Instrumentation for Wide Bandwidth Radio Astronomy

Thesis by

Glenn Evans Jones

In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy



California Institute of Technology Pasadena, California

2010

(Defended September 25, 2009)

© 2010 Glenn Evans Jones All Rights Reserved To my family.

Acknowledgements

Since I was very young I have been fascinated by radio telescopes, but despite a keen interest in other areas of electronics, I never learned how they were used until I began working with my research advisor, Sandy Weinreb. As a fellow electronics and technology enthusiast, I cannot imagine a better person to introduce me to the field of radio astronomy which has steadily become my career. From the very beginning, he has been a tireless advocate for me while sharing his broad experience across the gamut of radio astronomy instrumentation and observations. One of the most important ways in which Sandy has influenced me professionally is by pursuing every available opportunity to introduce me to members of the radio astronomy community.

I also very much appreciate the support provided by my academic advisor, Dave Rutledge. Dave provided a wide range of useful suggestions and encouragement during group meetings, and looked out for me from the very beginning.

Each of the other members of my thesis committee also deserves special thanks apart from my gratitude to them for serving on the committee. P.P. Vaidyanathan introduced me to digital signal processing and his clear, insightful lectures helped to inspire my interest in the field. Tony Read-head's courses on radio astronomy instrumentation and radiative processes significantly improved my understanding of the details of radio astronomy. Tom Kuiper oversaw the GAVRT digital processing subsystem, and provided me with a wide range of advice on scientific projects and observing techniques. Robert Jarnot has been extremely supportive and has provided another perspective on digital microwave spectroscopy from the field of atmospheric science.

I am fortunate to have met Tim Hankins who has shared his vast knowledge of pulsars and instrumentation with me. His initial invitation to watch over his shoulder during observations of the Crab Pulsar at Arecibo really cemented my interest in pulsars and transient phenomena. I have thoroughly enjoyed working with him on subsequent projects, as well as the many discussions we have had on pulsars, signal processing, and many other topics. Tim also introduced me to his colleague Jean Eilek and his student Jim Sheckard, both of whom provided useful discussions on this project.

The CASPER group at the University of California, Berkeley has certainly played an influential role in this thesis. I have thoroughly enjoyed working with them and learning from them. I thank Dan Werthimer for his insight and his hospitality, both in his research group and his home. Many of the current and past members of the CASPER collaboration have been extremely helpful including Henry Chen, Aaron Parsons, Peter MacMahon, Andrew Siemion, and Terry Filiba.

I would also like to thank the pulsar community as a whole for being so welcoming and willing to answer my questions. Jim Cordes, Don Backer, Joeri van Leeuwen, Paul Demorest, Willem van Straten, and Scott Ransom have been particularly encouraging and graciously provided careful answers and advice for all of my questions.

The excellent work Steve Smith and his technicians did assembling and testing the GAVRT receiver was crucial to this thesis. I look forward to continuing to work with Steve at Caltech.

Larry Teitlebaum managed the GAVRT project at the Jet Propulsion Laboratory. He always kept my thesis project as a high priority in deciding various logistical and scheduling issues, for which I am very thankful. He also provided excellent advice from his own experience on theses in general. I would also like to thank other members of the JPL community who worked on this project, including Mark Hofstadter, Dayton Jones, Farohk Bahar, and Art Freiley.

Two summer students, Robert Karl and Eythan Familier also made important contributions to this project. Robert helped design the software architecture used to control the spectrometers for the initial GAVRT system tests and Eythan helped implement and test a kurtosis spectrometer.

I would also like to thank Chris Walker and his group at the University of Arizona for providing me with the opportunity to work on the Supercam project with them. While this work did not end up in my thesis, it was an invaluable opportunity to learn about sub-millimeter astronomy as well as RF circuits and construction techniques.

Many past and current members of the Caltech MMIC group were also influential throughout this project, and have become good friends. I would particularly like to thank Hamdi Mani, Joe Bardin, Rohit Gawande, Moto Arii, Sanggeun Jeon, Feiyu Wang, Dale Yee, and Niklas Wadefalk.

I sincerely appreciate the financial support of the Lewis Center for Educational Research (LCER). In addition, the LCER staff was extremely helpful throughout the project. In particular, I thank Ryan Dorcey and John Leflang for their technical assistance and Kim Bunnell for help with administrative matters. I also appreciate the dedicated operators who watched over the telescope during long observations.

I gratefully acknowledge the Xilinx Inc. University Program for their generous donation of FPGA hardware and development software, which was essential to this project.

I appreciate the use of Arecibo Observatory, which is operated under a cooperative agreement with Cornell University. The Arecibo staff was extremely helpful during my time there, and I look forward to spending more time at the observatory in the future.

Finally, I would like to thank my family and close friends who have always inspired me to strive for my best and have supported me along the way. While I would like to list everyone here, I will opt instead to thank each of you in person. Two people deserve special mention. First, my mother, who has always made my education (along with my sister's) her top priority, for which I am always grateful. And second, my wonderful wife Arthi, who has been extremely encouraging and supportive throughout this project. I hope to return the favor as she finishes her thesis.

Abstract

Centimeter wavelength radio astronomy spans approximately two decades in frequency, from roughly 500 MHz to 50 GHz. In contrast, radio astronomy instruments have traditionally been limited to at most octave bandwidths, necessitating multiple instruments on a given telescope to cover a large fraction of the spectrum. This paradigm is infeasible for the next generation of telescopes, which will likely consist of hundreds to thousands of small dishes combined together in an array, because each receiver must be replicated for each element in the array. Therefore, wide bandwidth instrumentation must be developed for radio astronomy.

This thesis presents a novel radio telescope with excellent system noise temperature and reasonable efficiency across an instantaneous fractional bandwidth greater than 4:1; amongst the widest ever demonstrated. This instrument illustrates that extremely wide bandwidth instruments are feasible, even in the presence of terrestrial interference. To make use of the enormous instantaneous bandwidth, a flexible, high performance, and cost effective digital signal processing system is also presented. Theory, design, and measurements of special purpose digital spectrometers built to minimize the effects of terrestrial interference are included.

Aside from the practical advantages offered by wide bandwidth instrumentation, new scientific applications are also made possible. A special purpose system for performing detailed studies of giant radio pulses from rotating neutron stars (pulsars) is presented, along with demonstrations including some of the largest fractional bandwidth observations of these pulses made to date. This system includes a sensitive trigger which corrects for the dispersive effects of the interstellar medium in real time and a deep capture buffer optimized for observing repetitive transient phenomena. Measurements made using the trigger system at Arecibo Observatory are also presented to demonstrate the portability of the instrument.

Contents

Α	Acknowledgements iv							
Abstract								
1	Background and Motivation							
	1.1	Introd	uction	1				
	1.2 Wide Bandwidth Radio Astronomy			2				
		1.2.1	The Need for Wide Bandwidth Instruments for Radio Astronomy	3				
	1.3	Astro	nomical Applications	3				
		1.3.1	Continuum	4				
		1.3.2	Spectral Line	5				
			1.3.2.1 Example of a Spectral Line Observation	7				
		1.3.3	Pulsars	8				
			1.3.3.1 The Interstellar Medium	8				
			1.3.3.2 Giant Pulses	13				
2	GA	VRT I	OSS-28: A Novel Wide Bandwidth Radio Telescope	16				
	2.1	Optics	;	17				
		2.1.1	Feeds	17				
			2.1.1.1 The High Frequency Feed	17				
			2.1.1.2 The Low Frequency Feed	20				
		2.1.2	HFF Efficiency Results	21				
		2.1.3	Tertiary Response	23				
		2.1.4	Beam Patterns	24				
	2.2	Wide	Bandwidth Low Noise Amplifiers	26				
	2.3	The W	Vide Bandwidth Receiver	28				
	2.4	The D	igital Signal Processing System	31				
		2.4.1	Detailed architecture description	32				
			2.4.1.1 Digitizers	32				

			2.4.1.2 The iBOB FPGA Interface Boards	34
			2.4.1.3 The BEE2 FPGA Processing Board	34
			2.4.1.4 Clocking and synchronization	35
	2.5	Testin	g Novel Wide Bandwidth Technologies at DSS-28	36
3	Wie	le Bar	dwidth Spectral Line and Continuum Observations in the Presence o	f
	RFI	[37
	3.1	Radio	Frequency Interference	37
		3.1.1	The RFI situation at DSS-28	37
		3.1.2	In-situ RFI Surveys	38
		3.1.3	DSS-27	38
	3.2	Stabil	ity Measurements	41
	3.3	Demo	nstration of a Spectrometer with Real-time I/Q Imbalance Correction \ldots .	43
		3.3.1	Hardware Implementation and Measurements	53
	3.4	Towar	ds High Sensitivity Continuum Observations in the Presence of RFI \hdots	57
		3.4.1	RFI Detection with the Spectral Kurtosis Estimator $\hdots \hdots \hdot$	57
		3.4.2	Results	59
4	Wie	le Ban	dwidth Pulsar Observations	63
	4.1	Pulsa	Signal Processing	63
		4.1.1	Incoherent Dedispersion	63
		4.1.2	Coherent Dedispersion	64
	4.2	A Sen	sitive Trigger for Giant Pulses from a Known Pulsar	66
		4.2.1	Theoretical Discussion and Sensitivity Analysis	67
		4.2.2	System Description	70
		4.2.3	Initial Tests at DSS-28	72
		4.2.4	Results from Arecibo	73
			4.2.4.1 Testing on the Crab Pulsar	73
			4.2.4.2 Observations of B1937+21	74
			4.2.4.3 Conclusions	74
	4.3	A Dee	ep Capture Buffer for Detailed Multi-GHz Analysis of Single Pulses	76
		4.3.1	System Description	76
			4.3.1.1 The Incoherently Dedispersed Trigger	76
			4.3.1.2 The Deep Capture Buffer	77
		4.3.2	Results from Initial Testing on DSS-28	79
		4.3.3	Interpretation	83
		4.3.4	Summary and Future Improvements	91

	4.4	General Purpose Multi-Frequency Pulsar Observations at DSS-28	92						
5	Ong	Ongoing and Future Work							
	5.1	A Joint Study of Variability of the Crab Pulsar in Conjunction with the ${\rm Fermi}/{\rm GLAST}$							
		Gamma Ray Satellite	99						
	5.2	Hardware for Real-Time Searches for Radio Transients with Unknown Dispersion							
		Measure	100						
	5.3	Conclusion	101						
Bi	bliog	graphy	102						

Chapter 1 Background and Motivation

1.1 Introduction

This thesis presents the end-to-end design of a novel wide bandwidth instrument for radio astronomy. Details of the front-end portion of the system, including the optics, antenna feed, low noise amplifiers, and receiver electronics, are presented in the first sections of chapter 2, along with initial measurements of the performance of the telescope. The remainder of chapter 2 focuses on the digital back-end components, including the digitizers, FPGA processing boards, and high-speed digital interconnections.

Chapter 3 describes the problems that plague radio astronomy in general, and wide bandwidth instruments in particular, including terrestrial radio frequency interference (RFI) and sources of instability. Techniques to minimize the effect of RFI are also discussed and demonstrated, including digital quadrature imbalance correction to improve image rejection and real time statistical identification of RFI.

After a brief discussion on pulsar signal processing, chapter 4 focuses on the unique pulsar observations that this instrument is capable of making, and the digital hardware necessary to make them possible. This includes a sensitive trigger system for detecting giant radio pulses while excluding impulsive interference and a deep capture buffer for recording the baseband data of the radio pulse for off-line processing. In addition to giant pulse observations, measurements of the time and frequency evolution of pulsar signals over wide bandwidths are demonstrated.

Finally, chapter 5 concludes with a description of some ongoing projects which will make use of this instrument, including a long term study of the radio and gamma-ray variability from the Crab pulsar in conjunction with the *Fermi/GLAST* gamma-ray telescope.

The remainder of this chapter provides background on trends in radio astronomy instrumentation, as well as the various kinds of radio astronomy signals which are studied, with emphasis on the advantages offered by wide bandwidth instruments.

1.2 Wide Bandwidth Radio Astronomy

Radio astronomy provides a unique window into the universe. One unique feature of radio astronomy is the broad range of wavelengths it covers, roughly five orders of magnitude from 10 MHz to 1 THz. In comparison, ground-based optical astronomy spans merely an octave in wavelength. This breadth leads to some divisions in the field based on the different practical requirements of observing at a given wavelength range. For example, at meter wavelengths, the ionosphere presents a problem, while atmospheric moisture is of paramount importance in sub millimeter observations. This thesis focuses on centimeter wavelength radio astronomy. The radio window is also unique because many phenomena are most readily observed at these frequencies, and in some cases are unobservable by any other means. For example, neutral hydrogen which permeates galaxies can only be directly observed at radio wavelengths (e.g., Condon and Ransom, 2008). Radio waves can propagate through dust and gas that obscures other wavelengths. Finally, some rotating neutron stars, called pulsars, are most easily detected by their radio emission.

Radio astronomy signals are in general extremely weak, necessitating large amounts of collecting area and the most sensitive receivers. The development of radio astronomy has proceeded largely in step with the development of microwave communication and radar technology. However, radio astronomy instrumentation is consistently pushing the envelope as new scientific problems demand more advanced instruments. The bandwidth of radio astronomy receivers has grown from a few hundred kilohertz, through megahertz, to gigahertz today. As described later, many radio astronomy observations do not strictly require large bandwidths, but in general wider bandwidths are more useful. The bandwidth of these instruments has been limited by the receivers and antennas, and in some cases by the back-end system used to process the data. Improvements to the bandwidth of antenna systems has largely come about by new techniques enabled by advances in computer aided design (CAD) and manufacturing. Receivers have benefited both from advances in amplifier technology, in particular the exponential growth of transistor performance, and from design techniques such as microstrip which again grew in popularity and practicality with CAD. Back-end systems have progressed from strip chart recorders driven by square-law total power detectors and huge banks of handmade analog filters to digital spectrometers and correlators producing gigabytes of data per second.

Despite these incredible advances, new scientific problems demand further advances in instrumentation. In particular, one of the major pushes in radio astronomy today is the development of the Square Kilometer Array.¹ As its name implies, this project aims to construct a radio telescope made up of thousands or more small elements with a total effective collecting area of a square kilometer which will be two orders of magnitude more sensitive than existing instruments. The scientific

¹http://www.skatelescope.org

case for such an instrument has been carefully detailed, and includes such important problems as understanding dark energy, strong field gravity, and the origin of cosmic magnetism, to name a few (Carilli and Rawlings, 2004). Realizing this impressive instrument will be impractical without significant improvements in wide bandwidth instrumentation.

This thesis describes a wide bandwidth system which addresses many of these challenges, from the antenna feed to the receiver and digital back-end.

1.2.1 The Need for Wide Bandwidth Instruments for Radio Astronomy

Radio telescopes are generally equipped with a large number of feeds and receivers optimized for various frequency ranges. The feeds are designed to have the lowest possible system temperatures over as wide a bandwidth as practical. As technology has progressed, these front end receivers have moved from narrow band systems with less than 20% bandwidth to today's receivers which are commonly octave bandwidth. As an example, the world's largest fully steerable radio telescope, the Green Bank Telescope (GBT) in West Virginia, covers over 3 decades of frequency, from 100 MHz to 100 GHz, but to do so requires more than 10 front end receivers. In contrast, DSS-28, the telescope which is the subject of this thesis, uses just two receivers to cover approximately 500 MHz to 2 GHz and 2 to 14 GHz. Figure 1.1 shows the system temperature (which is equivalent to the noise figure) and bandwidths of the feeds available at the GBT up to 16 GHz, and of the two feeds DSS-28. The 2 to 14 GHz receiver, labeled HFF, covers nearly the full range of three available receivers, plus the gaps where no receiver is currently available. The system temperature, while not as good as attained by the GBT receivers, is impressively competetive. The lower frequency receiver, labeled LFF, cannot compete with the system temperatures of the GBT feeds at comparable frequencies because of engineering constraints discussed later. However, the sensitivity is still useful for many applications, particularly in light of the large instantaneous fractional bandwidth available.

Aside from the scientific advantages, such wide band feeds also offer many practical advantages. For large single dish instruments, having fewer feed systems reduces the weight in the receiver cabin which is a concern at many large telescopes. For large arrays, wide band feeds are a requirement to maximize the science per dollar return on the instrument. The Allan Telescope Array is an example of an instrument utilizing this advantage, though the wide band feed in that system requires refocusing for best performance at each frequency, limiting the instantaneous bandwidth (Welch et al., 2009).

1.3 Astronomical Applications

There are three broad categories of radio astronomy observations: continuum, spectroscopic, and pulsar or transient observations. These categories are defined by the characteristics of the physical



Figure 1.1: Comparison of system temperature and bandwidth of the wide bandwidth DSS-28 feeds to those available at the Green Bank Telescope. (Data from http://www.gb.nrao.edu/~fghigo/gbtdoc/greg.html.)

emission process, which in turn influences the instrumentation required. All of these areas can benefit from wide instantaneous bandwidth.

1.3.1 Continuum

In continuum observations, one is generally interested in measuring the emission flux density over at least a decade of bandwidth to characterize the physical conditions of the emission region. For example, the microwave spectrum of a synchrotron radiation source is directly related to the spectrum of energies of the electrons responsible for the emission. In general such sources do not display fine spectral detail, or such detail is not of importance to the continuum measurement, so at most a few measurements per decade of frequency are necessary to characterize the emission processes. In addition, these sources do not vary on short timescales: only a few vary on timescales as short as an hour, the so called intra-day variables (IDVs). Therefore the main advantage offered by a wide instantaneous bandwidth for observing these sources is throughput, being able to monitor many sources relatively rapidly. Additionally, there is great interest in monitoring the evolution of IDV sources across a wide frequency range to gain information about the interstellar medium (Lovell et al., 2007). Measuring multiple frequencies simultaneously can also help ensure consistent calibration.

Single dish continuum measurements have traditionally been made by simply measuring the total power in a given bandwidth. Radio astronomers typically use the Raleigh-Jeans approximation, $P \approx kT$, which is very accurate in the radio regime, to relate power to the temperature of a black body emitter. An astronomical signal causes an increase in the total power measured, $T_{sys} = T_a + T_{rx} + T_{other}$, where the system temperature T_{sys} is the total noise power of the system including the contribution from the source, T_a , and the noise of the receiver itself, T_{rx} . In practice, there will be other contributions to the system temperature, including spillover radiation from the ground, atmospheric emission, and the cosmic microwave background radiation, which we have grouped together into the T_{other} term. To measure T_a one generally subtracts measurements of T_{sys} made with the antenna on and off the astronomical source. Therefore, the sensitivity and precision of the measurement of T_a is limited by the fluctuations in the total power measurement. This fluctuation is given by the famous radiometer equation as

$$\Delta T = \frac{\kappa T_{sys}}{\sqrt{B\tau}} \tag{1.1}$$

where ΔT is the RMS fluctuation of the measurement of T_{sys} . κ is a constant, greater than or equal to one, which depends on the receiver architecture and is equal to one in the case of an ideal squarelaw detector (Rohlfs and Wilson, 1996). The product of the bandwidth B and integration time τ is effectively the number of independent Nyquist sampled measurements. The widest bandwidths (typically 10 to 1000 MHz) and longest integration times practical are used to improve sensitivity. The main limitation is stability: receiver gain fluctuations, atmospheric fluctuations, and RFI all contribute to instability. In particular, fractional receiver gain fluctuations expressed as $\Delta G/G$ add directly to yield

$$\Delta T = \kappa T_{sys} \left[\frac{1}{B\tau} + \left(\frac{\Delta G}{G} \right)^2 \right]^{1/2} \,. \tag{1.2}$$

Notice that gain fluctuations of just 1% will dominate with a $B\tau$ product of only 10⁴, corresponding to integrating for 10 ms with a bandwidth of 1 MHz. Gain and atmospheric fluctuations can be combated using various observing strategies, such as beam switching and noise switching. In a traditional total power receiver, the only defense against RFI is using a narrower band or a different center frequency. With modern digital polyphase filter banks (Vaidyanathan, 1992), it is often possible to isolate RFI in a range of frequencies which are then discarded before forming an equivalent total power by summing the clean frequency channels, taking care to keep track of the change in statistics from such an operation (e.g., Weber et al., 2005; Klein et al., 2006). This technique will be discussed further in chapter 3.

1.3.2 Spectral Line

Spectral line observations are similar to continuum observations in that one is again interested in measuring the emitted flux versus frequency. However, in this case the frequency structure is due to spectral lines, and is thus much finer. The structure of the spectral line gives information about



Figure 1.2: Spectrum around the 21 cm HI spectral line versus time as seen by DSS-28 pointed at zenith as the Milky Way passes overhead. The frequency resolution is 7.81 kHz or 1.65 km/s. The vertical line at 1420.0 MHz is likely caused by locally generated interference.

the temperature, pressure, and velocity of various populations of gas. Additionally, the relative strength of one spectral line to another can give insight into the excitation processes present. A classic example of this is the four hyperfine spectral lines of the hydroxyl molecule, OH, at 1612, 1665, 1667, and 1720 MHz. In local thermodynamic equilibrium (LTE), the ratio of intensities of these lines is expected to vary between 1:5:9:1 and 1:1:1:1, respectively, as the optical depth of the source increases (Elitzur, 1992). In practice, the observed ratios vary enormously from these expected values, indicating that the sources are not in LTE. In fact, maser emission is often observed in OH sources, clearly evidence that the source is far from LTE.

Originally, spectral line observations were undertaken with banks of narrow band analog filters. Such filters were limited in resolution, total number of channels, and stability. The advent of the digital autocorrelation spectrometer provided a much more stable and flexible system (Weinreb, 1963). Atmospheric and receiver instabilities are somewhat easier to remove from a spectroscopic observation because the observed bandwidth will generally contain a baseline level away from the spectral line which allows the gain variations to be measured and subtracted directly.



(a) The full 64 MHz analyzed bandwidth. The dashed lines show the region plotted in figure 1.2. The overall band pass shape is dominated by the digital filter used to select the center 64 MHz of the sampled 128 MHz bandwidth.



(b) Detail of the band around the galactic HI emission. The dotted line shows the fit to the bandpass structure that was subtracted.

Figure 1.3: A single spectrum from the measurement presented in figure 1.2

1.3.2.1 Example of a Spectral Line Observation

An example of a spectral line observation is shown in figure 1.2. The figure shows emission from the so-called 21 cm line from neutral hydrogen gas in our Milky Way Galaxy. The emission is due to the hyperfine spin flip transition between the state where the proton and electron in a hydrogen atom have the same spin and the lower energy state where they have opposite spin. At rest, this energy difference corresponds to a photon at 1420.40575177 MHz. The observed frequency differs from this frequency due to the Doppler effect. In particular, different populations of gas along the line of sight with different velocities will be observed at different radio frequencies. Figure 1.2 was produced by continuously measuring the spectrum while DSS-28 was pointing at the zenith. The vertical axis is time which corresponds to sweeping in right ascension at a constant declination as the Earth rotates. The declination is simply equal to the latitude of DSS-28. Therefore, each pixel in the image represents the intensity of emission at a given velocity at a point along a line of constant declination. The full bandpass is shown in figure 1.3(a). Notice that there are sources of interference even in the protected 1400 to 1427 MHz radio astronomy band. Some of this is genuine interference, but others are locally generated. The slow ripple in the band reflects the shape of the digital filter used to select the center 64 MHz of the sampled 128 MHz band. By allowing more ripple, greater out of band rejection was obtained for the same complexity of filter. The ripple has no impact on the measurement because it is easily removed by fitting and subtracting a polynomial baseline to the region of interest as shown in figure 1.3(b). Unfortunately, the measurement was not calibrated, so the intensity scale is given roughly in units of the system temperature, which is approximately 100K at 1.4 GHz. Note that this measurement was made with the High Frequency Feed, described in section 2.1.1.1, even though this is outside of its optimal frequency range.

1.3.3 Pulsars

Pulsars are understood to be rapidly rotating neutron stars, the resulting dense core from the supernova explosion of a star which originally had a mass in the range of 8 to 12 solar masses. Neutron stars have a density comparable to nuclear matter, with a typical diameter of 10–20 km and a mass of 1.4 solar masses (Lyne and Graham-Smith, 2006; Lorimer and Kramer, 2004). In general, the emission can be thought of as originating from beams attached to opposite poles of the star. When the rotation of the star causes one of the beams to sweep past the Earth, we see a pulse of radiation, as in the case of a lighthouse. Pulsar emission is remarkable in that many pulsars are observed from meter wavelength radio waves through high energy gamma rays, spanning 18 decades (Lyne and Graham-Smith, 2006). Despite years of research, no satisfactory model for the pulsar emission process exists, though various models explain some observed phenomena. Individual pulses at radio wavelengths show marked variation, as seen in figure 1.4, but amazingly for most pulsars, averaging many pulses together synchronous with the pulse period yields a stable average profile. The enormous angular inertia (of order 10^{45} g cm²) results in a very steady train of pulses, with stability rivaling atomic clocks. Accurate timing of these pulses has lead to many important results, including the best measurements confirming general relativity to date (Kramer and Stairs, 2008). One major thrust in pulsar timing research is attempting to detect or constrain the predicted stochastic gravitational wave background by long term precision timing of an array of pulsars (Jenet et al., 2006). The gravitational wave background is expected to slightly perturb pulse time of arrival at Earth. Pulsar observations at radio wavelengths can thus be broken down into two major areas: timing, and single pulse or transient observations.

1.3.3.1 The Interstellar Medium

One of the most important factors in pulsar observations is the effect of the interstellar medium. Figure 1.6 shows schematically the most important aspects of the interstellar medium and the various effects they have on the pulsar signal, along with the characteristic timescales over which they are measured. While generally thought of as empty space, the ISM is actually a cold tenuous plasma with roughly 0.03 electrons per cm³. As such, the index of refraction is frequency dependent. In addition, inhomogeneities in the density of the ISM cause scintillation and scattering of the radiation.



Figure 1.4: Spectrometer output showing individual dispersed pulses from the Vela pulsar observed with DSS-28. The spectrometer channels are 1 MHz wide and the integration time is 8 ms. The time series in the lower plot was produced by delaying the higher frequency spectrometer channels by the appropriate amount to align them with the lower frequencies, then adding.



Figure 1.5: Example of dispersion from the Vela pulsar. 447 pulses were added coherently with the pulsar period. The horizontal features are due to RFI. The dispersion measure is 67.99 pc cm⁻³. Notice that because the dispersion sweep time across this band is equal to the pulsar period, if this band were taken as a whole without correcting for the dispersion, no pulsed emission would be seen.



Figure 1.6: Schematic illustration of the effects of the interstellar medium. Based on a similar figure by Adam Deller.

Dispersion Neglecting the small effect of the Galactic magnetic field,² at an observing frequency f, the index of refraction of the interstellar medium is given by

$$\mu = \sqrt{1 - \left(\frac{f_p}{f}\right)^2}.\tag{1.3}$$

 f_p is the plasma frequency which is

$$f_p = \sqrt{\frac{e^2 n_e}{\pi m_e}} \approx 8.5 \text{ kHz} \sqrt{n_e}$$
(1.4)

where n_e is the number density of electrons per cm³, and e and m_e are the charge and mass of an electron. For $n_e \approx 0.03 \text{ cm}^{-3}$, this yields $f_p \approx 1.5 \text{ kHz}$. This frequency dependent dispersive delay causes the higher frequency components of a broadband pulse to arrive before the lower frequency components, because the group velocity is $v_g = c\mu$. Figures 1.4 and 1.5 both show very clearly the dispersion sweep of pulses from the Vela pulsar.

To calculate the dispersive delay due to the ISM, we start with the time of arrival at distance d given as

$$t = \int_0^d v_g^{-1} \mathrm{d}l \tag{1.5}$$

For pulsar observations at microwave frequencies, $f \gg f_p$, so we can use the Taylor expansion to make the simplification

$$\mu^{-1} = 1 + \frac{f_p^2}{2f^2} \tag{1.6}$$

yielding

$$t = \frac{1}{c} \int_0^d 1 + \frac{f_p^2}{2f^2} dl = \frac{d}{c} + \frac{e^2}{2\pi m_e c} \frac{\int_0^d n_e dl}{f^2} = \frac{\mathrm{DM}}{\mathcal{D}f^2}$$
(1.7)

Here we introduce the dispersion constant defined by Manchester and Taylor (1972) as

$$\mathcal{D} \equiv 2.41 \times 10^{-4} \text{ MHz}^{-2} \text{ pc cm}^{-3} \text{ s}^{-1}$$
(1.8)

and the dispersion measure defined as

$$DM = \int_0^d n_e dl \,. \tag{1.9}$$

The dispersion measure is thus the integrated column density of electrons along the line of sight and is usually expressed in units of $pc cm^{-3}$.

 $^{^{2}}$ The Galactic magnetic field causes birefringence, that is, the index of refraction differs for left- versus right-handed circularly polarized waves. This is certainly an important effect for pulsar emission, which is often highly polarized, and is seen as Faraday rotation versus frequency. This thesis does not include polarization calibrated data, thus this effect is not considered further.

Scintillation Inhomogeneities in the electron density act to distort the wavefront of the pulsed signal. If the scale size of the inhomogeneity, a at a distance D from the observer satisfies $D\lambda/d^2 \ge 1$, the observer lies in the focus of the inhomogeneity. This focusing region causes a strong increase in the observed flux known as a *scintle*. The transverse motion of the pulsar, ISM, and the earth causes these regions of focusing and defocussing to move past the observer, deeply modulating the flux of the pulsar on time scales of minutes, a phenomenon known as *diffractive interstellar scintillation* (DISS). Large scale density fluctuations will instead satisfy $D\lambda/d^2 \ll 1$, in which case the intensity is only weakly modulated. Because density fluctuations occur on many scale sizes, all of these phenomena are observed to varying degrees.

Scattering The diffractive scintillation also causes multi-path scattering of the pulsed signal. Following Lorimer and Kramer (2004), the bent paths arriving at the observer at an angle θ will cause the pulsar to appear blurred to an angular extent of approximately θ_d with an angular intensity distribution given by

$$I(\theta) \mathrm{d}\theta \propto e^{-\theta^2/\theta_d^2} 2\pi\theta \mathrm{d}\theta.$$
(1.10)

 θ_d is related to the distance to the scattering screen and the scale size of the inhomogeneity, and is proportional to f^{-2} . The excess time delay for the bent paths can be calculated geometrically using a small angle approximation for cosine as

$$\Delta t(\theta) = \frac{\theta^2 d}{c}.\tag{1.11}$$

Combining these two equations yields an expression for the observed intensity as a function of time

$$I(t) \propto e^{-c\Delta t/(\theta_d^2 d)} = e^{-\Delta t/\tau_s}$$
(1.12)

where

$$\tau_s \equiv \frac{\theta_d^2 d}{c} \propto f^{-4}.$$
(1.13)

This means that the intrinsic pulse emitted from the pulsar is convolved with a one-sided exponential with a decay time that is extremely frequency dependent, scaling as f^{-4} .

1.3.3.2 Giant Pulses

While the emission from most pulsars is so weak that individual pulses can only be detected with the largest telescopes, if at all, a few pulsars are known to emit *giant pulses* which are many times stronger than the average pulses. The pulsar at the heart of the Crab Nebula is the best known example, and in fact, was first discovered because of these very strong pulses (Staelin and Reifenstein, 1968). In the years since, these giant pulses have been extensively studied because they provide the most detailed information about individual pulses. To study these pulses in greatest detail, the raw dispersed voltage data is recorded from the telescope, and then processed using the technique of coherent dedispersion, described in section 4.1.2, to recover the original waveform from the dispersive effects of the interstellar medium (Hankins and Rickett, 1975).

Recording baseband voltage data clearly puts a constraint on the length and bandwidth of an observation. Therefore, the widest bandwidth observations are made using digital oscilloscopes with deep capture buffers which are triggered when a pulse is detected. Fortunately, pulses from the Crab pulsar are strong enough that a total power detector can be used to trigger the oscilloscope when a pulse occurs. Giant pulses from the Crab pulsar have been detected from 23 MHz to 43 GHz (Popov et al., 2006). Individual pulses are made up of *microbursts* which are between 1 to 100 μs long above 1 GHz (Crossley et al., 2009). The sub-structure of some microbursts can be resolved into *nanoshots* some of which show temporal structure that is unresolved on 0.4 ns timescales, and flux densities as high as 2 MJy (Hankins and Eilek, 2007).

The other pulsar notable for its giant pulses is the millisecond pulsar B1937+21 (Cognard et al., 1996). Measurements of its giant pulses have been reported up to 2.25 GHz (Soglasnov et al., 2004). While giant pulses from B1937+21 have also been found to have flux densities up to 100 kJy, the pulses are generally extremely narrow, less than 1 μ s. Because of this, triggering based on an increase in total power has not been very effective at capturing giant pulses over wide bandwidths.

The transient universe has become a hot topic in all areas of astronomy in recent years. Transient radio astronomy remains as one of the largest areas of unexplored parameter space, largely because the synoptic instruments used in the optical and high energy regimes are not available at radio wavelengths. Detecting and characterizing transient phenomena is one of the major science drivers for large arrays currently being built including LOFAR³, MWA⁴, ASKAP ⁵, and meerKAT⁶, and of course the future SKA.

Studying the unique emission phenomena from the Crab pulsar can provide insight (a sort of lower bound) into the types of transient emission future large scale surveys may encounter. For example, the giant pulses are so bright at lower frequencies (\sim 430 MHz) that they could be detected in other galaxies (Cordes et al., 2004). Giant pulses provide a unique opportunity to study transient emission because they are bright enough that an unknown object emitting by the same mechanism would be detectable in a blind survey over a large volume of space, yet they occur regularly enough to enable detailed study (McLaughlin and Cordes, 2003). Each successively detailed study of the Crab Pulsar has revealed surprising new phenomena.

To date, giant pulse observations have been limited in instantaneous bandwidth, and thus ul-

³http://www.lofar.org

⁴http://www.mwatelescope.org

⁵http://www.atnf.csiro.au/projects/askap/

⁶http://www.ska.ac.za/

timate time resolution, by the feeds and receiver systems available at large radio telescopes. Traditionally, feeds have been limited to fractional bandwidths of generally less than an octave. Thus observations must be done at higher frequencies to obtain large bandwidths. Until recently, spectrometers and pulsar backends at many observatories were capable of processing only a few hundred megahertz of bandwidth, thus intermediate frequency processors have been designed with limited bandwidths. Furthermore, radio frequency interference typically limits usable bandwidths to less than one gigahertz. Lower frequency observations have been limited by capture depth of the digital oscilloscopes, which must be longer than the dispersion smearing time across the desired bandwidth. Digital oscilloscopes also limit the duty cycle of observations because they require substantial time to off-load captured data before being ready for the next event. Many giant pulses can be missed while data is being off-loaded to a computer for storage and later processing. The DSS-28 system described herein addresses all of these areas.

Chapter 2

GAVRT DSS-28: A Novel Wide Bandwidth Radio Telescope

The antenna dish at the heart of this thesis is the 34 meter DSS-28 located at the Goldstone Deep Space Communication Complex near Barstow, California. This antenna was mothballed for many years before it was offered to the Lewis Center for Educational Research (LCER)¹. The LCER operates the Goldstone Apple Valley Radio Telescope (GAVRT) program which aims to enrich K-12 science education through radio astronomy. Since its inception, the GAVRT program has allowed students from all over the world to remotely control via the Internet the 34 meter DSS-12 telescope to perform radio astronomy observations. The LCER partners with radio astronomers with scientific projects which would benefit from the large amount of observing time available through GAVRT. The astronomer is responsible for helping to develop a curriculum around the observations, and also interacts with the students to help them understand the importance of the measurements they are helping to make. DSS-12 is equipped with a legacy total power receiver providing up to 50 MHz of bandwidth at around 2 and 8 GHz. The observational campaigns have generally been limited to monitoring flux variations of quasars and the outer planets. The students measure these sources using cross scans through the source and by making raster scanned maps.

When DSS-28 was offered to the LCER, it was realized that the program could be greatly expanded by getting more astronomers involved if a more versatile receiver system were available. Instead of equipping DSS-28 with a traditional narrow band total power radiometer as available on their original antenna, DSS-12, it was proposed that a novel wide band radiometer could be constructed to fit the same budget, yet provide a decade of frequency coverage. A digital signal processing system was also proposed to take full advantage of the capabilities of this instrument. This system employs many novel techniques, from the optics through the digital back end.

¹http://www.lcer.org

2.1 Optics

DSS-28 was originally constructed as a high power transmitting antenna (Imbriale, 2007). To this end, the Cassegrain optics of this antenna were designed such that the beam does not come to focus until deep into the beam waveguide. As a fortunate consequence, the radiation is approximately a plane wave at the entrance to the beam waveguide at the vertex of the dish. This allows an offset paraboloid to be used to focus the beam onto a feed mounted to a frame on the surface of the dish. Moreover, this offset paraboloidal tertiary can be rotated to direct the beam onto one of many feeds which can be mounted around the vertex of the dish. Currently the system is equipped with two feeds, but in the future other feeds could be added to test their performance or to expand the capabilities of the system. The optics and feed system are shown in figure 2.1.

2.1.1 Feeds

As with any radio telescope, wider bandwidth feeds help maximize the capabilities of the instrument for a given budget. The premier radio observatories all use highly optimized feeds to get the maximum sensitivity because the scale of the project is so large that the incremental cost of another feed is relatively insignificant. In contrast, the limited budget of DSS-28 provided a strong incentive to use the widest bandwidth feeds practical. This is a very relevant constraint to future large arrays, where each feed must be replicated across hundreds or thousands of small dishes. As such, DSS-28 provides a unique opportunity to demonstrate the kind of wide band technology necessary for large arrays like the SKA. While the original hope was to cover a decade of bandwidth with a single feed, at the time of the design, the best choice in terms of maximizing performance given the budget was to use two feeds to cover roughly 0.5 GHz to 4 GHz, called the Low Frequency Feed (LFF) and 2 to 14 GHz called the High Frequency Feed (HFF).

2.1.1.1 The High Frequency Feed

The High Frequency Feed is based around a commercially available feed designed for antenna pattern measurements and manufactured by ETS-Lindgren (Rodriguez, 2006). The beam pattern of the feed was measured at Caltech and found to have good performance. In particular, surrounding the feed with a metal can as a model of the Dewar and placing absorbing strips along the outer edge of the feed fins yielded a very constant beam width versus frequency. This is important to provide the best performance across frequency because a narrow beam will under-illuminate the reflector, yielding reduced efficiency, while a wide beam will see more spillover radiation from the ground and degrade the system noise (Baars, 2007). For an uncooled feed, the absorbing strips would significantly impact the noise of the system, but fortunately the High Frequency Feed is physically small enough that it could be cryogenically cooled.



Figure 2.1: The DSS-28 optical system as seen from the surface of the dish. The dashed red line shows the approximate path of an incoming ray. The dashed yellow line shows the axis of rotation of the tertiary. The box containing the receiver electronics can be seen in between the High and Low Frequency Feed support towers. The entrance of the original beam waveguide is also visible.



Figure 2.2: View of the HFF with the surrounding Dewar removed. Photo courtesy of Hamdi Mani.

The HFF cryogenics package is shown in 2.2. The feed is mounted directly to the 12 kelvin stage of a standard CTI-350 cryogenic refrigerator. The output from each polarization is fed into a 30 dB coupler used to inject calibration signals into the LNA which immediately follows the coupler. In the prototype system, the calibration noise diodes were housed in the receiver box. This configuration adversely affected the calibration signals because of the significant length of cabling required. In the final system, it is likely that the noise diodes will also be mounted in the cryogenics package to improve stability. A second coupler will be included to allow external injection of the phase calibration comb generator signal.

Designing a window for the HFF cryogenic Dewar proved to be a significant challenge. The window must be made large enough to avoid interfering with the pattern of the feed. This requires that the window material be sufficiently strong to withstand approximately 1200 pounds of atmospheric pressure, yet be transparent to microwaves. In other applications, a number of rigid materials could be used by applying an anti-reflective impedance matching coating to the window. However, such techniques are not readily applicable to feeds covering more than an octave of bandwidth. After trying a number of plastic films and foams, a 15 mil mylar film proved to be the best solution. The window was found to be largely transparent to infrared radiation, which significantly loaded the cold stage causing the temperature to rise. The window also became very cold itself and accumulated moisture and ice, degrading the system noise temperature. To solve this problem, a blanket made of layers of Teflon film and paper was placed between the feed and the window.

2.1.1.2 The Low Frequency Feed

The LFF is also based around an ETS-Lindgren feed, which is essentially a scaled version of the HFF. Because of its significantly larger size, it was deemed impractical to cool the feed itself. However, the LNAs are cooled to 50 K in a small Dewar by a Stirling cycle refrigerator. These refrigerators require no maintainance and are extensively used by the mobile phone industry. The architecture of the LFF is exactly the same as that of the HFF, with a cooled LNA and coupler for calibration signals. During the design phase, the efficiency of the LFF was predicted by scaling the calculations of the HFF performance, made using measured data. Unfortunately, after measuring the performance of the LFF it was realized that secondary reflector on the telescope was undersized for the frequency range of the LFF. Physical optics calculations were then performed and found to agree with the observed linear decrease in efficiency with increasing wavelength. The aperture efficiency was measured using the source Cygnus A (3C405) and is presented in figure 2.3 along with the system temperature.

Another problem with the LFF is the 50 K refrigerator is considerably less effective at cryopumping² than the 12 K refrigerator, making the Dewar extremely sensitive to the effectiveness of

 $^{^{2}}$ Cryopumping refers to the effectiveness of cryogenic refrigerators as vacuum pumps by condensing out residual gases.



Figure 2.3: Aperture and system temperature of the Low Frequency Feed measured on Cygnus A (3C405). The sharp cutoff in the efficiency measured on polarization Y is due to a low pass filter installed between the feed and LNA to combat RFI. The calibration signal was injected after the filter so the system temperature could be measured above the filter cut-off frequency. The linear drop in efficiency below 1 GHz is a result of the secondary reflector being too few wavelengths across. The large scatter in the points is largely due to RFI, which is especially problematic in this frequency range.

the vacuum seal.

These issues, along with the large amount of RFI present in the lower frequency band make for an uncertain future for the LFF. It should be noted, however, that the LFF should still be quite effective for observing giant pulses from the Crab pulsar because the flux from the surrounding nebula is so great that it will still dominate the system temperature despite the low efficiency.

2.1.2 HFF Efficiency Results

The efficiency of the High Frequency Feed was calculated by measuring the response of the system to the radio galaxy 3C405 and Venus. The 3C405 measurements were made at an elevation of approximately 23 degrees, while those on Venus were made at approximately 62 degrees. After correcting any pointing offsets by manually peaking the response of the telescope on the source, the telescope was moved off source by one degree in cross elevation. The high power noise calibration signal was pulsed to measure the system gain, then the telescope was commanded to scan through the source in cross elevation at a fixed rate. Finally, at the end of the scan, the noise source was pulsed again. To process the data, the system gain was estimated from the calibration signal responses using a lookup table to determine the calibration signal value in kelvin. A Gaussian was then fit to the scan data including a constant offset for the system temperature and a linear baseline to account for drift during the scan. The peak of the Gaussian provided the antenna temperature of the source, and the width of the Gaussian was used to compute the beamwidth. The beamwidth is shown in



Figure 2.4: Beamwidth of the DSS-28 High Frequency Feed measured on Venus and the radio galaxy Cygnus A (3C405). The standard rule of thumb for beamwidth, $1.2\lambda/d$, is plotted for comparison. Notice that the beamwidth becomes comparable to the angular extent of Cygnus A at around 10 GHz.



(a) Aperture enciency of both polarizations versus re- (b) Aperture enciency of polarization Y compared to quency and elevation angle.

Figure 2.5: Aperture efficiency measurements of the DSS-28 High Frequency Feed. Each point was measured using 100 MHz of bandwidth.

figure 2.4. Notice that 3C405 has an angular extent of approximately 170 by 45 arcseconds, and is thus partially resolved above 10 GHz (Ott et al., 1994). The beamwidth shows reasonable agreement to the standard estimate of $1.2\lambda/d$, which is based on the Bessel function response expected of a circular aperture (Baars, 2007).

The efficiency results are presented in figure 2.5(a). The measurements of 3C405 have been corrected for the source size as described by Baars (2007). As expected, the efficiency is generally lower for the 3C405 measurements because they were taken at a lower elevation, and thus reflect the deformation of the dish due to gravity. Both sets of measurements also show a general decrease at higher frequencies as expected for random surface errors. We expect the loss in efficiency from random surface errors to follow the 'Ruze formula,'

$$\frac{\eta_A}{\eta_{A0}} \approx e^{-\left(\frac{4\pi\epsilon}{\lambda}\right)^2} \tag{2.1}$$

(Ruze, 1966). Here η_A/η_{A0} represents the relative change in aperture efficiency from a perfect surface, ϵ is the RMS surface error, and λ is the wavelength. Figure 2.5(b) shows fits of the polarization Y data from both sources to this formula, with reasonable agreement in the general trend of the data. Interestingly, the best fit RMS error is similar at the two elevations. This may imply that the reduced efficiency is instead due to mechanical flexing of the feed support towers causing optical misalignment. Further measurements of the efficiency versus frequency and elevation are required to fully characterize the telescope.

2.1.3 Tertiary Response

To properly align the tertiary with the feed, the response of the system was measured versus tertiary angle. These measurements were performed on the methanol maser source W3OH (Menten, 1991). This source exhibits strong, narrow, spectral lines at 6.67 and 12.179 GHz. By measuring a spectral source, we can use the frequency channels away from the spectral line to measure the system noise simultaneously. The maser emission in W3OH is also extremely compact compared to the beamwidth of DSS-28, avoiding the need to apply source size corrections (Minier et al., 2001). Figure 2.6 shows the measured antenna temperature of the source and the system noise versus tertiary angle at both 6.67 and 12.179 GHz. The antenna temperature was calculated using the noise calibration diodes which were in turn calibrated using a Y-factor measurement between microwave absorber at ambient temperature and cold sky (Kraus, 1986). The antenna temperature of the source was measured by subtracting the local spectral baseline. Two features are striking in the tertiary response. First, the system noise temperature has a local minimum near the optimal angle, increasing to either side. This increase is understood to be the result of increased spillover radiation being scattered off the secondary reflector support legs. The asymmetry from one side to the other is partly due to



Figure 2.6: Measured response as a function of tertiary angle for the methanol maser lines in W3OH. The system temperatures are plotted as dotted lines while the antenna temperatures of the maser lines are solid. The 6.67 GHz measurements are plotted in blue, while the 12.179 GHz measurements are in red.

inherent asymmetry in the feed itself (W. Imbriale, priv. comm.). The asymmetry may also be partly attributed to asymmetry in the location of the feed relative to the support legs. When the angle is far from optimal, the system noise drops below the local minimum at the optimal position. This is presumably because spillover radiation seen by the secondary is now no longer focused on the feed. Further measurements are needed to characterize this effect versus antenna azimuth and elevation position. The second striking feature is the asymmetry in the antenna temperature response, which exhibits a small shoulder on the positive side. This may be due to the same asymmetric geometry which causes the asymmetric system temperature response, but may also point to some residual misalignment in the optical system.

Unfortunately, the servo system that controls the tertiary position is currently quite compliant because of a limitation in the servo controller. The wind is frequently strong enough at DSS-28 to significantly perturb the tertiary position, by angles of 2 degrees of more. This causes fluctuations in the system temperature, and is currently one of the major sources of instability in the system, as discussed in section 3.2.

2.1.4 Beam Patterns

Detailed beam patterns were also measured using the 6.67 and 12.179 GHz methanol maser emission from W3OH, shown in figures 2.7 and 2.8. The patterns were measured by scanning through the



Figure 2.7: Measured beam pattern on W3OH Methanol maser at 6.67 GHz. The heavy red curve is polarization X and the thin green curve is polarization Y.



Figure 2.8: Measured beam pattern on W3OH Methanol maser at 12 GHz. The heavy red curve is polarization X and the thin green curve is polarization Y.

source in elevation and in cross elevation. As in the tertiary measurements, variations in the system noise temperature or gain are removed by subtracting the baseline spectrum around the maser line. The patterns show some asymmetric shoulders which are characteristic of the feed itself. After the telescope control software is more mature, it will be possible to make raster scanned beam maps to further understand, and possibly optimize the sidelobe structure. Detailed measurements and understanding of the beam pattern are important for making maps of astronomical sources because the structure of the source will be convolved with the beam pattern. Routines have been developed to deconvolve the beam pattern from such maps, but their success depends critically on a detailed model of the beam pattern (e.g., Emerson et al., 1979). Masers provide an excellent means of measuring the detailed beam pattern, but are limited to selected emission frequencies. Characterizing the beam pattern over a wide bandwidth will be a difficult challenge.

2.2 Wide Bandwidth Low Noise Amplifiers

The wide band cryogenic low noise amplifiers used for the High Frequency Feed are some of the most mature technology in the system. The amplifiers, designed by Weinreb and Wadefalk are based around indium phosphide (InP) high electron mobility transistors (HEMTs), and have been the work horses of the radio astronomy and physics communities for the past ten years.

In general, amplifier design requires a trade-off between low noise and linearity at large signal levels. Radio astronomy applications demand amplifiers with the absolute lowest noise, which generally results in relatively poor linearity and therefore degradation in the presence of strong interfering signals. To this end, an RFI survey was conducted to evaluate the potential impact of the interference environment on the receiver. The results of the survey and the RFI situation are discussed in detail in 3.1. While HEMTs achieve by and large the best noise performance, like any field effect transistor, they suffer from higher flicker noise compared to bipolar devices. This noise modulates the gain of the amplifier and degrades the stability on timescales of \sim 1 second for a 1 GHz bandwidth observation (Jarosik, 1996; Gallego et al., 2004).

The LNA used for the LFF is a new design based around silicon germanium (SiGe) bipolar transistors. Recent work by Weinreb et al. (2007) has demonstrated cryogenic SiGe amplifiers with noise performance rivaling or exceeding that of HEMTs at lower microwave frequencies.

In narrow bandwidth applications, the deleterious effect of non-linearity on system performance is usually due to third-order intermodulation products which land within the band of interest, and thus pass through the intermediate frequency filters in a heterodyne receiver. With the wide bandwidth of the GAVRT system, it was found that the second-order intermodulation products are also of great concern, because the larger analyzed bandwidth will encompass them. Laboratory measurements of this effect were made with the SiGe LFF LNA and presented by Weinreb et al. (2009). When a


Figure 2.9: DSS-28 system temperature derived from the same measurements described in section 2.1.2



Figure 2.10: Photograph of the receiver box.

large interfering signal is present, the amplifier acts like a mixer with the large signal acting as the local oscillator. This effect is seen in the RFI measurements presented in 3.1.

2.3 The Wide Bandwidth Receiver

In keeping with the goal of maximizing the variety of science observations possible with DSS-28, a very flexible receiver was designed to take advantage of the unique wide bandwidth front ends. Details of the receiver architecture are shown in figures 2.11 through 2.14. Designing such a wide bandwidth receiver presents numerous challenges. For example, the local oscillator needs to be able to tune over the entire frequency range. While laboratory synthesizers are capable of meeting this requirement, they are prohibitively expensive. A relatively inexpensive module was found with a tuning range of 5.5 to 10 GHz. By quadrupling the output of this module, the roughly 22 to 40 GHz



Figure 2.11: Schematic of the High Frequency Feed showing the low noise amplifiers and including the calibration system. The Low Frequency Feed has the same configuration, except that the feed itself is uncooled.

tuning range needed for the receiver was achieved. A pair of high performance active doublers were used to quadruple the signal.

Deciding on the baseband bandwidth required many considerations. At the time, 1 GHz bandwidth digitizers were readily available, so the output is low pass filtered to 1 GHz. However, to take full advantage of the wide band front end, four independently tunable, dual polarization receivers are provided, each of which down-converts a 2 GHz band to baseband which is output as an upper and lower sideband or as quadrature signals. The 2 GHz bandwidth is determined by a waveguide IF filter centered at 22 GHz. The second mixer has a fixed LO at 22 GHz and natively produces quadrature signals, which can be optionally combined in an analog hybrid to produce the upper and lower sideband outputs. In addition, a second IF filter with a passband of 22.1 to 22.5 GHz can be used to eliminate images from the other sideband at frequencies where RFI is a problem. Using the analog hybrid reduces the processed band to 100–1000 MHz and the sideband rejection is limited to about 25 dB. On the other hand, the quadrature signals can be directly digitized and combined to access 5–1000 MHz with better than 35 dB image rejection as described in section 3.3.

While a maximum of 1000 MHz is available from each output, for many observations such wide bandwidths are not necessary and can even be detrimental by allowing RFI to corrupt the measurement. As such, filters are also available covering 0–500 MHz, 500–1000 MHz, and 270–370



Figure 2.12: The polarization selection portion of the receiver. This circuit provides an optional hybrid to convert to circular polarization and a transfer switch for rapidly swapping the polarizations. The receiver contains two such circuits, one for the HFF and one for the LFF, which are identical in design but contain components appropriate for their respective frequency ranges.



Figure 2.13: The up-down converter at the heart of the receiver system. The radiometer contains four copies of this circuit, each capable of converting a 2 GHz band to baseband. Notice that the quadrature mixer is driven by a fixed local oscillator, thus the quadrature imbalance should not depend on the receiver operating frequency. The 5.5 to 11 GHz synthesizer is locked to the 100 MHz station reference, as is the common 11 GHz oscillator.



Figure 2.14: The baseband processor circuit provides gain adjustment and several selectable filters for defining the final bandwidth for various applications. A selectable hybrid is also available to convert the native quadrature outputs to upper and lower sidebands if desired.

MHz. The 270–370 MHz filter is designed to be compatible with legacy instruments, such as the Wide-band VLBI Science Recorder (WVSR).

Unfortunately, the photonic modules necessary to bring the baseband signals down to the back end via optical fibers ended up being a major cost limitation. At the same time, this aspect of the system can be easily upgraded in the future to improve the capabilities of the system. As built, 8 fiber optic links are available, which allow 8 of the 16 baseband signals to be connected to the desired back end. A matrix switch is available in the receiver to allow any one of the 8 fibers to be connected to any one of the 16 baseband signals.

The receiver also provides the option of switching in a hybrid to change the polarization basis from linear to circular. In addition, a transfer switch allows rapidly swapping the two polarization inputs to perform switched polarimetry, a technique for sensitive polarimetry measurements.

Finally, the receiver includes noise calibration diodes and a comb generator for phase calibration. To make the system as flexible as possible, two noise levels are available for each polarization input. All of these signal sources are combined and coupled into the LNA input. The architecture of the DSS-28 receiver provides a unique advantage for digital quadrature imbalance correction in that aside from imbalance in the RF splitter driving the quadrature down-conversion mixer, everything is static with respect to frequency because the tuning occurs at the up-conversion mixer. Therefore the system should not require frequent imbalance calibration provided it is stable.

2.4 The Digital Signal Processing System

While the GAVRT receiver was designed to be compatible with existing baseband spectrometers and signal recorders, including the Wide-band VLBI Science Recorder, none of the available systems would take full advantage of the enormous instantaneous bandwidth provided by the receiver. Therefore a new system was desired which would provide versatility and enable unique science to be performed with this instrument.

The primary goals of the digital signal processing subsystem were to make use of as much of the bandwidth available while meeting the tight cost requirements of the project. Many science campaigns were foreseen as desirable applications of the new radio telescope, each with different signal processing requirements. Therefore, the signal processor was required to be highly configurable. To best meet these goals, it was decided to implement the signal processor using field programmable gate arrays (FPGAs), as is becoming increasingly common in radio astronomy. The enormous capability and versatility of FPGAs comes at the cost of engineering complexity, which in the past has required such large design efforts that any given system is obsolete by the time it is complete. To help alleviate this problem, the Center for Astronomical Signal Processing Equipment Research (CASPER)³ group at the University of California, Berkeley, has built up a collection of open-source hardware and 'gateware' (firmware for FPGAs) specifically oriented towards radio astronomy applications (Parsons et al., 2006). By standardizing on a set of boards, software tools, and libraries, they have effectively distributed the large non-recoverable engineering costs associated with complex FPGA-based signal processors.

Designing the GAVRT signal processor using the CASPER tools provided a far more versatile system at a much lower cost than would have been possible using closed-source, commercial solutions. Additionally, CASPER hardware has rapidly proliferated to observatories around the world which provides the benefit that signal processor designs developed at GAVRT can be used on other telescopes around the world, and similarly, designs originally made for other telescopes can be used at GAVRT. Inevitably, the digital hardware around which the GAVRT signal processor was built has already been superseded by the next generation of FPGA technology. Fortunately, new boards are already available from the CASPER group and it will be a relatively simple matter to port the designs described herein to the new hardware.

2.4.1 Detailed architecture description

The GAVRT Radio Astronomy Signal Processor (GRASP) is composed of digitizers, two types of interconnected FPGA boards, and a small cluster of PCs. The overall architecture is shown in figure 2.15

2.4.1.1 Digitizers

The eight baseband outputs from the matrix switch are sent down to the pedestal room over fiber optic links. The output of each fiber optic receiver is connected to a CASPER digitizer board, called the iADC. This board contains an E2V (formerly Atmel) AT84AD001B ADC. Internally, this chip

³http://casper.berkeley.edu/



Figure 2.15: Overall block diagram of the GAVRT Radio Astronomy Signal Processor. There are also 100 Mb Ethernet connections to all of the iBOB and BEE2 boards for control and low rate data transfer which are not shown for clarity.

contains two ADCs which can sample at up to 1 GSps each. The ADCs can also be interleaved to provide an effective sampling rate of 2 GSps and be capable of digitizing a 1 GHz bandwidth signal.

2.4.1.2 The iBOB FPGA Interface Boards

Each digitizer is connected to a board containing a single Xilinx Virtex II Pro FPGA, called an iBOB.⁴ The iBOB has sufficient resources to independently process the baseband signal. For example, for continuum observations, an iBOB can be configured with a polyphase filterbank spectrometer capable of dividing the band into 1024 channels. The low spectral leakage provided by the polyphase filterbank allows channels contaminated with RFI to be eliminated. The mean of the remaining uncorrupted channels then provides an equivalent total power output free of RFI. As another example, the iBOB can digitally down-convert a 64 MHz band and then use a polyphase filterbank to break this band into 8192 channels for an effective spectral resolution of 8 kHz. This is sufficient resolution for some spectral line studies. In particular, the Galactic hydrogen observation presented in figure 1.2 was made using this spectrometer. One of the chief limitations in using an iBOB independently is that the maximum practical data rate through the 100 Mbit Ethernet interface is only about 7 Mbit/s, which is insufficient for most pulsar observations.

Fortunately, the iBOB provides two 10 Gbit/s CX-4 XAUI ports which provide ample I/O bandwidth. Each CX-4 port can be used as a point-to-point XAUI link, or as a standard 10 Gbit/s Ethernet (10 GbE) link. In GRASP, all of the iBOB CX-4 ports are used as point-to-point XAUI links because the FPGA resource requirements are much smaller than for 10 GbE. However, this could be upgraded in the future, which would provide dynamic connections between all of the FPGAs in the system through a commercially available 10 GbE network switch, enabling a multitude of packetized signal processing architectures described by Parsons et al. (2008).

2.4.1.3 The BEE2 FPGA Processing Board

To provide more FPGA resources, the CASPER group has also developed the Berkeley Emulation Engine 2 (BEE2) board, which contains five interconnected Virtex II Pro FPGAs (Chang et al., 2005). The five FPGAs are arranged as four 'User' FPGAs and one 'Control' FPGA. The Control FPGA is able to reprogram each of the User FPGAs. The User FPGAs are connected in a ring and to the Control FPGA via digital I/O lines. One unique feature of the BEE2 compared to many commercial products which provide large numbers of interconnected FPGAs for computational acceleration is the I/O bandwidth provided. The BEE2 provides 18 CX-4 XAUI ports for 180 Gbit/s bidirectional I/O. These are allocated as 4 XAUI ports per User FPGA and 2 for the Control FPGA. Finally, each FPGA is connected to 4 DDR2 DRAM memory modules, providing up to a total of

 $^{^{4}}$ iBOB stands for Internet Break-Out Board, so named because it provides two 10 Gbps XAUI ports, and was originally designed to simply transfer digitized data to other FPGA boards for processing.

20 Gbytes of DRAM.

In the GRASP design, each User FPGA is connected to two iBOBs via its 4 XAUI links. The two XAUI links from the Control FPGA are configured as 10 GbE ports and connected to a commercial Ethernet switch which provides 24 1 GbE ports to distribute the data flowing from the BEE2 to a cluster of computers.

To simplify writing control and interface software, the BEE2 uses another technology developed by the CASPER group called BORPH⁵ (So and Brodersen, 2007). BORPH is a modified version of the Linux operating system which runs on the PowerPC embedded in the Control FPGA. BORPH provides tools to convert the FPGA programming bitfile to a special executable file which can be run in the same way as any other Linux program. When this special executable file is run, the Control FPGA configures the specified User FPGA with the bitfile, and then maps all of the shared registers and memory spaces from the User FPGA design into files under the Linux file system. This allows the designer to treat the registers and memory spaces as any other file, and thus can program the FPGA control and interface software in any programming language and without having to worry about the low-level details of interfacing with these shared memory regions.

2.4.1.4 Clocking and synchronization

Many scientific goals and instrument designs require careful synchronization between the digitizers and accurate timestamps. Additionally, high speed digitization places stringent requirements on the sampling clock.

The DSS-28 system uses standard Deep Space Network Time Code Translator (TCT) modules as reference clocks and oscillators. The TCT is connected by fiber to the station maser from which it produces synchronous 100 MHz and 1 pulse per second reference signals. The local oscillators in the GAVRT wide bandwidth Radiometer will all be locked to this 100 MHz reference. For the prototype system used for this thesis, however, they were locked to a free running 100 MHz oscillator.

The ADC sampling clocks will also be locked to the 100 MHz reference. The sampling clocks are generated by Analog Devices AD9517 synthesizers which include an on chip integrated VCO, PLL, and dividers. They are capable of producing clock rates from 10 MHz to 1100 MHz with a gap between 750 and 880 MHz. The clock signals have excellent phase noise characteristics. This is important for high speed samplers because the sampling process is equivalent to multiplying by the sampling clock in the time domain which results in convolution of the spectrum of the sampled signal with the phase noise spectrum of the sampling clock. They are designed to drive positive emitter coupled logic (PECL) but are quite effective at driving the 50 ohm ADC clock input.

For synchronization, each ADC contains a 50 ohm terminated input connected to a logic buffer designed to be connected to a PPS signal. The PPS signal is sampled at the ADC clock rate and 4

⁵BORPH stands for Berkeley Operating System for Re-Programmable Hardware.

successive samples are available at the FPGA per FPGA clock. This should allow synchronization to within ± 1 ADC sample. When the iBOB is powered up or otherwise reconfigured with a new FPGA design, the internal PLL on the FPGA is locked to the ADC output data clock. This synchronization can cause an ambiguity of up to 4 ADC clocks. However, experience at the Allan Telescope Array has shown this to be very stable, and can thus be calibrated out as a fixed delay (M. Dexter, private communication).

2.5 Testing Novel Wide Bandwidth Technologies at DSS-28

Several unique features of the DSS-28 system make it an excellent tool for developing technologies for future wide bandwidth radio telescopes. The rotatable tertiary allows new feeds to be installed with relative ease. There is currently extensive research being done on cryogenically cooled decade bandwidth (covering 1 to 10 GHz) feed candidates for the Square Kilometer Array, and it is possible one of these feeds may eventually be used at DSS-28.

The receiver system also provides a port to tap into the full bandwidth output from the feed and LNA. This could be used to test wide bandwidth (10 GHz or more) optical links which could transport the entire RF band of the LNA to the pedestal room. This would allow the receiver electronics to be relocated to a motionless, climate-controlled environment to improve stability. Alternatively, a large portion of the RF band could potentially be digitized in one swath using the multi-gigahertz bandwidth ADCs which have become available recently. Dynamic range would likely be the major limiting factor in either of these cases, due to the RFI environment.

Finally, the receiver and signal processing systems have been designed for commensal observations. New signal processing algorithms can be tested using the existing hardware in parallel with standard observations. Additionally, the versatility of the four independent receivers and the matrix switch makes it easy to connect new signal processing hardware to the system.

Chapter 3

Wide Bandwidth Spectral Line and Continuum Observations in the Presence of RFI

3.1 Radio Frequency Interference

The explosion in wireless technology in recent years has resulted in increasingly troublesome interference for radio astronomy. At many frequencies, RFI limits the practical usable bandwidth of a radio telescope. As with any radio telescope, DSS-28 is affected by the local interference environment, but fortunately its versatility provides various means of observing in the presence of many types of interference.

3.1.1 The RFI situation at DSS-28

Before the decision to equip DSS-28 with a wide bandwidth radiometer was finalized, an RFI survey was carried out to determine if any sources of RFI were strong enough to saturate or even damage the broadband LNAs. The survey was conducted with an omni-directional broad band antenna connected to a preamplifier and a spectrum analyzer. Laboratory measurements indicate that the gain of the LNAs in the HFF compresses by 1 dB when the input power is approximately -60 dBm. The far-out sidelobes of a parabolic reflector can be roughly approximated as having the same gain as an isotropic antenna. Therefore any signals above this level would be cause for concern. Another important level to compare against is the the system noise. Across much of the HFF range, the system noise temperature is below 40 K. This corresponds to -122 dBm/MHz or -92 dBm/GHz at the LNA input. Many radio astronomy observations require measuring changes in the system noise to the 0.01% level. Therefore, even though no signals were detected in the survey above the -60 dBm level that would cause saturation in the LNA, signals are often present that exceed the total noise power in a 1 GHz band (-92 dBm).

3.1.2 In-situ RFI Surveys

One benefit of the wide instantaneous bandwidth available at DSS-28 is it provides a superb instrument to survey the RFI environment at the site. Measurements were typically made by stepping through adjacent 500 MHz bands, recording spectra with an integration time of 40 ms and dwelling for 10 to 100 seconds at each step. In this way, the entire spectrum is repeatedly scanned to characterize the various sources of interference. An example of such a measurement are shown in figure 3.1. With a spectrometer channel width of 1 MHz, scanning from 1 to 14 GHz yields 13000 frequency bins which quickly becomes an overwhelming amount of data. To present the data meaningfully, the standard deviation divided by the mean of each channel is computed for each dwell time and compared to the theoretical value expected from the radiometer equation. The fraction of measurements exceeding some threshold can then be computed to estimate the spectral occupancy and the usability of each channel. The spectral occupancy computed from the dataset shown in figure 3.1 is shown in figure 3.2. The most notable feature in these plots is the broadband interference which is caused by a nearby transmitter as discussed in the next section.

3.1.3 DSS-27

One source of interference which has a major impact on DSS-28, which was not noticed until after the radiometer was installed, is the DSS-27 station 200 meters from DSS-28. DSS-27 is used for uplink communications to near-earth satellites with a 200 watt S-band transmitter. The DSS-27 transmitter is scheduled for operation approximately 12 hours per day and has an enormous effect on the DSS-28 receiver. The effect is very clear in the RFI survey conducted using the DSS-28 receiver and digital spectrometer shown in figures 3.1 and 3.2 as 'splatter' across the whole band when DSS-27 is transmitting. The splatter is the result of the LNA being driven into a non-linear regime and acting as a mixer with the DSS-27 carrier acting as the local oscillator. Other strong sources of RFI are thus smeared across the band with concentrations around harmonics of the DSS-27 carrier. The strong cell phone bands around 850 and 1900 MHz are amongst the signals that mix with the DSS-27 carrier. As expected, the level of RFI splatter varies strongly as a function of time as sidelobes of DSS-27 move as it tracks a spacecraft across the sky.

The different response in polarization X versus polarization Y is consistently observed, and can be attributed to the frequency response of the two LNAs. For the prototype system, different LNAs were installed in each polarization. Polarization X optimized for 4–12 GHz while polarization Y is optimized for 0.5–11 GHz. The 4–12 GHz amplifier has thus been selected for the final system.

While some observations are possible with the HFF while DSS-27 is transmitting, the extreme variability of the gain and system temperature makes observations unpleasant. In addition, the LFF is completely unusable while DSS-27 is transmitting. Because of this, filters will be installed in the



(b) Polarization Y

2000

Figure 3.1: Examples of RFI scans performed at DSS-28. The color scale is such that dark blue regions are within a factor of 1.5 of the RMS fluctuation predicted by the radiometer equation, while red regions suffer from severe interference. The right subplot indicates when the DSS-27 carrier was detected and the approximate frequency. Approximately half of the full 15 hour data set is shown.



Figure 3.2: Spectral occupancy computed from RFI scans presented in figure 3.1. Each channel was marked as occupied if the RMS exceeded 5% in the one minute dwell time. This represents 10 times the RMS fluctuation predicted by the radiometer equation for a 1 MHz bandwidth and a 40 ms integration time. Notice the main effect of DSS-27 on polarization X is at harmonics of the \sim 2 GHz carrier, while polarization Y shows a wider array of intermodulation products between the carrier and other low frequency interference.

final system. High pass filters with a cut-off of around 2.7 GHz will be used for the HFF, while 1.7 GHz low pass filters will be used for the LFF. These filters reduce the total bandwidth of the system, but remove the troublesome cell phone bands around 1800–1900 MHz in addition to the DSS-27 carrier from both feeds. The filters will be installed in the cryogenic packages to minimize the impact on system temperature.

3.2 Stability Measurements

As we have seen, the stability of a receiver is of paramount importance for sensitive continuum measurements. The stability of a radiometer is traditionally evaluated using measurements of the Allan Variance (Allan, 1987). This provides an estimate of the variance expected when differencing two consecutive integrations of length L. Ideally, one could make arbitrarily sensitive measurements with longer integration times, but in practice, fluctuations in the gain and noise temperature of the system, which generally follow a 1/f spectral distribution, dominate after some integration time. This integration time is referred to as the Allan time of the system, and represents the maximum useful continuous integration time. A detailed treatment of Allan variance computation for spectrometric characterization has been given by Ossenkopf (2008). As described therein, the formal definition of Allan variance involves a computationally wasteful convolution. In practice, the Allan variance can be approximated much more efficiently at each time lag L by first breaking up a long series of data into segments of l samples where $L = l\Delta t$ and Δt is the sampling time. The time series is first normalized by computing

$$s(t_k) = \frac{c(t_k) - c_{zero}}{\langle c(t_k) - c_{zero} \rangle_k}$$
(3.1)

where $c(t_k)$ is the output of a spectrometer channel at the kth time step and c_{zero} is the value out of the spectrometer channel when no input is present. We then compute the average of each segment to form the sequence S(K), written mathematically as

$$S(K) = \frac{1}{l} \sum_{k=Kl+k_1}^{(K+1)l+k_0-1} s(t_k).$$
(3.2)

Here k_0 is an arbitrary starting offset in the time series. As discussed by Ossenkopf (2008), different choices of k_0 will lead to somewhat different Allan variance results, but they generally fall within the error bars of the full convolution method. We then difference consecutive elements to obtain

$$Z(K) = S(K) - S(K+1)$$
(3.3)



Figure 3.3: Plots of total power deviations from the spectrometer (blue) and from the square-law detector (green) along with the deviation of the angle of the tertiary reflector from its nominal position (red)

the variance of which yields the Allan variance at time lag L:

$$\sigma_A^2(L) = \left\langle (Z(K) - \left\langle Z(K) \right\rangle_K)^2 \right\rangle_K \,. \tag{3.4}$$

On windy days, the stability of the DSS-28 system is currently limited by wind buffeting the tertiary. Changes in the tertiary position cause changes in the system temperature, as seen in section 2.1.3. To investigate sources of instability at DSS-28, measurements were taken with a 512 MHz, 512 channel spectrometer with the antenna at zenith. The data spans both a windy and a calm period so sections from both periods were analyzed separately to examine the effect of wind on the stability. During the windy period, deflections of the tertiary from its nominal position of up to $\pm 2^{\circ}$ were recorded. During the calm period, the tertiary was essentially stationary to the resolution of the angular encoder, approximately 5 millidegrees. The voltage readings from the square-law power detectors located in the baseband processor modules in the receiver box were also recorded to compare the total power measured directly at the receiver box with the spectral data measured after traversing the fiber optic link. The integration time of the square-law detectors was 0.1 seconds while the spectrometer integration time was 40 ms. A portion of the time series of these signals is shown in figure 3.3. In both windy and calm conditions, the total power measured by the square law detector and by the spectrometer track each other very well. There is clear correlation between the tertiary position and the total power measurements.

The waterfall plots in figure 3.4 show the power deviation from the average power in each spectrometer channel versus time. The residual non-Gaussian fluctuations in the calm data set are seen to effect the whole band. The lower frequencies show somewhat more variation than the higher frequencies. The exact cause of these instabilities is presently unknown. It's possible that strong



Figure 3.4: Waterfall plot of spectrometer output. For a perfectly stable system, these plots would look completely featureless. The variations in the system during the calm period apparently affect the whole band, though the effect is stronger at the low frequency end. Some minor RFI is apparent as vertical features.

out-of-band RFI, such as from cell phone base stations, is modulating the gain of the LNA. The 1/f gain fluctuations intrinsic to HEMT amplifiers is another possible source.

The wind related instability is also seen clearly by comparing the Allan variance plots in figures 3.5 and 3.6. The calm data shows a drift component that interestingly begins to integrate down again on time scales greater than about 20 seconds. On the other hand, the windy data shows a strong humped form peaking around 2–3 seconds. This is related to the response time of the tertiary servo. This effect can also be seen by computing the power spectral density of these data, shown in figure 3.7. In these plots, the expected white noise floor is seen as the data flattens out to the right. The calm data shows the 1/f component towards lower frequencies, while the windy data again shows a humped feature peaking around 0.1 Hz. The power spectral density of the tertiary angle measurements is shown as well, and clearly identifies the tertiary motion as the source of the instability.

3.3 Demonstration of a Spectrometer with Real-time I/Q Imbalance Correction

While sideband separating receivers are commonly used in millimeter and sub-millimeter wavelength radio astronomy, they are rarely encountered at centimeter wavelengths. The reasons for this are that traditionally the advantages offered—wider bandwidths and simpler receiver architectures—are outweighed by inevitability of interfering signals falling within a given band, which can be more easily rejected with a good image reject IF filter. In the past, achieving image rejection of 30 dB



Figure 3.5: Allan Variance measured under calm conditions. The legend indicates the center frequency (CF) and equivalent bandwidth (BW) of the spectrometer data used to make each curve.



Figure 3.6: Allan Variance measured under windy conditions. The legend indicates the center frequency (CF) and equivalent bandwidth (BW) of the spectrometer data used to make each curve.



Figure 3.7: Power spectral density plots of the total power computed using different numbers of channels of the spectrometer versus the square-law detector and the tertiary angular position error. The center frequency (CF) and equivalent bandwidth (BW) of the individual spectrometer curves is indicated in the legend. The power spectral density estimates were computed with a Hamming window and a linear baseline was removed from each segment of the time series data. The square-law detector and tertiary angular position error were sampled more quickly than their underlying sample rate, leading to the apparent high frequency roll-off.

or more with a quadrature mixer has required extremely careful design, as even 1 degree of phase or 0.1 dB of amplitude imbalance will result in degraded image rejection. This is a problem in centimeter wave radio astronomy because a strong interfering signal in one sideband will only be rejected by ~ 30 dB in the other sideband, thus corrupting twice as many spectrometer channels. In a standard sideband separating receiver, the incoming RF signal is split between two mixers driven by quadrature phases of a local oscillator. The baseband quadrature signals from the mixers are then low-pass filtered to define the RF and baseband bandwidths, then combined again in a quadrature hybrid to recover the upper and lower sideband signals.

The quadrature phases of the LO are generally obtained using an analog quadrature hybrid, which generally limits the tuning range of the local oscillator to octave, or at best decade bandwidths. Recently, integrated designs using digital flip-flops have been demonstrated yielding exceptional performance over tuning ranges exceeding 40 to 1 (Bardin and Weinreb, 2008). This problem is avoided in the DSS-28 receiver by employing an up-down conversion heterodyne architecture, which uses a fixed second LO for the quadrature down-converter. The bandwidth of the final sideband separation hybrid, however, is specified for only 100 MHz to 1000 MHz. In practice it shows reasonable performance down to about 60 MHz, but this still leaves a ~ 120 MHz gap in the center of the band. On the other hand, when the hybrid is bypassed, the baseband signals show good response down to just a few MHz. Careful matching in the DSS-28 receiver yields a measured sideband rejection of ~ 20 dB. In general, the analog hybrid can be eliminated in systems employing digital spectrometers, as the operation of the hybrid can be easily synthesized mathematically. In practice, the baseband quadrature signals travel down fiber optic links which were not carefully length matched and exhibit path length differences greater than 2 cm, corresponding to $\sim 4^{\circ}$ at 1 GHz, before being digitized. Additional imbalance is present as well, as shown below. Therefore, we must employ some compensation scheme in the spectrometer to make use of the quadrature signals from the DSS-28 receiver.

To date, most research on quadrature imbalance correction algorithms has focused on narrow band communications applications requiring correction at only a single frequency. This is generally accomplished by measuring the amplitude and phase error and then using essentially the Gram-Schmidt orthogonalization process to restore the quadrature relationship between the signals (Churchill et al., 1981). Ideally, the phasors representing a pair of signals in quadrature should be orthogonal and equal in amplitude. Orthogonality of the measured phasors can be restored by adding a scaled version of one to the other, where the scale factor is given by the Gram-Schmidt process.

Pun et al. (2000) extended this technique to allow for frequency dependent imbalance by using a polynomial to interpolate between measurements of the imbalance versus frequency, then casting this interpolating polynomial in the form of an FIR correction filter. This method becomes impractical



Figure 3.8: Sources of imbalance in quadrature down-conversion

if the imbalance has a significant ripple to it because many calibration points will be necessary to characterize the ripple, and thus the interpolating polynomial will be of a high order and likely poorly behaved. Additionally, the specific formulation treats the IF frequency response as symmetric about zero, which may not be the case if there is imbalance in the RF path.

In radio astronomy observations, the end goal is generally to estimate the power spectral density versus frequency in the band of interest. This application lends itself to a unique method of quadrature imbalance correction described by Fisher and Morgan (2008), and developed contemporaneously by the author (see also Morgan and Fisher, 2009).

To understand how we can digitally correct for imbalances in the receiver hardware, we consider the model of a general quadrature receiver shown in figure 3.8. In this receiver, the incoming RF signal is split with a power divider and fed into two mixers. The imbalance of the power divider is modeled by the transfer function H_{RF} . The local oscillator signal is also split to drive the two mixers. Ideally the two resulting LO signals should be 90° out of phase, but for generality we model the power division with an arbitrary transfer function H_{LO} . The output of the mixers is then low pass filtered to retain only the difference products and to reduce leakage of the LO into the baseband signals. Here we model the low pass filters as ideal and lump any imbalance after the mixers into H_{BB} . Notice that all of the imbalance can be lumped into the 'Q' signal path without loss of generality because it is only the ratio of the transfer functions in the 'I' and 'Q' paths that is important to imbalance correction, as we shall see.

To analyze this system, we consider a CW RF signal at the input at frequency ω_{RF} with arbitrary amplitude A

$$V_{RF}(t) = A \left[e^{j\omega_{RF}t} + e^{-j\omega_{RF}t} \right] .$$
(3.5)

Before the mixers, the RF signals will be

$$V_{RF,I}(t) = \frac{V_{RF}(t)}{\sqrt{2}} \tag{3.6}$$

$$V_{RF,Q}(t) = \frac{A}{\sqrt{2}} \left[|H_{RF}(\omega_{RF})| e^{j\omega_{RF}t + \phi_{RF}(\omega_{RF})} + |H_{RF}(-\omega_{RF})| e^{-j\omega_{RF}t + \phi_{RF}(-\omega_{RF})} \right]$$
$$= \frac{A|H_{RF}(\omega_{RF})|}{\sqrt{2}} \left[e^{j\omega_{RF}t + \phi_{RF}(\omega_{RF})} + e^{-j\omega_{RF}t - \phi_{RF}(\omega_{RF})} \right], \qquad (3.7)$$

where we have used the fact that H_{RF} is a real transfer function and thus $|H_{RF}(-\omega)|e^{j\phi_{RF}(-\omega)} = |H_{RF}(\omega)|e^{-j\phi_{RF}(\omega)}$. The local oscillator signals will be

$$V_{LO,I}(t) = e^{j\omega_{LO}t} + e^{-j\omega_{LO}t}$$

$$(3.8)$$

$$V_{LO,Q}(t) = |H_{LO}(\omega_{LO})|e^{j\omega_{LO}t + \phi_{LO}(\omega_{LO})} + |H_{LO}(-\omega_{LO})|e^{j\omega_{LO}t + \phi_{LO}(-\omega_{LO})}$$
$$= |H_{LO}(\omega_{LO})| \left[e^{j\omega_{LO}t + \phi_{LO}(\omega_{LO})} + e^{-j\omega_{LO}t - \phi_{LO}(\omega_{LO})}\right].$$
(3.9)

Thus, after mixing and low pass filtering, the baseband quadrature signals will be given by

$$V_I(t) = A \left[e^{j(\omega_{BB}t)} + e^{-j(\omega_{BB}t)} \right]$$
(3.10)

$$V_{Q}(t) = A |H_{RF}(\omega_{RF})| |H_{LO}(\omega_{LO})| |H_{BB}(\omega_{BB})| \left[e^{j(\omega_{BB}t - \phi_{LO}(\omega_{LO}) + \phi_{RF}(\omega_{RF}) + \phi_{BB}(\omega_{BB}))} + e^{j(-\omega_{BB}t + \phi_{LO}(\omega_{LO}) - \phi_{RF}(\omega_{RF}) - \phi_{BB}(\omega_{BB}))} \right],$$
(3.11)

where we have defined $\omega_{BB} \equiv \omega_{RF} - \omega_{LO}$. To reconstruct the original signal in a single sideband we consider the sum

$$V_{I}(t) + \Gamma V_{Q}(t) = A \left[e^{j\omega_{BB}t} + e^{-j\omega_{BB}t} + \Gamma |H_{RF}(\omega_{RF})| |H_{LO}(\omega_{LO})| |H_{BB}(\omega_{BB})| \left(e^{j(\omega_{BB}t - \phi_{LO}(\omega_{LO}) + \phi_{RF}(\omega_{RF}) + \phi_{BB}(\omega_{BB}))} + e^{j(-\omega_{BB}t + \phi_{LO}(\omega_{LO}) - \phi_{RF}(\omega_{RF}) - \phi_{BB}(\omega_{BB}))} \right) \right],$$

$$(3.12)$$

where Γ is a complex number determined as follows. For a fixed ω_{LO} we can collect the terms of

the transfer functions to obtain

$$V_{I}(t) + \Gamma V_{Q}(t) = A \left[e^{j\omega_{BB}t} + \Gamma | H(\omega_{BB}) | e^{j(\omega_{BB}t + \phi(\omega_{BB}))} + e^{-j\omega_{BB}t} + \Gamma | H(-\omega_{BB}) | e^{-j(\omega_{BB}t - \phi(\omega_{BB}))} \right]$$

$$(3.13)$$

where we define the complex transfer function $H(\omega_{BB})$ to represent the combined effects of H_{RF} , H_{LO} , and H_{BB} . Thus,

$$|H(\omega_{BB})| = |H_{RF}(\omega_{BB} + \omega_{LO})||H_{LO}(\omega_{LO})||H_{BB}(\omega_{BB})|$$
(3.14)

$$\phi(\omega_{BB}) = \phi_{RF}(\omega_{BB} + \omega_{LO}) - \phi_{LO}(\omega_{LO}) + \phi_{BB}(\omega_{BB}).$$
(3.15)

Now it is clear that if we set

$$\Gamma = \frac{1}{|H(-\omega_{BB})|} e^{j(\phi(\omega_{BB}) + \pi)}, \qquad (3.16)$$

the terms involving $-\omega_{BB}$ cancel leaving

$$V_{I}(t) + \Gamma V_{Q}(t) = A \left[e^{j\omega_{BB}t} + \frac{|H(\omega_{BB})|}{|H(-\omega_{BB})|} e^{j(\omega_{BB}t + 2\phi(\omega_{BB}) + \pi)} \right].$$
 (3.17)

The choice of sign in defining Γ is arbitrary; the opposite sign will simply swap the positive and negative baseband frequencies of the spectrum. Note also that if we had not lumped all of the transfer functions into the 'Q' signal path, any transfer function in the 'I' path would be absorbed in the choice of Γ . In an ideal system, $|H(\omega)| = 1$ and $\phi(\omega_{BB}) = -\pi/2$ which implies $\Gamma = e^{j\pi/2} = j$, and we obtain simply

$$V_I(t) + \Gamma V_Q(t) = 2Ae^{j\omega_{BB}t}.$$
(3.18)

Thus we see that for a single frequency we can always perfectly reject the image tone at $-\omega_{BB}$, but we potentially lose efficiency since the two desired components in equation 3.17 will not always add exactly in phase.

To calculate the power spectral density estimate, the spectrum is divided up into channels with a digital filter bank, and then the average power is computed in each channel. For a quadrature input spectrometer, the power spectrum estimate can be expressed mathematically as

$$\hat{\vec{P}_k} = \left\langle \left| \mathcal{FFT} \left(\vec{i_k} + j\vec{q_k} \right) \right|^2 \right\rangle_N, \qquad (3.19)$$

where $\vec{P_k}$ is the a vector of power spectral density estimates for the *k*th spectral channels of the filterbank, formed by integrating N spectra. $\vec{i_k}$ and $\vec{q_k}$ are vectors of k consecutive samples of the in-phase and quadrature baseband signals. We have used the FFT here as a filterbank, but the same

idea holds for any linear filterbank operation, such as a polyphase filterbank. Notice that because the FFT is a linear operation, we can trade a single complex FFT for two real FFTs as

$$\hat{\vec{P}_{k}} = \left\langle \left| \mathcal{FFT}\left(\vec{i_{k}}\right) + j\mathcal{FFT}\left(\vec{q_{k}}\right) \right|^{2} \right\rangle_{N} = \left\langle \left| \vec{V_{I,k}} + j\vec{V_{Q,k}} \right|^{2} \right\rangle_{N}.$$
(3.20)

Here each element in $V_{I,k}$ and $V_{Q,k}$ represents a sample from the narrow band quadrature signals in each of the k spectrometer channels. These signals are thus analogous to the quadrature time series discussed above, $V_I(t)$ and $V_Q(t)$. In general we can represent the quadrature imbalance by identifying

$$V_{Q,k}^{\vec{}} = \vec{H_k} V_{Q,k}^{\vec{}} \,, \tag{3.21}$$

where $V_{Q,k}^{\vec{r}}$ represents the uncorrupted signal and each element of $\vec{H_k}$ is the transfer function evaluated at each of the k center frequencies of the filterbank. Therefore, if we find the appropriate value of Γ for each frequency, we can arrange them in a vector $\vec{\Gamma_k}$ and compute the corrected spectrum as

$$\hat{\vec{P}}_{k}' = \left\langle \left| \vec{V}_{I,k} + \vec{\Gamma}_{k}^{T} \vec{V}_{Q,k} \right|^{2} \right\rangle_{N} .$$
(3.22)

It is interesting to note that this is similar to applying an FIR filter representing the inverse of H in the frequency domain using the FFT. This leads to the temptation to compute the inverse FFT of $\vec{\Gamma_k}^T V_{Q,k}$ to recover the original uncorrupted timeseries. However, doing so will not yield the desired result because this sequence of operations actually results in the cyclic convolution of h(t) and q(t), as described by Brigham (1988). Computing the true convolution of a timeseries using the FFT requires the use of the overlap and add or overlap and save algorithms.

The problem of determining $\vec{\Gamma}_k$ at each frequency falls in the broad field of system identification. Numerous methods have been proposed to measure quadrature imbalance in narrow bandwidth communications systems, with emphasis on using *a priori* information about the expected communications signals, thereby reducing hardware complexity for mass-market devices. The incremental cost of adding specific calibration hardware to a single dish radio telescope is completely reasonable, though this assumption may have to be reconsidered for large arrays. In anticipation of the need for phase calibration measurements, the DSS-28 receiver includes a 100 MHz comb signal that can be injected into the input of the LNA. The comb signal is locked to the station 100 MHz reference, and is coherent with the sampling clock of the ADCs. As mentioned in section 2.3, the architecture of the DSS-28 receiver has the advantage of a fixed local oscillator for the quadrature down-conversion mixer, thus the H_{LO} and H_{RF} transfer functions should be essentially fixed regardless of receiver configuration and observing frequency. Note that H_{RF} represents the relative response of the power splitter after the IF filter and thus the frequency range of interest is fixed to 21–23 GHz. This should be compared to the simpler direct down-conversion architecture studied by Fisher and Morgan (2008).

While system identification can be carried out using time domain or frequency domain methods, we find frequency domain methods yield more direct insight into this particular problem, and are more directly related to modern FPGA spectrometers which compute the spectrum using polyphase or FFT filterbanks as opposed to autocorrelation spectrometers.

At frequency $f_k = k/T$, the frequency of the kth bin of the filterbank, we measure the signals

$$I(k) = I_0(k) + N_I(k)$$
(3.23)

$$Q(k) = Q_0(k) + N_Q(k), \qquad (3.24)$$

where $I_0(k)$ and $Q_0(k)$ represent the signals of interest while $N_I(k)$ and $N_Q(k)$ represent the noise which is inevitably present. We can then estimate the relative transfer function as

$$H(k) = \frac{I(k)}{Q(k)} = \frac{I_0(k) + N_I(k)}{Q_0(k) + N_Q(k)} = H_0(k) \frac{1 + N_I(k)/I_0(k)}{1 + N_Q(k)/Q_0(k)}.$$
(3.25)

It has been shown that the variance of H(k) does not exist in general because of the chance that $1 + N_Q(k)/Q_0(k)$ is zero (Broersen, 1995). However, for high SNR measurements, we can make a reasonable estimate using the following the method presented by Pintelon and Schoukens (2001). The first order terms of the Taylor expansion yield

$$H(k) \approx H_0(k) \left(1 + \frac{N_I(k)}{I_0(k)}\right) \left(1 - \frac{N_Q(k)}{Q_0(k)}\right)$$
 (3.26)

from which we can recognize the noise in H(k) as

$$N_H(k) = \Gamma_0(k) \left(\frac{N_I(k)}{I_0(k)} - \frac{N_Q(k)}{Q_0(k)} \right) .$$
(3.27)

It is readily apparent that the expectation of N_H is zero, thus the variance is given by

$$\sigma_H^2 = E\left\{ |N_H(k)|^2 \right\} = |H_0(k)|^2 \left[\frac{\sigma_I^2(k)}{|I_0(k)|^2} + \frac{\sigma_Q^2(k)}{|Q_0(k)|^2} - 2\Re\left(\frac{\sigma_{IQ}^2(k)}{I_0(k)\overline{Q_0}(k)}\right) \right].$$
(3.28)

Interestingly, Pintelon and Schoukens (2001) point out that noise injected along with the stimulus signal, which corresponds to the system noise which is dominated by the LNA, has approximately zero effect on the uncertainty of the estimate of H(k). Therefore, the uncertainty is dominated by the noise of components after the quadrature signals are formed. We can thus obtain a worst-case estimate by assuming $\sigma_{IQ}^2(k) = 0$. Notice that to obtain the best estimate of H(k) we must maximize the signal to noise ratio of the measured I(k) and Q(k) signals as expected. Since the stimulus signal is coherent with the sampling rate, if we use a longer sequence of values to compute the FFT, we

effectively reduce the noise bandwidth in the channel that contains the CW tone, and thus the noise will reduce linearly with increasing data set length. This would not be possible if the stimulus were not locked to the sampling rate, as any relative frequency drift would smear the tone across multiple frequency bins.

The comb signal can then be used to determine H(k) and thus $\vec{\Gamma_k}$ by stepping the first LO over a range of 100 MHz. At each step the relative phase and frequency of the components of the comb that fall within the passband can be measured to derive $\vec{\Gamma_k}$. The smallest step size of the synthesizer is 0.1 MHz which corresponds to a step size of the first LO of 0.4 MHz. In the present system, a short segment of raw digitizer data of length L is captured at each frequency step and analyzed off-line. First, the FFT of each segment is computed. The length of the FFT is chosen such that the comb components all fall exactly in the center of an FFT bin. For example, at a sampling rate of 1024 megasamples per second, if the frequencies of the comb components f_i satisfy

$$f_i \equiv 0 \mod 1 \text{ MHz} \tag{3.29}$$

then they will fall in the center of an FFT bin for FFT length 2^k for $k \ge 10$. For each bin m that contains a comb component, the corresponding component of $\vec{\Gamma_k}$ is determined by

$$\vec{\Gamma_k} \left[\frac{mk}{L} \right] = \frac{I(m)\overline{Q(m)}}{|Q(m)|^2} \,. \tag{3.30}$$

Since the step size of the local oscillator is somewhat large, not all components of $\vec{\Gamma_k}$ can be measured directly. Therefore, real and imaginary parts of the remaining components are independently interpolated from the measured components using a cubic splines.

3.3.1 Hardware Implementation and Measurements

To demonstrate the effectiveness of this technique, a simple spectrometer targeting a single iBOB and iADC with a total bandwidth of 1024 MHz was designed. A functional block diagram of the spectrometer is shown in figure 3.9. The in-phase and quadrature signals are sampled at 1024 megasamples per second by the two inputs of the iADC. The FPGA thus processes four samples from each input at a clock rate of 256 MHz. Each input is fed into a real input, 4-tap polyphase filterbank with 512 channels. The resulting spectra emerge from the polyphase filterbank as complex samples from two consecutive (even and odd) frequency bins per clock. Thus one complete 512 channel spectrum is produced every 256 clock cycles. The samples from the 'Q' spectrum are fed into two sets of complex multipliers, one set for each sideband. Each set of multipliers contains two complex multipliers, one for each of the even and odd frequencies bins. A runtime programmable RAM block supplies the necessary coefficients to the multipliers. The address counter for the coefficient



Figure 3.9: Block diagram of the quadrature correcting spectrometer

memory is synchronized with the start of a spectrum so that the appropriate coefficient arrives at the multiplier along with its respective frequency bin. The corrected outputs from the 'Q' spectrum are then added with the respective outputs from the 'I' spectrum to form the sideband separated spectra, which are then squared and accumulated as in a standard spectrometer. The spectrometer also includes a memory buffer for capturing up to 131,072 eight-bit ADC samples from each input for calibration purposes. Notice that this design only requires four additional complex multipliers and the coefficient memory over a similar two independent input spectrometer that might be used with an analog hybrid.

This spectrometer was tested in the laboratory with the DSS-28 receiver using the comb generator for calibration. For simplicity, coax cables were used to connect the baseband signals from the matrix switch to the spectrometer, rather than fiber optic links. A typical spectrum taken after calibration with the comb generator turned on is shown in figure 3.10. The system was calibrated across the band in steps of 2 MHz corresponding to a synthesizer step size of 0.5 MHz. Only 50 frequency steps were needed because the 100 MHz comb allows approximately 10 measurements to be taken across the 1024 MHz band at each frequency step. At each frequency step, 128 sets of raw ADC data were captured from the iBOB using the 2^{17} sample buffer. A 2^{17} point FFT was computed for each set of data from which the phase and amplitude imbalance was extracted at the frequency of each component of the comb signal, as described in the previous section. The measured phase and amplitude imbalances are shown in figure 3.11. The phase imbalance shows a slope characteristic of a delay imbalance as well as more complicated structure. Approximately one day after the calibration data was measured, the calculated correction coefficients were loaded into the spectrometer and the image rejection was measured by injecting a single CW tone at 6500 MHz into the RF input of the receiver. The synthesizer was stepped across the band in 300 kHz steps (corresponding to 1.2 MHz steps at RF) to characterize the response in the center as well as at the edges of the spectrometer



Figure 3.10: Sample spectrum from the quadrature imbalance correction spectrometer with the 100 MHz comb calibration signal turned on. The vertical lines at the top of the plot indicate components from the comb signal, while those at the bottom of the plot indicate the image frequencies of the comb. The good image rejection demonstrated is due to the digital quadrature imbalance correction.



Figure 3.11: Measured amplitude and phase of I(k)/Q(k). The nominal phase difference of $\pm 90^{\circ}$ is shown by the dotted red line. The measured phase difference is shown by the blue dashed line. Applying a phase slope equivalent to a one sample-time delay results in the green curve.



Figure 3.12: Measured image rejection using the digital imbalance correction. The color of each point indicates the 100 kHz step of the test frequency. Frequencies which are multiples of 1 MHz are indicated by 0.0 in the legend and plotted in blue. These points lie in the center of the bin, where the calibration coefficient was measured. The frequencies furthest from the center of the bin are offset by 0.4 and 0.6 MHz, and show the poorest image rejection. This data was measured one full day after the data used to calculate the imbalance correction coefficients was obtained, and thus does not represent the ultimate performance of the system.

channels. The measured image rejection shown in figure 3.12, is considerably better than the 20 dB obtained with the analog hybrid, and the gap in the spectrum around 0 MHz is only about 8 MHz wide, compared to 120 MHz for the analog hybrid. Since these measurements were made a full day after the calibration measurements, they do not highlight the ultimate performance attainable using this spectrometer, but instead show that the image rejection remains usable over long periods of time. The phase and amplitude stability of the fiber optic links while the telescope is in motion has not yet been characterized, and may ultimately limit the performance while observing. The length of the fibers is of order 100 meters, so ripples in the frequency response due to standing waves may have frequency structure as fine as ~ 1.5 MHz, thus it may be necessary to use a spectrometer with finer frequency resolution to effectively correct for this ripple.

The real-time quadrature imbalance correction demonstrated by this spectrometer will be very useful for simplifying future radio astronomy receivers by making wide bandwidth quadrature down-conversion a practical option at centimeter wavelengths. This expands on the work of Fisher and Morgan (2008) by demonstrating a hardware implementation, noting in particular the minimal additional resources required for a pipelined implementation, and by increasing the processed bandwidth to 1 GHz.

As we have seen, the presence of RFI is an inevitable reality for wide bandwidth microwave receivers. In particular, the location of DSS-28 is vulnerable to a wide variety of signals, from ubiquitous cellular base stations and satellite radio to various sources of radar. Coping with this environment requires instrumentation that is capable of both isolating and identifying interfering signals. Polyphase filterbanks provide a excellent method of isolating frequency channels which contain interference from those that do not. By performing short sub-integrations, short segments of data containing impulsive RFI can be removed. Identification presents a more challenging task, and has often required manual intervention to obtain the best results. In this section we discuss an FPGA implementation of a recently developed algorithm which has been demonstrated to be extremely effective at automatically identifying many types of interference.

3.4.1 RFI Detection with the Spectral Kurtosis Estimator

One recent development in RFI excision is the use of the so called *spectral kurtosis* estimator which provides a means of detecting non-Gaussian signals, as described below (Nita et al., 2007). Such kurtosis tests have been around for a while, but they are enjoying increasing popularity because they are well suited to implementation in FPGA based spectrometers. In particular Liu (2008) implemented a spectral kurtosis spectrometer running on an iBOB which provides 4096 channels across 512 MHz of bandwidth using the CASPER FPGA tools. We created a modified version of this design which trades off spectral resolution for bandwidth, providing 1024 channels across 1024 MHz. A second version was also created that is capable of processing two independent 512 MHz polarization channels with 512 spectral channels each. Following the formalism used by Nita et al. (2007), the spectral kurtosis estimator is calculated for spectrometer channel k as

$$\hat{V_k}^2 = \frac{\hat{\sigma_k^2}}{\hat{\mu_k}^2} \tag{3.31}$$

where $\hat{\mu_k}^2$ and $\hat{\sigma_k}^2$ are unbiased estimates of the expected value of the power and of square of the power, respectively, in channel k. In practice, we compute the quantity $S_{1,k}$ which is the sum of the magnitude squared of M consecutive samples of the complex voltage output of channel k, as in a standard spectrometer. In parallel we compute the quantity $S_{2,k}$ which is the sum of the square of the magnitude squared over M consecutive samples. Thus

$$S_{1,k} = \sum_{i}^{M} \hat{P_{k,i}}, \quad S_{2,k} = \sum_{i}^{M} \hat{P_{k,i}}^{2}$$
(3.32)



Figure 3.13: Simplified block diagram of the spectral kurtosis spectrometer

from which we compute estimates of the mean and variance as

$$\hat{\mu_k} = \frac{S_{1,k}}{M}, \quad \hat{\sigma_k^2} = \frac{MS_{2,k} - S_{1,k}^2}{M(M-1)}.$$
(3.33)

Thus the spectral kurtosis estimator can be expressed in terms of measured quantities as

$$\hat{V_k}^2 = \frac{M}{M-1} \left(\frac{MS_{2,k}}{S_{1,k}^2} - 1 \right) \,. \tag{3.34}$$

The utility of this estimator is that for a Gaussian input signal we have

$$V_k^2 = \begin{cases} 2, & k = 0, N/2 \\ 1, & \text{otherwise} \end{cases}$$
(3.35)

Therefore, non-Gaussian RFI signals can be detected by excursions of this estimator. Nita et al. (2007) have studied the response of this estimator versus the duty cycle of the interfering signal and show that for channels $k \neq 0, N/2, V_k^2$ falls below 1 for signals with duty cycles greater than 50% and exceeds 1 for more impulsive signals with duty cycles smaller than 50%. They also point out the interesting fact that the estimator is insensitive to signals with duty cycles near 50%, for which the value of the estimator is 1.

To identify a confidence interval outside of which a signal may be considered non-Gaussian, Nita et al. (2007) have calculated the variance of this estimator in detail, which strictly speaking depends on the window function used in the spectrometer. However, for practical applications, they approximate the variance as

$$\operatorname{Var}\left(\hat{V_k}^2\right) \approx \begin{cases} 24/M, & k = 0, N/2\\ 4/M, & \text{otherwise} \end{cases}$$
(3.36)



Figure 3.14: Plots of equivalent total power. In top plots, the total power is computed by simply taking the mean of all spectrometer channels. In the middle plots the passband has been normalized before taking the mean. The bottom plots show the total power computed after removing data flagged as RFI by the spectral kurtosis estimator. Note that the RFI signals exceed the maximum limits of the top two left plots.

3.4.2 Results

A simplified block diagram of the 1024 MHz bandwidth, 1024 channel spectral kurtosis spectrometer described above is shown in figure 3.13. This spectrometer was tested briefly at DSS-28 to investigate its effectiveness for continuum observations. The telescope was commanded to scan through the radio galaxy Virgo A (3C274) and the quasar 3C286 at 10 mdeg/s in the cross-elevation direction. The observed band spanned 8576 to 9600 MHz, using the lower sideband output of polarization Y. There was considerable strong RFI present during both scans, particularly above 9200 MHz.

Results from the scans are presented in figure 3.14. The top panel in each plot shows the signal that would be obtained using a square-law total power detector to detect the full 1024 MHz band. Notice that the bandpass, shown in the lower panel of figure 3.15, exhibits about 15 dB of slope across it. This results in reduced efficiency of the total power detector, as the effective noise bandwidth is reduced. Provided sufficient dynamic range is available, a digital spectrometer allows one to correct for this effect by dividing out the average passband shape. This has been done to the scans to produce the plot in the middle panels of figure 3.14. The improvement is noticeable, but not overwhelming as RFI is still the dominant feature. The lower panels of figure 3.14 show the results of RFI excision using a simple threshold of the spectral kurtosis estimator. The RFI is very effectively eliminated such that the instability of the receiver now dominates the residual fluctuations. A single sample of the spectral kurtosis estimator is shown in the top panel of figure 3.15. In this case, the estimator is significantly above 1 across the whole band. According to Gary et al. (2009), this effect is a quantization artifact from rounding the lowest bits when computing the power squared term. In the FPGA design there is a provision for selecting which bits to use when computing the power squared term to trade off bit growth and accuracy. This feature was not taken advantage of during these initial tests, but will be used during routine observations. The expected



Figure 3.15: A single 40 ms integration from the spectral kurtosis spectrometer during the 3C274 observation. The lower plot shows the spectrum on a log scale while the upper plot shows the value of the SK estimator for each channel during this integration. The calculated 3σ limits are shown by the dashed lines at the left of the plot. The dashed lines at the right of the plot show the threshold used to reject RFI from this observation.

 3σ bounds on the estimator, calculated using equation 3.36, are shown by the dashed lines on the left hand side, while the dashed lines at the right shows the threshold used for producing figure 3.14, which were chosen by hand. Notice that apart from the offset, much of the band seems to fall within the 3σ range as expected. The 9600 MHz end of the band, which corresponds to the lower frequency end of the baseband signal since this is the lower sideband, shows the roll off of the analog hybrid. In this region, the noise is dominated by the intrinsic noise of the ADC, which is distinctly non-Gaussian as shown by the excursion of the SK estimator. Figure 3.16 shows a waterfall plot of the 3C274 observation after RFI excision. The blanked measurements are white.

An interesting view of the data suggested by Gary et al. (2009) is shown in figures 3.17 and 3.18. The upper panel shows both the average bandpass and the maximum value detected in each channel, with each frequency channel color coded. The lower scatter plot shows the power versus SK estimator for each 40 ms integration at each frequency. If the analyzed signal were purely Gaussian, all of the points would cluster in a horizontal line around SK=1, and indeed this is largely the case for the 3C286 observation shown in figure 3.18. These plots provide insight into the types of RFI experienced. Nita et al. (2007) showed that the SK estimator is generally above 1 for low duty cycle, impulsive interference, while it drops below 1 for more continuous signals. As might be expected, we do not see many points having both low power and SK values below 1, because CW signals would be concentrated into a single frequency bin and thus have higher power. Indeed, most of the points showing an SK value of less than 1 have powers above 10^8 . Impulsive RFI causes arc-like features, the curvature of which Gary et al. (2009) propose is related to the duty cycle of the interference. This makes intuitive sense if one considers a modulated signal that appears as a train of pulses with



Figure 3.16: Waterfall plot of a scan through the radio galaxy 3C274 using the spectral kurtosis spectrometer. The source passes through the beam at 15 seconds. Spectrometer data which has been flagged as RFI appears white in this image. Some remaining RFI is evident.

some length distribution at a given frequency. Since the integration time of the spectrometer is not related to this pulse length distribution, the effective duty cycle will vary within any given 40 ms integration. Integrations that happen to have a low effective duty cycle will logically contain less power but will exhibit higher SK values. As the duty cycle increases and the SK value lessens, the power increases, causing the curved locus observed.



Figure 3.17: Plot of total power versus spectral kurtosis estimator value for the 3C274 observation. The top left plot shows the mean and maximum power in each spectral channel, as well as the color code for each channel. The lower left plot shows the value of the spectral kurtosis estimator versus the total power for every time sample of every spectral channel. In the absence of interference, all of the points would cluster along the line where the spectral kurtosis is 1. The arcs are the result of impulsive interference. The plot on the right is a zoomed-in view showing that many samples are free from RFI. The horizontal dashed lines indicate the predicted 3σ limits on the SK estimator. The offset from 1 of the SK estimator is due to quantization as discussed in the text.



Figure 3.18: Plot of total power versus spectral kurtosis estimator value for the 3C286 observation. See figure 3.17 and the text for details.
Chapter 4

Wide Bandwidth Pulsar Observations

The combination of extremely wide instantaneous bandwidth and large amounts of available observing time make DSS-28 a powerful instrument for studying pulsars. The first two sections of this chapter focus on capturing individual giant pulses at baseband for detailed study. The final section presents how DSS-28 may be used in the future for studying pulsar profile evolution versus frequency and the interstellar medium over unprecedented bandwidths.

4.1 Pulsar Signal Processing

As discussed in chapter 1, the interstellar medium greatly distorts the signal emitted from a pulsar by the time it is received by a telescope on Earth. In particular, dispersion causes the pulse to be stretched out in time, significantly reducing its signal to noise ratio. Fortunately, this effect can be corrected for as described below.

4.1.1 Incoherent Dedispersion

The simplest method for correcting the dispersive effect of the interstellar medium is to first break up the incoming signal into small frequency channels using a filterbank, then to simply delay the detected power series from each channel by the appropriate amount before adding again to form the corrected total power series. This technique is referred to as incoherent dedispersion. The delay for the *n*th frequency channel can be calculated as

$$t_n \approx 4.15 \text{ ms DM}\left(\left(\frac{f_n}{1 \text{ GHz}}\right)^{-2} - \left(\frac{f_0}{1 \text{ GHz}}\right)^{-2}\right)$$

$$(4.1)$$

where f_n is the center frequency of the *n*th channel and f_0 is an arbitrary reference channel which the delays are calculated relative to. If the channel bandwidth, Δf is much less than the observing frequency f, we can calculate the residual dispersion smearing within the channel, Δt , using the differential form of 4.1 as

$$\Delta t_{DM} = 8.3 \ \mu \text{s DM} \frac{\Delta f}{1 \text{ MHz}} \left(\frac{f}{1 \text{ GHz}}\right)^{-3} . \tag{4.2}$$

This shows that to improve the ultimate time resolution by reducing the residual dispersion smearing, one should use the smallest channels possible. Unfortunately, a filter with bandwidth Δf inherently has an impulse response time of

$$\Delta t_{bw} = 1/\Delta f \,. \tag{4.3}$$

Thus the time resolution is ultimately limited to

$$\Delta t = \left(\Delta t_{DM}^2 + \Delta t_{bw}^2\right)^{1/2} . \tag{4.4}$$

The optimal time resolution will be achieved when $\Delta t_{DM} = \Delta t_{bw}$ which implies the optimal channel width is given by

$$\Delta f_{opt} = \sqrt{\frac{f_{\rm GHz}^3}{8.3 \ \mu \text{s DM}}} \,. \tag{4.5}$$

For a total bandwidth B, this implies that the optimum number of spectrometer channels is

$$N_{opt} = \frac{B}{1 \text{ MHz}} \sqrt{8.3DM \left(\frac{f}{1 \text{ GHz}}\right)^{-3}}.$$
(4.6)

4.1.2 Coherent Dedispersion

An alternative strategy for dispersion removal, first introduced by Hankins (1971) and described in detail by Hankins and Rickett (1975), is called coherent dedispersion. This method recognizes the fact that the interstellar dispersion can be modeled as a linear transfer function, and thus application of the inverse transfer function to the raw voltage signal from the telescope before detection will completely undo the dispersive effects. Starting with the expression for the index of refraction for the interstellar medium given in equation 1.3 we can calculate the propagation constant as a function of frequency as

$$k(f) = \frac{2\pi}{c} f \sqrt{1 - \frac{f_p^2}{f^2}}$$
(4.7)

where again, $f_p \approx 1.5 \text{kHz}$ is the plasma frequency. Thus, since $f \gg f_p$, we can write

$$k(f) \approx \frac{2\pi}{c} f\left(1 - \frac{f_p^2}{2f^2}\right).$$

$$(4.8)$$

As a wave propagates a distance d through a medium with propagation constant k(f), its phase change can be calculated as

$$\Delta \phi = -k(f)d. \tag{4.9}$$

We can represent the effect of this phase change as a linear transfer function h(t) acting on the intrinsic waveform $v_{int}(t)$ to yield the observed waveform v(t), presented in the frequency domain as

$$V(f) = V_{int}(f)H(f) \tag{4.10}$$

with

$$H(f) = e^{-jk(f)d} = e^{-j\frac{2\pi}{c}f\left(1 - \frac{f_p^2}{2f^2}\right)d} = e^{-j\frac{2\pi}{c}\left(f - \frac{f_p^2}{2f}\right)d}.$$
(4.11)

At this point it is convenient to make the substitution $f = f_0 + f_{bb}$ where f_0 is the center frequency of the observation and f_{bb} is the baseband frequency yielding

$$H(f_0 + f_{bb}) = e^{-j\frac{2\pi}{c}d\left[f_0 + f_{bb} - \frac{f_p^2}{2(f_0 + f_{bb})}\right]}.$$
(4.12)

Now we can use the identity

$$\frac{1}{f_0 + f_{bb}} = \frac{1}{f_0} - \frac{f_{bb}}{f_0^2} + \frac{f_{bb}^2}{(f_0 + f_{bb})f_0^2}$$
(4.13)

to obtain

$$H(f_0 + f_{bb}) = e^{-j\frac{2\pi}{c}d\left[(f_0 - \frac{f_p^2}{2f_0}) + (1 + \frac{f_p^2}{f_0^2})f_{bb} - \frac{f_p^2}{2(f_0 + f_{bb})f_0^2}f_{bb}^2\right]}.$$
(4.14)

The first term of the exponent does not depend on f_{bb} and represents a constant phase offset which cannot be determined to within 2π and does not affect the final detected signal. The second term is linear in f_{bb} and thus represents a time delay, which can be recognized as the same dispersive time delay calculated in equation 4.1. This can be ignored by correcting the arrival time appropriately. The final term is roughly quadratic in f_{bb} and represents the phase chirp which needs to be unwound to recover the original signal. Therefore the transfer function of interest is

$$H(f_0 + f_{bb}) = e^{j\frac{2\pi}{c}d\frac{f_p^2}{2(f_0 + f_{bb})f_0^2}f_{bb}^2}.$$
(4.15)

As before, we can relate d and f_p to the dispersion measure and dispersion constant (equation 1.8) as

$$\frac{\mathrm{DM}}{\mathcal{D}} = \frac{f_p^2 d}{2c} \tag{4.16}$$

to yield

$$H(f_0 + f_{bb}) = e^{j2\pi \text{DM} \frac{f_{bb}^2}{\mathcal{D}(f_0 + f_{bb})f_0^2}}.$$
(4.17)

To correct the dispersion, we thus need to convolve the measured voltage with the inverse of this transfer function. In practice, we perform this convolution in the frequency domain using the FFT. To do so, we take the FFT of N samples of the measured voltage signal, v(t) to obtain V_n , $n = 0 \dots N - 1$. We then multiply by H_n^{-1} which is $H^{-1}(f_0 + f_{bb})$ evaluated at the center frequency of each bin of the FFT. Finally we take the inverse FFT of the result to obtain the corrected waveform $v_{n,int}$. Note that this operation actually results in a cyclic convolution of H_n and V_n (Brigham 1988). This means that the beginning and ending points are corrupted. We can also see that we must take N large enough to ensure that the phase of H_n does not wrap by more than π between samples, otherwise the phase is ambiguous. Therefore, N should be at least as large as the dispersion sweep time across the band, as calculated by equation 4.1 divided by the sample time. To ameliorate the effects of cyclic convolution, either of the overlap-and-save or overlap-and-add algorithms should be used (Brigham 1988). Finally, to reduce the effects of ringing due to the finite window length, it is common practice to multiply H_n by an apodization window, such as the Hamming window. Any ripples in the bandpass can also be removed by multiplying by a bandpass flattening function.

In the case where the sampled signal is real, the negative frequency components of the FFT can be discarded to yield the analytic version of the signal (this is equivalent to applying an approximation of the Hilbert transform). After multiplying by the dedispersion kernel, the real signal can be recovered by setting the negative frequency components of the FFT to the complex conjugates of $H_n^{-1}V_n$ to ensure that the result of the inverse FFT is real.

4.2 A Sensitive Trigger for Giant Pulses from a Known Pulsar

The full DSS-28 radiometer provides eight 1 GHz signals which are Nyquist sampled and digitized with 8 bits of resolution. This represents an aggregate data rate of 16 GB/s. It would be prohibitively complex and costly to continuously record this data rate to disk for later analysis. Therefore, we need to identify the most interesting portions of the data stream to record, subject to the constraints of the equipment available. Previous studies of giant pulses have used three measurement techniques. Cordes et al. (2004), for example, recorded the output of a filterbank covering 100 MHz and used incoherent dedispersion to process the data. This provided sufficient bandwidth to study scintillation properties, but the time resolution was too low to study individual pulses at high time resolution. Popov and Stappers (2007) continuously recorded 2-bit quantized baseband data from two 10 MHz bands and used coherent dedispersion to process the data. Baseband data recorders designed for Very Long Baseline Interferometry (VLBI), such as the S2, have also been used in a similar fashion (e.g., Popov et al., 2002; Soglasnov et al., 2004). Modern VLBI recorders such as the Mark 5C are

capable of continuously recording up to 4096 megabits per second for 24 hours, and thus provide another useful tool for studying individual pulses (Whitney, 2008). Finally, Hankins and Eilek (2007) used a total power trigger to trigger a high speed digital oscilloscope to capture the widest bandwidth available at an observatory. This technique allows pulses to be studied with the highest level of detail, but the time required to offload the data from the oscilloscope leads to a low usable duty cycle so pulses are often missed. Additionally, the total power trigger is only sensitive to the strongest pulses, and does not provide discrimination against false triggers due to impulsive RFI. Since the dispersion measure of the Crab pulsar is known a priori, it is possible to dedisperse the received signal in real time to form a more sensitive trigger.

Capturing Giant Pulses from B1937+21

While giant pulses from the Crab Pulsar are strong enough to be detected using a total power trigger, as successfully demonstrated by Hankins and Eilek (2007) and others, those from the millisecond pulsar B1937+21 are generally much shorter, so the total energy in a pulse is less, making it difficult to detect in total power (Soglasnov et al., 2004). By capturing these giant pulses at wide bandwidth, the temporal and spectral signatures could be compared in detail to those from the Crab pulsar, potentially leading to a better understanding of giant pulses and pulsar emission in general. To this end, an observing campaign was undertaken to capture giant pulses from B1937+21 at the Arecibo observatory using Hankins' digital oscilloscope triggered by the incoherent dedispersion trigger described below.

4.2.1 Theoretical Discussion and Sensitivity Analysis

As described in section 4.1.1 incoherent dedispersion can be used to partially correct for the dispersive effects of the interstellar medium. In this section, we calculate the expected improvement in signal to noise obtained using this technique. At an observing frequency of 1 GHz, the minimum time resolution possible with incoherent dedispersion calculated from equation 4.5 is $\sim 22 \ \mu s$, and at 10 GHz is $\sim 0.7 \ \mu s$. Giant pulses from the Crab pulsar and B1937+21 generally have similar or shorter timescales (Hankins, 2000; Soglasnov et al., 2004).

The signal to noise of a radio astronomy signal can be expressed as

$$SNR = \frac{T_{psr}}{\Delta T} \tag{4.18}$$

where ΔT is the system RMS fluctuation given by the radiometer equation as

$$\Delta T = \frac{\kappa T_{sys}}{\sqrt{B\tau}} \tag{4.19}$$



Figure 4.1: Illustration of a dispersed signal sweeping through a spectrometer with N channels

and T_{psr} is the increase in total power due to the source. In the case of detecting a pulsed signal, the optimal integration time determined by analogy to matched filtering, and should be equal to the length of the pulse. Now consider a spectrometer with N channels read out with this optimal integration time, as depicted in figure 4.1. When no signal from the pulsar is present, each sample from each channel will have a mean value of T_{sys} . The samples which contain the pulsar signal will have a mean value of $T_{sys} + T_{psr}$. The pulsar will occupy a total of R time-frequency bins, where R is determined by the dispersion sweep time, the channel bandwidth, and the channel sampling time. At each time sample, the pulsar signal will be present in m_i channels of the spectrometer, so that summing over the full length of the pulse, $R = \sum_i m_i$. The RMS noise is additive, and is ΔT in each time sample from each channel.

In the case of no dedispersion, shown in figure 4.1, it can be seen that the total power given by averaging all spectrometer channels together is $T_{sys} + m_i T_{psr}$ and the RMS is $\Delta T \sqrt{N}$. Therefore the pulsar total power signal to noise at time sample *i* would be

$$SNR_{tp,no\ int} = \frac{m_i T_{psr}}{\Delta T \sqrt{N}} \,. \tag{4.20}$$

In this case, the pulse will be spread over L_{tp} time samples, so we can improve the SNR by integrating



Figure 4.2: Illustration of a dispersed signal sweeping through a spectrometer with N channels after delaying each channel by the appropriate amount to correct for the dispersion

by a factor of L_{tp} . This yields an integrated total power SNR of

$$SNR_{tp,int} = \frac{\frac{RT_{psr}}{NL_{tp}}}{\frac{\Delta T}{\sqrt{NL_{tp}}}}$$
(4.21)

$$=\frac{RT_{psr}}{\Delta T\sqrt{NL_{tp}}}\,.\tag{4.22}$$

In contrast, when we use incoherent dedispersion to align all of the bins that contain the pulse energy, as shown in figure 4.2, we obtain a dedispersed total power signal with mean $T_{sys} + RT_{psr}/NL_{idd}$ and RMS = $\Delta T/\sqrt{NL_{idd}}$. In this case, the SNR is

$$SNR_{idd} = \frac{\frac{RT_{psr}}{NL_{idd}}}{\frac{\Delta T}{\sqrt{NL_{idd}}}}$$
(4.23)

$$=\frac{RT_{psr}}{\Delta T\sqrt{NL_{idd}}}\,.$$
(4.24)

Therefore we can calculate the improvement in SNR as

$$\frac{\text{SNR}_{idd}}{\text{SNR}_{tp}} = \sqrt{\frac{L_{tp}}{L_{idd}}}.$$
(4.25)

By definition, the optimal channel bandwidth and time resolution given by equation 4.5 minimizes L_{idd} . If the dispersion sweep were a purely linear chirp and the filterbank channel spectral leakage is approximately zero, the dispersed pulse will be present in each channel for exactly one integration

69



Figure 4.3: Schematic of the iBOB-based real-time incoherent dedispersion system

time, τ_{opt} . In this case, $L_{idd} = \tau_{opt}$ and $L_{tp} = N_{opt}\tau_{opt}$ so the SNR is improved by $\sqrt{N_{opt}}$, where N_{opt} is the number of spectrometer channels that results in the optimal channel width $f_{opt} = 1/\tau_{opt}$. In practice, this level of improvement will not be realizable, in particular because the number of channels in a digital filterbank is generally constrained to be a power of two. Therefore, the maximum SNR is often found experimentally to occur when $L_{idd} = k\tau_{bw}$ for k ranging from 2 to 4 or more. Here τ_{bw} is the integration time per spectrum, given by the channel width.

4.2.2 System Description

The design of the incoherent dedispersion trigger used at Arecibo is shown in figure 4.3. The first section is a basic polyphase filterbank spectrometer, which takes in four real voltage samples per FPGA clock cycle and computes the complex voltage output from two filterbank channels per FPGA clock. Therefore, every 512 ADC clocks, representing 128 FPGA clocks, a 256 point complex spectrum is calculated. The real and imaginary parts of the complex voltage are then scaled by an equalization constant which can be set independently per channel. The resulting equalized signal is then detected and accumulated in the same fashion as a normal spectrometer in a vector accumulator that contains one accumulator per spectral channel. This provides a diagnostic view of the average spectrum used for calculating the necessary equalization constants, detecting channels corrupted by RFI, and checking the quality of the passband.

The detected filterbank outputs are also routed to the incoherent dedispersion engine. First the bit width is reduced to 8 bits to reduce the amount of FPGA resources required. The core of the incoherent dedispersion is similar to the vector accumulator, except that the relative offset of the incoming channel is read from a lookup table instead of simply being incremented by one each cycle. To account for necessary latency in the system, the core only accepts a new sample every fourth clock period. Therefore four cores are used to process each of the two outputs from the filterbank necessitating a total of eight cores, labeled 'IDD' in figure 4.3. The resulting eight signals will still have some relative delay which is then removed by a set of programmable delay blocks. The aligned signals are finally summed together to produce the dedispersed total power output signal which has a data rate of one sample per spectrum, thus 1/512th of the ADC clock rate. This dedispersed signal is then compared against a threshold to provide a trigger when a pulse arrives. In the stand alone design, the threshold is computed on the FPGA in the following way. The dedispersed time series is fed into a simple one stage cascade integrator comb (CIC) filter which calculates the approximate running mean with a time constant of approximately 1 second. The mean is subtracted from the dedispersed total power and the absolute value is taken resulting in the absolute deviations. The absolute deviations drive another CIC filter with a 1 second time constant which computes the running mean absolute deviation (MAD). The MAD is a statistical measure of variation defined as

$$MAD = \langle |x_i - \bar{x}| \rangle_i . \tag{4.26}$$

This number is multiplied by a run-time programmable value and added back to the mean to provide a trigger. For Gaussian variables, the MAD is related to the standard deviation by

$$\frac{\text{MAD}}{\sigma} = \sqrt{\frac{2}{\pi}} \approx 0.8 \tag{4.27}$$

(Krishnamoorthy, 2006). The dedispersed time series was experimentally found to obey this relation to better than a few percent, even in the presence of moderate (non-Gaussian) RFI. This is to be expected by the Central Limit Theorem since the dedispersed time series is the sum of a large number of approximately independent random variables. Calculating the MAD is considerably simpler on an FPGA than the standard deviation, as it does not require multiplications or square roots.

For the stand alone version, the ADC values are scaled by a run-time programmable coefficient, clipped to 4-bits, and stored in an on-chip circular buffer. The resources available on an iBOB allowed for 2^{17} samples to be stored per polarization on-chip. This corresponds to 512 μs for 128 MHz bandwidth. The iBOB also has two external 18 megabit SRAM chips which could be used for a stand alone system with an even deeper buffer. In the version designed to be used with the BEE2, described in section 4.3, the full 8-bit ADC data, representing 16 Gbit/s at 2 Gsps is split between the two XAUI ports and sent to the BEE2 for buffering. The stand alone version also used a simple R-2R resistor ladder connected to spare I/O pins as a digital to analog converter to provide an analog representation of the dedispersed total power for diagnostic purposes and to trigger the



Figure 4.4: An example of a giant pulse captured during initial testing of the iBOB based incoherent dedispersion system on DSS-28. The short segment of raw voltage data from each polarization has been rendered as a spectrogram in the top two panels. The lower left panel shows the sum of these two spectrograms, while the lower right panel shows the time history of the dedispersed total power around the trigger event.

high speed oscilloscope.

4.2.3 Initial Tests at DSS-28

Before going to Arecibo, the trigger system was tested extensively in the laboratory using a swept synthesizer to generate a chirped pulse. Additionally, the system was tested on DSS-28 using just the iBOB on its own. Limited telescope time was available, so the system was only tested on the Crab pulsar. The telescope did not have a pointing error model at the time, so the Crab pulsar was a much easier target because the telescope could be peaked on the bright, co-located, Crab Nebula. The tests were performed with the lower sideband output from each polarization connected to one of the two iBOB ADC inputs. The 22.1–22.5 GHz narrow band IF filter was selected since only one sideband was used. The baseband filter was set to 270–370. The ADC was clocked at 256 MHz, so the third Nyquist zone, 256–384 MHz, contained the 270–370 MHz baseband signal. The limited observing time precluded testing at a wide range of frequencies, but several pulses were captured in the range of 4–5 GHz. An example of one pulse is shown in figure 4.4



Figure 4.5: Comparison of total power trigger and incoherently dedispersed trigger on the Crab pulsar at 6500 MHz. The top trace shows the incoherently dedispersed total power as computed by the iBOB. The middle trace shows the total output from a square law detector previously used as a trigger. The bottom trace is a gate signal which was used to trigger this oscilloscope so that the main pulse and interpulse regions remain stationary on the screen. The oscilloscope was set to have infinite persistence, so the traces shown are many sweeps overlayed, and clearly shows the peak-to-peak noise level of the signals. Notice that the total power trigger is only sensitive to the largest pulses while the incoherently dedispersed signal detects many more pulses. The strongest pulses go off the top of the screen.

4.2.4 Results from Arecibo

4.2.4.1 Testing on the Crab Pulsar

Our Arecibo proposal was awarded three full 4 hour passes observing the Crab pulsar for system testing and four full passes for the actual observations on B1937+21. During the tests on the Crab pulsar, we extensively vetted the system, capturing pulses at various bands between 2.5 and 10 GHz. The enormous improvement in sensitivity of the incoherent dedispersion trigger versus the total power trigger is illustrated in figure 4.5. The figure is a screen capture of a diagnostic oscilloscope¹ which was triggered using a locally generated signal which was synchronous with the pulsar period. The top trace shows the dedispersed total power signal from the R-2R DAC on the iBOB, while the lower trace shows the total power signal traditionally used as a trigger. Clearly the incoherent dedispersion system is sensitive to a much wider range of pulse amplitudes.

¹This is not the wide bandwidth oscilloscope used to actually capture the baseband data.



Figure 4.6: Average profile of B1937+21 obtained by averaging the incoherent dedispersion output synchronous to the pulsar period

4.2.4.2 Observations of B1937+21

Confident that our trigger system was working well, the first day of observing B1937+21 we tried our lowest planned frequency band, around 2.75 GHz, expecting to see numerous pulses. We were quickly disappointed when we did not observe any pulses. Eventually we captured a single pulse at this frequency. We tried a few higher frequencies and still did not see anything. Finally, at the end of the first day, we switched to a bright pulsar B2020 to check the trigger system, which proved to be working still. In light of this behavior, we devised a new plan to observe at 1.4 GHz where the pulsar is known to be strong. At this frequency, we immediately were able to measure an average profile of the pulsar, shown in figure 4.6 by simply setting the diagnostic oscilloscope to average traces synchronous with the pulsar period. Giant pulses were readily detected as shown in Figure 4.7. We also observed the pulsar at 1.8 GHz. Unfortunately, limitations of the oscilloscope prevented us from capturing more than 125 MHz of bandwidth at baseband at these lower frequencies.

4.2.4.3 Conclusions

While our observing campaign was unsuccessful at capturing giant pulses from B1937+21 with bandwidths greater than 250 MHz, the utility and great potential of the real time incoherent dedispersion trigger were immediately obvious. We suspect we did not manage to capture pulses reliably above 2



Figure 4.7: The incoherently dedispersed trigger signal on B1937+21 at 1.4 GHz. In this case, the square law detector signal was not monitored at all because the dispersion smearing time across the 22 MHz trigger bandwidth is greater than the pulsar period. The main pulse and interpulse are both visible as slight rises in the dedispersed signal. In particular, since the oscilloscope was set to infinite persistence mode, the fact that the lower edge of the dedispersed signal is raised relative to the off pulse peak to peak level implies that *every* individual pulse from this pulsar is being detected.

GHz because our estimate of the spectral index of pulsar was overly optimistic and because the pulses are so narrow that despite the advantage provided by the incoherently dedispersed trigger system, the residual smearing significantly reduced the signal to noise of any giant pulses. In the future it would be very interesting to repeat the experiment with a real time coherent dedispersion trigger system. Real time coherent dedispersion over a modest $\sim 10-50$ MHz bandwidth should be possible in real time with modern graphics processing units (GPUs), and would provide a very sensitive trigger to extremely narrow dispersed pulses like those from B1937+21. This has been demonstrated using several 4 MHz sub-bands across 128 MHz, but is generally used for timing studies rather than giant pulse detection (Cognard, 2009). This pulsar has been studied in the past using a hardware based real-time coherent dedispersion system based around a special purpose integrated circuit designed for implementing chirped filters (Hankins and Rajkowski, 1987). However, this system only provided 2 MHz of bandwidth and was used for specifically for timing studies rather than giant pulse detection.

The limitations of the oscilloscope which prevented capturing more than 250 MHz bandwidth at 2500 MHz are precisely the motivation for the design of the deep capture transient buffer described in the next section.

4.3 A Deep Capture Buffer for Detailed Multi-GHz Analysis of Single Pulses

In this section we present a special purpose instrument optimized for capturing giant radio pulses from known pulsars, in particular the Crab pulsar. This instrument exhibits numerous advantages over other hardware used for studying giant pulses, and in particular, is specifically designed to take advantage of the wide instantaneous bandwidth available at DSS-28. At the same time, it is flexible enough that it could be easily deployed at another telescope in the future, particularly if CASPER hardware is already available at the site.

4.3.1 System Description

4.3.1.1 The Incoherently Dedispersed Trigger

The incoherent dedispersion trigger logic is very similar to that described for the stand-alone system in section 4.2.2. The iBOB implementation is shown schematically in figure 4.8. When capturing data at the full 1 GHz bandwidth, the ADC presents 8 samples per FPGA clock. Thus the polyphase filterbank produces 4 complex samples per clock. At full time resolution, this would require 16 incoherent dedispersion cores. In addition, for a given observing frequency, the number of dedispersion lags required increases quadratically with trigger bandwidth. Therefore, to demonstrate the design,



Figure 4.8: Functional diagram of the iBOB portion of the Deep Capture Buffer showing the incoherent dedispersion trigger and the path of the raw ADC data to the 10 Gb/s XAUI links which transfer the data to the BEE2 board. Note that a digital down conversion can proceed the polyphase filterbank to reduce the trigger bandwidth if needed.

a scaled back system was designed. The chief difference is the data rate and trigger bandwidth is reduced by 1/4th by digitally down converting the center quarter of the sampled bandwidth. For a sampling rate of 2048 MHz, this corresponds to the band from 384 MHz to 640 MHz. The digital down converter takes in 8 ADC samples per clock and outputs 1 complex sample. This single complex data stream is then analyzed by a complex input polyphase filterbank. The rest of the trigger proceeds as in the stand alone version, but requiring only four incoherent dedispersion cores. As in the stand-alone system, the combination of the polyphase filterbank and equalizer allows one to excise any channels with RFI to avoid false triggers by setting the equalizer value for those channels to zero. In addition, to simplify the design, the threshold level is calculated by the host computer instead of on chip. Finally, the iBOB also transmits the baseband ADC data via the two XAUI links to the BEE2 portion of the system.

4.3.1.2 The Deep Capture Buffer

As described in the overall architecture section, two iBOBs are connected to each BEE2 User FPGA. As shown in figure 4.9 User FPGAs implement one deep capture buffer (DCB) per iBOB which operate essentially independently. Each DCB is composed of two 1 Gbyte DRAM modules with only one being written to at a time. The 1 Gbyte memory space is logically divided into a number of equal sized circular buffers which can be configured at run time. The number of buffers must be between 1 and 256 and must be a power of two in the current design.

The operation of the DCB is shown in the cartoon in figure 4.10. In the example depicted, each DRAM is divided into 4 circular buffers, each of 256 megabytes. In quiescent operation, shown in the first frame, the incoming data is read from the XAUI ports and is continuously written into circular buffer 0 on DRAM0. At any given time, this circular buffer holds the past 256 megasamples of ADC data, corresponding to 128 ms at a sample rate of 2 Gsps. When a trigger is received, the buffer controller immediately stops writing to DRAM0 and starts writing to a new circular buffer on DRAM1 as shown in the second frame. The readout controller then detects that DRAM0 is idle and contains data ready to be dumped to the 10 GbE port, so it requests access to the 10 GbE port



Figure 4.9: Diagram of the BEE2 portion of the Deep Capture Buffer. The trigger distribution system is not shown and resides in the Center FPGA. In the initial version of the system, two quadrants of the BEE2 were used to buffer data from the four iBOBs. The data arbiter and 10 Gb Ethernet links, shown here in the Center FPGA, and the trigger distribution system were instead implemented in a third User FPGA.



Figure 4.10: Operational flow of the deep capture transient buffer. The two DRAM modules, DRAM0 and DRAM1, are each divided into 4 circular buffers.

from the arbiter. When the 10 GbE port becomes available, the readout controller starts sending the data in the buffer to the 10 GbE port (frame 3). While this is happening, if another trigger comes in, the buffer controller detects that DRAM0 is busy being read, so instead of switching to another buffer on DRAM0, it looks for an available buffer on DRAM1 and begins writing to that buffer (frame 4). Buffer 0 on DRAM1 is then queued to be read out when possible. If successive triggers arrive while DRAM0 is still being read, this process will repeat, with new buffers being used on DRAM1 until all of the available buffers have been exhausted. After the readout controller is finished reading dumping the initial buffer from DRAM0, it sees that there is a buffer ready for reading in DRAM1, but that DRAM1 is busy being written to (frame 5). Thus it must wait until the next trigger which will cause the buffer controller to begin writing to a new buffer in DRAM0, freeing up DRAM1 to be read out (frame 6). After Buffer 0 on DRAM1 is read out, the readout controller retains control of the 10 GbE port so it can read out Buffer 1 on DRAM1 (frame 7). Finally, the readout controller returns control of the 10 GbE port to the arbiter, and the system returns to the quiescent state awaiting the next trigger (frame 8).

This scheme allows bursts of giant pulses to be captured without exceeding the maximum data rate sustainable by the computer receiving the data from the BEE2. The trigger system aggregates trigger signals from all eight iBOBs. Each trigger has a programmable hold-off value to allow some amount of data to be recorded after a trigger event. This is particularly important in the case of triggering using the upper sideband; the hold-off can be used to account for the dispersion sweep time between the trigger frequency and the bottom of the lower sideband. A programmable intertrigger delay is also available to avoid the potential situation where the time between trigger events is shorter than the length of the circular buffer. Without this feature, the circular buffer might not have a chance to fill with valid data before a trigger forces it to be read out. The versatility of the trigger system is designed to enable a pulse to be detected in any sub-band to be captured in its entirety across all sub-bands.

4.3.2 Results from Initial Testing on DSS-28

During the DSS-28 prototype testing period a reduced version of this deep capture transient buffer system was installed and tested. Since the prototype receiver provides a total of four baseband signals from one receiver chain, only four iBOBs were necessary. These were connected to User FPGAs 2 and 4 on the BEE2. In this reduced version, User FPGA 1 was used for the 10 GbE connection, arbiter, and trigger aggregation because reconfiguring the User FPGAs is simpler than the Center FPGA. Approximately 4 full days of observing were available: July 31 and August 04–06, 2009. The primary goal of the observations was demonstrating the functionality of the system with real giant pulses. A second goal was getting a feel for the sensitivity of the system at various observing frequencies, with an eye towards future scientific observing campaigns with the full system.



Figure 4.11: Measured system temperature of DSS-28 on (dots) and off (solid red line) the Crab Nebula. The higher set of T_{sys} values between 5 and 7 GHz were due to non-linear effects caused by DSS-27 transmitting at around 2 GHz. The solid blue line is the expected value of T_{sys} on the Crab. See text for details.

Unfortunately, a problem with one of the fiber optic links significantly reduced the gain in the upper sideband of polarization X. Despite this limitation, some pulses were bright enough to show up in this polarization. The gain of the upper sideband of polarization Y was also somewhat reduced, but not as significantly. For simplicity, the most sensitive channel, polarization Y, lower sideband, was used alone to provide the trigger for the results presented here. The system was also tested successfully using both sidebands of polarization Y as triggers, but the relatively primitive control software used for these initial tests made this mode somewhat difficult to work with for the time being.

A glitch in the software on the control computer prevented the time tags of the triggers from being recorded properly, thus the time of arrival of the pulses is not known with sufficient precision to determine the rotational phase of the pulses.

A crude calibration was performed at most observing frequencies to determine the system temperature on the nebula. The results of these measurements are presented in figure 4.11. The system temperature off source is shown by the solid red line. The different colors of dots represent the system temperature on source during different observations. The expected value of system temperature was calculated by scaling expected contribution from the nebula by an aperture efficiency of 40% and by the ratio of the calculated solid angle of the antenna divided by a solid angle of 50 square arcminutes. The efficiency and solid angle value were determined experimentally to give the best fit, and seems to match the data well. This is somewhat surprising as the efficiency was shown before

Center Frequency	Number of Pulses	Time Span	MJD
(MHz)		(\min)	
2940	30	37	55049.929
3140	15	99	55058.875
3620	7	104	55048.599
3820	28	63	55049.753
4540	7	87	55047.706
6420	13	99	55047.772
6672	12	143	55049.487
6820	1	61	55047.863
7200	1	94	55049.649
7940	4	97	55049.855
11200	9	33	55043.897

 Table 4.1: Summary of initial wide bandwidth giant pulse capture observations with DSS-28

 Center Frequency Number of Pulses Time Span

 MID

to be somewhat higher than 40% and the geometric mean of the angular extent of the Crab Nebula given by Cordes et al. (2004) is 5.5 arcminutes instead of the 8 arcminutes used here. Nonetheless, it is apparent that the system temperature is dominated by the nebula below about 8 GHz. Since the spectrum of the Crab Nebula is relatively flat, given in janskys by $955f^{-0.27}$, dividing out the bandpass shape effectively cancels any variation in gain across the band.

Observations were performed at several frequencies between approximately 1.9 and 12 GHz, summarized in table 4.1. The receiver was configured to use the 2 GHz wide IF filter and the analog hybrid was used to provide upper and lower sideband outputs. Both sidebands from both polarizations were sampled with a sampling frequency of 2048 MHz.

A selection of the captured pulses are presented in figures 4.3.2 through 4.3.2. Each subfigure contains a single pulse. All of the pulses are presented on the same time scale for ease of comparison. The pulses are arranged in order of increasing center frequency, and were generally selected to show the variety of pulses observed at each frequency. Each pulse was coherently dedispersed as described in section 4.1.2, and the off-pulse bandpass was flattened. A range of dispersion measures was tried for each pulse, and the DM which resulted in the least residual slope in a spectrogram plot of the data was chosen. The best DMs range from 56.7945 to 56.7970 pc cm⁻³. This agrees favorably with the value given by the monthly ephemeris provided by Jodrell Bank Observatory² of 56.7949 pc cm⁻³ for August, 2009 (Lyne et al., 1993).

The threshold used to detect giant pulses was nominally set to 6 times the RMS value of the dedispersed signal above the mean calculated over a period of \sim 5 seconds. However, in many of the observations, the mean value of the dedispersed signal varied considerably on time scales of \sim 1–10 seconds because of interference from DSS-27 or windy conditions buffeting the tertiary, as

²http://www.jb.man.ac.uk/~pulsar/crab.html

discussed in section 3.2. Thus the trigger did not accurately track the mean signal and the effective threshold level was often significantly higher than 6σ . Numerous false triggers occurred when DSS-27 was transmitting and when sharp gusts of wind perturbed the tertiary. Despite this, during calm periods relatively weak pulses were captured, for example figures 4.14(b) and 4.15(f).

The number of giant pulses detected also depends heavily on interstellar scintillation. For example, in much more sensitive observations made using the NRAO Very Large Array by Moffett and Hankins (1996), they report that the Crab pulsar was only detected in two out of nine 3-hour observing sessions at 8.4 GHz, and only four out of nine at 4.9 GHz. Additionally, they mention that giant pulses typically occurred in bursts of 10–20 minutes during which pulses were detected almost every rotation of the pulsar. Since these scintillation effects are strongly frequency dependent, the sensitivity of the final DSS-28 system will be greatly improved by being able to trigger on giant pulses occurring anywhere in a 4 to 8 GHz bandwidth.

The dedispersed time series from each pulse is presented in several ways. For the pulses presented in figures 4.3.2 through 4.3.2 both polarizations were recorded and are presented. At the top of these figures, the detected dedispersed total power time series is presented, with polarization X on the left and polarization Y on the right. In these top plots, the blue curve represents the total power in the lower 1024 MHz sideband, while the green curve represents the upper sideband. Each curve has been divided by the off-pulse RMS level and smoothed to 31.2 ns time resolution. Directly below the total power time series plots are spectrograms of the raw dedispersed data for each polarization. The spectrograms were computed by dividing the time series from each 1024 MHz sideband into segments of length 128, multiplying each segment by a Hamming window, and computing the FFT and squaring. Since the original signal is real, only the positive frequencies of the FFT are shown. The spectrogram data has been divided by the off-pulse RMS level of the respective polarization and sideband. The dotted white lines on the spectrogram indicate the bandwidth used to trigger the system. To provide a view of the on-pulse spectrum, the samples of the spectrogram between the vertical dotted red lines at the top of the figure have been averaged over time and plotted at the far left of each figure. The blue curve is polarization X and the green curve is polarization Y. Similarly, the frequency channels of the spectrograms have been divided into eight groups which span the width of the eight subplots at the far right of each figure. The samples from each group have been averaged over frequency and are plotted in the corresponding subplot at right. Again, the blue curves represent polarization X and the green curves are polarization Y. This provides a view of the pulse shape versus frequency. Note that each of these subplots has been 'auto-scaled,' so care must be taken when interpreting the relative strength of the pulse components at each frequency. For the observations at 3620 MHz, only the upper sideband is presented because the buffer was incorrectly configured, so the full extent of the dispersed pulse was not captured in the lower sideband.

In the final set (figure 4.3.2), only one polarization was available. The plots presented are

essentially the same, but only include polarization Y. The spectrum in the far left plots includes the on-pulse spectrum in red and the off-pulse spectrum in blue.

Various artifacts are noticeable in the plots. In some plots, the dip in the center of the band caused by the roll-off of the analog baseband hybrid is visible. Since the two sidebands were normalized separately for the spectral plots, the relative levels in the two sidebands can differ considerably. Some residual temporal misalignment between the sidebands may exist because the local oscillator was not locked to the station reference, so the offset had to be determined by eye.

4.3.3 Interpretation

While we must stress again that this small set of observations was aimed solely at demonstrating the functionality of the instrument, the scientific possibilities are readily apparent. This data set will probably raise more questions than it answers, but hopefully the ongoing observations discussed in section 5.1 will provide a more cohesive picture.

The most notable feature of this data set is the incredible variability in pulse width, amplitude, shape, and bandwidths exhibited by giant pulses across the radio spectrum. Such variability has been discussed previously (e.g. Crossley et al. (2009), Hankins and Eilek (2007)), but the wide fractional bandwidth of the lower frequency observations in particular provides confirmation of previous suggestions and provokes new questions. For example, Crossley et al. (2009) demonstrated that at frequencies above 1 GHz, pulse broadening scales as f^{-2} instead of as f^{-4} which would be expected from scattering in the interstellar medium as discussed in section 1.3.3.1. This result was obtained by examining statistics of the pulse width in observations over 50 MHz of bandwidth at 1.4, 1.7, and 4.8 GHz. A cursory glance at the pulses captured over a ~2:1 bandwidth suggests that the pulses do not broaden by the factor of 16 expected from the scattering theory presented in section 1.3.3.1. A detailed plot of the pulse from figure 4.12(d) is presented in figure 4.17, which seems to show some frequency-width relationship, but not nearly as strong as f^{-4} . Following the method suggested by Crossley et al. (2009), we have fit a 'fast-rise exponential decay' curve to the pulses to analyze the frequency-width relationship in more detail. The fitting function is given by

$$F(t) = \begin{cases} A(t - t_0)e^{-(t - t_0)/\tau}, & t > t_0 \\ 0, & \text{otherwise} \end{cases}.$$
 (4.28)

Figure 4.18 presents the widths extracted from the curve fits versus frequency for five pulses captured at 2940 MHz. The pulses were chosen to have a single dominant component for simplicity. The f^{-4} frequency relationship expected from the ISM and the f^{-2} relationship which Crossley et al. (2009) found to be in better agreement with the widths they observed and attributed to the pulsar itself are also plotted for comparison. The two lines intersect at the 333 MHz, 200 μs data point presented by



Figure 4.12: Giant pulses from the Crab Pulsar



Figure 4.13: Giant pulses from the Crab Pulsar



Figure 4.14: Giant pulses from the Crab Pulsar



Figure 4.15: Giant pulses from the Crab Pulsar

87



Figure 4.16: Giant pulses from the Crab Pulsar

Crossley et al. (2009). While this initial analysis is far from definitive, the results generally confirm the conclusion by Crossley et al. (2009) that the pulse width above 1 GHz is not dominated by the ISM.

Cordes et al. (2004) presented statistics of the number of giant pulses detected versus pulsar rotational phase and radio frequency. This data shows a striking evolution of giant pulses at the phase of the main pulse and the interpulse. In particular, giant pulses are far more common at the phase of the main pulse from 430 MHz to 4.8 GHz. However, they suddenly become much less frequent at 8.8 GHz as the interpulse takes over. While infrequent, giant main pulses above 6 GHz are known to be very bright (Hankins and Eilek, 2007). The nature of the interpulse giant pulses also changes dramatically. At 430 MHz, giant interpulses are nearly as strong and as frequent as the giant main pulses, but appear to decrease in strength at higher radio frequencies, almost disappearing completely at 3.5 GHz. Then, around 4.15 GHz, a new component begins to appear at a slightly earlier rotational phase than the low frequency interpulse (Moffett and Hankins, 1996). This new interpulse component is the dominant source of giant pulses above ~ 8 GHz. The data presented by Cordes et al. (2004) hints at spectral features of the giant pulse components, but without simultaneous observations across the frequency range, it is impossible to fully characterize the emission. These are precisely the kinds of observations that are now possible with DSS-28, and already we see some tantalizing hints in this initial data set. For example the pulse in figure 4.14(a) seems to show components with negative (first component) and positive (second component) spectral indices. While unlikely, such a phenomenon could potentially be explained by overlapping giant pulses from the low frequency and high frequency interpulse components. The peak of the high frequency interpulse moves to an earlier rotational phase, but the range of phases at which



Figure 4.17: Detail of pulse width versus frequency for the pulse shown in figure 4.12(d). Each solid curve represents the total power time series from one 256 MHz sub-band smoothed to 31.2 ns. The top trace is centered at 3836 MHz, while the bottom trace is centered at 2044 MHz. The dashed curves show 'fast-rise exponential decay' curve fits to the pulses.



Figure 4.18: Pulse width versus frequency plotted for 5 pulses captured at 2940 MHz, including that from figure 4.12(d) (red dots). Each pulse is plotted in a different color. The upper dashed blue line represents the f^{-2} relationship found by Crossley et al. (2009), while the lower dashed red line represents the f^{-4} relationship predicted from the simple scattering theory presented in section 1.3.3.1. The pulse widths were extracted from the 'fast-rise exponential decay' curve fits.

high frequency interpulse giant pulses are observed slightly overlaps the phases of the low frequency interpulse (Cordes et al., 2004). Several of the pulses we observe (for example, figures 4.12(c), 4.12(e), and 4.13(e)) show a high frequency cutoff that may reflect the sudden decrease in occurrence of giant pulses at the main pulse phase above 5 GHz. This is unlikely to be an instrumental effect, because pulses captured with the same receiver configuration and processed identically do not show this dramatic roll-off at high frequencies (see figures 4.12(d), 4.12(f), and 4.13(f), respectively). Also notice that even though the upper sideband is less sensitive, the pulse in figure 4.12(c) has very high SNR even in the lowest sub-band of the upper sideband, then sharply decays in the higher subbands. Again, we recognize these explanations may not be convincing, especially since the rotational phase of this data set is not available, but they point to the kinds of phenomena that can now be investigated in detail.

The pulses observed at 11.2 GHz in figures 4.3.2 illustrate the need for wider trigger bandwidths. Many pulses were probably missed during the 11.2 GHz observing session because they peaked outside of the trigger band. These giant pulses are unlikely to be interpulses because, as Hankins and Eilek (2007) point out, all interpulses observed above ~6 GHz exhibit frequency bands, the widths of which scale as as $\Delta f = 0.06f$. At 11 GHz, we would expect to see bands spaced approximately 660 MHz apart. On the other hand, main pulses are known to show relatively narrow bandwidth bursts of emission such as we see here (Hankins and Eilek, 2007). Note that the wide instantaneous bandwidth presented here is somewhat disorienting as one might be inclined to suggest that these relatively narrow bandwidth bursts could be RFI. This is extremely unlikely, however, because even at this high frequency, the dispersion smearing time across even 100 MHz of bandwidth is ~ 35 μs , much longer than the timespan shown here. Additionally, there appears to be a hint of emission in the upper sideband near 12 GHz in figure 4.16(a).

Figure 4.19 presents a closer look at the pulse shown in figure 4.14(c). The layout of this plot is the same as in the original, but now the time series from each sideband is displayed at full 1 ns resolution over a span of approximately 120 ns. Fewer channels were used for the spectrogram to trade off frequency resolution for higher time resolution. At this time resolution, errors in the dispersion measure at the level of 0.001 pc cm⁻³ are readily apparent. Therefore, to produce this plot, this pulse was dedispersed over a finer range of DMs, and the best plot, judged by residual slope in the spectrogram, was chosen. The pulse exhibits structure which is unresolved on timescales of 1 ns. This is characteristic of giant main pulses at these frequencies (Hankins and Eilek, 2007). In the future, we hope to capture up to 8 GHz of contiguous bandwidth with this system and recombine the 1 GHz sub-bands to attain a time resolution of ~0.125 ns to investigate these nanoshots in further detail. Even if this proves infeasible, the 8 GHz dynamic spectra will provide unprecedented information about the giant pulse phenomenon. A commercial digital oscilloscope could also be connected to the full bandwidth port available on the receiver to digitize the full bandwidth of the



Figure 4.19: Full time resolution plot of the pulse in figure 4.14(c) showing nanosecond time scale structure

LNA to provide the ultimate time resolution.

4.3.4 Summary and Future Improvements

The initial results from the incoherent dedispersion trigger and deep capture buffer show great potential for performing unique scientific observations of giant pulses. The lowest frequency observations, centered at 2940 MHz, represent the widest instantaneous fractional bandwidth observations to date. Aside from the intrinsic wide bandwidth of the receiver, two main features of the digital capture system were critical to the success of these measurements.

First, the incoherent dedispersion trigger both increased sensitivity to the dispersed giant pulses, and at the same time reduced sensitivity to broadband impulsive interference. Additionally, channels with narrow bandwidth impulsive interference could be blanked using the polyphase filterbank and equalizer. This reduced the occurrence of false alarms to a tolerable level, which would not have been possible otherwise at these frequencies and bandwidths.

Second, the limited capture depth of digital oscilloscopes has previously prevented recording baseband voltage data at such low frequencies and wide bandwidths because of the long dispersion sweep time. The observations at 2940 MHz required nearly 256 megasamples per polarization, whereas the oscilloscope used for previous observations by Hankins and Eilek (2007) was limited to 128 megasamples. This oscilloscope has an additional disadvantage of an artificial limitation imposed

91

by the manufacturer on the combination of bandwidths and sampling rates available (T. Hankins, private communication). Additional capture depth carries a heavy cost premium for commercial oscilloscopes, whereas the BEE2-based digital capture buffer uses inexpensive commodity DRAM. Finally, the ability to partition the BEE2 capture memory into several sub-buffers makes false triggers caused by RFI more tolerable, and ensures that a higher fraction of the received pulses are captured.

By enabling real-time detection of giant pulses, this system is capable of combining the best aspects of sensitive statistical studies made by continuously recording a modest bandwidth (e.g., Popov and Stappers, 2007; Soglasnov et al., 2004) with the extreme detail and bandwidth information obtained by oscilloscope based studies (e.g., Hankins and Eilek, 2007). Additionally, the hardware available at DSS-28 should also be capable of simultaneously running the deep capture buffer system, as well as providing continuous incoherent filterbank data, thus allowing studies like that of Cordes et al. (2004) to be performed while capturing baseband data from each giant pulse.

When the full system is installed, it will be possible to make yet wider observations. In particular, with the availability of four independently tuned receivers at DSS-28, one could simultaneously monitor, for example, 3–5 GHz, 5–7 GHz, 9–11 GHz, and 12–14 GHz to study individual pulses over the widest possible frequency range. Alternatively, continuous dual-polarization frequency coverage over any 4 GHz range will also be possible.

Finally, to improve RFI rejection it will be useful to provide a gating mechanism for the dedispersed trigger, so that only triggers which occur in the specified portion of the pulse period (for example during the main or interpulse regions) are considered valid. To achieve this, the TEMPO2 pulsar timing program can be used to predict the topocentric pulsar period in real time (Hobbs et al., 2006). A programmable counter can then be added to the FPGA design with the period determined by a run-time programmable look-up table. The control computer will then update the look-up table with the value obtained from TEMPO2. A set of comparators can then be used to define a number of gates based on the value of the period counter. Such a system, based on a numerically controlled oscillator, is already in use by Hankins and Eilek (2007).

4.4 General Purpose Multi-Frequency Pulsar Observations at DSS-28

Aside from the extremely wide bandwidth transient studies already demonstrated, DSS-28 promises to be a very useful instrument for other kinds of pulsar studies. For example, to improve the precision of pulsar timing observations, it is necessary to measure the minute changes in dispersion measure that occur as the ISM changes along the line of sight to the pulsar (You et al., 2007). By measuring the dispersion sweep over more frequencies, a better estimate of the dispersion measure



Figure 4.20: Frequency versus phase diagram of B1133+16



Figure 4.21: Typical bandpass in the data set used to create figure 4.22, showing large amounts of RFI



Figure 4.22: Dynamic spectrum of B1133+16 around 1.4 GHz

can be obtained. In addition, pulsar timing has traditionally used a single profile template which is cross correlated with the observed data to determine the time of arrival of the pulse. However, most pulsars exhibit changes in their profile shape versus frequency. Therefore, it should be possible to obtain more precise arrival times by measuring pulses across as wide a bandwidth as possible, and correlating them against a frequency dependent profile template. These frequency dependent profile measurements may also be useful for gaining insight into the pulsar emission mechanism.

A second important area where the wide bandwidth of DSS-28 is useful is studying the interstellar medium. Dynamic spectra of pulsar observations have provided some of the most detailed information about the ISM (see, for example, Stinebring et al., 2001), as well as providing a means to infer the transverse velocity of the pulsar (Cordes, 1986). So far these measurements have been limited to bandwidths of approximately 10 MHz at 430 MHz and less than 300 MHz at 1400 MHz.

As a preliminary demonstration, observations were made of the strong pulsar, B1133+16. This pulsar has a period of ~ 1.19 seconds. The observations were made with an iBOB based polyphase filterbank spectrometer with 1024 channels spanning 1024 MHz. The minimum integration time was limited to 8 ms by the bandwidth of the iBOB 100 MbE link (which in practice is limited to 7 Mb/s). The low DM of the pulsar makes the dispersion smearing time across the band negligible compared to the 8 ms sampling time. The 1024 channels of the spectrometer were independently folded at the pulsar period into sub-integrations of 10 pulses, corresponding to 11 seconds. For the first observation, only the lower sideband of polarization Y was sampled, covering approximately 900 to 1900 MHz. Figure 4.20 shows a plot of intensity versus pulsar phase and sky frequency, made by averaging all of the sub-integrations of this observation together. The frequency channels have been binned to 4 MHz resolution. The pulsar appears as a vertical stripe. The spectrum of the on pulse region (between the red arrows) and an off pulse region (between the blue arrows) is shown in the left subplot. If the measurement were calibrated to account for variations in the gain and efficiency of the system versus frequency, this would provide a direct look at the spectral index of the pulsar. As can be readily seen, there is significant RFI in this band. Figure 4.21 shows the average bandpass on a log scale to illustrate that these measurements were made in the presence of RFI 30 dB above the noise floor.

The dynamic spectrum from this observation is shown in figure 4.22. To differentiate the effects due to the pulsar and those caused by interference, the top portion of the figure shows the dynamic spectrum formed using the on pulse region as defined in figure 4.20, while the bottom portion shows the dynamic spectrum of the off pulse region. While RFI dominates below 1200 MHz and above 1800 MHz, relatively straight forward excision algorithms, perhaps including the spectral kurtosis method discussed in section 3.4, should provide more than 500 MHz of usable bandwidth. This observation should be compared to observations presented by Wang et al. (2008), which were made with a 25 meter telescope with a bandwidth of 320 MHz at 1540 MHz.

Observations of B1133+16 were also carried out using the upper and lower sidebands of polarization Y. In this case, each sideband fed a separate iBOB spectrometer, each with the same specifications described previously. The observation spanned 2136 to 4184 MHz, as shown in figure 4.23. The pulsar is clearly visible across the band, and may show some hint of profile evolution, though the time resolution of this spectrometer is too coarse to say for sure. The RFI situation is significantly better than at 1.4 GHz, but is still problematic.

The dynamic spectrum from this observation, shown in figure 4.24, is less impressive than at 1.4 GHz. The cyan arrow points to one scintle from the pulsar. The scintillation bandwidth and timescale both scale strongly with frequency, so we see fewer scintles across the band. Less time was spent on this observation, so the variability is less apparent. Still, the results are promising.

While pulsars generally have a steep spectrum, with flux proportional to f^{α} where α is generally around -1.5, many have been successfully observed at 4.5 GHz and higher (Lorimer and Kramer, 2004). The relatively small collecting area of DSS-28 will limit observations to the stronger pulsars, but this subset should still prove to be interesting.

These preliminary observations demonstrate the feasibility of wide bandwidth, general purpose pulsar observations with DSS-28. Future pulsar observations will require more capable spectrometers than the 1024 channel iBOB-based design used for these observations. In particular, finer frequency resolution will be required for observations of pulsars with higher dispersion measures at lower sky frequencies. The dump time of the spectrometer is currently the biggest limitation. Other groups have demonstrated more appropriate instruments based around CASPER hardware, in particular the GUPPi instrument developed and deployed at NRAO, which provides up to 8192 channels across 800MHz, and dump times as short as 64 μ s (DuPlain et al., 2008). It will be relatively straightforward to implement this design at DSS-28 because of the common digital hardware platform.

Another goal is to implement a 'coherent filterbank' type pulsar processor similar to the ASP, GASP, and NBPP machines described by Demorest et al. (2004). In this design, a polyphase filterbank is used to divide the observing band into several narrow channels. The raw complex time series from each channel is transmitted to a node in a computer cluster which coherently dedisperses the narrow bandwidth data and folds the detected data modulo the pulsar period. In their original implementation, each node could only process a single 4 MHz channel. With the rapid pace of technology development, a modified system has been demonstrated that uses just two computers each with two GPUs to process two polarizations at 128 MHz (Cognard, 2009). Similar servers are already available at DSS-28 in anticipation of this development.



Figure 4.23: Frequency versus phase diagram of B1133+16 at 3 GHz. The on pulse and off pulse regions used to compute the dynamic spectrum in figure 4.24 are indicated between the red and blue arrows, respectively. The on and off pulse average spectra are shown in the plot on the right.



Figure 4.24: Wide bandwidth dynamic spectrum of B1133+16 around 3 GHz. The cyan arrow indicates an enhancement in flux received from the pulsar by interstellar scintillation.
Chapter 5

Ongoing and Future Work

5.1 A Joint Study of Variability of the Crab Pulsar in Conjunction with the Fermi/GLAST Gamma Ray Satellite

One of the most remarkable aspects of the Crab Pulsar is that its detectable emission spans almost the entire electromagnetic spectrum, from the longest radio wavelengths to the highest energy gamma rays. The radio emission is inferred to be coherent because the short duration and high intensity of the radiation imply such high brightness temperatures $(> 10^{35} \text{ K})$ that any incoherent process is ruled out (Hankins and Eilek, 2007). On the other hand, the infrared through gamma ray emission is caused by an incoherent process. The alignment of the pulsed emission across the spectrum implies that both emission processes are occurring within approximately the same region. Evidence exists that the emission processes do interact. For example, Shearer et al. (2003) detected a 3% enhancement in optical emission from the Crab Pulsar coincident with giant pulse emission. Studies have been performed to look for correlation between radio variability (including giant pulses) and X-ray or gamma ray emission, but limited sensitivity produced upper bounds that, while consistant with the optical enhancement detection, are not conclusive (Patt et al., 1999). The most detailed study of radio to gamma ray correlation to date is that of Lundgren et al. (1995). The limited radio bandwidth, observing time, and gamma ray sensitivity of their study allowed them to place an upper limit of 250% on gamma ray enhancement during giant pulses, which is far above the level of enhancement detected in the optical study. Various models of pulsar emission can be distinguished by their prediction of the presence or absence of modulation between the radio and higher energy emission. Coincident gamma ray and giant pulse emission has been predicted by a model proposed by Lyutikov (2007), in which giant pulses are emitted in the outer magnetosphere of the pulsar.

We have proposed a joint variability study of the Crab Pulsar with radio data from DSS-28 in conjunction with the gamma ray telescope, *Fermi*/GLAST. As we have seen from the preliminary observations presented herein, giant pulses from the Crab Pulsar exhibit enormous variability and a wide range of morphologies. While studying correlation between giant pulse and gamma ray emission is the primary goal of this project, the giant pulse deep capture buffer system will be run in parallel, capturing giant pulses across at least 2 GHz of bandwidth at various center frequencies. In the 500+ hours of observing per year, we will thus build up an enormously valuable database of high time resolution, wide bandwidth giant pulse observations. This will allow detailed studies of the array of pulse morphologies and finally provide conclusive information on the bandwidth, time duration, and energy distribution of the pulses. We will be able to statistically separate effects due to the interstellar medium from properties intrinsic to the pulsar.

5.2 Hardware for Real-Time Searches for Radio Transients with Unknown Dispersion Measure

With the growing interest in transient radio astronomy there is a strong push to develop sensitive new instrumentation to detect and observe transient events. One crucially important aspect of transient astronomy is that there is only one opportunity to observe a given transient. Because of this limitation, wide bandwidth instruments such as DSS-28 will be of critical importance to fully understanding transient phenomena. A large observing bandwidth provides much greater ability to distinguish the terrestrial or astronomical nature of a single event. As a case in point, the famous pulse discovered by Lorimer et al. (2007) in the Parkes Multi-beam Survey¹ data was recorded with an incoherent filterbank over a bandwidth of only 288 MHz. This tantalizing single event created huge excitement in the astronomy community, but with the limited data available, there is still considerable dispute as to whether it is of astronomical origin or not. If the event had been recorded with several GHz of baseband voltage data, it is likely much more stringent constraints on its origin could be obtained.

As mentioned in chapter 4, it is impractical to continuously record the huge bandwidth available from GRASP, thus the key to successful transient observations is a means to detect the transient event in real time to trigger the deep capture buffer. The incoherently dedispersed trigger presented in this thesis works very well for an object with known dispersion measure, but of course this information is not available for an unknown transient source. Therefore it will be necessary to perform a real time dispersion measure search. Dispersion measure searching is a common technique in surveys for new pulsars, but is generally not used in real time. Traditionally the limitation has been computational power. Taylor (1974) has suggested an efficient algorithm for dispersion measure searching which is well suited to implementation in massively parallel computer architectures such as FPGAs or modern graphics processing units (GPUs). The primary limitation to an FPGA based implementation would be the memory requirement which scales as N^2 for N DM steps. GPUs have

¹http://www.atnf.csiro.au/research/pulsar/pmsurv/

ample memory and I/O bandwidth to implement the algorithm for modest bandwidths of up to 50 MHz in real time. The main limitation to a GPU implementation would be the latency between detecting a dispersed signal and triggering the deep capture buffer. Fortunately, since the DCB provides up to one second of capture depth at full bandwidth, ample headroom is available for this latency.

5.3 Conclusion

Overall, the DSS-28 system presented here represents a major step forward in demonstrating the kinds of wide bandwidth technologies that will be crucial to the success of future large radio astronomy arrays. The wide bandwidth feed and LNA provide an unprecedented combination of low system noise and a 4 to 1 bandwidth. The receiver architecture provides enormous flexibility to support a wide range of observations. For the most part, the components that make up the receiver could be integrated together into a few mass-produced modules suitable for an array application.

The flexible FPGA-based architecture of the digital signal processing system provides an excellent test bed for new spectrometer modes, RFI excision algorithms, and transient detection mechanisms. The CASPER tools provide a migration path for future upgrades as digital technology advances. Additionally, the CASPER hardware is becoming common place at observatories around the world. This makes sharing instrument designs between observatories almost as easy as transferring files.

The effectiveness of a wide bandwidth spectrometer employing the spectral kurtosis algorithm for RFI detection for continuum observations has been demonstrated. This technique is invaluable at observatories which are not protected from strong sources of RFI.

The hardware implementation of a method for digital quadrature imbalance correction demonstrates that quadrature down-conversion is a viable architecture for centimeter wavelength radio astronomy instruments. The image rejection attained by this spectrometer is much better than from the analog hybrid. The low frequency roll-off of the hybrid is also avoided, meaning less of the center of the band is lost.

In addition, the potential to perform unique scientific observations of transient events, in particular giant pulses from the Crab pulsar, has been demonstrated. The special purpose instrument designed for this task overcomes many of the limitations faced by previous studies. In particular, the dedispersed trigger provides much greater sensitivity and immunity to interference than the total power trigger used previously. The full bandwidth capture buffer with ample memory depth provides much greater bandwidth than continuously recorded observations. It also overcomes the duty cycle limitations of the digital oscilloscopes used previously by providing multiple memory buffers to capture bursts of pulses.

The wide bandwidth available with DSS-28 will also be useful for studying pulsar profiles ver-

sus frequency and for dynamic spectra observations of the interstellar medium. A glimpse of this emerging capability has been demonstrated with initial observations.

Bibliography

- D. W. Allan. Time and frequency (time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 34(6):647–654, November 1987.
- Jacob W. M. Baars. The Paraboloidal Reflector Antenna in Radio Astronomy and Communication, volume 348 of Astrophysics and Space Science Library. Springer, 2007.
- J. C. Bardin and S. Weinreb. A 0.5–20 GHz quadrature downconverter. In Proc. IEEE Bipolar/BiCMOS Circuits and Technology Meeting BCTM 2008, pages 186–189, October 13–15, 2008. doi: 10.1109/BIPOL.2008.4662740.
- E. O. Brigham. The fast Fourier transform and its applications. Englewood Cliffs, N.J.: Prentice-Hall, 1988.
- P. M. T. Broersen. A comparison of transfer function estimators. *IEEE Transactions on Instrumen*tation and Measurement, 44(3):657–661, June 1995. doi: 10.1109/19.387302.
- C. Carilli and S. Rawlings, editors. *Science with the Square Kilometre Array*, volume 48 of *New Astronomy Reviews*. Elsevier, 2004.
- C. Chang, J. Wawrzynek, and R. W. Brodersen. BEE2: a high-end reconfigurable computing system. IEEE Design Test of Computers, 22(2):114–125, March 2005. doi: 10.1109/MDT.2005.30.
- F. E. Churchill, G. W. Ogar, and B. J. Thompson. The correction of I and Q errors in a coherent processor. *IEEE Transactions on Aerospace and Electronic Systems*, 1(1):131–137, January 1981. doi: 10.1109/TAES.1981.309045.
- I. Cognard, J. A. Shrauner, J. H. Taylor, and S. E. Thorsett. Giant radio pulses from a millisecond pulsar. ApJ, 457:L81 ff., February 1996. doi: 10.1086/309894.
- Ismael Cognard. BON upgrade: evolution of the cluster dedicated to the coherent pulsar dedispersor. http://lpce.cnrs-orleans.fr/~icognard/science/GPU-report.pdf, June 2009.
- J. J. Condon and S. M. Ransom. Essential Radio Astronomy. http://www.cv.nrao.edu/course/ astr534/ERA.shtml, 2008.

- J. M. Cordes. Space velocities of radio pulsars from interstellar scintillations. ApJ, 311:183–196, December 1986. doi: 10.1086/164764.
- J. M. Cordes, N. D. R. Bhat, T. H. Hankins, M. A. McLaughlin, and J. Kern. The brightest pulses in the universe: Multifrequency observations of the crab pulsar's giant pulses. *ApJ*, 612:375–388, September 2004. doi: 10.1086/422495.
- J. H. Crossley, J. A. Eilek, T. H. Hankins, and J. S. Kern. Short-lived radio bursts from the crab pulsar. *Submitted to ApJ*, 2009.
- P. Demorest, R. Ramachandran, D. Backer, R. Ferdman, I. Stairs, and D. Nice. Precision pulsar timing and gravity waves: Recent advances in instrumentation. In *Bulletin of the American* Astronomical Society, volume 36 of *Bulletin of the American Astronomical Society*, pages 1598–+, December 2004.
- R. DuPlain, S. Ransom, P. Demorest, P. Brandt, J. Ford, and A. L. Shelton. Launching GUPPI: the Green Bank Ultimate Pulsar Processing Instrument. In Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 7019, August 2008. doi: 10.1117/12.790003.
- M. Elitzur. Astronomical masers, volume 170 of Astrophysics and Space Science Library. Springer-Verlag New York, LLC, 1992.
- D. T. Emerson, U. Klein, and C. G. T. Haslam. A multiple beam technique for overcoming atmospheric limitations to single-dish observations of extended radio sources. A&A, 76:92–105, June 1979.
- J. R. Fisher and M. A. Morgan. Analysis of a single-conversion, analog/digital sideband-separating mixer prototype. Electronic Division Internal Report 320, National Radio Astronomy Observatory, June 2008. URL http://www.gb.nrao.edu/electronics/edir/edir320.pdf.
- J. D. Gallego, I. Lopez-Fernandez, C. Diez, and A. Barcia. Experimental results of gain fluctuations and noise in microwave low-noise cryogenic amplifiers. In Francois Danneville, Fabrizio Bonani, M. Jamal Deen, and Michael E. Levinshtein, editors, *Noise in Devices and Circuits II*, volume 5470, pages 402–413. SPIE, 2004. doi: 10.1117/12.547097. URL http://link.aip.org/link/ ?PSI/5470/402/1.
- D. E. Gary, Z. Liu, and G. M. Nita. A wideband spectrometer with RFI detection. In URSI National Radio Science Meeting, Boulder, CO, January 2009.
- T. H. Hankins. Microsecond intensity variations in the radio emissions from CP 0950. ApJ, 169:487 ff., November 1971. doi: 10.1086/151164.

- T. H. Hankins. Giant pulses from the Crab pulsar. In M. Kramer, N. Wex, & R. Wielebinski, editor, IAU Colloq. 177: Pulsar Astronomy—2000 and Beyond, volume 202 of Astronomical Society of the Pacific Conference Series, page 165 ff., 2000.
- T. H. Hankins and J. A. Eilek. Radio emission signatures in the Crab pulsar. ApJ, 670:693–701, November 2007. doi: 10.1086/522362.
- T. H. Hankins and J. M. Rajkowski. Wide bandwidth signal processor for removing dispersion distortion from pulsar radio signals. *Review of Scientific Instruments*, 58:674–680, April 1987. doi: 10.1063/1.1139652.
- T. H. Hankins and B. J. Rickett. Pulsar signal processing. Methods in Computational Physics, 14: 55–129, 1975.
- G. B. Hobbs, R. T. Edwards, and R. N. Manchester. TEMPO2, a new pulsar-timing package—I. An overview. MNRAS, 369:655–672, June 2006. doi: 10.1111/j.1365-2966.2006.10302.x.
- W. A. Imbriale. Design of a wideband radio telescope. In Proc. IEEE Aerospace Conference, pages 1–12, March 3–10, 2007. doi: 10.1109/AERO.2007.352858.
- N. C. Jarosik. Measurements of the low-frequency-gain fluctuations of a 30-ghz high-electronmobility-transistor cryogenic amplifier. *IEEE Transactions on Microwave Theory and Techniques*, 44(2):193–197, February 1996. doi: 10.1109/22.481567.
- F. A. Jenet, G. B. Hobbs, W. van Straten, R. N. Manchester, M. Bailes, J. P. W. Verbiest, R. T. Edwards, A. W. Hotan, J. M. Sarkissian, and S. M. Ord. Upper bounds on the low-frequency stochastic gravitational wave background from pulsar timing observations: Current limits and future prospects. ApJ, 653:1571–1576, December 2006. doi: 10.1086/508702.
- B. Klein, S. D. Philipp, R. Güsten, I. Krämer, and D. Samtleben. A new generation of spectrometers for radio astronomy: fast Fourier transform spectrometer. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6275, July 2006. doi: 10.1117/12.670831.
- M. Kramer and I. H. Stairs. The double pulsar. ARA&A, 46:541–572, September 2008. doi: 10.1146/annurev.astro.46.060407.145247.
- J. D. Kraus. Radio astronomy. Cygnus-Quasar Books, 2 edition, 1986.
- K. Krishnamoorthy. Handbook of statistical distributions with applications. Number 188 in Statistics, a series of textbooks & monographs. Chapman & Hall/CRC Press Taylor & Francis Group, 2006.
- Z. Liu. Kurtosis Spectrometer. http://casper.berkeley.edu/wiki/Kurtosis_Spectrometer, 2008.

- D. R. Lorimer and M. Kramer. Handbook of Pulsar Astronomy. Cambridge University Press, December 2004.
- D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford. A bright millisecond radio burst of extragalactic origin. *Science*, 318:777 ff., November 2007. doi: 10.1126/science. 1147532.
- J. E. J. Lovell, D. L. Jauncey, C. Senkbeil, S. Shabala, H. E. Bignall, T. Pursimo, R. Ojha, J.-P. Macquart, B. J. Rickett, M. Dutka, and L. Kedziora-Chudczer. MASIV: The Microarcsecond Scintillation-Induced Variability Survey. In M. Haverkorn and W. M. Goss, editors, SINS—Small Ionized and Neutral Structures in the Diffuse Interstellar Medium, volume 365 of Astronomical Society of the Pacific Conference Series, page 279 ff., July 2007.
- S. C. Lundgren, J. M. Cordes, M. Ulmer, S. M. Matz, S. Lomatch, R. S. Foster, and T. Hankins. Giant pulses from the Crab pulsar: a joint radio and gamma-ray study. *ApJ*, 453:433 ff., November 1995. doi: 10.1086/176404.
- A. G. Lyne and F. Graham-Smith. *Pulsar astronomy*. Cambridge University Press, 3rd edition, 2006.
- A. G. Lyne, R. S. Pritchard, and F. Graham-Smith. Twenty-three years of Crab pulsar rotational history. MNRAS, 265:1003 ff., December 1993.
- M. Lyutikov. On generation of Crab giant pulses. MNRAS, 381:1190–1196, November 2007. doi: 10.1111/j.1365-2966.2007.12318.x.
- R. N. Manchester and J. H. Taylor. Parameters of 61 pulsars. Astrophys. Lett., 10:67 ff., January 1972.
- M. A. McLaughlin and J. M. Cordes. Searches for giant pulses from extragalactic pulsars. ApJ, 596: 982–996, October 2003. doi: 10.1086/378232.
- K. M. Menten. The discovery of a new, very strong, and widespread interstellar methanol maser line. ApJ, 380:L75–L78, October 1991. doi: 10.1086/186177.
- V. Minier, J. E. Conway, and R. S. Booth. VLBI observations of 6.7 and 12.2 GHz methanol masers toward high mass star-forming regions. II. Tracing massive protostars. A&A, 369:278–290, April 2001. doi: 10.1051/0004-6361:20010124.
- D. A. Moffett and T. H. Hankins. Multifrequency radio observations of the Crab pulsar. ApJ, 468: 779 ff., September 1996. doi: 10.1086/177734.

- M. A. Morgan and J. R. Fisher. Next generation radio astronomy receiver systems. ArXiv e-prints, August 2009. URL http://arxiv.org/abs/0908.3849.
- G. M. Nita, D. E. Gary, Z. Liu, G. J. Hurford, and S. M. White. Radio frequency interference excision using spectral-domain statistics. *PASP*, 119:805–827, July 2007. doi: 10.1086/520938.
- V. Ossenkopf. The stability of spectroscopic instruments: a unified Allan variance computation scheme. A&A, 479:915–926, March 2008. doi: 10.1051/0004-6361:20079188.
- M. Ott, A. Witzel, A. Quirrenbach, T. P. Krichbaum, K. J. Standke, C. J. Schalinski, and C. A. Hummel. An updated list of radio flux density calibrators. A&A, 284:331–339, April 1994.
- A. Parsons, D. Backer, C. Chang, D. Chapman, H. Chen, P. Crescini, C. de Jesus, C. Dick, P. Droz, D. MacMahon, K. Meder, J. Mock, V. Nagpal, B. Nikolic, A. Parsa, B. Richards, A. Siemion, J. Wawrzynek, D. Werthimer, and M. Wright. Petaop/second FPGA signal processing for SETI and radio astronomy. In *Proc. Fortieth Asilomar Conference on Signals, Systems and Computers* ACSSC '06, pages 2031–2035, October 2006. doi: 10.1109/ACSSC.2006.355123.
- A. Parsons, D. Backer, A. Siemion, H. Chen, D. Werthimer, P. Droz, T. Filiba, J. Manley, P. McMahon, A. Parsa, D. MacMahon, and M. Wright. A Scalable Correlator Architecture Based on Modular FPGA Hardware, Reuseable Gateware, and Data Packetization. *PASP*, 120:1207–1221, November 2008. doi: 10.1086/593053.
- B. L. Patt, M. P. Ulmer, W. Zhang, J. M. Cordes, and Z. Arzoumanian. A search for pulse-to-pulse variability of the crab pulsar. ApJ, 522:440–443, September 1999. doi: 10.1086/307620.
- R. Pintelon and J. Schoukens. System identification: a frequency domain approach. IEEE Press, 2001.
- M. V. Popov and B. Stappers. Statistical properties of giant pulses from the Crab pulsar. A&A, 470:1003–1007, August 2007. doi: 10.1051/0004-6361:20066589.
- M. V. Popov, N. Bartel, W. H. Cannon, A. Y. Novikov, V. I. Kondratiev, and V. I. Altunin. Pulsar microstructure and its quasi-periodicities with the S2 VLBI system at a resolution of 62.5 nanoseconds. A&A, 396:171–187, December 2002. doi: 10.1051/0004-6361:20021402.
- M. V. Popov, A. D. Kuz'min, O. M. Ul'yanov, A. A. Deshpande, A. A. Ershov, V. V. Zakharenko, V. I. Kondrat'ev, S. V. Kostyuk, B. Y. Losovskiä, and V. A. Soglasnov. Instantaneous radio spectra of giant pulses from the crab pulsar from decimeter to decameter wavelengths. *Astronomy Reports*, 50:562–568, July 2006. doi: 10.1134/S1063772906070067.
- K.-P. Pun, J. E. Franca, and C. Azeredo-Leme. Wideband digital correction of I and Q mismatch in quadrature radio receivers. In *Proc. ISCAS 2000 Geneva The 2000 IEEE International Symposium*

on Circuits and Systems, volume 5, pages 661–664, May 28–31, 2000. doi: 10.1109/ISCAS.2000. 857556.

- V. Rodriguez. An open-boundary quad-ridged guide horn antenna for use as a source in antenna pattern measurement anechoic chambers. *IEEE Antennas and Propagation Magazine*, 48(2):157– 160, April 2006. doi: 10.1109/MAP.2006.1650857.
- K. Rohlfs and T. L. Wilson. Tools of Radio Astronomy. Springer Verlag, 2nd edition, 1996.
- J. Ruze. Antenna tolerance theory: A review. Proceedings of the IEEE, 54(4):633-640, April 1966.
- A. Shearer, B. Stappers, P. O'Connor, A. Golden, R. Strom, M. Redfern, and O. Ryan. Enhanced optical emission during crab giant radio pulses. *Science*, 301:493–495, July 2003. doi: 10.1126/ science.1084919.
- H. K. So and R. W. Brodersen. BORPH: An operating system for FPGA-based reconfigurable computers. Technical Report EECS-2007-92, University of California Berkeley, July 2007. URL http://techreports.lib.berkeley.edu/accessPages/EECS-2007-92.html.
- V. A. Soglasnov, M. V. Popov, N. Bartel, W. Cannon, A. Y. Novikov, V. I. Kondratiev, and V. I. Altunin. Giant pulses from PSR B1937+21 with widths ≤ 15 nanoseconds and $T_b \geq 5 \times 10^{39}$ K, the highest brightness temperature observed in the universe. *ApJ*, 616:439–451, November 2004. doi: 10.1086/424908.
- D. H. Staelin and E. C. Reifenstein, III. Pulsating radio sources near the crab nebula. Science, 162: 1481–1483, December 1968.
- D. R. Stinebring, M. A. McLaughlin, J. M. Cordes, K. M. Becker, J. E. E. Goodman, M. A. Kramer, J. L. Sheckard, and C. T. Smith. Faint scattering around pulsars: Probing the interstellar medium on solar system size scales. *ApJ*, 549:L97–L100, March 2001. doi: 10.1086/319133.
- J. H. Taylor. A sensitive method for detecting dispersed radio emission. A&AS, 15:367 ff., June 1974.
- P.P. Vaidyanathan. Multirate systems and filter banks. Prentice Hall PTR, 1992.
- N. Wang, Z. Yan, R. N. Manchester, and H. X. Wang. Daily observations of interstellar scintillation in PSR B0329+54. MNRAS, 385:1393–1401, April 2008. doi: 10.1111/j.1365-2966.2008.12864.x.
- R. Weber, C. Viou, A. Coffre, L. Denis, P. Zarka, and A. Lecacheux. DSP-enabled radio astronomy: towards IIIZW35 reconquest. *EURASIP J. Appl. Signal Process.*, 2005:2686–2693, 2005. ISSN 1110-8657. doi: http://dx.doi.org/10.1155/ASP.2005.2686.

- S. Weinreb. A digital spectral analysis technique and its application to radio astronomy. Technical Report 412, MIT Research Laboratory of Electronics, 1963. URL http://hdl.handle.net/1721. 1/4413.
- S. Weinreb, J. C. Bardin, and H. Mani. Design of cryogenic SiGe low-noise amplifiers. *IEEE Transactions on Microwave Theory and Techniques*, 55(11):2306–2312, November 2007. doi: 10. 1109/TMTT.2007.907729.
- S. Weinreb, J. Bardin, H. Mani, and G. Jones. Matched wideband low-noise amplifiers for radio astronomy. *Review of Scientific Instruments*, 80(4):044702 ff., April 2009. doi: 10.1063/1.3103939.
- J. Welch, D. Backer, L. Blitz, D. Bock, G. Bower, C. Cheng, S. Croft, M. Dexter, G. Engargiola, E. Fields, J. Forster, C. Gutierrez-Kraybill, C. Heiles, T. Helfer, S. Jorgensen, G. Keating, J. Lugten, D. MacMahon, O. Milgrome, D. Thornton, L. Urry, J. van Leeuwen, D. Werthimer, P. Williams, M. Wright, J. Tarter, R. Ackermann, S. Atkinson, P. Backus, W. Barott, T. Bradford, M. Davis, D. DeBoer, J. Dreher, G. Harp, J. Jordan, T. Kilsdonk, T. Pierson, K. Randall, J. Ross, S. Shostak, M. Fleming, C. Cork, A. Vitouchkine, N. Wadefalk, and S. Weinreb. The Allen Telescope Array: The first widefield, panchromatic, snapshot radio camera for radio astronomy and SETI. In J. Baars, R. Thompson, and L. D'Addario, editors, Special Issue of Proceedings of the IEEE: "Advances in Radio Telescopes", April 2009. URL http://arxiv.org/ftp/arxiv/papers/0904/0904.0762.pdf.
- A. Whitney. Mark 5C VLBI data system. In Proceedings of the 9th European VLBI Network Symposium on the role of VLBI in the Golden Age for Radio Astronomy and EVN Users Meeting. September 23-26, 2008. Bologna, Italy., 2008. URL http://pos.sissa.it/cgi-bin/reader/ conf.cgi?confid=72.
- X. P. You, G. Hobbs, W. A. Coles, R. N. Manchester, R. Edwards, M. Bailes, J. Sarkissian, J. P. W. Verbiest, W. van Straten, A. Hotan, S. Ord, F. Jenet, N. D. R. Bhat, and A. Teoh. Dispersion measure variations and their effect on precision pulsar timing. *MNRAS*, 378:493–506, June 2007. doi: 10.1111/j.1365-2966.2007.11617.x.