A Ku-Band Polarimeter for the Owens Valley Radio Observatory 40-Meter Telescope

Thesis by

Kirit Sukrit Karkare

In Partial Fulfillment of the Requirements

for the Degree of

Bachelor of Science

California Institute of Technology

Pasadena, California

2011

(Submitted May 3, 2011)

© 2011
Kirit Sukrit Karkare
All Rights Reserved

Acknowledgments

I am grateful to everyone in the Caltech Radio Astronomy Laboratory for their help during this project. Oliver King, with whom I've worked very closely for the last year, has been incredibly willing to explain the gory details of radio labwork, and I've learned a huge amount from his guidance. Glenn Jones is the reason I was able to get anything CASPER-related working in the first place, and has been extremely patient with all of my questions. Matthew Stevenson introduced me to CASPER and set me up to start working two summers ago. Rodrigo Reeves and Kieran Cleary have been great labmates. Walter Max-Moerbeck and Joey Richards were very helpful in explaining the science behind the instrument.

The staff at the Owens Valley Radio Observatory were excellent hosts during my several visits.

Finally, thanks to my research advisor, Tony Readhead, for giving me the opportunity to spend the last two years with the Caltech Radio Astronomy Laboratory. He has been an amazing teacher and mentor and I could not have asked for a better person to introduce me to radio instrumentation.

This work was supported by the Caltech Summer Undergraduate Research Fellowship Office and by the Rose Hills Foundation. Travel to the 2010 CASPER Workshop was provided by the Harvard-Smithsonian Center for Astrophysics.

Abstract

Blazars are active galactic nuclei – small, extremely luminous objects at the center of galaxies powered by material accreting around a supermassive black hole – which emit relativistic jets of highly energetic plasma along our line of sight. There is no accepted model for jet composition, acceleration, and confinement; observations at different wavelengths will help us understand these emission mechanisms. Since 2007, the 40-Meter Telescope at the Owens Valley Radio Observatory has been monitoring over 1100 blazars every two days. The variability in radio light curves is likely to be correlated with gamma-rays, which we are observing with the Fermi Gamma-Ray Space Telescope.

A new Ku-band receiver for the 40-Meter Telescope is in development. It will offer increased sensitivity, MHz spectral resolution from 12-18 GHz, and calculation of Stokes I, Q, and U parameters. In this thesis, I present the design, assembly, and testing of various components in the receiver chain. I evaluate the suitability of a commercial radio frequency over optical link for use in the receiver, describe the fabrication and verification of bandpass filters and sideband separating modules, and present the design and testing of a digital back end spectrometer which uses field-programmable gate array devices.

When complete in late 2011, the new receiver will allow the continuation of the blazar monitoring program with the addition of spectral and polarization information, which will be invaluable in understanding blazar jet emission.

Contents

List of Figures

Chapter 1

Introduction

1.1 Blazars

Active Galactic Nuclei (AGN) are compact regions at the centers of galaxies that emit much higher radiation than normal across the entire electromagnetic spectrum. They are believed to derive their power from rotation and accretion disks surrounding a rotating supermassive black hole. Those that emit strongly in the radio region are termed radio-loud, and show emission from two highly collimated relativistic jets which emerge along the spin axis of the accretion disk and black hole in opposite directions (Figure 1.1).

Blazars are those AGN in which the jets are aligned at a small angle to our line of sight. Small variations in this orientation angle are responsible for their wide range of observational properties, including rapid variability, high polarization, and apparent superluminal features. Despite decades of observations, there is no accepted model for the composition, acceleration, and confinement of these relativistic jets. Multi-wavelength observations of blazars will allow us to better understand these important features.

1.2 Monitoring Program

Since mid-2007, the 40-Meter Telescope at the Owens Valley Radio Observatory (OVRO) has been monitoring nearly 1200 gamma-ray blazar candidates at 15 GHz [1]. These sources from the Candidate Gamma-Ray Blazar Survey have been sampled about twice per week, giving 3-year light curves which are being cross-correlated with gamma-ray light curves (100 MeV to a few hundred GeV) of the same sources from the Large Area Telescope on the Fermi Gamma-Ray Space Telescope. This monitoring program will allow us to derive the intrinsic radio and gamma-ray characteristics of the blazar population, including the dependence of variability on redshift and luminosity and the combination of radio properties that could predict gamma-ray luminosity. The essential physics behind the correlation between radio and gamma-ray data is to characterize the time lag between peaks

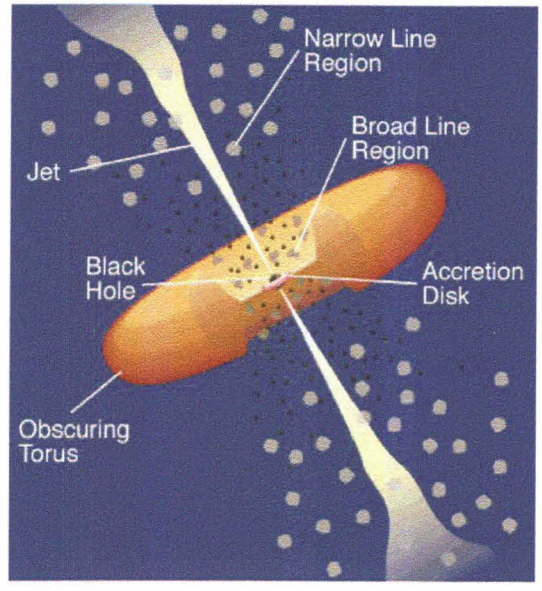

Figure 1.1: Schematic of an Active Galactic Nucleus.

in the light curves from the two regimes. This will help to constrain the location of the emission mechanisms within the jet, and thus aid in understanding the jet composition.

1.3 The OVRO 40-Meter Telescope

The OVRO 40-Meter Telescope (Figure 1.2), built in 1966, is an f/0.4 parabolic reflector that is actually 130 feet in diameter. The present receiver features a Dicke-switched cryogenic high electron mobility transistor (HEMT) receiver with a noise temperature of 30 K that observes a single band from 13.5 to 16.5 GHz. The two symmetric off-axis feed horns produce two Gaussian beams of 157" FWHM, separated by 12.95' in azimuth, and select left circular polarization. The overall aperture efficiency is $\eta_A = 0.25$. We are building a new pseudo-correlation polarimeter for this telescope that will give Stokes I, Q, and U measurements with megahertz spectral resolution in twelve 500 MHz bands between 12 and 18 GHz (Ku-band). The new receiver will feature a noise temperature of 20 K, allowing for better-sampled blazar light curves in addition to the polarization measurements. In addition, the use of reprogrammable digital electronics allows users to tailor the back end to their needs, providing flexibility in case future observational requirements change.

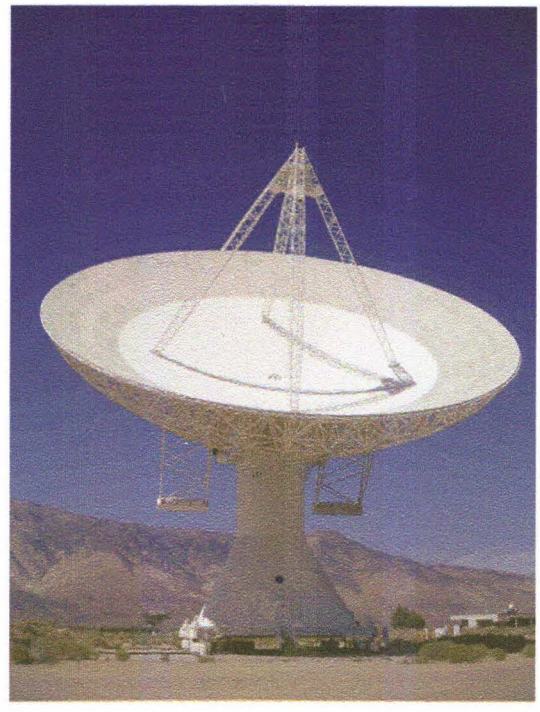

Figure 1.2: OVRO 40-Meter Telescope.

1.4 Overview

In this work, I present the design, assembly, and testing of various components in the receiver chain. I first review the necessary background to understand radio polarimetry, vector network analysis, and digital signal processing in the context of the new receiver. I then evaluate the suitability of a commercial radio frequency (RF) over optical link for use in the receiver, describe the fabrication and verification of bandpass filters and sideband separating modules, and present the design and testing of a digital back end spectrometer which uses field-programmable gate array (FPGA) devices to digitize analog data, calculate polarization information, and output the Stokes parameters to an Ethernet connection.

Chapter 2

Background

2.1 Radio Polarimetry

An astronomical receiver detects electromagnetic radiation, which consists of transverse vibrations of the electromagnetic field. If there is a lasting phase and amplitude relation between these vibrations, the wave is polarized. In general, astronomical sources are partially polarized. Polarization can be caused by several mechanisms; in active galaxies, for example, synchrotron emission gives rise to almost entirely linear polarization.

The Stokes parameters provide a convenient framework in which to describe the polarization of an electromagnetic field. Consider a time-varying electric field detected by orthogonal linear components, $E_x(t) = a_x(t)\cos(\omega t + \delta_x)$ and $E_y(t) = a_y(t)\cos(\omega t + \delta_y)$ [2]. Then the Stokes parameters are defined as

$$\begin{aligned} I &= \langle a_x^2(t)\rangle + \langle a_y^2(t)\rangle \\ Q &= \langle a_x^2(t)\rangle - \langle a_y^2(t)\rangle \\ U &= 2\langle a_x(t)a_y(t)\cos(\delta_x - \delta_y)\rangle \\ V &= 2\langle a_x(t)a_y(t)\sin(\delta_x - \delta_y)\rangle, \end{aligned}$$

where the brackets indicate time averaging. Stokes I represents the total power detected, while Q, U, and V together specify the polarization state. Since light is often unpolarized, $I^2 \geq Q^2 + U^2 + V^2$ [3]. Stokes Q and U together completely specify the linear polarization, and Stokes V represents the circular polarization state. The total polarized fraction is $p = \sqrt{Q^2 + U^2 + V^2}/I$, the total linear fraction is $p_l = \sqrt{Q^2 + U^2}/I$, and the circular fraction is $p_c = V/I$. The position angle of the linear polarization is given by $\chi = \frac{1}{2}\arctan(U/Q)$. Representing the fields as complex voltages, we have $E_x(t) = a_x(t)e^{i(\omega t + \delta_x)}$ and $E_y(t) = a_y(t)e^{i(\omega t + \delta_y)}$ and thus

$$I = \langle |E_x(t)|^2\rangle + \langle |E_y(t)|^2\rangle$$

$$
\begin{aligned}
Q &= \langle |E_x(t)|^2 \rangle - \langle |E_y(t)|^2 \rangle \\
U &= 2\langle \Re\{E_x(t)E_y^*(t)\} \rangle = \langle E_x(t)E_y^*(t) \rangle + \langle E_x^*(t)E_y(t) \rangle \\
V &= 2\langle \Im\{E_x(t)E_y^*(t)\} \rangle = -i[\langle E_x(t)E_y^*(t) \rangle - \langle E_x^*(t)E_y(t) \rangle].
\end{aligned}
$$

The new receiver for the 40-Meter Telescope uses an orthomode transducer to detect circular polarization, so transforming to a new orthogonal basis of left and right circular polarization, $\hat{l} = \frac{1}{\sqrt{2}}(\hat{x} - i\hat{y})$ and $\hat{r} = \frac{1}{\sqrt{2}}(\hat{x} + i\hat{y})$, we obtain

$$
\begin{aligned}
I &= \langle |E_l(t)|^2 \rangle + \langle |E_r(t)|^2 \rangle \\
Q &= 2\langle \Re\{E_l(t)E_r^*(t)\} \rangle = \langle E_l(t)E_r^*(t) \rangle + \langle E_l^*(t)E_r(t) \rangle = \langle |E_l(t)| \rangle \langle |E_r(t)| \rangle \cos(\delta_l - \delta_r) \\
U &= -2\langle \Im\{E_l(t)E_r^*(t)\} \rangle = \langle E_l(t)E_r^*(t) \rangle - \langle E_l^*(t)E_r(t) \rangle = \langle |E_l(t)| \rangle \langle |E_r(t)| \rangle \sin(\delta_l - \delta_r) \\
V &= \langle |E_l(t)|^2 \rangle - \langle |E_r(t)|^2 \rangle.
\end{aligned}
$$

Since active galaxies do not emit circular polarization, the receiver does not compute V.

2.2 Vector Network Analysis

One of the most common measuring tasks in RF engineering involves the analysis of circuits. A network analyzer achieves this by generating a sinusoidal test signal that is applied to the device under test (DUT) [4]. Assuming that the DUT is linear, the analyzer measures it response, which is also sinusoidal. A scalar network analyzer measures only the amplitude difference between wave quantities, while a vector network analyzer (VNA) measures both the amplitude and phase of the wave quantities. We first consider a one-port device, and distinguish between the incident wave a and the reflected wave b. The incident wave propagates from the analyzer to the DUT, and the reflected wave travels in the opposite direction. The wave quantities have units of $\sqrt{\mathrm{W}}$, so to detect true power we must measure $|a|^2$ and $|b|^2$. The reflection coefficient $\Gamma = \frac{b}{a}$ represents the ratio of the incident to reflected wave.

In general, Γ is complex and can be calculated from the complex impedance Z. The standard reference impedance is $Z_0 = 50\ \Omega$. The normalized impedance is $z = Z/Z_0$, and can be used to determine the reflection coefficient

$$\Gamma = \frac{z-1}{z+1}.$$

We now consider a two-port device. Besides the reflection at the two ports, we also admit the possibility of transmission in the forward and reverse directions. In comparison to the reflection coefficient, the scattering parameters (S-parameters, Figure 2.1) S_{11}, S_{12}, S_{21}, and S_{22} are defined as the ratios of the respective wave quantities. For forward measurements, we attach a reflection-free

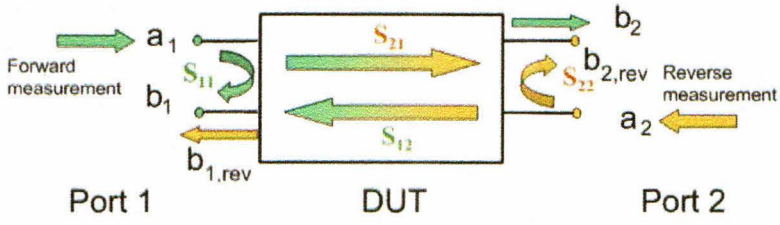

Figure 2.1: Two-port device with its wave quantities.

termination $\Gamma = 0$ on port 2, implying that $a_2 = 0$. We then measure the input reflection coefficient S_{11} (return loss) on port 1 and the forward transmission coefficient S_{21} (insertion loss) between ports 1 and 2:

$$\begin{aligned} S_{11} &= \left.\frac{b_1}{a_1}\right|_{a_2=0} \\ S_{21} &= \left.\frac{b_2}{a_1}\right|_{a_2=0}. \end{aligned}$$

The same measurement may be done in the reverse direction, terminating port 1. Then

$$\begin{aligned} S_{12} &= \left.\frac{b_1}{a_2}\right|_{a_1=0} \\ S_{22} &= \left.\frac{b_2}{a_2}\right|_{a_1=0}. \end{aligned}$$

In general, both incident waves can be non-zero ($a_1 \neq 0$ and $a_2 \neq 0$). We consider the superposition of the two measurement situations $a_1 = 0$ and $a_2 \neq 0$ with $a_1 \neq 0$ and $a_2 = 0$. Then we have

$$\begin{aligned} b_1 &= S_{11}a_1 + S_{12}a_2 \\ b_2 &= S_{21}a_1 + S_{22}a_2. \end{aligned}$$

We can group the parameters to obtain the S-matrix and the wave quantities to obtain the vectors $\mathbf{a}$ and $\mathbf{b}$, giving $\mathbf{b} = \mathbf{S}\mathbf{a}$:

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}.$$

A VNA measures the generally complex S-matrix as a function of frequency. The magnitude of the S-parameter corresponds to the amplitude ratio of the wave quantities, while the phase corresponds to the phase difference between the wave quantities.

2.3 Digital Signal Processing in Radio Astronomy

Radio astronomy instrumentation is highly specialized. Individual instruments can take over five years to design and build, and as a result are often obsolescent even before they are deployed. Additionally, due to the advent of arrays of large numbers of small antennas, a growing trend in radio astronomy is that of high-performance real-time digital signal processing (DSP). New instrumentation must be able to process large amounts of continuous data while performing such operations as spatial correlation and wideband spectroscopy. General computing clusters, which have traditionally been used for this purpose, are not suited for these high-bandwidth calculations due to their inherently non-parallel architectures.

FPGAs are reconfigurable logic devices in wide commercial use, occupying a place between custom-designed hardware and software. They are much better suited for real-time DSP due to their spatially parallel execution model; scaling on a faster track than Moore's Law, they can offer over 100 times the computing throughput of a microprocessor-based system. The Collaboration for Astronomy Signal Processing and Electronics Research (CASPER) group at UC Berkeley has been developing FPGA-based devices specifically for radio astronomy [5]. By using commercial FPGAs and open-source DSP libraries, CASPER stays up to date with the latest hardware advances, introduces modularity in instrumentation design, and has developed a large support community for its hardware. Their boards are in use at radio observatories around the world.

2.4 Receiver Design

The new Ku-band receiver for the 40-Meter telescope was designed by Dr. Oliver King. Two horns, slightly off axis at the prime focus, detect left and right circular polarization. The dual beam structure allows subtraction of atmospheric fluctuation and LNA (low noise amplifier) gain drift. These voltages are then combined using broadband 180° hybrids before passing through LNAs provided by Dr. Sandy Weinreb. The RF signals then leave the cryostat and are transported from the prime focus to the base of the telescope using an RF over Optical link. In the base, the signals are recovered and undergo quadrature mixing to separate the 12-18 GHz signal into 500 MHz subbands. The sub-band signals are then digitized and processed using FPGA-based boards developed by CASPER. They are correlated to produce beam-differenced total intensity and linear polarization from each horn with megahertz spectral resolution. Finally, the spectra are stored for later analysis.

Chapter 3

RF over Optical Link

The suitability of an RF over Optical Link was analyzed for use in the receiver. The link is intended to route the radio signal at sky frequency from the front end, located at the prime focus of the telescope, down the feed legs to the base, where the rest of the electronics reside. This method is advantageous because optical fibers offer very low loss – sometimes on the order of 0.1 dB/km – whereas traditional coaxial cable can lose up to 20 dB/100 feet. It is also possible to multiplex the four RF signals onto one fiber. Therefore, assuming noise figures and phase shifts can be kept low, it is advisable to use optical fibers to transmit signals over the long (100 meters) distance from the prime focus to the base. We obtained a test system from Optilab, consisting of a transmitter to transform an RF signal from DC to 20 GHz to an optical signal, which is sent over optical fiber until it arrives at a receiver, at which point it is again converted into an RF signal. An amplifier may also be inserted between the transmitter and receiver should the signal levels be too low.

3.1 Results

The Optilab RF over Optical link consists of an Optilab LT-20 transmitter, an EDFA optical amplifier, and an LR-30 receiver. We used a Rohde & Schwarz ZVA-24 VNA to measure the S-parameters of this system over the range of 10 MHz to 24 GHz. We first looked at the magnitude of S_{21} and S_{11}, insertion loss and return loss respectively. Figure 3.1 shows S_{21} for a wide range of input powers, from 0 to -50 dBm. Without the amplifier, the system generally exhibits 35 dB of loss over the 100 meters of cable, and is relatively constant across the band of interest (12-18 GHz). Figure 3.2 shows S_{11} for an input power of -10 dBm. Less than -10 dB are reflected across the Ku-band.

The group delay of the system is the phase difference between the inserted and received waves, $\arg(S_{21}) - \arg(S_{11})$. We looked at the group delay with and without the EDFA amplifier (Figure 3.3). When a best-fit trend is removed, the residual group delay is less than 0.05 ns across the band. Finally, we measured the 1 dB compression point. Although amplifiers maintain constant gain for low-level input signals, they will saturate at higher power levels and become nonlinear. The 1 dB

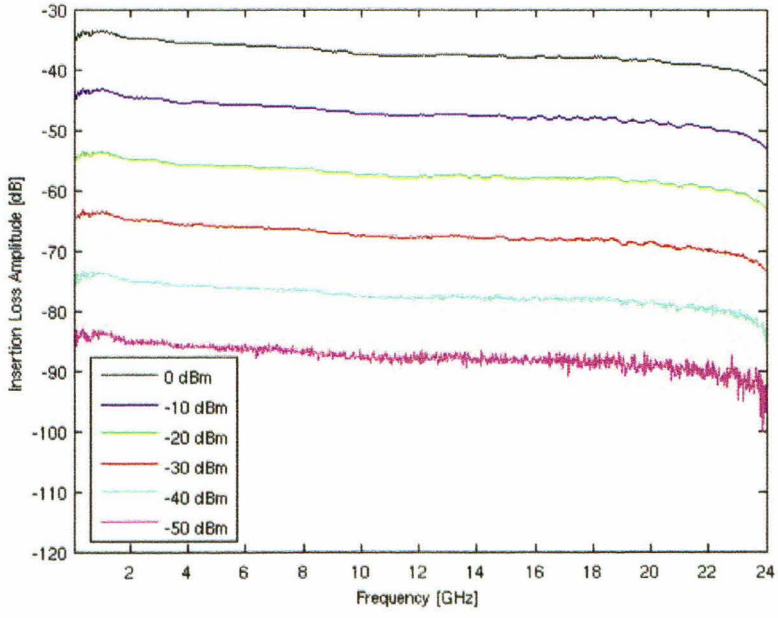

Figure 3.1: S_{21} (insertion loss) of the Optilab system.

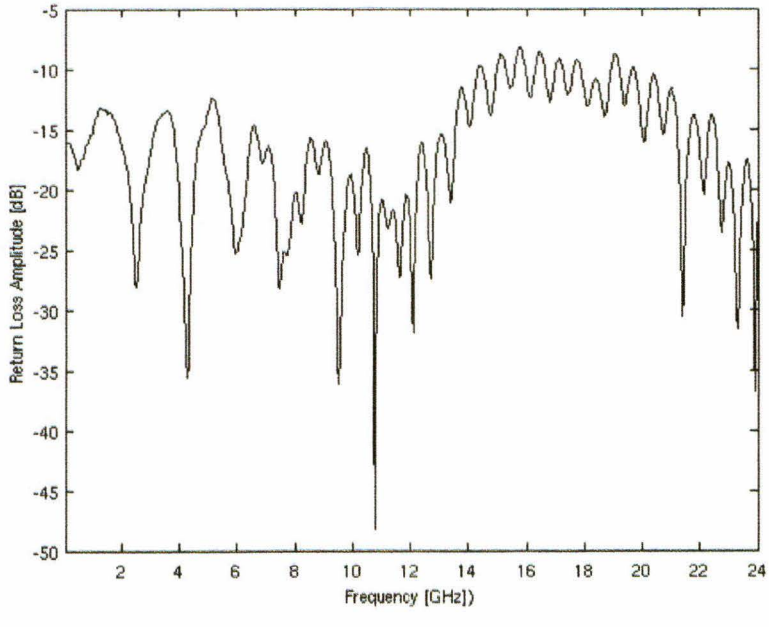

Figure 3.2: S_{11} (return loss) of the Optilab system.

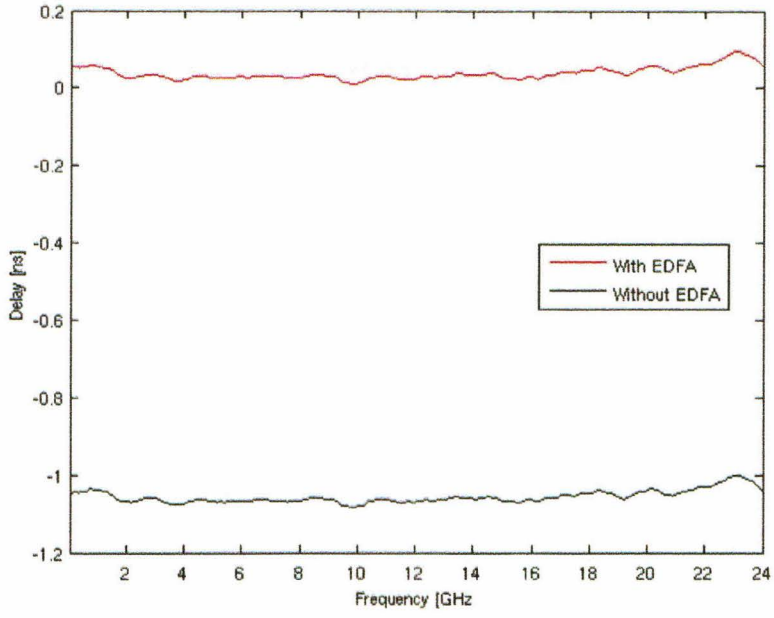

Figure 3.3: Group delay of the Optilab system, with and without EDFA.

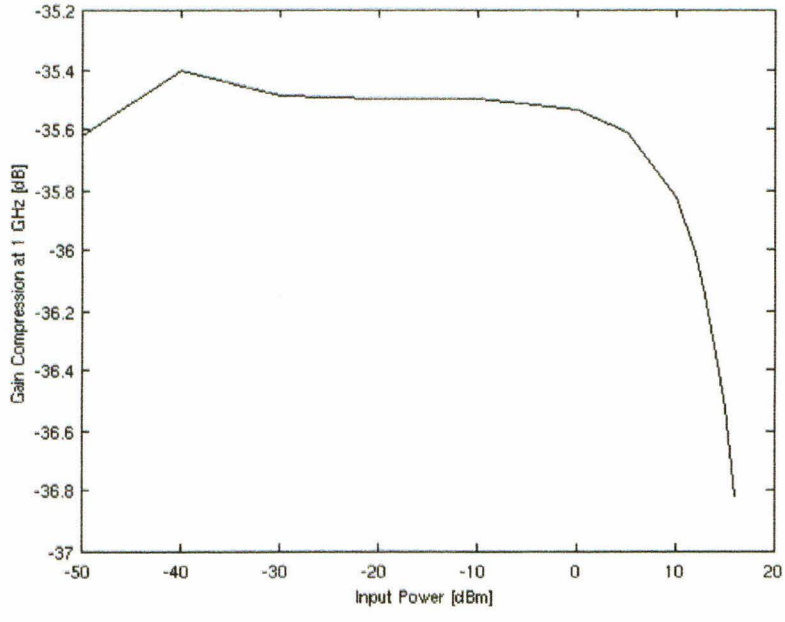

Figure 3.4: Compression at 1 GHz.

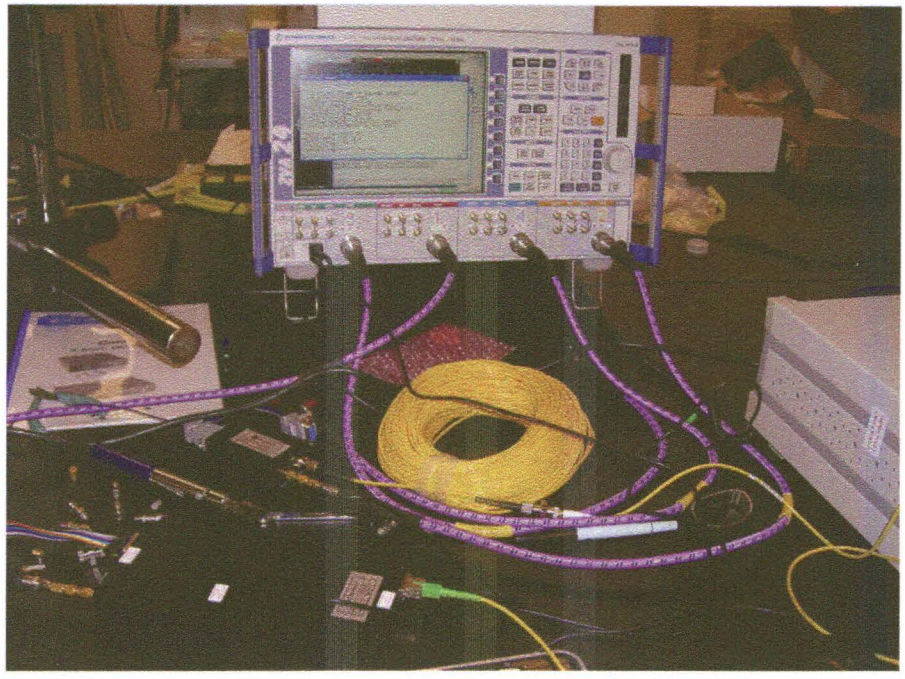

Figure 3.5: Optilab RF over Optical Link setup: VNA, LT-20 transmitter, 100 m cable, EDFA amplifier, and LR-30 receiver.

compression point is the power level that causes the system to drop in gain by 1 dB. Compression occurs around 15 dBm (Figure 3.4) at 1 GHz.

The RF over Optical link (Figure 3.5) may be considered suitable for use in the new receiver if it exhibits several properties: there must be enough dynamic range for the variety of expected signals, and there must be less than 5° variation in phase across the band. From these measurements, in addition to others made by Dr. Oliver King, the transmission amplitude, noise level, and linearity meet the requirements for the receiver. Other tests of stability over time have shown that phase is extremely dependent on the temperature of the cable; therefore, a suitable temperature control system must be constructed. This is a secondary issue, however, since all four signals are multiplexed onto the same cable and the same variations will be induced in each.

Chapter 4

Bandpass Filters and Sideband Separating Modules

Once the Ku-band signal is transported from the prime focus to the base of the telescope using the RF over Optical link, it must be converted to baseband and separated into 500 MHz sub-bands in order to be read by the digital backend. Consider an RF time-domain signal $x(t)$ in the range of 12-18 GHz arriving in the base of the telescope. It is mixed with a local oscillator (LO) of a certain frequency ω_0 and divided into two, one with a 90° phase shift. We then have two signals, $i(t) = x(t)\cos(\omega_0 t)$ and $q(t) = x(t)\sin(\omega_0 t)$. In the spectral domain, this is

$$\begin{aligned} I(\omega) &= X(\omega) \otimes \pi[\delta(\omega - \omega_0) + \delta(\omega + \omega_0)] \\ Q(\omega) &= X(\omega) \otimes \frac{\pi}{j}[\delta(\omega - \omega_0) + \delta(\omega + \omega_0)]. \end{aligned}$$

The signal is composed of the sum and difference of the RF signal and LO frequency. In order to pick out a single sideband for each signal, I and Q are sent through a 90° hybrid which performs the following:

$$\begin{aligned} I(\omega) + jQ(\omega) &= X(j\omega) \otimes \pi\delta(\omega - \omega_0) \\ jI(\omega) + Q(\omega) &= X(j\omega) \otimes \pi\delta(\omega + \omega_0). \end{aligned}$$

The signal is now separated into the upper and lower sideband. However, since the spectrum has been brought down to baseband, power will also appear reflected in the frequency axis. Therefore, a signal in the upper sideband will appear faintly in the lower sideband, and vice versa. Finally, the signals are sent through a 500 - 1000 MHz bandpass filter to isolate the frequencies outside of the sampling range of the back end. In the frequency domain, this is a convolution with the rectangular

Sideband separating mixing	
Band	LO Frequency
1	13.0
2	13.5
3	14.0
4	16.0
5	16.5
6	17.0

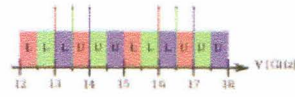

Figure 4.1: Local Oscillator schematic.

function $\Pi(\omega)$:

$$\begin{aligned} \text{Upper Sideband} &= X(\omega) \otimes \pi\delta(\omega + \omega_0) \otimes \Pi\left(\frac{\omega - 750\text{ MHz}}{500\text{ MHz}}\right) \\ \text{Lower Sideband} &= X(\omega) \otimes \pi\delta(\omega - \omega_0) \otimes \Pi\left(\frac{\omega + 750\text{ MHz}}{500\text{ MHz}}\right). \end{aligned}$$

The signals are now in the range of 500 - 1000 MHz, and can be Nyquist sampled by the back end running at 1 GHz. The spectrum calculated in the backend electronics is reversed in the frequency axis. By choosing different values of ω_0, we may pick out various bands to sample. In the new receiver design, LOs of 13.0, 13.5, 14.0, 16.0, 16.5, and 17.0 GHz cover the entire Ku-band (Figure 4.1).

4.1 Construction

Spare unmounted 500 - 1000 MHz bandpass filters were obtained from the Combined Array for Research in Millimeter-Wave Astronomy (CARMA). The filters were then soldered onto custom printed circuit boards (PCBs), connected to RF SMA connectors (Figure 4.2), and tested with a VNA to ascertain that the passband was correct. The filters are passive and require no external power.

The sideband separating modules were built from scratch. Capacitors, digital amplifiers, and a 90° hybrid were surface-mount soldered onto custom PCBs designed by Dr. Oliver King. Once the components were placed (Figure 4.3), the modules were heated in an infrared oven for 7 minutes to secure the pieces to the board. The board requires 5 volts to power the digital amplifiers. The user controls the gain by connecting to the board through a LabJack controller and Python script. Up to 32 dB of gain may be applied by controlling 5 bits from the Python script.

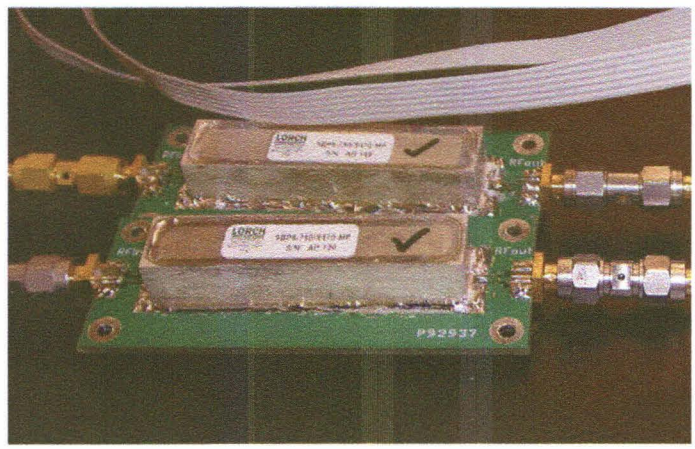

Figure 4.2: Two soldered bandpass filters.

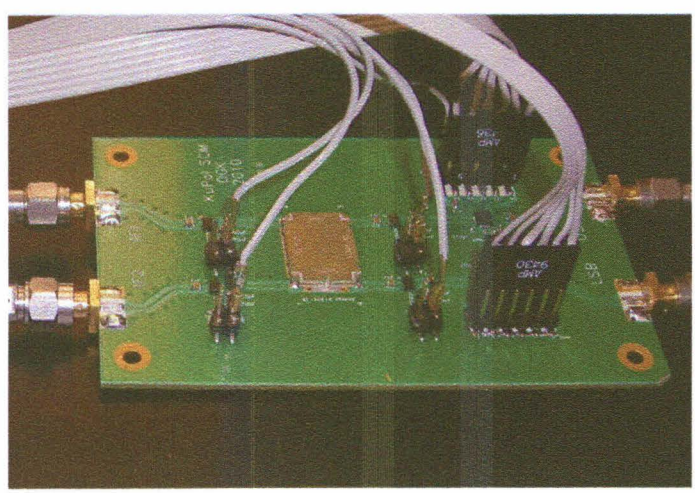

Figure 4.3: A surface-mount soldered sideband separating module.

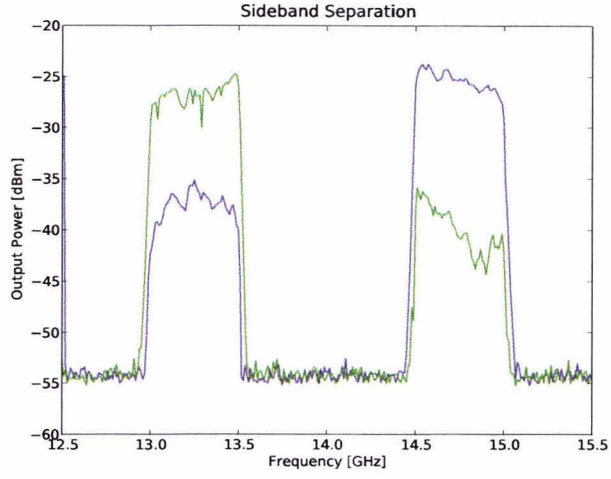

Figure 4.4: Spectral response of a 500-1000 MHz bandpass filter and sideband separating module at $\omega_0 = 14$ GHz (lower sideband in green, upper sideband in blue).

4.2 Results

To test the setup, a 0-20 GHz signal generator and 14 GHz LO were connected to a Ku-band mixer. The I and Q outputs were then sent into the two inputs on the sideband separating modules, through the bandpass filters, and finally detected on a spectrum analyzer. We measured the power output of the system by inserting a -40 dBm sinusoidal wave at sky frequency, and observing the output power from both the upper and lower sidebands. A script was written in GPIB to automate the testing procedure. The LO at 14 GHz produces two sidebands, from 13.0-13.5 GHz and from 14.5-15.0 GHz. Due to the downconversion, the output power appears in the 500 - 1000 MHz range and not in the 13.0 - 15.0 GHz range.

Figure 4.3 shows the results of both sidebands overlaid. The higher-power signal in each band corresponds to the intended domain, while the lower-power signal is the ghosting effect of the opposite sideband. The sideband separation, or difference in power between these signals, is about 10-15 dB. These figures are somewhat low for use in the receiver. It has recently been discovered that the system is starved for power from the LO. When the LO signal is amplified by 7 dB, image rejection improves to greater than 20 dB across the band.

Chapter 5

Back End Electronics

5.1 Hardware

We use the Reconfigurable Open Architecture Computing Hardware (ROACH, Figure 5.1) FPGA processing boards developed by the CASPER group at UC Berkeley to digitize the analog radio-frequency signal, compute the Stokes parameters, and output spectra to an ethernet connection. The ROACH features a Xilinx Virtex5 FPGA for signal processing and a PowerPC running Linux to control the board (programmining the FPGA, interfacing between external Ethernet devices and software registers). One may remotely log into the PowerPC through the Ethernet connection to communicate with the ROACH. Two quad data rate SRAMS provide high-speed memory, and one DDR2 SDRAM module provides lower-speed memory for the FPGA. Input is provided by two Z-DOK ports, through which two Analog to Digital Converters (iADCs, Figure 5.2) may be connected. The iADCs are also developed by CASPER and feature two input analog ports in addition to an optional external clock source. The iADCs may sample each of the analog inputs at up to 1 GHz; or, if the two inputs are interleaved, at 2 GHz.

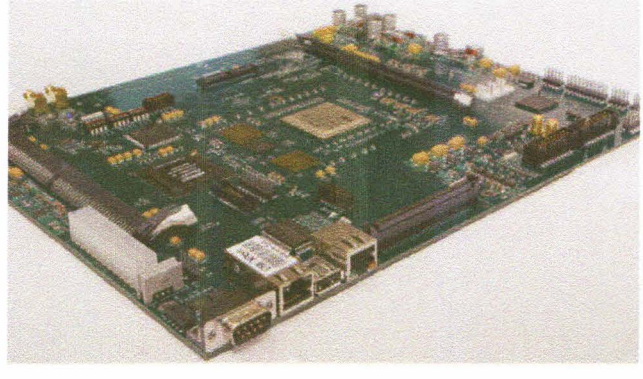

Figure 5.1: Reconfigurable Open Architecture Computing Hardware (ROACH).

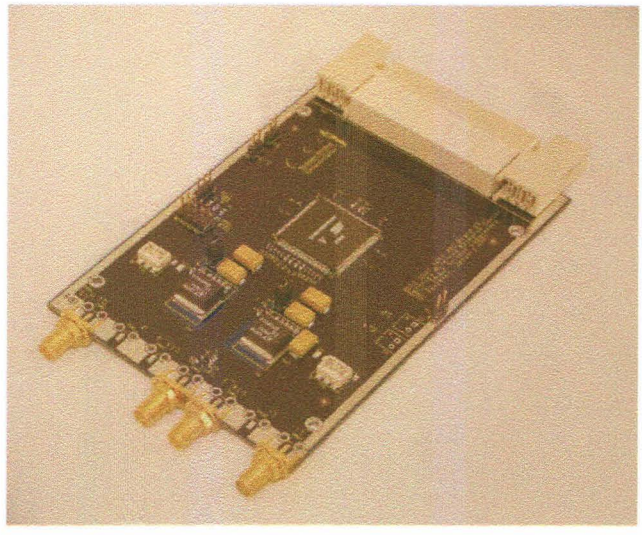

Figure 5.2: Analog to Digital Converter (iADC).

5.2 Software

The ROACH is programmed using a combination of MatLab/Simulink and libraries developed by the CASPER group. Simulink is a graphical block diagraming tool which is powered by a customizable set of block libraries, and is widely used in digital signal processing. In practice, one opens libraries developed by Xilinx (specific for the Virtex5 FPGA) and CASPER (specific for communication with the ROACH and astronomy applications), drags and drops individual blocks onto a template, connects the blocks, adjusts block parameters, and then compiles the design. Once complete, MatLab outputs a bitfile which can be copied onto the PowerPC, and then programmed onto the FPGA. In this section we walk through the FPGA design (Figure 5.3).

5.2.1 System Generator and iADCs

The System Generator block sets the clock speed of the FPGA, while the iADC block controls the parameters of the digitizers attached to the ROACH. In general, the FPGA speed is based on the desired sampling speed. For normal use, the iADC samples two analog signals at up to 1 GSample/s, and the FPGA runs at one quarter of the iADC speed. Therefore, for each clock cycle of the FPGA, the iADC feeds four parallel inputs into the FPGA.

The sideband separating modules and bandpass filters output in the 500-1000 MHz range. To Nyquist sample the 500 MHz bandwidth analog signal, we clock the iADCs at 1 GHz and the FPGA at 250 MHz. The band of interest is the 500 MHz above baseband, so the spectrum will be reversed in the frequency axis. The iADC inputs consist of combinations of circular polarizations: $A_L - B_L$ and $A_L + B_L$ into the first; $A_R - B_R$ and $A_R + B_R$ into the second.

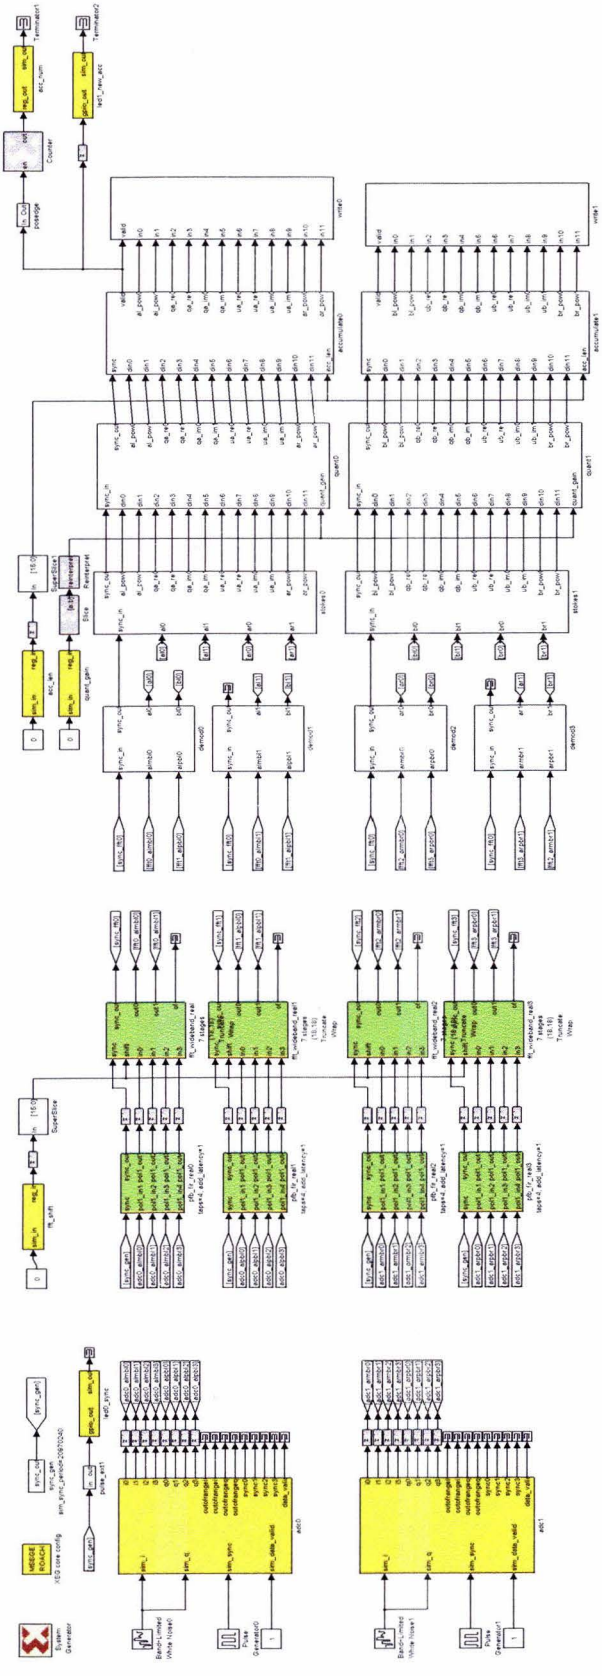

Figure 5.3: ROACH Polarimeter.

5.2.2 PFB and FFT

After digitization, the signal is sent to the Polyphase Filter Bank (PFB) and Fast Fourier Transform (FFT). These two blocks work in conjunction to transform the time-domain signal into a frequency-domain spectrum. The basic Discrete Fourier Transform (DFT) for an input time-domain signal $x(n)$ is given by

$$X(m) = \sum_{n=0}^{N-1} x(n) e^{-j2\pi nm/N},$$

where $X(m)$ is the mth output frequency component, $m = 0, 1, 2, 3, ..., N-1$ is the output frequency index (the spacing of which is determined by the sampling rate), n is the time-domain index of the input samples, $j = \sqrt{-1}$, and N is the number of samples in the input sequence [6]. Unfortunately, the DFT is extremely computationally intensive: the number of calculations scales as the square of the length of the input time series. However, the FFT radix-2 algorithm, which reduces various redundant operations, only takes $\frac{N}{2} \log_2 N$ calculations for the same size DFT. The only caveat is that they must be performed with a number of points equal to a power of two. The CASPER FFT block may output $2, 4, 8, 16, 32, 64...$ point spectra, and is only limited by the resources on the FPGA.

In most cases, the input frequencies will not fall exactly on one of the FFT bins, and thus power will leak out into the surrounding bins. This effect may be reduced by windowing, a process which reduces the sidelobes of the leakage by multiplying the input sequence by a function of the same size. For a window function $w(n)$, the FFT is now

$$X_w(m) = \sum_{n=0}^{N-1} w(n) \cdot x(n) e^{-2j\pi nm/N}.$$

A standard window function is the Hamming window, given by

$$w(n) = 0.54 - 0.46 \cos\left(\frac{2\pi n}{N}\right)$$

for $n = 0, 1, 2, ..., N-1$. After digitization by the iADC blocks, the time-domain signals are fed into the PFB block. The PFB further divides the signal into parallel "taps" and applies a finite impulse response filter to the timestream, followed by multiplication with a Hamming window. More taps used implies less inter-channel spectral leakage, but more logic is used. Each filter is lined up so that there is approximately a 3 dB amplitude decrease when they overlap. They are then fed into the FFT block, which performs the radix-2 algorithm and outputs frequency-domain spectra. We select a $2^7 = 128$-point FFT and a 4-tap PFB since we do not require extremely high spectral resolution for blazar monitoring.

When an FFT is performed on real data (as is the case for astronomical sources), it outputs

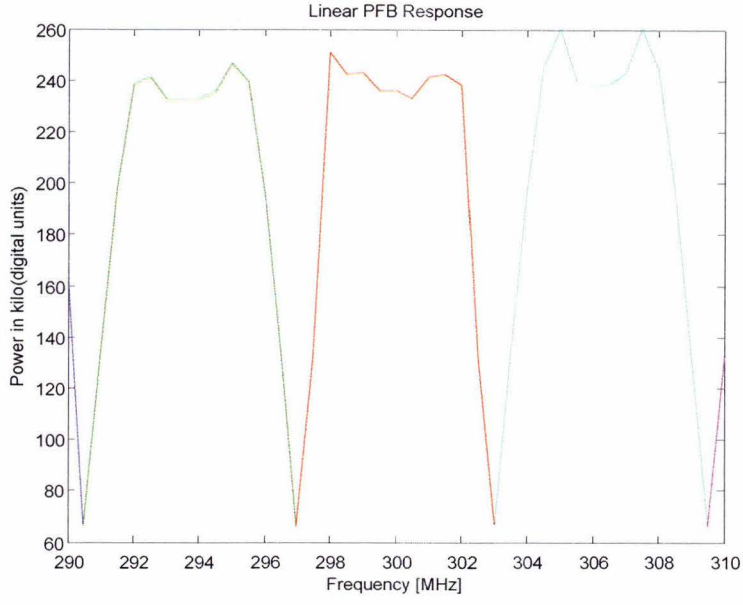

Figure 5.4: Fine spectral response of PFB/FFT (each of the 5 bins surrounding 300 MHz is a different color).

a spectrum which is symmetric in positive and negative frequencies. Therefore, in order to save resources and reduce redundant calculations, the FFT block only outputs the positive frequencies. In order to keep the timing correct, there are only two output channels – odd and even positive frequencies. Therefore, every 32 clocks of the FPGA, the PFB/FFT blocks read in 128 time-domain data points (due to the four parallel inputs from the iADC, clocked at four times the FPGA speed) and output 64 frequency-domain points. By clocking the iADC at 1000 MHz, we Nyquist sample up to 500 MHz, resulting in a spectral resolution of

$$\frac{500 \text{ MHz}}{64 \text{ bins}} = 7.8 \text{ MHz/bin}.$$

The response of the PFB/FFT combination to a sinusoidal signal over several frequency bins is seen in Figure 5.4. Each color corresponds to one of the 5 bins surrounding the 300 MHz bins. As input frequency is varied, power drops to about half of its nominal level as it crosses between bins; this occurs at 290.5, 297, 303, and 309.5 MHz, where power levels are about 65+65 = 130, approximately half of the power contained when the signal is in a single bin. This corresponds to the expected 3 dB falloff in between the PFB filter taps.

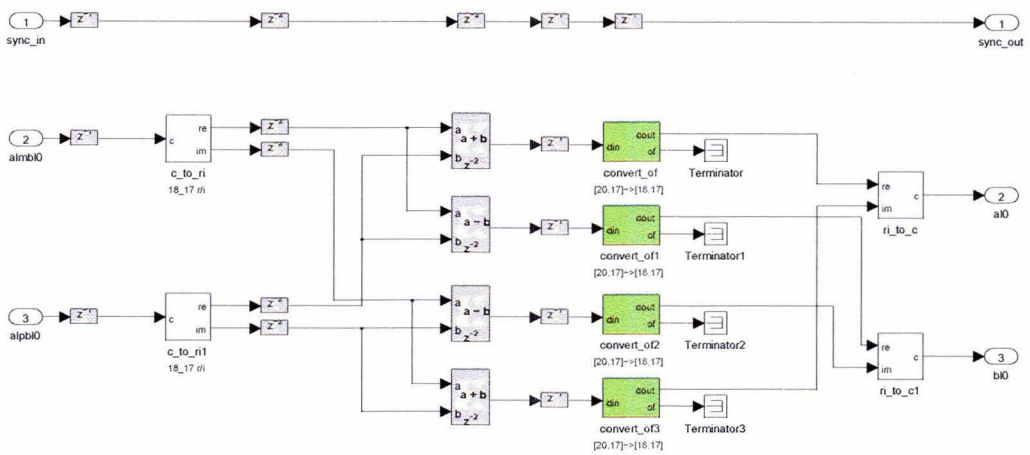

Figure 5.5: Demodulation block.

5.2.3 Demodulation Blocks

The PFB/FFT blocks output $A_L - B_L$, $A_L + B_L$, $A_R - B_R$, and $A_R + B_R$ in the frequency domain. The demodulation blocks (Figure 5.5) take these signals and recombine them to isolate the left and right circular polarizations for both horns. The left polarization block takes in $A_L - B_L$ and $A_L + B_L$, separates them into their real and imaginary components (due to the FFT), and performs

$$\begin{aligned} \Re(A_L - B_L) + \Re(A_L + B_L) &= 2\Re(A_L) \\ \Im(A_L - B_L) + \Im(A_L + B_L) &= 2\Im(A_L) \\ \Re(A_L - B_L) - \Re(A_L + B_L) &= -2\Re(B_L) \\ \Im(A_L - B_L) - \Im(A_L + B_L) &= -2\Im(B_L). \end{aligned} \tag{5.1}$$

The same process is performed for right circular polarization. The signals are then recombined into complex A_L, B_L, A_R, and B_R. Since the PFB/FFT block outputs both odd and even frequencies, in order to keep timing correct we use one block for even demodulation and one for odd, resulting in four total demodulation blocks.

5.2.4 Stokes Blocks

The Stokes blocks take in the left and right polarizations of each horn and cross-correlate them to obtain the Stokes parameters (Figure 5.6). In the horn A block, the following operations are performed:

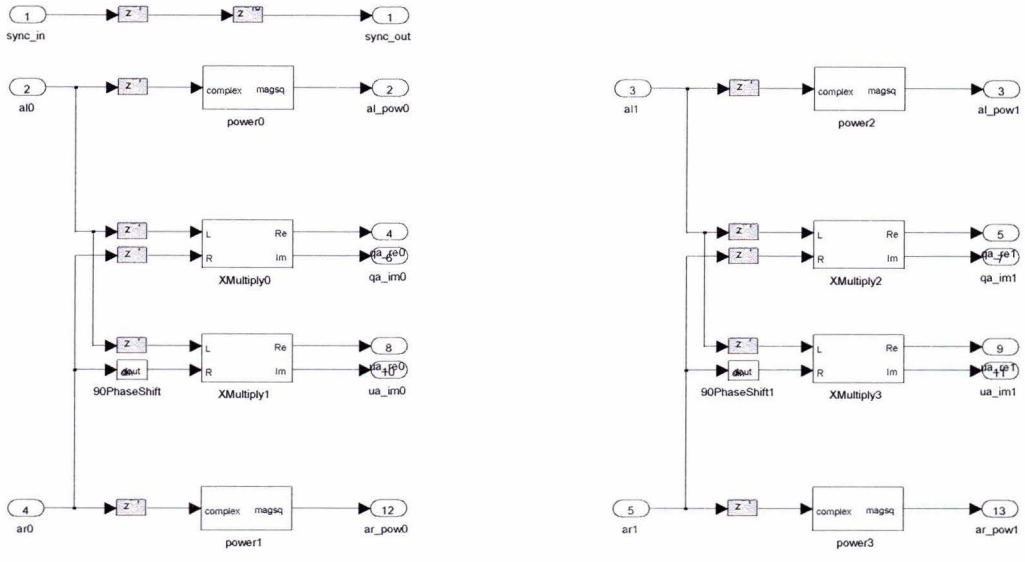

Figure 5.6: Stokes detection block.

$$
\begin{aligned}
A_L \otimes A_L &= |A_L|^2 \\
A_L \otimes A_R &= Q_A \\
A_L \otimes R(A_R) &= U_A \\
A_R \otimes A_R &= |A_R|^2,
\end{aligned}
$$

Where R is a 90° phase shift, operating on the separated real and imaginary values of A_R:

$$
R = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}.
$$

The spectra out of $|A_L|^2$ and $|A_R|^2$ are simply the the power detected, and the sum of the two gives Stokes I. The cross-correlations of A_L and A_R give complex spectra for Stokes Q and U. The spectra of interest corresponding to the expected Stokes Q and U from a real source, as calculated in Chapter 2.1, are computed in addition to other cross-correlations; all are stored for later analysis. Specifically, Stokes Q corresponds to the real part of Q_A, and Stokes U corresponds to the imaginary

part of U_A. Once we separate the inputs into their real and imaginary components, $A_L \otimes A_R$ yields

$$\begin{aligned} \Re(A_L A_R) &= \Re(A_L)\Re(A_R) - \Im(A_L)\Im(A_R) \\ \Im(A_L A_R) &= \Re(A_R)\Im(A_L) + \Re(A_L)\Im(A_R). \end{aligned}$$

The same process occurs for horn B. The Stokes blocks output 12 different spectra:

$$|A_L|^2,\ \Re(Q_A),\ \Im(Q_A),\ \Re(U_A),\ \Im(U_A),\ |A_R|^2,$$

$$|B_L|^2,\ \Re(Q_B),\ \Im(Q_B),\ \Re(U_B),\ \Im(U_B),\ |B_R|^2.$$

However, since we process odd and even frequencies simultaneously, 24 spectra emerge from the Stokes blocks to be interleaved later.

5.2.5 Quantization and Accumulation

Each spectrum is then multiplied by a factor that the user may enter into a software register. This digital gain allows the user to control the height of the spectrum, and thus find a balance between low values (yielding large artifacts) and high values (yielding large data files).

Next, the spectra are accumulated. The accumulation system is based on the FFT size: since we use a 128-point FFT, cut in half to 64 channels, and spread out over 2 parallel outputs, we obtain a full spectrum every $64/2 = 32$ clocks. There is a minimum sync period required of any CASPER PFB/FFT system determined by

$$T = \text{PFB taps} \cdot \text{LCM(reorder orders)} \cdot \frac{\text{FFT size}}{\text{no. inputs}},$$

where the reorder orders are characteristic of the FFT size. Our minimum sync period given reorders of size 2 and 5 is

$$T = 4 \cdot \text{LCM}(2,5) \cdot \frac{128}{4} = 1280 \text{ clocks}.$$

This minimum sync period is an integer multiple of the number of clocks it takes to read out a full spectrum – here, one sync accumulates $1280/32 = 40$ spectra in one period. We may accumulate spectral periods in integer multiples of the sync period; that is, the user may input an accumulation time of $T = 1280n$ clocks, where n is the number of multiples of 40 spectra to be accumulated. Since the period of a 250 MHz ROACH is 4 ns, the minimum accumulation time is $1280 \cdot 4$ ns $= 5.12$ μs. Accumulation lengths, which must be integer multiples of these, may be set in real-time by the user by writing to a software register.

5.2.6 Readout

After accumulation, each spectrum is fed into a block random access memory (BRAM) which is visible from the PowerPC. The spectrum remains on the BRAM until a new sequence is received, at which point the previous one is overwritten. Based on an interface written by Dr. Glenn Jones, a program to display the spectra was written in Python. Consisting of a graphical user interface (GUI) using the Qt backend, it runs on the host PC and connects remotely to the ROACH. It displays a window with the spectrum which can be scaled along both axes, in addition to a "waterfall plot" which shows stacked spectra in a time series. The software retrieves both the odd and even channels from any BRAM, which must be specified by the user in the GUI, and interleaves them to form a full 64-point spectrum. Every 100 ms, the GUI retrieves the latest spectrum from the BRAMs. Therefore, when used in display mode, most spectra are discarded. An upgrade is in development to retrieve and save each spectrum to disk in real time.

Chapter 6

Integration and Testing

6.1 ROACH Clocking

Although the back end is intended to run at 250 MHz, compiling a design with many multiplications and FFTs is difficult at this speed. At 200 MHz, however, almost any design will compile as long as there is enough space on the FPGA. It is therefore more efficient to make sure that the ROACH output makes sense at a lower speed, and only afterwards bring it up to the desired clock rate. In this chapter, the ROACH is clocked at 200 MHz and the iADCs are clocked at 800 MHz, meaning that only 400 MHz are Nyquist sampled.

6.2 Integration and Polarization Testing

The most important item in testing the ROACH is making sure that, given an polarized input, the software correctly computes the Stokes parameters. To test polarization, the sideband separating modules and bandpass filters are integrated with the ROACH to create a complete signal chain for the base of the telescope (Figure 6.1). Test signals are mixed with the 14 GHz LO and sent into the sideband separating modules. We see one twelfth of the total receiver bandwidth using one of the six LOs in total. Since each ROACH samples 500 MHz, i.e. one sideband, only four out of the eight possible outputs from the LO are tested. Recall that in a circular basis, the Stokes parameters calculated are given by

$$\begin{aligned} I &= \langle |E_l(t)|^2 \rangle + \langle |E_r(t)|^2 \rangle \\ Q &= \langle |E_l(t)| \rangle \langle |E_r(t)| \rangle \cos(\delta_l - \delta_r) \\ U &= \langle |E_l(t)| \rangle \langle |E_r(t)| \rangle \sin(\delta_l - \delta_r). \end{aligned}$$

Before entering the bandpass filters and sideband separating modules, the left and right polarizations from horns A and B enter two 180° hybrids, which output the sum and difference of the

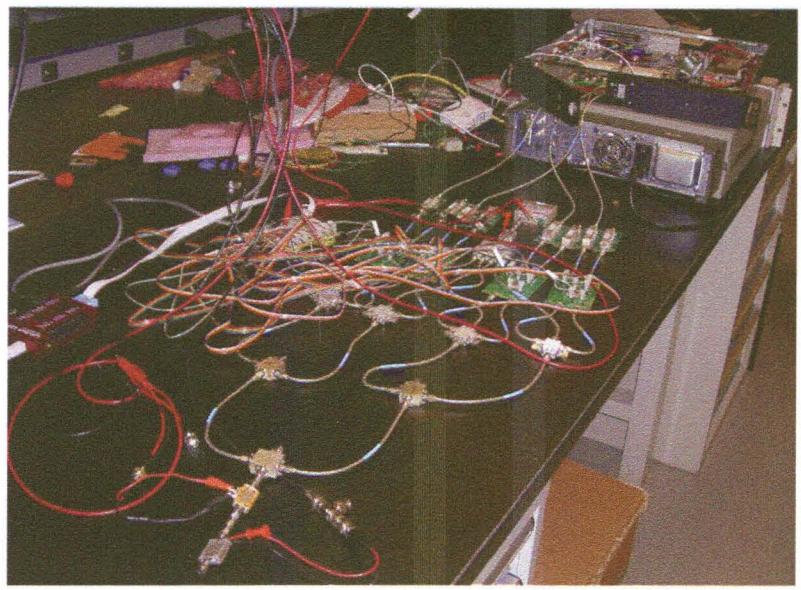

Figure 6.1: The entire signal chain. A noise diode is connected to two 180° hybrids, mixed with a 14 GHz LO, and sent into four sideband separating modules, eight bandpass filters, and the ROACH.

signals. For simplicity when testing the roach, we terminate one input on each of the hybrids so that the outputs of $A_L - B_L, A_L + B_L, A_R - B_R, A_R + B_R$ simplify to A_L, A_L, A_R, A_R.

To generate an unpolarized signal, Q = U = 0, the inputs must be uncorrelated. We use two Advanced Microwave broadband noise diodes, one for A_L and one for A_R, to generate white noise in the range of 0-18 GHz. By virtue of their separation the output spectrum should show only noiselike characteristics in Q and U. Figure 6.2 shows the spectrum of A_L given this noise source. Since the bandpass filter picks out 500 - 1000 MHz and we are only Nyquist sampling 400 MHz, a portion of the spectrum is zero. The spectrum of A_R is identical, and the sum of the two gives Stokes I. The spectra for both B_L and B_R are zero, as expected.

To generate a linearly polarized signal, it suffices to route the output of a single noise diode into both A_L and A_R. We then have $E_l = E_r$, and if the cable lengths are equal, $\delta_l - \delta_r = 0$, which corresponds to a pure Q signal. To change Q into U, a phase delay must be inserted; this may easily be accomplished by varying the length of cable to any one input.

6.3 ROACH Gain Relations

Total intensity spectra are calculated as expected. For such a signal we expect Q = 0; instead, however, the spectrum saturates (Figure 6.3). Before further polarization testing can proceed, the reason for this saturation must be determined. A likely culprit is the fact that with many the many multiplications necessary for calculation of the Stokes parameters, numbers may overflow the hard-

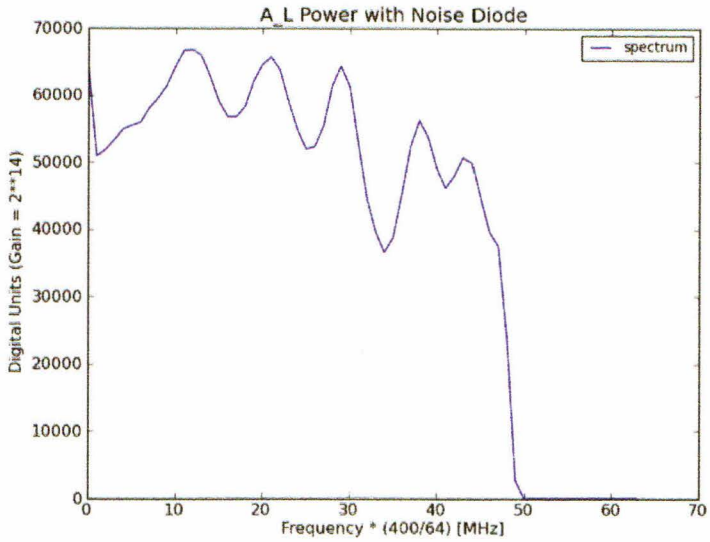

Figure 6.2: Power in horn A, left circular polarization from a noise diode.

coded bit widths. After several blocks in the design, numbers that have been growing in bitwidth are recast to lower widths in order to conserve resources. We suspect that some of these recastings choose the wrong bits and therefore discard significant numbers, resulting in an erroneous measurement.

To determine where in the design the overflow occurs, several more designs were compiled. The first is a simple spectrometer. After the FFT, power is detected and immediately accumulated into a spectrum without any further arithmetic. Since it has had much less time for bit growth, it is expected to have the most dynamic range in digital units. The second compiled design is a spectrometer with the demodulation blocks added.

Both of these designs also contain a digital gain block before accumulation. In order to determine the useful ceiling and find the range over which the designs output linearly, power into the ROACH was held constant and the gain was allowed to vary from 2^{10} to 2^{18}. Figure 6.3 shows the gain curve for the two designs, plotted just before the spectrum saturated.

By determining the saturation point as a function of depth into the ROACH design, it will be possible to correct the bit overflow errors in the full design and obtain reasonable spectra.

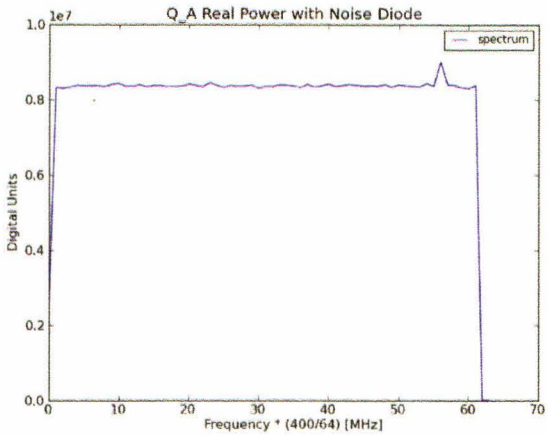

Figure 6.3: $\Re(Q_A)$, Stokes Q from a noise diode.

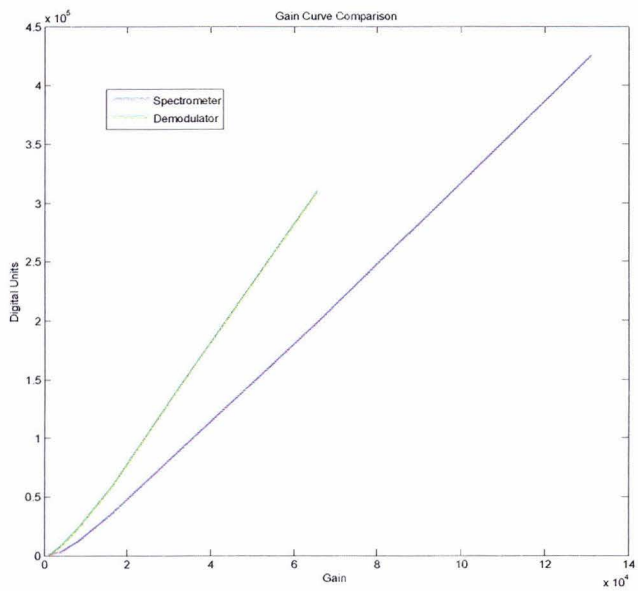

Figure 6.4: Gain curves for a simple spectrometer and a demodulator.

Chapter 7

Conclusion

In this thesis, I have presented the design, assembly, and testing of various components of a new Ku-band receiver for the OVRO 40-Meter Telescope. I tested an Optilab RF over Optical link to bring the sky signal into the base of the telescope, and determined that apart from phase variations with temperature, its noise figure and S-parameters are suitable for use in the receiver. I built several 500 - 1000 MHz bandpass filters and sideband separating modules, intended to downconvert the sky signal and separate it into two 500 MHz sidebands. The modules exhibited better than 20 dB of separation between the upper and lower sidebands from 12-18 GHz. I designed and wrote code to program the Virtex-5 FPGA on a ROACH to calculate the Stokes parameters given the input from the front end of the receiver, although it has not yet successfully compiled at the target speed of 250 MHz. The Stokes I total intensity calculation works correctly. I am now testing the dynamic range of the output given various bitwidths at different points in the design, in order to remove the saturation errors that appear in polarized spectra.

Once the ROACH code is complete and debugged, I will write software to control and visualize the back end data interactively. I will then replicate the design onto the remaining 11 ROACHes, construct the 5 other local oscillator assemblies, and connect the back and to the front end when the horns, cryostat, and amplifiers are complete. The new receiver will be on the sky and taking polarized spectra of blazars by the end of 2011.

Bibliography

[1] Richards, J., et al. 2011, accepted for publication.

[2] King, O. 2009, Doctoral Dissertation, Oxford University.

[3] Tinbergen, J., 1996, Cambridge University Press.

[4] Hiebel, M., 2005, Rohde & Schwarz GmbH & Co. KG.

[5] Parsons, A., et al. 2009, Astro2010 Decadal Survey Committee.

[6] Lyons, R., 2004, Prentice Hall.