

©Copyright 2024

Pyxie Star

Digital Nonlinearities and the first Epoch of Reionization Power
Spectrum Limit from MWA Phase III

Pyxie Star

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

Doctor of Philosophy

University of Washington

2024

Reading Committee:
Miguel Morales, Chair
Marilena Loverde
Marcel den Nijs

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

Digital Nonlinearities and the first Epoch of Reionization Power Spectrum Limit from
MWA Phase III

Pyxie Star

Chair of the Supervisory Committee:

Miguel Morales

Physics

This work advances our goal of understanding the Epoch of Reionization (EoR) by measuring the 21-centimeter line from neutral hydrogen. As the very first stars formed during the EoR, they released ultraviolet light which ionized the surrounding hydrogen until the whole universe was reionized about a billion years after the Big Bang. We probe the EoR by measuring the 21-centimeter line from neutral hydrogen. Detecting this faint signal will unlock mysteries of the universe, identifying primary drivers of the evolution into the structures we see today, as well as constraining theories of the very beginning of the universe and informing mechanics of present-day star and galaxy formation.

We use measurements from the Murchison Widefield Array (MWA) to detect the signal from the EoR. In order to uncover the faint signal of the 21-centimeter line we must remove the instrument response from our data with a precision of a part in 10^{-5} . While our calibration procedures remove the linear instrument response, nonlinear artifacts remain. We show the effects of these artifacts and identify a source: quantization in the digital signal pathway.

Like most instruments, the MWA employs requantization throughout the signal path to combat bit growth between digital signal processing components. We use the Van Vleck correction to remove the effects of one of these requantization stages: the 4-bit requantization immediately before correlation. Then, we model the signal path to develop a correction for

an incorrectly implemented requantization stage within the correlator polyphase filterbank between the subfilters and the FFT. These two corrections result in promising improvements to our data quality and analyses.

Partly motivated by this work, the MWA collaboration decided to eliminate requantization stages in the new MWA correlator. The first upgrade introduced for MWA Phase III, the MWAX correlator, was commissioned in 2021 and officially replaced the Legacy correlator in 2022. We inspect data from the 2021 EoR commissioning data taken with the new MWAX correlator for artifacts from digital nonlinearities and show that there are significant improvements to data quality.

We then analyze the 2021 data to obtain the first power spectrum limit for MWA Phase III. This requires revalidating our calibration approach, updating our flagging and averaging settings for the new correlator data, and retuning our radio frequency interference (RFI) detection for the 2021 RFI environment. We perform data selection quality cuts to obtain our final set of 180 observations with a total of ~ 5.569 hours of data.

Our power spectrum limit has promising results, revealing a drop in contamination from residual nonlinearities in the data. We calculate a lowest upper limit of $\Delta^2 \leq 1.76 \cdot 10^4 \text{ mK}^2$ with our relatively small 2021 commissioning data set, and are hopeful that adding Phase III data from later observing seasons to this analysis would lead to state of the art constraints on the Epoch of Reionization.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	vii
Glossary	viii
Chapter 1: Introduction	1
1.1 The early universe	1
1.2 Observing the Epoch of Reionization with the 21 centimeter hydrogen line	3
1.3 The Murchison Widefield Array	7
1.4 Precision EoR science with the MWA	9
Chapter 2: Precision calibration for EoR science	11
2.1 Calibration overview	11
2.2 Calibration precision requirements	13
Chapter 3: Nonlinear quantization effects in MWA data	17
3.1 Digital quantization nonlinearities	17
3.2 Antenna	18
3.3 Receiver	18
3.4 Correlator	23
3.5 MWA quantization stages	26
3.6 Effects of nonlinear quantization in MWA data	29
Chapter 4: Correcting quantization nonlinearities	36
4.1 Quantization effects on signal statistics	37
4.2 Implementation of quantization corrections	45
4.3 Results with Van Vleck corrections	49

Chapter 5:	Correcting quantization of real-world systems	55
5.1	Asymmetric quantization in the fine PFB	56
5.2	A correction for the PFB asymmetric quantization	56
5.3	Developing a correction for the asymmetric 8-bit requantization	59
Chapter 6:	MWA Phase III and the new MWAX correlator	67
6.1	MWAX design	68
6.2	Understanding MWAX	72
6.3	A first look at MWAX data	72
Chapter 7:	Analyzing MWAX performance through autocorrelations	74
7.1	MWAX antenna response	74
7.2	Antenna temporal stability	78
7.3	Antenna rogues' gallery	81
Chapter 8:	Preparing for a limit with MWAX data	95
8.1	MWAX limit data set	95
8.2	Preprocessing pipeline	96
8.3	FHD/ ϵ analysis pipeline	99
Chapter 9:	Compute systems for EoR data analysis	108
9.1	Pipeline computing requirements	108
9.2	Cloud computing	109
Chapter 10:	Fine Channel Flagging and Averaging with MWAX data	113
10.1	Flagging and averaging with MWAX	113
Chapter 11:	MWAX data quality analysis	129
11.1	The SSINS of MWAX	130
11.2	Calibration polynomial fitting parameter analysis	141
11.3	Power spectra quality analysis	145
Chapter 12:	A first power spectrum limit with MWA Phase III data	158
12.1	The 1D power spectrum limit	159
12.2	Discussion	164

12.3 Conclusion	165
Bibliography	169

LIST OF FIGURES

Figure Number	Page
1.1 History of the universe	2
1.2 Hydrogen ground state hyperfine splitting	5
1.3 Image of MWA array	8
2.1 Cartoon 2D power spectrum	15
3.1 MWA signal path	19
3.2 Image of MWA tile and beamformer	20
3.3 Image of MWA receivers	21
3.4 Conceptual polyphase filterbank model	24
3.5 Polyphase filterbank implementation	25
3.6 MWA quantization stages	28
3.7 MWA polyphase filterbank response	31
3.8 Nonlinearities in autocorrelations	32
3.9 Calibration nonlinear artifacts as a function of temperature	33
3.10 Coarse band lines in the 2D power spectrum	35
4.1 Standard deviations before and after quantization	39
4.2 Ratio of covariances before and after quantization	41
4.3 Difference between Van Vleck correction for crosscorrelations and Chebyshev approximation	48
4.4 Autocorrelation with the 4-bit Van Vleck correction	52
4.5 1D power spectrum with the 4-bit Van Vleck correction	53
4.6 Van Vleck correction 2D power spectrum comparison	54
5.1 The asymmetric rounding scheme of the MWA fine PFB	57
5.2 Fitting an standard deviation correction for the asymmetric 8-bit requantization	62
5.3 Fitting a covariance correction for the asymmetric 8-bit requantization	64

5.4	Autocorrelation with 4-bit and asymmetric 8-bit quantization corrections	64
6.1	Digital systems of MWA Phase I/II and MWA Phase III	70
7.1	MWAX autocorrelation digital corrections	77
7.2	MWAX autocorrelations from each pointing	79
7.3	MWAX autocorrelations from two nights	80
7.4	MWAX autocorrelations for 90 meter cable	84
7.5	MWAX autocorrelation lumpy structure	86
7.6	MWAX autocorrelation polarization differences	88
7.7	MWAX autocorrelation undetected RFI	90
7.8	MWAX autocorrelation coarse band dips	92
7.9	MWAX autocorrelation coarse band spikes	94
8.1	Calibration solution terms	105
10.1	Fine channel flagging schemes	118
10.2	Edge channel flagging comparison	121
10.3	Comparison of even vs uneven frequencies after averaging	124
10.4	Comparison of center-out averaging schemes	126
10.5	Comparison with averaging by a factor of 8	127
10.6	Evaluation of k_{\parallel} sampling	128
11.1	SSINS data rampage	134
11.2	SSINS RFI miscategorized as DTV	136
11.3	SSINS settings comparison	138
11.4	SSINS RFI cuts	140
11.5	SSINS cuts per pointing for each night	144
11.6	Means of calibration amplitude fit coefficients	146
11.7	Means of calibration phase fit coefficients	147
11.8	Standard deviations of calibration amplitude fit coefficients	148
11.9	Standard deviations of calibration phase fit coefficients	149
11.10	Wall of Shame	154
11.11	Power spectra of integrated observations with narrowband RFI	156
11.12	2D power spectra of the limit set	157
12.1	2D power spectra of the limit set for three redshifts	160

12.2	1D power spectrum upper limit	162
12.3	1D power spectrum from Wilensky et al. (2023)	163

LIST OF TABLES

Table Number		Page
8.1	MWA instrument pointings	97
9.1	Azure compute cluster virtual machine specifications	112
12.1	Table of upper limits for the East-West polarization	167
12.2	Table of upper limits for the North-South polarization	168

GLOSSARY

DFT: Discrete Fourier Transform

DTV: Digital Television

EOR: Epoch of Reionization

ϵ PPSILON: error propagated power spectrum with interleaved observed noise

FFT: Fast Fourier Transform

FHD: Fast Holographic Deconvolution

MWA: Murchison Widefield Array

PFB: polyphase filterbank

SSINS: Sky-Subtracted Incoherent Noise Spectra

ACKNOWLEDGMENTS

It has been an honor to pursue my graduate degree at the University of Washington. I first acknowledge my advisor Miguel Morales, who offered me excellent and challenging projects and who modeled an admirable dedication to getting the correct result rather than the ‘best’ result. I am a far better scientist after working with the UW EoR group.

My co-advisor Bryna Hazelton patiently explained concepts until I had fully grocked them and spent countless hours helping me wrangle Git, python, and IDL. I am so fortunate to have learned best practices in scientific analysis and software development from such a skilled role model and mentor.

Many members of the UW physics community have offered me invaluable support and instruction through my graduate career. A special thanks to my supervisory committee Miguel Morales, Sarah Tuttle, Alvaro Chavarria, Anna Goussiou, Marilena Loverde, and Marcel den Nijs for supporting my pursuit of a graduate degree. I am grateful to Catherine Provost for all the help navigating logistics of the degree process.

My first year advisor Ann Nelson encouraged me to persist despite difficulties. Her impact was great and she will be forever missed.

I am grateful for the collaboration and camaraderie with my fellow research group members: Ruby Byrne, Mike Wilensky, Dara Storer, Nicel Mohammed-Hinds, and Eli Lilleskov. My cohort contained too many excellent colleagues to list here whose impact on my graduate school experience is unmeasured.

It was an honor to mentor and learn from undergraduate students who all did such excellent work—Joseph Chaghouri, Cyrus Goodman, Ellie McLaurin, Adele Meyers, Aiden Morson, Matthew Rosencrans, Audrey Warn, and Catalina Zamora.

I am fortunate to have been a member of the international MWA EoR collaboration. In particular I would like to thank Brian Crosse, Jake Jones, Sam McSweeney, Dev Null, Greg Sleaf, Randall Wayth, and Andrew Williams who made this work possible by answering my many questions about the MWA digital system.

Nichole Barry offered me the invaluable opportunity to work alongside our MWA collaborators at Curtin University for two weeks, for which I am grateful. There are too many people at Curtin to recognize individually who welcomed me and shared their knowledge. A special thanks to fellow visiting students Theo Kunicki and Aman Choksi for being excellent companions during my stay in Australia.

Throughout my education I have had the opportunity to learn physics from many brilliant professors. In particular I recognize Krishna Chowdary who shepherded my academic path and Neil Switz who tirelessly mentored my graduate school application process.

There is not enough space to list the friends and family who have supported my journey. I will forever appreciate Sasha Porter for traveling the road with me and David Porter for believing in my ability to write this accounting of my work.

The vast amount of compute resources that went into this work were generously sponsored by the collaboration between the University of Washington's eScience Institute, UW Research Computing, and Microsoft Azure.

This scientific work makes use of *Inyarrimanha Ilgari Bundara*, the Murchison Radio-astronomy Observatory operated by CSIRO. I acknowledge the Wajarri Yamatji people as the traditional owners of the Observatory site. Support for the operation of the MWA is provided by the Australian Government (NCRIS), under a contract to Curtin University administered by Astronomy Australia Limited. I acknowledge the Pawsey Supercomputing Centre which is supported by the Western Australian and Australian Governments.

DEDICATION

To eight year-old Pyxie, who would have loved to hear the story of the very first stars.

Chapter 1

INTRODUCTION

When I joined the Epoch of Reionization group at the University of Washington, I was told two things. First, that we were looking for the signals from the very first stars. And second, that we were probably not going to find them while I completed my graduate work. These statements immediately drew me to this area of scientific discovery. I love a difficult problem, and the idea of attempting to peer into the very earliest beginnings of the structures that surround us is endlessly appealing. The Epoch of Reionization is a mysterious time, and discovering its secrets is the goal of many groups across a range of institutions and continents.

1.1 The early universe

To understand the Epoch of Reionization, we must start at the very beginning and walk through the early universe. This history is visualized in Figure 1.1. After the Big Bang, the universe was a hot primordial soup of subatomic particles. After a period of rapid inflation, the universe continued to expand and cool, quickly cooling enough that helium and hydrogen nuclei were able to form. Due to the density of photons and matter, the hydrogen was primarily ionized. If a hydrogen managed to capture an electron in its outer shell, an imminent photon collision would knock the electron off before it was able to descend to the hydrogen's ground state orbital. These absorption and emission interactions kept the hydrogen ionized and the electrons free. The photons could not move freely through space without scattering off electrons.




Figure 1.1: Diagram showing the evolution of the universe from an early primordial soup to the reionized universe we see today. The Cosmic Microwave background marks the beginning of the Dark Ages, around 375,000 years after the Big Bang. We seek to probe the Epoch of Reionization, the transition between the Dark Ages and the fully reionized universe. During this era the very first stars began to form, beginning the web of bright objects that surrounds us today. credit: C. Faucher-Giguère, A. Lidz, and L. Hernquist, *Science* 319, 5859 (47).

As the universe expanded and cooled, however, enough space opened up that the hydrogen nuclei and free electrons were able to combine without being separated by a photon collision. With this recombination, the number of free electrons plummeted and the photons were able to stream freely. This process was almost instantaneous and we can see the moment in which the photons were freed; it is encoded in the Cosmic Microwave Background (CMB) around 375,000 years after the Big Bang. At this moment the universe, which had been bright with photon emission, became relatively dark.

During these Dark Ages, regions with slight overdensities, first identifiable in the CMB, became more pronounced. These pockets of clumping matter continued to be fed by gravity until they became dense enough to trigger the formation of stars and black holes. The stars began to ignite and the black hole accretion disks started to glow, releasing ultraviolet radiation which ionized the surrounding hydrogen. This beginning of the Epoch of Reionization contained the inception of the immense web of stars and galaxies we see in the universe today. The ‘bubbles’ of ionized hydrogen grew until all the hydrogen between the stars and structures became ionized. About ~ 1 billion years after the universe began, the Epoch of Reionization ended, leaving in its wake the fundamental structure of our universe today.

Our goal is to peer into this Epoch of Reionization and witness the transition to an ionized universe. This would not only allow us to map the early universe but also distinguish between stars and black holes as the primary drivers of reionization¹, constrain theories of the early universe, and inform mechanics of star and galaxy formation. In order to explore this pivotal moment in the early universe, we measure the 21 centimeter photon signal emitted by the neutral hydrogen.

1.2 Observing the Epoch of Reionization with the 21 centimeter hydrogen line

We probe the Epoch of Reionization by detecting 21 centimeter photons emitted by ground state neutral hydrogen during an extremely rare change of quantum state. By measuring emission from neutral hydrogen we reveal the negative space: the dark growing bubbles of ionized hydrogen.

¹While theorists are fairly confident that reionization was primarily due to stars, it is important for observers to check that this was the case.

1.2.1 *The 21 centimeter emission*

Neutral hydrogen can emit low energy photons during a ‘forbidden’ quantum transition in the ground state. Specifically, the magnetic moment of the ground state electron interacts with the magnetic field generated by the proton’s magnetic moment. This interaction generates four possible quantum spin states, visualized in Figure 1.2. There are three degenerate spin states with higher energy and one with lower energy, splitting the ground state hydrogen energy level into two ‘hyperfine’ energy levels. When hydrogen transitions from an upper energy state to the lower energy state, it emits a photon with frequency 1420 MHz and wavelength 21 cm. Though this transition is extremely rare, the concentration of hydrogen during the Epoch of Reionization is so high that the 21 centimeter emission is potentially detectable. Measuring this signal will allow us to map the distribution of neutral hydrogen during the formation of the bright universe.

1.2.2 *Measuring the 21 centimeter signal*

The 21 centimeter photon wavelength is redshifted as the photon travels through time and space so that what we observe today has a wavelength of $\sim 1 - 2$ meters, putting it in the radio regime at $\sim 150 - 300$ MHz. The time period of the Epoch of Reionization is not well constrained, though analysis of high redshift quasars suggests reionization ended by a redshift of $z \lesssim 6$ (Bosman et al., 2022).

The signal from the neutral hydrogen spin transition is very faint, especially when compared to the bright radio emission from foreground sources such as active galactic nuclei (AGN), radio galaxies, and galactic synchrotron emission. These foreground radio signals are 4-5 orders of magnitude brighter than the EoR signal and must be separated from the faint 21 centimeter emission to make a detection possible. Additionally, just the foreground noise can be high enough to overwhelm the 21-centimeter signal, so we require a highly sensitive instrument in order to make a measurement.

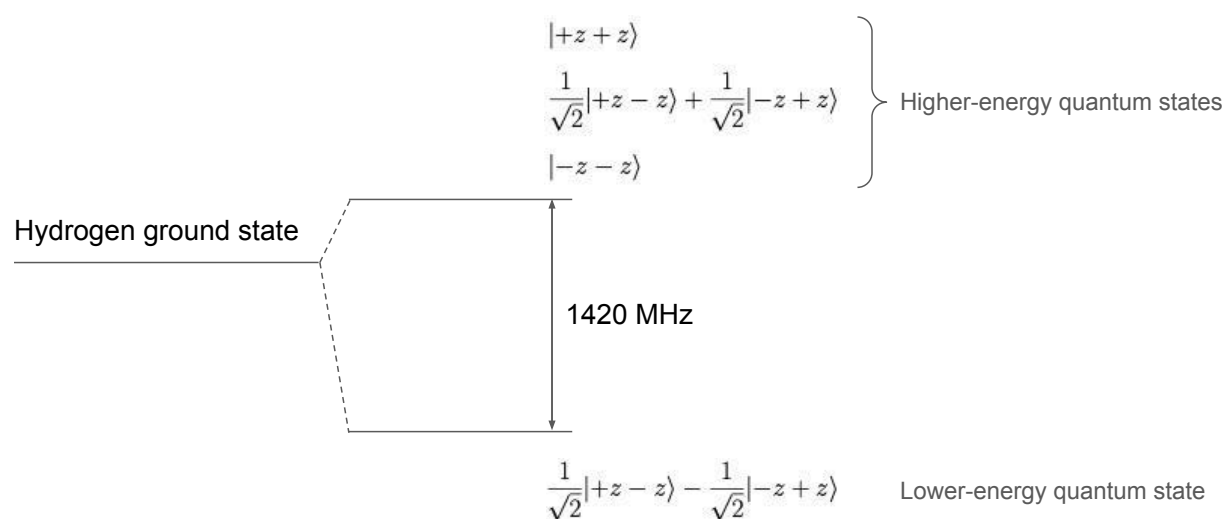


Figure 1.2: The ground state of hydrogen is split into two ‘hyperfine’ energy levels. The electron magnetic moment interacts with the magnetic field generated by the proton magnetic moment, creating four quantum spin states. Three of these states are degenerate, that is, having the same energy. These higher energy states are separated from a lower energy state by 1420 MHz. When the ground state hydrogen transitions from a higher energy state to the lower energy state, it emits a photon with frequency 1420 MHz and wavelength 21 cm.

Ideally, we would like to generate an image of space during the Epoch of Reionization. However, EoR instruments have not yet reached the sensitivity necessary to reconstruct that image. An alternate avenue is to make a statistical measurement from which we can infer the size and distribution of the ionized ‘bubbles’ as they form. This statistical analysis is what we perform here, and is described in more detail in Chapter 8.

Rather than a direct detection—for which EoR experiments have not yet reached the necessary sensitivity—EoR analyses typically report ‘upper limits’ which are a measure of the sensitivity limits of the respective experiment and analysis. An upper limit describes how bright the EoR signal would need to be in order for it to be detected. That is, an upper limit constrains the brightness of the EoR signal; if it were brighter than the upper limit, it would be detected by our experiments. An analysis which sets a new lowest upper limit has taken the science into a more sensitive regime for EoR detection.

The sensitivity required for EoR science makes radio interferometers a natural choice for measuring the 21 centimeter signal. A radio interferometer is a collection of antennas which measure the electric field signal from the sky. The measurements from these antennas are correlated, or cross-multiplied. Having multiple antennas increases the collecting area of the telescope, thus increasing sensitivity. If the noise on the antenna measurements is independent, cross-correlation will result in a measurement noise which is zero-mean. Averaging the measurements together will then cause the noise to drop. The sensitivity from high collecting area and low measurement noise has placed radio interferometer telescopes at the forefront of EoR detection.

Multiple experiments have been used to produce limits on the EoR including the Donald C. Backer Precision Array for Probing the Epoch of Reionization (PAPER) (Kolopanis et al., 2019), the LOw Frequency ARray (LOFAR) (Mertens et al., 2020), the Long Wavelength Array (LWA) (Eastwood et al., 2018), and the Hydrogen Epoch of Reionization Ar-

ray (HERA)(Collaboration et al., 2022b,a). Additionally, a possible detection of the global (sky-averaged) 21 centimeter signal was made by the Experiment to Detect the Global Eor Signature (EDGES) (Bowman et al., 2018).

Our telescope of interest, the Murchison Widefield Array (MWA), has been at the forefront of EoR science since being commissioned in 2013. Some of the most competitive limits in the field have come from the international EoR MWA collaboration (Wilensky et al., 2023; Rahimi et al., 2021; Trott et al., 2020; Barry et al., 2019b; Li et al., 2019; Beardsley et al., 2016; Ewall-Wice et al., 2016b; Dillon et al., 2015).

1.3 The Murchison Widefield Array

The Murchison Widefield Array (MWA) (Tingay et al., 2013; Wayth et al., 2018), shown in Figure 1.3, is a radio interferometer installed at *Inyarrimanha Ilgari Bundara*, the Murchison Radio-astronomy Observatory. It is located on traditional Wajarri Yamaji lands in the Western Australia Murchison shire. This radio-quiet region with expanses of flat landscape is ideal for sprawling radio arrays such as the MWA, and is the future home of SKA-Low. The first SKA-Low antennas were installed in March 2024, beginning the construction of what will be the largest radio telescope ever built. Results from the MWA have directly contributed to the design of this upcoming instrument (Trott and Wayth, 2016).

An MWA antenna, or tile, comprises 16 dipoles anchored to a 5 meter x 5 meter mesh ground screen. There are 256 antenna tiles, though due to current hardware limitations only 128 are actively used in any single observation. Future upgrades (described in Chapter 6) will make it possible to utilize all 256 antenna tiles at once. As we will discuss in detail in Chapter 3, the antenna electric field measurements are sent through analog conditioning and digital signal processing before arriving at the correlator, where the measurements are cross-correlated and output data sets are written.



Figure 1.3: Several antenna tiles of the Murchison Widefield Array (MWA), installed at *Inyarrimanha Ilgari Bundara*, the Murchison Radio-astronomy Observatory. The MWA is located on Wajarri Yamaji Country in Western Australia. The upper image shows a group of antenna tiles arranged in a regular hexagonal pattern. The lower image shows kangaroos visiting the telescope. Image credit: ICRAR/Curtin (accessed from the MWA multimedia archive).

Since the MWA was commissioned in 2013, analyses of MWA EoR data have not only led to some of the most stringent upper limits but have also been used to make important contributions to intersecting fields including radio frequency interference detection, calibration, analysis software development, and ionospheric metrics (Wilensky et al., 2019; Byrne et al., 2021a; Barry et al., 2019a; Jordan et al., 2017). Additionally, the MWA supports multiple science objectives in addition to EoR detection including understanding solar bursts, mapping galaxies, and detecting transient sources such as pulsars, fast radio bursts, and cosmic rays. The work we present here benefits not only the international MWA EoR collaboration but also a wide range of science research goals.

1.4 Precision EoR science with the MWA

In this work, we focus on the digital system response of the MWA. Removing the telescope response from the data is particularly crucial given the sensitivity required to measure the 21 cm signal. We must be able to describe the antenna response with exquisite precision as we discuss in Chapter 2. While our data processing pipeline removes linear antenna responses from the data, corruption from nonlinear digital effects remain. We describe impacts of nonlinear responses due to MWA digital quantization stages in Chapter 3. These effects must be removed before we analyze our data. In Chapter 4 and Chapter 5 we present corrections we developed to remove nonlinear digital effects from the Legacy correlator used during MWA Phase I and II (2013-2021).

Our work demonstrating the impacts of digital nonlinearities and developing corrections for their effects influenced the design of an exciting upgrade to the MWA digital system, the new MWAX correlator, described in Chapter 6. One of the key specifications for the new correlator was eliminating the digital quantization stages which were a primary source of nonlinear antenna response. The first EoR data with the new MWAX correlator was taken in 2021, and we evaluate the antenna response in this novel data set in Chapter 7.

Finally, we prepare to obtain the first upper limit for MWA Phase III with the 2021 MWAX data set in Chapter 8. This process requires setting up compute power (Chapter 9), making analysis decisions (Chapter 10), and selecting quality data (Chapter 11). We present the results of our limit analysis in Chapter 12, where we are able to see the impacts of our work with telescope digital systems on our ability to detect the signals of the very first stars.

Chapter 2

PRECISION CALIBRATION FOR EOR SCIENCE

For an interferometer such as the MWA, the antenna measurements of the sky signal travel through the instrument electronic systems, which we will describe in Chapter 3. They are then correlated, a process in which the electric field measurement from one antenna is multiplied by the complex conjugate of the measurement from another antenna (or itself). The products are summed to a desired integration time resolution to produce complex cross powers called visibilities. The visibility between antenna 1 and antenna 2 is thus defined as follows, where N is the number of samples necessary to achieve the integration time:

$$V_{12} = \sum_{n=1}^N E_{1n} E_{2n}^*. \quad (2.1)$$

For example, to achieve an integration time of 0.5 seconds in an MWA frequency channel of width 10 kHz, 5,000 samples are summed. The integration times we use for EoR science, typically between 0.5-2 seconds, are short enough that the electric field from the sky can be treated as stable during the time interval. For each integration time interval, the visibilities are recorded and written out as the telescope's final data product.

Visibilities contain not only the antennas' measurements of the sky but also the signature of the electrical system those measurements passed through. We use calibration to remove this instrument response so that our data is as close to the true sky measurement as possible.

2.1 Calibration overview

Calibration is our process for removing the instrument response to recover the sky signal. We do this by relating our measured visibility between antennas 1 and 2 (V_{12}) to what we

would get if we had measured only the sky (u_{12}). This relationship can be described by the equation

$$V_{12} = g_1 g_2^* u_{12} + n_{12} \quad (2.2)$$

where g_1 is the (complex) instrument response, or gain, of antenna 1, g_2^* is the complex conjugate of the gain of antenna 2, and n_{12} is the measurement noise. With this equation, we are essentially stating that our measured visibilities are the true sky signal multiplied by the responses of each of the relevant antennas. In order to get our data as close to the true sky measurement as possible, we need to divide out the antenna responses. The goal of calibration is to obtain the best possible estimation of the antenna responses, or gains, so that we can remove them from the data.

Generally, three types of calibration are used for 21-centimeter EoR science: sky-based, redundant, and hybrid or combined. Sky-based calibration generates a model of the sky in order to solve Equation (2.2) and is sensitive to errors in that model (Barry et al., 2016; Ewall-Wice et al., 2016a). Redundant calibration can be used for interferometers with regularly spaced antennas, and constrains Equation (2.2) by assuming that visibilities from redundant antenna pairs, or baselines, will be equal (Wieringa, 1992; Parsons et al., 2012; Liu and Shaw, 2020). This approach is vulnerable to sky model errors as well as variations between antennas within the same redundant baseline group (Byrne et al., 2019). Hybrid or combined calibration approaches blend sky-based and redundant calibration (Sievers, 2017; Li et al., 2018). Calibration is an active field of study, with novel approaches such as unified calibration (Byrne et al., 2021a) currently being developed.

Our work utilizes sky-based calibration, which has been the most successful method for MWA data. We describe our implementation in detail in Section 8.3.1, and here we give a brief overview. In Equation (2.2), we don't know u_{12} . For sky-based calibration, we use a list of sources on the sky along with our knowledge of the instrument's collecting area (essentially, where the antennas are physically located as well as their extent) to generate

a model m_{12} of what we think our measurements would be if they contained no instrument response. That is, we are simulating a set of clean data, or model visibilities. With these model visibilities, we can use mathematical methods (Salvini and Wijnholds, 2014) to solve for g_1 and g_2^* by minimizing

$$\chi^2 = |V_{12} - g_1 g_2^* m_{12}|^2, \quad (2.3)$$

where the above equation assumes the noise n_{12} is complex with an expectation value of zero. Our science requires that calibration estimate the antenna response with exquisite precision.

2.2 Calibration precision requirements

The precision of our calibration ultimately impacts our ability to make a measurement of the EoR. Any difference between the calibration gains and the true antenna response results in residual antenna artifacts in the data. The effects of residual antenna response are particularly evident in one of the powerful diagnostics of our analysis—the 2D power spectrum.

2.2.1 The 2D power spectrum

Our instruments are not yet sensitive enough for us to detect the EoR in images of the sky. Instead, we perform a statistical analysis to identify the size and distribution of the bubbles of ionized hydrogen. We provide a brief overview of our analysis here, describing the full process in detail in Chapter 8.

Using our knowledge of the MWA antenna locations and extents, we grid visibility data to a single plane of evenly spaced uv pixels at each frequency. We then take an angular Fourier transform of each uv -plane to obtain an image of the sky at every frequency. While we cannot detect the EoR in image space, this space is conducive to combining data from different observations.

For each frequency, we sum together images from multiple observations. Then we take a Fourier transform of these integrated images along the image axes θ_x, θ_y to obtain a k_x, k_y

plane at each frequency. Frequency corresponds to redshift, so the set of k_x, k_y planes shows the evolution of the universe over time. We next take a Fourier transform along the frequency axis to k_z , also referred to as k_{\parallel} as this axis is parallel to the line of sight. We then compute the power spectrum and average spherically over k_x, k_y, k_z to obtain a statistical measurement at each k mode.

If instead of averaging spherically we average k_x and k_y cylindrically, we obtain a 2D power spectrum with axes parallel and perpendicular to the line of sight, k_{\parallel} and k_{\perp} , respectively. This provides a powerful diagnostic space where we can examine the behavior along the k_{\parallel} axis after the frequency Fourier transform.

The EoR signal we are attempting to measure is orders of magnitude fainter than the bright foreground radio emission from sources such as active galactic nuclei (AGN), radio galaxies, and galactic synchrotron emission. The Fourier transform in frequency separates signals into spectral modes based on how quickly they vary in frequency. The bright radio foreground signals are quite smooth, or slowly changing in frequency, and so they fall into the lowest, near constant emission k_{\parallel} modes. The EoR signal has more fluctuations and so is allocated to higher spectral modes. However, a radio interferometer is inherently chromatic, having a frequency-dependent response. This response mixes with the foreground power, spreading it into higher modes (Furlanetto et al. 2006; Datta et al. 2010; Morales et al. 2012; Parsons et al. 2012; Trott et al. 2012; Thyagarajan et al. 2013; Hazelton et al. 2013). We visualize this in the cartoon 2D power spectrum in Figure 2.1. The bleeding of foreground power is mostly limited by the instrument’s primary field of view (dashed black line), and is ultimately contained by the horizon (solid black line), creating a ‘Wedge’ of contamination. The uncontaminated region above the the Wedge is where we expect to make an EoR measurement, so we refer to it as the ‘EoR Window’.

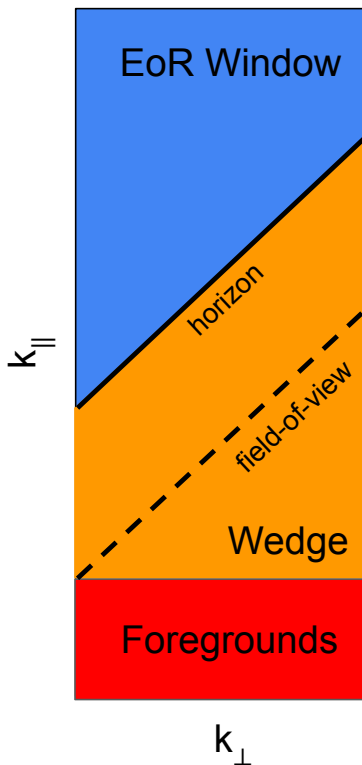


Figure 2.1: Our primary diagnostic tool for EoR science is the 2-dimensional power spectrum. A cartoon version of this plot is shown here. After Fourier transforming from frequency to k_z , or k_{\parallel} , the power spectrum is computed. The coordinates perpendicular to the line-of-sight, k_x and k_y , are averaged to give k_{\perp} . Astronomical foregrounds, which are relatively smooth in frequency, occupy the lowest k_{\parallel} bins after the Fourier transform. However, the frequency response of the instrument moves this power into higher k_{\parallel} bins, resulting in a ‘Wedge’ bounded by the horizon, indicated by the solid diagonal line. The dashed diagonal line marks the region of the instrument’s primary field-of-view. Above the horizon line is the relatively uncontaminated region we refer to as the ‘EoR Window’. In our analysis, we look for improvements to sensitivity in this EoR window.

2.2.2 Antenna response precision

If our antenna response is variable in frequency, then it mixes with the bright foreground power and causes it to leak into the EoR Window after the Fourier transform. This power bleed can completely overwhelm the subtle EoR signal. In order to make an EoR detection, we need to compute the antenna response with high enough precision to identify and remove any frequency variations. Specifically, Barry et al. (2016) found that our calibration must capture the antenna spectral structure with a precision of a part in 10^{-5} .

An impediment to identifying the antenna response to the required precision of a part in 10^{-5} is a fundamental assumption we make in calibration. Specifically, Equation (2.2) is a linear approximation of the full equation

$$V_{12} = n_{12} + g_1 g_2^* u_{12} + p_1 p_2^* u_{12}^2 + q_1 q_2^* u_{12}^3 + \dots \quad (2.4)$$

where p_1, q_1, \dots are the quadratic, cubic, and higher order terms of the instrument response. That is, we assume that the antenna responses are linear, ignoring the higher order terms in the above equation. This is always an imperfect assumption, though engineering efforts are dedicated to minimizing nonlinear terms in the analog system. While efforts have been focused on eliminating analog nonlinearities, the digital system has been presumed to have a linear response. Any residual nonlinear terms are small enough that for most science experiments it is reasonable to ignore them. For our EoR science, however, the calibration precision requirements mean that digital nonlinear terms have non-negligible impacts.

In order to make an EoR measurement, we must remove nonlinear antenna responses before attempting to calibrate. As we shall discuss in the next chapter, quantization stages in the telescope signal path generate a nonlinear antenna response which we clearly see in our data. Understanding and eliminating these nonlinear effects is the primary goal of this work, and could greatly enhance the science precision of the MWA.

Chapter 3

NONLINEAR QUANTIZATION EFFECTS IN MWA DATA

As discussed in the previous chapter, our science requires calibration to capture spectral structure to a precision of a part in 10^{-5} . And as seen in Equation (2.2), calibration approximates the instrument response as linear. In many science cases, this approximation is sufficient. For our precision needs, however, small nonlinearities in the instrument response have detrimental impacts on our analysis. A known source of nonlinear effects is digital quantization. The MWA, as do most instruments, has multiple requantization stages throughout the signal path. The evidence of nonlinear antenna response from these requantization stages is visible in autocorrelations, calibration solutions from crosscorrelations, and 2D power spectra. In Chapters 4 and 5 we will describe corrections for these effects.

3.1 *Digital quantization nonlinearities*

The process of quantization in a digital system involves truncating a signal into some specific number of bits. This may be an initial quantization stage, in which an analog signal is captured into a set of bits at some time cadence to create a series of measurements. Subsequent requantization stages also occur throughout the digital signal chain. For example, integers are often promoted to floats with more bits before calculations and then re-quantized to integers after the computations are complete. Requantization is also common when transmitting data between various signal processing components, as transferring smaller numbers is more efficient.

We can think of quantization as a binning process. For quantization into 4 bits, the possible bins are integers in the inclusive range $[-7, 7]$ (or $[-7, 8]$, or $[-8, 7]$, depending on

the implementation). Based on some sorting criteria, numbers are binned into these integers. This binning fundamentally alters the probability distribution of the measured signal, as we will discuss in more detail in Chapter 4. Importantly, the changes to the signal statistics are nonlinear and cannot be captured and removed by calibration.

It is useful to describe the MWA signal path in order to understand how quantization is taking place. Figure 3.1 shows an overview of the system components. The system can be coarsely divided into three sections: antennas, receivers containing both analog and digital hardware, and the correlator which ingests digital signals from the receivers and outputs our final data product, visibilities.

3.2 Antenna

Each MWA antenna, or tile, comprises 16 dipoles anchored in a regular grid to a mesh ground screen (Tingay et al., 2013; Lonsdale et al., 2009). Each dipole measures two polarizations of sky signal and an analog beamformer collects the measurements from all the tile’s dipoles for each polarization. Figure 3.2 shows an image of an MWA tile and beamformer. Within the beamformer, analog delays are applied to the dipole measurements which effectively ‘point’ the tile at a predetermined location in the sky (see Table 8.1 for more information on MWA pointings). This results in a single radio frequency beamformed signal for each antenna tile and each polarization. Measurements of the two polarizations are then sent over coaxial cable to the receiver boxes which house both an analog conditioning module and a digital signal processing module referred to as the ‘digital receiver’. An optional whitening filter is installed on the longer of these coaxial cables to mitigate frequency-dependent signal loss.

3.3 Receiver

Each receiver processes signals from both polarizations of 8 antennas. For a detailed description of the MWA receiver implementation, see Prabu et al. (2015). Figure 3.3 shows two receivers in the field. In addition to the analog and digital components which will be

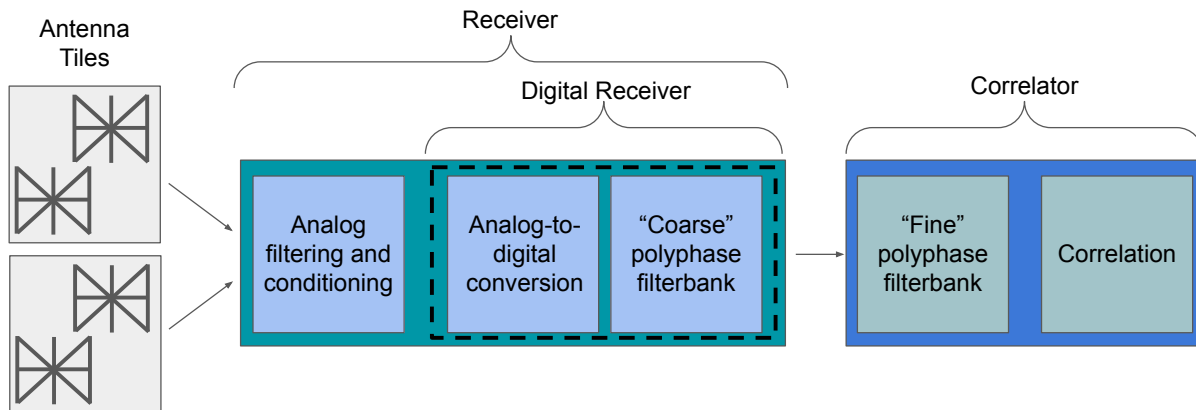


Figure 3.1: An overview of the MWA signal path. Antenna electric field measurements travel to receivers which house both analog and digital components. After analog signal conditioning (amplification and band limiting), the signal is quantized in the analog-to-digital conversion. The quantized signal is channelized by the ‘coarse’ polyphase filterbank (PFB) before being sent to the correlator. In the correlator the data is further channelized by the ‘fine’ PFB before being correlated to produce visibilities. These visibilities are written out as the final data product.



Figure 3.2: An MWA antenna tile, beamformer, and solar panel. The antenna comprises 16 dipoles attached to a mesh screen. They are connected to a beamformer, the small white box to the right of the tile. The beamformer applies delays to each dipole's measurement which effectively 'points' the antenna at some location on the sky. The beamformer then collects the dipole electric field measurements into a single antenna measurement for each polarization. This particular tile is located far from the center of the array, so its beamformer is powered by a solar panel. Image credit: ICRAR/Curtin (accessed from the MWA multimedia archive).



Figure 3.3: Two MWA receivers in sun shelters. Each receiver processes signals from 8 antenna tiles and houses analog and digital signal processing modules. Image credit: ICRAR/Curtin (accessed from the MWA multimedia archive).

explained in more detail below, the receiver contains other components key for processing, including an air conditioning system to protect its contents from overheating, a clock for sampling and synchronization, connections to monitor and control systems, and a power supply for its internal components as well as the antenna analog beamformers and amplifiers.

3.3.1 Analog conditioning

Within the receiver, the analog signals are first amplified. A filter is then applied which effectively limits the bandwidth to 80-300 MHz. This conditioned analog signal travels to the digital receiver.

3.3.2 Digital receiver

The MWA digital receiver converts the signal from analog to digital, sampling the analog signal at a rate of 655.36 MHz and quantizing it into 8-bit digital samples. This is the initial

quantization stage of the MWA. The 8-bit digitized streams are sent to a polyphase filterbank where they are subdivided into frequency channels. Below we provide a brief description of polyphase filterbanks before discussing the specific MWA PFB implementation.

Polyphase filterbanks

Channelization of data into frequency bins can be performed by Fourier transforming a sequence which is a function of time into a sequence which is a function of frequency. However, selecting a sequence of data to be transformed effectively applies a square window before the transform. Multiplication of a top hat function in the time domain is equivalent to convolving with a sinc function in the frequency domain and thus results in ‘ringing’, or strong spectral response, in the Fourier transform output. To avoid this, convolution windows or filters may be applied before the Fourier transform to smooth the output response. A common filtering and Fourier transform implementation which reduces spectral structure is the polyphase filterbank.

For visualizing and understanding polyphase filterbanks, I have found Harris and Haines (2011), Price (2016), and Price’s PFB tutorials¹ to be especially useful. Generally, a PFB entails a filtering and summing stage followed by a Fourier transform. We can consider a toy model of a PFB with $M = 2$ ‘taps’ and $P = 3$ ‘points’. This PFB has a set of filter coefficients of length $N = M \times P$: $\{h_1, h_2, \dots, h_6\}$. These coefficients can be generated by sampling a continuous filter at regular intervals. A set of N consecutive samples is selected from the data stream to give $\{x_1, x_2, \dots, x_6\}$. Conceptually, the data and the filter coefficients are both split into M subsets of length P . These are multiplied together to give, in our example, the subsets $\{h_1x_1, h_2x_2, h_3x_3\}$ and $\{h_4x_4, h_5x_5, h_6x_6\}$. These products are summed together to generate a set of length P : $y = \{h_1x_1 + h_4x_4, h_2x_2 + h_5x_5, h_3x_3 + h_6x_6\}$. A Fourier transform is performed on y to generate the output $\{X_1, X_2, X_3\}$. For our discussion we assume that a Fast Fourier Transform (FFT) is used.

¹https://github.com/telegraphic/pfb_introduction

The conceptual model described above is shown in Figure 3.4. An actual implementation is represented in Figure 3.5, in which the initial set of samples is selected into P groups of length M : $\{x_1, x_4\}$, $\{x_2, x_5\}$, $\{x_3, x_6\}$. These are each sent to a respective subfilter where they are multiplied by the corresponding filter coefficients and then summed. The outputs of the subfilters are described by the set y above, and are sent to the 3-point FFT.

If the PFB then selects the next (in our toy model case) 6 samples from the data stream for processing, the PFB is said to be critically sampled. A PFB can be oversampled by sampling at a faster rate, causing the windows of processed data to overlap. In our toy model, for example, this could mean that the next set of samples started at, say, x_5 rather than x_7 .

The MWA digital receiver employs a PFB with 8 taps and a 512-point FFT. The PFB input is a stream of real-valued numbers sampled at a rate of 655.36 MHz, and the output is 256 frequency channels of real-imaginary complex pairs sampled at a rate of 1.28 MHz. Since the PFB FFT input comprises real numbers, half of the output spectrum is redundant and is discarded.

After channelization, a subset of twenty-four 1.28 MHz channels are selected. In each of these channels, data are multiplied by an integer gain specific to the respective coarse channel. The data are then requantized into 5 + 5 bit real/imaginary pairs. The band is subdivided into 3 groups of 8 coarse channels each which are respectively sent via 3 fiber optic cables to the correlator.

3.4 Correlator

Within the correlator, antenna measurements are further channelized via PFB and then correlated, or cross-multiplied and summed. A detailed description of the correlator is provided by Ord et al. (2015), and in-depth discussion of the correlator PFB can be found in McSweeney et al. (2020).

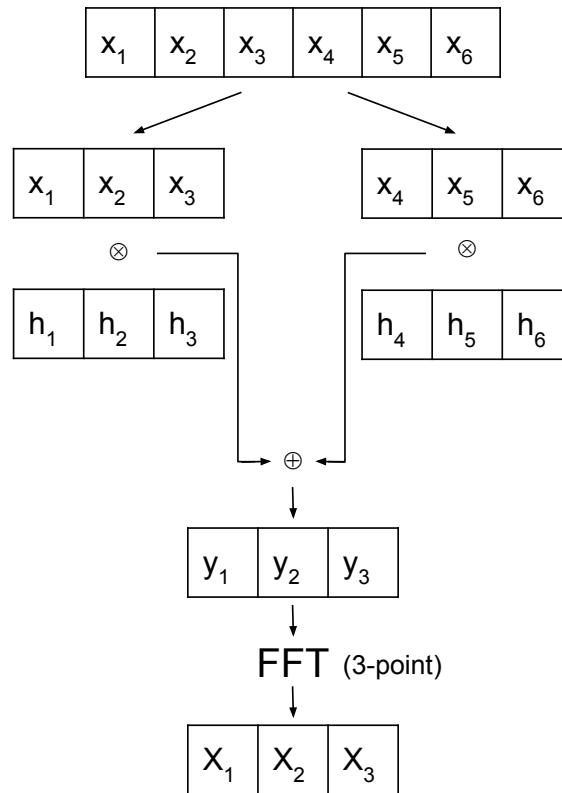


Figure 3.4: A conceptual model of a polyphase filterbank (PFB) with 2 taps and 3 points. A subset of the data stream of length 6, here called $\{x_1, x_2, x_3, x_4, x_5, x_6\}$, is sent to a PFB with a filter of coefficients $\{h_1, h_2, h_3, h_4, h_5, h_6\}$. The data and filter are each divided into 2 groups of length three, and the data and filter are multiplied together in each of these subgroups. The subgroups are then summed with each other, and the output $\{y_1, y_2, y_3\} = \{h_1x_1 + h_4x_4, h_2x_2 + h_5x_5, h_3x_3 + h_6x_6\}$ sent to a 3-point FFT, which produces the final PFB output $\{X_1, X_2, X_3\}$.

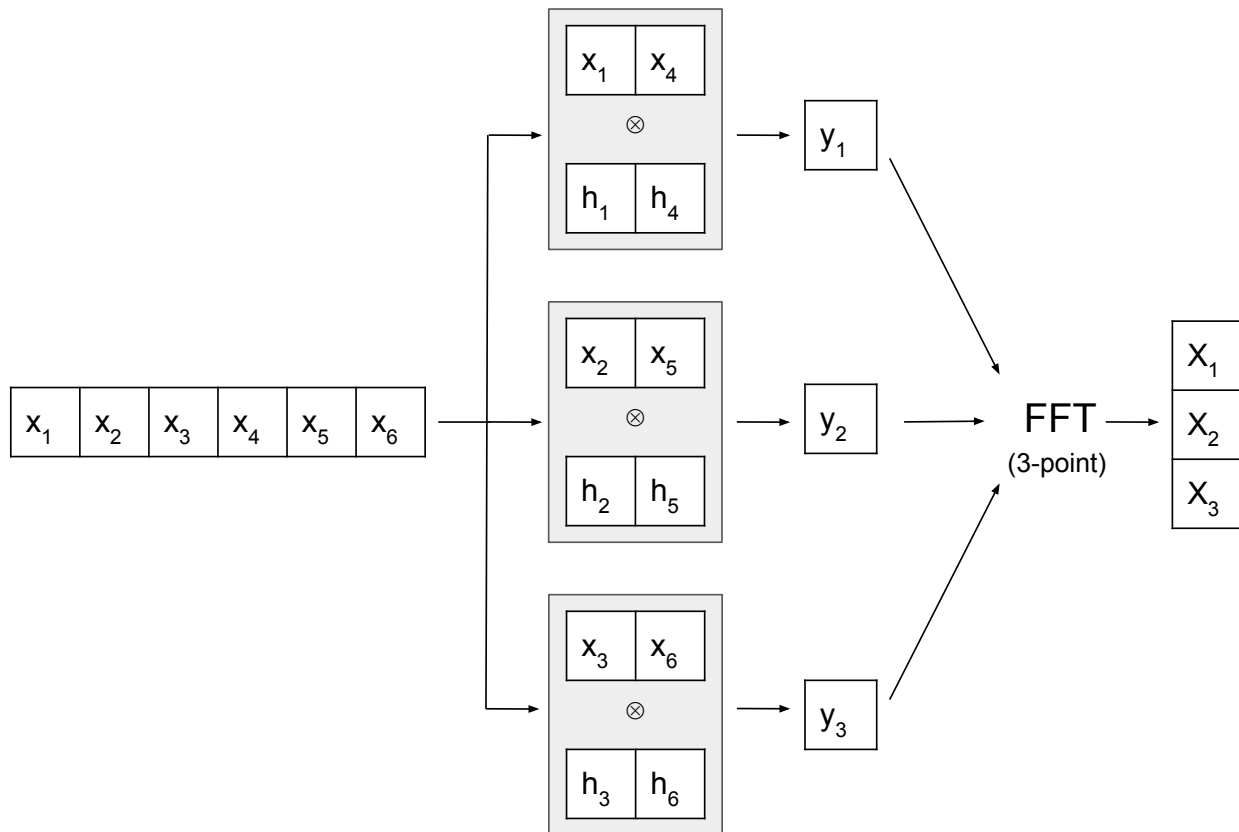


Figure 3.5: An implementation of the conceptual polyphase filterbank (PFB) model shown in Figure 3.4. The PFB has $M = 2$ taps and $P = 3$ points. A set of data of length 6, here called $\{x_1, x_2, x_3, x_4, x_5, x_6\}$, is sent to the polyphase filterbank. There it is distributed between P subfilters (demarcated by the grey boxes) to create subsets of length M . Within each of the three subfilters, the data is multiplied by the appropriate filter coefficients h_i . The products are summed to produce the subfilter outputs $\{y_1, y_2, y_3\} = \{h_1x_1 + h_4x_4, h_2x_2 + h_5x_5, h_3x_3 + h_6x_6\}$. These are sent to the 3-point FFT where they are transformed to give the PFB output $\{X_1, X_2, X_3\}$.

In the correlator the frequency band of 1.28 MHz frequency channels is further divided into finer frequency channels by a second PFB stage. Four PFB boards each process data from 4 receivers, respectively. These 12-tap, 512-point critically sampled PFBs produce frequency channels of width 10 kHz. Outputs from the PFBs are requantized into 4 + 4 bit real/imaginary pairs and sent to GPUs for correlation.

Due to the two-stage channelization implementation of the MWA, we typically describe the digital receiver PFB as the ‘coarse’ PFB and the second PFB in the correlator as the ‘fine’ PFB. The 24 1.28 MHz channels selected after the coarse PFB are referred to as coarse channels or coarse bands, and the 128 10 kHz channels they each respectively comprise after the fine PFB are referred to as fine channels.

Correlation is performed on 24 GPUs (referred to as ‘gpuboxes’), each of which processes data from all antennas for a single coarse band. Before reaching the gpuboxes, each real or imaginary number is promoted from a 4 bit integer to an 8 bit integer. Within the gpuboxes these integers are further promoted to 16 bit floating-point samples. Correlation multiplies the complex signal from some antenna with the complex conjugate of the signal from another antenna (or from itself). As the signals are correlated, the correlations are summed to achieve the desired time resolution of the output data. The correlation products, or visibilities, are either crosscorrelations (between different antennas) or autocorrelations (between the same antenna). With two instrumental polarizations X and Y, the output visibilities have polarizations XX, YY, XY, and YX. These are transferred back to host memory after correlation, where they are averaged in frequency and written out to a bespoke FITS file format referred to as ‘gpubox files’.

3.5 MWA quantization stages

Due to the nonlinear effects of digital quantization, we are particularly concerned with the MWA digital system quantization stages. After the initial quantization of 8 bits within the

ADC, requantization happens throughout the signal path. This is not unique to the MWA; it is standard for instruments to combat bit growth by requantizing signals as they pass between digital processing modules. We show the various quantization and requantization stages within the MWA signal path in Figure 3.6 and discuss their implementations below.

After the ADC, the 8 bit quantized numbers are sign-extended to 9 bit signed integers. Represented in the two's complement style, these quantized numbers are in the inclusive range $[-256, 255]$ and are the inputs to the coarse PFB. The outputs of the coarse PFB are $16 + 16$ bit real/imaginary pairs. After digital gains are applied, bits 15-11 (inclusively) are selected for each real or imaginary number. This generates integers in the inclusive range $[-16, 15]$. These requantized $5 + 5$ bit real/imaginary pairs are sent to the correlator.

After the filtering stage of the fine PFB in the correlator, the data is represented as $24 + 24$ bit real/imaginary pairs of signed integers. These are requantized before the PFB FFT stage as $8 + 8$ bit integers by selecting bits 21-14, inclusively. Importantly, as we shall describe in Chapter 5 in more detail, this requantization stage was implemented incorrectly, resulting in different rounding schemes for positive and negative numbers. The asymmetric rounding implementation has a non-negligible nonlinear impact on the statistics of the PFB output data. We note that it is likely that the coarse PFB also contains a rounding stage between the filter and FFT. However, since it is correctly implemented the nonlinear effects are small enough that we ignore them for this work.

The fine PFB FFT utilizes a fixed point algorithm, after which outputs are downselected to $4 + 4$ bit real/imaginary pairs. These span the inclusive range $[-7, 7]$, with ± 8 denoting corrupted data. Before correlation these integers are promoted and correlation performed at $16 + 16$ bit float resolution.

The MWA employs multiple quantization and requantization stages along the signal path. The nonlinear effects of these stages result in residual spectral structure in the data after

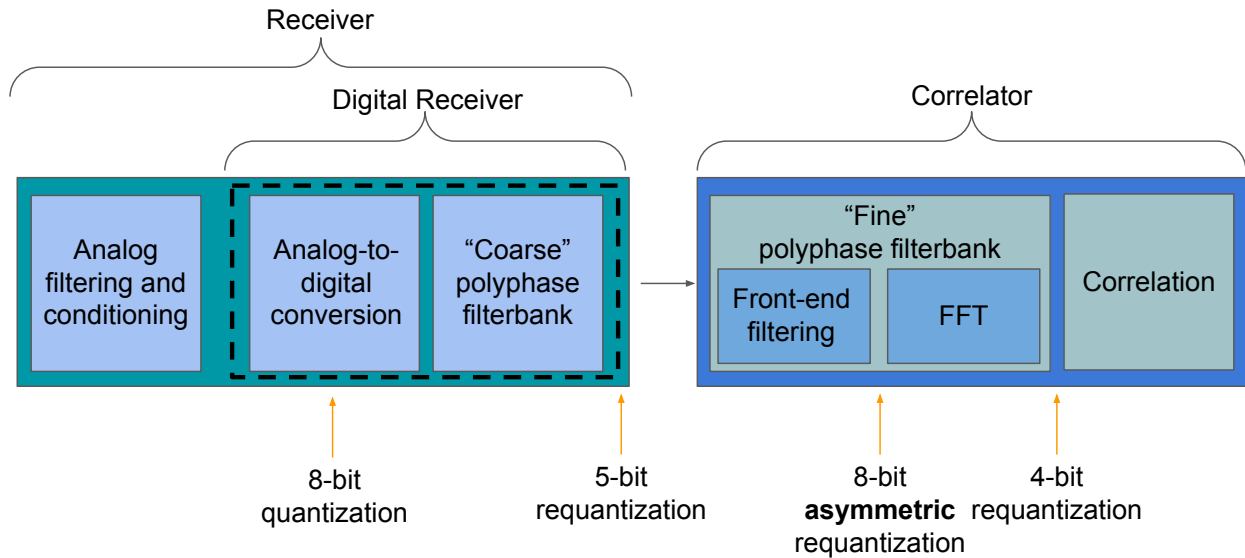


Figure 3.6: We show key quantization and requantization stages in the MWA digital signal path. The first stage is the analog-to-digital conversion which quantizes the analog signal into 8 bit digital numbers. This digital stream is processed through the coarse polyphase filterbank, which outputs 16 + 16 bit real/imaginary pairs. These are then requantized to 5 + 5 bit real/imaginary pairs before being sent to the correlator. Within the correlator’s fine polyphase filterbank is an 8-bit requantization stage after filtering and before the FFT. This stage is implemented incorrectly, resulting in different binning schemes for positive and negative numbers. This asymmetry has significant nonlinear effects on the the data. Finally, after the fine polyphase filterbank FFT, the data is requantized into 4 + 4 bit real/imaginary pairs before correlation. There is likely also a requantization stage between the filter and FFT of the coarse polyphase filterbank, but because it is implemented correctly the effects are small and we do not consider them in this work.

the linear antenna response is removed.

3.6 Effects of nonlinear quantization in MWA data

In MWA data, nonlinear effects from digital quantization interact with the linear digital gains and polyphase filterbank response. During preprocessing of MWA data, we divide out the digital gains as well as the spectral structure from the MWA’s two-stage polyphase filterbank implementation. After removing this linear antenna response, residual artifacts remain in the data. We see clear evidence of this interaction between nonlinear effects and the linear antenna response in autocorrelations, in calibration gains derived using crosscorrelations, and in 2D power spectra.

3.6.1 Nonlinear effects in autocorrelations

An autocorrelation is the correlation of an antenna with itself, and so encodes the antenna response. In the top panel of Figure 3.8 we show an uncorrected (‘raw’) autocorrelation averaged in time across a single 2-minute observation of data from 2013, the first year EoR data was recorded with the MWA. This data has a frequency resolution of 40 kHz. The autocorrelation is ‘raw’, meaning that we have not divided out the digital gains or polyphase filterbank response.

As shown by Levine (2012a,b), the 128 10 kHz fine frequency channels within each 1.28 MHz coarse band carry the signature of the coarse PFB. Levine derived the coefficients describing the bandpass for this response for each of the 128 fine channels, which we have plotted in Figure 3.7. This response is apparent in the raw autocorrelation coarse bands, which are delineated by vertical grey lines in Figure 3.8.

In 2013, a digital gain of 64 was applied to the lower 16 coarse bands, and a gain of 128 applied to the remaining 8 coarse bands. In the raw autocorrelation we can clearly see the

discontinuity between these gains separating the lower two-thirds and upper third of the frequency band.

The bottom panel of Figure 3.8 shows the results after dividing the digital gains and PFB response out of the raw autocorrelation. In the corrected autocorrelation we see residual peaks at the coarse band edges which are artifacts of the PFB spectral structure. These peaks are stronger in the upper 8 channels of the band, where the digital gains are higher. This dependence of residual peak height on the autocorrelation power points to a nonlinear response. We also clearly see a residual ‘jump’ after removing the digital gains, which implies an interaction between the linear gains and a nonlinear antenna response.

3.6.2 Nonlinear effects in crosscorrelation calibration solutions

Residual artifacts are also present in the calibration solutions, which are computed using crosscorrelation data. The fact that calibration cannot fully fit and remove the spectral antenna structure confirms that the residuals visible in the autocorrelations are not due to inaccuracies in the corrections we are applying but rather have a nonlinear source. Close inspection of the calibration artifacts revealed both time- and temperature-dependence (N. Barry, personal communication, March 14, 2017). Barry computed the mean of the artifacts at the coarse band edges for several days of 2013 MWA data and plotted them as a function of beamformer temperature. We show Barry’s plot in Figure 3.9, where each panel represents a different type and length of cable between the tile and receiver. The analog gains are temperature-dependent, and this plot points to a strong correlation between the residual PFB artifacts and the analog gains. The dependence of the residual PFB structure on the temperature-dependent analog gains indicates an interaction with nonlinear effects due to quantization stages.




Figure 3.7: Levine derived the response of the MWA’s two-stage polyphase filterbank architecture. In each 1.28 MHz coarse band, 128 10 kHz fine channels carry the distinct power response which, while relatively flat through the center of the coarse band, drops off steeply at the edges. As we show in autocorrelations in Figure 3.8, this response repeats across the entire frequency band, resulting in a distinctive toothed pattern. Ideally, we could simply divide this response out of the data to smooth in frequency. However, interactions with nonlinearities result in residual artifacts at the coarse band edges.




Figure 3.8: An autocorrelation from a single antenna averaged in time across a single 2-minute observation is shown as a function of frequency. Effects from the two-stage polyphase filterbank (PFB) and the digital gains are obvious in the ‘raw’ autocorrelation (upper panel) which has had no digital corrections applied. Each coarse band, delineated by vertical grey lines, carries the polyphase filterbank response shown in Figure 3.7. For this data the applied digital gain values are 64 for the lower 16 coarse bands and 128 for the upper 8 bands, respectively, resulting in a drastic discontinuity in power. By dividing out the digital gains and the derived linear PFB response, we would ideally remove these effects. After applying digital corrections (lower panel), residual PFB artifacts appear as peaks in power at the coarse band edges. These are more apparent at the upper third of the band, indicating a nonlinear dependence of these artifacts on the raw autocorrelation power. The digital gain removal leaves a residual gain ‘jump’ before the upper third of the frequency band. Due to their nonlinear source, these artifacts are not captured and removed by calibration, impeding the precision that we require in calibration for our science.

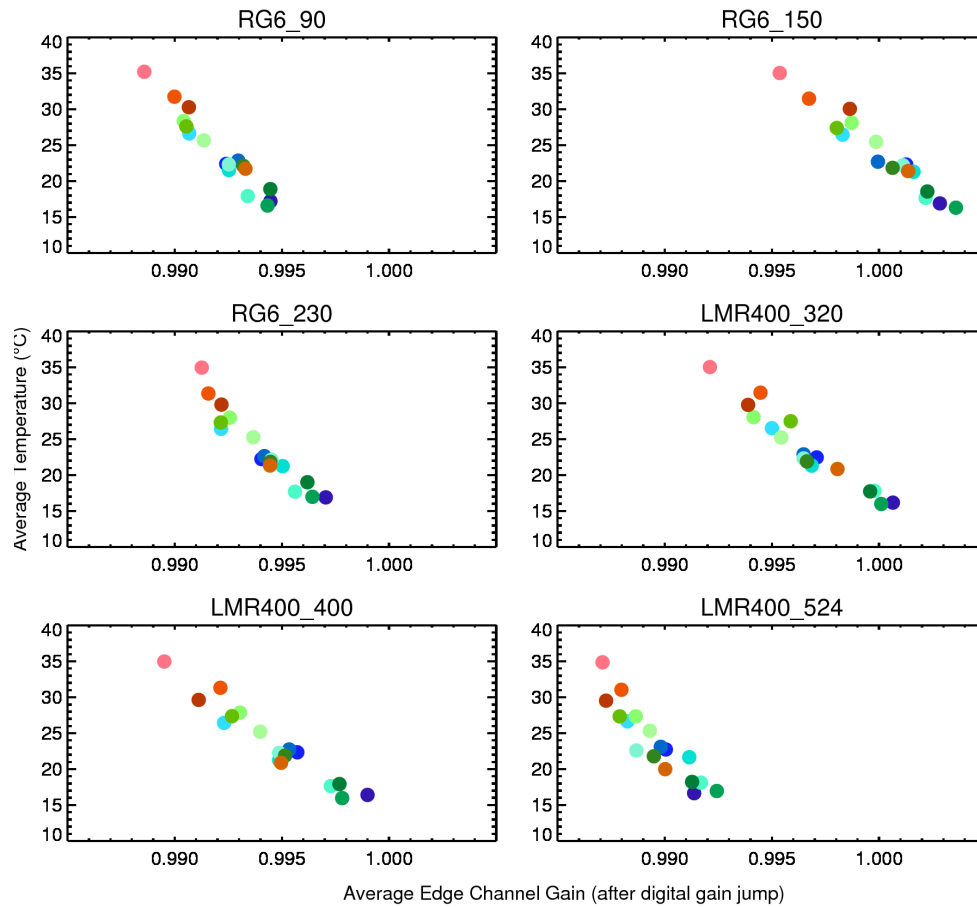


Figure 3.9: We show a plot by N. Barry of the residual PFB response artifacts at the coarse band edges in calibration solutions derived from crosscorrelations. Barry computed the mean of these artifacts for several nights of 2013 MWA data (each night is represented by a different color) and plotted the means as a function of beamformer temperature. Each of the six panels corresponds to a type and length of cable between the tile and digital receiver. The clear interaction between the temperature-dependent analog gain and the PFB response indicates that the artifacts have a nonlinear source. Image credit: N. Barry, personal communication, March 14, 2017

3.6.3 *Nonlinear effects in 2D power spectra*

The effects of nonlinear antenna response are also visible in the 2D power spectrum. After a Fourier transform, regularly spaced structures in frequency result in copies of the foreground power at the specific spectral modes associated with the frequency spacing. Corrupted data appearing at the same location in each coarse band, such as artifacts of the PFB response at the coarse band edges, results in peaks of power in the EoR Window. In Figure 3.10 we show a 2D power spectrum from a subset of the 2013 MWA data processed by Barry et al. (2019b). In the EoR window ‘coarse band lines’ from the residual PFB artifacts are apparent. These coarse band lines have historically been an unfortunate feature of MWA data and reducing their impact is a goal of this work.

To diminish the coarse band lines, we must remove the nonlinear digital quantization response from the data. Our science requires that we capture and remove this spectral response with a precision of a part in 10^{-5} . We describe the process of removing nonlinear effects in the next two chapters, first using mathematical methods to correct the 4-bit requantization stage before correlation in Chapter 4. Then, we model the digital signal pathway in order to develop a correction for the asymmetric 8-bit rounding that occurs between the filter and FFT stages of the fine polyphase filterbank (Chapter 5).

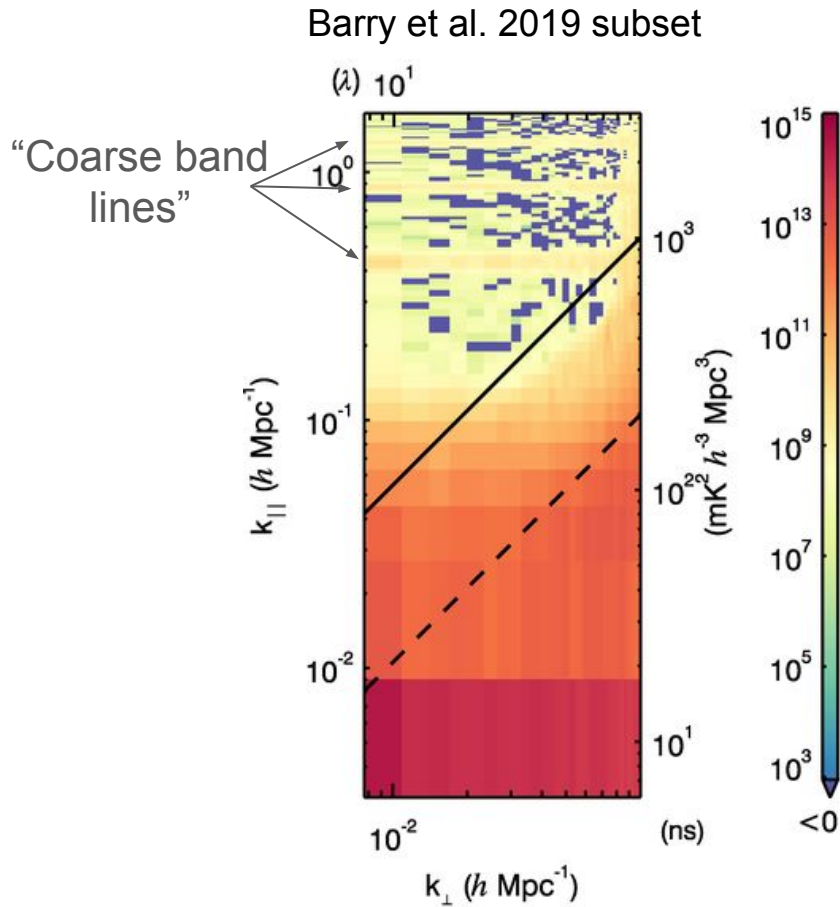


Figure 3.10: We show a 2D power spectrum from a subset of the 2013 MWA data analyzed by Barry et al. (2019b). In the EoR Window (above the black line, as described in Figure 2.1) we see horizontal lines of relatively high power. These peaks, referred to as ‘coarse band lines’, are the result of corruptions in the data which occur at the same location in each coarse band. For example, residual artifacts at the coarse band edges after the removal of the polyphase filterbank response will place copies of the foreground power at each of the coarse band lines. To eliminate these lines we must remove the nonlinear digital quantization response from our data.

Chapter 4

CORRECTING QUANTIZATION NONLINEARITIES

Digital quantization changes the statistics of the measured signal in a nonlinear way. The quantization effects depend on the number of allotted bits such that nonlinear effects become more extreme as fewer bits are used in the quantization stage. The lowest bit-depth quantization in the MWA signal path is the 4-bit requantization before correlation. This makes it likely the primary source of the nonlinear effects we wish to remove from our data, and thus is the first requantization stage we wish to correct.

Van Vleck and Middleton (1966) derived mathematical relations between the statistics of an analog signal before quantization and the resultant quantized signal. These equations are frequently referred to as ‘Van Vleck corrections’. Benkevitch et al. (2016) used the Van Vleck corrections to write down equations for the specific case of the MWA’s 4-bit requantization stage. We follow that derivation to show the change in statistics of an analog signal before and after a 4-bit quantization. Then we use these equations to implement corrections which remove the effects of the 4-bit requantization stage from MWA data, and we examine the impacts of this correction on our data quality.

As we will motivate below, we cannot use the Van Vleck corrections for quantization stages deeper in the MWA signal path. Instead, we must model the digital system to derive the relationships between the statistics of the pre- and post-quantization signals. In Chapter 5 we will use this process to develop a correction for the asymmetric 8-bit requantization stage within the correlator fine polyphase filterbank.

4.1 Quantization effects on signal statistics

We assume that before quantization the analog signal Z_j from the j th antenna is a complex circular Gaussian random variable. That is, the real and imaginary parts of Z_j are each Gaussian distributed with zero mean and the same variance, and are uncorrelated with each other. An implication of this is that there is no preferred phase in Z_j . Thus before quantization we have

$$Z_j = X_j + iY_j \quad (4.1)$$

where X_j and Y_j are drawn from a zero mean Gaussian distribution described by standard deviation σ_j where

$$\sigma_j^2 = E[X_j^2] = E[Y_j^2]. \quad (4.2)$$

We are interested in the effects of quantization on the standard deviation σ_j and on the covariance $E[X_j X_k]$ between analog signals measured by the j th and k th antennas. Since the statistics of X_j and Y_j are the same, it is sufficient to address only X_j below.

4.1.1 Quantization impacts on standard deviations

For the case of the 4-bit requantization stage in the MWA correlator (Section 3.4), the measured signal is binned as integers ranging from -7 to 7, inclusively. These bins are centered on the integers such that numbers between 0.5 and 1.5 are binned as 1, numbers between 1.5 and 2.5 are binned as 2, and so on. The probability that the quantized value \hat{X}_j has a particular integer value is equal to the probability of X_j falling into the corresponding quantization bin. Thus the expectation value of \hat{X}_j^2 can be written

$$\begin{aligned} E[\hat{X}_j^2] = & 0^2 \cdot P(-0.5 < X_j < 0.5) + (-1)^2 \cdot P(-1.5 < X_j < -0.5) \\ & + 1^2 \cdot P(0.5 < X_j < 1.5) + \cdots + (-7)^2 \cdot P(X_j < -6.5) + 7^2 \cdot P(X_j > 6.5). \end{aligned} \quad (4.3)$$

Taking advantage of the symmetries of both the distribution of X_j and the quantization bins,

$$\begin{aligned} E[\hat{X}_j^2] = & 0^2 \cdot P(-0.5 < X_j < 0.5) + 1^2 \cdot [P(-1.5 < X_j < 1.5) - P(-0.5 < X_j < 0.5)] \\ & + \cdots + 7^2 \cdot [1 - P(-6.5 < X_j < 6.5)]. \end{aligned} \quad (4.4)$$

With X_j drawn from a zero-mean Gaussian distribution with σ_j , the probability that $X_j \in [-a, a]$

$$P(X_j \in [-a, a]) = \operatorname{erf}\left(\frac{a}{\sigma_j\sqrt{2}}\right). \quad (4.5)$$

Substituting this probability and combining terms gives the compact expression

$$\begin{aligned} E[\hat{X}_j^2] &= 0^2 \cdot \operatorname{erf}\left(\frac{0.5}{\sigma_j\sqrt{2}}\right) + 1^2 \cdot \left[\operatorname{erf}\left(\frac{1.5}{\sigma_j\sqrt{2}}\right) - \operatorname{erf}\left(\frac{0.5}{\sigma_j\sqrt{2}}\right)\right] \\ &\quad + \dots + 7^2 \cdot \left[1 - \operatorname{erf}\left(\frac{6.5}{\sigma_j\sqrt{2}}\right)\right] \\ &= (-1)\operatorname{erf}\left(\frac{0.5}{\sigma_j\sqrt{2}}\right) + (-3)\operatorname{erf}\left(\frac{1.5}{\sigma_j\sqrt{2}}\right) + \dots + (-13)\operatorname{erf}\left(\frac{6.5}{\sigma_j\sqrt{2}}\right) + 7^2 \\ &= 7^2 - \sum_{m=0}^6 (2m+1)\operatorname{erf}\left(\frac{m+0.5}{\sigma_j\sqrt{2}}\right). \end{aligned} \quad (4.6)$$

Taking the square root of both sides gives the Van Vleck correction for standard deviations

$$\sqrt{\hat{\sigma}_j} = \left[7^2 - \sum_{m=0}^6 (2m+1)\operatorname{erf}\left(\frac{m+0.5}{\sqrt{\sigma_j}\sqrt{2}}\right)\right]^{1/2}. \quad (4.7)$$

We use Equation (4.7) to plot the standard deviation after the 4-bit quantization $\hat{\sigma}_j$ as a function of σ_j , the standard deviation before quantization in the left panel of Figure 4.1. For comparison, we plot a line with a slope of one in grey. In the right panel we show the residual $\hat{\sigma}_j - \sigma_j$. The nonlinear behavior of Equation (4.7) is evident in both plots.

4.1.2 Quantization impacts on covariance

We now consider the covariance between the real and/or imaginary parts of Z_j and Z_k , the analog signals measured by the j th and k th antennas, respectively. We consider here only the covariance $E[X_j X_k]$, as $E[Y_j X_k]$, $E[X_j Y_k]$, and $E[Y_j Y_k]$ have the same behavior. To relate the covariances $E[X_j X_k]$ and $E[\hat{X}_j \hat{X}_k]$ of the signals before and after quantization, respectively, we follow Benkevitch et al. (2016) and use Price's theorem. This states that for two random variables X and Y ,

$$\frac{\partial \langle f(X, Y) \rangle}{\partial \langle XY \rangle} = \left\langle \frac{\partial f}{\partial X} \frac{\partial f}{\partial Y} \right\rangle. \quad (4.8)$$




Figure 4.1: Using Equation (4.22), we plot the standard deviation $\hat{\sigma}_j$ of a signal after quantization as a function of the standard deviation σ_j of the analog signal measured by the j th antenna (left panel). In grey we plot a line with a slope of one for comparison. We also show the residual $\hat{\sigma}_j - \sigma_j$ in the right panel. The relationship between $\hat{\sigma}_j$ and σ_j is clearly nonlinear.

where brackets indicate expectation values. Using the function $f(X_j, X_k) = \hat{X}_j \hat{X}_k$, we obtain

$$\frac{\partial E[\hat{X}_j \hat{X}_k]}{\partial E[X_j X_k]} = \left\langle \frac{\partial \hat{X}_j}{\partial X_j} \frac{\partial \hat{X}_k}{\partial X_k} \right\rangle. \quad (4.9)$$

For our quantization scheme, the derivative is composed of delta functions

$$\frac{\partial \hat{X}}{\partial X} = \delta(X - (-6.5)) + \delta(X - (-5.5)) + \dots + \delta(X - 5.5) + \delta(X - 6.5), \quad (4.10)$$

which gives

$$\frac{\partial \hat{X}_j}{\partial X_j} \frac{\partial \hat{X}_k}{\partial X_k} = \sum_{m=-7}^6 \sum_{n=-7}^6 \delta(X_j - (m + 0.5)) \delta(X_k - (n + 0.5)). \quad (4.11)$$

To find the expectation value, we integrate with the joint normal probability density function

$$\begin{aligned} \left\langle \frac{\partial \hat{X}_j}{\partial X_j} \frac{\partial \hat{X}_k}{\partial X_k} \right\rangle &= \sum_{m=-7}^6 \sum_{n=-7}^6 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} dX_j dX_k \delta(X_j - (m + 0.5)) \delta(X_k - (n + 0.5)) \\ &\quad \times \frac{1}{2\pi\sigma_j\sigma_k\sqrt{1-\rho^2}} \exp \left[-\frac{1}{2(1-\rho^2)} \left(\frac{X_j^2}{\sigma_j^2} + \frac{X_k^2}{\sigma_k^2} - \frac{2\rho X_j X_k}{\sigma_j\sigma_k} \right) \right], \end{aligned} \quad (4.12)$$

where the correlation coefficient $\rho = E[X_j X_k]/\sigma_j\sigma_k$. Integration gives

$$\begin{aligned} \left\langle \frac{\partial \hat{X}_j}{\partial X_j} \frac{\partial \hat{X}_k}{\partial X_k} \right\rangle &= \sum_{m=-7}^6 \sum_{n=-7}^6 \frac{1}{2\pi\sigma_j\sigma_k\sqrt{1-\rho^2}} \exp \left[-\frac{1}{2(1-\rho^2)} \left(\frac{(m+0.5)^2}{\sigma_j^2} + \frac{(n+0.5)^2}{\sigma_k^2} \right. \right. \\ &\quad \left. \left. - \frac{2\rho(m+0.5)(n+0.5)}{\sigma_j\sigma_k} \right) \right]. \end{aligned} \quad (4.13)$$

Here we can use $\left\langle \frac{\partial \hat{X}_j}{\partial X_j} \frac{\partial \hat{X}_k}{\partial X_k} \right\rangle = \frac{\partial E[\hat{X}_j \hat{X}_k]}{\partial E[X_j X_k]} = \frac{\partial E[\hat{X}_j \hat{X}_k]}{\sigma_j\sigma_k\partial\rho}$ to write

$$\begin{aligned} \frac{\partial E[\hat{X}_j \hat{X}_k]}{\partial\rho} &= \sum_{m=-7}^6 \sum_{n=-7}^6 \frac{1}{2\pi\sqrt{1-\rho^2}} \exp \left[-\frac{1}{2(1-\rho^2)} \left(\frac{(m+0.5)^2}{\sigma_j^2} + \frac{(n+0.5)^2}{\sigma_k^2} \right. \right. \\ &\quad \left. \left. - \frac{2\rho(m+0.5)(n+0.5)}{\sigma_j\sigma_k} \right) \right]. \end{aligned} \quad (4.14)$$

An integration over both sides gives the Van Vleck correction for covariances:

$$\begin{aligned} E[\hat{X}_j \hat{X}_k] &= \sum_{m=-7}^6 \sum_{n=-7}^6 \int_0^\rho d\rho' \frac{1}{2\pi\sqrt{1-\rho'^2}} \exp \left[-\frac{1}{2(1-\rho'^2)} \left(\frac{(m+0.5)^2}{\sigma_j^2} \right. \right. \\ &\quad \left. \left. + \frac{(n+0.5)^2}{\sigma_k^2} - \frac{2\rho'(n+0.5)(m+0.5)}{\sigma_j\sigma_k} \right) \right]. \end{aligned} \quad (4.15)$$

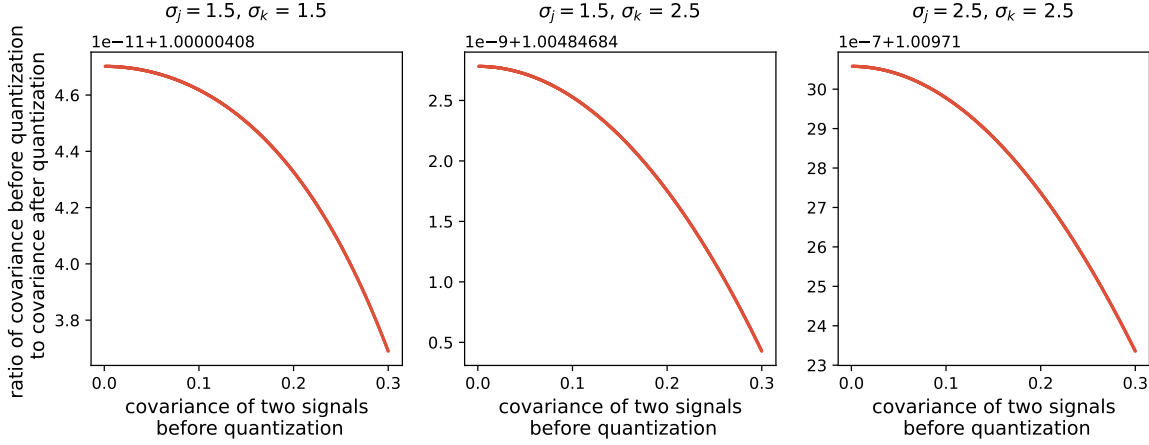


Figure 4.2: We use Equation (4.15) to plot the ratio of the covariances before and after quantization $E[X_j X_k]/E[\hat{X}_j \hat{X}_k]$ as a function of the covariance $E[X_j X_k]$. While the covariance could theoretically take on greater values, when we inspect our data we see $E[X_j X_k] \lesssim 0.3$ and so use that range here. This relationship depends on the standard deviations of the measured signal before quantization, σ_j and σ_k . We show the ratio for three pairs of standard deviations. The nonlinearity of the relationship between $E[X_j X_k]$ and $E[\hat{X}_j \hat{X}_k]$ is difficult to detect by eye.

With $\rho = E[X_j X_k]/\sigma_j \sigma_k$ we could use Equation (4.32) to plot $E[\hat{X}_j \hat{X}_k]$ as a function of $E[X_j X_k]$, but the nonlinearity is so small it is difficult to see by eye. Instead, we plot the ratio $E[X_j X_k]/E[\hat{X}_j \hat{X}_k]$ as a function of $E[X_j X_k]$ in Figure 4.2. Since the relationship between the covariances before and after quantization depends on the standard deviations of the measured signals, we plot the ratio for three pairs of standard deviations: $[1.5, 1.5]$, $[1.5, 2.5]$, and $[2.5, 2.5]$. While $E[X_j X_k]$ and $E[\hat{X}_j \hat{X}_k]$ are clearly not equal to each other, the nonlinearity is difficult to detect by eye.

By approximating $\hat{\sigma}_j$ and $E[\hat{X}_j \hat{X}_k]$ with our visibility data, we can use Equations (4.7) and (4.15) as corrections for the MWA 4-bit requantization stage. However, it is important to note that the 4-bit requantization stage is not quantizing an analog signal but rather the

digital outputs of the MWA correlator polyphase filter bank (PFB). In order to use the Van Vleck corrections on this data, we are explicitly assuming that we can treat the PFB output data as a complex circular Gaussian random variable.

4.1.3 Visibilities as approximations of signal statistics

In order to use the Van Vleck corrections we must relate digital visibilities to the statistics of the measured signal. As described below, we can use autocorrelations to approximate signal standard deviations and crosscorrelations to approximate covariances between signals measured by two antennas. These approximations allow us to use the Van Vleck corrections to remove the effects of the 4-bit requantization stage.

Autocorrelations

For the analog signal Z_j , the autocorrelation V_{jj} is calculated by summing N samples of $Z_j Z_j^*$. That is,

$$\begin{aligned} V_{jj} &= \sum_{n=1}^N Z_{j_n} Z_{j_n}^* \\ &= \sum_{n=1}^N (X_{j_n} + iY_{j_n})(X_{j_n} - iY_{j_n}) \\ &= \sum_{n=1}^N X_{j_n}^2 + \sum_{n=1}^N Y_{j_n}^2. \end{aligned} \tag{4.16}$$

For our typical EoR data resolutions of 0.5 seconds or 2 seconds, $N = 20,000$ or $80,000$, respectively. In the limit of large N , the above sums approximate the second moments of X_j and Y_j with

$$\frac{1}{N} \sum_{n=1}^N X_{j_n}^2 \rightarrow E[X_j^2], \tag{4.17}$$

$$\frac{1}{N} \sum_{n=1}^N Y_{j_n}^2 \rightarrow E[Y_j^2]. \tag{4.18}$$

Since $E[X_j^2] = E[Y_j^2] = \sigma_j$, we can write the autocorrelation in the large N limit as

$$\begin{aligned} \frac{V_{11}}{N} &\simeq E[X_j^2] + E[Y_j^2] \\ &= 2E[X_j^2] \\ &= 2\sigma_j^2 \end{aligned} \quad (4.19)$$

For use in the quantization correction, it is convenient to rewrite the above expression in the form

$$\sqrt{V_{jj}/2N} \simeq \sigma_j. \quad (4.20)$$

Using a similar process, we can write our autocorrelations of the quantized signal \hat{Z}_j as

$$\sqrt{\hat{V}_{jj}/2N} \simeq E[\hat{X}_j^2]. \quad (4.21)$$

Substituting Equations (4.20) and (4.21) into Equation (4.7) gives

$$\boxed{\sqrt{\hat{V}_{jj}/2N} \simeq \left[7^2 - \sum_{m=0}^6 (2m+1) \operatorname{erf} \left(\frac{m+0.5}{\sqrt{V_{jj}/2N}\sqrt{2}} \right) \right]^{1/2}}. \quad (4.22)$$

With Equation (4.22) we use our measured autocorrelations \hat{V}_{jj} after the 4-bit requantization stage to obtain estimates of the true autocorrelations V_{jj} had the 4-bit requantization not occurred. A similar process is used for crosscorrelations.

Crosscorrelations

The crosscorrelation V_{jk} of signals from the j th and k th antennas is calculated by summing N samples of $Z_j Z_k^*$ to obtain

$$\begin{aligned} V_{jk} &= \sum_{n=1}^N Z_{j_n} Z_{k_n}^* \\ &= \sum_{n=1}^N (X_{j_n} + iY_{j_n})(X_{k_n} - iY_{k_n}) \\ &= \sum_{j=1}^N (X_{j_n} X_{k_n} + Y_{j_n} Y_{k_n} + iY_{j_n} X_{k_n} - iX_{j_n} Y_{k_n}). \end{aligned} \quad (4.23)$$

The limit of large N can again be taken to give, for any combination of X_j, Y_j with X_k, Y_k :

$$\frac{1}{N} \sum_{n=1}^N X_{j_n} X_{k_n} \rightarrow E[X_j X_k]. \quad (4.24)$$

Thus in this limit, the crosscorrelation becomes

$$\begin{aligned} \frac{V_{jk}}{N} &\simeq E[X_j X_k] + E[Y_j Y_k] + iE[Y_j X_k] - iE[X_j Y_k] \\ &= 2E[X_j X_k] + 2iE[Y_j X_k], \end{aligned} \quad (4.25)$$

where the last step has used the following relations for circular complex random variables:

$$E[X_j X_k] = E[Y_j Y_k] \quad (4.26)$$

$$E[Y_j X_k] = -E[X_j Y_k] \quad (4.27)$$

Rewriting in terms of the real and imaginary parts of V_{jk} gives, in the limit of large N ,

$$\Re V_{jk}/2N \simeq E[X_j X_k], \quad (4.28)$$

$$\Im V_{jk}/2N \simeq E[Y_j X_k]. \quad (4.29)$$

For the case of quantized signals \hat{X}_j and \hat{X}_k , we can write the relations

$$\Re \hat{V}_{jk}/2N \simeq E[\hat{X}_j \hat{X}_k], \quad (4.30)$$

$$\Im \hat{V}_{jk}/2N \simeq E[\hat{Y}_j \hat{X}_k]. \quad (4.31)$$

Substituting Equations (4.28) and (4.30) into Equation (4.15) gives

$$\boxed{\Re \hat{V}_{jk}/2N \simeq \sum_{m=-7}^6 \sum_{n=-7}^6 \int_0^\rho d\rho' \frac{1}{2\pi\sqrt{1-\rho'^2}} \exp \left[-\frac{1}{2(1-\rho'^2)} \left(\frac{(m+0.5)^2}{\sigma_j^2} + \frac{(n+0.5)^2}{\sigma_k^2} - \frac{2\rho'(n+0.5)(m+0.5)}{\sigma_j\sigma_k} \right) \right]} \quad (4.32)$$

where the correlation coefficient $\rho \simeq \Re V_{jk}/(2N\sigma_j\sigma_k)$.

We can use the above equation to estimate the real(imaginary) part of $V_{jk}/(2N\sigma_j\sigma_k)$ from the real(imaginary) part of our requantized visibility \hat{V}_{jk} . However, this estimate requires knowledge of σ_j and σ_k . We first use Equation (4.22) to correct the autocorrelations and then take advantage of Equation (4.20) to estimate σ_j and σ_k . With these estimates we use Equation (4.32) to correct the crosscorrelations. Our implementation of these corrections is described below.

4.2 Implementation of quantization corrections

I implemented the corrections for the 4-bit requantization stage in `pyuvdata` (Hazelton et al., 2017), an open-source python package which is an integral component of our preprocessing pipeline described in Section 8.2. Because the 4-bit requantization happens immediately before correlation, it must be corrected before the linear digital gains and PFB response are removed and before calibration is performed. I needed access to the raw data from the MWA correlator in order to apply the Van Vleck corrections, so I first built functionality in `pyuvdata` to ingest the correlator data files. I then added capability to the file read function to apply the Van Vleck corrections, as well as other corrections described in Section 8.2.

4.2.1 MWA correlator *gpubox* files

The MWA correlator outputs raw data for each observation in a bespoke FITS file format. Each coarse frequency band is processed by an individual correlator server (known as a ‘*gpubox*’) and the resulting visibilities written out into a coarse band-specific file. If the number of time samples is sufficiently high, two *gpubox* data files will be written for each coarse band. An accompanying metadata (‘*metafits*’) file is written for the observation, containing details about the instrument, antenna tiles, and observation criteria. Later pipeline stages (such as FHD/ ϵ psilon (Barry et al., 2019a; Sullivan et al., 2012), used by many United States collaborators and Hyperdrive¹/Chips (Trott et al., 2016), used by many Australian

¹https://github.com/MWATelescope/mwa_hyperdrive

collaborators) typically require the metafits file and gpubox files for an observation to be consolidated by preprocessing into either a UVFITS file (Greisen, Accessed June 5 2024) or a CASA measurement set (Bean et al., 2022).

Historically, the process of converting correlator gpubox files to a standardized file format was performed by the COTTER² software, which was used behind the scenes to perform file conversion. While users could download raw data, there were no accessible options for interacting with that raw data directly. In order to create a more flexible user experience, I took advantage of the open-source software package `pyuvdata`.

In addition to allowing easy conversion between file types, `pyuvdata` creates an in-memory object for the astronomical observation. This object has attributes which hold the observation metadata as well as, optionally, arrays of data, flags, and sample numbers. Through both `pyuvdata` built-in methods as well as other python code, this object can easily be used to interrogate and manipulate the data. I added capability for `pyuvdata` to ingest MWA correlator gpubox files, providing a user-friendly interface for interacting with the raw data. COTTER’s behind-the-scenes processing has since been superceded by Birli³. We compared `pyuvdata` outputs with uvfits files produced by COTTER and later Birli to validate the `pyuvdata` processing implementation.

4.2.2 Van Vleck correction implementation

After building functionality for ingesting the correlator gpubox files in `pyuvdata`, I added options for applying Van Vleck corrections to the raw data. These corrections first scale the visibilities by $1/2N$. For the case of data taken at 40 kHz and 0.5 second resolution, $N = 20,000$. We then take the square root of the autocorrelations to obtain $\sqrt{\hat{V}_{jj}/2N}$.

²<https://github.com/MWATelescope/cotter>

³<https://github.com/MWATelescope/Birli>

Equation (4.22) is not invertible, so we use Newton's method to iteratively solve the equation

$$0 = \sqrt{\hat{V}_{jj}/2N} - \left[7^2 - \sum_{m=0}^6 (2m+1) \operatorname{erf} \left(\frac{m+0.5}{\sqrt{V_{jj}/2N} \sqrt{2}} \right) \right]^{1/2}. \quad (4.33)$$

This process returns an estimate of $\sqrt{V_{jj}/2N}$ which we square and rescale to obtain our corrected autocorrelation V_{jj} .

For each crosscorrelation visibility the correction to obtain an estimate of the real (imaginary) part of $V_{jk}/2N$ depends on three values: the scaled real (imaginary) part of the uncorrected visibility $\hat{V}_{jk}/2N$ as well as estimates of σ_j and σ_k . We use the corrected autocorrelations and Equation (4.21) to estimate σ_j and σ_k . We can then use Equation (4.32) to solve for the correlation coefficient ρ which we multiply by our estimated σ_j and σ_k to obtain the estimate of $\Re \hat{V}_{jk}/2N$.

A single 2 minute MWA observation taken at 0.5 second time resolution and 40 kHz frequency resolution has $\sim 10^{10}$ crosscorrelations which each require a unique correction. Initially, we used Newton's method to solve Equation (4.32). However, this algorithm had a high computational cost due to the 14^2 exponential terms, which were each evaluated ~ 10 times to compute the integral. By combining redundant terms, this number was reduced by a factor of 2. Though only ~ 5 evaluations of Equation (4.32) were required for the root-finding function to converge, correcting data for one observation took ~ 5 hours. We used cython to implement an approximation of Equation (4.32) using Chebyshev polynomials, which brought correction time for a single observation to ~ 30 minutes.

To develop a Chebyshev polynomial approximation, Equation (4.32) was fit to the following function

$$\rho = c_1 \cdot T_1(\kappa) + c_3 \cdot T_3(\kappa) + c_5 \cdot T_5(\kappa), \quad (4.34)$$

where T_k is the k th Chebyshev polynomial and κ is either the real or imaginary part of $\hat{V}_{jk}/2N$. A two-dimensional grid of (σ_j, σ_k) was created with values ranging between 0.9-4.5




Figure 4.3: We take the difference between corrected values obtained by solving Equation (4.32) and those obtained using a Chebyshev approximation of the equation. Values are computed for three pairs of σ_j, σ_k . The largest difference is of order 10^{-12} , which is sufficient considering we need to remove the antenna response with a precision of a part in 10^{-5}

and grid spacing of 0.01. For each pair (σ_j, σ_k) , a fitting function solved for the coefficients c_1, c_3, c_5 . These coefficients were saved and are imported into the Van Vleck correction function in `pyuvdata`. For each crosscorrelation in the data, a bisection search finds the indices for the nearest σ_j and σ_k in the grid for the real and imaginary parts, respectively. A bilinear interpolation algorithm uses these indices to obtain the coefficients c_1, c_3 , and c_5 . These coefficients are used in equation Equation (4.34) to obtain ρ , which we multiply by σ_j and σ_k to get an estimate of the real (or imaginary) part of $V_{jk}/2N$. The estimates are rescaled to give the corrected visibilities V_{jk} .

In Figure 4.3 we show agreement between Equation (4.32) and the Chebyshev approximation, plotting the difference between values corrected with each method for three different σ_j, σ_k pairs. In each case, the difference is small, with the largest difference of order 10^{-12} . This is sufficient for our precision requirements of capturing the antenna response to a part in 10^{-5} .

4.3 Results with Van Vleck corrections

As an initial evaluation of the performance of the Van Vleck correction for the 4-bit MWA requantization stage we examine autocorrelations. In Figure 4.4 we plot the XX instrumental polarization of an autocorrelation from antenna Tile028 which has been averaged over a single observation from 2013 MWA data. In the upper panel we have applied corrections for the digital gains and the polyphase filterbank (PFB) response as discussed in Figure 3.8. These linear instrument responses should divide out cleanly, but instead we see residual artifacts. At the coarse band edges, demarcated by vertical dashed grey lines, we see peaks in power after dividing out the PFB response. Also, dividing out the digital gains leaves a residual ‘jump’ which is highlighted by a red oval. In the lower panel we have applied the 4-bit Van Vleck correction before dividing out the gains and PFB response, and here we see a marked improvement in these corrections. The power has dropped at the coarse band edges and the digital gain jump has shrunk. This is very encouraging, and shows that we can mitigate nonlinearities in our data by correcting the effects of quantization. However, there are clearly still residual artifacts which suggests that correcting quantization stages deeper in the digital system could further improve our data quality.

We note that while we do not use autocorrelations in our power spectrum analysis, the calibration methods used in recent MWA analyses (Barry et al., 2019b; Li et al., 2019; Wilensky et al., 2023) as well as in this work (Section 8.3.1) incorporate autocorrelations to best capture the antenna spectral structure. If the Van Vleck corrections are not used, not only will crosscorrelations carry a nonlinear response, but also the autocorrelations will propagate their different nonlinear artifacts into the calibration solutions.

We applied the corrections for the 4-bit quantization stage to a subset of the 2013 MWA data used for the power spectrum limit analysis by Barry et al. (2019b). We compute the 1D power spectrum with our corrected data and show a comparison with Barry’s data in Figure 4.5. There is an overall drop in power between the data processed by Barry (black

line) and our data (blue line), which is very encouraging. In particular, we see a drop in the coarse band line peak power (located at $\sim k = 0.41, 0.9, \text{ and } 1.4$), evidence of the reduction in nonlinear artifacts at the coarse band edges. With these promising results, we decided to use the Van Vleck corrections in a power spectrum limit analysis.

For Wilensky’s analysis of 2014 MWA data (Wilensky et al., 2023), we applied the 4-bit Van Vleck correction to the entire data set. To evaluate the impacts on the 2D power spectrum coarse band lines, we compare with a set of 2013 MWA data from Barry et al. (2019b) in Figure 4.6. On the left is a 2D power spectrum from Barry’s analysis of 2013 data and on the right we show the 2D power spectrum of Van Vleck corrected 2014 data from Wilensky’s work. Detailed comparisons between these two power spectra are difficult because they contain different numbers of observations, making their noise levels different. The coarse band lines aren’t impacted by the noise levels, however, and we see a clear drop in the coarse band line power in the 2014 Van Vleck corrected power spectrum.

The results with the 4-bit Van Vleck correction show clear improvement to our data quality. However, as we saw in Figure 4.4, there are still residual nonlinear artifacts present in the data. To remove this contamination, we turn our attention to quantization stages deeper in the MWA digital system.




Figure 4.4: We show an autocorrelation averaged over a single zenith observation from 2013 MWA data. The XX instrumental polarization of antenna Tile028 is plotted. In the upper panel we have applied corrections for the digital gains and polyphase filterbank (PFB) response. We are left with clear residual artifacts. In particular, at the edges of the coarse bands—which are demarcated by dashed grey lines—we see peaks after dividing out the PFB response. Also, we see a digital gain ‘jump’ (as discussed in Chapter 3) which here is highlighted by the red oval. These features point to nonlinearities in the data. In the lower panel we have applied our implementation of the 4-bit Van Vleck correction. The peaks at the coarse band edges have dropped in power and the gain jump has shrunk. These exciting improvements confirm that we can reduce or eliminate nonlinearities in our data by correcting the effects of digital quantization. However, it is clear there are still lingering nonlinear effects in the data from other requantization stages deeper in the digital system.

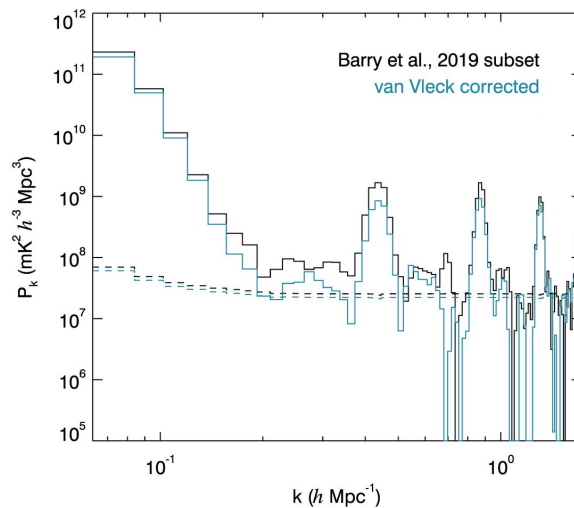


Figure 4.5: We applied the Van Vleck correction to a subset of 2013 MWA data analyzed by Barry et al. (2019b) and computed the 1D power spectrum. Here we plot the 1D spectrum of the Van Vleck corrected data (blue line) along with the 1D spectrum of the Barry data subset (black line). The respective 1σ thermal noise is plotted with the blue and black dashed lines. We see an overall drop in power when the Van Vleck correction is applied. Particularly, the coarse band line power (seen in the peaks at $\sim k = 0.41, 0.9,$ and 1.4) has dropped with the Van Vleck correction, which demonstrates the reduction of residual artifacts at the coarse band edges.

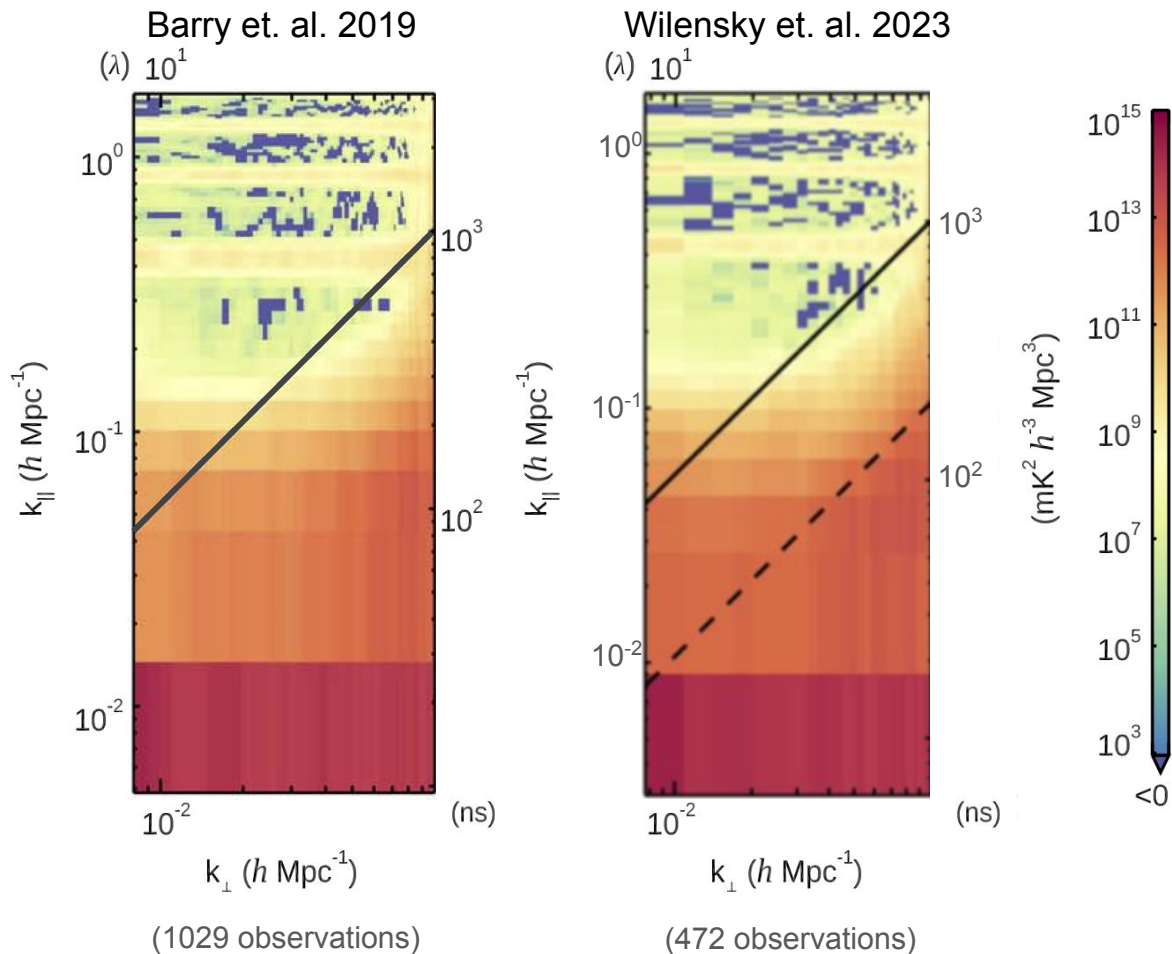


Figure 4.6: To evaluate the effects of the Van Vleck correction on the coarse band line power, we plot 2D power spectra. On the left we show a power spectrum from Barry et al. (2019b), computed using 2013 MWA data. We compare this with the 2014 analysis by Wilensky et al. (2023) (right) which used 2014 Van Vleck corrected data. These power spectra include a different number of observations and thus have different noise levels. While an overall comparison is not meaningful, the coarse band lines are not impacted by the noise levels. There is a clear drop in coarse band line power with the Van Vleck correction. This is very promising, and suggests that a reanalysis of Barry’s 2013 limit set using Van Vleck corrections could improve on that result.

Chapter 5

CORRECTING QUANTIZATION OF REAL-WORLD SYSTEMS

After developing and applying a Van Vleck correction for the 4-bit requantization stage in the correlator, it is clear in Figure 4.4 that some lingering nonlinear effects are still present in the data. While I was working on the 4-bit correction, a paper associated with pulsar observations was published which identified and described an asymmetric quantization stage in the correlator fine polyphase filterbank (PFB) implementation (McSweeney et al., 2020). There is an 8-bit requantization stage between the PFB subfilters and the FFT taps. Higher bit-depth quantization stages have smaller nonlinear effects, and normally the effects of an 8-bit requantization stage would be small enough that we would instead concentrate our efforts on first correcting the smaller 5-bit requantization stage. In the correlator implementation, however, negative numbers are rounded using a different scheme than positive numbers. The impacts of this scheme are nontrivial, and we would like to correct them.

We cannot use the Van Vleck correction for the 8-bit asymmetric requantization because that correction assumes correlation happens immediately after quantization. In this case, the requantization is followed by an FFT which channelizes the data. In order to develop a correction we model the digital signal path. To account for shot noise in the simulated values, we fit functions to describe the change in statistics, and thus to visibilities, from the requantization stage.

After creating a model and calculating initial fits to that model, this project was interrupted by the arrival of an exciting new set of data from MWA Phase III. We present our work correcting the effects of the asymmetric 8-bit requantization stage in hopes that it

benefits future students continuing this work.

5.1 Asymmetric quantization in the fine PFB

The 8-bit requantization scheme within the correlator FPGA fine polyphase filterbank was implemented incorrectly, resulting in different binning schemes for positive and negative numbers. This is illustrated in Figure 5.1 from McSweeney. A typical symmetric binning scheme is shown in blue while the PFB asymmetric bins are shown in orange. In the PFB scheme, the bin edges for negative numbers fall on integer boundaries. Numbers between -1 and 0 are binned as -1, numbers between -1 and -2 are binned as -2, and so on. Bins for positive numbers are centered between integers, so numbers between 0 and 0.5 are binned as zero, numbers between 0.5 and 1.5 are binned as 1, etc. This asymmetry results in a 0 bin which is not centered on zero but rather extends from 0 to 0.5 and is half the width of all other bins.

The asymmetric rounding scheme changes the signal standard deviation and covariance in a nonlinear way, and, importantly, results in a distribution with a nonzero mean after quantization. This distribution is then channelized in the fine PFB FFT, and the nonzero mean becomes excess power in the DC bin after the FFT. The DC power, or ‘spike’ in the fine channel at the center of each coarse channel was a known data quality problem, and McSweeney et al. (2020) identified its source. This periodic excess power not only corrupts each DC fine frequency channel but also contributes to coarse band line contamination in the EoR window. To avoid this we have historically excluded data in that frequency channel from MWA data analyses (Beardsley et al., 2016; Barry et al., 2019b; Li et al., 2019).

5.2 A correction for the PFB asymmetric quantization

The 4-bit requantization was relatively straightforward to correct because correlation occurred immediately after quantization. This meant that the correlator outputs were a good approximation of the post-quantization statistics and we could use the Van Vleck formal-

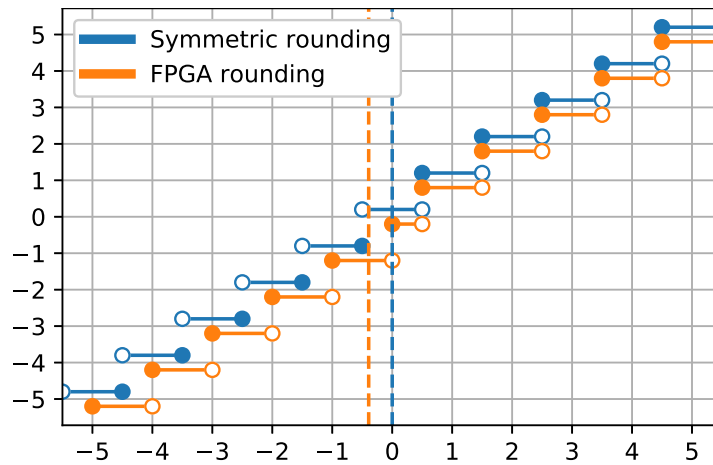


Figure 5.1: This plot, shown by McSweeney et al. (2020), illustrates a symmetric rounding scheme (blue) as compared to the asymmetric rounding scheme (orange) as implemented in the MWA correlator FPGA fine polyphase filterbank. The asymmetric binning scheme results in a zero bin with half the width of the other bins and has nontrivial impacts on our data quality. Figure credit: (McSweeney et al., 2020)

ism. However, in the case of the PFB asymmetric rounding, the data after quantization is channelized in a FFT, then requantized to 4 bits and correlated. Instead of using the mathematical formulas from Van Vleck, we must create a model of the signal path to determine a relationship between the pre- and post-quantization statistics.

We consider the statistics within each fine frequency channel after the FFT. We need to develop a correction that estimates the standard deviation σ_j^f and covariance $E[X_j^f X_k^f]$ of the fine channel data for the j th and k th antennas without the asymmetric quantization. These depend not only on the standard deviation $\hat{\sigma}_j^f$ and covariance $E[\hat{X}_j^f \hat{X}_k^f]$ of the fine channel after the quantization and FFT but also on the standard deviation σ_j^c of the respective coarse frequency channel before quantization.

The outputs of the 4-bit Van Vleck correction give us estimates of $\sqrt{V_{jj}/2N}$ and $V_{jk}/2N$ which we rescale to obtain our visibilities. Before rescaling these numbers, we can use them as estimates of the standard deviation $\hat{\sigma}_j^f$ and covariance $E[\hat{X}_j^f \hat{X}_k^f]$. We cannot measure the coarse channel standard deviation σ_j^c before the 8-bit requantization stage. We, can however, obtain an estimate of the coarse channel standard deviation $\hat{\sigma}_j^c$ after quantization by taking the square root of the mean of the fine channel variances $(\hat{\sigma}_j^f)^2$ across the coarse band fine frequency channels. We then need to develop corrections which estimate the standard deviation and covariance without quantization, σ_j^f and $E[X_j^f X_k^f]$, from the fine channel standard deviation $\hat{\sigma}_j^f$, covariance $E[\hat{X}_j^f \hat{X}_k^f]$, and coarse channel standard deviation $\hat{\sigma}_j^c$ after the asymmetric 8-bit requantization and FFT. The estimates of σ_j^f and $E[X_j^f X_k^f]$ give us approximations of $\sqrt{V_{jj}/2N}$ and $V_{jk}/2N$ without requantization effects.

While we have done work to estimate the excess DC power in each coarse channel, more work is needed to understand any possible dependence of the DC power on the spectral structure across the coarse channel. This must be well modeled in order to recover the DC channel with confidence.

5.3 *Developing a correction for the asymmetric 8-bit requantization*

By examining visibility data, I determined the ranges of the parameters for the simulation. In our data we see the real and imaginary parts of the scaled crosscorrelations (which we use to approximate covariance) with values $\lesssim 0.3$. From autocorrelations we see that the standard deviations of the fine and coarse channels, respectively, have ranges $1 \lesssim \hat{\sigma}_j^f \lesssim 4$ and $1.2 \lesssim \hat{\sigma}_j^c \lesssim 3.8$.

To find the relationships between the covariances and standard deviations with and without errors from the asymmetric 8-bit requantization, I simulated fine channel data streams sampled from a zero mean Gaussian distribution. For covariances I generated pairs of coarse channel standard deviations σ_j^c and σ_k^c using 61 values ranging between 1-4, inclusive. Combining without replacement resulted in 1891 pairs. For each of 45 initial covariance values between 0-0.35, inclusive, and each σ_j^c, σ_k^c pair I simulated a data stream of $20,000 \times 128$ values. For standard deviations I simulated a data stream of $20,000 \times 128$ values over a grid of 21 coarse channel standard deviations σ_j^c between 0.9-4.4, inclusive, and 21 fine channel standard deviations σ_j^f from 0.5-4.5, inclusive.

I took each simulated data stream backward through the fine channelization FFT to obtain a set of simulated coarse channel data. Then I applied an asymmetric 8-bit requantization to the coarse channel data before taking it back through the FFT. From the FFT output data streams I obtained estimates of the fine channel standard deviations $\hat{\sigma}_j^f$ and covariances $E[\hat{X}_j^f \hat{X}_k^f]$ after the requantization stage. To do this, I masked the DC channel containing excess power after the FFT and calculated the statistics of the remaining $127 \times 20,000$ values of each simulated data stream. Rather than using these simulated values directly to correct the data, I used these estimates to fit initial functions to approximate the relationships between the statistics with and without effects for the asymmetric 8-bit requantization. Fitting a function can reduce the effects of the simulation shot noise and produce a better approximation.

To correct standard deviations, and thus autocorrelations, I fit a 2-dimensional polynomial for σ_j^f as a function of $\hat{\sigma}_j^f$ and $\hat{\sigma}_j^c$, using a third order polynomial in $\hat{\sigma}_j^f$ and a fourth order polynomial in $\hat{\sigma}_j^c$. In Figure 5.2 I show the results of this fit. For three different iterations of σ_j^c the left column shows the standard deviation of the simulated data σ_j^f as a function of the standard deviation of the simulated quantized data $\hat{\sigma}_j^f$. The function fit to this data is overplotted as an orange line. In the center we show the difference between these two sets of standard deviations. As the coarse channel standard deviation σ_j^c increases, this residual becomes less steep. On the right we take the difference of the standard deviations of simulated data without quantization minus the corrected values of the standard deviations with 8-bit requantization effects.

For the covariance (crosscorrelation) correction I fit a 3-dimensional Chebyshev polynomial using only the first, third, fifth, and seventh Chebyshev polynomial terms for $E[\hat{X}_j^f \hat{X}_k^f]$, setting coefficients for even polynomials to zero. In this function $\hat{\sigma}_j^c$ and $\hat{\sigma}_k^c$ were fit with fifth-order Chebyshev polynomials. The results of this fit are shown in Figure 5.3 for three pairs of the coarse channels sigmas σ_j^c and σ_k^c of the j th and k th antennas, respectively. Similarly to Figure 5.2, the first column shows the covariance of the simulated data without quantization as a function of the covariance of the simulated quantized data. The fit function is overplotted in orange. The center and right columns show the residuals between the covariances of non-quantized and quantized simulated data and between the covariances of non-quantized and corrected quantized simulated data, respectively.

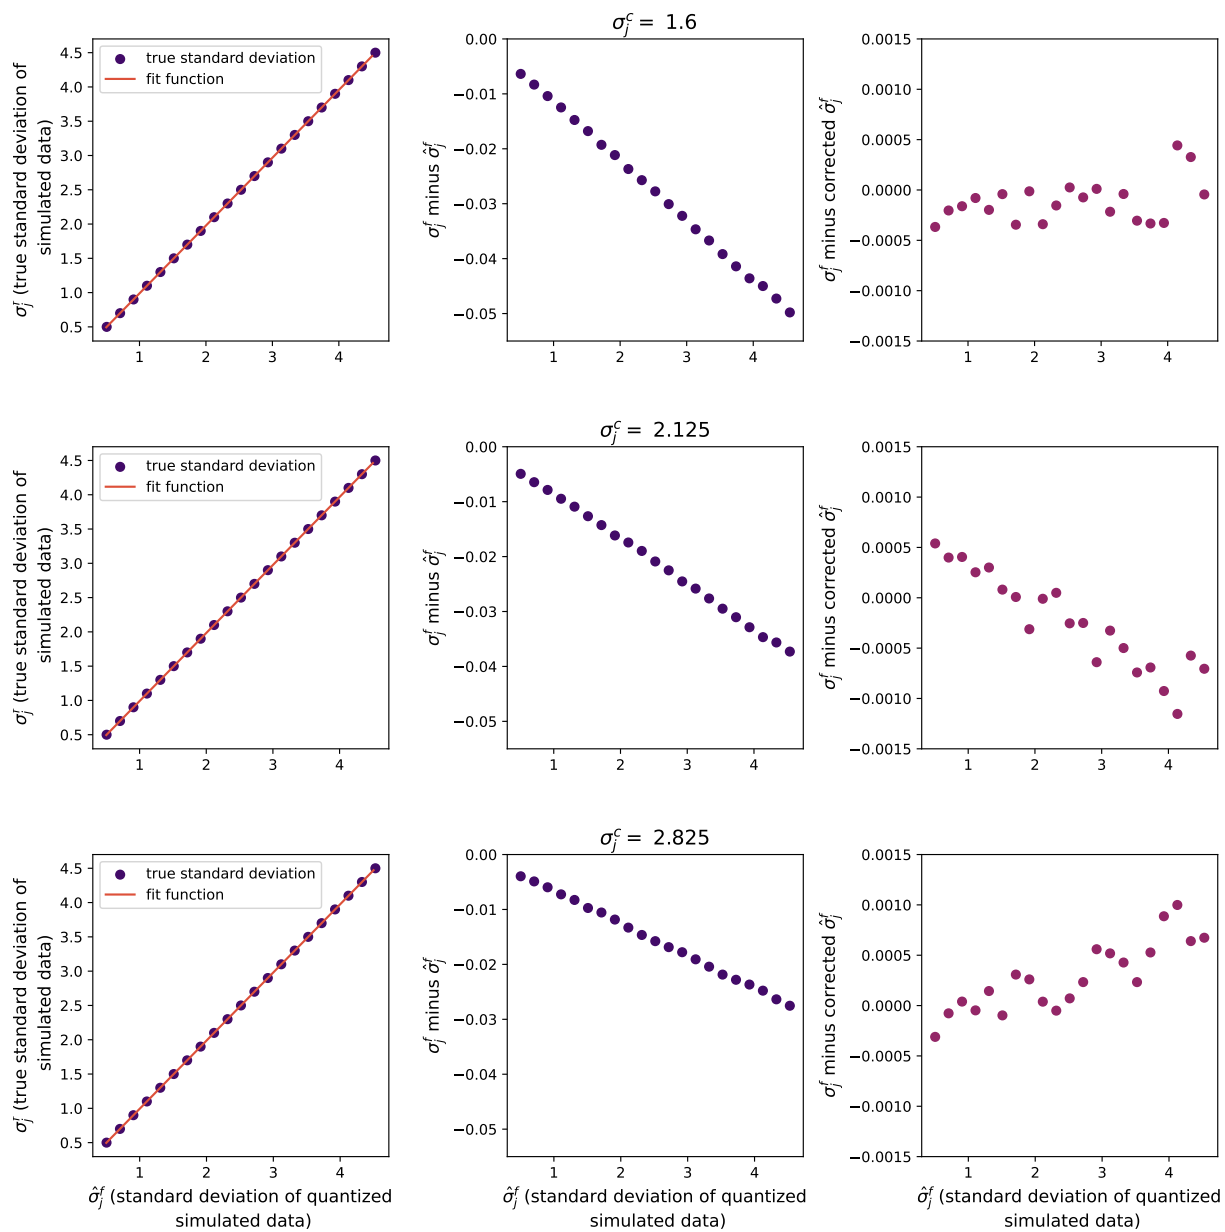


Figure 5.2: We show the results of fitting a function to correct the nonlinear effects of the asymmetric 8-bit requantization on the signal standard deviation. The left column shows the true standard deviation of simulated data without the 8-bit requantization as a function of the standard deviation of simulated data with the requantization nonlinear effects. Each row shows data simulated with a different coarse channel standard deviation σ_j^c . In the center column we plot the residual between the true standard deviation and the standard deviation with quantization nonlinearities. We see a dependence between the steepness of the bias and the coarse channel standard deviation σ_j^c . The third column shows the residual between the standard deviation of the simulated data without quantization effects and the simulated data after correction with the fit function. While some residual bias may be present, there is evident improvement.

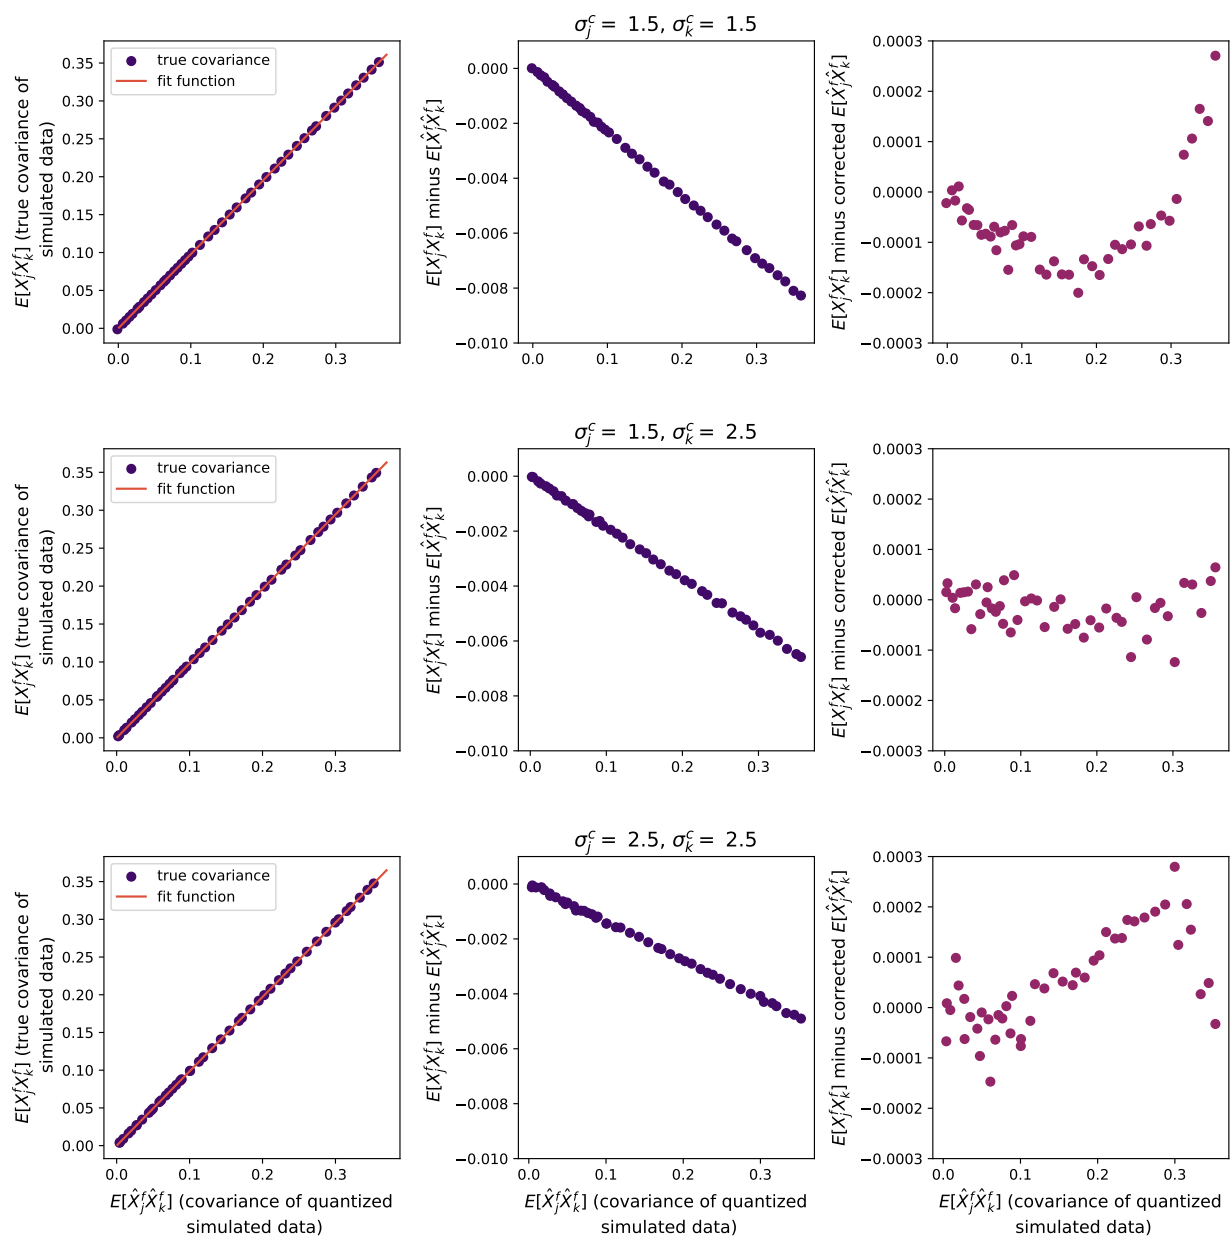


Figure 5.3: We show the results of fitting a function to correct the nonlinear effects of the asymmetric 8-bit requantization in covariances and thus crosscorrelations. The left column shows the covariance of simulated data without the 8-bit requantization as a function of the covariance of simulated data with the requantization nonlinear effects. Each row uses data simulated with different pair of coarse channel standard deviations σ_j^c and σ_k^c . In the center column we plot the difference between the covariance of simulated data with and without the 8-bit requantization effects. Here there is an evident bias with a clear dependence on the values of σ_j^c and σ_k^c . The third column shows the residual after correcting the covariances with the fit function. The correction brings the residual closer to zero, which is promising.

In both Figure 5.2 and Figure 5.3 the results after applying corrections to the data show improvement. The residuals in the right column after correction are significantly lower than those in the center column before correction and more clustered around zero. There is some structure in the residual after the data is corrected which suggests some lingering bias may be present. Further work improving these fits may reduce this structure.

We check the performance of the initial correction functions by analyzing an autocorrelation. We take the same autocorrelation plotted in Figure 4.4 and add the asymmetric 8-bit correction in the bottom panel of Figure 5.4. We see a further decrease in the digital gain jump, highlighted by the red oval. Barry’s analysis of 2013 MWA data excluded data above the gain jump (Barry et al., 2019b), and this improvement indicates that correcting the 8-bit asymmetric requantization could smooth the discontinuity sufficiently to allow reanalysis of 2013 data including the full frequency band. While this is promising, we still see residual artifacts at the coarse band edges. These might be mitigated by improving the functions for the 8-bit correction, or might require further modeling of the 5-bit requantization stage between the digital receiver and the correlator.

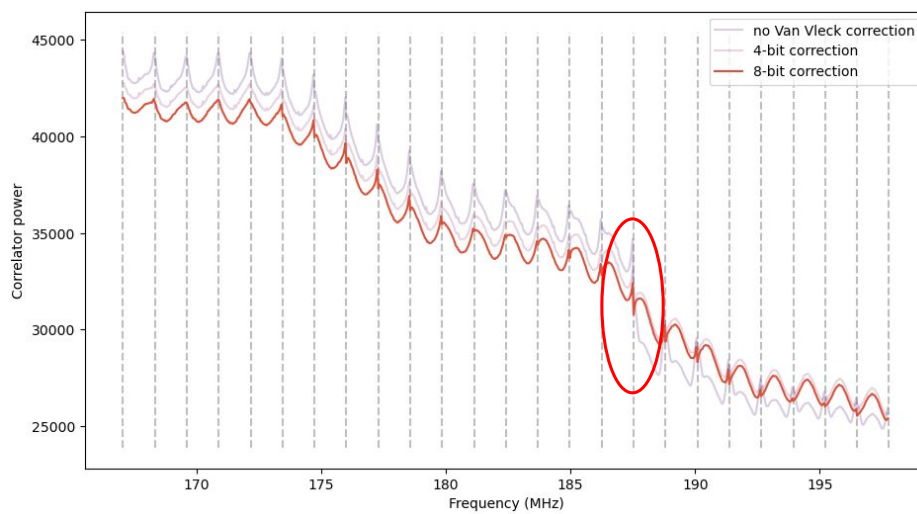
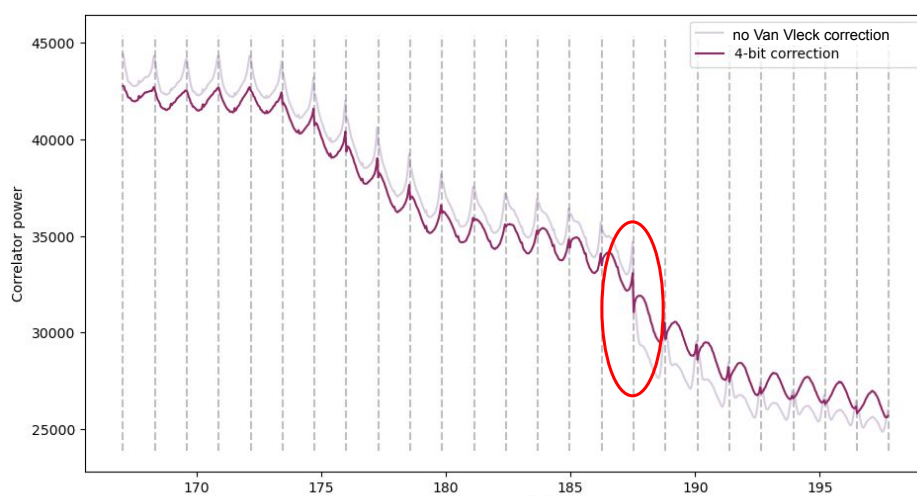



Figure 5.4: An autocorrelation from a single zenith observation from 2013 MWA data is plotted for antenna Tile028. In the upper panel we have applied corrections for the digital gains and polyphase filterbank (PFB) response, and in the center panel we have applied our implementation of the 4-bit Van Vleck correction. After the 4-bit correction we still see the residual digital gain jump highlighted by the red ovals, though it has improved. Also, we see residual artifacts at the coarse band edges, which are delineated by dashed grey lines. In the bottom panel we have applied the 8-bit quantization correction after the 4-bit correction. The gain jump has shrunk further after this correction, indicating that correcting the asymmetric 8-bit requantization might allow recovery of the upper third of the band which was excluded from Barry’s analysis of 2013 data (Barry et al., 2019b). The lingering artifacts at the coarse band edges suggest further work should be done. This residual spectral structure might be improved by further work on the 8-bit requantization correction as well as developing corrections for the 5-bit requantization stage between the receiver and correlator.

The corrected autocorrelations are promising and suggest it would be valuable for future students to continue this work. During my work developing requantization corrections for MWA data, a new correlator for the MWA was being developed. This new MWAX correlator is the first in a series of upgrades for MWA Phase III, as we will discuss in Chapter 6. Partly due to our work demonstrating the improvements from correcting nonlinear effects, a design goal of the MWAX correlator was to eliminate requantization stages, including both quantization stages explored in this thesis. To evaluate the effects of the new correlator on nonlinearities in the data, I pivoted to analyzing the first EoR data taken with the MWAX correlator.

Chapter 6

MWA PHASE III AND THE NEW MWAX CORRELATOR

In 2021, a new correlator was commissioned for the MWA (Morrison et al., 2023). This correlator, “MWAX”, in 2022 officially replaced the “Legacy” correlator described in Section 3.4. Largely as a result of our work in Chapters 4 and 5, one of the primary design goals of the MWAX correlator was eliminating the nonlinear effects resulting from quantization stages in the Legacy correlator. Other design goals included expanding correlation capacity from 128 to 256 antennas, adding capability for implementing cable delay corrections and fringe-stopping, increasing flexibility in time and frequency resolution options for output data, and eventually enabling oversampled coarse PFBs with envisioned digital receiver upgrades.

The MWAX correlator is the first upgrade of Phase III of the MWA. During MWA Phase I, from 2013-2016, the telescope comprised 128 antennas. The telescope was upgraded to Phase II in 2016, a process which added 128 more tiles to the array. This took the telescope from 128 to 256 tiles, adding 72 tiles in two hexagonal grids within the ‘core’, that is, within the central area encompassed by the Phase I configuration, as well as 56 distant tiles outside the Phase I area. However, the digital system only had processing capacity for 128 tiles at a time. Two configurations of 128 tiles were developed: compact and extended (Wayth et al., 2018). Switching between these configurations requires disconnecting tiles from receivers, moving receivers to other areas of the array, and then connecting the desired tiles. During Phase II (2016-2021), this reconfiguration was done about two times each year.

Phase III upgrades to the digital system will ultimately allow for processing data from all 256 tiles at once. The first stage of Phase III is the implementation of the new MWAX correlator. The next stage adds upgraded National Instrument (NI) digital receivers. A receiver

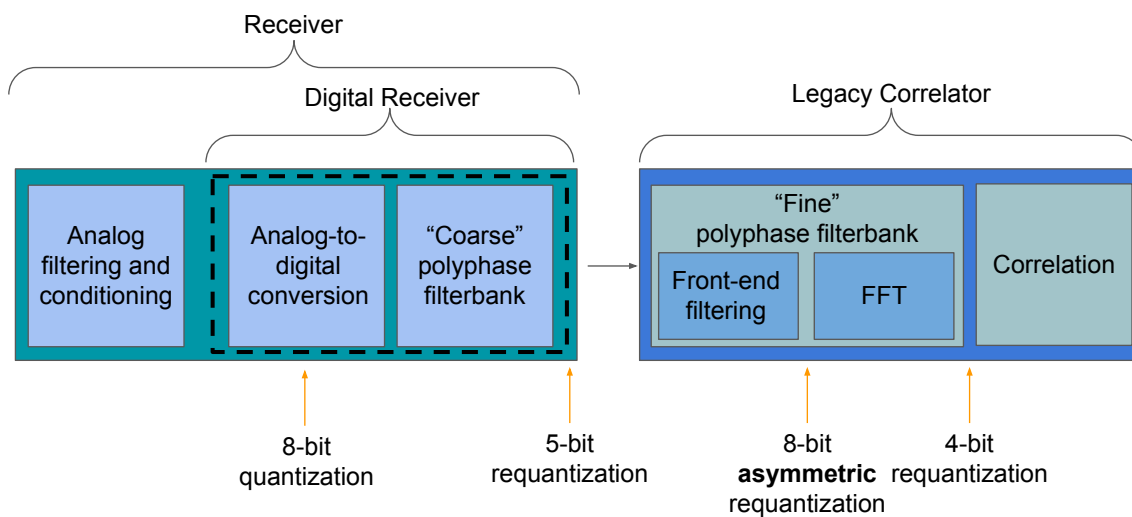
processes data from 8 antennas, and as each new receiver is added 8 more antenna tiles are able to be connected to the digital system until all 256 tiles can be in use simultaneously. As of January 2024, two new receivers have been added, raising the MWA signal processing capacity to 144 antennas (Williams, Accessed June 6 2024).

In Figure 6.1 we show the MWA Phase I/II digital system in the upper panel. Below we show the upgraded Phase III digital system. As we will discuss in more detail below, the new MWAX correlator maintains a minimum depth of 16 bits for each real and imaginary value. This holds exciting promise for recovering regions of the EoR Window contaminated by nonlinear effects from the Legacy correlator. Additionally, the new NI receivers use an 8-bit requantization stage before sending data to the correlator rather than a 5-bit requantization stage, which will further reduce the nonlinear response. The NI receivers can also employ an oversampled polyphase filterbank, which will flatten the strong spectral response of the old receiver PFB implementation.

6.1 MWAX design

Historically, MWA digital receivers have used the Xilinx RocketIO protocol to output their data, which in the Legacy correlator was received by Xilinx FPGAs. MWAX is designed to ingest User Datagram Protocol (UDP) packets, so the fiber optic cables from the receivers (Section 3.3.2) now arrive at a set of media conversion servers which perform the conversion from RocketIO to UDP. Fiber optic cables connect the media conversion servers to a Nexus switch, which routes packets of data to GPU-accelerated MWAX ‘Server’ compute nodes. The upgraded NI digital receivers output UDP data directly to the Nexus switch, eliminating the need for the media conversion servers.

Phase I & II



Phase III

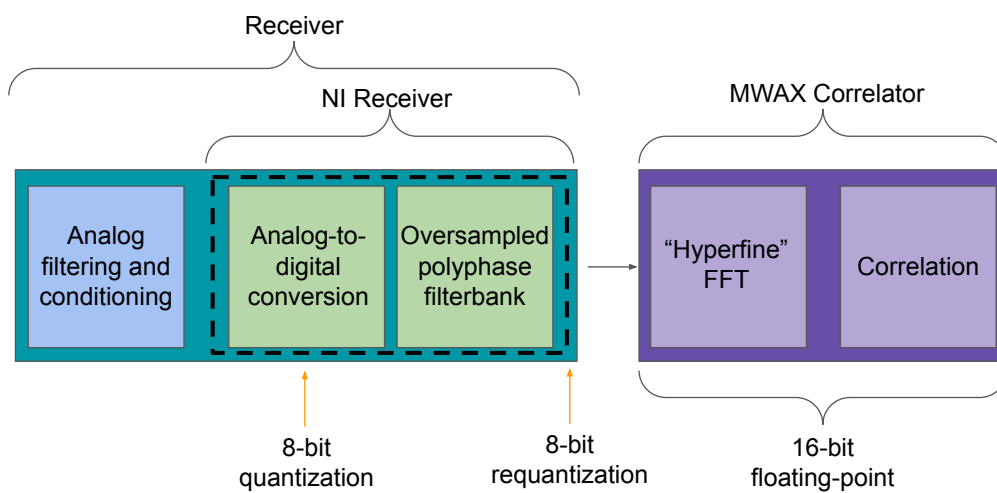


Figure 6.1: We show overviews of the MWA Phase I/II (top) and MWA Phase III (bottom) digital systems, respectively. The first upgrade of MWA Phase III is the new MWAX correlator, which officially replaced the Legacy correlator in early 2022. The MWAX correlator eliminates the 4-bit and asymmetric 8-bit requantization stages of the Legacy correlator, instead maintaining a bit-depth of 16 floats through its components. This will remove both the nonlinear effects from the Legacy correlator we detailed in Chapters 4 and 5. The NI receivers requantize data at 8 bits rather than 5 bits before sending it to the correlator, which will further reduce quantization nonlinearities. Replacing the coarse polyphase filterbank with an oversampled PFB will smooth the strong spectral response in the Phase I/II coarse band. The Legacy correlator fine PFB is replaced with an ultrafine FFT in MWAX, which reduces spectral leakage between modes. As of January 2024, two NI receivers have been added to the array, bringing the MWA processing capacity to 144 antennas. As more receivers are added, Phase III will ultimately allow for correlation of all 256 tiles at once.

The data received by the Nexus switch comprises twenty-four 1.28 MHz frequency coarse channels. Before routing data to the servers, the Nexus switch groups the data by these channels such that each server receives and processes a single frequency coarse band. There are currently 24 MWAX Servers, and a possible future upgrade for MWAX would increase the frequency bandwidth by adding more servers. Upgrades to MWAX Server components could also allow for correlation of more than 256 antennas. The servers each consolidate the data from their respective coarse channel into 8 second packets for three processing stages: channelization, correlation and time accumulation, and frequency averaging. These three stages are implemented by a GPU F-engine, a GPU X-engine, and a GPU function, respectively. Finally, the server CPU reorders the output visibilities and writes them to a buffer, from which a separate process reads the data and writes it to output files.

In MWAX, channelization occurs via an ultrafine FFT implemented on the MWAX Server F-engine GPU. Each 8 second packet is processed 50 ms (64,000 samples) at a time. First, the data is promoted to 32-bit floats. Next, the 64,000 samples are split into 10 groups of 6,400 samples each. These groups are taken through a 6,400 point FFT, resulting in output channels of 200 Hz width and a time resolution of 5 ms. The hyperfine 200 Hz frequency resolution after the FFT allows for a greater range of frequency averaging options than the Legacy correlator PFB output of 10 kHz channels.

The FFT outputs are transposed and written to GPU memory, where they can be directly accessed by the X-engine GPU. Keeping the data in GPU memory buffers allows bypassing an intermediate data capture processes between the F-engine and X-engine which would require requantization. In the memory buffer, data are collected into 250 ms chunks, limiting the possible final output integration times to be multiples of 250 ms which divide evenly into 8 seconds.

Correlation is implemented using the xGPU library (Clark et al., 2012; Ord et al., 2015), and the correlation products summed to the desired integration time resolution. The resulting visibilities are written to GPU memory where frequency averaging is performed by a separate GPU function. This is a key step, as averaging the fine channels smooths the response of the ultrafine FFT and mitigates spectral leakage. The frequency averaging schemes differ between Legacy and MWAX, as we will discuss in Section 10.1.1.

Finally, the averaged visibilities are written to the server host memory. A CPU process reorders the visibilities into [time, baseline, frequency, polarization] triangular order, where the four axes are ordered from slowest-changing to fastest. The reordered data is written to a buffer, from which a separate process reads it and writes it to the output FITS data files.

MWAX maintains a minimum of 16 + 16 bit floating point precision throughout its data processing. This is made possible by utilizing a GPU FFT implementation rather using

FPGAs as is typical for polyphase filterbanks. The GPU FFT can write its outputs to GPU buffers where they can be accessed directly by the correlation X-engine. This bypasses any need to utilize requantization in order to efficiently transfer data between components. Performing frequency averaging via a GPU function similarly avoids potential requantization stages. The data volume is decreased before it is written back into host memory, removing the need to reduce bloat with requantization. This design results in an exciting new instrument with great potential for removing nonlinear systematics from the EoR Window.

6.2 *Understanding MWAX*

In addition to the comprehensive MWAX correlator paper (Morrison et al., 2023), the MWA wiki has a few pages I have found particularly useful for understanding various nuances surrounding the introduction of the MWAX correlator. Discussions of changes to the output visibility files (Sleap et al., Accessed May 7 2024) and metafits files (Williams and Sleap, Accessed May 7 2024) were invaluable when I upgraded the `pyuvdata` reader to ingest files from MWAX as well as the Legacy correlator. Also, while not MWAX-specific, I have found the general page on metafits files to be extremely useful (Williams, Accessed May 7 2024a). Code for calculating the digital receiver PFB bandpass for MWAX fine channels was written by Jones (Accessed May 7 2024). I have reread the page on correlator frequency averaging (Sleap and Null, Accessed May 7 2024) countless times in order to understand the frequency averaging behaviors of both MWAX and Legacy.

6.3 *A first look at MWAX data*

The first set of MWAX EoR data was taken during commissioning in late 2021. In order to understand the performance of the new correlator for our science, we perform a set of investigations on this novel data set. First, in Chapter 7 we carefully inspect autocorrelations from the MWAX correlator for signs of digital nonlinearities and other systematics. Next, we prepare to take the 2021 MWAX data through our entire analysis pipeline (Chapter 8). This requires re-evaluating the analysis choices we have historically used for data from the

Legacy correlator and retuning our RFI detection algorithms for the 2021 RFI environment. We determine optimal flagging and averaging schemes for MWAX data in Chapter 10 and perform data selection in Chapter 11 before integrating the 2021 commissioning set and obtaining the first power spectrum limit with Phase III data in Chapter 12.

Chapter 7

ANALYZING MWAX PERFORMANCE THROUGH AUTOCORRELATIONS

Autocorrelations—visibilities in which the signal measured by a tile is correlated with itself—are a particularly instructive tool for understanding instrument response. The autocorrelations are real-valued measurements, as the complex signal is multiplied by its complex conjugate. Thus we do not preserve information regarding the phase response of the antennas but instead take advantage of the amplitude to reveal characteristics. The antenna response is encoded within this amplitude, and so to understand the performance of the MWAX correlator we first examine autocorrelations to evaluate the spectral structure of the new instrument.

7.1 *MWAX antenna response*

In order to develop a clear picture of the antenna response, we average the data in time over a single two-minute observation. We use an observation from a zenith pointing in which no radio frequency interference is detected by SSINS (Section 11.1). Then, we remove known effects from the polyphase filterbank, digital gains, and cable reflections (Section 3.3.2), as these can obscure more subtle artifacts. Additionally, subtracting a second-order polynomial fit helps to reveal spectral structure. To visualize the impact of each of these effects, we remove them one-by-one and examine the autocorrelations at each stage. This process is visualized in Figure 7.1. The top panel shows the raw autocorrelation after time averaging, before removal of any digital signatures. The first fine channel of each coarse band has been flagged, as it is known to be corrupted (Section 10.1.1). The frequency band is subdivided into 24 coarse bands of thirty-two 40kHz channels each due to the two channelization stages

in the digital receiver and correlator, respectively. Dashed grey lines in the figure delineate these coarse bands.

Perhaps the most dramatic digital effect visible in the autocorrelations is the polyphase filterbank (PFB) spectral structure. We use code by Jones (Accessed May 7 2024) to compute and divide out the PFB bandpass for the 40kHz resolution of our data. The PFB spectral response drops off steeply at the coarse band edges, and digital nonlinearities result in residual dips or peaks at the coarse band edges after the response is removed. The second panel from the top in Figure 7.1 shows the autocorrelation after this correction. Some residual peaking seems evident at the coarse band edges, but it is difficult to distinguish. Behavior at the edges is somewhat obstructed by the digital gains, which we divide out next.

Digital gains are applied as a single number for each coarse channel. In the center panel of Figure 7.1 we divide out these numbers and examine the resulting output for residual discontinuities, or ‘jumps’ in coarse band power, which indicate nonlinearities in the digital system. The resulting spectrum is relatively smooth across the entire band. However, after the digital gain removal, PFB artifacts become more evident. At the rightmost coarse band edges in particular we see an uptick in power resulting from mis-subtraction of the PFB spectral shape. This could be due to the 5 + 5 bit quantization stage in the digital receivers (Section 3.3.2), but could also indicate nontrivial effects due to digital nonlinearities in the MWAX correlator. In order to better reveal other antenna structures, we increase flagging to three fine channels (120 kHz in total) at each edge of each coarse band for the final two panels of the plot.

We now scale and flatten the autocorrelation by fitting and dividing out a second order polynomial. In the output, shown in the second panel from the bottom, the antenna analog structure finally begins to reveal itself.




Figure 7.1: Digital corrections applied to autocorrelations. The top panel shows the autocorrelation for Tile022 for a single zenith observation, averaged in time with no corrections applied. The first fine frequency channel of each coarse band has been flagged. Coarse channels are delineated by vertical dashed grey lines. In the second panel the polyphase filterbank bandpass has been divided out. Digital gains are divided out in the center panel. Here artifacts at the coarse band edges due to aliasing and/or digital nonlinearities are evident, particularly as upticks in power at the right-hand coarse band edges. To better reveal the antenna response, these artifacts are removed by flagging three fine channels (120 kHz in total) at each coarse band edge in subsequent panels. Next, a second order polynomial is fit and removed. Finally, the true lumpy bumpy antenna structure is evident in the bottom panel after dividing out a cable reflection term.

Prominent in this view is the ripple due to cable reflections, which we fit and remove. Reflections at the cable ends between antenna tiles and receivers result in ripples across the frequency band. The ripples depend on the type and length (‘flavor’) of cable used, which include 90, 150, and 230 m RG-6 coaxial cables, and 320, 400, and 524 m LMR-400 coaxial cables. The cable flavor for each antenna is recorded in the metafits file. In an autocorrelation, the spectral ripple due to cable reflection can be described by the following term:

$$1 + a^2 + 2a \cos(2\pi\tau f + \phi) \tag{7.1}$$

where f is frequency, and we fit for a , τ , and ϕ as described by Barry et al. (2019a). Using the cable flavor information, we estimate τ . We take a Fourier transform in frequency to obtain initial estimates of the amplitude a and phase ϕ of the peak at our estimated delay τ . Then we take a hyperfine transform around the peak to improve our estimates of a and τ . With these final values, we perform a least-squares fit of Equation (7.1) to the data to solve for our final estimate of ϕ . We compute Equation (7.1) with the values found for τ , a ,

and ϕ and divide it out of the autocorrelation.

After this correction, the antenna structure is uncovered in the final panel of Figure 7.1. Two features are immediately evident: larger-scale lumps which span several coarse bands, and smaller bumps of width 1-2 coarse bands. We now apply the averaging and correcting process described above to all autocorrelations from our observation, allowing us to carefully examine and evaluate the response of each antenna.

7.2 *Antenna temporal stability*

If the autocorrelation frequency structure was coming from a sky source moving through the instrument sidelobes, particularly the galactic center or other very bright off-axis radio sources, we would expect it to change over time. To determine if we have truly revealed the antenna response, we analyze a single tile's autocorrelation from each of the four pointings for a single night: Minus One, Zenith, Plus One, and Plus Two (these pointings are described in Table 8.1). The autocorrelations are shown in Figure 7.2, where it is clear the antenna structure is steady through the night. Additionally, the large-scale fluctuations in power are similar between the instrumental polarizations XX and YY, indicating this response is resulting from signal processing components common between both polarizations. More variation between polarizations is evident in structure on the 1-2 coarse band width scale.

We also compare two autocorrelations for a single tile taken from zenith pointings separated by two nights. These are plotted in Figure 7.3, where again the antenna structure demonstrates temporal stability. These tests both indicate that the autocorrelation shape is largely independent from the time the observation was taken, whether within a single night or between nights. Thus the structure revealed appears to truly be from the instrument itself rather than the signal being measured.




Figure 7.2: An autocorrelation from Tile022 is averaged over an observation from each pointing in a single night, roughly thirty minutes apart. The lumpy bumpy structure across the band is consistent between the observations, indicating that it is not due to the sky signal but rather the instrument response. Additionally, we see that the antenna structure is similar between polarizations for larger-scale fluctuations but smaller-scale lumpiness of the 1-2 coarse band width shows more variation.




Figure 7.3: An autocorrelation from Tile022 for two observations taken two nights apart. The lumpy bumpy structure across the band demonstrates temporal stability, indicating that the autocorrelation is clearly capturing the antenna's persistent response.

7.3 *Antenna rogues' gallery*

We now provide a sampling of autocorrelations from our zenith pointing observation which represent trends or anomalies within the full set. For each antenna shown, both XX and YY instrumental polarizations are plotted. Boundaries between frequency coarse bands are delineated by dashed grey lines. The antenna name is labeled in the upper right corner of each panel. This rogues' gallery is intended to provide a summary of important takeaways from this particular observation as well as generally demonstrate how autocorrelations can be utilized to detect patterns in our data.

One effect the averaged and corrected autocorrelations reveal is the frequency dependence of the cable reflection ripple. As seen in Figure 7.4, which shows an antenna (Tile081) with a ninety meter cable, dividing out the reflection term did not capture the ripple structure across the entire frequency band. Residual cable rippling is particularly apparent at the bottom of the band. This mis-subtraction indicates that there is actually a frequency dependence not contained in Equation (7.1). Future work could update the reflection term to capture this effect.

For our science, we have strict calibration accuracy needs (Chapter 2). In particular, variations in frequency within the antenna response must be captured with a precision of a part in 10^{-5} . Ideally, our antenna responses are smooth in frequency. Additionally, if our antenna responses are uniform, it is easier to describe their structure with a single common term during calibration. We examine the autocorrelations for sameness and smoothness in antenna structure. As seen in Figure 7.1, the antenna response is lumpy and bumpy rather than smooth across the band. An examination of all the antennas from this observation reveals that lumpiness and bumpiness is a typical feature. However, the size and shape of both large- and small-scale power fluctuations vary widely between antennas. Some antennas show more smoothness in their structure than others, as we see in Figure 7.5. Here we plot two antennas, Tile082 and Tile087, which are both relatively smooth in their small-scale

frequency structure. The larger-scale structure of Tile082 in the upper two panels is also relatively stable in frequency, while Tile087 in the lower two panels has three large bumps along the frequency band. Many autocorrelations from this observation reveal a similar pattern of 2-3 large-scale fluctuations, though the size and location varies from antenna to antenna. The necessity of accurately capturing this frequency dependent structure in our calibration solutions provides impetus for utilizing the autocorrelations themselves within the calibration gains (Section 8.3.1).

While generally quite different from each other, most antennas have similar large-scale structure in their own XX and YY polarizations. Though variation is common between an antenna's polarizations for structure on the scale of 1-2 coarse band widths, the overall large-scale fluctuations in power tend to have the same signature. There are a few exceptions to this, as shown in Figure 7.6. Here we have plotted two antennas, Tile041 and Tile063, which exhibit this behavior. For Tile063 in particular, the YY polarization has almost periodic fluctuations in power which are not present in the XX polarization. This indicates some instrumental effect interacting differently with the two polarizations.

Autocorrelations can reveal radio frequency interference (RFI) not captured by SSINS (Section 11.1). Two examples of this are plotted in Figure 7.7 which shows autocorrelations from HexE33 and HexS5. Both polarizations of HexE33 have a spike in power near 184 MHz, with a stronger peak in the XX polarization. A spike near 177.5 MHz is seen in both polarizations of HexS5, though in this case the spike power is greater in the YY polarization. In this observation, spikes at these locations are visible only in these two antennas respectively. This implies a localized RFI source, either generated by the respective antenna signal chain or by something at or near the respective antenna location. A later examination of calibration solutions (Section 11.2) reveals that the spike near 177.5 MHz is persistent in HexS5 data across multiple nights.

An interesting artifact revealed in the autocorrelations is an apparent drop in power in the third coarse band from the top of the frequency band. Shown in Figure 7.8, this effect is seen on several antennas, all of which have 230 meter cables. The offending coarse band has been circled in red in the figure. Since the digital gains are the same for all 230 meter antennas in this observation, this coarse band power artifact could be due to a discrepancy between the digital gain applied to the channel and the gain recorded in the metafits file.

We also find in multiple autocorrelations a fascinating spike in power in the center of the same coarse band (Figure 7.9). Unlike the drops in power, these spikes are not limited to antennas with 230 meter cables. However, they are not visible in all antennas and are frequently barely detectable in the autocorrelations from this observation. It is perhaps coincidental that this narrowband RFI signal appears within the same coarse band that also demonstrates the drop in power for antennas with 230 meter cables. The source(s) of both the coarse band discontinuity and the spike at ~ 194 MHz provides an intriguing mystery which has yet to be solved.

Now that we have performed a diagnostic inspection of antenna responses with the new MWAX correlator we are prepared to move forward with taking a limit. The visible decrease in artifacts after removing the linear digital gains and PFB spectral structure is very encouraging, and points to the elimination of many of the digital nonlinear effects seen in the Legacy correlator. This is promising for our power spectrum limit analysis.








Figure 7.4: An autocorrelation from Tile011 which has a 90 meter cable reflection. Residual ripples after fitting and removing the 90 meter cable reflection are evident in the bottom of the frequency band. This shows that there is in fact a frequency dependence to the cable ripple which is not captured in the reflection term (Equation (7.1)). Future work could derive and fit for this frequency dependence to better remove the cable reflection from the autocorrelations.



frequency (MHz)

Figure 7.5: Autocorrelations from a single zenith observation are plotted for two antennas: Tile082 and Tile087. While most autocorrelations reveal large-scale lumps spanning several coarse channels, as shown in the lower two panels, the size and location of this structure varies antenna to antenna. Additionally, many are much smoother in their structure on this scale, as seen in the upper two panels. This has the important implication that the antenna structure varies greatly throughout the telescope, not just in the bumpy structure seen at the 1-2 coarse channel width but also in the larger-scale fluctuations across the band. This variability between antennas makes autocorrelations particularly useful in calibration, as they provide a clear signature of each antenna's unique response. Most of the autocorrelations reveal fluctuations in power across the frequency band on some scale. Our science requires that calibration accurately capture frequency-dependent structure to a part in 10^{-5} (Chapter 2), so the lack of smoothness evidenced by the autocorrelations is particularly concerning to us.







Figure 7.6: Autocorrelations from a single zenith observation are plotted for two antennas: Tile041 and Tile063. Typically the antenna structure revealed in autocorrelations from this observation is very similar for each instrumental polarization XX and YY, particularly the large-scale fluctuations in frequency. More variability is frequently seen in structures of the 1-2 coarse band width scale. However, there are a few exceptions, two of which are shown here. The upper panel shows Tile041, for which the YY polarization structure (shown in the second panel from the top) has more dramatic large-scale structure in the center third of the band than the corresponding XX polarization (top panel). The bottom two panels show the XX and YY polarizations for Tile063, respectively. The YY polarization structure is markedly different from the XX, likely revealing some polarization-specific instrumental effect.















Figure 7.7: Autocorrelations from a single zenith observation are plotted for two antennas: HexE33 and HexS5. The averaged and corrected autocorrelations can reveal narrowband radio frequency interference (RFI) undetected by SSINS (Section 11.1), as shown here. The upper two panels show a sharp peak near 184 MHz which is strongest in the XX polarization (topmost panel). The bottom two panels reveal a peak at ~ 177.5 MHz which is much stronger in the YY polarization of antenna HexS5. For this observation, peaks at these frequencies are only observed in these two antennas, respectively. This indicates that the antennas are picking up localized signals from themselves or nearby in the array. Later examination of calibration solutions for multiple observations through the observing season reveals the peak near 177.5 MHz to be persistently present in antenna HexS5.



frequency (MHz)

Figure 7.8: Autocorrelations from a single zenith observation are plotted for two antennas: Tile086 and HexS21. Several antennas from this observation with 230 m cables show dips in power in the third coarse band from the top (circled in red in this plot). As the digital gains are the same for all antennas with 230 m cables, this discontinuity might indicate a miscalculation or misapplication in the digital gains for this coarse channel. As shown in the next plot (Figure 7.9), this particular coarse band is also sometimes the location of a narrow spike in power.



170 175 180 185 190 195

frequency (MHz)

Figure 7.9: Autocorrelations from a single zenith observation are plotted for two antennas: Tile096 and HexE21. Several antennas from this observation with 230 m cables show dips in power in the third coarse band from the top (circled in red in this plot). In this coarse band there is also frequently a narrow RFI spike near the center of the band. Unlike the power drops, which are found in antennas with 230 meter cables, these spikes are also evident in antennas with different cable lengths. It is likely a coincidence that the location of this narrowband RFI signal is within this particular coarse band which has its own unique artifact. An examination of autocorrelations from an observation two nights later reveals that these spikes in power are present for that night as well, showing that whatever is the source of this signal, it is not limited to this single observation. Future work could involve plotting the locations of antennas which detect this RFI in an attempt to discover the location of the source.

Chapter 8

PREPARING FOR A LIMIT WITH MWAX DATA

Now that we have evaluated MWAX performance by carefully inspecting autocorrelations, we are ready to take the MWAX data set through our FHD/ ϵ analysis pipeline and obtain an upper limit. First we describe the data set (Section 8.1) and then provide an overview of the preprocessing (Section 8.2) and analysis (Section 8.3) pipelines.

8.1 MWAX limit data set

For the 2021 MWAX data set, the MWA is in the Phase II compact configuration. This configuration incorporates two regularly spaced hexagons, allowing for analysis tools which take advantage of the redundancy of baseline length and orientation as demonstrated by Li et al. (2019) and Kolopanis et al. (2019). Since there is less diversity in baseline length and orientation with this antenna layout, more complete uv coverage is not obtained until observations from multiple times and pointings are averaged together.

The data processed here was taken in the 2021 season. Typically, we record EoR data during September-December. This set was taken relatively late in the season over 13 nights spanning November 16-December 10. It covers the ‘High Band’ frequency range (167 - 197 MHz) and is pointed at ‘EoR0’, a relatively dark region of the sky centered at 0 hours right ascension and -27 degrees declination. To stay directed at this field, the instrument is re-pointed approximately every 30 minutes. The 2021 EoR High Band data was taken at four pointings: Minus One, Zenith, Plus One, and Plus Two, corresponding to telescope gridpoint numbers of 2, 0, 4, and 12, respectively. For EoR0, these gridpoints are summarized in Table 8.1. A complete set of gridpoints available for the MWA can be found at Williams

(Accessed May 7 2024b). Our observations were taken at a time resolution of 2 seconds and a frequency resolution of 40 kHz. The set contains a total of 319 observations, each of which are 2 minutes long.

8.2 Preprocessing pipeline

Our preprocessing pipeline is implemented in two stages. First, radio frequency interference is detected and flagged using Sky Subtracted Integrated Noise Spectra (SSINS). Then we use the open-source software `pyuvdata` to ingest raw correlator files, implement corrections, apply SSINS flags, and write out a UVFITS file for each observation. We discuss SSINS flagging in Chapter 11 and describe our preprocessing with `pyuvdata` below.

8.2.1 pyuvdata preprocessing

In order to process MWAX output files with `pyuvdata`, I upgraded my existing functionality for ingesting bespoke MWA correlator `gpubox` files and the `metafits` file (as discussed in Chapter 4) to accommodate the differences in MWAX data formatting. For MWA data we typically utilize corrections for digital gains, the polyphase filterbank response, and cable length differences. A review of the digital system (Section 3.3.2) may be instructive for understanding these corrections, which are discussed in more detail below. Other optional preprocessing steps implemented in `pyuvdata` include averaging in time and/or frequency, removing flagged antennas, and phasing to the observation pointing center.

Digital gain correction

In the receivers, digital gains are applied to the data. These gains, stored in the `metafits` file, are a set of integer numbers for each antenna and each coarse frequency band. The correction is applied by first dividing the gains by a factor of 64, and then dividing the visibilities by the two scaled gains for the relevant antennas. To understand the division by 64, we take a quick diversion into the history of MWA digital gains.

Pointing name	Sweet gridpoint	EOR1 gridpoint	Azimuth (degrees)	Altitude (degrees)	Initial observations	Final observations
Minus Three	26	5	90	69.1655	-	-
Minus Two	10	3	90	76.2838	-	-
Minus One	2	1	90	83.1912	31	25
Zenith	0	0	-	90	81	55
Plus One	4	2	270	83.1912	135	70
Plus Two	12	4	270	76.2838	72	30
Plus Three	28	6	270	69.1655	-	-

Table 8.1: A summary of MWA instrument pointings for the EoR0 field (0/27 degrees right ascension/declination). Pointings which are included in the 2021 MWAX data set are shown in bold. A pointing describes the altitude and azimuth on which the instrument beam is centered. By applying delays to the dipoles, the MWA can be pointed to various locations on the sky. The dipole delays are restricted to be integer multiples of 435 picoseconds, in which the integers range from 0-31. This results in a grid of 196 possible sky locations, known as the ‘sweet’ grid. A subset of eleven sweet gridpoints was initially selected for EoR science. This subgrid is referred to as EOR1. Older EoR observations were submitted with EOR1 gridpoint numbers, while currently sweet gridpoint numbers are used. The metafits file “GRIDNAME” and “GRIDNUM” fields indicate which grid and gridpoint were used for the observation. For EoR0 observations, up to seven pointings can be used. The pointing names, in the leftmost column of the table, are described in relation to the Zenith pointing, which corresponds to delays of zero. Minus Three, Minus Two, and Minus One are earlier than Zenith, while Plus One, Plus Two, and Plus Three are later. The corresponding sweet gridpoints and EOR1 gridpoints are in the next two columns. The third and fourth columns give the azimuth and elevation coordinates of the pointing in degrees. The initial number of observations in the 2021 MWAX data set for each pointing is in the fifth column, and the number of observations remaining after quality assessment cuts is in the final column.

During the initial 2013 commissioning of the MWA, the digital receiver polyphase filterbank’s output of 16 + 16 bit real and imaginary pairs was downselected to 5 + 5 bit pairs by choosing bits numbered 09-05, inclusively. Digital gains are multiplied before the 5 + 5 bit selection, and within this bit range, applying digital gains > 7 resulted in oversaturation. Thus early in the telescope commissioning, the digital gains were restricted to integers in the inclusive range 1-7 (see (Crosse and Williams, Accessed June 9 2024) for a discussion of digital gain history). Researchers discovered that this was not sufficiently flattening the spectrum and that a greater resolution in gain choice was necessary in order to smooth in frequency. To allow for this, the bit selection was moved up by 6 bits, from bits 09-05 to bits 15-11. This increased the granularity of the gains, allowing for more fine-tuning across the frequency band. Increasing the original gains by a factor of 64 with the new bit selection resulted in the same 5 + 5 bit outputs sent from the digital receivers to the correlator. Thus dividing the gains by a factor of 64 before removing them gave the same visibilities as the original commissioning data before the bit selection change. Until mid-2014, the initial 1-7 integer gains were simply multiplied by a factor of 64 and applied, resulting in the digital gain ‘jump’ described in Chapter 3. While the gains used now are no longer linked to the original gains by a factor of 64 but rather take advantage of the increased granularity, the convention of rescaling the gains by a factor of 64 before dividing them out is still utilized.

Polyphase filterbank correction

As described in Section 3.3.2, the MWA has a polyphase filterbank (PFB) in the digital receiver which channelizes the frequency band into 1.28 MHz channels and takes the data from real to complex values. A second channelization stage occurs in the correlator (discussed in detail in Section 3.4 and Chapter 6), resulting in a frequency resolution of 10 kHz for the Legacy correlator and 200 Hz for the MWAX correlator. These fine channels carry the response of the initial PFB channelization stage (shown in Figure 3.7). For the Legacy system, this response was derived by Levine (2012b), and for the MWAX system, the response was derived by Jones (Accessed May 7 2024). In `pyuvdata`, we store the Legacy response at a

10 kHz resolution and average it as needed to match the data resolution. We also calculated and stored the MWAX response at 40 and 80 kHz resolutions using the code developed by Jones. If necessary in the future, the MWAX response for other frequency resolutions can also be derived and applied.

Cable length correction

The metafits “LENGTH” field contains the corrected electrical length of the cable to each antenna. That is, the cable length has been multiplied by a velocity factor accounting for the speed of light, and the length of the clock cable to the receiver has been subtracted. To account for the varying cable lengths between antennas (and thus the differing times at which signals are received), we apply a phase correction to each visibility for antennas j and k of $\exp(-2i\pi(l_k - l_j)f/c)$ where l_k and l_j are the corrected cable lengths for the k th and j th antennas, respectively, f is the frequency array, and c is the speed of light.

In addition to these corrections, the data is phased to the observation pointing center. That is, we apply a rotation to the phase of each visibility to account for the fact that the signal from a particular point on the sky reaches each antenna at a slightly different time. Then RFI flagging from SSINS is incorporated into the flag array. The flag array indicates any times, frequencies, antennas, and/or polarizations in the data to exclude from analysis. Data is then optionally averaged in time and/or frequency before being written along with the flag array, a sample array which counts the number of samples contained in each visibility, and relevant metadata to a UVFITS file for ingestion by the analysis pipeline.

8.3 FHD/ ϵ psilon analysis pipeline

8.3.1 FHD

After preprocessing, UVFITS files are ingested by Fast Holographic Deconvolution (FHD), our open-source calibration and imaging pipeline (Sullivan et al., 2012; Barry et al., 2019a). For MWA EoR data we run FHD in two stages: calibration and imaging. This is due

to differences in the beam kernel used when gridding visibilities to the uv -plane and/or degridting visibilities from the uv -plane. For the calibration stage, we use a beam which is accurate to our instrument when generating the model visibilities. Barry et al. (2019b) found that for power spectra analysis it was beneficial to limit the response from the chromatic sidelobes by using a tapered beam kernel. Following that work, we use a tapered kernel when gridding calibrated and model visibilities to the uv -plane before imaging. To allow for these different kernels, we run FHD in two stages, calibration and imaging.

FHD calibration

The core principle of calibration is solving for the linear time-dependent instrument response, or gain g , for each antenna. Our measured visibility V_{jk} for antennas j and k can be written in terms of the true sky signal u_{jk} for each frequency f as

$$V_{jk}(f) = g_j(f)g_k^*(f)u_{jk}(f) + n_{jk}(f), \quad (8.1)$$

where n_{jk} is a noise term. For MWA EoR calibration with FHD, we first simulate the true sky signal by generating a set of model visibilities m using $\sim 29,000$ sources from the GLEAM sky catalog (Hurley-Walker et al., 2017) supplemented with 9 additional bright sources as described by Byrne (Accessed June 4 2024). For each source with a true flux greater than 0.1 Jy, a Direct Fourier Transform (DFT) is performed to grid the source onto a uv -plane of $\frac{1}{2}\lambda$ resolution. This uv -plane is then multiplied by the instrumental beam for each antenna pair to simulate model visibilities for each polarization p , frequency f , and time t : $m_{jk,p}(f, t)$. For our analysis we use unpolarized sources only, ignoring contributions from diffuse emission (Byrne et al., 2021b) and only using baselines longer than 50λ for which the diffuse emission is sub-dominant. We simulate model visibilities for the instrumental XX and YY polarizations, ignoring the cross XY and YX polarizations. Assuming the noise is Gaussian and has the same amplitude across all baselines, we generate a set of ‘raw’ calibration gains for each frequency f by iteratively solving (Salvini and Wijnholds, 2014)

for each antenna j and each instrumental polarization p :

$$\chi_{j,p}^2 = \sum_{jk} |V_{jk,p}(f, t) - g_{j,p}^{raw}(f) g_{k,p}^{raw*}(f) m_{jk,p}(f, t)|^2. \quad (8.2)$$

Each iteration loops over frequency and then over antenna, obtaining a least-squares solution for each antenna at each frequency. Using an initial estimate for the per-frequency gain of the k th antenna, $g_{k,p}^{raw*}(f)$, we solve for the j th antenna gain $g_{j,p}^{raw}(f)$. Once a solution has been obtained for each antenna, the estimates of the gains are updated and convergence is tested for each frequency using the gains for all antennas. This process is repeated until either a convergence threshold of $1e - 7$ or the maximum number of iterations is reached. If the solution begins to diverge, iteration is stopped and the gains from the previous step are used. The per-frequency gains are complex, and before calibration a particular reference antenna is chosen to have zero phase. All other phases are solved relative to the reference antenna phase. By default, FHD chooses the second entry in its list of antennas as the reference antenna, which for this analysis is Tile012.

To better capture instrumental effects, we next perform a conditioning process of the raw per-frequency gains. We implement the autocorrelation ratio calibration process developed by Li et al. (2019) as follows. First, autocorrelations are averaged over the 2 minute observation. An ‘auto ratio’ is then computed for each antenna by dividing its averaged autocorrelation by the reference antenna’s averaged autocorrelation and taking the square root to give

$$A_{ratio} = \frac{\sqrt{\bar{A}_{j,p}(f)}}{\sqrt{\bar{A}_{ref,p}(f)}}. \quad (8.3)$$

For each antenna, the raw gain is divided by the respective auto ratio, removing antenna-dependent structures such as cable reflection ripples. This results in a modified gain g'

$$g'_{j,p}(f) = g_{j,p}^{raw}(f) * \frac{\sqrt{\bar{A}_{ref,p}(f)}}{\sqrt{\bar{A}_{j,p}(f)}}. \quad (8.4)$$

A ‘global bandpass’¹ is next obtained by computing means of g' such that any values over 2 sigma from the median are not included in the mean. First a mean is taken over frequency of each antenna’s amplitude; these means are divided out of the respective antennas’ modified gains. Then a mean is taken over these normalized modified antenna gain amplitudes at each frequency. This gives a bandpass for each instrumental polarization $b_p(f)$ such that

$$b_p(f) = \left\langle \frac{|g'_{j,p}(f)|}{\langle |g'_{j,p}(f)| \rangle_f} \right\rangle_j. \quad (8.5)$$

After dividing the modified gains by the bandpass, we fit a second-degree polynomial in amplitude and a first-degree polynomial in phase, as well as solve for a mode τ , amplitude c_c , and phase ϕ_c for the cable reflection. These three fits provide a set of 8 parameters for each antenna: $\{c_{2j}, c_{1j}, c_{0j}\}$ for the amplitude polynomial, $\{\phi_{1j}, \phi_{0j}\}$ for the phase polynomial, and $\{\tau_j, c_{c_j}, \phi_{c_j}\}$ for the cable reflection. The contribution of the cable reflections to the gain amplitudes is captured using autocorrelations, so we simply adjust the phase with the following:

$$\theta_{c_j} = \text{Im}(c_{c_j} * e^{i(-2\pi\tau_j f_{ind}/N_f + \phi_{c_j})}) \quad (8.6)$$

where f_{ind} is the frequency index and N_f is the number of frequencies. For an in-depth discussion of the fitting of the cable reflection parameters, please see Barry et al. 2019a and Li et al. 2019.

Finally, we use the autocorrelation ratios to capture antenna-dependent structures in the gain amplitudes. To produce each antenna’s gain, we multiply the corresponding auto ratio by the bandpass and the respective fitting terms. This gives the following gains which are

¹While this term described is referred to as a ‘global bandpass’ in the relevant FHD code and in other publications, it carries a different meaning than is commonly assumed. It does not imply that the antenna spectral structure can be described by a universal term; in fact, for the MWA we are not able to use a global term to describe the response of all antennas but must instead utilize the autocorrelation-derived term specific to each antenna in the calibration solutions. The ‘global bandpass’ we calculate with FHD is so called due to historical precedence unrelated to the typical usage of the term.

applied to the visibilities during the FHD imaging stage:

$$g_{j,p}(f) = \frac{\sqrt{A_{j,p}(f)}}{\sqrt{A_{ref,p}(f)}} * b_p(f) * [c_{2j}f_{ind}^2 + c_{1j}f_{ind} + c_{0j}] * e^{i(\phi_{1j}f_{ind} + \phi_{0j})} * e^{i\theta_{c_j}} \quad (8.7)$$

where f_{ind} is the frequency index. The terms in the above equation are illustrated in Figure 8.1, which shows the raw gain amplitude and phase along with the final gain amplitude and phase. The individual components of the final gain are plotted as well, showing how each contributes to capturing the structure of the individual antenna. To compare the final conditioned gain with the raw gain, we divide the final gain by the raw gain. We show the amplitude of this complex ratio in the bottom panel of the figure.

FHD imaging

After obtaining the final antenna gains, the FHD calibration run ends. We then run FHD again for imaging, transferring in the calibration solutions and uv -plane of model sources generated by the previous run. The calibration solutions are divided out of the visibilities and the calibrated visibilities are split into sets of interleaved even and odd time samples, with flagging consistent across both sets. We grid the calibrated visibility sets to the uv -plane using a modified beam kernel which tapers smoothly at the edges of the primary beam and reduces artifacts in the power spectra. Following Barry et al. (2019b), the transferred model uv -plane is used to construct model visibilities which are then also gridded using the modified beam kernel. Weights and ‘variance’ uv -planes are generated by gridding 1’s with the beam kernel and square of the beam kernel, respectively. The calibrated data (‘dirty’), model, weights, and variance uv -planes are then taken from uv -space to image space by a 2-dimensional FFT to slant orthographic sky coordinates. Interpolation to a Healpix grid (Gorski et al., 2005) results in cubes with coordinates (θ_x, θ_y, f) . The data can be easily integrated in the Healpix space. For our analysis of 2021 MWAX data we integrate cubes over each pointing for each night to perform power spectra quality assessment (see Section 11.3 for details), and surviving cubes are integrated for the final limit. These integrated Healpix cubes are the input product for ϵ psilon, the final stage of the analysis pipeline.




Figure 8.1: The gain solutions produced by FHD are illustrated for Tile022 in the XX polarization. The top panel shows the amplitude of the ‘raw’ per-frequency solution produced by the iterative minimization solving of (Equation (8.2)). Next are shown the three amplitude terms used in the final solution (Equation (8.7)): the auto ratio, computed from the square roots of time-averaged autocorrelations; the bandpass, computed from the raw gain amplitudes weighted by the auto ratios; and a polynomial fit. These three terms are multiplied to give the final gain amplitude as shown in the third panel. Comparison of this panel with the raw gain amplitude shown at the top illustrates the decrease in calibration noise achieved by the final gain solution. The next panel shows the phase of the raw gain solution. The phase terms derived from fitting a polynomial and a cable reflection, respectively, are shown in the panel below. These two terms are summed to give the final solution phase, shown next. Finally, the raw and final gain solutions are compared by taking a ratio of the complex gains and then plotting the amplitude of the ratio.

8.3.2 *epsilon*

To take our image cubes to power spectra we use our open-source software *epsilon* (error-propagated power spectrum with interleaved observed noise). This processing can be broken down into three stages: 2-dimensional DFT to $k_x k_y f$ space, 1-dimensional FT along frequency to k_z , and calculation of power spectra.

Healpix to $k_x k_y f$

The integrated Healpix cubes are ingested by *epsilon* and a 2-dimensional DFT is performed to go from image space to a uv space in which pixels are chosen to be of roughly the same scale as the beam kernel, preventing excess correlation between pixels. While in FHD the uv pixels are measured in wavelengths, here we do a transformation to physical coordinates, measuring the pixels in inverse co-moving megaparsecs. This remapping takes the coordinates from u, v

to k_x, k_y . The resultant dirty, model, residual, weights, and variance cubes are saved for each instrumental polarization and even/odd pair as intermediate outputs.

Frequency to k_z

Before the frequency Fourier transform, new products are created from the $k_x k_y f$ cubes. Weighted dirty and model cubes are generated by dividing by the weights cube. A residual cube is created by subtracting the weighted model cube from the weighted dirty cube. The original variance cube is divided by the original weights cube squared. A true variance cube is then created by multiplying the new variance cube by an estimate of the visibility variance as a function of frequency (Beardsley, Accessed June 4 2024).

With these new dirty, model, residual, and variance cubes, we make any desired adjustments in frequency such as slicing, averaging, and flagging. Then we subtract the mean from the data and apply a Blackman-Harris windowing function before taking the frequency transform.

Power spectra

Our science data product is the power spectrum, which provides a statistical measurement of the size and distribution of ionized regions during the EoR. The sampling pattern of our data is uneven and so we use a Lomb-Scargle periodogram (Lomb et al., 1976; Scargle et al., 1982; VanderPlas, 2018) to transform from frequency to k_z and compute the 3D cross-power spectrum. After the transform, the subtracted mean power is added back into the zeroth mode. For 1D power spectrum analysis, we take advantage of the statistical symmetry of the universe and boost sensitivity by performing a spherical average of the 3D power over k_x, k_y, k_z to calculate the observed power spectrum at each k mode. In addition to computing the 1D observed power for our scientific analysis, `epsilon` also provides useful diagnostic plots.

Diagnostic plots

Two particularly powerful diagnostic plots produced by `epsilon` are the 2D cross-power spectrum and the 2D power spectrum of the difference cubes. `epsilon` computes the 2D power spectrum by taking a cylindrical average over k_x and k_y , which results in a coordinate perpendicular to the line of sight k_{\perp} as a function of k_z . As k_z is parallel to the line of sight, in this space we typically refer to it as k_{\parallel} . In 2D power spectra we clearly see the behaviors of foreground power and instrumental systematic contamination as a function of the k_{\parallel} and k_{\perp} modes as we showed in Figure 2.1 and Figure 3.10.

In addition to the 2D power spectra of the cross-power, we also compute 2D power spectra of the difference cubes. The difference cubes are computed by subtracting the even and odd time-interleaved cubes, and their power provides a measure of the observed noise. We expect the observed noise to be cylindrically symmetric, so we use the 2D power spectrum to analyze the behavior of this noise.

Processing our ~ 5 TB of data through the preprocessing and analysis pipelines just described requires significant computation. Before proceeding with our power spectrum analysis, we briefly detour into a description of the compute systems we utilize for this work.

Chapter 9

COMPUTE SYSTEMS FOR EOR DATA ANALYSIS

Obtaining a limit with any MWA EoR dataset has particular compute requirements. We take hundreds of individual observations, typically 10-30 GB each depending on time and frequency resolution, through the preprocessing and analysis pipelines (Section 8.2, Section 8.3). Each pipeline stage has specific memory and central processing unit (CPU) needs, as well as necessary dependencies which must be installed. Since we process each observation individually, parallel compute resources are key. Rather than undertaking the cost and maintenance of an in-house compute system, our group instead leverages cloud computing to implement our data analysis pipelines.

9.1 Pipeline computing requirements

The preprocessing pipeline is implemented in python and so requires a python environment including specific packages. The analysis pipeline is written in IDL and, in addition to needing access to an IDL distribution and license, has dependencies which must be installed including Healpix, Imagemagick, and ghostscript. Both the preprocessing and analysis pipelines require users to clone a number of Git repositories and to keep their local copies of these software repositories up to date.

In the preprocessing pipeline, both SSINS and `pyuvdata` are expected to have RAM requirements of at least twice the incoming data size. Disk memory of roughly twice the incoming data size is also required to hold both the raw data and the output UVFITS file, in addition to SSINS output files. With our ~ 30 GB UVFITS file input for this MWAX analysis, the calibration stage of the analysis pipeline utilizes ~ 150 GB of RAM. FHD has the

capacity to perform its computation in parallel, and so processing time is dependent on CPU access. Given access to 12 CPUs on our cluster of high performance machines (specifications in Table 9.1), the calibration stage takes ~ 200 minutes per observation to complete. The imaging stage employs ~ 90 GB of RAM and takes ~ 580 minutes given 8 CPUs. The number of CPUs we allocate is largely dictated by the amount of storage required. The storage needs of the imaging stage are much higher than the calibration stage to accommodate output products such as Healpix, beam, and visibility files. We don't currently have access to RAM usage data for Healpix integration and ϵ psilon's 2-dimensional Fourier transforms; for these stages we allocate 12 CPUs/192 GiB and 8 CPUs/36 GiB, respectively. The final step of producing power spectra has low RAM demand and so, provided there is sufficient storage capacity for the relevant Healpix and $k_x k_y f$ -cubes, is often performed on personal machines to enable quick feedback implementation by the user.

The parallel computation we perform is conducive to a compute cluster. Typical cluster architecture includes a scheduler node from which jobs are submitted to a scheduling software—Slurm, in our case—as well as compute nodes where jobs are then executed. Slurm manages the allocation of jobs to nodes. This allows for the execution of hundreds of jobs in parallel without necessitating the user administer each job individually, thus providing an efficient use of time and resources.

9.2 Cloud computing

Our research group's strategy is to take advantage of cloud computing to create compute clusters on an as-needed basis rather than maintaining a long-term compute system in house. This has the advantage of only paying for high-powered compute resources during big data processing pushes. Additionally, we are able to design our clusters to meet our current RAM, disk space, and CPU needs. Historically, our pipeline was implemented through Amazon Web Services (AWS) cloud computing (Byrne and Jacobs, 2021). In this iteration of our data processing, users each generated a cluster for their personal usage. A custom virtual machine

image containing the pipeline requirements was used for the cluster nodes, so users did not have to install all requirements from scratch upon cluster creation.

While effective, there are particular disadvantages to this approach. Virtual machine images are somewhat brittle in that updates to a user's software or environment require creation of a new virtual image in order to propagate those changes to a new cluster. More importantly, the images are tied to specific versions of both the AWS software used to create clusters and the scheduling software implemented by the cluster. To use existing images, software is required to be pinned to old and sometimes outdated versions. Also, running multiple clusters simultaneously is not an efficient use of compute resources. With access to funding from Microsoft Azure, we had the opportunity to reimagine our compute cluster implementation.

In each year 2020-2023 we applied for and received Microsoft Azure computing grants through a collaboration between the University of Washington's eScience Institute, UW Research Computing, and Microsoft Azure. We were awarded \$27,000, \$15,142, \$15,000, and \$22,000 for each year, respectively. An estimated 50% of the grants have been spent on data storage, 35% on virtual compute machines, and 15% on bandwidth used to migrate data between storage and virtual machines.

Leveraging Azure's CycleCloud service, I designed and generated a customized Slurm compute cluster which allows access to multiple users. This allows for more efficient cluster management and means that updates and upgrades to the cluster benefit all users simultaneously. However, a disadvantage to this approach is that, due to the nature of the Azure computing grants, work is required to spin the cluster up and down for each funding cycle. In order to preserve the pipeline dependencies and user data between funding cycles, I created a Network File Share (NFS) which I mount on the scheduler node. The NFS hosts user home directories as well as all processing pipeline requirements. In between cycles the NFS as well

as data storage accounts are migrated to a paid Azure subscription where they persist until more compute grants are obtained. When restarting the cluster, we do not have to reinstall the pipeline from scratch. Additionally, we are able to update software without becoming limited by version conflicts.

Along with the NFS and data storage accounts, I created necessary cluster infrastructure. A virtual network was designed with specific subnets for the CycleCloud application and the cluster itself, respectively. I also created a managed identity which can be given access to data storage accounts and cluster nodes, allowing passage of data between these entities. Utilizing the managed identity means that data access no longer has to be managed by each individual user.

Our compute cluster consists of a scheduler node as well as two types of compute nodes. The specific virtual machine types used for each of these is determined by a cluster template and json file which can be updated at any point. The compute node partitions are designated as high performance and high throughput, respectively. For our high performance nodes we prioritize a higher RAM/CPU ratio, making this partition particularly useful for the preprocessing and FHD pipeline stages. The high throughput nodes are efficient for jobs with lower memory requirements such as the epsilon pipeline portion. Specifications for each machine type are shown in Table 9.1. We will note for future Azure cluster users that machine specifications are given in virtual CPUs; allocating 8 CPUs to a Slurm job will give the job access to 16 virtual CPUs.

Files and instructions for setting up the Azure cluster are located in an EoRImaging collaboration Git repository¹. The collaboration also maintains a repository² for Slurm submission scripts for all stages of the pipeline as well as wrappers which store settings for

¹https://Github.com/EoRImaging/azure_cluster_setup

²https://github.com/EoRImaging/pipeline_scripts

Partition	Machine type	Virtual CPUs	RAM (GiB)	Storage (GiB)
High performance	Standard_E48s_v3	48	384	768
High throughput	Standard_F32s_v2	32	64	256

Table 9.1: Specifications of virtual machines used as compute nodes in our customized Azure CycleCloud cluster.

various pipeline runs. Using Git repositories to host these pipeline components allows for consistency and collaboration between users, as well as preserving the history of changes.

We now proceed to interacting directly with the products of our Azure compute cluster, as we process data for making analysis decisions, evaluating data quality, and finally obtaining a limit with MWAX data in the following sections.

Chapter 10

FINE CHANNEL FLAGGING AND AVERAGING WITH MWAX DATA

In analysis of data from the Legacy correlator, typical practice has been to obtain data at a frequency resolution of 40 kHz from the correlator and then to average by a factor of 4 through the processing pipeline. Before averaging, two channels were flagged at each edge of each coarse band. For MWAX data analysis, the boundaries of frequency channels have changed as have the motivations for flagging specific channels. In this chapter we explore the best flagging and averaging choices for data from the new MWAX correlator.

10.1 Flagging and averaging with MWAX

To explore flagging and averaging options for MWAX data, we take data from a single night through preprocessing and FHD and integrate it together. This set consists of 42 observations taken across all four pointings. Rather than flagging and averaging within the pipeline, we keep the data at 40 kHz resolution and flag only the corrupted first fine channel of each coarse channel until we reach ϵ .

For this analysis, I upgraded ϵ to allow for flagging and averaging in frequency after the two-dimensional Fourier transform converting FHD Healpix cubes to $k_x k_y f$ cubes and before the frequency Fourier transform (Section 8.3). Historically, data has been averaged to 80 kHz in preprocessing before ingestion by FHD and then averaged to 160 kHz in the Healpix stage. We can now take data through preprocessing and FHD at 40 kHz resolution and implement frequency flagging and averaging in ϵ . This allows for experimenting with a number of various schemes without rerunning costly FHD stages.

Flagging and averaging schemes have different spectral point spread functions (PSF) which cause power to bleed both from the foregrounds into the EoR window as well as from the EoR window into the foregrounds. Since the foregrounds are several orders of magnitude brighter than the EoR signal, we are particularly concerned about bleed into the window. Also, our flagging schemes tend to be periodic at the coarse band frequency. This causes power to particularly bleed into the coarse band lines and nearby channels after the transform from frequency to k_z .

To analyze the results of our tests, we apply flagging and averaging schemes in ϵ psilon and then calculate power spectra. To examine the outputs, we essentially average over the range $10\lambda < k_{\perp} < 50\lambda$ in the 2D power spectrum to compute power as a function of k_{\parallel} and plot the result. A key feature in these plots are the coarse band lines, whose locations we demarcate with vertical grey lines. Our goal as we evaluate flagging and averaging schemes is to minimize power bleeding into these peaks and thus into the EoR window as a whole. We note that the y-axis on these plots is truncated and does not show the full power in the lowest k_{\parallel} bins, which reaches to $\sim 10^{14}$. Before implementing this analysis, we will explain some key differences between Legacy and MWAX fine channel averaging.

10.1.1 Fine channel averaging in MWAX

After channelization and before fine channel averaging, each coarse band of the Legacy correlator comprised 128 10 kHz channels. In the MWAX correlator, each coarse band contains 6,400 200 Hz channels. For both correlators, these channels are symmetric to the coarse band center. That is, a fine channel is centered at the center of the coarse band. This center channel contains the DC signal after the respective FFT stages. The centering means that at each edge of each coarse band is a half-size channel: 5 kHz for Legacy and 100 Hz for MWAX. For our typical case of averaging to 40 kHz fine channels, these two edge channels are combined within the channel at each coarse band's respective bottom edge, which we call the zeroth channel.

In the MWAX correlator, averaging is performed from the ‘center-out’ such that there is always a resultant averaged channel centered at the coarse band center frequency. We note that this particular channel would contain the DC bin, which is instead excluded from the average and the channel re-scaled accordingly. The center frequency of the zeroth channel aligns with the bottom coarse band edge, and in the case of 40 kHz fine channels the remaining 31 channels are symmetric around the center channel of the coarse band. When averaging to 40 kHz fine channels, the zeroth channel contains 20 kHz from the top of the band combined with 20 kHz from the bottom of the band. Since the zeroth channel includes information from both ends of the coarse band, we flag it in our analysis.

Legacy averaging was performed ‘bottom-up’. That is, channels were collected starting at the bottom of each coarse band and averaged in groups. Subsequent averaging stages in our pipeline have always also utilized a bottom-up averaging scheme. The center-out averaging performed by MWAX suggests that implementing a center-out scheme instead might be optimal, as we shall explore later in this chapter¹.

Finally, it is useful to note that the 40 kHz averaging schemes of Legacy and MWAX result in different frequency channel boundaries for their output data, which can impact the comparisons we make between these data sets.

10.1.2 *Flagging coarse band edge channels*

As we revealed in Figure 7.1, residual PFB structure is still evident at the coarse band edges in MWAX data. When we consider flagging schemes, we are comparing the impacts of two competing effects which cause bleed into the coarse band lines. The first of these is due to corrupted data at the coarse band edges. However, when we apply flagging at the coarse

¹A possibly far-fetched but interesting idea for future work is to explore requesting our data from the correlator as twenty-five 51.2 kHz channels per each coarse band. In MWAX, averaging schemes which result in an odd number of fine channels within each coarse band don’t combine data from each coarse band edge into the zeroth channel. This would allow us to recover the flagged zeroth channel and would open up new flagging and averaging options which might prove optimal.

band edges we create an uneven sampling pattern whose PSF will also cause foreground power to bleed into the coarse band lines. Our goal is to determine how much flagging will be optimal for minimizing bleed into the EoR window. This will tell us which effect is stronger: the corruption of data in a particular channel or the uneven sampling PSF due to flagging that channel. In Legacy data, the artifacts present in data at the coarse band edges were strong enough that flagging was required. While clearly some artifacts remain in MWAX data, it may be that they are small enough that we are able to recover channels at the coarse band edges.

As we develop our flagging schemes, one consideration is the symmetry of MWAX output frequency channels as compared to Legacy. Each MWAX coarse band contains the flagged zeroth channel along with 31 remaining channels which are symmetric about the center of the coarse band. Thus PFB artifacts and aliasing at the coarse band edges will also be symmetric within those 31 channels. This suggests the possibility that instead of flagging the same number of channels at each edge of each coarse band, it might be optimal in MWAX data to flag, say, two channels at each edge in addition to the corrupted zeroth channel.

We implement various flagging schemes in ϵ psilon before averaging from 40 kHz to 160 kHz, forcing the frequency array to be evenly spaced regardless of flagging. This isolates flagging effects from any PSF due to changes in the output frequency array. Two types of flagging schemes are implemented—symmetric and asymmetric—and tested with flagging incrementally deeper into the coarse channel.

Examples of symmetric and asymmetric flagging schemes are illustrated in Figure 10.1. As shown, symmetric flagging is symmetric with regards to the 32 channels within each coarse channel. That is, flagging two fine channels at each edge of each coarse band is a symmetric flagging scheme, while flagging two fine channels at each edge in addition to the flagged zeroth channel is an asymmetric flagging scheme. In Figure 10.1, fine frequency channels

are represented by square cells. Delineations separating coarse bands are marked by bold lines between cells such that 32 cells representing 32 40 kHz channels are within the coarse band. Cells which are flagged are colored grey, and averaging is represented by coloring cells in alternate groups of green and blue. These groups illustrate the fine channels comprising each frequency channel after averaging.

In Figure 10.2 we show a comparison of four of the flagging schemes shown in Figure 10.1: one channel at each edge in addition to the zeroth channel, two channels at each edge of each coarse channel, two channels at each edge in addition to the zeroth channel, and three channels at each edge of each coarse channel. The last of these is the maximum flagging we can apply without resulting in fully flagged channels after averaging. We did extensive exploration with fully flagged channels which revealed they always result in drastic foreground bleed. We do not present them here but rather focus on flagging schemes with more potential.

For each of the flagging schemes, we plot power as a function k_{\parallel} for each instrumental polarization XX and YY, and for the dirty (data) and model. In Figure 10.2, the coarse band line locations are marked by vertical grey lines, and peaks in power due to periodic coarse band structure are evident. At times a dip appears in the center of the coarse band line peaks; this is a result of subtracting the mean before taking the frequency Fourier transform.

We see that as we increase the number of flagged channels, power in the model coarse band lines increases as well. This is expected as the model does not have any knowledge of data corruption at the edge channels, and so simply responds to the PSF of increased periodic flagging. However, in the dirty coarse band lines, power drops as flagging increases. This indicates that residual nonlinearities after dividing out the PFB response outweigh the effects due to periodic flagging. Notably, no benefit is seen from using asymmetric flagging; the best scheme is to flag as many edge channels as is possible without resulting in fully flagged averaged channels. Unfortunately, we are not yet able to recover coarse band edge

channels with MWAX data. As digital receiver upgrades are implemented (Chapter 6), it is recommended to repeat this analysis.

10.1.3 Frequency averaging

To obtain a power spectrum EoR measurement, we take a Lomb-Scargle transform across the frequency axis (Section 8.3). Periodic flagging in frequency creates regular gaps in the frequency array, and these gaps have a PSF which bleeds power into the EoR window. Averaging in frequency can mitigate this by combining flagged and unflagged channels into a single averaged channel to result in a frequency array in which no frequency channel is without data. However, this also results in an uneven spacing of the frequency array—that is, the centers of the frequency bins are not equidistant from each other due to the locations of the unflagged data. This uneven spacing has its own PSF which causes power bleed, albeit less power than fully flagged frequency channels. With Legacy data, this effect was mitigated by forcing the frequency centers of averaged channels to be evenly spaced, essentially ignoring the fact that some averaged channels contained flagged data. However, since the incoming signal is not actually flat, lying about the frequency channel centers smears power into neighboring bins. We wish to determine whether our best practice is to correctly label the frequency centers of averaged bins, resulting in an unevenly spaced frequency array, or to continue lying about the channel centers in favor of evenly spaced frequencies.




Figure 10.2: We compare the impacts of flagging channels at each edge of each coarse channel. The zeroth channel contains data from both ends of the coarse channel and thus is always flagged. We compare four flagging schemes: one channel at each edge in addition to the zeroth channel, two channels at each edge of each coarse channel, two channels at each edge in addition to the zeroth channel, and three channels at each edge of each coarse channel. The last of these is the maximum flagging we can apply without resulting in fully flagged channels after averaging. Flagged channels are represented by grey cells, and blue and green cells indicate groups of channels which are averaged together, in this case, by a factor of 4. The effects of flagging are strongly seen in the peaks of power due to structure at the coarse band cadence, whose locations are demarcated by vertical grey lines. In the model (blue and purple lines, for the XX and YY instrumental polarizations, respectively) power increases with increased flagging due to unevenness of sampling. In the dirty (data, shown as orange and red lines), however, increased flagging decreases power, indicating corruption in coarse band edge channels due to mis-subtraction of the polyphase filterbank bandpass. To better show effects in the coarse band lines, the y-axis has been truncated. Power in the lowest k_{\parallel} bins reaches $\sim 10^{14}$.

Frequency spacing

We compare the effects of unevenly spaced frequencies versus mislabeling the frequency centers of averaged channels in Figure 10.3. For this comparison we flag three fine channels at each edge of each coarse channel before averaging by a factor of 4. As in the flagging analysis, we look at power integrated over a range of k_{\perp} as a function of k_{\parallel} for each instrumental polarization XX and YY, for both the dirty (data) and model. In the upper panel, the true frequencies are used for the averaged channels. This results in uneven frequency spacing with larger gaps between channels containing flagged data. In the lower panel, the frequencies of the averaged channels are forced to be evenly spaced. It is evident that the foreground

power (which in this case has amplitude $\sim 10^{14}$ in the lowest k_{\parallel} bin) is bleeding excessively into the coarse band lines when the true, uneven frequencies are used. Power in the coarse band lines in the bottom panel is ~ 2 orders of magnitude lower than in the top panel. Thus our optimal choice is to sacrifice accuracy in our frequency array in favor of a cleaner EoR window by using evenly spaced frequencies even though they do not fully describe the distribution of data within the averaged channels.

Averaging schemes

We next explore center-out averaging schemes in which an averaged channel is centered on the center of the coarse band's 31 symmetric channels. First, we bin channels in groups of three. As seen in the top panel of Figure 10.4, this results in bins of two channels at each coarse band edge. With this scheme only one channel can be flagged at each coarse band edge (in addition to the zeroth channel) without resulting in gaps in the frequency array. We see much higher power in the coarse band lines than we see from the scheme of flagging three channels at each coarse band edge and averaging by a factor of 4 with frequencies forced to be evenly spaced (shown in the bottom panel). This is unsurprising considering that we showed in Figure 10.2 that flagging more edge channels decreases power in the coarse band lines. In the second panel we implement binning channels in groups of five with the maximum possible flagging of two channels at each coarse band edge in addition to the zeroth channel. This scheme shows a drop in power as compared to binning by three, but is still higher power than seen in the bottom panel. We next modify binning by five such that averaged channels span the group of flagged channels at each coarse band edge, as shown in the third panel. This has the benefit of decreasing the discrepancy between frequency channel spacing after averaging, and the drop in coarse band line power reflects this. For this scheme, we additionally truncate the averaged channels at the very beginning and end of the entire frequency band. None of these schemes taking advantage of the possibility of center-out averaging with MWAX data perform better than the scheme of averaging by a factor of four from the bottom-up, likely due to the fact that bottom-up averaging is conducive to deeper flagging in the coarse band

edges.

In order to both execute maximum flagging and keep frequencies evenly spaced, we implement a final averaging scheme in which we bin eight frequency channels, shown in the top panel of Figure 10.5. Instead of averaging from the bottom-up, we center averaged channels across the coarse band edges. This results in naturally evenly spaced frequencies when we drop the very first and last four fine channels of the frequency band before averaging. This scheme does show an improvement over bottom-up averaging by a factor of 4 shown in the bottom panel, though averaging by 8 restricts k_{\parallel} to half its typical range. However, since sensitivity to the EoR signal increases with decreasing k_{\parallel} , losing the extent in k_{\parallel} might not be detrimental to our science.

A final concern we address is our choice of sampling in k_{\parallel} in the frequency Fourier transform. Sampling k_{\parallel} inappropriately runs the risk of missing key features in the power spectrum. To check this, we sample at 100 times our usual rate and compare the features with those seen at our typical sampling pattern. In Figure 10.6 we see that the features are gross enough that small changes to our choice of k_{\parallel} sampling do not run the risk of missing important structures.

Our flagging and averaging exploration reveals that an optimal scheme will flag as many edge channels as is possible and will keep averaged frequency channel centers evenly spaced. For our analysis we flag three channels at each edge of each coarse channel. Binning by eight such that coarse band edges are spanned has an improvement in lowering coarse band peak power but restricts the range of k_{\parallel} . Further explorations of binning by 8 could determine the usefulness of this scheme.




Figure 10.3: Effects due to forcing an evenly spaced frequency array (bottom) versus using the true uneven frequencies after averaging (top). As indicated by the colored cells, in both cases three channels are flagged at each edge of each coarse channel before averaging by a factor of 4. In the upper panel, the final frequency of each averaged bin is centered on the group of blue or green cells contributing the the bin. In the lower panel, the final frequency of each averaged bin is centered on each group of four cells, regardless of flagging. Mislabelling the averaged bins in this way results in a smearing of power after the frequency Fourier transform. However, it is clear in the bottom panel that power in the coarse band peaks (whose locations are demarcated by grey vertical lines) is lower when output frequencies are evenly spaced, while using the true uneven frequency spacing in the top panel results in an increase in power.




Figure 10.4: The typical practice of averaging in frequency by a factor of 4 from the ‘bottom-up’ of each coarse channel, in which channels are binned starting from the lowest frequency channel in each coarse band (bottom panel), is compared with ‘center-out’ averaging schemes which center an averaged channel in the symmetric 31 channels remaining in each coarse band after the zeroth is flagged. The top panel shows binning by 3 from the center-out. As illustrated by the grey cells, only one additional channel is flagged at each coarse band edge, as the frequency gaps due to flagging deeper result in excess power in the spectrum. This results in a single channel binned at each edge of each coarse channel, as seen in the blue and green cells. The peaks in power due to coarse band structure, demarcated by the vertical grey lines, are much higher with this scheme as compared to bottom panel. This is likely largely due to the decrease in flagging (Figure 10.2). In the next two plots we bin by 5 from the center-out, which allows for more flagging at the coarse band edges. In the third panel from the top, we center the averaged bins on the coarse band edges (indicated by the bold lines between boxes) which decreases the discrepancy in the frequency spacing of averaged bins as compared to the panel above. This results in a drop in power particularly evident in the coarse band peaks at higher k_{\parallel} . None of the center-out averaging schemes outperform the bottom-up scheme in the bottom panel, likely due to less flagging at each coarse channel edge as well as the uneven frequency spacing after averaging.




Figure 10.5: Averaging by a factor of 8 (upper panel) is compared with the typical Legacy scheme of averaging by a factor of four (lower panel). As illustrated by the grey cells, in both cases three fine channels are flagged at each edge of each coarse channel (coarse channel edges are indicated by the bold lines between flagged channels). The green and blue cells in the top panel shows that channels are binned to be centered on coarse channel edges, resulting in naturally evenly spaced frequency channels after averaging. Examination of the peaks in power due to coarse band features (demarcated by the vertical grey lines) shows a drop in power at the peaks for both the dirty (data) and model power when averaging by a factor of 8 (top panel). Significant improvement in the model power is likely due to the true even spacing of the averaged power. However, averaging by a factor of 8 cuts the extent in k_{\parallel} in half, suggesting that it would be advantageous to examine both averaging by 4 and by 8 in a future analysis.




Figure 10.6: Choice of k_{\parallel} values to sample after the frequency DFT is illustrated. A set of data is averaged by a factor of four and the frequencies of averaged channels forced to be evenly spaced regardless of the channels flagged at each coarse band edge (flagged channels are represented by grey boxes, and coarse channel edges demarcated by the bold lines between boxes). The bottom panel shows the default k_{\parallel} sampling after the frequency Fourier transform. The upper panel samples at 100 times the density of the lower panel, and shows that features in the power spectra are correctly represented in the more coarse sampling of the lower panel. This alleviates concerns that our choice of k_{\parallel} values might be causing us to miss key features in the power spectra.

Chapter 11

MWAX DATA QUALITY ANALYSIS

As we take our dataset through the analysis pipeline, we perform assessments at multiple pipeline stages to identify and exclude any data with quality issues. We first use the open-source software package Sky-Subtracted Incoherent Noise Spectra (SSINS) (Wilensky et al., 2019) developed by the UW team to identify radio frequency interference (RFI) from sources such as digital television (DTV), satellites, or even the telescope itself. Within each 2 minute observation SSINS identifies any timestamps with RFI and then flags every frequency for those timestamps. This produces a flag array for each observation indicating which timestamps to include or exclude from our data analysis.

We then use `pyuvdata` to ingest the raw correlator `gpubox` files and apply corrections as described in (Section 8.2). This process writes a flag array in which any antennas identified as problematic in the observation metadata file are excluded. We import the flag array from SSINS, and `pyuvdata` combines the flags into a single array before writing the visibilities to a UVFITS file. Next, the FHD calibration algorithm described in Section 8.3.1 excludes antennas from further analysis based on outlying behavior in the amplitude and phase of their calibrations solutions. At this stage we also develop a new metric to analyze the quality and consistency of calibration for each observation.

Finally, we integrate data over each telescope pointing of each night and compute 2-dimensional power spectra for each integrated set. We examine these spectra for residual RFI as described by Wilensky et al. (2023), as well as for overall data quality. All the observations in any integration exhibiting residual RFI are excluded from the final power spectrum analysis.

These flagging stages identify times contaminated with radio frequency interference, data from antennas which are not functioning correctly, and integrations of observations with residual RFI or overall data quality issues. Excluding this corrupted or contaminated data is crucial for making a measurement of the EoR signal. We will discuss a few of these flagging stages in more detail below: the RFI flagging by SSINS, our new calibration metric, and power spectra analysis of integrated pointings.

11.1 *The SSINS of MWAX*

SSINS detects temporally variable RFI contamination by looking for time dependent emission from sources with known frequency fingerprints. SSINS uses `pyuvdata` to ingest the raw MWA correlator data and then removes signal from the sky and from other time-stable sources by subtracting adjacent time samples of visibilities in each 2 minute observation. This creates a sky-subtracted spectrum sensitive to RFI that is variable in time from sources such as satellite signals, DTV transmissions reflecting off of airplanes, or flickering bright emitters (Wilensky et al., 2019). The time-subtracted complex crosscorrelation visibilities are then incoherently averaged to increase signal to noise, creating a single spectrum for each polarization of the observation. Subtracting the mean across time for each frequency channel normalizes and standardizes the spectrum across the frequency band, allowing for statistical detection of any RFI signals.

An iterative detection process then looks for specific spectral RFI patterns and flags the brightest RFI found above a set significance threshold, boosting sensitivity to any remaining RFI. This flagging process continues until no signals remain above the threshold. Between iterations, if any frequency channel is found to be contaminated for more than 60% of the observation timestamps, the entire observation will be flagged. Any timestamp with RFI detected in one instrumental polarization will be flagged in all polarizations. An optional step in the iterations is to flag every frequency for any time at which RFI is detected. The point spread function due to flagging some frequency channels but not others can result in

foreground power bleeding into the EoR window (Wilensky et al., 2022). For our analysis we implement the extension of the flags across the frequency band, which allows us to avoid the effects from uneven spectral flagging.

By default SSINS will detect two types of RFI: broadband streaks which span the entire frequency bandwidth and narrowband emission in single frequency channels. Custom settings can be used to detect RFI that spans multiple frequencies by running SSINS with a dictionary of RFI shapes. For MWA data taken over the frequency range 167-197 MHz, we historically use a dictionary containing four digital television (DTV) bands which each span 7 MHz.

Wilensky et al. (2020) found that MWA observations in which SSINS detected DTV contamination were likely to contain ultrafaint RFI below SSINS’s detection level. This was not the case for observations with broadband RFI, which Wilensky showed are free from RFI contamination outside timestamps where the broadband signal was detected. Wilensky’s analysis did not include enough observations with narrowband emission to conclusively determine whether or not residual RFI was likely to be lurking after the narrowband signal was flagged. Recent analyses of MWA EoR data (Barry et al., 2019b; Wilensky et al., 2023) have conservatively excluded all observations in which any RFI was detected by SSINS, calculating the power spectrum limit using only observations without any RFI detections.

We ran SSINS on the 2021 MWAX data set with the same settings used for Wilensky’s analysis of 2014 Phase I MWA data (Wilensky et al., 2023). Of the 319 observations in our data set, only 56 had no timestamps flagged by SSINS. This is an unprecedented amount of data loss; in Wilensky’s initial set of 1285 observations, 591 observations had no RFI detected by SSINS¹. Due to the likely presence of residual RFI, for our analysis we knew we wanted to discard the entire observation if any DTV was detected by SSINS. If, in an attempt to recover data in our set, we retained observations with flagged broadband streaks or narrowband

¹Some of these observations were identified by Wilensky as having RFI contamination below SSINS’s detection level

emission, the number of observations we could keep increases from 56 to 157. Though this is an improvement, over half the data set is still lost, and we would prefer to recover more observations if possible. Also, while we know the observations with flagged broadband RFI are not likely to contain residual contamination outside of the flagged timestamps, it was not clear whether or not this is the case with narrowband RFI. Recovering data by including observations with narrowband signals runs the risk of introducing contamination into our final power spectrum analysis. We decided to take a closer look at what SSINS was detecting in the data to understand and evaluate the RFI in our data set.

11.1.1 A SSINS rampage

To identify patterns in SSINS flagging of this data set, we held a ‘data rampage’: studying physical images of the mean-subtracted spectra, both with and without SSINS flagging applied, for all 319 observations. See Figure 11.1 for a photo of our SSINS data rampage. Examining the full set of data in this way, with a team of people making observations and identifying patterns, is a powerful tool for developing a rich understanding of what is actually happening and why.

When examining the SSINS spectra we frequently saw RFI contamination in the 175.00 MHz frequency channel. This emission is likely from Starlink satellites and has been identified by other telescopes including LOFAR (Vruno et al., 2023). In April 2022 LOFAR observed Starlink satellites emitting a signal within a frequency channel spanning 12.21 kHz from 174.994 to 175.006 MHz. A inspection of MWA EoR 2021-2024 data by Dev Null revealed contamination from this Starlink emission in almost every observation (D. Null, personal communication, February 1, 2024). Due to the persistence of this Starlink RFI we flag the 175.00 MHz channel in all observations of our data set, regardless of SSINS detection.

Our data rampage also revealed three patterns in the MWAX mean-subtracted spectra which are routinely miscategorized as DTV: a time-dependent ‘breathing’ of the instrument gains, dropped correlator packets, and RFI that is actually broadband. SSINS spectra of these three effects are shown in Figure 11.2, and we discuss their true natures and our mitigation strategies below.

In the MWA, each receiver is housed in a box which is periodically cooled by a refrigerator system. The cycling on and off of this cooling process is apparent in small fluctuations of the amplifier gain in the analog signal conditioning board within the receiver. When SSINS averages the baselines together, this ‘breathing’ of the gains can show up in the mean-subtracted spectra as a smooth change in amplitude over time in which the spectra transitions from primarily negative to primarily positive values or vice versa (shown in upper panel of Figure 11.2). The more extreme positive or negative values may rise over the detection threshold of SSINS, leading to flagging and miscategorization of this effect as broadband streaks or, importantly, as DTV. This amplitude change does not need to be flagged out of the data, so it is appropriate to remove it from the spectra by having SSINS fit and subtract a second-order polynomial rather than calculating and subtracting the mean. It should be noted, however, that a polynomial is fit for each individual frequency channel. This can have a detrimental impact on detecting narrowband RFI occupying a single channel, and we shall return to this point later in our discussion.

In several observations we see a negative block in the center third of the band for a single time stamp, as shown in the center panel of Figure 11.2. This feature, which we had not seen in SSINS analyses of data from the Legacy correlator, indicates packet loss in the MWAX correlator. Each digital receiver sends data from 24 selected coarse channels to the correlator, split into groups of 8 coarse channels which are each sent over 3 separate optical fibers to the correlator. Packet loss from one of these groups results in a drop in power in the corresponding third of the frequency band. SSINS does not distinguish between



Figure 11.1: The UW EoR group held a data rampage to inspect hundreds of SSINS spectra. A rampage allows us to identify and characterize patterns in our data. During this data rampage, we discovered emission likely from Starlink satellites as well as miscategorization of RFI by SSINS (shown in Figure 11.2) and developed mitigation strategies. Photo taken by Will Sutherland.

positive or negative deviations from the mean, instead taking an absolute value and flagging if that value is above the desired threshold. In order to allow differentiation of negative RFI patterns and detect packet loss, we upgraded SSINS to handle negative detection thresholds. We also added a shape for the packet loss to the dictionary which is only flagged if the mean-subtracted spectrum is negative over the central 10.28 MHz of the frequency band.

Finally, we see instances of broadband streaks detected as DTV. The bottom panel of Figure 11.2 shows a true broadband streak which extends below the DTV bands but has been labelled by SSINS as DTV. In order to prevent unnecessary flagging of breathing, our standard settings for running SSINS on MWA data have included setting the broadband streak detection threshold at 10σ , twice the value of the DTV thresholds. This has the unintended consequence of sometimes mislabelling true broadband emission as DTV. To mitigate this, we use the second-order polynomial subtraction to avoid flagging breathing and set the broadband streak detection threshold to the same 5σ value of the other RFI category thresholds to more accurately distinguish broadband RFI from DTV. Of course, we do not wish to lose sensitivity to any RFI, so a careful examination of SSINS flagging with these settings was required to ensure the desired results.

After careful validation we ran SSINS with the updated settings of second order polynomial subtraction, negative packet loss detection, and adjusted broadband streak sensitivity. A comparison with the standard settings is shown in Figure 11.3. The RFI categories described in this plot are not mutually exclusive; an observation can be contaminated with more than one type of RFI and thus be counted in multiple categories in the plot. When using the updated settings the number of observations with DTV detections decreased from 162 to 40. Observations with broadband detections also dropped from 119 to 92. Our data inspection confirms that these drops are due to reducing unnecessary flagging of breathing and to correctly categorizing broadband streaks, DTV, and packet loss. However, decreasing SSINS's sensitivity to breathing had an unintended consequence. The second order




Figure 11.2: SSINS computes sky-subtracted incoherently-averaged noise spectra for RFI detection. Here we show three spectra with phenomena which SSINS incorrectly categorized as digital television (DTV). The topmost panel shows ‘breathing’ of the analog amplifier gain as the cooling system in the receiver boxes cycles on and off. In SSINS spectra, the breathing results in a shift between overall negative pixels and overall positive pixels over the time of the observation. In the center panel, we see evidence of dropped correlator packets. Within the receiver the frequency band is divided into three sections and sent over three optical fibers to the correlator. If any of these packets are dropped, it will result in an anomalous drop in power (blue on this plot), as we see in the center third of the frequency band. The lowest panel shows broadband RFI as a red streak extending across the entire frequency band. Due to SSINS sensitivity settings, this RFI was labelled as DTV though it extends below the DTV bands.

polynomial subtraction is computed separately for each individual frequency channel. If a particular frequency channel contains narrowband RFI, the polynomial can fit and remove the RFI signal, preventing detection. This resulted in an erroneous drop in narrowband RFI detection when using the updated settings. While the number of observations without any RFI detections increased from 56 to 153 with the updated settings, 13 of these observations actually contained narrowband RFI. Within observations in which other RFI was present, narrowband detections also decreased. For our final flagging scheme we ran SSINS once with the original mean subtraction for narrowband detection and once with second order polynomial subtraction to optimize detection of the other types of RFI.

11.1.2 Data selection from SSINS analysis

With RFI flagging completed we make cuts to form our final data set, as shown in Figure 11.4. From our original set of 319 observations, 29 are highly contaminated. That is, RFI was present in at least one frequency channel for more than 60% of the observation time. We discard these from the data set. We then cut 37 observations in which DTV was detected from our analysis data set due to the likelihood of contamination by residual RFI. These cuts are shown in the white bars of the figure.

In the mean-subtraction SSINS run, narrowband emission was detected in 42 of the remaining observations, shown as purple bars. We split the narrowband observations into two sets. If narrowband emission was detected in a particular frequency channel for more than one timestamp over the entire dataset, we classify it as ‘recurring’ narrowband emission. If SSINS detected a single hot pixel at a unique frequency for a single timestamp in our entire data set, we categorize it as ‘random’ narrowband. Most RFI detected by SSINS appears repeatedly at particular frequency channels or bands as would be expected from frequency stable emitters. It is not clear if the random narrowband emission flagged by SSINS has a physical source or is the result of SSINS catching uncontaminated pixels just above the 5σ flagging cutoff. We process observations with recurring and random narrowband flagging




Figure 11.3: We ran SSINS on 319 MWAX observations and compared RFI detection between two versions of settings. The left panel shows observation counts when running SSINS using standard settings employed in the analysis of 2014 MWA data by Wilensky et al. (2023), while the right panel shows counts using updated settings designed to improve RFI categorization in the 2021 MWA RFI environment. An observation may contain more than one type of RFI and so be counted more than once. Using updated settings, the number of observations with DTV detections drop from 162 to 40. Data inspection confirms that this drop is due to improved RFI categorization. With the updated settings we see a drop in observations with narrowband RFI detections. This drop is erroneous; the updated settings can fit out RFI which is isolated to a single frequency channel. With the updated settings the number of observations with no RFI detections increases from 56 to 153 but some of these observations actually contain narrowband RFI as identified in the standard run. Correcting the miscategorization of ‘breathing’ (Figure 11.2) as broadband RFI caused the number of observations with broadband detections to drop from 119 to 92. Packet loss, which was not detectable with SSINS’s standard settings, was discovered in 15 observations. For our power spectrum analysis we run SSINS twice in order to optimize for both narrowband detection and improved RFI categorization.

separately for further analysis in Section 11.3.

In the remaining 211 observations, shown as blue bars in Figure 11.4, broadband streaks were detected and flagged in 58 observations. Since we do not expect residual undetected RFI to be present in observations with broadband contamination, we retain these observations for our final limit analysis. Correlator packet loss will only affect the timestamps of dropped packets and does not result in residual RFI in other timestamps. Thirteen observations had packet loss detection and flagging, and we include these in our limit data set. One of the observations with packet loss also had broadband emission, and is counted twice in the plot. Two observations, one of which has broadband RFI and one which is not shown in Figure 11.4, were identified in the polynomial-subtraction SSINS run as each having a single timestamp of random narrowband RFI. We include these observations in our final limit dataset because close inspection suggests this flagging is a product of the second order frequency fit rather than true narrowband contamination.

After these cuts we retain 211 2-minute observations from our original set of 319 observations. The first two seconds of each observation are known to be corrupted and are dropped before SSINS flagging, leaving 118 seconds of data per observation. After SSINS flagging, in our 211 observations we are left with ~ 6.537 hours of the initial 10.633 hour dataset.

Our 2021 MWAX data is taken over 4 pointings, as described in Table 8.1. To determine if the RFI detected by SSINS has any dependence on pointing or on the observing night, in Figure 11.5 we histogram three categories of observations for each pointing of each of the 13 nights of observation. Data retained after the SSINS analysis, which includes observations without any SSINS flagging as well as those with flagged broadband RFI or packet loss, is indicated by blue. Observations with narrowband detections from the mean-subtracted SSINS run are colored purple. While a few pointings have more narrowband detections than others, the narrowband emission overall appears to be evenly distributed throughout the




Figure 11.4: Based on radio frequency interference detection by SSINS, we make cuts on our data set for power spectrum analysis. We first discard observations which were highly contaminated and then those in which SSINS detected digital television (DTV). These are represented by the white bars. Observations with narrowband RFI are colored purple and split into two sets based on whether the detection occurred in a frequency channel which is contaminated in multiple timestamps (‘recurring’) or a channel which is only contaminated for a single timestamp in the entire data set (‘random’). Two of the recurring narrowband observations also contain random narrowband emission. For further analysis we retain 211 observations which are colored blue and include observations in which no RFI was detected as well as those in which broadband RFI or correlator packet loss was detected and flagged. Within this set, 51 observations have broadband emission, 12 have packet loss, and 1 has both. Two observations, one with broadband emission and one which is not included in this plot, do have random RFI but close inspection suggests it is an artifact of the SSINS second-order polynomial fit.

observing season, with no dependence on pointing or night.

White represents observations which we cut due to being fully flagged by SSINS or to containing DTV contamination. For these observations we see no RFI dependence on pointing in Figure 11.5. However, on December 4 we lose the entire Plus Two pointing, and on the subsequent observing night, December 6, almost all observations are excluded by this cut. One observation from the December 6 Plus One pointing is included in the narrowband analysis and the others are discarded completely. Visual inspection of SSINS spectra from December 4 and December 6 observations reveals DTV as well as strange streaking behavior. Future work is recommended to characterize these RFI sources. The final night of our observing season, December 10, only contributes a single observation for further analysis as the remaining observations contain narrowband emission. It is unclear why the data quality is apparently degrading at the end of the observing season.

We will continue the hunt for residual RFI in the power spectra space. First, we pause at the calibration stage and use a new metric to assess the quality and stability of calibration between observations.

11.2 Calibration polynomial fitting parameter analysis

Before calibrating the entire set of 2021 MWAX observations, we calibrated the observations from November 20 and November 22. An inspection of the amplitudes of the calibration solutions revealed a persistent RFI spike in antenna HexS5 (shown in the autocorrelation in Figure 7.7). We exclude this antenna from our power spectrum limit analysis.

As described in Section 8.3.1, during calibration we fit a second-order polynomial in amplitude and a first-order polynomial in phase to the scaled raw gains. These fits provide 5 parameters for each tile: $\{c_{2_j}, c_{1_j}, c_{0_j}\}$ for the amplitude polynomial and $\{\phi_{1_j}, \phi_{0_j}\}$ for the phase polynomial. The zeroth order terms, c_{0_j} and ϕ_{0_j} , describe an overall offset in the

amplitude and phase, respectively. The higher order terms describe the slope and curvature of the antenna response as a function of frequency.




Figure 11.5: We plot counts of the observations shown in Figure 11.4 for each pointing and each night of observation. This data set was taken late in the typical EoR observing season and as the season progresses, fewer observations are taken at earlier pointings. The observations we retain for further analysis, indicated by blue bars, comprise observations free from RFI as well as those with broadband streak or packet loss contamination. Purple bars demarcate observations with either recurring or random narrowband RFI. Observations which have been discarded due to being fully flagged or to containing DTV contamination are indicated by the white bars. We see no trends by pointing in the RFI indicated by purple and white boxes. However, the entire December 4 Plus Two pointing as well as almost every observation of December 6 are discarded, though one December 6 observation is included in the random narrowband set. Inspection of the SSINS spectra from these excluded observations reveals unusual RFI shapes from unidentified sources.

We created a new metric by computing the mean and standard deviation of the polynomial fit coefficients across all unflagged antennas over an observation to obtain a single mean and standard deviation for each parameter per observation. These values are calculated for all the observations we retained after the SSINS flagging analysis as well as the observations with narrowband RFI.

We histogrammed the computed means and standard deviations and inspected the plots for outlier observations. In the histograms, observations retained after the SSINS analysis are blue while those in the narrowband sets are purple. Grey dashed lines are plotted at $\pm 2\sigma$ for each histogram.

The means of the amplitude coefficients, shown in Figure 11.6, have no obvious outliers. While all parameters have some values that fall outside of $\pm 2\sigma$, they are still clustered with the main distribution. The means of the phase coefficients in Figure 11.7 show a similar

behavior for the zeroth order parameter while the first order parameter is more sharply distributed with a clear set of outlying observations to the right. These outliers comprise all the observations from December 8 and December 10.

While the standard deviations of the first and second order amplitude coefficients are distributed similarly to the means, the zeroth order parameters exhibit a different behavior as we show in Figure 11.8. A tail of observations extends to the left of a sharply clustered distribution. The observations with standard deviations below -2σ include all of December 8 as well as eight observations from December 4 and the single observation from December 6 which survived the SSINS flagging cuts. The standard deviations of the phase coefficients exhibit similar behavior as the means, with the same dozen observations lying outside of the sharply clustered distribution of the standard deviations of the first order parameters.

Our new calibration fit coefficient metrics identify 4 pointings containing unflagged data with anomalous behavior: December 4 Plus One, December 8 Plus One, December 8 Plus Two, and December 10 Plus Two. In the following analysis we will compute power spectra for each integrated pointing and inspect for residual RFI contamination. As part of this analysis we will pay particular attention to the 4 pointings listed above.

11.3 Power spectra quality analysis

We perform three quality analyses in the power spectrum space. As shown by the blue bars in Figure 11.5, there are 32 pointings over 13 nights with data remaining after the SSINS analysis cuts. Following the procedure in Wilensky et al. (2023), we integrate the observations for each of these 32 pointings and inspect power spectra of the integrations for residual RFI contamination undetected by SSINS.

We also take a close look at the 4 anomalous pointings we identified with the new calibration metric. Power spectra from pointings we remove through these analyses are plotted




Figure 11.6: We histogram means of the coefficients derived by the calibration second-order polynomial fit to amplitude. The means are computed across all unflagged antennas in each observation to produce a value for each fit parameter per observation. Observations retained after the SSINS analysis are colored blue and observations in which narrowband RFI was detected are colored purple. Grey dashed lines indicate $\pm 2\sigma$ from the histogram mean. While some values fall outside of these bounds, there are no significant outliers. Also, there is no discernible pattern for the observations with narrowband RFI, which are sprinkled throughout the distribution.




Figure 11.7: We histogram means of the coefficients derived by the calibration first-order polynomial fit to phase. Observations retained after the SSINS analysis are colored blue and observations in which narrowband RFI was detected are colored purple. Grey dashed lines indicate $\pm 2\sigma$ from the histogram mean. While we see no obvious outliers in the means of the zeroth order coefficient, the means of the first order coefficient are tightly clustered with a dozen observations falling far from the central distribution. These dozen observations comprise all of the December 8 and December 10 observations, including those in the narrowband RFI sets, and are the only observations for which the mean of the first order phase parameter is greater than zero. We will look more closely at the pointing integrations for these two nights in our power spectra quality assessment.




Figure 11.8: We histogram standard deviations of the coefficients derived by the calibration second-order polynomial fit to amplitude. The standard deviations are computed across all unflagged antennas in each observation to produce a value for each fit parameter per observation. Observations retained after the SSINS analysis are colored blue and observations in which narrowband RFI was detected are colored purple. Grey dashed lines indicate $\pm 2\sigma$ from the histogram mean. The means of the amplitude fit coefficients were clustered, and the standard deviations of the first and second order phase coefficients follow the same pattern. The zeroth order coefficient, however, is more sharply distributed with a tail to the left that contains all the observations from December 8 as well as eight observations from December 4. A single observation from December 10 with narrowband RFI contamination is also in this group of outliers.




Figure 11.9: We histogram standard deviations of the coefficients derived by the calibration first-order polynomial fit to phase. Observations retained after the SSINS analysis are colored blue and observations in which narrowband RFI was detected are colored purple. Grey dashed lines indicate $\pm 2\sigma$ from the histogram mean. The standard deviations of the phase fit parameters are distributed similarly to the means, with clustering for the zeroth order coefficient and outliers for the first order coefficient. The same dozen observations from December 8 and December 10 make up the group of outliers for both the means and the standard deviations of the first order phase fit parameters.

in a ‘Wall of Shame’ (Figure 11.10). In this figure pointings removed for residual RFI are demarcated by green boxes, and pointings indicated as potentially anomalous by the new calibration metric are demarcated by blue boxes.

Finally, we integrate our sets of recurring and random narrowband RFI. While we know that observations with DTV are likely to contain residual undetected contamination, it is not clear whether this is true for observations with narrowband emission. To answer this question, we examine power spectra of our recurring and random narrowband integrations for residual RFI. Our three power spectra analyses—hunting for residual RFI, examining integrations indicated by the calibration metric, and evaluating integrations of observations with narrowband RFI—are described in more detail below.

11.3.1 Residual RFI power in the EoR window

Wilensky et al. (2023) identified a signature RFI footprint in the EoR window from ultrafaint emissions below SSINS’s detection level. This residual contamination is characterized by smooth power with little to no dependence on k_{\perp} which extends up to or past the first coarse band line. Our inspection of the power spectra reveals four pointing integrations with this signature.

In these four pointings we observe the residual RFI power in the East-West (XX) instrumental polarization but not the North-South (YY). This polarization dependence was also seen by Wilensky in 2014 data, and in Figure 11.10 we plot only the spectra from the East-West polarizations for these pointing. One of these pointings, December 4 Plus One, is also in the set we identified for closer inspection due to unusually low standard deviations of calibration amplitude fit coefficients, and one (December 10 Plus Two) contains only a single observation which had an unusually high mean and standard deviation of its first order calibration phase fit coefficients. The other two pointings in which we identified residual RFI are November 22 Plus One and Plus Two.

11.3.2 Calibration metric anomalies

Of the four pointings flagged by the calibration metric, two were also identified as having residual RFI power in the EoR window as described above. The other two, December 8 Plus One and Plus Two, were identified for the means and standard deviations of their respective first order phase coefficients as well as the standard deviations of the respective zeroth order amplitude coefficients. In the power spectra of these two pointings we see high power extending throughout the bin at $k_{\perp} \simeq 13\lambda$ and $k_{\perp} \simeq 16\lambda$, respectively. While this power is present in both polarizations, it is strongest in the North-South polarization for Plus One and in the East-West polarization for Plus Two, and so we show the spectra for these polarizations in Figure 11.10. For these two pointings we also examined the 2D plots generated by computing the power spectra of the difference cubes as described in Section 8.3.2. We saw strange features in that diagnostic space and exclude these pointing integrations from our final limit analysis.

11.3.3 Narrowband residual RFI contamination.

Figure 11.11 shows 2D power spectra for the two integrated sets of observations with narrowband RFI detections, recurring (left) and random (right). We plot both instrumental polarizations, and in the East-West polarization of both sets we clearly see the signature of RFI power in the EoR window. This residual RFI suggests that observations in which SSINS detects narrowband RFI should be entirely excluded from our analyses. Though we typically expect RFI to appear at the same sets of frequency channels and bands, the presence of residual contamination in the random narrowband set indicates that the seemingly random pixels SSINS is flagging appear to be true RFI emission.

11.3.4 Data selection from power spectrum analysis

We discard 6 of the 32 integrated pointings we examine in our power spectrum analysis. All pointings from December 4, 8, and 10 are excluded by this cut, as well as two pointings from

November 22. Since the December 6 pointings were cut in our SSINS analysis, our final limit set only contains data from the first 9 nights data was recorded during our observing season.

From the initial set of 319 observations, 180 observations remain after our data quality assessments. These observations contain ~ 5.569 hours of unflagged data which we integrate for our final power spectrum limit analysis. The 2D power spectrum from this final limit set is shown in Figure 11.12. In both polarizations we see no evidence of the residual RFI signature, which indicates that our data quality analysis was successful. Also, the coarse band line power is relatively low, which is encouraging for the power spectrum limit we will show next.

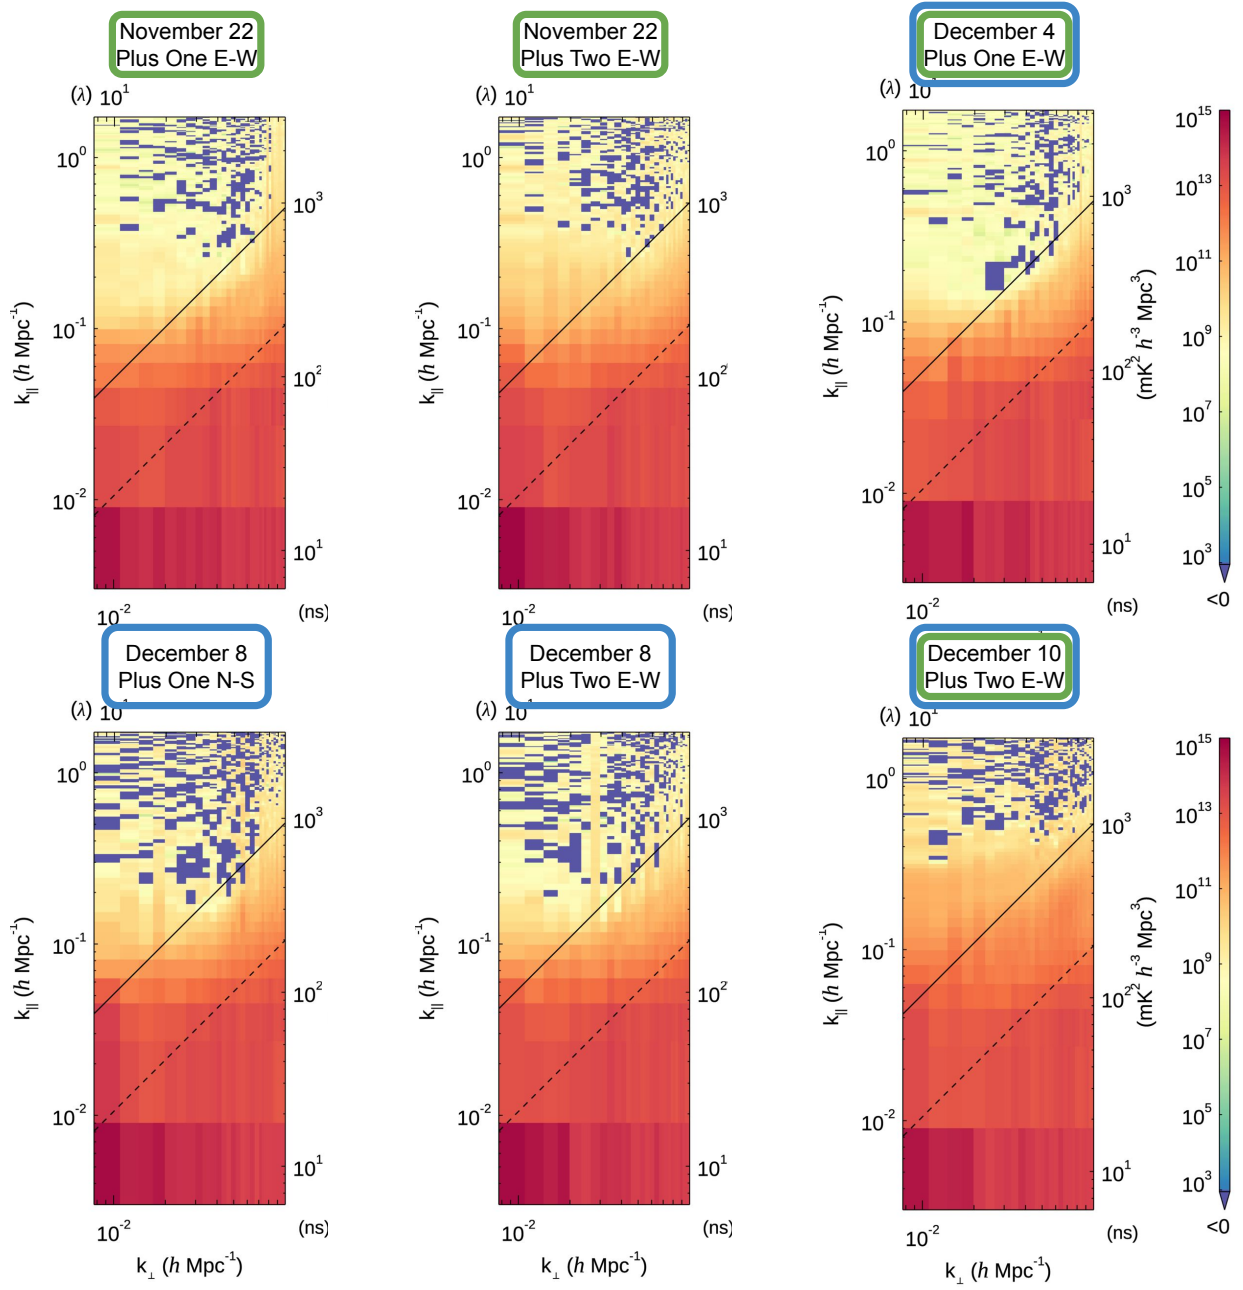


Figure 11.10: A ‘Wall of Shame’ of power spectra from the 6 integrated pointings we exclude from our final limit set. We show a power spectrum from a single instrumental polarization for each pointing. Four integrated pointings are eliminated due to the presence of residual RFI power in the EoR window, as indicated by the green box surrounding the pointing label. These exhibit smooth power extending up to or past the first coarse band. Residual RFI power in the December 4 Plus One pointing is less smooth in k_{\perp} , which may be due to the decrease in sensitivity with increasing k_{\perp} . For these four pointings the RFI signature was only present in the East-West polarization, which we plot here. In addition to containing residual RFI, December 4 Plus One and December 10 Plus Two were also under suspicion due to the calibration fit coefficient metric, as were December 8 Plus One and Plus Two. Labels for these pointings identified by the calibration metric are surrounded by blue boxes. The December 8 pointings were discovered to have odd features in their difference power spectra as well as high power extending into the EoR window at $k_{\perp} \simeq 13\lambda$ and $k_{\perp} \simeq 16\lambda$, respectively. While this power is present in both polarizations for both of these observations, it is brighter in the EoR window in the North-South polarization of the Plus One pointing and in the East-West polarization of the Plus Two pointing, so we plot those polarizations here.

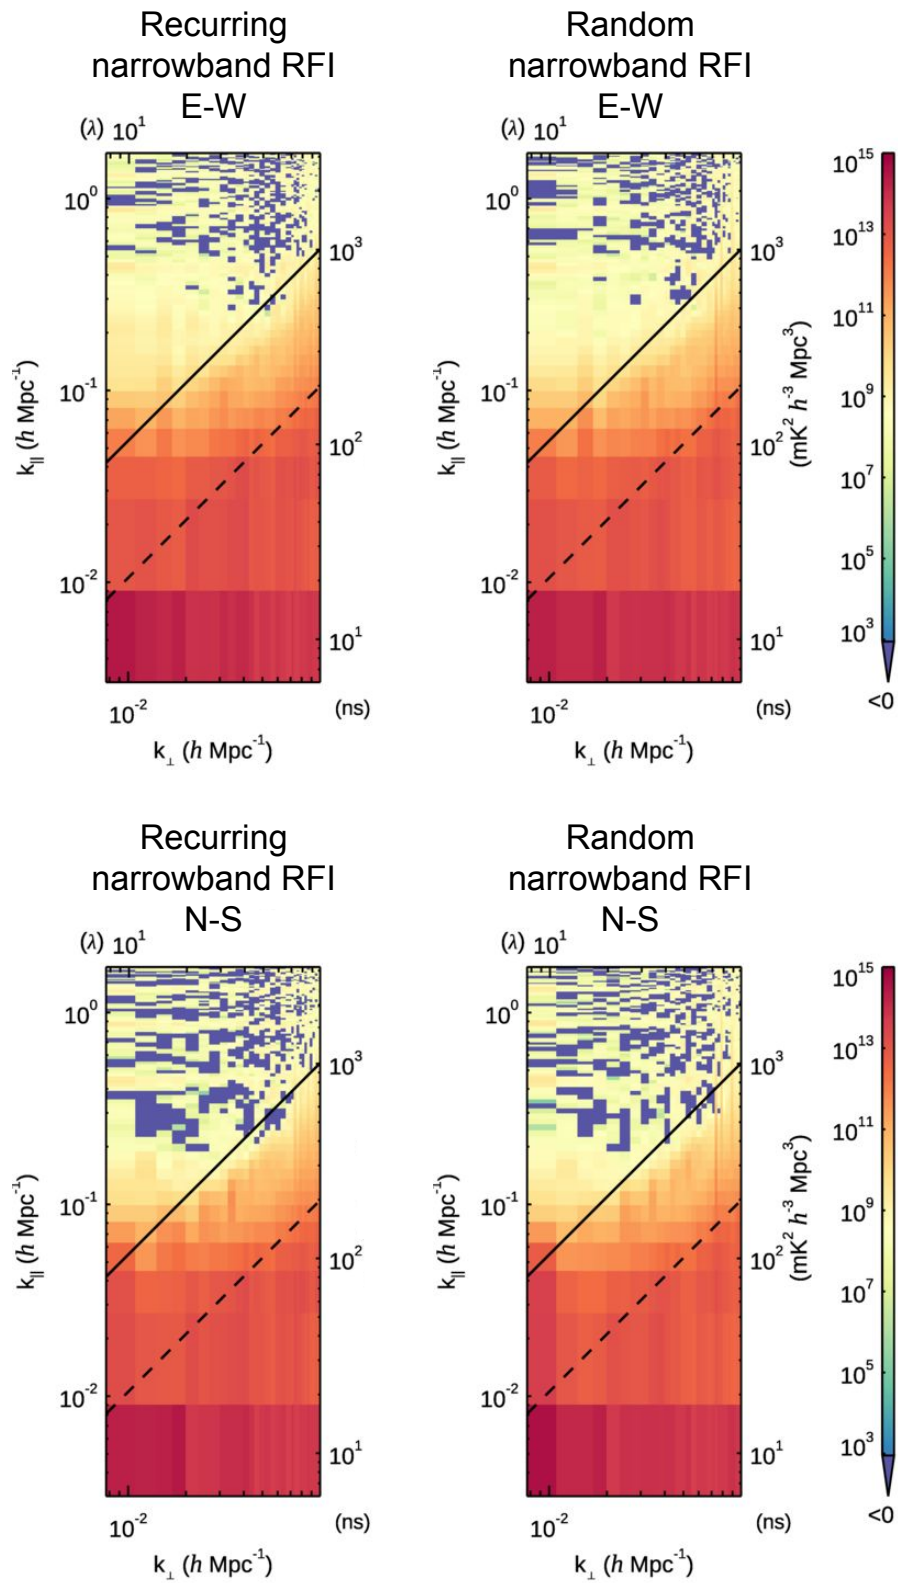


Figure 11.11: We show power spectra from both the East-West and North-South instrumental polarizations of two integrated sets of observations in which narrowband RFI has been identified and flagged. Residual RFI power is clearly evident in the East-West power spectrum for both sets. In one set, the frequency channel in which the RFI was detected was flagged for multiple timestamps throughout the data set. This ‘recurring’ narrowband set is shown in the left panels. The other ‘random’ set, shown in the right panels, contained narrowband RFI which appeared in a specific frequency channel only once through the entire data set. The random set does not show markedly different behavior from the recurring set. This suggests that the pixels flagged in this set may be true RFI emission. In both sets the residual RFI power occurs in the East-West polarization, following the pattern identified by Wilensky. The presence of this residual power indicates that observations in which SSINS has detected and flagged narrowband RFI should be excluded from analyses.




Figure 11.12: Our data quality analysis excised 139 observations from our initial data set of 319. The remaining 180 observations make up our final limit set and contain ~ 5.569 hours of unflagged data. We show the 2D power spectra of our integrated limit set for the East-West (left) and North-South (right) instrumental polarizations. There is no evidence of residual RFI power in the EoR window, which suggests our RFI excision methods were successful. Also, we see relatively low power in the coarse band lines, which is encouraging for our power spectrum limit.

Chapter 12

A FIRST POWER SPECTRUM LIMIT WITH MWA PHASE III DATA

We obtain the first power spectrum limit with the new MWAX correlator from the 180 observations that survived our quality analysis cuts. This final limit set contains ~ 5.569 hours of data taken from November 16 - December 2, 2021.

The EoR power spectrum evolves with redshift. To make an EoR measurement, we must analyze a frequency bandwidth encompassing a redshift range within which the EoR signal is statistically stable. This is expected at a redshift range of ~ 0.5 corresponding to a frequency range of ~ 8 MHz. Rather than taking the Fourier transform of our full ~ 30 MHz bandwidth and computing the power spectrum, we need to utilize shorter frequency subbands for our analysis.

We process three overlapping subsets of the frequency band: 167.02-182.38 MHz, 174.70-190.06 MHz, and 182.38-197.74 MHz. While these subbands are each of length 15.36 MHz, we apply a Blackman-Harris window before the frequency Fourier transform. This tapering function creates an effective bandwidth of ~ 7.68 MHz. Though the frequency subbands overlap, the windowing function ensures that the resultant redshift bins, centered at 7.1, 6.8, and 6.5 respectively, are mostly uncorrelated. We show 2D power spectra for each redshift and instrumental polarization in Figure 12.1. The region over which we compute the 1D power spectrum and upper limit is delineated in the EoR window.

When computing the 1D power spectrum we wish to exclude foreground power from the integration, as it would overwhelm the faint EoR signal. To isolate the region in 3D power

spectrum space over which we will integrate, we employ the cuts developed by Li et al. (2019). Li’s analysis was done on Phase II data which had the same antenna layout as our Phase III MWAX data. This means our data has the same uv sampling, and thus the same distribution of sensitivity. These cuts, while performed in the 3D space, are shown for each redshift and instrumental polarization in Figure 12.1. We set a lower k_{\parallel} bound of 0.15 hMpc^{-1} and increase the horizon slope by 15%. This excludes the foreground power in the Wedge and lower k_{\parallel} bins. We restrict k_{\perp} to a range of $12\text{-}50 \lambda$. Li found that foreground power in modes $k_{\perp} < 12 \lambda$ coupled with higher k_{\parallel} modes. As k_{\perp} increases, not only do the coarse band lines widen but also our sensitivity drops. We exclude modes higher than 50λ to reduce the coarse band contamination extent without a significant loss in sensitivity to the EoR signal.

Li also employed cuts excluding the coarse band lines, but we retain that data here. We wish to see the impact of the MWAX correlator on these systematics, and including them does not contaminate other data. After the spherical average, the coarse band line power is confined to specific k modes. Our hope is that the reduction of nonlinear effects in MWAX correlator data will reduce not only the height but also the width of the coarse band lines, potentially opening up previously contaminated regions.

12.1 The 1D power spectrum limit

We compute the 1D power spectrum limit for each redshift and polarization using the method outlined by Beardsley (2015) and by Li et al. (2019). In Figure 12.2 we plot the observed power spectrum (solid black line), 2σ upper limit (solid blue line), and 1σ thermal noise (dashed blue line). The 2σ error bars on the power spectrum measurement are shaded grey. A fiducial theoretical EoR model (solid brown line) is shown with its 95% confidence interval (dashed brown line). The calculation of this model is described in Barry et al. (2019b). Full tables of upper limits for each polarization at every redshift are shown in Tables 12.1 and 12.2. We report a lowest upper limit of $\Delta^2 \leq 1.76 \cdot 10^4 \text{ mK}^2$ or $\Delta^2 \leq (133 \text{ mK})^2$ at

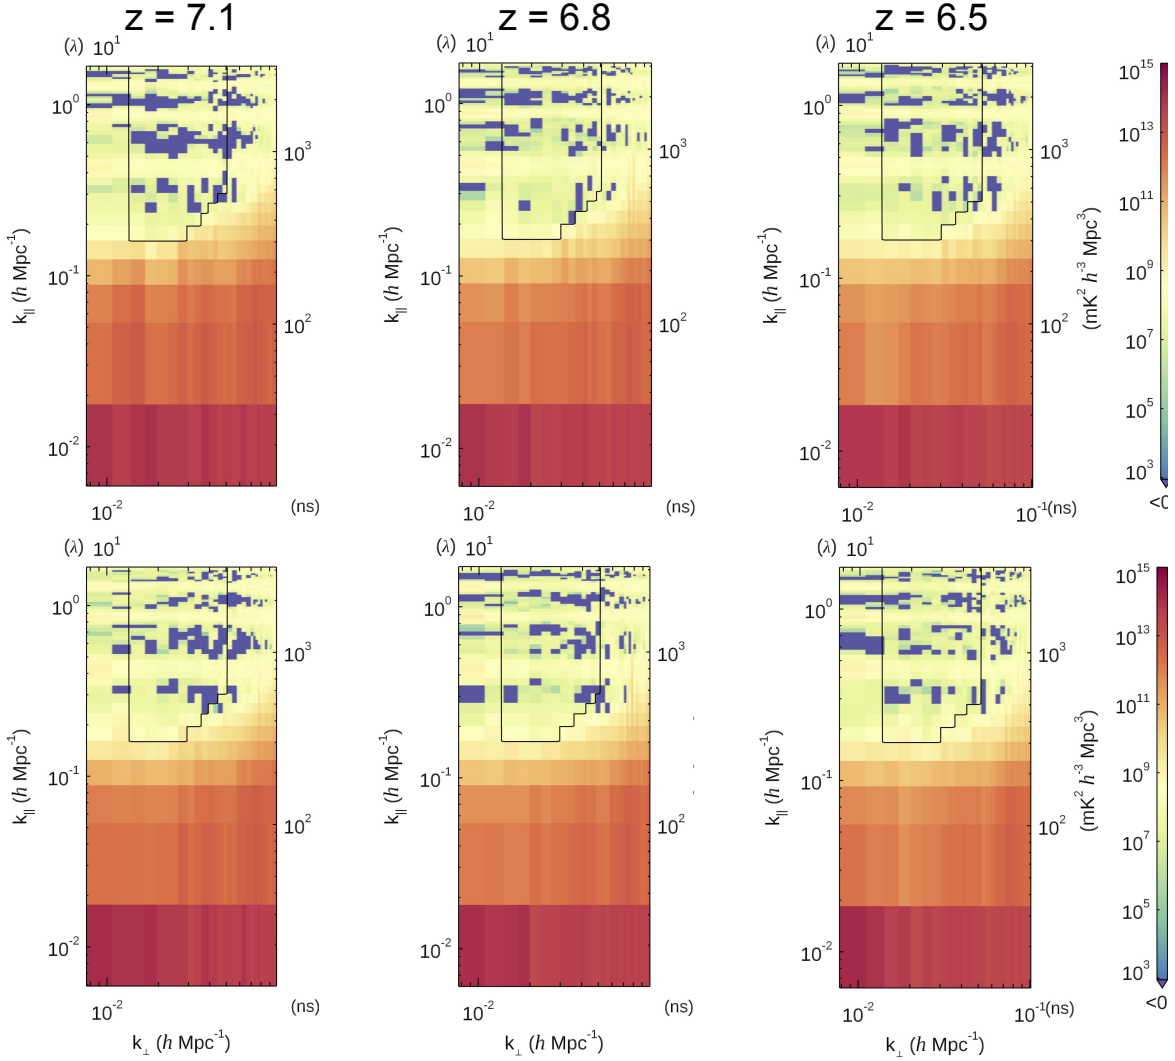


Figure 12.1: We show the 2D power spectra of our 180 observation limit set at three different redshifts 7.1, 6.8, and 6.5. The East-West instrumental polarization is shown in the upper row and the North-South polarization is in the lower row. In the EoR window we delineate the region over which we compute our 1D power spectrum limit.

$k = 0.332h \text{ Mpc}^{-1}$ in the North-South polarization data of redshift $z = 6.5$.

In Figure 12.2 we see relatively low power in the coarse band line peaks, located at $k \simeq 0.4$, 0.9 , and $1.1 h \text{ Mpc}^{-1}$. Comparison with the 1D power spectrum shown by Wilensky et al. (2023) (which we show here in Figure 12.3), shows a drop in coarse peak power of about half an order of magnitude. This is likely a direct effect of the reduction of nonlinear artifacts achieved by eliminating quantization stages within the MWAX correlator. We showed in Figure 4.5 that using the Van Vleck correction reduced coarse band line power in the analysis of Phase I MWA data (Wilensky et al., 2023). Further comparison with the 2D power spectrum of our 2021 MWAX data (Figure 11.12) shows even more improvement. The width of the coarse band lines bleeds power into the regions between the peaks, where we would like to make an EoR measurement. The improvements we see in the coarse band line power have potential to reduce this bleed and open up more k bins for EoR science.

Where the 2σ limit error bars (shaded grey in Figure 12.2) go to zero the associated bins are noise-dominated. This suggests that adding more data for a deeper integration could lower the upper limit in these bins. However, interpreting the noise-dominated bins in the context of other analyses is complicated with the 2021 MWAX data set due to our limited data quantity. As more data is included in a power spectrum analysis, the noise floor drops. The presence of noise-dominated bins becomes more compelling with deeper integrations which have a lower noise floor. Recent MWA limit analyses have used $\sim 500 - 700$ observations (Barry et al. 2019b; Li et al. 2019; Wilensky et al. 2023). Our limit set was taken during the MWAX commissioning season, and so even our initial data set before quality control cuts was much smaller at 319 observations. This means our noise floor is relatively high, making it difficult to directly compare with other limits.

However, the presence of noise-dominated bins in Figure 12.2 is encouraging. In particular, in Figure 12.2 we consistently see noise domination in the gaps between the coarse band

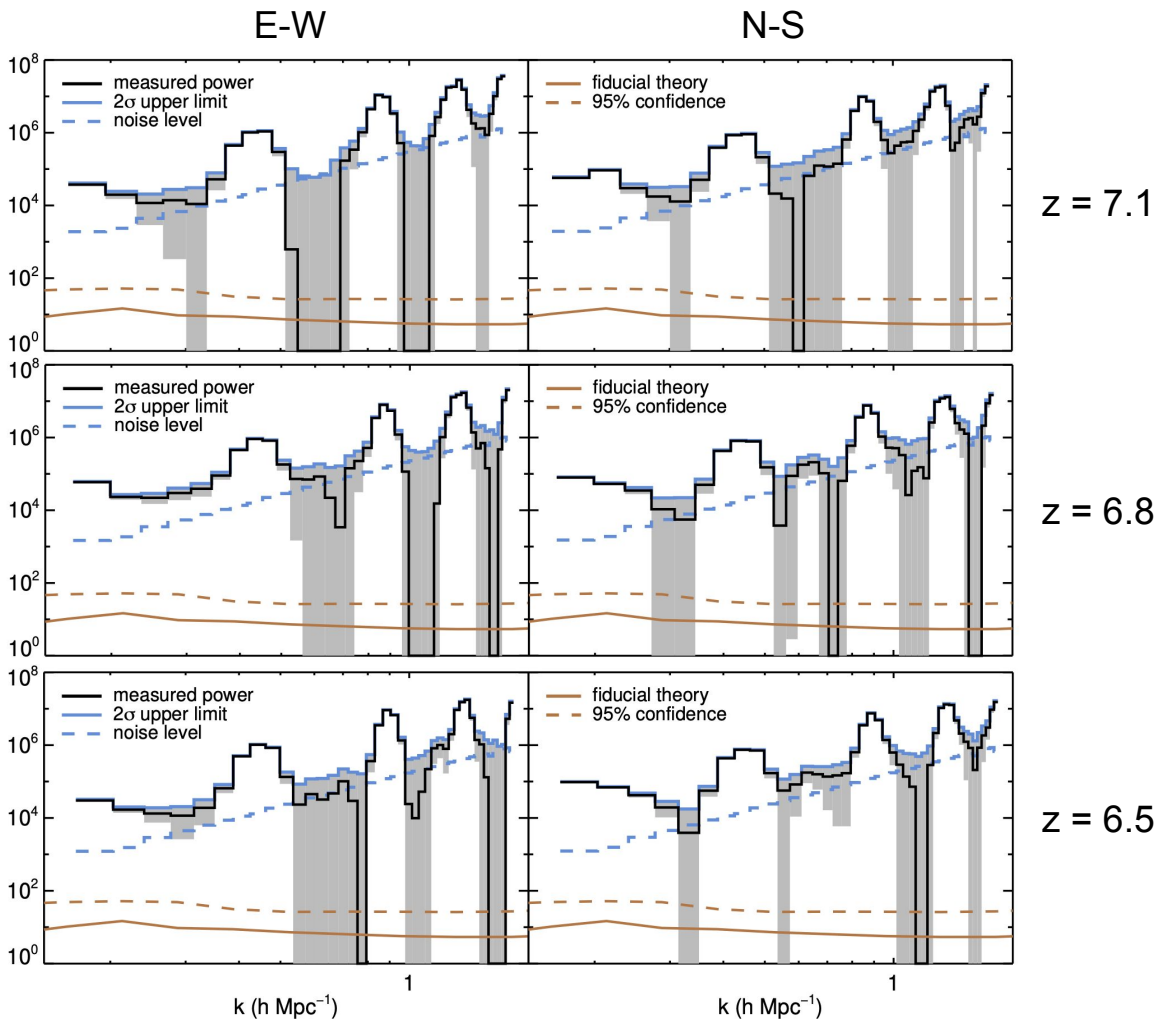


Figure 12.2: We plot the 1D observed power (solid black line), upper limit (solid blue line), and 1σ thermal noise (dashed blue line) for redshifts 7.1 (upper), 6.8 (center), and 6.5 (lower). The 2σ error bars are shown in grey, and where these errors go to zero the associated bins are noise-dominated. A fiducial EoR model (solid brown line) is shown with its 95% confidence interval (dashed brown line). While we see noise-dominated bins, this is difficult to interpret in the context of other upper limit analyses due to the low number of observations in our data set. However, the presence of these bins suggests that adding more data might lower the upper limit.

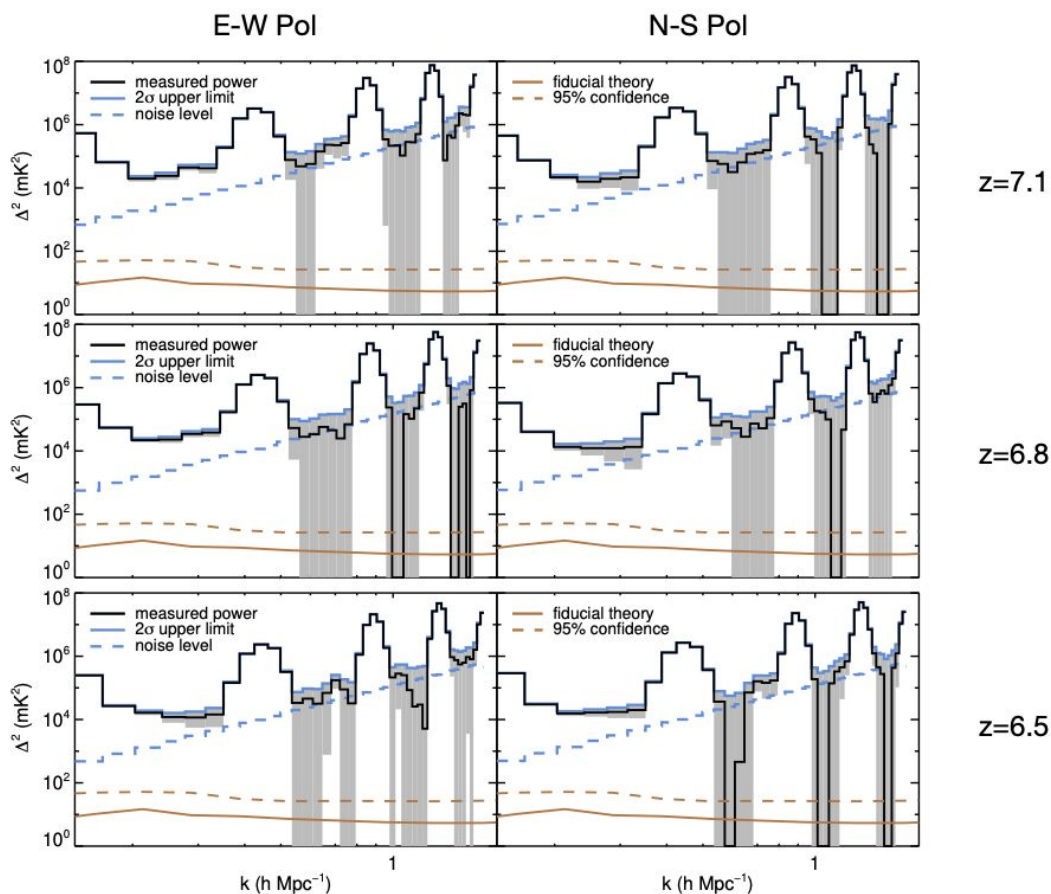


Figure 12.3: We show the 1D power spectrum from Wilensky et al. (2023). Wilensky analyzed MWA Phase I 2014 data with Van Vleck corrections for the 4-bit requantization stage. Comparison with our analysis of MWA Phase III 2021 data (Figure 12.2) shows that the new MWAX correlator decreases power in the coarse band line peaks (located at $k \simeq 0.4, 0.9, \text{ and } 1.1 h \text{ Mpc}^{-1}$).

line, especially in the North-South polarization. This difference between polarizations is unsurprising since the North-South polarization is likely to be less contaminated by residual RFI compared to the East-West polarization, as shown in Section 11.3. We also see noise-dominated bins below the first coarse band peak. Since the EoR signal increases with lower k , this region is of particular interest for making an EoR measurement.

12.2 Discussion

To place our upper limit into context, we compare with the results from Wilensky et al. (2023). Wilensky analyzed 2014 MWA data which used the Legacy correlator. We applied Van Vleck corrections to that data and it is interesting to see how MWAX data compares to the corrected Legacy data. However, the comparison is complicated by the fact that Wilensky’s data used the Phase I antenna layout while our analysis uses the updated Phase II antenna layout. Also, the number of observations included in our final limit is 180 while Wilensky’s analysis used 472 observations.

The lowest upper limit reported by Wilensky is $\Delta^2 \leq 1.61 \cdot 10^4 \text{ mK}^2$ at $k = 0.258 \text{ h Mpc}^{-1}$ with the East-West data in $z = 7.1$. Our upper limit of $\Delta^2 \leq 1.76 \cdot 10^4 \text{ mK}^2$ is similar and is noise-dominated, which indicates that data from the MWAX correlator is performing at least as well as data from the Legacy correlator with corrections for quantization nonlinearities.

The fact that our upper limit is noise-dominated and is close to that reported by Wilensky suggests that adding more data to this analysis has potential to further reduce the limit. A second season of MWAX data was taken in late 2023, too late to be included in this work. While it is impossible to predict the outcome if more data is added, the promising limit results presented above suggest the possibility that adding 2023 MWAX data to the set processed here could lead to state of the art limits on the EoR.

12.3 Conclusion

In this work we have advanced 21-centimeter Epoch of Reionization science by addressing the detrimental impacts of digital nonlinearities in our data.

To make a measurement of the 21-centimeter signal from neutral hydrogen during the EoR, we must remove the antenna response from our data with a precision of a part in 10^{-5} . As we discussed in Chapter 2, calibration removes only a linear term of the antenna gain. Any nonlinear terms of the antenna response must be removed before calibration.

Digital quantization stages are known to result in a nonlinear instrument response (Chapter 3). Like most instruments, the MWA digital signal path contains several quantization stages. In Chapter 4 and Chapter 5 we showed that implementing corrections for MWA quantization stages improved data quality and analyses. It is possible that existing MWA upper limits could be deepened by re-analyzing data after applying corrections for the 4-bit requantization stage as well as requantization stages deeper in the signal path.

Largely as a result of our work demonstrating the impact of correcting digital nonlinearities, the new correlator being developed for MWA Phase III was designed to eliminate quantization stages. We analyzed the first EoR data taken during the commissioning season of the new MWAX correlator, described in Chapter 6. An inspection of autocorrelations in Chapter 7 revealed reduced nonlinear artifacts in this data set. This suggests future instrument design should prioritize the minimization of requantization stages.

We then prepared to calculate a power spectrum limit with the MWAX commissioning data set (Chapter 8) to analyze the effects of the new correlator on our EoR science. This required setting up high-powered compute systems (Chapter 9), retuning our flagging and averaging settings for MWAX data (Chapter 10), and performing data selection based on quality analyses (Chapter 11).

In Chapter 12 we present the first upper limit with MWA Phase III data from the MWAX correlator. Our lowest upper limit of $\Delta^2 \leq 1.76 \cdot 10^4 \text{ mK}^2$ is very promising, and suggests that data from the new MWAX correlator is performing at least as well as Legacy correlator data with nonlinear corrections applied. We are hopeful that adding MWAX data from later observing seasons to our commissioning data set might deepen this result.

Through this work we advanced EoR science by showing that removing effects from digital nonlinearities is crucial for making an EoR measurement, whether by applying corrections to existing data or by building and upgrading instruments to eliminate requantization stages. This is not only beneficial for the international MWA collaboration, but also supports all science groups working towards an EoR measurement.

Upper limits for the East-West polarization

z=7.1			z=6.8			z=6.5		
k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)	k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)	k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)
0.177	4.128e+04	1.900e+03	0.181	6.201e+04	1.479e+03	0.185	3.262e+04	1.214e+03
0.212	2.449e+04	2.358e+03	0.217	2.691e+04	1.861e+03	0.222	1.998e+04	1.542e+03
0.248	2.056e+04	4.418e+03	0.253	2.905e+04	3.516e+03	0.258	1.905e+04	2.901e+03
0.283	2.760e+04	6.801e+03	0.289	4.062e+04	5.421e+03	0.295	2.050e+04	4.472e+03
0.318	3.075e+04	9.615e+03	0.325	5.439e+04	7.671e+03	0.332	3.167e+04	6.342e+03
0.354	7.876e+04	1.318e+04	0.362	1.109e+05	1.052e+04	0.369	8.284e+04	8.698e+03
0.389	4.773e+05	1.709e+04	0.398	4.791e+05	1.362e+04	0.406	5.148e+05	1.126e+04
0.424	1.057e+06	1.975e+04	0.434	9.446e+05	1.566e+04	0.443	1.050e+06	1.295e+04
0.460	1.141e+06	2.823e+04	0.470	8.632e+05	2.248e+04	0.480	8.768e+05	1.858e+04
0.495	3.689e+05	3.616e+04	0.506	2.386e+05	2.887e+04	0.517	1.819e+05	2.387e+04
0.531	1.017e+05	4.447e+04	0.542	1.439e+05	3.552e+04	0.554	8.496e+04	2.936e+04
0.566	6.450e+04	5.398e+04	0.579	1.572e+05	4.310e+04	0.591	1.180e+05	3.563e+04
0.601	5.952e+04	6.433e+04	0.615	1.882e+05	5.128e+04	0.628	1.213e+05	4.239e+04
0.637	6.871e+04	7.393e+04	0.651	1.501e+05	5.879e+04	0.665	1.481e+05	4.860e+04
0.672	1.771e+05	8.977e+04	0.687	1.655e+05	7.158e+04	0.702	2.212e+05	5.917e+04
0.707	3.838e+05	1.054e+05	0.723	3.128e+05	8.418e+04	0.738	1.803e+05	6.959e+04
0.743	5.869e+05	1.220e+05	0.759	4.217e+05	9.745e+04	0.775	1.630e+05	8.057e+04
0.778	1.310e+06	1.403e+05	0.795	7.410e+05	1.120e+05	0.812	5.497e+05	9.263e+04
0.814	4.749e+06	1.603e+05	0.832	3.697e+06	1.280e+05	0.849	3.717e+06	1.058e+05
0.849	1.113e+07	1.822e+05	0.868	8.194e+06	1.454e+05	0.886	9.370e+06	1.202e+05
0.884	9.909e+06	2.059e+05	0.904	5.778e+06	1.644e+05	0.923	6.929e+06	1.359e+05
0.920	3.852e+06	2.316e+05	0.940	1.571e+06	1.850e+05	0.960	1.651e+06	1.529e+05
0.955	1.040e+06	2.594e+05	0.976	5.586e+05	2.071e+05	0.997	4.072e+05	1.712e+05
0.990	5.490e+05	2.893e+05	1.012	4.398e+05	2.310e+05	1.034	4.420e+05	1.910e+05
1.026	4.372e+05	3.193e+05	1.049	3.970e+05	2.546e+05	1.071	5.173e+05	2.104e+05
1.061	4.331e+05	3.424e+05	1.085	4.081e+05	2.723e+05	1.108	6.858e+05	2.251e+05
1.097	6.879e+05	3.899e+05	1.121	5.053e+05	3.108e+05	1.145	1.333e+06	2.570e+05
1.132	1.685e+06	4.318e+05	1.157	7.962e+05	3.448e+05	1.182	1.578e+06	2.850e+05
1.167	3.626e+06	4.736e+05	1.193	1.782e+06	3.782e+05	1.219	1.418e+06	3.126e+05
1.203	9.167e+06	5.179e+05	1.229	5.027e+06	4.135e+05	1.255	2.684e+06	3.419e+05
1.238	1.855e+07	5.508e+05	1.265	1.363e+07	4.387e+05	1.292	9.971e+06	3.625e+05
1.273	1.934e+07	5.351e+05	1.302	1.568e+07	4.240e+05	1.329	1.582e+07	3.503e+05
1.309	2.921e+07	6.488e+05	1.338	1.788e+07	5.173e+05	1.366	1.835e+07	4.277e+05
1.344	1.673e+07	7.228e+05	1.374	7.770e+06	5.773e+05	1.403	6.570e+06	4.773e+05
1.380	5.835e+06	7.817e+05	1.410	2.633e+06	6.242e+05	1.440	2.416e+06	5.160e+05
1.415	3.491e+06	8.429e+05	1.446	1.921e+06	6.733e+05	1.477	2.171e+06	5.566e+05
1.450	3.047e+06	8.646e+05	1.482	2.134e+06	6.907e+05	1.514	1.395e+06	5.708e+05
1.486	2.855e+06	7.591e+05	1.519	1.482e+06	6.052e+05	1.551	1.115e+06	4.998e+05
1.521	2.940e+06	9.918e+05	1.555	1.631e+06	7.941e+05	1.588	1.372e+06	6.567e+05
1.556	5.549e+06	1.122e+06	1.591	1.223e+06	8.960e+05	1.625	9.301e+05	7.408e+05
1.592	1.254e+07	1.201e+06	1.627	2.537e+06	9.589e+05	1.662	1.042e+06	7.927e+05
1.627	3.250e+07	1.279e+06	1.663	1.274e+07	1.022e+06	1.699	6.961e+06	8.449e+05
1.663	3.762e+07	9.990e+05	1.699	2.176e+07	8.026e+05	1.735	1.596e+07	6.624e+05

Table 12.1: We show upper limits for the East-West polarization. For each redshift, the columns report the bin center k ($h \text{ Mpc}^{-1}$), the 2σ upper limit δ_{UL}^2 (mK^2) and the 1σ thermal noise σ (mK^2)

Upper limits for the North-South polarization

z=7.1			z=7.1			z=7.1		
k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)	k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)	k ($h \text{ Mpc}^{-1}$)	δ_{UL}^2 (mK^2)	σ (mK^2)
0.177	6.108e+04	1.944e+03	0.181	8.220e+04	1.514e+03	0.185	9.911e+04	1.236e+03
0.212	9.714e+04	2.411e+03	0.217	5.689e+04	1.905e+03	0.222	7.233e+04	1.568e+03
0.248	3.890e+04	4.514e+03	0.253	4.213e+04	3.600e+03	0.258	4.806e+04	2.947e+03
0.283	3.150e+04	6.948e+03	0.289	2.184e+04	5.549e+03	0.295	2.856e+04	4.543e+03
0.318	3.292e+04	9.823e+03	0.325	2.216e+04	7.854e+03	0.332	1.764e+04	6.443e+03
0.354	7.757e+04	1.346e+04	0.362	7.162e+04	1.077e+04	0.369	7.372e+04	8.837e+03
0.389	4.235e+05	1.747e+04	0.398	4.549e+05	1.395e+04	0.406	4.566e+05	1.144e+04
0.424	8.940e+05	2.018e+04	0.434	8.392e+05	1.604e+04	0.443	7.728e+05	1.315e+04
0.460	9.661e+05	2.884e+04	0.470	8.250e+05	2.301e+04	0.480	7.467e+05	1.888e+04
0.495	2.861e+05	3.694e+04	0.506	2.111e+05	2.956e+04	0.517	2.191e+05	2.425e+04
0.531	1.192e+05	4.544e+04	0.542	8.550e+04	3.636e+04	0.554	1.165e+05	2.983e+04
0.566	1.369e+05	5.515e+04	0.579	1.770e+05	4.413e+04	0.591	1.559e+05	3.620e+04
0.601	1.484e+05	6.572e+04	0.615	2.840e+05	5.250e+04	0.628	2.623e+05	4.307e+04
0.637	2.229e+05	7.553e+04	0.651	3.302e+05	6.019e+04	0.665	2.576e+05	4.937e+04
0.672	3.115e+05	9.171e+04	0.687	2.535e+05	7.328e+04	0.702	2.572e+05	6.012e+04
0.707	3.395e+05	1.077e+05	0.723	1.608e+05	8.618e+04	0.738	2.893e+05	7.070e+04
0.743	3.974e+05	1.247e+05	0.759	2.763e+05	9.977e+04	0.775	3.340e+05	8.185e+04
0.778	9.372e+05	1.434e+05	0.795	8.239e+05	1.147e+05	0.812	8.282e+05	9.411e+04
0.814	4.315e+06	1.638e+05	0.832	3.830e+06	1.311e+05	0.849	3.716e+06	1.075e+05
0.849	1.004e+07	1.861e+05	0.868	7.846e+06	1.489e+05	0.886	7.696e+06	1.221e+05
0.884	7.679e+06	2.104e+05	0.904	4.929e+06	1.683e+05	0.923	5.103e+06	1.381e+05
0.920	2.465e+06	2.366e+05	0.940	1.438e+06	1.894e+05	0.960	1.694e+06	1.553e+05
0.955	1.162e+06	2.650e+05	0.976	9.494e+05	2.121e+05	0.997	9.758e+05	1.740e+05
0.990	8.895e+05	2.956e+05	1.012	9.853e+05	2.365e+05	1.034	6.835e+05	1.940e+05
1.026	1.105e+06	3.262e+05	1.049	8.616e+05	2.606e+05	1.071	5.910e+05	2.138e+05
1.061	1.254e+06	3.498e+05	1.085	6.535e+05	2.787e+05	1.108	5.734e+05	2.287e+05
1.097	1.378e+06	3.983e+05	1.121	8.140e+05	3.182e+05	1.145	4.897e+05	2.610e+05
1.132	2.009e+06	4.412e+05	1.157	9.171e+05	3.530e+05	1.182	6.066e+05	2.896e+05
1.167	3.257e+06	4.838e+05	1.193	9.363e+05	3.872e+05	1.219	9.473e+05	3.176e+05
1.203	6.166e+06	5.291e+05	1.229	3.438e+06	4.234e+05	1.255	2.842e+06	3.473e+05
1.238	1.471e+07	5.628e+05	1.265	1.051e+07	4.491e+05	1.292	1.123e+07	3.684e+05
1.273	1.860e+07	5.467e+05	1.302	1.276e+07	4.341e+05	1.329	1.377e+07	3.560e+05
1.309	2.005e+07	6.629e+05	1.338	1.401e+07	5.296e+05	1.366	1.345e+07	4.345e+05
1.344	6.999e+06	7.385e+05	1.374	7.324e+06	5.910e+05	1.403	5.721e+06	4.848e+05
1.380	2.060e+06	7.986e+05	1.410	3.761e+06	6.391e+05	1.440	3.119e+06	5.243e+05
1.415	2.363e+06	8.611e+05	1.446	3.035e+06	6.893e+05	1.477	2.374e+06	5.654e+05
1.450	3.188e+06	8.834e+05	1.482	2.698e+06	7.071e+05	1.514	2.016e+06	5.800e+05
1.486	3.878e+06	7.757e+05	1.519	9.222e+05	6.196e+05	1.551	1.311e+06	5.079e+05
1.521	4.563e+06	1.014e+06	1.555	1.015e+06	8.129e+05	1.588	2.214e+06	6.671e+05
1.556	4.024e+06	1.146e+06	1.591	1.857e+06	9.173e+05	1.625	3.356e+06	7.525e+05
1.592	5.214e+06	1.227e+06	1.627	4.083e+06	9.817e+05	1.662	4.635e+06	8.053e+05
1.627	1.481e+07	1.307e+06	1.663	1.083e+07	1.046e+06	1.699	1.110e+07	8.584e+05
1.663	2.098e+07	1.021e+06	1.699	1.633e+07	8.213e+05	1.735	1.659e+07	6.738e+05

Table 12.2: We show upper limits for the North-South polarization. For each redshift, the columns report the bin center k ($h \text{ Mpc}^{-1}$), the 2σ upper limit δ_{UL}^2 (mK^2) and the 1σ thermal noise σ (mK^2)

BIBLIOGRAPHY

- N Barry, B Hazelton, I Sullivan, M F Morales, and J C Pober. Calibration Requirements for Detecting the 21 cm Epoch of Reionization Power Spectrum and Implications for the SKA. *Mon. Not. R. Astron. Soc.*, 461:3135–3144, 2016. doi: 10.1093/mnras/stw1380.
- N Barry, A P Beardsley, R Byrne, B Hazelton, M F Morales, J C Pober, and I Sullivan. The FHD/ ϵ ppsi Epoch of Reionisation Power Spectrum Pipeline. *Publ. Astron. Soc. Aust.*, 36:E026, 2019a. doi: 10.1017/pasa.2019.21.
- N Barry, M Wilensky, C M Trott, B Pindor, A P Beardsley, B J Hazelton, I S Sullivan, M F Morales, J C Pober, J Line, B Greig, R Byrne, A Lanman, W Li, C H Jordan, R C Joseph, B McKinley, M Rahimi, S Yoshiura, J D Bowman, B M Gaensler, J N Hewitt, D C Jacobs, D A Mitchell, N Udaya Shankar, S K Sethi, R Subrahmanyam, S J Tingay, R L Webster, and J S B Wyithe. Improving the Epoch of Reionization Power Spectrum Results from Murchison Widefield Array Season 1 Observations. *Astrophys. J.*, 884:1, 2019b. doi: 10.3847/1538-4357/ab40a8.
- Ben Bean, Sanjay Bhatnagar, Sandra Castro, Jennifer Donovan Meyer, Bjorn Emonts, Enrique Garcia, Robert Garwood, Kumar Golap, Justo Gonzalez Villalba, Pamela Harris, Yohei Hayashi, Josh Hoskins, Mingyu Hsieh, Preshanth Jagannathan, Wataru Kawasaki, Aard Keimpema, Mark Kettenis, Jorge Lopez, Joshua Marvil, Joseph Masters, Andrew McNichols, David Mehringer, Renaud Miel, George Moellenbrock, Federico Montesino, Takeshi Nakazato, Juergen Ott, Dirk Petry, Martin Pokorny, Ryan Raba, Urvashi Rau, Darrell Schiebel, Neal Schweighart, Srikrishna Sekhar, Kazuhiko Shimada, Des Small, Jan Willem Steeb, Kanako Sugimoto, Ville Suoranta, Takahiro Tsutsumi, Ilse M. Van Bemmelen, Marjolein Verkouter, Akeem Wells, Wei Xiong, Arpad Szomoru,

Morgan Griffith, Brian Glendenning, and Jeff Kern. CASA, the Common Astronomy Software Applications for Radio Astronomy. *Publications of the Astronomical Society of the Pacific*, 134:114501, 11 2022. ISSN 1538-3873. doi: 10.1088/1538-3873/AC9642. URL <https://iopscience.iop.org/article/10.1088/1538-3873/ac9642><https://iopscience.iop.org/article/10.1088/1538-3873/ac9642/meta>.

A P Beardsley, B J Hazelton, I S Sullivan, P Carroll, N Barry, M Rahimi, B Pindor, C M Trott, J Line, Daniel C Jacobs, M F Morales, J C Pober, G Bernardi, Judd D Bowman, M P Busch, F Briggs, R J Cappallo, B E Corey, A de Oliveira-Costa, Joshua S Dillon, D Emrich, A Ewall-Wice, L Feng, B M Gaensler, R Goeke, L J Greenhill, J N Hewitt, N Hurley-Walker, M Johnston-Hollitt, D L Kaplan, J C Kasper, H S Kim, E Kratzenberg, E Lenc, A Loeb, C J Lonsdale, M J Lynch, B McKinley, S R McWhirter, D A Mitchell, E Morgan, A R Neben, Nithyanandan Thyagarajan, D Oberoi, A R Offringa, S M Ord, S Paul, T Prabu, P Procopio, J Riding, A E E Rogers, A Roshi, N Udaya Shankar, Shiv K Sethi, K S Srivani, R Subrahmanyam, M Tegmark, S J Tingay, M Waterson, R B Wayth, R L Webster, A R Whitney, A Williams, C L Williams, C Wu, and J S B Wyithe. First Season MWA EoR Power Spectrum Results at Redshift 7. *Astrophys. J.*, 833:102, 2016. doi: 10.3847/1538-4357/833/1/102.

Adam Beardsley. *The Murchison Widefield Array 21cm Epoch of Reionization Experiment: Design, Construction, and First Season Results*. PhD thesis, University of Washington, 2015. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works. Last updated - 2023-03-04.

Adam Beardsley. epsilon error propagation, Accessed June 4 2024. URL https://github.com/EoRImaging/Memos/blob/main/PDFs/005_epsilon_error_prop.pdf.

L. V. Benkevitch, A. E. E. Rogers, C. J. Lonsdale, R. J. Cappallo, D. Oberoi, P. J. Erickson, and K. A. V. Baker. Van Vleck correction generalization for complex correlators with multilevel quantization, 8 2016. URL <http://arxiv.org/abs/1608.04367>.

Sarah E.I. Bosman, Frederick B. Davies, George D. Becker, Laura C. Keating, Rebecca L. Davies, Yongda Zhu, Anna Christina Eilers, Valentina D'Odorico, Fuyan Bian, Manuela Bischetti, Stefano V. Cristiani, Xiaohui Fan, Emanuele P. Farina, Martin G. Haehnelt, Joseph F. Hennawi, Girish Kulkarni, Andrei Mesinger, Romain A. Meyer, Masafusa Onoue, Andrea Pallottini, Yuxiang Qin, Emma Ryan-Weber, Jan Torge Schindler, Fabian Walter, Feige Wang, and Jinyi Yang. Hydrogen reionization ends by $z = 5.3$: Lyman- α optical depth measured by the XQR-30 sample. *Monthly Notices of the Royal Astronomical Society*, 514:55–76, 6 2022. ISSN 0035-8711. doi: 10.1093/MNRAS/STAC1046. URL <https://dx.doi.org/10.1093/mnras/stac1046>.

Judd D. Bowman, Alan E.E. Rogers, Raul A. Monsalve, Thomas J. Mozdzen, and Nivedita Mahesh. An absorption profile centred at 78 megahertz in the sky-averaged spectrum. *Nature* 2018 555:7694, 555:67–70, 3 2018. ISSN 1476-4687. doi: 10.1038/nature25792. URL <https://www.nature.com/articles/nature25792>.

R Byrne and D Jacobs. Development of a High Throughput Cloud-Based Data Pipeline for 21 cm Cosmology. *Astron. Comput.*, 34:100447, 2021. doi: 10.1016/j.ascom.2021.100447.

R Byrne, M F Morales, B Hazelton, W Li, N Barry, A P Beardsley, R Joseph, J Pober, I Sullivan, and C Trott. Fundamental Limitations on the Calibration of Redundant 21 cm Cosmology Instruments and Implications for HERA and the SKA. *Astrophys. J.*, 875:70, 2019. doi: 10.3847/1538-4357/ab107d.

R Byrne, M F Morales, B J Hazelton, and M Wilensky. A Unified Calibration Framework for 21 cm Cosmology. *Mon. Not. R. Astron. Soc.*, 503:2457–2477, 2021a. doi: 10.1093/mnras/stab647.

Ruby Byrne. The Calibration Catalog, Accessed June 4 2024. URL https://github.com/EoRImaging/Memos/blob/main/PDFs/009_Calibration_Catalog.pdf.

Ruby Byrne, Miguel F. Morales, Bryna Hazelton, Ian Sullivan, Nichole Barry, Christene

Lynch, Jack L.B. Line, and Daniel C. Jacobs. A Map of Diffuse Radio Emission at 182 MHz to Enhance Epoch of Reionization Observations in the Southern Hemisphere. *Monthly Notices of the Royal Astronomical Society*, 510:2011–2024, 7 2021b. ISSN 13652966. doi: 10.1093/mnras/stab3276. URL <https://arxiv.org/abs/2107.11487v2>.

M. A. Clark, P. C. La Plante, and L. J. Greenhill. Accelerating radio astronomy cross-correlation with graphics processing units. <http://dx.doi.org/10.1177/1094342012444794>, 27:178–192, 5 2012. ISSN 10943420. doi: 10.1177/1094342012444794. URL <https://journals.sagepub.com/doi/10.1177/1094342012444794>.

The HERA Collaboration, Zara Abdurashidova, Tyrone Adams, James E. Aguirre, Paul Alexander, Zaki S. Ali, Rushelle Baartman, Yanga Balfour, Rennan Barkana, Adam P. Beardsley, Gianni Bernardi, Tashalee S. Billings, Judd D. Bowman, Richard F. Bradley, Daniela Breitman, Philip Bull, Jacob Burba, Steve Carey, Chris L. Carilli, Carina Cheng, Samir Choudhuri, David R. DeBoer, Eloy de Lera Acedo, Matt Dexter, Joshua S. Dillon, John Ely, Aaron Ewall-Wice, Nicolas Fagnoni, Anastasia Fialkov, Randall Fritz, Steven R. Furlanetto, Kingsley Gale-Sides, Hugh Garsden, Brian Glendenning, Adélie Gorce, Deepthi Gorthi, Bradley Greig, Jasper Grobbelaar, Ziyaad Halday, Bryna J. Hazelton, Stefan Heimersheim, Jacqueline N. Hewitt, Jack Hickish, Daniel C. Jacobs, Austin Julius, Nicholas S. Kern, Joshua Kerrigan, Piyanat Kittiwisit, Saul A. Kohn, Matthew Kolopanis, Adam Lanman, Paul La Plante, David Lewis, Adrian Liu, Anita Loots, Yin-Zhe Ma, David H. E. MacMahon, Lourence Malan, Keith Malgas, Cresshim Malgas, Matthys Maree, Bradley Marero, Zachary E. Martinot, Lisa McBride, Andrei Mesinger, Jordan Mirocha, Mathakane Molewa, Miguel F. Morales, Tshegofalang Mosiane, Julian B. Muñoz, Steven G. Murray, Vighnesh Nagpal, Abraham R. Neben, Bojan Nikolic, Chuneeta D. Nunhokee, Hans Nuwegeld, Aaron R. Parsons, Robert Pascua, Nipanjana Patra, Samantha Pieterse, Yuxiang Qin, Nima Razavi-Ghods, James Robnett, Kathryn Rosie, Mario G. Santos, Peter Sims, Saurabh Singh, Craig Smith, Hilton Swarts, Jianrong Tan, Nithyanandan Thyagarajan, Michael J. Wilensky, Peter K. G. Williams, Pieter van Wyngaarden,

and Haoxuan Zheng. Improved Constraints on the 21 cm EoR Power Spectrum and the X-Ray Heating of the IGM with HERA Phase I Observations. *The Astrophysical Journal*, 945:124, 10 2022a. ISSN 0004-637X. doi: 10.3847/1538-4357/acaf50. URL <https://arxiv.org/abs/2210.04912v2>.

The HERA Collaboration, Zara Abdurashidova, James E. Aguirre, Paul Alexander, Zaki S. Ali, Yanga Balfour, Adam P. Beardsley, Gianni Bernardi, Tashalee S. Billings, Judd D. Bowman, Richard F. Bradley, Philip Bull, Jacob Burba, Steve Carey, Chris L. Carilli, Carina Cheng, David R. DeBoer, Matt Dexter, Eloy de Lera Acedo, Taylor Dibblee-Barkman, Joshua S. Dillon, John Ely, Aaron Ewall-Wice, Nicolas Fagnoni, Randall Fritz, Steven R. Furlanetto, Kingsley Gale-Sides, Brian Glendenning, Deepthi Gorthi, Bradley Greig, Jasper Grobbelaar, Ziyaad Halday, Bryna J. Hazelton, Jacqueline N. Hewitt, Jack Hickish, Daniel C. Jacobs, Austin Julius, Nicholas S. Kern, Joshua Kerrigan, Piyanat Kittiwisit, Saul A. Kohn, Matthew Kolopanis, Adam Lanman, Paul La Plante, Telalo Lekalake, David Lewis, Adrian Liu, David MacMahon, Lourence Malan, Cresshim Malgas, Matthys Maree, Zachary E. Martinot, Eunice Matsetela, Andrei Mesinger, Mathakane Molewa, Miguel F. Morales, Tshegofalang Mosiane, Steven G. Murray, Abraham R. Neben, Bojan Nikolic, Chuneeta D. Nunhokee, Aaron R. Parsons, Nipanjana Patra, Robert Pascua, Samantha Pieterse, Jonathan C. Pober, Nima Razavi-Ghods, Jon Ringuette, James Robnett, Kathryn Rosie, Peter Sims, Saurabh Singh, Craig Smith, Angelo Syce, Nithyanandan Thyagarajan, Peter K. G. Williams, and Haoxuan Zheng. First Results from HERA Phase I: Upper Limits on the Epoch of Reionization 21 cm Power Spectrum. *The Astrophysical Journal*, 925:221, 2 2022b. ISSN 0004-637X. doi: 10.3847/1538-4357/ac1c78. URL <https://ui.adsabs.harvard.edu/abs/2022ApJ...925..221A/abstract>.

Brian Crosse and Andrew Williams, Accessed June 9 2024. URL <https://github.com/RadioAstronomySoftwareGroup/pyuvdata/pull/908>.

A Datta, J D Bowman, and C L Carilli. Bright Source Subtraction Requirements for Red-

shifted 21 cm Measurements. *Astrophys. J.*, 724:526–538, 2010. doi: 10.1088/0004-637X/724/1/526.

Joshua S. Dillon, Abraham R. Neben, Jacqueline N. Hewitt, Max Tegmark, N. Barry, A. P. Beardsley, J. D. Bowman, F. Briggs, P. Carroll, A. De Oliveira-Costa, A. Ewall-Wice, L. Feng, L. J. Greenhill, B. J. Hazelton, L. Hernquist, N. Hurley-Walker, D. C. Jacobs, H. S. Kim, P. Kittiwisit, E. Lenc, J. Line, A. Loeb, B. McKinley, D. A. Mitchell, M. F. Morales, A. R. Offringa, S. Paul, B. Pindor, J. C. Pober, P. Procopio, J. Riding, S. Sethi, N. Udaya Shankar, R. Subrahmanyam, I. Sullivan, Nithyanandan Thyagarajan, S. J. Tingay, C. Trott, R. B. Wayth, R. L. Webster, S. Wyithe, G. Bernardi, R. J. Cappallo, A. A. Deshpande, M. Johnston-Hollitt, D. L. Kaplan, C. J. Lonsdale, S. R. McWhirter, E. Morgan, D. Oberoi, S. M. Ord, T. Prabu, K. S. Srivani, A. Williams, and C. L. Williams. Empirical covariance modeling for 21 cm power spectrum estimation: A method demonstration and new limits from early Murchison Widefield Array 128-tile data. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 91:123011, 6 2015. ISSN 15502368. doi: 10.1103/PHYSREVD.91.123011/FIGURES/9/MEDIUM. URL <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.91.123011>.

M W Eastwood, M M Anderson, R M Monroe, G Hallinan, B R Barsdell, S A Bourke, M A Clark, S W Ellingson, J Dowell, H Garsden, L J Greenhill, J M Hartman, J Kocz, T J W Lazio, D C Price, F K Schinzel, G B Taylor, H K Vedantham, Y Wang, and D P Woody. The Radio Sky at Meter Wavelengths: m-Mode Analysis Imaging with the OVRO-LWA. *Astron. J.*, 156:32, 2018. doi: 10.3847/1538-3881/aac721.

A Ewall-Wice, J S Dillon, A Liu, and J Hewitt. The Impact of Modeling Errors on Interferometer Calibration for 21 cm Power Spectra. *Mon. Not. R. Astron. Soc.*, 470:1849–1870, 10 2016a. doi: 10.1093/mnras/stx1221.

A. Ewall-Wice, Joshua S. Dillon, J. N. Hewitt, A. Loeb, A. Mesinger, A. R. Neben, A. R. Offringa, M. Tegmark, N. Barry, A. P. Beardsley, G. Bernardi, Judd D. Bowman, F. Briggs,

- R. J. Cappallo, P. Carroll, B. E. Corey, A. de Oliveira-Costa, D. Emrich, L. Feng, B. M. Gaensler, R. Goeke, L. J. Greenhill, B. J. Hazelton, N. Hurley-Walker, M. Johnston-Hollitt, Daniel C. Jacobs, D. L. Kaplan, J. C. Kasper, H. S. Kim, E. Kratzenberg, E. Lenc, J. Line, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, Nithyanandan Thyagarajan, D. Oberoi, S. M. Ord, S. Paul, B. Pindor, J. C. Pober, T. Prabu, P. Procopio, J. Riding, A. E.E. Rogers, A. Roshi, N. Udaya Shankar, Shiv K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, S. J. Tingay, C. M. Trott, M. Waterson, R. B. Wayth, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, C. Wu, and J. S.B. Wyithe. First Limits on the 21 cm Power Spectrum during the Epoch of X-ray heating. *Monthly Notices of the Royal Astronomical Society*, 460:4320–4347, 4 2016b. ISSN 13652966. doi: 10.1093/mnras/stw1022. URL <https://arxiv.org/abs/1605.00016v2>.
- Steven R Furlanetto, S Peng Oh, and Frank H Briggs. Cosmology at low frequencies: The 21 cm transition and the high-redshift Universe. *Phys. Rep.*, 2006. ISSN 03701573. doi: 10.1016/j.physrep.2006.08.002.
- K M Gorski, 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. *Astrophys. J.*, 622:759–771, 2005. doi: 10.1086/427976.
- Eric W Greisen. AIPS FITS File Format, Accessed June 5 2024. URL https://library.nrao.edu/public/memos/aips/memos/AIPSM_117.pdf.
- Christopher Harris and Karen Haines. A Mathematical Review of Polyphase Filterbank Implementations for Radio Astronomy. *Publications of the Astronomical Society of Australia*, 28:317–322, 2011. ISSN 1448-6083. doi: 10.1071/AS11032.
- Bryna J Hazelton, Miguel F Morales, and Ian S Sullivan. The Fundamental Multi-Baseline Mode-Mixing Foreground in 21 cm Epoch of Reionization Observations. *Astrophys. J.*, 770:156, 2013. doi: 10.1088/0004-637X/770/2/156.

Bryna J. Hazelton, Daniel C. Jacobs, Jonathan C. Pober, and Adam P. Beardsley. `pyuvdata`: An Interface for Astronomical Interferometric Datasets in Python. *J. Open Source Softw.*, 2:140, 2 2017. doi: 10.21105/joss.00140.

N Hurley-Walker, J R Callingham, P J Hancock, T M O Franzen, L Hindson, A D Kapińska, J Morgan, A R Offringa, R B Wayth, C Wu, Q Zheng, T Murphy, M E Bell, K S Dwarkanath, B For, B M Gaensler, M Johnston-Hollitt, E Lenc, P Procopio, L Staveley-Smith, R Ekers, J D Bowman, F Briggs, R J Cappallo, A A Deshpande, L Greenhill, B J Hazelton, D L Kaplan, C J Lonsdale, S R McWhirter, D A Mitchell, M F Morales, E Morgan, D Oberoi, S M Ord, T Prabu, N Udaya Shankar, K S Srivani, R Subrahmanyan, S J Tingay, R L Webster, A Williams, and C L Williams. GaLactic and Extragalactic All-Sky Murchison Widefield Array (GLEAM) Survey - I. A Low-Frequency Extragalactic Catalogue. *Mon. Not. R. Astron. Soc.*, 464:1146–1167, 2017. doi: 10.1093/mnras/stw2337.

Jake Jones. RRI Receiver PFB Filter, Accessed May 7 2024. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24972979/RRI+Receiver+PFB+Filter>.

C H Jordan, S Murray, C M Trott, R B Wayth, D A Mitchell, M Rahimi, B Pindor, P Procopio, and J Morgan. Characterization of the Ionosphere Above the Murchison Radio Observatory Using the Murchison Widefield Array. *Mon. Not. R. Astron. Soc.*, 471:3974–3987, 2017. doi: 10.1093/mnras/stx1797.

Matthew Kolopanis, Daniel C Jacobs, Carina Cheng, Aaron R Parsons, Saul A Kohn, Jonathan C Pober, James E Aguirre, Zaki S Ali, Gianni Bernardi, Richard F Bradley, Chris L Carilli, David R DeBoer, Matthew R Dexter, Joshua S Dillon, Joshua Kerrigan, Pat Klima, Adrian Liu, David H E MacMahon, David F Moore, Nithyanandan Thyagarajan, Chuneeta D Nunhokee, William P Walbrugh, and Andre Walker. A Simplified, Lossless Reanalysis of PAPER-64. *Astrophys. J.*, 883, 2019. doi: 10.3847/1538-4357/ab3e3a.

Alan M Levine. The Effects of the Two-Stage PFB Architecture on the Response of the MWA, 2012a. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/>

24971988/Memos?preview=/24971988/24994523/MEMO_CascadedFT_2012_05_25.pdf.

Accessed June 5 2024.

Alan M Levine. An Analysis of Two Cascaded Fourier Transforms and Application to the Coarse and Fine PFBs for the MWA, 2012b. URL https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24971988/Memos?preview=/24971988/24994523/MEMO_CascadedFT_2012_05_25.pdf. Accessed June 5 2024.

W Li, J C Pober, B J Hazelton, N Barry, M F Morales, I Sullivan, A R Parsons, Z S Ali, J S Dillon, A P Beardsley, J D Bowman, F Briggs, R Byrne, P Carroll, B Crosse, D Emrich, A Ewall-Wice, L Feng, T M O Franzen, J N Hewitt, L Horsley, D C Jacobs, M Johnston-Hollitt, C Jordan, R C Joseph, D L Kaplan, D Kenney, H Kim, P Kittiwisit, A Lanman, J Line, B McKinley, D A Mitchell, S Murray, A Neben, A R Offringa, D Pallot, S Paul, B Pindor, P Procopio, M Rahimi, J Riding, S K Sethi, N Udaya Shankar, K Steele, R Subrahmanian, M Tegmark, N Thyagarajan, S J Tingay, C Trott, M Walker, R B Wayth, R L Webster, A Williams, C Wu, and S Wyithe. Comparing Redundant and Sky-model-based Interferometric Calibration: A First Look with Phase II of the MWA. *Astrophys. J.*, 863:170, 2018. doi: 10.3847/1538-4357/aad3c3.

W Li, J C Pober, N Barry, B J Hazelton, M F Morales, C M Trott, A Lanman, M Wilensky, I Sullivan, A P Beardsley, T Boller, J D Bowman, R Byrne, B Crosse, D Emrich, T M O Franzen, K Hasegawa, L Horsley, M Johnston-Hollitt, D C Jacobs, C H Jordan, R C Joseph, T Kaneuji, D L Kaplan, D Kenney, K Kubota, J Line, C Lynch, B McKinley, D A Mitchell, S Murray, D Pallot, B Pindor, M Rahimi, J Riding, G Slep, K Steele, K Takahashi, S J Tingay, M Walker, R B Wayth, R L Webster, A Williams, C Wu, J S B Wyithe, S Yoshiura, and Q Zheng. First Season MWA Phase II Epoch of Reionization Power Spectrum Results at Redshift 7. *Astrophys. J.*, 887:141, 2019. doi: 10.3847/1538-4357/ab55e4.

- Adrian Liu and J Richard Shaw. Data Analysis for Precision 21cm Cosmology. *Publ. Astron. Soc. Pacific*, 132, 2020. doi: 10.1088/1538-3873/ab5bfd.
- N. R. Lomb, Lomb, and N. R. Least-Squares Frequency Analysis of Unequally Spaced Data. *Astrophysics and Space Science*, 39:447–462, 2 1976. ISSN 0004-640X. doi: 10.1007/BF00648343. URL <https://ui.adsabs.harvard.edu/abs/1976Ap&SS..39.447L/abstract>.
- C J Lonsdale, R J Cappallo, M F Morales, F H Briggs, Leonid Benkevitch, J D Bowman, J D Bunton, Steven Burns, B E Corey, L DeSouza, S S Doeleman, Mark Derome, Avinash Deshpande, M R Gopala, L J Greenhill, D E Herne, J N Hewitt, P A Kamini, J C Kasper, B B Kincaid, Jonathan Kocz, Errol Kowald, Eric Kratzenberg, Deepak Kumar, M J Lynch, S Madhavi, Michael Matejek, D A Mitchell, Edward Morgan, Divya Oberoi, Steven Ord, Joseph Pathikulangara, T Prabu, A Rogers, Anish Roshni, J E Salah, R J Sault, N U Shankar, K S Srivani, Jamie Stevens, Steven Tingay, Annino Vaccarella, Mark Waterson, R B Wayth, R L Webster, A R Whitney, A Williams, and Christopher Williams. The Murchison Widefield Array: Design Overview. *Proc. IEEE*, 97:1497–1506, 2009. doi: 10.1109/JPROC.2009.2017564.
- S. J. McSweeney, S. M. Ord, D Kaur, N. D. R. Bhat, B. W. Meyers, S. E. Tremblay, J Jones, B Crosse, and K. R. Smith. Mwa tied-array processing iii: Microsecond time resolution via a polyphase synthesis filter. *Publications of the Astronomical Society of Australia*, 37: e034, 8 2020. doi: 10.1017/pasa.2020.24.
- F G Mertens, M Mevius, L V E Koopmans, A R Offringa, G Mellema, S Zaroubi, M A Brentjens, H Gan, B K Gehlot, V N Pandey, A M Sardarabadi, H K Vedantham, S Yatawatta, K M B Asad, B Ciardi, E Chapman, S Gazagnes, R Ghara, A Ghosh, S K Giri, I T Iliev, V Jelić, R Kooistra, R Mondal, J Schaye, and M B Silva. Improved Upper Limits on the 21 cm Signal Power Spectrum of Neutral Hydrogen at z 9.1 from LOFAR. *Mon. Not. R. Astron. Soc.*, 493:1662–1685, 2020. doi: 10.1093/mnras/staa327.

- Miguel F Morales, Bryna Hazelton, Ian Sullivan, and Adam Beardsley. Four Fundamental Foreground Power Spectrum Shapes for 21cm Cosmology Observations. *Astrophys. J.*, 752:137, 2012. doi: 10.1088/0004-637X/752/2/137.
- I. S. Morrison, B. Crosse, G. Slep, R. B. Wayth, A. Williams, M. Johnston-Hollitt, J. Jones, S. J. Tingay, M. Walker, and L. Williams. MWAX: A New Correlator for the Murchison Widefield Array. *Publications of the Astronomical Society of Australia*, 40:1–22, 3 2023. ISSN 14486083. doi: 10.1017/pasa.2023.15. URL <https://arxiv.org/abs/2303.11557v1>.
- S. M. Ord, B. Crosse, D. Emrich, D. Pallot, R. B. Wayth, M. A. Clark, S. E. Tremblay, W. Arcus, D. Barnes, M. Bell, G. Bernardi, N. D. R. Bhat, J. D. Bowman, F. Briggs, J. D. Bunton, R. J. Cappallo, B. E. Corey, A. A. Deshpande, L. deSouza, A. Ewell-Wice, L. Feng, R. Goeke, L. J. Greenhill, B. J. Hazelton, D. Herne, J. N. Hewitt, L. Hindson, H. Hurley-Walker, D. Jacobs, M. Johnston-Hollitt, D. L. Kaplan, J. C. Kasper, B. B. Kincaid, R. Koenig, E. Kratzenberg, N. Kudryavtseva, E. Lenc, C. J. Lonsdale, M. J. Lynch, B. McKinley, S. R. McWhirter, D. A. Mitchell, M. F. Morales, E. Morgan, D. Oberoi, A. Offringa, J. Pathikulangara, B. Pindor, T. Prabu, P. Procopio, R. A. Remillard, J. Riding, A. E. E. Rogers, A. Roshi, J. E. Salah, R. J. Sault, N. Udaya Shankar, K. S. Srivani, J. Stevens, R. Subrahmanyan, S. J. Tingay, M. Waterston, R. L. Webster, A. R. Whitney, A. Williams, C. L. Williams, and J. S. B. Wyithe. The Murchison Widefield Array Correlator. *Publications of the Astronomical Society of Australia*, 1 2015. doi: 10.1017/pasa.2015.5. URL <http://arxiv.org/abs/1501.05992><http://dx.doi.org/10.1017/pasa.2015.5>.
- 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 21cm Cosmic Reionization Signature. *Astrophys. J.*, 756:165, 2012. doi: 10.1088/0004-637X/756/2/165.

- Thiagaraj Prabu, K. S. Srivani, D Anish Roshi, P. A. Kamini, S Madhavi, David Emrich, Brian Crosse, Andrew J Williams, Mark Waterson, Avinash A Deshpande, N Udaya Shankar, Ravi Subrahmanyam, Frank H Briggs, Robert F Goeke, Steven J Tingay, Melanie Johnston-Hollitt, Gopalakrishna M R, Edward H Morgan, Joseph Pathikulangara, John D Bunton, Grant Hampson, Christopher Williams, Stephen M Ord, Randall B Wayth, Deepak Kumar, Miguel F Morales, Ludi deSouza, Eric Kratzenberg, D Pallot, Russell McWhirter, Bryna J Hazelton, Wayne Arcus, David G Barnes, Gianni Bernardi, T Boole, Judd D Bowman, Roger J Cappallo, Brian E Corey, Lincoln J Greenhill, David Herne, Jacqueline N Hewitt, David L Kaplan, Justin C Kasper, Barton B Kincaid, Ronald Koenig, Colin J Lonsdale, Mervyn J Lynch, Daniel A Mitchell, Divya Oberoi, Ronald A Remillard, Alan E Rogers, Joseph E Salah, Robert J Sault, Jamie B Stevens, S Tremblay, Rachel L Webster, Alan R Whitney, and Stuart B Wyithe. A digital-receiver for the Murchison Widefield Array. *Experimental Astronomy*, 39:73–93, 3 2015. doi: 10.1007/s10686-015-9444-3.
- Danny C. Price. Spectrometers and Polyphase Filterbanks in Radio Astronomy. In *The WSPC Handbook of Astronomical Instrumentation, Volume 1: Radio Astronomical Instrumentation*, pages 159–179. World Scientific, 7 2016. doi: 10.1142/9789811203770_0007.
- M. Rahimi, B. Pindor, J. L.B. Line, N. Barry, C. M. Trott, R. L. Webster, C. H. Jordan, M. Wilensky, S. Yoshiura, A. Beardsley, J. Bowman, R. Byrne, A. Chokshi, B. J. Hazelton, K. Hasegawa, E. Howard, B. Greig, D. Jacobs, R. Joseph, M. Kolopanis, C. Lynch, B. McKinley, D. A. Mitchell, S. Murray, M. F. Morales, J. C. Pober, K. Takahashi, S. J. Tingay, R. B. Wayth, J. S.B. Wyithe, and Q. Zheng. Epoch of Reionization Power Spectrum Limits from Murchison Widefield Array Data Targeted at EoR1 Field. *Monthly Notices of the Royal Astronomical Society*, 508:5954–5971, 10 2021. ISSN 13652966. doi: 10.1093/mnras/stab2918. URL <https://arxiv.org/abs/2110.03190v1>.
- Stefano Salvini and Stefan J. Wijnholds. Fast gain calibration in radio astronomy using alternating direction implicit methods: Analysis and applications. *Astronomy and Astrophysics*, 571:A97, 11 2014. ISSN 0004-6361. doi: 10.1051/0004-6361/201424487. URL

https://www.aanda.org/articles/aa/full_html/2014/11/aa24487-14/aa24487-14.html
<https://www.aanda.org/articles/aa/abs/2014/11/aa24487-14/aa24487-14.html>.

J. D. Scargle, Scargle, and J. D. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, 12 1982. ISSN 0004-637X. doi: 10.1086/160554. URL <https://ui.adsabs.harvard.edu/abs/1982ApJ...263..835S/abstract>.

J L Sievers. Calibration of Quasi-Redundant Interferometers. *arXiv e-prints*, 2017.

Greg Slep and Dev Null. MWA Fine Channel Centre Frequencies, Accessed May 7 2024. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24972939/MWA+Fine+Channel+Centre+Frequencies>.

Greg Slep, Brian Crosse, and Ian Morrison. MWAX Visibility File Format, Accessed May 7 2024. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24970474/MWAX+Visibility+File+Format>.

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. *Astrophys. J.*, 759:17, 2012. doi: 10.1088/0004-637X/759/1/17.

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 Rosh, 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 HI from the Epoch of Reionization. *Astrophys. J.*, 776:6, 2013. doi: 10.1088/0004-637X/776/1/6.
- 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 Wicenc, N Kudryavtseva, W Arcus, D Barnes, G Bernardi, F Briggs, S Burns, J D Bunton, R J Cappallo, B E Corey, A Deshpande, L Desouza, B M Gaensler, L J Greenhill, P J Hall, B J Hazelton, D Herne, J N Hewitt, M Johnston-Hollitt, D L Kaplan, J C Kasper, B B Kincaid, R Koenig, E Kratzenberg, M J Lynch, B McKinley, S R McWhirter, E Morgan, D Oberoi, J Pathikulangara, T Prabu, R A Remillard, A E E Rogers, A Rosh, J E Salah, R J Sault, N Udaya-Shankar, F Schlagenhauser, K S Srivani, J Stevens, R Subrahmanyam, M Waterson, R L Webster, A R Whitney, A Williams, C L Williams, and J S B Wyithe. The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies. *Publ. Astron. Soc. Aust.*, 30:E007, 2013. doi: 10.1017/pasa.2012.007.
- C. M. Trott, B. Pindor, P. Procopio, R. B. Wayth, D. A. Mitchell, B. McKinley, S. J. Tingay, N. Barry, A. P. Beardsley, G. Bernardi, Judd D. Bowman, F. Briggs, R. J. Cappallo, P. Carroll, A. de Oliveira-Costa, Joshua S. Dillon, A. Ewall-Wice, L. Feng, L. J. Greenhill, B. J. Hazelton, J. N. Hewitt, N. Hurley-Walker, M. Johnston-

- Hollitt, Daniel C. Jacobs, D. L. Kaplan, H. S. Kim, E. Lenc, J. Line, A. Loeb, C. J. Lonsdale, M. F. Morales, E. Morgan, A. R. Neben, Nithyanandan Thyagarajan, D. Oberoi, A. R. Offringa, S. M. Ord, S. Paul, J. C. Pober, T. Prabu, J. Riding, N. Udaya Shankar, Shiv K. Sethi, K. S. Srivani, R. Subrahmanyam, I. S. Sullivan, M. Tegmark, R. L. Webster, A. Williams, C. L. Williams, C. Wu, and J. S. B. Wyithe. CHIPS: THE COSMOLOGICAL H α POWER SPECTRUM ESTIMATOR. *The Astrophysical Journal*, 818:139, 2 2016. ISSN 0004-637X. doi: 10.3847/0004-637X/818/2/139. URL <https://iopscience.iop.org/article/10.3847/0004-637X/818/2/139><https://iopscience.iop.org/article/10.3847/0004-637X/818/2/139/meta>.
- Cathryn M Trott and Randall B Wayth. Spectral Calibration Requirements of Radio Interferometers for Epoch of Reionisation Science with the SKA. *Publ. Astron. Soc. Aust.*, 33: E019, 2016. doi: 10.1017/pasa.2016.18.
- Cathryn M Trott, Randall B Wayth, and Steven J Tingay. The Impact of Point-Source Subtraction Residuals on 21cm Epoch of Reionization Estimation. *Astrophys. J.*, 757:101, 2012. doi: 10.1088/0004-637X/757/1/101.
- Cathryn M Trott, C H Jordan, S Midgley, N Barry, B Greig, B Pindor, J H Cook, G Sleap, S J Tingay, D Ung, P Hancock, A Williams, J Bowman, R Byrne, A Chokshi, B J Hazelton, K Hasegawa, D Jacobs, R C Joseph, W Li, J L B Line, C Lynch, B McKinley, D A Mitchell, M F Morales, M Ouchi, J C Pober, M Rahimi, K Takahashi, R B Wayth, R L Webster, M Wilensky, J S B Wyithe, S Yoshiura, Z Zhang, and Q Zheng. Deep Multiredshift Limits on Epoch of Reionization 21 cm Power Spectra from Four Seasons of Murchison Widefield Array Observations. *Mon. Not. R. Astron. Soc.*, 493, 2020. doi: 10.1093/mnras/staa414.
- J. H. Van Vleck and David Middleton. The Spectrum of Clipped Noise. *Proceedings of the IEEE*, 54, 1966. ISSN 15582256. doi: 10.1109/PROC.1966.4567.
- Jacob T VanderPlas. Understanding the Lomb-Scargle Periodogram. *The Astrophysical Journal Supplement Series*, 236:16, 5 2018. doi: 10.3847/1538-4365/aab766.

F. Di Vruno, B. Winkel, C. G. Bassa, G. I.G. Józsa, M. A. Brentjens, A. Jessner, and S. Garrington. Unintended electromagnetic radiation from Starlink satellites detected with LOFAR between 110 and 188 MHz. *Astronomy and Astrophysics*, 676:A75, 8 2023. ISSN 0004-6361. doi: 10.1051/0004-6361/202346374. URL https://www.aanda.org/articles/aa/full_html/2023/08/aa46374-23/aa46374-23.html<https://www.aanda.org/articles/aa/abs/2023/08/aa46374-23/aa46374-23.html>.

Randall B Wayth, Steven J Tingay, Cathryn M Trott, David Emrich, Melanie Johnston-Hollitt, Ben McKinley, B M Gaensler, A P Beardsley, T Booler, B Crosse, T M O Franzen, L Horsley, D L Kaplan, D Kenney, M F Morales, D Pallot, G Sleap, K Steele, M Walker, A Williams, C Wu, Iver H Cairns, M D Filipovic, S Johnston, T Murphy, P Quinn, L Staveley-Smith, R Webster, and J S B Wyithe. The Phase II Murchison Widefield Array: Design Overview. *Publ. Astron. Soc. Aust.*, 35:E033, 2018. doi: 10.1017/pasa.2018.37.

Mark H Wieringa. An Investigation of the Telescope Based Calibration Methods ‘Redundancy’ and ‘Self-Cal’. *Exp. Astron.*, 2:203–225, 1992. doi: 10.1007/BF00420576.

Michael J Wilensky, Miguel F Morales, Bryna J Hazelton, Nichole Barry, Ruby Byrne, and Sumit Roy. Absolving the SSINS of Precision Interferometric Radio Data: A New Technique for Mitigating Faint Radio Frequency Interference. *Publ. Astron. Soc. Pacific*, 131:114507, 2019. doi: 10.1088/1538-3873/ab3cad.

Michael J Wilensky, Nichole Barry, Miguel F Morales, Bryna J Hazelton, and Ruby Byrne. Quantifying Excess Power from Radio Frequency Interference in Epoch of Reionization Measurements. *Mon. Not. R. Astron. Soc.*, 498:265–275, 2020. doi: 10.1093/mnras/staa2442.

Michael J. Wilensky, Bryna J. Hazelton, and Miguel F. Morales. Exploring the consequences of chromatic data excision in 21-cm epoch of reionization power spectrum observations. *Monthly Notices of the Royal Astronomical Society*, 510:5023–5034, 1 2022. ISSN 0035-

8711. doi: 10.1093/MNRAS/STAB3456. URL <https://dx.doi.org/10.1093/mnras/stab3456>.

Michael J. Wilensky, Miguel F. Morales, Bryna J. Hazelton, Pyxie L. Star, Nichole Barry, Ruby Byrne, C. H. Jordan, Daniel C. Jacobs, Jonathan C. Pober, and C. M. Trott. Evidence of Ultra-faint Radio Frequency Interference in Deep 21 cm Epoch of Reionization Power Spectra with the Murchison Widefield Array. *The Astrophysical Journal*, 957:78, 10 2023. doi: 10.3847/1538-4357/acffbd. URL <http://arxiv.org/abs/2310.03851><https://dx.doi.org/10.3847/1538-4357/acffbd>.

Andrew Williams. MWAX correlator and VCS, Accessed June 6 2024. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24970246/MWAX+Correlator+and+VCS>.

Andrew Williams. Metafits files, Accessed May 7 2024a. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24969594/Metafits+files>.

Andrew Williams. MWA 'sweet spots' and gridpoint numbers, Accessed May 7 2024b. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24969505/MWA+sweet+spots+and+gridpoint+numbers>.

Andrew Williams and Greg Slep. MWAX Metafits Changes, Accessed May 7 2024. URL <https://mwatelescope.atlassian.net/wiki/spaces/MP/pages/24969941/MWAX+Metafits+Changes>.