The Omniscope: Mapping the Universe in 3D With Neutral Hydrogen

by

Ashley Nicole Perko

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

Bachelor of Science in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

Author ....
.....
Department of Physics
May 6, 2011

Certified by
.....
Max Tegmark
Professor of Physics
Thesis Supervisor

Accepted by .....
Professor Nergis Mavalvala
Senior Thesis Coordinator, Department of Physics

The Omniscope: Mapping the Universe in 3D With Neutral Hydrogen

by

Ashley Nicole Perko

Submitted to the Department of Physics
on May 6, 2011, in partial fulfillment of the
requirements for the degree of
Bachelor of Science in Physics

Abstract

21 cm tomography has the potential to become the most powerful cosmological probe yet. The Omniscope is a novel radio telescope being built to take advantage of this signal. This thesis describes my work on integrating, testing, and characterizing all modules of the Omniscope and identifying opportunities for further improving their sensitivity.

Thesis Supervisor: Max Tegmark
Title: Professor of Physics

Acknowledgments

I would like to thank Max, for taking a chance on me and helping me get so far; Adrian, for his advice and wisdom; all the Omniscopers including Jon, Nevada, Eben, Josh, Hrant, and Katelin for being a great team; Dave Zou, for making psets fun; Mike and Marianne for giving me a surrogate home; Mom and Dad for their constant support; Will for making it all worth it; and Phil for being awesome. Most of all, this is for Nana.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 A (Very) Brief History of Time

The early universe shortly after the Big Bang was composed of hot, dense plasma and photons. In this state, photons could not travel far before being scattered, and the universe was opaque. However, as the universe expanded it reached a period called recombination. At that time, the universe had cooled to the point where neutral hydrogen could form. The photons that had formerly scattered off of charged particles in the plasma were free to travel past the neutral hydrogen atoms. We see the photons from this surface of last scattering in the Cosmic Microwave Background (CMB). Much of cosmology to this point has relied on the CMB. The CMB has been mapped to great precision by WMAP [1], but the CMB is only a picture of the universe at one instant of time. Following recombination, the neutral hydrogen remained mostly neutral for about 150 million years until the period of reionization. Reionization occurred when the first stars and quasars formed and the radiation they emitted began to ionize the hydrogen around them. The period between recombination and reionization is called the cosmic dark ages because the universe emitted no light during this time.

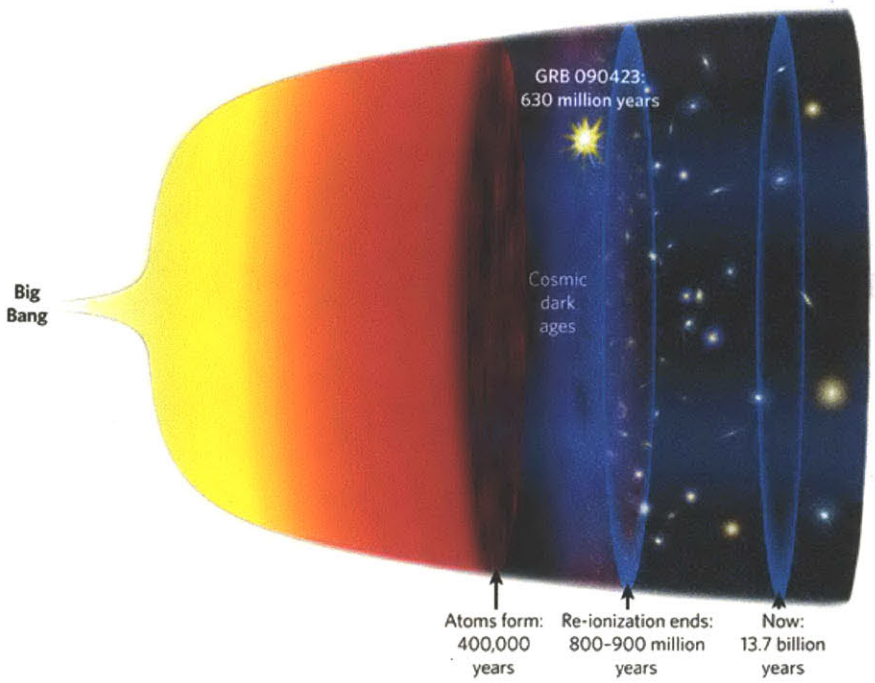

Figure 1-1: A schematic picture of the evolution of the universe, from [2]

1.2 21 Centimeter Cosmology

An emerging technique in cosmology that should be able to shed light on the dark ages for the first time is 21 centimeter cosmology, which uses the radiation emitted by neutral hydrogen at a wavelength of 21 cm to map the universe at redshifts that have never been probed. The basis of this technique is the fact that the electron and the proton, the two constituents of the hydrogen atom, both have intrinsic spin. The spins of the proton and the electron can be either aligned or anti-aligned. The latter state has lower energy so it is the preferred orientation but there is a small probability that the atom will be in the aligned state. When the atom transitions to the lower energy state, it releases a photon with wavelength of 21 cm.

The expansion of the universe after a particular 21 cm photon has been emitted stretches out the wavelength of the photon, so photons emitted by neutral hydrogen at different times in the evolution of the universe will end up with different wavelengths. When visible light is stretched out by the expansion of the universe, the

longer wavelength λ makes the light look red, so astronomers term the wavelength change as the *redshift*, denoted z:

$$z = \frac{\lambda_{now}}{\lambda_{emit}} - 1 \tag{1.1}$$

Since older photons are more redshifted, time is often referred to in terms of the redshift; today is redshift $z = 0$, the CMB was emitted at redshift $z \approx 1100$, and reionization is thought to have occurred between $z = 6$ and $z = 10$ [6]. 21 cm cosmology can map the universe across a large range of redshifts by tuning to different frequencies (Fig. 1-2). Because it can map a large fraction of the history of our universe, the 21 cm line offers a unique opportunity to probe the cosmic dark ages and the first structure formation. It has the potential to overtake the CMB as our most sensitive cosmological probe of the epoch of reionization, inflation, dark matter, dark energy, and neutrino masses [3].

1.3 The Omniscope

1.3.1 A Fast Fourier Transform Telescope

21-centimeter cosmology is an exciting development with the potential to map the structure of the universe over large periods of its evolution. The complication is that while hydrogen comprises a majority of the matter in the universe, the 21-centimeter transition is extraordinarily rare. This means that the cosmological signal is quite weak. In fact, it is about four orders of magnitude fainter than the signal in the same frequency range emitted by our own galaxy.

For such a faint signal, a very sensitive instrument is required. The expansion of the universe since the dark ages has redshifted its 21 cm photons to the approximate range 30-200 MHz, which is in the radio frequency range. The Omniscope focuses on the 100-200 MHz range, where the signal is expected to be easiest to detect. In traditional radio astronomy, signals are focused in a metal dish antenna. The sensitivity of the telescope scales with the area of the dish. Building a square kilometer

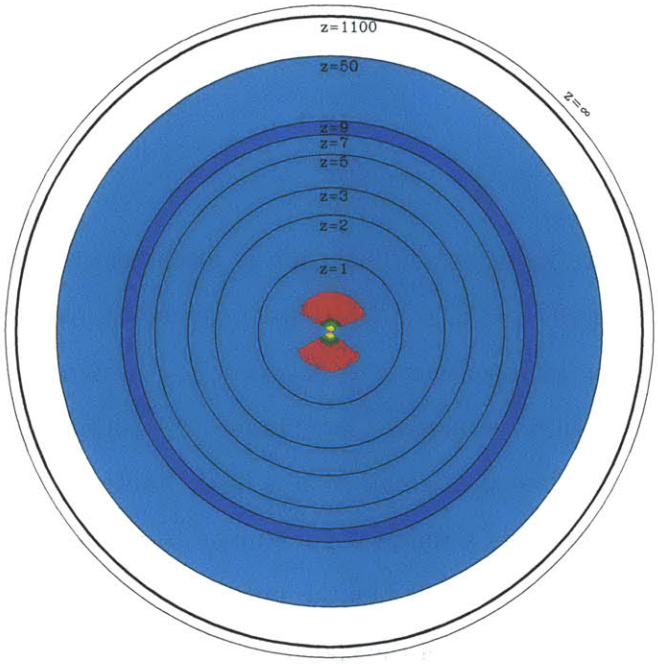

Figure 1-2: The 21 cm line can potentially map most of our observable universe (light blue/grey), whereas the CMB probes mainly a thin shell at $z \approx 1100$ and current large-scale structure maps map only small volumes near the center. Half of the comoving volume lies at $z > 29$. Even the convenient $7 < z < 9$ region (dark blue/grey) can eclipse the CMB in cosmological precision, probing the nature of neutrinos, dark energy, dark matter, reionization and early universe [3].

pointable dish would be prohibitively expensive with current technology. Moreover, a single dish can observe only a tiny fraction of the sky at any one time. This has led to the alternative approach known as interferometry. The signal is collected in an array of separate antennas and "focused" in software by a so-called correlator, which multiplies the signals from every two pairs of antennas and reconstructs the sky image. In a radio telescope array, the sensitivity increases with the number of antennas (the effective area of the array). In principle, the array could have an arbitrarily large number of antennas, so the limit of sensitivity lies in the computational cost of correlating the antennas. In traditional radio astronomy, the correlator works by multiplying each pair of signals together so the computational cost scales like N^2, where N is the number of antennas.

A key innovation of the Omniscope is to arrange the antennas in a rectangular

grid, which enables these correlations to be performed with Fast Fourier Transforms (FFTs), which cuts the computational cost scaling from N^2 to $N \log N$. Thus using an FFT correlator, the Omniscope can have greater sensitivity at a lower cost than traditional telescopes ([4], [5]).

1.3.2 Omniscope Data Analysis Chain

The data flow through the Omniscope is summarized in Fig. 1-3. A grid of dual-polarization antennas receives the sky signal, which is amplified and filtered in the analog chain and then converted to digital signals in the ADC. The F-engine performs a temporal Fourier transform and the X-engine correlates the signals between antennas using Fourier transforms. A small fraction of all antenna pairs is used for real-time calibration [7]. The correlated signals are then combined in software to produce 3D maps of the sky.

The work presented in this thesis was primarily focused on the first part of the data analysis chain depicted in Fig. 1-3: the analog chain. This thesis presents each of the modules in the analog chain (low noise amplifier, line driver, receiver board, and analog to digital converter), discussing the work done to characterize their performance and the improvements made. I will also discuss correlations produced by the digital correlator and how they can be used to perform an absolute calibration of the entire system.

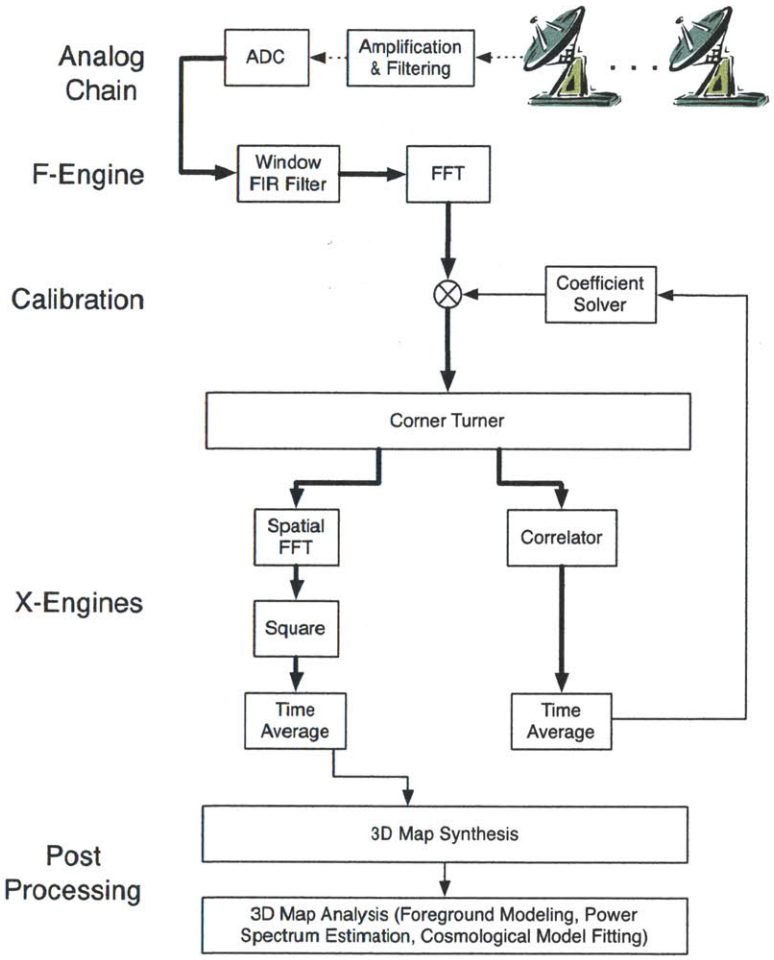

Figure 1-3: Omniscope data analysis pipeline

Chapter 2

Antenna and Low Noise Amplifier

2.1 Design

Figure 2-1: Array of MWA antennas in Green Bank, WV

The antennas are the first step in the path of the cosmological signal. It is important for a precision cosmological probe that the antennas be sensitive to the frequency band of the redshifted 21-cm signal that we are targeting (100-200 MHz) and that they be easily scalable so that the number of antennas can be large enough for high sensitivity. For the Omniscope we use the antennas designed by the Murchison Widefield Array experiment (MWA) shown in Fig. 2-1). The MWA is a 21-cm array led by MIT and an international consortium and is deployed in radio-quiet Western Australia. The MWA antennas are dual dipole antennas with integrated low noise

amplifiers (LNAs). The low noise amplifiers provide an amplification of approximately 20 dB over 80-300 MHz [8]. They are powered by DC biasing the antenna through the same 50Ω cable that carries the measured signals.

The MWA antennas suit the needs of this project because they can be mass-produced at low cost and include a low-noise amplifier (LNA) designed for a radio-quiet environment. The antennas are a dual-polarization bow-tie design. The bow-tie design was chosen instead of a simple dipole because it is sensitive over a wide frequency range and it has large angular coverage [8]. The bow-tie reduces the strong frequency dependence of the impedance so that the antenna impedance better matches the low noise amplifier impedance across the frequency range of interest. The open-frame design of the bow-tie means that it can be manufactured at low cost (each antenna costs about $100, including the LNA). The MWA bow-tie antennas each house two LNAs, one for each polarization. Each LNA has two Agilent ATF-54143 amplifiers. The signal out of the LNAs travels over a 7 meter SMA cable which also carries DC power to the LNA.

2.2 Faraday Cage

2.2.1 Motivation

Since our amplification and filtering chain as well as our deployment site differ from MWA, it was necessary to test the antenna and LNA for use in the Omniscope. However, because our amplifiers and filters were designed to significantly amplify extremely small signals in a radio quiet site, it was difficult to test our analog equipment in the lab. Boston is home to several radio stations less than 2 miles away from our lab in building 37 (see the peaks around 100 MHz on the blue curve in Fig. 2-2(a)). The power measured by the antenna is listed in dBm, which is the power ratio with respect to one milliwatt in decibels. The relation between the power in dBm d and power in watts P is given by Eq. 2.1.

$$d = 10 \log_{10} \frac{P}{1mW} \tag{2.1}$$

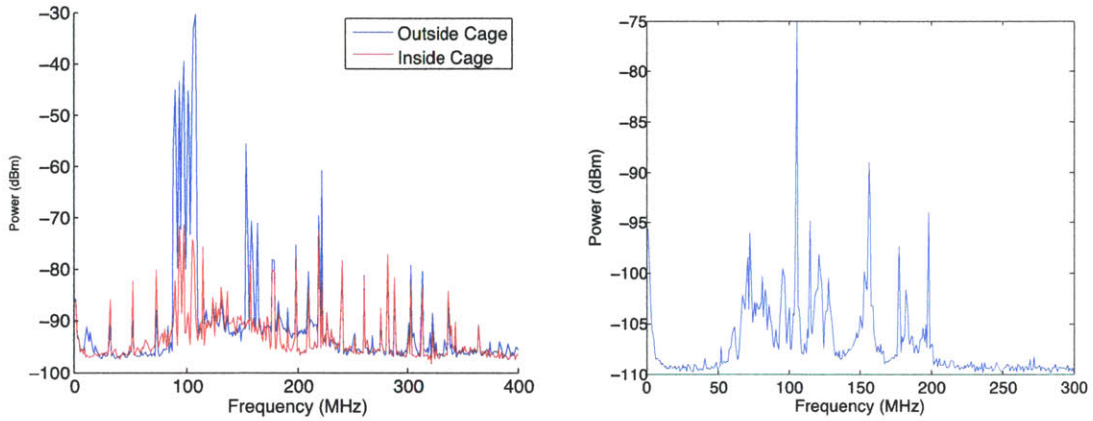

(a) RFI in the lab (blue) and the reduced RFI inside the Faraday cage (red)

(b) RFI in The Forks, ME

Figure 2-2: RFI in (a) the lab and (b) The Forks, ME

The strongest of these radio stations in the lab easily saturate our line drivers (the main source of amplification in the analog chain, described in chapter 3). Our tests of the Omniscope are done in The Forks, ME, a small town (population 47) with no cellphone signal and only one radio station at 101.5 MHz. Fig. 2-2(b) shows that The Forks is considerably quieter in radio frequencies than Boston. The Forks is only around a 5 hour drive from MIT. Fig. 2-3 is a map showing the location of The Forks relative to MIT.

In order to test our analog chain in a radio environment more comparable to the deployment site, I suggested that we build a Faraday cage. Faraday cages are used to shield from electromagnetic radiation. Basic electromagnetism tells us that a shell of a perfect conductor has its cavity completely free from electromagnetic waves. In fact, even a shell with holes (or a conducting mesh) provides good shielding from radiation as long as the holes are much smaller than the wavelength of interest. The range of our instrumentation is between 80 and 300 MHz, which corresponds to wavelengths on the order of a meter.

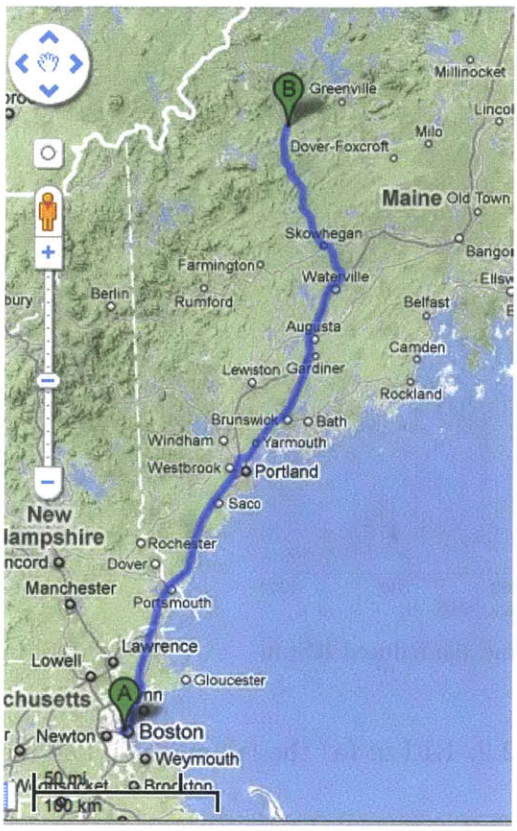

Figure 2-3: Location of (A) MIT and (B) The Forks, ME

2.2.2 Design and Construction

Our objectives in building the Faraday cage were to achieve a considerable noise reduction while keeping the cost reasonable and still allowing the cage to be large enough for the Omniscope signal chain, test equipment, and a UROP student. For materials, copper mesh would provide the best shielding of readily available mesh, but would be prohibitively expensive for a cage as large as we needed. We decided to use reuse the steel ground-screen that we used on a previous field expedition for the Faraday cage. This material has half-inch spacing, which is sufficiently small for our purposes. The frame of the cage was constructed primarily out of 1x2 pine and staples with a base of a reused pallet covered with a plywood sheet. The door has a simple hinged design with a latch that can be closed from the inside. The ground-screen mesh was stapled to the inside of the frame and the edges were sealed with aluminum

tape to ensure conduction. The gaps between the door and the frame were covered in aluminum tape and steel wool. The final Faraday cage (shown in Fig. 2-4) resulted in approximately 40 dB of attenuation in the RF (radio frequency) power from radio stations. Fig. 2-2(a) shows this reduction. Note that the RFI inside the Faraday cage is reduced so that the power in the loudest radio station in the lab is reduced to the level of the loudest radio station in The Forks. Since the radio environment inside the Faraday cage is similar to that of our test site, we were able to test for saturation in the lab.

Figure 2-4: Faraday Cage

2.3 LNA Frequency Response

Using the Faraday cage, we were able to test the response of the LNA in the lab. We had initially designed the rest of our signal chain assuming that the LNA cut off

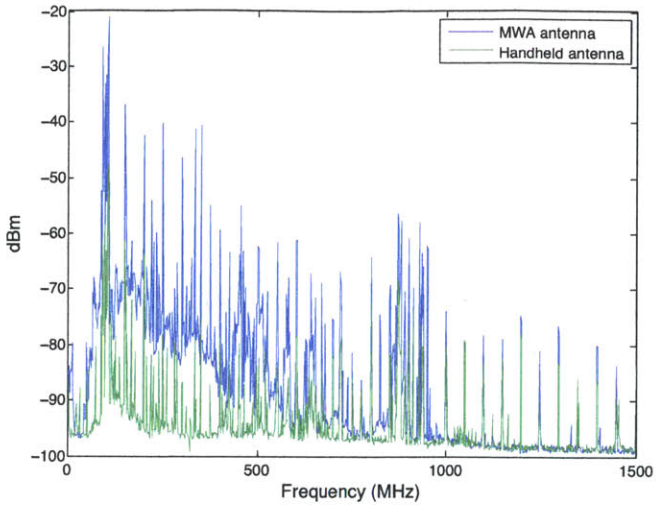

Figure 2-5: RFI picked up by LNA (blue) compared to RFI that is not amplified (green)

around 300 MHz (referring to an MWA plot), but we can see in Fig. 2-5 that this is not the case. There is a large amount of power near 1 GHz and significant power at even higher frequencies. To confirm this behavior, we measured the frequency response of the LNA on our vector network analyzer (VNA).

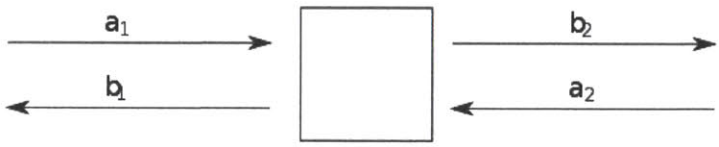

Figure 2-6: Incident and reflected waves in an electrical system

A VNA characterizes the transmission and reflection in an electrical system. At the both the input and output of the system there is both incident and reflected power. These four values are related to each other by a complex scattering matrix given in Eq. 2.2:

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \tag{2.2}$$

Here the a_i and b_i refer to the power traveling in the directions indicated in Fig. 2-6. A VNA measures two of the scattering parameters, S_{11} (the reflection coefficient) and S_{21} (the transmission coefficient). The total gain of the system is obtained from these coefficients as $S_{21}/(1 - S_{11})$. Since the VNA actually reports the log magnitude of the scattering parameters, we obtain the gain from the measured S_{11} and S_{21} as:

$$A = S_{21} - 10\log_{10}(1 - 10^{\frac{S_{11}}{10}}) \tag{2.3}$$

The gain for the LNA is plotted in Fig. 2-7(a).

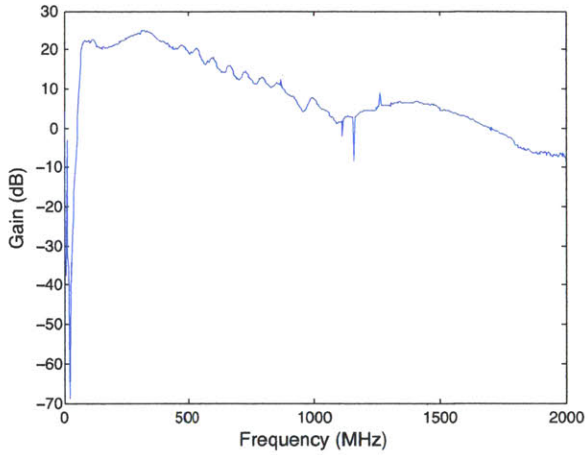

(a) LNA response measured with the VNA

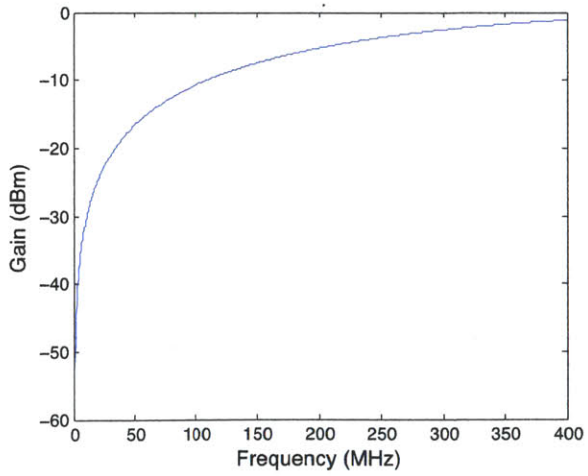

(b) Antenna-LNA impedance mismatch

Figure 2-7: Frequency response due to (a) LNA and (b) LNA-antenna impedance mismatch

This frequency response is of the LNA alone, so it does not include the impedance mismatch between the antenna and the LNA. Although the antenna's bow-tie design was chosen to minimize the variation in impedance, there is still a slight impedance mismatch with the LNA impedance. For a crude approximation of the mismatch, modeling the antenna as a center-fed dipole gives an impedance of $Z_A = 20\pi^2(L/\lambda)^2$, where L is the length of the dipole and λ is the wavelength of the received signal. Taking the LNA impedance to be $z_L = 250\Omega$, the gain G due to the impedance mismatch is then given by Eq. 2.4. Using $L = 0.5$m, this gives the frequency response in Fig. 2-7(b).

$$G = -10\log_{10}\left(\frac{4z_L z_A}{(z_L + z_A)^2}\right) \tag{2.4}$$

2.4 Saturation

The MWA antenna and LNA were designed for use in the radio-quiet Australian site, so we were concerned about possibly saturating the LNA with RFI (radio frequency interference) during our first deployments in Maine, which has more RFI contamination. Using the Faraday cage, the MWA antenna was tested for saturation in the lab. A tone was broadcast by sending a tone from our Agilent N9301A Signal Generator through a 25 dB amplifier and broadcasting it through a dipole antenna inside the Faraday cage. By increasing the tone from the Agilent to high enough power, we were able to see dramatic saturation in the LNA. Fig. 2-8 shows the behavior of the output of the LNA as it approaches saturation. At full saturation (bottom right of Fig. 2-8) you can clearly see the harmonics of the broadcast frequency (137.6 MHz). When saturation begins, the first harmonic can be seen at 265 MHz (top left of Fig. 2-8).

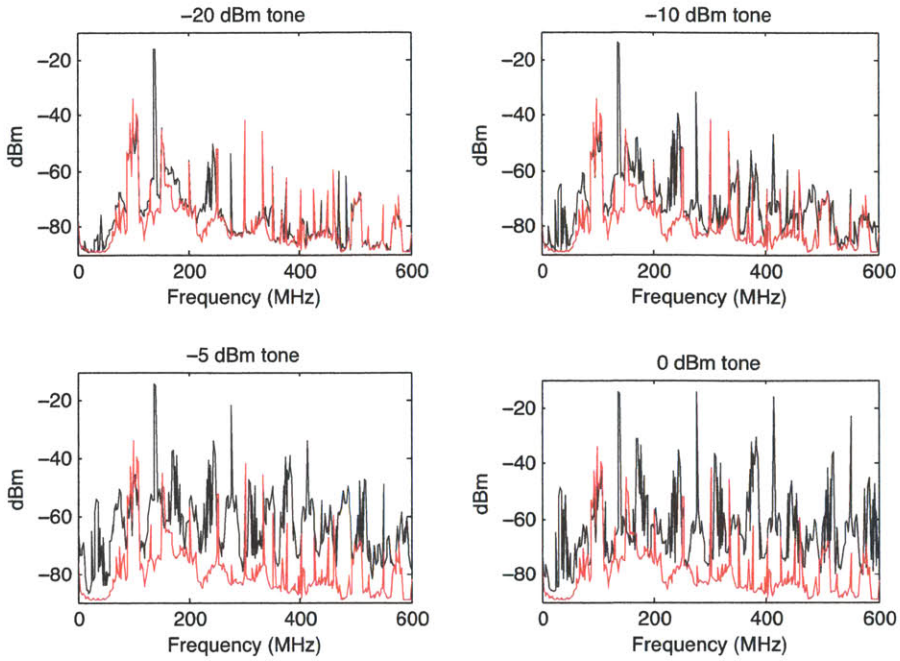

Figure 2-8: Output of LNA for different broadcast tone levels (black) compared to the output with no broadcast tone (red)

2.5 Improvements

From tests in the lab we determined that the LNA was able to amplify in the Faraday cage without saturating. We also found that the LNA did not cut out frequencies above 300 MHz like we had initially expected. This is significant because the rest of the analog chain was designed assuming these frequencies were excluded, so amplifiers down the line were susceptible to saturation at these high frequencies. Because of this discovery, we added a Mini-circuits SLP-200+ low pass filter (DC to 190 MHz) after each LNA, solving the problem. In the future, the low noise amplifier will be redesigned for this project to include a low pass filter.

Chapter 3

Line Driver

3.1 Design

After being amplified by 20 dB in the LNA, the signal is amplified by 51 dB across our band of interest in the line driver pictured in Fig. 3-1(a). Transmission coefficients for relevant frequencies are given in Table 3.1. It is necessary to amplify the signal by a large amount before sending it down the long 75 Ω cables to the receiver boards to ensure that the antenna signal dominates the noise picked up in the cable and analog components further down the chain. The schematic for the line driver is given in Fig. 3-1(b).

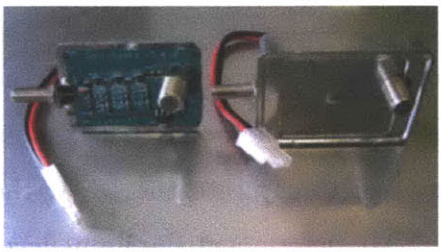

(a) Line driver with and without metal shielding

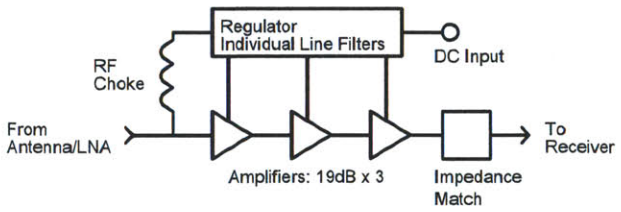

(b) Line driver schematic

Figure 3-1: Line drivers

The line drivers run on 6-12 Volts DC power and in turn provide DC bias power to the LNA. The amplification comes from three broadband 19 dB amplifiers. There are filters between each amplifier for the DC power to avoid feedback. Because of the

Table 3.1: S21 Transmission

Frequency (MHz)	Gain (dB)
100	51.1
150	51.3
200	51.5
500	51.1

large amount of amplification, any output power would get re-amplified, potentially causing a feedback loop and dominating the signal. The maximum total input power for the line drivers is -55 dBm and the maximum power they will output is -4 dBm. The line driver takes in signal from the LNA at 50 Ω but the cables out of the line drivers are 75 Ω. These cables were chosen because they need to be long enough to reach the receiver boards and digital equipment housed up to 50 meters away. The line drivers must be kept close to the antennas because the signal must be greatly amplified out of the LNA before it can travel through the long cable or it will be attenuated away. Because the output of the line driver is sent into a 75Ω cable, they include resistive impedance matching to reduce reflections. Finally, line drivers are shielded in metal cases to prevent excess amplification of exterior RFI.

3.2 Saturation

The line driver is a crucial component of the Omniscope because it is where the cosmological signal is amplified to detectable levels. It is also the point that is most vulnerable to saturation. Even if the RFI is too weak to saturate the LNA, the output of the LNA can still saturate the line driver. We determined the input level that saturates the line driver by sending a tone from the Agilent through 50 dB line driver and 40 dB of attenuator pads. We changed the level of the tone on the Agilent and recorded the output level on the handheld. The line driver behaves linearly until about -50 dBm of input power (Fig. 3-2). When 100 ft long 75 Ω cable was put in after the line driver, the saturation behavior did not change but the output was reduced due to loss in the cable. Fig. 3-2 shows the saturation range of the line

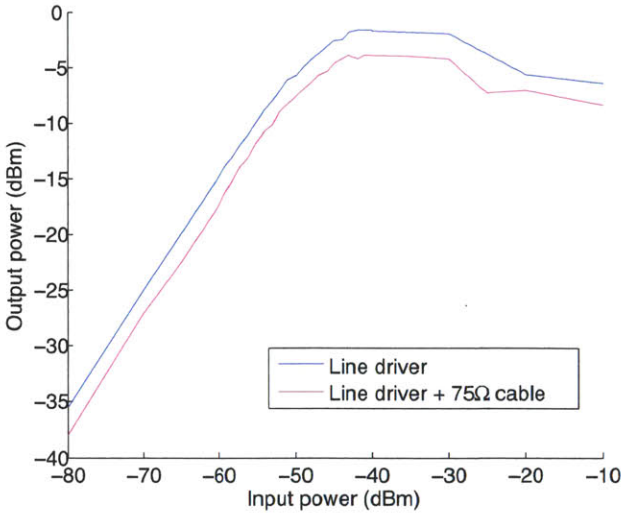

Figure 3-2: Output of line driver versus input amplitude

driver, but we also need to characterize the saturation behavior of the output of the line driver so we can recognize saturation in the field. A good way to visualize frequency spectra is with "waterfall" plots that have frequency on the horizontal axis, time on the vertical, and the power represented with color.

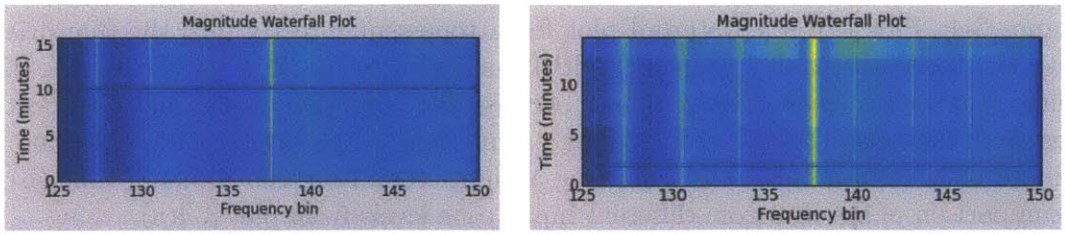

(a) Not saturating

(b) Saturating

Figure 3-3: Waterfall plots for a tone from the Agilent sent directly into 50 dB line driver and receiver board with full attenuation (shift schedule 100100001) with tone level increased in steps of 5 dBm fromand (a) -70 dBm to -55 dBm and (b) -55 dBm to -45 dBm. The horizontal dark blue stripes are a result of packet drops in the computer receiving the data out of the correlator, not an analog effect.

The waterfall plots in this thesis were produced using data analysis software designed by Nevada Sanchez. In order to see how the line driver behaved when it was saturating, we sent a tone from the Agilent at a certain frequency (137.6 MHz) and increased the power by 5 dB every few minutes. The level of the tone increases at

this interval, but we can also see the power level across the band increase when the power of the input tone is increased. This indicates that there is power mixing into different frequencies in the line driver, an indication of saturation. The saturation behavior of the waterfall plot also matches the saturation curve in Fig. 3-2. When the input power is below the level that saturates the line driver (Fig. 3-3(a)), the overall gain does not change when the level of the input tone is increased, but when the line driver is saturated (Fig. 3-3(b)), the overall gain does change.

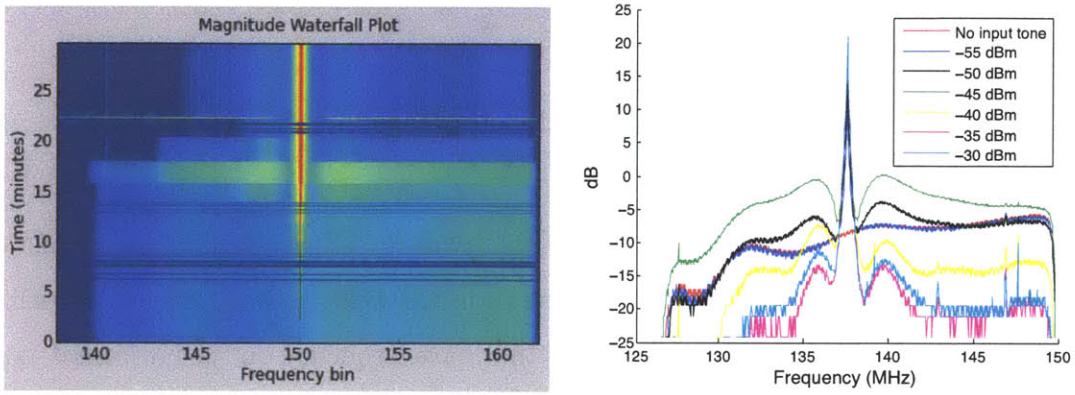

(a) (b)

Figure 3-4: Plot (a) is a waterfall plot for a tone sent directly into the line driver with the tone increased in steps of 5 dB from -80 dBm to -30 dBm (at minute 22). There were many zeroed integrations because the computer capturing the spectra was simultaneously being used for long calculations. The bright horizontal line at minute 22 is from accidentally inputting a higher tone from the Agilent. Plot (b) is a comparison of spectra from different input tone levels from plot 3-4(a).

We can see more directly the effect of increasing the tone on the noise floor by plotting individual spectra for different tone levels. Fig. 3-4(a) was produced in the same way as the waterfall plots in Fig. 3-4 except the input power steps from -80 dBm to -30 dBm. Individual spectra for each of the input power levels are plotted in Fig. 3-4(b). All of the spectra for input tone level below -55 have noise levels that lie on the red curve corresponding to no input tone in Fig. 3-4(b). The saturated curves (those with input power above -50 dBm) have noise levels that vary unpredictably from the true noise level, first increasing and then decreasing.

3.3 Improvements

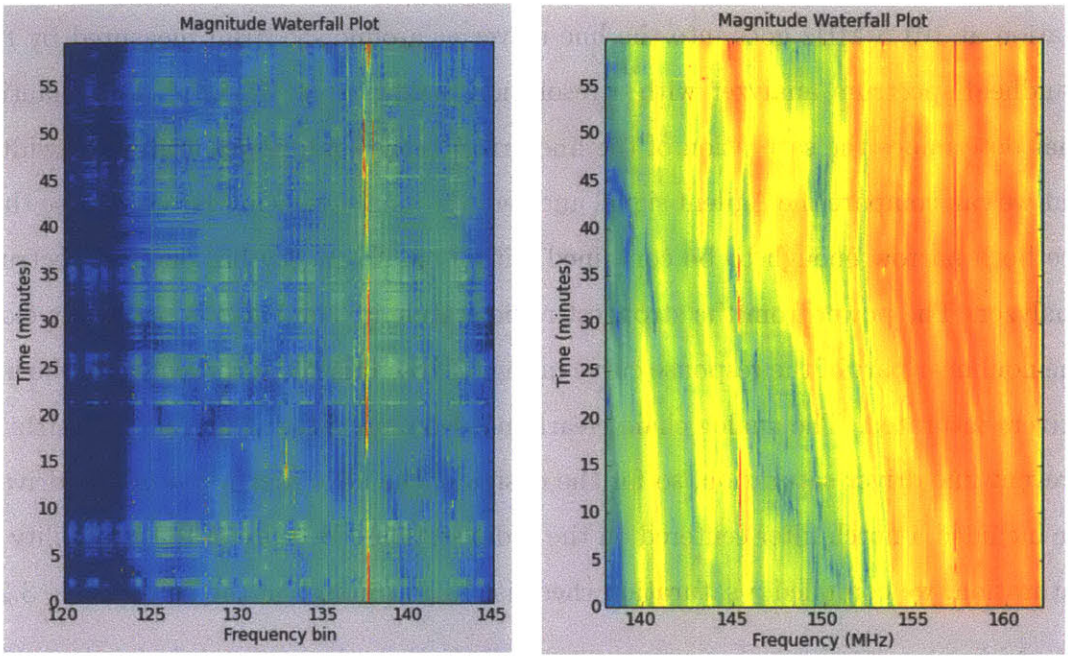

(a) First Maine trip

(b) Second Maine trip

Figure 3-5: Waterfall plots comparing the peformance of (a) first- and (b) second-generation line drivers

The effect of line driver saturation on our data is very nonlinear. If a radio station were to saturate the line driver in the field after being amplified by the LNA, the saturation would not only produce an incorrect value for that frequency, but would also unpredictably shift the base level of signal across the band. Luckily, it is very easy to see in the waterfall plots that saturation has occurred. In fact, examination of waterfall plots from a field test of the Omniscope in The Forks in October 2010 shows that the line driver was saturating there. There are clearly times in the waterfall plots in Fig. 3-5(a) where the level across the frequency band suddenly changes, creating horizontal stripes in both the magnitude and phase. The saturation is likely caused by the radio station at 101.5 MHz. The saturation is apparent in the noise floor because there is an ORBCOMM satellite at 137.6 MHz (the red stripe in Fig. 3-5(a)). This power from this satellite changes as it rises and sets, so the noise floor out of the saturated line drivers changes as well.

We can see that this saturation should have been expected by looking at the RFI in Maine (Fig. 2-2(b)). Adding the 20 dB for the gain of the LNA, the radio station at 101.5 MHz goes into the line driver at around -55 dBm measured by the handheld spectrum analyzer with a resolution bandwidth of 15 kHz. The quantity that determines the saturation of the line driver is the total power in its bandwidth, but we can compare the Agilent signal and radio station signals in Maine because they are both narrow enough to be contained within one frequency bin on the spectrum analyzer. The power from the strongest radio station in The Forks is perilously close the nonlinear part of the response curve in Fig. 3-2, so it is not surprising that the line drivers saturated. The analog chain clearly needed more filtering around 100 MHz to prevent line driver saturation, so for the design of the next generation of line drivers we included a notch filter centered on the radio stations. Because of the possibility of saturation, we compiled a saturation checklist to go through in the field (Table 3.2).

Table 3.2: Saturation Checklist

1. Check output of LNA on spectrum analyzer to make sure it is not saturating (see Fig. 2-8).
2. Check that input to line driver is under -50 dBm per 15 kHz on the handheld spectrum analyzer so that the line driver does not saturate.
3. Check that line driver really is not saturating by broadcasting a tone and seeing if the gain changes across all frequencies when the tone strength is increased (as in Fig. 3-4).

We performed a second test of the Omniscope at the test site in April 2011 using the second generation line drivers. We used the saturation checklist after setting up the signal chain to make sure the line drivers were behaving correctly before collecting data. Indeed, a waterfall plot from that test (show in Fig. 3-5(b)) does not show any signs of saturation. The notch filters in the second generation line drivers and the screw-on low pass filters successfully reduced the radio station power and power above 200 MHz below saturation levels.

Chapter 4

Receiver Board

4.1 Design

The last step for the sky signal before it is converted to digital bits is the receiver board. The receiver shifts the frequency of the signal from the band of interest (80-200 MHz) down to be centered around 0 MHz so that it can be input to the analog-to-digital converter. The receiver also amplifies the signal further. It has a variable attenuator that allows the gain to be set from -8 dBm to +32 dBm. IQ demodulation is used, which splits the signal and mixes with a local oscillator (LO) signal, resulting in the two output signals (I and Q) having a 90° phase shift between them [10]. Using IQ demodulation has the advantage that low speed ADCs can be used.

Each receiver board has four signal paths to reduce cost and reduce the number of times the LO must be split. The boards are powered by 6V DC and the maximum input power is -28 dBm. The receiver first impedance matches the 75 Ohm input from the long cable back to 50 Ohms for input into the ADC, then uses a band-pass filter to select only the band of interest and prevent saturation of the amplifiers. The gain on the variable attenuator can be adjusted so that the output of the two fixed amplifiers is well matched to the range of the ADC. The IQ demodulation module takes in an LO that is double the mixing frequency and reduces it before mixing. After the IQ demodulation, there is a final set of low-pass filters which eliminates the signal above 20 MHz. The final signal has a bandwidth of 40 MHz centered around

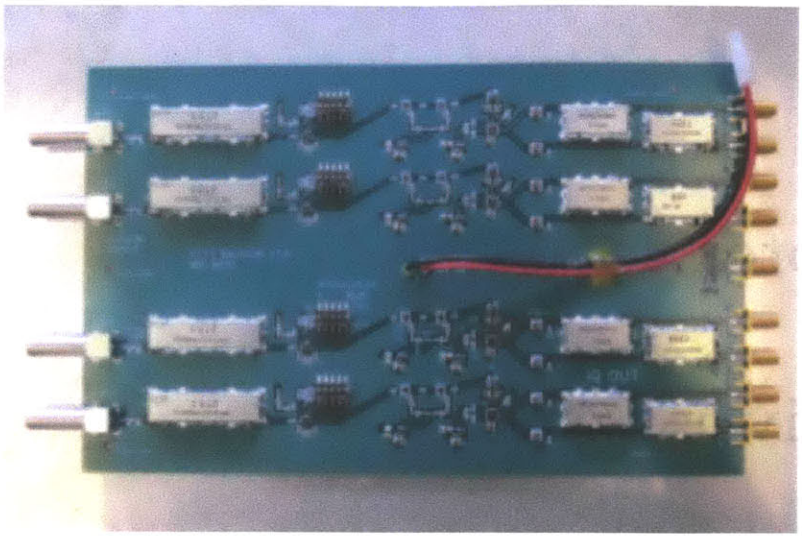

Figure 4-1: Receiver Board

0 MHz with a 2-3 MHz gap in the center. For example, if the LO is se to 150 MHz, we obtain measurements from 130 MHz to 170 MHz.

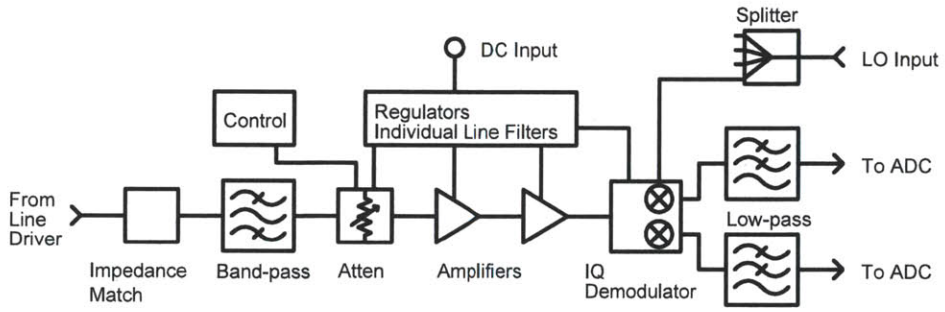

Figure 4-2: Each receiver board takes in four polarizations. The impedance is matched back to 50 Ω for input into the ADC. The signal is bandpassed and passes through variable attenuators, amplifier, I/Q demodulation, and low-pass filter. There are two outputs (I and Q) for each polarization

Chapter 5

Analog to Digital Conversion

5.1 Design

An analog-to-digital converter (ADC) converts a continuous analog signal to a discrete digital representation. The resolution of the ADC is expressed in bits. For example, a 12-bit ADC represents the range of its allowed input signal in $2^{12} = 4096$ integer values. The analog to digital conversion for the Omniscope is done on commercial ADCs. Our ADC consists of eight 12 bit, 8-channel ADCs from Analog Devices (model ADS5272) that run at 50 megasamples per second, and handle 9 input channels each. Eight such ADCs are mounted on a single board sold by Rick Raffanti that is designed to be compatible with our digital hardware, thus digitizing 64 channels at once.

5.2 ADC Quantization Noise

The conversion of continuous analog signals to discrete digital signals introduces error due to the finite number of bits in the ADC [11]. The input analog signal is continuous but the ADC can only represent the input value by a discrete value with a resolution depending on the number of bits in the ADC. The difference between the actual signal value and the discretized value results in quanitzation noise. Since this error is just due to rounding, it should be random and uncorrelated between channels.

Thus we expect it to decrease when we average the cross spectrum between any two channels over time. This is one of the reasons an accumulator is included in the digital correlator (averaging correlations also reduces thermal noise from the analog chain). If the correlator averages correlations for a long enough time, in theory the noise introduced by the ADC would decrease below the level of the systematic noise from the rest of the signal chain.

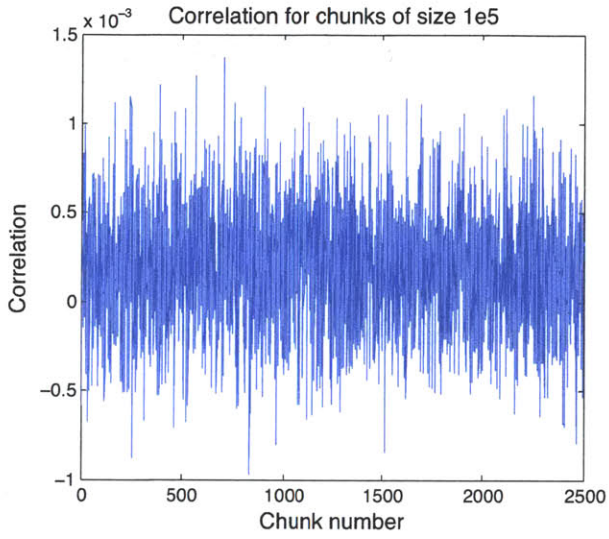

Figure 5-1: Example of the variation in the correlation in each section of size 10^5

In reality, there is some noise produced in the ADC that is correlated between channels due to electromagnetic coupling between the cables for the different channels. If there is such cross-talk we expect that at some point the averaged cross spectrum will hit a noise floor. We would also expect the cross talk to be greater between some channel pairs than others, for example those adjacent on one of the eight ADC chips and those sent on neighboring wires to the ADC. To examine the noise floor due to cross-talk between channels in the ADC, we took long sections of time-stream data (about 5 seconds, or 250 million samples at 50 million samples per second) over 10 Gigabit Ethernet from two terminated channels inside the Faraday cage.

If the noise is random, then the correlation in sections of size 10 samples will vary more than the average correlation in sections of size 10^5 samples because in the larger

sections, the noise has cancelled more. Fig. 5-1 shows some example correlations for a section size of 10^5. To quantify how much the correlation has stabilized for each length section, we compute the following statistic. We split the correlation into sections of length n samples. For each section of size n, we computed the correlation, giving $m = N/n$ such correlations if $N = 2.5 \times 10^8$ is the total number of samples. Then we computed the standard deviation of these m values. Repeating this procedure for different values of n gave the behavior of the noise with increasing accumulation length.

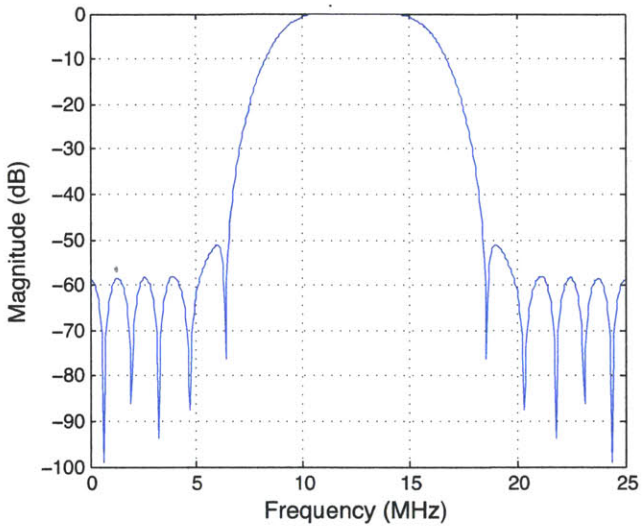

Figure 5-2: FIR bandpass filter

Taking the correlation between two channels involves taking a length n FFT of each channel and complex multiplying them. Taking an FFT becomes computationally intensive as n becomes large, so we instead took correlations in the time domain. This gives us the same answer as calculating the correlation in the frequency domain because the Fourier transform is a unitary operation. Because the cross power may be frequency-dependent, we want to take the mean cross power over a narrow bandwidth. In order to accomplish both goals, we can first bandpass the time signal and then calculate the mean cross correlation of the sections in the time domain. One way to bandpass the signal and remain in the time domain would be to divide it into small sections and for each section Fourier transform, then multiply by the bandpass,

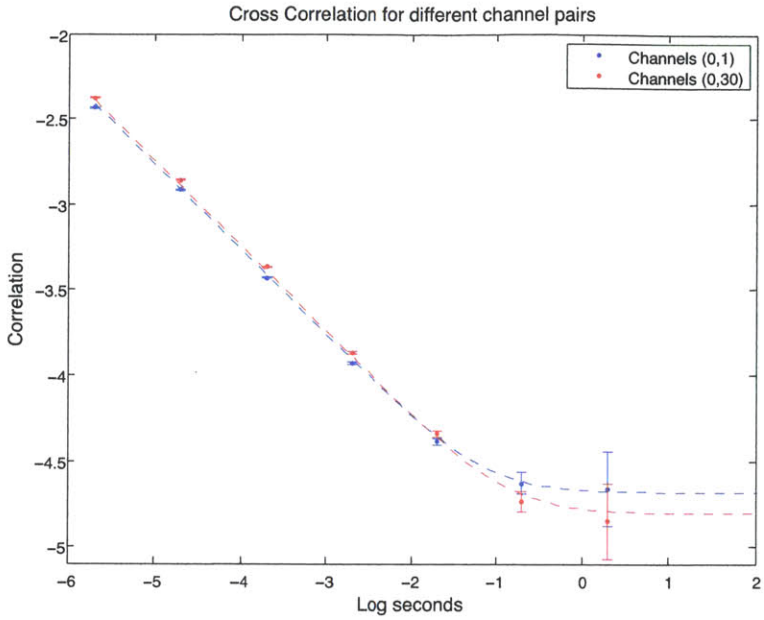

Figure 5-3: Correlation between two distant channels (red) and two neighboring channels (blue)

and then inverse Fourier transform back to the time domain. Computationally in our case it is more efficient to apply an FIR (finite impluse response) filter to the whole time-stream at once. FIR filters are implemented by multiplying the input signal by a set of coefficients that implement the impulse response corresponding to the windowing. The frequency response of the FIR filter used is given in Fig. 5-2.

Once the timestream was bandpass filtered, the correlation function for each section was calculated by multiplying the time-streams together and the standard deviation σ of the correlations in each section was found. Plotting the correlation over time for different pairs of channels (Fig. 5-3), we see that the channels that are furthest apart (0 and 30) have a lower cross correlation noise floor that neighboring channels (0 and 1). The error bars on this plot are the errors on the standard deviations, which are given by $\Delta\sigma = \frac{0.71\sigma}{\sqrt{m}}$ if we assume the signals are random. Since the mean of noise correlation data is not exactly zero due to cross talk in the ADC and other effects, the decrease in correlation with section size scales not like $1/\sqrt{N}$, but like $\sqrt{\mu^2 + \sigma^2/N}$. The fit curves in Fig. 5-3 are given by this function. The time at

which the correlations level off in Fig. 5-3 (around 2 seconds) tells us how long to average the correlations in the digital correlator. Averaging the correlations for any longer would not result in any further noise reduction.

5.3 Improvements

These noise accumulations show a significant coupling effect between channels that are on the same ADC chips. We are currently developing a device called a swapper that will reduce this cross-talk. The swapper works by multiplying the signal at the very beginning of the signal chain by a known series of -1 and 1. Once the signal travels through the analog chain and into the ROACH it is multiplied by the same series of zeros and ones in software, before the signal is averaged down. Any part of the signal that was picked up after the first swap has been multiplied once by either -1 or 1, so the positive and negative parts will sum to zero in the accumulations and this signal will cancel out. However, any signal that was present before the first swap (including the cosmological signal) has been multiplied by either -1 or 1 twice, so it will not cancel in accumulations. As long as the swapping is applied on a timescale smaller than the time at which the correlation reaches a floor in Fig. 5-3, the swapper will effectively eliminate the cross-talk in the ADC.

Chapter 6

Digital Correlator

6.1 Hardware

Once the analog signal is converted to digital format, the computations for the Omniscope are performed on Field Programmable Gate Arrays (FPGAs). The advantage of FPGAs over traditional computers is that FPGAs can be programmed at the gate level, which allows for highly efficient parallel computation. The FPGAs used in the Omniscope were designed by the Center for Astronomy Signal Processing and Electronics Reseach (CASPER). This group was formed to develop open-source hardware for radio astronomy [9]. The Omniscope's computations are performed on their ROACH boards (Reconfigurable Open Architecture Computing Hardware). We program the ROACH boards using Linux-based design tools developed by CASPER.

6.2 F-engine Design

The F-engine is the part of the digital design that performs the temporal FFT. Fig. 6-1 shows a block diagram of the F-engine. Before the digital signal from the ADC is Fourier transformed, it must be filtered (step two in Fig. 6-1). When a finite length discrete signal is Fourier transformed, frequencies in the signal that do not correspond to integral multiples of the sampling frequency divided by the FFT length are represented in many output frequencies in the FFT. This effect is termed spectral

leakage, and can be minimized by first multiplying the input signal by a windowing function that tapers the signal to zero at the edges [10]. The F-engine actually uses a more complicated windowing scheme called a polyphase filter bank [12].

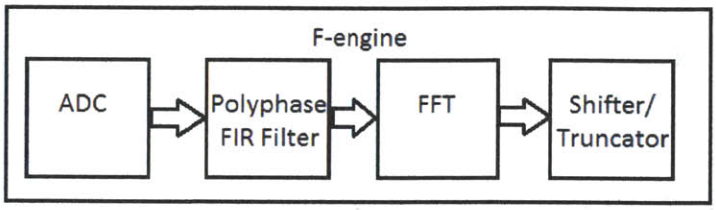

Figure 6-1: Overview of F-engine provided by Omniscoper and candidate for Masters in electrical engineering, Jon Losh

During the windowing, the signal is multiplied so it needs to be represented with more than 12 bits. The signal is represented in 18 bits at the output of the polyphase filter bank as a compromise between accuracy and conserving hardware. After the filtering, an FFT is performed on each signal. This FFT is actually done in multiple stages. In each stage, there is a possibility that the result may be larger than can be represented in 18 bits, so it must be shifted down. Thus the FFT includes a shift schedule that specifies for each stage whether to shift or not. The shift schedule depends on the type of input signal. For a tone of a single frequency it would be advantageous to shift at every step, so the shift schedule would be all ones. For a sharp pulse in time, which has equal power in all frequencies, which has equal power at all frequencies, the ideal shift schedule would be all zeros. For white noise, the best shift schedule is alternating zeros and ones.

After the FFT, the result must be truncated to 4 bits in the shifter/truncator before being correlated so there will be enough room on the ROACH for the computations. However, there is some subtlety here because loud signals have a lot of information in the higher bits and quiet signals have all their information in lower bits. In order to preserve the quiet bits, it would be good to first shift them to higher bits and then truncate. However, using this same shifting scheme for frequency bins with high power would cut off those values, resulting in digital saturation. The so-

lution is a frequency-dependent shifting scheme. In this scheme, the number of bits to shift for each frequency bin is determined by the average amount of power in that bin. This frequency-dependent shifting in the shifter/truncator is separate from the frequency-independent FFT shift schedule.

6.3 X-engine Design

In the X-engine, the FFTs from the F-engine are correlated. In the current design that uses 32 polarizations, the correlation is done by complex multiplying all antenna pairs, but in the future when the number of antennas is large this will be done by FFT multiplication to reduce computational cost. We currently have 8 ROACHes but for 32 polarizations we only need one to implement the digital design. After being correlated, the complex values are averaged together for a certain length of time to reduce noise. The output of the correlator is viewed in software as the magnitude in dB and the phase.

6.3.1 Testing

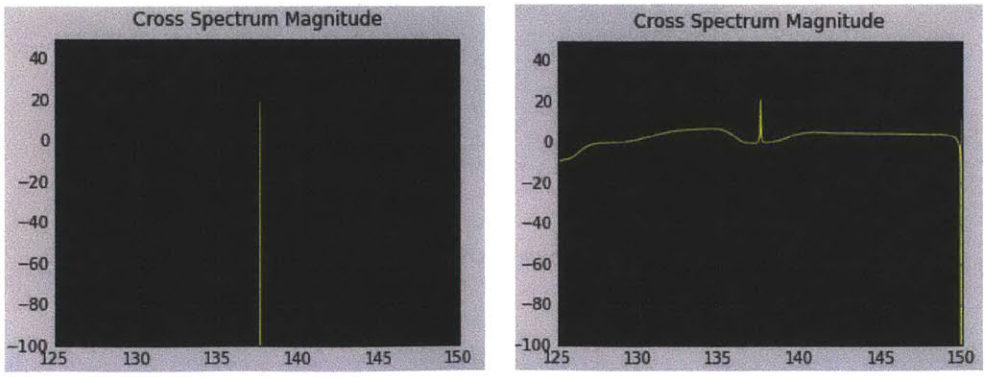

(a) (b)

Figure 6-2: Tone and noise sent directly into line driver and receiver board with shift schedule (a) 111111111 and (b) 1010101010

We can see the importance of having the correct FFT shift schedule by looking at the spectra of the same input with different shift schedules. The test signal we used

was broadband noise output by an amplifier added to a tone from the Agilent at 137.6 MHz. It is useful to test the optimal shift schedule for this the type of signal because the signal we see at The Forks will have a combination of strong tones from point sources (like the ORCOMM satellite at 137.6) and broadband noise and astrophysical sources like the Sun and our galaxy. The test signal was sent through the line driver and the receiver board and then processed in the F-engine and correlator on the ROACH. With a full shift schedule, the noise floor is shifted away so only the peak can be seen Fig. 6-2(a). However, shifting only every other step raises the noise floor artificially so there seems to be power below 130 MHz even though the band-pass filter in the line driver actually cuts out the signal below that frequency (Fig. 6-2(b)).

6.4 Sawtooth in Cross Correlations

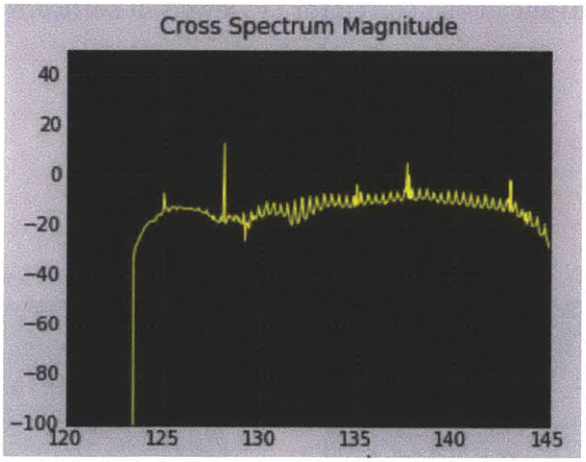

Figure 6-3: Sawtooth in The Forks

When we tested the digital correlator for the first time in The Forks, we found a surprising issue with the correlations. We expect the spectra to be smooth because noise contributions from the signal chain and the sky would be broadband. What we see, however, is a sawtooth pattern in the noise floor of cross correlations between different antennas. Fig. 6-3 shows a cross correlation from The Forks exhibiting the sawtooth pattern. This pattern appears in cross correlations of every pair of antennas,

but in no auto-correlations. The source of this signal must be correlated among every pair of antennas. The sawtooth has a period of about 0.5 MHz, which corresponds to a timescale of 2 μs. The light travel time in a 50 meter cable has a round trip travel time of $\frac{100m}{c} \approx 0.3\mu s$ where c is the speed of light, so the sawtooth is not likely due to reflections in the cables.

In order to reproduce this effect in the lab, we measured cross correlations using the LNA and line drivers, using a DC filter that allowed us to attenuate the signal out of the LNA without having leakage into the line driver. In this setup we sent 6V DC power from the lab bench power supply through the DC filter into a bias-tee, which was connected on the RF+DC end to the antenna. The RF end of the bias-tee was connected to two 20 dB screw-on attenuators. After the attenuators there was a bandpass filter, a DC block, and finally a line driver. The DC block was used because it contains a capacitor that prevents feedback into the line drive. The line driver then was connected to a 75 Ω cable leading to the receiver board as usual.

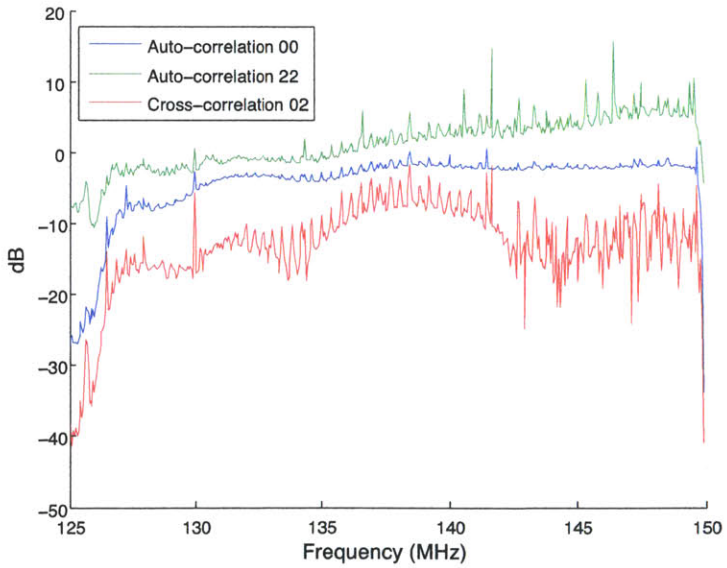

Figure 6-4: Auto-correlations and cross-correlations from sawtooth correlations in the lab

Using this setup, we were able to reproduce the sawtooth pattern in some cross correlations (Fig. 6-5(a)). Fig. 6-4 shows the auto-correlations and cross-correlation

from the the same data that produced Fig. 6-5(a). As we can see in Fig. 6-4, the magnitude of the auto correlations is generally higher than the magnitude of the cross correlations, because noise that is uncorrelated between two antennas is averaged down. The sawtooth noise must be of a lower magnitude than other noise sources, so it is dominated by them in auto correlation but it can be seen when the other noise is averaged down in cross correlation. Physically moving the antenna changed whether the sawtooth appeared in correlations. Moving it away from the window 6-5(b) or shielding it by someone holding it in their arms it made the sawtooth disappear. Moving or shielding the antenna both reduce the power received by the antenna, so the sawtooth is possibly a saturation effect. However, putting attenuation before or after the line drivers did not reduce the sawtooth effect. Investigation into this effect is ongoing.

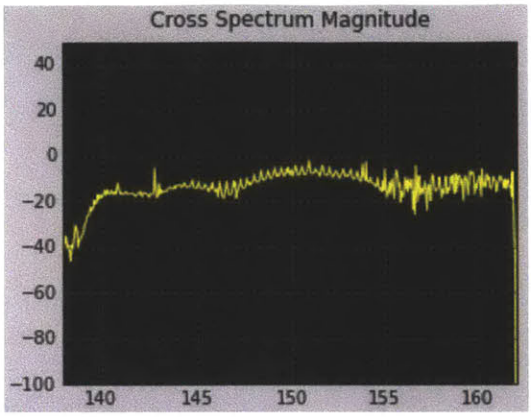

(a) Antenna near window

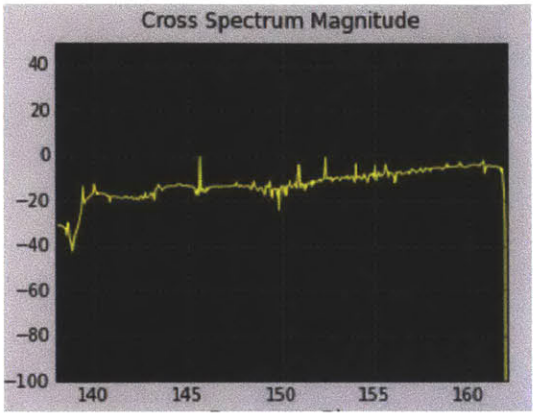

(b) Away from window

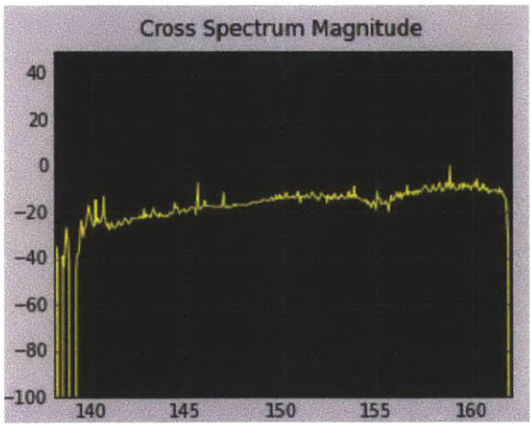

(c) Holding the antenna

Figure 6-5: (a) Cross-correlation from antenna in the lab, (b) away from the window, (c) holding the antenna

Chapter 7

Calibration

7.1 Relative calibration of analog chain

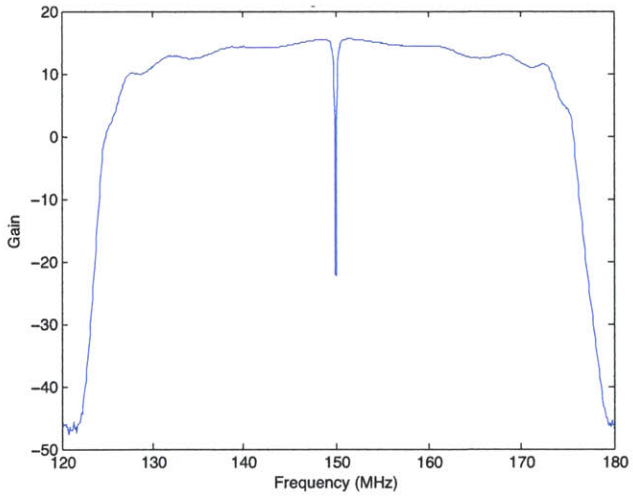

Figure 7-1: Frequency response of line driver and receiver boards with arbitrary offset for an LO of 150 MHz. The gain is symmetric about the LO frequency.

All of the spectra up to this point have been uncalibrated. That is, when we measure the power coming out of the Omniscope, we are not measuring the actual sky power, but this value multiplied by the gain of all the filters and amplifiers in the signal chain. The first step in calibrating our results is to divide out the known frequency-dependent gain of the LNA, line drivers, and receivers. Fig. 7-1 was

produced by measuring the output with a signal of known power into the line drivers and receiver boards and dividing by the input power. This measurement, performed by Andy Lutomirski, gives the shape of the frequency response due to the the line drivers and receivers. Combining this gain with the gain of the LNA from Fig. 2-7(a) and the antenna impedance mismatch from Fig. 2-7(b), we can correct for all the known sources of frequency-dependent gain from the analog chain to find a relative calibration of the sky signal.

7.2 Absolute calibration to solar flux

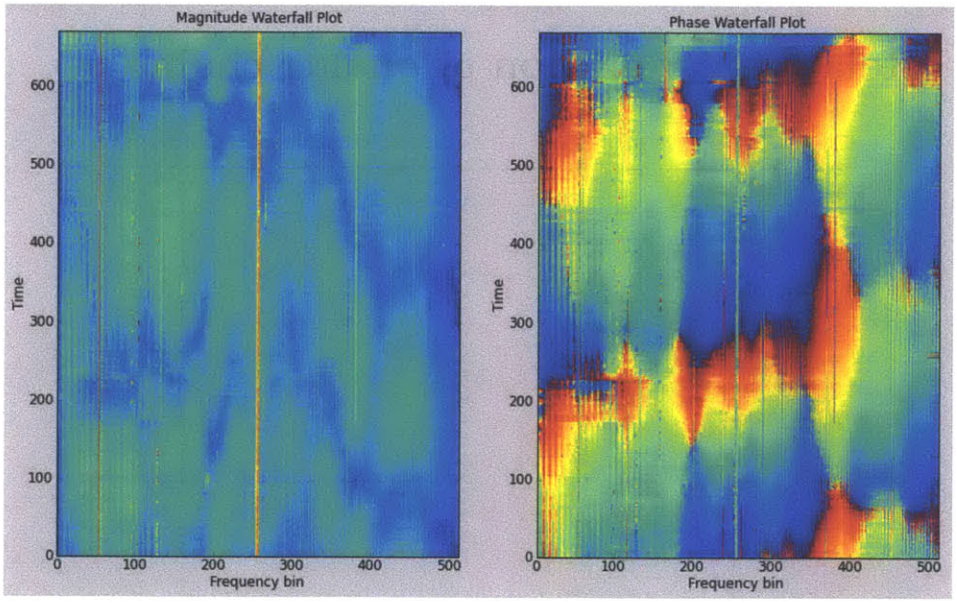

Figure 7-2: Waterfall plot from The Forks showing solar fringes

In order to calibrate the The Forks data in absolute units, we first divide out by the frequency response of the antenna and the analog chain described in the previous section. However, there still remains an unknown gain in the system from the antenna beam, and we would like to cross-check our calibration to an absolute reference such as the Sun. To compare the magnitude of the cross correlations from The Forks to the solar flux, it is necessary to characterize the quality of the data so that we can find spectra that are dominated by the Sun. Since the Sun continually traverses the

sky during the day, the angle between the Sun and the line connecting two pairs of antennas is constantly changing. This means that a waterfall plot that is dominated by solar power will have fringes in the phase. Fig. 7-2 is an example of such fringes seen at The Forks.

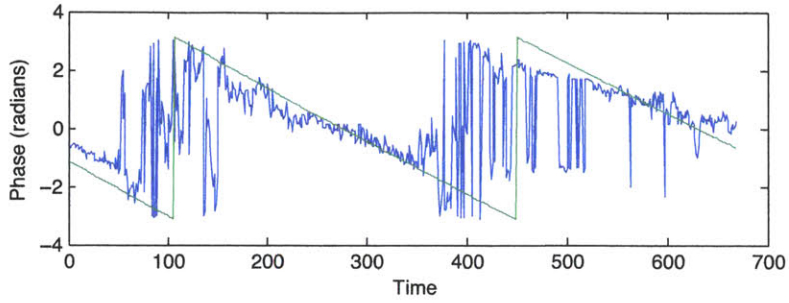

Figure 7-3: Cross spectrum over time compared to linear fit

Using the waterfall plots as a guide, spectra suitable for calibration can be determined by looking at individual plots versus time (Fig. 7-3). As the Sun moves in the sky over time, the phase of the cross spectrum is expected to vary as $a\cos(\vec{d}\cdot\vec{k})$, where $\vec{d}$ is the vector between the two antennas and $\vec{k}$ is the wave vector of radiation from the Sun. This is apparent in Fig. 7-3 (the phase actually zigzags because 2π wraps around to 0).

Once we find spectra that are dominated by solar power, we can calibrate them by comparing their magnitude to measured solar flux. Solar flux varies considerably from day to day depending on weather conditions and solar events such as sunspots, but the National Oceanic and Atmospheric Administration Space Weather Prediction Center makes daily measurements of this value at several frequencies including 245 MHz and 410 MHz [13]. The nearest NOAA observatory for which these measurements are available is Sagamore Hill Solar Radio Observatory in Hamilton, MA. We used a power law extrapolation of the noon solar flux at these frequencies as a model of solar flux. The solar power is given the unit of Janskys, which is equal to 10^{-26} Watts/m^2/Hz. After subtracting the LNA, line driver, and receiver gain the noon cross power spectra taken at The Forks were fit to this model.

To calibrate the spectra, they are saved from the Omniviewer and converted to

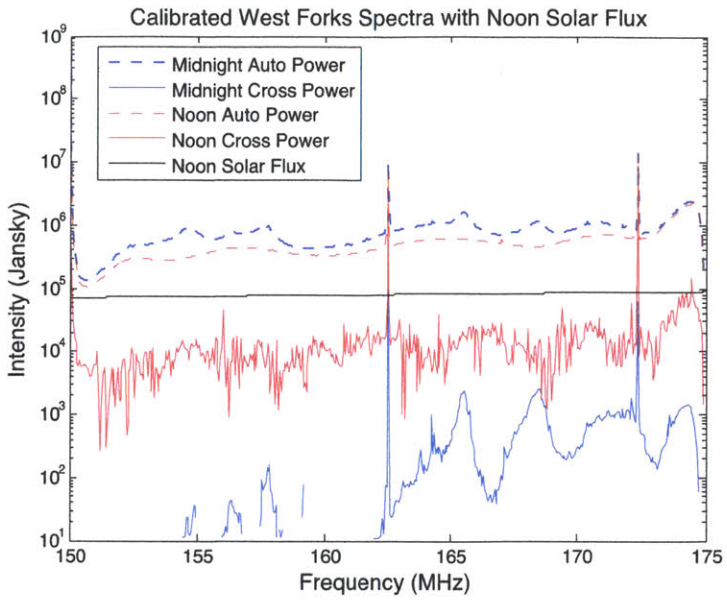

Figure 7-4: Calibrated median The Forks spectra

dB using Eq. 2.1. Then the frequency response of the line drivers (Fig. 7-1) and the frequency response of the LNA (Fig. 2-7(a) are subtracted, as well as the gain due to impedance mismatch (Fig. 2-7(b)). Finally, the spectra are inverted back to power and multiplied by 10^{26} to obtain units of Janskys that can be compared to the solar fluxes. The spectra are calibrated by fitting the level of the noon cross power to the extrapolation of the NOAA solar fluxes. This gives a constant offset that can be applied to all of the spectra. The resulting calibrated spectra are given in Fig 7-4.

Chapter 8

Conclusions

In the short time that I have been a part of this project, I have seen the Omniscope go from a collection of analog parts and digital designs to an integrated system producing RF spectra in the field. My role has been to integrate the analog and digital development so that problems are quickly identified and improvements can be made. Since the first system test in The Forks, when the system was already capable of measuring signals such as radio stations, satellites, and the Sun moving across the sky, much progress has been made. From the plots from this trip I identified several issues. Through my investigations in the lab, I found that the LNA responded at frequencies above 300 MHz and that an insufficient shift schedule in the F-engine resulted in digital saturation for some signals. I also found that the sudden gain increases across the band were due to line drivers saturating on radio station and high frequency power.

Using these insights, we have already made several improvements to the Omniscope signal chain. We included low pass filters at the output of the antennas to prevent power over 200 MHz from saturating the line drivers and included a frequency-dependent shifter in the F-engine to prevent digital saturation. We also included a notch filter in a new version of the line drivers to prevent them from saturating on power from a radio station at 105.1 MHz. We have already seen the effect of these improvements on our signal in a second test at The Forks and there is still progress to be made. Future low noise amplifiers could include low pass filtering in the an-

tenna itself, and a swapper is being built to reduce system noise. Finally, we must make further investigations into the source of the sawtooth pattern seen in some cross correlations.

The Omniscope has now made much progress toward the goal of becoming a high-resolution measurement of the cosmic dark ages and reionization. For the next step we plan to build a 128-polarization instrument that will include the FFT correlator. The Omniscope is an exciting part of the current 21 cm experiments to probe new domains of cosmology.

Bibliography

[1] D. Larson *et al*, *Seven-year Wilkinson microwave anisotropy probe (WMAP) observations: Power spectra and WMAP-derived parameters.* Astrophys. J. Suppl. **192**, 16 (2011) [arXiv:1001.4635 [astro-ph.CO]].

[2] Bing Zhang, *Astrophysics: Most distant cosmic blast seen.* Nature **461**, 1221 (2009) [http://dx.doi.org/10.1038/4611221a].

[3] Y. Mao, M. Tegmark, M. McQuinn, M. Zaldarriaga and O. Zahn, *How accurately can 21 cm tomography constrain cosmology?* Phys. Rev. D **78**, 023529 (2008) [arXiv:0802.1710 [astro-ph]].

[4] M. Tegmark and M. Zaldarriaga, *The Fast Fourier Transform Telescope.* Phys. Rev. D **79**, 083530 (2009) [arXiv:0805.4414 [astro-ph]].

[5] M. Tegmark and M. Zaldarriaga, *Omniscopes: Large area telescope arrays with only N log N computational cost.* Phys. Rev. D **82**, 103501 (2009) [arXiv:0910.1351 [astro-ph.CO]].

[6] S. Furlanetto *et al*, *Cosmology from the Highly-Redshifted 21 cm Line.* (2009) [arXiv:0902.3259 [astro-ph.CO]].

[7] A. Liu, M. Tegmark, S. Morrison, A. Lutomirski and M. Zaldarriaga, *Precision Calibration of Radio Interferometers Using Redundant Baselines.* Mon. Not. Roy. Astron. Soc. **408**, 1029 (2010) [arXiv:1001.5268 [astro-ph.IM]].

[8] J. D. Bowman *et al*, *Field Deployment of Prototype Antenna Tiles for the Mileura Widefield Array-Low Frequency Demonstrator.* Astron. J. **133**, 1505 (2007) [arXiv:astro-ph/0611751].

[9] A. Parsons *et al*, *Digital Instrumentation for the Radio Astronomy Community.* [arXiv:0904.1181 [astro-ph.IM]].

[10] R. G. Lyons, *Understanding Digital Signal Processing.* Prentice Hall Professional Technical Reference. Upper Saddle River, New Jersey (2004).

[11] S. W. Smith, *Digital signal processing: a practical guide for engineers and scientists.* Newnes, Burlington, MA (2003).

[12] Weiss, S.; Stewart, R.W., *Fast implementation of oversampled modulated filter banks.* Electronics Letters **37** 1502 (2000) [doi: 10.1049/el:20001068].

[13] U.S. Dept. of Commerce, NOAA, Space Weather Prediction Center, [http://www.swpc.noaa.gov/].

Appendix A

Assembling the Omniscope

Before taking data, use Table. A.1 to check that the Omniscope is set up correctly.

Table A.1: Omniscope Assembly Checklist

1. For each antenna, screw each of the two SMA cables into input of a line driver.
2. Screw one end of a 75Ω cable to the output of each line driver and the other end to one input of a receiver board.
3. Connect the power cable for each line driver to the power source plugged into an outlet.
4. Make sure that all of the 32 SMA cables from the 64-input ADC are screwed into the receiver board.
5. Plug in the receiver board power supply and turn it on by flipping the switch.
6. Make sure that a cable from a signal generator that is providing the LO frequency is connected to the splitter and that each of the four outputs of the splitter goes into a receiver board.
7. Make sure the signal generator is outputting signal.
8. Plug in and turn on the ROACH containing the 64-input ADC (it will hum and the green light will stay steady instead of blinking).
9. Make sure that an ethernet cable is connected to the back of the ROACH (and the 10 gigabit ethernet cable, if using).

Appendix B

Taking Data

B.1 Taking data from the ADC

To take timestream data from the ADC on Idun (the lab computer), you need to connect the 10 gigabit ethernet cable to the back of the ROACH with the 64-input ADC inside and to the back of Idun. Make sure that the ROACH is also connected to the regular ethernet. Note the IP address of the ROACH (it should be labelled on the case). You should have an account on Idun and a copy of the SVN. On Idun, type:

```
~/fftt/models/adc_to_net_x64.py <roachIP>
```

This instructs the ROACH to start sending the data over 10 gigabit ethernet. `<roachIP>` is the IP address of the ROACH. To capture the data on Idun, type:

```
python2.6 ~/fftt/utils/adc64_to_raw.py <roachIP> <#samples> <filename>
 <channel0> <channel1> <channel2> <channel3>
```

Where `<#samples>` is the number of samples (this must be less than 250000000 or Idun will crash), `<filename>` is the name of the file to be created, and `<channel0>` through `<channel3>` are the four ADC channels (from 0 to 31) that you want to take data from. This creates a binary file that you can read in MATLAB.

If you only want to take a small number of samples you can use the capture script that writes out ASCII format. To use this script, type:

```
python2.6 ~/fftt/utils/adc64_to_ascii.py <roachIP> <#samples>
<channel0> <channel1> <channel2> <channel3> > <filename>
```

To read the binary output of the `adc64_to_raw.py script`, the MATLAB commands are:

```
fid = fopen('filepath')
f=idivide(fread(fid,[4,Inf],'int16=>int16',0,'b'),int16(16),'floor');
fclose(fid)
```

This makes a MATLAB array f, with 4 columns, one for each channel, that have length equal to the number of samples.

B.2 Taking spectra from the digital correlator

The following commands save correlations from the roach in miriad format, which can be converted to .odf format and viewed on the Omniviewer. First type:

```
cd fftt/models/fx_s_tests/sw
```

Next, open the file `mit_pock8_1k.conf`. In this file, set `fft_shift` to the desired FFT shift schedule and `t_per_file` to the length in seconds you want the file to be. To take data, type:

```
python2.6 /fxs_start.py mit_pock8_1k.conf
```

This should result in a positive message about transmitting data. To collect data, type:

```
python2.6 /cn_rx_new.py -s mit_pock_1k.conf
```

The data is saved in `~/fftt/models/fx_s_tests/sw/data/` as a miriad file named with a timestamp.