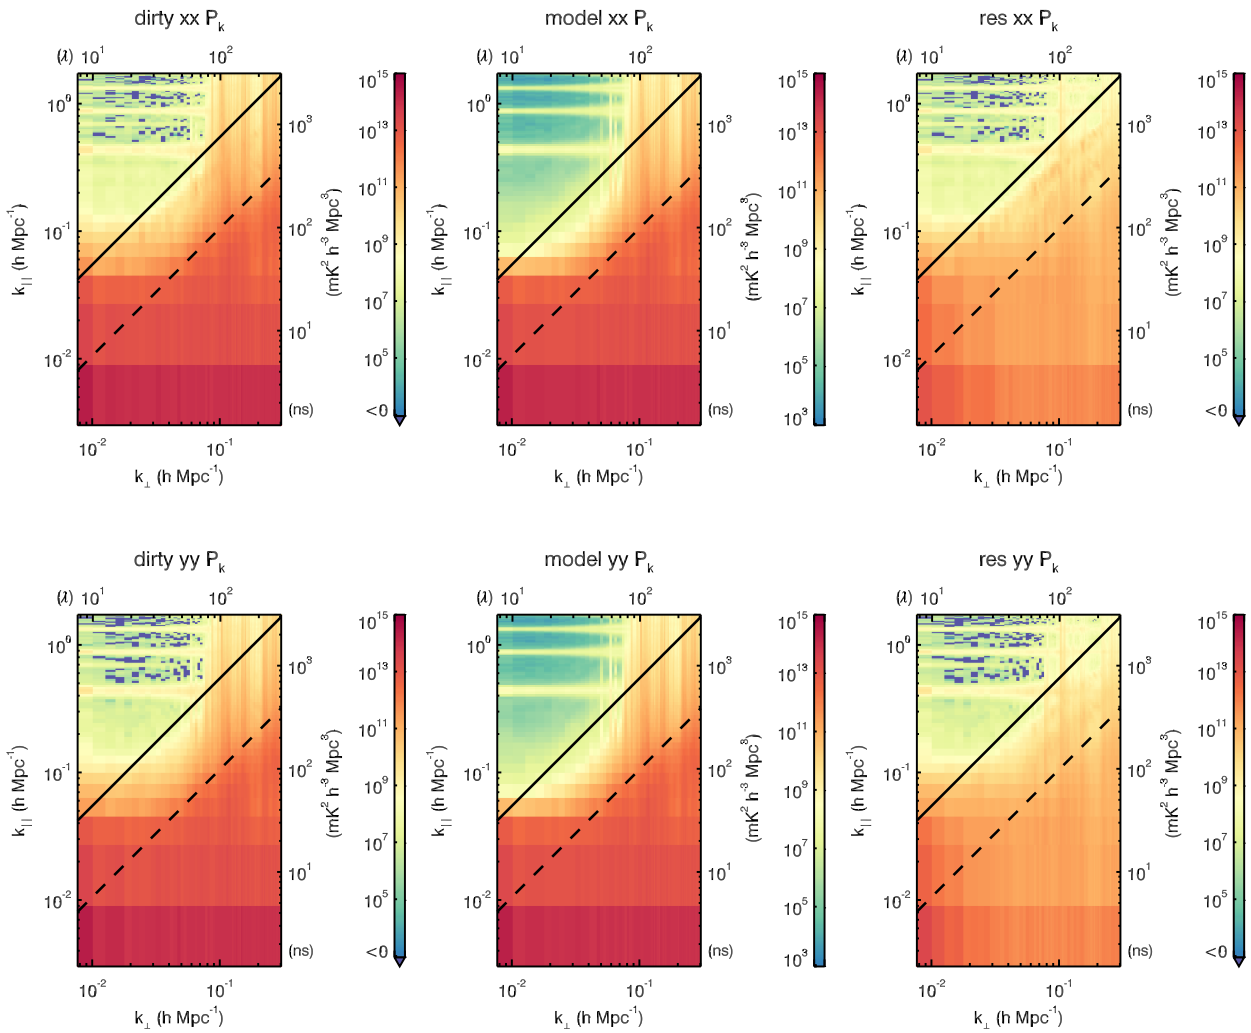

Figure 8.5: The 2D power spectra after all data cuts. This integration includes 1,029 snapshots, or just over 32 hours of data. The lowest region of the window is clearly dominated by some contamination, but noise-like regions remain above the first coarse band harmonic. The fourth line has also reemerged at a lower level.

are of concern. When the noise causes completely different modes to be fit our calibration is doing more harm than good. However, we can see that the true mode does not change over time, so selecting it a priori for future analysis will improve our calibration.

8.1.2 Window Trimming

Now that we have our deep integration, we face the final task of trimming the EoR window to minimize foreground contamination and average to one dimension. The first step is to divide the frequency band. Up until now we have used the full 30.72 MHz bandwidth of the MWA for diagnostic purposes, but a cosmological measurement must be limited in bandwidth to not wash out the signal due to cosmic evolution. We choose to divide the bandwidth into three overlapping sections. This results in three 15.36 MHz bands, but because we implement a Blackman-Harris window function in our frequency transform the effective bandwidth is 7.68 MHz for each, and the overlapping occurs in the wings of the weights (see Figure 8.7). The two dimensional residual power spectra for the three sub-bands are shown in Figure 8.8.

Next we select the region of k space to use in our final averaging to one dimension. Our selection is depicted in Figure 8.9. The lower edge of the window, $k_{||} = 0.1$ h Mpc^{-1}, is a rough estimate of the $k_{||}$ mode where foregrounds are contaminating our signal. In reality this contamination extends all the way to the first coarse band harmonic, but we choose to include part of the contaminated region because the anticipated signal is strongest there and our most stringent limit may very well come from a foreground detection. Because $k_{||} \gg k_{\perp}$ for most pixels in our power spectra, $k \approx k_{||}$, so it does not harm us to include additional $k_{||}$.

The left edge of the window shown in Figure 8.9, $k_{\perp} \approx 0.01$ h Mpc^{-1}, is chosen because of excess power seen at low $k_{\perp}$ both in the low $k_{||}$ modes as well as in the coarse bands. This is likely due to bright galactic or other diffuse emission, so we avoid it by excluding in our average. The upper end of the $k_{\perp}$ range, $k_{\perp} \approx 0.05$ h Mpc^{-1}, is determined by the uv coverage of the MWA. We have seen that when the uv coverage drops off, power is artificially thrown into high $k_{||}$ modes. We have found that this range is ideal for avoiding this effect.

The last exclusion made in Figure 8.9 is the wedge line. We have seen that foregrounds

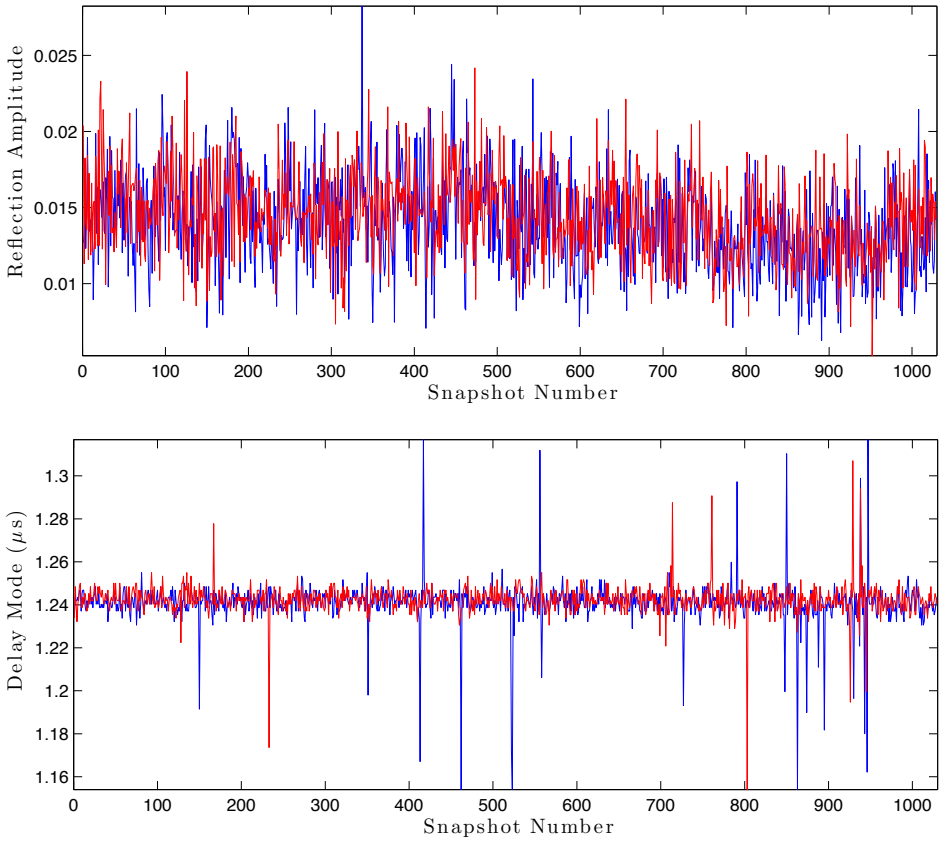

Figure 8.6: Reflection modes fit during deep integration on example tile 13. *Top:* The reflection amplitude fit over the coarse of our observing semester, where the snapshots are simply indexed on the x-axis. The amplitudes are noisy, but relatively steady with a hint of a trend downward at the end of the semester. *Bottom:* The delay mode used for each snapshot. While the mode appears very steady with some jittering noise, there are clear outliers which are likely mis-calibrated. By cutting outliers and using running averages we may be able to improve the fourth line in future analysis runs.

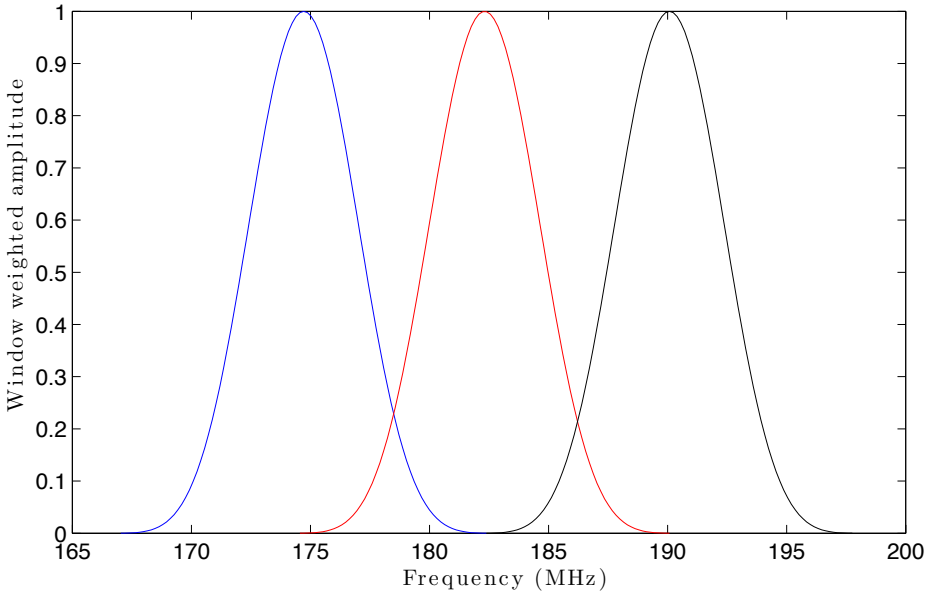

Figure 8.7: Three sub-bands, weighted by the Blackman-Harris window function, used for our cosmological measurement. The full MWA band is divided into three overlapping sub-bands, each of which is weighted by the Blackman-Harris window, shown with the blue line ("low"), red line ("middle"), and black line ("high"). The resulting weighted bands have little overlap and can be considered independent.

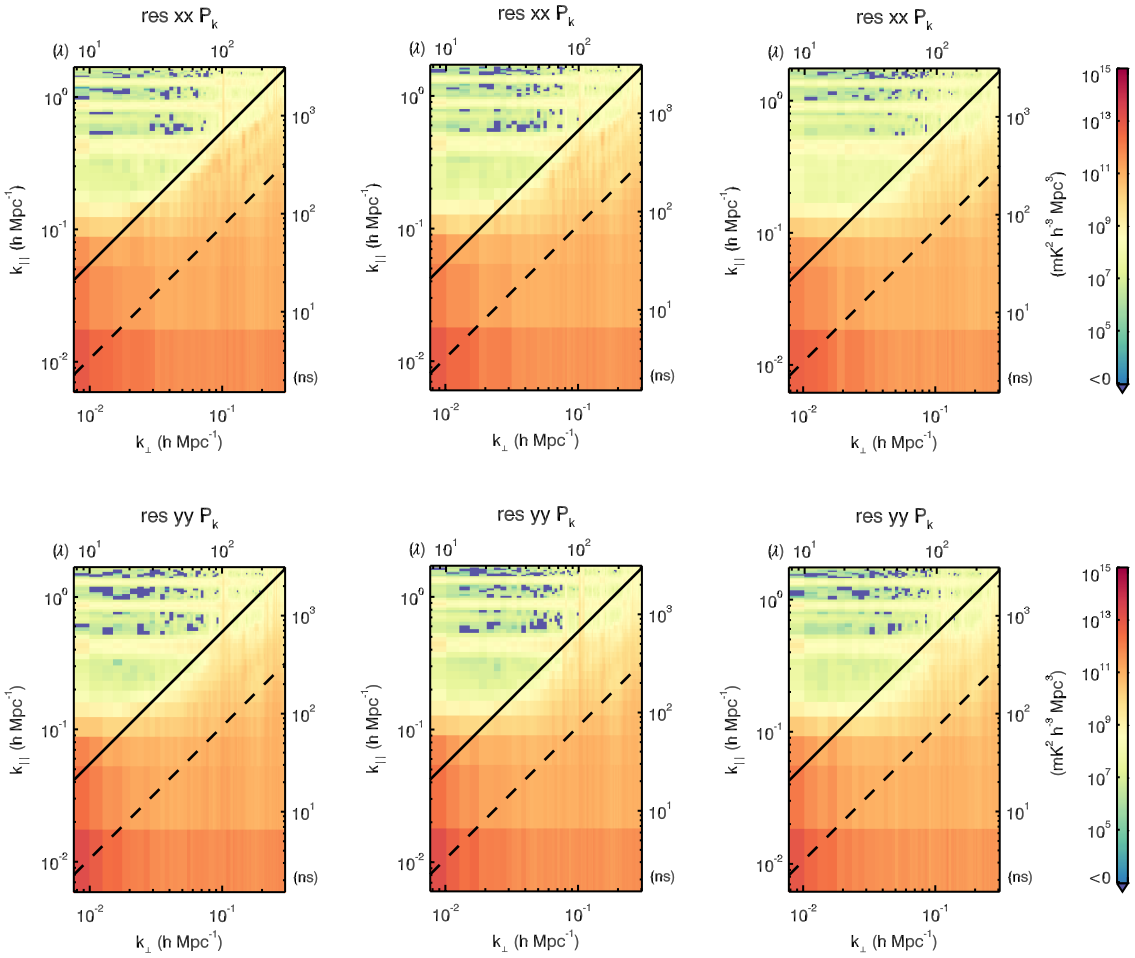

Figure 8.8: Residual power spectra for the three sub-bands used to place limits on the cosmological signal. The left panels show the low band, centered at 174.7 MHz, or a redshift of 7.1. The middle panels show the mid band, centered at 182.4 MHz, or a redshift of 6.8. The right panels show the high band, centered at 190.1 MHz, or a redshift of 6.5.

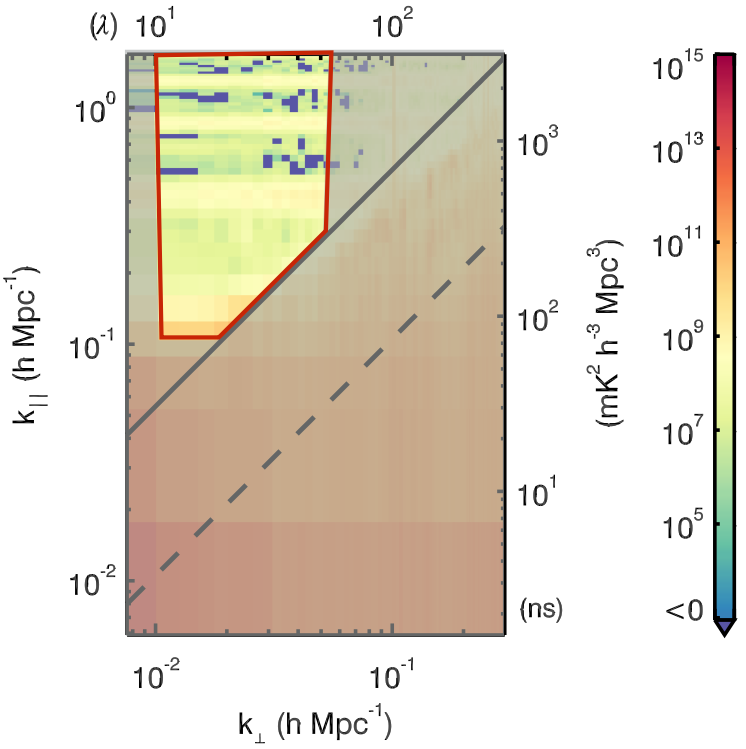

Figure 8.9: Trimming the EoR window. The highlighted region was used when averaging to one dimension. The exclusions are due to various foreground contaminations and low uv coverage at high $k_\perp$. Because the averaging first returns to the 3D power spectrum, and because the different sub-bands have different k-pixels, the highlighted box does not necessarily line up with exact bin edges here. Instead I show where the cut is defined, and used for each power spectrum separately. Bins whose center is within the 3D version of the highlighted region are included in 1D averaging, and bins whose center is outside are omitted.

in the sidelobes, and indeed all the way down to the horizon, are present in the wedge. Therefore we choose to exclude any modes below the solid black horizon wedge line.

While in principle we should exclude the coarse band harmonics, it turns out they will not affect our useful data points in one dimension. Again, because $k_{||} \gg k_{\perp}$ in our space, a spherical shell of constant k is very close to constant $k_{||}$. So the coarse bands remain fairly isolated in the one dimensional average, and including them does not interfere with our best limits.

8.2 A Limit on the Epoch of Reionization

We conclude this chapter by producing a one dimensional power spectrum, and quoting a final upper limit on the EoR signal. Using the window highlighted in Figure 8.9, we average in constant k spherical shells and arrive at our three sub-band 1-D power spectra shown in Figures 8.10 - 8.12.

In Figures 8.10 - 8.12 we plot the quantity $\Delta^2(k) \equiv P(k)k^3/(2\pi^2)$, which has units of mK^2. The advantage of this quantity is that the theoretical EoR signal is approximately flat, which makes for easier comparisons. The actual measured power spectra in these figures is shown with a blue solid step function, where the width of the steps is equal to the k bin size that went into the average. Where the signal becomes negative (which happens when the noise is on order the signal), we show the signal in green (see bottom panel of Figure 8.10, around $k \approx 0.6$ h Mpc^{-1}). The 1-σ noise levels are shown with a dashed red line. The noise is approximately flat in $P(k)$, so it is proportional to k^3 in $\Delta^2(k)$. A reference theoretical power spectrum is shown with a solid black line. We use the models from Lidz et al. [46], with ionizations fractions of $\overline{x}_i = 0.54$ for the low band ($z \sim 7.1$), $\overline{x}_i = 0.82$ for the high band ($z \sim 6.8$), and $\overline{x}_i = 0.96$ for the high band ($z \sim 6.5$). Because the actual redshift of reionization is largely unknown, we also plot with dashed black the reference theoretical model from Furlanetto et al. [26], which we used as a fiducial in our sensitivity calculation in Chapter 4.

The dashed magenta line is the equivalent of a 2-σ upper limit for each k value, which

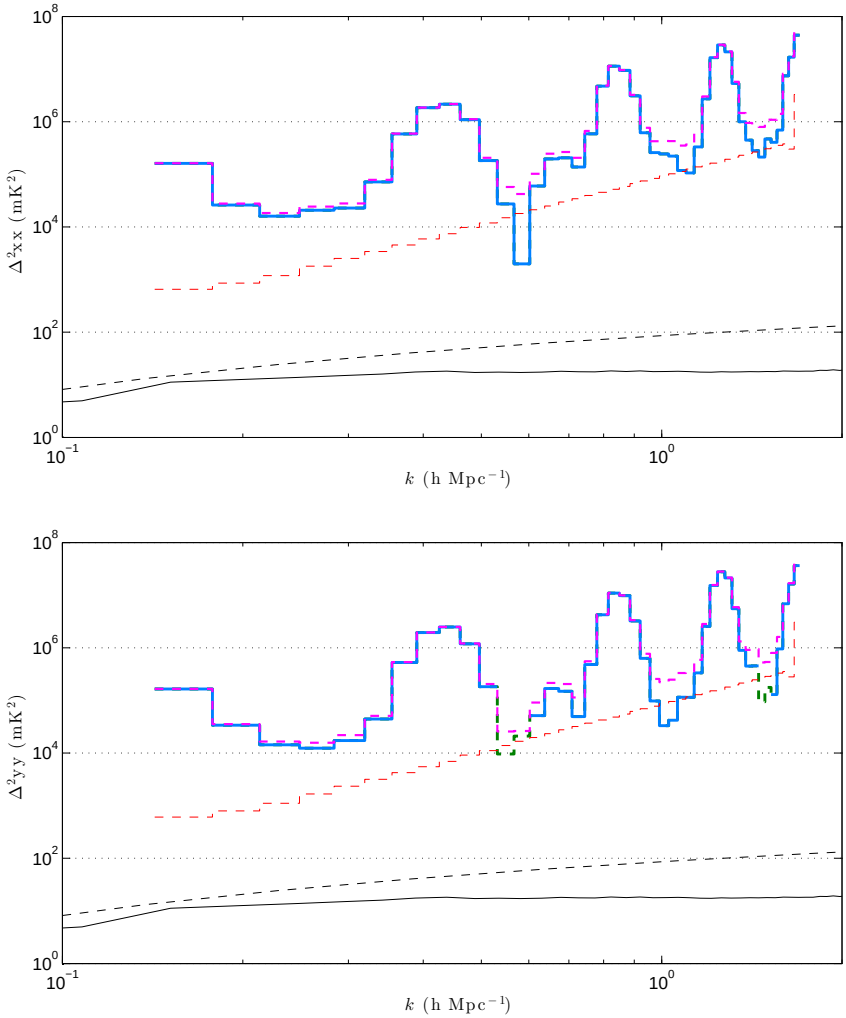

Figure 8.10: One dimensional power spectra for the two polarizations of the low sub-band ($z \approx 7.1$). The solid blue line shows the measured power spectrum with step widths corresponding to the bin size used in the average. Where the measured signal is negative we plot the absolute value in green. The red dashed line is the 1-σ noise level, and the dashed magenta line is the 2-σ upper limit for each k bin. A theoretical model for $\overline{x}_i = 0.54$ from Lidz et al. [46] is shown in solid black. An additional theoretical model for a fully neutral IGM from Furlanetto et al. [26] is shown with a dashed black line.

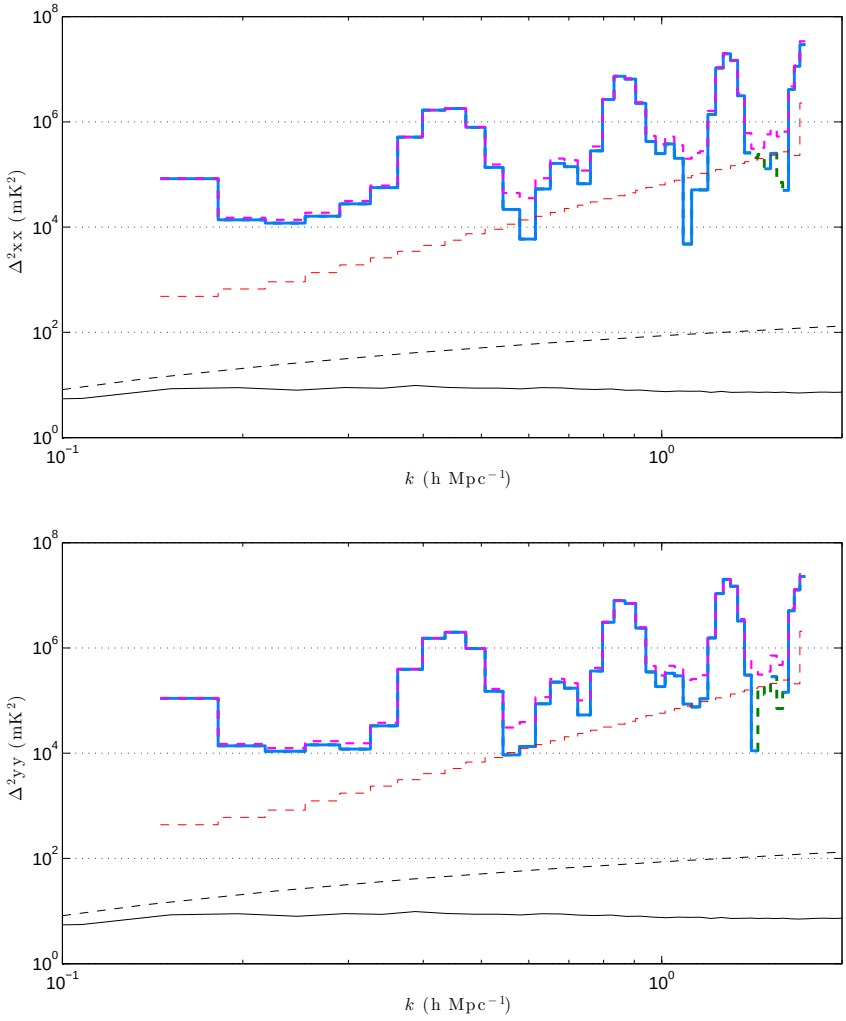

Figure 8.11: One dimensional power spectra for the two polarizations of the middle sub-band ($z \approx 6.8$). The line colors are the same as for Figure 8.10, except that the solid black theoretical model is for $\overline{x}_i = 0.82$.

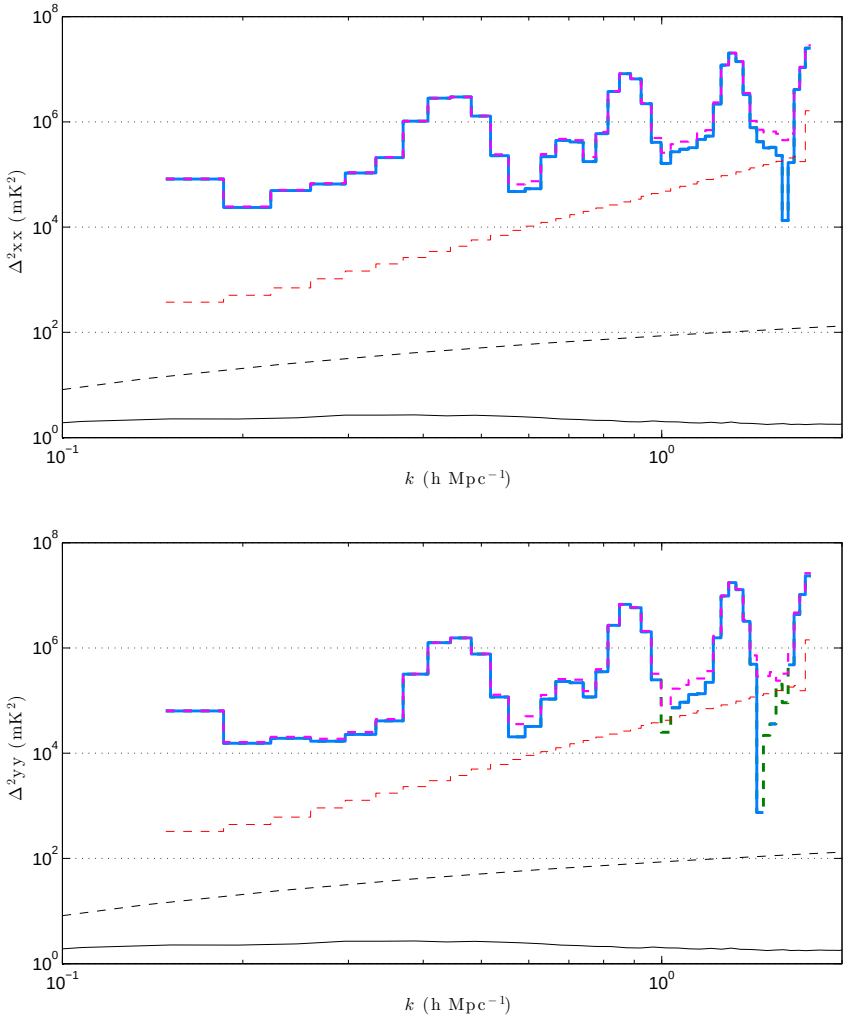

Figure 8.12: One dimensional power spectra for the two polarizations of the high sub-band ($z \approx 6.5$). The line colors are the same as for Figure 8.10, except that the solid black theoretical model is for $\overline{x}_i = 0.96$.

we arrive at using the following argument. First, we are clearly dominated by signal in most bins. Though we certainly do not claim that the signal we detect is cosmological, we can use it as an upper limit to the EoR signal because we do not expect any foregrounds or RFI to be coherent with the cosmological signal, and therefore any contamination should only increase the measured power. So an upper limit on the sky power spectrum is also an upper limit on the cosmological signal. Next we account for the zero-mean noise and the prior that the signal power must be greater than zero. We can express our posterior probability of a true signal, x, in terms of the measured value, x'.

$$\Pr(x|x') = N \Pr(x'|x)\Pr(x) \tag{8.1}$$

Here N is a normalization constant, $\Pr(x|x')$ is the probability that x is the true value, given the measured value x', $\Pr(x'|x)$ is the probability of measuring x' given the true value x, and $\Pr(x)$ is our prior distribution for x. We can express the pieces of the right-hand-side as

$$\Pr(x'|x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x'-x)^2}{2\sigma^2}}, \tag{8.2}$$

where σ^2 is the variance on our measurement, and we have assumed a Gaussian distribution. Although the measured values truly follow the effective Erlang function described in Section 6.3.2, we can approximate it as Gaussian due to the central limit theorem and the relatively large number of pixels that go into our 1-D average. We adopt a flat prior probability for our signal.

$$\Pr(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases} \tag{8.3}$$

We can set our normalization constant by forcing the probability to integrate to unity.

$$\int_{-\infty}^{\infty} \mathrm{d}x \Pr(x|x') = 1 \rightarrow N = \frac{\sqrt{2/(\pi\sigma^2)}}{1 + \mathrm{erf}\left(\frac{x'}{\sqrt{2}\sigma}\right)} \tag{8.4}$$

We can now determine an upper bound given some confidence interval, c. For a 2-σ equivalent, we adopt a confidence interval of $c = 0.977$. To find this upper limit, x_{UL}, we

simply integrate the probability until we reach our confidence level.

$$\int_{-\infty}^{x_{\rm UL}} \mathrm{d}x \, \Pr(x|x') = c \tag{8.5}$$

After some algebra, this yields a limit of

$$x_{\rm UL} = \sqrt{2\sigma^2}\, \mathrm{erf}^{-1}\left(c - (1-c)\mathrm{erf}\left(\frac{x'}{\sqrt{2\sigma^2}}\right)\right) + x', \tag{8.6}$$

which we use to plot the dashed magenta line in Figures 8.10 & 8.12. In the regions where we are signal dominated this detailed expression converges to simply adding 2-σ to the detected value, which makes little difference when our SNR $\gg 1$, and the magenta line closely follows the blue. However, it does have an effect for the SNR ~ 1 bins, and this correctly accounts for our prior that the signal is greater than 0.

There are many interesting features in our 1-D power spectra figures. The coarse band harmonics are clearly present at their expected regular intervals. We can also see the fourth line on the left shoulder of the second coarse band. But the most important things to observe are where our best limit is, and where we have the greatest opportunity for improvement. The best limit from each sub-band and polarization comes from low k, just before the increase of the first coarse band. In all three sub-bands the YY polarization has a slightly lower contamination, which may be a clue to how to isolate and remove it. This region is clearly signal dominated, so longer integrations will not improve the limit. A summary of the best limits from this data is presented in Table 8.1.

However, there is a small region where we are noise limited and reasonably close to matching the limit before the first coarse band. This is the narrow window between the first coarse band and the fourth line. We may be able to widen this window with improved calibrations on the reflection modes. In addition efforts are already underway to utilize inverse covariance weighting techniques to suppress the contamination due to the coarse band gaps, which may also help to widen this window. With another order of magnitude more data already observed, we can hope to improve the limit in this window without much effort beyond the effort of actually performing the analysis. However, the data volume will

Table 8.1: Upper limits on the EoR power spectrum for our three sub-bands and two polarizations. Upper limits, $\Delta^2_{\rm UL}$, are at 97.7% confidence level.

Sub-band	z_0	Polarization	k (h Mpc^{-1})	$\Delta^2_{\rm UL}$ (mK^2)
Low	7.1	XX	0.231	1.84×10^4
Low	7.1	YY	0.266	1.57×10^4
Mid	6.8	XX	0.236	1.37×10^4
Mid	6.8	YY	0.236	1.25×10^4
High	6.5	XX	0.204	2.46×10^4
High	6.5	YY	0.204	1.63×10^4

eventually become a major obstacle when considering the requirement to beat the noise down to the theoretical level.

We therefore see two major strategies for proceeding and improving the cosmological limits from the MWA. One is to further isolate and remove contamination from foregrounds to unlock the sensitivity potential of low-k modes. The other is to continue integration with the expectation that the noise-like modes will continue to integrate down. The latter can also be improved with some minor analysis improvements. However it is my opinion that in the interest of long-term advances the focus of energy should be placed on unlocking the low-k modes. While there may be larger obstacles to overcome, especially because the source of the contamination is yet unknown, the potential for an eventual detection is far greater in these modes. Furthermore, there is a good chance that the contamination we see at low-k is also leaking into higher k modes, but at a level such that we are yet insensitive to it. There may be opportunity for a short term improvement at the larger k bins, but a detection will be greatly dependent on solving our problems at low-k.

Finally we place our upper limit in context with published best limits on the EoR power spectrum from the MWA [23], PAPER [1, 37, 67], and the GMRT [65]. Each 2-σ upper

Table 8.2: Current best published 2-σ upper limits on the EoR power spectrum.

Reference	Instrument	z	k (h Mpc^{-1})	$\Delta^2_{\rm UL}$ (mK^2)
Beardsley, 2015	MWA	7.1	0.266	1.57×10^4
Beardsley, 2015	MWA	6.8	0.236	1.25×10^4
Beardsley, 2015	MWA	6.5	0.204	1.63×10^4
Dillon and others [23]	MWA	6.4	0.18	3.80×10^4
Dillon and others [23]	MWA	6.8	0.18	3.69×10^4
Dillon and others [23]	MWA	7.25	0.16	4.67×10^4
Ali et al. [1]	PAPER	8.4	0.325	5.02×10^2
Jacobs et al. [37]	PAPER	10.29	0.2	1.64×10^4
Jacobs et al. [37]	PAPER	8.54	0.1	6.90×10^3
Jacobs et al. [37]	PAPER	7.94	0.2	6.90×10^3
Jacobs et al. [37]	PAPER	7.55	0.2	2.40×10^3
Parsons et al. [67]	PAPER	7.7	0.27	1.68×10^3
Paciga et al. [65]	GMRT	8.6	0.5	6.15×10^4

limit is shown in Table 8.2[1]. Each experiment probes different scales and redshifts, but we can glean the status of the field with a rough comparison. Because the theoretical power spectrum is somewhat flat in $\Delta^2(k)$ and because the k values of the different experiments are within an order of magnitude, we can approximate them as probing the same scales. We plot these limits as a function of redshift in Figure 8.13, with an additional theoretical model from Lidz et al. [46] for reference.

Our results are a marginal improvement over the previous best results from the MWA.

[1]The limits in Jacobs et al. [37] were expressed as 1-σ upper limits, however through personal communication we have noticed that there was a typo in the article, and they should have been quoted as 2-σ. In other words, the limits are correct, but are actually 2-σ, not 1-σ. This will be corrected in an erratum soon.

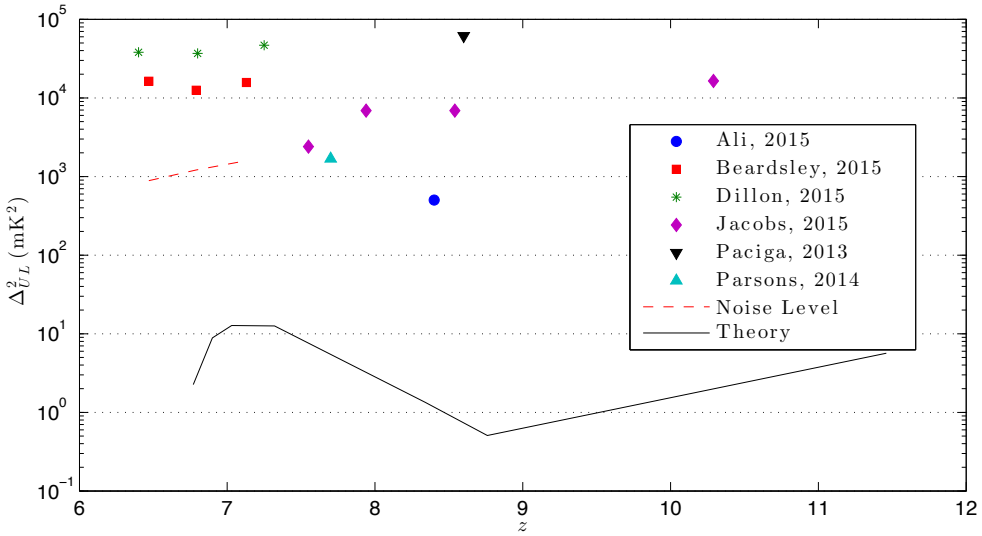

Figure 8.13: The current best published upper limits on the EoR power spectrum, shown as 2-σ upper limits as a function of redshift. The red dashed line shows the 2-σ noise level for our current integration – the best limit we can produce if we remove systematics through improved analysis. For reference we show the theoretical model from Lidz et al. [46], and at the scale $k = 0.2$ h Mpc^{-1}.

This is not surprising because the previous results were based on the same golden data set that is a subset of our current data, and their power spectrum pipeline used as an input the FHD image cubes generated for this work. Using the same data, and because we are significantly signal dominated, it is expected that the upper limits should be similar. Our results are also consistent with the GMRT limit, though probe significantly different periods of the EoR – the latter probing what is expected to be the very beginning of the "rise" of the EoR. The clear frontrunner in power spectrum constraints is the PAPER experiment, specifically Ali et al. [1], which has in turn been able to place scientific constraints on the kinetic temperature of the IGM and X-ray heating models of the EoR [73]. While the MWA analysis is not yet at that level of maturity, there is much hope for rapid improvement. If we can overcome our systematic limits and drive our limit to the noise, already in 32 hours of observations we will be competitive with PAPER. In addition the MWA has many more hours of data already collected, including a separate "low" band, and two other EoR fields. There is lots of room for improvement which will quickly make the MWA limits very competitive.

Chapter 9
CONCLUSIONS

"When one writes a novel about grown people, he knows exactly where to stop - that is, with a marriage; but when he writes about juveniles, he must stop where he best can."

– Mark Twain, *The Adventures of Tom Sawyer*

If the ultimate goal of the MWA 21cm EoR experiment is the detection of the cosmological power spectrum, then we can think of the current state as its juvenile period. While the telescope itself has collected hundreds of hours of data, we as a collaboration are coming of age through learning about the analysis of this data and overcoming challenges. Many of the obstacles were foreseen and built into the design of the instrument (e.g. Chapter 3), while others were discovered in the early stages of analysis (e.g. Chapter 7). The unique challenges involved in our ambitious goal in realizing the sensitivity of the MWA (Chapter 4) required a complete reevaluation of traditional analysis tools, and for several components to be built from scratch (Chapter 6).

While we have not yet reached the goal of an EoR detection, the results from the first season of observing is a sensible place to end this particular story. We have taken the MWA EoR experiment from concept to a realized telescope, and developed an analysis pipeline that placed an upper limit on the EoR power spectrum with 32 hours of data. Our 2-σ limit of $\Delta^2 < 1.57 \times 10^4$ mK^2 at $z = 7.1$, $k = 0.266$ h Mpc^{-1} and $\Delta^2 < 1.63 \times 10^4$ mK^2 at $z = 6.5$, $k = 0.204$ h Mpc^{-1} represents the best limit from the MWA to date. While there are other more stringent limits in the field as a whole, our limit is an excellent starting place to continue improvement of the MWA analysis.

More broadly speaking, our limit is the deepest to date produced by an imaging pipeline.

The imaging style of analysis has had many obstacles to overcome (e.g. efficient gridding and mapmaking, foreground modeling), but has potential to extend beyond the capabilities of more pointed analysis styles such as the delay spectrum. For example, we are seeing that the primary beam models will quickly become a major limiting factor in continuing to integrate down, and soon we will require individual models per antenna. This is a fundamental limitation of delay spectrum analysis as it is incapable of incorporating non-identical beams, whereas this feature is built into our imaging pipeline. Ultimately we believe the imaging analysis will be necessary to overcome many more challenges to come – some expected like polarized foregrounds, some unforeseen. The real value in this work is pushing forward on this analysis technique, and showing the potential for it to be the leading method in power spectrum detection.

We have outlined two strategies for improving our limit: to identify and remove the contamination at the low end of the EoR window, allowing our low-k modes to integrate down with the noise; or to leverage the higher-k modes which are currently noise-like and should integrate down with more data and modest analysis improvements. While the latter strategy is likely to yield quicker results, the former is ultimately necessary to realize the full potential of the instrument and make an eventual detection. The MWA EoR experiment was built with the consideration that the signal is strongest at large scaled (small k), and that a tightly packed layout will yield the highest sensitivity to these modes. We are faced with obstacles in the path to utilize these modes, but will continue to improve our understanding and analysis to overcome these challenges.

9.1 Future Directions

While our analysis efforts are ongoing the instruments themselves are moving forward. The MWA currently has plans to increase the number of tiles to 256, and upgrade the front end signal chain to digital (which will, amongst other benefits, remove the nasty gaps in our bandpass). The additional tiles will be placed with a significantly different strategy from the original 128 tiles. Instead of an imaging array, the EoR-dedicated additional tiles will

be placed in two highly redundant hexagonal cores. The resulting array will be a hybrid between the highly redundant strategy used by PAPER, and the imaging strategy used by the current MWA. This will allow for interesting experiments such as direct comparison between the current calibration algorithm and a redundant calibration which is fast and requires no sky model [47]. Retaining the non-redundant baselines of the original MWA will also allow us to maintain a clean PSF for foreground removal.

The HERA experiment is also in the prototyping stage, with preliminary funding to build a subarray which will eventually become a part of the full 331 dish HERA. The large dishes (14 m diameter) will allow for faster sensitivity, but also better rejection of unwanted signal outside the primary field of view. We have seen with the MWA that foregrounds in the sidelobes, even down to the horizon, are the limiting factor for opening the EoR window as far as possible, so a more pointed response function should aid us in realizing sensitivity.

With several international efforts to characterize the EoR using the 21cm line, it is important to also leverage opportunities to combine these measurements with experiments at other wavelengths. I have outlined one such strategy in Chapter 5 to combine the observations of MWA and HERA with galaxy surveys from JWST. While these different experiments probe separate tracers and seemingly disparate scales, being able to combine them will ultimately provide deeper and more significant insights into the processes that drove the EoR.

While EoR experiments are beginning to mature, 21cm tomography as a broad field is in its infancy. There remains many avenues to explore our cosmos through the 3D mapping of the neutral hydrogen in and between galaxies. The cosmic Dark Ages will be explored first with the global 21cm signal by experiments like EDGES [11], LEDA [30], and DARE [15]. The EoR continues to be pursued with the MWA, PAPER, LOFAR, GMRT, and future arrays such as HERA and SKA Low. And the properties of Dark Energy can be probed through BAOs by experiments like CHIME [80]. Each of these experiments will push the limits of radio observations, improving technology and analysis tools along the way. As the techniques are refined and the experiments mature we will be afforded a brand new look into the evolution of the cosmos through the eyes of the emergent field of 21cm tomography.

BIBLIOGRAPHY

[1] 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 21cm Power Spectrum at z=8.4. *ArXiv e-prints*, February 2015.

[2] Rennan Barkana and Abraham Loeb. A method for separating the physics from the astrophysics of high-redshift 21 centimeter fluctuations. *The Astrophysical Journal Letters*, 624(2):L65, 2005. URL `http://stacks.iop.org/1538-4357/624/i=2/a=L65`.

[3] A. P. Beardsley, B. J. Hazelton, M. F. Morales, R. J. Capallo, R. Goeke, D. Emrich, C. J. Lonsdale, W. Arcus, D. Barnes, G. Bernardi, J. D. Bowman, J. D. Bunton, B. E. Corey, A. Deshpande, L. deSouza, B. M. Gaensler, L. J. Greenhill, D. Herne, J. N. Hewitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, 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. Subrahmanyan, S. J. Tingay, R. B. Wayth, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. 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. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2012.20878.x. URL `http://dx.doi.org/10.1111/j.1365-2966.2012.20878.x`.

[4] A. P. Beardsley, B. J. Hazelton, M. F. Morales, W. Arcus, D. Barnes, G. Bernardi, J. D. Bowman, F. H. Briggs, 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, Shankar N. Udaya, K. S. Srivani, J. Stevens, R. Subrahmanyan, S. J. Tingay, R. B. Wayth, M. Waterson, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The eor sensitivity of the murchison widefield array. *Monthly Notices of the Royal Astronomical Society: Letters*, 429(1):L5–L9, 2013. doi: 10.1093/mnrasl/sls013. URL `http://mnrasl.oxfordjournals.org/content/429/1/L5.abstract`.

[5] A. P. Beardsley, M. F. Morales, A. Lidz, M. Malloy, and P. M. Sutter. Adding context to james webb space telescope surveys with current and future 21 cm radio observations. *The Astrophysical Journal*, 800(2):128, 2015. URL `http://stacks.iop.org/0004-637X/800/i=2/a=128`.

[6] C. L. Bennett, D. Larson, J. L. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. M. Smith, R. S. Hill, B. Gold, M. Halpern, E. Komatsu, M. R. Nolta, L. Page, D. N. Spergel, E. Wollack, J. Dunkley, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, and E. L. Wright. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results. ApJS, 208:20, October 2013. doi: 10.1088/0067-0049/208/2/20.

[7] G. Bernardi, D. A. Mitchell, S. M. Ord, L. J. Greenhill, B. Pindor, R. B. Wayth, and J. S. B. Wyithe. Subtraction of point sources from interferometric radio images through an algebraic forward modelling scheme. *Monthly Notices of the Royal Astronomical Society*, 413(1):411–422, 2011. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2010.18145.x. URL `http://dx.doi.org/10.1111/j.1365-2966.2010.18145.x`.

[8] Bhatnagar, S., Cornwell, T. J., Golap, K., and Uson, J. M. Correcting direction-

dependent gains in the deconvolution of radio interferometric images. *A&A*, 487(1): 419–429, 2008. doi: 10.1051/0004-6361:20079284. URL http://dx.doi.org/10.1051/0004-6361:20079284.

[9] F. Boone. Interferometric array design: Interferometric array design: Distributions of fourier samples for imaging. *Astronomy and Astrophysics*, 386(3):1160–1171, 2002.

[10] R. J. Bouwens, L. Bradley, A. Zitrin, D. Coe, M. Franx, W. Zheng, R. Smit, O. Host, M. Postman, L. Moustakas, I. Labbé, M. Carrasco, A. Molino, M. Donahue, D. D. Kelson, M. Meneghetti, N. Benítez, D. Lemze, K. Umetsu, T. Broadhurst, J. Moustakas, P. Rosati, S. Jouvel, M. Bartelmann, H. Ford, G. Graves, C. Grillo, L. Infante, Y. Jimenez-Teja, O. Lahav, D. Maoz, E. Medezinski, P. Melchior, J. Merten, M. Nonino, S. Ogaz, and S. Seitz. A Census of Star-forming Galaxies in the Z ˜ 9-10 Universe based on HST+Spitzer Observations over 19 Clash Clusters: Three Candidate Z ˜ 9-10 Galaxies and Improved Constraints on the Star Formation Rate Density at Z ˜ 9.2. ApJ, 795:126, November 2014. doi: 10.1088/0004-637X/795/2/126.

[11] Judd D. Bowman and Alan E. E. Rogers. A lower limit of [dgr]z[thinsp]>[thinsp]0.06 for the duration of the reionization epoch. *Nature*, 468(7325):796–798, 12 2010. URL http://dx.doi.org/10.1038/nature09601.

[12] Judd D Bowman, Miguel F Morales, and Jacqueline N Hewitt. The sensitivity of first-generation epoch of reionization observatories and their potential for differentiating theoretical power spectra. *The Astrophysical Journal*, 638:20, Jan 2006. doi: 10.1086/498703. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006ApJ...638...20B&link_type=ABSTRACT. (c) 2006: The American Astronomical Society.

[13] Judd D. Bowman, Miguel F. Morales, and Jacqueline N. Hewitt. Foreground contamination in interferometric measurements of the redshifted 21 cm power spectrum. *The Astrophysical Journal*, 695(1):183, 2009. URL http://stacks.iop.org/0004-637X/695/i=1/a=183.

[14] Judd D. Bowman, Iver Cairns, David L. Kaplan, Tara Murphy, Divya Oberoi, Lister Staveley-Smith, Wayne Arcus, David G. Barnes, Gianni Bernardi, Frank H. Briggs, Shea Brown, John D. Bunton, Adam J. Burgasser, Roger J. Cappallo, Shami Chatterjee, Brian E. Corey, Anthea Coster, Avinash Deshpande, Ludi deSouza, David Emrich, Philip Erickson, Robert F. Goeke, B. M. Gaensler, Lincoln J. Greenhill, Lisa Harvey-Smith, Bryna J. Hazelton, David Herne, Jacqueline N. Hewitt, Melanie Johnston-Hollitt, Justin C. Kasper, Barton B. Kincaid, Ronald Koenig, Eric Kratzenberg, Colin J. Lonsdale, Mervyn J. Lynch, Lynn D. Matthews, S. Russell McWhirter, Daniel A. Mitchell, Miguel F. Morales, Edward H. Morgan, Stephen M. Ord, Joseph Pathikulangara, Thiagaraj Prabu, Ronald A. Remillard, Timothy Robishaw, Alan E. E. Rogers, Anish A. Roshi, Joseph E. Salah, Robert J. Sault, N. Udaya Shankar, K. S. Srivani, Jamie B. Stevens, Ravi Subrahmanyan, Steven J. Tingay, Randall B. Wayth, Mark Waterson, Rachel L. Webster, Alan R. Whitney, Andrew J. Williams, Christopher L. Williams, and J. Stuart B. Wyithe. Science with the murchison widefield array. *PASA - Publications of the Astronomical Society of Australia*, 30, 1 2013. ISSN 1448-6083. doi: 10.1017/pas.2013.009. URL `http://journals.cambridge.org/article_S132335801300009X`.

[15] 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.

[16] E. Chapman, F. B. Abdalla, J. Bobin, J.-L. Starck, G. Harker, V. Jelić, P. Labropoulos, S. Zaroubi, M. A. Brentjens, A. G. de Bruyn, and L. V. E. Koopmans. The scale of the problem: recovering images of reionization with Generalized Morphological Component Analysis. MNRAS, 429:165–176, February 2013. doi: 10.1093/mnras/sts333.

[17] T. R. Choudhury, E. Puchwein, M. G. Haehnelt, and J. S. Bolton. Lyman-α emitters gone missing: evidence for late reionization? *ArXiv e-prints*, December 2014.

[18] Babak Cohanim, Jacqueline Hewitt, and Olivier De Weck. The design of radio telescope array configurations using multiobjective optimization: Imaging performance versus cable length. *The Astrophysical Journal Supplement Series*, 154:705–705–719–719, Jan 2010. doi: doi:10.1086/422356. URL `http://adsabs.harvard.edu/abs/2004ApJS..154..705C`.

[19] T. J. Cornwell. A novel principle for optimization of the instantaneous Fourier plane coverage of correlation arrays. *IEEE Transactions on Antennas and Propagation*, 36: 1165–1167, August 1988. doi: 10.1109/8.7233.

[20] T. J. Cornwell, K. Golap, and S. Bhatnagar. The Noncoplanar Baselines Effect in Radio Interferometry: The W-Projection Algorithm. *IEEE Journal of Selected Topics in Signal Processing*, 2:647–657, November 2008. doi: 10.1109/JSTSP.2008.2005290.

[21] A. Datta, J. D. Bowman, and C. L. Carilli. Bright source subtraction requirements for redshifted 21 cm measurements. *The Astrophysical Journal*, 724(1):526, 2010. URL `http://stacks.iop.org/0004-637X/724/i=1/a=526`.

[22] 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.

[23] J. S. Dillon and others. 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, in review*, 2015.

[24] X. Fan, M. A. Strauss, R. H. Becker, R. L. White, J. E. Gunn, G. R. Knapp, G. T. Richards, D. P. Schneider, J. Brinkmann, and M. Fukugita. Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with z˜6 Quasars. II. A Sample of 19 Quasars. *AJ*, 132:117–136, July 2006. doi: 10.1086/504836.

[25] Steven R. Furlanetto, Matias Zaldarriaga, and Lars Hernquist. The growth of h ii regions during reionization. *The Astrophysical Journal*, 613(1):1, 2004. URL `http://stacks.iop.org/0004-637X/613/i=1/a=1`.

[26] Steven R Furlanetto, S. Peng Oh, and Frank H Briggs. Cosmology at low frequencies: The 21 cm transition and the high-redshift universe. *Physics Reports*, 433:181, Sep 2006. doi: 10.1016/j.physrep.2006.08.002. URL `http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006PhR...433..181F&link_type=ABSTRACT`. Elsevier B.V.

[27] 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.

[28] M. Giavalisco, H. C. Ferguson, A. M. Koekemoer, M. Dickinson, D. M. Alexander, F. E. Bauer, J. Bergeron, C. Biagetti, W. N. Brandt, S. Casertano, C. Cesarsky, E. Chatzichristou, C. Conselice, S. Cristiani, L. Da Costa, T. Dahlen, D. de Mello, P. Eisenhardt, T. Erben, S. M. Fall, C. Fassnacht, R. Fosbury, A. Fruchter, J. P. Gardner, N. Grogin, R. N. Hook, A. E. Hornschemeier, R. Idzi, S. Jogee, C. Kretchmer, V. Laidler, K. S. Lee, M. Livio, R. Lucas, P. Madau, B. Mobasher, L. A. Moustakas, M. Nonino, P. Padovani, C. Papovich, Y. Park, S. Ravindranath, A. Renzini, M. Richardson, A. Riess, P. Rosati, M. Schirmer, E. Schreier, R. S. Somerville, H. Spinrad, D. Stern, M. Stiavelli, L. Strolger, C. M. Urry, B. Vandame, R. Williams, and C. Wolf. The Great Observatories Origins Deep Survey: Initial Results from Optical and Near-Infrared Imaging. ApJ, 600:L93–L98, January 2004. doi: 10.1086/379232.

[29] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke,

and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. ApJ, 622:759–771, April 2005. doi: 10.1086/427976.

[30] L. J. Greenhill and G. Bernardi. HI Epoch of Reionization Arrays. *ArXiv e-prints*, January 2012.

[31] E.W. Greisen. AIPS FITS File Format. AIPS Memo 117, 2012.

[32] J. E. Gunn and B. A. Peterson. On the Density of Neutral Hydrogen in Intergalactic Space. ApJ, 142:1633–1641, November 1965. doi: 10.1086/148444.

[33] Bryna J. Hazelton, Miguel F. Morales, and Ian S. Sullivan. The fundamental multi-baseline mode-mixing foreground in 21 cm epoch of reionization observations. *The Astrophysical Journal*, 770(2):156, 2013. URL `http://stacks.iop.org/0004-637X/770/i=2/a=156`.

[34] D. W. Hogg. Distance measures in cosmology. *ArXiv Astrophysics e-prints*, May 1999.

[35] N. Hurley-Walker, J. Morgan, R. B. Wayth, P. J. Hancock, M. E. Bell, G. Bernardi, R. Bhat, F. Briggs, A. A. Deshpande, A. Ewall-Wice, L. Feng, B. J. Hazelton, L. Hindson, D. C. Jacobs, D. L. Kaplan, N. Kudryavtseva, E. Lenc, B. McKinley, D. Mitchell, B. Pindor, P. Procopio, D. Oberoi, A. Offringa, S. Ord, J. Riding, J. D. Bowman, R. Cappallo, B. Corey, D. Emrich, B. M. Gaensler, R. Goeke, L. Greenhill, J. Hewitt, M. Johnston-Hollitt, J. Kasper, E. Kratzenberg, C. Lonsdale, M. Lynch, R. McWhirter, M. F. Morales, E. Morgan, T. Prabu, A. Rogers, A. Roshi, U. Shankar, K. Srivani, R. Subrahmanyan, S. Tingay, M. Waterson, R. Webster, A. Whitney, A. Williams, and C. Williams. The Murchison Widefield Array Commissioning Survey: A Low-Frequency Catalogue of 14 110 Compact Radio Sources over 6 100 Square Degrees. PASA, 31:e045, November 2014. doi: 10.1017/pasa.2014.40.

[36] J.D. Jackson. *Classical electrodynamics*. Wiley, 1999. ISBN 9780471309321. URL http://books.google.com/books?id=U3LBQgAACAAJ.

[37] 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.

[38] N. Kashikawa, K. Shimasaku, M. A. Malkan, M. Doi, Y. Matsuda, M. Ouchi, Y. Taniguchi, C. Ly, T. Nagao, M. Iye, K. Motohara, T. Murayama, K. Murozono, K. Nariai, K. Ohta, S. Okamura, T. Sasaki, Y. Shioya, and M. Umemura. The End of the Reionization Epoch Probed by Lyα Emitters at z = 6.5 in the Subaru Deep Field. ApJ, 648:7–22, September 2006. doi: 10.1086/504966.

[39] L. Kogan and A. Cohen. Optimization of the lwa antenna station configuration minimizing side lobes. LWA Memo 21, 2005.

[40] Leonid Kogan. Optimizing a large array configuration to minimize the sidelobes. *IEEE Transactions on Antennas and Propagation*, 48(7):1075, 2000. URL http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=876326.

[41] 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. *The Astrophysical Journal Supplement Series*, 192(2):18, 2011. URL http://stacks.iop.org/0067-0049/192/i=2/a=18.

[42] D.V. Lal, A.P. Lobanov, and S. Jiménez-Monferrer. Array configuration studies for the

square kilometre array - implementation of figures of merit based on spatial dynamic range. SKA Memo 107, 2009.

[43] 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.

[44] J. W. Lazio, A. Kimball, A. J. Barger, W. N. Brandt, S. Chatterjee, T. E. Clarke, J. J. Condon, R. L. Dickman, M. T. Hunyh, M. J. Jarvis, M. Jurić, N. E. Kassim, S. T. Myers, S. Nissanke, R. Osten, and B. A. Zauderer. Radio Astronomy in LSST Era. PASP, 126:196–209, February 2014. doi: 10.1086/675262.

[45] A. Lidz, O. Zahn, S. R. Furlanetto, M. McQuinn, L. Hernquist, and M. Zaldarriaga. Probing Reionization with the 21 cm Galaxy Cross-Power Spectrum. ApJ, 690:252–266, January 2009. doi: 10.1088/0004-637X/690/1/252.

[46] Adam Lidz, Oliver Zahn, Matthew McQuinn, Matias Zaldarriaga, and Lars Hernquist. Detecting the rise and fall of 21 cm fluctuations with the murchison wide-field array. *The Astrophysical Journal*, 680:962, Jun 2008. doi: 10.1086/587618. URL `http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2008ApJ...680..962L&link_type=ABSTRACT`. (c) 2008: The American Astronomical Society.

[47] 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.

[48] Adrian Liu and Max Tegmark. A method for 21 cm power spectrum estimation in the presence of foregrounds. *Physical Review D*, 83(10):103006–103006, Jan 2011. doi: doi: 10.1103/PhysRevD.83.103006. URL `http://adsabs.harvard.edu/abs/2011PhRvD..83j3006L`.

[49] Adrian Liu, Aaron R. Parsons, and Cathryn M. Trott. Epoch of reionization window. i.

mathematical formalism. *Phys. Rev. D*, 90:023018, Jul 2014. doi: 10.1103/PhysRevD.90.023018. URL http://link.aps.org/doi/10.1103/PhysRevD.90.023018.

[50] Adrian Liu, Aaron R. Parsons, and Cathryn M. Trott. Epoch of reionization window. ii. statistical methods for foreground wedge reduction. *Phys. Rev. D*, 90:023019, Jul 2014. doi: 10.1103/PhysRevD.90.023019. URL http://link.aps.org/doi/10.1103/PhysRevD.90.023019.

[51] C Lonsdale, R Cappallo, M Morales, F Briggs, L Benkevitch, J Bowman, J Bunton, S Burns, B Corey, L deSouza, S Doeleman, M Derome, A Deshpande, M Gopala, L Greenhill, D Herne, J Hewitt, P Kamini, J Kasper, B Kincaid, J Kocz, E Kowald, E Kratzenberg, D Kumar, M Lynch, S Madhavi, M Matejek, D Mitchell, E Morgan, D Oberoi, S Ord, J Pathikulangara, T Prabu, A Rogers, A Roshi, J Salah, R Sault, N Shankar, K Srivani, J Stevens, S Tingay, A Vaccarella, M Waterson, R Wayth, R Webster, A Whitney, A Williams, and C Williams. The murchison widefield array: Design overview. *Proceedings of the IEEE*, 97(8):1497 – 1506, Aug 2009. doi: 10.1109/JPROC.2009.2017564. URL http://ieeexplore.ieee.org/search/srchabstract.jsp?arnumber=5164979&isnumber=5165124&punumber=5&k2dockey=5164979@ieeejrns.

[52] Matthew Malloy and Adam Lidz. Identifying ionized regions in noisy redshifted 21 cm data sets. *The Astrophysical Journal*, 767(1):68, 2013. URL http://stacks.iop.org/0004-637X/767/i=1/a=68.

[53] Matthew McQuinn, Oliver Zahn, Matias Zaldarriaga, Lars Hernquist, and Steven R Furlanetto. Cosmological parameter estimation using 21 cm radiation from the epoch of reionization. *The Astrophysical Journal*, 653:815, Nov 2006. doi: 10.1086/505167. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006ApJ...653..815M&link_type=ABSTRACT. (c) 2006: The American Astronomical Society.

[54] Andrei Mesinger, Steven Furlanetto, and Renyue Cen. 21cmfast: a fast, seminumerical

simulation of the high-redshift 21-cm signal. *Monthly Notices of the Royal Astronomical Society*, 411(2):955–972, 2011. doi: 10.1111/j.1365-2966.2010.17731.x. URL http://mnras.oxfordjournals.org/content/411/2/955.abstract.

[55] 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. *Selected Topics in Signal Processing, IEEE Journal of*, 2(5):707–717, Oct 2008. ISSN 1932-4553. doi: 10.1109/JSTSP.2008.2005327.

[56] M. F. Morales and M. Matejek. Software holography: interferometric data analysis for the challenges of next generation observatories. *Monthly Notices of the Royal Astronomical Society*, 400(4):1814–1820, 2009. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2009.15537.x. URL http://dx.doi.org/10.1111/j.1365-2966.2009.15537.x.

[57] Miguel Morales and J Wyithe. Reionization and cosmology with 21-cm fluctuations. *Annual Review of Astronomy and Astrophysics*, 48:127–171, Sep 2010. doi: doi:10.1146/annurev-astro-081309-130936. URL http://adsabs.harvard.edu/abs/2010ARA%2526A..48..127M.

[58] Miguel F Morales. Power spectrum sensitivity and the design of epoch of reionization observatories. *The Astrophysical Journal*, 619:678, Jan 2005. doi: 10.1086/426730. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2005ApJ...619..678M&link_type=ABSTRACT. (c) 2005: The American Astronomical Society.

[59] Miguel F Morales and Jacqueline N Hewitt. Toward epoch of reionization measurements with wide-field radio observations. *The Astrophysical Journal*, 615:7, Oct 2004. doi: 10.1086/424437. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2004ApJ...615....7M&link_type=ABSTRACT. (c) 2004: The American Astronomical Society.

[60] Miguel F Morales, Judd D Bowman, and Jacqueline N Hewitt. Improving fore-

ground subtraction in statistical observations of 21 cm emission from the epoch of reionization. *The Astrophysical Journal*, 648:767, Aug 2006. doi: 10.1086/506135. URL `http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006ApJ...648..767M&link_type=ABSTRACT`. (c) 2006: The American Astronomical Society.

[61] Miguel F. Morales, Bryna Hazelton, Ian Sullivan, and Adam Beardsley. Four fundamental foreground power spectrum shapes for 21 cm cosmology observations. *ApJ*, 752 (2):137, 2012. URL `http://stacks.iop.org/0004-637X/752/i=2/a=137`.

[62] A. R. Offringa, A. G. de Bruyn, M. Biehl, S. Zaroubi, G. Bernardi, and V. N. Pandey. Post-correlation radio frequency interference classification methods. *Monthly Notices of the Royal Astronomical Society*, 405(1):155–167, 2010. doi: 10.1111/j.1365-2966.2010.16471.x. URL `http://mnras.oxfordjournals.org/content/405/1/155.abstract`.

[63] 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. *ArXiv e-prints*, January 2015.

[64] S. Ord, L. Greenhill, R. Wayth, D. Mitchell, K. Dale, H. Pfister, and R. G. Edgar. GPUs for data processing in the MWA. *ArXiv e-prints*, February 2009.

[65] Gregory Paciga, Joshua G. Albert, Kevin Bandura, Tzu-Ching Chang, Yashwant Gupta, Christopher Hirata, Julia Odegova, Ue-Li Pen, Jeffrey B. Peterson, Jayanta Roy, J. Richard Shaw, Kris Sigurdson, and Tabitha Voytek. A simulation-calibrated limit on the hâĂŁi power spectrum from the gmrt epoch of reionization experiment. *Monthly Notices of the Royal Astronomical Society*, 433(1):639–647, 2013. doi: 10.1093/mnras/stt753. URL `http://mnras.oxfordjournals.org/content/433/1/639.abstract`.

[66] 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.

[67] A. R. Parsons, A. Liu, J. E. Aguirre, Z. S. Ali, R. F. Bradley, C. L. Carilli, D. R. DeBoer, M. R. Dexter, N. E. Gugliucci, D. C. Jacobs, P. Klima, D. H. E. MacMahon, J. R. Manley, D. F. Moore, J. C. Pober, I. I. Stefan, and W. P. Walbrugh. 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.

[68] Aaron Parsons, Jonathan Pober, Matthew McQuinn, Daniel Jacobs, and James Aguirre. A sensitivity and array-configuration study for measuring the power spectrum of 21 cm emission from reionization. *ApJ*, 753(1):81, 2012.

[69] Aaron R. Parsons, Donald C. Backer, Griffin S. Foster, Melvyn C. H. Wright, Richard F. Bradley, Nicole E. Gugliucci, Chaitali R. Parashare, Erin E. Benoit, James E. Aguirre, Daniel C. Jacobs, Chris L. Carilli, David Herne, Mervyn J. Lynch, Jason R. Manley, and Daniel J. Werthimer. The precision array for probing the epoch of re-ionization: Eight station results. *The Astronomical Journal*, 139(4):1468, 2010. URL `http://stacks.iop.org/1538-3881/139/i=4/a=1468`.

[70] Aaron R Parsons, Jonathan C Pober, James E Aguirre, Christopher L Carilli, Daniel C

Jacobs, and David F Moore. A per-baseline, delay-spectrum technique for accessing the 21 cm cosmic reionization signature. *The Astrophysical Journal*, 756(2):165, 2012.

[71] Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, and et al. Planck 2013 results. XV. CMB power spectra and likelihood. A&A, 571:A15, November 2014. doi: 10.1051/0004-6361/201321573.

[72] 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.

[73] 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. *ArXiv e-prints*, February 2015.

[74] Jonathan C. Pober, Aaron R. Parsons, James E. Aguirre, Zaki Ali, Richard F. Bradley, Chris L. Carilli, Dave DeBoer, Matthew Dexter, Nicole E. Gugliucci, Daniel C. Jacobs, Patricia J. Klima, Dave MacMahon, Jason Manley, David F. Moore, Irina I. Stefan, and William P. Walbrugh. Opening the 21 cm epoch of reionization window: Measurements of foreground isolation with paper. *The Astrophysical Journal Letters*, 768(2):L36, 2013. URL `http://stacks.iop.org/2041-8205/768/i=2/a=L36`.

[75] Jonathan C. Pober, Adrian Liu, Joshua S. Dillon, James E. Aguirre, Judd D. Bowman, Richard F. Bradley, Chris L. Carilli, David R. DeBoer, Jacqueline N. Hewitt, Daniel C. Jacobs, Matthew McQuinn, Miguel F. Morales, Aaron R. Parsons, Max Tegmark, and Dan J. Werthimer. What next-generation 21 cm power spectrum measurements can

teach us about the epoch of reionization. *The Astrophysical Journal*, 782(2):66, 2014. URL http://stacks.iop.org/0004-637X/782/i=2/a=66.

[76] Thiagaraj Prabu, K.S. Srivani, D.Anish Roshi, P.A. Kamini, S. Madhavi, David Emrich, Brian Crosse, AndrewJ. Williams, Mark Waterson, AvinashA. Deshpande, N.Udaya Shankar, Ravi Subrahmanyan, FrankH. Briggs, RobertF. Goeke, StevenJ. Tingay, Melanie Johnston-Hollitt, GopalakrishnaM R, EdwardH. Morgan, Joseph Pathikulangara, JohnD. Bunton, Grant Hampson, Christopher Williams, StephenM. Ord, RandallB. Wayth, Deepak Kumar, MiguelF. Morales, Ludi deSouza, Eric Kratzenberg, D. Pallot, Russell McWhirter, BrynaJ. Hazelton, Wayne Arcus, DavidG. Barnes, Gianni Bernardi, T. Booler, JuddD. Bowman, RogerJ. Cappallo, BrianE. Corey, LincolnJ. Greenhill, David Herne, JacquelineN. Hewitt, DavidL. Kaplan, JustinC. Kasper, BartonB. Kincaid, Ronald Koenig, ColinJ. Lonsdale, MervynJ. Lynch, DanielA. Mitchell, Divya Oberoi, RonaldA. Remillard, AlanE. Rogers, JosephE. Salah, RobertJ. Sault, JamieB. Stevens, S. Tremblay, RachelL. Webster, AlanR. Whitney, and StuartB. Wyithe. A digital-receiver for the murchisonwidefield array. *Experimental Astronomy*, pages 1–21, 2015. ISSN 0922-6435. doi: 10.1007/s10686-015-9444-3. URL http://dx.doi.org/10.1007/s10686-015-9444-3.

[77] 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.

[78] 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.

[79] J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. ApJ, 263:835–853, December 1982. doi: 10.1086/160554.

[80] 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.

[81] O. B. Slee. Culgoora-3 list of radio source measurements. *Australian Journal of Physics Astrophysical Supplement*, 43:1–123, November 1977.

[82] 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. Wide-Field InfraRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA Final Report. *ArXiv e-prints*, May 2013.

[83] 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. Subrahmanyan, 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. *The Astrophysical Journal*, 759(1): 17, 2012. URL `http://stacks.iop.org/0004-637X/759/i=1/a=17`.

[84] A. Sutinjo, J. O'Sullivan, E. Lenc, R. B. Wayth, S. Padhi, P. Hall, and S. J. Tingay.

Understanding instrumental stokes leakage in murchison widefield array polarimetry. *Radio Science*, 50(1):52–65, 2015. ISSN 1944-799X. doi: 10.1002/2014RS005517. URL http://dx.doi.org/10.1002/2014RS005517.

[85] M. Tegmark. How to Make Maps from Cosmic Microwave Background Data without Losing Information. *ApJ*, 480:L87, May 1997. doi: 10.1086/310631.

[86] M. Tegmark. How to measure cmb power spectra without losing information. *Phys. Rev. D*, 55(10):5895–5907, May 1997. doi: 10.1103/PhysRevD.55.5895.

[87] A. Richard Thompson. Fundamentals of radio interferometry. In *Synthesis Imaging in Radio Astronomy II*, volume 180, page 11, 1999.

[88] 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. *ArXiv e-prints*, February 2015.

[89] Nithyanandan Thyagarajan, N. Udaya Shankar, Ravi Subrahmanyan, Wayne Arcus, Gianni Bernardi, Judd D. Bowman, Frank Briggs, John D. Bunton, Roger J. Cappallo, Brian E. Corey, Ludi deSouza, David Emrich, Bryan M. Gaensler, Robert F. Goeke, Lincoln J. Greenhill, Bryna J. Hazelton, David Herne, Jacqueline N. Hewitt, Melanie Johnston-Hollitt, David L. Kaplan, Justin C. Kasper, Barton B. Kincaid,

Ronald Koenig, Eric Kratzenberg, Colin J. Lonsdale, Mervyn J. Lynch, S. Russell McWhirter, Daniel A. Mitchell, Miguel F. Morales, Edward H. Morgan, Divya Oberoi, Stephen M. Ord, Joseph Pathikulangara, Ronald A. Remillard, Alan E. E. Rogers, D. Anish Roshi, Joseph E. Salah, Robert J. Sault, K. S. Srivani, Jamie B. Stevens, Prabu Thiagaraj, Steven J. Tingay, Randall B. Wayth, Mark Waterson, Rachel L. Webster, Alan R. Whitney, Andrew J. Williams, Christopher L. Williams, and J. Stuart B. Wyithe. A study of fundamental limitations to statistical detection of redshifted h i from the epoch of reionization. *The Astrophysical Journal*, 776(1):6, 2013. URL `http://stacks.iop.org/0004-637X/776/i=1/a=6`.

[90] S. J. Tingay, R. Goeke, J. D. Bowman, D. Emrich, S. M. Ord, D. A. Mitchell, M. F. Morales, T. Booler, B. Crosse, R. B. Wayth, C. J. Lonsdale, S. Tremblay, D. Pallot, T. Colegate, A. Wicenec, 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. Schlagenhaufer, K. S. Srivani, J. Stevens, R. Subrahmanyan, 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 - Publications of the Astronomical Society of Australia*, 30(e007), 1 2013. ISSN 1448-6083. doi: 10.1017/pasa.2012.007. URL `http://journals.cambridge.org/article_S1323358012000070`.

[91] Cathryn M. Trott, Randall B. Wayth, and Steven J. Tingay. The impact of point-source subtraction residuals on 21 cm epoch of reionization estimation. *The Astrophysical Journal*, 757(1):101, 2012. URL `http://stacks.iop.org/0004-637X/757/i=1/a=101`.

[92] Harish Vedantham, N. Udaya Shankar, and Ravi Subrahmanyan. Imaging the epoch

of reionization: Limitations from foreground confusion and imaging algorithms. *The Astrophysical Journal*, 745(2):176, 2012. URL `http://stacks.iop.org/0004-637X/745/i=2/a=176`.

[93] R. B. Wayth, L. J. Greenhill, and F. H. Briggs. A GPU-based Real-time Software Correlation System for the Murchison Widefield Array Prototype. PASP, 121:857–865, August 2009. doi: 10.1086/605334.

[94] D. C. Wells, E. W. Greisen, and R. H. Harten. FITS - a Flexible Image Transport System. A&AS, 44:363, June 1981.

[95] S. Yatawatta, A. G. de Bruyn, M. A. Brentjens, Labropoulos, P., Pandey, V. N., Kazemi, S., Zaroubi, S., Koopmans, L. V. E., Offringa, A. R., Jelić, V., Martinez Rubi, O., Veligatla, V., Wijnholds, S. J., Brouw, W. N., Bernardi, G., Ciardi, B., Daiboo, S., Harker, G., Mellema, G., Schaye, J., Thomas, R., Vedantham, H., Chapman, E., Abdalla, F. B., Alexov, A., Anderson, J., Avruch, I. M., Batejat, F., Bell, M. E., Bell, M. R., Bentum, M., Best, P., Bonafede, A., Bregman, J., Breitling, F., van de Brink, R. H., Broderick, J. W., Brüggen, M., Conway, J., de Gasperin, F., de Geus, E., Duscha, S., Falcke, H., Fallows, R. A., Ferrari, C., Frieswijk, W., Garrett, M. A., Griessmeier, J. M., Gunst, A. W., Hassall, T. E., Hessels, J. W. T., Hoeft, M., Iacobelli, M., Juette, E., Karastergiou, A., Kondratiev, V. I., Kramer, M., Kuniyoshi, M., Kuper, G., van Leeuwen, J., Maat, P., Mann, G., McKean, J. P., Mevius, M., Mol, J. D., Munk, H., Nijboer, R., Noordam, J. E., Norden, M. J., Orru, E., Paas, H., Pandey-Pommier, M., Pizzo, R., Polatidis, A. G., Reich, W., Röttgering, H. J. A., Sluman, J., Smirnov, O., Stappers, B., Steinmetz, M., Tagger, M., Tang, Y., Tasse, C., ter Veen, S., Vermeulen, R., van Weeren, R. J., Wise, M., Wucknitz, O., and Zarka, P. Initial deep lofar observations of epoch of reionization windows. *Astronomy and Astrophysics*, 550:A136, 2013. doi: 10.1051/0004-6361/201220874. URL `http://dx.doi.org/10.1051/0004-6361/201220874`.

[96] E. Zackrisson, C.-E. Rydberg, D. Schaerer, G. Östlin, and M. Tuli. The Spectral Evolution of the First Galaxies. I. James Webb Space Telescope Detection Limits and Color Criteria for Population III Galaxies. ApJ, 740:13, October 2011. doi: 10.1088/0004-637X/740/1/13.

[97] Oliver Zahn, Adam Lidz, Matthew McQuinn, Suvendra Dutta, Lars Hernquist, Matias Zaldarriaga, and Steven R. Furlanetto. Simulations and analytic calculations of bubble growth during hydrogen reionization. *The Astrophysical Journal*, 654(1):12, 2007. URL `http://stacks.iop.org/0004-637X/654/i=1/a=12`.

Appendix A

INTERPRETING DISCRETE FOURIER TRANSFORMS

Often times it is necessary to approximate a continuous Fourier transform (CFT) with a discrete Fourier transform (DFT) when an analytic form of the data is not available. For example, a DFT can be performed on a sampling of time-ordered electric field measurements to approximate the frequency spectrum of the field. In order to perform these transforms, we utilize the numerical packages of programming languages such as IDL, MATLAB, or Mathematica. However, the definition of a Fourier transform is not completely constrained, and freedoms exist that give rise to multiple conventions in the normalization[1], and scaling of the Fourier domain variables. In particular, the numerical package being used may not follow the desired convention. Below we outline a method for interpreting the results of a DFT, in terms of the CFT desired. Then we provide several examples of common conventions.

A.1 Derivation

We start by assuming we have a continuous function, $f_c(x)$, of which we would like to compute a continuous Fourier transform of the form

$$F_c(u; C_c, p_c) = C_c \int_{-\infty}^{\infty} dx \; f_c(x) e^{ip_c ux} \tag{A.1}$$

where the constants C_c and p_c are determined by convention of the transform. However, in reality we often have a discrete sample of the continuous function, $f_d(j)$, where j indexes the samples and runs from 0 to $N-1$, N being the number of samples. Then it is necessary to use a discrete Fourier transform. Numerical packages have various conventions for their

[1]There are of course constraints limiting these freedoms, but for this argument we will leave them completely free.

DFTs, but in general they can be expressed in the form

$$F_d(k; C_d, p_d) = C_d \sum_{j=0}^{N-1} f_d(j) e^{ip_d jk} \tag{A.2}$$

where k indexes the Fourier domain samples, and also runs from 0 to $N-1$. In some cases, the indices can be shifted. For example, MATLAB counts arrays from 1 rather than 0, so the sum would run from 1 to N, and the values in the exponent are adjusted accordingly. For simplicity, we assume counting from 0. The translation is straighforward, and an example is given in Section A.3. We wish to approximate the CFT with the DFT.

In order to proceed, we assume the function $f_c(x)$ is only significantly non-zero within a domain $\mathcal{D}$. Furthermore, we will assume $\mathcal{D} = [0, x_{max}]$. This constraint is not actually necessary, and a translation of the domain results in a phase in the Fourier transform, as seen in Section A.2. Then we can relate the indices in Eq. (A.2) to the axis values with the following expressions:

$$x(j) = j\Delta x \tag{A.3a}$$

$$u(k) = k\Delta u \tag{A.3b}$$

where $\Delta x = x_{max}/(N-1)$, and Δu will be derived below. We can now approximate the CFT with a finite sum.

$$\begin{aligned} F_c(u; C_c, p_c) &\approx C_c \int_0^{x_{max}} dx\ f_c(x) e^{ip_c ux} \\ &\approx C_c \sum_{j=0}^{N-1} \Delta x f_c(x(j)) e^{ip_c uj\Delta x} \\ &= C_c \Delta x \sum_{j=0}^{N-1} f_d(j) e^{ip_c uj\Delta x} \end{aligned} \tag{A.4}$$

In order to force the sum in Eq. (A.4) to the form of the DFT (Eq. (A.2)), we write $u(k)$ in the following way.

$$\begin{aligned} u(k) = k\Delta u &= \frac{p_d k}{p_c \Delta x} \\ \Rightarrow \Delta u &= \frac{p_d}{p_c \Delta x} \end{aligned} \tag{A.5}$$

This expression for Δu defines our Fourier axis, which is an extremely important part of interpreting a Fourier transform and in particular comparing to other computational conventions.

We can now simplify Eq. A.4.

$$\begin{aligned} F_c(u; C_c, p_c) &\approx C_c \Delta x \sum_{j=0}^{N-1} f_d(j) e^{i p_d j k} \\ &= \frac{C_c}{C_d} \Delta x F_d(k; C_d, p_d) \\ &= \frac{C_c}{C_d} \Delta x F_d(u/\Delta u; C_d, p_d) \end{aligned} \tag{A.6}$$

In the first line of Eq. A.6 we introduced an intermediate index $k = u/\Delta u$ for clarity in replacing the sum with the DFT. Of course any computer package will compute the DFT for integer values of k, and it should be noted that this approximation for the CFT is only used for u equal to integer multiples of Δu.

To summarize, we have related an arbitrary convention of the continuous Fourier transform to an arbitrary convention of the discrete Fourier transform (Eq. A.6). In addition, we have shown the resulting Fourier axis is described by Eq. A.5. For reference, several examples are discussed in Section A.3.

A.2 Shifting Domains

The above derivation applies to functions significantly non-zero only within a domain $\mathcal{D} = [0, x_{max}]$. We can relax that condition slightly by allowing $\mathcal{D} \rightarrow [x_{min}, x_{max}]$. The same process can be followed as in Section A.1, with a simple translation of the variable of

integration. The result is a mode-dependent phase, and we can generalize Eq. A.6 as

$$F_c(u; C_c, p_c) \approx \frac{C_c}{C_d} \Delta x e^{ip_c u x_{min}} F_d(u/\Delta u; C_d, p_d) \tag{A.7}$$

where now $\Delta x = (x_{max} - x_{min})/(N-1)$.

A particular case of interest is to shift the real and Fourier domains such that they are centered at zero. This can be achieved by a couple shifts of the arrays. In this case, $x_{min} = -\Delta x(N-1)/2$. The phases in Eq. A.7 are now negative, but can be shifted properly to make them positive. This amounts to exchanging the positive and negative x data in $f_d(j)$. This is illustrated in Step A in Fig. A.1.

Once the DFT is performed, our approximation for $F_c(u)$ is still only valid for positive (or in some cases only negative) values of u. But we note that by the periodic nature of a complex phase, for $k > (N-1)/2$, and for $p_d = 2\pi/N$ (as is normally the case), we can make the substituion $k \to -N + k$ and the sum in Eq. A.2 remains unchanged. In other words, we can relabel and shift our u axis to center our Fourier axis on $u = 0$, as shown in Steps C and D in Fig. A.1. All these operations can be easily implemented using shifting functions built into many software packages. In MATLAB, for example, `ifftshift` (Step A) and `fftshift` (Step D) are provided for this very reason.

A.3 Examples

Here we apply the interpretation of Section A.1 specifically to the default MATLAB discrete Fourier transform. Further examples are listed in Table A.1. For the forward DFT, MATLAB uses the convention

$$F_d(k) = \sum_{j=1}^{N} f_d(j) e^{-2\pi i(j-1)(k-1)/N},$$

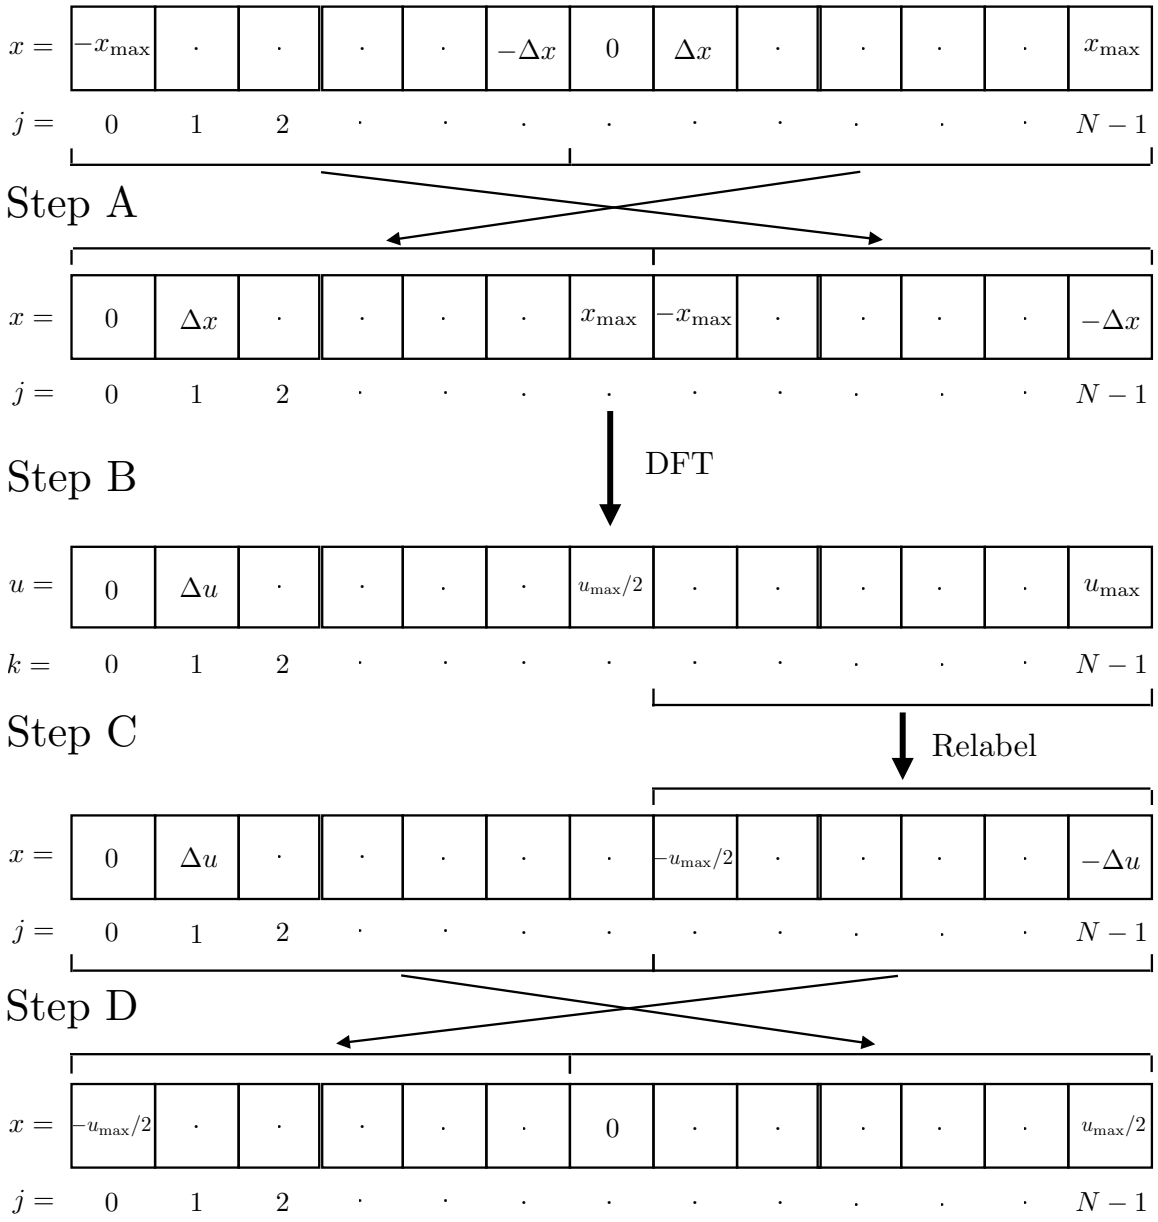

Figure A.1: Centering axes on 0. *Step A*: Shift the real axis to be fed into DFT package. *Step B*: Perform discrete Fourier transform. *Step C*: Relabel Fourier modes to equivalent negative modes. *Step D*: Shift Fourier axis to recover correct order. This figure is shown for the odd N case, but the general method also applies for even N, except that there is one fewer negative bins than positive, and one must be careful to index accordingly when shifting.

so that $C_d = 1$ and $p_d = -2\pi/N$. Note the shift in indices due to the fact that MATLAB counts from 1, rather than zero. This can be accounted for by writing the axes as

$$x(j) = (j-1)\Delta x$$
$$u(k) = (k-1)\Delta u$$

where $\Delta u = p_d/(p_c \Delta x) = -2\pi/(p_c N \Delta x)$.

Let us now assume that we have sampled a function of position, $f_d(j)$, with resolution $\Delta x = 3\text{m}$, and $N = 1001$ ($x_{range} = 3000\text{m}$). We are interested in the Fourier transform of the form

$$F_c(u) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dx \; f_c(x) e^{iux}$$

so that $C_c = 1/\sqrt{2\pi}$ and $p_c = 1$. Then we use Eq. A.6 to approximate the CFT as

$$F_c(u) \approx \frac{3\text{m}}{\sqrt{2\pi}} F_d(-u\frac{3003\text{m}}{2\pi} + 1; 1, 2\pi/N)$$

where $F_d(k, 1, 2\pi/N)$ is the array output from MATLAB's `fft` function, given $f_d(j)$. Note that because the sign of p_c is opposite the sign of p_d, the approximation for the CFT will only be valid for negative Fourier modes. In most applications this difference will not be significant, but in some cases it may be important to account for the axis reversal.

Similar steps can be taken for any DFT convention, and some common default parameters are listed in Table A.1. We also list several common CFT conventions and their corresponding conversions from common numerical packages in Table A.2.

A.4 Inverse Fourier Transform

The above argument holds for the inverse Fourier transform as well. Define the inverse CFT as

$$f_c(x; C_c', p_c') = C_c' \int_{-\infty}^{\infty} du \; F_c(k) e^{ip_c' ux} \tag{A.8}$$

and the inverse DFT as

$$f_d(j; C_d', p_d') = C_d' \sum_{k=0}^{N-1} F_d(k) e^{ip_d' jk}. \tag{A.9}$$

Table A.1: Example default parameters for the forward discrete Fourier transform for various numerical packages

Package	$F_d(k)$	C_d	p_d	$x(j)$	$u(k)$
IDL	$\frac{1}{N}\sum_{j=0}^{N-1} f_d(j)e^{-2\pi ijk/N}$	$1/N$	$-2\pi/N$	$j\Delta x$	$-k\frac{2\pi}{p_c N\Delta x}$
MATLAB	$\sum_{j=1}^{N} f_d(j)e^{2\pi i(j-1)(k-1)/N}$	1	$-2\pi/N$	$(j-1)\Delta x$	$-(k-1)\frac{2\pi}{p_c N\Delta x}$
Mathematica	$\frac{1}{\sqrt{N}}\sum_{j=1}^{N} f_d(j)e^{2\pi i(j-1)(k-1)/N}$	$\frac{1}{\sqrt{N}}$	$2\pi/N$	$(j-1)\Delta x$	$(k-1)\frac{2\pi}{p_c N\Delta x}$

Table A.2: Common CFT conventions listed with the normalization factor and axis scaling factors for sample numerical packages.

DFT Package	$F_c(u)$ Desired	Amplitude Factor	$\lvert\Delta u\rvert$
IDL	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{iux}$	$\Delta x N/\sqrt{2\pi}$	$2\pi/(N\Delta x)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{iux}$	$\Delta x N/2\pi$	$2\pi/(N\Delta x)$
	$\int dx\ f_c(x)e^{2\pi iux}$	$\Delta x N$	$1/(N\Delta x)$
MATLAB	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{iux}$	$\Delta x/\sqrt{2\pi}$	$2\pi/(N\Delta x)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{iux}$	$\Delta x/2\pi$	$2\pi/(N\Delta x)$
	$\int dx\ f_c(x)e^{2\pi iux}$	Δx	$1/(N\Delta x)$
Mathematica	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{iux}$	$\Delta x\sqrt{N/2\pi}$	$2\pi/(N\Delta x)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{iux}$	$\Delta x\sqrt{N}/2\pi$	$2\pi/(N\Delta x)$
	$\int dx\ f_c(x)e^{2\pi iux}$	$\Delta x\sqrt{N}$	$1/(N\Delta x)$

Table A.3: Example default parameters for the inverse discrete Fourier transform for various numerical packages

Package	$f_d(j)$	C'_d	p'_d	$u(k)$	$x(j)$
IDL	$\sum_{k=0}^{N-1} F_d(k)e^{2\pi ijk/N}$	1	$2\pi/N$	$k\Delta u$	$j\frac{2\pi}{p'_c N\Delta u}$
MATLAB	$\frac{1}{N}\sum_{k=1}^{N} F_d(k)e^{-2\pi i(j-1)(k-1)/N}$	$1/N$	$2\pi/N$	$(k-1)\Delta u$	$(j-1)\frac{2\pi}{p'_c N\Delta u}$
Mathematica	$\frac{1}{\sqrt{N}}\sum_{k=1}^{N} F_d(k)e^{-2\pi i(j-1)(k-1)/N}$	$\frac{1}{\sqrt{N}}$	$-2\pi/N$	$(k-1)\Delta u$	$-(j-1)\frac{2\pi}{p'_c N\Delta u}$

Then the result for the approximation of the inverse CFT is

$$f_c(x; C'_c, p'_c) \approx \frac{C'_c}{C'_d}\Delta u f_d(x/\Delta x; C'_d, p'_d) \tag{A.10}$$

where $\Delta x = p'_d/p'_c\Delta u$.

We summarize the default parameters for the inverse DFT for various numerical packages in Table A.3, and provide a few applied examples to common CFTs in Table A.4.

Table A.4: Common inverse CFT conventions listed with the normalization factor and axis scaling factors for sample numerical packages.

DFT Package	$f_c(x)$ Desired	Amplitude Factor	$\lvert\Delta x\rvert$
IDL	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{-iux}$	$\Delta u/\sqrt{2\pi}$	$2\pi/(N\Delta u)$
	$\int dx\ f_c(x)e^{-iux}$	Δu	$2\pi/(N\Delta u)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{-2\pi iux}$	$\Delta u/2\pi$	$1/(N\Delta u)$
MATLAB	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{-iux}$	$\Delta uN/\sqrt{2\pi}$	$2\pi/(N\Delta u)$
	$\int dx\ f_c(x)e^{-iux}$	ΔuN	$2\pi/(N\Delta u)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{-2\pi iux}$	$\Delta uN/2\pi$	$1/(N\Delta u)$
Mathematica	$\frac{1}{\sqrt{2\pi}}\int dx\ f_c(x)e^{-iux}$	$\Delta u\sqrt{N/2\pi}$	$2\pi/(N\Delta u)$
	$\int dx\ f_c(x)e^{-iux}$	$\Delta u\sqrt{N}$	$2\pi/(N\Delta u)$
	$\frac{1}{2\pi}\int dx\ f_c(x)e^{-2\pi iux}$	$\Delta u\sqrt{N}/2\pi$	$1/(N\Delta u)$

Appendix B

DETAILED WINDOW FUNCTION CALCULATION

We can calculate the integrated window function for zenith pointing analytically using a MWA beam model. We treat a tile as a set of short dipoles sitting above a ground plane. The dipoles will have relative phases with respect to one another based on the direction of some source. We will calculate the response for an individual dipole at location (x_i, y_i, h) relative to the center of the tile. The vertical component, h, is the height of the dipole, and we will assume it is the same for all dipoles on a single tile.

The phase difference for the dipole is found using simple geometry.

$$\Delta\phi_i = \frac{2\pi}{\lambda}(x_i \sin\theta_x + y_i \sin\theta_y) \tag{B.1}$$

In addition to the phase differences, the ground plane also serves as an amplitude modulation. If the reflected wave is in phase with the incoming wave, the dipole will have twice the response as if the plane was not there. On the other hand, if the reflection is completely out of phase, the dipole will have zero response. We can derive the form of the amplitude modulation using geometry again. This time we find the phase difference between the incident wave and the reflected wave. The reflected part can be found by modeling the path difference between the true dipole above the plane and a virtual dipole below the plane.

$$\Delta D = 2h\cos\theta_z \tag{B.2}$$

Here θ_z is the zenith angle. The phase will be given by this path difference as well as a π flip at the ground plane interface.

$$\Delta\phi_{r,i} = \pi + \frac{2\pi}{\lambda}(2h\cos\theta_z) \tag{B.3}$$

Because this phase is independent of dipole (assumed h is the same for all dipoles), we can solve for the amplitude of the incident plus reflected wave.

$$A = \left|1 + e^{i\Delta\phi_r}\right| \tag{B.4a}$$
$$= \left((1+\cos(\Delta\phi_r))^2 + (\sin(\Delta\phi_r))^2\right)^{1/2} \tag{B.4b}$$
$$= \left(1 + 2\cos(\Delta\phi_r) + \cos^2(\Delta\phi_r) + \sin^2(\Delta\phi_r)\right)^{1/2} \tag{B.4c}$$
$$= \left(2(1+\cos(\Delta\phi_r))\right)^{1/2} \tag{B.4d}$$
$$= 2\left|\cos(\Delta\phi_r/2)\right| \tag{B.4e}$$
$$= 2\left|\cos\left(\pi/2 + \frac{2\pi}{\lambda}h\cos\theta_z\right)\right| \tag{B.4f}$$
$$= 2\sin\left(\frac{2\pi}{\lambda}h\cos\theta_z\right) \tag{B.4g}$$

The dipole will be sensitive only to radiation with polarization aligned along its axis. Due to the orthogonality between direction of propagation and polarization, a non polarized source will appear polarized to the instrument if it is off zenith. For a dipole with single polarization, the amplitude attenuation is given by

$$P = 1 - \sin(\theta_z)\cos(\theta_{az}), \tag{B.5}$$

where θ_{az} is the azimuth angle.

The response for the dipole is then given by

$$R_i = APe^{i\Delta\phi_i} \tag{B.6a}$$
$$= 2\sin\left(\frac{2\pi}{\lambda}h\cos\theta_z\right)\exp\left[\frac{2\pi i}{\lambda}(x_i\sin\theta_x + y_i cos\theta_y)\right](1 - \sin(\theta_z)\cos(\theta_{az})) \tag{B.6b}$$

Next we construct the response for the full tile by summing over the dipoles. We can leverage the symmetry of the tile by defining a horizontal space, δ_x, which we will assume to

be uniform.

$$R_{tile} = \sum_i R_i \tag{B.7a}$$

$$= \sum_i 2\sin\left(\frac{2\pi}{\lambda}h\cos\theta_z\right)\exp\left[\frac{2\pi i}{\lambda}(x_i\sin\theta_x + y_i cos\theta_y)\right](1-\sin(\theta_z)\cos(\theta_{az})) \tag{B.7b}$$

$$= 2\sin\left(\frac{2\pi}{\lambda}h\cos\theta_z\right)\sum_j\sum_k\exp\left[\frac{2\pi i}{\lambda}(j\delta_x\sin\theta_x + k\delta_x cos\theta_y)\right](1-\sin(\theta_z)\cos(\theta_{az})) \tag{B.7c}$$

The j and k sums are over the half integers -3/2, -1/2, 1/2, and 3/2.

In the approximation that the tiles are identical (and thus have the same R_{tile}), the window function is simply the magnitude squared of the response. We also normalize the window so that at phase center it is equal to unity. Further, the window function in general is a function of the frequency, but we will assume it is constant within a channel.

$$W(\theta_x,\theta_y,f) = \frac{|R_{tile}(\theta_x,\theta_y)|^2}{|R_{tile}(0,0)|^2} \tag{B.8}$$

Next we can use the transformations from Eq. 2.8 to rewrite the window function in terms of cosmological axes ($\mathbf{r}$). With some simplification we arrive at the following for the window function squared.

$$\begin{aligned}|W(\mathbf{r})|^2 = \frac{1}{256\left|\sin\left(\frac{2\pi h}{\lambda}\right)\right|^4}\Bigg|\sin\left(\frac{2\pi h\cos\left(\sqrt{r_x^2+r_y^2}/D_m\right)}{\lambda}\right)\times \\ \left(\cos\left(\frac{\pi\delta_x\sin(r_x/D_m)}{\lambda}\right)+\cos\left(\frac{3\pi\delta_x\sin(r_x/D_m)}{\lambda}\right)\right)\times \\ \left(\cos\left(\frac{\pi\delta_x\sin(r_y/D_m)}{\lambda}\right)+\cos\left(\frac{3\pi\delta_x\sin(r_y/D_m)}{\lambda}\right)\right)\Bigg|^4\times \\ \left(1-\sin\left(\frac{\sqrt{r_x^2+r_y^2}}{D_m}\right)\frac{r_y}{\sqrt{r_x^2+r_y^2}}\right)^4\end{aligned} \tag{B.9}$$

This can then be integrated (probably numerically) to arrive at our estimate of the integrated window function (with the assumptions of identical tiles phased at zenith). As an

Parameter	Symbol	Value
Redshift	z	8
Wavelength	λ	1.90 m
Dipole spacing	δ_x	1.1 m
Dipole height	h	0.35 m
Bandwidth	B	8 MHz

Table B.1: Observing and instrument parameters for example calculation

example I compute this integral for observing parameters shown in Table B.1, and find

$$\int |W(\mathbf{r})|^2 d^3\mathbf{r} \approx 7.0 \times 10^8 \text{ Mpc}^3 \tag{B.10}$$

VITA

Adam Beardsley was born in Milwaukee, Wisconsin to Peter and Crystal Beardsley. He grew up in the small town of Oconomowoc, and attended the University of Wisconsin-Madison for his undergraduate studies. There he became active in physics research working on the Thomson scattering experiment on the Madison Symmetric Torus. He was introduced to 21cm radio astronomy by Professor Snežana Stanimirović in a hands-on astronomy course. Adam graduated from Wisconsin with a triple major in physics, mathematics, and astronomy before heading to the University of Washington to begin his graduate career. There he joined Professor Miguel Morales' radio cosmology group in his first year and has worked on the MWA since. Following graduation from Washington, Adam will continue working on EoR experiments by joining the Low-Frequency Cosmology group at Arizona State University as a postdoc working on HERA.