University of Texas Rio Grande Valley

ScholarWorks @ UTRGV

4-2014

A digital backend for the low frequency all sky monitor

Louis Percy Dartez
The University of Texas Rio Grande Valley

Follow this and additional works at: https://scholarworks.utrgv.edu/leg_etd

Part of the Cosmology, Relativity, and Gravity Commons, and the Instrumentation Commons

A Digital Backend for the Low Frequency All Sky Monitor

by

Louis Percy Dartez

A Thesis Presented to the Graduate Faculty of the College of Science, Mathematics, and Technology in

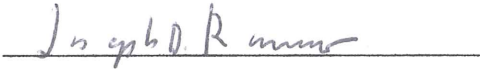Partial Fulfillment

of the Requirements for the Degree of

Master of Science

in the Field of Physics

Approved by:
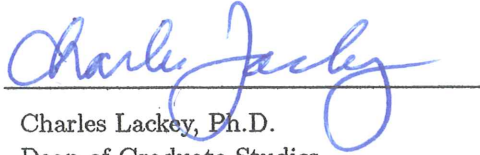
Teviet Creighton, Ph.D.
Thesis Director

Joe Romano, Ph.D.

Volker Quetschke, Ph.D.

Charles Lackey, Ph.D.
Dean of Graduate Studies

Graduate School

University of Texas at Brownsville

April 2014

A Digital Backend for the Low Frequency All Sky Monitor

by

Louis Percy Dartez

Teviet Creighton, Advisor

A thesis submitted in partial fulfillment
of the requirements for the
Degree of Master of Science
in the Field of
Physics

University of Texas at Brownsville

Brownsville, Texas

April 2014

Acknowledgements

My work with the Low Frequency All Sky Monitor (LoFASM) would not have been possible if it weren't for all of the guidance that I have received over the past few years. First of all, I acquired much of my digital design knowledge and experience while working with Luis Quintero on a novel radar system during a summer residency at Arecibo Observatory, and for that I will always be thankful. Many thanks go directly to Dan Werthimer of SETI at UC-Berkeley for taking me in at the Berkeley Wireless Research Laboratory and helping me piece together the LoFASM Correlator in his electronics lab. Our numerous mini-brainstorming sessions were invaluable to the design of LoFAM's correlator implementation. I would also like to thank LoFASM collaborators at the University of New Mexico, the Green Bank Observatory, and the Goldstone Deep Space Science Network, and the Jet Propulsion Laboratory for all of their help in setting up the various LoFASM stations and electronics at each installation site. This project would not have been possible without the tremendous amount of help I have received along the way from my mentors at home. Thanks go to Dr. Teviet Creighton for his indispensible guidance in completing this thesis. I am eternally grateful for the privilege of having Dr. Fredrick Jenet as my mentor, advisor, and teacher. Essentially all of what I have learned of radio astronomy, signal processing, and radio frequency instrumentation can be directly attributed to the mentorship of Dr. Jenet and Dr. Creighton. Furthermore, I would like to thank Dr. Jenet for the fountain of opportunity that he has brought forth in my academic career. It has been my honor to be your pupil and I very much look forward to continuing to work together on new and exciting projects for the foreseeable future.

Abstract

The Low Frequency All Sky Monitor (LoFASM) is a distributed array of dipole antennas that are sensitive to radio frequencies from 10 to 88 MHz. The primary science goals of LoFASM are the detection and study of low-frequency radio transients, a high priority science goal as deemed by the National Research Council's decadal survey. LoFASM consists of antennas and front-end electronics that were originally developed for the Long Wavelength Array (LWA) by the U.S. Naval Research Lab, the University of New Mexico, Virginia Tech, and the Jet Propulsion Laboratory. LoFASM, funded by the U.S. Department of Defense, will initially consist of four stations, each consisting of 12 dual-polarization dipole antennas.

In a single station, RF signals from each of the individual LoFASM dipoles are combined in phase in order to synthesize LoFASM's beam. The LoFASM RF signals are phased up so that the resulting beam is sensitive to radio emission that originates from the zenith and RF signals approaching from the horizon are attenuated. Digitally, this is achieved using a full Stokes 100MHz correlating spectrometer constructed using field programmable gate array (FPGA) technology. In this thesis I will describe the design and usage of the LoFASM Correlator.

Table of Contents

List of Figures

1   Introduction

1.1   A Brief History of Radio Astronomy

Radio astronomy is a field within astronomy that focuses on studying celestial bodies throughout the known universe by analyzing extraterrestrial emissions at radio frequencies (typically 3 Hz – 300GHz). Radio waves from extraterrestrial origins were first detected by Karl G. Jansky (known as the father of radio astronomy) at ~20 MHz in the early 1930's while working at Bell Telephone Laboratories, see Figure 1 [14]. A few years later Grote Reber, a radio engineer, was able to take Jansky's work and make improvements that allowed him to carry out measurements at 160 MHz (see Figure 2). Since these initial detections were made, many scientists from all over the world have taken the work of Jansky and Reber and expanded upon it. Over the next 80 years, the development of radio telescopes has leaned towards acheiving greater angular resolution and higher sensitiviy, typically at shorter wavelengths (higher frequencies).

Historically, most radio telescopes have been designed to use a single receiver and dish aimed towards specific points of interest in the sky. Since the sensitivity of a radio telescope is directly related to the amount of waves captured, most radio telescopes are built with as large a dish as possible. At the time of this writing the most sensitive operational radio telescope on the planet is the Arecibo Telescope located in Arecibo, Puerto Rico (see Figure 3). The Arecibo dish has a diameter of 305 m, and though it was built in the 1950's, it remains among the most-used research telescopes in the world today.

Building gargantuan structures with huge metal dishes like the Arecibo Telescope can be quite expensive. In addition to the engineering hurdles associated with such a project, one also had to consider the computational resources needed to record and analyze the radio telescope's output. Most early instruments used paper chart recorders to collect information about the amount of power or intensity in an antenna's signal. However, mechanical paper recorders are slow, difficult to maintain, and require constant attention. In addition, because of the immense costs associated with manufacturing larger and larger dishes, it soon became apparent that simply building gargantuan Arecibos would be impractical.

As general purpose computers became faster, more reliable, and easier to program, radio telescopes began to incorporate digitization techniques that allowed scientists to obtain information about radio emission at

a higher temporal rate. In addition, the need to calculate a signal's power (an irreversible process through which information about the phase of the signal is lost) in real time no longer existed, although it is still done today because of practicality. Many radio telescopes today make use of a backend system that provides the option to record full temporal (time series) voltage information, which allows astronomers to retain all of the radio signal's phase information.

The increasingly low cost of computers allowed the radio astronomy community to take on a different (yet old!) approach of obtaining higher angular resolution. Instead of using a single large reflector, 21st century radio astronomers began using arrays consisting of numerous individual antennas (either dishes or dipoles) that were geographically separated by great distances, resulting in much larger baselines than could possibly be acheived by a single element.[1] These early ground-breaking aperture arrays include ASTRON's LOw Frequency ARray (LOFAR), the Murchison Widefield Array (MWA), the University of New Mexico's Long Wavelength Array (LWA), and the Precision Array for Probing the Epoch of Re-ionization (PAPER).

Major advances in signal processing, high-speed networking, and high-peformance computing systems led to huge increases in the speed and reliabilty of aperture arrays [10]. With absolutely no moving parts (except maybe the data recording equipment, e.g. magnetic tapes, hard drive disks, etc.), aperture arrays can be steered electronically by digitally manipulating the transmission delay of any number of dipoles within the array in real time; there is no need to physically rotate or point a massive heavy structure. Since the signal from each array element is processed independently, the number of simultaneously pointings (or beams) that a telescope can handle depends only the on the computational prowess of the instrument's backend (i.e., the combined performace of the data-taking computer and the ability of the networking scheme that connects all of the array's elements). Thus, most modern aperture arrays take advantage of the fact that they can track multiple sources in the sky at once, maximizing their efficiency.

---

[1]It is probably worth noting here that aperture synthesis, in its early form, has been around since the 1940's during WWII. However, after the end of World War II, most radio engineers went back to their civilian lives. Those who stayed working in the infant field of radio astonomy mostly turned to using large parabolic (or, in the case of Arecibo, spherical) dishes. These were more suited to the study of atomic and molecular emisions at higher frequencies than could be achieved by aperture arrays consisting of large dipoles. It is only recently that radio astronomy has begun to transform itself into focusing on interferometric arrays. For a more detailed discussion on the fall and rise of the aperture array, see [10].

1.2 Overview

The number of radio interferometers using dipole based technologies to study the radio sky at low frequencies (<300 MHz) is growing [3]. Two such telescopes in use today are LOFAR, which operates within the radio window 10–240MHz, and the LWA, which operates from 10 to 88 MHz [8, 12]. While there have been a few studies in the largely unexplored sub–100 MHz region, most of our knowledge of the radio emission mechanisms of pulsars and other transient sources originates from radio surveys conducted at frequencies above 300 MHz. This may be attributed to four hurdles that one typically has to deal with when studying radio pulses of extraterrestrial origins (see Chapter 4): (i) the deleterious effects of the interstellar medium (ISM) on pulsed radio signals, (ii) the effective background sky temperature of the Galaxy due to synchrotron emission, (iii) Earth's ionospheric effects, and (iv) the unmanageably high levels of human-generated radio frequency interference (RFI) below approximately 25 MHz (FM ranges from about 88–107.9 MHz) [8, 12]. The first three of these effects depend on steep power laws in frequency that cause the effects to worsen at lower frequencies [23]. However, despite all of the challenges associated with observing at low frequencies, there are numerous reasons why it is important and interesting to search for and study radio transients in a much lower frequency regime than what is more commonly used today; these are discussed in more detail in Chapter 2.

The Low Frequency All Sky Monitor (LoFASM) was designed, constructed, and installed by the University of Texas at Brownsville's Center for Advanced Radio Astronomy (CARA) and has four independent stations across the continental United States in facilities that are owned and operated by various parties. The first of these stations, LoFASM I, located in Port Mansfield, Texas, is the LoFASM prototype station. The prototype station's main purpose is to provide a testbench for designing, testing, and characterizing LoFASM instrumentation. LoFASM II, III, and IV are located, respectively, at the North Arm of the LWA1 Radio Observatory outside of Socorro, New Mexico; outside of the Interferometer Control Building at the Green Bank Observatory in Green Bank, West Virgina; and at the Venus site of the Goldstone Deep Space Communications Complex outside of Barstow, California. LoFASM has the potential to greatly improve our understanding of the origins of radio transients as well as their emission mechanisms by providing a large

field of view and enhanced localization techniques that will better determine the source of interesting radio emissions. In addition, the LoFASM stations will provide the perfect platform for performing tests on novel RFI mitigation techniques at low frequencies (see Chapter 4).

The backend data acquisition system for the LoFASM array consists of a digital 4-input correlator implemented using Field Programmable Gate Array (FPGA) technology. The LoFASM Correlator, which runs on a Reconfigurable Open Architecture Computing Hardware (ROACH) board, calculates and averages the spectral cross-power densities from the four independent LoFASM signals and sends them to a nearby control computer for analysis. The correlator data is used to synthesize a $40°$ principal beam centered at zenith.

The topic of this thesis is the design and implementation of the LoFASM Correlator as the inaugural digital backend system for LoFASM. The basic design parameters and considerations are discussed in Chapter 3. The mitigation of problems due to local RFI and other complications that can arise when observing at low frequencies is presented in Chapter 4 on page 20. LoFASM's beam synthesis techniques are covered in Chapter 5. The LoFASM Correlator firmware (low level code that is executed on the FPGA chip) was designed and compiled using a suite of open source software tools made available by the University of California at Berkeley's Wireless Research Center. The use of these tools to produce the digital logic for LoFASM and their implementation are presented in Chapter 6. Chapter 7 describes how to use the LoFASM Correlator control software and contains detailed information about the native LoFASM data format that is used internally within the LoFASM network.
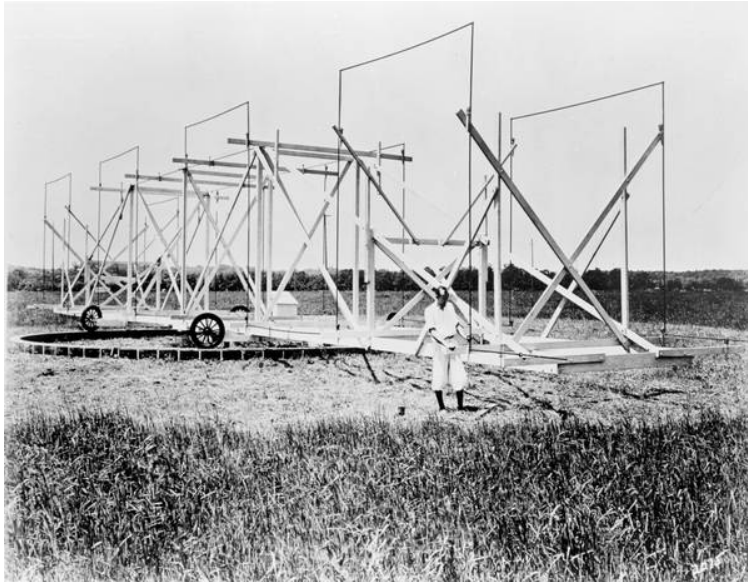
Figure 1: Karl G. Janksy standing next to his 20 MHz radio telescope.

Figure 2: Grote Reber standing next to the Reber telescope.

Figure 3: Aerial view of Arecibo Observatory

2 Science Goals

LoFASM will be dedicated to continuous, long–term drift observations of the radio sky in the low-frequency band. LoFASM's primary scientific goals will be to monitor and study astrophysical radio transient events, and to develop, test, and characterize novel techniques to mitigate strong RFI at low frequencies. Of particular interest are transient events due to coalescences of neutron star-neutron star (NS-NS) binaries and neutron star-black hole (NS-BH) binaries. These unique and ellusive celestial systems are good candidates for being the origins of huge cosmological bursts of energy that span the entire electromagnetic spectrum. A coincident (triggered) detection of a transient event with either a gravitational observeratory (e.g., the Laser Interferometer Gravitational wave Observatory [1]) or an observatory operating at a different frequency in the EM spectrum (e.g., the SWIFT satellite [9]) could significantly increase the chances of a true gravitational-wave detection. In this way LoFASM could play an important role in the race for the first detection of gravitational waves.

Gravitational waves are minute fluctuations in the curvature of spacetime that are caused by the accelerated motion of very massive objects (e.g., binary systems of a NS-NS or NS-BH). Gravitational waves were originally predicted by Einstein in his general theory of relativity in 1918. In 1974, Hulse & Taylor found evidence for gravitational waves while observing PSR1913+16 [13]. Gravitational waves will provide us with detailed information about the bulk, coherent motion of matter, which complements the electromagnetic spectrum. The successful detection of gravitational waves will open new doors in the field of astronomy. A successful gravitational-wave detection would result in a transformation (or evolution) in the way we currently understand the known universe [2].

In addition to aiding with the search for gravitational-wave bursts and other transients in the sky, part of the aim of LoFASM is to facilitate the development, design, and implementation of novel RFI mitigation techniques. The arrangement of the LoFASM antennas into a tight dual-ring configuration originates from one such consideration (see Chapter 4.1).

3 LoFASM Design

LoFASM is a phased array of dipole antennas that are sensitive to radio frequencies from roughly 10–88MHz. The lower limit of 10 MHz stems from the effects of Earth's ionosphere at low frequencies (see Chapter 4.5). The upper limit of 88 MHz is mostly due to the fact that local FM stations, which generally broadcast from 88–107 MHz, are too loud for the study of faint extraterrestrial radio emissions. The LoFASM dipole antennas were originally designed by the U.S. Naval Research Laboratory (NRL) and Virginia Tech to be used by the LWA1 in New Mexico. Ellingson et al. [8] and references therein demonstrate that Galactic noise-limited sensitivity of the radio sky is achievable in most of the 10–88 MHz range by using low-gain dipoles that are matched with receivers having a system temperature of less than ≈ 500 K.

Each LoFASM station is composed of 12 antennas that are grouped into a close-packed configuration consisting of two concentric rings, each containing six dipole stands (see Figure 4). As is discussed further in Chapter 4.1, when the antennas are appropriately phased the tight double–ring antenna configuration helps attenuate unwanted radio signals originating from the horizon. Figure 5 shows the 12 LoFASM I antennas after their installation. All of the LoFASM antennas are oriented in such a way that one orthogonal polarization is aligned East-West and the other is aligned North-South. The signals for each polarization of each ring (a total of 4 independent signals) are handled separately and contain no information about each other until they are combined by the digital backend. However, all of the signals from the same ring of antennas of like-polarization are summed in phase before reaching the analog receiving systems. The four signals entering the analog receiving system correspond to the summed North-South and East-West polarized signals from each of the two LoFASM rings. The signals are summed using a set of in-phase DC-passing 6-to-1 RF combiners (see Figure 6).

After being summed in phase, the four remaining LoFASM signals continue propagating along ~1,000 feet of low-loss LMR-400 coaxial cable that runs to an electronics warehouse, where the analog receivers (ARX) treat each of the four channels individually. The current version of the LoFASM ARX is a modular system called the mARX in which each analog component (e.g., attenuator, RF filter, or low noise amplifier) is kept on its own printed circuit board (PCB). This maintains a manageable level of modularity in order to

facilitate the tracking and eradication of bugs in the mARX[2]. As depicted below in Figure 7, the mARX is where analog adjustments are imposed on the LoFASM signals. These include various stages of RF filters, two stages of signal amplification, and a stage of signal attenuation. The LoFASM RF signal path ends at the pair of analog to digital converters (ADCs) located on the ROACH board (see Figure 7, right) on which the LoFASM Correlator firmware is hosted.

---

[2]The original version of the LoFASM ARX was a complex system containing all of the analog components for all four LoFASM signals on a single PCB that was digitally controlled using C++ software and an Arduino board. It turned out that the added flexibility of being able to digitally control the ARX was simply not worth the extensive amount of effort needed to find and fix issues that frequently arose with the ARX.
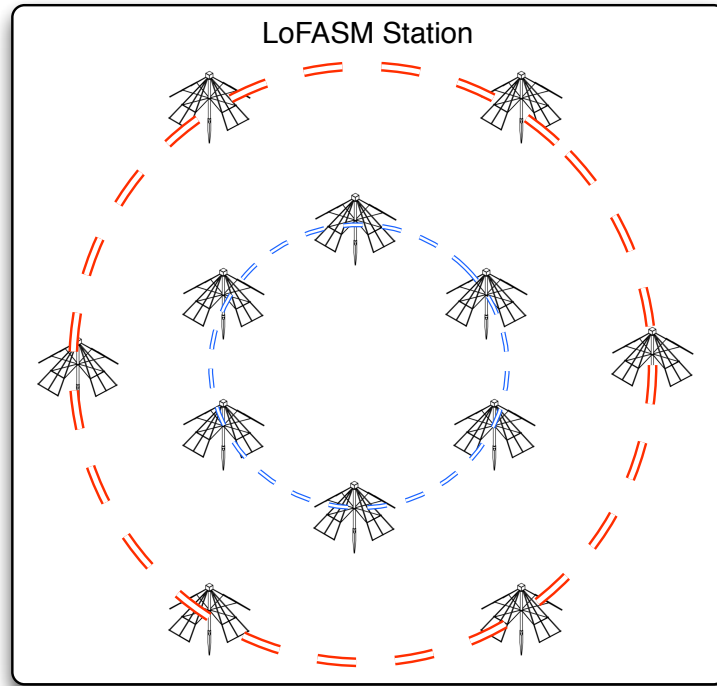
Figure 4: The antenna configuration of a single LoFASM station. Each LoFASM station consists of two concentric rings of 6 dipole antennas, for a total of 12. The radius of the inner ring is 441cm and the outer radius is $\sqrt{3}R_{inner} = 763.83\,\mathrm{cm}$ [18].

Figure 5: An image of the 12 antennas installed at the LoFASM I site in Port Mansfield, Texas
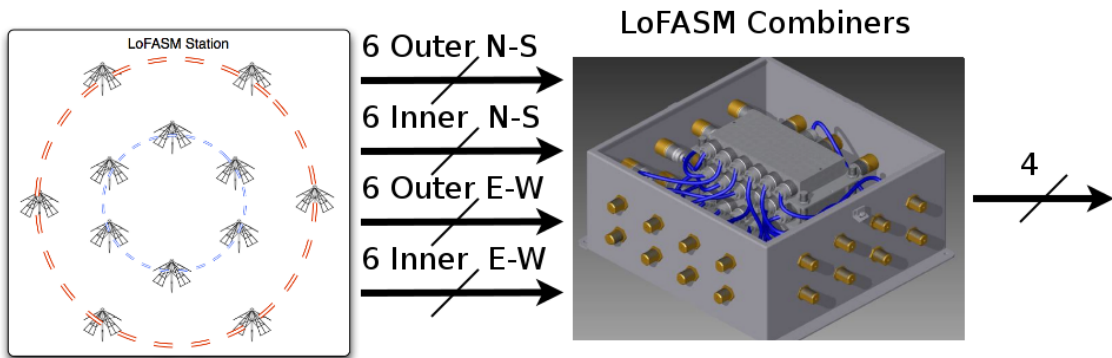
Figure 6: A schematic depicting the first half of the LoFASM signal path. The in-phase combiner box on the right is a shielded container for four separate DC-passing RF combiners. Once the LoFASM signals are appropriately summed to produce only 4 independent lines, the four signals are sent to the analog receiving system.
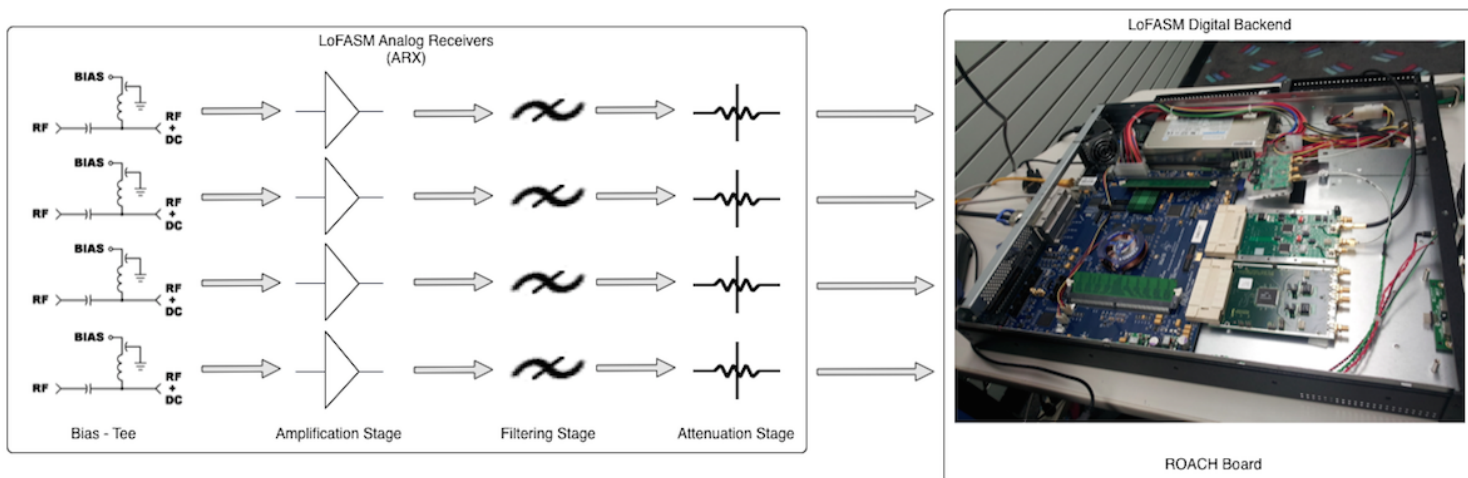
Figure 7: Flow chart portraying the second half of the signal path of a single LoFASM station. The four RF signals are operated on by the mARX and then sent to the LoCo system for digitization and processing.

4 Mitigation of RFI and other complications at low frequencies

In this chapter we discuss the challenges associated with observing the radio sky at low frequencies ($\nu <$ 300 MHz) as well as the measures that are being taken to mitigate them. The system noise of a LoFASM station is essentially dominated by the Galaxy and local RFI. The design parameters of the LoFASM antenna configuration, as well as the arrangement's contributions to RFI suppression are presented in Chapter 4.1. Chapter 4.2 covers the placement of the four separate LoFASM stations across the continental United States and how this will facilitate the detection and removal of local RFI. The effects due to the Galaxy, the ionosphere, and the ISM that were mentioned in Chapter 1.2 are expanded upon in Chapters 4.3-4.5.

4.1 RFI Horizon Rejection and the LoFASM Array Configuration

The mitigation of local RFI is of critical importance for many ground-based radio telescopes. The presence of RFI can be detrimental to the detection of radio sources in the same way that light pollution is the foe of optical astronomy. Faint astronomical emissions are often distorted by nearby, relatively strong, radio sources that are transmitting in (or near) the operating frequency band of the radio telescope. The origin of RFI signals will vary from telescope to telescope, depending on the surrounding region in which the telescope resides. For example, the Arecibo Telescope has to deal with unwanted radio emissions from four different communications towers in the vicinity, as well as various radar signals from offshore ships [25]. The Green Bank Telescope (GBT), on the other hand, is more affected by RFI from nearby populations, like Pocahontas County, WV [29]. It turns out that local emergency services are the cause of infrequent but very strong radio emissions that are picked up by the GBT. Other RFI sources of concern to the GBT include National Oceanic and Atmospheric Administration (NOAA) weather radio broadcasts, digital cameras used by tourists at a tour bus stop about 150 m from the GBT receivers, and on one occasion, a doorbell [11, 29].

With four stations across North America, each LoFASM station will likely be exposed to its own set of RFI signatures in the bands of interest. The first line of defense against unwelcome RFI signals for LoFASM will stem from the Horizon Rejection properties provided by the unique architecture of the antenna array configuration (see Figure 4). The arrangement of the LoFASM dipole antennas emerged from considerations

of the normalized beam pattern response of a 12-element interferometer (in which all the signals are combined in phase) to radio signals approaching the telescope from the horizon. The angle subtended by the horizon was taken to be $10°$ from the azimuthal plane. The geometry of the LoFASM station, depicted in Figure 4, is the result of using minimization techniques to acheive the lowest possible gain response to radio signals originating from sources within $10°$ of the horizon (see Figure 8) [18]. This allows the principal LoFASM beam to remain sensitive to radio signals coming from the sky (near zenith) while simultaneously applying a strong attenuation ($\sim 40\,\mathrm{dB}$) to incoming signals from the horizon, which are presumed to be human-generated RFI.

In order to study RFI mitigation at low frequencies and to explore a largely untouched region in the radio spectrum, the LoFASM array has been designed to be most sensitive at zenith in an narrow band centered at 20MHz. A plot of the normalized beam pattern of a LoFASM station that is optimzed for horizon rejection at $20\,\mathrm{MHz}$ is shown in Figure 8. The horizon rejection a LoFASM station optimized for $20\,\mathrm{MHz}$ signals as a function of frequency is depicted in Figure 9.

4.2   Anti-Coincidence RFI Detection and LoFASM Station Placement

In the interest of exercising new RFI exclusion and suppression techniques, the first four LoFASM stations were placed in distinct locations spanning the continental United States. The LoFASM stations are separated by distances on the order of thousands of kilometers, with the longest distance between any two stations being roughly $3{,}910\,\mathrm{km}$ (LoFASM IV & LoFASM III) and the smallest being about $1{,}218\,\mathrm{km}$ (LoFASM II & LoFASM IV). Figure 10 presents a map of the United States, marking the locations of each of the four LoFASM stations. The fact that the LoFASM stations are separated by such large distances ensures that each of the stations will have its own local RFI environment (i.e., no two LoFASM stations should be affected by the same terrestrial RFI source).

LoFASM, with four identical stations across the continent, generally observing the same area of the sky, will have the capability to triangulate the source of a burst detection. If a burst is detected by all four LoFASM stations then the differences in arrival times at each of the LoFASM stations can be used to

determine the direction of the source on the sky. On the other hand, anti-coincidence techniques will allow LoFASM observers to excise RFI signals that were only detected by one of the four LoFASM sites. A radio signal detected only by one station and not by the other three is most likely some sort of local RFI [4].

4.3   Dispersive effects of the interstellar medium

The effects of the interstellar medium (ISM) between the observer and the source of pulsed radio emission deteriorate the qualily of the data. In particular, one will see a degradation in the time resolution of the data. These effects can completely mask short timescale variations like those of a rapidly rotating pulsar. Free electrons along the line of sight between the observer and the source will cause a pulsed signal to become dispersed and somewhat scattered. Of particular interest is the effect of dispersion on transient signals. The effect of dispersion is directly related to the number of free electrons along the line of sight (the path traveled by the signal), which is called the dispersion measure (DM). DM is defined as the column-density of the electrons in the intervening space between the source and the observer. Equation 1 provides the general expression for dispersion measure as it is used in radio astronomy:

$$DM = \int_0^\infty N(l)\mathrm{d}l. \tag{1}$$

The dispersion measure is conventionally expressed in units of parsecs per cubic centimeter (pc/cm^3) [27]. Dispersion is a frequency-dependent effect; lower-frequency signals are affected more than signals at higher frequencies. Dispersion generally results in lower-frequency signals being delayed in time more than their higher-frequency counterparts. Figure 11 from [17] shows the uncorrected dispersive delays from a pulsar observation centered at 1380 MHz. It can easily be seen that the higher the frequency of the signal, the sooner its time of arrival. We will be using coherent and incoherent dedispersion algorithms developed at CARA to correct for the issues introduced by dispersive effects.

4.4 Galactic background temperatures at low frequencies

The diffuse radio continuum emission in the Galaxy is dominated by synchrotron radiation up to frequencies of a few GHz. This is especially true for any source within a few degrees of the Galactic plane. This radio emission follows a very strong frequency dependence, and is especially strong in our band of interest [16]. It's worth noting that, due to the frequency dependence of the Galaxy's diffuse radio emission, transient objects (e.g., pulsars and rotating radio transients) that are located in or near regions of bright synchrotron emission (like the Galactic plane) are much more difficult to detect at low frequencies [23].

4.5 Ionospheric effects on signals under 10 MHz

The ionized component of the Earth's atmosphere is mostly composed of free electrons and positive ions. This mixture of ions and electrons essentially acts like a mirror for very low frequency RF signals about 300 km above the ground. By the same notion, RFI can be reflected *back to Earth* by the ionosphere and pollute observations. One countermeasure for such events will be the use of the anti-coincidence of RFI signals across the four LoFASM stations (see Chapter 4.2). The hard frequency cutoff, below which ground-based observing is futile, turns out to be around 10 MHz, which is just below the operating band of LoFASM. At these frequencies, inhomogeneities in the ionosphere can contribute additional frequency-dependent phase delays to incoming signals. This makes adding together signals from multiple LoFASM stations much more difficult, since the signals detected at each station will have undergone different ionospheric experiences and, therefore, have different phase delays.
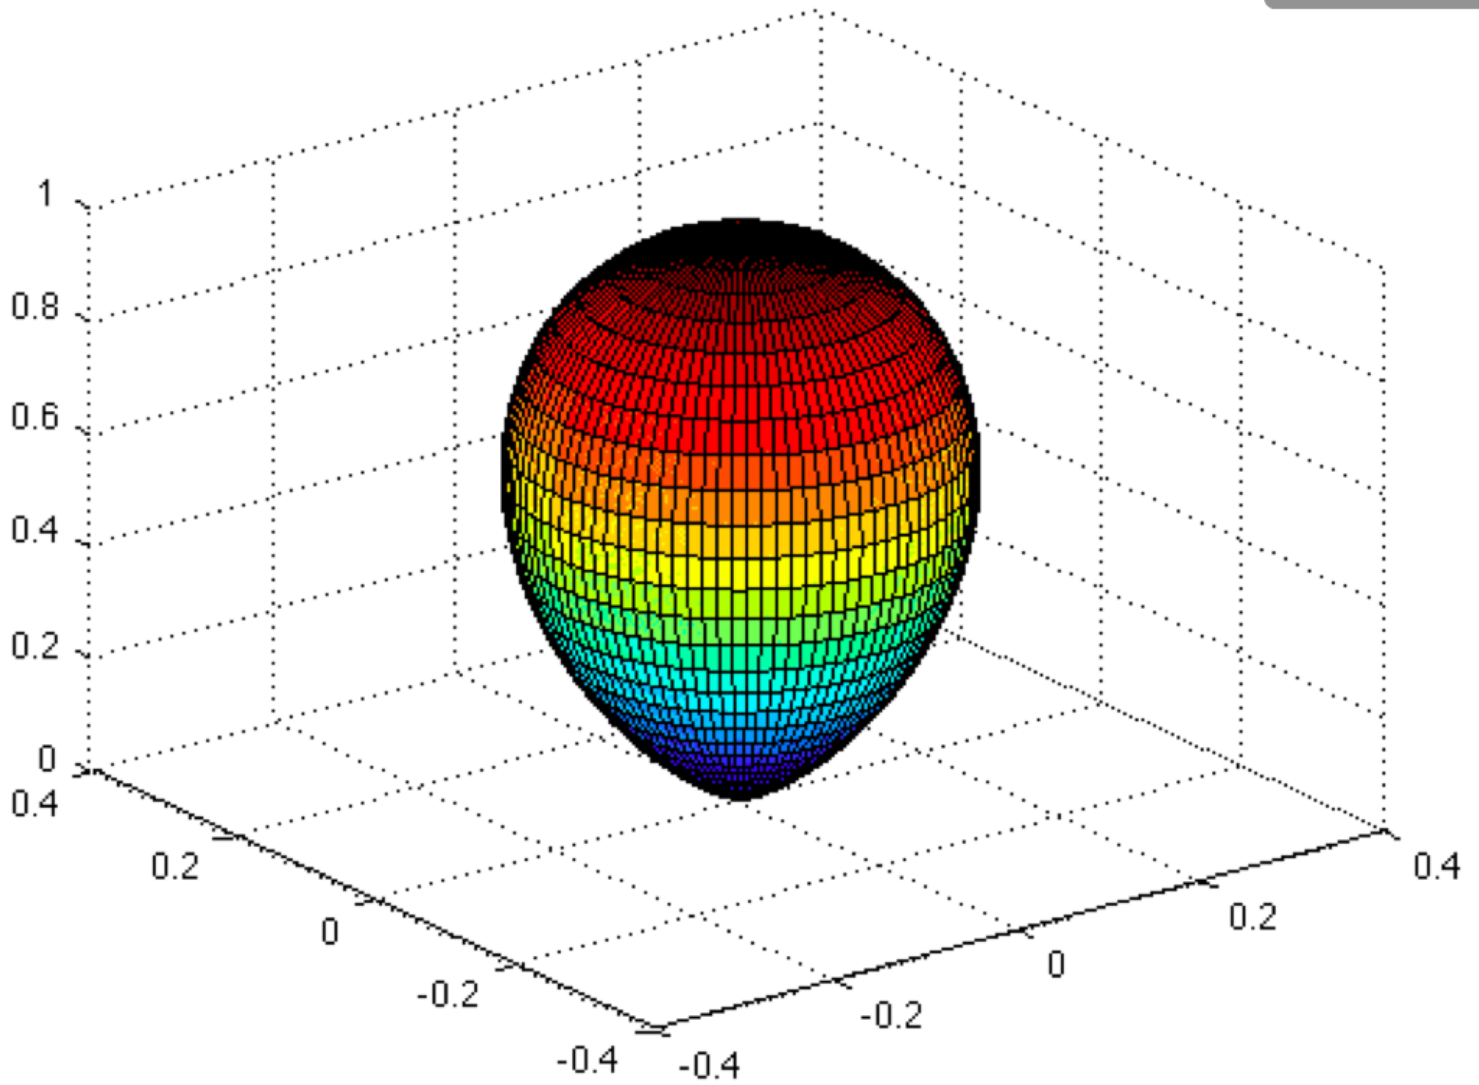
Figure 8: The normalized beam pattern for an optimized LoFASM array. The color pattern represents the telescope's sensitivity, with red being the highest sensitivity and blue being the lowest. For a more detailed discussion on LoFASM's normalized beam pattern and antenna configuration, see [18].
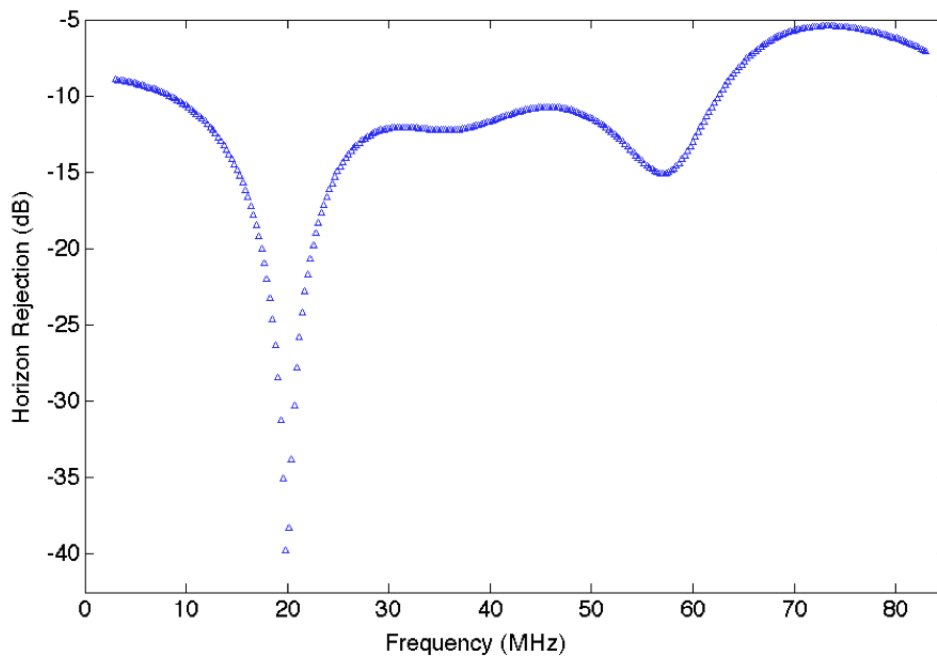
Figure 9: Plot of LoFASM's horizon rejection across the operating band for a LoFASM ring geometrically optimized for 20 MHz. For more info, see [18].

Figure 10: Map depicting the locations of all four LoFASM stations across the continental United States.
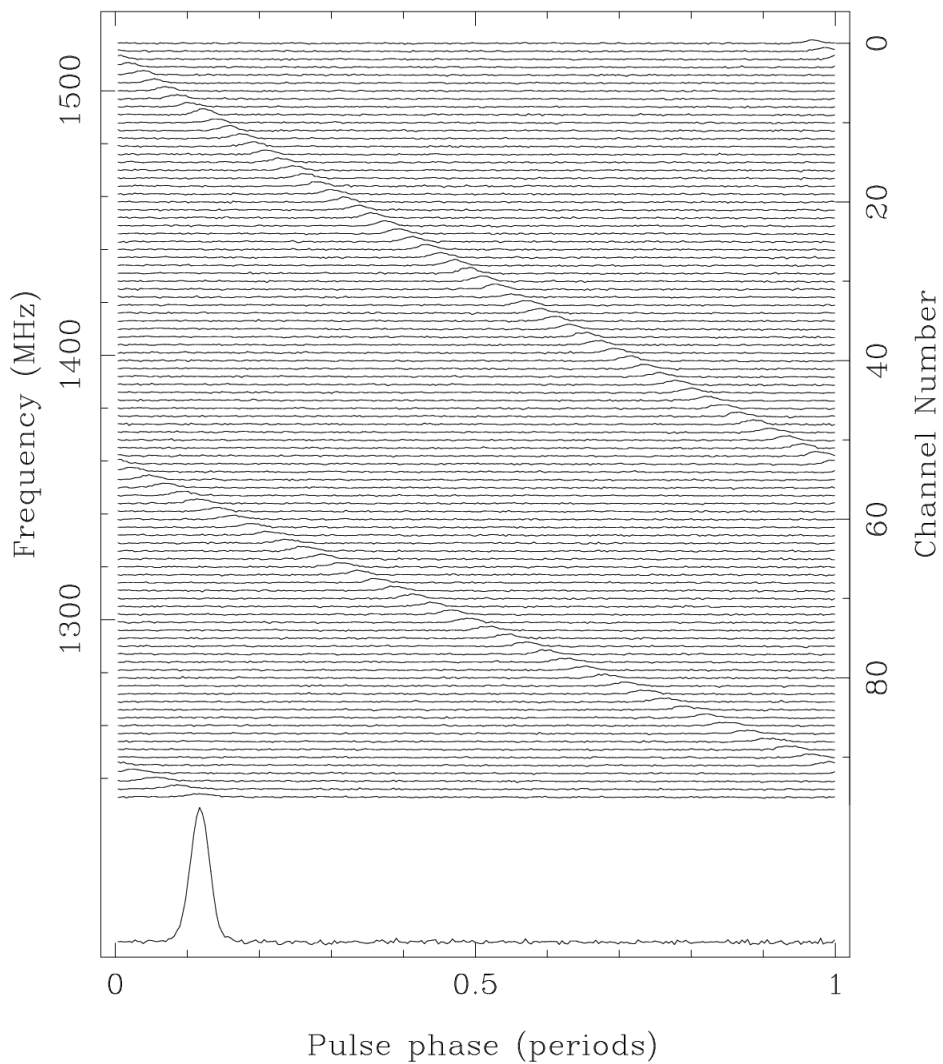
Figure 11: Pulse dispersion. Uncorrected dispersive delays for a pulsar observation over a bandwidth of 288 MHz (96 channels of 3 MHz width each), centered at 1380 MHz. The delays wrap since the data are folded (i.e. averaged) modulo the pulse period (see [17]).

5   Beam Synthesis

In order to harness the RFI mitigation power of LoFASM's dual ring configuration, the signals corresponding to both of the LoFASM antenna rings must summed together and correlated. As mentioned in Chapter 3, the LoFASM backend is a 4-input correlator. These inputs correspond to the four independent RF signals that are provided by the LoFASM array. Correlators are devices or mechanisms designed to measure the complex cross-correlation function $r(\tau)$ as presented in Equation 2. The cross-power terms that are produced by the LoFASM Correlator are used to synthesize the LoFASM beam.

If we let $V_1(t)$ and $V_2(t)$ be the input voltages of two signals to be correlated, then the cross-correlation function is

$$r(\tau) = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} V_1(t) V_2(t + \tau) \mathrm{d}t, \tag{2}$$

where τ is the time by which the voltage signal V_2 lags behind V_1. By making use of the Wiener-Kinchin relation and the convolution theorem we find that the cross-correlation function and the cross-power spectrum (which is needed for beam synthesis) are Fourier Transform pairs:

$$r(\tau) = V_1(t) \star V_2(t) \rightleftharpoons \widehat{V}_1(\nu) \widehat{V}_2^*(\nu), \tag{3}$$

where $\star$ represents the cross-correlation, $*$ represents the complex conjugate, and $\widehat{V}(\nu)$ is the Fourier Transform of $V(t)$. The LoFASM Correlator calculates all 10 (4 real and 6 complex) of the cross-power terms associated with the cross-correlation of four signals. The power detected by the LoFASM array in either of its two polarizations is modeled as

$$P_{sum} = |V_{inner} + V_{outer}|^2, \tag{4}$$

where V_{inner} and V_{outer} are the analog RF signals corresponding to the contributions of both LoFASM rings

in a polarization. We define

$$V_i = A_i V_{inner} \tag{5}$$

$$V_q = A_q V_{outer} \tag{6}$$

$$P_i = |V_i|^2 \tag{7}$$

$$P_q = |V_q|^2 \tag{8}$$

where A_i and A_q are frequency-dependent factors that can be attributed to the unique gains of the two distict digitizers in the RF signal path. V_i and V_q are the digitized LoFASM voltage signals, with P_i and P_q being the power detected in each signal. In being digitized, each of the LoFASM signals is multiplied by a frequency-dependent coefficient that must be corrected. Expanding Equation 4, we have

$$P_{sum} = (V_{inner}^* + V_{outer}^*)(V_{inner} + V_{outer}) \tag{9}$$

$$= |V_{inner}|^2 + |V_{outer}|^2 + V_{inner}^* V_{outer} + V_{outer}^* V_{inner} \tag{10}$$

$$\rightarrow P_{sum} = \frac{P_i}{A_i^2} + \frac{P_q}{A_q^2} + \frac{2}{A_i A_q} \text{Re}(V_i^* V_q). \tag{11}$$

The LoFASM Correlator calculates the values for the LoFASM beam in real time using Equation 9 and stores the resulting cross-power spectrum to disk, along with all the other cross-power terms. It is easy to see that the only terms in Equation 9 that do not depend on the cross-power terms are the gain factors due to the two LoFASM ADC's, A_i and A_q. The fact that LoCo stores all the cross-power terms individually will allow LoFASM data analysts to correct any calibration errors or simply recalibrate (for a specific purpose) at a later time. The LoFASM Correlator will retain and archive all relative phase information of the four input signals. Once implemented, LoFASM's baseband recording features will allow for the storage of absolute phase information that can subsequently be used for anti-coincidence across multiple LoFASM stations.

6　Digital Design

6.1　Field Programmable Gate Array Technologies

The 20[th] century was revolutionized by the advent of electronics. The growth of the computer industry in mobile telephony and the general digitization of many radio services and technologies have boosted many changes that continue to impact modern life in the 21[st] century. In the latter half of the 20[th] century most electronic systems were built and designed by accumulating multitudes of basic components such as microprocessors and memory chips. In general, specific combinations of such components in a single unit were implemented as integrated circuits (ICs) featuring dedicated input/output (I/O) components on printed circuit boards (PCBs). As electronic needs and designs grew more sophisticated, PCBs became increasingly more complicated and difficult. As the number of interconnected components on each PCB grew, so did the probability of incorrectly connecting these components. In particular, the feasibility of properly designing and ensuring that a system was working properly before the production stage was diminishing at alarming rates [28]. Most issues arose from the fact that system specifications and designs would often evolve as boards were being developed. Increasing pressures to design systems that could either meet evolving standards or that were sufficiently malleable to sustain fundamental changes after board construction due to changes in design specifications meant that the concept of having a fully-complete design layout prior to construction or implementation was becoming increasingly unlikely.

Electronic design reached an era that called for the ability to incorporate system design changes after physical device construction. Modern field programmable gate arrays (FPGAs) have filled this void quite nicely. Most modern FPGA chips consist of multiple layers of arrays of logic blocks, memory blocks, I/O regions, and routing channels. These firmware resources can be digitally mapped to satisfy the requirements of applications made up of multiple circuits. The beautiful flexibility of an FPGA stems from being able to reprogram an FPGA after installation, i.e. "in the field," hence the name.

6.2 General Digital Logic Design in MATLAB Simulink using the CASPER Toolflow

The LoFASM Correlator firmware was designed for and implemented on a model XC5VSX95T Virtex-5 FPGA, provided by FPGA vendor Xilinx. The Virtex-5 features 640 48-bit digital signal processing (DSP) blocks and 1,520kb of distributed RAM that are spread out in various regions of the chip. The abundance of DSP regions in the FPGA chip provides digital engineers with the ability to implement very computationally-intensive algorithms (e.g. like those used in correlators) with ease. The Virtex-5 FPGA sits atop the first rollout version of a Reconfigurable Open Architecture Computing Hardware (ROACH I) board. The ROACH board, whose underlying technology was originally developed by the Collaboration for Astronomical Signal Processing (CASPER) from the Berkeley Wireless Research Center, combines the power of a leading FPGA (the Virtex-5) with a PowerPC, and 4 ethernet ports that are capable of transmitting and receiving data rates of up to 10Gbps. Two ZDOK connectors, to which various analog to digital converters (ADCs) can be connected, are also part of the ROACH I design [19]. A schematic of the ROACH I board can be found in Figure 12.

While an explicit firmware mapping in a hardware description language such as Verilog or HDL is usually needed to program an FPGA, a project out of the CASPER produced a specific version of the Linux operating system called the Berkeley Operating system for ReProgrammable Hardware (BORPH) that was designed for FPGA-based reconfigurable computers [15]. Through BORPH, firmware mapping and routing schemes can be implemented on a ROACH board as a normal executable BORPH Object File (or BOF file). BORPH is fully integrated with the ROACH I system and is essentially an extended version of Linux that is able to manage and handle the FPGA's resources as if they were part of the CPU itself. Along with BORPH, the CASPER group also released an entire toolflow that provides system designers with all the tools needed to graphically design a large variety of purpose-driven algorithms. When using this CASPER toolflow, an FPGA programmer needs only interact with MATLAB's Simulink, a block diagram environment for model-based design of intricate systems. An example of an FPGA algorithm designed in Simulink using the CASPER tools to make a LED blink at a periodic interval is shown in Figure 13.

The yellow block labeled "gpio" represents a general purpose I/O device, which in this case is simply a LED. The blue "counter_led" Xilinx block is a 32-bit counter whose output value increases by a step size of 1 on each consecutive FPGA clock cycle. The "slice" block simply picks out a specific bit or range of bits from the counter's output and feeds the boolean version of its input into the gpio block. The gpio block's input (black line on the left) dictates whether or not the particular LED it is mapped to should turn on or off at any point in time. The two leftmost blocks in Figure 13 ("System Generator" and "XSG core config") are necessary to ensure that the system firmware mapping being compiled is compatible with the specific version of Xilinx's Virtex-5 chip that is mounted on LoFASM's ROACH board [22].

The entirety of the Xilinx Simulink blockset, as well as all of the Simulink FPGA blocks developed by CASPER programmers, can be browsed and selected via the Simulink Browser (see Figures 14 and 15).

6.3 Digital Implementation of the LoFASM Correlator

The LoFASM Correlator computes the spectral power density in each of the system's four ADC input channels (each corresponds to a unique polarization) as well as all of the possible cross-power terms (see

Chapter 5).

6.3.1   Master Sync Pulse Generation

The entire LoFASM Correlator firmware is triggered and kept from drifting temporally via a sync pulse generated by an automated sync pulse circuit implemented using Simulink (see Figure 16 below).

The four input lines ("sync_arm", "soft_sync", "adc0_sync", and "adc1_sync") allow for the LoFASM data to be precisely timestamped by a GPS-enabled rubidium oscillation clock. The "adc0_sync" and "adc1_sync" sync inputs shown in Figure 16 are active once per second, corresponding to a 1 pulse-per-second (1PPS) signal that is injected into each ADC card. The "soft_sync" trigger is activated if the appropriate software register on the ROACH board is set high. This can be done by using the LoFASM Python control software. A network of ROACH boards (the LoFASM network has only a single node) can be initiated simultaneously by broadcasting the "sync_arm" to each ROACH board on the LoFASM network. The "sync_arm" flag will cause the firmware to wait for the next 1PPS pulse to arrive before continuing to process the incoming RF signals. This way an observer using LoFASM data will be able to know precisely when an integration has begun by using the resulting data timestamp.

6.3.2   Control Register

The LoFASM Correlator (and also the LoFASM baseband recorder) can be reset and triggered by using a specified group of bits known as the *control register* (see Figure 17). Currently, only the 3 least significant bits of the 32-bit control register are used. This design has the flexibility to add more features to better manipulate the LoFASM Correlator. All software register flags can be set by using the Python-based LoFASM control software library.

6.3.3   Analog to Digital Converters

Because of the high levels of RFI pollution that we expect to encounter with LoFASM (refer to Chapter 4), each LoFASM ROACH board is equipped with two dual-input ADC cards, each of which hosts a pair of 14-bit digitizers (one for each input). An in-depth analysis of the frequency response of these cards is presented

in [20] and [26]. While each ADC card is capable of a sampling rate of up to 400 Msps, the LoFASM ADCs are clocked at 200 Msps. The Simulink blocks for the 14-bit ADCs are demux 1-to-4 (refer to Figure 18), providing four successive time samples for each FPGA clock cycle (the FPGA clock must be clocked at a rate that is slower than the ADCs by a factor of 4 for this to work).

6.3.4 Polyphase FilterBank and Fast Fourier Transforms

The Polyphase Filterbank Finite Impulse Response (PFB FIR) and Fast Fourier Transform (FFT) are the heart and soul of the first stage of the LoFASM Correlator. Before the cross power spectra can be calculated, the temporal voltages at each input must be windowed and converted into frequency space. This is done by applying the Polyphase Filter Bank Technique (refer to [7]) on the voltage time series.

Although there are actually four signals being processed simultaneously, there is only one FIR and one FFT block shown in Figure 19. This is because the CASPER toolflow provided a way to maximize resource utilization on the FPGA by sharing many of the logical components across the four distinct channels. As a result, the PFB and FFT Simulink blocks are able to accept all of the signal inputs in parallel and process them alongside each other. In addition, the Butterfly FFT algorithm being used in the FFT implementation outputs two FFT channels at a time per time series input. In other words, FFT channels 0 and 1 are provided on the same clock cycle by two different output ports, followed by channels 2 and 3, etc. The FFT outputs are broken up into even and odd channels. This allows the firmware to produce a 2048 channel spectrum in only 1024 FPGA clock cycles.

Since the LoFASM ADCs are sampling at 200 Msps we will have a 100 MHz useable bandwidth. The frequency resolution of the resulting spectrum can be calculated as follows:

$$RBW = \frac{\text{Bandwidth}}{N_{\text{channels}}} = \frac{100\,\text{MHz}}{2048 \ \text{chan}} = 48.8\,\text{kHz/chan} \tag{12}$$

6.3.5   Cross Power Calculation: Direct-X Blocks

The latest version of the LoFASM Correlator features a parallel direct multiply and accumulation design flow in which all cross-power pairs are directly calculated by multiplying to produce each baseline pair and accumulating for a user-specified number of integrations. Figure 20 shows the Simulink implementation scheme. The resulting accumulated spectra are stored in block random access memory (BRAM) onboard the FPGA until they are overwritten with the next integration's values. The "acc_ctrl" block controls the integration duration and ensures that all channels are sychronized throughout the design by resetting the appropriate counters and BRAM addresses when necessary. The "dir_x0" and "dir_x1" blocks refer to the even channels and odd channels of each spectrum, respectively.

6.4   Data Rates

At the end of each accumulation, the LoFASM Correlator packages the resulting spectra into 16 8 kB User Datagram Packet (UDP) frames (refer to Chapter 7.2) to be sent to a dedicated data recording machine.

Accompanying the 16 UDP frames is a single 8kB frame that consists solely of header information pertaining to the latest integration. The 17 UDP frames are then transmitted using a 10 Gb Ethernet connection after each integration. The LoFASM Correlator can be set up to produce 1 ms integrations and ship them via UDP in real time. The limiting factor in this process will be the receiving end of the data. The recording machine must be able to write the LoFASM data to disk fast enough to keep up with the FPGA's dump rate. At 1 ms integrations the resulting data rate will be ∼1.038 Gb/s. The recording feature of the LoFASM Control Computer (LoFASM's dedicated recording machine) is controlled using the LoFASM Correlator Python control software (see Chapter 7).
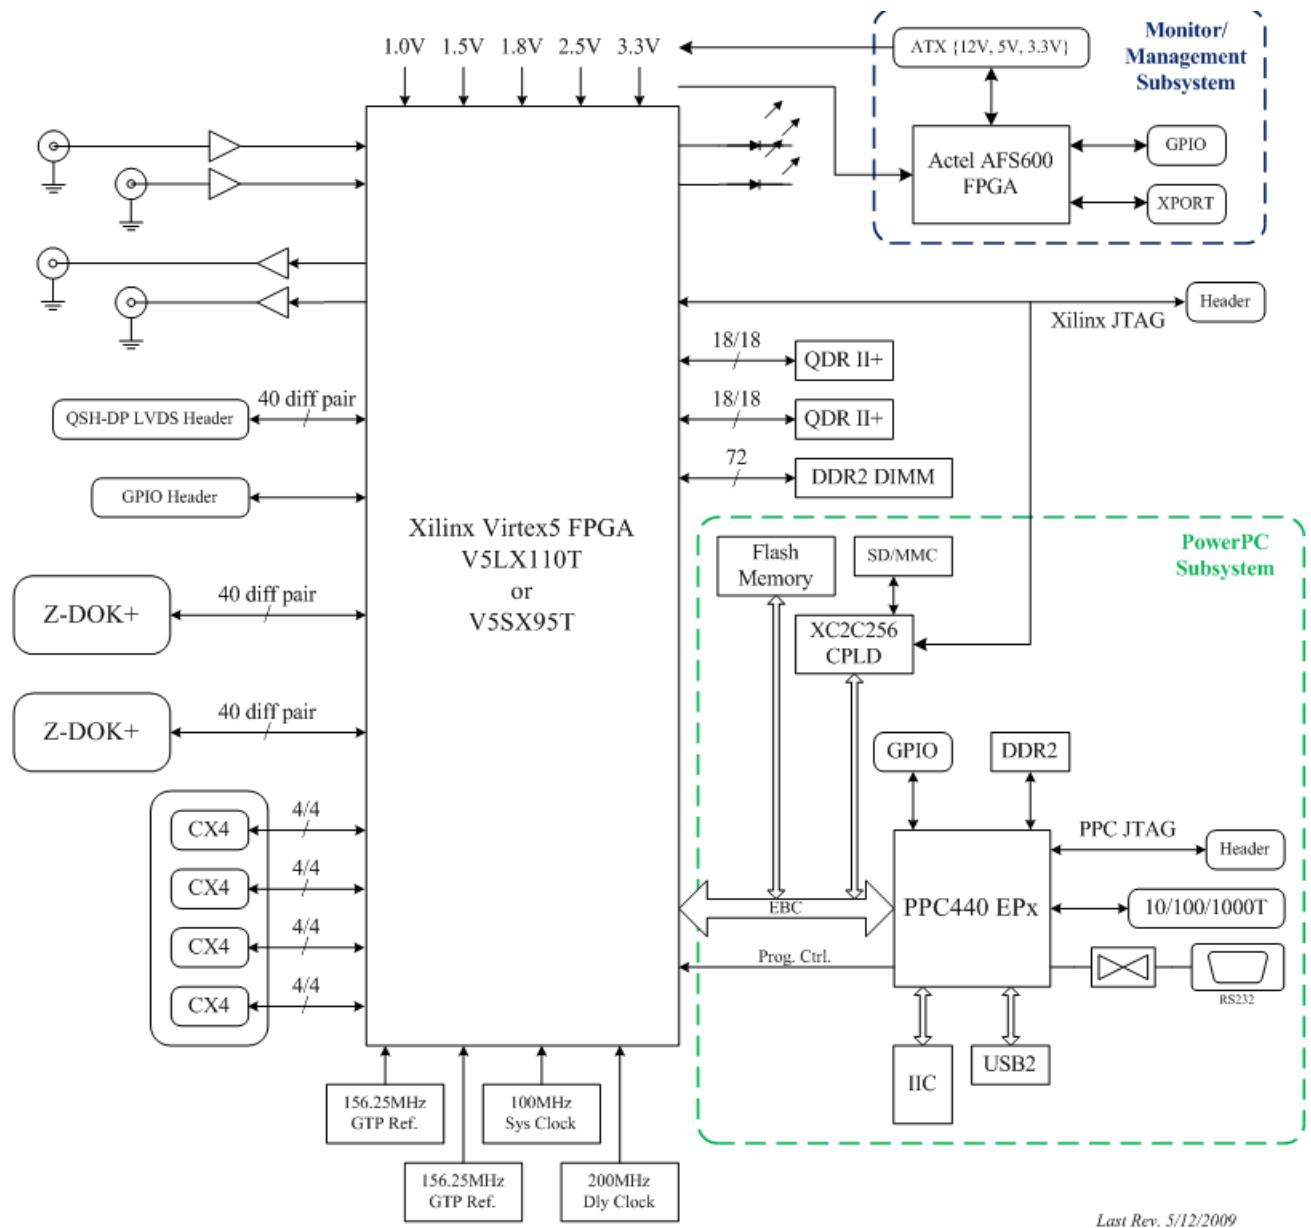
Figure 12: Schematic of ROACH I board. For more detailed information on the architecture of the ROACH I board refer to [5].
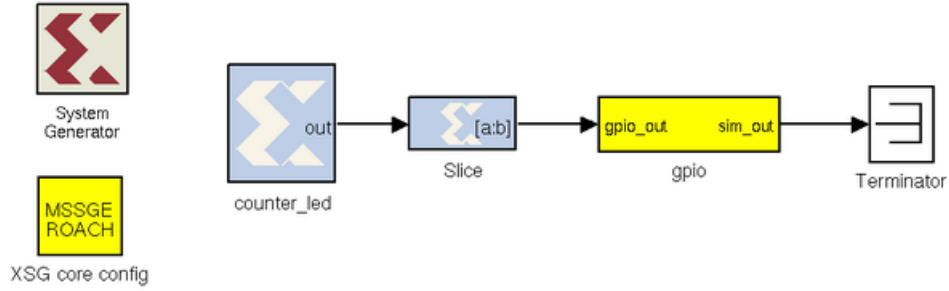
Figure 13: Simulink block diagram of a FPGA program that blinks one of the ROACH I's four on-board LEDs at a pre-defined periodic interval.
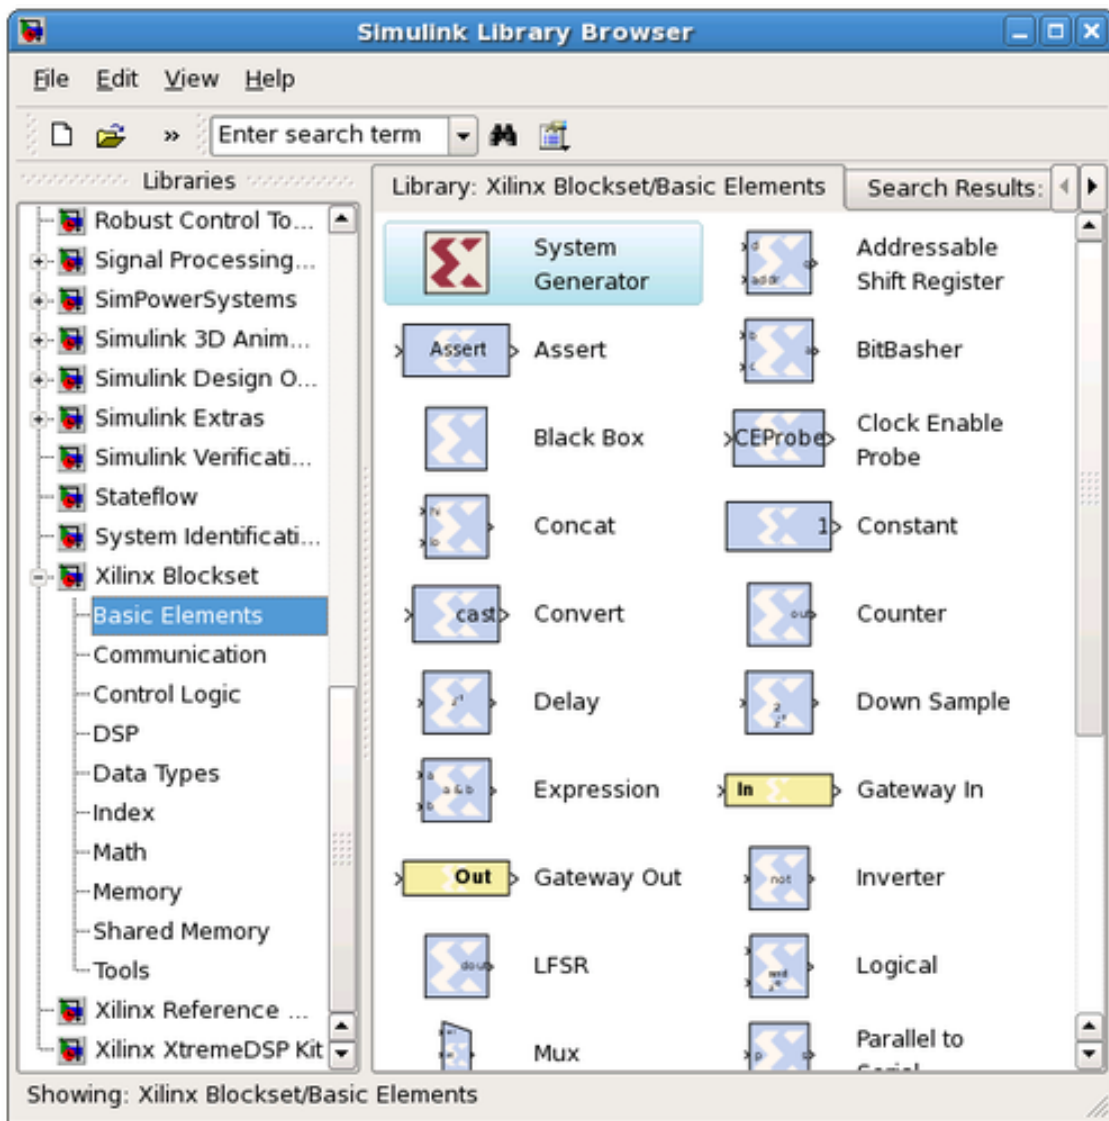
Figure 14: Simulink Library Browser depicting the general Xilinx Simulink blocks for FPGA logic design. All of the Xilinx blocks are blue with an embedded "X".
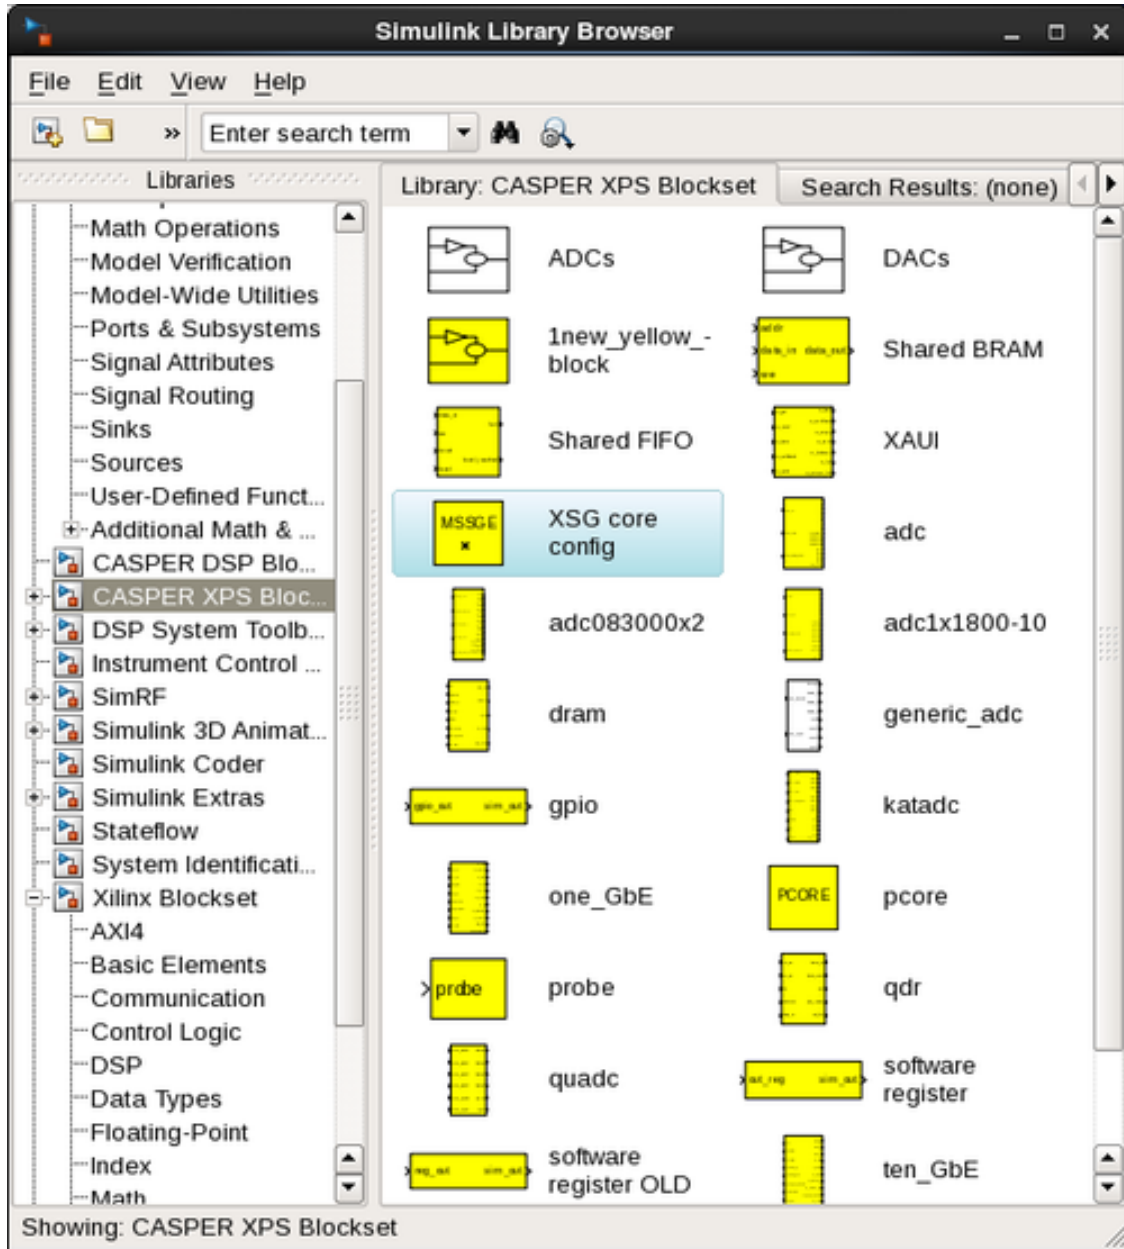
Figure 15: Simulink Library Browser depicting the custom CASPER blocks that were designed specifically for ROACH board implementation. All CASPER blocks are either yellow (interface modules) or green (larger algorithms, e.g. FFT and FIR).
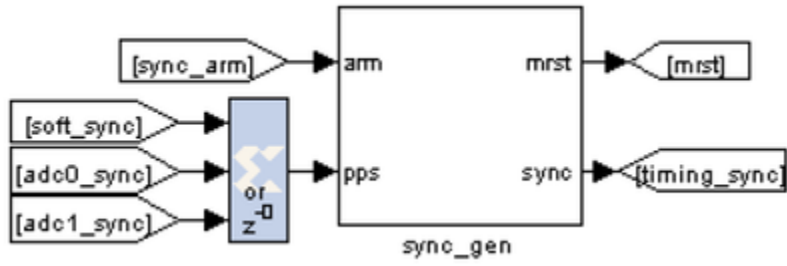
Figure 16: Simulink circuit of the LoFASM Correlator sync pulse generator. For more information about sync pulse usage, refer to [6].
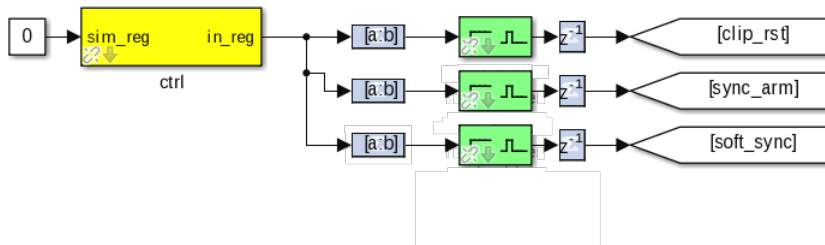
Figure 17: LoFASM firmware control register. This register can be used to arm the ROACH firmware or to reset it if it's already running.
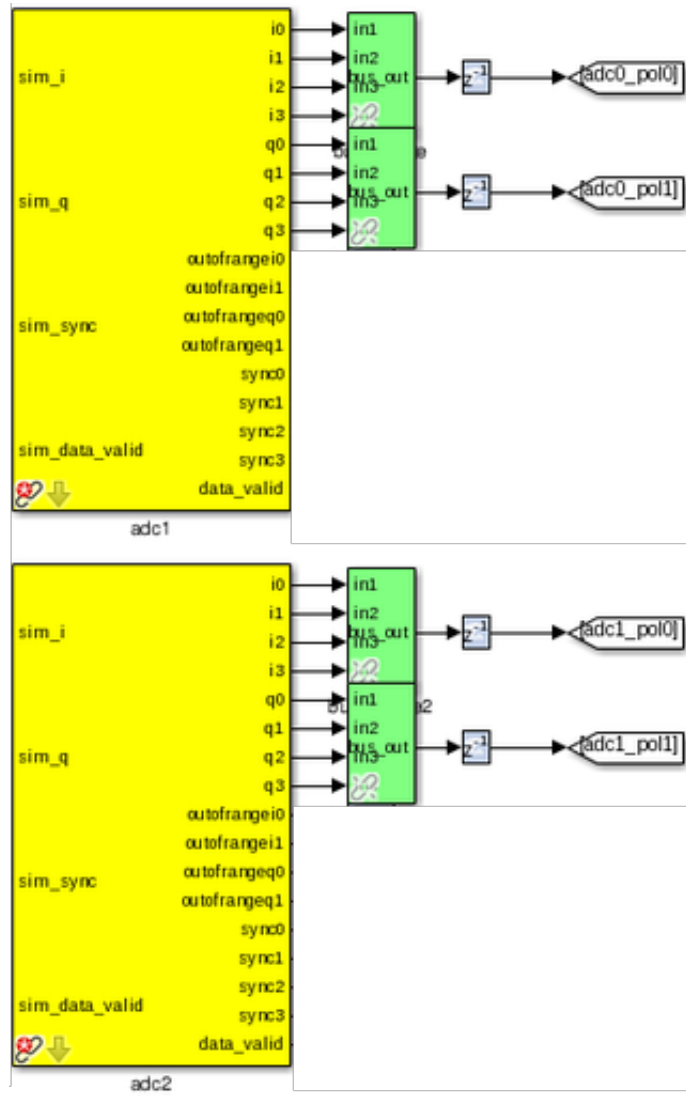
Figure 18: LoFASM ADC blocks. The Simulink blocks for the LoFASM ADCs were provided by the CASPER toolflow. The blocks are yellow because they represent the interface layer between the FPGA logic and the physical ADC digitizers. The output datatype of the yellow ADC blocks is UFIX_14_13, meaning that the resulting numbers are 14 bits in length with their binary point in position 13.
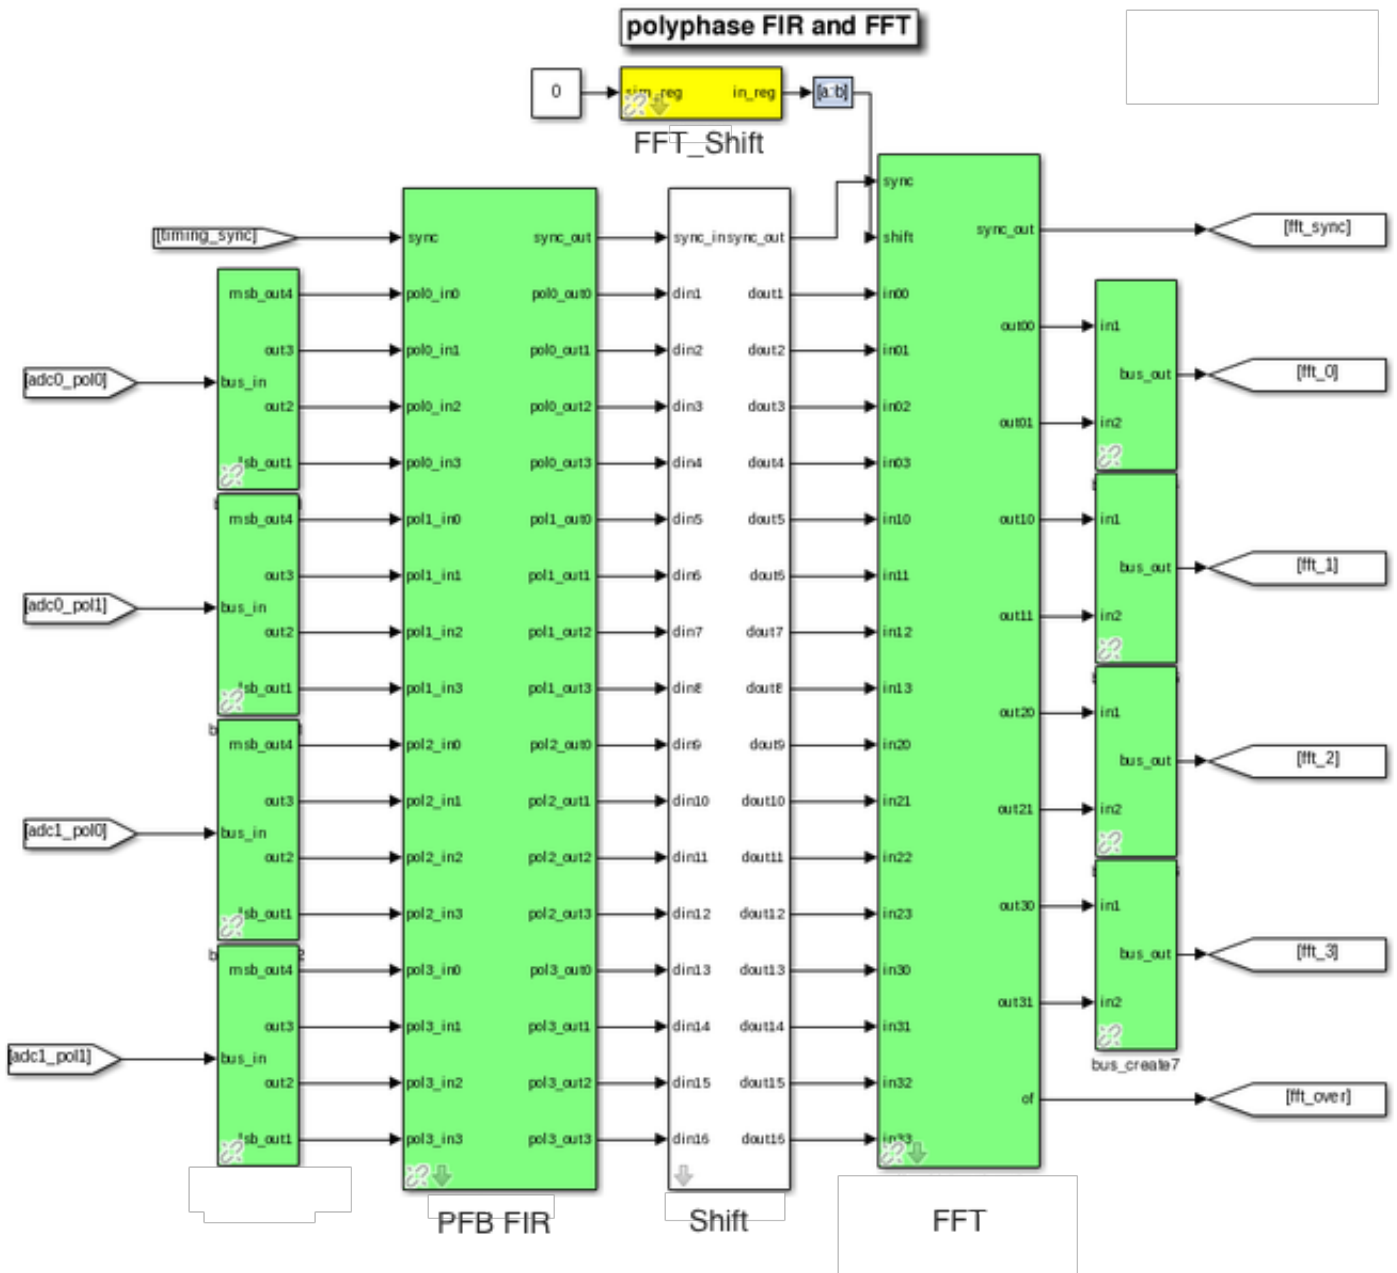
Figure 19: Polyphase FIR and FFT blocks with a single bit shift in between to avoid overflows.
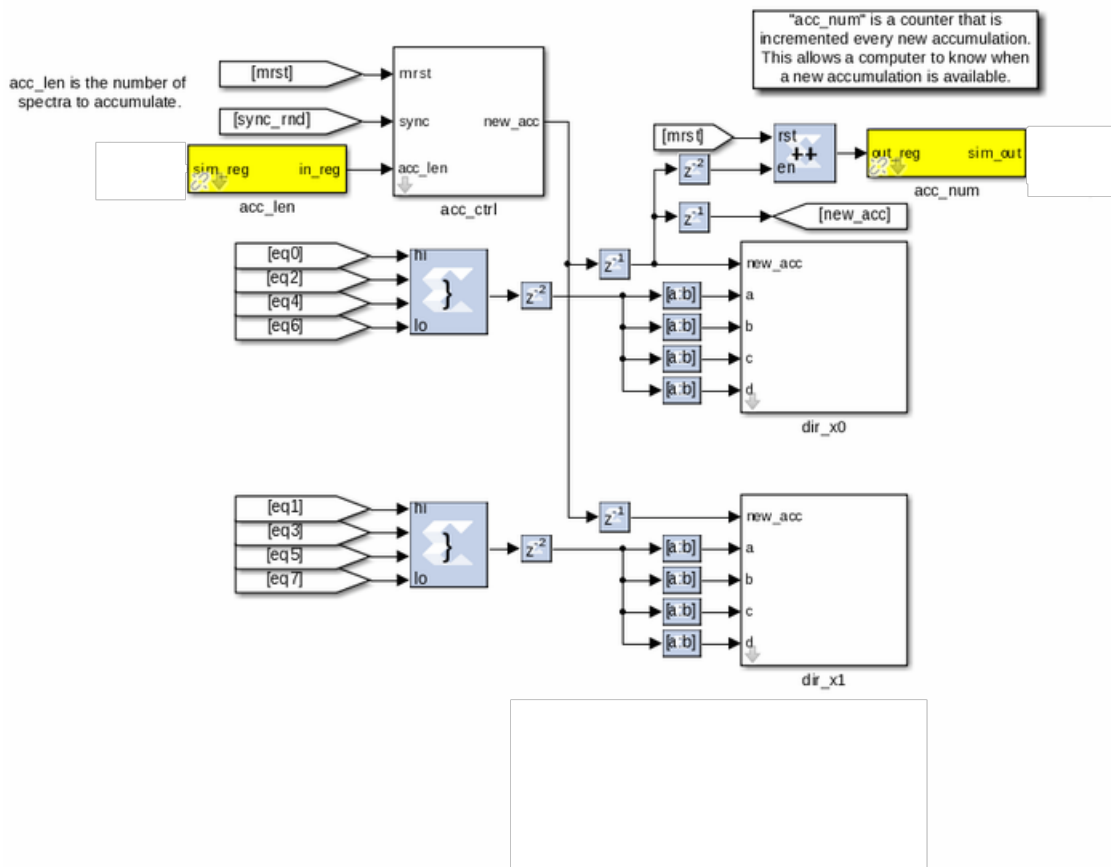
Figure 20: Simulink implementation for multiplying and accumulating all possible cross power pairs for the LoFASM signals. The acc_ctrl block controls the integration length by sending a pulsed signal to the vector accumulators exactly one FPGA clock cycle before a new integration begins.

7 Software Implementation

The LoFASM Correlator can be controlled using a set of Python tools that facilitate the programming and initialization of LoFASM's FPGA. The LoFASM Correlator Python toolset contains two independent scripts designed to provide LoFASM's operator complete control over the digital correlator (see Chapter 7.1.1) and the control computer's data recording function (see Chapter 7.1.2). Chapter 7.2 describes how the LoFASM Correlator data is packaged when it is transmitted over the network as a series of UDP packets.

7.1 LoFASM Correlator Controller Usage

7.1.1 Initialization of the Correlator Firmware

The `initialize.py` Python script is designed to provide the LoFASM operator the ability to program the LoFASM FPGA with the desired version of the LoFASM firware and initialize all of the necessary LoFASM parameters. The current version of the LoFASM Correlator initializer allows the operator to specify the correlator's integration length, digital gains, firmware version, and the destination transmission port and internet protocol (IP) address (see Figure 21).

The `initialize.py` script requires the ROACH board's IP address to be passed as a command-line argument. This IP address can either be provided in the usual format (e.g., 192.168.4.21) or as the name of a host that has been saved in the control computer's `/etc/hosts` file. An example usage of the initialization script using a hostname, along with the corresponding output, is shown in Figure 22. Figure 23 shows the contents of a typical `/etc/hosts` file.

The length of an individual integration performed on the FPGA can be set according to the value of the `acc_len` parameter given to the `initialize.py` script. The integration time (in seconds) is proportional to the number of vector accumulations by

$$T_{\mathrm{int}} = \frac{N_{\mathrm{acc}} N_{\mathrm{chan}}}{2 f_{\mathrm{FPGA}}}, \qquad (13)$$

where T_{int} is the integration period in seconds, N_{acc} is the number of vector accumulations that are performed by the FPGA before the spectra are sent off to be recorded (this is represented by register `acc_len` in the

FPGA firmware), N_{chan} is the number of FFT channels in each spectrum, and f_{FPGA} is the frequency of the FPGA onboard clock in Hz. The factor of 1/2 accounts for the fact that the FPGA calculates two FFT channels in one clock cycle. So, for example, the default `acc_len` value of 2^{18} would result in 2.68 s integrations.

The LoFASM Correlator firmware features a frequency-dependent gain factor that can be set during the initialization of the FPGA. The digital gain factor, which is implemented as a voltage multiplier in frequency-space, can be used to perform interferometry among the four RF signals. Currently, this feature is only used to adjust the correlator's readout levels and explore the dynamic range of the instrument. An example of how to set the gain value is shown in lines 2–4 of Figure 22. Once LoFASM's initialization script has been run, the correlator will continously provide data packets to the recording machine. The reception of the correlator data is controlled by using the `ten_gbe_recorder.py` script (see Chapter 7.1.2).

7.1.2 Recording LoFASM Correlator Data

The process of writing LoFASM data to disk is handled by LoFASM's

`ten_gbe_recorder.py` script. The LoFASM recorder provides the ability to either capture a user-defined

number of UDP packets from the LoFASM network or to write continuously until the program is killed by

the operator. Figure 24 shows the recorder's help menu. Upon execution, the LoFASM recorder resets all the

counter registers within the FPGA in order to obtain accurate timestamps while data is accumulated. An

example of how to execute the recording script to capture and record 1500 UDP packets from the LoFASM

Correlator is shown in Figure 25.

7.2 LoFASM Data Format

Onboard the FPGA, the LoFASM data is packaged in real time into a series of individual UDP frames in

order to facilitate the transmission of the entire set of spectra over the network. The UDP frames are stored

in memory buffers until they are ready to be transmitted by the ROACH 10 Gb Ethernet core. All network

transmissions are sent in 17-packet bursts over the network in between integration dumps. Each LoFASM

integration cycle produces three different architectures of UDP frames. Two of these data architectures are

used to store the auto-correlation spectra and the cross-correlation spectra. The third type of packet that is

produced is a simple header packet. Currently, the LoFASM header packet keeps track of each integration

by storing the values of an accumulation counter and shipping ahead of the 16 data UDP frames.

Each integration consists of 4 auto-correlation packets. Let the four LoFASM FFT inputs be labeled A,

B, C, and D (complex voltages in frequency-space). Then the four auto-correlation terms are A^*A, B^*B,

C^*C, and D^*D, respectively. Each auto-correlation term is saved as a 32 bit (4 byte) long integer. In order

to store all of the FFT spectral information corresponding to the four auto-correlation terms, the data are

arranged in 8kB serial UDP packets that can be visualized as tables with 1024 rows x 64b words, as shown

in Figure 26.

LoFASM's cross-correlation terms are stored in a similar manner as the auto-correlation terms. The

major difference between these two data types arises from the fact that the cross-correlation terms are

complex and may contain negative numbers. For this reason, each cross-correlation value (e.g. either real or imaginary part of a single cross-correlation value) is stored as a 32 bit two's complement signed integer. The LoFASM cross-term UDP frames are organized in such a way that a single 64 bit-wide row will store a single complex value: 32 bits for the real part and 32 bits for the imaginary part. The six individual cross-terms are A^*B, A^*C, A^*D, B^*C, B^*D, and C^*D. Table 27 shows the format of the cross-correlation UDP frames.

When data are being written to disk by the LoFASM data recorder, consecutive spectra are separated only by each integration's header UDP frame. All LoFASM UDP frames have a length of 8192 bytes are in the order shown in Figure 28.

```
──────────────────────── initialize.py help menu ────────────────────────
1  controller@lcc2:~/LoCo/poco_ctrl
2  $ ./initialize.py -h
3  Usage: poco_init.py <ROACH_HOSTNAME_or_IP> [options]
4
5  Options:
6    -h, --help            show this help message and exit
7    -l ACC_LEN, --acc_len=ACC_LEN
8                          Set the number of vectors to accumulate
9                          between dumps. Default is (2^28)/1024.
10   -g GAIN, --gain=GAIN  Set the digital gain (4bit quantisation scalar).
11                         Default is 1000.
12   -s, --skip            Skip reprogramming the FPGA and configuring EQ.
13   -b BOFFILE, --bof=BOFFILE
14                         Specify the bof file to load
15   --10gbe_dest_port=TEN_GBE_DEST_PORT
16                         Specify the 10GbE destination port.
17                         Default is 60001.
18   --10gbe_dest_ip=TEN_GBE_DEST_IP
19                         Specify the 10GbE destination IP address
20                         as a string. Default is '192.168.4.11'.
21   --list_boffiles       Print list of available boffiles and exit.
──────────────────────────────────────────────────────────────────────────
```

Figure 21: initialize.py's help menu.

```
──────────────────────── initialize.py output ────────────────────────
1  controller@lcc2:~/LoCo/poco_ctrl
2  $ ./initialize.py roach \
3          -b poco_14b_fine_tge_2014_Feb_03_1842.bof \
4          -l 4096 --10gbe_dest_ip=192.168.4.11 -g350;
5  Connecting to server roach on port 7147...  ok
6
7  -----------------------
8  Programming FPGA:
9            poco_14b_fine_tge_2014_Feb_03_1842.bof...done
10 Configuring fft_shift... done
11 Configuring accumulation period... done
12 Resetting software triggering and resetting error counters... done
13 Setting gains of all channels on all inputs to 350... done
14 ==========================================
15 10GbE Configuration on ROACH @ 192.168.4.44:
16 ROACH IP: 192.168.4.5    LCC IP: 192.168.4.11
17 ROACH Port: 60000        LCC Port: 60001
18 ==========================================
19 Configuring 10GbE Packet Transmitter...
20 Starting 10GbE core...
21 Cable verified connected to port 0.
22 core_name gbe0
23 mac_base: 2207613190144
24 fabric_ip: 192.168.4.5
25 mac_base+fabric_ip: 2210845426693
26 fabric_port: 60000
27 ok, all set up.
──────────────────────────────────────────────────────────────────────────
```

Figure 22: An example depicting how initialize.py is typically used to program the LoFASM firmware.

```
                                  ─── /etc/hosts ─────────────────────────
1  controller@lcc2:~/LoCo/poco_ctrl
2  $ cat /etc/hosts
3  146.88.201.45   pdu
4
5  # ROACH 100mb port
6  192.168.4.44    roach
7
8  # ROACH 10GbE ports
9  10.10.10.1  roach10g1
10 10.10.10.2  roach10g2
11 10.10.10.3  roach10g3
12 10.10.10.4  roach10g4
13
14
15
16 127.0.0.1       localhost
17 127.0.1.1       lcc2
18
19 # The following lines are desirable for IPv6 capable hosts
20 ::1     ip6-localhost ip6-loopback
21 fe00::0 ip6-localnet
22 ff00::0 ip6-mcastprefix
23 ff02::1 ip6-allnodes
24 ff02::2 ip6-allrouters
25 You have new mail in /var/mail/controller
```

Figure 23: LoFASM IV's /etc/hosts file.

```
                                  ─── ten_gbe_recorder.py help menu ──────────
1  $ ./ten_gbe_recorder.py -h
2  Usage: ten_gbe_recorder.py
3
4  Options:
5    -h, --help            show this help message and exit
6    --max_packets=MAX_PACKETS
7                          specify number of packets to write to disk.
8    -g FPGA_GAIN, --gain=FPGA_GAIN
9                          fpga gain
10   -C CONTINUOUS_WRITE   Set to write data continously.
```

Figure 24: The LoFASM Recorder's help menu.

```
                                  ─── LoFASM recording example ───────────────
1  $ python ten_gbe_recorder.py  --max_packets=1500 -g275
2  Setting up the socket interface
3  Created bound socket on 192.168.4.11:60001
4  receiving packets
5  done writing packets to:
6  /data/20140331/20140331_014943_275.lofasm
```

Figure 25: An example execution of the LoFASM data recording script.

32 bits	32 bits
$A_0^* A_0$	$B_0^* B_0$
$A_2^* A_2$	$B_2^* B_2$
$A_4^* A_4$	$B_4^* B_4$
$\vdots$	$\vdots$
$A_{1022}^* A_{1022}$	$B_{1022}^* B_{1022}$

Figure 26: Table depicting the contents of a UDP packet containing auto-correlation data for inputs A and B. The subscripts indicate the FFT channel that is being stored. The packet depicted above only stores the data corresponding to the even channels of the auto-correlations for inputs A and B. Two UDP frames are needed to store the entire spectrum for A and B, and another two for C and D. The width of each row in the packet is 64 bits. The length of the table is 1024 rows or 8192 Bytes.

32 bits	32 bits
$\mathrm{Re}(A_0^* B_0)$	$\mathrm{Im}(A_0^* B_0)$
$\mathrm{Re}(A_2^* B_2)$	$\mathrm{Im}(A_2^* B_2)$
$\mathrm{Re}(A_4^* B_4)$	$\mathrm{Im}(A_4^* B_4)$
$\vdots$	$\vdots$
$\mathrm{Re}(A_{1022}^* B_{1022})$	$\mathrm{Im}(A_{1022}^* B_{1022})$

Figure 27: Table depicting the contents of a UDP packet containing cross-correlation data for the $A^* B$ cross-term. Each value is signed and must be treated accordingling when parsing the raw LoFASM data. Two UDP frames are required to store the information for a single cross-term: one for even FFT channels and one for odd FFT channels. The length of each UDP frame is 8kB.

Packet No.	Frame	FFT Channels	Data Type
1	Header	N/A	Unsigned Long Int
2	$A^* A$ & $B^* B$	Even	Unsigned Long Int
3	$C^* C$ & $D^* D$	Even	Unsigned Long Int
4	$A^* A$ & $B^* B$	Odd	Unsigned Long Int
5	$C^* C$ & $D^* D$	Odd	Unsigned Long Int
6	$A^* B$	Even	Signed Long Int
7	$A^* C$	Even	Signed Long Int
8	$A^* D$	Even	Signed Long Int
9	$B^* C$	Even	Signed Long Int
10	$B^* D$	Even	Signed Long Int
11	$C^* D$	Even	Signed Long Int
12	$A^* B$	Odd	Signed Long Int
13	$A^* C$	Odd	Signed Long Int
14	$A^* D$	Odd	Signed Long Int
15	$B^* C$	Odd	Signed Long Int
16	$B^* D$	Odd	Signed Long Int
17	$C^* D$	Odd	Signed Long Int

Figure 28: Table showing the UDP frame order of transmission after each LoFASM integration. The first frame is always the header and is used to identify the spectra that follow. Frames 2–5 contain all of the auto-correlation data and the remaining frames contain all of the cross-correlation terms. The size of a full integration, including the header frame, is 136 kB.

8 Conclusion

The Low Frequency All Sky Monitor is installed and will soon be running at four discrete locations across the continental United States. LoFASM's digital backend consists of a 4-input full Stokes correlator that executes the synthesis of a wide beam sensitive to radio emission approaching from the zenith with horizon rejection at 20 MHz. The LoFASM Correlator is currently in its pre-commissioning stages at LoFASM IV, at NASA's Deep Space Network Goldstone Communications Complex, and will soon be installed at every other LoFASM station. With all four LoFASM stations taking data, we will be able to use anti-coincidence techniques to excise unwanted RFI in our bands of interest and focus on radio transients overhead. The next version of the LoFASM backend will feature a voltage baseband recorder, enabling the use of coherent dedispersion algorithms to study LoFASM data.

References

[1] B. Abbott and the LIGO Scientific Collaboration. LIGO: The Laser Interferometer Gravitational-Wave Observatory. *ArXiv*, November 2007.

[2] Nils Andersson, John Baker, Kris Belczynski, and et al. The Transient Gravitational-Wave Sky. *ArXiv*, May 2013.

[3] M.E. Bell, T. Murphy, D. L. Kaplan, and et al. A survey for transients and variables with the Murchison Widefield Array 32-tile prototype at 154 MHz. *ArXiv e-prints*, November 2013.

[4] N.D.R. Bhat, J.M. Cordes, S. Chatterjee, and T.J.W. Lazio. RFI Identification and Mitigation Using Simultaneous Dual Station Observations. *ArXiv*, May 2013.

[5] H. Chen. ROACH Architecture. `https://casper.berkeley.edu/wiki/ROACH\_Architecture`, May 2009.

[6] H. Chen, D. MacMahon, and A. Parsons. Sync Pulse Usage in CASPER DSP Blocks. `https://casper.berkeley.edu/memos/sync\_memo\_v1.pdf`, August 2008.

[7] Jayanth Chennamangalam. The Polyphase Filter Bank Technique. `https://casper.berkeley.edu/wiki/The\_Polyphase\_Filter\_Bank\_Technique`, August 2011.

[8] S.W. Ellingson, G.B. Taylor, J. Craig, J. Hartman, J. Dowell, C.N. Wolfe, T.E. Clarke, B.C. Hicks, N.E. Kassim, P.S. Ray, L. J. Rickard, F.K. Schinzel, and K.W. Weiler. The LWA1 Radio Telescope. *ArXiv e-prints*, January 2013.

[9] P.A. et al. Evans. Swift follow-up observations of candidate gravitational-wave transient events. *ArXiv*, May 2012.

[10] M.A. Garrett. Radio Astronomy Transformed: Aperture Arrays Past, Present & Future. *Proceedings of Science*, July 2013.

[11] Frank Ghigo, John Ford, and Carla Beaudet. The Experiment. `http://www.gb.nrao.edu/IPG/Interference/Showme.pdf`.

[12] M.P. van Haarlem, M.W. Wise, A.W. Gunst, G. Heald, J.P. Mckean, J.W.T. Hessels, and et al. LOFAR: The LOw-Frequency ARray. *ArXiv e-prints*, May 2013.

[13] R.A. Hulse and J.H. Taylor. Discovery of a pulsar in a binary system. *ApJ*, 1974.

[14] K. G. Jansky. Electrical Disturbances Apparently of Extraterrestrial Origin. *IEEE*, 1998.

[15] Hayden Kwok-Hay So. BORPH: Berkley Operating System for ReProgrammable Hardware. `http://www.eee.hku.hk/~hso/borph.html`, 2013.

[16] K.D. Lawson, C.J. Mayer, J.L. Osborne, and M.L. Parkinson. Variations in the spectral index of the galactic radio continuum emission in the northern hemisphere. *MNRAS*, 1987.

[17] D. Lorimer and M. Kramer. *Handbook of Pulsar Astronomy*. Cambridge University Press, 2005.

[18] Rossina B. Miller. The Low Frequency All Sky Monitor for the Study of Radio Transients: Array Configuration and Sensitivity. Undergraduate Thesis, 2012.

[19] A. Parsons, D. Backer, C. Chang, D. Chapman, H Chen, P. Droz, C. de Jesus, D. MacMahon, A. Siemion, D. Werthimer, and M. Wright. A New Approach to Radio Astronomy Signal Processing: Packet Switched, FPGA-based, Upgradeable, Modular Hardware and Reusable, Platform-Independent Signal Processing Libraries. `https://casper.berkeley.edu/papers/200509URSI.pdf`, September 2005.

[20] Rick Raffanti. Casper Memo #34. `http://www.gb.nrao.edu/IPG/Interference/Showme.pdf`.

[21] Nan Rendong, Li Di, Chengjin Jin, Qiming Wang, Wenbai Zhu, Lichun Zhu, Haiyan Zhang, Youling Yue, and Lei Qian. The Five-Hundred-Meter Aperture Spherical Radio Telescope (FAST) Project. *ArXiv*, May 2011.

[22] A. Siemion, J. Manley, W. New, and M Wagner. CASPER Tutorial 1. `https://casper.berkeley.edu/wiki/Introduction\_to\_Simulink`, September 2013.

[23] B.W. Stappers, J.W.T. Hessels, and et al. Observering pulsars and fast transients with LOFAR. *ArXiv*, April 2011.

[24] A.R. Thompson, J.M. Moran, and G.W. Swenson Jr. *Interferometry and Synthesis in Radio Astronomy*. WILEY-VCH, 2004.

[25] Reinaldo Veldez. RFI at Arecibo. `http://www.naic.edu/~rfiuser/`, May 2007. Accessed: 2014-01-02.

[26] M. Wagner and R. Raffanti. ADC2x400-14. `https://casper.berkeley.edu/wiki/ADC2x400-14`, July 2009.

[27] T.L. Wilson, K. Rohlfs, and S. Hüttemeister. *Tools of Radio Astronomy*. Springer, 5 edition, 2009.

[28] Roger Woods, John McAllister, Ying Yi, and Gaye Lightbody. *FPGA-based Implementation of Signal Processing Systems*. Wiley, 2008.

[29] Paulette Woody. Table of Known RFI Sources. `https://safe.nrao.edu/wiki/bin/view/GB/Projects/RFIReportsTable`, October 2013. Accessed: 2014-01-02.